

Computing Price Trajectories in Combinatorial Auctions with Proxy Bidding

Jie Zhong, Gangshu Cai, and Peter R. Wurman
Computer Science Department
North Carolina State University
Raleigh, NC 27695-7535

{jzhong,gcai,wurman}@ncsu.edu

ABSTRACT

Proxy bidding has proved useful in a variety of real auction formats, such as eBay, and has been proposed for combinatorial auctions. Previous work on proxy bidding in combinatorial auctions requires the auctioneer essentially run the auction with myopic bidders to determine the outcome. In this paper we present a radically different approach that computes the bidders' allocation of their attention across the bundles only at the inflection points. Inflections are caused by the introduction of a new bundle into an agent's demand set, a change in the set of currently competitive allocations, or the withdrawal of an agent from the set of active bidders. This algorithm has several advantages over alternatives, including that it computes exact solutions and is invariant to the magnitude of the bids.

Categories and Subject Descriptors

F.2.2 [Theory of Computation]: Analysis of Algorithms and Problem Complexity; J.4 [Computer Applications]: Social and Behavioral Science—*Economics*; K.4.4 [Computing Milieux]: Computers and Society—*Electronic Commerce*

General Terms

Combinatorial Auctions, Proxy Bidding, Auction Algorithms

1. INTRODUCTION

Iterative Combinatorial Auctions are auctions that combine the expressive power of combinational bids with the progressive march towards a solution common in iterative, single-item mechanisms like the English auction (see [8] for an introduction). Iterative combinatorial auctions are attractive for several reasons, but particularly because they require bidders to determine exact values on items, or combinations of items, only if those items are relevant to the final allocation. In cases where value determination is costly, bidders need not compute the value of bundles that they do not expect to win in the auction.

Recently, mechanism designers have expressed interest in enabling

proxy bidding in iterative combinatorial auctions. Proxy bidding is the process in which the bidder expresses a bid—typically greater than strictly necessary to be accepted by the auction—to an agent that then bids incrementally on behalf of the bidder until it either wins or exhausts the authority granted it by the bidder. In a combinatorial auction, the bidder sends a message to the proxy that expresses value for some or all of the bundles. The proxy bidder, guided by the message, bids incrementally on behalf of the user by following some prescribed bidding policy.

Proxy bidding has the advantage of speeding up the auction by allowing bidders to place larger bids which will be executed only to the extent necessary to outbid competitors [2, 9, 10]. Also, by restricting the strategic flexibility of the bidders, mechanism designers may be better able to design successful auctions and predict their outcomes. Indeed, Ausubel and Milgrom [2] show that a semi-sincere equilibrium always exists under proxy bidding in their Ascending Package Auction (APA).

In addition, enabling an auction with proxy bidding may reduce the need for the bidders to accurately estimate the valuations of the other participants in the auction. For example, an equilibrium strategy in a first-price sealed-bid auction requires estimating the value of the second highest bidder. However, when a first-price, sealed-bid auction is enhanced with proxy bidding, as on eBay when snipers bid at the end of the auction, it reduces to a Vickrey auction and each bidder's equilibrium strategy is to submit her true value [11]. Iterative combinatorial auctions are more complex than the English auction format used on eBay, and proxy bidding does not necessarily produce the incentive compatibility property of the Generalized Vickrey Auction (GVA). Nevertheless, there may be many problems for which proxy bidding simplifies the bidders' strategy selection problems, though we do not address strategic issues in this paper.

Once the auction has collected the proxy bids, it must compute the final prices and allocation entailed by the bids and the proxy bidding policy. We refer to this task as solving the *Proxy Auction Problem* (PAP). A natural approach, which we refer to as solving by simulation, is to instantiate the proxy bidders as algorithms that, using the bidders' messages, incrementally bid until the winners and prices are determined. Recently [13], we provided an alternative algorithm that more directly computes the outcome of the auction by computing the bidding patterns engendered by the proxy statements, and the subsequent price trajectories. The previous version of the algorithm provides the basic framework of the approach but was limited by a simplifying assumption on the events that cause

changes in bidding behavior. In this paper, we relax that assumption and present a technique that solves the general problem.

In the next section we define a simple combinatorial auction with proxy bidding. Section 3 describes our novel algorithm for solving PAP, with particular attention to the task of computing the bidders’ behavior. Section 4 walks through a portion of the calculations necessary to solve the running example problem. The final section contains observations on aspects of the algorithm and a discussion of future work.

2. A SIMPLE COMBINATORIAL PROXY AUCTION

We consider an iterative CA that accepts bids on bundles of items, and generates non-linear bundle-prices, that is, the price of a bundle may not be the sum of the prices of individual items. Several proposed iterative combinatorial auctions fit into this category, including various versions of *iBundle* [7, 9], *AIBA* [14], and, with some interpretation, *APA* [2]. The set of rules we employ here define a simplified iterative CA that combines aspects of the existing proposals. The details of the auction format described in this paper were chosen because of their natural interpretation within the framework of the proxy solution algorithm.

The problem faced by the auctioneer is to allocate n heterogeneous items to m buyers. Let \mathcal{J} be the set of items, and \mathcal{I} be the set of buyers, and index the sets with j and i , respectively. There are 2^n different combinations of items, including the empty set, \emptyset . Let $\mathbf{b} \in \{0, 1\}^n$ where $b^j = 1$ implies that item j is an element of the bundle \mathbf{b} .

The auction proceeds in rounds, and in each round the buyers may place offers on a subset of bundles. We denote buyer i ’s bid at round t as a collection of mutually exclusive offers on bundles of the form $r_i^t(\mathbf{b})$, where $r_i^t(\mathbf{b}) \in \mathcal{R}_+$,¹ and we make the standard assumption that $r_i^t(\emptyset) = 0$. The auction remembers each bidder’s last offer on each item.

It is convenient for our reduction to the algorithm for PAP to assume that the bidding rules require that an agent either pass or improve its bid on exactly one bundle by increasing it to the current price plus a small increment, δ . This restriction could also be implemented as part of the proxy agent’s bidding strategy.

After the bids are received, the auction computes the winning combination of bids, a task called solving the *Winner Determination Problem* (WDP). Algorithms for solving the WDP have been widely studied in recent years [1, 3, 12]. At the end of each round the new bundle-prices are announced and the auction tells the bidders which bundles, if any, they are winning.

The announced prices are simply the highest bid received on a particular item and are *anonymous* in that all bidders are given the same information. However, the prices are not necessarily *separating* [16] because the optimal allocation may include bidders whose last offer on their winning bundle is less than the current price; hence we need to directly inform bidders of their winning status. Let $\pi_{\mathbf{b}}$ denote the price associated with bundle \mathbf{b} , and

$$\pi_{\mathbf{b}} = \max_i r_i^t(\mathbf{b}).$$

¹This standard XOR format is fully expressive, although not the most concise method of expressing bids [6].

In the proxy-enabled auction, the bidder submits to the proxy agent a (not necessarily truthful) value statement, v_i , that defines a willingness to pay for some subset of bundles. The proxy agent submits the incremental bids, $r_i^t(\mathbf{b})$, on behalf of the user by following a *straightforward bidding policy* [5]. If it is told by the auction that it is winning, the agent does not increase its bid. Otherwise, the agent bids on bundle, \mathbf{b}' , that maximizes its real surplus at the given prices,

$$\mathbf{b}' = \arg \max_{\mathbf{b}} \{v_i(\mathbf{b}) - (\pi_{\mathbf{b}} + \delta)\}. \quad (1)$$

where δ is the minimum bid increment and π is the set of announced bundle-prices. If more than one bundle satisfies (1), the agent selects one randomly to bid on. We refer to the set of bundles that satisfy (1) as the agent’s *demand set*² and denote it D_i^t .³

Note that agents will often reach a point at which no bundle provides positive surplus at the current prices. When this occurs, the agent will stop bidding in the auction, though its most recent bids remain. The auction terminates when no new bids are received in a round.

We designate the auction defined by the combination of proxy bidding and these simple rules as the *Simple Combinatorial Proxy Auction* (SCPA). Given a set of value statements, the auctioneer must now solve the *Proxy Auction Problem* (PAP), that is, it must compute the prices and allocation that result when all of the agents follow straightforward bidding. A natural approach is to iteratively compute incremental bids until the termination criteria is reached. However, this *simulated bidding* approach has several undesirable properties. First, the outcome is dependent upon implementation details such as the tie-breaking rule, the order of bidding, and the bid increment. In fact, the randomizations described in the SCPA description were chosen to avoid biases and produce a consistent outcome over a sufficient number of iterations. Second, the accuracy is a function of the bid increment. We can improve the accuracy by decreasing the bid increment, but doing so will greatly increase the number of iterations and the amount of time the process takes. This is particularly undesirable because each iteration requires the auction solve an NP-complete WDP and compute new prices. Although some researchers [4] have begun looking at techniques for reusing previous solutions in iterative combinatorial auctions, the techniques are at an early stage of development.

We advocate a novel approach to solving PAPs rooted in the observation that the straightforward bidding policy generates predictable patterns of bidding behavior that results in complex, but computable, trajectories for bundle-prices. Figure 1 shows the trajectories taken by prices when proxy agents bid in the SCPA with the valuations shown in Table 1. The approach outlined in the next section allows us to compute the results by taking large steps similar to the way a proxy bid in eBay can directly be computed to jump to δ above the second highest bid. Moreover, we can compute an exact outcome, one that is independent of the bid increment or tie-breaking rules, and whose run time is independent of the magnitude of the bids.

3. SOLVING PAP

3.1 Concepts

²The demand set is also referred to as a *best response set* in the literature.

³For expository clarity, we leave the t superscript off the symbols when it is not ambiguous.

	A	B	AB	C	AC	BC	ABC
Buyer 1	10	3	18	2	18	10	20
Buyer 2	4	9	15	3	12	18	20
Buyer 3	1	3	11	9	16	17	25
Buyer 4	7	7	16	7	16	16	20

Table 1: An example with four buyers bidding on the combinations of three objects.

In the framework we presented previously [15], the central concept was to compute each bidder’s allocation of *attention* among the bundles. Attention represents the proportion of its bidding opportunities the agent will use to increase its offer on a bundle. We denote agent i ’s attention to bundle b as $\theta_{i,b}$. If the agent is told it is winning by the auctioneer, it will pass. Denote the proportion of time agent i spends passing as $\theta_{i,\text{pass}}$. Each agent has one unit of attention per unit of time to allocate among the bundles and the passing action. Attention could also be thought of as dollars per unit time; for the purposes of this exposition we will use a normalized value of one dollar per unit time.

Figure 1 shows the progression of the price of each bundle through the course of the proxied auction with the agents given buyer values as shown in Table 1. It should be noted that both the simulation—when δ is sufficiently small—and our new algorithm generates the price curves in the figure. The figure has obvious structure: prices increase at steady rates with occasional events that cause them to change trajectories.

The trajectory of a bundle is simply the sum of the attention being paid to it. If during a unit of time, three agents each give half their attention to bundle b , the price of b will increase by 1.5. In Figure 1, at time zero, every agent’s demand set is the singleton ABC so they spend their attention outbidding one another. However, although the price of ABC initially increases rapidly, it does not increase at slope of 4. The bidding pattern that is established during the initial period has one (randomly selected) bidder passing because it is announced as the winner of ABC, and the other three bidders increasing their bids. Thus, each agent has a .25 probability of being announced as the winner during any given round, and a .75 probability of increasing its bid because it is not winning. Four agents bidding with .75 units of attention creates the slope of 3 seen in the first 2/3 of a time unit in Figure 1. Let θ_b^t , the trajectory of bundle b at time t , be computed as

$$\theta_b^t = \sum_i \theta_{i,b}^t.$$

In order to correctly determine the price trajectories, the algorithm must keep track of the *competitive allocations* (CAs) as well as the allocation of attention to the demand sets. In the first 2/3 of a time unit in the example there are four CAs: $\{ABC, -, -, -\}$, $\{-, ABC, -, -\}$, $\{-, -, ABC, -\}$, and $\{-, -, -, ABC\}$, using a lexical ordering of the agents. Competitive allocations are simply solutions to the WDP.

More formally, let f denote a feasible allocation of the objects, that is, $f : \mathcal{J} \rightarrow \mathcal{I}$. Let F be the set of all feasible allocations. Denote agent i ’s allocation in f as f_i . The value of allocation f at time t is

$$V(f) = \sum_i r_i^t(f_i). \quad (2)$$

Let F^* denote the set of allocations that maximize (2).

The CA notion is critical because agents will pass if and only if they are a member of a CA, and they will pass as often as a CA to which they belong is selected by the auctioneer. The values of the CAs that are managed by the proxy algorithm in the example are depicted in Figure 2. The CAs on the envelope are competitive. Notice that the four allocations in which ABC is allocated begin as CAs, but at time 2/3 the allocation $\{ABC, -, -, -\}$ flatlines and agent 1 does not bid on ABC again during the duration of the auction. $\{-, ABC, -, -\}$ and $\{-, -, -, ABC\}$ also flatline when the price of ABC reaches 11 and 14, respectively. The allocation that eventually wins the auction, $\{A, BC, -, -\}$, does not even become competitive until around time 17, and remains on the envelope until time 40 at which the bidders supporting $\{-, -, ABC, -\}$ and $\{-, -, C, AB\}$ give up.

The proxy solution algorithm is an iterative process in which each step involves (1) computing the allocation of attention, and (2) computing the duration for which the former holds. In the following sections, we deal with these two issues in more detail, starting with the latter.

3.2 Computing Intervals

We use the same methodology to compute the duration of an interval described elsewhere [15], but here extend it to the general case. There are two potential events that define the end of an interval:

1. the prices of bundles reach a point where one or more bidders become attracted to one or more bundles that were not previously in their demand sets, or
2. a allocation that was not formerly competitive reaches a value that makes it competitive.

Unlike our previous work, we do not restrict these events to consist of a single allocation or a single agent and bundle. The complexity that arises when we lift this restriction is dealt with primarily in the next section.

To establish the duration of the interval starting at time t , we need the first change in demand sets among all agents; the trajectories will change at the inflection point, rendering the other computed collisions obsolete. The amount of time it will take for c to become as attractive to i as b is

$$\Delta t_i^{b,c} = \frac{v_i(b) - \pi_b^t - v_i(c) + \pi_c^t}{\theta_b^t - \theta_c^t}. \quad (3)$$

Note that cases where $\Delta t_i^{b,c}$ is negative or undefined indicate that the two trajectories do not collide in the future because they are diverging or parallel. Such pairs are excluded.

The duration of time until the next demand set change is

$$\Delta t^{\text{DS}} = \min_i \left\{ \min_{b \in D_i^t, c \notin D_i^t} \Delta t_i^{b,c} \right\} \quad (4)$$

Similarly, we need to compute the first allocation that will become competitive under the current allocation of attention. First, recognize that the trajectory of an allocation is the sum of the trajectories of the components that are actively being bid upon. Let γ_f^t be the trajectory of f , computed as

$$\gamma_f^t = \sum_{i|f_i \in D_i^t} \theta_{f_i}^t.$$

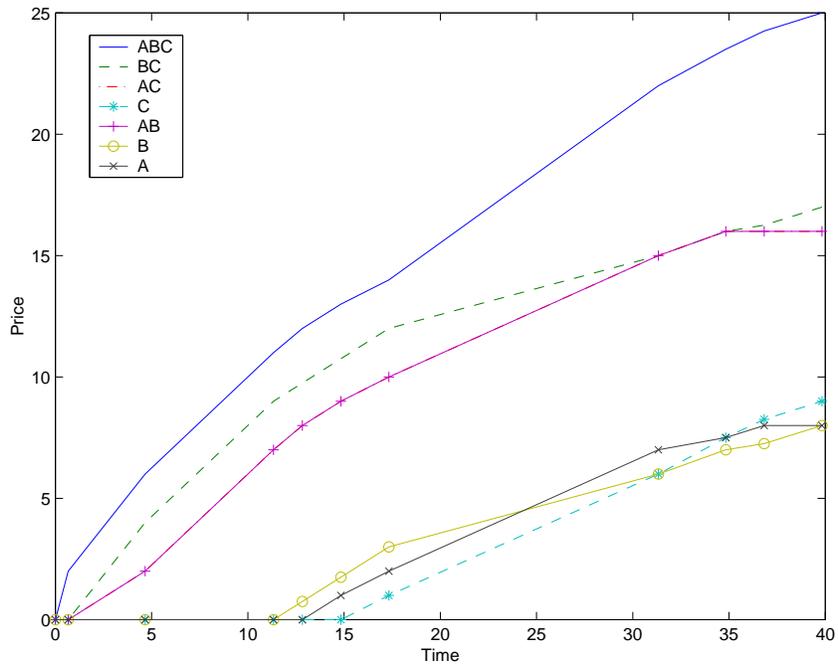


Figure 1: The prices over time when proxy agents bid on the example in Table 1.

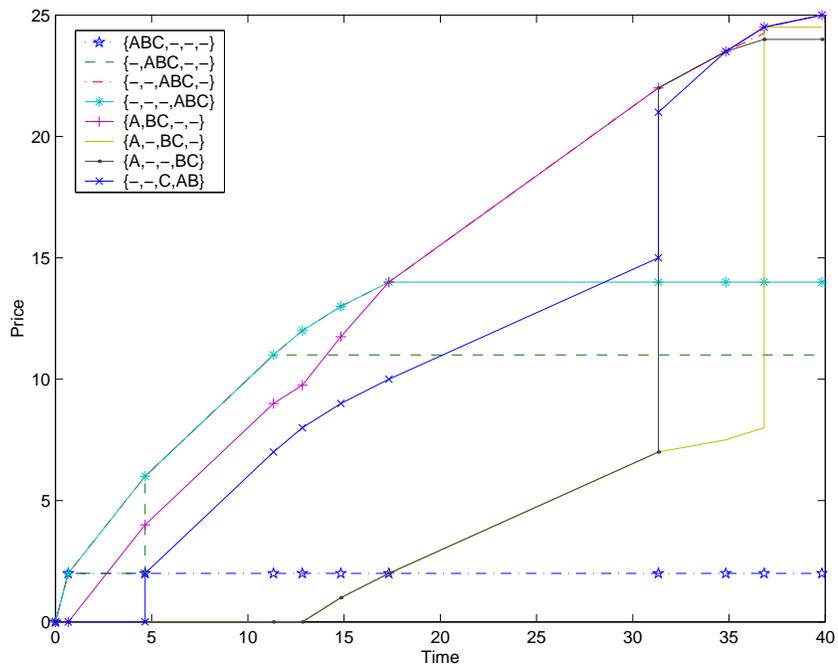


Figure 2: The values of the competitive allocations for the example in Table 1.

Competitive allocation, \mathbf{f} , will collide with non-competitive allocation, $\hat{\mathbf{f}}$, when the values of the two become equal. If the current time is t , the collision will occur in $\Delta t^{\mathbf{f},\hat{\mathbf{f}}}$ time increments, where

$$\Delta t^{\mathbf{f},\hat{\mathbf{f}}} = \frac{V^t(\hat{\mathbf{f}}) - V^t(\mathbf{f})}{\gamma_{\hat{\mathbf{f}}}^t - \gamma_{\mathbf{f}}^t}. \quad (5)$$

As in the computation of collisions in Section 3.3, a non-positive value for $\Delta t^{\mathbf{f},\hat{\mathbf{f}}}$ means that, with the current allocation of attention, $\hat{\mathbf{f}}$ will not become competitive with \mathbf{f} . Finally, among the positively valued $\Delta t^{\mathbf{f},\hat{\mathbf{f}}}$'s,

$$\Delta t^{\text{CA}} = \min_{\mathbf{f} \in F^*, \hat{\mathbf{f}} \notin F^*} \Delta t^{\mathbf{f},\hat{\mathbf{f}}}. \quad (6)$$

The next inflection point is determined by taking the minimum of Δt^{DS} and Δt^{CA} . In addition, we keep track of which allocations and combinations of agents and bundles triggered the inflection point.

3.3 Computing the Allocation of Attention

The main contribution of this paper is in defining how to compute the allocation of attention in the general case where an arbitrary number of bundles are added to demand sets or several allocations become competitive. The assumptions made in the original formulation of the algorithm [13] allowed us to simplify the maintenance of the demand set information at the inflection points. In the general case, when we add a set of bundles to the agents' demand sets, bundles that were previously in the demand sets may be dropped. It is necessary to determine which bundles are dropped in order to establish which allocations will remain competitive.

This section details a mixed integer-linear program (MILP) that simultaneously computes which allocations remain competitive, which bundles should be kept in the demand set for each agent, and how much attention each agent allocates to each item in its demand set. The domain of the MILP is defined by the *potential* demand sets and the *potential* competitive allocations. Although all of the identified bundles and allocations are instantaneously active at the inflection at time t , we are concerned with determining which will remain active going forward.

The potential demand set for each agent is the union of the bundles that were previously in the agent's demand set and those bundles being introduced by the agent at this inflection point; in other words, the set of bundles that satisfy (1) at time t . The potential CA set, F^* , is the union of the previous CAs and those allocations being introduced at this inflection point; no allocation can become competitive without being the cause of the inflection.

The integer variable $y_{i,b}$ takes the value 1 if bundle b will remain in agent i 's demand set during the following interval, and zero otherwise. The integer variable $x_{\mathbf{f}}$ takes the value one if allocation \mathbf{f} will remain competitive during the interval, and zero otherwise. Let N be a sufficiently large constant.

It is clear that $\forall i \in \mathcal{I}$,

$$y_{i,b} \geq \theta_{i,b}, \quad \forall b \in D_i \quad (7)$$

$$y_{i,b} \leq N\theta_{i,b}, \quad \forall b \in D_i \quad (8)$$

$$\theta_{i,b} = y_{i,b} = 0, \quad \forall b \notin D_i.$$

The straightforward bidding policy that each agent implements implies that if bundles b and c are both in D_i over an interval, $\theta_b = \theta_c$. It follows that when $y_{i,b} = 1$

$$\text{if } y_{i,c} = 1, \text{ then } \theta_b = \theta_c,$$

$$\text{if } y_{i,c} = 0, \text{ then } \theta_b \leq \theta_c.$$

Expressed as integer constraints, $\forall i \in \mathcal{I}, \mathbf{b}, \mathbf{c} \in D_i$,

$$\theta_c - \theta_b + Ny_{i,b} + Ny_{i,c} \leq 2N, \quad (9)$$

$$\theta_c - \theta_b + Ny_{i,c} - Ny_{i,b} \leq N. \quad (10)$$

Active agents have one unit of attention to allocate, while agents that no longer achieve positive surplus on any bundle (and therefore have empty demand sets) will allocate no attention. Let K_i be a constant used during problem construction where $K_i = 0$ if D_i is empty, and $K_i = 1$ otherwise. We conserve attention with the constraint

$$\sum_{b \in B} \theta_{i,b} + \theta_{i,\text{pass}} = K_i, \quad \forall i \in \mathcal{I}. \quad (11)$$

We now turn our attention to the constraints that capture the influence of the competitive allocations. Let $\beta_{\mathbf{f}} \in [0, 1]$ be the frequency with which allocation \mathbf{f} is announced as the winner, enabling the members of that allocation to pass. At every iteration, one of the CAs must be announced as the winning bundle. Thus,

$$\sum_{\mathbf{f} \in F^*} \beta_{\mathbf{f}} = 1. \quad (12)$$

If allocation \mathbf{f} is not competitive then $\beta_{\mathbf{f}} = 0$. In other words,

$$\beta_{\mathbf{f}} \leq x_{\mathbf{f}}. \quad (13)$$

An allocation cannot be competitive unless at least one agent is actively bidding on it.

$$x_{\mathbf{f}} \leq \sum_{i \in \mathcal{I}} y_{i,\mathbf{f}_i}, \quad \forall \mathbf{f} \in F^* \quad (14)$$

By definition, competitive allocations increase their value at the same rate, while if a potential CA turns out to be not competitive, its slope must be less than those allocations that are competitive. Consider the case where \mathbf{f} is competitive, that is, $x_{\mathbf{f}} = 1$.

$$\text{If } x_{\mathbf{f}} = 1, \text{ then } \gamma_{\mathbf{f}} = \gamma_{\hat{\mathbf{f}}};$$

$$\text{if } x_{\hat{\mathbf{f}}} = 0, \text{ then } \gamma_{\mathbf{f}} \geq \gamma_{\hat{\mathbf{f}}}.$$

A MILP form of this logical constraint is

$$\gamma_{\hat{\mathbf{f}}} - \gamma_{\mathbf{f}} + Nx_{\mathbf{f}} + Nx_{\hat{\mathbf{f}}} \leq 2N, \quad (15)$$

$$\gamma_{\mathbf{f}} - \gamma_{\hat{\mathbf{f}}} - Nx_{\mathbf{f}} + Nx_{\hat{\mathbf{f}}} \leq N. \quad (16)$$

Every competitive allocation will have some probability of being selected as the winning allocation. Let $\beta_{\mathbf{f}}$ be the frequency with which the auctioneer selects solution \mathbf{f} . It is not true that all competitive allocations are selected equally often during an interval; they are only equally likely to be chosen when they are tied. Consider the bidding pattern in a two-agent, two-item scenario where Agent 1 is bidding only on AB, and Agent 2 is alternating between bidding on A and bidding on AB. In those rounds in which Agent 2 bids on A, $\{-, AB\}$ is the only CA. In the alternate rounds, Agent 2 may increase its bid on AB, thereby tying Agent 1 and having a

0.5 probability of being declared the current winner. Thus, Agent 2 will bid on A half the time, on AB a quarter of the time, and pass the remaining quarter of the time.

This analysis suggest a connection between agent i 's behavior and the frequency with which i is not a member of the winning allocation. Agent i bids whenever one of the other competitive allocations pass. We employ another constant that is used during problem construction: let $G_{f,i} = 1$ if i is allocated a bundle in f , and $G_{f,i} = 0$ otherwise.

$$1 - \theta_{i,\text{pass}} = \sum_{f \in F^*} (1 - G_{f,i})\beta_f. \quad (17)$$

The final piece of the MILP is the objective function. We take as our objective the maximization of the number of competitive allocations.

$$\text{maximize } \sum_{f \in F^*} x_f.$$

Solving the MILP outlined above will compute the bundles that are in each agents best response set going forward (those that have $y_{i,b} = 1$), which allocations will be competitive (those for which $x_f = 1$), and how much attention each agent pays to each bundle in its demand set. From the attention we can compute the slope of the price of each bundle and the slope of the allocations, both of which are necessary to determine the duration of the interval. The complete MILP is shown in Figure 3.3.

4. WORKED EXAMPLE

The result of applying the algorithm to the example in Table 1 is shown in Figure 1. In this section, we highlight some of the steps in the process for expository purposes. Table 2 shows the computations involved at six of the ten steps involved in solving the auction. Each step shows the prices at the designated time, the potentially competitive allocations (with those that remain competitive in the interval designated with an asterisk), each agent's potential demand set, and each agent's allocation of attention. The attention columns are summed to give the trajectories of the prices going forward.

At $t = 0$, the prices are zero and ABC is the sole element in each agent's demand set. Since there are four CAs, and no other bundles distracting the bidders, each CA wins one fourth of the time. At time $2/3$, the bundles AB and AC enter Agent 1's demand set, and BC enters Agent 2's. Both agents focus their attention on the new elements, although ABC remains in Agent 2's demand set even though the agent allocates no attention to it because the slope of ABC matches the slope of BC.

We then skip to step 7 where we see a complex allocation of attention across the bundles, yet one that continues to satisfy the constraints of the problem and exactly matches the simulation pattern. It is also interesting to note that the inflection at step 7 is caused by the allocation $\{A, BC, -, -\}$ entering the CA set. The effect of this is clearly visible in Figure 2. At step 8, we see an interesting effect when BC enters Agent 4's demand set. The price of BC at the time is 15, so when Agent 4 suddenly joins the bidding, the allocation $\{A, -, -, BC\}$ suddenly becomes competitive, as seen by the dramatic rise in its value. At step 9, we have our first agent drop out because no bundles provide positive surplus, and we see another inflection caused by the allocation $\{-, -, C, AB\}$. It is also interesting that the new competitive allocation includes Agent 4 even though

it has stopped bidding. At step 11, Agent 3 drops out when it can no longer achieve positive surplus, and the allocation $\{A, BC, -, -\}$ wins the auction.

5. DISCUSSION AND FUTURE WORK

Aside from research on algorithms to solve singular instances of the WDP [1, 3, 12], there is little research that studies the computation over the life of an iterative combinatorial auction. One exception is the recent work by Kastner, et al. [4], in which the authors study the costs of maintaining previous solutions to the WDP as bidders incrementally increase their bids. While relevant to scenarios in which bidders have unconstrained bidding strategies, when proxy bidding is allowed, maintaining previous solutions to the WDP is unnecessary given the algorithm outlined in this paper.

We intend to continue with this line of research and to expand the framework to encompass the other auctions in the literature. For example, the SCPA presented in this paper differs from the Ausubel and Milgrom's APA mechanism [2] principally in the amount of attention agents are permitted to allocate. In APA, the bidders raise their offer on *all* elements of their best-response set, whereas in our mechanism the bidder randomly selects one element of its demand set on which to bid. We feel there is a strong correspondence between SCPA and AIBA [14], and posit that AIBA's distinctive method of price determination can be applied *ex-post*. We would like to try to represent at least the anonymous price variations of iBundle [7] in the framework, and it may be possible to capture the discriminatory mechanisms.

We also plan to study the computational complexity of the process. We expect the MILP to be NP-complete, and at least as hard as the WDP. While we clearly benefit by solving for a limited number of inflection points, the value of the overall approach will depend upon how much harder it is to compute the allocation of attention than to solve the WDP. At the same time, we have made no effort yet to improve the tractability of the problem, and our experience so far suggests that there are many ways in which to improve the computational costs. For example, we described the interval computations as a comparison between all combinations of elements in a set with those not in the set. With some clever bookkeeping, a much smaller number of comparisons need to be made. There are also obvious ways in which a computer program could avoid creating some instances of the constraints in the MILP. Because these optimizations obfuscate the central ideas in the approach, we left them out.

Acknowledgments

This research was funded by NSF CAREER award 0092591 and the E-Commerce program NC State. The views and conclusions contained herein are those of the authors, as are any errors. We thank Ashish Sureka, who helped with the earlier version of this algorithm, and David Parkes, for his encouragement and suggestions.

6. REFERENCES

- [1] A. Andersson, M. Tenhunen, and F. Ygge. Integer programming for combinatorial auction winner determination. In *Fourth International Conference on Multiagent Systems*, pages 39–46, 2000.
- [2] L. M. Ausubel and P. R. Milgrom. Ascending auctions with package bidding. *Frontiers of Theoretical Economics*, 1(1):1–42, 2002.

$$\left. \begin{aligned}
& \max \sum_{f \in F^*} x_f \\
& \text{s.t. } y_{i,b} \geq \theta_{i,b}, \quad \forall i \in \mathcal{I}, \mathbf{b} \in D_i \\
& y_{i,b} \leq N\theta_{i,b}, \quad \forall i \in \mathcal{I}, \mathbf{b} \in D_i \\
& \sum_{i \in \mathcal{I}} \theta_{i,b} - \sum_{i \in \mathcal{I}} \theta_{i,c} + Ny_{i,b} + Ny_{i,c} \leq 2N, \quad \forall i \in \mathcal{I}, \mathbf{b}, \mathbf{c} \in D_i \\
& \sum_{i \in \mathcal{I}} \theta_{i,b} - \sum_{i \in \mathcal{I}} \theta_{i,c} + Ny_{i,b} - Ny_{i,c} \leq N, \quad \forall i \in \mathcal{I}, \mathbf{b}, \mathbf{c} \in D_i \\
& \beta_f \leq x_f, \quad \forall f \in F^* \\
& x_f \leq \sum_{i \in \mathcal{I}} y_{i,f_i}, \quad \forall f \in F^* \\
& \sum_{f \in F^*} \beta_f = 1 \\
& \sum_{i \in \mathcal{I}} \sum_{j \in \mathcal{I}} \theta_{j,f_i} - \sum_{i \in \mathcal{I}} \sum_{j \in \mathcal{I}} \theta_{j,\hat{f}_i} + Nx_f + Nx_{\hat{f}} \leq 2N, \quad \forall f, \hat{f} \in F^* \\
& \sum_{i \in \mathcal{I}} \sum_{j \in \mathcal{I}} \theta_{j,f_i} - \sum_{i \in \mathcal{I}} \sum_{j \in \mathcal{I}} \theta_{j,\hat{f}_i} - Nx_f + Nx_{\hat{f}} \leq N, \quad \forall f, \hat{f} \in F^* \\
& (1 - \theta_{i,\text{pass}}) = \sum_{f \in F^*} (1 - G_{f,i})\beta_f \quad \forall i \in \mathcal{I} \\
& \theta_{i,\text{pass}} + \sum_{\mathbf{b} \in B} \theta_{i,b} = K_i, \quad \forall i \in \mathcal{I} \\
& \theta_{i,\text{pass}} \geq 0, \\
& \beta_f \geq 0, \quad x_f \in \{0, 1\}, \quad \forall f \in F^* \\
& \theta_{i,b} \geq 0, \quad y_{i,b} \in \{0, 1\}, \quad \forall i \in \mathcal{I}, \mathbf{b} \in D_i
\end{aligned} \right\}$$

Figure 3: The complete mixed-integer program for computing the allocation of attention in the proxy auction.

- [3] Y. Fujishima, K. Leyton-Brown, and Y. Shoham. Taming the computational complexity of combinatorial auctions: Optimal and approximate approaches. In *Sixteenth International Joint Conference on Artificial Intelligence*, pages 548–553, 1999.
- [4] R. Kastner, C. Hsieh, M. Potkonjak, and M. Sarrafzadeh. On the sensitivity of incremental algorithms for combinatorial auctions. In *IEEE International Workshop on Advanced Issues of E-Commerce & Web-Based Information Systems (WECWIS)*, 2002.
- [5] P. Milgrom. Putting auction theory to work: The simultaneous ascending auction. *Journal of Political Economy*, 108(2):245–272, 2000.
- [6] N. Nisan. Bidding and allocation in combinatorial auctions. In *Second ACM Conference on Electronic Commerce*, pages 1–12, 2000.
- [7] D. C. Parkes. *iBundle*: An efficient ascending price bundle auction. *First ACM Conference on Electronic Commerce*, pages 148–157, November 1999.
- [8] D. C. Parkes. *Iterative Combinatorial Auctions*, chapter 3, page to appear. 2003.
- [9] D. C. Parkes and L. H. Ungar. Preventing strategic manipulation in iterative auctions: Proxy agents and price-adjustment. In *Seventeenth National Conference on Artificial Intelligence*, pages 82–89, 2000.
- [10] J. Rodríguez, P. Noriega, C. Sierra, and J. Padget. FM96.5 a java-based electronic auction house. In *Second International Conference on the Practical Application of Intelligent Agents and Multi-Agent Technology (PAAM '97)*, 1997.
- [11] A. E. Roth and A. Ockenfels. Last-minute bidding and the rules for ending second-price auctions: Evidence from eBay and Amazon auctions on the Internet. *American Economic Review*, 92(4):1093–1103, June 2002.
- [12] T. Sandholm, S. Suri, A. Gilpin, and D. Levine. Cabob: A fast optimal algorithm for combinatorial auctions. In *International Joint Conference on Artificial Intelligence*, Seattle, WA, 2001.
- [13] P. R. Wurman, G. Cai, J. Zhong, and A. Sureka. An algorithm for computing the outcome of combinatorial auctions with proxy bidding. In *Fifth International Conference on Electronic Commerce*, pages 1–8, Pittsburgh, PA, 2003.
- [14] P. R. Wurman and M. P. Wellman. AkBA: A progressive, anonymous-price combinatorial auction. In *Second ACM Conference on Electronic Commerce*, pages 21–29, 2000.
- [15] P. R. Wurman and M. P. Wellman. Nonlinear pricing in combinatorial auctions. Technical report, North Carolina State University, November 2002.
- [16] P. R. Wurman, M. P. Wellman, and W. E. Walsh. A parametrization of the auction design space. *Games and Economic Behavior*, 35(1-2):304–338, April-May 2001.

Step 1, $t = 0$			Prices:							
$f V(f) = 0$	Agent	D_i	A	B	AB	C	AC	BC	ABC	pass
*{ABC, -, -, -}	1	ABC							0.75	0.25
*{-, ABC, -, -}	2	ABC							0.75	0.25
*{-, -, ABC, -}	3	ABC							0.75	0.25
*{-, -, -, ABC}	4	ABC							0.75	0.25
θ_b			0	0	0	0	0	0	3	

Step 2, $t = 2/3$			Prices:							
$f V(f) = 2$	Agent	D_i	A	B	AB	C	AC	BC	ABC	pass
{ABC, -, -, -}	1	AB, AC, ABC			0.5		0.5			
*{-, ABC, -, -}	2	BC, ABC						1		
*{-, -, ABC, -}	3	ABC							0.5	0.5
*{-, -, -, ABC}	4	ABC							0.5	0.5
θ_b			0	0	0.5	0	0.5	1	1	

⋮

Step 7, $t = 17\frac{1}{3}$			Prices:							
$f V(f) = 14$	Agent	D_i	A	B	AB	C	AC	BC	ABC	pass
*{A, BC, -, -}	1	A, AB, AC	5/14				1/14			4/7
{-, ABC, -, -}	2	B, BC		3/14				3/14		4/7
*{-, -, ABC, -}	3	ABC							4/7	3/7
{-, -, -, ABC}	4	AB, C, AC, ABC			5/14	5/14	4/14			
θ_b			5/14	3/14	5/14	5/14	5/14	3/14	4/7	

Step 8, $t = 31\frac{1}{3}$			Prices:							
$f V(f) = 22$	Agent	D_i	A	B	AB	C	AC	BC	ABC	pass
*{A, BC, -, -}	1	A, AB, AC	1/7							6/7
*{A, -, -, BC}	2	B, BC		1/7				1/7		5/7
*{-, -, ABC, -}	3	C, ABC				3/7			3/7	1/7
	4	B, AB, C, AC, BC		1/7	2/7		2/7	1/7		1/7
θ_b			1/7	2/7	2/7	3/7	2/7	2/7	3/7	

Step 9, $t = 34\frac{5}{6}$			Prices:							
$f V(f) = 23.5$	Agent	D_i	A	B	AB	C	AC	BC	ABC	pass
*{A, BC, -, -}	1	A	1/4							3/4
{A, -, -, BC}	2	B, BC		1/8				1/8		3/4
*{-, -, ABC, -}	3	C, ABC				3/8			3/8	1/4
*{-, -, C, AB}	4	\emptyset								
θ_b			1/4	1/8	0	3/8	0	1/8	3/8	

Step 10, $t = 36\frac{5}{6}$			Prices:							
$f V(f) = 24.25$	Agent	D_i	A	B	AB	C	AC	BC	ABC	pass
*{A, BC, -, -}	1	A, AB, AC								1
{A, -, -, BC}	2	B, BC		1/4				1/4		1/2
*{-, -, ABC, -}	3	C, BC, ABC				1/4			1/4	1/2
*{-, -, C, AB}	4	\emptyset								
θ_b			0	1/4	0	1/4	0	1/4	1/4	

Step 11, $t = 39\frac{5}{6}$			Prices:							
	Agent	D_i	A	B	AB	C	AC	BC	ABC	pass
	1	A, AB, AC								
	2	B, BC								
	3	C, BC, ABC								
	4	\emptyset								
θ_b			8	8	16	9	16	17	25	

Agent 3 stops bidding and the allocation {A, BC, -, -} wins.

Table 2: Some of the steps in the computation of the auction result for the example in Table 1. Asterisks indicate which of the potential CAs are determined to be competitive.