Preference Elicitation in Proxied Multiattribute Auctions

Aditya V. Sunderam*   David C. Parkes†

ABSTRACT

We consider the problem of minimizing preference elicitation in efficient multiattribute auctions, that support dynamic negotiation over non-price based attributes such as quality, time-of-delivery, and processor speed. We introduce asynchronous price-based multiattribute auctions, with proxy bidding agents that sit between the auctioneer and the participants. Empirical results demonstrate the preference elicitation savings that are provided with minimal price spaces, asynchronous updates, and proxy agents.

Categories and Subject Descriptors

F.2 [**Theory of Computation**]: Analysis of Algorithms and Problem Complexity; J.4 [**Computer Applications**]: Social and Behavioral Sciences—*Economics.*

General Terms

Algorithms, Economics.

1. INTRODUCTION

Electronic auction mechanisms are becoming increasing popular as mechanisms for procurement, sourcing, and logistics in the supply chain. Multiattribute auctions [1] extend the traditional auction setting to allow negotiation over non-price attributes, such as quality, delivery time, color, speed, etc.

There are many markets in which preference elicitation is costly for participants, for example when participants must consider alternative business plans, collect additional information, or solve hard computational problems in order to refine their value for different outcomes. Iterative auctions are useful in these settings, because they allow participants to consider the accuracy to which they should refine their values, and in which parts of the outcome space to focus, all in response to feedback about the bids from other participants [4]. Preference elicitation has previously been considered in the context of iterative combinatorial auctions (e.g. [5, 2]). In this paper, we examine the preference elicitation properties of iterative multiattribute auctions, and in particular we consider the effect that the size of the price space, asynchronous updates, and proxy bidding agents can have on preference elicitation.

In the multiattribute allocation problem (MAP) there are n sellers, one buyer, and m attributes. Let $\mathcal{I}$ denote the set of sellers, and $\mathcal{J}$ denote the set of attributes. Each attribute, $j \in \mathcal{J}$, has a domain of possible attribute values (or *levels*), denoted with abstract set Θ_j; for example $\Theta_1 = \{\text{red,yellow,green}\}$ if attribute 1 is the color of an item. The joint domain, across all attributes, is denoted $\Theta = \Theta_1 \times \ldots \times \Theta_M$. Each seller, $i \in \mathcal{I}$, has a *cost function*, $c_i(\theta) \geq 0$, for an attribute bundle, $\theta \in \Theta$, and the buyer has a *valuation function*, $v(\theta) \geq 0$. The *efficient outcome* (θ^*, i^*) solves $\max_{\theta \in \Theta, i \in \mathcal{I}} [v(\theta) - c_i(\theta)]$. We assume *quasilinear* utility functions, with seller i receiving utility $u_i(\theta, p) = p - c_i(\theta)$ for selling bundle θ at price p, and the buyer receiving utility $u(\theta, p) = v(\theta) - p$.

We focus on the preferential independence (PI) special case in the experimental analysis [3]. With PI, a seller cost function is expressed as $c_i(\theta) = \sum_{j \in \mathcal{J}} w_{ij} f_{ij}(\theta_j)$, where θ_j is the level of the jth attribute, $w_{ij} \geq 0$ is the weight of that attribute, and $f_{ij}(\theta_j)$ is the marginal cost function for attribute j. Similarly, the buyer's valuation is expressed as $v(\theta) = \sum_{j \in \mathcal{J}} w_j f_j(\theta_j)$, where w_j and f_j are the weights and functions for the buyer for attribute j.

In experiments, we assume knowledge of the marginal functions of the sellers and buyer, but not the weights. With this, partial information about preference information is captured with a convex simplex of feasible weights. Every time a participant provides information, via bids, we add new constraints to this simplex. The residual volume, calculated using a simple Monte Carlo method, and normalized by taking the mth root (with m attributes) is used to measure the information revelation. A small volume indicates little uncertainty, and a high degree of information revelation.

2. THE AUCTION MECHANISMS

The iterative multiattribute auctions, NONLINEAR&DISCRETE (NLD) and ADDITIVE & DISCRETE (AD), introduced in Parkes & Kalagnanam [6], can be extended to include mandatory proxy agents [7], that sit between the auction and the participants. The role of these proxy agents is to maintain partial information about preferences, based on revealed preferences in bids. The proxy agents submit bids automatically (following an equilibrium strategy) whenever there is enough information. Otherwise, the proxy agents go back to participants and request additional preference information.

Auction NLD maintains nonlinear prices, $p(\theta)$, on attribute bundles. Auction AD, designed for the PI special case, maintains linear prices, $p(j,k)$, on level $k \in \Theta_j$ of attribute j, and an additional price penalty term, Δ. The overall price in AD is defined as $p(\theta) = \sum_{j \in \mathcal{J}} p(j, \theta_j) - \Delta$. The auctions proceed in rounds, maintaining a provisional allocation and decreasing ask prices by a

*Harvard College, Cambridge, MA 02138. Email: sunderam@fas.harvard.edu. Supported in part by the Harvard College Research Fund.

†Corresponding author. Division of Engineering and Applied Sciences, 33 Oxford Street, Harvard University, Cambridge, MA 02138. Email: parkes@eecs.harvard.edu. Supported in part by NSF grant IIS-0238147. Thanks to Jayant Kalagnanam for many useful discussions related to this work.

EC'03, June 9–12, 2003, San Diego, California, USA.
ACM 1-58113-679-X/03/0006.

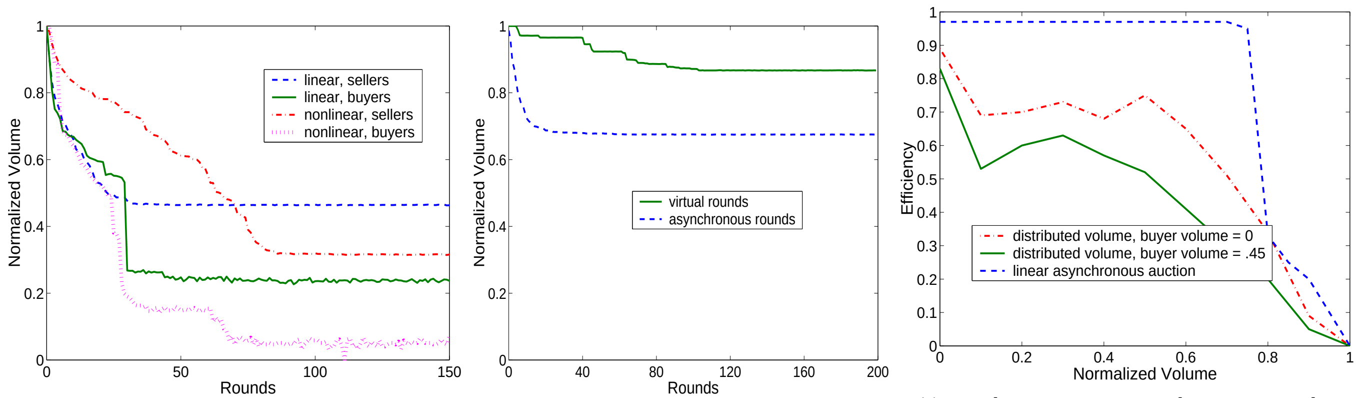

(a) synchronous NLD vs. synchronous AD

(b) asynchronous AD vs. virtual-round AD

(c) asynchronous AD vs. asynchronous AD with a static preference elicitation policy

Figure 1: Total preference elicitation (measured in terms of residual volume in feasible preference space) vs. Number of rounds in the auction. Careful auction design reduces preference elicitation while maintaining economic efficiency (a, b), and improves efficiency for constrained preference elicitation (c).

minimal bid increment ϵ based on unsuccessful bids, with dynamic feedback from the buyer about which bids she currently prefers. The auctions terminate when they reach quiescence.

The auctions have several useful theoretical properties. For instance, myopic best-response (MBR), in which sellers respond to prices in each round by submitting a bid for the bundle with maximal utility at the price, is a sequential equilibrium strategy in both auctions against a reasonable class of buyer strategies [6]. Simple *asynchronous* variations are also defined, with the provisional allocation, and prices, updated whenever another bid is received.

In addition to verifying the informational benefits of iterative multiattribute auctions, we identify a useful interaction between proxy agents and asynchronous auctions. The auction can continuously *poll* the proxy agents at random, until no proxy is able to bid without additional preference information. Only then does one of the proxy agents fall back and collect additional preference information. In the current context, we assume that query is simply "what bundle will you bid for at the current prices?" This is precisely the information that is necessary and sufficient to make progress in the auction. This rapid interchange between proxies and the auction, until all information is exhausted, is termed a *virtual round*.

3. EXPERIMENTS

We focus our empirical study on the simplified case in which the buyer functions, $f_j(\theta_j)$, and the seller functions, $c_{ij}(\theta_j)$, are *known* to the proxy agents. Partial information about preferences is represented as a convex simplex in weight space, with additional constraints on weights inferred from the bids submitted by sellers. As discussed in the introduction, information revelation is measured as the residual normalized volume of this simplex. The model used to generate valuations and costs, along with complete experimental results, are presented in the longer version of this paper [8]. By default in all experiments, the number of sellers, attributes, and attribute levels is five, and results are averaged over 20 trials.

Figure 1 illustrates some results. From (a), observe that the linear synchronous auction (AD) requires less information revelation than the nonlinear synchronous auction (NLD), for both sellers and the buyer. Both auctions are efficient in this PI setting, but the price space is minimal in AD, leading to more effective preference elicitation. Comparing (a) with (b), observe that the asynchronous AD auction requires less information revelation than the synchronous AD auction. Also, notice a further saving in preference elicitation in moving to the proxied virtual-rounds variation.[1]

Finally, in (c) we plot the efficiency of the asynchronous AD auction for different levels of total information revelation, as the bid increment ϵ is varied. For comparison, for each level of information revelation we also ran the asynchronous AD auction with an *ex ante* fixed preference elicitation strategy. In particular, for each seller we grow a symmetric simplex around the true weight vector, until the volume equals the average seller volume in the iterative auction. Adaptive preference elicitation is seen to have a significant efficiency advantage in this setting, with constrained preference elicitation.

4. REFERENCES

[1] Y. K. Che. Design competition through multidimensional auctions. *RAND Journal of Economics*, 24:668–680, 1993.

[2] W. Conen and T. Sandholm. Preference elicitation in combinatorial auctions. In *Proc. 3rd ACM Conf. on Electronic Commerce (EC-01)*, pages 256–259, 2001.

[3] R. L. Keeney and H. Raiffa. *Decision Making with Multiple Objectives: Preferences and Value Tradeoffs*. CUP, 1993.

[4] D. C. Parkes. Optimal auction design for agents with hard valuation problems. In *Proc. IJCAI-99 Workshop on Agent Mediated Electronic Commerce*, pages 206–219, July 1999. Stockholm.

[5] D. C. Parkes. Price-based information certificates for minimal-revelation combinatorial auctions. In *Agent Mediated Electronic Commerce IV*, LNAI 2531, pages 103–122, 2002.

[6] D. C. Parkes and J. Kalagnanam. Iterative Multiattribute Vickrey Auctions. Submitted for publication, 2002.

[7] D. C. Parkes and L. H. Ungar. Preventing strategic manipulation in iterative auctions: Proxy agents and price-adjustment. In *Proc. 17th Nat. Conf. on Artificial Intelligence (AAAI-00)*, pages 82–89, July 2000.

[8] A. V. Sunderam and D. C. Parkes. Preference elicitation in proxied multiattribute auctions. In *Fourth ACM Conf. on Electronic Commerce (EC'03)*, 2003. Shorter version. Extended version at http://www.eecs.harvard.edu/˜parkes/pubs/proxy.pdf.

[1] In total, the residual volume of information not revealed by an average seller changes from 31.7% to 46.3% to 77.0% to 86.9%, moving from synchronous NLD to synchronous AD to asynchronous AD to proxied/virtual-rounds AD.