Reducing Costly Information Acquisition in Auctions

Kate Larson
School of Computer Science
University of Waterloo
Waterloo, Ontario, Canada
klarson@cs.uwaterloo.ca

ABSTRACT

Most research on auctions assumes that potential bidders have private information about their willingness to pay for the item being auctioned, and that they use this information strategically when formulating their bids. In reality, bidders often have to go through a costly information-gathering process in order to learn their valuation for the item being auctioned. Recent attempts at modelling this phenomena has brought to light complex strategic behavior arising from information-gathering, and has shown that traditional approaches to auction and mechanism design are not able to overcome it. In this paper, we show that if the auction designer has some information about the agents' information-gathering processes, then it is possible to create an auction where, in equilibrium, agents have incentive to only gather information on their own valuation problems and to reveal the results truthfully to the auctioneer. Additionally, simulation results show that, from a system-level perspective, the overall cost of information acquisition is substantially lower in this new auction when it is compared to a classic auction mechanism.

Categories and Subject Descriptors: I.2.11 Multiagent Systems
General Terms: Design, Theory
Keywords: Mechanism Design, Game Theory, Resource-bounded Agents

1. INTRODUCTION

Most research on auctions assumes that the problem facing bidders is how to bid given their private preferences for the item(s) being auctioned. In reality, however, bidders often have to go through a costly information-gathering process in order to determine what their actual preferences are. This process may involve such things as asking questions about a product in order to learn whether it is of high or low quality, solving optimization problems in order to ensure that the bidder bids on the minimum amount of materials needed to complete some job, or conducting research on a firm in order to learn whether its assets complement the assets currently held by the potential bidder.

Recently, researchers have started studying how computational and informational constraints influence both the auctioneer and the bidders. One direction of research has focused on what happens if the auctioneer does not have infinite computational powers. In many interesting auction domains, such as combinatorial auctions, the auctioneer is required to solve (possibly multiple) $\mathcal{NP}$-hard problems. Determining how to replace it with approximation algorithms, while still maintaining the desirable game-theoretic properties, as well as characterizing the domains of truthful social choice functions, has been a vibrant research area [5, 11, 13, 10, 14].

Researchers have also started studying computational and informational constraints faced by bidders. Most of the auction literature has assumed that the bidders know how they value the items being auctioned. However, in many settings bidders do not have adequate information in order to formulate their valuations and instead must learn them by computing or acquiring information (at some cost). One model that has been studied allows a bidder the choice between participating in the auction without knowing its true valuations or paying a fee to learn them. Questions asked using this model include what sort of incentives are required for bidders to acquire information about their valuations [1], and how does information acquisition depend on the rules of the auction [3, 15, 16].

It has also been noted that a bidder's decision as to whether to compute or gather information about it's valuations can depend on the preferences of others [18]. We have proposed explicitly modeling the information-gathering actions of bidders along with the decisions they make when deciding how to use their information-gathering resources. We placed this *deliberative-agent* model into a game theoretic framework, and have analyzed classic auctions in order to gain an understanding of the impact that computational and information-gathering constraints have on agents' strategic behavior [7]. This brought to light new forms of strategizing on the part of the bidders where they actively gathered information on their competitors, and provided a possible explanation for behavior seen in practice [17].

In recent work we studied the problem of designing auctions specifically for bidders who must gather information [9] in order to determine their actual valuations and formulate their bids. We proposed three intuitive properties for auctions. We argued that the auction should not have to know

AAMAS'06 May 8–12 2006, Hakodate, Hokkaido, Japan.

the details of the bidders information-gathering processes (*preference-formation independence*), bidders should not have incentive to gather information about other bidders (*deliberation-proof*), and that the auction should support some level of honesty among the bidders (*non-misleading*). However, we showed that it is impossible to design an (interesting) auction that has these three properties.[1] The contribution of this paper is that we show by that relaxing one of the properties, preference-formation independence, and allowing the auction to use some minimal amount of information about the bidders, it is possible to avoid strategic-deliberation and misleading behavior.

The rest of the paper is organized as follows. In the next section we describe our deliberative-bidder model, and discuss the problems faced in designing auctions for such bidders. In Section 3 we describe an optimal-search procedure and discuss how it can be incorporated into an auction. We then show that this auction provides the appropriate incentives for deliberative bidders. In particular it reduces the strategic burden placed on these bidders as they decide how to use their (limited) deliberation resources. Finally, Section 5 reports on experimental results which illustrate the feasibility of running such a procedure, as well as the effects it has on reducing the total cost of information gathering.

2. DELIBERATIVE BIDDING AGENTS

In this section we provide the background and motivation for the rest of the paper. We start by presenting our model of a deliberative agent. We discuss the implications of having deliberative agents participate in auctions and, in particular, we highlight the complex strategic behavior that may arise when bidders are deliberative. We then discuss some of the challenges in designing auctions for such agents.

In the rest of this paper we assume that the reader has a basic understanding of elementary game theory and auction theory. While we do explain key concepts as they are needed, a more thorough treatment on this topic can be found in microeconomic texts or books specializing on auctions [12, 6]. We pause now to comment on the terminology used in this paper. We use the terms "bidder" and "agent" interchangeably. That is, every agent is a bidding agent. We also use the terms "deliberative" and "information-gathering" interchangeably.

2.1 Deliberative agents

In this paper we assume that bidders are *deliberative*. A deliberative bidder is one who must compute or gather information in order to determine how much it values the item(s) being auctioned, has restrictions on its computing or information-gathering capabilities, and who carefully considers how to use its available resources given its restrictions.

We assume that a deliberative agent has a set of deliberation resources, which we will simply refer to as the agent's resources. We denote the resources of agent i by R_i. An agent is able to apply its resources to any preference-determination problem it wishes. If there are m possible problems, then we let $(r_1, \ldots, r_m) \in R_i^m$ denote the situation where the agent has devoted r_j resources to problem j. In particular, the agent is allowed to deliberate on its own preference-formation problems, as well as the problems of any other agent.

We model deliberation-resource limitations through cost functions. The cost function of agent i is $\text{cost}_i : R_i^m \to \mathbb{R}^+$. The only restriction placed on the cost functions are that they must be additive and non-decreasing.

A deliberative agent i is endowed with a multi-set of algorithms $\mathcal{A}_i = \{A_i^j\}$ where A_i^j is the algorithm agent i can use on problem j. We use the term *algorithm* in its broadest sense; algorithms are step-by-step procedures for solving some problem. In particular, we include information gathering processes in our set of algorithms. The algorithms of a deliberative agent have the anytime property; they can be stopped at any point and are guaranteed to return a solution, but if additional resources are allocated to the problem then a better solution (or at least no worse a solution) will be returned.

A deliberative agent carefully decides how to allocate its resources on algorithms given its cost function. To help with this process, deliberative agents are equipped with a set of *performance profiles*, $\mathcal{PP}_i = \{P_i^j\}$, one performance profile for each algorithm. A performance profile has two components. First, it describes how allocating resources to a problem changes the output of the algorithm. In particular, it describes, for any resource allocation to a problem, what possible solutions the algorithm may return, conditional on any and all features of the algorithm and problem instance which are deemed to be of importance to the agent. This is coupled with a procedural component, which, given the descriptive part of the performance profile, returns a deliberation policy that describes how the agent should optimally allocate deliberation resources. A performance profile can be as simple as a probability distribution over possible valuations coupled with a rule that states whether the agent should deliberate or not, or it may use a more sophisticated procedure that allows agents to condition their deliberation policies on current results [8].

To summarize, a deliberative agent is defined by

$$\langle R_i, \text{cost}_i(\cdot), \mathcal{A}_i, \mathcal{PP}_i \rangle$$

where R_i is the set of deliberation resources of agent i, $\text{cost}_i : R_i^m \to \mathbb{R}$ is a cost function which limits the amount of resources the agent can use, $\mathcal{A}_i$ is the multi-set of algorithms available to the agent, and $\mathcal{PP}_i$ is the associated set of performance profiles.

2.2 Auctions and Deliberative Bidders

A deliberative bidder needs to decide how and whether to gather information on valuations and then how to use the information it has in order to bid. Figure 1 illustrates this process. Motivating the bidders is the desire to maximize their utility. In particular, the utility of a deliberative bidding agent depends on its valuation, whether or not it has won the item that was up for auction, and the price that it pays if it was the winner. That is, if an agent i had used resources $r = (r_1, \ldots, r_m)$, had determined valuation $v_i(r_i)$ had won the item, and had to pay price p, then its (expected) utility would be

$$u_i = v_i(r_i) - p - \text{cost}_i(r).$$

If the agent had not won the auction then its utility is

$$u_i = -\text{cost}_i(r).$$

[1] Auctions where the outcome is completely independent of the bidding strategies of the agents trivially satisfy the three properties.

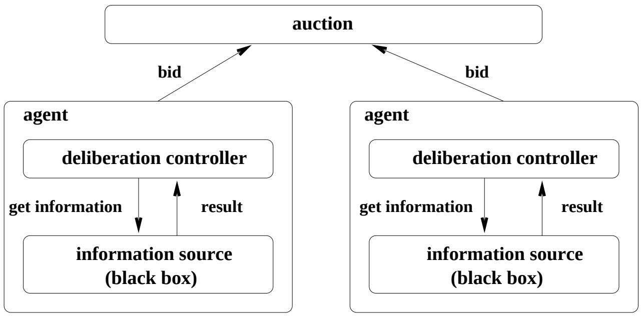

Figure 1: An auction with two deliberative bidders. Each bidder must gather information in order to determine how much it values the item that is being auctioned. A bidder may also have incentive to gather information on other participants in the auction.

Note that the bidder may use its deliberation resources on problems which do not directly affect its own valuation. However, this will still influence its utility since it incurs a cost for the total amount of resource usage.

A strategy for a deliberative bidder is a policy which specifies what actions (deliberative and bidding) to execute at every stage in the auction. We define a *history* at stage t, $H(t) \in \mathcal{H}(t)$, as the set which includes all actions (both deliberative and bidding) that the bidder itself has taken up to stage t, the results of the deliberation actions, as well as all actions other bidders may have taken (these may not have been observed by the agent). A (deliberation) *strategy* is a mapping from the set of histories to the set of actions (deliberative or bidding) for each stage in the game. That is, $S_i = (\sigma_i^t)_{t=0}^{\infty}$ and

$$\sigma_i^t : \mathcal{H}(t) \mapsto A_i$$

where A_i is the set of actions available to bidder i.

To clarify this definition we present a simple example. In a direct auction, bidders submit their bid directly to the auctioneer at some deadline T. The strategies of a deliberative bidder have the following form; $S_i = (\sigma_i^t)_{t=0}^{\infty}$ where

$$\sigma_i^t(H_i(t)) = \begin{cases} d_i^j & \text{if } t < T \text{ or } t > T \\ \hat{v} & \text{if } t = T \text{ and } \hat{v} \in \mathbb{R} \end{cases}$$

where $\hat{v}$ is the bid submitted to the auctioneer, and d_i^j is a deliberative action.

In this new, enlarged, strategy space we look for equilibria, which we call *deliberation equilibria*. When we studied the deliberation equilibria in standard auctions, we noted that the equilibrium strategies of the bidders could be very complex. In particular, deliberative bidders have incentive to deliberate on the valuation problems of other bidders, and to base their own deliberation actions on the results [7]. We call this behavior *strategic deliberation*.

We believe that strategic deliberation is a undesirable phenomenon. When bidders strategically deliberate they use their own costly information-gathering resources on valuation problems that do not directly relate to their own preferences, thus wasting these resources. Additionally, the strategic burden placed on the bidders is potentially overwhelming. They must take into consideration the bidding actions of competitors, the deliberation actions of competitors, as well as the results of their competitors deliberation actions, in order to determine how to best deliberate and bid for themselves.

Recently, we studied the problem of designing auctions explicitly for deliberative bidders [9]. To this end we proposed three desirable properties for such auctions. First, we argued that a well-designed auction should reduce the strategic burden of the bidders. In particular, a well-designed auction should not provide any incentive for bidders to strategically deliberate. That is, auctions should be *deliberation-proof*. Second, we proposed that an auction should encourage bidders to take actions which are consistent with their actual valuations. That is, bidders should not try to willfully mislead others concerning their valuations. We called this property *non-misleading*. Finally, we argued that an auction should not be involved in the information-gathering or deliberative process of the bidders. In particular, the auction should not depend on the actual form of the cost functions, anytime algorithms or performance profiles of the bidders. We called this property *preference-formation independence*.

While each of these properties is desirable in isolation, unfortunately it is impossible to design an (interesting) auction with all three [9]. In the rest of this paper we relax the preference-formation independent property and allow the auction to use some information about the information-gathering processes of the bidders. We show that by relaxing this one property it is possible to have an auction where the other two properties are achieved.

3. AN OPTIMAL-SEARCH AUCTION

The preference-formation independence property is very strong since it means that the auction is not able to use any information about the information-gathering processes of the bidders. In this section we slightly relax this property and allow the auction designer to use some information about the cost functions and performance profiles of the bidders. In Section 4 we show that this is enough to get a deliberation-proof and non-misleading auction.

The key insight is that if the auction designer has information about the bidders' cost functions and performance profiles, then it can use this information to sort the bidders. The auction can then process the bidders in order, getting each one in turn to gather information and bid until some criterion has been met. By using an optimal search procedure from the operations research literature [19], coupled with carefully designed reservation prices, it is possible to create an auction where the most promising bidders (i.e. the ones most likely to have a high valuation and thus win the auction) are asked to participate early in the process, and, secondly, where bidders are provided with enough information that they no longer have incentive to gather information about others.

This is not the first time that someone has proposed using a search procedure in an auction. In particular, Burguet [2] and Cremer, Spiegel and Zheng [4] have studied this idea and have proposed both auctions and revelation mechanisms. While the auction in this paper is based on the Cremer-Spiegel-Zheng auction, our bidder model, motivation, and the properties we prove differ substantially. Both Burguet and Cremer *et al.* assume that the bidders can only gather information on their own valuations, as opposed to our assumption that bidders are free to gather information on any problem they wish. Second, their goal is to design an auction which maximizes the revenue of the seller and so are interested in extracting the expected surplus from the bidders by imposing appropriate admission fees. Instead, we focus on the strategic behavior of the bidders and show that the auction has desirable properties from their perspective.

We now describe the assumptions that have been made in this paper. First, we assume that there are n deliberative bidders, and that the performance profiles and the cost functions of the bidders are known to the auction designer and to each other. Second, we assume that the cost functions and the performance profiles of the bidders are independent. Both these assumptions are required for the results.

To simplify the explanation of the auction, as well as the notation, we will make certain assumptions about the performance profiles of the agents. In particular, we assume that while a bidder i does not know its valuation *a priori*, it does know that it comes from some interval $[0, \overline{v}_i]$ with probability distribution P_i which is equivalent to the descriptive component of the bidder's performance profile. We argue that this assumption is not unduly restrictive. First, deliberative agents with performance profiles like this still have incentive to strategically deliberate in standard auctions. Second, the proposed auction will still work in settings where bidders have more sophisticated performance profiles. The only repercussion is that sorting the bidders, as will be described next, becomes more complex.

In the rest of this section we describe the optimal-search auction. We first give an overview of the optimal search procedure used by the auctioneer, and then describe how the auctioneer uses it when setting up an auction. In the next section (Section 4) we prove that this has the properties we desire.

3.1 Optimal Search

In the operations research literature a search problem is often formulated in the following way. Assume that there are n boxes (bidders) to open (ask to deliberate), and each box (bidder) i has some random value v_i which is drawn from distribution (performance profile) $P_i(v)$ with density f_i. If the searcher opens a box (queries a bidder) i then it incurs a cost of cost_i. The payoff to the searcher is the maximum value found up to the point when it stops. The question is, then, in what order should the boxes (bidders) be searched (asked to deliberate) and when should the process stop?

Weitzman proved that this problem can be solved optimally using a very simple procedure [19]. First, each box (or bidder) is assigned a *cutoff value* K_i where

$$K_i = \int_0^{K_i} K_i f_i(v) dv + \int_{K_i}^{\overline{v}_i} v f_i(v) dv - \text{cost}_i.$$

If the searcher had already opened a box with reward K_i then it would be ambivalent between opening box i or keeping reward K_i. If the best reward found so far is less than K_i then the searcher is best off opening box i.

Once the cutoff values have been computed, the optimal search procedure is completely characterized by the following rules:

1. (**Selection Rule**) If a box is to be opened, it should be that closed box with highest reservation price.

2. (**Stopping Rule**) Terminate search whenever the maximum sampled reward exceeds the reservation price of every closed box.

The cutoff values and the search procedure have some interesting properties which make them particularly useful for our application. The cutoff values are not the expected valuations of the bidders and, in fact, the optimal search procedure may well search a bidder with a low expected value before a bidder with a higher expected value. As Weitzman noted "Other things being equal, it is optimal to sample first from distributions that are more spread out or riskier in hopes of striking it rich early and ending the costly search" [19]. Interestingly, one of the main incentives driving bidders to strategically deliberate is very similar. Bidders are better off paying some small amount to learn whether or not a competitor has a high valuation for an item (and thus would be likely to win the auction) and only then deciding whether it is worth while to gather information on its own valuation problem [7]. Some other appealing properties of the cutoff values and the search procedure are that as the cost of learning a valuation increases, the cutoff value decreases, and as the search progresses the valuations needed to stop the search decrease.

We note that the cutoff value calculation is equivalent to finding a fixed-point, and may be a computational challenge for the searcher. However, in many settings, finding the fixed point is relatively straightforward. For example, if v_i is drawn from the uniform distribution over $[0, 1]$, and the cost of learning v_i is 0.25, then the cutoff value is

$$\begin{aligned} K_i &= \int_0^{K_i} K_i dv + \int_{K_i}^{1} v dv - 0.25 \\ &= K_1^2/2 + 0.25. \end{aligned}$$

Solving for K_i, we get $K_i \approx 0.293$.

3.2 The Auction

The optimal-search auction progresses through a series of stages. First, using its knowledge of the performance profiles and cost functions of each bidder, it computes the cutoff values. It labels the bidders in decreasing order of cutoff values so that $K_1 > K_2 > \ldots K_n$. It then asks each bidder in turn to gather information on its valuation problem and reveal the result. If, after asking the first t bidders to deliberate and reveal their results, the highest valuation is greater than K_{t+1} then the auction is halted, the highest bidder wins the item, and the auction has followed the optimal search procedure. However, it is not possible to simply get each bidder to truthfully reveal their valuations when asked. Instead, the auctioneer uses reserve prices and a sequence of second-price auctions to get the bidders to truthfully reveal whether or not their deliberations have resulted in valuations which are larger than a specific cutoff value.

In stage 1, the auctioneer tells bidder 1 that it should gather information on its valuation problem. The auctioneer then makes a take-it-or-leave offer of price p_1^1 to the agent. If the offer is accepted then the auction stops, bidder 1 gets the item and pays p_1^1. If the offer is rejected then the auction continues to the next stage. If the auction reaches stage t then bidder t is asked to deliberate on its valuation problem. Once it has done so then bidders 1 to t participate in a second-price auction with bidder and stage specific reserve prices. If bidder i submits a bid which is greater than its reserve price p_i^t then the auction stops. If only one bidder submitted a bid which was greater than its reserve price then it is given the item and it pays its reserve price. If there are multiple bidders who submitted bids that were higher than their reserve prices then the highest bidder i wins the item and pays $p = \max[p_i^t, b_{-i}]$ where $b_{-i} = \max\{b_j | b_j > p_j^t$ and $j \neq i\}$. If the auction reaches the final state, n, then all bidders have been asked to gather information on their valuations. All bidders then participate in a second-price auction with no reserve prices.

The auctioneer must be careful when it comes to setting the reserve prices. If it is not done properly then bidders may have incentive to lie about their valuation after they have learned them. If the reserve prices are too low then a bidder may want to submit a bid which is higher than its true valuation in order to stop the auction early so as to avoid competition from others. If the reserve prices are set too high then a bidder may submit a bid which is lower than its true valuation in the hopes that as the auction progresses the reserve prices will drop. Cremer *et al* showed that it is possible to set the reserve prices so that in (Bayes-Nash) equilibrium the bidders bid truthfully [4]. This is done by working backwards through the stages of the auction.

If the auction reaches the last stage, then all bidders know their valuation. Since they participate in a second-price auction, each bidder is best off bidding truthfully. In the second last stage the first $n-1$ bidders know their valuations. Assume all of these bidders but i bid truthfully. Our goal is to set the reserve price of bidder i, p_i^{n-1} so that bidder i will only submit a bid greater than p_i^{n-1} if its valuation is greater than or equal to K_n. If $v_i = K_n$ then the only way the bidder can win the auction is if all bidders who know their valuations submitted a bid less than K_n. That is $v_{-i} < K_n$. If bidder i bid truthfully, the auction would stop, it would win, and it would pay its reserve price. Its expected payoff is

$$(K_n - p_i^{n-1})\Pi_j P_j(K_n) \qquad (1)$$

where $1 \leq j < n$, $j \neq i$. If bidder i lied about its valuation at stage $n-1$ then the auction would continue to the final stage where all bidders are informed and all participate in a second-price auction. Bidder i will only win if all bidders, including bidder n, have valuations which are less than $v_i = K_n$. Since, by assumption, all other bidders are telling the truth about their valuations, agent i knows that the first $n-1$ bidders (minus itself) must have valuations which are less than K_n. Thus, its payoff if it waits for the last round is

$$\Pi_{1\leq j\leq n-1} P_j(K_n) \int_0^{K_n} (K_n - v_{-i}^n) d\frac{\Pi_{1\leq j\leq n} P_j(v_{-i}^n)}{\Pi_{1\leq j<n-1} P_j(K_N)} \qquad (2)$$

where $j \neq i$ and v_{-i}^n is the highest valuation of the bidders from 1 to n, not including bidder i. Bidder i is ambivalent between declaring its true valuation in round $n-1$ and waiting for the second-price auction in the next round whenever (1)=(2). If the reserve price p_i^{n-1} is set such that (1)=(2) then bidder i will truthfully reveal its valuation, given that all other bidders are also doing the same.

The appropriate reserve prices for earlier rounds are determined in a similar fashion. If in stage t bidder i has valuation $v_i = K_{t+1}$, and assuming that all bidders j, $1 \leq j \leq t$, $j \neq i$ reveal their valuations truthfully, then, if bidder j bids truthfully, the auction will stop and its payoff will be

$$(K_{t+1} - p_i^t)\Pi_j P_j(K_{t+1}) \qquad (3)$$

where $1 \leq j < t$, $j \neq i$. If the bidder waits until the next stage and then reveals its true valuation, its expected payoff is

$$(K_{t+1} - p_i^{t+1})\Pi_j P_i(K_{t+2}) + \qquad (4)$$
$$\Pi_j P_j(K_{t+1}) \int_{K_{t+2}}^{K_{t+1}} (K_{t+1} - v_{-i}^{t+1}) d\frac{\Pi_j P_j(v_{-i}^{t+1})}{\Pi_j P_j(K_{t+1})}$$

where $1 \leq j < t+1$ and $j \neq i$, and where bidder i will pay p_i^{t+1} if $v_{-i}^{t+1} < K_{t+2}$ and v_{-i}^{t+1} if $v_{-i}^{t+1} > K_{t+2}$. If p_i^t is set such that (3)=(4) then bidder i is best off bidding truthfully in stage t, assuming that all other possible bidders also bid truthfully.

4. PROPERTIES

The optimal-search auction described in the previous section has some properties which make it appealing. First, there exists a Bayes-Nash equilibrium where the bidders truthfully reveal their valuations at the right time. Secondly, and of particular interest, it is structured so that bidders have no incentive to strategically deliberate. That is, the optimal-search auction is deliberation-proof and non-misleading.

Proposition 1 *The optimal-search auction is non-misleading. There exists a Bayes-Nash equilibrium where bidders truthfully reveal their valuations when asked by the auctioneer.*

Proof: (Sketch) The full proof is quite technical and follows the proof described in Cremer *et al* [4]. Thus, we only provide an overview of the proof. The main idea is that the reserve prices used by the auctioneer have been carefully chosen so that a Bayes-Nash equilibrium, where the bidders truthfully reveal the results of their deliberation, exists. If the auction does not stop before all bidders have been asked to deliberate on their own valuation problems, then a second-price auction is run. At this point, all the bidders know their valuations and so are best off bidding truthfully. At stage $t < n$, using the equilibrium hypothesis that all other bidders who have learned their valuations truthfully reveal them, the reserve prices for bidder i were carefully designed so that it is in it's best interest to truthfully reveal the valuation it obtained via deliberation. This is because in stage t bidder i is participating in a second-price auction with $t-1$ other bidders and so it's bid does not change the price it will pay if it wins in stage t. Additionally, the reserve prices are such that if $v_i \geq K_{t+1}$ then bidder i is best off revealing this and stopping the auction in stage t, while if $v_i < K_{t+1}$ then the bidder is best off letting the auction continue. In either situation, truth-telling is bidder i's best response, given its beliefs about the potential valuations of all other bidders, conditional on the fact that the auction has reached stage t and the hypothesis that all other bidders truthfully reveal their valuation once they have learned it.

□

Proposition 2 *The optimal-search auction is deliberation-proof. Deliberative bidders have no incentive to use their deliberation resources on the valuation problems of others.*

Before we present the proof of this proposition, we provide some insight into why bidders strategically deliberate in auctions. Strategic deliberation is caused when there is some form of asymmetry between the bidders in terms of the cost functions or the performance profiles. Strategic deliberation often occurs when one bidder gathers information on a competitor in order to check whether the competitor has a high valuation or not. Based on this information, the deliberating bidder will then decide whether or not it is worthwhile to deliberate on its own valuation problem. Interestingly, the optimal search procedure uses a similar motivation for sorting agents. In particular, it queries agents in such an order so that agents that other bidders may want to gather information on (i.e those with potentially high valuations and low cost) are asked to deliberate and reveal their results early in the process, thus removing the need for other agents to gather information on them.

Proof: (Sketch) If the auction stops before a bidder is asked to deliberate, then the bidder has no incentive to deliberate on any problem and thus will not strategically deliberate. If a bidder already knows its valuation then it has no incentive to learn the valuations of other bidders in the optimal-search auction, since it will participate in second-price auctions. Therefore, the only situation of interest is when the auction asks a bidder to gather information and it has not yet done so on its own valuation problem.

Another useful observation is that if $\text{cost}_i < \text{cost}_j$ then bidder i has no incentive to deliberate on the valuation problem of bidder j. That is, if it is more expensive to deliberate on another bidder's valuation problem compared to ones own, then one will not deliberate on the other's problem.

Case 1: Assume that $K_i > K_j$, and assume that the auction has reached stage i. Since the performance profiles of all bidders are common knowledge, bidder i can determine $E[v_i]$ and $E[v_j]$. If $E[v_i] \leq E[v_j]$ then it must be the case that $K_i < K_j$ since $\text{cost}_i > \text{cost}_j$. This is a contradiction, and so it must be the case that $E[v_j] < E[v_i]$.

If $E[v_i] > K_j$, then in expectation the auction will stop before it reaches bidder j and thus the valuation of bidder j is irrelevant. Therefore, bidder i has no incentive to incur a cost from learning it. Assume that $E[v_i] < K_j$. Let $t = \max\{t | K_t < E[v_i]\}$. That is, in expectation, stage t is the earliest stage in which bidder i could influence the outcome of the auction. If $K_t > E[v_j]$ then we are in a similar situation as before and so for the same reason bidder i has no incentive to deliberate on the valuation problem of bidder j. If $K_t < E[v_j] < E[v_i]$ then it is possible that v_j will effect what price bidder i may pay, but in expectation it will not effect the allocation decision of the auction. Since strategic-deliberation is driven by allocation uncertainty as opposed to price uncertainty, bidder i has no incentive to deliberate on player j.

Case 2: Now assume that $K_j > K_i$. If the auction never reaches stage i then bidder i has no incentive to deliberate on any valuation, and thus certainly not on the problem of bidder j. Assume that the auction does reach stage i. This can only occur if $v_j < K_i$. Since, the auction is implementing an optimal search process, we know that it must be the case that

$$v_j - \sum_{1<l<i} \text{cost}_l < E[v_i] - \text{cost}_i - \sum_{1<l<i} \text{cost}_l.$$

Therefore, by simple algebra,

$$v_j < E[v_i] - \text{cost}_i$$

Clearly bidder i now has no incentive to deliberate on bidder j's problem. In expectation bidder i has a higher valuation than bidder j and thus will be allocated the item before bidder j would. Secondly, even if bidder j influenced the price that bidder i had to pay, learning the actual valuation of bidder j will not change the deliberation strategy of bidder i.

□

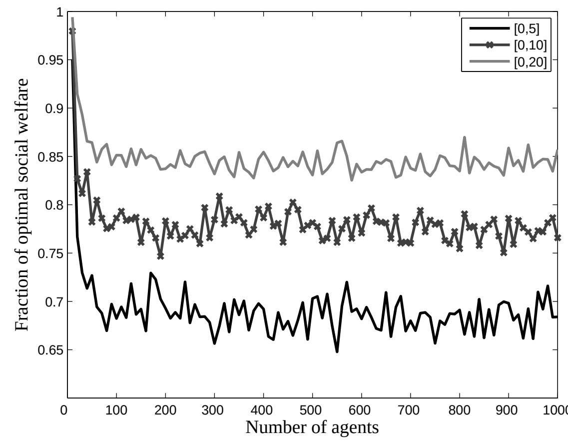

Figure 2: Fraction of the optimal allocation the search auction achieves. In this example all agents have their cost functions fixed to 1.0.

5. EXPERIMENTAL RESULTS

We investigated how the optimal-search auction worked in practice. In particular, we were interested in the economic efficiency of this approach, and whether it significantly reduced the amount of deliberation among the agents. To this end, we conducted two sets of simulations. The first set of simulations compared the allocation made by the search auction to the optimal allocation (made by a second-price auction). The second set of simulations compared the total cost of deliberation of all agents in the search auction with that of a second-price auction.

We set the parameters of the two sets of experiments the same way. We divided the bidders into six different classes, and then ran second-price auctions and optimal-search auctions in each class as we varied the number of bidders.

Each bidder had some valuation v_i which was drawn independently from the uniform distribution over $[0, V]$, where V was either 5, 10 or 20. For each of these valuation intervals, we set the cost functions in one of two ways. We first set the cost function of all agents in each class to be $\text{cost} = 1.0$. This meant that the cutoff values used in the optimal-search auction were the same for all bidders who's value had been drawn from the same interval. This put the optimal-search auction at a disadvantage since it had to run with minimal information available to it. We varied the cost functions for the bidders in other experiments. In particular, for any bidder who had a valuation drawn from the interval $[0, V]$, we randomly drew a cost constant cost_i from the uniform distribution over $[0, V/2]$. This meant that with high probability each bidder had a different cutoff value allowing the optimal-search auction to order the bidders to its advantage.

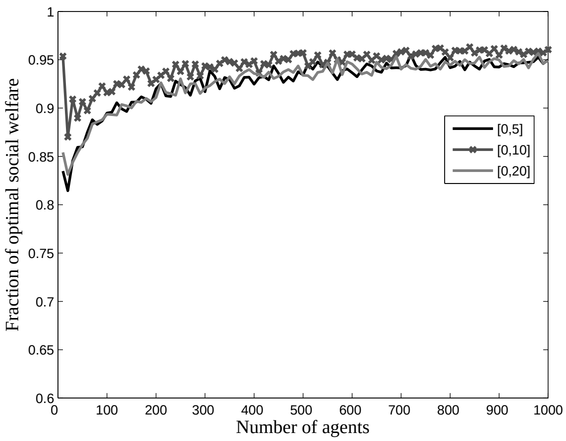

Figure 3: Fraction of the optimal allocation the search auction achieves. In this example, cost functions are drawn uniformly and independently from an interval related to their valuations.

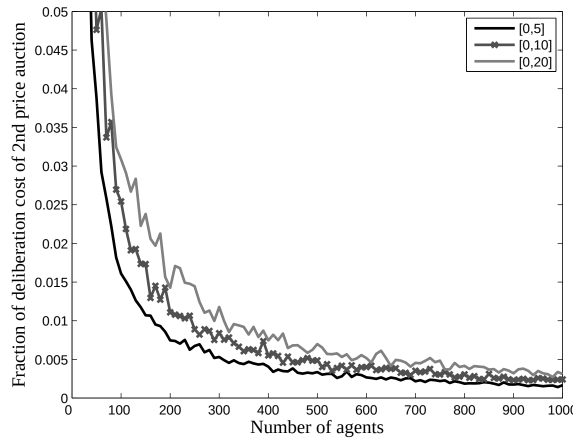

Figure 4: Cost incurred by all agents in the search auction compared to cost incurred in a second price auction. As the number of agents increase, the fraction becomes vanishingly small. Cost functions were fixed so that they were all equal to 1.0. This means that all agents (in each experiment) had the same cutoff value.

For each of the six parameter configurations we ran a series of second-price auctions and optimal-search auctions. We compared the efficiency of the optimal-search auction with that of the efficient second-price auction in order to determine what the efficiency loss was. We also compared the total cost of deliberation done by all bidders in the second-price auction with that done in the optimal-search auction in order to determine whether there was a substantial overall savings with the optimal-search auction. In each experiment we varied the number of bidders from 2 to 1000, and repeated each experiment 100 times. We present the average of our results.

In our first set of experiments we measured the efficiency of the search auction with that of the second-price auction. The second-price auction is guaranteed to allocate the item to the agent who has the highest valuation, while the search-auction allocates the item to the agent who has the highest valuation amongst the agents who have been asked to deliberate before the auction is stopped. It is possible that the search-auction is stopped before the agent with the actual highest valuation is asked to deliberate. Figures 2 and 3 show our findings.

In Figure 2, the search auction is at a disadvantage since all the cutoff values for the agents in each experiment are the same. However, we still observe that the search auction is within at least 65% of optimal (the allocation of the second-price auction) for all parameter settings. The allocation does depend on the interval from which the valuations are being drawn, with values from $[0, 5]$ leading to a lower allocation than the others. If the auction had more than 100 bidders, then the allocation did not seem to be affected by the number of bidders (this was expected). In Figure 3 the search auction was able to make better allocations when it was able to distinguish between the agents based on their cutoff values. Once we had at least 200 bidders, the allocation made by the search auction was always within 90% of optimal.

In our second set of experiments we studied whether or not the search auction reduced the total amount of deliberation done by all agents significantly, compared to what would happen in a second-price auction. In a second-price auction, given our setup, each agent would have incentive to gather information on their own valuation problem, while in the search auction only a subset of the agents will likely be asked to gather information (and thus incur a cost). Figures 4 and 5 show our results. Even if the search auction has little information by which to sort the agents (Figure 4) the overall savings in deliberation cost from a system-wide perspective is substantial. This is even more marked in the experiments where agents had different cutoff values, allowing the search auction to query the most promising agents first (Figure 5).

To conclude, our experiments show that the search auction does quite well when it comes to allocating the item. It also greatly reduces the amount of deliberation done by all agents in the auction.

6. CONCLUSION

A common assumption in auction research is that bidders have private information about their willingness to pay for the item being auctioned and that they use this information strategically when formulating their bids. In reality, bidders often have to undergo a costly information-gathering process in order to determine their valuations for the item being auctioned. Recent attempts at modelling this phenomenon has brought to light complex strategic behavior that had previously been overlooked. In particular, bidders may either have incentive to actively gather information on others or may have incentive to "fool" other bidders about their results. Earlier work proved that if the auction designer had no information about the bidders' information-gathering processes then this strategic behavior on the part of the agents was impossible to avoid.

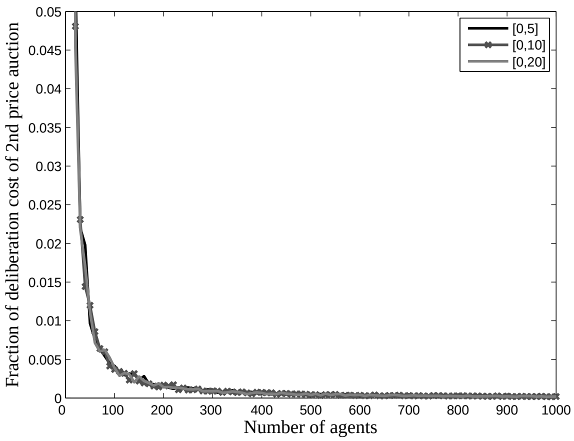

Figure 5: Cost incurred by all agents in the search auction compared to cost incurred in a second price auction. As the number of agents increase, the fraction becomes vanishingly small. Cost functions were drawn uniformly and independently from intervals related to their valuations.

In this paper we showed that by providing the auction designer with some minimal amount of information about the information-gathering processes of the bidders, it is possible to design an auction such that the bidders only gather information on their own valuation problems, and, when asked, report the results truthfully to the auctioneer. The insight is that by carefully ordering the agents using an optimal search procedure, and then asking them sequentially to gather information and reveal their results, it is possible to remove the uncertainty that drives the complex strategic behavior of deliberative agents in standard auctions. We additionally showed, via simulation results, that the savings in terms of overall information-gathering costs is significant.

There are several directions in which this work can be taken. This paper focussed solely on single-item auctions, and so generalizing the approach to multi-item auctions is an obvious next step. Second, the search auction relied on the assumption that the performance profiles and cost functions of the agents are known to everyone. While it seems likely that the auction also provides the correct incentives so that bidders will truthfully reveal their performance profiles and cost functions, this still needs further analysis. Finally, it is our belief that optimal search procedures from the operations research literature may find use in other settings, for example in situations where the total amount of communication has to be minimized.

7. REFERENCES

[1] D. Bergemann and J. Välimäki. Information acquisition and efficient mechanism design. *Econometrica*, 70:1007–1034, 2002.

[2] R. Burguet. Optimal repeated purchases when sellers are learning about costs. *Journal of Economic Theory*, 68:440–455, 1996.

[3] O. Compte and P. Jehiel. Auctions and information acquisition: Sealed-bid or dynamic formats?, 2001. Working Paper.

[4] J. Cremer, Y. Spiegel, and C. Zheng. Optimal selling mechanisms with costly information acquisition. Working paper, August 2003.

[5] N. Kfir-Dahav, D. Monderer, and M. Tennenholtz. Mechanism design for resource bounded agents. In *Proceedings of the Fourth International Conference on Multi-Agent Systems (ICMAS)*, pages 309–315, Boston, MA, 2000.

[6] V. Krishna. *Auction Theory*. Academic Press, 2002.

[7] K. Larson and T. Sandholm. Costly valuation computation in auctions. In *Theoretical Aspects of Rationality and Knowledge (TARK VIII)*, pages 169–182, Siena, Italy, July 2001.

[8] K. Larson and T. Sandholm. Using performance profile trees to improve deliberation control. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, San Jose, CA, USA, July 2004.

[9] K. Larson and T. Sandholm. Mechanism design and deliberative agents. In *Proceedings of the Fourth International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS'05)*, Utrecht, Netherlands, 2005.

[10] R. Lavi, A. Mu'Alem, and N. Nisan. Towards a characterization of truthful combinatorial auctions. In *Proceedings of the Annual Symposium on Foundations of Computer Science (FOCS)*, 2003.

[11] D. Lehmann, L. I. O'Callaghan, and Y. Shoham. Truth revelation in rapid, approximately efficient combinatorial auctions. *Journal of the ACM*, 49(5):577–602, 2002. Early version appeared in ACMEC-99.

[12] A. Mas-Colell, M. Whinston, and J. R. Green. *Microeconomic Theory*. Oxford University Press, 1995.

[13] N. Nisan and A. Ronen. Computationally feasible VCG mechanisms. In *Proceedings of the ACM Conference on Electronic Commerce (ACM-EC)*, pages 242–252, Minneapolis, MN, 2000.

[14] D. Parkes. *Iterative Combinatorial Auctions: Achieving Economic and Computational Efficiency.* PhD thesis, University of Pennsylvania, 2001.

[15] D. Parkes. Auction design with costly preference elicitation. *Annals of Mathematics and AI*, 44:269–302, 2004.

[16] N. Perisco. Information acquisition in auctions. *Econometrica*, 68(1):135–148, January 2000.

[17] M. Rothkopf. On misusing auctions to value stranded assets. *The Electricity Journal*, 10(10):10–17, December 1997.

[18] T. Sandholm. Issues in computational Vickrey auctions. *International Journal of Electronic Commerce*, 4(3):107–129, 2000. Special Issue on Applying Intelligent Agents for Electronic Commerce. A short, early version appeared at the Second International Conference on Multi–Agent Systems (ICMAS), pages 299–306, 1996.

[19] M. Weitzman. Optimal search for the best alternative. *Econometrica*, 47(3):641–654, 1979.