

Using value queries in combinatorial auctions

Benoît Hudson* Tuomas Sandholm
Carnegie Mellon University
Computer Science Department
5000 Forbes Avenue
Pittsburgh, PA 15213, USA
{bhudson,sandholm}@cs.cmu.edu

ABSTRACT

Combinatorial auctions, where bidders can bid on bundles of items are known to be desirable auction mechanisms for selling items that are complementary and/or substitutable. However, there are $2^k - 1$ bundles, and each agent may need to bid on all of them to fully express its preferences. We address this by showing how the auctioneer can recommend to the agents incrementally which bundles to bid on so that they need to only place a small fraction of all possible bids. In particular, we design preference elicitation algorithms that provably determine the optimal allocation while asking a vanishingly small fraction of all possible value queries. The elicitors impose a great computational burden on the auctioneer. We show how to speed them up dramatically. In order to provide an instance-specific lower bound on how well *any* elicitation algorithm can do, we develop a search-based method for finding the smallest certificate. We show that our best elicitor is almost as effective as this omniscient elicitor. We also present an optimal non-omniscient elicitor, but it is utterly intractable and thus impractical. Finally, we introduce the notion of a *universal revelation reducer*, demonstrate a randomized one, and prove that no deterministic one exists.

1. INTRODUCTION

Combinatorial auctions, where agents can submit bids on *bundles* of items, are economically efficient mechanisms for selling k items to n bidders, and are attractive when the bidders' valuations on bundles exhibit *complementarity* (a bundle of items is worth more than the sum of its parts) and/or *substitutability* (a bundle is worth less than the sum of its parts). Determining the winners in such auctions is a complex optimization problem that has recently received considerable attention (e.g., [1, 8, 15, 21]).

An equally important problem, which has received much less attention, is that of bidding. There are $2^k - 1$ bundles, and each agent may need to bid on all of them to fully express its preferences. This can be undesirable for any of several reasons: determining one's valuation for any given bundle can be computationally intractable [13, 18–20]; there is a huge number of bundles to evaluate; communicating the bids can incur prohibitive overhead (e.g.,

*The first author of this paper is a student. The material in this paper is based upon work supported by the National Science Foundation under CAREER Award IRI-9703122, Grant IIS-9800994, ITR IIS-0081246, and ITR IIS-0121678. Some of the results in this paper appeared in the AMEC 2002 workshop.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

Copyright 2002 ACM X-XXXXX-XX-X/XX/XX ...\$5.00.

network traffic); and agents may prefer not to reveal all of their valuation information due to reasons of privacy or long-term competitiveness. Appropriate bidding languages [8, 10, 15, 21, 22] can solve the communication overhead in some cases (when the bidder's utility function is compressible). However, they still require the agents to completely determine and transmit their valuation functions and as such do not solve all the issues. So in practice, when the number of items for sale is even moderate, the bidders will not bid on all bundles. Instead, they may wastefully bid on bundles which they will not win, and they may suffer reduced economic efficiency by failing to bid on bundles they would have won.

Selective incremental preference elicitation by the auctioneer was recently proposed to address these problems [6]. Some of the theoretical properties of the so-called rank-lattice based elicitors were recently studied [7]. That family of elicitors suffers from having to ask queries in an extremely rigid order. In this paper we study the more practical approach where the elicitor uses value queries in arbitrary order. We first develop a framework to describe our algorithms (Section 2), and an experimental setup for evaluating them (Section 3). Using this framework, we design a series of increasingly effective elicitation algorithms and study their properties theoretically and experimentally, showing that elicitation reduces revelation drastically, and that this benefit increases with problem size (Section 4). In Section 5 we provide an omniscient elicitor, which acts as an instance-specific lower bound on any real elicitor. We then discuss an optimal non-omniscient elicitor (Section 6) which is unfortunately not tractable. Finally, we introduce the notion of a universal revelation reducer, and demonstrate that while a randomized one exists, no deterministic one does (Section 7).

2. ELICITATION FRAMEWORK

We study a setting where one auctioneer is selling a set K , $|K| = k$, of items to n bidders. We work with the standard independent private values model, that is, we assume that each bidder's valuations for the bundles are independent of the other bidders' valuations.

The elicitation framework that we use is a slightly modified version of that of Conen & Sandholm [6]. The auctioneer asks questions of the agents (it *elicits* information, and thus it is an *elicitor*), and assimilates the information from the replies. Based on that information, the elicitor clears the auction if possible; if not, it elicits further information.

2.1 Value propagation

In this paper we focus on elicitation using *value queries*. In a value query, the elicitor asks an agent i for that agent's valuation $v_i(b)$ for bundle b .¹ Throughout this paper we make the usual as-

¹An important issue is that the agents might reveal their valuations




Figure 1: Constraint network with 3 items a , b , and c .

assumption of *free disposal*, that is, any agent can throw away extra items for free, so extra items never decrease an agent’s valuation. This allows the elicitor to infer an upper bound on any subbundle b^- of b : namely, no smaller bundle can be worth more than a larger bundle. Therefore $UB_i(b^-) = \min_{b \supset b^-} [v_i(b)]$. Similarly, no larger bundle can be worth less than a smaller bundle. Therefore, $LB_i(b^+) = \max_{b \subset b^+} [v_i(b)]$. The empty bundle is assumed to be worth 0.

To support fast inference of the upper and lower bounds, we implemented an interval constraint network. The elicitor maintains one network for each agent. In each network i , a node corresponds to a bundle b . The node is labeled by an upper bound $UB_i(b)$ and a lower bound $LB_i(b)$ which together define an interval in which the true value $v_i(b)$ is known to lie. A node for b has a directed link to a node for b' if $b \supset b'$. When a value for b becomes known, the elicitor sets $UB_i(b) = LB_i(b)$ to that value. It then propagates the value down the links to smaller bundles, enforcing the invariant that the $UB_i(b^-)$ in the network corresponds to the definition of the previous paragraph; similarly, it propagates the value up the links to larger bundles to enforce $LB_i(b^+)$.

2.2 Certificates

The auctioneer clears the auction if, given the information it has received, the auctioneer can infer that one allocation is worth at least as much as any other. That allocation is an optimal allocation. If the information the auctioneer has allows this inference, the information forms a *certificate* for that allocation. The certificate contains a set of value queries. A *minimal certificate* is a certificate that would cease to be a certificate if any query were removed from it. A *shortest certificate* is a certificate that has the smallest number of queries among all certificates. This depends on the problem instance. For any given instance, we call the length of the shortest certificate q_{min} .

2.3 Elicitation algorithm

The skeleton of the general elicitation algorithm is insincerely. Furthermore, the elicitor’s queries leak information: a query that an agent receives signals to that agent information about how the other agents have answered the queries so far. That opens the door for additional strategic manipulation by the agents. However, if *Vickrey-Clarke-Groves (VCG)* pricing [5, 9, 24] is used (this requires some additional elicitation, but not significantly more), and the bidders’ utility functions are quasilinear, then revealing one’s preferences truthfully is an *ex post* equilibrium [6].

```
SOLVE()
1  $C \leftarrow \text{INITIALCANDIDATES}(n, k)$ 
2 while not DONE( $C$ )
3    $q \leftarrow \text{SELECTQUERY}(C)$ 
4   ASKQUERY( $q$ )
5    $C \leftarrow \text{PRUNE}(C)$ 
```

Here, C is a set of candidate allocations, where a *candidate* is a vector $c = \langle c_1, c_2, \dots, c_n \rangle$ of bundles where the bundles contain no items in common. There are n agents bidding on k items. The value of a candidate is $v(c) = \sum_i v_i(c_i)$; $UB(c) = \sum_i UB_i(c_i)$ is an upper bound, and $LB(c) = \sum_i LB_i(c_i)$ a lower bound.

INITIALCANDIDATES generates the set of all candidates, which is the set of all n^k allocations of the k items to the n agents (some agents might not get any items).

PRUNE removes, one candidate at a time, each candidate that is dominated by a remaining candidate. A candidate c *dominates* another candidate c' —which we write as $c \succeq c'$ —if the elicitor can prove, using what it currently knows, that the value of c must be at least as high as the value of c' .²

DONE returns true if each candidate in the set of remaining candidates C is provably optimal. This is the case either if C has only one element, or if all candidates in C have known value (that is, $\forall c \in C, UB(c) = LB(c)$). Because the algorithm has just pruned, it knows that if all candidates have known value, then they have equal value.

SELECTQUERY selects the next query to be asked. This function can be instantiated in different ways to implement different elicitation policies, as we will show.

ASKQUERY takes a query (b, i) , asks agent i for its value of bundle b , then appropriately updates the constraint network. In particular, every superbundle b^+ has its lower bound be the tighter of $v_i(b)$ and the previous lower bound $LB_i(b^+)$; every subbundle b^- has its upper bound be the tighter of $v_i(b)$ and the previous upper bound $UB_i(b^-)$.

3. EXPERIMENTAL SETUP

We ran experiments to evaluate the effectiveness of the elicitation policies we developed. Each experimental plot shows what fraction of all value queries (of which there are $n(2^k - 1)$) are made before the auctioneer finds an optimal allocation and can prove that it is optimal (that no other allocation is better). We call this fraction the *elicitation ratio*. In each plot, each point represents an average over 50 runs, where each run is on a different problem instance (different draw of valuations for the agents). Each algorithm was tested on the same set of problem instances.

Unfortunately, real data for combinatorial auctions are not publicly available.³ Therefore, as in all of the other academic work on combinatorial auctions so far, we used randomly generated data. We first considered using existing benchmark distributions. However, the existing problem generators output instances with sparse bids, that is, each agent bids on a relatively small number of bundles. This is the case for the CATS suite of economically-motivated random problem instances [14] as well as for the other prior benchmarks [1, 8, 21]. However, in many real settings (spectrum auctions

²This is the case if $\sum_{i=1}^n \delta_i(c_i, c'_i) \geq 0$, where $\delta_i(b, b')$ is the greatest difference between $v_i(b)$ and $v_i(b')$ that can be proven given the information in the constraint network. Namely, if $b \supset b'$, then $\delta_i(b, b') = \max(0, LB_i(b) - UB_i(b'))$; otherwise, $\delta_i(b, b') = LB_i(b) - UB_i(b')$.

³Furthermore, even if the data were available, they would only have some bids, not the full valuation functions of the agents (because not all agents bid on all bundles).

for instance), every bidder has positive value on every bundle (at least due to renting and reselling possibilities). In addition, the instances generated by many of the earlier benchmarks do not honor the free disposal constraints.

To capture these considerations, we developed a new benchmark problem generator. In each problem instance we generate, each bidder has a nonzero valuation for almost every bundle, and all valuations honor free disposal. Specifically, the generator assigns, for each agent in turn, integer valuations using the following routine. We impose an arbitrary maximum bid value $\text{MAXBID} = 10^7$ in order to avoid integer arithmetic overflow issues, while at the same time allowing a wide range of values to be expressed. Valuations generated with this routine exhibit both complementarity and substitutability.

```

GENERATEBIDS( $K$ )
1  $G \leftarrow$  new constraint network
2  $S \leftarrow 2^K$  (the set of all bundles)
3 impose free disposal constraints on  $G$ 
4  $UB(K) \leftarrow \text{MAXBID}$ 
5 while  $S \neq \emptyset$ 
6   pick  $b$  uniformly at random from  $S$ 
7    $S \leftarrow S - b$ 
8   pick  $v(b)$  uniformly at random from  $[LB(b), UB(b)]$ 
9   propagate  $LB(b) = UB(b) = v(b)$  through  $G$ 

```

4. ELICITATION POLICIES

We designed and tested several elicitation policies. All of them compute a set of potential queries, and then randomly pick one. We present the policies starting from the simplest. We then incrementally add restrictions on the query set in an attempt to constrain the random choice towards hopefully better queries.

4.1 Random elicitation policy

The first policy we investigate simply asks random value queries. In the beginning, we generate the set of all $n(2^k - 1)$ value queries. Whenever it is time to ask a query, the policy chooses a random query from the set, ignoring those it has already asked or for which the value can already be inferred.

We can show that if it is possible to save elicitation, then on average this policy saves elicitation:

PROPOSITION 1. *Let $Q = n(2^k - 1)$ be the total number of queries, and let q_{\min} be the number of queries in the shortest certificate. For any given problem instance, the expected number of queries that the random elicitation policy asks is at most $\frac{q_{\min}}{q_{\min} + 1}(Q + 1)$.*

PROOF. Assume pessimistically that a query is either required to prove the optimal allocation or useless. Under this assumption, the analysis reduces to the following problem. We have r red “necessary” balls and b blue “useless” balls in a bag. We then randomly draw one ball at a time without replacement. The question is how many balls we expect to draw before all red balls have been drawn. Let $e(r, b)$ be this number. The base case is $e(0, b) = 0$, because there are no red balls to draw. In the general case, we pick one ball from the bag. With probability $r/(r + b)$, it is red, so the bag now has $r - 1$ red balls and b blue balls. Similarly, with probability $b/(r + b)$, it is blue, so the bag now has r red balls and $b - 1$ blue balls. Therefore, $e(r, b) = 1 + \frac{r}{r+b}e(r - 1, b) + \frac{b}{r+b}e(r, b - 1)$. It is easy to verify that $e(r, b) = \frac{r}{r+1}(b + r + 1)$ solves this recurrence. In the elicitation setting, $r = q_{\min}$ and $r + b = Q$. The result follows. \square

The upper bound given in the above proposition only guarantees relatively minor savings in elicitation (especially because q_{\min} in-




Figure 2: Random elicitation policy.

creases when the number of agents and items increases). However, we base the proof on the pessimistic assumption that there is only one short certificate, and that all other certificates are extensions of the shortest one. When there are several minimal certificates, it is more likely that the next query will complete a certificate. We do not know how to analyse how many minimal certificates we expect to be present. Therefore, we ran an experiment to see how well this policy does in practice. Figure 2 shows that the elicitation ratio q/Q is indeed less than 1. Also, the ratio slowly falls as the number of items increases, proving that indeed, the number of minimal certificates increases with the number of items. Nevertheless, we would hope to do better: on all the instance sizes we ran on, this algorithm asks more than half of all the queries.

4.2 Random allocatable elicitation policy

Essentially, the random elicitation policy is asking many queries which, as it turns out, are not useful. We will now present a useful restriction on the set of queries from which the elicitation policy should choose. The key observation is that the elicitor might already know that a bundle b will not be allocated to a particular bidder—even before the elicitor knows the bidder’s valuation for the bundle. This occurs when the elicitor knows that the value other agents have for the other items (those in $K - b$) is insufficient. On the other hand, if the elicitor cannot (yet) determine this, then the bundle-agent pair is deemed *allocatable*.

DEFINITION 1. *A bundle-agent pair (b, i) is allocatable if there exists a remaining candidate allocation $c \in C$ such that $c_i = b$.*

Now we can refine our random elicitation policy to ask queries on allocatable (b, i) only. This restriction is intuitively appealing, and we can characterize cases where it cannot hurt. We define the notation $\langle x, y \rangle$ to mean that revealing the value of a non-allocatable (b, i) would raise the lower bound on x allocatable super-bundles of b (that is, there are x allocatable (b', i) such that $b' \supseteq b$), and lower the upper bound on y allocatable sub-bundles of b . To affect a lower bound, $v_i(b)$ must be strictly greater than the currently-proven lower bound on any of the x super-bundles. Similarly, $v_i(b)$ must be strictly less than the currently-proven upper bound on the y sub-bundles.

Given this notation, we can examine the cases where eliciting a non-allocatable (b, i) is no more useful than eliciting some allocatable (b', i) . Because the elicitor does not know $v_i(b)$, it cannot know what case actually applies.

PROPOSITION 2. *Given the value queries the elicitor has asked so far, for every non-allocatable (b, i) in case $\langle x, y \rangle$ with $x + y < 2$, there is an allocatable pair (b', i) which the elicitor, in hindsight, would have preferred (perhaps not strictly) to ask—no matter what value queries the elicitor asked from then on.*

PROOF. We analyze each case separately.

Case $\langle 0, 0 \rangle$: If b is not allocatable, and no parent of b is allocatable, and no child of b is allocatable, then eliciting $v_i(b)$ will yield no information that the elicitor could not already infer about the optimality of any given allocation. Thus, in retrospect at the end of the auction, the elicitor would prefer not to have elicited (b, i) , since it would then have found a shorter certificate.

Cases $\langle 0, 1 \rangle$ and $\langle 1, 0 \rangle$: The two cases are symmetric; we assume $\langle 0, 1 \rangle$ here. Let the single allocatable sub-bundle of b be called b' . By eliciting $v_i(b)$, the upper bound on $v_i(b')$ will be tightened. However, eliciting $v_i(b')$ would reveal an upper bound on b' that is at least as tight. So eliciting the allocatable bundle b' would have been no worse. \square

While the idea of restricting the queries to allocatable bundles is intuitively appealing and can never hurt in the cases above, there are cases where this restriction forces the elicitor to ask a larger number of queries:

PROPOSITION 3. *Querying a non-allocatable (b, i) in case $\langle x, y \rangle$ with $x + y \geq 2$ can allow an elicitor to find a shorter certificate than an elicitor that is constrained to querying allocatable (b, i) only.*

PROOF. Assume there are two bidders. Further assume that given the information elicited previously, there are only three allocations that could be optimal ($|C| = 3$). One allocation involves giving bidder 1 the items in bundle b' , and the other items to bidder 2, and bidder 2 places a value of 50 on those items. Similarly, the second allocation gives bidder 1 the items in bundle b'' , and the other items to bidder 2, who again places a value of 50 on those items. Bundle b' is neither a super-bundle nor a sub-bundle of b'' . Given the information elicited so far, the elicitor knows $UB_i(b') = UB_i(b'') = 100$. A third allocation gives bidder 2 all the items, and bidder 2 places a value of 100 on this outcome; bidder 1 gets no items. Finally, assume the true value agent 1 has on bundle b is 40 ($v_1(b) = 40$), and that b is not allocatable for agent 1. This means that $(b, 1)$ is in case $\langle 0, 2 \rangle$.

EXAMPLE 1. *Example of case $\langle 0, x \rangle$ where revealing a non-allocatable bundle b is better than revealing any allocatable bundle. By eliciting $v_1(b) = 40$, the elicitor learns that the first two candidate allocations have a value of at most 90, and can therefore eliminate them.*

$v_1(b') = [0, 100]$	$v_2(K - b') = 50$	sum: [0, 150]
$v_1(b'') = [0, 100]$	$v_2(K - b'') = 50$	sum: [0, 150]
$v_1(\emptyset) = 0$	$v_2(K) = 100$	sum: 100

In this situation, by eliciting just $v_1(b)$, the elicitor would prove that the third allocation (giving bidder 2 all the items) is optimal. Restricted to revealing allocatable bundles b' and b'' but not b , the elicitor will instead need to use two queries instead of just one. \square

Proposition 3 does not mean that in all cases $\langle x, y \rangle$ with $x + y \geq 2$ it is a good idea to ask a non-allocatable (b, i) . Indeed, in some such cases, eliciting the non-allocatable (b, i) gives insufficiently tight bounds on the allocatable bundles it affects, and therefore the allocatable bundles need to be elicited anyway.⁴

⁴An open problem is whether, in the case of an oracle that chooses the best bundle to elicit every time, the bad cases would ever happen.

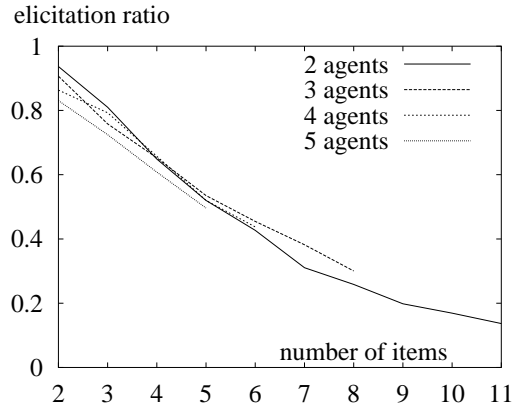


Figure 3: *Random allocatable elicitation policy.*

The case-by-case analysis of Propositions 2 and 3 indicates that restricting the elicitation policy to choosing only allocatable bundles will often help, but may sometimes also cause harm. However, the harm is limited, as we now show.

PROPOSITION 4. *Any bad case $\langle x, y \rangle$ with $x + y \geq 2$ causes the random query policy that restricts itself to allocatable queries to ask at most twice as many queries (in expectation) as the unrestricted random policy.*

PROOF. Assume that to prove an optimal allocation, it is necessary either to reveal the single non-allocatable bundle, or all of the $x + y$ allocatable bundles. If fewer than all $x + y$ bundles are needed, or if there is more than one subset of the $x + y$ that is sufficient, this only reduces the advantage of being allowed to ask the bad-case query.

In the restricted policy, which cannot ask the bad-case query, the elicitor will ask $x + y$ queries. In the unrestricted policy, the elicitor's task corresponds to removing balls one at a time from a bag that has 1 red ball and $b = x + y$ blue balls until either the red ball has been removed, or all the blue balls have been removed. Let $e(b)$ be the number of balls we expect to pick until we are done. We pick a red ball with probability $1/(b + 1)$ and are done immediately. Otherwise, we pick a blue ball with probability $b/(b + 1)$. Therefore, $e(b) = \frac{1}{b+1} + \frac{b}{b+1}(1 + e(b - 1))$. If only one blue ball is left, whether we pick the red ball or the blue ball, we are done, so $e(1) = 1$. It can be verified that $e(b) = \frac{b(3+b)}{2(1+b)}$ solves the recurrence relation. Remembering that $b = x + y$, the ratio of queries asked in the restricted policy to queries asked in the unrestricted policy is $b/e(b) = \frac{2(1+b)}{(3+b)}$. From this, it is clear that $\lim_{b \rightarrow \infty} b/e(b) = 2$; and furthermore, it approaches this limit from below, so the ratio never exceeds 2. \square

In summary, restricting value elicitation to allocatable bundles either helps, does not hurt, or at worst only causes the elicitor to ask (in expectation) twice as many queries. We ran experiments (Figure 3) to determine whether the restriction helps in practice. It clearly does help. For example, at $k = 8$ items and $n = 3$ agents, the elicitation ratio of the unrestricted random policy is 78% while that of the allocatable-only policy is merely 30%. That is, the random elicitation policy restricted to eliciting only allocatable (b, i) avoids much of the elicitation needed in the unrestricted random elicitation policy.

Furthermore, the elicitation ratio vanishes as the number of items increases. This validates the idea of preference elicitation in combinatorial auctions.

4.3 Stronger restrictions

We analyzed several policies which restricted the query set more strongly than simply to allocatable bundles, trying to cleverly account for many features. A list of negative results includes: counting the number of allocations that include (b, i) ⁵; counting the expected amount by which bounds will change in the constraint networks when we elicit (b, i) ; counting the expected number of bounds that will change; counting the expected number of candidates that will be pruned after eliciting (b, i) . For the policies that count the expected value of a quantity, we also tried the minimum and maximum value. All of these policies fared worse in terms of elicitation ratio than the random allocatable elicitation policy.

4.4 High-value candidate elicitation policy

We did find one policy that significantly outperforms the random allocatable elicitation policy. Let C_{\max} be the set of candidates of greatest upper bound (“high-value candidates”). That is, $C_{\max} = \{c \in C \text{ s.t. } UB(c) = \max_{c' \in C} UB(c')\}$. For each $(b, i) \in C_{\max}$ define subbundles (b, i) to be the number of other bundles in high-value candidates whose value might be affected upon eliciting (b, i) . The subbundles are those $(b', i) \in C_{\max}$ for which $b \supset b'$ and $LB_i(b) < UB_i(b')$. Finally, pick uniformly at random among the (b, i) with the most subbundles.

This policy is a modification of the elicitation policy that was recently used in a combinatorial exchange for allocating tasks in a multi-robot system [23]. The intuition is that to prove an allocation optimal, we must prove a sufficient high lower bound on it, while at the same time proving sufficient low upper bounds on all other allocations. By only picking from high-value candidates, we expect to be biasing toward asking questions that will need to be asked anyway. In addition, by picking from those queries that will reduce as many values as possible, we bias toward reducing upper bounds, which is desirable since there is typically only one optimum out of the n^k total candidates (the latter restriction was not present in the previous work).

In terms of elicitation ratio, this is the best policy we know. It achieves an elicitation ratio of only 24% with $k = 8$ items and $n = 3$ agents, as opposed to 30% for the random allocatable policy and 78% for the unrestricted random policy.

4.4.1 Representing candidates implicitly

The policy of the previous section works well in terms of elicitation, but in terms of time it scales poorly with the number of agents. The chief cost is due to representing the candidates explicitly: PRUNE runs in time quadratic in the number of candidates, while SELECTQUERY and DONE run in time linear in the number of candidates, of which there are as many as n^k . Since the policy chooses among a set of size at most $n2^k$, we might hope to save work by implicitly representing the candidates, thus avoiding the expensive operations (as long as $n > 2$).

We accomplish this by repeatedly solving an integer program (IP) every time a query is to be selected—rather than explicitly representing the set of candidates. We use the following IP to compute the value of the highest-valued candidate:

$$\begin{aligned} & \text{maximize} && \sum_{i \in N, b \in 2^K} UB_i(b) x_i(b) \\ & \text{subject to} && x_i(b) \in \mathbb{Z}_2 && \forall i \in N, \forall b \in 2^K \\ & && \sum_{b \in 2^K} x_i(b) \leq 1 && \forall i \in N \\ & && \sum_{i \in N} \sum_{b \supseteq j} x_i(b) \leq 1 && \forall j \in K \end{aligned}$$

⁵This is the value-query policy previously proposed by Conen & Sandholm [6].




Figure 4: High-value candidate elicitation policy, with implicit candidate set representation.

The first constraint of the IP states that each bundle is either allocated or not. The second constraint states that each agent only gets one bundle, and the third constraint states that each item is allocated to only one agent.

Upon solving the IP, the elicitor will know the value UB_{\max} of the candidate with greatest upper bound. Then, for each pair (b, i) in turn, we force $x_i(b) = 1$ and solve again. This returns the value $UB_{\max}(b, i)$ of the candidate with greatest upper bound, constrained to only those candidates which allocate b to agent i . If $UB_{\max}(b, i) = UB_{\max}$, then (b, i) is in a high-value candidate. The elicitor now has the set of (b, i) that are in high-value candidates, and can proceed as before (that is, for each such (b, i) , count the subbundles (b, i) and pick a random (b, i) among the ones with the greatest number of subbundles).

Implemented naively, the policy solves the IP for each pair (b, i) in each call to SELECTQUERY. One can sometimes avoid solving the IP by caching $cache(b, i) = UB_{\max}(b, i)$. Since the IP uses the current upper bounds on the true valuations, and upper bounds only decrease, the value of the IP solutions will only decrease. That is, in a later call to SELECTQUERY, it will be the case that $UB_{\max}(b, i) \leq cache(b, i)$. Therefore, if $cache(b, i) < UB_{\max}$, the elicitor can infer that (b, i) is not in a high-value candidate, without computing $UB_{\max}(b, i)$.

Our experiments (Figure 4) show that in our implementation, the implicit representation of candidates is faster than the explicit one already with three agents ($n = 3$). With 5 agents, the implicit approach is several orders of magnitude faster.

5. OMNISCIENCE

So far we have shown that the elicitors that we designed save most of the preference revelation. However, the question remains: what is the best that any elicitor could do?⁶ To answer this question, we examine the setting where the elicitor is *omniscient*: it knows every agent’s valuation function exactly. However, it must prove to a non-omniscient observer (say, the Federal Trade Commission) that it is conducting the auction correctly. That is, it must provide a certificate consisting of value queries and their answers. The elicitor tries to minimize the length of the certificate. This corresponds to the normal non-omniscient elicitor above which tries to

⁶It is known that *in the worst case*, exponential communication is required to find an (even approximately) optimal allocation in a combinatorial auctions—no matter what query types and elicitation policy is used [16].

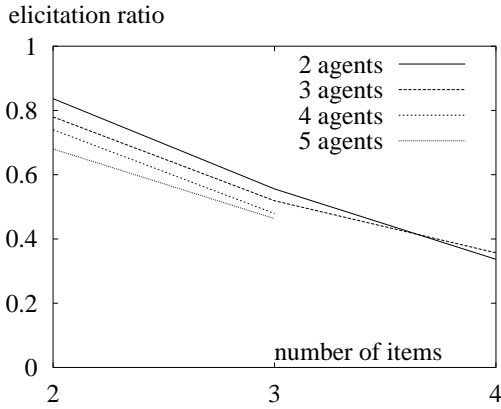


Figure 5: Omniscient elicitor.

ask as few queries as possible. Therefore, the performance of the omniscient elicitor provides a lower bound on that of any real—i.e., non-omniscient—elicitor.

We implemented a routine based on the IDA* search algorithm [12] to find the shortest certificate. Firstly, we use the fact that there exists a shortest certificate where the grand bundle value is queried from every agent (this was proven as Proposition 6 in [11]). Therefore, we immediately put (K, i) into the certificate for every agent i .

In the search, we define a lexicographic order on queries. This means that the search space is a tree, rather than a graph, eliminating any revisiting of nodes.

As is usual for IDA* search, we developed an admissible heuristic h for guiding the search. At any given node of the search, the h -value gives a lower bound on the number of further queries needed to obtain a certificate. We based our heuristic on the observation that for a certificate to prove that the optimal candidate dominates *all* candidates, it must at least prove that the optimal candidate dominates *any given* candidate. For every candidate c , we compute how many queries would be needed to separate c from the optimal candidate. The heuristic function h is the maximum over all c .⁷

This is clearly a weak heuristic: h is at most $2n$ and does not vary as the number of items grows, whereas we expect the certificate length to grow exponentially in the number of items! Instead, the major strength of the heuristic lies in its ability to discover that no certificate is reachable. Future work involves developing stronger heuristics.

If there are multiple optimal candidates, then we run IDA* for each optimum, in parallel, and use the best solution.

Figure 5 shows the elicitation ratio of the omniscient elicitor. Given these results, we can evaluate how much we might hope to improve the real (non-omniscient) elicitation algorithms presented above. The answer is: a bit, but not much. For example, with 3 agents and 4 items, the high-value candidate elicitation policy launches 50% of all queries, and the omniscient elicitor launches 33% of all queries. In addition to the difference being rather small, it is highly unlikely that any non-omniscient policy could do as well as the omniscient algorithm.

⁷Because of the lexicographic order on queries, it may be that there is a candidate c which cannot be proven suboptimal using only queries still reachable in the current branch of the search space. In that case, the heuristic returns ∞ immediately.

6. OPTIMAL ELICITATION ALGORITHM

As we show in this section, one can also design an optimal non-omniscient elicitation algorithm. Unfortunately, it will be completely intractable even on the smallest instances. However, we hope it will motivate new, better elicitation algorithms.

Consider the game where the elicitor is asking questions, and the adversary answers them. The game ends when the revealed information constitutes a certificate. The goal of the elicitor is to ask as few questions as possible; the goal of the adversary is to force the elicitor to ask as many as possible. In order to guarantee the fewest number of queries, the elicitor can look through the gametree and pick the shortest path (from the root to a leaf). As we will see in the next section, the elicitor has no pure strategy against the adversary that does better than blindly eliciting the value of every bundle, but better mixed (i.e., probabilistic) strategies exist—for example, the unrestricted random policy.

In the auction setting, the elicitor is not playing against an adversary, but rather against a random distribution. In its gametree search, it should therefore assume not the worst-case answer, but rather a probability distribution over possible answers.

Unfortunately, a full game tree search is infeasible: at the root, the elicitor can choose among $n(2^k - 1)$ queries. At the next level, the adversary can choose among MAXBID possible answers. For most of the answers, the elicitor has $n(2^k - 1) - 1$ queries left to ask. And so on. Even with very aggressive pruning, the combinatorics are almost certainly too large to solve this problem optimally except on the smallest instances.

7. UNIVERSAL REVELATION REDUCERS

So far we have presented elicitation policies that, on average over instances, save a large amount of preference revelation. Now we ask the question: Do there exist *universal* elicitors, that is, elicitors that save revelation on all instances (excepting those where even the omniscient elicitor must reveal everything)?

DEFINITION 2. (*universal revelation reducer*)

A universal revelation reducer is an elicitation policy with the following property: given a problem instance, it can guarantee (always in the deterministic case; in expectation over the random choices in the randomized case) saving some elicitation over full revelation—provided the shortest certificate is shorter than full revelation. More formally, if $q_{min} < Q$, the policy makes $q < Q$ queries.

THEOREM 1. The unrestricted random elicitation policy is a universal revelation reducer.

PROOF. From Proposition 1 we have a bound on the number of queries asked in expectation. This bound is less than Q provided $q_{min} < Q$, so in expectation, the policy asks fewer than all the queries. \square

Interestingly, *deterministic* universal revelation reducers do not exist:

THEOREM 2. No deterministic value query policy is a universal revelation reducer.

We prove that no deterministic policy can guarantee saving any elicitation by constructing a fooling set. A minimal fooling set can be built for the case where there are 2 items and 2 agents; we show a somewhat more general result, that an equivalent fooling set can also be built for the case where there are 2 items, but $n \geq 2$ agents.

The fooling set R_{fool} consists of valuation functions of the form:

$$\begin{aligned} \forall i \in N \quad v_i(ab) &= 2 \\ \forall a \in K, \forall i \in N \quad v_i(a) &\in \{0, 1\} \end{aligned}$$

And either (a):

$$\exists a \in K \text{ s.t. } \sum_i v_i(a) = 2 \quad \text{and} \\ \sum_i v_i(b) = 0$$

Or (b):

$$\exists i \in N \text{ s.t. } v_i(a) = v_i(b) = 1 \quad \text{and} \\ \forall j \neq i, \quad v_j(a) = v_j(b) = 0$$

That is, either one of the items has value 1 to two agents, and no agent wants the other item; or one of the agents is happy with either item and no other agent wants only a single item. Thus, there are exactly two 1 values.

The socially optimal allocation is to give the entire set of items K to one of the agents. Otherwise, we can only get value 1 from the allocation.

LEMMA 1. *Each instance in R_{fool} has a certificate that does not fully reveal the agents' valuations.*

PROOF. (of Lemma 1)

Any instance in the fooling set can be solved by revealing all the zero values and no 1 values. If we have an instance of type (a), we have now revealed that any allocation of item a to an agent i , and item b to another agent j , has value at most $UB_i(a) + UB_j(b) = 2 + 0 \leq v_i(ab)$. Similarly for instances of type (b). Thus, we have a certificate. We necessarily have two bundles of value 1, and after revealing the zeros, an observer knows only that those two bundles have value at most two. Therefore, we have not revealed the values of all bundles. \square

PROOF. (of Theorem 2)

The proof that no deterministic universal reducer exists operates under the model that the adversary can choose the instance during the execution of the policy. Because the policy is deterministic, this is equivalent to having the adversary examine the policy and choose an instance *a priori*.

To the first query $v_i(S)$ where S is either a or b , the adversary will return 1.

From then on, the adversary will return 0, until the policy asks either $v_i(\bar{S})$ or until it asks $v_j(S)$ having already asked all $v_l(S)$ for $i \neq l \neq j$, at which point the adversary will return 1. In other words, the adversary forces the policy to ask both 1 values.

If the certificate thus chosen reveals both $v_i(S) = 1$ and $v_i(\bar{S}) = 1$, then it must reveal $v_j(S) = v_j(\bar{S}) = 0$ for all $j \neq i$. Otherwise, the allocation of S (resp. \bar{S}) to agent i and \bar{S} (resp. S) to agent j has value up to $UB_i(S) + UB_j(\bar{S}) = 1 + 2 > 2$ which contradicts that we have a certificate. Thus the certificate must reveal the value of all bundles.

If the certificate instead reveals both $v_i(S) = 1$ and $v_j(S) = 1$, then it must reveal $v_l(\bar{S}) = 0$ for all l . Otherwise, we do not have a certificate. In addition, if the adversary chose to answer $v_j(S) = 1$, then the policy asked $v_l(S)$ for all $i \neq l \neq j$. Thus, the certificate reveals the value of all bundles. \square

The theorem, as stated, proves that no deterministic value query policy saves on all instance sizes: given an algorithm, the adversary can provide an instance that fools it. The proof is slightly stronger: even if we require the adversary to provide an instance

with n agents, the adversary can fool the elicitor. We conjecture the yet-stronger result that given any deterministic value query policy, any $n \geq 2$, and any $k \geq 2$, the adversary is still able to provide an instance of the appropriate size that fools the elicitor.

8. CONCLUSIONS AND FUTURE RESEARCH

The work presented here showed how to drastically reduce the amount of valuation information bidders in a combinatorial auction need to reveal (and thus the amount they have to compute). We showed that while our elicitors do quite well, there is still a gap between their current performance and the best possible, leaving us to consider developing even better elicitation policies. However, closing that gap may be highly computationally expensive. To scale to larger auctions, the techniques described here-in would need to be sped up dramatically, and doing so may require compromising on the elicitation ratio.

Further research should also focus on a tighter theoretical analysis on elicitation algorithms, both on upper bounds (we have an upper bound for the unrestricted random policy but not the others) and on lower bounds (we can compute one using IDA*, but that technique is intractable).

We focused here on value queries only. It is not unlikely that some other query types could be used to some advantage [6, 7]. Motivated by the same concerns as our research on elicitors, considerable research has been devoted to ascending auctions (e.g. [2–4, 17, 18, 25]). Those mechanisms can be viewed as a special case of the elicitation framework, where the queries are demand queries: If these were the prices, what items would you buy from the auction? Future research includes more deeply understanding the connection between preference elicitation and ascending auctions, and, if appropriate, designing auction mechanisms that combine features of both families.

9. REFERENCES

- [1] A. Andersson, M. Tenhunen, and F. Ygge. Integer programming for combinatorial auction winner determination. ICMAS, 2000.
- [2] L. M. Ausubel and P. Milgrom. Ascending auctions with package bidding. Technical report draft, June 7, 2001.
- [3] S. Bikhchandani, S. de Vries, J. Schummer, and R. V. Vohra. Linear programming and Vickrey auctions. Draft, 2001.
- [4] S. Bikhchandani and J. Ostroy. The package assignment model. UCLA Working Paper Series, mimeo, 2001.
- [5] E. H. Clarke. Multipart pricing of public goods. *Public Choice*, 11:17–33, 1971.
- [6] W. Conen and T. Sandholm. Preference elicitation in combinatorial auctions: Extended abstract. ACM-EC, 2001. A more detailed description of the algorithmic aspects appeared in the IJCAI-2001 Workshop on Economic Agents, Models, and Mechanisms, pp. 71–80.
- [7] W. Conen and T. Sandholm. Partial-revelation VCG mechanism for combinatorial auctions. AAAI, 2002.
- [8] Y. Fujishima, K. Leyton-Brown, and Y. Shoham. Taming the computational complexity of combinatorial auctions: Optimal and approximate approaches. IJCAI, 1999.
- [9] T. Groves. Incentives in teams. *Econometrica*, 41:617–631, 1973.
- [10] H. Hoos and C. Boutilier. Bidding languages for combinatorial auctions. IJCAI, 2001.
- [11] B. Hudson and T. Sandholm. Effectiveness of preference elicitation in combinatorial auctions. AAMAS-02 workshop

- on Agent-Mediated Electronic Commerce, 2002. Extended version: Carnegie Mellon Univ., Computer Science Department, CMU-CS-02-124, March 2002.
- [12] R. E. Korf. Depth-first iterative-deepening: An optimal admissible tree search. *Artificial Intelligence*, 27(1):97–109, 1985.
 - [13] K. Larson and T. Sandholm. Costly valuation computation in auctions. TARK-VIII, 2001.
 - [14] K. Leyton-Brown, M. Pearson, and Y. Shoham. Towards a universal test suite for combinatorial auction algorithms. ACM-EC, 2000.
 - [15] N. Nisan. Bidding and allocation in combinatorial auctions. ACM-EC, 2000.
 - [16] N. Nisan and I. Segal. The communication complexity of efficient allocation problems. Draft March 5th, 2002.
 - [17] D. C. Parkes. iBundle: An efficient ascending price bundle auction. ACM-EC, 1999.
 - [18] D. C. Parkes. Optimal auction design for agents with hard valuation problems. IJCAI workshop on Agent-Mediated Electronic Commerce, 1999.
 - [19] T. Sandholm. An implementation of the contract net protocol based on marginal cost calculations. AAAI, 1993.
 - [20] 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. Early version: ICMAS-96.
 - [21] T. Sandholm. Algorithm for optimal winner determination in combinatorial auctions. *Artificial Intelligence*, 135:1–54, Jan. 2002. Also: Washington Univ., Dept. of Computer Science, tech report WUCS-99-01, 1/28/1999; and IJCAI, 1999.
 - [22] T. Sandholm. eMediator: A next generation electronic commerce server. *Computational Intelligence*, 2002. Special issue on Agent Technology for Electronic Commerce. To appear. Early versions: AGENTS, 2000; AAAI Workshop on AI in Electronic Commerce, 1999; and Washington University, St. Louis, Dept. of Computer Science technical report WU-CS-99-02, Jan. 1999.
 - [23] T. Smith, T. Sandholm, and R. Simmons. Constructing and clearing combinatorial exchanges using preference elicitation. AAAI workshop on Preferences in AI and CP: Symbolic Approaches, 2002.
 - [24] W. Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, 16:8–37, 1961.
 - [25] P. R. Wurman and M. P. Wellman. AkBA: A progressive, anonymous-price combinatorial auction. ACM-EC, 2000.