

Approximating Revenue-Maximizing Combinatorial Auctions

Anton Likhodedov and Tuomas Sandholm

Carnegie Mellon University
Computer Science Department
5000 Forbes Avenue
Pittsburgh, PA 15213
{likh,sandholm}@cs.cmu.edu

Abstract

Designing revenue-maximizing combinatorial auctions (CAs) is a recognized open problem in mechanism design. It is unsolved even for two bidders and two items for sale. Rather than attempting to characterize the optimal auction, we focus on designing approximations (suboptimal auction mechanisms which yield high revenue).

Our approximations belong to the family of *virtual valuations combinatorial auctions (VVCA)*. VVCA is a Vickrey-Clarke-Groves (VCG) mechanism run on virtual valuations that are linear transformations of the bidders' real valuations.

We pursue two approaches to constructing approximately optimal CAs. The first is to construct a VVCA with worst-case and average-case performance guarantees. We give a logarithmic approximation auction for basic important special cases of the problem: 1) limited supply of items on sale with additive valuations and 2) unlimited supply. The second approach is to search the parameter space of VVCAs in order to obtain high-revenue mechanisms for the general problem. We introduce a series of increasingly sophisticated algorithms that use economic insights to guide the search and thus reduce the computational complexity. Our experiments demonstrate that in many cases these algorithms perform almost as well as the optimal VVCA, yield a substantial increase in revenue over the VCG mechanism and drastically outperform the straightforward algorithms in run-time.

1 Introduction

Combinatorial auctions (CAs), where agents can bid on bundles of items, are popular autonomy-preserving ways of allocating items (goods, tasks, resources, services, etc.). They are relatively efficient both in terms of process and outcome, and are extensively used in a variety of allocation problems in economics and computer science.

One of the main open problems in CAs (and the whole field of mechanism design) is designing *optimal auctions*, that is, auctions that maximize the seller's expected revenue. A major advance on the problem was the full characterization of 1-item auctions (Myerson 1981), later extended to the case of selling multiple units of the same item. However, the characterization of multi-item auctions has been obtained only for very specialized models (two

items, two agents drawing valuations for the items from the same binary distribution (Avery & Hendershott 2000; Armstrong 2000)).

Rather than attempting to characterize the optimal CA, we focus on designing approximations (suboptimal auction mechanisms which yield high revenue). Our approximations belong to the family of *virtual valuations combinatorial auctions (VVCA)* (Likhodedov & Sandholm 2004). VVCA is a Vickrey-Clarke-Groves (VCG) mechanism run on virtual valuations that are linear transformations of the bidders' real valuations. The coefficients of these linear transformations parameterize the family of VVCAs. The restriction to linear transformations is motivated by incentive compatibility.

In this paper we use the VVCA concept to design the mechanisms, which yield high revenue. In Section 3 we design a randomized mechanism, based on VVCAs that yields a logarithmic worst-case approximation and deterministic VVCAs that yield a logarithmic average case approximation to the optimal auction, for the basic settings of 1) items in limited supply and additive valuations (no complementary or substitutable items), and 2) items in unlimited supply and general valuations.

In Section 4 we pursue the approach of designing the high-revenue auctions automatically. We present increasingly sophisticated algorithms for searching the parametric families of VVCAs and a more general family of affine maximizer auctions (AMA) for good parameters for the specific setting (seller's prior over the bidders' valuations). The algorithms use economic insights to navigate the search space efficiently in order to enhance computational speed. The experiments show that they yield significantly higher revenue than the VCG, that they scale much better than the previous automated design algorithms for this problem (Conitzer & Sandholm 2003; Likhodedov & Sandholm 2004), and that the more sophisticated methods indeed drastically outperform the more obvious ones in both absolute run-time and anytime performance.

2 Notation and framework

We study a setting with one seller (index 0 refers to the seller), a set N of n bidders, and a set $G = (g_1, \dots, g_m)$ of heterogeneous items on sale.

In an auction, the bidders submit bids for the bundles of

items and the auction rules determine the allocation a and the payments t , where a_i is the bundle of goods that bidder i receives and t_i is the payment by bidder i .

2.1 Valuations and mechanism design principles

We make the standard assumption that each bidder i has a quasi-linear utility function $u_i = v_i(a) - t_i$, where $v_i(a)$ is the valuation of bidder i for allocation a . Each bidder's true valuations are private information. Thus a bidder might strategically misrepresent her valuations in order to gain higher utility.

As is standard in the (computer science) mechanism design literature, we focus on *ex-post incentive compatible (IC)* mechanisms, that is, mechanisms where each bidder maximizes her utility by bidding truthfully, regardless of what valuations the other bidders reveal. Such mechanisms are also called dominant-strategy mechanisms. They are robust in the sense that the bidders do not benefit from counter-speculating each others' valuations and rationality. Limiting the scope to truthful mechanisms is without loss of generality: the well-known *revelation principle* shows that anything that can be accomplished with an arbitrary mechanism can also be accomplished with a truth-promoting mechanism (Mas-Colell, Whinston, & Green 1995).

As usual, we also require that the mechanism be *ex post individually rational (IR)*: each bidder is no worse off by participating than not participating, for all possible valuation revelations of the other bidders.

2.2 Affine maximizer auctions (AMA)

An important family of auctions that satisfies the above conditions is the affine maximizer auction (AMA).

Definition 2.1 Affine maximizer auction (AMA). Each bidder i submits her valuations, v_i . The allocation, a , is computed so as to maximize¹

$$SW_{\lambda}^{\mu}(a) = \sum_{i=0}^n \mu_i v_i(a) + \lambda(a) \quad (2.1)$$

Here μ_i are positive numbers and $\lambda(a)$ is an arbitrary function of allocation. Both μ and λ are chosen by the auction designer, and they are common knowledge. The payments are

$$t_i = \frac{1}{\mu_i} \left[\sum_{j \neq i} \mu_j v_j(a_{-i}) + \lambda(a_{-i}) - \sum_{j \neq i} \mu_j v_j(a) - \lambda(a) \right]$$

where

$$a_{-i} = \operatorname{argmax}_{\tilde{a}} \sum_{j=0, j \neq i}^n \mu_j v_j(\tilde{a}) + \lambda(\tilde{a}) \quad (2.2)$$

An AMA with all $\mu = 1$ and $\lambda \equiv 0$ is the famous *Vickrey-Clarke-Groves (VCG)* mechanism (Vickrey 1961; Clarke 1971; Groves 1973), aka *Generalized Vickrey Auction*. The winning allocation of the VCG is efficient, that is, it maximizes the sum of the bidders' true valuations.

¹Throughout this paper, ties in allocation rules can be broken arbitrarily.

The AMA was introduced by Roberts (Roberts 1979). He proved that AMAs are the only ex-post incentive compatible mechanisms over unrestricted domains of valuations. The valuations in combinatorial auction (CA) domain are not unrestricted because they satisfy the following restrictions: 1) *no externalities*: the valuation of any bidder i for each allocation a depends only on the bundle a_i that the bidder receives, not on how the items that i does not receive get allocated, 2) *free disposal*: the value of a subset of a bundle is less than or equal to the value of a bundle ($\forall b' \subset b, v_i(b') \leq v_i(b)$), and 3) the valuation for the empty bundle is 0. However, even in the CA domain the set of incentive compatible auction mechanisms is almost limited to AMAs: (Lavi, Mu'Alem, & Nisan 2003) showed that under certain natural assumptions, every *ex post* incentive compatible CA is an AMA.

Therefore it is natural to look for high-revenue CAs within the AMA family. This approach reduces the problem of finding a good CA to a search in the space of AMA's parameters, as studied in (Likhodedov & Sandholm 2004).

In this paper we focus on an important subclass of AMA—*virtual valuations combinatorial auctions (VVCA)*—introduced in (Likhodedov & Sandholm 2004).

Definition 2.2 Virtual valuations combinatorial auction (VVCA). The mechanism computes an allocation a that maximizes

$$SW_{\lambda}^{\mu}(a) = \sum_{i=0}^n [\mu_i v_i(a) + \lambda_i(a)] \quad (2.3)$$

Here μ are positive, $\lambda_i(a) = c_{\{i,b\}}$ for all allocations that give bidder i exactly bundle b , and $\lambda_i(a) = 0$ otherwise. The μ and $c_{\{i,b\}}$ are parameters chosen by the auction designer, and they are common knowledge.

The payment rule is

$$t_i = \frac{1}{\mu_i} \left[\sum_{j \neq i} [\mu_j v_j(a_{-i}) + \lambda_j(a_{-i})] - \sum_{j \neq i} [\mu_j v_j(a) + \lambda_j(a)] - \lambda_i(a) \right] \quad (2.4)$$

The revenue of the seller is the sum of payments of bidders, that is $\sum_{i=1}^n t_i$.

The VVCA can be thought of as Vickrey-Clarke-Groves mechanism run on bidders' *virtual valuations* rather than their real valuations: the mechanism replaces the valuation of bidder i , $v_i(a)$, with the virtual valuation $\mu_i v_i(a) + \lambda_i(a)$. This technique allows one to apply the ideas of the Myerson revenue-maximizing single-item auction (Myerson 1981) to the case of CAs: the revenue can be increased by setting reserve prices and boosting the valuations of disadvantaged bidders (disadvantaged bidders are ones that are likely to have low true valuations). Both of these levers increase competition in the auction and thus increase the expected revenue of the seller. In the VVCA, these levers are controlled by setting the parameters μ and λ .

2.3 Average-case and worst-case framework

The problem of designing high-revenue CAs can be analyzed in two different frameworks:

1. Average case analysis is the standard approach in designing high-revenue auctions, both in economics and computer science. In this setup we assume that the valuations of the bidders are drawn from some underlying probability distributions (not necessarily the same for different bidders), and the auction designer knows the distributions, but not the exact draws, i.e. valuations, of the bidders. We do not assume that the bidders know each others' distributions. In this framework the goal is to construct an auction which yields high revenue on average with respect to the distributions.
2. Worst case analysis of the problem has sometimes been used in computer science: in that framework the objective is to construct an auction with worst case performance guarantees (Goldberg, Hartline, & Wright 2001; Guruswami *et al.* 2005). The advantage is that the design typically does not require complete knowledge of the underlying distributions, although the mechanisms are not completely prior-free. A disadvantage is lower expected revenue. An essential feature of auctions with worst-case performance guarantees is randomization: in many cases deterministic auctions perform far worse than randomized ones with respect to the worst case performance objective (Goldberg, Hartline, & Wright 2001).

In this paper we take the more standard approach of average case analysis. However, some of the results also hold in the worst case framework, as the propositions will point out.

In the next section we present theoretical results on designing CAs whose revenue is within a provable bound of optimal. In the section after that, we present iterative algorithms for automatically designing CAs with high expected revenue.

3 Logarithmic approximations to the optimal CA

In this section we study two basic subclasses of CA setting. For each setting, we derive VVCAs that guarantee average-case and worst-case revenue that are provably within a bound of optimal. It should be noted that the revenue performance of VCG auction can be arbitrarily bad (Conitzer & Sandholm 2004).

3.1 Additive valuations

In this section we study the special case where valuations are additive ($\forall i \in N, \forall b \in 2^G, v_i(b) = \sum_{g \in b} v_i(\{g\})$). In addition we make the following mild natural assumptions about the priors:

1. Define l to be the lowest possible valuation of any bidder for any *individual* item. We assume $l > 0$, and that the auction designer knows l . (l can be arbitrary small.)
2. Define h to be the highest possible valuation of any bidder for any *individual* item. We assume that the auction designer knows h .

We now construct a CA which guarantees a fraction $\frac{1}{2+2\lceil \log(h/l) \rceil}$ of the revenue of the optimal CA, even in the worst case. This generalizes the result in (Guruswami *et al.* 2005), which was for bidders that only demand one item each.

Proposition 3.1 *Let $VVCA^k$ be the virtual valuations combinatorial auction with the following parameters:*

1. $\mu_i = 1$ for all i .
2. $c_{\{i,b\}} = 0$ for all $i > 0$ and all b
3. $c_{\{0,g_j\}} = l \cdot 2^k$ for all items g_j in G
4. $c_{\{0,b\}} = |b| \cdot l \cdot 2^k$ for all bundles b , where $|b|$ is the number of items in b

In other words, $VVCA^k$ is a Vickrey-Clarke-Groves auction in which the seller submits a bid of $l \cdot 2^k$ for every item (and any number of those bids can be accepted).

Consider the mechanism M , which uniformly randomly selects k from $\{0, 1, \dots, \lceil \log(h/l) \rceil\}$ and runs $VVCA^k$. Then M is ex-post incentive compatible, ex-post individually rational and for any given set of valuations v yields the expected revenue of at least

$$\frac{R_{opt}}{2 + 2\lceil \log(h/l) \rceil}$$

where R_{opt} is the revenue of the optimal CA (note, that the bound hold for all sets of valuations v and the expectation is taken w.r.t. k , which is the only source of randomness).

Before giving the proof, we need to introduce the following notation. Let a^{eff} be an efficient allocation, a^k be the winning allocation of $VVCA^k$, and a_{-i}^k be the allocation that would have won had bidder i not submitted any bids. Let $v_N(g_j)$ be the highest bid for item g_j : $v_N(g_j) = \max_{i' \in N} v_{i'}(g_j)$. Also let $v_{N \cup \{0\}}^k(g_j)$ be the highest bid for item g_j , including the bid of the seller: $v_{N \cup \{0\}}^k(g_j) = \max\{v_N(g_j), l \cdot 2^k\}$. Finally, let $v_{N \cup \{0\} \setminus \{i\}}^k(g_j)$ be the highest bid for item g_j , including the bid of the seller, but excluding the bid of bidder i : $v_{N \cup \{0\} \setminus \{i\}}^k(g_j) = \max_{i' \in \{1 \dots n\} \setminus \{i\}} \{v_{i'}(g_j), l \cdot 2^k\}$.

Because the valuations are additive, a^{eff} allocates every item g_j according to $v_N(g_j)$, that is, to bidder i' ($1 \leq i' \leq n$) that submitted $v_N(g_j)$. Since the seller's bids are also additive, a^k allocates every item g_j according to $v_{N \cup \{0\}}^k(g_j)$ and a_{-i}^k allocates every item g_j according to $v_{N \cup \{0\} \setminus \{i\}}^k(g_j)$.

We will use the following lemma in the proof.

Lemma 3.1 *Consider a set of bidders' valuations v . If bidder i wins bundle b in $VVCA^k$, she pays at least $|b| \cdot l \cdot 2^k$.*

Proof. By Equation (2.4), the payment of bidder i is

$$\begin{aligned} t_i &= SW_{\lambda}^{\mu}(a_{-i}^k) - SW_{\lambda}^{\mu}(a^k) + v_i(b) \\ &= \left(\sum_{g_j \notin b} v_{N \cup \{0\} \setminus \{i\}}^k(g_j) + \sum_{g_j \in b} v_{N \cup \{0\} \setminus \{i\}}^k(g_j) \right) \\ &\quad - \sum_{j=1}^m v_{N \cup \{0\}}^k(g_j) + v_i(b) \end{aligned}$$

Obviously $v_i(b) = \sum_{g_j \in b} v_{N \cup \{0\}}^k(g_j)$. Thus the last two terms simplify to

$$-\sum_{j=1}^m v_{N \cup \{0\}}^k(g_j) + v_i(b) = -\sum_{g_j \notin b} v_{N \cup \{0\}}^k(g_j)$$

For the items which are not allocated to bidder i we have

$$\sum_{g_j \notin b} v_{N \cup \{0\} \setminus \{i\}}^k(g_j) = \sum_{g_j \notin b} v_{N \cup \{0\}}^k(g_j)$$

Therefore

$$t_i = \sum_{g_j \in b} v_{N \cup \{0\} \setminus \{i\}}^k(g_j)$$

which by definition of $v_{N \cup \{0\} \setminus \{i\}}^k$ is no less than $|b| \cdot l \cdot 2^k$. \square

Proof of Proposition 3.1. Since every $VVCA^k$ is *ex-post* incentive compatible and *ex-post* individually rational and M is a randomization over $VVCA^k$, M is also *ex-post* incentive compatible and *ex-post* individually rational.

We now prove the revenue bound. By Lemma 3.1, any bidder that wins a bundle, b , in $VVCA^k$, pays at least $|b| \cdot l \cdot 2^k$. Because valuations are additive, a^k allocates every item g_j to the same bidder as a^{eff} if $v_N(g_j) \geq l \cdot 2^k$, and leaves the item for the seller otherwise. Therefore the revenue in $VVCA^k$ is at least $n_k \cdot l \cdot 2^k$, where n_k is the number of such g_j that $v_N(g_j) \geq l \cdot 2^k$:

$$n_k = \sum_{j=1}^m I[v_N(g_j) \geq l \cdot 2^k],$$

where I is an indicator function which equals 1 if its argument is true and 0 otherwise.

So, when the valuations of bidders are given by v , the expected revenue of mechanism M , $E_k[R_M(v)]$, is at least

$$\begin{aligned} & \frac{1}{1 + \lfloor \log(h/l) \rfloor} \sum_{k=0}^{\lfloor \log(h/l) \rfloor} l \cdot 2^k \cdot \sum_{j=1}^m I[v_N(g_j) \geq l \cdot 2^k] = \\ & \frac{1}{1 + \lfloor \log(h/l) \rfloor} \sum_{j=1}^m \sum_{k=0}^{\lfloor \log(h/l) \rfloor} I[v_N(g_j) \geq l \cdot 2^k] l \cdot 2^k \quad (3.1) \end{aligned}$$

The sum on the right of (3.1) can be bounded as follows

$$\begin{aligned} v_N(g_j) & \leq l + \sum_{k=0}^{\lfloor \log(h/l) \rfloor} I[v_N(g_j) \geq l \cdot 2^k] \cdot l \cdot 2^k \quad (3.2) \\ & \leq 2 \cdot \sum_{k=0}^{\lfloor \log(h/l) \rfloor} I[v_N(g_j) \geq l \cdot 2^k] \cdot l \cdot 2^k \end{aligned}$$

Substituting (3.2) into (3.1) we obtain

$$E_k[R_M(v)] \geq \frac{1}{2 + 2\lfloor \log(h/l) \rfloor} \sum_{j=1}^m v_N(g_j) \quad (3.3)$$

Here, $\sum_{j=1}^m v_N(g_j)$ is the welfare of the efficient allocation. No individually rational auction can yield more revenue than that. Therefore the revenue of the optimal auction is bounded from above by $\sum_{j=1}^m v_N(g_j)$. It follows that

$$E_k[R_M(v)] \geq \frac{R_{opt}(v)}{2 + 2\lfloor \log(h/l) \rfloor}. \quad \square \quad (3.4)$$

The same bound can also be made to hold in the average-case framework with a *deterministic* CA:

Corollary 3.1 *There exists such k that $VVCA^k$ yields a fraction $\frac{1}{2 + 2\lfloor \log(h/l) \rfloor}$ of the revenue of the optimal auction on expected revenue basis.*

Proof. By construction of M in Proposition 3.1, we have

$$E_k[R_M(v)] = \frac{1}{1 + \lfloor \log(h/l) \rfloor} \sum_{k=0}^{\lfloor \log(h/l) \rfloor} R_{VVCA^k}(v)$$

Substituting $E_k[R_M(v)]$ into (3.4) and taking expectations over v we obtain

$$\frac{\sum_{k=0}^{\lfloor \log(h/l) \rfloor} E_v[R_{VVCA^k}(v)]}{1 + \lfloor \log(h/l) \rfloor} \geq E_v[R_{opt}(v)]$$

Since the sum of $VVCA^k$ contains exactly $1 + \lfloor \log(h/l) \rfloor$ terms, there exists such k_0 that

$$\frac{E_v[R_{VVCA^{k_0}}(v)]}{1 + \lfloor \log(h/l) \rfloor} \geq E_v[R_{opt}(v)]$$

k_0 can be found by enumeration of all $VVCA^k$ and evaluating their expected revenues. \square

The logarithmic bounds in Proposition 3.1 and Corollary 3.1 were obtained by comparing revenue of our auctions to the welfare of an efficient allocation, $SW(a^{eff})$, which obviously bounds the revenue of any individually rational auction. This proof technique cannot get us past the logarithmic approximation, as demonstrated by the following example.

Example 3.1 *Consider an n -item auction with n bidders. Assume the valuation of bidder b_i for item g_j is drawn from distribution F_i with the density*

$$f_i(v_i) = \begin{cases} \frac{h}{(h-1)v_i^2} & \text{for } v_i \in [1, h] \\ 0 & \text{otherwise} \end{cases}$$

The valuations of other bidders for item g_j are 0. The valuations for the bundles are additive.

In the above setup no incentive compatible individually rational mechanism can raise more than

$$\left(1 - \frac{1}{h}\right) \cdot \frac{E_v[SW(a^{eff})]}{\ln h}$$

on expected revenue basis. Due to limited space, we omit the proof.

3.2 Unlimited supply

Another special case of the optimal CA design problem is the case when items are available in unlimited supply: the auctioneer is still selling items g_1, \dots, g_m , but each item is now available in an infinite number of copies. In this setting we assume that each bidder is interested in at most one copy of every item. This is not a restrictive assumption, since the preferences of a bidder who wants several copies of the

same item can be expressed by adding these copies to the set of items G . As in Subsection 3.1, we assume that the lowest and highest possible valuation (for any bidder for any bundle), l and h , are known by the auction designer. We do not assume that valuations are additive.

Since items are available in unlimited supply, there is no competition among the bidders: under the efficient allocation every bidder is allocated her most wanted bid. Due to free disposal, the allocation a^{eff} which allocates a bundle b_G with all items in G to every bidder is also efficient. This allows us to prove the following:

Proposition 3.2 *Let $VVCA^{lk}$ be the virtual valuations combinatorial auction with*

1. $\mu_i = 1$ for all i .
2. $c_{\{i,b\}} = -\infty$ for all $i > 0$ and all $b \neq b_G$
3. $c_{\{i,b_G\}} = -l \cdot 2^k$ for all $i \in N$.

Consider the mechanism M' which uniformly randomly selects k from $\{0, \dots, \lfloor \log(h/l) \rfloor\}$ and runs $VVCA^{lk}$. M' is ex-post incentive compatible, ex-post individually rational and for every given set of valuations v yields expected revenue of at least

$$\frac{R_{opt}}{2 + 2 \lfloor \log(h/l) \rfloor}$$

where R_{opt} is the revenue of the optimal auction.

Proof. M' is ex post incentive compatible and individually rational because it is a randomization over ex post incentive compatible and individually rational auctions. Let a^k be the winning allocation in $VVCA^{lk}$ and a_{-i}^k be the allocation that would have been optimal had bidder i not submitted any bids. Since there is no competition, a_{-i}^k and a^k are the same for all bidders except for bidder i . By construction of $VVCA^{lk}$, bidder i wins b_G iff $v_i(b_G) \geq -l \cdot 2^k$ and wins nothing otherwise.

Since a_{-i}^k and a^k are equivalent for bidders other than i , the payment of bidder i for bundle b_G is

$$\begin{aligned} t_i &= (SW_{\lambda}^{\mu}(a_{-i}^k) - SW_{\lambda}^{\mu}(a^k)) + v_i(b_G) \\ &= (-v_i(b_G) + l \cdot 2^k) + v_i(b_G) = l \cdot 2^k \end{aligned}$$

Using the notation of Proposition 3.1, the expected revenue of mechanism M' , $E_k[R_M(v)]$, can be written as

$$\begin{aligned} \frac{1}{1 + \lfloor \log(h/l) \rfloor} \sum_{k=0}^{\lfloor \log(h/l) \rfloor} l \cdot 2^k \sum_{i=1}^n I[v_i(b_G) \geq l \cdot 2^k] &\geq \\ \frac{\sum_{i=1}^n v_i(b_G)}{2 + 2 \lfloor \log(h/l) \rfloor} = \frac{SW(a^{eff})}{2 + 2 \lfloor \log(h/l) \rfloor} &\geq \frac{R_{opt}(v)}{2 + 2 \lfloor \log(h/l) \rfloor}. \quad \square \end{aligned}$$

Again, the same bound can be obtained with a *deterministic* mechanism in the average-case model.

Corollary 3.2 *There exists such k that $VVCA^{lk}$ yields fraction $\frac{1}{2+2 \lfloor \log(h/l) \rfloor}$ of the revenue of the optimal auction on an expected revenue basis.*

4 Designing high-revenue auctions algorithmically

In Section 3 we designed several auctions with average-case performance guarantees. However, these guarantees are fairly weak and for most of the problem setups superior mechanisms exist. In this section we suggest several automated approaches for constructing such mechanisms numerically. We focus on average-case analysis.

We first consider general AMAs. The expected revenue is a function of the AMA parameters. Thus the problem of designing a high revenue auction is reduced to a search for the maximum of expected revenue in the AMA parameter space. We implement this search by sampling the valuations from the prior distributions (every sample point is the complete set of valuations of all bidders) and running a hill climbing algorithm in the parameter space. The expected revenue of the AMA with a given set of parameters is estimated by running that AMA on each sample and averaging. (Likhodedov & Sandholm 2004) state the following two main obstacles to this:

1. The revenue surface is not convex in the parameter space and can have many local maxima.
2. There is a large number of parameters: $(n+1)^m$ (one for every possible allocation).

Our experiments, summarized in the table in Subsection 4.1 suggest that the local maxima of the revenue surface are likely not to be significantly inferior that the global maximum. That justifies the following algorithm:

Algorithm 1 (Basic local optimization of AMA)

1. *Sample the valuations from the prior distributions.*
2. *Start at some known AMA (typically VCG or one of the AMAs with average-case performance guarantees from Section 3). Evaluate the mechanism at the sample points.*
3. *Run Fletcher-Reeves conjugate gradient ascent (Stoer & Bulirsch 1980) in the AMA parameter space from the starting point.*

However, Algorithm 1 is still susceptible to the second problem, i.e., the prohibitive number of optimization parameters. (For one, in order to compute the gradient for choosing the direction of the climb at every step, the algorithm must consider an exponential number of parameters.)

To address this problem we introduce new algorithms that guess the climbing direction based on insights drawn from the fact that we are in a CA domain. The idea of the first of these algorithms is from Equation (2.2), i.e., the payment rule of AMA. If the payment, t_i , of bidder i in allocation a is much lower than her valuation for a , one should expect that the her payment could have been increased. The payment can be increased directly only by 1) decreasing $\lambda(a)$, 2) increasing $\lambda(a_{-i})$, or 3) modifying the μ parameters.

Algorithm 2 (Allocation boosting of AMA)

1. *Sample the valuations from the prior distributions.*
2. *Start at some known AMA (typically VCG or one of the auctions from Section 3).*

3. For every sample point, compute the revenue loss on the winning allocation a (variant a) or the second-best allocation (variant b). (The revenue loss from a bidder is the difference between the bidder's valuation and her payment. The revenue loss is the sum of the bidders' revenue losses.) Note that each allocation may be associated with multiple samples. Let the revenue loss of an allocation be the sum of the revenue losses of the samples associated with the allocation. Make a list of allocations in decreasing order of revenue loss.
4. Choose the first allocation, a , from the list. If the list is empty, exit.
5. Run Fletcher-Reeves conjugate gradient ascent in the $\{\mu, \lambda(a)\}$ subspace of the AMA parameter space; leave the other parameters unchanged. If the values of $\{\mu, \lambda(a)\}$ did not change (i.e., we cannot further improve the revenue by modifying $\{\mu, \lambda(a)\}$), remove a from the list and go to step 4. Otherwise go to step 3.

The only parameters considered by Algorithm 2 at each step are the μ and λ corresponding to the winning or second-best allocations. In practice the number of those allocations is small, which dramatically decreases the number of parameters in consideration.

Another computational issue is that evaluating the revenue requires computing the optimal allocation of the AMA, i.e., solving a winner determination problem, which is NP-complete (Rothkopf, Pekeč, & Harstad 1998). While there are quite efficient tree search-based algorithms for the basic CA winner determination problem (for a review, see (Sandholm 2006)), with AMA the parameter λ can be different for every possible allocation, necessitating the explicit enumeration of all allocations in the winner determination. This further hinders the scalability.

To partially solve this problem, and to search in a smaller number of parameters than the number of parameters that AMAs have, we can focus on VVCAs instead (a VVCA has $(n + 1)2^m$ parameters, one for every bidder-bundle pair). The parameters of VVCA are valuation (and not allocation) specific, and all the methods for winner determination which apply to the standard VCG mechanism (such as search algorithms) also apply to VVCA. (In the experiments below we use the dynamic program of (Rothkopf, Pekeč, & Harstad 1998).) Therefore each iteration of the design algorithm will run faster. Below we present the design algorithm for VVCAs (which is similar to Algorithm 2 for AMAs).

Algorithm 3 (Bidder-bundle boosting of VVCA)

1. Sample the valuations from the prior distributions.
2. Start at some known VVCA (typically VCG or one of the auctions from Section 3).
3. For every sample point, compute the payments of winning bidders. For every bidder i winning bundle b and paying t_i , compute $v_i(b) - t_i$, i.e., the revenue loss for that bidder-bundle pair. Sum up the revenue losses over the sample and make a list of bidder-bundle pairs in decreasing order of the revenue loss.
4. Choose the first bidder-bundle pair, $\{i, b\}$, from the list. If the list is empty, exit.

5. Run Fletcher-Reeves conjugate gradient ascent in the $\{\mu, c_{\{i,b\}}\}$ subspace of the VVCA parameter space ($\{i, b\}$ is the bidder-bundle pair which incurs the highest revenue loss). Leave the values of all the other parameters unchanged. If the new values of $\{\mu, c_{\{i,b\}}\}$ do not change (i.e., we cannot improve the revenue further by modifying $\{\mu, c_{\{i,b\}}\}$), remove $\{i, b\}$ from the list and go to step 4. Otherwise go to step 3.

4.1 Experiments

We conducted experiments with the VCG, the "optimal" AMA obtained by grid enumeration of the parameter space, followed by gradient ascent from every grid point (AMA*), the "optimal" VVCA obtained by grid enumeration of that parameter space, followed by gradient ascent from every grid point (VVCA*), and the four algorithms described in this section: basic local optimization of AMA (BLAMA), allocation boosting AMA (ABAMA variants a and b), and bidder-bundle boosting VVCA (BBBVVCA).

The first experiment is with 2 items, g_1 g_2 , and 2 bidders with valuation functions v_1 and v_2 , respectively. Assume $v_1(g_1)$ and $v_1(g_2)$ are drawn from the distribution F_1 . $v_2(g_1)$, and $v_2(g_2)$ are drawn from the distribution F_2 . The valuation of bidder 1 for the bundle of two items is given by $v_1(g_{12}) = v_1(g_1) + v_1(g_2) + c_1$ where c_1 is a complementarity parameter drawn from distribution C . Similarly $v_2(g_{12}) = v_2(g_1) + v_2(g_2) + c_2$ where c_2 is also drawn from C . As test cases we used the three different settings from (Likhodedov & Sandholm 2004). The results for various distributions F_1, F_2, C are given in the following table.

	Example I	Example II	Example III
F_1	$U[0, 1]$	$U[1, 2]$	$U[1, 2]$
F_2	$U[0, 1]$	$U[1, 2]$	$U[1, 5]$
C	0	$U[-1, 1]$	$U[-1, 1]$
VCG	2/3	2.45	2.85
AMA*	0.88	2.79	4.22
VVCA*	0.87	2.79	4.20
BLAMA	0.78	2.78	3.76
ABAMAa	0.78	2.78	3.77
ABAMAb	0.78	2.78	3.76
BBBVVCA	0.79	2.79	3.75

The columns correspond to the three settings. The first three rows specify distributions F_1, F_2 , and C ; the last seven rows give the estimates of the expected revenue of the mechanisms, found by the different algorithms.

The second experiment tested scalability (thus we omit the details of the prior distributions due to lack of space). AMA* and VVCA* are obviously not scalable because the grid search suffers from a total combinatorial explosion. We thus conducted the scalability experiment with the other strategies only, Figure 4.1. All of the techniques yield significantly higher revenue than the VCG. As expected, the economically motivated methods are significantly faster (both in terms of absolute run-time and anytime performance) than the basic hill-climbing procedure. BBBVVCA is the fastest because it has fewest parameters, and does not require exhaustive allocation enumeration at each iteration (for winner determination).

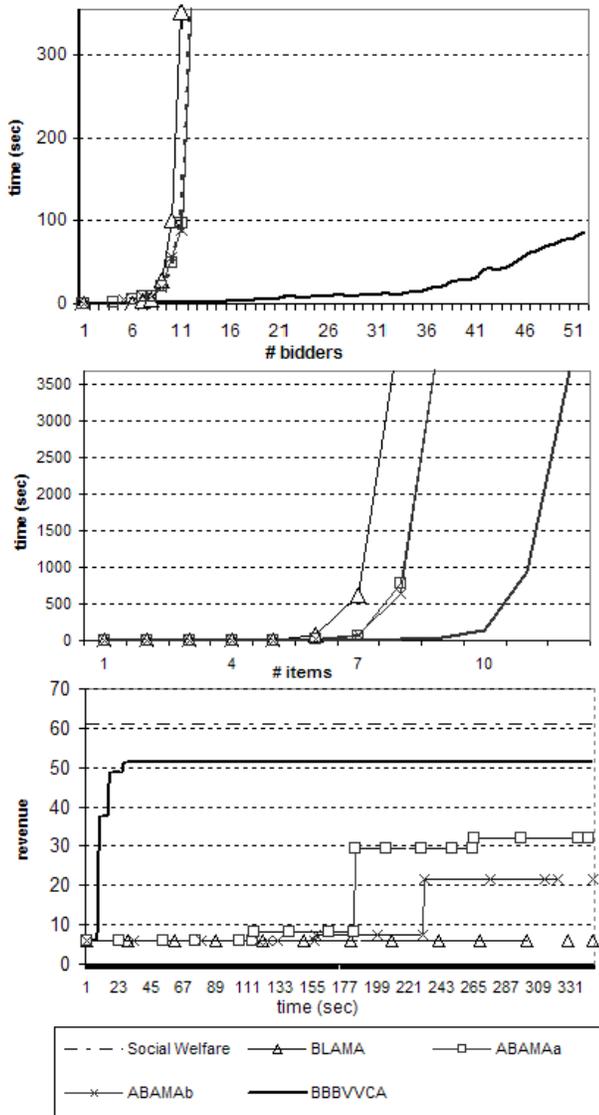


Figure 4.1: Top: Run-time as the number of bidders grows (3 items). (Note that the ABAMAA and ABAMAb curves overlap). Middle: Run-time as the number of items grows (3 bidders). Bottom: Anytime performance (7 items, 7 bidders).

5 Conclusions

The design of optimal (i.e., revenue-maximizing) combinatorial auctions (CAs) is a recognized open research problem. The characterization is open even for two items and two bidders. Our work was motivated by the desire to construct high-revenue CAs in problems beyond that tiny size. We designed randomized *virtual valuations CAs* (VVCAs) that yield a logarithmic worst-case approximation and deterministic VVCAs that yield a logarithmic average case bound from optimal revenue, for the basic settings of 1) items in limited supply and additive valuations, and 2) unlimited supply. We also presented increasingly sophisticated

algorithms for automatically designing high-revenue CAs for the general CA setting. The algorithms use economic insights to navigate the search space efficiently in order to enhance computational speed. The experiments showed that they yield significantly higher revenue than the VCG, that they scale much better than the previous automated design algorithms for this problem, and that the more sophisticated methods indeed drastically outperform the more obvious ones in both absolute run-time and anytime performance.

6 Acknowledgements

This material is based upon work supported by the National Science Foundation under ITR grants IIS-0121678 and IIS-0427858, and a Sloan Fellowship.

References

- Armstrong, M. 2000. Optimal multi-object auctions. *Review of Economic Studies* 67:455–481.
- Avery, C., and Hendershott, T. 2000. Bundling and optimal auctions of multiple products. *Review of Economic Studies* 67:483–497.
- Clarke, E. H. 1971. Multipart pricing of public goods. *Public Choice* 11:17–33.
- Conitzer, V., and Sandholm, T. 2003. Applications of automated mechanism design. In *UAI-03 workshop on Bayesian Modeling Applications*.
- Conitzer, V., and Sandholm, T. 2004. Revenue failures and collusion in combinatorial auctions and exchanges with vcg payments. In *Workshop on Agent-Mediated Electronic Commerce (AMEC VI)*.
- Goldberg, A.; Hartline, J.; and Wright, A. 2001. Competitive auctions and digital goods. In *Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*.
- Groves, T. 1973. Incentives in teams. *Econometrica* 41:617–631.
- Guruswami, V.; Hartline, J.; Karlin, A.; Kempe, D.; Kenyon, C.; and McSherry, F. 2005. On profit-maximizing envy-free pricing. In *Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*.
- Lavi, R.; Mu’Alem, A.; and Nisan, N. 2003. Towards a characterization of truthful combinatorial auctions. In *Proceedings of the Annual Symposium on Foundations of Computer Science (FOCS)*.
- Likhodedov, A., and Sandholm, T. 2004. Boosting revenue in combinatorial auctions. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*.
- Mas-Colell, A.; Whinston, M.; and Green, J. R. 1995. *Microeconomic Theory*. Oxford University Press.
- Mu’alem, A., and Nisan, N. 2002. Truthful approximate mechanisms for restricted combinatorial auctions. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, 379–384.
- Myerson, R. 1981. Optimal auction design. *Mathematics of Operation Research* 6:58–73.
- Roberts, K. 1979. The characterization of implementable choice rules. In *Aggregation and Revelation of Preferences. Papers presented at the first European Summer Workshop of the Econometric Society*, 321–349.
- Rothkopf, M. H.; Pekeč, A.; and Harstad, R. M. 1998. Computationally manageable combinatorial auctions. *Management Science* 44(8):1131–1147.
- Sandholm, T. 2006. Winner determination algorithms. In Cramton, P.; Shoham, Y.; and Steinberg, R., eds., *Combinatorial Auctions*. MIT Press.
- Stoer, J., and Bulirsch, R. 1980. *Introduction to numerical analysis*. Springer-Verlag.
- Vickrey, W. 1961. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance* 16:8–37.