

Methods for Boosting Revenue in 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

We study the recognized open problem of designing revenue-maximizing combinatorial auctions. It is unsolved even for two bidders and two items for sale. Rather than pursuing the pure economic approach of attempting to characterize the optimal auction, we explore techniques for automatically modifying existing mechanisms in a way that increase expected revenue. We introduce a general family of auctions, based on bidder weighting and allocation boosting, which we call *virtual valuations combinatorial auctions (VVCA)*. All auctions in the family are based on the Vickrey-Clarke-Groves (VCG) mechanism, executed on virtual valuations that are linear transformations of the bidders' real valuations. The restriction to linear transformations is motivated by incentive compatibility. The auction family is parameterized by the coefficients in the linear transformations.

The problem of designing a high revenue mechanism is therefore reduced to search in the parameter space of VVCA. We analyze the complexity of the search for the optimal such mechanism and conclude that the search problem is computationally hard. Despite that, optimal parameters for VVCA can be found at least in settings with few items and bidders (the experiments show that VVCA yield a substantial increase in revenue over the traditionally used VCG). In larger auctions locally optimal parameters, which still yield an improvement over VCG, can be found.

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 (Myerson 1981; Vohra 2001). This problem is unsolved even for auctions with two distinct items on sale (e.g. a TV and a VCR) and two agents. 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)).

A related line of research is concerned with bundling decisions of the seller, and the effect they have on revenue. Typically the starting point is the *Vickrey-Clarke-Groves (VCG)* mechanism (Vickrey 1961; Clarke 1971; Groves 1973) (aka. *Generalized Vickrey Auction*), which in the 1-item case is equivalent to the second-price sealed-bid auction. The VCG allocates the items in a way that maximizes the social welfare of the agents (sum of their valuations for the allocated items), and each winning agent pays the minimal valuation that she would have had to bid in order to win the bundle she won. Rather than running a standard VCG, the seller might want to bundle some items together and decide to either sell the whole bundle or not sell it at all (in the extreme case, all items are bundled together). Not surprisingly, a good bundling policy yields higher expected revenue than the VCG. (Palfrey 1983) proves that in certain models, the seller is better off (on an expected revenue basis) bundling the products on sale together when the number of bidders is small and should auction them separately when the number of bidders increases. (Jehiel, ter Vehn, & Moldovanu 2003) show that modifying the objective function of the VCG in favor of allocations that sell all the items together can yield higher expected revenue.

We exploit the benefits of both lines of research. We generalize the idea of artificially making weak bidders more competitive (from Myerson's 1-item auction) and ideas on tweaking the VCG allocation rule (from bundling research). Using these ideas we design a *parametric family* of CAs. Since all attempts to obtain an analytic characterization of the revenue-optimal CA have been unsuccessful, we instead advocate a computer science approach where we algorithmically search for good parameters in this CA family, for the specific setting (seller's prior over the bidders' valuations).¹ Although it may not construct an optimal auction, this practical approach yields significant revenue improvement over the currently used VCG (see Section 6).

¹In this sense, our approach parallels *automated mechanism design (AMD)* (Conitzer & Sandholm 2002). However, in the AMD work so far, the types of agents had to be discretized, and an optimal mechanism was searched for *without any use of exogenous ideas/results on how to improve the mechanism*. Thus that AMD work only scales to small CAs (Conitzer & Sandholm 2003).

2 Framework and notation

We index agents with numbers from 1 to n . Index 0 refers to the seller. Items are indexed with numbers from 1 to k . The set of all items is denoted by $G = (g_1, \dots, g_k)$.

In an auction, bidders submit bids for the items and the auction rules determine the allocation a and the payments t , where a_i is the bundle that bidder i receives and t_i is the payment by bidder i . The pair (a, t) is called the *outcome*.

2.1 Utilities and valuations

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 is the *valuation function* of bidder i . We also make the following standard assumptions: 1) *no externalities*: the valuation of any bidder i for each allocation a depends only on the bundle a_i that 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 an empty bundle is 0.

Let V_i denote the set of all valuation functions for bidder i . V denotes $\times_{i=1 \dots n} V_i$. We make the following standard assumptions from economics:

1. V_i is a convex, compact subset of $\mathbb{R}^{|2^k|}$.
2. Each valuation function v_i is generated from a continuous density f_i and f_i is positive on all V_i .
3. The valuations of different bidders are drawn independently of each other.

A simple example (and an important special case) is the Additive Valuations Model:

Definition 2.1 Additive Valuations Model (AVM).

1. The valuation of bidder i for item j , is a realization of the random variable X_{ij} with the density $f_{ij} : [l_i, r_i] \rightarrow \mathbb{R}$, and
2. The valuation of bidder i for any bundle b is the sum of i 's valuations for the individual items contained in b .

The AVM is the most basic model, with no substitutability or complementarity effects (valuation for the bundle of items equals the sum of the valuations for individual items).

In terms of the distributions from which valuations are drawn, two classes of models are considered in literature. In the *symmetric* case, $f_i = f_j$ for all bidders i and j . In the *asymmetric* case, valuations of different bidders are drawn from different f_i . We will consider both cases.

2.2 Mechanism design principles

Each bidder's valuation function is private information (although the auctioneer and other bidders may know the distribution from which it is drawn). Thus a concern is that a bidder might not bid her true valuations for the bundles—she might be able to obtain higher utility by submitting a different valuation function. As is common in much of mechanism design, especially within computer science, we focus on on *deterministic dominant-strategy mechanisms*, that is,

mechanisms where each bidder has a strategy that is optimal regardless of what the other bidders do. Such mechanisms are robust in the sense that the bidders do not benefit from counterspeculating each others' valuations and rationality. As is standard in mechanism design, we focus on *incentive-compatible* mechanisms, that is, mechanisms where each bidder's dominant strategy is to bid truthfully. This is without loss of generality in the sense that 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 (see e.g. (Krishna 2002)).

It is also important that the mechanism motivates the bidders to participate. A mechanism is *ex post individually rational* if each bidder is no worse off participating than not participating, for all possible valuations of other bidders. As does most of auction theory, we focus on such mechanisms.

2.3 Vickrey-Clarke-Groves (VCG) mechanism

A classic example that satisfies the above conditions is the following mechanism (Vickrey 1961; Clarke 1971; Groves 1973).

Definition 2.2 Vickrey-Clarke-Groves (VCG) mechanism. Each bidder i submits a valuation function v_i . The allocation, a , is computed to maximize social welfare

$$SW(v) = \sum_{i=0}^n v_i(a). \quad (2.1)$$

The payment by bidder i is

$$t_i = \left[\sum_{j \neq i} v_j(a_{-i}) - \sum_{j \neq i} v_j(a) \right]$$

where

$$a_{-i} = \operatorname{argmax}_a \sum_{j=0, j \neq i}^n v_j(a)$$

is the allocation that is optimal among the allocations where bidder i does not receive any items. One can also interpret t_i as the minimum valuation for a_i (the bundle won by i), which i would have had to bid in order to win a_i .

The VCG is the most popular mechanism in the literature. (In the 1-item case it is also equivalent to the widely used English auction). It maximizes the welfare of the bidders. However, it often yields poor revenue for the seller.

3 Revenue maximization in combinatorial auctions (CAs)

In this section we first explain the non-optimal revenue of the VCG, and then describe the general idea of revenue boosting.

3.1 Non-optimality of VCG

There are at least two reasons why the VCG mechanism may not yield maximal expected revenue.

1. **Bundling effect.** The following well-known simple example shows that bundling decisions of the seller may affect the revenue.

Example 3.1 Consider an auction with k items for sale (g_1, \dots, g_k) , and two bidders in the AVM model. The VCG would sell each item separately to the higher bidder, collecting payment equal to the valuation of the lower bidder for each item. So, the revenue is $\sum_{j=1}^k \min_{i \in \{1,2\}} v_i(g_j)$. However, should the seller decide to bundle all the items together and sell them as a whole via a second-price mechanism, she would receive revenue $\min_{i \in \{1,2\}} [\sum_{j=1}^k v_i(g_j)]$ which is clearly greater.

2. **Asymmetry of valuation distributions.** In the asymmetric case, it may happen that the distribution of valuations of bidder i for some bundle b stochastically dominates the distributions of other bidders (we call such bidders "strong" and "weak", respectively). For instance, consider a 1-item auction, where valuation of bidder i for the item is drawn uniformly from the $[2, 3]$ interval, while other bidders' valuations are drawn from $[0, 1]$ interval. Under the VCG allocation scheme bidder i wins the auction, paying just the second-highest bid price, while it would be easy to improve revenue beyond that by charging (at least) 2, which is the lowest possible valuation of bidder i .

3.2 Priors in mechanism design

Before discussing the methods for boosting revenue, we need to make a point about the use of priors, f_i , (on bidders' valuations) in mechanism design. The VCG is prior-free, since it does not use information about f_i 's in the allocation rule or payment rule. Although being prior-free is desirable (because the auctioneer can run the mechanism with less information), no prior-free mechanism can guarantee even a fraction of the optimal revenue:

Proposition 3.1 For every prior-free, incentive-compatible, individually-rational deterministic CA mechanism M , and every $\epsilon > 0$, there exist distributions of valuation functions V , such that
$$\frac{E_v(R_M(v))}{E_v(OPT(v))} < \epsilon$$

Here $OPT(v)$ denotes the revenue-optimal mechanism and E_v denotes the expectation over V .

Although in some special cases prior-free mechanisms may yield good revenue,² Proposition 3.1 shows that in order to construct a good general-purpose mechanism we need to use priors. Therefore, we focus on prior-dependent mechanisms (where the designer of the mechanism knows the distributions f_i).

3.3 Ideas from the Myerson auction

For boosting revenue in CAs, we will draw some ideas from the optimal 1-item auction (Myerson 1981):

Definition 3.1 Myerson 1-item auction. Each bidder i submits her valuation v_i for the item. The mechanism computes virtual valuations for the bidders:

²For example, in the symmetric case, the VCG revenue is asymptotically optimal as the number of bidders approaches infinity (Monderer & Tennenholtz 1999).

$$\tilde{v}_i(v_i) = v_i - \frac{1 - F_i(v_i)}{f_i(v_i)} \quad (3.1)$$

The allocation is computed so as to maximize the following objective:

$$SW(v) = \sum_{i=0}^n \tilde{v}_i. \quad (3.2)$$

Thus, the item is given to the bidder with highest virtual valuation.³ So, the allocation rule is the same as in the VCG (2.2), except that virtual valuations are used in place of real valuations. The payment by the winning bidder is equal to the minimal bid that she would have had to make in order to win (that is, $\tilde{v}_i^{-1}(v_j)$, where v_j is the second highest bid). The item is sold only if the virtual valuation of the winning bidder is above 0. All losing bidders pay nothing.

The intuition behind the mechanism is that it is biased in favor of weak bidders, creating an artificial competition between weak and strong bidders, and extracting more revenue from strong bidders. It is easy to check that the transformation (3.1) brings down the valuations of strong bidders more than those of weak bidders. Such a mechanism allows the auctioneer to set a high sell price for a strong bidder while motivating her to stay truthful (even if she is sure that her valuation exceeds valuation of any other bidder).

Drawing from this intuition, we propose increasing revenue in CAs by designing certain virtual valuations, and then running VCG on those valuations. We argue that virtual valuations are capable of improving VCG in both asymmetry handling and bundling aspects. In the next section we discuss the forms of virtual valuations that we use.

4 Techniques for boosting revenue in CAs

In this section we introduce two families of CAs that employ virtual valuations, and a hybrid family that combines the two. We then analyze the restrictions that individual rationality and incentive compatibility impose on virtual valuations.

4.1 Bidder weighting technique

Definition 4.1 Weighted auction. An allocation, a , is chosen that maximizes

$$SW^\mu(v) = \sum_{i=0}^n \mu_i v_i(a) \quad (4.1)$$

The parameters, μ , of the mechanism are positive real numbers, chosen in advance by the auctioneer, based on the priors on bidders' valuations. The payments by the bidders are

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

The mechanism effectively replaces the valuation function v_i of bidder i with $\mu_i v_i$. This is useful in asymmetric cases when valuations of some bidders are concentrated around higher values than those of other bidders. The proof

³For the mechanism to be incentive compatible, \tilde{v}_i should be increasing in v_i . If (3.1) does not satisfy this condition, an "ironing" technique is used to make \tilde{v}_i non-decreasing (Myerson 1981).

of incentive compatibility of this mechanism follows that of the VCG.

In many cases the same bidder can be strong w.r.t. some bundle and weak w.r.t. another bundle. It would thus seem to be helpful to allow the mechanism to assign a bidder different weights for different bundles. However it is easy to show that such a mechanism is not incentive compatible:

Proposition 4.1 *No mechanism that chooses an allocation, a , that maximizes $SW^{\mu(a)}(v) = \sum_{i=0}^n \mu_i(a_i)v_i(a)$ is incentive compatible for all possible valuations in the CA domain (unless $\mu_i(a_i)$ is constant over a_i).*

Proofs are omitted due to limited space.

4.2 Allocation boosting technique

Proposition 4.1 shows that there does not exist a general bundle- and bidder-specific multiplicative weighting mechanism. However, it turns out that it is possible to give a bidder bundle-specific advantage in an incentive-compatible way by using *additive* terms. Let $\lambda_{\{i,b\}}(a) = c_{\{i,b\}}$ for all allocations a that give bidder i exactly bundle b , and $\lambda_{\{i,b\}}(a) = 0$ otherwise (as before, λ 's are chosen *in advance* by the auctioneer, based on the priors on valuations of the bidders). Here, the $c_{\{i,b\}}$ values are real numbers that the auction designer sets. We call this the *allocation boosting technique*.

4.3 Bidder weighting and allocation boosting

Now, a mechanism that uses both *allocation boosting* and *bidder weighting* can be defined as follows:

Definition 4.2 Virtual valuations CA (VVCA). *The mechanism computes an allocation a that maximizes*

$$SW_{\lambda}^{\mu}(v) = \sum_{i=0}^n [\mu_i v_i(a) + \lambda_{\{i,a(i)\}}(a)] \quad (4.2)$$

where μ are positive. The payment rule is

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

VVCA is a family of mechanisms, parameterized by μ and λ . It also includes the VCG. By analogy to Myerson's auction, $\mu_i v_i(a) + \lambda_{\{i,a(i)\}}$ can be viewed as a virtual valuation $[\tilde{v}_i]_{\lambda}^{\mu}$ of bidder i for allocation a .

The mechanism adds $c_{\{i,b\}}$ to the value of the objective on allocations where bidder i gets bundle b .⁴ Obviously the

⁴The basic idea of adding an allocation-specific constant to the objective was introduced in (Jehiel, ter Vehn, & Moldovanu 2003) for the purpose of tuning the bundling policy. Their λ -auction can be viewed as a special case of the VVCA where 1) no weights are used ($\mu = 1$ for all i), and 2) the same fixed additive term is added to the objective whenever all items are sold to the same bidder (i.e., $\lambda_{\{i,b\}}(a) = c$ if all items are sold in one bundle b to bidder i , and $\lambda_{\{i,b\}}(a) = 0$ for all other bundles). In the symmetric AVM model, they show that the λ -auction can increase the revenue over both pure bundling auctions (which always sells all the items together) and separate auctioning of individual items (this is what the VCG does in the AVM model).

probability of bidder i winning b under the rule 4.2 is increasing in $c_{\{i,b\}}$. The proof of the truthfulness of the mechanism above also follows that of the VCG.

4.4 Impossibility of nonlinear virtual valuations

Incentive-compatibility imposes limitations on the virtual valuations that can be used in the mechanism. In 1-item auctions, it is sufficient for the virtual valuations \tilde{v}_i to be increasing in v_i (Myerson 1981). However, this is not sufficient in CAs. (Lavi, Mu'Alem, & Nisan 2003) recently showed that under certain natural assumptions, every (*ex post*) incentive-compatible CA is almost⁵ an affine maximizer.

Definition 4.3 Affine Maximizer Auction (AMA). *The allocation is computed so as to maximize*

$$SW_{\lambda}^{\mu}(v) = \sum_{i=0}^n \mu_i v_i(a) + \lambda(a) \quad (4.3)$$

Here μ_i are positive and $\lambda(a)$ is arbitrary. The payments are

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

It is easy to see that VVCA mechanisms are a strict subset of AMAs. The results of (Lavi, Mu'Alem, & Nisan 2003) imply that every "reasonable" (*ex post*) incentive-compatible and individually-rational *general* mechanism is an AMA. (Non-AMAs might be truthful for some specific distributions of valuations, only AMAs are truthful for *all* CA settings.) Therefore, among mechanisms that are based on virtual valuations, only those which use linear virtual valuations (\tilde{v}_i linear in v_i) are incentive compatible. Our VVCA captures all linear virtual valuations, and is thus the most general class of incentive-compatible CA mechanism that use virtual valuations.

4.5 Bidder-specific reserve prices

Despite the simple form of VVCA, manipulating the parameters (μ, λ) is a powerful tool. For instance, it allows the auction designer to enforce or prevent any bidder from receiving a certain bundle. Another important property of the VVCA is that it allows for bidder-specific reserve prices. (Bidder-specific reserve prices are also used in Myerson's revenue-optimal 1-item auction (3.2). Recall that the item is sold only if the virtual valuation of the winning bidder i is above 0, which sets a reserve price for this bidder to $\tilde{v}_i^{-1}(0)$)). However, the reserve-price mechanism does not generalize straightforwardly to arbitrary CAs. The standard way to set reserve prices in CAs is to submit fake bids by the seller. That approach does not support *bidder-specific* reserve prices. On the other hand, bidder-specific reserve prices can be achieved in the VVCA: to ensure that bidder i never gets the bundle b for a price below p_0 , set $\lambda_{\{i,b\}} = -p_0$.

⁵A mechanism is an almost affine maximizer if it is an affine maximizer for sufficiently high valuations. (Lavi, Mu'Alem, & Nisan 2003) conjecture that the "almost" qualifier is merely technical, and can be removed in future research.

5 Searching for good parameters

VVCA and AMA define families of mechanisms, parameterized by (λ, μ) . Depending on the value of the parameters, the seller's expected revenue may be greater or less than in the VCG. In this section we discuss the problem of choosing parameters that yield high revenue.

The seller's expected revenue in VVCA is

$$R(\mu, \lambda, v) = \sum_{i=1}^n t_i(\mu, \lambda, v) = \quad (5.1)$$

$$-(n-1)SW_{\lambda}^{\mu}(v) - \sum_{i=1}^n \lambda_{\{i, a(i)\}} + \sum_{i=1}^n [SW_{-i}]_{\lambda}^{\mu}(v)$$

The following Proposition proves that $E_v[R(\mu, \lambda, v)]$ is a "well behaved" function of (λ, μ) and therefore suggests the use of numerical methods (as hill-climbing) for estimating (locally) optimal values of those parameters. That requires evaluating $E_v[R(\mu, \lambda, v)]$ for given (λ, μ) , which can be estimated by sampling valuations from the distributions f_i .

Proposition 5.1 *The expected revenue of the AMA (and consequently VVCA) is continuous and almost everywhere differentiable in μ and λ .*

Finding $R(\mu, \lambda, v)$ for a given set of valuations v requires determining the affine maximizing allocation (4.2). This winner determination problem is known to be NP-complete, but can be optimally solved for relatively large instances in practice. Any optimal winner determination algorithm for CAs can be used here; the affine maximization problem can be converted into the standard combinatorial auction winner determination problem by preprocessing the bids with the multiplicative and additive terms. In the AVM model, an important special case, the optimal allocation is trivial to find: every item is sold to the bidder with highest $[\tilde{v}_i]_{\lambda}^{\mu}$.

The main problem in optimization is that the number of parameters in (λ, μ) is exponential in the number of items for sale: μ is just a vector of size n , but the length of λ is $n * 2^k$ (for every bidder we have one parameter for every bundle) for VVCA and n^k for AMA. It would be helpful if we could discard some choices of $\lambda_{\{i,b\}}$ beforehand (e.g. by setting them to 0), thereby simplifying the optimization process. Unfortunately, the theorem below shows that there cannot exist a polynomial-time algorithm capable of always determining the optimal value for any $\lambda_{\{i,b\}}$, even if the valuations of the bidders are given. Moreover, no polynomial-time algorithm can always determine whether the mechanism with $\lambda_{\{i,b\}}$ set to some particular value λ_1 yields higher revenue than the mechanism with $\lambda_{\{i,b\}}$ set to λ_2 .

Theorem 5.1 *For any bundling parameter in VVCA (and AMA) - $\lambda_{\{i,b\}}$ - and any pair of values of this parameter - $(\lambda_1$ and $\lambda_2)$ there does not exist an algorithm that determines whether $R(\mu, (\lambda_{-\{i,b\}}, \lambda_1)) > R(\mu, (\lambda_{-\{i,b\}}, \lambda_2))$ in polynomial time, even if the valuations v of the bidders are given, unless $P=NP$. (Here $\lambda_{-\{i,b\}}$ denotes the set of all λ parameters except for $\lambda_{\{i,b\}}$.)*

Theorem 5.1 shows that there is no easy general method to decide whether one set of parameters is better than another. Therefore, there is no easy way to fix some of the

parameters up front without compromising optimality. Any search algorithm that guarantees the optimum for *any* distribution of valuations must run optimization in all $n * 2^k$ parameters (and therefore uses an exponential number of optimization runs). A related problem is that the surface of $E_v[R(\mu, \lambda, v)]$ is non-convex even in simple mechanisms, see Figure 6.1. The latter makes optimization complicated.

However, the problem can be addressed as follows. If the number of items and bidders is small, run optimization in all $n * 2^k$ parameters and find a close to optimal solution (see Section 6). For larger problems we can find the local optimum (for example, by hill climbing with the VCG as the starting point, thus yielding higher revenue than the VCG). Also, for some *special* distributions, mechanisms for finding the optimal values of parameters might exist.

6 Experiments

In this section we demonstrate how the suggested methodology can be used. The approach of this paper is a form of automated mechanism design (AMD) (Conitzer & Sandholm 2002; 2003). The main difference is that the mechanisms' space is restricted to mechanisms of the VVCA form. Although this may not yield an optimal mechanism (while traditional AMD does), this approach drastically simplifies the computation. It thus scales to larger problem instances. It also handles settings where the distributions of valuations are continuous (unlike in traditional AMD).

We now consider several example problems. We compare the revenue of VVCA to the revenue of the VCG and AMA. (The parameters of VVCA and AMA were determined using a hill-climbing procedure, starting at all points of a fine multidimensional grid; this does not guarantee optimality of the solution, but it is extremely likely that this method finds the optimum for the examples considered). The expected revenues were computed by sampling the valuations from the specified distributions (100000 samples were drawn, where each sample included one valuation function for each bidder) and running the auction in question on each sample. *Consider an auction setting with 2 items, g_1 , g_2 , and 2 bidders with valuation functions v_1 and v_2 , correspondingly. 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 .*

The results for various distributions F_1, F_2, C are given in the following table. The columns correspond to the experiments with various valuation models: **Ex. I** - AMA model, **Ex. II** - symmetric model with substitutabilities/complementarities, **Ex. III** - asymmetric model with substitutabilities/complementarities. The first three rows specify distributions F_1, F_2, C , the last three - the expected revenue on these distributions in VCG, optimal AMA found and optimal VVCA found. We give the revenue estimate and 95% confidence intervals for the estimated values (in tiny font). $U[a, b]$ denotes a uniform distribution on $[a, b]$.

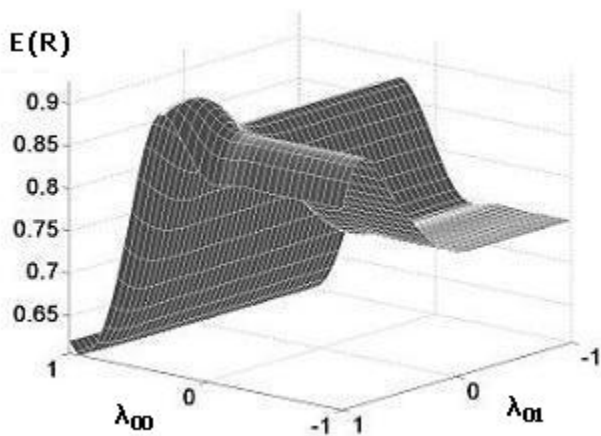


Figure 6.1: 3-dimensional projection in (μ, λ) space of the expected revenue surface of the AMA mechanism in Experiment I. All μ and λ are fixed, except for the following λ parameters: λ_{00} favors allocations where both items are kept by the seller, and λ_{10} favors allocations where item 1 is allocated to bidder 1 and the other item is kept. The analogous parameters λ_{01} (referring to bidder 1 and item 2), λ_{20} (bidder 2 and item 1) and λ_{02} (bidder 2 and item 2) are set equal to λ_{10} .

	Ex. I	Ex. II	Ex. 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 _(2.72, 2.78)	2.85 _(2.82, 2.88)
AMA	0.88 _(0.86, 0.9)	2.78 _(2.75, 2.81)	4.21 _(4.16, 4.26)
VVCA	0.87 _(0.85, 0.89)	2.78 _(2.75, 2.81)	4.20 _(4.15, 4.25)

The experiments show that both the general AMA mechanism and our VVCA mechanisms yield substantial improvement over the VCG even in symmetric auctions. Also, the VVCA is as good as AMA (within the confidence)—at least on these small problem instances.

7 Conclusions and future research

In this paper we developed a new approach to the problem of maximizing revenue in combinatorial auctions (CAs). Instead of attempting a full characterization (a recognized elusive problem), we developed methods for automatically modifying existing auction mechanisms to increase revenue.

We introduced two techniques: bidder weighting and allocation boosting and a general family of auctions that uses those techniques. We call such auctions *virtual valuations combinatorial auctions* (VVCA). All auctions in the family are based on the Vickrey-Clarke-Groves mechanism, executed on virtual valuations that are linear transformations of the bidders' valuations. The restriction to linear transformations is motivated by incentive compatibility (truthfulness). The auction family is parameterized by the multipliers and constants in the linear transformations. VVCAs are a subset of an even more general parametric family of *affine maximizer auctions* (AMAs).

For VVCA and AMA, the problem of mechanism design is therefore reduced to search in the parameter space (AMA has many more parameters). We proved that the revenue of

both VVCA and AMA is a well-behaved function, suggesting the use of hill-climbing methods for parameter search. However, finding the optimal parameters turned out to be a computationally hard problem for both families. Despite that, close to optimal parameters for VVCA and AMA can be found at least in settings with few items and bidders (the experiments on small auctions showed that VVCA yields a drastic increase in revenue over the VCG, and the same revenue (within tolerance) as AMA). In larger auctions, locally optimal parameters can be used. This still yields higher revenue than the VCG because the hill-climbing in parameter space can be started from the VCG auction.

We plan to pursue several extensions to this approach. The most important problem is to show that the proposed mechanism yields a good approximation of revenue compared to the (unknown) revenue-optimal mechanism (if the parameters are set optimally). Another line of research is concerned with the organization of the search of the parameter space: although we proved that any search algorithm which guarantees the optimum for *any* given distribution of valuations must run optimization in an exponential number of parameters, better algorithms might exist for specific distributions. Finally, the approach could be extended to combinatorial exchanges with multiple sellers.

8 Acknowledgements

This material 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.

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. 2002. Complexity of mechanism design. In *Proceedings of the 18th Annual Conference on Uncertainty in Artificial Intelligence (UAI-02)*, 103–110.
- Conitzer, V., and Sandholm, T. 2003. Applications of automated mechanism design. In *UAI-03 workshop on Bayesian Modeling Applications*.
- Groves, T. 1973. Incentives in teams. *Econometrica* 41:617–631.
- Jehiel, P.; ter Vehn, M. M.; and Moldovanu, B. 2003. Mixed bundling auctions. Working Paper.
- Krishna, V. 2002. *Auction Theory*. Academic Press.
- Lavi, R.; Mu'Allem, A.; and Nisan, N. 2003. Towards a characterization of truthful combinatorial auctions. In *Proceedings of the Annual Symposium on Foundations of Computer Science (FOCS)*.
- Monderer, D., and Tennenholtz, M. 1999. Asymptotically optimal multi-object auctions for risk-averse agents. Technical report, Faculty of Industrial Engineering and Management, Technion, Haifa, Israel.
- Myerson, R. 1981. Optimal auction design. *Mathematics of Operation Research* 6:58–73.
- Palfrey, T. 1983. Bundling decisions by a multiproduct monopolist with incomplete information. *Econometrica* 51:463–484.
- Vickrey, W. 1961. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance* 16:8–37.
- Vohra, R. V. 2001. Research problems in combinatorial auctions. Mimeo, version Oct. 29.