

An Approximate Truthful Mechanism for Combinatorial Auctions with Single Parameter Agents

Aaron Archer, Christos Papadimitriou, Kunal Talwar, and Éva Tardos

Abstract. Mechanism design seeks algorithms whose inputs are provided by selfish agents who would lie if it were to their advantage. Incentive-compatible mechanisms compel the agents to tell the truth by making it in their self-interest to do so. Often, as in combinatorial auctions, such mechanisms involve the solution of NP-hard problems. Unfortunately, approximation algorithms typically destroy incentive compatibility. Randomized rounding is a commonly used technique for designing approximation algorithms. We devise a version of randomized rounding that is incentive-compatible, giving a truthful mechanism for combinatorial auctions with single parameter agents (e.g., “single minded bidders”) that approximately maximizes the social value of the auction. We discuss two orthogonal notions of truthfulness for a randomized mechanism—truthfulness with high probability and in expectation—and give a mechanism that achieves both simultaneously.

We consider combinatorial auctions where multiple copies of many different items are on sale, and each bidder i desires a subset S_i . Given a set of bids, the problem of finding the allocation of items that maximizes total valuation is the well-known SETPACKING problem. This problem is NP-hard, but for the case of items with many identical copies, the optimum can be approximated very well. To turn this approximation algorithm into a truthful auction mechanism, we overcome two problems: We show how to make the allocation algorithm *monotone*, and give a method to compute the appropriate payments efficiently.

I. Introduction

Multiple-item auctions are a basic tool of electronic commerce. These auctions allow bidders to coordinate their purchases of complementary goods, such as airfare, lodging, and ground transportation for a vacation package, or mobile phone spectrum licenses in adjacent geographical regions. Many such large-scale auctions have indeed been used recently by the FCC and governmental bodies in Europe and elsewhere to allocate spectrum licenses to mobile phone providers. The FCC auctions alone granted thousands of licenses to hundreds of companies, raising more than \$40 billion [Cramton 02]. While the FCC decided to use an innovative simultaneous ascending auction for its various spectrum licenses, it also strongly considered various forms of combinatorial auctions, in which the participants bid on sets of items, rather than individual items. The sheer magnitude of these spectrum auctions and the rise of electronic commerce have both generated a surge of interest in designing good mechanisms for such combinatorial auctions.

We will consider combinatorial auction mechanisms (direct revelation auctions) where each bidder i bids a valuation b_i for a set S_i she is interested in. We will assume that each bidder i is bidding for a single set S_i , and this set is known to the auctioneer or can be inferred from context. Thus, each agent's only private information is her true valuation for that set. A standard desire in the design of combinatorial auctions is that they be *truthful* (or *incentive compatible*). The auction is truthful if each bidder's best strategy is always to reveal her true valuation, regardless of the other bidders' valuations, and regardless of how they decide to bid. That is, truthful bidding is a *dominant strategy* for each bidder.

It is known that an allocation algorithm leads to a truthful mechanism if and only if it is monotone (see [Mirrlees 71, Archer and Tardos 01, McAfee and McMillan 88] and [Fudenberg and Tirole 91, pages 257–262]). A randomized auction mechanism \mathcal{A} is said to be monotone if for every agent i , the probability that \mathcal{A} assigns the desired set S_i to agent i is increasing in her bid b_i . This characterization is very useful in designing computationally feasible truthful mechanisms for problems which are NP-hard—if we can come up with an approximation algorithm that is monotone, there exists an accompanying payment scheme that gives a truthful mechanism. Further, the whole mechanism is computationally efficient if the payments can also be computed efficiently.

In this paper, we develop a technique that makes randomized rounding-based approximation algorithms useful in designing truthful mechanisms. Randomized rounding of an LP solution is a commonly used technique for designing polynomial time approximation algorithms. Typically, such rounding algorithms

succeed with high probability. However, it is not clear what the associated mechanism should do to ensure incentive compatibility when the rounding fails to produce a feasible solution. We show a technique for obtaining a monotone allocation algorithm from such a rounding scheme, and also show how to compute payments in polynomial time.

There are two natural goals for designing good auctions: maximizing the revenue, and maximizing the total valuation, which is the sum of the valuations of the bidders who receive their desired sets. We will concentrate on the latter objective, which is referred to as *efficiency* in the economics terminology. *In some cases, maximizing efficiency is a more important objective than generating revenue.* For instance, one of the primary goals in the spectrum auctions was to get spectrum licenses into the hands of the companies that could best use them to build up a viable mobile phone network; it is widely believed that high valuation is a strong indicator of how well positioned the company is to make good use of the spectrum license [Cramton 02]. The well-known Vickrey-Clarke-Groves (VCG) mechanism [Vickrey 61, Clarke 71, Groves 73] is truthful and maximizes the total valuation. However, finding the VCG allocation often requires solving an NP-hard optimization problem (e.g., in the case of single-minded bidders, the optimization problem is the well-known SETPACKING problem), and simply replacing the exact optimization routine required in the VCG mechanism with an approximation algorithm causes the mechanism to lose its incentive compatibility properties [Nisan and Ronen 00].

Over the last 15 to 20 years, there has been a large amount of work on approximation algorithms for a huge array of hard optimization problems. However, so far, there are only very few examples when approximation algorithms have turned out to be useful for designing polynomial time truthful mechanisms. One of the first such examples was due to Lehmann, O’Callaghan, and Shoham [Lehmann et al. 02] for the case of *single-minded bidders* (i.e., each agent bids for a single set), who give a mechanism based on a greedy allocation. Their mechanism is truthful and attains a \sqrt{m} -approximation to the optimal allocation, where m is the number of items. Moreover, they prove their result is nearly the best possible, in the following sense. For every constant $c < \frac{1}{2}$, there is no polynomial time algorithm that achieves an approximation guarantee of m^c , unless $NP = ZPP$. Mu’alem and Nisan [Mu’alem and Nisan 02] consider the case of *known single-minded bidders*, where the sets are known, and each agent’s only private data is her valuation. They show how to combine certain truthful mechanisms into an improved mechanism, while preserving truthfulness. Using this technique, they improve the greedy mechanism of Lehmann et al. by adding a partial enumeration of the space of allocations. The resulting polynomial time mechanism yields an $\epsilon\sqrt{m}$ approximation, for any constant $\epsilon > 0$.

We consider the case of known single-minded bidders when there are $\Omega(\ln K)$ copies of each item available, where K is the maximum size among any of the sets S_i . For the corresponding optimization problem of finding an allocation maximizing total valuation, there is a good approximation algorithm that uses randomized rounding. Our randomized auction mechanism is based on this algorithm. To turn the approximation algorithm into a truthful auction mechanism, we overcome two difficulties: We show how to make this allocation monotone, and give a method to compute the appropriate payments efficiently. Our auction mechanism runs in polynomial time, is truthful, and attains a $(1 + \epsilon)$ approximation to the optimal valuation. It essentially implements the fractional version of the VCG mechanism, both in terms of expected allocation and expected revenue.

2. Basic Definitions

A combinatorial auction is designed to divide up a set \mathcal{G} of items among a set \mathcal{N} of n bidders. Each bidder i has a valuation function $v_i : 2^{\mathcal{G}} \rightarrow \mathbb{R}^+$ that describes her preferences over the various subsets of items. For $S \subseteq \mathcal{G}$, $v_i(S)$ represents the maximum amount of money bidder i is willing to pay for the set of items S . The function v_i is known only to player i . A *single-minded bidder* i is one who values only a particular set of items. More formally, there is a set S_i and a $c \geq 0$ such that $v_i(T) = c$ if $S_i \subseteq T$ and $v_i(T) = 0$ otherwise. We consider the case where all bidders are single-minded, and the auctioneer knows the sets S_i ahead of time. This is the case of *known single-minded bidders* considered by [Mu'alem and Nisan 02]. There can be multiple copies of each good, in which case the *multiplicity* m_j denotes the number of copies of item j that are available. For simplicity, we assume that each bidder desires only one copy of each good in her desired set; that is, S_i is a set, not a multiset. This assumption is not crucial to our auction mechanism. We make it only to simplify the statements of our performance guarantees. In particular, many of our results guarantee some small failure probability, which is valid so long as the multiplicities m_j are sufficiently high. If the S_i are allowed to be multisets, then the expressions for the multiplicities become more complicated.

We consider *direct revelation* auction mechanisms. Each player i submits a *bid* b_i to the mechanism. Player i 's bid is supposed to represent the maximum amount v_i that she is willing to pay for her desired set, but she may choose to lie. We assume that v_i is some given constant that does not depend on the outcome of the auction or on the other players' bids (i.e., a *private values* model). Based on the bids, the mechanism decides which players win and at what price. Denote the vector of all n bids by b . Formally, a mechanism \mathcal{M} is a collection of 0-1

functions $x_i(b)$ and real functions $P_i(b)$, where $x_i(b)$ is 1 if i wins her desired set and 0 otherwise, and $P_i(b)$ is the price i must pay. The functions x_i must be such that each item is sold to at most as many players as there are items available, i.e., $\sum_{i:j \in S_i} x_i \leq m_j$ for each item j . We define

$$\text{profit}_i(b) = v_i x_i(b) - P_i(b), \quad (2.1)$$

that is, i 's valuation for the goods she gets, minus the price she pays. We assume that each player's goal is to maximize her own profit. The allocation functions x_i and the price functions P_i are all publicly known. The only pieces of private information are the valuations—only player i knows the true value of v_i . We require that our mechanisms satisfy the *voluntary participation* condition, which says that a player is charged zero if she loses, and her (expected) payment is at most b_i if she wins. This guarantees that players who bid truthfully always obtain non-negative (expected) profit.

In the game theory literature, preference orderings over outcomes are often given by *utility functions*, and the term *utility* typically refers to von Neumann-Morgenstern utilities. This means that a player's utility for a probability distribution over several outcomes is the expected utility of the random outcome. One usually refers to an expression of the form (2.1) as a quasilinear utility function. We instead use the term *profit* to emphasize that one should not necessarily interpret it as a von Neumann-Morgenstern utility. The assumption that profits take the form (2.1) simply means that all the valuations can be expressed in a common unit of money. Whether or not the agents are risk-neutral (i.e., wish to maximize expected profit) is a separate assumption that we may or may not wish to make.

Given a publicly specified auction mechanism, how should a player bid to maximize her own profit? Let b_{-i} denote the vector of bids by all bidders besides i , so we can write b as (b_{-i}, b_i) . We say that *truthtelling* is a (*weakly*) *dominant strategy* for bidder i if, no matter what the other agents do, bidding her true valuation v_i will maximize her profit. That is, $v_i \in \operatorname{argmax}_{b_i} \text{profit}(b_{-i}, b_i)$, for all b_{-i} . In other words, even if player i knew the bids of the other agents ahead of time, still the best she could do is to tell the truth. If truthtelling is a dominant strategy for each agent, then we say the mechanism is *truthful* (or *incentive compatible*). For a deterministic mechanism to be truthful, it is necessary that bidder i 's price $P_i(b_{-i}, b_i)$ depends on her own bid b_i only to the extent that it determines whether she wins or loses.

Auction designers care about truthfulness for two main interrelated reasons. First, it makes life easy for the bidders. In order to determine an optimal bidding strategy, each bidder only has to figure out her own valuation. She does not have to make any assumptions about the other players' valuations, or what

bidding strategies they will use. In particular, she does not have to perform any difficult Nash equilibrium calculations, nor does she have to assume that the other agents are performing those same calculations to determine their own bids. The second reason is that, because truthful bidding is a dominant strategy, players are likely to follow it, so bidder behavior becomes much more predictable than in an auction without dominant strategies.

Sometimes it is useful for the mechanism to use randomization. A *randomized mechanism* can be viewed as a randomization over a collection of deterministic mechanisms. That is, a randomized mechanism flips some coins to select a random element ω from some probability space, then uses a deterministic mechanism \mathcal{M}_ω based on the coin flips. All details of the mechanism are public knowledge, except for the actual outcomes of the coin flips. There are several notions of truthfulness for randomized mechanisms. The strongest notion is for the mechanism to be *universally truthful*. This means that for *every* ω , the mechanism \mathcal{M}_ω is truthful. This concept has been used in [Nisan and Ronen 01, Bar-Yossef et al. 02, Goldberg et al. 01, Fiat et al. 02], but it is very restrictive. For example, [Archer and Tardos 01] considers a machine scheduling problem for which no universally truthful constant-factor approximation mechanism is known, but they attain a 3-approximation via a mechanism that is truthful in expectation.

Because strong truthfulness is so restrictive, there have been various attempts to find a weaker, but still useful, concept. One approach is to guarantee that truthful bidding always maximizes a player's expected profit [Archer and Tardos 01], i.e., the mechanism is *truthful in expectation*. Two orthogonal notions are that a player may benefit from lying, but not by much [Schummer 01], or only with a small probability. We pursue the first and third approaches.

We say a randomized mechanism is *truthful with error probability* ϵ if for each b_{-i} and each v_i we have $Pr[v_i \notin \operatorname{argmax}_{b_i} \operatorname{profit}_i(b_{-i}, b_i)] \leq \epsilon$. If ϵ is inverse polynomial in some specified parameters of the auction (such as the number of items or bidders), then we say the mechanism is *truthful with high probability*. Even in the rare event that a bad ω is chosen by the mechanism, computing an effective lie could be difficult and would require knowledge about the other bids. Moreover, such a lie may backfire in the probability $(1 - \epsilon)$ event that the mechanism selects a good ω . In using such a mechanism, one hopes that these factors combined will convince the agents not to bother lying. This notion may be preferable to that of truthfulness in expectation because it does not assume players are risk-neutral. In this paper, we design an auction mechanism that is simultaneously truthful in expectation and truthful with high probability.

3. Our Mechanism for Known Single-Minded Bidders

We design a randomized mechanism based on solving the natural linear programming relaxation of the SETPACKING problem, and randomly rounding the resulting fractional allocation. In the case that the number of copies of each item is $\Omega(\ln K)$ (where K is the maximum size of a set S_i), we prove that our mechanism achieves near-optimal total valuation, is truthful in expectation and strongly truthful with high probability, and has revenue that compares well with a natural variant of VCG. We describe our mechanism by successive refinement of a simple randomized rounding idea, sprinkling in motivation as we encounter and overcome various obstacles.

First, recall that a deterministic mechanism for known single-minded bidders is truthful if and only if the allocation rule is monotone and the price for a winning player equals her “threshold.” That is, if we fix the other bids b_{-i} , then player i has some threshold bid $T_i(b_{-i})$ such that she wins and pays $T_i(b_{-i})$ if $b_i > T_i(b_{-i})$, and loses if $b_i < T_i(b_{-i})$. (If $b_i = T_i(b_{-i})$, the player can win and pay $T_i(b_{-i})$, or lose and pay zero; it doesn’t matter which.) This characterization has been noted many places, such as [Goldberg et al. 01, Lehmann et al. 02, Archer and Tardos 01, Archer and Tardos 02, Ronen 00, Mu’alem and Nisan 02]. Analogously, a randomized mechanism is truthful in expectation if and only if for every agent i , the probability $p_i(b_{-i}, b_i)$ that the mechanism assigns her the desired set S_i is increasing in her bid b_i , and her expected payment is equal to a certain integral of the function p_i (see [Mirrlees 71, Archer and Tardos 01] or [Fudenberg and Tirole 91, pages 257–262]).

Our mechanism works as follows. First, collect the bids. Using some small fixed $\epsilon' \in (0, 1)$ that is publicly known, pretend that we have only $m'_j = \lfloor (1 - \epsilon')m_j \rfloor$ copies of each item j to distribute. Now, solve the following linear program to get an optimal fractional allocation, using the artificially reduced supply of goods.

$$\begin{aligned} & \text{maximize} && \sum_{i \in \mathcal{N}} b_i x_i && (3.1) \\ & \text{subject to:} && \sum_{i: j \in S_i} x_i \leq m'_j && \text{for all } j \in \mathcal{G} \\ & && 0 \leq x_i \leq 1 && \text{for all } i \in \mathcal{N} \end{aligned}$$

Denote the optimal fractional allocation by x . We assume that we always find a vertex solution to the linear program, and break ties in a particular fixed way independent of the bids b (e.g., between two vertex solutions, choose the solution with the higher value of x_i for the smallest index i in which they differ). Notice that a fractional value of x_i means that the LP allocates player i an x_i fraction of *each* good in her set S_i . Now we perform the standard trick of treating the

x_i as probabilities. We define a preliminary set of *initial winners* by selecting each bidder i independently with probability x_i . However, we may have tried to sell too many copies of some items, so we will need to modify this outcome by deleting certain selected bidders. The modified outcome will be feasible, yet we keep the auction monotone in the bids.

First, it is not hard to see that, with high probability, no item is oversold.

Chernoff Bound. [Motwani and Raghavan 95] *Let X_1, \dots, X_n be independent Poisson trials such that, for $1 \leq i \leq n$, $\Pr[X_i = 1] = p_i$. Then for $X = X_1 + \dots + X_n$, $\mu \geq p_1 + \dots + p_n$, and any $\alpha < 2e - 1$, we have*

$$\Pr[X > (1 + \alpha)\mu] < e^{-\mu\alpha^2/4}.$$

Proposition 3.1. *Let $c > 0$. Suppose that each item $j \in \mathcal{G}$ has multiplicity $m_j = \Omega(\ln K)$. Then the probability that a given item is oversold is at most $\frac{1}{K^{c+1}}$ (the multiplicative constant inside the Ω is $\frac{4(c+1)}{e^{c^2(1-e^c)}}$).*

It is easy to show that this randomized initial allocation is monotone, i.e., that the value x_i in the optimum is monotone in the bid b_i of agent i .

Lemma 3.2. *Let x be an optimal solution to a linear program when the objective function vector is b , and let x' be an optimal solution when the objective function vector is b' (where ties are broken independently of b and b'). Suppose $b_i = b'_i$ for all $i \neq i_0$ and $b'_{i_0} > b_{i_0}$. Then either $x' = x$ or $x'_{i_0} > x_{i_0}$.*

Proof. Since x is the optimal solution to the linear program, and x' is a feasible solution,

$$b \cdot x \geq b \cdot x'. \quad (3.2)$$

Similarly,

$$b' \cdot x' \geq b' \cdot x. \quad (3.3)$$

If $x = x'$, we are already done. Otherwise, the two vertices are distinct, and since the tie breaking rule is the same, one of the two inequalities must be strict. Adding (3.2) and (3.3), and noting that $b' \cdot x = b \cdot x + (b'_{i_0} - b_{i_0})x_{i_0}$ for all vectors x , we get

$$(b'_{i_0} - b_{i_0})(x'_{i_0} - x_{i_0}) > 0$$

from which the result follows. \square

From the above result, it follows that the probability of rounding any agent to 1 is monotone in her bid. Thus, if no items were ever initially oversold, we would get a truthful mechanism. However, there is some probability that an agent i is not allocated her set even though she is initially rounded to 1, because some item j in S_i may be oversold, in which case some or all of the agents desiring j must be dropped. For reasons of egalitarianism (and for lack of a better idea), we propose to drop every agent i whose set S_i contains an oversold item j . As the following example shows, the overall probability that agent i wins is not necessarily monotone in her bid.

Example 3.3. Consider an instance where there are 51 agents bidding for an item “A” with multiplicity 1. Suppose that when the first agent bids a value b_1 , the fractional allocation is $\frac{1}{2}$ for agent 1, and $\frac{1}{100}$ for agents 2, \dots 51. In this case, the probability that agent 1 is an initial winner is $\frac{1}{2}$, and the probability that she is finally allocated the item is $\frac{1}{2}(1 - \frac{1}{100})^{50} \approx 0.3$.

Now suppose she raises her bid to $b_2 > b_1$ and the fractional solution changes to 0.51 for agent 1, 0.49 for agent 2, and 0 for all the other agents. In this case, agent 1’s probability of being an initial winner increases to 0.51. However, her probability of being allocated the item is now $(0.51)(1 - 0.49) \approx 0.26$. Hence, the algorithm described above is not monotone.

As a result, we do not yet have a truthful mechanism. In the next two sections, we show how to achieve monotonicity. We give a very high level description in Section 3.1. and then fill in the necessary details in Section 3.2.

3.1. Dealing with Oversold Items: The Basic idea

While the probability that the rounding fails is small, it may depend on the agent’s bid and may not be monotone. Our approach is to drop each agent with some additional probability so as to ensure that the overall probability that agent i wins is directly proportional to the variable x_i . When the rounded solution is not feasible, it may still be possible to serve some agents without clashes. We use the following approach:

Step 1. *Solve* the scaled linear programming relaxation (3.1).

Step 2. *Round* each variable x_i to 1 with probability x_i , set to 0 otherwise.

Step 3. *Select* all agents i that are rounded to 1 and such that the supply constraints for all items in S_i are satisfied.

Step 4. *Drop* each agent with some additional probability (to be defined later).

Let \hat{x} denote the integer assignment resulting in Step 2. Consider an agent i_0 . The agent is selected in Step 3 if she is rounded to 1 in Step 2 and the constraints for all items in S_{i_0} are satisfied. That is, i_0 is selected if $\hat{x}_{i_0} = 1$ and \hat{x} satisfies

$$\sum_{i: j \in S_i, i \neq i_0} \hat{x}_i \leq m_j - 1 \quad \text{for all } j \in S_{i_0}.$$

Let $I_{i_0} = \{i : S_i \cap S_{i_0} \neq \emptyset\}$. The variables $x_i : i \in I_{i_0}$ form a feasible solution to the scaled linear program (3.1) induced on the items in S_{i_0} . Let q_{i_0} be the conditional probability that no item in S_{i_0} is oversold, given that $\hat{x}_{i_0} = 1$. Set $q^* = 1 - \frac{2}{K^c}$. Using Proposition 3.1 and the union bound on the items in S_{i_0} , we get that $q_{i_0} > 1 - \frac{1}{K^c} > q^*$. Thus the probability that agent i_0 is selected at Step 3 is $x_{i_0} q_{i_0} > x_{i_0} q^*$. Therefore, in Step 4 we would like to drop agent i_0 with probability $1 - (q^*/q_{i_0})$, so that the probability that agent i_0 survives through the end of Step 4 is exactly $x_{i_0} q^*$. We would then have a monotone allocation algorithm.

3.2. Dealing with Oversold Items: Important Details

Note that the algorithm described above requires us to exactly compute the probability q_{i_0} . However, it is NP-hard to compute this number exactly, so the above scheme cannot be implemented efficiently. We get around this problem by using an estimator for this probability. First, we need the following simple observation:

Lemma 3.4. *Let x be any vertex of the polytope $\{x : Ax \leq r, 0 \leq x \leq 1\}$, where $A \in \{0, 1\}^{m \times n}$ and $r \in \mathbb{Z}^m$. Then $x \in \mathbb{Q}^n$ and each x_i can be written with denominator D , for some $D \leq m!$.*

Proof. Any vertex of the polytope is given by a linear system $\tilde{A}\tilde{x} = \tilde{r}$ where \tilde{A} is a nonsingular square submatrix of A , and \tilde{x} and \tilde{r} are corresponding submatrices of x and r , respectively. Then $\tilde{x} = \tilde{A}^{-1}\tilde{r}$. Let $D = \det(\tilde{A})$. Since \tilde{A} is a 0-1 matrix, $D \leq m!$. Moreover, $D \cdot \tilde{A}^{-1}$ and \tilde{r} are integer matrices and hence $D\tilde{x}$ is an integer vector. Any x_i not in \tilde{x} is set to zero or one, and hence trivially satisfies the conclusion. Hence, the claim. \square

Corollary 3.5. *Let x', x'' be vertices of the polytope $\{x : Ax \leq r, 0 \leq x \leq 1\}$, where $A \in \{0, 1\}^{m \times n}$ and $r \in \mathbb{Z}^m$. Then for each i , either $x'_i = x''_i$ or $x'_i \geq x''_i(1 + \delta)$ or $x''_i \geq x'_i(1 + \delta)$, where $\delta = (1/m!)^2$.*

Corollary 3.5, along with Lemma 3.2, implies that whenever an agent i increases her bid, this either has no effect on the allocation or it increases x_i by a factor of at least $(1 + \delta)$.

The algorithm described above requires computing $1/q_{i_0}$ for each agent i_0 . Instead of the exact value, we use an estimator Y for this number, and in Step 4 we retain agent i_0 with probability q^*Y . Consider the following experiment: Round x_{i_0} to 1. For each $i \in I_{i_0}$, round x_i to 1 independently with probability x_i . Recall that q_{i_0} is defined to be the probability that this solution satisfies the constraints for all items in S_{i_0} . Let the random variable X denote the number of trials of the experiment required before this happens, so that $E[X] = 1/q_{i_0}$. Our estimator Y for $1/q_{i_0}$ is $\min(\frac{1+\delta\epsilon}{N} \sum_{\ell=1}^N X^\ell, 1/q^*)$, where $N = O(K^c \log \frac{1}{\delta\epsilon})$, the X^ℓ are independent trials of the above experiment, c is from Proposition 3.1, and $\epsilon \ll 1$ is some error parameter of our choosing. (Later, ϵ will be part of the error probability in our guarantee on strong truthfulness. We will typically set $\epsilon \ll 1/K^c$.) The reason we take the min is that in Step 4, we retain agent i_0 with probability q^*Y , so we must ensure that this number is at most 1. The expectation of $\frac{1+\delta\epsilon}{N} \sum_{\ell=1}^N X^\ell$ is exactly $(1+\delta\epsilon)/q_{i_0}$, which is less than and bounded away from $1/q^*$. Since $\sum_{\ell=1}^N X^\ell$ is a negative binomial random variable with success probability $q_{i_0} \approx 1$, its distribution is concentrated about its mean, so the expectation of Y will not be much smaller than $(1+\delta\epsilon)/q_{i_0}$. In particular, we can use moment generating functions to bound the upper tail's contribution to the expectation (see, e.g., [Motwani and Raghavan 95]), and obtain the following.

Lemma 3.6. *The estimator Y defined above is at most $1/q^*$, and $E[Y] \in [1/q_{i_0}, (1+\delta\epsilon)/q_{i_0}]$.*

This finishes the description of our algorithm, which we will call **RANDROUND**. Note that the probability that agent i_0 is not dropped in Step 4 of the algorithm above is exactly $q^*E[Y]$. We now argue that this mechanism is monotone.

Theorem 3.7. *The probability that an agent i is selected by the algorithm **RANDROUND** is monotone increasing in her bid b_i .*

Proof. Fix an agent i and a vector of bids b_{-i} for agents other than i . Let $b'_i > b_i$. Consider the corresponding LP optima x and x' . From Lemma 3.2, either $x = x'$ or $x'_i > x_i$. In the first case, the experiment is the same and hence agent i 's probability of succeeding is the same whether she bids b_i or b'_i . In the second case, her probability $p_i(b_{-i}, b_i)$ of winning if she bids b_i is given by

$$\begin{aligned} p_i(b_{-i}, b_i) &= x_i q_i q^* E[Y] \\ &\leq x_i q_i q^* (1 + \delta\epsilon) / q_i \\ &= x_i q^* (1 + \delta\epsilon). \end{aligned}$$

If she bids b'_i , then her probability of winning is

$$\begin{aligned} p_i(b_{-i}, b'_i) &= x'_i q'_i q^* E[Y'] \\ &\geq x'_i q'_i q^* / q'_i \\ &\geq x_i q^* (1 + \delta), \end{aligned}$$

where the last inequality follows from Corollary 3.5. Since $\epsilon < 1$, this shows that the above LP rounding algorithm is monotone. \square

Theorem 3.8. *Suppose that each item has multiplicity $\Omega(\ln K)$, as in Proposition 3.1. Let OPT denote the optimal total valuation achievable by any allocation. Then the expected total valuation achieved by the above algorithm is at least $(1 - \epsilon')q^* OPT$.*

Proof. Every feasible allocation gives a feasible solution to the LP with the actual multiplicities m_j . Scaling down any solution to this LP by a factor of $(1 - \epsilon')$ yields a solution to our LP (3.1) with the artificially reduced multiplicities m'_j . Therefore, the optimal solution to this LP has value $\sum_{i \in \mathcal{N}} b_i x_i \geq (1 - \epsilon') OPT$. The probability that agent i is selected is at least $x_i q^*$, hence the expected total valuation is at least $(1 - \epsilon')q^*$ times the LP optimum. \square

4. Computing Payments

It is not obvious how one should compute payments for the winners. There is a payment scheme given in [Archer and Tardos 01] that yields truthfulness in expectation (provided the allocation rule is monotone), but in our case it involves an integral of a step function with possibly exponentially many breakpoints. One approach is to use an appropriate unbiased estimator for this integral, which we explain briefly below. In Section 4.1, we show another method, which attains truthfulness with high probability (but not in expectation), using a simpler (non-monotone) allocation rule. Finally, in Section 4.2, we show how to combine this payment scheme with our monotone allocation rule of Section 3.1 to simultaneously obtain truthfulness in expectation and with high probability.

Recall that $p_i(b_{-i}, b_i)$ denotes the overall probability that i wins her desired set S_i . Note that b_{-i} is fixed throughout this discussion, so we suppress it in the notation where convenient. A result in [Archer and Tardos 01] says that to guarantee truthfulness in expectation, i 's payment should be

$$R_i(b_{-i}, b_i) = p_i(b_{-i}, b_i) b_i - \int_0^{b_i} p_i(b_{-i}, u) du. \quad (4.1)$$

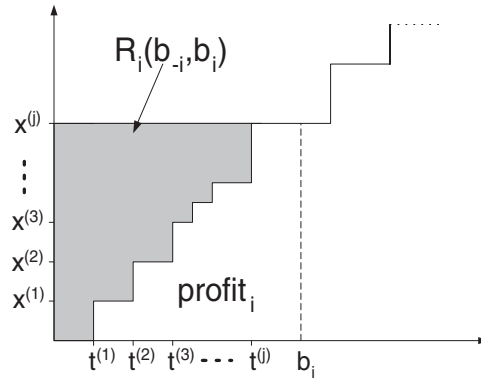


Figure 1. The graph shows i 's probability $p_i(b_{-i}, b_i)$ of winning as a function of her bid b_i . The gray area $R_i(b_{-i}, b_i)$ is her expected payment. If b_i is a truthful bid, then the white area is her expected profit.

Using the definition of incentive compatibility, the monotone property of the curve and a very simple calculation shows that this payment function induces truth-telling. This is a special case of a more complicated multidimensional result in [McAfee and McMillan 88]. A nice pictorial proof is given in [Archer and Tardos 01]. To gain some intuition for this, notice first that the payment formula yields the single-item Vickrey auction (highest bidder wins, and pays the second-highest bid) as a special case [Vickrey 61]. In the Vickrey auction, p_i is just a 0-1 step function. Now suppose p_i is a step function with multiple breakpoints. On the sections where p_i is flat, R_i does not change, so all bids along the flat section give the same result. For a truthful bidder, the first term in the payment function exactly cancels out the agent's expected valuation, so the second term is the profit. As an agent raises her bid above her true value, whenever a breakpoint of p_i is crossed, the payment increases by more than the agent's expected valuation, so overbidding is detrimental. The symmetric situation holds for underbidding.

In order to satisfy individual rationality, we must charge 0 to losing bidders. Thus, one truthful payment scheme is to charge $R_i(b_{-i}, b_i)/p_i(b_{-i}, b_i)$ to agent i if she wins. Note that $p_i(b_{-i}, u)$ is a step function which jumps whenever the selected vertex in the LP (3.1) changes and is flat elsewhere. See Figure 1. Thus, it could have exponentially many breakpoints, so it is not clear how to compute R_i efficiently. Instead, we can randomize the payment for winning bidder i by running the following experiment. Select some $u \in [0, b_i]$ uniformly at random, and run the allocation algorithm once, assuming i had bid u . If in the experiment i wins, set $Z = b_i$, else set $Z = 0$. Then Z is an unbiased estimator




Figure 2. The graph shows i 's fractional allocation as a function of her bid b_i (with b_{-i} fixed). It is a step function that is flat while one vertex of the LP (3.1) stays optimal, then jumps when another vertex becomes preferred. Since y_i lands in $(x^{(2)}, x^{(3)})$, i 's payment is $t^{(3)}$.

for the integral. Let X be an independent unbiased estimator for $1/p_i(b_{-i}, b_i)$. Then the random price $b_i - ZX$ for agent i has the correct expectation.¹

4.1. Threshold Payments

Consider the simpler allocation rule where we leave out Step 4 (the drop step). As previously noted, this allocation may not be monotone, in which case there is no payment scheme that is truthful in expectation. Here, we give payments that are instead truthful with high probability.

In the bidder selection step, let us perform the rounding by selecting n independent uniform $[0, 1]$ random variables y_1, \dots, y_n , and choosing i to be an initial winner if $y_i \leq x_i$. Note that each bidder i is selected by this experiment independently with probability x_i , as required by Step 2. For each winning bidder i , we compute a price that will depend on the outcome of the random variable y_i (and of course also on b_{-i}). Fix b_{-i} and a realization of the random cutoff y_i . There is a *threshold* value $T_i(b_{-i}, y_i)$ such that i will lose if $b_i < T_i(b_{-i}, y_i)$ and be an initial winner if $b_i > T_i(b_{-i}, y_i)$. See Figure 2. This threshold is the point at which $x_i(b_{-i}, b_i)$, considered as a function of b_i with b_{-i} fixed, first rises to y_i . We set player i 's price to be $T_i(b_{-i}, y_i)$ if she actually wins.

To compute $T(b_{-i}, y_i)$, we binary search on b_i , re-solving the LP each time. For the vector of bids $(b_{-i}, T_i(b_{-i}, y_i))$, there are two equally good fractional allocations x and x' . Therefore, $T_i(b_{-i}, y_i)(x_i - x'_i) = \sum_{j \neq i} b_j(x'_j - x_j)$. Assuming all

¹This payment has the peculiar property that it is often negative, i.e., the auctioneer must pay the buyer.

bids are given to d bits of precision, we can express $T_i(b_{-i}, y_i)$ as a fraction with denominator at most $2^d(m!)^4$ (by Lemma 3.4), so we can use binary search and the method of Diophantine approximation to compute it exactly in polynomial time (see [Schrijver 86]).

If our mechanism never had to throw away any initial winners, then our allocation algorithm would be universally truthful. Suppose we fix a particular realization ω of the vector of random variables y_1, \dots, y_n . Then the only circumstance under which agent i could benefit by lying is if i is selected as an initial winner, but is discarded because one of the items in S_i is oversold. Since this probability is at most $\frac{1}{K\epsilon}$, our mechanism is truthful with high probability.

This payment scheme also has the nice property that it satisfies the “no positive transfers” property—i.e., the bidders never get paid by the mechanism—and a stronger notion of individual rationality—if agent i wins, then she pays at most b_i for sure, not just in expectation.

4.2. Combining Threshold Payments with the Monotone Allocation

We now show how to modify this threshold scheme to get truthfulness in expectation, using the monotone allocation rule of Section 3.1. If it were the case that $p_i = q^* x_i$ for all i , i.e., each agent’s probability of winning was directly proportional to her fractional allocation from the LP, then the threshold payment scheme would give the correct expected payment, so it would already be truthful in expectation. The problem is that we just have $p_i = q_i^* x_i$ for some $q_i^* \in [q^*, q^*(1 + \delta\epsilon)]$, and moreover we cannot compute the q_i^* exactly. Our solution is to use the threshold payments as a first approximation, then add a small correction on a set of small probability.

Let the fractional solution values (the steps in Figure 2) be $x^{(1)}, x^{(2)}, \dots, x^{(j)}$ such that the solution for agent i ’s actual bid b_i is $x^{(j)}$. Moreover, let $q^{(k)}$ be the probability that the sale to agent i survives Steps 3 and 4 given that she is selected in Step 2, when the solution corresponding to $x^{(k)}$ is used. To get truthfulness in expectation, when agent i wins, her expected payment should be

$$\frac{1}{q^{(j)} x^{(j)}} \left(b_i x^{(j)} q^{(j)} - \sum_{k=1}^{j-1} (t^{(k+1)} - t^{(k)}) q^{(k)} x^{(k)} - (b_i - t^{(j)}) q^{(j)} x^{(j)} \right).$$

Suppose we use the threshold payment scheme. Given that agent i wins, y_i is distributed uniformly on $[0, x^{(j)}]$. Thus, agent i ’s expected payment is $\frac{1}{x^{(j)}} \sum_{k=1}^j t^{(k)} (x^{(k)} - x^{(k-1)})$, which rearranges to

$$\frac{1}{x^{(j)}} \left(x^{(j)} b_i - \sum_{k=1}^{j-1} (t^{(k+1)} - t^{(k)}) x^{(k)} - (b_i - t^{(j)}) x^{(j)} \right)$$

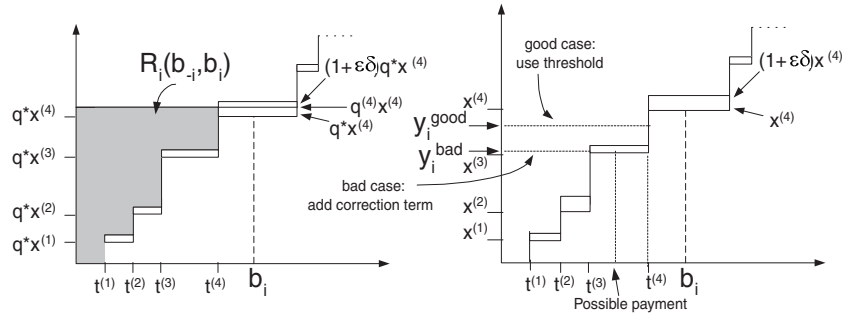


Figure 3. The left graph shows agent i 's probability of success as a function of her bid b_i . The boxes indicate our margin of uncertainty about this probability. Modulo this uncertainty, the shaded area denotes the truthful payment function. The right graph illustrates our payment scheme.

where $x^{(0)} = 0$. Therefore, we must add some correction term to increase this payment by

$$\sum_{k=1}^{j-1} (t^{(k+1)} - t^{(k)}) \frac{x^{(k)}}{x^{(j)}} \left(1 - \frac{q^{(k)}}{q^{(j)}}\right)$$

in expectation.

One way to do this is to add $(t^{(k+1)} - t^{(k)}) \frac{1 - q^{(k)}/q^{(j)}}{\epsilon \delta}$ whenever $y_i \in [x^{(k)}, (1 + \epsilon \delta)x^{(k)}]$, for $k = 1, \dots, j - 1$. See Figure 3. Since we do not know $q^{(k)}$ and $1/q^{(j)}$, we must replace them in the formula with independent unbiased estimators. These estimators can be obtained by running our allocation algorithm. The expected payment is now that given by formula (4.1), so the mechanism is truthful in expectation. When $y_i \notin [x^{(k)}, x^{(k)}(1 + \epsilon \delta)]$ for all k , our payments are just threshold payments. This happens with probability at least $(1 - \epsilon)$ since by Corollary 3.5, consecutive fractional allocations $x^{(k)}$ are spaced by factors of at least $(1 + \delta)$. As we argued in the last section, the threshold mechanism is strongly truthful with high probability. Hence, we get the following. (Recall $\epsilon \ll 1/K^c$.)

Theorem 4.1. *Assuming the item multiplicities are all $\Omega(\ln K)$ as in Proposition 3.1, allocation algorithm RANDROUND combined with the above payment scheme is truthful in expectation and is also truthful with error probability $\epsilon + \frac{1}{K^c}$.*

5. Revenue Considerations

Now we consider the revenue generated by our auction. We show that the expected revenue generated is very close to that generated by a natural fractional relaxation of the VCG mechanism. First, we must define this mechanism.

Recall the VCG mechanism: It chooses a feasible allocation that maximizes the *utilitarian objective function*, which is the total reported valuation of the winning players. That is, it selects some

$$x^*(b) \in \operatorname{argmax}_x \sum_{i \in \mathcal{N}} b_i x_i, \quad (5.1)$$

where x runs over all feasible allocations.

The mechanism computes a bonus for each bidder, based on that bidder's *marginal value*, which is the difference she made in the objective function by participating. Formally, let $V(\mathcal{N} - \mathcal{N}')$ denote the maximum total valuation achievable in (5.1) when the players in \mathcal{N}' are removed. Then bidder i 's marginal value is defined to be $V(\mathcal{N}) - V(\mathcal{N} - i)$. The mechanism charges $P_i(b) = b_i x_i - (V(\mathcal{N}) - V(\mathcal{N} - i))$ to each player i . (This formula evaluates to zero for players who lose.) The utilitarian allocation is clearly monotone. Moreover, the objective function is indifferent about satisfying bidder i when she bids exactly her threshold. So if she wins, then $V(\mathcal{N}) - V(\mathcal{N} - i) = b_i - T_i(b_{-i})$, so the VCG payment is $T_i(b_{-i})$. Notice that if player i bids truthfully, then her profit is equal to her marginal value.

The VCG mechanism is defined with respect to a set of feasible allocations. Usually we maximize over all feasible integer allocations, meaning that each bidder either wins or loses, and no item is oversold. However, we could consider enlarging the set of allowed allocations to permit fractional allocations. That is, we could allow ourselves to let player i win to a fractional extent x_i , which would mean that she receives an x_i fraction of each good in S_i . In other words, we would be maximizing the linear program (3.1), using the actual multiplicities m_j . Of course, we could implement such a mechanism only if the goods were divisible. We assume that player i attains a benefit of $v_i x_i$ from winning to the fractional extent x_i . Thus, she wishes to maximize $\operatorname{profit}_i(b) = v_i x_i(b) - P_i(b)$. Player i 's marginal value to the system is $V(\mathcal{N}) - V(\mathcal{N} - i)$, where $V(\mathcal{N}')$ is the optimal LP value using only the players in \mathcal{N}' . The VCG payment formula becomes $P_i(b) = b_i x_i - (V(\mathcal{N}) - V(\mathcal{N} - i))$. We refer to this resulting mechanism as *fractional VCG*, or FVCG for short.

Similarly, we can define an FVCG mechanism with respect to the artificially reduced multiplicities m'_j . We will show that the expected revenue of our mechanism is almost the same as the revenue generated by the fractional VCG mechanism using the reduced multiplicities.

First, we obtain an expression for the revenue generated by the fractional VCG mechanism. Fixing b_{-i} , how does the optimal allocation $x_i(b_{-i}, b_i)$ change as i increases her bid from 0 upward? Suppose the mechanism always selects some vertex solution of the LP. Initially, $x_i = 0$. The only part of the LP that changes is the direction of the objective function vector, not the polytope of feasible solutions. Thus, the optimal x_i remains zero for an interval until i 's bid hits some threshold $t^{(1)}$. At this point, some other vertex solution with $x_i = x^{(1)} > 0$ becomes optimal. Now this solution remains optimal for some interval, until x_i jumps again at $b_i = t^{(2)}$ to some higher level $x^{(2)}$. Suppose x_i jumps j times at $t^{(1)}, \dots, t^{(j)}$ to new levels $x^{(1)}, \dots, x^{(j)}$ as we raise i 's bid to its actual value b_i . For bids in $(t^{(k)}, t^{(k+1)})$, the LP value increases at rate $x^{(k)}$. Thus, i 's marginal value is $\sum_{k=1}^{j-1} x^{(k)}(t^{(k+1)} - t^{(k)}) + x^{(j)}(b_i - t^{(j)})$, and her price P_i is $b_i x_i$ minus this, which is

$$\sum_{k=0}^{j-1} (x^{(j)} - x^{(k)})(t^{(k+1)} - t^{(k)}),$$

where $x^{(0)} = t^{(0)} = 0$. To visualize this computation, consider Figure 1, with the curve denoting $x_i(b_{-i}, b_i)$ (whereas originally the curve in Figure 1 was the probability of being selected as a winner, which is $\approx q^* x_i(b_{-i}, b_i)$). In our mechanism, the expected payment by agent i is

$$\sum_{k=0}^{j-1} (q^{(j)} x^{(j)} - q^{(k)} x^{(k)})(t^{(k+1)} - t^{(k)}),$$

where $q^{(k)}$ is the probability that the sale to agent i is not cancelled in Steps 3 and 4, if i were to bid between $t^{(k)}$ and $t^{(k+1)}$ (just as in Section 4.2.). Thus, we are comparing vertical strips of equal width, and height $x^{(j)} - x^{(k)}$ for FVCG as opposed to height $q^{(j)} x^{(j)} - q^{(k)} x^{(k)}$ for our mechanism. But

$$\begin{aligned} q^{(j)} x^{(j)} - q^{(k)} x^{(k)} &\geq q^* x^{(j)} - (1 + \epsilon\delta) q^* x^{(k)} \\ &\geq (1 - \epsilon) q^* (x^{(j)} - x^{(k)}) \end{aligned}$$

because $q^{(k)} \in [q^*, q^*(1 + \epsilon\delta)]$ and $x^{(j)} \geq (1 + \delta)x^{(k)}$ for all $k < j$.

Theorem 5.1. *Suppose that each item $j \in \mathcal{G}$ has multiplicity $m_j = \Omega(\ln K)$, as in Proposition 3.1. Then the expected revenue generated by RANDROUND is at least $(1 - \epsilon)q^*$ times the revenue generated by the FVCG mechanism with multiplicities m'_j .*

Under the same conditions on the multiplicities, the probability that player i actually wins is also at least $q^* x_i$. Thus, the auction essentially implements the FVCG mechanism on the artificially reduced multiplicities.

5.1. Comparing Against the “Optimal” Mechanism

It is natural to ask how our revenue compares with that of an “optimal” truthful mechanism, but it turns out that even posing this question correctly is a tricky endeavor. One truthful mechanism is to arbitrarily select a feasible set \mathcal{W} of possible winners, set fixed prices P_i for every bidder in that set, and refuse to sell to any other players. Any player i with $v_i \geq P_i$ will then buy her set at price P_i . If we happen to get lucky and choose \mathcal{W} to be the feasible set of bidders that maximizes the total valuation, and happen to choose $P_i = v_i$ for each $i \in \mathcal{W}$, then we reap the entire valuation as revenue. However, this “omniscient mechanism” hardly seems a fair benchmark. In fact, it is well known that even when auctioning just a single copy of a single item, no truthful mechanism can always attain a guaranteed fraction of the optimal valuation, because there is no way to deal with a single astronomical bidder.

Therefore, in the single item case, [Goldberg et al. 01, Fiat et al. 02, Lavi and Nisan 00] suggest comparing against variants of the VCG mechanism. We have shown that our auction achieves expected revenue approximately equal to that of the FVCG mechanism with a slightly reduced supply of goods. It is easy to construct a pair of examples showing that neither the VCG nor the FVCG mechanism’s revenue dominates the other. Moreover, it is well known that artificially decreasing the supply of goods can sometimes dramatically increase revenues. (See [Goldberg et al. 01] for a striking example.) Therefore, it is unclear how the revenue compares with that of the VCG mechanism using the full supply. This could be an interesting direction for further work. A huge body of literature addresses how various auctions fare under probabilistic assumptions about the valuations, but we are not aware of any such work applied to combinatorial auctions. This is another interesting future direction.

6. Lying about the Set: An Example

It is natural to ask if we can extend our method to handle the case where the set S_i is part of agent i ’s bid (i.e., the case of single-minded bidders, instead of *known* single-minded bidders). The following example shows that, if we insist that agent i wins with probability roughly proportional to the fractional allocation x_i given by the LP, then it is impossible to obtain a mechanism for single-minded bidders that is truthful with high probability.

Suppose there are three items $\{A, B, C\}$ and three bidders $\{1, 2, 3\}$. One copy of each item is available. Suppose bidder 1 bids 2 units for set $\{B, C\}$ (the truth), bidder 2 bids $\frac{3}{2}$ units for set $\{A, B\}$, and bidder 3 bids $\frac{3}{2}$ units for set $\{A, C\}$. Then the LP solves to $x = (\frac{1}{2}, \frac{1}{2}, \frac{1}{2})$ (with total valuation $\frac{5}{2}$). If bidder

1 lies by increasing her set to $\{A, B, C\}$ (but still bidding 2 units for it), then the LP solves to $x = (1, 0, 0)$ (with total valuation 2). Suppose that we actually had three copies of each item available, but were just using 1 as the reduced multiplicities. Then no item is ever oversold. Thus, our mechanism (without the drop step) implements fractional VCG in expectation, so is still truthful in expectation. However, we show that with probability $\frac{1}{2}$, bidder 1 benefits by lying, and hence, we do not have truthfulness with high probability.

When player 1 bids her true set, her probability $p_1(b_{-1}, b_1)$ of winning stays constant at $\frac{1}{2}$ for $b_1 \in (0, 3]$. Thus, the threshold payment scheme of Section 4.1 would charge her a price of 0 (leading to a profit of 2) for any $y_1 \in [0, \frac{1}{2}]$. When y_1 falls in $(\frac{1}{2}, 1]$, however, she loses and has a profit of 0.

On the other hand, if she lies about her set as above, then $p_1(b_{-1}, b_1)$ jumps from 0 to 1 at $b_1 = \frac{3}{2}$. Thus, with probability 1, she wins and pays $\frac{3}{2}$, so her net profit is $\frac{1}{2}$, irrespective of y_1 . Note that when her rounding variable y_1 lands in $(\frac{1}{2}, 1]$ (which happens with probability $\frac{1}{2}$), her profit when she lies is more than her profit when she tells the truth. Thus, with probability $\frac{1}{2}$, the mechanism is not truthful.

We can extend this example to arbitrarily high item multiplicities by simply adding in appropriate bidders j who bid high enough that they are fully satisfied (i.e., the optimal solution has $x_j = 1$). Note that in this case the reduced multiplicities are smaller than the actual multiplicities only by an additive -2 , not a multiplicative $\frac{1}{3}$.

7. Conclusions

We have shown a general technique to modify a linear program rounding algorithm to make it monotone. This gives an approximately efficient truthful mechanism (in expectation and with high probability) for the combinatorial auction problem with single parameter agents.

The simple rounding algorithm can be derandomized using pessimistic estimators [Raghavan 88]. It would be interesting to see if the algorithm can be derandomized maintaining its monotonicity.

Finally, this scheme gives a truthful mechanism for known single-minded bidders; an open problem is to relax this constraint.

8. Acknowledgments

We thank the two anonymous referees for contributing many comments that helped us to clarify the exposition. Aaron Archer was supported by the Fannie and John Hertz Foundation. Christos Papadimitriou was supported by an NSF ITR grant, and by the

EU project OPUS, IST-2001-33464. Kunal Talwar was supported by NSF via grants CCR-9820897 and CCR-0105533. Éva Tardos was supported in part by NSF grant CCR-9700163 and ONR grant N00014-98-1-0589.

References

- [Archer and Tardos 01] Aaron Archer and Éva Tardos. “Truthful Mechanisms for One-Parameter Agents.” In *Proceedings of the 42nd IEEE Symposium on Foundations of Computer Science*, pp. 482–491, Los Alamitos, CA: IEEE, 2001.
- [Archer and Tardos 02] Aaron Archer and Éva Tardos. “Frugal Path Mechanisms.” In *Proceedings of the 13th Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 991–999, New York/Philadelphia: ACM/SIAM, 2002.
- [Bar-Yossef et al. 02] Ziv Bar-Yossef, Kirsten Hildrum, and Felix Wu. “Incentive-Compatible Online Auctions for Digital Goods.” In *Proceedings of the 13th Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 964–970, New York/Philadelphia: ACM Press/SIAM, 2002.
- [Clarke 71] Edward H. Clarke. “Multipart Pricing of Public Goods.” *Public Choice* 8 (1971), 17–33.
- [Cramton 02] Peter Cramton. “Spectrum Auctions.” In *Handbook of Telecommunications Economics*, edited by Martin Cave, Sumit Majumdar, and Ingo Vogelsang, pp. 605–639. Amsterdam: Elsevier, 2002.
- [Fiat et al. 02] Amos Fiat, Andrew Goldberg, Jason Hartline, and Anna Karlin. “Competitive Generalized Auctions.” In *Proceedings of the 34th ACM Symposium on Theory of Computing*, pp. 72–81, New York: ACM Press, 2002.
- [Fudenberg and Tirole 91] Drew Fudenberg and Jean Tirole. *Game Theory*. Cambridge, MA: MIT Press, 1991.
- [Goldberg et al. 01] Andrew Goldberg, Jason Hartline, and Andrew Wright. “Competitive Auctions and Digital Goods.” In *Proceedings of the 12th Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 735–744, New York/Philadelphia: ACM Press/SIAM, 2002.
- [Groves 73] Theodore Groves. “Incentives in Teams.” *Econometrica* 41:4 (1973), 617–631.
- [Lavi and Nisan 00] Ron Lavi and Noam Nisan. “Competitive Analysis of Incentive Compatible On-Line Auctions.” In *Proceedings of the 2nd ACM Conference on Electronic Commerce*, pp. 233–241, New York: ACM Press, 2000.
- [Lehmann et al. 02] Daniel Lehmann, Liadan O’Callaghan, and Yoav Shoham. “Truth Revelation in Approximately Efficient Combinatorial Auctions.” *Journal of the ACM* 49 (2002), 577–602.
- [McAfee and McMillan 88] R. Preston McAfee and John McMillan. “Multidimensional Incentive Compatibility and Mechanism Design.” *Journal of Economic Theory* 46 (1988), 335–354.

- [Mirrlees 71] James Mirrlees. “An Exploration in the Theory of Optimum Income Taxation.” *Review of Economic Studies* 38 (1971), 175–208.
- [Motwani and Raghavan 95] Rajeev Motwani and Prabhakar Raghavan. *Randomized Algorithms*. New York: Cambridge University Press, 1995.
- [Mu’alem and Nisan 02] Ahuva Mu’alem and Noam Nisan. “Truthful Approximation Mechanisms for Restricted Combinatorial Auctions.” In *Proceedings of the 18th National Conference on Artificial Intelligence and 14th Conference on Innovative Applications of Artificial Intelligence*, pp. 379–384. Menlo Park, CA: AAAI Press, 2002.
- [Nisan and Ronen 01] Noam Nisan and Amir Ronen. “Algorithmic Mechanism Design.” *Games and Economic Behavior* 35 (2001), 166–196.
- [Nisan and Ronen 00] Noam Nisan and Amir Ronen. “Computationally Feasible VCG Mechanisms.” In *Proceedings of the 2nd ACM Conference on Electronic Commerce*, pp. 242–252, New York: ACM Press, 2000.
- [Raghavan 88] Prabhakar Raghavan. “Probabilistic Construction of Deterministic Algorithms: Approximating Packing Integer Programs.” *Journal of Computer and System Sciences* 37 (1988), 130–143.
- [Ronen 00] Amir Ronen. “Algorithms for Rational Agents.” In *27th Conference on Current Trends in Theory and Practice of Informatics*, Lecture Notes in Computer Science 1963, pp. 56–70. Heidelberg: Springer-Verlag, 2000.
- [Schrijver 86] Alexander Schrijver. *Theory of Linear and Integer Programming*. New York: Wiley, 1986.
- [Schummer 01] James Schummer. “Almost-Dominant Strategy Implementation.” Unpublished manuscript, 2001.
- [Vickrey 61] William Vickrey. “Counterspeculation, Auctions and Competitive Sealed Tenders.” *Journal of Finance* 16 (1961), 8–37.

Aaron Archer, Operations Research Department, Cornell University, Ithaca, NY 14853 (aarcher@orie.cornell.edu)

Christos Papadimitriou, Division of Computer Science, University of California, Berkeley, CA 94720 (christos@cs.berkeley.edu)

Kunal Talwar, Division of Computer Science, University of California, Berkeley, CA 94720 (kunal@cs.berkeley.edu)

Éva Tardos, Computer Science Department, Cornell University, Ithaca, NY 14853 (eva@cs.cornell.edu)

Received March 10, 2003; accepted March 26, 2003.