An Approximate Truthful Mechanism for Combinatorial Auctions with Single Parameter Agents

Aaron Archer* Christos Papadimitriou† Kunal Talwar‡ Éva Tardos§

Abstract

Mechanism design seeks algorithms whose inputs are provided by selfish agents who would lie if advantageous. 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.

*Operations Research Department, Cornell University, Ithaca, NY 14853. Email: aarcher@orie.cornell.edu. Supported by the Fannie and John Hertz Foundation.

†Division of Computer Science, University of California, Berkeley, CA 94720. Email: christos@cs.berkeley.edu. Supported by an NSF ITR grant, and by the EU project OPUS, IST-2001-33464.

‡Division of Computer Science, University of California, Berkeley, CA 94720. Email: kunal@cs.berkeley.edu. Supported by NSF via grants CCR-9820897 and CCR-0105533.

§Computer Science Department, Cornell University, Ithaca, NY 14853. Email: eva@cs.cornell.edu. Research supported in part by NSF grant CCR-9700163 and ONR grant N00014-98-1-0589.

1 Introduction

The combinatorial auction is one of the basic mechanisms of electronic commerce. Many large-scale combinatorial auctions have also 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 over $40 billion [5]. 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 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. 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. In this paper 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. In this paper 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, and 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 [5]. The well-known Vickrey-Clarke-Groves (VCG) mechanism [18, 4, 9] 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 SET-PACKING 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 [14].

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 [11] 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. Mu'alem and Nisan [12] consider the case of *known single-minded bidders*, where the sets are known, and each agent's only private data is its 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 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 of items $\mathcal{I}$ among a set $\mathcal{N}$ of n bidders. Each bidder i has a valuation function $v_i : 2^{\mathcal{I}} \rightarrow \mathbb{R}^+$ that describes her preferences over the various subsets of items. For $S \subseteq \mathcal{I}$, $v_i(S)$ 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. In this paper, 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 [12]. 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. A single-minded bidder wants only one copy of each good in her desired set.

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

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, $v_i \in \text{argmax}_{b_i} \text{profit}(b_{-i}, b_i)$, 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 $P_i(b_{-i}, b_i)$ depend 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 $\mathcal{M}_\omega$ 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 *strongly truthful.* This means that for *every* ω the mechanism $\mathcal{M}_\omega$ is truthful. This concept has been used in [13, 3, 7, 6], but it is very restrictive.

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 [1], i.e., the mechanism is *truthful in expectation.* Two orthogonal notions are that a player may benefit from lying, but not by much [17], or only with a small probability. We pursue the first and third approaches.

We say a randomized mechanism is *strongly truthful with error probability* ϵ if for each b_{-i} and each v_i we have $Pr[v_i \notin \text{argmax}_{b_i} \text{profit}_i(b_{-i}, b_i)] \leq \epsilon$. If ϵ is inverse polynomial in some specified parameters of the auction (such as number of items or bidders) then we say the mechanism is *strongly 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 strongly 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 $\Omega(\ln K)$ (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 variant of VCG.

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 $b_i < T_i(b_{-i})$. This characterization has been noted many places, such as [7, 11, 1, 2, 16, 12]. 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 monotone in her bid b_i, and her expected payment is equal to a certain integral of the function p_i [1].

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.

$$\begin{array}{lll} \text{maximize} & \sum_{i\in\mathcal{N}} b_i x_i & \\ \text{subject to:} & \sum_{i:j\in S_i} x_i \le m'_j & \text{for all } j\in\mathcal{I} \\ & 0\le x_i\le 1 & \text{for all } i\in\mathcal{N} \end{array} \quad (1)$$

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 bids b_i (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's 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 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 over-sold.

Chernoff bound. *Let $X_1, \ldots, X_n$ be independent Poisson trials such that, for $1 \le i \le n$, $Pr[X_i = 1] = p_i$. Then for $X = X_1 + \ldots + X_n$, $\mu \ge p_1 + \ldots + p_n$, and any $\alpha < 2e - 1$ we have*

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

Proposition 3.1 *Suppose that each item $j \in \mathcal{I}$ has multiplicity $m_j = \Omega(\ln K)$. Then the probability that a given item is over-sold is at most $\frac{1}{K^{c+1}}$ (where the multiplicative constant inside the Ω is $\frac{4(c+1)}{\epsilon'^2(1-\epsilon')}$).*

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 $i \ne i_0$ 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,

$$b\cdot x \ \ \ge \ \ b\cdot x' \quad (2)$$

Similarly,

$$b'\cdot x' \ \ \ge \ \ b'\cdot x \quad (3)$$

If $x = x'$, we are already done. Now, since the tie breaking rule is the same, if the two vertices are distinct, one of the two inequalities must be strict. Adding (2) and (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. ■

From the above result, it follows that the probability of any agent being rounded 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 is not allocated the good even though she is initially rounded to 1, and moreover, this probability is not necessarily monotone in the agent's bid. As a result, we do not yet have a truthful mechanism.

3.1 Dealing with over-sold 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 make the overall probability of any agent i being allocated her set 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 (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 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

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

Let $I_{i_0} = \{i : S_i \cap S_{i_0} \ne \emptyset\}$. The variables $x_i : i \in I_{i_0}$ form a feasible solution to the scaled linear program (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 over-sold, 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 with probability $1-(q^*/q_i)$, so that the probability that agent i survives through the end of Step 4 is exactly $x_i q^*$. We would then have a monotone allocation algorithm.

3.2 Dealing with over-sold items: important details

Note that the algorithm described above requires us to exactly compute the probability q_i. 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 which follows from Cramer's rule and the fact that the reduced multiplicities m'_j are integers.

Observation 3.3 *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 $D \leq m!$.*

Corollary 3.4 *Let x^1, x^2 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 $x_i^1 = x_i^2$ or $x_i^1 \geq x_i^2(1+\delta)$ or $x_i^2 \geq x_i^1(1+\delta)$, where $\delta = (1/m!)^2$.*

Corollary 3.4 along with Lemma 3.2 imply 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's are independent trials of the above experiment, 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, and obtain the following.

Lemma 3.5 *The estimator Y defined above is at most $1/q^*$, and $E[Y] \in [1/q_{i_0}, (1+\delta\epsilon)/q_{i_0}]$.*

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

We first argue that this mechanism is monotone.

Theorem 3.6 *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

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

where the last inequality follows from Corollary 3.4. Since $\epsilon < 1$, this shows that the above LP rounding algorithm is monotone. ■

Theorem 3.7 *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

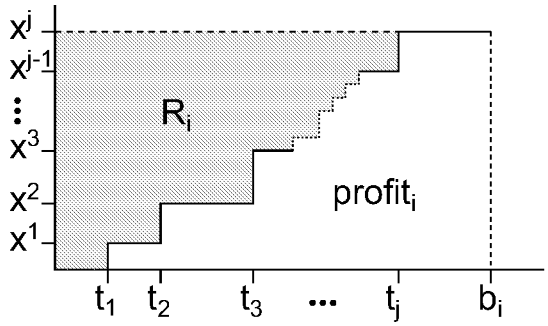

Figure 1: The graph shows i's probability $p_i(b_{-i}, b_i)$ of winning as a function of her bid b_i. The gray area R_i is her expected payment. If b_i is a truthful bid, then the white area is her expected profit.

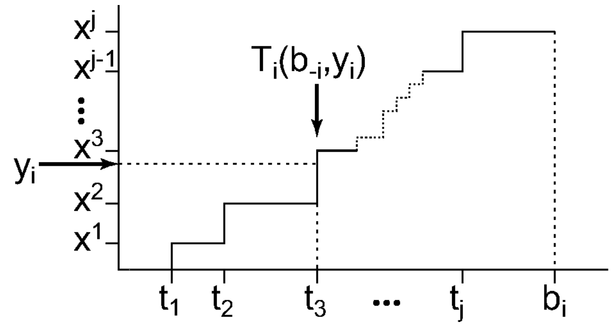

Figure 2: The graph shows i's fractional allocation as a function of her bid b_i. It is a step function that is flat while one vertex of the LP (1) stays optimal, then jumps when another vertex becomes preferred. Since y_i lands in (x^2, x^3), i's payment is t_3.

a solution to our LP (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. ■

4 Computing Payments

It is not obvious how one should compute payments for the winners. There is a payment scheme given in [1] 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 strong 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. Since b_{-i} is fixed throughout this discussion, we suppress it in the notation. A result in [1] says that to guarantee truthfulness in expectation, i's payment should be

$$R_i = p_i(b_i)b_i - \int_0^{b_i} p_i(u)du. \quad (4)$$

See Figure 1. In order to satisfy individual rationality, we must charge 0 to losing bidders. Thus, one truthful payment scheme is to charge $R_i/p_i(b_i)$ to agent i if she wins.

Note that $p_i(u)$ is a step function which jumps whenever the selected vertex in the LP (1) changes and is flat elsewhere. Thus, it could have exponentially many breakpoints, so we can not obviously 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 for the integral. Let X be an independent unbiased estimator for $1/p_i(b_i)$. 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 truthful with high probability.

In the bidder selection step, let us perform the rounding by selecting n independent uniform$[0, 1]$ random variables $y_1, \ldots, y_n$, 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

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

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^1 and x^2. Therefore, $T_i(b_{-i}, y_i)(x_i^1 - x_i^2) = \sum_{j \neq i} b_j(x_j^2 - x_j^1)$. Assuming all 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 Observation 3.3), so we can use binary search and the method of Diophantine approximation to compute it exactly in polynomial time.

If our mechanism never had to throw away any initial winners, then our allocation algorithm would be strongly truthful. Suppose we fix a particular realization ω of the vector of random variables $y_1, \ldots, y_n$. 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 $\frac{1}{K^c}$, our mechanism is strongly 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 were 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 $x^1, x^2, \ldots, x^j$ 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 is cancelled in Steps 3 and 4, 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

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

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 $k = 1, \ldots, j - 1$. 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), 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 3.4, 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, algorithm* RANDROUND *with the above payment scheme is truthful in expectation and is also strongly 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.

First 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

$$x^*(b) \in \operatorname{argmax}_x \sum_{i \in \mathcal{N}} b_i x_i, \tag{5}$$

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

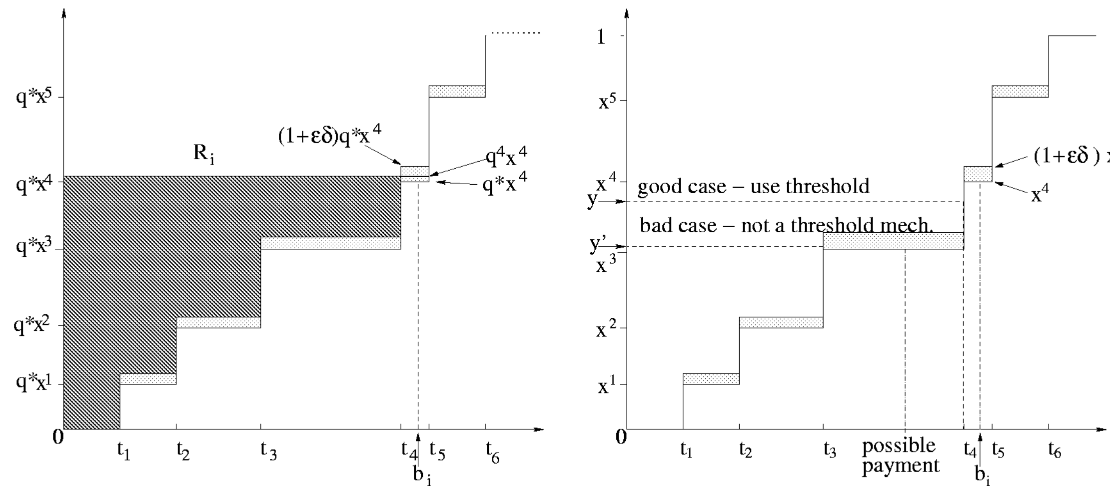

Figure 3: The left graph shows agent i's probability of success as a function of her bid b_i. The boxes indicate our margin of uncertainty about this probability. Modulo this uncertainty, the shaded area denotes the truthful payment function. The right graph illustrates our payment scheme.

valuation achievable in (5) 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 over-sold. 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 (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 $\text{profit}_i(b) = v_i x_i(b) - P_i(b)$. 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 $t_1, \ldots, t_j$ to new levels $x^1, \ldots, x^j$ 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 the 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 $j \in \mathcal{I}$ 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 of player i actually winning 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, [7, 8, 6, 10] 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 [7] for a striking example.) Therefore, it is unclear how the revenue compares with that of the VCG mechanism using the full supply.

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 it is impossible to obtain a mechanism for single-minded bidders that is strongly truthful with high probability, if we want the probability that agent i wins to be roughly proportional to the fractional allocation x_i given by the LP.

Suppose there are three items and three bidders. One copy of each item is available. Suppose bidder 1 bids 2 for set $\{2,3\}$ (the truth), bidder 2 bids $\frac{3}{2}$ for set $\{1,2\}$ and bidder 3 bids $\frac{3}{2}$ for set $\{1,3\}$. 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 $\{1,2,3\}$, 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 over-sold. Thus, our mechanism (without the drop step) implements fractional VCG in expectation, so is still truthful in expectation, but not truthful with high probability because the lie increased x_1 from $\frac{1}{2}$ to 1.

Let us flesh out the details. 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, she wins at price 0 with probability $\frac{1}{2}$, and loses with probability $\frac{1}{2}$, for an expected profit of 1.

If she lies about her set, 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 expected profit decreases to $\frac{1}{2}$. But when her rounding variable y_1 lands in $(\frac{1}{2}, 1]$ she does better by lying, whereas when her rounding variable y_1 lands in $[0, \frac{1}{2}]$ she does better by telling the truth.

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 (in expectation and with high probability) mechanism for the combinatorial auction problem with single parameter agents.

The simple rounding algorithm can be derandomized using pessimistic estimators [15]. 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.

References

[1] A. Archer and É. Tardos. Truthful mechanisms for one-parameter agents. In *FOCS 2001*, 482–491.

[2] A. Archer and É. Tardos. Frugal path mechanisms. In *SODA 2002*, 991–999.

[3] Z. Bar-Yossef, K. Hildrum, and F. Wu. Incentive-compatible online auctions for digital goods. In *SODA 2002*, 964–970.

[4] E. H. Clarke. Multipart pricing of public goods. *Public Choice*, 8:17–33, 1971.

[5] P. Cramton. Spectrum auctions. In M. Cave, S. Majumdar, and I. Vogelsang, editors, *Handbook of Telecommunications Economics*. Elsevier, Amsterdam, 2002.

[6] A. Fiat, A. Goldberg, J. Hartline, and A. Karlin. Competitive generalized auctions. In *STOC 2002*, 72–81.

[7] A. Goldberg, J. Hartline, and A. Wright. Competitive auctions and digital goods. In *SODA 2001*, 735–744.

[8] A. V. Goldberg and J. D. Hartline. Competitive auctions for multicast content. Unpublished manuscript, 2000.

[9] T. Groves. Incentives in teams. *Econometrica*, 41(4):617–631, Jul 1973.

[10] R. Lavi and N. Nisan. Competitive analysis of incentive compatible on-line auctions. In *ACM Conf. on Electronic Commerce 2000*, 233–241.

[11] D. Lehmann, L. O'Callaghan, and Y. Shoham. Truth revelation in rapid, approximately efficient combinatorial auctions. In *ACM Conf. on Electronic Commerce 1999*, 96–102.

[12] A. Mu'alem and N. Nisan. Truthful approximation mechanisms for restricted combinatorial auctions. To appear in *AAAI 2002*.

[13] N. Nisan and A. Ronen. Algorithmic mechanism design. In *STOC 1999*, 129–140.

[14] N. Nisan and A. Ronen. Computationally feasible VCG mechanisms. In *ACM Conf. on Electronic Commerce 2000*, 242–252.

[15] P. Raghavan. Probabilistic construction of deterministic algorithms: approximating packing integer programs. *J. Comp. and System Sciences*, 37:130–143, 1988.

[16] A. Ronen. Algorithms for rational agents. In *Conf. on Current Trends in Theory and Practice of Informatics 2000*, 56–70.

[17] J. Schummer. Almost-dominant strategy implementation. Unpublished manuscript, 2001.

[18] W. Vickrey. Counterspeculation, auctions and competitive sealed tenders. *J. Finance*, 16:8–37, 1961.