Truthful Randomized Mechanisms for Combinatorial Auctions

Shahar Dobzinski Noam Nisan Michael Schapira*

March 20, 2006

Abstract

We design two computationally-efficient incentive-compatible mechanisms for combinatorial auctions with general bidder preferences. Both mechanisms are randomized, and are incentive-compatible in the universal sense. This is in contrast to recent previous work that only addresses the weaker notion of incentive compatibility in expectation. The first mechanism obtains an $O(\sqrt{m})$-approximation of the optimal social welfare for arbitrary bidder valuations – this is the best approximation possible in polynomial time. The second one obtains an $O(\log^2 m)$-approximation for a subclass of bidder valuations that includes all submodular bidders. This improves over the best previously obtained incentive-compatible mechanism for this class which only provides an $O(\sqrt{m})$-approximation.

1 Introduction

1.1 Background

The field of Algorithmic Mechanism Design attempts to design efficient mechanisms for decentralized computerized settings. These mechanisms must take into account both the strategic behavior of the different participants and the usual algorithmic efficiency considerations. Target applications include many types of protocols for Internet environment that necessitate looking at both issues – strategic and algorithmic – together. For an introduction see [20].

The basic strategic notions are taken from the field of mechanism design – a subfield of economic theory (see [18, 23]), and in most of the work in computational settings, as in this one, the very robust notion of equilibrium in dominant strategies is used. It is well known ([18], see [20]) that without loss of generality, we can limit ourselves to looking at "incentive compatible" mechanisms, also known as "truthful" mechanisms or "strategy-proof" mechanisms. In such mechanisms participants are always rationally motivated to correctly report their private information.

The main difficulty in this field is the fact that the basic technique of mechanism design – namely VCG mechanisms [25, 4, 11] – can only be applied in cases where the exact optimal outcome is achieved. However, in most interesting computational applications, exact optimization is NP-hard, and computationally-speaking we must settle for approximations or heuristics. As was observed in [20, 17], the VCG technique cannot be applied in such cases, and in fact [21] showed that this inapplicability was essentially universal. Thus, the challenge is to design alternative incentive-compatible mechanisms for interesting applications.

The problem of combinatorial auctions has gained the status of the paradigmatic problem of this field. For a thorough overview see [5]. In a combinatorial auction, m items are auctioned to n players. Each player i has a valuation function v_i that describes his value $v_i(S)$ for each subset

*The School of Computer Science and Engineering, The Hebrew University of Jerusalem, {shahard, noam, mikesch}@cs.huji.ac.il.

S of items. The basic question is how to construct the auction mechanism that allocates all the items in a way that maximizes the social welfare $\Sigma_i v_i(S_i)$ where S_i is the set of items allocated to bidder i. This problem indeed exhibits the basic issues of algorithmic mechanism design: finding the exact optimum is computationally hard, even for the most interesting special cases, but several approximation algorithms, with varying approximation ratios, are known for the general case as well as for various interesting special cases [16, 6, 7, 8]. However, these approximation algorithms do not yield incentive compatible mechanisms.

In a landmark paper, Lehmann et al [17] were able to design an incentive-compatible, efficiently-computable, approximation mechanism – which achieves an approximation ratio that is as good as computationally possible $\Theta(\sqrt{m})$ [24] – for the special case of "single-minded bidders". This is the case in which each bidder is only interested in a single bundle of goods. For this special case, as well as some other single-parameter scenarios a host of incentive compatible mechanisms have been designed in the last few years (e.g., [19, 1, 10, 9]). However, almost nothing is known for more general cases in which bidders have complex multi-dimensional preferences. Only two results are known in multi-dimensional settings[1]: the first is a pair of algorithms that completely optimize over a very restricted range of allocations and then use the usual VCG mechanism. These get a barely better than trivial approximation ratio of $O(m/\sqrt{\log m})$ for the general case [12] and a weak $O(\sqrt{m})$ for the "complement-free" case [6] – both ratios being quite far from what is computationally possible. The second result is the mechanism of [2] that applies only to the special case of auctions with many duplicates of each good and indeed is *not* a VCG mechanism. Some evidence showing that obtaining a non-VCG incentive-compatible mechanism for combinatorial auctions and related problems would be difficult was given in [14].

1.2 Randomized Mechanisms

It was observed in [20] that randomized mechanisms can sometimes provide better approximation ratios than deterministic ones. There are two possible definitions for incentive compatibility of a randomized mechanism. The first and stronger one, defines an incentive-compatible randomized mechanism as a probability distribution over incentive compatible deterministic mechanisms. Thus, this definition requires that for any fixed outcome of the random choices made by the mechanism, players still maximize their utility by reporting their true valuations. This definition was used in [20, 10, 9], and will be called incentive compatible in the *universal sense*. The weaker definition only requires that players maximize their *expected* utility, where the expectation is over the random choices of the mechanism (but still for every behavior of the other players). This was used in [15, 8] (see below), and will be called incentive compatibility *in expectation*.

There are two major implications of the difference between these two notions:

1. **Attitude towards risk:** randomized mechanisms that are incentive compatible in expectation only motivate *risk-neutral* bidders to act truthfully. Risk-averse bidders may benefit from strategic behavior. In contrast, the universal sense of incentive compatibility applies to any attitude towards risk, as it applies to every possible realization of the random coins.

2. **Knowledge of the randomization results:** randomized mechanisms that are incentive compatible in expectation induce truthful behavior only as long as players have no information about the outcomes of the random coin flips before they need to act. Thus, in order to ensure truthful behavior the mechanism must utilize cryptography-grade randomness, and keep it secret from the players. In contrast, any randomization that is effective algorithmically suffices to ensure truthful behavior in the universal case. (In a similar vein, technically speaking,

[1] This is true not only for combinatorial auctions but also for any other computationally-hard problem.

using a pseudorandom generator will destroy the formal incentive compatibility properties of randomized mechanisms that are incentive compatible in expectation, due to the slight – sub-polynomial – change in probabilities of outcomes.)

In the recent [15] a rather general technique was developed for converting approximation algorithms into randomized mechanisms that are incentive compatible in *expectation*. The technique is based on randomized rounding of the LP relaxation, and relies on a clever representation of the LP solution as a scaled convex combination of integer solutions. In particular, they design a randomized mechanism for general combinatorial auctions that is incentive compatible in expectation and obtains the computationally-optimal approximation ratio of $O(\sqrt{m})$. Very recently, [8] used a different but somewhat related randomized rounding procedure to obtain another randomized mechanism for the case of combinatorial auctions with complement-free bidders. This mechanism is, again, incentive-compatible in expectation, and achieves an approximation ratio of $O(\frac{\log m}{\log \log m})$, which is worse than what he obtains algorithmically – a ratio of 2.

1.3 Our Results

We present the first randomized mechanism for combinatorial auctions that is incentive compatible is the universal sense. This is another step towards the "holy grail" of obtaining a deterministic one.

Theorem: There exists a polynomial-time computable randomized mechanism for combinatorial auctions with general bidders that is incentive compatible in the universal sense and obtains a $O(\sqrt{m})$ approximation ratio.[2]

The algorithm runs in time that is polynomial in the natural parameters of the problem: the number of players n and the number of items m. Access to the (exponentially long) valuation functions of the players is done using the usual demand queries [3, 6, 7], in which bidders are presented with a vector of item prices $p_1...p_m$ and reply with the set of items S that maximizes their utility under these prices $v(S) - \sum_{j \in S} p_j$. The approximation factor mentioned in the theorem is in expectation, however, our result is technically stronger: for any fixed $\epsilon > 0$ we provide a mechanism that obtains $\frac{\sqrt{m}}{poly(\epsilon)}$-approximation with probability of at least $1 - \epsilon$.

Our techniques are quite simple, completely different than the methods of [15, 8], and do not rely on the LP-relaxation of the problem. They are more in line with the random sampling methods that were used for auctioning "digital goods" [10, 9]. These techniques can be viewed as providing a general framework for obtaining randomized incentive compatible mechanisms in the universal sense. In particular, a significant property of this framework is that it provides, for any $\epsilon > 0$, a mechanism that achieves an approximation ratio not just in expectation, but with probability $1 - \epsilon$. We stress that this cannot be achieved by the usual techniques of amplification, since repetition can destroy the incentive properties. Using the same framework, we are also able to design improved mechanisms for the important special case of submodular valuations, and actually even for a more general class of valuations termed "XOS" in [16] and "fractionally-subadditive" in [8][3]. This improves over the truthful deterministic $O(\sqrt{m})$-approximation achieved in [6].

Theorem: There exists a polynomial-time computable randomized mechanism for combinatorial auctions with submodular bidders that is incentive compatible in the universal sense and obtains a $O(\log^2 m)$ approximation ratio.

[2] Somewhat unusually, the equilibrium obtained is in dominant strategies even for the adaptive query model which usually only supports ex-post equilibria.

[3] For the XOS class, the bidders must also be able to answer, so called, XOS queries [6].

Beyond the use of randomization, this theorem is sub-optimal in two other senses, which remain as open problems: first, the approximation ratio achieved is worse than the ratio of $\frac{e}{e-1}$ that is computationally possible [7]; second, our mechanism does not apply to the somewhat wider class of complement-free valuations that is handled in [8, 6].

The major open problem left is that of finding *deterministic* $O(\sqrt{m})$-approximation efficiently-computable incentive-compatible mechanisms for combinatorial auctions.

2 Preliminaries

In a combinatorial auction, a set M of items, $M = \{1, ..., m\}$, is sold to n bidders. Every bidder values *bundles* of items, rather than only assigning values to single items. The value that bidder i assigns to bundle S is defined by a valuation function $v_i : 2^M \to \mathbb{R}^+$. Two standard assumptions regarding each bidder i, are that v_i is normalized ($v_i(\emptyset) = 0$), and monotone (for every $S \subseteq T \subseteq M$, $v_i(S) \leq v_i(T)$). The *allocation problem* is to partition the items between the bidders in a way that maximizes the "total social welfare". I.e., to find a partition $S_1, ..., S_n$ of M, that maximizes $\Sigma_i v_i(S_i)$.

Even though the size of the "input" is exponential in m (each v_i is described by 2^m real numbers) we require algorithms to run in time polynomial in the natural parameters of the problem, m and n. An important issue is how the input can be accessed. In this paper we follow the "black box" approach: we assume that we are given an oracle for each valuation function. The oracle is limited to some predefined type of queries. A common type of query is the *demand query* (e.g., [6, 7, 3]). A demand query to a valuation v_i specifies a vector $p = (p_1 ... p_m)$ of "item prices". The answer to the query is a set that would be "demanded" by the queried bidder under these item prices. I.e., a subset S that maximizes $v_i(S) - \sum_{j \in S} p_j$.

In this paper we seek algorithms that are incentive compatible (a.k.a. truthful). That is, algorithms which ensure that it is in the best interest of each of the bidders to *always* reveal his true preferences when asked. In the case of randomized mechanisms this translates to being *incentive compatible in the universal sense* – randomized mechanisms that are a probability distribution over incentive compatible deterministic mechanisms. In other words, telling the truth is the *dominant strategy* of each bidder, regardless of the coins tossed by the mechanism. This is a much stronger requirement than *incentive compatibility in expectation* (see [15]).

Some special cases of combinatorial auctions have recently received great attention. In particular, combinatorial auctions in which all bidders are known to have submodular valuations are the subject of extensive research (e.g., [16, 6, 13, 7]). A valuation v is submodular if $v(S \cup T) + v(S \cap T) \leq v(S) + v(T)$ for all $S, T \subseteq M$. All submodular valuations are known to be strictly contained in the more general class of valuations termed "XOS" in [16], and "fractionally-subadditive" in [8]. A valuation v is said to be XOS if there are additive valuations $\{a_1, ..., a_t\}$, such that $v(S) = \max_k \{a_k(S)\}$ for all $S \subseteq M$[4]. See [7] for a more thorough explanation. For every XOS valuation $v = \max_k \{a_k\}$, and bundle S, we call an additive valuation a such that $a(S) = \arg\max_k \{a_k(S)\}$ a *maximizing clause* for S in v_i. We require XOS bidders to be able to answer *XOS queries*. In this type of queries the question is in the form of a bundle and the answer is a maximizing clause for that bundle.

3 A Framework For Designing Incentive-Compatible Mechanisms

The design of a randomized *approximation algorithm* comprises two basic steps: first, we are interested in making sure that the expected value of the solution produced by the algorithm is "not

[4]A valuation a is additive if for every $S \subseteq M$, $a(S) = \Sigma_{j \in S} a(\{j\})$

far" from the optimum. Second, we wish to be able to find a solution with a value "close" to the expectation with high probability. Usually, the main difficulty is in achieving the first goal and proving that a solution close to the expectation can be obtained with some (perhaps polynomially low) probability. Amplification of the probability of success is then easily attainable by running the algorithm a polynomial number of times and choosing the best solution.

In contrast, the design of a randomized *mechanism* is inherently different: in general, running a mechanism multiple times and choosing the best output does not preserve the truthfulness of the mechanism. In addition, it is well known that in order to ensure truthfulness, the price a bidder pays for the bundle he is allocated cannot depend on information he provides. The framework we introduce here helps us overcome these problems.

The framework relies on the examination of two distinct possible cases: either there is one bidder such that allocating all items to him is a good approximation to the welfare, or there is no such bidder. I.e., there is no "small" group of bidders that contributes "a lot" to the optimal solution. In the first case, achieving a good approximation is easy - allocate all items to that bidder. In the second and more complicated case, we will perform a fixed-price auction, and will have to prove that we get a good approximation. The key observation used in handling the second case is that two randomly chosen groups that consist (in expectation) of a constant fraction of the bidders have many properties in common (e.g., both hold a constant fraction of the total welfare.) This idea is similar to the main principle in random-sampling auctions for "digital goods" [10, 9]. However, our situation is much more complex due to the multi-parameter setting of combinatorial auctions, in contrast to the single-parameter setting of [10, 9]. In addition, our goal is to optimize the welfare, and not maximize revenue. Moreover, we do not assume that all the items are identical and that there is an unlimited supply of items, as in the case of "digital goods". From a computational point of view, another difference is that the problems we consider are NP-hard to approximate.

The framework allows us, with high probability, to distinguish between the two cases, and provides us with the tools for finding the price used in the fixed-price auction. The main difficulty in tailoring the framework to a specific setting is showing that the fixed-price auction guarantees a good approximation. Indeed, in the two mechanisms we are about to present in this paper the price used in the fixed-price auction is determined in a completely different manner.

The Framework:

`Phase I: Partitioning the Bidders`

We assign each bidder to exactly one of the following three sets: SEC-PRICE with probability $1-\epsilon$, FIXED with probability $\frac{\epsilon}{2}$, and STAT with probability $\frac{\epsilon}{2}$. Only bidders from SEC-PRICE will be allowed to participate in the second-price auction. Bidders in STAT will never get any items, so we can safely use this group to gather the necessary statistics (see next phase). The bidders in FIXED will be the only bidders who participate in the fixed-price auction.

`Phase II: Gathering Statistics`

The goal in this phase is to use the bidders in STAT in order to find prices for the second-price auction with reserve price, and the fixed-price auction. Both auctions will be conducted in the next phases. To ensure incentive compatibility, a bidder should have no influence on the price of the bundle he is offered. This is why the prices of bundles offered to bidders in SEC-PRICE and FIXED in the following phases will be determined using bidders in STAT only. The bidders in STAT never get any items and so have no incentive to misreport their preferences.

Finding the price of the fixed-price auction is mechanism specific. However, the reserve price for the second price auction is generally determined by applying an approximation algorithm to bidders from STAT. If no small groups of bidders contributes a large fraction of the optimal solution (the first case), we can prove that with high probability the reserve price we obtain is a good approximation to the optimal welfare. On the other hand, If there is one bidder with very high value for the bundle of all items (the second case), we will see that this reserve price has no effect on the result of the second-price auction.

`Phase III: A Second-Price Auction`

We now conduct a second-price auction with a reserve price for selling the bundle of all items to one of the bidders. Intuitively, one can think of this phase as handling the first case, where there is one bidder with a very high value for the bundle of all items. A second-price auction will allocate the bundle of all items to the bidder that values it the most. If there is one bidder with a very high value for this bundle, he will be placed in SEC-PRICE with probability $1-\epsilon$. We then get a good approximation to the welfare, and the algorithm terminates.

The purpose of the reserve price is to handle the second case, where no small group of bidders contributes a lot to the optimal solution[5]. If this is the situation, allocating all items to one bidder may provide a bad approximation. Fortunately, in the previous phase we obtained a reserve price which is a good approximation to the optimal welfare. Therefore, if there is a "winning bidder", we know that we have a good approximation because the revenue obtained in the second-price auction (which is at least the reserve price) is a lower bound on the welfare. If we do not have a winning bidder, we continue to the next phase.

`Phase IV: A Fixed-Price Auction`

We go over the bidders in FIXED one by one, in some arbitrary order, asking each one for his demand under a fixed price per item, obtained earlier from the bidders in STAT. We allocate each bidder his most demanded set, and charge him the appropriate price. We remove the set allocated to him from the set of items that are offered to the next bidders.

This phase is meant to handle the second case, where no small group of bidders contributes a lot to the optimal solution. Indeed, it can be shown that since FIXED is a randomly chosen group that consists of a constant fraction of all bidders, it also holds, with high probability, a constant fraction of the optimal welfare. In addition, we show that in the second case the bidders in STAT aid us in choosing a fixed-price that leads to a good approximation. The way this price is chosen is mechanism-specific, and is not the same in our two mechanisms.

For every possible tosses of coins the framework produces a truthful deterministic mechanism. First, bidders who are in STAT never get any items, and thus have no incentive to misreport their preferences. A bidder can get items in exactly one of the following ways: by participating in the second-price auction with the reserve price, or by participating in the fixed-price auction.

It is well known that second-price auctions with a reserve price are incentive compatible. The fixed-price auction is also clearly incentive compatible, as each bidder gets the bundle that maximizes his demand, given prices which he does not affect.

4 Combinatorial Auctions with General Valuations

In this section we exhibit an incentive-compatible mechanism for approximating combinatorial auctions with general valuations. The incentive compatibility of the mechanism is guaranteed by its use of the framework. As in all mechanisms built using the framework, the main difficulty is to

[5]Of course, both a second-price auction and a second-price auction with a reserve price are incentive-compatible.

analyze the case in which no "small" group of bidders contributes "a lot" to the optimal solution. In the case of general valuations this is translated to the case where no bidder assigns a value to M that is higher than the $\sqrt{m}$-fraction of the value of the optimal fractional solution.

In this case, our mechanism uses the bidders of STAT to approximate the value of the optimal fractional solution. We set the item price for the fixed-price auction to be (approximately) the value of the approximation we obtained, divided by the number of items. The important technical observation is that for each item we manage to sell at this price, we "lose" a value of at most $O(\sqrt{m})$ times this price (compared to the optimal fractional solution). The revenue we get in this case sets a lower bound on the welfare we achieve.

Although the mechanism does use the LP relaxation of the problem, LP does play a relatively minor role, and we mainly use it for the analysis. This is in contrast to previous related work [15, 8], where the technique itself is LP based. The reader is referred to the appendix for the standard LP relaxation of the problem.

The Algorithm:

Input: n bidders[6], each with a general valuation v_i that is represented by a demand oracle, a rational number $0 < \epsilon < 1$.

Output: An allocation of the items, which is a $O(\frac{\sqrt{m}}{\epsilon^3})$-approximation to the optimal allocation.

The Algorithm:

`Phase I: Partitioning the Bidders`

1. Assign each bidder to exactly one of the following three sets: SEC-PRICE with probability $1-\epsilon$, FIXED with probability $\frac{\epsilon}{2}$, and STAT with probability $\frac{\epsilon}{2}$.

`Phase II: Gathering Statistics`

2. Calculate the value of the optimal fractional solution in the combinatorial auction with all m items, but only with the bidders in STAT. Denote this value by OPT^*_{STAT}.

`Phase III: A Second-Price Auction`

3. Conduct a second-price auction with a reserve price of $\frac{OPT^*_{STAT}}{\sqrt{m}}$, in which the bundle M of all items is sold to the bidders in SEC-PRICE. If there is a "winning bidder", allocate all the items to that bidder and output this allocation. Otherwise, proceed to the next step.

`Phase IV: A Fixed-Price Auction`

4. Let $R = M$. Let $p = \frac{\epsilon OPT^*_{STAT}}{8m}$.
5. For each bidder $i \in FIXED$, in some arbitrary order:
   (a) Let S_i be the demand of bidder i given the following prices: p for each item in R, and ∞ for each item in $M - R$.
   (b) Allocate S_i to bidder i, and set his price to be $p \cdot |S_i|$.
   (c) Let $R = R \setminus S_i$.

Theorem 4.1 *For any constant $\epsilon > 0$, there exists a randomized and truthful polynomial-time mechanism that achieves an $O(\frac{\sqrt{m}}{\epsilon^3})$-approximation with probability $1-\epsilon$.*

[6] Both in this mechanism and in the XOS mechanism, we assume that n is not constant. If n is constant, one can easily get a truthful $\frac{1}{n}$-approximation by bundling all items together and performing a second price auction.

Proof: The algorithm produces a feasible allocation. In addition, the algorithm is clearly incentive compatible, since it was designed using the framework. It is left to prove that it obtains the desired approximation ratio with probability $1-\epsilon$.

Denote by OPT^* the optimal fractional solution. There are two possible cases:

1. There is a bidder i such that $v_i(M) \geq \frac{OPT^*}{\sqrt{m}}$.

2. For each bidder i, $v_i(M) < \frac{OPT^*}{\sqrt{m}}$.

We start by handling the first case. Let i be some bidder such that $v_i(M) \geq \frac{OPT^*}{\sqrt{m}}$. Observe that with probability $1-\epsilon$ bidder i is in SEC-PRICE. If there is no such bidder is in SEC-PRICE, then the algorithm fail to guarantee any approximation ratio. This happens with probability of at most ϵ. The next proposition shows that if bidder i is in SEC-PRICE the algorithm obtains an $O(\sqrt{m})$-approximation.

Proposition 4.2 *If there exists a bidder i in SEC-PRICE such that $v_i(M) \geq \frac{OPT^*}{\sqrt{m}}$, then the allocation generated by the algorithm is an $O(\sqrt{m})$-approximation to the optimal allocation.*

Proof: Let i' be the bidder in SEC-PRICE with the highest value for M. By the conditions of the lemma, $v_{i'}(M) \geq \frac{OPT^*}{\sqrt{m}}$. Clearly, since $STAT \subseteq N$, we have that:

$$\frac{OPT^*_{STAT}}{\sqrt{m}} \leq \frac{OPT^*}{\sqrt{m}} \leq v_{i'}(M)$$

Hence, due to the properties of second-price auctions with a reserve price, all items will be sold to i', and the algorithm will terminate in Step 3. Thus, we get an allocation that is an $O(\sqrt{m})$-approximation to the optimal one. □

The second case is more involved. For each bidder i, $v_i(M) < \frac{OPT^*}{\sqrt{m}}$. We will take advantage of the fact that no bidder contributes "a lot" to the optimal fractional solution, and see that, with high probability, OPT^*_{STAT} is a good approximation to the optimal fractional solution. We will see that the same holds for OPT^*_{FIXED}, which is the value of the optimal fractional solutions in the combinatorial auctions with all m items and with bidders from FIXED only.

Lemma 4.3 *If for each bidder i, $v_i(M) < \frac{OPT^*}{\sqrt{m}}$, then with probability $1-o(1)$:*

1. $\frac{\epsilon}{4} \cdot OPT^* \leq OPT^*_{STAT}$

2. $\frac{\epsilon}{4} \cdot OPT^* \leq OPT^*_{FIXED}$

Proof: We will start by proving that the probability that the first event does not occur is $o(1)$. The proof for the second is almost identical. The lemma will then follow, by applying the union bound.

Let A be the random variable that receives the value of OPT^*_{STAT}. For every bidder i we denote by A_i the random variable that receives the value of bidder i in OPT^*_{STAT}. Let $\{x_{i,S}\}_{1 \leq i \leq n, S \subseteq M}$ be the set of variables in the fractional solution, OPT^*. Since every bidder is placed in STAT with probability $\frac{\epsilon}{2}$, and $STAT \subseteq N$, we have that $E[A] = \Sigma_i \frac{\epsilon}{2} E[A_i] \geq \Sigma_i \frac{\epsilon}{2} \Sigma_S x_{i,S} v_i(S) = \frac{\epsilon}{2} OPT^*$. If the conditions of the lemma hold, we also have that for each i, $A_i < \frac{OPT^*}{\sqrt{m}}$. We can use this fact to set an upper bound on the probability that A gets a value that is substantially smaller than its expectation. We make use of the following corollary from Chebyshev's inequality:

Claim 4.4 *Let X be the sum of independent random variables, each of which lies in $[0, t]$. Then, for any $\alpha > 0$,* $\Pr[|X - E[X]| \geq \alpha] \leq \frac{tE[X]}{\alpha^2}$.

Since for each i, $A_i \in [0, \frac{OPT^*}{\sqrt{m}}]$, we have that

$$\Pr[A < \frac{\epsilon}{4} \cdot OPT^*] \leq \Pr[|A - \frac{\epsilon}{2} \cdot OPT^*| \geq \frac{\epsilon}{4} \cdot OPT^*]$$

$$\leq \frac{\frac{OPT^*}{\sqrt{m}} \cdot \frac{\epsilon}{2} \cdot OPT^*}{(\frac{\epsilon}{4} OPT^*)^2} \leq \frac{8}{\epsilon\sqrt{m}}$$

□

With probability of $1 - o(1)$ we have that the values of the optimal fractional solutions for FIXED and STAT are "close" to OPT^*. If this is the case, we will show that we manage to achieve an $O(\frac{\sqrt{m}}{\epsilon^2})$ approximation factor. With probability of at most $o(1)$ this is not the case, and the algorithm fails to provide any approximation ratio.

Although the second-price auction was designed to handle the first case, when there is one bidder that contributes "a lot" to the welfare, it is still possible that some bidder i in SEC-PRICE will be allocated the bundle M in Step 3. However, notice that bidder i was forced to pay at least $\frac{OPT^*_{STAT}}{\sqrt{m}}$. Therefore, that bidder's value for the bundle M is greater than $\frac{OPT^*_{STAT}}{\sqrt{m}}$, which by Lemma 4.3 is at least $\frac{\epsilon OPT^*}{4\sqrt{m}}$. Hence, allocating the bundle M to bidder i provides an $O(\frac{\sqrt{m}}{\epsilon})$ approximation to the optimal solution.

If no bidder in SEC-PRICE got the bundle M then the algorithm attempts to sell items to the bidders in FIXED (Step 5). As before, we claim that the revenue is a lower bound to the social welfare. The next lemma shows that in this case the revenue will be $\Omega(\frac{\epsilon^3 OPT^*}{\sqrt{m}})$. Hence, Step 5 will result in an allocation that is a $\Omega(\frac{\sqrt{m}}{\epsilon^3})$-approximation to the optimal allocation.

Lemma 4.5 *If the following conditions hold:*

1. *The algorithm reaches Step 5*
2. *For each bidder i, $v_i(M) < \frac{OPT^*}{\sqrt{m}}$*
3. *For the item-price p it holds that: $\frac{\epsilon^2 OPT^*}{16m} \leq p \leq \epsilon\frac{OPT^*}{8m}$*
4. *$OPT^*_{FIXED} \geq \frac{\epsilon}{4} \cdot OPT^*$*

then the revenue of the algorithm is $\Omega(\frac{\epsilon^3 OPT^*}{\sqrt{m}})$.

Proof: Let $\{y_{i,S}\}_{i \in FIXED, S \subseteq M}$ be the variables in the fractional solution OPT^*_{FIXED}. We will restrict our attention to bundles in OPT^*_{FIXED} that are profitable when setting a price of p for each item. That is, let T be the set of pairs (i, S) such that $y_{i,S} > 0$, and $v_i(S) - p \cdot |S| > 0$. The next claim shows that we do not lose too much by ignoring all other bundles in OPT^*_{FIXED}.

Claim 4.6 $\Sigma_{(i,S) \in T}\, y_{i,S}\, v_i(S) \geq \frac{1}{2} \cdot OPT^*_{FIXED}$

Proof: Define $\overline{T}$ to be the "complement" set of T. Formally, $\overline{T}$ consists of all pairs (i, S) such that $y_{i,S} > 0$ in OPT^*_{FIXED}, but $v_i(S) - p \cdot |S| \leq 0$. It is easy to see that $OPT^*_{FIXED} =$

$\Sigma_{(i,S)\in T}\ y_{i,S}v_i(S) + \Sigma_{(i,S)\in \overline{T}} y_{i,S}v_i(S)$. Since $OPT^*_{FIXED} \geq \frac{\epsilon}{4} \cdot OPT^*$ it is enough to bound from above the contribution of $\overline{T}$ to OPT^*_{FIXED} to prove the claim.

$$\Sigma_{(i,S)\in \overline{T}} y_{i,S}v_i(S) \leq \Sigma_{(i,S)\in \overline{T}}\ y_{i,S}\ p \cdot |S| \leq m \cdot p$$

$$\leq m \cdot \frac{\epsilon \cdot OPT^*}{8m} \leq \frac{OPT^*_{FIXED}}{2}$$

where the first inequality is because of the definition of $\overline{T}$ and the second inequality is due to the LP constraints. $\square$

Let us now calculate the revenue we get in Step 5. Without loss of generality, assume the bidders in FIXED are $1, ..., \frac{\epsilon}{2}n$. In the first iteration of Step 5, bidder 1 is asked for his most demanded set. The key observation is that if there is some S such that $x_{1,S} > 0$ and $(1, S) \in T$ then bidder 1's demand is not empty. Recall that for each item in S_1 we gain a revenue of p.

We will now upper bound what we "lose" by assigning S_1 to bidder 1 in comparison to OPT^*_{FIXED}. Notice, that by assigning S_1 to bidder 1 we lose both the value of all the fractional bundles assigned to bidder 1 in OPT^*_{FIXED}, and of all the bundles in OPT^*_{FIXED} that contain some item from S_1. The value of all the fractional bundles assigned to bidder 1 in OPT^*_{FIXED} is at most $\frac{OPT^*}{\sqrt{m}}$:

$$\Sigma_{(1,S)\in T}\ y_{1,S}v_1(S) \leq \frac{OPT^*}{\sqrt{m}}$$

because $v_1(M) < \frac{OPT^*}{\sqrt{m}}$ and $\Sigma_{(1,S)} y_{1,S} \leq 1$, due to the constraints of the LP formulation.

We will now bound the value of all the bundles in OPT^*_{FIXED} that contain some item from S_1. Fix some item $j \in S_1$. Again, using the constraints of the LP and $v_i(M) < \frac{OPT^*}{\sqrt{m}}$,

$$\Sigma_{(i,S)\in T | j \in S} y_{i,S}v_i(S) \leq \frac{OPT^*}{\sqrt{m}}$$

To conclude, for every item we sell to bidder 1 at price $p \geq \epsilon^2 \cdot \frac{OPT^*}{16m}$, we lose bundles in T that are together worth at most $2 \cdot \frac{OPT^*}{\sqrt{m}}$. The analysis continues by removing from OPT^*_{FIXED} all pairs (i, S) which can not be assigned now (either $i = 1$, or $j \in S_i$ and $j \in S$), and applying similar arguments to the rest of the bidders in FIXED.

The revenue achieved by the algorithm is an $O(\frac{\sqrt{m}}{\epsilon^2})$-approximation to the value of OPT^*_{FIXED}. Since $OPT^*_{FIXED} \geq \frac{\epsilon}{4} \cdot OPT^*$ we have that it is a $O(\frac{\sqrt{m}}{\epsilon^3})$ approximation to OPT^*, and the theorem follows.

$\square$

$\square$

5 Combinatorial Auctions with XOS Valuations

Like the mechanism for approximating combinatorial auctions with general valuations, the mechanism for XOS valuations is also based on the general framework. Again, the main challenge involved in designing this mechanism is analyzing the case in which no "small" group of bidders contributes "a lot" to the optimal solution. The way this is achieved for XOS valuations is entirely different from the way it is done with general valuations.

Suppose we assign a bundle S to a bidder with an XOS valuation v_i. By the definition of XOS valuations, $v(S)$ is determined by the value S gets under some additive valuation a (we will also refer to a as the maximizing clause). We can look at the whole process as implicitly assigning a "price" to each item in S (the price that a assigns to each item in S.) We will use this property for finding a price for the fixed-price auction. To see how this is done we must first introduce the following definition:

Definition 5.1 *We say that an allocation of the items $T = (T_1, ..., T_n)$ is* supported by a price p, *if for each bidder i and each possible bundle $S_i \subseteq T_i$, it holds that $v_i(S_i) \geq |S_i| \cdot p$. We call $\Sigma_i |T_i| \cdot p$ the* supported value *of T.*

We now show that for every allocation it is possible to find a "contained" allocation and a price that supports it, and holds a considerable part of the welfare of the original allocation.

Lemma 5.2 *For every allocation $T = (T_1, ..., T_n)$ it is possible to find in polynomial time an allocation $(S_1, ..., S_n)$ and a price p that supports it, such that for each i, $S_i \subseteq T_i$, and $\Sigma_i |S_i| \cdot p \geq \Omega(\frac{\Sigma_i v_i(T_i)}{\log m})$.*

Proof: Given an allocation T, we query each bidder i's XOS oracle for the maximizing XOS clause for T_i. We refer to the value of an item in T_i as the item's value in the maximizing clause of T_i. Let $W = \Sigma_i v_i(T_i)$ (i.e., the welfare value of T.) Define the set $P = \{\frac{W}{2m}, \frac{W}{m}, ..., \frac{W}{2}, W\}$. Notice that $|P| = O(\log m)$.

Round down each item's value in the maximizing clauses to the nearest value in P. Let $p \in P$ be the (rounded down) item value that "contributes the most" to the welfare. Notice that we ignore items with value lower than $\frac{W}{2m}$ – our "loss" is not too high since the sum of these items' values is less than $\frac{W}{2}$. We can now define $(S_1, ..., S_n)$ to be the allocation in which $S_i \subseteq T_i$ and the (rounded down) value of every item in T_i is at least p. □

There is still the matter of finding such a price that would enable us to get a good approximation in the fixed-price auction. We prove that one can use the bidders in STAT to find such a price for the bidders in FIXED with high probability.

We also note that if a valuation is known to be submodular, an XOS oracle for it can be simulated using a demand oracle [6]. Thus, if all bidders are known to be submodular our mechanism can be implemented using demand oracles only.

The Algorithm:

Input: n bidders, $v_1, ..., v_n$, each represented by a demand and a XOS oracle, a rational number $0 < \epsilon < \frac{1}{2}$.

Output: An allocation of the items, which is an $O(\frac{\log^2 m}{\epsilon^3})$-approximation to the optimal allocation.

The Algorithm:

`Phase I: Partitioning the Bidders`

1. Assign each bidder to exactly one of the following three sets: SEC-PRICE with probability $1 - \epsilon$, FIXED with probability $\frac{\epsilon}{2}$, and STAT with probability $\frac{\epsilon}{2}$.

`Phase II: Gathering Statistics`

2. Find an allocation that is an $O(1)$ approximation to the value of the optimal solution in the combinatorial auction with all m items, but only with the bidders in STAT (e.g., using the algorithms of [6, 7]). Denote this value by OPT_{STAT}.

3. Using the allocation obtained in the previous step, find a price p' and an allocation $T = (T_1, ..., T_{|STAT|})$, such that T is supported by p' (rounded down to the nearest power of 2), and $\Sigma_{i \in STAT} |T_i| p' \geq \Omega(\frac{OPT_{STAT}}{\log m})$.

`Phase III: A Second-Price Auction`

4. Conduct a second-price auction with a reserve price of $\frac{\epsilon^2}{100}\frac{OPT_{STAT}}{\log^2 m}$, in which the bundle M of all items is sold to the bidders in SEC-PRICE. If there is a "winning bidder", allocate all the items to that bidder and output this allocation. Otherwise, proceed to the next step.

`Phase IV: A Fixed-Price Auction`

5. Let $R = M$. Let $p = p'/2$.
6. For each bidder $i \in FIXED$, in some arbitrary order:
   (a) Let S_i be the demand of bidder i given the following prices: p for each item in R, and ∞ for each item in $M - R$.
   (b) Allocate S_i to bidder i, and set his price to be $p \cdot |S_i|$.
   (c) Let $R = R \setminus S_i$.

Theorem 5.3 *For any constant $0 < \epsilon < \frac{1}{2}$, there exists a randomized and truthful algorithm that achieves an $O(\frac{\log^2 m}{\epsilon^3})$-approximation with probability $1 - \epsilon$.*

Proof: The algorithm produces a feasible allocation. Incentive compatibility of the algorithm is guaranteed since it was built using the framework. It is left to prove that it obtains the desired approximation ratio with probability $1 - \epsilon$.

We will now prove that the the algorithm provides the approximation ratio. Let $R = \frac{\epsilon^2}{100}\frac{OPT}{\log^2 m}$. There are two possible cases:

1. There is a bidder i such that $v_i(M) \geq R$.
2. For each bidder i, $v_i(M) < R$.

We handle the first case in a way similar to the way we handled the first case in the correctness proof for the algorithm of Section 4. Let i be some bidder such that $v_i(M) \geq R$. Observe that with probability $1 - \epsilon$ bidder i is in SEC-PRICE. If there is no such bidder is in SEC-PRICE, then the algorithm fails to guarantee any approximation ratio. This happens with probability of at most ϵ. If bidder i is in SEC-PRICE the algorithm obtains a R-approximation. The next proof is similar to the proof of Lemma 4.2.

Proposition 5.4 *If there exists a bidder i in SEC-PRICE such that $v_i(M) \geq R$, then the allocation generated by the algorithm is a $O(R)$-approximation to the optimal allocation.*

Let us now examine the second case, where for each bidder i, $v_i(M) < R$. The basic idea is to use the bidders in STAT to find a price that, with high probability, will obtain an allocation of the items to the bidders in FIXED in Step 6.

To show this we need to prove that OPT_{FIXED}, the value of the optimal solution consisting of the bidders in FIXED only, has a value that is "close" to the value of the total welfare. This will be done in a similar way to the previous algorithm. However, unlike the previous algorithm, we have to prove that if a price is "good" in OPT_{FIXED} (i.e. supports an allocation that holds a substantial part of the welfare), then it can be found using the bidders in STAT. As in Lemma 5.2, we restrict our attention to prices which are greater than $\frac{OPT}{2m \log m}$.

Lemma 5.5 *If for each bidder i, $v_i(M) < R$, then with probability higher than $1 - 2\epsilon^2$:*

1. $\frac{\epsilon}{4} \cdot OPT \leq OPT_{STAT}$

2. $\frac{\epsilon}{4} \cdot OPT \leq OPT_{FIXED}$

3. *Let* $P = \{p | p$ *is a power of* 2*, and* $\frac{OPT}{2m \log m} \leq p \leq \frac{OPT}{\log m}$*, and there exists an allocation* T *that is supported by* p*, and the supported value of* T *is at least* $\frac{\epsilon}{4} \cdot \frac{OPT}{\log m}\}$*. Then, for every* $p_k \in P$ *there exists an allocation* T_k *of the items to the bidders in FIXED only such that* T_k *is supported by* p_k*, and the supported value of* T_k *is at least* $\frac{\epsilon^2}{16} \cdot \frac{OPT}{\log m}$*.*

Proof: The proof that the probability that one of the first two events does not occur is $o(1)$ is identical to that of Lemma 4.3. We now bound from above the probability that the third event occurs and use the union bound to complete the proof.

Let $T = (T_1, ..., T_n)$ be an allocation, and $p_k \in P$ a price such that the supported value of T is at least $\frac{\epsilon}{4} \cdot \frac{OPT}{\log m}$, and T is supported by p_k. We now turn our attention to the bidders in FIXED. Observe that for each bidder $i \in FIXED$, $v_i(T_i) \geq |T_i| \cdot p_t$. Therefore, we will prove that there exists a T_k with the desired value by looking at the expected value of T, restricted only to bidders in FIXED.

Let A_i be the random variable that gets the value of $p \cdot |T_i|$ with probability $\frac{\epsilon}{2}$, and 0 with probability $1 - \frac{\epsilon}{2}$. Let $A = \Sigma_i A_i$. Since every bidder i is placed in FIXED with probability $\frac{\epsilon}{2}$ we have that $E[A] = \Sigma_i E[A_i] = \frac{\epsilon}{2} \Sigma_i p \cdot |T_i| \geq \frac{\epsilon}{4} \cdot \frac{OPT}{\log m}\}$. Using Claim 4.4, and since for each i, $A_i \in [0, R]$, we have that

$$\Pr[A < \frac{\epsilon^2 OPT}{16 \log m}] \leq \Pr[|A - \frac{\epsilon^2 OPT}{8 \log m}| \geq \frac{\epsilon^2 OPT}{16 \log m}]$$

$$\leq \frac{R \cdot \frac{\epsilon^2 OPT}{8 \log m}}{(\frac{\epsilon^2 OPT}{16 \log m})^2} \leq \frac{32 R \cdot \log m}{\epsilon^2 \cdot OPT}$$

Since there are less than $\log m$ possible choices of p_k, we can apply the union bound to verify that the fourth event does not occur with probability $\frac{32 R \cdot \log m \log m}{\epsilon^2 \cdot OPT}$. By our choice of R, we get that they all hold simultaneously with probability of at least $1 - 2\epsilon^2$. □

Given that the conditions of Lemma 5.5 hold, we will show that we manage to achieve an $O(\frac{\log m}{\epsilon^3})$ approximation factor. With probability of at most $2\epsilon^2$ this is not the case, and the algorithm fails to provide any approximation ratio.

If some bidder i in SEC-PRICE was allocated M in Step 4, then he was forced to pay at least $\frac{\epsilon^2}{100} \frac{OPT_{STAT}}{\log^2 m}$. Therefore, that bidder's value for M is greater than $\frac{\epsilon^2}{100} \frac{OPT_{STAT}}{\log^2 m}$, which by Lemma 5.5 is at least $O(\frac{\epsilon^3 OPT}{\log^2 m})$. Hence, allocating M to bidder i provides a a $O(\frac{\log^2 m}{\epsilon^3})$ approximation to the optimal solution.

If no bidder in SEC-PRICE got the bundle M then the algorithm attempts to sell items to the bidders in FIXED (Step 6). The next two lemmas show that in this case we will get an allocation that is an $O(\frac{\log m}{\epsilon^3})$-approximation to the optimal allocation.

Lemma 5.6 *Let* $T_p = (T_1, ..., T_n)$ *be an allocation that maximizes* $\Sigma_i v_i(T_i)$ *such that*

1. T_p *is supported by* p*.*

2. *For each bidder* $i \notin FIXED$*,* $T_i = \emptyset$*.*

Then, if the algorithm reaches Step 6 the approximation ratio achieved is $O(\Sigma_i |T_i| \cdot p)$*.*

Proof: We first note that by assigning T_i to each bidder i and charging a price of $|T_i| \cdot p$, we gain a revenue of $\Sigma_i |T_i| \cdot p$, while all bidders are profitable. We will use this revenue as a lower bound to the welfare that can be achieved. Notice that we do not guarantee that the actual revenue the mechanism gets is a constant factor away from $\Sigma_i |T_i| \cdot p$.

We will now upper bound the revenue we lose by assigning S_1 to bidder 1, comparing to the allocation considered earlier. Without loss of generality, assume the bidders in FIXED are numbered $1, ..., \frac{\epsilon}{2} n$. In the first iteration of Step 5, bidder 1 is asked for his most demanded set. First, we could have assigned T_1 to bidder 1 and gain a revenue of $|T_1| \cdot \frac{p}{2}$. (Recall that the the price for item is $\frac{p}{2}$.) However, we did not lose too much because the value of T_1 is at most twice the value of S_1. The last statement is true since bidder 1 could gain a profit of at least $|T_1| \cdot \frac{p}{2}$ by choosing T_1, and S_1 has at least that value being bidder 1's most demanded set. We note again that the revenue we achieve in this case (but not the welfare) might be very small comparing to $v_i(T_i)$.

The second possible lose occurs when there is an item $j \in S_1$, and there exists another bidder i' with $j \in T_{i'}$. Because T_p is supported by p, we have that $v_{i'}(T_{i'} \backslash \{j\}) \geq (|T_{i'}| - 1) \cdot p$. Summing over all such items, we have that we lose a value of at most $|S_1| \cdot \frac{p}{2} \leq v_1(S_1)$. The inequality holds since S_1 is profitable to bidder 1 under a price per item of $\frac{p}{2}$.

To conclude, by assigning T_1 to bidder 1 we lose a revenue of $O(T_1)$. The analysis continues by removing from $T_2, ..., T_{\frac{\epsilon}{2}}$ all items which can not be assigned now, and using induction to apply similar arguments to the rest of the bidders in FIXED. $\square$

Lemma 5.7 *If the following conditions hold:*

1. *The algorithm reaches Step 6*
2. $\frac{\epsilon}{4} \cdot OPT \leq OPT_{STAT}$
3. $\frac{\epsilon}{4} \cdot OPT \leq OPT_{FIXED}$
4. *Let* $P = \{p | p$ *is a power of* 2, *and* $\frac{OPT}{2m \log m} \leq p \leq \frac{OPT}{\log m}$, *and there exists an allocation* T *that is supported by* p, *and the supported value of* T *is at least* $\frac{\epsilon}{4} \cdot \frac{OPT}{\log m}\}$. *Then, for every* $p_k \in P$ *there exists an allocation* T_k *of the items to the bidders in FIXED only such that* T_k *is supported by* p_k, *and the supported value of* T_k *is at least* $\frac{\epsilon^2}{16} \cdot \frac{OPT}{\log m}$.

Then the algorithm produces an allocation that is an $O(\frac{\log m}{\epsilon^3})$*-approximation to the welfare.*

Proof: Observe that in Step 3 we have found an allocation that is supported by p and worth more than $\frac{OPT_{STAT}}{\log m} \geq \frac{\epsilon OPT}{4 \log m}$. Obviously, an allocation restricted to bidders in STAT only is also an allocation for all bidders with the same value. We can therefore deduce that there exists an allocation T_p of the items to bidders in FIXED such that $T_p = (T_1, ..., T_n)$ is supported by p, and worth at least $\frac{\epsilon^2}{16} \cdot \frac{OPT}{\log m}$.

Clearly, all conditions of Lemma 5.6 hold. Therefore, the algorithm is an $O(\frac{\log m}{\epsilon^2})$-approximation to the value of OPT_{FIXED}. Since $OPT_{FIXED} \geq \frac{\epsilon}{4} \cdot OPT$ we have that it is an $O(\frac{\log m}{\epsilon^3})$ approximation to OPT. $\square$

$\square$

Acknowledgments

We thank Moshe Babaioff, Liad Blumrosen, Uri Feige, Ron Lavi, Ahuva Mu'alem, and Chaitanya Swamy for helpful discussions and comments. The authors are supported by grants from the Israel Science Foundation and the USA-Israel Bi-national Science Foundation.

References

[1] A. Archer, C. Papadimitriou, K. Talwar, and E. Tardos. An approximate truthful mechanism for combinatorial auctions with single parameter agent. In *Proceedings of the 14th Annual ACM Symposium on Discrete Algorithms (SODA)*, 2003.

[2] Yair Bartal, Rica Gonen, and Noam Nisan. Incentive compatible multi unit combinatorial auctions. In *TARK 03*, 2003.

[3] Liad Blumrosen and Noam Nisan. On the computational power of iterative auctions I: Demand queries. Working Paper. Preliminary version in EC 2005.

[4] E. H. Clarke. Multipart pricing of public goods. *Public Choice*, 2:19–33, 1971.

[5] P. Cramton, Y. Shoham, and R. Steinberg (Editors). *Combinatorial Auctions*. MIT Press. Forthcoming., 2005. http://www.cramton.umd.edu/papers2000-2004/cramton-shoham-steinberg-combinatorial-auctions.pdf.

[6] Shahar Dobzinski, Noam Nisan, and Michael Schapira. Approximation algorithms for combinatorial auctions with complement-free bidders. In *STOC '05: Proceedings of the thirty-seventh annual ACM symposium on Theory of computing*, pages 610–618, New York, NY, USA, 2005. ACM Press.

[7] Shahar Dobzinski and Michael Schapira. An improved approximation algorithm for combinatorial auctions with submodular bidders. In SODA 2006.

[8] Uriel Feige. On maximizing welfare where the utility functions are subadditive. To appear in STOC 2006.

[9] Amos Fiat, Andrew V. Goldberg, Jason D. Hartline, and Anna R. Karlin. Competitive generalized auctions. In *STOC*, pages 72–81, 2002.

[10] Andrew Goldberg, Jason Hartline, Anna Karlin, Mike Saks, and Andrew Wright. Competitive auctions. Submitted.

[11] T. Groves. Incentives in teams. *Econometrica*, pages 617–631, 1973.

[12] Ron Holzman, Noa Kfir-Dahav, Dov Monderer, and Moshe Tennenholtz. Bundling equilibrium in combinatrial auctions. *Games and Economic Behavior*, 47:104–123, 2004.

[13] Subhash Khot, Richard Lipton, Evangelos Markakis, and Aranyak Mehta. Inapproximability results for combinatorial auctions with submodular utility functions. To appear in WINE 2005.

[14] Ron Lavi, Ahuva Mu'alem, and Noam Nisan. Towards a characterization of truthful combinatorial auctions. In *FOCS '03: Proceedings of the 44th Annual IEEE Symposium on Foundations of Computer Science*, page 574, Washington, DC, USA, 2003. IEEE Computer Society.

[15] Ron Lavi and Chaitanya Swamy. Truthful and near-optimal mechanism design via linear programming. In FOCS 2005.

[16] Benny Lehmann, Daniel Lehmann, and Noam Nisan. Combinatorial auctions with decreasing marginal utilities. In *ACM conference on electronic commerce*, 2001.

[17] Daniel Lehmann, Liadan Ita O'Callaghan, and Yoav Shoham. Truth revelation in approximately efficient combinatorial auctions. In *JACM 49(5)*, pages 577–602, Sept. 2002.

[18] A. Mas-Collel, W. Whinston, and J. Green. *Microeconomic Theory.* Oxford university press, 1995.

[19] Ahuva Mua'lem and Noam Nisam. Truthful approximation mechanisms for restricted combinatorial auctions. In *AAAI-02*, 2002.

[20] Noam Nisan and Amir Ronen. Algorithmic mechanism design. In *STOC*, 1999.

[21] Noam Nisan and Amir Ronen. Computationally feasible vcg-based mechanisms. In *ACM Conference on Electronic Commerce*, 2000.

[22] Noam Nisan and Ilya Segal. The communication requirements of efficient allocations and supporting prices. To appear in Journal of Economic Theory.

[23] M. J. Osborne and A. Rubistein. *A Course in Game Theory.* MIT press, 1994.

[24] Tuomas Sandholm. Algorithm for optimal winner determination in combinatorial auctions. *Artificial Intelligence*, 135(1-2):1–54, 2002.

[25] William Vickrey. Counterspeculation, auctions and competitive sealed tenders. *Journal of Finance*, 16(1):8–37, March 1961.

A The Standard LP Formulation of a Combinatorial Auction

Maximize: $\Sigma_{i,S} x_{i,S} v_i(S)$
Subject to:

- For each item j: $\Sigma_{i,S|j \in S} x_{i,S} \leq 1$
- for each bidder i: $\Sigma_S x_{i,S} \leq 1$
- for each i, S: $x_{i,S} \geq 0$

We remark that the LP relaxation can be solved using demand oracles only [22].