

Truthfulness via Proxies

Shahar Dobzinski
Department of Computer Science
Cornell University
shahar@cs.cornell.edu

Hu Fu
Department of Computer Science
Cornell University
hufu@cs.cornell.edu

Robert Kleinberg
Department of Computer Science
Cornell University
rdk@cs.cornell.edu

February 3, 2011

Abstract

This short note exhibits a truthful-in-expectation $O(\frac{\log m}{\log \log m})$ -approximation mechanism for combinatorial auctions with subadditive bidders that uses polynomial communication.

1 Introduction

We consider the problem of maximizing the social welfare in combinatorial auctions with subadditive bidders. In this problem, we have a set M , $|M| = m$, of heterogeneous items, and n bidders. Each bidder i has a valuation function v_i , $v_i : 2^M \rightarrow \mathbb{R}$. We assume that each valuation v_i is normalized ($v_i(\emptyset) = 0$), non-decreasing, and subadditive: for any two bundles S and T , $v_i(S) + v_i(T) \geq v_i(S \cup T)$. This important class of valuations contains other interesting classes. In particular, it strictly contains the class of submodular valuations. Our goal is to find an allocation of the items (S_1, \dots, S_n) that maximizes the social welfare: $\sum_i v_i(S_i)$. The efficiency of our algorithms will be measured in terms of the natural parameters of the problem: n and m . Since the valuation functions are objects of exponential size, we assume that they are represented by black boxes that can answer any type of queries. In particular, we assume that bidders are computationally unbounded and measure only the amount of communication between them.

Feige [8] obtains a 2-approximation algorithm by applying an ingenious randomized rounding. This is the best possible [6] in polynomial communication. Much research was concerned with the design of truthful algorithms for the problem: algorithms in which a profit-maximizing strategy for each bid-

der is to truthfully answer the queries. The simplest and strongest notion considered is deterministic truthfulness, where no randomization is allowed. The VCG mechanism is truthful and optimally solves the problem, but is not computationally efficient. The best truthful approximation algorithm known achieves a poor approximation ratio of $O(\sqrt{m})$ [6], which is achieved by a maximal-in-range algorithm (see below). An evidence that deterministic algorithms cannot achieve an improved approximation ratio was given in [5]: maximal-in-range algorithms cannot achieve an approximation ratio better than $m^{\frac{1}{6}}$ with polynomial communication.

If randomization is allowed, there exists a universally truthful mechanism (a distribution over deterministic truthful mechanisms) that guarantees an approximation ratio of $O(\log m \log \log m)$ [3]. This paper relaxes the solution concept and considers mechanisms that are *truthful in expectation*: truth telling maximizes the *expected* profit of each bidder, where the expectation is taken over the internal random coins of the algorithm. We stress that truthfulness in expectation is much weaker than universal truthfulness: in particular, bidders that are not risk neutral may not be incentivized to bid truthfully. See [7] for a discussion.

In this note we show that there exists a truthful-in-expectation $O(\frac{\log m}{\log \log m})$ -approximation mechanism that uses only polynomial amount of communication. While this only slightly improves the best approximation ratio provided by universally truthful algorithms (at the cost of weakening the solution concept), we feel that this result is of interest for two main reasons (1) the random-sampling based techniques of [3] do not seem capable of achieving a better than a loga-

rithmic factor, so this result might hint at a performance gap between truthful-in-expectation and universally truthful mechanisms in our setting, and (2) our rounding technique is non-trivial and may be of independent interest.

Our algorithm is described in the communication complexity model. The model is best suited for proving strong, unconditional impossibility results. However, if we want algorithms designed in the model to be useful, they should not take advantage of the unlimited computational power of the bidders. Indeed, essentially all known algorithms known for combinatorial auctions need polynomially bounded computation power and use a restricted and natural communication form, e.g., demand queries (given prices per item p_1, \dots, p_m return a bundle that maximizes the profit). Unfortunately, we do not know how to implement our algorithm using demand queries, or any other type of natural query. In a sense, this is an abuse of the communication complexity model; hence we recommend viewing our result as demonstrating the falsehood of certain strong impossibility results for the problem, rather than as a “real” mechanism suitable for use in combinatorial auctions.

Let us briefly discuss the techniques we use. Almost all truthful deterministic mechanisms known in the literature are *maximal-in-range*, a scaled-down version of the classic VCG mechanism: for every set of valuations, select the welfare-maximizing allocation in a *predefined* set of allocations. However, as mentioned above, [5] shows that this method can only guarantee a poor approximation ratio in our setting. One way to overcome this obstacle was spelled out by [4] which suggested the use of *maximal in distributional range* (MIDR) mechanisms: for every set of valuations, select the welfare-maximizing distribution in a predefined set of *distributions over allocations*, then sample an allocation from this distribution. This results in a truthful-in-expectation mechanism. Furthermore, in [4] it is shown that for some problems MIDR mechanisms may be more powerful than any universally truthful mechanism.

Our mechanism is also MIDR¹. The basic idea is to represent each bidder by a proxy bidder, find the optimal fractional solution among these proxy bidders, and then round the fractional solution. Each proxy bidder is defined so that he “simulates” the expected value of the bundle after the randomized rounding. The main obstacle is to prove feasibility while still

¹In fact, our range will consist of distributions over *infeasible* allocations, but we will always output a distribution over feasible allocations with the same expected welfare as the best distribution in the range. Hence the mechanism is *equivalent* to MIDR (see [5]), and truthfulness in expectation follows.

being able to relate the value of the rounded solution (that was calculated with respect to the proxy bidders) to the original bidders.

We note that Lavi and Swamy [9] already implicitly used maximal-in-distributional-range algorithms together with the LP relaxation of the problem. Our solution requires a subtler rounding of the LP, and in particular overcomes one of the main barriers of [9]: their decomposition is based on an algorithm that “verifies” the integrality gap, and thus they can only provide a truthful-in-expectation $O(\sqrt{m})$ -approximation mechanism for our setting. Our mechanism uses a direct approach to “decompose” a linear program with proxy bidders. This is one of the main reasons for our success in guaranteeing a better approximation ratio.

The main question we leave open is the existence of a constant-ratio truthful-in-expectation mechanism for combinatorial auctions with subadditive bidders. A first step in this direction might be to prove that no MIDR mechanism can achieve an $O(1)$ approximation in polynomial time.

2 The Mechanism

2.1 The Range

We first remind the reader of the linear program relaxation of the problem:

Maximize: $\sum_{i,S} x_{i,S} v_i(S)$

Subject to:

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

Definition 2.1 A distribution D over (not necessarily feasible) allocations is called (c,p) -fractional if there exists a feasible fractional solution $\{x_{i,S}\}_{i,S}$ to the LP such that D is equal to the distribution produced by the following process: with probability p no bidder is allocated any item. With probability $1 - p$ each bidder i receives exactly one bundle S with probability $x_{i,S}$ and keeps each item $j \in S$ with probability c , independently at random.

Definition 2.2 (the range) The range $\mathcal{R}_{c,p}$ consists of all (c,p) -fractional distributions.

Before proving that there is one distribution in $\mathcal{R}_{c,p}$ that provides a good approximation, we require the following definition:

Definition 2.3 Given a valuation v , let the c -proxy valuation v' be:

$$v'(S) = E_{T \sim \mathcal{D}_S}[v(T)]$$

where \mathcal{D} is the distribution where each $j \in S$ is in T with probability c independently at random.

Lemma 2.4 The optimal distribution in a $\mathcal{R}_{c,p}$ range is an $O(\frac{1}{c \cdot p})$ -approximation to the optimal social welfare if the valuations are subadditive.

The proof of this lemma immediately follows by considering the optimal allocation and using the following proposition from [8]:

Proposition 2.5 (paraphrasing [8]) Let S be a bundle, v a subadditive valuation, and c such that $0 \leq c \leq 1$ and $\frac{1}{c}$ is an integer. Then, $v'(S) \geq c \cdot v(S)$, where v' is the c -proxy valuation of v .

2.2 The Algorithm

1. For each valuation v_i of bidder i , let v'_i be the c -proxy valuation.
2. Solve the linear program relaxation of the problem with respect to the c -proxy valuations v'_i .
3. Each bidder i is tentatively assigned exactly one bundle S_i , where bundle S is allocated to i with probability exactly $x_{i,S}$. If there is an item that is allocated more than $\frac{1}{c}$ times, the algorithm halts and no bidder is allocated any items. Otherwise, proceed to the next steps.
4. For each bidder i , let q_i be the probability that some item in S_i was allocated more than $\frac{1}{c} - 1$ times in the following random experiment: each bidder i' , $i' \neq i$, is allocated bundle S with probability $x_{i',S}$.
5. Independently for each item $j \in S$, select at most one bidder to receive j , so that each bidder that is tentatively allocated S_i where $j \in S_i$ receives item j with probability exactly c .
6. For each bidder i , with probability $1 - \frac{p}{1-q_i}$ he is not allocated any items, and with probability $\frac{p}{1-q_i}$ he keeps the items that he was assigned in the previous step.

We note that the last step of the algorithm is a re-implementation of the main idea behind the truthful-in-expectation mechanism for single-minded bidders of [1].

In our proofs we assume m is large enough (when m is a constant, VCG can be implemented in polynomial

time in n). Also, from now on we fix $c = \frac{\log \log m}{100 \log m}$ and $p = \frac{1}{20}$. We first have to make sure that the algorithm is correctly defined. For that we have to prove that $p \leq 1 - q_i$, for every q_i . This follows from the next claim [8, 6]:

Claim 2.6 ([8]) Fix any feasible solution of the linear program. Allocate each bidder i exactly one bundle where each bundle S is allocated with probability $x_{i,S}$. The probability that no item is allocated more than $\frac{1}{c}$ times is at least $1 - \frac{1}{m}$.

In other words, $1 - q_i \geq 1 - \frac{1}{m} \geq p$, as needed. We now show that the algorithm indeed finds the optimal distribution in the range and uses a polynomial amount of communication.

Lemma 2.7 The algorithm finds a distribution with value that equals the distribution with the maximum expected welfare in $\mathcal{R}_{(c,p)}$.

Proof: Notice that the optimal solution of the linear program with the proxy valuations is exactly the expected value of the optimal distribution in $\mathcal{R}_{c,1}$: the proxy valuations “simulate” the random process where each bidder keeps each item with probability c . Thus, it suffices to prove that the algorithm always finds a solution with value *exactly* $p \cdot OPT$, where OPT is the value of the optimal solution of the LP. This will guarantee us that the expected value of the solution equals $\mathcal{R}_{(c,p)}$.

Notice that after the tentative assignment in Step 3 the expected value of the sum of the tentative bundles (with respect to the proxy valuations) is exactly the value of the optimal solution of the linear program. After step 5 the expected value of the bundles assigned to the bidders (now with respect to the real valuations) is greater than $p \cdot OPT$. The last step “cancels” some of the allocations so that if bidder i was allocated bundle S , the probability he will be allocated some items from S (i.e., the probability that auction is not canceled at Step 3 and that he keeps some items at Step 6) is *exactly* $p \cdot x_{i,S}$. Thus the expected value of the solution is exactly $p \cdot OPT$, as needed. \square

Lemma 2.8 The communication complexity of the algorithm is polynomial in n and m .

Proof: The only two steps for which it is not obvious that only polynomial communication is required are Steps 2 and the calculation of the q_i 's. In Step 2 we solve a linear program that calculates the optimal fractional solution for some valuations. This can be done in polynomial communication, as long

as demand queries are available [2]. In our case we need to compute answers to demand queries with respect to the proxy valuations. This can be done by each bidder i with no additional communication since each proxy valuation v'_i depends only on v_i . We note that the support of the solution of the linear program consists only of polynomially many variables [2].

Calculating each q_i can be done using no communication cost simply by enumerating over all possible outputs of the random coins. \square

Together we have:

Theorem 2.9 *There exists a truthful-in-expectation $O(\frac{\log m}{\log \log m})$ -approximation mechanism for combinatorial auctions with submodular bidders. The algorithm uses polynomial communication.*

Acknowledgments

We thank Shaddin Dughmi and Tim Roughgarden for pointing out a bug in an earlier version of this note.

References

- [1] A. Archer, C. Papadimitriou, K. Talwar, and E. Tardos. An approximate truthful mechanism for combinatorial auctions with single parameter agent. In *SODA'03*.
- [2] Liad Blumrosen and Noam Nisan. On the computational power of demand queries. *SIAM J. Comput.*, 39(4):1372–1391, 2009.
- [3] Shahar Dobzinski. Two randomized mechanisms for combinatorial auctions. In *APPROX, 2007*.
- [4] Shahar Dobzinski and Shaddin Dughmi. On the power of randomization in algorithmic mechanism design. In *FOCS'09*.
- [5] Shahar Dobzinski and Noam Nisan. Limitations of vcg-based mechanisms. Preliminary version in *STOC'07*.
- [6] Shahar Dobzinski, Noam Nisan, and Michael Schapira. Approximation algorithms for combinatorial auctions with complement-free bidders. In *STOC'05*.
- [7] Shahar Dobzinski, Noam Nisan, and Michael Schapira. Truthful randomized mechanisms for combinatorial auctions. In *STOC'06*.
- [8] Uriel Feige. On maximizing welfare where the utility functions are subadditive. In *STOC'06*.
- [9] Ron Lavi and Chaitanya Swamy. Truthful and near-optimal mechanism design via linear programming. In *FOCS 2005*.