

# On the Computational Power of Demand Queries

Liad Blumrosen\* and Noam Nisan†

## Abstract

We study the computational power of iterative combinatorial auctions. Most existing iterative combinatorial auctions are based on repeatedly suggesting prices for bundles of items, and querying the bidders for their “demand” under these prices. We prove several results regarding such auctions that use a polynomial number of demand queries: (1) that such auctions can simulate several other natural types of queries; (2) that they can approximate the optimal allocation as well as generally possible using polynomial communication or computation, while weaker types of queries cannot do so; (3) that such auctions that only use item prices may solve allocation problems in communication cost that is exponentially lower than the cost incurred by auctions that use prices for bundles. For the latter result, we initiate the study of how prices of bundles can be represented when they are not linear, and show that the “default” representation has severe limitations. Our results hold for any series of demand queries with polynomial length, without any additional restrictions on the queries (e.g., to ascending prices).

**Keywords:** Combinatorial Auctions, Communication Complexity, Demand Queries, Iterative Auctions

**AMS Subject Classifications:** 91A28, 91B32, 94A05

## 1 Introduction

In a combinatorial auction, a set  $M$  of  $m$  non-identical items are sold in a single auction to  $n$  competing bidders. The bidders have preferences regarding the *bundles of items* that they may receive. The preferences of bidder  $i$  are specified by a valuation function  $v_i : 2^M \rightarrow R^+$ , where  $v_i(S)$  denotes the value that bidder  $i$  attaches to winning the bundle of items  $S$ . We assume “free disposal”, i.e., that the  $v_i$ ’s are monotone non-decreasing. The usual goal of the auctioneer is to optimize the social welfare  $\sum_i v_i(S_i)$ , where the allocation  $S_1, \dots, S_n$  must be a partition of the items. Applications include many complex resource allocation problems and, in fact, combinatorial auctions may be viewed as *the* common abstraction of many complex resource allocation problems. Combinatorial auctions face both economic and computational difficulties and are a central problem in the recently active border of economic theory and computer science.

---

\*blumrosen@gmail.com. Microsoft Research, Silicon Valley. 1065 La Avenida, Mountain View, CA, 94043. phone: 650-6932363

†noam@cs.huji.ac.il. School of Engineering and Computer Science, The Hebrew University of Jerusalem, Givat Ram, 91904, Israel. Phone: 972-2-6585503, Fax: 972-2-8585727. A preliminary version appeared in an extended abstract in the 6th ACM conference on Electronic Commerce (EC’05), and as a discussion paper no. 381 of the Center for the Study of Rationality, The Hebrew University. Supported by grants from the Israeli Academy of Sciences and the USA-Israel Binational Science Foundation. The authors thank Moshe Babaioff, Shahar Dobzinski, Ron Lavi, Daniel Lehmann, Ahuva Mu’alem, David Parkes, Michael Schapira and Ilya Segal for helpful discussions.

A forthcoming book [12] addresses many of the issues involved in the design and implementation of combinatorial auctions.

The design of a combinatorial auction involves many considerations. In this paper, we focus on just one central issue: the communication between bidders and the allocation mechanism – “preference elicitation”. Transferring all information about bidders’ preferences can require an infeasible (exponential in  $m$ ) amount of communication. Thus, “direct revelation” auctions in which bidders simply declare their preferences to the mechanism are only practical for very small auction sizes or for very limited families of bidder preferences. We have therefore seen a multitude of suggested “iterative auctions” in which the auction protocol repeatedly interacts with the different bidders, aiming to adaptively elicit enough information about the bidders’ preferences as to be able to find a good (optimal or close to optimal) allocation.

Most of the suggested iterative auctions proceed by maintaining temporary prices for the bundles of items and repeatedly querying the bidders as to their preferences between the bundles under the current set of prices, and then updating the set of bundle prices according to the replies received (e.g., [26, 14, 23, 44, 1]). Effectively, such an iterative auction accesses the bidders’ preferences by repeatedly making the following type of *demand query* to bidders: “Query to bidder  $i$ : a vector of bundle prices  $p = \{p(S)\}_{S \subseteq M}$ ; Answer: a bundle of items  $S \subseteq M$  that maximizes  $v_i(S) - p(S)$ ”. These types of queries are very natural in an economic setting as they capture the “revealed preferences” of the bidders. Some auctions, called *item-price* or *linear-price* auctions, specify a price  $p_i$  for each *item*, and the price of any given bundle  $S$  is always linear,  $p(S) = \sum_{i \in S} p_i$ . Other auctions, called *bundle-price* auctions, allow specifying arbitrary (non-linear) prices  $p(S)$  for bundles.

In this paper, we embark on a systematic analysis of the computational power of iterative auctions that are based on demand queries. We do not aim to present auctions for practical use but rather to understand the limitations and possibilities of these kinds of auctions. Our main question is what can be done using a polynomial number of these types of queries? That is, polynomial in the main parameters of the problem:  $n$ ,  $m$  and the number of bits  $t$  needed for representing a single value  $v_i(S)$ . Note that from an algorithmic point of view we are talking about sub-linear time algorithms: the input size here really consists of up to  $n(2^m - 1)$  numbers – the descriptions of the valuation functions of all bidders. There are two aspects to computational efficiency in these settings: the first is the communication with the bidders, i.e., the number of queries made, and the second is the “usual” computational tractability. Our hardness results (lower bounds) will depend only on the number of queries – and hold independently of any computational assumptions like  $P \neq NP$ . Our positive results (upper bounds) will always be computationally efficient both in terms of the number of queries and in terms of regular computation. As mentioned, this paper concentrates on the single aspect of preference elicitation and on its computational consequences and does not address issues of incentives. We also do not address the important issue of the complexity of calculating the response for such a demand query, and assume that it is tractable task for the users.<sup>1</sup> This strengthens our lower bounds, but means that the upper bounds also require evaluation from this perspective before being used in any real combinatorial auction.<sup>2</sup>

---

<sup>1</sup>Environments where realizing preferences is costly are discussed, e.g., in [29, 43].

<sup>2</sup>We do observe however that some weak incentive property comes for free in demand-query auctions since “myopic” players will answer all demand queries truthfully; i.e., if bidders only consider the (tentative) allocation decision at the current price level, they will have no incentive to misreport their true demand. We also note that in some cases the incentives issues can be handled orthogonally to the preference elicitation issues. For example, one can use Vickrey-Clarke-Groves (VCG) prices (e.g., [2, 41]) when the socially-optimal solution is chosen; this will not always work since such general schemes do not always exist, e.g., for approximate solutions or for objective functions other than welfare maximization, or because that calculating the appropriate prices requires additional

Valuation family	Upper bound	Reference	Lower bound	Reference
<b>General</b>	$\min(n, O(\sqrt{m}))$	[31], Section 5	$\min(n, m^{1/2-\epsilon})$	[40]
<b>Substitutes</b>	1	[22, 4]		
<b>Submodular</b>	$\frac{e}{e-1}$	[18]	$\frac{276}{275}$	[21]
<b>Subadditive</b>	2	[20]	$2-\epsilon$	[15]
<b>k-duplicates</b>	$O(m^{1/k+1})$	[17],[10]	$O(m^{1/k+1})$	[17]
<b>Procurement</b>	$\ln m$	[40]	$(\log m)/2$	[36, 40]

Figure 1: The best approximation factors currently achievable by computationally-efficient combinatorial auctions, for several classes of valuations. All lower bounds in the table apply to all iterative auctions; all upper bounds in the table are achieved with item-price demand queries. The table concerns the following classes of valuations: *Substitutes* valuations are defined in Appendix B. A valuation is *Submodular* if for every two bundles  $S, T$  we have  $v(S) + v(T) \geq v(S \cup T) + v(S \cap T)$ . A valuation is *Subadditive* if for every such two bundles  $v(S) + v(T) \geq v(S \cup T)$ . The term *k-duplicates* refers to combinatorial auctions where we have  $k$  identical copies of each item. *Procurement* auctions are “reversed” combinatorial auctions, where a single buyer needs to buy a set of  $m$  items from  $n$  suppliers, and the buyer tries to minimize her total payment.

In a companion paper ([9]) we study similar questions for the more restricted natural case of *ascending-price* combinatorial auctions.

## 1.1 Extant Work

Many iterative combinatorial auction mechanisms rely on demand queries (see the survey in [42]). For our purposes, two families of these auctions serve as the main motivating starting points: the first is the ascending item-price auctions of [26, 14, 23] that with computational efficiency find an optimal allocation among “(gross) substitutes” valuations<sup>3</sup>, and the second is the ascending bundle-price auctions of [44, 1] that find an optimal allocation among general valuations – but not necessarily with computational efficiency.<sup>4</sup> The main lower bound in this area, due to [40], states that indeed, due to inherent communication requirements, it is impossible for any iterative auction to find the optimal allocation among general valuations with sub-exponentially many queries. A similar exponential lower bound was shown by [40] also for even approximating the optimal allocation to within a factor of  $m^{1/2-\epsilon}$  for every  $\epsilon > 0$ . Several lower bounds and upper bounds for approximation are known for some natural classes of valuations – these are summarized in Figure 1.

In [40], the universal generality of demand queries is also shown: any *non-deterministic* communication protocol for finding an allocation that optimizes the social welfare can be converted into one that only uses demand queries (with bundle prices). In [46] this was generalized also to non-deterministic protocols for finding allocations that satisfy other natural types of economic requirements (e.g., approximate social efficiency, envy-freeness). However, in [39] it was demonstrated that this “completeness” of demand queries holds only in the nondeterministic setting, while in the usual deterministic setting, demand queries (even with bundle prices) may be exponentially weaker than general communication.

---

information ([19]).

<sup>3</sup>Informally, the substitutes property means that the bidder will continue to demand an item when the prices of some *other* items were raised. For completeness, the exact definition is given in Appendix B.

<sup>4</sup>Other auction designs that use bundle-price demand queries include [34, 13].

The organization of the rest of the paper is as follows: First, in Section 2, we present an informal exposition that describes our new results and their context. Section 3 describes our model. In Section 4 we discuss the power of different types of queries, and Section 5 studies the approximability of the social welfare with a polynomial number of queries. Finally, Section 6 compares the power of item-price demand queries to the power of demand queries that use bundle prices, focusing on the representation of bundle-price demand queries.

## 2 A Survey of Our Results

### 2.1 Comparison of Query Types

We first ask what other natural types of queries could we imagine iterative auctions using. Here is a list of such queries that are either natural, have been used in the literature, or that we found useful.

1. *Value query*: The auctioneer presents a bundle  $S$ , the bidder reports his value  $v(S)$  for this bundle.
2. *Marginal-value query*: The auctioneer presents a bundle  $A$  and an item  $j$ , the bidder reports how much he is willing to pay for  $j$ , given that he already owns  $A$ , i.e.,

$$v(j|A) = v(A \cup \{j\}) - v(A)$$

3. *Demand query (with item prices)*: The auctioneer presents a vector of item prices  $p_1, \dots, p_m$ ; the bidder reports his demand under these prices, i.e., some set  $S$  that maximizes  $v(S) - \sum_{i \in S} p_i$ .<sup>5</sup>
4. *Indirect-utility query*: The auctioneer presents a set of item prices  $p_1, \dots, p_m$ , and the bidder responds with his “indirect-utility” under these prices, that is, the highest utility he can achieve from a bundle under these prices:  $\max_{S \subseteq M} (v(S) - \sum_{i \in S} p_i)$ .<sup>6</sup>
5. *Relative-demand query*: the auctioneer presents a set of non-zero prices  $p_1, \dots, p_m$ , and the bidder reports the bundle that maximizes his value per unit of money, i.e., some set that maximizes  $\frac{v(S)}{\sum_{i \in S} p_i}$ . We apply this type of query, for example, in the design of the approximation algorithm described in Figure 4 in Section 5.<sup>7</sup>

**Theorem [see Section 4]:** Each of these queries can be efficiently (i.e., in time polynomial in  $n$ ,  $m$ , and the number of bits of precision  $t$  needed to represent a single value  $v_i(S)$ ) simulated by a sequence of demand queries with item prices.

In particular, this shows that demand queries can elicit all information about a valuation by simulating all  $2^m - 1$  value queries. We also observe that value queries and marginal-value queries can simulate each other in polynomial time and that demand queries and indirect-utility queries can also simulate each other in polynomial time. We prove that exponentially many value queries may be needed in order to simulate a single demand query.<sup>8</sup>

<sup>5</sup>A tie breaking rule should be specified. All of our results apply for any fixed tie breaking rule.

<sup>6</sup>This is exactly the utility achieved by the bundle which would be returned in a demand query with the same prices. This notion relates to the *indirect-utility function* studied in the Microeconomic literature (see, e.g., [32]).

<sup>7</sup>Note that when all the prices are 1, the bidder actually reports the bundle with the highest per-item value.

<sup>8</sup>It is interesting to note that for the restricted class of substitutes valuations, demand queries may be simulated by polynomial number of value queries [3].

Query type	Upper bound	Reference	Lower bound	Reference
<i>General Communication</i>	$\min(n, O(m^{1/2}))$	[31]	$\min(n, m^{1/2-\epsilon})$	[40]
<i>Demand Queries</i>	$\min(n, O(m^{1/2}))$	<b>new</b>	$\min(n, m^{1/2-\epsilon})$	[40]
<i>Value Queries</i>	$O(\frac{m}{\sqrt{\log m}})$	[25]	$O(\frac{m}{\log m})$	<b>new</b>

Figure 2: Achievable approximation factors for the social welfare using polynomially many value queries, demand queries (with item prices), and general queries (communication).

## 2.2 Welfare Approximation

The next question that we ask is how well can a computationally-efficient auction that uses only demand queries *approximate* the optimal allocation? Two separate obstacles are known: In [40], a lower bound of  $\min(n, m^{1/2-\epsilon})$ , for any fixed  $\epsilon > 0$ , was shown for the approximation factor obtained using any polynomial amount of communication. A computational bound with the same value applies even for the case of single-minded bidders under the assumption of  $NP \neq P$  [45, 31, 47]. As noted in [40], the computationally-efficient greedy algorithm of [31] can be adapted to become a polynomial-time iterative auction that achieves a nearly matching approximation factor of  $\min(n, O(\sqrt{m}))$ . This iterative auction may be implemented with bundle-price demand queries but, as far as we see, not as one with item prices. Since in a single bundle-price demand query an exponential number of prices can be presented, this algorithm can have an exponential communication cost. In Section 5, we describe a different item-price auction that achieves, for the first time, the same approximation factor with a polynomial number of demand queries (and thus polynomial communication).

**Theorem [Theorem 1, Section 5]:** There exists a computationally-efficient iterative auction with item-price *demand queries* that finds an allocation that approximates the optimal welfare between arbitrary valuations to within a factor of  $\min(n, O(\sqrt{m}))$ .

One may then attempt obtaining such an approximation factor using iterative auctions that use only the weaker value queries. However, we show that this is impossible:

**Theorem [Theorem 2, Section 5]:** Any iterative auction that uses a polynomial (in  $n$  and  $m$ ) number of *value queries* cannot achieve an approximation factor that is better than  $O(\frac{m}{\log m})$ .<sup>9</sup>

Note however that auctions with only value queries are not completely trivial in power: the bundling auctions of [25] can easily be implemented by a polynomial number of value queries and can achieve an approximation factor of  $O(\frac{m}{\sqrt{\log m}})$  by using  $O(\log m)$  equi-sized bundles. We do not know how to close the (tiny) gap between this upper bound and our lower bound. Figure 2 summarizes these upper and lower bounds.

## 2.3 Representing Bundle Prices

The different pricing methods for iterative combinatorial auctions generate much debate among auction designers, both in theory and in practice. One prominent example is the design of spectrum auctions by the FCC in the US. The final part of our paper compares the power of item prices and bundle prices. It is straightforward to see that bundle-price demand queries are a generalization of item price queries; we also observe that existing results actually imply that bundle-price queries are more powerful than item-price queries, since an exponential number of

<sup>9</sup>This was also proven independently by Shahar Dobzinski and Michael Schapira in [17].

item-price queries may be required for simulating a bundle-price demand query. As we shall see, the latter observation is not delicate enough and one should not reach strike conclusions when comparing the two query types. This can be shown by dealing with a critical issue with bundle-price auctions that was side-stepped by our model, as well as by most previous work that used bundle-price auctions: *how are the bundle prices represented?*

For item-price auctions this is not an issue since a query needs only to specify a small number,  $m$ , of prices. In bundle-price auctions that situation is more difficult since there are exponentially many bundles that require pricing. Our basic model (like all previous work that used bundle prices, e.g., [44, 41, 1]), ignores this issue, and only requires that the prices be determined, *somehow*, by the protocol. A finer model would fix a specific *language* for denoting bundle prices, force the protocol to represent the bundle-prices in this language, and require that the *representations of the bundle-prices* also be polynomial.

What could such a language for denoting prices for all bundles look like? First note that specifying a price for each bundle is equivalent to specifying a *valuation*. Second, as noted in [37], most of the proposed *bidding languages* are really just languages for representing valuations, i.e., a syntactic representation of valuations – thus we could use any of them. This point of view opens up the general issue of *which* language should be used in bundle-price auctions and what are the implications of this choice.

Here we initiate this line of investigation. We consider bundle-price auctions where the prices must be given as a XOR-bid, i.e., the protocol must explicitly indicate the price of every bundle whose value is different than that of all of its proper subsets. Note that all bundle-price auctions that do not explicitly specify a bidding language must implicitly use this language or a weaker one, since without a specific language one would need to list prices for all bundles, except perhaps for trivial ones (those with value 0, or more generally, those with a value that is determined by one of their proper subsets.) We show that once the representation length of bundle prices is taken into account (using the XOR-language), bundle-price auctions are no more strictly stronger than item-price auctions. Our proof relies on the sophisticated known lower bounds for constant depth circuits due to Håstad [24]. We were not able to find an elementary proof.

Define the *cost* of an iterative auction as the total length of the queries and answers used throughout the auction (in the worst case):

**Theorem [Theorem 3, Section 6]:** For some profile of valuations, bundle price auctions that use the XOR-language require an exponential cost for finding the optimal allocation. In contrast, item-price auctions can find the optimal allocation for this class within polynomial cost.

This puts doubts on the applicability of bundle-price auctions like [1, 44], and it may justify the use of “hybrid” pricing methods such as Ausubel, Cramton and Milgrom’s Clock-Proxy auction ([11]).

### 3 The Model

A single auctioneer is selling  $m$  indivisible non-homogeneous items in a single auction, and let  $M$  be the set of these items and  $N$  be the set of bidders. Each one of the  $n$  bidders in the auction has a valuation function  $v_i : 2^M \rightarrow \{0, 1, \dots, L\}$ , where for every bundle of items  $S \subseteq M$ ,  $v_i(S)$  denotes the value of bidder  $i$  for the bundle  $S$  and is an integer in the range  $0, \dots, L$ . We will sometimes denote the number of bits needed to represent an integer in the range  $0, \dots, L$  by  $t = \log L$ . We assume free disposal, i.e.,  $S \subseteq T$  implies  $v_i(S) \leq v_i(T)$  for all  $i \in N$  and that  $v_i(\emptyset) = 0$  for all bidders.

A valuation is called a *k-bundle XOR* if it can be represented as a XOR combination of (at most)  $k$  atomic bids [35], i.e., if there are at most  $k$  bundles  $S_j$  and prices  $p_j$ ,  $1 \leq j \leq k$ , such that for all  $S$ ,  $v(S) = \max_{j|S \supseteq S_j} p_j$ .<sup>10</sup> As mentioned, this is the standard method to represent general valuations. A discussion on this representation language and other natural languages can be found in [37].

### 3.1 Iterative Auctions

The auctioneer sets up a protocol (equivalently an “algorithm”), where at each stage of the protocol some information  $q$  – termed the “query” – is sent to some bidder  $i$ , and then bidder  $i$  replies with some reply that depends on the query as well as on his own valuation. In this paper, we assume that we have complete control over the bidders’ behavior, and thus the protocol also defines a reply function  $r_i(q, v_i)$  that specifies bidder  $i$ ’s reply to query  $q$ . In general, the term *iterative auction* refers to any interactive protocol of such form, but it is natural to have restrictions on the set of allowed queries. For example, one may study an iterative auction that uses only value queries, only demand queries, or both. The most studied family of iterative auctions is *ascending-price* auctions, that are actually iterative auctions that use only demand queries but have the additional restriction that prices presented to bidders cannot decrease over time (see [9] for a systematic analysis of ascending auctions). This paper studies the power of iterative auctions that use only demand queries whose prices are not necessarily ascending, and compare such auctions to auctions that use other types of queries. Note that the protocol may be adaptive: the query value as well as the queried bidder may depend on the replies received for past queries. At the end of the protocol, an *allocation*  $S_1, \dots, S_n$  must be declared, where  $S_i \cap S_j = \emptyset$  and  $S_i, S_j \subseteq M$  for all  $i \neq j$ . To summarize, combinatorial iterative auctions are protocols where the auctioneer publishes a series of queries (taken from a collection of allowed series of queries), and the outcome of the protocol is an allocation that is determined based on all the queries and the responses to them.

We say that the auction finds an *optimal allocation* if it finds the allocation that maximizes the social welfare  $\sum_i v_i(S_i)$  over all possible allocations  $S_1, \dots, S_n$ . We say that it finds a  $c$ -approximation if  $\sum_i v_i(S_i) \geq \sum_i v_i(T_i)/c$  where  $T_1, \dots, T_n$  is an optimal allocation. The running time of the auction on a given instance of the bidders’ valuations is the total number of queries made on this instance. The running time of a protocol is the worst case cost over all instances. Note that we impose no computational limitations on the protocol or on the players.<sup>11</sup> This of course only strengthens our hardness results. Yet, our positive results will not use this power and will be efficient also in the usual computational sense.

Our goal will be to design computationally-efficient protocols. We will deem a protocol computationally-efficient if its cost is polynomial in the relevant parameters: the number of bidders  $n$ , the number of items  $m$ , and  $t = \log L$ , where  $L$  is the largest possible value of a bundle. Note that all of our results give concrete bounds, where the dependence on the parameters is given explicitly; we use the standard big-Oh notation just as a shorthand.

---

<sup>10</sup>Following is an example for a 3-XOR valuation: consider a bidder with values of  $p_1 = 5, p_2 = 3, p_3 = 4$  for the atomic bundles  $S_1 = abcd, S_2 = ac, S_3 = b$ , respectively. The value of a bundle is the maximal value of an atomic bundle it contains, e.g.,  $v(ac) = 3, v(dcb) = 4$  but  $v(abcd) = 5$ .

<sup>11</sup>The running time really measures communication costs and not computational running time.

## 3.2 Demand Queries

Most of the paper will be concerned with a common special case of iterative auctions in which the queries that are sent to bidders are demand queries: the query specifies a price  $p(S) \in \mathbb{R}^+$  for each bundle  $S$ . The reply of bidder  $i$  is simply the set most desired – “demanded” – under these prices. Formally, bidder  $i$  replies with a set  $S$  that maximizes  $v_i(S) - p(S)$ . It may happen that more than one set  $S$  maximizes this value. In which case, ties are broken according to some fixed tie-breaking rule, e.g., the lexicographically first such set is returned. All of our results hold for any fixed tie-breaking rule.<sup>12</sup>

Note that even though in our model valuations are integral, we allow the demand query to use arbitrary real numbers. A practical issue here is how will the query be specified: in the general case, an exponential number of prices needs to be sent in a single query. Formally, this is not a problem as the model does not limit the length of queries in any way – the protocol must simply define what the prices are in terms of the replies received for previous queries. We look into this issue further in Section 6.

Many auctions in the literature restrict the prices’ representation to item prices (or linear prices):

**Definition 1. Item Prices:** The prices in each query are given by prices  $p_j$  for each item  $j$ . The price of a set  $S$  is additive:  $p(S) = \sum_{j \in S} p_j$ .

## 4 The Power of Different Types of Queries

In this section, we compare the power of the various types of queries defined in the introduction. We will present computationally-efficient simulations of these query types using item-price demand queries. The opposite, however, is false: we show that an exponential number of some of these queries may be needed for simulating demand queries. Figure 3 summarizes the relations between the different query types. Some parts of the following lemmas are elementary, and some are harder. These lemmas will be used in the analysis in the rest of this paper. All missing proofs can be found in Appendix A.

**Lemma 1.** *A value query can be simulated by  $m$  marginal-value queries. A marginal-value query can be simulated by two value queries.*

**Lemma 2.** *A value query can be simulated by  $mt$  demand queries (where  $t = \log L$  is the number of bits needed to represent a single bundle value).<sup>13</sup>*

As a direct corollary we get that demand auctions can always fully elicit the bidders’ valuations by simulating all possible  $2^m - 1$  queries and thus elicit enough information for determining the optimal allocation. Note, however, that this elicitation may be computationally inefficient.

The next lemma shows that demand queries can be exponentially more powerful than value queries.

---

<sup>12</sup>Much of the previous work on combinatorial auctions assumed that bidders respond with all the bundles in their demand sets. This can imply communicating an exponential number of bundles at a single step. All our results hold for this definition of demand queries (except Lemma 4, where it is clear that simulating such queries is a harder task). In order to measure the total communication rather than the number of queries, we assume that bidders report a single bundle at each step, and our results hold even for this less expressive query.

<sup>13</sup>Note that  $t$  bundle-price demand queries can easily simulate a value query by setting the prices of all the bundles except  $S$  (the bundle with the unknown value) to be  $L$ , and performing a binary search on the price of  $S$ .



	Value	Mar-value	Demand	Ind-util	Rel-demand
Value query	1	2	exp	exp	exp
Marginal-value query	$m$	1	exp	exp	exp
Demand query	$mt$	poly	1	$mt+1$	poly
Indirect-utility query	1	2	$m+1$	1	poly
Relative-demand query	-	-	-	-	1

Figure 3: Each entry in the table specifies how many queries of the relevant row are needed to simulate a query from the relevant column. A polynomial number of demand queries can simulate all other queries in the list. Relative-demand queries cannot simulate even a value query, therefore the last row of the table is empty.

**Lemma 3.** *An exponential number of value queries may be required for simulating a single demand query.*

*Proof.* We will actually show an example where a single demand query suffices for finding the optimal allocation, but an exponential number of value queries are required for that. Consider a player with a valuation of  $2|S|$  for any bundle  $S$ , except for some “hidden” bundle  $H$  of size  $\frac{m}{2}$  with a valuation of  $2|S| + 2$ , and a second player with a known valuation of  $2|S| + 1$  for every bundle  $S$ . The only optimal allocation gives the hidden set  $H$  to the first bidder. In a demand query with a price of  $2 + \epsilon$  for every item, the first bidder demands his “hidden” set, and thus reveals the optimal allocation.

However, consider any algorithm that uses only value queries. An adversary will answer each value query  $v(S)$  to the first bidder with  $v(S) = 2|S|$ . As long as two sets  $S$  of size  $\frac{m}{2}$  have not been queried any of them can be the hidden set  $H$  and the optimal allocation cannot be determined. Thus,  $\Omega(2^m)$  value queries will be needed in the worst case.  $\square$

Indirect utility queries are, however, equivalent in power to demand queries:

**Lemma 4.** *An indirect-utility query can be simulated by  $mt + 1$  demand queries. A demand query can be simulated by  $m + 1$  indirect-utility queries.*

Demand queries can also simulate relative-demand queries:<sup>14</sup>

**Lemma 5.** *Relative-demand queries can be simulated by a polynomial number of demand queries.*

According to our definition of relative-demand queries, they clearly cannot simulate even value queries. Consider, for example, two bidders 1,2 where for every bundle  $S$ ,  $v_2(S) = c \cdot v_1(S)$  for some  $c > 0$ . It is straightforward from the definition of relative-demand queries that the responses of the two bidders to such queries will always be identical, although they have different values, so no value query can be simulated by a relative-demand query. This is the reason why the lower row in Figure 3 is empty.

<sup>14</sup>Note: although in our model values are integral, we allow the query prices to be arbitrary real numbers, thus we may have bundles with arbitrarily close relative demands. In this sense the simulation above is only up to any given  $\epsilon$  (and the number of queries is  $O(\log L + \log \frac{1}{\epsilon})$ ). When the relative-demand query prices are given as rational numbers, exact simulation is implied when  $\log \epsilon$  is linear in the input length.

**An approximation algorithm:****Initialization:** Let  $T \leftarrow M$  be the current items for sale.Let  $K \leftarrow N$  be the currently participating bidders.Let  $S_1^* \leftarrow \emptyset, \dots, S_n^* \leftarrow \emptyset$  be the provisional allocation.**Repeat until  $T = \emptyset$  or  $K = \emptyset$ :**Ask each bidder  $i$  in  $K$  for the bundle  $S_i$  that maximizes her per-item value, i.e.,  $S_i \in \arg \max_{S \subseteq T} \frac{v_i(S)}{|S|}$ .Let  $i$  be the bidder with the maximal per-item value, i.e.,  $i \in \arg \max_{i \in K} \frac{v_i(S_i)}{|S_i|}$ , and set:  $S_i^* = S_i$ ,  $K = K \setminus i$ ,  $T = T \setminus S_i$ **Finally:** Ask all bidders for their values  $v_i(M)$  for all items.If allocating all the items to some bidder  $i$  improves the social welfare achieved so far (i.e.,  $\exists i \in N$  such that  $v_i(M) > \sum_{i \in N} v_i(S_i^*)$ ), then allocate all items to this bidder  $i$ .

Figure 4: This algorithm achieves a  $\min\{n, 4\sqrt{m}\}$ -approximation for the social welfare, which is asymptotically the best worst-case approximation possible with polynomial communication. This algorithm can be implemented with a polynomial number of demand queries.

## 5 Approximating the Social Welfare with Value and Demand Queries

We know from [40] that iterative combinatorial auctions that only use a polynomial number of queries (of any kind) cannot find an optimal allocation among general valuations and, in fact, cannot even approximate it to within a factor better than  $\min\{n, m^{1/2-\epsilon}\}$ . Can such an approximation ratio be attained using demand queries, or even using the weaker value queries? In this section, we show that this lower bound can be matched using demand queries, while value queries can only do much worse.

Figure 4 describes a polynomial-time algorithm that achieves a  $\min(n, O(\sqrt{m}))$  approximation ratio. This algorithm greedily picks the bundles that maximize the bidders' per-item value (using “relative-demand” queries, see Lemma 5). As a final step, it allocates all the items to a single bidder if it improves the social welfare (this can be checked using value queries). Since both value queries and relative-demand queries can be simulated by a polynomial number of demand queries with item prices (Lemmas 2 and 5), this algorithm can be implemented by a polynomial number of demand queries with item prices.

**Theorem 1.** *The auction described in Figure 4 can be implemented by a polynomial number of demand queries and achieves a  $\min\{n, 4\sqrt{m}\}$ -approximation for the social welfare.*

*Proof.* We first observe that the algorithm can be implemented by a polynomial number of value queries and relative demand queries: querying a bidder for the bundle that maximizes his per-item value is a relative-demand query when all the item prices are 1, and revealing the value of this bundle requires one value query. Querying a bidder for his value for all items can be done by an additional value query. In Section 4, we showed that any value query and any relative-demand query can be implemented by a polynomial number of demand queries. Each bidder is asked at most  $m$  relative demand queries, and exactly two value queries, thus a polynomial number of demand queries can implement this algorithm.

Next, we prove that the algorithm achieves an approximation ratio of at least  $\min\{n, 4\sqrt{m}\}$ . The algorithm will clearly achieve a  $\frac{1}{n}$ -approximation since we allocate the whole bundle  $M$  to

the bidder with the highest valuation if it improves the welfare achieved. Next, we prove that the algorithm achieves at least  $\frac{1}{4\sqrt{m}}$  of the optimal welfare.

Let  $OPT = \{T_1, \dots, T_k, Q_1, \dots, Q_l\}$  be an optimal allocation where for every  $i \in \{1, \dots, k\}$   $|T_i| \leq \sqrt{m}$  and for every  $j \in \{1, \dots, l\}$   $|Q_j| > \sqrt{m}$  ( $l, k \in \{0, \dots, n\}$ ). Let  $ALG$  be the allocation found by the algorithm, and let  $v(OPT)$  and  $v(ALG)$  be the optimal welfare and the welfare achieved by the algorithm, respectively. First, we analyze cases where “large” bundles contribute most of the optimal welfare, i.e.,  $\sum_{i=1}^l v_i(Q_i) \geq \sum_{i=1}^k v_i(T_i)$ . Then,

$$v(OPT) \leq 2 \sum_{i=1}^l v_i(Q_i) \leq 2 \sum_{i=1}^l v(ALG) = 2l \cdot v(ALG) \leq 2\sqrt{m} \cdot v(ALG)$$

Where the first inequality holds since  $v(OPT) = \sum_{i=1}^l v_i(Q_i) + \sum_{i=1}^k v_i(T_i)$  and the second holds since the last stage of the algorithm verifies that the welfare achieved by the algorithm is at least the valuation of every player for the whole bundle  $M$ . The last inequality holds since there are no more than  $\sqrt{m}$  bundles of size of at least  $\sqrt{m}$ .

The analysis of the case where “small” bundles contribute most of the optimal welfare (i.e.,  $\sum_{i=1}^l v_i(Q_i) < \sum_{i=1}^k v_i(T_i)$ ) is more involved. Let  $I \subseteq \{1, \dots, k\}$  be the set of bidders that receives a “small” bundle (i.e., bundles in  $\{T_1, \dots, T_k\}$ ) in  $OPT$  that does not intersect any bundle in  $ALG$ . Consider the following sum:

$$\sum_{i=1}^k \frac{v_i(T_i)}{|T_i|} = \sum_{i \in I} \frac{v_i(T_i)}{|T_i|} + \sum_{i \in \{1, \dots, k\} \setminus I} \frac{v_i(T_i)}{|T_i|} \quad (5.1)$$

In the two claims below, we show that each of the summands in the right side of Equation 5.1 is not greater than  $v(ALG)$ . This immediately derives that  $\sum_{i=1}^k \frac{v_i(T_i)}{|T_i|} \leq 2 \cdot v(ALG)$ . Since for every  $i \in 1, \dots, k$ ,  $|T_i| \leq \sqrt{m}$ :

$$\sum_{i=1}^k v_i(T_i) \leq 2\sqrt{m} \cdot v(ALG)$$

Since most of the optimal welfare is contributed by “small” bundles,

$$v(OPT) \leq 2 \sum_{i=1}^k v_i(T_i) \leq 4\sqrt{m} \cdot v(ALG)$$

What is left to be proved is that both summands in Equation 5.1 are not greater than  $v(ALG)$ :

*Claim 1.*  $\sum_{i \in I} \frac{v_i(T_i)}{|T_i|} \leq v(ALG)$

*Proof.* Consider a bidder  $i$  that receives a small bundle  $T_i$  in  $OPT$  such that  $T_i$  is disjoint to all bundles in  $ALG$ . We observe that this bidder surely receives a non-empty bundle  $S_i$  in  $ALG$ . This holds since the items in  $T_i$  are not allocated at the end of the algorithm (they are not in any bundle in  $ALG$ ), but player  $i$  has a non-zero value for  $T_i$ .

Since the algorithm picked some  $S_i \in ALG$  and not  $T_i$  for bidder  $i$ ,  $\frac{v_i(T_i)}{|T_i|} \leq \frac{v_i(S_i)}{|S_i|}$ . Therefore,

$$\sum_{i \in I} \frac{v_i(T_i)}{|T_i|} \leq \sum_{i \in I} \frac{v_i(S_i)}{|S_i|} \leq \sum_{i \in I} v_i(S_i) \leq \sum_{i=1}^n v_i(S_i) = v(ALG)$$

□

*Claim 2.*  $\sum_{i \in \{1, \dots, k\} \setminus I} \frac{v_i(T_i)}{|T_i|} \leq v(ALG)$

*Proof.* For every bidder  $i \in \{1, \dots, k\} \setminus I$ ,  $T_i$  intersects at least one bundle from  $ALG$ , and let  $F(i)$  be the first bidder for which the algorithm allocates a bundle that intersects  $T_i$ . Then,

$$\sum_{i \in \{1, \dots, k\} \setminus I} \frac{v_i(T_i)}{|T_i|} \leq \sum_{j=1}^n \sum_{i|F(i)=j} \frac{v_i(T_i)}{|T_i|} \leq \sum_{j=1}^n \sum_{i|F(i)=j} \frac{v_j(S_j)}{|S_j|} \leq \sum_{j=1}^n |S_j| \frac{v_j(S_j)}{|S_j|} \leq \sum_{j \in ALG} v_j(S_j)$$

Where the second leftmost inequality holds since bidder  $j = F(i)$  demands  $S_j \in ALG$  when all the items in  $T_i$  are still on sale and the third inequality holds since each  $S_j$  intersects at most  $|S_j|$  bundles from  $\{T_1, \dots, T_k\}$  (all  $T_i$ 's are disjoint).  $\square$

We showed before how the theorem follows from these two claims.  $\square$

We now ask how well can the optimal welfare be approximated by a polynomial number of *value queries*. First we note that value queries are not completely powerless: In [25] it is shown that if the  $m$  items are split into  $k$  fixed bundles of size  $m/k$  each, and these fixed bundles are auctioned as though each was indivisible, then the social welfare generated by such an auction is at least  $\frac{m}{\sqrt{k}}$ -approximation of that possible in the original auction. Notice that such an auction can be implemented by  $2^k - 1$  value queries to each bidder – querying the value of each bundle of the fixed bundles. Thus, if we choose  $k = \log m$  bundles we get an  $\frac{m}{\sqrt{\log m}}$ -approximation while still using a polynomial number of queries.

We show that not much more is possible using value queries:

**Lemma 6.** *Any iterative auction that uses only value queries and distinguishes between  $k$ -tuples of 0/1 valuations where the optimal allocation has value 1, and those where the optimal allocation has value  $k$  requires at least  $2^{\frac{m}{k}}$  queries.*

*Proof.* Consider the following family of valuations: for every  $S$ , such that  $|S| > m/2$ ,  $v(S) = 1$ , and there exists a single bundle  $T$  (the "hidden" bundle), such that for  $|S| \leq m/2$ ,  $v(S) = 1$  iff  $T \subseteq S$  and  $v(S) = 0$  otherwise. Now look at the behavior of the protocol when all valuations  $v_i$  have  $T = \{1, \dots, m\}$ . Clearly in this case the value of the best allocation is 1 since no set of size  $\frac{m}{2}$  or lower has non-zero value for any player. Fix the sequence of queries and answers received on this  $k$ -tuple of valuations.

Now consider the  $k$ -tuple of valuations chosen at random as follows: a partition of the  $m$  items into  $k$  sets  $T_1, \dots, T_k$  each of size  $\frac{m}{k}$  is chosen uniformly at random among all such partitions. Now consider the  $k$ -tuple of valuations from our family that correspond to this partition –  $T_i$  is the "hidden bundle" for player  $i$ . Clearly,  $T_i$  can be allocated to  $i$ , for each  $i$ , getting a total value of  $k$ . Now look at the protocol when running on these valuations and compare its behavior to the original case. Note that the answer to a query  $S$  to player  $i$  differs between the case of  $T_i$  and the original case of  $T = \{1, \dots, m\}$  only if  $|S| \leq \frac{m}{2}$  and  $T_i \subseteq S$ . Since  $T_i$  is distributed uniformly among all sets of size exactly  $\frac{m}{k}$ , we have that for any fixed query  $S$  to player  $i$ , where  $|S| \leq \frac{m}{2}$ ,

$$Pr[T_i \subseteq S] \leq \left(\frac{|S|}{m}\right)^{|T_i|} \leq 2^{-\frac{m}{k}}$$

Using the union-bound, if the original sequence of queries was of length less than  $2^{\frac{m}{k}}$ , then with positive probability none of the queries in the sequence would receive a different answer than for

the original input tuple. This is forbidden since the protocol must distinguish between this case and the original case – which cannot happen if all queries receive the same answer. Hence there must have been at least  $2^{\frac{m}{k}}$  queries for the original tuple of valuations.  $\square$

We conclude that a polynomial time protocol that uses only value queries cannot obtain a better than  $O(\frac{m}{\log m})$  approximation of the welfare.

**Theorem 2.** *An iterative auction that uses a polynomial number of value queries cannot achieve better than  $O(\frac{m}{\log m})$ -approximation for the social welfare.*

*Proof.* This follows immediately from Lemma 6: achieving any approximation ratio  $k$  which is asymptotically greater than  $\frac{m}{\log m}$ , i.e.,  $\omega\left(\frac{m}{\log m}\right)$ , requires by Lemma 6 at least  $2^{\omega(\log m)}$  queries which is exponential in  $m$ .  $\square$

We conclude this section by mentioning that demand-queries have one more strong property – they allow solving the linear program for combinatorial auctions in polynomial time. The winner determination problem in combinatorial auctions may be formulated as an integer program. In many cases solving the linear-program relaxation of this integer program is useful: for some restricted classes of valuations it finds the optimum of the integer program (e.g., substitute valuations [26, 23]) or helps approximating the optimum. However, the linear program has an exponential number of variables. Nisan and Segal [40] observed the surprising fact that despite the exponential number of variables, this linear program may be solved within polynomial communication. The basic idea is to solve the dual program using the Ellipsoid method. The dual program has a polynomial number of variables, but an exponential number of constraints. The Ellipsoid algorithm runs in polynomial time even on such programs, provided that a “separation oracle” is given for the set of constraints. Surprisingly, such a separation oracle can be implemented using a single demand query (with item prices) to each of the bidders. Interestingly, several recent state-of-the-art approximation results for combinatorial auctions use randomized-rounding techniques that rely on the above observation (e.g., [15, 18, 7, 20, 21]). The treatment of [40] was somewhat ad-hoc to the problem at hand (the case of substitute valuations), and a more detailed analysis can be found in a preliminary version of this manuscript [8] and also in [38].

## 6 Demand Queries with Bundle Prices

In previous sections, we argued that item-price demand queries are powerful: they can simulate other natural types of queries, they can achieve the best tractable approximation ratio, they do much better than value queries, etc. Nevertheless, item price queries are known to be limited in their expressiveness. One notable weakness of item prices relates to the existence of a *competitive equilibrium*. A competitive equilibrium is a set of prices and an allocation, such that each bidder is allocated with his most desirable bundle under these prices and all the items are allocated; such allocations are always socially-optimal. With item prices, competitive equilibria always exist for the restricted family of *substitutes* valuations (see [26] and Definition 3 in Appendix B), but do not exist for practically any wider preferences domain (see, e.g., [33, 5]). One important consequence of this fact is that, for general valuations, an iterative process cannot converge to prices that induce such an optimal allocation. Furthermore, even if the allocation can be determined based on all the information gathered in an ascending-price iterative auction, an optimal allocation, or even a reasonable approximation, cannot be determined (see [9]).

Interestingly, these expressiveness issues are solved when the auctioneer is allowed to use *bundle-price* demand queries: the seller publishes prices for every bundle, and the bidders respond with their demand under the published prices. With bundle prices, a competitive equilibrium always exists (see, e.g., [44, 38]), and there are ascending-price iterative auctions that converge to such competitive equilibria, e.g., [44, 1]. In this section, we study the connection between item-price auctions and bundle-price auctions. On one hand, it is straightforward to see that bundle-price demand queries are a more powerful generalization of item-price demand queries. On the other hand, when one carefully analyzes the representation of bundle-price demand queries, this conclusion no longer holds. Consequently, bundle-price demand queries and item-price queries can be regarded as incomparable in their power.

## 6.1 Bundle Prices are Exponentially More Powerful than Item prices

Without taking into account any representational issues, it is easy to see that bundle-price queries are at least as powerful as item-price queries, and in some settings they are better by an exponential factor. The first claim is trivial, and the latter follows, for example, from results by [28, 6].

**Proposition 1.** *A single bundle-price demand query can simulate any item-price demand query. Conversely, simulating a bundle-price demand query may require an exponential number of item-price demand queries (in the number of items  $m$ ), even for bidders whose valuations has polynomial representations (e.g., by  $k$ -XOR expressions where  $k$  is polynomial in  $m$ ).*

*Proof.* The first statement is trivial: we can simulate any demand query with the item prices  $p_1, \dots, p_m$  by bundle price query with  $p(S) = \sum_{i \in S} p_i$ .

The second statement follows from two existing results: Lahaie and Parkes [28] devise a bundle-price auction that always computes the optimal allocation using a polynomial number of bundle price demand queries for bidders with  $k$ -XOR valuations<sup>15</sup> whenever  $k$  is polynomial. In contrast, [6] show that there exist valuations that are XORs of  $k = \sqrt{m}$  bundles such that any item-price auction that finds an optimal allocation between them requires exponentially many queries.  $\square$

It is important to stress that using bundle-price demand queries with polynomial length (e.g., that present prices for a polynomial number of bundles) should still overcome existing hardness results. For example, achieving a reasonable approximation for the social welfare for general valuations requires communicating exponential amount of information [39] (even for restricted classes of valuations, e.g., submodular valuations), and, in particular, an exponential number of polynomial-length bundle-price demand queries is required in the worst case.

## 6.2 On the Representation of Bundle Prices Demand Queries

Although bundle-price queries generalize item-price queries, a single bundle-price demand query can use an exponential number of prices which is clearly impractical. One may try solving this communication problem by some clever compact representation of the bundle prices; however, the representation of bundle-price auctions is implicit in most of the existing literature.

In this section, we explicitly fix the language in which bundle prices are presented to the bidders in bundle-price auctions. This language requires the algorithm to explicitly list the price

---

<sup>15</sup>These are valuations where bidders have values for  $k$  specific bundles, and the value of each bundle is the maximal value of one of these bundles that it contains. See definition in Section 3

of every bundle with a non-trivial price. “Trivial” in this context is a price that is equal to that of one of its proper subsets (which was listed explicitly). This representation is equivalent to the XOR-language for expressing valuations. Formally, each query  $q$  is given by an expression:  $q = (S_1 : p_1) \oplus (S_2 : p_2) \oplus \dots \oplus (S_l : p_l)$ . In this representation, the price demanded for every set  $S$  is simply  $p(S) = \max_{\{k=1, \dots, l \mid S_k \subseteq S\}} p_k$ .

**Definition 2.** The *length* of the query  $q = (S_1 : p_1) \oplus (S_2 : p_2) \oplus \dots \oplus (S_l : p_l)$  is  $l$ . The *communication cost* of an algorithm is the sum of the lengths of the queries asked during the operation of the algorithm on the worst case input.

Note that under this definition, bundle-price auctions are not necessarily stronger than item-price auctions. An item-price query that prices each item for 1, is translated to an exponentially long bundle-price query that needs to specify the price  $|S|$  for each bundle  $S$ . But perhaps bundle-price auctions can still find optimal allocations whenever item-price auctions can, without directly simulating such queries? We show that this is not the case: indeed, when the representation length is taken into account, bundle price auctions are sometimes seriously inferior to item price auctions.

Consider the following family of valuations: Each item is valued at 3, except that for some single set  $S$ , its value is a bit more:  $3|S| + b$ , where  $b \in \{0, 1, 2\}$ . Note that an item price auction can easily find the optimal allocation between any two such valuations: Set the prices of each item to  $3 + \epsilon$ ; if the demand sets of the two players are both empty, then  $b = 0$  for both valuations, and an arbitrary allocation is fine. If one of them is empty and the other non-empty, allocate the non-empty demand set to its bidder, and the rest to the other. If both demand sets are non-empty then, unless they form an exact partition, we need to see which  $b$  is larger, which we can do by increasing the price of a single item in each demand set.

We will show that any bundle-price auction that uses only the XOR-language to describe bundle prices requires an exponential communication cost (which includes the sum of all description lengths of prices) to find an optimal allocation between any two such valuations.

The complication in the proof stems from the fact that using XOR-expressions, the length of the price description depends on the number of bundles whose price is strictly larger than each of their subsets – this may be significantly smaller than the number of bundles that have a non-zero price. (The proof becomes easy if we require the protocol to explicitly name every bundle with non-zero price.) We do not know of any elementary proof for this lemma (although we believe that one can be found). Instead we reduce the problem to a well known lower bound in boolean circuit complexity [24] stating that boolean circuits of depth 3 that compute the majority function on  $m$  variables require  $2^{\Omega(\sqrt{m})}$  size.

**Lemma 7.** *Every bundle-price auction that uses XOR-expressions to denote bundle prices requires communication cost of  $2^{\Omega(\sqrt{m})}$  in order to find the optimal allocation among two valuations from the above family.*

*Proof.* Consider the protocol running on the following two valuations: the first has  $b = 0$  (i.e. is simply additive), and the second has  $b = 1$  for the set  $S$  of all items. In this case the outcome must be to allocate all to the second bidder. Let  $e_1, \dots, e_t$  be the queries made on this input, where each  $e_i = E_i^1 \oplus E_i^2 \oplus \dots \oplus E_i^{l_i}$ . Now consider what happens when the first valuation is changed so that for some  $S$  of size exactly  $m/2$ , we get a bonus  $b = 2$  – clearly the allocation must change so that this  $S$  is allocated to the first player – hence one of the queries  $e_1, \dots, e_t$  must change its answer. We will see that the fact that this is true for every such  $S$  implies that  $\sum_{i=1}^t l_i$  is exponential.

First note that if in  $e_i$  there exists some set of size  $m/2 + 1$  that has price zero, then the answer will not change as this set will give a surplus of at least  $3m/2 + 3$  as opposed to at most  $3m/2 + 2$  that  $S$  gives. Let us focus on an  $e_i$  that does not have such a set. We build a boolean DNF formula from this expression as follows: the variable set will be  $x_1, \dots, x_m$  – a variable for each item. Consider a term (atomic bid)  $E_i^j = (B_i^j, p_i^j)$  in  $e_i$ . We call this term essential if there exists some bundle of size exactly  $m/2 + 1$  and the price of this bundle in  $e_i$  is exactly  $p_i^j$ . For every essential term  $(B_i^j, p_i^j)$  in  $e_i$  we build a conjunction of the variables in it (ignoring the price for this bundle). We then take the disjunction of all of these conjunctions. First notice that this DNF must accept all inputs with more than  $m/2$  1's in the input – since otherwise consider a set that is not accepted by this expression, and the value of this set in  $e_i$  must be zero.

Now notice that if an input with 1's in exactly the set  $S$  of size exactly  $m/2$  is accepted by this formula, then the answer to query  $e_i$  will not change. The reason is that an accepted set  $S$  contains some essential bundle  $B_i^j$ , and thus its price in  $e_i$  would be at least  $p_i^j$ . However, since the bundle is essential, there exists some set of size  $m/2 + 1$  that is priced at exactly  $p_i^j$  – this set would clearly be preferable to  $S$  – the only set whose value has changed. Since for every set  $S$  of size exactly  $m/2$  the answer to one of the queries must change, at least one of the formulas constructed must reject the input with 1's exactly in  $S$ .

We now take the conjunction of all boolean expressions built for all  $i$ . This formula accepts all inputs with exactly  $m/2 + 1$  1's, and rejects all inputs with exactly  $m/2$  1's. Note that this formula is a conjunction of disjunctions of conjunctions of variables – a, so called, monotone depth 3 formula. Since it is a monotone formula, it computes the majority function. Its size is clearly bounded from above by the total length of all expressions  $e_i$ . We are now ready to invoke the well known lower bound by Hastad [24] that states that a depth 3 formula for majority must have size at least  $2^{\Omega(\sqrt{m})}$ .  $\square$

The following theorem follows immediately from the above lemma.

**Theorem 3.** *There are classes of valuations for which the optimal allocation can be determined by a polynomial number of item-price demand queries, but this task requires exponential communication cost when using bundle-price demand queries in XOR representation.*

## 7 Conclusions and Future Work

Much work has been going on since the first version of this paper was published. Several papers presented approximate iterative auctions for combinatorial auctions that use demand queries. For general valuations, two papers presented randomized algorithms that achieve the same  $O(\sqrt{m})$  approximation for the social welfare as our algorithm does, but also in an incentive compatible manner. Lavi and Swamy [30] presented a general method to convert algorithms to dominant-strategy incentive-compatible mechanisms based on a randomized-rounding technique of the relevant linear program; their work resulted in a  $O(\sqrt{m})$  approximation that is incentive-compatible in expectation. Dobzinski et al. [16] presented a randomized  $O(\sqrt{m})$  approximation that is truthful for any coin flipping done by the algorithm. The most interesting open problem in this context is whether there exists a *deterministic* mechanism that achieves a  $O(\sqrt{m})$  approximation (like the one discussed in Section 5) that is also incentive compatible.

More iterative auctions based on demand queries were recently designed for approximating the welfare in several sub-families of sub-additive valuations, e.g., [17, 20, 21]. Several gaps still exist between lower and upper bounds for these problems, and larger gaps exist between the



approximation results and results that are known to be attainable by incentive-compatible algorithms. There are currently no results separating what can be achieved by incentive-compatible algorithms and unrestricted algorithms, and proving such a separation result is another important open problem. A recent survey [38] surveys the state-of-the-art approximation mechanism, with demand queries, value queries and with general communication. Interestingly, all these approximation results use item-price demand queries, and it is unknown whether there are natural algorithmic environments where other types of queries are optimal. ([39] gave an artificial instance where demand queries are inferior to unrestricted communication protocols.)

Finally, this paper embarks on a discussion of the representation of bundle-price demand queries, queries that seem to become more popular in recent auction designs. We discussed the limitations of the most natural representation method - by XOR formulae. Lately, few papers studied different bidding languages also in the context of price representation (e.g., [28, 27]), still focusing on their role as bidding languages. A more general and systematic analysis of price-representation languages is still missing.

## References

- [1] L. M. Ausubel and P. R. Milgrom. Ascending auctions with package bidding. *Frontiers of Theoretical Economics*, 1:1–42, 2002.
- [2] Lawrence Ausubel. An efficient dynamic auction for heterogeneous commodities. *American Economic Review*, 96(3):602–629, 2006.
- [3] Alejandro Bertelsen. Substitutes valuations and  $m^b$ -concavity”. M.Sc. Thesis, The Hebrew University of Jerusalem, completed under the supervision of Daniel Lehmann., 2005.
- [4] Sushil Bikhchandani and John W. Mamer. Competitive equilibrium in an exchange economy with indivisibilities. *Journal of economic theory*, 74:385 – 413, 1997.
- [5] Sushil Bikhchandani and Joseph M. Ostroy. The package assignment model. *Journal of Economic theory*, 107:377–406, 2002.
- [6] Avrim Blum, Jeffrey C. Jackson, Tuomas Sandholm, and Martin A. Zinkevich. Preference elicitation and query learning. *Journal of Machine Learning Research*, 5:649–667, 2004.
- [7] Liad Blumrosen and Shahar Dobzinski. Welfare maximization in congestion games. In *Proceedings of the 7th ACM conference on Electronic commerce*, pages 52–61, 2006.
- [8] Liad Blumrosen and Noam Nisan. On the computational power of iterative auctions I: demand queries. Discussion paper no. 381, The Center for the Study of Rationality, The Hebrew University., 2005. An extended abstract in EC’05 contained preliminary results.
- [9] Liad Blumrosen and Noam Nisan. On the computational power of iterative auctions II: ascending auctions. Discussion paper no. 382, The Center for the Study of Rationality, The Hebrew University., 2005. An extended abstract in EC’05 contained preliminary results.
- [10] Patrick Briest, Piotr Krysta, and Berthold Vocking. Approximation techniques for utilitarian mechanism design. In *the 37th ACM symposium on Theory of computing*, pages 39–48, 2005.

- [11] P. Cramton, L.M. Ausubel, and P.R. Milgrom. *In P. Cramton and Y. Shoham and R. Steinberg (Editors), Combinatorial Auctions. Chapter 5. The Clock-Proxy Auction: A Practical Combinatorial Auction Design.* MIT Press., 2006.
- [12] P. Cramton, Y. Shoham, and R. Steinberg. *Combinatorial Auctions.* MIT Press., 2006.
- [13] S. de Vries, J. Schummer, and R. V. Vohra. On ascending Vickrey auctions for heterogeneous objects. *Journal of Economic Theory*, 132(1):95–118, 2007.
- [14] G. Demange, D. Gale, and M. Sotomayor. Multi-item auctions. *Journal of Political Economy*, 94:863–872, 1986.
- [15] Shahar Dobzinski, Noam Nisan, and Michael Schapira. Approximation algorithms for combinatorial auctions with complement-free bidders. In *The 37th ACM symposium on theory of computing*, pages 610–618, 2005.
- [16] Shahar Dobzinski, Noam Nisan, and Michael Schapira. Truthful randomized mechanisms for combinatorial auctions. In *Proceedings of the thirty-eighth annual ACM symposium on Theory of computing*, pages 644–652, 2006.
- [17] Shahar Dobzinski and Michael Schapira. Optimal upper and lower approximation bounds for k-duplicates combinatorial auctions. Working paper, the Hebrew University.
- [18] Shahar Dobzinski and Michael Schapira. An improved approximation algorithm for combinatorial auctions with submodular bidders. In *Proceedings of the seventeenth annual ACM-SIAM symposium on Discrete algorithm*, pages 1064–1073, 2006.
- [19] Ronald Fadel and Ilya Segal. The communication cost of selfishness. *Journal of Economic Theory*, to appear. A preliminary version appeared in *TARK '05*, 2005.
- [20] Uriel Feige. On maximizing welfare where the utility functions are subadditive. In *38th ACM Symposium on Theory of Computing*, pages 41–50, 2006.
- [21] Uriel Feige and Jan Vondrak. Approximation algorithms for allocation problems: Improving the factor of  $1-1/e$ . In *47th Annual IEEE Symposium on Foundations of Computer Science*, 2006.
- [22] Faruk Gul and Ennio Stacchetti. Walrasian equilibrium with gross substitutes. *Journal of Economic Theory*, 87:95 – 124, 1999.
- [23] Faruk Gul and Ennio Stacchetti. The english auction with differentiated commodities. *Journal of Economic Theory*, 92(3):66 – 95, 2000.
- [24] J. Hastad. Almost optimal lower bounds for small depth circuits. In *18th STOC*, pages 6–20, 1986.
- [25] Ron Holzman, Noa Kfir-Dahav, Dov Monderer, and Moshe Tennenholtz. Bundling equilibrium in combinatorial auctions. *Games and Economic Behavior*, 47:104–123, 2004.
- [26] A.S. Kelso and V.P. Crawford. Job matching, coalition formation, and gross substitute. *Econometrica*, 50:1483–1504, 1982.

- [27] Sébastien Lahaie, Florin Constantin, and David C. Parkes. More on the power of demand queries in combinatorial auctions: Learning atomic languages and handling incentives. In *Proc. 19th International Joint Conference on Artificial Intelligence (IJCAI'05)*, 2005.
- [28] Sébastien Lahaie and David C. Parkes. Applying learning algorithms to preference elicitation. In *The 5th ACM Conference on Electronic Commerce*, pages 180–188, 2004.
- [29] Kate Larson and Tuomas Sandholm. Costly valuation computation in auctions. In *Proceedings of the 8th conference on Theoretical Aspects of Rationality and Knowledge*, pages 169–182, 2001.
- [30] Ron Lavi and Chaitanya Swamy. Truthful and near-optimal mechanism design via linear programming. In *46th Annual IEEE Symposium on Foundations of Computer Science*, pages 595–604, 2005.
- [31] 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.
- [32] A. Mas-Collel, W. Whinston, and J. Green. *Microeconomic Theory*. Oxford university press, 1995.
- [33] Paul Milgrom. Putting auction theory to work: the simultaneous ascending auction. *Journal of Political Economy*, 108(2):245 – 272, 2000.
- [34] Debasis Mishra and David Parkes. Ascending price Vickrey auctions for general valuations. *Journal of Economic Theory*, 132(1):335–366, 2007.
- [35] Noam Nisan. Bidding and allocation in combinatorial auctions. In *Proceedings of the 2nd ACM conference on Electronic commerce*, pages 1–12, 2000.
- [36] Noam Nisan. The communication complexity of approximate set packing and covering. In *Proceedings of the 29th International Colloquium on Automata, Languages and Programming*, pages 868–875, 2002.
- [37] Noam Nisan. In *P. Cramton and Y. Shoham and R. Steinberg (Editors), Combinatorial Auctions. Chapter 9. Bidding Languages for Combinatorial Auctions*. MIT Press., 2006.
- [38] Noam Nisan. In *Noam Nisan, Tim Roughgarden, Eva Tardos and Vijay Vazirani (Editors), Algorithmic Game Theory. Chapter 10. Combinatorial Auctions*. Cambridge University Press. To appear., 2007.
- [39] Noam Nisan and Ilya Segal. Exponential communication inefficiency of demand queries. In *Proceedings of the 10th conference on Theoretical aspects of rationality and knowledge*, pages 158–164, 2005.
- [40] Noam Nisan and Ilya Segal. The communication requirements of efficient allocations and supporting prices. *Journal of Economic Theory*, 129(1):192 – 224, 2006.
- [41] D. C. Parkes and L. H. Ungar. An ascending-price generalized Vickrey auction. Presented at the Stanford Institute for Theoretical Economics (SITE) Summer Workshop, June 2002.
- [42] David Parkes. In *P. Cramton and Y. Shoham and R. Steinberg (Editors), Combinatorial Auctions. Chapter 3. Iterative Combinatorial Auctions*. MIT Press., 2006.

- [43] David C. Parkes. Auction design with costly preference elicitation. *Annals of Mathematics and Artificial Intelligence*, 44(3):269–302, 2005.
- [44] David C. Parkes and Lyle H. Ungar. Iterative combinatorial auctions: Theory and practice. In *AAAI/IAAI*, pages 74–81, 2000.
- [45] Tuomas Sandholm. Algorithm for optimal winner determination in combinatorial auctions. In *Artificial Intelligence. Preliminary version appeared in IJCAI-99*, volume 135, pages 1–54, 2002.
- [46] Ilya Segal. The communication requirements of social choice rules and supporting budget sets. *Journal of Economic Theory*, 136:341–378, 2007.
- [47] David Zuckerman. Linear degree extractors and the inapproximability of max clique and chromatic number. In *Proceedings of the thirty-eighth annual ACM symposium on Theory of computing*, pages 681–690, 2006.

## A Missing Proofs

**Lemma 1:** *A value query can be simulated by  $m$  marginal-value queries. A marginal-value query can be simulated by two value queries.*

*Proof.* The simulation of a marginal value query by its value queries is direct from the definition:  $v(j|S) = v(S \cup \{j\}) - v(S)$ . The simulation of a value query  $S$  by  $|S| \leq m$  marginal value queries is given by the equation  $v(S) = \sum_{j \in S} v(j|\{j' \in S | j' < j\})$ .  $\square$

**Lemma 2:** *A value query can be simulated by  $mt$  demand queries (where  $t = \log L$  is the number of bits needed to represent a single bundle value).*

*Proof.* We will show that demand queries can simulate any marginal value query  $v(j|S)$  using  $t$  queries, and then invoke the previous lemma. Set the prices of all the items in  $S$  to zero, and the prices of all other items (except  $j$ ) to  $\infty$ . Then, we perform a binary search on  $p_j$  to find its lowest value for which the bidder demands  $v(S)$ . It is straightforward to see that this price is indeed the marginal value of item  $j$ : at this price, the utilities from the bundles  $S$  and  $S \cup \{j\}$  are equal, thus  $v(S) - 0 = v(S \cup \{j\}) - p_j$  and the claim follows.

A binary search makes  $t$  demand queries, and  $m$  marginal value queries are needed to simulate a single value query thus  $v(S)$  can be simulated by  $mt$  demand queries.  $\square$

**Lemma 4:** *An indirect-utility query can be simulated by  $mt + 1$  demand queries. A demand query can be simulated by  $m + 1$  indirect-utility queries.*

*Proof.* An indirect-utility query with prices  $\vec{p}$  can be answered by first querying for the demand  $D$  under these prices and then simulating the value query  $v(D)$ .

The following algorithm uses  $m + 1$  indirect-utility queries to simulate a demand query with some price vector  $\vec{p}$ :

Initialization: start with the price vector  $\vec{p}$  for which the player answers some utility  $x$ .

Repeat: for every item  $i = 1, \dots, m$ , raise the price of item  $i$  by some  $\epsilon \in (0, 1)$ . If the answer to the indirect-utility query now is other than  $x$ , we decrease its price back by  $\epsilon$  in all future queries. If the answer was  $x$ , we use the new price for  $i$  in all future queries.

Finally: After all the  $m + 1$  indirect-utility queries are done, return the bundle of all items for which the answer was changed when we increased their prices.

In the algorithm above, if we raised the price of some item  $i$ , and the reported maximal-utility did not change, then there would clearly be utility-maximizing bundles that do not contain  $i$ , thus we can ignore this item. If the maximal-utility changed, then any utility-maximizing bundle under the current prices clearly contains  $i$ , thus we include it in our answer. Leaving the price of item  $i$  (of the first kind) at  $p_i + \epsilon$ , ensures that any bundle that contains it will not be output (but we are guaranteed to have other utility-maximizing bundles).  $\square$

**Lemma 5:** *Relative-demand queries can be simulated by a polynomial number of demand queries.*

*Proof.* For any  $\epsilon > 0$ , we simulate  $RD(\vec{p})$  by the following binary search (up to an  $\epsilon$ , see below):

Initialization: start with a price vector  $c\vec{p}$  ( $c > 0$ ).

Binary search: find with a binary search the value  $c^* \in \mathfrak{R}^+$  for which the bidder has a non-empty demand for the price vector  $c^* \cdot \vec{p}$  and the bidder demands the empty set for  $(c^* + \epsilon) \cdot \vec{p}$ .

Finally: return the bundle  $S$  demanded under the price vector  $c^* \cdot \vec{p}$ .

Now we show that for the price vector  $\vec{p}$ , every other bundle  $T$  has a smaller weight than  $S$  (up to  $\epsilon$ ), i.e.,

$$\frac{v(S)}{p(S)} \geq \frac{v(T)}{p(T)} - \epsilon. \quad (\text{A.1})$$

Denote  $c^* = \epsilon t$  for some  $t \in \mathfrak{R}^+$ . The bundle  $S$  was demanded under the prices  $\epsilon t \cdot \vec{p}$ , therefore  $v(S) - \epsilon t p(S) \geq 0$ . Thus,  $\frac{v(S)}{p(S)} \geq \epsilon t$ . Assume that inequality A.1 does not hold, then it follows that  $\frac{v(T)}{p(T)} - \epsilon > \epsilon t$ , or  $v(T) > \epsilon(t + 1)p(T)$ . But for the price vector  $(c^* + \epsilon)\vec{p} = \epsilon(t + 1)\vec{p}$  no bundle achieved a positive utility. Contradiction.  $\square$

## B Missing definitions

**Definition 3.** A valuation  $v_i$  satisfies the *substitutes* (or *gross-substitutes*) property if for every pair of item-price vectors  $\vec{q} \geq \vec{p}$  (coordinate-wise comparison), we have that the demand at prices  $q$  contains all items in the demand at prices  $p$  whose price remained constant. Formally, for every  $A \in \arg \max_S \{v(S) - \sum_{j \in S} p_j\}$ , there exists  $D \in \arg \max_S \{v(S) - \sum_{j \in S} q_j\}$ , such that  $D \supseteq \{j \in A \mid p_j = q_j\}$ .