Optimal Upper and Lower Approximation Bounds for k-Duplicates Combinatorial Auctions

Shahar Dobzinski Michael Schapira*

December 3, 2005

Abstract

We study the allocation problem in combinatorial auctions with k duplicates of each item. We exhibit a $\min(\frac{n}{k}, O(m^{\frac{1}{k+1}}))$-approximation algorithm that uses demand queries only. We prove the optimality of this upper bound by presenting an unconditional and tight lower bound in the communication model.

We also prove nearly tight upper and lower approximation bounds for the value oracles model, thus showing that it is strictly less powerful than the demand oracles model.

In particular, all our results apply to standard combinatorial auctions (where $k = 1$).

1 Introduction

In a combinatorial auction a set of items $M = \{1, ..., m\}$, is being sold to n bidders. The combinatorial nature of this problem is due to the fact that bidders value *bundles* of items rather than only assigning values to single items. Bidder i is defined by a valuation function $v_i : 2^M \rightarrow R_+$. We assume that each valuation function is monotone (for each $S \subseteq T \subseteq M$ $v(T) \geq v(S)$), and normalized ($v(\emptyset) = 0$). The *allocation problem* is to find a partition of the m items that maximizes the "total social welfare". I.e. to find a partition $S_1, ..., S_n$ of M for which $\Sigma_i v_i(S_i)$ is maximized.

Even though the size of the "input" is exponential in m (each v_i is described by 2^m real numbers) we would like our algorithm to run in time polynomial in the natural parameters of the problem, m and n. That is, we wish our algorithm's running time to be sublinear (poly-logarithmic) in the input size. Because of the exponential size of the input we assume that we are given an oracle for accessing each valuation. The oracle can only answer a specific predefined type of query. There are three kinds of queries commonly used (e.g. [3, 6, 7, 10, 4]):

*The School of Computer Science and Engineering The Hebrew University of Jerusalem {shahard,mikesch}@cs.huji.ac.il. Supported by grants from the Israel Science Foundation and the USA-Israel Bi-national Science Foundation.

1. **Value queries:** The query specifies a subset $S \subseteq M$ of items and receives the value $v_i(S)$ as the reply. This query is very natural from a "computer science" point of view, but in general is quite weak.

2. **Demand queries:** The query specifies a vector $p = (p_1 ... p_m)$ of "item prices", and the reply is the set that would be "demanded" by the queried bidder under these item prices. I.e., the subset S that maximizes $v_i(S) - \sum_{j \in S} p_j$. This query is natural from an economic point of view as it corresponds to the "revealed preferences" of the bidders (i.e. what is directly observable from their behavior). It is known that these queries can simulate value queries in polynomial time [4].

3. **General queries:** In this model we allow *any* kind of query to each valuation function (but the query is always to a single valuation function). This model captures the communication complexity between the bidders. Due to its strength this model is mostly interesting for lower bounds.

The hardness of the the allocation problem is almost completely understood: achieving an approximation ratio of $min\{n^{1-\epsilon}, O(m^{\frac{1}{2}-\epsilon})\}$ for any $\epsilon > 0$ requires an amount of communication exponential in m and n [13]. In fact, even if the input can be succinctly described the problem is still NP-hard to approximate within this factor [16], assuming $NP \neq ZPP$.

A lot of effort has been invested in finding and solving subcases that can be better approximated. Examples include auctions in which we assume the valuations are restricted to some set of functions (e.g. subadditive [6], submodular [10, 7]), and auctions in which multiple units of each item exist [3, 2, 5].

In the latter case, there are k duplicates of each item. Every bidder is interested in at most one unit of each item and is therefore still defined by his value for all subsets of M. Obviously, k-duplicates combinatorial auctions are a generalization of "standard" combinatorial auctions (in which k=1).

In this paper we improve upon the known upper and lower bounds for k-duplicates combinatorial auctions in the demand oracles model. We also study the approximability of such combinatorial auctions in the value oracles model.

Several upper bounds for k-duplicates combinatorial auctions are known: for the demand oracles model, [3] proves an $O(km^{\frac{1}{k-1}})$ upper bound. In particular, an approximation factor of only $O(\log m)$ is achieved for combinatorial auctions with $\Omega(\log m)$ duplicates. For the special case $k = 1$, [4], independent of our work, proved an upper bound of $O(\sqrt{m})$ for approximation with demand queries. For the very restricted case of single minded bidders, [2] shows a $(1 + \epsilon)$-approximation for $k = \Omega(\log m)$[1].

We present the first algorithm for k-duplicates combinatorial auctions that achieves a $\min\{\frac{n}{k}, O(m^{\frac{1}{k+1}})\}$-approximation. Independent of our work, [5] presented an algorithm that achieves the same approximation ratio. However, while their upper bound requires general queries, our upper bound even holds for the

[1]Both the algorithm of [3] and the algorithm of [2] are incentive compatible. In this paper we do not address incentive compatibility issues.

demand oracles model. We also show that when $k = \Omega(\log m)$ one can achieve a $(1+\epsilon)$-approximation, for every $\epsilon > 0$.

Theorem: There is a polynomial time $\min\{\frac{n}{k}, O(m^{\frac{1}{k+1}})\}$-approximation algorithm for k-duplicates combinatorial auctions that uses demand queries only.

Our algorithm is based on solving the linear relaxation of the problem, and then applying subtle randomized rounding. We show that our upper bound is tight by proving the following theorem:

Theorem: For every $\epsilon > 0$, achieving an approximation ratio better than $\min\{(\frac{n}{k})^{1-\epsilon}, O(m^{\frac{1}{k+1}-\epsilon})\}$ requires exponential communication. This lower bound also applies for both randomized and nondeterministic settings.

In [3] a $O(km^{\frac{1}{k-1}})$ a computational lower bound of $O(m^{\frac{1}{k+1}-\epsilon})$, that applies to demand queries, is proved. This lower bound is conditional on $P \neq NP$, $NP \neq ZPP$ and $NP \not\subseteq DTIME(n^k)$. We improve upon this lower bound by using communication complexity tools to derive a lower bound that is unconditional in any computational assumptions and applies to all kinds of queries. In the communication model we assume all bidders have unlimited computational power and we only measure the amount of communication (bits) transferred between them. For a good introduction to communication complexity see [9]. Our lower bound generalizes the lower bound proved in [11, 13] for the case $k = 1$.

We show that value oracles are strictly less powerful than demand oracles by exhibiting nearly tight upper and lower approximation bounds:

Theorem: There is an algorithm for approximating k-duplicates combinatorial auctions that achieves an $O(\frac{m}{\sqrt{\log m}})$-approximation using polynomial number value queries only.

In [8] an $O(\frac{m}{\sqrt{\log m}})$-approximation algorithm that uses value queries only is shown for combinatorial auctions in which $k = 1$. We show that this result holds for all possible values of k. The proof of this theorem also yields a simpler proof of the result of [8].

Theorem: Using a polynomial number of value queries it is impossible to approximate a combinatorial auction with k-duplicates to a factor of $O(\frac{m}{\log m})$[2].

This lower bound holds even if the auctioneer is computationally unlimited (i.e. even if the algorithm may run in exponential time).

We note that when using demand oracles the optimal approximation ratio gets significantly better as k gets bigger. Interestingly, this is not the case when using value queries only.

[2]For the special case $k = 1$, this lower bound was independently proved by [4].

	Upper Bound	Lower Bound
Demand Queries	$\min(\frac{n}{k}, O(m^{\frac{1}{k+1}}))$, section 2	$\min((\frac{n}{k})^{1-\epsilon}, O(m^{\frac{1}{k+1}-\epsilon}))$, section 3
Value Queries	$O(\frac{m}{\sqrt{\log m}})$, section 4.1	$O(\frac{m}{\log m})$, section 4.2

Table 1: The table summarizes the approximation bounds known for k-duplicates combinatorial auctions.

2 An $O(m^{\frac{1}{k+1}})$ Upper Bound

Randomized rounding of an LP-relaxation of a problem is a common technique for designing approximation algorithms. Applying this technique for combinatorial auctions poses some difficulties: first, we have an exponential number (in n and m) of variables in the LP formulation. Second, a straightforward randomized rounding in which a random experiment is conducted for every bidder i and bundle S (similar to the one used in [15] for the set packing problem) fails to give good results. This is due to the limitation that each bidder is assigned at most one bundle.

The basic idea at the heart of our algorithm is a more subtle randomized rounding method in which a single independent random experiment is conducted for each bidder. However, a randomized choice will very likely yield unfeasible solutions, in the sense that too many items will be allocated. Thus, the probabilities taken are reduced by an $O(m^{\frac{1}{k+1}})$ factor, which leads to a reduction by a similar factor in the expected quality of solution. Let us now turn to the description of the algorithm:

Input: n demand oracles for the n valuations v_i.

Output: An allocation which is a $\min\{\frac{n}{k}, O(m^{\frac{1}{k+1}})\}$ approximation to the optimal allocation.

The Algorithm: We first describe the basic steps of the algorithm and then discuss important details necessary for its implementation.

1. Solve the linear relaxation of the problem:

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

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

2. Use careful randomized rounding to create a feasible allocation of the items $\{S_1, ..., S_n\}$ in which each good is assigned to at most k bidders.

3. Let I be the set of k bidders who offer the highest bids on M. If $\Sigma_{t\in I} v_t(M) > \Sigma_i v_i(S_i)$, assign M to each bidder in I. Otherwise, allocate according to $\{S_1, ..., S_n\}$.

Even though the linear program has an exponential number of variables, it may still be solved in polynomial time. This is done by solving the dual linear program using the ellipsoid method. The ellipsoid method requires a "separation" oracle, and this may be directly implemented using the demand oracles of the bidders. This was shown in [13] and in more details in [4].

The randomized rounding step is implemented as follows: let $\gamma = \Theta(m^{\frac{1}{k+1}})$ (to be chosen later). For each bidder i we independently choose a set S_i by performing the following random experiment: each set S is chosen with probability $\frac{x_{i,S}}{\gamma}$, and the empty set is chosen with probability $(1 - \frac{\Sigma_S x_{i,S}}{\gamma})$. Notice that the independence is across the bidders and that this randomized rounding technique ensures that every bidder is assigned at most one bundle of items. If any of the items was assigned to more than k bidders, then this stage is repeated from scratch. This randomized step may be converted to be deterministic by derandomizing using the generator of [12] as explained in [17].

Theorem 2.1 *The algorithm produces an allocation that is a* $\min\{\frac{n}{k}, O(m^{\frac{1}{k+1}})\}$*-approximation to the optimal allocation.*

Proof: First, notice that a $\frac{n}{k}$-approximation is easily achieved by giving M to each bidder $i \in I$.

Without loss of generality, we assume that $max_i\{v_i(M)\} = 1$ (otherwise we can simply divide all valuations by $max_i\{v_i(M)\}$).

The first step of the algorithm returns the optimal fractional solution $OPT^* = \Sigma_{i,S} x_{i,S} v_i(S)$, which is an upper bound on the value of the optimal allocation, OPT.

If $OPT^* \leq 3\gamma$ the algorithm ensures $O(m^{\frac{1}{k+1}})$ approximation by assigning M to the highest bidder ($max_i\{v_i(M)\} = 1$). Otherwise, we show that for $OPT^* > 3\gamma$ the allocation generated by the randomized rounding step is w.h.p. a $O(m^{\frac{1}{k+1}})$ approximation to the optimal one.

We prove that with constant probability the second step produces an allocation $\{S_1, ..., S_n\}$ in which $\Sigma_i v_i(S_i) \geq \frac{OPT^*}{3\gamma}$ and no item is assigned to more than k bidders. Let us denote the event that item j was assigned to more than k bidders by E_j and the event that $\Sigma_i v_i(S_i) < \frac{OPT^*}{3\gamma}$ by E_{m+1}. We are interested in showing that $Pr[\vee_{t=1}^{m+1} E_t] \leq c$ with $0 < c < 1$.

Define Y_i^j to be the random variable that indicates whether item j was assigned to bidder i, and $Y^j = \Sigma_i Y_i^j$. Due to the randomized rounding method we used $\{Y_i^j\}_i$ are independent (for j) and take values in $\{0, 1\}$. We have that $E[Y_i^j] = Pr[Y_i^j = 1] = \Sigma_{S|j \in S} \frac{x_{i,S}}{\gamma}$. Hence, $E[Y^j] = \Sigma_i E[Y_i^j] = \Sigma_{i,S|j\in S} \frac{x_{i,S}}{\gamma} \leq \frac{k}{\gamma}$ (the last inequality is based on the first condition of the LP formulation). Observe that $Pr[E_j] = Pr[Y^j \geq k+1]$. We use the following result from the Chernoff-Hoeffding Bounds (see e.g. [18]):

Claim: Let $X_1, ...X_l$ be independent random variables that take values in $[0, 1]$ and $R = \Sigma_{s \in [l]} X_s$. If $\mu = E[X]$ and $\delta \geq 0$ it holds that $Pr[R \geq \mu(1+\delta)] \leq (\frac{e}{(1+\delta)})^{(1+\delta)\mu}$

By applying the claim on $\{Y_i^j\}_i$ and Y^j, and for $\gamma = \Theta(m^{\frac{1}{k+1}})$ we can upper bound $Pr[E_j] = Pr[Y^j \geq k+1]$ by $\frac{1}{20m}$ (for example).

Now, define Z_i to be the random variable that receives the value $v_i(S_i)$ after the randomized rounding step and $Z = \Sigma_i Z_i$. Observe, that $\{Z_i\}_{i\in[n]}$ are independent (due to the randomized rounding method) and take values in $[0,1]$. $E[Z_i] = \frac{\Sigma_S x_{i,S} v_i(S)}{\gamma}$ and hence $E[Z] = \frac{\Sigma_{i,S} x_{i,S} v_i(S)}{\gamma} = \frac{OPT^*}{\gamma}$. We make use of the following corollary from Chebyshev's inequality (see e.g. [18]):

Claim: Let X be the sum of independent random variables, each of which lies in $[0,1]$, and let $\mu = E[X]$. Then, for any $\alpha >> 0$, $Pr[|X-\mu| \geq \alpha] \leq \frac{\mu}{\alpha 2}$.

$Pr[E_{m+1}] = Pr[Z < \frac{OPT^*}{3\gamma}] \leq Pr[|Z - \frac{OPT^*}{\gamma}| \geq \frac{2OPT^*}{3\gamma}] \leq \frac{9\gamma}{4OPT^*} \leq \frac{3}{4}$
(the last inequality is because $OPT^* > 3\gamma$). Therefore, using the Union Bound:

$$Pr[\vee_{t=1}^{m+1} E_t] \leq \Sigma_{j\in M} Pr[E_j] + Pr[E_{m+1}] \leq \frac{1}{20} + \frac{3}{4} = 0.8$$

and the theorem follows. □

Theorem 2.2 *For any $\epsilon > 0$, there is a $(1+\epsilon)$-approximation algorithm for combinatorial auctions with $\Omega(\log m)$-duplicates.*

Proof: Proof of this theorem is based on the randomized rounding technique presented earlier with ideas from [14] and theorem 2.1. Full proof will appear in the full version of this paper. □

3 An $O(m^{\frac{1}{k+1}-\epsilon})$ Communication Lower Bound

We will prove our lower bound by reduction from the approximate disjointness problem. In this problem, there are r bidders, and each bidder i holds a string A_i which specifies a subset of $\{1,...,t\}$. The bidders' goal is to distinguish between the following two extreme cases:

- $\cap_{i=1}^r A_i \neq \emptyset$
- for every $i \neq j$, $A_i \cap A_j = \emptyset$

The communication complexity of this problem is $\Omega(\frac{t}{r^4})$ [1]. This result holds also for randomized protocols with bounded 2-sided error, and for non-deterministic protocols.

Theorem 3.1 *Any protocol for approximating the social welfare of a k-duplicates combinatorial auction to a factor of $\min\{(\frac{n}{k})^{1-\epsilon}, O(m^{\frac{1}{k+1}-\epsilon})\}$, for any $\epsilon > 0$, requires exponential communication. This lower bound also applies for randomized and non-deterministic settings.*

Proof: We prove a communication lower bound for distinguishing between the case the social welfare is k and the case it is rk with $r = O(m^{\frac{1}{k+1}-\epsilon})$ in a combinatorial auction with $n = rk$ bidders. We will show a reduction from the approximate-disjointness problem on vectors of size $t = \frac{e^{m/(k+1)r^{(k+1)}}}{r}$.

We will create a set $F = \{P^s\}_{s=1...t}$, where each P^s is a partition of M into r disjoint subsets $\{P_1^{s_i}, ..., P_r^{s_i}\}$. This set of partitions will have the following property:

Definition 3.2 *A set of partitions $F = \{P^s\}_{s=1...t}$ is said to have the $(k+1)$-wise intersection property if for every choice of indices $1 \leq s_1, s_2, ...s_{k+1} \leq t$ such that no two are equal, and $1 \leq l_1, l_2, ..., l_{k+1} \leq r$, it holds that $\cap_{i=1}^{k+1} P_{l_i}^{s_i} \neq \emptyset$.*

Lemma 3.3 *There exists a set F of partitions with the $(k+1)$-wise intersection property of size $|F| = t = \frac{e^{m/(k+1)r^{(k+1)}}}{r}$.*

Proof: We use probabilistic construction to obtain such a set: Each partition P^s will be chosen independently at random (each element will be independently and randomly placed in one of the P_i^s). Fix indices: $1 \leq s_1, s_2, ..., s_{k+1} \leq t$ (s.t. no two are equal) and $1 \leq l_1, l_2, ..., l_{k+1} \leq r$. We have that:

$$Pr[\cap_{i=1}^{k+1} P_i^{s_i} = \emptyset] = \Pi_{j=1}^{m} Pr[j \notin \cap_{i=1}^{k+1} P_{l_i}^{s_i}]$$

$$= (1 - 1/r^{k+1})^m \leq e^{-m/r^{(k+1)}}$$

Since there are at most $r^{k+1}t^{k+1}$ choices of such indices we get that as long as $r^{k+1}t^{k+1} < e^{m/r^{(k+1)}}$ such a set of partitions exists. □

We are now left with describing the reduction. Assume an instance of the approximate-disjointness problem on vectors of size $t = \frac{e^{m/(k+1)r^{(k+1)}}}{r}$, in which bidder i receives the string $A_i \subseteq \{1...t\}$. We reduce it into a k-duplicates combinatorial auction with $n = rk$ bidders in the following manner:

Let the goods be $\{1...m\}$, and assume that there are k duplicates of each good. Each bidder will be defined by a pair (i, d), $i \leq i \leq r$, $1 \leq d \leq k$. Bidder (i, d) will construct the collection $B_i = \{P_i^s | s \in A_i\}$. Notice, that the valuation of a bidder (i, d) depends only on i:

$$v_{i,d}(S) = \begin{cases} 1, & \exists l \textit{ such that}: \ l \in B_i \textit{ and } l \subseteq S \ ; \\ 0, & \text{otherwise.} \end{cases}$$

Observe that if there exists $s \in \cap A_i$, then the total social welfare is rk. On the other hand, if for every $i \neq j$, $A_i \cap A_j = \emptyset$ then the total welfare is at most k. The second observation is due to the fact that social welfare cannot be higher than k since F has the $(k+1)$-wise intersection property, and there are only k duplicates of each commodity.

Since the communication complexity of the approximate-disjointness problem is $\Omega(t/r^4)$, in our case it is $\Omega(e^{m/(k+1)r^{(k+1)}-5ln(r)})$. In particular, as long as $r = O(m^{\frac{1}{k+1}-\epsilon})$ the communication complexity is exponential. The theorem follows. □

4 Approximation Bounds Using Value Queries

In this section we present a polynomial time algorithm that achieves and approximation ratio of $O(\frac{m}{\sqrt{\log m}})$ using value queries only. We also prove that any algorithm that uses a polynomial number of value queries cannot achieve an approximation factor within $O(\frac{m}{\log} m)$.

Thus, we show that value queries are strictly less powerful that demand queries.

4.1 An Upper Bound

In [8] an upper bound of $O(\frac{m}{\sqrt{\log m}})$ is proved for approximating a standard combinatorial auction using value queries only. In this section we generalize their algorithm, and prove the same results for the k-duplicates setting. As a by product, a simpler proof of the upper bound of [8] is obtained.

Input: n value oracles for the n valuations v_i.

Output: An allocation which is an $O(\frac{m}{\sqrt{\log m}})$ approximation to the optimal allocation.

The Algorithm:

1. Partition M into $\log m$ bundles $B_1, ..., B_{\log m}$, each of size $\frac{m}{\log m}$.
2. Query each bidder i for his value on all possible combinations of bundles in $B_1, ..., B_{\log m}$. I.e. for all $J \subseteq [\log m]$, query $v_i(\cup_{t \in J} B_t)$.
3. Find the partition of the km items' duplicates, $S_1, ..., S_n$, such that each S_i is a combination of bundles in $\{B_s\}_s$, and $\Sigma_i v_i(S_i)$ is maximized. Allocate accordingly.

Theorem 4.1 *The algorithm provides a $O(\frac{m}{\sqrt{\log m}})$ approximation to the optimal allocation, and uses a polynomial number of value queries.*

Proof: First, notice that the algorithm queries each bidder $2^{|\{B_1,...,B_{\log m}\}|} = 2^{\log m} = m$ times, and thus the total number of value queries is $m \cdot n$. Also notice that step 3 of the algorithm can be executed in polynomial time by using dynamic programming.

We will now prove that the algorithm provides the guaranteed approximation ratio. Denote the optimal allocation OPT by $T_1, ..., T_r, Q_1, ..., Q_l$, where for each t, $1 \le t \le r$, $|T_t| \ge \sqrt{\log m}$, and for each q, $1 \le q \le l$, $|Q_q| < \sqrt{\log m}$. Let $|OPT| = \Sigma_{t=1}^r v_t(T_t) + \Sigma_{q=1}^l v_q(Q_q)$.

The first case we consider is when $\Sigma_{t=1}^r v_t(T_t) \ge \frac{|OPT|}{2}$ (i.e. "large" bundles contribute most of the social welfare). Let $I \subseteq [r]$ be the set of indices of the k bidders who offer the highest bids on M. Since $r \le \frac{mk}{\sqrt{\log m}}$ (otherwise, too many items where allocated), it holds that $\Sigma_{i\in I} v_i(T_i) \geq \frac{k\sqrt{\log m}|OPT|}{2km} = \frac{\sqrt{\log m}|OPT|}{2m}$. From the monotonicity of the valuations, for each $i \in I$ $v_i(M) \geq v_i(T_i)$. Allocating all items to each bidder in I is one of the allocations over which the algorithm maximizes, and this case is completed.

The second case is when $\Sigma_{q=1}^{l} v_q(Q_q) > \frac{|OPT|}{2}$. Without loss of generality, assume that $v_1(Q_1) \geq v_2(Q_2) \geq ... \geq v_l(Q_l)$. Define $B = \{B_{(r,t)}\}_{1\leq r\leq \log m, 1\leq t\leq k}$, such that for every r and t, $B_{(r,t)} = B_r$. We will now describe a process which will generate a partition of B, and show that it provides the required approximation ratio. At the beginning, let N be the set of all bidders, and $R = B$ be the set of all the bundles in B which haven't been allocated yet.

1. Let $q = \min\{i | i \in N, \exists S \text{ s.t. } S \subseteq R \text{ and } Q_i \subseteq (\cup_{T\in S} T)\}$. let S_q be the minimal cardinality subset for which $Q_q \subseteq (\cup_{T\in S_q} T)$. If there is no such q, exit.

2. Allocate $\cup_{T\in S_q} T$ to bidder q.

3. Let $R = R \backslash S_q$, $N = N \backslash q$.

4. Return to step 1.

The allocation generated by the process is certainly one possible outcome of the algorithm. We will now see that $\Sigma_i v_i(\cup_{T\in S_i} T) \geq \frac{\sqrt{\log m}|OPT|}{2m}$.

At the first iteration of the process, the chosen q will be 1. Since for every i, $v_1(Q_1) \geq v_i(Q_i)$, and $v_1(\cup_{T\in S_1} T) \geq v_1(Q_1)$, it holds that $v_1(\cup_{T\in S_1} T) \geq \frac{\sqrt{\log m}}{m} \Sigma_{q=1}^{\frac{m}{\sqrt{\log m}}} v_q(Q_q)$.

$|Q_1| \leq \sqrt{\log m}$ therefore clearly $|\cup_{T\in S_1} T| \leq \frac{m}{\sqrt{\log m}}$. Hence, it must hold for the bidder q' chosen at the second iteration that $2 \leq q' \leq \frac{m}{\sqrt{\log m}} + 1$. Using a similar argument to the one we used in the first iteration, we can see that $v_1(\cup_{T\in S_1} T) + v_{q'}(\cup_{T\in S_{q'}} T) \geq \frac{\sqrt{\log m}}{m} \Sigma_{q=1}^{\frac{2m}{\sqrt{\log m}}} v_q(Q_q)$.

We repeat this argument until the process ends, and get that $\Sigma_i v_i(\cup_{T\in S_i} T) \geq \frac{\sqrt{\log m}|OPT|}{2m}$. The theorem follows. □

Remark 4.2 *Instead of partitioning M to bundles of size $\log m$, we could have partitioned it into bundles of an arbitrary size t. The approximation guarantee would have been $\frac{m}{\sqrt{t}}$, and the number of value queries would have been $n \cdot 2^t$.*

4.2 A Lower Bound

Theorem 4.3 *Using a polynomial number of value queries it is impossible to approximate a combinatorial auction with k-duplicates to a factor of $O(\frac{m}{\log m})$. In particular, this bound holds for the special case where $k = 1$.*

Proof: We will construct a combinatorial auction with k-duplicates with $n = k \cdot \frac{m}{t}$ bidders. Each bidder will be defined by the pair (i, d), $1 \leq i \leq n$, $1 \leq d \leq k$. The valuation of bidder (i, d) depends only on i:

$$v_{(i,d)}(\mathrm{S}) = \begin{cases} 1, & S_i \subseteq S; \\ 1, & |S| \geq \frac{m+1}{2}; \\ 0, & \text{otherwise.} \end{cases}$$

where $S_1, ..., S_{\frac{m}{t}}$ is some partition of M to bundles of size t. It is easy to see that the optimal total social welfare is $k\frac{m}{t}$. A total social welfare of at most $2k$ is easily achieved by giving any subsets of size $\frac{m}{2} + 1$ to as many bidders possible. To complete the proof, we will now show that finding a subset S, $|S| \leq \frac{m}{2}$, and a bidder (i, d) such that $v_{(i,d)}(S) = 1$ can not be done in polynomial number of value queries.

Without loss of generality, assume that all queries are on bundles of size $\frac{m}{2}$. Indeed, the value of bigger bundles is known in advance, and for each smaller bundle there is a bundle of size $\frac{m}{2}$ that contains it, and so has at least its value. Observe that every query to a bundle of size $\frac{m}{2}$, rules out at most $\binom{\frac{m}{2}}{t}\frac{(m-t)!}{t!^{\frac{m}{t}-1}}(\frac{m}{t}-1)!$ allocations of M. However, there are $\frac{m!}{t!^{\frac{m}{t}}}(\frac{m}{t})!$ such possible allocations. We use a simple adversary argument and find that the number of queries required to find a bundle of value 1 is:

$$\frac{\frac{m!}{t!^{\frac{m}{t}}}(\frac{m}{t})!}{\binom{\frac{m}{2}}{t}\frac{(m-t)!}{t!^{\frac{m}{t}-1}}(\frac{m}{t}-1)!} \geq \frac{m!}{\binom{\frac{m}{2}}{t}(m-t)!t!}$$

$$= \frac{m!(\frac{m}{2}-t)!}{(m-t)!(\frac{m}{2})!} = \frac{(m-t+1)\cdot(m-t+2)\cdot \ldots \cdot m}{(\frac{m}{2}-t+1)\cdot(\frac{m}{2}-t+1)\cdot \ldots \cdot(\frac{m}{2}-t)} \geq 2^t$$

By choosing $t = \Omega(\log^{1+\epsilon} m)$, for any $\epsilon > 0$, the number of value queries needed is not polynomial. In addition, for any algorithm that uses a polynomial number of queries, one can choose $t = \Theta(\log m)$, with a proper choice of the constants inside the Θ, to show (by adversary argument) that it doesn't achieve the required approximation. □

Acknowledgements

We warmly thank Noam Nisan for his guidance and for very interesting discussions. We would also like to thank Yair Bartal for helpful discussions, and Liad Blumrosen for comments on a draft of this paper.

References

[1] N. Alon, Y. Matias, and M. Szegedy. The space complexity of approximating the frequency moments. *Proc. 28th ACM Symp. on Theory of Computing*, pages 20–29, May 1996.

[2] 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.

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

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

[5] Patrick Briest, Piotr Krysta, and Berthold Voecking. Approximation techniques for utilitarian mechanism design. To appear in STOC 2005.

[6] Shahar Dobzinski, Noam Nisan, and Michael Schapira. Approximation algorithms for combinatorial auctions with complement-free bidders. To appear in STOC 2005.

[7] Shahar Dobzinski and Michael Schapira. An improved appoximation algorithm for combinatorial auctions with submodular bidders, 2005. Working paper.

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

[9] Eyal Kushilevitz and Noam Nisan. *Communication Complexity*. Cambridge University Press, 1997.

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

[11] Noam Nisan. The communication complexity of approximate set packing and covering. In ICALP 2002.

[12] Noam Nisan. $RL \subseteq SC$. In *Proceedings of the twenty-fourth annual ACM symposium on Theory of computing*, pages 619–623. ACM Press, 1992.

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

[14] Prabhakar Raghavan. Probabilistic construction of deterministic algorithms: Approximating packing integer programs. *J. Comput. Syst. Sci.*, 37(2):130–143, 1988.

[15] Prabhakar Raghavan and Clark D. Tompson. Randomized rounding: A technique for provably good algorithms and algorithmic proofs. *Combinatorica*, 7:365–374, 1987.

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

[17] D. Sivakumar. Algorithmic derandomization via complexity theory. In *Proceedings of the 34th annual ACM Symposium on Theory of Computing (STOC)*, 2002.

[18] Aravind Srinivasan. *Approximation Algorithms Via Randomized Rounding: A Survey. In M. Karonski and H. J. Promel (Editors), Lectures on Approximation and Randomized Algorithms.* Polish Scientific Publishers PWN, 1999.