Combinatorial Auctions: VC *v.* VCG

Elchanan Mossel* Christos Papadimitriou† Michael Schapira‡ Yaron Singer§

Abstract

The existence of incentive-compatible, computationally-efficient protocols for combinatorial auctions with decent approximation ratios is one of the most central and well studied open questions in mechanism design. The *only* universal technique known for the design of truthful mechanisms is the celebrated *Vickrey-Clarke-Groves (VCG)* scheme, which is "*maximal in range*", i.e., it always exactly optimizes over a subset of the possible outcomes. We present a first-of-its-kind technique for proving *computational-complexity* inapproximability results for maximal-in-range mechanism for combinatorial auctions (under the complexity assumption that NP has no polynomial circuits). We show that in some interesting cases the lower bounds obtained using this technique can be extended to hold for *all* truthful mechanisms. Our lower-bounding method is based on a generalization of the VC-dimension to k-tuples of disjoint sets. We illustrate our technique via the case of two-bidder combinatorial auctions. We believe that this technique is of independent interest, and has great promise for making progress on the general problem.

*Statistics and Computer Science, U.C. Berkeley, and Mathematics and Computer Science Weizmann Institute. Supported by Sloan fellowship in Mathematics, NSF Career award DMS 0548249, DOD grant N0014-07-1-05-06, and by ISF. mossel@stat.berkeley.edu

†Computer Science Division University of California at Berkeley, CA, 94720 USA. christos@cs.berkeley.edu

‡Department of Computer Science, Yale University, CT, USA, and Computer Science Division, University of California at Berkeley, CA, USA. Supported by NSF grant 0331548. michael.schapira@yale.edu.

§Computer Science Division University of California at Berkeley, CA, 94720 USA. Supported by grants XXXXX. yaron@cs.berkeley.edu.

1 Introduction

The field of *algorithmic mechanism design* [35] is about the reconciliation of bounded computational resources and strategic interaction between selfish participants. Traditional computer science theory handles intractable optimization problems by designing *approximation algorithms*. Mechanism design theory in economics handles strategic conflicts by designing *mechanisms* that incent agents to behave truthfully. Algorithmic mechanism design aims to bring these two research agendas together by asking the following: "Are algorithms that are *both* computationally-efficient *and* incentive-compatible *less powerful* than algorithms that only achieve one of these two desiderata?".

The only universal technique in mechanism design for providing agents with incentives to truthfully reveal their preferences are the celebrated *Vickrey-Clarke-Groves (VCG) mechanisms* [44, 9, 22].[1] However, in many settings the application of VCG is computationally intractable, as VCG necessarily involves exact optimization (approximations do not suffice [34]). To overcome the intractability of VCG mechanisms in such cases, the class of VCG-based mechanisms, called *maximal-in-range mechanisms*, was devised. Informally, maximal-in-range mechanisms do not necessarily output the globally optimal solution, but they do always *exactly* optimize over a fixed *subset* of the possible outcomes.

Maximal-in-range mechanisms are essentially the only general method known for the design of (deterministic) computationally-efficient and truthful mechanisms. Indeed, with very few exceptions [5, 28, 30] the best known computationally-efficient and truthful approximation-algorithms for *multi-parameter* mechanism design settings are obtained via maximal-in-range mechanisms [35, 8, 12, 13, 24, 41]. In fact, in several interesting cases, it is known that an algorithm is truthful *if and only if* it is maximal-in-range [7, 16, 36]. For this reason, understanding the limitations of maximal-in-range mechanisms, in terms of approximability, is an important research agenda.

We tackle this issue in the context of the paradigmatic problem of algorithmic mechanism design – *combinatorial auctions*. In a combinatorial auction, a set of items is to be sold to bidders with private preferences over *subsets* of the items. We present a first-of-its-kind technique for proving computational-complexity inapproximability results for maximal-in-range mechanisms for combinatorial auctions. We demonstrate its use by proving lower bounds for the two-bidder case[2]. We show that for some well-studied auction settings our results actually apply to *all* truthful algorithms. With the exception of a result in [26], these are the first such computational-complexity results for combinatorial auctions.

Dobzinski and Nisan [11] initiated the study of the inapproximability of maximal-in-range mechanisms for combinatorial auctions by taking a *communication complexity* [46, 25] approach — in this approach it is assumed that each bidder has an *exponentially large string of preferences*. However, in many real-life settings this assumption is problematic. In contrast, our intractability results deal with bidder preferences that are *succinctly described*, and therefore they relate to *computational complexity*.

For many interesting combinatorial auction settings good (constant-) approximation algorithms are known [3, 29, 13, 15, 17, 18, 45], yet the best known *truthful* approximation algorithms fail to obtain such approximation ratios [13, 14, 10]. This is commonly thought to be due to an inherent clash between the truthfulness and computational-efficiency requirements, that manifests itself in degraded algorithm performance. However, *the community currently lacks the machinery to prove this widely-held belief.* We believe that our approach is quite powerful, and holds great promise for

[1] For interesting *single-parameter* settings [4], techniques other than VCG are known (see, e.g., [20, 21]).

[2] We use the term lower bound as a general reference to an inapproximability result. Hence, a lower bound of $\frac{1}{2}$ means (as we are looking at a maximization problem) that no approximation better than $\frac{1}{2}$ is possible. This use is similar to that of Hastad in [23]

making progress on the general problem (we discuss this further in the open questions section). We also believe that this approach can be used for proving computational-complexity lower bounds for many of the auction settings studied in economics and computer-science literature.

From a technical perspective, our technique involves lower bounding the *VC dimension of collections of partitions*.[3] We consider a class of succinctly representable valuations and show that if a maximal-in-range mechanism approximates closely the optimal social welfare for these valuations, then it is implicitly solving *optimally* a smaller, but still relatively large, optimization problem of the same nature — an NP-hard feat. We establish this by showing that the subset of outcomes considered by the assumed approximate maximal-in-range mechanism must "shatter" a relatively large subset of the items.

The connection between VC dimension and maximal-in-range mechanisms underlies the work in [36]. [36] presented a general (i.e., not restricted to maximal-in-range mechanisms) inapproximability result, albeit in the context of a different mechanism design problem [19], called *combinatorial public project*. Like our analysis, the analysis in [36] was also based on lower bounding the VC dimension of the range of maximal-in-range mechanisms. However, we stress that the analysis in [36] was carried out within the standard VC framework, and so it relied on *existing* machinery, namely the Sauer-Shelah Lemma [39, 42] and its probabilistic version due to Ajtai [1]. In contrast, to handle the unique technical challenges posed by the combinatorial auctions setting, our results require *new* machinery. Our current technique can be interpreted as *an extension of the Sauer-Shelah Lemma to the case of partitions* (Lemma 1 in Section 4). We believe that further refinement and development of our technique are the key to many more such results.

1.1 The Auction

We prove inapproximability results for the following simple setting (used only to illustrate our technique): There are 2 bidders $1, 2$ and m items $1, ..., m$. Each bidder i has a *private valuation function* v_i (referred to as a valuation in short) that assigns a non-negative real value to every subset of the items. $v_i(S)$ can be regarded as i's maximum willingness to pay for the bundle of items S. The valuation v_i of each bidder i is an *additive valuation with a budget constraint* [29], defined as follows: i is assumed to have a *private value* for each item j, v_{ij}, and a *private budget* B_i. i's value for every bundle S is then $v_i(S) = \min\{\Sigma_{j\in S}\ v_{ij}, B_i\}$. The objective is to find a partition of the items into two disjoint subsets (S_1, S_2), for which the *social welfare*, i.e., $\Sigma_i\ v_i(S_i)$, is maximized.

A maximal-in-range mechanism M for this auction is defined by a collection R_M of partitions of items between the bidders (we formally define maximal-in-range mechanisms in Section 2). Rather than trying to find a partition of the items that maximizes the social-welfare out of *all* possible partitions, a maximal-in-range mechanism M always outputs the optimal partition in R_M (which is a subset of the collection of all possible partitions)[4]. That is, for every (v_1, v_2), M outputs the partition in R_M that maximizes the social welfare, i.e., $argmax_{(T_1,T_2)\in R_m}\ \Sigma_i v_i(T_i)$. We refer to R_M as M's *range*. Maximal-in-range mechanisms are known to always be truthful. As finding the solution for our auction setting that optimizes the social welfare is NP-hard [29, 3], the naive optimal maximal-in-range mechanism in which R_M contains *all* possible partitions of items is computationally intractable.

How well can we approximate the social welfare in a truthful manner? Observe that if we

[3]By a *partition* in this paper we mean an ordered pair of disjoint subsets which, however, may or may not exhaust the universe.

[4]Maximal-in-range mechanisms are a special case of a more general class of mechanisms called "affine maximizers" [38, 26]. All of the results in this paper actually apply to this more general class.

set $R_m = \{([m], \emptyset), (\emptyset, [m])\}$, that is, if M always assigns all items to one bidder, this results in a computationally efficient $\frac{1}{2}$-approximation mechanism. Is there a maximal-in-range mechanism that obtains a better approximation ratio? We prove the following theorem:

Theorem: *For any constant $\epsilon > 0$, no polynomial-time maximal-in-range mechanism obtains an approximation ratio of $\frac{3}{4} + \epsilon$ unless NP $\subseteq$ P/poly.*

We note that a non-truthful PTAS for our problem exists [3]. Hence, this result shows that if we are not only interested in computational efficiency, but *also* in truthfulness, then the standard technique for achieving truthfulness in auctions (VCG) leads to solutions of much lesser quality.

For the case in which we insist that all items be allocated (that is, the auctioneer is not allowed not to sell a certain item) we can improve our lower bound to $\frac{1}{2}$ (matching the trivial upper bound presented above).

Theorem: *For the allocate-all-items case, for any constant $\epsilon > 0$, no polynomial-time maximal-in-range mechanism obtains an approximation ratio of $\frac{1}{2} + \epsilon$ unless NP $\subseteq$ P/poly.*

We conjecture that the same negative result holds for the general case as well.

Conjecture: *No maximal-in-range mechanism can obtain an approximation ratio better than $\frac{1}{2}$ (unless NP $\subseteq$ P/poly).*

Our lower bound for the allocate-all-items case implies the an inapproximability result for *all* truthful mechanisms for the class of XOS valuations (defined in [33] as a "bidding language", and later studied in [13, 15, 10, 16, 17, 29, 41]):

Theorem: *For the allocate-all-items case of combinatorial auctions with 2 XOS bidders, and for any constant $\epsilon > 0$, no polynomial-time truthful mechanism obtains an approximation ratio of $\frac{1}{2} + \epsilon$ unless NP $\subseteq$ P/poly.*

Thus, our techniques enable us to show that while a non-truthful $\frac{3}{4}$-approximation algorithm exists for this class [15, 17], there is a lower bound of $\frac{1}{2}$ for truthful algorithms, that is unconditioned on the maximal-in-range assumption. This result relies on the recent characterization of truthful algorithms for this setting presented in [16]. We stress that with the exception of a result in [26], this is the first such computational-complexity lower bound for auction settings.

1.2 Organization of the Paper

In Section 2 we discuss the outline of our approach to proving inapproximability results for maximal-in-range mechanisms. In particular, we present the interesting connections between our goal and the classic notion of VC dimension. In Section 3 we present our lower bound for the special case of the 2-bidder auction described above in which all items must be allocated. In Section 4 we present our result for the case in which this assumption is removed. Finally, in Section 5 we present open questions and directions for future research.

2 The Relationship between Between VC and VCG

In this section we outline our general approach for proving inapproximability results for maximal-in-range mechanisms for auctions. We do so by pointing out the interesting connection between our goal and the classic notion of the VC dimension.

2.1 Maximal-In-Range Mechanisms for Combinatorial Auctions

Consider a combinatorial auction with 2 bidders, as in our introductory auction setting. The outcome of any algorithm for this setting is a partition of the items.

Definition 2.1 *A partition $T = (T_1, T_2)$ of a universe $U = \{1, ..., m\}$ is a pair of two disjoint subsets of U, i.e. $T_1, T_2 \subseteq [m]$ and $T_1 \cap T_2 = \emptyset$.*

Observe that a partition of a universe of items U need not necessarily be such that every item appears in one of the subsets that form the partition.

We now define the class of maximal-in-range mechanisms for this auction: Let P be the set of all possible partitions of the items into two disjoint subsets $P = \{(T_1, T_2)|\ T_1 \cap T_2 = \emptyset\}$ (observe that a partition in P does not necessarily contain all items). Rather than trying to find a partition of the items in P that maximizes the social-welfare, a maximal-in-range mechanism M fixes a *subset* $R_M \subseteq P$, and, for every pair of valuations, outputs the optimal solution *within this subset*. That is, for every (v_1, v_2), M outputs the partition in R_M that maximizes the social welfare, i.e., $argmax_{(T_1,T_2)\in R_m}\ \Sigma_i v_i(T_i)$. We refer to R_M as M's *range*.

Given a maximal-in-range mechanism M, how can we prove that it does not run in polynomial time? By definition, a maximal-in-range mechanism (exactly) optimizes over a collection of partitions of items. Of course, if the range of our mechanism M contains all possible partitions, then it is bound to reach the globally optimal solution – an NP-hard task. If, on the other hand, M optimizes over a strict subset of all possible partitions, we hope to be able to show that we can reduce from an NP-hard problem to the problem of optimizing over M's range. As we shall now show, the concept of VC dimension allows us to achieve just that.

2.2 The VC Dimension of Partitions

We now present our formal framework for analyzing the VC dimension of *partitions*:

Definition 2.2 *A partition (T_1, T_2) of a universe U is said to cover U if $T_1 \cup T_2 = U$. $C(U)$ is defined to be the set of all partitions that cover U.*

For every subset E of a universe U, we can define (in an analogous way) what a partition of E is, and denote by $P(E)$ the set of all partitions of E and by $C(E)$ the set of all partitions of E that cover E.

Definition 2.3 *The projection of a partition (S_1, S_2) of a universe U on $E \subseteq U$, denoted by $(S_1, S_2)_{|E}$ is the partition $(S_1 \cap E, S_2 \cap E)$ of E. For any collection of partitions R of a universe U we define R's projection on $E \subseteq U$, $R_{|E}$, to be $R_{|E} = \{(T_1, T_2)|\ \exists (S_1, S_2) \in R$ s.t. $(S_1, S_2)_{|E} = (T_1, T_2)\}$.*

Observation 2.4 *If a partition (S_1, S_2) of $E \subseteq U$ is in $C(E)$, then for any $E' \subseteq E$ $(S_1, S_2)_{|E'} \in C(E')$.*

Now, we are ready to define shattering and VC dimension in our context:

Definition 2.5 *A subset E of a universe U is said to be shattered by a collection R of n-partitions of U if $C(E) \subseteq R_{|E}$.*

Observation 2.6 *If $E \subseteq U$ is shattered by a collection R of partitions of U then so are all subsets of E.*

Definition 2.7 *The VC dimension $VC(R)$ of a collection R of partitions of a universe U is the cardinality of the biggest subset $E \subseteq U$ that is shattered by R.*

2.3 The Connection: The VC Dimension of the Range

Let M be a maximal-in-range mechanism for a combinatorial auction with 2 bidders. We can now explain our approach for proving lower bounds on the approximability of M: Assume that we are dealing with a class of valuations for which maximizing the social-welfare is NP-hard (as is indeed the case with our introductory auction setting).

Proposition 2.8 *Let M be a maximal-in-range mechanism with $VC(R_M) \geq m^\alpha$ for some $\alpha > 0$. Then M does not run in polynomial time unless NP $\subseteq$ P/poly.*

Proof: [Sketch] Assume that M is such that there is some constant $\alpha > 0$ for which $VC(R_M) \geq m^\alpha$. Then, there is some set of items E of size at least m^α that is shattered by R_M. Therefore, given an auction with m^α items and valuation functions v_1, v_2 we can identify each item in this smaller auction with some unique item in E, and construct valuation functions v'_1, v'_2, such that v'_i is identical to v_i on E and assigns 0 to all other items. Observe that this means that M will output for v'_1, v'_2 the optimal solution for v_1, v_2 (as M's range contains all partitions in $C(E)$). We now have a reduction from an NP-hard problem (social-welfare maximization in the smaller auction) to the optimization problem solved by M. However, this is a non-uniform reduction, since we are not guaranteed that we can find E in polynomial time. Indeed, in many cases E's existence proof is non-constructive (see Lemma 1 below), and so our complexity conclusion is dependent on the assumption that NP does not have polynomial-size circuits. □

Proposition 2.8 suggests a general way of proving inapproximability complexity results for maximal-in-range mechanisms for auctions: If we wish to prove a lower bound of r, we show that any maximal-in-range mechanism M that obtains an approximation ratio better than r must be such that there is a constant $\alpha > 0$ for which $VC(R_M) \geq m^\alpha$.

3 First Step: The Allocate-All-Items Case

We prove a lower bound of 2 for the special case of our introductory simple 2-bidder auction, in which we require that all items be allocated by the mechanism (the auctioneer can keep no items to himself). We then show that this lower bound can, in certain interesting cases, be extended to hold for all truthful algorithms. The proof of our lower bound for the allocate-all-items consists of three main steps:

- Show that the range of any maximal-in-range mechanism that obtains an approximation ratio better than $\frac{1}{2}$ must be of exponential size (in the number of items).
- Show (via the Sauer-Shelah Lemma [39, 42]) that if the range of a maximal-in-range mechanism is exponential, then there is a "large" subset E of the items that is shattered by this range.
- Use E to reduce from an NP-hard problem to the optimization problem that is solved (exactly) by the maximal-in-range mechanism.

Intuitively, the above steps will also be the outline of the proof of our lower bound for the case in which not all items are necessarily allocated. However, this will require new mathematical machinery, as explained in Section 4.

Theorem 3.1 *For any $\epsilon > 0$, any maximal-in-range mechanism that always allocates all items and obtains an approximation ratio of $\frac{1}{2} + \epsilon$ cannot be implemented in polynomial-time unless NP $\subseteq$ P/poly.*

Proof: [Sketch] Let M be some maximal-in-range mechanism that always allocates all items and obtains an approximation ratio of $\frac{1}{2} + \epsilon$. Let R_M be M's range. Because M always allocates all items $R_M \subseteq C([m])$.

Lemma 3.2 *[11] There is some constant $\alpha > 0$ such that $|R_M| = \Omega(e^{m^\alpha})$.*

Proof: [Sketch] We use a probabilistic construction. We construct the valuations of the bidders in the following way: For each item j we choose one bidder i uniformly at random and set $v_{ij} = 1$. We set the value of the other bidder for that item to be 0. (We set both budgets to be m so that they will not play any role.) Now, observe that the optimal social welfare if the valuations are v_1, v_2 is m (assign each item j to the bidder who values it as 1). However, let $S = (S_1, S_2)$ be some partition in R_M. Using standard Chernoff arguments it is easy to show that the probability that the social-welfare obtained by allocating the items as in S is greater than $\frac{(1+\epsilon')}{2}m$ is exponentially small in ϵ'. Hence, it requires exponentially many sets in the range R_M to ensure that an approximation better than 2 is obtained for every such pair of valuations. □

We denote by $R_{M,i}$ the collection of all sets of player i that appear in some partition in R_M. Because $|R_M|$ is exponential in m and there are two bidders we have that:

Corollary 3.3 *There is a bidder $i \in \{1, 2\}$ and a constant $\alpha' > 0$ such that $|R_{M,i}| = \Omega(e^{m^{\alpha'}})$.*

We now recall the Sauer-Shelah Lemma:

Lemma 3.4 *([39, 42]) For any family Z of subsets of a universe U, there is a subset E of U of size $\Theta(\frac{\log |Z|}{\log |U|})$ such that for each $E' \subseteq E$ there is a $Z' \in Z$ such that $E' = Z' \cap E$.*

The Sauer-Shelah Lemma, and the fact that for some bidder i $|R_{M,i}| = \Omega(e^{m^{\alpha'}})$ (for some constant α'), implies that there is a constant $\beta > 0$, and a subset of the items E of size at least m^β such that E is shattered by the *subsets* in $R_{M,i}$ (in the traditional sense). However, as $R^M \subseteq C([m])$, and by Observation 2.4 it must hold that E is also shattered by the *partitions* in R_M. Hence, $VC(R_M) \geq m^\beta$. Therefore, by Proposition 2.8 M cannot run in polynomial time unless NP $\subseteq$ P/*poly* (this relies on the fact that finding the global social-welfare-maximizing outcome is NP-hard even for $n = 2$, see [3]). □

It has recently been shown in [16] that for the class of XOS valuations [33, 17], if all items are allocated then truthful algorithms are bound to be maximal-in-range. As XOS valuations are a superclass of the valuations considered in this paper, we get the following:

Theorem 3.5 *For the allocate-all-items case of combinatorial auctions with 2 XOS bidders, and for any constant $\epsilon > 0$, no polynomial-time truthful mechanism obtains an approximation ratio of $\frac{1}{2} + \epsilon$ unless NP $\subseteq$ P/poly.*

We note that a non-truthful $\frac{3}{4}$-approximation algorithm exists for this class of valuations [15, 17]. Hence, we prove that there is a gap between what can be achieved by truthful and unrestricted (computationally-efficient) algorithms for this problem.

4 The General Case

We are now ready to prove our main result, that no maximal-in-range mechanism has an approximation-ratio better than $\frac{3}{4}$ unless SAT has polynomial-size circuits . It may seem, at first sight, like one can easily extend our proof for the case where all the items must allocated to prove the same lower bound for the case in which this assumption is removed. Indeed, from a purely computational viewpoint, any algorithm that does not allocate all items can easily be converted to one that does simply by assigning all unallocated items to a specific bidder (say, bidder 1). However, this simple solution does not work when it comes to maximal-in-range mechanisms. This is because arbitrarily assigning items that are not allocated by a maximal-in-range mechanism does *not* result in a maximal-in-range mechanism.

Similarly to the proof of our previous lower bound, the proof of our main result consists of three main steps:

- Show that the range of any maximal-in-range mechanism that obtains an approximation ratio better than $\frac{3}{4}$ must be "rich". Informally, the richness of a range means that it contains a subset of partitions of exponential-size, such that no two partitions in this subset agree on too many items (by agreeing on an item, we mean that it is allocated to the same bidder by both partitions).
- Once we know that the range of a maximal-in-range mechanism is rich, we prove that there is a "large" subset E of the items that is shattered by this range. Unlike the proof of our lower bound for the allocate-all-items case (in Section 3), this proof necessitates new mathematical machinery (we cannot simply apply the Sauer-Shelah Lemma). We present a lemma that can be regarded as the analogue of the Sauer-Shelah Lemma for the case of partitions (in Subsection 4.1).
- We use E to reduce from an NP-hard problem to the optimization problem that is solved (exactly) by the maximal-in-range mechanism.

We start (Subsection 4.1) by presenting a method for lower bounding the VC dimension of a collection of partitions. This is then used for proving our lower bound in Subsection 4.2.

4.1 Lower Bounding the VC Dimension of Partitions

We shall now present our way of lower bounding the VC dimension of partitions:

Definition 4.1 *Given a universe U, two partitions of it (T_1, T_2) and (T'_1, T'_2) are said to be b-far (or at distance b) if $|T_1 \cap T'_2| + |T'_1 \cap T_2| \geq b$.*

Definition 4.2 *Let $t(\epsilon, k, m)$ be the smallest possible number of sets $E \subset [m]$ that are shattered by a set R of partitions of size k, such that every two elements in R are at least ϵm-far.*

Observation 4.3 *Suppose $k \geq 1$ and $\epsilon m \geq 1$. Then if $t(\epsilon, k, m) > \sum_{i=1}^{r} \binom{m}{r}$ then the VC dimension of any collection of partitions of size at least k for which every two partitions are at least ϵm-far has to be at least $r+1$.*

Proof: The proof follows from the fact that $t(\epsilon, k, m) \geq \sum_{i=0}^{r} \binom{m}{r}$ is a bound on the number of sets of size at most r. □

Lemma 1 *For all $\epsilon > 0, k, m$, $t(\epsilon, k, m) \geq k^{\alpha}$ for some constant $\alpha > 0$.*

The proof follows the basic idea of [2, 6, 31]. Our novel observation is that the same proof strategy applies with our new definition of distance.

Proof: [Sketch] Fix $\epsilon > 0, k, m$. We wish to prove that $t(\epsilon, k, m) \geq k^{\alpha}$, for some constant $\alpha > 0$. We shall bound $t(\epsilon, k, m)$ by induction (ϵ shall remain fixed throughout the proof and the induction is on k and m). Let R be some set of partitions as in the statement of the lemma. Arbitrarily partition R into pairs. Since the partitions that make up each pair are at least ϵm-far there must exist (via simple counting) an element $e \in U$, such that in at least $\frac{\epsilon k}{2}$ pairs $(T_1, T_2), (T_1', T_2')$, $e \in T_1 \cap T_2'$ or $e \in T_1' \cap T_2$. Let $R' \subseteq R$ be the collection of all partitions (T_1, T_2) in R in which $e \in T_1$. Let $R'' \subseteq R$ be the collection of all partitions (T_1, T_2) in R in which $e \in T_2$. By the arguments above we are guaranteed that $|R'| \geq \frac{\epsilon k}{2}$ and $|R''| \geq \frac{\epsilon k}{2}$.

Let I be all the subsets of U that are shattered by R. We wish to lower bound $|I|$. Let R'_{-e} be all the partitions of $U \setminus \{e\}$ we get by removing e from T_1 for every partition $(T_1, T_2) \in R'$. Let I' be all the subsets of $U \setminus \{e\}$ shattered by R'_{-e}. As there are at least $\frac{\epsilon k}{2}$ sets in R', by definition $|I'| \geq t(\epsilon, \frac{\epsilon k}{2}, m-1)$. Similarly, let R''_{-e} be all the partitions of $U \setminus \{e\}$ we get by removing e from T_1 for every partition $(T_1, T_2) \in R''$. Let I'' be all the subsets of $U \setminus \{e\}$ shattered by R''_{-e}. As there are at least $\frac{\epsilon k}{2}$ sets in R'', by definition $|I''| \geq t(\epsilon, \frac{\epsilon k}{2}, m-1)$.

We claim that $|I| \geq |I'| + |I''|$. To see why this is true consider the following argument: All sets in $I' \setminus I''$ and in $I'' \setminus I'$ are distinct and belong to I. Let S be a set in $I' \cap I''$. Observe that this means that not only is S in I, but so is $S \cup \{e\}$. So, $I \geq |I' \setminus I''| + |I' \setminus I''| + 2|I' \cap I''| = |I'| + |I''|$. Hence, $t(\epsilon, k, m) \geq 2 \times t(\epsilon, \frac{\epsilon k}{2}, m-1)$. We now use the induction hypothesis to conclude the proof. □

Lemma 1 and Observation 4.3 lead to the following corollary, which will play a crucial role in the proof of our lower bound:

Corollary 4.4 *For every $\alpha, \epsilon > 0$ there exists a $\beta > 0$ such that, if R is a set of partitions of a universe U of size m with these properties: (1) R is of size at least $e^{m^{\alpha}}$. (2) two partitions in R are at least ϵm-far. Then, $VC(R) \geq m^{\beta}$.*

4.2 A Lower Bound for Maximal-In-Range Mechanisms

Using the machinery presented in the previous subsection we are able to prove the following theorem:

Theorem 4.5 *For any constant $\epsilon > 0$, no polynomial-time maximal-in-range mechanism obtains an approximation ratio of $\frac{3}{4} + \epsilon$ unless NP ⊆ P/poly.*

Proof: The main part of the proof of this theorem is showing that the range of any maximal-in-range mechanism that obtains an approximation ratio better than $\frac{3}{4}$ must be rich, in the sense that it contains a subset of exponential size of partitions that are not too close to each other. This is precisely what is shown by the following lemma:

Lemma 2 *Let $\epsilon > 0$. Let M be a maximal-in-range mechanism that obtains an approximation ratio of $\frac{3}{4} + \epsilon$, and let R_M be M's range. Then, there is a constant $\delta > 0$ and subset of R_M, R'_M of size exponential in m such that every two elements of R'_M are at least δm-far.*

Proof: Using probabilistic construction we show that, for our universe of items $U = [m]$, that there exists an exponential-sized set of partitions of all items that are at least $\frac{1-\epsilon'}{2}m$-far (for some arbitrarily small $\epsilon' > 0$).

Claim 4.6 *For every $\epsilon' > 0$, there is a family F of partitions (T_1, T_2) in $C([m])$ and a constant $\alpha > 0$ such that $|F| = e^{\alpha m}$ and every two partitions in F are at least $\frac{1-\epsilon'}{2}m$-far.*

Proof: [Sketch] We will prove the claim for partitions $T = (T_1, T_2)$ where $T_1 \cup T_2 = [m]$. For such partitions, the distance between partition $T = (T_1, T_2)$ and $T' = (T'_1, T'_2)$ is just the size of the symmetric difference of T_1 and T_2. The existence of the desired collection now follows from the existence of good codes, see e.g. [43]. For completeness we include the standard construction to show the existence of F.

Let $T = (T_1, T_2)$ and $T' = (T'_1, T'_2)$ be two partitions in $C([m])$ chosen at random in the following way: For each item $j \in [m]$ we choose, uniformly at random, whether it will be placed in T_1 or in T_2. Similarly, we choose, uniformly at random whether each item j shall be placed in T'_1 or T'_2. Using standard Chernoff arguments it is easy to show that the probability that there are at least $\frac{m+\epsilon'}{2}$ that appear in either $T_1 \cap T'_1$ or $T_2 \cap T'_2$ is exponentially small in ϵ'. Observe that this immediately implies (by our definition of distance) that the probability that the distance between T and T' is less than $\frac{1-\epsilon'}{2}m$ is exponentially small is ϵ'. Hence, a family F of exponential size must exist. □

Let F be a family of partitions as in Claim 4.6 (for some arbitrarily small ϵ'). For every partition $T = (T_1, T_2)$ in F we define valuations v_1^T, v_2^T as follows: v_1^T assigns a value of 1 to every item in T_1 and a value of 0 to every item in T_2. Similarly, v_2^T assigns a value of 1 to every item in T_2 and a value of 0 to every item in T_1. (We ignore the budgets by setting them to be high.) For every pair of valuations v_1^T, v_2^T, let $R^T = (R_1^T, R_2^T)$ the partition in R_M that M outputs for the input v_1^T, v_2^T.

Claim 4.7 *There is a constant $\delta > 0$ such that for every two partitions $T, T' \in F$ it holds that R^T and $R^{T'}$ are δm-far.*

Proof: Fix $T, T' \in F$. Observe that the optimal social welfare for v_1^T, v_2^T and for $v_1^{T'}, v_2^{T'}$ is m (we can always assign each item to the only bidder who wants it). T and T' are identical only for at most $\frac{1+\epsilon'}{2}m$ items (that is, only for at most $\frac{1+\epsilon'}{2}m$ items j, either $j \in T_1 \cap T'_1$ or $j \in T_2 \cap T'_2$). So, even if we assume that both R^T and $R^{T'}$ allocate all these items as in T and T', we are left with $\frac{1-\epsilon'}{2}m$ items that are allocated in different ways in T and T'. For each such item, if R^T and $R^{T'}$ are identical for it, then it will only contribute to the social welfare of one of the two valuation pairs, and not to both.

We are assuming that M has an approximation ratio of $\frac{3}{4} + \epsilon$. Hence, it must be that the social welfare obtained by R^T and $R^{T'}$ for v_1^T, v_2^T and $v_1^{T'}, v_2^{T'}$ (respectively) is at least $(\frac{3}{4} + \epsilon)m$. However, observe that this cannot be achieved without assigning a constant fraction of the items (say $\frac{\epsilon}{3}m$, given an arbitrarily small ϵ') to different bidders in R^T and $R^{T'}$. This implies that there is some $\delta > 0$ such that R^T and $R^{T'}$ are δm-far. □

So, for every partition T in F there exists a partition in R^M such that every two such partitions in R^M are δm-far. As F is of size exponential in m this concludes the proof of the lemma. □

We have shown that any maximal-in-range mechanism R that obtains an approximation ratio better than $\frac{3}{4}$ is such that its range R^M contains a subset R'_M of size exponential in m, and every two partitions in R'_M are δm-far (for some constant $\delta > 0$). We can now use Corollary 4.4 conclude that $VC(R_M) \geq VC(R'_M) \geq m^{\alpha}$ (for some constant $\alpha > 0$). By applying Proposition 2.8 we can now determine that M is not computable in polynomial-time unless SAT has polynomial circuits. $\square$

5 Discussion and Open Questions

We believe that our work opens a new avenue for proving complexity-theoretic inapproximability results for maximal-in-range mechanisms for auctions. In particular, the following important questions remain wide open:

1. **We already conjectured** that the trivial bound of $\frac{1}{2}$ is the best approximation possible for our two bidder auctions.

2. **Generalizing our result to the n-bidder case.** For the general n-bidder case the best (deterministic) approximation ratio known is $O(\min\{n, \sqrt{m}\})$ and is achieved via a simple maximal-in-range mechanism [13] (using randomization, improved, but still non-constant, approximation ratios are achievable [14, 10]). A straightforward extension of our techniques leads to the following result:

   Theorem 5.1 *For any constant number of bidders n, and any $\epsilon > 0$, no maximal-in-range mechanism can obtain an approximation ratio of $\frac{n+1}{2n} + \epsilon$ unless NP $\subseteq$ P/poly.*

   Conjecture: *No maximal-in-range mechanism can obtain a constant approximation ratio for the n-bidder case.*

   The proof for such a result will probably necessitate the development of more advanced tools for lower bounding the VC dimension of partitions.

3. **Relaxing the computational assumption.** Our results in this paper are dependent on the assumption that SAT does not have polynomial-size circuits. This is due to the fact that we do not know a constructive way of finding a large shattered set. In [36], this problem was overcome by using a probabilistic version of the Sauer-Shelah Lemma due to Ajtai [1] (see [37, 40, 32] on the complexity of making Sauer-Shelah Lemma constructive). This relaxes he assumption that NP is not contained in P/*poly* to the weaker assumption that NP is not contained in BPP. Is a similar approach possible in our more complex case? That is, is there a probabilistic version of our combinatorial lemmas that will enable up to find a shattered set with high probability?

4. **Characterizing truthfulness in auctions.** Our results in this paper hold for the important class of maximal-in-range mechanisms. Can similar results be obtained for all truthful mechanisms (as we show is indeed the case for XOS valuations)? To answer this question we must better understand the properties of truthful algorithms in complex mechanism design settings. So far, despite much work on this subject [38, 26, 27, 7, 16, 36], very little is known about characterizations of truthfulness in combinatorial auctions (and in other multi-parameter settings).

References

[1] Miklós Ajtai. The shortest vector problem in L_2 is NP-hard for randomized reductions (extended abstract). In *STOC*, pages 10–19, 1998.

[2] Noga Alon, Shai Ben-David, Nicolò Cesa-Bianchi, and David Haussler. Scale-sensitive dimensions, uniform convergence, and learnability. *J. ACM*, 44(4):615–631, 1997.

[3] Nir Andelman and Yishay Mansour. Auctions with budget constraints. In *SWAT*, pages 26–38, 2004.

[4] A. Archer and É. Tardos. Truthful mechanisms for one-parameter agents. In *FOCS '01: Proceedings of the 42nd IEEE symposium on Foundations of Computer Science*, page 482, Washington, DC, USA, 2001. IEEE Computer Society.

[5] Yair Bartal, Rica Gonen, and Noam Nisan. Incentive compatible multi unit combinatorial auctions. In *TARK '03: Proceedings of the 9th conference on Theoretical aspects of rationality and knowledge*, pages 72–87, New York, NY, USA, 2003. ACM.

[6] Peter L. Bartlett and Philip M. Long. Prediction, learning, uniform convergence, and scale-sensitive dimensions. *J. Comput. Syst. Sci.*, 56(2):174–190, 1998.

[7] Sushil Bikhchandani, Shurojit Chatterji, Ron Lavi, Ahuva Mu'alem, Noam Nisan, and Arunava Sen. Weak monotonicity characterizes deterministic dominant-strategy implementation. *Econometrica*, 74(4):1109–1132, 07 2006.

[8] Liad Blumrosen and Shahar Dobzinski. Welfare maximization in congestion games. In *EC '06: Proceedings of the 7th ACM conference on Electronic commerce*, pages 52–61, New York, NY, USA, 2006. ACM.

[9] E. H. Clarke. Multipart pricing of public goods. *Public Choice*, pages 17–33, 1971.

[10] Shahar Dobzinski. Two randomized mechanisms for combinatorial auctions. In *APPROX-RANDOM*, pages 89–103, 2007.

[11] Shahar Dobzinski and Noam Nisan. Limitations of vcg-based mechanisms. In *STOC '07: Proceedings of the thirty-ninth annual ACM symposium on Theory of computing*, pages 338–344, New York, NY, USA, 2007. ACM.

[12] Shahar Dobzinski and Noam Nisan. Mechanisms for multi-unit auctions. In *ACM Conference on Electronic Commerce*, pages 346–351, 2007.

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

[14] Shahar Dobzinski, Noam Nisan, and Michael Schapira. Truthful randomized mechanisms for combinatorial auctions. In *STOC '06: Proceedings of the thirty-eighth annual ACM symposium on Theory of computing*, pages 644–652, New York, NY, USA, 2006. ACM.

[15] Shahar Dobzinski and Michael Schapira. An improved approximation algorithm for combinatorial auctions with submodular bidders. In *SODA '06: Proceedings of the seventeenth annual ACM-SIAM symposium on Discrete algorithm*, pages 1064–1073, New York, NY, USA, 2006. ACM.

[16] Shahar Dobzinski and Mukund Sundararajan. On characterizations of truthful mechanisms for combinatorial auctions and scheduling. In *EC '08: Proceedings of the 9th ACM conference on Electronic commerce*, pages 38–47, New York, NY, USA, 2008. ACM.

[17] Uriel Feige. On maximizing welfare when utility functions are subadditive. In *STOC '06: Proceedings of the thirty-eighth annual ACM symposium on Theory of computing*, pages 41–50, New York, NY, USA, 2006. ACM.

[18] Uriel Feige and Jan Vondrak. Approximation algorithms for allocation problems: Improving the factor of 1 - 1/e. In *FOCS '06: Proceedings of the 47th Annual IEEE Symposium on Foundations of Computer Science*, pages 667–676, Washington, DC, USA, 2006. IEEE Computer Society.

[19] Joan Feigenbaum, Michael Schapira, and Scott Shenker. Distributed algorithmic mechanism design. In Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay Vazirani, editors, *Algorithmic Game Theory*, chapter 14. Cambridge University Press, 2007.

[20] Amos Fiat, Andrew V. Goldberg, Jason D. Hartline, and Anna R. Karlin. Competitive generalized auctions. In *34th Symposium on Theory of Computing*, pages 72–81, New York, 2002. ACM.

[21] A. V. Goldberg, J. D. Hartline, A. R. Karlin, M. Saks, and A. Wright. Competitive auctions. *Games and Economic Behavior*, 55(2):242 – 269, 2006.

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

[23] Johan Håstad. Some optimal inapproximability results. *J. ACM*, 48(4):798–859, 2001.

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

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

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

[27] Ron Lavi, Ahuva Mu'alem, and Noam Nisan. Two simplified proofs for Roberts' Theorem, 2004. Submitted.

[28] Ron Lavi and Chaitanya Swamy. Truthful mechanism design for multi-dimensional scheduling via cycle monotonicity. In *ACM Conference on Electronic Commerce*, pages 252–261, 2007.

[29] Benny Lehmann, Daniel Lehmann, and Noam Nisan. Combinatorial auctions with decreasing marginal utilities. *Games and Economic Behaviour. (A preliminary version appeared in EC'01)*, 55(2):270–296, 2006.

[30] Hagay Levin, Michael Schapira, and Aviv Zohar. Interdomain routing and games. In *STOC*, pages 57–66, 2008.

[31] Shahar Mendelson and Roman Vershynin. Entropy, combinatorial dimensions and random averages. In *COLT*, pages 14–28, 2002.

[32] Elchanan Mossel and Christopher Umans. On the complexity of approximating the VC dimension. *J. Comput. Syst. Sci.*, 65(4):660–671, 2002.

[33] Noam Nisan. Bidding and allocation in combinatorial auctions. In *ACM Conference on Electronic Commerce*, 2000.

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

[35] Noam Nisan and Amir Ronen. Algorithmic mechanism design. *Games and Economic Behaviour*, 35:166 – 196, 2001. A preliminary version appeared in STOC 1999.

[36] Christos Papadimitriou, Michael Schapira, and Yaron Singer. On the hardness of being truthful. In *FOCS '08: Proceedings of the 49th Annual IEEE Symposium on Foundations of Computer Science*, Philadelphia, PA, USA, 2008. IEEE Computer Society.

[37] Christos H. Papadimitriou and Mihalis Yannakakis. On limited nondeterminism and the complexity of the v-c dimension. *J. Comput. Syst. Sci.*, 53(2):161–170, 1996.

[38] Kevin Roberts. The characterization of implementable choice rules. In Jean-Jacques Laffont, editor, *Aggregation and Revelation of Preferences. Papers presented at the 1st European Summer Workshop of the Econometric Society*, pages 321–349. North-Holland, 1979.

[39] Norbert Sauer. On the density of families of sets. *J. Comb. Theory, Ser. A*, 13(1):145–147, 1972.

[40] M. Schäfer. Deciding the Vapnik-Cervonenkis dimension is Σ_3^p-complete. *Journal of Computer and System Sciences*, 58:177–182, 1999.

[41] Michael Schapira and Yaron Singer. Inapproximability of combinatorial public projects. In *WINE*, 2008.

[42] Saharon Shelah. A combinatorial problem; stability and order for models and theories in infinitary languages. *Pacific J Math*, 41:247–261, 1972.

[43] Madhu Sudan. Essential coding theory, lecture notes, 2008. (MIT 6.440).

[44] W. Vickrey. Counterspeculation, auctions and competitive sealed tenders. *Journal of Finance*, pages 8–37, 1961.

[45] Jan Vondrák. Optimal approximation for the submodular welfare problem in the value oracle model. In *STOC*, pages 67–74, 2008.

[46] Andrew Chi-Chih Yao. Some complexity questions related to distributive computing(preliminary report). In *STOC '79: Proceedings of the eleventh annual ACM symposium on Theory of computing*, pages 209–213, New York, NY, USA, 1979. ACM.