

VC *v.* VCG: Inapproximability of Combinatorial Auctions via Generalizations of the VC Dimension

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 the paradigmatic problem in computational mechanism design. It is believed that in many cases good approximations for combinatorial auctions may be unattainable due to an inherent clash between truthfulness and computational efficiency. However, to date, *researchers lack the machinery to prove such results*. In this paper, we present a new approach: We take the first steps towards the development of new technologies for lower bounding the *VC-dimension of k -tuples of disjoint sets*. We apply this machinery to prove the first *computational-complexity* inapproximability results for incentive-compatible mechanisms for combinatorial auctions. These results hold for the important class of VCG-based mechanisms, and are based on the complexity assumption that NP has no polynomial-size circuits. We believe that our approach holds great promise. Indeed, subsequently to our work, Buchfuhrer and Umans [10], and, independently, Dughmi *et al.* [20], strengthened our results via extensions of our techniques, thus proving conjectures we presented in [32]

*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. yaron@cs.berkeley.edu.

1 Introduction

The field of *algorithmic mechanism design* [34] is about the reconciliation of bounded computational resources and strategic interaction between selfish participants. In *combinatorial auctions*, the paradigmatic problem in this area, a set of items is sold to bidders with *private* preferences over *subsets* of the items, with the intent of maximizing the *social welfare* (*i.e.*, the sum of bidders' utilities from their allocated items). Researchers constantly seek auction protocols that are both *incentive-compatible* and *computationally-efficient*, and guarantee decent approximation ratios. Sadly, to date, huge gaps exist between the state of the art approximation ratios obtained by unrestricted, and by truthful, algorithms. It is believed that this could be due to an inherent clash between the truthfulness and computational-efficiency requirements, that manifests itself in greatly degraded algorithm performance. Such tension between the two desiderata was recently shown to exist in [36] for a different mechanism design problem called *combinatorial public projects* [40]. However, in the context of combinatorial auctions, due to their unique combinatorial structure, *the algorithmic game theory community currently lacks the machinery to prove this* [38].

The celebrated class of *Vickrey-Clarke-Groves (VCG) mechanisms* [43, 11, 23] is the only known universal technique for the design of *deterministic* incentive-compatible mechanisms (in certain interesting cases VCG mechanisms are the *only* truthful mechanisms [8, 19, 27, 37, 36]). While a naive application of VCG is often computationally intractable, more clever uses of VCG are the key to the best known (deterministic) approximation ratios for combinatorial auctions [16, 25]. For these reasons, the exploration of the computational limitations of such mechanisms is an important research agenda (pursued in [15, 19, 27, 33, 36]). Recently, it was shown [36] that the computational-complexity of VCG-based mechanisms is closely related to the notion of VC dimension. [36] was able to make use of *existing* VC machinery to prove computational hardness results for combinatorial public projects. However, for combinatorial auctions, these techniques are no longer applicable. This is because, unlike combinatorial public projects, the space of outcomes in combinatorial auctions does not consist of subsets of the universe of items, but rather of *partitions* of this universe (between the bidders). This calls for the development of new VC machinery for the handling of such problems.

We formally define the notion of the *VC dimension of collections of partitions of a universe*¹ and present techniques for establishing lower bounds on this VC dimension. We show how these results can be used to prove computational-complexity inapproximability results for VCG-based mechanisms for combinatorial auctions. (We note that these results actually hold for the more general class of mechanisms that are affine maximizers [27, 37].) Our inapproximability results depend on the computational assumption that SAT does not have polynomial-size circuits. Informally, our method of lower bounding² the approximability of VCG-based mechanisms via VC arguments is the following: We consider well-known auction environments for which *exact* optimization is NP-hard. We show that if a VCG-based mechanism *approximates closely* the optimal social welfare, 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 (*partitions* of items) considered by the VCG-based mechanism “shatters” a relatively large subset of the items.

Our results imply the first *computational complexity* lower bounds for VCG-based mechanisms, and truthful mechanisms in general, for combinatorial auctions (with the possible exception of a result in [27] for a related auction environment). In fact, we show that in some auction en-

¹By a *partition* in this paper we mean an ordered k -tuple of disjoint subsets which, however, may not exhaust the universe.

²We 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 [24].

vironments these results actually apply to *all* truthful algorithms. We illustrate our techniques via 2-bidder combinatorial auctions. It is also our belief that the notion of the VC dimension of k -tuples of disjoint sets is of independent interest, and suggests many new and exciting problems in combinatorics.

1.1 Related Work

Combinatorial auctions have been extensively studied in both the economics and the computer science literature [9, 12, 13]. It is known that if the preferences of the bidders are unrestricted then no constant approximation ratios are achievable (in polynomial time) [30, 35]. Hence, much research has been devoted to the exploration of restrictions on bidders' preferences that allow for good approximations, *e.g.*, for *subadditive*, and *submodular*, preferences *constant* approximation ratios have been obtained [16, 18, 21, 22, 29, 44]. In contrast, the known *truthful* approximation algorithms for these classes have *non-constant* approximation ratios [14, 16, 17]. It is believed that this gap may be due to the computational burden imposed by the truthfulness requirement. However, to date, this belief remains unproven. In particular, no *computational complexity* lower bounds for truthful mechanisms for combinatorial auctions are known.

Vickrey-Clarke-Groves (VCG) mechanisms [43, 11, 23], named after their three inventors, are the fundamental technique in mechanism design for inducing truthful behaviour of strategic agents. Nisan and Ronen [33, 34] were the first to consider the computational issues associated with the VCG technique. In particular, [33] defines the notion of VCG-Based mechanisms. VCG-based mechanisms have proven to be useful in designing approximation algorithms for combinatorial auctions [16, 25]. In fact, the best known (deterministic) truthful approximation ratios for combinatorial auctions were obtained via VCG-based mechanisms [16, 25] (with the notable exception of an algorithm in [5] for the case that many duplicates of each item exist). Moreover, Lavi, Mu'alem and Nisan [27] have shown that in certain interesting cases VCG-based mechanisms are essentially the only truthful mechanisms (see also [19]).

Dobzinski and Nisan [15] tackled the problem of proving inapproximability results for VCG-based mechanisms by taking a *communication complexity* [45, 26] approach. Hence, in the settings considered in [15], it is assumed that each bidder has an *exponentially large string of preferences* (in the number of items). However, real-life considerations render problematic the assumption that bidders' preferences are exponential in size. Our intractability results deal with bidder preferences that are *succinctly described*, and therefore relate to *computational complexity*. Thus, our techniques enable us to prove lower bounds even for the important case in which bidders' preferences can be concisely represented.

The connection between the VC dimension and VCG-based mechanisms was observed in [36], where a general (i.e., not restricted to VCG-based mechanisms) inapproximability result was presented, albeit in the context of a different mechanism design problem, called *combinatorial public projects* (see also [40]). 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, 41] and its probabilistic version due to Ajtai [1]). To handle the unique technical challenges posed by combinatorial auctions (specifically, the fact that the universe of items is *partitioned* between the bidders) *new* machinery is required. Indeed, our technique can be interpreted as *an extension of the Sauer-Shelah Lemma to the case of partitions* (Lemma 2.12 in Sec. 2).

The VC framework has received much attention in past decades (see, *e.g.*, [3, 6, 31] and references therein), and many generalizations of the VC dimension have been proposed and studied (*e.g.*, [2, 7]). To the best of our knowledge, none of these generalizations captures the case of k -tuples of disjoint subsets of a universe considered in this paper. In addition, no connection was previously

made between the the VC dimension and the approximability of combinatorial auctions.

1.2 Follow-Up Work

Subsequently to this work [32], Buchfuhrer and Umans [10], and, independently, Dughmi *et al.* [20], strengthened our results via extensions of our techniques. Specifically, both works show that no constant approximation ratios are achievable for n -bidder case of the combinatorial auction environment we discuss in Sec. 3, thus proving a conjecture we presented in [32]. The techniques in these two works are quite different, as we briefly discuss below.

One of the main challenges in dealing with partitions of items is the possibility of unallocated items (as we shall later explain). Our lower bounds for VCG-based mechanisms fall into two categories: (1) lower bounds that rely on the assumption that the auctioneer allocates all items. (2) lower bounds for the general case that items may remain unallocated. While case (1) enables a straightforward use of the Sauer-Shelah Lemma, case (2) requires heavier machinery.

The work of Dughmi *et al.* [20] strengthens our results for case (2) via a clever and subtle generalization of our proof techniques. The main ingredients in the proof in [20] are: (i) An analogue of the Sauer-Shelah Lemma that incorporates sets of partial k -colorings, generalizing our Lemma 2.12, that can be regarded as the case of partial 2-colorings. (ii) Our lower bounds are based on a collection of sets such that every *pair* of them are “epsilon-close to orthogonal” (see Theorem 2.14). [20] get sharper lower bounds by considering collections of sets such that every ℓ -tuple of them are epsilon-close to orthogonal.

Buchfuhrer and Umans [10] overcome the problem of unallocated items via a clever “reduction” from case (2) to case (1). They use a counting argument that shows that, for any VCG-based mechanism that obtains a “good” approximation ratio, there must be a reasonably large subset of items that are fully allocated in exponentially many different ways. This enables an indirect application of the Sauer-Shelah Lemma.

1.3 Organization of the Paper

In Sec. 2 we present our approach to analyzing the VC dimension of partitions. In Sec. 3 we prove our inapproximability results for combinatorial auctions. We conclude and present open questions in Sec. 4.

2 The VC Dimension of Partitions

The main hurdle in combinatorial auctions stems from the fact that the outcomes are *partitions* of the set of items. Approximation algorithms for combinatorial auctions can be thought of as functions that map bidders’ preferences to partitions of items. This motivates our extension of the standard notion of VC dimension to *the VC dimension of partitions*. This section presents this notion of VC dimension and lower bounding techniques. In Sec. 3 we harness this machinery to prove computational complexity lower bounds for combinatorial auctions.

2.1 The VC Dimension and α -Approximate Collections of Partitions

We focus on partitions that consist of two disjoint subsets (our definitions can easily be extended to k -tuples). Our formal definition of a partition of a universe is the following:

Definition 2.1 (partitions) A partition $T = (T_1, T_2)$ of a universe $U = \{1, \dots, m\}$ 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 we *do not* require that every element in the universe appear in one of the two disjoint subsets that form a partition. This definition of partitions will later enable us to address crucial aspects of combinatorial auctions. We refer to partitions that *do* exhaust the universe (i.e., cover all elements in the universe) as “*covering partitions*” (we shall refer to not-necessarily-covering partitions as “*general partitions*”).

Definition 2.2 (covering partitions) 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 (projections) The projection of a partition $(S_1, S_2) \in P(U)$ on $E \subseteq U$, denoted by $(S_1, S_2)|_E$, is the partition $(S_1 \cap E, S_2 \cap E) \in P(E)$. For any collection of partitions $R \subseteq P(U)$ we define R 's projection on $E \subseteq U$, $R|_E$, to be $R|_E = \{(T_1, T_2) \mid \exists (S_1, S_2) \in R \text{ s.t. } (S_1, S_2)|_E = (T_1, T_2)\}$.

Observe that if a partition (S_1, S_2) of $E \subseteq U$ is in $C(E)$, then for any $E' \subseteq E$ $(S_1, S_2)|_{E'} \in C(E')$. We are now ready to define the notions of shattering and VC dimension in our context:

Definition 2.4 (shattering) A subset $E \subseteq U$ is said to be shattered by a collection of partitions $R \subseteq P(U)$ if $C(E) \subseteq R|_E$.

Observe that that if $E \subseteq U$ is shattered by a collection of partitions $R \subseteq P(U)$ then so are all subsets of E .

Definition 2.5 (VC dimension) The VC dimension $VC(R)$ of a collection of partitions $R \subseteq P(U)$ is the cardinality of the biggest subset $E \subseteq U$ that is shattered by R .

We now introduce the useful concept of α -approximate collections of partitions. Informally, a collection R of partitions is α -approximate if, for every partition S of the universe (not necessarily in R), there is some partition in R that is “not far” (in terms of α) from S . We are interested in the connection between the value of α of an α -approximate collection of partitions and its VC dimension. This will play a major role in the proofs of our results for combinatorial auctions.

Definition 1 (α -approximate collections of partitions) Let R be a collection of partitions of a universe U . R is said to be α -approximate if for every partition $S = (S_1, S_2) \in P(U)$ there exists some partition $T = (T_1, T_2) \in R$ such that $|S_1 \cap T_1| + |S_2 \cap T_2| \geq \alpha(|S_1| + |S_2|)$.

2.2 Lower Bounding the VC Dimension of Collections of Covering Partitions

When dealing with collections of covering partitions it is possible to use existing VC machinery to lower bound their VC dimension. Specifically, a straightforward application of the Sauer-Shelah Lemma [39, 41] implies that:

Lemma 2.6 (lower bounding the VC dimension of covering partitions) *For every $R \subseteq C(U)$ it holds that $VC(R) = \Omega(\frac{\log |R|}{\log |U|})$.*

Proof: Let $R_1 = \{S_1 \mid \exists S_2 \text{ s.t. } (S_1, S_2) \in R\}$. Because R only consists of covering partitions it must be that $|R_1| = |R|$. We now recall the Sauer-Shelah Lemma:

Lemma 2.7 ([39, 41]) *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, when applied to R_1 implies the existence of a set E (as in the statement of the lemma) of size $\Omega(\frac{\log |R_1|}{\log |U|})$ (that is, a large set that is shattered in the traditional sense). The fact that all partitions in R are covering partitions now immediately implies that $VC(R) = \Omega(\frac{\log |R_1|}{\log |U|}) = \Omega(\frac{\log |R|}{\log |U|})$. \square

Lemma 2.6 enables us to prove a lower bounds on the VC dimension of $(\frac{1}{2} + \epsilon)$ -approximate collections of *covering* partitions.

Theorem 2.8 *Let $R \subseteq C(U)$. If R is $(\frac{1}{2} + \epsilon)$ -approximate (for any small constant $\epsilon > 0$) then there exists some constant $\alpha > 0$ such that $VC(R) \geq m^\alpha$.*

Proof: [sketch] We use a probabilistic construction argument: Consider a partition of U , $S = (S_1, S_2)$, that is chosen, uniformly at random, out of all possible partitions in $C(U)$. Observe that (by the Chernoff bounds), for every partition $T = (T_1, T_2) \in R$, the probability that $|S_1 \cap T_1| + |S_2 \cap T_2| \geq (\frac{1}{2} + \epsilon)(|S_1| + |S_2|) = (\frac{1}{2} + \epsilon)m$ is exponentially small in m . Hence, for R to be $(\frac{1}{2} + \epsilon)$ -approximate it must contain exponentially many partitions in $C(U)$. We can now apply Lemma 2.6 to conclude the proof. \square

2.3 Lower Bounding the VC Dimension of Collections of General Partitions

Observe that the proof Lemma 2.6 heavily relied on the fact that the partitions considered were covering partitions. Dealing with collections of general partitions necessitates the development of different techniques for lower bounding the VC dimension. We shall now present such a method, that can be regarded as an extension of the Sauer-Shelah Lemma [39, 41] to the case of collections of general partitions:

Definition 2.9 (distance between partitions) *Given a universe U , two partitions in $P(U)$, (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 2.10 *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 2.11 *Suppose $k \geq 1$ and $\epsilon m \geq 1$. Then if $t(\epsilon, k, m) > \sum_{i=1}^r \binom{m}{i}$ 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 $t(\epsilon, k, m) \geq \sum_{i=0}^r \binom{m}{i}$ is a bound on the number of sets of size at most r . \square

Lemma 2.12 (lower bounding the VC dimension of general partitions) *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 [3, 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 collection 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. \square

Lemma 2.12 and Observation 2.11 imply the following important corollary:

Corollary 2.13 *For every constant $\alpha > 0$, and (sufficiently small) constant $\epsilon > 0$, there exists a $\beta > 0$ such that, if $R \subseteq P(U)$ and it holds that: (1) $|R| \geq e^{m^\alpha}$, and (2) every two partitions in R are ϵm -far, then $VC(R) \geq m^\beta$.*

We shall now discuss the connection between the value of α of an α -approximate collection of partitions and its VC dimension. We prove the following theorem:

Theorem 2.14 *Let $R \subseteq P(U)$. If R is $(\frac{3}{4} + \epsilon)$ -approximate (for any constant $\epsilon > 0$) then there exists some constant $\alpha > 0$ such that $VC(R) \geq m^\alpha$.*

Proof: To prove the theorem we use the following claim:

Claim 2.15 *For every small constant $\delta > 0$, 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 $\frac{1-\delta}{2}m$ -far.*

Proof: 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. [42]. 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 $\frac{m+\delta}{2}$ 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 $\frac{1-\delta}{2}m$ is exponentially small in δ . Hence, a family F of exponential size must exist. \square

Lemma 2.16 *Let $\epsilon > 0$. Let $R \subseteq P(U)$ such that R is $(\frac{3}{4} + \epsilon)$ -approximate. Then, there is a subset of R , R' , of size exponential in m such that every two elements of R' are at least αm -far (for some constant $\alpha > 0$).*

Proof: By Claim 2.15 we know that, for our universe of U , there exists an exponential-sized family of partitions in $C(U)$, F , such that every two partitions in F are at least $\frac{1-\delta}{2}m$ -far (for some arbitrarily small $\delta > 0$). Fix some $T, T' \in F$. By definition of F , T and T' are identical only on at most $\frac{1+\delta}{2}m$ elements (that is, only for at most $\frac{1+\delta}{2}m$ items j , either $j \in T_1 \cap T'_1$ or $j \in T_2 \cap T'_2$). Let R_T and $R_{T'}$ represent two partitions in R that obtain $\frac{3}{4} + \epsilon$ “approximations” for T and T' , respectively (because R is $(\frac{3}{4} + \epsilon)$ -approximate such partitions must exist). Even if we assume that both R^T and $R^{T'}$ are identical on all elements on which T and T' are identical, we are still left with $\frac{1-\delta}{2}m$ elements. Observe that for each such element, if R^T and $R^{T'}$ are identical on it, it holds that it can only contribute to the approximation obtained by *one* of them. This implies that to obtain the promised approximation R^T and $R^{T'}$ must differ on quite a lot (a constant fraction) of the elements in U . This, in turn, implies that there is some $\alpha > 0$ such that R^T and $R^{T'}$ are αm -far. \square

Corollary 2.13 now concludes the proof of the theorem. \square

3 Implications For VCG-Based Mechanisms

In this section we present the connection between our results for collections of partitions in Sec. 2 and the problem of social-welfare-maximization in combinatorial auctions. We use the VC dimension framework developed in the previous section to present a general technique for proving computational complexity lower bounds for VCG-based mechanisms.

3.1 Maximal-In-Range Mechanisms for Combinatorial Auctions

2-bidder combinatorial auctions. We consider auction environments of the following form: There is a set of items $1 \dots, m$ that are sold to 2 bidders, 1 and 2. Each bidder i has a private *valuation function* (sometimes simply referred to as a *valuation*) v_i that assigns a nonnegative 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 . Each v_i is assumed to be *nondecreasing*, *i.e.*, $\forall S \subseteq T$ it holds that $v_i(S) \leq v_i(T)$. The objective is find a partition of the items (S_1, S_2) between the two bidders that maximizes the social welfare, *i.e.*, the expression $\Sigma_i v_i(S_i)$.

It is known that optimizing the social welfare value in 2-bidder combinatorial auctions is computationally intractable even for very restricted classes of valuation functions. In particular, [29] shows that this task is NP-hard even for the simple class of *capped additive valuations*:

Definition 2 (additive valuations) A valuation function a is said to be additive if there exist per-item values a_{i_1}, \dots, a_{i_m} , such that for every bundle $S \subseteq [m]$, $a(S) = \sum_{j \in S} a_{ij}$.

Definition 3 (capped additive valuations) A valuation function v is said to be a capped additive valuation if there exist an additive valuation a and a real value B , such that for every bundle $S \subseteq [m]$, $v(S) = \min\{a(S), B\}$.

Intuitively, a bidder has a capped additive valuation if his value for each bundle of items is simply the additive sum of his values for the items in it, up to some maximum amount he is willing to spend. This class of valuations shall be used throughout this section to illustrate our impossibility results (as we aim to prove inapproximability results the restrictedness of this class works to our advantage).

Maximal-in-range mechanisms. Mechanisms that rely on the VCG technique to ensure truthfulness (VCG-based mechanisms) are known to have the useful combinatorial property of being *maximal-in-range* [15, 33, 37]:³ Maximal-in-range mechanisms are mechanisms that always *exactly* optimize over a (fixed) set of outcomes. In our context, this means that for every maximal-in-range mechanism M there exists some $R_M \subseteq P([m])$ such that M always outputs an optimal outcome *in* R_M (with respect to social-welfare maximization). We refer to R_M as M 's *range*.

Definition 4 (maximal-in-range mechanisms) A mechanism M is maximal-in-range if there is a collection of partitions $R_M \subseteq P([m])$ such that for every pair of valuations, (v_1, v_2) , M outputs a partition $(T_1, T_2) \in \operatorname{argmax}_{(S_1, S_2) \in R_M} v_1(S_1) + v_2(S_2)$.

It is known that every maximal-in-range mechanism can be made incentive compatible via the VCG technique [33, 34]. This suggests a general way for the design of truthful mechanisms for combinatorial auctions: Fix the range R_M of a maximal-in-range mechanism M to be such that (1) optimizing over R_M can be done in polynomial time, and (2) the optimal outcome *in* R_M always provides a “good” approximation to the *globally-optimal* outcome. This approach was shown to be useful in [16, 25].

Observe that the maximal-in-range mechanism in which R_M contains *all* possible partitions of items is computationally intractable even for capped additive valuations. In contrast, the fact that bidders' valuations are nondecreasing implies the following general upper bound:

Observation 3.1 (the trivial upper bound) For any 2-bidder combinatorial auction, the maximal-in-range mechanism M for which $R_M = \{([m], \emptyset), (\emptyset, [m])\}$ provides a $\frac{1}{2}$ -approximation to the optimal social welfare.

That is, the maximal-in-range mechanism that bundles all items together and allocates them to the bidder with the highest value provides a $\frac{1}{2}$ -approximation to the optimal social welfare. This mechanism is easy to implement in a computationally-efficient manner as it only requires learning the value of each bidder for the bundle of all items.

³Maximal-in-range mechanisms are a special case of a more general class of mechanisms called “affine maximizers” [37, 27]. All of the results in this paper actually apply to this more general class.

Is the trivial upper bound optimal? Naturally, we are interested in the question of whether a more clever choice of range than $\{([m], \emptyset), (\emptyset, [m])\}$ can lead to better approximation ratios (without jeopardizing computational efficiency). Let us consider 2-bidder combinatorial auctions with capped additive valuations. For this restricted case, a non-truthful PTAS exists [4]. Can a similar result be obtained via a maximal-in-range mechanism? We show that the answer to this question is *No* by proving that the approximation ratios obtained by computationally-efficient VCG-based mechanisms are always bounded away from 1. We stress that these are the first computational complexity lower bounds on the approximability of VCG-based mechanisms for combinatorial auctions. In fact, as we shall later show, in certain cases these bounds extend to *all* incentive-compatible mechanisms.

3.2 Putting the VC in VCG

We now present our method of proving lower bounds on the approximability of VCG-based mechanisms using the VC framework. On a high level, our technique for proving that a maximal-in-range mechanism M cannot obtain an α -approximation consists of three steps:

- Observe that M 's range must be an α -approximate collection of partitions.
- Conclude (from our results in Sec. 2) the existence of a shattered set of items of size m^α (if α is sufficiently high).
- Show a non-uniform reduction from NP-hard 2-bidder combinatorial auctions with m^α items to the optimization problem solved by M .

We illustrate these three steps by proving a lower bound of $\frac{3}{4}$ for 2-bidder combinatorial auctions with capped additive valuations (which naturally extends to the more general classes of submodular, and subadditive, valuations). We stress that our proof technique can be applied to prove the *same* lower bound for practically any NP-hard 2-bidder combinatorial auction environment. Essentially, our only requirement from the class of valuations is that it be expressive enough to contain the class of 0/1-additive valuations defined below.

Definition 5 (0/1-additive valuations) *A valuation v is said to be 0/1-additive if it is an additive valuation in which all the per-item values are in $\{0, 1\}$.*

We make the following observation:

Observation 3.2 *Any α -approximation maximal-in-range mechanism for 2-bidder combinatorial auctions with 0/1-additive valuations has a range that is an α -approximate collection of partitions.*

Proof: A 0/1-additive valuation can be regarded as an indicator function that specifies some subset of the universe (that contains only the items that are assigned a value of 1). Hence, pairs of such valuations that specify disjoint subsets correspond to partitions of the universe. Now, it is easy to see that, by definition, the range of an α -approximation maximal-in-range mechanism must be an α -approximate collection of partitions. \square

Observation 3.2 enables us to make use of Theorem 2.14 to conclude that:

Theorem 3.3 *The range of any $(\frac{3}{4} + \epsilon)$ -approximation maximal-in-range mechanism for 2-bidder combinatorial auctions with 0/1-additive valuations shatters a set of items of size m^α (for some constant $\alpha > 0$).*

We can now exploit the existence of a large shattered set of items to prove our lower bound by showing a non-uniform reduction from an NP-hard optimization problem:

Theorem 3.4 *No polynomial-time maximal-in-range mechanism obtains an approximation ratio of $\frac{3}{4} + \epsilon$ for 2-bidder combinatorial auctions with capped additive valuations unless $NP \subseteq P/poly$.*

Proof: Let M be a mechanism as in the statement of the theorem. Since 0/1-additive valuations are a special case of capped additive valuations, by Theorem 3.3 there exists a constant $\alpha > 0$ such that R_m shatters a set of items E of size m^α . Therefore, given an auction with m^α items and capped additive 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 non-uniform reduction from an NP-hard problem (social-welfare maximization in the smaller auction) to the optimization problem solved by M . \square

Recall that the trivial upper bound provides an approximation ratio of $\frac{1}{2}$. We leave the problem of closing the gap between this upper bound and our lower bound open. We conjecture that the trivial upper bound is, in fact, tight. This conjecture is motivated by the following result.

The allocate-all-items case. We now consider the well-studied case that the auctioneer must allocate all items [19, 27]. Observe that, in this case, the range of a maximal-in-range mechanism can only consist of covering partitions, for which stronger results are obtained in Sec. 2. This enables us to use our technique to prove the following theorem:

Theorem 3.5 *For the allocate-all-items case, no polynomial-time maximal-in-range mechanism obtains an approximation ratio of $2 - \epsilon$ for 2-bidder combinatorial auctions with capped additive valuations unless $NP \subseteq P/poly$.*

If bidders have subadditive valuations, and all items are allocated, then maximal-in-range mechanisms are the *only* truthful mechanisms [19]. Since capped additive valuations are a special case of subadditive valuations, the lower bound in Theorem 3.5 holds for *all* truthful mechanisms in this more general environment. In Appendix A, we show that, for a superclass of capped additive valuations, it is possible to relax the computational assumption in the statement of Theorem 3.5 to the assumption that NP is not contained in BPP. This is achieved by using Ajtai's [1] probabilistic version of the Sauer-Shelah Lemma.

4 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) Can our inapproximability results be made to hold for all truthful mechanisms? So far, despite much work on this subject [37, 27, 28, 8, 19, 36], very little is known about characterizations of truthfulness in combinatorial auctions (and in other multi-parameter environments). (2) Our computational complexity results depend on the assumption that SAT does not have polynomial-size circuits. Can this assumption be relaxed by proving probabilistic versions of our VC machinery (see Appendix A)?

References

- [1] M. Ajtai. The shortest vector problem in L_2 is NP-hard for randomized reductions (extended abstract). In *Proceedings of the 30th annual ACM symposium on Theory of computing*, pages 10–19, 1998.
- [2] N. Alon. On the density of sets of vectors. *Discrete Mathematics*, 46(2):199–202, 1983.
- [3] N. Alon, S. Ben-David, N. Cesa-Bianchi, and D. Haussler. Scale-sensitive dimensions, uniform convergence, and learnability. *J. ACM*, 44(4):615–631, 1997.
- [4] N. Andelman and Y. Mansour. Auctions with budget constraints. In *Proceedings of the 9th Scandinavian Workshop on Algorithm Theory*, pages 26–38, 2004.
- [5] Y. Bartal, R. Gonen, and N. Nisan. Incentive compatible multi unit combinatorial auctions. In *Proceedings of the 9th conference on Theoretical aspects of rationality and knowledge*, pages 72–87, 2003.
- [6] P. L. Bartlett and P. M. Long. Prediction, learning, uniform convergence, and scale-sensitive dimensions. *J. Comput. Syst. Sci.*, 56(2):174–190, 1998.
- [7] S. Ben-David, N. Cesa-Bianchi, and P. M. Long. Characterizations of learnability for classes of $O, \$ \dots \$$, n -valued functions. In *COLT '92: Proceedings of the fifth annual workshop on Computational learning theory*, pages 333–340. ACM, 1992.
- [8] S. Bikhchandani, S. Chatterji, R. Lavi, A. Mu'alem, N. Nisan, and Arunava Sen. Weak monotonicity characterizes deterministic dominant-strategy implementation. *Econometrica*, 74(4):1109–1132, 2006.
- [9] L. Blumrosen and N. Nisan. Combinatorial auctions. In N. Nisan, T. Roughgarden, E. Tardos, and V. Vazirani, editors, *Algorithmic Game Theory*, chapter 11. Cambridge University Press, 2007.
- [10] D. Buchfuhrer and C. Umans. Limits on the social welfare of maximal-in-range auction mechanisms. 2009. Working paper.
- [11] E. H. Clarke. Multipart pricing of public goods. *Public Choice*, pages 17–33, 1971.
- [12] P. Cramton, Y. Shoham, and R. Steinberg (Editors). *Combinatorial Auctions*. MIT Press., 2006.
- [13] S. de Vries and R. Vohra. Combinatorial auctions: A survey. *INFORMS journal of computing*, 15(3):284–309, 2003.
- [14] S. Dobzinski. Two randomized mechanisms for combinatorial auctions. In *Proceedings of the 10th International Workshop on Approximation and the 11th International Workshop on Randomization, and Combinatorial Optimization. Algorithms and Techniques*, pages 89–103, 2007.
- [15] S. Dobzinski and N. Nisan. Limitations of VCG-based mechanisms. In *Proceedings of the 39th annual ACM symposium on Theory of computing*, pages 338–344, 2007.
- [16] S. Dobzinski, N. Nisan, and M. Schapira. Approximation algorithms for combinatorial auctions with complement-free bidders. In *Proceedings of the 37th annual ACM symposium on Theory of computing*, pages 610–618, 2005.

- [17] S. Dobzinski, N. Nisan, and M. Schapira. Truthful randomized mechanisms for combinatorial auctions. In *Proceedings of the 38th annual ACM symposium on Theory of computing*, 2006.
- [18] S. Dobzinski and M. Schapira. An improved approximation algorithm for combinatorial auctions with submodular bidders. In *Proceedings of the 17th annual ACM-SIAM symposium on Discrete algorithm*, pages 1064–1073, 2006.
- [19] S. Dobzinski and M. Sundararajan. On characterizations of truthful mechanisms for combinatorial auctions and scheduling. In *Proceedings of the 9th ACM conference on Electronic commerce*, pages 38–47, 2008.
- [20] S. Dughmi, H. Fu, and R. Kleinberg. Amplified hardness of approximation for VCG-based mechanisms. 2009. Working paper.
- [21] U. Feige. On maximizing welfare when utility functions are subadditive. In *Proceedings of the 38th annual ACM symposium on Theory of computing*, pages 41–50, 2006.
- [22] U. Feige and J. Vondrak. Approximation algorithms for allocation problems: Improving the factor of $1 - 1/e$. In *Proceedings of the 47th Annual IEEE Symposium on Foundations of Computer Science*, pages 667–676, 2006.
- [23] T. Groves. Incentives in teams. *Econometrica*, pages 617–631, 1973.
- [24] J. Håstad. Some optimal inapproximability results. *J. ACM*, 48(4):798–859, 2001.
- [25] R. Holzman, N. Kfir-Dahav, D. Monderer, and M. Tennenholtz. Bundling equilibrium in combinatorial auctions. *Games and Economic Behavior*, 47:104–123, 2004.
- [26] E. Kushilevitz and N. Nisan. *Communication Complexity*. Cambridge University Press, 1997.
- [27] R. Lavi, A. Mu’alem, and N. Nisan. Towards a characterization of truthful combinatorial auctions. In *Proceedings of the 44th Annual IEEE Symposium on Foundations of Computer Science*, page 574, 2003.
- [28] R. Lavi, A. Mu’alem, and N. Nisan. Two simplified proofs for Roberts’ Theorem, 2004. Submitted.
- [29] B. Lehmann, D. Lehmann, and N. Nisan. Combinatorial auctions with decreasing marginal utilities. *Games and Economic Behaviour*, 55(2):270–296, 2006.
- [30] D. Lehmann, L. I. O’Callaghan, and Y. Shoham. Truth revelation in approximately efficient combinatorial auctions. In *J.ACM 49(5)*, pages 577–602, Sept. 2002.
- [31] S. Mendelson and R. Vershynin. Entropy, combinatorial dimensions and random averages. In *COLT*, pages 14–28, 2002.
- [32] E. Mossel, C. H. Papadimitriou, M. Schapira, and Y. Singer. VC v. VCG: Inapproximability of combinatorial auctions via generalizations of the VC dimension. *CoRR*, abs/0905.1995, 2009.
- [33] N. Nisan and A. Ronen. Computationally feasible VCG-based mechanisms. In *Proceedings of the 2nd ACM conference on Electronic commerce*, pages 242–252, 2000.
- [34] N. Nisan and A. Ronen. Algorithmic mechanism design. *Games and Economic Behaviour*, 35:166 – 196, 2001.

- [35] N. Nisan and I. Segal. The communication requirements of efficient allocations and supporting prices. *Journal of Economic Theory*, 129:192–224, 2006.
- [36] C. H. Papadimitriou, M. Schapira, and Y. Singer. On the hardness of being truthful. In *Proceedings of the 49th Annual IEEE Symposium on Foundations of Computer Science*, pages 250–259, 2008.
- [37] K. 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.
- [38] T. Roughgarden. An algorithmic game theory primer. In *Proceedings of the 5th IFIP International Conference on Theoretical Computer Science. An invited survey*, 2008.
- [39] N. Sauer. On the density of families of sets. *J. Comb. Theory, Ser. A*, 13(1):145–147, 1972.
- [40] M. Schapira and Y. Singer. Inapproximability of combinatorial public projects. In *Proceedings of the Workshop on Internet Economics*, pages 351–361, 2008.
- [41] S. Shelah. A combinatorial problem; stability and order for models and theories in infinitary languages. *Pacific J Math*, 41:247–261, 1972.
- [42] Madhu Sudan. Essential coding theory, lecture notes, 2008. (MIT 6.440).
- [43] W. Vickrey. Counterspeculation, auctions and competitive sealed tenders. *Journal of Finance*, pages 8–37, 1961.
- [44] J. Vondrák. Optimal approximation for the submodular welfare problem in the value oracle model. In *Proceedings of the 40th annual ACM symposium on Theory of computing*, pages 67–74, 2008.
- [45] A. C. Yao. Some complexity questions related to distributive computing (preliminary report). In *Proceedings of the 11th annual ACM symposium on Theory of computing*, pages 209–213, 1979.

A Strengthening Theorem 3.5

We shall now show how, if the valuation functions of the bidders are slightly more expressive than capped additive valuations, one can obtain a lower bound as in 3.5 dependent on the weaker computational assumption that NP is not contained in BPP. This is achieved by using the probabilistic version of the Sauer-Shelah Lemma presented by Ajtai [1] to obtain a *probabilistic* polynomial-time reduction from an NP-hard problem to the problem solved by the maximal-in-range mechanism. We are currently unable to prove a similar result for capped additive valuations.

We consider the class of double-capped additive valuations. Informally, a bidder has a double-capped additive valuation if he has an additive valuation, but also has some upper bound on how much he is willing to spend on different subsets of items, as well as a global upper bound on how much he is willing to spend overall.

Definition 6 (double-capped additive valuations) *A valuation function v is said to be a double-capped additive valuation if there exists a partition of the set of items $[m]$ into disjoint subsets S_1, \dots, S_r (that cover all items), an additive valuation a , and real values B, B_1, \dots, B_r , such that for every bundle $S \subseteq [m]$, $v(S) = \min\{\sum_{t=1}^r \min\{\sum_{j \in S_t} a(j), B_t\}, B\}$.*

We prove the following theorem:

Theorem A.1 *For the allocate-all-items case, no polynomial-time maximal-in-range mechanism obtains an approximation ratio of $2 - \epsilon$ for 2-bidder combinatorial auctions with double-capped additive valuations unless $NP \subseteq BPP$.*

Proof: [sketch] Let R be the range of a $(2 - \epsilon)$ -approximation maximal-in-range mechanism. In the proof of Theorem 2.8 we show that R must be of exponential size. Let R_1 denote the subsets of items bidder 1 is assigned in R . Because the allocate-all-items assumption holds then, as explained in Lemma 2.6, $|R_1|$ is also of exponential size. In Lemma 2.6, we used the Sauer-Shelah Lemma to conclude that there must be a large set E that is shattered in the traditional sense. Now, we make use of Ajtai's probabilistic version of the Sauer-Shelah Lemma:

Lemma A.2 ([1]) *Let Z be a family of subsets of a universe R that is regular (i.e., all subsets in Z are of equal size) and $|Z| \geq 2^{|R|^\alpha}$ (for some $0 < \alpha \leq 1$). There are integers q, l (where $|R|, q$ and l are polynomially related) such that if we randomly choose q pairwise-disjoint subsets of R , Q_1, \dots, Q_q , each of size l , then, w.h.p., for every function $f : [q] \rightarrow \{0, 1\}$ there is a subset $Z' \in Z$ for which $|Z' \cap Q_j| = f(j)$ for all $j \in [q]$.*

Using Ajtai's Lemma (with $Z = R_1$), we conclude that there must be sets of items Q_1, \dots, Q_q as in the statement of the lemma. Now, a reduction similar to that in Theorem 3.5, in which each item in the smaller auction is identified with *all* items in a specific Q_s concludes the proof of the Theorem. \square