

CS369E: Communication Complexity (for Algorithm Designers)

Lecture #7: Lower Bounds in Algorithmic Game Theory*

Tim Roughgarden[†]

March 5, 2015

1 Preamble

This lecture explains some applications of communication complexity to proving lower bounds in *algorithmic game theory (AGT)*, at the border of computer science and economics. In AGT, the natural description size of an object is often exponential in a parameter of interest, and the goal is to perform non-trivial computations in time polynomial in the parameter (i.e., logarithmic in the description size). As we know, communication complexity is a great tool for understanding when non-trivial computations require looking at most of the input.

2 The Welfare Maximization Problem

The focus of this lecture is the following optimization problem, which has been studied in AGT more than any other.

1. There are k players.
2. There is a set M of m items.
3. Each player i has a *valuation* $v_i : 2^M \rightarrow \mathbb{R}_+$. The number $v_i(T)$ indicates i 's value, or willingness to pay, for the items $T \subseteq M$. The valuation is the private input of player i — i knows v_i but none of the other v_j 's. We assume that $v_i(\emptyset) = 0$ and that the valuations are *monotone*, meaning $v_i(S) \leq v_i(T)$ whenever $S \subseteq T$. To avoid bit complexity issues, we'll also assume that all of the $v_i(T)$'s are integers with description length polynomial in k and m .

*©2015, Tim Roughgarden.

[†]Department of Computer Science, Stanford University, 474 Gates Building, 353 Serra Mall, Stanford, CA 94305. Email: tim@cs.stanford.edu.

Note that we may have more than two players — more than just Alice and Bob. Also note that the description length of a player’s valuation is exponential in the number of items m .

In the *welfare-maximization problem*, the goal is to partition the items M into sets T_1, \dots, T_k to maximize, at least approximately, the welfare

$$\sum_{i=1}^k v_i(T_i),$$

using communication polynomial in n and m . Note this amount of communication is logarithmic in the sizes of the private inputs.

The main motivation for this problem is combinatorial auctions. Already in the domain of government spectrum auctions, dozens of such auctions have raised hundreds of billions of dollars of revenue. They have also been used for other applications such as allocating take-off and landing slots at airports. For example, items could represent licenses for wireless spectrum — the right to use a certain frequency range in a certain geographic area. Players would then be wireless telecommunication companies. The value $v_i(S)$ would be the amount of profit company i expects to be able to extract from the licenses in S .

Designing good combinatorial auctions requires careful attention to “incentive issues,” making the auctions as robust as possible to strategic behavior by the (self-interested) participants. Incentives won’t play much of a role in this lecture. Our lower bounds for protocols in Section 4 apply even in the ideal case where players are fully cooperative. Our lower bounds for equilibria in Section 5 effectively apply no matter how incentive issues are resolved.

3 Multi-Party Communication Complexity

3.1 The Model

Welfare-maximization problems have an arbitrary number k of players, so lower bounds for them follow most naturally from lower bounds for *multi-party* communication protocols. The extension from two to many parties proceeds as one would expect, so we’ll breeze through the relevant points without much fuss.

Suppose we want to compute a Boolean function $f : \{0, 1\}^{n_1} \times \{0, 1\}^{n_2} \times \dots \times \{0, 1\}^{n_k} \rightarrow \{0, 1\}$ that depends on the k inputs $\mathbf{x}_1, \dots, \mathbf{x}_k$. We’ll be interested in the *number-in-hand (NIH)* model, where player i only knows \mathbf{x}_i . What other model could there be, you ask? There’s also the stronger *number-on-forehead (NOF)* model, where player i knows everything except \mathbf{x}_i . (Hence the name — imagine the players are sitting in a circle.) The NOF model is studied mostly for its connections to circuit complexity; it has few direct algorithmic applications, so we won’t discuss it in this course. The NIH model is the natural one for our purposes and, happily, it’s also much easier to prove strong lower bounds for it.

Deterministic protocols are defined as you would expect, with the protocol specifying whose turn it is speak (as a function of the protocol’s transcript-so-far) and when the computation is complete. We’ll use the *blackboard model*, where we think of the bits sent by each

player as being written on a blackboard in public view.¹ Similarly, in a nondeterministic protocol, the prover writes a proof on the blackboard, and the protocol accepts the input if and only if all k players accept the proof.

3.2 The MULTI-DISJOINTNESS Problem

We need a problem that is hard for multi-party communication protocols. An obvious idea is to use an analog of DISJOINTNESS. There is some ambiguity about how to define a version of DISJOINTNESS for three or more players. For example, suppose there are three players, and amongst the three possible pairings of them, two have disjoint sets while the third have intersecting sets. Should this count as a “yes” or “no” instance? We’ll skirt this issue by worrying only about unambiguous inputs, that are either “totally disjoint” or “totally intersecting.”

Formally, in the MULTI-DISJOINTNESS problem, each of the k players i holds an input $\mathbf{x}_i \in \{0, 1\}^n$. (Equivalently, a set $S_i \subseteq \{1, 2, \dots, n\}$.) The task is to correctly identify inputs that fall into one of the following two cases:

- (1) “Totally disjoint,” with $S_i \cap S_{i'} = \emptyset$ for every $i \neq i'$.
- (0) “Totally intersecting,” with $\bigcap_{i=1}^k S_i \neq \emptyset$.

When $k = 2$, this is just DISJOINTNESS. When $k > 2$, there are inputs that are neither 1-inputs nor 0-inputs. We let protocols off the hook on such ambiguous inputs — they can answer “1” or “0” with impunity.

In the next section, we’ll prove the following communication complexity lower bound for MULTI-DISJOINTNESS, credited to Jaikumar Radhakrishnan and Venkatesh Srinivasan in Nisan [14].

Theorem 3.1 *The nondeterministic communication complexity of MULTI-DISJOINTNESS, with k players with n -bit inputs, is $\Omega(n/k)$.*

The nondeterministic lower bound is for verifying a 1-input. (It is easy to verify a 0-input — the prover just suggests the index of an element r in $\bigcap_{i=1}^k S_i$, the validity of which is easily checked privately by each of the players.)

In our application in Section 4, we’ll be interested in the case where k is much smaller than n , such as $k = \Theta(\log n)$. Intuition might suggest that the lower bound should be $\Omega(n)$ rather than $\Omega(n/k)$, but this is incorrect — a slightly non-trivial argument shows that Theorem 3.1 is tight for nondeterministic protocols (for all small enough k , like $k = O(\sqrt{n})$). See the Homework for details. This factor- k difference won’t matter for our applications, however.

¹In the weaker *message-passing model*, players communicate by point-to-point messages rather than via broadcast.

3.3 Proof of Theorem 3.1

The proof of Theorem 3.1 has three steps, all of which are generalizations of familiar arguments.

Step 1: *Every deterministic protocol with communication cost c induces a partition of $M(f)$ into at most 2^c monochromatic boxes.* By “ $M(f)$,” we mean the k -dimensional array in which the i th dimension is indexed by the possible inputs of player i , and an array entry contains the value of the function f on the corresponding joint input. By a “box,” we mean the k -dimensional generalization of a rectangle — a subset of inputs that can be written as a product $A_1 \times A_2 \times \dots \times A_k$. By “monochromatic,” we mean a box that does not contain both a 1-input and a 0-input. (Recall that for the MULTI-DISJOINTNESS problem there are also wildcard (“*”) inputs — a monochromatic box can contain any number of these.)

The proof of this step is the same as in the two-party case. We just run the protocol and keep track of the joint inputs that are consistent with the transcript. The box of all inputs is consistent with the empty transcript, and the box structure is preserved inductively: when player i speaks, it narrows down the remaining possibilities for the input \mathbf{x}_i , but has no effect on the possible values of the other inputs. Thus every transcript corresponds to a box, with these boxes partitioning $M(f)$. Since the protocol’s output is constant over such a box and the protocol computes f , all of the boxes it induces are monochromatic with respect to $M(f)$.

Similarly, every nondeterministic protocol with communication cost c (for verifying 1-inputs) induces a cover of the 1-inputs of $M(f)$ by at most 2^c monochromatic boxes.

Step 2: *The number of 1-inputs in $M(f)$ is $(k + 1)^n$.* This step and the next are easy generalizations of our second proof of our nondeterministic communication complexity lower bounds for DISJOINTNESS (from Lecture #6): first we lower bound the number of 1-inputs, then we upper bound the number of 1-inputs that can coexist in a single 1-box. In a 1-input $(\mathbf{x}_1, \dots, \mathbf{x}_k)$, for every coordinate ℓ , at most one of the k inputs has a 1 in the ℓ th coordinate. This yields $k + 1$ options for each of the n coordinates, thereby generating a total of $(k + 1)^n$ 1-inputs.

Step 3: *The number of 1-inputs in a monochromatic box is at most k^n .* Let $B = A_1 \times A_2 \times \dots \times A_k$ be a 1-box. The key claim here is: for each coordinate $\ell = 1, \dots, n$, there is a player $i \in \{1, \dots, k\}$ such that, for every input $\mathbf{x}_i \in A_i$, the ℓ th coordinate of \mathbf{x}_i is 0. That is, to each coordinate we can associate an “ineligible player” that, in this box, never has a 1 in that coordinate. This is easily seen by contradiction: otherwise, there exists a coordinate ℓ such that, for every player i , there is an input $\mathbf{x}_i \in A_i$ with a 1 in the ℓ th coordinate. As a box, this means that B contains the input $(\mathbf{x}_1, \dots, \mathbf{x}_k)$. But this is a 0-input, contradicting the assumption that B is a 1-box.

The claim implies the stated upper bound. Every 1-input of B can be generated by choosing, for each coordinate ℓ , an assignment of at most one “1” in this coordinate to one of the $k - 1$ eligible players for this coordinate. With only k choices per coordinate, there are at most k^n 1-inputs in the box B .

Conclusion: Steps 2 and 3 imply that covering of the 1s of the k -dimensional array of the MULTI-DISJOINTNESS function requires at least $(1 + \frac{1}{k})^n$ 1-boxes. By the discussion in Step 1, this implies a lower bound of $n \log_2(1 + \frac{1}{k}) = \Theta(n/k)$ on the nondeterministic communication complexity of the MULTI-DISJOINTNESS function (and output 1). This concludes the proof of Theorem 3.1.

Remark 3.2 (Randomized Communication Complexity of MULTI-DISJOINTNESS)
 Randomized protocols with two-sided error also require communication $\Omega(n/k)$ to solve MULTI-DISJOINTNESS [8, 2].² This generalizes the $\Omega(n)$ lower bound that we stated (but did not prove) in Lecture #4, so naturally we're not going to prove this lower bound either. Extending the lower bound for DISJOINTNESS to MULTI-DISJOINTNESS requires significant work, but it is a smaller step than proving from scratch a linear lower bound for DISJOINTNESS [10, 15]. This is especially true if one settles for the weaker lower bound of $\Omega(n/k^4)$ [1], which is good enough for our purposes in this lecture.

4 Lower Bounds for Approximate Welfare Maximization

4.1 General Valuations

We now put Theorem 3.1 to work and prove that it is impossible to obtain a non-trivial approximation of the general welfare-maximization problem with a subexponential (in m) amount of communication. First, we observe that a k -approximation is trivial. The protocol is to give the full set of items M to the player with the largest $v_i(M)$. This protocol can clearly be implemented with a polynomial amount of communication. To prove the approximation guarantee, consider a partition T_1, \dots, T_k of M with the maximum-possible welfare W^* . There is a player i with $v_i(T_i) \geq W^*/k$. The welfare obtained by our simple protocol is at least $v_i(M)$; since we assume that valuations are monotone, this is at least $v_i(T_i) \geq W^*/k$.

To apply communication complexity, it is convenient to turn the optimization problem of welfare maximization into a decision problem. In the WELFARE-MAXIMIZATION(k) problem, the goal is to correctly identify inputs that fall into one of the following two cases:

- (1) Every partition (T_1, \dots, T_k) of the items has welfare at most 1.
- (0) There exists a partition (T_1, \dots, T_k) of the items with welfare at least k .

Clearly, communication lower bounds for WELFARE-MAXIMIZATION(k) apply more generally to the problem of obtaining a better-than- k -approximation of the maximum welfare.

We prove the following.

Theorem 4.1 ([14]) *The communication complexity of WELFARE-MAXIMIZATION(k) is $\exp\{\Omega(m/k^2)\}$.*

²There is also a far-from-obvious matching upper bound of $O(n/k)$ [9, 2].

Thus, if the number of items m is at least $k^{2+\epsilon}$ for some $\epsilon > 0$, then the communication complexity of the WELFARE-MAXIMIZATION(k) problem is exponential. Because the proof is a reduction from MULTI-DISJOINTNESS, the lower bound applies to deterministic protocols, nondeterministic protocols (for the output 1), and randomized protocols with two-sided error.

The proof of Theorem 4.1 relies on Theorem 3.1 and a combinatorial gadget. We construct this gadget using the probabilistic method. As a thought experiment, consider t random partitions P^1, \dots, P^t of M , where t is a parameter to be defined later. By a random partition $P^j = (P_1^j, \dots, P_k^j)$, we just mean that each of the m items is assigned to exactly one of the k players, independently and uniformly at random.

We are interested in the probability that two classes of different partitions intersect: for all $i \neq i'$ and $j \neq \ell$, since the probability that a given item is assigned to i in P^j and also to i' in P^ℓ is $\frac{1}{k^2}$, we have

$$\Pr[P_i^j \cap P_{i'}^\ell = \emptyset] = \left(1 - \frac{1}{k^2}\right)^m \leq e^{-m/k^2}.$$

Taking a Union Bound over the k choices for i and i' and the t choices for j and ℓ , we have

$$\Pr[\exists i \neq i', j \neq \ell \text{ s.t. } P_i^j \cap P_{i'}^\ell = \emptyset] \leq k^2 t^2 e^{-m/k^2}. \quad (1)$$

Call P^1, \dots, P^t an *intersecting family* if $P_i^j \cap P_{i'}^\ell \neq \emptyset$ whenever $i \neq i', j \neq \ell$. By (1), the probability that our random experiment fails to produce an intersecting family is less than 1 provided $t < \frac{1}{k} e^{m/2k^2}$. The following lemma is immediate.

Lemma 4.2 *For every $m, k \geq 1$, there exists an intersecting family of partitions P^1, \dots, P^t with $t = \exp\{\Omega(m/k^2)\}$.*

A simple combination of Theorem 3.1 and Lemma 4.2 proves Theorem 4.1.

Proof of Theorem 4.1: The proof is a reduction from MULTI-DISJOINTNESS. Fix k and m . (To be interesting, m should be significantly bigger than k^2 .) Let (S_1, \dots, S_k) denote an input to MULTI-DISJOINTNESS with t -bit inputs, where $t = \exp\{\Omega(m/k^2)\}$ is the same value as in Lemma 4.2. We can assume that the players have coordinated in advance on an intersecting family of t partitions of a set M of m items. Each player i uses this family and its input S_i to form the following valuation:

$$v_i(T) = \begin{cases} 1 & \text{if } T \supseteq P_i^j \text{ for some } j \in S_i \\ 0 & \text{otherwise.} \end{cases}$$

That is, player i is either happy (value 1) or unhappy (value 0), and is happy if and only if it receives all of the items in the corresponding class P_i^j of some partition P^j with index j belonging to its input to MULTI-DISJOINTNESS. The valuations v_1, \dots, v_k define an input to WELFARE-MAXIMIZATION(k). Forming this input requires no communication between the players.

Consider the case where the input to MULTI-DISJOINTNESS is a 1-input, with $S_i \cap S_{i'} = \emptyset$ for every $i \neq i'$. We claim that the induced input to WELFARE-MAXIMIZATION(k) is a 1-input, with maximum welfare at most 1. To see this, consider a partition (T_1, \dots, T_k) in which some player i is happy (with $v_i(T_i) = 1$). For some $j \in S_i$, player i receives all the items in P_i^j . Since $j \notin S_{i'}$ for every $i' \neq i$, the only way to make a second player i' happy is to give it all the items in $P_{i'}^\ell$ in some other partition P^ℓ with $\ell \in S_{i'}$ (and hence $\ell \neq j$). Since P^1, \dots, P^t is an intersecting family, this is impossible — P_i^j and $P_{i'}^\ell$ overlap for every $\ell \neq j$.

When the input to MULTI-DISJOINTNESS is a 0-input, with an element r in the mutual intersection $\bigcap_{i=1}^k S_i$, we claim that the induced input to WELFARE-MAXIMIZATION(k) is a 0-input, with maximum welfare at least k . This is easy to see: for $i = 1, 2, \dots, k$, assign the items of P_i^r to player i . Since $r \in S_i$ for every i , this makes all k players happy.

This reduction shows that a (deterministic, nondeterministic, or randomized) protocol for WELFARE-MAXIMIZATION(k) yields one for MULTI-DISJOINTNESS (with t -bit inputs) with the same communication. We conclude that the communication complexity of WELFARE-MAXIMIZATION(k) is $\Omega(t/k) = \exp\{\Omega(m/k^2)\}$. ■

4.2 Subadditive Valuations

To an algorithms person, Theorem 4.1 is depressing, as it rules out any non-trivial positive results. A natural idea is to seek positive results by imposing additional structure on players' valuations. Many such restrictions have been studied. We consider here the case of *subadditive* valuations, where each v_i satisfies $v_i(S \cup T) \leq v_i(S) + v_i(T)$ for every pair $S, T \subseteq M$.

Our reduction in Theorem 4.1 immediately yields a weaker inapproximability result for welfare maximization with subadditive valuations. Formally, define the WELFARE-MAXIMIZATION(2) problem as that of identifying inputs that fall into one of the following two cases:

- (1) Every partition (T_1, \dots, T_k) of the items has welfare at most $k + 1$.
- (0) There exists a partition (T_1, \dots, T_k) of the items with welfare at least $2k$.

Communication lower bounds for WELFARE-MAXIMIZATION(2) apply to the problem of obtaining a better-than-2-approximation of the maximum welfare.

Corollary 4.3 ([4]) *The communication complexity of WELFARE-MAXIMIZATION(2) is $\exp\{\Omega(m/k^2)\}$, even when all players have subadditive valuations.*

Proof: Picking up where the reduction in the proof of Theorem 4.3 left off, every player i adds 1 to its valuation for every non-empty set of items. Thus, the previously 0-1 valuations become 0-1-2 valuations that are only 0 for the empty set. Such functions always satisfy the subadditivity condition ($v_i(S \cup T) \leq v_i(S) + v_i(T)$). 1-inputs and 0-inputs of MULTI-DISJOINTNESS now become 1-inputs and 0-inputs of WELFARE-MAXIMIZATION(2), respectively. The communication complexity lower bound follows. ■

	Rock	Paper	Scissors
Rock	0,0	-1,1	1,-1
Paper	1,-1	0,0	-1,1
Scissors	-1,1	1,-1	0,0

Figure 1: Player utilities in Rock-Paper-Scissors. The pair of numbers in a matrix entry denote the utilities of the row and column players, respectively, in a given outcome.

There is also a quite non-trivial matching upper bound of 2 for deterministic, polynomial-communication protocols [5].

5 Lower Bounds for Equilibria

The lower bounds of the previous section show that every protocol for the welfare-maximization problem that interacts with the players and then explicitly computes an allocation has either a bad approximation ratio or high communication cost. Over the past five years, many researchers have aimed to shift the work from the protocol to the players, by analyzing the equilibria of simple auctions. Can such equilibria bypass the communication complexity lower bounds proved in Section 4? The answer is not obvious, because equilibria are defined non-constructively, and not through a low-communication protocol.³

5.1 Game Theory

Next we give the world’s briefest-ever game theory tutorial. See e.g. [17], or the instructor’s CS364A lecture notes, for a more proper introduction. We’ll be brief because the details of these concepts do not play a first-order role in the arguments below.

5.1.1 Games

A (*finite, normal-form*) *game* is specified by:

1. A finite set of $k \geq 2$ *players*.
2. For each player i , a finite *action set* A_i .
3. For each player i , a *utility function* $u_i(\mathbf{a})$ that maps an action profile $\mathbf{a} \in A_1 \times \dots \times A_k$ to a real number. The utility of a player generally depends not only on its action, but also those chosen by the other players.

For example, in “Rock-Paper-Scissors (RPS),” there are two players, each with three actions. A natural choice of utility functions is depicted in Figure 1.

For a more complex and relevant example of a game, consider *simultaneous first-price auctions (S1As)*. There are k players. An action a_i of a player i constitutes a bid b_{ij} on

³This question was bothering your instructor back in CS364B (Winter ’14) — hence, Theorem 5.4.

each item j of a set M of m items.⁴ In a S1A, each item is sold separately in parallel using a “first-price auction” — the item is awarded to the highest bidder, and the price is whatever that player bid.⁵ To specify the utility functions, we assume that each player i has a valuation v_i as in Section 2. We define

$$u_i(\mathbf{a}) = \underbrace{v_i(S_i)}_{\text{value of items won}} - \underbrace{\sum_{j \in S_i} b_{ij}}_{\text{price paid for them}},$$

where S_i denotes the items on which i is the highest bidder (given the bids of \mathbf{a}).⁶ Note that the utility of a bidder depends both on its own action and those of the other bidders. Having specified the players, their actions, and their utility functions, we see that an S1A is an example of a game.

5.1.2 Equilibria

Given a game, how should one reason about it? The standard approach is to define some notion of “equilibrium” and then study the equilibrium outcomes. There are many useful notions of equilibria (see e.g. the instructor’s CS364A notes); for simplicity, we’ll stick here with the most common notion, (mixed) Nash equilibria.⁷

A *mixed strategy* for a player i is a probability distribution over its actions — for example, the uniform distribution over Rock/Paper/Scissors. A *Nash equilibrium* is a collection $\sigma_1, \dots, \sigma_k$ of mixed strategies, one per player, so that each player is performing a “best response” to the others. To explain, adopt the perspective of player i . We think of i as knowing the mixed strategies σ_{-i} used by the other $k - 1$ players (but not their coin flips). Thus, player i can compute the expected payoff of each action $a_i \in A_i$, where the expectation assumes that the other $k - 1$ players randomly and independently select actions from their mixed strategies. Every action that maximizes i ’s expected utility is a *best response* to σ_{-i} . Similarly, every probability distribution over best responses is again a best response (and these exhaust the best responses). For example, in Rock-Paper-Scissors, both players playing the uniform distribution yields a Nash equilibrium. (Every action of a player has expected utility 0 w.r.t. the mixed strategy of the other player, so everything is a best response.)

Nash proved the following.

Theorem 5.1 ([13]) *In every finite game, there is at least one Nash equilibrium.*

⁴To keep the game finite, let’s agree that each bid has to be an integer between 0 and some known upper bound B .

⁵You may have also heard of the *Vickrey* or *second-price* auction, where the winner does not pay their own bid, but rather the highest bid by someone else (the second-highest overall). We’ll stick with S1As for simplicity, but similar results are known for simultaneous second-price auctions, as well.

⁶Break ties in an arbitrary but consistent way.

⁷For the auction settings we study, “Bayes-Nash equilibria” are more relevant. These generalize Nash equilibria, so our lower bounds immediately apply to them.

Theorem 5.1 can be derived from, and is essentially equivalent to, Brouwer’s Fixed-Point Theorem. Note that a game can have a large number of Nash equilibria— if you’re trying to meet a friend in New York City, with actions equal to intersections, then every intersection corresponds to a Nash equilibrium.

An ϵ -Nash equilibrium is the relaxation of a Nash equilibrium in which no player can increase its expected utility by more than ϵ by switching to a different strategy. Note that the set of ϵ -Nash equilibria is nondecreasing with ϵ . Such approximate Nash equilibria seem crucial to the lower bound in Theorem 5.4, below.

5.1.3 The Price of Anarchy

So how good are the equilibria of various games, such as S1As? To answer this question, we use an analog of the approximation ratio, adapted for equilibria. Given a game (like an S1A) and a nonnegative maximization objective function on the outcomes (like welfare), the *price of anarchy (POA)* [11] is defined as the ratio between the objective function value of an optimal solution, and that of the worst equilibrium.⁸ If the equilibrium involves randomization, as with mixed strategies, then we consider its expected objective function value.

The POA of a game and a maximization objective function is always at least 1. It is common to identify “good performance” of a system with strategic participants as having a POA close to 1.⁹

For example, the equilibria of S1As are surprisingly good in fairly general settings.

Theorem 5.2 ([7]) *In every S1A with subadditive bidder valuations, the POA is at most 2.*

Theorem 5.2 is non-trivial and we won’t prove it here (see the paper or the instructor’s CS364B notes for a proof). This result is particularly impressive because achieving an approximation factor of 2 for the welfare-maximization problem with subadditive bidder valuations by any means (other than brute-force search) is not easy (see [5]).

A recent result shows that the analysis of [7] is tight.

Theorem 5.3 ([3]) *The worst-case POA of S1As with subadditive bidder valuations is at least 2.*

The proof of Theorem 5.3 is an ingenious explicit construction — the authors exhibit a choice of subadditive bidder valuations and a Nash equilibrium of the corresponding S1A so that the welfare of this equilibrium is only half of the maximum possible. One reason that proving results like Theorem 5.3 is challenging is that it can be difficult to solve for a (bad) equilibrium of a complex game like a S1A.

⁸Recall that games generally have multiple equilibria. Ideally, we’d like an approximation guarantee that applies to *all* equilibria — this is the point of the POA.

⁹An important issue, outside the scope of these notes, is the plausibility of a system reaching an equilibrium. A natural solution is to relax the notion of equilibrium enough so that it become “relatively easy” to reach an equilibrium. See e.g. the instructor’s CS364A notes for much more on this point.

5.2 POA Lower Bounds from Communication Complexity

Theorem 5.2 motivates an obvious question: can we do better? Theorem 5.3 implies that the analysis in [7] cannot be improved, but can we reduce the POA by considering a different auction? Ideally, the auction would still be “reasonably simple” in some sense. Alternatively, perhaps no “simple” auction could be better than S1As? If this is the case, it’s not clear how to prove it directly — proving lower bounds via explicit constructions auction-by-auction does not seem feasible.

Perhaps it’s a clue that the POA upper bound of 2 for S1As (Theorem 5.2) gets stuck at the same threshold for which there is a lower bound for protocols that use polynomial communication (Theorem 4.3). It’s not clear, however, that a lower bound for low-communication protocols has anything to do with equilibria. In the spirit of the other reductions that we’ve seen in this course, can we extract a low-communication protocol from an equilibrium?

Theorem 5.4 ([16]) *Fix a class \mathcal{V} of possible bidder valuations. Suppose there exists no nondeterministic protocol with subexponential (in m) communication for the 1-inputs of the following promise version of the welfare-maximization problem with bidder valuations in \mathcal{V} :*

- (1) *Every allocation has welfare at most W^*/α .*
- (0) *There exists an allocation with welfare at least W^* .*

Let ϵ be bounded below by some inverse polynomial function of n and m . Then, for every auction with sub-doubly-exponential (in m) actions per player, the worst-case POA of ϵ -Nash equilibria with bidder valuations in \mathcal{V} is at least α .

Theorem 5.4 says that lower bounds for nondeterministic protocols carry over to all “sufficiently simple” auctions, where “simplicity” is measured by the number of actions available to each player. These POA lower bounds follow from communication complexity lower bounds, and do not require any new explicit constructions.

To get a feel for the simplicity constraint, note that S1As with integral bids between 0 and B have $(B + 1)^m$ actions per player — singly exponential in m . On the other hand, in a “direct-revelation” auction, where each bidder is allowed to submit a bid on each bundle $S \subseteq M$ of items, each player has a doubly-exponential (in m) number of actions.¹⁰

The POA lower bound promised by Theorem 5.4 is only for ϵ -Nash equilibrium; since the POA is a worst-case measure and the set of ϵ -Nash equilibria is nondecreasing with ϵ , this is weaker than a lower bound for exact Nash equilibria. It is an open question whether or not Theorem 5.4 holds also for the POA of exact Nash equilibria. Arguably, Theorem 5.4 is good enough for all practical purposes — a POA upper bound that holds for exact Nash equilibria and does not hold (at least approximately) for ϵ -Nash equilibria with very small ϵ is too brittle to be meaningful.

Theorem 5.4 has a number of interesting corollaries. First, since S1As have only a singly-exponential (in m) number of actions per player, Theorem 5.4 applies to them. Thus,

¹⁰Equilibria can achieve the optimal welfare in direct-revelation mechanisms, so the bound in Theorem 5.4 on the number of actions is necessary. See the Exercises for further details.

combining it with Theorem 4.3 recovers the POA lower bound of Theorem 5.3 — modulo the exact vs. approximate Nash equilibria issue — and shows the optimality of the upper bound in Theorem 5.2 without an explicit construction. More interestingly, this POA lower bound of 2 (for subadditive bidder valuations) applies not only to S1As, but more generally to all auctions in which each player has a sub-doubly-exponential number of actions. Thus, S1As are in fact *optimal* among the class of all such auctions when bidders have subadditive valuations (w.r.t. the worst-case POA of ϵ -Nash equilibria).

We can also combine Theorem 5.4 with Theorem 4.1 to prove that no “simple” auction gives a non-trivial (better than k -) approximation for general bidder valuation. Thus with general valuations, complexity is essential to any auction format that offers good equilibrium guarantees.

5.3 Proof of Theorem 5.4

Presumably, the proof of Theorem 5.4 extracts a low-communication protocol from a good POA bound. The hypothesis of Theorem 5.4 offers the clue that we should be looking to construct a nondeterministic protocol. So what could we use an all-powerful prover for? We’ll see that a good role for the prover is to suggest a Nash equilibrium to the players.

Unfortunately, it’s too expensive for the prover to even write down the description of a Nash equilibrium, even in S1As. Recall that a mixed strategy is a distribution over actions, and that each player has an exponential (in m) number of actions available in a S1A. Specifying a Nash equilibrium thus requires an exponential number of probabilities. To circumvent this issue, we resort to ϵ -Nash equilibria, which are guaranteed to exist even if we restrict ourselves to distributions with small descriptions.

Lemma 5.5 ([12]) *For every $\epsilon > 0$ and every game with k players with action sets A_1, \dots, A_k , there exists an ϵ -Nash equilibrium with description length polynomial in k , $\log(\max_{i=1}^k |A_i|)$, and $\frac{1}{\epsilon}$.*

We give the high-level idea of the proof of Lemma 5.5; see the Exercises for details.

1. Let $(\sigma_1, \dots, \sigma_k)$ be a Nash equilibrium. (One exists, by Nash’s Theorem.)
2. Run T independent trials of the following experiment: draw actions $a_1^t \sim \sigma_1, \dots, a_k^t \sim \sigma_k$ for the k players independently, according to their mixed strategies in the Nash equilibrium.
3. For each i , define $\hat{\sigma}_i$ as the empirical distribution of the a_i^t ’s. (With the probability of a_i in $\hat{\sigma}_i$ equal to the fraction of trials in which i played a_i .)
4. Use Chernoff bounds to prove that, if T is at least a sufficiently large polynomial in k , $\log(\max_{i=1}^k |A_i|)$, and $\frac{1}{\epsilon}$, then with high probability $(\hat{\sigma}_1, \dots, \hat{\sigma}_k)$ is an ϵ -Nash equilibrium. Note that the natural description length of $(\hat{\sigma}_1, \dots, \hat{\sigma}_k)$ — for example, just by listing all of the sampled actions — is polynomial in n , $\log(\max_{i=1}^k |A_i|)$, and $\frac{1}{\epsilon}$.

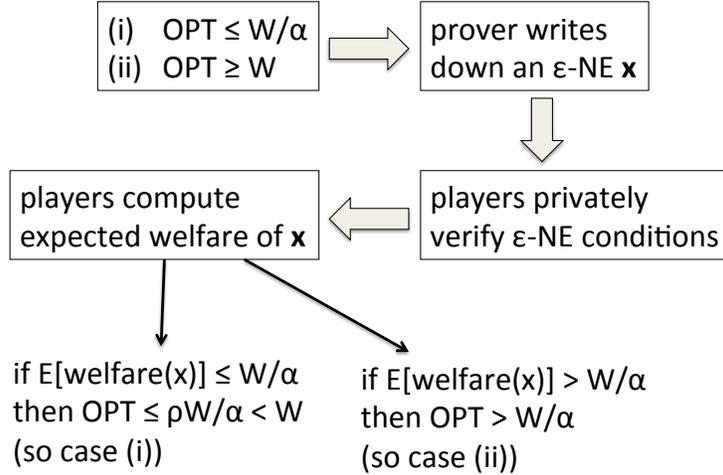


Figure 2: Proof of Theorem 5.4. How to extract a low-communication nondeterministic protocol from a good price-of-anarchy bound.

The intuition is that, for T sufficiently large, expectations with respect to σ_i and with respect to $\hat{\sigma}_i$ should be roughly the same. Since there are $|A_i|$ relevant expectations per player (the expected utility of each of its actions) and Chernoff bounds give deviation probabilities that have an inverse exponential form, we might expect a $\log |A_i|$ dependence to show up in the number of trials.

We now proceed to the proof of Theorem 5.4.

Proof of Theorem 5.4: Fix an auction with at most A actions per player, and a value for $\epsilon = \Omega(1/\text{poly}(k, m))$. Assume that, no matter what the bidder valuations $v_1, \dots, v_k \in \mathcal{V}$ are, the POA of ϵ -Nash equilibria of the auction is at most $\rho < \alpha$. We will show that A must be doubly-exponential in m .

Consider the following nondeterministic protocol for computing a 1-input of the welfare-maximization problem — for convincing the k players that every allocation has welfare at most W^*/α . See also Figure 2. The prover writes on a publicly visible blackboard an ϵ -Nash equilibrium $(\sigma_1, \dots, \sigma_k)$ of the auction, with description length polynomial in k , $\log A$, and $\frac{1}{\epsilon} = O(\text{poly}(k, m))$ as guaranteed by Lemma 5.5. The prover also writes down the expected welfare contribution $\mathbf{E}[v_i(S)]$ of each bidder i in this equilibrium.

Given this advice, each player i verifies that σ_i is indeed an ϵ -best response to the other σ_j 's and that its expected welfare is as claimed when all players play the mixed strategies $\sigma_1, \dots, \sigma_k$. Crucially, player i is fully equipped to perform both of these checks without any communication — it knows its valuation v_i (and hence its utility in each outcome of the game) and the mixed strategies used by all players, and this is all that is needed to verify the ϵ -Nash equilibrium conditions that apply to it and to compute its expected contribution to the welfare.¹¹ Player i accepts if and only if the prover's advice passes these two tests,

¹¹These computations may take a super-polynomial amount of time, but they do not contribute to the

and if the expected welfare of the equilibrium is at most W^*/α .

For the protocol correctness, consider first the case of a 1-input, where every allocation has welfare at most W^*/α . If the prover writes down the description of an arbitrary ϵ -Nash equilibrium and the appropriate expected contributions to the social welfare, then all of the players will accept (the expected welfare is obviously at most W^*/α). We also need to argue that, for the case of a 0-input — where some allocation has welfare at least W^* — there is no proof that causes all of the players to accept. We can assume that the prover writes down an ϵ -Nash equilibrium and its correct expected welfare W , since otherwise at least one player will reject. Since the maximum-possible welfare is at least W^* and (by assumption) the POA of ϵ -Nash equilibria is at most $\rho < \alpha$, the expected welfare of the given ϵ -Nash equilibrium must satisfy $W \geq W^*/\rho > W^*/\alpha$. Since the players will reject such a proof, we conclude that the protocol is correct. Our assumption then implies that the protocol has communication cost exponential in m . Since the cost of the protocol is polynomial in k , m , and $\log A$, A must be doubly exponential in m . ■

Conceptually, the proof of Theorem 5.4 argues that, when the POA of ϵ -Nash equilibria is small, every ϵ -Nash equilibrium provides a privately verifiable proof of a good upper bound on the maximum-possible welfare. When such upper bounds require large communication, the equilibrium description length (and hence the number of actions) must be large.

5.4 An Open Question

While Theorems 4.3, 5.2, and 5.4 pin down the best-possible POA achievable by simple auctions with subadditive bidder valuations, there are still open questions for other valuation classes. For example, a valuation v_i is *submodular* if it satisfies

$$v_i(T \cup \{j\}) - v_i(T) \leq v_i(S \cup \{j\}) - v_i(S)$$

for every $S \subseteq T \subset M$ and $j \notin T$. This is a “diminishing returns” condition for set functions. Every submodular function is also subadditive, so welfare-maximization with the former valuations is only easier than with the latter.

The worst-case POA of S1As is exactly $\frac{e}{e-1} \approx 1.58$ when bidders have submodular valuations. The upper bound was proved in [18], the lower bound in [3]. It is an open question whether or not there is a simple auction with a smaller worst-case POA. The best lower bound known — for nondeterministic protocols and hence, by Theorem 5.4, for the POA of ϵ -Nash equilibria of simple auctions — is $\frac{2e}{2e-1} \approx 1.23$. Intriguingly, there is an upper bound (slightly) better than $\frac{e}{e-1}$ for polynomial-communication protocols [6] — can this better upper bound also be realized as the POA of a simple auction? What is the best-possible approximation guarantee, either for polynomial-communication protocols or for the POA of simple auctions?

protocol’s cost.

References

- [1] N. Alon, Y. Matias, and M. Szegedy. The space complexity of approximating the frequency moments. *Journal of Computer and System Sciences*, 58(1):137–147, 1999.
- [2] A. Chakrabarti, S. Khot, and X. Sun. Near-optimal lower bounds on the multi-party communication complexity of set disjointness. In *Proceedings of the 18th Annual IEEE Conference on Computational Complexity*, pages 107–117, 2003.
- [3] G. Christodoulou, A. Kovacs, A. Sgoutitsa, and B. Tang. Tight bounds for the price of anarchy of simultaneous first price auctions. arXiv:1312.2371, 2013.
- [4] Shahar Dobzinski, Noam Nisan, and Michael Schapira. Approximation algorithms for combinatorial auctions with complement-free bidders. *Math. Oper. Res.*, 35(1):1–13, 2010.
- [5] Uriel Feige. On maximizing welfare where the utility functions are subadditive. *SIAM Journal on Computing*, 39(1):122–142, 2009.
- [6] Uriel Feige and Jan Vondrák. The submodular welfare problem with demand queries. *Theory of Computing*, 6(1):247–290, 2010.
- [7] Michal Feldman, Hu Fu, Nick Gravin, and Brendan Lucier. Simultaneous auctions are (almost) efficient. In *45th ACM Symposium on Theory of Computing (STOC)*, pages 201–210, 2013.
- [8] A. Gronemeier. Asymptotically optimal lower bounds on the NIH-multi-party information complexity of the AND-function and disjointness. In *Proceedings of the 26th Annual Symposium on Theoretical Aspects of Computer Science (STACS)*, pages 505–516, 2009.
- [9] J. Håstad and A. Wigderson. The randomized communication complexity of set disjointness. *Theory of Computing*, 3(11):211–219, 2007.
- [10] B. Kalyanasundaram and G. Schnitger. The probabilistic communication complexity of set intersection. *SIAM Journal on Discrete Mathematics*, 5(4):545–557, 1992.
- [11] E. Koutsoupias and C. H. Papadimitriou. Worst-case equilibria. In *Proceedings of the 16th Annual Symposium on Theoretical Aspects of Computer Science (STACS)*, volume 1563 of *Lecture Notes in Computer Science*, pages 404–413, 1999.
- [12] R. J. Lipton, E. Markakis, and A. Mehta. Playing large games using simple strategies. In *Proceedings of the 4th ACM Conference on Electronic Commerce (EC)*, pages 36–41, 2003.
- [13] J. F. Nash. Equilibrium points in N -person games. *Proceedings of the National Academy of Science*, 36(1):48–49, 1950.

- [14] N. Nisan. The communication complexity of approximate set packing. In *Proceedings of the 29th Annual International Colloquium on Automata, Languages, and Programming (ICALP)*, volume 2380 of *Lecture Notes in Computer Science*, pages 868–875, 2002.
- [15] A. A. Razborov. On the distributional complexity of disjointness. *Theoretical Computer Science*, 106(2):385–390, 1992.
- [16] T. Roughgarden. Barriers to near-optimal equilibria. In *Proceedings of the 55th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 71–80, 2014.
- [17] Y. Shoham and K. Leyton-Brown. *Multiagent Systems*. Cambridge, 2010.
- [18] V. Syrgkanis and É. Tardos. Composable and efficient mechanisms. In *Proceedings of the 45th ACM Symposium on Theory of Computing (STOC)*, pages 211–220, 2013.