2016 IEEE 57th Annual Symposium on Foundations of Computer Science

Computational Efficiency Requires Simple Taxation

Shahar Dobzinski
Computer Science and Applied Mathematics Department
Weizmann Institute of Science
Rehovot, Israel
Email: shahar.dobzinski@weizmann.ac.il

***Abstract*—We characterize the communication complexity of truthful mechanisms. Our departure point is the well known taxation principle. The taxation principle asserts that every truthful mechanism can be interpreted as follows: every player is presented with a menu that consists of a price for each bundle (the prices depend only on the valuations of the other players). Each player is allocated a bundle that maximizes his profit according to this menu. We define the *taxation complexity* of a truthful mechanism to be the logarithm of the maximum number of menus that may be presented to a player.**

Our main finding is that in general the taxation complexity essentially equals the communication complexity. The proof consists of two main steps. First, we prove that for rich enough domains the taxation complexity is at most the communication complexity. We then show that the taxation complexity is much smaller than the communication complexity only in "pathological" cases and provide a formal description of these extreme cases.

Next, we study mechanisms that access the valuations via value queries only. In this setting we establish that the menu complexity – a notion that was already studied in several different contexts – characterizes the number of value queries that the mechanism makes in exactly the same way that the taxation complexity characterizes the communication complexity.

Our approach yields several applications, including strengthening the solution concept with low communication overhead, fast computation of prices, and hardness of approximation by computationally efficient truthful mechanisms.

***Keywords*-Mechanism Design; Communication Complexity**

I. INTRODUCTION

The field of Communication Complexity studies settings in which n players are interested in computing some known function f. Each player i, holds some input x_i. The basic task is to determine the maximum number of bits that the parties need to exchange in order to compute $f(x_1, \dots, x_n)$.

One of the most successful applications of communication complexity is Algorithmic Mechanism Design, starting with the pioneering work of Nisan and Segal [26], [28]. Nisan and Segal proved lower bounds on the approximation ratios achievable by algorithms with low communication complexity for combinatorial auctions. Many other applications of communication complexity to mechanism design have been introduced since. For example, communication complexity is used to bound the power of certain computationally-efficient truthful mechanisms [4], [8], to understand the overhead of price computation [1], [14], and to prove bounds on the quality of equilibria [29].

Our goal in this paper is to answer a fundamental question in the intersection of communication complexity and algorithmic mechanism design: given a truthful mechanism A, how many bits do the parties need to exchange in order to determine the allocation and payments?

A bit more formally, the communication complexity of a protocol is the maximum number of bits that are exchanged in the protocol, where the maximum is taken over all inputs. The communication complexity of a function f, denoted $cc(f)$, is the communication complexity of the protocol that computes f with the smallest communication complexity.

A truthful mechanism A is composed of a social choice function that selects one alternative from a set $\mathcal{S}$ of alternatives and a payment function that specifies the payment of each player. Our results build on a basic concept of Mechanism Design, the *taxation principle*. Denote by $v_i(S)$ the value of player i for alternative $S \in \mathcal{S}$. The taxation principle asserts that A can be interpreted in the following simple form: every player i is (implicitly) presented with a menu $\mathcal{M}_{v_{-i}}$ that is a function that assigns a price (possibly ∞) for each alternative in $\mathcal{S}$. $\mathcal{M}_{v_{-i}}$ depends only on the valuations v_{-i} of the other players. The truthful mechanism A always outputs an alternative S that simultaneously maximizes the profit $v_i(S) - \mathcal{M}_{v_{-i}}(S)$ of each player i.

With this interpretation in mind, given a truthful mechanism A, for each player i denote by $M^i = \{\mathcal{M}_{v_{-i}}\}_{v_{-i}}$ the set of menus that might be presented to i. Denote by $tax(A)$ the *taxation complexity* of a truthful mechanism A – the number of bits needed to represent an index of a specific menu among the set of menus that may be presented to a player. That is, $tax(A) = \max_i \log |M^i|$.

Our main finding directly connects the semantics of the mechanism and its communication complexity by showing that the taxation complexity of every truthful mechanism essentially equals its communication complexity:

Informal take-home message of this paper: In "rich enough" domains, $tax(A) \approx cc(A)$.

Incumbent of the Lilian and George Lyttle Career Development Chair. Supported in part by the I-CORE program of the planning and budgeting committee and the Israel Science Foundation 4/11 and by EU CIG grant 618128.

0272-5428/16 $31.00 © 2016 IEEE
DOI 10.1109/FOCS.2016.30

We also apply the lens of the taxation principle in a more restricted model and prove an analogous result: if access to the valuations is restricted to value queries, then the menu complexity essentially equals the query complexity (see definitions below).

In the rest of the introduction we provide a more formal description of the setup[1], of the results, and of various implications.

The Setting

For concreteness, this paper considers only the setting of combinatorial auctions, although it should be possible to extend the results beyond that domain. In a combinatorial auction there is a set M ($|M| = m$) of heterogeneous items and a set N ($|N| = n$) of players. The output is an allocation $(S_1, \ldots, S_n)$ of the items to the players. The private information of each player i is his value for the set of items he receives: $v_i : 2^M \to \mathbb{R}$. As common in the literature, in this paper we assume that each v_i is normalized ($v_i(\emptyset) = 0$) and monotone (for each $S \subseteq T$, $v_i(T) \geq v_i(S)$).

This paper considers truthful mechanisms. The Algorithmic Mechanism Design literature usually defines truthful mechanisms to be those that implement some social choice function in a dominant-strategy equilibrium. In this paper a mechanism is *truthful* if it implements the social choice function in an ex-post Nash equilibrium. This solution concept applies to games with *incomplete* information and is closely related but less restrictive than dominant-strategy equilibrium. It is more appropriate for the iterative mechanisms that this paper considers. Very roughly speaking, in an ex-post Nash equilibrium a dominant strategy of every player is to play according to his true valuation, as long as the other players are not playing "crazy" strategies. We refer the reader to the full version for formal definitions and discussion.

The Taxation Principle in Algorithmic Mechanism Design

The taxation principle [17, 16] was already considered in Algorithmic Mechanism Design. Most notably, the crux of the impossibility results of [5, 13, 11] is showing that for every mechanism that approximately maximizes the welfare there must be an instance in which one player is presented a "complicated" menu. In particular, finding a profit-maximizing bundle in that menu is hard.

The paper [18] considers a setting with only a single player whose valuation is drawn from some known distribution. Since there is only one player, the taxation principle implies that all a truthful mechanism can do is to present a fixed menu to the player. The player then "selects" a profit-maximizing bundle. They show that to approximately maximize revenue the prices of many alternatives (equivalently, bundles) in that menu must be finite (high "menu complexity").

[1]See Section the full version for formal definitions.

We stress that the notion of taxation complexity does not measure the difficulty of finding a profit-maximizing bundle, nor how difficult it is to represent a specific menu. The taxation complexity takes a more "high-level" view of the mechanism and only measures the *number* of menus that might be presented to a player.

The Taxation Complexity is at most the Communication Complexity

Our first main result says that in rich enough domains the taxation complexity is at most the communication complexity:

$$tax(A) \leq cc(A)$$

We will shortly provide a more formal statement, but to better understand this result and its implications, we first discuss a domain for which this result *does not* hold. Consider a two-player combinatorial auction where the valuations of the players belong to the class of gross substitutes (GS) valuations[2]. Let us furthermore restrict the values of all bundles to be integers in $\{1, \ldots, m\}$. We study the VCG mechanism in this setting. Since the welfare maximizing allocation can be found with $poly(m)$ communication [28] (this is also implied by the algorithms mentioned in [24]), the communication complexity of the VCG mechanism is $poly(m)$ as well.

Let us now analyze the taxation complexity of the mechanism. Denote the valuation of player 1 by v. Player 1 presents a menu to player 2. By the definition of the VCG mechanism, the price of bundle S in that menu is $v(M) - v(M - S)$. Thus, there is a one-to-one and onto correspondence between the set of possible valuations of player 1 and the set of menus he presents to player 2. All that is left is to point out that the number of gross substitutes valuations is[3] doubly exponential (e.g, Knuth [19] shows that it is at least $2^{\frac{\binom{m}{m/2}}{2m}}/m!$). Therefore, for two players with gross substitutes valuations, the taxation complexity of the VCG mechanism $tax(VCG)$ is therefore exponential, whereas the communication complexity $cc(VCG)$ is polynomial.

In contrast, in richer domains the taxation complexity is not much larger than the communication complexity:

Theorem: Fix some mechanism A.

1) If A is truthful for general valuations, then $tax(A) \leq cc(A)$.
2) If A is truthful for subadditive valuations, then also $tax(A) \leq cc(A)$.
3) If A is truthful for XOS valuations, then $tax(A) \leq m \cdot (cc(A) + 1)$.

[2]The definition of this class is subtle; Since it will not be needed in this paper, we refer the interested reader to the survey [24] for a definition.

[3]In fact, Knuth [19] shows that the number of matroid rank functions on $\{1, \ldots, m\}$ is doubly exponential, and it is known that every matroid rank function is in particular gross substitutes.

4) If A is truthful for submodular valuations, then $tax(A) \leq d \cdot m \cdot (cc(A)+1)$, where $d = |\{\mathcal{M}(S)\}_{\mathcal{M},S}|$ is the total number of distinct prices that appear in some menu.

The cautious reader might wonder how it can be that the class of GS valuations is contained in all the above-mentioned classes, but in these classes the communication complexity severely limits the taxation complexity. The point is that implementations of mechanisms that are specifically tailored to GS valuations are able to find profit-maximizing bundles without learning the full menu. However, when running those implementations in richer domains, the set of possible deviations increases and truthfulness is lost.

Characterizing the Communication Complexity of Truthful Mechanisms

We have that in rich enough domains $tax(A) \leq poly(cc(A))$. Had we were able to prove that $cc(A) \leq poly(tax(A))$, this would immediately imply that the communication complexity of a truthful mechanism is completely determined (up to a polynomial factor) by $tax(A)$ – a well defined combinatorial property that depends only on the social choice function.

However, we show that $cc(A)$ cannot be bounded by $poly(tax(A))$. Towards this end, consider the following (naive and incorrect) implementation of a two-player mechanism A which is truthful for general valuations: since the menu that is presented to a player depends only on the other player's valuation, each player can send $tax(A)$ bits that denote the index of the menu he presents to the other player. The obvious next step is to ask each player to select the profit maximizing bundle from the menu that was presented to him, announce it (using additional m bits) and allocate accordingly. The total communication cost of this implementation is $2(tax(A) + m)$ as we wanted.

The above implementation is incorrect since this last step is not well defined because of tie-breaking: there might be several bundles that simultaneously maximize the profit. The tie-breaking rule that defines which profit-maximizing bundle each player receives can be in principle quite involved, and there is no way to avoid that: we show (in the full version) a two-player mechanism with taxation complexity 1 and communication complexity $exp(m)$, due to tie breaking.

This leads us to the following definition. Let $tie(A)$ be the communication complexity of determining the allocation in a truthful mechanism A, where the input of a player is a valuation v_i and in addition all players know the menu that is presented to each player. Notice that obviously $tie(A) \leq cc(A)$ since we can always ignore this extra information and simply run A to determine the allocation. By our discussion above and our bound on the taxation complexity we get that:

Theorem: Let A be a two-player truthful mechanism for rich enough domain. Then,

$$\frac{tax(A) + tie(A)}{2} \leq cc(A) \leq 2 \cdot tax(A) + tie(A)$$

In particular, whenever the communication complexity of the tie-breaking rule is low, we indeed get that the communication complexity almost equals the taxation complexity.

Can we extend this theorem to more than two players? The missing component for three players or more is that it is not clear whether it is possible to *explicitly* find the menu that $n-1$ players present to the remaining player with low communication (the taxation principle only guarantees the *existence* of such menu, but gives no guidance on how to find it). We provide a positive answer:

The Menu Reconstruction Theorem: Let A be an n-player truthful mechanism. Fix some player i. Then for every valuation profile v_{-i}, there is a protocol with communication complexity $poly(tax(A), price(A), m, n)$ that finds the menu that is presented to player i.

Where we denote by $price(A)$ the maximum number of bits that it takes for any $n-1$ players to find the price of a given bundle S in the menu that they present to the remaining player. Therefore:

Theorem: Let A be an n-player truthful mechanism in any domain. Then,

$$cc(A) \leq poly(tax(A), price(A), tie(A), m, n)$$

We also show that as long as the domain includes additive valuations, $price(A) \leq cc(A)$. We thus get that for rich enough domains $\frac{tax(A)+tie(A)+price(A)}{3} \leq cc(A)$, which allows us to completely determine the communication complexity of truthful mechanisms (up to a polynomial factor):

$$\frac{tax(A) + price(A) + tie(A)}{3} \leq cc(A)$$
$$\leq poly(tax(A), price(A), tie(A), m, n)$$

Our characterization is tight in the sense that if we drop one of the three main terms $(tax(A), price(A), tie(A))$ then the gap between the LHS and the RHS might be exponential. For instance, we have already mentioned an example of a truthful mechanism A with $tax(A) = 1$ (and thus $price(A) = 0$) in which $cc(A) = exp(m)$. We also provide other "pathological" examples with similar gaps when dropping either $price(A)$ or $tax(A)$.

Characterizing the Query Complexity

Up until now we imposed no restrictions on the communication between the parties. However, many of the truthful mechanisms in the literature assume that the valuations are represented as black boxes that answer only a specific type of queries. A simple type of query that was extensively studied is a *value query*: given a bundle S, what is $v(S)$? We now characterize the number of value queries that a truthful mechanism makes, again by applying the taxation principle.

Denote the query complexity of a truthful mechanism by $val(A)$ – this is the number of value queries that the most efficient implementation of A makes. Following [18], The *menu complexity* of A, denoted $mc(A)$, is roughly speaking the maximum number of bundles with finite price that appear in any menu that is used in[4] A. We establish that for mechanisms that use only value queries, the menu complexity characterizes the query complexity in exactly the same way that the taxation complexity characterizes the communication complexity:

Theorem: Let A be a mechanism that is truthful for general valuations and accesses the valuations via value queries only. Then:

$$\frac{mc(A) + price^{val}(A) + tie^{val}(A)}{3} \leq val(A)$$
$$\leq poly(mc(A), price^{val}(A), tie^{val}(A), m, n)$$

where $price^{val}(A)$ and $tie^{val}(A)$ are defined similarly to $price(A)$ and $tie(A)$ with the additional restriction that the communication is restricted to value queries. We note that the inequality $mc(A) \leq tax(A)$ is in fact implicit in [5] and [27, Theorem 11], whereas the right inequality (a menu reconstruction theorem that uses only value queries) is new and very different from the menu reconstruction theorem for unrestricted communication.

To strengthen the analogy between $tax(A)$ and $mc(A)$, consider the following two player menu optimization problem: Alice's input is some menu $\mathcal{M} \in U$, where the set of menus U is known in advance. Bob's input is some valuation v. The goal is to find a bundle that maximizes the profit $v(S) - \mathcal{M}(S)$. We restrict ourselves to one way protocols: Alice speaks first and then Bob. After Bob speaks, both parties know a profit maximizing bundle S.

We observe that if we let U be the set of menus presented to some player in a truthful mechanism A, the (one way) communication complexity of the menu optimization problem is $tax(A)$ (up to an additive factor of m bits). Interestingly, our results yield that when Bob is restricted to value queries (Alice sends an arbitrary message, then, based on this message, the center queries for the value of some bundles in v) then the communication complexity is essentially $mc(A)$. That is, both $tax(A)$ and $mc(A)$ capture the informational bottleneck of finding a profit-maximizing bundle in a non-interactive way in their respective models.

The other type of popular query is *demand query*: given prices $p_1, \ldots, p_m$ return a bundle $S \in \arg\max_T v(T) - \Sigma_{j \in T} p_j$. We identify the *affinity* of a mechanism as the communication complexity of the menu optimization problem when the communication is restricted to demand queries. Specifically, a menu $\mathcal{M}$ is *α-min affine* if there are α price vectors $p^1, \ldots, p^\alpha$ and α non-negative numbers $r^1, \ldots, r^\alpha$ such that for all S, $\mathcal{M}(S) = \min_{1 \leq k \leq \alpha} \Sigma_{j \in S}(p_j^k) + r^k$. The affinity of A, denoted $aff(A)$, is the minimal number α such that all menus presented by A are α-min affine.

Denote by $dem(A)$ the number of demand queries that the most efficient implementation of a truthful mechanism A makes. Just as $tax(A) \leq cc(A)$ and $mc(A) \leq val(A)$, we show that for general valuations $aff(A) \leq dem(A)$. However, unlike $tax(A)$ and $mc(A)$, $aff(A)$ cannot be used to characterize $dem(A)$. This is one particular consequence of an *impossibility* result:

Theorem (no menu reconstruction theorem for demand queries): There is a mechanism A that is truthful for some player i with a general valuation in which $dem(A) = poly(m)$ and $price^{dem}(A) = 1$, but $exp(m)$ demand queries are needed to find the menu presented to player i.

In other words, just as q value queries suffice to find a profit maximizing bundle when the menu complexity is q, when the affinity is q a profit maximizing bundle can be found with q demand queries. However, by our impossibility result it is not easy to figure out *which* q queries to make. Nevertheless, the affinity $aff(A)$ will be useful in some of the applications that we mention below.

The query complexity of mechanisms is studied in the full version.

Implications and Extensions

We view the study of these complexity measures as an investigation of the fundamentals of Algorithmic Mechanism Design that needs no further justification. Nevertheless, as is often the case, studying the foundations yields several interesting implications. We elaborate on some now, as well as on some open questions. Other open questions are stated in the technical parts of the paper.

From ex-post Nash to Dominant Strategies (Section IV): The revelation principle implies that any mechanism that implements a social choice function in an ex-post Nash equilibrium can be transformed to a mechanism that implements the same social choice function in a dominant-strategy equilibrium. However, the communication blow-up might be exponential, as in combinatorial auctions with GS valuations. Our theorems allow us to obtain a new transformation that takes any two-player mechanism that implements an ex-post Nash equilibrium in a rich enough domain to a mechanism that implements the same social choice function in dominant-strategy equilibrium with only a polynomial blow up in the communication complexity.

The Limits of Computationally-Efficient Truthful Mechanisms (full version): A major research direction in Algorithmic Mechanism Design studies the power of computationally efficient truthful mechanisms for welfare maximiza-

[4]The description of the menu complexity given here is inaccurate as it ignores tie breaking issues. We refer the reader to the full version of this paper for the precise definition.

tion in combinatorial auctions (e.g., [23, 21, 25, 22, 3, 10]). For example, VCG is a truthful mechanism that maximizes the welfare, but in combinatorial auctions with submodular valuations it requires exponential communication [28]. On the other hand, there is a 1.58-approximation algorithm that uses only polynomial communication [15], but it is not truthful. The best deterministic truthful mechanism that uses only polynomial communication achieves a poor approximation ratio of $O(\sqrt{m})$ [10]. Whether this is the best possible for a deterministic truthful mechanism that uses only polynomial communication is a major open question. This state of affairs is typical to other problems as well, and we basically completely lack tools for proving impossibility results for computationally efficient truthful mechanisms[5]. Our work gives rise to two different novel approaches for proving such impossibilities.

Approach I: simultaneous algorithms. The first approach for proving impossibilities is by a reduction to *simultaneous* algorithms. This model was introduced in [9]: each of the n players in a combinatorial auction simultaneously sends c bits that are a function of his valuation only. The center then determines the allocation using only those messages. We show that *impossibilities for two-player simultaneous algorithms imply impossibilities for computationally efficient truthful mechanisms*. Specifically, we show that if there is a truthful mechanism (for the domains discussed above) that provides an approximation ratio of α with communication complexity $cc(A)$, then there is an α-approximation *simultaneous* algorithm where the length of the simultaneous messages is $poly(cc(A))$. For example, a proof that no simultaneous algorithm with polynomially long messages for submodular players achieves a 1.59 approximation (ignoring incentives issues) immediately establishes the first ever gap between computationally efficient truthful mechanisms and their non-truthful counterparts. Note that strong impossibility results for simultaneous algorithms are known [9], but those unfortunately hold only for a large number of players. In particular, nothing is known for two players.

Approach II: lower bounds on $tax(A)$. The second approach involves handling the taxation complexity directly. The idea is simple: suppose one can prove that the taxation complexity of every truthful α-approximation mechanism for combinatorial auctions with submodular players is exponential (ignoring computational issues). Since $tax(a) \leq poly(cc(A))$, it follows that the communication complexity of every α-approximation truthful mechanism is exponential as well, which establishes a gap between the power of truthful and non-truthful computationally efficient algorithms.

[5]The papers [5], [11], [13] prove impossibility results when access is restricted to value queries, or when the valuations are given in a succinct and very specific form. For the general communication model, or even when access is restricted to demand queries, that are no impossibility results on the power of computationally efficient truthful mechanisms.

This approach can be extended to mechanisms with restricted access. For example, to prove impossibility results for mechanisms that use only demand queries (e.g., [6]) it suffices to show that the affinity of every α-approximation truthful mechanism for combinatorial auctions with submodular players (ignoring computational issues) is exponential. We note that the menu complexity of $m^{\frac{1}{2}-\epsilon}$-approximation mechanisms for submodular valuations was already proved to be exponential, which indeed yielded an impossibility on the power of truthful mechanisms that use value queries in the aforementioned setting (the direct hardness approach of [5]).

Extensions to randomized mechanisms. Before this paper, the only viable approach for proving impossibility results for randomized mechanisms in the general communication model was by characterizing all truthful mechanisms with a good approximation. Such characterizations are notoriously hard even for deterministic mechanisms. For randomized ones, it is probably fair to describe the possibility of obtaining such characterizations in the foreseeable future as almost hopeless.

There are two main notions of randomized truthfulness. The first is truthfulness in expectation, where each player maximizes his *expected* profit. We discuss this notion below, and here we focus on the other (stronger) notion: universal truthfulness. Universally truthful mechanisms are simply a probability distribution over deterministic mechanisms. Interestingly, universally truthful mechanisms achieve the best currently known approximation ratios in some important settings (e.g., combinatorial auctions with submodular players [6]), even if truthful in expectation mechanisms are considered.

Both approaches are capable of proving impossibility results for universally truthful mechanisms. First, an impossibility for two-player randomized simultaneous algorithms implies an impossibility for randomized truthful mechanisms. The second approach is also applicable: a lower bound on the taxation complexity is likely to be proved by obtaining a distribution over the input on which no mechanism with polynomial taxation complexity provides a good approximation ratio. Yao's principle and our results imply that for every universally truthful mechanism with polynomial communication there is an instance on which its (expected) approximation ratio is bad.

Efficient Price Computation: Every truthful mechanism has two tasks: the first is determining the outcome of the social choice function and the second is computing the players' payments. Fadel and Segal [14] ask whether the additional communication cost of computing the payments is significantly larger than the communication complexity of computing the social choice function. In deterministic settings, they show that the bound is at most exponential and ask whether this is tight. Single parameter domains are handled by [1] and various multi-parameter domains are handled by [2], [30] via a "single call' approach, but at the cost of introducing randomization.

We extend this line of work. Since we showed that in rich enough domains $tax(A), price(A) \leq cc(A)$, our menu reconstruction theorem immediately implies that fully presenting the players with the actual menu takes only $poly(cc(A))$ bits, thus finding the menu is essentially as easy as computing the allocation and the payments of winning bundles ($cc(A)$ bits).

Truthful in Expectation Mechanisms (full version): Our theorems were proved for deterministic mechanisms (and thus they also apply to randomized universally truthful mechanisms). Since truthful in expectation mechanisms were extensively studied (e.g., [7], [12], [22]), it is natural to ask whether in truthful in expectation mechanisms $tax(A)$ similarly characterizes $cc(A)$. We provide a negative answer: for combinatorial auctions with general valuations there is a truthful in expectation mechanism with polynomial communication and exponential taxation complexity.

Recall that another setting where the taxation complexity might be exponential in the communication complexity is when the valuations are gross substitutes. One can further make the following wild speculation, which lacks more evidence and formalization: in domains where the gap between the taxation complexity and the communication complexity is small, the performance of computationally efficient truthful mechanisms is poor, whereas for domains where it is large, truthfulness is not a severely limiting requirement.

II. Bounding the Taxation Complexity: $tax(A) \leq poly(cc(A))$

We would now like to bound the taxation complexity as a function of the communication complexity. As a warm up, in Subsection II-A we prove that the taxation complexity of a truthful mechanism for general valuations with communication complexity $cc(A)$ is $cc(A) + 1$, as long as all bundles get a finite price in every menu (for example, this is true for deterministic mechanisms that provide some finite approximation ratio to the welfare).

In Subsection II-B we strengthen this result in several aspects. First, we generalize the result to any truthful mechanism by allowing the prices of bundles to be ∞. Second, we slightly strengthen the bound on the taxation complexity to $cc(A)$ (and not just $cc(A)+1$). Finally, we extend our results for other classes of valuations, not just general valuations.

A. A Warm-up

We start with a warm-up theorem that is inferior to the result of Subsection II-B:

Theorem 2.1: Consider a mechanism A for combinatorial auction that is truthful for general valuations. Suppose that for each player i, bundle S, and menu $\mathcal{M}$ that might be presented to player i we have that $\mathcal{M}(S) < \infty$. Then, $tax(A) \leq cc(A) + 1$.

Proof: Fix some player i. Let the *taxation complexity of player* i be the logarithm of the number of menus that player i might be presented with in A. We will prove that the taxation complexity of player i is at most $cc(A) + 1$ and the theorem will follow.

Define a protocol A' that is completely identical to the most efficient implementation of A (in particular, the same allocation and payment functions), except that at the end of A' player i sends the bit 1 if his profit is positive (i.e., $A_i(v) - p_i(v) > 0$) and the bit 0 otherwise. Note that A' is truthful since A is truthful and that the communication complexity of A' is $cc(A) + 1$. In addition, for each set of valuations v_{-i} of the other players the menu presented to i is identical in both A and A'. Therefore it suffices to prove that the taxation complexity of player i in A' is at most $cc(A')$. This implies that the taxation complexity of A is $cc(A) + 1$.

With this in mind, let $M^i = \{\mathcal{M} | \exists (v_1, \ldots, v_{i-1}, v_{i+1}, \ldots v_n) \text{ s.t. } \mathcal{M} \text{ is presented to } i\}$ be the set of all possible menus that might be presented to player i. For each $\mathcal{M} \in M^i$, let $v_i^{\mathcal{M}}$ be the valuation in which for every S we have that $v_i^{\mathcal{M}}(S) = \mathcal{M}(S)$. Observe that each $v_i^{\mathcal{M}}$ is a valid valuation function: in the full version we show that $\mathcal{M}$ is monotone and normalized, thus $v_i^{\mathcal{M}}$ is monotone and normalized as well. Also notice that for each S, $v_i^{\mathcal{M}}(S) < \infty$ since $\mathcal{M}(S) < \infty$. In addition, for every $\mathcal{M} \in M^i$ choose an arbitrary set of valuations of the other players $v_{-i}^{\mathcal{M}} = (v_1^{\mathcal{M}}, \ldots, v_{i-1}^{\mathcal{M}}, v_{i+1}^{\mathcal{M}}, \ldots, v_n^{\mathcal{M}})$ such that if the players' valuations are $v_{-i}^{\mathcal{M}}$ the menu that player i is presented with is $\mathcal{M}$.

Now we get to the heart of the proof. We will show that for each $\mathcal{M}, \mathcal{M}' \in M^i$, $\mathcal{M} \neq \mathcal{M}'$, the transcript of A' in the instance $(v_i^{\mathcal{M}}, v_{-i}^{\mathcal{M}})$ differs from the transcript of A' in the instance $(v_i^{\mathcal{M}'}, v_{-i}^{\mathcal{M}'})$. Recall that the communication complexity of A' is $cc(A)+1$, thus there are at most $2^{cc(A)+1}$ different transcripts. The bound on the taxation complexity of A' will then follow since every instance of the form $(v_i^{\mathcal{M}}, v_{-i}^{\mathcal{M}})$ corresponds to exactly one menu $\mathcal{M} \in M^i$.

Claim 2.2: For every $\mathcal{M}, \mathcal{M}' \in M^i$, $\mathcal{M} \neq \mathcal{M}'$, the transcript of A' in the instance $(v_i^{\mathcal{M}}, v_{-i}^{\mathcal{M}})$ differs from the transcript of A' in the instance $(v_i^{\mathcal{M}'}, v_{-i}^{\mathcal{M}'})$.

Proof: Assume towards contradiction that there are $\mathcal{M}, \mathcal{M}' \in M^i$, $\mathcal{M} \neq \mathcal{M}'$ such that the transcript of the instance $(v_{-i}^{\mathcal{M}}, \mathcal{M})$ is identical to the transcript of the instance $(v_{-i}^{\mathcal{M}'}, \mathcal{M}')$. Using standard fooling-set arguments (e.g., [20]), this implies that the transcripts of $(v_{-i}^{\mathcal{M}}, \mathcal{M}')$ and $(v_{-i}^{\mathcal{M}'}, \mathcal{M})$ are identical as well. We will show that this is not the case and reach a contradiction. Towards this end, observe that the last bit that player i sends in both instances is by construction 0 (since $v_i^{\mathcal{M}}(S) = \mathcal{M}(S)$ and $v_i^{\mathcal{M}'}(S) = \mathcal{M}'(S)$ for every bundle S, so the profit in both instances is 0). However, we will show that in either $(v_i^{\mathcal{M}}, v_{-i}^{\mathcal{M}'})$ or $(v_i^{\mathcal{M}'}, v_{-i}^{\mathcal{M}})$ the last bit that player i sends is 1. In particular we get a different transcript, which is a contradiction.

To see that in one of the instances $(v_i^{\mathcal{M}}, v_{-i}^{\mathcal{M}'})$ and $(v_i^{\mathcal{M}'}, v_{-i}^{\mathcal{M}})$ the last bit that is communicated is 1, notice that since $\mathcal{M} \neq \mathcal{M}'$, there must be a bundle S such that $\mathcal{M}(S) \neq \mathcal{M}'(S)$. Assume without loss of generality that $\mathcal{M}(S) > \mathcal{M}'(S)$. Thus in the instance $(v_i^{\mathcal{M}'}, v_{-i}^{\mathcal{M}})$ the profit of player 1 for the bundle S is $v_i^{\mathcal{M}}(S) - \mathcal{M}'(S) > 0$. By the taxation principle, player i must win a bundle with at least that (positive) profit. Thus the last bit that player i communicates is 1, which gives us the desired contradiction. ■

This concludes the proof of Theorem 2.1. ■

Tightness: To see that the Theorem 2.1 is essentially tight, we present a mechanism with taxation complexity very close to the communication complexity. Consider the following truthful mechanism for combinatorial auctions with two players, Alice and Bob. Let $M^{Bob} = \{\mathcal{M}_1, \ldots, \mathcal{M}_{2^c}\}$ be a set of 2^c menus, where for each $\mathcal{M}_i \in M^{Bob}$ we have that $\mathcal{M}_i(\{a\}) = i$, $\mathcal{M}_i(\emptyset) = 0$, and $\mathcal{M}_i(S) = \infty$ for every $S \neq \emptyset, \{a\}$. Let t be Alice's value for item a rounded to the nearest integer in $\{1, 2, 3, \ldots, 2^c\}$. Alice now sends t using c bits of communication. If Bob's value for item a is at least t, Bob sends the bit 1, receives a, and pays t. Otherwise, he sends the bit 0, receives no items at all, and pays nothing. Alice always receives the empty bundle. The mechanism is clearly truthful, its communication complexity $c+1$, and its taxation complexity is c.

B. Bounding the Taxation Complexity: The Full Result

We now significantly strengthen the results of Subsection II-A. In particular, we give bounds on the taxation complexity also for mechanisms that are truthful for restricted classes of valuations (subadditive, XOS, and submodular). We will show that:

Theorem 2.3: Let $\mathcal{V}$ be some class of valuations. Fix a mechanism A for combinatorial auctions that is truthful when the valuations of the players are in $\mathcal{V}$. Then:

1) If $\mathcal{V}$ is the set of all normalized and monotone valuations then $tax(A) \leq cc(A)$.
2) If $\mathcal{V}$ is the set of subadditive valuations then $tax(A) \leq cc(A)$.
3) If $\mathcal{V}$ is the set of XOS valuations then $tax(A) \leq m \cdot (cc(A) + 1)$.
4) If $\mathcal{V}$ is the set of submodular valuations then $tax(A) \leq d \cdot m \cdot (cc(A) + 1)$, where $d = |\{\mathcal{M}(S)\}_{\mathcal{M} \in M^i, S}|$ is the total number of distinct prices that appear in some menu.

The proof of Theorem 2.3 is postponed to the full version. Interestingly, we do not know whether the bounds for XOS and submodular valuations are tight, since we have no example for a mechanism A that is truthful for these classes with $tax(A) > cc(A)$.

We note that the vast majority of the algorithms in the literature are truthful for general valuations (e.g., maximal in range algorithms, posted prices mechanisms). The restrictions on the valuations are typically used only for the performance analysis. The taxation complexity of those mechanisms is therefore at most their communication complexity. We also remark that the bound on the taxation complexity for submodular valuations depends on the number of possible prices. A natural open question is:

Open Question 1: Let A be a mechanism for combinatorial auctions that is truthful for submodular valuations. Is $tax(A) \leq poly(cc(A), m)$?

More generally, we have already mentioned two domains in which there is a truthful mechanism A with $tax(A) >> cc(A)$. The first was combinatorial auctions with gross substitute valuations. For combinatorial auctions with general valuations we mentioned a randomized truthful in expectation mechanism, but this mechanism (as well as all truthful in expectation mechanisms) can be thought as a deterministic one by letting the range be the set of all possible distributions over allocations and letting the value of a player for a distribution be the expected value of the bundle he receives in that distribution. This leads us to the following question:

Open Question 2: Characterize the set of domains in which for every truthful mechanism A we have that $tax(A) \leq poly(cc(A), m)$.

III. MENU RECONSTRUCTION: $cc(A) \leq poly(tax(A), price(A), tie(A), m, n)$

Our goal in this section is to provide a characterization of the communication complexity of truthful mechanisms. Our first task is to develop a low communication protocol that lets $n-1$ players find the menu they present to the remaining player.

Theorem 3.1 (The Menu Reconstruction Theorem): Fix a truthful mechanism A. Denote by v_{-i} the valuation profile of all players except i. The communication complexity of finding the index of the menu presented to i by v_{-i} is $poly(tax(A), price(A), m, n)$.

In the statement of the theorem, we denote by $price(A)$ the communication complexity of the following $(n-1)$-player problem: fix some truthful mechanism A, player i, and bundle S. The input of each player $i' \neq i$ is a valuation $v_{i'}$. Let $\mathcal{M}$ be the menu that is presented to i in A when the valuations are v_{-i}. $price(A)$ is the minimal number of bits that is required to compute $\mathcal{M}(S)$, considering all players i and bundles S.

Due to lack of space, we bring the proof of the theorem in the full version. We note that although for concreteness the menu reconstruction theorem is proved for combinatorial auctions, it actually applies to *any* domain. That is, fix any truthful n-player mechanism A whose range is a set of alternatives $\mathcal{A}$. Then, the menu presented to any player can be found using $poly(tax(A), price(A), \log|\mathcal{A}|, n)$ bits (again, $price(A)$ is the communication complexity of finding the price of an alternative $S \in \mathcal{A}$).

We now want to express the communication complexity of menu reconstruction in terms of $cc(A)$. In the full version we show that if A is truthful for additive valuations, then $price(A) \leq cc(A)$. Since Theorem 2.3 gives us that for general valuations $tax(A) \leq cc(A)$, we get that finding the full menu is not much harder than determining the allocation and the prices of the winning bundles:

Corollary 3.2: Fix a mechanism A that is truthful for general valuations. Denote by v_{-i} the valuation profile of all players except i. The communication complexity of finding the index of the menu presented to i by v_{-i} is $poly(cc(A), m, n)$.

Similar bounds also hold for other valuation classes, by using Theorem 2.3 appropriately.

More importantly, we can now characterize the communication complexity of truthful mechanisms. Let $tie(A)$ be the communication complexity of determining the allocation of A when the valuation of each player i is v_i and all players know the menu $\mathcal{M}_i$ player i is presented with by v_{-i}.[6] Notice that $tie(A) \leq cc(A)$, since we can always run A and ignore the extra information about the $\mathcal{M}_i$'s. This gives a characterization of the communication complexity of mechanisms that are truthful for general valuations (again, similar bounds hold by applying other parts of Theorem 2.3):

Theorem 3.3: (characterization of the communication complexity of truthful mechanisms) Fix a mechanism A that is truthful for general valuations. Then:

$$\frac{tax(A) + price(A) + tie(A)}{3} \leq cc(A) \leq poly(tax(A), price(A), tie(A), m, n)$$

Proof: We first prove the LHS. We always have that $tie(A) \leq cc(A)$. A is truthful for general valuations and hence it is also truthful for additive valuations. In the full version we use this to show that $price(A) \leq cc(A)$. To finish this part, observe that by Theorem 2.3, $tax(A) \leq cc(A)$.

The RHS is obtained by applying the menu reconstruction theorem n times, once for each player. Then, we need additional $tie(A)$ communication bits to determine the final allocation. ■

In the full version we show that our characterization is tight in the sense that if we drop at least one of the three main terms ($tax(A), price(A), tie(A)$) then the gap between the LHS and the RHS might be exponential. For instance, we have already mentioned an example of a truthful mechanism A with $tax(A) = 1$ (and thus $price(A) = 0$) in which $cc(A) = exp(m)$. We also provide examples with similar gaps when dropping $price(A)$ and $tax(A)$.

IV. FROM EX-POST NASH TO DOMINANT STRATEGY

The revelation principle implies that if there is a mechanism that implements some social choice function in an ex-post Nash equilibrium, there is also a mechanism that implements the same social choice function in a dominant strategy equilibrium. Unfortunately, the communication complexity of the latter mechanism might be exponential comparing to the communication complexity of the former (e.g., the already mentioned example of the VCG mechanism for gross substitutes valuations). We provide a more efficient transformation.

Proposition 4.1: Let A be a two player mechanism for combinatorial auctions that reaches an ex-post Nash equilibrium. Then, there is a mechanism A' that implements the same social choice function in dominant strategies with communication complexity $2(tax(A) + m) + tie(A) \leq 2(tax(A) + m) + cc(A)$. In particular, if A is truthful for general valuations, then the communication complexity of the new implementation is $3cc(A) + 2m$.

In particular, for general valuations we pay "almost nothing" (communication-wise) for strengthening the solution concept (simply using the fact that for general valuations by Theorem 2.3 we have that $tax(A) \leq cc(A)$). Similar transformations are possible of course for other classes of valuations using Theorem 2.3. Before formally proving Proposition 4.1 we provide some intuition. A naive proof for this proposition would be the following protocol:

1) Each player i simultaneously sends $tax(A)$ bits that denote the index of $\mathcal{M}_i$ – the menu he presents to the other player.
2) Each player i sends a description of some maximum profit bundle T_i in the menu presented to him (m bits for each player). Denote the price of T_i in the menu by p_i.
3) Each player i is assigned T_i and pays p_i.

This protocol "almost works" except that it is not clear how each player i chooses which maximum-profit bundle to report if there are several bundles that maximize the profit. To solve this we have to be able to break ties correctly, and make sure that each player has a dominant strategy.

Proof: (of Proposition 4.1) We start with some definitions. Let P be some protocol. Given strategies strategies $s_1(\cdot), \ldots, s_n(\cdot)$, we say that a (possibly partial) transcript T of P is *consistent* with a valuation profile $(v_1, \ldots, v_n)$ if the transcript is T when each player i is playing $s_i(v_i)$. Consider a transcript T of P that is not consistent with any valuation profile $(v_1, \ldots, v_n)$. Let T' be the minimal prefix of T that is not consistent and let i be the player that sent the last bit in T'. Player i is the *inconsistent player* of T'.

[6] The notation $tie(A)$ hints that this is a question about tie breaking: each player must be allocated a profit maximizing bundle and the set of profit maximizing bundles can be computed without additional communication as it depends on the presented menu and v_i only. However, deciding which specific bundle in the set the player is allocated might depend also on the valuations of the other players and might require extra communication.

For each player i, let s_i denote the equilibrium strategy of each player i in A. The mechanism A' is the following:

1) Each player i simultaneously sends $tax(A)$ bits that denote the index of $\mathcal{M}_i$ – the menu he presents to the other player in $s_i(v_i)$.
2) Each player i sends a description of some maximum profit bundle T_i according to the menu presented to him by the other player (m bits for each player). Denote the price of T_i in the menu by p_i.
3) Run the mechanism A.
4) For each player i let $s'_i(v_i)$ be the strategy where in the first step player i sends the index of $\mathcal{M}_i$ and in Step 3 player i plays as in $s_i(v_i)$.
5) If there exist valuations v_1, v_2 such that the transcript is consistent with $s'_1(v_1)$ and $s'_2(v_2)$ then the outcome of A' is identical to that of A. Otherwise, let i be the inconsistent player. In this case, player i is not allocated any bundle and pays nothing. The other player wins the bundle T_i and pays p_i.

Observe that if each player i with valuation v_i plays $s'_i(v_i)$ then the outcome of A' is identical to the outcome of A when each player i plays $s_i(v_i)$. The statement of the proposition regarding the communication complexity of A' is obvious as well. It remains to show that s'_i is a dominant strategy. We start with two helper claims.

Claim 4.2: If player i's strategy is $s'_i(v_i)$, for some v_i, then player i is not an inconsistent player.

Proof: Assume that $i = 1$, but the proof is essentially the same for $i = 2$. Let q_2 be the strategy of player 2. Consider a run of A'. If player 2 is an inconsistent player then there is nothing to prove. Therefore, we will consider the messages sent by player 1 one by one, in each point assuming that the transcript so far is consistent with some strategies $(s'_1(v_1), s'_2(v_2))$. Now consider player 1 sending his next message according to $s'_1(v_1)$. Notice that this message is identical to the message that is sent at this point in the transcript where the players use strategies $s'_1(v_1)$ and $s'_2(v_2)$. In particular the next message according to $s'_1(v_1)$ that player 1 sends does not make him inconsistent since the partial transcript is identical to the prefix of the final transcript when both players are playing according to $s'_1(v_1)$ and $s'_2(v_2)$. ∎

Claim 4.3: If player i with valuation v_i uses the strategy $s'_i(v_i)$ then his profit is $v_i(T_i) - p_i$.

Proof: As before, assume that $i = 1$, the proof is essentially the same for $i = 2$. By Claim 4.2 player 1 is not an inconsistent player. If player 2 is the inconsistent player, then by the definition of the protocol player 1 is assigned T_1 and pays p_1 so his profit his $v_1(T_1) - p_1$, as needed.

Therefore, assume that the strategies played are consistent with some strategies $(s'_1(v_1), s'_2(v_2))$. Denote the bundle that player 1 got by T'_1 and his payment by p'_1. Now recall that the outcome of A' with these strategies is identical to the outcome of A with the strategies $(s_1(v_1), s_2(v_2))$ and that since these strategies form an ex-post Nash equilibrium in A, it must be that T'_1 is a maximum-profit bundle according to the menu $\mathcal{M}_2$. However, T_1 is also a maximum profit bundle according to $\mathcal{M}_2$ and thus $v_1(T'_1) - p'_1 = v_1(T_1) - p_1$, which finishes the proof. ∎

To conclude the proof it suffices to show that:

Claim 4.4: For every player i with valuation v_i, $s'_i(v_i)$ is a dominant strategy in A'.

Proof: We prove the claim for $i = 1$, the claim for $i = 2$ is essentially the same. Fix a strategy q_i for every player i. If player 2 is inconsistent then by claim 4.3 the profit of player 1 is $v_1(T_1) - p_1$. If player 1 plays a different strategy that makes player 2 consistent, the menu $\mathcal{M}_2$ must still be consistent with the menu presented to him in the first step, otherwise player 2 is inconsistent. Since T_1 maximizes the profit, the profit of player 1 is at most $v_1(T_1) - p_1$.

Let S_1 be the bundle that player 1 is allocated in A'. Let $\mathcal{M}_2$ be the menu that player 1 is presented with at the first step of the protocol. We will show that the payment of player 1 is $\mathcal{M}_2(S_1)$ – the price of S_1 in $\mathcal{M}_2$. Claim 4.3 gives us that player 1 is not worse off playing $s'_1(v_1)$ in this case. If player 1 is inconsistent then his profit is 0 and, again, he is not worse off playing $s'_1(v_1)$, which completes the proof.

We now show that if player 1 is consistent then the payment of player 1 is $\mathcal{M}_2(S_1)$. This is obvious if player 2 is inconsistent. Else, the strategies are consistent with some $(s'_1(v'_1), s'_2(v'_2))$. Since player 2 is consistent, we have that when player 2 uses the strategy $s_2(v'_2)$ he presents to player 2 the menu $\mathcal{M}_2$ and player 1 chooses a profit-maximizing bundle from that menu. Since the outcome of A with strategies $(s_1(v'_1), s_2(v'_2))$ is identical to the outcome of A' with strategies $(s'_1(v'_1), s'_2(v'_2))$ we have that the price of S_1 in A' is determined according to $\mathcal{M}_2$. ∎

This concludes the proof of the proposition. ∎

Open Question 3: Is there a social choice function f (for three players or more) for combinatorial auctions with general valuations such that:

- There exists a protocol with communication complexity $poly(m)$ that implements f in ex-post Nash equilibrium.
- Any dominant-strategy implementation of f requires $exp(m)$ bits.

Acknowledgments

I thank Sigal Oren for valuable advice and countless discussions during the work on this paper. I would also like to thank Hu Fu, Nitzan Uziely and Omri Weinstein for comments on drafts of this paper. I am also grateful to the participants of the Hebrew University's AGT seminar for their very helpful feedback when this work was in a preliminary stage.

References

[1] Moshe Babaioff, Liad Blumrosen, Moni Naor, and Michael Schapira. Informational overhead of incentive compatibility. In *EC*, 2008.

[2] Moshe Babaioff, Robert Kleinberg, and Aleksandrs Slivkins. Multi-parameter mechanisms with implicit payment computation. In *EC*, pages 35–52. ACM, 2013.

[3] Patrick Briest, Piotr Krysta, and Berthold Vöcking. Approximation techniques for utilitarian mechanism design. In *STOC*, pages 39–48, 2005.

[4] Amit Daniely, Michael Schapira, and Gal Shahaf. Inapproximability of truthful mechanisms via generalizations of the VC dimension. In *STOC*, 2015.

[5] Shahar Dobzinski. An impossibility result for truthful combinatorial auctions with submodular valuations. In *STOC*, pages 139–148, 2011.

[6] Shahar Dobzinski. Breaking the logarithmic barrier for truthful combinatorial auctions with submodular bidders. *STOC*, 2016.

[7] Shahar Dobzinski and Shaddin Dughmi. On the power of randomization in algorithmic mechanism design. In FOCS'09.

[8] Shahar Dobzinski and Noam Nisan. Limitations of VCG-based mechanisms. In *STOC*, pages 338–344, 2007.

[9] Shahar Dobzinski, Noam Nisan, and Sigal Oren. Economic efficiency requires interaction. In *STOC*, 2014.

[10] Shahar Dobzinski, Noam Nisan, and Michael Schapira. Approximation algorithms for combinatorial auctions with complement-free bidders. In *STOC*, pages 610–618, 2005.

[11] Shahar Dobzinski and Jan Vondrák. The computational complexity of truthfulness in combinatorial auctions. In *EC*, pages 405–422, 2012.

[12] Shaddin Dughmi, Tim Roughgarden, and Qiqi Yan. From convex optimization to randomized mechanisms: Toward optimal combinatorial auctions for submodular bidders. In *STOC*, pages 149–158, 2011.

[13] Shaddin Dughmi and Jan Vondrák. Limitations of randomized mechanisms for combinatorial auctions. In *FOCS*, pages 502–511, 2011.

[14] Ronald Fadel and Ilya Segal. The communication cost of selfishness. *Journal of Economic Theory*, 144(5):1895–1920, 2009.

[15] Uriel Feige and Jan Vondrák. Approximation algorithms for allocation problems: Improving the factor of 1-1/e. In *FOCS*, pages 667–676, 2006.

[16] Roger Guesnerie. On taxation and incentives: further remarks on the limits to redistribution. *University of Bonn*, 89, 1981.

[17] Peter J Hammond. Straightforward individual incentive compatibility in large economies. *The Review of Economic Studies*, pages 263–282, 1979.

[18] Sergiu Hart and Noam Nisan. The menu-size complexity of auctions. In *EC*, 2013.

[19] Donald E Knuth. The asymptotic number of geometries. *Journal of Combinatorial Theory, Series A*, 16(3):398–400, 1974.

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

[21] Ron Lavi, Ahuva Mu'alem, and Noam Nisan. Towards a characterization of truthful combinatorial auctions. In *FOCS*, pages 574–583, 2003.

[22] Ron Lavi and Chaitanya Swamy. Truthful and near-optimal mechanism design via linear programming. In FOCS'05.

[23] Daniel Lehmann, Liadan Ita O'Callaghan, and Yoav Shoham. Truth revelation in approximately efficient combinatorial auctions. In *JACM 49(5)*, pages 577–602, Sept. 2002.

[24] R Paes Leme. Gross substitutability: An algorithmic survey. *preprint*, 2014.

[25] Ahuva Mu'alem and Noam Nisan. Truthful approximation mechanisms for restricted combinatorial auctions. In AAAI'02.

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

[27] Noam Nisan. Survey: Algorithmic mechanism design (through the lens of multi-unit auctions). *Handbook of Game Theory*, pages 263–282, 2014.

[28] Noam Nisan and Ilya Segal. The communication requirements of efficient allocations and supporting prices, 2006. In the Journal of Economic Theory.

[29] Tim Roughgarden. Barriers to near-optimal equilibria. In *FOCS*, 2014.

[30] Christopher A Wilkens and Balasubramanian Sivan. Single-call mechanisms. *ACM Transactions on Economics and Computation*, 3(2):10, 2015.