The Communication Complexity of Combinatorial Auctions
Draft – Comments welcome

Noam Nisan *

Institute of Computer Science,
The Hebrew University of Jerusalem

Abstract

We show that any implementation of combinatorial auctions that produces efficient allocations requires an exponential amount of information transfer. The lower bound is independent of any computational complexity considerations, holds even if only an approximately efficient outcome is achieved, holds whether or not the bidder strategies are in equilibrium, and holds even if all bidder valuations have decreasing marginal utilities. This is in contrast to Ausubel's efficient auction for heterogeneous goods that applies in the case that all bidder valuations satisfy the "gross substitutes" condition. The lower bound implies that mechanisms such as AUSM, iBundle or any of the suggested variants of ascending auctions with package bids cannot, in the general case, ensure both the efficiency of the outcome and sub-exponential communication.

1 Introduction

We have recently seen great interest in, so called, *combinatorial auctions* – auctions in which multiple heterogeneous items are concurrently sold, and bidders are able to express preferences for *combinations* of items. Such auctions may be useful in a host of situations where items may by complementary or substitutes, with the recent candidate "killer applications" of spectrum licenses and online procurement. The reader is referred to, e.g., [18, 20, 12] for background and an overview.

While auctions of a single good and of multiple homogenous goods are quite well understood, combinatorial auctions are not so. The underlying reason for this difficulty is the exponential amount of information that needs to be handled: every bidder may have a valuation for each subset of the items. Since there are an exponential number of subsets, agent types are given by an exponential amount

*E-mail:noam@cs.huji.ac.il. Supported by the Israeli Academy of Sciences.

of information. Simplistic mechanisms (such as direct revelation mechanisms) will thus require bidders to communicate an exponential amount of information to the mechanism. This exponential blowup in communication is problematic in several respects:

1. **Valuation Determination:** A bidder will normally require some effort for determining his valuation for a subset of items. Determining the valuation for all subsets may require an impossible amount of effort.

2. **Bid Communication:** The mere transfer of an exponential amount of information is technologically infeasible over any communication medium when the number of items is larger than 20 or 30.

3. **Information Revelation:** Bidders will normally not want to reveal to competitors their full valuation, but rather to only provide the minimal amount of information needed to determine their winning bundle.

It should be emphasized that the communication bottleneck is different from the often considered computational problem of "winner determination" – computing the optimal allocation once valuations are known. The computational problem remains theoretically intractable even if each bidder is interested in only a single package of items, while this case exhibits no communication bottleneck. However, it seems that the communication bottleneck is actually more severe than the computational one. First, it "kicks in" already when there are a dozen or two items, while it seems that the computational complexity can be handled for up to hundreds of items (and thousands of bids) optimally [20, 19] and thousand of items (with tens of thousand of bids) near-optimally [22]. Additionally, the computational burden may be side-stepped by transferring it to the bidders themselves, e.g. by requiring the bidders to suggest matches to their package [2] or by allowing them to suggest allocations [15].

One may be tempted to by-pass the communication bottleneck by restricting the number of bundles that a bidders' valuation may assign non-zero weight to. Such a restriction however is much too severe for most purposes. Even simple valuations such as "I'll pay 1 for any set of $n/2$ items", or "I'll pay 1 for any single frequency band in each geographic region" place valuations on an exponential number of packages. Even allowing combinations of a small number of package bids using either "OR-bidding" or "XOR-bidding" will not enable such simple valuations to be succinctly communicated, but rather more complicated combinations expressed in some "bidding language" will be needed [13].

Much work is currently being done to design multi-round combinatorial auctions, where at each stage bidders place a bid on just a single package (or reasonable number of packages). Several such mechanisms have been suggested to the FCC (e.g. [12]) as well as elsewhere [2, 5, 16, 17, 1]. Presumably, in such a multi-round auction, bidders will only send bids on subsets that "matter", allowing the mechanism to reach an efficient outcome using only a reasonable amount of communication. The only complete proof was given by Ausubel [1] who has designed such an auction for the special case where all valuations satisfy the "gross substitutes" condition. Our main theorem states that such a mechanism with provable properties cannot be designed for the general case!

Theorem 1: Any mechanism that obtains an efficient allocation in combinatorial auctions will require an exponential amount of communication from bidders for some valuations.

It should be noted that this impossibility result is totally independent of any computational issues. The proof of this theorem uses techniques from the field of *communication complexity* [8]. This field abstracts and quantifies the communication requirements that are implicit in many computational settings. Surprisingly, in many situations the communication itself is a major bottleneck – as it is in our case. For the sake of readers who are not familiar with the field of communication complexity we provide a complete self-contained proof. The impossibility result is very general and applies even with the following relaxations (some requiring more difficult proofs):

- **Equilibrium:** The lower bound holds for any tuple of strategies, whether they are in equilibrium or not.
- **Distribution:** The lower bound applies to certain a-priori distributions on valuations, and for these distributions, any mechanism that uses sub-exponential communication will produce non-efficient results on a significant fraction of valuations. The point is not that such distributions are realistic in any sense – they are not. Rather, the point is that any claim that some mechanism produces good results must explicitly rely on the valuation distribution.
- **Revenue:** The lower bound also applies in cases where the revenue of a VCG auction is equal to the total efficiency. Thus exponential communication is also needled for optimizing revenue.
- **Simplicity:** The lower bounds holds even if there are only two bidders and even if all valuations of subsets are either 0 or 1.
- **Approximation:** The lower bound holds even if the mechanism is only required to produce an "approximately efficient" allocation – e.g. with efficiency loss proportional to some minimum bid increment size. Even more, it holds for any non-trivial approximation – one with efficiency that is within a factor of k from optimal, where k is the number of bidders (for k sufficiently smaller than n). Interestingly, procurement auctions are different, and may be approximated to within a factor of $\log n$ (with myopic bidders) using polynomial communication, but no better. The approximate case relies on new results in the field of communication complexity that are described in a companion paper [14].
- **Decreasing marginal utilities:** The lower bound holds even if all valuation are known to have decreasing marginal utilities, i.e. are submod-

ular. Note that valuations satisfying the slightly stronger "gross substitutes" condition can already be handled with polynomial communication by Ausubel's mechanism [1]. The submodular case may be approximated to within a factor of 2 (with myopic bidders) using only polynomial communication.

This result implies that any formal treatment of mechanisms for combinatorial auctions must rely on some restrictions on the class of valuations. Without such restrictions any mechanism will be theoretically useless either in terms of the required communication or in terms of efficiency obtained. (Practically, of course, such mechanisms may still work reasonably well.) A case in point is Ausubel's mechanism that *provably* achieves efficiency using only a reasonable (polynomial) amount of communication in the case where all valuations satisfy the "gross substitutes" condition. In [4] it was noted that Ausubel's auction may be viewed as a primal-dual linear programming algorithm, and that in fact other primal-dual algorithms can also be converted into mechanisms. We show that in terms of communication this is completely general:

Theorem 2: A Truthful mechanism (with dominant strategies) that uses a polynomial amount of communication exists for every case where valuations are restricted so that the linear programming relaxation will find an optimal allocation.

In addition to the case of "gross substitutes" several other natural cases where the linear program produces optimal results are known [13, 20]. While this mechanism is very un-natural, mimicking a separation-based linear programming algorithm, it does show the possibility of constructing a mechanism in such cases, in contrast with the general impossibility result. One may take the following view of affairs: it is known that the linear program produces an optimal allocation if and only if there exists a Walrasian price equilibrium (with prices on single items) [3]. We thus get that when Walrasian prices exist, then indeed a mechanism may reach an efficient allocation within a reasonable amount of communication; otherwise, an exponential amount of communication is needed.

2 Preliminaries

We will assume that n items are auctioned among k bidders. Each bidder j holds a privately known valuation function v_j, that gives a positive real valuation for each subset of items S. We will assume the standard setting:

- $v_j(\emptyset) = 0$.
- Free disposal: $v_j(S) \leq v_j(T)$ for all $S \subseteq T$.

Thus v_j may be specified by a positive real vector of size $2^n - 1$. Note that v_j may have complementarities, where for some disjoint subsets $v(S \cup T) >$

$v(S) + v(T)$, or substitutabilities, where for some disjoint subsets $v(S \cup T) < v(S) + v(T)$.

The mechanism defines the rules of the auction, in particular specifying which bidder speaks at any point of time, how to interpret what the bidder says, the stopping rules, the final allocation, and the payments. Each player has a strategy σ_j that specifies its behavior as a function of its valuation v_j. We will not be assuming any type of equilibrium from the strategies, and make no further assumptions about them. Any fixed mechanism and fixed k-tuple of strategies completely defines a mapping from each k-tuple of valuation $v_1 ... v_k$ to a *run* of the protocol yielding a particular allocation $S_1 ... S_k$. We will be looking at the total information communicated from the bidders to the mechanism for this tuple of valuations, and call this the "communication pattern" of this tuple. The efficiency of the allocation is $\sum_j v_j(S_j)$, and an efficient mechanism maximizes this efficiency (implicitly we assume here that all items have 0 reservation value.)

3 The Basic Lower Bound

Overview:

The basic setting in the field of communication complexity has two players, each holding some kind of data, who together wish to compute some function of their combined data. Since the function depends on both parts of the data they will need to communicate with each other, and we wish to quantify the amount of communication required. The whole point of this field is to prove that *whatever these players do* in order to solve the problem, a certain amount of communication will be needed. While this setting is not identical to the case of combinatorial auctions, it does correctly abstract a key aspect: we may view the valuations of the players as their "data", the interactions of their strategies with the mechanism as their communication with each other, and the allocation produced by the mechanism as the "computed" output. Thus every mechanism for combinatorial auctions has to implicitly also solve the corresponding communication complexity problem. We will now attempt finding a certain sub-structure in combinatorial auctions that corresponds to a well studied problem in communication complexity.

We restrict ourselves to the case of two bidders, each of them interested in a certain collection of bundles, where each desired bundle contains exactly half of the items (assume that the number of items is even). Assume further that each such desired bundle is valued at 1, and that any subset of items that does not contain one of these desired bundles is valued at 0. An allocation with total efficiency of 2 can be obtained whenever the n auctioned items can be partitioned into two subsets each of size $n/2$, each one desired by one of the bidders. Let $N = \binom{n}{n/2}$ denote the total number of subsets of size $n/2$ from n elements. The desired subsets of each bidder may be described by a binary vector of length N, indexed by sets S, where each bit specifies whether the indexed subset is desired by the bidder. To obtain efficiency of 2, the mechanism now needs to find a location S in the vector where the first bidder has a 1-bit and the second bidder has a 1-bit in location S^c – this will provide a partition $(S : S^c)$ with efficiency 2. If we order the indexing of subsets for the two bidders such that the index of S for the first bidder is the same as the index of S^c for the second bidder, then we are now looking for an index where both bidders have a 1-bit in their vector.

Thus we have shown that any mechanism for a combinatorial auction will implicitly also solve the following communication problem: each bidder holds a binary vector of size N and they need to find an index in which they both hold a 1-bit. This is one of the most well studied problems in the field of communication complexity, and it is known [21] that N bits of communication are required. See [8] for an introduction as well as the result. Since N is exponential in the number of items n, we see that an exponential amount of communication must take place during the operation of the mechanism.

We will now prove this result formally, and fold into the proof the required lower bound on communication complexity[1]. In order to quantify the amount of information exchanged we will need to assume that communication with the mechanism is in sequences of bits.

Theorem 3.1 *Let M be a 2-player mechanism for a combinatorial auction on an even number n of elements and σ_1, σ_2 be any strategies of the two players. If the mechanism, when the players use these strategies, produces the optimal allocation for any pair of player valuations v_1, v_2 then for some pair of valuations a total of at least $\binom{n}{n/2}$ bits are communicated by the bidders to the mechanism.*

Proof:

Our first step is to restrict the set of valuations to a simple form that is easy to handle in the proof. Clearly any general mechanism will also work for these restricted cases. We will require the following notations:

Notation:

- Let H denote the set of all subsets of $\{1..n\}$ of size exactly $n/2$. Denote $N = |H| = \binom{n}{n/2}$.
- For $D \subseteq H$ denote by v_D the valuation where $v(S) = 1$ if for some $T \in D$, $T \subset S$, and $v(S) = 0$ otherwise. (I.e. the value is 1 if S contains a desired subset, where D specifies the collection of desired subsets.)
- Denote $V = \{v_D \mid D \subseteq H\}$ be the set of all valuations defined by such collections of desired subsets. Thus $|V| = 2^N$. We will only be considering valuations in V.

Now one may easily verify that for any player valuations $v_D, v_E \in V$, an allocation can be found with value 1 (give everything to the first player), but an allocation with value 2 can be found if and only if for some partition of $\{1...n\}$

[1] Readers familiar with communication complexity will notice we are using the fooling-set method with a cut-and-paste argument [8].

into two sets $(S : S^c)$ each of size $n/2$, we have that $S \in D$ and $S^c \in E$. The rest of the proof now shows that recognizing these cases requires exponential communication.

Notation: For a set $D \subseteq H$ define $D^* = \{S \in H | S^c \notin D\}$. I.e. those sets whose complement is *not* in D. Note that $D^* \subseteq H$. Denote $v_D^* = v_{D^*}$, and note that $v_D^* \in V$.

The definition of D^* ensures the following two simple claims:

Claim For every D, if the players' valuations are v_D and v_D^* then no allocation will have value more than 1.
Proof: Let S be allocated to player 1 and S^c to player 2. If $|S| < n/2$ then clearly $v_D(S) = 0$. Since $v_D^*(S^c) \leq 1$, the allocation value is at most 1. Similarly, if $|S| > n/2$ then $|S^c| < n/2$ and $v_D^*(S^c) = 0$, while $v_D(S) \leq 1$ and the allocation value is at most 1. So we are left with the case that $|S| = |S^c| = n/2$. If $v(S) > 0$ then $S \in D$, and thus by definition $S^c \notin D^*$ and thus $v_D^*(S^c) = 0$ and thus the total allocation value is at most 1. □

Claim 2: For every $D \neq E$, one of the following two cases holds: either (A) for player valuations v_D and v_E^* there exists an allocation with value 2 or (B) for player valuations v_E and v_D^* there exists an allocation with value 2.
Proof: If $D \neq E$ then either $D - E \neq \emptyset$ or $E - D \neq \emptyset$. Assume, first, $D - E \neq \emptyset$ and let $S \in D$ but $S \notin E$ and thus by definition $S^c \in E^*$. Thus $v_D(S) = 1$, and $v_E^*(S^c) = 1$, producing an allocation $(S : S^c)$ with value 2 for case (A). If, on the other hand, $E - D \neq \emptyset$ then let $S \in E$ but $S \notin D$, and similarly $v_E(S) = 1$ and $v_D^*(S^c) = 1$, producing an allocation $(S : S^c)$ with value 2 for case (B). □

We are now ready for the core of the argument: let us consider the operation of the mechanism (with the fixed player strategies) under all pairs of valuations of the form (v_D, v_D^*), for all possible $D \subseteq H$. For each such pair of player valuations, some information is transferred between the mechanism and the players (according to the mechanism rules and the fixed player strategies) and finally some allocation is reached. For every valuation pair, let us consider the complete sequence of bits communicated from the players to the mechanism throughout the run of the mechanism, and call it the "pattern of communication" for this pair of valuations. The main claim is that this pattern of communication must be different for any two different choices of D. Before proving this claim let us see why it implies the result. Since there are 2^N different choices of D, there must be 2^N different possible communication patterns. However, if the number of bits communicated is always at most t, then there can be at most 2^t different patterns of communication. It follows that the number of bits communicated, $t \geq N = \binom{n}{n/2}$.

Claim 3: For every $D \neq E$, the communication pattern on the pair of valuations (v_D, v_D^*) is different from the communication pattern on (v_E, v_E^*).
Proof: Assume to the contrary that the same communication pattern takes place in both cases, in particular reaching an identical allocation in the end,

whose value is, of course, 1. Now, we claim that if we consider a third case, where the pair of valuation is (v_D, v_E^*) then we would still get the same communication pattern. The reason for this is that player 1 will not be able to distinguish between this case and the case of (v_D, v_D^*) (since his valuation is the same, and player 2's communication pattern is exactly the same), and similarly player 2 will not be able to distinguish this case from the case (v_E, v_E^*). Thus the whole run of the mechanism will never deviate from the cases of (v_D, v_D^*) and (v_E, v_E^*), in particular yielding the same allocation value of 1 at the end. The same argument holds for a fourth case where the pair of valuations is (v_E, v_D^*). However, now we have reached a contradiction since claim 2 guaranteed that either for (v_D, v_E^*) or for (v_E, v_D^*) the efficient allocation has value 2, which is what the mechanism should find. □

This concludes the proof of the theorem.

□

4 Extensions

We discuss shortly here why the basic lower bound applies under all the relaxations mentioned in the introduction.

4.1 Distribution

The proof above essentially showed that an allocation with value 2 can be found on inputs (v_D, v_E) if and only if D and $(E^*)^c$ are not disjoint. Thus any efficient mechanism for combinatorial auctions must, in particular, require as least as much communication as the well known "disjointness problem" in communication complexity. It is known that for some distributions this problem has high "distributional complexity", i.e. that for such distributions, every mechanism that uses sub-exponential communication will err on a constant fraction of valuations. See sections 3.4 and 4.6 of [8] for details.

These sections also show that this type of lower bound implies a randomized lower bound, and thus one that applies to mixed player strategies.

4.2 Approximation

The proof given showed that a mechanism that produces an outcome that is closer than a factor of 2 from optimal already needs exponential communication (since we only used that the mechanism is not allowed an outcome with value 1 when the optimal one has value 2). For the case of 2 bidders, an approximation ratio of 2 is trivially possible, simply by bundling all items and selling them in a simple (single-item) Vickerey auction. For the more general case of k bidders, an approximation to within k is similarly trivial, and it turns out that no mechanism can do better without requiring $2^{\Omega(n/k^2 - k \log k)}$ communication. Thus for all $k \leq n^{1/3-\epsilon}$, any mechanism that achieves better than k-approximation requires exponential communication. (For larger number of bidders, we have $n^{1/3-\epsilon}$ as the upper bound to approximation.)

It should be noted that a $n^{1/2}$ approximation is possible in polynomial communication, but using myopic bidders, by adapting the algorithm of [10].

Somewhat surprisingly, the situation for procurement auctions is different: an approximation to within a factor of $\ln n$ is possible (with myopic bidders) in polynomial communication, adapting the greedy method of [11], but a better approximation ratio requires exponential communication.

The proofs use communication complexity types of reasoning, and are presented in a companion paper aimed at the computer science community [14].

4.3 Revenue

If one considers cases where the k players are partitioned into pairs, both players in each pair having the exact same valuation, then we also get a lower bound for mechanisms that approximate the revenue. In such cases a VCG mechanism will extract all surplus from the auction (i.e. the price for a winning bundle must be equal to the valuation of this bundle). Any mechanism that satisfies participation constraints will not be able to extract more than the efficiency is generates, and thus approximating the revenue in these cases requires approximating the efficiency as well.

4.4 Decreasing Marginal Utilities

The valuations v_D used in the proof exhibit complementarities. If instead we define valuations $u_D(S) = v_D(S) + 1$, then it is easy to verify that u_D is complementarity-free. Now, an allocation achieving value x on a certain pair of v's would achieve value $x + 2$ on the corresponding pair of u's. Thus, it is as hard for a mechanism to distinguish between allocation efficiency 3 and 4 when valuations are complement-free, as it is to distinguish between 1 and 2 for general valuations.

In order to get hard examples with decreasing marginal utilities, just define $w_D(S) = |S|$ for $|S| < n/2$, $w_D(S) = n/2$ for $|S| > n/2$, and for $|S| = n/2$, $w_D(S) = n/2$ if $S \in D$ and $w_D(S) = n/2 - 1$ if $S \notin D$. One may easily verify that such valuations have decreasing marginal utilities (sub-modular), and that an allocation between two w valuations will achieve value n if and only if the same allocation between the corresponding v valuations achieves value 2. Thus any mechanism that produces efficient allocations even where valuations are restricted to be submodular still requires exponential communication.

It should be noted that a 2-approximation for submodular valuations is possible (with myopic bidders) in polynomial communication, adapting the algorithm of [9].

5 Possibility results using Linear Programming

The communication lower bound shown still leaves room for mechanisms that handle restricted classes of valuations. In this section we sketch some possibility results along these lines. These mechanisms are not "pretty" in any sense, rather they only suggest that, for restricted classes of valuations, communication is not a bottleneck.

It is well known that the allocation problem of combinatorial auctions may be phrased as an integer programming problem (see [20] for a survey). This integer programming problem is commonly relaxed to a linear programming problem, and in some cases it is known that this linear program will indeed return integer allocations, solving the original problem as well. In particular it is known that this is the case if all valuations satisfy the "gross substitutes" property [6]. What we wish to point here is that in all such cases, a mechanism that only requires polynomial amount of communication may be implemented.

The basic correspondence between many auction mechanisms and primal-dual methods of linear programming was pointed out in [4]. However, to prove a theorem about communication we will need to use separation-based LP algorithms. A separation-based linear programming algorithm may be used in cases of linear programs that have an exponential number of constraints (inequalities) that are given implicitly. Specifically, such an algorithm does not receive the inequalities as an explicit input, but rather is provided with a "separation oracle" as its input: whenever an infeasible solution is presented to this oracle, it must be able to produce a violated inequality. Algorithms of this type can solve linear programs in polynomial time, given just this type of an oracle. The reader is referred to any textbook on linear programming (e.g. [7]) for more information.

The allocation problem itself is usually phrased as having an exponential number of variables (x^i_S specifying the allocation of the bundle S of goods to bidder i), but only a polynomial number of significant inequalities (that each good is sold only once, and that each bidder is allocated only one set):

Maximize: $\sum_{i,S} x^i_S v_i(S)$
Subject to:

- For all goods j: $\sum_{i,S|j \in S} x^i_S \leq 1$.
- For all bidders i: $\sum_S x^i_S \leq 1$.
- For all i, S: $x^i_S \geq 0$.

In order to use a separation-based linear programming algorithm, we move to the dual that has just a polynomial number of variables:

Minimize: $\sum_j p_j + \sum_i u_i$
Subject to:

- For all i, S: $u_i + \sum_{j \in S} p_j \geq v_i(S)$.

- For all j: $p_j \geq 0$.
- For all i: $u_i \geq 0$.

It is important to note that in the dual, each inequality specifies a condition that depends on the valuation of a single player! Now the separation-based LP is used, where the mechanism runs the separation-based LP algorithm, and whenever an oracle query needs to me made (i.e. a violated inequality is desired), the current solution is presented to all players, and each of them responds with a violated inequality, if any[2]. Since each inequality is held by some single player, if a violated equality exists it is found, and the mechanism can continue. The communication that takes place at each such stage is polynomial, and since the algorithm is known to terminate within a polynomial number of steps, we have that the whole mechanism requires polynomial communication.

In order to ensure truthfulness on the part of all the bidders, the mechanism should simply impose VCG payments on allocated bundles. The calculation of these payments will require running n different linear programs in the same way. This is clearly not very elegant, but still in such a mechanism truth remains a dominant strategy.

Acknowledgements

I thank Ilan Kremer, Daniel Lehmann, Motty Perry, Amir Ronen and Rakesh Vohra for related discussions.

References

[1] Lawrence Ausubel. An efficient dynamic auction for heterogeneous commodities, 2000. Mimeo, University of Maryland.

[2] Jeffrey Banks, John Ledyard, and David Porter. Allocating uncertain and unresponsive resources: An experimental approach. *Rand Journal of Economics*, 20:1 – 25, 1989.

[3] Sushil Bikhchandani and John W. Mamer. Competitive equilibrium in an exchange economy with indivisibilities. *Journal of economic theory*, 74:385 – 413, 1997.

[4] Sushil Bikhchandani, Rakesh Vohra, Sven de Vries, and James Schummer. Linear programming and vickrey auctions, 2001. Availailabe from http://www.kellogg.nwu.edu/faculty/vohra/htm/res.htm.

[5] C. DeMartini, A. Kwasnica, J. Ledyard, and D. Porter. A new and improved design for multi-object iterative auctions, 1998.

[2] A violated inequality is exactly the selection of a subset S of items that would give player i at least utility u_i under the prices p_j. Finding this set may be a computationally difficult problem depending on the representation of v_i, but this does not concern us here.

[6] Faruk Gul and Ennio Stacchetti. Walrasian equilibrium with gross substitutes. *Journal of Economic Theory*, 87:95 – 124, 1999.

[7] H. Karloff. *Linear Programming.* Birkhäuser Verlag, 1991.

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

[9] Benny Lehmann, Daniel Lehmann, and Noam Nisan. Combinatorial auctions with decreasing marginal utilities. In *To appear in ACM conference on electronic commerce*, 2001. Available from http://www.cs.huji.ac.il/˜noam/mkts.html.

[10] Daniel Lehmann, Liadan Ita O'Callaghan, and Yoav Shoham. Truth revelation in rapid, approximately efficient combinatorial auctions. In *1st ACM conference on electronic commerce*, 1999.

[11] L. Lov'asz. The ratio of optimal integral and fractional covers. *Discrete Mathematics*, 13, 1975.

[12] Paul Milgrom. Ascending auctions with package bidding, 2000.

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

[14] Noam Nisan. The communication complexity of approximate set packing and covering, 2001. Preliminary version available from http://www.cs.huji.ac.il/˜noam/mkts.html.

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

[16] D. Parkes. ibundle: An efficient ascending price bundle auction, 1999.

[17] David C. Parkes and Lyle H. Ungar. Iterative combinatorial auctions: Theory and practice. In *AAAI/IAAI*, pages 74–81, 2000.

[18] Michael H. Rothkhof, Aleksandar Pekeč, and Ronald M. Harstad. Computationally manageable combinatorial auctions. *Management Science*, 44(8):1131–1147, 1998.

[19] T. Sandholm, S. Suri, A. Gilpin, and Levine D. Cabob: A fast optimal algorithm for combinatorial auctions. In *IJCAI*, 2001.

[20] Rakesh Vohra and Sven de Vries. Combinatorial auctions: A survey, 2000. Availailabe from http://www.kellogg.nwu.edu/faculty/vohra/htm/res.htm.

[21] Andrew Chi-Chih Yao. Some complexity questions related to distributive computing. In *ACM Symposium on Theory of Computing*, pages 209–213, 1979.

[22] Edo Zurel and Noam Nisan. An efficient approximate allocation algorithm for combinatorial auctions. In *To appear in ACM conference on electronic commerce*, 2001. Available from http://www.cs.huji.ac.il/~noam/mkts.html.