The communication Requirements of Efficient Allocations and Supporting Lindahl Prices

A self-contained CS-friendly executive summary

Noam Nisan* Ilya Segal†

March 25, 2003

1 Introduction

Much recent work has been aimed at the design and implementation of combinatorial auctions. In our recent paper, "The communication Requirements of Efficient Allocations and Supporting Lindahl Prices", we present an analysis of the communication burden of combinatorial auctions. The paper is written mostly for an audience of economists and presents its analysis in a rather general way. The main result of the paper is really quite simple and is of direct relevance to computer scientists working on combinatorial auctions. This note aims to give a short self-contained presentation of its main result in a way that is accessible to computer scientists. More details, results, discussion, as well as references appear in the full paper.

In a combinatorial auction, a set of L items is auctioned concurrently to a set of N bidders. Each bidder i has a valuation function $v_i : 2^L \rightarrow R^+$ that assigns a private value $v_i(S)$ for each subset of items S. The valuation functions v_i are assumed to be monotone

*Institute of Computer Science, Hebrew University, Jerusalem, Israel, e-mail: noam@cs.huji.ac.il

†Department of Economics, Stanford University, Stanford, CA 94305; e-mail: ilya.segal@stanford.edu

non-decreasing, and $v_i(\emptyset) = 0$. The outcome of the combinatorial auction is a partition $S_1, ...S_N$ of the items, where S_i is the set of items allocated to bidder i. The main goal of the auction is optimizing the social welfare, i.e. finding the partition that maximizes $\sum_i v_i(S_i)$. We would like to do this by means of some protocol between the auctioneer and the bidders, and have this protocol be efficient: polynomial in L and N.

Computationally speaking, even without addressings incentive issues, this goal faces two separate difficulties. The first is the exponential size of the "input": the valuation of each bidder holds $2^L - 1$ values in it. The second is the hardness of the optimization problem: the problem is NP-complete even for very simple valuations. Much computational work on combinatorial auctions addressed the second issue, reaching a state of the art where the optimization problem can be practically solved (or approximated) for hundreds and even thousands of items.

Other work has focused on the first issue, that of eliciting enough information about bidders' valuations. Clearly full "direct revelation" of the valuations to the auctioneer would be exponential. Thus, there have been various suggestions of "iterative" auctions in which only partial information about the valuations need to be revealed to the auctioneer. The hope is that such iterative auctions could reach the optimal (or near-optimal) allocation while still using only a polynomial-length protocol. The main point of the paper is that this hope is not justified in general: every protocol for combinatorial auctions must be exponential in the worst case. The lower bound is independent of any computational limitations, provides exact bounds (rather than only asymptotic ones), and is unconditional (i.e. applies even if $P = NP$.)

2 Communication complexity and its meaning

The lower bound is obtained in Yao's standard model of two player communication complexity. We consider two players, Alice and Bob, each holding a valuation function. We can restrict ourselves to the special case where the value of each set is 0 or 1. Thus the inputs are monotone functions $v_1, v_2 : 2^L \rightarrow \{0, 1\}$. Alice and Bob must embark on a communication protocol whose final outcome is the declaration of an allocation (S, S^c) that maximizes $v_1(S) + v_2(S^c)$. The protocol specifies rules for exchanging bits of information, where Alice's message at each point may depend only on v_1 and on previous messages received from Bob, while Bob's message at each point may depend only on v_2 and on previous messages received from Alice. No computational constraints are put on Alice and Bob – only communication is measured. The main result shows that:

Theorem 1 *Every protocol that finds the optimal allocation for every pair of 0/1 valuations v_1, v_2 must use at least $\binom{L}{L/2}$ bits of total communication in the worst case.*

Since Yao's communication model is very powerful, the lower bound immediately applies to essentially all computational settings where v_1 and v_2 reside in "different places". In particular:

- A combinatorial auction with two bidders exchanging messages with an auctioneer rather than with each other. (A protocol with an auctioneer can be converted into one without an auctioneer, by sending all messages directly to each other and having Alice and Bob simulate the auctioneer.)
- Any larger number of bidders. (The 2-bidder case is a special case where all bidders but two have null valuations.)
- An algorithm where the auctioneer makes various types of queries to the bidders' valuation functions. (The bound applies to the total length, in bits, of the answers.)
- Iterative auctions where bidders' repeatedly make bids on various bundles. (The bound applies to the total length, in bits, of the bids made throughout the protocol.)

The lower bound may also be formulated in a setting where real numbers are communicated rather than bits (as is commonly done in economics). The lower bound proven in the paper applies also to "non-deterministic" communication protocols (where the correct allocation need only be verified rather than found.) For economists, the basic means of communication is prices, and indeed prices are one method of non deterministic communication. Thus we obtain a lower bound to the "number of prices" that need to be announced before an allocation can be found. Somewhat surprisingly, the paper also shows that prices, in a certain Lindahl price formulation, are in fact "complete" for non-deterministic communication. Specifically, Lindahl prices can be obtained as a side-effect of *any* communication protocol for finding efficient allocations (in a very general setting). In the paper this is used as a as the tool for proving the main lower bound. Here we give a direct proof of the main lower bound.

3 The proof

Fix a communication protocol that for every input valuation pair (v_1, v_2) finds an optimal allocation S, S^c. We will construct a "fooling set": a set of valuation pairs with the property that the communication patterns produced by the protocol must be different for different valuation pairs. Specifically, for every 0/1 valuation v, we define the *dual valuation* v^* to be $v^*(S) = 1 - v(S^c)$. Note that (i) v^* is indeed a 0/1 valuation, and (ii) for every partition (S, S^c), $S \subseteq L$, we have that $v(S) + v^*(S^c) = 1$.

Lemma 1 *Let $v \neq u$ be arbitrary 0/1 valuations. Then, the sequence of bits transmitted on inputs (v, v^*), is not identical to the sequence of bits transmitted on inputs (u, u^*).*

Before we prove the lemma let us see how the main theorem is implied. Since different input valuation pairs lead to different communication sequences, we see that the total possible number of communication sequences produced by the protocol is at least the number of valuation pairs (v, v^*), which is exactly the number of distinct 0/1 valuations v. The number of 0/1 valuations can be easily bounded from below by $2^{\binom{L}{L/2}}$ by counting only valuations such that $v(S) = 0$ for all $|S| < L/2$, $v(S) = 1$ for all $|S| > L/2$, and allowing $v(S)$ to be either 0 or 1 for $|S| = L/2$. There are $\binom{L}{L/2}$ sets of size $L/2$, so the total number of such valuations is exponential in this number. The protocol must thus be able to produce $2^{\binom{L}{L/2}}$ different communication sequences. Since these are binary sequences, at least one of the sequences must be of length at least $\binom{L}{L/2}$.

Proof. (of lemma) Assume, by way of contradiction, that the communication sequence on (v, v^*) is the same as on (u, u^*). We first show that the same communication sequence would also be produced for (v, u^*) and for (u, v^*). Consider the case of (v, u^*), i.e. Alice has valuation v and Bob has valuation u^*. Alice does not see u^* so she behaves and communicates exactly as she would in the (v, v^*) case. Similarly, Bob behaves as he would in the (u, u^*) case. Since the communication sequences in the (v, v^*) and the (u, u^*) cases are the same, neither Alice nor Bob ever notice a deviation from this common sequence, and thus never deviate themselves. In particular, this common sequence is followed also on the (v, u^*) case. Thus, the same allocation (S, S^c) is produced by the protocol in all 4 cases: $(v, v^*), (u, u^*), (v, u^*), (u, v^*)$. We will show that this is impossible, since a single allocation cannot be optimal for all 4 cases.

Since $u \neq v$, we have that for some set T, $v(T) \neq u(T)$. Without loss of generality, $v(T) = 1$ and $u(T) = 0$, and so $v(T) + u^*(T^c) = 2$. The allocation (S, S^c) produced by the protocol must be optimal on the valuation pair (v, u^*), thus $v(S) + u^*(S^c) \geq 2$. However, since $(v(S) + v^*(S^c)) + (u(S) + u^*(S^c)) = 1 + 1 = 2$, we get that $u(S) + v^*(S^c) \leq 0$. Thus (S, S^c) is not an optimal allocation for the input pair (u, v^*) – contradiction to the fact that the protocol produces it as the output in this case as well. ■

4 Other results in the paper

The paper discusses the following issues as well:

1. **Models of communication complexity:** The bounds in the paper apply to various variants of communication complexity: deterministic and non-deterministic, discrete and continuous, worst case and distributional. The paper discusses and compares these models.

2. **Lindahl prices:** Economists usually think of prices as communication. Computer scientists will normally view prices as only one method of communication. The paper shows that so called, Lindahl prices, are in fact a "complete" method of communication. In the paper this is used as the tool for proving the lower bound, as well as a way of gaining intuition.

3. **Sub-families of valuations:** For which sub-families of valuations, is it possible to use a small amount of communication? The paper shows that for submodular valuations exponential communication is needed, but for the more restricted family of "(gross) substitute" valuations, polynomial communication suffices.

4. **Approximation:** How well can we *approximate* the optimal allocation? It turns out that as long as the number of players is sufficiently smaller than the number of items, achieving an approximation that is better than selling all items as a bundle requires exponential communication. For the case of homogenous items, on the other hand, arbitrarily good approximations can be obtained using exponentially less communication than that required for obtaining the exact optimum.