Algorithmic Game Theory October 13, 2008

Lecture 6: Communication Complexity of Auctions

Lecturer: Sébastien Lahaie *Scribe: Rajat Dixit, Sébastien Lahaie*

In this lecture we examine the amount of communication performed in sealed-bid and iterative combinatorial auctions. A sealed-bid auction is clearly expensive in terms of communication—each agent must transmit a nonlinear valuation to the auctioneer. In general, an iterative auction need not have the agents completely reveal their valuations before an efficient allocation is reached. We would like to lower bound the communication needed by iterative auctions in the worst-case, to compare it to that of sealed-bid auctions. In this lecture, we do not consider incentives at all; we assess the communication needed by combinatorial auctions even if agents bid truthfully.

1 Sealed-Bid Auctions

As usual let $\mathcal{V}$ denote the set of general valuations, namely monotone and normalized valuations. The following outlines the steps in a generic sealed-bid auction.

Generic Sealed-Bid Auction

1. Collect a valuation $\tilde{v}_i \in \mathcal{V}$ from each bidder.
2. Select an efficient allocation $R \in \Gamma^*$.
3. Charge each agent $i \in N$ a payment q_i.

Step 1 dominates the communication and it is common to all sealed bid auctions. To characterize the amount of communication required, we consider finite sets of valuations; so consider the set of bounded, integer-valued general valuations, $v_i : 2^M \to \mathbf{Z}_+$ such that $v_i(M) \leq C$ for some constant C. To transmit such a valuation, we can simply transmit the value of every non-empty bundle. There are $2^m - 1$ such bundles, and communicating a value is on the order of $\log C$ bits, so the communication uses $O(2^m \log C)$ bits. (Throughout, log is always understood base 2.)

This counting argument is somewhat unrefined, because with so many bits we can transmit *any* valuation over bundles, not necessarily just monotone valuations. Thus we may need less bits than this to communicate a general valuation. In general, to encode the elements of a finite set of size k, we need on the order of $\log k$ bits (to be precise, $\lceil \log_2 k \rceil$ bits). Thus we would like to lower bound the size of the set of general valuations.

Consider the set of {0-1}-valued general valuations, denoted $\mathcal{V}_{01}$. Assume for the moment that m is even. Included in $\mathcal{V}_{01}$ is the set of all valuations of the form

$$v_i(S) = \begin{cases} 1 & \text{if } |S| > m/2 \\ 0 \text{ or } 1 & \text{if } |S| = m/2 \\ 0 & \text{if } |S| < m/2 \end{cases}$$

Such valuations are clearly monotone and normalized. There are $\binom{m}{m/2}$ sets of size $m/2$, and each can have take on one of two values (0 or 1), so there are $2^{\binom{m}{m/2}}$ valuations of this form. Therefore, the number of bits needed to communicate a {0-1}-valued general valuation is at least

$$\log 2^{\binom{m}{m/2}} = \binom{m}{m/2} \approx \sqrt{\frac{2}{\pi m}} 2^m$$

where the approximation holds for large m (Stirling's approximation).

To summarize, step 1 in the generic sealed-bid auction takes $\Omega(n\binom{m}{m/2})$ bits, step 2 takes $n \log m$ bits to communicate the allocation (you can check this by counting the number of possible feasible allocations), and step 3 takes $\Omega(n \log C)$ bits in general. This is all dominated by step 1, whose communication is exponential in m. Let us see how iterative auctions compare.

2 Iterative Auctions

The following outlines the steps in a generic iterative auction for several items. Recall that the iterative auctions we have covered all converge to a *competitive equilibrium*: an efficient allocation together with prices such that demand balances supply.

Generic Iterative Auction

1. Quote prices p^t in round t.
2. Collect the demand of each bidder.
3. Select an allocation R^t that satisfies as much demand as possible while maximizing revenue.
   - 3a. If each bidder is allocated a demanded bundle, the auction terminates: we have a competitive equilibrium $\langle R^t, p^t \rangle$.
   - 3b. Otherwise, update the prices and go to step 1 to start another round.

The prices in step 1 are nonlinear and non-anonymous in general. In step 2, the "demand" of a bidder is either a bundle that maximizes the bidder's utility at the given prices, or a set of such bundles (perhaps all of them). Different iterative combinatorial auctions are usually distinguished by their price update rule in step 3b. What virtually all iterative auctions share is that they converge to a competitive equilibrium.

In light of this, to lower bound the communication performed in an iterative auction, we will lower bound the communication needed to transmit a competitive equilibrium. This may seem like a very weak bound: we are only considering the communication of the very last round! However, the bound will turn out to be surprisingly strong, and there are currently no good techniques to lower bound communication across all rounds.

So the questions is: how much communication is needed to transmit a competitive equilibrium? This is a subtle question because for any profile of valuations, there may not necessarily be a unique competitive equilibrium—any will do. All we need is a set of allocation-price pairs (of the form $\langle R, p \rangle$) such that for each $v \in \mathcal{V}^n$, there is a competitive

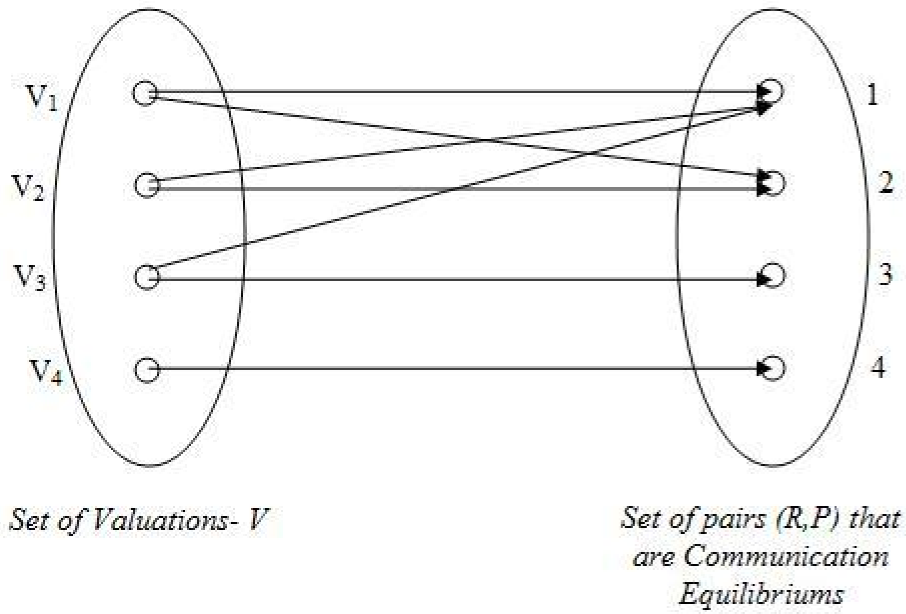

Figure 1: Mapping of valuation profiles to competitive equilibria.

equilibrium with respect to v in the set. Consider Figure 1, which shows how different *profiles* of valuations map to corresponding competitive equilibria. There may be multiple equilibria associated with a single profile, as with profiles v_1, v_2, and v_3. If we just consider the set of equilibria consisting of 1 and 4, then we are assured to have a competitive equilibrium in the set no matter what the profile of valuations turns out to be. So the communication needed to transmit an equilibrium is $\log 2 = 1$ bit in this example. To check whether you understand this clearly, you should ask yourself: is it always the case that at least 1 bit of communication is needed? (The answer is no—why?)

To lower bound the size of any set of competitive equilibria that is guaranteed to include a competitive equilibrium for each possible valuation profile, we use the concept of a "fooling set", fundamental in communication complexity.

Definition 1 *A* fooling set *is a subset of valuation profiles $V \subseteq \mathcal{V}^n$ such that no two distinct $v, v' \in V$ share a common competitive equilibrium.*

The size of a fooling set lower bounds the size of any complete set of competitive equilibria, because we need a distinct competitive equilibrium for each profile in the set. Therefore, we would like to find a large fooling set. We will use graphical intuition to do this. First, let us recall the formal definitions of an efficient allocation and competitive equilibrium.

Definition 2 *A feasible allocation $R \in \Gamma$ is* efficient *if it maximizes the sum of the agent's values,*

$$R \in \arg\max_{R' \in \Gamma} \sum_{i \in N} v_i(R'_i).$$

Recall that valuations are *normalized*: $v_i(\emptyset) = 0$. With respect to the efficient allocation problem, this assumption is without loss of generality, because note that we can add a constant c_i to each value $v_i(S)$, and the set of efficient allocations remains unchanged. Specifically, let $v'_i(S) = v_i(S) + c_i$ for all $S \subseteq M$ (including the empty bundle) and $i \in N$, where c_i is a constant. Then observe that

$$\arg\max_{R' \in \Gamma} \sum_{i \in N} v'_i(R'_i) = \arg\max_{R' \in \Gamma} \sum_{i \in N} v_i(R'_i) + \sum_{i \in N} c_i = \arg\max_{R' \in \Gamma} \sum_{i \in N} v_i(R'_i).$$

Therefore we can "renormalize" any valuation by translating it by a constant, without changing the set of efficient allocations. Recall now the definition of competitive equilibrium.

Definition 3 *A pair $\langle R, p \rangle$ consisting of a feasible allocation R and (nonlinear, non-anonymous) prices p is a* competitive equilibrium *if for any feasible allocation R' we have*

$$v_i(R_i) - p_i(R_i) \geq v_i(R'_i) - p_i(R'_i) \tag{1}$$

for all $i \in N$, as well as

$$\sum_{i \in N} p_i(R_i) \geq \sum_{i \in N} p_i(R'_i). \tag{2}$$

According to (1), the set of competitive equilibria does not change if we translate valuations by a constant. The same holds for prices: translating competitive equilibrium prices by a constant still results in prices that satisfy (1) and (2). We these facts in mind we can draw insightful depictions of competitive equilibrium situations.

From now on we restrict our attention to 2-agent profile (v_1, v_2); we assume that $v_3, \ldots, v_n$ are identically zero on all bundles. Let R be an efficient allocation for this profile. Consider Figure 2, which depicts profiles v_1 and v_2 normalized so that $v_1(R_1) = 0$ and $v_2(R_2) = 0$. In this drawing, the horizontal axis corresponds to the set of feasible allocations. (This is stylized, because the set of allocations is actually discrete rather than continuous; even if it were continuous, it would not necessarily be the case that valuations would be continuous.) Because the total value of R is 0 after normalization, and R is efficient, we must necessarily have v_1 lie entirely above v_2; this ensures that the total value of any other allocation is non-positive.

Figure 3 depicts a competitive equilibrium. Here prices are normalized so that $p_1(R_1) = 0$ and $p_2(R_2) = 0$. As just discussed, we can renormalize valuations and prices however we want when depicting a competitive equilibrium. Because R must maximize revenue at prices p, and R has revenue of 0, prices p_1 must lie above prices p_2 to ensure that all other allocations generate non-positive revenue. Also, the utility of R_1 to agent 1 is 0: $v_1(R_1) - p_1(R_1) = 0 - 0 = 0$. Since R_1 maximizes agent 1's utility at prices p, this means that prices p_1 must lie entirely below v_1, to ensure that all bundles give non-positive utility. By the same reasoning, prices p_2 must lie entirely above v_2.

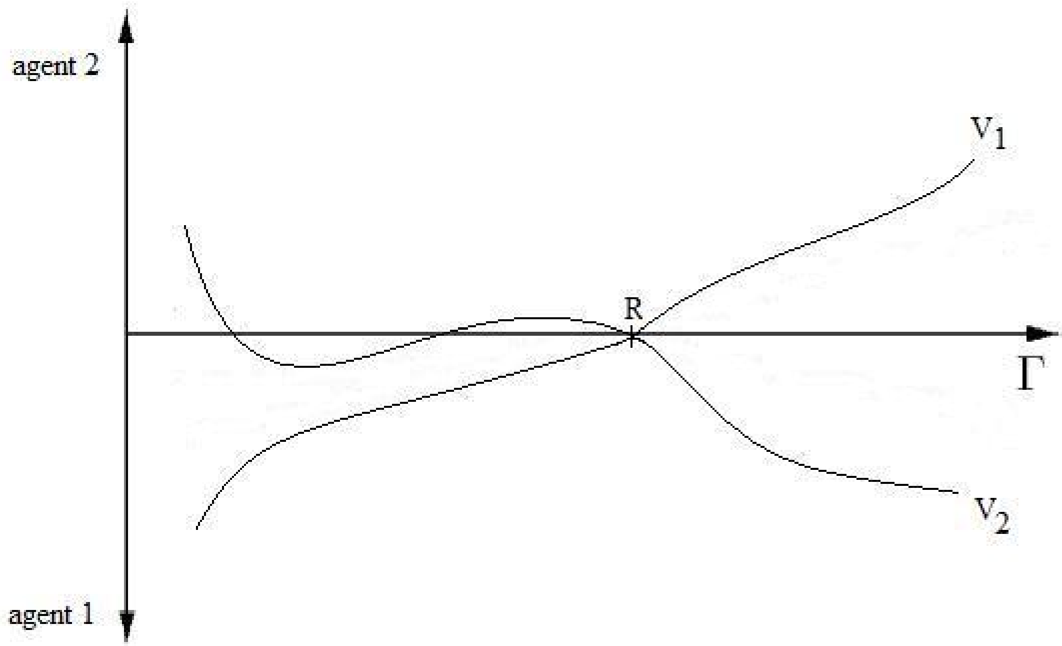

Figure 2: Valuations normalized so that efficient allocation R has value 0 to both agents.

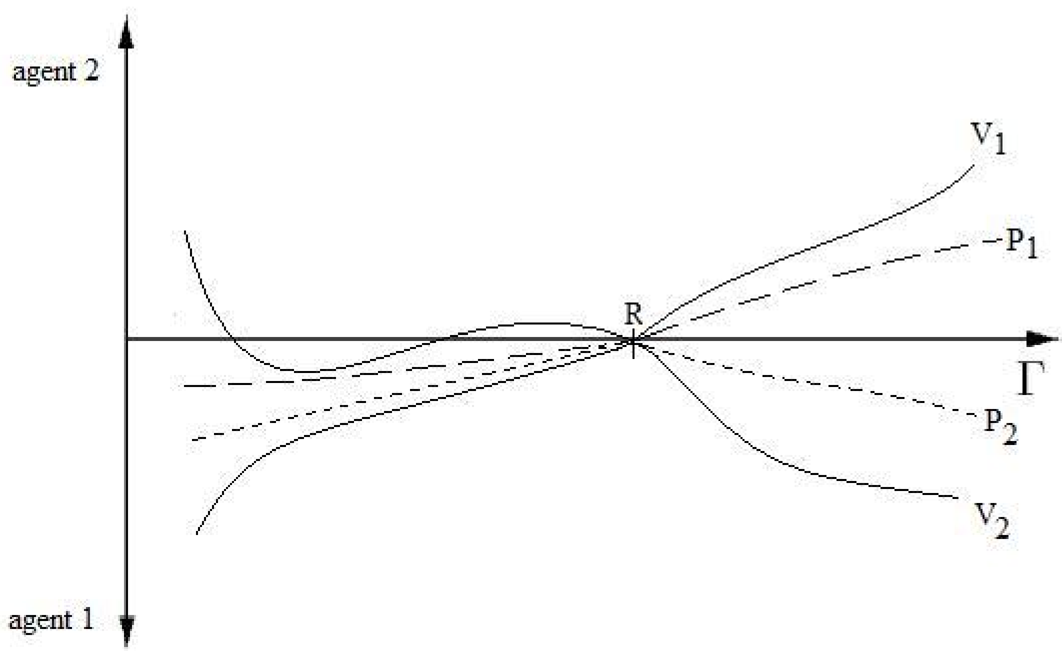

Figure 3: A competitive equilibrium, with valuations and prices normalized so that efficient allocation R has value and price 0 for both agents.

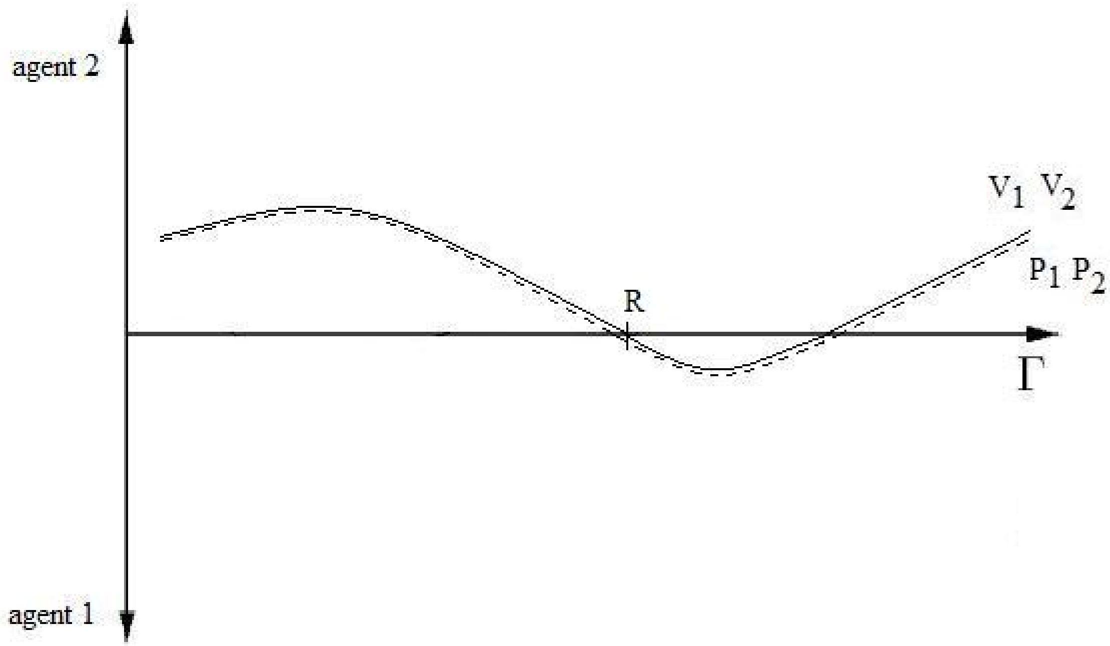

Figure 4: Competitive equilibrium when all allocations are efficient.

To summarize, in competitive equilibrium we must have that

- v_1 lies entirely above v_2, and they touch at R.
- p_1 lies entirely above p_2, and they touch at R.
- p_1 lies entirely below v_1, and they touch at R.
- p_2 lies entirely above v_2, and they touch at R.

By the four points just made, competitive equilibrium prices p_1 and p_2 have to lie between v_1 and v_2. Now suppose that *every* allocation is efficient. In this case v_1 and v_2 must touch everywhere, because every allocation must have a total value of 0. It then follows that competitive equilibrium prices p_1 and p_2 must be "squashed" between the two, and in particular, $v_1 = p_1$ and $v_2 = p_2$. The graphical intuition for this is given in Figure 4. Let's formally check that this is indeed the case.

Lemma 1 *Let $V \subseteq \mathcal{V}^n$ be a subset of valuation profiles such that for each profile $v \in V$, all allocations are efficient. Let p be competitive equilibrium prices with respect to v. Then for each $i \in N$ we have*

$$v_i(S) = p_i(S)$$

for all $S \subseteq M$.

Proof. Prices p support allocations where i gets $\emptyset$ and where i gets S for any $S \subseteq M$. Thus we have

$$v_i(S) - p_i(S) = v_i(\emptyset) - p_i(\emptyset) = 0.$$

□

A subset V satisfying the conditions of Lemma 1 is clearly a fooling set, because for distinct $v, v' \in V$, competitive equilibrium prices must be $p = v$ and $p' = v'$, and p and p' are thus necessarily distinct. This immediately gives the following corollary.

Corollary 1 *Let $V \subseteq \mathcal{V}^n$ be a subset of valuation profiles such that for each profile $v \in V$, all allocations are efficient. Then the number of bits needed to transmit a competitive equilibrium is at least $\log |V|$.*

We now proceed to construct 2-agent profiles that satisfy the conditions of the corollary. We can restrict our attention to allocations of the form $(S, M \backslash S)$; the remaining agents always implicitly get $\emptyset$. We want each such allocation to yield the same total value:

$$v_1(S) + v_2(M \backslash S) = c,$$

where c is a constant. To determine c, note that $v_1(M) + v_2(\emptyset) = c$ and thus $c = v_1(M)$. Thus,

$$v_2(S) = v_1(M) - v_1(M \backslash S).$$

If v_1 and v_2 are related in this way, then all the relevant allocations $(S, M \backslash S)$ are efficient.

Definition 4 *Given valuation v_1, its* dual valuation *v_1^* is defined as*

$$v_1^*(S) = v_1(M) - v_1(M \backslash S).$$

Since all allocations are efficient for profile (v_1, v_1^*), with appeal to Corollary 1 we get our main result.

Theorem 1 *If for each v_1 in domain $\mathcal{V}$ there is a dual valuation $v_1^* \in \mathcal{V}$, then the communication needed to transmit a competitive equilibrium is at least $\log |\mathcal{V}|$ bits in the worst-case.*

To apply this theorem to general valuations and complete our analysis, we should check that v_1^* is a general valuation (i.e., monotone and normalized) for each general valuation v_1; this is straightforward. Hence the communication needed to transmit a competitive equilibrium is at least the communication needed to fully reveal one valuation. Compare this with sealed-bid auctions, which require full revelation of all n valuations. Is it important to stress that these are worst-case results; in practice iterative auctions might do better than this worst-case bound. Nevertheless, the bound is quite strong, because we saw that full revelation of one general valuation requires on the order of 2^m bits, exponential in the number of items.