CS364B: Frontiers in Mechanism Design
Lecture #4: The Clinching Auction*

Tim Roughgarden†

January 15, 2014

1 Multi-Unit Auctions with Downward-Sloping Valuations

This lecture pursues good ascending auctions in a fourth scenario, *multi-unit auctions*.

Scenario #4:

- A set m *identical* items.
- Each bidder i has a private *marginal valuation* $\mu_i(j)$ for a jth item. Thus, bidder i's total valuation for ℓ units is $v_i(\ell) := \sum_{j=1}^{\ell} \mu_i(j)$.
- Valuations are *downward-sloping*, meaning that $\mu_i(1) \geq \mu_i(2) \geq \cdots \geq \mu_i(m)$ for every i. Thus, additional units provide diminishing returns.

The first scenario — identical items and unit-demand valuations — corresponds to special case in which $\mu_i(j) = 0$ for all i and $j \geq 2$. The second and third scenarios — which have non-identical goods but restrict to additive or unit-demand valuations, respectively — are incomparable to this one.

Our goals are the usual ones — a simple ascending auction that is EPIC and in which sincere bidding leads to a welfare-maximizing allocation. Most aspects of the following solution are simpler than in the unit-demand case we just studied, but we'll see that non-unit-demand valuations do introduce one important complication.

†Department of Computer Science, Stanford University, 462 Gates Building, 353 Serra Mall, Stanford, CA 94305. Email: tim@cs.stanford.edu.

2 A DSIC Solution

We begin with our usual "sanity check" that there is a good DSIC solution — i.e., that the VCG mechanism can be implemented efficiently. In scenario #4, the VCG mechanism has the following form.

1. Collect bid $\nu_i(j)$ for all bidder i and quantities j, with $\nu_i(j)$ nonincreasing in j for each i, allegedly the bidders' marginal valuations.

2. Compute quantities $x_1, \ldots, x_n$ maximizing $\sum_{i=1}^n \sum_{j=1}^{x_i} \nu_i(j)$ subject to $\sum_{i=1}^n x_i \leq m$.

3. Charge each bidder its externality (more details shortly).

A simple but important observation is that the second, welfare-maximization step can be computed using a simple greedy algorithm. First, identify the set of the top m $\mu_i(j)$'s. Since the bids are downward-sloping, each bidder i will have a prefix of its first x_i reported marginal valuations in this set. A straightforward exchange argument shows that giving x_i items to each bidder i maximizes the welfare with respect to the reported marginal valuations.[1]

3 Walrasian Equilibria and VCG Payments

The biggest challenge posed by non-unit-demand valuations, even with identical items, is that the VCG outcome need not correspond to a Walrasian equilibrium. Recall that this equivalence was our most useful tool for designing an ascending auction in scenario #3.

In more detail, for multi-unit auctions, a *Walrasian equilibrium* for a valuation profile is a pair $(q, \mathbf{x})$ with the following properties:

1. $q \geq 0$ is a nonnegative price per unit;

2. $x_1, \ldots, x_n$ are nonnegative integers with sum at most m;

3. each bidder i gets its preferred quantity at the price q, namely $x_i \in \text{argmax}_{j=0}^m \{v_i(j) - q \cdot j\}$

4. $\sum_{i=1}^n x_i < m$ only if $q = 0$.

Example 3.1 (VCG $\neq$ WE) Consider two bidders and two identical items. The first bidder is additive, with $\mu_1(1) = \mu_1(2) = 2$. The second bidder is unit-demand, with $\mu_2(1) = 1$ and $\mu_2(2) = 0$.

[1] This implementation of the VCG mechanism runs in time polynomial in n and m, which is polynomial in the input size if each bidder's valuation is an explicit list of marginal valuations. This running time can be improved if we assume that each bidder can answer a "value query" — given j, return $\mu_i(j)$ — in sublinear time. By cleverly using binary search, the welfare-maximizing allocation can be computed using a number of value queries that is polynomial in n, $\log m$, and the number of bits used to describe bidders' valuations. See [2, §4.2] for the details.

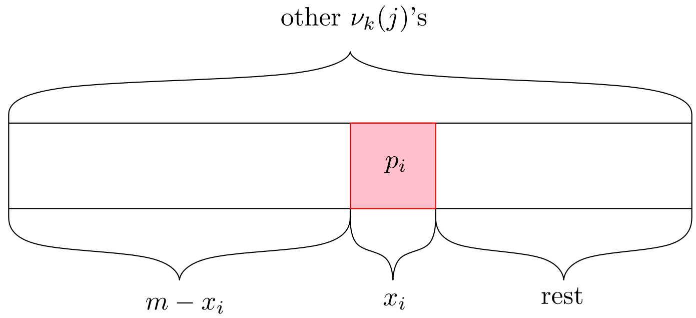

Figure 1: When the other bids are sorted, the x_i bids after the first $m - x_i$ bids form the price for player i.

In the truthful VCG outcome, the first bidder receives both items and pays the externality it creates, which is 1. This could only correspond to a WE with per-unit price $\frac{1}{2}$. But with $q = \frac{1}{2}$, the second bidder would prefer quantity 1 to its allocation of zero.

Example 3.1 notwithstanding, most of the nice properties of WE carry over from the third scenario to the present one. There always exists a WE; there is a smallest WE price; and payments under the VCG mechanism are at most those in the smallest WE (though equality need not hold). We leave the details to the Exercises.

In contrast to scenario #3, where we could attain the VCG outcome and a WE simultaneously, in this scenario we need to choose between two competing options.

1. Strive for an EPIC auction in which sincere bidding leads to a (non-VCG) WE outcome.

2. Strive for an EPIC auction in which sincere bidding leads to a (non-WE) VCG outcome.

We'll see later why the first option is impossible. For starters, we leave as an exercise a proof that the direct-revelation mechanism that charges the smallest WE price (with respect to the reported valuations) is not DSIC. A key issue is "demand reduction" — with a Walrasian price, bidders can have incentive to underbid, to obtain a smaller quantity at a much cheaper price (and thus larger utility).

The rest of this lecture gives a solution to the second option.

4 Characterization of VCG Payments

Analogous to the previous scenario, to design a good ascending auction we need to understand the special structure of the VCG payments in multi-unit auctions with downward-sloping valuations. Recall that, by definition, the payment of bidder i under the VCG mechanism is the maximum-possible welfare obtainable by the other $n-1$ minus the welfare those bidders obtain in the VCG outcome (with respect to the submitted bids). Because a greedy algorithm maximizes the welfare in scenario #4, these payments are relatively simple, and might plausibly be computed by an ascending auction.

For a bidder i, imagine sorting the $\nu_k(j)$'s submitted by other bidders from high to low (Figure 1). We can divide the $\nu_k(j)$'s into three regions: (i) the highest $m - x_i$ of them, where x_i is the number of units awarded to i in the VCG outcome; (ii) the next x_i highest $\nu_k(j)$'s; and (iii) the rest of the $\nu_k(j)$'s. The bids in the second category are precisely those that will not be allocated when i is present, but will be allocated when i is absent. That is, the VCG payment of bidder i is the sum of the x_i $\nu_k(j)$'s in category (ii). It is convenient to associate these x_i numbers with the x_i units that i wins in the VCG outcome as follows:

$$\text{VCG price } p_i(j) = (m - j + 1)\text{th largest } b_k(\ell) \text{ (with } k \neq i). \tag{1}$$

As one would expect in a DSIC mechanism, the right-hand side of (1) is a function only of the bids by bidders $k \neq i$. The following observations are immediate.

1. $p_i(1) \leq p_i(2) \leq \cdots \leq p_i(x_i)$;
2. by the downward-sloping restriction, $\nu_i(1) \geq \nu_i(2) \geq \cdots \geq \nu_i(x_i)$;
3. $$\underbrace{\nu_i(x_i)}_{\text{in the top } m} \geq \underbrace{p_i(x_i)}_{\text{not in the top } m};$$
4. $$\underbrace{\nu_i(x_i+1)}_{\text{not in the top } m} \leq \underbrace{p_i(x_i+1)}_{\text{in the top } m}.$$

These four inequalities imply that the VCG mechanism gives bidder i the quantity x_i that maximizes

$$\sum_{j=1}^{\ell} (\nu_i(j) - p_i(j)) \tag{2}$$

over all possibly quantities $\ell = 0, 1, 2, \ldots, m$. One way to think about this is that items are being doled out by a Pez dispenser at an increasing sequence of prices (which are independent of i's bids), and that the VCG mechanism keeps taking items on behalf of bidder i as long as they contribute positively to its utility.

While simple enough, the expression in (1) references all of the bids by others — information that is available in a direct-revelation mechanism but against the spirit of an ascending auction, where such values should be implicitly elicited only on a "need to know" basis. With this in mind, define the *demand* $D_i(q)$ of bidder i at the (per-unit) price q as the number of units that i wants at price q — also known as $\max\{j \leq m : v_i(j) > q\}$, or the length of the prefix of i's (downward-sloping) valuation that contains only number bigger than q. Since the $(m - j + 1)$th largest of the $\nu_k(j)$'s is also the threshold price at which the combined demand of bidders other than i passes from $m - j + 1$ to $m - j$, we can write

$$\text{VCG price } p_i(j) = \inf\{q : \sum_{k \neq i} D_k(q) \leq m - j\}. \tag{3}$$

The right-hand side of (3) is defined purely in terms of "demand queries," which have thus far been the method of communication between bidders and an ascending auction.

5 The Clinching Auction

The clinching auction [1] is an ascending-price auction, that stops once demand equals supply. To halt with the VCG payments (3), rather than the payments of some WE, units will effectively be sold (i.e., "clinched") as the auction proceeds, with the per-unit price increasing as the auction proceeds. Using the characterization of VCG payments in (3), we can compute the appropriate prices on the fly using demand queries.

Clinching Auction (Main Loop):

1. Initialize $p = 0$.
2. while (TRUE):
   (a) Ask each bidder i a demand query a collection the resulting $D_i(p)$'s.[2]
   (b) If $\sum_{i=1}^{n} D_i(p) \leq m$, then halt (see below for final allocation and payments).
   (c) Otherwise, increment p by ϵ.

Final allocation: Let p denote the price at termination. Give $x_i \in [D_i(p), D_i(p-\epsilon)]$ units to each bidder i, subject to $\sum_{i=1}^{n} x_i = m$. This is possible because $\sum_{i=1}^{n} D_i(p) \leq m < \sum_{i=1}^{n} D_i(p-\epsilon)$.[3]

Final payments: For $j = 1, 2, \ldots, m$, define

$$q_i(j) := -\epsilon + \min_{t \in \mathcal{Z}^+} \{\epsilon t \ : \ \sum_{k \neq i} D_k(\epsilon t) \leq m - j\}. \tag{4}$$

Bidder i pays $q_i(j)$ for its jth unit ($j = 1, 2, \ldots, x_i$). This is exactly the price during the first iteration in which $\sum_{k \neq i} D_k(p) \leq m$ — in this iteration, bidder i "clinches" its j unit at a price of $q_i(j)$.

Example 5.1 Recall Example 3.1, with $n = m = 2$, $\mu_1(1) = \mu_1(2) = 2$, $\mu_2(1) = 1$ and $\mu_2(2) = 0$. Observe that $D_1(q) = 2$ for $q \in [0, 2)$ while $D_2(q) = 1$ for $q \in [0, 1)$. In the clinching auction with sincere bidding, up to an ϵ, bidder 1 pays 0 for the first unit and 1 for the second unit. This matches the truthful VCG outcome.

6 Analysis of the Clinching Auction

We assume in this section that bidders bid sincerely; the next section addresses incentive issues.

---

[2] In contrast to the unit-demand case, the final payments are computed using the whole history of demand query answers, not merely the most recent answers. The fact that VCG payments are not generally a WE necessitates this complexity.

[3] In the trivial case where $\sum_{i=1}^{n} D_i(0) \leq m$, all the units can be given away to free to those who want them.

We first note that the clinching auction approximately maximizes welfare, and hence simulates the allocation of the VCG mechanism.

Lemma 6.1 *If all bidders bid sincerely, then the allocation computed by the clinching auction is within $m\epsilon$ of the maximum possible.*

Proof: Recall that, because valuations are downward-sloping, the maximum-possible welfare is the sum of the top m $\mu_i(j)$'s. The last two iterations of the clinching auction identify the set B_1 of $\mu_i(j)$'s that exceed p and the set B_2 of $\mu_i(j)$'s that exceed $p-\epsilon$, with $|B_1| \leq m < |B_2|$. The welfare of the clinching auction's allocation is the sum of all terms in $|B_1|$ plus the sum of $m-|B_1|$ terms from B_2. Since all terms of $B_2 \setminus B_1$ are within ϵ of each other, its allocation is at most $m\epsilon$ less than the maximum possible. ∎

We next show that the clinching auction also simulates the VCG payments, in the sense that all bidders' utilities are approximately the same in both mechanisms.

Theorem 6.2 ([1]) *The utility of every bidder in the sincere bidding outcome of the clinching auction is within $\pm m\epsilon$ of its utility in the truthful VCG outcome.*

Proof: Recall from (2) the utility of bidder i in the truthful VCG outcome:

$$\sum_{j=1}^{\ell}(\mu_i(j) - p_i(j)), \tag{5}$$

where $p_i(j)$ is defined in (3) and ℓ is the largest index for which $\mu_i(\ell) > p_i(\ell)$.

Assume that bidders bid sincerely in the clinching auction and let p^* denote the final price. The utility of bidder i is

$$\sum_{j=1}^{x_i}(\mu_i(j) - q_i(j)). \tag{6}$$

We first note that every term in (6) is nonnegative. If i wins at least j units in the clinching auction, then $j \leq D_i(p^* - \epsilon)$; since $q_i(j) \leq p^* - \epsilon$ (by (4)), $\mu_i(j) \geq p^* - \epsilon \geq q_i(j)$. Second, consider a term in (5) with value greater that ϵ — i.e., a unit j with $\mu_i(j) > p_i(j) + \epsilon$. Then, by the definition 3 of $p_i(j)$, there will be an iteration of the clinching auction with $p < \mu_i(j)$ (and hence $D_i(p) \geq j$) and $\sum_{k \neq i} D_k(p) \leq m - j$. This implies that bidder i receives at least j units when the auction terminates.

We've proved that (6) has only nonnegative terms and is only missing terms from (5) that contribute utility at most ϵ. Finally, comparing the definitions of $p_i(j)$ in (3) and $q_i(j)$ in (4) shows that $q_i(j) \leq p_i(j)$ for all i, j. It follows that i's utility (6) in the clinching auction is within $m\epsilon$ of that (5) in the VCG mechanism. ∎

7 Incentives in the Clinching Auction

We conclude this lecture by proving that the clinching auction is EPIC. As discussed last lecture, this does not follow automatically from Theorem 6.2 because of possible vulnerabilities to deviations to inconsistent actions. In the clinching auction, however, it is easy to show that such deviations are of no use.

Theorem 7.1 ([1]) *The clinching auction is EPIC (up to $m\epsilon$).*

Proof: Fix a valuation profile $\mathbf{v}$ and a bidder i. Assume that all bidders other than i bid sincerely. This means, in particular, that other bidders' responses to demand queries are independent of bidder i's action. Since the prices $q_i(j)$ paid by bidder i are a function only of the demand query responses of other bidders, these prices are independent of bidder i's action. The proof of Theorem 6.2 shows that sincere bidding maximizes i's utility (6) up to $m\epsilon$; thus, an arbitrary action can only improve i's utility by this amount. ■

References

[1] L. M. Ausubel. An efficient ascending-bid auction for multiple objects. *American Economic Review*, 94(5):1452–1475, 2004.

[2] N. Nisan. Algorithmic mechanism design: Through the lens of multi-unit auctions. In *Handbook of Game Theory*. North-Holland, 2014. Forthcoming.