CS364A: Algorithmic Game Theory
Lecture #7: Multi-Parameter Mechanism Design and the VCG Mechanism*

Tim Roughgarden†

October 14, 2013

1 General Mechanism Design Problems

Previous lectures only considered single-parameter mechanism design problems, where each participant has just one piece of private information, its valuation per unit of stuff. In many problems, a participant has different private valuations for different items. Once we are unsure about whether a participant prefers item A to item B, for example, we are in the realm of *multi-parameter* mechanism design.

Here are the ingredients of a general, multi-parameter mechanism design problem:

- n strategic participants, or "agents;"
- a finite set Ω of outcomes;
- each agent i has a private valuation $v_i(\omega)$ for each outcome $\omega \in \Omega$.

The outcome set Ω is abstract and could be very large. In a single-item auction, Ω has only $n + 1$ elements, corresponding to the winner of the item (if any). In the standard single-parameter model of a single-item auction, we assume that the valuation of an agent is 0 in all of the n outcomes in which it doesn't win, leaving only one unknown parameter per agent. In the more general multi-parameter framework above, an agent can have a different valuation for each possible winner of the auction. This example is not without plausible application: in a bidding war over a hot startup, for example, agent i's highest valuation might be for acquiring the startup, but if it loses it prefers that the startup be bought by a company in a different market, rather than by a direct competitor.

†Department of Computer Science, Stanford University, 462 Gates Building, 353 Serra Mall, Stanford, CA 94305. Email: tim@cs.stanford.edu. These lecture notes are provided for personal use only. See my book *Twenty Lectures on Algorithmic Game Theory*, published by Cambridge University Press, for the latest version.

2 The VCG Mechanism

Our next result is a cornerstone of mechanism design theory.

Theorem 2.1 (The Vickrey-Clarke-Groves (VCG) Mechanism [6, 1, 3]) *In every general mechanism design environment, there is a DSIC welfare-maximizing mechanism.*

Recall the three properties of "awesome" auctions from Lecture 2. Theorem 2.1 asserts the first two properties but not the third (polynomial running time). We already know that, even in single-parameter environments, we can't always have the second and third properties (unless $P = NP$). As we'll see, the VCG mechanism is highly non-awesome in many important applications.

As always, designing a (direct-revelation) DSIC mechanism is tricky because the allocation and payment rules need to be coupled carefully.[1] We apply the same two-step approach that served us so well in single-parameter environments.

The first step is to assume, without justification, that agents truthfully reveal their private information, and then figure out which outcome to pick. Since Theorem 2.1 demands welfare-maximization, the only solution is to pick the welfare-maximizing outcome, using bids as proxies for the true (and unknown) valuations. That is, given bids $\mathbf{b}_1, \ldots, \mathbf{b}_n$, where each $\mathbf{b}_i$ is indexed by Ω, we define the allocation rule $\mathbf{x}$ by

$$\mathbf{x}(\mathbf{b}) = \operatorname*{argmax}_{\omega \in \Omega} \sum_{i=1}^{n} b_i(\omega). \tag{1}$$

The second step is to define a payment rule that, when coupled with the above allocation rule, yields a DSIC mechanism. Last time we arrived at this point, for single-parameter environments, we formulated and proved Myerson's Lemma, which is a general solution to this second step for all such environments. Recall that Myerson's Lemma asserts that allocation rule monotonicity is necessary and sufficient for implementability and, for monotone rules, it gives an explicit formula for the unique payment rule that meets the DSIC condition. Myerson's Lemma does not hold beyond single-parameter environments — with an agent submitting bids in more than one dimension, it's not even clear how to define "monotonicity" of an allocation rule.[2] Similarly, the "critical bid" characterization of DSIC payments (for 0-1 problems) does not have an obvious analog in multi-parameter problems.

The key idea is to make use of an alternative characterization of DSIC payments for the welfare-maximizing allocation rule (proved in the exercises), as the "externality" caused by an agent i — the welfare loss inflicted on the other $n-1$ agents by i's presence. For example,

[1] The proof of the Revelation Principle in Lecture 4 holds without change in multi-parameter environments, so restricting to direct-revelation mechanisms is without loss of generality.

[2] There is an analogous characterization of implementable multi-parameter allocation rules in terms of "cycle monotonicity." This is an elegant result, analogous to the fact that a network admits well-defined shortest paths if and only if it possesses no negative cycle. Cycle monotonicity is far more unwieldy than single-parameter monotonicity, however. Because it is so brutal to verify, cycle monotonicity is rarely used to argue implementability or to derive DSIC payment rules in concrete settings.

in a single-item auction, the winning bidder inflicts a welfare loss of the second-highest bid to the others (assuming truthful bids), and this is precisely the Vickrey auction's payment rule. This idea of "charging an agent its externality" makes perfect sense in general mechanism design environments, and it corresponds to the payment rule

$$p_i(\mathbf{b}) = \underbrace{\max_{\omega \in \Omega} \sum_{j \neq i} b_j(\omega)}_{\text{without } i} - \underbrace{\sum_{j \neq i} b_j(\omega^*)}_{\text{with } i}, \tag{2}$$

where $\omega^* = \mathbf{x}(\mathbf{b})$ is the outcome chosen in (1). Note that $p_i(\mathbf{b})$ is always nonnegative (exercise).

We claim that this mechanism $(\mathbf{x}, \mathbf{p})$, the *VCG mechanism*, is DSIC. (By definition, it maximizes welfare assuming truthful bids.) For the first time since the Vickrey auction, we'll prove the DSIC property from scratch (i.e., without use of Myerson's Lemma). Recall this means that for every agent i and every set $\mathbf{b}_{-i}$ of bids by the other agents, agent i maximizes its quasilinear utility $v_i(\mathbf{x}(\mathbf{b})) - p_i(\mathbf{b})$ by setting $\mathbf{b}_i = \mathbf{v}_i$.

Fix i and $\mathbf{b}_{-i}$. When the chosen outcome $\mathbf{x}(\mathbf{b})$ is ω^*, i's utility is

$$v_i(\omega^*) - p_i(\mathbf{b}) = \underbrace{\left[v_i(\omega^*) + \sum_{j \neq i} b_j(\omega^*) \right]}_{\text{(A)}} - \underbrace{\left[\max_{\omega \in \Omega} \sum_{j \neq i} b_j(\omega) \right]}_{\text{(B)}}. \tag{3}$$

Observe that the term (B) is a constant, independent of what $\mathbf{b}_i$ is. Thus, the problem of maximizing agent i's payoff reduces to the problem of maximizing the first term (A). As a thought experiment, let's suppose agent i has the power to choose the outcome ω^* directly, rather than merely influencing the chosen outcome indirectly via its choice of bid $\mathbf{b}_i$. Agent i would, of course, use this extra power to choose an outcome that maximizes the term (A). If agent i sets $\mathbf{b}_i = \mathbf{v}_i$, then term (1) that the mechanism maximizes becomes identical to the term (A) that the agent wants maximized. Thus, bidding truthfully results in the mechanism choosing an outcome that maximizes agent i's utility; no other bid could be better. This completes the proof of Theorem 2.1.

Here is an alternative interpretation of the payments in the VCG mechanism. Rewrite the expression in (2) as

$$p_i(\mathbf{b}) = \underbrace{b_i(\omega^*)}_{\text{bid}} - \underbrace{\left[\sum_{j=1}^{n} b_j(\omega^*) - \max_{\omega \in \Omega} \sum_{j \neq i} b_j(\omega) \right]}_{\text{rebate}}. \tag{4}$$

We can thus think of agent i's payment as its bid minus a "rebate," equal to the increase in welfare attributable to i's presence. For example, in the Vickrey auction, the highest bidder pays its bid b_1 minus a rebate of $b_1 - b_2$ (where b_2 is the second-highest bid), the increase in welfare that the bidder brings to the table.

We leave it as an exercise to observe that the discount in (4) is always nonnegative, implying that $p_i(\omega^*) \leq b_i(\omega^*)$ and hence truthtelling agents are guaranteed nonnegative utility.

The upshot of the VCG mechanism is that, in general multi-parameter environments, DSIC welfare-maximization is always possible in principle. While it can be infeasible to implement in practice, the VCG mechanism nevertheless serves as a useful benchmark for other, more practical approaches.

3 Combinatorial Auctions

Combinatorial auctions are important in practice. Already in the domain of government spectrum auctions, dozens of such auctions have raised hundreds of billions of dollars of revenue. They have also been used for other applications such as allocating take-off and landing slots at airports. Combinatorial auctions are also notoriously difficult, in both theory and practice. Theoretical work has identified many impossibility results for what can be done with reasonable communication and computation. Practice has provided examples of badly designed combinatorial auctions with serious consequences, such as the 1990 New Zealand spectrum auction that raised merely \$36 million, a far cry from consultants' estimates of \$250 million (see [5, Chapter 1] for details).

3.1 The Model

A combinatorial auction has n bidders — for example, Verizon, AT & T, and several regional providers. There is a set M of m items, which are *not* identical — for example, a license awarding the right to broadcast on a certain frequency in a given geographic area. The outcome set Ω corresponds to n-vectors $(S_1, \ldots, S_n)$, with S_i denoting the set of items allocated to bidder i (its "bundle"), and with no item allocated twice. There are $(n+1)^m$ different outcomes. Each bidder i has a private valuation $v_i(S)$ for each bundle $S \subseteq M$ of items it might get. Thus, each bidder has 2^m private parameters. One generally assumes that $v_i(\emptyset) = 0$ and that $v_i(S) \leq v_i(T)$ whenever $S \subseteq T$ (i.e., "free disposal"). We'll discuss additional assumptions on valuations later.

The welfare of an outcome $(S_1, \ldots, S_n)$ is $\sum_{i=1}^{n} v_i(S_i)$. In principle, the VCG mechanism provides a DSIC solution for maximizing the welfare. This can be useful when bidders' valuations are sufficiently structured, as with "unit-demand" bidders (see the exercises). In general, however, there are several major impediments to implementing the VCG mechanism.

3.2 Challenges

The first major challenge of combinatorial auctions is that of *preference elicitation*. Each bidder has $2^m - 1$ private parameters, roughly a thousand when $m = 10$ and a million when $m = 20$. No bidder in their right mind would want to write down (or even figure out) that many bids. No seller would want to listen to that many bids. This exponential number of private parameters makes the VCG mechanism, and every other direct-revelation mechanism, a nonstarter for combinatorial auctions in practice. Note that this challenge never arises in single-parameter mechanism design, where each bidder only has to communicate one number.

The utter absurdity of direct-revelation combinatorial auctions motivates *indirect* mechanisms, which learn information about bidders' preferences only on a "need-to-know" basis. We have not yet discussed any such mechanisms. The canonical indirect auction is the ascending English auction. You're familiar with this auction format from the movies — an auctioneer keeps track of the current price and tentative winner, and the auction stops when only one interested bidder remains.[3] Each bidder has a dominant strategy, which is to stay in the auction as long as the current price is below its valuation (the player might win for positive utility) and to drop out once the current price reaches its valuation (after which winning can only lead to negative utility). Assuming all bidders play these strategies, the outcome of the English ascending auction is the same as that of the Vickrey (sealed-bid) auction. The Vickrey auction is what you get when you apply the Revelation Principle (Lecture 4) to the English auction.

Indirect auctions are unavoidable for all but the smallest combinatorial auctions. We'll discuss the auction formats used in practice for wireless spectrum auctions next lecture.[4] An important question is: can indirect mechanisms achieve, at least in principle, non-trivial welfare guarantees while eliciting only a small amount of information (say, polynomial in n and m) from the bidders? There is nice theoretical work on this question, and the answer, roughly, is "no" for complex valuations and "yes" for sufficiently simple valuations. We'll discuss this more in the next lecture, and at length in CS364B. In practice, one hopes that bidders' valuations are sufficiently simple and/or that auction performance will be much better than the worst-case bounds.

The second challenge in designing practice combinatorial auctions is familiar from our discussion of algorithmic mechanism design (Lecture 4). Even when the first challenge is not an issue — for example, when bidders are single-parameter and direct revelation is practical — welfare-maximization can be an intractable problem. We encountered this issue with Knapsack auctions and in a couple of the Problems. This challenge is fundamental, and cannot be avoided by choosing a clever auction format. In practice, one hopes that combinatorial auctions compute allocations that are reasonably close to welfare-maximizing. This is impossible to check directly, since bidders' valuations are unknown and the combinatorial

[3]There are a few variants. The movies, and auction houses like Christie's and Soethby's, use an "open outcry" auction in which bidders can drop out and return, and can make "jump bids" to aggressively raise the current price. When doing mathematical analysis, the "Japanese" variant is usually more convenient: the auction begins at some opening price, which is publicly displayed and increases at a steady rate. Each bidder either chooses "in" or "out," and once a bidder drops out it cannot return. The winner is the last bidder in, and the sale price is the price at which the second-to-last bidder dropped out.

[4]Indirect auctions can also be useful in single-parameter settings like single-item auctions. Empirical studies show that bidders are more likely to play their dominant strategy in an English auction than in a sealed-bid second-price auction, where some bidders inexplicably overbid [4]. Second, ascending auctions leak less valuation information to the auctioneer. In a Vickrey auction, the auctioneer learns the highest bid; in an English auction, the auctioneer only learns a lower bound on the highest bid (the second-highest bid).

auctions used in practice are not DSIC and offer some opportunities for strategizing. Nevertheless, there are various simple "sanity checks" that can be applied to an auction outcome to suggest good welfare maximization. For example, did bidders successfully acquire sensible packages (e.g., spectrum licenses that are adjacent geographically or in the spectrum)? Did similar items sell for similar prices?

The third challenge applies to the VCG mechanism — which turns out to be essentially the unique DSIC welfare-maximizing mechanism — even when the first two challenges are not relevant (e.g., a single-parameter problem small enough that welfare maximization can be done in a reasonable amount of time). Namely, the VCG mechanism can have bad revenue and incentive properties, despite being DSIC.

For instance, suppose there are two bidders and two items, A and B. The first bidder only wants both items, so $v_1(AB) = 1$ and is 0 otherwise. The second bidder only wants item A, so $v_2(AB) = v_2(A) = 1$ and is 0 otherwise. The revenue of the VCG mechanism is 1 in this example (exercise). But now suppose we add a third bidder who only wants item B, so $v_3(AB) = v_3(B) = 1$. The maximum welfare has jumped to 2, but the VCG revenue has dropped to 0 (exercise)! The fact that the VCG mechanism has zero revenue in seemingly competitive environments is a dealbreaker in practice. The revenue non-monotonicity in this example also implies numerous incentive problems, including vulnerability to collusion and false-name bids (see the exercises). None of these issues plague the single-item Vickrey auction.

The first challenge begets a fourth. Almost all combinatorial auctions used in practice are iterative, comprising a sequence of rounds; we'll discuss details next lecture. Iterative auctions offer new opportunities for strategic behavior. For example, Cramton and Schwatz [2] found that, in an early and relatively uncompetitive spectrum auction, bidders used the low-order digits of their bids to effectively send messages to other bidders. Let's consider license #378 in that auction, for spectrum use rights in Rochester, MN. USWest and McLeod were battling it out for this license, with each repeatedly outbidding the other. Apparently, USWest tired of this bidding war and switched to a retaliatory strategy — bidding on a number of licenses in other geographical areas on which McLeod was the standing high bidder, and on which USWest had shown no interest in previous rounds. McLeod ultimately won back all of these licenses, but had to pay a higher price due to USWest's bids. To make sure its message came through loud and clear, all of USWest's retaliatory bids were a multiple of 1000 *plus 378* — presumably warning McLeod to get the hell out of the market for Rochester, or else. This particular type of signalling can be largely eliminated by forcing all bids to be multiples of a suitably large number, but other opportunities for undesirable strategic behavior remain.

References

[1] E. H. Clarke. Multipart pricing of public goods. *Public Choice*, 11(1):17–33, 1971.

[2] P. Cramton and J. Schwartz. Collusive bidding: Lessons from the fcc spectrum auctions. *Journal of Regulatory Economics*, 17(3):229–252, 2000.

[3] T. Groves. Incentives in teams. *Econometrica*, 41(4):617–631, 1973.

[4] R. M. Harstad. Dominant strategy adoption and bidders' experience with pricing rules. *Experimental Economics*, 3:261–280, 2000.

[5] P.. Milgrom. *Putting Auction Theory to Work*. Cambridge, 2004.

[6] W. Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, 16(1):8–37, 1961.