Algorithms, Games, and Networks February 26, 2013

Lecture 1

Lecturer: Avrim Blum *Scribe: Jamie Morgenstern*

1 Mechanism Design, Chapters 9.3, 9.5

Suppose the department has a spare printer, and would like to give it away to the person who values it the most. Specifically, suppose there are n people, and each person has some private value $v_i \in \mathbb{R}^+$ for receiving the printer. We'd like to give the printer to $argmax_i v_i$. Moreover, we're going to allow a system with payments, that is to say each of the n persons will be told some p_i that they'll pay as the outcome of this mechanism. Further, we assume that people's utilities are *linear in money*, that is, the person who gets the printer will have value $v_i - p_i$ for getting the printer and paying p_i.

One potential mechanism is to just ask everyone to report their v_i's and give it away to the person with the highest v_i (that is, $p_i = 0$ for all i). Unfortunately, Avrim might say that he liked the printer enough to pay $\$1,000,000$, even if he only liked it enough to pay \$10, because he could win with that bid and wouldn't have to pay anything.

A second attempt would be to try to charge the highest bid to the highest bidder. In this case, there is a different problem. The highest bidder i would rather bid something just very slightly more than the second-highest bidder, so they can pay less than v_i and have positive utility for the outcome of the "auction". In particular, bidding v_i is a bit silly since then you have guaranteed your utility will be 0 no matter what.

Finally, there are two very related mechanisms that don't have either of these problems. First, suppose we have an in-person auction, where the bidding starts at \$1, where the bidders keep their hands raised as the price increases infintesimally, and a bidder lowers her hand once she wouldn't pay the price for the printer. Second, consider a second-price, or Vickrey auction, where everyone submits *sealed* bids, and the printer goes to the i with the highest bid, at the second-highest price. To give some notation, give the printer to $i = argmax_j v_j$ and charge price $p = max(v_{-i})$ (where v_{-i} denotes all the bids except the bid of player i).

In a Vickrey auction, it is a *dominant strategy* to give your true value as a bid. That is, regardless of what every other bidder does, any particular bidder is at least as well-off with the result if they tell the truth.

Claim 1 *The Vickrey auction is dominant-strategy truthful, or incentive compatible. For* *any i, any v_{-i}, $u_i(Vic(v_i, v_{-i})) \geq u_i(Vic(v_i', v_{-i}))$, where $Vic(\vec{v})$ is the outcome of running the Vickrey auction on vector of bids $\vec{v}$.*

There are several proofs of this claim. Here's one.

Proof: Consider bidder i and let p be the maximum bid out of all the *other* bidders.

If $v_i > p$, there would be several ways to deviate. If you announce any value more than p, including v_i, your utility will be $v_i - p$. If you announce any value less than p, you will get utility 0. So, you're at least as well off saying v_i.

If $v_i = p$, there is zero utility whether or not you win the auction. So you may as well tell the truth.

If $v_i < p$, if you bid anything less than p, you don't win and get utility 0, and if you bid more than p, you win but have to pay $p > v_i$ and have negative utility. ■

You can also think of this as the same (for bidder i) as a take-it-or-leave-it auction with price p. You can choose to pay p or not, and only should when $v_i > p$. Or, alternatively, your price does not depend on your bid. So, you can't manipulate your set of options (winning with a price which doesn't depend on your bid, or losing) by giving a different bid.

The social welfare of a 1-item auction is the value for the good of the person who gets the good.

Claim 2 *Vickrey auctions maximize social welfare: that is, the winner of the auction is the person who values the item the most.*

Maximizing social welfare is a somewhat intuitive notion. It makes sense that we want to get our good to whoever wants it the most. But why do we care about incentive compatibility? An introspective reason might be that these mechanisms are easier to analyze (if people bid truthfully, it's not so hard to prove theorems about a mechanism). If we are slightly less navel-gazing, from the auctioneer's perspective there is more predictability in how the mechanism behaves. Even more, from the bidder's perspective, playing the game has now become easy. They can do less strategizing and just tell the truth.

Now, suppose the department has two identical printers. Which of the following mechanisms is incentive compatible?

1. Everyone submits a bid v_i, and give to the highest bidder at the second-highest bid. Give to the second-highest bidder at the third-highest bid.
2. Give to the top 2 bidders at the third-highest bid.

The first one isn't incentive compatible. It's better to be the second-highest bidder, because you get to pay less than being the highest bidder for the same good. The second one is incentive compatible: there are exactly 2 bidders who will pay the third-highest price (or the highest price leaving those two bids out in the computation of "highest").

What if the two printers are not idential? Suppose one printer is much nicer, and the other one is an old, janky printer that often jams. So long as every bidders' utility is still linear in money, we can handle just such a case and much more general cases. We can also handle complements and substitutes (you might want both a right and left shoe, but only one printer and not both).

2 Mechanism design, general setup

We assume we have n players, and a set of "alternatives" A (these are also called outcomes or allocations). Player i has valuation function $v_i : A \to R$ (you can dislike allocations that give printers to your mortal enemies more than those that give printers to your friends). These can be arbitrary. The one assumption we will make is that players' utilities are *quasilinear*: The utility for player i of alternative a, paying p_i is $u_i = v_i(a) - p_i$.

Definition 3 *A* direct revelation mechanism *(f, p), $f : V \to A$ takes in a sequence v of valuation functions $v = (v_1, v_2, \ldots, v_n)$, selects some alternative a and some payments $p = (p_1, p_2, \ldots, p_n)$. Specifically, $f(v) = a$ and $p(v)$ is the vector of payments with player i making a payment of $p_i(v)$. A* truthful mechanism *has the property that for all v_i, v_i', v_{-i} , $v_i(f(v)) - p_i(v) \geq v_i(f(v_i', v_{-i})) - p_i(v_i', v_{-i})$.*

Definition 4 *The* social welfare *of some outcome a is $\sum_j v_j(a)$. A mechanism is* efficient *if it maximizes social welfare.*

Think of the pricing as being irrelevant to social welfare: if the auctioneer is part of the social welfare computation, any payment made to the auctioneer has net social welfare change of 0.

There is a direct revelation mechanism, which is incentive compatible and maximizes social welfare, for any setting where people's utilities are quasilinear in money. The maximium social welfare is $\max_{a \in A} \sum_j v_j(a)$. The key ideas is that we should set the payments so that, when individuals maximize their personal utility, that coincides with them maximizing our social welfare. The VCG mechanism does just that.

3 VCG, version 1

Given a vector v of valuation functions,

- Let $f(v) = argmax_{a \in A} \sum_j v_j(a)$

- Let $p_i = -\sum_{j \neq i} v_j(f(v))$ (e.g., we're going to pay each person the value everyone else gets from this outcome).

What is the utility of player i if i truthfully reports v_i?

$$v_i(f(v)) + \sum_{j \neq i} v_j(f(v)) = \sum_j v_j(f(v))$$

or the social welfare! Notice that we chose the $f(v)$ that maximizes social welfare. If player i instead reports some v_i', the utility is

$$v_i(f(v_i', v_{-i})) + \sum_{j \neq i} v_j(f(v_i', v_{-i})) = \sum_j v_j(f(v_i', v_{-i}))$$

This quantity can't be more than the utility in the previous case; it was the allocation that maximized the social welfare. Now, each player has the same incentives we have because their happiness is equal to the social welfare, which is what we are maximizing. So, this version of VCG is incentive compatible. In the case of the one printer setting, VCG # 1 gives the printer to the person who likes it best, and the amount he values it at is given in money to everyone who didn't win. This is somewhat unnatural, though, beause the department would have to give away a bunch of money just to give away a printer.

If, instead, we charge each player some additional amount that has nothing to do with what they bid, whether or not they win, that doesn't change the incentive compatibility of the mechanism. This gives us some idea for VCG # 2:

4 VCG, version 2

Let h_i be any function over v_{-i}. Given a vector v of valuation functions,

- Let $f(v) = argmax_{a \in A} \sum_j v_j(a)$
- Let $p_i = h_i(v_{-i}) - \sum_{j \neq i} v_j(f(v))$

This is still incentive compatible, for the reason mentioned above (the change in payment to bidders has nothing to do with their bid). We would like two other properties:

- The auctioneer never pays the bidder
- Assuming the v_i's themselves are non-negative, no player gets negative utility (in the case of one item, this is equivalent to saying no one who doesn't win has to pay money). This is called *ex-post individual rationality.*

VCG # 2 with appropriate payments can achieve these simultaneously. We instantiate h_i to get these properties with the Clarke pivot rule. Let

$$h_i(v_{-i}) = \max_a \sum_{j \neq i} v_j(a) = \max_a v_{-i}(a)$$

Then, the payments are just

$$p_i(v) = \max_a v_{-i}(a) - v_{-i}(f(v)) = \sum_{j \neq i} v_j(a) + \sum_{j \neq i} v_j(f(v))$$

The first part of the sum is at most as much as the second part (the second part of the sum maximizes the term), so we're not giving money away.

Why does no one get negative utility for her payment and allocation? We're charging everyone else their externality. By participating, you're making everyone else a different amount of happy, and you have to pay the difference in everyone else's happiness. This is non-negative:

$$v_i(f(v)) + v_{-i}(f(v)) - max_a v_{-i}(a)$$

$$= max_a v(a) - \max_a v_{-i}(a) \geq 0$$

This last inequality comes from the fact that no player can hurt the social welfare: if so, one can simply ignore that bidder and allocate to the other players as if that player weren't there.