Algorithms, Games, and Networks February 28, 2013

Lecture 14

Lecturer: Avrim Blum *Scribe: Hsu-Chun Hsiao*

1 Overview

In this lecture we review the standard version of VCG using the Clark pivot rule, discuss its applications and characteristics, and finally get a little bit into combinatorial auctions.

2 VCG Standard Version (Clark Pivot Rule)

In the previous lecture, we saw three versions of VCG. Today we will continue our discussion on the final version that uses the Clark pivot rule.

Recall the general setup of VCG. The setup considers n players and a set A of alternatives (or outcomes). Player i has a valuation function v_i, and v represents the vector of valuation functions of players. We denote by $p_i(v)$ the payment of player i and by $f(v)$ the allocation under valuation v. Although we mainly talked about positive values, e.g., the prices of a printer, the setting flexibly allows both positive and negative values. Given a vector of valuation functions v, what a VCG mechanism does is (1) compute the allocation and (2) charge people money.

Specifically, given a vector of valuation functions v, a VCG mechanism with Clark pivot payments determines $f(v)$ and $p_i(v)$ such that

- $f(v) = \arg\max_{a \in A} \sum_j v_j(a)$,
- $p_i(v) = \max_a \sum_{j \neq i} v_j(a) - \sum_{j \neq i} v_j(f(v))$.

The first equation computes the optimal allocation maximizing the social welfare; the second equation computes how much each player should pay to the mechanism. The Clark pivot rule essentially sets the first term in $p_i(v)$ to be $\max_a \sum_{j \neq i} v_j(a)$, or the social welfare for everybody else if player i were not there. Thus $p_i(v)$ is equal to *(maximum social welfare if i were absent)* - *(the social welfare of others when i is present)*, which is exactly player i's externality.

Therefore, player i's utility ($u_i(v)$) when he or she reports truthfully can be derived as

$$u_i(v) = v_i(f(v)) - p_i(v) \tag{1}$$
$$= \sum_j v_i(f(v)) - \max_a \sum_{j \neq i} v_j(a) \tag{2}$$
$$= \max_a \sum_j v_j(a) - \max_a \sum_{j \neq i} v_j(a) \tag{3}$$

2.1 Properties

The above design offers several nice properties:

1. **Incentive-compatible (IC).** This implies that people will report their true valuations and do not need to guess how much other players might want to bid. To see why, consider player i misreports and unilaterally changes the valuation from v_i to v_i', then the utility becomes $u_i(v_i', v_{-i}) = \sum_j v_j(f(v_i', v_{-i})) - \max_a \sum_{j \neq i} v_j(a)$, which is smaller than $u_i(v)$ because the first part in $u_i(v_i', v_{-i})$ (i.e., $\sum_j v_j(f(v_i', v_{-i}))$) is the social welfare of one allocation whereas the first part in $u_i(v)$ (i.e., $\max_a \sum_j v_j(a)$) is the social welfare of the *optimal* allocation. Intuitively, we can also argue that the above mechanism is incentive-compatible because the utility of player i does not depend on his or her own reported valuation.

2. **Maximize social welfare.** This is guaranteed by design as $f(v)$ is assigned the alternative maximizing the social welfare.

3. **No payment to bidders.** This is also guaranteed by design because the Clark pivot rule ensures $p_i(v) \geq 0$ for all i. The see why $p_i(v)$ never goes negative, note that $\max_a \sum_{j \neq i} v_j(a) \geq \sum_{j \neq i} v_j(f(v))$ because the left hand side is the optimal social welfare without i and the right hand side is the social welfare of a certain allocation without i.

4. **Individual rationality (IR) if v_js are non-negative.** This means no negative utility. Utilities never go below zero because in Equation 3, for all i $\max_a \sum_j v_j(a) \geq \max_a \sum_{j \neq i} v_j(a)$, if $v_j \geq 0$ for all j. We can also interpret this as adding one more player will never decrease the social welfare.

2.2 Examples

Let's instantiate VCG in the context of two interesting examples: *bilateral trade* and *public project*.

Bilateral Trade

Consider a car trading scenario with one seller S and one buyer B where v_s and v_b are how much S and B value the car, respectively. There are two alternatives: i.e., A ={seller has car (no trade), buyer has car (trade)}. Our goal is to design a VCG mechanism to help the trade happen.

One tricky thing about this setting is the seller begins with the car. So if we want to equate individual rationality (people participate willingly) with the notion of having nonzero utility, we should give the seller a valuation of 0 on the "no-trade" alternative and a valuation of $-v_s$ on the "trade" alternative. In particular, since there are now negative valuations, we will not want to use the Clarke pivot since that always uses non-negative payments and so will violate individual rationality. In fact, let's figure out what the right pivot rule is, if we enforce that no payments are made under the "no-trade" alternative. In fact, this will turn out to correspond to "VCG version 1" where the mechanism ensures that all players get utility equal to the social welfare.

Specifically, to maximize the social welfare (i.e., achieve social efficiency), the car goes to the buyer if $v_b > v_s$; and the seller keeps the car if $v_s \geq v_b$. So what does VCG do in this case? It first asks the two players for valuations, and then it determines their payments.

Now let's find out how much each player should pay when no payments are made in the case of no trade (i.e., when $v_s \geq v_b$). In other words, there exist a VCG mechanism, h_s and h_b such that

$$p_b(v) = h_b(v_s) - v_s(\text{no-trade}) = 0, \tag{4}$$
$$p_s(v) = h_s(v_b) - v_b(\text{no-trade}) = 0. \tag{5}$$

That is, $h_b(v_s) = v_s(\text{no-trade})$ and $h_s(v_b) = v_b(\text{no-trade})$. Therefore, the payments when the trade happens are

$$p_b(v) = h_b(v_s) - v_s(\text{trade}) = v_s(\text{no-trade}) - v_s(\text{trade}) = v_s, \tag{6}$$
$$p_s(v) = h_s(v_b) - v_b(\text{trade}) = v_b(\text{no-trade}) - v_b(\text{trade}) = -v_b. \tag{7}$$

Because when trade happens, S's valuation decreases by v_s (as the seller gives out the car) and B's valuation increases by v_b (as the buyer gets the car).

In sum, we have designed an incentive-compatible mechanism in which B pays v_s and S collects v_b if $v_b > v_s$, and no payments are made if $v_s \geq v_b$. However, the mechanism has to subsidize the cost. It can be shown that it is impossible to get all the desired properties (as described in Section 2.1) in bilateral trade.

Public Project

The department is considering getting a 3-D printer that costs \$$C$, and each potential user i has a valuation v_i on the printer.

To maximize the social welfare, the department will only get the printer if $\sum_j v_j \geq C$. (Formally we should view the department as an extra player whose valuation on obtaining the printer is $-C$.) But how much should the department charge each user in order to get the true valuation? Clearly, without any payment, users have incentives to overstate their values; whereas with payment linear to each user's valuation, for example, users have incentives to understate their values.

Notice that with the Clarke pivot rule, if the presence of player i does not change the alternative chosen, then $p_i(v) = 0$ since the two terms in the payment definition cancel out. Formally applying VCG with the Clark pivot rule, we obtain

$$p_i(v) = 0 \quad , \text{if } \sum_{j \neq i} v_j \geq C \text{ or } \sum_j v_j < C, \tag{8}$$

$$p_i(v) = C - \sum_{j \neq i} v_j \quad , \text{if } \sum_{j \neq i} v_j < C \text{ and } \sum_j v_j \geq C. \tag{9}$$

Equivalently, $p_i(v) = \max\{0, C - \sum_{j \neq i} v_j\}$. In other words, if player i is not pivotal (whether player i is present or not does not affect the result), player i has no incentives to pay any money. But if player i is pivotal, he or she will be willing to pay up to his or her own valuation of the printer. Also, nobody is charged if the printer is not purchased.

However, note that the payments collected by the department may be less than C. For example, if everybody's value is more than $\frac{C}{n-1}$ where n is the number of users, then the department gets the printer while nobody has to pay.

Again, it is impossible to maximize social welfare without the department subsidizing the cost.

3 Characterizations of Incentive Compatible Mechanisms

3.1 The Revelation Principle

Claim 1 *Any incentive-compatible non-direct revelation can be converted to incentive-compatible direct revelation.*

We are not going to formally prove this. Roughly let's think of "non-direct" as multi-round. The intuition is that if a non-direct mechanism is IC, then it means that there is a well-defined way to play the mechanism and the mechanism does not depend on other players' behaviors. Hence, we can have the mechanism executes the protocol for us.

This result allows us to focus on a simpler class (i.e., direct revelation than non-direct revelation) when we try to figure out what a mechanism is doing. Note that sometimes we may want to employ non-direct revelation, as direct revelation may be computational hard in some cases.

3.2 Direct Characterization

We saw in the examples of bilateral trade and public project that the mechanism cannot be incentive compatible without subsidizing the cost. It would be nice if we can characterize incentive-compatible VCG mechanisms such that we know the conditions under which incentive-compatible is achievable.

Theorem 2 *A direct revelation mechanism is IC iff for all v_i and v_{-i}:*

1. *If $f(v_i, v_{-i}) = f(v'_i, v_{-i})$, then $p_i(v_i, v_{-i}) = p_i(v'_i, v_{-i})$, and*
2. *$f(v_i, v_{-i}) = \arg\max_{a \in f(\cdot, v_{-i})}[v_i(a) - p_i(a, v_{-i})]$.*

The first condition basically says the payment p_i only depends on the output of $f(v_i, v_{-i})$ but not on v_i directly, and hence we can call this payment $p_i(a, v_{-i})$ where a is the alternative given by f. The second condition says that the mechanism maximizes the utility of player i given a fixed v_{-i}.

Sketch of Proof: We show these conditions are necessary and sufficient.

Necessity: If the first condition is violated, i.e., $\exists v_i, v'_i$, s.t. $f(v_i, v_{-i}) = f(v'_i, v_{-i})$ but $p_i(v_i, v_{-i}) > p_i(v'_i, v_{-i})$, then player i is better off by switching to v'_i. Also, if the second condition is violated, i.e., $f(v_i, v_{-i}) = a$ but $\arg\max_{x \in f(\cdot, v_-)}[v_i(x) - p_i(x, v_{-i})] = a' \neq a$, then $\exists v'_i$ s.t. $f(v'_i, v_{-i}) = a'$ and player i is better off by claiming v'_i.

Sufficiency: Suppose by contradiction the mechanism is not IC, and thus there exists a player i who can get a higher utility by switching to v'_i from v_i. However, when $f(v_i, v_{-i}) = f(v'_i, v_{-i})$, the utility remains the same; and when $f(v_i, v_{-i}) = a \neq f(v'_i, v_{-i}) = a'$, the new utility $(v_i(a') - p_i(a', v_{-i}))$ is not higher than the old utility $(v_i(a) - p_i(a, v_{-i}))$ either, because the second condition maximizes the old utility over $(\cdot, v_{-i})$. ■

We can prove many other theorems based on this nice result.

4 Combinatorial Auctions

Now we detour a little bit to talk about combinatorial auctions.

Consider n items and m bidders. Each bidder i has a valuation function v_i: subsets of $\{1, 2, \cdots, n\} \to \mathbb{R}$, and $v_i(\emptyset) = 0$.

It is natural to consider the following two-dimensional problem spaces with four categories:

	limited supply	**unlimited supply**
maximize welfare	VCG, but might be computationally hard	give things for free
maximize revenue	may talk about later	(*) our focus

Limited supply means we have n items and one copy of each item; unlimited supply means we have n items and unlimited copies of each item, e.g., software. Generally, mechanisms with unlimited supply are easier to handle than those with limited supply.

This following example illustrates why combinatorial auctions with limited supply may be computationally hard.

Example 1 *Assume each bidder i wants a specific set of items S_i, $v_i(S_i) = 1$, and for all $S \neq S_i$, $v_i(S) = 0$.*

Maximizing the social welfare here is equivalent to maximizing the number of non-overlapping sets (the maximum set packing problem), which is known to be NP-complete.

For combinatorial auctions with unlimited supply, if the goal is to maximize social welfare, we can simply give out items for free as they are unlimited. But it gets complicated when the goal is to maximize revenue. Two typical approaches are: (1) assume the valuation functions follow certain probability distributions, and (2) compare the worst case to the optimal social welfare or to the case with the best fixed prices on items. We will study the second approach next.

Consider a simple scenario with one item, such as water. To compare the maximum revenue to the OPT social welfare, we can apply the randomized weighted majority algorithm that we learned a couple of lectures ago. For example, we can imagine a mechanism that asks each person to submit his or her bid and then adjusts the price adaptively based on the observations so far. Since the value a player says will not affect the player's payment, such a mechanism is incentive compatible.

Formally, assume values are between 1 and H, we will prove that

- it cannot beat $\frac{1}{\log H} \times OPT$, but
- it can get $O(\frac{1}{\log H}) \times OPT$.

Can you come up with a probability distribution over valuations (i.e., a randomized adversary) as an example of the first bullet point?

We will continue on this topic in the next lecture.