CS787: Advanced Algorithms

Topic: Sponsored Search Auction Design **Presenter(s):** Nilay, Srikrishna, Taedong

17.6.1 Introduction to Auction Design

The Internet, which started of as a research project in 1960's, today influences almost every aspect of our daily lives. This has increasingly led marketers to focus on online advertising. Search engines like Google and Yahoo! source their revenues almost entirely from selling advertisement space. In this lecture we will study profit-maximization problems related to selling advertisement space on Internet web sites.

The online advertisement market is highly dynamic in nature, with advertisers arriving and leaving all the time. Moreover, the number of hits received by a website changes over time, as does the value of an Internet user clicking on an advertisement. Clearly the selling price of an advertisement depends on the benefit derived by the advertisers from the display of their advertisements. In most cases, the web site owner does not have a good estimate of the value of this benefit, as the advertisers needs often change with time. This is unlike classical optimization problems where the optimizer (the advertisement auctioneer in our case) knows the values of all the variables over which the optimization is to be carried out. This makes auctions the logical choice for selling advertisement space on web sites. The web site owners ask the advertisers themselves for their valuations of the benefit they will receive. The value reported by the advertiser is called her bid. However, the advertisers, being selfish agents, need not reveal their valuations truthfully if it is in their self-interest to do so.

The goal of the web site owner is to set up a selling mechanism that uses these reported values to determine the price(s) to be offered to the advertisers, in such a way that maximizes profit when each agent bids according to her best interest. At the same time, the web site owner would like to keep the process of finding the optimal bidding strategy simple in order to attract more advertisers to the online advertising market.

17.6.2 Mathematical Framework

Mechanism design deals with the construction of mechanisms, or games which are designed to get the agents to behave in a certain way so as to achieve some desired outcome. [2] provides a framework for algorithmic mechanism design.

- A mechanism deals with a set of agents $\mathcal{N}$ and wishes to choose from a collection of outcomes $\mathcal{O}$. Let $n = |\mathcal{N}|$ be the number of agents.
- Each agent i is assumed to have a *valuation function* $v_i : \mathcal{O} \to \mathbb{R}$. Intuitively, the valua-

tion function describes agent i's intrinsic preference for each outcome. An agent's valuation function is known only to that agent.

- $\mathcal{T}$ denotes the set of all possible valuation functions and $v \in \mathcal{T}^n$ is the vector of valuation functions of all agents.

The mechanism works by asking each agent to report her valuation function and computing an outcome and a set of payments based on the reported functions.

- b_i denotes the valuation function reported by agent i (also known as i's *bid*), $\mathbf{b} \in \mathcal{T}^n$ denotes the vector of bids.

- $P_i : \mathcal{T}^n \to \mathbb{R}$ be the payment made by agent i, and $P = (P_i)_{i \in \mathcal{N}}$ be the *payment scheme*.

Thus, a mechanism $\mathcal{M}$ consists of a pair (O, P), where $O : \mathcal{T}^n \to \mathcal{O}$ is the *output function* and P is the payment scheme.

Each agent's goal is to maximize her utility function, which is assumed to be of the form

$$u_i(O, P_i) = v_i(O) - P_i$$

Since the mechanism determines the outcome and the payments based on the bid vector $\mathbf{b}$, we will abbreviate $u_i(O(\mathbf{b}), P_i(\mathbf{b}))$ and $v_i(O(\mathbf{b}))$ to $u_i(\mathbf{b})$ and $v_i(\mathbf{b})$ respectively. Clearly, an agent's utility depends not only on her valuation function, but also on the bid vector.

- Let $\mathbf{b}_{-i} = (b_1, \ldots, b_{i-1}, ?, b_{i+1}, \ldots b_n)$ denote the vector of bids with agent i's bid hidden by a question mark. We will refer to this as the masked bid vector. We will write $\mathbf{b}$ as $(\mathbf{b}_{-i}, b_i)$.

- Similarly, let $\mathbf{v}_{-i}$ denote the vector of all other agents valuation functions, $\mathcal{T}_{-i} = \mathcal{T}^{n-1}$ be the space of those vectors.

A strategy $S_i : \mathcal{T} \to \mathcal{T}$ is said to be a *dominant strategy* for agent i if $u_i(\mathbf{b}_{-i}, S_i(v_i)) \geq u_i(\mathbf{b}_{-i}, b_i)$ for all $\mathbf{b}_{-i} \in \mathcal{T}_{-i}$ and $b_i, v_i \in \mathcal{T}$. In other words, agent i's best strategy is to report her valuation as $S_i(v_i)$ whenever her true valuation is v_i. A mechanism is said to be a *dominant strategy mechanism* if every agent has a dominant strategy.

Definition 17.6.2.1 (Truthful Mechanism) *A truthful mechanism is a dominant strategy mechanism in which truth-telling is a dominant strategy for every agent, i.e.*

$$u_i(\boldsymbol{b}_{-i}, v_i) \geq u_i(\boldsymbol{b}_{-i}, b_i) \quad \forall\ \boldsymbol{b}_{-i} \in \mathcal{T}_{-i} \text{ and } b_i,\ v_i \in \mathcal{T}$$

Theorem 17.6.2.2 (Bid-independence principle) *If the mechanism (O,P) is truthful, and $O(\boldsymbol{b}_{-i}, b_i) = O(\boldsymbol{b}_{-i}, b'_i)$, then $P_i(\boldsymbol{b}_{-i}, b_i) = P_i(\boldsymbol{b}_{-i}, b'_i)$.*

Proof: Suppose to the contrary, i.e. $P_i(\mathbf{b}_{-i}, b_i) > P_i(\mathbf{b}_{-i}, b_i')$, while $O(\mathbf{b}_{-i}, b_i) = O(\mathbf{b}_{-i}, b_i')$. When $v_i = b_i$ and all the other agents bid $\mathbf{b}_{-i}$, agent i is better off lying and bidding b_i', contradicting truthfulness. ■

This principle asserts that in a truthful mechanism, the bid of an agent affects the payment made by the agent only through its effect on the outcome of the mechanism.

17.6.3 Truthfulness as a Means to Simplified Bidding

As mentioned earlier, we would like to construct mechanisms for selling advertising space on Internet web sites under which it is easy for the advertisers to determine their optimal bidding strategies. Unless an advertiser has a dominant strategy, she would be forced to speculate (or hire someone to speculate for her) on how the other advertisers are going to bid in order to determine her optimal strategy. Thus, in order to get rid of speculation and keep the process of bidding simple, we would like each advertiser to have a dominant strategy, i.e. we would like to use a dominant strategy mechanism. The *Revelation principle* stated below allows us to restrict our attention to truthful mechanisms without missing any possible combination of outcome and payment functions (when we view the outcome and payment as a function of valuation rather than bids).

Theorem 17.6.3.1 (Revelation principle [5, 6]) *Every dominant strategy mechanism can be converted to a truthful mechanism without changing the outcome or the payments on vector of valuation functions of all agents.*

Proof: Given a dominant strategy mechanism M, construct a truthful mechanism M' that simulates each bidder's dominant strategy in M. Let S_i be agent i's dominant strategy under M. Then, M' produces the same outcome and payments on bid vector $\mathbf{b}$ as the mechanism M produces on the bid vector $(S_i(b_i))_{i \in \mathcal{N}}$. Clearly, agent i's dominant strategy under M' is to bid $b_i = v_i$, because this is equivalent to playing the dominant strategy $S_i(v_i)$ in M. ■

17.6.4 The Vickrey Auction

We next give an example of a classical truthful mechanism due to Vickrey [7]. Given the nature of its payment scheme, it is also called the *second-price auction*. It is used for selling a single item; so the outcome consists of the item being given to one of the agents. The valuation functions of the agents take a very simple form: if an agent gets the item, her valuation of the outcome is v_i (here we are slightly abusing notation by using the term v_i to refer to a single number); otherwise her valuation of the outcome is 0. The auction consists of inviting bids for the item, and giving the item to the highest bidder and charging her an amount equal to the second-highest bid. If two bidders tie for the highest bid, the item goes to the bidder with the lower index.

Theorem 17.6.4.1 *The Vickrey auction is truthful.*

Proof: In order to prove truthfulness, we have to show that none of the bidders can benefit by not revealing her true valuation function. Fix an agent i and let h be the highest bid among the other agents. If $v_i > h$, then the agent gets the item for an amount h whenever she bids an amount higher than h (which includes bidding her true valuation function). In this case, her utility is $v_i - h$.

The only possible way she can change the outcome is by bidding no more than h in which case her utility is 0.
On the other hand, if $v_i \leq h$, she makes a profit of 0 as long as she bids h or less (which includes bidding her true valuation). She can possibly change the outcome by raising her bid above h, in which case her utility becomes negative. This shows that bidder i cannot benefit by not bidding her true valuation. ■

17.6.5 The Vickrey-Clarke-Groves Mechanism

The most celebrated result in truthful mechanism design is the Vickrey-Clarke-Groves (VCG) mechanism [7]. It is a generalization of the Vickrey auction and can be used when the goal of a mechanism is to maximize the total valuation of all the agents. The VCG mechanism (O, P) is given by:

$$\begin{aligned} O(b) &= \quad o^* \qquad \textit{where } o^* \in \operatorname*{argmax}_{o \in \mathcal{O}} \sum_{i \in \mathcal{N}} b_i(o) \\ P_i(b) &= -\sum_{j \neq i} b_j(o^*) + h_i(\mathbf{b}_{-i}) \end{aligned}$$

where the functions h_i are arbitrary. That is, VCG selects the outcome that maximizes the total reported valuation, and charges agent i an amount P_i that depends on b_i only through its influence on the outcome o^*, just as required by bid-independence principle. Since the h_i terms in the payment are completely independent of agent i's bid, they are irrelevant to truthfulness. The term $\sum_{j \neq i} b_j(o^*)$ in the payment is quite special though – it aligns the utility function of agent i with the utilitarian objective function. This makes the mechanism truthful, as asserted by the following theorem.

Theorem 17.6.5.1 *The VCG mechanism is truthful.*

The basic VCG mechanism can be augmented by weighting the agents differently and adding a bias to the outcome function, while preserving truthfulness [3, 4]. More formally, let $w \in \mathbb{R}^n_+$ be a set of non-negative weights. Let $H : \mathcal{O} \to \mathbb{R}$ be a "bias" function. The resulting weighted, biased VCG mechanism is defined by :

$$\begin{aligned} O(b) &= \quad o^* \qquad \textit{where } o^* \in \operatorname*{argmax}_{o \in \mathcal{O}} \left(\sum_{j \in \mathcal{N}} w_j b_j(o) + H(o) \right) \\ P_i(b) &= -\frac{1}{w_i} \left(\sum_{j \neq i} w_j b_j(o^*) + H(o^*) \right) + h_i(\mathbf{b}_{-i}) \qquad \textit{when } w_i > 0 \\ P_i(b) &= \quad H(o^*) + h_i(\mathbf{b}_{-i}) \qquad \textit{when } w_i = 0 \end{aligned}$$

where the functions h_i are arbitrary.

Theorem 17.6.5.2 *For every choice of weights and bias function, the weighted, biased VCG mechanism is truthful.*

Proof: Clearly when the weight of agent i is zero, agent i has no incentive for being untruthful. When $w_i > 0$, then there are two possibilities – agent i is truthful, or agent i is untruthful. Assume that the output of the mechanism is o^* when i is truthful and it is o' when i is untruthful. Then,

$$o^* \in \underset{o \in \mathcal{O}}{\text{argmax}} \left(w_i v_i(o) + \sum_{j \in \mathcal{N},\ j \neq i} w_j b_j(o) + H(o) \right) \textit{ and } o' \in \underset{o \in \mathcal{O}}{\text{argmax}} \left(w_i b_i(o) + \sum_{j \in \mathcal{N},\ j \neq i} w_j b_j(o) + H(o) \right)$$

where v_i is the true valuation function of agent i and b_i is any other valuation function. Next we compute the utilities of agent i for outcomes o^* and o'.

$$\begin{aligned} u_i(o^*) &= v_i(o^*) - P_i(\mathbf{b}) \\ &= v_i(o^*) + \frac{1}{w_i} \left(\sum_{j \neq i} w_j b_j(o^*) + H(o^*) \right) + h_i(\mathbf{b}_{-i}) \\ &= \frac{1}{w_i} \left(w_i v_i(o^*) + \sum_{j \neq i} w_j b_j(o^*) + H(o^*) \right) + h_i(\mathbf{b}_{-i}) \end{aligned}$$

Similarly,

$$u_i(o') = \frac{1}{w_i} \left(w_i v_i(o') + \sum_{j \neq i} w_j b_j(o') + H(o') \right) + h_i(\mathbf{b}_{-i})$$

The manner in which o^* has been chosen makes it clear that $u_i(o^*) \geq u_i(o')$. Therefore, the agent has no incentive for bidding untruthfully. ■

17.6.6 Weaknesses of VCG Mechanism

Despite the attractiveness of the dominant-strategy property, the VCG mechanism also has several possible weaknesses [8].

- low(or zero) auctioneer revenues.
- non-monotonicity of the auctioneer's revenues in the set of bidders and the amount of bid.
- vulnerability to collusion by a coalition of losing bidders.
- vulnerability to the use of multiple bidding identities by a single bidder.

17.6.7 The Search Engine Problem

Consider a web page with slots where advertisements are displayed. Whenever the web page is accessed by an Internet user, the web site owner can choose to show her advertisements.

Definition 17.6.7.1 (Impression) *The process of displaying an advertisement is called impression.*

Since each of these advertisers invoke a different level of interest from the Internet users visiting the web pages we account for it using a parameter known as *click-through rate*.

Definition 17.6.7.2 (Click-through rate(CTR)) *Click-through rate of an advertisement is the fraction of its impressions that result in a click by an Internet user.*

In the auctions currently being used, search engines first pick the subset of advertisements to be displayed and match them to slots based on the submitted bids. The matching criteria is referred to as the *ranking function*.

Definition 17.6.7.3 (Ranking function) *A ranking function is a function mapping advertisements to slots by evaluating the rank of each advertisements.*

The ranking function is an integral component of the existing keyword auctions. Once the ranking function is determined, the auctioneer decides the pricing for each merchant based on both, the bids as well as the slot allocated.

There two popular ranking methods in use are:

- The *Overture*(or Yahoo!) method (Direct Ranking) : Merchants are ranked in the decreasing order of the submitted bids.
- The *Google* method (Revenue Ranking) : Merchants are ranked in the decreasing order of the *ranking scores*, where the ranking score of a merchant is defined as the product of the merchant's bid and a quantity known popularly as page rank which is determined by Google. Page ranks are usually not known to merchants.

In the following sections, we formulate a framework to help better understand the sponsored search engine auction.

17.6.8 Model and Notation

- N denotes the number of merchants.
- K denotes the number of slots on a specific keyword($K < N$).
- $\text{CTR}_{i,j}$ denotes click-through rate of the i^{th} merchant if placed at slot $j \leq K$.
- v_i denotes the true valuation of a click-through to merchant i.
- b_i denotes the bid of the i^{th} merchant for a click-through.
- p_i denotes a price-per-click charged for merchant i. $p_i \leq b_i$.
- w_i denotes a priori weight assigned to merchant i by the auctioneer. This is independent of her bid.

17.6.9 Assumptions

- In sponsored search, we typically work with $K \leq N$. However, this can be easily generalized to the case with $K \geq N$ by reducing the number of slots K to N and adding dummy merchants with all relevant parameters set to 0.
- Slots are ordered so that the probability of clicks going through is higher for slots ranked higher.
- $\text{CTR}_{i,j}$ is arbitrary and non-increasing in j. In most cases, the auctioneer is aware of the Click Through Rates.
- v_i is known to merchant i, but not to the auctioneer.
- The merchants are ranked in the order of decreasing $w_i b_i$.
- The merchants are risk-neutral.

We formally define the next-price auctions as follows

Definition 17.6.9.1 (Next-price Auction) *Given the rank function, $R = (w_1, w_2, \cdots, w_n)$ and the bid vector $\mathbf{b} = (b_1, \cdots, b_n)$, the next-price auction ranks the merchants in the decreasing order of $w_i b_i$ and charges the merchant ranked i an amount-per-click equal to the minimum bid she needs to have submitted in order to retain rank i. Then the price charged to the merchant ranked i is*

$$p_i = \frac{w_{i+1} b_{i+1}}{w_i}.$$

Note : Setting $w_i = 1$ for all i is equivalent to the direct ranking function (the Overture model), while setting $w_i = \text{CTR}_{i,1}$ reduces to the revenue-ranking function (the Google model).

In order to model the Click-through rates, we assume that they can be separated into a merchant-specific factor and a position-specific factor.

Definition 17.6.9.2 (Separable Click-through Rates) *The click-through rates are said to be separable if there exist $\mu_1, \mu_2, \cdots, \mu_n > 0$ and $\theta_0 \geq \theta_2 \geq \cdots \geq \theta_K > 0$ such that the click-through rate $CTR_{i,j}$ of the i^{th} merchant at the j^{th} slot is given by $\mu_i \theta_j$*

17.6.10 Need for a New Auction

In this section, we provide the motivation for designing a new auction by showing that the next-price auctions being currently used by Google and Yahoo! are not truthful. In order to reiterate the shortcomings of the VCG mechanism, we give instances of ranking functions for which there does not exist any set of weights and biases for which the ranking output by the VCG mechanism is always the same as the output by the given ranking function.

17.6.10.1 Next-price Auction is not Truthful

Consider the following example.

- Three merchants A, B and C bidding for two slots.
- All of them have a click through rate of 0.5 at the top slot and 0.4 at the bottom slot.
- The true valuations per click of the three merchants be 200, 180 and 100 respectively.

If all the merchants bid truthfully, merchant A ends up paying a price of 180 per click, making an expected profit of $(200 - 180) \times 0.5 = 10$. In this case, she has an incentive to under cut B by lowering her bid to 110, and make a net profit of $(200 - 100) \times 0.4 = 40$.

Clearly, there is no incentive to bid higher than ones true valuation under the next-price auction. This is because the price-per-click charged is the minimum bid required to retain one's rank. Therefore in cases where bidding higher improves one's rank, the price-per-click charged is higher than one's true valuation.

17.6.10.2 Weighted VCG may not Always Apply

In this section, we show by means of a counter-example that even for the simple case of direct ranking, there does not exist any set of (bid-independent) weights and biases for which the VCG solution achieves the same allocation as direct ranking. This will show that, in general, VCG does not apply to our problem. Consider the following example.

- Two merchants A and B bidding for two slots.
- Merchant A has a click-through rate of 0.4 at both the first and the second slot.
- Merchant B has a click-through rate of 0.4 at the first slot and 0.2 at the second slot.

Since any of the merchants can bid the highest and get the top slot in direct ranking, both the merchants must have non-zero weight in order for weighted VCG to achieve the same allocation as direct ranking.

- Let $w_A > 0$ and $w_B > 0$ be the weights assigned by the VCG mechanism to merchants A and B respectively.
- $H(x, y)$ denotes the bias assigned to ranking merchant x followed by y for $x, y \in \{A, B\}$.

Then the VCG mechanism will rank B before A if

$$w_A(0.4b_A) + w_B(0.4b_B) + H(B, A) > w_A(0.4b_A) + w_B(0.2b_B) + H(A, B).$$

This is true whenever

$$b_B > \frac{H(A, B) - H(B, A)}{0.2w_B},$$

Hence, the ranking of the merchants by VCG does not depend on A's bid. However, the direct ranking scheme will rank A before B whenever A's bid is higher than B's bid. Thus, the VCG mechanism does not apply to this example. We formally state (without proof) a few more shortcomings of the VCG auction. (For proof, see [1].)

Theorem 17.6.10.1 *Let the number of merchants with non-zero click-through rates be $n > K$. If the click-through rates are not separable, then there exists a ranking function $R = (w_1, w_2, \cdots, w_n)$ for which there does not exist any set of weights for which unbiased, weighted VCG always yields the same ranking as the ranking function R.*

Theorem 17.6.10.2 *Let the number of merchants with non-zero click-through rates be $n > K$. If the click-through rates are not separable, then there exists a ranking function $R = (w_1, w_2, \cdots, w_n)$ for which there does not exist any set of weights for which biased. weighted VCG always yields the same ranking as the R.*

Theorem 17.6.10.3 *Let the click-through rates be separable. Then the VCG mechanism having merchant i's VCG weight set to $\dfrac{w_i}{CTR_{i,1}}$ always produces the same ordering as the ranking function $(w_1, \cdots, w_n)$.*

The above theorem implies that with the separability assumption, the ranking functions maximize a certain global utility function. In particular, the revenue-ranking scheme maximize the total utility obtained by the merchants and the auctioneer.

17.6.11 The Truthful Auction

We will assume without loss of generality that the i^{th} merchant also has the i^{th} rank in the auction. The truthful auction is quite simple: For $j \le i \le K$, set the price-per-click p_i charged to merchant i as:

$$p_i = \sum_{j=i}^{K} (\mathrm{CTR}_{i,j} - \mathrm{CTR}_{i,j+1}) \frac{w_{j+1}}{w_i} b_{j+1} \tag{17.6.11.1}$$

In other words,

1. For those clicks which merchant i would have received at position $i+1$, she pays the same price as she would have paid at position $i+1$.
2. For the additional clicks, merchant i pays an amount equal to the bid value of merchant $i+1$.

Since $w_i b_i \ge w_j b_j$ for $j > i$, it follows that $p_i \le b_i$. Hence the price charged per click-through can be no larger than the submitted bid. We will refer to this auction as *Laddered_Auction*$(w_1, \cdots, w_n)$.

17.6.12 Analysis

For the Laddered_Auction$(w_1, \cdots, w_n)$, we state the following theorem without proof. (For proof, see [1].)

Theorem 17.6.12.1 *Given fixed $w_1, \cdots, w_n$, the Laddered_Auction$(w_1, \cdots, w_n)$ is truthful. Further, it is the unique truthful auction that ranks according to decreasing $w_i b_i$.*

Corollary 17.6.12.2 *For any fixed $w_1, \cdots, w_n$, the Laddered_Auction$(w_1, \cdots, w_n)$ is the profit-maximizing truthful auction that ranks merchants by decreasing $w_i b_i$.*

References

[1] Gagan Aggarwal. Privacy Protection and Advertising in a Networked World. PhD. Thesis, 2005

[2] N. Nisan, and A. Ronen. Algorithmic mechanism design. In *Proceedings of the 31st Annual ACM Symposium on Theory of Computing*, 1999.

[3] J. Green, and J. Laffont. Characterization of satisfactory mechanisms for the revelation of preferences for public goods. In *Econometrica*, 1977.

[4] K. Roberts. The characterization of implementable choice rules. In *J. Laffont, editor, Aggregation and Revelation of Preferences, North Holland Publishing Company*, 1979.

[5] P. Dasgupta, P. Hammond, and E. Maskin. The implementation of social choice rules: Some results on incentive compatibility. In Review of Economic Studies, 1979

[6] R. Myerson. Incentive compatibility and the bargaining problem. In Econometrica, 1979.

[7] W. Vickrey. Counterspeculation, auctions and competitive sealed tenders. In Journal of Finance, 1961

[8] Lawrence M. Ausubel, and Paul Milgrom. The Lovely but Lonely Vickrey Auction. In Combinatorial Auctions, 2006