CS 405/6001: Game Theory and Algorithmic Mechanism Design **Jul–Nov 2026**

Lecture 20: Internet Advertising and the Limits of VCG

Lecturer: Swaprava Nath *Scribe(s): Claude Scribe*

Disclaimer: *These notes aggregate content from several texts and have not been subjected to the usual scrutiny deserved by formal publications. If you find errors, please bring to the notice of the Instructor.*

20.1 Recap

Last lecture introduced the **Vickrey-Clarke-Groves (VCG) mechanism**, the Groves mechanism with $h_i(\theta_{-i}) = \max_{a\in A} \sum_{j\neq i} v_j(a, \theta_j)$:

$$p_i^{VCG}(\theta) = \left(\max_{a\in A} \sum_{j\neq i} v_j(a, \theta_j) \right) - \left(\sum_{j\neq i} v_j(f^{AE}(\theta_i, \theta_{-i}), \theta_j) \right).$$

- **Payment**: agent i pays the **loss in the maximum social welfare of the other agents** caused by its presence.
- **Utility**: agent i's utility equals its **marginal contribution to the overall social welfare**.
- **Single object**: VCG reduces exactly to the **second price auction**. Only **pivotal** agents pay.
- **Combinatorial allocation of goods**: an agent who receives no object pays 0, and VCG is **ex-post individually rational**.

Today we see VCG at work in its largest real-world application, and then where it breaks.

20.2 Application Domain: Internet Advertising

Internet advertising is defined as the method of delivering promotional marketing messages called 'ads' to consumers when a user searches something on a search engine or visits a website. In this lecture we will consider the methods of efficiently placing an ad and deciding the payment to the advertiser.

20.2.1 Advantages of Internet Advertising

Internet advertising is much better than conventional advertising in newspapers, articles, or radio because of the following advantages.

1. **User Data**: using Internet ads, the advertiser can gain information about which set of buyers are interested in which set of ads. This enables the advertisers to target certain products to the interested users.

2. **Measurable Actions**: classification of buyers is possible using Internet advertising. The buyers can be classified into a set of groups and be shown ads according to their interest.
3. **Low Latency**: Internet advertising allows the auctions to happen on the fly. The auction takes place just before the ad is to be shown. For example in search engines like Google, bidding takes place after a user enters a set of keywords to search, and ads for the organizations that win the auction are shown. This also allows the advertisers to enter any time and enables automation of the whole process.

Each of these three is a statement about *information*: the platform knows more, measures more, and acts faster than any newspaper ever could. That is precisely the setting in which a mechanism, rather than a posted price, earns its keep.

20.2.2 Types of Ads on the Internet

1. **Sponsored Search Ads**: in this type of ads, the advertisers bid on the keywords that are entered by the user while searching on a search engine. Some details of this method are as follows.
   - Some words can have multiple meanings based on the context. Like 'jaguar' can mean a car manufacturer or an animal. So to prevent the car ads when the animal jaguar is searched, stop words are used. Stop words are a set of words which if used with a particular word (here, jaguar), the ad (for car jaguar) is not presented.
   - One user can search for the same keyword multiple times. So to prevent activity of a malicious user, or to prevent draining of money from the advertiser, or to prevent irritating the user with the same ad every time, a cap is set on the number of times an ad can be shown to a user.
2. **Contextual Ads**: these type of ads are presented to a user based on the context of a page. For example, gmail can present certain type of ads based on the content of the email.
3. **Display Ads**: it is the classical way of displaying the ads just like the banner ads in the newspapers or magazines. Display ads are typically placed in dedicated areas of a webpage and are often image-based, animated, or interactive.

Ads are complex – modern internet advertising is handled via ad exchanges, which also allow small businesses to customize these ads.

20.2.3 Position Auction

This type of auction is used to sell multiple ad slots in a webpage. Let $N = \{1, 2, \ldots, n\}$ denote the set of bidders and $M = \{1, 2, \ldots, m\}$ be the set of slots available to be auctioned. For simplicity, we assume $m \geqslant n$, i.e., every advertiser gets a slot – the question is who gets which slot. Here among the set of slots, 1 represents the 'best' position and m represents the 'worst' position.

20.2.3.1 Deciding the Auction Mechanism

The auction mechanism decides the allocation of resources (e.g., the position at which an ad is displayed) and the payment to be made by each advertiser. The following are some standard methods of auctions.

1. **Early position auctions**: in this type of auction, the received bids are sorted in decreasing order of value and the positions are allotted according to this order – the highest bidder receiving the first

position and lowest bidder receiving last position – and they pay some fixed amount according to their position upfront. This type of bidding puts all risk on the advertiser: now an advertiser has to estimate the probability of getting a click since they are paying even before the ad is clicked.

2. **Pay per click model**: in this mechanism, the price charged to the advertiser is proportional to the number of clicks the users made on the ad. This mechanism puts all the risk on the website because there could be an ad that is at one of the top positions but never gets clicked and hence does not produce any revenue.
3. **Rank by expected revenue**: modern sponsored search auctions use a middle ground between these two extreme approaches. In this approach, the advertisers are ranked by product of *estimated* CTR and the bid of the advertiser. This mechanism shares the risk between the advertiser and the website.
4. **Today's approach**: rank advertisers based on the product of the probability of a click and the bid value.
   - Probability of click is called click through rate (CTR)
   - rank by expected revenue

20.2.4 Advertiser Value

Advertiser value is the valuation received by the advertiser from her ad on the Internet. We make the following assumptions in this setting.

1. Only clicks generate the value to the bidders.
2. All the clicks are valued equally no matter which position it came from. This means that the role of position is only in the probability of being clicked. This assumption helps in decoupling the position effect from the value effect.

The expected value of agent i, when her ad is shown at position j, is

$$v_{ij} = CTR_{ij} \cdot v_i$$

where CTR_{ij} is the click through rate (CTR) that denotes the probability that bidder i's ad gets clicked at position j. This quantity CTR_{ij} is again assumed to be decomposable into two terms: (1) the quality of the ad of the advertiser i, denoted by $\rho_i \in [0,1]$, that reflects how good the ad is, and (2) the effect of position on the CTR, denoted by $p_j \in [0,1]$. Therefore

$$CTR_{ij} = \rho_i \cdot p_j, \qquad \rho_i \in [0,1],\ p_j \in [0,1].$$

Thereby the value of agent i when her ad is in position j is decomposable into the position effect and agent effect:

$$v_{ij} = \underbrace{p_j}_{\text{position effect}} \cdot \underbrace{(\rho_i v_i)}_{\text{agent effect}}.$$

Note that both p_j and ρ_i are measurable and can be estimated by the auctioneer, but v_i is private information (type) of the agent. (A caution on symbols: in this section p_j denotes the *position component* of the CTR, while a payment always carries a superscript, as in p_i^{VCG}.) The position effect is assumed to be decreasing with position: $p_1 = 1$ and $p_j > p_{j+1}$ for all $j = 1, \ldots, m-1$.

We use the following notations:

- $\hat{\rho}_i$: the value of the click through rate of the Ad i estimated by the search engine
- b_i: the amount that agent i is willing to pay if a click occurs
- $a_i \in M$: the position assigned to agent i s.t. $a_i \neq a_k$, $i, k \in N$
- $a = (a_1, a_2, \ldots, a_n)$: allocation of ads to positions

Here we assume that the ads are allocated according to the decreasing order of their expected revenue, i.e., according to the decreasing order of $\hat{\rho}_i \cdot b_i$. Suppose the agent i reports his bid as b_i. Then

$$\hat{v}_i(a) = p_{a_i}(\hat{\rho}_i \cdot b_i)$$

is the reported value of the agent for the allocation a.

Allocation rule: standard search engines and advertisement websites pick the allocation a^* such that it maximizes the sum of the reported valuations of the advertisers, i.e.,

$$a^* \in \arg\max_{a \in A} \sum_{i \in N} \hat{v}_i(a).$$

This method of allocation is exactly same as the VCG allocation.

Remark 20.1 (Why the allocation is easy here) *An allocation is efficient if and only if it is a 'rank-by-expected-revenue' mechanism, ranking by $\hat{\rho}_i \cdot b_i$. The reason is that this is a moment maximization problem: a sum of products is maximized when the largest weight sits on the largest value. So slot allocation is a* sorting *problem, computationally trivial, in sharp contrast to general combinatorial allocation where computing f^{AE} is NP-hard.*

Claim 20.2 (Why sorting is efficient) *Write $w_i = \hat{\rho}_i b_i$ for the agent effect, and recall $p_1 = 1 > p_2 > \cdots > p_m > 0$. The value of an allocation a is $V(a) = \sum_{i \in N} p_{a_i} w_i$. Then a maximizes V if and only if it ranks agents by w_i: a larger w_i gets the better slot (ties broken arbitrarily).*

Proof: Suppose $w_i > w_k$ but i sits *below* k, i.e., $p_{a_i} < p_{a_k}$. Swap the two. Only their terms change, and

$$\underbrace{(p_{a_k} w_i + p_{a_i} w_k)}_{\text{after swap}} - \underbrace{(p_{a_i} w_i + p_{a_k} w_k)}_{\text{before swap}} = \underbrace{(p_{a_k} - p_{a_i})}_{>0} \underbrace{(w_i - w_k)}_{>0} > 0.$$

So any such inversion can be swapped for a strict gain, and an efficient allocation has none: it is sorted. Conversely, each swap removes at least one inversion, so finitely many swaps turn any allocation into the sorted one without losing value; hence the sorted allocation is efficient. ■

Remark 20.3 *This is the **rearrangement inequality**: a sum of products is largest when both sequences are sorted the same way. Computing f^{AE} here is just* sorting *by $\hat{\rho}_i b_i$.*

20.2.5 VCG Payments in a Position Auction

Given bids $(b_1, \ldots, b_n)$, assume without loss of generality that the order is such that $\hat{\rho}_1 b_1 \geqslant \hat{\rho}_2 b_2 \geqslant \cdots \geqslant \hat{\rho}_n b_n$. Then the efficient allocation a^* is $a^*_i = i$, and in the allocation $a^*_{-i} \in \arg\max_a \sum_{j \neq i} \hat{v}_j(a)$, the slots allocated

to the agents after i, i.e., from $i+1$ to n, each get one slot better than in a^*. Hence

$$p_i^{\text{VCG}} = \sum_{j \neq i} \hat{v}_j(a^*_{-i}) - \sum_{j \neq i} \hat{v}_j(a^*) = \sum_{j=i}^{n-1} p_j(\hat{\rho}_{j+1} b_{j+1}) - \sum_{j=i}^{n-1} p_{j+1}(\hat{\rho}_{j+1} b_{j+1})$$
$$= \sum_{j=i}^{n-1} (p_j - p_{j+1})(\hat{\rho}_{j+1} b_{j+1}), \quad \forall i = 1, \ldots, n-1, \qquad p_n^{\text{VCG}}(b) = 0.$$

Read the formula: each bidder pays, for every slot below them, the *drop in click probability* they inflicted on the bidder occupying it, priced at that bidder's expected revenue. The last bidder displaces nobody, so pays nothing, which is Lecture 19's empty-handed claim showing up again in a different costume. This is the total expected payment; to convert it to a pay-per-click price, divide by $p_i \cdot \hat{\rho}_i$:

$$\frac{1}{p_i \cdot \hat{\rho}_i} p_i^{\text{VCG}}(b).$$

Note: actual implementation in practice might be different; here we discuss only the principle of its computation.

20.3 Advantages of the VCG Mechanism

The VCG mechanism possesses several powerful and desirable properties.

- **DSIC: hence very low cognitive load on the bidders.** The mechanism is Dominant Strategy Incentive Compatible. This means that for every agent, reporting their true valuation is a weakly dominant strategy. Participants do not need to strategize or speculate about others' actions; their optimal strategy is simply to be truthful. This significantly lowers the cognitive burden on bidders.

- **No subsidy (and therefore, deficits) under certain conditions if items are goods.** The VCG payment rule ensures that the mechanism does not run at a loss. The sum of payments collected is always greater than or equal to zero, meaning the auctioneer never has to pay out of pocket to finance the allocation. This is a crucial property for any practical auction system.

- **Never charges an agent who gets no items.** In an allocation of goods, if an agent does not win any item in the efficient allocation, their VCG payment is zero. This is because their presence did not change the optimal allocation for the other agents, meaning the externality they imposed was zero. This is an intuitive and fair property.

- **Individually rational to participate: nobody loses money.** Under common conditions (such as the allocation of goods), the VCG mechanism is individually rational. This guarantees that every participating agent will receive a non-negative utility. No agent is made worse off by participating, which ensures voluntary participation in the mechanism.

Every one of these guarantees, however, is about a *single, honest, isolated* run of the mechanism. The rest of the lecture is about what happens when agents talk to each other, when an irrelevant cost moves, or when one more bidder walks in.

20.4 Disadvantages of the VCG Mechanism

Despite its theoretical strengths, VCG faces significant practical challenges that often limit its implementation.

20.4.1 Privacy and Transparency Issues

The DSIC property requires agents to fully reveal their private valuations. This lack of privacy can be detrimental in competitive environments: a bid is a statement about your business, and two competing companies would rather not make that public. It also makes the mechanism vulnerable to a **malicious auctioneer**, who could introduce **fake bidders (shills)** to manipulate the externality calculation and extract higher payments. This is precisely why on-chain auction designs worry about who orders the transactions, and why ad exchanges are audited: the mechanism is only as honest as the party computing $\max_a \sum_{j \neq i} v_j(a, \theta_j)$.

20.4.2 Susceptibility to Collusion

VCG is not group-strategyproof and is known to be **vulnerable to collusion**. A coalition of agents can coordinate their reports to manipulate the outcome, lower their combined payments, and increase their joint utility.

Figure 20.1: Example of Collusion in a Public Good Auction (A vs. B)

(a) No Collusion

Agent	Valuation for A	Valuation for B
1	200	0
2	100	0
3	0	250

Outcome: A wins ($300 > 250$).

$$\begin{aligned} p_1 &= (v_2(A) + v_3(B)) - (v_2(A) + v_3(A)) \\ &= (100 + 250) - (100 + 0) = 150, \\ p_2 &= (v_1(A) + v_3(B)) - (v_1(A) + v_3(A)) \\ &= (200 + 250) - (200 + 0) = 50. \end{aligned}$$

(b) With Collusion

Agent	Valuation for A	Valuation for B
1 (lies)	250	0
2 (lies)	150	0
3	0	250

Outcome: A wins ($400 > 250$).

$$\begin{aligned} p_1 &= (v_2(A) + v_3(B)) - (v_2(A) + v_3(A)) \\ &= (150 + 250) - (150 + 0) = 100, \\ p_2 &= (v_1(A) + v_3(B)) - (v_1(A) + v_3(A)) \\ &= (250 + 250) - (250 + 0) = 0. \end{aligned}$$

The outcome is unchanged – A wins in both cases – but the colluding pair's total payment falls from 200 to 100, so both are strictly better off. Individually, neither agent can gain by lying, which is exactly what DSIC promises. Together, they can, which is exactly what DSIC does not promise.

20.4.3 Not Frugal

VCG payments can be disproportionately large and sensitive to the valuations of irrelevant alternatives. In a cost-minimization setting (where costs are negative valuations), an agent's payment can be dramatically influenced by the cost of another agent who is not even part of the optimal solution. This can make payments seem unfair and unpredictable.

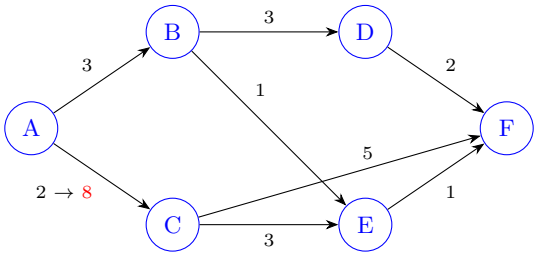

Figure 20.2: Item delivery network where edges are agents with costs (negative valuations). The efficient path is A → B → E → F.

Consider the item delivery network shown in Figure 20.2, where each edge is an agent and their cost is their private valuation. The efficient allocation is the shortest path from A to F, which is A→B→E→F, with a total cost of $3 + 1 + 1 = 5$. Let us analyze the payment to agent A→B (p_{AB}).

- **Scenario 1: Cost of A→C is 2.**

$$p_{AB} = \left(\max_{a \in A} \sum_{j \neq AB} v_j(a) \right) - \left(\sum_{j \neq AB} v_j(f^{AE}) \right)$$

  The first term is the best outcome for all other agents without A→B:

$$\text{Path: } A \to C \to E \to F, \quad \text{Cost} = 2 + 3 + 1 = 6, \quad \text{Valuation} = -6.$$

  The second term is the sum of valuations for the other agents in the efficient path:

$$\text{Agents: } B \to E,\ E \to F, \quad \text{Valuation} = -1 + (-1) = -2.$$

  Hence

$$p_{AB} = (-6) - (-2) = -4.$$

  This is a payment *from* the mechanism to the agent: the edge receives 4.

- **Scenario 2: Cost of irrelevant edge A→C increases to 8.** The efficient path remains unchanged (A→B→E→F). Without A→B:

$$\text{Path: } A \to C \to E \to F, \quad \text{Cost} = 8 + 3 + 1 = 12, \quad \text{Valuation} = -12.$$

  The second term remains -2. Thus

$$p_{AB} = (-12) - (-2) = -10.$$

The payment to agent A→B has more than doubled, from 4 to 10, purely because the cost of an unused, irrelevant edge changed. This demonstrates that VCG is not frugal: its payments are hostage to irrelevant alternatives.

20.4.4 Revenue Monotonicity Violation

Revenue monotonicity requires that the auctioneer's revenue does not decrease when more bidders join. VCG **violates this property**. Total revenue can decrease with the addition of a new agent, which is counter-intuitive for an auction. Consider the following example.

Figure 20.3: Example of Revenue Monotonicity Violation

(a) Scenario with 2 Agents

Agent	Val. for F	Val. for M
1	100	0
2	0	90

(b) Scenario with 3 Agents

Agent	Val. for F	Val. for M
1	100	0
2	0	90
3 (new)	0	90

Outcome: F wins.

$$p_1 = v_2(M) - v_2(F) = 90 - 0 = 90,$$
$$p_2 = v_1(F) - v_1(F) = 100 - 100 = 0.$$
$$\textbf{Total Revenue} = \textbf{90}.$$

Outcome: M wins.

$$\begin{aligned} p_1 &= \max(v_2(F) + v_3(F),\, v_2(M) + v_3(M)) - (v_2(M) + v_3(M)) \\ &= \max(0, 180) - 180 = 0, \\ p_2 &= \max(v_1(F) + v_3(F),\, v_1(M) + v_3(M)) - (v_1(M) + v_3(M)) \\ &= \max(100, 90) - (0 + 90) = 10, \\ p_3 &= \max(v_1(F) + v_2(F),\, v_1(M) + v_2(M)) - (v_1(M) + v_2(M)) \\ &= \max(100, 90) - (0 + 90) = 10. \end{aligned}$$
$$\textbf{Total Revenue} = \textbf{20}.$$

When agent 3 joins, the efficient outcome switches from F to M, and the total revenue for the auctioneer drops sharply from 90 to 20. This is a direct violation of revenue monotonicity. The mechanism is easy to read once stated in externality terms: agent 1 is no longer pivotal at all, because M wins even without them, and agents 2 and 3 each displace very little, because the other one would have carried M anyway.

20.4.5 Not Budget Balanced

While VCG guarantees no deficit, it is almost never **budget-balanced**, typically generating a surplus. This surplus cannot be redistributed to the participants, as this would break the DSIC property. If the players are partitioned into two groups and the surplus of one group is redistributed over the other group, then it is budget balanced, but the overall efficiency is compromised. This phenomenon is often called "money burning," as the surplus must be removed from the system.

Theorem 20.4 (Green–Laffont) *In a general quasi-linear setting, no mechanism can be allocatively efficient, DSIC, and budget balanced at the same time.*

The natural question is then: how much to burn, and what efficiency do we compromise? This is studied in [NS19].

20.5 Concluding Remark

These are certain limitations that are good to know for effective use of VCG; however, it is the most widely used mechanism in the literature. Next week we take up what remains possible once one of the three properties in the Green–Laffont theorem is given up.

References

[GL79] J. Green and J.-J. Laffont, *Incentives in Public Decision Making*, North-Holland, 1979.

[NS19] S. Nath and T. Sandholm, "Efficiency and budget balance in general quasi-linear domains," *Games and Economic Behavior*, 113, 2019, pp. 673–693.

[V61] W. Vickrey, "Counterspeculation, auctions, and competitive sealed tenders," *Journal of Finance*, 16(1), 1961, pp. 8–37.