CS532L Final Project
A survey of online advertisement auction research

Anton Likhtarov

April 22, 2009

1 Introduction

In recent years, online advertising has become a huge market. It is a major source of income for such companies as Google, Yahoo!, or Microsoft. Accordingly, these companies are interested in maximizing the revenues they can get by selling advertising spots, and a lot of research has been done on this topic.

Edelman at al. [6] provide a short outline of the evolution of Internet advertising. In the early days of Internet advertising, ads were sold per impression. An advertiser would pay once to show their ad a fixed number of times.

In 1997, Overture (later acquired by Yahoo!) developed a new model. They would auction off several ad spots shown on a search results page for a particular keyword. Advertisers submitted bids for a number of different keywords, and the spots would be allocated according to the bids — highest bidder for the particular keyword searched would receive the first (best) spot, second highest would receive the second-best spot, and so on. Advertisers would also pay the amount of their bid every time their ad was clicked on.

This model had a huge advantage over the earlier per-impression models, since now advertisers could target their ads. A user who searched for a "honda" and clicked on a car dealership's ad is on average more valuable to the dealership than a user who simply saw their ad on some page. The particular auction design used had its problems however. As we will see, this "Generalized First-Price auction" is unstable, and the search engine's revenues are lower than under other models.

Google came up with another keywords auction system (AdWords) in 2002. Instead of paying the amount of their own bid, advertisers would now pay the next highest bidder's amount plus a small constant. We will see that this "Generalized Second-Price" auction allows higher revenues and more stable bids. Indeed, a number of equilibria is possible under some conditions.

We study the auction for a single keyword. Note that in real applications, the bidders have limited budgets and allocating these budgets across different keywords they bid on to maximize revenue is in itself an interesting problem [9, 10]. In the same way, the advertisers will attempt to optimize their bids under the constraints to get the best return on investment [3, 4]. We will assume that the budgets are unlimited.

Some of other interesting problems to consider are how the ad spot location on the web page affects the probability that a user will click on it; whether or not other displayed ads affect the probability of clicking on some particular ad; how much the contents of the ad itself affect its probability of being clicked on (the Click-Through Rate, CTR); and how to estimate these probabilities. These problems are beyond the scope of this survey.

2 Model

We have N advertisers and M ad *slots* — locations where the ads can be placed on the search results web page. In the unweighted model [6], we assume that the only thing affecting the probability of an ad being clicked on is its position on the page. The number of clicks per second received by an ad in slot i (the Click-Through Rate) is c_i. We order the slots so that if $i < j$ then $c_i > c_j$. In a weighted model, CTR will generally also depend on which advertiser got the slot. CTR in [12, 1] is assumed to be separable; that is $c_{ij} = x_j c_i$ where c_i is the component of the CTR that is only affected by the slot's position on the page, and x_j is the component that is associated with the bidder, but is independent of the position.

Each advertiser has their own value for a click, and we call the j-th advertiser's value v_j. Thus, the j-th advertiser's payoff if his ad is in position i is in general $c_{ij}(v_j - p_j)$, where p_j is his payment per click to the search engine under the auction payment rules.

An auction is modelled as follows: at any point in time, an advertiser can submit a bid. When the keyword is searched for, only advertisers' most recently submitted bids are considered. We call the i-th advertiser's current bid b_i. Then the advertiser with the highest bid gets the first (best) spot, the second-highest bidder gets the second spot, and so on; though other ranking functions are possible, e.g. [1, 7] (in particular, this is not the function Google uses). The bidders are also made to pay according to the particular auction payment rules.

3 Generalized first-price auction

Under the GFP rules (used by Overture/Yahoo! from 1997 to 2002 [6, 11]), the bidders pay the amount of their bids, that is $p_i = b_i$. Consider an example with two slots and three bidders. Suppose that $c_1 = 10$ and $c_2 = 5$.

Note that this auction is not truthful - if everyone bid their true values, each bidder would gain by reducing their own bid to just above the next highest. Also note that this auction exhibits cycling behaviours [2]. Suppose the bids are as above. Eventually, Bob will notice that he doesn't have to pay more than 3 to stay in the same spot, and will reduce his bid. Alice will likewise reduce her bid to 4. Now Bob has to pay only 5 to get into the first spot. Then Alice

Table 1: GFP example

Advertiser	v_i	b_i	Slot	c_i	$p_i c_i$	Total payoff
Alice	50	40	1	10	400	100
Bob	20	19	2	5	95	5
Charlie	2	2	None	0	0	0

will increase her bid to 6, and so on.

Note that in these bidding wars, all bids will stay close to the lowest bid necessary to get the last spot. Thus the search engine's revenue will never get too high. The other loss is the cost to the advertisers, who now have to invest into designing effective bidding systems (especially if the bids are allowed to change very often) to have any hope of getting a reasonable ad spot.

4 Generalized second-price auction

To solve this problem, a new mechanism was adopted by Google in 2002 [6]. In GSP, the bidders pay the amount of the next highest bid (in practice, a small constant is usually added, e.g. 1 cent). Using the above example,

Table 2: GSP example

Advertiser	v_i	b_i	Slot	c_i	$p_i c_i$	Total payoff
Alice	50	40	1	10	190	310
Bob	40	19	2	5	10	105
Charlie	2	2	None	0	0	0

Note that this auction is also not truthful. Suppose everyone bids their true value:

Table 3: GSP example - true bids

Advertiser	v_i	b_i	Slot	c_i	$p_i c_i$	Total payoff
Alice	50	50	1	10	400	100
Bob	40	40	2	5	10	190
Charlie	2	2	None	0	0	0

Now Alice could reduce her bid to 3 and get a higher payoff:

Table 4: GSP example - Alice's strategy

Advertiser	v_i	b_i	Slot	c_i	$p_i c_i$	Total payoff
Alice	50	3	2	5	10	240
Bob	40	40	1	10	30	370
Charlie	2	2	None	0	0	0

GSP auction was researched extensively [6, 11, 8, 12]. In certain ways, it behaves like VCG. VCG is used a lot in the following discussion, to let's consider it explicitly. Under VCG, the allocation rules are the same as in GSP, but each bidder j would be made to pay the sum of $(c_{i-1} - c_i)b_i$ for each position i below him - since that bidder would get the spot $i-1$ if not for the bidder j. (Bidders in higher positions are not affected). Here's the same example with payoffs under VCG, assuming truth-telling (since it's a dominant strategy in VCG):

Table 5: VCG payoffs

Advertiser	v_i	b_i	Slot	c_i	$p_i c_i$	Total payoff
Alice	50	50	1	10	210	290
Bob	40	40	2	5	10	190
Charlie	2	2	None	0	0	0

Note that the advertisers' payments are smaller in this case than they were under the GSP rules. It's easy to show by induction that if the bids are the same, each bidder's GSP payment is at least as large than their VCG payment. For the last bidder who gets a slot the two payments are the same, and for any smaller i we have

$$c_i p_i^{VCG} - c_{i+1} p_{i+1}^{VCG} = (c_i - c_{i+1}) b_{i+1} \leq c_i b_{i+1} - c_{i+1} b_{i+2} = c_i p_i - c_{i+1} p_{i+1}$$

This is a justification (though weak) for using GSP, since the revenues to the search engine are guaranteed to be at least as large as in VCG.

We now consider a special kind of equilibrium, known as a "locally envy-free equilibrium" [6]. In this equilibrium, no bidder will want to swap bids with the bidder one position higher.

Consider the following strategy profile. Assume that if $i < j$ then $v_i \geq v_j$. Let $b_1 = v_1$ and for $i > 1$, $b_i = p_{i-1}^{VCG}$, that is, the payment per click of bidder $i-1$ under VCG if all advertisers bid truthfully. A result proven in [6] states that this strategy profile is a locally envy-free equilibrium; each advertiser's position and payment are the same as they would be in the equilibrium of VCG; and in any other locally envy-free equilibrium, the revenue of the search engine is at least as large as under this profile.

Using our example,

Table 6: A locally envy-free equilibrium

Advertiser	v_i	b_i	Slot	c_i	$p_i c_i$	Total payoff
Alice	50	50	1	10	210	290
Bob	40	21	2	5	10	190
Charlie	2	2	None	0	0	0

This result explains what happens if all bidders converge to the equilibrium strategy profile, but does not explain how they get here. The main result of [6] answers this question. It considers a generalized English (or "Japanese") auction which is analogous to GSP. The allocation rules are identical to GSP. The auction itself is run as follows: the price is continuously raised, and bidders can drop out at any time. Once they drop out, they cannot go back in and their bid is set to the current price. The strategy of bidder i can be described with a function $p_i(k, b, v_i)$, where p_i is the price at which he drops out; k is the number of other bidders left; and b is the bidding history of the bidders who dropped out before him. The theorem states that if p_i is continuous in v_i, there exists a unique ex-post perfect Bayesian equilibrium of this auction. In particular, each bidder i drops out when $p_i(k, b, v_i) = v_i - \frac{c_k}{c_{k-1}}(v_i - b_{k+1})$. In this equilibrium, the positions and payoffs are exactly the same as in the equilibrium of VCG. Since it's ex-post, the strategy of each advertiser wouldn't change if they knew exactly what the other advertisers' strategies were.

More general models are possible. Varian [12] analyzes the Weighted GSP where c_i are replaced with $c_{ij} = x_j c_i$, where x_j is the CTR component that results if the j-th advertiser takes the slot. He showed that in a globally envy-free equilibrium (that is, no bidder would want to swap bids with any other bidder) this more general GSP is also efficient and has revenue at least as high as VCG.

A natural question that arises is how often do these efficient equilibria occur. Thompson and Leyton-Brown answer this question in [11]. It was found that in practice, Weighted GSP outperforms both GSP and GFP in terms of efficiency, but is "a less clear winner in terms of revenue" [11].

5 Truthful mechanisms

It's often desirable to have a truthful mechanism. It would guarantee that no cycling behaviours would occur, as was the case with GFP. It also means there are no losses associated with investing into developing best strategies by advertisers, and corresponding costs to the search engine (i.e. prevention mechanisms that include not allowing to bid too often and restrictions on "bidding robots").

A VCG auction would seem as a natural choice; however, Aggarwal et al. [1] explain why VCG doesn't apply in the case when the ranking function orders bidders not by their bids, but by expected revenue (the ranking used by Google), as well as expose some other problems with VCG. Apart from these theoretical issues, it's unknown how well a VCG auction would work in practice, since it's harder for the customers to understand the payment rules than those of a GSP auction.

Aggarwal et al. instead present a simple truthful auction (the *laddered auction*) which is the unique auction for the general class of ranking functions. They also show that there's a pure-strategy Nash equilibrium for GSP that has exactly the same revenue as the laddered auction.

In the laddered auction, the bidder i in position j pays:

1. For the number of clicks that would be received in position $j+1$, same price as would be paid in position $j+1$.

2. For the additional clicks (the advantage of being in position j), the amount equal to the minimum bid required to stay in position j.

6 Conclusion

Ad keyword auctions are an interesting example of mechanisms that appeared in the industry first and were later studied and found to have good theoretical justification. Currently, GSP is the most widely used mechanism. While it's not truthful and it's not clear how often equilibria are reached in practice, empirically it's a good alternative to some other mechanisms, like GFP and VCG.

In practical applications, it's important to prevent non-truthful strategies [5]. This reduces the costs for everyone involved. That is why a truthful auction is desired in general. Even though GSP is not truthful, it was found to have many properties which make it a good choice.

On the other hand, it's not clear how a theoretically "better" but a more complicated mechanism would do in practice. Auction rules which are hard to understand might have a negative effect on the revenues. This might be another reason the simple GSP is so successful.

References

[1] Gagan Aggarwal, Ashish Goel, and Rajeev Motwani. Truthful auctions for pricing search keywords. In *EC '06: Proceedings of the 7th ACM conference on Electronic commerce*, pages 1–7, New York, NY, USA, 2006. ACM.

[2] Kursad Asdemir. Bidding patterns in search engine auctions. In *In Second Workshop on Sponsored Search Auctions (2006), ACM Electronic Commerce*. Press, 2006.

[3] Christian Borgs, Jennifer Chayes, Nicole Immorlica, Kamal Jain, Omid Etesami, and Mohammad Mahdian. Dynamics of bid optimization in online advertisement auctions. In *WWW '07: Proceedings of the 16th inter-*

national conference on World Wide Web, pages 531–540, New York, NY, USA, 2007. ACM.

[4] Matthew Cary, Aparna Das, Ben Edelman, Ioannis Giotis, Kurtis Heimerl, Anna R. Karlin, Claire Mathieu, and Michael Schwarz. Greedy bidding strategies for keyword auctions. In *EC '07: Proceedings of the 8th ACM conference on Electronic commerce*, pages 262–271, New York, NY, USA, 2007. ACM.

[5] Benjamin Edelman and Michael Ostrovsky. Strategic bidder behavior in sponsored search auctions. *Decis. Support Syst.*, 43(1):192–198, 2007.

[6] Benjamin Edelman, Michael Ostrovsky, and Michael Schwarz. Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords. *American Economic Review*, 97(1):242–259, March 2007.

[7] Juan Feng, Hemant K. Bhargava, and David Pennock. Comparison of allocation rules for paid placement advertising in search engines. In *ICEC '03: Proceedings of the 5th international conference on Electronic commerce*, pages 294–299, New York, NY, USA, 2003. ACM.

[8] Sébastien Lahaie. An analysis of alternative slot auction designs for sponsored search. In *EC '06: Proceedings of the 7th ACM conference on Electronic commerce*, pages 218–227, New York, NY, USA, 2006. ACM.

[9] Mohammad Mahdian, Hamid Nazerzadeh, and Amin Saberi. Allocating online advertisement space with unreliable estimates. In *EC '07: Proceedings of the 8th ACM conference on Electronic commerce*, pages 288–294, New York, NY, USA, 2007. ACM.

[10] Aranyak Mehta, Amin Saberi, Umesh Vazirani, and Vijay Vazirani. Adwords and generalized online matching. *J. ACM*, 54(5):22, 2007.

[11] David R.M. Thompson and Kevin Leyton-Brown. Computational analysis of perfect-information position auctions. In *ACM Conference on Electronic Commerce (ACM-EC)*, 2009.

[12] Hal R. Varian. Position auctions. *International Journal of Industrial Organization*, 25(6):1163–1178, December 2007.