

Algorithmic Game Theory

Edited by

Noam Nisan, Tim Roughgarden, Éva Tardos, and Vijay Vazirani

Contents

1 Sponsored Search Auctions

S. Lahaie, D. Pennock, A. Saberi, R. Vohra

page 4

Attachment converted: Macintosh HD:overture-lvtravel-results.png (PNGf/prvw)
(00376E01)

1

Sponsored Search Auctions

Sébastien Lahaie, David M. Pennock, Amin Saberi, and Rakesh Vohra

Abstract

1.1 Introduction

Sponsored search is a form of advertising typically sold at auction where merchants bid for positioning alongside web search results. For example, when a user queries a web search engine like Google or Yahoo! for “iPod”, advertisers (Apple Computer, Best Buy, eBay, etc.) may bid to have their listings featured alongside the standard “algorithmic” search listings. The advertisements appear in a separate section of the page designated as “sponsored” above or to the right of the algorithmic results. The sponsored search results are displayed in a format similar to algorithmic results: as a list of items each containing a title, a text description, and a hyperlink to a corresponding web page. We call each position in the list a *slot*. Generally, advertisements that appear in a higher ranked slot (higher on the page) garner more attention and more clicks from users. Thus, all else being equal, merchants generally prefer higher ranked slots to lower ranked slots. Figure 1.1(a) shows an example layout of sponsored search results for the query “las vegas travel”. Figure 1.1(b) shows the advertisers’ bids in the corresponding auction.

Advertisers bid for placement on the page in an auction-style format where the larger their bid the more likely their listing will appear above other ads on the page. By convention, sponsored search advertisers generally pay *per*

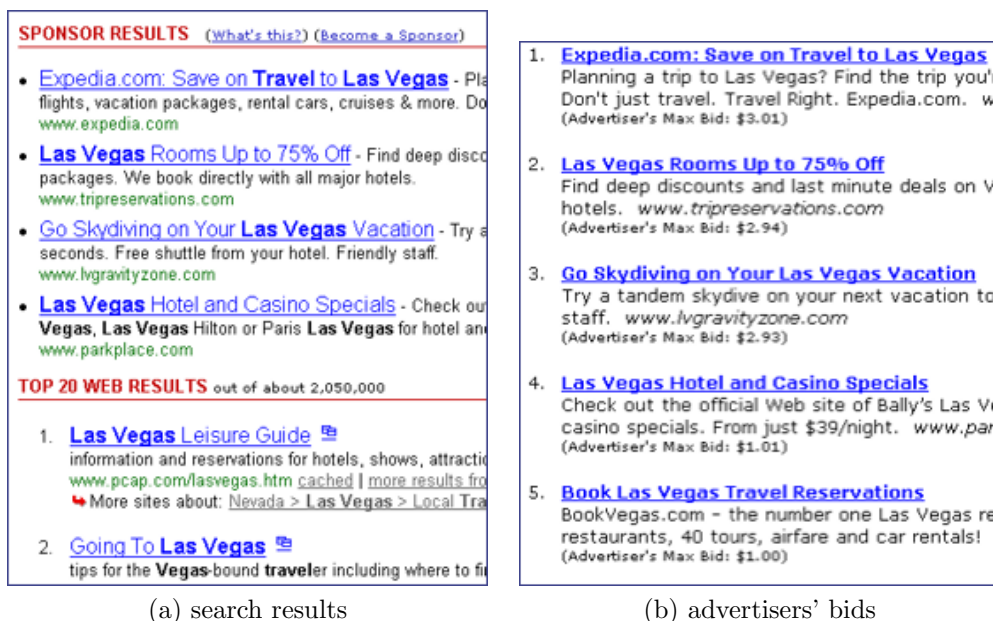


Fig. 1.1. (a) An example display of sponsored search listings above the regular algorithmic listings for the query “las vegas travel”. The ordering of sponsored listings is determined via a continuous auction mechanism. (b) The top five advertisers’ bids (maximum willingness to pay per click) in the auction.

click, meaning that they pay only when a user clicks on their ad, and do not pay if their ad is displayed but not clicked. Overture Services, formerly GoTo.com and now owned by Yahoo! Inc., is credited with pioneering sponsored search advertising. Overture’s success prompted a number of companies to adopt similar business models, most prominently Google, the leading web search engine today. Sponsored search is one of the fastest growing, most effective and profitable forms of advertising, generating roughly \$7 billion in revenue in 2005 after nearly doubling every year for the previous five years.

The sponsored search industry typically runs separate auctions for different search queries: for example, the queries “plasma television” and “investment advice” are associated with two distinct auctions. The entity being sold in each auction is the right to appear alongside the results of that search query. As mentioned, bids are expressed as a maximum willingness to pay per click. For example, a forty-cent bid by HostRocket for “web hosting” means HostRocket is willing to pay up to forty cents every time a user clicks

on their ad. Advertisers may also set daily or monthly budget caps. In practice, hundreds of thousands of advertisers compete for positions alongside several millions of search queries every day. Generally the auctions are continuous and dynamic, meaning that advertisers can change their bids at any time, and a new auction clears every time a user enters a search query. In this way advertisers can adapt to changing environments, for instance by boosting their bids for the query “buy flowers” during the week before Valentine’s Day. The auctioneer (the search engine[†]) evaluates the bids and allocates slots to advertisers. Notice that, although bids are expressed as payments per click, the search engine cannot directly allocate clicks, but rather allocates *impressions*, or placements on the screen. Clicks relate only stochastically to impressions.

Advertising in traditional media (e.g. magazines and television) as well as banner advertising is typically sold on a per-impression basis, or according to the (estimated) number of people exposed to the ad, in part because of the difficulty of measuring and charging based on the actual effectiveness of the ad. Traditional (offline) advertising, and to a large extent banner advertising on the Web, is usually priced via an informal process of estimation and negotiation. The Web’s capability for two-way communication makes it easy to track some measures of effectiveness, in particular user clicks. Many advertisers, especially direct marketers looking to close a sale as opposed to brand advertisers, prefer to pay per click rather than per impression, alleviating some of the uncertainty inherent in an impression. More direct performance-based pricing is possible by charging per “action” or per conversion (sale) on the merchant’s site.

Search engines are an information gateway to many search and decision-making tasks. As a result, entire niche industries exist touting services to boost a Web page’s ranking on the popular search engines, in part by reverse engineering the search engines’ information retrieval algorithms. Research has shown that good placement on a search page leads to high traffic, and eventually an increased financial payoff. Paying for sponsored slots is an alternative means of obtaining prominent positioning. Sponsored search works because users often tolerate or even welcome targeted ads directly related to what they are actively searching for. For example, Majestic Research

[†] In the sponsored search industry, the auctioneer and search engine are not always the same entity. Both Google and Yahoo! syndicate ads on a variety of partner search engines and content sites, with revenue being shared.

reports that as many as 17% of Google searches result in a paid click, and that Google earns roughly nine cents on average for every search query they process.†

Sponsored search lies at the core of the business models of Internet giants such as Yahoo! and Google. A number of other companies—including LookSmart, FindWhat, InterActiveCorp (Ask Jeeves), and eBay (Shopping.com)—earn hundreds of millions of dollars of sponsored search revenue annually.

The goal of this chapter is to formally model and analyze various mechanisms used in this domain and to study potential improvements. In section 1.2, we briefly describe existing mechanisms used to allocate and price sponsored search advertisements. Subsequently in sections 1.3 and 1.4 we discuss formal models used to analyze the properties of these auctions. Section 1.5 discusses further extensions and open problems.

1.2 Existing Models and Mechanisms

Typically, in sponsored search mechanisms, the advertisers specify a list of pairs of keywords and bids as well as a total maximum daily or weekly budget. Then, every time a user searches for a keyword, an auction takes place amongst the set of interested advertisers who have not exhausted their budgets.

Focusing on a single auction, let n be the number of bidders and $m < n$ the number of slots. The search engine estimates α_{ij} , the probability that a user will click on the i^{th} slot when it is occupied by bidder j . The quantity α_{ij} is called a *click through rate* (CTR). It is usually presumed for all j that $\alpha_{ij} \geq \alpha_{i+1,j}$ for $i = 1, \dots, m - 1$.†

The search engine also assigns a weight w_j to each advertiser j . The weight can be thought of as a relevance or quality metric. If agent j bids b_j , his corresponding *score* is $s_j = w_j b_j$. The search engine allocates slots in decreasing order of scores, so that the agent with highest score is ranked first, and so on. We assume throughout that agents are numbered so that agent j obtains slot j . An agent pays per click the lowest bid necessary to retain his position, so that the agent in slot j pays s_{j+1}/w_j . This weighted bid ranking

† <http://battellemedia.com/archives/001102.php>

† The assumption that clickthrough rate decays monotonically with lower slots is a distinguishing feature of keyword auctions; in particular, it implies that all bidders prefer the first slot to the second, the second slot to the third, etc. This allows for more refined equilibrium analyses than in the more general multi-item case.

mechanism includes the two most prominent keyword auction designs that have been used in practice: Overture introduced a “rank by bid” mechanism ($w_j = 1$) whereas Google uses a “rank by revenue” mechanism ($w_j = \alpha_{1j}$). Both variants are sometimes called *generalized second price* (GSP) auctions. Prior to 2004, Yahoo! used what is now known as a *generalized first price* (GFP) auction. Agents are ranked by bid but each bidder who secures a slot pays their bid per click.

1.3 A Static Model

The most popular model used to analyze keyword auctions is a static one where the private information of bidder j , the expected payoff from a click, v_j , is one dimensional. The expected payoff to a bidder from not obtaining a slot is assumed to be zero.

Four features of the model deserve comment. The first is its static nature. A substantial departure from reality. Since the study of recurrent auctions is rather daunting, one may be disposed to accept this simplification. Second, the expected payoff per click to a bidder is slot independent. This is tied to the assumption that all bidders prefer the top slot to the second slot to the third slot and so on. Observations from the field do not always support this assumption. It is believed that the probability of a click being converted into a purchase is lower in the top slot than in the second slot because many clicks on the top slot are made in error. A second story is that a searcher who clicks on a lower ranked slot is more serious in their intent to purchase. Third, a bidder’s value and CTR for a slot does not depend on the identity of other bidders. It seems plausible that Avis might value the fact that Hertz is not present in any slot when Avis is present. Fourth, the CTR’s are assumed to be common knowledge. In practice the CTR is estimated by the auctioneer and is not reported to the bidder. It is not clear that CTR’s can be estimated to a reasonable degree of accuracy, particularly by the bidder. The dynamic nature of the environment means that CTR’s can fluctuate dramatically over small periods.

As usual we assume that bidders are risk neutral and that their utility for a slot can be denominated on a common monetary scale. Supplied with copious amounts of salt, let us see where this model takes us.

1.3.1 Revenue Maximization and Efficiency

An auctioneer usually has one of two objectives: revenue maximization or allocative efficiency. In the static model one knows exactly what auction design will achieve either objective.

If the goal is revenue maximization, the classic result of Myerson (described in chapter 13) applies directly. One simply relabels the allocation variables. In chapter 13 section ??, the allocation variable, $x_j(b)$, is defined to be the expected quantity received by bidder i who bids b . For our setting, $x_j(b)$ becomes the expected click through rate for a bidder who bids b . Basically the generalized Vickrey auction is applied not to the actual values, v_j , but to the corresponding virtual values. The upshot is that the revenue maximizing auction is a generalized Vickrey auction with reserve prices.

If the goal is allocative efficiency, the generalized Vickrey auction will do the trick. The auction is described in chapters 9 and 11 of this book. The underlying problem of finding the efficient allocation in this case is an instance of the maximum weight assignment problem. For each slot i and bidder j let $x_{ij} = 1$ if bidder j is assigned to slot i and zero otherwise. The object is to choose x_{ij} 's to solve the following:

$$\max \sum_{i=1}^k \sum_{j=1}^n \alpha_{ij} v_j x_{ij} \quad (1.1)$$

$$\text{s.t.} \quad \sum_{j=1}^n x_{ij} \leq 1 \quad \forall i = 1, \dots, k \quad (1.2)$$

$$\sum_{i=1}^k x_{ij} \leq 1 \quad \forall j = 1, \dots, n \quad (1.3)$$

$$x_{ij} \geq 0 \quad \forall i = 1, \dots, k, \forall j = 1, \dots, n \quad (1.4)$$

This is equivalent to finding a maximum-weight perfect matching in a bipartite graph and hence can be solved in polynomial time. In fact, because the constraint matrix of this linear program is totally unimodular, it will have an optimal solution that is integral. Any feasible integer solution is called an assignment.

A single computation of the maximum weight assignment is sufficient to determine both the allocation and the generalized Vickrey payments. This is because the Vickrey payments lie in the dual to the above linear program. To write down the dual, let p_i be the dual variable associated with (1.2) and

q_j the dual associated with (1.3).

$$\min \sum_{i=1}^k p_i + \sum_{j=1}^n q_j \quad (1.5)$$

$$\text{s.t. } p_i + q_j \geq \alpha_{ij} v_j \quad \forall i = 1, \dots, k, \forall j = 1, \dots, n \quad (1.6)$$

$$p_i, q_j \geq 0 \quad \forall i = 1, \dots, k, \forall j = 1, \dots, n \quad (1.7)$$

Here p_i can be interpreted as the expected payment (CTR times price per click) of the bidder obtaining slot i , and q_j as the profit of bidder j . The objective in this program is to minimize the bidders' and auctioneer's profits combined. Amongst all optimal dual solutions pick the one that minimizes $\sum_{i=1}^k p_i$. The corresponding p_i is the price that the generalized Vickrey auction would set for slot i .

In the special case when the CTR's are bidder independent (i.e. $\alpha_{ij} = \mu_i$) there is a particularly simple algorithm, called the Northwest corner rule, to find the maximum weight assignment. Assign the bidder with the highest value per click to the top slot, the bidder with the second highest value per click to the second slot, and so on. In the Economics literature this is called an *assortative* assignment.

If one objects to the sealed bid nature of the generalized Vickrey auction there are ascending implementations available.

Interestingly, neither of these auctions correspond to the GFP or GSP auctions. In particular, bidding truthfully is not an equilibrium of either the GFP or GSP auctions. It is interesting to observe that Google's promotional material touts their auction as a modification of Vickrey's sealed bid auction for a single item (which it is) and concluding, therefore, that bidding sincerely is the correct thing to do (which it is not). A similar claim was made with respect to their auction used to sell shares of their IPO. They are not the first and quite possibly not the last to make such claims. For example, the financial services firm Hambrecht which pioneered the use of auctions to sell IPO's in 1998, says that their auction design is based on the Vickrey auction for a single good. While the Hambrecht auction does specialize to the Vickrey auction for a single good, it does not inherit the attractive properties of the Vickrey auction when applied to multiple units.†

To see why one must be careful when generalizing the Vickrey auction

† All of this reminds one of what is known as the freshman binomial theorem: $(a+b)^n = a^n + b^n$. True for $n = 1$ but not for $n > 2$.

to the sale of more than one unit, suppose there are three bidders with $v_1 > v_2 > v_3$ and two slots. Also suppose that $\alpha_{ij} = \mu_i$ with $\mu_1 > \mu_2$. If one were to auction off the top slot only, by an English ascending auction, each bidder would remain in as long as at the current price their surplus is non-negative. So, if the current price on the top slot is p_1 , bidder j remains active if $\mu_1(v_j - p_1) \geq 0$. Hence the auction ends at a price p_1 where $\mu_1(v_2 - p_1) = 0$, i.e., $p_1 = v_2$. Now suppose both slots are available but we will auction off the top slot first followed by the second slot. Let p_1 be the current price of slot 1, $p_2 = 0$ the current price of slot 2. Now bidder j will remain active in the auction for the top slot provided their surplus from the top slot is at least as large the surplus they could get from the second slot (which is currently priced at zero). That is

$$\mu_1(v_j - p_1) \geq \mu_2(v_j - 0) \Rightarrow p_1 \leq \left(1 - \frac{\mu_2}{\mu_1}\right)v_j.$$

Therefore the auction on the top slot terminates at a price of $\left(1 - \frac{\mu_2}{\mu_1}\right)v_2 < v_2$. The point is that the presence of a second slot lowers the price at which a bidder on the top slot will drop out of the auction on the top slot. The generalized Vickrey auction incorporates this change in the outside option of a bidder to ensure truthful bidding. The GSP auction does not. The generalized Vickrey auction, however, would allocate the top slot to bidder 1 and charge her $\left(1 - \frac{\mu_2}{\mu_1}\right)v_2$ and the second slot to bidder 2 and charge her v_3 .

As noted above, the GFP and GSP are special cases of what have been called ranking auctions. Bids (the reported v_j 's) are weighted (weights are independent of the bids) and then ranked in descending order. The highest ranked bidder gets the top slot, the second highest ranked bidder gets the second slot, and so on. The higher the bid the higher the slot one obtains (other bids held fixed). Since the assignment of bidders to slots is monotonic in the bid (other bids held fixed) it follows from standard results (see section 6 of chapter 9 for example) that there exists a payment rule that will make truthful bidding an equilibrium of the resulting auction. That payment rule is described, for example, in section ?? of chapter 13. Let $x_j(b|b_{-j})$ denote the expected click through rate for agent j when she bids b , given the profile of other bids is b_{-j} . Then the payment $P_j(b|b_{-j})$ she must make to ensure

incentive compatibility is given by

$$P_j(b|b_{-j}) = bx(b|b_{-j}) - \int_0^b x(t|b_{-j})dt. \quad (1.8)$$

These ranking auctions are, in general, neither efficient nor revenue maximizing. The payment rules associated with the GFP and GSP are not such as to induce truthful bidding as an equilibrium.

1.3.2 Equilibrium Properties

The fact that neither the GFP nor GSP are incentive compatible does not imply that they are inefficient or suboptimal in terms of revenue. It is possible that the equilibrium outcomes of both these auctions may be efficient or revenue maximizing. To identify the revenue and efficiency properties of these auctions it is necessary to determine their equilibria.

The GFP auction does not admit a pure strategy full information equilibrium but does admit a pure strategy Bayes-Nash symmetric equilibrium. The argument is identical to that of the sealed bid first price auction for a single good. The equilibrium bid functions are monotonic in the value. Therefore the equilibrium allocation of bidders to slots is the same as in the efficient allocation. Hence, by the revenue equivalence theorem, the symmetric equilibrium is efficient.

The efficiency of the GFP (in a Bayesian setting) lends it some appeal but this is where the ‘static’ assumption has bite. In a dynamic setting, the absence of a pure strategy full information equilibrium encourages bidders to constantly adjust their bids from one period to the next. This produces fluctuations in the bids over time and it has been argued that these fluctuations resulted in significant inefficiencies.

To date nothing is known about the Bayesian equilibrium of the GSP auction. Assume for simplicity that CTR’s are bidder-independent, so $\alpha_{ij} = \mu_i$, and that all weights are set to 1. †

In this case one can show that the GSP is efficient under full information and a restricted notion of equilibrium called **locally envy-free**. An assignment x is called locally envy-free if there exist prices, $\{p_i\}$, one for each slot,

† The analysis in this section generalizes straightforwardly to the case where CTR’s are separable (i.e. $\alpha_{ij} = \mu_i\beta_j$) and agents are assigned arbitrary weights w_j . These extensions are developed in the exercises.

such that for all i, j with $x_{ij} = 1$

$$\mu_i v_j - p_i \geq \mu_{i-1} v_j - p_{i-1} \quad (1.9)$$

and

$$\mu_i v_j - p_i \geq \mu_{i+1} v_j - p_{i+1} \quad (1.10)$$

In words, if bidder j is assigned to slot i then she prefers slot i to the slot just above her and the slot just below her.

Theorem 1.1 *An assignment x^* is optimal if and only if it is locally envy free.*

Proof Suppose first that x^* is locally envy-free and let p^* be the corresponding price vector. It suffices to prove that the assignment x^* is assortative. Let j be such that $x_{ij}^* = 1$ and j' such that $x_{i+1, j'}^* = 1$. To show that the assignment is assortative we must show that $v_j \geq v_{j'}$. From the property of being locally envy-free we have

$$\mu_i v_j - p_i^* \geq \mu_{i+1} v_j - p_{i+1}^*$$

and

$$\mu_{i+1} v_{j'} - p_{i+1}^* \geq \mu_i v_{j'} - p_i^*.$$

Adding them together yields

$$(\mu_i - \mu_{i+1})(v_j - v_{j'}) \geq 0.$$

Since $\mu_i \geq \mu_{i+1}$ it follows from this inequality that $v_j \geq v_{j'}$.

Now let x^* be an optimal assignment. Let (p^*, q^*) denote an optimal dual solution. It suffices to show that (x^*, p^*) is locally envy-free. Consider a pair (r, j) such that $x_{rj}^* = 1$. Complementary slackness and dual feasibility implies that $\mu_r v_j - p_r^* = q_j^* = \max_i \{\mu_i v_j - p_i^*\}$. Therefore

$$\mu_r v_j - p_r^* \geq \max\{\mu_{r-1} v_j - p_{r-1}^*, \mu_{r+1} v_j - p_{r+1}^*\}.$$

□

Theorem 1.2 *The GSP has a full information equilibrium that yields an allocation that is locally envy-free.*

Proof Order the bidders so that $v_1 \geq v_2 \geq \dots \geq v_n$. Let p_i^* be the Vickrey price of slot i . Let bidder 1 bid $b_1 = v_1$ and each bidder $j \geq 2$ bids $b_j = \frac{p_{j-1}^*}{\mu_{j-1}}$. First we show that under the rules of the GSP, bidder 1 is assigned to slot 1, bidder 2 to slot 2, and so on. To do this it suffices to show that $b_{j-1} \geq b_j$. Since the optimal assignment is locally envy-free we have

$$\mu_j v_j - p_j^* \geq \mu_{j-1} v_j - p_{j-1}^*.$$

Therefore

$$v_j - \frac{p_j^*}{\mu_j} \geq \frac{\mu_{j-1}}{\mu_j} v_j - \frac{p_{j-1}^*}{\mu_j}$$

which implies

$$b_{j-1} = \frac{p_{j-1}^*}{\mu_{j-1}} \geq \frac{p_{j-1}^*}{\mu_j} \geq \frac{p_j^*}{\mu_j} + \left(\frac{\mu_{j-1}}{\mu_j} - 1\right)v_j \geq \frac{p_j^*}{\mu_j} = b_j.$$

Hence if each bidder j bids b_j the GSP returns the optimal assignment. It is also easy to see that bidder $j \leq m$ pays p_j^* for their slot. Bidder $j > m$ pays zero. Since each bidder pays their Vickrey price and receives the slot they would have under the efficient allocation, no bidder has a unilateral incentive to change their bid. Therefore we have an equilibrium that, from Theorem 1, is envy-free. \square

Absent the recurrent nature of keyword auctions, they are similar to what are known as condominium auctions. In a condominium auction, bidders are interested in purchasing a condominium in a building. The condominiums are identical except for their height above the ground, the side of the building they are located on, etc. If all bidders have identical preferences over the condominiums, i.e., everyone prefers to be on a higher floor, they coincide with keyword auctions.

1.4 Dynamic Aspects

Since these auctions are repeated with great frequency one should properly model them as repeated games of incomplete information. The set of equilibria of such games is quite rich and complicated, even when restricted to the setting considered here. A full treatment of this case will not be given here. Rather we mention two phenomena that arise in this setting.

One is known as bid rotation. This occurs when competing bidders take turns at winning the auction. In our context this might mean bidders take

turns at occupying the top slot. If bidders are short lived this is unlikely to be a problem, if not, this will lower the auctioneers revenue.

Another possibility that repetition makes possible is vindictive bidding. In the GSP auction one's bid determines the payment of the bidder in the slot above and not one's own. Therefore one can increase the payment of the bidder in the slot above by raising one's bid without affecting one's own payment. This may be beneficial if the bidder in the slot above is a competitor with a limited budget for advertising. In a dynamic environment this encourages a bidder to constantly adjust their bids so as to inflict or avoid damage upon or from their competitor.

Even if one could ignore strategic considerations, a problem remains. The online nature of the auctions in sponsored search complicates the computation of an efficient allocation. Below we describe one model that addresses this difficulty.

1.4.1 The Online Allocation Problem

In this model, the search engine receives the bids of advertisers and their maximum budget for a certain period (e.g. a day). As users search for these keywords during the day, the search engine assigns their ad space to advertisers and charges them the value of their bid for the impression of the ad.† For simplicity of notation we assume that each page has only one slot for ads. The objective is to maximize total revenue while respecting the budget constraint of the bidders. Note that in this model bidders pay their bid which is counter to practice. On the other hand, budget constraints that apply across a set of keywords, a real world feature, are part of the model.

Let n be the number of advertisers and m the number of keywords. Suppose advertiser j has a bid of b_{ij} for keyword i and a total budget of B_j . In this context, it is reasonable to assume that bids are small compared to budgets i.e. $b_{ij} \ll B_j$.

If the search engine has an accurate estimate of r_i , the number of people searching for keyword i for all $1 \leq i \leq m$, then it is easy to approximate the optimal allocation using a simple linear program. Let x_{ij} be the total number of queries on keyword i allocated to bidder j . The linear program

† If one scales the bids by the CTR, the model would accommodate pay per click.

is

$$\begin{aligned}
\max \quad & \sum_{i=1}^m \sum_{j=1}^n b_{ij} x_{ij} & (1.11) \\
\text{s.t.} \quad & \sum_{j=1}^n x_{ij} \leq r_i \quad \forall 1 \leq i \leq m \\
& \sum_{i=1}^m b_{ij} x_{ij} \leq B_j \quad \forall 1 \leq j \leq n \\
& x_{ij} \geq 0 \quad \forall 1 \leq i \leq m, \forall 1 \leq j \leq n \\
\min \quad & \sum_{j=1}^n B_j \beta_j + \sum_{i=1}^m r_i \alpha_i \\
\text{s.t.} \quad & \alpha_i + b_{ij} \beta_j \geq b_{ij} \quad \forall 1 \leq i \leq m, \forall 1 \leq j \leq n \\
& \beta_j \geq 0 \quad \forall 1 \leq j \leq n \\
& \alpha_i \geq 0 \quad \forall 1 \leq i \leq m
\end{aligned}$$

By complementary slackness, in an optimal solution, advertiser j is assigned to keyword i if $(1 - \beta_j)b_{ij} = \max_{1 \leq k \leq n} (1 - \beta_k)b_{ik}$. Using this property, the search engine can use the solution of the dual linear program to find the optimum allocation: every time a user searches for keyword i , the search engine allocates its corresponding ad space to the bidder j with the highest $b_{ij}(1 - \beta_j)$. In other words, the bid of advertiser j will be scaled down by $1 - \beta_j$.

Now β_j represents rate of change of the optimal objective function value of (1.11) for a sufficiently small change in the right hand side of the corresponding constraint. In other words, if advertiser j 's budget were to increase by Δ , the optimal objective function value would increase by $\beta_j \Delta$. Equivalently, it is the opportunity cost of consuming agent j 's budget. Hence, if we allocate keyword i to agent now we obtain an immediate 'payoff' of b_{ij} . However, this consumes b_{ij} of the budget, which imposes an opportunity cost of $\beta_j b_{ij}$. Therefore, it makes sense in the optimal solution to (1.11) to assign keyword i to j provided $b_{ij} - \beta_j b_{ij} > 0$.

In practice, a good estimate of the frequencies of all search queries is unavailable. Queries arrive sequentially and the search engine must instantly decide to allocate their ad space to bidders without knowledge of the future queries. Therefore, what is needed is a dynamic procedure for allocat-

ing bidders to keywords that are queried. We describe one such procedure and analyze its performance within the usual competitive ratio framework. Specifically, we compare the revenue achieved by a dynamic procedure that does not know the r_i 's in advance, with the revenue that could be achieved knowing the r_i 's advance. The revenue in this second case is given by the optimal objective function value of the program (1.11).

The obvious dynamic procedure to consider is a greedy one: amongst the bidders whose budgets are not exhausted, allocate the query to the one with the highest bid. It is easy to see that this approach is equivalent to setting all β_j 's to 0.

The greedy procedure is not guaranteed to find the optimum solution. It is easy to construct a simple example with two bidders and two keywords in which the revenue of the greedy algorithm is as small as half of the optimum revenue. For example, suppose two bidders each with a budget of \$2. Assume that $b_{11} = 2$, $b_{12} = 2 - \epsilon$, $b_{21} = 2$ and $b_{22} = \epsilon$. If query 1 arrives before query 2, it will be assigned to bidder 1. Then bidder 1's budget is exhausted. When query 2 arrives it is assigned to bidder 2. This produces an objective function value of $2 + \epsilon$. The optimal solution would assign query 2 to bidder 1 and query 1 to bidder 2 yielding an objective function value of 4. The problem with the greedy algorithm is that, unlike the solution to (1.11), it ignores the opportunity cost of assigning a query to a bidder.

One can prove that the revenue of greedy algorithm is at least half of the optimum revenue for any instance. In the standard terminology of online algorithms, the competitive ratio of greedy algorithm is $1/2$. Can one do better in terms of competitive ratio? Yes. One does so by trying to dynamically estimate the opportunity cost, i.e., the β_j 's, of assigning a query to a bidder. This has the effect of spreading the bidders expenditures over time. The effect is called "budget smoothing", and is a feature that some search engines offer their advertisers.

The following modification of the greedy algorithm adaptively updates the β_j 's as a function of the bidders spent budget. Let

$$\phi(x) = 1 - e^{x-1}.$$

The algorithm sets $\beta_j = 1 - \phi(f_j)$ where f_j is the fraction of the budget of bidder j which has been spent.

Algorithm 1. Every time a query i arrives, allocate its ad space to the

bidder j who maximizes $b_{ij}\phi(f_j)$, where f_j is the fraction of the bidder j 's budget which has been spent so far.

The revenue of this algorithm is at least $1 - 1/e$ of the optimum revenue. It is also possible to prove that no deterministic or randomized algorithm can achieve a better competitive ratio.

Theorem 1.3 *The competitive ratio of Algorithm 1 is $1 - 1/e$*

We outline the main ideas in the proof of the theorem. Let k be a sufficiently large number used for discretizing the budgets of the bidders. We say that an advertiser is of type j , if she has spent within $(\frac{j-1}{k}, \frac{j}{k}]$ fraction of her budget so far. Let s_j be the total budget of type j bidders. For $i = 0, 1, \dots, k$, define w_i to be the amount of money spent by all the bidders from the interval $(\frac{i-1}{k}, \frac{i}{k}]$ of their budgets. Also define the discrete version of function ϕ ,

$$\Phi(s) = 1 - \left(1 - \frac{1}{k}\right)^{k-s}. \quad (1.12)$$

It is easy to see that when k tends to infinity $\Phi(s) \rightarrow \phi(\frac{s}{k})$. Let OPT be the solution of the optimal off-line algorithm (i.e. the solution of the optimization program (1.11)). For simplicity assume that the optimal algorithm spends all of the budget of the bidders. We have the following lemma:

Lemma 1.4 *At the end of the algorithm, this inequality holds:*

$$\sum_{i=0}^k \Phi(i) s_i \leq \sum_{i=0}^k \Phi(i) w_i \quad (1.13)$$

Proof Consider the time that query q arrives. Suppose OPT allocates q to a bidder of current type t whose type at the end of the algorithm will be t' . Let b_{opt} and b_{alg} be the amount of money that OPT and the algorithm gets from bidders for q . Let i be the type of the bidder that the algorithm allocates the query. We have:

$$\Phi(t') b_{opt} \leq \Phi(t) b_{opt} \leq \Phi(i) b_{alg}. \quad (1.14)$$

Now summing the inequality above over all the queries, the left hand side of (1.14) contributes to the sum $\sum_i \Phi(i) s_i$, and the right hand side contributes to $\sum \Phi(i) w_i$. So the lemma follows. \square

Now, we are ready to prove the Theorem 1.3.

Proof By definition $w_i \leq \frac{1}{k} \sum_{j=i}^k s_j$. Using lemma 1.4,

$$\sum_{i=0}^k \Phi(i) s_i \leq \frac{1}{k} \sum_{i=0}^k \Phi(i) \sum_{j=i}^k s_j.$$

Changing the order of the sums and computing the sum of the geometric series, we have:

$$\begin{aligned} \sum_{i=0}^k \Phi(i) s_i &\leq \frac{1}{k} \sum_{i=0}^k \Phi(i) \sum_{j=i}^k s_j \\ &\leq \frac{1}{k} \sum_{i=0}^k \left(\sum_{j=0}^i \Phi(j) \right) s_i \\ &\leq \sum_{i=0}^k \left(\frac{i}{k} + \Phi(i) - \Phi(0) + O\left(\frac{1}{k}\right) \right) s_i \\ &\leq \sum_{i=0}^k \frac{i}{k} s_i - \left(\Phi(0) - O\left(\frac{1}{k}\right) \right) \sum_{i=0}^k s_i + \sum_{i=0}^k \Phi(i) s_i \end{aligned}$$

which yields

$$\left(\Phi(0) - O\left(\frac{1}{k}\right) \right) \sum_{i=0}^k s_i \leq \sum_{i=0}^k \frac{i}{k} s_i.$$

Note that as k goes to infinity the left hand side tends to $(1 - \frac{1}{e})OPT$. The right hand side is equal to the revenue of the algorithm. So the theorem follows. \square

The same algorithm can be applied when multiple ads can appear with the result of a query or when advertisers enter at different times. At present, the equilibrium properties of this allocation rule are unknown.

1.5 Open Questions

We close this chapter with a brief review of important issues not directly addressed in this chapter.

While our discussion has focused on existing mechanisms one should not conclude that there is no room for improvement in their design. For example, there is debate over the role of the budget constraints in these auction. In many cases they don't appear to be hard constraints as bidders frequently adjust them. A bidder can also 'expand' their budget simply by lowering their bid and paying less per click. Some argue that the budget constraint is merely a convenient way to express other desires. For example, limiting ones exposure or spreading ones advertising over a longer period. All of this suggests the need for richer bidding models. Ones that might allow bidders to express decreasing marginal value for clicks, or distinct values for traffic from certain geographic regions, demographic profiles, etc., support greater allocative efficiency, though pose a significant burden in terms of computational and elicitation costs.

Unhappiness has also been expressed about paying per click. This forces search engines to invest in the task of detecting and ignoring robot clicks, spam clicks as well as clicks from an advertiser trying to impose costs on their competitor or from an affiliate who actually benefits monetarily from additional clicks. For this reason there is interest in exploring alternate pricing conventions. The most compelling is pay per activity. The advertiser pays only if a click results in a sale, for example. This raises new incentive issues associated with tracking sales.

The models in this chapter as do most analyses in the literature, assume a monopoly search engine with a static user base. This would be an appropriate model if switching costs for advertisers and users were high. In fact, switching costs for many advertisers are low; many advertisers work with both Google and Yahoo! simultaneously, or work with third-party *search engine marketers* to manage their account across multiple search engines. † Switching costs for users are essentially zero: to patronize a different search engine, users need merely type a new address into their web browser.‡ The competitive pressures to retain advertisers able to switch ad networks or use multiple networks may cause firms to focus less on extracting the maximum revenue from advertisers possible and more on attracting and retaining

† Market share in the sponsored search industry is not necessarily protected by a "network effect" as in the case of eBay, since advertisers can easily utilize multiple sponsored search providers, unlike the seller of a physical object who must choose a single venue.

‡ Personalization features may begin to introduce moderate switching costs for users. For now, reputation and branding seem to play a major role in search engine loyalty, as blind relevance tests show little or no difference between major search engines in terms of results quality.

advertisers. Similarly, search engines must make tradeoff decisions between maximizing current period revenue and attracting and retaining users in the long term. For this reason it would be very instructive to understand the properties of keyword auctions in competition with each other.

1.6 Bibliographic Notes

The growth of paid placement has attracted recent research on this topic. Hoffman and Novak [HN00] discuss the trend in Internet advertising towards per-click pricing rather than the traditional per-impression model. A good discussion of the practice of sponsored search is available on the web at <http://searchenginewatch.com/webmasters/paid.html>.

Computing the explicit form of incentive compatible payments for ranking auctions is carried out in [AGM06] and [IK06]. The Bayesian equilibrium of the GFP is derived in see [Lah06]. The details of the revenue maximizing auction for (static) slot auctions is derived in [Fen05] and [IK06]. The envy-free analysis of the static model is due to Edelman et al. [EOS]. A similar analysis can be found in Varian [Var]. The latter paper shows how upper and lower bounds on bidders' actual values can be derived given their bids. All of these results would apply to condominium auctions as well. See [Bur05] for a discussion of condominium auctions.

The Northwest corner rule for the assignment problem dates back to Monge [Mon81]. Ascending implementations of the Vickrey auction for the static model can be found in Crawford and Knoer [CK81] and Demange, Gale, and Sotomayor [DGS86] (which is a variant of the Hungarian algorithm for solving the assignment problem). The auction of Demange, Gale, and Sotomayor was dubbed, in [EOS], the generalized English auction.

The online allocation problem studied in Section 1.4.1 is proposed and analyzed by Mehta et al. [MSVV05]. This problem is a generalization of the online bipartite matching problem studied by Karp et al. [KVV90] and Kalyanasundaram and Pruhs [KP00]. More recently Buchbinder et al. [BJN06] gave a primal-dual algorithm and analysis for the problem given in Mehta et al. They also extended that framework to scenarios in which additional information is available, yielding improved worst case competitive factors.

Mahdian et al [MNS06] study the online allocation problem when the

search engine has a somewhat reliable estimate of the number of users searching for a keyword everyday. Mahdian and Saberi [MS06] study multi-unit auctions for perishable goods, in a setting where the supply arrives online. They motivate their model by its application to sponsored search. Abrams [Abr06] and Borgs et al. [BCI⁺05] design multi-unit auctions for budget-constrained bidders, which can be interpreted as slot auctions, with a focus on revenue optimization and truthfulness. For a discussion of vindictive bidding and some of the dynamic aspects of slot auctions see [Asd06] and [ZL06].

Weber and Zheng [WZ06] study the implementation of paid placement strategies, and find that the revenue-maximizing search engine design bases rankings on a weighted average of relative quality performance and bid amount. Hu [Hu03] uses contract theory to show that performance-based pricing models can give the publisher proper incentives to improve the effectiveness of advertising campaigns. Rolland and Patterson [RP03] propose a methodology using expert systems to improve the matching between advertisers and Web users.

Besides the optimal ranking mechanism, the search engine must also choose the number of paid slots by finding the optimal tradeoff between sponsorship and user retention. Bhargava and Feng [BF02] provide a theoretical model to explain and analyze this tradeoff.

Exercises

- 1.1 Consider the model of keyword auctions where the CTR of agent j in slot i is μ_i . Is every full information equilibrium of the GSP locally envy-free?
- 1.2 Consider the model of keyword auctions where the CTR of agent j in slot i is $\mu_i\beta_j$; that is, the CTR is separable into a bidder effect β_j and a position effect μ_i . Suppose also that $\mu_1 > \mu_2 > \dots > \mu_m$. Give a simple algorithm for determining the efficient allocation of bidders to slots. Derive the payment rule implied by the VCG mechanism for this environment.
- 1.3 In the model of the previous exercise, suppose also that the auctioneer assigns a weight $w_j \equiv w_j(\beta_j)$ to each bidder; weights may depend on the bidder effects, but not on their bids. Suppose bidders

are assigned to slots by decreasing order of their scores $w_j b_j$. Use formula (1.8) to derive the payment rule that combined with the allocation rule just described would yield an incentive compatible mechanism.

- 1.4 Consider the model of keyword auctions where the CTR of agent j in slot i is $\mu_i \beta_j$; that is, the CTR is separable into a bidder effect β_j and a position effect μ_i . The auctioneer sets weights $w_j = \beta_j$, and a bidder pays the lowest amount necessary to retain his position.
- (a) Give the inequalities that characterize a full-information (Nash) equilibrium in this model. Strengthen them to give the inequalities for a locally envy-free equilibrium.
 - (b) Show that in a locally envy-free equilibrium, bidders are ranked in order of decreasing $\beta_j v_j$.
 - (c) From amongst the set of locally envy-free equilibria, exhibit the one that yields the smallest possible revenue to the auctioneer.
- 1.5 Consider the model of keyword auctions where the CTR of agent j in slot i is μ_i . Give an example of where the GFP auction does not admit a pure strategy full information equilibrium. For simplicity, you may assume a discretized set of allowable bids.
- 1.6 Consider the online allocation problem discussed in section 1.4. Show that the competitive ratio of the algorithm remains the same even if the optimum solution does not exhaust all the budgets.

Bibliography

- [Abr06] Zoë Abrams. Revenue maximization when bidders have budgets. In *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms*, Miami, FL, 2006.
- [AGM06] Gagan Aggarwal, Ashish Goel, and Rajeev Motwani. Truthful auctions for pricing search keywords. In *Proceedings of the 7th ACM Conference on Electronic Commerce*, Ann Arbor, MI, 2006.
- [Asd06] Kursad Asdemir. Bidding patterns in search engine auctions. In *Proceedings of the 2nd Workshop on Sponsored Search Auctions*, Ann Arbor, MI, 2006.
- [BCI⁺05] Christian Borgs, Jennifer Chayes, Nicole Immorlica, Mohammad Mahdian, and Amin Saberi. Multi-unit auctions with budget-constrained bidders. In *Proceedings of the 6th ACM Conference on Electronic Commerce*, Vancouver, BC, 2005.

- [BF02] Hemant K. Bhargava and Juan Feng. Preferential placement in internet search engines. In *Proceedings of the 11th World Wide Web Conference*, Honolulu, HI, 2002.
- [BJN06] Niv Buchbinder, Kamal Jain, and Joseph (Seffi) Naor. Online primal-dual algorithms for maximizing ad-auctions revenue. Preprint, 2006.
- [Bur05] Roberto Burguet. The condominium problem; auctions for substitutes. *Review of Economic Design*, 9, 2005.
- [CK81] Vincent P. Crawford and Elsie Marie Knoer. Job matching with heterogeneous firms and workers. *Econometrica*, 49(2):437–450, 1981.
- [DGS86] Gabrielle Demange, David Gale, and Marilda Sotomayor. Multi-item auctions. *Journal of Political Economy*, 94(4):863–872, 1986.
- [EOS] 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*, forthcoming.
- [Fen05] Juan Feng. Optimal mechanism for selling a set of commonly ranked objects. Working paper, University of Florida, February 2005.
- [HN00] Donna L. Hoffman and Thomas P. Novak. How to acquire customers on the Web. *Harvard Business Review*, 78(3), May–June 2000.
- [Hu03] Yu J. Hu. Performance-based pricing models in online advertising. Technical report, Sloan School of Management, MIT, 2003.
- [IK06] Garud Iyengar and Anuj Kumar. Characterizing optimal keyword auctions. In *Proceedings of the 2nd Workshop on Sponsored Search Auctions*, Ann Arbor, MI, 2006.
- [KP00] Bala Kalyanasundaram and Kirk R. Pruhs. An optimal deterministic algorithm for online b -matching. *Theoretical Computer Science*, 233(1–2):319–325, 2000.
- [KVV90] Richard Karp, Umesh Vazirani, and Vijay Vazirani. An optimal algorithm for online bipartite matching. In *Proceedings of the 22nd ACM Symposium on Theory of Computing*, Baltimore, MD, 1990.
- [Lah06] Sébastien Lahaie. An analysis of alternative slot auction designs for sponsored search. In *Proceedings of the 7th ACM Conference on Electronic Commerce*, Ann Arbor, MI, 2006.
- [MNS06] Mohammad Mahdian, Hamid Nazerzadeh, and Amin Saberi. Allocating online advertisement space with unreliable estimates. preprint, 2006.
- [Mon81] Gaspard Monge. *Sur la théorie des déblais et des remblais*. Mémoires de l’académie de Paris, 1781.
- [MS06] Mohammad Mahdian and Amin Saberi. Multiunit auctions with unknown supply. In *Proceedings of the 7th ACM Conference on Electronic Commerce*, Ann Arbor, MI, 2006.
- [MSVV05] Aranyak Mehta, Amin Saberi, Umesh Vazirani, and Vijay Vazirani. AdWords and generalized on-line matching. In *Proceedings of the 46th Annual Symposium on Foundations of Computer Science*, 2005.
- [RP03] Erik Rolland and Raymond A. Patterson. Classification in online pay-for-performance advertising. In *Proceedings of The Thirteenth Annual Workshop On Information Technologies And Systems*, Seattle, WA, 2003.
- [Var] Hal R. Varian. Position auctions. *International Journal of Industrial Organi-*

zation, forthcoming.

- [WZ06] Thomas A. Weber and Zhiqiang Zheng. A model of search intermediaries and paid referrals. OPIM Working Paper 02-12-01, Wharton School, April 2006.
- [ZL06] Yunhong Zhou and Rajan Lukose. Vindictive bidding in keyword auctions. In *Proceedings of the 2nd Workshop on Sponsored Search Auctions*, Ann Arbor, MI, 2006.