MC758/MO758 - Algorithmic Game Theory

Sponsored search auctions

Erik de Godoy Perillo, RA135582

Abstract. *Search engines are very useful as an Internet tool today. Such platforms also generate big volumes of revenue by means of advertisement auctions. In this work, we present a brief overview on the topic, focusing on game theory principles underlying in current search systems. We review most used mechanisms today and some of their most important properties.*

Introduction

It is hard to imagine using the Internet without eventually reaching to some search engine at least once a day. Modern engines are an efficient tool to quickly obtain answers about almost any topic in such a way that it transformed human interactions with information in a variety of ways.

Companies such as Google, Bing (and older Yahoo) have built very lucrative businesses on top of search engines by means of advertisement. The fact that search engines are broadly used and user-provided queries are a very powerful piece of information for targetting make the platform very attractive to advertisers. In fact, Internet advertisement corresponds to more than 21% of all advertisement as of 2018, and search advertisement is responsible for 44% of the revenue of this type of advertisement, or approximately $26.1 billions of dollars. Companies such as Google and Facebook get the majority of their revenue from ads.

The question that naturally arises is: how do people pay for advertisement? The answer is that search engines apply what is known as sponsored search auctions, where advertisers bid for getting to appear in search queries related to specific terms. Game Theory plays an important role in this: there are multiple ways by which auctions can work regarding organization, ranking, pricing etc. In further sections of this document, we cover the most relevant topics and theoretic aspects of this interesting subject.

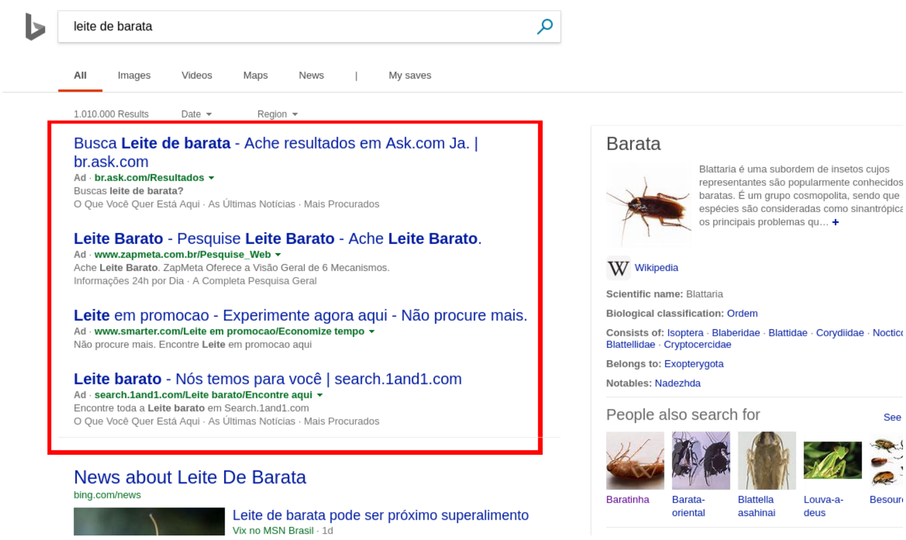

Figure 1. A typical search result with four slots being occupied by ads.

Key concepts

The auction procedure

In high-level terms, typical auctions are carried out in the following manner: Advertisers are the agents (or bidders) involved in sponsored search auctions. Before any auction takes place, agents specify terms (or keywords) they are interested in being related to. They also specify a (usually monthly or weekly) budget. Every search query issued by an user triggers an auction. Such query is related to a set of terms that are of interest to m agents (with enough budget) and there are k available slots (usually $k < m$). The search engine determines allocation – that is, what agents will be placed in the k slots and how – as well as the price that each agent will pay. Pricing is usually issued on a per-click basis, that is, advertisers only pay the specified price if the user clicks on their slot.

Click-Through Rate

Click-Through Rate (CTR) is the probability α_{ij} that the user is going to click on the i-th slot when it is occupied by bidder j. The CTR is estimated by the engine and is usually obtained via statistical methods and Machine Learning. An usual assumption is that $\alpha_{ij} > \alpha_{i+1,j}$ for every bidder j, that is, there's always a higher probability to be clicked on the topmost slots.

Ranking

The search engine must determine a way to give a score s_j and rank each bidder j since in most cases there are more bidders than available slots. It is typically done as follows: the bidder j offers to pay b_j for each click. The engine assigns a weight w_j to the bidder and the final score s_j is given as $s_j = w_j b_j$. Bidders are then ranked in decreasing order and allocated from top to bottom on the k available slots. Bidders not appearing in any slot pay nothing. A method to assign weights to bidders is rank by bid, where $w_j = 1$, that is, the final score of the bidder is their own bid. Another method is rank by revenue in which $w_j = \alpha_{1j}$. The latter is usually more reasonable since it attends to the expected value to be paid by the bidder.

Pricing

As mentioned before, bidders usually pay on a per-click basis. The price p_j each bidder j pays, however, is an important component of any auction mechanism. Two main pricing methods are applied: in the Generalized First Price (GFP), the price to be paid by the winner(s) on each click is their own bid. A more common pricing scheme is Generalized Second Price (GSP) in which the price to be paid by the bidder on slot i is the bid offered by the agent ranked on the position $i+1$ below. The GSP is today broadly used by companies such as Google and thus we focus on the GSP in following sections.

Table 1. Example of an auction with $m = 3$ bidders and $k = 2$ slots using rank-by-revenue and applying GFP and GSP.

j	b_j	w_j	s_j	Allocation	p_j (GFP)	p_j (GSP)
1	10	0.10	1.0	2	10	5
2	5	0.50	2.5	1	5	2
3	50	0.01	0.5	-	0	0

In example illustrated by table 1, an auction with 3 bidders and 2 slots is simulated. Using rank-by-revenue, the first slot is given to bidder $j = 2$ even though it doesn't bid the highest value. If the GFP is applied, each bidder pays their own bid. The GSP makes sure the expected value paid by each bidder is the expected value bid by the bidder ranked below: Bidder $j = 2$, for instance, pays $p_2 = b_1 w_1 / w_2 = 10 * 0.1/0.5 = 2$.

First and second price auctions

It's fairly straightforward to see why in sealed-bid – i.e. static, only one bid for each agent and all bids are revealed at once – auctions, first-price methods are not very convenient. Let v_k be the real valuation a player k gives to the item being auctioned. Let b_k be the bid it places and p_k the price it pays in the end. The utility u_k is then $u_k = v_k - p_k$ if it wins the auction and 0 otherwise. If the price to be paid by the winner is the own winner's bid, the utility is given by $u_k = v_k - b_k$. This implies that, in order to have a positive utility, the bidder must necessarily bid less than the real value it gives to the item. Thus, the first-price sealed-bid single-item auction is not strategy-proof: agents can ''lie" in order to achieve higher utility. It's also easy to see that this results extend naturally to the GFP: In example 1, if agent $j = 1$ is bidding their real valuation, their utility per click is $10 - 10 = 0$.

Second-price sealed-bid single-item auctions, however, are incentive compatible (or strategy-proof): bidding your real valuation for the item is the best strategy, maximizing the bidder j's utility b_j. Let $j \in [n]$ be a bidder with real item valuation v_j. Let $b_k = max\left(\{b_i : i \in [n] \setminus \{j\}\}\right)$. Now, there are two options for b_k: If $b_k > v_j$, j can bet $b_j > b_k$ and get utility $u_j = v_j - b_j < 0$. j can also bet $b_j \leq b_k$ and get utility $u_j = 0$. On the other hand, if $b_k \leq v_j$ j bets $b_j \geq b_k$ and gets $u_j = v_j - b_k \geq 0$ or it bets $b_j < b_k$, getting $u_j = 0$. We can conclude that betting $b_j = v_j$ maximizes utility. This does not imply, however, that multi-items second-price auctions are all strategy-proof. Consider the english ascending auction with three bidders and two slots: $v_1 \geq v_2 \geq v_3$ and $\alpha_{ij} = \mu_i$ and $\mu_1 > \mu_2$. When auctioning the two slots and starting from the top slot, bidder j keeps interest while $\mu_1(v_j - p_1) \geq \mu_2(v_j - 0) \implies p_1 \leq \left(1 - \frac{\mu_2}{\mu_1}\right) v_j$. So bidder 2 stops when $p_1 = \left(1 - \frac{\mu_2}{\mu_1}\right) v_2$. In this moment, auction stops and bidder 1 pays $p_1 = \left(1 - \frac{\mu_2}{\mu_1}\right) v_2 < v_2$! So bidder 2 could have lied a higher valuation in order to try to make bidder 1 give up on the item. This shows that second price not necessarily implies in the mechanism being strategy-proof.

Generalized Second Price auction properties

Desirable properties

If the goal of the mechanism is to maximize **allocative efficiency**, one can convert the allocation problem into a linear program. Let $x_{ij} = 1 \iff$ bidder j is in slot i, v_j be the bidder j's value and α_{ij} the CTR of bidder j on slot i. Then we obtain the allocation from the following problem (and pricing p_i from the dual):

$$max \sum_{i=1}^{k} \sum_{j=1}^{n} \alpha_{ij} v_j x_{ij}, \quad s.t. \qquad \begin{aligned} \sum_{j=1}^{n} x_{ij} &\leq 1, \ \forall i = 1, \ldots, k \\ \sum_{i=1}^{k} x_{ij} &\leq 1, \ \forall j = 1, \ldots, n \\ x_{ij} \geq 0, \ \forall j = 1, \ldots, n, &\ \forall i = 1, \ldots, k \end{aligned}$$

$$min \sum_{i=1}^{k} p_i + \sum_{j=1}^{n} q_j, \quad s.t. \qquad \begin{aligned} p_i + q_j \geq \alpha_{ij} v_{ij}, \ \forall j = 1, \ldots, n, \ \forall i = 1, \ldots, k \\ p_i, q_j \geq 0 \ \forall j = 1, \ldots, n, \ \forall i = 1, \ldots, k \end{aligned}$$

If the CTRs are bidder-independent ($\alpha_{ij} = \mu_i$), a simpler approach is to sort the bidders in decreasing order of their values per click (assortative assignment).

If we aim at maximizing **revenue**, the Myerson mechanism applies. Let F_j and f_j be, respectively, the distribution and density functions for v_j. One can apply the well-known VCG mechanism to the virtual valuations $\phi_j(v_j) = v_j - \frac{1-F_j(v_j)}{f_j(v_j)}$. It is also ensured – from the fact that assignment of bidders is monotonic in the bid – that there is a payment method that ensures **incentive compatibility**. The problem is that one does necessarily guarantee the other: an incentive-compatible pricing is generally not revenue/efficiency maximizing. The usual GFP/GSP payment rules also do not induce truthful equilibrium.

Equilibrium properties

In spite of the usually unfortunate equilibrium properties mentioned earlier, the GSP can indeed be efficient if we assume a more restricted equilibrium called *locally envy-free*: it is the notion that bidders are ''locally satisfied" with their allocations. More precisely, for bidder j that bids v_j and slot i with price p_i and CTR μ_i, an assignment is locally envy-free if $\mu_i v_j - p_i \geq \mu_{i-1} v_j - p_{i-1}$ and $\mu_i v_j - p_i \geq \mu_{i+1} v_j - p_{i+1}$, that is, the bidder does not get a greater utility by moving to a slot up or down. We now show that efficiency holds under these circumstances.

Lemma 1: *An assignment x^* is optimal if and only if it's locally envy-free.*
Proof: Let's first assume that it is locally envy-free. Let p^* be the pricing. Let i, j, k such that $x^*_{ij} = 1$ and $x^*_{i+1,k} = 1$. Then we have that $\mu_i v_j - p^*_i \geq \mu_{i+1} v_j - p^*_{i+1}$ and $\mu_{i+1} v_k - p^*_{i+1} \geq \mu_i v_k - p^*_i$. If we add both expressions, we get:

$$(\mu_i - \mu_{i+1})(v_j - v_k) \geq 0$$

This implies that $v_j \geq v_k$ (since $\mu_i > \mu_{i+1}$). Thus, we have an assortative assignment, which is optimal. If we now suppose x^* is optimal, let (p^*, q^*) be an optimal dual solution to the allocation LP problem. Let r, j such that $x^*_{rj} = 1$. The properties of complementary slackness and dual feasibility imply that $\mu_u v_j - p^*_r = q^*_j = max\{\mu_i v_j - p^*_i : i \in [n]\}$ and thus

$$\mu_u v_j - p^*_r \geq max\{\mu_{r-1} v_j - p^*_{r-1}, \mu_{r+1} v_j - p^*_{r+1}\}$$

Which implies that assignment is locally envy-free. We have thus proved Lemma 1.

Theorem: *The GSP has a full information equilibrium that yields a locally envy-free and optimal allocation.*
Proof: Let there be bidders $1, \ldots, n$ sorted so that $v_1 \geq v_2 \geq \ldots \geq v_n$. This assortative assignment, with Vickrey (second) pricing p^*, is optimal as seen previously. We now show that, if bidders place a certain bet and we follow the GSP rules, we will obtain the optimal (and locally envy-free) allocation. Let such bet by each agent j be b_j such that $b_1 = v_1$ and $b_j = \frac{p^*_{j-1}}{\mu_{j-1}}$. If we show that $b_j \geq b_{j+1}$, this will imply that the GSP will order bidders $1, \ldots, n$ from top to bottom and, since $v_1 \geq v_2 \geq \ldots \geq v_n$, this is an assortative and optimal allocation. Indeed, since our initial assignment is optimal, it follows from Lemma 1 that it's locally envy-free: $\mu_j v_j - p^*_j \geq \mu_{j-1} v_j - p^*_{j-1}$, which implies:

$$v_j - \frac{p^*_j}{\mu_j} \geq \frac{\mu_{j-1}}{\mu_j} v_j - \frac{p^*_{j-1}}{\mu_j} \implies b_{j-1} = \frac{p^*_{j-1}}{\mu_{j-1}} \geq \frac{p^*_{j-1}}{\mu_j} \geq \frac{p^*_j}{\mu_j} + \left(\frac{\mu_{j-1}}{\mu_j} - 1\right) v_j \geq \frac{p^*_j}{\mu_j} = b_j$$

Hence indeed, $b_j \geq b_{j+1}$ and the GSP will produce the optimal assignment which is also locally envy-free. We have thus proved the theorem.

References

- Noam Nisan, Tim Roughgarden, Eva Tardos, e Vijay V. Vazirani. Algorithmic Game Theory, Cambridge University Press, 2007.
- Google. Challenge in the Design of Sponsored Search Auctions. 2007.
  URL: `https://youtu.be/JzRHiJdE0mk`
- PWC. IAB Internet advertising revenue report. 2017.