Creative ID Recommendation

AlephD

May 2014

1 Presentation

We study ad-auctions, auctions on which ad spaces are sold by publishers to advertisers in milliseconds. We focus on the simulation of the gain of publishers, should they have set a different floor price for a specific set of auctions.

1.1 Second price ad-auctions

We denote:

- by $\mathcal{B} = \{b_1, \ldots, b_B\}$ the set of advertisers (buyers or bidders);
- and by $\mathcal{A} = \{a_1, \ldots, a_A\}$ the set of successive auctions, happening at time

$$\{t_{a_1}, \ldots, t_{a_A}\}$$

The advertisers (also called bidders) are sent *bid requests* for every auction a. We assume they take part in all of them by simultaneously announcing a price. Note that a price of 0, can be considered equivalent to a decision not to participate to a specific auction. The bidder with the highest price for auction a[1] can display his ad and pays the maximum of the second price and the reserve price for the auction (provided his price is above the reserve price).

Let us consider an auction a happening at time t_a, let

$$p_a = (p_{a,b_1}, p_{a,b_2}, \ldots, p_{a,b_B})$$

be the bids of the B bidders in the auction a. We denote:

- by $\rho_a : \mathcal{B} \to \{1, \ldots, B\}$ the function which gives the rank of an advertiser in the auction a (1 will be the rank of the highest price);
- by $\beta_a : \{1, \ldots, B\} \to \mathcal{B}$ the function which gives the advertiser corresponding to a rank in the auction a. We have: $\beta_a = \rho_a^{-1}$.

We denote:

$$(p_{a,1}, p_{a,2}, \ldots, p_{a,B}) = \left(p_{a,\beta_a(1)}, p_{a,\beta_a(2)}, \ldots, p_{a,\beta_a(B)}\right)$$

[1] The case of equality is extremely rare. We assume it does not happen

1.2 The publisher

The publisher: s (seller) can act on an auction a by setting a reserve-price r_a before it happens. The payoff of s at auction a is a function of $p_{a,1}$, $p_{a,2}$ and r_a:

$$\pi_a(p_{a,1}, p_{a,2}, r_a) = \begin{cases} p_{a,2} & \text{if } p_{a,2} \geq r_a \\ r_a & \text{if } p_{a,1} \geq r_a > p_{a,2} \\ d_a & \text{if } r_a > p_{a,1} \end{cases}$$

Where d_a is the value received by the publisher for the display of a *default ad*. It can be assumed to be 0.

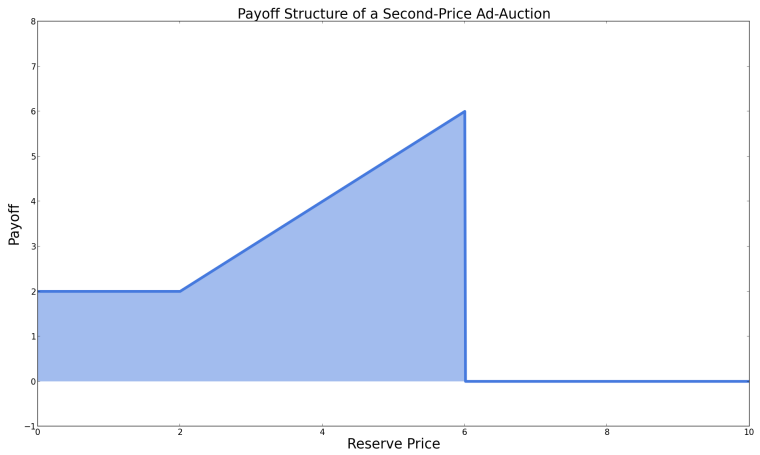

Figure 1: Payoff Structure of a Second-Price Ad-Auction

2 The objective

The advertiser tries to target his campaigns to the most relevants users and at the most relevant placements. Each time an auction occurs, an advertiser will have the option to bid for one of its campaigns to be sent to the user at a specific placement. Therefore, the advertiser must choose to which users and placements his campaigns should be targeted. The objective is to recommend users and placements that are similar to those already targeted by the advertiser for a given campaign.

2.1 The information available

To analyze the similarities between users and placements with respect to campaigns, we have logs that contain the following fields:

- The auction id a, that uniquely characterizes an auction, including the parameters below;

- The user id u, that characterizes the user that is loading the page;
- The tag id p, that characterizes the placement where the impression will appear;
- The timestamp t, unix epoch;
- The creative id c, that characterizes the ad that won the auction, it includes the advertiser, the campaign and the format of the ad;
- The first-bid b.

2.2 The problem

The objective is to retreive bundles (sets of items) that verify the condition of complementarity and maximize the similarity of the items it contains. The complementarity constraint means that two items in the same bundle should have different values for the same property. A property is a defined attribute for an item. The similarity is a function of two items that expresses how close these items are.

The hypotheses underlying our problem are the following:

- Advertisers target their campaigns using the information of user id and tag id, they need to select those that will maximize their revenue;
- The set of items $\mathcal{I}$ is the set of auction ids a;
- All other things being equal, two items with the common creative id c are more similar than these two items without this common creative id c;
- All other things being equal, two items with the common creative id c and the same first-bid b are more similar than these two items with this common creative id c but different first-bid b;
- The pair (user id, tag id) is the property of the item, a bundle should not contain more than one occurence of a pair;

Therefore we consider that when an advertiser sends the same campaign with the same format to a set of users and tags, these pairs of user and tag are similar. The property implies that we want to group different users but also different tags for the same user. It means that we are interested in finding other users that have similar characteristics as well as other placements where the same user will respond similarly to the campaign. The bundles will allow an advertiser to find for his campaign complementary users and placements that are similar to those that he has already targeted.