Electronic Commerce Research and Applications 13 (2014) 178–187

Contents lists available at ScienceDirect

Electronic Commerce Research and Applications

journal homepage: www.elsevier.com/locate/ecra

Optimal reserve prices in weighted GSP auctions ☆

Yang Sun [a,*], Yunhong Zhou [b], Xiaotie Deng [c]

[a] City University of Hong Kong, Kowloon, Hong Kong
[b] Microsoft Corporation, Redmond, WA, USA
[c] Shanghai Jiaotong University, Shanghai, PR China

ARTICLE INFO

Article history:
Received 30 April 2013
Received in revised form 16 February 2014
Accepted 18 February 2014
Available online 6 March 2014

Keywords:
Generalized second price auction
Weighted GSP
Reserve price
Optimal auction
Sponsored search auction

ABSTRACT

Most search engines use the weighted Generalized Second Price (wGSP) auction to sell keyword-based text ads, generating billions of dollars of advertising revenue every year. Designing and implementing near-optimal reserve prices for these wGSP auctions are naturally important problems for both academia and industry.

In this paper, we show how to calculate and implement the near-optimal reserve price of the wGSP mechanism in realistic settings. Unlike reserve prices in standard single-item auctions, optimal reserve prices in wGSP auctions are discriminatory, different even for advertisers bidding on the same keyword. The optimal reserve price results can be extended to support CPA/CPC/CPM[1] hybrid auctions.

Our simulations indicate that setting a proper reserve price will transfer some bidder utility (payoff) to auctioneer utility, resulting in higher revenue for the search engine. We describe a practical methodology to implement optimal reserve prices in production systems.

© 2014 Elsevier B.V. All rights reserved.

1. Introduction

Search advertising places keyword-based text ads alongside search results, generating billions of dollars of advertising revenue annually, and is the primary factor in the successful commercialization of search engines (Edelman et al. 2007, Varian 2007). Keyword-based text ads typically contain a title, a text description, and a display URL. Advertisers generate revenue when search engine users click on keyword ads and subsequently purchase advertised goods or services. Search providers conduct auctions to allocate advertising positions and decide per-click prices. Early keyword auctions implemented the Generalized First Price (GFP) mechanism, which was pioneered by Overture (later acquired by Yahoo!). The major stages of keyword auction evolution (Thompson and Leyton-Brown 2009) are as follows:

1. **GFP**: The unweighted Generalized First Price auction. Agents are ranked by their bids and each bidder who wins a slot pays her bid per click. This mechanism was first used by Overture, 1997–2002.
2. **uGSP**: The unweighted Generalized Second Price auction. Agents are ranked by their bids and each bidder who wins a slot pays the next highest bid per click. This mechanism was first used by Yahoo!, 2002–2007.
3. **wGSP**: The weighted Generalized Second Price auction. The mechanism assigns a weight w_i to each bidder; agent i is ranked by the product of her bid b_i and her weight w_i, winner i pays $p_i = b_{i+1} \cdot w_{i+1}/w_i$ per click (where b_{i+1} and w_{i+1} are the next highest bidder's bid and weight). This mechanism was first used by Google and adopted by all major search engines gradually.

The wGSP mechanism is widely used in the industry nowadays and naturally is the focus in both academia and industry. Instead of rank-by-bid in uGSP, we refer to wGSP as rank-by-revenue. When users search for the keyword, the search engine will rank the ads in descending order of the product of each advertiser's weight (ad quality) and bid (per-click). Under the wGSP mechanism, search engines charge advertisers according to their bids as well as their ad quality factors (primarily the estimated click-through rate). Our work focuses on wGSP auctions and our results allow the search engine to set discriminatory reserve prices for each advertiser and each keyword. Our settings are as follows: each bidder knows her private value and her expected click-through rate (eCTR) of the keyword, and only bidders who bid greater than their reserve prices can participate in the auction.

☆ A preliminary version of this paper was presented in the Seventh Ad Auctions Workshop, in conjunction with the 12th ACM Conference on Electronic Commerce, 2011.

* Corresponding author. Tel.: +1(403)3992839.
E-mail addresses: sun.yang@my.cityu.edu.hk (Y. Sun), yunhongz@microsoft.com (Y. Zhou), deng-xt@cs.sjtu.edu.cn (X. Deng).

[1] CPM: cost per mille (impressions), CPA: cost per action, CPC: cost per click.

http://dx.doi.org/10.1016/j.elerap.2014.02.003
1567-4223/© 2014 Elsevier B.V. All rights reserved.

1.1. Related work

Myerson (1981) proposed the general optimal auction framework in a Bayesian setting known as the classical Myerson optimal auction, which maximizes auctioneer revenue. Myerson proved that the optimal auction problem is equivalent to the problem of maximizing bidders' expected virtual surplus, which can be solved by an efficient auction with respect to virtual surplus under the technical condition of monotone non-decreasing hazard rate (MHR) on the private value distribution. Thus, Myerson converted the problem of optimal auction design into one of efficient auction design on a perturbed problem.

In the case of single-item auctions, the Myerson optimal auction is equivalent to the Vickrey auction with a reserve price. For sponsored search auctions, the goal is to compute the allocation and payment rules that result in a revenue optimal mechanism for the multi-item auction. This calls for extending the Myerson optimal auction to the case of the sponsored search auction.

Varian (2007) and Edelman et al. (2007) modeled the keyword auction problem as a Nash equilibrium problem with complete information about all the bids. They defined a refined Nash equilibrium called *locally envy-free equilibrium* (Edelman et al. 2007) or *symmetric Nash equilibrium* (Varian 2007). These two equilibria are actually equivalent: if each advertiser bids at an envy-free point, then she is exactly indifferent between remaining in her current position and trading places with a bidder above her.

Gomes and Sweeney (2009) characterized the efficient Bayes-Nash equilibrium in GSP auctions with a Bayesian setting. Their incomplete information assumption was different from the above-mentioned full-information setting. They assumed that each bidder can only estimate the overall bidders' value distribution, instead of knowing the details of others' bids. Their assumption was closer to actual conditions in keywords auctions, since most search engines run sealed-bid auctions nowadays and bidders cannot observe others' bids. Their results provided the foundation for our study: assume that bidders' values were independently drawn from a distribution F, each bidder privately observed her value v_i and simultaneously submitted her bid to the search engine following the same function $b_i = \beta(v_i, F)$.

Edelman and Schwarz (2010) generalized Myerson's theory to the multi-unit environment and simulated uGSP auctions. They proved that the uGSP auction with reserve price r^* is an optimal mechanism. Ostrovsky and Schwarz (2011) followed their work and conducted a large scale experiment at Yahoo! Search with a new reserve price system. Below are excerpts from their paper:

> "The theoretically optimal reserve prices were computed under the assumption that all bidders have the same quality factors and are ranked solely on the basis of their bids, which is a simplification. In practice, the ads on Yahoo! are ranked based largely on the product of their quality factors and bids, and the amount each advertiser pays is lower when his ad's quality is higher. Thus, in order to keep the implementation of reserve prices consistent with the company's ranking and pricing philosophy, the theoretical reserve prices were converted into advertiser-specific reserve prices that reflected the quality factors of the ads: ads with higher quality factors faced lower per-click reserve prices, and vice versa."
> *Note that this is a deviation from the theoretically optimal auction design with asymmetric bidders.*" (Ostrovsky and Schwarz 2011).

Their new reserve price was the average of the theoretically optimal reserve price and the original uniform reserve price (10 cents), and it was reported that the new reserve price improved the average revenue per search by almost 13%.

Liu et al. (2010) studied a scheme where the auctioneer can weigh advertisers' bids differently and require different minimum bids based on advertisers' quality factors. They studied the impact and design of such weighting schemes and minimum-bid policies. They proposed the revenue-maximizing optimal reserve price for wGFP auctions, where all bidders' quality factors e_i are discrete numbers. For example, n bidders bid on one keyword, and there are only 2 possible values of e_i: the h-type bidder with e_h, and l-type bidder with e_l. In the setting without considering the distortion effect, this requires the revenue-maximizing reserve prices to satisfy

$$\underline{b_l} - \frac{1 - F_l(\underline{b_l})}{f_l(\underline{b_l})} = 0 \quad \text{and} \quad \underline{b_h} - \frac{1 - F_h(\underline{b_h})}{f_h(\underline{b_h})} = 0 \qquad (1)$$

While Liu et al. (2010) proved the revenue equivalence between wGFP and wGSP under their settings, their results cannot extend to the setting where e_i follows a continuous distribution. In realistic wGSP auctions, bidders' quality factors e_i are real numbers from $(0, 1]$, and it is possible that all e_i are distinct from each other. For example, there are 30 bidders bidding on one keyword. Since there are 30 different e_i, 30 distributions F need to be estimated and 30 equations need to be solved. But there is only one sample for each distribution estimation, implementation is thus impossible. Even if we reduce those e_i to 10 discrete intervals approximately, it is hard to estimate well each distribution F from 3 samples on average. Obviously, these results could not work in production systems.

Recent related work of Thompson and Leyton-Brown (2013) analyzed a more general GSP ranking scheme, where bidders were ranked by $b_i \cdot q_i^{\alpha}$ with $\alpha \in [0, 1]$ a constant. They simulated the revenue optimal parameters (coefficient α and reserve prices) by searching all possible parameter values, for both single-slot and multi-slot cases.

1.2. Our contributions

Obviously, the optimal reserve prices in wGSP auctions were still an unsolved issue when Yahoo! did the field experiment, and the previous theoretically optimal reserve price is not directly applicable to the case where bidders have different quality factors. Thus, Ostrovsky and Schwarz (2011) implemented it with a compromise between the theoretically optimal uGSP results and practical wGSP auctions, even though it was a deviation from the theoretically optimal results.

In this paper, we show how to calculate provably optimal reserve prices in wGSP auctions where bidders have *different* quality factors drawn from a *continuous* interval (such as $(0, 1]$). This model is based on the realistic keyword auctions implemented by most search engines nowadays. We also show that this optimal reserve price can support CPA/CPC/CPM hybrid auctions.

The rest of the paper is organized as follows: Section 2 introduces the details of the wGSP mechanism and our model. Section 3 discusses how to compute the optimal reserve price, and extends this optimal reserve price to CPA/CPC/CPM hybrid auctions. Section 4 presents simulation results from estimated distributions, considering the search engine revenue, bidders' payments and utilities. Section 5 describes a practical technique to compute the optimal reserve price for each keyword and each advertiser.

2. Preliminaries

2.1. Ranking and pricing in wGSP auctions

There are k positions to be allocated among n bidders, with $k \leqslant n$. We denote the ad positions as $t \in \{1, \ldots, k\}$, and the bidders as $i \in \{1, \ldots, n\}$. We also denote v_i and b_i as the private value and bid of bidder i. For a specific bidder i, his valuation v_i for one click is identical through all positions.

Download English Version:

https://daneshyari.com/en/article/379665

Download Persian Version:

https://daneshyari.com/article/379665

Daneshyari.com