A Unified Optimization Framework for Auction and Guaranteed Delivery in Online Advertising

Konstantin Salomatin
Carnegie Mellon University
5000 Forbes Ave
Pittsburgh, PA, USA 15213
ksalomat@cs.cmu.edu

Tie-Yan Liu
Microsoft Research Asia
5 Danling St, Haidian District
Beijing, P.R. China 100080
tyliu@microsoft.com

Yiming Yang
Carnegie Mellon University
5000 Forbes Ave
Pittsburgh, PA 15213, USA
yiming@cs.cmu.edu

ABSTRACT

This paper proposes a new unified optimization framework combining pay-per-click auctions and guaranteed delivery in sponsored search. Advertisers usually have different (and sometimes mixed) marketing goals: brand awareness and direct response. Different mechanisms are good at addressing different goals, e.g., guaranteed delivery was often used to build brand awareness and pay-per-click auctions was widely used for direct marketing. Our new method accommodates both in a unified framework, with the search engine revenue as an optimization objective. In this way, we can target a guaranteed number of ad clicks (or impressions) per campaign for advertisers willing to pay a premium and enable keyword auctions for all others. Specifically, we formulate this joint optimization problem using linear programming and a column generation strategy for efficiency. To select the best column (a ranked list of ads) given a query, we propose a novel dynamic programming algorithm that takes the special structure of the ad allocation and pricing mechanisms into account. We have tested the proposed framework and the algorithms on real ad data obtained from a commercial search engine. The results demonstrate that our proposed approach can outperform several baselines in guaranteeing the number of clicks for the given advertisers, and in increasing the total revenue for the search engine.

Keywords

Online advertising, sponsored search, optimization, linear programming

1. INTRODUCTION

Sponsored search is an important means of Internet monetization, and is the driving force of major search engines today. It has sustained a market of tens of billions of dollars, and the market size is still growing very fast [4]. The success of sponsored search is in part due to its business model, and in part due to its strong technical foundation in information retrieval, data mining, machine learning, and algorithmic economics.

Currently, sponsored search works in the following manner. When a user submits a query to a search engine, in addition to the organic search results, a list of advertisements (ads) are also presented to the user. The advertiser is charged only if the user clicks the ad. The selection, ranking, and pricing of these ads are determined by an auction mechanism. The most popular mechanism is generalized second price (GSP) [3], where the ads are ranked according to their expected values (the predicted click probability of an ad times the bid price given by the advertiser), and advertisers are charged according to the second-price rule. That is, the cost per click is the minimum bid price required to keep a given ad in its current rank position.

Another popular marketing strategy in online advertising is called guaranteed delivery (GD) where ads that share a single idea and theme are grouped into campaigns, and charged on a pay-per-campaign basis for the prespecified number of deliveries (clicks or impressions). Most popular GD solutions are based on offline optimization algorithms, adjusted for online setup [5]. In [10] the problem is formulated as a suplly-demand Linear Program on a bipartite graph (supply corresponds to page visits, demand - to impression counts of the campaigns requested by the advertisers and the edge denotes relevance of a campaign for a particular page visit).

Comparing GD with the pay-per-click auction, the latter focuses on the winning of each individual auction (i.e., a micro-level optimization), while the former concentrates on the campaign-level (macro-level) success for grouped ad impressions or clicks. It has been commonly believed that the pay-per-click auction is more important for sponsored search, and that GD is more important in display advertising. Can the strengths of both strategies be combined in sponsored search? Furthermore, how to model both in a unified optimization framework? Those meaningful questions have not been studied before. Answering these questions with a principled solution is our aimed contribution in this paper. By doing so, we will be able to provide the advertisers with more powerful and flexible means to achieve their marketing goals, as well as help search engines to attract more advertisers in a long run.

First of all, we would like to point out that GD is also very important in sponsored search, due to the following reasons. Given the high user traffic in search, many advertisers want to promote their brands through sponsored search, [9] [6] [11]. Such advertisers typically set campaign-

CIKM'12, October 29–November 2, 2012, Maui, HI, USA.

level budgets for pre-defined numbers of reaches to users (i.e., ad impressions) or the numbers of leads (i.e. clicks) to the advertiser websites. For transaction-oriented advertisers whose primary focus is on individual auctions, the success of search advertising is usually measured at a campaign level as well, e.g., the total number of ad impressions and clicks, and the average cost per click, in a certain period of time [8]. Thus, those advertisers also care about the guaranteed performance on macro-level campaigns. This supports the need for joint optimization of auctions and GD on the basis of bid keywords, which is our focus in this paper.

Current approaches in pay-per-click auctions cannot provide guaranteed macro-level performance because those auctions are run in a greedy manner per individual query, not taking into account the total budgets and total number of targeted deliveries (ad impressions or clicks) as global optimization criteria. As a result, often a significant portion of the budgets allocated by advertisers remains unused in the auction process, meaning that the advertisers cannot reach their expected campaign goals. Our analysis of one-month auctions of a commercial search engine shows that the used-budget ratio is only around 10% on average. Some recent works propose modifications to the greedy approach by accommodating budgeted bidders (e.g., [2, 1]). Those approaches, however, have only focused on the advertisers whose budgets are running out, and not on the advertisers who cannot reach their targets in spending and click-through numbers. As a result, those methods are insufficient either for optimizing macro-level performance, or for best combining GD and auctions.

How to best combine GD and auctions in sponsored search, to satisfy the needs of all kinds of advertisers, and to optimize the revenue for search engines is the open challenge for research that we focus in this paper. We propose a unified optimization framework for combining GD and pay-per-click auctions. Specifically, we formulate the optimization problem using linear programming, which maximizes the revenue of the search engine subject to a number of constraints, including maximum budgets in auctions and guaranteed click numbers in GD. We also propose to use a column generation strategy to efficiently find the optimal solution in linear programming. To select the best columns (a ranked list of ads) given a query, we propose a novel dynamic programming algorithm that takes the special structure of the ad allocation and pricing mechanisms into account. To our knowledge, this is the first work in the literature of sponsored search that addresses advertisers' needs of both direct-response marketing and brand marketing using joint optimization of GD and auctions.

We have tested the proposed framework and the algorithms on real ad data obtained from a commercial search engine. The results show that our proposal method significantly improved the number of clicks of selected ads, and significantly increased the revenue of the search engine, compared to representative baseline methods, including using the pay-per-click auction only and allocating pre-specified proportions of query submissions to GD and auction respectively.

The remainder of the paper is organized as follows. In section 2 we introduce the unified auction and GD allocation mechanism, formulate the problem as a linear program and describe our solution based on delayed column generation (2.1) and dynamic programming (2.2). We present our dataset, evaluation setup and experimental results in section 3 and the conclusion with future work in section 4.

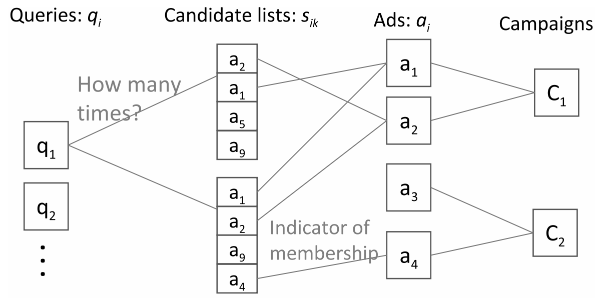

Figure 1: Supply-demand graph that represents our formulation. Supply nodes on the left-hand side correspond to query submissions. Demand nodes on the right-hand side correspond to campaign objectives. The decision variables determine how often to display a particular candidate ranked list of ads in response to a query.

2. UNIFIED AUCTION AND GUARANTEED DELIVERY FOR SPONSORED SEARCH

In this section we describe our novel approach to sponsored search that combines together keyword auctions and guaranteed delivery in a unified optimization framework. In our proposed mechanism we allow two types of advertisers. Some advertisers are interested in auctions and bidding on keywords, they select relevant keywords and bids for their ads, we call them Auction-type advertisers and their campaigns accordingly - Auction campaigns. The others prefer to specify the number of clicks as the explicit goal of the campaign, they select relevant keywords for their ads, we call them GD advertisers and their campaigns - GD campaigns[1]. Whenever a query is sumbitted a ranked list of ads is displayed. This ranked list can contain both GD and auction ads. GD ads can take any positions and Auction ads are ranked by expected value (product of predicted click probability and the bid). GD advertisers are charged per campaign, the penalty is imposed on the search engine if the advertiser's goal (number of clicks) is not met. In scientific literature this penalty is usually taken to be linear. The Auction advertisers are charged per click according to the updated GSP pricing rule: the smallest bid that would preserve his/her position in the *partial* ranking of auction ads. Displaying the list of ads contributes to the goals of all the campaigns that have ads in it, Figure 1.

Let us assume that we can enumerate all possible ranked lists of ads consistent with the ranking criteria above. How often do we need to display a particular ranked list of ads in order to satisfy the goals of the advertisers, the search engine and the users? To answer this question we formulate and solve the following Linear Program (see the notation in Table 1):

[1]Actually, one advertiser can run several campaigns: some of them can be Auction-type campaigns and some of them can be GD-type campaigns. This situation can be addressed by our solution.

Table 1: Problem notation

Variable	Description
q_i	a query
n_i	number of times we expect q_i to be issued
s_{ik}	a ranked list of ads that can be displayed in response to a query q_i
$x_{ik} \in [0, 1]$	the frequency of displaying the list s_{ik}
$a_j \in C$	an individual advertisement that is a part of a campaign C
$C(j)$	the campaign of an ad a_j, i.e. $C(j) = C : a_j \in C$
A, GD	the sets of auction and GD type campaigns, for example $C \in A$ or $C \in GD$
d_C	a budget of a campaign C
m_C	a click requirement for a GD campaign
z_{ikj}	an indicator variable, $z_{ikj} = 1$ if an ad a_j is a member of a list s_{ik}
p_{ikj}	the payment if an ad a_j in the list s_{ik} is clicked; $p_{ikj} = 0$ for GD ads. For the auction ads the payment is determined by the updated GSP pricing rule
c_{ikj}	the click-through rate (CTR) for an ad a_j in the list s_{ik}

$$\max_{x_{ik},\xi_C} \sum_i \sum_k \alpha_{ik} x_{ik} + \sum_{C \in GD} (d_C - \mu_C \xi_C) \tag{1}$$

$$\text{s.t.} \quad \sum_i \sum_k \sum_{j:a_j \in C} \beta_{ikj} x_{ik} \le d_C \quad \forall C \in A$$

$$\sum_i \sum_k \sum_{j:a_j \in C} \beta_{ikj} x_{ik} - \xi_C \le -m_C \quad \forall C \in GD$$

$$\sum_k x_{ik} \le 1 \quad \forall i$$

$$x_{ik}, \xi_C \ge 0$$

We use i to index queries, k to index ranked lists and j to index ads. The decision variable of the problem x_{ik} is the frequency of displaying the list s_{ik} in response to a query q_i. The objective function is the revenue of the search engine. The first summation of the objective function is the expected revenue from the auction campaigns. The coefficient α_{ik} is the expected revenue for n_i displays of the list s_{ik} (the summation over member ads):

$$\alpha_{ik} = n_i \sum_j c_{ikj} p_{ikj} z_{ikj} \tag{2}$$

The second summation of the objective function is the revenue from GD campaigns penalized by underdelivery. We express underdelivery using slack variables ξ_C. The first constraint guarantees that the expected spending of the Auction campaign C is limited by the campaign budget d_C. The second constraint guarantees that the GD campaign C gets the required amount of clicks m_C. The coefficients of the constraints matrix β_{ikj} represent the expected contribution of an ad a_j in the list s_{ik} to the campaign goals.

$$\beta_{ikj} = n_i c_{ikj} p_{ikj} z_{ikj} \quad \forall j : C(j) \in A \tag{3}$$

$$\beta_{ikj} = -n_i c_{ikj} z_{ikj} \quad \forall j : C(j) \in GD \tag{4}$$

The formulation (1) explicitly satisfies the following global goals: 1) maximize the search engine revenue, 2) guarantee the click requirements of GD campaigns and 3) provide equal opportunities to auction campaigns. We also expect our solution to 4) improve user's satisfaction expressed by the average click-thru rate (CTR). Note, that unlike display advertisers who are interested in a guaranteed number of impressions, sponsored search advertisers are interested in a guaranteed number of clicks. This means that the user's satisfaction is an implicit objective.

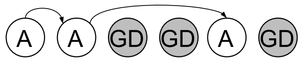

Figure 2: Dependencies in a ranked list of ads, that affect the objective.

The dimensionality of (1) is huge because of $\sum_k$ - summation over feasible ranked lists of ads. There are exponentially many candidate lists. Our remedy to the size of the problem is the fact that 1) only few "good" candidate lists are worth considering; and 2) these "good" candidate lists can be efficiently generated on the fly. To identify these good candidates we formulate a series of auxiliary optimization problems in section 2.1. To solve these problems we exploit the special structure of (1) in section 2.2

2.1 Delayed Column Generation

Linear programs with a huge number of variables and a small number of constraints can be solved using delayed column generation technique [7]. The number of non-zero (active) variables in the solution is small and to find them it is not necessary to consider all the variables. Delayed column generation finds the active variables iteratively: start with an LP with a small subset of variables, solve this LP (also called restricted LP), find the variables that can improve the objective and add them to the restricted LP, solve again etc. The variables that can improve the objective are the variables with positive reduced cost (see [7] for details). The reduced cost of a variable x_{ik} of (1) is:

$$\begin{aligned} cost(x_{ik}) &= \sum_{j:C(j)\in A} n_i c_{ikj} p_{ikj} z_{ikj} (1 - y_{C(j)}) \\ &+ \sum_{j:C(j)\in GD} n_i c_{ikj} z_{ikj} y_{C(j)} - \gamma_i \end{aligned} \tag{5}$$

Where y and γ is the dual solution for restricted LP. y is dual to campaign constraints of (1) and γ is dual to query constraints of (1). Usually the variables with the highest reduced cost are added. To find the maximizer of (5) we do not need to inspect all the possible ranked lists of ads s_{ik}, it is sufficient to construct the best list. What makes this task challenging is the auction advertisements whose payments depend on the next auction advertisement in the list through the GSP pricing rule. We describe the solution to this problem in the next section.

2.2 Dynamic Programming

In this section we provide the solution to maximizing (5). Optimization of the objectives defined over the chains of objects, when the contribution of a current element depends not only on the element itself, but also on the next element in the chain, can be done by dynamic programming (DP). However (5) is more challenging because of the presence of two types of objects: GD ads (unchained) and auction ads (chained). If an auction ad is followed by one or more GD ads than its contribution depends not on the next object in the chain, but on the object located further down the list, Figure 2.

Our goal is to find a sequence of ads:

$$s = \{a_{j_1}, a_{j_2}, ...a_{j_p}\}$$

that maximizes (5). Let us assume that we have assigned the first $r-1$ positions in this sequence:

$$\tilde{s} = \{a_{j_1}, a_{j_2}, ...a_{j_{r-1}}\}$$

and reached a state that is a function of our decisions:

$$T_r = State(a_{j_1}...a_{j_{r-1}})$$

Then our residual goal is to assign the ads to the positions $r \ldots p$ given the state T_r. The sufficient condition of optimality and applicability of dynamic programming is Bellman's principle: the decisions at steps $r \ldots p$ that constitute optimal solution depend only on the state that we reach after steps $1 \ldots r-1$. Then we can use backwards induction to solve this problem with complexity $O(M^2L)$ where M is the number of possible states and L is the length of a chain. We want to find a compact state representation that satisfies the Bellman's principle.

We claim that the allocation of future auction ads only depends on the last auction ad in the list $a_{j_1}...a_{j_{r-1}}$ and its position. GD ads in the list are ranked by $n_i c_{ij_r k} y_{C(j_r)}$ and the allocation of future GD ads depends on how many GD ads were already allocated. I.e. the state function is the following tuple:

$$T_r = \{last\ auction\ ad, its\ position, number\ of\ GD\ ads\}$$

The size of this state space is bounded by ZLL, where Z is the size of the pool of relevant auction ads and L is the length of the chain - usually a small integer (maximum number of ads displayed on a page). If we ignore the positional bias (dependency of click probabilities on position in the list) then the state space can be reduced to ZL.

The necessary condition for the optimal sequence is expressed by Bellman's equation that connects the residual objective at state T_r with immediate gain from selecting an ad for the slot $r : g(a_{j_r}|T_r)$ and the residual objective for the next state $T_r + a_{j_r} \to T_{r+1}$:

$$F(T_r) = \max_{a_{j_r}} \{g(a_{j_r}|T_r) + F(T_{r+1})\} \quad (6)$$

The immediate gain if $a_{j_r} \in GD$ is simply

$$g(a_{j_r}|T_r) = n_i c_{ikj_r} y_{C(j_r)} \quad (7)$$

i.e., the expected number of clicks scaled by $y_{C(j_r)}$. To compute the immediate gain if $a_{j_r} \in A$ assume that the position of the last auction ad in the list T_r is m. If we assign an ad $a_{j_r} \in A$ to the slot r then the advertiser who got the position m for his/her ad a_{j_m} would pay the amount to maintain his/her position against a_{j_r}:

$$p_{ikj_m} = b_{ij_r} c_{ikj_r} / c_{ikj_m} \quad (8)$$

where b_{ij} is the bid value for a_j on a keyword in query q_i. This gives us the gain:

$$g(a_{j_r}|T_r) = n_i c_{ikj_r} b_{ij_r} (1 - y_{C(j_m)}) \quad (9)$$

Note that the CTR and the bid correspond to the ad in position r and the scaling factor $1 - y_{C(j_m)}$ corresponds to the ad in position m.

The Bellman's equation (6) can be solved with the backward induction algorithm.

3. EVALUATION

The evaluation was performed on the log data from a commercial search engine.To make the problem manageable we use only the head queries. This is a standard technique (for example [1]) since few head queries cover a large proportion of query submissions. We discuss how to generalize our approach to the tail queries in Section 4. Our dataset contains 1000 unique query strings and 436K query submissions. A total of 37,864 unique ads from 2,801 unique campaigns match the selected queries.

Our data comes from the keyword auction log of the search engine and only contains Auction campaigns. We have no control of which advertisers and how many of them will switch their campaigns to GD - this will be driven by the market demand and will not be controlled by the search engine.To simulate GD campaigns we randomly select $r\%$ of auction campaigns and convert them to GD campaigns. In the simulation we investigate the performance of the algorithm for different values of r between 0 and 1. We do not impose extra payments (premiums) for using Guaranteed Delivery in the current experimental setup. Normally such payments would exist and would provide more opportunities for optimization and higher expected revenue for the search engine. We set the per click underdelivery penalty μ_C equal to the average per click payment of the campaign.

The number of ads that we display for a query is the same for all three methods.

Baseline 1: We mimic the GSP mechanism currently employed in the search engine. This baseline does not consider GD at all. By comparing to this baseline, we can see the value of our proposed concept, i.e., considering both auction and GD in sponsored search.

Baseline 2: For a more representative baseline we split the available resources (query submissions) proportionally between auction and GD campaigns. Then we run GD and auction allocation mechanisms independently. This baseline considers both auction and GD, but not in a unified framework. It mimics the approach applied in display advertising in which GD and Auctions coexist but do not compete for the same inventory. By comparing to this baseline, we can see the value of our proposed joint optimization framework.

We evaluate our approach from the perspective of the search engine (by revenue), the advertisers (by the number of clicks, average cost per click and the ability to satisfy GD constraints) and the users (by the average quality of displayed ads which is proportional to the total number of clicks). The experimental results are plotted in Figures 3 and 4. From these figures, we have the following observations.

First, both expected revenue (Figure 3, Left) and expected number of clicks (Figure 3, Middle) of our proposed approach significantly outperform the two baselines. For example, the improvement in revenue over Baseline 1 ranges from 3.6% to 7.6%. The improvement in the number of clicks reaches 12.2%. This clearly shows the advantage of jointly optimizing GD and auction. Inferior performance of Baseline2 signals us about sub-optimality of straightforward resource distribution between auction and GD mechanisms and again emphasizes the importance of joint optimization.

The average CPC (Figure 3, Right) of our approach is lower than that of Baseline 1 when the proportion of GD campaigns is more than 0.2. We do not impose the premium payments for GD service in the current setup and therefore

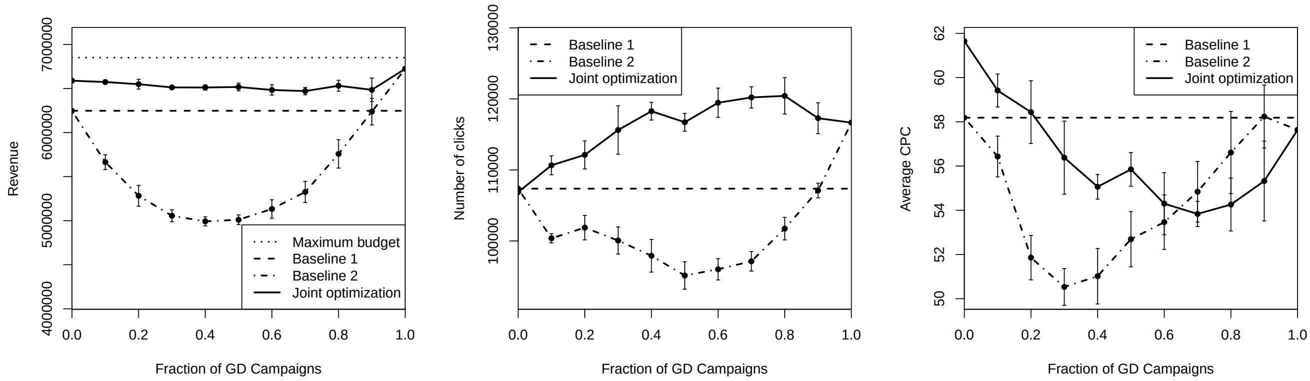

Figure 3: Left: search engine revenue. Middle: expected number of clicks. Right: Average cost per click.

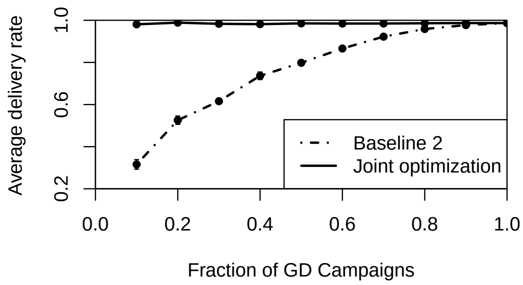

Figure 4: Delivery rate for GD ads.

have a margin for improving the revenue. Baseline 2 has the lowest CPC resulting from its low revenue efficiency.

The last performance metric, average delivery rate, characterizes our ability to satisfy the click constraints for GD campaigns. Delivery rate for a GD campaign is equal to $\min\{recieved\ clicks/requested\ clicks, 1\}$. Since Baseline 1 has no guarantee on the delivery, we only make comparison with Baseline 2. From Figure 4, we can see that our framework maintains the average delivery rate above 0.98 and significantly outperforms Baseline 2.

To sum up, all the above experimental results demonstrate the effectiveness of our proposed joint optimization framework.

4. CONCLUSIONS AND FUTURE WORK

In this paper we proposed the first optimization framework for combining pay-per-click auctions and ad guaranteed delivery (GD) in sponsored search, satisfying advertisers' diverse needs and maximizing revenue for search engine providers simultaneously. Our experimental findings demonstrate that the proposed approach can outperform several representative baseline methods.

For future work we plan to:

(i) evaluate the proposed approach on much larger sponsored search datasets;

(ii) investigate which advertisers are more likely to choose Guaranteed Delivery over Auctions, and vice versa, and to improve our optimization formulation accordingly,

(iii) incorporate personalized and/or affinity-group-based click probabilities in order to model the user utility function,

(iv) generalize our method to the tail queries by combining them into a manageable number of coarse clusters in offline optimization and dispatching the solution in the online stage,

(v) and investigate the effect of the error in forecasting the number of query submissions and predicting the click-through rates on the proposed approach.

5. REFERENCES

[1] Z. Abrams, O. Mendelevitch, and J. Tomlin. Optimal delivery of sponsored search advertisements subject to budget constraints. In *Proceedings of the 8th ACM conference on Electronic commerce*, EC '07, pages 272–278, New York, NY, USA, 2007. ACM.

[2] N. R. Devenur and T. P. Hayes. The adwords problem: online keyword matching with budgeted bidders under random permutations. In *Proceedings of the 10th ACM conference on Electronic commerce*, EC '09, pages 71–78, 2009.

[3] B. Edelman, M. Ostrovsky, and M. Schwarz. Internet advertising and the generalized second price auction: Selling billions of dollars worth of keywords. *American Economic Review*, 97:242–259, 2007.

[4] IAB. http://www.iab.net/adrevenuereport.

[5] R. M. Karp, U. V. Vazirani, and V. V. Vazirani. An optimal algorithm for on-line bipartite matching. In *Proceedings of the twenty-second annual ACM symposium on Theory of computing*, STOC '90, pages 352–358, New York, NY, USA, 1990. ACM.

[6] C. Lake. http://econsultancy.com/us/blog/7345-how-to-use-paid-search-as-a-branding-tool.

[7] M. Lubbecke and J. Desrosiers. Selected topics in column generation. *Operations Research*, 53:1007–1023, 2004.

[8] Microsoft. http://advertising.microsoft.com/small-business/support-center/search-advertising/measure-your-success.

[9] P. Parker. www.clickz.com/clickz/news/1694929/sempo-most-search-advertisers-want-brand-impact.

[10] E. Vee, S. Vassilvitskii, and J. Shanmugasundaram. Optimal online assignment with forecasts. In *Proceedings of the 11th ACM conference on Electronic commerce*, EC '10, pages 109–118, New York, NY, USA, 2010. ACM.

[11] L. Weiman. http://searchengineland.com/top-5-ways-to-use-search-for-branding-66332.