

Profit Maximization for Online Advertising Demand-Side Platforms

Paul Grigas, Alfonso Lobos
University of California, Berkeley
pgrigas,alobos@berkeley.edu

Zheng Wen
Adobe Research
zwen@adobe.com

Kuang-chieh Lee
Yahoo Inc.
kclee@yahoo-inc.com

ABSTRACT

We develop an optimization model and corresponding algorithm for the management of a demand-side platform (DSP), whereby the DSP aims to maximize its own profit while acquiring valuable impressions for its advertiser clients. We formulate the problem of profit maximization for a DSP interacting with ad exchanges in a real-time bidding environment in a cost-per-click/cost-per-action pricing model. Our proposed formulation leads to a nonconvex optimization problem due to the joint optimization over both impression allocation and bid price decisions. We use Lagrangian relaxation to develop a tractable convex dual problem, which, due to the properties of second-price auctions, may be solved efficiently with subgradient methods. We propose a two-phase solution procedure, whereby in the first phase we solve the convex dual problem using a subgradient algorithm, and in the second phase we use the previously computed dual solution to set bid prices and then solve a linear optimization problem to obtain the allocation probability variables. On several synthetic examples, we demonstrate that our proposed solution approach leads to superior performance over a baseline method that is used in practice.

1 INTRODUCTION

In targeted online advertising, the main goal is to figure out the best opportunities by showing an advertisement to an online user, who is most likely to take a desired action, such as ordering a product or signing up for an account. The complexity of realizing this goal is so high that advertisers need specialized technology solutions called demand-side platforms (DSP).

In a DSP, each individual advertiser usually sets up a list of campaigns that can be thought of as plans for delivering advertisements. For each campaign, the advertiser specifies the characteristics of the audience segments that she would like to target (e.g., males, ages 18-35, who view news articles on *espn.com*) along with the particular media that she would like to display to the target audience (e.g., a video ad for beer). In addition, the advertiser specifies a budget amount, time schedule, pacing details, and performance goals for each campaign. The performance goals typically can be specified by minimizing cost-per-click (CPC) or cost-per-action (CPA).

The DSP manages its active campaigns for many different advertisers simultaneously across multiple ad exchanges where ad impressions can be acquired through a real-time bidding (RTB) process. In the RTB process, the DSP interacts with several ad exchanges where bids are placed for potential impressions on behalf of those advertisers. This interaction happens in real time when an ad request is submitted to an ad exchange (which may happen, for example, when a user views a news story on a webpage). In this scenario, the DSP needs to offer a solution to decide, among the list of all campaigns associated with its advertiser clients, which

campaign to bid on behalf of and how much to offer for the corresponding bid. The fundamental problem we consider here is how to make these decisions in real time to maximize the profit for the DSP while ensuring that all of its advertiser clients are satisfied.

Generally speaking, today most DSPs offer different pricing models and enhancement schemes to help advertisers manage their campaigns. Those pricing models include cost per thousand impressions (CPM), cost per click (CPC) and cost per action (CPA). Advertisers often like to choose the CPC/CPA pricing model because the return on investment (ROI) is always positive. However, this pricing model might introduce revenue loss for the DSP since the DSP only earns revenue when a click or action occurs. Therefore, in CPC/CPA pricing model, the DSP needs to convert the CPC/CPA bid to an expected cost per thousand impressions (eCPM) bid in order to sensibly participate in auctions in the RTB exchanges. In this paper, we focus on CPC/CPA pricing model as it is a very challenging problem for the DSP.

It is challenging for a DSP to perform such profit optimization with CPC/CPA pricing model in a RTB environment for several reasons. First, top DSPs typically receive as many as a million ad requests per second. The short latency and high throughput requirements introduce extreme time sensitivity on the process. Second, a large amount of information is missing in the real time evaluation of the individual ad requests, e.g., the feedback on previous decisions normally has a long delay in practice. Therefore, most of the DSPs today only apply a greedy approach by selecting the ad with the highest bid among all the qualified ads for each incoming request.

In this paper, we propose a novel approach based on a precise mathematical formulation to optimize the overall DSP profit. We appropriately model the uncertainty in impression arrival, auction, and click/action processes and develop an optimization formulation to maximize profit for the DSP while ensuring that each campaign remains under budget. Our formulation is aimed at optimizing with respect to both impression allocation and bid price decisions, and due to the additional complexity of accounting for both of these decisions the formulation is a large-scale, nonconvex model. However, due to the properties of second-price auctions we are able to effectively use the technique of Lagrangian relaxation. We construct a dual problem and establish that subgradients of the dual function may be efficiently computed. Our overall approach is based on a two-phase procedure, wherein we solve the dual problem in the first phase and use the dual solution to naturally recover a primal solution in the second phase. We conduct several computational experiments on synthetic datasets and demonstrate that our Lagrangian relaxation based approach is able to significantly increase DSP profits relative to a baseline greedy approach.

Revenue optimization in online advertising has been extensively studied in recent literature from different perspectives, such as

optimization [2, 4, 9, 13], game theory and mechanism design [1, 7], and contract design [10]. Due to space limitations, we only review several directly relevant papers. [2–4] focus on the publisher’s revenue management problem. Specifically, [2, 3] study how publishers should optimally trade off guaranteed contracts with RTB. [4] studies how a publisher should optimally allocate impressions and set up bid prices for campaigns, under the implicit assumption that the publisher is a “central planner”. On the other hand, [5] studies ad networks’ revenue management problem based on model predictive control. Finally, [13] studies advertisers’ optimal bidding problem in RTB. Unlike our paper, [13] focuses on optimal bidding and its framework does not consider impression allocation.

The rest of the paper is organized as follows. In Section 2, we describe the notation and problem statement and we set up the model. The profit optimization formulation is presented in Section 3. Our proposed Lagrangian relaxation based algorithm to approximately solve the profit optimization problem is specified in Section 4. Experimental results are presented in Section 5, we conclude with some remarks about possible future work.

2 MODEL FOUNDATIONS

We assume that the planning takes place over a fixed time horizon (e.g., 24 hours). To simplify the presentation, we also assume that the DSP interacts with a single ad exchange. For our purposes, the ad exchange simply represents a pool of potential impressions that the DSP may bid on. Thus, in the likely scenario that the DSP interacts with more than one ad exchange, we may simply group all of the different ad exchanges together into a representative ad exchange.

Let \mathcal{I} denote the set of *impression types* and let \mathcal{K} denote the set of all campaigns associated with advertisers managed by the DSP. Before discussing the details of our assumptions, let us describe the basic flow of events in the model. When an impression of type $i \in \mathcal{I}$ is submitted to the ad exchange, a real-time second-price auction is held for which the DSP has an opportunity to bid. Thus the DSP has an opportunity to make two strategic decisions related to each real-time auction: (i) how to select a campaign $k \in \mathcal{K}$ to bid on behalf of in the auction, and (ii) how to set the corresponding bid price b_{ik} . If the DSP wins the auction on behalf of campaign k , then the DSP must pay the ad exchange an amount equal to the second largest price and an ad from campaign k is displayed. The advertiser corresponding to campaign k is charged only if a user clicks on the ad.

Impression Types. It is important that the set of impression types \mathcal{I} represents a partition of all possible impressions that are submitted to the ad exchange. Thus, every impression submitted to the ad exchange is associated with a particular impression type $i \in \mathcal{I}$. It is most natural to define \mathcal{I} in terms of features associated with impressions. For example, if the DSP determines that there are only two relevant attributes associated with each impression – say gender and whether or not the viewer is 18 years or older – then the DSP would choose $\mathcal{I} = \{(M, 18-), (M, 18+), (F, 18-), (F, 18+)\}$. Then, all impressions corresponding to male viewers who are under 18 years of age would be assigned to impression type $i = (M, 18-)$, etc. Note that the construction of the set \mathcal{I} is part of the modeling process and consideration should be given to the trade-off

between computational limitations and the potential for higher profits due to a more fine-grained construction of \mathcal{I} . Nevertheless, the algorithmic schemes we propose in Section 3 are scalable to problem instances where the size of \mathcal{I} is extremely large. We use the following notation and make the following assumptions about the impression types:

- Let S_i denote the number of impressions of type i that arrive during the planning horizon, and assume that S_i is a random variable with mean s_i .

Campaigns. Recall that \mathcal{K} denotes the set of all campaigns that are managed by the DSP. That is, \mathcal{K} is the union over all advertisers (who are managed by the DSP) of the sets of campaigns run by each advertiser. We use the following notation and make the following assumptions about the campaigns:

- m_k denotes the (advertiser selected) budget for campaign k during the planning horizon.
- \mathcal{I}_k denotes the set of impression types that campaign k targets. For example, if an advertiser wishes to create a particular campaign to target female users, then in this case \mathcal{I}_k would denote the set of all impression types corresponding to female users (e.g., $\mathcal{I}_k = \{(F, 18-), (F, 18+)\}$ in the example described above). Note that each advertiser can create multiple campaigns to achieve different targeting goals.
- In this model, it is assumed that advertisers are charged on a *CPC (cost per click)* basis. That is, campaign k is charged an amount $q_k > 0$, called the CPC price, each time a user clicks on an advertisement from campaign k . (Note that a “click” may also be thought of more generally as an “action” whereby our model easily extends to campaigns that are charged on a CPA basis. Moreover, our model may be easily extended to allow for multiple actions, each with their own rewards.)

Auctions. When an impression is submitted to the ad exchange, an instantaneous real-time auction occurs to determine who gets to display an advertisement. We assume that these are *second-price* auctions, which are very common in practice. In a second-price auction, the bidder who submits the highest bid is the winner, but the amount that the winner pays is the amount of the *second highest* bid. It is well known that, in a second-price auction, a dominant strategy for each participant is to bid truthfully [12]. Herein we assume that the DSP takes a *probabilistic* approach to modeling the behavior of the other bidders in the auction. Namely, we make the following assumptions:

- For each impression type $i \in \mathcal{I}$, let B_i^{\max} be a random variable representing the maximum, among all other bidders excluding the DSP, of the bid prices entered in an auction for an impression of type i . It is assumed that $B_i^{\max} > 0$ with probability one. Let $\rho_i(\cdot) : \mathbb{R} \rightarrow [0, 1]$ denote the cumulative distribution function (CDF) of B_i^{\max} , so that $\rho_i(b) := \mathbb{P}(B_i^{\max} \leq b)$ is the probability that the DSP wins an auction for an impression of type i when the DSP enters the bid b . Note that the functions $\rho_i(\cdot)$ are typically estimated using a bid landscape model (see, e.g., [6]). (In the case of a tie, here we assume that the DSP automatically wins the auction. Our framework may be easily adapted to the case of a fair tie-breaking mechanism.)
- Furthermore, let $\beta_i^{\max}(b) := \mathbb{E}[B_i^{\max} \mid B_i^{\max} \leq b]$ denote the expected value of the second highest bid price (i.e., the value of

the payment to the ad exchange) given that the DSP enters a bid price of b and b is the largest bid price entered.

Click Events. After the DSP has won an auction on behalf of campaign k , an ad for campaign k is displayed to the user corresponding to the impression for which the auction was held. For a given impression type $i \in \mathcal{I}$ and a given campaign $k \in \mathcal{K}$, let $\theta_{ik} \in [0, 1]$ denote the *click-through-rate* for users corresponding to impression type i and when the ad corresponds to campaign k . That is, θ_{ik} represents the fraction of users corresponding to impression type i that click on an ad associated with campaign k , i.e., the probability that the user clicks on the ad that is shown. Although the true click-through-rates are not available, the DSP is typically able to leverage a vast amount of historical data and use predictive models to produce accurate predictions of these values, even when \mathcal{K} and \mathcal{I} are extremely large (see, e.g., [8]).

Finally, given an impression type $i \in \mathcal{I}$ and a campaign $k \in \mathcal{K}$, let r_{ik} denote the *expected cost per impression (eCPI)* value, namely $r_{ik} := q_k \theta_{ik}$ where q_k is the CPC price defined earlier. Note that r_{ik} is the expected amount of revenue that the DSP earns each time an ad for campaign k is shown to an impression of type i , and r_{ik} also corresponds to the optimal bid price when campaign k has unlimited budget.

Decision Variables and Additional Notation. As mentioned previously, when an auction for impression type $i \in \mathcal{I}$ arrives to the ad exchange, the DSP decides which campaign $k \in \mathcal{K}$ to bid of behalf of and also selects the value of the corresponding bid price. Let $\mathcal{E} \subseteq \mathcal{I} \times \mathcal{K}$ denote the edges of an undirected bipartite graph between \mathcal{I} and \mathcal{K} , whereby there is an edge $e = (i, k) \in \mathcal{E}$ whenever campaign k targets impression type i , i.e., $\mathcal{E} := \{(i, k) : i \in \mathcal{I}_k\}$. Let $\mathcal{K}_i := \{k \in \mathcal{K} : (i, k) \in \mathcal{E}\}$ denote the set of campaigns that target impression type i .

When a new auction for impression type i arrives to the ad exchange, we say that the DSP selects campaign k for the auction if the DSP chooses to bid on behalf campaign k in the auction. For each edge $(i, k) \in \mathcal{E}$, we define two decision variables as follows: (i) x_{ik} is the probability that the DSP selects campaign k , and (ii) b_{ik} is the corresponding bid price that the DSP submits in the auction. Interpreted differently, x_{ik} represents a proportional allocation, i.e., the fraction of auctions for impression type i that are allocated to campaign k on average. Note that b_{ik} represents the bid price that the DSP submits to an auction for impression type i conditional on the fact that the DSP has selected campaign k for the auction. Related approaches (e.g., as in [4]) also use bid prices to rank advertisers – in our approach, the selection of which campaign to bid on behalf of is completely captured by the x_{ik} decision variables and thus the b_{ik} decision variables only determine the actual bid price decisions. Let $\mathbf{x}, \mathbf{b} \in \mathbb{R}^{|\mathcal{E}|}$ denote vectors of these quantities, which will represent decision variables in our model.

Let us also define some additional notation used herein. For a given set S and a function $f(\cdot) : S \rightarrow \mathbb{R}$, let $\arg \max_{x \in S} f(x)$ denote the (possibly empty) set of maximizers of the function $f(\cdot)$ over the set S . If $f(\cdot) : \mathbb{R}^n \rightarrow \mathbb{R}$ is a convex function then, for a given $x \in \mathbb{R}^n$, $\partial f(x)$ denotes the set of subgradients of $f(\cdot)$ at x , i.e., the set of vectors g such that $f(y) \geq f(x) + g^T(y - x)$ for all $y \in \mathbb{R}^n$. Finally, let $\mathbf{1}(\cdot)$ denote an indicator function that is equal to 1 whenever the argument of $\mathbf{1}(\cdot)$ is true and equal to 0 otherwise.

3 OPTIMIZATION FORMULATION

Problem (1) presents our formulation of the allocation and real-time bidding planning problem faced by the DSP.

$$\begin{aligned} & \underset{\mathbf{x}, \mathbf{b}}{\text{maximize}} && \sum_{(i,k) \in \mathcal{E}} [r_{ik} - \beta_i^{\max}(b_{ik})] s_i x_{ik} \rho_i(b_{ik}) \\ & \text{subject to} && \sum_{i \in \mathcal{I}_k} r_{ik} s_i x_{ik} \rho_i(b_{ik}) \leq m_k \quad \forall k \in \mathcal{K} \\ & && \sum_{k \in \mathcal{K}_i} x_{ik} \leq 1 \quad \forall i \in \mathcal{I} \\ & && \mathbf{x}, \mathbf{b} \geq 0. \end{aligned} \quad (1)$$

Herein, let $\pi(\mathbf{x}, \mathbf{b}) := \sum_{(i,k) \in \mathcal{E}} [r_{ik} - \beta_i^{\max}(b_{ik})] s_i x_{ik} \rho_i(b_{ik})$ denote the objective function of (1). Let us now briefly describe the interpretation of each part of the formulation (1) above. First, note that the formulation is based on the idea of “deterministic approximation,” whereby we assume that all random quantities deterministically take on their expected values. In this formulation, the DSP seeks to maximize its total profit over the planning horizon, while ensuring that each campaign does not spend more than its budget. Indeed, the expected number of times during the planning horizon that the DSP selects campaign k for an auction of impression type i is $s_i x_{ik}$ and the expected number of such auctions that the DSP wins is $s_i x_{ik} \rho_i(b_{ik})$. Furthermore, for each instance that the DSP selects campaign k and wins the corresponding auction for impression type i , the expected profit for the DSP is $r_{ik} - \beta_i^{\max}(b_{ik})$. Therefore, the objective function of (1) represents the expected total profit earned by the DSP throughout the planning horizon. The first set of constraints in (1) represent the *budget constraints* for the campaigns, which ensure that, in expectation, each campaign does not spend more than its pre-specified budget level. Finally, the second set of constraints in (1) are referred to as the *supply constraints* for the impression types, which (along with the nonnegativity constraints on \mathbf{x}) ensure that the variables \mathbf{x} represent valid probabilities. Note also that these probabilities may sum to a value strictly less than 1, in which case $1 - \sum_{k \in \mathcal{K}_i} x_{ik}$ represents the probability of electing not to bid when an impression of type i arrives to the ad exchange. For ease of notation, let us denote the supply and nonnegativity constraints on \mathbf{x} using $\mathcal{S} := \{\mathbf{x} : \sum_{k \in \mathcal{K}_i} x_{ik} \leq 1 \text{ for all } i \in \mathcal{I}, \mathbf{x} \geq 0\}$.

Note that, due to the joint optimization over both \mathbf{x} and \mathbf{b} , problem (1) is generally a nonconvex optimization problem (this is clearly evident, for example, when β_i^{\max} is taken to be uniformly distributed on $[0, 1]$ and hence $\rho_i(b_{ik}) = b_{ik}$ for $b_{ik} \in [0, 1]$). Despite its nonconvexity, problem (1) does have some important structural properties that we now highlight. First, if we consider the bid prices \mathbf{b} to be fixed, then the resulting problem in \mathbf{x} is a *linear optimization* problem – in other words the objective function and constraints can all be expressed as linear functions of \mathbf{x} – and may be solved very efficiently using off-the-shelf solvers or perhaps a specialized algorithm. Conversely, if we consider \mathbf{x} to be fixed, then the resulting problem in \mathbf{b} is generally still nonconvex but the main “difficulty” arises from the budget constraints. Indeed, due to the presence of budget constraints, it may be optimal for the DSP to underbid on a relatively less valuable impression due to the possibility of a more valuable impression arriving in the future. Therefore, whenever a campaign has unlimited budget, it is optimal for the DSP to set $b_{ik} = r_{ik}$, i.e., to bid truthfully. The following Proposition, which will be useful in the development of

the Lagrangian relaxation algorithm in Section 4, formalizes this intuition.

PROPOSITION 3.1. *Consider the following modification of (1) without budget constraints:*

$$\begin{aligned} & \underset{\mathbf{x}, \mathbf{b}}{\text{maximize}} && \sum_{(i,k) \in \mathcal{E}} [r_{ik} - \beta_i^{\max}(b_{ik})] s_i x_{ik} \rho_i(b_{ik}) \\ & \text{subject to} && \mathbf{x} \in \mathcal{S} \quad \text{and} \quad \mathbf{b} \geq 0, \end{aligned} \quad (2)$$

where $\mathcal{S} := \{\mathbf{x} : \sum_{k \in \mathcal{K}_i} x_{ik} \leq 1 \text{ for all } i \in \mathcal{I}, \mathbf{x} \geq 0\}$, as defined earlier. Define $(\mathbf{x}^*, \mathbf{b}^*)$ by $b_{ik}^* := r_{ik}$ for all $(i, k) \in \mathcal{E}$ and by letting \mathbf{x}^* be an arbitrary optimal solution of the resulting linear optimization problem, i.e., $\mathbf{x}^* \in \arg \max_{\mathbf{x} \in \mathcal{S}} \{\sum_{(i,k) \in \mathcal{E}} \pi_{ik} x_{ik}\}$, where $\pi_{ik} := [r_{ik} - \beta_i^{\max}(r_{ik})] s_i \rho_i(r_{ik})$ for all $(i, k) \in \mathcal{E}$. Then, (2) is finite and $(\mathbf{x}^*, \mathbf{b}^*)$ is an optimal solution of (2).

Due to space limitations, we do not include the proof of Proposition 3.1 nor the proof of Theorem 4.1 in Section 4. Both proofs are contained in a supplementary document available on the first author's webpage.

REMARK 3.1. *Given coefficients $\pi_{ik} \in \mathbb{R}$ for each $(i, k) \in \mathcal{E}$, an optimal solution \mathbf{x}^* of the linear optimization problem $\max_{\mathbf{x} \in \mathcal{S}} \{\sum_{(i,k) \in \mathcal{E}} \pi_{ik} x_{ik}\}$ may be computed efficiently in $O(|\mathcal{E}|)$ time using a simple greedy algorithm. Namely, for each $i \in \mathcal{I}$, we compute $k^*(i) \in \arg \max_{k \in \mathcal{K}_i} \{\pi_{ik}\}$, set $x_{ik^*(i)} = 1$ ($\pi_{ik^*(i)} > 0$), and set $x_{ik} = 0$ for all other $k \in \mathcal{K}_i$.*

4 LAGRANGIAN DUAL AND ALGORITHMIC SCHEME

We begin this section with a high-level description of our approach for solving (1). Our algorithmic approach is based on a two phase procedure. In the first phase, we construct a suitable dual of (1), which turns out to be a convex optimization problem that can be efficiently solved with most subgradient based algorithms. A solution of the dual problem naturally suggests a way to set the bid prices \mathbf{b} . In the second phase, we set the bid prices using the previously computed dual solution then we solve the linear optimization problem that results when \mathbf{b} is fixed in order to recover allocation probabilities \mathbf{x} .

Let us now construct a Lagrangian dual of the deterministic approximation problem (1) by relaxing the ‘‘difficult’’ budget constraints. We show that the resulting dual problem is a convex optimization problem with only very simple box constraints and that subgradients of the objective function may be efficiently computed. Since our formulation is quite general (for example, there are no strong assumptions made about the distribution of B_i^{\max}), we are unable to exploit any special structure of the dual function $L^*(\cdot)$ and must resort to simple subgradient based algorithms to solve the dual problem. Nevertheless, subgradient methods offer the advantage of being highly scalable and parallelizable, and moreover our overall two phase procedure does not necessitate a high accuracy solution of the dual problem.

To start, we introduce multipliers $\lambda \in \mathbb{R}_+^{|\mathcal{K}|}$ for the budget constraints in (1) and form the Lagrangian function:

$$\begin{aligned} L(\mathbf{x}, \mathbf{b}, \lambda) := & \sum_{(i,k) \in \mathcal{E}} [r_{ik} - \beta_i^{\max}(b_{ik})] s_i x_{ik} \rho_i(b_{ik}) \\ & + \sum_{k \in \mathcal{K}} \lambda_k [m_k - \sum_{i \in \mathcal{I}_k} r_{ik} s_i x_{ik} \rho_i(b_{ik})]. \end{aligned} \quad (3)$$

After rearranging, we may re-express the Lagrangian function:

$$\begin{aligned} L(\mathbf{x}, \mathbf{b}, \lambda) = & \sum_{(i,k) \in \mathcal{E}} [(1 - \lambda_k) r_{ik} - \beta_i^{\max}(b_{ik})] s_i x_{ik} \rho_i(b_{ik}) \\ & + \sum_{k \in \mathcal{K}} \lambda_k m_k \end{aligned} \quad (4)$$

Note that the above implies a natural interpretation of the dual variables as related to scaling factors, namely $1 - \lambda_k$ for each $k \in \mathcal{K}$, to reduce the bid prices based on the fact that each campaign has limited budget. The dual function is defined in the standard way:

$$L^*(\lambda) := \max_{\mathbf{x} \in \mathcal{S}, \mathbf{b} \geq 0} L(\mathbf{x}, \mathbf{b}, \lambda), \quad (5)$$

and we define the dual problem as:

$$\begin{aligned} & \underset{\lambda}{\text{minimize}} && L^*(\lambda) \\ & \text{subject to} && 0 \leq \lambda_k \leq 1 \text{ for all } k \in \mathcal{K}. \end{aligned} \quad (6)$$

Note that the nonnegativity constraints above are standard in Lagrangian duality to ensure that (6) provides a valid upper bound on (1). The upper bound constraints, i.e., $\lambda_k \leq 1$ for all $k \in \mathcal{K}$, are without loss of generality since whenever $\lambda_k > 1$ the dual function $L^*(\lambda)$ is only improved by instead setting $\lambda_k = 1$. It is always the case that $L^*(\cdot)$ is a convex function, but in general it may not be differentiable. Nevertheless, Procedure 1 precisely describes how to compute a subgradient of $L^*(\cdot)$ at λ and is based on Proposition 3.1. Theorem 4.1 demonstrates that Procedure 1 computes valid subgradients and also summarizes the most important properties of the dual function $L^*(\cdot)$ and the dual problem (6).

Procedure 1 Computing a Subgradient of $L^*(\cdot)$

Input: $\lambda \in \mathbb{R}^{|\mathcal{K}|}$ such that $0 \leq \lambda_k \leq 1$ for all k .

1. For each $(i, k) \in \mathcal{E}$, set:

$$b^*(\lambda)_{ik} \leftarrow (1 - \lambda_k) r_{ik}, \text{ and}$$

$$\pi(\lambda)_{ik} \leftarrow [b^*(\lambda)_{ik} - \beta_i^{\max}(b^*(\lambda)_{ik})] s_i \rho_i(b^*(\lambda)_{ik})$$

2. Compute $\mathbf{x}^*(\lambda) \in \arg \max_{\mathbf{x} \in \mathcal{S}} \{\sum_{(i,k) \in \mathcal{E}} \pi(\lambda)_{ik} x_{ik}\}$ using the greedy algorithm described in Remark 3.1.

3. For each $k \in \mathcal{K}$, set:

$$g(\lambda)_k \leftarrow m_k - \sum_{i \in \mathcal{I}_k} r_{ik} s_i x_{ik}^*(\lambda) \rho_i(b^*(\lambda)_{ik}).$$

Output: $g(\lambda) \in \partial L^*(\lambda)$.

THEOREM 4.1. *We have the following properties:*

- (i) $L^*(\cdot)$ is finite and convex everywhere on $\mathbb{R}^{|\mathcal{K}|}$.
- (ii) For any $\lambda \in \mathbb{R}^{|\mathcal{K}|}$ such that $0 \leq \lambda_k \leq 1$ for all k , Procedure 1 computes a valid subgradient $g(\lambda) \in \partial L^*(\lambda)$.
- (iii) Let P^* denote the optimal objective function value of the primal problem (1), and let D^* denote the optimal objective function value of the dual problem (6). Then, for any (\mathbf{x}, \mathbf{b}) that is feasible for (1) and any λ that is feasible for (6), it holds that:

$$\pi(\mathbf{x}, \mathbf{b}) \leq P^* \leq D^* \leq L^*(\lambda).$$

Algorithm 2 presents our two phase procedure for obtaining an approximate solution $(\hat{\mathbf{x}}, \hat{\mathbf{b}})$ of problem (1). In Phase 1, we solve the dual problem (6). As mentioned previously, we suggest the use of simple subgradient methods (see, for example, [11] and the references therein), with the use of Procedure 1 to compute subgradients, in order to solve this problem. In our experiments in Section 5, we use the basic method of projected subgradient descent with

step-sizes proportional to $1/\sqrt{T}$ where T is the iteration counter. In this case, as evident from Procedure 1, the subgradients will remain bounded and therefore we may apply classical convergence results for this method, which state that the objective function value optimality gap converges to zero at the rate of $O(1/\sqrt{T})$ [11]. Moreover, the per iteration cost of this method is dominated by the cost of computing a subgradient, which, as is clear from Procedure 1 and Remark 3.1, is $O(|\mathcal{E}|)$.

Algorithm 2 Two Phase Lagrangian Relaxation-based Scheme for Problem (1)

Phase 1: Solve Lagrangian Relaxation

Solve the dual problem (6) to near global optimality using a subgradient method, and return dual variables $\hat{\lambda}$ and dual objective value $\hat{D} \leftarrow L^*(\hat{\lambda})$.

Phase 2: Primal Recovery

1. Set bid prices $\hat{\mathbf{b}}: \hat{b}_{ik} \leftarrow (1 - \hat{\lambda}_k)r_{ik}$ for all $(i, k) \in \mathcal{E}$.
2. Consider the primal problem (1) with the \mathbf{b} variables fixed at the values $\hat{\mathbf{b}}$, and solve the resulting linear optimization problem to obtain allocation probabilities $\hat{\mathbf{x}}$.

Output: Approximate primal solution $(\hat{\mathbf{x}}, \hat{\mathbf{b}})$, primal objective value $\pi(\hat{\mathbf{x}}, \hat{\mathbf{b}})$, and dual upper bound \hat{D} .

In Phase 2 of Algorithm 2, we suggest a heuristic to construct an approximate primal solution $(\hat{\mathbf{x}}, \hat{\mathbf{b}})$ based on the previously computed dual solution $\hat{\lambda}$. First, we use the natural correspondence suggested by (4) to set the bid prices $\hat{\mathbf{b}}$. Then, we fix these bid prices and solve the resulting linear optimization problem from (1) to obtain allocation probabilities $\hat{\mathbf{x}}$. In our experiments, we use the off-the-shelf solver Gurobi although more sophisticated and scalable approaches may also be employed. Finally, Algorithm 2 outputs the approximate primal solution $(\hat{\mathbf{x}}, \hat{\mathbf{b}})$ along with the dual objective value \hat{D} . Using item (iii) of Theorem 4.1, we may use \hat{D} to obtain a useful bound on the suboptimality of $(\hat{\mathbf{x}}, \hat{\mathbf{b}})$, namely $\pi(\hat{\mathbf{x}}, \hat{\mathbf{b}}) \leq P^* \leq \hat{D}$.

5 COMPUTATIONAL EXPERIMENTS

In this section, we present the results of several computational experiments wherein we applied our two-phase solution procedure to synthetic data examples and compared its performance to a baseline policy. Before discussing the experimental results, it is important to clarify how our model, and in particular the output of Algorithm 2, should be applied in a practical, online environment. Policy 3 precisely describes the sequence of events that occur when a new impression of type i arrives to ad exchange and also describes how the decision variables $(\hat{\mathbf{x}}, \hat{\mathbf{b}})$ resulting from Algorithm 2 would be used to make decisions in real-time. (Recall that the decision variables \mathbf{x} represent probabilities of selecting campaigns, and that $1 - \sum_{k \in \mathcal{K}_i} x_{ik}$ represents the probability of not participating in an auction for a new impression of type i . In Policy 3 below, the symbol “0” is used to encode a “null campaign” that represents this option of refraining from participating in the auction.)

We also refer to Policy 3 as the “Lagrangian policy.” Notice that Policy 3 is particularly conservative in dealing with what happens when a campaign depletes its budget. Indeed, if a campaign with

Policy 3 Online Policy Implied by Algorithm 2

Input: Approximate primal solution $(\hat{\mathbf{x}}, \hat{\mathbf{b}})$ from Algorithm 2 and new impression arrival $i \in \mathcal{I}$.

1. Sample a campaign $\tilde{k} \in \mathcal{K}_i \cup \{0\}$ according to the distribution implied by the values x_{ik} for $k \in \mathcal{K}_i$ and $1 - \sum_{k \in \mathcal{K}_i} x_{ik}$. If $\tilde{k} = 0$ (the sampled campaign is null) or the sampled campaign \tilde{k} has depleted its budget, then break.
 2. Enter bid price $\hat{b}_{i\tilde{k}}$. If the auction is won, then pay the ad exchange an amount to equal to the second price. If the auction is not won, then break.
 3. Show an ad for campaign \tilde{k} . If a click happens, then deduct $q_{\tilde{k}}$ from the budget of campaign \tilde{k} and earn revenue $q_{\tilde{k}}$.
-

depleted budget is sampled then the DSP does not participate in the auction. Policy 3 may be improved by incorporating the idea of model predictive control whereby Algorithm 2 is rerun every time a campaign depletes its budget, or possibly at periodic time intervals. We compare Algorithm 2 and correspondingly Policy 3 against a simple “greedy policy” that is often employed in practice. Indeed, the greedy policy has the same basic flow of events as Policy 3 with two major differences in how decisions are made: (i) at Step (1.) the greedy policy selects, among those campaigns in \mathcal{K}_i with budgets that are not yet depleted, the campaign \tilde{k} with the largest eCPI value of r_{ik} , and (ii) at Step (2.) the greedy policy uses $r_{i\tilde{k}}$, the eCPI value of the selected campaign, as the bid price.

Synthetic Data Examples. Let us now describe how the synthetic data examples were generated and how the corresponding simulations were conducted. Throughout this discussion, all relevant random variables are generated independently unless otherwise mentioned. Furthermore, throughout our experiments, it is assumed that the optimization model developed herein is *correctly specified* in that all distributional information used by our model (e.g., in (1) and Algorithm 2) accurately reflect the corresponding random variables in our simulations.

Now, to generate our synthetic instances, we first fix the sizes of the sets \mathcal{K} and \mathcal{I} , and for each campaign $k \in \mathcal{K}$ (and also for each impression type $i \in \mathcal{I}$) we generate a “quality score” Q_k (resp., Q_i) that is uniformly distributed on $[0, 1]$. The quality scores are intended to reflect the “desirability” of each impression type and each campaign and are used to generate the main parameters of our model. Indeed, for each impression type $i \in \mathcal{I}$, the set \mathcal{K}_i is constructed by sampling edges independently with probability Q_i . Hence $|\mathcal{K}_i| \sim \text{Bin}(|\mathcal{K}|, Q_i)$, where $\text{Bin}(n, p)$ denotes a binomial random variable with n trials and success probability parameter p . Moreover, for each impression type $i \in \mathcal{I}$, B_i^{\max} is taken to be the maximum of $\text{Bin}(M, Q_i)$ independent random variables that are uniformly distributed on $[0, 1]$, where M is an integer parameter dictating the “size of the market.” The click-through-rate values θ_{ik} are defined by $\theta_{ik} := Q_i \cdot Q_k$.

In all of our experiments, we set $|\mathcal{K}| = 100$, $M = 10$, $s_i = 5000$ for all $i \in \mathcal{I}$, and the CPC value $q_k = 1$ for all $k \in \mathcal{K}$. In our first experiment, we generated a single problem instance, referred to as Example A, that additionally had $|\mathcal{I}| = 100$ and the budget

parameters set to $m_k = 50$ for all $k \in \mathcal{K}$. In this case, as verified by the dual upper bound \hat{D} , Algorithm 2 was able to solve problem (1) to within 13% of optimality. We compared the Policy 3 implied by Algorithm 2 to the “greedy policy” described earlier by simulating the impression arrival, real-time bidding, and click processes. In our simulations, we assume that impressions arrive to the ad exchange according to $|\mathcal{I}|$ independent Poisson processes that are “merged” together. This assumption implies that S_i is a Poisson random variable, and the time horizon T of the overall arrival process is set so that $s_i = 5000$ for each $i \in \mathcal{I}$.

Results. The top left table in Figure 1 presents the main results of our first experiment on Example A. We ran 500 simulation runs comparing our Lagrangian relaxation approach, i.e., Policy 3 to the greedy baseline policy. Each policy saw the same exact sequence of impression arrivals and the same sequence of realized B_i^{\max} values during each individual simulation run. For each simulation run, we computed the relative profit, relative cost, and relative revenue of the two policies, where each relative statistic is computed as the Lagrangian statistic relative to the greedy baseline, e.g., Relative Profit := $\frac{\text{Lagrangian Profit}}{\text{Greedy Profit}}$. The results in the tables are averaged over 500 such simulation runs. As the top left table demonstrates, the Lagrangian policy is able to achieve significantly higher profit levels and lower costs than the greedy baseline. Interestingly, the Lagrangian policy also achieves lower revenue levels than the greedy policy. This makes good intuitive sense since the Lagrangian policy uses bid prices that are shaded down by a factor of $1 - \lambda_k$ as compared to the greedy policy, and moreover the Lagrangian policy should make smarter allocation decisions whereby, under the Lagrangian policy, a particularly valuable campaign would wait for better opportunities before depleting its budget as compared to the greedy policy. The bottom left table in Figure 1 reports the budget utilization (defined as total revenue divided by $\sum_{k \in \mathcal{K}} m_k$) and profit margin (defined as total profit divided by total revenue) statistics for each policy, which confirms our intuition and the results presented in top left table. Example B constitutes our second experiment, whereby we took the same exact problem instance and made one modification, namely instead of using constant budgets across the different campaigns we allowed the budget to be correlated with the quality score of each campaign so that $m_k := 50Q_k$. As Figure 1 demonstrates, the profit improvement of the Lagrangian policy over the greedy baseline is even more dramatic in this case.




Figure 1: Lagrangian two-phase policy vs. greedy policy

In our third experiment, we generated a fresh problem instance using all of the same parameters as before except for two differences: (i) we set $|\mathcal{I}| = 10$, and (ii) we varied the budget parameters $m_k \in \{5, 10, 15, 20, 25, 30, 35, 40, 45, 50\}$ (also we took m_k to be constant

across the different campaigns) and reran Algorithm 2 for each of the 10 budget values. The right side of Figure 1 plots the relative profit of the two policies versus these budget values, and for each of the 10 budget values the relative profit statistic was averaged over 500 simulation runs. Clearly, the Lagrangian policy has significantly larger profits over the greedy policy for small budget values, but as the budget increases this improvement is diminished. This makes sense since the Lagrangian policy is based on accounting for the budget constraints in (1) via dual variables. Indeed, as the budget values become larger, the budget constraints are less active and Proposition 3.1 implies that Policy 3 and the greedy policy are exactly the same when m_k is large enough for each $k \in \mathcal{K}$.

Let us conclude this section by mentioning a few directions for future research. First, it would be very valuable to also perform some computational experiments comparing the Lagrangian policy to the greedy policy using a real advertising dataset. Second, it would be very interesting to extend our methodology, in particular problem (1) and Algorithm 2 to different pricing models, such as the CPM pricing model, with performance constraints. Finally, it would be interesting to examine the benefits of more sophisticated stochastic or robust optimization approaches that more carefully account for the uncertainty in the impression arrivals and the real-time bidding environment. The authors plan to pursue all of these directions in future research.

REFERENCES

- [1] Santiago R Balseiro, Omar Besbes, and Gabriel Y Weintraub. 2015. Repeated auctions with budgets in ad exchanges: Approximations and design. *Management Science* 61, 4 (2015), 864–884.
- [2] Santiago R Balseiro, Jon Feldman, Vahab Mirrokni, and S Muthukrishnan. 2014. Yield optimization of display advertising with ad exchange. *Management Science* 60, 12 (2014), 2886–2907.
- [3] Bowei Chen, Shuai Yuan, and Jun Wang. 2014. A dynamic pricing model for unifying programmatic guarantee and real-time bidding in display advertising. In *Proceedings of the Eighth International Workshop on Data Mining for Online Advertising*. ACM, 1–9.
- [4] Ye Chen, Pavel Berkhin, Bo Anderson, and Nikhil R Devanur. 2011. Real-time bidding algorithms for performance-based display ad allocation. In *Proceedings of the 17th ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 1307–1315.
- [5] Dragos Florin Ciocan and Vivek Farias. 2012. Model predictive control for dynamic resource allocation. *Mathematics of Operations Research* 37, 3 (2012), 501–525.
- [6] Ying Cui, Ruofei Zhang, Wei Li, and Jianchang Mao. 2011. Bid landscape forecasting in online ad exchange marketplace. In *Proceedings of the 17th ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 265–273.
- [7] Takanori Maehara, Akihiro Yabe, and Ken-ichi Kawarabayashi. 2015. Budget Allocation Problem with Multiple Advertisers: A Game Theoretic View. In *Proceedings of the 32nd International Conference on Machine Learning (ICML-15)*. 428–437.
- [8] H Brendan McMahan, Gary Holt, David Sculley, Michael Young, Dietmar Ebner, Julian Grady, Lan Nie, Todd Phillips, Eugene Davydov, Daniel Golovin, and others. 2013. Ad click prediction: a view from the trenches. In *Proceedings of the 19th ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 1222–1230.
- [9] Aranyak Mehta, Amin Saberi, Umesh Vazirani, and Vijay Vazirani. 2007. Adwords and generalized online matching. *Journal of the ACM (JACM)* 54, 5 (2007), 22.
- [10] Vahab Mirrokni and Hamid Nazerzadeh. 2017. Deals or no deals: Contract design for online advertising. In *Proceedings of the 26th International Conference on World Wide Web*. International World Wide Web Conferences Steering Committee, 7–14.
- [11] Yurii Nesterov. 2013. *Introductory lectures on convex optimization: A basic course*. Vol. 87. Springer Science & Business Media.
- [12] William Vickrey. 1961. Counterspeculation, auctions, and competitive sealed tenders. *The Journal of finance* 16, 1 (1961), 8–37.
- [13] Weinan Zhang, Shuai Yuan, and Jun Wang. 2014. Optimal real-time bidding for display advertising. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 1077–1086.