Optimal Bidding, Allocation and Budget Spending for a Demand Side Platform Under Many Auction Types

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

Zheng Wen
Adobe Research
zwen@adobe.com

Kuang-chih Lee
Alibaba Inc.
kuang-chih.lee@alibaba-inc.com

ABSTRACT

We develop a novel optimization model to maximize the profit of a Demand-Side Platform (DSP) while ensuring that the budget utilization preferences of the DSP's advertiser clients are adequately met. Our model is highly flexible and can be applied in a Real-Time Bidding environment (RTB) with *arbitrary* auction types, e.g., both first and second price auctions. Our proposed formulation leads to a non-convex optimization problem due to the joint optimization over both impression allocation and bid price decisions. Using Fenchel duality theory, we construct a dual problem that is convex and can be solved efficiently to obtain feasible bidding prices and allocation variables that can be deployed in a RTB setting. With a few minimal additional assumptions on the properties of the auctions, we demonstrate theoretically that our computationally efficient procedure based on convex optimization principles is guaranteed to deliver a globally optimal solution. We conduct experiments using data from a real DSP to validate our theoretical findings and to demonstrate that our method successfully trades off between DSP profitability and budget utilization in a simulated online environment.

KEYWORDS

Demand-Side Platforms; Real-Time Bidding; Online Advertising; Optimization

ACM Reference format:
Alfonso Lobos, Paul Grigas, Zheng Wen, and Kuang-chih Lee. 2017. Optimal Bidding, Allocation and Budget Spending for a Demand Side Platform Under Many Auction Types. In *Proceedings of ADKDD'17, Halifax, NS, Canada, August 14, 2017,* 6 pages.
DOI: 10.1145/3124749.3124761

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. Advertisers usually use the service of companies called demand-side platforms (DSP) to achieve this goal.

In a DSP, each individual advertiser 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 it would like to target (e.g., males, ages 18-35, who view news articles on espn.com) along with the particular media that it would like to display to the target audience (e.g., a video ad for beer). In this work we will call an impression type to a specific collection of those attributes (*e.g.*, male, California, interested in sports). 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 the bid values.

The advertisers who work with the DSP expect its budget to be spent fully or at least in an adequate amount as their marketing areas count on it. Failure to do so may motivate an advertiser to stop working with the DSP in the future, which is unacceptable for its business. In addition, they would like their budget to be spend smoothly if possible. Then, the DSP faces the problem of maximizing its profit while ensuring an adequate budget spending for its advertisers clients.

DSPs can charge their clients using several pricing schemes, for example in a CPM format advertisers are charged a fixed amount per thousand of impressions showed to users (which is mostly used for branding of products). If the advertisers are interested in some click or action of interest, they may pay in CPM scheme, but requiring that no more than certain amount per click or action of interest is paid (action of interest could be filling a form, purchasing a product, etc.). In this work, we we will assume that the DSP gets paid only when a click or action of interest occurs, but has to pay to the ad exchanges for each impression it acquires. This is a challenging payment setting as the DSP may have a negative operation if the actions or clicks of interest don't occur at the rates the DSP expects. It is important to mention that DSPs usually receive millions of ad requests opportunities per minute, and their bidding systems needs to respond to each of this ad request in matter of milliseconds making most companies apply simple heuristics to bid in the RTB systems. To simplify notation we will assume in this work that the advertisers are interested in clicks of interest, while this work apply in general to any action of interest.

ADKDD'17, Halifax, NS, Canada

DOI: 10.1145/3124749.3124761

As a final remark ad exchanges may use different auctions types to sell advertisement opportunities. As an example, several ad exchanges such as OpenX, AppNexus have announced that they use first price auctions, *i.e.* the highest bidder pays the ad exchange the amount it offered, while others like Google's AdX use second price auctions which is that the highest bidder pays the second highest bid submitted to the ad exchange. This add an extra layer to any general DSP optimization algorithm that may want to bid in different ad exchanges for the same advertisers.

In this paper, we propose a novel approach to maximize the DSP profit while ensuring an adequate budget spending for its advertisers clients. We take into account that the DSP may bid in different ad exchanges who may use different auction rules. Appropriately modeling the impression arrival, auction, and click/action processes our non-convex model gives as an output bidding and allocation vectors that can be used in real time by a DSP to bid in RTB environments. To solve our model we propose a dual formulation using Fenchel conjugates and derive a two-phase primal-dual procedure to solve our non-convex problem. We show that the solutions given by our solution procedure are optimal for several first and second price auctions, results that up to our knowledge are novel in the literature. Experimental results show how our methodology is able to trade off DSP profitability for better budget spending for first and second price auctions in synthetic data, and data based on a real DSP operation.

Due to space limitations we only review works very close to ours, and of those who we take ideas from. In terms of finding biding and allocation schemes different schemes have been suggested in the literature from the ad exchange point of view [1, 3, 6], and from the DSP side [4, 9]. In terms of spending the advertisers budget adequately [8, 12] set smart pacing strategies. Strategies for bidding using Lagrangian schemes for DSPs have appeared [10, 13] and who use the Ipinyou dataset to validate their results [14] as us. Here we formulate a dual problem using the concept of Fenchel conjugates [2, p. 91], which we solve using standard subgradients methods. Our results are similar in spirit to the recent work [11]. The latter studies a non-convex multi-agent optimization problem and also uses Fenchel conjugates to construct a dual problem. Our work differs from the latter as we are able to obtain stronger theoretical results in comparison to [11] using the structure of the online advertising problem studied here (which makes our proofs unique).

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. Section 3 show our proposed optimization problem. In Section 4 we propose a dual for our problem of interest, showing several properties of it and proposing a two-phase primal-dual scheme. In Section 5 we show important optimality results and propose two-phase primal-dual scheme to solve our problem of interest. Experimental results using the Ipinyou data [14] are presented in Section 6, and we conclude with some remarks about possible future work.

2 MODEL FOUNDATIONS

Let us begin by describing the basic structure and flow of events in the model. Let $\mathcal{K}$ denote the set of all campaigns associated with advertisers managed by the DSP. The DSP interacts with several ad exchanges, and recall that each auction held by one of these ad exchanges represents an opportunity to show an ad to a particular impression (i.e., a user). Although there may be billions of possible impression opportunities each day, we assume that the DSP uses a procedure for mapping each impression opportunity to an *impression type*. Let $\mathcal{I}$ denote the set of all such impression types. Whenever an opportunity for an impression of type $i \in \mathcal{I}$ arrives to one of the ad exchanges, the DSP has to make two real-time strategic decisions related to the corresponding 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 pays the corresponding market price (which depends on the auction type) to the ad exchange, and an ad from campaign k is displayed to the user. The advertiser corresponding to campaign k is charged only if the user clicks on the ad.

Key Parameters for Impression Types and Campaigns. Our model presumes that the DSP has knowledge (or estimates) of the following parameters:

- s_i: the expected number of impressions of type i that will arrive during the planning horizon.
- m_k: the (advertiser selected) budget for campaign k during the planning horizon.
- $\mathcal{I}_k$: the set of impression types that campaign k targets. (Note that each advertiser can create multiple campaigns to achieve different targeting goals.)
- $q_k > 0$: the CPC (cost per click) price for campaign k, i.e., the amount charged to the associated advertiser each time a user clicks on an advertisement from campaign k.

Auction Modeling. We take a flexible approach to auction modeling. In particular, we simply presume that, for each impression type $i \in \mathcal{I}$, the DSP has constructed the following two bid landscape [5] functions (which include first and second price auctions):

- $\rho_i(b)$: the probability of winning an auction for an impression of type $i \in \mathcal{I}$ given that the DSP submitted a bid of b.
- $\beta_i(b)$: the expected amount the DSP pays the ad exchange, conditional on the DSP winning the auction with a submitted bid of b.

We will assume $\rho_i(b)$ and $\beta_i(b)$ to be non-decreasing functions, and $\beta_i(b) \leq b$ for all $b \geq 0$ and $i \in \mathcal{I}$.

Click Events. Whenever an ad of campaign $k \in \mathcal{K}$ is shown to an impression of type $i \in \mathcal{I}$ (after the DSP wins the corresponding auction), we presume that a click event happens with probability θ_{ik}. In other words, θ_{ik} is the expected *click-through-rate*. In addition, given an impression type $i \in \mathcal{I}$ and a campaign $k \in \mathcal{K}$, let r_{ik} denote the corresponding *expected revenue* earned by the DSP, which is the same as the *expected cost per impression (eCPI)* to the advertiser. Namely, it holds that $r_{ik} := q_k \theta_{ik}$ where q_k is the CPC price defined earlier.

Decision Variables and Additional Notation. When the DSP has the opportunity to participate in an auction for an impression of type $i \in \mathcal{I}$ it needs to decide which campaign $k \in \mathcal{K}$ to bid on behalf of and the bid value to submit. 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. For each edge $(i, k) \in \mathcal{E}$, we define two decision variables: *(i)* x_{ik} the probability that the DSP selects campaign k, and *(ii)* b_{ik} the bid value to submit to 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. (The fraction of impression type i auctions for which the DSP decides to not bid is $1 - \sum_{k \in \mathcal{K}_i} x_{ik}$.) 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. 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 $f(\cdot) : S \to \mathbb{R}$, 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 $f(\cdot) : \mathbb{R}^n \to \mathbb{R}$ 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 $[\cdot]_+$ be the function that returns the maximum between the input and 0, and $'$ denote a derivative in the right context.

3 OPTIMIZATION FORMULATION[1]

Let us begin by recalling the model developed in [6] (proposed only for second price auctions there), which aims to maximize the profit of the DSP under budget constraints:

$$\begin{aligned}
\max_{\mathbf{x},\mathbf{b}} \quad & \sum_{(i,k)\in\mathcal{E}} [r_{ik} - \beta_i(b_{ik})] s_i x_{ik} \rho_i(b_{ik}) \\
\text{subject to} \quad & \sum_{i \in \mathcal{I}_k} r_{ik} s_i x_{ik} \rho_i(b_{ik}) \leq m_k \quad \text{for all } k \in \mathcal{K} \\
& 0 \leq b_{ik} \leq \bar{b}_i \quad \text{for all } (i,k) \in \mathcal{E} \\
& \sum_{k \in \mathcal{K}_i} x_{ik} \leq 1 \quad \text{for all } i \in \mathcal{I} \\
& x_{ik} \geq 0 \quad \text{for all } (i,k) \in \mathcal{E}
\end{aligned} \tag{1}$$

The first set of constraints above specify that the expected budget spent by each campaign should be less than the total available budget. The second set of constraints bounds the range of the bid prices, and the third and fourth set of constraints ensure that $\mathbf{x}$ represents a valid probability vector when restricted to each impression type. The objective function is the expected DSP profit, which we aim to maximize. Indeed, note that for each pair $(i,k) \in \mathcal{E}$ the quantity $r_{ik} - \beta_i(b_{ik})$ is the expected profit earned by the DSP whenever an ad of campaign k is show to an impression of type i, and $s_i x_{ik} \rho_i(b_{ik})$ is the expected number of impression of type i that we will acquire on behalf of campaign k. Therefore, $[r_{ik} - \beta_i(b_{ik})] s_i x_{ik} \rho_i(b_{ik})$ is the expected profit due to bidding for impressions of type i on behalf of campaign k, and summing these quantities over all pairs $(i,k) \in \mathcal{E}$ yields the total expected profit for the DSP, which we call $\pi(\mathbf{x}, \mathbf{b})$.

Notice that the previous formulation does *not* ensure or even encourage an adequate budget spending for the campaigns, it only ensures that each campaign does not spend in expectation more than its total budget. In reality, advertisers are not satisfied by merely ensuring that their spending on each campaign is below the specified budget level. Rather, most advertisers view the budget value m_k as a "target" and may have complex preferences regarding their spending behaviors. For example, an advertiser may be very dissatisfied with underspending behavior and may in fact prefer slightly overspending above the budget value m_k instead of severely underspending well below m_k.

In order to greatly enhance the flexibility of our model as well as its ability to capture complicated budget spending preferences, we replace the budget constraints in (1) with a more general utility function model as follows. First, note that the expected total spending of campaign $k \in \mathcal{K}$, as a function of the decision variables, is given by $v_k(\mathbf{x}, \mathbf{b}) := \sum_{i \in \mathcal{I}_k} r_{ik} s_i x_{ik} \rho_i(b_{ik})$. Now, let $u_k(\cdot) : \mathbb{R}_+ \to \mathbb{R}$ be a concave utility function describing the budget spending preferences of campaign k, whereby $u_k(v_k)$ is the "utility" of campaign k when its expected spending level is v_k. Furthermore, define the vector of expected spending levels $v(\mathbf{x}, \mathbf{b}) \in \mathbb{R}^{|\mathcal{K}|}$ whose k^{th} coordinate is $v_k(\mathbf{x}, \mathbf{b})$, and let $u(\cdot) : \mathbb{R}^{|\mathcal{K}|} \to \mathbb{R}$ be the overall budget spending utility function whereby $u(v(\mathbf{x}, \mathbf{b})) = \sum_{k \in \mathcal{K}} u_k(v_k(\mathbf{x}, \mathbf{b}))$. Finally, as an extra way to simplify notation let's define the feasible set of allocation and bidding variables:

$$\mathcal{S} := \left\{ (\mathbf{x}, \mathbf{b}) : \sum_{k \in \mathcal{K}_i} x_{ik} \ \leq \ 1 \text{ for all } i \in \mathcal{I}, \ 0 \leq b_{ik} \leq \bar{b}_i \text{ for all } (i,k) \in \mathcal{E}, \ \mathbf{x} \geq 0 \right\}$$

We are now ready to state our proposed optimization model:

$$F^* := \max_{(\mathbf{x},\mathbf{b}) \in S} \sum_{(i,k)\in\mathcal{E}} [r_{ik} - \beta_i(b_{ik})] s_i x_{ik} \rho_i(b_{ik}) + u(v(\mathbf{x}, \mathbf{b})) \tag{2}$$

Note that problem (2) is non-convex, and in Section 4 we propose a computationally efficient procedure based on convex duality. We finish this section by showing three examples of utility functions that illustrate the improved generality and flexibility of model (2).

Examples of Utility Functions

(1) Formulation (1) may be recovered as a special case of the more general problem (2) by letting $u_k(\cdot)$ be the (extended real valued) concave function such that $u_k(v_k)$ equals $-\infty$ if v_k is strictly greater than m_k, and 0 otherwise.

(2) If we want to maximize the DSP profit but also try to enforce an appropriate target spending for a campaign $k \in \mathcal{K}$, we can take $u_k(\cdot)$ to be the concave function such that $u_k(v_k)$ equals $-\infty$ if v_k is strictly greater than m_k, and $-\frac{\tau_k}{2}\|v_k - m_k\|_2^2$ otherwise. Here $\tau_k \geq 0$ is a user defined penalization constant.

(3) If we want to maximize the DSP profit while requiring both a minimum and maximum expected spending for campaign $k \in \mathcal{K}$, we can take $u_k(\cdot)$ to be the concave function such that $u_k(v_k)$ equals $-\infty$ if v_k is strictly greater than m_k or strictly less than $\alpha_k m_k$, and 0 otherwise. Here, the parameter $\alpha_k \in [0, 1]$ is user defined and represents the minimum percentage of expected budget spending.

Note that the model (2) allows for each campaign to have its own distinct utility function $u_k(\cdot)$, and therefore the three examples above may be combined together across the different campaigns, for example. Finally, note also that the separable structure of the utility function $u(\cdot)$, whereby $u(\cdot) = \sum_{k \in \mathcal{K}} u_k(\cdot)$, is actually not needed for all of the results that we develop herein. Indeed, the only crucial assumption about $u(\cdot)$ is that $u(\cdot)$ is a concave function. However, the separable structure is quite natural and all of our examples do have this separable structure as well, so for ease of presentation we present the model in this way.

[1]THE PROOFS OF THIS ARTICLE CAN BE FOUND IN ARXIV:1805.11645.

4 DUAL OPTIMIZATION PROBLEM AND SCHEME

We begin this section with a high-level description of our approach for solving (2). Our algorithmic approach is based on a two phase procedure. In the first phase, we construct a suitable dual of (2), 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 a convex optimization problem that results when $\mathbf{b}$ is fixed in order to recover allocation probabilities $\mathbf{x}$. A very mild assumption we make for the rest of the paper is that $F^* > -\infty$, otherwise there would be no optimization problem to optimize.

As we mentioned before we have assumed that our utility function $u(\cdot)$ is concave, therefore $-u(\cdot)$ is a convex function and we can define $p(\cdot) : \mathbb{R}^{|\mathcal{K}|} \to \mathbb{R}$ its convex conjugate as $p(\lambda) := \sup_{v \in \mathbb{R}^{|\mathcal{K}|}} \{\lambda^T v + u(v)\}$ which is a convex function. The Fenchel Moreau Theorem [2, p. 91] ensures that $u(v) = \inf_{z \in \mathbb{R}^{|\mathcal{K}|}} \{p(z) - z^T v\}$. Using the latter we can re-write our primal problem as (here we use $\pi(\mathbf{x}, \mathbf{b}) := \sum_{(i,k)\in\mathcal{E}} [r_{ik} - \beta_i(b_{ik})] s_i x_{ik} \rho_i(b_{ik})$):

$$F^* := \max_{(\mathbf{x},\mathbf{b})\in S} F(\mathbf{x},\mathbf{b}) := \left\{ \pi(\mathbf{x}, \mathbf{b}) + \inf_{\lambda \in \mathbb{R}^{\mathcal{K}}} \left(-\lambda^T v(\mathbf{x}, \mathbf{b}) + p(\lambda) \right) \right\}$$

For a given $\lambda \in \mathbb{R}^{|\mathcal{K}|}$ we will define the dual function as:

$$Q(\lambda) := \sup_{(\mathbf{x},\mathbf{b})\in S} \left(\pi(\mathbf{x}, \mathbf{b}) - \lambda^T v(\mathbf{x}, \mathbf{b}) \right) + p(\lambda)$$

And the dual problem as:

$$Q^* := \min_{\lambda \in \mathbb{R}^{|\mathcal{K}|}} Q(\lambda)$$

Then, the following inequalities hold for our primal and dual formulations (they follow from the max-min inequality [2, p. 238]):

$$Q(\lambda) \geq Q^* \geq F^* \geq F(\mathbf{x},\mathbf{b}) \text{ for all } \lambda \in \mathbb{R}^{|\mathcal{K}|}, (\mathbf{x}, \mathbf{b}) \in S$$

Let's now define $h_i(z, b) := (z - \beta_i(b)\rho_i(b)$ the expected profit the DSP receives from showing an ad of campaign k to an impression of type i which has an expected revenue of z when submitting a bid of value b. Then, given a fixed $\lambda \in \mathbb{R}^{|\mathcal{K}|}$ the dual function $Q(\lambda)$ can re-defined as:

$$Q(\lambda) := \underset{(\mathbf{x},\mathbf{b})\in S}{\text{maximize}} \sum_{k\in\mathcal{K}} \sum_{i\in\mathcal{I}} h_i(r_{ik}(1 - \lambda_k), b) s_i x_{ik} + p(\lambda)$$

Proposition 4.1 (Efficient computation of $Q(\lambda)$). *Given a dual variable λ, an optimal solution $(\boldsymbol{x}(\lambda), \boldsymbol{b}(\lambda))$ for the dual function $Q(\lambda)$ can be found using Algorithm 1.*

Theorem 1 shows that the dual problem can be solved in a parallel fashion, and furthermore finding $\mathbf{b}(\lambda)$ can be a simple operation. For example, in the case of a second price auction it is known that $[r_{ik}(1 - \lambda_k)]_+ \in \arg\max_{b\in[0,\bar{b}_i]} \; h_i(r_{ik}(1 - \lambda_k), b)$, and some examples for first price auctions have nice close forms as shown in the next section. Being able to solve $Q(\lambda)$ efficiently is of great importance as it is a core component to find a subgradient of $Q(\lambda)$ as the following theorem shows:

Proposition 4.2. *Given $\lambda \in \mathbb{R}^{|\mathcal{K}|}$ the output of Algorithm 2 is a vector $g \in \partial Q(\lambda)$.*

Algorithm 1 Computing $(\mathbf{x}(\lambda), \mathbf{b}(\lambda)) \in \arg Q(\lambda)$ for a fixed λ

Input: $\lambda \in \mathbb{R}^{|\mathcal{K}|}$.

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

$$b(\lambda)_{ik} \leftarrow \arg \max_{b\in[0,\bar{b}_i]} \; h_i(r_{ik}(1 - \lambda_k), b), \text{ and}$$

$$\pi(\lambda)_{ik} \leftarrow h_i(r_{ik}(1 - \lambda_k), b(\lambda)_{ik}) s_i$$

2. For each $i \in \mathcal{I}$, pick $k_i^* \in \max_{k\in\mathcal{K}_i} \{\pi(\lambda)_{ik}\}$ arbitrary. Set $x(\lambda)_{ik} = 0$ for all $k \in \mathcal{K}_i \neq k_i^*$, and if $\pi(\lambda)_{ik_i^*} > 0$ make $x(\lambda)_{ik_i^*} = 1$ otherwise $x(\lambda)_{ik_i^*} = 0$.

Output: $(\mathbf{x}(\lambda), \mathbf{b}(\lambda))$.

Algorithm 2 Computing a subgradient of $Q(\lambda)$

Input: $\lambda \in \mathbb{R}^{|\mathcal{K}|}$

1. Obtain $(\mathbf{x}(\lambda), \mathbf{b}(\lambda)) \in \arg\max Q(\lambda)$ using Algorithm 1.
2. Obtain $p' \in \partial p(\lambda)$

3 Set:

$$g(\lambda) \leftarrow p' - v(\mathbf{x}(\lambda), \mathbf{b}(\lambda))$$

Output: $g(\lambda) \in \partial Q(\lambda)$.

Notice that with STEP 1. in Algorithm 1 we can recover bidding prices $\mathbf{b}(\lambda)$ from dual variables $\lambda \in \mathbb{R}^{|\mathcal{K}|}$. The final Theorem of this section shows that for fixed bidding prices $\mathbf{b}$ is easy to obtain allocation probabilities $\mathbf{x}$ by solving problem (2).

Proposition 4.3. *For fixed bidding prices* $\mathbf{b}$*, problem* (2) *is a convex problem.*

Proposition 4.3 tell us that we could use a sub-gradient method to find to find an allocation vector $\mathbf{x}$ given a fixed $\mathbf{b}$. Better than the previous, depending on the utility function used (2) can have a nice structure, for example for the utility function examples shown in the previous section, examples 1. and 3. transform problem (2) in a linear program and example 2. in a quadratic problem. Problems that could be solved directly using solvers like Gurobi [7]. We finish this section by presenting Algorithm 3 which formalize our approach to solve problem (2).

Algorithm 3 Two Phase primal-dual Scheme

Phase 1: Solve the Dual problem.

Solve $\min_\lambda Q(\lambda)$ to near global optimality
using a subgradient method obtaining dual variables $\hat{\lambda}$.

Phase 2: Primal Recovery

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

$$b_{ik}(\hat{\lambda}) \leftarrow \arg \max_{b\in[0,\bar{b}_i]} \; h_i(r_{ik}(1 - \hat{\lambda}_k), b)$$

2. Using bid prices $\mathbf{b}(\hat{\lambda})$ solve (2) obtaining allocation probabilities $\mathbf{x}(\hat{\lambda})$.

Output: Feasible primal solution $(\mathbf{x}(\hat{\lambda}), \mathbf{b}(\hat{\lambda}))$.

---

5 ZERO DUALITY GAP RESULTS

Algorithm 3 can always be used as long as the parameters and functions of problem (2) are well defined. Here we we will go further and show that our dual formulation and dual scheme are the right methods to solve (2). In particular, we have strong duality results which to the best of our knowledge are novel and have important applications to first and second price auctions by showing optimal bidding prices to be used in an RTB environment. These will be derived from the following theorem:

Theorem 5.1. *If for all $i \in \mathcal{I}$ we have that $\rho_i(\cdot)$ and $\beta_i(\cdot)$ are differentiable and:*

- $\rho_i'(b) > 0$ *for all* $b \in [0, \bar{b}_i]$.
- $g_i(b) = \frac{(\rho_i(b)\beta_i(b))'}{\rho_i'(b)}$ *is strictly increasing for all* $b \in [0, \bar{b}_i]$.

for all $i \in \mathcal{I}$, then for any $\lambda^* \in \arg\max_{\lambda \in \mathbb{R}^{|\mathcal{K}|}} Q(\lambda)$, we have $F(\mathbf{b}^*, \mathbf{x}^*) = Q(\lambda^*)$ for any $(\mathbf{x}^*, \mathbf{b}^*) \in \arg\max_{(\mathbf{b},\mathbf{x}) \in S} Q(\lambda^*)$.

Notice that Theorem 5.1 ensures that no duality gap exists, but furthermore for an optimal dual variable λ^* it gives a form of the variables $(\mathbf{x}^*, \mathbf{b}^*)$ such that $F(\mathbf{x}^*, \mathbf{b}^*) = Q(\lambda^*)$. Also, notice that the second condition of the theorem is a form of diminishing returns.

Applications of Theorem 5.1

(1) If for all $i \in \mathcal{I}$ their auctions are second price and $\rho_i(\cdot)$ is a strictly increasing function in $[0, \bar{b}_i]$, then Theorem 5.1 holds. Also, for an optimal dual variable λ^* optimal bidding prices are $b(\lambda^*)_{ik} = [\min\{\bar{b}_i, r_{ik}(1 - \lambda_k^*)\}]_+$ for all $(i,k) \in \mathcal{E}$.

(2) If for all $i \in \mathcal{I}$ all auctions are first price auctions or more generally scaled first price auctions in which the winning DSP pays an $\alpha \in (0,1]$ percentage of the bid it offered, Theorem 5.1 holds when $\rho_i(\cdot)$ is a strictly increasing twice differentiable concave function. Example of the later are: 1) $\rho(b) = \frac{b}{c+b}$ for some fixed $c > 0$, 2) square roots and logarithms, 3) $\rho(\cdot)$ representing the cumulative distribution function of an exponential or logarithm-exponential random variable.

(3) If for all $i \in \mathcal{I}$ all auctions are first scaled price auctions with $\alpha \in (0,1]$ in which for each impression i there is a fixed number $n_i \geq 1$ of other DSPs who bid independently as $Uniform(0, \bar{b}_i)$, then Theorem 5.1 holds. Furthermore, for an optimal dual variable λ^* optimal bidding prices are $b(\lambda^*)_{ik} = \left[\min\left\{\bar{b}_i, \frac{n_i r_{ik}(1-\lambda_k^*)}{\alpha(n_i+1)}\right\}\right]_+$ for all $(i,k) \in \mathcal{E}$.

(4) Any combination of the above, or cases in which each impression type satisfies the conditions of Theorem 5.1.

We finish this section by making three important comments. First, to obtain the form of the optimal bidding prices is only needed to solve $\arg\max_{b \in [0, \bar{b}_i]} h_i(r_{ik}(1-\lambda_k^*), b)$. In many cases, like second price auctions, this will have a close form, but for many others the DSP can have tables with approximate solutions that can be used instead of solving the problem in real time. Second, many ad-exchanges use what are called hard reserve prices that consider a bid valid only if it is higher than the reserve price. This poses a problem for Theorem 5.1 as the condition of $\rho(\cdot)$ being strictly non-decreasing would not be true. If the impression types had fixed hard reserve prices this is not a major issue as we can change the model to bid between the reserve price and $\bar{b}_i$ (if the reserve price were higher than $\bar{b}_i$ the model wouldn't bid for that impression type). In the case that hard reserve prices change dynamically, heuristics could be used, *e.g.* considering bidding in real time only for those campaigns with bid values higher than the reserve price, putting levels of reserve price as a field in the impression types, and others which we don't explore here. Third, it can be proven that Theorem 5.1 guarantees that Algorithm 3 would converge to an optimal solution for (2) as we get better λ solutions of the dual problem. For space reasons we don't extend on this topic here.

6 COMPUTATIONAL EXPERIMENTS

Here we present computational results using the Ipinyou DSP data [14] to which we applied our two-phase solution procedure comparing its performance w.r.t. a greedy heuristic. The way we suggest and apply in our experiments an allocation and bidding variables $(\mathbf{x}, \mathbf{b})$ in a practical RTB environment is shown in Policy 4 and the heuristic to which we compare our method is shown in Policy 5.

The Greedy Heuristic is optimal for the case of infinite budgets and second price auctions (is easy to extend it to arbitrary auction types). As an important remark our method should be used in a real DSP operation inside a Model Predictive Control Scheme, which calls Algorithm 3 as budgets get used and different model inputs gets updated as time progresses.

Policy 4 Online Policy Implied by **(x,b)**

Input: Allocation and bidding variables $(\hat{\mathbf{x}}, \hat{\mathbf{b}})$ and a new impression arrival of type $i \in \mathcal{I}$.

1. Sample a campaign $\tilde{k} \in \mathcal{K}_i$ according to the distribution implied by the values x_{ik} for $k \in \mathcal{K}_i$ or break with probability $1 - \sum_{k \in \mathcal{K}_i} x_{ik}$ (*i.e.* choosing no campaign). If some campaigns have depleted its budget adjust the probabilities to take this fact into account.
2. Enter bid price $\hat{b}_{i\tilde{k}}$. If the auction is won, then pay the ad exchange the amount is asking. If the auction is not won, then break.
3. Show an ad for campaign $\tilde{k}$. If a click or action of interest happens, then deduct $q_{\tilde{k}}$ from the budget of campaign $\tilde{k}$ and earn revenue $q_{\tilde{k}}$.

Policy 5 Greedy Policy

Input: New impression arrival $i \in \mathcal{I}$.

1. Bid for a campaign $k \in \mathcal{K}_i$ with remaining budget bigger than q_k with the highest r_{ik} value.

The public available Ipinyou data [14] contains information about real bidding made by the chinese DSP Ipinyou in 2013. It contains different features including the bidding prices of the impressions for which Ipinyou bid for, and the price paid by Ipinyou to the ad-exchange in case an impression was won and if a click or conversion occurred (we did not use conversion data). Ipinyou assumes that ad-exchanges use second price auctions. The data is already divided in train and test sets and it has been used to test bidding strategies for DSPs [10, 13] but we haven't found a paper that use it for both bidding and allocation strategies, reason why we compare to the Greedy Heuristic 5. Ipinyou data is divided in three different time periods in 2013, of those we decided to use the third as in the first there is no information about the campaigns Ipinyou bid for, and in the second Ipinyou assumed that all impression types could serve all campaigns which make the impression-campaign graph non-interesting. The third season contains 3.15M and 1.5M logs of impressions won by Ipinyou in the train and test set resp. in behalf of four advertisers, which have 2716 and 1155 clicks associated to them. Here we use the different advertisers as our campaigns.

To create "impression types" we divided the impressions by the visibility feature which has a strong correlation with CTR, and then by the regions, homepage url, and "width x height" of the ad to be shown (features that appear in all impression logs). The last three features have a high dimensionality, for example homepage url have 54,108 unique urls. For that reason we created mutually exclusive sets of the form all urls that were targeted only by advertiser 1, all that were targeted only by advertisers 1 and 3, etc. With this technique we partition all impressions for which Ipinyou bid for in 160 clusters of impressions which we used to create our final partition of 23 impression types. Of those, 19 corresponds to the clusters with a minimum of 30,000 impressions won in the train set and the 4 left are the union of all clusters having different visibility attribute (we grouped together the second, third, fourth and fifth view as if they were the same visibility type). Our final graph is composed of 4 campaigns, 23 impression types, it has 43 edges, and the CTR values were taken as the empirical rates for each pair of (impression type, advertiser). Using only the impressions won in the train set for each impression type i we fitted a beta distribution using the python Scipy package (imposing the location parameter to be equal to zero) to obtain parameters to estimate the bid landscape functions (for a given b value the calculation of $\rho_i(b)$ is just a Scipy function call, while $\beta_i(b)$ had to be approximated using Monte-Carlo). Finally, we count the times each impression type appears in the test set to create the s_i values, and the budgets correspond to the total amount of money that Ipinyou paid for the impressions assigned to each advertiser in the test set. To simulate a real time environment we used the empirical train CTR to train our models and the greedy heuristic (we took the average CTR per impression type for the pairs of (impression type, advertiser) that appear less than 5000 times in the train set). To test our model we use the impressions won by Ipinyou in the actual order saving their impression type. Then, one by one we read each impression log and we assume that the impression was won if the proposed bid for it is higher than the amount Ipinyou paid for it. A click for the proposed advertiser occurs with probability equal to the empirical CTR from the test data for the pair (impression type, advertiser). We tried the utility functions 1. and 2. from Section 3, using $\tau_k = 1/m_k$ for all $k \in \mathcal{K}$ for the latter u. function.

Results. Our results are shown in Figure 1. Here we define budget utilization (b.u.) as the percentage of the total budget that was used at the end of one simulation, and u.f. stands for utility function. We performed two experiments, the first corresponding to the tables and the second to the graph in Figure 1. The first studies the sensitivity of our model w.r.t. to the budget. We tried 1/32, 1/8, 1/4, 1/2 and 1/1 of the budgets Ipinyou used for each advertiser in the test data, running 100 simulations for each setting. We report the relative average profit and b.u. obtained by the utility functions w.r.t. to the ones obtained by greedy heuristic. What makes one simulation different from the other is that the CTRs are random variables. In the second experiment we multiplied the penalization parameters $\tau_k = 1/m_k$ that appear in the u.f. 2 by 0.1, then by 0.3, and so on until 2.1 running 100 simulations for each case. Then, we obtain the relative average profit and b.u. w.r.t. the u.f. 1., also obtaining these values for the greedy heuristic. Our results show that our methodology works very well for cases in which the budget is tight, but when is not the case the greedy heuristic is a good alternative. From our second experiment we can see that a better b.u. utilization can be obtained at the cost of having a worst profit (it can even be negative).

Relative Profit

(Perc./Met.)	1/32	1/8	1/4	1/2	1/1
U.F. 1	2.10	1.38	1.08	0.92	1.21
U.F. 2	1.83	1.34	0.94	0.76	0.8

Relative Budget Utilization

(Perc./Met.)	1/32	1/8	1/4	1/2	1/1
U.F. 1	0.95	0.88	0.8	0.95	1.07
U.F. 2	0.84	0.94	1.07	1.10	1.38

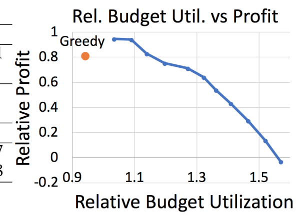

Figure 1: Two-phase policy vs. greedy policy

Let us conclude this section by mentioning a few directions for future research. It would be very valuable to perform experiments in which impressions are auctioned in both first and second price auctions and in which the cardinality of the impression types and campaigns are higher. Also, several of the quantities we assume as known in this work are hard to estimate in practice. We will study robust approaches to our model.

REFERENCES

[1] 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.

[2] Stephen Boyd and Lieven Vandenberghe. 2004. *Convex optimization*. Cambridge university press.

[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] 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.

[6] Paul Grigas, Alfonso Lobos, Zheng Wen, and Kuang-chih Lee. 2017. Profit Maximization for Online Advertising Demand-Side Platforms. In *Proceedings of the ADKDD'17 (ADKDD'17)*. ACM, New York, NY, USA, Article 11, 7 pages. DOI: http://dx.doi.org/10.1145/3124749.3124761

[7] Inc. Gurobi Optimization. 2016. Gurobi Optimizer Reference Manual. (2016). http://www.gurobi.com

[8] Kuang-Chih Lee, Ali Jalali, and Ali Dasdan. 2013. Real time bid optimization with smooth budget delivery in online advertising. In *Proceedings of the Seventh International Workshop on Data Mining for Online Advertising*. ACM, 1.

[9] Claudia Perlich, Brian Dalessandro, Rod Hook, Ori Stitelman, Troy Raeder, and Foster Provost. 2012. Bid optimizing and inventory scoring in targeted online advertising. In *Proceedings of the 18th ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 804–812.

[10] Kan Ren, Weinan Zhang, Ke Chang, Yifei Rong, Yong Yu, and Jun Wang. 2018. Bidding Machine: Learning to Bid for Directly Optimizing Profits in Display Advertising. *IEEE Transactions on Knowledge and Data Engineering* 30, 4 (2018), 645–659.

[11] Mengdi Wang. 2017. Vanishing Price of Decentralization in Large Coordinative Nonconvex Optimization. *SIAM Journal on Optimization* 27, 3 (2017), 1977–2009.

[12] Jian Xu, Kuang-chih Lee, Wentong Li, Hang Qi, and Quan Lu. 2015. Smart pacing for effective online ad campaign optimization. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, 2217–2226.

[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.

[14] Weinan Zhang, Shuai Yuan, Jun Wang, and Xuehua Shen. 2014. Real-Time Bidding Benchmarking with iPinYou Dataset. *arXiv preprint arXiv:1407.7073* (2014).