Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence (AAAI-17)

A Study of Compact Reserve Pricing Languages

MohammadHossein Bateni,[†] Hossein Esfandiari,[‡] Vahab Mirrokni,[†] Saeed Seddighin[‡] [*]
†Google, ‡University of Maryland
{bateni,mirrokni}@google.com, {esfandiari, seddighin}@cs.umd.edu

Abstract

Online advertising allows advertisers to implement fine-tuned targeting of users. While such precise targeting leads to more effective advertising, it introduces challenging multidimensional pricing and bidding problems for publishers and advertisers. In this context, advertisers and publishers need to deal with an exponential number of possibilities. As a result, designing efficient and *compact* multidimensional bidding and pricing systems and algorithms are practically important for online advertisement. Compact bidding languages have already been studied in the context of multiplicative bidding. In this paper, we study the compact pricing problem.

More specifically, we first define the *multiplicative reserve price optimization problem* (MRPOP) and show that unlike the unrestricted reserve price system, it is NP-hard to find the best reserve price solution in this setting. Next, we present an efficient algorithm to compute a solution for MRPOP that achieves a logarithmic approximation of the optimum solution of the unrestricted setting, where we can set a reserve price for each individual impression type (i.e., one element in the Cartesian product of all features). We do so by characterizing the properties of an optimum solution. Furthermore, our empirical study confirms the effectiveness of multiplicative pricing in practice. In fact, the simulations show that our algorithm obtains 90–98% of the value of the best solution that sets the reserve prices for each auction individually (i.e., the optimum set of reserve prices).

Finally, in order to establish the tightness of our results in the adversarial setting, we demonstrate that there is no *compact* pricing system (i.e., a pricing system using $O(n^{1-\epsilon})$ bits to set n reserve prices) that loses, in the worst case, less than a logarithmic factor compared to the optimum set of reserve prices. Notice that this hardness result is not restricted to the multiplicative setting and holds for any compact pricing system. In summary, not only does the multiplicative reserve price system show great promise in our empirical study, but it is also theoretically optimal up to a constant factor in the adversarial setting.

[*]Supported in part by NSF CAREER award CCF-1053605, NSF BIGDATA grant IIS-1546108, NSF AF:Medium grant CCF-1161365, DARPA GRAPHS/AFOSR grant FA9550-12-1-0423, and another DARPA SIMPLEX grant.
Copyright © 2017, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

1 Introduction

As a main advantage over traditional advertising, online advertising allows advertisers to target specific subsets of users via a very fine-tuned and descriptive targeting criteria. In these settings, both publishers and advertisers face challenging pricing and bid optimization problems in multidimensional settings. As a result, the space of possibilities for setting the price or declaring the bids are exponential, even when restricted to the few most important features. This leads to interesting problems of designing efficient and *compact* multidimensional bidding and pricing systems and algorithms, which are practically important for online advertisement. While compact bidding languages have already been studied, our goal in this work is to study the compact pricing problem.

More specifically, in the context of online advertising for sponsored search or display ads, the space of impression types of interest to the advertiser is usually very big, specially because it is the Cartesian product of several features (such as geographic location, time of day/week), domains of which have sizes typically ranging from thousands to millions. These features have a big impact on the quality and desirability of the impression, hence their importance in determining the bid. *Compact* bidding languages (and, in particular, multiplicative bidding) were proposed and are now widely used as a solution to the challenge of maneuvering in this large exponential space. Not only do these make campaign management easier and more focused, but they also mitigate the issue of over-reliance on sparse data. Such bidding languages provide a middle ground between a uniform bidding strategy ((Feldman et al. 2007b; Muthukrishnan, Pál, and Svitkina 2007a)) over all interesting impressions and the unmanageable strategy of setting up a customized campaign for each type of impression. More specifically, the advertiser specifies, in addition to targeting criteria for interesting impressions and a base bid, certain multiplicative adjustments for each feature associated with the impression. In the wake of adoption of multiplicative bidding languages by several major search engine-based advertising platforms ((Bing Ads 2014; Google Support 2014; HubSpot 2013)), (Bateni et al. 2014) studied the impact of employing such a bidding language on advertisers' welfare. They focus on two questions: (1) how expressive is this language? (2) how difficult is it to optimize over?

In this work, we ask similar questions regarding the specification of reserve prices for auctions. As was the case with optimal bids for a single advertiser, the characterizing features of an impression may have significant impact on its typical value for all advertisers and hence it plays a big role in determining the optimal reserve price.

Next we formally define the *multiplicative reserve price system*. Let $\mathcal{D}$ denote the set of all (bid) distributions over $\mathbb{R}^+$. The space of auctions is characterized by a d-dimensional type vector $\langle a_1, a_2, \ldots, a_d \rangle$ where $a_i \in A_i$ for $i \in [d]$, and $[d] = \{1, 2, \ldots, d\}$. We let $A = A_1 \times A_2 \times \cdots \times A_d$. Given as input is a mapping $B : A \mapsto \mathcal{D}$ denoting the bid distribution (where B denotes the distribution of the top bid in the auction). A multiplicative reserve price solution is specified by d vectors $v_1, v_2, \ldots, v_d$, where for a feature $a_i \in A_i$, the adjustment factor of a_i is $v_i(a_i)$. The reserve price of an auction with type $\langle a_1, a_2, \ldots, a_d \rangle$ is set to $r(\langle a_1, a_2, \ldots, a_d \rangle) = \prod_i v_i(a_i)$.

The *multiplicative reserve price optimization problem* (MRPOP) is as follows: Given mappings B and s as described above, the goal is to find a multiplicative reserve price solution $v_1, v_2, \ldots, v_d$ that maximizes the total profit $\sum_{a \in A} \pi(B(a), r(a))$, where $\pi(B, r)$ denotes the expected profit of the auction with bid distribution B and reserve price r; see Equation (1). In this paper, we first study the computation of the multiplicative reserve price optimization problem. We show that this problem is NP-hard.

Theorem 1.1. *The multiplicative reserve price optimization problem is NP-hard.*

This rules out the existence of a polynomial-time algorithm that finds an optimum solution to MRPOP, unless P = NP. Next, we investigate the properties of the optimum solutions in order to design an approximation algorithm for MRPOP. We say that a solution $(v_1, v_2, \ldots, v_d)$ is *stable* if changing any single vector v_i does not increase the total profit (i.e., it is a locally maximum solution). Notice that any optimum solution is a stable solution, as well. We say that a solution $(v_1, v_2, \ldots, v_d)$ is *polynomially bounded* if all of the elements of the vectors $v_1, v_2, \ldots, v_d$ are upper-bounded and lower-bounded by two polynomial functions. In the following theorem we bound the efficiency of any polynomially bounded stable solution.

Theorem 1.2. *The profit of any polynomially bounded stable solution is within a logarithmic factor of that of the optimum set of reserve prices.*

Next based on this structural property we design an almost efficient polynomial-time algorithm that works well both in theory and practice. In particular, this algorithm achieves 90–98% of the profit of the optimum set of reserve prices. Our results are based on a large number of two-dimensional instances where each of the two features have millions of possibilities. We design an algorithm to proves the following.

Theorem 1.3. *There exists a polynomial-time algorithm that finds a solution with profit within a logarithmic factor of that of the optimum set of reserve prices.*

Finally, we study the question of whether one can design an algorithm that works better than logarithmic in the worst case, or even, if one can design a new pricing system admitting such an algorithm. Recall that the main purpose of a compact reserve pricing language is to simplify the process of setting reserve prices for n different type of auctions simultaneously. Clearly, by using $\tilde{O}(n)$ bits of communication, one can set a single independent reserve price on each auction. Hence, the number of communication bits for a desirable pricing system should be significantly less than $\tilde{\Theta}(n)$. In the following theorem, we provide an upper bound on the approximation factor of any pricing system as a function of the number of communication bits it requires.

Theorem 1.4. *Pick an arbitrary number $\sigma \in (0, 1)$. Any $\frac{1}{4}\sigma \log n$-approximation pricing system requires $\Omega(n^{1-\sigma})$ communication bits. This also holds for the restricted case of one deterministic bid per auction.*

Plugging $\sigma = \epsilon/2$ into the above theorem gives the following corollary which shows the tightness of our results in the adversarial setting. Recall that both Theorem 1.4 and Corollary 1.4.1 hold for all pricing systems. Therefore, this lower bound rules out the existence of a better approximation factor for other pricing systems such as a pricing system that applies another operation instead of multiplication.

Corollary 1.4.1. *For any small constant ϵ there is no $o(\log(n))$-approximation pricing system using $O(n^{1-\epsilon})$ communication bits.*

Note that the above result rules out the possibility of improving the logarithmic approximation factor for a wide variety of compact reserve price settings including the compact reserve pricing systems defined in preliminaries section.

Related Work. To the best of our knowledge, this work initiates the study of compact reserve pricing systems and the multiplicative language for reserve prices. Moreover, our information-theoretic lower bound is the first such result for selling a single item in the multidimensional setting. Previous inapproximability result have been achieved for other cases like the multi-item settings (Hart and Nisan 2013; Briest et al. 2010; Chawla, Malec, and Sivan 2010). Initially, (Bateni et al. 2014) study multiplicative bidding languages on the advertiser's welfare. Indeed, this provide a middle ground between a uniform bidding strategy ((Feldman et al. 2007b; Muthukrishnan, Pál, and Svitkina 2007a)) over all interesting impressions and the unmanageable strategy of setting up a customized campaign for each type of impression. The multiplicative bidding languages has been adopted by several major search engine-based advertising platforms (Bing Ads 2014; Google Support 2014; HubSpot 2013). In the recent decade, motivated by online advertisements several works has been initiated on optimization under budget constraints, both from publishers' point of view ((Mehta et al. 2007; Karande, Mehta, and Srikant 2013; Goel, Mirrokni, and Leme 2015; Devanur and Hayes 2009; Charles et al. 2013)) and advertisers' point of view ((Archak, Mirrokni, and Muthukrishnan 2010; Borgs et al. 2007; Feldman et al. 2007a; Muthukrishnan, Pál, and Svitkina 2007b)).

Due to space constraints, some of the proofs and figures are omitted in this version and included in the full version.

2 Preliminaries

In this section we formally define the *multiplicative reserve price optimization problem* (MRPOP). As mentioned earlier, in this problem we seek to find the optimal multiplicative reserves to maximize the expected revenue of several posted-price auctions in total. Our focus is on the Bayesian setting, that is, for each auction we assume we're given the distribution of the bids in advance. Recall that, for a bid distribution D, the (expected) revenue of a posted-price auction with reserve price ρ is

$$\pi(D,\rho) = \rho \cdot \Pr_{b\sim D}[b \geq \rho]. \tag{1}$$

The space of auctions is characterized by a d-dimensional auction type where the set of all possible realizations of each type i is given as A_i. More precisely, every auction corresponds to a type vector $\langle a_1, a_2, \ldots, a_d\rangle$ where every $a_i \in A_i$. We let $A = A_1 \times A_2 \times \cdots \times A_d$ be the set of all type vectors for which we run a posted price auction. Given as input is a mapping $B : A \mapsto \mathcal{D}$ denoting the bid distribution of each auction type. For a reserve mapping function $\mathcal{M} : {\mathbb{R}^+}^d \mapsto \mathbb{R}^+$, we represent a compact solution of the problem by a list of d vectors $v = \langle v^1, v^2, \ldots, v^d\rangle$, where the size of each v^i matches that of A_i. For every auction type $\langle a_1, a_2, \ldots, a_d\rangle$ then, the corresponding reserve price would be equal to $\mathcal{M}(v^1_{a_1}, v^2_{a_2}, \ldots, v^d_{a_d})$. The emphasis of our work is on the case where

$$\mathcal{M}(r_1, r_2, \ldots, r_n) = \prod v^i_{a_i},$$

or in words, the corresponding reserve price of an auction type is the product of the pertinent coefficients. In this case, we call the problem MRPOP and refer to a compact solution of this type as a *multiplicative reserve price solution*.

Given a mapping B, our goal is to find a multiplicative reserve price solution $v^1, v^2, \ldots, v^d$ that maximizes the total profit

$$\mathsf{Rev}(v) = \sum_{a\in A} \pi(B(a), r(v,a)),$$

where $r(v,a)$ is the reserve price of the multiplicative solution v for auction type a.

When $d = 2$, one can think of $A = A_1 \times A_2$, as a table of auctions where every type vector corresponds to a cell of the table. Moreover, a multiplicative reserve price solution $v = \langle v^1, v^2\rangle$ can be represented as a list of multiplicative reserves on the rows and the columns of the table. As declared, the reserve price of every cell is the multiplication of the reserve prices of its corresponding row and column.

We denote by R, an upper bound on the range of the distributions. In other words, we assume all of the bid distributions are over the integer numbers in interval $[0, R]$. Also, n denotes the total number of type vectors in our setting.

3 $\Omega(\log n)$ Gap

A pricing system $\Psi : \Gamma \to \mathbb{R}^n$ maps a key $\gamma \in \Gamma$ from a set of predefined keys Γ to a vector of reserve prices $\vec{r} = \langle r_1, r_2, \ldots, r_n\rangle$. Clearly, the number of bits required to represent a key in Γ is lower bounded by $\log_2|\Gamma|$. This is the number of bits one needs to communicate with the pricing system to set the reserve prices.

We say a pricing system $\Psi : \Gamma \to \mathbb{R}^n$ is an α-approximation pricing system if for any vector of bid distributions $\vec{D} = \langle D_1, D_2, \ldots, D_n\rangle$, there exists a key $\gamma \in \Gamma$ such that the profit of reserve prices $\Psi(\gamma)$ on $\vec{D}$ is at least α times that of the optimum reserve prices.

Theorem 3.1 bounds the number of bits one needs to communicate with an α-approximation pricing system.

Theorem 3.1. *Pick an arbitrary constant $\sigma \in (0,1)$. Any $\frac{1}{4}\sigma\log n$-approximation pricing system $\Psi : \Gamma \to \mathbb{R}^n$ requires $\Omega(n^{1-\sigma})$ communication bits. This also holds for a restricted case in which there is one deterministic bid on each item.*

To prove the theorem, we use the Bernstein's bound.

Lemma 3.1 (Bernstein's bound)**.** *Let $x_1, x_2, \ldots, x_n$ be a sequence of independent random variables such that for all $1 \leq i \leq n$ we have $0 \leq x_i \leq C$. If we define $X = \sum_{i=1}^n x_i$, we have $\Pr\Big(|X - E[X]| \geq \lambda\Big) \leq 2\exp\big(-\frac{\lambda^2/2}{\mathsf{Var}(X)+Ct/3}\big)$.*

We can now give the proof of the theorem.

Proof of Theorem 3.1. In order to prove this theorem we provide a set $\mathcal{B}$ of vectors of bids $\vec{b} = \langle b_1, b_2, \ldots, b_n\rangle$ such that for any arbitrary vector of reserve prices $\vec{r} = \langle r_1, r_2, \ldots, r_n\rangle$ we have the following two properties.

1. For at least half the vectors $\vec{b} = \langle b_1, b_2, \ldots, b_n\rangle \in \mathcal{B}$, we have $\sum_{i=1}^n b_i \geq \frac{1}{2}\sigma n \log n$.
2. For $1 - 2\exp\big(-\frac{3}{8}n^{1-\sigma}\big)$ fraction of the vector of bids $\vec{b}$ in $\mathcal{B}$, the profit of assigning reserve prices $\vec{r}$ is less than $2n$.

To show the second property, equivalently, we prove that if one picks a vector of bids $\vec{b}$ from $\mathcal{B}$ uniformly at random, the profit of reserve prices $\vec{r}$ on $\vec{b}$ is less than $2n$ with probability at least $1 - 2\exp\big(-\frac{3}{8}n^{1-\sigma}\big)$.

Let $\Psi : \Gamma \to \mathbb{R}^n$ be a $\frac{1}{4}\sigma\log n$-approximation pricing system. This means that for each vector of bids $\vec{b} \in \mathcal{B}$, there exists a key $\gamma \in \Gamma$ such that the profit of $\Psi(\gamma)$ on $\vec{b}$ is at least $2n$. Notice that given an arbitrary $\gamma \in \Gamma$, for at most $2\exp\big(-\frac{3}{8}n^{1-\sigma}\big)$ fraction of the bids $\vec{b} \in \mathcal{B}$, the reserve prices $\Psi(\gamma)$ has profit more than or equal to $2n$. Therefore, we have

$$|\Gamma| \geq \frac{1/2}{2\exp\big(-\frac{3}{8}n^{1-\sigma}\big)} = 0.25\exp\big(\frac{3}{8}n^{1-\sigma}\big).$$

This directly implies that Ψ requires at least $\log_2\Big(0.25\exp\big(\frac{3}{8}n^{1-\sigma}\big)\Big) \in \Omega(n^{1-\sigma})$ communication bits, as desired. Now, we just need to provide the set $\mathcal{B}$ with the promised properties.

We define $\mathcal{B} = \{\vec{b} = \langle b_1, b_2, \dots, b_n\rangle | \forall_i b_i \in \{n^\sigma, n^\sigma/2, n^\sigma/3, \dots, 1\}\}$. Remark that selecting one vector $\vec{b} = \langle b_1, b_2, \dots, b_n\rangle$ from $\mathcal{B}$ uniformly at random is equivalent to selecting each element b_i of $\vec{b}$ uniformly at random and independently from the set $\{n^\sigma, n^\sigma/2, n^\sigma/3, \dots, 1\}$.

We start with proving the first property of the set $\mathcal{B}$. To prove this property we show that if one draws a vector $\vec{b} = \langle b_1, b_2, \dots, b_n\rangle$ from $\mathcal{B}$ uniformly at random, with probability at least 0.5 we have $\sum_{i=1}^n b_i \geq \frac{1}{2}\sigma n \log n$. Let's define $B = \sum_{i=1}^n b_i$. Recall that b_i's are independent random variables drawn form $\{n^\sigma, n^\sigma/2, n^\sigma/3, \dots, 1\}$ uniformly at random. Thus, for any $1 \leq i \leq n$ we have $E[b_i] = \frac{\sum_{j=1}^{n^\sigma} \frac{n^\sigma}{j}}{n^\sigma} = \sum_{j=1}^{n^\sigma} \frac{1}{j} \log n^\sigma = \sigma \log n$. Thus, we have $E[B] \geq \sigma n \log n$. Also, for any $1 \leq i \leq n$ we upper bound the variance of b_i as follow.

$$\begin{aligned}\mathsf{Var}(b_i) &= E[b_i^2] - (E[b_i])^2 \leq E[b_i^2] \\ &= \frac{\sum_{j=1}^{n^\sigma} \left(\frac{n^\sigma}{j}\right)^2}{n^\sigma} = n^\sigma \sum_{j=1}^{n^\sigma} \frac{1}{j^2} \leq 2n^\sigma,\end{aligned}$$

Thus we have $\mathsf{Var}(B) \leq 2nn^\sigma \leq n^\sigma(\sigma n \log n)$, when $n \geq \exp(1/2\sigma)$. Now, by applying Bernstein's bound we have

$$\begin{aligned}&\Pr\Big(E[B] - B \geq \frac{1}{2}\sigma n \log n\Big) \leq \\ &\Pr\Big(\sigma n \log n - B \geq \frac{1}{2}\sigma n \log n\Big) \leq \\ &2\exp\big(-\frac{(0.5\sigma n \log n)^2/2}{n^\sigma(\sigma n \log n) + n^\sigma(0.5\sigma n \log n)/3}\big) = \\ &2\exp\big(-\frac{(1/4)\sigma n \log n}{(2+1/3)n^\sigma}\big) = \\ &2\exp\big(-\frac{3}{28}\sigma n^{1-\sigma} \log n\big) \leq 1/2,\end{aligned}$$

where the last inequality holds for large enough n. This means that $\Pr(B \geq \frac{1}{2}\sigma n \log n) \geq 1/2$.

Next we prove the second property of the set $\mathcal{B}$. Let $\vec{r} = \langle r_1, r_2, \dots, r_n\rangle$ be an arbitrary vector of reserve prices. For all $1 \leq i \leq n$, let $x_i = r_i$ if $r_i \leq b_i$ and let $x_i = 0$ otherwise. Note that the profit of assigning reserve prices $\vec{r} = \langle r_1, r_2, \dots, r_n\rangle$ on $\vec{b}$ is $\sum_{i=1}^n x_i$. Let's define $X = \sum_{i=1}^n x_i$. Below, for any arbitrary $0 \leq i \leq n$ we bound $E[x_i]$. Here we pick j such that $\frac{n^\sigma}{j+1} \leq r_i \leq \frac{n^\sigma}{j}$. Note that we have $E[x_i] = r_i Pr(r_i \leq b_i) \leq \frac{n^\sigma}{j}\frac{j}{n^\sigma} = 1$. Therefore, we have $E[X] = n$. Also, for any $1 \leq i \leq n$ we upper bound the variance of x_i as follow.

$$\begin{aligned}\mathsf{Var}(x_i) &= E[x_i^2] - (E[x_i])^2 \leq E[x_i^2] = \frac{j}{n^\sigma} r_i^2 \\ &\leq \frac{j}{n^\sigma}\Big(\frac{n^\sigma}{j}\Big)^2 = \frac{n^\sigma}{j} \leq n^\sigma.\end{aligned}$$

Thus, we have $\mathsf{Var}(X) \leq nn^\sigma = n^{1+\sigma}$. Now, by applying Bernstein's bound we have

$$\begin{aligned}\Pr\Big(X - E[X] \geq n\Big) &\leq 2\exp\big(-\frac{n^2/2}{n^{1+\sigma} + n^\sigma n/3}\big) \\ &= 2\exp\big(-\frac{3}{8}n^{1-\sigma}\big).\end{aligned}$$

Therefore we have $\Pr\big(X \geq 2n\big) \leq 2\exp\big(-\frac{3}{8}n^{1-\sigma}\big)$, which proves the second property of the set $\mathcal{B}$ and completes the proof. □

Setting $\sigma = \epsilon/2$ in Theorem 3.1 gives us the following corollary.

Corollary 3.1.1. *For any small constant ϵ there is no $o(\log(n))$-approximation pricing system using $O(n^{1-\epsilon})$ communication bits.*

4 Hardness Result

We show that MRPOP is NP-hard even when the space of the auction is 2-dimensional. Indeed this hardness result carries over to the general problem. We obtain this result via a reduction from MaxCut. In the MaxCut problem, we are given an undirected graph G, whose vertices we wish to partition into two disjoint sets to maximize the number of crossing edges. It has been shown that the MaxCut problem is NP-hard and cannot be even approximated within a factor better than 0.878 in polynomial time unless a widely believed conjecture fails (Christofides 1975; Khot and Vishnoi 2005). We provide the proof of the following theorem in the full version.

Theorem 4.1. *The MRPOP problem is NP-hard.*

5 Polynomial Time Approximation

In this section, we propose a structure for a set of desirable solutions for the MRPOP problem, namely *stable solutions*. Recall that a multiplicative reserve price solution is specified by d vectors $\langle v^1, v^2, \dots, v^d\rangle$, where every v^i is itself a vector of multiplicative reserves of size $|A_i|$. The reserve price of every type vector $(\alpha_1, \alpha_2, \dots, \alpha_d)$ is then determined as

$$\prod_{1 \leq i \leq d} v^i_{\alpha_i}.$$

We defined a solution $v = \langle v^1, v^2, \dots, v^d\rangle$ to be stable if every vector v^i is optimal given that the other vectors are fixed. More precisely, v is stable if and only if for any $i \in [d]$ and any u^i such that $|u^i| = |v^i|$ we have

$$\mathsf{Rev}(v) \geq \mathsf{Rev}(v^{-i}, u^i).$$

In addition to this, we assume stable solutions are always non-zero.

Stable solutions have several desirable properties. For instance, one can show that the revenue of a polynomially bounded stable solution is at least $1/\log(R)$ of that of the optimal solution.

Theorem 5.1. *Let $v = \langle v^1, v^2, \dots, v^d\rangle$ be a stable solution for an instance of the problem wherein the reserve prices are $O(R)$. Then, $\mathsf{Rev}(v) \geq \mathsf{Opt}/O(\log(R))$ where Opt is the maximum possible revenue any solution can achieve for this instance of the problem.*

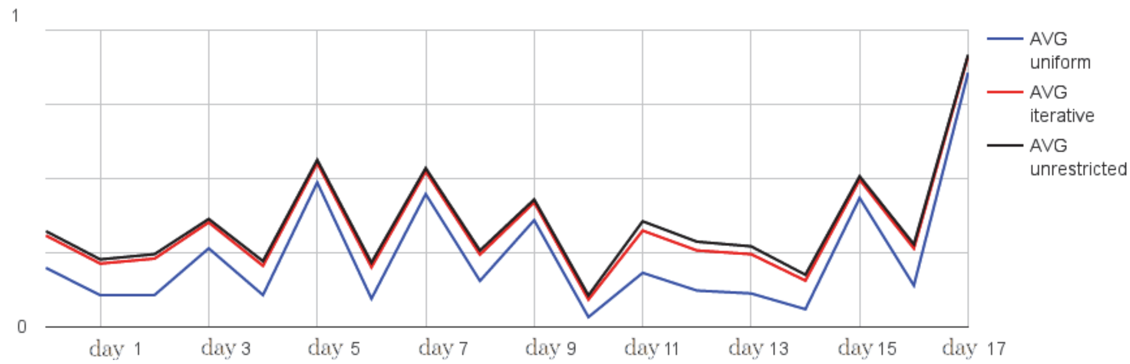

Figure 1: The horizontal line denotes the date and the vertical line illustrates the average accumulated bid per impression for every algorithm. Due to privacy issues, we scaled down the revenues to numbers between 0 and 1. The black and red segments are very close and they sometimes overlap.

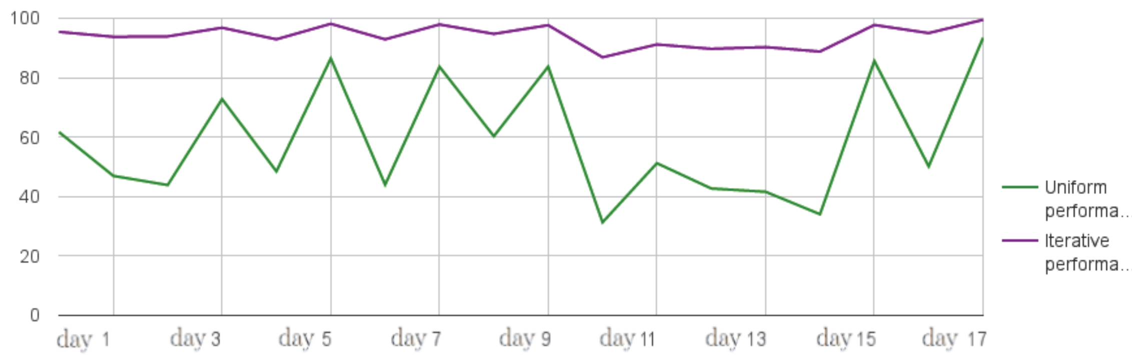

Figure 2: The horizontal line denotes the date and the vertical line illustrates the ratio of the revenue achieved by an algorithm that that of the optimal unrestricted algorithm in percentage.

In what follows, we give an algorithm that finds an $O(\log R)$ approximation solution for the MRPOP problem with dimension d, in polynomial time.

At first, we set the multiplicative solutions for all reserve price vectors equal to 1. Since all bids are integer numbers bounded by R, this guarantees a $1/R$ fraction of the optimal revenue. Then we iteratively improve the revenue by updating the solution vectors. More precisely, in every step, for each $1 \leq i \leq d$, we find a vector u^i that maximizes $\text{Rev}(v^{-i}, u^i)$. Notice that this can be computed in polynomial time by iterating over all possible choices of each index. Note that we only consider multiplicative reserves between $1/R$ and R, so that our solution always remains polynomially bounded. Next we calculate $\text{Rev}(v^{-i}, u^i) - \text{Rev}(v)$ for each index i and find the vector that maximizes this expression. Then we update our solution by replacing the vector in our solution. Note that since $\text{Rev}(v^{-i}, u^i) \geq \text{Rev}(v)$, our solution improves in every step. We stop when the improvement in revenue is no more than an arbitrary small threshold.

Theorem 5.2. *There exists an algorithm that finds a $O(1/\log R)$ approximation solution to MRPOP in polynomial time.*

6 Empirical Study

In this section, we present an empirical evaluation of the algorithm we present in Section 5. We compare the performance of our algorithm with that of the *uniform reserve pricing* and *unrestricted* algorithms described as follows.

uniform reserve pricing: we find the reserve price p that maximizes $\sum \pi(D, p)$ for all auctions and set the multiplicative reserves in a way that the reserve price of every auction is equal to p.

unrestricted: We set reserve prices for each auction separately. As declared, the revenue achieved in this case is an upper bound on the optimal revenue we can get in the restricted setting.

Our experiments are based on millions of bids submitted by advertisers to Google Advertising Exchange, AdX. More precisely, once a user visits a webpage with an ad supported by the exchange, this opportunity is reported to the advertisers as a set of features. Every feature reveals a property of the ad to the advertiser; the advertiser submits a bid based on these. Therefore, in our dataset, for every combination of the (advertiser, values of features), we have a submitted bid. To run the experiments, we select two features that have a meaningful correlation with bids and create a table where each cell denotes a pair of values for the features. Note that one is related to contextual property of ad slot ID and one of them is related to the demographic of user watching the page. For every cell, we find the bid distribution corresponding to it and evaluate the algorithms performance on the table. Each row of the table corresponds to a value for feature 1 and each column represents a unique value for feature 2. Tables have millions of rows and columns. The experiments are run on the data of 17 consecutive days.

We would like to note that although the numbers are close for the last day, there is still a gap between revenues. Of course, the gap fluctuates based on several factors that have high impacts on the revenues of the algorithms. For instance, depending on the time that the query arrives, some query features might be reset, and hence are less informative in comparison to other dates. Another example is that on particular dates, both the number of queries and the number of demands for ad slots increase. Analogously, sometimes the advertisers are reluctant to bid, affecting the results. What is clear from the charts is that the iterative algorithm is clearly providing better revenues; nonetheless, the ratio of improvement is affected by several factors.

As shown in Figures 1 and 2, the iterative algorithm works almost optimally in practice. The experiments are run on 18 different data sets, each containing millions of distributions in average. According to our experiments, the iterative algorithm achieves more than 94% (in average) of the total revenue achievable in the unrestricted version of the problem. Note that, as multiplicative pricing is inherently weaker than and unrestricted algorithm, revenue loss is inevitable. However, our experiments show that the revenue loss in practice is substantially better than the theoretical bounds.

Uniform reserve pricing algorithm exhibits a poor performance in the experiments. The average revenue loss of this algorithm is more than 40% over the data sets. Although this is better than the theoretical guarantees, this algorithm is not competitive to the iterative algorithm.

References

Archak, N.; Mirrokni, V.; and Muthukrishnan, S. 2010. Budget optimization for online advertising campaigns with carryover effects. In *Sixth Ad Auctions Workshop*. Citeseer.

Bateni, M.; Feldman, J.; Mirrokni, V. S.; and Wong, S. C. 2014. Multiplicative bidding in online advertising. In *ACM Conference on Economics and Computation, EC '14, Stanford , CA, USA, June 8-12, 2014*, 715–732.

2014. Bing ads bid adjustments. https://help.bingads.microsoft.com/apex/index/3/en-au/51004.

Borgs, C.; Chayes, J.; Immorlica, N.; Jain, K.; Etesami, O.; and Mahdian, M. 2007. Dynamics of bid optimization in online advertisement auctions. In *Proceedings of the 16th international conference on World Wide Web*, 531–540. ACM.

Briest, P.; Chawla, S.; Kleinberg, R.; and Weinberg, S. M. 2010. Pricing randomized allocations. In *Proceedings of the twenty-first annual ACM-SIAM symposium on Discrete Algorithms*, 585–597. Society for Industrial and Applied Mathematics.

Charles, D.; Chakrabarty, D.; Chickering, M.; Devanur, N. R.; and Wang, L. 2013. Budget smoothing for internet ad auctions: a game theoretic approach. In *Proceedings of the fourteenth ACM conference on Electronic commerce*, 163–180. ACM.

Chawla, S.; Malec, D. L.; and Sivan, B. 2010. The power of randomness in bayesian optimal mechanism design. In *Proceedings of the 11th ACM conference on Electronic commerce*, 149–158. ACM.

Christofides, N. 1975. *Graph Theory, An algorithmic approach*. New York: Academic Press Inc.

Devanur, N. R., and Hayes, T. P. 2009. The adwords problem: online keyword matching with budgeted bidders under random permutations. In *Proceedings of the 10th ACM conference on Electronic commerce*, 71–78. ACM.

Feldman, J.; Muthukrishnan, S.; Pal, M.; and Stein, C. 2007a. Budget optimization in search-based advertising auctions. In *Proceedings of the 8th ACM conference on Electronic commerce*, 40–49. ACM.

Feldman, J.; Muthukrishnan, S.; Pál, M.; and Stein, C. 2007b. Budget optimization in search-based advertising auctions. In *EC*, 40–49. ACM.

Goel, G.; Mirrokni, V.; and Leme, R. P. 2015. Polyhedral clinching auctions and the adwords polytope. *Journal of the ACM (JACM)* 62(3):18.

2014. Setting bid adjustments. http://support.google.com/adwords/answer/2732132.

Hart, S., and Nisan, N. 2013. The menu-size complexity of auctions. In *ACM Conference on Electronic Commerce, EC '13, Philadelphia, PA, USA, June 16-20, 2013*, 565–566.

2013. The most important changes to google adwords in 2013. http://blog.hubspot.com/marketing/google-adwords-changes-2013-list.

Karande, C.; Mehta, A.; and Srikant, R. 2013. Optimizing budget constrained spend in search advertising. In *Proceedings of the sixth ACM international conference on Web search and data mining*, 697–706. ACM.

Khot, S., and Vishnoi, N. K. 2005. On the unique games conjecture. In *Annual Symposium on Foundations of Computer Science*, volume 46, 3. IEEE COMPUTER SOCIETY PRESS.

Mehta, A.; Saberi, A.; Vazirani, U.; and Vazirani, V. 2007. Adwords and generalized online matching. *Journal of the ACM (JACM)* 54(5):22.

Muthukrishnan, S.; Pál, M.; and Svitkina, Z. 2007a. Stochastic models for budget optimization in search-based advertising. In *WINE*.

Muthukrishnan, S.; Pál, M.; and Svitkina, Z. 2007b. Stochastic models for budget optimization in search-based advertising. In *International Workshop on Web and Internet Economics*, 131–142. Springer.