Selective Call Out and Real Time Bidding

Tanmoy Chakraborty[1,⋆], Eyal Even-Dar[4,⋆⋆], Sudipto Guha[1],
Yishay Mansour[2], and S. Muthukrishnan[3]

[1] University of Pennsylvania, Philadelphia, PA
{tanmoy,sudipto}@cis.upenn.edu
[2] Google Israel and The Blavatnik School of Computer Science,
Tel-Aviv University, Tel-Aviv, Israel
mansour.yishay@gmail.com
[3] Google Research, 76 Ninth Ave, New York, NY
muthu@google.com
[4] evendar.eyal@gmail.com

Abstract. Ads on the Internet are increasingly sold via ad exchanges such as RightMedia, AdECN and Doubleclick Ad Exchange. These exchanges allow real-time bidding, that is, each time the publisher contacts the exchange, the exchange "calls out" to solicit bids from ad networks. This solicitation introduces a novel aspect, in contrast to existing literature. This suggests developing a joint optimization framework which optimizes over the allocation and well as solicitation.

We model this selective call out as an online recurrent Bayesian decision framework with bandwidth type constraints. We obtain natural algorithms with bounded performance guarantees for several natural optimization criteria. We show that these results hold under different call out constraint models, and different arrival processes. Interestingly, the paper shows that under MHR assumptions, the expected revenue of generalized second price auction with reserve is constant factor of the expected welfare. Also the analysis herein allow us prove adaptivity gap type results for the adwords problem.

1 Introduction

A dominant form of advertising on the Internet involves *display* ads; these are images, videos and other ad forms that are shown on a web page when viewers navigate to it. Each such showing is called an *impression*. Increasingly, display ads are being sold through exchanges such as RightMedia, AdECN and DoubleClick Ad Exchange. On the arrival of an impression, the exchange solicits bids and runs an auction on that particular impression. This allows *real time bidding* where ad networks can determine their bids for each impression individually in real time (for an example, see [24]), and more importantly where the creative (advertisement) can be potentially produced on-the-fly to achieve

⋆ Part of this work was done while the authors were visiting Google Research.
⋆⋆ Part of this work was done while the author was at Google Research.

A. Saberi (Ed.): WINE 2010, LNCS 6484, pp. 145–157, 2010.

better targeting [22]. This potential targeting comes hand in hand with several challenges. The Exchange and the networks face a mismatch in infrastructure and capacities and objectives. From an infrastructure standpoint, the volume of impressions that come to the exchange is very large comparison to a smaller ad network limited in servers, bandwidths, geographic location preferences. This implies a bound on the number of auctions the network can participate in effectively. A network would prefer to be solicited only on impressions which are of interest to it, and in practice use a descriptive languages to specify features of impressions (say, only impressions from NY). However this is an offline feature and runs counter to the attractiveness of real time bidding. Therefore the exchange has to "call out" to the networks selectively, simultaneously trying to balance the objective of soliciting as many networks as possible and increasing total value, as well as not creating congestion. This leads to a host of interesting questions in *developing a joint optimization framework that optimizes over the allocation objective as well as the decisions to solicit the bids.*

1.1 Selective Call Out: The Model

Let n be the number of ad networks $1, 2 \ldots n$. We assume that impressions arrive from a fixed (unknown) distribution over a finite set U_I, and that there exists a finite set of bid values U_B, where $L = \max\{u \mid u \in U_B\}$. In the following, ad networks will be indexed by $i \in \{1, 2 \ldots n\}$, impressions by $j \in U_I$ and bid values by $k \in U_B$. The problem setting involves several steps:

1. An impression (or keyword) j, assumed drawn from a distribution $\mathcal{D}$, comes to the exchange. There may be *multiple slots* associated with a single impression, corresponds to text ads being blocked together, different locations in the page, which are often characterized by different *discount rate*. Let there be M slots, with discount rates $1 \geq \varrho_1 \geq \varrho_2 \ldots \geq \varrho_M \geq 0$. If a bidder bids v, then it is assumed that the bid for the ℓ^{th} slot is $v\varrho_\ell$. The case of $M = 1$ is common and correspond to a basic pay-per-impression mechanism with discount rate 1.
2. Given an impression j, the bid of ad network i for impression j is drawn from a fixed distribution $\mathbf{V}_{ij}$ such that the bid is v with probability p_{ijv}. Note that the bids of different networks are likely to be correlated based on the perceived value of the impression, however, conditioned on j, the specific dynamics of different bidders can be construed to be independent. We assume that the exchange has learned or can predict these p_{ijv} given the impression j. This is an assumption similar to estimation processes used by search engines to predict click through rates.
3. The exchange decides on the subset S_j of networks to call out, subject to the *Call-Out constraints*, which roughly bounds the rate at which the exchange can send impressions to solicit bid from an ad network. This decision is executed before seeing the next impression.

To define a specific problem in the above framework, we need to specify (i) an objective function (ii) a model for call out constraints, and (iii) the comparison class. The **goal** is to design a call-out policy that satisfies (ii) and is near optimal in the objective function (i), when compared to other algorithms in class (iii). We discuss the instantiations of (i),(ii), and (iii) in the following.

(i) Objective Functions. We consider three different objective functions. (a) *Total Value*: The sum of the maximum bids in S_j over the arriving impressions j. (b) *GSP-Reserve* (defined above) and (c) *Revenue under posted price mechanism (take-it-or-leave-it prices).* All the quantities are in expectation. Total value corresponds to the welfare. The GSP with an uniform reserve price is a common mechanism used in these settings. Posted Prices (different networks may get different prices) are also used in this context. Note, unlike [8,4], the mechanism is *parallel* posted price because all the prices are posted first.

(ii) Call Out Constraint Models. The simplest model for the call out constraints is a model where one impression arrives at the exchange at each time step and if the total number of arrivals is m, ad network i can be solicited at most times for some known $\rho_i \leq 1$. We will refer to this as *time average model under uniform arrival* – which describes the constraints at the outgoing and incoming sides of the exchange respectively. Most of the paper will focus on this model – *primarily because we can show that other common models reduce to this variant.* The non-initiated reader can skip the description of these models and proceed to (iii). On the incoming side of the exchange, standard practice is to assume bursty (Poisson) arrivals. We consider this generalization. On the outgoing side, the simple model allow the possibility that the call-outs to a network are made on contiguous subset of impressions. This misses the original goal that the ad network would receive the impressions at a "smooth" rate. A common model used for behavior is the *token bucket model* [26]. A token bucket has two parameters, bucket size σ and token generation rate ρ. The tokens represent sending rights, and the bucket size is the maximum number of tokes we can store. The tokens are generated at a rate of ρ per unit time, but the number of tokens never exceeds σ. In order to send, one needs to use a token, and if there are no tokens, one can not send. The output stream of a (σ, ρ) token bucket can be handled by a buffer of size σ and a time average rate of ρ – the buffer is initially full. Unlimited buffer size corresponds to the time average model.

(iii) Comparison Class. Given the call out constraints, we define the class of admissible policies. An *admissible call out policy* specifies (possibly with randomization), for each arriving impression j, the subset S_j of ad networks to call out, while satisfying all call out constraints over the entire sequence of impressions. The policy bases its decision on the prior information about the bid distributions, and has no knowledge of the actual bid values. In the case of GSP-Reserve mechanism, the call out policy also decides the reserve price for each impression (and likewise for the case of a posted price mechanism). Our comparison class is the set of all admissible call out policies which know the bid distributions for every impression, but do not know the actual realization of the bids. The performance of a policy is measured as the expected (over the bid distributions and the impression arrivals) objective value obtained per arriving impression, when impressions are drawn from $\mathcal{D}$. *A policy is α-approximate if it achieves at least α times the performance of the optimal policy for the corresponding objective.*

1.2 Our Results, Roadmap and Related Work

We provide three algorithms LP-Val, LP-GSP, and LP-Post for the objectives discussed for the time average uniform arrival model. We then prove that the results translate naturally to other constraint models. Recall, L is the largest possible bid. The algorithms will have a natural two-phase approach where we use the t initial impressions as a sample as exploration and subsequently use/exploit this algorithm (see Section 2 for more discussion). We show that:

Theorem 1. *Suppose the optimal policy has expected total value at least $\delta > 0$. For any $\epsilon > 0$,* LP-Val *with a sample of $t = \tilde{O}(\frac{n^2 L}{\delta\epsilon})$ impressions gives a $(1 - \frac{1}{e} - \epsilon)$-approximate policy.*

Theorem 2. *Suppose the optimal GSP-Reserve policy has expected revenue at least $\delta > 0$. For any $\epsilon > 0$,* LP-GSP *with a sample of $t = \tilde{O}(\frac{n^2 L}{\delta\epsilon})$ impressions gives a $O(1)$-approximate policy, if all bid distributions satisfy the* monotone hazard rate (MHR) *property. Moreover, the call outs of the policy derived from* LP-GSP *are identical to those of the policy derived from* LP-Val.

Theorem 3. *Suppose the optimal posted price policy has expected revenue at least $\delta > 0$. For any $\epsilon > 0$,* LP-Post *with a sample of $t = \tilde{O}(\frac{n^2 L}{\delta\epsilon})$ impressions gives a $(1 - \frac{1}{e} - \epsilon)$-approximate policy.*

In particular, we show that when every bidder is solicited – GSP-Reserve achieves a revenue that is $O(1)$ factor of optimal welfare, when all bid distributions satisfy the MHR property. This is a common distributional assumption in economic theory, and is satisfied by many distributions [3]. This result is in the same spirit as (but immediately incomparable to) the result in [4], which relates the optimum sequential posted price revenue to the optimal welfare under the same assumptions. We are unaware of such results about GSP-Reserve. We do **not** need the MHR assumption unlike the result in [4] for sequential posted prices, since the comparison classes are different (the prices are posted in parallel). We discuss the realization of the above algorithms for different network models next.

Theorem 4. *For a given distribution on impressions $\mathcal{D}$, suppose we have an α-approximate policy for an objective which is additive given the allocations and the realizations (and is at least $\delta > 0$) in the time average uniform arrival call out model. Let $\sigma_i > \sigma\ \forall i$. Then we can convert the policy to a $(\alpha - \frac{1}{\sigma - 1})$-approximate policy in the token bucket model. The result extends to Poisson arrivals.*

Comparison with the Adwords Model: The call out optimization framework is similar to the online ad allocation framework for search ads, or the Adwords problem [19,5,7], and its stochastic variants [9,27]. However there are significant differences which we discuss below. The Adwords problem is posed in the deterministic setting where the expected revenue is treated as a known deterministic reward of allocating an impression j to an advertiser i. The call out framework has no deterministic analogue; the rationale of the exchange is that the bids and the participation are not known. In this regard, the call-out framework is similar to the Bayesian mechanism design [21].

This unknown participation has broad conceptual implications, first of which is the notion of "adaptivity gap". An ad-allocation policy may choose to react to realization of the random variables. This aspect is central in the exchange setting. Interestingly, the analysis in this paper also allows us to consider adaptivity gap for in the Adwords setting, if the allocation policy is allowed to adapt to the occurence or absence of a click (instead of using a deterministic quantity which is the product of the click-through-rate and the cost of a click). This is a significant issue for low click through rates and large bids, such that a payout affects the budget substantially. To the best of our knowledge, no analysis of adaptivity gap existed for the adwords problem previously,

Second, many objective functions such as generalized second price with reserve (henceforth GSP-Reserve), for one or multiple slots, have very different behaviors in the Bayesian and deterministic settings. Consider running Myerson's (or similar) mechanism on the expected bids instead of the distributions. For *known* deterministic bids, reserve prices can be made equal to the bid, and are not useful. Strong lower bounds hold for GSP without reserves [1]. Note that in GSP-Reserve we announce an uniform reserve price, *before* the bids are solicited.

Third, the notion of a comparison class in case of call out optimization framework requires more care than the Adwords framework. In the setting of these large exchanges, a comparison class with full foreknowledge of all information (in particular, the realization of the bids) is unrealistic. Moreover, the realizations of the bids depends on the networks which are called out, and two different strategies that call out to two different subsets will have completely different information. Thus to compare two algorithms, we should compare their expected outcome – *but each algorithm is allowed to be adaptive.*

We use Lagrangian decoupling techniques for separable convex optimization pioneered by Rockafellar [25]. This has been used in the stochastic variants of the Adwords problem in [9,27]. But the similarity ends there. *The different possible objectives of the call-out framework are not convex.* In fact, in the case of optimizing revenue in GSP (with reserve) or in posted price mechanisms, the objective is not submodular for all prices (as in welfare maximization). Submodular maximization with linear constraints has been studied, and while good approximation algorithms exist [6,17], they are inherently offline – the key aspect of call out optimization is that the decision has to be made in an online fashion. The same is true for sequential posted price mechanisms analyzed in [8,4] (albeit with more general matroid setting), the posted prices in the call out setting need to be announced in parallel (and the eventual allocation is sequential).

Other Related Work: A combination of stochastic and online components appear in many different settings [14,15,2,13] which are not immediately relevant to the call-out problem. We note that the bandwidth-like constraints (where the constraint is on a parameter different than the obtained value, as is the case for call-outs) has not been studied in the bandit setting (see [16,23]). Finally, bidding and inventory optimization problems [11,10,20], are not immediately relevant.

Roadmap: We summarize the results on online stochastic convex optimization in Section 2. We subsequently discuss the the total value problem in Section 3. We discuss the GSP-Reserve problem in Section 4. The posted price problem is discussed in the full version. The token bucket model and other arrival assumptions are also discussed in the full version.

2 Preliminaries

Consider a maximizing a "separable" linear program (LP) $\mathcal{L}$ defined on Q *global* constraints with right hand side b_i, such that the Lagrangian relaxation produced by the transferring these constraints to the objective function decouples into a collection of independent non-negative smaller LPs $\mathcal{L}_j$ over n' variables and *local* constraints. This implies that the objective function of $\mathcal{L}$ is a weighted linear combination of $\mathcal{L}_j$. The uniqueness of the optimum solutions for $\mathcal{L}_j$ implies that $\mathcal{L}$ reduces to finding the Lagrangian multipliers. The unique solution is achieved by adding "small perturbations", see Rockafellar [25]. However, this approach only provides a certificate of optimality and a solution, once we are given the Lagrangians. The approach does not give us an algorithm to find the Lagrangian multipliers themselves. Devanur and Hayes [9] showed that if the smaller LPs could be sampled with the same probability as their contribution to the objective of $\mathcal{L}$, and the derivatives of the Lagrangian can be bounded, then the Lagrangians derived from a small number of samples (suitably scaled) can be used to solve the overall LP. The weighted sampling reduces to the prefix of the input if the $\mathcal{L}_j$s arrive in random order (see [12]). This was extended to convex programs in [27]. The number of sample bound requires several (easy) Lipschitz type properties:

1. The optimum value of $\mathcal{L}$ is at least $\delta > 0$ and the optimum solution of $\mathcal{L}_j$ is at most R.
2. For each setting of the Lagrangians, every $\mathcal{L}_j$ has a unique optimum solution.
3. Reducing b_i by a factor of $1-\epsilon$ reduces $\mathcal{L}$ by at most $(1-\epsilon)$.
4. $\mathcal{L}$ does not change by more than a constant times $1+\epsilon$ if we alter the value of the optimum Lagrangian multipliers by a factor of $1+\epsilon$.

Theorem 5. *Sample $t = \tilde{O}(\frac{n'QR}{\delta\epsilon})$ of the smaller linear programs and consider the linear program $\mathcal{L}'$ which corresponds to the union of these smaller linear programs and suitably scaled global constraints. If we use the optimum Lagrangian multipliers corresponding to the global constraints of $\mathcal{L}'$ to solve the decoupled instances of $\mathcal{L}_j$ as they are available (in an online fashion) then we produce a $1+\epsilon$ approximation to the optimum solution of $\mathcal{L}$.*

In the setting of Adwords, the smaller LPs correspond to the arrival of an impression, and the associated assignment. Thus the stochastic framework of the Adwords problem is obviously of relevance to the call out optimization framework. However the focus shifts on solving the smaller LPs which encode the call out decision. A nice outcome of the approach is a simple two phase algorithm;

an *Exploration phase* where the samples are drawn, and an *Exploitation phase* where the Lagrangian multipliers are used. If Q, n', R are small, then the exploration phase can be (relatively) short and this yields a natural algorithm. Thus the goal of the rest of the paper would be to formulate separable convex relaxations and achieve the mentioned properties.

3 The Total Value Problem

In this section, we prove Theorem 1, and describe LP-Val. Let q_j denote the probability that impression j arrives. We shall add infinitesimal random perturbations to p_{ijv} which shall not affect the performance of any policy but ensure p_{ijv} are in general positions, that is, any combination of them will almost surely create a non-singular matrix.

The LP Relaxation: Let x_{ij} be the (conditional) probability that advertiser i was called out on impression j. Let $y_{ijv\ell}$ be the probability that advertiser i bid the value t and was assigned the slot ℓ (also conditioned on j). The constraints are named as $A(x, y)$ and $B(x, y)$ as shown.

$$LP1 = \max \sum_j q_j \sum_v \sum_i \sum_\ell v\varrho_\ell y_{ijv\ell} \quad \text{s.t.} \quad \left.\begin{array}{l} \sum_j q_j x_{ij} \quad \le \rho_i \\ \left.\begin{array}{l} \sum_i \sum_v y_{ijv\ell} \le 1 \\ x_{ij} \qquad\quad \le 1 \\ \sum_\ell y_{ijv\ell} \le p_{ijv} x_{ij} \\ x_{ij}, y_{ijv\ell} \ge 0 \end{array}\right\} B(x,y) \end{array}\right\} A(x,y)$$

Decoupling: Let λ_i^* be the optimum Lagrangian variable for the constraint $\sum_j q_j x_{ij} \le \rho_i$. $LP1$ then decouples to smaller LPs, $LP2(j, \boldsymbol{\lambda_i^*})$ subject to the constraints $A(x, y)$, that is, $LP1 = LP1(\boldsymbol{\lambda_i^*}) = \sum_i \lambda_i^* \rho_i + \sum_j q_j LP2(j, \boldsymbol{\lambda_i^*})$ where $LP2(j, \boldsymbol{\lambda_i^*}) = \max\left(\sum_v \sum_i \sum_\ell v\varrho_\ell y_{ijv\ell} - \sum_i \lambda_i^* x_{ij}\right)$

Solving $\mathbf{LP2(j, \boldsymbol{\lambda_i^*})}$. We begin by considering the dual. Let $\tau_{j\ell}$ be the dual of the constraint $\sum_i \sum_v y_{ijv\ell} \le 1$. Let ξ_{ijv} correspond to the dual of the constraint $\sum_\ell y_{ijv\ell} \le p_{ijv} x_{ij}$. Let ζ_{ij} correspond to the dual of $x_{ij} \le 1$.

$$DualLP2(j, \boldsymbol{\lambda_i^*}) = \min \sum_i \zeta_{ij} + \sum_\ell \tau_{j\ell} \quad \text{s.t.} \quad \begin{array}{ll} \tau_{j\ell} + \xi_{ijv} & \ge v\varrho_\ell \\ \zeta_{ij} - \sum_v \xi_{ijv} p_{ijv} & \ge -\lambda_i^* \\ \tau_{j\ell}, \xi_{ijv}, \zeta_{ij} & \ge 0 \end{array}$$

Lemma 1. *Let $\tau_{j\ell}^*$ be the optimum dual variables for $LP2$. Then (i) For all ℓ there exists i and $v \ge \tau_{j\ell}^*/\varrho_\ell$ s.t. $\xi_{ijv}^* = v\varrho_\ell - \tau_{j\ell}^*$ (ii) $\tau_{j\ell}^*/\varrho_\ell$ is non-increasing.*

Proof. For every ℓ there must be some i, v such that we have $\tau_{j\ell}^* + \xi_{ijv}^* = v\varrho_\ell$. Otherwise we can keep decreasing $\tau_{j\ell}^*$, keeping all other variables the same and contradict the optimality of the dual solution. Now $\xi_{ijv}^* \ge 0$ and the condition on t follows. The condition corresponds to the set of points (v, ξ_{ijv}^*) in the two dimensional (x, y) plane being above the lines $\{y = \varrho_\ell x - \tau_{j\ell}^*\}$. For the second part, consider $\tau_{j\ell}^*$ and the i, v such that we have $\tau_{j\ell}^* + \xi_{ijv}^* = v\varrho_\ell$. Define t to be the support of ℓ. Let $v \ge \tau_{j\ell}^*/\varrho_\ell$ be the largest such support of ℓ. Consider $\tau_{j(\ell-1)}^*$. We have $(\tau_{j(\ell-1)}^* + \xi_{ijv}^*)/\varrho_{\ell-1} \ge v = (\tau_{j\ell}^* + \xi_{ijv}^*)/\varrho_\ell$. But $\varrho_{\ell-1} \ge \varrho_\ell$ and thus $\tau_{j\ell}^*/\varrho_\ell$ are non-increasing in ℓ. Moreover, if $\varrho_\ell = \varrho_{\ell-1}$ then $\tau_{j\ell}^* = \tau_{j(\ell-1)}^*$.

Decoupling LP2$(\mathbf{j}, \boldsymbol{\lambda}_{\boldsymbol{i}}^*)$ itself. Consider $LP2(j, \boldsymbol{\lambda}_{\boldsymbol{i}}^*)$ with the Lagrangians $\tau_{j\ell}^*$. The problem decomposes under the constraints $B(x, y)$, to $LP2(j, \boldsymbol{\lambda}_{\boldsymbol{i}}^*) = \sum_\ell \tau_{j\ell}^* + \sum_i LP3(j, \boldsymbol{\lambda}_{\boldsymbol{i}}^*, \boldsymbol{\tau}_{\boldsymbol{j\ell}}^*, i)$ where $LP3(j, \boldsymbol{\lambda}_{\boldsymbol{i}}^*, \boldsymbol{\tau}_{\boldsymbol{j\ell}}^*, i) = \max\left(\sum_\ell \sum_v \left[v\varrho_\ell - \tau_{j\ell}^*\right] y_{ijv\ell} - \lambda_i^* x_{ij}\right)$.

Lemma 2. *Define $\ell(v) = \arg\max_{\ell'} \left\{\varrho_{\ell'} v - \tau_{j\ell'}^* | \varrho_{\ell'} v > \tau_{j\ell'}^*\right\}$ and $\ell(t) = M+1$ if the set is empty. Set $y_{ijv\ell}^* = p_{ijv}$ if $\ell = \ell(v)$ and 0 otherwise. If $\sum_v \sum_\ell v\varrho_\ell y_{ijv\ell}^* \geq \lambda_i^*$ we set $x_{ij} = 1$ and $y_{ijv\ell} = y_{ijv\ell}^*$. Otherwise we set $x_{ij} = y_{ijv\ell} = 0$.*

Proof. $LP3(j, \boldsymbol{\lambda}_{\boldsymbol{i}}^*, \boldsymbol{\tau}_{\boldsymbol{j\ell}}^*, i)$ is optimized at $x_{ij} = 1$ or $x_{ij} = 0$. This is because if $0 < x_{ij} < 1$ and $\sum_\ell \sum_v \left[v\varrho_\ell - \tau_{j\ell}^*\right] y_{ijv\ell} - \lambda_i^* x_{ij} > 0$ then we can multiply all the variables by $1/x_{ij}$ and have a better solution. If the latter condition is not true then $x_{ij} = 0$ is an equivalent solution. If $x_{ij} = 1$ the optimal setting for $y_{ijv\ell}$ is $y_{ijv\ell}^*$. (Note that $y_{ijv\ell}^*$ is uniquely determined for a fixed t.) Thus the overall optimization follows from comparing the $x_{ij} = 1$ and $x_{ij} = 0$ case.

Interpretation and the Call Out Algorithm: Given $\{\tau_{j\ell}^*\}_{\ell=1}^M$, the distribution $\{p_{ijv}\}$ for i, is divided into at most $M + 1$ pieces (some of the pieces can be a single point) given by the upper envelope (the constraint max) of the lines $\{\varrho_\ell x - \tau_{j\ell}^*\}_{\ell=1}^M$ and the line $y = 0$, in the x–y coordinate plane. Intuitively, seeing the value $x = t$, if the upper envelope corresponds to the equation $\varrho_\ell x - \tau_{j\ell}^*$ then we are "interested" in the slot ℓ. If the weighted (by ϱ_ℓ) sum of interests, given by $\sum_v v\varrho_\ell y_{ijv\ell}^*$ exceeds λ_i^*, then it is beneficial to call out i. We call out based on this condition and allocate the slots in decreasing order of bids.

Analysis: The LP2 solution satisfies: $LP1 = \sum_{i,v} v p_{ijv} \varrho_{\ell(v)}$ and $\sum_{i,v:\ell(v)=\ell} p_{ijv} = 1$. For each slot ℓ, let $w_i(\ell) = \sum_{v:\ell(v)=\ell} v p_{ijv}$ and $u_i(\ell) = \sum_{v:\ell(v)=\ell} p_{ijv}$. Order the i in non-increasing order of $w_i(\ell)/u_i(\ell)$ inside the slot. If we call out to i and get t, then for the sake of analysis we will consider its contribution to slot $\ell(t)$ only. Moreover, we stop the contribution to a slot ℓ if any any of the i return a value t with $\ell(t) = \ell$. The best M ordered bids outperform the analyzed contribution in every scenario. Therefore it suffices to bound the contribution.

Lemma 3. *Suppose we are given a set of independent variables Y_i such that $Pr[Y_i \neq 0] = u_i$ and $E[Y_i] = w_i$. Consider the random variable Y corresponding to the process which orders the variables $\{Y_i\}$ in non-increasing order of w_i/u_i, and stops as soon as the first non-zero value is seen. Then $E[Y] = \sum_i \prod_{i'<i}(1 - u_{i'})w_i \geq \sum_i w_i\left(1 - e^{-1}\right)$.*

Proof. Let $F(\{(w_i, u_i)\}) = \sum_i \prod_{i'<i}(1 - u_{i'})w_i$. Let $\lambda = \sum_i w_i / \sum_i u_i$. Given the sequence $\{(w_i, u_i)\}$ where w_i/u_i are non-increasing, if there exists an i such that $w_i/u_i \neq w_{i+1}/u_{i+1}$, then define a new sequence $\{(w_i', u_i)\}$ as follows:

$$w_{i'}' = \begin{cases} w_{i'} & \text{if } i' \neq i, i+1 \\ w_i - \Delta & \text{if } i' = i \\ w_{i+1} + \Delta & \text{if } i' = i+1 \end{cases} \qquad \text{where} \qquad \Delta = \frac{\frac{w_i}{u_i} - \frac{w_{i+1}}{u_{i+1}}}{\frac{1}{u_i} + \frac{1}{u_{i+1}}}$$

Note that $\sum_i w_i = \sum_i w'_i$ and w'_i/u_i remains non-increasing. Now $F(\{(w_i,u_i)\}) - F(\{(w'_i,u_i)\}) = \prod_{i'<i}(1-u_{i'})\Delta - \prod_{i'<i+1}(1-u_{i'})\Delta = \prod_{i'<i}(1-u_{i'})u_i\Delta > 0$. Thus, we can repeatedly perform the above steps till we get a sequence such that w'_i/u_i remains the same for all i and $\sum_i w_i = \sum_i w'_i$. Clearly $w'_i = \lambda u_i$ in this case. The function F continues to decrease, $F(\{(w_i,u_i)\}) \geq F(\{(w'_i,u_i)\})$ and

$$F(\{(w'_i,u_i)\}) = \sum_i \prod_{i'<i}(1-u_{i'})\lambda u_i = \lambda\left(1 - \prod_i (1-u_i)\right) \geq \lambda\left(1 - e^{-\sum_i u_i}\right)$$

But $\frac{1}{x}(1-e^{-x})$ is decreasing over $[0,1]$ and the worst case is $x = 1$.

In slot ℓ (renumbering the advertisers in the order of $w_i(\ell)/u_i(\ell)$) we get an expected reward of $w_i(\ell)$ if we reach i. But the events are independent in a particular slot. Thus the expected reward in a slot is bounded by (using independence and Claim 3) to be $(1-e^{-1})$ times $\sum_i w_i(\ell)$. We now apply linearity of expectation across the slots – *observe that the events across the slots are quite correlated.* The expected reward is at least $(1-e^{-1})$ times $\sum_\ell \sum_i w_i(\ell) = LP1$. Theorem 1 follows from Lemmas 2 and 3 and the application of Theorem 5.

4 Generalized Second Price with Reserve (GSP-Reserve)

The call outs for this problem would be exactly the same as the algorithm in Section 3. *We will however adjust the reserve prices.* The reserve price will be the same for all the advertisers being called out on that impression. In fact either we will run a single slot auction with a reserve price, or simply GSP for the M slots. The decision will depend on the LP solution found for this specific impression (and the contributions of different parts of the LP). Recall that the bid distribution $\mathbf{V}_{ij}$ of advertiser i on impression j is assumed to satisfy the MHR property. We use the following:

Lemma 4. *(Lemma 3.3 in [4]) For any random variable V following an MHR distribution, let $v^* = \arg\min_v\{v | v\Pr[V \geq v] \geq \frac{1}{2}\sum_{v' \geq v} v' \Pr[V = v']\}$. Then $\Pr[V \geq v^*] \geq e^{-2}$.*

The next lemma is a restatement of Lemma 2 and the subsequent analysis.

Lemma 5. *Given an impression $j(t)$, and define $v_1(t) = \max_{\ell:\varrho_1 \neq \varrho_\ell}(\tau^*_{j1} - \tau^*_{j\ell})/(\varrho_1 - \varrho_\ell)$. The call out to a set $S(t) \neq \emptyset$, ensures that $\sum_i \sum_{v \geq v_1(t)} p_{ijv} \geq 1$.*

Definition 1. *Given an impression $j(t)$, and the call out decision to a set $S(t) \neq \emptyset$ at time t, let $v^*_m(i,t) = \min\{v | 2v\Pr[V_{ij} \geq v] \geq \sum_{v' \geq v} v' p_{ijv'}\}$ and $\Psi(t) = \{i | i \in S(t)$ and $v_1(t) \leq v^*_m(i,t)\}$. Note that using Lemmas 4, and 5, we have $|\Psi(t)| \leq \lfloor e^2 \rfloor = 7$ since $i \in \Psi(t)$ contributes a probability mass of at least e^{-2}.*

Lemma 6. *Given an impression $j(t)$, and the call out decision to a set $S(t) \neq \emptyset$ at time t, we can set a single threshold $v^*(t) \geq v_1(t)$ such that if we set a reserve price $v^*(t)$ for a single slot then the revenue (ignoring the multiplicative discount factor ϱ_1) is at least $\frac{1}{4(7e^2+1)}\sum_{i \in S(t)} \sum_{v \geq v_1(t)} v p_{ijv}$.*

Proof. Let $\sum_{i\in S(t)}\sum_{v\geq v_1(t)} vp_{ijv} = Z$. We have two cases, (i) $\sum_{i\in\Psi(t)}\sum_{v\geq v_1(t)} vp_{ijv} \geq 7e^2Z/(7e^2+1)$ or (ii) otherwise. In case (i), pick the $i\in\Psi(t)$ such that $\sum_{v\geq v_1(t)} vp_{ijv}$ is maximized, which is at least $e^2Z/(7e^2+1)$ since $|\Psi(t)|\leq 7$. Let V_{ij} be the random variable that corresponds to the bid of advertiser i on impression j. Now since $v^*_m(i,t)\geq v_1(t)$ we have that

$$\begin{aligned}\sum_{v\geq v^*_m(i,t)} vp_{ijv} &\geq \Pr\left[V_{ij}\geq v^*_m(i,t)|V_{ij}\geq v_1(t)\right]\sum_{v\geq v_1(t)} vp_{ijv}\\ &\geq \Pr\left[V_{ij}\geq v^*_m(i,t)\right]\sum_{v\geq v_1(t)} vp_{ijv}\leq \frac{1}{e^2}\sum_{v\geq v_1(t)} vp_{ijv}\end{aligned}$$

which is at least $Z/(7e^2+1)$. Now, if we set $v^*(t)=v^*_m(i,t)$ then just from i we have $\sum_{i\in S(t)}\sum_{v\geq v^*(t)} p_{ijv} \geq \frac{1}{2}\sum_{v\geq v^*_m(i,t)} vp_{ijv}$ and therefore in this case the lemma is true.

In case (ii), we have $\sum_{i\in S(t)\setminus\Psi(t)}\sum_{v\geq v_1(t)} vp_{ijv}\geq Z/(7e^2+1)$. But for each $i\in S(t)\setminus\Psi(t)$ we have $v_1(t)\sum_{v\geq v_1(t)} p_{ijv}\geq\frac{1}{2}\sum_{v\geq v_1(t)} vp_{ijv}$ and as a consequence, $v_1(t)\sum_{i\in S(t)\setminus\Psi(t)}\sum_{v\geq v_1(t)} p_{ijv}$ is at least $Z/(2(7e^2+1))$. Consider setting $v^*(t)=v_1(t)$. Let $p=\sum_{i\in S(t)\setminus\Psi(t)}\sum_{v\geq v_1(t)} p_{ijv}$. Since $p\leq 1$ (from definition of $v_1(t)$, see Lemma 5) the probability of sale is at least $(1-\frac{1}{e})p$ which is bounded below by $p/2$. The Lemma follows in this case as well.

Lemma 7. *Given an impression $j(t)$, and the call out decision to a set $S(t)\neq\emptyset$ at time t, consider (i) If $\sum_{i\in S(t)}\sum_v v\varrho_{\ell(v)}y^*_{ijv\ell(v)}\geq 3\sum_{i\in S(t)}\sum_{v\geq v_1(t)} v\varrho_1 y^*_{ijv1}$ then call-out to $S(t)$ and run regular GSP. (ii) Otherwise call-out to $S(t)$ and run a single slot auction with the threshold $v^*(t)$ given by Lemma 6. This algorithm gives a revenue which is $\Omega(1)$ factor of the LP bound on efficiency which is given by $\sum_v\sum_{i\in S(t)}\sum_v v\varrho_\ell y^*_{ijv\ell(v)}$. Note that the call out decisions are based on optimizing the total value/ efficiency of the slots, and thus are feasible.*

Proof. Let the non-increasing ordered list of values that are returned for a time step be $a_1(t)$. Suppose we are in case (i). Then the revenue of GSP is at least $\sum_{r=1}^M \varrho_r a_{r+1}(t)$. Now since ϱ_r are decreasing, and $a_r(t)$ are non-increasing, we have $\sum_{r=1}^M \varrho_r a_{r+1}(t)\geq\sum_{r=1}^M \varrho_r a_r(t)-\varrho_1 a_1(t)$, which implies that

$$\mathbf{E}\left[\sum_{r=1}^M \varrho_r a_{r+1}(t)\right]\geq\mathbf{E}\left[\sum_{r=1}^M \varrho_r a_r(t)\right]-\varrho_1\mathbf{E}\left[a_1(t)\right] \quad (1)$$

We know from Section 3 that $\mathbf{E}\left[\sum_{r=1}^M \varrho_r a_r(t)\right]\geq\left(1-\frac{1}{e}\right)\sum_{i\in S(t)}\sum_v v\varrho_\ell y^*_{ijv\ell(v)}$. We observe that $\mathbf{E}\left[a_1(t)\right]\leq\sum_{i\in S(t)}\sum_{v\geq v_1(t)} vy^*_{ijv1}$. This is easily seen if we write an LP for the maximum value seen (this LP is for analysis only). Let x_{iv} be the probability that $i\in S(t)$ is the maximum with value v. Then (we drop the index j for convenience):

$$\mathbf{E}\left[a_1(t)\right]\leq LPMAX=\max\sum_i\sum_v vx_{iv} \qquad \text{s.t.} \qquad \begin{array}{l}\sum_i\sum_v x_{iv}\leq 1\\ x_{iv}\leq p_{iv}\\ x_{iv}\geq 0\end{array}$$

The optimum solution of $LPMAX$ is $x^*_{iv} = p_{iv}$ for $v > \tau$ and $x^*_{iv} \leq p_{iv}$ for one i and $v = \tau$. Here τ is the optimum dual variable for the constraint $\sum_{i \in S(t)} \sum_v x_{iv} \leq 1$. Note that $\sum_{i \in S(t)} \sum_v x^*_{iv} = 1$. For $v > v_1(t)$ we have $y^*_{ijv1} = p_{ijv}$ and $v < v_1(t)$ we have $y^*_{ijv1} = 0$. Moreover $\sum_{i \in S(t)} \sum_{v \geq v_1(t)} y^*_{ijv1} = 1$. Likewise for $v > \tau$ we have $x^*_{iv} = p_{iv}$ and $v < \tau$ we have $x^*_{iv} = 0$ and $\sum_{i \in S(t)} \sum_{v \geq \tau} x^*_{iv} = 1$.

Suppose that $\tau < v_1(t)$. We arrive at a contradiction because $\sum_{i \in S(t)} \sum_{v \geq \tau} x^*_{iv} > \sum_{i \in S(t)} \sum_{v \geq v_1(t)} y^*_{ijv1} = 1$ which implies that we are exceeding the probability mass of 1 for the maximum. On the other hand if $\tau > v_1(t)$, then we again have a contradiction that $\sum_{i \in S(t)} \sum_{v \geq v_1(t)} y^*_{ijv1} > \sum_{i \in S(t)} \sum_{v \geq \tau} x^*_{iv} = 1$ which implies $\{y^*_{ijv1}\}$ were not feasible.

As a consequence, $\tau = v_1(t)$ and for $v > \tau = v_1(t)$ we have $x^*_{iv} = y^*_{ijv1} = p_{ijv} = p_{iv}$. For $v < \tau = v_1(t)$ we have $x^*_{iv} = y^*_{ijv1} = 0$. Therefore $\mathbf{E}\left[a_1(t)\right] \leq LPMAX = \sum_{i \in S(t)} \sum_v v x^*_{iv} = \sum_{i \in S(t)} \sum_v v y^*_{ijv1}$ as claimed. Applying this claim to Equation 1, and the fact that $\sum_{i \in S(t)} \sum_v v y^*_{ijv1} \leq \frac{1}{3} \sum_{i \in S(t)} \sum_v v \varrho_\ell y^*_{ijv\ell(v)}$, we get

$$\mathbf{E}\left[\sum_{r=1}^{M} \varrho_r a_{r+1}(t)\right] \geq \left(1 - \frac{1}{e} - \frac{1}{3}\right) \sum_{i \in S(t)} \sum_v v \varrho_\ell y^*_{ijv\ell(v)}$$

Thus in this case the expected revenue is $\Omega(1)$ of the LP bound on the efficiency.

Suppose we are in case (ii). By Lemma 6 we are guaranteed an expected revenue of $\Omega(1)$ times $\varrho_1 \sum_{i \in S(t)} \sum_{v \geq v_1(t)} v p_{ijv} \geq \sum_{i \in S(t)} \sum_{v \geq v_1(t)} v \varrho_1 y^*_{ijv1} \geq \frac{1}{3} \sum_{i \in S(t)} \sum_v v \varrho_{\ell(v)} y^*_{ijv\ell(v)}$. In this case also the expected revenue is $\Omega(1)$ of the LP bound; the lemma follows.

Proof. (Of Theorem 2). Let App be the policy that approximately maximizes the efficiency. Let OPT_G be the optimum GSP-Reserve policy. Let OPT_B be the optimum policy which maximizes the total value. Given a policy P let $\textsc{Gsp}(P)$ denote the expected revenue of the policy if the charged as GSP, $\textsc{BestKW}(P)$ denote the expected (weighted) efficiency. Then for any policy $\textsc{Gsp}(P) \leq \textsc{BestKW}(P)$. Let $R(\textsc{App})$ be the revenue of the policy in Lemma 7. Therefore, for some absolute constant $\alpha \geq 1$,

$$\textsc{Gsp}(\mathrm{OPT}_G) \leq \textsc{BestKW}(\mathrm{OPT}_G) \leq \textsc{BestKW}(\mathrm{OPT}_B) \leq LP1 \leq \alpha\ R(\textsc{App})$$

The theorem follows (again appealing to Theorem 5).

References

1. Azar, Y., Birnbaum, B.E., Karlin, A.R., Nguyen, C.T.: On revenue maximization in second-price ad auctions. In: Fiat, A., Sanders, P. (eds.) ESA 2009. LNCS, vol. 5757, pp. 155–166. Springer, Heidelberg (2009)
2. Babaioff, M., Immorlica, N., Kempe, D., Kleinberg, R.: Knapsack secretary problem with applications. In: Charikar, M., Jansen, K., Reingold, O., Rolim, J.D.P. (eds.) RANDOM 2007 and APPROX 2007. LNCS, vol. 4627, pp. 16–28. Springer, Heidelberg (2007)

3. Barlow, R.E., Marshall, A.W., Proschan, F.: Properties of probability distributions with monotone hazard rate. The Annals of Mathematical Statistics 34(2), 375–389 (1963)
4. Bhattacharya, S., Goel, G., Gollapudi, S., Munagala, K.: Budget constrained auctions with heterogeneous items. In: STOC (2010), http://arxiv.org/abs/0907.4166
5. Buchbinder, N., Jain, K., Naor, J.: Online primal-dual algorithms for maximizing ad-auctions revenue. In: Arge, L., Hoffmann, M., Welzl, E. (eds.) ESA 2007. LNCS, vol. 4698, pp. 253–264. Springer, Heidelberg (2007)
6. Calinescu, G., Chekuri, C., Pál, M., Vondrák, J.: Maximizing a submodular set function subject to a matroid constraint (extended abstract). In: Fischetti, M., Williamson, D.P. (eds.) IPCO 2007. LNCS, vol. 4513, pp. 182–196. Springer, Heidelberg (2007)
7. Chakrabarty, D., Goel, G.: On the approximability of budgeted allocations and improved lower bounds for submodular welfare maximization and gap. In: Proceedings of FOCS, pp. 687–696 (2008)
8. Chawla, S., Hartline, J., Malec, D., Sivan, B.: Sequential posted pricing and multiparameter mechanism design. In: STOC (2010), http://arxiv.org/abs/0907.2435
9. Devanur, N., Hayes, T.: The adwords problem: Online keyword matching with budgeted bidders under random permutations. In: Proc. 10th Annual ACM Conference on Electronic Commerge, EC (2009)
10. Ghosh, A., McAfee, P., Papineni, K., Vassilvitskii, S.: Bidding for representative allocations for display advertising. In: Leonardi, S. (ed.) WINE 2009. LNCS, vol. 5929, pp. 208–219. Springer, Heidelberg (2009)
11. Ghosh, A., Rubinstein, B.I.P., Vassilvitskii, S., Zinkevich, M.: Adaptive bidding for display advertising. In: WWW, pp. 251–260 (2009)
12. Guha, S., McGregor, A.: Approximate quantiles and the order of the stream. In: Proc. of PODS, pp. 273–279 (2006)
13. Guha, S., Munagala, K.: Adaptive uncertainty resolution in bayesian combinatorial optimization problems. ACM Transactions of Algorithms (to appear)
14. Kleinberg, J., Rabani, Y., Tardos, E.: Allocating bandwidth for bursty connections. SIAM J. Of Computing 30(1), 171–217 (2001)
15. Kleinberg, R.: A multiple-choice secretary problem with applications to online auctions. In: Proceedings of SODA, pp. 630–631 (2005)
16. Kleinberg, R., Slivkins, A., Upfal, E.: Multi-armed bandits in metric spaces. In: STOC, pp. 681–690 (2008)
17. Kulik, A., Shachnai, H., Tamir, T.: Maximizing submodular set functions subject to multiple linear constraints. In: SODA, pp. 545–554 (2009)
18. Marchetti-Spaccamela, A., Vercellis, C.: Stochastic on-line knapsack problems. Mathematical Programming 68, 73–104 (1995)
19. Mehta, A., Saberi, A., Vazirani, U., Vazirani, V.: Adwords and generalized on-line matching. In: FOCS, pp. 264–273 (2005)
20. Muthukrishnan, S.: Adexchanges: Research issues. In: Leonardi, S. (ed.) WINE 2009. LNCS, vol. 5929, pp. 1–12. Springer, Heidelberg (2009)
21. Myerson, R.B.: Optimal auction design. Math. of OR 6(1), 58–73 (1981)
22. Nolet, M.: http://www.mikeonads.com/2009/08/30/rtb-part-i-what-is-it/
23. Pandey, S., Olston, C.: Handling advertisements of unknown quality in search advertising. In: Proceedings of NIPS, pp. 1065–1072 (2006)

24. P. release, http://www.doubleclick.com/insight/blog/archives/doubleclick-advertising-exchange/announcing-the-new-doubleclick-adexchange.html
25. Rockafellar, R.T.: Convex Analysis. Princeton University Press, Princeton (1970)
26. Tanenbaum, A.S.: Computer Networks, 3rd edn. Prentice-Hall, Englewood Cliffs (1996)
27. Vee, E., Vassilvitskii, S., Shanmugasundaram, J.: Optimal online assignment with forecasts. In: EC (2010)