Autobidding with Constraints

Gagan Aggarwal, Ashwinkumar Badanidiyuru(✉), and Aranyak Mehta

Google, Mountain View, USA
ashwinkumarbv@gmail.com

Abstract. Autobidding is becoming increasingly important in the domain of online advertising, and has become a critical tool used by many advertisers for optimizing their ad campaigns. We formulate fundamental questions around the problem of bidding for performance under very general affine cost constraints. We design optimal single-agent bidding strategies for the general bidding problem, in multi-slot truthful auctions. The novel contribution is to show a strong connection between bidding and auction design, in that the bidding formula is optimal if and only if the underlying auction is truthful.

Next, we move from the single-agent view to a full-system view: What happens when all advertisers adopt optimal autobidding? We prove that in general settings, there exists an equilibrium between the bidding agents for all the advertisers. As our main result, we prove a *Price of Anarchy* bound: For any number of general affine constraints, the total value (conversions) obtained by the advertisers in the bidding-agent equilibrium is no less than 1/2 of what we could generate via a centralized ad allocation scheme, one which does not consider any auction incentives or provide any per-advertiser guarantee.

Keywords: Automated bidder · Price of anarchy · Constrained optimization

1 Introduction

Autobidding is taking on an increasingly important role in online advertising [5] and has already become a critical tool used by many advertisers for optimizing their ad campaigns. Given its importance in the ad ecosystem, autobidding deserves fundamental investigation into algorithms and properties. In this paper, we formulate the questions of designing optimal bidding algorithms, study the interaction of bidding with the underlying auction, and study system equilibrium properties when all advertisers adopt autobidding.

The motivation behind autobidding is performance and product simplicity. The main idea is that instead of asking advertisers for fine-grained bids (e.g., a bid per keyword), the ad platform asks for higher level goals and higher level constraints. An *Autobidding agent* then converts these goals and constraints into per-query bids at serving time, based on predictions of performance of each

I. Caragiannis et al. (Eds.): WINE 2019, LNCS 11920, pp. 17–30, 2019.
https://doi.org/10.1007/978-3-030-35389-6_2

potential ad impression. Besides increased performance, these products also provide for a much simpler interaction with the ad system. For example, in some settings like Google's Universal App Campaigns (UAC) product [6], advertisers do not target at all, but only provide targets for cost-per-install and other goals. In other cases advertisers continue targeting via keywords to identify what queries are of interest, but let the system adjust bids based on predicted performance. In either case, the bidding agent will automatically adjust bids so as to give maximum performance for the campaign, even in a dynamically changing environment, as query volume and features change over time.

There are several autobidding products in the market. The simplest and oldest is that of budget optimization, in which the advertiser provides targeting and a (daily) budget, and the system bids on its behalf. This is a well-studied topic with significant related work. We now have increasingly sophisticated products which allow for performance-based optimization of campaigns, based on goals that advertisers may care more about, by leveraging predicted conversions (sales). For example, Target Cost-per-acquisition (tCPA), Enhanced-CPC (ECPC), and products aiming for deeper optimizations, such as Return on Ad Spend (ROAS), and post-install-value optimization (see, e.g., [5] for more detailed description of these products).

In this paper we formulate and answer several fundamental questions in autobidding. Specifically, (1) find an optimal bidding formula for very general constraints and connect it to the truthfulness of the underlying auction, and (2) quantify the price of anarchy in equilibrium when all advertisers adopt the optimal autobidding.

Remark: These are critical questions from an ad platform's point of view, but they are also interesting and novel from a purely theoretical view as several of the important autobidding products go beyond the classic profit-maximization setting and instead, follow the framework of maximizing value (e.g., number of conversions) under constraints on the average cost (of clicks or conversions) and a budget on total spend. One can consider such objectives and constraints as generalizations of the well-studied budget constraint – the difference now is that the cap on spend is not a fixed number (i.e., budget), but is a function of the specific items allocated (see Sect. 2 for details).

1.1 Overview, Results, and Techniques

In Sect. 2 we formulate the single agent bidding problem, for the general setting of value maximization under general affine constraints (in a multi-slot truthful auction). Specifically, given an advertiser's goals and constraints, and given predictions at query time, how should the bidding agent bid on behalf of the advertiser? This formulation generalizes all the autobidding products we mentioned above. Our two main results are:

– In Sect. 3 we show how to derive an optimal bidding formula assuming we have access to the cost-value landscape. In particular, we show that there is

a simple bidding formula which takes in the value for an item (including the predictions for probability of click or conversions), and converts it into a bid into the auction. While the technique of using LP duality to find an optimal allocation is not entirely new, our novel contribution here is to connect it to bidding, and more specifically, back to the truthfulness of the auction and show that the bidding formula is optimal if and only if the auction is truthful.
- In Sect. 5, we prove a "price-of-anarchy" result which is the most technically challenging and novel portion of the paper: For an autobidding setting with any number of general affine constraints, if all advertisers adopt autobidding, then the total value generated for all advertisers in equilibrium is at least a factor 1/2 of the total value we could generate via a centralized ad allocation scheme – one which does not need to consider any pricing or auction incentives constraint, or have any per-advertiser optimization guarantee.

  For this result, we extend the definition of liquid welfare [3,11] from the budgeted setting to the general affine constraints setting. Then, we use a charging argument, in which we use the structure of equilibrium bids, the truthfulness property of the underlying auction, as well as the nature of the affine constraints to bound the liquid welfare of global allocation in terms of the liquid welfare at equilibrium.

For the sake of completeness, we also provide two additional results, which may be considered as using somewhat standard techniques from the literature: Firstly, in Sect. 4, we show how a Multiplicative Weights Update based method of control feedback can help find the optimal parameters of the bidding formula assuming full access to the cost-value distribution. While the algorithm is a simple instantiation of MWU for solving LPs, which is standard, we do this by interpreting the hyperplanes generated by MWU with dual weights as the realization of a truthful auction, thus connecting truthfulness to the bidding formula.

Secondly, we show in Sect. 6 that, for multi-slot, general constraints setting, there exists a pure strategy bidding equilibrium, under certain technical smoothness assumptions. This result follows by defining a map from the space of dual variables to itself, using the optimal bidding formula, so that the fixed point of the map are the optimal parameters of the formula.

1.2 Related Work

Bidding algorithms have been studied previously in various forms and we describe some related work below. We have specifically two new contributions: the connection of bidding in auction with the truthfulness of the underlying auction, and analyzing the setting of multiple advertisers bidding optimally in a truthful auction and bounding the price of anarchy.

As mentioned above, perhaps the simplest autobidding product is budget optimization. This has been a well-studied topic, in particular [13] provided a formulation for this problem, and gave simple practical uniform bid strategies which achieve a constant factor of the optimal bidding under any auction –

however, they do not consider multiple bidders or equilibrium properties. Also somewhat related is work on back-end system optimization for budget management (not as a bidding agent); this includes work on ad allocation, budget throttling, and bid lowering, e.g., [4,8,12,15–17].

Besides budget optimization bidding, previous literature includes several results on the so-called real-time bidding (RTB) in the context of display ads. A related paper is [9], which considers the problem of bidding algorithms for performance advertising. Similar to the work here, they use a Primal Dual formulation to find a bidding formula. However, there are some salient differences compared to our work: Firstly, their objective is *global value* maximization (sum over all bidders values) under volume constraints. Secondly, the pricing is simply first price, and it is not immediately clear how to extend this to second price auction (or a truthful auction for multiple slots). Bidding into an auction is a more difficult question, as bidders set prices for each other and thus have to be in equilibrium. Indeed, we show that no bidding formula can work in a non-truthful multi-slot auction, and even in a second price auction, the global value generated in an equilibrium solution can be bounded away from the global optimum by a factor of as much as 2.

There are several other interesting papers on RTB, e.g., [14,19,20]. The latter paper focuses on learning the underlying traffic distributions and using them to find a bid. The bidding question considered there is simple if the distribution is known (due to the simple nature of the constraints), but the innovation lies in learning the distributions from possibly partial feedback, which is not the focus of our work.

In another related work [7], the authors consider a different but related equilibrium question, in the setting of backend budget throttling (*aka* pacing, in which a budget constrained advertiser is throttled out of some subset of unprofitable auctions). The authors consider the question of whether there is a regret-free stable solution if we use optimal budget throttling for all budget constrained bidders in a single slot auction. However, they do not further analyze the price of anarchy in such a stable solution.

Finally, a very relevant line of work is that of solving online stochastic linear programs [10] and online stochastic convex programs [1]. The specific problems we study are actually instantiations of the more general problem they study, and they also use duality theory to find optimal allocations (in more general settings with stochastic input). We note two novel contributions in our work: Firstly, the connection to truthfulness, i.e., the dual based allocation gives rise to a bidding formula which is optimal if and only if the auction is truthful. And secondly, we study the equilibrium properties if every bidder uses this algorithm in a truthful auction, and prove a bound on the price of anarchy.

2 Preliminaries

There is a set of advertisers A bidding into an ad auction. There is a large set of queries I, each with potentially multiple slots (*aka* positions) S. For each query i, there is an auction which takes in bids and determines which advertisers show in which slots, and determines a cost-per-click (cpc^a_{is}) for each advertiser $a \in A$ and slot $s \in S$.

We define indices into different sets as follows. Let $i \in I$ be an query, $s \in S$ be a slot and $c \in C$ be a constraint. Let the click through rate (CTR), the probability that a user clicks on an ad of advertiser a on slot s of query i be ctr^a_{is}. Let the cost per click for winning slot s of query i be cpc^a_{is}. Let the value to a of a click on impression i be v^a_i – this is the estimate of the total downstream value accrued by a after the click, which we assume is independent of the slot s that the ad was in. Let x^a_{is} be indicator variable if slot s of impression i was allocated to a. Note that we will also abuse notation and refer to query i as an impression (which makes sense for the case when $|S| = 1$).

We study the problem of finding an optimal bidding strategy for each advertiser $a \in A$, assuming that the bids of all other advertisers are fixed. For this problem, even though all the parameters in the problem definitions are indexed by a, for simplicity we drop the index a when we study the optimal bidding strategy for a (here, and in Sects. 3 and 4). In Sects. 5 and 6 we will reintroduce the index a as we study what happens when all advertisers bid according to the proposed bidding formula.

The goal of an advertiser is to maximize its total value i.e. $\sum_{i,s} x_{is} ctr_{is} v_i$. Additionally we have several affine constraints on the spend of the advertiser. This can be formalized by integer program in Fig. 1, in which the index c stands for the constraints, and the v_{ic} are non-negative constants (one per query and constraint).

$$\max \sum_{i,s} x_{is} ctr_{is} v_i$$
$$\forall c, \quad \sum_{i,s} x_{is} ctr_{is} cpc_{is} \leq B_c + \sum_{i,s} x_{is} ctr_{is} v_{ic}$$
$$\forall i, \sum_{s} x_{is} \leq 1$$
$$\forall i, s, x_{is} \in \{0, 1\}$$

Fig. 1. The Integer Program for value maximization for an advertiser under general affine constraints.

Next we show how many products in the industry can be modeled with the above set of constraints.

Budget Optimization: In this case there is a single constraint c with $v_{ic} = 0\ \forall i$, and B_c is the budget B. So the constraint is simply $\sum_{i,s} x_{is} ctr_{is} cpc_{is} \leq B$. Here, v_i is sometimes taken to be 1 for all i, which means the goal would be to maximize clicks.

Target CPA: In the TCPA product the goal is to maximize the number of conversions subject to the constraint that average cost per conversion does not exceed an advertiser given target value T. Here once again we have a single constraint c. Here v_i represents the predicted number of conversions (*aka* acquisitions, or sale events) that the advertiser gets after a click on impression i (usually called $pcvr_i$, and also assumed to be independent of the slot s). We take $B_c = 0$ and $v_{ic} = T \cdot v_i$, $\forall i$. Note that we can rewrite the constraint which becomes $\frac{\sum_{i,s} x_{is} ctr_{is} cpc_{is}}{\sum_{i,s} x_{is} ctr_{is} v_i} \leq T$. This means that the ratio of the total expected spend to the total expected number of conversions should be at most T, as required.

Target on CPA and CPC: In some bidding products the goal is to maximize number of conversions, but we have two constraints. One is to ensure that average cost per conversion does not exceed T (the same as in TCPA) and the other is to ensure that average cost per click is at most M (both T and M are given by the advertiser). For the second constraint we set $B_c = 0$ and $v_{ic} = M$.

Note that the last two settings above can also be accompanied by a separate budget constraint.

We will also make an assumption throughout that the parameters of this problem are in general position.

3 Bidding Formula

In this section we show that there is an optimal bidding formula of the form $b(i) = \frac{v_i + \sum_c \alpha_c v_{ic}}{\sum_c \alpha_c}$, and this holds if and only if the auction is truthful. If an advertiser bids according to this bidding formula then they violate their constraints by at most $|C|$ impressions (where $|C|$ is the number of constraints) and get a value which is at least the value of an optimal bidding strategy minus the value of at most $|C|$ impressions.

To prove the result we consider Integer Program 1 and relax it to a Linear Program and also consider the corresponding dual LP.

Primal Linear Program

$$\max \sum_{i,s} x_{is} ctr_{is} v_i$$

$$\forall c, \ \sum_{i,s} x_{is} ctr_{is} cpc_{is} \leq B_c + \sum_{i,s} x_{is} ctr_{is} v_{ic}$$

$$\forall i, \ \sum_{s} x_{is} \leq 1$$

$$\forall i, s, x_{is} \ \geq 0$$

Dual Linear Program

$$\min \sum_i \delta_i + \sum_c \alpha_c B_c$$

$$\forall i, s \left| \begin{matrix} \delta_i + \\ \sum_c \alpha_c ctr_{is}(cpc_{is} - v_{ic}) \end{matrix} \right\} \geq ctr_{is} v_i$$

$$\forall i, \ \delta_i \geq 0$$

$$\forall c, \ \alpha_c \geq 0$$

Let $\{\alpha_c\}$ be the optimal dual solution. Define

$$\Delta_{is} = ctr_{is} v_i - \sum_c \alpha_c ctr_{is}(cpc_{is} - v_{ic})$$

Then the dual constraints can be written as $\delta_i \geq \Delta_{is}$ $\forall i, s$.

Now define the bidding formula as

$$b(i) := \frac{v_i + \sum_c \alpha_c v_{ic}}{\sum_c \alpha_c}$$

This is the bid that the bidder puts into the auction for query i. The auction determines the slot and price for the bidder. We will assume that ties are broken arbitrarily. Note that the dual constraints can also be written as

$$\frac{\delta_i}{\sum_c \alpha_c} \geq ctr_{is}(b_i - cpc_{is}),\ \forall\ i, s \tag{2}$$

We first note some properties of the optimal solution to the primal and dual programs.

Lemma 1. *Let $\{x_{is}\}$ and $\{\alpha_c\}$ be optimal solutions to primal and dual linear program. Then they satisfy the following properties.*

1. $\delta_i = max(0, max_s(\Delta_{is}))$
2. *If $\delta_i > 0$ and there is a unique s such that $\delta_i = \Delta_{is}$ then $x_{is} = 1$*
3. *If $\delta_i = 0$ with $b(i) < cpc_{is}, \forall s$ then $x_{is} = 0$.*
4. *There can be at most $|C|$ impressions i such that $\exists s \neq s'$ with $\delta_i = \Delta_{is} = \Delta_{is'}$.*
5. *There can be at most $|C|$ impressions i such that $\delta_i = 0$ and $\exists s, b(i) = cpc_{is}$.*

Proof. The proof follows from linear programming complementary slackness and will be included in the full version of the paper.

Theorem 1. *Bidding strategy $b(i)$ gives a solution which has value at least OPT minus value of $2|C|$ impressions and violates each constraint by at most $2|C|$ impressions if and only the auction is truthful.*

Proof. Fix a query i, for which the bidder's bid is b_i. A truthful auction will allocate the bidder to the slot which maximizes its profit given the bid, i.e., the slot s which maximizes $ctr_{is}(b_i - cpc_{is})$ (of course the cpcs are derived during the auction itself from other bidders). A non-truthful auction will, for some value of the bid b_i (and some values of other bidder bids) allocate the bidder to some other slot $s' \neq s$.

Now consider the LP. By Lemma 1 point 2 $x_{is} = 1$, precisely for the tight dual constraint. But, by Eq. (2), this is precisely the one which maximizes the profit. Therefore the LP allocation solution matches the solution that a truthful auction would choose with the same bids, and would not match the allocation of a non-truthful auction, for at least some values of bids.

Next we prove the bound on the value achieved by the bidding strategy to compare it to the optimal achievable value.

$$\begin{aligned}
&\text{Value of bidding strategy}\\
&\geq \sum_{\delta>0,\text{unique s with } \delta_i=\Delta_{is}} ctr_{is}v_i\\
&= \sum_{\delta>0,\text{unique s with } \delta_i=\Delta_{is}} x_{is}ctr_{is}v_i\\
&= \sum_{i,s} x_{is}ctr_{is}v_i - \sum_{i,s,\delta=0 \text{ or } \Delta_{is}=\Delta_{is'}} x_{is}ctr_{is}v_i\\
&\geq \sum_{i,s} x_{is}ctr_{is}v_i - |C| \cdot max_{is}ctr_{is}v_i\\
&= OPT - 2|C| \cdot max_{is}ctr_{is}v_i
\end{aligned}$$

Here Eq. 1 is because bidder wins at least these impressions, Eq. 2 is from Lemma 1 point 2 and Eq. 4 is from Lemma 1 points 3, 4, 5. Next we show that the constraints are violated by at most $2|C|$ impressions. This is simple because by Lemma 1, Points 3, 4, 5 there are at most $2|C|$ impressions for which $x_{is} = 1$ and bidder doesn't win it or $x_{is} < 1$ and bidder wins it.

We will include an intuitive example illustrating the bidding formula in full version of the paper.

4 Bidding Algorithm

In this section we will give a bidding algorithm which computes the bidding formula and bids accordingly. The algorithm is an application of multiplicative weight update (MWU) method. While it is well known how to use MWU to solve a linear program we note that we specifically need a bidding formula which does not depend on other bidders bids. We show how feasibility oracle used in MWU to solve linear program translates to bidding formula and hence the bidding formula can be used to answer the separation oracle.

We use the MWU algorithm to solve $Ax \geq b$ subject $x \in P$ when a feasibility oracle for any $c, d, \exists ?x \in P : c^T x \geq d$ is given. We borrow this from Sect. 3.2 of [2] (Included in full version of the paper). MWU in each step maintains a weight vector w of same number of rows as A and in each step multiplicatively updates w based on how much each constraint is violated from solution in previous step. Then in each step it asks oracle question of the form $w^T Ax \geq w^T b$.

Let $\mathcal{V}$ be upper bound on what the OPT and let $\mathcal{V}_c$ be an upper bound on $|B_c + \sum_{i,s} x_{is}ctr_{is}(v_{ic} - cpc_{iks})|$.

Theorem 2. *In* $T \geq O(\frac{1}{\delta^3})$ *steps Algorithm 1 converges to a solution which satisfies the following.*

1. *Value*$\geq OPT - \delta \cdot \mathcal{V}$.

2. For each constraint we have $\sum_{i,s} x_{is} ctr_{is}(cpc_{is} - v_{ic}) \leq B_c + \delta \mathcal{V}_c$

Algorithm 1. Bidder

1: **for** $i = 1, \ldots, O(\frac{1}{\delta})$ **do**
2: $\quad \mathcal{V}_{OPT} = i \cdot \delta \cdot \mathcal{V}$ (Guess for the value of OPT)
3:

$$A = \begin{pmatrix} \cdots\cdots\cdots & ctr_{is} v_i/\mathcal{V} & \cdots \\ \cdots\cdots\cdots & \cdots & \cdots \\ \cdots\cdots\cdots & ctr_{is}(v_{ic} - cpc_{is})/\mathcal{V}_c & \cdots \\ \cdots\cdots\cdots & \cdots & \cdots \end{pmatrix} \qquad b = \begin{pmatrix} \mathcal{V}_{OPT}/\mathcal{V} \\ \cdots \\ -B_c/\mathcal{V}_c \\ \cdots \end{pmatrix}$$

Where each row (except first one) of A corresponds to constraint c and column for impression/slot i, s.
4: $\quad P$ be a convex constraints on x_{is} denoting $0 \leq x_{is}$ and $\sum_s x_{is} \leq 1$.
5: $\quad$ Run algorithm MWU to check feasibility of $Ax \geq b$ such that $x \in P$.
6: $\quad$ **for** $t = 1, \ldots T$ (Each step of MWU) **do**
7: $\quad\quad$ Let $w^t = (w_1^t, w_2^t, \ldots)$ be the weight vector maintained by MWU.
8: $\quad\quad$ Let $\alpha_c = (w_{c+1}^t/\mathcal{V}_c)/(w_1^t/\mathcal{V})$
9: $\quad\quad$ Define oracle O for $F = \exists ?x \in P : w^T Ax \geq w^T b$.
- Run bidder with bidding strategy $\frac{v_i + \sum_c \alpha_c v_{ic}}{\sum_c \alpha_c}$
- Let $x_{is} = 1$ if bidder won the impression at slot s, otherwise $x_{is} = 0$.
- Check if this solution $\{x_{is}\}$ satisfies $w^T Ax \geq w^T b$

10: $\quad$ **end for**
11: $\quad$ If algorithm MWU returns infeasibility then break
12: **end for**

Proof. Consider the value of i such that $OPT \geq i \cdot \delta \cdot \mathcal{V} \geq OPT - \delta \cdot \mathcal{V}$. We will fix this iteration for the remaining part of the proof. We know that for such i we have feasible solution for $Ax \geq b, x \in P$. By proof of MWU we know that as long as oracle O is implemented such that feasibility of $F = \exists ?x \in P : w^T Ax \geq w^T b$ correctly then MWU returns a feasible solution. We show this by showing that this is equivalent to the bidding strategy in step 9.

First note that bidder in step 9 which produces solution $x_i = 1$ if and only if $\frac{v_i + \sum_c \alpha_c v_{ic}}{\sum_c \alpha_c} \geq cpc_{is}$. We will show that checking $w^T Ax \geq w^T b$ for the output of bidder is equivalent to solving F.

$$w^T Ax - w^T b = \left(\sum_{i,s} \frac{w_1^t x_{is} ctr_{is} v_i}{\mathcal{V}} + \sum_{i,s,c} \frac{w_{c+1}^t x_{is} ctr_{is}(v_{ic} - cpc_{is})}{\mathcal{V}_c} \right) - \frac{w_1^t \mathcal{V}_{OPT}}{\mathcal{V}} + \sum_c \frac{w_{c+1}^t B_c}{\mathcal{V}_c}$$
$$= \frac{w_1^t}{\mathcal{V}} \sum_c \alpha_c \sum_{i,s} x_{is} ctr_{is} \left(\frac{v_i + \sum_c \alpha_c v_{ic}}{\sum_c \alpha_c} - cpc_{is} \right) - \frac{w_1^t \mathcal{V}_{OPT}}{\mathcal{V}} + \sum_c \frac{w_{c+1}^t B_c}{\mathcal{V}_c}$$

It is easy to see that the right hand side is maximized when $x_{is} = 1$ when $\frac{v_i + \sum_c \alpha_c v_{ic}}{\sum_c \alpha_c} \geq cpc_{is}$ which is the exact set of impressions/slots won by the bidder. Hence it is enough to check for this vector if $w^T Ax \geq w^T b$ which completes the proof.

5 Price of Anarchy

In this section we show a factor 1/2 price of anarchy of any bidder which optimizes for each bidder separately as opposed to optimizing globally for everyone. We consider the special case when each impression has a single slot and hence we will drop the subscript s from remaining part of this section.

5.1 Price of Anarchy Objective

We consider Liquid Welfare as defined in [3] as our objective function. This is defined as sum over all advertisers of the maximum revenue that can be got from an advertiser. This turns out to be the following.

$$\sum_a \left(\min_c B_a^c + \sum_i x_i^a ctr_i^a v_{ic}^a \right)$$

Let OPT be the welfare objective achieved by OPT and let ALG be the welfare objective achieved at equilibrium. Also for any subset S of bidders, define $OPT(S)$ to be the contribution of bidders in S to OPT's welfare objective. Define $ALG(S)$ analogously.

5.2 Price of Anarchy Is Bounded by 2

Let $c'(a)$ be one of the indices that decides the contribution of bidder a of OPT's welfare objective function. Let $C(a)$ be the set of constraints that are tight for bidder a (let $C(a)$ be empty if no constraints are tight). Let A_1 be the set of bidders who are completely unconstrained at equilibrium and let A_2 be the remaining bidders.

Bidders in A_1 are bidding infinity at equilibrium and are winning everything they are interested in. So for $a \in A_1$, a's contribution to ALG is $\sum_i v_i^a$ which is the maximum possible contribution that bidder a can make to the welfare objective.

$$ALG(A_1) >= OPT(A_1)$$

Next we split $OPT(A_2)$ into two parts and bound each separately. For this, define $O(a)$ to be the set of impressions allocated to bidder a in OPT, and let $A(a)$ be the set of impressions allocated to bidder a at equilibrium.

The proof is by a charging argument to bound the liquid welfare of the global allocation in terms of global welfare at equilibrium. For this, we consider two types of impressions – impressions where the global optimal allocation overlaps with the allocation at equilibrium (i.e. $O(a) \cap A(a)$), and the impressions where the two allocations differ (i.e. $O(a) - A(a)$). At first glance, it would appear that the efficiency contribution of the overlapping impressions is trivially equal, and that we need to only worry about the non-overlapping impressions. But interestingly, because the efficiency contribution of each bidder is the minimum over several "types" of value, the efficiency contribution of the same subset of impressions with identical winning bidders may not be equal, and may in fact be incomparable, for the two allocations. To overcome this difficulty, we identify a subset $C(a)$ of constraints, which can be used to characterize both the contribution to liquid welfare as well as the query-level equilibrium bids for bidder a. In particular, this subset has the following three properties:

– For a given bidder a, its bid at equilibrium is no less than a particular convex linear combination of the RHS of bidder a's constraints indexed by $C(a)$.

$$b^a(i) = \frac{v_i^a + \sum_c \alpha_c^a v_{ic}^a}{\sum_c \alpha_c^a} = \frac{v_i^a + \sum_{c \in C(a)} \alpha_c^a v_{ic}^a}{\sum_{c \in C(a)} \alpha_c^a} \geq \frac{\sum_{c \in C(a)} \alpha_c^a v_{ic}^a}{\sum_{c \in C(a)} \alpha_c^a}$$

Using this bound on bids for each impression we can bound the "portion" of OPT from $O(a) - A(a)$.

$$\begin{aligned} ALG &\geq Total\_ALG\_Spend \\ &\geq \sum_{a \in A_2} \sum_{i \in O(a) \cap A(a)} ALG\_Spend(i) \\ &\geq \sum_{a \in A_2} \sum_{i \in O(a) - A(a)} ctr_i^a b^a(i) \\ &\geq \sum_{a \in A_2} \sum_{i \in O(a) - A(a)} ctr_i^a \frac{\sum_{c \in C(a)} \alpha_c^a v_{ic}^a}{\sum_{c \in C(a)} \alpha_c^a} \\ &= \sum_{a \in A_2} \frac{\sum_{c \in C(a)} \alpha_c^a \sum_{i \in O(a) - A(a)} ctr_i^a v_{ic}^a}{\sum_{c \in C(a)} \alpha_c^a} \end{aligned} \tag{3}$$

– A bidder a's contribution to the welfare at equilibrium is equal to the sum of its "c-type" values for the impressions it gets at equilibrium for any $c \in C(a)$. This in turn implies that its contribution is also equal to the sum (over its equilibrium allocation) of any convex linear combination of its c-type values for $c \in C(a)$.

$$\begin{aligned} ALG(A_2) &= \sum_{a \in A_2} \frac{\sum_{c \in C(a)} \alpha_c^a (B_c^a + \sum_{i \in A(a)} ctr_i^a v_{ic}^a)}{\sum_{c \in C(a)} \alpha_c^a} \\ &\geq \sum_{a \in A_2} \frac{\sum_{c \in C(a)} \alpha_c^a (B_c^a + \sum_{i \in O(a) \cap A(a)} ctr_i^a v_{ic}^a)}{\sum_{c \in C(a)} \alpha_c^a} \end{aligned} \tag{4}$$

– For a given bidder a, its contribution to global optimal allocation's welfare is no more than the sum over its global optimal allocation of any convex linear combination of its c-type values (over any subset of C including $C(a)$).

$$\begin{aligned} OPT(A_2) &= \sum_{a \in A_2} (B_{c'(a)}^a + \sum_i x_i^a ctr_i^a v_{ic}^a) = \sum_{a \in A_2} (B_{c'(a)}^a + \sum_{i \in O(a)} ctr_i^a v_{ic}^a) \\ &\leq \sum_{a \in A_2} \frac{\sum_{c \in C(a)} \alpha_c^a (B_c^a + \sum_{i \in O(a)} ctr_i^a v_{ic}^a)}{\sum_{c \in C(a)} \alpha_c^a} \end{aligned} \tag{5}$$

We now split the right hand side into two parts and then use all the other inequalities to get the final bound.

$$OPT(A_2) \leq \sum_{a \in A_2} \frac{\sum_{c \in C(a)} \alpha_c^a (B_c^a + \sum_{i \in O(a) \cap A(a)} ctr_i^a v_{ic}^a)}{\sum_{c \in C(a)} \alpha_c^a} + \sum_{a \in A_2} \frac{\sum_{c \in C(a)} \alpha_c^a \sum_{i \in O(a) - A(a)} ctr_i^a v_{ic}^a}{\sum_{c \in C(a)} \alpha_c^a}$$
$$\leq ALG(A_2) + ALG$$

Where the first inequality is due to 5 and second is due to 3 and 4. Summing $ALG(A_2)+ALG+ALG(A_1) >= OPT(A_1)+OPT(A_2)$ giving $2ALG >= OPT$.

5.3 Tight Example for Factor 2

Here we give an example showing that factor 2 is tight. We have two advertisers $A = \{a1, a2\}$ and two impressions $I = \{i1, i2\}$. ctr is 1 for all ad impression pairs. Value for advertiser $a1$ are $v_1^{a1} = \epsilon + \epsilon^2, v_2^{a1} = 1 - \epsilon$ and for second advertiser are $v_2^{a1} = 1, v_2^{a2} = 0$. We have one constraint (special case of TCPA constraint) with $B_c^{a1} = B_c^{a2} = 0$ and $v_{ic}^a = v_i^a$ for $i \in I, a \in A$.

One can show that $\alpha_c^1 = \epsilon, \alpha_c^2 = \frac{2}{\epsilon^2}$ is a locally optimal bidding strategy. This gives allocation of both $i1$ and $i2$ to $a1$ giving it liquid welfare of $\epsilon + \epsilon^2 + 1 - \epsilon = 1 + \epsilon^2$. But globally optimal solution allocates $c1$ to $a2$ and $c2$ to $a1$ giving it liquid welfare of $1 - \epsilon + 1 = 2 - \epsilon$.

6 Equilibrium

In this section we consider special case when the space of impressions/slots is a measure space. We further assume that there is no point mass distribution except a special impression,slot i, s for each advertiser a which has small ϵ positive value, 0 cost, $v_{ic}^a = \epsilon$ and always allocated to advertiser a. Then we show that there is an equilibrium bidding given by our bidding formula and no advertiser wants to deviate. We use the special impression to upper bound the dual variables. We use the no point mass distribution to make sure that slack in each constraint is a continuous function of the dual variables. Based on these two we can invoke Brower's fixed point theorem to show the existence of an equilibrium.

Lemma 2. *In any optimal solution to the dual linear program the dual variables* α_c^a *are bounded by* $\int_{i,s} ctr_{is}^a v_i^a d(i,s)/\epsilon$. *Further the slack in primal constraint i.e* $slack_c^a = B_c^a + \int_{i,s} x_{is}^a ctr_{is}^a (v_{ic}^a - cpc_{is}^a) d(i,s)$ *is a continuous function of the bidding formula* $\frac{v_i^a + \sum_c \alpha_c^a v_{ic}^a}{\sum_c \alpha_c^a}$.

Proof. We first note that slack in primal constraint being a continuous function of bidding formula is just a manifestation of assumption that we don't have any point mass distribution. Next to prove a upper bound on α_c^a we first prove an upper bound on each δ_i^a. Note that dual LP objective is lower bounded by δ_i^a and primal objective is upper bounded by $\int_{i,s} ctr_{is}^a v_{ik}^a d(i,s)$. Hence we get $\delta_i^a \leq \int_{i,s} ctr_{is}^a v_i^a d(i,s)$.

Now consider the dual constraint corresponding to the special impression, slot i, s for advertiser a. Then we know that $cpc_{is}^a = 0$ and $v_{ic}^a = \epsilon$. Then consider the corresponding dual constraint. $\delta_i^a + \sum_c \alpha_c^a (cpc_{is}^a - v_{ic}^a) \geq v_i^a$. Substituting the values and rewriting we get $\sum_c \alpha_c^a \epsilon \leq \delta_i^a$ which implies $\alpha_c^a \leq \delta_i^a / \epsilon$. Now using the upper bound on δ_i^a we get the upper bound on α_c^a.

We next define a map from α_c^a to itself. Define it as follows.

$$\phi(\alpha_c^a) = min\left(\frac{\int_{i,s} ctr_{is}^a v_{is}^a}{\epsilon}, \alpha_c^a (1+\eta)^{-slack_c^a}\right)$$

Here $0 < \eta < 1$ is any positive number. Since this map is continuous and bounded, by Brower's fixed point theorem we have a fixed point. At fixed point $\alpha_c^a > 0$ if and only if the constraint is tight. Hence by primal dual complementary slackness we have that the solution is also locally optimal for each advertiser.

References

1. Agrawal, S., Devanur, N.R.: Fast algorithms for online stochastic convex programming. In: Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2015, San Diego, CA, USA, 4–6 January 2015, pp. 1405–1424 (2015). https://doi.org/10.1137/1.9781611973730.93
2. Arora, S., Hazan, E., Kale, S.: The multiplicative weights update method: a meta-algorithm and applications. Theory Comput. **8**(1), 121–164 (2012)
3. Azar, Y., Feldman, M., Gravin, N., Roytman, A.: Liquid price of anarchy. In: Bilò, V., Flammini, M. (eds.) SAGT 2017. LNCS, vol. 10504, pp. 3–15. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-66700-3_1
4. Buchbinder, N., Jain, K., Naor, J.S.: 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). https://doi.org/10.1007/978-3-540-75520-3_24
5. Google Ads Help Center: About automated bidding. https://support.google.com/google-ads/answer/2979071. Accessed 12 Feb 2019
6. Google Ads Help Center: About universal app campaigns. https://support.google.com/google-ads/answer/6247380. Accessed Feb 13 2019
7. Charles, D.X., Chakrabarty, D., Chickering, M., Devanur, N.R., Wang, L.: Budget smoothing for internet ad auctions: a game theoretic approach. In: Proceedings of the Fourteenth ACM Conference on Electronic Commerce, EC 2013, Philadelphia, PA, USA, 16–20 June 2013, pp. 163–180 (2013)
8. Chen, P., et al.: Ad serving using a compact allocation plan. In: Proceedings of the 13th ACM Conference on Electronic Commerce, EC 2012, Valencia, Spain, 4–8 June 2012, pp. 319–336 (2012)

9. Chen, Y., Berkhin, P., Anderson, B., Devanur, N.R.: 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, San Diego, CA, USA, 21–24 August 2011, pp. 1307–1315 (2011)
10. Devanur, N.R., Jain, K., Sivan, B., Wilkens, C.A.: Near optimal online algorithms and fast approximation algorithms for resource allocation problems. J. ACM **66**(1), 7:1–7:41 (2019). https://doi.org/10.1145/3284177
11. Dobzinski, S., Leme, R.P.: Efficiency guarantees in auctions with budgets. In: Esparza, J., Fraigniaud, P., Husfeldt, T., Koutsoupias, E. (eds.) ICALP 2014, Part I. LNCS, vol. 8572, pp. 392–404. Springer, Heidelberg (2014). https://doi.org/10.1007/978-3-662-43948-7_33
12. Feldman, J., Korula, N., Mirrokni, V., Muthukrishnan, S., Pál, M.: Online ad assignment with free disposal. In: Leonardi, S. (ed.) WINE 2009. LNCS, vol. 5929, pp. 374–385. Springer, Heidelberg (2009). https://doi.org/10.1007/978-3-642-10841-9_34
13. Feldman, J., Muthukrishnan, S., Pál, M., Stein, C.: Budget optimization in search-based advertising auctions. In: Proceedings 8th ACM Conference on Electronic Commerce (EC-2007), San Diego, California, USA, 11–15 June 2007, pp. 40–49 (2007)
14. Ghosh, A., Rubinstein, B.I.P., Vassilvitskii, S., Zinkevich, M.: Adaptive bidding for display advertising. In: Proceedings of the 18th International Conference on World Wide Web, WWW 2009, Madrid, Spain, 20–24 April 2009, pp. 251–260 (2009)
15. Karande, C., Mehta, A., Srikant, R.: Optimizing budget constrained spend in search advertising. In: Sixth ACM International Conference on Web Search and Data Mining, WSDM 2013, Rome, Italy, 4–8 February 2013, pp. 697–706 (2013)
16. Mehta, A.: Online matching and ad allocation. Found. Trends® Theor. Comput. Sci. **8**(4), 265–368 (2013)
17. Mehta, A., Saberi, A., Vazirani, U., Vazirani, V.: Adwords and generalized online matching. J. ACM **54**(5), 22 (2007)
18. Plotkin, S.A., Shmoys, D.B., Tardos, É.: Fast approximation algorithms for fractional packing and covering problems. In: 32nd Annual Symposium on Foundations of Computer Science, San Juan, Puerto Rico, 1–4 October 1991, pp. 495–504 (1991)
19. Ren, K., Zhang, W., Chang, K., Rong, Y., Yu, Y., Wang, J.: Bidding machine: learning to bid for directly optimizing profits in display advertising. IEEE Trans. Knowl. Data Eng. **30**(4), 645–659 (2018)
20. Zhang, W., Yuan, S., Wang, J.: Optimal real-time bidding for display advertising. In: The 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD 2014, New York, NY, USA, 24–27 August 2014, pp. 1077–1086 (2014)