

## AdWords in a Panorama

Zhiyi Huang  
 Computer Science  
 The University of Hong Kong  
 Hong Kong, China  
 zhiyi@cs.hku.hk

Qiankun Zhang  
 Computer Science  
 The University of Hong Kong  
 Hong Kong, China  
 qkzhang@cs.hku.hk

Yuhao Zhang  
 Computer Science  
 The University of Hong Kong  
 Hong Kong, China  
 yhzhang2@cs.hku.hk

**Abstract**—Three decades ago, Karp, Vazirani, and Vazirani (STOC 1990) defined the online matching problem and gave an optimal  $(1-1/e)$ -competitive (about 0.632) algorithm. Fifteen years later, Mehta, Saberi, Vazirani, and Vazirani (FOCS 2005) introduced the first generalization called *AdWords* driven by online advertising and obtained the optimal  $(1-1/e)$  competitive ratio in the special case of *small bids*. It has been open ever since whether there is an algorithm for *general bids* better than the 0.5-competitive greedy algorithm. This paper presents a 0.5016-competitive algorithm for AdWords, answering this open question on the positive end. The algorithm builds on several ingredients, including a combination of the online primal dual framework and the configuration linear program of matching problems recently explored by Huang and Zhang (STOC 2020), a novel formulation of AdWords which we call the *panorama view*, and a generalization of the online correlated selection by Fahrback, Huang, Tao, and Zadimorghaddam (FOCS 2020) which we call the *panoramic online correlated selection*.

**Keywords**-Adwords, online matching, primal-dual

### I. INTRODUCTION

Consider an *ad platform* in online advertising, e.g., a search engine in the case of sponsored search. Each advertiser on the platform provides its *bids* for a set of keywords, for which it likes its ad to be shown. It further has a *budget* which upper bounds its payment in a day. When a user submits a request, often referred to as an *impression*, the platform sees the bids of the advertisers for it. The platform then selects an advertiser, who pays either its bid or its remaining budget, whichever is smaller. The goal of the platform is to allocate impressions to advertisers to maximize the total payment of the advertisers. The revenues of online advertising in the US have surpassed those of television advertising in 2016 [5], and have totaled \$57.9 billions in the first half of 2019 [6].

While online advertising acquires growing importance in practice, it has also been extensively studied in theoretical computer science. The first related research dates back to three decades ago, when Karp et al. [30] introduced the online matching problem and designed an online algorithm with the optimal  $1 - \frac{1}{e} \approx 0.632$  competitive ratio. It can be viewed as the special case

with unit bids and unit budgets. Fifteen years later, Mehta et al. [36] formally formulated it as the *AdWords* problem. They introduced an optimal  $1 - \frac{1}{e}$ -competitive algorithm under the *small-bid assumption*: an advertiser's bid for any impression is much smaller than its budget.

Subsequently, AdWords has been studied under stochastic assumptions. Goel and Mehta [16] showed that *assuming a random arrival order of the impressions* and small bids, a  $1 - \frac{1}{e}$  competitive ratio can be achieved using the greedy algorithm: allocate each impression to the advertiser who would make the largest payment. Mirrokni et al. [38] analyzed the algorithm of Mehta et al. [36] in the more restricted *unknown iid model*, and obtained an improved ratio of 0.76 for small bids. Finally, Devanur et al. [8] proved that the greedy algorithm is  $1 - \frac{1}{e}$ -competitive for general bids in the unknown iid model. For small bids, they proposed a  $(1 - \epsilon)$ -competitive algorithm. We refer readers to the survey by Mehta [34] for further references.

Little is known, however, about the most general case of AdWords, i.e., with general bids and without stochastic assumptions. On the positive end, we only have the greedy algorithm and the trivial 0.5 competitive ratio. On the negative end, there is no provable evidence that the optimal  $1 - \frac{1}{e}$  competitive ratio of online matching cannot be achieved in AdWords. It has been open since Mehta et al. [36] whether there is an online algorithm that achieves a competitive ratio strictly better than 0.5.

#### A. Our Contributions and Techniques

The main result of the paper is the first online algorithm for AdWords that breaks the 0.5 barrier.

**Theorem 1.** *There is a 0.5016-competitive algorithm for AdWords.*

We develop the algorithm under the online primal dual framework. In a nutshell, by considering an appropriate linear program (LP) of the problem, the online primal dual framework designs the online algorithm according to the optimality conditions of LPs, and uses

the objective of the dual LP as the benchmark in the analysis. Buchbinder et al. [4] applied it to AdWords with small bids, using the standard matching LP, to obtain an alternative analysis of the  $1 - \frac{1}{e}$  competitive algorithm by Mehta et al. [36]. Later, Devanur and Jain [7] and Devanur et al. [9] found further applications of the framework in other online matching problems. Recently, Huang and Zhang [20] demonstrated an advantage of using the configuration LP instead of the standard matching LP in online matching with stochastic rewards. The current paper also builds on the strength of the configuration LP, echoing the message of Huang and Zhang [20]. See Section II for details.

Our second ingredient is a novel formulation of AdWords which we call the *panorama view*. Recall that an advertiser's payment in the original formulation is either the sum of its bids for the assigned impressions or its budget, whichever is smaller. The panorama view further associates each advertiser with an interval whose length equals the budget, and requires the algorithms to assign each impression to not only an advertiser, but further a subset of its interval with size equal to the bid. For example, consider an impression  $i$  and an advertiser  $a$  whose budget is 2 and whose bid for  $i$  is 1. The panorama view associates advertiser  $a$  with an interval  $[0, 2)$ . Further, when an algorithm assigns  $i$  to  $a$ , it must further assign  $i$  to a subset of size at most 1, e.g.,  $[0.5, 1.5)$ . Define an advertiser's payment in the panorama view to be the size of the union of the assigned subsets, which lower bounds the payment in the original formulation. The panorama view allows a fine-grained characterization on how the assignment of an impression  $i$  to an advertiser  $a$  affects the marginal gains of the other impressions assigned to  $a$ . Concretely, suppose we shortlist two advertisers for each impression, and then assign it to one of them with a fresh random bit. In the original formulation, having advertiser  $a$  in impression  $i$ 's shortlist decreases the marginal gain of *all* other impressions that shortlist  $a$  in a *complicated manner*. In the panorama view, however, it decreases the marginal gain *only for those whose assigned subsets intersect with  $i$ 's*; more precisely, it *decreases the contribution of the intersection by half*. See Section III for a formal definition of the panorama view and some examples.

Finally, instead of using a fresh random bit to select a shortlisted advertiser for each impression, our algorithm selects one with negative correlation. If a previous impression which shortlists advertiser  $a$  with an overlapping subset does *not* select  $a$ , the current one will be more likely to select  $a$ . Given the same shortlists, negatively correlated selections get larger expected gains in the panorama view than independent selections. An algorithmic ingredient called *online correlated se-*

*lection* (OCS) by Huang and Tao [18, 19] provides a quantitative control of such negative correlation in the special case when bids equal budgets. The final piece of our algorithm is a generalization of OCS which applies to the general case of AdWords in the panorama view. We refer to it as the panoramic OCS (PanOCS). Section II includes a formal definition of OCS, Section IV-A defines the PanOCS and sketches the main ideas behind it, and the details are provided in the full version [26].

Building on these ingredients, we get a 0.50005-competitive online primal dual algorithm for AdWords in Section IV, weaker than the ratio in Theorem 1 yet breaking the 0.5 barrier nonetheless. To obtain the final ratio, we observe that the above algorithm works better for larger bids while the algorithm of Mehta et al. [36] is better for smaller bids. Hence, we design a 0.5016-competitive hybrid algorithm by unifying both approaches under the online primal dual framework in the full version [26] and analyze the algorithm of Mehta et al. [36] for small bids using online primal dual and configuration LP, which may serve as a warmup for readers unfamiliar with the framework.

## B. Other Related Works

AdWords is closely related to the literature of online matching started by Karp et al. [30]. Aggarwal et al. [1] studied the vertex-weighted problem and obtained the optimal  $1 - \frac{1}{e}$  competitive ratio with a generalization of the algorithm by Karp et al. [30]. Feldman et al. [13] investigated edge-weighted online matching in the free-disposal model, where the algorithm may dispose a previous matched edge for free to make room for a new one. They called it the *display ads* problem, and achieved the optimal  $1 - \frac{1}{e}$  competitive ratio assuming large capacities, i.e., each offline vertex can be matched to a large number of online vertices. The analysis was simplified by Devanur et al. [10] under the online primal dual framework. Further, Fehrbach et al. [11, 12, 18, 19] obtained a better than 0.5-competitive edge-weighted algorithm without assuming large capacities. In doing so, they introduced the OCS which directly inspired this paper. Finally, there are generalized models which allow all vertices to be online and even consider general graphs [3, 15, 21, 23, 24, 25, 39].

Online matching problems are also widely investigated under different stochastic assumptions. First, consider random arrivals of online vertices. Karande et al. [29] and Mahdian and Yan [32] showed that the algorithm of Karp et al. [30] is strictly better than  $1 - \frac{1}{e}$ -competitive in this model. Huang et al. [22] gave a better than  $1 - \frac{1}{e}$ -competitive algorithm for the vertex-weighted problem. Kesselheim et al. [31] showed that the greedy algorithm is  $\frac{1}{e}$ -competitive for

the edge-weighted problem even without free-disposal. Under the stronger assumption that online vertices are drawn iid from an unknown distribution, Kapralov et al. [28] proved that greedy is  $1 - \frac{1}{e}$ -competitive for a more general problem called online submodular welfare maximization which captures both the edge-weighted problem with free-disposal and AdWords as special cases. Further assuming that the distribution is known leads to better competitive ratios [14, 17, 27, 33]. We leave for future research if the algorithm in this paper is better than  $1 - \frac{1}{e}$ -competitive under random arrivals.

Finally, Mehta and Panigrahi [35] proposed online matching with stochastic rewards, where an edge chosen by the algorithm is successfully matched only with some probability. They focused on the special case of equal success probabilities and gave algorithms that are 0.567-competitive if the success probability is vanishing, and better than 0.5-competitive in general. Later, Mehta et al. [37] showed a 0.534-competitive algorithm for vanishing unequal success probabilities. Recently, Huang and Zhang [20] improved the competitive ratios to 0.576 and 0.572 for vanishing equal and unequal success probabilities respectively. In doing so, they showed an advantage of the configuration LP over the standard matching LP under online primal dual. This paper echoes the above message.

## II. PRELIMINARIES

Consider a bipartite graph  $G = (A, I, E)$ , where  $A$  and  $I$  are sets of vertices corresponding to the advertisers and impressions in AdWords respectively, and  $E \subseteq A \times I$  is the set of edges between them. Further, each edge  $(a, i)$  is associated with a non-negative real number  $b_{ai}$  which represents advertiser  $a$ 's bid for impression  $i$ .<sup>1</sup> By allowing zero bids, we may assume without loss of generality (wlog) that  $G$  is a complete bipartite graph, i.e.,  $E = A \times I$ . Finally, each advertiser  $a$  is associated with a positive budget  $B_a$  which upper bounds the payment of the advertiser. Concretely, assigning a subset of impressions  $S \subseteq I$  to an advertiser  $a$  leads to a budget-additive payment:

$$b_a(S) \stackrel{\text{def}}{=} \min \left\{ \sum_{i \in S} b_{ai}, B_a \right\}.$$

By this definition, we may assume wlog that  $b_{ai} \leq B_a$  for any advertiser  $a$  and any impression  $i$ .

The advertisers are given upfront, while the impressions arrive one at a time. We write  $i < i'$  if an impression  $i$  arrives before another impression  $i'$ . On the arrival of an impression, the algorithm must immediately and irrevocably assign it to an advertiser.

<sup>1</sup>AdWords as an online algorithm problem does not consider the strategic behaviors of the advertisers. We merely inherit the term *bid* from the original paper of Mehta et al. [36].

The objective is to maximize the sum of the above payments from all advertisers. Following the standard competitive analysis of online algorithms, an algorithm is  $\Gamma$ -competitive for some *competitive ratio*  $0 \leq \Gamma \leq 1$  if its expected objective is at least  $\Gamma$  times the offline optimal in hindsight for any AdWords instance.

*Configuration Linear Program.*: The algorithms in this paper and their analyses rely on the LP relaxations of the problem. Instead of the standard matching LP, this paper considers the more expressive advertiser-side configuration LP and its dual:

$$\begin{aligned} \max \quad & \sum_{a \in A} \sum_{S \subseteq I} b_a(S) x_{aS} \\ \text{s.t.} \quad & \sum_{S \subseteq I} x_{aS} \leq 1 \quad \forall a \in A \\ & \sum_{a \in A} \sum_{S \ni i} x_{aS} \leq 1 \quad \forall i \in I \\ & x_{aS} \geq 0 \quad \forall a \in A, \forall S \subseteq I \\ \min \quad & \sum_{a \in A} \alpha_a + \sum_{i \in I} \beta_i \\ \text{s.t.} \quad & \alpha_a + \sum_{i \in S} \beta_i \geq b_a(S) \quad \forall a \in A, \forall S \subseteq I \\ & \alpha_a \geq 0 \quad \forall a \in A \\ & \beta_i \geq 0 \quad \forall i \in I \end{aligned}$$

Let  $P$  and  $D$  denote the objectives of the primal and dual LPs respectively. Throughout the paper we will always let  $x_{aS}$  be the probability that  $S$  is the subset of impressions assigned to advertiser  $a$ . Then, the primal objective  $P$  equals the objective of the algorithm.

*Online Primal Dual Framework.*: We build on the online primal dual framework which uses the dual objective as an upper bound of the offline optimal in the competitive analyses of online algorithms. In particular, this paper applies it to the configuration LP of AdWords.

**Lemma 2.** *Suppose an online algorithm is coupled with a dual algorithm which maintains a dual assignment such that for some  $0 \leq \Gamma \leq 1$ :*

- 1) *Approximate dual feasibility:*  $\alpha_a + \sum_{i \in S} \beta_i \geq \Gamma \cdot b_a(S)$  for any  $a \in A$  and any  $S \subseteq I$ .
- 2) *Reverse weak duality:*  $P \geq D$ ;

*Then, it is  $\Gamma$ -competitive.*

*Proof:* By the first condition, scaling the dual assignment by a factor of  $\Gamma^{-1}$  makes it feasible while changing the dual objective by the same factor. Therefore, by weak duality of LPs, the offline optimal is at most  $\Gamma^{-1}D$ . Putting together with the second condition proves the lemma. ■

*Online Correlated Selection.*: The algorithms in this paper further utilize a recent algorithmic ingredient called *online correlated selection* (OCS) by Huang and Tao [18, 19]. Consider a set of ground elements, and further a sequence of pairs of these elements arriving one at a time. Suppose we randomly select one element from each pair with a fresh random bit. Then, an element will be selected at least once with probability  $1 - 2^{-k}$  after appearing in  $k$  pairs. The OCS correlates the randomness to achieve better efficiency. We state below a simplified definition, removing some aspects irrelevant to AdWords.

**Definition 1.** For any  $0 \leq \gamma \leq 1$ , a  $\gamma$ -OCS is an online algorithm ensuring that any element that appears in  $k$  pairs is selected at least once with probability at least:

$$1 - 2^{-k}(1 - \gamma)^{\max\{k-1, 0\}} .$$

### III. PANORAMA VIEW

The algorithms in this paper are based on a novel viewpoint of the AdWords problem which we call the *panorama view*. Recall that the payment of an advertiser  $a$  is budget-additive in AdWords: assigning a subset of impressions  $S$  to an advertiser  $a$  gives  $b_a(S) = \min \{ \sum_{i \in S} b_{ai}, B_a \}$ . Let  $\mu(\cdot)$  denote the Lebesgue measure. In the panorama view, we further associate each advertiser  $a$  with an interval  $[0, B_a)$ ; each impression  $i$  assigned to  $a$  is further assigned to a subset  $Y_{ai} \subseteq [0, B_a)$  whose Lebesgue measure  $\mu(Y_{ai})$  is at most  $b_{ai}$ . In fact, we will always choose  $Y_{ai}$  to be a finite union of disjoint left-closed, right-open intervals, for which the Lebesgue measure is simply the sum of their lengths. Further define the payment of an advertiser  $a$  in the panorama view as:

$$\mu(\cup_{i \in S} Y_{ai}) .$$

Correspondingly, the objective in the panorama view is the sum of the above payment from all advertisers. Importantly, it lower bounds the original objective of AdWords.

**Lemma 3.** For any advertiser  $a$ , any subset of impressions  $S$  assigned to  $a$ , and any subsets  $Y_{ai} \subseteq [0, B_a)$  with Lebesgue measure at most  $b_{ai}$  for impressions  $i \in S$ , we have:

$$\mu(\cup_{i \in S} Y_{ai}) \leq \min \left\{ \sum_{i \in S} b_{ai}, B_a \right\} .$$

*Proof:* On the one hand, by subadditivity of the Lebesgue measure function  $\mu$ , and by the measure upper bounds of the subsets  $Y_{ai}$ 's, we have  $\mu(\cup_{i \in S} Y_{ai}) \leq \sum_{i \in S} \mu(Y_{ai}) \leq \sum_{i \in S} b_{ai}$ . On the other hand, because  $Y_{ai}$ 's are subsets of  $[0, B_a)$ , we have  $\mu(\cup_{i \in S} Y_{ai}) \leq \mu([0, B_a)) = B_a$ . ■

*Example 1 (Deterministic Algorithms).*: Consider an arbitrary deterministic algorithm. Then, whenever it assigns an impression  $i$  to an advertiser  $a$ , we may wlog further assign it to the leftmost unassigned interval. For instance, suppose impressions 1, 2, 3, and so on are assigned to advertiser  $a$  in this order; we may further assign 1 to  $[0, b_{a1})$ , 2 to  $[b_{a1}, b_{a1} + b_{a2})$ , 3 to  $[b_{a1} + b_{a2}, b_{a1} + b_{a2} + b_{a3})$ , and so forth. In doing so, the objectives in the panorama view is identical to the original one.

*Oblivious Semi-randomized Algorithms.*: The panorama view of AdWords separates itself from the original one when it comes to a special family of randomized algorithms which we call the *oblivious semi-randomized algorithms*. They are semi-randomized in that for every impressions  $i$ , they either assign it deterministically to an advertiser-subset combination, or choose two advertiser-subset combinations and assign it to one of them with equal marginal probability. We shall refer to the former as a *deterministic round* and the latter as a *randomized round*. If an impression  $i$  corresponds to a randomized round, we say that it is *semi-assigned* to the advertiser-subset combinations. For the time being, readers may think of using a fresh random bit in every randomized round for a concrete understanding of the panorama view, although our algorithms will correlate the decisions in different rounds negatively.

Further, the algorithms are oblivious: neither the decisions of deterministic versus randomized rounds, nor the choices of advertiser-subset combinations depend on the realization of random bits in previous rounds. Hence, the semi-assignments to the same advertiser may have overlapping subsets, and the objective in the panorama view no longer equals the original one in general.

*Example 2 (Oblivious Semi-randomized Algorithms).*: Let there be two advertisers whose budgets equal 2, and three impressions for which both advertisers bid 1. Further suppose that we select with a fresh random bit for each impression. In the original budget-additive payments, with probability  $\frac{1}{4}$  all impressions are assigned to the same advertiser and thus the objective equals 2; otherwise, the objective equals 3. Hence, the expected objective equals  $\frac{1}{4} \cdot 2 + \frac{3}{4} \cdot 3 = \frac{11}{4}$ . In the panorama view, however, the algorithm must further assign each impression to a subset of  $[0, 2)$  for both advertisers. It is wlog to assign the first impression to  $[0, 1)$ , contributing 1 to the objective. Further, it is reasonable to further assign the second impression to  $[1, 2)$  so that it is disjoint with the first one and contributes 1 to the objective. However, the third impression only contributes  $\frac{1}{2}$  to the objective regardless of the choices of subsets because the entire interval  $[0, 2)$  has been semi-assigned once. Therefore, the expected objective is only  $1 + 1 + \frac{1}{2} = \frac{5}{2}$ .

It may seem odd to restrict ourselves to oblivious algorithms. Would it not be better if we first check the realized assignments in earlier rounds and then pick an advertiser-subset combination disjoint with the previous ones? In a nutshell, we focus on oblivious algorithms to separate the algorithmic component for choosing assignments and semi-assignments, and that for correlating the decisions in different randomized rounds. Importantly, we can achieve negative correlation: a semi-assignment is more likely to get selected if an earlier overlapping one is not. By contrast, screening the options based on the realization of earlier random bits could lead to positive correlations (e.g., weighted sampling without replacements [2]). That said, there may be AdWords algorithms with controlled positive correlations which are better than 0.5-competitive. The study of such algorithms, however, is beyond the scope of this paper and is left for future research.

*Bookkeeping at the Point-level.*: It is more convenient to account for the primal objective at the point-level as follows. We say that a point  $y \in [0, B_a)$  of an advertiser  $a$  is assigned if there is an impression  $i$  assigned to  $a$  and a subset  $Y_{ai}$  containing  $y$ , either due to a deterministic round, or due to a semi-assignment in a randomized round which selected  $a$ . For randomized algorithms, let  $0 \leq x_a(y) \leq 1$  denote the probability that  $y$  is assigned. Then, the primal objective equals:

$$P = \sum_{a \in A} \int_0^{B_a} x_a(y) dy . \quad (1)$$

Similarly, we say that  $y$  is semi-assigned whenever an impression is semi-assigned to  $a$  and a subset containing  $y$ . Let  $k_a(y)$  denote the number of times that  $y$  is semi-assigned. Further define  $k_a(y) = \infty$  if  $y$  has been assigned in a deterministic round, driven by the fact that semi-assignments on their own take finitely many rounds to make a point  $y$  assigned with certainty.

We further introduce point-level dual variables  $\alpha_a(y)$ ,  $a \in A$ ,  $y \in [0, B_a)$ , and let:

$$\alpha_a = \int_0^{B_a} \alpha_a(y) dy . \quad (2)$$

Then, approximate dual feasibility becomes:

$$\int_0^{B_a} \alpha_a(y) dy + \sum_{i \in S} \beta_i \geq \Gamma \cdot b_a(S) \quad (3)$$

*Panoramic Interval-level Assignments.*: We first introduce some notations which are useful throughout the paper. For any point  $y \in [0, B_a)$ , any subset  $Y \subseteq [0, B_a)$ , and any  $0 \leq b \leq B_a$ , let  $y \oplus_Y b$  denote the point in  $[0, B_a)$  such that the interval  $[y, y \oplus_Y b)$  excluding  $Y$  has Lebesgue measure  $b$ . Here, we abuse notation and allow  $y \oplus_Y b$  to be smaller than  $y$ , in

which case  $[y, y \oplus_Y b)$  denotes the union of  $[y, B_a)$  and  $[0, y \oplus_Y b)$ . In the boundary case when the subset  $[0, B_a) \setminus Y$  has a measure strictly less than  $b$ , i.e.,  $B_a - \mu(Y) < b$ , define  $y \oplus_Y b = y$ . Further define the reverse operation  $y \ominus_Y b$  such that  $[y \ominus_Y b, y) \setminus Y$  has measure  $b$ .

For any advertiser  $a$  and the set of impressions assigned and semi-assigned to it, the algorithm will further select subsets of  $[0, B_a)$  greedily as follows. Maintain a point  $y^*$  initially at 0 which represents the start of the next subset. For each impression  $i$ , further assign or semi-assign it to  $[y^*, y^* \oplus_{Y_D} b_{ai}) \setminus Y_D$ , where  $Y_D$  is the subset of  $[0, B_a)$  that has already been assigned deterministically. Further update  $y^* = y^* \oplus_{Y_D} b_{ai}$ . To this end, think of  $[0, B_a)$  as a circle by gluing its endpoints; the algorithm scans along the circle to find a subset with measure  $b_{ai}$  that has not been deterministically assigned. It is similar to taking a panorama and hence the name of the alternative view of AdWords.

The panoramic interval-level assignments equalize the numbers of times the points  $y \in [0, B_a)$  are semi-assigned, among those that have not been deterministically assigned. We omit the proof since it follows by the definition of the algorithm. See Figure 1 for an illustrative example.

**Lemma 4.** *For any  $a \in A$  and any  $y \in [0, B_a)$ ,  $k_a(y)$  equals either (1)  $k_{\min} = \min_{z \in [0, B_a)} k_a(z)$ , or (2)  $k_{\min} + 1$ , or (3)  $\infty$ . Further, the first kind satisfies  $y \geq y^*$ , and the second kind satisfies  $y < y^*$ .*

#### IV. BASIC ALGORITHM

This section presents an oblivious semi-randomized algorithm that is better than 0.5-competitive for AdWords. Section IV-A presents a brief introduction to an algorithmic ingredient called *panoramic online correlated selection* (PanOCS), which correlates the randomized decisions in different rounds negatively. Section IV-B then demonstrates an online algorithm powered by PanOCS, and Section IV-C analyzes it under the online primal dual framework. Finally, Section IV-D optimizes the parameters of the algorithm to achieve a 0.5041 competitive ratio in the crux of AdWords, i.e., when all nonzero bids are large,  $\frac{1}{2}B_a \leq b_{ai} \leq B_a$ , and a smaller 0.50005 competitive ratio in the general case. The latter is smaller than the ratio in the main theorem but breaks the 0.5 barrier nonetheless.

Formal descriptions of the PanOCS algorithms and their analyses may be of independent interest and are therefore deferred to the full version [26]. The 0.5016 ratio in the main theorem requires a hybrid approach which treats large and small bids differently, which is also presented in the full version [26]




Figure 1: example of interval-level assignment

### A. Panoramic Online Correlated Selection at a Glimpse

Recall the oblivious semi-randomized algorithms. In each randomized round, such an algorithm chooses a pair of advertiser-subset combinations, oblivious to the random bits in previous rounds. Then, the combinations are passed on to an algorithmic component which selects one of them with equal marginal probability, and *correlates across different randomized rounds negatively*. We call it the PanOCS since it is a generalization of the OCS [18, 19] in the panorama view of AdWords.

For a formal definition, recall that  $x_a(y)$  is the probability a point  $y \in [0, B_a)$  of an advertiser  $a$  has been assigned, and  $k_a(y)$  is the number of randomized rounds in which  $y$  is semi-assigned.

**Definition 2.** A PanOCS is an online algorithm which takes a sequence of pairs of advertiser-subset combinations as input, and for each pair selects one combination. It is a  $\gamma$ -PanOCS for some  $0 \leq \gamma \leq 1$  if for any advertiser  $a$ , and any point  $y \in [0, B_a)$ , we have:

$$x_a(y) \geq 1 - 2^{-k_a(y)}(1 - \gamma)^{\max\{k_a(y)-1, 0\}}. \quad (4)$$

Observe that using an independent random bit in every randomized round is a 0-PanOCS, since the probability of being assigned after  $k$  semi-assignments with independent random bits is precisely  $1 - 2^{-k}$ . The parameter  $\gamma$  quantifies the advantage over independent random bits.

The intuition behind the inequality is best explained with a thought experiment. Suppose that whenever  $y$  is semi-assigned other than the first time, there is a  $\gamma$  chance to be perfectly negatively correlated with the last semi-assignment of  $y$ :  $a$  is chosen this time if it is not chosen last time, and vice versa. Further suppose that the above events are negatively dependent for the  $k_a(y) - 1$  different pairs of adjacent semi-assignments of  $y$ . Then,  $y$  is never assigned only if none of the events happens, whose probability is at most  $(1 - \gamma)^{k_a(y)-1}$ , and further when none of the  $k_a(y)$

independent selections picks  $a$ , which equals  $2^{-k_a(y)}$ . Our analysis will substantiate this intuition.

To see the connection with OCS, consider a special case when the bids equal the budgets, i.e.,  $b_{ai} = B_a$ . Then, since  $x_a(y)$  and  $k_a(y)$  are independent of  $y$ , it suffices to consider if advertiser  $a$  has been assigned. As a result, the above definition coincides with the definition of OCS, taking the advertisers as the ground elements. The extra challenge of PanOCS is to ensure the inequality simultaneously for all points  $y \in [0, B_a)$  when the bids are arbitrary.

**Theorem 5.** *Suppose all nonzero bids are large, i.e.,  $\frac{1}{2}B_a \leq b_{ai} \leq B_a$  or  $b_{ai} = 0$  for any advertiser  $a \in A$  and any impression  $i \in I$ . Then, there is a 0.05144-PanOCS.*

**Theorem 6.** *Suppose the algorithm makes at most  $k_{\max}$  semi-assignments to any point  $y \in [0, B_a)$  of any advertiser  $a \in A$ . Then, there is a  $0.01245 \cdot k_{\max}^{-1}$ -PanOCS.*

The proofs of the theorems are deferred to the full version [26]. We include below a proof sketch of a weaker  $\frac{1}{64}$ -PanOCS for large bids to foreshadow the arguments in our PanOCS analyses.

*Proof Sketch of a Weaker Theorem 5 ( $\gamma = \frac{1}{64}$ ):* For any impression  $i$  semi-assigned to an advertiser  $a$ , we write  $(i + j)_a$  to denote the  $j$ -th impression semi-assigned to  $a$  after impression  $i$  (or the  $(-j)$ -th impression before  $i$ , if  $j < 0$ ).

We next explain the algorithm. Consider an impression  $i$  in a randomized round. Suppose it is semi-assigned to advertisers  $a_1$  and  $a_2$ . Then, sample  $a^* \in \{a_1, a_2\}$  and  $j \in \{-2, -1, 1, 2\}$  uniformly at random. If  $j > 0$ , select  $a_1$  or  $a_2$  and the corresponding subsets with a fresh random bit; further, pass the result to future impression  $(i + j)_{a^*}$ . If  $j < 0$ , check if the past impression  $(i + j)_{a^*}$  passes its result to  $i$ . If so, makes the opposite choice: select  $a^*$  and the corresponding subset if  $a^*$  was *not* selected in round  $(i + j)_{a^*}$ , and vice versa. Otherwise, select with a fresh random bit.

Finally, we argue this is a  $\frac{1}{64}$ -PanOCS. Fix any advertiser  $a$  and any point  $y \in [0, B_a)$ . Suppose  $i_1 < i_2 < \dots < i_k$  are the impressions semi-assigned to  $a$  and subsets containing  $y$ . Consider any neighboring  $i_\ell$  and  $i_{\ell+1}$ . We now use the assumption of large bids to get that  $i_{\ell+1}$  is the first or second impression semi-assigned to advertiser  $a$  after  $i_\ell$ . Hence,  $i_\ell$  may pass its result to  $i_{\ell+1}$ , which may then make the opposite choice. When it happens,  $y$  is assigned exactly once in rounds  $i_\ell$  and  $i_{\ell+1}$ . More precisely, this is when  $i_\ell$  samples  $a^* = a$  and  $j > 0$  so that  $(i_\ell + j)_a = i_{\ell+1}$ , and  $i_{\ell+1}$  samples  $a^* = a$  and  $j < 0$  so that  $(i_{\ell+1} + j)_a = i_\ell$ , which happens with probability  $\frac{1}{64}$ . Moreover, we claim that the events are negatively dependent for  $k - 1$  neighboring pairs of  $i_\ell$  and  $i_{\ell+1}$ . Hence, the probability of having no such pair is at most  $(1 - \frac{1}{64})^{k-1}$ . Finally, even if  $i_1 < i_2 < \dots < i_k$  are independent, the probability that  $a$  is never selected is only  $2^{-k}$ . Putting together,  $y$  has been assigned in at least one of these  $k$  rounds with probability no less than  $1 - 2^{-k}(1 - \frac{1}{64})^{k-1}$ . ■

### B. Online Primal Dual Algorithm

We demonstrate an online primal dual Algorithm 1, taking a  $\gamma$ -PanOCS as a blackbox. Recall that an oblivious semi-randomized algorithm either deterministically assigns  $i$  to an advertiser-subset combination, or semi-assigns it to two combinations in a randomized round. In the latter case, let the  $\gamma$ -PanOCS select a combination. Let  $\bar{x}_a(y)$  be the lower bound of  $x_a(y)$  given by a  $\gamma$ -PanOCS in Eqn. (4), and let  $\bar{P}$  be the corresponding lower bound of the primal objective in Eqn. (1), i.e.:

$$\bar{x}_a(y) \stackrel{\text{def}}{=} 1 - 2^{-k_a(y)} \cdot (1 - \gamma)^{\max\{k_a(y)-1, 0\}},$$

$$\bar{P} \stackrel{\text{def}}{=} \sum_{a \in A} \int_0^{B_a} \bar{x}_a(y) dy.$$

Let  $\Delta x$  denote the increment of  $\bar{x}_a(y)$  as  $k_a(y)$  increases, i.e.:

$$\Delta x(k) \stackrel{\text{def}}{=} \begin{cases} 2^{-1} & k = 1; \\ 2^{-k}(1 - \gamma)^{k-2}(1 + \gamma) & k \geq 2. \end{cases} \quad (5)$$

An online primal dual algorithm's decision for each impression is driven by maximizing  $\beta_i$ . For each advertiser  $a$ , compute two quantities  $\Delta_a^D \beta_i$  and  $\Delta_a^R \beta_i$  which we shall detail shortly. The former denotes how much  $\beta_i$  would gain if  $i$  is assigned to  $a$ . The latter denotes how much  $\beta_i$  would gain if  $i$  is semi-assigned to  $a$ . The corresponding subsets are decided by the panoramic interval-level assignment in Section III. Then, find advertisers  $a_1$  and  $a_2$  with the largest  $\Delta_a^R \beta_i$ , and advertiser  $a^*$  with the largest  $\Delta_a^D \beta_i$ . If  $\Delta_{a_1}^R \beta_i + \Delta_{a_2}^R \beta_i$  is greater than  $\Delta_{a^*}^D \beta_i$ , semi-assign  $i$  to  $a_1$  and  $a_2$  in a randomized round. Otherwise, assign  $i$  to  $a^*$  deterministically.

Next, we define  $\Delta_a^D \beta_i$  and  $\Delta_a^R \beta_i$  from two invariants below. First, let the lower bound of primal equal the dual, i.e.,  $\bar{P} = D$ . It ensures reverse weak duality in Lemma 2 because  $P \geq \bar{P} = D$ . Second, recall that  $\alpha_a(y)$ 's account for dual variable  $\alpha_a$  at the point-level as explained in Eqn. (2). For a set of parameters  $\Delta\alpha(\ell)$ ,  $\ell \geq 1$ , which will be optimized in the analysis, let:

$$\alpha_a(y) = \sum_{\ell=1}^{k_a(y)} \Delta\alpha(\ell). \quad (6)$$

We first derive  $\Delta_a^R \beta_i$  from the invariants. Suppose  $i$  is semi-assigned to  $a$  and a subset  $Y_{ai}$ . For any point  $y \in Y_{ai}$ , the primal increment due to  $y$  is  $\Delta x(k_a(y) + 1)$ , where  $k_a(y)$  denotes the value before the semi-assignment. The dual increment in  $\alpha_a(y)$  is  $\Delta\alpha(k_a(y) + 1)$  by the second invariant. Finally, by the first invariant, the increment in  $\beta_i$  due to point  $y \in Y_{ai}$  shall equal the difference between  $\Delta x(k_a(y) + 1)$  and  $\Delta\alpha(k_a(y) + 1)$ . For convenience of notations, define:

$$\Delta\beta(k) \stackrel{\text{def}}{=} \Delta x(k) - \Delta\alpha(k). \quad (7)$$

Our choice of  $\Delta\alpha(k)$  will ensure non-negativity of  $\Delta\beta(k)$ . Putting together we get that:

$$\Delta_a^R \beta_i \stackrel{\text{def}}{=} \int_{Y_{ai}} \Delta\beta(k_a(y) + 1) dy. \quad (8)$$

Similarly, suppose  $i$  is assigned deterministically to  $a$  and a subset  $Y_{ai}$ . For any point  $y \in Y_{ai}$ , the primal increment due to  $y$  is  $\sum_{\ell > k_a(y)} \Delta x(\ell)$  since  $k_a(y)$  becomes  $\infty$ ; the dual increment in  $\alpha_a(y)$  is  $\sum_{\ell > k_a(y)} \Delta\alpha(\ell)$  by the second invariant. Thus, together with the first invariant, we let:

$$\Delta_a^D \beta_i \stackrel{\text{def}}{=} \int_{Y_{ai}} \sum_{\ell > k_a(y)} \Delta\beta(\ell) dy. \quad (9)$$

### C. Online Primal Dual Analysis

Recall that reverse weak duality always holds because of the first invariant. Next, we derive a set of conditions on the parameters which imply approximate dual feasibility. The conditions will be numbered. Then, we optimize the competitive ratio  $\Gamma$  and  $\Delta\alpha(k)$ 's through an LP. For any advertiser  $a$  and any subset of impressions  $S \subseteq I$ , recall approximate dual feasibility in Eqn. (3):

$$\int_0^{B_a} \alpha_a(y) dy + \sum_{i \in S} \beta_i \geq b_a(S) \cdot \Gamma.$$

1) *Warmup*: First consider a special case when  $S$  has only one impression  $i$  who bids  $b_{ai} = B_a$ . This warmup case is simple enough to be analyzed in around one page, yet is also general enough to derive the binding conditions on the parameters that are still sufficient in

---

**Algorithm 1** Basic Online Primal Dual Algorithm  
 (Parameterized by  $\Delta\alpha(k)$ ,  $k \geq 1$ )
 

---

**state variables:**

$k_a(y) \geq 0$ , number of times  $y$  is semi-assigned;  
 $k_a(y) = \infty$  if  $y$  is assigned in a deterministic round  
**for all** impression  $i$  **do**

**for all** advertiser  $a \in A$  **do**

computer subset  $Y_{ai} \subseteq [0, B_a)$  using panoramic  
 interval-level assignment (Section III)

compute  $\Delta_a^R \beta_i$  and  $\Delta_a^D \beta_i$  according to Equations (7), (8), and (9)

**end for**

find  $a_1, a_2$  that maximize  $\Delta_a^R \beta_i$ , and  $a^*$  that  
 maximizes  $\Delta_a^D \beta_i$

**if**  $\Delta_{a_1}^R \beta_i + \Delta_{a_2}^R \beta_i \geq \Delta_{a^*}^D \beta_i$  **# randomized**

assign  $i$  to what PanOCS selects between  $a_1$  and  
 $a_2$  and the corresponding subsets

**else** ( $\Delta_{a_1}^R \beta_i + \Delta_{a_2}^R \beta_i < \Delta_{a^*}^D \beta_i$ ) **# deterministic**

assign  $i$  to  $a^*$  and the corresponding subset

**endif**
**end for**


---

the general case. The full analysis of the general case is deferred to the full version [26].

We will divide into four subcases, depending on whether impression  $i$  is assigned to the advertiser  $a$ , and whether  $i$  is a deterministic or randomized round. In each case, we will lower bound both  $\alpha_a(y)$ 's and  $\beta_i$  as functions of the  $k_a(y)$ 's. To avoid ambiguity, let  $k_a(y)$  be the final value at the end, and  $k_a^i(y)$  be the value right before the arrival of impression  $i$ .

(Case 1) Round of  $i$  is randomized, and  $i$  is not semi-assigned to  $a$ : Both  $a_1$  and  $a_2$  chosen by the algorithm give at least  $\Delta_a^R \beta_i$  to  $\beta_i$ . By the definition of  $\Delta_a^R \beta_i$  in Eqn. (8) and the invariant about  $\alpha_a(y)$  in Eqn. (6), approximate dual feasibility reduces to:

$$\int_0^{B_a} \sum_{\ell=1}^{k_a(y)} \Delta\alpha(\ell) dy + 2 \int_0^{B_a} \Delta\beta(k_a^i(y)+1) dy \geq \Gamma \cdot B_a .$$

Since  $k_a(y) \geq k_a^i(y)$ , and the first term is increasing in  $k_a(y)$ 's, it suffices to prove it when  $k_a(y)$  equals  $k_a^i(y)$ . We ensure it pointwise for every  $y \in [0, B_a)$ :

$$\forall k \geq 0 : \quad \sum_{\ell=1}^k \Delta\alpha(\ell) + 2 \cdot \Delta\beta(k+1) \geq \Gamma . \quad (10)$$

(Case 2) Round of  $i$  is deterministic, and  $i$  is not assigned to  $a$ : We reduce to the previous case by introducing a condition about the superiority of randomized rounds.

$$\forall k \geq 1 : \quad \Delta\beta(k) \geq \sum_{\ell=k+1}^{\infty} \Delta\beta(\ell) . \quad (11)$$

**Lemma 7.** Assuming Eqn. (11), for any advertiser  $a$  and any impression  $i$ ,  $2 \cdot \Delta_a^R \beta_i \geq \Delta_a^D \beta_i$ .

We remark that the lemma holds in the general case as well. Adding  $\Delta\beta(k)$  to both sides of Eqn. (11) gives  $2 \cdot \Delta\beta(k) \geq \sum_{\ell=k}^{\infty} \Delta\beta(\ell)$ . It then follows by the definition of  $\Delta_a^R \beta_i$  and  $\Delta_a^D \beta_i$  in Equations (8) and (9). Intuitively, it means that a randomized round with two equally good advertisers in terms of  $\Delta_a^R \beta_i$  is better than a deterministic round with only one of them.

By the definition of the algorithm, the advertiser  $a^*$  to which the algorithm deterministically assigns  $i$  satisfies  $\Delta_{a^*}^D \beta_i \geq \Delta_{a^*}^R \beta_i + \Delta_a^R \beta_i$ . Further by  $2 \cdot \Delta_{a^*}^R \beta_i \geq \Delta_{a^*}^D \beta_i$  because of Lemma 7, we have  $\Delta_{a^*}^R \beta_i \geq \Delta_a^R \beta_i$ . Thus, we get  $\beta_i = \Delta_{a^*}^D \beta_i \geq \Delta_a^R \beta_i + \Delta_{a^*}^R \beta_i \geq 2 \cdot \Delta_a^R \beta_i$ . The rest is verbatim.

(Case 3) Round of  $i$  is randomized, and  $i$  is semi-assigned to  $a$ : Since the algorithm does not deterministically assign  $i$  to  $a$ ,  $\beta_i \geq \Delta_a^D \beta_i$ . By the definition of  $\Delta_a^D \beta_i$  in Eqn. (9) and the invariant about  $\alpha_a(y)$ 's in Eqn. (6), approximate dual feasibility reduces to:

$$\int_0^{B_a} \sum_{\ell=1}^{k_a(y)} \Delta\alpha(\ell) dy + \int_0^{B_a} \sum_{\ell > k_a^i(y)} \Delta\beta(\ell) dy \geq \Gamma \cdot B_a .$$

Importantly, since  $i$  is semi-assigned to  $a$ , we have  $k_a(y) \geq k_a^i(y) + 1$ . By contrast, the previous cases only have  $k_a(y) \geq k_a^i(y)$ . It suffices to ensure the inequality pointwise when  $k_a(y) = k_a^i(y) + 1$ :

$$\forall k \geq 1 : \quad \sum_{\ell=1}^k \Delta\alpha(\ell) + \sum_{\ell=k}^{\infty} \Delta\beta(\ell) \geq \Gamma . \quad (12)$$

(Case 4) Round of  $i$  is deterministic, and  $i$  is assigned to  $a$ : For any  $y \in [0, B_a)$ ,  $k_a(y) = \infty$  after round  $i$  because  $b_{ai} = B_a$  in the warmup case. Since  $k_a(y)$  may already be large before  $i$  arrives, we do not have any nontrivial lower bound of  $\beta_i$ . Hence,  $\alpha_a(y)$ 's on their own must satisfy approximate dual feasibility. By the invariant of  $\alpha_a(y)$ 's in Eqn. (6), it reduces to:

$$\sum_{\ell=1}^{\infty} \Delta\alpha(\ell) \geq \Gamma . \quad (13)$$

*Optimizing the competitive ratio.*: We shall solve an LP, whose variables are the competitive ratio  $\Gamma$  and parameters  $\Delta\alpha(k)$ 's and  $\Delta\beta(k)$ 's, and whose constraints are the first invariant about  $\Delta\alpha(k)$ 's and  $\Delta\beta(k)$ 's in Eqn. (7) and the sufficient conditions for approximate dual feasibility in Equations (10) to (13).

$$\begin{aligned} \max \quad & \Gamma \\ \text{s.t.} \quad & \text{Eqn. (7), (10), (11), (12), and (13);} \\ & \Delta\alpha(k), \Delta\beta(k) \geq 0 \quad \forall k \geq 1 \end{aligned}$$



2) *General Case*: The proof of the approximate dual feasibility in the general case is presented in the full version [26], which is under the invariant in Eqn. (7), and the conditions derived from the warmup, i.e., Equations (10) to (13). To simplify the argument, we further assume *monotonicity of  $\Delta\beta$* :

$$\forall k \geq 1 : \quad \Delta\beta(k) \geq \Delta\beta(k+1) . \quad (14)$$

We remark that it would be satisfied automatically by the solution of the LP even if it was not stated explicitly. The LP becomes:

$$\begin{aligned} \max \quad & \Gamma \\ \text{s.t.} \quad & \text{Eqn. (7), (10), (11), (12), (13), and (14); (15)} \\ & \Delta\alpha(k), \Delta\beta(k) \geq 0 \quad \forall k \geq 1 \end{aligned}$$

**Lemma 8.** *Suppose  $\Gamma$ ,  $\Delta\alpha(k)$ 's, and  $\Delta\beta(k)$ 's form a solution of the LP in Eqn. (15). Then, Algorithm 1 satisfies approximate dual feasibility, as restated below:*

$$\int_0^{B_a} \alpha_a(y) dy + \sum_{i \in S} \beta_i \geq b_a(S) \cdot \Gamma .$$

We defer the proof of Lemma 8 to the full version [26]. Recall that reverse weak duality always holds. By Lemma 2, Lemma 8 leads to the next corollary.

**Corollary 9.** *Algorithm 1 is  $\Gamma$ -competitive for any solution of the LP in Eqn. (15).*

#### D. Optimizing the Parameters

It remains to optimize the parameters  $\Delta\alpha(k)$ 's and  $\Delta\beta(k)$ 's and the competitive ratio  $\Gamma$  by solving the LP in Eqn. (15). Observe that the LP has countably infinitely many variables and constraints and therefore, cannot be directly solved with an LP solver. One possible strategy is to solve a finite restriction by setting  $\Delta\alpha(k) = \Delta\beta(k) = 0$  for  $k > k_{\max}$  with some sufficiently large  $k_{\max}$ . This is indeed the strategy we use for the hybrid algorithm in the full version. Fortunately, the LP in Eqn. (15) admits benign structures. As a result, we can provide an explicit solution.

**Lemma 10.** *For any  $0 \leq \gamma \leq 1$ , the following is a solution of the LP in Eqn. (15):*

$$\begin{aligned} \Gamma &= \frac{3+2\gamma}{6+3\gamma} . \\ \Delta\alpha(k) &= \begin{cases} \frac{3+\gamma}{6+3\gamma} \Delta x(1) & k=1 ; \\ \frac{1+\gamma}{2+\gamma} \Delta x(k) & k \geq 2 . \end{cases} \quad (16) \\ \Delta\beta(k) &= \begin{cases} \frac{3+2\gamma}{6+3\gamma} \Delta x(1) & k=1 ; \\ \frac{1}{2+\gamma} \Delta x(k) & k \geq 2 . \end{cases} \end{aligned}$$

The proof is deferred to the full version [26].

*Large Bids: the Crux of AdWords*: In light of the positive results by Mehta et al. [36] for small bids that are at most half the budgets, the case of large bids, i.e.,  $\frac{1}{2}B_a < b_{ai} \leq B_a$ , can be viewed as the crux of the AdWords problem. As a direct corollary of Lemma 10 and the 0.05144-PanOCS for large bids in Theorem 5, we get the first online algorithm that breaks the 0.5 barrier in the crux.

**Theorem 11.** *If all nonzero bids are large, i.e.,  $\frac{1}{2}B_a < b_{ai} \leq B_a$  or  $b_{ai} = 0$  for any advertiser  $a \in A$  and any impression  $i \in I$ , Algorithm 1 with the  $\gamma = 0.05144$ -PanOCS in Theorem 5 is  $\Gamma$ -competitive for:*

$$\Gamma = \frac{3+2\gamma}{6+3\gamma} > 0.5041 .$$

#### General Bids: a Weaker Version of Theorem 1:

Next, consider a restricted version of Algorithm 1 such that for any advertiser  $a$  and any point  $y \in [0, B_a)$ ,  $y$  is semi-assigned at most  $k_{\max}$  times for some positive integer  $k_{\max}$ . This can be achieved by letting  $\Delta x(k) = \Delta\alpha(k) = \Delta\beta(k) = 0$  for any  $k > k_{\max}$ . The restriction allows us to use the PanOCS for general bids in Theorem 6. We need, however, a solution to the LP in Eqn. (15) under the restriction. A natural choice is adopting the solution in Lemma 10 directly for  $k \leq k_{\max}$ , and decreasing the competitive ratio  $\Gamma$  accordingly to preserve feasibility.

**Lemma 12.** *For any  $0 \leq \gamma \leq 1$ , the following is a solution of the LP in Eqn. (15):*

$$\begin{aligned} \Gamma &= \frac{3+2\gamma}{6+3\gamma} - 2^{-k_{\max}}(1-\gamma)^{k_{\max}-1} . \\ \Delta\alpha(k) &= \begin{cases} \frac{3+\gamma}{6+3\gamma} \cdot \Delta x(1) & k=1 ; \\ \frac{1+\gamma}{2+\gamma} \cdot \Delta x(k) & 2 \leq k \leq k_{\max} ; \\ 0 & k > k_{\max} . \end{cases} \quad (17) \\ \Delta\beta(k) &= \begin{cases} \frac{3+2\gamma}{6+3\gamma} \cdot \Delta x(1) & k=1 ; \\ \frac{1}{2+\gamma} \cdot \Delta x(k) & 2 \leq k \leq k_{\max} ; \\ 0 & k > k_{\max} . \end{cases} \end{aligned}$$

*Proof*: It follows by Lemma 10 that the contributions of  $\Delta\alpha(k)$ 's and  $\Delta\beta(k)$ 's,  $k > k_{\max}$ , to the approximate dual feasibility constraints, i.e., Eqn. (10), (12), and (13), are at most:

$$\sum_{\ell=k_{\max}+1}^{\infty} \Delta x(\ell) = 2^{-k_{\max}}(1-\gamma)^{k_{\max}-1} .$$

Hence, decreasing the competitive ratio  $\Gamma$  by this amount restores approximate dual feasibility even after setting  $\Delta\alpha(k) = \Delta\beta(k) = 0$  for  $k > k_{\max}$ .

Finally, the other constraints, i.e., Eqn. (7), (11), and (14), are trivially preserved after letting  $\Delta x(k) = \Delta \alpha(k) = \Delta \beta(k) = 0$  for  $k > k_{\max}$ . ■

Consider the  $\gamma$ -PanOCS in Theorem 6 where  $\gamma = 0.01245 \cdot k_{\max}^{-1}$ . The competitive ratio from the above solution is:

$$\begin{aligned} \Gamma &= \frac{3 + 2\gamma}{6 + 3\gamma} - 2^{-k_{\max}}(1 - \gamma)^{k_{\max}-1} \\ &= \frac{1}{2} + \Omega(k_{\max}^{-1}) - 2^{-k_{\max}}(1 - \gamma)^{k_{\max}-1}. \end{aligned}$$

The second term is inverse proportional to  $k_{\max}$  while the third term decreases exponentially in  $k_{\max}$ . Hence, by choosing a sufficiently large  $k_{\max}$ , the competitive ratio is strictly larger than half. Indeed, letting  $k_{\max} = 18$  gives  $\Gamma > 0.50005$ .

#### ACKNOWLEDGMENT

Zhiyi Huang was supported in part by an RGC grant 17203717E.

#### REFERENCES

- [1] Gagan Aggarwal, Gagan Goel, Chinmay Karande, and Aranyak Mehta. Online vertex-weighted bipartite matching and single-bid budgeted allocations. In *Proceedings of the 22nd Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1253–1264. SIAM, 2011.
- [2] Kenneth S Alexander. A counterexample to a correlation inequality in finite sampling. *The Annals of Statistics*, pages 436–439, 1989.
- [3] Itai Ashlagi, Maximilien Burq, Chinmoy Dutta, Patrick Jaillet, Amin Saberi, and Chris Sholley. Edge weighted online windowed matching. In *Proceedings of the 2019 ACM Conference on Economics and Computation*, pages 729–742, 2019.
- [4] Niv Buchbinder, Kamal Jain, and Joseph Seffi Naor. Online primal-dual algorithms for maximizing ad-auctions revenue. In *Proceedings of the 15th Annual European Symposium on Algorithms*, pages 253–264. Springer, 2007.
- [5] Internet Advertising Bureau. Iab internet advertising revenue report: 2016 full year results, 2017.
- [6] Internet Advertising Bureau. Iab internet advertising revenue report: 2019 first six months results, 2019.
- [7] Nikhil R Devanur and Kamal Jain. Online matching with concave returns. In *Proceedings of the 44th Annual ACM Symposium on Theory of Computing*, pages 137–144, 2012.
- [8] Nikhil R Devanur, Kamal Jain, Balasubramanian Sivan, and Christopher A Wilkens. Near optimal online algorithms and fast approximation algorithms for resource allocation problems. In *Proceedings of the 12th ACM conference on Electronic commerce*, pages 29–38, 2011.
- [9] Nikhil R Devanur, Kamal Jain, and Robert D Kleinberg. Randomized primal-dual analysis of ranking for online bipartite matching. In *Proceedings of the 24th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 101–107. SIAM, 2013.
- [10] Nikhil R Devanur, Zhiyi Huang, Nitish Korula, Vahab S Mirrokni, and Qiqi Yan. Whole-page optimization and submodular welfare maximization with online bidders. *ACM Transactions on Economics and Computation (TEAC)*, 4(3):1–20, 2016.
- [11] Matthew Fahrbach and Morteza Zadimoghaddam. Online weighted matching: breaking the  $\frac{1}{2}$  barrier. *arXiv*, pages arXiv-1704, 2017.
- [12] Matthew Fahrbach, Zhiyi Huang, Runzhou Tao, and Morteza Zadimoghaddam. Edge-weighted online bipartite matching. In *Proceedings of the 61st Annual IEEE Symposium on Foundations of Computer Science*. IEEE, 2020.
- [13] Jon Feldman, Nitish Korula, Vahab Mirrokni, Shanmugavelayutham Muthukrishnan, and Martin Pál. Online ad assignment with free disposal. In *Proceedings of the 5th International Workshop on Internet and Network Economics*, pages 374–385. Springer, 2009.
- [14] Jon Feldman, Aranyak Mehta, Vahab Mirrokni, and Shan Muthukrishnan. Online stochastic matching: beating  $1 - \frac{1}{e}$ . In *Proceedings of the 50th Annual IEEE Symposium on Foundations of Computer Science*, pages 117–126. IEEE, 2009.
- [15] Buddhima Gamlath, Michael Kapralov, Andreas Maggiori, Ola Svensson, and David Wajc. Online matching with general arrivals. In *Proceedings of the 60th Annual IEEE Symposium on Foundations of Computer Science*, pages 26–37. IEEE, 2019.
- [16] Gagan Goel and Aranyak Mehta. Online budgeted matching in random input models with applications to AdWords. In *Proceedings of the 19th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 982–991, 2008.
- [17] Bernhard Haeupler, Vahab S Mirrokni, and Morteza Zadimoghaddam. Online stochastic weighted matching: Improved approximation algorithms. In *International Workshop on Internet and Network Economics*, pages 170–181. Springer, 2011.
- [18] Zhiyi Huang. Understanding Zadimoghaddam’s edge-weighted online matching algorithm: weighted case. *arXiv preprint arXiv:1910.03287*, 2019.
- [19] Zhiyi Huang and Runzhou Tao. Understanding Zadimoghaddam’s edge-weighted online matching algorithm: unweighted case. *arXiv preprint arXiv:1910.02569*, 2019.

- [20] Zhiyi Huang and Qiankun Zhang. Online primal dual meets online matching with stochastic rewards: configuration LP to the rescue. In *Proceedings of the 52nd ACM Symposium on Theory of Computing*, 2020.
- [21] Zhiyi Huang, Ning Kang, Zhihao Gavin Tang, Xiaowei Wu, Yuhao Zhang, and Xue Zhu. How to match when all vertices arrive online. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*, pages 17–29, 2018.
- [22] Zhiyi Huang, Zhihao Gavin Tang, Xiaowei Wu, and Yuhao Zhang. Online vertex-weighted bipartite matching: beating  $1 - \frac{1}{e}$  with random arrivals. In *Proceedings of the 45th International Colloquium on Automata, Languages, and Programming*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2018.
- [23] Zhiyi Huang, Binghui Peng, Zhihao Gavin Tang, Runzhou Tao, Xiaowei Wu, and Yuhao Zhang. Tight competitive ratios of classic matching algorithms in the fully online model. In *Proceedings of the 30th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2875–2886. SIAM, 2019.
- [24] Zhiyi Huang, Ning Kang, Zhihao Gavin Tang, Xiaowei Wu, Yuhao Zhang, and Xue Zhu. Fully online matching. *Journal of the ACM*, 67(3):1–25, 2020.
- [25] Zhiyi Huang, Zhihao Gavin Tang, Xiaowei Wu, and Yuhao Zhang. Fully online matching ii: beating ranking and water-filling. In *Proceedings of the 61st Annual IEEE Symposium on Foundations of Computer Science*. IEEE, 2020.
- [26] Zhiyi Huang, Qiankun Zhang, and Yuhao Zhang. Adwords in a panorama, 2020.
- [27] Patrick Jaillet and Xin Lu. Online stochastic matching: New algorithms with better bounds. *Mathematics of Operations Research*, 39(3):624–646, 2014.
- [28] Michael Kapralov, Ian Post, and Jan Vondrák. Online submodular welfare maximization: Greedy is optimal. In *Proceedings of the 24th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1216–1225. SIAM, 2013.
- [29] Chinmay Karande, Aranyak Mehta, and Pushkar Tripathi. Online bipartite matching with unknown distributions. In *Proceedings of the 43rd Annual ACM Symposium on Theory of Computing*, pages 587–596, 2011.
- [30] Richard M Karp, Umesh V Vazirani, and Vijay V Vazirani. An optimal algorithm for on-line bipartite matching. In *Proceedings of the 22nd Annual ACM Symposium on Theory of Computing*, pages 352–358, 1990.
- [31] Thomas Kesselheim, Klaus Radke, Andreas Tönnis, and Berthold Vöcking. An optimal online algorithm for weighted bipartite matching and extensions to combinatorial auctions. In *In Proceedings of the 21th European Symposium on Algorithms*, pages 589–600. Springer, 2013.
- [32] Mohammad Mahdian and Qiqi Yan. Online bipartite matching with random arrivals: an approach based on strongly factor-revealing lps. In *Proceedings of the 43rd Annual ACM Symposium on Theory of Computing*, pages 597–606, 2011.
- [33] Vahideh H Manshadi, Shayan Oveis Gharan, and Amin Saberi. Online stochastic matching: Online actions based on offline statistics. *Mathematics of Operations Research*, 37(4):559–573, 2012.
- [34] Aranyak Mehta. Online matching and ad allocation. *Foundations and Trends in Theoretical Computer Science*, 8(4):265–368, 2013.
- [35] Aranyak Mehta and Debmalya Panigrahi. Online matching with stochastic rewards. In *Proceedings of the 53rd IEEE Annual Symposium on Foundations of Computer Science*, pages 728–737. IEEE, 2012.
- [36] Aranyak Mehta, Amin Saberi, Umesh Vazirani, and Vijay Vazirani. Adwords and generalized online matching. *Journal of the ACM*, 54(5):22:1–22:19, 2007.
- [37] Aranyak Mehta, Bo Waggoner, and Morteza Zadimoghaddam. Online stochastic matching with unequal probabilities. In *Proceedings of the 26th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1388–1404. SIAM, 2014.
- [38] Vahab S Mirrokni, Shayan Oveis Gharan, and Morteza Zadimoghaddam. Simultaneous approximations for adversarial and stochastic online budgeted allocation. In *Proceedings of the 23rd Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1690–1701. SIAM, 2012.
- [39] Yajun Wang and Sam Chiu-wai Wong. Two-sided online bipartite matching and vertex cover: Beating the greedy algorithm. In *Proceedings of the 42nd International Colloquium on Automata, Languages, and Programming*, pages 1070–1081. Springer, 2015.