

Near-Optimum Online Ad Allocation for Targeted Advertising

Joseph (Seffi) Naor^{*1} and David Wajc^{†2}

¹CS Dept., Technion , naor@cs.technion.ac.il

²CS Dept., Carnegie Mellon University , dwajc@cs.cmu.edu

Abstract

Motivated by Internet targeted advertising, we address several ad allocation problems. Prior work has established these problems admit no randomized online algorithm better than $(1 - \frac{1}{e})$ -competitive (Karp et al. [14], Mehta et al. [19]), yet simple heuristics have been observed to perform much better in practice. We explain this phenomenon by studying a generalization of the bounded-degree inputs considered by Buchbinder et al. [4], graphs which we call (k, d) -bounded. In such graphs the maximal degree on the online side is at most d and the minimal degree on the offline side is at least k . We prove that for such graphs, these problems' natural greedy algorithms attain competitive ratio $1 - \frac{d-1}{k+d-1}$, tending to one as d/k tends to zero. We prove this bound is tight for these algorithms.

Next, we develop deterministic primal-dual algorithms for the above problems achieving competitive ratio $1 - (1 - \frac{1}{d})^k > 1 - \frac{1}{e^{k/d}}$, or *exponentially* better loss as a function of k/d , and strictly better than $1 - \frac{1}{e}$ whenever $k \geq d$. We complement our lower bounds with matching upper bounds for the vertex-weighted problem. Finally, we use our deterministic algorithms to prove by dual-fitting that simple randomized algorithms achieve the same bounds in expectation. Our algorithms and analysis differ from previous ad allocation algorithms, which largely scale bids based on the spent fraction of their bidder's budget, whereas we scale bids according to the number of times the bidder could have spent as much as her current bid. Our algorithms differ from previous online primal-dual algorithms, as they do not maintain dual feasibility, but only primal-to-dual ratio, and only attain dual feasibility upon termination. We believe our techniques could find applications to other well-behaved online packing problems.

Categories and Subject Descriptors: F.2.2 Nonnumerical Algorithms and Problems.

Keywords: Online Matching, Online Ad Allocation, Targeted Advertising, Sponsored search.

1 Introduction

Internet advertising is ubiquitous. Forecast to surpass the 50 billion dollar/year mark in 2015 in the United States alone, it has become, to a large extent, the driving economic force behind much of the content of the world wide web. How is this advertising space sold and bought? Most ads fall either under sponsored search or targeted advertising, both of which are sold in what constitute instances of the online ad allocation problem.

In online ad allocation, we are faced with the following problem: advertisers announce to an advertising platform (e.g. Yahoo, Google, Microsoft) what their advertising budgets are, and their bids for an ad to be displayed to every kind of user. The user "type" is determined, for example, by search terms searched, in the case of sponsored search, or user-demographics, in the case of targeted advertising. When a user visits a web-page with an ad slot managed by the ad platform, the latter needs to decide immediately and irrevocably which (if any) of the advertisers' ads to display to the user. The advertising platform's goal is to maximize its revenues, despite

^{*}Supported in part by United States-Israel BSF Grant No. 2010-246 and ISF Grant No. 954/11.

[†]Part of this research conducted while the author was working for Yahoo! Labs, Israel.

uncertainty concerning future page-views. This problem can be formulated as a generalization of online bipartite matching, with advertisers as the offline vertices and ad slots as online vertices.¹ See Section 2 for a formal definition of this and other problems we consider.

The theoretical interest in online allocations can be traced back to 1990, when Karp et al. [14] considered the fundamental problem of bipartite maximum matching in an online setting. In their seminal paper, Karp et al. proved that randomized online algorithms cannot in general achieve competitive ratio above $1 - \frac{1}{e} \approx 0.632$, and presented the RANKING algorithm, which matches this upper bound and is thus optimal.²

The online maximum matching problem was generalized, first by Kalyanasundaram and Pruhs [12], and later by Aggarwal et al. [1], who presented algorithms achieving optimal $1 - \frac{1}{e}$ competitive ratio for the b -matching and vertex-weighted matching problems, respectively. The AdWords problem, first proposed by Mehta et al. [19], is the more general ad allocation problem, but subject to the realistic *small bid assumption*, i.e. assuming every advertiser i has budget B_i much larger than its bids b_{ij} . (This assumption is necessary to achieve non-trivial results. See 6.4). For this problem too the natural greedy algorithm has competitive ratio $\frac{1}{2}$. Mehta et al. gave an algorithm for this problem with competitive ratio $1 - \frac{1}{e}$. Buchbinder et al. [4] achieved the same results using an online primal-dual approach. See Mehta [18] for an in-depth survey of prior art and techniques used to tackle these problems.

We will address the problems discussed above, but first, we start with motivation.

1.1 Motivation

As is to be expected of a problem for which a loss of $1/e \approx 36.7\%$ can translate itself to billions of dollars in potential revenue lost yearly, researchers have studied weaker models than the adversarial model for the ad allocation problem, in the hope that these may permit better guarantees. (See 1.4.) In this paper we revisit the stronger adversarial model, for graphs with structural characteristics met by many ad allocation instances arising from targeted advertising. Specifically, we assume advertisers are interested in a large number of ad slots (at least k), and that every ad slot is of interest to a relatively small number of advertisers (at most d). As with the small bid assumption $B_i \gg b_{ij}$ for the AdWords problem, assumption of the above structure is not only useful in order to obtain better bounds (as we will show), but also constitutes a reasonable assumption for targeted advertising, for the following twin reasons:

online side: advertisers typically target their advertising campaigns at *specific* segments of the population (e.g. young Californians who ski often); while these segments may be large in absolute terms, they are mostly small in relative terms (e.g., less than four percent of Californians ski often). Consequently, users tend to belong to relatively few segments. Coupled with the fact that the number of active campaigns at any given time is limited, this implies a restricted pool of ads that might be displayed to any particular user, justifying the small degree assumption for ad slots.

offline side: advertisers typically target *large* segments of the population (as in the example above), while not allocating a budget high enough to display ads to all users in a segment. Coupled with the fact that every page-view of a particular targeted user corresponds to a vertex in the graph, this implies the high degree assumption on the offline side, and more generally for the ad allocation problem, the assumption that $\sum_{i,j} b_{ij} \geq k \cdot B_i$ for some large k .

We call the graphs displaying these characteristics (k, d) -bounded graphs.

¹In this paper, without loss of generality, we assume advertisers only pay for impressions, and not e.g. clicks.

²The original proof of RANKING's competitive ratio was found to contain a mistake nearly twenty years later by Krohn and Varadarajan, but the algorithm's performance has since been re-proven by Birnbaum and Mathieu [3], Goel and Mehta [9], and recently by Devanur et al. [6].

Definition 1.1 ((k, d) -bounded graphs). We say a bipartite graph $G = (L, R, E)$ is (k, d) -bounded if every left vertex $i \in L$ has degree $d(i) \geq k$ and every right vertex $j \in R$ has degree $d(j) \leq d$. For ad allocations, we replace $d(i) \geq k$ with the property $\sum_j b_{ij} \geq k \cdot B_i$.

We concern ourselves with such graphs with k large and d small. For brevity’s sake, as all graphs in this paper will be bipartite, we refrain from stating the fact explicitly, and refer to (k, d) -bounded graphs as (k, d) -graphs henceforth. As the problems studied in this paper are all maximization problems, we adopt the convention that a lower bound indicates a positive result and an upper bound indicates a negative result.

1.2 Our Results

By focusing on (k, d) -graphs, we justify the observed success of greedy algorithms “in the wild”, and propose algorithms that are exponentially better, and provably optimal under these structural assumptions. Finally, we leverage our deterministic algorithms to prove simple randomized algorithms achieve the same bounds in expectation. Our results hold for the maximum matching, vertex-weighted matching and AdWords problems (with the exception of the matching upper bound for the latter). Table 1 delineates our results for these problems on (k, d) -graphs. We obtain similar results for the general ad allocation problem, even with large-ish bids (see 1.5.)

Table 1: Best results for general and (k, d) -graphs

Algorithms	General Graphs	(k, d) -Graphs
Greedy	$\frac{1}{2}$ (Tight)	$1 - \frac{d-1}{k+d-1}$ (Tight)
	Folklore	This work
Deterministic	$\frac{1}{2}$ (Tight)	$1 - (1 - \frac{1}{d})^k$ (Tight)
	Folklore	This work
Randomized	$1 - \frac{1}{e}$ (Tight)* [14, 3, 9, 1, 19, 4, 6]	$1 - (1 - \frac{1}{d})^k$ This work

* can be achieved *deterministically* for AdWords.

We begin by explaining the empirical success of greedy algorithms for the above problems, proving their loss is proportional to the ratio of the maximal degree in the online side to the minimal degree in the offline side; i.e., their competitive ratio tends to *one* as this ratio tends to zero. We complement this lower bound with a family of examples for which these algorithms do no better.

Theorem 1.2. Greedy algorithms achieve a competitive ratio of $\frac{k}{k+d-1}$ on (k, d) -bounded graphs. This analysis is tight for all $k \geq d - 1$.

We improve on the above, designing deterministic algorithms with *exponentially* smaller loss. We prove this is optimal for deterministic algorithms.

Theorem 1.3. There exist deterministic online algorithms for the unweighted and vertex-weighted matching problems with competitive ratio $1 - (1 - \frac{1}{d})^k > 1 - (\frac{1}{e})^{k/d}$ on (k, d) -bounded graphs. Moreover, these algorithms gain at least a $1 - (1 - \frac{1}{d})^k$ fraction of the total sum of weights. This is optimal whenever $k \geq d$.

Corollary 1.4. (Structural Corollary) For every bipartite graph G with the minimal degree of its left side at least $\ln c$ times larger than the maximal degree of its right side, G has a matching with at least a $(1 - \frac{1}{c})$ -fraction of G ’s left side matched.

In stating our bounds for general ad allocation, we follow the notation of Buchbinder et al. [4] and denote the maximum bid-to-budget ratio by $R_{\max} = \max_{(i,j) \in E} \left\{ \frac{b_{ij}}{B_i} \right\}$.

Theorem 1.5. *There exists a deterministic algorithm which gains total revenue at least*

$$\left(\sum_{i \in L} B_i \right) \cdot \left((1 - R_{\max}) \cdot \left(1 - \left(1 - \frac{1}{d} \right)^k \right) \right)$$

for ad allocation on (k, d) -graphs with $k \geq d - 1$, and is thus $((1 - R_{\max}) \cdot (1 - (1 - \frac{1}{d})^k))$ -competitive. This is optimal – no deterministic algorithm can do better for $k \geq d$.

To contrast our results with the state-of-the-art, we note that the algorithms of Mehta et al. [19], Buchbinder et al. [4], Devanur et al. [6] achieve competitive ratio $(1 - R_{\max}) \cdot (1 - 1/(1 + R_{\max})^{1/R_{\max}})$. This bound tends to $1 - \frac{1}{e}$ from below as R_{\max} tends to zero, but is far from this value for larger R_{\max} . Our algorithms fare better whenever $k \geq d$ even for large-ish R_{\max} . As stated in 1.1, we expect k to be significantly larger than d , but in order to emphasize the strength of our bound, let us assume only that $d/k = R_{\max}$. Table 2 displays the resulting competitive ratios in this case. Note that in this regime our algorithm is already better at $R_{\max} = \frac{1}{3}$ than prior algorithms are at the limit (i.e. when $R_{\max} \rightarrow 0$).

Table 2: Results for Ad Allocation with large-ish bids in (k, d) -graphs with $d/k = R_{\max}$

R_{\max}	$\frac{1}{2}$	$\frac{1}{3}$	$\frac{1}{4}$	$\frac{1}{5}$	$\frac{1}{6}$	$\frac{1}{8}$	$\frac{1}{16}$	$\frac{1}{32}$	$\frac{1}{100}$	$\rightarrow 0$
State-of-the-art	0.278	0.385	0.443	0.478	0.503	0.534	0.582	0.607	0.624	0.632
Our Work	0.432	0.633	0.736	0.795	0.831	0.875	0.938	0.969	0.99	1

Better still, our algorithms are robust to a few outlying advertisers increasing R_{\max} , as the $(\sum_i B_i) \cdot (1 - R_{\max})$ term in the above bound is rather $(\sum_i B_i - \max_{j \in N(i)} b_{ij})$. This is the first such result in the adversarial setting. To the best of our knowledge only the algorithm of Devanur et al. [7] for the iid model holds this desired property. Likewise, our algorithms are robust to few outlying advertisers making the input not (k, d) -bounded (alternatively, increasing k), as the following theorem asserts.

Theorem 1.6 (Outliers). *If every advertiser i satisfies $\sum_j b_{ij} \geq k \cdot B_i$, except for a subset $S \subset L$ with total budget at most an α -fraction of the total sum of budgets, $\sum_{i \in S} B_i \leq \alpha \cdot \sum_{i \in L} B_i$, then the algorithms of Theorems 1.3 and 1.5 gain revenue at least $(1 - \alpha)$ times the bounds guaranteed by the above theorems. In particular, these algorithms achieve competitive ratio at least $(1 - \alpha) \cdot (1 - (1 - \frac{1}{d})^k)$ and $(1 - \alpha) \cdot ((1 - R_{\max}) \cdot (1 - (1 - \frac{1}{d})^k))$.*

Finally, we prove that several easy-to-implement randomized algorithms match the bounds of our optimal deterministic algorithms in expectation, despite making no use of the input's structure.

Theorem 1.7. *Several simple randomized algorithms achieve expected competitive ratio matching those of Theorems 1.3, 1.5, and 1.6.*

1.3 Techniques

As many previous ad allocation algorithms, our algorithms can be seen as bid-scaling algorithms. That is, matches are chosen greedily based on the bids b_{ij} of each advertiser i , times a scaling factor. However, contrary to previous algorithms Mehta et al. [19], Buchbinder et al. [4], Devanur et al. [6] that scale bids according to the function $1 - e^{f(i)-1}$, where $f(i)$ is the fraction of i 's budget spent so far, our algorithms essentially scale bids according to an exponential in u , the

number of *unused opportunities* for spending the current bid b_{ij} , specifically, $\left(\frac{d}{d-1}\right)^u$. Other differences can be seen in our algorithms’ primal-dual interpretation: we make no use of the ad slots’ dual variables, and only update the dual variables of each arriving ad slot’s neighbors. Interestingly, our online primal-dual algorithms do not guarantee dual feasibility throughout their execution, but only upon termination. To the best of our knowledge, ours are the first online primal-dual algorithms with this behavior.

The above approach works directly for vertex-weighted matching. To generalize our approach to ad allocations, we first consider an intermediary problem – equal-bids ad allocation – where every advertiser i bids the same bid b_i for all neighbors $j \in N(i)$. We reduce this problem in (k, d) -graphs to the vertex-weighted problem in (k, d) -graphs in an online manner. We then rewrite this reduction along with our vertex-weighted online algorithm as a single online primal-dual algorithm for the equal-bids problem. Guided by this algorithm we devise a primal-dual algorithm for general-bids ad allocation on (k, d) -graphs, using a bounded fraction of the advertisers’ dual variables to guide our choice of matches and dual updates. This allows us to simulate the bid-scaling described above also in the case where each advertiser has different bids.

Finally, our randomized results stem from our deterministic primal-dual algorithms, whose dual updates we use in our dual-fitting analysis of the randomized algorithms. Dual feasibility follows as it does for our algorithms. The dual costs are bounded in expectation by the primal cost times the required constant, conditioned over the random algorithm’s previous choices. Taking total expectation over the possible previous choices yields the expected competitive ratio.

1.4 Related Work

Several stochastic models have been studied for the problems we address. Most prominent among these are the random arrival order and i.i.d model with known/unknown distribution. Our algorithms beat all of these bounds *in the worst case* for sufficiently small d/k and R_{\max} , replacing stochastic assumptions by structural ones.

For the random order model a line of work beginning with Goel and Mehta [9] has shown the optimal competitive ratio for maximum matching lies in the range $(0.696, 0.823)$ Feldman et al. [8], Karande et al. [13], Mahdian and Yan [16], Manshadi et al. [17]. For the known distribution model Feldman et al. [8] were the first to show the optimal competitive ratio is strictly greater than $1 - \frac{1}{e}$ and bounded away from 1. Subsequent work Bahmani and Kapralov [2], Haeupler et al. [10], Jaillet and Lu [11] showed the optimal competitive ratio for bipartite matching in this setting lies in the range $(0.706, 0.823)$, and $(0.729, 0.823)$ if the expected number of arrivals of each ad slot type is integral. For the vertex-weighted problem under the previously-mentioned integrality assumptions Haeupler et al. [10] and Jaillet and Lu [11] showed a lower bound of 0.667 and 0.725, respectively. For the AdWords problem under the random order model, Devanur and Hayes [5] give a $(1 - \epsilon)$ -competitive algorithm, assuming the online side’s size is known in advance and no bid is higher than roughly $\epsilon^3/|L|^2$ times the optimum value. Devanur et al. [7] gave an algorithm in the unknown distribution model achieving asymptotically optimal competitive ratio of $1 - O(\sqrt{R_{\max}})$.

In a different vein, Mahdian et al. [15] considered the AdWords problem given black-box estimates of the input. They show how to obtain performance trading-off between the worst-case optimal and the black-box’s performance on the given input. We require no such algorithm be available, but rather rely on domain-specific structure.

Closer to our work, Buchbinder et al. [4] considered $(1, d)$ -graphs for equal-bids ad allocation. We obtain more general results, and strictly better bounds for all $k > d$.

1.5 Paper Outline

In Section 2 we formally define the problems considered throughout the paper. In Section 3 we give a tight analysis of algorithm GREEDY in (k, d) -graphs. In Section 4 we build on the hard examples of Section 3 and present optimal algorithms for the online maximum matching and vertex-weighted matching problems in (k, d) -graphs. In Section 5 we extend these results to the general ad allocation problem. In Section 6 we present hardness results for the problems considered. In Section 7 we extend our analysis to prove competitiveness of several simple randomized algorithms. We conclude with a discussion of future work and open questions in Section 8.

2 Problem Definitions

An instance of the *ad allocation* problem consists of a bipartite graph $G = (L, R, E)$. The left-hand L side corresponds to advertisers, and the right-hand side R to ad slots. Each advertiser $i \in L$ has some budget B_i and is willing to bid some value $b_{ij} \leq B_i$ for every neighboring ad slot $j \in N(i)$ (the bids of advertiser i need not be equal for all $j \in N(i)$). Each ad slot $j \in R$ can be allocated to (up to) one advertiser i , yielding a profit of b_{ij} . The bids for ad slots allocated to an advertiser i may not exceed i 's budget, B_i . Figure 1 presents the ad allocation problem's LP relaxation and its dual.

Primal (Packing)		Dual (Covering)	
maximize	$\sum_{(i,j) \in E} b_{ij} \cdot x_{ij}$	minimize	$\sum_{i \in L} B_i \cdot z_i + \sum_{j \in R} y_j$
subject to:		subject to:	
$\forall j \in R:$	$\sum_{(i,j) \in E} x_{ij} \leq 1$	$\forall (i,j) \in E:$	$b_{ij} \cdot z_i + y_j \geq b_{ij}$
$\forall i \in L:$	$\sum_{(i,j) \in E} b_{ij} \cdot x_{ij} \leq B_i$	$\forall i \in L:$	$z_i \geq 0$
$\forall (i,j) \in E:$	$x_{ij} \geq 0$	$\forall j \in R:$	$y_j \geq 0$

Figure 1: The fractional ad allocation LP and the corresponding dual

An instance of the *online* ad allocation problem consists of an ad allocation instance; the advertisers given up-front, along with their budgets, and the ad slots arriving one-by-one, together with their edges and bids. An online ad allocation algorithm must, upon arrival of an ad slot j , determine to which advertiser (if any) to allocate the ad slot. Allocations are irrevocable, and so must be made to *feasible* advertisers, whose residual budget is sufficient to pay their actual bid.

We will consider several interesting special cases of the above problem throughout this paper. These problems are both interesting in their own right (theoretically as well as practically), in addition to providing some insight towards achieving a solution to the general problem.

The *equal-bids* online ad allocation problem is the above problem with each advertiser i bidding the same value for all neighboring ad slots; i.e., $b_{ij} = b_i$ for all $j \in N(i)$.

The online *vertex-weighted matching* problem is the above problem with every advertiser i bidding all its budget for every neighboring ad slot; i.e., $b_{ij} = B_i$ for all $j \in N(i)$.

The online *maximum matching* problem is the above problem with all budgets and bids equal to 1; i.e., $b_{ij} = B_i = 1$ for all $j \in N(i)$.

3 Warm-up: Greediness in (k, d) -Graphs

In this section we show that the natural greedy algorithms for the problems considered, which in general graphs are only $1/2$ -competitive, achieve on (k, d) -graphs a competitive ratio tending to one as d/k tends to zero. We prove this result by applying dual-fitting, and prove our analysis is tight.

Algorithm GREEDY for the online ad allocation problem matches an ad slot $j \in R$ to a feasible neighbor i with highest bid b_{ij} . For the vertex-weighted case, where $b_{ij} = B_i$ for all $j \in N(i)$, this reduces to picking an unmatched neighbor of highest weight. Our analysis relies on a dual-fitting formulation, given in Algorithm 1 below.

Algorithm 1 AD ALLOCATION GREEDY (Dual-Fitting Formulation)

```

1: Init: set  $z_i \leftarrow 0$  for all  $i \in L$ 
2: for all  $j \in R$  do
3:   if  $j$  has a feasible neighbor then
4:     match  $j$  to a feasible neighbor maximizing  $b_{ij}$ 
5:     set  $x_{ij} \leftarrow 1$ 
6:     set  $z_i \leftarrow \min\{1, z_i + \frac{b_{ij}}{B_i}\}$ 
7:     set  $z_{i'} \leftarrow \min\{1, z_{i'} + \frac{b_{i'j}}{k \cdot B_i}\}$  for every feasible neighbor of  $j$ ,  $i' \neq i$ 
8:   for all  $i \in L$  do
9:     if  $i$ 's residual budget is less than  $R_{\max} \cdot B_i$  then
10:      set  $z_i \leftarrow 1$ 

```

Theorem 3.1. *Algorithm GREEDY is $\left(\frac{k}{k+d-1}\right)$ -competitive for the unweighted, vertex-weighted maximum matching and equal-bids ad allocation problems on (k, d) -graphs.*

Theorem 3.2. *Algorithm GREEDY is $\frac{(1-R_{\max}) \cdot k}{k+(d-1) \cdot (1-R_{\max})} > (1-R_{\max}) \cdot \frac{k}{k+d-1}$ competitive for online ad allocation on (k, d) -graphs with $k \geq 1$ and $R_{\max} = \max_{(i,j) \in E} \{b_{ij}/B_i\} < 1$.*

Proof. We prove the following claims: (a) z, y form a feasible dual solution (b) for every $j \in R$ the changes to the primal and dual solutions' values, ΔP and ΔD , satisfy $\Delta D/\Delta P \leq \frac{k+d-1}{k}$. (c) for the vertex-weighted and unweighed matching problems and equal-bids problem Lines 8-10 incur no dual cost, and (c') for the general ad allocation problem Lines 8-10 cost the dual solution no more than $R_{\max}/(1-R_{\max})$ times the primal profit. As x forms an integral feasible primal solution, claims (a,b,c) combined entail Theorem 3.1. Similarly, claims (a,b,c') entail Theorem 3.2, as claims (b) and (c') imply the ratio of the programs' overall values is at least

$$\frac{P}{D} \geq \frac{P}{\frac{k+d-1}{k} \cdot P + \frac{R_{\max}}{1-R_{\max}} \cdot P} = \frac{(1-R_{\max}) \cdot k}{k + (d-1) \cdot (1-R_{\max})}$$

Claim (a): For every advertiser $i \in L$, if over a $(1-R_{\max})$ -fraction of i 's budget is spent then z_i is set to one in Line 10. Otherwise, i is a feasible match of all of its neighbors j , each such j causing z_i to increase by at least $\frac{b_{ij}}{k \cdot B_i}$. As $\sum_j b_{ij} \geq k \cdot B_i$ then $z_i = 1$ by the algorithm's termination. Consequently, all dual inequalities are satisfied.

Claim (b): For each ad slot $j \in R$, by the choice of j 's match i , and the fact that j has degree $d(j)$ at most d , the primal value increases by $\Delta P = b_{ij}$ and the dual cost increases by at most $\Delta D = b_{ij} + \sum_{i' \in F_j \setminus \{i\}} b_{i'j}/k \leq b_{ij} \cdot (1 + \frac{d-1}{k})$.

Claim (c): For an advertiser $i \in L$ to have spent over $(1-R_{\max})B_i$ for all but the general problem, it must and have z_i set to one. Thus Lines 8-10 incur no dual cost.

Claim (c'): For an advertiser $i \in L$ to be affected by Lines 8–10, it must spend up to a $(1-R_{\max})$ -fraction of its budget. However, whenever i spends an f -fraction of its budget, the dual variable z_i increases by f in Line 6, and so the cost of increasing z_i in line 10 is at most $R_{\max} \cdot B_i$, while i garnered a primal profit of at least $(1-R_{\max}) \cdot B_i$. The total dual cost of Lines 8–10 is thus at most $\frac{R_{\max}}{1-R_{\max}} \cdot P$, for P the primal profit. \square

3.1 Tight Examples for Algorithm GREEDY

We show that our analysis of algorithm GREEDY for the unweighted and vertex-weighted matching is tight whenever $k \geq d - 1$.³

Theorem 3.3. *For all $k \geq d - 1$ there exist (k, d) -graphs G with maximal matchings that achieve a competitive ratio no better than $\frac{k}{k+d-1}$ on G .*

Proof. The tight example, along with a poor choice of matching, is defined as follows: The offline side contains $k + d - 1$ advertisers. The first k ad slots by order of arrival j_1, j_2, \dots, j_k each have degree exactly d , with the t -th ad slot j_t having as neighbors the t -th advertiser i_t (to which it is matched), as well as the last $d - 1$ advertisers, $i_{k+1}, i_{k+2}, \dots, i_{k+d-1}$. After the first k ad slots arrive all $d - 1$ unmatched advertisers have degree exactly k . (By this stage the ad slots and the last $d - 1$ advertisers form a copy of $K_{k,d-1}$. See Figure 2.) The following ad slots are used to increase the degree of the first k advertisers to at least k , while guaranteeing that the first k advertisers can be simultaneously matched to the last ad slots (for example, by having all the latter ad slots have degree one). The resulting graph is a (k, d) -graph with k of its advertisers matched for which all $k + d - 1$ advertisers can be matched simultaneously. \square




Figure 2: The graph and matching after the arrival of the first k ad slots, for $k = 7$ and $d = 4$. Matching edges are marked in **bold**.

For any R_{\max} a unit fraction, gluing $1/R_{\max}$ copies of the above tight example at the advertisers, with each advertiser having a budget of $1/R_{\max}$, yields an equal-bid ad allocation instance and greedy allocation for which the same $\frac{k}{k+d-1}$ performance holds, proving tightness of our analysis for equal-bid allocations. We now state a theorem implying our analysis' tightness for GREEDY in general ad allocations. For a proof of this theorem, see Appendix A.

Theorem 3.4. *For all $k \geq d - 1$ and $R_{\max} \leq \frac{1}{2}$ there exist (k, d) -graphs G and greedy allocations achieving a competitive ratio at most $\frac{(1-R_{\max}) \cdot k}{k+(d-1) \cdot (1-R_{\max})}$ on G .*

4 Optimal Vertex-Weighted Matching on (k, d) -graphs

The previous section shows our analysis of GREEDY is tight, though for a particular(ly bad) instantiation of the input, and more importantly of the algorithm. The family of tight examples suggests the following improved algorithm: match every arriving ad slot to an unmatched neighbor of highest degree. This algorithm, which we call HIGH-DEGREE, is given below. The intuition behind it, substantiated by the above examples, is that unmatched advertisers with higher degree may have fewer chances to be matched later. This approach fares better on the above examples (actually yielding an *optimal* solution), but can it do better than GREEDY for *all* (k, d) -graphs? We answer this question in the affirmative, proving a lower bound with *exponentially* smaller loss. In Section 6 we prove a matching lower bound, implying the algorithm's *optimality*.

Algorithm 2 HIGH-DEGREE

- 1: **for all** $j \in R$ **do**
 - 2: **if** j has an unmatched neighbor **then**
 - 3: match j to unmatched neighbor of highest degree.
-

³For $k < d - 1$ the $\frac{k}{k+d-1}$ bound is strictly less than the $\frac{1}{2}$ bound obtained by all maximal matchings, and so the bound cannot be tight for $k < d - 1$. We therefore turn our attention to the case $k \geq d - 1$.

4.1 Analysis of HIGH-DEGREE

Theorem 4.1. *Algorithm HIGH-DEGREE is $1 - (1 - \frac{1}{d})^k$ competitive for all (k, d) -graphs.*

Corollary 4.2. *On d -regular graphs algorithm HIGH-DEGREE is $1 - (1 - \frac{1}{d})^d$ competitive.*

This is the first result for maximum online matching in regular graphs in the adversarial setting, beating the $1 - \frac{1}{e}$ “barrier” *deterministically*.

Theorem 4.1 can be proven directly (see 4.2), but in order to set the groundwork for proofs of our more general results, we generalize this algorithm and rewrite it as a primal-dual algorithm, below (3). The constant C will be chosen during the analysis.

Algorithm 3 Vertex-Weighted HIGH-DEGREE (Primal-Dual Formulation)

- 1: **Init:** set $z_i \leftarrow 0$ for all $i \in L$.
 - 2: **for all** $j \in R$ **do**
 - 3: **if** j has an unmatched neighbor i **then**
 - 4: match j to an unmatched neighbor i maximizing $(z_i + C) \cdot b_{ij}$.
 - 5: set $x_{ij} \leftarrow 1$.
 - 6: set $z_i \leftarrow 1$.
 - 7: set $z_{i'} \leftarrow \min\{1, z_{i'} \cdot (\frac{d}{d-1}) + \frac{1}{d-1} \cdot C\}$ for every feasible neighbor of j , $i' \neq i$.
-

Theorem 4.3. *Algorithm 3 generalizes HIGH-DEGREE and is $1 - (1 - \frac{1}{d})^k$ competitive. Moreover, it gains revenue at least $(1 - (1 - \frac{1}{d})^k) \cdot (\sum_i B_i)$.*

Proof. We rely on the following observation, verifiable by induction: All unmatched advertisers i satisfy $z_i = C \cdot ((\frac{d}{d-1})^{d(i)} - 1)$. Hence Algorithm 3 matches each ad slot j to an unmatched neighbor i maximizing $b_{ij} \cdot C \cdot ((\frac{d}{d-1})^{d(i)} - 1)$. For the unweighted problem, $b_{ij} = 1$. By monotonicity of exponentiation, picking such i is tantamount to picking an advertiser of highest degree. We proceed to bound the algorithm’s gain.

Let $j \in R$ be some ad slot matched to i . The incurred change to the primal profit equals $\Delta P = b_{ij}$. By our choice of j ’s match, the change to the dual cost satisfies

$$\begin{aligned} \Delta D &= (1 - z_i) \cdot b_{ij} + \sum_{i' \in N(j) \setminus \{i\}} \left(\left(\frac{1}{d-1} \right) \cdot (z_{i'} + C) \cdot b_{i'j} \right) \\ &\leq (1 - z_i) \cdot b_{ij} + (d-1) \cdot \left(\frac{1}{d-1} \right) \cdot (z_i + C) \cdot b_{ij} \\ &= (1 + C) \cdot b_{ij}. \end{aligned}$$

Given dual feasibility, the above would imply a competitive ratio of $1/(1 + C)$. Hence, we choose the minimal C ensuring $z_i = 1$ by the algorithm’s end for all advertisers i (matched and unmatched alike). Recall all unmatched advertisers i satisfy $z_i = C \cdot ((\frac{d}{d-1})^{d(i)} - 1)$. As such i have degree at least k by the algorithm’s end (but possibly no higher), the minimal C ensuring $z_i = 1$ is $C = 1/((\frac{d}{d-1})^k - 1)$. As the dual solution has $z_i = 1$ for all i by the algorithm’s termination, the dual cost is exactly $D = \sum_{i \in L} B_i$. Consequently, the primal gain satisfies $P \geq \frac{1}{1+C} \cdot (\sum_i B_i)$. The theorem follows. \square

The above algorithm implies structural Corollary 1.4 and the following corollary.

Corollary 4.4. *For (k, d) -graphs with $k \geq d \cdot \ln |L|$, by integrality of number of vertices matched, HIGH-DEGREE successfully matches all of L , obtaining a maximum matching.*

We can extend our analysis to handle the possible existence of *outlying* advertisers i , that do not satisfy $\sum_j b_{ij} \geq k \cdot B_i$, and so may not satisfy $z_i = 1$, ruining dual feasibility. Let $S \subseteq L$ be the set of outlying advertisers, and assume $\sum_{i \in S} B_i \leq \alpha \cdot \sum_{i \in L} B_i$. As $z_i = 1$ for all $i \notin S$, we have $D \geq (1 - \alpha) \cdot \sum_{i \in L} B_i$, implying the following theorem.

Theorem 4.5 (Outliers). *Let $S \subseteq L$ be the set of outlying advertisers, and α be such that $\sum_{i \in S} B_i \leq \alpha \cdot \sum_{i \in L} B_i$. Then Algorithm 3 gains at least $(1 - \alpha) \cdot (1 - (1 - \frac{1}{d})^k) \cdot (\sum_i B_i)$, and in particular is $(1 - \alpha) \cdot (1 - (1 - \frac{1}{d})^k)$ -competitive.*

4.2 Potential-based Analysis of HIGH-DEGREE

In this subsection we present a potential-based proof of Theorem 4.1. We note that this proof can easily be extended to provide alternative proofs of Theorems 4.3 and 4.5.

Theorem 4.6. *Algorithm HIGH-DEGREE achieves value at least $(1 - (1 - \frac{1}{d})^k) \cdot |L|$ for all (k, d) -graphs $G = (L, R, E)$, and it is therefore $(1 - (1 - \frac{1}{d})^k)$ -competitive.*

Proof. Let $U_L \subseteq L$ denote the set of unmatched advertisers. Consider the following potential:

$$\phi = \sum_{i \in U_L} \left(\frac{d}{d-1} \right)^{d(i)}.$$

Algorithm HIGH-DEGREE outputs a matching that effectively strives to greedily minimize ϕ .⁴ The initial and final values of the potential function hold $\phi_{start} = |L|$ and $\phi_{final} \geq (d/(d-1))^k \cdot |U_L|$, respectively. Denote by $\Delta\phi_j$ the change to ϕ incurred by the arrival of ad slot $j \in R$. Clearly, if j is unmatched we have $\Delta\phi_j = 0$. On the other hand, if j is matched to a neighbor i , previously of degree $d(i)$, we find that i 's matching results in ϕ decreasing by $(d/(d-1))^{d(i)}$, and the degree of j 's remaining unmatched neighbors increase each cause ϕ to increase by at most $(d/(d-1))^{d(i)+1} - (d/(d-1))^{d(i)}$. Therefore, if j is matched to i we have

$$\begin{aligned} \Delta\phi_j &\leq - \left(\frac{d}{d-1} \right)^{d(i)} + (d-1) \cdot \left(\left(\frac{d}{d-1} \right)^{d(i)} \cdot \left(\frac{d}{d-1} - 1 \right) \right) \\ &= - \left(\frac{d}{d-1} \right)^{d(i)} + \left(\frac{d}{d-1} \right)^{d(i)} = 0. \end{aligned}$$

In other words $\Delta\phi_j \leq 0$, irrespective of whether or not j is matched. By this fact and our bounds on the initial and final potential, we find that

$$\left(\frac{d}{d-1} \right)^k \cdot |U_L| \leq \phi_{final} \leq \phi_{start} = |L|.$$

The theorem follows. □

5 Online Ad Allocation

In this section we solve the ad allocation problem. We consider first the equal-bids case, where each advertiser i offers the same bid for all its neighbors; i.e., $b_{ij} = b_i \forall j \in N(i)$. This will prove to be a useful stepping-stone towards a solution for general bids, in 5.1.

One way to solve equal-bids ad allocation is via an online reduction to vertex-weighted matching in (k, d) -graphs. As each advertiser i bids $\sum_{j \in N(i)} b_i \geq k \cdot B_i$ in total, we have $d(i) \geq k \cdot B_i/b_i$. Without loss of generality, B_i/b_i is integral. The reduction splits each i into B_i/b_i copies, each of value b_i and receiving up to k distinct edges of i , stopping if the copy is matched. The obtained graph G is (k, d) -bounded (perhaps after adding inconsequential neighbors to matched advertisers), and matchings in G induce allocations of same value for the ad allocation instance. As Algorithm 3 gains $1 - (1 - \frac{1}{d})^k$ of the sum of vertex weights, or

⁴The sum of unmatched advertisers' degrees may seem like a more natural potential function to consider, but it turns out that it cannot be used to derive tight bounds. E.g., it does not yield a bound significantly better than $\frac{5}{8} = 0.625$ for $k = d$.

equivalently the sum of budgets, applying it yields a $1 - (1 - \frac{1}{d})^k$ competitive solution to the original ad allocation instance.

We restate the above as a primal-dual algorithm for equal-bids ad allocation. (See Algorithm 4 below). In this algorithm, z_i^c serves the role of z_i in Algorithm 3 for i 's "current copy" (hence the c in the notation), weighted to reflect the copy contributes b_i/B_i of i 's budget. Intuitively, when i is matched we imagine its current copy is matched, and set z_i^c to b_i/B_i . Conversely, we ensure that once the copy has k edges $z_i^c = b_i/B_i$. Either way, once $z_i^c = b_i/B_i$, we add z_i^c to z_i and nullify z_i^c (moving to i 's next copy, whose dual variable would be zero in Algorithm 3.) The number of copies of i guarantees dual feasibility and the choice of match and dual updates guarantee the desired bound.

Algorithm 4 Equal-Bid Ad Allocation in (k, d) -graphs

- 1: **Init:** set $z_i \leftarrow 0, z_i^c \leftarrow 0$ for all $i \in L$
 - 2: **for all** $j \in R$ **do**
 - 3: **if** j has a feasible neighbor i **then**
 - 4: match j to feasible neighbor i maximizing $z_i^c \cdot B_i + C \cdot b_i$.
 - 5: set $x_{i,j} \leftarrow 1$.
 - 6: set $z_i^c \leftarrow b_i/B_i$.
 - 7: **for all** feasible neighbor of $j, i' \neq i$ **do**
 - 8: set $z_{i'}^c \leftarrow \min\{b_{i'}/B_{i'}, z_{i'}^c \cdot (\frac{d}{d-1}) + \frac{1}{d-1} \cdot C \cdot b_{i'}/B_{i'}\}$
 - 9: **for all** $i' \in N(j)$ with $z_{i'}^c = b_{i'}/B_{i'}$ **do**
 - 10: set $z_{i'} \leftarrow z_{i'} + z_{i'}^c$.
 - 11: set $z_{i'}^c \leftarrow 0$.
-

Theorem 5.1. *Algorithm 4 with $C = 1/((\frac{d}{d-1})^k - 1)$ gains revenue $(1 - (1 - \frac{1}{d})^k) \cdot \sum_i B_i$, and is thus $(1 - (1 - \frac{1}{d})^k)$ -competitive for the equal-bid problem on (k, d) -graphs.*

Proof. To bound the primal-dual ratio, we bound increases of $z_i^c \cdot B_i$, as all dual costs can be traced back to past increases of z_i^c . Consider some ad slot j matched to i . The primal gain is $\Delta P = b_i$, whereas the dual cost satisfies

$$\begin{aligned} \Delta D &\leq (b_i/B_i - z_i^c) \cdot B_i + \sum_{i' \in N(j) \setminus \{i\}} \left(\frac{1}{d-1}\right) \cdot (z_{i'}^c + C \cdot b_{i'}/B_{i'}) \cdot B_{i'} \\ &\leq b_i - z_i^c \cdot B_i + (d-1) \cdot \left(\frac{1}{d-1}\right) \cdot (z_i^c \cdot B_i + C \cdot b_i) \leq (1 + C) \cdot b_i \end{aligned}$$

As in Theorem 4.3's proof, $z_i^c = C \cdot \frac{b_i}{B_i} \left(\left(\frac{d}{d-1}\right)^{d^c(i)} - 1\right)$, where $d^c(i)$ is the degree of i 's current copy, or equivalently, the number of i 's edges since z_i^c was last nullified. Hence, by our choice of C , after at most k i -edges, $z_i^c = \frac{b_i}{B_i}$ (whether or not i is matched), and z_i is increased by $\frac{b_i}{B_i}$. As $d(i) \geq k \cdot \frac{B_i}{b_i}$ by the end, $z_i \geq 1$ for all i . The theorem follows. \square

5.1 General Bids

A natural way to extend Algorithm 4 to general bids would be to replace for every ad slot j and every neighbor i (or i') all appearances of b_i (or $b_{i'}$) by b_{ij} (resp., $b_{i'j}$) in the choice of j 's match and updates to $z_i^c, z_{i'}^c$ and z_i . Such dual updates would guarantee, similarly to our prior algorithms, that an advertiser i with budget B_i and rejected bids $b_{i0}, b_{i1}, \dots, b_{it}$ since its last match (ordered chronologically) would have dual variable

$$z_i^c = \frac{1}{d-1} \cdot C \cdot \sum_{r=0}^t \frac{b_{ir}}{B_i} \cdot \left(\frac{d}{d-1}\right)^{t-r} \quad (1)$$

Unfortunately, replacing b_i by b_{ij} in the updates for matched i could result in z_i arbitrarily small. Worse still, since previously-rejected bids may be greater than the current bid, setting z_i^c to $\frac{b_{ij}}{B_i}$

could even *decrease* z_i^c , complicating the task of bounding the primal-dual ratio. Algorithm 5 below sidesteps these issues by considering bounded fractions of z_i^c , and using the following notation, motivated by Equation 1, to represent variables z_i^c , and z_i^f (the f in the notation refers to a bounded fraction of z_i^c “used”). This notation’s use will become apparent during the algorithm’s analysis.

Definition 5.2. Let $z = \frac{1}{d-1} \cdot C \cdot \sum_{r=0}^t b_r \cdot \left(\frac{d}{d-1}\right)^r$. We think of z as a number in base $\frac{d}{d-1}$, denoting it by $z = [b_t, \dots, b_1, b_0]$, disregarding the $\frac{1}{d-1} \cdot C$ term for simplicity. Addition and subtraction of numbers in this notation is done place-wise, disallowing carries/borrows. In particular, if $z = [b_t, \dots, b_1, b_0]$, then $z \cdot \frac{d}{d-1} + \frac{1}{d-1} \cdot C \cdot b = [b_t, \dots, b_1, b_0, b]$. Comparisons involving numbers in this notation refer to their numerical value.

Algorithm 5 Online Ad Allocation in (k, d) -graphs with general bids.

```

1: Init: set  $z_i \leftarrow 0, z_i^c \leftarrow 0$  for all  $i \in L$ 
2: for all  $j \in R$  do
3:   if  $j$  has a feasible neighbor  $i$  then
4:     for all feasible neighbors  $i$  do
5:       let  $z_i^c = [b_{k-1}, \dots, b_1, b_0]$ .
6:       set  $z_i^f \leftarrow [\min\{b_{k-1}, b_{ij}/B_i\}, \dots, \min\{b_1, b_{ij}/B_i\}, \min\{b_0, b_{ij}/B_i\}]$ .
7:       set  $z_i^c \leftarrow z_i^c - z_i^f$ .
8:       match  $j$  to feasible neighbor  $i$  maximizing  $z_i^f \cdot B_i + C \cdot b_{ij}$ .
9:       set  $x_{i,j} \leftarrow 1$ .
10:      set  $z_i^f \leftarrow 0$ .
11:      set  $z_i \leftarrow z_i + b_{ij}/B_i$ .
12:      for all feasible neighbor of  $j, i' \neq i$  do
13:        set  $z_{i'}^f \leftarrow z_{i'}^f \cdot \left(\frac{d}{d-1}\right) + \frac{1}{d-1} \cdot C \cdot b_{i'j}/B_{i'}$ 
14:         $z_{i'}^c \leftarrow z_{i'}^c + z_{i'}^f$ .
15:        if  $z_{i'}^c = [b_k, b_{k-1}, \dots, b_1, b_0]$  with  $b_k \neq 0$  then
16:          set  $z_{i'} \leftarrow z_{i'} + b_k \cdot \frac{1}{k}$ .
17:          set  $z_{i'}^c \leftarrow [b_{k-1}, \dots, b_1, b_0]$ .
18:          if  $z_{i'}^c = [b_{k-1}, \dots, b_1, b_0]$  with all digits  $b_r \neq 0$  then
19:            let  $b = \min\{b_r\}_{r=0}^{k-1}$ .
20:            set  $z_{i'} \leftarrow z_{i'} + b$ .
21:            set  $z_{i'}^c \leftarrow [b_{k-1} - b, \dots, b_1 - b, b_0 - b]$ .
22:      for all  $i \in L$  do
23:        set  $z_i \leftarrow \max\{1, z_i\}$ .

```

The algorithm for the general bids setting is Algorithm 5, below. The algorithm’s primal feasibility is trivial, as is its dual feasibility, due to Lines 22-23. It remains to bound the ratio of the cost of the dual solution to the value of the primal solution.

High-Level intuition: The algorithm asserts three invariants. The first guarantees increases in z_i are “paid for” by increases in z_i^c , allowing us to focus on bounding changes to z_i^c . A second invariant guarantees every increase of z_i by some value b/B_i can be accredited to previous bids (or fractions thereof) of total value at most $k \cdot b/B_i$. As the graph is (k, d) , if every bid of i of value b were to cause z_i to increase (by at least $b/(k \cdot B_i)$, by the above), then eventually $z_i \geq 1$. However, some bids may not incur an increase in z_i . The third and last invariant guarantees the total value of fractions of bids that do not cause z_i to increase is at most $k \cdot R_{\max}$, and so $z_i \geq (1 - R_{\max})$ before Lines 22-23. Thus, the cost of rounding each z_i to one in these lines is at

most $R_{\max}/(1 - R_{\max})$ of the previously-paid dual cost. The bound will follow. The following four lemmas formalize the above, allowing us to derive our sought-after bound.

Lemma 5.3. *Before every ad slot's arrival and before Line 22, every z_i^c is a number in the above numeral system satisfying the following three properties:*

- (i) z_i^c is a k -digit number; i.e., $z_i^c = [b_{k-1}, \dots, b_1, b_0]$.
- (ii) z_i^c has at most $k - 1$ non-null digits.
- (iii) Each digit of z_i^c is no greater than $\max_j \{b_{ij}/B_i\}$.

Proof. Properties i and ii are enforced explicitly by Lines 15-17 and 18-21, respectively. Property iii follows by induction: When z_i^f is subtracted from z_i^c , every digit of z_i^c is either nullified, if it was smaller than b_{ij}/B_i , or decreased by b_{ij}/B_i . After z_i^f is updated and added to z_i^c , each digit of z_i^c is increased by at most b_{ij}/B_i . Thus each digit is no greater than its previous value and b_{ij}/B_i , both of which are at most $\max_j \{b_{ij}/B_i\}$. \square

Lemma 5.4. *If $k \geq d - 1$ and $C = 1/((\frac{d}{d-1})^k - 1)$, every increase in z_i by some b in Lines 15-17 and 18-21 goes hand-in-hand with both*

- (i) a decrease of the same value or higher in z_i^c , and
- (ii) a decrease of k times this value or less in the sum of digits of z_i^c .

Proof. In Lines 15-17, z_i is increased by b_k/k . On the other hand, we remove b_k , the k -th digit of z_i^c in this numeral system, resulting in a decrease of z_i^c by

$$\frac{1}{d-1} \cdot C \cdot b_k \cdot \left(\frac{d}{d-1}\right)^k \geq \frac{1}{k} \cdot b_k.$$

Thus, Properties i and ii both hold for Lines 15-17. In Lines 18-21, the value of z_i^c is decreased by $\frac{1}{d-1} \cdot C \cdot \sum_{r=0}^{k-1} b \cdot (\frac{d}{d-1})^r = C \cdot ((\frac{d}{d-1})^k - 1) \cdot b$, which is exactly b , by our choice of C . The decrease in the sum of digits of z_i^c on the other hand is exactly $k \cdot b$. \square

Lemma 5.5. *Taking $C = 1/((\frac{d}{d-1})^k - 1)$ guarantees every increase in z_i by b_{ij}/B_i in Line 11 coincides with a decrease of at most b_{ij}/B_i in z_i^c . Moreover, Δdigit , the decrease in sum of digits of z_i^c , satisfies $\Delta\text{digit} + b_{ij}/B_i \leq k \cdot b_{ij}/B_i$.*

Proof. In Line 11, z_i^f , which was subtracted from z_i^c , is nullified. Both bounds follow similarly to our proof of Lemma 5.4 relying on z_i^f being a k -digit number with at most $k - 1$ non-null digits, by Lemma 5.3, and each digit of z_i^f being no greater than b_{ij}/B_i , by initialization of z_i^f . \square

Lemma 5.6. *By Line 22 each i satisfies $z_i \geq \frac{\sum_j b_{ij} - k \cdot \max_j \{b_{ij}\}}{k \cdot B_i} \geq 1 - \frac{\max_j b_{ij}}{B_i} \geq 1 - R_{\max}$.*

Proof. Throughout the algorithm, every edge (i, j) causes the sum of digits of z_i^c to increase by b_{ij}/B_i (again ignoring the $\frac{1}{d-1} \cdot C$ term), unless (i, j) are matched. Moreover, the sum of digits does not decrease due to carries. On the other hand, every increase in z_i by b coincides with a decrease in the sum of digits of z_i^c plus $\sum_{(i,j) \text{ matched}} b_{ij}/B_i$, of at most $k \cdot b$, by Lemmas 5.4 and 5.5. Put otherwise, the increase in z_i is at least $1/k$ times the total sum of i 's bids so far, minus the sum of digits of z_i^c . By Lemma 5.3, the sum of digits of z_i^c by Line 22 cannot exceed $k \cdot \max_j \{b_{ij}/B_i\}$. The lemma follows. \square

Given the above we can now prove our main result.

Theorem 5.7. *On general-bid ad allocations on (k, d) -graphs with $k \geq d - 1$ Algorithm 5 gains $\sum_i (B_i - \max_j b_{ij}) \cdot (1 - (1 - \frac{1}{d})^k)$, and is thus $(1 - R_{\max}) \cdot (1 - (1 - \frac{1}{d})^k)$ -competitive.*

Proof. Lemmas 5.4 and 5.5 imply increases in z_i can be traced back to a previous increase in z_i^c of the same value or higher. We therefore bound increases of $z_i^c \cdot B_i$ in order to bound the total dual cost. For each online $j \in R$, by our choice of match i , the change to the dual cost is at most $(1 + C)$ times the change to the primal value, as in Algorithm 4. However, by Lemma 5.6, by Line 22 each i satisfies $z_i \geq (1 - \max_j b_{ij}/B_i)$. Consequently, we have that before Line 22 the primal value P and dual cost D satisfy

$$P \geq \frac{1}{1+C} \cdot D \geq \sum_i (B_i - \max_{j \in N(i)} b_{ij}) \cdot \left(1 - \left(1 - \frac{1}{d}\right)^k\right)$$

As the primal value is unaffected by Lines 22-23, P above is our algorithm's gain. The competitive ratio follows from $OPT \leq \sum_i B_i$ and the definition of R_{\max} . \square

Finally, we note that Lemmas 5.3, 5.4, 5.5 and 5.6 hold for all advertisers i satisfying $\sum_j b_{ij} \geq k \cdot B_i$, irrespective of outliers who don't hold this property, implying the following.

Theorem 5.8 (Outliers). *Let $S \subseteq L$ be the set of outlying advertisers (advertisers i with $\sum_j b_{ij} < k \cdot B_i$), and α be such that $\sum_{i \in S} B_i \leq \alpha \cdot \sum_{i \in L} B_i$. Then Algorithm 3 gains at least*

$$(1 - \alpha) \cdot \left((1 - R_{\max}) \cdot \left(1 - \left(1 - \frac{1}{d}\right)^k\right) \right) \cdot \left(\sum_i B_i\right),$$

and in particular is $(1 - \alpha) \cdot \left((1 - R_{\max}) \cdot \left(1 - \left(1 - \frac{1}{d}\right)^k\right) \right)$ -competitive.

6 Upper Bounds

6.1 Maximum Matching and Vertex-Weighted Matching

In order to construct hard examples, we start by showing that the optimal matching in (k, d) -graphs matches *all* the advertisers whenever $k \geq d$.

Lemma 6.1. *Every (k, d) -graph $G = (L, R, E)$ with $k \geq d$ has a matching matching all of L .*

Proof. By Hall's Theorem G has a matching with all of L matched if and only if every subset $A \subseteq L$ satisfies $|\Gamma(A)| \geq |A|$. But, as G is a (k, d) -graph we have

$$k \cdot |A| \leq |E[G[A]]| \leq d \cdot |\Gamma(A)|$$

Consequently, we find that $|\Gamma(A)| \geq \frac{k}{d} \cdot |A| \geq |A|$, and the lemma follows. \square

Equipped with Lemma 6.1 we may now prove this section's main result – an upper bound matching the lower bounds of Section 4, implying algorithm HIGH-DEGREE's optimality. To this end we cause HIGH-DEGREE to be effectively indistinguishable from any other algorithm.

Theorem 6.2. *For all $k \geq d$ no deterministic online algorithm for bipartite matching can achieve competitive ratio better than $1 - \left(1 - \frac{1}{d}\right)^k$ on (k, d) -graphs.*

Proof. Let \mathcal{A} be some online matching algorithm. The adversarial input consists of d^{k+1} advertisers, with the ad slots arriving in k phases, numbered 0 to $k-1$. During the i -th phase, which begins with $d^{k+1} \cdot \left(1 - \frac{1}{d}\right)^i$ unmatched advertisers each of degree i , the arriving ad slots each have exactly d neighbors, all unmatched; every unmatched advertiser neighbors exactly one new ad slot per phase. Every phase causes unmatched advertisers to have their degree increase by one, and exactly a $\left(1 - \frac{1}{d}\right)$ -fraction of the advertisers unmatched at the phase's beginning remain unmatched. (If algorithm \mathcal{A} does not match some ad slot to one of its d unmatched neighbors, we

consider it matched to an arbitrary neighbor; this can only serve to improve \mathcal{A} 's performance.) After the k phases additional ad slots of degree exactly d arrive in order to increase the degree of the matched advertisers to k . The resulting graph is k -regular and d -regular on the offline and online sides respectively, and is thus a (k, d) -graph. Moreover, exactly $d^{k+1} \cdot (1 - \frac{1}{d})^k$ of the d^{k+1} advertisers are unmatched. However, by Lemma 6.1 all d^{k+1} advertisers can be matched simultaneously. The theorem follows. \square

Corollary 6.3. *The bound of Theorem 6.2 holds for d -regular graphs with $d^{d+1} \leq n$, where $n = |L| = |R|$. In particular, Algorithm HIGH-DEGREE is optimal for d -regular graphs with $d = O\left(\frac{\log n}{\log \log n}\right)$.*

6.2 Upper Bound for Ad Allocation

In this subsection we prove an upper bound on the possible competitive ratio of deterministic algorithms in (k, d) -graphs. We start by showing a simple weaker bound, useful in proving this section's main result.

Lemma 6.4. *For all ratio R_{\max} no deterministic algorithm can achieve competitive ratio better than $(1 - R_{\max})$ for the ad allocation problem under the adversarial model. This bound holds even for (k, d) -graphs for all k and d .*

Proof. The hard input consists of disjoint stars with advertisers for internal vertices and ad slots for leaves. Every advertiser i has budget $B_i = 1$, with i 's bids given by

$$b_{ij} = \begin{cases} R_{\max} & \text{if } i\text{'s remaining budget is less than } R_{\max} \\ \epsilon & \text{else} \end{cases}$$

Given enough ad slots, an optimal allocation exhausts all advertisers' budgets, but every advertiser i gains at most $1 - R_{\max} + \epsilon$, whether or not i has neighbors j with $b_{ij} = R_{\max}$. Summing over all advertisers, the lemma follows. \square

Using the above and extending Theorem 6.2's proof, we can now prove the following.

Theorem 6.5. *For all $k \geq d$ no deterministic online algorithm for ad allocation is better than $(1 - R_{\max}) \cdot \left(1 - \left(1 - \frac{1}{d}\right)^{k/R_{\max}}\right)$ -competitive on (k, d) -graphs with $R_{\max} \leq \frac{1}{2}$ a unit fraction.*

Proof. The offline side consists of $d^{k/R_{\max}}$ advertisers, each with a budget $B_i = 1$. For the first phase, all edges have bids R_{\max} . During k/R_{\max} rounds ad slots arrive, each neighboring d distinct advertisers, and a $\frac{1}{d}$ -fraction of the advertisers are matched. The next round is as the last, but restricted to the previously unmatched advertisers. There are $(1 - \frac{1}{d})^{k/R_{\max}}$ unmatched advertisers by this phase's termination; these advertisers now satisfy the offline side's constraints for (k, d) -graphs, and receive no more neighbors. All of these advertisers' potential profit is lost. For the matched advertisers we now apply the construction of Lemma 6.4 to guarantee that at most a $(1 - R_{\max})$ -fraction of their potential profit in an optimal solution is gained, for a total gain of $(1 - R_{\max}) \cdot \left(1 - \left(1 - \frac{1}{d}\right)^{k/R_{\max}}\right) \cdot |L|$. Applying Lemma 6.1 repeatedly we find that there exists an allocation with all advertisers unmatched by algorithm \mathcal{A} matched $1/R_{\max}$ times (to neighbors for which they bid R_{\max}), and all advertisers matched by \mathcal{A} also exhausting their budgets simultaneously. The theorem follows. \square

7 Randomized Algorithms

By relying on the dual updates of Algorithm 3, we prove competitiveness of algorithm RANDOM, which matches every arriving ad slot to some feasible neighbor.

Theorem 7.1. *Algorithm RANDOM achieves expected competitive ratio of $1 - (1 - \frac{1}{d})^k$ for both unweighted and vertex-weighted matching problems.*

Proof. Consider some maximally-matching algorithm \mathcal{A} .⁵ We maintain and update a dual solution as in our deterministic Algorithm 3 while choosing matches according to Algorithm \mathcal{A} . As observed in the proof of Theorem 4.3, such dual updates guarantee all unmatched advertisers $i \in L$ with current degree $d(i)$ satisfy

$$z_i = C \cdot \left(\left(\frac{d}{d-1} \right)^{d(i)} - 1 \right)$$

Consequently, these dual update rules guarantee dual feasibility for any maximally-matching algorithm, including RANDOM, provided $C = 1 / ((\frac{d}{d-1})^k - 1)$. We need only bound the expected ratio between the dual and primal solutions' values.

Consider some ad slot j matched to some i . We recall that in the vertex-weighted matching problem the bid b_{ij} is exactly $b_{ij} = B_i$. Therefore, given the current state (determined by the previous random choices), including the set $N_F(j)$ of j 's unmatched (feasible) neighbors, j 's match is chosen uniformly among $N_F(j)$ by RANDOM, and consequently

$$\mathbb{E}[\Delta P | \text{state}] = \frac{1}{|N_F(j)|} \cdot \sum_{i \in N_F(j)} B_i.$$

On the other hand, by the same argument

$$\begin{aligned} \mathbb{E}[\Delta D | \text{state}] &= \frac{1}{|N_F(j)|} \cdot \sum_{i \in N_F(j)} \left((1 - z_i) \cdot B_i + \sum_{i' \in N_F(j) \setminus \{i\}} \left(\frac{1}{d-1} \cdot (z_{i'} + C) \cdot B_{i'} \right) \right) \\ &= \frac{1}{|N_F(j)|} \cdot \sum_{i \in N_F(j)} \left((1 - z_i) \cdot B_i \right) + \frac{1}{|N_F(j)|} \cdot \sum_{i' \in N_F(j)} \left(\frac{|N_F(j)|-1}{d-1} \cdot (z_{i'} + C) \cdot B_{i'} \right) \\ &\leq \frac{1}{|N_F(j)|} \cdot \sum_{i \in N_F(j)} B_i \cdot (1 + C). \end{aligned}$$

With the last inequality following from $|N_F(j)| \leq |N(j)| \leq d$. Taking total expectation over the possible states, we obtain $\mathbb{E}[\Delta D] \leq (1 + C) \cdot \mathbb{E}[\Delta P]$. The theorem follows. \square

We note that Theorem 7.1 can also be proved using the potential-based proof of Subsection 4.2, observing that the expected potential change incurred by the processing of every online arrival is non-negative. In addition, in the same way that Theorem 4.3 is extended in Theorem 4.5, we can show that RANDOM is also robust to outliers. We omit the details for brevity. Finally, we show that RANDOM also performs well for the general online ad allocation problem.

Theorem 7.2. *Algorithm RANDOM achieves expected competitive ratios of $1 - (1 - \frac{1}{d})^k$ and $(1 - R_{\max}) \cdot (1 - (1 - \frac{1}{d})^k)$ for the equal-bids and general-bids ad allocation problems.*

Proof (sketch). The proof resembles that of Theorem 7.1, relying on Algorithms 4 and 5 respectively for the dual-fitting analysis. Dual feasibility is guaranteed by the dual updates. On the other hand, linearity of expectation implies the expected primal-dual ratio matches that of Algorithms 4 and 5 (for the latter, this requires showing Lemmas 5.3–5.6 all hold in expectation). The claimed bounds follow. \square

⁵By “maximally-matching” we mean algorithm \mathcal{A} always matches an arriving ad slot, if at all possible.

8 Future Work and Open Questions

This paper attempts to give a theoretical explanation of the empirical success of simple heuristic algorithms for online ad allocation in practice and further proposes better algorithms under assumptions that could explain above-mentioned success. The paper further poses several interesting follow-up questions.

Optimality for Adwords. We proved optimality of our algorithms for online maximum and vertex-weighted matching. However, for the general ad allocation problem our lower and upper bounds differ by a factor of $(1 - (1 - \frac{1}{d})^k) / (1 - (1 - \frac{1}{d})^{k/R_{\max}})$. For small R_{\max} (i.e., the AdWords problem), this discrepancy is large. Can better algorithms be obtained for this problem, or can the upper bounds be tightened (or both)?

Randomization. We are able to prove that a multitude of randomized algorithms, which we do not discuss for the sake of brevity, match our deterministic algorithms' expected competitive on (k, d) -graphs. However, we have no upper bounds on randomized algorithms' performance. We believe that randomization does allow for better results, and have some preliminary results that indicate this is indeed the case.

Stochastic Models. An interesting direction would be to extend our analysis of (k, d) -graphs to stochastic models, in which it seems plausible that even better competitiveness guarantees should be achievable.

9 Acknowledgments

This work was supported in part by United States-Israel BSF Grant No. 2010-246 and ISF Grant No. 954/11.

References

- [1] AGGARWAL, G., GOEL, G., KARANDE, C., AND MEHTA, A. 2011. Online vertex-weighted bipartite matching and single-bid budgeted allocations. In *Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms*. SIAM, 1253–1264.
- [2] BAHMANI, B. AND KAPRALOV, M. 2010. Improved bounds for online stochastic matching. In *Algorithms-ESA 2010*. Springer, 170–181.
- [3] BIRNBAUM, B. AND MATHIEU, C. 2008. On-line bipartite matching made simple. *ACM SIGACT News* 39, 1, 80–87.
- [4] BUCHBINDER, N., JAIN, K., AND NAOR, J. S. 2007. Online primal-dual algorithms for maximizing ad-auctions revenue. In *Algorithms-ESA 2007*. Springer, 253–264.
- [5] DEVANUR, N. R. AND HAYES, T. P. 2009. The adwords problem: online keyword matching with budgeted bidders under random permutations. In *Proceedings of the 10th ACM conference on Electronic commerce*. ACM, 71–78.
- [6] DEVANUR, N. R., JAIN, K., AND KLEINBERG, R. D. 2013. Randomized primal-dual analysis of ranking for online bipartite matching. In *SODA*. SIAM, 101–107.
- [7] DEVANUR, N. R., SIVAN, B., AND AZAR, Y. 2012. Asymptotically optimal algorithm for stochastic adwords. In *Proceedings of the 13th ACM Conference on Electronic Commerce*. ACM, 388–404.
- [8] FELDMAN, J., MEHTA, A., MIRROKNI, V., AND MUTHUKRISHNAN, S. 2009. Online stochastic matching: Beating $1-1/e$. In *Foundations of Computer Science, 2009. FOCS'09. 50th Annual IEEE Symposium on*. IEEE, 117–126.

- [9] GOEL, G. AND MEHTA, A. 2008. Online budgeted matching in random input models with applications to adwords. In *Proceedings of the nineteenth annual ACM-SIAM symposium on Discrete algorithms*. SIAM, 982–991.
- [10] HAEUPLER, B., MIRROKNI, V. S., AND ZADIMOGHADDAM, M. 2011. Online stochastic weighted matching: Improved approximation algorithms. In *Internet and Network Economics*. Springer, 170–181.
- [11] JAILLET, P. AND LU, X. 2013. Online stochastic matching: New algorithms with better bounds. *Mathematics of Operations Research*.
- [12] KALYANASUNDARAM, B. AND PRUHS, K. R. 2000. An optimal deterministic algorithm for online b-matching. *Theoretical Computer Science* 233, 1, 319–325.
- [13] KARANDE, C., MEHTA, A., AND TRIPATHI, P. 2011. Online bipartite matching with unknown distributions. In *Proceedings of the 43rd annual ACM symposium on Theory of computing*. ACM, 587–596.
- [14] KARP, R. M., VAZIRANI, U. V., AND VAZIRANI, V. V. 1990. An optimal algorithm for on-line bipartite matching. In *Proceedings of the twenty-second annual ACM symposium on Theory of computing*. ACM, 352–358.
- [15] MAHDIAN, M., NAZERZADEH, H., AND SABERI, A. 2007. Allocating online advertisement space with unreliable estimates. In *Proceedings of the 8th ACM conference on Electronic commerce*. ACM, 288–294.
- [16] MAHDIAN, M. AND YAN, Q. 2011. 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*. ACM, 597–606.
- [17] MANSHADI, V. H., GHARAN, S. O., AND SABERI, A. 2012. Online stochastic matching: Online actions based on offline statistics. *Mathematics of Operations Research* 37, 4, 559–573.
- [18] MEHTA, A. 2012. Online matching and ad allocation. *Theoretical Computer Science* 8, 4, 265–368.
- [19] MEHTA, A., SABERI, A., VAZIRANI, U., AND VAZIRANI, V. 2007. Adwords and generalized online matching. *Journal of the ACM (JACM)* 54, 5, 22.

A Bad Instance for Algorithm GREEDY for General Ad Allocation

We now prove our analysis of Algorithm GREEDY for the ad allocation problem is tight.

Theorem A.1. *For all $k \geq d - 1$ and $R_{\max} \leq \frac{1}{2}$ a unit fraction, there exist (k, d) ad allocation instances for which algorithm GREEDY can achieve competitive ratio exactly $\frac{(1-R_{\max}) \cdot k}{k+(d-1) \cdot (1-R_{\max})}$.*

Proof. Let $1 - R_{\max} = \frac{a}{b}$ for $0 < a < b$ and a and b integers. The hard instance will consist of $k \cdot b + (d - 1) \cdot a = b \cdot (k + (d - 1) \cdot (1 - R_{\max}))$ advertisers. Each advertiser has budget exactly one. We designate $k \cdot b$ advertisers to be the “lucky” advertisers, from which we will achieve revenue of $(1 - R_{\max})$ and the remaining $(d - 1) \cdot a$ “unlucky” advertisers will garner no profit. The theorem will follow by constructing the instance such that all budgets can be exhausted simultaneously.

All edges have bids either R_{\max} or some arbitrarily small positive ϵ . At first, each arriving ad slot will have d edges with bids R_{\max} , one to some lucky advertiser of lowest degree (to whom the ad slot is matched), and $(d - 1)$ edges to some unlucky advertisers of lowest degree. After $\frac{a \cdot k \cdot b}{b - a}$ ad slots arrive (this value is integral, as is $\frac{1}{R_{\max}} = \frac{b}{b - a}$), the following holds

- (i) each of the unlucky advertisers are unmatched and have degree exactly $k \cdot \frac{b}{b - a} = k \cdot \frac{1}{R_{\max}}$.
- (ii) each of the lucky ad slots are matched to all of their neighbors, and have degree exactly $\frac{a}{b - a} = \frac{1}{R_{\max}} - 1$.

The remaining ad slots recreate the construction of Lemma 6.4, thus guaranteeing each of the lucky advertisers gain no more than $1 - R_{\max} + \epsilon$. On the other hand all the lucky advertisers can exhaust their budgets without using any of the R_{\max} -bid edges of ad slots neighboring unlucky advertisers, which, as can be readily verified (using, e.g. Lemma 6.1 repeatedly), allows both lucky and unlucky advertisers to exhaust their budgets simultaneously whenever $k \geq d - 1$. The described instance is (k, d) and the theorem follows. \square

The above bound holds for any $R_{\max} \leq \frac{1}{2}$, as the following theorem asserts.

Theorem A.2. *For all $k \geq d - 1$ and $R_{\max} \leq \frac{1}{2}$ there exist (k, d) ad allocation instances for which algorithm GREEDY can achieve competitive ratio exactly $\frac{(1-R_{\max}) \cdot k}{k+(d-1) \cdot (1-R_{\max})}$.*

Proof (sketch). In order to generalize the above, we rely on the fact that every number R_{\max} in the range $(0, \frac{1}{2}]$ can be written as a convex combination of two unit fractions $\frac{1}{a} \leq R_{\max} \leq \frac{1}{b}$. That is, $w_a \cdot \frac{1}{a} + w_b \cdot \frac{1}{b} = R_{\max}$ and $w_a + w_b = 1$. We glue $2n$ copies of the above construction at the advertisers, n of the copies with budget w_a/n (w_b/n) for each advertiser, and highest bid-to-budget ratio in the copy being $1/a$ (resp. $1/b$). In this case the overall budget from all copies is $n \cdot (w_a/n + w_b/n) = w_a + w_b = 1$, and for large enough n each bid is at most $w_b/(b \cdot n) < R_{\max}$. On the other hand, all unlucky advertisers are completely unmatched and garner no profit, and all lucky advertisers gain a total of $n \cdot (w_a/n + w_b/n - w_a/(a \cdot n) + w_b/(b \cdot n)) = 1 - (w_a/a + w_b/b) = 1 - R_{\max}$. As in Lemma 6.4 we can guarantee each such lucky advertiser yields at most ϵ additional revenue. \square