Online Stochastic Matching: Beating $1 - \frac{1}{e}$

Jon Feldman [*]   Aranyak Mehta [†]   Vahab Mirrokni [‡]   S. Muthukrishnan [§]

May 26, 2009

Abstract

We study the online stochastic bipartite matching problem, in a form motivated by display ad allocation on the Internet. In the online, but adversarial case, the celebrated result of Karp, Vazirani and Vazirani gives an approximation ratio of $1 - \frac{1}{e} \simeq 0.632$, a very familiar bound that holds for many online problems; further, the bound is tight in this case. In the online, stochastic case when nodes are drawn repeatedly from a known distribution, the greedy algorithm matches this approximation ratio, but still, no algorithm is known that beats the $1 - \frac{1}{e}$ bound.

Our main result is a 0.67-approximation online algorithm for stochastic bipartite matching, breaking this $1 - \frac{1}{e}$ barrier. Furthermore, we show that no online algorithm can produce a $1 - \epsilon$ approximation for an arbitrarily small ϵ for this problem.

Our algorithms are based on computing an optimal offline solution to the expected instance, and using this solution as a guideline in the process of online allocation. We employ a novel application of the idea of the power of two choices from load balancing: we compute two disjoint solutions to the expected instance, and use both of them in the online algorithm in a prescribed preference order. To identify these two disjoint solutions, we solve a max flow problem in a boosted flow graph, and then carefully decompose this maximum flow to two edge-disjoint (near-)matchings. In addition to guiding the online decision making, these two offline solutions are used to characterize an upper bound for the optimum in any scenario. This is done by identifying a cut whose value we can bound under the arrival distribution.

At the end, we discuss extensions of our results to more general bipartite allocations that are important in a display ad application.

---

[*]Google, Inc., New York, NY. `jonfeld@google.com`

[†]Google, Inc., Mountain View, CA. `aranyak@google.com`

[‡]Google, Inc., New York, NY. `mirrokni@google.com`

[§]Google, Inc., New York, NY. `muthu@google.com`

1 Introduction

Bipartite matching problems are central in combinatorial optimization with many applications. Our motivating application is the allocation of display advertisements on the Internet,[1] and so we will use the language of this application to define and discuss the problem:

(Online Bipartite Matching) There is a bipartite graph $G(A, I, E)$ with advertisers A and impressions I, and a set E of edges between them. Advertisers in A are fixed and known. Impressions (or requests) in I (along with their incident edges) arrive online. Upon the arrival of an impression $i \in I$, we must assign i to any advertiser $a \in A$ where $(i, a) \in E(G)$. At all times, the set of assigned edges must form a matching (that is, no end points coincide). □

If the online algorithm knows nothing about I or E beforehand, and the impressions arrive in an arbitrary order, we have the *adversarial model*. Then, Karp, Vazirani and Vazirani [14] solved this problem by presenting an online algorithm with an approximation ratio of $1 - 1/e \simeq 0.632$, and further showed that no algorithm can achieve a better ratio.

A different model is the online, stochastic one called the *iid* model, where impressions $i \in I$ arrive online according some *known* probability distribution (with repetition). In other words, in addition to G, we are given a probability distribution $\mathcal{D}$ over the elements of I. Our goal is then to compute a maximum matching on $\hat{G} = (A, \hat{I}, \hat{E})$, where $\hat{I}$ is drawn from $\mathcal{D}$.[2] In this *iid* model, the greedy algorithm achieves an approximation ratio of $1 - 1/e$ [12, 1]. Nothing better is known.

Another stochastic model is the *random order model* where we assume that I is unknown, but impressions in I arrive in a random order. This has proved be an important analytical construct for other problems such as secretary-type problems where worst cases are inherently difficult. It is known that in this case even the greedy algorithm has a (tight) competitive ratio of $1 - \frac{1}{e}$ [12]. Further, no deterministic algorithm can achieve approximation ratio better than 0.75 and no randomized algorithm better than 0.83 [12]. Currently the best known approximation ratio remains $1 - 1/e$.

Can one beat the $1 - 1/e$ bound? We address this main question.

1.1 Our Results and Techniques. We present two results for the online stochastic bipartite matching problem under the *iid* model.

- We present an algorithm with an approximation factor of $\frac{1 - \frac{2}{e^2}}{\frac{4}{3} - \frac{2}{3e}} \simeq 0.67$, breaking past the $1 - 1/e$ bottleneck. We also show that our analysis is tight, by providing an example for which our algorithm achieves exactly this factor.

- We show that there is no $1 - o(1)$-approximation algorithm for this problem. Specifically, we show that any online algorithm will be off by at least $26/27$ (or $\approx .99$ if one requires a family of instances that grows with n).

Our algorithms are based on computing an optimal offline solution, and using it to guide online allocation. An intuitive approach under this paradigm is to compute a matching M_{OFF} on the "expected graph"—that is, the one that would result if all impressions occurred exactly as many times as expected. Thereafter, one can use this matching online, that is, when node $i \in I$ arrives, match it with $a \in A$ iff $(i, a) \in M_{\text{OFF}}$. One expects this to perform well if the empirical probability of occurrence of each node $i \in I$ is very close to its value in the distribution. This can be shown if all $i \in I$ occur very frequently

---

[1]For details of this application, see Section 1.3.

[2]We give more details on this model in Section 2, including a discussion of different ways to characterize an approximation ratio in this context.

using for example the Chernoff bound. However in general, many $i \in I$ will have very low frequency. In this paper, we show that this first attempt achieves (you guessed it) $1 - 1/e$, and this is tight.

To get our main result and beat $1 - 1/e$, we compute *two* disjoint offline solutions and use them as follows: when a request arrives, we try to assign it based on the first offline solution, and if that assignment fails, we try the second. In order to identify these two disjoint offline solutions, we solve a max flow problem in a boosted flow graph, and then carefully decompose this maximum flow to two edge-disjoint (near)-matchings. Other than guiding the online decision making, these offline solutions are used to characterize an upper bound for the optimum in each scenario. This bound is determined by identifying an appropriate cut in each scenario that is guided by a cut in the offline solution. This is the main technical part of the analysis, and we hope this technique proves useful for analyzing heuristic algorithms for other stochastic optimization problems.[3]

The idea of using two solutions is inspired by the idea of power of two choices in online load balancing [3, 17]. Power of two choices has traditionally meant choosing between two *random* choices for online allocation; in contrast, we use two *deterministic* choices, carefully computed offline to guide online allocation.[4]

Our results are somewhat more general as shown in the technical sections, and the problem itself was motivated from an Internet ad application described later.

1.2 Other Related Work. Our online stochastic matching problem is an example of online decision making problems studied in the Operations Research literature as stochastic approximate dynamic programming problems [5, 4, 8, 10]. Several heuristic methods have been proposed for such problems (e.g., see Rollout algorithms for stochastic dynamic programming in [4]), but we are not aware of any rigorous analysis of the performance of the heuristics. Recently other online stochastic combinatorial optimization problems like Steiner tree and set cover problems have been studied in the *iid* model [13, 11]; one can achieve an approximation factor better than the best bound for the adversarial online variant.

A related ad allocation problem is the *Adwords assignment* problem [16] that was motivated by sponsored search auctions. When modeled as an online bipartite assignment problem, here, each edge has a *weight*, and there is a *budget* on each ad representing the upper bound on the total weight of edges that may be assigned to it. In the offline setting, this problem is NP-Hard, and several approximations have been designed [7, 19, 2]. For the online setting, it is typically assumed that every weight is very small compared to the corresponding budget, in which case there exist $1 - 1/e$ factor online algorithms [16, 6, 12, 1]. Recently, it has been brought to our attention that an online algorithm [9] gives a $1 - \epsilon$-approximation, for any ϵ, for Adwords assignment when *opt* is larger than $O(\frac{n^2}{\epsilon^3})$ times each bid in the iid and random permutation models. Thus, technically, our problem is different from their problem in two ways: the edges are unweighted (making it easier), but OPT is not necessarily much larger than each bid (making it harder – in the bipartite graph case, OPT can be $O(n)$). Moreover, our offline problem is solvable in polynomial time, and we show that no $1 - \epsilon$-approximation can be achieved for our problem for some fixed ϵ. In fact, their algorithm, along with other previously studied algorithms (e.g, algorithms based on greedy, greedy bid-scaling, and primal-dual techniques) does not achieve a factor better than $1 - \frac{1}{e}$ for our problem, and we beat $1 - \frac{1}{e}$ factor using a different technique. An interesting related model for combining stochastic-based and online solutions for the Adwords problem is considered in [15], but their approach does not give an improved approximation algorithm for the *iid* model.

---

[3]For example, this technique might be applicable for proving performance guarantees for heuristics for approximate dynamic programming problems studied in the OR literature [5, 4, 8, 10].

[4]Previously, power of two choices has been used in various congestion control and load balancing settings. Our work is a novel adaptation of this idea to a stochastic bipartite matching setting.

1.3 Applied Motivation: Display Ad Allocation. Our motivation is in part applied and arises from allocation of "display ads" on the Internet. Here is a high level view. Websites have multiple pages (e.g., sports, real estate, etc), and several slots where they can display ads (say an image or video or a block of text). Each user who views one of these pages is shown ads, i.e., the ads get what is called an "impression." Advertisers pay the website per impression and buy them (typically in lots of one thousand) ahead of time, often specifying a subset of pages on which they would like their ad to appear, or a type of user they wish to target. All such sales are entered into an ad delivery system (ADS).

Since the ADS serves ads on the same web pages from day to day, they have an idea of the traffic that occurs on these websites. While there are inaccuracies and indeed it is nearly impossible to forecast the number of viewers of a webpage in the future, it is standard industry practice to use these estimates at the time of selling inventory to various advertisers (to judge whether a new sale can be accommodated).

When a user visits one of the pages, the ADS determines the set of eligible ads for that slot, and selects an ad to be shown. Since not all ads are suitable for each page or slot, we have an online (in two senses of the word) bipartite matching scenario. The ADS would like to maximize the number of impressions that are filled with ads in order to satisfy their contracts, and thus maximize their revenue.

The underlying problem is an online bipartite matching problem in the *iid* model. Each $i \in I$ is an "impression type," which may represent a particular web page, or even a cross product of targeting criteria (location, demographic, etc.). Edges (a, i) then capture the fact that advertiser a was interested in an impression of type i. Using past traffic data, the ADS defines e_i to be the typical number of impressions they get of type i. Then, the distribution $\mathcal{D}$ over I is given by $\Pr[i^*] = \frac{e_{i^*}}{\sum_i e_i}$.

In contrast to sponsored search, the display ad business is easier to model, since currently, display ads are not sold via auctions, and prices are the same for different impressions of an advertiser (so we do not need to worry about the underlying auction pricing schemes). Differing values of ad slots to different advertisers is handled exogenously via sales contracts, and the online problem is just to assign edges to meet the contracted sales. Still, we note that there are many aspects of online ad serving that deserve a richer model than the one we give here, and indeed there is more work to be done in this area. For example, the ADS may want to maximize the value of the contracts fulfilled, rather than the total number of impressions, or may want to maximize some notion of quality of ads served. One extension that we address is frequency capping, which we discuss in the conclusion. As such, display ad selection problems are solved routinely by ADSs, and any insights or solutions we develop for our problem are likely to be useful in practice.

2 Preliminaries

Consider the following online stochastic matching problem in the i.i.d model: We are given a bipartite graph $G = (A, I, E)$ over advertisers A and impression types I. Let $k = |A|$ and $m = |I|$. We are also given, for each impression type $i \in I$, an integer number e_i of impressions we expect to see. Let $n = \sum_{i \in I} e_i$. We use $\mathcal{D}$ to denote the distribution over I defined by $\Pr[i] = e_i/n$.

An instance $\Gamma = (G, \mathcal{D}, n)$ of the *online stochastic matching* problem is as follows: We are given offline access to G and the distribution $\mathcal{D}$. Online, n i.i.d. draws of impressions $i \sim \mathcal{D}$ arrive, and we must immediately assign ad impression i to some advertiser a where $(a, i) \in E$, or not assign i at all. Each advertiser $a \in A$ may only be assigned at most once[5]. Our goal is to assign arriving impressions to advertisers and maximize the total number of assigned impressions. In the following, we will formally define the objective function of the algorithm.

---

[5]All results in this paper hold for a more general case that each advertiser a has a capacity c_a and advertiser a can be assigned at most c_a times. This more general case can be reduced easily to the degree one case by repeating each node a c_a number of times in the instance.

3

Let $D(i)$ be the set of draws of impression type i that arrive during the run of the algorithm. We let a scenario $\hat{I} = \cup_{i \in I} D(i)$ be the set of impressions. Let $\hat{G}(\hat{I})$ be the "realization" graph, i.e., with node sets A and $\hat{I}$, and edges $\hat{E} = \{(a, i') : (a, i) \in E, i' \in D(i)\}$.

Given an instance $\Gamma = (G, \mathcal{D}, n)$ of the online matching problem, we wish an algorithm ALG for which for any instance Γ of the online matching problem, with high probability $\frac{\text{ALG}(\hat{I})}{\text{OPT}(\hat{I})} \geq \alpha$. In this case, we say that the algorithm achieves approximation factor α with high probability. One could also study weaker notions of approximation, namely $\frac{E[\text{ALG}(\hat{I})]}{E[\text{OPT}(\hat{I})]}$ (the approximation factor in expectation), or $E[\frac{\text{ALG}(\hat{I})}{\text{OPT}(\hat{I})}]$ (the expected approximation factor). Note that if one proves a high-probability factor of α, it implies an approximation factor in expectation, and an expected approximation factor of at least $\alpha - o(1)$.

2.1 Balls in Bins. In this section we characterize two useful extensions of the standard balls-in-bins problem, where we are interested in the distribution of certain functions of the bins. We characterize the expectations of these functions, and use Azuma's inequality on appropriately defined Doob's Martingales to establish concentration results as needed. In particular, we will use the following facts. The proofs are left to the appendix.

Fact 1. *Suppose n balls are thrown into n bins, i.i.d. with uniform probability over the bins. Let B be a particular subset of the bins, and S be a random variable that equals the number of bins from B with at least one ball. With probability at least $1 - 2e^{-\epsilon n/2}$, for any $\epsilon > 0$, we have $|B|(1 - \frac{1}{e}) - \epsilon n \leq S \leq |B|(1 - \frac{1}{e} + \frac{1}{en}) + \epsilon n$*

Fact 2. *Suppose n balls are thrown into n bins, i.i.d. with uniform probability over the bins. Let $B_1, B_2, \ldots, B_\ell$ be ordered sequences of bins, each of size c, where no bin is in more than d such sequences. Fix some arbitrary subset $\mathcal{R} \subseteq \{1, \ldots, c\}$. We say that a bin sequence $B_a = (b_1, \ldots, b_c)$ is "satisfied" if (i) at least one of its bins b_i with $i \notin \mathcal{R}$ has at least one ball in it; or, (ii) at least one of its bins b_i with $i \in \mathcal{R}$ has at least two balls in it. Let S be a random variable that equals the number of satisfied bin sequences. With probability at least $1 - 2e^{-\epsilon^2 n/2}$, we have $S \geq \ell(1 - \frac{2^{|\mathcal{R}|}}{e^c}) - \epsilon dn - \frac{2^{|\mathcal{R}|}c^2}{e^c} \frac{\ell}{n-c^2}$.*

3 Hardness

In this section, we show that the expected approximation factor of every (randomized) online algorithm is bounded strictly away from 1.

Consider the 6-cycle G defined by $A = \{a, b, c\}$, $I = \{x, y, z\}$, and $E = \{(x, a), (y, a), (y, b), (z, b), (z, c), (x, c)\}$. The distribution $\mathcal{D}$ is the uniform distribution $(1/3, 1/3, 1/3)$ on I, and $n = 3$. We show that no (randomized) algorithm can achieve an expected approximation factor better than $26/27$ on this instance. Without loss of generality (from the symmetry of the 6-cycle), assume that the first impression to arrive is x and that it gets assigned to advertiser a. Now, if the next two arrivals are both of impression y, then any algorithm will only be able to assign one of these. The optimal assignment for the scenario (x, y, y) is to assign x to c, and the two y impressions to a and b. Since the probability of (x, y, y) is $1/9$, the expected approximation factor is at most $(1/9)(2/3) + (8/9)1 = 26/27$.

To get a family of instances on which no algorithm can do better than a constant bounded away from 1, we will have to construct an instance consisting of a large number k copies of 6-cycles. Using this idea, we can prove the following theorem. The details of the proof are left to the appendix.

Theorem 3. *There is an instance of the online stochastic matching problem in which no algorithm can achieve an expected approximation factor better than $\frac{26}{27}$. Moreover, there exists a family of instances with $n \to \infty$ for which no algorithm can achieve an expected approximation of $1 - o(1)$.*

4   Offline Algorithms for Online Matching

In this section, we present our improved online algorithms guided by offline solutions. Before stating the improved approximation result, we "warm up" with a simple, natural algorithm that uses the idea of computing an offline solution to "guide" our online choices. This algorithm will only achieve a $1 - \frac{1}{e}$-approximation (which is tight). The proof of this part illustrates the framework we will use in the second section to beat $1 - \frac{1}{e}$; however we will need a new idea to achieve this—namely, the use of a *second* offline solution.

4.1   "Suggested Matching" Algorithm: a $1 - \frac{1}{e}$-Approximation. The *suggested matching* algorithm is a first attempt at the approach of using an offline solution for online matching. In this algorithm, we simply find a maximum matching in the graph we "expect" to arrive, then restrict our online choices to this matching.

Offline Algorithm. We will describe this algorithm more formally in terms of the standard characterization of b-matching as a max-flow problem, since we will later use this flow graph explicitly to bound OPT. Given an instance $\Gamma = (G(A, I, E), \mathcal{D}, n)$ of the problem, we will find a max-flow in a graph G_f constructed from G as follows: define a new source node s and an edge (s, a) with capacity 1 to all $a \in A$, direct all edges in E from A to I, and add a sink node t with edges (i, t) from all $i \in I$ with capacity e_i. Let $f_{ai} \in \{0, 1\}$ be the flow on edge (a, i) in this max flow (since all the capacities are integers, we may assume that the resulting flow is integral [18]). For ease of notation, we say $f_{ai} = 0$ if edge $(a, i) \notin E$.

Online Algorithm. When an impression $i' \in D(i)$ arrives online, we choose a random ad a' according to the distribution defined by the flow; i.e., the probability of choosing a' is $\frac{f_{a'i}}{e_i}$. (Note that if $\sum_a f_{ai} < e_i$ there is some probability that no a' is chosen.) If a' is already taken, we do not match i to any ad.[6]

Bounding ALG. The performance of this algorithm is easily characterized with high probability in terms of the computed max-flow. Define $F_a = \sum_i f_{ai}$, and note that $F_a \in \{0, 1\}$; this indicates whether ad a was chosen in the max flow. Let $A^* = \{a \in A : F_a = 1\}$. When an impression $i \in \hat{I}$ arrives online, a particular ad $a : f_{a,i} = 1$ has probability $1/e_i$ of being chosen by the online algorithm; since each impression i has probability e_i/n of arriving, we conclude that each $a \in A^*$ has probability $1/n$ of being chosen by the online algorithm upon each arrival. Thus, to bound the total number of ads chosen we have a balls-in-bins problem with n balls and n bins, and we are interested in lower-bounding the number of bins (among a subset of size $|A^*|$) that have at least one ball. Applying concentration results for balls-in-bins (Fact 1), we get that with probability $1 - e^{-\Omega(n)}$, ALG $\geq (1 - \frac{1}{e})|A^*| - \epsilon n$.

Bounding OPT. To bound the optimal solution, we will construct a cut in the realization graph $\hat{G} = (A, \hat{I}, \hat{E})$ using a min-cut of G_f (constructed using the max-flow found by the algorithm) as a "guide." Let (S, T) be a min $s-t$ cut in the graph G_f using the canonical "reachability" cut in G_f; i.e., S is defined as the set of nodes reachable from s using paths in the residual graph after sending the flow f found by the algorithm. This is always a min-cut.[18] Let $A_S = A \cap S$ and define A_T, I_S and I_T similarly.

We claim that there are no edges in E from A_S to I_T; suppose there is such an edge (a, i). Then, a much be reachable from s since $a \in S$, but i must not be reachable since $i \in T$. This implies that there is no residual capacity along (a, i); i.e., $f_{ai} = 1$. However this also implies that there is no residual capacity along (s, a) since (s, a) is the only edge entering a and it has capacity 1, and that there is no other flow leaving a. This implies that a is not reachable in the residual graph, a contradiction. Thus the only edges in the cut (S, T) are from s to A_T (capacity 1) and from $i \in I_S$ to t (capacity e_i). We may conclude using max-flow min-cut that $|A^*| = \sum_a F_a = |A_T| + \sum_{i \in I_S} e_i$.

---

[6]Clearly, making an arbitrary available match is always as good (and in some cases better) than doing nothing; we present the algorithm this way for ease of presentation.

Now consider the "realization" graph $\hat{G} = (A, \hat{I}, \hat{E})$, and define a max-flow instance $\hat{G}_f$ whose solution has size equal to the maximum matching in $\hat{G}$; i.e., create a source s with edges to all $a \in A$, direct edges of $\hat{E}$ toward $\hat{I}$, and create a sink t with edges from all $i' \in \hat{I}$. Set the capacity of every edge to one. Note that any $s-t$ cut in $\hat{G}_f$ is a bound on OPT.

We define an $s-t$ cut in $(\hat{S}, \hat{T})$ in $\hat{G}_f$ as follows. Let $\hat{I}_S = \cup_{i \in I_S} D(i)$ and $\hat{I}_T = \cup_{i \in I_T} D(i)$. Define $\hat{S} = A_S \cup \hat{I}_S$ and $\hat{T} = A_T \cup \hat{I}_T$. Note that since there are no edges from A_S to I_T in G_f, there are also no edges from A_S to $\hat{I}_T$ in $\hat{G}_f$. Thus the size of the cut $(\hat{S}, \hat{T})$ is equal to $|\hat{I}_S| + |A_T|$. An online impression ends up in the set $\hat{I}_S$ with probability $\sum_{i \in I_S} e_i / n$, independent of the other impressions. Using a Chernoff bound, we can conclude that for any $\epsilon > 0$, with probability $1 - e^{-\Omega(n)}$ (over the scenarios), the size of the cut (and therefore OPT) obeys $\text{OPT} \leq |A_T| + \sum_{i \in I_S} e_i + \epsilon n = |A^*| + \epsilon n$.

Tightness of the Analysis. Consider a special case of the online matching problem $\Gamma(G, \mathcal{D}, n)$ where $e_i = 1$ for each $i \in I$ and the underlying graph G is a complete bipartite graph. The algorithm will find a perfect matching between I and A, and so each ad is matched with probability at least $1 - \frac{1}{e}$. Using Fact 1, the algorithm achieves $\approx (1 - \frac{1}{e})n$ with high probability. However, the optimum is n. Therefore:

Theorem 4. *The approximation factor of the* suggested matching *algorithm is* $1 - \frac{1}{e}$ *with high probability, and this is tight, even in expectation.*

4.2 "Two Suggested Matchings" (TSM) Algorithm: Beating $1 - \frac{1}{e}$.
To improve upon the *suggested matching* algorithm, we will instead use *two* disjoint (near-)matchings to guide our online algorithm. To find these matchings, we boost the capacities of the flow graph and then decompose the resulting solution into disjoint solutions. The second solution allows to to break the $1 - \frac{1}{e}$ barrier and prove:

Theorem 5. *For any $\epsilon > 0$, with probability at least $1 - e^{-\Omega(n)}$, as long as $\text{OPT} = \Omega(n)$, the* two suggested matchings *algorithm achieves approximation ratio*

$$\frac{\text{ALG}}{\text{OPT}} - \epsilon \geq \alpha := \frac{1 - \frac{2}{e^2}}{\frac{4}{3} - \frac{2}{3e}} \approx 0.67029 > 1 - \frac{1}{e}.$$

Moreover, this ratio is tight; specifically, there is a family of instances for which the two suggested matchings *algorithm has expected approximation factor at most $\alpha + \epsilon$.*

Throughout the section, until the final proof of Theorem 5, we assume $e_i = 1$ for all $i \in I$, which also implies $m = n$. Extending to integer e_i is a simple reduction to this case.

4.2.1 The TSM Algorithm.
In this algorithm, we construct a *boosted flow graph* G_f, built from G in the standard reduction of matching to max-flow; i.e., create a source s with edges to all $a \in A$, direct the edges of G towards nodes in I, and create a sink t with edges from all $i \in I$. However, we set the capacities of the edges differently than in the max-flow reduction: (i) Edges (s, a) from the source get capacity 2, (ii) edges $(a, i) \in E$ get capacity 1, and (iii) edges (i, t) from I to t get capacity 2.

We find a max-flow in this graph from s to t. Since all the capacities are integers, we may assume that the resulting flow is integral [18]. Let E_f be the set of edges $(a, i) \subseteq E$ with non-zero flow on them, which must be unit flow. Since the capacities of edges (s, a) and (i, t) are all 2, we know that the graph induced by E_f is a collection of paths and cycles. Using this structure, we assign colors blue and red to the edges of E_f as follows:

- Color the cycle edges alternating blue and red.
- Color the edges of the odd-length paths alternating blue and red, with more blue than red.

6

- For the even-length paths that start and end with nodes $a \in A$, alternate blue and red.
- For the even-length paths that start and end with impressions $i \in I$, color the first two edges blue, and then alternate red, blue, red, blue, etc., ending in blue.

Note that all $i \in I$ are incident to either no colored edges, one blue edge, or a blue and a red edge.

The TSM algorithm for serving online ad impressions is simple: For each $i \in I$, the first time i arrives try the blue edge; the second time i arrives try the red edge. More formally, for all $i \in I$ maintain a count x_i of the number of impressions $i' \in D(i)$ that have arrived so far. When $i' \in D(i)$ arrives: if $x_i = 0$, set a' to be the ad along i's blue edge (if i has a blue edge); if $x_i = 1$, set a' to be the ad along i's red edge (if i has a red edge). Now assign i to a' if a' is unassigned. If this a' is already assigned, or if $x_i > 1$, do not make an assignment.[7]

4.2.2 Performance of the TSM Algorithm.
To analyze the performance of this algorithm, we first derive a lower bound on the number of ads assigned during the run of the algorithm. We do so in terms of the incidence pattern of the different ads with respect to the edges E_f. Specifically, let A_{BR} be the ads that are incident to a blue and a red edge, and A_{B} be the ads that are incident to only a blue edge. Similarly define A_{BB} and A_{R}. We have

$$|E_f| = 2A_{\mathrm{BR}} + 2A_{\mathrm{BB}} + A_{\mathrm{B}} + A_{\mathrm{R}}. \tag{1}$$

Consider some $a \in A_{\mathrm{B}}$ with blue edge (a, i). The event that a is ever chosen is exactly the event that some $i' \in D(i)$ is ever drawn from D, since then we will choose a (and no other impression will choose a). Since $e_i = 1$, this is exactly the probability that a particular bin is non-empty in a balls-in-bins problem with n balls (the online impressions), and $m = n$ bins (the impression types I). Applying Fact 1, we get that with high probability the number of ads chosen from A_{B} is at least $|A_{\mathrm{B}}|(1 - \frac{1}{e}) - \epsilon n$. Now consider some $a \in A_{\mathrm{BR}}$ with blue edge (a, i_b) and red edge (a, i_r). If $|D(i_b)| \geq 1$, or if $|D(i_r)| \geq 2$, then a will definitely be chosen. Thus we can apply Fact (2) with n balls, $m = n$ bins, $c = 2$, bin sequences equal to the neighborhood sets of A_{BR} along the blue and red edges (ordered blue, red), $d = 2$ (since each impression is incident to at most 2 edges of E_f), and $\mathcal{R}$ set to the second (red) bin of the bin sequence. We conclude that with high probability, the number of ads chosen from A_{BR} is at least $|A_{\mathrm{BR}}|(1 - \frac{2}{e^2}) - \epsilon n$. Similar reasoning gives bounds with coefficients of $(1 - \frac{1}{e^2})$ for A_{BB} and $(1 - \frac{2}{e})$ for A_{R}. We may conclude that with high probability (over the scenarios), $\mathrm{ALG} \geq (1 - \frac{1}{e^2})|A_{\mathrm{BB}}| + (1 - \frac{2}{e^2})|A_{\mathrm{BR}}| + (1 - \frac{1}{e})|A_{\mathrm{B}}| + (1 - \frac{2}{e})|A_{\mathrm{R}}| - 4\epsilon n$. Note that since $|A_B| \geq |A_R|$, we can also assert

$$\mathrm{ALG} \geq (1 - \tfrac{1}{e^2})A_{\mathrm{BB}} + (1 - \tfrac{2}{e^2})A_{\mathrm{BR}} + (1 - \tfrac{3}{2e})(A_{\mathrm{B}} + A_{\mathrm{R}}) - 4\epsilon n. \tag{2}$$

4.2.3 Bound on the optimal solution.
Let (S, T) be a particular min $s - t$ cut of the flow graph G_f defined as follows. First start with the canonical "reachability" min $s - t$ cut of the flow graph G_f, where S is defined as the set of nodes reachable from s in the residual graph G_f left after finding the max-flow E_f. Then, we do a small bit of "surgery" to this cut: for all $i \in I \cap T$, if i is incident to more than one $a \in A \cap S$, we move i over to S. Note that this does not increase the value of the cut, since we save at least 2 for the two edges from $A \cap S$, and pay exactly 2 for the edge (i, t). Let $A_S = A \cap S$, and define A_T, I_S and I_T similarly. Let E_δ be the set of edges $(a, i) \in E$ that cross the cut (from A_S to I_T).

Some observations: **(i)** We have $E_\delta \subseteq E_f$, since otherwise, if some $(a, i) \in E_\delta$ has no flow across it, then i would be reachable from s, and would not be in the set T. (And we did not introduce any such edges in our surgery.) **(ii)** All $i \in I_T$ have at most one incident edge in E_δ (follows from the surgery). **(iii)**

---

[7]A slight improvement to this algorithm is to try to match along the red edge if matching along the blue edge fails; we do not make use of this in the analysis so we leave it out for clarity.

All $a \in A_S$ have at most one incident edge in E_δ. To see this, suppose it had two such edges (it cannot have more than 2 since $E_\delta \subseteq E_f$). Then, since a is reachable from s (since it is in S), it must have either residual capacity from s directly, or residual capacity from I_S; but it cannot have either, since (s, a) is saturated and both flow edges from a go to I_T.

Let A_δ, I_δ be the ads and impressions, respectively, that are incident to edges in E_δ. We may conclude from the observations above that the graph $(A_\delta, I_\delta, E_\delta)$ induced by E_δ is a matching. The min-cut of G_f is made up of the edges E_δ, the $|A_T|$ edges from s to A_T (with capacity 2), and the $|I_S|$ edges from I_S to t (also capacity 2). Thus, by max-flow-min-cut, we have

$$|E_f| = 2(|A_T| + |I_S|) + |E_\delta|. \tag{3}$$

We are interested in bounding the value of the optimal matching in the realization graph $\hat{G} = (A, \hat{I}, \hat{E})$. To do this, we will use the min-cut (S, T) of the graph G_f as a "guide" to construct a (not necessarily min) cut in a flow graph built from $\hat{G}$, and prove a high-probability bound on the size of this cut.

More precisely, we let $\hat{G}_f$ be a directed version of $\hat{G}$, constructed as before with a source and a sink, and edges corresponding to $\hat{G}$; but now we put capacity 1 on all edges. Note that any $s-t$ cut in this graph constitutes an upper bound on OPT, the maximum matching in $\hat{G}$. We construct such a cut $(\hat{S}, \hat{T})$ as follows. We let $\hat{I}_S = \cup_{i \in I_S} D(i)$ and $\hat{I}_T = \cup_{i \in I_T} D(i)$. For the ads, we will use almost the same partition (A_S, A_T) as in G_f but we will perform some "surgery" on this partition as well. Let $A_\delta^* \subseteq A_\delta$ be the set of ads $a \in A$ that are incident (in $\hat{G}_f$) to some $i' \in D(i) \subseteq \hat{I}_T$. Note that $i \in I_\delta$ and $(a, i) \in E_\delta$. We set $\hat{S} = \hat{I}_S \cup (A_S \setminus A_\delta^*)$ and $\hat{T} = \hat{I}_T \cup A_T \cup A_\delta^*$.

Now we will measure the size of the cut $(\hat{S}, \hat{T})$ in G_f. We pay 1 for each $a \in \hat{I}_S$, $i \in A_T$ and $i \in A_\delta^*$. But note that there are no edges in G_f from $A \cap \hat{S}$ to $\hat{I}_T$, since we got rid of them in our surgery. Thus we have OPT $\leq |\hat{I}_S| + |A_T| + |A_\delta^*|$.

Using a Chernoff bound, with probability $1 - e^{-\Omega(n)}$ we have $|\hat{I}_S| \leq |I_S| + \epsilon n$ for any $\epsilon > 0$. To bound $|A_\delta^*|$, consider some $a \in A_\delta$, and the impression $i \in I_\delta$ along the edge (a, i) in the matching $(A_\delta, I_\delta, E_\delta)$. The ad a appears in A_δ^* iff impression i is drawn during the run of the algorithm. Thus we have a balls-in-bins problem with n balls, $m = n$ bins, uniform bin probabilities and a bin subset of size $|A_\delta^*|$, and we are concerned with an upper bound on the number of bins in that subset that get at least one ball. Using Fact 1 we may conclude that with high probability $|A_\delta^*| \leq (1 - \frac{1}{e})|E_\delta| + \epsilon n + O(1)$.

Summarizing the previous arguments, we get, for any $\epsilon > 0$, with probability $1 - e^{-\Omega(n)}$, OPT $\leq |I_S| + |A_T| + (1 - \frac{1}{e})|E_\delta| + \epsilon n$. Applying Equations (3) then (1), we get

$$\begin{aligned} \text{OPT} \quad &\leq \quad \tfrac{1}{2}|E_f| + (\tfrac{1}{2} - \tfrac{1}{e})|E_\delta| + \epsilon n \\ &= \quad |A_{\text{BR}}| + |A_{\text{BB}}| + \tfrac{1}{2}(|A_{\text{B}}| + |A_{\text{R}}|) + (\tfrac{1}{2} - \tfrac{1}{e})|E_\delta| + \epsilon n \end{aligned} \tag{4}$$

In order to use this bound on OPT together with the bound on ALG in Equation 2, we must bound the size of E_δ in terms of the sets A_{BR}, A_{BB}, A_{B} and A_{R}. The following lemma takes a deeper look at the two matchings constructed by the algorithm, and their relationship to the min-cut (S, T) in G_f, in order to achieve this bound.

Lemma 1. $|E_\delta| \leq \frac{2}{3}|A_{\text{BR}}| + \frac{4}{3}|A_{\text{BB}}| + |A_{\text{B}}| + \frac{1}{3}|A_{\text{R}}|$.

Proof. It suffices to show that the inequality holds for every connected component (path or cycles) of the graph induced by E_f. We thus assume notationally that the graph induced by E_f consists of a single such connected component.

Consider an arbitrary pair of edges $(a_1, i_1), (a_2, i_2) \in E_\delta \subseteq E_f$. Since the edges of E_δ are independent, (a_1, i_1) and (a_2, i_2) cannot occur consecutively in this component (path or cycle); we claim further

that (a_1, i_1) and (a_2, i_2) must have at least two edges between them. Suppose not, then wlog (a_2, i_1) is in the component; but since $a_2 \in A_S$ and $i_1 \in I_T$ (by the definition of E_δ) we must have $(a_2, i_1) \in E_\delta$, contradicting the fact that the edges E_δ are independent.

If the component is a cycle of length k, we can use the reasoning above to conclude that there are at most $\lfloor \frac{k}{3} \rfloor$ edges of E_δ in the cycle. The ads in the cycle are all in A_{BR} and there are exactly $\frac{k}{2}$ of them. Thus $|E_\delta| \leq \frac{2}{3}|A_{\mathrm{BR}}|$, which implies the inequality.

If the component is a path, we can conclude that $|E_\delta| \leq \lceil \frac{k}{3} \rceil$ by the reasoning above—the worst case is when the path starts and ends in a E_δ edge. We have three cases for this path, depending on the parity of its length, and (in the case of even-length paths) whether it starts and ends in A or I.

- For odd paths of length k, by construction of the edge colors, we have one ad in A_{B} and $\frac{k-1}{2}$ ads in A_{BR}. Thus $|E_\delta| \leq \lceil \frac{k}{3} \rceil = \lceil \frac{2|A_{\mathrm{BR}}|}{3} + \frac{1}{3} \rceil \leq \frac{2}{3}|A_{\mathrm{BR}}| + 1 = \frac{2}{3}|A_{\mathrm{BR}}| + |A_{\mathrm{B}}|$.
- For even paths of length k that start and end with ads, we have $|A_{\mathrm{B}}| = 1$, $|A_{\mathrm{R}}| = 1$ and $|A_{\mathrm{BR}}| = \frac{k}{2} - 1$. Thus $|E_\delta| \leq \lceil \frac{k}{3} \rceil = \lceil \frac{2|A_{\mathrm{BR}}|}{3} + \frac{2}{3} \rceil$. We bound this using a case analysis on $|A_{\mathrm{BR}}| \mod 3$, as follows: **(i)** If $|A_{\mathrm{BR}}| \equiv 0 \mod 3$ then we get $|E_\delta| \leq \frac{2}{3}|A_{\mathrm{BR}}| + 1$. **(ii)** If $|A_{\mathrm{BR}}| \equiv 1 \mod 3$ then we get $|E_\delta| \leq \frac{2}{3}|A_{\mathrm{BR}}| + \frac{4}{3}$. **(iii)** If $|A_{\mathrm{BR}}| \equiv 2 \mod 3$ then we get $|E_\delta| \leq \frac{2}{3}|A_{\mathrm{BR}}| + \frac{2}{3}$. In all cases this is less than $\frac{2}{3}|A_{\mathrm{BR}}| + \frac{4}{3} = \frac{2}{3}|A_{\mathrm{BR}}| + |A_{\mathrm{B}}| + \frac{1}{3}|A_{\mathrm{R}}|$.
- For even length paths that start and end in impressions, we have $|A_{\mathrm{BB}}| = 1$ and $|A_{\mathrm{BR}}| = \frac{k}{2} - 1$. As in the previous case we can say $|E_\delta| \leq \lceil \frac{k}{3} \rceil = \lceil \frac{2|A_{\mathrm{BR}}|}{3} + \frac{2}{3} \rceil$, and reason by the same case analysis that $|E_\delta|$ is at most $\frac{2}{3}A_{\mathrm{BR}} + \frac{4}{3}$. This is equal to $\frac{2}{3}A_{\mathrm{BR}} + \frac{4}{3}A_{\mathrm{BB}}$. $\qquad \square$

4.2.4   Proof of Theorem 5. We first prove the approximation ratio for $e_i = 1$. The bounds in equations (2) and (4) each hold with probability $1 - e^{-\Omega(n)}$, and so using a union bound they both hold with probability $1 - e^{-\Omega(n)}$. Using Lemma 1 (ignoring the $\frac{1}{3}$ in front of the A_R) and Equation (4), we get

$$\mathrm{OPT} \leq \left(\frac{4}{3} - \frac{2}{3e}\right)A_{\mathrm{BR}} + \left(\frac{5}{3} - \frac{4}{3e}\right)A_{\mathrm{BB}} + \left(1 - \frac{1}{e}\right)(A_{\mathrm{B}} + A_{\mathrm{R}}) + \epsilon' n.$$

Since $\mathrm{OPT} = \Omega(n)$, we can choose ϵ' small enough such that when we apply Equation (2) (also using ϵ') we may conclude

$$\frac{\mathrm{ALG}}{\mathrm{OPT}} + \epsilon \geq \min\left\{\frac{1 - \frac{1}{e^2}}{\frac{5}{3} - \frac{4}{3e}}, \frac{1 - \frac{2}{e^2}}{\frac{4}{3} - \frac{2}{3e}}, \frac{1 - \frac{3}{2e}}{1 - \frac{1}{e}}\right\} = \min\{.735\ldots, .670\ldots, .709\ldots\} = \frac{1 - \frac{2}{e^2}}{\frac{4}{3} - \frac{2}{3e}} \approx .670.$$

The tightness of this analysis is proved in Section 4.2.5. For arbitrary integer e_i, we give a reduction to the case $e_i = 1$. Given a set of instance $\Gamma = (G, \mathcal{D}, n)$, we reduce to a new instance $\Gamma' = (G', \mathcal{D}', n)$ with $e_i' = 1$ by making e_i copies of each impression type i. Then, when an impression of type i arrives online, "name" it randomly according to one of its copies. The resulting distribution $\mathcal{D}'$ is uniform over the impression types I' in the new instance.

Let $\hat{I}$ be the impressions that are drawn from $\mathcal{D}$ in one run of the algorithm, and let $\hat{I}'$ be the resulting draws from $\mathcal{D}'$. By the arguments above, we achieve the desired bound on $\mathrm{ALG}/\mathrm{OPT}'$ with high probability, where OPT' is with respect to $\hat{I}'$; however we have $\mathrm{OPT}' = \mathrm{OPT}$, since the realization graphs $\hat{G} = (A, \hat{I}', \hat{E}')$ and $\hat{G}' = (A, \hat{I}, \hat{E})$ are in fact the same graph.

4.2.5   Tightness of the analysis for the TSM Algorithm. In this section we demonstrate a family of instances for which the TSM algorithm achieves a factor no better than $\frac{1 - 2/e^2}{4/3 - 2/(3e)}$, thus showing that the analysis in Section 4 is tight.

The family is parameterized by n, which is the number of advertisers, the number of impression types, as well as the number of impression arrivals. We shall take n to be a multiple of 4. The set A of advertisers consists of the following parts: a set K of size $\frac{n}{4}$ and, for $i \in [1, \frac{n}{4}]$, advertisers $\{u_i, v_i, w_i\}$. The set I of impressions consists of the following parts: a set L of size $\frac{n}{4}$ and, for $i \in [1, \frac{n}{4}]$, impressions $\{x_i, y_i, z_i\}$. Define $U = \{u_i : i \in [1, n/4]\}$, and similarly, V, W, X, Y, Z. Thus $A = K \cup U \cup V \cup W$ and $I = L \cup X \cup Y \cup Z$. Draws are from the uniform distribution on I.

The edges E are as follows: **(i)** For $i \in [1, \frac{n}{4}]$, the 6-cycle $\{u_i - x_i - v_i - y_i - w_i - z_i - u_i\}$, **(ii)** a complete bipartite graph between K and X, and **(iii)** a complete bipartite graph between L and W.

We now describe the max-flow and min-cut in G_f found during the algorithm. The only edges with (unit) flow are the edges of the 6-cycles, i.e., for $i \in [1, \frac{n}{4}]$, $\{u_i - x_i - v_i - y_i - w_i - z_i - u_i\}$. Thus all vertices in U, V, W, X, Y and Z have a flow of 2 each, and the vertices in K and L have a flow of 0. The reachability cut (S, T) obtained from this flow has $S = K \cup X \cup U \cup V \cup \{s\}$ (where s is the source vertex). The flow and the cut both have size $\frac{3n}{2}$. Using Fact 1, one can easily check that the algorithm achieves the total matching size of $\frac{3n}{4}(1 - \frac{2}{e^2})$ with high probability.

The following assignment can be made with high probability, and is a lower bound on OPT. **(i)** With high probability there will be $\frac{n}{4}$ draws of impressions from X (with repeats). These can be matched to the $\frac{n}{4}$ advertisers in K (in any order). **(ii)** With high probability there will be $\frac{n}{4}$ draws of impressions from L. These can be matched to the $\frac{n}{4}$ advertisers in W. **(iii)** With high probability there will be $(1 - \frac{1}{e})\frac{n}{4}$ unique draws of impressions from Y (counting each y_i only once, even if it is drawn multiple times). For every such y_i, its first draw is matched to v_i, and the repeat draws of y_i are left unmatched. Similarly, with high probability, there are $(1 - \frac{1}{e})\frac{n}{4}$ unique draws of z_i's, and these are matched to the corresponding u_i's. Thus, this assignment has size $\frac{n}{4} + \frac{n}{4} + (1 - \frac{1}{e})\frac{n}{2} = n(1 - \frac{1}{2e})$. This means that the TSM algorithm cannot achieve a factor better than $(1 - \frac{2}{e^2})/(\frac{4}{3} - \frac{2}{3e})$.

5 Concluding Remarks

Applying the insights to the display ads application. The approach of using the offline solution to allocate ads online may be quite useful in practice because while one can invest some time offline to find the guiding solutions, the online allocation has to be done very quickly in this application. One can use this approach to model other objective functions such as fairness in quality of ad slots assigned to ads, which may be solvable offline with some computational effort. As an example, we elaborate on the extension of our algorithm to the following problem. In the display ads business, advertisers have "frequency caps;" i.e., they do not want the same user to see their ad more than some fixed (constant) number of times. We can extend our approach here to get a $1 - 1/e$ approximation as shown in the appendix.

Generalizing the algorithm. One can generalize the two-matching algorithm to a "k-matching" algorithm by computing k matchings instead of 2 matchings, and then using them online in a prescribed order. We can easily show that if the underlying expected graph G admits k edge-disjoint perfect matchings, the approximation factor of such an algorithm is $1 - \frac{2}{e^2} \simeq 0.72$ and $1 - \frac{5}{e^3} \simeq 0.75$ for $k = 2$ and $k = 3$ respectively, however for $k = 3$, we do not know how to generalize our result for to graphs. One natural question left open by this work is what constant $c(k)$ is achieved by extending to k matchings, where $.67 \le c(k) \le .99$.

Fractional version. A theoretical version of online stochastic matching problem that may be of interest is the case in which e_i's are not necessarily integers, but arbitrary rational numbers. We observe the analysis of the "one suggested matching" algorithm can be generalized to this case, but do not know how to generalize the analysis of the "two suggested matchings" algorithm. The details are in the appendix.

10

References

[1] S. Alaei and A. Malekian. Maximizing sequence-submodular functions, manuscript. 2009.

[2] Y. Azar, B. Birnbaum, A.R. Karlin, C. Mathieu, and C.T. Nguyen. Improved Approximation Algorithms for Budgeted Allocations. In *Proceedings of the 35th international colloquium on Automata, Languages and Programming, Part I*, pages 186–197. Springer-Verlag Berlin, Heidelberg, 2008.

[3] Yossi Azar, Andrei Z. Broder, Anna R. Karlin, and Eli Upfal. Balanced allocations. *SIAM J. Comput.*, 29(1):180–200, 1999.

[4] D. Bertsekas. *Dynamic programming and optimal control*. 2007.

[5] Dimitri P. Bertsekas and David A. Castanon. Rollout algorithms for stochastic scheduling problems. *Journal of Heuristics*, 5(1):89–108, 1999.

[6] N. Buchbinder, K. Jain, and J.S. Naor. Online Primal-Dual Algorithms for Maximizing Ad-Auctions Revenue. In *Algorithms–ESA 2007 15th Annual European Symposium, Eilat, Israel, October 8-10, 2007: Proceedings*, page 253. Springer, 2007.

[7] D. Chakrabarty and G. Goel. On the approximability of budgeted allocations and improved lower bounds for submodular welfare maximization and GAP. In *Proceedings of the 2008 49th Annual IEEE Symposium on Foundations of Computer Science-Volume 00*, pages 687–696. IEEE Computer Society Washington, DC, USA, 2008.

[8] Daniela Pucci de Farias and Benjamin Van Roy. On constraint sampling in the linear programming approach to approximate dynamic programming. *Math. Oper. Res.*, 29(3):462–478, 2004.

[9] Nikhil Devanur and Thomas Hayes. The adwords problem: Online keyword matching with budgeted bidders under random permutations, manuscript.

[10] Vivek F. Farias and Benjamin Van Roy. Approximation algorithms for dynamic resource allocation. *Oper. Res. Lett.*, 34(2):180–190, 2006.

[11] Naveen Garg, Anupam Gupta, Stefano Leonardi, and Piotr Sankowski. Stochastic analyses for online combinatorial optimization problems. In *SODA*, pages 942–951, 2008.

[12] Gagan Goel and Aranyak Mehta. Online budgeted matching in random input models with applications to adwords. In *SODA*, pages 982–991, 2008.

[13] Fabrizio Grandoni, Anupam Gupta, Stefano Leonardi, Pauli Miettinen, Piotr Sankowski, and Mohit Singh. Set covering with our eyes closed. In *FOCS*, pages 347–356, 2008.

[14] R.M. Karp, U.V. Vazirani, and V.V. Vazirani. An optimal algorithm for online bipartite matching. In *Proceedings of the 22nd Annual ACM Symposium on Theory of Computing*, 1990.

[15] Mohammad Mahdian, Hamid Nazerzadeh, and Amin Saberi. Allocating online advertisement space with unreliable estimates. In *ACM Conference on Electronic Commerce*, pages 288–294, 2007.

11

[16] Aranyak Mehta, Amin Saberi, Umesh Vazirani, and Vijay Vazirani. Adwords and generalized online matching. In *FOCS*, 2005.

[17] Michael Mitzenmacher. The power of two choices in randomized load balancing. *IEEE Trans. Parallel Distrib. Syst.*, 12(10):1094–1104, 2001.

[18] Alexander Shrijver. *Combinatorial Optimization-Polyhedra and Efficiency*. 2003.

[19] A. Srinivasan. Budgeted Allocations in the Full-Information Setting. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques: 11th International Workshop, Approx 2008 and 12th International Workshop, Random 2008, Boston, Ma, USA, August 25-27, 2008*, page 247. Springer, 2008.

A Balls in Bins

In this part, we prove the concentration facts we used throughout the paper.

Fact 1. *Suppose n balls are thrown into n bins, i.i.d. with uniform probability over the bins. Let B be a particular subset of the bins, and S be a random variable that equals the number of bins from B with at least one ball. With probability at least $1 - 2e^{-\epsilon n/2}$, for any $\epsilon > 0$, we have*

$$|B|(1 - \frac{1}{e}) - \epsilon n \leq S \leq |B|(1 - \frac{1}{e} + \frac{1}{en}) + \epsilon n$$

Proof of Fact 1: We have $E[S] = \sum_{a \in B} 1 - (1 - \frac{1}{n})^n$ and so using standard identities, we obtain

$$\sum_{a \in B} 1 - e^{-1} \leq E[S] \leq \sum_{a \in B} 1 - e^{-1}(1 - \frac{1}{n}).$$

Since S, as a function of the placements of the n balls, satisfies the Lipschitz condition, we may apply Azuma's inequality to the Doob Martingale and obtain

$$\Pr[|S - E[S]| \geq \epsilon n] \leq 2e^{-\epsilon^2 n/2}.$$

$\square$

Fact 2. *Suppose n balls are thrown into n bins, i.i.d. with uniform probability over the bins. Let $B_1, B_2, \ldots, B_\ell$ be ordered sequences of bins, each of size c, where no bin is in more than d such sequences. Fix some arbitrary subset $\mathcal{R} \subseteq \{1, \ldots, c\}$. We say that a bin sequence $B_a = (b_1, \ldots, b_c)$ is "satisfied" if*

- *at least one of its bins b_i with $i \notin \mathcal{R}$ has at least one ball in it; or,*

- *at least one of its bins b_i with $i \in \mathcal{R}$ has at least two balls in it.*

Let S be a random variable that equals the number of satisfied bin sequences. With probability at least $1 - 2e^{-\epsilon^2 n/2}$, we have

$$S \geq \ell(1 - \frac{2^{|\mathcal{R}|}}{e^c}) - \epsilon dn - \frac{2^{|\mathcal{R}|}c^2}{e^c}\frac{\ell}{n - c^2}$$

Proof of Fact 2: First, we claim

$$E[S] \geq \ell\left(1 - \frac{2^{|\mathcal{R}|}}{e^c}\left(1 + \frac{c^2}{n - c^2}\right)\right)$$

To see this, fix some bin sequence B_a. The probability that B_a is not satisfied is

$$\sum_{\mathcal{R}' \subseteq \mathcal{R}} \binom{n}{|\mathcal{R}'|} |\mathcal{R}'|! \, n^{-|\mathcal{R}'|} \left(1 - \frac{c}{n}\right)^{n - |\mathcal{R}'|} \leq \sum_{\mathcal{R}' \subseteq \mathcal{R}} \left(1 - \frac{c}{n}\right)^{n-c} \leq \frac{2^{|\mathcal{R}|}}{e^c}\left(1 + \frac{c^2}{n - c^2}\right).$$

The bound on $E[S]$ follows by linearity of expectation. Now, consider S as a function of the placements of the n balls. Moving one ball can affect S by at most d, since each bin is in at most d sequences. Thus we may apply Azuma's inequality and obtain, for all $\epsilon > 0$,

$$\Pr[|S - E[S]| \geq \epsilon dn] \leq 2e^{-\epsilon^2 n/2}.$$

$\square$

B   Details of the proof for Hardness Result

Consider the instance which consists of a large number k copies of 6-cycles, the uniform distribution on the union of the impressions, and $n = 3k$. Let $\gamma_1, \gamma_2, \gamma_3$ and γ_+ be the fraction of the cycles that receive $1, 2, 3$ and more than 3 impressions, respectively. We have (using a simple application of Azuma's inequality) that with high probability,

$$\begin{aligned}
\gamma_1 &= 3k\frac{1}{k}(1 - \frac{3}{3k})^{3k-1} \simeq \frac{3}{e^3}, \\
\gamma_2 &= \binom{3k}{2}\frac{1}{k^2}(1 - \frac{3}{3k})^{3k-2} \simeq \frac{9}{2e^3}, \\
\gamma_3 &= \binom{3k}{3}\frac{1}{k^3}(1 - \frac{3}{3k})^{3k-3} \simeq \frac{27}{6e^3}, \\
\gamma_+ &= 1 - \gamma_1 - \gamma_2 - \gamma_3
\end{aligned}$$

For cycles that receive 1 or 2 impressions, we can assume that both ALG and OPT match 1 or 2 ads, respectively. As we are upper-bounding $E[\text{ALG/OPT}]$, we may assume that on cycles that receive more than 3 impressions, both ALG and OPT achieve 3 matches, which maximizes the contribution of these cycles to the ratio ALG/OPT.

For cycles that receive exactly 3 impressions, we have the same situation as in the single cycle above. We assume wlog that x arrives first and is matched to ad a. If the other two impressions are also both x, then both ALG and OPT match two ads (a and c) for this cycle. If the other two impressions are both y, we have that ALG matches at most two ads but OPT matches three. In all other scenarios, we assume that both ALG and OPT match three ads. By a Chernoff bound, with high probability the scenarios (x, x, x) and (x, y, y) each happen $\simeq \gamma_3 k/9$ times.

Summarizing, we have argued that with high probability,

$$\frac{\text{ALG}}{\text{OPT}} \simeq \frac{\gamma_1 + 2\gamma_2 + 3\gamma_+ + (2 \cdot \frac{2}{9} + 3 \cdot \frac{7}{9})\gamma_3}{\gamma_1 + 2\gamma_2 + 3\gamma_+ + (2 \cdot \frac{1}{9} + 3 \cdot \frac{8}{9})\gamma_3} \simeq \frac{6e^3 - 23}{6e^3 - 22} \approx .9898.$$

This establishes Theorem 3.

C Non-integral Impression Arrival Rates

One natural extension of the online stochastic matching problem is the case in which e_i's are not necessarily integers, but arbitrary rational numbers. We observe that the "Suggested Matching" algorithm, with $1 - \frac{1}{e}$-approximation factor, easily generalizes to this case, as follows: instead of computing a maximum matching, we can compute a maximum flow, f, on the corresponding flow graph, and upon the arrival of an impression i, assign the impression i to an ad a with probability f_{ia}, i.e., proportional to the fractional edge from i to a. Given the total fraction F_a on each ad a, we can argue that this algorithm achieves value $\sum_{a \in A}(1 - e^{-F_a})$ with high probability. Moreover, one can show that optimum is at most $\sum_{a \in A} F_a$ with high probability. As a result, the approximation factor of the algorithm can be captured by the ratio $\frac{\sum_{a \in A}(1 - e^{-F_a})}{\sum_{a \in A} F_a}$ where $0 \le F_a \le 1$ for all $a \in A$. Since $0 \le F_a \le 1$, we can characterize this bound as the solution to the following mathematical program:

$$
\begin{aligned}
\min \quad & \sum_{a \in A}(1 - e^{-F_a}) \\
s.t. \quad & \sum_{a \in A} F_a = 1 \\
& 0 \le F_a \le 1 \qquad \forall a \in A
\end{aligned}
$$

This mathematical program can be solved analytically. Consider the vector Φ of values $F_1, \ldots, F_{|A|}$ in nonincreasing order of F's, and let $f(\Phi) = \sum_{a \in A}(1 - e^{-F_a})$. For any vectors Φ_1 and Φ_2 subject to $||\Phi||_1 = 1$, if Φ_1 majorizes Φ_2, then clearly $f(\Phi_1) \ge f(\Phi_2)$. Since the uniform vector $\bar{\Phi} = [1/|A|, \ldots, 1/|A|]$ is majorized by all the vectors, $f(\bar{\Phi}) = |A|(1 - e^{-1/|A|})$ is the minimum value attainable. When $|A| = 1$, $f(\bar{\Phi}) = 1 - 1/e$. We derive that $\frac{df(\bar{\Phi})}{d|A|} = 1 - e^{-1/|A|} + |A|(-e^{-1/|A|} \times 1/|A|^2) = 1 - e^{-1/|A|} - e^{-1/|A|}/|A|$. Now, $1 - e^{-1/|A|} - e^{-1/|A|}/|A| \ge 0$ because multiplying by $e^{1/|A|}$, we get $e^{1/|A|} \ge 1 + 1/|A|$ which follows from the Maclaurin series expansion of e^x. Thus, $df(\bar{\Phi})/d|A| \ge 0$ and this implies that the solution to the mathematical program is attained at $1 - \frac{1}{e}$. Therefore, the approximation factor of the algorithm is equal to $1 - \frac{1}{e}$ with high probability.

Generalizing the TSM algorithm to non-integer e_is needs a proper decomposition of the flow on the corresponding flow graph to two edge-disjoint flows each with large values. Unlike the integral case, such edge-disjoint decomposition is not possible for the non-integer e_i's and one need to exploit other ideas to analyze the algorithm. We leave this as an open question.

D Frequency Capping

A useful generalization of the online matching problem that is well-motivated by the ad allocation application is when the advertisers have "frequency caps;" i.e., they do not want the same user to see their ad more than some fixed (constant) number of times. We can regard the user as a "feature" of the impression; i.e., that an "impression" i as we've used it in this paper is in fact a pair $\langle i, u \rangle$, where u is a particular user, and we have a distribution that gives us $e_{\langle i, u \rangle}$, the expected number of impressions of each type from each user. Also as part of the input, we are given, for each advertiser a, a total number of impressions d_a and a cap c per user. We could also regard these caps as operating as impression limits on other features, e.g., demographic or geographic.

Our $1 - \frac{1}{e}$-approximation algorithm (the *suggested matching* algorithm) from Section 4.1 is easily extended to this generalization of the problem. Here we give a sketch of this extension. For the algorithm, we simply make another layer U of nodes in our max-flow computation, with one node $\langle a, u \rangle$ for each (advertiser, user) pair. We make edges from each $a \in A$ to this layer with capacity c, and set the capacity of the edge edge (s, a) to d_a. The algorithm proceeds as before, and one can easily show with the same

argument that the number of impressions matched is $\simeq F(1-1/e)$, where F is the value of the flow. Then, by reasoning about the min-cut in this graph, with some simple reasoning about where this new layer sits in the min-cut, one can still show that OPT is bounded by F with high probability, giving the desired approximation ratio.

Interestingly, it is more challenging to generalize the TSM algorithm. Setting the capacities to $2d_a$ and 2, respectively, of the top and mid-layer edges does not work as desired, since then the flow could be spread among more than d_a nodes in the middle layer.