

A New Impossibility Result for Online Bipartite Matching Problems

Flavio Chierichetti ✉ 

Sapienza University of Rome, Italy

Mirko Giacchini ✉ 

Sapienza University of Rome, Italy

Alessandro Panconesi ✉ 

Sapienza University of Rome, Italy

Andrea Vattani ✉ 

Reddit, San Francisco, CA, US

Abstract

Online Bipartite Matching with random user arrival is a fundamental problem in the online advertisement ecosystem. Over the last 30 years, many algorithms and impossibility results have been developed for this problem. In particular, the latest impossibility result was established by Manshadi, Oveis Gharan and Saberi [50] in 2011. Since then, several algorithms have been published in an effort to narrow the gap between the upper and the lower bounds on the competitive ratio.

In this paper we show that no algorithm can achieve a competitive ratio better than $1 - \frac{e}{e^e} = 0.82062\dots$, improving upon the 0.823 upper bound presented in [50]. Our construction is simple to state, accompanied by a fully analytic proof, and yields a competitive ratio bound intriguingly similar to $1 - \frac{1}{e}$, the optimal competitive ratio for the fully adversarial Online Bipartite Matching problem.

Although the tightness of our upper bound remains an open question, we show that our construction is extremal in a natural class of instances.

2012 ACM Subject Classification Theory of computation → Online algorithms

Keywords and phrases Bipartite Matching, Random Graphs, Competitive Ratio

Digital Object Identifier 10.4230/LIPIcs.ICALP.2025.58

Category Track A: Algorithms, Complexity and Games

Related Version *Full Version*: 10.48550/arXiv.2504.14251

Funding Flavio Chierichetti, Mirko Giacchini, and Alessandro Panconesi were supported in part by a Google Focused Research Award, by BiCi – Bertinoro international Center for informatics, and by the PRIN project 20229BCXNW funded by the European Union - Next Generation EU, Mission 4 Component 1, CUP B53D23012910006.

Acknowledgements We thank Will Ma, David Wajc, Pan Xu, and the anonymous reviewers for several comments and suggestions.

1 Introduction

Online Bipartite Matching problems are central to the online advertising ecosystem [45, 49, 50, 52] and have numerous other applications, including efficient packet switching and routing [2, 5], crowdsourcing tasks [62], market clearing problems [10] and ride-sharing platform optimization [22]. In this paper, we establish a new upper bound of $1 - e^{1-e}$ on the competitive ratio for several well-studied variants of the problem. This result marks the first improvement of the upper bound for this classic problem in more than a decade.



© Flavio Chierichetti, Mirko Giacchini, Alessandro Panconesi, and Andrea Vattani; licensed under Creative Commons License CC-BY 4.0

52nd International Colloquium on Automata, Languages, and Programming (ICALP 2025).

Editors: Keren Censor-Hillel, Fabrizio Grandoni, Joël Ouaknine, and Gabriele Puppis

Article No. 58; pp. 58:1–58:20



Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany



To explain our result in context, we begin by recalling the problem. We are given a bipartite graph where one side of the bipartition consists of a known set of “ads” (or “advertisers”), while the other side consists of “users” (or “ad slots”), whose neighbor sets are initially unknown. Users arrive sequentially, and the generic user’s neighbors are revealed upon arrival. When a user arrives, the online algorithm can match, irrevocably, the user to one of its available neighbors; the algorithm’s aim is to maximize the size of the matching at the end of the sequence. The competitive ratio of an algorithm, in this case, is defined as the expected size of the matching produced by the algorithm divided by the expected size of the maximum possible matching. (We consider expectations since the algorithm can be randomized, and the input graph can be sampled from a distribution).

A strong motivation for studying this problem comes from online advertising, where edges represent ads that a user is interested in. Budget constraints limit how often each ad can be shown; without loss of generality, we assume that each ad can be matched to at most one user (ads with budgets allowing for multiple impressions can be reproduced multiple times in the graph). In this setting, a matching corresponds to a set of ads that will be shown to users, and the competitive ratio indicates how effectively the algorithm maximizes revenue.

This problem was introduced by Karp, Vazirani and Vazirani in a classic paper [47]. In their setup, the graph and the ordering of the users were chosen adversarially, and they presented a simple algorithm called RANKING that achieves a competitive ratio of $1 - e^{-1}$ (see [34, 47]), which was also proven to be tight [47]. The assumptions of adversarial graphs and adversarial user orderings have since been relaxed to develop algorithms with better competitive ratios; in particular, an important variant of the problem assumes that the user ordering is sampled uniformly at random, instead of being produced by an adversary [34]. Furthermore, researchers have relaxed the notion of adversarial graphs, exploring probabilistic models of user behavior [53]. In particular, consider the power set 2^A of the set of ads A , and let P be a probability distribution over 2^A . When a user arrives, a set S of ads is sampled according to P , and the ads in S become the user’s neighbors. Both the cases where P is known to the algorithm and where it is not have been considered in the literature. This model of random bipartite graphs has been studied extensively and is motivated by modern machine learning-driven methods used to identify promising ads for users.

A particularly interesting special case of this last model, known as the “irregular cuckoo hashing” model,¹ arises when a sample from P is obtained by first sampling the degree of a user from a random variable D taking values over the non-negative integers, and then sampling uniformly-at-random with replacement a number of ads equal to the sampled degree. This process generates a multiset of ads, but duplicate ads can be removed, leaving the remaining set as the user’s neighbors.

We can now state our result more precisely. We show that no online algorithm can achieve a competitive ratio greater than $1 - e^{1-e} = 0.82062\dots$ across all the models discussed above. We establish our bound in the irregular cuckoo hashing model and so, by definition, the bound extends to all the other models and is the strongest known bound for all those variants. The previous strongest upper bound, established in 2011 by Manshadi, Oveis Gharan, and Saberi [50], was 0.823. (Table 1 summarizes the state of the art). While our improvement is small, it is comparable in magnitude to many previous algorithmic advancements for this problem.

¹ We recall that in (a standard variant of) cuckoo hashing [24, 56], given an integer $k \geq 1$, each item x is hashed to k uniform-at-random slots of a hashtable; the item is stored in an empty slot, if one is available. In irregular cuckoo hashing [23], one obtains k by applying another hash function to x – this way, different items can be hashed to a different number of slots.

■ **Table 1** Bounds for various Online Bipartite Matching problem variants. Each row corresponds to a different variant, listed in order from the easiest to the hardest.

Setting	Competitive Ratio		
	Best Algorithm	Our UB	Previous Best UB
Irregular Cuckoo Hashing	0.716 [37]	$1 - \frac{e}{e^e} = 0.820\dots$	0.823 [50]
IID Users (Known Distr.)			
IID Users (Unknown Distr.)	0.696 [49]		
Adv. Graph, UAR Ordering			
Adv. Graph and Ordering	$1 - \frac{1}{e} = 0.632\dots$ [34, 47]		$1 - \frac{1}{e} = 0.632\dots$ [47]

The upper bound of 0.823 on the competitive ratio from [50] relies on the irregular cuckoo hashing constructions of Dietzfelbinger et al. [23] and was derived through numerical optimization, without yielding a closed-form expression. In contrast, our construction enables an exact analytical computation of its optimal competitive ratio. Furthermore, our construction uses simple rational probabilities, whereas the optimal probabilities in [50] remain unspecified and likely irrational.

The random graphs used to prove our bound are based on a simple random variable D taking values over the non-negative integers:

$$\Pr[D = 0] = \Pr[D = 1] = 0, \quad \Pr[D = d] = \frac{1}{d \cdot (d - 1)}, \text{ for } d \geq 2,$$

or, equivalently, $\Pr[D > d] = 1/d$ for each positive integer d . As discussed earlier, when a user arrives, a value d is drawn from D , and d ads are chosen uniformly at random (with replacement) as neighbors. The graph consists of n ads and n users. As we will clarify later, this distribution greedily ensures the “maximum possible” number of vertices of degree d for $d = 0, 1, 2, \dots$, while maintaining a quasi-complete matching, i.e., one matching a $1 - o(1)$ fraction of the users.

Intuitively, selecting user neighbors uniformly at random while keeping user degrees low impairs the performance of any online algorithm: over time, newly arriving users are likely to find all their randomly assigned neighbors already matched. However, establishing a strong upper bound also requires ensuring that the maximum matching remains large – this is the main technical challenge in our analysis.

Whether our upper bound is tight for some or all of the four problem variants mentioned above remains an open question. Notably, the bound involves the Euler’s number, in a manner analogous to the optimal competitive ratio of $1 - \frac{1}{e}$ of the fully-adversarial Online Bipartite Matching problem. While we cannot determine whether our upper bound can be improved, we show that it possesses a certain extremal property and that it is the strongest bound possible for a natural class of instances.

Paper Organization. Section 2 gives an overview of our main construction, of its relation to previous upper bounds, and of the techniques we used to study it. Section 3 reviews related work. Section 4 introduces key tools, techniques, and notation. In Section 5, we present our construction and prove the main result: the $1 - e^{1-e}$ upper bound on the competitive ratio of Online Bipartite Matching Problems. Section 6 establishes the extremality of our

construction, showing that it greedily maximizes the number of low-degree users while ensuring a quasi-complete matching. Section 7 extends our construction to arbitrary user-to-ad ratios, demonstrating that equibipartite graphs yield the strongest bound. We conclude with open questions in Section 8.

All the missing proofs can be found in the full version (10.48550/arXiv.2504.14251).

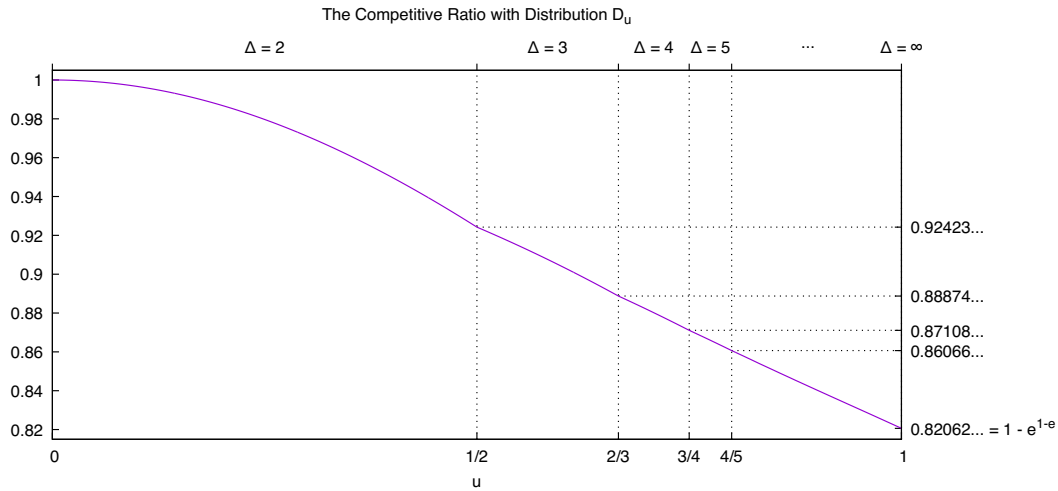
2 Overview

Our irregular cuckoo hashing instance contains n users and n ads, and assigns probability $\frac{1}{d \cdot (d-1)}$ to each user-degree $d \geq 2$ – in other words, it posits that the probability that a user has degree strictly larger than the generic positive integer d is exactly $1/d$. We obtained this distribution by greedily maximizing the number of users of degree $d = 0, 1, 2, 3, \dots$ while guaranteeing that the random graph admits a quasi-complete matching from the user side – a matching that pairs all users except for at most $o(n)$ of them. Observe that, in a cuckoo hashing instance, the smaller the user-degrees the less likely it is for an online algorithm to match users. The goal of our distribution is to minimize the number of online matches while guaranteeing that the graph admits a quasi-complete matching – in other words, the distribution aims to minimize the numerator of the competitive ratio while making its denominator as large as possible. (We remark that most previous upper bounds were also based on graphs admitting quasi-complete matchings; see, e.g., [6, 34, 47, 50].)

Our Main Construction. It is easy to see that, in order for a cuckoo hashing instance containing n ads, and no more than n users, to be quasi-complete from the user side, the instance must contain at most $o(n)$ users of degree 0 and $o(n)$ users of degree 1. Furthermore, classical results can be used to prove that, if one aims to be quasi-complete from the user side, and all users have degree 2, then the maximum number of users is $c_2 \cdot n \pm o(n)$, for $c_2 = \frac{1}{2}$.² Now, suppose that we must have $c_2 \cdot n \pm o(n)$ users of degree 2. What is the largest number of users of degree 3 that we can add while still guaranteeing that there exists a matching covering a $1 - o(1)$ fraction of the users? We prove that the answer is “ $c_3 \cdot n \pm o(n)$, for $c_3 = \frac{1}{6}$ ”. Next, if we are bound to have $c_2 \cdot n \pm o(n)$ users of degree 2 and $c_3 \cdot n \pm o(n)$ users of degree 3, what is the maximum c_4 such that we can add $c_4 \cdot n \pm o(n)$ users of degree 4 and match a $1 - o(1)$ fraction of users? And so on and so forth.

In general, we set $c_d = \frac{1}{d \cdot (d-1)}$ for each degree $d \geq 2$, and $c_0 = c_1 = 0$. These fractions result in a total of n users, to be matched to the n ads. To establish that these c_d 's give rise to a graph admitting a quasi-complete matching we couple the irregular cuckoo hashing graph with a random graph in the configuration model having the same degree distribution on the user side, and the same Poissonian degree distribution on the ads side; then, by means of the Karp-Sipser algorithm [46] (which we analyze with a variant of the techniques presented in [7] by Balister and Gerke), we establish that this configuration model graph admits a quasi-complete matching; using the coupling, we derive the existence of a quasi-complete

² The feasibility of $c_2 = \frac{1}{2}$ can be obtained as a corollary of a classical result of Erdős and Rényi [27] which states that, for all small enough $\epsilon > 0$, $G(n, (\frac{1}{2} - \epsilon) \cdot n)$ has, with probability $1 - o(1)$, a $1 - o(1)$ fraction of its vertices in connected components that are trees. By taking each vertex of $G(n, (\frac{1}{2} - \epsilon) \cdot n)$ to be an ad, and each of its edges $\{u, v\}$ to be a user of degree 2 connected to the ads u and v , one deduces the claim from the observation that the edges of a forest can be oriented so that no vertex has in-degree larger than 1.



■ **Figure 1** The Competitive Ratios induced by the D_u instances, obtained numerically for $u > 0$. The plot shows how $u = 1$ (which corresponds to the case of unbounded user-degree) gives the strongest upper bound. In Theorem 27, we analytically establish the optimality of $u = 1$.

matching in the irregular cuckoo hashing graph, as well. Finally, we upper bound the fraction of ads matched by any online algorithm with $1 - \frac{e}{e^e} = 0.82062\dots$, thus proving our new upper bound on the competitive ratio of several Online Bipartite Matching problems.

As an extra step, aimed at establishing the *extremality* of our construction, we show that if one increases any of the c_d 's by any constant $\epsilon > 0$, then a constant fraction of the vertices of degree at most d become unmatched. (In particular, if the extra ϵ fraction assigned to degree d is taken from vertices of higher degree, then the quasi-complete matching disappears).

A Class of Constructions. Our construction has users of unbounded degree, and contains the same number of users and ads. To get a more thorough understanding of the construction, and of the problem, we study a class of constructions that generalizes our main one; the instances in this class have n ads, and $u \cdot n$ users. Just like our main construction, we populate the generic instance in the class by greedily adding users with the smallest possible degree, under the constraint that the instance admits a quasi-complete matching. For a given $u \in (0, 1)$, the generalized instance D_u has a maximum user degree of $\lceil \frac{1}{1-u} \rceil$. We prove that, for all $u \in (0, 1)$, these instances give a weaker upper bound on the competitive ratio than the case $u = 1$, that is, the case with n users and n ads (see Figure 1 and Theorem 27).

Previous Constructions. We point out that studying instances of varying user degrees is necessary to improve over the state of the art. In particular, if an irregular cuckoo hashing instance has n users and n ads, and all the users have degree 1, any greedy online algorithm will necessarily return a maximum matching (which matches a fraction of $1 - \frac{1}{e} = 0.63212\dots$ users). Thus, the best upper bound provable with this instance is 1. If, instead, each user has degree 2, a greedy online algorithm will return a matching with, roughly, $\tanh(1) \cdot n = 0.76159\dots \cdot n$ edges. The maximum matching has size approximately equal to $\left(1 + \frac{1}{2}W(-2e^{-2}) + \frac{1}{4}W(-2e^{-2})^2\right) \cdot n = 0.83809\dots \cdot n$ (see, e.g., [43]; here, W is the Lambert W function), so that the best upper bound provable with this instance is

0.90871... Finally, if each user has degree 3 or more, one can show that the number of users matched by any greedy online algorithm is at least $0.82304\dots \cdot n$. It follows that having a uniform user degree in equibipartite graphs makes it impossible to improve upon – or even match – the 0.823 upper bound of Manshadi et al. [50].

Indeed, Manshadi et al. [50] derived their upper bound by employing a mixture of users with varying degrees. In their construction, the user degree distribution is supported on $\{2, 3, n\}$.³ Manshadi et al.’s analysis leverages on an irregular cuckoo hashing construction of Dietzfelbinger et al. [23], whose parameters were obtained numerically; as a result, the optimal fractions of vertices of degrees 2 and 3 in the bipartite matching result of [50] are not given explicitly (and might very well be irrational numbers). Our extremal construction, instead, allows us to compute the competitive ratio analytically and exactly, and is made up of simple, rational, probabilities.

3 Related Work

Online Bipartite Matching problems have a long history. We mention here the results that are most relevant to our paper; see the book of Mehta [52], or the recent survey of Huang, Tang and Wajc [38], for a more comprehensive introduction to these problems.

As we mentioned before, if the graph and the ordering are adversarial, then a competitive ratio of $1 - \frac{1}{e}$ can be achieved via the RANKING algorithm of Karp, Vazirani and Vazirani [47] (see, Goel and Mehta [34], or Birnbaum and Mathieu [8], for proofs; simpler analyses were given in, e.g., Devanur, Jain and Kleinberg [21] and Eden et al. [26]): this algorithm begins by sampling a uniform-at-random ordering of the ads; whenever a user arrives, if at least one ad is available for that user, the user is matched to the first available ad in the ordering. As shown in [47], the competitive ratio of the RANKING algorithm, $1 - \frac{1}{e}$, is optimal in the fully-adversarial case.

Determining the optimal competitive ratio if the set of users is permuted uniformly-at-random is, possibly, the most important open problem in this area. The best known upper bound on the competitive ratio of this problem was proved by Manshadi, Oveis Gharan and Saberi [50, 51] and has value 0.823 – this upper bound holds regardless of whether the bipartite graph is generated adversarially or randomly. The RANKING algorithm is still the best known algorithm for this problem on adversarial graphs; Mahdian and Yan [49] show that it achieves a competitive ratio of at least 0.696, and Karande, Mehta and Tripathi [45] show that its competitive ratio is not larger than 0.727.

The case of randomized bipartite graphs has also been studied extensively; in this setting the adversary produces a distribution over subsets of ads, and each user samples independently a set of ads from that distribution. In the “known distribution” case, this distribution is assumed to be known to the algorithm; the “unknown distribution” variant has also been studied extensively. As mentioned in the previous paragraph, Mahdian and Yan [49] gave an algorithm for randomly-permuted adversarial graphs with a competitive ratio of 0.696; this algorithm applies to both the known, and the unknown, distribution cases. Several improvements on the competitive ratio for the known distribution case have been proved in the last fifteen years. In 2011, Manshadi et al. [50], proposed an algorithm with a competitive

³ Specifically, Manshadi et al. [50] have n ads and n users, and assign probability 0.40517 to user-degree 2, 0.40517 to user-degree 3, and 0.18966 to user-degree n . Using the results in [23], they prove that the graph admits a maximum matching of size $n - o(n)$. They conclude their argument by numerically computing the size of the online matching, which they find to be no larger than $0.823n$.

ratio of 0.702. In 2014, Jaillet and Lu [39] had an algorithm achieving a competitive ratio of 0.706. Then, in 2021, Huang and Shu [36] gave an algorithm with a competitive ratio of 0.711; finally, in 2022, Huang, Shu and Yan [37] achieved a competitive ratio of 0.716.

The best impossibility result so far – which holds in the cases of adversarial graph with UAR ordering, of random graphs with unknown distribution, and known distribution, as well as in the irregular cuckoo hashing case – is the aforementioned upper bound of 0.823 published by Manshadi et al. in 2011 [50]. Our result improves the bounds in each of these four cases from 0.823 to $1 - \frac{e}{e^e} = 0.82062\dots$. We remark that our numerical improvement, of value roughly 0.002, is of the same order of all the algorithmic improvements mentioned in the previous paragraph.

Several special cases of online bipartite matching problems have been studied in the last few years. In particular, the “integral arrival rates” special case of the “known distribution” setting requires that each set in the support of the distribution has an integral expected number of occurrences.⁴ For this case, the best known algorithm – due to Brubach et al. [17] – achieves a competitive ratio of 0.7299; the best known upper bound for the integral arrival rates case has value $1 - e^{-2} = 0.864\dots$, and is due to Manshadi et al. [51].

We conclude this section by noting that the interest in online bipartite matching is further evidenced by the attention given in recent years to several related problems. These include some of its weighted [29, 37, 44], and fairness-aware [48], variants; in particular, Huang et al. [37] show an upper bound of 0.703 for the (known distribution) edge-weighted online bipartite matching problem, whereas Ma and Xu [48] show an upper bound of $\sqrt{3} - 1 = 0.732\dots$ for a variant whose goal is to maximize the minimum matching rate across different groups of users. Other related problems include fractional variants and their rounding schemes [9, 18, 55, 58, 61], variants with different probing models [12–14], those with time-dependent user distributions [15, 16, 57, 60], with users arriving in batches [30, 31, 42] or with different arrival models [19, 28, 32], variants that leverage machine-learned advice [3, 20, 42], as well as generalized selection tasks [33, 35].

4 Preliminaries

We present here the objects, the tools and the notation that we will use in this paper.

Notation. Given a predicate P , the Iverson bracket $[P]$ has value 1 if P is true, and value 0 otherwise. With a slight abuse of notation, for a positive integer n , we use $[n]$ to denote the set $[n] = \{1, \dots, n\}$. The support of a multiset S is the set of elements whose multiplicity in S is at least 1. For an integer $k \geq 1$, we let $H_k = \sum_{i=1}^k \frac{1}{i}$ be the k -th harmonic number. Finally, we let $\Phi(z, s, \alpha) = \sum_{k=0}^{\infty} \frac{z^k}{(k+\alpha)^s}$ be the Lerch transcendent, a series that generalizes the Riemann ζ function, and which converges if $\alpha > 0$, $|z| < 1$ and $s \in \mathbf{R}$.

Bipartite (Multi)Graphs and Matchings. An undirected bipartite (multi)graph $G(V, \hat{V}, E)$ is defined by two disjoint sets of vertices V and \hat{V} , and by a (multi)set of edges E having a support set which is a subset of $\{\{v, \hat{v}\} \mid v \in V \wedge \hat{v} \in \hat{V}\}$. If E coincides with its support set, then $G(V, \hat{V}, E)$ has no parallel edges, and we say that $G(V, \hat{V}, E)$ is a simple graph. We will

⁴ We observe that our construction, like the 0.823 construction of Manshadi et al. [50], is an equibipartite irregular cuckoo hashing instance with a maximum user degree larger than 1. Consequently, the total number of “user types” in the support of the distribution is at least quadratic in the number of users. Thus, these two constructions do not exhibit integral arrival rates.

typically use $V = \{v_1, v_2, \dots\}$ to denote the set of users, and $\hat{V} = \{\hat{v}_1, \hat{v}_2, \dots\}$ to denote the set of ads. We will also use n_i (resp., \hat{n}_i) to denote the number of vertices of degree $i \geq 0$ in V (resp., \hat{V}).

A matching M in the (multi)graph $G(V, \hat{V}, E)$ is a subset $M \subseteq E$ containing pairwise disjoint edges. We say that $G(V, \hat{V}, E)$ admits a *complete matching* if there exists a matching $M \subseteq E$ of $G(V, \hat{V}, E)$ such that $|M| = \min(|V|, |\hat{V}|)$, and that it admits a *perfect matching* if there exists a matching $M \subseteq E$ such that $|M| = |V| = |\hat{V}|$.

Given a bipartite multigraph $G(V, \hat{V}, E)$, if E' is the support set of E , we have that $G(V, \hat{V}, E')$ is a bipartite simple graph. Observe that each matching of $G(V, \hat{V}, E')$ is also a matching of $G(V, \hat{V}, E)$; moreover, given an arbitrary matching M of $G(V, \hat{V}, E)$, we can transform M into a matching of $G(V, \hat{V}, E')$ by substituting, for each $e \in M$, the edge e with its unique parallel edge $e' \in E'$.

We say that an infinite sequence of bipartite (multi)graphs $G(V_1, \hat{V}_1, E_1), G(V_2, \hat{V}_2, E_2), \dots$ of increasing side sizes ($\min(|V_1|, |\hat{V}_1|) < \min(|V_2|, |\hat{V}_2|) < \dots$), and with maximum matchings $M_1 \subseteq E_1, M_2 \subseteq E_2, \dots$, admits a *quasi-complete matching* if $|M_i| \geq (1 - o(1)) \cdot \min(|V_i|, |\hat{V}_i|)$ – that is, if for each $\epsilon > 0$ there exists i^* such that, for each $i \geq i^*$, $|M_i| \geq (1 - \epsilon) \cdot \min(|V_i|, |\hat{V}_i|)$.

The Karp-Sipser Algorithm. The Karp-Sipser algorithm [46] is a greedy algorithm for computing a matching of a graph. A run of this algorithm is divided in two phases; for our purposes, only the first phase will be necessary. During the first phase, the algorithm looks for any vertex of degree 1; as long as such a vertex exists, the algorithm adds its incident edge e to a set S , removes the two endpoints of e from the graph, and iterates. If no such vertex exists, the first phase ends. It is easy to prove that, at any point during the execution of the first phase, the set S can be extended to a maximum matching of the original graph. In fact, it is also easy to prove that if the original graph $G(V, \hat{V}, E)$ is bipartite and, at any point during the execution of the first phase, $V_0 \subseteq V$ (resp. $\hat{V}_0 \subseteq \hat{V}$) is the subset of V (resp. \hat{V}) containing vertices that have current degree 0 and which are still unmatched, then the maximum matching cannot be larger than $\min(|V| - |V_0|, |\hat{V}| - |\hat{V}_0|)$.

The Karp-Sipser algorithm was introduced to bound the size of the maximum matching in sparse $G(n, p)$ random graphs; since its introduction, the algorithm has been used to bound the size of the maximum matching of many other random graph models.

The Irregular Cuckoo Hashing Model. For simplicity of exposition, the random graph models that we will consider in this paper might produce non-simple graphs, that is, multigraphs with parallel edges; these can be easily transformed into simple graphs through the removal of parallel edges. One of the two random graph models that we will consider is the irregular cuckoo hashing model, which is defined as follows:

► **Definition 1** (Irregular Cuckoo Hashing Model). *Let n, m be two positive integers, and P be a probability distribution over the non-negative integers. The Irregular Cuckoo Hashing random multigraph $\text{ICH}(n, m, P)$ is defined as follows. This bipartite graph has user set V and ads set \hat{V} , with $n = |V|$ and $m = |\hat{V}|$. Each user $v_i \in V$ samples independently an integer d_i from P and, for each $j = 1, 2, \dots, d_i$, samples independently an ad $\hat{v}_{i,j}$ uniformly at random from \hat{V} . The multiset of neighbors of v_i is $\{\hat{v}_{i,1}, \dots, \hat{v}_{i,d_i}\}$. (Note that $\text{ICH}(n, m, P)$ might not be simple).*

The Configuration Model. In order to analyze the irregular cuckoo hashing graph of our construction, we will couple it with a random graph in the configuration model.

► **Definition 2** (Configuration Model). Consider two sequences of non-negative integers $\mathbf{d} = (d_1, \dots, d_n)$ and $\hat{\mathbf{d}} = (\hat{d}_1, \dots, \hat{d}_m)$ such that $\sum_{i=1}^n d_i = \sum_{i=1}^m \hat{d}_i$. To sample a bipartite graph $G = (V, \hat{V}, E)$ from the configuration model $\text{CM}(\mathbf{d}, \hat{\mathbf{d}})$, for each $i \in [n]$ (resp. $j \in [m]$), assign to vertex v_i (resp. \hat{v}_j) d_i (resp. \hat{d}_j) configuration points, and then sample a uniform at random perfect bipartite matching between the configuration points. (Note that $\text{CM}(\mathbf{d}, \hat{\mathbf{d}})$ might not be simple).

To compute the size of the maximum matching in our configuration model distribution, we will study the behavior of the first phase of the Karp-Sipser algorithm using the techniques presented, e.g., by Balister and Gerke in [7] (see also [4, 11]). In particular, we will show that our distribution admits a quasi-complete matching through the following modification⁵ of a result in [7]:

► **Theorem 3.** Let $\rho \in (0, 1]$ be any constant. Let $f(x) = \sum_{i=0}^{\infty} z_i \cdot x^i$, and $\hat{f}(x) = \sum_{i=0}^{\infty} \hat{z}_i \cdot x^i$, where $(z_i)_{i=0}^{\infty}, (\hat{z}_i)_{i=0}^{\infty}$ are non-negative sequences such that $f'(1) = \hat{f}'(1) = \mu < \infty$ and $f'(2), \hat{f}'(2) < \infty$. Consider any configuration model $\text{CM}(\mathbf{d}, \hat{\mathbf{d}})$, with $|\mathbf{d}| = f(1)n \pm o(n)$, $|\hat{\mathbf{d}}| = \hat{f}(1)n \pm o(n)$, and such that (i) $n_i = (z_i \pm o(n^{-\rho}))n$ for $i \geq 0$, (ii) $\hat{n}_i = (\hat{z}_i \pm o(n^{-\rho}))n$ for $i \geq 0$, (iii) $|E| = (\mu \pm o(n^{-\rho}))n$, and (iv) for $M = M(n) = \Theta(\log n)$, $\sum_{i=M}^{\infty} i^2 \cdot z_i = o(n^{-\rho})$ and $\sum_{i=M}^{\infty} i^2 \cdot \hat{z}_i = o(n^{-\rho})$, where n_i (resp. \hat{n}_i) is the number of vertices of degree i in \mathbf{d} (resp. $\hat{\mathbf{d}}$).

Let $w_2, \hat{w}_1 \in [0, 1]$ be the smallest solutions to the simultaneous equations $w_2 = 1 - \frac{f'(1-\hat{w}_1)}{\mu}$, $\hat{w}_1 = \frac{\hat{f}'(w_2)}{\mu}$. Then, with probability $1 - o(1)$, the first phase of Karp-Sipser, on the graph $G(V, \hat{V}, E) \sim \text{CM}(\mathbf{d}, \hat{\mathbf{d}})$, returns a matching of size at least

$$(f(1) - f(1 - \hat{w}_1) - f'(1 - \hat{w}_1)\hat{w}_1 - o(1))n.$$

Moreover, with probability $1 - o(1)$, after the first phase of Karp-Sipser, at least $(\hat{f}(w_2) - f(1) + f(1 - \hat{w}_1) + f'(1 - \hat{w}_1)\hat{w}_1 - o(1))n$ vertices in \hat{V} will be isolated (i.e., of degree 0 and unmatched). Equivalently, with probability $1 - o(1)$, the maximum matching cannot be larger than,

$$(\hat{f}(1) - \hat{f}(w_2) + f(1) - f(1 - \hat{w}_1) - f'(1 - \hat{w}_1)\hat{w}_1 + o(1))n.$$

In the full version, we prove Theorem 3 by following the approach in [7], with minor modifications to handle distributions of unbounded support.

Random Variables and Concentration Inequalities. We denote with $\text{Bin}(n, p)$ the binomial distribution with success probability p and n trials; with $\text{Poisson}(\lambda)$ the Poisson distribution of parameter $\lambda > 0$; and with $\text{Geom}(p)$ the geometric distribution with success probability p , in particular $\Pr_{X \sim \text{Geom}(p)}[X = k] = p(1-p)^{k-1}$ for $k \geq 1$, and $\mathbb{E}_{X \sim \text{Geom}(p)}[X] = \frac{1}{p}$. For two random variables X, Y taking values in the natural numbers \mathbf{N} , we let $d_{\text{TV}}(X, Y) = \sup_{A \subseteq \mathbf{N}} |\Pr[X \in A] - \Pr[Y \in A]|$ be their total variation distance.

⁵ The Theorems of Balister and Gerke require the degree distributions to have finite support. Observe that any degree distribution with a finite mean can be truncated to a finite prefix altering the distribution's expectation by no more than any $\epsilon > 0$; in the setting of the configuration model, this truncation results in a change of no more than an ϵ fraction in the maximum matching size. Balister and Gerke use this approach to apply their Theorems to generic distributions of finite average degree. In our case, however, the ads degree distribution makes it possible to apply several algebraic shortcuts that become infeasible once that distribution is transformed into one having finite support. Therefore, to simplify our analysis, we reproved certain Lemmas and Theorems of Balister and Gerke [7], so to accommodate distributions with unbounded support.

58:10 A New Impossibility Result for Online Bipartite Matching Problems

In our arguments, we will make use of the following well-known concentration bounds (whose proofs can be found, e.g., in [25]):

► **Fact 4** (Chernoff-Hoeffding's inequality). *Let X_1, \dots, X_n be independent random variables such that $X_i \in [a, b]$ almost surely. Let $X = \sum_{i=1}^n X_i$, then, for all $t > 0$,*

$$\Pr[|X - \mathbb{E}[X]| \geq t] \leq 2 \exp\left(-\frac{2 \cdot t^2}{n \cdot (b-a)^2}\right).$$

Additionally, we will employ McDiarmid's inequality, a generalization of the Chernoff-Hoeffding bound:

► **Fact 5** (McDiarmid's inequality). *Let X_1, X_2, \dots, X_n be independent random variables with $X_i \in \mathcal{X}$. Let $f : \mathcal{X}^n \rightarrow \mathbb{R}$ be such that for each $i \in [n]$ and for each $x_1, x_2, \dots, x_n, x'_i \in \mathcal{X}$, it holds $|f(x_1, \dots, x_{i-1}, x_i, x_{i+1}, \dots, x_n) - f(x_1, \dots, x_{i-1}, x'_i, x_{i+1}, \dots, x_n)| \leq d$. Then, for each $t > 0$,*

$$\Pr[|f(X_1, \dots, X_n) - \mathbb{E}[f(X_1, \dots, X_n)]| \geq t] \leq 2 \exp\left(-\frac{2 \cdot t^2}{n \cdot d^2}\right).$$

We will also leverage on the following concentration inequality for the sum of geometric random variables:

► **Fact 6** ([40], Theorem 2.1 and Theorem 3.1). *Let X_1, \dots, X_n be independent random variables with $X_i \sim \text{Geom}(p_i)$. Let $X = \sum_{i=1}^n X_i$, and let $p^* = \min_i p_i > 0$. For each $\lambda \geq 1$, $\Pr[X \geq \lambda \mathbb{E}[X]] \leq e^{-p^* \mathbb{E}[X](\lambda-1-\ln \lambda)}$, and for each $\lambda \in (0, 1]$, $\Pr[X \leq \lambda \mathbb{E}[X]] \leq e^{-p^* \mathbb{E}[X](\lambda-1-\ln \lambda)}$.*

Finally, we will use the following approximation, by means of Poisson variables, of binomial random variables having small expectation.

► **Fact 7** ([1, 59]). *Let $X \sim \text{Bin}(n, p)$, $Y \sim \text{Poisson}(\lambda)$ for $\lambda > 0$, then $d_{\text{TV}}(X, Y) \leq np^2 + |\lambda - np|$.*

5 Impossibility Result

In this Section we present, and analyze, an irregular cuckoo hashing instance of the online bipartite matching problem; by means of this instance, we will show that no online algorithm can achieve a competitive ratio better than $1 - \frac{e}{e^e} = 0.8206259\dots$. In Section 6 we will show how this construction can be obtained by greedily maximizing the probabilities of the degrees $0, 1, 2, 3, \dots$, while guaranteeing that the resulting graph has a quasi-complete matching.

► **Definition 8** (Main Instance). *Our main instance is given by the irregular cuckoo hashing distribution $\text{ICH}(n, n, D)$, where D is defined as,*

$$\Pr[D = d] = \frac{1}{d \cdot (d-1)} \quad \forall d \in \{2, 3, 4, \dots\}.$$

Observe that, for each positive integer d , $\Pr[D > d] = \frac{1}{d}$. Recall that in our irregular cuckoo hashing instances the ads are sampled with replacement.⁶

We will bound the performance of any online algorithm on the instance of Definition 8 in Section 5.1; we will then show in Section 5.2 that the instance admits, with high probability, a quasi-complete (and, thus, quasi-perfect) matching.

⁶ We point out that our result could also be proved if the sampling of the D ads was done without replacement (with a suitable truncation of the distribution), since the probability that a user chooses

5.1 Online Matching

In this Section we upper bound the expected size of the matching found by any online algorithm when presented with the instance of Definition 8.

► **Theorem 9.** *No online algorithm can match more than*

$$\left(1 - \frac{e}{e^e}\right) \cdot n + O(\sqrt{n}) = 0.8206259 \dots \cdot n + O(\sqrt{n})$$

users in expectation with the instance of Definition 8.

Proof. Consider the greedy algorithm that, whenever possible, matches the current user to any of its neighboring ads that are still unmatched. This algorithm is optimal for the instance of Definition 8, given that each (multi)set of ads is chosen independently of the others, and that, after conditioning on the (multi)set size s , the (multi)set is chosen uniformly at random among those of cardinality s . We then restrict our analysis to this algorithm.

Let S be the number of users sampled up until, and including, the point where the greedy algorithm matches m ads, for $m = \lceil \alpha \cdot n \rceil + 1 = \alpha \cdot n + O(1)$, with $\alpha = 1 - e^{1-e}$. Note that it may be that $S \geq n$. Without loss of generality, we can assume that n is larger than any constant, otherwise the statement is trivially true – in our case, we will take $n \geq 100$.

The crux of the analysis of the algorithm lies in bounding the expected number of users necessary to achieve a matching of size $m = (1 - e^{1-e})n + O(1)$. First, we split the value S in m parts. For $i \in \{0, 1, \dots, m-1\}$, let S_i be the number of users sampled in order to match the $(i+1)$ -th user, counting from the first user after the i -th user is matched (or, if $i = 0$, counting from the first user). We then have $S = \sum_{i=0}^{m-1} S_i$. Let $p_j = \frac{1}{j(j-1)}$ for $j \geq 2$. Observe that $S_i \sim \text{Geom}(q_i)$, where $q_i = 1 - \sum_{j \geq 2} \left(p_j \cdot \binom{i}{n}^j\right)$. For each $i \in \{0, 1, \dots, m-1\}$, let $x_i = i/n$; then, $0 \leq x_i < 1$, and:

$$\begin{aligned} q_i &= 1 - \sum_{j \geq 2} \left(\frac{1}{j \cdot (j-1)} \cdot x_i^j\right) = 1 - \sum_{j \geq 2} \left(\left(\frac{1}{j-1} - \frac{1}{j}\right) \cdot x_i^j\right) \\ &= 1 - \sum_{j \geq 2} \left(\frac{1}{j-1} \cdot x_i^j\right) + \sum_{j \geq 2} \left(\frac{1}{j} \cdot x_i^j\right) = 1 - x_i \cdot \sum_{j \geq 1} \left(\frac{1}{j} \cdot x_i^j\right) + \sum_{j \geq 1} \left(\frac{1}{j} \cdot x_i^j\right) - x_i \\ &= 1 + x_i \cdot \ln(1 - x_i) - \ln(1 - x_i) - x_i = (1 - x_i)(1 - \ln(1 - x_i)), \end{aligned} \tag{1}$$

where we used $\sum_{j \geq 1} \left(\frac{1}{j} \cdot x^j\right) = -\ln(1 - x)$ for each $x \in [-1, 1)$.

We can now bound the expectation of S :

$$\frac{\mathbb{E}[S]}{n} = \frac{1}{n} \sum_{i=0}^{m-1} \frac{1}{(1 - x_i)(1 - \ln(1 - x_i))} \geq \frac{1}{n} \sum_{i=1}^{m-1} \frac{1}{(1 - x_i)(1 - \ln(1 - x_i))}$$

the same ad more than once using the model of Definition 8 – that is, the probability that the chosen multiset of ads is not a set – is at most $O(n^{-1/2})$. Indeed, if S is the multiset of the neighbors of a given user, we have that $\Pr[S \text{ is not a set}] = \sum_{d=2}^{\infty} (\Pr[D = d] \cdot \Pr[S \text{ is not a set} \mid D = d]) \leq \sum_{d=2}^{\sqrt{n}} (\Pr[D = d] \cdot \Pr[S \text{ is not a set} \mid D = d]) + \sum_{d=\sqrt{n}+1}^{\infty} \Pr[D = d] \leq \sum_{d=2}^{\sqrt{n}} \left(\frac{1}{d \cdot (d-1)} \cdot \binom{d}{2} \cdot \frac{n}{n^2}\right) + \Pr[D > \sqrt{n}] \leq \frac{\sqrt{n}}{2n} + \frac{1}{\sqrt{n}} = \frac{3}{2\sqrt{n}}$. Therefore a sampling with replacement creates, in expectation and with high probability, at most $O(\sqrt{n})$ users with parallel edges; removing these users reduces the size of any matching by no more than $O(\sqrt{n})$ edges.

$$\begin{aligned} &\geq \int_0^{(m-1)/n} \frac{1}{(1-x)(1-\ln(1-x))} dx \geq \int_0^\alpha \frac{1}{(1-x)(1-\ln(1-x))} dx \\ &= [\ln(1-\ln(1-x))]_{x=0}^\alpha = 1, \end{aligned}$$

where the second inequality follows from the fact that $\frac{1}{(1-x)(1-\ln(1-x))}$ is increasing in $[0, 1)$ and $\frac{1}{n} \sum_{i=1}^{m-1} \frac{1}{(1-x_i)(1-\ln(1-x_i))}$ is a right Riemann sum on the interval $[0, (m-1)/n]$. The third inequality holds since $\frac{1}{(1-x)(1-\ln(1-x))} > 0$ in $[0, 1)$ and $(m-1)/n \geq \alpha$. Finally, the solution to the integral can be easily verified. Thus, $E[S] \geq n$.

We move on to showing the concentration of S :

► **Lemma 10.** *Let $\delta \in [e^{-n/9}, e^{-1}]$, and $\epsilon(\delta) = \frac{1}{n} \lceil \sqrt{6n \ln(1/\delta)} \rceil$. It holds, $\Pr[S \leq (1 - \epsilon(\delta))n] \leq \delta$.*

Now, let A_k be the size of the matching produced by the greedy algorithm after having processed the first k users (in particular, the matching found by the online algorithm has size A_n). It holds $A_k \geq m \iff S \leq k$. Moreover, $A_n \leq A_{n-k} + k$. Note also that $\epsilon(\delta) \cdot n \leq 3\sqrt{n \ln(1/\delta)}$ since $\delta \leq e^{-1}$. By Lemma 10,

$$\begin{aligned} \Pr[A_n \geq m + 3\sqrt{n \cdot \ln(1/\delta)}] &\leq \Pr[A_n \geq m + \epsilon(\delta) \cdot n] \leq \Pr[A_{(1-\epsilon(\delta))n} \geq m] \\ &= \Pr[S \leq (1 - \epsilon(\delta))n] \leq \delta. \end{aligned}$$

We can finally bound the expected number of matched users,

$$\begin{aligned} E[A_n] &= \sum_{k=1}^n \Pr[A_n \geq k] \leq m + 3\sqrt{n} + \sum_{k=m+\lceil 3\sqrt{n} \rceil}^n \Pr[A_n \geq k] \\ &\leq m + 4\sqrt{n} + \sum_{k=3}^{\lceil e^{1-e} \sqrt{n} \rceil} \Pr[A_n \geq m + k \cdot \sqrt{n}] \cdot \sqrt{n} \\ &= m + 4\sqrt{n} + \sqrt{n} \cdot \sum_{k=3}^{\lceil e^{1-e} \sqrt{n} \rceil} \Pr \left[A_n \geq m + 3\sqrt{n \cdot \ln \frac{1}{e^{-\frac{k^2}{9}}}} \right] \\ &\leq m + 4\sqrt{n} + \sqrt{n} \cdot \sum_{k=3}^{\lceil e^{1-e} \sqrt{n} \rceil} e^{-\frac{k^2}{9}} \leq (1 - e^{1-e})n + O(\sqrt{n}), \end{aligned}$$

where we used the fact that $\sum_{k=0}^\infty e^{-\frac{k^2}{9}}$ converges to a positive constant. ◀

A similar argument can be used to prove that our analysis is tight up to lower order terms.

► **Lemma 11.** *The online greedy algorithm matches at least $(1 - \frac{e}{e^e}) \cdot n - O(\sqrt{n})$ ads in expectation, with the instance of Definition 8.*

5.2 Maximum Matching

We now prove that the instance of Definition 8 admits a matching of $(1 - o(1)) \cdot n$ edges.

We will analyze our instance using the Karp-Sipser algorithm; first, we will reduce our instance to a configuration model; then, we will study this configuration model with a modification of the analysis in [7] of the first phase of the Karp-Sipser algorithm. In order to implement this plan, we need to modify our instance so that its average degree is bounded.

► **Definition 12.** Given any integer $\Delta \geq 2$, we consider the irregular cuckoo hashing distribution $\text{ICH}(n, n, D_\Delta)$, where D_Δ is defined as

$$\Pr[D_\Delta = 0] = \frac{1}{\Delta}, \quad \text{and} \quad \Pr[D_\Delta = i] = \frac{1}{i \cdot (i-1)} \quad \forall i \in \{2, \dots, \Delta\}.$$

Observe that, to sample an instance of Definition 12, one can first sample an instance of Definition 8 and then remove all edges incident on users of degree larger than Δ . We will prove the following result in Section 5.2.1.

► **Lemma 13.** With probability $1 - o(1)$, the instance of Definition 12 admits a matching of size $(1 - \frac{1}{\Delta})n - o(n)$.

The existence of a quasi-complete matching in our original instance can be easily derived from Lemma 13.

► **Theorem 14.** With probability $1 - o(1)$, the instance of Definition 8 admits a matching of size $(1 - o(1))n$.

We also note that our main Theorem follows from Theorems 9 and 14:

► **Theorem 15.** The optimal competitive ratio for the online matching problem with irregular cuckoo hashing distributions is no better than $1 - \frac{e}{e^e} + o(1) \approx 0.8206259 \dots$

Clearly, Theorem 15 directly applies to the cases of “IID users with known distribution”, “IID users with unknown distribution” and “adversarial graphs whose users are permuted uniformly at random”.

Moreover, given that Theorem 14 guarantees that the maximum matching has size $(1 - o(1)) \cdot n$ with probability at least $1 - o(1)$, our upper bound of $1 - \frac{e}{e^e} + o(1)$ holds for both the ratios $\frac{\mathbb{E}[\text{ALG}]}{\mathbb{E}[\text{OPT}]}$ and $\mathbb{E}[\frac{\text{ALG}}{\text{OPT}}]$,⁷ that is, for both the ratio-of-the-expectations and the expectation-of-the-ratio definitions of competitive ratio.

5.2.1 Reduction to the Configuration Model

In this Subsection, we aim to prove Lemma 13. We start by showing that, if we condition on the degrees of the vertices, then our instance is distributed like a configuration model. The following Lemma is the analogue of many that have been proved for various random graph models; a famous example is that of random regular graphs [41].

► **Lemma 16.** For any non-negative integers $\mathbf{d} = (d_1, d_2, \dots, d_n)$, $\hat{\mathbf{d}} = (\hat{d}_1, \hat{d}_2, \dots, \hat{d}_n)$ such that $\sum_{i=1}^n d_i = \sum_{i=1}^n \hat{d}_i$, and for any bipartite graph G , it holds,

$$\Pr_{X \sim \text{ICH}(n, n, P)} \left[X = G \mid \deg(v_i) = d_i \wedge \deg(\hat{v}_i) = \hat{d}_i \text{ for each } i \in [n] \right] = \Pr_{X \sim \text{CM}(\mathbf{d}, \hat{\mathbf{d}})} [X = G],$$

where P is such that $\Pr[P = d_i] > 0$ for each $i \in [n]$, so that the event $\{\deg(v_i) = d_i \wedge \deg(\hat{v}_i) = \hat{d}_i \text{ for each } i \in [n]\}$ happens with positive probability.

⁷ If $\text{OPT} = 0$, the generic online algorithm necessarily returns the maximum (empty) matching; it is typical to define $\frac{\text{ALG}}{\text{OPT}} = 1$ in this borderline case. Moreover, if $\mathbb{E}[\text{OPT}] = 0$ then necessarily $\Pr[\text{OPT} = 0] = 1$, so that it is natural to define $\frac{\mathbb{E}[\text{ALG}]}{\mathbb{E}[\text{OPT}]} = 1$ in this other case, as well.

We now prove that, in our instance, the number of vertices with a certain degree is concentrated around the mean. This will allow us to restrict our attention to a specific configuration model. Define, for $i \geq 0$,

$$z_i = \begin{cases} \frac{1}{\Delta} & \text{if } i = 0 \\ 0 & \text{if } i = 1 \text{ or } i > \Delta \\ \frac{1}{i(i-1)} & \text{if } 2 \leq i \leq \Delta \end{cases} \quad \hat{z}_i = \frac{H_{\Delta-1}^i e^{-H_{\Delta-1}}}{i!}. \quad (2)$$

The following Observation can be proved by applying the concentration bounds in Section 4.

► **Observation 17.** *Consider an irregular cuckoo hashing instance $\text{ICH}(n, n, P)$ such that P is upper bounded by a constant $\Delta > 0$. Let $p_i = \Pr[P = i]$ for $0 \leq i \leq \Delta$, $\mu = \mathbb{E}[\deg(v)] = \sum_{i=0}^{\Delta} ip_i$, and let $\hat{p}_i = \frac{e^{-\mu} \cdot \mu^i}{i!}$, $i \geq 0$, be a Poisson distribution with mean μ . We have that, with probability $1 - o(1)$,*

1. $|E| = (\mu \pm o(n^{-1/5}))n$,
2. $n_i = (p_i \pm o(n^{-1/5}))n$ for $0 \leq i \leq \Delta$,
3. $\hat{n}_i = (\hat{p}_i \pm o(n^{-1/5}))n$ for $i \geq 0$.

Note that Observation 17 applies to our graph with $p_i = z_i$ and $\hat{p}_i = \hat{z}_i$. Therefore, we can now leverage on Theorem 3, that is, a variant of the result in [7], which uses the Karp-Sipser algorithm to bound the maximum matching in configuration models whose degrees follow the conditions of Observation 17. Note that Theorem 3 requires a bounded average degree; we made sure that our random graph has bounded average degree by limiting its maximum user degree. Formally,

► **Lemma 18.** *For any constant $\Delta \geq 2$, let $\mathbf{d} \in \mathbb{Z}_{\geq 0}^n$ and $\hat{\mathbf{d}} \in \mathbb{Z}_{\geq 0}^n$ be any two valid degree sequences such that (i) $|E| = (H_{\Delta-1} \pm o(n^{-1/5}))n$, (ii) $n_i = (z_i \pm o(n^{-1/5}))n$, and (iii) $\hat{n}_i = (\hat{z}_i \pm o(n^{-1/5}))n$ for $i \geq 0$, where z_i and \hat{z}_i are defined in Equation (2). Then, with probability $1 - o(1)$, $\text{CM}(\mathbf{d}, \hat{\mathbf{d}})$ has a matching of size at least $(1 - 1/\Delta - o(1))n$.*

Proof. We apply Theorem 3 with z_i and \hat{z}_i as defined in Equation (2). Let $M = \ln(n)$, and without loss of generality, take $n > e^{\Delta+2} = O(1)$, so that $M > \Delta + 2 \geq H_{\Delta-1} + 3$. We have, $\sum_{i=M}^{\infty} i^2 \hat{z}_i \leq \sum_{i=M}^{\infty} \frac{e^{-H_{\Delta-1}} (H_{\Delta-1})^i}{(i-3)!} = (H_{\Delta-1})^3 \sum_{i=M-3}^{\infty} \frac{e^{-H_{\Delta-1}} (H_{\Delta-1})^i}{i!} = (H_{\Delta-1})^3 \Pr[\text{Poisson}(H_{\Delta-1}) \geq M-3] \leq \frac{(H_{\Delta-1})^M \cdot e^{M-3-H_{\Delta-1}}}{(M-3)^{M-3}} = \frac{2^{\Theta(\log n)}}{\Theta(\log n)^{\Theta(\log n)}} \leq o(n^{-1/5})$, where we used the standard tail bound $\Pr[\text{Poisson}(\lambda) \geq x] \leq \frac{(e\lambda)^x e^{-\lambda}}{x^x}$ for $x > \lambda$ (see, e.g., [54, Theorem 5.4]). We have,

$$\begin{aligned} f(x) &= \frac{1}{\Delta} + \sum_{i=2}^{\Delta} \left(\frac{1}{i(i-1)} \cdot x^i \right) \\ f'(x) &= \sum_{i=1}^{\Delta-1} \frac{x^i}{i} = \begin{cases} H_{\Delta-1} & \text{if } x = 1 \\ -\ln(1-x) - \sum_{i=\Delta}^{\infty} \frac{x^i}{i} & \text{if } |x| < 1 \end{cases} \\ \hat{f}(x) &= \sum_{i=0}^{\infty} \frac{H_{\Delta-1}^i e^{-H_{\Delta-1}}}{i!} \cdot x^i = e^{(x-1)H_{\Delta-1}} \cdot \sum_{i=0}^{\infty} \frac{(x \cdot H_{\Delta-1})^i e^{-x \cdot H_{\Delta-1}}}{i!} = e^{(x-1) \cdot H_{\Delta-1}} \\ \hat{f}'(x) &= H_{\Delta-1} \cdot e^{(x-1) \cdot H_{\Delta-1}}, \end{aligned}$$

where we used $\sum_{i=0}^{\infty} \frac{\lambda^i e^{-\lambda}}{i!} = 1$ for all λ and $\sum_{i=1}^{\infty} \frac{x^i}{i} = -\ln(1-x)$ for all $|x| < 1$. Note that $f'(1) = \hat{f}'(1) = H_{\Delta-1} < \infty$, $\hat{f}'(2) = H_{\Delta-1} e^{H_{\Delta-1}} = O(1)$, and clearly $f'(2) =$

$\sum_{i=1}^{\Delta} z_i \cdot i \cdot 2^{i-1} = O(1)$. Moreover, $|V| = |\hat{V}| = n = f(1)n = \hat{f}(1)n$. Therefore, the hypotheses of Theorem 3 are met, we now compute \hat{w}_1 and w_2 to apply the Theorem. Recall that $\hat{w}_1 = \frac{\hat{f}'(w_2)}{H_{\Delta-1}} = e^{(w_2-1) \cdot H_{\Delta-1}}$. Since $w_2 \in [0, 1]$, we have that $\hat{w}_1 > 0$, and therefore $1 - \hat{w}_1 \in [0, 1)$. Substituting this expression for \hat{w}_1 into the relation $w_2 = 1 - \frac{f'(1-\hat{w}_1)}{H_{\Delta-1}}$, we get,

$$w_2 = 1 - \frac{-\ln(e^{(w_2-1)H_{\Delta-1}}) - \sum_{i=\Delta}^{\infty} \frac{(1-e^{(w_2-1)H_{\Delta-1}})^i}{i}}{H_{\Delta-1}} = w_2 + \sum_{i=\Delta}^{\infty} \frac{(1 - e^{(w_2-1)H_{\Delta-1}})^i}{i \cdot H_{\Delta-1}}$$

Equivalently, $\sum_{i=\Delta}^{\infty} \frac{(1-e^{(w_2-1)H_{\Delta-1}})^i}{i} = 0$ and this equation is satisfied only for $w_2 = 1$. Indeed, for $x \in [0, 1)$, $\frac{(1-e^{(x-1)H_{\Delta-1}})^{\Delta}}{\Delta} > 0$. Therefore, $w_2 = \hat{w}_1 = 1$. Now, Theorem 3 ensures that the configuration model admits a matching of size at least,

$$(f(1) - f(0) - f'(0) - o(1))n = (1 - 1/\Delta - o(1))n,$$

and the proof is complete. \blacktriangleleft

We now apply Lemma 16, Observation 17 and Lemma 18, to prove Lemma 13, and conclude the proof of our impossibility result.

Proof of Lemma 13. By Observation 17, there exists $h(n) = o(n^{4/5})$ such that, in the instance of Definition 12, with probability $1 - o(1)$, (i) $|n_i - nz_i| \leq h(n)$, (ii) $|\hat{n}_i - n\hat{z}_i| \leq h(n)$ for all $i \geq 0$, and (iii) $||E| - nH_{\Delta-1}| \leq h(n)$. Let D_n be the set of all sequences $\mathbf{d} \in \mathbb{Z}_{\geq 0}^n$, $\hat{\mathbf{d}} \in \mathbb{Z}_{\geq 0}^n$ that are compatible with conditions (i), (ii), and (iii). Note that $D_n \neq \emptyset$ given that with probability $1 - o(1)$, the sampled graph respects the conditions. Let L be the distribution defined in Definition 12. For a graph G , let $M(G)$ be the event that the maximum matching of G has size at least $(1 - 1/\Delta - o(1))n$ (where the $o(1)$ term is the one given by Lemma 18), let $\text{good}(G)$ be the event that G respects conditions (i), (ii), and (iii), and let $D(G, \mathbf{d}, \hat{\mathbf{d}})$ be the event that the vertices of G have degrees \mathbf{d} and $\hat{\mathbf{d}}$. Note that $\{D(G, \mathbf{d}, \hat{\mathbf{d}})\}_{(\mathbf{d}, \hat{\mathbf{d}}) \in D_n}$ is a partition of $\text{good}(G)$. We have,

$$\begin{aligned} \Pr_{G \sim L} [M(G)] &\geq \Pr_{G \sim L} [\text{good}(G)] \Pr_{G \sim L} [M(G) \mid \text{good}(G)] \\ &\geq (1 - o(1)) \sum_{(\mathbf{d}, \hat{\mathbf{d}}) \in D_n} \Pr_{G \sim L} [M(G) \mid D(G, \mathbf{d}, \hat{\mathbf{d}})] \Pr_{G \sim L} [D(G, \mathbf{d}, \hat{\mathbf{d}}) \mid \text{good}(G)] \\ &= (1 - o(1)) \sum_{(\mathbf{d}, \hat{\mathbf{d}}) \in D_n} \Pr_{G \sim \text{CM}(\mathbf{d}, \hat{\mathbf{d}})} [M(G)] \Pr_{G \sim L} [D(G, \mathbf{d}, \hat{\mathbf{d}}) \mid \text{good}(G)] \\ &\geq (1 - o(1))^2 \sum_{(\mathbf{d}, \hat{\mathbf{d}}) \in D_n} \Pr_{G \sim L} [D(G, \mathbf{d}, \hat{\mathbf{d}}) \mid \text{good}(G)] = 1 - o(1), \end{aligned}$$

where in the first equality we use Lemma 16 and in the last inequality we apply Lemma 18. \blacktriangleleft

6 Extremality

We prove in this section that the construction of Definition 8 is extremal among the irregular cuckoo hashing instances in the following sense: for each $d = 0, 1, 2, \dots$, our distribution (greedily) maximizes the total number of users of degree at most d while guaranteeing that a $1 - o(1)$ fraction of these users can be matched to ads, that is, while guaranteeing that the resulting graph admits a quasi-complete matching. The distribution, then, aims to make

the user degrees as small as possible (so to make the task of the online algorithm as hard as possible), while guaranteeing that offline algorithms can match a $1 - o(1)$ fraction of the users.

We will prove that, for each fixed $\Delta \geq 2$, if we increase the probability that Definition 12 assigns to user-degree Δ by any small enough positive constant $\epsilon > 0$, while keeping the probabilities of user-degrees $1, \dots, \Delta - 1$ fixed, then a constant fraction of the users having degree in $\{1, 2, \dots, \Delta\}$ cannot be matched. In particular, we will consider the following class of distributions.

► **Definition 19.** Choose an integer constant $\Delta \geq 2$ and a small enough constant $\epsilon > 0$. Consider the irregular cuckoo hashing instance $\text{ICH}(n, n, D_{\Delta, \epsilon})$, where $D_{\Delta, \epsilon}$ is defined as,

$$\begin{aligned} \Pr[D_{\Delta, \epsilon} = 0] &= \Pr[D_{\Delta} = 0] - \epsilon \\ \Pr[D_{\Delta, \epsilon} = i] &= \Pr[D_{\Delta} = i] \quad \forall i \in \{1, 2, \dots, \Delta - 1\} \\ \Pr[D_{\Delta, \epsilon} = \Delta] &= \Pr[D_{\Delta} = \Delta] + \epsilon, \end{aligned}$$

where D_{Δ} is given in Definition 12.

We will again analyze the first phase of the Karp-Sipser algorithm by using Theorem 3.

► **Lemma 20.** With probability $1 - o(1)$, the instance of Definition 19 for any constant $\Delta \geq 2$ and for a small enough constant $\epsilon > 0$ with respect to Δ , does not admit matchings of size larger than $(1 - \frac{1}{\Delta} + \epsilon - \psi_{\Delta}(\epsilon) + o(1)) \cdot n$ where $\psi_{\Delta}(\epsilon) > 0$ for ϵ sufficiently small compared to Δ . Specifically, $\psi_{\Delta}(\epsilon) = \frac{\Delta^{2\Delta}}{\Delta+1} \cdot \epsilon^{\Delta+1} \pm O(c_{\Delta} \cdot \epsilon^{\Delta+2})$, where c_{Δ} depends only on Δ . In particular, with probability $1 - o(1)$, each matching will leave at least $(\psi_{\Delta}(\epsilon) - o(1)) \cdot n$ users of positive degree unmatched.

We can now use the previous Lemma to prove the main result of this Section.

► **Theorem 21.** Suppose that there exists a constant $i \geq 1$ such that $\Pr[P \leq j] = 1 - \frac{1}{j}$ for each $j \in \{1, 2, \dots, i - 1\}$, and $\Pr[P \leq i] > 1 - \frac{1}{i}$. Then, with probability $1 - o(1)$, the instance $\text{ICH}(n, n, P)$ admits no matching that matches at least $\Pr[P \leq i] \cdot n - o(n)$ users of degree $\leq i$.

Our distribution, then, is extremal among the irregular cuckoo hashing distributions that make it possible to match $n - o(n)$ users to the n ads; indeed, it has $\Pr[P \leq i] = 1 - \frac{1}{i}$ for each integer $i \geq 1$, and thus it greedily maximizes, for each $d = 0, 1, 2, \dots$, the number of users of degree d under the constraint that a $1 - o(1)$ fraction of the users of degree at most d can be matched.

7 Generalization to Non-Equibipartite Graphs

In this Section, we consider a more general case, where we have n ads and $u \cdot n$ users, for a constant $u \in (0, 1]$. As in the instance of Definition 8, we will greedily add users of smallest degree, while guaranteeing the existence of a quasi-complete matching. Our generalized instances have varying competitive ratios; we will show that the worst competitive ratio is obtained at $u = 1$,⁸ that is, with the original instance of Definition 8 (see Figure 1 and Theorem 27).

⁸ We do not consider explicitly the case where there are more users than ads, that is, the case $u > 1$. This case can be easily reduced to the case of $u = 1$, i.e., to the case of Definition 8. Indeed, if $u > 1$ and one

► **Definition 22** (Generalized Instance). *Let $u \in (0, 1]$ be a constant. Consider the irregular cuckoo hashing instance $\text{ICH}(\lceil u \cdot n \rceil, n, D_u)$. The distribution D_u is defined as,*

- *if $u = 1$, D_u is equal to the instance of Definition 8, that is, $\Pr[D_u = d] = \frac{1}{d \cdot (d-1)}$ for $d \geq 2$,*
- *if $u \in (0, 1)$, let Δ be the minimum integer such that $u \leq 1 - \frac{1}{\Delta}$, that is, let $\Delta = \left\lceil \frac{1}{1-u} \right\rceil$. Then, $\Delta \geq 2$. For each $d \in \{2, 3, \dots, \Delta - 1\}$, we set $\Pr[D_u = d] = \frac{1}{u \cdot d \cdot (d-1)}$, and $\Pr[D_u = \Delta] = 1 - \frac{1}{u} \cdot \left(1 - \frac{1}{\Delta-1}\right)$.*

In particular, the distribution D_u is obtained by conditioning the D of Definition 8 on an event of probability u , chosen to include the smallest possible values of D 's support. Observe that, for each $u \in (0, \frac{1}{2}]$, the distribution D_u assigns probability 1 to degree 2. Observe, further, that as u increases from $\frac{d-1}{d}$ to $\frac{d}{d+1}$, for any integer $d \geq 2$, D_u assigns more and more probability to degree $d + 1$.

The following claim is a corollary of the analysis of the maximum matching for the case $u = 1$.

► **Corollary 23.** *The maximum matching of the instance of Definition 22 has size $u \cdot n - o(n)$ with probability $1 - o(1)$; thus, the resulting graph admits a quasi-complete matching.*

One can prove that the D_u construction is extremal using the same argument we developed in the last section for the case $u = 1$ (i.e., moving any constant mass from higher degrees to smaller degrees makes the quasi-complete matching lose a constant fraction of the edges).

► **Corollary 24.** *Suppose that there exists a constant $i \geq 1$ such that $\Pr[P \leq j] = \Pr[D_u \leq j]$ for each $j \in \{1, 2, \dots, i - 1\}$, and $\Pr[P \leq i] > \Pr[D_u \leq i]$. Then in the instance $\text{ICH}(\lceil u \cdot n \rceil, n, P)$, with probability $1 - o(1)$, there exists no matching that matches at least $\Pr[P \leq i] \cdot u \cdot n - o(n)$ users of degree $\leq i$.*

We now move on to studying the performance of online algorithms on the generalized instance of Definition 22. We start with a technical lemma that relates the expected number of users to capture a new ad with distribution D_u , and with distribution D . This Lemma will be used to bound the number of users needed to capture a new ad with distribution D_u , in terms of the number of users needed with D – we need this Lemma because the integral required to determine the number of users to capture an ad with D_u is not easy to solve, while the one with D has already been solved in the proof of Theorem 9.

► **Lemma 25.** *Fix $u \in (\frac{1}{2}, 1)$ and let $\Delta = \left\lceil \frac{1}{1-u} \right\rceil$, $c_i = \frac{1}{u \cdot i \cdot (i-1)}$ for $i \in \{2, \dots, \Delta - 1\}$ and $c_\Delta = 1 - \sum_{i=2}^{\Delta-1} \frac{1}{u \cdot i \cdot (i-1)} = 1 - \frac{\Delta-2}{u \cdot (\Delta-1)}$.⁹ Moreover, let $x \in [0, u)$. Then,*

$$\frac{1}{1 - \sum_{i=2}^{\Delta} (c_i \cdot x^i)} \leq \frac{1 - \frac{1}{8(1+\ln 2)\Delta} \cdot \left[x \geq \frac{u}{2}\right]}{\left(1 - \frac{x}{u}\right) \left(1 - \ln\left(1 - \frac{x}{u}\right)\right)}.$$

We now study the performance of online algorithms on the instance D_u for $u < 1$.

aims to greedily add (fractions of) users of the minimum degree that allow a quasi-complete matching to exist, the first n users will have the degrees given by the distribution of Definition 8. As we proved in Theorem 14, these first n users induce a matching that is quasi-perfect, and thus quasi-complete from the ads side. The next minimum user degree which guarantees a quasi-complete matching is then 0. Thus, if $u > 1$, the greedy approach would produce $(u - 1) \cdot n$ users of degree 0, and $\frac{1}{d \cdot (d-1)} \cdot n$ users of degree d for each $d \geq 2$. The resulting graph will then induce online and offline matching sizes equal to those of Definition 8.

⁹ Observe that $c_i = \Pr[D_u = i] \forall i \in \{2, \dots, \Delta\}$, where D_u is the distribution of Definition 22.

► **Lemma 26.** Fix $u \in (0, 1)$. Then, the online greedy algorithm, if run on the D_u distribution of Definition 22, will match at least $(1 - e^{1-e} + \frac{1-u}{500}) \cdot u \cdot n - O(\sqrt{n \ln n})$ ads in expectation.

We can now conclude that $u = 1$ gives the strongest bound of $1 - \frac{e}{e^e}$:

► **Theorem 27.** The strongest impossibility result for the competitive ratio of online matching that can be obtained with the limiting constructions of Definition 22 is $1 - \frac{e}{e^e} \pm o(1)$. This bound can be obtained by setting $u = 1$, that is, with the construction of Definition 8.

8 Conclusion

We proved that the optimal competitive ratio of several online bipartite matching problems cannot be larger than $1 - \frac{e}{e^e}$. The simplicity of our upper bound expression, and its similarity to the optimal competitive ratio for adversarial graphs that are also adversarially permuted, naturally raise the question of its potential optimality. We leave as open questions whether this upper bound is tight in (i) irregular cuckoo hashing instances, in (ii) IID instances with known distribution, in (iii) IID instances with unknown distribution, and in the hardest case, that of (iv) adversarial graphs with users permuted uniformly at random.

References

- 1 Jose Adell and Pedro Jodra. Exact Kolmogorov and total variation distances between some familiar discrete distributions. *Journal of Inequalities and Applications*, 2006. doi:10.1155/JIA/2006/64307.
- 2 Mohit Agarwal and Anuj Puri. Base station scheduling of requests with fixed deadlines. In *INFOCOM*, 2002. doi:10.1109/INFCOM.2002.1019293.
- 3 Antonios Antoniadis, Themis Gouleakis, Pieter Kleer, and Pavel Kolev. Secretary and online matching problems with machine learned advice. *Discrete Optimization*, 48, 2023. doi:10.1016/J.DISOPT.2023.100778.
- 4 Jonathan Aronson, Alan Frieze, and Boris G. Pittel. Maximum matchings in sparse random graphs: Karp–Sipser revisited. *Random Structures & Algorithms*, 12(2), 1998. doi:10.1002/(SICI)1098-2418(199803)12:2<3C111::AID-RSA1%3E3.0.CO;2-%23.
- 5 Yossi Azar and Yoel Chaiutin. Optimal node routing. In *STACS*, 2006. doi:10.1007/11672142_49.
- 6 Bahman Bahmani and Michael Kapralov. Improved bounds for online stochastic matching. In *ESA*, 2010. doi:10.1007/978-3-642-15775-2_15.
- 7 Paul Balister and Stefanie Gerke. Controllability and matchings in random bipartite graphs. *Surveys in Combinatorics*, 2015. doi:10.1017/CB09781316106853.004.
- 8 Benjamin Birnbaum and Claire Mathieu. On-line bipartite matching made simple. *SIGACT News*, 39(1), 2008. doi:10.1145/1360443.1360462.
- 9 Joakim Blikstad, Ola Svensson, Radu Vintan, and David Wajc. Online edge coloring is (nearly) as easy as offline. In *STOC*, 2024. doi:10.1145/3618260.3649741.
- 10 Avrim Blum, Tuomas Sandholm, and Martin Zinkevich. Online algorithms for market clearing. *J. ACM*, 53(5), 2006. doi:10.1145/1183907.1183913.
- 11 Tom Bohman and Alan Frieze. Karp–Sipser on random graphs with a fixed degree sequence. *Combinatorics, Probability and Computing*, 20(5), 2011. doi:10.1017/S0963548311000265.
- 12 Allan Borodin and Calum MacRury. Online bipartite matching in the probe-commit model. *Mathematical Programming*, 2025. doi:10.1007/s10107-024-02184-y.
- 13 Allan Borodin, Calum MacRury, and Akash Rakheja. Secretary Matching Meets Probing with Commitment. In *APPROX/RANDOM*, 2021. doi:10.4230/LIPIcs.APPROX/RANDOM.2021.13.
- 14 Allan Borodin, Calum MacRury, and Akash Rakheja. Prophet Matching in the Probe-Commit Model. In *APPROX/RANDOM*, 2022. doi:10.4230/LIPIcs.APPROX/RANDOM.2022.46.
- 15 Mark Braverman, Mahsa Derakhshan, and Antonio Molina Lovett. Max-weight online stochastic matching: Improved approximations against the online benchmark. In *EC*, 2022. doi:10.1145/3490486.3538315.

- 16 Mark Braverman, Mahsa Derakhshan, Tristan Pollner, Amin Saberi, and David Wajc. New philosopher inequalities for online bayesian matching, via pivotal sampling. In *SODA*, 2025. doi:10.1137/1.9781611978322.98.
- 17 Brian Brubach, Karthik Abinav Sankararaman, Aravind Srinivasan, and Pan Xu. New Algorithms, Better Bounds, and a Novel Model for Online Stochastic Matching. In *ESA*, 2016. doi:10.4230/LIPIcs.ESA.2016.24.
- 18 Niv Buchbinder, Joseph Naor, and David Wajc. Lossless online rounding for online bipartite matching (despite its impossibility). In *SODA*, 2023. doi:10.1137/1.9781611977554.ch78.
- 19 Niv Buchbinder, Danny Segev, and Yevgeny Tkach. Online algorithms for maximum cardinality matching with edge arrivals. *Algorithmica*, 81(5), 2019. doi:10.1007/s00453-018-0505-7.
- 20 Davin Choo, Themistoklis Gouleakis, Chun Kai Ling, and Arnab Bhattacharyya. Online bipartite matching with imperfect advice. In *ICML*, 2024.
- 21 Nikhil R. Devanur, Kamal Jain, and Robert D. Kleinberg. Randomized primal-dual analysis of ranking for online bipartite matching. In *SODA*, 2013. doi:10.1137/1.9781611973105.7.
- 22 John P. Dickerson, Karthik A. Sankararaman, Aravind Srinivasan, and Pan Xu. Allocation problems in ride-sharing platforms: Online matching with offline reusable resources. *ACM Trans. Econ. Comput.*, 9(3), 2021. doi:10.1145/3456756.
- 23 Martin Dietzfelbinger, Andreas Goerdts, Michael Mitzenmacher, Andrea Montanari, Rasmus Pagh, and Michael Rink. Tight thresholds for cuckoo hashing via xorsat. In *ICALP*, 2010. doi:10.1007/978-3-642-14165-2_19.
- 24 Martin Dietzfelbinger and Christoph Weidling. Balanced allocation and dictionaries with tightly packed constant size bins. *Theoretical Computer Science*, 380(1), 2007. doi:10.1016/j.tcs.2007.02.054.
- 25 Devdatt Dubhashi and Alessandro Panconesi. *Concentration of Measure for the Analysis of Randomized Algorithms*. Cambridge University Press, 2009. doi:10.1017/CB09780511581274.
- 26 Alon Eden, Michal Feldman, Amos Fiat, and Kineret Segal. An economics-based analysis of ranking for online bipartite matching. In *SOSA*, 2021. doi:10.1137/1.9781611976496.12.
- 27 Paul Erdős and Alfréd Rényi. On the evolution of random graphs. *Publ. math. inst. hung. acad. sci.*, 5(1), 1960.
- 28 Tomer Ezra, Michal Feldman, Nick Gravin, and Zhihao Gavin Tang. Prophet matching with general arrivals. *Math. Oper. Res.*, 47(2), 2022. doi:10.1287/moor.2021.1152.
- 29 Matthew Fahrbach, Zhiyi Huang, Runzhou Tao, and Morteza Zadimoghaddam. Edge-weighted online bipartite matching. *J. ACM*, 69(6), 2022. doi:10.1145/3556971.
- 30 Yiding Feng, Rad Niazadeh, and Amin Saberi. Two-stage stochastic matching with application to ride hailing. In *SODA*, 2021. doi:10.1137/1.9781611976465.170.
- 31 Yiding Feng, Rad Niazadeh, and Amin Saberi. Two-stage stochastic matching and pricing with applications to ride hailing. *Oper. Res.*, 72(4), 2024. doi:10.1287/opre.2022.2398.
- 32 Buddhima Gamlath, Michael Kapralov, Andreas Maggiori, Ola Svensson, and David Wajc. Online Matching with General Arrivals . In *FOCS*, 2019. doi:10.1109/FOCS.2019.00011.
- 33 Ruiquan Gao, Zhongtian He, Zhiyi Huang, Zipei Nie, Bijun Yuan, and Yan Zhong. Improved online correlated selection. In *FOCS*, 2022. doi:10.1109/FOCS52979.2021.00123.
- 34 Gagan Goel and Aranyak Mehta. Online budgeted matching in random input models with applications to adwords. In *SODA*, 2008. URL: <http://dl.acm.org/citation.cfm?id=1347082.1347189>.
- 35 Daniel Hathcock, Billy Jin, Kalen Patton, Sherry Sarkar, and Michael Zlatin. The online submodular assignment problem. In *FOCS*, 2024. doi:10.1109/FOCS61266.2024.00026.
- 36 Zhiyi Huang and Xinkai Shu. Online stochastic matching, poisson arrivals, and the natural linear program. In *STOC*, 2021. doi:10.1145/3406325.3451079.
- 37 Zhiyi Huang, Xinkai Shu, and Shuyi Yan. The power of multiple choices in online stochastic matching. In *STOC*, 2022. doi:10.1145/3519935.3520046.
- 38 Zhiyi Huang, Zhihao Gavin Tang, and David Wajc. Online matching: A brief survey. *SIGecom Exch.*, 22(1), 2024. doi:10.1145/3699824.3699837.

- 39 Patrick Jaillet and Xin Lu. Online stochastic matching: New algorithms with better bounds. *Mathematics of Operations Research*, 39(3), 2014. doi:10.1287/moor.2013.0621.
- 40 Svante Janson. Tail bounds for sums of geometric and exponential variables. *Statistics & Probability Letters*, 135, 2018. doi:10.1016/j.spl.2017.11.017.
- 41 Svante Janson, Tomasz Luczak, and Andrzej Rucinski. *Random Graphs*. Wiley, 2000. doi:10.1002/9781118032718.
- 42 Billy Jin and Will Ma. Online bipartite matching with advice: Tight robustness-consistency tradeoffs for the two-stage model. In *NeurIPS*, 2022.
- 43 Yossi Kanizo, David Hay, and Isaac Keslassy. Maximum bipartite matching size and application to cuckoo hashing, 2011. arXiv:1007.1946.
- 44 Haim Kaplan, David Naori, and Danny Raz. Online weighted matching with a sample. In *SODA*, 2022. doi:10.1137/1.9781611977073.52.
- 45 Chinmay Karande, Aranyak Mehta, and Pushkar Tripathi. Online bipartite matching with unknown distributions. In *STOC*, 2011. doi:10.1145/1993636.1993715.
- 46 Richard M. Karp and Michael Sipser. Maximum matching in sparse random graphs. In *SFCS*, 1981. doi:10.1109/SFCS.1981.21.
- 47 Richard M. Karp, Umesh V. Vazirani, and Vijay V. Vazirani. An optimal algorithm for on-line bipartite matching. In *STOC*, 1990. doi:10.1145/100216.100262.
- 48 Will Ma and Pan Xu. Promoting fairness among dynamic agents in online-matching markets under known stationary arrival distributions. In *NeurIPS*, 2024.
- 49 Mohammad Mahdian and Qiqi Yan. Online bipartite matching with random arrivals: an approach based on strongly factor-revealing LPs. In *STOC*, 2011. doi:10.1145/1993636.1993716.
- 50 Vahideh H. Manshadi, Shayan Oveis Gharan, and Amin Saberi. Online stochastic matching: online actions based on offline statistics. In *SODA*, 2011. doi:10.1137/1.9781611973082.98.
- 51 Vahideh H. Manshadi, Shayan Oveis Gharan, and Amin Saberi. Online stochastic matching: Online actions based on offline statistics. *Mathematics of Operations Research*, 37(4), 2012. doi:10.1287/MOOR.1120.0551.
- 52 Aranyak Mehta. *Online Matching and Ad Allocation*. Now Publishers, 2013.
- 53 Aranyak Mehta, Amin Saberi, Umesh V. Vazirani, and Vijay V. Vazirani. Adwords and generalized on-line matching. In *FOCS*, 2005. doi:10.1109/SFCS.2005.12.
- 54 Michael Mitzenmacher and Eli Upfal. *Probability and Computing: Randomized Algorithms and Probabilistic Analysis*. Cambridge University Press, 2005. doi:10.1017/CB09780511813603.
- 55 Joseph (Seffi) Naor, Aravind Srinivasan, and David Wajc. Online dependent rounding schemes for bipartite matchings, with applications. In *SODA*, 2025. doi:10.1137/1.9781611978322.100.
- 56 Rasmus Pagh and Flemming Friche Rodler. Cuckoo hashing. *Journal of Algorithms*, 51(2), 2004. doi:10.1016/j.jalgor.2003.12.002.
- 57 Christos Papadimitriou, Tristan Pollner, Amin Saberi, and David Wajc. Online stochastic max-weight bipartite matching: Beyond prophet inequalities. In *EC*, 2021.
- 58 Barna Saha and Aravind Srinivasan. A new approximation technique for resource-allocation problems. *Random Structures & Algorithms*, 52(4), 2018. doi:10.1002/rsa.20756.
- 59 J. Michael Steele. Le Cam's Inequality and Poisson Approximations. *The American Mathematical Monthly*, 101(1), 1994. doi:10.1080/00029890.1994.11996904.
- 60 Zhihao Gavin Tang, Jinzhao Wu, and Hongxun Wu. (Fractional) online stochastic matching via fine-grained offline statistics. In *STOC*, 2022. doi:10.1145/3519935.3519994.
- 61 Zhihao Gavin Tang and Yuhao Zhang. Improved bounds for fractional online matching problems. In *EC*, 2024. doi:10.1145/3670865.3673459.
- 62 Yongxin Tong, Jieying She, Bolin Ding, Libin Wang, and Lei Chen. Online mobile Micro-Task Allocation in spatial crowdsourcing. In *ICDE*, 2016. doi:10.1109/ICDE.2016.7498228.