Online Matching with Stochastic Rewards: Advanced Analyses Using Configuration Linear Programs*

Zhiyi Huang[1], Hanrui Jiang[2], Aocheng Shen[2], Junkai Song[1], Zhiang Wu[3], and Qiankun Zhang[2](✉)

[1] The University of Hong Kong, Hong Kong SAR, China
zhiyi@cs.hku.hk, dsgsjk@connect.hku.hk
[2] Huazhong University of Science and Technology, Wuhan, China
{hanry, awsonshen, qiankun}@hust.edu.cn
[3] Hong Kong University of Science and Technology, Hong Kong SAR, China
zwube@connect.ust.hk

Abstract. Mehta and Panigrahi (2012) proposed Online Matching with Stochastic Rewards, which generalizes the Online Bipartite Matching problem of Karp, Vazirani, and Vazirani (1990) by associating the edges with success probabilities. This new feature captures the pay-per-click model in online advertising. Recently, Huang and Zhang (2020) studied this problem under the online primal dual framework using the Configuration Linear Program (LP), and got the best known competitive ratios of the Stochastic Balance algorithm. Their work suggests that the more expressive Configuration LP is more suitable for this problem than the Matching LP.

This paper advances the theory of Configuration LP in two directions. Our technical contribution includes a characterization of the joint matching outcome of an offline vertex and *all its neighbors*. This characterization may be of independent interest, and is aligned with the spirit of Configuration LP. By contrast, previous analyses of Ranking generally focus on only one neighbor. Second, we designed a Stochastic Configuration LP that captures a stochastic benchmark proposed by Goyal and Udwani (2020), who used a Path-based LP. The Stochastic Configuration LP is smaller and simpler than the Path-based LP. Moreover, using the new LP we improved the competitive ratio of Stochastic Balance from 0.596 to 0.611 when the success probabilities are infinitesimal, and to 0.613 when the success probabilities are further equal.

1 Introduction

Suppose that Alice is planning an upcoming trip to Hawaii and she just types "Hawaii resort" in a search engine. Once she presses the return key, the search

* Zhiyi Huang is supported by the Research Grants Council of Hong Kong (Grant No. 17203022). Qiankun Zhang is supported by National Natural Science Foundation of China (Grant No. 62302183) and National Science Foundation of Hubei Province (Grant No. 2023AFB258).

engine will provide a list of related websites, among which several websites have a small "Ad" tag next to them to indicate that they are sponsored results. We will refer to such websites as advertisers. If Alice clicks on a sponsored result, the advertiser would pay the search engine, where the amount depends on its bid for the keywords among other factors. This is one of many scenarios of online advertising, which contributes hundreds of billions of US dollars to the annual revenue of major IT companies.

Online advertising presents many unique challenges, each of which has led to a long line of research. For example, the search engine must decide which advertiser shall be returned to a search query without knowing what search queries may come next. The essence of this online decision-making problem is captured by Online Bipartite Matching, which amazingly was introduced by Karp, Vazirani, and Vazirani [22] well before online advertising even existed. To see the connection, consider the advertisers and search queries as vertices on the two sides of a bipartite graph. The advertisers are known from the beginning and we call them *offline vertices*. The search queries come one by one and we call them *online vertices*. The algorithm needs to decide how to match each online vertex on its arrival, with the goal of maximizing the matching size. We measure an algorithm by the worst-case ratio of its expected matching size to the size of the optimal matching in hindsight, a.k.a., the *competitive ratio*.

Readers may notice a disparity, however, if they compare our example scenario of online advertising with the Online Bipartite Matching problem. In online advertising, the search engine cannot control whether a user clicks on the ad or not. The best that a search engine could do is to estimate how likely the user will click, based on the keywords and the information it gathers about the user. To this end, an attempt to match an online vertex to an offline neighbor only succeeds with some probability, rather than with certainty like in Online Bipartite Matching. Because of this disparity, Mehta and Panigrahi [25] proposed Online Matching with Stochastic Rewards, extending Online Bipartite Matching by associating the edges with success probabilities.

Much progress has been made in Online Matching with Stochastic Rewards in the last decade. Most related to this paper is the work of Huang and Zhang [17], who successfully applied the online primal-dual framework to this problem. Their analysis of the Stochastic Balance algorithm, which matches each online vertex to a neighbor that has been matched the least number of times so far, yielded the best competitive ratios to date. Conceptually, they found that the Matching Linear Program (LP) is not expressive enough for an online primal-dual analysis of this problem, and instead one could consider the Configuration LP.

There are still unanswered questions, however, that cast some doubt on the usefulness of the Configuration LP. For instance, Mehta and Panigrahi [25] showed that the Ranking algorithm, which randomly ranks the offline vertices at the beginning and matches each online vertex to the neighbor with the highest rank, is another competitive algorithm for Online Matching with Stochastic Rewards. However, Huang and Zhang [17] failed to offer any non-trivial competitive analysis for the Ranking algorithm using the online primal-dual analysis with

Configuration LP. Moreover, Goyal and Udwani [13] introduced an alternative benchmark, which we will refer to as the stochastic benchmark (see Section 2 for details). They proved that the competitive ratios of both Stochastic Balance and Ranking would be strictly larger if they were compared against the stochastic benchmark. To do so, they designed a Path-based LP which is completely different than the Configuration LP. Can we reproduce or even improve these results by further developing the theory of online primal-dual with Configuration LP?

1.1 Our Results

Our first contribution is *an online primal-dual analysis of Ranking based on the Configuration LP*, improving the competitive ratio of Ranking from 0.534 to 0.572. To present the technical novelty that enables this analysis, we need to first explain the intuition behind the usefulness of the Configuration LP in an online primal-dual analysis compared to the simple Matching LP. The essence of an online primal-dual analysis is to design an appropriate gain-splitting rule. When the algorithm matches an online vertex v to an offline neighbor u and increases the expected objective accordingly by success probability p_{uv}, imagine that we further split the gain of p_{uv} between its two vertices u and v. We need the gain splitting rule to satisfy some LP-dependent conditions which essentially say that the total expected gain of neighboring vertices are sufficiently large. More precisely, if we use the Matching LP, the condition is:

$$p_{uv} \cdot \mathbf{E}\big[\text{gain of } u\big] + \mathbf{E}\big[\text{gain of } v\big] \geq \Gamma \cdot p_{uv} \ , \tag{1}$$

where Γ is the competitive ratio that we seek to prove for the algorithm.

By contrast, if we use the Configuration LP, the critical condition no longer considers just a single edge (u, v). Instead, it examines an offline vertex u and *a subset of neighbors* S whose success probabilities sum to about 1, and requires that:

$$\mathbf{E}\big[\text{gain of } u\big] + \sum_{v \in S} \mathbf{E}\big[\text{gain of } v\big] \geq \Gamma \ . \tag{2}$$

In other words, the configuration LP only needs an amortized version of Eqn. (1), which sums over $v \in S$. Even if Eqn. (1) fails in the worst-case scenario with respect to (w.r.t.) an offline vertex u and a single neighbor v, Eqn. (2) may still hold if such worst-case cannot happen to all $v \in S$ at the same time. Note that, however, to use the power of the weaker condition, we need to characterize the joint matching outcome of u and a subset of neighbors S, rather than doing so for each $v \in S$ separately. Huang and Zhang [17] gave such a characterization for Stochastic Balance.

This paper finds that even the joint matching outcome of u and S is insufficient for getting a useful characterization of the worst-case scenario of Eqn. (2) w.r.t. Ranking. Instead, we consider the joint outcome of u and *all its neighbors*, including those outside S. Characterizing the entire neighborhood's matching outcome may seem counter-intuitive *a priori*, which may be why Huang and

Zhang [17] failed to get a non-trivial competitive ratio for Ranking. The characterization and the new approach in general may be useful for other online matching problems whose offline vertices can be matched more than once, e.g., the AdWords problem [26].

Our second contribution is *a new Stochastic Configuration LP*, which we design to bound the stochastic benchmark of Goyal and Udwani [13]. The stochastic benchmark refers to the best objective achievable by an algorithm which knows the graph and the edges' success probabilities but not the realization of edge successes and failures, and which still needs to match online vertices by their arrival order. It is natural to consider LPs whose decision variables depend on the observed edge successes and failures to capture the stochastic benchmark. This is the approach of Goyal and Udwani [13], who referred to the observed edge successes and failures as a sample path and gave a Path-based LP. Their LP is huge, however, because the number of sample paths is exponential in the number of edges by definition. Further, compared to the Matching and Configuration LPs, the Path-based LP has an extra set of consistency constraints: if a sample path is a sub-path of another sample path, then the relevant matching outcome must be consistent. These constraints lead to new dual variables that are outside the scope of existing gain-splitting methods. Goyal and Udwani [13] still found a canonical gain-splitting rule and sufficient conditions for proving competitive ratios for Ranking and Stochastic Balance, but their conditions do not correspond to the constraints of Path-based LP's dual. Therefore, it is unclear how to apply their approach to other problems.

We use a different strategy to design Stochastic Configuration LP for the stochastic benchmark. Instead of using variables x_{uS} for the probability an offline vertex u is matched to a subset of online vertices S as in the Configuration LP, we consider variables y_{uS} that represents the probability an offline vertex u *would be matched to subset S if all matches to u failed*. The size of the Stochastic Configuration LP is therefore exponential only in the number of vertices, like the original Configuration LP. Further, these variables implicitly capture the consistency requirements and therefore no additional consistency constraints are needed. As a result, the online primal dual analysis can directly work with the (approximate) dual constraints. Compared to condition (2) given by the Configuration LP, the condition w.r.t. the stochastic variant scales the gains of online vertices and the competitive ratio on the right-hand-side by some multipliers smaller than 1, with the multiplier of competitive ratio being smaller. Hence, the resulting condition is indeed easier to satisfy.

Using the Stochastic Configuration LP, we improved the competitive ratios of Stochastic Balance w.r.t. the stochastic benchmark from 0.596 to 0.611 when the success probabilities are infinitesimal, and to 0.613 if the success probabilities

Table 1: Summary of Results. The non-stochastic and stochastic benchmarks are denoted as OPT and S-OPT respectively. The results for Stochastic Balance apply to infinitesimal success probabilities. We round the competitive ratios down to the third digit after the decimal point.

	OPT	S-OPT	
Ranking (Equal Prob.)	0.534 [25] $\to$ **0.572**	$1 - \frac{1}{e} \approx 0.632$ [13]	
Stochastic Balance (Equal Prob.)	0.576 [17]	0.596 [13]	$\to$ **0.613**
Stochastic Balance (General)	0.572 [17]	0.596 [13]	$\to$ **0.611**

are further equal.[4][5] We also reproduce the optimal $1 - \frac{1}{e}$ competitive ratio of Ranking w.r.t. the stochastic benchmark.

1.2 Related Works

Our analyses follow the (randomized) online primal dual framework by Devanur, Jain, and Kleinberg [7], who were inspired by Birnbaum and Mathieu [2] and Goel and Mehta [12].

An online advertising platform's objective is rarely as simple as maximizing the matching size. For instance, advertisers usually have different values for different keywords. Hence, the literature has subsequently introduced many variants of the original Online Bipartite Matching problem by Karp et al. [22]. Aggarwal et al. [1] generalized Ranking and the optimal $1 - \frac{1}{e}$ competitive ratio to the (offline) vertex-weighted problem. Feldman et al. [9] considered the edge-weighted problem with free disposal, and gave a $1 - \frac{1}{e}$ competitive algorithm assuming large capacities of offline vertices; see also Devanur et al. [5] for an online primal dual version of this result. Their algorithm may be viewed as a generalization of the Balance algorithm for unweighted matching by Kalyanasundaram and Pruh [20], which also inspired the Stochastic Balance algorithm by Mehta and Panigrahi [25] for Online Matching with Stochastic Rewards. Recently, Fahrbach et al .[8] gave the first algorithm that is strictly better than $\frac{1}{2}$-competitive for the edge-weighted problem without large-capacity assumption. The ratio was later improved independently to 0.509 by Shin and An [27], to 0.519 by Gao et al. [11], and to 0.536 by Blanc and Charikar [3]. Another important variant is the AdWords problem by Mehta et al. [26], where each offline vertex may be matched multiple times and has a budget-additive valuation. Under a small bid assumption, Mehta et al. [26] gave an optimal $1 - \frac{1}{e}$ competitive algorithm. See

[4] We remark that the competitive ratio of Goyal and Udwani [13] holds for the optimal value of the Path-based LP, which may be strictly larger than the optimal value of our Stochastic Configuration LP in some cases. Nevertheless, both LPs' optimal values upper bound the stochastic benchmark.

[5] We thank an anonymous reviewer for pointing out that a preliminary arXiv version (v5) of [13] claims a competitive ratio of 0.605 for the general infinitesimal case without proof. Nevertheless, our bound is also larger than that.

also Buchbinder, Jain, and Naor [4] for an online primal dual analysis of this algorithm, and Devanur and Jain [6] for a generalization called Online Matching with Concave Returns. Huang, Zhang, and Zhang [18] recently broke the $\frac{1}{2}$ barrier for AdWords with general bids and gave a 0.501-competitive algorithm, using an approach similar to Fahrbach et al. [8]. A recent line of results [23,29,30] studied budget-oblivious algorithms for AdWords, whose allocation policies do not depend on the budgets of offline vertices.

Besides the stochastic edge successes and failures in Online Matching with Stochastic Rewards, the literature has also studied other stochastic models of online matching problems. Mahdian and Yan [24] and Karande, Mehta, and Tripathi [21] independently showed that Ranking is strictly better than $1 - \frac{1}{e}$ competitive if online vertices arrive by random order. Huang et al. [16] extended this result to the vertex-weighted problem, which was further improved upon by Jin and Williamson [19]. The problem with stochastically generated graphs was first studied by Feldman et al. [10], who named it Online Stochastic Matching. We also refer readers to the most recent works on this problem by Huang and Shu [14], Huang, Shu, and Yan [15], and Tang, Wu, and Wu [28].

1.3 Paper Outline

Section 2 presents a formal definition of online matching with stochastic rewards, the benchmarks we will compare against, and existing linear programs and algorithms. Section 3 introduces our new stochastic Configuration LPs, and their properties. Section 4 shows the online primal-dual analyses of Ranking. Finally, our results for Stochastic Balance are deferred to full paper because of page limitation and the fact that the improvements mainly come from a combination of the structural lemmas by Huang and Zhang [17] and the new stochastic Configuration LP proposed in this paper. All proofs in this paper are deferred to full paper.

2 Preliminaries

We write x^+ for function $\max\{x, 0\}$. For any vector x and any index i, we write x_{-i} for the vector with the i-th entry removed, and (x_i, x_{-i}) for a vector whose i-th entry is x_i, and whose other entries are x_{-i}.

2.1 Model

Consider a bipartite graph $G = (U, V, E)$. Vertices in U are referred to as *offline vertices* because they are known to the algorithm from the beginning. Vertices in V arrive one by one, and are referred to as *online vertices*. We will write $v \prec v'$ to denote that v arrives before v'. Each edge (u, v) is associated with a success probability $0 \le p_{uv} \le 1$; further define $p_{uS} = \sum_{v \in S} p_{uv}$ for any $u \in U$ and any $S \subseteq V$. When an online vertex $v \in V$ arrives, the algorithm sees v's incident edges and the corresponding success probabilities. Then, the algorithm makes an irrevocable matching decision about v. Should the algorithm choose to match v to an offline vertex u, the match would succeed with probability p_{uv}. If a match is unsuccessful, the offline vertex u can still be matched to future online vertices, but the online vertex v will not get a second chance. An offline vertex $u \in U$ is *successful* if there is an edge that is matched to u and is successful; u is *unsuccessful* otherwise. We want to maximize the number of successful offline vertices.

We remark that the results in this paper generalize to the vertex-weighted problem, where offline vertices have positive weights and we want to maximize the total weight of successful offline vertices. Readers familiar with the online matching literature shall not find this surprising since almost all known results for unweighted online matching problems generalize to the vertex-weighted problems. This version of the paper only presents the unweighted case to keep the exposition simple.

By allowing $p_{uv} = 0$, we may assume without loss of generality that $E = U \times V$ which will simplify the exposition in some parts of this paper. We say that u and v are neighbors if $p_{uv} > 0$. We say that the success probabilities are equal if p_{uv} is either 0 and p for some $0 \le p \le 1$.

Stochastic Budgets. Mehta and Panigrahi [25] showed that when the success probabilities are infinitesimal, we can consider an equivalent model in which the randomness plays a different role. At the beginning, each offline vertex $u \in U$ independently draws a budget θ_u from the exponential distribution with mean 1. When an online vertex $v \in V$ arrives, the algorithm may match it to any offline neighbor and collect a gain that equals the success probability of the edge. However, the total gain from an offline vertex u is capped by its budget θ_u, i.e., it is either the sum of success probabilities of the edges matched to u, denoted as ℓ_u and referred to as its *load*, or its budget θ_u, whichever is smaller. Further, the budget θ_u is unknown to the algorithm until the moment when the load ℓ_u exceeds the budget, which corresponds to u's succeeding.

Benchmarks and Competitive Analysis. The *offline (non-stochastic) optimum*, denoted as OPT, refers to the optimal objective achievable by a computationally unlimited offline algorithm that knows the graph and success probabilities, and when the budgets are non-stochastically equal to 1. The offline optimum upper bounds the objective achievable by any (online or offline) algorithm in the model with stochastic budgets [25].

The *offline stochastic optimum*, denoted as S-OPT, refers to the optimal objective achievable by a computationally unlimited offline algorithm that knows the graph and success probabilities, but can only match the online vertices one by one by their arrival order, and can only observe the stochastic budget of an offline vertex when its load exceeds the budget like online algorithms.

An online algorithm is Γ-competitive w.r.t. one of these benchmarks if for *any instance* of Online Matching with Stochastic Rewards, the expected objective given by the algorithm is at least a Γ fraction of the benchmark. Here $0 \le \Gamma \le 1$ is called the *competitive ratio*.

2.2 Existing Linear Programs

Consider $0 \le x_{uv} \le 1$ as the probability that u is matched to v (successful or not) for any $u \in U$ and any $v \in V$. The *Matching LP* is defined as:

$$\begin{aligned} \max_{x \ge 0} \quad & \sum_{u \in U} \sum_{v \in V} p_{uv} x_{uv} & \\ \text{s.t.} \quad & \sum_{v \in V} p_{uv} x_{uv} \le 1 & \forall u \in U \\ & \sum_{u \in U} x_{uv} \le 1 & \forall v \in V \end{aligned}$$

The first constraint states that the expected load of an offline vertex u, which is equal to the probability that u succeeds, is at most 1. The second constraint says that each online vertex v can be matched to at most one offline vertex. An offline allocation yields a feasible solution to the Matching LP by the aforementioned interpretation of x_{uv}. Hence, the optimal objective of Matching LP upper bounds the offline optimum OPT.

For any offline vertex $u \in U$ and any subset of online vertices $S \subseteq U$, let $\bar{p}_{uS} = \min\{p_{uS}, 1\}$, where recall that $p_{uS} = \sum_{v \in S} p_{uv}$. Consider $0 \le x_{uS} \le 1$ as the probability that S is the subset of online vertices matched to u by the algorithm. Huang and Zhang [17] used the following *Configuration LP* and its dual LP:

$$\begin{aligned} \max_{x \ge 0} \quad & \sum_{u \in U} \sum_{S \subset V} \bar{p}_{uS} x_{uS} & \textbf{(Primal)} \\ \text{s.t.} \quad & \sum_{S \subseteq V} x_{uS} \le 1 & \forall u \in U \\ & \sum_{u \in U} \sum_{S \subseteq V : v \in S} x_{uS} \le 1 & \forall v \in V \end{aligned}$$

$$\begin{aligned} \min_{\alpha, \beta \ge 0} \quad & \sum_{u \in U} \alpha_u + \sum_{v \in V} \beta_v & \textbf{(Dual)} \\ \text{s.t.} \quad & \alpha_u + \sum_{v \in S} \beta_v \ge \bar{p}_{uS} & \forall u \in U, \forall S \subseteq V \end{aligned} \tag{3}$$

The optimal objective of Configuration LP also upper bounds the offline optimum OPT. In fact, it coincides with the optimal objective of Matching LP when the success probabilities are infinitesimal (c.f., Huang and Zhang [17]). Nevertheless, its dual structure is better suited for an online primal dual analysis.

Lemma 1 (c.f., Lemma 3 of Huang and Zhang [17]). *Suppose that an algorithm for Online Matching with Stochastic Rewards can further split p_{uv} between α_u and β_v each time it matches an edge (u, v), such that the dual constraint (3) holds up to factor Γ* in expectation*:*

$$\mathbf{E}\Big[\alpha_u + \sum_{v \in S} \beta_v \Big] \ \ge \ \Gamma \cdot \bar{p}_{uS} \ .$$

Then, the algorithm is Γ-competitive w.r.t. OPT.

Finally, Goyal and Udwani [13] introduced the *Path-based LP* which is tailored for the offline stochastic optimum S-OPT. A *sample path* refers to all information of what has happened so far, including the subset of edges matched by the algorithm, and the realization of whether these matches are successful. This LP considers variable $0 \le x_{uv}^{\omega} \le 1$ for any edge (u, v) which represents the probability that the algorithm matches v to u conditioned on a sample path ω. We omit the details of Path-based LP because it is not directly related to the LPs in this paper, and its constraints are too complicated to be covered concisely here. We remark that the number of possible sample paths and thus the number of variables of Path-based LP are exponential in the number of *edges*. By contrast, the number of variables of Configuration LP is exponential in the number of *vertices*, which is typically much smaller than the number of edges.

2.3 Existing Algorithms

Ranking. We will consider the *Ranking* algorithm in the case of equal success probabilities. At the beginning, draw a rank $\rho_u \in [0, 1]$ uniformly at random for each offline vertex $u \in U$. Then, on the arrival of each online vertex $v \in V$, match it to the unsuccessful neighbor that has the smallest rank. We can break ties arbitrarily since they occur with zero probability.

Stochastic Balance (Equal Probability). On the arrival of each online vertex $v \in V$, match it to the unsuccessful neighbor that has the smallest load, breaking ties arbitrarily, e.g., by the lexicographical order. Recall that the load of an offline vertex is the sum of success probabilities of the edges that have been matched to it so far.

Stochastic Balance (General). We first describe a fractional algorithm. On the arrival of each online vertex $v \in V$, consider a continuous process that matches v fractionally to the unsuccessful neighbor $u \in U$ with the largest $p_{uv}\big(1 - g(\ell_u)\big)$ where $g : [0, 1] \to [0, 1]$ is a non-decreasing function chosen by the algorithm designer. If the success probabilities are infinitesimal, one can further convert it to an integral algorithm with no loss in the competitive ratio, through a reduction by Huang and Zhang [17] based on randomized rounding. Hence, it suffices to analyze the fractional algorithm's competitive ratio.

3 Stochastic Configuration Linear Programs

3.1 Stochastic Thresholds

Even when the success probabilities are arbitrary, we may still consider an equivalent model that generalizes the viewpoint of stochastic budgets for infinitesimal success probabilities. We call this the *stochastic thresholds viewpoint.* At the beginning, each offline vertex u independently draws a threshold τ_u uniformly from $[0, 1]$. When an online vertex v arrives, the algorithm may match it to any offline neighbor and collect a gain which equals the success probability of the edge. For each offline vertex u and any subset of online vertices S, define:

$$\tilde{p}_{uS} = 1 - \prod_{v \in S} \big(1 - p_{uv}\big) \ .$$

An offline vertex u is successful (and can no longer be matched) if the subset of online vertices S matched to it satisfies $\tilde{p}_{uS} \geq \tau_u$. Further, the threshold τ_u is unknown to the algorithm until the moment that offline vertex u becomes successful.

Observe that when the success probabilities are infinitesimal, we have $\tilde{p}_{uS} = 1 - e^{-p_{uS}}$, and the stochastic budget and threshold of an offline vertex $u \in U$ satisfy $\theta_u = -\ln(1 - \tau_u)$.

3.2 Reduced-form Stochastic Configuration Linear Program

Consider $0 \leq y_{uS} \leq 1$ as the probability that a subset of online vertices S would be matched to an offline vertex u if it has an infinite stochastic budget $\theta_u = \infty$ (i.e., when $\tau_u = 1$ if the success probabilities are not infinitesimal), over the randomness of the stochastic budgets θ_{-u} of other offline vertices and the randomness of the algorithm. Further, for any $S \subseteq V$ and $v \in V$, let:

$$S(v) = \big\{ v' \in S : v' \prec v \big\}$$

denote the subset of online vertices in S that arrive before v. We consider the following *Reduced-form Stochastic Configuration LP* and its dual:

$$\begin{aligned}
\max_{y \geq 0} \quad & \sum_{u \in U} \sum_{S \subseteq V} \tilde{p}_{uS}\, y_{uS} && \textbf{(Primal)} \\
\text{s.t.} \quad & \sum_{S \subseteq V} y_{uS} \leq 1 && \forall u \in U \qquad (4) \\
& \sum_{u \in U} \sum_{S \subseteq V : v \in S} \big(1 - \tilde{p}_{uS(v)}\big) y_{uS} \ \leq 1 && \forall v \in V \qquad (5)
\end{aligned}$$

$$\begin{aligned}
\min_{\alpha, \beta \geq 0} \quad & \sum_{u \in U} \alpha_u + \sum_{v} \beta_v && \textbf{(Dual)} \\
\text{s.t.} \quad & \alpha_u + \sum_{v \in S} \big(1 - \tilde{p}_{uS(v)}\big) \beta_v \geq \tilde{p}_{uS} && \forall u \in U, \forall S \subseteq V \qquad (6)
\end{aligned}$$

Theorem 1. *The optimal objective of Reduced-form Stochastic Configuration LP is greater than or equal to* S-OPT.

Lemma 2. *Suppose that an algorithm for Online Matching with Stochastic Rewards can further split p_{uv} between α_u and β_v each time it matches an edge (u, v), such that the dual constraint (6) holds up to factor Γ* in expectation, *i.e.:*

$$\mathbf{E}\Big[\alpha_u + \sum_{v \in S} \big(1 - \tilde{p}_{uS(v)}\big) \beta_v \Big] \ \geq \ \Gamma \cdot \tilde{p}_{uS} \ .$$

Then, the algorithm is Γ-competitive w.r.t. S-OPT.

3.3 Stochastic Configuration Linear Program

Consider $0 \le y_{uS}(\theta_{-u}) \le 1$ (respectively, $y_{uS}(\tau_{-u})$ if the success probabilities are *not* infinitesimal) as the probability that a subset of online vertices S would be matched to an offline vertex u if it has an infinite stochastic budget $\theta_u = \infty$ (respectively, if $\tau_u = 1$), *conditioned on the stochastic budgets* θ_{-u} (respectively, the stochastic thresholds τ_{-u}) of other offline vertices, and over the randomness of the algorithm. This paper will further consider an even more expressive Stochastic Configuration LP and its dual in some analyses.

$$
\begin{aligned}
\underset{y \ge 0}{\text{maximize}} \quad & \sum_{u \in U} \sum_{S \subseteq V} \mathbf{E}_{\theta_{-u}} \Big[\tilde{p}_{uS} \, y_{uS}(\theta_{-u}) \Big] && \textbf{(Primal)} \\
\text{subject to} \quad & \sum_{S \subseteq V} y_{uS}(\theta_{-u}) \le 1 && \forall u \in U, \forall \theta_{-u} \qquad (7) \\
& \sum_{u} \sum_{S \subseteq V : v \in S, \theta_u \ge p_{uS(v)}} y_{uS}(\theta_{-u}) \le 1 && \forall v \in V, \forall \theta \qquad (8)
\end{aligned}
$$

$$
\begin{aligned}
\underset{\alpha, \beta \ge 0}{\text{minimize}} \quad & \mathbf{E}_{\theta} \Big[\sum_{u} \alpha_u(\theta_{-u}) + \sum_{v \in V} \beta_v(\theta) \Big] && \textbf{(Dual)} \\
\text{subject to} \quad & \alpha_u(\theta_{-u}) + \mathbf{E}_{\theta_u} \Big[\sum_{v \in S : \theta_u \ge p_{uS(v)}} \beta_v(\theta_u, \theta_{-u}) \Big] \ge \tilde{p}_{uS} && \forall u \in U, \forall S \subseteq V, \forall \theta_{-u} \qquad (9)
\end{aligned}
$$

Theorem 2. *The optimal objective of Stochastic Configuration LP is at least* S-OPT.

Lemma 3. *Suppose that an algorithm for Online Matching with Stochastic Rewards can further split p_{uv} between α_u and β_v each time it matches an edge (u, v). Let $\alpha_u(\theta_{-u})$ be the expectation of α_u conditioned on any stochastic budgets θ_{-u} of offline vertices other than u, and let $\beta(\theta)$ be the expectation of β_v conditioned on any stochastic budgets θ of all vertices. If the dual constraint (6) holds up to factor Γ, i.e.:*

$$
\alpha_u(\theta_{-u}) + \mathbf{E}_{\theta_u} \Big[\sum_{v \in S : \theta_u \ge p_{uS(v)}} \beta_v(\theta_u, \theta_{-u}) \Big] \;\ge\; \Gamma \cdot \tilde{p}_{uS} \;,
$$

then the algorithm is Γ-competitive w.r.t. S-OPT.

4 Ranking

4.1 Basics

We will assume throughout the section that the ranks of offline vertices are distinct, since that happens with probability 1. Below we first develop some basic elements that will be used by the proofs of both results in this section.

Dual Updates. We start by explaining the dual update rule associated with Ranking, which is identical to the existing one in the literature. The dual variables are initially 0. Let $g : [0,1] \to [0,1]$ be a non-decreasing gain splitting function to be determined. We will assume that $g(1) = 1$ which will allow us to handle a boundary case of the analysis under a unified framework. Recall that ρ_u denotes the rank of u which is uniformly drawn between 0 and 1 at the beginning by the Ranking algorithm. When the algorithm matches an online vertex v to an offline vertex u, increase α_u by $p \cdot g(\rho_u)$, and set β_v as $p \cdot \big(1 - g(\rho_u)\big)$, where recall that p denotes the equal success probability of matching neighboring vertices.

Characterization of Matching. All analyses in this section will fix an offline vertex $u \in U$ and the ranks ρ_{-u} and stochastic thresholds τ_{-u} of other offline vertices. A canonical analysis of Ranking, such as those for Online Bipartite Matching (c.f., Devanur et al. [7]), would further fix a neighboring online vertex $v \in V$ and characterize the matching outcome of u and v. For the problem of Online Matching with Stochastic Rewards and in the spirit of Configuration LPs, however, we need to characterize the joint matching outcome of u and *all its neighbors*.

We first define some notations. Consider an imaginary run of Ranking with vertex u removed, while keeping the ranks of other offline vertices as ρ_{-u}. Further consider an online vertex v in u's neighborhood. If v is matched to an offline vertex u' in the imaginary run, define the v's *critical rank* as $\mu_v = \rho_{u'}$. If v is not matched, define its critical rank as $\mu_v = 1$. Finally, for any $0 \le \mu \le 1$, let $N_u(\mu)$ be the set of u's neighbors whose critical ranks are greater than or equal to μ.

Lemma 4. *Suppose that u's rank is ρ_u and u's stochastic threshold is such that u succeeds after i matches. Then, the subset of online vertices matched to u is the first $\min\big\{i, |N_u(\rho_u)|\big\}$ vertices in $N_u(\rho_u)$ according to the arrival order.*

Expectation of Dual Variables. The next invariant about dual variable α_u follows from the above gain splitting rule.

Lemma 5. *For any offline vertex $u \in U$, any ranks ρ and stochastic thresholds τ of offline vertices, and the corresponding load ℓ_u of u, we have:*

$$\alpha_u \;=\; \ell_u \cdot g(\rho_u) \ .$$

As a corollary, we can further bound the expectation of α_u conditioned on the ranks ρ_{-u} and stochastic thresholds τ_{-u} of other offline vertices.

Lemma 6. *For any offline vertex $u \in U$, any ranks ρ_{-u} and stochastic thresholds τ_{-u} of offline vertices other than u, we have:*

$$\mathbf{E}_{\rho_u,\tau_u}\big[\,\alpha_u \mid \rho_{-u}, \tau_{-u}\,\big] = \int_0^1 \Big(1 - (1-p)^{|N_u(\rho_u)|}\Big) g(\rho_u)\,\mathrm{d}\rho_u \ .$$

Recall that $N_u(\mu)$ is the subset of u's neighbors whose critical ranks are greater than or equal to μ. Further let $N_u(\mu, v)$ be the subset of vertices in $N_u(\mu)$ that arrive before v.

Lemma 7. *For any offline vertex $u \in U$, any u's neighbor v, any ranks ρ_{-u} and stochastic thresholds τ_{-u} of offline vertices other than u, we have:*

$$\mathbf{E}_{\rho_u,\tau_u}\big[\,\beta_v \mid \rho_{-u},\tau_{-u}\,\big] \;\geq\; p\bigg(1-g(\mu_v)+\int_0^{\mu_v}\big(1-p\big)^{|N_u(\rho_u,v)|}\big(g(\mu_v)-g(\rho_u)\big)\,\mathrm{d}\rho_u\bigg)\;.$$

4.2 Non-stochastic Benchmark

Theorem 3. *Ranking is at least* 0.572*-competitive w.r.t. the offline (non-stochastic) benchmark* OPT *for any instance with equal success probabilities.*

The rest of the subsection is devoted to proving Theorem 3 by an online primal dual analysis with the Configuration LP. By Lemma 1, it suffices to show that $\mathbf{E}\big[\,\alpha_u+\sum_{v\in S}\beta_v\big]\geq\Gamma\cdot\bar{p}_{uS}$ for any offline vertex $u\in U$ and any subset of online vertices $S\subseteq V$, and with the stated competitive ratio $\Gamma=0.572$. We will prove it further conditioning on any ranks ρ_{-u} and stochastic thresholds τ_{-u} of other vertices, i.e.:

$$\mathbf{E}_{\rho_u,\tau_u}\bigg[\,\alpha_u+\sum_{v\in S}\beta_v \mid \rho_{-u},\tau_{-u}\,\bigg] \;\geq\; \Gamma\cdot\bar{p}_{uS}\;, \tag{10}$$

First we apply the bounds for dual variables' expectation from Lemmas 6 and 7 to the LHS of the inequality:

$$\mathbf{E}_{\rho_u,\tau_u}\bigg[\,\alpha_u+\sum_{v\in S}\beta_v \mid \rho_{-u},\tau_{-u}\,\bigg] \geq \int_0^1\big(1-\underbrace{(1-p)^{|N_u(\rho_u)|}}_{(a)}\big)g(\rho_u)\,\mathrm{d}\rho_u \tag{11}$$

$$+\underbrace{\sum_{v\in S}}_{(b)} p\bigg(1-g(\mu_v)+\int_0^{\mu_v}\underbrace{\big(1-p\big)^{|N_u(\rho_u,v)|}}_{(c)}\big(g(\mu_v)-g(\rho_u)\big)\,\mathrm{d}\rho_u\bigg)\;.$$

Assumptions. The rest of the analysis will treat the above as a pure inequality problem, where the arrival order and critical ranks of u's neighbors may be chosen arbitrarily by an adversary. In particular, when we characterize the worst-case scenario of the inequality we shall not concern about how to construct an instance to get the specified critical ranks. With this treatment, we can make several assumptions that on the one hand are without loss of generality, and on the other hand simplify the subsequent analysis.

First recall the assumption that the offline vertices' ranks are distinct, because the exceptions happen with zero probability. In other words, we may relax and simplify the RHS assuming distinct ranks, which for example implies that u's rank ρ_u does not equal the critical rank μ_v of any v. Although the resulting bound may be violated for some zero measure subset of ranks ρ_u, the bound still holds after the integration.

Further, it suffices to consider the above inequality for infinitesimal success probabilities. For $p=1$ it is the online bipartite matching problem by Karp et al. [22], for which Ranking is $1-\frac{1}{e}$ competitive. For any instance with $0<p<1$ and any large integer n, we can instead consider a new instance such that (i) the equal success probabilities is $p'=1-\sqrt[n]{1-p}$, (ii) each online vertex in the original instance has n copies in the new instance, which arrive consecutively and have the same critical ranks as the original vertex, and (iii) the subset S' of concerned in the new instance contains $m=\lfloor\frac{p}{p'}\rfloor$ copies of each vertex $v\in S$ in the original one. By this construction, the RHS of Eqn. (10) is the same in the two instances, up to an error from the rounding of m that diminishes as n tends to infinity. Further, the RHS of Eqn. (11) is smaller in new instance because (i) and (ii) ensure that part (a) stays the same and part (c) decreases, and (iii) ensures that the changes to part (b) weakly decreases the expression. Hence, it suffices to establish the inequality in the new instance which satisfy the assumption of infinitesimal success probability. We may therefore rewrite the expression as:

$$\mathbf{E}_{\rho_u,\tau_u}\left[\alpha_u+\sum_{v\in S}\beta_v \mid \rho_{-u},\tau_{-u}\right]\geq\int_0^1\left(1-e^{-p|N_u(\rho_u)|}\right)g(\rho_u)\,\mathrm{d}\rho_u \qquad (12)$$
$$+\sum_{v\in S}p\big(1-g(\mu_v)\big)+\sum_{v\in S}\int_0^{\mu_v}e^{-p|N_u(\rho_u,v)|}\left(1-e^{-p}\right)\big(g(\mu_v)-g(\rho_u)\big)\,\mathrm{d}\rho_u\ .$$

Finally, it is sufficient to establish the stated approximate dual feasibility condition when the critical ranks μ_v are distinct for different v. Otherwise, we may slightly perturb the critical ranks, e.g., by a random noise drawn from $[-\epsilon,\epsilon]$ for a sufficiently small $\epsilon>0$, to get a new instance satisfying our assumption. The resulting RHS of Eqn. (12) converges to the original one when the magnitude of perturbations tends to zero.

For any $0\leq\mu\leq 1$, let:

$$P(\mu)\;=\;p\cdot\big|N_u(\mu)\cap S\big| \quad , \qquad \bar{P}(\mu)\;=\;p\cdot\big|N_u(\mu)\setminus S\big|$$

denote the sums of success probabilities of u's neighbors with critical ranks greater than or equal to μ, inside S and outside S respectively.

Characterization of Worst-case. We next present a series of lemmas that characterize the worst-case of various aspects of the instance, including the arrival order, the size of S, and critical ranks, etc.

Lemma 8 (Worst Arrival Order). *We have:*

$$\mathbf{E}_{\rho_u,\tau_u}\left[\alpha_u+\sum_{v\in S}\beta_v \mid \rho_{-u},\tau_{-u}\right]\;\geq\;\int_0^1\left(1-e^{-\bar{P}(\rho_u)}\right)g(\rho_u)\,\mathrm{d}\rho_u$$
$$-\int_0^1\big(1-g(\mu_v)\big)\,\mathrm{d}P(\mu_v)-\int_0^1 e^{-P(\rho_u)-\bar{P}(\rho_u)}\int_{\rho_u}^1 g(\mu_v)\,\mathrm{d}e^{P(\mu_v)}\,\mathrm{d}\rho_u\ . \qquad (13)$$

The lemma's inequality would hold with equality if online vertices in S arrive by increasing order of critical ranks, and vertices outside S arrive before those in

S. We remark that the negative signs in front of the second and third integration come from the fact that $P(\mu)$ is non-increasing. All three parts make positive contribution to the RHS.

Lemma 9 (Worst Critical Ranks Outside S). *Given any non-increasing function $P : [0,1] \to [0,\infty)$, the RHS of Eqn. (13) is minimized when $\bar{P}(\mu)$ is a step function that equals ∞ for $\mu < \mu_0$ and equals 0 for $\mu > \mu_0$, for some threshold μ_0.*

The lemma indicates that letting all online vertices outside S have the same critical rank μ_0 is the worst-case scenario.

Given Lemma 9, and noting that $p_{uS} = P(0)$ and thus $\bar{p}_{uS} = \min\big\{P(0), 1\big\}$, it suffices to find a gain splitting function g such that for any non-increasing function $P : [0,1] \to [0,\infty)$ we have:

$$\int_0^{\mu_0} g(\rho_u)\,\mathrm{d}\rho_u - \int_0^1 \big(1 - g(\mu_v)\big)\,\mathrm{d}P(\mu_v) - \int_{\mu_0}^1 e^{-P(\rho_u)} \int_{\rho_u}^1 g(\mu_v)\,\mathrm{d}e^{P(\mu_v)}\,\mathrm{d}\rho_u \geq \Gamma \cdot \min\big\{P(0), 1\big\} . \tag{14}$$

Lemma 10 (Worst Size of S). *If Eqn. (14) holds for all non-increasing function P that satisfies $p_{uS} = P(0) = 1$, then it also holds for an arbitrary non-increasing function P.*

Therefore, our task further simplifies to finding a gain splitting function g such that for any non-increasing function $P : [0,1] \to [0,1]$ with $P(0) = 1$ we have:

$$\int_0^{\mu_0} g(\rho_u)\,\mathrm{d}\rho_u - \int_0^1 \big(1 - g(\mu_v)\big)\,\mathrm{d}P(\mu_v) - \int_{\mu_0}^1 e^{-P(\rho_u)} \int_{\rho_u}^1 g(\mu_v)\,\mathrm{d}e^{P(\mu_v)}\,\mathrm{d}\rho_u \;\geq\; \Gamma . \tag{15}$$

This is the moment when we finally specify the gain splitting function g:

$$g(\rho) = \begin{cases} \min\big\{\frac{c}{e-(e-1)\rho}, 1 - \frac{1}{e}\big\} & 0 \leq \rho < 1; \\ 1 & \rho = 1. \end{cases} \tag{16}$$

for a constant $c \approx 1.161$ such that $\int_0^1 g(x) = 1 - g(0) > 0.572$.

Lemma 11 (Worst Critical Ranks Inside S). *For the function g in Eqn. (15) is minimized when P approaches a step function in the limit, i.e., when for some $\rho_0 > \mu_0$:*

$$P(\rho) = \begin{cases} 1 & 0 \leq \rho < \rho_0 - \epsilon; \\ \frac{\rho_0 - \rho}{\epsilon} & \rho_0 - \epsilon \leq \rho \leq \rho_0; \\ 0 & \rho_0 < \rho \leq 1. \end{cases}$$

and let $\epsilon \to 0$.

Despite its complex look, the lemma actually gives a simple characterization of the worst-case critical ranks inside S: all online vertices in S have the same critical rank $\rho_0 > \mu_0$. We have the complex form because we write the lower bounds of the LHS of approximate dual feasibility through a function P, which help simplify the proofs of previous lemmas but is inconvenient when we need to represent identical critical ranks.

Finally, we apply the above worst-case function P and focus on optimizing the gain splitting function g w.r.t. the resulting differential inequality.

Lemma 12. *For any $\mu_0 < \rho_0$, and for the function g in Eqn. (16), we have:*

$$\int_0^{\mu_0} g(\rho_u)\,\mathrm{d}\rho_u + \Big(1 - g(\rho_0)\Big) + \Big(1 - \frac{1}{e}\Big)(\rho_0 - \mu_0) g(\rho_0) \;\geq\; \Gamma \;,$$

for the stated competitive ratio $\Gamma = 0.572$

4.3 Stochastic Benchmark

Theorem 4. *Ranking is $1-\frac{1}{e}$-competitive w.r.t. the offline stochastic benchmark* S-OPT *for any instance with equal success probabilities.*

The rest of the subsection is devoted to proving Theorem 4 by an online primal dual analysis with the Reduced-form Stochastic Configuration LP. By Lemma 2, it is sufficient to show that $\mathbf{E}\big[\alpha_u + \sum_{v \in S}(1 - \tilde{p}_{uS(v)})\beta_v\big] \geq \Gamma\, \tilde{p}_{uS}$ for any offline vertex $u \in U$ and any subset of online vertices $S \subseteq V$, and with the stated competitive ratio $\Gamma = 1 - \frac{1}{e}$. We will prove it further conditioning on any ranks ρ_{-u} and stochastic thresholds τ_{-u} of other vertices, i.e.:

$$\mathbf{E}_{\rho_u, \tau_u}\left[\alpha_u + \sum_{v \in S}(1 - \tilde{p}_{uS(v)})\beta_v \mid \rho_{-u}, \tau_{-u}\right] \geq \Big(1 - \frac{1}{e}\Big) \cdot \tilde{p}_{uS} \tag{17}$$

Let $v_1, v_2, \ldots, v_n$ be the online vertices in S, sorted by their critical ranks, which we denote as $\mu_1 \geq \mu_2 \geq \cdots \geq \mu_n$. We next apply the lower bounds for the dual variables' expectation to the LHS of Eqn. (17). First we have:

$$\begin{aligned}
\mathbf{E}_{\rho_u, \tau_u}\big[\,\alpha_u \mid \rho_{-u}, \tau_{-u}\,\big] &\geq \int_0^1 \Big(1 - (1-p)^{|N_u(\rho_u)|}\Big) g(\rho_u)\,\mathrm{d}\rho_u \qquad \text{(Lemma 6)} \\
&\geq \int_0^1 \Big(1 - (1-p)^{|N_u(\rho_u) \cap S|}\Big) g(\rho_u)\,\mathrm{d}\rho_u \\
&= \int_0^1 \sum_{i=1}^{|N_u(\rho_u) \cap S|} p(1-p)^{i-1} g(\rho_u)\,\mathrm{d}\rho_u \\
&= \sum_{i=1}^{n} p(1-p)^{i-1} \int_0^{\mu_i} g(\rho_u)\,\mathrm{d}\rho_u \;.
\end{aligned}$$

To bound the expected contribution by β_v's, we apply a weaker version of Lemma 7, dropping the second part on the RHS of the lemma's inequality.

Suppose that the vertices in S arrive by order $v_{\pi(1)} \prec v_{\pi(1)} \prec \ldots \prec v_{pi(n)}$. We get that:

$$\begin{aligned}
&\mathbf{E}_{\rho_u,\tau_u}\left[\sum_{v\in S}(1-\tilde{p}_{uS(v)})\beta_v \mid \rho_{-u},\tau_{-u}\right] \\
&\geq \sum_{i=1}^{n} p(1-p)^{i-1}\big(1-g(\mu_{\pi(i)})\big) && \text{(Lemma 7)} \\
&\geq \sum_{i=1}^{n} p(1-p)^{i-1}\big(1-g(\mu_i)\big)\ . && \text{(rearrangement inequality)}
\end{aligned}$$

Combining the two bounds, we get that the LHS of Eqn. (17) is at least:

$$\sum_{i=1}^{n} p(1-p)^{i-1}\Big(\underbrace{\int_0^{\mu_i} g(\rho_u)\,\mathrm{d}\rho_u + 1 - g(\mu_i)}_{(\star)}\Big)$$

Next we choose $g(x) = e^{x-1}$ just like the analysis of Ranking for the original Online Bipartite Matching problem. This choice ensures that $(\star)$ equals $1-\frac{1}{e}$. The above bound is therefore:

$$\left(1-\frac{1}{e}\right)\sum_{i=1}^{n} p(1-p)^{i-1} = \left(1-\frac{1}{e}\right)\Big(1-(1-p)^n\Big) = \left(1-\frac{1}{e}\right)\cdot \tilde{p}_{uS}\ .$$

References

1. Aggarwal, G., Goel, G., Karande, C., Mehta, A.: Online vertex-weighted bipartite matching and single-bid budgeted allocations. In: SODA. pp. 1253–1264. SIAM (2011)
2. Birnbaum, B., Mathieu, C.: On-line bipartite matching made simple. ACM SIGACT News **39**(1), 80–87 (2008)
3. Blanc, G., Charikar, M.: Multiway online correlated selection. In: FOCS. pp. 1277–1284. IEEE (2022)
4. Buchbinder, N., Jain, K., Naor, J.S.: Online primal-dual algorithms for maximizing ad-auctions revenue. In: ESA. pp. 253–264. Springer (2007)
5. Devanur, N.R., Huang, Z., Korula, N., Mirrokni, V.S., Yan, Q.: Whole-page optimization and submodular welfare maximization with online bidders. ACM Transactions on Economics and Computation **4**(3), 1–20 (2016)
6. Devanur, N.R., Jain, K.: Online matching with concave returns. In: STOC. pp. 137–144 (2012)
7. Devanur, N.R., Jain, K., Kleinberg, R.D.: Randomized primal-dual analysis of ranking for online bipartite matching. In: SODA. pp. 101–107. SIAM (2013)
8. Fahrbach, M., Huang, Z., Tao, R., Zadimoghaddam, M.: Edge-weighted online bipartite matching. Journal of the ACM (forthcoming)
9. Feldman, J., Korula, N., Mirrokni, V., Muthukrishnan, S., Pál, M.: Online ad assignment with free disposal. In: WINE. pp. 374–385. Springer (2009)

10. Feldman, J., Mehta, A., Mirrokni, V., Muthukrishnan, S.: Online stochastic matching: beating $1 - \frac{1}{e}$. In: FOCS. pp. 117–126. IEEE (2009)
11. Gao, R., He, Z., Huang, Z., Nie, Z., Yuan, B., Zhong, Y.: Improved online correlated selection. In: FOCS. pp. 1265–1276. IEEE (2021)
12. Goel, G., Mehta, A.: Online budgeted matching in random input models with applications to adwords. In: SODA. pp. 982–991 (2008)
13. Goyal, V., Udwani, R.: Online matching with stochastic rewards: Optimal competitive ratio via path-based formulation. Operations Research (2022)
14. Huang, Z., Shu, X.: Online stochastic matching, Poisson arrivals, and the natural linear program. In: STOC. pp. 682–693 (2021)
15. Huang, Z., Shu, X., Yan, S.: The power of multiple choices in online stochastic matching. In: STOC. pp. 91–103 (2022)
16. Huang, Z., Tang, Z.G., Wu, X., Zhang, Y.: Online vertex-weighted bipartite matching: Beating $1 - \frac{1}{e}$ with random arrivals. ACM Transactions on Algorithms **15**(3), 1–15 (2019)
17. Huang, Z., Zhang, Q.: Online primal dual meets online matching with stochastic rewards: configuration LP to the rescue. In: STOC. pp. 1153–1164. ACM (2020)
18. Huang, Z., Zhang, Q., Zhang, Y.: AdWords in a panorama. In: FOCS. pp. 1416–1426. IEEE (2020)
19. Jin, B., Williamson, D.P.: Improved analysis of ranking for online vertex-weighted bipartite matching in the random order model. In: WINE. pp. 207–225. Springer (2021)
20. Kalyanasundaram, B., Pruhs, K.R.: An optimal deterministic algorithm for online b-matching. Theoretical Computer Science **233**(1-2), 319–325 (2000)
21. Karande, C., Mehta, A., Tripathi, P.: Online bipartite matching with unknown distributions. In: STOC. pp. 587–596 (2011)
22. Karp, R.M., Vazirani, U.V., Vazirani, V.V.: An optimal algorithm for on-line bipartite matching. In: STOC. pp. 352–358. ACM (1990)
23. Liang, J., Tang, Z.G., Xu, Y., Zhang, Y., Zhou, R.: On the perturbation function of ranking and balance for weighted online bipartite matching. ArXiv **abs/2210.10370** (2022)
24. Mahdian, M., Yan, Q.: Online bipartite matching with random arrivals: an approach based on strongly factor-revealing LPs. In: STOC. pp. 597–606 (2011)
25. Mehta, A., Panigrahi, D.: Online matching with stochastic rewards. In: FOCS. pp. 728–737. IEEE (2012)
26. Mehta, A., Saberi, A., Vazirani, U., Vazirani, V.: Adwords and generalized online matching. Journal of the ACM **54**(5), 22–es (2007)
27. Shin, Y., An, H.C.: Making three out of two: three-way online correlated selection. In: ISAAC. Schloss Dagstuhl-Leibniz-Zentrum für Informatik (2021)
28. Tang, Z.G., Wu, J., Wu, H.: (fractional) online stochastic matching via fine-grained offline statistics. In: STOC. pp. 77–90 (2022)
29. Udwani, R.: Adwords with unknown budgets and beyond. ArXiv **abs/2110.00504** (2021)
30. Vazirani, V.V.: Towards a practical, budget-oblivious algorithm for the adwords problem under small bids. ArXiv **abs/2107.10777** (2022)