

Attenuate Locally, Win Globally: An Attenuation-based Framework for Online Stochastic Matching with Timeouts *

Brian Brubach[†], Karthik A. Sankararaman[‡], Aravind Srinivasan[§] and Pan Xu[¶]

Department of Computer Science, University of Maryland, College Park, MD 20742, USA

April 21, 2018

Abstract

Online matching problems have garnered significant attention in recent years due to numerous applications in e-commerce, online advertisements, ride-sharing, etc. Many of them capture the uncertainty in the real world by including stochasticity in both the arrival process and the matching process. The Online Stochastic Matching with Timeouts problem introduced by Bansal, *et al.*, (*Algorithmica*, 2012) models matching markets (e.g., E-Bay, Amazon). Buyers arrive from an independent and identically distributed (i.i.d.) known distribution on buyer profiles and can be shown a list of items one at a time. Each buyer has some probability of purchasing each item and a limit (timeout) on the number of items they can be shown.

Bansal *et al.*, (*Algorithmica*, 2012) gave a 0.12-competitive algorithm which was improved by Adamczyk, *et al.*, (*ESA*, 2015) to 0.24. We present an online attenuation framework that uses an algorithm for offline stochastic matching as a black box. On the upper bound side, we show that this framework, combined with a black-box adapted from Bansal *et al.*, (*Algorithmica*, 2012), yields an online algorithm which nearly doubles the ratio to 0.46. On the lower bound side, we show that no algorithm can achieve a ratio better than 0.632 using the standard LP for this problem. This framework has a high potential for further improvements since new algorithms for offline stochastic matching can directly improve the ratio for the online problem.

Our online framework also has the potential for a variety of extensions. For example, we introduce a natural generalization: Online Stochastic Matching with Two-sided Timeouts in which both online and offline vertices have timeouts. Our framework provides the first algorithm for this problem achieving a ratio of 0.30. We once again use the algorithm of Adamczyk *et al.*, (*ESA*, 2015) as a black-box and plug-it into our framework.

1 Introduction

Consider a typical problem in matching markets (e.g. E-Bay, Amazon). We have a certain number of buyer profiles and items. Let us denote the set of buyer profiles by V and the set of items by U . In our initial problem, we assume that an item is present in the market until bought. We have a n -round process. In each round, a buyer, sampled uniformly at random (with replacement) from the buyer profiles, arrives. We assume that we have n buyer profiles (this assumption is called *integral arrival rates* in the literature [27]). For every item $u \in U$ and buyer $v \in V$, let $e = (u, v)$ and p_e denote

*This is the full version of the paper that appeared in AAMAS-2017 [10]. There was an error in one of the offline black-box. This version fixes all the ratios in the theorems.

[†]**Email:** bbrubach@cs.umd.edu

Supported in part by NSF Awards CNS 1010789 and CCF 1422569.

[‡]**Email:** kabinav@cs.umd.edu

Supported in part by NSF Awards CNS 1010789 and CCF 1422569.

[§]**Email:** srin@cs.umd.edu

Supported in part by NSF Awards CNS-1010789, CCF-1422569 and CCF-1749864, and by research awards from Adobe, Inc.

[¶]**Email:** panxu@cs.umd.edu

Supported in part by NSF Awards CNS 1010789 and CCF 1422569.

the the probability that buyer v will buy the item u . If v buys u , then we obtain a reward of w_e . When a buyer arrives, she will be shown items one-by-one until she chooses to buy one. Since every buyer v has a limited attention-span, she can be shown at most t_v items where t_v is typically called a *timeout* or *patience constraint* for v . The goal is to design an algorithm such that the expected reward at the end of the n round process is maximized. Usually the buyer profiles are gathered based on historical information. Hence, we will assume that the system knows all the buyer profiles as well as the buying probabilities for each item profile pair.

We model this as the *Online Stochastic Matching with Timeouts* problem introduced by Bansal, Gupta, Li, Mestre, Nagarajan, and Rudra [7]. Formally, this is a *probe-commit* model for online bipartite stochastic matching. Let us represent an item-buyer pair by an edge $e = (u, v)$. The algorithm can “probe” e to see if the buyer v “buys” the item u . If she does, the decision of selling u to v is made irrevocably (commit). We now define the model in an abstract setting. We are given a bipartite graph $G = (E, U, V)$ as input. Each edge e has a probability p_e (independent of other edges) of existing (modeling a buyer’s interest in an item) and a weight w_e . Each vertex $v \in V$ has a timeout t_v ; however, the vertices in U have no timeout restrictions (equivalently, we can say they have timeouts of infinity). These values are all known a priori. The algorithm proceeds in n rounds. In each round, a vertex v arrives and we can probe at most t_v neighbors in an attempt to match v . Arrivals are drawn with replacement from a known i.i.d. distribution on V . For simplicity, we will consider the uniform distribution.¹ If a probed edge (u, v) is found to exist, we must match v to u and no more probing is allowed for that round. The vertices in U can be matched at most once. The vertices in V are called *types* (a buyer profile) and two or more arrivals of the same type $v \in V$ are considered distinct vertices (two different buyers of a particular profile) which can each be probed up to t_v times and matched to separate neighbors in U . The objective is to maximize the expected weight (or profit) of the final matching obtained. Unsurprisingly, this model captures problems beyond the buyer/seller scenario described above. Various online stochastic matching problem have been considered for online advertising and many other applications [27].

We give new algorithms for the Online Stochastic Matching with Timeouts problem that improve the competitive ratio over the previous work. We then introduce a *new model* wherein the seller selling item u has a limited patience. In this generalization, called *Online Stochastic Matching with Two-sided Timeouts*, we have an additional constraint that every vertex $u \in U$ has a timeout t_u and the algorithm can probe at most t_u neighbors of u across the n rounds. We give the first constant factor approximation algorithm for this generalized setting.

Related Work: Online bipartite matching and its variants has been an active area of study beginning with the seminal work of Karp, Vazirani and Vazirani [24]. They studied weighted bipartite matching in the adversarial arrival model and gave an optimal $(1 - 1/e)$ competitive algorithm. The advent of e-commerce and ad-allocation brought more variants of this problem. For an exhaustive literature survey, refer to the book by Aranyak Mehta [27]. Many variants study arrivals in a random or adversarial order on an unknown set of online vertices. In the I.I.D. arrival model, Feldman *et al.* [18], Bahmani and Kapralov [6], Manshadi *et al.* [26], Haeupler *et al.* [21], Jaillet and Lu [22], and Brubach *et al.* [9] gave improved algorithms to the Online Stochastic Matching problem. The term stochastic here refers to the known i.i.d. arrival model, although some of those papers also address stochastic edge models.

Beyond online matching, other related problems have been studied. The Adwords problem was introduced by Mehta *et al.* [28] and subsequently studied by Buchbinder *et al.* [11] and Devanur and Hayes [13]. More variants have been considered by Devanur *et al.* [15], Devanur *et al.* [16], and Devanur and Jain [14]. Other generalizations that capture the online matching problem are Online Packing Linear Programs by Feldman *et al.* [17] and Agrawal *et al.* [3] and the study of Online Convex Programs by Agrawal and Devanur [2].

Bansal *et al.* [7] introduced the problem of Online Stochastic Matching with Timeouts and gave the first constant factor competitive ratio of 0.12. This was later improved to 0.24 by Adamczyk *et al.* [1]. In both works, they considered the notion of timeouts only on the online vertices. The original motivation for timeouts came from the patience constraints in the Offline Stochastic Matching problem. This offline problem was first introduced by Chen *et al.* [12] and later studied by Bansal *et al.* [7], Adamczyk *et al.* [1], and Baveja *et al.* [8]. A generalization to packing problems was studied by Gupta and Nagarajan [20].

¹Our results are readily applicable to any distribution which yields integral arrival rates.

The Online Stochastic Matching with Two-sided Timeouts problem which we introduce has no direct previous work. One related problem is Online b -matching wherein the offline vertices can each be matched at most b times. This is somewhat similar to having timeouts on the offline vertices in online stochastic matching. In the adversarial setting, b -matching was first studied by Kalyanasundaram and Pruhs [23] and they gave an optimal algorithm. Alaei *et al.* [4] studied the prophet-inequality problem and considered the stochastic i.i.d. setting. They gave an algorithm whose competitive ratio is $1 - \frac{1}{\sqrt{b+3}}$. Alaei *et al.* [5] studied the Online b -matching Stochastic Matching problem in the i.i.d. setting and gave a ratio of $1 - O(1/\sqrt{b})$.

1.1 Preliminaries

We now describe the theoretical preliminaries and the required background for this paper. Henceforth, we refer to the items U as the *offline vertices* and the buyers V as the *online vertices*. We refer to the reward w_e as the *weight* of the edge while we refer to the buying probability p_e as the *edge probability*. We use the words time and round of the process interchangeably. When we say at time $t \in [n]$, it refers to the beginning of step t in the process. Also, following many of the related works, we assume that the arrival rate of each online vertex type is integral and hence, WLOG assume it to be 1². Additionally, the competitive ratio for this class of problems is defined slightly differently from usual online algorithms (see [27]). In particular for a given instance \mathcal{I} , let $\mathbb{E}[\text{ALG}(\mathcal{I})]$ denote the expected reward obtained by the algorithm on this instance and let $\mathbb{E}[\text{OPT}(\mathcal{I})]$ denote the expected value of the offline optimal solution. Then the competitive ratio can be defined as $\min_{\mathcal{I}} \mathbb{E}[\text{ALG}(\mathcal{I})]/\mathbb{E}[\text{OPT}(\mathcal{I})]$. A commonly used technique in the literature is to construct an appropriate linear program, called the benchmark LP, whose optimal value upper bounds $\mathbb{E}[\text{OPT}(\mathcal{I})]$. Hence comparing the expected reward obtained by the algorithm to the optimal value of this LP immediately leads to a lower bound on the competitive ratio. We say that a vertex in the offline set U is *safe* at some time t if it has not been matched in a previous round. Similarly we say an edge $e = (u, v)$ is *safe* in round i iff u is available in an arrival of v at round i . We say a safe edge is *safely probed* if it is probed before v is matched or the timeout t_v is reached. Finally, we denote $\partial(u), \partial(e)$ to denote the edges incident to vertex u and edge e respectively.

Benchmark Linear Program (LP): As in related works, we use the following Linear Program as a benchmark for our competitive ratios.

$$\text{maximize } \sum_{e \in E} w_e f_e p_e \quad (1.1)$$

$$\text{subject to } \sum_{e \in \partial(u)} f_e p_e \leq 1 \quad \forall u \in U \quad (1.2)$$

$$\sum_{e \in \partial(v)} f_e p_e \leq 1 \quad \forall v \in V \quad (1.3)$$

$$\sum_{e \in \partial(u)} f_e \leq t_u \quad \forall u \in U \quad (1.4)$$

$$\sum_{e \in \partial(v)} f_e \leq t_v \quad \forall v \in V \quad (1.5)$$

$$0 \leq f_e \leq 1 \quad \forall e \in E \quad (1.6)$$

The variable f_e in the above LP refers to the expected number of probes on e in the offline optimal. For the problem with timeouts on only the online vertices, WLOG we can assume $t_u = n, \forall u \in U$. Hence constraints 1.2 and 1.3 denotes the matching constraint in the offline graph and constraints 1.4 and 1.5 on the seller and buyer side respectively.

Overview of Attenuation Framework: We present a general online *attenuation framework* for the design and analysis of algorithms for the Online Stochastic Matching with Timeouts problem. In simple terms, an attenuation framework

²If the rate is greater than 1, we can split it into multiple copies each with a rate of 1

is a method for balancing the performance of all edges over all rounds to improve worst case analysis. We analyze the expected performance for every edge individually, across the n rounds and use linearity of expectation to compute the total expected reward. It is easy to see that a lower bound on the competitive ratio can be obtained by considering the edge $e = (u, v)$ with the lowest ratio. In all algorithms, it is common that the neighboring edges $e' \in \partial(e)$ of e have much higher ratios. Moreover, performance of e is negatively affected by that of $\partial(e)$. Thus we can improve the performance of e by attenuating those $e' \in \partial(e)$ and in turn improve the overall competitive ratio.

In each online round when a vertex v arrives, we need to probe at most t_v neighbors sequentially until v is matched. This is essentially an offline stochastic matching problem studied in [7] on a star graph $G(v)$, which consists of v and its safe neighbors. Hence we can view our setting as follows. Suppose we have a *black box*, which is an LP-based algorithm solving an offline stochastic matching problem with timeouts on a general star graph. Then any online framework takes as input a *black box* with a designated property and outputs an online algorithm. The final competitive ratio is determined jointly by the online framework and the input *black box* (see Theorems 3.2, 3.3, 3.4, 3.5).

The idea of *edge-attenuation*, first proposed in [1], aims to balance the performance of all edges. Suppose our black box guarantees that each edge e will be probed with probability at least αf_e , where α is some constant and f_e is the value assigned by an LP. Edge-attenuation will guarantee each edge is probed with probability equal to αf_e . If we know that the black box probes e with probability $\alpha' f_e > \alpha f_e$, we can achieve our goal by setting a probability $1 - \alpha/\alpha'$ with which we “pretend” to probe e , but don’t actually probe it. Unfortunately, the exact value of α' is hard to find since it is jointly determined by multiple sources of randomness in the algorithm and the input. Therefore, we resort to a Monte-Carlo based simulation for a sharp estimate, e.g., simulating the algorithm on the input instance many times and taking the sample mean as an estimate of α' . Simulation-based attenuation has been used before to attack stochastic knapsack problem [25] and offline stochastic matching [1]. As shown in those works, we can ensure that the simulation errors add up to give at most an additive factor of ϵ in the final ratio.

Our novel *vertex-attenuation* approach, applies a similar idea to all the edges incident to an offline vertex. Notice that the star graph centered on an arriving vertex v will have a smaller number of incident edges in later rounds as the offline neighbors of v have been matched in earlier rounds. Intuitively, each edge in this smaller graph has less competition and therefore a greater chance to be probed. Since the probability of being matched in a previous round can differ significantly among the offline vertices, we apply attenuation to the offline vertices to bound these probabilities in our analysis.

Our Contributions: One of our main contributions is a general framework for solving the Online Stochastic Matching with Timeouts problem which gives rise to a class of algorithms. Notably, we decouple the offline subproblem of which edges to probe when a vertex arrives from the online problem of handling a series of arrivals.

The offline subproblem addressed in Section 2 takes as input a stochastic star graph and an LP solution on that graph. The output is a probing strategy which preserves the LP values up to some factor in expectation. Our framework allows the offline subproblem to be solved by any black box algorithm provided it satisfies one of the three basic properties described in Section 3. Thus any new algorithm for this subproblem can easily be plugged into our overall algorithm and hence the analysis. For the purposes of this paper, we plug-in the prior work of [1] as the “black-box” and obtained improved ratios.

For the online framework (Section 3), we bring tighter, cleaner analysis to the edge-attenuation approach in [1] and generalize it to work with a broad class of algorithms for the offline subproblem. For example, [1] uses only the first arrival of each online vertex type, discarding subsequent arrivals of the same type while we use every arrival. We then present a new vertex-attenuation approach which can be combined with edge-attenuation to achieve further improvement in the competitive ratio. The previous algorithm of [1] gave a ratio of 0.24, while we almost double this to give a ratio of 0.46.

We introduce the more general *Online Stochastic Matching with Two-sided Timeouts* problem in Section 3.2. This new problem is well-motivated from applications and theoretically interesting. Using our framework, we give a constant factor 0.30-competitive algorithm for this problem.

Finally, we show in Section 4 that no algorithm using the LP in this paper and prior work can achieve a ratio better than $1 - 1/e \approx 0.632$.

2 Offline Black Box

The online process consists of n offline rounds. In each round, we have an offline stochastic matching instance as studied in [7, 1] on a star graph. Consider a single round at time t : Let v be the arriving vertex and $G(v)$ be the star graph of v and its safe neighbors. For ease of notation, we overload $G(v)$ to denote both the set of edges and the set of neighbors of v . Consider the following polytope:

$$\left\{ \sum_{e \in G(v)} f_e p_e \leq 1, \sum_{e \in G(v)} f_e \leq t_v, 0 \leq f_e \leq 1, \forall e \in G(v) \right\} \quad (2.1)$$

An offline *black box* is any algorithm that transfers a *feasible* solution to the polytope (2.1) to a feasible probing strategy on $G(v)$ with a guaranteed performance for each edge. By a *feasible* probing strategy, we mean one that does not violate the matching and patience constraints on v . Now we present one concrete example of a black box which was known in prior work ([1]). Let $\mathbf{g} = \{g_e | e \in G(v)\}$ be any given feasible solution to the polytope (2.1). We will use g_e to refer to the value of \mathbf{g} for edge e .

Uniform Random Black Box: The Uniform Random Black Box, denoted by BB_{UR} , is a direct application of the algorithm in [7] to the star graph $G(v)$. To be consistent with the notation in [7], we use GKPS to denote the dependent rounding techniques developed in Gandhi *et al* [19].

Algorithm 1: BB_{UR}

- 1 Apply GKPS [19] to $G(v)$ where edge e is associated with a value g_e . Let $\hat{G}(v)$ be the set of edges that gets rounded;
 - 2 Choose a random permutation π over $\hat{G}(v)$. Probe each edge $e \in \hat{G}(v)$ in the order π until v is matched.
-

The performance of BB_{UR} is presented in lemma 2.1. For each given $e \in G(v)$, let $R_{e,\mathbf{g}} = \sum_{e' \neq e} g_{e'} p_{e'}$.

Lemma 2.1. *In BB_{UR} , each edge e will be safely probed with probability at most g_e and at least $\left(1 - \frac{R_{e,\mathbf{g}}}{2}\right) g_e$.*

Proof. Each edge e can be probed only when e appears in \hat{G}_v which occurs with probability g_e according to GKPS [19]. Therefore, we have that each edge e will be probed with probability at most g_e .

Consider an edge e . The random permutation order π can be viewed as: each e' uniformly draws a real number $a_{e'}$ from $[0, 1]$ and we sort all $\{a_{e'}\}$ in increasing order. Condition on $a_e = x \in [0, 1]$ and for each $e' \neq e$, let $X_{e'}$ be the indicator random variable that is 1 if $a_{e'} < x$ (i.e., e' falls before e in π). For each e' , let $Y_{e'}$ and $Z_{e'}$ be the respective indicator random variables corresponding to the event e' being rounded and if e' is present when probed. Let 1_e be the indicator for the event e is probed. Therefore we have

$$\begin{aligned} \Pr[e \text{ is probed} \mid a_e = x, Y_e = 1] &= \mathbb{E}[1_e \mid a_e = x, Y_e = 1] \\ &\geq \mathbb{E}\left[1 - \sum_{e' \neq e} X_{e'} Y_{e'} Z_{e'} \mid a_e = x, Y_e = 1\right] \\ &\geq 1 - x \sum_{e' \neq e} g_{e'} p_{e'} = 1 - x R_{e,\mathbf{g}} \end{aligned}$$

Thus we have

$$\begin{aligned} \Pr[e \text{ is probed}] &= g_e \int_0^1 \Pr[e \text{ is probed} \mid a_e = x, Y_e = 1] dx \\ &\geq g_e(1 - R_{e,g}/2) \end{aligned} \quad \square$$

3 Attenuation Framework for Online Matching with Timeouts

The main idea of our attenuation framework is to decouple the offline and online subproblems. The offline problem is what to do with an arriving vertex once it has arrived and we must choose which edges to probe. This is handled by a black box offline algorithm. The black box is only restricted by one of the three properties listed below. The online problem is how to manage a series of arrivals and that is the primary focus of this section.

Throughout this section, we assume that through Monte-Carlo simulations (referred to as simulations in short, henceforth) we can get an accurate estimate of our target probabilities. As shown in [1, 25], we can manipulate the simulation error appropriately such that we lose at most an additive factor of $\epsilon = o(1)$ in the final ratio.

The three properties of a black box: Property A states that the black box BB is guaranteed to probe each edge with probability at least αg_e for some constant $\alpha \in (0, 1)$. It gives a lower bound on the performance of each edge without any further restrictions on the black box. More formally:

Property A: For any feasible solution \mathbf{g} to LP (2.1), BB outputs a feasible probing strategy $\text{BB}[\mathbf{g}]$ such that every edge e will be probed with probability at least αg_e for some constant $\alpha \in (0, 1)$.

Recall that $R_{e,g} = \sum_{e' \neq e} g_{e'} p_{e'}$ and this value expresses the amount of competition e will face from its neighbors. Properties B and C add the restriction that the probability of probing a given edge will be a function of both g_e and $R_{e,g}$. This allows us to take advantage of the fact that $R_{e,g}$ may decrease as the number of arrivals increases. The conditions of *non-increasing and convexity* on the function R_{BB} are required to ensure that the offline ratio of BB can be used to bound the overall competitive ratio. The condition of *finitely bounded first derivative* guarantees that the error accumulated from simulation is bounded.

Property B: For any feasible \mathbf{g} to LP (2.1), BB outputs a feasible probing strategy $\text{BB}[\mathbf{g}]$ such that each edge e is probed with probability at least $g_e R_{\text{BB}}[R_{e,g}]$, where R_{BB} is a non-increasing and convex function and has finitely bounded first derivative on $[0, 1]$.

Property C adds a further restriction that each edge is probed with probability at most g_e , which states as follows:

Property C: For any feasible \mathbf{g} to LP (2.1), BB outputs a feasible probing strategy $\text{BB}[\mathbf{g}]$ such that every edge e is probed with probability at most g_e and at least $g_e R_{\text{BB}}[R_{e,g}]$, where R_{BB} is a non-increasing and convex function and has finitely bounded first derivative on $[0, 1]$.

Observation 3.1. For any BB satisfying **Property B** or **Property C** with R_{BB} , we have $R_{\text{BB}}[x] \leq R_{\text{BB}}[0] \leq 1$ for all $x \in [0, 1]$.

The fact that $R_{\text{BB}}[0] \leq 1$ can be seen from this example: consider the graph $G(v)$ which has exactly one edge $e = (u, v)$. Clearly, $g_e = 1$ is a feasible solution to LP (2.1). Then, $\text{BB}[\mathbf{g}]$ will probe e with probability at least $R_{\text{BB}}[0]g_e = R_{\text{BB}}[0]$, implying $R_{\text{BB}}[0] \leq 1$.

Note that the black box BB_{UR} satisfies all three properties.

3.1 Attenuation

Our black box properties give us lower bounds on the probability that an edge or vertex will be matched at any given time during the online phase. Attenuation allows us to make those bounds tight by reducing the performance of any

edge or vertex which is exceeding the lower bound. The intuition is that weakening the over-performing edges will increase the performance of the lowest performing edges that provide the worst case competitive ratio.

We define three distinct attenuation frameworks: *edge attenuation* which requires an offline black box satisfying **Property A**, *vertex-attenuation* which requires a black box satisfying **Property C**, and *edge and vertex-attenuation* which requires an offline black box satisfying **Property B**. The edge-attenuation framework generalizes and clarifies the edge-attenuation approach of [1]. We also give an improved result due to tighter analysis. Vertex-attenuation is a novel approach introduced in this paper that upper bounds the probability that a vertex in U will be safe at time t . This lets us exploit the fact that the star graph $G(v)$ will be smaller in later rounds leading to a higher probability of safely probing each of the remaining edges. It can be combined with edge-attenuation to get the best known result for this problem.

Let $\mathbf{f} = \{f_e | e \in E\}$ be an optimal solution to the LP (1.1). Let v be the vertex arriving at time $t \in [n]$ and $G_t(v)$ be the star graph consisting of v and its safe neighbors. Throughout this section we assume $\mathbf{f}_{t,v} = \{f_e | e \in G_t(v)\}$, which is a feasible solution to the LP (2.1) on $G_t(v)$.

Edge-attenuation: The most basic form of attenuation we consider is edge attenuation. Suppose we are given a black box BB satisfying **Property A** and guaranteeing that each edge is probed with probability at least αg_e . This attenuation will guarantee that in each offline subproblem, each edge is probed with probability equal to αg_e .

From **Property A**, we know that $\text{BB}[\mathbf{f}_{t,v}]$ will probe each edge e with probability at least αf_e . *In this framework, we maintain that each safe edge e is probed with probability exactly equal to αf_e in all rounds via appropriate edge-attenuation.* Algorithm 2 gives a formal description of the algorithm.

Algorithm 2: $\text{ATTN}_1[\text{BB}]$

- 1 For each $t \in [n]$, let v be the vertex arriving at time t and $G_t(v)$ be the graph consisting of v and its safe neighbors.
 - 2 Let $\mathbf{f}_{t,v} = \{f_e | e \in G_t(v)\}$ be the induced feasible solution to LP (2.1).
 - 3 Apply $\text{BB}[\mathbf{f}_{t,v}]$ and simulation-based edge-attenuation (see Sec. 1.1) to $G_t(v)$, such that each e is probed with probability exactly equal to αf_e .
-

Theorem 3.2. *For any BB satisfying **Property A**, $\text{ATTN}_1[\text{BB}]$ has an online competitive ratio of $1 - e^{-\alpha}$.*

Proof. Consider an edge $e = (u, v)$ and let $F_u = \sum_{e \in \partial(u)} f_e p_e$. From Algorithm 2, we know that during any round $t \in [n]$, u will be matched with probability exactly equal to $\alpha F_u / n$ (conditioned on u being safe at the beginning of round t). Therefore u will be safe at t with probability equal to $(1 - \alpha F_u / n)^{t-1}$. Thus we have

$$\begin{aligned} \Pr[e \text{ is probed}] &= \sum_{t=1}^n \frac{1}{n} \alpha f_e \left(1 - \frac{\alpha F_u}{n}\right)^{t-1} \\ &\geq f_e \left(1 - \left(1 - \frac{\alpha}{n}\right)^n\right) > f_e (1 - e^{-\alpha}) \end{aligned}$$

We claim that after incorporating the simulation error as shown in Section A.1, we can get a ratio of $1 - (1 - \alpha/n)^n - \epsilon$ for any given ϵ . Thus by setting $\epsilon = e^{-\alpha} - (1 - \alpha/n)^n = O(1/n)$, we get the result in theorem 3.2. \square

Notice that BB_{UR} satisfies **Property A** with $\alpha = 1/2$. Plugging those values into the above theorem, we get ratios of 0.3934.

Corollary 3.1. *When combined with BB_{UR} , the first framework will yield an algorithm which achieves a competitive ratio of 0.3934, for the Online Stochastic Matching with Timeouts problem.*

Although this approach does not give our best result, we note that it places fewer restrictions on the black box than the other approaches presented in this paper. Thus, developing a stronger black box satisfying only **Property A** could lead to the edge-attenuation framework yielding the best result for this problem in the future.

Vertex-attenuation: Applying vertex-attenuation without any edge-attenuation requires an offline black box BB satisfying our most strict property, **Property C**. Here is the intuition behind vertex-attenuation. Notice that over time, the offline vertices in U will be matched and therefore removed from the graph. Suppose we apply $\text{BB}[\mathbf{f}_{t,v}]$ to $G_t(v)$ on each round t when v arrives. Thus when t gets larger and $G_t(v)$ gets smaller, $R_{e,\mathbf{f}_{t,v}}$ will decrease for each safe edge $e = (u, v)$. This means the lower bound on probing an edge, $f_e R_{\text{BB}}[R_{e,\mathbf{f}_{t,v}}]$, will increase with time. We can think of previous approaches to this problem as using a weak bound of $R_{e,\mathbf{f}_{t,v}} \leq 1$. Vertex-attenuation lets us take advantage of $R_{e,\mathbf{f}_{t,v}}$ decreasing by guaranteeing that every offline node is safe with a uniformly decreasing probability at the start of each round t .

Consider a specific round t when v comes. Let $\mathcal{S}_{u,t}$ be the event that u is safe at t for each $u \in U$. We have the below Lemma 3.1.

Lemma 3.1. *Suppose we apply $\text{BB}[\mathbf{f}_{v,t}]$ to $G_t(v)$ during each round t when v arrives. Then for each round $t \in [n]$, we have (1) $\Pr[\mathcal{S}_{u,t}] \geq \Pr[\mathcal{S}_{u,t-1}](1 - 1/n)$ and (2) $\Pr[\mathcal{S}_{u,t} \wedge \mathcal{S}_{u',t} | \mathcal{S}_{u,t-1}, \mathcal{S}_{u',t-1}] \leq \Pr[\mathcal{S}_{u,t} | \mathcal{S}_{u,t-1}] \Pr[\mathcal{S}_{u',t} | \mathcal{S}_{u',t-1}]$.*

Proof. First, we show the proof of inequality (1). Assume u is safe at (the beginning of) $t - 1$. Notice that in the round $t - 1$, every edge $e \in \partial(u)$ will be matched with probability at most $f_e p_e$. Therefore, we have that u will be matched in round $t - 1$ with probability at most $1 - F_u/n$, where $F_u = \sum_{e \in \partial(u)} f_e p_e$. Thus we have $\Pr[\mathcal{S}_{u,t}] \geq \Pr[\mathcal{S}_{u,t-1}](1 - F_u/n) \geq \Pr[\mathcal{S}_{u,t-1}](1 - 1/n)$.

Now we show the proof of inequality (2). Assume both u and u' are safe at time $t - 1$. Consider the round $t - 1$ and assume each edge e is probed with probability $\alpha_e f_e$ with $\alpha_e \in [0, 1]$. Notice that $\Pr[\mathcal{S}_{u,t} \wedge \mathcal{S}_{u',t} | \mathcal{S}_{u,t-1}, \mathcal{S}_{u',t-1}] = 1 - \sum_{e \in \partial(u) \cup \partial(u')} f_e p_e \alpha_e / n$ and $\Pr[\mathcal{S}_{u,t} | \mathcal{S}_{u,t-1}] = 1 - \sum_{e \in \partial(u)} f_e p_e \alpha_e / n$, $\Pr[\mathcal{S}_{u',t} | \mathcal{S}_{u',t-1}] = 1 - \sum_{e \in \partial(u')} f_e p_e \alpha_e / n$. Therefore we get the inequality (2), since $\partial(u)$ and $\partial(u')$ are disjoint. \square

Lemma 3.1 implies that adding vertex-attenuation *independently* to every offline u at the start of each round $t \in [n]$ ensures every u is safe at time t with probability *equal to* $(1 - 1/n)^{t-1}$. Additionally, for two vertices u and u' , this event is negatively correlated. Algorithm 3 describes this online framework, denoted by ATTN_2 .

Algorithm 3: $\text{ATTN}_2[\text{BB}]$

- 1 For each $t \in [n]$, let v be the vertex arriving at time t and $G_t(v)$ be the star graph consisting of v and its safe neighbors.
 - 2 Compute and add attenuation factors to each offline u such that each u is safe at time t with probability *equal to* $(1 - 1/n)^{t-1}$.
 - 3 Let $\mathbf{f}_{t,v} = \{f_e | e \in G_t(v)\}$ be the induced feasible solution to LP (2.1) and $\text{BB}[\mathbf{f}_{t,v}]$ be the feasible probing strategy of BB.
 - 4 Apply $\text{BB}[\mathbf{f}_{t,v}]$ to $G_t(v)$.
-

Theorem 3.3. *For any BB satisfying **Property C** with function R_{BB} , $\text{ATTN}_2[\text{BB}]$ has a competitive ratio of $\int_0^1 e^{-x} R_{\text{BB}}[e^{-x}] dx - \epsilon$.*

Proof. Consider a single edge $e = (u, v)$. Let $\mathcal{A}_{e,t}$ be the event that e is *effectively* probed during round t , i.e., v comes at t , u is safe, and e is probed. Let $\mathcal{S}_{u,t}$ be the event that u is safe at (the beginning of) t , i.e., $e \in G_t(v)$. From ATTN_2 , we have that $\Pr[\mathcal{S}_{u,t}] = (1 - 1/n)^{t-1}$ and $\Pr[\mathcal{A}_{e,t} | \mathcal{S}_{u,t}] \leq (1 - 1/n)^{t-1}$. Let v be the vertex arriving at t . Recall that $R_{e,\mathbf{f}_{t,v}}$ is the sum of $f_{e'} p_{e'}$ over all edges in $G_t(v)$ excluding e itself. Therefore, we have $\mathbb{E}[R_{e,\mathbf{f}_{t,v}} | \mathcal{S}_{u,t}] \leq (1 - 1/n)^{t-1}$.

Thus, from **Property C**, we have

$$\begin{aligned}\Pr[\mathcal{A}_{e,t}] &\geq (f_e/n) \Pr[\mathcal{S}_{u,t}] \mathbb{E}[\mathbb{R}_{\text{BB}}[R_{e,\mathbf{f}_{t,v}}]] \\ &\geq (f_e/n)(1-1/n)^{t-1} \mathbb{R}_{\text{BB}}[\mathbb{E}[R_{e,\mathbf{f}_{t,v}}]] \\ &\geq (f_e/n)(1-1/n)^{t-1} \mathbb{R}_{\text{BB}}[(1-1/n)^{t-1}]\end{aligned}$$

and (The right-most equality below is obtained by letting $n \rightarrow \infty$),

$$\begin{aligned}\Pr[e \text{ is probed}] &= \sum_{t=1}^n \Pr[\mathcal{A}_{e,t}] \\ &\geq \sum_{t=1}^n \frac{f_e}{n} (1-\frac{1}{n})^{t-1} \mathbb{R}_{\text{BB}}[(1-\frac{1}{n})^{t-1}] = \int_0^1 e^{-x} \mathbb{R}_{\text{BB}}[e^{-x}] dx\end{aligned}$$

Incorporating simulation errors (see Section A.2), we get an online ratio of

$$\int_0^1 e^{-x} \mathbb{R}_{\text{BB}}[e^{-x}] dx - \epsilon \text{ for any given } \epsilon > 0. \quad \square$$

Plugging the BB_{UR} function $\mathbb{R}_{\text{BB}_{UR}}[x] = 1 - x/2$ into the above formula, we get a ratio of 0.4159.

Corollary 3.2. *The second framework combined with BB_{UR} yields an algorithm, which achieves a competitive ratio of 0.4159 for the Online Stochastic Matching with Timeouts problem.*

Edge and Vertex-attenuation Combined: Our final and currently most powerful framework combines both edge and vertex-attenuation. Notice that by design, edge-attenuation upper bounds the probability an edge will be probed in an offline step. Therefore, our black box only needs to satisfy **Property B** which is slightly less restrictive than **Property C**.

At the start of each round t , let every u be safe with a target probability equal to $\gamma_t \in [0, 1]$. From **Property B**, we have that each safe edge $e = (u, v)$ is probed during round t with probability at least $f_e \alpha_t$, where $\alpha_t = \mathbb{E}[\mathbb{R}_{\text{BB}}[R_{e,\mathbf{f}_{t,v}}]] \geq \mathbb{R}_{\text{BB}}[\gamma_t]$ (same analysis as Theorem 3.3). Using edge-attenuation, each safe edge e is probed with probability equal to $\mathbb{R}_{\text{BB}}[\gamma_t] f_e$. Consequently, each safe u at time t will remain safe at $t+1$ with probability at least $1 - \mathbb{R}_{\text{BB}}[\gamma_t]/n$. Through vertex-attenuation, each u remains safe at $t+1$ with probability equal to $\gamma_{t+1} = \gamma_t(1 - \mathbb{R}_{\text{BB}}[\gamma_t]/n)$. Thus, we ensure each edge is probed with a uniformly increasing ratio and every offline node is safe with a uniformly decreasing probability. Algorithm 4 describes this online framework denoted ATTN_3 .

Algorithm 4: $\text{ATTN}_3[\text{BB}]$

- 1 For time steps $1, 2, \dots, t$ do
 - 2 Let each u be safe with probability equal to γ_t .
 - 3 Let v arrive at time t and $G_t(v)$ be the graph of v and its safe neighbors. Let $\mathbf{f}_{t,v}$ be the feasible solution to LP (2.1).
 - 4 Apply $\text{BB}[\mathbf{f}_{t,v}]$ and edge-attenuation to $G_t(v)$ such that each edge e is probed with probability equal to $\alpha_t f_e$, where $\alpha_t = \mathbb{R}_{\text{BB}}[\gamma_t]$.
 - 5 Apply vertex-attenuation to each u such that each u is safe at time $t+1$ with probability equal to $\gamma_{t+1} = \gamma_t(1 - \alpha_t/n)$.
-

We can express a recurrence relation for (γ_t, α_t) as follows.

$$\gamma_1 = 1, \quad \alpha_t = \mathbb{R}_{\text{BB}}[\gamma_t]; \quad \gamma_{t+1} = \gamma_t \left(1 - \frac{\alpha_t}{n}\right) \quad (3.1)$$

Theorem 3.4. *For any BB satisfying **Property B**, $\text{ATTN}_3[\text{BB}]$ has an online competitive ratio of $(1 - h(1) - \epsilon)$, where h is the unique function satisfying $h' = -h\mathbb{R}_{\text{BB}}[h]$ with boundary condition $h(0) = 1$. Here, h' represents the first-order derivative of function h .*

Proof. Consider an edge $e = (u, v)$. It will be probed with probability equal to $f_e \sum_{t=1}^n \frac{\gamma_t \alpha_t}{n}$. From Observation 3.1, $R_{BB}[x] \in [0, 1]$ for all $x \in [0, 1]$. From **Property B** and Equation (3.1), we know that $\{\alpha_t\}$ is an increasing sequence and $\{\gamma_t\}$ is a decreasing sequence with $\gamma_t \geq 1/e$ and $\alpha_t \leq R_{BB}[1/e]$ for all t .

Define a function $h : [0, 1] \rightarrow [0, 1]$ such that $h((t-1)/n) = \gamma_t$ for all $t \in [n]$. Thus we have $h(0) = 1$. Equation (3.1) implies that

$$\frac{h(t/n) - h((t-1)/n)}{n} = -h((t-1)/n)R_{BB}[h((t-1)/n)]$$

Letting $x = (t-1)/n$ and the above equation yields $\frac{h(x+1/n) - h(x)}{n} = -h(x)R_{BB}[h(x)]$. Now letting $n \rightarrow \infty$, we can see that h satisfies the differential equation $h' = -hR_{BB}[h]$ with boundary condition $h(0) = 1$.

Given h , we have

$$\begin{aligned} \sum_{t=1}^n \frac{\alpha_t \gamma_t}{n} &= \frac{1}{n} \sum_{t=1}^n h((t-1)/n)R_{BB}[h((t-1)/n)] \\ &= \int_0^1 h(x)R_{BB}[h(x)]dx = h(0) - h(1) = 1 - h(1) \end{aligned}$$

Simulation error subtracts at most $O(\epsilon)$ in the final ratio (see Section A.3). Hence, this completes the proof of the theorem. \square

BB_{UR} satisfies **Property B** with $R_{BB_{UR}}[x] = 1 - x/2$. Plugging $R_{BB_{UR}}$ into the above theorem, we get $h(x) = 2/(1 + e^x)$, which implies $ATTN_3[BB_{UR}]$ has an online ratio of $1 - h(1) - \epsilon \geq 0.4621$.

Corollary 3.3. *The third framework combined with BB_{UR} yields an algorithm, which achieves a competitive ratio of 0.4621 for the Online Stochastic Matching with Timeouts problem.*

3.2 Extensions to a More General Model

The online attenuation framework combined with an offline black box can be extended to more general models. In this section, we give an example by showing how the first attenuation framework together with an offline black box BB satisfying **Property A** can be used for the generalization of Stochastic Matching with timeouts on both offline and online vertices. We do believe the other two frameworks can be used to attack the generalized model as well.

In this model, in addition to our previous setting each offline vertex u has a timeout constraint of t_u , i.e., each u can be probed at most t_u times over the n rounds. Hence, the constraint 1.4 in LP (1.1) is a valid constraint in the benchmark.

Theorem 3.5. *For any BB satisfying **Property A** with α , $ATTN_1[BB]$ has an online competitive ratio of $\alpha e^{-\alpha} - \epsilon$ for the Online Stochastic Matching with Two-sided Timeouts problem.*

Recall that $S_{u,t}$ is the probability that u is safe at time t . In this new setting, u is safe if u is not matched and the timeout of u has not been exhausted. The lemma 3.2 gives a lower bound on $\Pr[S_{u,t}]$ when we apply $ATTN_1[BB]$ using any BB satisfying **Property A** with α .

Lemma 3.2.

$$\Pr[S_{u,t}] \geq \left(1 - \frac{\alpha}{n}\right)^{t-1} \left(1 - \frac{\alpha(t-1)}{n}\right)$$

Proof. Consider a node u . For each $e \in \partial(u)$ and each $t' \in [n]$, let $X_{e,t'}$ be the indicator for the event: e comes at t' ; $Y_{e,t'}$ be the indicator for the event: e is probed when e comes at t' ; $Z_{e,t'}$ be the indicator for the event: e is present when probed. Notice that $\{X_{e,t'}, Y_{e,t'}, Z_{e,t'}\}$ are all independent for each given (t', e) .

Let $S_{u,t}^1$ be the event that u is not matched at time t and $S_{u,t}^2$ be the event that u is probed at most $t_u - 1$ at the beginning of time t . Define A_1 to be the event that $\sum_{t'=1}^{t-1} \sum_{e \in \partial(u)} X_{e,t'} Y_{e,t'} Z_{e,t'} = 0$ and A_2 to be the event that $\sum_{t'=1}^{t-1} \sum_{e \in \partial(u)} X_{e,t'} Y_{e,t'} \leq t_u - 1$. Observe that $\Pr[S_{u,t}^1 \wedge S_{u,t}^2] \geq \Pr[A_1 \wedge A_2]$. Let us now lower bound the value of $\Pr[A_1 \wedge A_2]$.

Recall that $F_u = \sum_{e \in \partial(u)} f_e p_e$. For each given $t' < t$, we know the following — $\Pr[\sum_{e \in \partial(u)} X_{e,t'} Y_{e,t'} Z_{e,t'} = 0] = 1 - \alpha F_u / n \geq 1 - \alpha / n$. Therefore we have $\Pr[A_1] \geq (1 - \alpha / n)^{t-1}$. Notice that for each given $t' < t$,

$$\begin{aligned} \mathbb{E}\left[\sum_{e \in \partial(u)} X_{e,t'} Y_{e,t'} | A_1\right] &= \mathbb{E}\left[\sum_{e \in \partial(u)} X_{e,t'} Y_{e,t'} \mid \sum_{e \in \partial(u)} X_{e,t'} Y_{e,t'} Z_{e,t'} = 0\right] \\ &= \sum_{e \in \partial(u)} \Pr[X_{e,t'} = Y_{e,t'} = 1 \mid \sum_{e \in \partial(u)} X_{e,t'} Y_{e,t'} Z_{e,t'} = 0] \\ &= \sum_{e \in \partial(u)} \frac{\alpha f_e (1 - p_e) / n}{1 - \alpha F_u / n} \\ &\leq \frac{\alpha(t_u - F_u) / n}{1 - \alpha F_u / n} \leq \frac{\alpha t_u}{n} \end{aligned}$$

Thus we get $\mathbb{E}[\sum_{t'=1}^{t-1} \sum_{e \in \partial(u)} X_{e,t'} Y_{e,t'} | A_1] \leq \frac{\alpha t_u}{n} (t-1)$, which implies that $\Pr[A_2 | A_1] \geq (1 - \frac{\alpha}{n}(t-1))$. Therefore

$$\Pr[\mathcal{S}_{u,t}] = \Pr[\mathcal{S}_{u,t}^1 \wedge \mathcal{S}_{u,t}^2] \geq \Pr[A_1 \wedge A_2] \geq (1 - \frac{\alpha}{n})^{t-1} (1 - \frac{\alpha t - \alpha}{n}) \quad \square$$

Let us now prove Theorem 3.5.

Proof. The proof is very similar to that of Theorem 3.2. Consider a single edge $e = (u, v)$, we have

$$\Pr[e \text{ is probed}] \geq \sum_{t=1}^n \frac{1}{n} \alpha f_e \left(1 - \frac{\alpha}{n}\right)^{t-1} \left(1 - \frac{\alpha(t-1)}{n}\right) \geq f_e \alpha e^{-\alpha} \quad \square$$

Plugging BB_{UR} with $\alpha = 0.5$, we get a ratio of 0.303 for this generalized model.

Corollary 3.4. *The first framework combined with BB_{UR} yields an algorithm, which achieves a competitive ratio of 0.303 for the Online Stochastic Matching with Two-sided Timeouts problem.*

4 Lower Bound to the Benchmark LP

Here we present an unconditional lower bound for this LP due to the stochasticity of the problem. We call this lower bound a *stochasticity gap*, similar to the concept of an integrality gap.

Consider a complete bipartite graph with $|U| = |V| = n$. Let the edge probabilities p_e for all edges be $1/n$ and the rewards w_e be 1. Let the patience values for all vertices be n . Notice that assigning $f_e = 1$ for every edge is a feasible solution to LP (1.1). Hence, the optimal LP value is at least n . However, we will show that any online algorithm cannot perform better than $(1 - 1/e)n$. Therefore, the stochasticity gap for this LP is at least $(1 - 1/e) \approx 0.63$.

Consider any vertex $u \in U$. We have the following:

$$\Pr[u \text{ is matched}] = 1 - \Pr[\wedge_{t=1}^n u \text{ is not matched at } t] \quad (4.1)$$

$$= 1 - \prod_{t=1}^n \Pr[u \text{ was not matched at } t] \quad (4.2)$$

$$\leq 1 - \prod_{t=1}^n \left(1 - \frac{1}{n} \times \frac{1}{n} \times n\right) \quad (4.3)$$

$$\leq 1 - 1/e - o(1) \quad (4.4)$$

Equality (4.2) is due to independence. Inequality (4.3) uses union bound, and the facts that (1) $p_e = 1/n$ and (2) with probability $1/n$ each v is drawn in each round. By applying linearity of expectation, we claim that no algorithm can do better than $(1 - 1/e)n$.

5 Conclusion and Future Directions

We gave a general framework for the Online Stochastic Matching with Timeouts problem and its extension. This led to improved competitive ratios for the former and first constant factor ratio for the latter. More importantly, the frameworks are general enough to obtain further improvements by simply finding a better black box for the offline problem on star graphs. One future direction is to increase the competitive ratio by designing better black boxes. Another future direction is to design similar framework(s) for the various other online stochastic matching problems, such as b -matching. We believe these frameworks have the potential to give a unified framework for many of the stochastic matching problems.

References

- [1] ADAMCZYK, M., GRANDONI, F., AND MUKHERJEE, J. Improved approximation algorithms for stochastic matching. In *Algorithms - ESA 2015: 23rd Annual European Symposium, Patras, Greece, September 14-16, 2015, Proceedings* (Berlin, Heidelberg, 2015), N. Bansal and I. Finocchi, Eds., Springer Berlin Heidelberg, pp. 1–12.
- [2] AGRAWAL, S., AND DEVANUR, N. R. Fast algorithms for online stochastic convex programming. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms* (Philadelphia, PA, USA, 2015), SODA '15, Society for Industrial and Applied Mathematics, pp. 1405–1424.
- [3] AGRAWAL, S., WANG, Z., AND YE, Y. A dynamic near-optimal algorithm for online linear programming. *Operations Research* 62, 4 (2014), 876–890.
- [4] ALAEI, S., HAJIAGHAYI, M., AND LIAGHAT, V. Online prophet-inequality matching with applications to ad allocation. In *Proceedings of the 13th ACM Conference on Electronic Commerce* (New York, NY, USA, 2012), EC '12, ACM, pp. 18–35.
- [5] ALAEI, S., HAJIAGHAYI, M., AND LIAGHAT, V. The online stochastic generalized assignment problem. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques: 16th International Workshop, APPROX 2013, and 17th International Workshop, RANDOM 2013, Berkeley, CA, USA, August 21-23, 2013. Proceedings* (Berlin, Heidelberg, 2013), Springer Berlin Heidelberg, pp. 11–25.
- [6] BAHMANI, B., AND KAPRALOV, M. Improved bounds for online stochastic matching. In *Algorithms-ESA 2010*. Springer, 2010, pp. 170–181.
- [7] BANSAL, N., GUPTA, A., LI, J., MESTRE, J., NAGARAJAN, V., AND RUDRA, A. When LP is the cure for your matching woes: Improved bounds for stochastic matchings. In *Algorithms – ESA 2010: 18th Annual European Symposium, Liverpool, UK, September 6-8, 2010. Proceedings, Part II* (Berlin, Heidelberg, 2010), Springer Berlin Heidelberg, pp. 218–229.
- [8] BAVEJA, A., CHAVAN, A., NIKIFOROV, A., SRINIVASAN, A., AND XU, P. Improved bounds in stochastic matching and optimization. In *APPROX-RANDOM 2015, LIPIcs-Leibniz International Proceedings in Informatics* (2015), vol. 40, Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik.
- [9] BRUBACH, B., SANKARARAMAN, K. A., SRINIVASAN, A., AND XU, P. New algorithms, better bounds, and a novel model for online stochastic matching. *European Symposium on Algorithms (ESA)* (2016).
- [10] BRUBACH, B., SANKARARAMAN, K. A., SRINIVASAN, A., AND XU, P. Attenuate locally, win globally: An attenuation-based framework for online stochastic matching with timeouts. In *Proceedings of the 16th Conference*

- on *Autonomous Agents and MultiAgent Systems* (2017), International Foundation for Autonomous Agents and Multiagent Systems, pp. 1223–1231.
- [11] BUCHBINDER, N., JAIN, K., AND NAOR, J. S. Online primal-dual algorithms for maximizing ad-auctions revenue. In *Algorithms – ESA 2007: 15th Annual European Symposium, Eilat, Israel, October 8-10, 2007. Proceedings* (Berlin, Heidelberg, 2007), Springer Berlin Heidelberg, pp. 253–264.
 - [12] CHEN, N., IMMORLICA, N., KARLIN, A. R., MAHDIAN, M., AND RUDRA, A. Approximating matches made in heaven. In *In Proceedings of the 36th International Colloquium on Automata, Languages and Programming, LNCS 5555* (2009), pp. 266–278.
 - [13] DEVANUR, N. R., AND HAYES, T. P. The adwords problem: online keyword matching with budgeted bidders under random permutations. In *Proceedings of the 10th ACM conference on Electronic commerce* (2009), ACM, pp. 71–78.
 - [14] DEVANUR, N. R., AND JAIN, K. Online matching with concave returns. In *Proceedings of the forty-fourth annual ACM symposium on Theory of computing* (2012), ACM, pp. 137–144.
 - [15] DEVANUR, N. R., JAIN, K., SIVAN, B., AND WILKENS, C. A. Near optimal online algorithms and fast approximation algorithms for resource allocation problems. In *Proceedings of the 12th ACM Conference on Electronic Commerce* (New York, NY, USA, 2011), EC '11, ACM, pp. 29–38.
 - [16] DEVANUR, N. R., SIVAN, B., AND AZAR, Y. Asymptotically optimal algorithm for stochastic adwords. In *Proceedings of the 13th ACM Conference on Electronic Commerce* (2012), EC '12.
 - [17] FELDMAN, J., HENZINGER, M., KORULA, N., MIRROKNI, V. S., AND STEIN, C. Online stochastic packing applied to display ad allocation. In *Proceedings of the 18th Annual European Conference on Algorithms: Part I* (Berlin, Heidelberg, 2010), ESA'10, Springer-Verlag, pp. 182–194.
 - [18] FELDMAN, J., MEHTA, A., MIRROKNI, V., AND MUTHUKRISHNAN, S. Online stochastic matching: Beating $1-1/e$. In *Foundations of Computer Science, 2009. FOCS'09. 50th Annual IEEE Symposium on* (2009), IEEE, pp. 117–126.
 - [19] GANDHI, R., KHULLER, S., PARTHASARATHY, S., AND SRINIVASAN, A. Dependent rounding and its applications to approximation algorithms. *Journal of the ACM (JACM)* 53, 3 (2006), 324–360.
 - [20] GUPTA, A., AND NAGARAJAN, V. A stochastic probing problem with applications. In *Integer Programming and Combinatorial Optimization: 16th International Conference, IPCO 2013, Valparaíso, Chile, March 18-20, 2013. Proceedings* (Berlin, Heidelberg, 2013), Springer Berlin Heidelberg, pp. 205–216.
 - [21] HAEUPLER, B., MIRROKNI, V. S., AND ZADIMOGHADDAM, M. Online stochastic weighted matching: Improved approximation algorithms. In *Internet and Network Economics*, vol. 7090 of *Lecture Notes in Computer Science*. Springer Berlin Heidelberg, 2011, pp. 170–181.
 - [22] JAILLET, P., AND LU, X. Online stochastic matching: New algorithms with better bounds. *Mathematics of Operations Research* 39, 3 (2013), 624–646.
 - [23] KALYANASUNDARAM, B., AND PRUHS, K. R. An optimal deterministic algorithm for online b-matching. *Theoretical Computer Science* 233 (2000), 2000.
 - [24] KARP, R. M., VAZIRANI, U. V., AND VAZIRANI, V. V. An optimal algorithm for on-line bipartite matching. In *Proceedings of the twenty-second annual ACM symposium on Theory of computing* (1990), ACM, pp. 352–358.
 - [25] MA, W. Improvements and generalizations of stochastic knapsack and multi-armed bandit approximation algorithms. In *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms* (2014), Society for Industrial and Applied Mathematics, pp. 1154–1163.
 - [26] MANSHADI, V. H., GHARAN, S. O., AND SABERI, A. Online stochastic matching: Online actions based on offline statistics. *Mathematics of Operations Research* 37, 4 (2012), 559–573.

- [27] MEHTA, A. Online matching and ad allocation. *Foundations and Trends in Theoretical Computer Science* 8, 4 (2012), 265–368.
- [28] MEHTA, A., SABERI, A., VAZIRANI, U., AND VAZIRANI, V. Adwords and generalized online matching. *J. ACM* 54, 5 (Oct. 2007).

A Appendix

A.1 Error Accumulation in the First Attenuation Framework

Our simulation-based edge-attenuation approach is very similar to that shown in [1]. For more details, please refer to the Appendix B in [1]. Here we assume that with probability $(1 - \epsilon)$, we can similarly obtain that all safe edges with $f_e \geq \epsilon/n$ should be probed with probability $[\alpha f_e/(1 + \epsilon), \alpha f_e/(1 - \epsilon)]$ in all rounds. For those edges with $f_e < \epsilon/n$, we add no attenuation.

Now we show that the error from simulation accumulates at most $O(1)$ times in the final ratio. Consider an edge $e = (u, v)$. From the analysis in Theorem 3.2, we have that in each round, u will be matched with probability at most $F'_u = \frac{1}{n} \sum_{e \sim \partial(u)} \frac{\alpha f_e p_e}{1 - \epsilon} + \frac{\epsilon}{n} \leq \frac{\alpha}{n} + \frac{2\epsilon}{n}$, if u is safe. Therefore e will be probed with probability at least

$$\Pr[e \text{ is probed}] \geq (1 - \epsilon) \sum_{t=1}^n \frac{1}{n} \frac{\alpha f_e}{1 + \epsilon} \left(1 - \frac{\alpha + 2\epsilon}{n}\right)^{t-1} \geq f_e \left(1 - (1 - \alpha/n)^n - O(\epsilon)\right)$$

A.2 Simulation-Based Vertex-Attenuation in the Second Attenuation Framework

The vertex-attenuation approach is slightly more involved compared to edge-attenuation. Consider a node u and let β_t and β'_t be the probability that u is safe at time t before and after attenuation, respectively. Define the event E_t as $\beta'_t \in [(1 - 1/n)^{t-1} \frac{1}{1+\epsilon}, (1 - 1/n)^{t-1} \frac{1}{1-\epsilon}]$. Here we show how to achieve the goal that $(\bigwedge_{t \in [n]} E_t)$ occurs with probability at least $1 - \epsilon$.

For $t = 1$, we do not need any attenuation. From Lemma 3.1, we have that $\beta_2 \geq (1 - 1/n)$. Let $\hat{\beta}_2$ be the estimation obtained from N experiments. Thus we see

$$\Pr[|\hat{\beta}_2 - \beta_2| \geq \epsilon \beta_2] \leq 2 \exp\left(-\frac{\epsilon^2}{3e} N\right)$$

Thus by setting $N = \ln(2/\epsilon)(3e/\epsilon^2)$, we claim that with probability at least $1 - \epsilon$, $\hat{\beta}_2 \in [\beta_2(1 - \epsilon), \beta_2(1 + \epsilon)]$. The resulting attenuation factor is defined as follows: $\sigma_2 = 1$ if $\hat{\beta}_2 < (1 - 1/n)$ and $\sigma_2 = \frac{1-1/n}{\hat{\beta}_2}$ if $\hat{\beta}_2 \geq (1 - 1/n)$. At the beginning of $t = 2$, we keep u with probability σ_2 and throw away u otherwise, and do this independently for other LHS nodes.

Now assume $\hat{\beta}_2 \in [\beta_2(1 - \epsilon), \beta_2(1 + \epsilon)]$ occurs. Note that $\beta'_2 = \beta_2 * \sigma_2$. We have two cases.

- $\hat{\beta}_2 < (1 - 1/n)$. In this case, we have $\beta_2 < (1 - 1/n) \frac{1}{1-\epsilon}$. And hence, $\beta'_2 = \beta_2 \in [(1 - 1/n), (1 - 1/n) \frac{1}{1-\epsilon}]$.
- $\hat{\beta}_2 \geq (1 - 1/n)$. In this case we have $\beta'_2 = \beta_2 \frac{1-1/n}{\hat{\beta}_2} \in [(1 - 1/n) \frac{1}{1+\epsilon}, (1 - 1/n) \frac{1}{1-\epsilon}]$.

Therefore at $t = 2$ with probability $1 - \epsilon$, the event E_2 occurs. Now assume E_2 occurs. From Lemma 3.1, we see $\beta_3 \geq (1 - 1/n)\beta'_2 \geq (1 - 1/n)^2 \frac{1}{1+\epsilon}$. Using the same analysis as above and the same value of N , we have that with probability $1 - \epsilon$, the event E_3 occurs. Similarly the analysis carries through from E_3, E_4, \dots, E_n . Therefore we claim that $\Pr[\bigwedge_{t \in [n]} E_t] \geq (1 - \epsilon)^n \geq 1 - n\epsilon$. Finally, in each round we scale down the error probability by a factor of $1/n$.

Error Accumulation in the Second Attenuation Framework

From the analysis above, we can safely assume that in the second attenuation framework, with probability at least $(1 - \epsilon)$, $\Pr[\mathcal{S}_{u,t}] \in [(1 - 1/n)^{t-1}(1 - \epsilon), (1 - 1/n)^{t-1}(1 + \epsilon)]$ for all u and t . We can achieve this by setting the simulation

number in each round as $N = \ln(2nm/\epsilon)(3e/\epsilon^2)$. Condition on this, we show how the error accumulates in the final ratio. Consider a given edge $e = (u, v)$. Notice that $\mathbb{E}[R_{e,t,v} | \mathcal{S}_{u,t}] \leq (1 - 1/n)^{t-1}(1 + \epsilon)$, which implies that

$$\mathbb{E}[R_{\text{BB}}[R_{e,t,v}]] \geq R_{\text{BB}}[(1 - 1/n)^{t-1}(1 + \epsilon)] \geq R_{\text{BB}}[(1 - 1/n)^{t-1}] - \epsilon M$$

where M is a constant upper bound for absolute value of the first derivative of R_{BB} over $[0, 1]$. Recall that $\mathcal{A}_{e,t}$ is the event that e is *effectively* probed during round t . Applying the same analysis in Theorem 3.3, we have

$$\begin{aligned} \Pr[e \text{ is probed}] &\geq (1 - \epsilon) \sum_{t=1}^n \Pr[\mathcal{A}_{e,t}] \\ &\geq (1 - \epsilon)^2 \sum_{t=1}^n \frac{f_e}{n} (1 - \frac{1}{n})^{t-1} (R_{\text{BB}}[(1 - \frac{1}{n})^{t-1}] - \epsilon M) = \int_0^1 e^{-x} R_{\text{BB}}[e^{-x}] dx - O(\epsilon) \end{aligned}$$

Therefore by setting ϵ small enough, we can get an online ratio of $\int_0^1 e^{-x} R_{\text{BB}}[e^{-x}] dx - \epsilon$ for any given $\epsilon > 0$.

A.3 Simulation-Based Attenuation in the Third Attenuation Framework

For the third attenuation framework, we need the following key ingredient. Suppose we have a random variable X with $\mathbb{E}[X] = \mu \in [\beta - \epsilon, 1]$ where $0 < \beta < 1$. The random variable models the event that an edge e is probed in some round or a LHS node is safe at some round t . Through analytical analysis, we know a good lower bound β for the unknown mean value μ with error ϵ . Now we need to compute a proper attenuation factor $\sigma \in [0, 1]$ such that $\sigma\mu$ is very close to β with high probability. Consider the following simulation-based approach: we sample the random variable X for N times and let $\hat{\mu}$ be the sample mean; define $\sigma = \beta/\hat{\mu}$ if $\hat{\mu} \geq \beta$ and $\sigma = 1$ otherwise. Assume $\mu \geq \beta - \epsilon \geq \beta/2$.

Lemma A.1. *When $N = \frac{6}{\epsilon^2\beta} \ln \frac{2}{\delta}$, we have that with probability at least $1 - \delta$, $\sigma\mu \in [\beta - \epsilon, \frac{\beta}{1-\epsilon}]$*

Proof. By applying Chernoff Bound, we see

$$\Pr[|\hat{\mu} - \mu| \geq \epsilon\mu] \leq 2 \exp\left(-\frac{\epsilon^2}{3} N\mu\right) \leq 2 \exp\left(-\frac{\epsilon^2}{3} \frac{\beta N}{2}\right) = \delta$$

Thus with probability $1 - \delta$, $\hat{\mu} \in [(1 - \epsilon)\mu, (1 + \epsilon)\mu]$. Assume this occurs. Consider the first case $\hat{\mu} \geq \beta$. Then $\sigma\mu = \beta \frac{\mu}{\hat{\mu}} \in [\beta/(1 + \epsilon), \beta/(1 - \epsilon)]$. We are done since $\beta/(1 + \epsilon) \geq \beta - \epsilon$. For the second case $\hat{\mu} < \beta$, we see that $\mu \leq \beta/(1 - \epsilon)$. Thus we have $\sigma\mu = \mu \in [\beta - \epsilon, \beta/(1 - \epsilon)]$. □

WLOG assume all $f_e \geq \epsilon/n$ and R_{BB} have finite first derivative and upper bounded by $1/2$. Recall that in Equation (3.1), $\alpha_t \geq \alpha_1 > 0$ and $\gamma_t \in [1/e, 1]$ for all $t \in [n]$.

Consider the first round $t = 1$. Consider an edge e and let $\beta'_{e,1}$ be the probability that e is probed during the round $t = 1$ before attenuation. Through the analysis in Section 3.1, we see $\beta'_{e,1} \geq \alpha_1 f_e$. Let $\beta''_{e,1}$ be the probability that e is probed during the round $t = 1$ after attenuation. From Lemma A.1 and by setting $N = O(\ln(1/\delta)n/\epsilon^3)$, we have that with probability $1 - \delta$, $\beta''_{e,1} \in [\alpha_1 f_e - \epsilon, \alpha_1 f_e/(1 - \epsilon)]$. Let A_1 be the event that during round $t = 1$, $\beta''_{e,1} \in [\alpha_1 f_e - \epsilon, \alpha_1 f_e/(1 - \epsilon)]$ for all e . By union bound and by setting $N = O(\ln(nm/\delta)n/\epsilon^3)$ (where $m = |U|$ and $|E| \leq mn$), we can ensure A_1 occurs with probability $1 - \delta$.

Now condition on A_1 . Consider a node u at the beginning of $t = 2$. Let $\gamma'_{u,2}$ and $\gamma''_{u,2}$ be the probability that u is safe at the beginning of $t = 2$ before and after attenuation. We have $\gamma'_{u,2} \geq 1 - \frac{\alpha_1(1-\epsilon)}{n} \geq \gamma_2 - \epsilon_2$ where $\epsilon_2 \doteq 2\epsilon/n$. Here WLOG assume $\epsilon < 1/2$. Let B_2 be the event that $\gamma''_{u,2} \in [\gamma_2 - \epsilon_2, \gamma_2/(1 - \epsilon_2)]$ for all $u \in U$. Similarly we can ensure B_2 occurs with probability $1 - \delta$ by setting $N = O(\ln(m/\delta)n^2/\epsilon^2)$.

Now condition on both A_1 and B_2 . For a given safe edge e , let $\beta'_{e,2}$ be the probability that e is probed during the round $t = 2$ before attenuation. From previous analysis, we have

$$\beta'_{e,2} \geq f_e \mathbf{R}_{\text{BB}}[\gamma_2''] \geq f_e (\mathbf{R}_{\text{BB}}[\gamma_2 + 2\epsilon_2]) \geq f_e \alpha_2 - \epsilon_2$$

Let $\beta''_{e,2}$ be the probability that e is probed during the round $t = 2$ after attenuation and A_2 is the event that $\beta''_{e,2} \in [f_e \alpha_2 - \epsilon_2, f_e \alpha_2 / (1 - \epsilon_2)]$ for all safe e at $t = 2$. Applying Lemma A.1 and union bound, we can make sure that A_2 occurs with probability $1 - \delta$, by setting $N = O(\ln(mn/\delta)n^3/\epsilon^3)$.

Now condition on A_1, B_2, A_2 . Similarly, let $\gamma'_{u,3}$ and $\gamma''_{u,3}$ be the probability that u is safe at time $t = 3$ before and after attenuation. Note that

$$\gamma'_{3,u} \geq \gamma''_{u,2} \left(1 - \frac{\alpha_2/(1 - \epsilon_2)}{n}\right) \geq (\gamma_2 - \epsilon_2) \left(1 - \frac{\alpha_2}{n} - \frac{2\epsilon_2}{n}\right) \geq \gamma_3 - \epsilon_2 \left(1 + \frac{2}{n}\right) \doteq \gamma_3 - \epsilon_3$$

Define B_3 as the event that $\gamma''_{u,3} \in [\gamma_3 - \epsilon_3, \gamma_3 / (1 - \epsilon_3)]$ for all $u \in U$. Similarly we can ensure B_3 occurs with probability $1 - \delta$ by setting $N = O(\ln(m/\delta)n^2/\epsilon^2)$.

Let $\gamma''_{u,t}$ be the probability that u is safe at time t and $\beta''_{e,t}$ the probability that e is probed during round t when it is safe, after attenuation. Define $\epsilon_t = \frac{2\epsilon}{n}(1 + \frac{2}{n})^{t-2}$ for each $t \geq 2$ and $\epsilon_1 = \epsilon$ for $t = 1$. Similarly, let B_t with $t > 1$ be the event that $\gamma''_{u,t} \in [\gamma_t - \epsilon_t, \gamma_t / (1 - \epsilon_t)]$ for all u and A_t with $t \geq 1$ be the event that $\beta''_{e,t} \in [f_e \alpha_t - \epsilon_t, f_e \alpha_t / (1 - \epsilon_t)]$ for all safe e . Doing a similar analysis as above we have the following two observations.

- Condition on $\{A_{t'}, B_{t'} | t' < t\}$. We can ensure that B_t occurs with probability $1 - \delta$ by setting $N = O(\ln(m/\delta)n^2/\epsilon^2)$.
- Condition on $\{A_{t'}, B_{t'} | t' < t\}$ and B_t . We can ensure that A_t occurs with probability $1 - \delta$ by setting $N = O(\ln(mn/\delta)n^3/\epsilon^3)$.

Therefore by setting $\delta = \epsilon/(2n)$, we achieve that with probability $1 - \epsilon$, all events in $\{A_t, B_t | t \in [n]\}$ occur. Note that during each round and for each edge or LHS node, our sampling size is $N = O(\ln(mn^2/\epsilon)n^3/\epsilon^2)$.

Error Accumulation in the Third Attenuation Framework

Now assume all events in $\{A_t, B_t | t \in [n]\}$ occur. Consider an edge $e = (u, v)$. The performance should be at least

$$(1 - \epsilon) \sum_{t=1}^n \frac{\gamma''_{u,t} \beta''_{e,t}}{n} \geq (1 - \epsilon) \sum_{t=1}^n \frac{(\gamma_t - \epsilon)(\alpha_t f_e - \epsilon)}{n} = \sum_{t=1}^n \frac{\gamma_t \alpha_t f_e}{n} - O(\epsilon)$$

This completes the description of the error analysis.