When LP is the Cure for Your Matching Woes: Improved Bounds for Stochastic Matchings

Nikhil Bansal [*] Anupam Gupta [†] Jian Li [‡] Julián Mestre [§]
Viswanath Nagarajan[*] Atri Rudra [¶]

Abstract

Consider a random graph model where each possible edge e is present independently with some probability p_e. Given these probabilities, we want to build a large/heavy matching in the randomly generated graph. However, the only way we can find out whether an edge is present or not is to query it, and if the edge is indeed present in the graph, we are forced to add it to our matching. Further, each vertex i is allowed to be queried at most t_i times. How should we *adaptively* query the edges to maximize the expected weight of the matching? We consider several matching problems in this general framework (some of which arise in kidney exchanges and online dating, and others arise in modeling online advertisements); we give LP-rounding based constant-factor approximation algorithms for these problems. Our main results are the following:

- We give a 4 approximation for weighted stochastic matching on general graphs, and a 3 approximation on bipartite graphs. This answers an open question from [Chen *et al.* ICALP 09].
- Combining our LP-rounding algorithm with the natural greedy algorithm, we give an improved 3.46 approximation for unweighted stochastic matching on general graphs.
- We introduce a generalization of the stochastic *online* matching problem [Feldman *et al.* FOCS 09] that also models preference-uncertainty and timeouts of buyers, and give a constant factor approximation algorithm.

1 Introduction

Motivated by applications in kidney exchanges and online dating, Chen *et al.* [7] proposed the following stochastic matching problem: we want to find a maximum matching in a random graph G on n nodes, where each edge $(i, j) \in [\binom{n}{2}]$ exists with probability p_{ij}, independently of the other edges. However, all we are given are the probability values $\{p_{ij}\}$. To find out whether the random graph G has the edge (i, j) or not, we have to try to add the edge (i, j) to our current matching (assuming that i and j are both unmatched in our current partial matching)—we call this "probing" edge (i, j). As a result of the probe, we also find out if (i, j) exists or not—and if the edge (i, j) indeed exists in the random graph G, it gets irrevocably added to M. Such policies make sense, e.g., for dating agencies, where the only way to find out if two people are actually compatible is to send them on a date; moreover, if they do turn out to be compatible, then it makes sense to match them to each other. Finally, to model the fact that there might be a limit on the number

---

[*] IBM T.J. Watson Research Center, Yorktown Heights, NY 10598, USA.
[†] Computer Science Department, Carnegie Mellon University, Pittsburgh, PA 15213, USA.
[‡] Computer Science Department. University of Maryland College Park, MD, USA.
[§] Max-Planck-Institut für Informatik, Saarbrücken, Germany.
[¶] University at Buffalo, SUNY, Buffalo NY 14260, USA.

of unsuccessful dates a person might be willing to participate in, “timeouts” on vertices are also provided. More precisely, valid policies are allowed, for each vertex i, to only probe at most t_i edges incident to i. Similar considerations arise in kidney exchanges, details of which appear in [7].

Chen *et al.* asked the question: how can we devise probing policies to maximize the expected cardinality (or weight) of the matching? They showed that the greedy algorithm that probes edges in decreasing order of p_{ij} (as long as their endpoints had not timed out) was a 4-approximation to the cardinality version of the stochastic matching problem. This greedy algorithm (and other simple greedy schemes) can be seen to be arbitrarily bad in the presence of weights, and they left open the question of obtaining good algorithms to *maximize the expected weight* of the matching produced. In addition to being a natural generalization, weights can be used as a proxy for revenue generated in matchmaking services. (The unweighted case can be thought of as maximizing the social welfare.) In this paper, we resolve the main open question from Chen *et al.* [7]:

Theorem 1 *There is a* 4*-approximation algorithm for the weighted stochastic matching problem. For bipartite graphs, there is a* 3*-approximation algorithm.*

Our main idea is to use the knowledge of edge probabilities to solve a linear program where each edge e has a variable $0 \leq y_e \leq 1$ corresponding to the probability that a strategy probes e (over all possible realizations of the graph). This is similar to the approach for stochastic packing problems considered by Dean et al. [9, 8]. We then give two different rounding procedures to attain the bounds claimed above.

- The first algorithm (§2.1) is very simple: it considers edges in a uniformly random order and probes each edge e with probability proportional to y_e; the analysis uses Markov’s inequality and a Chernoff-type bound (Lemma 2).
- The second algorithm (§2.2) is more nuanced and achieves a better bound: we use *dependent rounding* [11] on the y-values to obtain a set $\hat{E}$ of edges to be probed, and then probe edges of $\hat{E}$ in a uniformly random order.

Though the first algorithm has a weaker approximation ratio, we still present it since it is useful in the online stochastic matching problem (Section 3).

The second rounding algorithm has an additional advantage: The probing strategy returned by the algorithm can be made *matching-probing* [7]. In this alternative (more restrictive) probing model we are given an additional parameter k and edges need to be probed in k rounds, each round being a matching. It is clear that this matching-probing model is more restrictive than the usual *edge-probing* model (with timeouts $\min\{t_i, k\}$) where one edge is probed at a time. Our algorithm obtains a matching-probing strategy that is only a small constant factor worse than the optimal edge-probing strategy; hence, we also obtain the same constant approximation guarantee for weighted stochastic matching in the matching-probing model. It is worth noting that previously only a logarithmic approximation in the unweighted case was known [7].

Theorem 2 *There is a* 4*-approximation algorithm for the weighted stochastic matching problem in the* **matching-probing** *model. For bipartite graphs, there is a* 3*-approximation algorithm.*

Notice that for general graphs our algorithm matches the performance of the greedy algorithm shown by Chen *et al.* [7] for the unweighted case. Interestingly, even though their individual analyses show that they are 4-approximations, they can be combined to obtain better approximations.

Theorem 3 *There is a* 3.46*-approximation algorithm for the unweighted stochastic matching problem in general graphs.*

Apart from solving these open problems and yielding improved approximations, our LP-based analysis turns out to be applicable in a wider context.

***Online* Stochastic Matching Revisited.** In a bipartite graph $(A, B; E)$ of items $i \in A$ and potential buyer types $j \in B$, p_{ij} denotes the probability that a buyer of type j will buy item i. A sequence of n buyers are to arrive online, where the type of each buyer is an i.i.d. sample from B according to some pre-specified distribution—when a buyer of type j appears, he can be shown a list L of up to t_j as-yet-unsold items, and the buyer buys the *first* item on the list according to the given probabilities $p_{\cdot,j}$. (Note that with probability $\prod_{i \in L}(1 - p_{ij})$, the buyer leaves without buying anything.) What items should we show buyers when they arrive online, and in which order, to maximize the expected weight of the matching? Building on the algorithm for stochastic matching in §2.3, we prove the following in Section 3.

Theorem 4 *There is a 7.92-approximation algorithm for the above online stochastic matching problem.*

This question is an extension of similar online stochastic matching questions considered earlier in [10]—in that paper, $w_{ij}, p_{ij} \in \{0, 1\}$ and $t_j = 1$. Our model tries to capture the facts that buyers may have a limited attention span (using the timeouts), they might have uncertainties in their preferences (using edge probabilities), and that they might buy the first item they like rather than scanning the entire list.

A New Proof for Greedy. The proof in [7] that the greedy algorithm for stochastic matching was a 4-approximation in the unweighted case was based on a somewhat delicate charging scheme involving the decision trees of the algorithm and the optimal solution. We show (Appendix B) that the greedy algorithm, which was defined without reference to any LPs, admits a simple LP-based analysis.

Theorem 5 *The greedy algorithm is a 5-approximation for the unweighted stochastic matching problem.*

Cardinality Constrained Matching in Rounds. We also consider the model from [7] where one can probe as many as C edges in parallel, as long as these C edges form a matching; the goal is to maximize the expected weight of the matched edges after k rounds of such probes. We improve on the $\min\{k, C\}$-approximation offered in [7] (which only works for the unweighted version), and show in Appendix A:

Theorem 6 *There is a constant-factor approximation algorithm for weighted cardinality constrained multiple-round stochastic matching.*

Extension to Hypergraphs. We extend our analysis to a much more general situation where we try to pack k-hyperedges with random sizes into a d-dimensional knapsack of a given size; this is just the stochastic knapsack problem of [8], but where we consider the situation where $k \ll d$. For this setting of parameters, we improve on the $\sqrt{d}$-approximation of [8] to prove the following (Section 4).

Theorem 7 *There is a $2k$-approximation algorithm for the weighted stochastic k-set-packing problem.*

We note that the stochastic k-set-packing problem is a direct generalization of the stochastic matching problem; so an 8-approximation for stochastic matching follows from Theorem 7. However, using more structure in the matching problem, we could obtain the better approximation ratios in Theorems 1 and 2.

1.1 Related Work.

As mentioned above, perhaps the work most directly related to this work is that on stochastic knapsack problems (Dean et al. [9, 8]) and multi-armed bandits (see [13, 14] and references therein). Also related is some recent work [4] on budget constrained auctions, which uses similar LP rounding ideas.

In recent years stochastic optimization problems have drawn much attention from the theoretical computer science community where stochastic versions of several classical combinatorial optimization problems have been studied. Some general techniques have also been developed [15, 23]. See [24] for a survey.

The online bipartite matching problem was first studied in the seminal paper by Karp *et al.* [17] and an optimal $1 - 1/e$ competitive online algorithm was obtained. Katriel *et al.* [18] considered the two-stage stochastic min-cost matching problem. In their model, we are given in a first stage probabilistic information about the graph and the cost of the edges is low; in a second stage, the actual graph is revealed but the costs are higher. The original online stochastic matching problem was studied recently by Feldman *et al.* [10]. They gave a 0.67-competitive algorithm, beating the optimal $1-1/e$-competitiveness known for worst-case models [17, 16, 21, 5, 12]. Our model differs from that in having a bound on the number of items each incoming buyer sees, that each edge is only present with some probability, and that the buyer scans the list linearly (until she times out) and buys the first item she likes. Recently, some improved bounds on this model were obtained [2, 20].

Our problem is also related to the Adwords problem [21], which has applications to sponsored search auctions. The problem can be modeled as a bipartite matching problem as follows. We want to assign every vertex (a query word) on one side to a vertex (a bidder) on the other side. Each edge has a weight, and there is a budget on each bidder representing the upper bound on the total weight of edges that may be assigned to it. The objective is to maximize the total revenue. The stochastic version in which query words arrive according to some known probability distribution has also been studied [19].

1.2 Preliminaries.

For any integer $m \geq 1$, define $[m]$ to be the set $\{1, \ldots, m\}$. For a maximization problem, an α-approximation algorithm is one that computes a solution with expected objective value at least $1/\alpha$ times the expected value of the optimal solution.

We must clarify here the notion of an optimal solution. In standard worst case analysis we would compare our solution against the optimal *offline* solution, e.g. the value of the maximum matching, where the offline knows all the edge instantiations in advance (i.e. which edge will appear when probed, and which will not). However, it can be easily verified that due to the presence of timeouts, this adversary is too strong [7]. Consider the following example. Suppose we have a star where each vertex has timeout 1, and each edge has $p_{ij} = 1/n$. The offline optimum can match an edge whenever the star has an edge i.e. with probability about $1 - 1/e$, while our algorithm can only get expected $1/n$ profit, as it can only probe a single edge. Hence, for all problems in this paper we consider the setting where even the optimum does not know the exact instantiation of an edge until it is probed. This gives our algorithms a level playing field. The optimum thus corresponds to a "strategy" of probing the edges, which can be chosen from an exponentially large space of potentially adaptive strategies.

We note that our algorithms in fact yield *non-adaptive* strategies for the corresponding problems, that are only constant factor worse than the adaptive optimum. This is similar to previous results on stochastic packing problems: knapsack (Dean *et al.* [9, 8]) and multi-armed bandits (Guha-Munagala [13, 14] and references therein).

2 Stochastic Matching

We consider the following stochastic matching problem. The input is an undirected graph $G = (V, E)$ with a weight w_e and a probability value p_e on each edge $e \in E$. In addition, there is an integer value t_v for each vertex $v \in V$ (called *patience parameter*). Initially, each vertex $v \in V$ has patience t_v. At each step in the algorithm, any edge $e(u, v)$ such that u and v have positive remaining patience can be probed. Upon probing edge e, one of the following happens: (1) with probability p_e, vertices u and v get *matched* and are removed from the graph (along with all adjacent edges), or (2) with probability $1 - p_e$, the edge e is removed and the remaining patience numbers of u and v get reduced by 1. An algorithm is an adaptive strategy for probing edges: its performance is measured by the expected weight of matched edges. The *unweighted* stochastic matching problem is the special case when all edge-weights are uniform.

Consider the following linear program: as usual, for any vertex $v \in V$, $\partial(v)$ denotes the edges incident to v. Variable y_e denotes the probability that edge $e = (u, v)$ gets probed in the adaptive strategy, and $x_e = p_e \cdot y_e$ denotes the probability that u and v get matched in the strategy. (This LP is similar to the LP used for general stochastic packing problems by Dean, Goemans and Vondrák [8].)

$$\text{maximize} \sum_{e \in E} w_e \cdot x_e \tag{LP1}$$

subject to

$$\sum_{e \in \partial(v)} x_e \leq 1 \qquad \forall v \in V \tag{1}$$

$$\sum_{e \in \partial(v)} y_e \leq t_i \qquad \forall v \in V \tag{2}$$

$$x_e = p_e \cdot y_e \qquad \forall e \in E \tag{3}$$

$$0 \leq y_e \leq 1 \qquad \forall e \in E \tag{4}$$

The following claim shows that the LP above is a valid relaxation for the stochastic matching problem.

Claim 1 *The optimal value for LP* (LP1) *is an upper bound on any (adaptive) algorithm for stochastic matching.*

Proof: To show this, it suffices to show that any adaptive strategy satisfies the constraints of the LP. Conditioned on any instantiation of all edges in E (i.e. each edge $e \in E$ is present with probability p_e), the expected number of probes involving any vertex $v \in V$ is at most t_v (the patience parameter). Similarly conditioning on edges E, the expected number of matched edges involving $v \in V$ is at most 1. Hence these constraints hold unconditionally as well, which implies that any valid strategy satisfies (1) and (2). ■

2.1 Weighted Stochastic Matching: General Graphs

Our algorithm first solves (LP1) to optimality and uses the optimal solution (x, y) to obtain a non-adaptive strategy achieving expected value $\Omega(1) \cdot (w \cdot x)$. Next, we present the algorithm. Let (x, y) denote an optimal solution to the above LP, which by Claim 1 gives an upper-bound on any adaptive strategy. Let $\alpha \geq 1$ be a constant to be set later. The algorithm first fixes a uniformly random permutation π on edges E. It then inspects edges in the order of π, and *probes* only a subset of the edges. A vertex $v \in V$ is said to have *timed out* if t_v edges incident to v have already been probed (i.e. its remaining patience reduces to 0); and vertex v is said to be *matched* if it has already been matched to another vertex. An edge (u, v) is called *safe* at the time it is considered if (A) neither u nor v is matched, and (B) neither u nor v has timed out. The algorithm is the following:

```
1. Pick a permutation $\pi$ on edges $E$ uniformly at random
2. For each edge $e$ in the ordering $\pi$, do:
    a. If $e$ is safe then probe it with probability $y_e/\alpha$, else do not probe it.
```

In the rest of this section, we prove that this algorithm achieves a 5.75-approximation for the weighted stochastic matching problem. Even though this is slightly worse that the approximation factors claimed in Theorem 1, this first algorithm is significantly simpler, it readily illustrates the power of the LP approach, and, as we shall see in $\S$ 4, it can handle a much more general version of the basic problem.

We begin with the following property:

Lemma 1 *For any edge $(u, v) \in E$, at the point when (u, v) is considered under π,*
(a) the probability that vertex u has timed out is at most $\frac{1}{2\alpha}$, and
(b) the probability that vertex u is matched is at most $\frac{1}{2\alpha}$.

Proof: We begin with the proof of part (a). Let random variable U denote the number of probes incident to vertex u by the time edge (u, v) is considered in π.

$$\begin{aligned}
\mathbb{E}[U] &= \sum_{e \in \partial(u)} \Pr[\,\text{edge } e \text{ appears before } (u,v) \text{ in } \pi \text{ AND } e \text{ is probed}\,], \\
&\leq \sum_{e \in \partial(u)} \Pr[\,\text{edge } e \text{ appears before } (u,v) \text{ in } \pi\,] \cdot \frac{y_e}{\alpha}, \\
&= \sum_{e \in \partial(u)} \frac{y_e}{2\alpha}, \\
&\leq \frac{t_u}{2\alpha}.
\end{aligned}$$

The first inequality above follows from the fact that the probability that edge e is probed (conditioned on π) is at most y_e/α. The second equality follows since π is a u.a.r. permutation on E. The last inequality is by the LP constraint (2). The probability that vertex u has timed out when (u, v) is considered equals $\Pr[\,U \geq t_u\,] \leq \frac{\mathbb{E}[U]}{t_u} \leq \frac{1}{2\alpha}$, by the Markov inequality. This proves part (a). The proof of part (b) is identical (where we consider the event that an edge is matched instead of being probed and replace y_e and t_u by x_e and 1 respectively and use the LP constraint (1)) and is omitted. ■

Now, a vertex $u \in V$ is called *low-timeout* if $t_u = 1$, else u is called a *high-timeout* vertex if $t_u \geq 2$. We next prove the following bound for high-timeout vertices that is stronger than the one from Lemma 1(a).

Lemma 2 *Suppose $\alpha \geq \mathrm{e}$. For a high-timeout vertex $u \in V$, and any edge f incident to u, the probability that u has timed out when f is considered in π is at most $\frac{2}{3\alpha^2}$.*

Proof: Let $t = t_u \geq 2$ denote the patience parameter for vertex u, and $F = \partial(u) \setminus \{f\}$ the set of edges incident to u excluding f. Then the probability that u has timed out when f is considered under π is upper bounded by:

$$\sum_{\{p_1, \cdots, p_t\} \subseteq F} \Pr[\,\text{edges } p_1, \cdots, p_t \text{ appear before } f \text{ in } \pi \text{ AND are all probed}\,], \tag{5}$$

$$\leq \frac{1}{t!} \cdot \sum_{p_1, \cdots, p_t \in F} \Pr[\,\text{edges } p_1, \cdots, p_t \text{ appear before } f \text{ in } \pi \text{ AND are all probed}\,], \tag{6}$$

$$\leq \frac{1}{t!} \cdot \sum_{p_1, \cdots, p_t \in F} \Pr[\text{ edges } p_1, \cdots, p_t \text{ appear before } f \text{ in } \pi\,] \cdot \prod_{\ell=1}^{t} \frac{y_{p_\ell}}{\alpha}, \tag{7}$$

$$= \frac{1}{(t+1)!} \cdot \sum_{p_1, \cdots, p_t \in F} \prod_{\ell=1}^{t} \frac{y_{p_\ell}}{\alpha}, \tag{8}$$

$$= \frac{1}{(t+1)!} \cdot \left(\sum_{p \in F} \frac{y_p}{\alpha} \right)^{t} \tag{9}$$

$$\leq \frac{1}{(t+1)!} \cdot \left(\frac{t}{\alpha} \right)^{t}. \tag{10}$$

In the above, the summation in (5) is over unordered t-tuples whereas the subsequent ones (6)-(8) are over ordered tuples (with repetition). Inequality (7) uses the fact that for any edge g, the probability of probing g *conditioned* on π and the outcomes until g is considered, is at most y_g/α (and the fact that the probability of probing an edge is independent of the probability of probing other edges). Equation (8) follows from the fact that probability that f is the last to appear among $\{p_1, \cdots, p_t, f\}$ in a random permutation π is $\frac{1}{t+1}$. Finally, (10) follows from the LP constraint (2) at u.

Let $f(t) := \frac{1}{(t+1)!} \cdot \left(\frac{t}{\alpha}\right)^t$. We claim that $f(t) \leq \frac{2}{3\alpha^2}$ when $\alpha \geq \mathrm{e}$ and $t \geq 2$, which would prove the claim. Note that this is indeed true for $t = 2$ (in fact with equality). Also $f(t+1) \leq f(t)$ for all $t \geq 2$ due to:

$$\frac{f(t+1)}{f(t)} = \left(\frac{t+1}{t}\right)^t \cdot \frac{t+1}{t+2} \cdot \frac{1}{\alpha} \leq \frac{\mathrm{e}}{\alpha} \leq 1.$$

Thus we obtain the desired upper bound. ■

Using this, we can analyze the probability that an edge is safe.

Lemma 3 *For $\alpha \geq \mathrm{e}$, an edge $f = (u, v)$ is safe with probability at least $(1 - \frac{1}{\alpha} - \frac{4}{3\alpha^2})$ when f is considered under a random permutation π.*

Proof: The analysis proceeds by considering the following cases.

1. *Both u and v are low-timeout.* Since $t_u = t_v = 1$, the event that u (resp. v) is matched at any point is a subset of the event that u (resp. v) has timed out. Thus by Lemma 1, the probability that edge f is not safe (when it is considered) is $\leq \frac{2}{2\alpha}$.

2. *Both u and v are high-timeout.* Lemma 2 implies that the probability that u (resp. v) has timed out is at most $\frac{2}{3\alpha^2}$. Again by Lemma 1, the probability that u (resp. v) is matched is at most $\frac{1}{2\alpha}$. Thus the probability that f is not safe is at most $\frac{1}{\alpha} + \frac{4}{3\alpha^2}$.

3. *u is low-timeout and v is high-timeout.* Using the argument in Step (1) for vertex u, the probability that vertex u has timed out or matched is at most $\frac{1}{2\alpha}$. And using Step (2) for vertex v, the probability that vertex v has timed out or matched is at most $\frac{1}{2\alpha} + \frac{2}{3\alpha^2}$. So the probability that edge (u, v) is not safe is at most $\frac{1}{\alpha} + \frac{2}{3\alpha^2}$.

Hence every edge is safe (when considered in π) with probability $\geq (1 - \frac{1}{\alpha} - \frac{4}{3\alpha^2})$. ■

Theorem 8 *Setting $\alpha = 1 + \sqrt{5}$ in the above algorithm gives an 5.75-approximation for the weighted stochastic matching problem.*

Proof: Theorem 1 Given that an edge $e \in E$ is safe when considered, the expected profit for the algorithm is $w_e \cdot p_e \frac{y_e}{\alpha} = w_e \cdot x_e/\alpha$. Now using Lemma 3, the algorithm gets expected profit at least $(\frac{1}{\alpha} - \frac{1}{\alpha^2} - \frac{4}{3\alpha^3})$ times the optimal LP value. Plugging in $\alpha = 1 + \sqrt{5}$ gives an approximation ratio of $\frac{3(16+8\sqrt{5})}{11+3\sqrt{5}} < 5.75$, as desired. ■

2.2 Weighted Stochastic Matching: Bipartite Graphs

In this section, we obtain an improved bound for stochastic matching on bipartite graphs via a different rounding procedure. In fact, the algorithm produces a *matching-probing strategy* whose expected value is a constant fraction of the optimal value of (LP1) (which was for edge-probing).

Algorithm. First, we find an optimal fractional solution (x, y) to (LP1) and round y to identify a set of interesting edges $\widehat{E}$. Then we use König's Theorem [22, Ch. 20] to partition $\widehat{E}$ into a small collection of matchings $M_1, \ldots, M_h$. Finally, these matchings are then probed in random order. If we are only interested in edge-probing strategies, probing the edges in $\widehat{E}$ in random order would suffice. We will refer to this algorithm as ROUND-COLOR-PROBE:

```
1. $(x, y) \leftarrow$ optimal solution to (LP1)
2. $\widehat{y} \leftarrow$ round $y$ to an integral solution using GKSP
3. $\widehat{E} \leftarrow \{e \in E : \widehat{y}_e = 1\}$
4. $M_1, \ldots, M_h \leftarrow$ optimal edge coloring of $\widehat{E}$
5. For each $M$ in $\{M_1, \ldots, M_h\}$ in random order, do:
     a. probe $\{(u, v) \in M : u$ and $v$ are unmatched$\}$
```

Besides the edge coloring step, the key difference from the algorithm of the previous subsection is in the choice of $\widehat{E}$. For this we use the GKSP procedure of Gandhi *et al.* [11], which we describe next.

The GKSP algorithm. We state some properties of the dependent rounding framework of Gandhi *et al.* [11] that are relevant in our context.

Theorem 9 ([11]) *Let $(A, B; E)$ be a bipartite graph and $z_e \in [0, 1]$ be fractional values for each edge $e \in E$. The GKSP algorithm is a polynomial-time randomized procedure that outputs values $Z_e \in \{0, 1\}$ for each $e \in E$ such that the following properties hold:*

P1. Marginal distribution. For every edge e, $\Pr[Z_e = 1] = z_e$.

P2. Degree preservation. For every vertex $u \in A \cup B$, $\sum_{e\in\partial(u)} Z_e \leq \lceil \sum_{e\in\partial u} z_e \rceil$.

P3. Negative correlation. For any vertex u and any set of edges $S \subseteq \partial(u)$:

$$\Pr[\bigwedge_{e\in S} (Z_e = 1)] \leq \prod_{e\in S} \Pr[Z_e = 1].$$

We note that the GKSP algorithm in fact guarantees stronger properties than the ones stated above. For the purpose of analyzing ROUND-COLOR-PROBE, however, the properties stated above will suffice.

Feasibility. Let us first argue that our algorithm outputs a feasible strategy. If we care about feasibility in the edge-probing model, we only need to show that each vertex u is not probed more than t_u times. The following lemma shows that:

Lemma 4 *For every vertex u,* ROUND-COLOR-PROBE *probes at most t_u edges incident on u.*

Proof: Vertex u is matched in $\left|\left\{e \in \partial_{\widehat{E}}(u)\right\}\right|$ matchings. This is an upper bound on the number of times edges incident on u probed. Hence we just need to show that this quantity is at most t_u. Indeed,

$$\left|\left\{e \in \partial_{\widehat{E}}(u)\right\}\right| = \sum_{e \in \partial(u)} \hat{y}_e \leq \left\lceil \sum_{e \in \partial(u)} y_e \right\rceil \leq t_u,$$

where the first inequality follows from the degree preservation property of Theorem 9 and the second inequality from the fact that y is a feasible solution to (LP1). ■

Let us argue that the strategy is also feasible under the matching-probing model. Recall that in the latter model we are given an additional parameter k (which without loss of generality we can assume to be at most $\max_{v \in V} t_u$) and we can probe edges in k round, with each round forming a matching. Let $\widehat{E}$ be the set of edges in the support of $\widehat{y}$, i.e., $\widehat{E} = \{e \in E \mid \widehat{y}_e = 1\}$. Let $h = \max_{v \in V} \mathsf{deg}_{\widehat{E}}(v) \leq \max_{v \in V} t_v$. König's Theorem allows us to decomposed $\widehat{E}$ into h matchings. Therefore, the probing strategy devised by the algorithm is also feasible in the matching-probing model.

Performance guarantee. Let us focus our attention on some edge $e = (u, v) \in E$. Our goal is to show that there is good chance that the algorithm will indeed probe e. We first analyze the probability of e being probed conditional on $\widehat{E}$. Notice that the algorithm will probe e if and only if all previous probes incident on u and v were unsuccessful; otherwise, if there was a successful probe incident on u or v, we say that e was *blocked*.

Let π be a permutation of the matchings $M_1, \ldots, M_h$. We extend this ordering to the set $\widehat{E}$ by listing the edges within a matching in some arbitrary but fixed order. Let us denote by $B(e, \pi) \subseteq \widehat{E}$ the set of edges incident on u or v that appear before e in π. It is not hard to see that

$$\Pr\left[\, e \text{ was not blocked } \mid \widehat{E}\,\right] \geq \mathbb{E}_\pi \left[\prod_{f \in B(e,\pi)} (1 - p_f) \;\middle|\; \widehat{E} \right]; \tag{11}$$

here we assume that $\prod_{f \in B(e,\pi)} (1 - p_f) = 1$ when $B(e, \pi) = \emptyset$.

Notice that in (11) we only care about the order of edges incident on u and v. Furthermore, the expectation does not range over all possible orderings of these edges, but only those that are consistent with some matching permutation. We call this type of restricted ordering *random matching ordering* and we denote it by π; similarly, we call an unrestricted ordering *random edge ordering* and we denote it by σ. Our plan is to study first the expectation in (11) over random edge orderings and then to show that the expectation can only increase when restricted to range over random matching orderings.

The following simple lemma is useful in several places.

Lemma 5 *Let r and p_{max} be positive real values. Consider the problem of minimizing $\prod_{i=1}^{t}(1 - p_i)$ subject to the constraints $\sum_{i=1}^{t} p_i \leq r$ and $0 \leq p_i \leq p_{\max}$ for $i = 1, \ldots, t$. Denote the minimum value by $\eta(r, p_{\max})$. Then,*

$$\eta(r, p_{\max}) = (1 - p_{\max})^{\left\lfloor \frac{r}{p_{\max}} \right\rfloor} \left(1 - \left(r - \left\lfloor \frac{r}{p_{\max}} \right\rfloor p_{\max}\right)\right) \geq (1 - p_{\max})^{r/p_{\max}}.$$

Proof: Suppose the contrary that the quantity is minimized but there are two p_is that are strictly between 0 and $p_{\max}$. W.l.o.g, they are p_1, p_2 and $p_1 > p_2$ Let $\epsilon = \min(p_{\max} - p_1, p_2)$. It is easy to see that

$$(1-(p_1+\epsilon))(1-(p_2-\epsilon))\prod_{i=3}^{t}(1-p_i) - \prod_{i=1}^{t}(1-p_i) = \epsilon(p_2-p_1-\epsilon)\prod_{i=3}^{t}(1-p_i) < 0.$$

This contradicts the fact the quantity is minimized. Hence, there is at most one p_i which is strictly between 0 and $p_{\max}$.

The last inequality holds since $1-b \geq (1-a)^{b/a}$ for any $0 \leq b \leq a \leq 1$. ■

Let $\partial_{\widehat{E}}(e)$ be the set of edges in $\widehat{E}$ incident on either endpoint of e excluding e itself.

Lemma 6 *Let e be an edge in $\widehat{E}$ and let σ be a random edge ordering. Let $p_{\max} = \max_{f\in\widehat{E}} p_f$. Assume that $\sum_{f\in\partial_{\widehat{E}}(e)} p_f = r$. Then,*

$$\mathbb{E}_\sigma\Big[\prod_{f\in B(e,\sigma)}(1-p_f) \;\Big|\; \widehat{E}\Big] \geq \int_0^1 \eta(xr, xp_{\max})\,\mathrm{d}x.$$

Proof: We claim that the expectation can be written in the following continuous form:

$$\mathbb{E}_\sigma\Big[\prod_{f\in B(e,\sigma)}(1-p_f) \;\Big|\; \widehat{E}\Big] = \int_0^1 \prod_{f\in\partial_{\widehat{E}}(e)}(1-xp_f)\,\mathrm{d}x.$$

The lemma easily follows from this and Lemma 5.

To see the claim, we consider the following random experiment: For each edge $f \in \partial(e)$, we pick uniformly at random a real number a_f in $[0,1]$. The edges are then sorted according to these numbers. It is not difficult to see that the experiment produces uniformly random orderings. For each edge f, let the random variable $A_f = 1-p_f$ if $f \in B(e,\sigma)$ and $A_f = 1$ otherwise. Hence, we have

$$\begin{aligned}
\mathbb{E}_\sigma\Big[\prod_{f\in B(e,\sigma)}(1-p_f) \;\Big|\; \widehat{E}\Big] &= \int_0^1 \mathbb{E}\Big[\prod_{f\in\partial_{\widehat{E}}(e)} A_f \;\Big|\; a_e = x\Big]\,\mathrm{d}x \\
&= \int_0^1 \prod_{f\in\partial_{\widehat{E}}(e)} \mathbb{E}\Big[A_f \;\Big|\; a_e = x\Big]\,\mathrm{d}x \\
&= \int_0^1 \prod_{f\in\partial_{\widehat{E}}(e)} \big(x(1-p_f) + (1-x)\big)\,\mathrm{d}x \\
&= \int_0^1 \prod_{f\in\partial_{\widehat{E}}(e)} (1-xp_f)\,\mathrm{d}x
\end{aligned}$$

The second equality holds since the A_f variables, conditional on $a_e = x$, are independent. ■

Lemma 7 *Let $\rho(r, p_{\max}) = \int_0^1 \eta(xr, xp_{\max})\,\mathrm{d}x$. For any $r, p_{\max} > 0$, we have*

1. *$\rho\,(r, p_{\max})$ is convex and decreasing on r.*

2. $\rho\left(r, p_{\max}\right) \geq \frac{1}{r+p_{\max}} \cdot \left(1-\left(1-p_{\max}\right)^{1+\frac{r}{p_{\max}}}\right) > \frac{1}{r+p_{\max}} \cdot \left(1-e^{-r}\right)$

Proof: To see the first part, let us consider the function values on discrete points $r = p_{\max}, 2p_{\max}, \ldots$. Let $F(x) = \frac{1}{x}(1-c^x)$ where $c = 1 - p_{\max}$. From the above derivation, we can easily get that for integral t,

$$\rho\left(tp_{\max}, p_{\max}\right) = \int_0^1 \left(1 - xp_{\max}\right)^t \mathrm{d}x = \frac{1}{p_{\max}(t+1)}\left(1-c^{t+1}\right) = \frac{1}{p_{\max}} F(t+1).$$

The function $F(x)$ is a convex function for any $0 < c < 1$. Indeed, it is not hard to prove that $\frac{\mathrm{d}^2}{\mathrm{d}x^2}F(x) = \frac{2}{x^3} + c^x\left(-\frac{2}{x^3} + \frac{2\ln a}{x^2} - \frac{\ln^2 a}{x}\right) > 0$ for any $0 < c < 1$. However, $\rho\left(tp_{\max}, p_{\max}\right)$ only coincides with $\frac{1}{p_{\max}}F(t+1)$ at integral values of t. Now, let us consider the value of $\rho(r, p_{\max})$ for $\gamma p_{\max} < r < (\gamma+1)p_{\max}$:

$$\rho\left(r, p_{\max}\right) = \int_0^1 \left(1 - xp_{\max}\right)^\gamma \Big(1 - x(r - \gamma p_{\max})\Big)\,\mathrm{d}x \tag{12}$$

The key observation is that for fixed values of $p_{\max}$ and γ the right hand side of (12) is a just linear function of r. The dependency of ρ in terms of r then becomes clear: it is a piecewise linear function that takes the value $F(t+1)$ at abscissa points $tp_{\max}$ for $t \in \mathbb{Z}_0$. Therefore, ρ is a convex decreasing function of r.

The second part follows easily from Lemma 5:

$$\begin{aligned}\rho\left(r, p_{\max}\right) &= \int_0^1 \eta(xr, xp_{\max})\,\mathrm{d}x \geq \int_0^1 \left(1 - xp_{\max}\right)^{r/p_{\max}} \mathrm{d}x \\ &= \frac{1}{r+p_{\max}} \cdot \left(1-\left(1-p_{\max}\right)^{1+\frac{r}{p_{\max}}}\right) \geq \frac{1}{r+p_{\max}} \cdot \left(1-e^{-r}\right)\end{aligned}$$

■

Lemma 8 *Let $e = (u,v) \in \widehat{E}$. Let π be a random matching ordering and σ be a random edge ordering of the edges adjacent to u and v. Then*

$$\mathbb{E}_\pi\Big[\prod_{f \in B(e,\pi)} (1-p_f) \;\Big|\; \widehat{E}\Big] \geq \mathbb{E}_\sigma\Big[\prod_{f \in B(e,\sigma)} (1-p_f) \;\Big|\; \widehat{E}\Big].$$

Proof: We can think of π as a permutation of bundles of edges: For each matching, if there are two edges incident on e, we bundle the edges together; if there is a single edge incident on e this edge is in a singleton bundle by itself. The random edge ordering σ can be thought as having all edges incident on e in singleton bundles.

Consider the same random experiment as in Lemma 6 except that we only pick one random number for each bundle. Let $G(e)$ be the set of all bundles incident on e. Using the same argument, we have

$$\mathbb{E}_\pi\Big[\prod_{f \in B(e,\pi)} (1-p_f) \;\Big|\; \widehat{E}\Big] = \int_0^1 \prod_{g \in G(e)} \Big(x \cdot \prod_{f \in g}(1-p_f) + (1-x)\Big)\,\mathrm{d}x.$$

But for any bundle $g \in G(e)$ and $0 \leq x \leq 1$, we claim that

$$x \cdot \prod_{f \in g}(1-p_f) + (1-x) \geq \prod_{f \in g}(1-xp_f).$$

For singleton bundles we actually have equality. For a bundle $g = \{f_1, f_2\}$, we have $x(1-p_{f_1})(1-p_{f_2}) + (1-x) = 1 - xp_{f_1} - xp_{f_2} + xp_{f_1}p_{f_2} \geq 1 - xp_{f_1} - xp_{f_2} + x^2 p_{f_1}p_{f_2} = (1-xp_{f_1})(1-xp_{f_2})$. This completes the proof. ■

As we shall see shortly, if $\sum_{f \in \partial_{\widehat{E}}(e)} p_e$ is small then the probability that e is not blocked is large. Because of the marginal distribution property of the GKSP rounding procedure, we can argue that this quantity is small *in expectation* since $\sum_{f \in \partial(e)} p_e y_e \leq 2$ due to the fact that y is a feasible solution to (LP1). This, however, is not enough; in fact, for our analysis to go through, we need a slightly stronger property.

Lemma 9 *For every edge e,*

$$\mathbb{E}\Big[\sum_{f \in \partial_{\widehat{E}}(e)} p_f \mid e \in \widehat{E} \Big] \leq \sum_{f \in \partial(e)} p_f \, y_f .$$

Proof: Let u be an endpoint of e.

$$\begin{aligned}
\mathbb{E}\Big[\sum_{f \in \partial_{\widehat{E}}(u)-e} p_f \mid e \in \widehat{E} \Big] &= \sum_{f \in \partial(u)-e} \Pr[\, \widehat{y}_f = 1 \mid \widehat{y}_e = 1] \cdot p_f, \\
&\leq \sum_{f \in \partial(u)-e} \Pr[\, \widehat{y}_f = 1\,] \cdot p_f, && \text{[by Theorem 9 P3]} \\
&= \sum_{f \in \partial(u)-e} y_f \, p_f. && \text{[by Theorem 9 P1].}
\end{aligned}$$

The same bound holds for the other endpoint of e. Adding the two inequalities we get the lemma. ■

Everything is in place to derive a bound the expected weight of the matching found by our algorithm.

Theorem 10 *If G is bipartite then* ROUND-COLOR-PROBE *is a $1/\rho(2, p_{\max})$ approximation under the edge- and matching-probing model, where ρ is defined in Lemma 7. The worst ratio is attained at $p_{\max} = 1$, where it is 3. The ratio tends to $\frac{2}{1-e^{-2}}$ as $p_{\max}$ tends to 0.*

Proof: Recall that the optimal value of (LP1) is exactly $\sum_{e \in E} w_e y_e x_e$. The expected weight of the matching found by the algorithm is

$$\begin{aligned}
\mathbb{E}[\, \mathsf{ALG}\,] &= \sum_{e \in E} w_e \, p_e \, \Pr[\, e \in \widehat{E}\,] \cdot \Pr[\, e \text{ was not blocked } \mid e \in \widehat{E}] \\
&= \sum_{e \in E} w_e \, p_e \, y_e \cdot \Pr[\, e \text{ was not blocked } \mid e \in \widehat{E}] && \text{[by Theorem 9 P1]} \\
&\geq \sum_{e \in E} w_e \, p_e \, y_e \cdot \mathbb{E}_\pi \Big[\prod_{f \in B(e,\pi)} (1-p_f) \mid e \in \widehat{E} \Big] && \text{[by (11)]} \\
&\geq \sum_{e \in E} w_e \, p_e \, y_e \cdot \mathbb{E}_\sigma \Big[\prod_{f \in B(e,\sigma)} (1-p_f) \mid e \in \widehat{E} \Big] && \text{[by Lemma 8]} \\
&\geq \sum_{e \in E} w_e \, p_e \, y_e \cdot \mathbb{E}\Big[\rho\Big(\sum_{f \in \partial_{\widehat{E}}(e)} p_f, p_{\max} \Big) \mid e \in \widehat{E} \Big] && \text{[by Lemma 6]} \\
&\geq \sum_{e \in E} w_e \, p_e \, y_e \cdot \rho\Big(\mathbb{E}\Big[\sum_{f \in \partial_{\widehat{E}}(e)} p_f \mid e \in \widehat{E} \Big], p_{\max} \Big) && \text{[by Jensen's inequality]}
\end{aligned}$$

$$\geq \sum_{e\in E} w_e\, p_e\, y_e \cdot \rho\Big(\sum_{f\in\partial(e)} y_f\, p_f, p_{\max}\Big) \qquad \text{[by Lemma 9]}$$

$$\geq \sum_{e\in E} w_e\, p_e\, y_e \cdot \rho(2, p_{\max}) \qquad [y \text{ is feasible for (LP1)}].$$

Notice that we are able to use Jensen's inequality because, as shown in Lemma 7, $\rho(r, p_{\max})$ is a convex and decreasing function of r. The last inequality also uses the fact that ρ is decreasing.

It can be checked directly (using the first inequality in Lemma 7(2)) that $\rho(2, p_{max})$ is maximized at $p_{max} = 1$ where it is 3. Moreover $\rho(2, p_{max}) \to \big(1 - e^{-2}\big)/2$ as p_{max} tends to 0. ■

2.3 Weighted Stochastic Matching: General Graphs Redux

We present an alternative algorithm for weighted stochastic matching in general graphs that builds on the algorithm for the bipartite case. The basic idea is to solve (LP1), randomly partition the vertices of G into two sets A and B, and then run ROUND-COLOR-PROBE on the bipartite graph induced by (A, B). For the analysis to go through, it is crucial that we use the already computed fractional solution instead of solving again (LP1) for the new bipartite graph in the call to ROUND-COLOR-PROBE.

```
1. $(x, y) \leftarrow$ optimal solution to (LP1)
2. randomly partition vertices into $A$ and $B$
3. run ROUND-COLOR-PROBE on the bipartite graph and the fractional solution induced by $(A, B)$
```

Theorem 11 *For general graphs there is a $2/\rho(1, p_{\max})$ approximation under the edge- and matching-probing model, where ρ is defined in Lemma 7. The worst ratio is attained at $p_{\max} = 1$, where it is 4. The ratio tends to $\frac{2}{1-e^{-1}}$ as $p_{\max}$ tends to 0.*

Proof: The analysis is very similar to the bipartite case. Essentially, conditional on a particular outcome for the partition (A, B), all the lemmas derived in the previous section hold. In other words, the same derivation done in the proof of Theorem 11 yields:

$$\mathbb{E}[\,\mathsf{ALG} \mid (A, B)\,] \geq \sum_{e\in(A,B)} w_e p_e y_e \cdot \rho\Big(\sum_{f\in\partial_{A,B}(e)} p_f\, y_f, p_{\max}\Big),$$

where $\partial_{A,B}(e) = \partial(e) \cap (A, B)$.

Hence, the expectation of algorithm's performance is:

$$\mathbb{E}[\,\mathsf{ALG}\,] \geq \sum_{e\in E} w_e\, p_e\, y_e \, \Pr[e \in (A, B)] \cdot \mathbb{E}\Big[\rho\Big(\sum_{f\in\partial_{A,B}(e)} p_f\, y_f,\, p_{\max}\Big) \mid e \in (A, B)\Big],$$

$$\geq \sum_{e\in E} w_e\, p_e\, y_e\, \frac{1}{2} \cdot \rho\Big(\mathbb{E}\Big[\sum_{f\in\partial_{A,B}(e)} p_f\, y_f \mid e \in (A, B)\Big],\, p_{\max}\Big),$$

$$\geq \sum_{e\in E} w_e\, p_e\, y_e\, \frac{1}{2} \cdot \rho\Big(\sum_{f\in\partial(e)} \frac{p_f\, y_f}{2},\, p_{\max}\Big),$$

$$\geq \sum_{e\in E} w_e\, p_e\, y_e\, \frac{1}{2} \cdot \rho\big(1,\, p_{\max}\big),$$

where the second inequality follows from Jensen's inequality and the fact that $\rho(r, p_{\max})$ is a convex decreasing function of r. Finally, noting that $\sum_{e\in E} w_e\, p_e\, y_e$ is a lower bound on the value of the optimal strategy, the theorem follows. ■

2.4 Unweighted Stochastic Matching

In this subsection, we consider the unweighted stochastic matching problem, and show that our algorithm from §2.3 can be combined with the natural greedy algorithm [7] to obtain a better approximation guarantee than either algorithm can achieve on their own. Basically, our algorithm attains its worst ratio when $p_{\max}$ is large and greedy attains its worst ratio when $p_{\max}$ is small. Therefore, we can combine the two algorithms as follows: We probe edges using the greedy heuristic until the maximum edge probability in the remaining graph is less than a critical value p_c, at which point we switch to our algorithm from §2.3. We denote by ALG this combined algorithm and by OPT the optimal probing strategy.

Lemma 10 *Suppose that we use the greedy rule until all remaining edges have probability less than p_c, at which point we switch to an algorithm with approximation ratio $\gamma(p_c)$. Then the approximation ratio of the overall scheme is $\alpha(p_c) = \max\{4 - p_c, \gamma(p_c)\}$.*

Proof: First, let us review some facts from the work of Chen *et al.* [7]. Let (G, t) be an instance of the edge-probing model. Suppose $e = (u, v)$ is the edge with the largest probability. Denote by (G_L, t_L) and (G_R, t_R) the instances resulted from the success and failure for the probe to e, respectively. In other words, $G_L = G \setminus \{u, v\}, t_L = t$ and $G_R = G \setminus \{e\}, t_R(u) = t(u) - 1, t_R(v) = t(v) - 1, t_R(w) = t(w)\ \forall w \neq u, v$. Denote the expected value generated by algorithm ALG on instance (G, t) by $\mathbb{E}_{\mathsf{ALG}}(G, t)$. Suppose $p_{max} > p_c$. It is easy to see that, for any ALG that first probes e,

$$\mathbb{E}_{\mathsf{ALG}}(G, t) = p_e + p_e \mathbb{E}_{\mathsf{ALG}}(G_L, t_L) + (1 - p_e)\mathbb{E}_{\mathsf{ALG}}(G_R, t_R). \tag{13}$$

Moreover, Chen *et al.* showed that

$$\mathbb{E}_{\mathsf{OPT}}(G, t) \le p_e(4 - p_e) + p_e \mathbb{E}_{\mathsf{OPT}}(G_L, t_L) + (1 - p_e)\mathbb{E}_{\mathsf{OPT}}(G_R, t_R). \tag{14}$$

Now, we prove the theorem by induction on the size (the number of vertices and edges) of the instance. The base cases are all instances in which the maximum probability is at most p_c. Then $\mathbb{E}_{\mathsf{ALG}}(G, t) \le \gamma(p_c)\mathbb{E}_{\mathsf{OPT}}(G, t) \le \alpha(p_c)\mathbb{E}_{\mathsf{OPT}}(G, t)$ for any base instance. The inductive step only concerns instances where greedy is used. Thus, following from (13) and (14) and the inductive hypothesis, we get

$$\begin{aligned}\mathbb{E}_{\mathsf{OPT}}(G, t) &\le p_e\alpha(p_c) + p_e\alpha(p_c)\mathbb{E}_{\mathsf{ALG}}(G_L, t_L) + (1 - p_e)\alpha(p_c)\mathbb{E}_{\mathsf{ALG}}(G_R, t_R)\\ &\le \alpha(p_c)\mathbb{E}_{\mathsf{ALG}}(G, t).\end{aligned}$$

This completes the inductive proof. ■

We are ready to derive the improved bounds for unweighted stochastic matching.

Proof of Theorem 3: The stated approximation guarantee can be obtained by setting the cut-off point to $p_c = 0.541$ and then using Lemma 10 in combination with Theorem 11 for bounding the performance of the second algorithm at $p_{\max} = p_c$. ■

We remark that the approximation ratio of the algorithm in Section 2.1 does not depend on $p_{\max}$, thus we can not combine that algorithm with the greedy algorithm to get a better bound. Furthermore, the result of this subsection only holds for the unweighted version of the problem since greedy has an unbounded approximation ratio in the weighted case.

3 Stochastic *Online* Matching (Revisited)

As mentioned in the introduction, the stochastic online matching problem is best imagined as selling a finite set of goods to buyers that arrive over time. The input to the problem consists of a bipartite graph $G = (A, B, A \times B)$, where A is the set of *items* that the seller has to offer, with exactly one copy of each item, and B is a set of *buyer types/profiles*. For each buyer type $b \in B$ and item $a \in A$, p_{ab} denotes the probability that a buyer of type b will like item a, and w_{ab} denotes the revenue obtained if item a is sold to a buyer of type b. Each buyer of type $b \in B$ also has a patience parameter $t_b \in \mathbb{Z}_+$. There are n buyers arriving online, with $e_b \in \mathbb{Z}$ denoting the expected number of buyers of type b, with $\sum e_b = n$. Let $\mathcal{D}$ denote the induced probability distribution on B by defining $\Pr_{\mathcal{D}}[b] = e_b/n$. All the above information is given as input.

The stochastic online model is the following: At each point in time, a buyer arrives, where her type $\mathbf{b} \in_{\mathcal{D}} B$ is an i.i.d. draw from $\mathcal{D}$. The algorithm now shows her *up to* $t_{\mathbf{b}}$ *distinct items one-by-one*: the buyer likes each item $a \in A$ shown to her independently with probability p_{ab}. The buyer purchases the first item that she is *offered and likes*; if she buys item a, the revenue accrued is w_{ab}. If she does not like any of the items shown, she leaves without buying. The objective is to maximize the expected revenue.

We get the stochastic online matching problem of Feldman *et al.* [10] if we have $w_{ab} = p_{ab} \in \{0, 1\}$, in which case we need only consider $t_b = 1$. Their focus was on beating the $1 - 1/e$-competitiveness known for worst-case models [17, 16, 21, 5, 12]; they gave a 0.67-competitive algorithm that works for the unweighted case with high probability. On the other hand, our results are for the weighted case (with preference-uncertainty and timeouts), but only in expectation. Furthermore, in our extension, due to the presence of timeouts (see §1.2), any algorithm that provides a guarantee whp must necessarily have a high competitive ratio.

By making copies of buyer types, we may assume that $e_b = 1$ *for all* $b \in B$*, and* $\mathcal{D}$ *is uniform over* B*.* For a particular run of the algorithm, let $\hat{B}$ denote the actual set of buyers that arrive during that run. Let $\hat{G} = (A, \hat{B}, A \times \hat{B})$, where for each $a \in A$ and $\hat{b} \in \hat{B}$ (and suppose its type is some $b \in B$), the probability associated with edge $(a, \hat{b})$ is p_{ab} and its weight is w_{ab}. Moreover, for each $\hat{b} \in \hat{B}$ (with type, say, $b \in B$), set its patience parameter to $t_{\hat{b}} = t_b$. We will call this the *instance graph*; the algorithm sees the vertices of $\hat{B}$ in random order, and has to adaptively find a large matching in $\hat{G}$.

It now seems reasonable that the algorithm of §2.1 should work here. But the algorithm does not know $\hat{G}$ (the actual instantiation of the buyers) up front, it only knows G, and hence some more work is required to obtain an algorithm. Further, as was mentioned in the preliminaries, we use OPT to denote the optimal adaptive strategy (instead of the optimal offline matching in $\hat{G}$ as was done in [10]), and compare our algorithm's performance with this OPT.

The Linear Program. For a graph $H = (A, C, A \times C)$ with each edge (a, c) having a probability p_{ac} and weight w_{ac}, and vertices in C having patience parameters t_j, consider the LP(H):

$$\text{maximize} \sum_{a \in A,\, c \in C} w_{ac} \cdot x_{ac} \tag{LP2}$$

subject to

$$\sum_{c \in C} x_{ac} \le 1 \qquad \forall a \in A \tag{15}$$

$$\sum_{a \in A} x_{ac} \le 1 \qquad \forall c \in C \tag{16}$$

$$\sum_{a \in A} y_{ac} \le t_c \qquad \forall c \in C \tag{17}$$

$$x_{ac} = p_{ac} \cdot y_{ac} \qquad \forall a \in A,\, c \in C \tag{18}$$

$$y_{ac} \in [0,1] \qquad \forall a \in A,\, c \in C \tag{19}$$

Note that this LP is very similar to the one in §2, but the vertices on the left do not have timeout values. Let $\mathsf{LP}(H)$ denote the optimal value of this LP.

The algorithm:

1. Before any buyers arrive, solve the LP on the expected graph G to get values y^*.
2. When any buyer $\hat{b}$ (of type b) arrives online:
    a. If $\hat{b}$ is the first buyer of type b, consider the items $a \in A$ in u.a.r. order. One by one, offer each item a (that is still unsold) to $\hat{b}$ independently with probability y^*_{ab}/α; stop if either t_b offers are made or $\hat{b}$ purchases any item.
    b. If $\hat{b}$ is not the first arrival of type b, do not offer any items to $\hat{b}$.

In the following, we prove that our algorithm achieves a constant approximation to the stochastic online matching problem. The first lemma show that the expected value obtained by the best online adaptive algorithm is bounded above by $\mathbb{E}[\mathsf{LP}(\hat{G})]$.

Lemma 11 *The optimal value* OPT *of the given instance is at most* $\mathbb{E}[\mathsf{LP}(\hat{G})]$*, where the expectation is over the random draws to create* $\hat{G}$*.*

Proof: Consider an algorithm that is allowed to see the instantiation $\hat{B}$ of the buyers before deciding on the selling strategy—the expected revenue of the best such algorithm is clearly an upper bound on OPT. Given any instantiation $\hat{B}$, the expected revenue of the optimal selling strategy is at most $\mathsf{LP}(\hat{G})$ (see e.g. Claim 1). The claim follows by taking an expectation over $\hat{B}$. ■

The proof of the next lemma is similar to the analysis of Theorem 1 for weighted stochastic matching.

Lemma 12 *Our expected revenue is at least* $\left(1 - \frac{1}{e}\right) \frac{1}{\alpha} \left(1 - \frac{1}{\alpha} - \frac{2}{3\alpha^2}\right) \cdot \mathsf{LP}(G)$.

Proof: For any buyer-type $b \in B$, in this proof, $\hat{b}$ refers to the first type-b buyer (if any). For each $b \in B$, let r.v. $T_b \in [n] \cup \{\infty\}$ denote the earliest arrival time of a type-b buyer; if there is no type-b arrival then $T_b = \infty$. Note that our algorithm obtains positive revenue only for buyers $\{\hat{b} \mid b \in B, T_b < \infty\}$; let R_b denote the revenue obtained from buyer $\hat{b}$ (if any). The expected revenue of the algorithm is $\mathbb{E}[\sum_{b \in B} R_b]$. We now estimate $\mathbb{E}[R_b]$ for a fixed $b \in B$.

Let $\mathcal{A}_b \equiv (T_b < \infty)$ denote the event that there is some type-b arrival in the instantiation $\hat{B}$. Since each arrival is i.i.d. from the uniform distribution over B, $\Pr[\mathcal{A}_b] = 1 - (1 - 1/n)^n \geq 1 - \frac{1}{e}$. In the following, we condition on $\mathcal{A}_b$ and bound $\mathbb{E}[R_b \mid \mathcal{A}_b]$. Hence we assume that buyer $\hat{b}$ exists.

For any vertex $a \in A$, let M_a denote the indicator r.v. that a is already matched before time T_b; and O_a (resp. M'_a) the indicator r.v. that $\hat{b}$ is timed-out (resp. already matched) when item a is considered for offering to $\hat{b}$. Now,

$$\begin{aligned} \Pr[\text{ item } a \text{ offered to } \hat{b} \mid \mathcal{A}_b\,] &= (1 - \Pr[\,M_a \cup M'_a \cup O_a \mid \mathcal{A}_b\,]) \cdot \frac{y_{ab}}{\alpha} \\ &\geq (1 - \Pr[\,M_a \mid \mathcal{A}_b\,] - \Pr[\,M'_a \cup O_a \mid \mathcal{A}_b\,]) \cdot \frac{y_{ab}}{\alpha} \end{aligned} \tag{20}$$

Claim 2 *For any* $a \in A$ *and* $b \in B$*,* $\Pr[\,M_a \mid \mathcal{A}_b\,] \leq \frac{1}{2\alpha}$.

Proof: For any $v \in B \setminus \{b\}$, let I_b^v denote the indicator r.v. for the event $T_v < T_b$. We have:

$$\begin{aligned}
\Pr[\, M_a \mid \mathcal{A}_b \,] &= \sum_{v \in B \setminus \{b\}} \Pr[\,\text{type-}v \text{ buyer is matched to } a \text{ before time } T_b \mid \mathcal{A}_b \,] && (21)\\
&= \sum_{v \in B \setminus \{b\}} \Pr[\, I_b^v \mid \mathcal{A}_b \,] \cdot \Pr[\, \hat{v} \text{ matched to } a \mid I_b^v, \mathcal{A}_b \,] && (22)\\
&\le \sum_{v \in B \setminus \{b\}} \Pr[\, I_b^v \mid \mathcal{A}_b \,] \cdot \frac{x_{av}}{\alpha} \quad \le \quad \frac{1}{2} \sum_{v \in B \setminus \{b\}} \frac{x_{av}}{\alpha} \quad \le \quad \frac{1}{2\alpha}, && (23)
\end{aligned}$$

where the first inequality follows from the fact that even after the algorithm has considered an edge (a, v), the probability of matching (a, v) is $\frac{y_{av}}{\alpha} \cdot p_{av}$, the last inequality uses LP-constraint (15) for graph G, and the second last inequality uses $\Pr[\, I_b^v \mid \mathcal{A}_b \,] \le \frac{1}{2}$ (for $v \in B \setminus \{b\}$), which we show next.

Note that event $I_b^v \wedge \mathcal{A}_b$ corresponds to $(T_v < T_b < \infty)$; and event $\mathcal{A}_b$ contains both $(T_v < T_b < \infty)$ and $(T_b < T_v < \infty)$. By symmetry, $\Pr[\, T_v < T_b < \infty \,] = \Pr[\, T_b < T_v < \infty \,]$, which implies:

$$\Pr[\, I_b^v \mid \mathcal{A}_b \,] = \frac{\Pr[\, T_v < T_b < \infty \,]}{\Pr[\mathcal{A}_b]} \le \frac{\Pr[\, T_v < T_b < \infty \,]}{\Pr[\, (T_v < T_b < \infty) \vee (T_b < T_v < \infty) \,]} = \frac{1}{2}.$$

This completes the proof of Claim 2. ■

Claim 3 *For any $a \in A$ and $b \in B$,* $\Pr[\, M'_a \cup O_a \mid \mathcal{A}_b \,] \le \frac{1}{2\alpha} + \frac{2}{3\alpha^2}$.

Proof: This is a direct application of Lemmas 1 and 2, since items offered to $\hat{b}$ are considered in u.a.r. order. As in §2.1, there are two cases:

- Suppose $t_b = 1$. Here we have $M'_a \subseteq O_a$, so $\Pr[\, M'_a \cup O_a \mid \mathcal{A}_b \,] = \Pr[\, O_a \mid \mathcal{A}_b \,] \le \frac{1}{2\alpha}$, by the proof of Lemma 1 using LP-constraint (17).
- Suppose $t_b \ge 2$. Using the proof of Lemma 2 and LP-constraint (17), we have $\Pr[\, O_a \mid \mathcal{A}_b \,] \le \frac{2}{3\alpha^2}$. Again by the proof of Lemma 1 and LP-constraint (16), $\Pr[\, M'_a \mid \mathcal{A}_b \,] \le \frac{1}{2\alpha}$.

In both cases above, the statement in Claim 3 holds. ■

Now applying Claims 2 and 3 to (20), we obtain:

$$\Pr[\,\text{item } a \text{ offered to } \hat{b} \mid \mathcal{A}_b \,] \ge \frac{1}{\alpha} \left(1 - \frac{1}{\alpha} - \frac{2}{3\alpha^2}\right) \cdot y_{ab}.$$

This implies:

$$\begin{aligned}
\mathbb{E}[R_b \mid \mathcal{A}_b] &= \sum_{a \in A} w_{ab} \cdot p_{ab} \cdot \Pr[\,\text{item } a \text{ offered to } \hat{b} \mid \mathcal{A}_b \,]\\
&\ge \frac{1}{\alpha} \left(1 - \frac{1}{\alpha} - \frac{2}{3\alpha^2}\right) \sum_{a \in A} w_{ab} \cdot x_{ab}.
\end{aligned}$$

Since $\Pr[\, \mathcal{A}_b \,] \ge 1 - \frac{1}{e}$, we also have $\mathbb{E}[R_b] \ge (1 - \frac{1}{e}) \frac{1}{\alpha} \left(1 - \frac{1}{\alpha} - \frac{2}{3\alpha^2}\right) \sum_{a \in A} w_{ab} \cdot x_{ab}$.

Finally, the expected revenue obtained by the algorithm is:

$$\sum_{b\in B} \mathbb{E}[R_b] \geq \left(1-\frac{1}{e}\right)\frac{1}{\alpha}\left(1-\frac{1}{\alpha}-\frac{2}{3\alpha^2}\right)\cdot \mathsf{LP}(G).$$

This proves Lemma 12. ■

Note that we have shown that $\mathbb{E}[\mathsf{LP}(\hat{G})]$ is an upper bound on OPT, and that we can get a constant fraction of $\mathsf{LP}(G)$. The final lemma relates these two, namely the LP-value of the expected graph G (computed in Step 1) to the expected LP-value of the instantiation $\hat{G}$; the proof uses a simple but subtle duality-based argument.

Lemma 13 $\mathsf{LP}(G) \geq \mathbb{E}[\mathsf{LP}(\hat{G})]$.

Proof: Consider the dual of the linear program (LP2).

$$\min \sum_{a\in A} \alpha_a + \sum_{c\in C} (\alpha_c + t_c\cdot\beta_c) + \sum_{a\in A,\, c\in C} z_{ac} \tag{24}$$

$$z_{ac} + p_{ac}\cdot(\alpha_a+\alpha_c) + \beta_c \geq w_{ac}\cdot p_{ac} \qquad \forall a\in A,\, c\in C \tag{25}$$

$$\alpha,\beta,z \geq 0 \tag{26}$$

Let (α,β,z) denote the optimal dual solution corresponding to graph G; note that its objective value equals $\mathsf{LP}(G)$ by strong duality. For any instantiation $\hat{G}$, define dual solution $(\hat{\alpha},\hat{\beta},\hat{z})$ as follows:

- For all $a\in A$, $\hat{\alpha}_a = \alpha_a$.
- For each $c\in\hat{B}$ (of type b), $\hat{\alpha}_c = \alpha_b$ and $\hat{\beta}_c = \beta_b$.
- For each $a\in A$ and $c\in\hat{B}$ (of type b), $\hat{z}_{ac} = z_{ab}$.

Note that $(\hat{\alpha},\hat{\beta},\hat{z})$ is a feasible dual solution corresponding to the LP on $\hat{G}$: there is constraint for each $a\in A$ and $c\in\hat{B}$, which reduces to a constraint for (α,β,z) in the dual corresponding to G. By weak duality, the objective value for $(\hat{\alpha},\hat{\beta},\hat{z})$ is an upper-bound on $\mathsf{LP}(\hat{G})$. For each $b\in B$, let N_b denote the number of type b buyers in the instantiation $\hat{B}$; note that $\mathbb{E}[N_b]=1$ by definition of distribution $\mathcal{D}$. Then the dual objective for $(\hat{\alpha},\hat{\beta},\hat{z})$ satisfies:

$$\sum_{a\in A}\alpha_a + \sum_{b\in B} N_b\cdot(\alpha_b + t_b\cdot\beta_b) + \sum_{a\in A,\, b\in B} N_b\cdot z_{ab} \quad \geq \quad \mathsf{LP}(\hat{G}).$$

Taking an expectation over $\hat{B}$, we obtain:

$$\begin{aligned}\mathbb{E}[\mathsf{LP}(\hat{G})] &\leq \sum_{a\in A}\alpha_a + \sum_{b\in B}\mathbb{E}[N_b]\cdot\left(\alpha_b + t_b\cdot\beta_b + \sum_{a\in A} z_{ab}\right)\\ &= \sum_{a\in A}\alpha_a + \sum_{b\in B}(\alpha_b + t_b\cdot\beta_b) + \sum_{a\in A,\, b\in B} z_{ab} = \mathsf{LP}(G).\end{aligned}$$

This proves the lemma. ■

Applying Lemmas 11, 12 and 13, and setting $\alpha = \frac{2}{\sqrt{3}-1}$, completes Theorem 4's proof.

4 Stochastic k-Set Packing

We now consider a generalization of the stochastic matching problem to hypergraphs, where each edge has size at most k. Formally, the input to this *stochastic k-set packing* problem consists of

- n items/columns, where each item has a random profit $v_i \in \mathbb{R}_+$, and a random d-dimensional size $S_i \in \{0,1\}^d$; these random values and sizes are drawn from a probability distribution specified as part of the input. The probability distributions for different items are independent, as are the probability distributions for the value and the size for any of the items. Additionally, for each item, there is a set C_i of at most k coordinates such that each size vector takes positive values only in these coordinates; i.e., $S_i \subseteq C_i$ with probability 1 for each item i.
- A capacity vector $b \in \mathbb{Z}_+^d$ into which the items must be packed.

The parameter k is called the *column sparsity* of the problem. The instantiation of any column (i.e., its size and profit) is known only when it is probed. The goal is to compute an adaptive strategy of choosing items until there is no more available capacity such that the expectation of the obtained profit is maximized.

Note that the stochastic matching problem can be modeled as a stochastic 4-set packing problem in the following way: we set $d = 2n$, and associate the i^{th} and $(n+i)^{th}$ coordinate with the vertex i—the first n coordinates capture whether the vertex is free or not, and the second n coordinates capture how many probes have been made involving that vertex. Now each edge (i,j) is an item whose value is w_{ij}; if $e_t \in \{0,1\}^d$ denotes the indicator vector with a single 1 in the t^{th} position, then the size of the edge (i,j) is either $e_i + e_j + e_{n+i} + e_{n+j}$ (with probability p_i) or $e_{n+i} + e_{n+j}$ (with probability $1 - p_i$). If we set the capacity vector to be $b = (1, 1, \cdots, 1, t_1, t_2, \cdots, t_n)$, this precisely captures the stochastic matching problem. Thus, each size vector has $\leq k = 4$ ones.

This stochastic k-set packing problem was studied (among many others) as the "stochastic b-matching" problem in Dean *et al.* [8]; however the authors of that work did not consider the 'column sparsity' parameter k and instead gave an $O(\sqrt{d})$-approximation algorithm for the general. Here we consider the performance of algorithms for this problem specifically as a function of the column sparsity k, and prove Theorem 7.

A quick aside about "safe" and "unsafe" adaptive policies: a policy is called *safe* if it can include an item only if there is *zero* probability of violating any capacity constraint. In contrast, an *unsafe* policy may attempt to include an item even if there is non-zero probability of violating capacity—however, if the random size of the item causes the capacity to be violated, then no profit is received for the overflowing item, and moreover, no further items may be included by the policy. The model in Dean et al. [8] allowed unsafe policies, whereas we are interested (as in the previous sections) in safe policies. However, due to the discreteness of sizes in stochastic k-set packing, it can be shown that our approximation guarantee is relative to the optimal unsafe policy.

For each item $i \in [n]$ and constraint $j \in [d]$, let $\mu_i(j) := \mathbb{E}[S_i(j)]$, the expected value of the j^{th} coordinate in size-vector S_i. For each column $i \in [n]$, the coordinates $\{j \in [d] \mid \mu_i(j) > 0\}$ are called the *support* of column i. By column sparsity, the support of each column has size at most k. Also, let $w_i := \mathbb{E}[v_i]$, the mean profit, for each $i \in [n]$. We now consider the natural LP relaxation for this problem, as in [8].

$$\text{maximize} \sum_{i=1}^{n} w_i \cdot y_i \tag{LP3}$$

subject to

$$\sum_{i=1}^{n} \mu_i(j) \cdot y_i \leq b_j \qquad \forall j \in [d] \tag{27}$$

$$y_i \in [0,1] \qquad \forall i \in [n] \tag{28}$$

Let y^* denote an optimal solution to this linear program, which in turn gives us an upper bound on any adaptive (safe) strategy. Our rounding algorithm is a natural extension of the one for stochastic matching in §2.1. Fix a constant $\alpha \geq 1$, to be specified later. The algorithm picks a uniformly random permutation $\pi : [n] \to [n]$ on all columns, and probes only a subset of the columns as follows. At any point in the algorithm, column c is *safe* iff there is positive residual capacity in *all* the coordinates in the support of c—in other words, irrespective of the instantiation of S_c, it can be feasibly packed with the previously chosen columns. The algorithm inspects columns in the order of π, and whenever it is safe to probe the next column $c \in [n]$, it does so with probability $\frac{y_c}{\alpha}$. Note that the algorithm skips all columns that are unsafe at the time they appear in π.

We now prove Theorem 7 by showing that this algorithm is a $2k$-approximation. The analysis proceeds similar to that in §2.1. For any column $c \in [n]$, let $\{\mathbf{I}_{c,\ell}\}_{\ell=1}^{k}$ denote the indicator random variables for the event that the ℓ^{th} constraint in the support of c is tight at the time when c is considered under the random permutation π. Note that the event "column c is safe when considered" is precisely $\bigwedge_{\ell=1}^{k} \overline{\mathbf{I}_{c,\ell}}$. By a trivial union bound, the $\Pr[\, c \text{ is safe} \,] \geq 1 - \sum_{\ell=1}^{k} \Pr[\mathbf{I}_{c,\ell}]$.

Lemma 14 *For any column $c \in [n]$ and index $\ell \in [k]$,* $\Pr[\mathbf{I}_{c,\ell}] \leq \frac{1}{2\alpha}$.

Proof: Let $j \in [d]$ be the ℓ^{th} constraint in the support of c. Let U_c^j denote the usage of constraint j, when column c is considered (according to π). Then, using argument similar to those used to prove Lemma 1, we have

$$\begin{aligned}
\mathbb{E}[U_c^j] &= \sum_{a=1}^{n} \Pr[\,\text{column } a \text{ appears before } c \text{ AND } a \text{ is probed}\,] \cdot \mu_a(j), \\
&\leq \sum_{a=1}^{n} \Pr[\,\text{column } a \text{ appears before } c\,] \cdot \frac{y_a}{\alpha} \cdot \mu_a(j), \\
&= \sum_{a=1}^{n} \frac{y_a}{2\alpha} \cdot \mu_a(j), \\
&\leq \frac{b_i}{2\alpha}.
\end{aligned}$$

Since $\mathbf{I}_{c,\ell} = \{U_c^i \geq b_i\}$, Markov's inequality implies that $Pr[\mathbf{I}_{c,\ell}] \leq \mathbb{E}[U_c^i]/b_i \leq \frac{1}{2\alpha}$. ■

Again using the trivial union bound, the probability that a particular column c is safe when considered under π is at least $1 - \frac{k}{2\alpha}$, and thus the probability of actually probing c is at least $\frac{y_c}{\alpha}(1 - \frac{k}{2\alpha})$. Finally, by linearity of expectations, the expected profit is at least $\frac{1}{\alpha}(1 - \frac{k}{2\alpha}) \cdot \sum_{c=1}^{n} w_c \cdot y_c$. Setting $\alpha = k$ implies an expected profit of at least $\frac{1}{2k} \cdot \sum_c w_c y_c$, which proves Theorem 7.

5 Final Remarks

An extended abstract of this paper appeared in the Proceedings of the 18th Annual European Symposium on Algorithms [3]. The bounds presented here in §2 are slightly better than those claimed in the extended abstract. Quite recently, Adamczyk has proved that the greedy algorithm is a 2-approximation for unweighted stochastic matching [1], improving our bounds from Theorem 3. It remains an open question whether the stochastic matching problem is NP-complete.

Acknowledgements. We would like to thank Aravind Srinivasan for helpful discussions.

References

[1] M. Adamczyk. Greedy algorithm for stochastic matching is a 2-approximation, 2010. arXiv:1007.3036.

[2] B. Bahmani and M. Kapralov. Improved bounds for online stochastic matching. In *ESA*, 2010.

[3] N. Bansal, A. Gupta, J. Li, J. Mestre, V. Nagarajan, and A. Rudra. When LP is the Cure for Your Matching Woes: Improved Bounds for Stochastic Matchings. In *ESA*, 2010.

[4] S. Bhattacharya, G. Goel, S. Gollapudi, and K. Munagala. Budget constrained auctions with heterogeneous items. In *STOC*, 2009. Arxiv:abs/0907.4166.

[5] B. E. Birnbaum and C. Mathieu. On-line bipartite matching made simple. *SIGACT News*, 39(1):80–87, 2008.

[6] R. Carr and S. Vempala. Randomized metarounding. *Random Structures Algorithms*, 20(3):343–352, 2002. Probabilistic methods in combinatorial optimization.

[7] N. Chen, N. Immorlica, A. R. Karlin, M. Mahdian, and A. Rudra. Approximating matches made in heaven. In *ICALP (1)*, pages 266–278, 2009.

[8] B. C. Dean, M. X. Goemans, and J. Vondrák. Adaptivity and approximation for stochastic packing problems. In *SODA*, pages 395–404, 2005.

[9] B. C. Dean, M. X. Goemans, and J. Vondrák. Approximating the stochastic knapsack problem: the benefit of adaptivity. *Math. Oper. Res.*, 33(4):945–964, 2008.

[10] J. Feldman, A. Mehta, V. S. Mirrokni, and S. Muthukrishnan. Online stochastic matching: Beating $1 - 1/e$. In *FOCS*, 2009. *Arxiv:abs/0905.4100*.

[11] R. Gandhi, S. Khuller, S. Parthasarathy, and A. Srinivasan. Dependent rounding and its applications to approximation algorithms. *JACM*, 53(3):360, 2006.

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

[13] S. Guha and K. Munagala. Approximation algorithms for partial-information based stochastic control with markovian rewards. In *FOCS*, pages 483–493, 2007.

[14] S. Guha and K. Munagala. Multi-armed bandits with metric switching costs. In *ICALP*, pages 496–507, 2009.

[15] A. Gupta, M. Pál, R. Ravi, and A. Sinha. Boosted sampling: approximation algorithms for stochastic optimization. In *STOC*, pages 417–426. ACM, 2004.

[16] B. Kalyanasundaram and K. Pruhs. Online weighted matching. *J. Algorithms*, 14(3):478–488, 1993.

[17] R. M. Karp, U. V. Vazirani, and V. V. Vazirani. An optimal algorithm for on-line bipartite matching. In *STOC*, pages 352–358, 1990.

[18] I. Katriel, C. Kenyon-Mathieu, and E. Upfal. Commitment under uncertainty: Two-stage stochastic matching problems. *Theoretical Computer Science*, 408(2-3):213–223, 2008.

[19] M. Mahdian, H. Nazerzadeh, and A. Saberi. Allocating online advertisement space with unreliable estimates. In *EC*, page 294, 2007.

[20] V. H. Manshadi, S. O. Gharan, and A. Saberi. Online stochastic matching: Online actions based on offline statistics, 2010. arXiv:1007.1673.

[21] A. Mehta, A. Saberi, U. V. Vazirani, and V. V. Vazirani. Adwords and generalized on-line matching. In *FOCS*, pages 264–273, 2005.

[22] A. Schrijver. *Combinatorial Optimization*. Springer-Verlag, 2003.

[23] D. Shmoys and C. Swamy. An approximation scheme for stochastic linear programming and its application to stochastic integer programs. *J. ACM*, 53(6):1012, 2006.

[24] C. Swamy and D. Shmoys. Approximation algorithms for 2-stage stochastic optimization problems. *ACM SIGACT News*, 37(1):46, 2006.

A Cardinality Constrained Multiple Round Stochastic Matching

We now consider stochastic matching with a different objective in mind; this was also defined in [7]. In this problem, we arrange for many pairs to date each other simultaneously (constrained by the fact that each person is involved in at most one date at any time), and have k days in which all these dates must happen—again, we want to maximize the expected weight of the matched pairs.

More formally, we can probe several edges concurrently—a "round" may involve probing any set of edges that forms a matching of size at most C. Given k and C, the goal is to find an adaptive strategy for probing edges in rounds such that we use at most k rounds, and maximize the expected weight of matched edges during these k rounds. As before, one can probe edges involving individual i at most t_i times, and only if i is not already matched by the algorithm. In this section, we give a constant-factor approximation for this problem, improving over the previously known $O(\min\{k, C\})$-approximation [7] (which only works for the unweighted case).

Our approach, as in the previous sections, is based on linear programming. The following LP captures adaptive strategies, and hence is a relaxation of the multiple round stochastic matching problem; moreover, it can be solved in poly-time. Below, $\mathcal{M}_C(G)$ denotes the convex hull of all matchings in G having size at most C.

$$\text{maximize} \sum_{(i,j)\in E} w_{ij} \cdot \sum_{h=1}^{k} x_{ij}^h \tag{LP4}$$

subject to

$$\sum_{h=1}^{k} y_{ij}^h \leq 1 \qquad \forall (i,j) \in E \tag{29}$$

$$\sum_{j \in \partial(i)} \sum_{h=1}^{k} y_{ij}^h \le t_i \qquad \forall i \in V \tag{30}$$

$$y^h \in \mathcal{M}_C(G) \qquad \forall h \in [k] \tag{31}$$

$$x_{ij}^h = p_{ij} \cdot y_{ij}^h \qquad \forall (i,j) \in E,\ h \in [k] \tag{32}$$

$$\sum_{j \in \partial(i)} \sum_{h=1}^{k} x_{ij}^h \le 1 \qquad \forall i \in V \tag{33}$$

Since there is a linear description for $\mathcal{M}_C(G)$, for which we can separate in polynomial time [22, Corollary 18.10a]), the above LP can be solved in polynomial time using, say, the Ellipsoid algorithm. To see that this LP is indeed a relaxation of the original adaptive problem, observe that setting y_{ij}^h to be "probability that (ij) is probed in round h by the optimal strategy" defines a feasible solution to the LP with objective equal to the optimal value of the stochastic matching instance.

Our algorithm first solves the LP to optimality and obtains solution (x, y). Note that for each $h \in [k]$, using the fact that polytope $\mathcal{M}_C(G)$ is integral and that the variables $y^h \in \mathcal{M}_C(G)$, we can write y^h as a convex combination of matchings of size at most C; i.e., we can find matchings $\{M_\ell^h\}_\ell$ and positive values $\{\lambda_\ell^h\}_\ell$ such that each M_ℓ^h is a matching in G of size at most C and $y^h = \sum \lambda_\ell^h \cdot \chi(M_\ell^h)$, where $\chi(M_\ell^h)$ denotes the characteristic vector corresponding to the edges that the are present in the matching. (See, e.g. [6], for a polynomial-time procedure.) Fixing the parameter α to a suitable value to be specified later, the algorithm does the following.

```
1. for each round $h = 1, \cdots, k$ do
    a. define the $h^{\text{th}}$ matching
       $$\mathbb{P}^h := \begin{cases} \emptyset & \text{with probability } 1 - \frac{1}{\alpha} \\ M_\ell^h & \text{with probability } \frac{\lambda_\ell^h}{\alpha} \end{cases}$$
    b. Probe all edges in $\mathbb{P}^h$ that are safe.
```

We show that this algorithm is a 20-approximation for $\alpha = 10$, which proves Theorem 6.

As before, an edge $(i, j) \in E$ is said to be *safe* iff (a) (i, j) has not been probed earlier, (b) neither i nor j is matched, and (c) neither i nor j has timed out.

Lemma 15 *For any edge $(i, j) \in E$, and at round $h \in [k]$,* $\Pr[\,(i, j)$ *is safe in round* $h\,] \ge 1 - \frac{5}{\alpha}$.

Proof: We will show that the following three statements hold at round h:

i. $\Pr[\,(i, j)$ has probed$\,] \le \frac{1}{\alpha}$.
ii. $\Pr[\,$vertex i is already timed out$\,] \le \frac{1}{\alpha}$.
iii. $\Pr[\,$vertex i is already matched$\,] \le \frac{1}{\alpha}$.

Since $\Pr[\,(i, j)$ is *not* safe in round $h\,]$ is at most

$$\Pr[\,(i,j) \text{ been probed}\,] + \Pr[\,i \text{ matched}\,] + \Pr[\,i \text{ timed out}\,] + \Pr[\,j \text{ matched}\,] + \Pr[\,j \text{ timed out}\,]$$

by the trivial union bound, proving (i)-(iii) will prove the lemma. To prove (i), observe that for any edge $e \in E$ and round g, $\Pr[\,e$ probed in round $g\,] \le \Pr[e \in \mathbb{P}^g\,] = \frac{1}{\alpha}\, y_e^g$, and hence $\Pr[\,(i, j)$ probed before round $h\,] \le \frac{1}{\alpha} \sum_{g<h} y_e^g \le \frac{1}{\alpha}$, where the last inequality uses LP constraint (29).

The proof for (iii) is identical, using the LP constraint (33). The proof for statement (ii) is also similar, though one upper bounds the expected value of the number of times the vertex i is probed (in this step one needs to use the LP constraints (30)) and then uses Markov inequality. ■

Theorem 12 *Setting $\alpha = 10$ gives a 20-approximation for multiple round stochastic matching.*

Proof: Using Lemma 15, we have for any edge $(i,j) \in E$ and round $h \in [k]$,

$$\begin{aligned}
\Pr[\,(i,j) \text{ probed in round } h\,] &= \Pr[\,(i,j) \text{ safe in round } h\,] \cdot \Pr[\,(i,j) \in \mathbb{P}^h \mid (i,j) \text{ safe in round } h\,] \\
&\geq \left(1 - \frac{5}{\alpha}\right) \cdot \Pr[\,(i,j) \in \mathbb{P}^h \mid (i,j) \text{ safe in round } h\,] \\
&= \left(1 - \frac{5}{\alpha}\right) \cdot \frac{y_{ij}^h}{\alpha},
\end{aligned}$$

where the equality follows from the fact that events $(i,j) \in \mathbb{P}^h$ and (i,j) is safe in round h are independent. Thus the expected value accrued by the algorithm is

$$\sum_{e \in E} w_e \cdot \sum_{h=1}^{k} \Pr[\,e \text{ probed in round } h\,] \cdot p_e \geq \frac{1}{\alpha}\left(1 - \frac{5}{\alpha}\right) \cdot \sum_{e \in E} w_e \cdot \sum_{h=1}^{k} y_e^h \cdot p_e,$$

which is $\frac{1}{\alpha}\left(1 - \frac{5}{\alpha}\right)$ times the optimal LP-value. Setting $\alpha = 10$ completes the proof. ■

B Unweighted Stochastic Matching: A Greedy Algorithm

In this section we consider a greedy algorithm for the *unweighted* stochastic matching problem: in this unweighted version, all edges have unit weight, and the goal is to maximize the expected number of matched edges. The greedy algorithm was proposed by Chen et al. [7], and they gave an analysis proving it to be a 4-approximation; however, the proof was fairly involved. Here, we give a significantly simpler analysis showing an approximation guarantee of 5. The greedy algorithm we consider is the following:

```
1. Let $\sigma$ denote the ordering of the edges in $E$ by non-increasing $p_e$-values.
2. Consider the edges $e \in E$ in the order given by $\sigma$
      a. If edge $e$ is safe then probe it, else do not probe $e$.
```

Recall that an edge is safe if neither of its endpoints have been matched or timed out. Note that the expected value of the greedy algorithm is

$$\mathsf{ALG} = \sum_{e \in E} \Pr[\,e \text{ is matched}\,] = \sum_{e \in E} \Pr[\,e \text{ is probed}\,] \cdot p_e.$$

B.1 The Analysis

While the algorithm does not have anything to do with the linear programming relaxation we presented in the previous section, we will use that LP for our analysis. Consider the optimal LP solution (x^*, y^*), and recall that (x^*, y^*) satisfy the conditions (1)-(4). (Alternatively, use the fractional solution $y_e^* :=$

$\Pr[\,e \text{ is probed in the optimal strategy}\,]$ and $x_e^* := \Pr[\,e \text{ is matched in the optimal strategy}\,]$.) For each $e = (i,j) \in E$, define the following three events:

$$\begin{aligned} M_e &:= \text{ either } i \text{ or } j \text{ is matched when } e \text{ is considered in } \sigma, \\ R_e &:= \text{ either } i \text{ or } j \text{ has timed out when } e \text{ is considered in } \sigma, \text{ and} \\ B_e &:= M_e \vee R_e. \end{aligned}$$

By the algorithm, it follows that $\Pr[\,e \text{ is probed}\,] = 1 - \Pr[B_e]$ for all $e \in E$. So,

$$\mathsf{ALG} = \sum_{e\in E} (1 - \Pr[B_e])\, p_e \geq \sum_{e\in E} (1 - \Pr[B_e]) \cdot y_e^*\, p_e \tag{34}$$

The following two lemmas charge the value accrued by the algorithm in two different ways to the optimal LP solution.

Lemma 16 $2\mathsf{ALG} \geq \sum_{g\in E} \Pr[M_g] \cdot y_g^* \cdot p_g$.

Proof: In the greedy algorithm, whenever edge $e = (i,j)$ gets matched, *write* value of $\frac{y_f^* \cdot p_f}{2}$ on each edge $f \in \partial(i) \bigcup \partial(j)$. Note that the total value written when edge $e = (i,j)$ gets matched is at most:

$$\sum_{f\in\partial(i)} \frac{y_f^*\, p_f}{2} + \sum_{f\in\partial(j)} \frac{y_f^*\, p_f}{2} = \frac{1}{2} \sum_{f\in\partial(i)} x_f^* + \frac{1}{2} \sum_{f\in\partial(j)} x_f^* \leq 1,$$

where the inequality follows from (1). Recall that in any possible execution of Greedy, an edge is matched at most once. Thus the expected total value written (on all edges) is at most $\sum_{e\in E} \Pr[\,e \text{ is matched}\,] = \mathsf{ALG}$.

On the other hand, whenever event M_g occurs in the greedy algorithm (at some edge $g = (a,b) \in E$), *read* $\frac{y_g^* \cdot p_g}{2}$ value from g. Consider any outcome where event M_g occurs: it must be that either a or b was already matched (say via edge e); this in turn means that $\frac{y_g^* \cdot p_g}{2}$ value was written on edge g at the time when e got matched (since g is adjacent to e). Thus the value read from an edge (at any point) is at most the value already written on it. Thus the expected total value read from all edges is $\sum_{g\in E} \Pr[M_g] \cdot \frac{y_g^*\, p_g}{2} \leq \mathbb{E}[\text{total value written}] \leq \mathsf{ALG}$. ■

Lemma 17 $2\mathsf{ALG} \geq \sum_{g\in E} \Pr[R_g] \cdot y_g^* \cdot p_g$.

Proof: Consider the execution of the greedy algorithm, with a *value* α_e defined on each edge $e \in E$ (initialized to zero). Whenever an edge $e = (i,j)$ gets probed, do (where σ_e denotes the edges in E that appear after e in σ):

1. For each $f \in \partial(i) \cap \sigma_e$, increase α_f by $\frac{y_f^*\, p_f}{2t_i}$.
2. For each $f \in \partial(j) \cap \sigma_e$, increase α_f by $\frac{y_f^*\, p_f}{2t_j}$.

Let $A := \sum_{e\in E} \alpha_e$. Note that the increase in A when edge $e = (i,j)$ gets probed is:

$$\sum_{f\in\partial(i)\cap\sigma_e} \frac{y_f^*\, p_f}{2t_i} + \sum_{f\in\partial(j)\cap\sigma_e} \frac{y_f^*\, p_f}{2t_j} \leq \frac{p_e}{2} \left(\frac{1}{t_i} \sum_{f\in\partial(i)\cap\sigma_e} y_f^* + \frac{1}{t_j} \sum_{f\in\partial(j)\cap\sigma_e} y_f^* \right) \leq p_e,$$

where for the first inequality we use the greedy property that $p_e \geq p_f$ for all $f \in \sigma_e$ and the second inequality follows from (2). Thus the expected value of A at the end of the greedy algorithm is $\mathbb{E}[A$ at the end of Greedy$] \leq \sum_{e \in E} \Pr[\, e$ is probed$\,] \cdot p_e = \mathsf{ALG}$. (Recall that in any possible execution of Greedy, an edge is probed at most once.)

On the other hand, whenever event R_g occurs in the greedy algorithm (at some edge $g = (a, b) \in E$), *read* the value α_g from g. Consider any outcome where event R_g occurs: it must be that either a or b was already timed out (say vertex a). This means that t_a edges from $\partial(a)$ have already been probed. By the updates to α-values defined above, since g is adjacent to each edge in $\partial(a)$, the current value $\alpha_g \geq t_a \cdot \frac{y_g^* \, p_g}{2\, t_a} = y_g^* \, p_g/2$. So whenever R_g occurs, the value read $\alpha_g \geq y_g^* \, p_g/2$. I.e. the expected total value read is at least $\sum_{g \in E} \Pr[R_g] \cdot \frac{y_g^* \, p_g}{2}$. However, the total value read is at most the value A at the end of the greedy algorithm. This implies that $\sum_{g \in E} \Pr[R_g] \cdot \frac{y_g^* \, p_g}{2} \leq \mathbb{E}[$total value read$] \leq \mathbb{E}[A$ at the end of Greedy$] \leq \mathsf{ALG}$. ■

Proof of Theorem 5: Adding the expressions from Lemmas 16 and 17, we get

$$4\, \mathsf{ALG} \geq \sum_{e \in E} (\Pr[M_e] + \Pr[R_e]) \cdot y_e^* \, p_e \geq \sum_{e \in E} \Pr[B_e] \cdot y_e^* \, p_e,$$

where the second inequality uses the definition $B_e = M_e \vee R_e$. Adding this to (34), we obtain $5\, \mathsf{ALG} \geq \sum_{e \in E} y_e^* \cdot p_e$, which is the optimal LP objective. Thus, the greedy algorithm is a 5-approximation. ■