

# Improved Approximation Algorithms for Stochastic Matching <sup>\*</sup>

Marek Adamczyk<sup>1</sup>, Fabrizio Grandoni<sup>2</sup>, and Joydeep Mukherjee<sup>3</sup>

<sup>1</sup> Department of Computer, Control, and Management Engineering, Sapienza University of Rome, Italy, [adamczyk@dis.uniroma1.it](mailto:adamczyk@dis.uniroma1.it).

<sup>2</sup> IDSIA, University of Lugano, Switzerland, [fabrizio@idsia.ch](mailto:fabrizio@idsia.ch).

<sup>3</sup> Institute of Mathematical Sciences, CIT, India, [joydeepm@imsc.res.in](mailto:joydeepm@imsc.res.in)

**Abstract.** In this paper we consider the *Stochastic Matching* problem, which is motivated by applications in kidney exchange and online dating. We are given an undirected graph in which every edge is assigned a probability of existence and a positive profit, and each node is assigned a positive integer called *timeout*. We know whether an edge exists or not only after probing it. On this random graph we are executing a process, which one-by-one probes the edges and gradually constructs a matching. The process is constrained in two ways: once an edge is taken it cannot be removed from the matching, and the timeout of node  $v$  upper-bounds the number of edges incident to  $v$  that can be probed. The goal is to maximize the expected profit of the constructed matching.

For this problem Bansal et al. [4] provided a 3-approximation algorithm for bipartite graphs, and a 4-approximation for general graphs. In this work we improve the approximation factors to 2.845 and 3.709, respectively.

We also consider an online version of the bipartite case, where one side of the partition arrives node by node, and each time a node  $b$  arrives we have to decide which edges incident to  $b$  we want to probe, and in which order. Here we present a 4.07-approximation, improving on the 7.92-approximation of Bansal et al. [4].

The main technical ingredient in our result is a novel way of probing edges according to a random but non-uniform permutation. Patching this method with an algorithm that works best for large probability edges (plus some additional ideas) leads to our improved approximation factors.

## 1 Introduction

In this paper we consider the *Stochastic Matching* problem, which is motivated by applications in kidney exchange and online dating. Here we are given an

---

<sup>\*</sup> This work was partially done while the first and last authors were visiting IDSIA. The first and second authors were partially supported by the ERC StG project NEWNET no. 279352, and the first author by the ERC StG project PAAL no. 259515. The third author was partially supported by the ISJRP project Mathematical Programming in Parameterized Algorithms.

undirected graph  $G = (V, E)$ . Each edge  $e \in E$  is labeled with an (existence) probability  $p_e \in (0, 1]$  and a weight (or profit)  $w_e > 0$ , and each node  $v \in V$  with a *timeout* (or *patience*)  $t_v \in \mathbb{N}^+$ . An algorithm for this problem probes edges in a possibly adaptive order. Each time an edge is probed, it turns out to be *present* with probability  $p_e$ , in which case it is (irrevocably) included in the matching under construction and provides a profit  $w_e$ . We can probe at most  $t_u$  edges among the set  $\delta(u)$  of edges incident to node  $u$  (independently from whether those edges turn out to be present or absent). Furthermore, when an edge  $e$  is added to the matching, no edge  $f \in \delta(e)$  (i.e., incident on  $e$ ) can be probed in subsequent steps. Our goal is to maximize the expected weight of the constructed matching. Bansal et al. [4] provide an LP-based 3-approximation when  $G$  is bipartite, and via reduction to the bipartite case a 4-approximation for general graphs (see also [3]).

We also consider the *Online Stochastic Matching with Timeouts* problem introduced in [4]. Here we are given in input a bipartite graph  $G = (A \cup B, A \times B)$ , where nodes in  $B$  are *buyer types* and nodes in  $A$  are *items* that we wish to sell. Like in the offline case, edges are labeled with probabilities and profits, and nodes are assigned timeouts. However, in this case timeouts on the item side are assumed to be unbounded. Then a second bipartite graph is constructed in an online fashion. Initially this graph consists of  $A$  only. At each time step one random buyer  $\tilde{b}$  of some type  $b$  is sampled (possibly with repetitions) from a given probability distribution. The edges between  $\tilde{b}$  and  $A$  are copies of the corresponding edges in  $G$ . The online algorithm has to choose at most  $t_b$  unmatched neighbors of  $\tilde{b}$ , and probe those edges in some order until some edge  $a\tilde{b}$  turns out to be present (in which case  $a\tilde{b}$  is added to the matching and we gain the corresponding profit) or all the mentioned edges are probed. This process is repeated  $n$  times, and our goal is to maximize the final total expected profit<sup>4</sup>.

For this problem Bansal et al. [4] present a 7.92-approximation algorithm. In his Ph.D. thesis Li [8] claims an improved 4.008-approximation. However, his analysis contains a mistake [9]. By fixing that, he still achieves a 5.16-approximation ratio improving over [4].

## 1.1 Our Results

Our main result is an approximation algorithm for bipartite Stochastic Matching which improves the 3-approximation of Bansal et al. [4] (see Section 2).

**Theorem 1.** *There is an expected 2.845-approximation algorithm for Stochastic Matching in bipartite graphs.*

Our algorithm for the bipartite case is similar to the one from [4], which works as follows. After solving a proper LP and rounding the solution via a rounding technique from [7], Bansal et al. probe edges in uniform random order. Then they show that every edge  $e$  is probed with probability at least  $x_e \cdot g(p_{max})$ , where

<sup>4</sup> As in [4], we assume that the probability of a buyer type  $b$  is an integer multiple of  $1/n$ .

$x_e$  is the fractional value of  $e$ ,  $p_{max} := \max_{f \in \delta(e)} \{p_f\}$  is the largest probability of any edge incident to  $e$  ( $e$  excluded), and  $g(\cdot)$  is a decreasing function with  $g(1) = 1/3$ .

Our idea is to rather consider edges in a carefully chosen *non-uniform* random order. This way, we are able to show (with a slightly simpler analysis) that each edge  $e$  is probed with probability  $x_e \cdot g(p_e) \geq \frac{1}{3}x_e$ . Observe that we have the same function  $g(\cdot)$  as in [4], but depending on  $p_e$  rather than  $p_{max}$ . In particular, according to our analysis, small probability edges are more likely to be probed than large probability ones (for a given value of  $x_e$ ), regardless of the probabilities of edges incident to  $e$ . Though this approach alone does not directly imply an improved approximation factor, it is not hard to patch it with a simple greedy algorithm that behaves best for large probability edges, and this yields an improved approximation ratio altogether.

We also improve on the 4-approximation for general graphs in [4]. This is achieved by reducing the general case to the bipartite one as in prior work, but we also use a refined LP with blossom inequalities in order to fully exploit our large/small probability patching technique.

**Theorem 2.** *There is an expected 3.709-approximation algorithm for Stochastic Matching in general graphs.*

Similar arguments can also be successfully applied to the online case. By applying our idea of non-uniform permutation of edges we would get a 5.16-approximation (the same as in [8], after correcting the mentioned mistake). However, due to the way edges have to be probed in the online case, we are able to finely control the probability that an edge is probed via *dumping factors*. This allows us to improve the approximation from 5.16 to 4.16. Our idea is similar in spirit to the one used by Ma [10] in his neat 2-approximation algorithm for correlated non-preemptive stochastic knapsack. Further application of the large/small probability trick gives an extra improvement down to 4.07 (see Section 3).

**Theorem 3.** *There is an expected 4.07-approximation algorithm for Online Stochastic Matching with Timeouts.*

## 1.2 Related work

The Stochastic Matching problem falls under the framework of adaptive stochastic problems presented first by Dean et al. [6]. Here the solution is in fact a process, and the optimal one might even require larger than polynomial space to be described.

The Stochastic Matching problem was originally presented by Chen et al. [5] together with applications in kidney exchange and online dating. The authors consider the unweighted version of the problem, and prove that a greedy algorithm is a 4-approximation. Adamczyk [1] later proved that the same algorithm is in fact a 2-approximation, and this result is tight. The greedy algorithm does not provide a good approximation in the weighted case, and all known algorithms

for this case are LP-based. Here, Bansal et al. [4] showed a 3-approximation for the bipartite case. Via a reduction to the bipartite case, Bansal et al. [4] also obtain a 4-approximation algorithm for general graphs (see also [3]).

## 2 Stochastic Matching

### 2.1 Bipartite graphs

Let us denote by  $OPT$  the optimum probing strategy, and let  $\mathbb{E}[OPT]$  denote its expected outcome. Consider the following LP:

$$\begin{aligned} \max \quad & \sum_{e \in E} w_e p_e x_e && \text{(LP-BIP)} \\ \text{s.t.} \quad & \sum_{e \in \delta(u)} p_e x_e \leq 1, && \forall u \in V; && (1) \\ & \sum_{e \in \delta(u)} x_e \leq t_u, && \forall u \in V; && (2) \\ & 0 \leq x_e \leq 1, && \forall e \in E. && (3) \end{aligned}$$

The proof of the following Lemma is already quite standard [3,4,6] — just note that  $x_e = \mathbb{P}[OPT \text{ probes } e]$  is a feasible solution of LP-BIP.

**Lemma 1.** [4] *Let  $LP_{bip}$  be the optimal value of LP-BIP. It holds that  $LP_{bip} \geq \mathbb{E}[OPT]$ .*

Our approach is similar to the one of Bansal et al. [4] (see also Algorithm 1 in the figure). We solve LP-BIP: let  $x = (x_e)_{e \in E}$  be the optimal fractional solution. Then we apply to  $x$  the rounding procedure by Gandhi et al. [7], which we shall call just GKPS. Let  $\hat{E}$  be the set of rounded edges, and let  $\hat{x}_e = 1$  if  $e \in \hat{E}$  and  $\hat{x}_e = 0$  otherwise. GKPS guarantees the following properties of the rounded solution:

1. (Marginal distribution) For any  $e \in E$ ,  $\mathbb{P}[\hat{x}_e = 1] = x_e$ .
2. (Degree preservation) For any  $v \in V$ ,  $\sum_{e \in \delta(v)} \hat{x}_e \leq \lceil \sum_{e \in \delta(v)} x_e \rceil \leq t_v$ .
3. (Negative correlation) For any  $v \in V$ , any subset  $S \subseteq \delta(v)$  of edges incident to  $v$ , and any  $b \in \{0, 1\}$ , it holds that  $\mathbb{P}[\wedge_{e \in S} (\hat{x}_e = b)] \leq \prod_{e \in S} \mathbb{P}[\hat{x}_e = b]$ .

Our algorithm sorts the edges in  $\hat{E}$  according to a random permutation and probes each edge  $e \in \hat{E}$  according to that order, but provided that the endpoints of  $e$  are not matched already. It is important to notice that, by the degree preservation property, in  $\hat{E}$  there are at most  $t_v$  edges incident to each node  $v$ . Hence, the timeout constraint of  $v$  is respected even if the algorithm probes all the edges in  $\delta(v) \cap \hat{E}$ .

Our algorithm differs from [4] and subsequent work in the way edges are randomly ordered. Prior work exploits a random uniform order on  $\hat{E}$ . We rather use the following, more complex strategy. For each  $e \in \hat{E}$  we draw a random

---

**Algorithm 1** Approximation algorithm for bipartite Stochastic Matching.
 

---

1. Let  $(x_e)_{e \in E}$  be the solution to LP-BIP.
  2. Round the solution  $(x_e)_{e \in E}$  with GKPS; let  $(\hat{x}_e)_{e \in E}$  be the rounded 0-1 solution, and  $\hat{E} = \{e \in E | \hat{x}_e = 1\}$ .
  3. For every  $e \in \hat{E}$ , sample a random variable  $Y_e$  distributed as  $\mathbb{P}[Y_e \leq y] = \frac{1 - e^{-yp_e}}{p_e}$ .
  4. For every  $e \in \hat{E}$  in increasing order of  $Y_e$ :
    - (a) If no edge  $f \in \hat{\delta}(e) := \delta(e) \cap \hat{E}$  is yet taken, then probe edge  $e$
- 

variable  $Y_e$  distributed on the interval  $\left[0, \frac{1}{p_e} \ln \frac{1}{1-p_e}\right]$  according to the following cumulative distribution:  $\mathbb{P}[Y_e \leq y] = \frac{1}{p_e} (1 - e^{-p_e y})$ . Observe that the density function of  $Y_e$  in this interval is  $e^{-yp_e}$  (and zero otherwise). Edges of  $\hat{E}$  are sorted in increasing order of the  $Y_e$ 's, and they are probed according to that order. We next let  $Y = (Y_e)_{e \in \hat{E}}$ .

Define  $\hat{\delta}(v) := \delta(v) \cap \hat{E}$ . We say that an edge  $e \in \hat{E}$  is *safe* if, at the time we consider  $e$  for probing, no other edge  $f \in \hat{\delta}(e)$  is already taken into the matching. Note that the algorithm can probe  $e$  only in that case, and if we do probe  $e$ , it is added to the matching with probability  $p_e$ .

The main ingredient of our analysis is the following lower-bound on the probability that an arbitrary edge  $e$  is safe.

**Lemma 2.** *For every edge  $e$  it holds that  $\mathbb{P}[e \text{ is safe} | e \in \hat{E}] \geq g(p_e)$ , where*

$$g(p) := \frac{1}{2+p} \left( 1 - \exp \left( - (2+p) \frac{1}{p} \ln \frac{1}{1-p} \right) \right).$$

*Proof.* In the worst case every edge  $f \in \hat{\delta}(e)$  that is before  $e$  in the ordering can be probed, and each of these probes has to fail for  $e$  to be safe. Thus

$$\mathbb{P}[e \text{ is safe} | e \in \hat{E}] \geq \mathbb{E}_{\hat{E} \setminus e, Y} \left[ \prod_{f \in \hat{\delta}(e): Y_f < Y_e} (1 - p_f) \middle| e \in \hat{E} \right].$$

Now we take expectation on  $Y$  only, and using the fact that the variables  $Y_f$  are independent, we can write the latter expectation as

$$\mathbb{E}_{\hat{E} \setminus e} \left[ \int_0^{\frac{1}{p_e} \ln \frac{1}{1-p_e}} \left( \prod_{f \in \hat{\delta}(e)} (\mathbb{P}[Y_f \leq y] (1 - p_f) + \mathbb{P}[Y_f > y]) \right) e^{-p_e \cdot y} dy \middle| e \in \hat{E} \right]. \quad (4)$$

Observe that  $\mathbb{P}[Y_f \leq y] (1 - p_f) + \mathbb{P}[Y_f > y] = 1 - p_f \mathbb{P}[Y_f \leq y]$ . When  $y > \frac{1}{p_f} \ln \frac{1}{1-p_f}$ , then  $\mathbb{P}[Y_f \leq y] = 1$ , and moreover,  $\frac{1}{p_f} (1 - e^{-p_f \cdot y})$  is an increasing function of  $y$ . Thus we can upper-bound  $\mathbb{P}[Y_f \leq y]$  by  $\frac{1}{p_f} (1 - e^{-p_f \cdot y})$  for any

$y \in [0, \infty]$ , and obtain that  $1 - p_f \mathbb{P}[Y_f \leq y] \geq 1 - p_f \frac{1}{p_f} (1 - e^{-p_f \cdot y}) = e^{-p_f \cdot y}$ . Thus (4) can be lower bounded by

$$\begin{aligned} & \mathbb{E}_{\hat{E} \setminus e} \left[ \int_0^{\frac{1}{p_e} \ln \frac{1}{1-p_e}} e^{-\sum_{f \in \delta(e)} p_f \cdot y - p_e \cdot y} dy \middle| e \in \hat{E} \right] \\ &= \mathbb{E}_{\hat{E} \setminus e} \left[ \frac{1}{\sum_{f \in \delta(e)} p_f + p_e} \left( 1 - e^{-\left(\sum_{f \in \delta(e)} p_f + p_e\right) \frac{1}{p_e} \ln \frac{1}{1-p_e}} \right) \middle| e \in \hat{E} \right]. \end{aligned}$$

From the negative correlation and marginal distribution properties we know that  $\mathbb{E}_{\hat{E} \setminus e} [\hat{x}_f | e \in \hat{E}] \leq \mathbb{E}_{\hat{E} \setminus e} [\hat{x}_f] = x_f$  for every  $f \in \delta(e)$ , and therefore  $\mathbb{E}_{\hat{E} \setminus e} \left[ \sum_{f \in \delta(e)} p_f \middle| e \in \hat{E} \right] \leq \sum_{f \in \delta(e)} p_f x_f \leq 2$ , where the last inequality follows from the LP constraints. Consider function  $f(x) := \frac{1}{x+p_e} \left( 1 - e^{-(x+p_e) \frac{1}{p_e} \ln \frac{1}{1-p_e}} \right)$ . This function is decreasing and convex. From Jensen's inequality we know that  $\mathbb{E}[f(x)] \geq f(\mathbb{E}[x])$ . Thus

$$\begin{aligned} \mathbb{E}_{\hat{E} \setminus e} \left[ f \left( \sum_{f \in \delta(e)} p_f \right) \middle| e \in \hat{E} \right] &\geq f \left( \mathbb{E}_{\hat{E} \setminus e} \left[ \sum_{f \in \delta(e)} p_f \middle| e \in \hat{E} \right] \right) \\ &\geq f(2) = \frac{1}{2+p_e} \left( 1 - e^{-(2+p_e) \frac{1}{p_e} \ln \frac{1}{1-p_e}} \right) = g(p_e). \quad \square \end{aligned}$$

From Lemma 2 and the marginal distribution property, the expected contribution of edge  $e$  to the profit of the solution is

$$w_e p_e \cdot \mathbb{P} \left[ e \in \hat{E} \right] \cdot \mathbb{P} \left[ e \text{ is safe} \middle| e \in \hat{E} \right] \geq w_e p_e x_e \cdot g(p_e) \geq w_e p_e x_e \cdot g(1) = \frac{1}{3} w_e p_e x_e.$$

Therefore, our analysis implies a 3 approximation, matching the result in [4]. However, by playing with the probabilities appropriately we can do better.

*Patching with Greedy.* We next describe an improved approximation algorithm, based on the patching of the above algorithm with a simple greedy one. Let  $\delta \in (0, 1)$  be a parameter to be fixed later. We define  $E_{large}$  as the (*large*) edges with  $p_e \geq \delta$ , and let  $E_{small}$  be the remaining (*small*) edges. Recall that  $LP_{bip}$  denotes the optimal value of LP-BIP. Let also  $LP_{large}$  and  $LP_{small}$  be the fraction of  $LP_{bip}$  due to large and small edges, respectively; i.e.,  $LP_{large} = \sum_{e \in E_{large}} w_e p_e x_e$  and  $LP_{small} = LP_{bip} - LP_{large}$ . Define  $\gamma \in [0, 1]$  such that  $\gamma LP_{bip} = LP_{large}$ . By refining the above analysis, we obtain the following result.

**Lemma 3.** *Algorithm 1 has expected approximation ratio  $\frac{1}{3}\gamma + g(\delta)(1 - \gamma)$ .*

*Proof.* The expected profit of the algorithm is at least:

$$\begin{aligned} & \sum_{e \in E} w_e p_e x_e \cdot g(p_e) \geq \sum_{e \in E_{large}} w_e p_e x_e \cdot g(1) + \sum_{e \in E_{small}} w_e p_e x_e \cdot g(\delta) \\ &= \frac{1}{3} LP_{large} + g(\delta) LP_{small} = \left( \frac{1}{3}\gamma + g(\delta)(1 - \gamma) \right) LP_{bip}. \quad \square \end{aligned}$$

Consider the following greedy algorithm. Compute a maximum weight matching  $M_{grd}$  in  $G$  with respect to edge weights  $w_e p_e$ , and probe the edges of  $M_{grd}$  in any order. Note that the timeout constraints are satisfied since we probe at most one edge incident to each node (and timeouts are strictly positive by definition and w.l.o.g.).

**Lemma 4.** *The greedy algorithm has expected approximation ratio  $\delta\gamma$ .*

*Proof.* It is sufficient to show that the expected profit of the obtained solution is at least  $\delta \cdot LP_{large}$ . Let  $x = (x_e)_{e \in E}$  be the optimal solution to LP-BIP. Consider the solution  $x' = (x'_e)_{e \in E}$  that is obtained from  $x$  by setting to zero all the variables corresponding to edges in  $E_{small}$ , and by multiplying all the remaining variables by  $\delta$ . Since  $p_e \geq \delta$  for all  $e \in E_{large}$ ,  $x'$  is a feasible fractional solution to the following matching LP:

$$\begin{aligned} \max \quad & \sum_{e \in E} w_e p_e z_e && \text{(LP-MATCH)} \\ \text{s.t.} \quad & \sum_{e \in \delta(u)} z_e \leq 1, && \forall u \in V; \\ & 0 \leq z_e \leq 1, && \forall e \in E. \end{aligned} \quad (5)$$

The value of  $x'$  in the above LP is  $\delta \cdot LP_{large}$  by construction. Let  $LP_{match}$  be the optimal profit of LP-MATCH. Then  $LP_{match} \geq \delta \cdot LP_{large}$ . Given that the graph is bipartite, LP-MATCH defines the matching polyhedron, and we can find an integral optimal solution to it. But such a solution is exactly a maximum weight matching according to weights  $w_e p_e$ , i.e.  $\sum_{e \in M_{grd}} w_e p_e = LP_{match}$ . The claim follows since the expected profit of the greedy algorithm is precisely the weight of  $M_{grd}$ .  $\square$

The overall algorithm, for a given  $\delta$ , simply computes the value of  $\gamma$ , and runs the greedy algorithm if  $\gamma\delta \geq (\frac{1}{3}\gamma + g(\delta)(1 - \gamma))$ , and Algorithm 1 otherwise<sup>5</sup>.

The approximation factor is given by  $\max\{\frac{\gamma}{3} + (1 - \gamma)g(\delta), \gamma\delta\}$ , and the worst case is achieved when the two quantities are equal, i.e., for  $\gamma = \frac{g(\delta)}{\delta + g(\delta) - \frac{1}{3}}$ , yielding an approximation ratio of  $\frac{\delta \cdot g(\delta)}{\delta + g(\delta) - \frac{1}{3}}$ . Maximizing (numerically) the latter function in  $\delta$  gives  $\delta = 0.6022$ , and the final 2.845-approximation ratio claimed in Theorem 1.

## 2.2 General graphs

For general graphs, we consider the linear program LP-GEN which is obtained from LP-BIP by adding the following *blossom inequalities*:

$$\sum_{e \in E(W)} p_e x_e \leq \frac{|W| - 1}{2} \quad \forall W \subseteq V, |W| \text{ odd.} \quad (6)$$

<sup>5</sup> Note that we cannot run both algorithms, and take the best solution.

Here  $E(W)$  is the subset of edges with both endpoints in  $W$ . We remark that, using standard tools from matching theory, we can solve LP-GEN in polynomial time despite its exponential number of constraints; see the book of Schrijver for details [11]. Also in this case  $x_e = \mathbb{P}[\text{OPT probes } e]$  is a feasible solution of LP-GEN, hence the analogue of Lemma 1 still holds.

Our Stochastic Matching algorithm for the case of a general graph  $G = (V, E)$  works via a reduction to the bipartite case. First we solve LP-GEN; let  $x = (x_e)_{e \in E}$  be the optimal fractional solution. Second we randomly split the nodes  $V$  into two sets  $A$  and  $B$ , with  $E_{AB}$  being the set of edges between them. On the bipartite graph  $(A \cup B, E_{AB})$  we apply the algorithm for the bipartite case, but using the fractional solution  $(x_e)_{e \in E_{AB}}$  induced by LP-GEN rather than solving LP-BIP. Note that  $(x_e)_{e \in E_{AB}}$  is a feasible solution to LP-BIP for the bipartite graph  $(A \cup B, E_{AB})$ .

The analysis differs only in two points w.r.t. the one for the bipartite case. First, with  $\hat{E}_{AB}$  being the subset of edges of  $E_{AB}$  that were rounded to 1, we have now that  $\mathbb{P}[e \in \hat{E}_{AB}] = \mathbb{P}[e \in E_{AB}] \cdot \mathbb{P}[e \in \hat{E}_{AB} | e \in E_{AB}] = \frac{1}{2}x_e$ . Second, but for the same reason, using again the negative correlation and marginal distribution properties, we have

$$\mathbb{E} \left[ \sum_{f \in \hat{\delta}(e)} p_f \mid e \in \hat{E}_{AB} \right] \leq \sum_{f \in \delta(e)} p_f \mathbb{P}[f \in \hat{E}_{AB}] = \sum_{f \in \delta(e)} \frac{p_f x_f}{2} \leq \frac{2 - 2p_e x_e}{2} \leq 1.$$

Repeating the steps of the proof of Lemma 2 and including the above inequality we get the following.

**Lemma 5.** *For every edge  $e$  it holds that  $\mathbb{P}[e \text{ is safe} | e \in \hat{E}_{AB}] \geq h(p_e)$ , where*

$$h(p) := \frac{1}{1+p} \left( 1 - \exp \left( - (1+p) \frac{1}{p} \ln \frac{1}{1-p} \right) \right).$$

Since  $h(p_e) \geq h(1) = \frac{1}{2}$ , we directly obtain a 4-approximation which matches the result in [4]. Similarly to the bipartite case, we can patch this result with the simple greedy algorithm (which is exactly the same in the general graph case). For a given parameter  $\delta \in [0, 1]$ , let us define  $\gamma$  analogously to the bipartite case. Similarly to the proof of Lemma 3, one obtains that the above algorithm has approximation factor  $\frac{\gamma}{4} + \frac{1-\gamma}{2}h(\delta)$ . Similarly to the proof of Lemma 4, the greedy algorithm has approximation ratio  $\gamma\delta$  (here we exploit the blossom inequalities that guarantee the integrality of the matching polyhedron). We can conclude similarly that in the worst case  $\gamma = \frac{h(\delta)}{2\delta+h(\delta)-1/2}$ , yielding an approximation ratio of  $\frac{\delta \cdot h(\delta)}{2\delta+h(\delta)-1/2}$ . Maximizing (numerically) this function over  $\delta$  gives, for  $\delta = 0.5580$ , the 3.709 approximation ratio claimed in Theorem 2.

### 3 Online Stochastic Matching with Timeouts

Let  $G = (A \cup B, A \times B)$  be the input graph, with items  $A$  and buyer types  $B$ . We use the same notation for edge probabilities, edge profits, and timeouts as

in Stochastic Matching. Following [4], we can assume w.l.o.g. that each buyer type is sampled uniformly with probability  $1/n$ . Consider the following linear program:

$$\begin{aligned}
\max \quad & \sum_{a \in A, b \in B} w_{ab} p_{ab} x_{ab} && \text{(LP-ONL)} \\
\text{s.t.} \quad & \sum_{b \in B} p_{ab} x_{ab} \leq 1, && \forall a \in A \\
& \sum_{a \in A} p_{ab} x_{ab} \leq 1, && \forall b \in B \\
& \sum_{a \in A} x_{ab} \leq t_b, && \forall b \in B \\
& 0 \leq x_{ab} \leq 1, && \forall ab \in E.
\end{aligned}$$

The above LP models a bipartite Stochastic Matching instance where one side of the bipartition contains exactly one buyer per buyer type. In contrast, in the online case several buyers of the same buyer type (or none at all) can arrive, and the optimal strategy can allow many buyers of the same type to probe edges. Still, that is not a problem since the following lemma from [4] allows us just to look at the graph of buyer types and not at the actual realized buyers.

**Lemma 6.** ([4], Lemmas 9 and 11) *Let  $\mathbb{E}[OPT]$  be the expected profit of the optimal online algorithm for the problem. Let  $LP_{onl}$  be the optimal value of LP-ONL. It holds that  $\mathbb{E}[OPT] \leq LP_{onl}$ .*

We will devise an algorithm whose expected outcome is at least  $\frac{1}{4.07} \cdot LP_{onl}$ , and then Theorem 3 follows from Lemma 6.

*The algorithm.* We initially solve LP-ONL and let  $(x_{ab})_{ab \in A \times B}$  be the optimal fractional solution. Then buyers arrive. When a buyer of type  $b$  is sampled, then 1) if a buyer of the same type  $b$  was already sampled before we simply discard her, do nothing, and wait for another buyer to arrive, 2) if it is the first buyer of type  $b$ , then we execute the following *subroutine for buyers*. Since we take action only when the first buyer of type  $b$  comes, we shall denote such a buyer simply by  $b$ , as it will not cause any confusion.

*Subroutine for buyers.* Let us consider the step of the online algorithm in which the first buyer of type  $b$  arrived, if any. Let  $A_b$  be the items that are still available when  $b$  arrives. Our subroutine will probe a subset of at most  $t_b$  edges  $ab$ ,  $a \in A_b$ . Consider the vector  $(x_{ab})_{a \in A_b}$ . Observe that it satisfies the constraints  $\sum_{a \in A_b} p_{ab} x_{ab} \leq 1$  and  $\sum_{a \in A_b} x_{ab} \leq t_b$ . Again using GKPS, we round this vector in order to get  $(\hat{x}_{ab})_{a \in A_b}$  with  $\hat{x}_{ab} \in \{0, 1\}$ , and satisfying the marginal distribution, degree preservation, and negative correlation properties<sup>6</sup>. Let  $\hat{A}_b$  be the

<sup>6</sup> Actually in this case we have a bipartite graph where one side has only one vertex, and here GKPS reduces to Srinivasan's rounding procedure for level-sets [12].

set of items  $a$  such that  $\hat{x}_{ab} = 1$ . For each  $ab$ ,  $a \in \hat{A}_b$ , we independently draw a random variable  $Y_{ab}$  with distribution:  $\mathbb{P}[Y_{ab} < y] = \frac{1}{p_{ab}} (1 - \exp(-p_{ab} \cdot y))$  for  $y \in \left[0, \frac{1}{p_{ab}} \ln \frac{1}{1-p_{ab}}\right]$ . Let  $Y = (Y_{ab})_{a \in \hat{A}_b}$ .

Next we consider items of  $\hat{A}_b$  in increasing order of  $Y_{ab}$ . Let  $\alpha_{ab} \in [\frac{1}{2}, 1]$  be a *dumping factor* that we will define later. With probability  $\alpha_{ab}$  we probe edge  $ab$  and as usual we stop the process (of probing edges incident to  $b$ ) if  $ab$  is present. Otherwise (with probability  $1 - \alpha_{ab}$ ) we *simulate* the probe of  $ab$ , meaning that with probability  $p_{ab}$  we stop the process anyway — like if edge  $ab$  were probed and turned out to be present. Note that we do not get any profit from the latter simulation since we do not really probe  $ab$ .

*Dumping factors.* It remains to define the dumping factors. For a given edge  $ab$ , let

$$\beta_{ab} := \mathbb{E}_{\hat{A}_b \setminus a, Y} \left[ \prod_{a' \in A_b: Y_{a'b} < Y_{ab}} (1 - p_{a'b}) \middle| a \in \hat{A}_b \right].$$

Using the inequality  $\sum_{a \in A_b} p_{ab} x_{ab} \leq 1$ , by repeating the analysis from Section 2 we can show that

$$\beta_{ab} \geq h(p_{ab}) = \frac{1}{1 + p_{ab}} \left( 1 - \exp \left( - (1 + p_{ab}) \frac{1}{p_{ab}} \ln \frac{1}{1 - p_{ab}} \right) \right) \geq \frac{1}{2}.$$

Let us assume for the sake of simplicity that we are able to compute  $\beta_{ab}$  exactly. We set  $\alpha_{ab} = \frac{1}{2\beta_{ab}}$ . Note that  $\alpha_{ab}$  is well defined since  $\beta_{ab} \in [1/2, 1]$ .

*Analysis.* Let us denote by  $\mathcal{A}_b$  the event that at least one buyer of type  $b$  arrives. The probability that an edge  $ab$  is probed can be expressed as:

$$\mathbb{P}[\mathcal{A}_b] \cdot \mathbb{P}[\text{no } b' \text{ takes } a \text{ before } b | \mathcal{A}_b] \cdot \mathbb{P}[b \text{ probes } a | \mathcal{A}_b \wedge a \text{ is not yet taken}].$$

The probability that  $b$  arrives is  $\mathbb{P}[\mathcal{A}_b] = 1 - \left(1 - \frac{1}{n}\right)^n \geq 1 - \frac{1}{e}$ . We shall show first that

$$\mathbb{P}[b \text{ probes } a | \mathcal{A}_b \wedge a \text{ is not yet taken}]$$

is exactly  $\frac{1}{2}x_{ab}$ , and later we shall show that  $\mathbb{P}[\text{no } b' \text{ takes } a \text{ before } b | \mathcal{A}_b]$  is at least  $\frac{1}{1 + \frac{1}{2}(1 - \frac{1}{e})}$ . This will yield that the probability that  $ab$  is probed is at least

$$\left(1 - \frac{1}{e}\right) \frac{1}{1 + \frac{1}{2}(1 - \frac{1}{e})} \cdot \frac{1}{2}x_{ab} = \frac{e-1}{3e-1}x_{ab} > \frac{1}{4.16}x_{ab}.$$

Consider the probability that some edge  $a'b$  appearing before  $ab$  in the random order *blocks* edge  $ab$ , meaning that  $ab$  is not probed because of  $a'b$ . Observe that each such  $a'b$  is indeed considered for probing in the online model, and the probability that  $a'b$  blocks  $ab$  is therefore  $\alpha_{a'b}p_{a'b} + (1 - \alpha_{a'b})p_{a'b} = p_{a'b}$ . We can conclude that the probability that  $ab$  is not blocked is exactly  $\beta_{ab}$ .

Due to the dumping factor  $\alpha_{ab}$ , the probability that we actually probe edge  $ab \in \hat{A}_b$  is exactly  $\alpha_{ab} \cdot \beta_{ab} = \frac{1}{2}$ . Recall that  $\mathbb{P}[a \in \hat{A}_b] = x_{ab}$  by the marginal distribution property. Altogether

$$\mathbb{P}[b \text{ probes } a | \mathcal{A}_b \wedge a \text{ is not yet taken}] = \frac{1}{2}x_{ab}. \quad (7)$$

Next let us condition on the event that buyer  $b$  arrived, and let us lower bound the probability that  $ab$  is not blocked on the  $a$ 's side in such a step, i.e., that no other buyer has taken  $a$  already. The buyers, who are first occurrences of their type, arrive uniformly at random. Therefore, we can analyze the process of their arrivals as if it was constructed by the following procedure: every buyer  $b'$  is given an independent random variable  $Y_{b'}$  distributed exponentially on  $[0, \infty]$ , i.e.,  $\mathbb{P}[Y_{b'} < y] = 1 - e^{-y}$ ; buyers arrive in increasing order of their variables  $Y_{b'}$ . Once buyer  $b'$  arrives, it probes edge  $ab'$  with probability (exactly)  $\alpha_{ab'}\beta_{ab'}x_{ab'} = \frac{1}{2}x_{ab'}$  — these probabilities are independent among different buyers. Thus, conditioning on the fact that  $b$  arrives, we obtain the following expression for the probability that  $a$  is safe at the moment when  $b$  arrives:

$$\begin{aligned} & \mathbb{P}[\text{no } b' \text{ takes } a \text{ before } b | \mathcal{A}_b] \\ & \geq \mathbb{E} \left[ \prod_{b' \in B \setminus b: Y_{b'} < Y_b} (1 - \mathbb{P}[\mathcal{A}_{b'} | \mathcal{A}_b] \mathbb{P}[b' \text{ probes } ab' | \mathcal{A}_{b'}] p_{ab'}) \middle| \mathcal{A}_b \right] \\ & = \int_0^\infty \prod_{b' \in B \setminus b} (1 - \mathbb{P}[\mathcal{A}_{b'} | \mathcal{A}_b] \cdot \mathbb{P}[Y_{b'} < y | \mathcal{A}_{b'}] \cdot \mathbb{P}[b' \text{ probes } ab' | \mathcal{A}_{b'}] p_{ab'}) e^{-y} dy. \end{aligned}$$

Now let us upper-bound each of the probability factors in the above product. First of all  $\mathbb{P}[\mathcal{A}_{b'} | \mathcal{A}_b] = 1 - (1 - \frac{1}{n})^{n-1} \leq 1 - \frac{1}{e}$ . Second,  $\mathbb{P}[Y_{b'} < y | \mathcal{A}_{b'}] = 1 - e^{-y}$  just by definition<sup>7</sup>. Third, from (7) we have that  $\mathbb{P}[b' \text{ probes } ab' | \mathcal{A}_{b'}] = \frac{x_{ab'}}{2}$ .

Thus the above integral can be lower bounded by

$$\begin{aligned} & \int_0^\infty \prod_{b' \in B \setminus b} \left( 1 - \left( 1 - \frac{1}{e} \right) (1 - e^{-y}) \cdot \frac{1}{2} x_{ab'} \cdot p_{ab'} \right) e^{-y} dy \\ & \geq \int_0^\infty \prod_{b' \in B \setminus b} \exp \left( - \left( 1 - \frac{1}{e} \right) \frac{1}{2} x_{ab'} \cdot p_{ab'} \cdot y \right) e^{-y} dy \\ & = \frac{1}{1 + \left( 1 - \frac{1}{e} \right) \frac{1}{2} \left( \sum_{b' \in B \setminus b} p_{ab'} \cdot x_{ab'} \right)} \geq \frac{1}{1 + \frac{1}{2} \left( 1 - \frac{1}{e} \right)} = \frac{2e}{3e - 1}. \end{aligned}$$

Above in the first inequality we used the fact that  $1 - c(1 - e^{-y}) \geq e^{-cy}$  for  $c \in [0, 1]$  and any  $y \in \mathbb{R}$ : here  $c = \left( 1 - \frac{1}{e} \right) \frac{1}{2} x_{ab'} \cdot p_{ab'}$ . In the first equality we used  $\int_0^\infty e^{-ax} dx = \frac{1}{a}$ . In the last inequality we used the LP constraint  $\sum_{b' \in B \setminus b} p_{ab'} \cdot x_{ab'} \leq 1$ .

<sup>7</sup> The  $\mathcal{A}_{b'}$  event in the condition simply indicates that  $Y_{b'}$  was drawn.

Altogether, as anticipated earlier,

$$\mathbb{P}[ab \text{ is probed}] \geq \left(1 - \frac{1}{e}\right) \frac{x_{ab}}{2} \cdot \frac{2e}{3e-1} = x_{ab} \cdot \frac{e-1}{3e-1} > \frac{1}{4.16} \cdot x_{ab}.$$

*Technical details.* Recall that we assumed that we are able to compute the quantities  $\beta_{ab}$ , hence the desired dumping factors  $\alpha_{ab}$ . Indeed, for our goals it is sufficient to estimate them with large enough probability and with sufficiently good accuracy. This can be done by simulating the underlying random process a polynomial number of times. This way the above probability can be lower bounded by  $(\frac{e-1}{3e-1} + \varepsilon)x_e$  for an arbitrarily small constant  $\varepsilon > 0$ . In particular, by choosing a small enough  $\varepsilon$  the factor 4.16 is still guaranteed. The approximation factor can be further reduced to 4.07 via the technique based on small and big probabilities that we introduced before. The omitted technical details will be given in the full version of the paper (see also [2]). Theorem 3 follows.

## References

1. M. Adamczyk. Improved analysis of the greedy algorithm for stochastic matching. *Information Processing Letters*, 111:731–737, 2011.
2. M. Adamczyk, F. Grandoni, and J. Mukherjee. Improved approximation algorithms for stochastic matching. *CoRR*, abs/1505.01439, 2015.
3. M. Adamczyk, M. Sviridenko, and J. Ward. Submodular stochastic probing on matroids. In *STACS 2014*, pages 29–40.
4. 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. *Algorithmica*, 63(4):733–762, 2012.
5. N. Chen, N. Immorlica, A. R. Karlin, M. Mahdian, and A. Rudra. Approximating matches made in heaven. In *ICALP 2009*, pages 266–278.
6. 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.
7. R. Gandhi, S. Khuller, S. Parthasarathy, and A. Srinivasan. Dependent rounding and its applications to approximation algorithms. *Journal of the ACM*, 53(3):324–360, 2006.
8. J. Li. *Decision making under uncertainty*. PhD thesis, University of Maryland, 2011.
9. J. Li. Private communication, 2015.
10. W. Ma. Improvements and generalizations of stochastic knapsack and multi-armed bandit approximation algorithms: Extended abstract. In *SODA 2014*, pages 1154–1163.
11. A. Schrijver. *Combinatorial Optimization - Polyhedra and Efficiency*. Springer, 2003.
12. A. Srinivasan. Distributions on level-sets with applications to approximation algorithms. In *FOCS 2001*, pages 588–597.