The Stochastic Matching Problem with (Very) Few Queries

SEPEHR ASSADI, SANJEEV KHANNA, and YANG LI, University of Pennsylvania

Motivated by an application in *kidney exchange*, we study the following *stochastic matching* problem: We are given a graph $G(V, E)$ (not necessarily bipartite), where each edge in E is *realized* with some constant probability $p > 0$, and the goal is to find a maximum matching in the realized graph. An algorithm in this setting is allowed to make *queries* to edges in E to determine whether or not they are realized.

We design an *adaptive* algorithm for this problem that, for any graph G , computes a $(1 - \epsilon)$ -approximate maximum matching in the realized graph G_p with high probability, while making $O(\frac{\log(1/\epsilon p)}{\epsilon p})$ queries per vertex, where the edges to query are chosen adaptively in $O(\frac{\log(1/\epsilon p)}{\epsilon p})$ rounds. We further present a *non-adaptive* algorithm that makes $O(\frac{\log(1/\epsilon p)}{\epsilon p})$ queries per vertex and computes a $(\frac{1}{2} - \epsilon)$ -approximate maximum matching in G_p with high probability.

Both our adaptive and non-adaptive algorithms achieve the same approximation factor as the previous best algorithms of Blum et al. (EC 2015) for this problem, while requiring an *exponentially smaller* number of per-vertex queries (and rounds of adaptive queries for the adaptive algorithm). Our results settle an open problem raised by Blum et al. by achieving only a polynomial dependency on both ϵ and p . Moreover, the approximation guarantee of our algorithms is *instance-wise* as opposed to only being competitive in *expectation* as is the case for Blum et al. This is of particular relevance to the key application of stochastic matching in kidney exchange. We obtain our results via two main techniques—namely, *matching-covers* and *vertex sparsification*—that may be of independent interest.

CCS Concepts: • **Theory of computation** → **Graph algorithms analysis**;

Additional Key Words and Phrases: Stochastic matching, kidney exchange

ACM Reference format:

Sepehr Assadi, Sanjeev Khanna, and Yang Li. 2019. The Stochastic Matching Problem with (Very) Few Queries. *ACM Trans. Econ. Comput.* 7, 3, Article 16 (September 2019), 19 pages.
<https://doi.org/10.1145/3355903>

1 INTRODUCTION

We study the problem of finding a maximum matching in presence of *uncertainty* in the input graph. Specifically, we consider the *stochastic matching* problem in which we are given an undirected graph $G(V, E)$ where each edge $e \in E$ is *realized* with some constant probability $p > 0$ and the goal is to find a maximum matching in the realized graph. To find a large matching, an algorithm is allowed to *query* edges in E to determine whether or not they are realized.

This work was supported in part by National Science Foundation grants CCF-1116961, CCF-1552909, and IIS-1447470.

Authors' addresses: S. Assadi, S. Khanna, and Y. Li; emails: {sassadi, sanjeev, yangli2}@cis.upenn.edu.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2019 Copyright held by the owner/author(s). Publication rights licensed to ACM.

2167-8375/2019/09-ART16 \$15.00

<https://doi.org/10.1145/3355903>

There is a trivial solution for the stochastic matching problem: Simply query all edges in E and compute a maximum matching over the realized graph. However, in many applications, determining whether or not an edge is realized could be both costly and time-consuming (we will elaborate more on this in the next section). Consequently, to minimize cost, it is preferable that an algorithm queries as few edges as possible, and to minimize the time consumed in the query process, an algorithm should have only a few rounds of *adaptivity* whenever possible, meaning that it is preferable if the decision of “which edges to query next” does not depend on the outcome of previous queries, since in this case, many edges can be queried in parallel. We refer to the number of rounds of parallel querying as the *degree of adaptivity* of the algorithm. The formal definition of this model is presented in Section 1.2.

1.1 Kidney Exchange

A canonical and arguably the most important application of the stochastic matching problem appears in *kidney exchange*. Typically, organ donation comes from deceased donors, since many organs cannot be harvested from a living donor without jeopardizing the health of the donor. Fortunately, kidney is an exception, and a healthy living donor (typically a family member) is able to donate one of his/her two kidneys without facing major life-threatening consequences.

However, patients that cannot be matched to a living donor (due to reasons like incompatible blood type, etc.) still have to deal with long waiting periods on the deceased donor list, and as of January 2016, over 100K patients await kidney transplants, and the median wait time for an individual’s first kidney transplant is 3.6 years [22]. This leads to the idea of kidney exchange (e.g., References [28–30]): Two patients *swap* their incompatible donors to each get a compatible donor.

This setting can be modeled as a maximum matching problem as follows: Create a graph $G(V, E)$ where each patient-donor pair is a vertex and there is an edge between two vertices iff the two patient-donor pairs can perform a kidney exchange. Consequently, a maximum matching in this graph identifies the maximum number of patient-donor pairs that are able to perform an exchange.

To construct a graph for kidney exchange, typically, we only have access to the medical record of the patients and the donors, which contains information like blood type, tissue type, and so on. Therefore, $O(|V|)$ tests, such as blood and tissue type, are performed to rule out the patient-donor pairs where donation is impossible (e.g., different blood types), but does not provide a conclusive answer for whether or not a donation is indeed feasible. To be (more) certain that a donor can donate to a patient, more accurate tests must be performed before the transplant, which include crossmatching, antibody screen, and so on.¹ Note that performing such tests for all pairs of donor-patient pairs could require $\Omega(|V|^2)$ number of tests, each of which is typically costly and time-consuming.

The stochastic matching problem captures the essence of the need of extra tests for kidney exchange: An algorithm selects a set of patient-donor pairs to perform the extra costly and time-consuming tests (i.e., query the edges), while making sure that w.h.p. there is a large matching among the pairs that pass the extra tests (i.e., in the realized graph). The objective of querying as few edges as possible captures the essence of minimizing the total cost, and the objective of having small degree of adaptivity captures the essence of minimizing patients’ waiting time by performing the extra exams between many patient-donor pairs in parallel.

The kidney exchange problem has been extensively studied in the literature, particularly under stochastic settings (see, e.g., References [2–4, 10, 18–20, 26, 31]); we refer the interested reader to Reference [13] for a more detailed discussion. Moreover, we should also note that swapping kidneys between two donor-patient pairs is only a part of a kidney exchange system, and in more

¹ American Transplant Foundation: <http://www.americantransplantfoundation.org/>.

recent years, there has been kidney exchanges performed in longer cycles or chains. The focus of this article is solely on kidney swapping (i.e., pair-wise exchange) and the techniques developed in this article do not seem to be directly applicable to longer chains. We refer the interested reader to Reference [30] for more details on three-way exchanges and to Reference [5] for exchanges with even longer chains.

1.2 Model

We now formally define the model for the stochastic matching problem. For any graph $G(V, E)$, let $\text{OPT}(G)$ denote the *maximum matching size* in G , i.e., the number of edges in a maximum matching. With a slight abuse of notation, we sometimes also use $\text{OPT}(X) := \text{OPT}(G(V, X))$ for $X \subseteq E$, i.e., X is a set of edges instead of a graph. Throughout, we use n to denote $|V|$.

In the stochastic setting, for the input graph $G(V, E)$, each edge in E is realized *independently* w.p.² p . When sampling each edge w.p. p , for any set of edges $X \subseteq E$, we slightly abuse the notation and use X_p to denote both a random variable that corresponds to sampling edges in X w.p. p , as well as a specific realization of the random variable X_p . We call each possible realized graph $G(V, E_p)$ a *realization* of G .

In the *stochastic matching problem*, we are given a graph $G(V, E)$ and our goal is to compute a matching M in $G(V, E_p)$, such that w.h.p. (taken over both the randomness of the algorithm and the randomness of the realization $E_p \subseteq E$), the size of M is close to $\text{OPT}(E_p)$. An algorithm is allowed to *query* any edge $e \in E$ to determine whether or not $e \in E_p$, and we consider the following two classes of algorithms:

- **Non-adaptive algorithms.** A non-adaptive algorithm specifies a subset of edges $Q \subseteq E$, queries all edges in Q in parallel, and outputs a matching among the edges realized in Q .
- **Adaptive algorithms.** An adaptive algorithm proceeds in *rounds* where in each round, based on the edges queried and realized thus far, the algorithm chooses a new set of edges to query in parallel. We say the *degree of adaptivity* of an algorithm is d if the algorithm makes at most d rounds of adaptive queries.

In general, the goal is to design algorithms where the number of per-vertex queries is independent of n . For adaptive algorithms, the degree of adaptivity is further required to be independent of n .

We remark that throughout, we always assume $\text{OPT}(G) = \omega(1/p)$ to obtain the desired concentration bounds. This assumption essentially says that the expected matching size in the realized graph is bounded from below by a sufficiently large constant.

1.3 Related Work

Prior to our work, the state-of-the-art adaptive and non-adaptive algorithms for stochastic matching are that of Blum et al. [13], which is an adaptive (resp. non-adaptive) algorithm that achieves a $(1 - \epsilon)$ -approximate (resp. $(\frac{1}{2} - \epsilon)$ -approximate) matching *in expectation*, while both the number of per-vertex queries and the degree of adaptivity (for the adaptive algorithm) is $\frac{1}{p^{O(1/\epsilon)}}$. Note that while in these algorithms, the number of per-vertex queries and degree of adaptivity is independent of n , the exponential dependence on $\frac{1}{\epsilon}$, limiting the practical appeal of the algorithm. Blum et al. [13] raised an open problem regarding the possibility of avoiding the exponential dependency on $\frac{1}{\epsilon}$ for both the number of per-vertex queries and the degree of adaptivity.

Other variants of the stochastic matching setting have also been studied in the literature. Blum et al. [14] considered the setting where each vertex can only pick two incident edges to query,

²Throughout, we use *w.p.* and *w.h.p.* to abbreviate “with probability” and “with high probability,” respectively.

and the goal is to find the optimal set of edges to query. Another well-studied setting is the *query-commit* model, whereby if an algorithm decides to query an edge e , then e must be part of the matching the algorithm outputs in case e is realized [1, 11, 15, 17, 24].

1.4 Our Results

We provide algorithms for the stochastic matching problem with exponentially smaller number of queries and degree of adaptivity (for the adaptive algorithm) compared to the best previous bounds of Blum et al. [13]. In particular,

THEOREM 1.1 (INFORMAL). *For any $\varepsilon > 0$, there exists a poly-time adaptive $(1 - \varepsilon)$ -approximation algorithm for the stochastic matching problem that queries $O(\frac{\log(1/\varepsilon p)}{\varepsilon p})$ edges per vertex and has degree $O(\frac{\log(1/\varepsilon p)}{\varepsilon p})$ of adaptivity.*

The formal statement of Theorem 1.1 is presented in Section 4 (as Theorem 4.1).

THEOREM 1.2 (INFORMAL). *For any $\varepsilon > 0$, there exists a poly-time non-adaptive $(\frac{1}{2} - \varepsilon)$ -approximation algorithm for the stochastic matching problem that queries $O(\frac{\log(1/\varepsilon p)}{\varepsilon p})$ edges per vertex.*

The formal statement of Theorem 1.2 is presented in Section 5 (as Theorem 5.1).

These results provide an affirmative answer to an open question raised by Blum et al. [13] regarding the possibility of avoiding the exponential dependency on $1/\varepsilon$ for both the number of per-vertex queries and the degree of adaptivity.

Furthermore, for obtaining an $(1 - \varepsilon)$ approximation using a non-adaptive algorithm, no non-trivial lower bound on the number of per-vertex queries is known (not even for any arbitrary dependency on p and ε). We illustrate in Section 6 a barrier to achieve such a result.

One of the key properties of our results is that we provide *instance-wise* approximation guarantees, i.e., for $1 - o(1)$ fraction of the realizations of G , the algorithm outputs a competitive solution. This is a *stronger* guarantee than the *expectation* guarantee provided in Blum et al. [13], which states that expected size of the matching output by the algorithm is competitive with the expected size of the maximum matching among all realizations. We remark that our instance-wise guarantee is of particular interest to the key application of the kidney exchange problem.

Finally, it is worth mentioning that even when the input graph is a complete graph, one needs to query $\Omega(\log(1/\varepsilon)/p)$ edges per each vertex to simply ensure that the number of isolated vertices in the realized graph is at most εn . This is due to the fact that if one queries less than $\ln(1/\varepsilon)/2p$ edges per vertex, the probability that no edge incident on a vertex is realized is $(1 - p)^{\ln(1/\varepsilon)/2p} \geq \exp(-2p \cdot \ln(1/\varepsilon)/2p) = \varepsilon$. However, it is easy to see that for any constant $p > 0$, any realization of a complete graph has a perfect matching w.h.p. Hence, $\Omega(\log(1/\varepsilon)/p)$ is a simple lower bound on the number of per-vertex queries for any $(1 - \varepsilon)$ -approximation algorithm even on complete graphs. Our per-vertex query bounds in Theorem 1.1 and Theorem 1.2 only ask for slightly more than this simple lower bound.

1.5 Our Techniques

To explain the high-level idea underlying our algorithms, it will be convenient to focus on the case when G_p has a perfect matching; however, we emphasize that our algorithms do not require this property. The idea behind both of our algorithms is to construct a *matching-cover* of the input graph G and query the edges of the cover in the algorithm. Roughly speaking, a γ -matching-cover of a graph $G(V, E)$ is a collection of matchings of G of size $\gamma \cdot (|V|/2)$ that are essentially edge-disjoint (see Section 3.1 for a formal definition). One of the main technical ingredients of

our work is a structural result proving that: For any algorithm that outputs a γ -matching cover with $\Theta(1/\varepsilon p)$ matchings, w.h.p. the set of realized edges in the cover contains a matching of size $(1 - \varepsilon) \cdot \gamma \cdot (|V|/2)$. We prove this result through a constructive argument based on the Tutte-Berge formula (see Section 2 for more details on Tutte-Berge formula).

Next, we show that there is a simple adaptive (resp. non-adaptive) algorithm that computes a 1-matching-cover (resp. 1/2-matching-cover) with $\Theta(1/\varepsilon p)$ matchings, which immediately implies that w.h.p. the algorithm achieves a $(1 - \varepsilon)$ -approximation (resp. $(\frac{1}{2} - \varepsilon)$ -approximation).

Finally, to eliminate the assumption that G_p has a perfect matching, we establish a *vertex sparsification lemma* (see Section 3.2) that allows us to reduce the number of vertices in any instance G from $|V|$ to $O(\text{OPT}(G)/\varepsilon)$, while w.h.p. preserving the maximum matching size to within a factor of $(1 - \varepsilon)$. In the sparsified graph $\mathcal{G}(\mathcal{V}, \mathcal{E})$, although we only have $\text{OPT}(G_p) = \Omega(|\mathcal{V}|)$ instead of having a perfect matching, we can show that, with some more care in the analysis, the constant gap between $\text{OPT}(G_p)$ and $|\mathcal{V}|$ is enough for us to establish the approximation ratios of our algorithms.

Comparison with Reference [13]. The adaptive algorithm of Blum et al. [13] can be summarized as follows: Maintain a matching M , and at each round find a collection of vertex-disjoint *augmenting paths* of length $O(1/\varepsilon)$ in the input graph $G(V, E)$ with respect to M ; query the edges of the augmenting paths and augment M if possible. Using the well-known fact that a matching with no augmenting path of length $O(1/\varepsilon)$ is a $(1 - \varepsilon)$ -approximate matching, the authors show that the found matching of the algorithm is a $(1 - \varepsilon)$ -approximation in expectation.

In this process, the probability that an augmenting path of length $O(1/\varepsilon)$ “survives” the querying process is only $p^{O(1/\varepsilon)}$; hence, one needs to repeat the whole process roughly $\frac{1}{p^{O(1/\varepsilon)}}$ times, which leads to the same degree of adaptivity and per-vertex queries. The non-adaptive algorithm of Blum et al. [13] is designed based on similar framework of using augmenting paths.

However, our algorithms exploit the structure of matchings in a global way (using the Tutte-Berge formula) instead of locally searching for short augmenting paths. In particular, through the use of matching covers, we completely eliminate the need of searching for the augmenting paths and hence avoid the exponential dependency on $\frac{1}{\varepsilon}$, which is essentially the length of the augmenting paths. It is worth mentioning, however, that most of these differences only appear in the analysis; the description of our algorithms and algorithms of Blum et al. [13] are both simple and similar (modulo the extra sparsification part of our algorithm).

1.6 Subsequent Work

Subsequent to the conference version of this article [7], the stochastic matching problem has been studied further in References [8, 12, 32]. In a follow-up work [8], we were able to break the 1/2 approximation barrier for the non-adaptive case in previous work [7, 13] and achieve a non-adaptive algorithm with approximation ratio strictly larger than 1/2. This was subsequently improved to a (2/3)-approximation for non-adaptive algorithms in a very recent work [6]. Our work in Reference [8] built on the techniques in the current article plus many new ideas; in particular, a structural result on graphs with large matchings in expectation in their realizations. However, the techniques in Reference [6] are entirely disjoint from the current article.

Our techniques in this article were also extended to achieve efficient adaptive and non-adaptive algorithms for the more general problem of stochastic integer programs [32]. Interestingly, as pointed out in Reference [32], the proposed strategy in Reference [32] coincides with our approach when applied to the stochastic matching problem. However, while the results in Reference [32] span a wider set of problems, the bounds achieved in Reference [32] for the stochastic matching problem are weaker than ours by a factor of $\Theta(\log n)$, i.e., Reference [32] achieves algorithms with $O(\frac{\log(n/\varepsilon)}{\varepsilon p})$ per vertex queries as opposed to $O(\frac{\log(1/\varepsilon p)}{\varepsilon p})$ achieved in the current article. We

emphasize that in general, the goal in the stochastic matching problem is to design algorithms where the number of per-vertex queries and degree of adaptivity is independent of n , and hence the loss of this extra $O(\log n)$ factor is significant in this context.

Finally, the weighted variant of the stochastic matching problem is very recently studied in Behnezhad and Reyhani [12]. The authors in Reference [12] developed an adaptive $(1 - \varepsilon)$ -approximation algorithm for the stochastic weighted matching problem that requires $(1/p)^{O(1/\varepsilon)}$ per vertex queries and degree of adaptivity. The techniques in Reference [12] are more along the lines of the previous approach in Reference [13] (compared to our article) and as such suffer from the exponential dependence on the parameter $1/\varepsilon$. It is an intriguing question to understand whether our techniques can be used for the stochastic weighted matching problem as well to avoid the exponential dependence on $1/\varepsilon$.

2 PRELIMINARIES

Notation. Throughout, we use n to denote $|V|$, where V is the set of vertices in the input graph $G(V, E)$. For any set of edges $X \subseteq E$, $V(X)$ denotes the set of vertices incident on X . For two integers $a \leq b$, $[a, b]$ denotes the set $\{a, \dots, b\}$ and $[b] := [1, b]$.

Tutte-Berge formula. In our proofs, we crucially rely on the Tutte-Berge formula, which generalizes the *Hall's marriage theorem* for characterizing perfect matchings in bipartite graphs to maximum matchings in general graphs. For any graph $G(V, E)$ and any $U \subseteq V$, $\text{odd}(V - U)$ denotes the number of connected components with *odd* number of vertices in $G(V \setminus U, E)$. We have,

LEMMA 2.1 (TUTTE-BERGE FORMULA). *The size of a maximum matching in a graph $G(V, E)$ is*

$$\frac{1}{2} \min_{U \subseteq V} (|U| + |V| - \text{odd}(V - U)).$$

See, e.g., Reference [25] (Chapter 3) for a proof of this lemma.

Finally, we have the following simple concentration result on the size of a maximum matching in G_p ; the proof is a standard application of the Chernoff bound:

CLAIM 2.2. *For any graph $G(V, E)$ with $\text{OPT}(G) = \omega(1/p)$, $\Pr(\text{OPT}(G_p) \geq p \cdot \text{OPT}(G)/2) = 1 - o(1)$.*

3 MATCHING COVERS AND VERTEX SPARSIFICATION FOR STOCHASTIC MATCHING

In this section, we present our main technical results, namely the *matching-cover lemma* and the *vertex sparsification lemma*, which lie in the core of both our adaptive and non-adaptive algorithms. We start by describing the matching-cover lemma. As explained earlier in Section 1.5, the matching-cover lemma is already sufficient for establishing the approximation guarantee of the algorithms as long as $\text{OPT}(G) = \Omega(n)$. To tackle the case where $\text{OPT}(G)$ is much smaller than the number of vertices, we later introduce a simple vertex sparsification lemma that provides a way of reducing the number of vertices in any graph $G(V, E)$ from $|V|$ to $O(\text{OPT}(G))$ while preserving the maximum matching size approximately, for *any* realization of $G(V, E_p)$.

3.1 Matching-cover Lemma

We start by defining the following process, which takes any graph as an input and outputs a list of matchings (i.e., a *matching-cover*).

Definition 3.1 (Incremental Matching Selection Process). We say an algorithm is an *incremental matching selection process (IMSP)* iff for any input graph G the algorithm selects a sequence of matchings $M_1, M_2, \dots, M_r$ one-by-one from G such that for any $i \in [r]$, for any edge e selected in the i th matching M_i where e also appears in M_j for some $j < i$, the edge e must be realized.

We refer to the matchings an IMSP outputs as a *matching-cover*, and we say that an IMSP outputs a γ -*matching-cover* iff for all $i \in [r]$, $|M_i| \geq \frac{\gamma n}{2}$. Note that an IMSP could be either an adaptive algorithm or a non-adaptive algorithm. For instance, one could, (a) adaptively, pick a maximum matching M from G , query the edges of M , remove the edges that are not realized, and repeat: In this case, an edge e appearing in more than one matching implies that in the first matching that contains e , e must be realized; or (b) non-adaptively, pick a maximum matching M from G , remove M from G , and repeat: In this case, no edge appears in more than one matching. The following claim states the key property of any IMSP, which will be used in our proofs.

CLAIM 3.2. *For any IMSP A and any graph G , denote by $M_1, M_2, \dots, M_r$ the matching-cover that A outputs on G ; then for any $i \in [r]$, conditioned on any realization of $M_1, M_2, \dots, M_{i-1}$, and any choice of the matching M_i , each edge $e \in M_i$ is realized w.p. at least p , independent of any other edges in M_i .*

PROOF. For the edges $e \in M_i$ that appear in previous matchings, by the definition of IMSP, e is realized w.p. $1(\geq p)$, which is trivially independent of any other edges in M_i . For the set of edges E' that do not appear in any previous matching, E' is disjoint with $M_1, M_2, \dots, M_{i-1}$ and the set of realized edges in $M_1, M_2, \dots, M_{i-1}$ is independent of realization of edges in E' . Therefore, by the definition of the stochastic setting, each edge in E' is realized w.p. p independently. $\square$

We are now ready to state the matching-cover lemma, which is the main result of this section.

LEMMA 3.3 (MATCHING-COVER LEMMA). *For any parameter $0 < \varepsilon, p < 1$, any graph G , and any IMSP A , denote by $M_1, \dots, M_r$ the γ -matching-cover of G that A outputs. If $r \geq \frac{32 \log(2e/\gamma)}{\varepsilon p}$ and $\gamma \cdot n = \omega(1)$, then, w.p. $1 - o(1)$, there is a matching of size at least $(1 - \varepsilon) \frac{\gamma n}{2}$ among the realized edges in the matching-cover.*

We remark that Lemma 3.3 holds even for multi-graphs, which is a property required by our algorithms (due to the usage of vertex sparsification). We first provide a high-level summary of the proof. Suppose by contradiction that the output of IMSP A does not contain a large matching; then, by Tutte-Berge formula, there should exist a set of vertices $U \subseteq V$ where removing U from the graph results in many connected components (CC) with odd number of vertices; namely, U is an *odd-sets-witness* (see Figure 1(a) for an example of an odd-sets-witness). Our strategy is to show that, for any fixed set U , the probability that U ends up being an odd-sets-witness is sufficiently small, and then use a union bound over all possible choices of U to argue that w.h.p. no such odd-sets-witness can arise.

To see that w.h.p. each U does not lead to an odd-sets-witness, we again use the Tutte-Berge formula: If the edges realized in the first i matchings leave many odd-size CC's, then the large matching M_{i+1} must eliminate most of them. Note that this is not enough to show that the number of *odd-size CC's* will decrease w.h.p., since it is possible that M_{i+1} eliminates two odd-size CC's by connecting them through a chain of even-size CC's (see the long chain in Figure 1(b) for an illustration). The length of the chain could be arbitrarily long and even if one edge on the chain is not realized, the two odd-size CC's will still end up being disconnected. A key observation here is that though the number of odd-size CC's might not decrease (requiring all edges on the chain to be realized), the total number of CC's will decrease w.h.p. (any realized edge on the chain reduces the number of CC's). Using this fact, we show that after enough rounds, the total number of CC's will drop significantly; and even if all of them are odd-size CC's, it is not enough for being an odd-sets-witness.

PROOF OF LEMMA 3.3. If the edges realized in the matching-cover do not contain a matching of size more than $(1 - \varepsilon) \frac{\gamma n}{2}$, then, by Lemma 2.1, there exists a set of vertices U where the number

Fig. 1. An example of adding a large matching M_i to an odd-sets-witness (the red/thick edges). The number of odd components does not decrease, but the total number of connected components indeed decreases.

of odd-size connected components after removing U , i.e., $odd(V - U)$, satisfies

$$(1 - \varepsilon) \frac{\gamma n}{2} \geq \frac{1}{2}(|U| + |V| - odd(V - U)),$$

which implies that $odd(V - U) \geq |U| + (1 - \gamma + \varepsilon\gamma)n$. In this case, we say U leads to an *odd-sets-witness* and denote this event by E_U . Using this fact, we only need to prove the following lemma. $\square$

CLAIM 3.4. For any $U \subseteq V$, $Pr(E_U) \leq 2^{-2\gamma n \log(2e/\gamma)}$.

We first show that Claim 3.4 implies Lemma 3.3. We will apply a union bound over all candidate sets U to show the probability that *there exists* some U where E_U happens is $o(1)$. To do so, we argue that the number of different choices of U 's that need to be considered is at most $2^{\gamma n \log(2e/\gamma)}$. To see this, note that every odd-size set must contain at least one vertex, and there are only n vertices that could be part of the odd-size sets. Thus, we have $n \geq odd(V - U)$. However, as we just established, $odd(V - U) \geq |U| + (1 - \gamma + \varepsilon\gamma)n$, and combining the two inequalities, we have

$$|U| \leq n - (1 - \gamma + \varepsilon\gamma)n \leq \gamma n. \quad (1)$$

Therefore, it suffices to consider the sets U with cardinality at most γn , and only

$$\sum_{i=0}^{\gamma n} \binom{n}{i} \leq \binom{n + \gamma n}{\gamma n} \leq \left(\frac{e(1 + \gamma)n}{\gamma n} \right)^{\gamma n} \leq \left(\frac{2e}{\gamma} \right)^{\gamma n} = 2^{\gamma n \log(2e/\gamma)}$$

such choices of U exists. Now, we can apply a union bound over all such choices of U , and by Claim 3.4 w.p. at least $1 - 2^{-\gamma n \log(2e/\gamma)} = 1 - o(1)$ (recall that $\gamma n = \omega(1)$), there is no odd-sets-witness, proving Lemma 3.3.

The main step is to prove Claim 3.4, and as stated in Equation (1) we can assume, w.l.o.g., that $|U| \leq \gamma n$.

PROOF OF CLAIM 3.4. Recall that the goal is to show that w.h.p., U does not lead to an odd-sets-witness (i.e., E_U does not happen). We first define some notation. For any $i \in [r]$, $M_p(i)$ denotes the set of edges in M_i that are realized and are not incident on vertices in U , and G_i denotes the graph formed by edges $M_p(1) \cup M_p(2) \cup \dots \cup M_p(i)$. We use $cc(G_i)$ to denote the number of connected components in G_i and use E_i to denote the event that the number of odd-size connected components in G_i is at least $|U| + (1 - \gamma + \varepsilon\gamma)n$.

Let Y_i be a random binary variable where $Y_i = 1$ iff $cc(G_{i-1}) - cc(G_i) < \varepsilon p \cdot \gamma n / 16$ (and $Y_i = 0$ otherwise), which is the event that the edges of $M_p(i)$ do not reduce the number of connected components in G_{i-1} by more than $\varepsilon p \cdot \gamma n / 16$. Suppose for at least half of the rounds, $Y_i = 0$; then, the number of connected components after the IMSP selects the r matchings is less than

$$n - \frac{r}{2} \cdot \frac{\varepsilon p \cdot \gamma n}{16} \leq n - \frac{32 \log(2e/\gamma)}{\varepsilon p} \cdot \frac{1}{2} \cdot \frac{\varepsilon p \cdot \gamma n}{16} \leq (1 - \gamma)n.$$

However, since having $|U| + (1 - \gamma + \varepsilon \gamma)n$ odd-size connected components (i.e., E_U happens) implies that there are more than $(1 - \gamma)n$ connected components,

$$\Pr(E_U) = \Pr\left(E_U, \sum_{i \in [r]} Y_i \geq \frac{r}{2}\right). \quad (2)$$

Hence, we can focus on upper bounding the probability that E_U happens *and* for more than half of the rounds, $Y_i = 1$. In the following, we first establish a key property regarding the probability that $Y_i = 1$ (Claim 3.5) and then show how to use this property to bound the probability of the target event in Equation (2) (Claim 3.6). To simplify the presentation, we use $M_p(i_1, i_2, \dots, i_j)$ to denote the set of edges in $M_p(i_1) \cup M_p(i_2) \cup \dots \cup M_p(i_j)$, and, with a slight abuse of notation, use Y_i to denote the event that $Y_i = 1$. The proofs of the following two claims are postponed to Section 3.1.1 to preserve the flow of the current argument. $\square$

CLAIM 3.5. For any $M_p(1, \dots, i-1)$, $\Pr(Y_i, E_{i-1} \mid M_p(1, \dots, i-1)) \leq \exp(-\frac{3\varepsilon p \cdot \gamma n}{16})$.

CLAIM 3.6. For any set of $r/2$ rounds $i_1, i_2, \dots, i_{r/2}$, $\Pr(Y_{i_1}, Y_{i_2}, \dots, Y_{i_{r/2}}, E_U) \leq 2^{-3\gamma n \log(2e/\gamma)}$.

By Claim 3.6, for each collection of $\frac{r}{2}$ rounds, Y_i happens to all of them w.p. at most $2^{-3\gamma n \log(2e/\gamma)}$. There are at most 2^r (which is independent of n) choices of different (at least) $\frac{r}{2}$ rounds, hence using union bounds, for n sufficiently large, the probability that Y_i happens in more than $\frac{r}{2}$ rounds is at most $2^{-2\gamma n \log(2e/\gamma)}$, proving Claim 3.4.

3.1.1 Missing Technical Details in the Proof of Claim 3.4. We now present the missing proofs in Claim 3.4, which was the key part of the proof of Lemma 3.3. The following is a restatement of Claim 3.5 in the proof of Claim 3.4.

CLAIM. For any $M_p(1, \dots, i-1)$,

$$\Pr(Y_i, E_{i-1} \mid M_p(1, \dots, i-1)) \leq \exp\left(-\frac{3\varepsilon p \cdot \gamma n}{16}\right).$$

PROOF. Recall that E_{i-1} is the event that the number of odd-size connected components in G_{i-1} is at least $|U| + (1 - \gamma + \varepsilon \gamma)n$. We have,

$$\begin{aligned} & \Pr(Y_i, E_{i-1} \mid M_p(1, \dots, i-1)) \\ &= \Pr(Y_i \mid M_p(1, \dots, i-1), E_{i-1}) \cdot \Pr(E_{i-1} \mid M_p(1, \dots, i-1)) \quad (\text{Chain rule}) \\ &\leq \Pr(Y_i \mid M_p(1, \dots, i-1), E_{i-1}), \end{aligned}$$

hence, we only need to show that $\Pr(Y_i \mid M_p(1, \dots, i-1), E_{i-1}) \leq \exp(-\frac{3\varepsilon p \cdot \gamma n}{16})$, which is, roughly speaking, the probability that the i th matching M_i does not reduce the number of connected components by a lot, given that the graph G_{i-1} contains many odd-size connected components. $\square$

To proceed, we need the following definition: For any graph H , we say a set of edges E' form a *component-based spanning forest* of H , if E' is a spanning forest of the graph obtained by contracting each connected component in H into a single vertex (and any edge $(u, v) \in E'$ becomes an edge

between the connected components that u and v , respectively, resides in). It is straightforward to verify that if we add any component-based spanning forest $E' \subset E$ to H , the number of connected components in H would reduce by $|E'|$. The following claim is the key to obtain the target upper bound on $\Pr(Y_i \mid M_p(1, \dots, i-1), E_{i-1})$:

CLAIM 3.7. *Whenever E_{i-1} happens, there exist at least $\frac{\varepsilon \gamma n}{2}$ edges of M_i that form a component-based spanning forest of G_{i-1} .*

PROOF. Since the matching M_i has size at least $\gamma n/2$ (by definition of being in a γ -matching-cover), the edges of M_i can reduce the number of odd-size sets from at least $|U| + (1 - \gamma + \varepsilon \gamma)n$ (i.e., E_{i-1} happens) down to at most $|U| + (1 - \gamma)n$ (by Lemma 2.1: $(|U| + n - \text{odd}(V - U))/2 \leq \gamma n/2 \Rightarrow \text{odd}(V - U) \geq |U| + (1 - \gamma)n$). Therefore, after adding edges of M_i to G_{i-1} , at least $\varepsilon \gamma n$ odd-size sets will disappear. For each odd-size set S that disappears, M_i must contain at least one edge between S and another connected component. Therefore, for the largest component-based spanning forest obtained by the edges in M_i , at least $\varepsilon \gamma n$ vertices (i.e., connected components) have degree at least one, which implies the number of edges in the forest is at least $\frac{\varepsilon \gamma n}{2}$. $\square$

Let $T_i \subseteq M_i$ be the set of (at least) $\frac{\varepsilon \gamma n}{2}$ edges promised by Claim 3.7 (conditioned on E_{i-1}) and t_i be the number of edges realized in T_i . By Claim 3.2, each edge in T_i is realized w.p. at least p independent of the other edges, hence, $E[t_i \mid E_{i-1}] \geq \varepsilon p \cdot \gamma n/2$. Note that t_i is a lower bound on $cc(G_{i-1}) - cc(G_i)$ (i.e., the decrement of the number of connected components from G_{i-1} to G_i), and $Y_i = 1$ iff $cc(G_{i-1}) - cc(G_i) \leq \varepsilon p \cdot \gamma n/16$, which implies that no more than $\varepsilon p \cdot \gamma n/16$ edges are realized in T_i (which is $1/8$ of the expectation). Hence, by Chernoff bound,

$$\begin{aligned} \Pr(Y_i \mid M_p(1, \dots, i-1), E_{i-1}) &\leq \Pr\left(t_i \leq \frac{\varepsilon p \cdot \gamma n}{16} \mid E_{i-1}\right) \\ &\leq \exp\left(-\frac{1}{2} \cdot \left(\frac{7}{8}\right)^2 \cdot \frac{\varepsilon p \cdot \gamma n}{2}\right) = \exp\left(-\frac{49}{64} \cdot \frac{\varepsilon p \cdot \gamma n}{4}\right) \\ &\leq \exp\left(-\frac{48}{64} \cdot \frac{\varepsilon p \cdot \gamma n}{4}\right) \leq \exp\left(-\frac{3\varepsilon p \cdot \gamma n}{16}\right), \end{aligned}$$

which concludes the proof of Claim 3.5.

The following is a restatement of Claim 3.6 in the proof of Claim 3.4:

CLAIM. *For any collection of $r/2$ rounds $i_1, i_2, \dots, i_{r/2}$,*

$$\Pr(Y_{i_1}, Y_{i_2}, \dots, Y_{i_{r/2}}, E_U) \leq 2^{-3\gamma n \log(2e/\gamma)}.$$

PROOF. Assume w.l.o.g. that $i_1 < i_2 < \dots < i_{r/2}$. First, E_U happening implies that $E_{i_1-1}, E_{i_2-1}, \dots, E_{i_{r/2}-1}$ should all happen³; therefore,

$$\Pr(Y_{i_1}, Y_{i_2}, \dots, Y_{i_{r/2}}, E_U) \leq \Pr(Y_{i_1}, Y_{i_2}, \dots, Y_{i_{r/2}}, E_{i_1-1}, E_{i_2-1}, \dots, E_{i_{r/2}-1}). \quad (3)$$

After reorganizing the terms, we have,

$$\begin{aligned} &\Pr(Y_{i_1}, Y_{i_2}, \dots, Y_{i_{r/2}}, E_{i_1-1}, E_{i_2-1}, \dots, E_{i_{r/2}-1}) \\ &= \Pr(Y_{i_1}, E_{i_1-1}, Y_{i_2}, E_{i_2-1}, \dots, Y_{i_{r/2}}, E_{i_{r/2}-1}) \\ &= \prod_{j \in [r/2]} \Pr(Y_{i_j}, E_{i_j-1} \mid Y_{i_1}, E_{i_1-1}, Y_{i_2}, E_{i_2-1}, \dots, Y_{i_{j-1}}, E_{i_{j-1}-1}). \end{aligned} \quad (\text{Chain rule})$$

We will upper bound each of the $r/2$ terms separately. Fix a $j \in [r/2]$, denote the event $(Y_{i_1}, E_{i_1-1}, Y_{i_2}, E_{i_2-1}, \dots, Y_{i_{j-1}}, E_{i_{j-1}-1})$ by E^* . Note that E^* is completely determined by

³It is straightforward to verify that the number of odd-size connected components is monotonically decreasing.

$M_p(1, \dots, i_{j-1})$, which is also determined by $M_p(1, \dots, i_j - 1)$ since $i_j - 1 \geq i_{j-1}$. We have

$$\begin{aligned}
 \Pr(Y_{i_j}, E_{i_j-1} \mid Y_{i_1}, E_{i_1-1}, Y_{i_2}, E_{i_2-1}, \dots, Y_{i_{j-1}}, E_{i_{j-1}-1}) &= \Pr(Y_{i_j}, E_{i_j-1} \mid E^*) \\
 &= \sum_{M_p(1, \dots, i_{j-1})} \Pr(M_p(1, \dots, i_j - 1) \mid E^*) \cdot \Pr(Y_{i_j}, E_{i_j-1} \mid E^*, M_p(1, \dots, i_j - 1)) \\
 &= \sum_{M_p(1, \dots, i_j-1) \text{ s.t. } E^* \text{ happens}} \Pr(M_p(1, \dots, i_j - 1) \mid E^*) \cdot \Pr(Y_{i_j}, E_{i_j-1} \mid M_p(1, \dots, i_j - 1)) \\
 &\quad (E^* \text{ is determined by } M_p(1, \dots, i_j - 1)) \\
 &\leq \sum_{M_p(1, \dots, i_j-1) \text{ s.t. } E^* \text{ happens}} \Pr(M_p(1, \dots, i_j - 1) \mid E^*) \cdot \exp\left(-\frac{3\epsilon p \cdot \gamma n}{16}\right) \quad (\text{By Claim 3.5}) \\
 &= \exp\left(-\frac{3\epsilon p \cdot \gamma n}{16}\right).
 \end{aligned}$$

Therefore,

$$\Pr(Y_{i_1}, Y_{i_2}, \dots, Y_{i_{r/2}}, E_U) \leq \exp\left(-\frac{3\epsilon p \cdot \gamma n}{16} \cdot \frac{r}{2}\right) \leq \exp(-3\gamma n \log(2e/\gamma)) \leq 2^{-3\gamma n \log(2e/\gamma)},$$

where the second inequality is by the choice of r . $\square$

3.2 Vertex Sparsification Lemma

In the following, we give an algorithm that for any $0 < \epsilon < 1$ reduces the number of vertices in any graph G from $|V|$ to $O(\text{OPT}(G)/\epsilon)$, while preserving the maximum matching size to within a factor of $(1 - \epsilon)$ for *any* realization G_p of G w.h.p.

For inputs G and ϵ and a sparsification parameter τ to be determined later, SPARSIFY(G, τ, ϵ) works as follows: We create and output a multi-graph $\mathcal{G}(\mathcal{V}, \mathcal{E})$ where: (i) $|\mathcal{V}| = \tau$, (ii) each vertex v in G is mapped to a vertex $\mathcal{V}(v)$ in $\mathcal{G}$ chosen *uniformly at random* from $\mathcal{V}$, and (iii) for each edge (u, v) , there is a corresponding edge between $\mathcal{V}(u)$ and $\mathcal{V}(v)$. A pseudo-code of the SPARSIFY algorithm is presented in Algorithm 1. We point out that similar ideas of randomly grouping vertices for matchings have also been recently used in References [9, 16] for the purpose of reducing *space requirement* of algorithms in graph streams.

ALGORITHM 1: SPARSIFY(G, τ, ϵ). A Matching-Preserving Sparsification Algorithm

Input: Graph $G(V, E)$, sparsification parameter τ , and input parameter $\epsilon > 0$.

Output: A multi-graph $\mathcal{G}(\mathcal{V}, \mathcal{E})$ with $|\mathcal{V}| = \tau$.

- (1) Partition the vertices in V into τ groups $\mathcal{V} := (\mathcal{V}_1, \dots, \mathcal{V}_\tau)$ by assigning each vertex *independently* to one of the τ groups chosen *uniformly at random*.
 - (2) For any edge $(u, v) \in E$, add an edge $e_{u,v}$ between the vertices $\mathcal{V}(u)$ and $\mathcal{V}(v)$ in $\mathcal{E}$.
 - (3) Return the multi-graph $\mathcal{G}(\mathcal{V}, \mathcal{E})$.
-

The following lemma states the main property of the SPARSIFY algorithm:

LEMMA 3.8. *For any graph $G(V, E)$, suppose $\mathcal{G}(\mathcal{V}, \mathcal{E}) := \text{SPARSIFY}(G, \tau, \epsilon)$ for the parameter $\tau \geq \frac{8 \cdot \text{OPT}(G)}{\epsilon}$, and let M be any fixed matching in G with $|M| = \omega(1)$; then, w.p. $1 - o(1)$, there exists a matching $\mathcal{M}$ of size $(1 - \epsilon) \cdot |M|$ in $\mathcal{G}$. Moreover, the edges of $\mathcal{M}$ correspond to a unique matching M' in G of the same size.*

PROOF. Let $\epsilon' := \epsilon/4$, and let s denote the number of vertices matched in M (i.e., $s = 2|M|$). Note that $\tau \geq 8\text{OPT}(G)/\epsilon = 4s/\epsilon = s/\epsilon'$. For the sake of analysis, we first merge the τ groups $\mathcal{V}$ into

s/ε' *super-groups* in the following manner: Fix any partition that evenly breaks $[\tau]$ into s/ε' non-empty parts $P_1, P_2, \dots, P_{s/\varepsilon'}$ (i.e., $|P_i| = \frac{\varepsilon'\tau}{s}$ for any $i \in [s/\varepsilon']$). For each part P_i , define a super-group $\mathcal{S}_i := \cup_{j \in [P_i]} \mathcal{V}_j$. Then, since in Algorithm 1 each vertex $v \in V$ is assigned to exactly one group chosen uniformly at random, the probability that $v \in \mathcal{S}_i$ is $\frac{\varepsilon'}{s}$.

We say a super-group $\mathcal{S}_i$ is *good* iff $\mathcal{S}_i$ contains at least one vertex in $V(M)$, and is otherwise *bad*. For each super-group $\mathcal{S}_i$, let X_i be a random variable where $X_i = 1$ iff $\mathcal{S}_i$ is bad (otherwise, $X_i = 0$). Let $X := \sum_{i \in [s/\varepsilon']} X_i$ be the number of bad super-groups. In the following, we first show that X is small w.h.p., which implies that there are many good super-groups, and then show that there is a large matching between the good super-groups.

To see that $X = \sum_{i \in [s/\varepsilon']} X_i$ is small w.h.p., first notice that

$$\Pr(X_i = 1) = \left(1 - \frac{\varepsilon'}{s}\right)^s \leq e^{-\varepsilon'} \leq 1 - \varepsilon' + \frac{\varepsilon'^2}{2}. \quad (\forall x \geq 0, e^{-x} \leq 1 - x + x^2/2)$$

Therefore, we have $E[X] = \sum_{i \in [s/\varepsilon']} E[X_i] \leq (1 - \varepsilon' + \frac{\varepsilon'^2}{2}) \frac{s}{\varepsilon'}$. However, since at most $|V(M)| (= s)$ super-groups could contain a vertex from $V(M)$ (i.e., could be good), $X \geq s/\varepsilon' - s = \Omega(s) = \omega(1)$, and hence $E[X] = \omega(1)$. To continue, observe that our setting can be viewed as a standard balls-and-bins experiment: Each vertex in $V(M)$ is a ball; each super-group is a bin; and X_i denotes the event that the i th bin is empty. Therefore, the random variables X_i 's are *negatively correlated*, and we can apply Chernoff bound (see Reference [27] for more details on negative correlation):

$$\Pr(X \geq (1 + \varepsilon'^2)E[X]) \leq e^{-\Omega_{\varepsilon'}(E[X])} = o(1) \quad (E[X] = \omega(1)).$$

Since $E[X] \leq (1 - \varepsilon' + \frac{\varepsilon'^2}{2}) \frac{s}{\varepsilon'}$ (as shown above), we further have

$$\Pr\left(X \geq (1 + \varepsilon'^2) \left(1 - \varepsilon' + \frac{\varepsilon'^2}{2}\right) \frac{s}{\varepsilon'}\right) \leq \Pr(X \geq (1 + \varepsilon'^2)E[X]) = o(1).$$

Hence, w.p. $1 - o(1)$, the number of bad super-groups is at most

$$(1 + \varepsilon'^2) \left(1 - \varepsilon' + \frac{\varepsilon'^2}{2}\right) \frac{s}{\varepsilon'} \leq \left(1 - \varepsilon' + \frac{\varepsilon'^2}{2} + \varepsilon'^2 - \varepsilon'^3 + \frac{\varepsilon'^4}{2}\right) \frac{s}{\varepsilon'} \leq (1 - \varepsilon' + 2\varepsilon'^2) \frac{s}{\varepsilon'},$$

which implies that the number of good super-groups is at least $\frac{s}{\varepsilon'} - (1 - \varepsilon' + 2\varepsilon'^2) \frac{s}{\varepsilon'} = (\varepsilon' - 2\varepsilon'^2) \frac{s}{\varepsilon'} = (1 - 2\varepsilon')s$.

It remains to show that if at least $(1 - 2\varepsilon')s$ super-groups are good (i.e., contain a vertex from $V(M)$), then the edges in M form a matching $\mathcal{M}$ in $\mathcal{G}$ of size at least $(1 - 4\varepsilon')|M| = (1 - \varepsilon)|M|$. To see this, for each good super-group $\mathcal{S}_i$, we fix one vertex $v \in V(M) \cap \mathcal{S}_i$ and remove all other vertices in $\mathcal{S}_i$. For the matching M , at most $2\varepsilon's$ vertices in $V(M)$ are removed and hence at least $s/2 - 2\varepsilon's = (1 - 4\varepsilon')|M|$ edges in M remain. Since all endpoints of these edges are assigned to distinct super-groups, these edges form a matching $\mathcal{M}$ of size $\geq (1 - 4\varepsilon')|M| = (1 - \varepsilon)|M|$ in $\mathcal{G}$.

To see the second part of the lemma, simply note that each edge of $\mathcal{M}$ comes from a distinct edge in M . $\square$

For any $G_p := G(V, E_p)$, define $\mathcal{G}_p$ as the graph obtained from $\mathcal{G}$ by considering only the edges that correspond to edges in E_p . We are now ready to prove our vertex sparsification lemma.

LEMMA 3.9 (VERTEX SPARSIFICATION LEMMA). *Let $G(V, E)$ be a graph with maximum matching size $\omega(1/p)$ and let $\mathcal{G} = \text{SPARSIFY}(G, \tau, \varepsilon)$ for $\tau = \frac{8 \cdot \text{OPT}(G)}{\varepsilon}$; then,*

$$\Pr(\text{OPT}(\mathcal{G}_p) \geq (1 - \varepsilon) \cdot \text{OPT}(G_p)) = 1 - o(1),$$

where the probability is taken over both the inner randomness of the SPARSIFY algorithm as well as the realization $E_p \subseteq E$.

PROOF. By Claim 2.2, the maximum matching size in G_p is $\omega(1)$ w.p. $1 - o(1)$. Now, for any realization G_p with maximum matching size of $\omega(1)$, by construction, $\mathcal{G}_p = \text{SPARSIFY}(G_p, \tau, \epsilon)$. Hence, by applying Lemma 3.8 on any maximum matching M in G_p , we have that $\mathcal{G}_p$ has a matching of size $(1 - \epsilon)|M|$ w.p. $1 - o(1)$. $\square$

4 A $(1 - \epsilon)$ -APPROXIMATION ADAPTIVE ALGORITHM

We now present our adaptive algorithm and prove Theorem 1.1. The following is a formal restatement of Theorem 1.1.

THEOREM 4.1. *There is an adaptive algorithm that for any input graph $G(V, E)$, and any input parameter $\epsilon > 0$, outputs a matching of size $\text{ALG} := \text{ALG}(G_p)$ such that,*

$$\Pr(\text{ALG} \geq (1 - \epsilon) \cdot \text{OPT}) = 1 - o(1),$$

where $\text{OPT} := \text{OPT}(G_p)$ is the maximum matching size in $G_p(V, E_p)$. The probability is taken over both the inner randomness of the algorithm and the realization $E_p \subseteq E$.

Moreover, the algorithm makes only $O(\frac{\log(1/\epsilon p)}{\epsilon p})$ rounds of adaptive queries, and queries only $O(\frac{\log(1/\epsilon p)}{\epsilon p})$ edges per vertex.

Our adaptive algorithm in Theorem 4.1 works as follows: We first use our vertex sparsification lemma (Lemma 3.9) to compute a graph $\mathcal{G}(\mathcal{V}, \mathcal{E}) := \text{SPARSIFY}(G, \tau, \epsilon)$ where $\text{OPT}(\mathcal{G}) = \Omega(|\mathcal{V}|)$. Next, we repeat for $O(\frac{\log(1/\epsilon p)}{\epsilon p})$ rounds the following operation: Pick a maximum matching M from $\mathcal{G}$, query the edges of M , and remove the edges that are not realized. Finally, return a maximum matching among the realized edges. The pseudo-code of this algorithm is presented as Algorithm 2.

ALGORITHM 2: A $(1 - \epsilon)$ -Approximation Adaptive Algorithm for Stochastic Matching

Input: Graph $G(V, E)$ and input parameters $0 < \epsilon, p < 1$.

Output: A matching M in $G(V, E_p)$.

- (1) Let $\mathcal{G}(\mathcal{V}, \mathcal{E}) := \text{SPARSIFY}(G, \tau, \epsilon)$ for $\tau := \left\lceil \frac{8\text{OPT}(G)}{\epsilon} \right\rceil$.
 - (2) Let $R := \left\lceil \frac{32 \log(8e/\epsilon p)}{\epsilon \cdot p} \right\rceil$, and $\mathcal{E}^* \leftarrow \mathcal{E}$.
 - (3) For $i = 1, \dots, R$, do:
 - (a) Pick a maximum matching M_i in $\mathcal{G}(\mathcal{V}, \mathcal{E}^*)$.
 - (b) Query the edges in M_i and remove the non-realized edges from $\mathcal{E}^*$.
 - (4) Output a maximum matching among realized edges in $M_1, M_2, \dots, M_R$.
-

It is straightforward to verify the number of queries and the degree of adaptivity used in Algorithm 2: Each round of the algorithm queries edges of a matching, and hence each vertex is incident on at most one queried edge. We now prove the bound on the approximation ratio of the algorithm.

PROOF OF THEOREM 4.1. Let $\widehat{\text{OPT}} := \widehat{\text{OPT}}(\mathcal{G}_p)$ denote the maximum matching size in the graph $\mathcal{G}_p$. By Lemma 3.9, w.p. $1 - o(1)$, $\widehat{\text{OPT}} \geq (1 - \epsilon)\text{OPT}(G_p)$. Now, it suffices to show that w.p. $1 - o(1)$, Algorithm 2 outputs a matching of size at least $(1 - \epsilon)\widehat{\text{OPT}}$, since, with a union bound, it would imply w.p. $1 - o(1)$,

$$\text{ALG}(G_p) \geq (1 - \epsilon) \cdot \widehat{\text{OPT}} \geq (1 - \epsilon)^2 \cdot \text{OPT}(G_p) \geq (1 - 2\epsilon) \cdot \text{OPT}(G_p),$$

and we can replace ϵ with $\epsilon/2$ in Algorithm 2 to obtain a $(1 - \epsilon)$ -approximation.

To see that $\text{ALG}(G_p) \geq (1 - \varepsilon)\widehat{\text{OPT}}$ w.h.p., let $L := \min_{i \in [R]} |M_i|$, i.e., the minimum size of a matching chosen by Algorithm 2. Since all M_i 's are maximum matchings in $\mathcal{E}^*$ while $\mathcal{E}^*$ always contains all edges of the optimum matching in $\mathcal{G}_p$, we have $L \geq \widehat{\text{OPT}}$ and we can focus on showing $\text{ALG}(G_p) \geq (1 - \varepsilon)L$.

It is straightforward to verify that the way Algorithm 2 selects the matchings $M_1, M_2, \dots, M_R$ satisfies the condition of IMSP (Definition 3.1). Moreover, by Claim 2.2 and Lemma 3.9, we have, $\Pr(\widehat{\text{OPT}} \geq p \cdot \text{OPT}(G)/2) = 1 - o(1)$. Therefore, we can use Lemma 3.3 with parameters:

$$\gamma = \frac{2L}{|\mathcal{V}|} \geq \frac{\varepsilon \cdot p}{4}, \quad r = \frac{32 \log(2e/\gamma)}{\varepsilon p} \leq \frac{32 \log(8e/\varepsilon p)}{\varepsilon p} \leq R,$$

which states that w.p. $1 - o(1)$, the realized edges in matchings $M_1, M_2, \dots, M_R$ contain a matching of size at least $(1 - \varepsilon)\frac{\gamma|\mathcal{V}|}{2} = (1 - \varepsilon)L$, which completes the proof. $\square$

5 A $(\frac{1}{2} - \varepsilon)$ -APPROXIMATION NON-ADAPTIVE ALGORITHM

In this section, we present our non-adaptive algorithm and prove Theorem 1.2. The following is a formal restatement of Theorem 1.2.

THEOREM 5.1. *There is a non-adaptive algorithm that for any input graph $G(V, E)$, any input parameter $\varepsilon > 0$, outputs a matching of size $\text{ALG} := \text{ALG}(G_p)$ such that,*

$$\Pr\left(\text{ALG} \geq \left(\frac{1}{2} - \varepsilon\right) \cdot \text{OPT}\right) = 1 - o(1),$$

where $\text{OPT} := \text{OPT}(G_p)$ is the maximum matching size in $G_p(V, E_p)$. The probability is taken over both the inner randomness of the algorithm and the realization $E_p \subseteq E$.

Moreover, the algorithm non-adaptively queries $O(\frac{\log(1/\varepsilon p)}{\varepsilon p})$ edges per vertex.

Note that any non-adaptive algorithm works in the following framework:

- (1) Compute a subgraph $H(V, Q)$ of the input graph $G(V, E)$ for some $Q \subseteq E$.
- (2) Query all edges in Q and compute a maximum matching in $H(V, Q_p)$.

Therefore, the main task of any non-adaptive algorithm is to choose a “good” subgraph H . Our non-adaptive algorithm chooses a subgraph H as follows: We first use our vertex sparsification lemma to compute a graph $\mathcal{G}(\mathcal{V}, \mathcal{E}) := \text{SPARSIFY}(G, \tau, \varepsilon)$ where $\text{OPT}(\mathcal{G}) = \Omega(|\mathcal{V}|)$. Next, we repeat for $O(\frac{\log(1/\varepsilon p)}{\varepsilon p})$ times the process of picking a maximum matching from $\mathcal{G}(\mathcal{V}, \mathcal{E})$ and removing the edges of the matching from $\mathcal{E}$. Let Q be the set of edges in these matchings; the algorithm returns $H(V, Q)$ as the subgraph H . A pseudo-code of this algorithm is presented as Algorithm 3.

ALGORITHM 3: A Non-Adaptive $(\frac{1}{2} - \varepsilon)$ -Approximation Algorithm for Stochastic Matching

Input: Graph $G(V, E)$ and input parameters $0 < \varepsilon, p < 1$.

Output: A matching M in $G(V, E_p)$.

- (1) Let $\mathcal{G}(\mathcal{V}, \mathcal{E}) := \text{SPARSIFY}(G, \tau, \varepsilon)$ for $\tau := \left\lceil \frac{8\text{OPT}(G)}{\varepsilon} \right\rceil$.
 - (2) Let $R := \left\lceil \frac{32 \log(16e/\varepsilon p)}{\varepsilon \cdot p} \right\rceil$.
 - (3) Initially $Q \leftarrow \emptyset$. For $i = 1, \dots, R$, do:
 - (a) Pick a maximum matching M_i in $\mathcal{G}(\mathcal{V}, \mathcal{E} \setminus Q)$.
 - (b) Let $Q \leftarrow Q \cup M_i$.
 - (4) Query all edges in Q and return a maximum matching in $M_1, M_2, \dots, M_R$.
-

We now briefly provide the intuition behind Algorithm 3. Similar to the adaptive case, using the vertex sparsification lemma (Lemma 3.9), our task reduces to approximating the maximum matching in $\mathcal{G}_p$ (as opposed to G_p). For the adaptive case, our algorithm guarantees that every selected matching is of size at least $\text{OPT}(\mathcal{G}_p)$, which allows us to use the matching-cover lemma directly to complete the argument. However, Algorithm 3 does not have such a strong guarantee, since it is non-adaptive. To address this issue, we establish a weaker guarantee that allows us to obtain a $1/2$ -approximation. The idea is as follows: On one hand, if the smallest matching selected by the algorithm is of size at least $\text{OPT}(\mathcal{G}_p)/2$, we can still invoke the matching-cover lemma (Lemma 3.3) and have that Algorithm 3 outputs a matching of size $(1 - \varepsilon)\text{OPT}(\mathcal{G}_p)/2$. On the other hand, we show that if the smallest selected matching has size less than $\text{OPT}(\mathcal{G}_p)/2$, then for *any* maximum matching M^* in $\mathcal{G}_p$, we must have selected in Q at least half the edges of M^* , which immediately results in a matching of size $\text{OPT}(\mathcal{G}_p)/2$.

We now present the formal proof. In the following, let $L := \min_{i \in [R]} |M_i|$, i.e., the minimum size of a matching chosen by Algorithm 3. Note that, equivalently, $L = |M_R|$, since the size of matchings chosen by the algorithm is a *non-increasing* sequence by construction.

LEMMA 5.2. *Let $L = |M_R|$; if $L \geq p \cdot \text{OPT}(G)/4$, then, $\Pr(\text{ALG} \geq (1 - \varepsilon) \cdot L) = 1 - o(1)$.*

PROOF. Since $M_1, \dots, M_R$ are edge disjoint matchings, the process of choosing $M_1, \dots, M_R$ is an IMSP (Definition 3.1). Hence, by Lemma 3.3 with parameters:

$$\gamma = \frac{2L}{|V|} \geq \frac{\varepsilon \cdot p}{8}, \quad r = \frac{32 \log(2e/\gamma)}{\varepsilon p} \leq \frac{32 \log(16e/\varepsilon p)}{\varepsilon p} = R,$$

there exists a matching of size $(1 - \varepsilon) \cdot \frac{\gamma|V|}{2} = (1 - \varepsilon) \cdot L$ in Q_p w.p. $1 - o(1)$. Noting that any matching of $\mathcal{G}_p$ in Q_p corresponds to a matching in G_p completes the proof. $\square$

We define $\widehat{\text{OPT}} := \widehat{\text{OPT}}(\mathcal{G}_p)$ as the maximum matching size in $\mathcal{G}_p$.

LEMMA 5.3. $\Pr(\text{ALG} \geq (\frac{1}{2} - \varepsilon) \cdot \widehat{\text{OPT}} \mid \widehat{\text{OPT}} > 2L) = 1$.

PROOF. Let M be any arbitrary matching in $\mathcal{G}$. We have, $|M| - |Q \cap M| \leq L$, since otherwise, for the matching $M' := M \setminus Q$, we have $|M'| = |M| - |Q \cap M| > L > |M_R|$, contradicting the fact that M_R is a maximum matching in the remaining graph.

Now let M be any maximum matching in $\mathcal{G}_p$. Since $\widehat{\text{OPT}} > 2L$, we have $|M| > 2L$. Consequently, $|Q \cap M| \geq |M| - L \geq |M| - |M|/2 = |M|/2$. Hence, at least half of the edges in M are also present in Q , implying that in this case, $\text{ALG} \geq \widehat{\text{OPT}}/2$ w.p. 1. $\square$

We now prove the bound on the approximation ratio of Algorithm 3.

LEMMA 5.4. $\Pr(\text{ALG} \geq (\frac{1}{2} - 2\varepsilon) \cdot \text{OPT}) = 1 - o(1)$.

PROOF. Recall that $\widehat{\text{OPT}} := \widehat{\text{OPT}}(\mathcal{G}_p)$. By Claim 2.2 and Lemma 3.9, we have,

$$\Pr(\widehat{\text{OPT}} \geq p \cdot \text{OPT}(G)/2) = 1 - o(1). \quad (4)$$

Let E_{win} be the event that $\text{ALG} \geq (\frac{1}{2} - \varepsilon) \cdot \widehat{\text{OPT}}$. We argue that $\Pr(E_{\text{win}}) = 1 - o(1)$. This, together with the fact that w.p. $1 - o(1)$, $\widehat{\text{OPT}} \geq (1 - \varepsilon) \cdot \text{OPT}$ (Lemma 3.9) completes the proof.

Consider two cases, (i) $L < p \cdot \text{OPT}(G)/4$, and (ii) $L \geq p \cdot \text{OPT}(G)/4$. For case (i),

$$\begin{aligned} \Pr(E_{\text{win}}) &\geq \Pr(\widehat{\text{OPT}} > 2L) \cdot \Pr(E_{\text{win}} \mid \widehat{\text{OPT}} > 2L) \\ &\geq \Pr(\widehat{\text{OPT}} > p \cdot \text{OPT}(G)/2) \cdot \Pr(E_{\text{win}} \mid \widehat{\text{OPT}} > 2L) = (1 - o(1)) \cdot 1, \end{aligned}$$

Fig. 2. An example of a barrier to $(1 - \varepsilon)$ -approximation non-adaptive algorithms. The edges in the gadget graph are *not* presented in this figure. In part (b), solid red edges (resp. dashed edges) are the edges that are realized (resp. not realized).

where the last equality is by Equation (4) and Lemma 5.3. For case (ii),

$$\begin{aligned}
 \Pr(E_{win}) &= \Pr(\widehat{\text{OPT}} > 2L) \cdot \Pr(E_{win} \mid \widehat{\text{OPT}} > 2L) + \Pr(\widehat{\text{OPT}} \leq 2L) \cdot \Pr(E_{win} \mid \widehat{\text{OPT}} \leq 2L) \\
 &= \Pr(\widehat{\text{OPT}} > 2L) + \Pr(\widehat{\text{OPT}} \leq 2L) \cdot \Pr(E_{win} \mid \widehat{\text{OPT}} \leq 2L) \quad (\text{By Lemma 5.3}) \\
 &\geq \Pr(\widehat{\text{OPT}} > 2L) + \Pr(\widehat{\text{OPT}} \leq 2L) \cdot \Pr(\text{ALG} \geq (1 - \varepsilon) \cdot L \mid \widehat{\text{OPT}} \leq 2L) \\
 &\geq \Pr(\widehat{\text{OPT}} > 2L) \cdot \Pr(\text{ALG} \geq (1 - \varepsilon) \cdot L \mid \widehat{\text{OPT}} > 2L) \\
 &\quad + \Pr(\widehat{\text{OPT}} \leq 2L) \cdot \Pr(\text{ALG} \geq (1 - \varepsilon) \cdot L \mid \widehat{\text{OPT}} \leq 2L) \\
 &= \Pr(\text{ALG} \geq (1 - \varepsilon) \cdot L) = 1 - o(1),
 \end{aligned}$$

where the last equality is by Lemma 5.2. $\square$

Theorem 5.1 now follows from Lemma 5.4 (by replacing ε with $\varepsilon/2$ in Algorithm 3), and the fact that Algorithm 3 queries $O(\frac{\log(1/\varepsilon p)}{\varepsilon p})$ matchings and hence $O(\frac{\log(1/\varepsilon p)}{\varepsilon p})$ incident edges are queried for each vertex.

6 A BARRIER TO OBTAINING A NON-ADAPTIVE $(1 - \varepsilon)$ -APPROXIMATION ALGORITHM

The approximation ratio of our non-adaptive algorithm is $(\frac{1}{2} - \varepsilon)$ as opposed to the near-optimal ratio of $(1 - \varepsilon)$ achieved by our adaptive algorithm. A natural question, first raised by Reference [13], is if one can obtain a $(1 - \varepsilon)$ -approximation using a non-adaptive algorithm, even by allowing arbitrary dependence on p and ε . In the following, we highlight a possible barrier to obtaining such a result.

Consider a bipartite graph $G(L, R, E)$ constructed as follows: (i) the vertex sets are $L = V_1 \cup V_3$, $R = V_2 \cup V_4$, and $|V_i| = N$ for $i \in [4]$, (ii) there is a perfect matching between V_1 and V_2 , and a perfect matching between V_3 and V_4 , and (iii) there is a *gadget graph* $\widehat{G}(V_2, V_3, \widehat{E})$, to be determined later, between V_2 and V_3 (see Figure 2(a)).

Suppose we want to design a non-adaptive $(1 - \varepsilon)$ -approximation algorithm for the instance $G(V, E)$ with the parameter $p = 2/3$. In this case, for any graph G_p , w.h.p., there is a matching M_1 between V_1 and V_2 , and another matching M_2 between V_3 and V_4 , each of size $(2/3)N - o(N)$. Hence,

$$\text{OPT}(G_p) \geq |M_1| + |M_2| + (2/3) \cdot m(A, B) - o(N), \quad (5)$$

where A (resp. B) is the set of vertices in V_2 (resp. V_3) that are not matched by M_1 (resp. M_2), and $m(A, B)$ denotes the size of a maximum matching between A and B . A few observations are

in order. First, picking edges of M_1 and M_2 is crucial for having any large matching in G_p ; and second, for a uniformly at-random chosen realization of M_1 and M_2 , the sets A and B are chosen uniformly at random from V_2 and V_3 (see Figure 2(b)). Based on these observations, we define the following problem:

PROBLEM 1. *Given a bipartite graph $G(L, R, E)$, choose a subgraph $H(L, R, Q)$ such that given two subsets $A \subseteq L$ and $B \subseteq R$, if $m(A, B) \geq N/3 - o(N)$ in G , then H contains at least $\Omega(N)$ edges between A and B .*

The goal is to solve Problem 1 using a graph H with small number of edges. Note that any non-adaptive $(1 - \varepsilon)$ -approximation algorithm has to solve Problem 1 for the gadget graph $\bar{G}$ when the two sets A and B are chosen *uniformly at random*, since otherwise, for the maximum matching size $\text{ALG}(G_p)$ in $H(V, Q_p)$ and maximum matching size $\text{OPT}(G_p)$ in G_p , we have:

$$\begin{aligned} \text{ALG}(G_p) &\leq |M_1| + |M_2| + o(N) \leq (4/3)N + o(N) \\ \text{OPT}(G_p) &\geq (4/3)N + (2/3)(N/3) - o(N) = (14/9)N - o(N). \end{aligned} \quad (\text{by Equation (5)})$$

Hence, the approximation ratio of the algorithm in this instance is at most $6/7 + o(1)$, bounded away from being a $(1 - \varepsilon)$ approximation for $\varepsilon < 1/7$.

Although for randomly chosen subsets A and B , no lower bound on the size of H is known, we show in the following that if A and B are chosen *adversarially*, then there exist graphs for which solving Problem 1 requires choosing a subgraph with *super linear* in n number of edges. Note that the number of queries of any non-adaptive algorithm is at least the number of edges in H and hence this bound on the number of edges in H implies that $\omega(n)$ queries are needed; or in other words, the number of per-vertex queries needs to be a *function of n* . The existence of such a graph indicates a barrier to obtain a non-adaptive $(1 - \varepsilon)$ -approximation algorithm.

To continue, we need a few definitions: For a graph $G(V, E)$, a matching M is called an *induced matching* if there is no edge between the vertices matched in M , i.e., $V(M)$, except for the edges in M . A graph $G(V, E)$ is called an (r, t) -Ruzsa-Szemerédi graph $((r, t)$ -RS graph for short) if the edge set E can be partitioned into t *induced matchings* each of size r . Note that the number of edges in any (r, t) -RS graph is $r \cdot t$.

Suppose $G(L, R, E)$ is an (r, t) -RS graph with the parameter $r = N/3$ and induced matchings $M_1, \dots, M_t$. For each $i \in [t]$, define A_i (resp. B_i) as $V(M_i) \cap L$ (resp. $V(M_i) \cap R$). Suppose we choose the pair (A, B) only from the set of pairs $\mathcal{F} := \{(A_1, B_1), \dots, (A_t, B_t)\}$. Note that between any pair in $\mathcal{F}$, there is a matching of size $r = N/3$, and moreover all edges of G are partitioned between these matchings. If the subgraph $H(V, Q)$ has only $o(r \cdot t)$ edges, a simple counting argument suggests that for a $1 - o(1)$ fraction of pairs in $\mathcal{F}$, only $o(r)$ edges between the pairs are present in H . Hence, H cannot be a solution to Problem 1.

To complete the argument, we point out that there are (r, t) -RS graphs on $2N$ vertices with parameters $r = N/3$ and $t = N^{\Omega(1/\log \log N)}$ [21, 23]. These constructions certify that to solve Problem 1 when the sets A and B are chosen adversarially, one needs to choose a subgraph with $n^{1+\Omega(1/\log \log N)} = \omega(N \cdot \text{polylog}(N))$ edges. In conclusion, while this result does not rule out the possibility of a non-adaptive $(1 - \varepsilon)$ -approximation algorithm where the number of per-vertex queries is independent of N , it suggests that any such algorithm has to crucially overcome Problem 1 using the fact that the two sets A and B are chosen randomly instead of adversarially.

7 CONCLUSIONS

We studied the stochastic matching problem in this article. We showed that there exists an adaptive $(1 - \varepsilon)$ -approximation algorithm for this problem with $O(\frac{\log(1/\varepsilon p)}{\varepsilon p})$ per-vertex queries and

degree of adaptivity. We further presented a non-adaptive $(\frac{1}{2} - \epsilon)$ -approximation algorithm with $O(\frac{\log(1/\epsilon p)}{\epsilon p})$ per-vertex queries. These results represent an exponential improvement over the previous best bounds of Blum et al. [13], answering an open problem in that work.

An interesting direction for future research is to design or formally rule out a non-adaptive algorithm that obtains a $(1 - \epsilon)$ -approximation while maintaining the property that the number of per-vertex queries is independent of n . Toward this direction, we highlighted a potential barrier to achieve a $(1 - \epsilon)$ -approximation non-adaptively.

ACKNOWLEDGMENTS

We would like to thank Nika Haghtalab for introducing us to the stochastic matching problem.

REFERENCES

- [1] Marek Adamczyk. 2011. Improved analysis of the greedy algorithm for stochastic matching. *Inf. Process. Lett.* 111, 15 (2011), 731–737.
- [2] Mohammad Akbarpour, Shengwu Li, and Shayan Oveis Gharan. 2014. Dynamic matching market design. In *Proceedings of the ACM Conference on Economics and Computation (EC'14)*. 355.
- [3] Ross Anderson, Itai Ashlagi, David Gamarnik, and Yash Kanoria. 2015. A dynamic model of barter exchange. In *Proceedings of the 26th ACM-SIAM Symposium on Discrete Algorithms (SODA'15)*. 1925–1933.
- [4] Ross Anderson, Itai Ashlagi, David Gamarnik, and Alvin E. Roth. 2015. Finding long chains in kidney exchange using the traveling salesman problem. *Proc. Nat. Acad. Sci.* 112, 3 (2015), 663–668.
- [5] Itai Ashlagi and Alvin E. Roth. 2012. New challenges in multi-hospital kidney exchange. *Amer. Econ. Rev.* 102, 3 (May 2012), 354–359.
- [6] Sepehr Assadi and Aaron Bernstein. 2018. Towards a unified theory of sparsification for matching problems. Retrieved from: [CoRR abs/1811.02009](https://arxiv.org/abs/1811.02009) (2018).
- [7] Sepehr Assadi, Sanjeev Khanna, and Yang Li. 2016. The stochastic matching problem with (very) few queries. In *Proceedings of the ACM Conference on Economics and Computation (EC'16)*. 43–60.
- [8] Sepehr Assadi, Sanjeev Khanna, and Yang Li. 2017. The stochastic matching problem: Beating half with a non-adaptive algorithm. In *Proceedings of the ACM Conference on Economics and Computation, (EC'17)*. 99–116.
- [9] Sepehr Assadi, Sanjeev Khanna, Yang Li, and Grigory Yaroslavtsev. 2016. Maximum matchings in dynamic graph streams and the simultaneous communication model. In *Proceedings of the 27th ACM-SIAM Symposium on Discrete Algorithms (SODA'16)*. 1345–1364.
- [10] Pranjali Awasthi and Tuomas Sandholm. 2009. Online stochastic optimization in the large: Application to kidney exchange. In *Proceedings of the 21st International Joint Conference on Artificial Intelligence (IJCAI'09)*. 405–411.
- [11] Nikhil Bansal, Anupam Gupta, Jian Li, Julián Mestre, Viswanath Nagarajan, and Atri Rudra. 2012. When LP is the cure for your matching woes: Improved bounds for stochastic matchings. *Algorithmica* 63, 4 (2012), 733–762.
- [12] Soheil Behnezhad and Nima Reyhani. 2018. Almost optimal stochastic weighted matching with few queries. In *Proceedings of the ACM Conference on Economics and Computation*. 235–249.
- [13] Avrim Blum, John P. Dickerson, Nika Haghtalab, Ariel D. Procaccia, Tuomas Sandholm, and Ankit Sharma. 2015. Ignorance is almost bliss: Near-optimal stochastic matching with few queries. In *Proceedings of the 16th ACM Conference on Economics and Computation (EC'15)*. 325–342.
- [14] Avrim Blum, Anupam Gupta, Ariel D. Procaccia, and Ankit Sharma. 2013. Harnessing the power of two crossmatches. In *Proceedings of the 14th ACM Conference on Electronic Commerce (EC'13)*. 123–140.
- [15] Ning Chen, Nicole Immorlica, Anna R. Karlin, Mohammad Mahdian, and Atri Rudra. 2009. Approximating matches made in heaven. In *Proceedings of the 36th International Colloquium on Automata, Languages and Programming (ICALP'09)*. 266–278.
- [16] Rajesh Chitnis, Graham Cormode, Hossein Esfandiari, Mohammad Taghi Hajiaghayi, Andrew McGregor, Morteza Monemizadeh, and Sofya Vorotnikov. 2016. Kernelization via sampling with applications to finding matchings and related problems in dynamic graph streams. In *Proceedings of the 27th ACM-SIAM Symposium on Discrete Algorithms (SODA'16)*. 1326–1344.
- [17] Kevin P. Costello, Prasad Tetali, and Pushkar Tripathi. 2012. Stochastic matching with commitment. In *Proceedings of the 39th International Colloquium on Automata, Languages, and Programming (ICALP'12)*. 822–833.
- [18] John P. Dickerson, Ariel D. Procaccia, and Tuomas Sandholm. 2012. Dynamic matching via weighted myopia with application to kidney exchange. In *Proceedings of the 26th AAAI Conference on Artificial Intelligence*.

- [19] John P. Dickerson, Ariel D. Procaccia, and Tuomas Sandholm. 2013. Failure-aware kidney exchange. In *Proceedings of the 14th ACM Conference on Electronic Commerce (EC'13)*. 323–340.
- [20] John P. Dickerson and Tuomas Sandholm. 2015. FutureMatch: Combining human value judgments and machine learning to match in dynamic environments. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence*. 622–628.
- [21] Eldar Fischer, Eric Lehman, Ilan Newman, Sofya Raskhodnikova, Ronitt Rubinfeld, and Alex Samorodnitsky. 2002. Monotonicity testing over general poset domains. In *Proceedings of the 34th ACM Symposium on Theory of Computing (STOC'02)*. 474–483.
- [22] National Kidney Foundation. 2016. Organ Donation and Transplantation Statistics. Retrieved from: <https://www.kidney.org/news/newsroom/factsheets/Organ-Donation-and-Transplantation-Stats>.
- [23] Ashish Goel, Michael Kapralov, and Sanjeev Khanna. 2012. On the communication and streaming complexity of maximum bipartite matching. In *Proceedings of the 23rd ACM-SIAM Symposium on Discrete Algorithms (SODA'12)*. 468–485.
- [24] Anupam Gupta and Viswanath Nagarajan. 2013. A stochastic probing problem with applications. In *Proceedings of the 16th International Conference on Integer Programming and Combinatorial Optimization (IPCO'13)*. 205–216.
- [25] L. Lovász and D. Plummer. 2009. *Matching Theory*. American Mathematical Society. Retrieved from: <https://books.google.com/books?id=yW3WSVq8ygcC>.
- [26] David F. Manlove and Gregg O'Malley. 2014. Paired and altruistic kidney donation in the UK: Algorithms and experimentation. *ACM J. Exper. Algor.* 19, 1, Article 2.6 (2014), 2.6:1–2.6:21 pages.
- [27] Alessandro Panconesi and Aravind Srinivasan. 1997. Randomized distributed edge coloring via an extension of the Chernoff-Hoeffding bounds. *SIAM J. Comput.* 26, 2 (1997), 350–368.
- [28] Alvin E. Roth, Tayfun Sönmez, and M. Utku Ünver. 2004. Kidney exchange. *Quart. J. Econ.* 119, 2 (May 2004), 457–488.
- [29] Alvin E. Roth, Tayfun Sönmez, and M. Utku Ünver. 2005. Pairwise kidney exchange. *J. Econ. Theory* 125, 2 (2005), 151–188.
- [30] Alvin E. Roth, Tayfun Sönmez, and M. Utku Ünver. 2007. Efficient kidney exchange: Coincidence of wants in markets with compatibility-based preferences. *Amer. Econ. Rev.* 97, 3 (June 2007), 828–851.
- [31] Utku Ünver. 2010. Dynamic kidney exchange. *Rev. Econ. Stud.* 77, 1 (2010), 372–414.
- [32] Yutaro Yamaguchi and Takanori Maehara. 2018. Stochastic packing integer programs with few queries. In *Proceedings of the 29th ACM-SIAM Symposium on Discrete Algorithms (SODA'18)*. 293–310.

Received February 2017; revised March 2018; accepted December 2018