Towards a Unified Theory of Sparsification for Matching Problems

Sepehr Assadi[1]
Department of Computer and Information Science, University of Pennsylvania
Philadelphia, PA, US
sassadi@cis.upenn.edu

Aaron Bernstein
Department of Computer Science, Rutgers University
Piscataway, NJ, US
bernstei@gmail.com

Abstract

In this paper, we present a construction of a "matching sparsifier", that is, a sparse subgraph of the given graph that preserves large matchings approximately and is robust to modifications of the graph. We use this matching sparsifier to obtain several new algorithmic results for the maximum matching problem:

- An almost (3/2)-approximation one-way communication protocol for the maximum matching problem, significantly simplifying the (3/2)-approximation protocol of Goel, Kapralov, and Khanna (SODA 2012) and extending it from bipartite graphs to general graphs.
- An almost (3/2)-approximation algorithm for the stochastic matching problem, improving upon and significantly simplifying the previous 1.999-approximation algorithm of Assadi, Khanna, and Li (EC 2017).
- An almost (3/2)-approximation algorithm for the fault-tolerant matching problem, which, to our knowledge, is the first non-trivial algorithm for this problem.

Our matching sparsifier is obtained by proving new properties of the edge-degree constrained subgraph (EDCS) of Bernstein and Stein (ICALP 2015; SODA 2016)—designed in the context of maintaining matchings in dynamic graphs—that identifies EDCS as an excellent choice for a matching sparsifier. This leads to surprisingly simple and non-technical proofs of the above results in a unified way. Along the way, we also provide a much simpler proof of the fact that an EDCS is guaranteed to contain a large matching, which may be of independent interest.

2012 ACM Subject Classification Theory of computation → Sparsification and spanners, Theory of computation → Graph algorithms analysis.

Keywords and phrases Maximum matching, matching sparsifiers, one-way communication complexity, stochastic matching, fault-tolerant matching

Digital Object Identifier 10.4230/OASIcs.SOSA.2019.11

1 Introduction

A common tool for dealing with massive graphs is sparsification. Roughly speaking, a sparsifier of a graph G is a subgraph H that (approximately) preserves certain properties of G while having a smaller number of edges. Such sparsifiers have been studied in great detail

[1] Supported in part by the National Science Foundation grant CCF-1617851.

for various properties: for example, a spanner [6, 29] or a distance preserver [18, 20] preserves pairwise distances, a cut sparsifier [26, 11, 22] preserves cut information, and a spectral sparsifier [32, 8] preserves spectral properties of the graph. An additional property that we often require of a graph sparsifier is *robustness*: it should continue to be a good sparsifier even as the graph changes. Some sparsifiers are robust by nature (e.g cut sparsifiers), but others (e.g spanners) are not, and for this reason there is an extensive literature on designing sparsifiers that can provide additional robustness guarantees.

In this paper, we study the problem of designing robust sparsifiers for the prominent problem of *maximum matching*. Multiple notions of sparsification for the matching problem have already been identified in the literature. One example is a subgraph that preserves the largest matching inside any given subset of vertices in G approximately. This notion is also known as a matching cover or a matching skeleton [23, 27] in the literature and is closely related to the communication and streaming complexity of the matching problem. Another example of a sparsifier is a subgraph that can preserve the largest matching on random subsets of edges of G, a notion closely related to the stochastic matching problem [15, 5]. An example of a robust sparsifier for matching is a fault-tolerant subgraph, namely a subgraph G that continue to preserve large matchings in G even after a fraction of the edges is deleted by an adversary. As far as we know, the fault-tolerant matching problem has not previously been studied, but it is a natural model to consider as it has received lots of attention in the context of spanners and distance preservers (see e.g. [19, 28, 7, 17, 16]).

Our first contribution is a subgraph H that we show is a robust matching sparsifier in *all* of the senses above. Our result is thus the first to unify these notions of sparsification for the maximum matching problem. In addition to unifying, our construction yields improved results for each individual notion of sparsification and the corresponding problems, namely, the one-way communication complexity of matching, stochastic matching, and fault-tolerant matching problems. Interestingly, our unified approach allows us to also provide much simpler proofs than all previously existing work for these problems. The subgraph we use as our sparsifier comes from a pair of papers by Bernstein and Stein on dynamic matching [13, 14]—they refer to this subgraph as an edge-degree constrained subgraph (EDCS for short). The EDCS was also very recently used in [2] to design sublinear algorithms for matching across several different models for massive graphs. Our applications of the EDCS in the current paper, as well as the new properties we prove for the EDCS, are quite different from those in [13, 14, 2]. Our first contribution thus takes an existing subgraph, and then provides the first proofs that it satisfies the three notions of sparsification described above.

Our second contribution is a much simpler (and even slightly improved) proof of the main property of an EDCS in previous work proved in [13, 14], namely that an EDCS contains a large matching of the original graph. Our new proof significantly simplifies the analysis of [14] and allows for simple and self-contained proofs of the results in this paper.

Definition of the EDCS. Before stating our results, we give a definition of the EDCS from [13, 14], as this is the subgraph we use for all of our results (see Section 2 for more details).

▶ **Definition 1 ([13]).** For any graph $G(V, E)$ and integers $\beta \geq \beta^- \geq 0$, an *edge-degree constrained subgraph (EDCS)* (G, β, β^-) is a subgraph $H := (V, E_H)$ of G with the following two properties:

(P1) For any edge $(u, v) \in E_H$: $\deg_H(u) + \deg_H(v) \leq \beta$.

(P2) For any edge $(u, v) \in E \setminus E_H$: $\deg_H(u) + \deg_H(v) \geq \beta^-$.

It is not hard to show that an EDCS of a graph G always exists for any parameters $\beta > \beta^-$ and that it is sparse, i.e., only has $O(n\beta)$ edges. A key property of EDCS proven previously [13, 14] (and simplified in our paper) is that for any reasonable setting of the parameters (e.g. β^- being sufficiently close to β), any EDCS H of G contains an (almost) $3/2$ approximate matching of G.

1.1 Our Results and Techniques

We now give detailed definitions of the notions of sparsification and the corresponding problems addressed in this paper, as well as our results for each one. Our second contribution—a significantly simpler proof that an EDCS contains an almost (3/2)-approximate matching—is left for Section 3.

One-Way Communication Complexity of Matching. Consider the following two-player communication problem: Alice is given a graph $G_A(V, E_A)$ and Bob holds a graph $G_B(V, E_B)$. The goal for Alice is to send a single message to Bob such that Bob outputs an approximate maximum matching in $E_A \cup E_B$. What is the minimum length of the message, i.e., the one-way communication complexity, for achieving a certain fixed approximation ratio on all graphs? One can show that the message communicated by Alice to Bob is indeed a matching skeleton, namely a data structure (but not necessarily a subgraph), that allows Bob to find a large matching in a given subset of vertices in Alice's input (see [23] for more details).

This problem was first studied by Goel, Kapralov, and Khanna [23] (see also the subsequent paper of Kapralov [25]), owing to its close connection to one-pass streaming algorithms for matching. Goel *et al.* [23] designed an algorithm that achieves a (3/2)-approximation in bipartite graphs using only $O(n)$ communication and proved that any better than (3/2)-approximation protocol requires $n^{1+\Omega(\frac{1}{\log\log n})}$ communication even on bipartite graphs (see, e.g. [23, 4] for further details on this lower bound). A follow-up work by Lee and Singla [27] further generalized the algorithm of [23] to general graphs, albeit with a slightly worse approximation ratio of $5/3$ (compared to $3/2$ of [23]).

We extends the results in [23] to general graphs with almost no loss in approximation.

▶ Result 1. For any constant $\varepsilon > 0$, the protocol where Alice computes an EDCS of her graph with $\beta = O(1)$ and $\beta^- = \beta - 1$ and sends it to Bob is a $(3/2 + \varepsilon)$-approximation one-way communication protocol for the maximum matching problem with uses $O(n)$ communication.

We remark that both the previous algorithm of [23] as well as its extension in [27] are quite involved and rely on a fairly complicated graph decomposition as well as an intricate primal-dual analysis. As such, we believe that the main contribution in Result 1 is in fact in providing a simple and self-contained proof of this result.

Stochastic Matching. In the stochastic matching problem, we are given a graph $G(V, E)$ and a probability parameter $p \in (0, 1)$. A realization of G is a subgraph $G_p(V, E_p)$ obtained by picking each edge in G independently with probability p to include in E_p. The goal in this problem is to find a subgraph H of G with max-degree bounded by a function of p (independent of number of vertices), such that the size of maximum matching in realizations of H is close to size of maximum matching in realizations of G. It is immediate to see that H in this problem is simply a sparsifier of G which preserves large matchings on random subsets of edges.

This problem was first introduced by Blum *et al.* [15] primarily to model the kidney exchange setting and has since been studied extensively in the literature [3, 5, 10, 34]. Early algorithms for this problem in [15, 3] (and the later ones for the weighted variant of the problem [10, 34]) all had approximation ratio at least 2, naturally raising the question that whether 2 is the best approximation ratio achievable for this problem. Assadi, Khanna, and Li [5] ruled out this perplexing possibility by obtaining a slightly better than 2-approximation algorithm for this problem, namely an algorithm with approximation ratio close to 1.999 (which improves to 1.923 for small p).

We prove that an EDCS results in a significantly improved algorithm for this problem.

▶ Result 2. For any constant $\varepsilon > 0$, an *EDCS* of G with $\beta = O(\frac{\log(1/p)}{p})$ and $\beta^- = \beta - 1$ achieves a $(3/2+\varepsilon)$-approximation algorithm for the stochastic matching problem with a subgraph of maximum degree $O(\frac{\log(1/p)}{p})$.

We remark that our bound on the maximum degree in Result 2 is optimal (up to an $O(\log(1/p))$ factor) for any constant-factor approximation algorithm (see [5]). In addition to significantly improving upon the previous best algorithm of [5], our Result 2 is much simpler than that of [5], in terms of the both the algorithm and (especially) the analysis.

Remark. Independently and concurrently, Behnezhad *et al.* [9] also presented an algorithm for stochastic matching with a subgraph of max-degree $O(\frac{\log(1/p)}{p})$ that achieves an approximation of almost $(4\sqrt{2}-5)$ (≈ 0.6568 compared to 0.6666 in Result 2). They also provided an algorithm with approximation ratio strictly better than half for weighted stochastic matching (our result does not work for weighted graphs). In terms of techniques, our paper and [9] are entirely disjoint.

Fault-Tolerant Matching. Let $f \geq 0$ be an integer, $G(V,E)$ be a graph, and H be any subgraph of G. We say that H is an α-approximation f-tolerant subgraph of G iff for any subset $F \subseteq E$ of size $\leq f$, the maximum matching in $H \setminus F$ is an α-approximation to maximum matching in $G \setminus F$ – that is, H is a robust sparsifier of G. This definition is a natural analogy of other fault-tolerant subgraphs, such as fault-tolerant spanners and fault-tolerant distance preservers (see, e.g. [19, 28, 7, 17, 16]), to the maximum matching problem. Despite being such fundamental objects, quite surprisingly fault-tolerant subgraphs have not previously been studied for the matching problem.

We complete our discussion of applications of EDCS as a robust sparsifier by showing that it achieves an optimal size fault-tolerant subgraph for the matching problem.

▶ Result 3. For any constant $\varepsilon > 0$ and any $f \geq 0$, there exists a $(3/2+\varepsilon)$-approximation f-tolerant subgraph H of any given graph G with $O(f+n)$ edges in total.

The number of edges used in our fault-tolerant subgraph in Result 3 is clearly optimal (up to constant factors). In Appendix A.2, we show that by modifying the lower bound of [23] in the communication model, we can also prove that the approximation ratio of $(3/2)$ is optimal for any f-tolerant subgraph with $O(f)$ edges, hence proving that Result 3 is optimal in a strong sense. We also show that several natural strategies for this problem cannot achieve better than 2-approximation, hence motivating our more sophisticated approach toward this problem (see Appendix A.3).

The qualitative message of our work is clear: *An EDCS is a robust matching sparsifier under all three notions of sparsification described earlier, which leads to simpler and improved algorithms for a wide range of problems involving sparsification for matching problems in a unified way.*

Overall Proof Strategy

Recall that our algorithm in all of the results above is simply to compute an EDCS H of the input graph G (or G_A in the communication problem). The analysis then depends on the specific notion of sparsification at hand, but the same high- level idea applies to all three cases. In each case, we have an original graph G, and then a modified graph G^* produced by changes to G: G^* is $G_A \cup G_B$ in the communication model, the realized subgraph G_p in the stochastic matching, and the graph $G \setminus F$ after adversarially removing edges F in the fault-tolerant matching problem. Let H be the EDCS that our algorithm computes in G, and let H^* be the graph that results from H due to the modifications made to G. If we could show that H^* is an EDCS of G^* then the proof would be complete, since we know that an EDCS is guaranteed to contain an almost (3/2)-approximate matching. Unfortunately, in all the three problems that we study it might not be the case that H^* is an EDCS of G^*. Instead in each case we are able to exhibit subgraphs $\widetilde{H} \subseteq H^*$ and $\widetilde{G} \subseteq G^*$ such that $\widetilde{H}$ is an EDCS of $\widetilde{G}$, and size of maximum matching of $\widetilde{G}$ and G^* differ by at most a $(1+\varepsilon)$ factor. This guarantees an approximation ratio of almost $(3/2)(1+\varepsilon)$ (precisely what we achieve in all three results above), since the EDCS $\widetilde{H}$ preserves the maximum matching in $\widetilde{G}$ to within an almost (3/2)-approximation and $\widetilde{H}$ is a subgraph of H.

Organization. The rest of the paper is organized as follows. Section 2 includes notation, simple preliminaries, and existing work on the EDCS. In Section 3, we present a significantly simpler proof of the fact that an EDCS contains an almost (3/2)-approximation matching (originally proved in [14]). Sections 4, 5, and 6 prove the sparsification properties of the EDCS in, respectively, the one-way communication complexity of matching (Result 1), the stochastic matching problem (Result 2), and the fault-tolerant matching problem (Result 3). These three sections are designed to be self-contained (beside assuming the background in Section 2) to allow the reader to directly consider the part of most interest. The appendix contains some secondary observations.

2 Preliminaries and Notation

Notation. For any integer $t \geq 1$, $[t] := \{1, \ldots, t\}$. For a graph $G(V, E)$ and a set of vertices $U \subseteq V$, $N_G(U)$ denotes the neighbors of vertices in U in G and $E_G(U)$ denotes the set of edges incident on U. Similarly, for a set of edges $F \subseteq E$, $V(F)$ denotes the set of vertices incident on these edges. For any vertex $v \in V$, we use $\deg_G(v)$ to denote the degree of $v \in V$ in G (we may drop the subscript G in these definitions if it is clear from the context). We use $\mu(G)$ to denote the size of the maximum matching in the graph G.

Throughout the paper, we use the following two standard variants of the Chernoff bound.

▶ Proposition 1 (Chernoff Bound). Suppose $X_1, \ldots, X_t$ are t independent random variables that take values in $[0, 1]$. Let $X := \sum_{i=1}^{t} X_i$ and assume $\mathbb{E}[X] \leq \lambda$. For any $\delta > 0$ and integer $k \geq 1$,

$$\Pr\Big(|X - \mathbb{E}[X]| \geq \delta \cdot \lambda\Big) \leq 2 \cdot \exp\Big(-\frac{\delta^2 \cdot \lambda}{3}\Big),$$

$$\Pr\Big(|X - \mathbb{E}[X]| \geq k\Big) \leq 2 \cdot \exp\Big(-\frac{2k^2}{t}\Big).$$

We also need the following basic variant of Lovasz Local Lemma (LLL).

▶ Proposition 2 (Lovasz Local Lemma; cf. [21, 1]). Let $p \in (0,1)$ and $d \geq 1$. Suppose $\mathcal{E}_1, \ldots, \mathcal{E}_t$ are t events such that $\Pr(\mathcal{E}_i) \leq p$ for all $i \in [t]$ and each $\mathcal{E}_i$ is mutually independent of all but (at most) d other events $\mathcal{E}_j$. If $p \cdot (d+1) < 1/e$ then $\Pr\left(\cap_{i=1}^{n} \overline{\mathcal{E}_i}\right) > 0$.

Hall's Theorem. We use the following standard extension of the Hall's marriage theorem for characterizing maximum matching size in bipartite graphs.

▶ Proposition 3 (Extended Hall's marriage theorem; cf. [24]). Let $G(L, R, E)$ be any bipartite graph with $|L| = |R| = n$. Then, $\max\Big(|A| - |N(A)|\Big) = n - \mu(G)$, where A ranges over L or R. We refer to such set A as a *witness set*.

Proposition 3 follows from Tutte-Berge formula for matching size in general graphs [33, 12] or a simple extension of the proof of Hall's marriage theorem itself

Previously Known Properties of the EDCS

Recall the definition of an EDCS in Definition 1. It is not hard to show that an EDCS always exists as long as $\beta > \beta^-$ (see, e.g. [2]). For completeness, we repeat the proof in the Appendix A.1.

▶ Proposition 4 (cf. [13, 14, 2]). Any graph G contains an EDCS(G, β, β^-) for any parameters $\beta > \beta^-$, which can be found in polynomial time.

The key property of an EDCS, originally proved in [13, 14], is that it contains an almost (3/2)-approximate matching.

▶ **Lemma 2** ([13, 14]). *Let $G(V, E)$ be any graph and $\varepsilon < 1/2$ be a parameter. For parameters $\lambda \leq \frac{\varepsilon}{100}$, $\beta \geq 32\lambda^{-3}$, and $\beta^- \geq (1-\lambda) \cdot \beta$, in any subgraph $H :=$ EDCS(G, β, β^-), $\mu(G) \leq \left(\frac{3}{2} + \varepsilon\right) \cdot \mu(H)$.*

Another particularly useful (technical) property of an EDCS is that it "balances" the degree of vertices and their neighbors in the EDCS; this property is implicit in [13] but we explicitly state and prove it here as it shows a main distinction in the properties of EDCS compared to more standard (and less robust) subgraphs in this context such as b-matchings.

▶ Proposition 5. Let $H :=$ EDCS(G, β, β^-) and U be any subset of vertices. If average degree of U in H is $\bar{d}$ then average degree of $N_H(U)$ from edges incident on U is $\leq \beta - \bar{d}$.

Proof. Let H' be a subgraph of H containing the edges incident on U. Let $W := N_{H'}(U) = N_H(U)$ and $E' = E_H(U, W) = E_{H'}(U, W)$. We are interested in upper bounding the quantity $|E'| / |W|$. Firstly, by Property (P1) of EDCS, we have that $\sum_{(u,v) \in E'} \deg_{H'}(u) + \deg_{H'}(v) \leq \beta \cdot |E'|$. We write the LHS in this equation as:

$$\sum_{(u,v) \in E'} \deg_{H'}(u) + \deg_{H'}(v) = \sum_{u \in U} (\deg_{H'}(u))^2 + \sum_{w \in W} (\deg_{H'}(w))^2$$

$$\geq \sum_{u \in U} (\frac{|E'|}{|U|})^2 + \sum_{w \in W} (\frac{|E'|}{|W|})^2$$

(as $\sum_u \deg_{H'}(u) = \sum_w \deg_{H'}(w) = |E'|$ and each is minimized when the summands are equal.)

$$= |E'| \cdot \left(\bar{d} + |E'| / |W|\right).$$

By plugging in this bound in LHS above, we obtain $|E'| / |W| \leq \beta - \bar{d}$, finalizing the proof. ◀

3 A Simpler Proof of the Key Property of an EDCS

In this section we provide a much simpler proof of the key property that an EDCS contains an almost (3/2)-approximate matching. This lemma was previously used in [13, 14, 2]. Our proof is self-contained to this section, and for general graphs, our new proof even improves the dependence of β on parameter λ from $1/\lambda^3$ to (roughly) $1/\lambda^2$, thus allowing for an even sparser EDCS.

The proof contains two steps. We first give a simple and streamlined proof that an EDCS contains a (3/2)-approximate matching in bipartite graphs. Our proof in this part is similar to [13] but instead of modeling matchings as flows and using cut-flow duality, we directly work with matchings by using Hall's theorem. The main part of the proof however is to extend this result to general graphs. For this, we give a simple reduction that extends the result on bipartite graphs to general graphs by taking advantage of the "robust" nature of EDCS. This allows us to bypass the complicated arguments in [14] specific to non-bipartite graphs and to obtain the result directly from the one for bipartite graphs (the paper of [14] explicitly acknowledges the complexity of the proof and asks for a more "natural" approach).

A Slightly Simpler Proof for Bipartite Graphs

Our new proof should be compared to Lemma 2 in Section 4.1 of the Arxiv version of [13].

▶ **Lemma 3.** *Let $G(L,R,E)$ be any bipartite graph and $\varepsilon < 1/2$ be a parameter. For $\lambda \leq \frac{\varepsilon}{4}$, $\beta \geq 2\lambda^{-1}$, and $\beta^- \geq (1-\lambda)\cdot\beta$, in any subgraph $H := \text{EDCS}(G,\beta,\beta^-)$, $\mu(G) \leq \left(\frac{3}{2}+\varepsilon\right)\cdot\mu(H)$.*

Proof. Fix any $H := \text{EDCS}(G,\beta,\beta^-)$ and let A be any of its witness sets in extended Hall's marriage theorem of Proposition 3 and $B := N_H(A)$. Without loss of generality, let us assume A is a subset of L. Define $\overline{A} := L \setminus A$, $\overline{B} := R \setminus B$ (see Figure 1). By Proposition 3,

$$|\overline{A}| + |B| = n - (|A| - |B|) \leq n - (n - \mu(H)) = \mu(H). \tag{1}$$

On the other hand, since G has a matching of size $\mu(G)$, we need to have a matching M of size $(\mu(G) - \mu(H))$ between A and $\overline{B}$ as otherwise by Proposition 3, A would be a witness set in G that implies the maximum matching of G is smaller than $\mu(G)$ (to see why the set of edges between A and $\overline{B}$ is a matching simply apply Proposition 3 to a subgraph of G containing only a maximum matching of G). Let $S \subseteq A \cup \overline{B}$ be the end points of this matching (see Figure 1). As edges in M are all missing from H, by Property (P2) of EDCS H, we have that,

$$\sum_{v\in S} \deg_H(v) = \sum_{(u,v)\in M} (\deg_H(u) + \deg_H(v)) \geq (\mu(G) - \mu(H)) \cdot \beta^-. \tag{2}$$

Consequently, as $|S| = 2(\mu(G) - \mu(H))$, the average degree of S is $\geq \beta^-/2$. As such, by Proposition 5, the average degree of of $N_H(S)$ (from S) is at most $\beta - \beta^-/2 \leq (1+\lambda)\beta/2$. Finally, note that $N_H(S) \subseteq \overline{A} \cup B$ as there are no edges between A and $\overline{B}$ in H, and hence by Eq (1), $|N_H(S)| \leq \mu(H)$. By double counting the number of edges between S and $N_H(S)$, i.e., $E_H(S)$:

$$|E_H(S)| \geq |S| \cdot \beta^-/2 \geq 2(\mu(G) - \mu(H)) \cdot \beta^-/2,$$
$$|E_H(S)| \leq |N_H(S)|\,(1+\lambda)\beta/2 \leq \mu(H) \cdot (1+\lambda)\beta/2.$$

This implies that,

$$2\mu(G) \leq 2\mu(H) + \mu(H) \cdot (1+\lambda)(\beta/2) \cdot (2/\beta^-) \leq 3\mu(H) \cdot \frac{1+\lambda}{1-\lambda} \leq 3\mu(H)(1+\varepsilon).$$

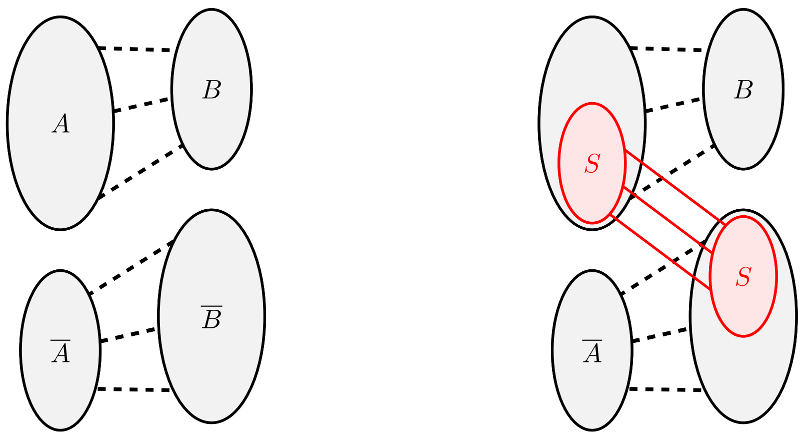

(a) A and $B := N_H(A)$ form a Hall's theorem witness set in EDCS H and $\left|\overline{A} \cup B\right| \leq \mu(H)$.

(b) There is a matching of size $\mu(G) - \mu(H)$ between A and $\overline{B}$ (i.e., the set S) in $G \setminus H$.

Figure 1 The partitioning of vertices used in the proof of Lemma 3.

Reorganizing the terms above finalizes the proof. ◀

A Much Simpler Proof for Non-bipartite Graphs

Our new proof in this part should be compared to Lemma 5.1 on page 699 in [14]: see Appendix B of their paper for the full proof, as well Section 4 for an additional auxiliary claim needed.

▶ **Lemma 4.** *Let $G(V, E)$ be any graph and $\varepsilon < 1/2$ be a parameter. For $\lambda \leq \frac{\varepsilon}{32}$, $\beta \geq 8\lambda^{-2}\log(1/\lambda)$, and $\beta^- \geq (1-\lambda)\cdot\beta$, in any subgraph $H := \mathrm{EDCS}(G, \beta, \beta^-)$, $\mu(G) \leq \left(\frac{3}{2} + \varepsilon\right)\cdot\mu(H)$.*

Proof. The proof is based on the probabilistic method and Lovasz Local Lemma. Let $M^\star$ be a maximum matching of size $\mu(G)$ in G. Consider the following randomly chosen bipartite subgraph $\widetilde{G}(L, R, \widetilde{E})$ of G with respect to $M^\star$, where $L \cup R = V$:

- For any edge $(u, v) \in M^\star$, with probability $1/2$, u belongs to L and v belongs to R, and with probability $1/2$, the opposite (the choices between different edges of $M^\star$ are independent).
- For any vertex $v \in V$ not matched by $M^\star$, we assign v to L or R uniformly at random (again, the choices are independent across vertices).
- The set of edges in $\widetilde{E}$ are all edges in E with one end point in L and the other one in R.

Define $\widetilde{H} := H \cap \widetilde{G}$. We argue that as H is an EDCS for G, $\widetilde{H}$ also remains an EDCS for $\widetilde{G}$ with non-zero probability. Formally,

▶ Claim 1. $\widetilde{H}$ is an $\mathrm{EDCS}(\widetilde{G}, \widetilde{\beta}, \widetilde{\beta}^-)$ for $\widetilde{\beta} = (1+4\lambda)\beta/2$ and $\widetilde{\beta}^- = (1-5\lambda)\beta^-/2$ with probability strictly larger than zero (over the randomness of $\widetilde{G}$).

Before we prove Claim 1, we argue why it implies Lemma 4. Let $\widetilde{G}$ be chosen such that $\widetilde{H}$ is an $\mathrm{EDCS}(\widetilde{G}, \widetilde{\beta}, \widetilde{\beta}^-)$ for parameters in Claim 1 (by Claim 1, such a choice of $\widetilde{G}$ always exist). By construction of $\widetilde{G}$, $M^\star \subseteq \widetilde{E}$ and hence $\mu(\widetilde{G}) = \mu(G)$. On the other hand, $\widetilde{G}$ is now a bipartite graph and $\widetilde{H}$ is its EDCS with appropriate parameters. We can hence apply Lemma 3 and obtain that $\mu(\widetilde{G}) \leq (3/2+\varepsilon)\mu(\widetilde{H})$. As $\widetilde{H} \subseteq H$, $\mu(\widetilde{H}) \leq \mu(H)$, and hence

$(\mu(\widetilde{G})=)\mu(G) \leq (3/2+\varepsilon)\mu(H)$, proving the assertion in the lemma statement. It thus only remains to prove Claim 1.

Proof of Claim 1. Fix any vertex $v \in V$, let $d_v := \deg_H(v)$ and $N_H(v) := \{u_1, \ldots, u_{d_v}\}$ be the neighbors of v in H. Let us assume v is chosen in L in $\widetilde{G}$ (the other case is symmetric). Hence, degree of v in $\widetilde{H}$ is exactly equal to the number of vertices in $N_H(v)$ that are chosen in R. As such, by construction of $\widetilde{G}$, $\mathbb{E}\left[\deg_{\widetilde{H}}(v)\right] = d_v/2$ ($+1$ iff v is incident on $M^\star \cap H$). Moreover, if two vertices u_i, u_j in $N_H(v)$ are matched by $M^\star$, then exactly one of them appears as a neighbor to v in $\widetilde{H}$ and otherwise the choices are independent. Hence, by Chernoff bound (Proposition 1),

$$\Pr\left(\left|\deg_{\widetilde{H}}(v) - d_v/2\right| \geq \lambda \cdot \beta\right) \leq \exp\left(-\frac{2\lambda^2 \cdot \beta^2}{\beta}\right) \leq \exp\left(-4\log\beta\right) \leq \frac{1}{\beta^4}.$$
$$\text{(as } \beta \geq 8\lambda^{-2}\log(1/\lambda) \text{ and hence } \beta \geq 2\lambda^{-2}\cdot\log\beta)$$

Define $\mathcal{E}_v$ as the event that $\left|\deg_{\widetilde{H}}(v) - d_v/2\right| \geq \lambda \cdot \beta$. Note that $\mathcal{E}_v$ depends only on the choice of vertices in $N_H(v)$ and hence can depend on at most β^2 other events $\mathcal{E}_u$ for vertices u which are neighbors to $N_H(v)$ (recall that for all $u \in V$, $\deg_H(u) \leq \beta$ in H by Property (P1) of EDCS). As such, we can apply Lovasz Local Lemma (Proposition 2) to argue that with probability strictly more than zero, $\cap_{v\in V}\overline{\mathcal{E}_v}$ happens. In the following, we condition on this event and argue that in this case, $\widetilde{H}$ is an EDCS of $\widetilde{G}$ with appropriate parameters. To do this, we only need to prove that both Property (P1) and Property (P2) hold for the EDCS $\widetilde{H}$ (with the choice of $\widetilde{\beta}$ and $\widetilde{\beta}^-$).

We first prove Property (P1) of EDCS $\widetilde{H}$. Let (u,v) be any edge in $\widetilde{H}$. By events $\overline{\mathcal{E}_v}$ and $\overline{\mathcal{E}_u}$,

$$\deg_{\widetilde{H}}(u) + \deg_{\widetilde{H}}(v) \leq \frac{1}{2}\cdot(\deg_H(u)+\deg_H(v)) + 2\lambda\beta \leq \beta/2 + 2\lambda\beta = (1+4\lambda)\cdot\beta/2,$$

where the second inequality is by Property (P1) of EDCS H as (u,v) belongs to H as well.

We now prove Property (P2) of EDCS $\widetilde{H}$. Let (u,v) be any edge in $\widetilde{G}\setminus\widetilde{H}$. Again, by $\overline{\mathcal{E}_v}$ and $\overline{\mathcal{E}_u}$,

$$\deg_{\widetilde{H}}(u) + \deg_{\widetilde{H}}(v) \geq \frac{1}{2}\cdot(\deg_H(u)+\deg_H(v)) - 2\lambda\beta \geq \beta^-/2 - 2\lambda(1-\lambda)\beta^- \geq (1-5\lambda)\cdot\beta/2,$$

where the second inequality is by Property (P2) of EDCS H as $(u,v) \in G \setminus H$. ◀ Claim 1

◀

Lemma 4 now follows immediately from Claim 1 as argued above. ◀ Lemma 4

◀

4 One-Way Communication Complexity of Matching

In the one-way communication model, Alice and Bob are given graphs $G_A(V,E_A)$ and $G_B(V,E_B)$, respectively, and the goal is for Alice to send a small message to Bob such that Bob can output a large approximate matching in $E_A \cup E_B$. In this section, we show that if Alice communicates an appropriate EDCS of G_A, then Bob is able to output an almost (3/2)-approximate matching.

▶ **Theorem 5** (Formalizing Result 1). *There exists a deterministic poly-time one-way communication protocol that given any $\varepsilon > 0$, computes a $(3/2+\varepsilon)$-approximation to maximum matching using $O(\frac{n\cdot\log(1/\varepsilon)}{\varepsilon^2})$ communication from Alice to Bob.*

Theorem 5 is based on the following protocol:

> A one-way communication protocol for maximum matching.
>
> 1. Alice sends $H := \text{EDCS}(G_A, \beta, \beta - 1)$ for $\beta := 32 \cdot \varepsilon^{-2} \cdot \log(1/\varepsilon)$ to Bob.
> 2. Bob computes a maximum matching in $H \cup G_B$ and outputs it as the solution.

By Proposition 4, the EDCS H computed by Alice always exists and can be found in polynomial time. Moreover, by Property (P1) of EDCS H, the total number of edges (and hence the message size) sent by Alice is $O(n\beta)$. We now prove the correctness of the protocol which concludes the proof of Theorem 5.

▶ **Lemma 6.** $\mu(G_A \cup G_B) \leq (3/2 + \varepsilon) \cdot \mu(H \cup G_B)$.

Proof. Let $M^\star$ be a maximum matching in $G_A \cup G_B$ and $M^\star_A$ and $M^\star_B$ be its edges in G_A and G_B, respectively. Let $\widetilde{G} := G_A \cup M^\star_B$ and note that $\mu(\widetilde{G}) = \mu(G)$ simply because $M^\star$ belongs to $\widetilde{G}$. Define the following subgraph $\widetilde{H} \subseteq H \cup M^\star_B$ (and hence $\subseteq H \cup G_B$): $\widetilde{H}$ contains all edges in H and any edge $(u, v) \in M^\star_B$ such that $\deg_H(u) + \deg_H(v) \leq \beta$. In the following, we prove that $(\mu(G) =)\mu(\widetilde{G}) \leq (3/2 + \varepsilon) \cdot \mu(\widetilde{H})$, which finalizes the proof as $\mu(\widetilde{H}) \leq \mu(H \cup G_B)$.

We show that $\widetilde{H}$ is an $\text{EDCS}(\widetilde{G}, \beta + 2, \beta - 1)$ and apply Lemma 4 to argue that $\widetilde{H}$ contains a (3/2)-approximate matching of $\widetilde{G}$. We prove the EDCS properties of $\widetilde{H}$ using the fact that for $v \in V$, $\deg_{\widetilde{H}}(v) \in \{\deg_H(v), \deg_H(v) + 1\}$ as $\widetilde{H}$ is obtained by adding a matching ($\subseteq M^\star_B$) to H.

- Property (P1) of EDCS $\widetilde{H}$: For an edge $(u, v) \in \widetilde{H}$,

$$\text{if } (u,v) \in H \text{ then:} \qquad \deg_{\widetilde{H}}(u) + \deg_{\widetilde{H}}(v) \leq \deg_H(u) + \deg_H(v) + 2 \leq \beta + 2,$$
(by Property (P1) of EDCS H of G_A)

$$\text{if } (u,v) \in M^\star_B \text{ then:} \qquad \deg_{\widetilde{H}}(u) + \deg_{\widetilde{H}}(v) \leq \deg_H(u) + \deg_H(v) + 2 \leq \beta + 2.$$
(as $(u, v) \in M^\star_B$ is inserted to $\widetilde{H}$ iff $\deg_H(u) + \deg_H(v) \leq \beta$)

- Property (P2) of EDCS $\widetilde{H}$: For an edge $(u, v) \in \widetilde{G} \setminus \widetilde{H}$,

$$\text{if } (u,v) \in G_A \setminus H \text{ then:} \qquad \deg_{\widetilde{H}}(u) + \deg_{\widetilde{H}}(v) \geq \deg_H(u) + \deg_H(v) \geq \beta - 1,$$
(by Property (P2) of EDCS H of G_A)

$$\text{if } (u,v) \in M^\star_B \setminus \widetilde{H} \text{ then:} \qquad \deg_{\widetilde{H}}(u) + \deg_{\widetilde{H}}(v) \geq \deg_H(u) + \deg_H(v) > \beta.$$
(as $(u, v) \in M^\star_B$ is not inserted to $\widetilde{H}$ iff $\deg_H(u) + \deg_H(v) > \beta$)

As such, $\widetilde{H}$ is an $\text{EDCS}(\widetilde{G}, \beta + 2, \beta - 1)$. By Lemma 4 and the choice of parameter β, we obtain that $\mu(\widetilde{G}) \leq (3/2 + \varepsilon) \cdot \mu(\widetilde{H})$, finalizing the proof. ◀

5 The Stochastic Matching Problem

Recall that in the stochastic matching problem, the goal is to compute a bounded-degree subgraph H of a given graph G, such that $\mathbb{E}[\mu(H_p)]$ is a good approximation of $\mathbb{E}[\mu(G_p)]$, where G_p is a realization of G (i.e a subgraph where every edge is sampled with probability p), and $H_p = H \cap G_p$. In this section, we formalize Result 2 by proving the following theorem.

▶ **Theorem 7** (Formalizing Result 2). *There exists a deterministic poly-time algorithm that given a graph $G(V, E)$ and parameters $\varepsilon, p > 0$ with $\varepsilon < 1/4$, computes a subgraph*

$H(V, E_H)$ of G with maximum degree $O(\frac{\log(1/\varepsilon p)}{\varepsilon^2 \cdot p})$ such that the ratio of the expected size of a maximum matching in realizations of G to realizations of H is at most $(3/2+\varepsilon)$, i.e., $\mathbb{E}[\mu(G_p)] \leq (3/2+\varepsilon) \cdot \mathbb{E}[\mu(H_p)]$.

We note that while in Theorem 7, we state the bound in expectation, the same result also holds with high probability as long as $\mu(G) = \omega(1/p)$ (i.e., just barely more than a constant), by concentration of maximum matching size in edge-sampled subgraphs (see, e.g. [2], Lemma 3.1). The algorithm in Theorem 7 simply computes an EDCS of the input graph as follows:

<u>An algorithm for the stochastic matching problem.</u>

Output the subgraph $H := \text{EDCS}(G, \beta, \beta - 1)$ for $\beta := \frac{C \log(1/\varepsilon p)}{\varepsilon^2 p}$, for large enough constant C.

By Proposition 4, the EDCS H in the above algorithm always exists and can be found in polynomial time. Moreover, by Property (P1) of EDCS H, the total number of edges in this subgraph is $O(n\beta)$. We now prove the bound on the approximation ratio which concludes the proof of Theorem 7 (by re-parametrizing ε to be a constant factor smaller).

▶ **Lemma 8.** *Let $H_p := H \cap G_p$ denote a realization of H; then $\mathbb{E}[\mu(G_p)] \leq (3/2 + O(\varepsilon)) \cdot \mathbb{E}[\mu(H_p)]$ where the randomness is taken over the realization G_p of G.*

Suppose first that H_p were an EDCS of G_p; we would be immediately done in this case as we could have applied Lemma 4 directly and prove Lemma 8. Unfortunately, however, this might not be the case. Instead, we exhibit subgraphs $\widetilde{H}_p \subseteq H_p$ and $\widetilde{G}_p \subseteq G_p$ with the following properties:

1. $\mathbb{E}[\mu(G_p)] \leq (1+\varepsilon)\, \mathbb{E}\left[\mu(\widetilde{G}_p)\right]$, where the expectation is taken over realizations G_p.
2. $\widetilde{H}_p$ is an $\text{EDCS}(G, (1+\varepsilon)p \cdot \beta, (1-2\varepsilon)p \cdot \beta)$ for $\widetilde{G}_p$.

Showing these properties concludes the proof of Lemma 8, as for the EDCS in item (2) above, we have $\frac{(1+\varepsilon)p\cdot\beta}{(1-2\varepsilon)\cdot p\beta} = 1 + O(\varepsilon)$, so by Lemma 4 we get that $\mu(\widetilde{G}_p) \leq (3/2 + O(\varepsilon)) \cdot \mu(\widetilde{H}_p)$. Combining this with item (1) then concludes $\mathbb{E}[\mu(G_p)] \leq (1+\varepsilon) \cdot (3/2+\varepsilon)\, \mathbb{E}[\mu(H_p)]$.

It now remains to exhibit $\widetilde{H}_p$ and $\widetilde{G}_p$ that satisfy the main properties stated above. Note that for any vertex $v \in V$, we have $\mathbb{E}\left[\deg_{H_p}(v)\right] = p \cdot \deg_H(v)$ by definition of a realization G_p (and hence H_p). We now want to separate out vertices that deviate significantly from this expectation.

▶ **Definition 9.** Let $V^+ \subseteq V$ contain all vertices v for which $\deg_{H_p}(v) > p \cdot \deg_H(v) + \varepsilon p\beta/2$. Similarly, let V^- contain all vertices v such that $\deg_{H_p}(v) < p \cdot \deg_H(v) - \varepsilon p\beta/2$ OR there exists an edge $(v, w) \in H$ such that $w \in V^+$, i.e., if v is neighbor to V^+.

▶ Claim 2. $\mathbb{E}[|V^+|] \leq \varepsilon^7 p^7 \mu(G)$ and $\mathbb{E}[|V^-|] \leq \varepsilon^4 p^4 \mu(G)$, where the expectation is over the realization G_p of G. As we a result we also have $\mathbb{E}[|V^+| + |V^-|] \leq \varepsilon^3 p^3 \mu(G)$.

Before proving this claim, let us consider why it completes the larger proof.

Proof of Lemma 8 (assuming Claim 2). To prove Lemma 8 it is enough to show the existence of subgraphs $\widetilde{G}_p$ and $\widetilde{H}_p$ that satisfy the properties above. We define $\widetilde{G}_p$ as follows: the vertex set is V and the edge-set is the same as G_p, except we remove all edges incident to V^+ and all edges $(u, v) \notin H$ that are incident to V^-. We define $\widetilde{H}_p$ to be the subgraph of H_p induced by the vertex set $V \setminus V^+$, that is, $\widetilde{H}_p$ contains all edges of H_p except those incident to V^+; see Figure 2.

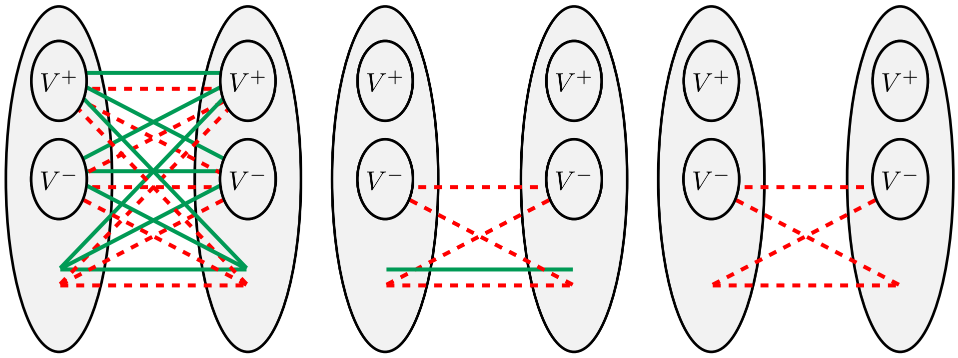

(a) Realized Graph G_p. **(b)** Subgraph $\widetilde{G}_p \subseteq G_p$. **(c)** Subgraph $\widetilde{H}_p \subseteq H_p$.

Figure 2 Illustration of the sets V^+, V^- and the subgraphs $\widetilde{G}_p$ and $\widetilde{H}_p$ in the proof of Lemma 8 on a bipartite graph G. Here, (green) solid lines denote the edges of G_p that appear in each subgraph and (red) dashed lines denote the edges of H_p.

For item (1), note that $\widetilde{G}_p$ differs from G_p by vertices in $V^+ \cup V^-$, so $\mu(\widetilde{G}_p) \geq \mu(G_p) - |V^+| - |V^-|$. It is also clear that $\mathbb{E}[\mu(G_p)] \geq p \cdot \mu(G)$ (as each edge in G is sampled w.p. p in G_p). By Claim 2,

$$\mathbb{E}\left[\mu(\widetilde{G}_p)\right] \geq \mathbb{E}[\mu(G_p)] - \mathbb{E}\left[|V^+| - |V^-|\right] \geq \mathbb{E}[\mu(G_p)] - p^3\varepsilon^3\mu(G) \geq (1-\varepsilon^3)\,\mathbb{E}[\mu(G_p)]\,.$$

The above equation then implies the desired $\mathbb{E}[\mu(G_p)] \leq (1+\varepsilon)\,\mathbb{E}\left[\mu(\widetilde{G}_p)\right]$.

For item (2), let us verify Property (P1) and Property (P2) for EDCS $\widetilde{H}_p$ of $\widetilde{G}_p$. Neither $\widetilde{H}_p$ nor $\widetilde{G}_p$ have any edge incident on V^+ and hence we can ignore these vertices entirely. Thus, for all vertices v we have $\deg_{\widetilde{H}_p}(v) \leq p \cdot \deg_H(v) + \varepsilon p\beta/2$, and for all $v \notin V^-$ we have $\deg_{\widetilde{H}_p}(v) \geq p \cdot \deg_H(v) - \varepsilon p\beta/2$. Moreover, recall that $\widetilde{G}_p \setminus \widetilde{H}_p$ contains no edges incident to V^-. As such,

- Property (P1) of EDCS $\widetilde{H}_p$: For an edge $(u,v) \in \widetilde{H}_p$,

$$\deg_{\widetilde{H}_p}(u) + \deg_{\widetilde{H}_p}(v) \leq p \cdot \deg_H(u) + p \cdot \deg_H(v) + \varepsilon p\beta \leq (1+\varepsilon)p\beta.$$

(by Property (P1) of EDCS H of G)

- Property (P2) of EDCS $\widetilde{H}_p$: For any edge $(u,v) \in \widetilde{G}_p \setminus \widetilde{H}_p$, we have $u, v \notin V^-$ so:

$$\deg_{\widetilde{H}_p}(u) + \deg_{\widetilde{H}_p}(v) \geq p \cdot \deg_H(u) + p \cdot \deg_H(v) - \varepsilon p\beta \geq (1-2\varepsilon)p\beta.$$

(by Property (P2) of EDCS H of G)

This concludes the proof of Lemma 8 (assuming Claim 2). ◀ Lemma 8

◀

All that remains is to prove Claim 2.

Proof of Claim 2. Let us start by bounding the size of V^+. Consider any vertex $v \in V$. We know that $\deg_H(v) \leq \beta$. Each edge then has probability p of appearing in H_p, so $\mathbb{E}\left[\deg_{H_p}(v)\right] = p \cdot \deg_H(v) \leq p\beta$. By the multiplicative Chernoff bound in Proposition 1 with $\lambda = p\beta$:

$$\Pr[v \in V^+] = \Pr[\deg_{H_p}(v) \geq p{\cdot}\deg_H(v) + \varepsilon p\beta/2] \leq e^{-O(\varepsilon^2 p\beta)} \leq e^{-O(\log(\varepsilon^{-1}p^{-1}))} \leq K^{-2}\varepsilon^{10}p^{10},$$

where K is a large constant and the last two inequalities follow from the fact that we set $\beta := \frac{C \log(1/\varepsilon p)}{\varepsilon^2 p}$, for large enough constant C. (Note that since constant C is in the exponent, we can easily set C large enough to achieve the final probability with a constant $K > C$.) This probability bound shows that $\mathbb{E}[|V^+|] \leq nK^{-2}\varepsilon^{10}p^{10}$, but that is not quite good enough since we want a dependence on $\mu(G)$ instead of on n. To achieve this, we observe that the total number of edges in H is at most $\beta\mu(G)$: the reason is that G has a vertex cover of size at most $2\mu(G)$, and all vertices in H have degree at most β (by Property (P1) of EDCS H). There are thus at most $2\beta\mu(G)$ vertices that have non-zero degree in H, each of which has at most a $\varepsilon^{10}p^{10}$ probability of being in V^+; all vertices with zero degree in H are clearly not in V^+ by definition. We thus have $\mathbb{E}[|V^+|] \leq 2\beta\mu(G) \cdot K^{-2}\varepsilon^{10}p^{10} \leq K^{-1}\varepsilon^7 p^7 \mu(G)$, where in the last inequality we use that $K > C$.

Let us now consider V^-. First let us bound the number of vertices $v \in V^-$ for which $\deg_{H_p}(v) < p \cdot \deg_H(v) - \varepsilon p\beta/2$. By an analogous argument to the one above, we have that the expected number of such vertices is at most $\varepsilon^7 p^7 \mu(G)$. A vertex can also end up in V^- because it has a neighbor in V^+ in H. But each vertex in H has degree at most β so we have

$$\mathbb{E}\left[|V^-|\right] \leq \varepsilon^7 p^7 \mu(G) + \beta\, \mathbb{E}\left[|V^+|\right] \leq \varepsilon^4 p^4 \mu(G),$$

where the last inequality again uses that $K > C$. ◀

▶ Remark. Interestingly, our result in Theorem 7 continues to hold as it is even when the edges sampled in realizations of G_p are only $\Theta(1/p)$-wise independent, by simply using a Chernoff bound for bounded-independence random variables (see, e.g. [31]) in the proof of Claim 2. Allowing correlation in the process of edge sampling is highly relevant to the main application of this problem to the kidney exchange setting (see [15]). To our knowledge, our algorithm is the first to work with such a little amount of independence between the edges.

6 A Fault-Tolerant Subgraph for Matching

In the fault-tolerant matching problem, we are given a graph $G(V, E)$ and an integer $f \geq 1$, and our goal is to compute a subgraph H of G, named an f-tolerant subgraph, such that for any subset $F \subseteq E$ of size f, $H \setminus F$ contains an approximate maximum matching of $G \setminus F$. We show that,

▶ **Theorem 10** (Formalizing Result 3). *There exists a deterministic poly-time algorithm that given any $\varepsilon > 0$ and integer $f \geq 1$, computes a $(3/2 + \varepsilon)$-approximate f-tolerant subgraph H of any given graph G with $O(\varepsilon^{-2} \cdot (n \log(1/\varepsilon) + f))$ edges.*

The algorithm in Theorem 10 simply computes an EDCS of the input graph as follows:

> An algorithm for the fault-tolerant matching problem.
>
> **1.** Define $\mu_{\min} := \min_F \left(\mu(G \setminus F)\right)$, where F is taken over all subsets of E with size f.
>
> **2.** Output $H := \text{EDCS}(G, \beta, \beta - 1)$ for $\beta := \frac{C \cdot f}{\varepsilon^2 \cdot \mu_{\min}} + \frac{C \cdot \log(1/\varepsilon)}{\varepsilon^2}$ for a constant $C > 0$.

By Proposition 4, the EDCS H in the above algorithm always exists and can be found in polynomial time. The above algorithm as stated however is not a polynomial time algorithm because it is not clear how to compute the quantity $\mu_{\min}$. Nevertheless, for simplicity, we work with the above algorithm throughout this section, and at the end show how to fix this problem and obtain a poly-time algorithm. We start by proving that the subgraph H only has $O(f + n)$ edges.

▶ **Lemma 11.** *The total number of edges in H is $O(\frac{f}{\varepsilon^2} + n \cdot \frac{\log(1/\varepsilon)}{\varepsilon^2})$.*

Proof. Let $F^\star$ be a subset of E with size f such that $\mu_{\min} = \mu(G \setminus F^\star)$. Let $M^\star$ be a maximum matching of size $\mu_{\min}$ in $G \setminus F^\star$. Note that $V(M^\star)$ is a vertex cover for $G \setminus F^\star$. This means that all edges in G except for f of them are incident on $V(M^\star)$. As no vertex in the EDCS H can have degree more than β by Property (P1) of EDCS, the degree of vertices in $V(M^\star)$ in $E \setminus F^\star$ is at most β. This implies that:

$$|E_H| \leq |V(M^\star)| \cdot \beta + |F^\star| \leq 2\mu_{\min} \cdot \left(\frac{C \cdot f}{\varepsilon^2 \cdot \mu_{\min}} + \frac{C \cdot \log(1/\varepsilon)}{\varepsilon^2}\right) + f$$
$$= O(\frac{f}{\varepsilon^2} + n \cdot \frac{\log(1/\varepsilon)}{\varepsilon^2}),$$

finalizing the proof. ◀

We now prove the correctness of the algorithm in the following lemma.

▶ **Lemma 12.** *Fix any subset $F \subseteq E$ of size f and define $G_F := G \setminus F$ and $H_F := H \setminus F$. Then, $\mu(G_F) \leq (3/2 + O(\varepsilon)) \cdot \mu(H_F)$.*

We first need some definitions. We say that a vertex $v \in V$ is *bad* iff $\deg_{H_F}(v) < \deg_H(v) - \varepsilon\beta$, i.e., at least $\varepsilon\beta$ edges incident on v (in H) are deleted by F. We use B_F to denote the set of bad vertices with respect to F, and bound $|B_F|$ in the following claim.

▶ Claim 3. Number of bad vertices in H_F is at most $|B_F| \leq \varepsilon \cdot \mu(G_F)$.

Proof. Any deleted edge can decrease the degrees of exactly two vertices. Any vertex becomes bad iff at least $\varepsilon\beta$ edges incident on it from H_F are removed. As such, $|B_F| \leq \frac{2f}{\varepsilon \cdot \beta} \leq \frac{2f \cdot \varepsilon^2 \cdot \mu_{\min}}{\varepsilon \cdot C \cdot f} \leq \varepsilon \cdot \mu(G_F)$, for sufficiently large $C > 0$, and since $\mu(G_F) \geq \mu_{\min}$ by definition. ◀ Claim 3

◀

Proof of Lemma 12. Define a subgraph $\widetilde{G}_F \subseteq G_F$ as follows: $V(\widetilde{G}_F) = V(G_F)$ $(= V(G))$ and edges in $\widetilde{G}_F$ are all edges in G_F except that we remove any edge $(u, v) \in G_F$ such that $(u, v) \notin H_F$ and either of u or v is a bad vertex. We prove that $\mu(\widetilde{G}_F)$ is at least $(1 - \varepsilon)$ fraction of $\mu(G_F)$, and moreover, H_F is an EDCS of $\widetilde{G}_F$ with appropriate parameters. We can then apply Lemma 4 to obtain that $\mu(G_F) \leq (1+2\varepsilon)\mu(\widetilde{G}_F) \leq (1+\varepsilon) \cdot (3/2+O(\varepsilon))\mu(H_F)$, finalizing the proof.

We first prove the bound on $\mu(\widetilde{G}_F)$. Fix any maximum matching M in G_F. It can have at most $|B_F|$ edges incident on vertices of B_F. Hence, even if we remove all edges incident on B_F, the size of this matching would be at least $\mu(G_F) - \varepsilon \cdot \mu(G_F)$, by the bound of $|B_F| \leq \varepsilon \cdot \mu(G_F)$ in Claim 3. However, this matching belongs to $\widetilde{G}_F$ entirely by the definition of this subgraph, and hence we have, $\mu(G_F) \leq (1 + 2\varepsilon)\mu(\widetilde{G}_F)$.

We now prove that H_F is an EDCS$(\widetilde{G}_F, \beta, (1 - 2\varepsilon)\beta - 1)$ of $\widetilde{G}_F$. It suffices to prove the two properties of EDCS for H_F using the fact that $\deg_{H_F}(v) \in [\deg_H(v) - \varepsilon\beta, \deg_H(v)]$ for vertices in $V \setminus B_F$, and that all edges incident on B_F in $\widetilde{G}_F$ also belong to H_F.

- Property (P1) of EDCS H_F of $\widetilde{G}_F$: For any edge $(u, v) \in H_F$:

$$\deg_{H_F}(u) + \deg_{H_F}(v) \leq \deg_H(u) + \deg_H(v) \leq \beta.$$
(by Property (P1) of EDCS H of G)

- Property (P2) of EDCS H_F of $\widetilde{G}_F$: For any edge $(u,v) \in \widetilde{G}_F \setminus H_F$ both $u, v \in V \setminus B_F$ and so:

$$\deg_{H_F}(u) + \deg_{H_F}(v) \geq \deg_H(u) + \deg_H(v) - 2\varepsilon\beta \geq (1-2\varepsilon)\beta - 1.$$
(by Property (P2) of EDCS H of G as (u,v) is missing from H)

As such, H_F is an EDCS$(\widetilde{G}_F, \beta, (1-2\varepsilon)\beta - 1)$ of $\widetilde{G}_F$ and by the lower bound on value of β in the algorithm (the second term in definition of β), we can apply Lemma 4, and obtain that $\mu(\widetilde{G}_F) \leq (3/2 + O(\varepsilon)) \cdot \mu(H_F)$, finalizing the proof. ◀

Theorem 10 now follows from Lemmas 11 and 12 by re-parametrizing ε to a sufficiently smaller constant factor of ε (by picking the integer C large enough) modulo the fact that the algorithm designed in this section is not a polynomial time algorithm. To make the algorithm polynomial time, we only need to make a simple modification: instead of finding $\mu_{\min}$ explicitly, we find the smallest value of β (by searching over all n possible choices of β, or by doing a binary search) such that the EDCS H has at least $\frac{2 \cdot C \cdot f}{\varepsilon^2} + \frac{n \cdot C \cdot \log(1/\varepsilon)}{\varepsilon^2}$ many edges. By the proof of Lemma 11, any EDCS of G can have at most $2\mu_{\min} \cdot \beta + f$ edges. This implies that the chosen $\beta \geq \frac{C \cdot f}{\varepsilon^2 \cdot \mu_{\min}} + \frac{C \cdot \log(1/\varepsilon)}{\varepsilon^2}$ as needed in the algorithm. This concludes the proof, as by definition of β, H has $O(\frac{C \cdot f}{\varepsilon^2} + \frac{n \cdot C \cdot \log(1/\varepsilon)}{\varepsilon^2})$ many edges, and hence satisfies the sparsity requirements of Theorem 10.

Acknowledgements

Sepehr Assadi is grateful to his advisor Sanjeev Khanna for many helpful discussions, and to Soheil Behnezhad for sharing a write-up of [9].

References

1. Noga Alon and Joel H Spencer. *The probabilistic method.* John Wiley & Sons, 2004.
2. Sepehr Assadi, MohammadHossein Bateni, Aaron Bernstein, Vahab S. Mirrokni, and Cliff Stein. Coresets meet EDCS: algorithms for matching and vertex cover on massive graphs. *CoRR*, abs/1711.03076. To appear in SODA 2019, 2017.
3. Sepehr Assadi, Sanjeev Khanna, and Yang Li. The stochastic matching problem with (very) few queries. In *Proceedings of the 2016 ACM Conference on Economics and Computation, EC '16, Maastricht, The Netherlands, July 24-28, 2016*, pages 43–60, 2016.
4. Sepehr Assadi, Sanjeev Khanna, and Yang Li. On estimating maximum matching size in graph streams. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2017, Barcelona, Spain, Hotel Porta Fira, January 16-19*, pages 1723–1742, 2017.
5. Sepehr Assadi, Sanjeev Khanna, and Yang Li. The stochastic matching problem: Beating half with a non-adaptive algorithm. In *Proceedings of the 2017 ACM Conference on Economics and Computation, EC '17, Cambridge, MA, USA, June 26-30, 2017*, pages 99–116, 2017.
6. Baruch Awerbuch. Complexity of network synchronization. *J. ACM*, 32(4):804–823, 1985.
7. Surender Baswana, Keerti Choudhary, and Liam Roditty. Fault tolerant subgraph for single source reachability: generic and optimal. In *Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2016, Cambridge, MA, USA, June 18-21, 2016*, pages 509–518, 2016.
8. Joshua D. Batson, Daniel A. Spielman, and Nikhil Srivastava. Twice-ramanujan sparsifiers. In *Proceedings of the 41st Annual ACM Symposium on Theory of Computing, STOC 2009, Bethesda, MD, USA, May 31 - June 2, 2009*, pages 255–262, 2009.

9 Soheil Behnezhad, Alireza Farhadi, MohammadTaghi Hajiaghayi, and Nima Reyhani. Stochastic matching with few queries: New algorithms and tools. In *Manuscript. To appear in SODA 2019.*, 2018.

10 Soheil Behnezhad and Nima Reyhani. Almost optimal stochastic weighted matching with few queries. In *Proceedings of the 2018 ACM Conference on Economics and Computation, Ithaca, NY, USA, June 18-22, 2018*, pages 235–249, 2018.

11 András A. Benczúr and David R. Karger. Approximating s-t minimum cuts in $\tilde{O}(n^2)$ time. In *Proceedings of the Twenty-Eighth Annual ACM Symposium on the Theory of Computing, Philadelphia, Pennsylvania, USA, May 22-24, 1996*, pages 47–55, 1996.

12 Claude Berge. *The theory of graphs*. Courier Corporation, 1962.

13 Aaron Bernstein and Cliff Stein. Fully dynamic matching in bipartite graphs. In *Automata, Languages, and Programming - 42nd International Colloquium, ICALP 2015, Kyoto, Japan, July 6-10, 2015, Proceedings, Part I*, pages 167–179, 2015.

14 Aaron Bernstein and Cliff Stein. Faster fully dynamic matchings with small approximation ratios. In *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2016, Arlington, VA, USA, January 10-12, 2016*, pages 692–711, 2016.

15 Avrim Blum, John P. Dickerson, Nika Haghtalab, Ariel D. Procaccia, Tuomas Sandholm, and Ankit Sharma. Ignorance is almost bliss: Near-optimal stochastic matching with few queries. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation, EC '15, Portland, OR, USA, June 15-19, 2015*, pages 325–342, 2015.

16 Greg Bodwin, Michael Dinitz, Merav Parter, and Virginia Vassilevska Williams. Optimal vertex fault tolerant spanners (for fixed stretch). In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2018, New Orleans, LA, USA, January 7-10, 2018*, pages 1884–1900, 2018.

17 Greg Bodwin, Fabrizio Grandoni, Merav Parter, and Virginia Vassilevska Williams. Preserving distances in very faulty graphs. In *44th International Colloquium on Automata, Languages, and Programming, ICALP 2017, July 10-14, 2017, Warsaw, Poland*, pages 73:1–73:14, 2017.

18 Béla Bollobás, Don Coppersmith, and Michael Elkin. Sparse distance preservers and additive spanners. In *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms, January 12-14, 2003, Baltimore, Maryland, USA.*, pages 414–423, 2003.

19 Shiri Chechik, Michael Langberg, David Peleg, and Liam Roditty. Fault-tolerant spanners for general graphs. In *Proceedings of the 41st Annual ACM Symposium on Theory of Computing, STOC 2009, Bethesda, MD, USA, May 31 - June 2, 2009*, pages 435–444, 2009.

20 Don Coppersmith and Michael Elkin. Sparse sourcewise and pairwise distance preservers. *SIAM J. Discrete Math.*, 20(2):463–501, 2006.

21 Paul Erdős and László Lovász. Problems and results on 3-chromatic hypergraphs and some related questions. In *COLLOQUIA MATHEMATICA SOCIETATIS JANOS BOLYAI 10. INFINITE AND FINITE SETS, KESZTHELY (HUNGARY)*. Citeseer, 1973.

22 Wai Shing Fung, Ramesh Hariharan, Nicholas J. A. Harvey, and Debmalya Panigrahi. A general framework for graph sparsification. In *Proceedings of the 43rd ACM Symposium on Theory of Computing, STOC 2011, San Jose, CA, USA, 6-8 June 2011*, pages 71–80, 2011.

23 Ashish Goel, Michael Kapralov, and Sanjeev Khanna. On the communication and streaming complexity of maximum bipartite matching. In *Proceedings of the Twenty-third Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '12, pages 468–485. SIAM, 2012. URL: http://dl.acm.org/citation.cfm?id=2095116.2095157.

24 Philip Hall. On representatives of subsets. *Journal of the London Mathematical Society*, 1(1):26–30, 1935.

25 Michael Kapralov. Better bounds for matchings in the streaming model. In *Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2013, New Orleans, Louisiana, USA, January 6-8, 2013*, pages 1679–1697, 2013. URL: http://dx.doi.org/10.1137/1.9781611973105.121, doi:10.1137/1.9781611973105.121.
26 David R. Karger. Random sampling in cut, flow, and network design problems. In *Proceedings of the Twenty-Sixth Annual ACM Symposium on Theory of Computing, 23-25 May 1994, Montréal, Québec, Canada*, pages 648–657, 1994.
27 Euiwoong Lee and Sahil Singla. Maximum matching in the online batch-arrival model. In *Integer Programming and Combinatorial Optimization - 19th International Conference, IPCO 2017, Waterloo, ON, Canada, June 26-28, 2017, Proceedings*, pages 355–367, 2017.
28 David Peleg. As good as it gets: Competitive fault tolerance in network structures. In *Stabilization, Safety, and Security of Distributed Systems, 11th International Symposium, SSS 2009, Lyon, France, November 3-6, 2009. Proceedings*, pages 35–46, 2009.
29 David Peleg and Alejandro A. Schäffer. Graph spanners. *Journal of Graph Theory*, 13(1):99–116, 1989.
30 Imre Z Ruzsa and Endre Szemerédi. Triple systems with no six points carrying three triangles. *Combinatorics (Keszthely, 1976), Coll. Math. Soc. J. Bolyai*, 18:939–945, 1978.
31 Jeanette P. Schmidt, Alan Siegel, and Aravind Srinivasan. Chernoff-hoeffding bounds for applications with limited independence. *SIAM J. Discrete Math.*, 8(2):223–250, 1995.
32 Daniel A. Spielman and Shang-Hua Teng. Spectral sparsification of graphs. *SIAM J. Comput.*, 40(4):981–1025, 2011.
33 William T Tutte. The factorization of linear graphs. *Journal of the London Mathematical Society*, 1(2):107–111, 1947.
34 Yutaro Yamaguchi and Takanori Maehara. Stochastic packing integer programs with few queries. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2018, New Orleans, LA, USA, January 7-10, 2018*, pages 293–310, 2018.

A Missing Details and Proofs

A.1 Proof of Proposition 4

We give the proof of this proposition following the argument of [2], which itself was based on [14].

Proof. We give a polynomial local search algorithm for constructing an EDCS H of the graph G which also implies the existence of H. The algorithm is as follows. Start with empty graph H. While there exists an edge in H or $G \setminus H$ that violates Property (P1) or Property (P2) of EDCS, respectively, fix this edge by removing it from H for the former or inserting it to H for the latter.

We prove that this algorithm terminates after polynomial number of steps which implies both the existence of the EDCS as well as give a polynomial time algorithm for computing it. We define the following potential function Φ for this task:

$$\Phi_1(H) := (\beta - 1/2) \cdot \sum_{u \in V(H)} \deg_H(u), \qquad \Phi_2(H) := \sum_{(u,v) \in E(H)} \left(\deg_H(u) + \deg_H(v)\right),$$
$$\Phi(H) := \Phi_1(H) - \Phi_2(H).$$

We claim that after fixing each edge in H in the algorithm, Φ increases by at least 1. Since max-value of Φ is $O(n \cdot \beta^2)$, this implies that this procedure terminates in $O(n \cdot \beta^2)$ steps.

Let (u, v) be the fixed edge at this step, H_1 be the subgraph before fixing the edge (u, v), and H_2 be the resulting subgraph. Suppose first that the edge (u, v) was violating

Property (P1) of EDCS. As the only change is in the degrees of vertices u and v, Φ_1 decreases by $(2\beta - 1)$. On the other hand, $\deg_{H_1}(u) + \deg_{H_1}(v) \geq \beta + 1$ originally (as (u, v) was violating Property (P1) of EDCS) and hence after removing (u, v), Φ_2 also decreases by $\beta + 1$. Additionally, for each neighbor w of u and v in H_2, after removing the edge (u, v), $\deg_{H_2}(w)$ decreases by one. As there are at least $\deg_{H_2}(u) + \deg_{H_2}(v) = \deg_{H_1}(u) + \deg_{H_1}(v) - 2 \geq \beta - 1$ choices for w, this means that in total, Φ_2 decreases by at least $(\beta + 1) + (\beta - 1) = 2\beta$. As a result, in this case $\Phi = \Phi_1 - \Phi_2$ increases by at least 1 after fixing the edge (u, v).

Now suppose that the edge (u, v) was violating Property (P2) of EDCS instead. In this case, degree of vertices u and v both increase by one, hence Φ_1 increases by $2\beta - 1$. Additionally, since edge (u, v) was violating Property (P2) we have $\deg_{H_1}(u) + \deg_{H_1}(v) \leq \beta^- - 1$, so the addition of edge (u, v) decreases Φ_2 by at most $\deg_{H_2}(u) + \deg_{H_2}(v) = \deg_{H_1}(u) + \deg_{H_1}(v) + 2 \leq \beta^- + 1$. Moreover, for each neighbor w of u and v, after adding the edge (u, v), $\deg_{H_2}(w)$ increases by one and since there are at most $\deg_{H_1}(u) + \deg_{H_1}(v) \leq \beta^- - 1$ choices for w, Φ_2 decreases in total by at most $(\beta^- + 1) + (\beta^- - 1) = 2\beta^-$. Since $\beta^- \leq \beta - 1$, we have that Φ increases by at least $(2\beta - 1) - (2\beta^-) \geq 1$ after fixing the edge (u, v), finalizing the proof. ◀

A.2 Optimality of the $(3/2)$-Approximation Ratio in Result 3

Our argument is a simple modification of the one in [23] for proving a lower bound on the one-way communication complexity of approximating matching and is provided for the sake of completeness.

Let $G_1(V_1, E_1)$ be a graph on N vertices such that its edges can be partitioned into $t := N^{\Omega(1/\log\log N)}$ *induced* matchings $M_1, \ldots, M_t$ of size $(1 - \delta)N/4$ for arbitrarily small constant $\delta > 0$. These graphs are referred to as (r, t)-Ruzsa-Szemerédi graphs [30] ((r, t)-RS graphs for short) and have been studied extensively in the literature (see [4, 23] for more details). In particular, the existence of such graphs with parameters mentioned above is proven in [23].

Let $G(V, E)$ be a graph with $n = 2N$ vertices consisting of $G_1(V_1, E_1)$ plus N additional vertices U that are connected via a perfect matching M_U to V_1. In the following, we prove that any f-fault tolerant subgraph H of G that achieves a $(3/2 - \varepsilon)$-approximation for some constant $\varepsilon > 0$ when $f = \Theta(n)$ requires $n^{1+\Omega(1/\log\log n)} = \omega(f)$ edges.

Suppose towards a contradiction that H contains $o(m)$ edges where m is the number of edges in the graph G. As edges in G_1 are partitioned into induced matchings $M_1, \ldots, M_t$, it means that there exists some induced matching M_i such that only $o(1)$ fraction of its edges belong to H. Let the set of deleted edge F be only the set of edges in the perfect matching between U and V_1, namely, M_U, which are incident to $V(M_i)$. The number of deleted edges is $O(n)$ and after deletion, M_U has size $N - (1 - \delta)N/2 = (1 + \delta)N/2$. As such, $\mu(G \setminus F) \geq (1 + \delta)N/2 + (1 - \delta)N/4 \geq 3N/4$, by picking the remainder of the matching M_U and the induced matching M_i (which is not incident on remainder of M_U by construction). However, we argue that $\mu(H \setminus F) \leq (1 + \delta)N/2 + o(N)$, simply because only $o(N)$ edges of M_i belong H and all other matchings are incident to the remaining edges of M_U (we can assume remaining edges of M_U belong to *any* maximum matching of $H \setminus F$ because they are incident on degree one vertices). As such, $\mu(H \setminus F) < (2/3 + 2\delta)\mu(G \setminus F)$. By picking $\delta < \varepsilon/4$, we obtain that H is not a $(3/2 - \varepsilon)$-approximate f-fault tolerant subgraph of G.

A.3 Other Standard Algorithms for Fault-Tolerant Matching

Since the goal in fault-tolerant matching is to prepare for adversarial deletions, the most natural approach seem to be adding many different matchings by a finding maximum matching in G, adding it to the subgraph H, deleting it from G, and repeating until we have $O(f+n)$ edges. A similar approach would be to let H be a maximum b-matching, with b set appropriately to end up with $O(f+n)$ edges. We show a lower bound of 2 on the approximation ratio of these approaches.

Consider the following approach first: find a maximum matching M in G, add all the edges of M to the fault-tolerant subgraph H, remove all the edges of M from G, and repeat until the graph contains $C(f+n)$ edges for some large constant C. For $f = n/5$, we present a graph G where this approach yields a graph H where $\mu(H) = \mu(G)/2$. The graph is bipartite and the vertex set is partitioned into 5 sets X, Y, Y', Z, Z', each of size $n/5$. There is an edge in G from every vertex in X to every vertex in Y or Z, and there are also exactly $n/5$ vertex-disjoint edges from Y to Y', and similarly from Z to Z'; those are all the edge of G. The fault tolerant algorithm might choose the following subgraph H: H contains a perfect matching from Y to Y' and from Z to Z', as well as many edges from X to Y, but no edges from X to Z. (The algorithm can end up with such an H by first choosing the maximum matching in G that consists of the edges from Y to Y' and from Z to Z'; then for all future iterations the maximum matching size is only $|X| = n/5$, so the algorithm might always pick a maximum matching that only contains edges between X and Y.) Now consider the set of failures F which consists of the $n/5$ edges from Z to Z'. It is clear that $\mu(G \setminus F) = 2n/5$, while $\mu(H \setminus F) = n/5$. Note also that allowing H to contain more than $O(n+f)$ edges would still not allow this approach to break through the 2-approximation: in this lower-bound instance, even if H was allowed to have up to $n^2/100$ edges, H might still not contain any edges from X to Z, and so we would still have $\mu(H \setminus F) = n/5 = \mu(G \setminus F)/2$.

The other natural approach is to let H contain the edges of a maximum b-matching in G, where b is set to a value for which the resulting b-matching still contains $\Theta(f+n)$ edges. The lower-bound graph G is exactly the same as above, though in this case we use $f = 2n/5$. The maximum b-matching H might then contain the edges from Y to Y' and Z to Z', a single matching of size $n/5$ from X to Z, and then many edges from X to Y. It is easy to see that this is a maximum b-matching. Now consider the following set F of deletions: F contains all edges from Z to Z', as well as the $n/5$ edges in H from X to Z. It is easy to see that we once again have $\mu(H) = n/5$ and $\mu(G) = 2n/5$. Also as above, setting B to be very large and allowing H to have $n^2/100$ edges would still not break through the 2-approximation.