

# 1 Dynamic Matching with Better-than-2 Approximation in 2 Polylogarithmic Update Time\*

3 SAYAN BHATTACHARYA<sup>†</sup>, University of Warwick, United Kingdom

4 PETER KISS<sup>‡</sup>, University of Warwick, United Kingdom

5 THATCHAPHOL SARANURAK<sup>§</sup>, University of Michigan, United States

6 DAVID WAJC<sup>¶</sup>, Technion – Israel Institute of Technology, Israel

7 We present dynamic algorithms with *polylogarithmic* update time for estimating the size of the maximum  
8 matching of a graph undergoing edge insertions and deletions with approximation ratio *strictly better than 2*.  
9 Specifically, we obtain a  $1 + \frac{1}{\sqrt{2}} + \epsilon \approx 1.707 + \epsilon$  approximation in bipartite graphs and a  $1.973 + \epsilon$  approximation  
10 in general graphs. We thus answer in the affirmative the value version of the major open question repeatedly  
11 asked in the dynamic graph algorithms literature. Our randomized algorithms' approximation and worst-case  
12 update time bounds both hold w.h.p. against adaptive adversaries.

13 Our algorithms are based on simulating new two-pass streaming matching algorithms in the dynamic  
14 setting. Our key new idea is to invoke the recent sublinear-time matching algorithm of Behnezhad (FOCS'21)  
15 in a white-box manner to efficiently simulate the second pass of our streaming algorithms, while bypassing  
16 the well-known vertex-update barrier.

17 **CCS Concepts:** • **Theory of computation** → Streaming, sublinear and near linear time algorithms; Approximation  
18 algorithms analysis; **Dynamic graph algorithms**.

19 Additional Key Words and Phrases: dynamic algorithms, approximate matching

## 20 **ACM Reference Format:**

21 Sayan Bhattacharya, Peter Kiss, Thatchaphol Saranurak, and David Wajc. 2018. Dynamic Matching with  
22 Better-than-2 Approximation in Polylogarithmic Update Time. *J. ACM* 37, 4, Article 111 (August 2018), 32 pages.  
23 <https://doi.org/XXXXXX.XXXXXXX>

## 24 **1 INTRODUCTION**

25 The maximum matching problem is a cornerstone of combinatorial optimization and theoretical  
26 computer science more broadly. (We recommend [41] for a brief history of this problem.) The  
27 study of this problem and its extensions has contributed foundational advances and concepts

28 \*A preliminary version of this paper appeared in the Proceedings of the 34th Annual ACM-SIAM Symposium on Discrete  
29 Algorithms (SODA), 2023 [31]. This full version contains (in Sections 1.3.2 and 6) a discussion of the numerous developments  
30 in the area following our conference publication, as well as more detailed follow-up questions, including concerning  
31 applications of our adversarially-robust fully-dynamic almost-maximal matching (AMM) algorithms.

32 <sup>†</sup>Supported by Engineering and Physical Sciences Research Council, UK (EPSRC) Grant EP/S03353X/1.

33 <sup>‡</sup>Work was partially conducted while the author was visiting Max-Planck-Institut für Informatik.

34 <sup>§</sup>Supported by NSF grant CCF-22238138.

35 <sup>¶</sup>Supported in part by a Taub Family “Leaders in Science & Technology” Fellowship. Work done in part while the author  
36 was at Stanford University.

37 Authors’ addresses: Sayan Bhattacharya, S.Bhattacharya@warwick.ac.uk, University of Warwick, United Kingdom; Peter  
38 Kiss, peter.kiss@warwick.ac.uk, University of Warwick, United Kingdom; Thatchaphol Saranurak, University of Michigan,  
39 United States, thsa@umich.edu; David Wajc, Technion – Israel Institute of Technology, Israel, david.wajc@gmail.com.

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

45 © 2018 Association for Computing Machinery.

46 0004-5411/2018/8-ART111 \$15.00

47 <https://doi.org/XXXXXX.XXXXXXX>

50 to the theory of computing, from the introduction of the primal-dual method [67], impact on  
 51 polyhedral combinatorics [44], and the advocacy for polynomial-time computability as the measure  
 52 of efficiency (in static settings) [45].

53 The maximum matching problem has also been intensely studied in *dynamic* settings. Here,  
 54 the graph undergoes edge *updates* (insertions and deletions), and we wish to approximate the  
 55 maximum matching, while spending little computation time between updates, referred to as *update*  
 56 *time*. Polynomial update time is trivial to achieve by running exact static algorithms (e.g., [45])  
 57 after each update. However, intuitively, such minor changes to the graph should allow for much  
 58 faster algorithms, with possibly even exponentially smaller, *polylogarithmic* update times.

59 The first sublinear (i.e.,  $o(m) = o(n^2)$ ) update time dynamic matching algorithm was given 15  
 60 years ago by Sankowski [79], who used fast dynamic matrix inversion to maintain the maximum  
 61 matching size in update time  $O(n^{1.495})$ , recently improved to  $O(n^{1.407})$  [34]. Unfortunately, a number  
 62 of fine-grained complexity results rule out fast, and even sublinear-in- $n$  update time [1, 2, 40, 58, 66]  
 63 for (exact) maximum matching size estimation. This motivates the wealth of work on computing  
 64 *approximate* matchings dynamically.

65 The first polylogarithmic update time dynamic matching algorithm is due to an influential work  
 66 of Onak and Rubinfeld [74], who gave a (large) constant approximation in polylog update time. This  
 67 was later improved by Baswana et al. [11] to a 2-approximation in logarithmic update time, later  
 68 improved to constant time by Solomon [80]. Numerous other algorithms achieving a 2- or  $(2 + \epsilon)$ -  
 69 approximation in polylog update time were subsequently developed, with expected amortized  
 70 update time improved to worst-case w.h.p.,<sup>1</sup> and oblivious randomized algorithms improved to  
 71 adversarially-robust ones, and then to deterministic ones [4, 14, 21, 26, 28, 33, 35, 36, 62, 82].<sup>2</sup>

72 A complementary line of work studied better-than-two-approximate dynamic matching, pro-  
 73 viding a number of small polynomial (even sublinear in  $n$ ) update times for approximation ra-  
 74 tios below the natural bound of 2 achieved by inclusion-wise maximal matchings. This includes  
 75  $(1 + \epsilon)$ -approximate algorithms with  $O_\epsilon(\sqrt{m}) = O_\epsilon(n)$  update time [57, 77],  $(\frac{3}{2} + \epsilon)$ -approximate  
 76 algorithms with  $O_\epsilon(\sqrt[4]{m}) = O_\epsilon(\sqrt{n})$  update time [22, 23, 56, 62] and a number of tradeoffs between  
 77 approximation in the range  $(3/2, 2)$  and sublinear-in- $n$  polynomial update times [16, 17, 28, 78, 82].<sup>3</sup>

78 This state of affairs leaves open a key question, repeatedly raised in the literature [16, 17, 23, 26,  
 79 35, 68, 82], first posed by Onak and Rubinfeld in their aforementioned groundbreaking work [74]:

80     *How small can [approximation factors] be made with polylogarithmic update time?*  
 81     *[...] Can the approximation constant be made smaller than 2 for maximum matching?*

## 83 1.1 Our Results

84 We resolve the question of polylogarithmic update time better-than-two-approximate dynamic  
 85 matching algorithms in the affirmative, *for the value version of the problem*. That is, letting  $\mu(G)$   
 86 denote the maximum matching size in  $G$ , we maintain an estimate  $v$  that is  $\alpha < 2$  approximate, i.e.,  
 87 it satisfies  $v \leq \mu(G) \leq \alpha \cdot v$  at every point in time. Our main result is the following.

88  
 89  
 90  
 91  
 92 <sup>1</sup>An algorithm has *amortized* update time  $f(n)$  if every sequence of  $t$  updates starting from an empty graph takes at most  
 93  $t \cdot f(n)$  update time. If each operation takes at most  $f(n)$  time, it has *worst-case* update time  $f(n)$ .

94 <sup>2</sup>An algorithm *works against an adaptive adversary* if its guarantees hold even when future updates depend on the algorithm's  
 95 previous output. We also say that such an algorithm is *adversarially robust*, or *robust* for short. The importance of robustness  
 96 for *static* applications has motivated a recent concentrated effort to design robust dynamic algorithms for myriad problems  
 97 (see, e.g., discussions in [19, 28, 38, 49, 71, 82]).

98 <sup>3</sup>Throughout the paper, we use  $O_\epsilon(\cdot)$  to suppress  $\text{poly}(1/\epsilon)$  factors and  $\tilde{O}(\cdot)$  to suppress  $\text{poly}(\log n)$  factors.

99  
 100 **THEOREM 1.1.** *For every  $\epsilon \in (0, 1)$ , there exists a randomized  $(1.973 + \epsilon)$ -approximate dynamic  
 101 matching size estimation algorithm with  $\text{poly}(\log n, 1/\epsilon)$  worst-case update time. Both the  
 102 algorithm's approximation ratio and update time hold w.h.p against an adaptive adversary.*

103  
 104 For bipartite graphs, we obtain a stronger approximation guarantee of  $1 + \frac{1}{\sqrt{2}} + \epsilon \approx 1.707 + \epsilon$ .

105 **Secondary results.** Our approach is versatile, and yields the following generic reduction.

106  
 107 **THEOREM 1.2.** *For any  $\alpha > 1.5$ , a dynamic  $\alpha$ -approximate matching algorithm with update time  
 108  $t_u$  implies a dynamic  $\left(\alpha - \Omega\left((1 - 6\left(\frac{1}{\alpha} - \frac{1}{2}\right))^2\right)\right)$ -approximate matching size estimator with update  
 109 time  $\tilde{O}(t_u)$ .*

110  
 111 Very recently, Behnezhad and Khanna [16] presented new dynamic matching algorithms trading  
 112 off approximations  $\alpha \in (1.5, 2]$  and small polynomial update times. Applying Theorem 1.2 to their  
 113 algorithms, we obtain improved approximation for dynamic matching size estimation, within the  
 114 same update time up to polylog factors.

115 To obtain our main results, we design several two-pass *semi-streaming* algorithms (see Section 1.3),  
 116 including a deterministic  $(1 + 1/\sqrt{2} + \epsilon)$ -approximate algorithm on bipartite graphs. This matches  
 117 the prior state-of-the-art [63, 65] up to an  $\epsilon$  term, while removing the need for randomization.

## 118 1.2 Our Techniques

119  
 120 We take the following high-level approach to prove Theorem 1.1: (1) compute a maximal (and  
 121 hence 2-approximate) matching  $M_1$ , and (2) augment  $M_1$  if it is no better than 2-approximate, using  
 122 the myriad short augmenting paths  $M_1$  must have in this case. This approach is common in many  
 123 computational models, including the two-pass semi-streaming model (see Section 1.3). Implementing  
 124 this approach in a dynamic setting, however, faces several challenges. The first challenge if we want  
 125 robust algorithms with low worst-case update times is that no robust (near-)maximal matching  
 126 algorithms with worst-case  $\tilde{O}_\epsilon(1)$  update time are known. Of possible independent interest, we  
 127 resolve this first challenge in Section 5, by leveraging the robust fast matching sparsifiers of [82].

128 The more central challenge when trying to implement the above approach is that the search for  
 129 augmenting paths requires us to find many (disjoint) edges between matched and unmatched nodes  
 130 in  $M_1$ . In a (multi-pass) streaming setting, this can be done by computing a large  $(b\text{-})$ matching in the  
 131 bipartite graph induced by edges in  $V(M_1) \times \overline{V(M_1)}$ . In a dynamic setting, however, this requires us  
 132 to deal with *vertex updates*, which are notoriously challenging in the context of dynamic matching,  
 133 and all algorithms to date require reading all  $\Omega(n)$  edges of each updated vertex [68].

134 To overcome the above key challenge, we first note that we do not need to handle vertex updates  
 135 individually, but may instead process these in *batches* of  $\Theta(\epsilon n)$  vertex updates, building on the  
 136 periodic recomputation and sparsification techniques common in the literature (see Theorem 2.1).  
 137 Our main observation is that these batches of vertex updates, which need to be handled if we  
 138 wish to maintain the  $b$ -matchings from the second pass of our semi-streaming algorithms, can be  
 139 implemented in  $\tilde{O}_\epsilon(n)$  time using the sublinear-time algorithm of Behnezhad [12]. This leads to an  
 140 amortized  $\tilde{O}_\epsilon(n)/(\epsilon n) = \tilde{O}_\epsilon(1)$  additive overhead in the update time (easily deamortized), implying  
 141 our main result. This approach is versatile, and similarly underlies our secondary results.

## 142 1.3 Further Related Work

143 Having discussed the rich literature on the dynamic matching problem above, we do not elaborate  
 144 on it further here. We do, however, highlight some connections to the literature on matching in  
 145 other computational models that is closely related to our work.

**Streaming Matching.** In the (semi-)streaming model, an  $n$ -node graph is revealed in a stream, edge by edge, and we wish to compute a large matching using only (optimal)  $\tilde{O}(n)$  space. A line of work studying the problem of computing an approximately-maximum weighted matching [39, 47, 50, 53, 70, 76] culminated in a  $(2 + \epsilon)$ -approximation [53, 76]. For unweighted graphs, lower bounds are known [54, 60, 61], but it remains a major open question whether one can break the barrier of 2-approximation achievable by a trivial maximal matching algorithm. Striving for better approximation (and insights to break this barrier), several works designed algorithms using *multiple passes* over the stream [3, 3, 6, 46, 46, 48, 51, 52, 59, 70]. For 2 passes, the state-of-the-art approximation ratios are 1.857 [51], and  $1 + \frac{1}{\sqrt{2}} \approx 1.707$  for bipartite graphs using the randomized algorithms of [63, 65], with the best prior deterministic bound being  $\frac{12}{7} \approx 1.714$  [48].

**Sublinear-Time Matching.** Computation of large matchings in sublinear *time* has also been the subject of great interest. In regular bipartite graphs, a maximum matching can be computed in  $\tilde{O}(n)$  time [55]. In general graphs with bounded degrees, it was known how to achieve a  $(2 + \epsilon)$ -approximation in sublinear time [72, 73, 75, 83]. This was recently improved to an  $\tilde{O}(n)$  time algorithm for *any* general graph [12]. As discussed in Section 1.2, we use this latter algorithm in a white-box manner to obtain our main result.

**1.3.1 Concurrent work.** Independently and concurrently, Behnezhad [13] (in a work in the same conference) obtained the same main qualitative result as ours: a better-than-two-approximate polylogarithmic-xwtime dynamic matching size estimation algorithm. The basic approach to achieve this qualitative result is the same in both papers: Simulate the second pass of a two-pass streaming algorithm using the sublinear-time algorithm of [12], together with batched computation. The quantitative differences in the papers' approximation ratios are due to the two-pass streaming algorithms used—our new maximal-b-matching-based algorithms here, and an algorithm inspired by [64] in [13]. We note that [13] also achieves  $(3/2 - \Omega(1))$ -approximate size estimation algorithm in time  $O(\sqrt{n})$  (the best update times for  $(3/2 + \epsilon)$ -approximate explicit matching [22, 23, 56, 62]). This result also uses the high-level approach of batched computation using sublinear-time algorithms, building on a new characterization of tight examples for the  $3/2$ -approximate matching sparsifiers (EDCS) of [22].

**1.3.2 Subsequent work.** The conference versions of this work and the concurrent work of [13] have sparked an interest in the question of dynamic matching size maintenance. [10] show that our  $(1 + \frac{1}{\sqrt{2}} + \epsilon)$ -approximate two-pass streaming algorithm also works in general graphs, and show how to implement this in dynamic settings, using distributed algorithms. [30] show that  $(1 + \epsilon)$ -approximate size estimation can be maintained in truly sublinear-in- $n$  update time for any constant  $\epsilon > 0$ . This should be contrasted with a conditional lower bound of [69], who show that such an update time is unlikely if one wishes to *maintain* such an approximately maximum matching. Complementing this conditional lower bound, [7, 15] show conditional upper bounds. Specifically, they provide  $(1 + \epsilon)$ -approximate matching algorithms with sublinear-in- $n$  update time, provided some generalization of Rusza-Szemerédi graphs cannot be too dense. Proving or ruling out algorithms for maintenance of large matchings mirroring the recent results for dynamic size estimation is an exciting research direction. One direction to obtain such algorithmic results would be to obtain algorithms for dynamic *fractional* matchings, since the latter are known to be efficiently roundable dynamically, both in bipartite graphs [4, 28, 32, 82] and in general graphs [37, 43].

## 197 2 PRELIMINARIES

198 Our input is a graph  $G$  on  $n$  nodes  $V$ , with an initially empty edge set  $E$ , undergoing edge updates  
 199 (insertions and deletions). Our objective is to approximate the maximum matching size  $\mu(G)$   
 200 well, while spending little update time (computation between updates). In addition, we want  
 201 our algorithms to work in the strictest settings: they should have small *worst-case* update time  
 202 guarantees, and they should be (adversarially) *robust*, i.e., they should work against an adaptive  
 203 adversary, and thus their guarantees hold for any update sequence.

204 **Matching theory basics.** A *matching* is a vertex-disjoint subset of edges. A *maximal matching* is  
 205 an inclusionwise-maximal matching. A maximum matching is a matching of largest cardinality.  
 206 In a weighted graph with edge weights  $w_e \in \mathbb{R}$ , a maximum weight matching is a matching  $M$  of  
 207 largest total weight,  $w(M) := \sum_{e \in M} w_e$ . An *augmenting path*  $P$  with respect to a matching  $M$  is a  
 208 simple path starting and ending with distinct nodes unmatched in  $M$ , with the edges alternatingly  
 209 outside and inside  $M$ . Setting  $M \leftarrow M \oplus P$ , where  $\oplus$  denotes the symmetric difference, referred  
 210 to as *augmenting*  $M$  along  $P$ , increases the cardinality of  $M$  by one. A  $b$ -*matching* with capacities  
 211  $\{b_v\}_{v \in V}$  is a collection of *multi-edges*  $F$  of  $E$  (that is, edges of  $E$  may appear multiple times in  $F$ )  
 212 with no vertex  $v$  having more than  $b_v$  multi-edges in  $F$ . A *fractional matching*  $x : E \rightarrow \mathbb{R}_{\geq 0}$  assigns  
 213 non-negative values to edges so that each vertex  $v$  has *fractional degree*  $\sum_{e \ni v} x_e$  at most one. In  
 214 bipartite graphs, the existence of a fractional matching of size  $k$  implies the existence of an integral  
 215 matching of cardinality  $\lceil k \rceil$ . In general graphs, this fractional relaxation has a maximum integrality  
 216 gap of  $3/2$ , attained by a triangle graph with values  $x_e = 1/2$  for each edge  $e$ .

217 **Notation:** Let  $V(M)$  denote the set of all endpoints of edges in a matching  $M$ , and let  $\overline{V(M)} :=$   
 218  $V \setminus V(M)$ . For any disjoint vertex sets  $A, B \subseteq V$ , we let  $G[A, B]$  denote the bipartite subgraph  
 219 induced by the edges in  $G$  with one endpoint in  $A$  and another in  $B$ . Finally, for any subset of edges  
 220  $E' \subseteq E$ , we let  $G[E']$  denote the subgraph of  $G$  induced by  $E'$ .

### 222 2.1 Previous building blocks

223 A ubiquitous paradigm in the approximate dynamic matching literature is *periodic recomputation*,  
 224 introduced by Gupta and Peng [57]. This approach is particularly useful in conjunction with  
 225 sparsification techniques. We will use the vertex sparsification technique introduced by Assadi  
 226 et al. [8] in the context of stochastic optimization, and adapted to dynamic settings by Kiss [62].  
 227 Combined, these approaches yield the following “reduction” from dynamic matching algorithms  
 228 with *immediate* queries to ones with slower query time.

229 **Proposition 2.1.** *Let  $\epsilon \in (0, 1)$  and  $\alpha \geq 1$ . Suppose there exists an algorithm  $\mathcal{A}$  on a dynamic  
 230  $n$ -node graph  $G$  with update time  $t_u$ , that, provided  $\mu(G) \geq \epsilon \cdot n$ , supports  $t_q$ -time  $\alpha$ -approximate  
 231 size estimate queries w.h.p. Then, there is another algorithm  $\mathcal{A}'$  on  $G$  that always maintains an  
 232  $(\alpha + O(\epsilon))$ -approximate estimate  $v'$  in  $\tilde{O}_\epsilon(t_u + t_q/n)$  update time. Moreover if the update time of  $\mathcal{A}$  is  
 233 worst-case, so is that of  $\mathcal{A}'$ , and if  $\mathcal{A}$  works against an adaptive adversary, then so does  $\mathcal{A}'$ .*

234 The above proposition, implicit in prior work, serves as a useful abstraction, and so we provide a  
 235 proof of this proposition for completeness in Appendix A. As discussed in Section 1.2, this reduction  
 236 is one of the crucial ingredients that allows us to bypass the vertex-update barrier.

237 Another key ingredient we use is the sublinear-time (approximate) maximal matching algorithm  
 238 of Benhezad [12], whose guarantees are captured by the following proposition. (See Lemma 4.6  
 239 for a proof of a detailed description of that algorithm, adapted for a variant of this proposition.)

240 **Proposition 2.2.** *Let  $\epsilon \in (0, 1/2)$ . Using  $\tilde{O}_\epsilon(n)$  time and  $\tilde{O}_\epsilon(n)$  adjacency matrix queries w.h.p. in an  
 241  $n$ -node graph  $G$ , one can compute a value  $v$  which approximates  $\tilde{\mu}$ , the size of some maximal matching  
 242 in  $G$ , within additive error  $\epsilon n$ . Namely,  $\tilde{\mu} \geq v \geq \tilde{\mu} - \epsilon n$ .*

246 A simple combination of propositions 2.1 and 2.2 (with  $t_q = \tilde{O}_\epsilon(n)$ ) immediately yields (yet)  
 247 another  $\tilde{O}_\epsilon(1)$ -time  $(2 + \epsilon)$ -approximation algorithm. As we will show, these propositions are also  
 248 useful ingredients for breaking the barrier of 2-approximation within the same update time.  
 249

## 250 2.2 New algorithmic primitive: Robust Almost-Maximal Matchings

251 To make our algorithms robust against adaptive adversaries we need an algorithm for maintaining  
 252 *almost-maximal matchings* (AMM), which are defined as follows.  
 253

254 **Definition 2.3** ([77]). A matching  $M$  is an  $\epsilon$ -almost maximal matching ( $\epsilon$ -AMM) in graph  $G$  if  $M$  is  
 255 maximal in some subgraph obtained by removing at most  $\epsilon \cdot \mu(G)$  nodes of  $G$ .  
 256

257 **Observation 2.4.** If  $M$  is an  $\epsilon$ -AMM in  $G$ , then  $|M| \geq \frac{1}{2}(1 - \epsilon) \cdot \mu(G) = \left(\frac{1}{2} - \frac{\epsilon}{2}\right) \cdot \mu(G)$ .  
 258

259 Peleg and Solomon [77] showed how to deterministically maintain an  $\epsilon$ -AMM quickly in bounded-  
 260 arboricity (i.e., globally sparse) graphs. In Section 5 we show how to robustly maintain such  
 261 matchings quickly in arbitrary graphs, proving the following.  
 262

263 **Lemma 2.5.** For any  $\epsilon \in (0, 1/2)$ , there exists a robust dynamic algorithm that w.h.p. maintains an  
 264  $\epsilon$ -AMM in worst-case update time  $\tilde{O}_\epsilon(1)$ .  
 265

266 A well-known fact is that a maximal matching that is close to 2-approximate must admit many  
 267 length-three augmenting paths (see e.g., [64, Lemma 1]). Our interest in AMMs is in part motivated  
 268 by the following slight generalization of this fact, also presented in Appendix A for completeness.  
 269

270 **Proposition 2.6.** Let  $\epsilon > 0$  and  $c \in \mathbb{R}$  and let  $M$  be an  $\epsilon$ -AMM in  $G$  such that  $|M| \leq \left(\frac{1}{2} + c\right) \cdot \mu(G)$ .  
 271 Then  $M$  admits a collection of at least  $\left(\frac{1}{2} - 3c - \frac{7\epsilon}{2}\right) \cdot \mu(G)$  node-disjoint 3-augmenting paths.  
 272

## 273 3 ALGORITHMS ON BIPARTITE GRAPHS

274 In this section we illustrate our techniques for the special case of bipartite graphs, for which we  
 275 obtain an improved approximation ratio of  $1 + \frac{1}{\sqrt{2}} + \epsilon \approx 1.707 + \epsilon$ .  
 276

### 277 3.1 Two-Pass Streaming Algorithm

278 Here we present our deterministic two-pass streaming algorithm. We first compute an approximately-  
 279 maximal matching  $M_1$  from the first pass.<sup>4</sup> Then, in the second pass, we compute a maximal  $b$ -  
 280 matching  $M_2$  in the graph between matched and unmatched vertices, with capacities  $k$  and  $\lfloor k \cdot b \rfloor$ ,  
 281 respectively, where we set the parameters  $k \in \mathbb{N}$ ,  $\delta \in \mathbb{R}$  and  $b = 1/\delta \in \mathbb{R}$  later. Finally, we output  
 282 an estimate  $(1 - \delta) \cdot |M_1| + (\delta/k) \cdot |M_2|$ . Our pseudocode is given in Algorithm 1.  
 283

---

#### 284 Algorithm 1 Bipartite Two-Pass Streaming Algorithm

---

285 **Parameters:**  $\epsilon \in \mathbb{R}_{\geq 0}$ ,  $k \in \mathbb{N}$ ,  $\delta \in \mathbb{R}_{\geq 0}$  and  $b = 1/\delta \in \mathbb{R}_{\geq 0}$ .  
 286

287 1:  $M_1 \leftarrow (\epsilon/8)$ -AMM in  $G$  computed from first pass

288 2: assign each vertex  $v$  capacity  $b_v = \begin{cases} k & v \in V(M_1) \\ \lfloor k \cdot b \rfloor & v \notin V(M_1) \end{cases}$

289 3:  $M_2 \leftarrow$  maximal  $b$ -matching in  $G[V(M_1), \overline{V(M_1)}]$  computed from second pass

290 4: **Output**  $(1 - \delta) \cdot |M_1| + (\delta/k) \cdot |M_2|$ .

---

291 First, we prove that the above algorithm's output estimate corresponds to a matching in  $G$ .  
 292

293 <sup>4</sup>We suggest to the reader to think of  $M_1$  as a maximal matching (i.e.,  $\epsilon = 0$ ). We relax  $M_1$  to be an  $(\epsilon/8)$ -AMM since this  
 294 will be useful in our dynamic implementation that works against adaptive adversaries.

295 **Observation 3.1.** We have that  $\mu(G[M_1 \cup M_2]) \geq (1 - \delta) \cdot |M_1| + (\delta/k) \cdot |M_2|$ .

296 PROOF. Since  $G$  is bipartite, by the integrality of the bipartite fractional matching polytope, to  
 297 prove that  $G' := G[M_1 \cup M_2]$  contains a large matching witnessing the desired inequality, it suffices  
 298 to prove that  $G'$  contains a *fractional* matching  $\vec{x}$  of value  $\sum_e x_e = (1 - \delta) \cdot |M_1| + (\delta/k) \cdot |M_2|$ .  
 299 Indeed, such a fractional matching is obtained by assigning edge values  
 300

$$301 \quad x_e = \begin{cases} 1 - \delta & e \in M_1 \\ 302 \quad \delta/k & e \in M_2 \setminus M_1. \end{cases}$$

303 This is indeed a fractional matching, since each vertex  $v$  has bounded fractional degree,  $\sum_{e \ni v} x_e \leq 1$ :  
 304 every vertex  $v \in V(M_1)$  has one incident  $M_1$  edge and at most  $k$  many incident  $M_2$  edges, and so  
 305  $\sum_{e \ni v} x_e \leq (1 - \delta) + (\delta/k) \cdot k = 1$ , while every vertex  $v \notin V(M_1)$  has no incident  $M_1$  edge and has  
 306 at most  $k \cdot b$  incident  $M_2$  edges, and so  $\sum_{e \ni v} x_e \leq k \cdot b \cdot (\delta/k) = 1$ .  $\square$   
 307

308 By Lemma 3.1, Algorithm 1 outputs a valid estimate for the matching size,  $v \leq \mu(G)$ . It remains  
 309 to prove that  $v$  provides a good approximation of  $\mu(G)$ . For this, we require the following.

311 **Lemma 3.2.** Let  $M$  be a maximal  $b$ -matching in a bipartite graph  $G = (L \cup R, E)$ , with positive  
 312 integral capacities  $b_v = \ell$  for all  $v \in L$  and  $b_v = r$  for all  $v \in R$ . Then

$$314 \quad |M| \geq \mu(G) \cdot \frac{\ell \cdot r}{\ell + r}. \\ 315$$

316 PROOF. Fix a maximum matching  $M^*$  in  $G$ . Next, we define the subset of matched nodes in  $M^*$   
 317 that are also saturated in  $M$ . That is, if  $d_M(v)$  is  $v$ 's degree in  $M$ , we let  
 318

$$319 \quad L_{sat}^* := \{u \in L \cap V(M^*) \mid d_M(u) = \ell\} \\ 320 \quad R_{sat}^* := \{v \in R \cap V(M^*) \mid d_M(v) = r\}.$$

321 Let  $\alpha := |L_{sat}^*|/|M^*|$  and  $\beta := |R_{sat}^*|/|M^*|$  denote the fraction of  $M^*$  edges with a saturated  $L$  and  
 322  $R$  node, respectively. Since  $M$  is a maximal  $b$ -matching in  $G$ , each edge has at least one saturated  
 323 endpoint, and so  $\alpha + \beta \geq 1$ . By double counting the edges of  $M$ , relying on  $\alpha + \beta \geq 1$ , and noting  
 324 that  $\alpha \cdot r + (1 - \alpha) \cdot \ell$  attains its minimum of  $2\frac{\ell \cdot r}{\ell + r}$  at  $\alpha = \frac{r}{\ell + r}$ , we obtain the claimed inequality.

$$325 \quad |M| = \frac{1}{2} \left( \sum_{v \in L} d_M(v) + \sum_{v \in R} d_M(v) \right) \\ 326 \quad \geq \frac{1}{2} \left( \sum_{v \in L_{sat}^*} d_M(v) + \sum_{v \in R_{sat}^*} d_M(v) \right) \\ 327 \quad = \frac{1}{2} \cdot (|M^*| \cdot \alpha \cdot \ell + |M^*| \cdot \beta \cdot r) \\ 328 \quad \geq \frac{1}{2} \cdot \mu(G) \cdot (\alpha \cdot \ell + (1 - \alpha) \cdot r) \\ 329 \quad \geq \mu(G) \cdot \frac{\ell \cdot r}{\ell + r}. \quad \square$$

330 We are now ready to bound the approximation ratio of Algorithm 1.

331 **Lemma 3.3.** For any  $\epsilon \in (0, 1)$ , Algorithm 1 with  $b = 1 + \sqrt{2}$  and  $k \geq \frac{8}{\epsilon b}$  run on bipartite graph  $G$   
 332 computes a  $(1 + \frac{1}{\sqrt{2}} + \epsilon) \approx (1.707 + \epsilon)$ -approximation to  $\mu(G)$ .

344 PROOF. Fix a maximum matching  $M^*$  in  $G$ . Next, for  $i \in \{0, 1, 2\}$ , let  $M_i^*$  denote the edges of  $M^*$   
 345 with  $i$  endpoints matched in  $M_1$ . By definition, and since  $|M_1| = \frac{1}{2} \cdot |V(M_1)|$ , we have that  
 346

$$347 \quad |M_1| = |M_2^*| + (1/2) \cdot |M_1^*|. \quad (1)$$

348 Furthermore, since  $M_1$  is an  $\epsilon'$ -AMM in  $G$  for  $\epsilon' = \epsilon/8$ , we have that  $|M_0^*| \leq \epsilon' \cdot \mu(G)$ , since at most  
 349  $\epsilon' \cdot \mu(G)$  nodes of  $G$  must be removed from  $G$  to make  $M_1$  maximal, and at least one endpoint of  
 350 each  $M_0^*$  edge must be removed to achieve the same effect. But then, since  $M_0^*, M_1^*, M_2^*$  partition  $M^*$ ,  
 351 whose cardinality is  $|M^*| = \mu(G)$ , this implies that  
 352

$$353 \quad |M_1^*| + |M_2^*| \geq (1 - \epsilon') \cdot \mu(G). \quad (2)$$

354 Now, by Lemma 3.2, since  $M_1^*$  is a matching in graph  $G' := G[V(M_1), \overline{V(M_1)}]$  and  $k \geq \frac{1}{\epsilon' b}$ , we have  
 355

$$356 \quad |M_2| \geq \mu(G') \cdot \frac{k \cdot \lfloor kb \rfloor}{k + \lfloor kb \rfloor} \geq \frac{\lfloor kb \rfloor}{b+1} \cdot |M_1^*| \geq \frac{kb(1 - \epsilon')}{b+1} \cdot |M_1^*|. \quad (3)$$

359 Combining equations (1), (2) and (3), we obtain the following lower bound on our output estimate.  
 360

$$\begin{aligned} 361 \quad (1 - \delta) \cdot |M_1| + (\delta/k) \cdot |M_2| &\stackrel{(1),(3)}{\geq} (1 - \delta) \cdot (|M_2^*| + (1/2) \cdot |M_1^*|) + (\delta/k) \cdot \frac{kb(1 - \epsilon')}{b+1} \cdot |M_1^*| \\ 362 \quad &= (1 - 1/b) \cdot |M_2^*| + \left( \frac{1 - 1/b}{2} + \frac{1 - \epsilon'}{b+1} \right) \cdot |M_1^*| \\ 363 \quad &\geq (|M_1^*| + |M_2^*|) \cdot \min \left\{ 1 - 1/b, \frac{1 - 1/b}{2} + \frac{1 - \epsilon'}{b+1} \right\} \\ 364 \quad &\stackrel{(2)}{\geq} (1 - \epsilon') \cdot \mu(G) \cdot \min \left\{ 1 - 1/b, \frac{1 - 1/b}{2} + \frac{1 - \epsilon'}{b+1} \right\} \\ 365 \quad &\geq (1 - 2\epsilon') \cdot \mu(G) \cdot \min \left\{ 1 - 1/b, \frac{1 - 1/b}{2} + \frac{1}{b+1} \right\} \\ 366 \quad &= (1 - 2\epsilon') \cdot (2 - \sqrt{2}) \cdot \mu(G), \end{aligned}$$

374 where the last equality follows by our choice of  $b = 1 + \sqrt{2}$ .  
 375

376 Thus, combining with Lemma 3.1, and using that  $\epsilon' = \epsilon/8 < 1/8$ , we find that the output  
 377 matching size estimate  $v := (1 - \delta) \cdot |M_1| + (\delta/k) \cdot |M_2|$  is indeed a  $\left(1 + \frac{1}{\sqrt{2}} + \epsilon\right)$ -approximation,  
 378 since

$$379 \quad v \leq \mu(G) \leq v \cdot \left( \frac{1}{(2 - \sqrt{2}) \cdot (1 - 2\epsilon')} \right) \leq v \cdot \left( \left(1 + \frac{1}{\sqrt{2}}\right) \cdot (1 + 4\epsilon') \right) \leq v \cdot \left(1 + \frac{1}{\sqrt{2}} + \epsilon\right),$$

382 as desired. □

383

384 **Remark 3.4.** A direct extension of the tight example of [65] proves that this analysis is tight, up to  
 385 the exact dependence on  $\epsilon$ .

386 Lemma 3.1 implies a two-pass streaming algorithm for computing a  $\left(1 + \frac{1}{\sqrt{2}} + \epsilon\right)$ -approximate  
 387 maximum matching: simply store  $G[M_1 \cup M_2]$  and output a maximum matching in this subgraph  
 388 by the stream's end. The space used in the first and second passes are  $\tilde{O}(n)$  and  $\tilde{O}(nkb) = \tilde{O}(n/\epsilon)$ ,  
 389 respectively. More interestingly for our goals, we show in the next section that Lemma 3.3 can be  
 390 used to obtain a *dynamic* approximation of the same quality, in polylogarithmic update time.  
 391

### 393 3.2 Dynamic Algorithm

394 In this section, we show how to (approximately) implement Algorithm 1 in polylogarithmic update  
 395 time in a dynamic setting.

396 THEOREM 3.5. *Let  $\epsilon \in (0, 1)$ . There exists a robust dynamic algorithm  $\mathcal{A}$  with worst-case update time  
 397  $t_u = \tilde{O}_\epsilon(1)$  w.h.p. and query time  $t_q = \tilde{O}_\epsilon(n)$  that outputs w.h.p. a value  $v \in [\mu(G)/(1 + \frac{1}{\sqrt{2}} + \epsilon), \mu(G)]$ .  
 398 That is, it answers  $(1 + \frac{1}{\sqrt{2}} + \epsilon)$  approximate matching size estimate queries.*

400 PROOF. The dynamic algorithm  $\mathcal{A}$  is based on Lemma 1. Let  $\epsilon' = \epsilon/12$ . Throughout the updates,  
 401 Algorithm  $\mathcal{A}$  simply maintains an  $(\epsilon'/8)$ -AMM in  $G$ , denoted by  $M_1$ , invoking Lemma 2.5. This  
 402 immediately implies the desired update time of  $t_u = \tilde{O}_\epsilon(1)$ .

403 We now describe how Algorithm  $\mathcal{A}$  responds to a query about the maximum matching size.  
 404 To answer this query, the algorithm considers a new auxiliary graph  $G^* = (V^*, E^*)$ , which is  
 405 defined as follows. Set  $b = 1 + \sqrt{2}$ . For each  $u \in V(M_1)$ , create  $k := \lceil \frac{8}{\epsilon' b} \rceil$  copies of the node  $u$  in  
 406  $G^*$ . Next, for each  $v \in \overline{V(M_1)}$ , create  $\lfloor k \cdot b \rfloor$  copies of the node  $v$  in  $G^*$ . Finally, for every edge  
 407  $(u, v) \in G[V(M_1), \overline{V(M_1)}]$ , create an edge in  $G^*$  between every pair  $(u^*, v^*)$  of copies of the nodes  
 408  $u, v$ . Note that there is a one-to-one mapping between maximal matchings in the new graph  $G^*$  and  
 409 maximal  $b$ -matchings in  $G[V(M_1), \overline{V(M_1)}]$ .

410 We emphasize that our dynamic algorithm  $\mathcal{A}$  does not explicitly maintain the auxiliary graph  
 411  $G^*$ . When we receive a query about the maximum matching size in  $G$ , we explicitly construct only  
 412 the node-set  $V^*$  of  $G^*$ , based on the matching  $M_1$ . This takes only  $O_\epsilon(n)$  time since  $|V^*| = O_\epsilon(n)$ .  
 413 We can, however, access the edges of  $G^*$  by using adjacency matrix queries: there exists an edge  
 414  $(u^*, v^*)$  in  $G^*$  iff there exists an edge  $(u, v)$  between the corresponding nodes in  $G$ .

415 At this point, we invoke Proposition 2.2 with  $G = G^*$  and precision parameter  $\epsilon'' = (\epsilon')^3$ . This  
 416 gives us a value  $\psi$ , which is an estimate of  $|M_2|$ . We now return  $v := (1 - 1/b) \cdot |M_1| + (1/bk) \cdot \psi$  as  
 417 our estimate of  $\mu(G)$ . Clearly, this entire procedure for answering a query takes  $\tilde{O}_\epsilon(n)$  time. It now  
 418 remains to analyze the approximation ratio. Towards this end, we again appeal to Proposition 2.2.  
 419 This proposition asserts that the value  $\psi$  satisfies

$$420 |M_2| \geq \psi \geq |M_2| - \epsilon'' \cdot |V^*| \geq |M_2| - (\epsilon')^2 \cdot 16n \geq |M_2| - \epsilon' \cdot \mu(G), \quad (4)$$

421 Therefore, our estimate  $v$  satisfies that  $v' \geq v \geq v' - \epsilon \cdot \mu(G)$ , where

$$422 v' := (1 - 1/b) \cdot |M_1| + (1/bk) \cdot |M_2|. \quad (5)$$

423 Now, by Lemma 3.3 and our choice of  $k = \lceil \frac{8}{\epsilon' b} \rceil$  and  $b = 1 + \sqrt{2}$ , we have that  $v' \leq \mu(G) \leq$   
 424  $v' \cdot (1 + \frac{1}{\sqrt{2}} + \epsilon)$ , from which we obtain that  $v \leq v' \leq \mu(G)$  and moreover

$$425 \mu(G) \leq v' \cdot \left(1 + \frac{1}{\sqrt{2}} + \epsilon'\right) \leq (v + \epsilon' \cdot \mu(G)) \cdot \left(1 + \frac{1}{\sqrt{2}} + \epsilon'\right) \leq v \cdot \left(1 + \frac{1}{\sqrt{2}} + \epsilon'\right) + 3\epsilon' \cdot \mu(G).$$

426 Rearranging terms, and using that  $\epsilon' = \epsilon/16 \leq 1/12$ , we have that

$$427 v \leq \mu(G) \leq v \cdot \left(1 + \frac{1}{\sqrt{2}} + \epsilon'\right) / (1 - 3\epsilon) \leq v \cdot \left(1 + \frac{1}{\sqrt{2}} + \epsilon\right) \cdot (1 + 4\epsilon) \leq v \cdot \left(1 + \frac{1}{\sqrt{2}} + 16\epsilon'\right).$$

428 That is, since  $\epsilon' = \epsilon/16$ , the estimate  $v$  output after a query is  $(1 + \frac{1}{\sqrt{2}} + \epsilon)$ -approximate, w.h.p.  $\square$

429 Combining Theorem 3.5 and Proposition 2.1, we obtain our result for bipartite graphs.

430 THEOREM 3.6. *For any  $\epsilon \in (0, 1)$ , there exists a  $(1 + \frac{1}{\sqrt{2}} + \epsilon) \approx (1.707 + \epsilon)$ -approximate randomized  
 431 dynamic bipartite matching size algorithm with  $\tilde{O}_\epsilon(1)$ -update time. The algorithm's approximation  
 432 ratio holds w.h.p. against an adaptive adversary.*

442 **Remark 3.7.** *By the same approach as the recent dynamic weighted matching framework of [20]*  
443 *restricted to bipartite graphs, Theorem 3.6 implies a  $(1 + \frac{1}{\sqrt{2}} + \epsilon)$  robust approximation for weighted*  
444 *bipartite matching with the same update time, up to an exponential blowup in the dependence on  $\epsilon$ .*<sup>5</sup>  
445

## 4 ALGORITHMS ON GENERAL GRAPHS

448 In this section we present our main result: a robust dynamic algorithm maintaining a  $(2 - \Omega(1))$ -  
449 approximation to the size of the maximum matching in a general graph in worst-case  $\tilde{O}_\epsilon(1)$  update  
450 time. As with our algorithm for bipartite graphs, we start with a two-pass semi-streaming algorithm  
451 in Section 4.1, and then show how to approximately implement it dynamically in Section 4.2.1. Note  
452 that unlike in our implementation for bipartite graphs, we will not be able to refer to Proposition 2.2  
453 in a black-box way to estimate the size of the second matching, defined by the matched and  
454 unmatched node sets of the dynamic matching we explicitly maintain. Instead, we unbox the proof  
455 of Proposition 2.2 in Lemma 4.6, so that it can handle queries suitable for our implementation for  
456 non-bipartite graphs. Finally, in Section 4.2.2 we show that our approach allows us to improve the  
457 approximation of any algorithm with approximation ratio in the range  $(1.5, 2]$ .  
458

### 4.1 Two-Pass Streaming Algorithm

460 The key challenge in extending Algorithm 1 and its analysis to non-bipartite graphs is its reliance  
461 on the integrality of the fractional matching polytope in bipartite graphs. This allowed us to focus  
462 on proving the existence of a large fractional matching, which guarantees the existence of a large  
463 integral matching of (at least) the same size. For general graphs this argument fails, and so instead  
464 we search for length-three augmenting paths (3-augmenting paths, for short) with respect to our  
465 first matching,  $M_1$ , by computing some large  $b$ -matching  $M_2$  in the edge set  $V(M_1) \times \overline{V(M_1)}$ . The  
466 main difficulty with this approach in general graphs is that both the endpoints of an edge  $(u, v) \in M_1$   
467 might get matched (in  $M_2$ ) to the same node  $w$ , and the resulting triangle  $w - u - v - w$  does not  
468 help us in any way to create a 3-augmenting path involving the edge  $(u, v) \in M_1$ .  
469

470 We overcome this difficulty using *random bipartitions* (see Algorithm 2). As before, in the first  
471 pass we compute an  $(\epsilon/4)$ -AMM  $M_1$  in the input graph  $G$ .<sup>6</sup> Next, we define the following random  
472 bipartition  $(L, R)$  of the node-set  $V$ . For each matched edge  $(u, v) \in M_1$ , we arbitrarily include one  
473 of its endpoints in  $L$  and the other in  $R$ . Next, for each unmatched node  $v \in \overline{V(M_1)}$ , we include  
474 the node  $v$  in into one of  $L$  and  $R$  chosen uniformly at random. Subsequently, we assign a capacity  
475  $b_v := 1$  to all nodes  $v \in V(M_1)$  and a capacity  $b_v := b$  to all nodes  $v \in \overline{V(M_1)}$ , for some integer  $b$   
476 to be chosen later. Let  $B = (V, E_2)$  be the bipartite subgraph spanned by edges with a single node  
477 matched in  $M_1$  and endpoints in opposite sides, i.e.,  
478

$$E_2 := \{(u, v) \in E \mid u \in V(M_1), v \in \overline{V(M_1)}, |\{u, v\} \cap L| = |\{u, v\} \cap R| = 1\}.$$

481 In the second pass, we compute a maximal  $b$ -matching  $M_2$  in  $B$  w.r.t. the capacities  $\{b_v\}$ . Finally,  
482 we return the maximum matching in the subgraph  $G[M_1 \cup M_2]$ .  
483

487<sup>5</sup>This extension is not obtained by using the framework of [20] directly, as the latter requires explicit dynamic matchings.  
488 Nonetheless, their arguments can be extended to the value version of the problem.

489<sup>6</sup>As with the bipartite Algorithm 1, we suggest the reader think of  $M_1$  as being maximal for now (i.e.,  $\epsilon = 0$ ).

**Algorithm 2** General Two-Pass Streaming Algorithm

---

```

491 1:  $M_1 \leftarrow (\epsilon/4)$ -AMM computed from first pass
492 2: for edge  $(u, v) \in M_1$  do
493 3:    $s(u) \leftarrow \ell$  and  $s(v) \leftarrow r$ 
494 4: for vertex  $w \in \overline{V(M_1)}$  do
495 5:    $s(w) \sim \text{Uni}\{\ell, r\}$ 
496 6: let  $L \leftarrow \{v \mid s(v) = \ell\}$  and  $R \leftarrow \{v \mid s(v) = r\}$ 
497 7: Assign each vertex  $v$  capacity  $b_v = \begin{cases} 1 & v \in V(M_1) \\ b & v \in \overline{V(M_1)}. \end{cases}$ 
498 8: let  $B \leftarrow (V, E_2)$ , for  $E_2 = \{(u, v) \in E \mid u \in V(M_1), v \in \overline{V(M_1)}, |\{u, v\} \cap L| = |\{u, v\} \cap R| = 1\}$ .
499 9:  $M_2 \leftarrow$  maximal  $b$ -matching in  $B$  computed from second pass
500 10: Output maximum matching in  $G[M_1 \cup M_2]$ 
501

```

---

**Intuition:** The intuition behind Algorithm 2 is as follows: if  $M_1$  is only roughly 2-approximate, then many 3-augmenting paths exist in  $G$  w.r.t.  $M_1$ , by Proposition 2.6. Now, a constant fraction of these (specifically, a quarter) “survive” the random bipartition and have their extreme edges belong to  $B$ . Now, for each augmenting path  $u' - u - v - v'$  that survives, either an augmenting path containing  $u - v$  is found in  $M_1 \cup M_2$ , or at least one of  $u'$  or  $v'$  is matched  $b$  times to nodes in  $V(M_1)$  other than  $u$  or  $v$ . Next, since nodes in  $V(M_1)$  can only be matched once in  $M_2$ , this limits the number of paths where  $u$  and  $v$  do not participate in an augmenting path. This implies a large number of augmenting paths in  $M_1 \cup M_2$  that are disjoint in their  $V(M_1)$  nodes. Finally, since each node in  $\overline{V(M_1)}$  belongs to at most  $b$  such paths, some large  $\Omega(1/b)$  fraction of these augmenting paths are also disjoint in their  $\overline{V(M_1)}$  nodes, from which we conclude that  $M_1 \cup M_2$  contains a large set of node-disjoint augmenting paths w.r.t.  $M_1$ , and that  $G[M_1 \cup M_2]$  contains a large matching.

We now substantiate the above intuition. The first lemma in this vein asserts that  $M_1 \cup M_2$  contains many 3-augmenting paths w.r.t.  $M_1$  (assuming  $M_1$  is not already near maximum in size).

**Lemma 4.1.** *If  $|M_1| = (\frac{1}{2} + c) \cdot \mu(G)$ , then  $G[M_1 \cup M_2]$  contains a set  $\mathcal{P}$  of 3-augmenting paths w.r.t.  $M_1$  that are disjoint in their  $V(M_1)$  nodes, with expected cardinality at least*

$$\mathbb{E}[|\mathcal{P}|] \geq \left( \frac{b}{b+1} \right) \cdot \left( \frac{1}{4} \cdot \left( \frac{1}{2} - 3c \right) - \frac{1}{b} \cdot \left( \frac{1}{2} + c \right) - \frac{7\epsilon}{8} \right) \cdot \mu(G).$$

As the proof of Lemma 4.1 is a little calculation heavy, we defer its proof to Appendix C, and instead prove the following slightly weaker but simpler bound here.

**Lemma 4.2.** *If  $|M_1| = (\frac{1}{2} + c) \cdot \mu(G)$ , then  $G[M_1 \cup M_2]$  contains a set  $\mathcal{P}$  of 3-augmenting paths w.r.t.  $M_1$  that are disjoint in their  $V(M_1)$  nodes, with expected cardinality at least*

$$\mathbb{E}[|\mathcal{P}|] \geq \left( \frac{1}{4} \cdot \left( \frac{1}{2} - 3c - \frac{7\epsilon}{2} \right) - \frac{2}{b} \cdot \left( \frac{1}{2} + c \right) \right) \cdot \mu(G).$$

**PROOF.** Fix a maximum set of node-disjoint 3-augmenting paths in  $G$  w.r.t.  $M_1$ , denoted by  $\mathcal{P}^*$ . By Proposition 2.6, we have  $|\mathcal{P}^*| \geq (\frac{1}{2} - 3c - \frac{7\epsilon}{2}) \cdot \mu(G)$ . Next, let  $S \subseteq \mathcal{P}^*$  be the paths  $u' - u - v - v'$  who “survive” the bipartition, in that  $(u, u'), (v, v') \in E_2$ . By construction, each path in  $\mathcal{P}^*$  survives with probability exactly  $\frac{1}{4}$ . Therefore,  $\mathbb{E}[|S|] \geq \frac{1}{4} \cdot (\frac{1}{2} - 3c - \frac{7\epsilon}{2}) \cdot \mu(G)$ .

Now, for each survived path  $u' - u - v - v' \in S$ , either both  $u$  and  $v$  are matched (exactly once) in  $M_2$ , thus contributing an augmenting path, or at least one of  $u'$  and  $v'$  must be matched in  $M_2$  to  $b$  distinct nodes in  $V(M_1)$ . But since each vertex in  $V(M_1)$  is matched at most once in  $M_2$ ,

there are thus at most  $|V(M_1)|/b = 2|M_1|/b$  paths in  $S$  whose middle edges do not belong to a 3-augmenting path in  $M_1 \cup M_2$ . Therefore, there are at least  $|S| - 2|M_1|/b$  many edges  $(u, v)$  in  $M_1$  whose endpoints are both matched in  $M_2$  to some (different) nodes  $u''$  and  $v''$ , respectively. Each such edge contributes an augmenting path to a set  $\mathcal{P}$  of the desired size,

$$\mathbb{E}[|\mathcal{P}|] = \mathbb{E}[|S|] - \frac{2|M_1|}{b} \geq \left( \frac{1}{4} \cdot \left( \frac{1}{2} - 3c - \frac{7\epsilon}{2} \right) - \frac{2}{b} \cdot \left( \frac{1}{2} + c \right) \right) \cdot \mu(G). \quad \square$$

The preceding two lemmas imply the existence of a multitude of 3-augmenting paths that are disjoint in their  $V(M_1)$  nodes. We now use these augmenting paths to prove the existence of numerous (though possibly fewer) augmenting paths that are disjoint in *all* their nodes. Since each of the two  $V(M_1)$  nodes of a 3-augmenting path belong to at most  $b$  such paths, it is easy to find some  $1/(2b - 1)$  fraction of these augmenting paths that are disjoint in all their nodes. The following lemma, resembling [48, Lemma 6], increases this fraction to  $1/b$ .

**Lemma 4.3.** *Let  $\mathcal{P}$  be a set of 3-augmenting paths w.r.t.  $M_1$  in  $G[M_1 \cup M_2]$  such that each  $V(M_1)$  (resp.  $V(M_1)$ ) node belongs to at most one (resp.,  $b$ ) paths in  $\mathcal{P}$ . Then  $\mathcal{P}$  contains a set of node-disjoint 3-augmenting paths  $\mathcal{P}' \subseteq \mathcal{P}$  of cardinality at least  $|\mathcal{P}'| \geq \frac{1}{b} \cdot |\mathcal{P}|$ .*

**PROOF.** Consider the graph  $G' = (\overline{V(M_1)}, E')$  obtained by replacing each path  $u' - u - v - v'$  in  $\mathcal{P}$  with a single edge  $u' - v'$ . This graph  $G'$  is bipartite, by virtue of our random bipartition of  $G$ . Now, since this bipartite graph  $G'$  has maximum degree  $b$ , it contains a matching of size at least  $|E'|/b = |\mathcal{P}|/b$ : the fractional matching assigning values  $1/b$  to each edge has value  $|E'|/b$ , and so  $G'$  contains an *integral* matching of at least the same value. On the other hand, disjoint edges in  $G'$  have a one-to-one mapping to node-disjoint paths in  $\mathcal{P}$ , since each node in  $V(M_1)$  belongs to at most one such path. Thus, the maximum matching in  $G'$  corresponds to a collection  $\mathcal{P}' \subseteq \mathcal{P}$  of node-disjoint augmenting paths in  $G[M_1 \cup M_2]$  w.r.t.  $M_1$ , of cardinality at least  $|\mathcal{P}'| \geq |\mathcal{P}|/b$ .  $\square$

The three preceding lemmas imply that  $G[M_1 \cup M_2]$  contains a large set of vertex-disjoint 3-augmenting paths w.r.t.  $M_1$ , assuming this latter matching is not already large. As we now show, this implies that  $G[M_1 \cup M_2]$  contains a better-than-2-approximate matching.

**THEOREM 4.4.** *Let  $\epsilon \in (0, 1/4)$ . Then, Algorithm 2 with  $b = 9$  satisfies  $\mu(G) \geq \mathbb{E}[\mu(G[M_1 \cup M_2])] \geq (\frac{1}{2} + \frac{1}{144} - \epsilon) \cdot \mu(G)$ , and is thus  $(\frac{1}{2} + \frac{1}{144} - \epsilon)^{-1} < 1.973(1 + 2\epsilon)$ -approximate in expectation.*

**PROOF.** Let  $|M_1| = (\frac{1}{2} + c) \cdot \mu(G)$ , where  $c \in [-\epsilon/8, 1/2]$ , with the lower bound on  $c$  following from Observation 2.4 and  $M$  being an  $(\epsilon/4)$ -AMM. Let  $\mathcal{P}'$  be a maximum set of vertex-disjoint 3-augmenting paths w.r.t.  $M_1$  in  $G[M_1 \cup M_2]$ . Then, augmenting along these paths, we find that  $M_1 \bigoplus \mathcal{P}'$  contains a matching (hence of size at most  $\mu(G)$ ) of the desired expected cardinality.

$$\begin{aligned} \mathbb{E}[|M_1 \bigoplus \mathcal{P}'|] &= \mathbb{E}[|M_1| + |\mathcal{P}'|] \\ &\geq \left( \frac{1}{2} + c \right) \cdot \mu(G) + \left( \frac{1}{b+1} \cdot \left( \frac{1}{4} \cdot \left( \frac{1}{2} - 3c \right) - \frac{1}{b} \cdot \left( \frac{1}{2} + c \right) - \frac{7\epsilon}{8} \right) \right) \cdot \mu(G) \\ &\geq \left( \frac{1}{2} - \frac{\epsilon}{8} + \frac{1}{b+1} \cdot \left( \frac{1}{4} \cdot \frac{1}{2} - \frac{1}{b} \cdot \frac{1}{2} \right) - \frac{7\epsilon}{8} \right) \cdot \mu(G) \\ &= \left( \frac{1}{2} + \frac{1}{144} - \epsilon \right) \cdot \mu(G). \end{aligned}$$

Above, the first inequality follows from Lemma 4.1 and Lemma 4.3, the second inequality mainly relies on the parenthetical expression being increasing in  $c \geq -\epsilon/8$  (for our choice of  $b = 9$ ). Finally, the equality holds by our choice of  $b = 9$ .  $\square$

## 589 4.2 Dynamic Algorithms

590 In this section we provide the dynamic algorithms yielding our main results, theorems 1.1 and 1.2.  
 591 As with the bipartite case, our general approach is to approximately implement our two-pass  
 592 streaming algorithm in a dynamic setting. Unlike the algorithm for bipartite graphs, here we need  
 593 to (slightly) unbox the sublinear-time algorithm of [12] to find a large set of edges in  $M_1$  which  
 594 belong to 3-augmenting paths in  $G[M_1 \cup M_2]$ , as explained below.

595 **4.2.1** *Breaking the barrier of two in polylog time.* In this section, we present a robust dynamic  
 596 (1.973 +  $\epsilon$ )-approximate maximum matching size with worst case update time of  $t_u = \tilde{O}_\epsilon(1)$ , and a  
 597 query time of  $t_q = \tilde{O}_\epsilon(n)$ , provided  $\mu(G) \geq \epsilon n$ . This, combined with Proposition 2.1, implies our  
 598 main result, Theorem 1.1.

600 For our dynamic (approximate) implementation of Algorithm 1, which works on bipartite graphs,  
 601 all we needed was to estimate  $|M_2|$ . In contrast, for our dynamic (approximate) implementation of  
 602 Algorithm 2, we will need to estimate the size of the set  $\mathcal{P}$  as guaranteed by Lemma 4.1. Specifically,  
 603 we note that the proofs of Lemmas 4.1, 4.3 and 4.4 imply the following observation.

604 **Observation 4.5.** *Let  $\widehat{M}_1 \subseteq M_1$  be the set of edges in  $M_1$  whose two endpoints are matched in  $M_2$  in  
 605 Algorithm 2 run with  $b = 9$ . Then,  $|M_1| + \frac{1}{b} \cdot \mathbb{E} [|\widehat{M}_1|] \geq \frac{\mu(G)}{1.973 \cdot (1+2\epsilon)}$ , and also  $\mu(G) \geq |M_1| + \frac{1}{b} \cdot |\widehat{M}_1|$ .*

606 To estimate  $|\widehat{M}_1|$  efficiently, we make use of the following extension of the algorithm of [12].

607 **Lemma 4.6.** *Consider a graph  $G' = (V', E')$  with  $|V'| = n'$ , and a matching  $M$  with  $V(M) \subseteq V'$   
 608 that is not necessarily part of  $G'$  (i.e., we might have  $M \not\subseteq E'$ ). For any matching  $M'$  in  $G'$ , let  $k_{M'}$   
 609 denote the number of edges in  $M$  both of whose endpoints are matched in  $M'$ . There is an algorithm  
 610 which, given adjacency matrix query access to the edges of  $G'$ , w.h.p. runs in  $\tilde{O}_\epsilon(n')$  time and returns  
 611 an estimate  $\kappa \in [k_{M'} - \epsilon^2 n', k_{M'}]$  for some maximal matching  $M'$  in  $G'$ .*

612 This lemma follows from the work of [12] rather directly, though it requires some unboxing of  
 613 the results there, due to the organization of that work. We substantiate this lemma in Appendix B.

614 Given the above, we are now ready to prove the main result of this section, which is summarized  
 615 in the theorem below.

616 **THEOREM 4.7.** *For any  $\epsilon \in (0, 1/4)$ , there exists a robust dynamic matching size estimator algorithm  
 617  $\mathcal{A}$  with worst-case update time  $\tilde{O}_\epsilon(1)$  that, provided  $\mu(G) \geq \epsilon \cdot n$ , supports  $\tilde{O}_\epsilon(n)$ -time queries and  
 618 outputs a  $(1.973 + \epsilon)$ -approximate estimate w.h.p.*

619 **PROOF.** The dynamic algorithm  $\mathcal{A}$  is based on Algorithm 2. For its updates, it maintains an  
 620  $(\epsilon/4)$ -AMM  $M_1$  in the input graph  $G$ , using Lemma 2.5, and a balanced binary search tree (BST)  
 621 of edges in the graph, allowing for logarithmic-time insertion, deletion and edge queries. This  
 622 immediately implies a worst-case update time of  $t_u = \tilde{O}_\epsilon(1)$ .

623 We now describe how Algorithm  $\mathcal{A}$  responds to a query about the maximum matching size. To  
 624 answer this query, the algorithm considers a new auxiliary graph  $G^* = (V^*, E^*)$ , which is defined as  
 625 follows. For each node  $u \in V(M_1)$ , create a node  $0_u$  in  $G^*$ . Next, for each node  $v \in \overline{V(M_1)}$ , create  $b$   
 626 nodes  $1_v, \dots, b_v$  in  $G^*$ . Finally, for every edge  $(u, v) \in E_2$ , with  $u \in V(M_1)$  and  $v \in \overline{V(M_1)}$ , create an  
 627 edge  $(0_u, i_v)$  in  $G^*$  for all  $i \in \{1, \dots, b\}$ . Note that there is a one-to-one mapping between maximal  
 628 matchings in the new graph  $G^*$  and maximal  $b$ -matchings in  $B$ .

629 We emphasize that our dynamic algorithm  $\mathcal{A}$  does not explicitly maintain the auxiliary graph  
 630  $G^*$ . When we receive a query about the maximum matching size in  $G$ , we explicitly construct only  
 631 the node-set  $V^*$  of  $G^*$ , based on the matching  $M_1$ . This takes only  $O(n)$  time. We can, however,  
 632 simulate adjacency matrix queries in  $G^*$  efficiently: there exists an edge  $(0_u, i_v)$  in  $G^*$  iff there exists

638 an edge  $(u, v)$  between the corresponding nodes in  $G$ , verifiable in  $O(\log n)$  time using our edge-set  
 639 BST.

640 At this point, we estimate the size of  $\widehat{M}_1$  by invoking Lemma 4.6 with  $G' = G^*$  and  $M = M_1$ .  
 641 This gives us, in time  $\tilde{O}_\epsilon(n)$  a value  $\kappa$  satisfying  $\kappa \in [|\widehat{M}_1| - \epsilon^2 n, |\widehat{M}_1|]$ , w.h.p. We now return  
 642  $v := |M_1| + \frac{1}{b} \cdot \kappa$  as our estimate of  $\mu(G)$ . All in all, our algorithm has query time  $t_q = \tilde{O}_\epsilon(n)$ .

643 It remains to analyze the approximation ratio. For this, we observe that, by our hypothesis that  
 644  $\mu(G) \geq \epsilon \cdot n$ ,

$$\begin{aligned}
 645 \mathbb{E}[v] &\geq \mathbb{E} \left[ |M_1| + \frac{1}{b} \cdot (|\widehat{M}_1| - \epsilon^2 n) \right] \\
 646 &= |M_1| + \frac{1}{b} \cdot \mathbb{E} \left[ \widehat{M}_1 \right] - \frac{\epsilon^2 n}{b} \\
 647 &\geq \frac{\mu(G)}{1.973 \cdot (1 + 2\epsilon)} - \epsilon^2 n \\
 648 &\geq \frac{\mu(G)}{1.973 \cdot (1 + 2\epsilon)} - \epsilon \cdot \mu(G) \\
 649 &\geq \frac{\mu(G)}{1.973 \cdot (1 + 5\epsilon)}. \tag{6}
 \end{aligned}$$

657 In the above derivation, the second inequality follows from Observation 4.5. Similarly, we have  
 658 w.h.p.:

$$659 v \leq |M_1| + \frac{1}{b} \cdot |\widehat{M}_1| \leq \mu(G). \tag{7}$$

660 The second inequality in the above derivation again follows from Observation 4.5. From (6) and (7),  
 661 we conclude that we return in response to each query a  $1.973(1 + 5\epsilon)$ -approximation to  $\mu(G)$  in  
 662 expectation. Therefore, by standard Chernoff bounds, running  $O_\epsilon(\log n)$  copies of this algorithm  
 663 (increasing update and query time appropriately) and taking the average of these will then result  
 664 in a  $1.973(1 + O(\epsilon))$  approximation of the desired value, w.h.p. Reparameterizing  $\epsilon$  appropriately,  
 665 the theorem follows.  $\square$

666 Combined with Observation 2.1, the above theorem implies our main result, Theorem 4.7.

667 **4.2.2 New time/approximation tradeoffs.** In this section we show our secondary result: a black-  
 668 box method to improve dynamic matching algorithm's approximation ratio, at the cost of only  
 669 outputting a size estimate. We start with the following observation.

670 **Proposition 4.8.** *Let  $G$  be an  $n$ -node graph,  $\epsilon \in (0, 1)$  and  $\alpha \geq 1$ . Then, given an  $\epsilon$ -AMM  $M'$  and  
 671  $\alpha$ -approximate maximum matching  $M''$  in  $G$ , one can compute in  $O(n)$  time a matching  $M$  in  $G$   
 672 which is both  $\alpha$ -approximate and an  $\epsilon$ -AMM.*

673 **PROOF.** The subgraph  $G[M' \cup M'']$  has maximum degree two, and is thus the union of paths  
 674 and cycles. Let  $M$  be the matching obtained by taking from each connected component  $C$  in  
 675  $G[M' \cup M'']$  either the set of edges of  $M'$  or  $M''$  that are most plentiful in  $C$ , breaking ties in favor  
 676 of  $M'$ . By construction, it is clear that  $M$  is a matching, and that moreover  $|M| \geq |M''|$ , and so  $M$  is  
 677  $\alpha$ -approximate. On the other hand,  $M$  matches all nodes of  $M'$  in each component, and therefore  
 678 overall. That is, after removing at most  $\epsilon\mu(G)$  nodes in  $V \setminus V(M) \subseteq V \setminus V(M')$ , we obtain a graph  
 679 in which  $M$  is maximal. That is, the matching  $M$  is also an  $\epsilon$ -AMM.  $\square$

680 We are now ready to prove Theorem 1.2, restated below for ease of reference.

687 THEOREM 1.2. For any  $\alpha > 1.5$ , a dynamic  $\alpha$ -approximate matching algorithm with update time  
 688  $t_u$  implies a dynamic  $\left(\alpha - \Omega\left((1 - 6(\frac{1}{\alpha} - \frac{1}{2}))^2\right)\right)$ -approximate matching size estimator with update  
 689 time  $\tilde{O}(t_u)$ .  
 690

691 PROOF. Let  $\epsilon > 0$  be some sufficiently small constant. We describe how to obtain a new dynamic  
 692 matching size estimator  $\mathcal{A}$  for  $G$ , with update time  $\tilde{O}_\epsilon(t_u)$  that, provided  $\mu(G) \geq \epsilon \cdot n$ , supports  $\tilde{O}_\epsilon(n)$ -  
 693 time queries and outputs a  $\beta$ -approximate estimate w.h.p., for some  $\beta = \alpha - \Omega((1 - 6(1/\alpha - 1/2))^2)$ .  
 694 The theorem then immediately follows from Proposition 2.1.  
 695

696 The algorithm  $\mathcal{A}$  works as follows. It maintains an  $\epsilon$ -AMM  $M'_1$ , invoking Lemma 2.5, taking  $\tilde{O}_\epsilon(1)$   
 697 update time. It also maintains  $\alpha$ -approximate matching  $M''_1$ , by running the dynamic algorithm  
 698 guaranteed by the theorem's hypothesis, taking  $t_u$  update time. Therefore, Algorithm  $\mathcal{A}$  has an  
 699 overall update time of  $t_u + \tilde{O}_\epsilon(1) = \tilde{O}_\epsilon(t_u)$ .

700 Upon receiving a query, Algorithm  $\mathcal{A}$  first invokes Proposition 4.8 to obtain a matching  $M_1$  (based  
 701 on  $M'_1$  and  $M''_1$ ) that is simultaneously an  $\epsilon$ -AMM and an  $\alpha$ -approximate maximum matching in  $G$ .  
 702 This takes  $O(n)$  time. The rest of the query algorithm remains exactly the same as in Section 4.2.  
 703 This implies that Algorithm  $\mathcal{A}$  has an overall query time of  $\tilde{O}_\epsilon(n)$ .

704 We now analyze the approximation guarantee of  $\mathcal{A}$ . Towards this end, observe that as  $\alpha > 1.5$ ,  
 705 we can write  $\frac{1}{\alpha} = \frac{1}{2} + c$  where  $0 < c < 1/6$ . So, we have that  $|M_1| = \left(\frac{1}{2} + z\right) \cdot \mu(G)$  for some  $z \geq c$ .  
 706 Therefore, by Lemma 4.2 and Lemma 4.3, there exists some set  $\mathcal{P}'$  of node-disjoint length-three  
 707 augmenting paths w.r.t.  $M_1$  in  $G[M_1 \cup M_2]$  whose cardinality satisfies

$$\frac{|\mathcal{P}'|}{\mu(G)} \geq f_b(z) := \frac{1}{b} \cdot \left( \frac{1}{4} \cdot \left( \frac{1}{2} - 3z - \frac{7\epsilon}{2} \right) - \frac{2}{b} \cdot \left( \frac{1}{2} + z \right) \right).$$

712 Augmenting along these paths with respect to  $M_1$ , we obtain a new matching in  $G[M_1 \cup M_2]$  of  
 713 cardinality at least  $(\frac{1}{2} + z + f_b(z)) \cdot \mu(G)$ . Now, for  $b \geq 2$  (as we will choose), this matching size  
 714 is decreasing in  $z$ , as observed by taking the derivative of  $1/2 + z + f_b(z)$  w.r.t.  $z$ . Therefore, the  
 715 matching size is minimized at  $z = c$ , and we find that  $\mu(G[M_1 \cup M_2]) \geq 1/2 + c + f_b(c)$ . Taking  
 716 another derivative, this time with respect to  $b$ , we find that this expression is minimized (ignoring  
 717 the  $\epsilon$  dependence) at  $b^* = \frac{16(1+2c)}{1-6c}$ . Note that  $b \geq 16$ , as  $c \in (0, 1/6)$ . This optimal  $b^*$  need not be an  
 718 integer, however, and so we take  $b = \lceil b^* \rceil \leq \frac{17}{16}b^*$  in our algorithm, and find that  $M_1 \cup M_2$  contains a  
 719 matching of size at least  $\mu(G)$  times  $1/2 + z + f_b(z) \geq 1/2 + c + f_{\frac{17}{16}b^*}(c) \geq 1/2 + c + \frac{9(1-6c)^2}{2312(1+2c)} - O(\epsilon/b)$ .  
 720 Moreover, some  $b \cdot f_b(z)$  many edges  $\hat{M}_1 \subseteq M_1$  have both of their endpoints matched in the  $b$ -  
 721 matching  $M_2$ .  
 722

723 We conclude that  $\mathbb{E}[\mu(G[M_1 \cup M_2])]$  gives a strictly-better-than- $\alpha$  approximation to  $\mu(G)$  (again  
 724 using that  $c < 1/6$ ). Specifically, the gain we get in the approximation ratio is of the order of  
 725  $\Theta((1 - 6c)^2) = \Theta((1 - 6(1/\alpha - 1/2))^2)$ . Now, using the fact that  $\mu(G) \geq \epsilon n$  and we are running  
 726 the same query algorithm as in Section 4.2, our estimation using the sub-linear-time algorithm  
 727 (Lemma 4.6) gives a strictly-better-than- $\alpha$  approximation to  $\mu(G)$  in expectation. As before, taking  
 728 the average of  $O(\log n)$  copies of this algorithm will provide the same bound w.h.p., at an additional  
 729 logarithmic multiplicative overhead to the update and query times.  $\square$   
 730

731  
 732 **Remark 4.9.** We note that the reduction of Theorem 1.2 preserves robustness and worst-case update  
 733 time.  
 734

## 736 5 AMMS AGAINST ADAPTIVE ADVERSARIES

737 In this section we prove Lemma 2.5. That is, we provide a robust dynamic algorithm for maintaining  
 738 an  $\epsilon$ -AMM in worst-case polylogarithmic update time. But first, we motivate our algorithm, and  
 739 characterize the kind of matching we wish to compute.

740 We first recall a useful tool in the literature, namely *edge sparsification*: maintaining a sparse  
 741 subgraph of  $G$  containing a large matching—a so-called *matching sparsifier*. Such sparsifiers naturally  
 742 allow to achieve speedups in the algorithms needed for Proposition 2.1, as a large matching in  
 743 a sparsifier can be computed quickly. One influential such sparsifier that we will use are *kernels*,  
 744 introduced by Bhattacharya et al. [25].

745 **Definition 5.1.** For  $\epsilon \geq [0, 1]$  and  $d \in \mathbb{N}$ , a subgraph  $K = (V, E_K)$  of graph  $G = (V, E)$  is an  
 746  $(\epsilon, d)$ -kernel if  $K$ 's maximum degree is at most  $d$  and each edge  $e \in E \setminus E_K$  has at least one endpoint  
 747 of degree at least  $d(1 - \epsilon)$  in  $K$ .

748 These sparsifiers will play an integral role in robustly and efficiently maintaining an AMM in  
 749 this section. We start by motivating their use in computing AMMs in a *static* setting.

### 750 5.1 From kernels to AMMs

751 To motivate the interest in bounded-degree graphs, we recall the following observation, which  
 752 follows from the fact that the  $2\mu(G)$  endpoints of a maximum matching form a vertex cover (i.e.,  
 753 are incident on each edge of the graph).

754 **Fact 5.2.** Let  $G = (V, E)$  be a graph of maximum degree  $\Delta$ . Then  $|E| \leq 2\mu(G) \cdot \Delta$ .

755 When  $d$  and  $\mu(G)$  are small, Fact 5.2 and Definition 5.1 imply that  $(\epsilon, d)$ -kernels of  $G$  are quite  
 756 sparse subgraphs compared to the size of a maximum matching contained within them. In particular,  
 757 since maximal matchings are computable in linear time, in kernel  $K$  computing a maximal matching  
 758 requires only  $|E_K| \leq 2\mu(K) \cdot d = O(\mu(G) \cdot d)$  time. The following result of [41] implies that  
 759 essentially the same amount of time is needed to compute a *near-maximum-weight* matching in  $K$ .

760 **Proposition 5.3.** Let  $G = (V, E, w)$  be a weighted graph. There exist a deterministic algorithm for  
 761 computing a  $(1 + \epsilon)$ -approximate maximum weight matching of  $G$  in time  $O_\epsilon(|E|)$ .

762 We now explain how to find large matchings in a kernel  $K$  that allow us to obtain an AMM of  
 763  $G$ . For this, we will need to upper bound the number of high-degree nodes in  $K$ . Specifically, for  
 764 an  $(\epsilon, d)$ -kernel  $K$  of graph  $G$ , we denote by  $H_K := \{v \mid d_K(v) \geq d(1 - \epsilon)\}$  the set of high-degree  
 765 nodes in  $K$ . We will wish to argue that a removal of few high-degree nodes in the kernel yields  
 766 a subgraph in which our matching is maximal. We therefore need to prove that the number of  
 767 high-degree nodes is itself small in terms of  $\mu(G)$ .

768 **Lemma 5.4.** Let  $K = (V, E_K)$  be an  $(\epsilon, d)$ -kernel  $K$  of  $G$  with  $\epsilon \leq 1/4$ . Then  $|H_K| \leq 4\mu(G)$ .

769 **PROOF.** We consider the fractional matching  $x \in \mathbb{R}^E$  where  $x_e = \mathbb{1}[e \in E_K]/d$ . By the degree  
 770 bound of  $K$ , this is a feasible fractional matching in  $G$ . Using this fractional matching, we can show  
 771 that

$$772 \frac{1}{2} \cdot (1 - \epsilon) \cdot |H_K| \leq \sum_{v \in H_K} \sum_{e \ni v} x_e \leq \sum_e x_e \leq \frac{3}{2} \cdot \mu(G),$$

773 where the first inequality follows from the definition of  $H_K$  and possible double counting of edges,  
 774 and the last inequality follows from the integrality gap of  $\frac{3}{2}$  of the fractional matching polytope.  
 775 Simplifying the above and using  $\epsilon \leq 1/4$ , we have that indeed  $|H_K| \leq (3/(1-\epsilon)) \cdot \mu(G) \leq 4\mu(G)$ .  $\square$

785 We now show that it is sufficient to match the  $(1 - \Theta_\epsilon(1))$ -fraction of vertices of  $H_K$  in order to  
 786 find an AMM in  $G$  through Lemma 5.5. Lemma 5.6 then proves that there exists a static algorithm  
 787 with running time  $\tilde{O}_\epsilon(k \cdot \mu(G))$  for computing an AMM of  $G$  given access to kernel  $K$ .

788 **Lemma 5.5.** *Let  $K = (V, E_K)$  be an  $(\epsilon, d)$ -kernel of  $G = (V, E)$ , for  $\epsilon \in (0, 1/4)$  and  $d \geq \frac{1}{\epsilon}$ . Then,  
 789 a maximal matching  $M$  in  $K$  that matches at least a  $(1 - c \cdot \epsilon)$ -fraction of  $H_K$  is a  $4c\epsilon$ -AMM in  $G$ .  
 790 Moreover, such a matching exists for  $c = 2$ .*

791 **PROOF.** First, we argue that such a matching  $M$ , if it exists, is indeed a  $4c\epsilon$ -AMM in  $G$ . We recall  
 792 that every edge in  $E \setminus E_K$  has a high-degree endpoint in  $K$ . Therefore, if we remove the  $c\epsilon$  fraction  
 793 of high-degree nodes  $H_K$  unmatched by  $M$ , each edge in  $E \setminus E_K$  in the resulting graph  $G'$  has at  
 794 least one endpoint matched in  $M$ . On the other hand, every edge in  $E_K$  has an endpoint matched in  
 795  $M$ , by maximality of  $M$  in  $K$ . We conclude that after removing  $c\epsilon \cdot |H_K| \leq 4c\epsilon \cdot \mu(G)$  nodes in  $G$   
 796 (with the inequality relying on Lemma 5.4), we obtain a graph  $G'$  where  $M$  is maximal. That is,  $M$   
 797 is a  $4c\epsilon$ -AMM.  
 798

799 We now argue the existence of such a matching  $M$  for  $c = 2$ . Since  $H$  has maximum degree  $d \geq \frac{1}{\epsilon}$ ,  
 800 by Vizing's theorem [81] it can be  $(d + 1)$ -edge-colored, i.e., decomposed into  $(d + 1)$  matchings.  
 801 A randomly-chosen color in this edge coloring is a matching  $M'$  that matches each edge with  
 802 probability  $\frac{1}{d+1}$ , and thus it matches each high-degree vertex  $v$  with probability at least

$$803 \Pr[v \text{ matched}] \geq d(1 - \epsilon)/(d + 1) \geq (1 - \epsilon)/(1 + \epsilon) \geq 1 - 2\epsilon.$$

804 Finally, extending this matching  $M'$  to also be maximal in  $K$  by adding edges of  $K$  greedily then  
 805 proves the existence of the desired  $8\epsilon$ -AMM contained in the kernel  $K$ .  $\square$

806 We now turn to making the above proof of existence constructive.

807 **Lemma 5.6.** *Given an  $(\epsilon, d)$ -kernel  $K = (V, E_K)$  of  $G = (V, E)$ , one can compute an  $\epsilon$ -AMM in  $G$  in  
 808 deterministic time  $O_\epsilon(d \cdot \mu(G))$ .*

809 **PROOF.** Let  $\epsilon' = \epsilon/12$ . By Fact 5.2, the number of edges in  $K$  is at most  $|E_K| = O(d \cdot \mu(G))$ . We  
 810 then compute a  $(1 + \epsilon')$ -max weight matching  $M'$  in the graph  $G$  with edge weights equaling the  
 811 number of high-degree nodes incident on them,  $w_e = \sum_{v \in e} \mathbb{1}[v \in H_K] \in \{0, 1, 2\}$ . By Proposition 5.3,  
 812 this can be done in deterministic time  $O_\epsilon(d \cdot \mu(G))$ . By Lemma 5.5, this guarantees that at least a  
 813  $(1 - 2\epsilon')/(1 + \epsilon') \geq (1 - 3\epsilon')$  fraction of high-degree nodes in  $K$  are unmatched by this dynamic  
 814 subroutine. We then extend  $M'$  to also be maximal in  $K$ , by scanning over the  $|E_K| = O(d \cdot \mu(G))$   
 815 edges of  $K$  (in the same deterministic time) and adding them to  $M'$  where possible. By Lemma 5.5,  
 816 this results in a  $12\epsilon'$ -AMM, i.e., an  $\epsilon$ -AMM, after a total of  $O_\epsilon(d \cdot \mu(G))$  deterministic time.  $\square$

817 So far, we have provided a *static* AMM algorithm with deterministic time  $O_\epsilon(d \cdot \mu(G))$ , provided we  
 818 have access to a kernel. To dynamize the above, we first show how to maintain a kernel dynamically  
 819 through periodic recomputation.

## 820 5.2 Periodic kernels and AMMs

821 In [82], Wajc provided a method for rounding dynamic fractional matchings to matching sparsifiers,  
 822 and from these (by methods underlying Algorithm 3), we can obtain integral matchings. Crucially  
 823 for our needs, his framework was robust, and allowed for worst-case update times. Unfortunately  
 824 for us, the lemma statements in his work do not immediately imply a robust dynamic kernel  
 825 maintenance. However, they do allow for *kernel queries*, with running time  $\tilde{O}(d \cdot \mu(G))$ .

826 **Lemma 5.7.** *Let  $\epsilon \in (0, 1)$  and  $d = \tilde{O}_\epsilon(1)$  be sufficiently large. Then, there exists a robust algorithm  
 827 with worst-case update time  $t_u = \tilde{O}_\epsilon(1)$  allowing for  $(\epsilon, d)$ -kernel and  $\epsilon$ -AMM queries in worst-case  
 828 query time  $t_q = \tilde{O}_\epsilon(d \cdot \mu(G))$ . The query's outputs are a kernel and an  $\epsilon$ -AMM w.h.p.*

Given the ability to query a kernel, the ability to query an AMM then follows directly from Lemma 5.6. As the proof and presentation of an algorithm allowing for kernel queries essentially requires repeating verbatim numerous lemmas in [82], we defer its proof to Appendix D.

We now turn to designing a robust dynamic algorithm that *always* maintains an AMM.

### 5.3 Robust dynamic AMMs

So far, we have provided a method to answer AMM queries in a dynamic setting. Lemma 5.8 proves that once an AMM is computed it remains an AMM for  $\sim \epsilon \cdot \mu(G)$  updates.

**Lemma 5.8.** *Let  $\epsilon \in (0, 1/2)$ . If  $M$  is an  $\epsilon$ -AMM in  $G$ , then the non-deleted edges of  $M$  during any sequence of at most  $\epsilon \cdot \mu(G)$  updates constitute a  $6\epsilon$ -AMM in  $G$  (during the updates).*

**PROOF.** Let  $G, M$  and  $G', M'$  be the graph and matching before and after the updates, respectively. Since each update can decrease the size of the maximum matching size by at most one, we have

$$\frac{1}{2} \cdot \mu(G) \leq (1 - \epsilon) \cdot \mu(G) \leq \mu(G').$$

Now, recall that for some set of vertices  $U \subseteq V$  of size at most  $|U| \leq \epsilon \cdot \mu(G)$  nodes from  $G$ , the matching  $M$  is maximal in  $G[V \setminus U]$ . Now, after these  $\epsilon \cdot \mu(G)$  updates, it might be that  $2\epsilon \cdot \mu(G)$  edges in  $G'$  are now not incident on edges in  $M'$ . (The factor of two arises due to edges of  $M$  that are deleted leaving two uncovered edges, addressable by removing two more nodes). That is, after removing a node set  $U' \subseteq V$  of size at most  $|U'| \leq 3\epsilon \cdot \mu(G) \leq 6\epsilon \cdot \mu(G')$  nodes from  $G'$ , we obtain a graph  $G'[V \setminus U']$  where  $M'$  is maximal. That is,  $M'$  is an  $O(\epsilon)$ -AMM in  $G'$ .  $\square$

This “stability” of AMMs again lends itself to the periodic re-computation framework of [57], which, together with our algorithms for querying AMMs, allow us to maintain AMMs (always).

**Lemma 2.5.** *For any  $\epsilon \in (0, 1/2)$ , there exists a robust dynamic algorithm that w.h.p. maintains an  $\epsilon$ -AMM in worst-case update time  $\tilde{O}_\epsilon(1)$ .*

**PROOF.** We will run the dynamic AMM query algorithm  $\mathcal{A}$  of Lemma 5.7, whose update fits within our update time budget. We will periodically query  $\mathcal{A}$ , and spread this computation over these periods to guarantee low worst-case update time. Specifically, we will divide the update sequence into *epochs*, where if the graph  $G$  at the start of epoch  $i$  is  $G_i$ , then the epoch has length  $\ell_i \in [\epsilon \cdot \mu(G)/3, \epsilon \cdot \mu(G_i)]$ . In order to determine the length of the epochs, we run the deterministic dynamic  $(2 + \epsilon)$ -approximate fractional matching algorithm of [27], which in particular gives us a  $2 + \epsilon \leq 3$ -approximation of  $\mu(G_i)$  in worst-case update time  $\tilde{O}_\epsilon(1)$ , again fitting within our time budgets. Now, during phase  $i$ , we spend the time  $t_q$  for the  $\epsilon$ -AMM query subroutine of  $\mathcal{A}$ , so as to finish computing  $M_i$ . The amount of time spent per update to achieve this goal is at most  $\frac{\tilde{O}_\epsilon(\mu(G_i))}{[\epsilon \cdot \mu(G_i)/10]} = \tilde{O}_\epsilon(1)$ , again fitting within our time updates. We now describe and analyze the matchings maintained by this algorithm (these are not always  $M_i$ ).

By Lemma 5.8, we need to provide a matching  $M'_{i+1}$  at the start of each phase  $i + 1$  which is an  $O(\epsilon)$ -AMM in  $G_{i+1}$ , thus guaranteeing that the non-deleted edges of  $M'_{i+1}$  remain an  $O(\epsilon)$ -AMM. Reparameterizing appropriately will then yield the desired result. It remains to define our matchings  $M'_i$ . Using our estimate of  $\mu(G)$  obtained by the dynamic fractional matching, we test whether  $\mu(G_i) \in [1/\epsilon, 10/\epsilon]$ . If this is the case, then  $M'_{i+1}$  is obtained by querying the AMM algorithm  $\mathcal{A}$  at the beginning of phase  $i + 1$ , in time  $O_\epsilon(1)$ . (This relied on  $\mu(G_{i+1}) \leq \mu(G_i) + \ell_i \leq \mu(G_i) \cdot (1 + \epsilon) = \tilde{O}_\epsilon(1)$ .) By the properties of  $\mathcal{A}$ , the matching  $M'_{i+1}$  is an  $\epsilon$ -AMM in  $G_{i+1}$  w.h.p. Now, if conversely  $\mu(G_i) \geq 10/\epsilon$ , then we have that

$$\frac{1}{2} \cdot \mu(G_i) \leq (1 - \epsilon) \cdot \mu(G_i) \leq \mu(G_i) - \ell_i \leq \mu(G_{i+1}).$$

883 Now,  $M_i$ , is an  $\epsilon$ -AMM in  $G_i$ , which is obtained from  $G_{i+1}$  by at most  $\epsilon \cdot \mu(G_i) \leq 2\epsilon \cdot \mu(G_{i+1})$   
 884 updates. Therefore, by Lemma 5.8  $M_i$  is a  $12\epsilon$ -AMM in  $G_{i+1}$ . We therefore take  $M'_{i+1}$  to be  $M_i$ .  
 885 Reparameterizing  $\epsilon$  appropriately, the lemma follows.  $\square$

## 887 6 CONCLUSION AND FUTURE DIRECTIONS

888 We presented the first dynamic matching (size estimation) algorithm breaking the approximation  
 889 barrier of 2 in polylogarithmic update time. While this presents a major advance in our understand-  
 890 ing of the dynamic matching problem, many questions remain. We mention a few such questions  
 891 which we find particularly intriguing, and some progress on these following the publication of this  
 892 paper's conference version.

893 **Explicit fast matching.** In our work we show how to maintain a better-than-two approximate  
 894 estimate of the maximum matching size. Can one also maintain an explicit matching of similar  
 895 approximation ratio within the same time bounds? We note that unless all greater than 1.5 approxi-  
 896 mation ratios are possible for explicit matching in polylog time, then Theorem 1.2 would imply a  
 897 separation between the attainable time/approximate tradeoffs for explicit dynamic matching and  
 898 its size estimation counterpart.

900 **Better approximation in  $o(n)$  update time?** Known conditional impossibility results rule out an  
 901 exact algorithm with  $n^{1-\Omega(1)}$  update time [2, 40, 58], but the best approximation ratios currently  
 902 known are  $\frac{3}{2} + \epsilon$  [22, 23, 56, 62]. Can one do better in  $o(n)$  time? Following the publication of the  
 903 conference version of this paper, the sublinear+streaming connection to dynamic algorithms has  
 904 resulted in better time/approximation tradeoffs for the value version of the problem [10, 18, 29, 30].  
 905 Again, one may ask, can explicit matching algorithms with such improved guarantees be obtained?  
 906 In a similar vein, sub-logarithmic speedups were obtained for  $(1 + o(1))$ -approximate matching [5].  
 907 Can one obtain stronger speedups for  $(1 + \epsilon)$ -approximate matching? On the flip side, can we show  
 908 any (conditional) hardness of *approximate* dynamic matching, for any approximation ratio?

909 **Unconditional impossibility results.** With this work we bring dynamic matching with better-  
 910 than-two approximation into the polylogarithmic update time regime—the range where *uncondi-  
 911 tional* impossibility are known for numerous data structures and dynamic algorithms. Can such  
 912 unconditional impossibility results be proven for (approximate) dynamic matching?

913 **More applications of robust AMMs.** A plethora of recent developments in fast (static) graph  
 914 algorithms rely on invocations of adversarially-robust dynamic algorithms (i.e., ones that work  
 915 against an adaptive adversary). Our work adds almost-maximal matchings to the list of algorithmic  
 916 subroutines useful for this approach. (Subsequently, [32] derandomized this result, though at the  
 917 cost of amortization.) Robust AMMs played a key role for dynamic problems, both here, in the follow-  
 918 up [10], and in the bounded-arboricity matching algorithms of [77]. What further applications  
 919 can our robust dynamic AMM algorithms for arbitrary graphs find for other (dynamic and static)  
 920 algorithmic problems?

922 **Acknowledgements.** We thank the anonymous SODA and J.ACM reviewers for helpful comments.

## 925 REFERENCES

- [1] Amir Abboud and Søren Dahlgaard. 2016. Popular conjectures as a barrier for dynamic planar graph algorithms. In *Proceedings of the 57th Symposium on Foundations of Computer Science (FOCS)*. 477–486.
- [2] Amir Abboud and Virginia Vassilevska Williams. 2014. Popular conjectures imply strong lower bounds for dynamic problems. In *Proceedings of the 55th Symposium on Foundations of Computer Science (FOCS)*. 434–443.
- [3] Kook Jin Ahn and Sudipto Guha. 2013. Linear programming in the semi-streaming model with application to the maximum matching problem. *Information and Computation* 222 (2013), 59–79.

[4] Moab Arar, Shiri Chechik, Sarel Cohen, Cliff Stein, and David Wajc. 2018. Dynamic Matching: Reducing Integral Algorithms to Approximately-Maximal Fractional Algorithms. In *Proceedings of the 45th International Colloquium on Automata, Languages and Programming (ICALP)*. 79:1–79:16.

[5] Sepehr Assadi, Soheil Behnezhad, Sanjeev Khanna, and Huan Li. 2023. On regularity lemma and barriers in streaming and dynamic matching. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing (STOC)*. 131–144.

[6] Sepehr Assadi, Arun Jambulapati, Yujia Jin, Aaron Sidford, and Kevin Tian. 2022. Semi-Streaming Bipartite Matching in Fewer Passes and Optimal Space. In *Proceedings of the 33rd Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 627–669.

[7] Sepehr Assadi and Sanjeev Khanna. 2024. Improved Bounds for Fully Dynamic Matching via Ordered Ruzsa-Szemerédi Graphs. *arXiv preprint arXiv:2406.13573* (2024).

[8] Sepehr Assadi, Sanjeev Khanna, and Yang Li. 2019. The stochastic matching problem with (very) few queries. *ACM Transactions on Economics and Computation (TEAC)* 7, 3 (2019), 1–19.

[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 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 1345–1364.

[10] Amir Azarmehr, Soheil Behnezhad, and Mohammad Roghani. 2024. Fully Dynamic Matching:  $(2 - \sqrt{2})$ -Approximation in Polylog Update Time. In *Proceedings of the 35th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 3040–3061.

[11] Surender Baswana, Manoj Gupta, and Sandeep Sen. 2015. Fully Dynamic Maximal Matching in  $O(\log n)$  Update Time. *SIAM Journal on Computing (SICOMP)* 44, 1 (2015), 88–113.

[12] Soheil Behnezhad. 2022. Time-optimal sublinear algorithms for matching and vertex cover. In *Proceedings of the 62nd Symposium on Foundations of Computer Science (FOCS)*. 873–884.

[13] Soheil Behnezhad. 2023. Dynamic Algorithms for Maximum Matching Size. In *Proceedings of the 34th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 129–162.

[14] Soheil Behnezhad, Mahsa Derakhshan, MohammadTaghi Hajiaghayi, Cliff Stein, and Madhu Sudan. 2019. Fully dynamic maximal independent set with polylogarithmic update time. In *Proceedings of the 60th Symposium on Foundations of Computer Science (FOCS)*. 382–405.

[15] Soheil Behnezhad and Alma Ghafari. 2024. Fully Dynamic Matching and Ordered Ruzsa-Szemerédi Graphs. In *Proceedings of the 65th Symposium on Foundations of Computer Science (FOCS)*. To appear.

[16] Soheil Behnezhad and Sanjeev Khanna. 2022. New Trade-Offs for Fully Dynamic Matching via Hierarchical EDCS. In *Proceedings of the 33rd Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 3529–3566.

[17] Soheil Behnezhad, Jakub Łącki, and Vahab Mirrokni. 2020. Fully Dynamic Matching: Beating 2-Approximation in  $\Delta^\epsilon$  Update Time. In *Proceedings of the 31st Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 2492–2508.

[18] Soheil Behnezhad, Mohammad Roghani, and Aviad Rubinstein. 2023. Sublinear time algorithms and complexity of approximate maximum matching. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing (STOC)*. 267–280.

[19] Amos Beimel, Haim Kaplan, Yishay Mansour, Kobbi Nissim, Thatchaphol Saranurak, and Uri Stemmer. 2022. Dynamic algorithms against an adaptive adversary: Generic constructions and lower bounds. In *Proceedings of the 54th Annual ACM Symposium on Theory of Computing (STOC)*. 1671–1684.

[20] Aaron Bernstein, Aditi Dudeja, and Zachary Langley. 2021. A Framework for Dynamic Matching in Weighted Graphs. In *Proceedings of the 53rd Annual ACM Symposium on Theory of Computing (STOC)*. 668–681.

[21] Aaron Bernstein, Sebastian Forster, and Monika Henzinger. 2019. A deamortization approach for dynamic spanner and dynamic maximal matching. In *Proceedings of the 30th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 1899–1918.

[22] Aaron Bernstein and Cliff Stein. 2015. Fully Dynamic Matching in Bipartite Graphs. In *Proceedings of the 42nd International Colloquium on Automata, Languages and Programming (ICALP)*. 167–179.

[23] Aaron Bernstein and Cliff Stein. 2016. Faster fully dynamic matchings with small approximation ratios. In *Proceedings of the 27th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 692–711.

[24] Sayan Bhattacharya, Deeparnab Chakrabarty, and Monika Henzinger. 2020. Deterministic dynamic matching in  $O(1)$  update time. *Algorithmica* 82, 4 (2020), 1057–1080.

[25] Sayan Bhattacharya, Monika Henzinger, and Giuseppe F Italiano. 2018. Deterministic fully dynamic data structures for vertex cover and matching. *SIAM Journal on Computing (SICOMP)* 47, 3 (2018), 859–887.

[26] Sayan Bhattacharya, Monika Henzinger, and Danupon Nanongkai. 2016. New deterministic approximation algorithms for fully dynamic matching. In *Proceedings of the 48th Annual ACM Symposium on Theory of Computing (STOC)*. 398–411.

[27] Sayan Bhattacharya, Monika Henzinger, and Danupon Nanongkai. 2017. Fully Dynamic Approximate Maximum Matching and Minimum Vertex Cover in  $O(\log^3 n)$  Worst Case Update Time. In *Proceedings of the 28th Annual*

981        *ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 470–489.

982 [28] Sayan Bhattacharya and Peter Kiss. 2021. Deterministic Rounding of Dynamic Fractional Matchings. In *Proceedings of the 48th International Colloquium on Automata, Languages and Programming (ICALP)*. 27:1–27:14.

983 [29] Sayan Bhattacharya, Peter Kiss, and Thatchaphol Saranurak. 2023. Sublinear Algorithms for  $(1.5 + \epsilon)$ -Approximate

984 Matching. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing (STOC)*. 254–266.

985 [30] Sayan Bhattacharya, Peter Kiss, and Thatchaphol Saranurak. 2024. Dynamic  $(1 + \epsilon)$ -Approximate Matching Size

986 in Truly Sublinear Update Time. In *Proceedings of the 56th Symposium on Foundations of Computer Science (FOCS)*. 1563–1588.

987 [31] Sayan Bhattacharya, Peter Kiss, Thatchaphol Saranurak, and David Wajc. 2023. Dynamic matching with better-than-2

988 approximation in polylogarithmic update time. In *Proceedings of the 2023 Annual ACM-SIAM Symposium on Discrete*

989 Algorithms (SODA). SIAM, 100–128.

990 [32] Sayan Bhattacharya, Peter Kiss, Aaron Sidford, and David Wajc. 2024. Near-Optimal Dynamic Rounding of Fractional

991 Matchings in Bipartite Graphs. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing (STOC)*. 59–70.

992 [33] Sayan Bhattacharya and Janardhan Kulkarni. 2019. Deterministically Maintaining a  $(2 + \epsilon)$ -Approximate Minimum

993 Vertex Cover in  $O(1/\epsilon^2)$  Amortized Update Time. In *Proceedings of the 30th Annual ACM-SIAM Symposium on Discrete*

994 *Algorithms (SODA)*. 1872–1885.

995 [34] Jan van den Brand, Danupon Nanongkai, and Thatchaphol Saranurak. 2019. Dynamic matrix inverse: Improved

996 algorithms and matching conditional lower bounds. In *Proceedings of the 60th Symposium on Foundations of Computer*

997 *Science (FOCS)*. 456–480.

998 [35] Moses Charikar and Shay Solomon. 2018. Fully Dynamic Almost-Maximal Matching: Breaking the Polynomial

999 Barrier for Worst-Case Time Bounds. In *Proceedings of the 45th International Colloquium on Automata, Languages and*

1000 *Programming (ICALP)*. 33:1–33:14.

1001 [36] Shiri Chechik and Tianyi Zhang. 2019. Fully dynamic maximal independent set in expected poly-log update time. In *Proceedings of the 60th Symposium on Foundations of Computer Science (FOCS)*. 370–381.

1002 [37] Jiale Chen, Aaron Sidford, and Ta-Wei Tu. 2023. Entropy Regularization and Faster Decremental Matching in General

1003 Graphs. *arXiv preprint arXiv:2312.09077* (2023).

1004 [38] Julia Chuzhoy and Sanjeev Khanna. 2019. A new algorithm for decremental single-source shortest paths with

1005 applications to vertex-capacitated flow and cut problems. In *Proceedings of the 51st Annual ACM Symposium on Theory*

1006 *of Computing (STOC)*. 389–400.

1007 [39] Michael Crouch and Daniel M Stubbs. 2014. Improved Streaming Algorithms for Weighted Matching, via Unweighted

1008 Matching. In *Proceedings of the 17th International Conference on Approximation Algorithms for Combinatorial Optimization*

1009 *Problems (APPROX)*. 96.

1010 [40] Søren Dahlgaard. 2016. On the Hardness of Partially Dynamic Graph Problems and Connections to Diameter. In *Proceedings of the 43rd International Colloquium on Automata, Languages and Programming (ICALP)*. 48:1–48:14.

1011 [41] Ran Duan and Seth Pettie. 2014. Linear-time approximation for maximum weight matching. *Journal of the ACM* (JACM) 61, 1 (2014), 1.

1012 [42] Devdatt P. Dubhashi and Desh Ranjan. 1998. Balls and bins: A study in negative dependence. *Random Struct. Algorithms* 13, 2 (1998), 99–124.

1013 [43] Aditi Dudeja. 2024. A Note on Rounding Matchings in General Graphs. *arXiv preprint arXiv:2402.03068* (2024).

1014 [44] Jack Edmonds. 1965. Maximum matching and a polyhedron with 0, 1-vertices. *Journal of research of the National*

1015 *Bureau of Standards B* 69, 125–130 (1965), 55–56.

1016 [45] Jack Edmonds. 1965. Paths, trees, and flowers. *Canadian Journal of mathematics* 17, 3 (1965), 449–467.

1017 [46] Sebastian Eggert, Lasse Kliemann, and Anand Srivastav. 2009. Bipartite graph matchings in the semi-streaming model.

1018 In *Proceedings of the 16th Annual European Symposium on Algorithms (ESA)*. 492–503.

1019 [47] Leah Epstein, Asaf Levin, Danny Segev, and Oren Weimann. 2013. Improved Bounds for Online Preemptive Matching.

1020 In *Proceedings of the 30th International Symposium on Theoretical Aspects of Computer Science (STACS)*. 389.

1021 [48] Hossein Esfandiari, MohammadTaghi Hajiaghayi, and Morteza Monemizadeh. 2016. Finding large matchings in

1022 semi-streaming. In *2016 IEEE 16th International Conference on Data Mining Workshops (ICDMW)*. 608–614.

1023 [49] Matthew Fahrbach, Gary L Miller, Richard Peng, Saurabh Sawlani, Junxing Wang, and Shen Chen Xu. 2018. Graph

1024 sketching against adaptive adversaries applied to the minimum degree algorithm. In *Proceedings of the 59th Symposium*

1025 *on Foundations of Computer Science (FOCS)*. 101–112.

1026 [50] Joan Feigenbaum, Sampath Kannan, Andrew McGregor, Siddharth Suri, and Jian Zhang. 2005. On graph problems in a

1027 semi-streaming model. *Theoretical Computer Science (TCS)* 348, 2-3 (2005), 207–216.

1028 [51] Moran Feldman and Ariel Szarf. 2022. Maximum Matching sans Maximal Matching: A New Approach for Finding

1029 Maximum Matchings in the Data Stream Model. In *Proceedings of the 25th International Conference on Approximation*

1030 *Algorithms for Combinatorial Optimization Problems (APPROX)*. 33:1–33:24.

1030 [52] Manuela Fischer, Slobodan Mitrović, and Jara Uitto. 2022. Deterministic  $(1 + \epsilon)$ -approximate maximum matching  
 1031 with  $\text{poly}(1/\epsilon)$  passes in the semi-streaming model and beyond. In *Proceedings of the 54th Annual ACM Symposium  
 1032 on Theory of Computing (STOC)*. 248–260.

1033 [53] Mohsen Ghaffari and David Wajc. 2019. Simplified and Space-Optimal Semi-Streaming  $(2 + \epsilon)$ -Approximate Matching.  
 1034 In *Proceedings of the 2nd Symposium on Simplicity in Algorithms (SOSA)*. 13:1–13:8.

1035 [54] Ashish Goel, Michael Kapralov, and Sanjeev Khanna. 2012. On the communication and streaming complexity of  
 1036 maximum bipartite matching. In *Proceedings of the 23rd Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*.  
 1037 468–485.

1038 [55] Ashish Goel, Michael Kapralov, and Sanjeev Khanna. 2013. Perfect Matchings in  $O(n \log n)$  Time in Regular Bipartite  
 1039 Graphs. *SIAM Journal on Computing (SICOMP)* 42, 3 (2013), 1392–1404.

1040 [56] Fabrizio Grandoni, Chris Schwiegelshohn, Shay Solomon, and Amitai Uzrad. 2022. Maintaining an EDCS in General  
 1041 Graphs: Simpler, Density-Sensitive and with Worst-Case Time Bounds. *Proceedings of the 5th Symposium on Simplicity  
 1042 in Algorithms (SOSA)* (2022), 12–23.

1043 [57] Manoj Gupta and Richard Peng. 2013. Fully dynamic  $(1 + \epsilon)$ -approximate matchings. In *Proceedings of the 54th  
 1044 Symposium on Foundations of Computer Science (FOCS)*. 548–557.

1045 [58] Monika Henzinger, Sebastian Krinninger, Danupon Nanongkai, and Thatchaphol Saranurak. 2015. Unifying and  
 1046 strengthening hardness for dynamic problems via the online matrix-vector multiplication conjecture. In *Proceedings of  
 1047 the 47th Annual ACM Symposium on Theory of Computing (STOC)*. 21–30.

1048 [59] Sagar Kale and Sumedh Tirodkar. 2017. Maximum Matching in Two, Three, and a Few More Passes Over Graph  
 1049 Streams. In *Proceedings of the 20th International Conference on Approximation Algorithms for Combinatorial Optimization  
 1050 Problems (APPROX)*. 15:1–15:21.

1051 [60] Michael Kapralov. 2013. Better bounds for matchings in the streaming model. In *Proceedings of the 24th Annual  
 1052 ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 1679–1697.

1053 [61] Michael Kapralov. 2021. Space lower bounds for approximating maximum matching in the edge arrival model. In  
 1054 *Proceedings of the 32nd Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 1874–1893.

1055 [62] Peter Kiss. 2022. Improving update times of dynamic matching algorithms from amortized to worst case. *Proceedings  
 1056 of the 13th Innovations in Theoretical Computer Science Conference (ITCS)* (2022), 94:1–94:21.

1057 [63] Christian Konrad. 2018. A simple augmentation method for matchings with applications to streaming algorithms. In  
 1058 *Proceedings of the 43rd International Symposium on Mathematical Foundations of Computer Science (MFCS)*. 74:1–74:16.

1059 [64] Christian Konrad, Frédéric Magniez, and Claire Mathieu. 2012. Maximum matching in semi-streaming with few  
 1060 passes. In *Proceedings of the 15th International Conference on Approximation Algorithms for Combinatorial Optimization  
 1061 Problems (APPROX)*. 231–242.

1062 [65] Christian Konrad and Kheeran K Naidu. 2021. On Two-Pass Streaming Algorithms for Maximum Bipartite Matching. In  
 1063 *Proceedings of the 24th International Conference on Approximation Algorithms for Combinatorial Optimization Problems  
 1064 (APPROX)*. 19:1–19:18.

1065 [66] Tsvi Kopelowitz, Seth Pettie, and Ely Porat. 2016. Higher lower bounds from the 3SUM conjecture. In *Proceedings of  
 1066 the 27th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 1272–1287.

1067 [67] Harold W Kuhn. 1955. The Hungarian method for the assignment problem. *Naval research logistics quarterly* 2, 1-2  
 1068 (1955), 83–97.

1069 [68] Hung Le, Lazar Milenković, Shay Solomon, and Virginia Vassilevska Williams. 2022. Dynamic Matching Algorithms  
 1070 Under Vertex Updates. In *Proceedings of the 13th Innovations in Theoretical Computer Science Conference (ITCS)*.  
 96:1–96:24.

1071 [69] Yang P Liu. 2024. On approximate fully-dynamic matching and online matrix-vector multiplication. In *Proceedings of  
 1072 the 65th Symposium on Foundations of Computer Science (FOCS)*. To appear.

1073 [70] Andrew McGregor. 2005. Finding graph matchings in data streams. In *Proceedings of the 8th International Conference  
 1074 on Approximation Algorithms for Combinatorial Optimization Problems (APPROX)*. 170–181.

1075 [71] Danupon Nanongkai and Thatchaphol Saranurak. 2017. Dynamic spanning forest with worst-case update time:  
 1076 adaptive, Las Vegas, and  $O(n^{1/2-\epsilon})$ -time. In *Proceedings of the 49th Annual ACM Symposium on Theory of Computing  
 1077 (STOC)*. 1122–1129.

1078 [72] Huy N Nguyen and Krzysztof Onak. 2008. Constant-time approximation algorithms via local improvements. In  
 1079 *Proceedings of the 49th Symposium on Foundations of Computer Science (FOCS)*. 327–336.

1080 [73] Krzysztof Onak, Dana Ron, Michal Rosen, and Ronitt Rubinfeld. 2012. A near-optimal sublinear-time algorithm for  
 1081 approximating the minimum vertex cover size. In *Proceedings of the 23rd Annual ACM-SIAM Symposium on Discrete  
 1082 Algorithms (SODA)*. 1123–1131.

1083 [74] Krzysztof Onak and Ronitt Rubinfeld. 2010. Maintaining a large matching and a small vertex cover. In *Proceedings of  
 1084 the 42nd Annual ACM Symposium on Theory of Computing (STOC)*. 457–464.

1079 [75] Michal Parnas and Dana Ron. 2007. Approximating the minimum vertex cover in sublinear time and a connection to  
 1080 distributed algorithms. *Theoretical Computer Science (TCS)* 381, 1-3 (2007), 183–196.

1081 [76] Ami Paz and Gregory Schwartzman. 2018. A  $(2 + \epsilon)$ -Approximation for Maximum Weight Matching in the Semi-  
 1082 streaming Model. *ACM Transactions on Algorithms (TALG)* 15, 2 (2018), 18.

1083 [77] David Peleg and Shay Solomon. 2016. Dynamic  $(1 + \epsilon)$ -approximate matchings: a density-sensitive approach. In  
 1084 *Proceedings of the 27th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 712–729.

1085 [78] Mohammad Roghani, Amin Saberi, and David Wajc. 2022. Beating the Folklore Algorithm for Dynamic Matching. In  
 1086 *Proceedings of the 13th Innovations in Theoretical Computer Science Conference (ITCS)*. 111:1–111:23.

1087 [79] Piotr Sankowski. 2007. Faster dynamic matchings and vertex connectivity. In *Proceedings of the 18th Annual ACM-SIAM  
 1088 Symposium on Discrete Algorithms (SODA)*. 118–126.

1089 [80] Shay Solomon. 2016. Fully dynamic maximal matching in constant update time. In *Proceedings of the 57th Symposium  
 1090 on Foundations of Computer Science (FOCS)*. 325–334.

1091 [81] Vadim G Vizing. 1964. On an estimate of the chromatic class of a p-graph. *Diskret analiz* 3 (1964), 25–30.

1092 [82] David Wajc. 2020. Rounding dynamic matchings against an adaptive adversary. In *Proceedings of the 52nd Annual  
 1093 ACM Symposium on Theory of Computing (STOC)*. 194–207.

1094 [83] Yuichi Yoshida, Masaki Yamamoto, and Hiro Ito. 2012. Improved constant-time approximation algorithms for maximum  
 1095 matchings and other optimization problems. *SIAM Journal on Computing (SICOMP)* 41, 4 (2012), 1074–1093.

## APPENDIX

### A PROOFS OF BASIC BUILDING BLOCKS

1097 Here we substantiate some key propositions implied by prior work. We stress that we provide  
 1098 proofs mostly for completeness, due to our propositions being slight variants or being differently  
 1099 organized than their previous counterparts. That is, we do not claim novelty of the underlying  
 1100 ideas of this section.

#### A.1 Proof of Proposition 2.1

1103 A key component of Proposition 2.1 is the following vertex sparsification technique for dynamic  
 1104 settings by Kiss [62], adapted from such a vertex sparsification of Assadi et al. [9] in the context of  
 1105 stochastic matching.

1106 **Proposition A.1.** *There exists a randomized algorithm which for each update to  $G$  makes an update  
 1107 to  $O\left(\frac{\log^2 n}{\epsilon^3}\right)$  contracted subgraphs, such that w.h.p. throughout any (possibly adaptively generated)  
 1108 update sequence, one subgraph  $G'$  has a matching of cardinality  $\mu(G) \cdot (1 - O(\epsilon))$  and nodeset of  
 1109 size  $n' \leq \mu(G)/\epsilon$ . Moreover, any matching  $M'$  in  $G'$  can be transformed into a matching in  $G$  of  
 1110 cardinality  $|M'|$  in time  $O(|M'|)$ . For any matching  $M'$  in any  $G'$  undergoing edge updates we can  
 1111 maintain a matching of cardinality  $|M'|$  in  $G$  with  $O(1)$  worst-case update time.*

1113 **PROOF.** Consider a random graph  $G'$  obtained by hashing each node into one of  $k/\epsilon$  buckets, for  
 1114 some integer  $k$ , and contracting all nodes that are hashed into the same bin. That is, two contracted  
 1115 nodes neighbor in  $G'$  if their corresponding bins contain neighboring nodes in  $G$ . By storing for each  
 1116 edge  $e$  in  $G'$  a list of edges inducing  $e$ , we can easily transform a matching  $M'$  in  $G'$  to a matching in  
 1117  $G$  of the same cardinality in time  $O(|M'|)$ . The majority of this proof is thus dedicated to showing  
 1118 that  $O\left(\frac{\log n}{\epsilon^2}\right)$  such contractions for each value  $k = \lceil(1 + \epsilon)^i\rceil$  with  $i \in [\log_{1+\epsilon}(n)] \subseteq \left[O\left(\frac{\log n}{\epsilon}\right)\right]$   
 1119 suffice to guarantee that one of these  $G'$  contains a matching of cardinality at least  $\mu(G) \cdot (1 - 3\epsilon)$ .  
 1120

1121 Fix an integer  $i$  and  $k = \lceil(1 + \epsilon)^i\rceil \leq n$ . Fix a matching  $M$  in  $G$  of cardinality  $|M| \leq k$ . The  
 1122 probability that a vertex  $v$  incident on some edge of  $M$  is contracted into a separate bin than the  
 1123 other  $2|M| - 1$  endpoints can be expressed as follows:

$$1124 (1 - 1/(k/\epsilon - 1))^{2|M|-1} \geq \left(1 - \frac{2 \cdot \epsilon}{k}\right)^{2 \cdot k} \geq (1 - 5 \cdot \epsilon).$$

1128 Thus, by linearity, the number  $X$  of such endpoints of edges of  $M$  satisfy that  $\mathbb{E}[X] \geq (1 - 5\epsilon) \cdot 2|M|$ .  
 1129 Observe that  $X$  is the sum of negatively associated random variables, by [42], since the hashing of  
 1130 vertices is equivalent to the folklore balls and bins experiment, so by standard Chernoff Bounds,

$$1131 \quad 1132 \quad \Pr[X \leq 2 \cdot |M| \cdot (1 - 6\epsilon)] \leq \exp(-\Theta(\epsilon^2 |M|)).$$

1133 If at least  $2|M| \cdot (1 - 6\epsilon)$  endpoints of  $M$  are hashed to unique vertices then at least  $|M| - 2|M| \cdot 6\epsilon \geq$   
 1134  $|M| \cdot (1 - 12\epsilon)$  edges of  $M$  had both of their endpoints assigned to unique vertices in  $G'$  hence are  
 1135 present in  $G'$ .

1136 We say that the contraction is *bad* if for *some* matching  $M$  of cardinality in the range  $[k, k(1+\epsilon)+1]$   
 1137 if the number of edges of  $M$  that are not present in  $G'$  is lesser than  $|M| \cdot (1 - 12\epsilon)$ . Otherwise,  
 1138 it is *good*. Now, there are  $\sum_{i=k}^{k(1+\epsilon)+1} \binom{n}{i} \leq k\epsilon \cdot n^{k(1+\epsilon)+1} \leq n^{k(1+\epsilon)+2}$  possible matchings of size  
 1139  $|M| \in [k, k(1 + \epsilon)]$ . Therefore, by randomly contracting the graph for range  $[k, k(1 + \epsilon) + 1]$  some  
 1140  $\frac{C \log n}{\epsilon^2}$  many times, for a sufficiently large  $C$ , we have that the probability that all contractions for  
 1141 range  $[k, k(1 + \epsilon) + 1]$  are bad is

$$1143 \quad 1144 \quad \Pr[\text{all contractions are bad}] \leq n^{k(1+\epsilon)+2} \cdot \exp\left(-\Theta(\epsilon^2 k) \cdot \frac{C \log n}{\epsilon^2}\right) \leq n^{-3}.$$

1145 Therefore, taking union bound over the  $\log_{1+\epsilon}(n)$  possible value of  $k$ , we find that with high  
 1146 probability, each range  $[k, k(1 + \epsilon) + 1]$  has some good contraction.

1147 We conclude that, w.h.p., among the  $O\left(\frac{\log^2 n}{\epsilon^3}\right)$  contracted graphs, there exists a good contraction  
 1148 for every  $k = \lceil (1 + \epsilon)^i \rceil$ , and in particular for  $k \leq \mu(G) \leq k(1 + \epsilon) + 1$ . That is, one of the contracted  
 1149 graphs contains a large matching,  $\mu(G') \geq \mu(G) \cdot (1 - 12\epsilon)$ , and has few nodes,  $n' \leq k/\epsilon \leq \mu(G)/\epsilon$ ,  
 1150 as desired.  $\square$

1153 We now proceed towards proving Proposition 2.1, restated below for ease of reference.  
 1154

1155 **Proposition 2.1.** *Let  $\epsilon \in (0, 1)$  and  $\alpha \geq 1$ . Suppose there exists an algorithm  $\mathcal{A}$  on a dynamic  
 1156  $n$ -node graph  $G$  with update time  $t_u$ , that, provided  $\mu(G) \geq \epsilon \cdot n$ , supports  $t_q$ -time  $\alpha$ -approximate  
 1157 size estimate queries w.h.p. Then, there is another algorithm  $\mathcal{A}'$  on  $G$  that always maintains an  
 1158  $(\alpha + O(\epsilon))$ -approximate estimate  $v'$  in  $\tilde{O}_\epsilon(t_u + t_q/n)$  update time. Moreover if the update time of  $\mathcal{A}$  is  
 1159 worst-case, so is that of  $\mathcal{A}'$ , and if  $\mathcal{A}$  works against an adaptive adversary, then so does  $\mathcal{A}'$ .*

1160 **PROOF.** Let  $\epsilon' = \alpha' \cdot \epsilon \cdot 2$  (here  $\alpha'$  is some  $O(1)$  factor). Using the algorithm described by  
 1161 Proposition A.1 we can generate  $T = \tilde{O}_\epsilon(1)$  graphs  $G_i : i \in [T]$  with the following properties: A)  
 1162  $\mu(G_i) \leq \mu(G)$  for all  $i \in [T]$ , B) There is an  $i \in [T]$  satisfying that  $\mu(G_i) \geq (1 - \epsilon') \cdot \mu(G)$  and  
 1163  $\mu(G_i) \geq n \cdot \epsilon'$ , C) All sub-graphs  $G_i$  undergo a single update when  $G$  undergoes an update.

1164 Our algorithm proceeds as follows: on all  $T$  generated sub-graphs we run algorithm  $\mathcal{A}$  at all  
 1165 times. Furthermore, on each sub-graph we maintain an  $O(1) = \alpha'$ -approximate estimate on the  
 1166 maximum matching size  $\tilde{\mu}_i$  using algorithms from literature (randomized against an adaptive  
 1167 adversary) in  $\tilde{O}_\epsilon(1)$  worst-case time. For all sub-graphs we monitor the relationship of  $\tilde{\mu}_i$  and  $|V_i|$ .  
 1168 If  $\tilde{\mu}_i$  increases above the threshold of  $|V_i| \cdot \epsilon$  we start a run of the query algorithm on  $G_i$  returning  
 1169 us an  $\alpha$ -approximate estimate of  $\mu(G_i)$  which will define  $v'_i$ . We distribute the work of this query  
 1170 over  $|V_i| \cdot (\epsilon)^2$  updates and re-initiate the query every  $|V_i| \cdot (\epsilon)^2$  updates. The matching size queries  
 1171 of  $G_i$  always run on the state of  $G_i$  at the start of the query (even though  $G_i$  undergoes updates  
 1172 during its run). If  $\tilde{\mu}_i$  decreases below the threshold of  $|V_i| \cdot \epsilon$  we stop the querying process and set  
 1173  $v'_i = 0$ . Note that at initialization we just set  $v'_i = \mu(G_i)$  for all  $i \in [T]$  statically.

1174 At all times we maintain the output  $\max_{i \in R} v'_i$ , the maximum of our matching size estimates.

1177 **Algorithm 3** Vertex set sparsification

---

```

1178 1: Initialize  $v'_i = \mu(G_i)$ 
1179 2: Maintain contracted sub-graphs  $G_i$  and  $\alpha'$ -approximate matching size estimates  $\tilde{\mu}_i$ 
1180 3: Run algorithm  $\mathcal{A}$  on every  $G_i$ 
1181 4: for  $i \in [T]$  do
1182 5:   if  $\tilde{\mu}_i$  becomes at least  $|V_i| \cdot \epsilon$  then
1183 6:     Initiate a matching size query of  $G_i$  in  $O(t_q)$  time on the current state of  $G_i$ 
1184 7:     Distribute the work over the next  $|V_i| \cdot \epsilon^2$  updates
1185 8:     Repeatedly recompute distributed over every  $|V_i| \cdot \epsilon^2$  updates
1186 9:     Let  $v'_i$  be the latest finished estimate
1187 10:   if  $\tilde{\mu}_i$  reduces below  $|V_i| \cdot \epsilon$  then
1188 11:     Terminate the querying process of  $\mu(G_i)$ 
1189 12:     Set  $v'_i \leftarrow 0$ 
1190 13: At all times return  $\max_{i \in [T]} v'_i$ 

```

---

1193 We first discuss the update time of Algorithm 3. The maintenance of the  $T$  contracted sub-graphs  
 1194 and matching size estimates  $\tilde{\mu}_i$  takes  $\tilde{O}_\epsilon(1)$  w.c. time. Running algorithm  $\mathcal{A}$  on each of the contracted  
 1195 sub-graphs takes update time  $\tilde{O}_\epsilon(t_u)$  (and is worst case if  $\mathcal{A}$  has worst-case update time). A matching  
 1196 size query will only be initiated and run on contracted sub-graph  $G_i$  if  $\tilde{\mu}_i \geq \mu(G_i)/\alpha' \geq |V_i| \cdot \epsilon$ , that  
 1197 is if  $\mu(G_i) \geq |V_i| \cdot \epsilon/2$ . Each re-computation of the estimate  $v'_i$  will be distributed over some  $|V_i| \cdot \epsilon^2$   
 1198 updates, that is, it will take  $O(t_q/(|V_i| \cdot \epsilon^2)) = O_\epsilon(t_q/|V_i|)$  worst-case time. Finding and returning  
 1199 the maximum matching size estimate  $v'_i$  takes  $\tilde{O}_\epsilon(1) = O(T)$  time. Therefore, the total update time  
 1200 of the algorithm is  $\tilde{O}_\epsilon(t_u + t_q \cdot \beta/n)$  and is worst-case if  $\mathcal{A}$  has worst-case update time. Furthermore,  
 1201 all components of the algorithm but  $\mathcal{A}$  are randomized against an adaptive adversary..

1202 It remains to argue that the algorithm maintains  $v'$  such that  $v' \leq \mu(G) \leq v' \cdot (\alpha + O(\epsilon))$  at all  
 1203 times. Say that  $G_i$  is a 'successful' contraction if  $G_i$  satisfies property B). By Proposition A.1, w.h.p.,  
 1204 there is a successful contraction at all times, at time point  $\tau_1$  let that contraction be  $G_i$ . We will  
 1205 separate two instances:

1206 i) Throughout the run of the algorithm at all times it held that  $\tilde{\mu}_i \geq |V_i| \cdot \epsilon$ : The algorithm has  
 1207 ran the matching size query sub-routine on  $G_i$  after every  $|V_i| \cdot \epsilon^2$  edge updates. Let  $G_i^{\tau_0}$  be the past  
 1208 state of the graph  $G_i$  when the algorithm started calculating the current estimate  $(v_i^{\tau_1})'$ . By the  
 1209 scheduling of this calculation we know that  $\tau_0 \geq \tau_1 - \epsilon^2 \cdot |V_i|$ . Hence,  $\mu(G_i^{\tau_0}) \geq \mu(G_i^{\tau_1}) - \epsilon^2 \cdot |V_i|$ ,  
 1210 where  $\mu(G_i^{\tau_1}) \geq \mu(G) \cdot (1 - \epsilon')$  and  $\mu(G_i^{\tau_1}) \geq |V_i| \cdot \epsilon'$ . Hence,  $(v_i^{\tau_1})' \cdot \alpha \cdot (1 + O(\epsilon)) \geq \mu(G)$ .

1211 ii) At time  $\tau_1$   $G_i$  is a successful contraction but at some prior point during the run of the algorithm  
 1212  $\tilde{\mu}_i$  became less than  $|V_i| \cdot \epsilon$ : we know that at some point  $\tau_0$  prior to  $\tau_1$   $\tilde{\mu}_i$  must have increased above  
 1213  $|V_i| \cdot \epsilon$ . Define the state of  $G_i$  at the two time points as  $G_i^{\tau_0}$  and  $G_i^{\tau_1}$  respectively. As at  $\tau_1$   $G_i^{\tau_1}$  is  
 1214 a successful contraction we know that  $\mu(G_i^{\tau_1}) \geq |V_i| \cdot \epsilon'$ . When  $\tilde{\mu}_i$  crossed the threshold at  $\tau_0$  it  
 1215 held that  $\tilde{\mu}_i = |V_i| \cdot \epsilon$  that is  $\mu(G_i^{\tau_0}) \leq |V_i| \cdot \epsilon \cdot \alpha$ . As per each update the maximum matching size  
 1216 may only change by 1 we have that  $\tau_1 - \tau_0 \geq |V_i| \cdot \epsilon \cdot \alpha$ . Hence, by time  $\tau_1$  the algorithm already  
 1217 had an updated estimate of  $v'_i$  (that is one calculated in the previous  $\epsilon^2 \cdot |V_i|$  updates such that  
 1218  $\mu(G_i) \geq |V_i| \cdot \epsilon$  during these updates). Here we can refer back to the previous case (pretending the  
 1219 algorithm initialized at  $\tau_0$ ).  $\square$

1220 **A.2 Proof of Proposition 2.6**

1221 We now give a proof extending standard arguments that small maximal matchings contain many  
 1222 length-three augmenting paths to showing that small  $\epsilon$ -AMM likewise contain many such paths.

1226 **Proposition 2.6.** Let  $\epsilon > 0$  and  $c \in \mathbb{R}$  and let  $M$  be an  $\epsilon$ -AMM in  $G$  such that  $|M| \leq (\frac{1}{2} + c) \cdot \mu(G)$ .  
 1227 Then  $M$  admits a collection of at least  $(\frac{1}{2} - 3c - \frac{7\epsilon}{2}) \cdot \mu(G)$  node-disjoint 3-augmenting paths.

1228 PROOF. The above bound for  $\epsilon = 0$  is well-known (see, e.g., [64]). We reduce to this case by  
 1229 removing the at most  $\epsilon \cdot \mu(G)$  nodes in  $V \setminus V(M)$  needed to make  $M$  maximal. This yields a graph  
 1230  $G'$  with  $\mu(G') \geq \mu(G) \cdot (1 - \epsilon)$ , and therefore  
 1231

$$1232 |M| \leq \left(\frac{1}{2} + c\right) \cdot \mu(G) \leq \frac{\frac{1}{2} + c}{1 - \epsilon} \cdot \mu(G') \leq \left(\frac{1}{2} + c + \epsilon\right) \cdot \mu(G'),$$

1234 Consequently, by the special case of this proposition with  $\epsilon = 0$ , we have that the maximum number  
 1235 of disjoint 3-augmenting paths that  $M$  admits in  $G'$  (and hence also in  $G$ ) is at least

$$1237 \left(\frac{1}{2} - 3(c + \epsilon)\right) \cdot \mu(G') \geq \left(\frac{1}{2} - 3(c + \epsilon)\right) (1 - \epsilon) \cdot \mu(G) \geq \left(\frac{1}{2} - 3(c + \epsilon) - \frac{\epsilon}{2}\right) \cdot \mu(G),$$

1239 as claimed.  $\square$

## 1240 B PROOF OF LEMMA 4.6

1242 In this section we prove Lemma 4.6 which we re-state for convenience.

1243 **Lemma 4.6.** Consider a graph  $G' = (V', E')$  with  $|V'| = n'$ , and a matching  $M$  with  $V(M) \subseteq V'$   
 1244 that is not necessarily part of  $G'$  (i.e., we might have  $M \not\subseteq E'$ ). For any matching  $M'$  in  $G'$ , let  $k_{M'}$   
 1245 denote the number of edges in  $M$  both of whose endpoints are matched in  $M'$ . There is an algorithm  
 1246 which, given adjacency matrix query access to the edges of  $G'$ , w.h.p. runs in  $\tilde{O}_\epsilon(n')$  time and returns  
 1247 an estimate  $\kappa \in [k_{M'} - \epsilon^2 n', k_{M'}]$  for some maximal matching  $M'$  in  $G'$ .

1248 Our proof of this lemma is a minor modification of the argument from [12, Section 5]. We claim  
 1249 no novelty for this proof. To make our notations consistent with the ones used by [12], we will focus  
 1250 on an  $n$ -node graph  $G = (V, E)$  (different from our dynamic input graph). Let  $\pi$  be a permutation  
 1251 of the edges of graph  $G = (V, E)$ . Let  $GMM(G, \pi)$  stand for the output of the greedy maximum  
 1252 matching algorithm when run on graph  $G$  with edge ordering  $\pi$ .

### 1254 B.1 Building blocks

1255 Lemma B.1 is explicitly concluded by [12], whereas Lemma B.2 is a slight modification of a con-  
 1256 struction appearing in Section 5 of [12] we need to fit our arguments.

1258 **Lemma B.1.** There is a randomized algorithm that in  $\tilde{O}(|E|/|V|)$  expected time returns the matched  
 1259 status of a random  $v$  under  $GMM(G, \pi)$ , for random  $\pi$ . This algorithm relies on list access to the edges  
 1260 of  $G$ .

1261 In order to prove Lemma 4.6 we have to work with adjacency matrix queries. Based on a slight  
 1262 modification of Section 5 of [12] we can derive the following tool for this purpose.

1263 **Lemma B.2.** Let  $\delta \in (0, 1/2)$ . For a given  $n$ -node graph  $G = (V, E)$  there exists a supergraph  
 1264  $H = (V_H, E_H)$  of  $G$  (i.e.,  $V_H \supseteq V$  and  $E_H \supseteq E$ ) satisfying the following:

- 1266 •  $|E_H| = \Theta_\delta(n^2)$ .
- 1267 •  $|V_H| = \Theta_\delta(n^2)$ .
- 1268 • At most  $\delta \cdot n$  nodes of  $V$  are matched to nodes in  $V_H \setminus V$  by  $GMM(H, \pi)$ , w.h.p. over  $\pi$ .
- 1269 •  $GMM(H, \pi) \cap E$  is a maximal matching in  $G[V \setminus V_{slack}]$ , where  $V_{slack} \subseteq V$  are nodes in  $V$   
 1270 that are matched to nodes in  $V \setminus V_H$ .
- 1271 • Any adjacency list query to  $E_H$  (querying the  $i$ -th neighbour of a vertex according to some  
 1272 ordering of neighbours) can be implemented using one adjacency matrix query to  $E$  (querying  
 1273 the existence of any edge  $(u, v)$ ).

1274

Informally, the main change in our construction compared to that of [12] is that our construction will allow us to argue that the random matching in the constructed graph  $H$  is, w.h.p., a maximal matching after ignoring a small set of nodes. In contrast, the construction in [12] resulted in an “expected” version of this guarantee. As the high-probability bounds will simplify our discussion later, we modify this construction below. The second change we make is in externalizing the fact that the matching computed this way is maximal, rather than 2-approximate, as stated in [12]. We now turn to proving the above lemma.

PROOF OF LEMMA B.2. The node-set of  $H$  is  $V_H := V \cup V^* \cup (W_1, \dots, W_n) \cup (U_1, \dots, U_n)$ , where  $V = \{v_1, v_2, \dots, v_n\}$  (note that  $G = (V, E_H[V])$ ),  $V^* = \{v_2^*, v_2^*, \dots, v_n^*\}$ ,  $W_i = \{w_i^1, w_i^2, \dots, w_i^n\}$ , and the set  $U_i = \{u_i^1, u_i^2, \dots, u_i^s\}$  is of size  $s := 10n/\delta$  for all  $i \in [n]$ . To specify the edge-set  $E_H$ , we now define the *ordered* adjacency list for every node  $v \in V_H$ .

- Every node  $v_i \in V$  has degree exactly  $n$ : For any  $j \in [n]$ , if  $(v_i, v_j) \in E$  then the  $j^{th}$  neighbor of  $v_i$  is the node  $v_j \in V$ , otherwise it is the node  $v_j^* \in V^*$ .
- Every node  $v_i^* \in V^*$  has degree exactly  $n+s$ : For any  $j \in [n]$ , if  $(v_i, v_j) \in E$  then the  $j^{th}$  neighbor of  $v_i^*$  is the node  $w_i^j \in W_i$ , otherwise it is the node  $v_j \in V$ . Furthermore, for all  $j \in [s]$ , the  $(n+j)^{th}$  neighbor of  $v_i^*$  is the node  $u_i^j \in U_i$ .
- Each node in  $U_j$ , for any  $j \in [n]$ , has only one neighbor (which is  $v_j^*$ ).
- Node  $w_i^j$  may have degree at most one: if  $(v_i, v_j) \in E$  then  $w_i^j$  is a neighbour of  $v_i^* \in V^*$ , otherwise it is an isolated vertex of  $H$ .

Note that  $|V_H| = 2n + n^2 + ns = \Theta(n^2/\delta)$  and that similarly  $|E_H| = n^2 + |E| + ns = \Theta(n^2/\delta)$ . Furthermore, from the above discussion it is immediate that an adjacency list query to  $E_H$  (i.e., querying for the  $j$ -th neighbor of a vertex) can be implemented using at most one adjacency matrix query to  $E$ . It remains to prove the remaining two properties of  $H$ .

To this end, recall that  $V_{slack}$  denotes the set of vertices in  $V$  matched to  $V^*$  nodes. Then, by maximality of  $GMM(H, \pi)$ , we have that  $GMM(H, \pi) \cap E$  is indeed maximal matching of  $G[V \setminus V_{slack}]$ . We now turn to bound  $|V_{slack}|$ . To this end, we say a node  $v^* \in V^*$  is *occupied* if its earliest edge in  $\pi$  has its other endpoint in  $W_i$  or  $U_i$ . Trivially, such an occupied vertex  $v^* \in V^*$  is matched to a vertex of  $U_{v^*} \cup W_{v^*}$  under  $GMM(H, \pi)$ . The following simple claim, which follows by a Chernoff bound together with the simple observation that it is unlikely for a node in  $V^*$  to be matched in  $V$  (and thus contribute to  $|V_{slack}|$ ).

**Claim B.3.** Let  $\pi$  be a uniformly random permutation of  $E_H$ . Let  $X_{v^*} : v^* \in V^*$  represent the indicator variable of  $v^*$  being occupied and  $X_O = \sum X_{v^*}$ . Then  $X \geq n \cdot (1 - \delta)$  w.h.p.

PROOF. Note that each  $v^* \in V^*$  has at most  $n$  edges with vertices of  $V'$  and has at least  $10n/\delta$  edges with vertices in  $U_{v^*}$  and  $W_{v^*}$ . Therefore,

$$\mathbb{E}[X_{v^*}] = \Pr(X_{v^*} = 1) \geq \frac{n \cdot 10/\delta}{n \cdot (10/\delta + 1)} = 1 - \frac{1}{10/\delta + 1} \geq 1 - \delta/10.$$

On the other hand, the variables  $\{X_{v^*} \mid v \in V\}$  are independent binary variables. Therefore, by Chernoff's bound, we have that

$$\begin{aligned} \Pr(X_O \leq n \cdot (1 - \delta)) &\leq \Pr\left(X_O \leq n \cdot (1 - \delta/10) - \frac{n \cdot \delta}{2}\right) \\ &\leq \Pr\left(X_O \leq \mathbb{E}[X_O] - \mathbb{E}[X_O] \cdot \frac{\delta}{2}\right) \end{aligned} \tag{8}$$

$$\leq 2 \cdot \exp\left(-\frac{(\delta/2)^2 \cdot \mathbb{E}[X_O]}{3}\right) \tag{9}$$

$$\leq n^{-\Theta(1)} \quad (10)$$

Inequality 8 follows from the fact that  $n \geq E[X_O] \geq n \cdot (1 - \delta)$ . Inequality 9 is an application of Chernoff's bound. Inequality 10 follows as long as  $n \cdot \delta^2 \in \Omega(\log(n))$ .  $\square$

The above claim completes the proof of the last requirement of Lemma B.2.  $\square$

## B.2 The algorithm

We now introduce the algorithm that will build on the previous two lemmas and inform the proof of Lemma 4.6, given in Algorithm 4. Recall that we wish to estimate the number of edges in some input matching, which here, to avoid confusion, we denote by  $M^*$ , that are both matched in some maximal matching in  $G$ .

Let  $H = (V_H, E_H)$  be the supergraph of  $G$  defined by Lemma B.2 of  $G$  with  $\delta = \epsilon^2/8$ . For a permutation  $\pi$  of  $E_H$ , define  $M'(\pi) := \text{GMM}(H, \pi) \cap (V \times V)$  to be the set of edges in  $\text{GMM}(H, \pi)$  both of whose endpoints are in  $V$ . Let  $M(\pi)$  be a maximal matching in  $G$  that is obtained by augmenting  $M'(\pi)$ , i.e., we start with  $M = M'(\pi)$ , visit the edges  $e \in E$  in an arbitrarily fixed order, and obtain the matching  $M(\pi)$  by greedily adding as many edges to  $M$  as possible. Note that  $M'(\pi) \subseteq M(\pi) \subseteq E'$ . Note also that  $M(\pi)$  will be the maximal matching that Lemma 4.6 refers to as  $M'$ . We now slightly overload our notations and let  $k_{M'(\pi)}$  denote the number of edges in  $M^*$  both of whose endpoints are matched in  $M'(\pi)$ .<sup>7</sup>

---

### Algorithm 4 Extended Sub-Linear Algorithm

---

```

1: if  $|M^*| \leq \epsilon^2 \cdot n$  then
2:   Return  $\kappa = 0$ 
3: (Implicitly) construct  $H = (V_H, E_H)$  as in Lemma B.2 with  $\delta = \epsilon^2/8$ 
4: Sample a permutation  $\pi$  of  $E_H$  uniformly at random
5:  $L \leftarrow \frac{10^5 \cdot \log(n)}{\epsilon^5}$ 
6: Sample  $L$  edges  $e_1, \dots, e_L \in M^*$  uniformly at random with replacement
7: Let  $X_i$  be one if both endpoints of edge  $e_i$  are matched by  $\text{GMM}(H, \pi)$  and  $X = \sum_i X_i$ 
8: Return  $\kappa = \frac{X \cdot |M|}{L} - \frac{n \cdot \epsilon^2}{2}$ 

```

---

**Claim B.4.** *Algorithm 4 can be implemented in time  $\tilde{O}_\epsilon(n)$  in expectation.*

**PROOF.** The construction of  $H$  is implicit, and as such takes no time. Let  $T_H(v, \pi)$  stand for the time it takes to calculate the matched status of vertex  $v \in V_H$  in  $\text{GMM}(H, \pi)$  using the algorithm of [16]. By Lemma B.1 we have that  $\mathbb{E}_{v \sim V_H} [T_H(v, \pi)] = \tilde{O}_\epsilon(|E_H|/|V_H|) = \tilde{O}_\epsilon(1)$ . Therefore, since the endpoints of the sampled edges  $S = \bigcup_{i=1}^L e_i \subseteq V_H$  are a subset of vertices of cardinality  $|S| \geq \epsilon^2 \cdot n$ , and since  $|V_H| = \Theta_\epsilon(n^2)$  we have the expected time to calculate their matched status (using adjacency matrix queries, using the construction of  $H$ ) is

$$\mathbb{E}_{v \sim S} [T_H(v, \pi)] \leq \mathbb{E}_{v \sim V_H} [T_H(v, \pi)] \cdot \frac{|V_H|}{|S|} \leq \tilde{O}_\epsilon(n)$$

$\square$

We now argue that Algorithm 4 provides a good approximation of the number of nodes in  $M^*$  both of whose endpoints are matched by  $\text{GMM}(H, \pi)$ . But first, we recall the basic Chernoff bounds that we will rely on here.

<sup>7</sup>Recall that in the statement of Lemma 4.6 we defined the notation  $k_{M'}$  only if  $M'$  is a matching in  $G'$ , which is not the case with  $M'(\pi)$ . Nevertheless, for ease of exposition, we use the notation  $k_{M'(\pi)}$ .

1373 **Lemma B.5.** *Chernoff bound: Let  $X$  be the sum of independently distributed (or negatively associated) 1374 random variables  $X_1, \dots, X_m$  with  $X_i \in [0, 1]$  for each  $i \in [m]$ . Then for all  $\delta \in (0, 1)$ :*

$$1376 \quad \Pr(|X - E[X]| \geq \delta \cdot E[X]) \leq 2 \cdot \exp\left(-\frac{\delta^2 \cdot E[X]}{3}\right). \\ 1377$$

1378 **Lemma B.6.** *W.h.p., The output  $\kappa$  of Algorithm 4 satisfies*

$$1380 \quad \kappa_{M(\pi)} \geq \kappa \geq \kappa_{M(\pi)} - n \cdot \epsilon^2. \\ 1381$$

1382 PROOF. First, by Lemma B.2, we have that w.h.p., the set of nodes  $V_{slack} \subseteq V$  that are matched to 1383 nodes in  $V_H \setminus V$  have cardinality at most  $|V_{slack}| \leq n\epsilon^2/8$ . Moreover,  $GMM(H, \pi) \cap E$  is a maximal 1384 matching in  $G[V \setminus V_{slack}]$ .

1385 Observe that whenever  $\kappa = 0$  is returned by the algorithm due to  $|M^*|$  being small, the algorithm 1386 returns a trivially correct solution. As  $M'(\pi)$  is an  $\epsilon^2/8$ -AMM, we conclude that:

$$1387 \quad |M'(\pi)| \leq |M(\pi)| \leq |M'(\pi)| + \frac{n \cdot \epsilon^2}{8}. \quad (11) \\ 1388 \\ 1389$$

1390 Define  $M_H^*$  to be the set of edges of  $M^*$  such that both of their endpoints are matched by 1391  $GMM(H, \pi)$ . By the guarantees of the construction of  $H$  we know that there can be at most  $n \cdot \epsilon^2/8$  1392 vertices of  $V$  matched by an edge not in  $M'(\pi)$ . Therefore,

$$1393 \quad |M_H^*| \geq \kappa_{M'(\pi)} \geq |M_H^*| - \frac{n \cdot \epsilon^2}{8}. \quad (12) \\ 1394 \\ 1395$$

1396 Note that using the Algorithm 4 is sampling from edges of  $M^*$  and determining if they are in 1397  $M_H^*$  (hence approximating  $\kappa_{M_H^*}$ ). Specifically, by inequalities (11) and (12), we get the following.

$$1398 \quad \kappa_{M_H^*} \in \left[ \kappa_{M'(\pi)} \pm \frac{n \cdot \epsilon^2}{8} \right] \subseteq \left[ \kappa_{M(\pi)} \pm \frac{n \cdot \epsilon^2}{8} \pm |M(\pi)| - |M'(\pi)| \right] \subseteq \left[ \kappa_{M(\pi)} \pm \frac{n \cdot \epsilon^2}{4} \right]. \\ 1399 \\ 1400$$

1401 We will argue that with high probability  $\frac{X \cdot |M|}{L} \in \left[ \kappa_{M_H^*} \pm \frac{n \cdot \epsilon^2}{8} \right]$ , dependent on the randomization 1402 of  $M_L^*$ . Observe that  $X_i$  are independently distributed random variables taking values in  $[0, 1]$  and 1403  $X$  is a binomial variable with parameters  $k, |\kappa_{M_H^*}|/|M|$ . We will consider two cases:

1404 **Case (A):**  $\kappa_{M_H^*} \leq \frac{n \cdot \epsilon^3}{8}$ . In this case, we derive that

$$1405 \quad \Pr\left(\frac{X \cdot |M|}{L} \notin \left[ \kappa_{M_H^*} \pm \frac{n \cdot \epsilon^2}{8} \right]\right) = \Pr\left(\frac{X \cdot |M|}{L} \geq \kappa_{M_H^*} + \frac{n \cdot \epsilon^2}{8}\right) \\ 1406 \\ 1407 \quad \leq \Pr\left(\frac{X \cdot |M|}{L} \geq \frac{n \cdot \epsilon^2}{8}\right) \\ 1408 \\ 1409 \quad = \Pr\left(B(L, \kappa_{M_H^*}/|M|) \geq \frac{n \cdot \epsilon^2}{8}\right) \quad (13) \\ 1410 \\ 1411$$

$$1412 \quad \leq \Pr\left(B(L, \epsilon) \geq \frac{n \cdot \epsilon^2}{8}\right) \quad (14) \\ 1413 \\ 1414$$

$$1415 \quad \leq \Pr(B(L, \epsilon) \geq 2 \cdot \mathbb{E}[B(L, \epsilon)]) \\ 1416 \\ 1417$$

$$1418 \quad \leq 2 \cdot \exp\left(-\frac{\mathbb{E}[B(L, \epsilon)]}{3}\right) \quad (15) \\ 1419 \\ 1420$$

$$1421 \quad \leq n^{-\Theta(1)}.$$

In the above derivation, (13) holds as  $\kappa_{M_H^*}/|M| \leq \epsilon$  (as otherwise  $\kappa = 0$  would have been returned by the algorithm), and (14) is true assuming  $n \cdot \epsilon^2/8 \geq 2 \cdot L \cdot \epsilon = \tilde{O}(1)$ . Finally, (15) follows from Chernoff bound (Lemma B.5) on a binomial random variable.

**Case (B):**  $\kappa_{M_H^*} \geq n \cdot \frac{\epsilon^3}{8}$ . In this case, we derive that

$$\begin{aligned} \Pr\left(\frac{X \cdot |M|}{L} \notin \left[\kappa_{M_H^*} \pm \frac{n \cdot \epsilon^2}{8}\right]\right) &= \Pr\left(|X - E[X]| \geq \frac{n \cdot \epsilon^2}{8} \cdot \frac{L}{|M|}\right) \\ &= \Pr\left(|X - \mathbb{E}[X]| \geq \mathbb{E}[X] \cdot \frac{n \cdot \epsilon^2}{8 \cdot \kappa_{M_H^*}}\right) \\ &\leq \Pr\left(|X - \mathbb{E}[X]| \geq \mathbb{E}[X] \cdot \frac{n \cdot \epsilon^2}{8 \cdot n}\right) \\ &\leq 2 \cdot \exp\left(-\frac{(\epsilon^2/8)^2 \cdot \mathbb{E}[X]}{3}\right) \end{aligned} \quad (16)$$

$$\begin{aligned} &= \exp\left(-\frac{L \cdot \kappa_{M_H^*} \cdot \epsilon^4}{|M| \cdot 194}\right) \\ &\leq n^{-\Theta(1)}. \end{aligned} \quad (17)$$

In the above derivation, (16) follows from Chernoff bound (Lemma B.5), and (17) holds due to our assumptions on  $|M|$  and  $\kappa_{M_H^*}$ . Therefore, with high probability  $X \cdot |M|/k \in [\kappa_{M_H^*} \pm \epsilon^2 \cdot n/8] \in [\kappa_{M(\pi)} \pm \epsilon^2 \cdot n \cdot (1/8 + 1/4)]$  (recall that  $\kappa_{M(\pi)} = \kappa_M$ ). This implies that  $\kappa \leq \kappa_M + \epsilon^2 \cdot n \cdot (1/8 + 1/4 - 1/2) \leq \kappa_M$  and  $\kappa \geq \kappa_M - \epsilon^2 \cdot n \cdot (1/8 + 1/4 - 1/2) \geq \kappa_M - \epsilon^2 \cdot n$ .  $\square$

Having concluded that Algorithm 4 can be implemented in low expected time, and is correct w.h.p., we are now ready to prove Lemma 4.6, restated below for ease of reference. (Note that here  $G = (V, E)$  are renamed  $G' = (V', E')$ , and  $n'$  and  $k_{M'}$  correspond respectively to  $n$  and  $\kappa_{M(\pi)}$ , whereas  $M^*$  in Algorithm 4 is renamed  $M$ .)

**Lemma 4.6.** *Consider a graph  $G' = (V', E')$  with  $|V'| = n'$ , and a matching  $M$  with  $V(M) \subseteq V'$  that is not necessarily part of  $G'$  (i.e., we might have  $M \not\subseteq E'$ ). For any matching  $M'$  in  $G'$ , let  $k_{M'}$  denote the number of edges in  $M$  both of whose endpoints are matched in  $M'$ . There is an algorithm which, given adjacency matrix query access to the edges of  $G'$ , w.h.p. runs in  $\tilde{O}_\epsilon(n')$  time and returns an estimate  $\kappa \in [k_{M'} - \epsilon^2 n', k_{M'}]$  for some maximal matching  $M'$  in  $G'$ .*

**PROOF.** By Claim B.4, Algorithm 4 runs in expected  $\tilde{O}_\epsilon(n)$  time and returns a correct solution with high probability. To improve its running time guarantee to a high probability bound we only need to incur a blowup of  $O(\log(n))$  in running time: run the algorithm  $O(\log(n))$  time in parallel and output the solution given by the first terminating copy. One of these algorithms will terminate within at most twice the expected time, by Markov's inequality, and so w.h.p., one of these completes after  $\tilde{O}_\epsilon(n)$  time. Finally, by union bound and Lemma B.6, all of the  $\log n$  algorithms' output satisfies the desired bounds with probability  $1 - 1/\text{poly}(n)$ , and so w.h.p., we obtain a solution satisfying the desired bounds after  $\tilde{O}_\epsilon(n)$  time.  $\square$

## C OMITTED PROOFS FROM SECTION 4

Here we prove the tighter bound on the number of  $V(M_1)$ -disjoint 3-augmenting paths in the subgraph  $M_1 \cup M_2$  as output by Algorithm 2, restated below.

1471 **Lemma 4.1.** *If  $|M_1| = (\frac{1}{2} + c) \cdot \mu(G)$ , then  $G[M_1 \cup M_2]$  contains a set  $\mathcal{P}$  of 3-augmenting paths*  
 1472 *w.r.t.  $M_1$  that are disjoint in their  $V(M_1)$  nodes, with expected cardinality at least*

$$1473 \quad 1474 \quad 1475 \quad \mathbb{E}[|\mathcal{P}|] \geq \left( \frac{b}{b+1} \right) \cdot \left( \frac{1}{4} \cdot \left( \frac{1}{2} - 3c \right) - \frac{1}{b} \cdot \left( \frac{1}{2} + c \right) - \frac{7\epsilon}{8} \right) \cdot \mu(G).$$

1476 PROOF. Fix a maximum set of disjoint length-three augmenting paths w.r.t.  $M_1$  in  $G$ , denoted  
 1477 by  $\mathcal{P}^*$ . By Proposition 2.6, we have  $|\mathcal{P}^*| \geq (\frac{1}{2} - 3c - \frac{7\epsilon}{2}) \cdot \mu(G)$ . Next, let  $S \subseteq \mathcal{P}^*$  be the paths  
 1478  $u' - u - v - v'$  that “survive” the bipartition, in the sense that  $(u, u'), (v, v') \in E_2$ . By construction,  
 1479 each path in  $\mathcal{P}^*$  survives with probability exactly  $\frac{1}{4}$ . Therefore,  $\mathbb{E}[|S|] \geq \frac{1}{4} \cdot (\frac{1}{2} - 3c - \frac{7\epsilon}{2}) \cdot \mu(G)$ .  
 1480 Let  $D := \mathcal{P}^* \setminus S$  be the set of paths that did not survive this bipartition.

1481 For  $i \in \{0, 1, 2\}$ , let  $S_i \subseteq S$  and  $D_i \subseteq D$  be the sets of paths  $u' - u - v - v'$  in  $S$  and  $D$  (respectively)  
 1482 with  $i$  of their  $V(M_1)$  nodes  $u$  and  $v$  matched in  $M_2$ . Now, by our bipartition, if  $u' - u - v - v' \in S_2 \cup D_2$ ,  
 1483 i.e., if  $u$  and  $v$  are both matched in  $M_2$ , then they are matched to distinct nodes. Therefore,  $G[M_1 \cup M_2]$   
 1484 contains a set of augmenting paths  $\mathcal{P}$  w.r.t.  $M_1$  that are disjoint in their  $V(M_1)$  nodes, of cardinality  
 1485  $|\mathcal{P}| = |S_2| + |D_2|$ . We now turn to lower bounding  $|S_2| + |D_2|$ .

1486 To bound  $|S_2| + |D_2|$ , we will double count the edges of  $M_2$ , once from their  $V(M_1)$  endpoints,  
 1487 and once from their  $\overline{V(M_1)}$  endpoints. First, by definition, since each edge in  $M_2$  has exactly  
 1488 one endpoint in  $V(M_1)$  and each node in  $V(M_1)$  is matched at most once in  $M_2$ , we have that  
 1489  $|M_2| = 2|S_2| + |S_1| + 2|D_2| + |D_1| \leq |S_2| + |D_2| + |M_1|$ , where the inequality follows from  $|M_1| \geq$   
 1490  $\sum_{i=0}^2 (|S_i| + |D_i|)$ , by definition. On the other hand, for each of the  $|S| - |S_2| = |S_0| + |S_1|$  survived  
 1491 paths  $u' - u - v - v' \in S_0 \cup S_1$  that does not have both its internal nodes matched in  $M_2$ , we  
 1492 have by maximality of  $M_2$  that  $u'$  and/or  $v'$  must contribute  $b$  distinct edges to  $M_2$ . Therefore,  
 1493  $|M_2| \geq b \cdot (|S_0| + |S_1|)$ . Combining the above, we obtain

$$1494 \quad 1495 \quad b \cdot (|S| - |S_2|) \leq |M_2| \leq |S_2| + |D_2| + |M_1|,$$

1496 which after rearranging, yields

$$1497 \quad 1498 \quad b \cdot |S| - |M_1| \leq (b+1) \cdot |S_2| + |D_2| \leq (b+1) \cdot (|S_2| + |D_2|).$$

1499 Simplifying and combining with the lower bound on  $\mathbb{E}[|S|]$ , we obtain the claimed bound, as  
 1500 follows.

$$1501 \quad 1502 \quad \mathbb{E}[|\mathcal{P}|] = \mathbb{E}[|S_2| + |D_2|] \geq \frac{b}{b+1} \cdot \left( \mathbb{E}[|S|] - \frac{1}{b} \cdot |M_1| \right) \\ 1503 \\ 1504 \quad \geq \frac{b}{b+1} \cdot \left( \frac{1}{4} \cdot \left( \frac{1}{2} - 3c - \frac{7\epsilon}{2} \right) - \frac{1}{b} \cdot \left( \frac{1}{2} + c \right) \right) \cdot \mu(G). \\ 1505 \\ 1506 \quad = \frac{b}{b+1} \cdot \left( \frac{1}{4} \cdot \left( \frac{1}{2} - 3c \right) - \frac{1}{b} \cdot \left( \frac{1}{2} + c \right) - \frac{7\epsilon}{8} \right) \cdot \mu(G). \\ 1507 \\ 1508$$

□

## 1510 D OMITTED PROOFS OF SECTION 5

1511 We stress that the following is essentially implied by the work of [82], from which we now repeat  
 1512 significant amount of text essentially verbatim. The only difference here will be our final proof of  
 1513 Lemma 5.7, allowing us to efficiently periodically compute an  $\epsilon$ -AMM, and the use of this lemma in  
 1514 the subsequent section. Readers familiar with [82] are encouraged to read ahead to that lemma.

1515 **Overview.** Briefly, [82] identified an edge-coloring-based approach to compute, based on the  
 1516 efficient maintenance of edge colorings and a particular fractional matching of [25], a kernel. (See  
 1517 Algorithm 5.) We start by recalling the type of fractional matching needed here, due to [4].

1519

1520 **Definition D.1.** For  $c \geq 1$  and  $d \geq 1$ , a fractional matching  $\vec{x}$  is  $(c, d)$ -approximately-maximal  
 1521  $(c, d)$ -AMfM if every edge  $e \in E$  either has fractional value  $x_e > 1/d$  or it has one endpoint  $v$  with  
 1522  $\sum_{e \ni v} x_e \geq 1/c$  with all edges  $e'$  incident on this  $v$  having value  $x_{e'} \leq 1/d$ .

1523 As proven in [4, Appendix A], the dynamic fractional matching of [27] is precisely such an  
 1524 approximately-maximal matching.  
 1525

1526 **Lemma D.2.** For all  $\epsilon \leq \frac{1}{2}$ , there is a deterministic dynamic  $(1 + 2\epsilon, \max\{54 \log n/\epsilon^3, (3/\epsilon)^{21}\})$ -  
 1527 AMfM algorithm with  $t_u = O(\log^3 n/\epsilon^7)$  worst-case update time, changing at most  $O(\log n/\epsilon^2)$  edges'  
 1528 fractions per update in the worst case.

1529 Now, we turn to the sparsification procedure of [82], given in Algorithm 5. Briefly, this algorithm  
 1530 decomposes the graph into a logarithmic number of subgraphs, based on grouped  $x$ -values, edge  
 1531 colors these subgraphs using at most  $\gamma = 2$  times their maximum degree, and then outputs the  
 1532 union of these subgraphs.  
 1533

---

1534 **Algorithm 5** Edge-Color and Sparsify [82]

---

1536 1: **for**  $i \in \{1, 2, \dots, \lceil 2 \log_{1+\epsilon}(n/\epsilon) \rceil\}$  **do**  
 1537 2:   let  $E_i \triangleq \{e \mid x_e \in ((1 + \epsilon)^{-i}, (1 + \epsilon)^{-i+1}]\}$ .  
 1538 3:   compute a  $2\lceil(1 + \epsilon)^i\rceil$ -edge-coloring  $\chi_i$  of  $G_i \triangleq G[E_i]$ .                   ▷ Note:  $\Delta(G_i) < (1 + \epsilon)^i$   
 1539 4:   Let  $S_i$  be a sample of  $\min\{2\lceil d(1 + \epsilon) \rceil, 2\lceil(1 + \epsilon)^i\rceil\}$  colors without replacement in  $\chi_i$ .  
 1540 5: **Return**  $K \triangleq (V, \bigcup_i \bigcup_{M \in S_i} M)$ .

---

1542 The following lemma of [82] allows us to compute kernels from AMfMs using Algorithm 5.

1544 **Lemma D.3.** Let  $c \geq 1$ ,  $\epsilon > 0$  and  $d \geq \frac{9c(1+\epsilon)^2 \cdot \log n}{\epsilon^2}$ . If  $\vec{x}$  is a  $(c, d)$ -AMfM, then the subgraph  $K$   
 1545 output by Algorithm 5 when run on  $\vec{x}$  with  $\epsilon$  and  $d$  is a  $(c(1 + O(\epsilon), d(1 + O(\epsilon), 0)$ -kernel, w.h.p.  
 1546

1547 We are now ready to prove our (periodic) algorithmic kernel and AMM algorithm's guarantees,  
 1548 restated below for ease of reference.

1549 **Lemma 5.7.** Let  $\epsilon \in (0, 1)$  and  $d = \tilde{O}_\epsilon(1)$  be sufficiently large. Then, there exists a robust algorithm  
 1550 with worst-case update time  $t_u = \tilde{O}_\epsilon(1)$  allowing for  $(\epsilon, d)$ -kernel and  $\epsilon$ -AMM queries in worst-case  
 1551 query time  $t_q = \tilde{O}_\epsilon(d \cdot \mu(G))$ . The query's outputs are a kernel and an  $\epsilon$ -AMM w.h.p.  
 1552

1553 **PROOF.** We maintain the dynamic  $(1 + 2\epsilon, \tilde{O}_\epsilon(1))$ -AMfM of Lemma D.2, using  $\tilde{O}_\epsilon(1)$  deterministic  
 1554 w.c. update time and number of changes to edges per update. In addition, we maintain the subgraphs  
 1555  $G_i$  in Algorithm 5. In each such subgraph we maintain  $2\lceil(1 + \epsilon)^i\rceil$ -color edge colorings in each  
 1556  $G_i$  in  $O(\log n)$  deterministic w.c. time per change to  $\vec{x}$ , using the logarithmic-time  $(2\Delta - 1)$ -edge  
 1557 coloring algorithm of [24]. This concludes the description of the updates, which by the above take  
 1558 deterministic w.c. update time  $t_u = \tilde{O}_\epsilon(1)$ .

1559 Next, to compute a kernel, we run the sampling step of Algorithm 5. As this is bottlenecked by the  
 1560 time to write down the  $O(\log^2 n)$  colors (matchings), each of size no greater than  $\mu(G)$  (by definition),  
 1561 this query takes deterministic  $\tilde{O}(\mu(G))$ . Finally, this output graph  $K$  is an  $(O(\epsilon), d(1 + O(\epsilon)))$ -kernel  
 1562 w.h.p., by Lemma D.3. Finally, to output an  $\epsilon$ -AMM, we appeal to the static algorithm Lemma 5.6,  
 1563 which runs in deterministic time  $\tilde{O}_\epsilon(d \cdot \mu(G)) = \tilde{O}(\mu(G))$  and outputs an  $\epsilon$ -AMM, provided  $K$  is a  
 1564 kernel, i.e., it also succeeds w.h.p. □