

# Maintaining an EDCS in General Graphs: Simpler, Density-Sensitive and with Worst-Case Time Bounds

Fabrizio Grandoni\*, Chris Schwiegelshohn†, Shay Solomon‡, Amitai Uzdad‡

## Abstract

In their breakthrough ICALP'15 paper, Bernstein and Stein presented an algorithm for maintaining a  $(3/2 + \varepsilon)$ -approximate maximum matching in fully dynamic *bipartite* graphs with a *worst-case* update time of  $O_\varepsilon(m^{1/4})$ ; we use the  $O_\varepsilon$  notation to suppress the  $\varepsilon$ -dependence. Their main technical contribution was in presenting a new type of bounded-degree subgraph, which they named an *edge degree constrained subgraph (EDCS)*, which contains a large matching — of size that is smaller than the maximum matching size of the entire graph by at most a factor of  $3/2 + \varepsilon$ . They demonstrate that the EDCS can be maintained with a worst-case update time of  $O_\varepsilon(m^{1/4})$ , and their main result follows as a direct corollary. In their followup SODA'16 paper, Bernstein and Stein generalized their result for general graphs, achieving the same update time of  $O_\varepsilon(m^{1/4})$ , albeit with an amortized rather than worst-case bound. To date, the best *deterministic* worst-case update time bound for *any* better-than-2 approximate matching is  $O(\sqrt{m})$  [Neiman and Solomon, STOC'13], [Gupta and Peng, FOCS'13]; allowing randomization (against an oblivious adversary) one can achieve a much better (still polynomial) update time for an approximation slightly below 2 [Behnezhad, Lacki and Mirrokni, SODA'20].

In this work we<sup>1</sup> simplify the approach of Bernstein and Stein for bipartite graphs, which allows us to generalize it for general graphs while maintaining the same bound of  $O_\varepsilon(m^{1/4})$  on the *worst-case* update time. Moreover, our approach is *density-sensitive*: If the *arboricity* of the dynamic graph is bounded by  $\alpha$  at all times, then the worst-case update time of the algorithm is  $O_\varepsilon(\sqrt{\alpha})$ .

**Recent related work:** Independently and concurrently to our work, Roghani, Saberi and Wajc [arXiv'21] obtained two dynamic algorithms for approximate maximum matching with worst-case update time bounds. Their first algorithm achieves approximation factor slightly better than 2 within  $O(\sqrt{n} \cdot m^{1/8})$  update time, and their second algorithm achieves approximation factor  $(2 + \varepsilon)$  for any  $\varepsilon > 0$  within  $O_\varepsilon(\sqrt{n})$  update time. In terms of techniques, the two works are entirely disjoint.

## 1 Introduction

Dynamic matching algorithms have been subject to extensive research attention for more than a decade, starting with the pioneering work of Onak and Rubinfeld [OR10]. We do not aim to cover here the entire literature, but rather to briefly survey most of the state-of-the-art results that concern approximation factor 2 or less — indeed the 2-approximation barrier is a central one in the context of graph matching. One may try to optimize the *amortized* (i.e., average) update time of an algorithm or its *worst-case* (i.e., maximum) update time, over a worst-case sequence of graphs, and we will put special emphasis on this distinction. Indeed, there is a strong separation between the state-of-the-art amortized versus worst-case time bounds for dynamic matching algorithms; a similar separation exists for various other dynamic graph problems, such as spanning tree, minimum spanning tree and two-edge connectivity.

Before starting the background survey, we summarize the key question around which our work revolves:

QUESTION 1. *Is there a deterministic algorithm for maintaining a better-than-2 (approximate) maximum cardinality matching (MCM) with a low worst-case update time? Further, is it possible to push the approximation factor well below 2 (possibly using a randomized algorithm)?*

---

\*IDSIA, USI-SUPSI

†Aarhus University

‡Tel Aviv University

<sup>1</sup>*quasi nanos, gigantium humeris insidentes*

**Maximal matching.** In their seminal paper, Baswana et al. [BGS11] showed that a maximal matching, which provides a 2-MCM, can be maintained with an *amortized* update time of  $O(\log n)$ . A constant amortized update time was given by Solomon [Sol16], and a *worst-case* polylog( $n$ ) update time was given by Bernstein et al. [BFH19]. All these algorithms are randomized, and as common in the area of dynamic graph algorithms, they operate under the *oblivious adversarial model*, in which the adversary (the entity adding/deleting edges to/from the graph) may know all the edges in the graph and their arrival order, as well as the algorithm to be used, but is not aware of the random bits used by the algorithm, and so cannot choose updates adaptively in response to the randomly guided choices of the algorithm.

Neiman and Solomon [NS13] gave a deterministic algorithm for maintaining a maximal matching with a worst-case update time of  $O(\sqrt{m})$ , where  $m$  is the (dynamically changing) number of edges in the graph. This  $\sqrt{m}$  bound remains the state-of-the-art for dynamic maximal matching, even allowing amortization and even allowing randomization under a non-oblivious adversary.

**Crossing the 2-approximation barrier.** Neiman and Solomon [NS13] were the first to cross the 2-approximation barrier: 3/2-MCM with a worst-case deterministic update time of  $O(\sqrt{m})$ . Soon afterwards Gupta and Peng [GP13] showed that a  $(1 + \varepsilon)$ -MCM can be maintained with a worst-case deterministic update time of  $O(\sqrt{m}/\varepsilon^2)$ , for any  $0 < \varepsilon < 1/2$ ; moreover, if the maximum degree is always upper bounded by  $\Delta$ , then a worst-case update time of  $O(\Delta/\varepsilon^2)$  can be achieved [GP13].

Bernstein and Stein [BS15] showed that a  $(3/2 + \varepsilon)$ -MCM can be maintained in *bipartite* graphs with a *worst-case* deterministic update time of  $O(m^{1/4}\varepsilon^{-2.5})$ . Their main technical contribution was in presenting a new type of bounded-degree subgraph, an *edge degree constrained subgraph (EDCS)*, which contains a large matching — of size that is smaller than the maximum matching size of the entire graph by at most a factor of  $3/2 + \varepsilon$ . They demonstrate that the EDCS can be maintained with a worst-case update time of  $O(m^{1/4}\varepsilon^{-2.5})$  (and also with a worst-case *recourse*, which is the number of edge changes per update step, of  $O(1/\varepsilon)$ ), and their main result follows by running the bounded degree version of Gupta-Peng algorithm [GP13] on top of the EDCS. In their followup paper [BS16], they generalized their result for general graphs, using an inherently different algorithm for the EDCS maintenance; they achieved the same update time of  $O_\varepsilon(m^{1/4})$ , albeit with an amortized rather than worst-case bound.

Behnezhad et al. [BLM20] showed that a (slightly better-than-2)-MCM can be maintained with arbitrarily small polynomial worst-case update time, that is, the update time of their algorithm is  $O(\Delta^\varepsilon + \text{polylog}(n))$  and their approximation is  $2 - \Omega_\varepsilon(1)$ , where  $\Omega_\varepsilon(1)$  is a tiny constant that shrinks rapidly as  $\varepsilon$  reduces. This result is achieved via a randomized algorithm that assumes an oblivious adversary; for bipartite graphs, similar (but somewhat weaker) results can be achieved without making the oblivious adversary assumption: Via a randomized algorithm against an adaptive adversary [Waj20, BHN16] and via a deterministic algorithm but with an amortized update time bound [BK21, BHN16].

**Bounded arboricity graphs.** The *arboricity* of an  $m$ -edge graph is the minimum number of forests into which it can be decomposed, and it ranges from 1 to  $\sqrt{m}$ . The family of bounded arboricity graphs can be viewed as the family of “sparse everywhere” graphs, containing bounded-degree graphs, all minor-closed graph classes (e.g., planar graphs and graphs of bounded treewidth), and randomly generated preferential attachment graphs. Moreover, many natural and real world graphs, such as the world wide web graph, social networks and transaction networks, are believed to have bounded arboricity. Thus, it is only natural to try and come up with better dynamic matching algorithms for bounded arboricity graphs.

Bernstein and Stein [BS15] showed that by using a *weighted* EDCS, one can maintain a  $(1 + \varepsilon)$ -MCM in *bipartite* graphs of arboricity at most  $\alpha$  with a *worst-case* update time of  $O(\alpha(\alpha + \log n) + \varepsilon^{-4}(\alpha + \log n) + \varepsilon^{-6})$ . In their followup paper [BS16], Bernstein and Stein generalized their result for general graphs using an ordinary EDCS, achieving a similar update time of  $O(\alpha(\alpha + \log n + \varepsilon^{-2}) + \varepsilon^{-6})$ , albeit with an amortized rather than worst-case bound and with approximation factor  $3/2 + \varepsilon$  rather than  $1 + \varepsilon$ . Peleg and Solomon [PS16] showed that a  $(1 + \varepsilon)$ -MCM can be maintained in graphs of arboricity at most  $\alpha$  with a *worst-case* update time of  $O(\alpha/\varepsilon^2)$ ; since  $\alpha$  ranges between 1 and  $\sqrt{m}$ , this result generalizes the result of [GP13] for general graphs with update time  $O(\sqrt{m}/\varepsilon^2)$ .

**1.1 Our contribution** In this work we simplify the approach of Bernstein and Stein [BS15] for bipartite graphs, which allows us to generalize it for general graphs while maintaining the same upper bound on the *worst-case* update time.

**THEOREM 1.1.** *For any dynamic graph  $G$  subject to edge updates and for any  $\varepsilon < 1/2$ , one can maintain a  $(3/2 + \varepsilon)$ -MCM for  $G$  with a deterministic worst-case update time of  $O(m^{1/4}\varepsilon^{-2.5} + \varepsilon^{-6})$ .*

Note that 1.1 resolves 1 in the affirmative. In particular, it provides the first *deterministic* algorithm achieving approximation better-than-2 with a worst-case update time that is strictly sublinear in the number of vertices for general graphs; it also provides the first (possibly randomized) algorithm achieving approximation well below 2 with such a worst-case update time bound.

Importantly, our approach is *density-sensitive*: If the *arboricity* of the dynamic graph is bounded by  $\alpha$  at all times, then the worst-case update time of the algorithm reduces to  $O(\sqrt{\alpha}\varepsilon^{-2.5} + \varepsilon^{-6})$ .

**THEOREM 1.2.** *Fix any parameter  $\alpha \geq 1$ . For any dynamic graph  $G$  subject to edge updates whose arboricity is upper bounded by  $\alpha$  at all times and for any  $\varepsilon < 1/2$ , one can maintain a  $(3/2 + \varepsilon)$ -MCM for  $G$  with a deterministic worst-case update time of  $O(\sqrt{\alpha}\varepsilon^{-2.5} + \varepsilon^{-6})$ .*

We note that the previous state-of-the-art arboricity-dependent update time, due to Peleg and Solomon [PS16], is  $O(\alpha/\varepsilon)$  — which is quadratically higher than the update time provided by 1.2, ignoring the  $\varepsilon$ -dependence. Although the approximation guarantee provided by the matching of [PS16] is  $1 + \varepsilon$  rather than  $3/2 + \varepsilon$ , no better arboricity-dependent update time bounds were known prior to this work, even for much larger approximation guarantee and even for amortized bounds.

**1.2 Technical Overview** Our argument for maintaining an EDCS efficiently consists of the following three steps.

Suppose that the maximum degree in the dynamic graph never exceeds  $\Delta$ . In the first step, we demonstrate that an EDCS can be maintained with a worst-case update time of  $O_\varepsilon(\Delta)$ ; we use the  $O_\varepsilon$  notation to suppress the  $\varepsilon$ -dependence. Our key observation for this step is that the argument used by [BS15] for the case of *bipartite graphs of bounded arboricity* can be employed — with minimal changes — for *general graphs*, to restore a valid EDCS following any edge update by computing and augmenting a short (possibly non-simple) alternating path in the EDCS. The changes that we introduce only simplify the argument, as we only trim parts from it: (1) While the argument of [BS15] proves that the alternating path is simple, which is true only in bipartite graphs, we make do without relying on simple paths, which is crucial for coping with general graphs. (2) We don't maintain a dynamic edge orientation.

In the second step, we reduce the update time from  $O_\varepsilon(\Delta)$  to  $O_\varepsilon(\sqrt{\Delta})$ . Here too we use an idea from [BS15], of scanning just a small subset of neighbors following a change of degree (in the EDCS), which then triggers a discrepancy between the true degrees and their estimations by their neighbors. The main difference to [BS15] is that we do not try to preserve an EDCS with the same parameters as those achieved in step 1, but rather allow them to degrade by a small (constant) factor. This has no effect whatsoever on any of the guarantees, yet it simplifies the algorithm and its analysis quite a bit.

Finally, we demonstrate that  $\Delta$  can be substituted with either  $\Theta(\sqrt{m}/\varepsilon)$  or  $\Theta(\alpha/\varepsilon)$ , where  $\alpha$  is a fixed upper bound on the arboricity of the dynamic graph. This follows easily from the following result:

**THEOREM 1.3.** *[Sol18] Let  $G$  be a graph with arboricity  $\alpha$ . Suppose that each vertex  $v$  in  $G$  “marks” (up to)  $\delta := c(\alpha/\varepsilon)$  arbitrary incident edges, for an appropriate constant  $c$ . The graph  $G'$  obtained as the union of all edges marked twice (by both endpoints) is a  $(1 + \varepsilon)$ -MCM sparsifier for  $G$ , i.e.,  $\mu(G) \leq (1 + \varepsilon)\mu(G')$ , where  $\mu(H)$  denotes the maximum matching size of any graph  $H$ .*

It is straightforward to dynamically maintain the graph  $G'$  defined by 1.3 within constant worst-case update time. The only subtlety arises in case we would like to substitute  $\Delta$  with  $\Theta(\sqrt{m}/\varepsilon)$ , where  $m$  dynamically changes over time; this issue can be resolved quite easily, as we show in 5. This sparsification step is very simple, whereas if instead we resort to dynamic edge orientations, as done by [BS15], life would become more complicated; the following was written in Section 5 of [BS15], and is mostly attributed to the interplay between dynamic edge orientations and the discrepancy between the true and estimated degrees in the EDCS: “The details, however, are quite involved, especially since we need a worst-case update time.”

From the above discussion, it is clear that each of the steps above simplify a respective ingredient from the argument of [BS15]. Another source of simplification stems from the fact that our argument is broken into three rather separate steps, which are then combined together in a rather natural way, which stands in contrast to the

argument of [BS15], in which all the ingredients are intertwined together. Interestingly, our improved bounds over [BS15, BS16] are achieved as a direct by-product of this simplification.

**Related work.** The EDCS was introduced in [BS15] for dynamic matching algorithms, but it turned out to be a very useful graph structure also outside the area of dynamic graph algorithms. In particular, EDCSs are especially useful in space and communication constrained settings, see e.g. [AB19, ABB<sup>+</sup>19, Ber20].

For any fixed  $\varepsilon > 0$ , a  $(1+\varepsilon)$ -MCM can be maintained in constant (respectively,  $\text{polylog}(n)$ ) amortized update time in incremental (resp., decremental) graphs [GLS<sup>+</sup>19] (resp., [BK21]). Importantly, these results only concern amortized bounds, and no better worst-case time bounds for incremental or decremental graphs are known than the aforementioned results.

The improvement to the update time in bounded arboricity graphs versus general graphs becomes less significant as the arboricity grows. Milenković and Solomon [MS20] showed that a  $(1+\varepsilon)$ -MCM can be maintained with a worst-case update time of  $O(\frac{\beta}{\varepsilon^3} \log \frac{1}{\varepsilon})$ , where  $\beta = \beta(G)$  is the *neighborhood independence number* of the graph  $G$ , i.e., the size of the largest independent set in the neighborhood of any vertex. Graphs with bounded neighborhood independence, already for constant  $\beta$ , constitute a wide family of *possibly dense graphs*, including line graphs, unit-disk graphs and graphs of bounded growth.

**Concurrent work.** Independently and concurrently to our work, Roghani, Saberi and Wajc [RSW21] obtained two dynamic algorithms for approximate maximum matching with worst-case update time bounds. Their first algorithm achieves approximation factor slightly better than 2 within  $O(\sqrt{n} \cdot m^{1/8})$  update time, and their second algorithm achieves approximation factor  $(2+\varepsilon)$  for any  $\varepsilon > 0$  within  $O_\varepsilon(\sqrt{n})$  update time. In terms of techniques, the two works are entirely disjoint.

**Organization.** In 2 we present the basic notation and preliminaries used throughout. 3 describes how to achieve an  $O_\varepsilon(\Delta)$  deterministic worst-case update time, where  $\Delta$  is an upper bound on the maximum degree in the dynamic graph. In 4 we improve this update time to  $O_\varepsilon(\sqrt{\Delta})$ . Finally, in 5, we combine the algorithm from 4 with a simple sparsification technique to reduce the update time to  $O_\varepsilon(m^{1/4})$ , or, in case of graphs with arboricity bounded by  $\alpha$ , to  $O_\varepsilon(\sqrt{\alpha})$ .

## 2 Preliminaries

Let  $G = (V, E)$  be an undirected, unweighted graph, where  $|V| = n$  and  $|E| = m$ . Let  $H = (V, E_H)$  be a subgraph of  $G$ , with  $E_H \subseteq E$ ; denote by  $d_H(v)$  the *degree* of vertex  $v$  in the subgraph  $H$ . For technical convenience, we shall use a weight function over the edges of  $G$ ,  $w : E \rightarrow \mathbb{N}$ , which assigns a nonnegative integer for each edge of the graph; the weights of edges in  $H$  are inherited from their weights in  $G$ . Specifically, the weight of any edge  $(u, v) \in E$ , denoted by  $w(u, v)$ , is given by  $d_H(u) + d_H(v)$ . Following [BS15, BS16], we say that  $H$  is an  $(\beta, \beta^-)$ -EDCS for  $G$ , for any pair  $\beta, \beta^-$  of real numbers such that  $\beta^- < \beta$ , if  $H$  is a spanning subgraph of  $G$  that satisfies the following two properties:

(P1) For each  $(u, v) \in H$ ,  $w(u, v) \leq \beta$ .

(P2) For each  $(u, v) \notin H$ ,  $w(u, v) \geq \beta^-$ .

For a vertex  $v$ , denote by  $N(v)$  its adjacency list. Denote by  $\mu(G)$  the maximum matching size of a graph  $G$ . We say that a subgraph  $G'$  of  $G$  is a  $(1+\varepsilon)$ -approximate matching sparsifier of  $G$  if  $\mu(G) \leq (1+\varepsilon)\mu(G')$ . For any sequence of edge updates that defines a dynamic graph  $G$  and for any sparsifier  $G'$  for  $G$ , we say that an algorithm maintaining  $G'$  has a worst-case *recourse* (or a worst-case *update ratio*) of  $r$  if, for each edge update in  $G$ , the number of edge changes made to  $G'$  is at most  $r$ .

Finally, we rely on the following key property that relates the EDCS to  $\mu(G)$ .

LEMMA 2.1. (LEMMA 2.5 OF [AB19], PROVED FIRST BY [BS16, BS15]) *Let  $G = (V, E)$  be any graph and  $\varepsilon < 1/2$  be a parameter. For any parameters  $\lambda \leq \frac{\varepsilon}{100}$ ,  $\beta \geq 32\lambda^{-3}$ , and  $\beta^- \geq (1-\lambda) \cdot \beta$ , in any subgraph  $H := \text{EDCS}(G, \beta, \beta^-)$ , it holds that  $\mu(G) \leq (\frac{3}{2} + \varepsilon) \mu(H)$ .*

## 3 Maintaining an EDCS, Part I: Update time $O_\varepsilon(\Delta)$

In this section we show that an  $(\beta, (1-\varepsilon)\beta)$ -EDCS can be maintained dynamically with a worst-case update time of  $O(1/\varepsilon + \Delta)$ . This section follows along similar lines as those in Appendix C.1 of [BS15].

To dynamically maintain an EDCS, it is instructive to define the following two types of edges:

(P1) A full edge  $(u, v)$  is in  $H$  and satisfies  $w(u, v) = \beta$ .

(P2) A deficient edge  $(u, v)$  is not in  $H$  and satisfies  $w(u, v) = (1 - \epsilon)\beta$ .

Upon insertion of edge  $(u, v)$  in  $G$ , if  $w(u, v) \geq (1 - \epsilon)\beta$ , then we do not add the edge to  $H$ , and Properties (P1) and (P2) continue to hold. On the other hand, if  $w(u, v) < (1 - \epsilon)\beta$ , we need to add  $(u, v)$  to  $H$ . Doing so will increase  $d_H(u)$  and  $d_H(v)$  by 1, which may lead to a violation of Property (P1) for other edges in  $H$  that are incident to either  $u$  or  $v$ . We will then fix an arbitrary such violating edge (if any) in a similar way, and thus we proceed to finding and augmenting an *alternating path* in  $H$  (edges along the path are in  $H$  and out of  $H$ , alternately), which ultimately allows us to add  $(u, v)$  to  $H$  and still maintain Properties (P1) and (P2) for all vertices. Next, we describe this process in detail.

We say that a vertex  $x$  is *increase-safe* (respectively, *decrease-safe*) if it has no incident full (resp., deficient) edges. Returning to adding  $(u, v)$  to  $H$ , let us focus on vertex  $v$ ; we will then treat vertex  $u$  analogously. If  $v$  is increase-safe, then, by definition, adding  $(u, v)$  to  $H$  does not violate Property (P1) for any edge incident on  $v$ . If  $v$  is not increase-safe, then it must have at least one incident full edge incident on it, say  $(v, p_1)$ . We would like to add  $(u, v)$  to  $H$  and remove  $(v, p_1)$  from  $H$ , thereby leaving  $v$ 's degree (in  $H$ ) unchanged. Doing so would decrease  $d_H(p_1)$ , which we can do only if  $p_1$  is decrease-safe. If  $p_1$  is decrease safe, then adding  $(u, v)$  to  $H$  and removing  $(v, p_1)$  leaves  $v$ 's degree unchanged, decreases  $p_1$ 's degree and reestablishes Properties (P1) and (P2) for all vertices (except possibly  $u$ ). However, if  $p_1$  is not decrease-safe, it must have an incident deficient edge, say  $(p_1, p_2)$ . We can add this edge to  $H$  and continue from  $p_2$ , just like we did from  $v$ . We can continue in this manner, stopping when we find either an increase-safe or a decrease-safe vertex. We will prove below that this process terminates, and when it does, all the degrees (in  $H$ ) have returned to their value prior to the edge insertion, except for at most two vertices, which we will need to handle separately. We remark that the alternating path that we augment throughout the process is not necessarily a simple path (a simple path is obtained only in bipartite graphs), but this fact has no consequences whatsoever on the validity of the argument. An alternating path from  $v$  in the case of edge insertion can be seen in Figure 1:

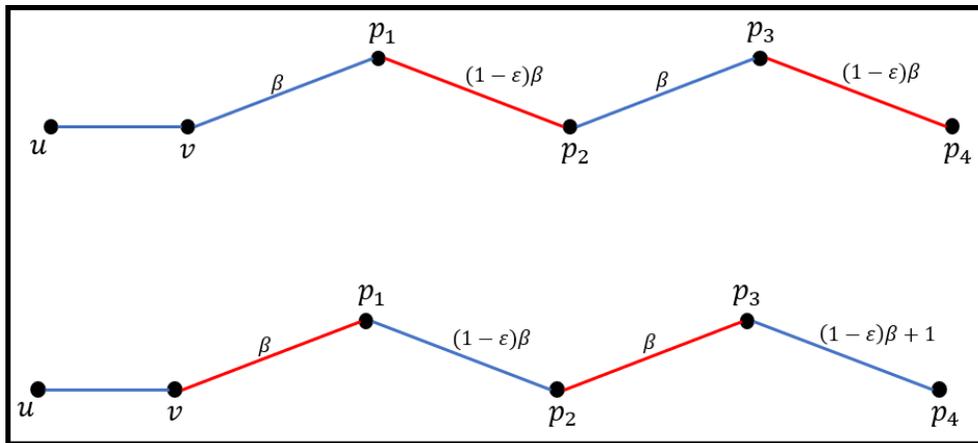


Figure 1:  $v$ 's alternating path upon insertion of edge  $(u, v)$  to  $H$ . Blue edges designate edges in the  $EDCS$ , and the red ones are not. On the top we have the path at the beginning of the alternating process, and on the bottom we have the path at the end. Notice that  $w(p_3, p_4)$  grows by one since  $d_H(p_4)$  grows by one.

Upon deletion of edge  $(u, v)$  from  $G$ , we can reestablish Properties (P1) and (P2) via a symmetric process. If  $(u, v)$  is not in  $H$ , then we do not need to change  $H$ . If  $(u, v)$  is in  $H$  and both  $u$  and  $v$  are decrease-safe, we just remove the edge  $(u, v)$ . Otherwise, we find an alternating path in the same way that we did for an edge insertion. For completeness, we provide the pseudo-codes for edge insertion and for edge deletion at the end of this section (where  $E_H$  is the set of edges that belong to  $H$ ). Along with the code for insert and delete, two auxiliary codes are added, which are used to find the next full/deficient edge on the path and update the relevant data structures.

Observe that any alternating path that is found and augmented following an edge update is a path that alternates between full and deficient edges. The following lemma and its proof are similar (but not identical) to

Lemmas 7 and 10 from [BS15] and their respective proofs.

LEMMA 3.1. *For any path  $P$  of alternating full and deficient edges, we have  $|P| \leq \frac{2}{\epsilon}$ .*

*Proof.* Consider first the case that  $P = (p_0, p_1, \dots, p_k)$  starts with a full edge. Let  $d = d_H(p_0)$ . Since  $(p_0, p_1)$  is full,  $d_H(p_1) = \beta - d$ . Since  $(p_1, p_2)$  is deficient,  $d_H(p_2) = (1 - \epsilon)\beta - d_H(p_1) = d - \epsilon\beta$ . Continuing, we get that  $d_H(p_3) = \beta - d + \epsilon\beta$ ,  $d_H(p_4) = d - 2\epsilon\beta$ , and in general we inductively get  $d_H(p_{2i}) = d - i\epsilon\beta$ , for any  $0 \leq i \leq k/2$ . Since each vertex in  $P$  has an incident full edge, all vertices have a positive degree, and thus  $d_H(p_{2i}) = d - i\epsilon\beta > 0$ . Since  $d \leq \beta$ , it follows that  $|P| \leq \frac{2}{\epsilon} - 1$ . If  $P$  starts with a deficient edge, let  $d = d_H(p_1)$ . Since  $(p_1, p_2)$  is full,  $d_H(p_2) = \beta - d$ . Since  $(p_2, p_3)$  is deficient,  $d_H(p_3) = (1 - \epsilon)\beta - d_H(p_2) = d - \epsilon\beta$ . Continuing, we get that  $d_H(p_4) = \beta - d + \epsilon\beta$ ,  $d_H(p_5) = d - 2\epsilon\beta$ , and in general we inductively get  $d_H(p_{2i+1}) = d - i\epsilon\beta$ , for any  $0 \leq i < k/2$ . Since all vertices have a positive degree, we have  $d_H(p_{2i+1}) = d - i\epsilon\beta > 0$ . Since  $d \leq \beta$ , it follows that  $|P| \leq \frac{2}{\epsilon}$ .  $\square$

After finding and augmenting an alternating path starting at  $v$  as described above, we repeat the same update procedure starting at  $u$  (we execute this update procedure sequentially).

**Remark.** When  $(u, v)$  is inserted to  $H$  — although no specific changes are needed — it is instructive to consider the case where the alternating path starting at  $v$  intersects  $u$ . (Recall that we first handle  $v$  and only later  $u$ .) If  $v$ 's path reaches  $u$  with a full edge (which was the only possibility in the bipartite case), then we would be done since the degree in  $H$  of each vertex would be the same as it was prior to the edge insertion, which reestablishes Properties (P1) and (P2) of an EDCS. (When we later handle  $u$ , nothing else will be done.) If  $v$ 's path reaches  $u$  with a deficient edge, then we just continue the path as we normally would — the fact that  $u$ 's degree in  $H$  is "momentarily" raised by 2 does not affect the correctness of the algorithm and its analysis; in such a case we still need to handle  $u$  in the end, and the fact that the alternating path reached it does not change that. Symmetrically, when  $(u, v)$  is deleted from  $H$ , the same claim holds as for insertion, only that here if  $v$ 's path reaches  $u$  with a deficient edge then we would be done, and if  $v$ 's path reaches  $u$  with a full edge, then we just continue the path normally.

We have thus shown the following:

LEMMA 3.2. *Following each edge update, we can reestablish Properties (P1) and (P2) using at most  $\frac{4}{\epsilon}$  insertions and deletions to and from  $H$ .*

Thus far we have explained how to find the alternating paths starting at  $v$  and  $u$  in lay terms, and also proved that they contain together at most  $\frac{4}{\epsilon}$  edges. However, we still haven't described the data structures required for finding and augmenting these paths efficiently. In order to find the alternating paths, we need to maintain the necessary data structures to identify full and deficient edges. Each vertex  $x$  will maintain its degree in  $H$ ,  $d_H(x)$ , as well as a partition of its adjacent edges into three lists: (1)  $F(x)$ , list of  $x$ 's full adjacent edges, (2)  $D(x)$ , list of  $x$ 's deficient adjacent edges, and (4)  $R(x)$ , list of  $x$ 's remaining edges. Consider the process of constructing the alternating path and assume we are currently at vertex  $x$ . If we are looking for a full (resp., deficient) edge, we check  $F(x)$  (resp.,  $D(x)$ ). Each test takes  $O(1)$  in the worst case (whether the list is empty or not). If the list is non-empty, we remove from it one arbitrary edge (say the first one on the list), again in  $O(1)$  time. After we find a full (resp., deficient) edge and continue on our path, we must then remove the previous edge from the list  $F(x)$  (resp.,  $D(x)$ ) and add it to  $R(x)$ , yet again in  $O(1)$  time. Since we spend  $O(1)$  time for processing each vertex along the path, computing and augmenting the entire alternating path (which also concerns updating  $F$ ,  $D$  and  $R$  for the vertices along the way) can be done in total  $O(\frac{1}{\epsilon})$  time. Moreover, at the end of each such update, we have (up to) 2 vertices whose degree in  $H$  has changed, i.e., the final vertex of  $v$ 's alternating path and the final vertex in  $u$ 's alternating path. Consequently, for each neighbor  $w$  of any of these (up to) 2 vertices whose degree in  $H$  has changed, we spend  $O(1)$  time to update the lists  $F(w)$ ,  $D(w)$  and  $R(w)$  accordingly. Since the degree in  $G$  is upper bounded by  $\Delta$ , the worst-case update time of this process is  $O(\Delta)$ . Summarizing, the total worst-case update time is  $O(\Delta + \frac{1}{\epsilon})$ .

#### 4 Maintaining an EDCS, Part II: Improved Update Time $O_\epsilon(\sqrt{\Delta})$

In this section we achieve roughly a quadratic improvement over the worst-case update time achieved in 3. To this end, we tweak the algorithm of 3 slightly: instead of *notifying* all the neighbors of a vertex following a *permanent*

---

**Algorithm 1** Insert  $(u, v)$ 

---

**if**  $d_H(u) + d_H(v) < (1 - \epsilon)\beta$  **then**  
     $E_H \leftarrow E_H \cup \{(u, v)\}$   
    HandleFull  $(v)$   
    HandleFull  $(u)$   
**end if**

---

---

**Algorithm 2** Delete  $(u, v)$ 

---

**if**  $(u, v) \in E_H$  **then**  
     $E_H \leftarrow E_H \setminus \{(u, v)\}$   
    HandleDeficient  $(v)$   
    HandleDeficient  $(u)$   
**end if**

---

---

**Algorithm 3** HandleFull  $(p_i)$ 

---

**if**  $\text{EMPTY}(F(p_i)) = \text{FALSE}$  **then**  
     $x \leftarrow \text{POP}(F(p_i))$  // i.e., we set  $x$  as the first element of  $F(p_i)$ , and then remove  $x$  from  $F(p_i)$   
     $E_H \leftarrow E_H \setminus \{(p_i, x)\}$   
    HandleDeficient  $(x)$   
**end if**

---

---

**Algorithm 4** HandleDeficient  $(p_i)$ 

---

**if**  $\text{EMPTY}(D(p_i)) = \text{FALSE}$  **then**  
     $x \leftarrow \text{POP}(D(p_i))$  // i.e., we set  $x$  as the first element of  $D(p_i)$ , and then remove  $x$  from  $D(p_i)$   
     $E_H \leftarrow E_H \cup \{(p_i, x)\}$   
    HandleFull  $(x)$   
**end if**

---

change in its degree in  $H$ , we only notify a subset of its neighbors; if this vertex is at the end of an alternating path, we refer to its degree change as permanent, to distinguish from a momentary degree change of a vertex not in the end of that path — though of course the degree of that vertex in  $H$  might change again in subsequent update steps. By “notify” a neighbor we mean that the algorithm updates all the relevant data structures of that neighbor concerning the degree change. Next, we explain this tweak in detail. For each vertex  $x$ , we will maintain a *cyclic queue*  $Q_x$  that holds all its neighbors, and each time there is a permanent change in  $x$ 's degree in  $H$  (since it is the last vertex of an alternating path), it will only notify the first  $\frac{10\Delta}{\epsilon\beta}$  neighbors in  $Q_x$  (so that they can update their relevant data structures), and then these neighbors will be moved to the end of  $Q_x$ ; new neighbors will be added to the end of  $Q_x$  as well, and the point is that once a neighbor (old or new) has been moved to the end of  $Q_x$ , all its data structures are up-to-date with respect to  $x$ 's current degree in  $H$ . Edge deletions incident on  $x$  are also reflected in  $Q_x$ , by simply removing these vertices from  $Q_x$ , wherever they might be there (via appropriate pointers). Clearly, this tweak reduces the worst-case update time from  $O(\Delta + 1/\epsilon)$  to  $O(\frac{\Delta}{\epsilon\beta} + 1/\epsilon)$ , which we record in the next lemma for further use:

LEMMA 4.1. *The worst-case update time of this tweaked algorithm is  $O(\frac{\Delta}{\epsilon\beta} + 1/\epsilon)$ .*

It remains to analyze the ramifications of notifying only a fraction of the neighbors of a vertex following its permanent degree change in  $H$ . For a vertex  $x$ , recall that  $d_H(x)$  denotes  $x$ 's degree in  $H$ ; for a neighbor  $w$  of  $x$ , i.e.,  $w \in N(x)$ , denote by  $\tilde{d}_H^w(x)$  the estimation that  $w$  has on  $x$ 's degree in  $H$ . We are interested in upper bounding  $\text{Dis}(H) := \max_{x \in V, w \in N(x)} \{d_H(x) - \tilde{d}_H^w(x)\}$ , i.e., the maximum discrepancy, over all vertices, between the degree of a vertex in  $H$  and its estimated degree by any of its neighbors. The maximum degree in  $G$  is  $\leq \Delta$ , thus the number of vertices in the cyclic queue  $Q_x$  of  $x$  is at most  $\Delta$ . Since  $Q_x$  is cyclic, and it updates one batch of  $\frac{10\Delta}{\epsilon\beta}$  vertices at a time, it follows that  $\frac{\epsilon\beta}{10}$  would be the maximum number of batches before  $w$  gets notified by  $v$  again by its degree change in  $H$ , hence  $\text{Dis}(H) \leq \frac{\epsilon\beta}{10}$ .

The following lemma shows that we can still upper bound the length of any alternating path by  $O(1/\epsilon)$ .

LEMMA 4.2. *For any path  $P$  of alternating full and deficient edges, we have  $|P| \leq \frac{5}{2\epsilon}$ .*

*Proof.* Consider first the case that  $P = (p_0, p_1, \dots, p_k)$  starts with a full edge. Recall that  $\tilde{d}_H^{p_i}(p_{i+1})$  is the estimation of  $p_i$  on the degree of  $p_{i+1}$  in  $H$ . Let  $d = d_H(p_0)$ . Since  $(p_0, p_1)$  is full "in the eyes" of  $p_0$  (i.e., according to the degree estimation on  $p_1$  that  $p_0$  has), we get that  $\tilde{d}_H^{p_0}(p_1) = \beta - d_H(p_0) = \beta - d$ . Since  $\text{Dis}(H) \leq \frac{\epsilon\beta}{10}$ , we get that  $d_H(p_1)$  and  $\tilde{d}_H^{p_0}(p_1)$  differ by at most an additive term of  $\frac{\epsilon\beta}{10}$ , i.e.,  $d_H(p_1) = \tilde{d}_H^{p_0}(p_1) \pm \frac{\epsilon\beta}{10}$ ; in what follows we shall use the  $\pm$  operator to “absorb” these additive terms. Similarly, since  $(p_1, p_2)$  is deficient in the eyes now of  $p_1$ , we get that

$$\tilde{d}_H^{p_1}(p_2) = (1 - \epsilon)\beta - d_H(p_1) = (1 - \epsilon)\beta - \tilde{d}_H^{p_0}(p_1) \pm \frac{\epsilon\beta}{10} = (1 - \epsilon)\beta - \beta + d \pm \frac{\epsilon\beta}{10} = d - \epsilon\beta \pm \frac{\epsilon\beta}{10}.$$

Continuing in the same way, we get  $\tilde{d}_H^{p_2}(p_3) = \beta - d + \epsilon\beta \pm \frac{2\epsilon\beta}{10}$ , and then  $\tilde{d}_H^{p_3}(p_4) = d - 2\epsilon\beta \pm \frac{3\epsilon\beta}{10}$ , and in general we inductively get  $\tilde{d}_H^{p_{2i-1}}(p_{2i}) = d - i\epsilon\beta \pm \frac{(2i-1)\epsilon\beta}{10}$ , which implies that  $d_H(p_{2i}) \leq d - i\epsilon\beta + \frac{i\epsilon\beta}{5} = d - \frac{4i\epsilon\beta}{5}$ . Since each vertex in  $P$  has an incident full edge, all vertices have a positive degree, and since  $d \leq \beta$ , it follows that  $|P| \leq \frac{5}{2\epsilon} - 1$ . If  $P$  starts with a deficient edge, we can go through the same argument again (here too following similar lines as those in the proof of 3.1), to conclude that  $|P| \leq \frac{5}{2\epsilon}$ .  $\square$

COROLLARY 4.1. *The worst-case recourse bound for maintaining  $H$  is  $O(1/\epsilon)$ .*

We next argue that, while Properties (P1) and (P2) may no longer hold — and thus  $H$  may no longer provide a  $(\beta, (1 - \epsilon)\beta)$ -EDCS — similar properties, for a slightly different choice of parameters, do hold.

LEMMA 4.3. *Define  $\gamma = \beta(1 + \epsilon/10)$ . The tweaked algorithm of this section maintains a  $(\gamma, (1 - 2\epsilon)\gamma)$ -EDCS, or in other words it satisfies the following two properties:*

(P1') *If  $(u, v) \in H$ , then  $w(u, v) \leq \gamma$ .*

(P2') *If  $(u, v) \notin H$ , then  $w(u, v) \geq (1 - 2\epsilon)\gamma$ .*

*Proof.* If  $(u, v) \in H$ , it no longer holds that  $w(u, v) \leq \beta$ , due to the discrepancy  $\text{Dis}(H) \leq \frac{\epsilon\beta}{10}$  between the degrees of vertices and the estimations of those by their neighbors. Following a permanent increase in the degree of some vertex  $u$  in  $H$ , it is increase-safe by definition, meaning that none of its adjacent edges in  $H$  are full. However, the increase-safe condition is defined with respect to the estimated degree at the other endpoints, so in fact it only holds at this stage that

$$w(u, v) = d_H(u) + \tilde{d}_H^u(v) \leq d_H(u) + d_H(v) + \epsilon\beta/10 \leq \beta(1 + \epsilon/10).$$

Symmetrically, it no longer holds that  $w(u, v) \geq (1 - \epsilon)\beta$ , but rather  $w(u, v) \geq (1 - \epsilon)\beta - \epsilon\beta/10 = \beta(1 - 11\epsilon/10)$ . We next argue that  $H$  provides a  $(\gamma, (1 - 2\epsilon)\gamma)$ -EDCS, for  $\gamma = \beta(1 + \epsilon/10)$ . Indeed, for any  $(u, v) \in H$ , we have  $w(u, v) \leq \beta(1 + \epsilon/10) = \gamma$ , whereas for any  $(u, v) \notin H$ , we have  $w(u, v) \geq \beta(1 - 11\epsilon/10) \geq \gamma(1 - 2\epsilon)$ , where the last inequality holds for any  $\epsilon > 0$ .  $\square$

**Putting it all together.** We need to calibrate the parameters properly, so as to be able to apply 2.1. In 3 we did not bother ourselves with such details, since the bound achieved there just served as a stepping stone towards the one achieved in this section; we simply mentioned there that we maintain a  $(\beta, (1 - \epsilon)\beta)$ -EDCS. Next, we will be more precise. Our aim here is to argue that the maintained EDCS  $H$  is a  $(3/2 + \epsilon)$ -approximate matching sparsifier for  $G$ , for a parameter  $\epsilon > 0$ ; thus  $\epsilon$  is now fixed. We take  $\lambda' = \lambda/2 = \frac{\epsilon}{200}$  and we will take  $\beta$  so that it is at least  $32\lambda^{-3}$ . Now the algorithm of 3 will maintain a  $(\beta, (1 - \lambda')\beta)$ -EDCS (so  $\epsilon$  as used in 3 is replaced by  $\lambda'$ ), but more importantly, by 4.3, the algorithm of 4 maintains a  $(\gamma, \gamma^-)$ -EDCS, for  $\gamma^- = (1 - 2\lambda')\gamma = (1 - \lambda)\gamma$ . It remains to verify that these parameters satisfy the conditions of 2.1, indeed:  $\lambda = 2\lambda' = \frac{\epsilon}{100}$ ,  $\gamma > \beta \geq 32\lambda^{-3}$  and  $\gamma^- = (1 - \lambda)\gamma$ . Thus  $H$  provides a  $(3/2 + \epsilon)$ -approximate matching sparsifier for  $G$ , as required.

By 4.1, the worst-case update time of maintaining  $H$  is  $O(\frac{\Delta}{\epsilon\beta} + 1/\epsilon)$ . Moreover, note that  $H$  has a maximum degree of  $O(\beta)$  at all times. We use the Gupta-Peng [GP13] algorithm, whose worst-case update time per change in  $H$  is  $O(\frac{\beta}{\epsilon^2})$ , to maintain a  $(1 + \epsilon)$ -MCM on top of  $H$ . Combining the approximation ratio of the EDCS from Lemma 2.1 with the  $(1 + \epsilon)$  approximation of the Gupta-Peng [GP13] algorithm yields a  $(3/2 + \epsilon)(1 + \epsilon) \leq (3/2 + 7/2\epsilon)$  approximation. Rescaling  $\epsilon$  by a factor  $7/2$  yields the desired  $(3/2 + \epsilon)$ -approximation, which will only have a constant factor effect on the update time. By 4.1, the worst-case recourse bound of  $H$  is  $O(1/\epsilon)$ , thus a single update in  $G$  is handled within a worst-case update time of  $O(\frac{\beta}{\epsilon^3})$ . It follows that the worst-case update time for maintaining a  $(1 + \epsilon)$ -MCM on top of  $H$  is  $O(\frac{\Delta}{\epsilon\beta} + \frac{\beta}{\epsilon^3})$ , for any parameter  $\beta \geq 32\lambda^{-3}$ . We can optimize by substituting  $\beta$  with  $\epsilon\sqrt{\Delta}$  (assuming  $\epsilon\sqrt{\Delta} \geq 32\lambda^{-3}$ , otherwise we need to increase  $\beta$  accordingly), which yields:

LEMMA 4.4. *The worst-case update time for maintaining a  $(1 + \epsilon)$ -MCM for  $H$  is  $O(\sqrt{\Delta}/\epsilon^2 + \epsilon^{-6})$ .*

## 5 Input Sparsification, Part III: Final Update Time $O_\epsilon(m^{1/4})$ or $O_\epsilon(\sqrt{\alpha})$

In this section we show how to combine a simple *sparsification* step with the  $O_\epsilon(\sqrt{\Delta})$  update time algorithm of 4 (Part II) to achieve the desired update time of  $O_\epsilon(\sqrt{\alpha})$ , which is  $O_\epsilon(m^{1/4})$  in general graphs.

The key component is a matching sparsification algorithm by Solomon [Sol18], which works as follows. Suppose the arboricity of a given  $m$ -edge graph  $G$  is bounded above by  $\alpha$ ; as mentioned,  $\alpha$  is no greater than  $\sqrt{m}$ . Every vertex marks up to  $\eta$  arbitrary incident edges; that is, if a vertex has at most  $\eta$  incident edges then it marks all of them, otherwise it marks exactly  $\eta$  arbitrary incident edges. An edge that is marked twice (by both endpoints) is added to the sparsified graph, denoted by  $G' = G'_\eta$ . In [Sol18] it was shown that for  $\eta = 5(5/\epsilon + 1)2\alpha$ ,  $G'$  is a  $(1 + \epsilon)$ -approximate matching sparsifier for  $G$ , i.e.,  $\mu(G) \leq (1 + \epsilon)\mu(G')$ . Moreover, the maximum degree of  $G'$  is trivially bounded by  $\eta = O(\alpha/\epsilon)$ .

Thus, our goal is to dynamically maintain  $G'$  on top of the input dynamic graph  $G$ , and feed  $G'$  (rather than  $G$ ) to the algorithm from Part II. In what follows we fix  $\epsilon > 0$  to be an arbitrary parameter.

**A density-sensitive degree bound.** First, we shall assume that  $\alpha$  is a fixed upper bound on the arboricity of the dynamic graph  $G$ , and show that the sparsifier  $G'$  can be maintained with a constant worst-case update time and recourse bound, where the maximum degree of  $G'$  is always  $O_\epsilon(\alpha)$ .

LEMMA 5.1. *We can maintain  $G' = G'_\eta$ , for  $\eta := 5(5/\epsilon + 1)2\alpha$ , with a constant worst-case (deterministic) update time and recourse bound.*

*Proof.* For every vertex  $u$ , we maintain a doubly-linked list of incident marked and unmarked edges, denoted by  $LM(u)$  and  $LU(u)$ , respectively; we also maintain a doubly-linked list of its incident edges in  $G'$ , denoted by  $G'(u)$ . Furthermore, every edge has mutual pointers to its various incarnations in all the lists it is contained in, as well as to the endpoints of the edge.

When a new edge  $e = (u, v)$  is added to  $G$ , we add  $e$  to  $LM(u)$  (resp.,  $LM(v)$ ) if  $|LM(u)| < \eta$  (resp.,  $|LM(v)| < \eta$ ), otherwise we add  $e$  to  $LU(u)$  (resp.,  $LU(v)$ ). If  $e$  belongs to both  $LM(u)$  and  $LM(v)$ , it is also added to  $G'(u)$  and  $G'(v)$ , and thus to  $G'$ . Clearly, the update time is constant and  $G'$  can change by at most one edge.

When an edge  $e = (u, v)$  is deleted from  $G$ , if  $e \in LU(u)$  (resp.  $LU(v)$ ), we simply remove it. If  $e$  is in  $G'$ , we remove it from there by removing it from the corresponding lists  $G'(u)$  and  $G'(v)$ . If  $e \in LM(u)$  (resp.,  $e \in LM(v)$ ), we remove an arbitrary edge from  $LU(u)$  (resp.  $LU(v)$ ) if one exists and insert it into  $LM(u)$  (resp.,  $LM(v)$ ); in the latter case, if an edge that is moved to  $LM(u)$  (resp.,  $LM(v)$ ) is also marked by the other endpoint, it is also added to  $G'$ . Here too the update time is constant, but now  $G'$  can change by at most three edges:  $e$  may be deleted from  $G'$ , and the edges moved from  $LU(u)$  and  $LU(v)$  to  $LM(u)$  and  $LM(v)$ , respectively, if any, may be inserted to  $G'$ .  $\square$

**COROLLARY 5.1.** *One can dynamically maintain a  $(1 + \varepsilon)$ -approximate matching sparsifier  $G' = G'_\eta$  for any dynamic graph  $G$  with a constant worst-case update time and recourse bound, where the maximum degree of  $G'$  is always upper bounded by  $\eta = 5(5/\varepsilon + 1)2\alpha$  and  $\alpha$  is a fixed upper bound on the arboricity.*

**A degree bound of  $O(\sqrt{m}/\varepsilon)$ , for a dynamic  $m$ .** Next, we shall drop the assumption that an upper bound  $\alpha$  on the arboricity is given. Starting from a graph with no edges, we will show how to maintain the sparsifier  $G'$  with a constant worst-case update time and recourse, where the maximum degree of  $G'$  is always  $O_\varepsilon(\sqrt{m})$ , where  $m$  is the dynamically changing number of edges in the graph.

Clearly, in the degenerate case that the number of edges always remains the same up to, say, a factor of 2, we can use the argument of 5.1 verbatim, by setting  $\eta$  to be larger than  $5(5/\varepsilon + 1)2\sqrt{m}$  by a small constant factor. The general case is handled using a standard trick of *periodic restarts*. When the number of edges grows from  $m$  to  $2m$  or shrinks from  $m$  to  $m/2$ , we shall restart the sparsifier  $G'$  to be in accordance with the new value of the number of edges. For concreteness, if the number of edges at the moment of restart is denoted by  $m_R$ , we set  $\eta$  to be  $5(5/\varepsilon + 1)4\sqrt{m_R}$ , which is a factor 2 larger than the bound used in 5.1. If we aimed for an amortized bound, a restart would be straightforward. Indeed, if the number of edges grows from  $m$  to  $m_R = 2m$ , we set  $\eta$  to  $5(5/\varepsilon + 1)4\sqrt{m_R} = 5(5/\varepsilon + 1)4\sqrt{2}\sqrt{m}$ , and then scan, for each non-isolated vertex  $v$  in the graph, all its incident edges, in order to mark its incident edges until (at most)  $\eta$  are "marked", which involves moving edges from  $LU(v)$  to  $LM(v)$  until  $|LM(v)| = \min\{\eta, |N(v)|\}$  and updating  $G'(v)$  accordingly. The case that the number of edges decreases from  $m$  to  $m_R = m/2$  is handled symmetrically: we set  $\eta$  to  $5(5/\varepsilon + 1)4\sqrt{m_R} = 5(5/\varepsilon + 1)2\sqrt{2}\sqrt{m}$ , and then scan for each non-isolated vertex  $v$ , all its incident edges, in order to "unmark" incident edges until (at most)  $\eta$  are marked, which involves moving edges from  $LM(v)$  to  $LU(v)$  until  $|LM(v)| = \min\{\eta, |N(v)|\}$  and updating  $G'(v)$  accordingly. In both cases this restart takes  $O(m)$  time, assuming we keep track of the non-isolated vertices in a separate list from the isolated ones, which can be amortized over all the edge updates occurred since the previous restart to achieve a constant *amortized* update time and recourse. Moreover, at any point in time, since the values of  $m_R$  and  $m$  (the number of edges at the time of the last restart and at the current time, respectively) differ by at most a factor of 2, the value of  $\eta$ , namely  $5(5/\varepsilon + 1)4\sqrt{m_R}$ , is the same as  $5(5/\varepsilon + 1)4\sqrt{m}$  up to a factor of  $\sqrt{2}$ . As  $\eta$  upper bounds the maximum degree in  $G'$ , we conclude that the maximum degree of  $G'$  is at most  $5(5/\varepsilon + 1)4\sqrt{2}\sqrt{m}$ . Moreover,  $\eta$  is at least  $5(5/\varepsilon + 1)2\sqrt{2}\sqrt{m}$ , which allows us to apply 5.1; in fact, we have an extra factor of  $\sqrt{2}$  on the value  $\eta$ , which is not needed here, but will be needed next for achieving a worst-case bound.

Since we are interested in achieving a constant *worst-case* update time and recourse, we must implement the restart process *gradually* rather than instantaneously, simulating  $c$  "computational steps" rather than 1 per each update step, for a large constant  $c$  of our choice, where each computational step corresponds to a sufficiently large constant running time. Of course, each edge insertion and deletion to and from the graph should also be processed, as it normally would, so we process it as in 5.1 but according to the new value of  $\eta$  that was set at the start of the last restart — which absorbs 1 computational step out of  $c$ . Then we use the remaining  $c - 1$  computational steps to continue the restart process from where we left off in the previous edge update. Since we are simulating  $c$  computational steps per edge update and as we can control the value of  $c$ , the entire restart

process will be completed within less than  $\rho \cdot m$  update steps, where  $\rho \ll 1$  is a small constant of our choice, at which stage the number of edges is between  $m_R - \rho \cdot m$  and  $m_R + \rho \cdot m$ . The worst-case update time and recourse bound are both constants by design, and it is easy to see that the maximum degree of the maintained sparsifier  $G'$  is equal, up a constant factor, to  $5(5/\varepsilon + 1)4\sqrt{m}$ , where  $m$  is the dynamic number of edges. The aforementioned extra factor of  $\sqrt{2}$  on the value of  $\eta$  allows us to apply 5.1 in this case too. Summarizing, we have proved:

**COROLLARY 5.2.** *One can dynamically maintain a  $(1 + \varepsilon)$ -approximate matching sparsifier  $G' = G'_\eta$  for any dynamic graph  $G$  with a constant worst-case update time and recourse bound, where the maximum degree of  $G'$  is within a constant factor of  $\eta = 5(5/\varepsilon + 1)4\sqrt{m}$  and  $m$  is the dynamic number of edges.*

**Completing the proofs of 1.1 and 1.2.** Combining 5.1 and 5.2 in conjunction with 4.4, we maintain in this way a  $((3/2 + \varepsilon) \cdot (1 + \varepsilon))$ -MCM for the original graph  $G$ , which provides a  $(3/2 + \varepsilon)$ -MCM after rescaling  $\varepsilon$  by a constant factor. By 5.1 and 5.2, the worst-case recourse bound of maintaining  $G'$  is constant, hence by 4.4 the overall worst-case update time is  $O(\sqrt{\Delta}/\varepsilon^2 + \varepsilon^{-6})$ , which is  $O(m^{1/4}\varepsilon^{-2.5} + \varepsilon^{-6})$  for general graphs and  $O(\sqrt{\alpha}\varepsilon^{-2.5} + \varepsilon^{-6})$  for graphs with arboricity bounded by  $\alpha$ . 1.1 and 1.2 follow.

## 6 Acknowledgements

Fabrizio Grandoni was partially supported by the SNSF Excellence Grant 200020B\_182865\1. Shay Solomon and Amitai Uzrad are supported by the Israel Science Foundation (ISF) grant No.1991\1, and by a grant from the United States-Israel Binational Science Foundation (BSF), Israel, and the United States National Science Foundation (NSF).

## References

- [AB19] Sepehr Assadi and Aaron Bernstein. Towards a unified theory of sparsification for matching problems. In Jeremy T. Fineman and Michael Mitzenmacher, editors, *2nd Symposium on Simplicity in Algorithms, SOSA 2019, January 8-9, 2019, San Diego, CA, USA*, volume 69 of *OASICS*, pages 11:1–11:20. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019.
- [ABB<sup>+</sup>19] Sepehr Assadi, MohammadHossein Bateni, Aaron Bernstein, Vahab S. Mirrokni, and Cliff Stein. Coresets meet EDCS: algorithms for matching and vertex cover on massive graphs. In Timothy M. Chan, editor, *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2019, San Diego, California, USA, January 6-9, 2019*, pages 1616–1635. SIAM, 2019.
- [Ber20] Aaron Bernstein. Improved bounds for matching in random-order streams. In Artur Czumaj, Anuj Dawar, and Emanuela Merelli, editors, *47th International Colloquium on Automata, Languages, and Programming, ICALP 2020, July 8-11, 2020, Saarbrücken, Germany (Virtual Conference)*, volume 168 of *LIPICs*, pages 12:1–12:13. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020.
- [BFH19] Aaron Bernstein, Sebastian Forster, and Monika Henzinger. A deamortization approach for dynamic spanner and dynamic maximal matching. In Timothy M. Chan, editor, *Proc. 50th SODA*, pages 1899–1918, 2019.
- [BGS11] S. Baswana, M. Gupta, and S. Sen. Fully dynamic maximal matching in  $O(\log n)$  update time. In *Proceedings of the 52nd IEEE Annual Symposium on Foundations of Computer Science, FOCS 2011, Palm Springs, CA, October 23-25, 2011*, pages 383–392, 2011. See also the *SICOMP'15* version, and the subsequent erratum.
- [BHN16] S. Bhattacharya, M. Henzinger, and D. Nanongkai. New deterministic approximation algorithms for fully dynamic matching. In *STOC*, 2016.
- [BK21] Sayan Bhattacharya and Peter Kiss. Deterministic rounding of dynamic fractional matchings. In Nikhil Bansal, Emanuela Merelli, and James Worrell, editors, *Proc. 48th ICALP*, volume 198, pages 27:1–27:14, 2021.
- [BLM20] Soheil Behnezhad, Jakub Lacki, and Vahab S. Mirrokni. Fully dynamic matching: Beating 2-approximation in  $\Delta^\varepsilon$  update time. In Shuchi Chawla, editor, *Proc. 51th SODA*, pages 2492–2508, 2020.
- [BS15] Aaron Bernstein and Cliff Stein. Fully dynamic matching in bipartite graphs. In *Automata, Languages, and Programming - 42nd International Colloquium, ICALP 2015, July 6-10, 2015, Proceedings, Part I*, pages 167–179, 2015.
- [BS16] Aaron Bernstein and Cliff Stein. Faster fully dynamic matchings with small approximation ratios. In *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2016, January 10-12, 2016*, pages 692–711, 2016.
- [GLS<sup>+</sup>19] Fabrizio Grandoni, Stefano Leonardi, Piotr Sankowski, Chris Schwiegelshohn, and Shay Solomon.  $(1 + \varepsilon)$ -approximate incremental matching in constant deterministic amortized time. In Timothy M. Chan, editor, *Proc. 50th SODA*, pages 1886–1898, 2019.

- [GP13] M. Gupta and R. Peng. Fully dynamic  $(1 + \epsilon)$ -approximate matchings. In *Proceedings of the 54th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2013, Berkeley, CA, USA, October 26-29, 2013*, pages 548–557, 2013.
- [MS20] Lazar Milenkovic and Shay Solomon. A unified sparsification approach for matching problems in graphs of bounded neighborhood independence. In Christian Scheideler and Michael Spear, editors, *Proc. of 32nd SPAA*, pages 395–406. ACM, 2020.
- [NS13] Ofer Neiman and Shay Solomon. Simple deterministic algorithms for fully dynamic maximal matching. In *Proceedings of the 45th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2013, Palo Alto, CA, USA, June 1-4, 2013*, pages 745–754, 2013.
- [OR10] Krzysztof Onak and Ronitt Rubinfeld. Maintaining a large matching and a small vertex cover. In *Proceedings of the 42nd ACM Symposium on Theory of Computing, STOC 2010, 5-8 June 2010*, pages 457–464, 2010.
- [PS16] D. Peleg and S. Solomon. Dynamic  $(1 + \epsilon)$ -approximate matchings: A density-sensitive approach. In *Proceedings of the 27th Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2016, Arlington, VA, USA, January 10-12, 2016*, 2016.
- [RSW21] Mohammad Roghani, Amin Saberi, and David Wajc. Beating the folklore algorithm for dynamic matching. *CoRR*, abs/2106.10321, 2021.
- [Sol16] S. Solomon. Fully dynamic maximal matching in constant update time. In *Proceedings of the 57th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2016, New Brunswick, NJ, USA, October 9-11, 2016*, pages 325–334, 2016.
- [Sol18] Shay Solomon. Local algorithms for bounded degree sparsifiers in sparse graphs. In *9th Innovations in Theoretical Computer Science Conference (ITCS 2018)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2018.
- [Waj20] David Wajc. Rounding dynamic matchings against an adaptive adversary. In Konstantin Makarychev, Yury Makarychev, Madhur Tulsiani, Gautam Kamath, and Julia Chuzhoy, editors, *Proc. 52nd STOC*, pages 194–207, 2020.