



Deterministic Dynamic Matching in Worst-Case Update Time

Peter Kiss¹

Received: 8 July 2022 / Accepted: 27 June 2023
© The Author(s) 2023

Abstract

We present deterministic algorithms for maintaining a $(3/2 + \epsilon)$ and $(2 + \epsilon)$ -approximate maximum matching in a fully dynamic graph with worst-case update times $\hat{O}(\sqrt{n})$ and $\tilde{O}(1)$ respectively. The fastest known deterministic worst-case update time algorithms for achieving approximation ratio $(2 - \delta)$ (for any $\delta > 0$) and $(2 + \epsilon)$ were both shown by Roghani et al. (Beating the folklore algorithm for dynamic matching, 2021) with update times $O(n^{3/4})$ and $O_\epsilon(\sqrt{n})$ respectively. We close the gap between worst-case and amortized algorithms for the two approximation ratios as the best deterministic amortized update times for the problem are $O_\epsilon(\sqrt{n})$ and $\tilde{O}(1)$ which were shown in Bernstein and Stein (in: Proceedings of the twenty-seventh annual ACM-SIAM symposium on discrete algorithms, 2016) and Bhattacharya and Kiss (in: 48th international colloquium on automata, languages, and programming, ICALP 2021, 12–16 July, Glasgow, 2021) respectively. The algorithm achieving $(3/2 + \epsilon)$ approximation builds on the EDCS concept introduced by the influential paper of Bernstein and Stein (in: International colloquium on automata, languages, and programming, Springer, Berlin, 2015). Say that H is a (α, δ) -approximate matching sparsifier if at all times H satisfies that $\mu(H) \cdot \alpha + \delta \cdot n \geq \mu(G)$ (define (α, δ) -approximation similarly for matchings). We show how to maintain a locally damaged version of the EDCS which is a $(3/2 + \epsilon, \delta)$ -approximate matching sparsifier. We further show how to reduce the maintenance of an α -approximate maximum matching to the maintenance of an (α, δ) -approximate maximum matching building based on an observation of Assadi et al. (in: Proceedings of the twenty-seventh

Independent Work: Independently and concurrently to our work Grandoni et al. [arXiv'2021] has presented a fully dynamic algorithm for maintaining a $(3/2 + \epsilon)$ -approximate maximum matching with deterministic worst-case update time $O_\epsilon(\sqrt{n})$.

This work is supported by Engineering and Physical Sciences Research Council, UK (EPSRC) GrantEP/S03353X/1.

✉ Peter Kiss
peter.kiss@warwick.ac.uk

¹ Department of Computer Science, University of Warwick, Coventry, UK

annual (ACM-SIAM) symposium on discrete algorithms, (SODA) 2016, Arlington, VA, USA, January 10–12, 2016). Our reduction requires an update time blow-up of $\hat{O}(1)$ or $\tilde{O}(1)$ and is deterministic or randomized against an adaptive adversary respectively. To achieve $(2 + \epsilon)$ -approximation we improve on the update time guarantee of an algorithm of Bhattacharya and Kiss (in: 48th International colloquium on automata, languages, and programming, ICALP 2021, 12–16 July, Glasgow, 2021). In order to achieve both results we explicitly state a method implicitly used in Nanongkai and Saranurak (in: Proceedings of the twenty-seventh annual ACM symposium on theory of computing, 2017) and Bernstein et al. (Fully-dynamic graph sparsifiers against an adaptive adversary, 2020) which allows to transform dynamic algorithms capable of processing the input in batches to a dynamic algorithms with worst-case update time.

Keywords Dynamic algorithms · Matching · Approximate matching · EDCS

1 Introduction

In the dynamic setting our task is to maintain a 'good' solution for some computational problem as the input undergoes updates [1–6]. Our goal is to minimize the update time we need to spend in order to update the output when the input undergoes updates. One of the most extensively studied computational problems in the dynamic setting is approximate maximum matching. Our task is to maintain an α -approximate matching M in G , which is a matching which satisfies that $|M| \cdot \alpha \geq \mu(G)$ (where $\mu(G)$ represent the size of a maximum size matching of graph G). Due to the conditional lower bound of [7] the maintenance of an exact maximum matching (a 1-approximate maximum matching) requires at least $O(\text{poly}(n))$ update time. Hence, a long line of papers were focused on the possible approximation ratio-update time trade-offs achievable for $\alpha > 1$ [8–15].¹

If a dynamic algorithm computes the updated output after at most $O(T)$ time following any single change in the input we say that its update time is worst-case $O(T)$. A slight relaxation of this bound is that the algorithm takes at most $O(T \cdot k)$ total time to maintain the output over $k > 0$ consecutive updates to the input, for any k , in this case the update time of the algorithm is amortized $O(T)$.

A number of dynamic algorithms in literature utilize different levels of randomization [16–21]. However, currently all known techniques for proving update time lower bounds fail to differentiate between randomized and deterministic dynamic algorithms [7, 22–25]. Hence, understanding the power of randomization in the dynamic setting is an important research agenda. In the case of dynamic matching getting rid of randomization has proven to be difficult within the realm of $\tilde{O}(1)$ update time. While as early as the influential work of Onak and Rubinfeld [26] a randomized algorithm with $\tilde{O}(1)$ update time has been found the first deterministic algorithm with the same update time was first shown by Bhattacharya et al. [27]. For achieving $(2 + \epsilon)$ -approximation with worst-case update time there is still an $O(\text{poly}(n))$ factor difference between the fastest randomized and deterministic implementations ([17, 18] and [28] respectively).

¹ Throughout the paper \tilde{O} hides $\text{poly}(\log n, 1/\epsilon)$ factors and \hat{O} hides $\text{poly}(n^{o(1)}, 1/\epsilon)$ factors.

While amortized update time bounds don't tell us anything about worst-case update time some problems in the dynamic setting have proven to be difficult to solve efficiently without amortization. Notably, for the dynamic connectivity problem the first deterministic amortized update time solution by Holm et al. [29] has long preceded the first worst-case update time implementation of Kapron et al. [30] which required randomization.

Both of the algorithms presented by this paper carry the best of both worlds as they are deterministic and provide new worst-case update-time bounds.

Many dynamic algorithms such as [31, 32] rely on the robustness of the output of the output. To consider this in a context of matching as an example observe that if a matching M is α -approximate it remains $(\alpha \cdot (1 + O(\epsilon)))$ -approximate even after $\epsilon \cdot |M|$ edge updates. Hence, if we are to rebuild M after the updates we can amortize its reconstruction cost over $\epsilon \cdot |M|$ time steps. However, such an approach initially inherently results in amortized update time bound. In some cases with additional technical effort de-amortization was shown to be achievable for these algorithms [31, 33]. A natural question to ask is weather an amortized update time bound is always avoidable for amortized rebuild based dynamic algorithms.

To answer this question we explicitly present a versatile framework for improving the update time bounds of amortized rebuild based algorithms to worst-case while incurring only a $\tilde{O}(1)$ blowup in update time. Our framework was implicitly shown in Bernstein et al. [33] and Nanongkai and Saranurak [34]. To demonstrate the framework we present two new results:

Theorem 1 *There is a deterministic algorithm for maintaining a $(2 + \epsilon)$ -approximate matching in a fully dynamic graph with worst-case update time $O_\epsilon(\log^7(n)) = \tilde{O}(1)$ (where O_ϵ hides $O(\text{poly}(1/\epsilon))$ factors).*

For the approximation ratio of $(2 + \epsilon)$ the best known worst-case update time algorithm of $\tilde{O}(\sqrt{n})$ was show recently in [28]. However, $\tilde{O}(1)$ amortized update time algorithms were previously shown by [27, 32]. We show that an $O(\text{poly}(n))$ blowup in update time is not necessary to improve these bounds to worst-case.

Theorem 2 *There is a fully dynamic algorithm for maintaining a $(3/2 + \epsilon)$ -approximate maximum matching in worst-case deterministic $\hat{O}\left(\frac{m}{n \cdot \beta} + \beta\right)$ (for our choice of β) or $\hat{O}(\sqrt{n})$ update time (where \hat{O} hides $O(\text{poly}(n^{o(1)}, 1/\epsilon))$ factors).*

For achieving better than than 2-approximation the fastest known worst-case update time of $\tilde{O}(\sqrt{n} \sqrt[3]{m})$ was shown in [28]. Similar to the case of $(2 + \epsilon)$ -approximation there is an $\tilde{O}(\text{poly}(n))$ faster algorithm achieving the same approximation ratio shown in [14] using amortization. We again show that such a large blowup is not necessary in order to achieve worst-case update times.

In order to derive the later result we first show an amortized rebuild based algorithm for maintaining the widely utilized [14, 35–41] matching sparsifier EDCS introduced by Bernstein and Stein [35]. At the core of amortized rebuild based algorithms there is a static algorithm for efficiently recomputing the underlying data-structure. As the EDCS matching sparsifier (as far as we are aware) doesn't admit a deterministic near-linear time static algorithm, we introduce a relaxed version of the EDCS we refer to as

'damaged EDCS'. For constructing a damaged EDCS we show a deterministic $\tilde{O}(m)$ static algorithm. Say that matching sparsifier (or matching) H is (α, δ) -approximate if $\mu(H) \cdot \alpha + n \cdot \delta \geq \mu(G)$. A damaged EDCS is a $(3/2 + \epsilon, \delta)$ -approximate matching sparsifier as opposed to the EDCS which is $(3/2 + \epsilon)$ -approximate. To counter this we show new reductions from $(\alpha + \epsilon)$ to (α, δ) -approximate dynamic matching algorithms based on ideas of [42, 43]. Previous such reductions relied on the *oblivious adversary* assumption that the input sequence is independent from the choices of the algorithm and is fixed beforehand. Our reductions work against an *adaptive adversary* whose decisions may depend on the decisions and random bits of the algorithm. The update time blowup required by the reductions is $\tilde{O}(1)$ or $\hat{O}(1)$ if the reduction step is randomized or deterministic respectively. These reductions and the static algorithm for constructing a damaged EDCS might be of independent research interest. Using the randomized reduction we receive the following corollary:

Corollary 3 *The update time bound of Theorem 2 can be improved to $\tilde{O}\left(\frac{m}{n \cdot \beta} + \beta\right)$ (or $\tilde{O}(\sqrt{n})$) if we allow for randomization against an adaptive adversary (where \tilde{O} hides $O(\text{poly}(\log(n), 1/\epsilon))$ factors).*

1.1 Techniques

We base our approach for improving an amortized rebuild based algorithm to worst-case update time on an observation implicitly stated in Bernstein et al. [33] (Lemma 6.1). Take an arbitrary input sequence of changes I for a dynamic problem and arbitrarily partition it into k continuous sub-sequences $I_i : i \in [k]$. If a dynamic algorithm with update time $O(T)$ is such that (knowing the partitionings) it can process the input sequence and the total time of processing sub-sequence I_i is $O(|I_i| \cdot T)$ then call it k batch-dynamic. Note that the update time guarantee of a batch-dynamic algorithm is stronger than of an amortized update time algorithm but it is weaker than a worst-case update time bound.

Building on the framework of [33] we show that an $O(\log(n))$ batch-dynamic algorithm Alg can be used to maintain $\tilde{O}(1)$ parallel output tapes with *worst-case update time* such that at all times at least one output tape contains a valid output of Alg while only incurring a blowup of $\tilde{O}(1)$ in update-time. If Alg is an α -approximate dynamic matching algorithm then each of the $O(\log(n))$ output tapes each contain a matching. Therefore, the union of the output tapes is an α -approximate matching sparsifier with maximum degree $O(\log(n))$ on which we can run the algorithm of Gupta and Peng [31] to maintain an $(\alpha + \epsilon)$ -approximate matching.

Therefore, in order to find new worst-case update time dynamic matching algorithms we only have to find batch-dynamic algorithms. We show a framework (building on [33]) for transforming amortized rebuild based dynamic algorithms to batch-dynamic algorithms. On a high level an amortized rebuild based algorithm allows for a slack of ϵ factor damage to its underlying data-structure before commencing a rebuild. To turn such an algorithm k batch-dynamic during the progressing of the i -th batch we ensure a slack of $\frac{i \cdot \epsilon}{k}$ instead. This way once the algorithm finishes processing a batch it has $\frac{\epsilon}{k}$ factor of slack it is allowed to take before commencing a rebuild

meaning that the next rebuild operation is expected to happen well into the proceeding batch.

With this general method and some technical effort we show a batch-dynamic version of the $(2 + \epsilon)$ -approximate dynamic matching algorithm of [32] and prove Theorem 1.

In order to generate a batch-dynamic algorithm for maintaining a $(3/2 + \epsilon)$ -approximate maximum matching more work is required as algorithms currently present in literature for this approximation ratio are not conveniently amortized rebuild based. We introduce a relaxed version of the matching sparsifier EDCS (initially appeared in [35]) called 'damaged EDCS'. We further show that a damaged EDCS can be found in $\tilde{O}(m)$ time. We show that a damaged EDCS is robust against $\tilde{O}(n \cdot \beta)$ edge updates and has maximum degree β for our choice of β . This means we can maintain the damaged EDCS in $\tilde{O}\left(\frac{m}{n \cdot \beta}\right)$ amortized update time with periodic rebuilds. We can then run the algorithm of [31] to maintain a matching in the damaged EDCS in $\tilde{O}(\beta)$ update time.

1.2 Independent Work

Independently from our work Grandoni et al. [36] presented a dynamic algorithm for maintaining a $(3/2 + \epsilon)$ -approximate matching with deterministic worst-case update time $O_\epsilon(m^{1/4})$, where O_ϵ is hiding $O(\text{poly}(1/\epsilon))$ dependency.

2 Notations and Preliminaries

Throughout this paper, we let $G = (V, E)$ denote the input graph and n will stand for $|V|$ and m will stand for the maximum of $|E|$ as the graph undergoes edge updates. $\text{deg}_E(v)$ will stand for the degree of vertex v in edge set E while $N_E(v)$ stand for the set of neighbouring vertices of v in edge set E . We will sometimes refer to $\text{deg}_E(u) + \text{deg}_E(v)$ as the degree of edge (u, v) in E . A matching M of graph G is a subset of vertex disjoint edges of E . $\mu(G)$ refers to the size of a maximum cardinality matching of G . A matching M is an α -approximate maximum matching if $\alpha \cdot |M| \geq \mu(G)$. Define a matching to be (α, δ) -approximate if $|M| \cdot \alpha + \delta \cdot n \geq \mu(G)$.

In the maximum dynamic matching problem the task is to maintain a large matching while the graph undergoes edge updates. In this paper we will be focusing on the fully dynamic setting where the graph undergoes both edge insertions and deletions over time. An algorithm is said to be a dynamic α (or (α, δ))-approximate maximum matching algorithm if it maintains an α (or (α, δ))-approximate matching at all times. A sub-graph $H \subseteq E$ is said to be an α (or (α, δ))-approximate matching sparsifier if it contains an α (or (α, δ))-approximate matching. We will regularly be referring to the following influential result from literature:

Lemma 4 *Gupta and Peng [31]: There is a $(1 + \epsilon)$ -approximate maximum matching algorithm for fully dynamic graph G with deterministic worst-case update time $O(\Delta/\epsilon^2)$ given the maximum degree of G is at most Δ at all times.*

Throughout the paper the notations $\tilde{O}()$, $\hat{O}()$ and $O_\epsilon()$ will be hiding $O(\text{poly}(\log(n), \epsilon))$, $O(\text{poly}(n^{o(1)}, \epsilon))$ and $O(\text{poly}(\frac{1}{\epsilon}))$ factors from running times respectively.

The update time of a dynamic algorithm is worst-case $O(T)$ if it takes at most $O(T)$ time for it to update the output each time the input undergoes a change. An algorithm update time is said to be amortized $O(T)$ if there is some integer $k > 0$ such that over k consecutive changes to the input the algorithm takes $O(k \cdot T)$ time steps to maintain the output. The recourse of a dynamic algorithm measures the changes the algorithm makes to its output per change to the input. Similarly to update time recourse can be amortized and worst-case.

We call a dynamic algorithm k batch-dynamic with update time $O(T)$ if for any partitioning of the input sequence I into k sub-sequences $I_i : i \in [k]$ during the processing of I the algorithm can process input sub-sequence I_i in $O(T \cdot |I_i|)$ total update time. Note that this implies that the worst-case update time during the progressing of I_i is $O(T \cdot |I_i|)$. The definition is based on [33]. A k -batch dynamic algorithm provides slightly better update time bounds than an amortized update time algorithm as we can select k sub-sequences to amortize the update time over.

We will furthermore be referring to the following recent result from Solomon and Solomon [44]:

Lemma 5 *Theorem 1.3 of Solomon and Solomon [44] (slightly phrased differently and trivially generalized for (α, δ) -approximate matchings): Any fully dynamic α (or (α, δ) -approximate maximum matching algorithm with update time $O(T)$ can be transformed into an $(\alpha + \epsilon)$ (or $(\alpha + \epsilon, \delta)$ -approximate maximum matching algorithm with $O(T + \frac{\alpha}{\epsilon})$ update time and worst-case recourse of $O(\frac{\alpha}{\epsilon})$ per update. The update time of the new algorithm is worst-case if so is the underlying matching algorithm.*

Definition 6 Random variables X_1, \dots, X_n are said to be negatively associated if for any non-decreasing functions g, f and disjoint subsets $I, J \subseteq [n]$ we have that:

$$\text{Cov}(g(X_i : i \in I), h(X_j : j \in J)) \leq 0$$

We will make use of the following influential result bounding the probability of a sum of negatively associated random variables falling far from their expectation.

Lemma 7 *(Chernoff bound for negatively associated random variables [45]): Let $\bar{X} = \sum_{i \in [n]} X_i$ where $X_i : i \in [n]$ are negatively associated and $\forall i \in [n] : X_i \in [0, 1]$. Then for all $\delta \in (0, 1)$:*

$$\Pr[\bar{X} \leq (1 - \delta) \cdot \mathbb{E}[\bar{X}]] \leq \exp\left(-\frac{\mathbb{E}[\bar{X}] \cdot \delta^2}{2}\right)$$

3 Batch Dynamic to Worst Case Update Time

3.1 Improving a Batch-Dynamic Algorithm to Amortized Update Time

Lemma 8 *Given an α approximate (or (α, ϵ) -approximate) dynamic matching algorithm Alg is $O(\log(n))$ batch-dynamic with update time $O(T(n))$ and dynamic graph*

G undergoing edge insertions and deletions. There is an algorithm Alg' which maintains $O(\log(n))$ matchings of G such that at all times during progressing an input sequence of arbitrarily large polynomial length one of the matchings is α approximate (or (α, ϵ) -approximate). The update time of Alg' is worst-case $O(T(n) \cdot \log^3(n))$ and it is deterministic if Alg is deterministic.

As this lemma was implicitly stated in [33] and [34] in a less general setting we defer the proof to Appendix A.

Corollary 9 *If there exists an α (or (α, δ))-approximate dynamic matching algorithm (where $\alpha = O(1)$) Alg which is $O(\log(n))$ batch-dynamic with update time $O(T(n))$ then there is an $(\alpha + \epsilon)$ (or $(\alpha + \epsilon, \delta)$)-approximate matching algorithm Alg' with worst case update time $O\left(\frac{T(n) \cdot \log^3(n)}{\epsilon^3}\right)$. If Alg is deterministic so is Alg' .*

Proof Maintain $O(\log(n))$ parallel matchings of G using the algorithm from Lemma 8 in $O(T(n) \cdot \log^3(n))$ worst case update time. Their union, say H , is a graph with maximum degree $O(\log(n))$ and is an α (or (α, δ))-approximate matching sparsifier and is a union of the output of $O(\log(n))$ dynamic matching algorithms with worst-case update time $O(T \cdot \log^2(n))$. By Lemma 5 ([44]) these approximate matching algorithms can be transformed into $(\alpha + \epsilon/2)$ (or $(\alpha + \epsilon/2, \delta)$)-approximate matching algorithms with $O(T \cdot \log^2(n) + \frac{\alpha}{\epsilon})$ update time and $O(\frac{\alpha}{\epsilon})$ worst-case recourse. This bounds the total recourse of the sparsifier at $O\left(\frac{\log(n) \cdot \alpha}{\epsilon}\right)$. Therefore, with slack parameter $\frac{\epsilon}{2 \cdot \alpha}$ we can run the algorithm of Lemma 4 ([31]) to maintain an $(\alpha + \epsilon)$ (or $(\alpha + \epsilon, \delta)$)-approximate matching in the sparsifier with worst-case update time $O\left(T \cdot \log^3(n) + \frac{\log(n) \cdot \alpha}{\epsilon} + \frac{\log^2(n) \cdot \alpha^2}{\epsilon^3}\right) = O\left(\frac{T \cdot \log^3(n)}{\epsilon^3}\right)$. \square

Observe that the framework outlined by Lemma 8 has not exploited any property of the underlying batch-dynamic algorithm other than the nature of its running time. This allows for a more general formulation of Lemma 8.

Corollary 10 *If there is a $O(\log(n))$ batch-dynamic algorithm Alg with deterministic (randomized) update time $O(T(n))$ and $poly(n)$ length input update sequence I then there is an algorithm Alg' such that*

- *The update time of Alg' is worst-case deterministic (randomized) $O(T(n) \cdot \log^3(n))$*
- *Alg' maintains $\log(n)$ parallel outputs and after processing update sequence $I[0, \tau]$ one of Alg' -s maintained outputs is equivalent to the output of Alg after processing $I[0, \tau]$ partitioned into at most $\log(n)$ batches*

4 Vertex Set Sparsification

An (α, δ) -approximate matching sparsifier satisfies that $\mu(H) \cdot \alpha + n \cdot \delta \geq \mu(G)$. Selecting $\delta = \frac{\epsilon \cdot \mu(H)}{n}$ results in a $(\alpha + \epsilon)$ -approximate sparsifier. The algorithm we present in this paper has a polynomial dependence on $1/\delta$ therefore we can't select the required δ value to receive an $(\alpha + \epsilon)$ -approximate sparsifier assuming $\mu(H)$ is

significantly lower than $\mu(G)$. To get around this problem we sparsify the vertex set to a size of $\hat{O}(\mu(H))$ while ensuring that the sparsified graph contains a matching of size $(1 - O(\epsilon)) \cdot \mu(G)$.

Let V^k be a partitioning of the vertices of $G = (V, E)$ into k sets $v^i : i \in [k]$. Define the concatenation of G based on V^k to be graph G_{V^k} on k vertices corresponding to vertex subsets v^i where there is an edge between vertices v^i and v^j if and only if there is $u \in v^i$ and $w \in v^j$ such that $(u, w) \in E$. Note that maintaining V^k as G undergoes edge changes can be done in constant time. Also note that given a matching M_{V^k} of G_{V^k} is maintained under edge changes to G_{V^k} in constant update time per edge changes to M_{V^k} we can maintain a matching of the same size in G .

4.1 Vertex Sparsification Against an Oblivious Adversary

Assume we are aware of $\mu(G)$ (note we can guess $\mu(G)$ within an $1 + \epsilon$ multiplicative factor through running $O(\frac{\log(n)}{\epsilon})$ parallel copies of the algorithm). Choose a partitioning of G 's vertices into $O(\mu(G)/\epsilon)$ vertex subsets V' uniformly at random. Define G' to be the concatenation of G based on V' .

Consider a maximum matching M^* of G . Its edges have $2 \cdot \mu(G)$ endpoints. Fix a specific endpoint v . With probability $(1 - \frac{2 \cdot \mu(G)}{\mu(G)/\epsilon})^{2 \cdot \mu(G) - 1} \sim (1 - o(\epsilon))$ it falls in a vertex set of V' no other endpoint of M^* does. Hence, in expectation $2 \cdot \mu(G) \cdot (1 - O(\epsilon))$ endpoints of M^* fall into unique vertex subsets of V' with respect to other endpoints. This also implies that $\mu(G) \cdot (1 - O(\epsilon))$ edges of M^* will have both of their endpoints falling into unique vertex sets of V' , hence $\mu(G') \geq \mu(G) \cdot (1 - O(\epsilon))$. This observation motivates the following lemma which can be concluded from [31] and [42, 43].

Lemma 11 *Assume there is a dynamic algorithm Alg which maintains an (α, δ) -approximate maximum matching where $\alpha = O(1)$ in graph $G = (V, E)$ with update time $O(T(n, \delta))$. Then there is a randomized dynamic algorithm Alg' which maintains an $(\alpha + \epsilon)$ -approximate maximum matching in update time $O\left(T(n, \epsilon^2) \cdot \frac{\log^2(n)}{\epsilon^4}\right)$. If the running time of Alg is worst-case (amortized) so will be the running time of Alg'.*

(Stated without proof as it concludes from [31, 42, 43])

4.2 Vertex Set Sparsification Using (k, ϵ) Matching Preserving Partitionings

A slight disadvantage of the method described above is that if the adversary is aware of our selection of V' they might insert a maximum matching within the vertices of a single subset in V' which would be completely lost after concatenation. In order to counter this we will do the following: we will choose some L different partitionings of the vertices in such a way that for any matching M of G most of M 's vertices fall into unique subsets in at least one partitioning.

Definition 12 Call a set of partitionings \mathcal{V} of the vertices of graph $G = (V, E)$ into d vertex subsets is (k, ϵ) matching preserving if for any matching of size k in G there

is a partitioning V_i^d in \mathcal{V} such that if G' is a concatenation of G based on V' then G' satisfies that $\mu(G') \geq (1 - \epsilon) \cdot k$.

We will show that using randomization we can generate a (k, ϵ) matching preserving set of partitionings of size $O\left(\frac{\log(n)}{\epsilon^2}\right)$ into $O(k/\epsilon)$ vertex subsets in polynomial time. Furthermore, we will show how to find an (k, ϵ) matching preserving set of partitionings of size $O(n^{(1)})$ into $O(k \cdot n^{o(1)})$ vertex subsets deterministically in polynomial time.

Lemma 13 *Assume there exists a dynamic matching algorithm Alg_M maintaining an (α, δ) -approximate matching in update time $O(T(n, \delta))$ for $\alpha = O(1)$ as well as an algorithm Alg_S generating an (k, ϵ) matching preserving set of vertex partitionings into $O(k \cdot C)$ vertex subsets of size L . Then there exists an algorithm Alg maintaining an $(\alpha + \epsilon)$ -approximate matching with update time $O\left(T(n, \epsilon/C) \cdot \frac{L^2 \cdot \log^2(n)}{\epsilon^4}\right)$. If both Alg_S and Alg_M are deterministic then so is Alg . If Alg_M is randomized against an adaptive adversary then so is Alg . If the update time of Alg_M is worst-case then so is of Alg . Alg makes a single call to Alg_S .*

The proof of the lemma is deferred to Appendix C.2. The intuition is as follows: through running $O\left(\frac{\log(n)}{\epsilon}\right)$ parallel copies of the algorithm guess $\mu(G)$ within a $1 + \epsilon$ factor. In the knowledge of $\mu(G)$ run Alg_M on the L concatenations of G we generate with Alg_S . Each of these concatenated sub-graphs are of size $O(\mu(G) \cdot C)$ and have maximum matching size $(1 - O(\epsilon)) \cdot \mu(G)$. Therefore running Alg_M with δ parameter $\Theta(\epsilon/C)$ yields an $(\alpha + O(\epsilon))$ -approximate matching in on of these L graphs. Using the algorithm of [31] find an approximate maximum matching in the union of the $O_\epsilon(L \cdot \log(n))$ concatenated graphs. Note that with an application of Lemma 5 ([44]) the update time can be changed into $O\left(\frac{T(n, \epsilon/C) \cdot L \cdot \log(n)}{\epsilon} + \frac{L^2 \cdot \log^2(n)}{\epsilon^5}\right)$ as shown in the appendix.

4.3 Generating Matching Preserving Partitionings Through Random Sampling

Lemma 14 *There is a randomized algorithm succeeding with $1 - 1/\text{poly}(n)$ probability for generating a (k, ϵ) matching preserving set of partitionings of graph G into $O(k/\epsilon)$ vertex subsets of size $O\left(\frac{\log(n)}{\epsilon^2}\right)$ running in polynomial time.*

We defer the proof to Appendix C.2. Essentially, $O\left(\frac{\log(n)}{\epsilon^2}\right)$ random chosen vertex partitionings into $O(k/\epsilon)$ vertex subsets are (k, ϵ) matching preserving. Note, that in unbounded time we can find an appropriate set of partitionings deterministically as we can iterate through all possible sets of partitionings and test each separately.

4.4 Generating Matching Preserving Partitionings Using Expanders

We will define expander graphs as follows. Such expanders are sometimes called unbalanced or lossless expanders in literature.

Definition 15 Define a (k, d, ϵ) -expander graph as a bipartite graph $G = ((L, R), E)$ such that $\forall v \in L : \text{deg}_E(v) = d$ and for any $S \subseteq L$ such that $|S| \leq k$ we have that $|N_E(S)| \geq (1 - \epsilon) \cdot d \cdot |S|$.

Graph expanders are extensively researched and have found a number of different applications. We will now show how an expander graph can be used to be the bases of an (k, ϵ) preserving set of partitionings.

Lemma 16 Assume there exists an algorithm Alg which outputs a (k, d, ϵ) -expander $G_{\text{exp}} = ((L, R), E)$ in $O(T(k, d, \epsilon))$ time. There is an algorithm Alg' which outputs a set of (k, ϵ) matching preserving vertex partitionings of a vertex set of size $|L|$ into $|R|$ subsets of size d with running time $O(T(k, d, \epsilon))$. Alg' is deterministic if Alg is deterministic.

Proof Take graph $G = (V, E)$ and bipartite $(2 \cdot k, d, \epsilon/2)$ expander graph $G_{\text{Exp}} = ((V, R), E')$ such that vertices of the left partition of G_{Exp} correspond to vertices of V . For each $v \in V$ define an arbitrary ordering of it's neighbours in R according to E' and let $N_{E'}(v)_i$ be it's i -th neighbour according to this ordering ($i \in [d]$). For each $i \in [d]$ and $v \in R$ define $V_{i,v} \subseteq V$ to be the set of vertices in V whose i -th neighbour is v (or $V_{i,v} = \{v' \in V : N_{E'}(v')_i = v\}$).

Define set of vertex partitionings $\mathcal{V} = \{V_i^{|R|} : i \in [d]\}$ where $V_i^{|R|}$ contain vertex sets $V_{i,v} : v \in R$. Fix a matching M in G of size k and call it's endpoints V_M . By the definition of the expander we have that $|N_{E'}(V_M)| \geq (1 - \epsilon/2) \cdot d \cdot 2k$. Hence by the pigeonhole principle we have that $|N_{E'}(V_{M^*})_i| \geq (1 - \epsilon/2) \cdot 2k$ for some $i \in [d]$. Define G' as the concatenation of G based on $V_i^{|R|}$. By the definition of $V_i^{|R|}$ at least $(1 - \epsilon/2) \cdot 2k$ endpoints of M are concatenated into vertices of G' containing exactly one vertex of V_M . Therefore, $(1 - \epsilon) \cdot k$ edges of M will have both their endpoints concatenated into unique vertices of G' within M . Hence, $\mu(G') \geq (1 - \epsilon) \cdot k$ and \mathcal{V} is a (k, ϵ) matching preserving set of partitionings. \square

Lemma 17 (Theorem 7.3 of [46] and Proposition 7 of [47]): Given $n \geq k$ and $\epsilon > 0$. There exists a (k, d, ϵ) -expander graph $G_{\text{exp}} = ((L, R), E)$ such that $|L| = n, |R| = \frac{k \cdot 2^{O(\log^3(\log(n)/\epsilon))}}{\text{poly}(\epsilon)} = \hat{O}(k), d = 2^{O(\log^3(\log(n)/\epsilon))} = \hat{O}(1)$ which can be deterministically computed in $\hat{O}(n)$ time.

4.5 Black-Box Implications

The following statements are black-box statements which can be concluded based on this section.

Corollary 18 [42, 43]: If there is a dynamic algorithm for maintaining an (α, δ) -approximate maximum matching for dynamic graphs in update time $O(T(n, \delta))$ then there is a randomized algorithm (against oblivious adversaries) for maintaining an $(\alpha + \epsilon)$ -approximate maximum matching with update time $O\left(T(n, \epsilon^2) \cdot \frac{\log^2(n)}{\epsilon^4}\right)$.

Corollary 19 *If there is a dynamic algorithm for maintaining an (α, δ) -approximate maximum matching for dynamic graphs in update time $O(T(n, \delta))$ then there is a randomized algorithm for maintaining an $(\alpha + \epsilon)$ -approximate maximum matching with update time $O\left(T(n, \epsilon^2) \cdot \frac{\log^4(n)}{\epsilon^8}\right)$ which works against adaptive adversaries given the underlying algorithm also does.*

Proof Follows from Lemmas 13 and 14. □

Corollary 20 *If there is a dynamic algorithm for maintaining an (α, δ) -approximate maximum matching for dynamic graphs in update time $O(T(n, \delta))$ then there is a deterministic algorithm for maintaining an $(\alpha + \epsilon)$ -approximate maximum matching with update time $\hat{O}\left(T\left(n, \frac{\text{poly}(\epsilon)}{n^{o(1)}}\right)\right)$ which is deterministic given the underlying matching algorithm is also deterministic.*

Proof Follows from Lemmas 13, 17 and 16. □

5 $(3/2 + \epsilon)$ -Approximate Fully Dynamic Matching in $\hat{O}(\sqrt{n})$ Worst-Case Deterministic Update Time

5.1 Algorithm Outline

In this section we present an amortized rebuild based algorithm for maintaining a locally relaxed EDCS we refer to as 'damaged EDCS'. The following definition and key-property originates from [35] and [38].

Definition 21 From Bernstein and Stein [35]:

Given graph $G = (V, E)$, $H \subseteq E$ is a (β, λ) -EDCS of G if it satisfies that:

- $\forall e \in H : \text{deg}_H(e) \leq \beta$
- $\forall e \in E \setminus H : \text{deg}_H(e) \geq \beta \cdot (1 - \lambda)$

Lemma 22 From Assadi and Stein [38]:

If $\epsilon < 1/2$, $\lambda \leq \frac{\epsilon}{32}$, $\beta \geq 8 \cdot \lambda^2 \cdot \log(1/\lambda)$ and H is a (β, λ) -EDCS of G then $\mu(G) \leq \mu(H) \cdot (\frac{3}{2} + \epsilon)$

The intuition behind the algorithm is as follows: take a (β, λ) -EDCS H . Relax it's parameter bounds slightly through observing that H is also a $(\beta \cdot (1 + \lambda), 4\lambda)$ -EDCS. As H is a (β, λ) -EDCS for every edge e in it's local neighbourhood $\Theta(\beta \cdot \lambda)$ edge updates may occur in an arbitrary fashion before either of the two edge degree bounds of a $(\beta \cdot (1 + \lambda), 4\lambda)$ -EDCS is violated on e .

Therefore, after $\tilde{\Theta}(n \cdot \beta)$ edge updates the properties of a $(\beta \cdot (1 + \lambda), 4\lambda)$ -EDCS should only be violated in the local neighbourhood of $O(\delta \cdot n)$ vertices for some small δ of our choice. At this point the EDCS is locally 'damaged' and it's approximation ratio as a matching sparsifier is reduced to $(3/2 + O(\epsilon), \delta)$. However, the reductions appearing in Sect. 4 allows us to improve this approximation ratio to $(3/2 + O(\epsilon))$. At this point we commence a rebuild, the cost of which can be amortized over $\tilde{\Theta}(n \cdot \beta)$ edge updates.

We then proceed to turn this amortized rebuild based algorithm into a batch-dynamic algorithm which we improve to worst-case update time using Lemma 8.

5.2 Definition and Properties of (β, λ, δ) Damaged EDCS

In order to base an amortized rebuild based dynamic algorithm on the EDCS matching sparsifier we need an efficient algorithm for constructing an EDCS. As far as we are aware there is no known deterministic algorithm for constructing an EDCS in $\tilde{O}(n)$ time. In order to get around this we introduce a locally relaxed version of EDCS.

Definition 23 For graph $G = (V, E)$ a (β, λ, δ) -damaged EDCS is a subset of edges $H \subseteq E$ such that there is a subset of 'damaged' vertices $V_D \subseteq V$ and the following properties hold:

- $|V_D| \leq \delta \cdot |V|$
- $\forall e \in H : deg_H(e) \leq \beta$
- All $e \in E \setminus H$ such that $e \cap V_D = \emptyset$ satisfies $deg_H(e) \geq \beta \cdot (1 - \lambda)$

Lemma 24 If $\epsilon < 1/2$, $\lambda \leq \frac{\epsilon}{32}$, $\beta \geq 8\lambda^{-2} \log(1/\lambda)$ and H is a (β, λ, δ) -damaged EDCS of graph $G = (V, E)$ then H is an $(3/2 + \epsilon, \delta)$ -approximate matching sparsifier.

Proof Define the following edge-set: $E' = \{e \in E : e \cap V_D = \emptyset\}$. Observe, that H is a (β, λ) -EDCS of $E' \cup H$. Fix a maximum matching M^* of G . At least $\mu(G) - |V_D| = \mu(G) - \delta \cdot |V|$ edges of M^* appear in E' as each vertex of V_D can appear on at most one edge of M^* . Therefore, $\mu((V, E')) \geq \mu(G) - |V| \cdot \delta$. Now the lemma follows from Lemma 22. □

5.3 Constructing a Damaged EDCS in Near-linear Time

Lemma 25 Algorithm 1 returns H_{fin} as a (β, λ, δ) -damaged EDCS of G .

The potential function Φ used in proof of the following lemma is based on [14].

Lemma 26 Algorithm 1 runs in deterministic $O\left(\frac{m}{\delta \cdot \lambda^2}\right)$ time.

The proofs of the lemmas are deferred to Appendix B. The intuition is the following: at the start of each iteration we add all edges of the graph to H which have $deg_H(e) < \beta \cdot (1 - \lambda/2)$. If we fail to add at least $O(\lambda \cdot \delta \cdot \beta \cdot n)$ such edges we terminate with H stripped of some edges. At the end of each iteration we remove all edges such that $deg_H(e) > \beta$. Consider what happens if we fail to add $\Omega(\lambda \cdot \delta \cdot \beta \cdot n)$ edges in an iteration. That means that only in the local neighbourhood of $\Theta(\delta \cdot n)$ 'damaged' vertices could we have added $\Omega(\beta \cdot \lambda)$ edges in the last iteration. We strip away the edges around damaged vertices to get H . The running time argument is based on a potential function Φ from [14]. Initially it is 0 and has an upper bound of $O(n \cdot \beta^2)$ and grows by at least $\Omega(n \cdot \beta^2 \cdot \lambda^2 \cdot \delta)$ in each iteration bounding the number of iterations by $O\left(\frac{1}{\delta \cdot \lambda^2}\right)$.

Algorithm 1: StaticDamagedEDCS

Input: $G = (V, E), \beta, \lambda, \delta$
Output: $H_{fin} \subseteq E: (\beta, \lambda, \delta)$ -damaged EDCS of G

```

1  $H = \emptyset$ 
2 repeat
3    $E' = \emptyset$ 
4   for  $e \in E/H$  do
5     if  $deg_H(e) < \beta \cdot (1 - \lambda/2)$  then
6        $H \leftarrow H \cup \{e\}$ 
7        $E' \leftarrow E' \cup \{e\}$ 
8   if  $|E'| \leq \frac{\delta \cdot \lambda \cdot \beta \cdot n}{16}$  then
9      $V_D \leftarrow \{v \in V : deg_{E'}(v) > \frac{\lambda \cdot \beta}{8}\}$ 
10     $E_D \leftarrow \{e \in E' : |e \cap V_D| > 0\}$ 
11     $H_{fin} \leftarrow H \setminus E_D$ 
12    Return  $H_{fin}$ 
13  for  $e \in H$  do
14    if  $deg_H(e) > \beta \cdot (1 - \lambda/4)$  then
15       $H \leftarrow H \setminus \{e\}$ 

```

5.4 Maintaining a Damaged EDCS in $\tilde{O}(\frac{m}{n \cdot \beta})$ pdate Time with Amortized Rebuilds

Algorithm 2: DynamicDamagedEDCS

Input: $G = (V, E), \beta, \lambda, \delta$
Output: $H \subseteq E: (\beta, \lambda, \delta)$ damaged-EDCS of G

```

1  $\alpha \leftarrow \frac{n \cdot \delta \cdot \lambda \cdot \beta}{64}$ 
2 Initially and after every  $\alpha$  edge updates
3    $H \leftarrow \text{StaticDamagedEDCS}(G, \frac{\beta}{1 + \lambda/4}, \lambda/4, \delta/2)$ 
4    $E_D \leftarrow \emptyset$ 
5    $E_I \leftarrow \emptyset$ 
6 Function InsertEdge( $(u, v)$ ):
7    $E_I \leftarrow E_I \cup \{(u, v)\}$ 
8   if  $\max\{deg_{E_I}(u), deg_{E_I}(v)\} < \frac{\beta \cdot \lambda}{16} - 1$  and  $deg_H((u, v)) \leq \beta - 2$  then
9      $H \leftarrow H \cup \{(u, v)\}$ 
10 Function DeleteEdge( $e$ ):
11    $E_D \leftarrow E_D \cup \{e\}$ 
12    $H \leftarrow H \setminus \{e\}$ 

```

Note that E_D is defined for the purposes of the analysis.

Lemma 27 *The sparsifier H maintained by Algorithm 2 is a (β, λ, δ) -damaged EDCS of G whenever the algorithm halts.*

Lemma 28 *The amortized update time of Algorithm 2 over a series of α updates is $O\left(\frac{m}{n \cdot \beta \cdot \lambda^3 \cdot \delta^2}\right)$ and the sparsifier H undergoes $O\left(\frac{1}{\lambda \cdot \delta}\right)$ amortized recourse.*

The lemmas are proven in the appendix. On a high level, a $(\beta, O(\lambda), O(\delta))$ damaged EDCS will gain $O(n \cdot \delta)$ damaged vertices in the span of $O(n \cdot \delta \cdot \lambda \cdot \beta)$ edge updates as for a vertex to be damaged there has to be $O(\beta \cdot \lambda)$ edge updates in its local neighbourhood. At this point we can call a rebuild of the EDCS in $\tilde{O}(m)$ time to get an amortized update time of $\tilde{O}\left(\frac{m}{\beta \cdot n}\right)$.

5.5 k Batch-Dynamic Algorithm for Maintaining an Approximate EDCS

Lemma 29 *Given fully dynamic graph G with n vertices and m edges. There is a k batch-dynamic dynamic algorithm which maintains a (β, λ, δ) -damaged EDCS of this graph with deterministic update time $O\left(\frac{k \cdot m}{n \cdot \beta \cdot \delta^2 \cdot \lambda^3}\right)$ and recourse $O\left(\frac{k}{\delta \cdot \lambda}\right)$.*

Proof Define an alternative version of Algorithm 2 where α is simply set to $\alpha_i = i \cdot \frac{\alpha}{k}$ during the processing of the i -th batch. Observe that in the proof of Lemma 27 the only detail which depends on the choice of α is the size of $V_{E_D} \cup V_{E_I}$. At any point in this batch modified version of the algorithm $\alpha_i \leq \alpha$ therefore the correctness of the algorithm follows.

The running time of the algorithm will be affected by this change. As every edge update is processed in constant time by the algorithm the running time is dominated by calls to StaticDamagedEDCS. By definition for every batch at least α/k edge updates will occur between the start of the batch and the first rebuild (if there is one) yielding an amortized update time of at most $O\left(\frac{k \cdot m}{n \cdot \beta \cdot \delta^2 \cdot \lambda^3}\right)$ over the first rebuild (due to Lemma 28). After the first rebuild the algorithm simply proceeds to run with α -parameter α_i therefore the amortized update time for the remainder of batch i is $O\left(\frac{i \cdot m}{n \cdot \beta \cdot \delta^2 \cdot \lambda^3}\right) = O\left(\frac{k \cdot m}{n \cdot \beta \cdot \delta^2 \cdot \lambda^3}\right)$. □

Corollary 30 *For fully dynamic graph G there is a deterministic k batch-dynamic algorithm for maintaining a $(3/2 + \epsilon, \delta)$ -approximate maximum matching with update time $\tilde{O}\left(\frac{k \cdot m}{n \cdot \beta} + k \cdot \beta\right)$.*

Proof Set $\lambda = \frac{\epsilon}{128}$ and β large enough to satisfy the requirements of Lemma 24 such that the resulting sparsifier is $(3/2 + \epsilon/4, \delta)$ -approximate. Use the algorithm of Lemma 29. The resulting damaged-EDCS sparsifier will have maximum degree $O(\beta)$, undergo $\tilde{O}(k)$ recourse per update and will take $\tilde{O}\left(\frac{k \cdot m}{n \cdot \beta}\right)$ time to maintain. By Lemma 24 it will be a $(3/2 + \epsilon/4, \delta)$ -approximate matching sparsifier. Hence, if we apply the algorithm of Lemma 4 to maintain a $(1 + \epsilon/4)$ -approximate maximum matching within the sparsifier we can maintain a $(3/2 + \epsilon, \delta)$ approximate matching in $\tilde{O}\left(\frac{m \cdot k}{n \cdot \beta} + \beta \cdot k\right)$ update time and recourse. □

5.6 Proof of Theorem 2

Proof Take the algorithm of Corollary 30. Set $k = \log(n)$ and apply Corollary 9 to receive a deterministic $(3/2 + \epsilon, \delta)$ -approximate dynamic matching algorithm with worst-case update time $\tilde{O}\left(\frac{m}{n \cdot \beta} + \beta\right)$. Finally, transform this algorithm into a $(3/2 + \epsilon)$ -approximate matching algorithm using either Corollary 19 or Corollary 20. \square

6 $(2 + \epsilon)$ -Approximate Fully Dynamic Maximum Matching in $\tilde{O}(1)$ Worst-Case Update Time

In the appendix we present a deterministic worst-case $O(\text{poly}(\log(n), 1/\epsilon))$ -update time $(2 + \epsilon)$ -approximate fully dynamic matching algorithm. Currently, the only deterministic $O(\text{poly}(\log(n), 1/\epsilon))$ -update time algorithms [27, 32] have amortized update time bounds, while the fastest worst-case algorithm runs in $\tilde{O}(\sqrt{n})$ update time from [28]. We will first improve the running time bounds of the algorithm presented in [32] to k batch-dynamic using the same technique as presented previously. [32] similarly bases the algorithm on amortized rebuilds which are triggered when ϵ factor of change occurs within the data-structure. In order to improve the update time to batch-dynamic we define $\epsilon_i = \frac{i \cdot \epsilon}{k}$ to be the slack parameter during the progressing of batch i . Firstly, this ensures that $\epsilon_i \leq \epsilon$ during any of the batches progressed guaranteeing the approximation ratio. Secondly, whenever a new batch begins the slack parameter increases by $\frac{\epsilon}{k}$ which insures that there will be enough time steps before next rebuild occurs to amortize the rebuild time over.

Lemma 31 *There is a deterministic k batch amortized $O_\epsilon(k \cdot \log^4(n))$ update time $(2 + \epsilon)$ -approximate fully dynamic matching algorithm.*

Proof The proof of this lemma is only available [online](#) due to length requirements. \square

Theorem 1 follows from Lemma 31 and Corollary 9.

7 Open Questions

Worst-Case Update Time Improvement Through Batch-Dynamization: We have shown two applications on how batch-dynamization can be used to improve amortized rebuild based algorithm update times to worst-case. As amortized rebuild is a popular method for dynamizing a data-structure not just in the context of matching it would be interesting to see if the batch-dynamization based framework has any more applications.

(α, δ) -Approximate Dynamic Matching: In current dynamic matching literature most algorithms focus on maintaining an α -approximate matching or matching sparsifier both for the integral and fractional version of the problem. However, a more relaxed (α, δ) -approximate matching algorithm using the reductions presented in this paper (or [42, 43]) allow for the general assumption that $\mu(G) = \Theta(n)$ at all times. This assumption has proven to be useful in other settings for the matching problem

such as the stochastic setting ([42, 43]) but largely seems to be unexplored in the dynamic setting.

Damaged EDCS: The EDCS matching sparsifier [35] has found use in a number of different settings for the matching problem [14, 28, 37–41]. In contrast with the EDCS (as far as we are aware) a damaged EDCS admits a deterministic near-linear time static algorithm. This might lead to new results in related settings.

Acknowledgements We would like to thank Sayan Bhattacharya and Thatchaphol Saranurak for helpful discussions. Further we would like to thank Thatchaphol Saranurak for suggesting to use lossless expanders to deterministically generate ϵ -matching preserving partitionings.

Author Contributions PK is the sole author of the paper. Acknowledged researchers SB and TS supported the development of the paper through pointing at relevant literature and advising on some aspects of the presentation.

Data Availability Data sharing not applicable to this article as no datasets were generated or analysed during the current study.

Declarations

Conflict of interest The authors have no conflicts of interest to declare. All co-authors have seen and agree with the contents of the manuscript and there is no financial interest to report. We certify that the submission is original work and is not under review at any other publication.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

A Proof of Lemma 8

We restate the lemma for the readers convenience:

Lemma 32 *Given (α) approximate (or (α, ϵ) -approximate) dynamic matching algorithm Alg is $O(\log(n))$ batch-dynamic with update time $O(T(n))$ and dynamic graph G undergoing edge insertions and deletions. There is an algorithm Alg' which maintains $O(\log(n))$ matchings of G such that at all times during processing an input sequence of arbitrarily large polynomial length one of the matchings is (α) approximate (or (α, ϵ) -approximate). The update time of Alg' is worst-case $O(T(n) \cdot \log^3(n))$ and it is deterministic if Alg is deterministic.*

Proof Fix some integer $k = O(\log(n))$. Alg' will be running k instances of Alg in parallel on graph G , call them $A_i : i \in \{0, \dots, k - 1\}$. Assume that Alg' 's running time is k batch-dynamic. We will describe what state each instance of Alg will take during processing specific parts of the input, then argue that at least one of them will be outputting an α (or (α, δ))-approximate matching at all times.

Assume that the input sequence I is k^k long. Let $I[i]$ represent the i -th element of the input sequence and $I[i, j]$ represents elements $i, i + 1, \dots, j - 1$ for $j > i$. Let $I[i, j]$ represent $I[i, j] \cup I[j]$. Fix a specific instance of Alg say A_i . Call the input batches of A_i as $B_i^j : j \in [k]$. At a given point in time let $|B_i^j|$ refer to the number of input elements instance A_i has progressed as it's j -th batch. Note that we will assume that in update time $O(T(n) \cdot |B_i^j|)$ instance A_i can revert back to a state where input batch B_i^j was empty given the elements of B_i^j where the last elements of I progressed by A_i .

Represent the input elements of I as k -long k -airy strings starting from $\{0\}^k$. Choose $I[\lambda]$ such that λ -s k -airy representation ends with an ' i ' followed by $\gamma > 0$ ' 0 '-s and contains a single i digit. We will now describe what instance A_i will be doing while Alg' is processing input elements $I[\lambda, \lambda + k^\gamma]$. We will call this process as the *resetting of batches* $B_i^\lambda, \dots, B_i^1$.

Resetting the Contents of Batches $B_i^\lambda, \dots, B_i^1$:

With a slight overload of notation partition the input sub-sequence $I[\lambda, \lambda + k^\gamma]$ into $\gamma + 1$ sub-sequences $I_j : j \in \{0, \dots, \gamma\}$. Let $\lambda_j = \lambda + \sum_{x=j}^{\gamma-1} k^x \cdot (k - 1)$ for $\gamma \geq j \geq 0$. Let $I_j = I[\lambda_j, \lambda_{j-1})$ for $\gamma \geq j > 0$ and $I_0 = I[\lambda_0]$. Observe that $|I_j| = \Theta(k^j)$.

- While Alg' is processing input elements I_λ instance A_i will revert to the state it was in before processing the contents of the batches $B_i^{\gamma+1}, \dots, B_i^1$. Then it proceeds to place all these elements into batch $B_i^{\gamma+1}$ a single batch.
- While Alg' is processing input elements $I_j : \gamma > j > 0$ instance A_i will progress input elements I_{j+1} as batch B_i^{j+1} .
- While Alg' is processing the input element I_0 instance A_i will place input elements $I_1 \cup I_0$ into B_i^1 .

If A_i is not resetting batches it is processing single elements of the input string.

Processing Single Elements of The Input String:

If the first $k - 1$ digits of the k -airy representation of λ don't contain a single i digit then while Alg is processing $I[\lambda]$ instance A_i will extend it's last batch B_i^1 with input element $I[\lambda]$.

These two instances describe the behaviour of A_i over the whole of I . If A_i is processing a single input element at any point in time it's output is an α (or (α, δ))-approximate matching. Also observe, that for any λ there is a digit $i \in [k]$ in it's k -airy representation which is not one of it's first $k - 1$ digits. By definition, this implies that A_i will be processing $I[\lambda]$ as a single input element. Hence, the output of A_i will be an (α) (or (α, δ))-approximate matching for some i at all time steps.

Claim 33 At all times for all $j \in [k]$ and $i \in \{0, \dots, k - 1\}$ it holds that $|B_i^j| \leq (j + 1) \cdot k^j$.

Proof We will proof the claim through induction on j . Fix i . Whenever the contents of B_i^1 are reset it will be set to contain exactly k input elements. If the contents of B_i^1 are not reset while $I[\lambda]$ is progressed by Alg' then B_i^1 is extended by $I[\lambda]$. However, over the course of k consecutive input elements being progressed by Alg' batch B_i^1 must be reset. Therefore, B_i^1 will never contain more than $2 \cdot k - 1$ elements.

Assume that $|B_i^j| \leq (j + 1) \cdot (k^j - k^{j-1})$ at all times as an inductive hypothesis. Consider how many elements may B_i^{j+1} contain. Whenever B_i^{j+1} is reset it will be set to contain exactly $(k - 1) \cdot k^j$ elements. Furthermore, whenever B_i^j is reset B_i^{j+1} is extended by the contents of B_i^j, \dots, B_i^1 . These are the only cases when B_i^{j+1} may be extended by any input elements. B_i^j is reset at most $k - 1$ times between two resets of B_i^{j+1} . Therefore, at all times $|B_i^{j+1}| \leq (k - 1) \cdot (k^j + \sum_{x=1}^j (x + 1) \cdot (k^x - k^{x-1})) \leq (k - 1) \cdot (j + 2) \cdot k^j = (j + 2) \cdot (k^{j+1} - k^j)$. This finishes the inductive argument. \square

Claim 34 The worst-case running time of A_i is $O(T(n) \cdot k^2)$ for all $i \in \{0, \dots, k - 1\}$.

Proof To bound worst case running times differentiate two cases. Firstly, if $I[\lambda]$ is progressed as a single input element by A_i then A_i will extend its smallest batch B_i^1 with $I[\lambda]$. As at all times $|B_i^1| \leq 2 \cdot k$ due to Claim 33 this can be done in worst-case update time $O(T(n) \cdot k)$.

Fix λ as described previously, such that its k -airy representation contains a single i digit followed by $\gamma > 0$ 0-s so that A_i will be resetting batches B_i^γ, \dots, B_i^1 while Alg' is processing $I[\lambda, \lambda + k^\gamma]$. Define $I_j : \gamma \geq j \geq 0$ as before. While Alg' is processing I_γ instance A_i has to revert to the state before processing any of $B_i^{\gamma+1}, \dots, B_i^1$ and progress their contents as a single batch into $B_i^{\gamma+1}$. This concerns the backtracking and processing of $O(k^{\gamma+1} \cdot \gamma)$ input elements by Claim 33. The computational work required to complete this can be distributed over the time period Alg' is handling I_γ evenly as this computation doesn't require A_i to know the contents of I_γ . Hence, it can be completed in $O(T(n) \cdot k \cdot \gamma) = O(T(n) \cdot k^2)$ worst-case update time.

Similarly, over the course of Alg' processing I_j which consists of $\Theta(k^j)$ elements we can distribute the $O(T(n) \cdot k^{j+1})$ total work of processing I_{j+1} into batch B_i^{j+1} evenly resulting in $O(T(n) \cdot k)$ worst case update time. Finally, for instance A_i processing $I_1 \cup I_0$ while Alg' progresses I_0 will take $O(T(n) \cdot k)$ time. \square

Therefore, each instance A_i runs in $O(T(n) \cdot k^2)$ worst-case update time. As there are k instances of Alg running like as described in parallel, this takes a total of $O(T(n) \cdot k^2)$ worst case update time. It remains to select $k = O(\log(n))$ so the algorithm can progress an input of length $O(\log^{\log(n)}(n)) = O(n^{\log(\log(n))})$, that is of input sequences of arbitrarily large polynomial length for large enough n . \square

B Missing Proofs of Section 5

B.1 Proof of Lemma 25

Proof Let E'_{fin} represent the state of E' at termination. First let's argue that $\forall e \in E'_{fin} : deg_{E'_{fin}} \leq \beta$. At the end of the penultimate iteration of the outer loop H must have maximum edge degree of $\beta \cdot (1 - \lambda/4)$. H then will be extended with edges of $E' \setminus E_D$ which has a maximum degree of $\beta \cdot \lambda/8$. Therefore, $\max_{e \in E'_{fin}} deg_{E'_{fin}}(e) \leq \beta \cdot (1 - \lambda/4) + 2 \cdot \lambda/8 \leq \beta$.

As $\sum_{v \in V} \text{deg}_{E'_{fin}}(v) \leq \frac{\delta \cdot \lambda \cdot \beta \cdot n}{8}$ it must hold that $|V_D| \leq \delta \cdot n = |V| \cdot \delta$. Take an edge $e \in E \setminus H_{fin}$ which doesn't intersect V_D . As all such edges with lower than $\beta \cdot (1 - \lambda/2)$ edge degree in E were added to E'_{fin} it must hold that $\text{deg}_{H_{fin} \cup E'_{fin}}(e) \geq \beta \cdot (1 - \lambda/2)$. As neither endpoints of e are in V_D it must hold that $\text{deg}_{E_D}(e) \leq \lambda \cdot \beta/4$. This implies that $\text{deg}_{H_{fin}}(e) \geq \text{deg}_{H_{fin} \cup E'_{fin}}(e) - \text{deg}_{E_D}(e) \geq \beta \cdot (1 - \lambda/2) - \lambda \cdot \beta/4 \geq \beta \cdot (1 - \lambda)$. Hence, H_{fin} is a (β, λ, δ) -damaged EDCS of G . \square

B.2 Proof of Lemma 26

Proof Observe that every iteration of the repeat loop runs in $O(m)$ time as each iteration can be executed over a constant number of passes over the edge set. Define $\Phi(H) = \Phi_1(H) - \Phi_2(H)$ where $\Phi_1(H) = \sum_{v \in V} \text{deg}_H(v) \cdot (\beta - 1/2) = |E(H)| \cdot (2 \cdot \beta - 1)$ and $\Phi_2(H) = \sum_{e \in H} \text{deg}_H(e)$. Initially $\Phi(H) = 0$ and $\Phi(H) \leq \beta^2 \cdot n$. We will show that $\Phi(H)$ monotonously increases over the run of the algorithm and each iteration of the repeat loop (except for the last one) increases it by at least $\Omega(\beta^2 \cdot \lambda^2 \cdot \delta \cdot n)$ which implies the lemma.

$\Phi(H)$ may change at times when edges are added to or removed from H . Whenever e is removed from H we know that $\text{deg}_H(e) > \beta \cdot (1 - \lambda/4)$ (before the deletion). This means that $\Phi_1(H)$ decreases by $2\beta \cdot (1 - \lambda/4) - 1$ but $\Phi_2(H)$ also decreases by at least $2 \cdot \beta \cdot (1 - \lambda/4)$. This is because $\text{deg}_H(e)$ disappears from the sum of $\Phi_2(H)$ and $\text{deg}_H(e) - 2$ elements of the sum (degrees of edges neighbouring e) reduce by 1 and $\text{deg}_H(e) \geq \beta \cdot (1 - \lambda/4) + 1$. Hence, $\Phi(H)$ increases by at least 1.

Whenever an edge e is added to H we know that $\text{deg}_H(e) < \beta \cdot (1 - \lambda/2)$ (before the insertion). Due to the insertion $\Phi_1(H)$ increases by exactly $2 \cdot \beta - 1$. $\Phi_2(H)$ increases by at most $2 \cdot \beta \cdot (1 - \lambda/2)$ as a term of at most $\beta \cdot (1 - \lambda/2) + 1$ is added to it's sum and at most $\beta \cdot (1 - \lambda/2) - 1$ elements of it's sum increase by 1. Therefore, $\Phi(H)$ increases by at least $\lambda \cdot \beta$. In every iteration but the last one of the repeat loop at least $\frac{\lambda \cdot \beta \cdot \delta \cdot n}{16}$ edges were added to H . This means every iteration increases $\Phi(H)$ by at least $\frac{\lambda^2 \cdot \beta^2 \cdot \delta \cdot n}{16} = \Omega(\lambda^2 \cdot \beta^2 \cdot \delta \cdot n)$ finishing the lemma. \square

B.3 Proof of Lemma 27

Proof Every time H is reset rebuilt through StaticDamagedEDCS the lemma statement is satisfied (by Lemma 25) Focus on one period of α updates after a rebuild. Define E_D and E_I to be the set of edges deleted and inserted over these updates respectively (note that $E_D \cap E_I$ might not be empty). Define $V_{E_D} = \{v \in V | \text{deg}_{E_D}(v) \geq \frac{\beta \cdot \lambda}{16}\}$ and $V_{E_I} = \{v \in V | \text{deg}_{E_I}(v) \geq \frac{\beta \cdot \lambda}{16}\}$. Note, that $|V_{E_D} \cup V_{E_I}| \leq \frac{2 \cdot \alpha}{\frac{\beta \cdot \lambda}{16}} \leq \frac{\delta \cdot n}{2}$.

As after a call to Algorithm 1 the sparsifier H is a $(\frac{\beta}{1+\lambda/4}, \lambda/4, \delta/2)$ -damaged EDCS (follows from Lemma 25) and the following holds for some $V_D \subseteq V$ with $|V_D| \leq |V| \cdot \delta/2$:

- $\forall e \in H : \text{deg}_H(e) \leq \frac{\beta}{1+\lambda/4}$
- All $e \in E \setminus H$ such that $e \cap V_D = \emptyset$ satisfies that $\text{deg}_H(e) \geq \frac{\beta \cdot (1-\lambda/4)}{1+\lambda/4}$

Define $V'_D = V_D \cup V_{E_D} \cup V_{E_I}$. Note that $|V'_D| \leq |V_D| + |V_{E_D} \cup V_{E_I}| \leq n \cdot \delta$. Also note that after a rebuild $\max_{e \in E} \deg_H(e) \leq \frac{\beta}{1+\lambda/4}$. As edges will only be inserted between vertices u and v if their degrees is at most $\frac{\beta \cdot \lambda}{16} - 2$ in H we can be certain that at any point $\max_{e \in E} \deg_H(e) \leq \frac{\beta}{1+\lambda/4} + \frac{\beta \cdot \lambda}{16} \leq \beta$ (for small enough values of λ).

At any point during the phase take an arbitrary $e \in E \setminus H \wedge e \cap V'_D = \emptyset$. If $e \in E_I$ at the time of it's (last) insertion one of its endpoints, say v had $\deg_{E_I}(v) \geq \frac{\lambda \cdot \beta}{16}$ or $\deg_H(e) > \beta - 2$. The former would imply $v \in V'_D$. Therefore, we can assume that if $e \notin H$ then either $e \in E_I$ and at time of it's insertion $\deg_H(e) > \beta - 2$ or $e \notin E_I$ and at the start of the phase $\deg_H(e) \geq \beta \frac{(1-\lambda/4)}{1+\lambda/4}$. Either way, during the phase the edge degree of e may have reduced by at most $\frac{\beta \cdot \lambda}{8}$ as none of it's endpoints are in V_{E_D} . Therefore, $\deg_H(E) \geq \frac{\beta \cdot (1-\lambda/4)}{1+\lambda/4} - \frac{\beta \cdot \lambda}{8} \geq \beta \cdot (1 - \lambda)$. This concludes the proof. \square

B.4 Proof of Lemma 28

Proof Edge insertions and deletions are handled in $O(1)$ time apart from the periodic rebuilds. The rebuilds run in $O\left(\frac{m}{\delta \cdot \lambda^2}\right)$ deterministic time by Lemma 26 therefore over α insertions the amortized update time is $O\left(\frac{m}{\delta \cdot \lambda^2 \cdot \alpha}\right) = O\left(\frac{m}{n \cdot \beta \cdot \lambda^3 \cdot \delta^2}\right)$. The recourse of the sparsifier is also constant apart from rebuild operations. When a rebuild occurs the sparsifier goes under at most $O(n \cdot \beta)$ edge updates. Therefore, the amortized recourse is $O\left(\frac{n \cdot \beta}{\alpha}\right) = O\left(\frac{1}{\lambda \cdot \delta}\right)$. \square

C Missing Proofs from Section 4

C.1 Proof of Lemma 13

Claim 35 Algorithm 3 maintains an $(\alpha + \epsilon)$ -approximate maximum matching.

Proof Fix $i = \lfloor \log_{1+\epsilon/(8\alpha)}(\mu(G)) \rfloor$ and let $\mu_{1+\epsilon/(8\alpha)}(G) = (1+\epsilon/(8\alpha))^i$. Note that G contains a matching of size $\mu_{1+\epsilon/(8\alpha)}(G)$ (assume integrality for sake of convenience) and $\mu(G) \leq \mu_{1+\epsilon/(8\alpha)}(G) \cdot (1 + \epsilon/(8\alpha))$. By the definition of matching preserving vertex partitionings there is a $j \in [L]$ such that $\mu(G_j^i) \geq (1 - \epsilon/(8\alpha))\mu_{1+\epsilon/8}(G)$.

Hence, $\mu(G_j^i) \geq \mu(G) \cdot (1 - \epsilon/(4\alpha))$. As the vertex set of G_j^i is of size $C \cdot \mu_{1+\epsilon/8}(G)$ we have that $|M_j^i| \cdot \alpha + \frac{\epsilon}{8C} \cdot C \cdot \mu_{1+\epsilon/(8\alpha)}(G) \geq \mu(G_j^i) \geq \mu(G) \cdot (1 - \epsilon/(4\alpha))$ as M_j^i is an $(\alpha, \frac{\epsilon}{8C})$ -approximate maximum matching of G_j^i . This simplified states that $|M_j^i| \cdot \frac{\alpha}{1 - \frac{3\epsilon}{\alpha \cdot 8}} \geq \mu(G)$.

As $M_j^i \subseteq E'$ we have $|M^*| \cdot (1 + \epsilon/(8\alpha)) \geq |M_j^i|$ and therefore $|M^*| \cdot \frac{\alpha \cdot (1 + \frac{\epsilon}{8\alpha})}{1 - \frac{\epsilon \cdot 3}{8\alpha}} \geq \mu(G)$. This can be simplified to $|M^*| \cdot (\alpha + \epsilon) \geq \mu(G)$. \square

Claim 36 Algorithm 3 has an update time of $O(T(n, \epsilon/C) \cdot \frac{L^2 \cdot \log^2(n)}{\epsilon^4})$.

Algorithm 3: Vertex Sparsification

Input: $G = (V, E), Alg_M, Alg_S$
Output: $(\alpha + \epsilon)$ -approximate maximum matching of G

- 1 $MMSize \leftarrow 1$
- 2 $i \leftarrow 1$
- 3 **while** $MMSize \leq n$ **do**
- 4 $i \leftarrow i + 1$
- 5 $MMSize \leftarrow MMSize \cdot (1 + \epsilon/(8\alpha))$
- 6 $\mathcal{V}^i : V_j^i : j \in [L] \leftarrow$ Set of $(MMSize, \epsilon/(8\alpha))$ matching preserving vertex partitionings of G
 into $C \cdot MMSize$ vertex subsets of size L output by Alg_S
- 7 $G_j^i \leftarrow$ Vertex concatenation of G based on V_j^i
- 8 $M_j^i \leftarrow$ Maintain $(\alpha, \frac{\epsilon}{8C})$ -approximate matching of G_j^i with Alg_M
- 9 $G_{M_j^i} \leftarrow$ Maintain edges of M_j^i in G
- 10 $E' \leftarrow \bigcup_{i,j} G_{M_j^i}$
- 11 $M^* \leftarrow$ Maintain a $1 + \epsilon/(8\alpha)$ -approximate maximum matching of (V, E') with Lemma 4

Proof The maintenance of M_j^i will take $O(T(n, \epsilon/C))$ update time for specific values of i, j . As $\alpha = O(1)$ i will range in $\left[O\left(\frac{\log(n)}{\epsilon}\right)\right]$. Therefore, the algorithm maintains $O\left(\frac{L \cdot \log(n)}{\epsilon}\right)$ matchings in parallel using Alg_M . This means E' has maximum degree $O\left(\frac{L \cdot \log(n)}{\epsilon}\right)$ and can be maintained in update time $O\left(T(n, \epsilon/C) \cdot \frac{L \cdot \log(n)}{\epsilon}\right)$ and may undergo the same amount of recourse. Hence, with the invocation of the algorithm from Lemma 4 the total update time is $O\left(T(n, \epsilon/C) \cdot \frac{L^2 \cdot \log^2}{\epsilon^4}\right)$.

The two claims conclude Lemma 13 □

Do note, that the update time can be slightly improved to $O\left(T\left(n, \frac{\epsilon}{C}\right) \cdot \frac{L \cdot \log(n)}{\epsilon} + \frac{L^2 \cdot \log^2(n)}{\epsilon^5}\right)$ using Lemma 5 ([44]). The update time of the sparsifier is $O\left(T\left(n, \frac{\epsilon}{C}\right) \cdot \frac{L \cdot \log(n)}{\epsilon}\right)$. Using the lemma it's recourse can be bounded at $O\left(\frac{L \cdot \log(n)}{\epsilon}\right)$. Applying Lemma 4 ([31]) yields the slightly different update time.

C.2 Proof of Lemma 14

Proof For graph $G = (V, E)$ generate $L = \left\lceil \frac{512 \cdot \log(n)}{\epsilon^2} \right\rceil$ vertex partitionings into $d = \left\lceil 4 \cdot \frac{(2k)}{\epsilon} \right\rceil$ sets at random. Call the set of partitionings $\mathcal{V} = \{\mathcal{V}^j : j \in [L]\}$ and let V_i^j stand for the i -th vertex set of the j -th partitioning. Fix $2k$ vertices S of V arbitrarily to represent the endpoints of a matching of size k in G and note that this can be done $\binom{n}{2k} \leq n^{2 \cdot k} \leq e^{\ln(n) \cdot 2 \cdot k} \leq e^{4 \cdot \log_2(n) \cdot k}$ number of ways.

Fix a specific vertex partitioning \mathcal{V}^j with vertex sets $V_i^j : i \in [d]$. Let the random variable $X_i^j : i \in [d]$ be an indicator variable of $S \cap V_i^j \neq \emptyset$ and $\bar{X}^j = \sum_{i \in [d]} X_i^j$.

Claim 37 $X_i^j : i \in [d]$ are negatively associated random variables.

Proof Define $B_i^l : i \in [d], l \in [2 \cdot k]$ be the indicator variable of the l -th vertex of S falling into the i -th subset V_i^j . This turns the random variables into the well known balls and bins experiment. By [45] (this can also be considered a folklore fact) random variables $B_i^l : i \in [d], l \in [2 \cdot k]$ are negatively associated. By definition $X_i^j = \max_{l \in [2 \cdot k]} \{B_i^l\}$. By Theorem 2 of [48] monotonously increasing functions defined on disjoint subsets of a set of negatively associated random variables are negatively associated. As max is monotonously increasing this implies that $X_i^j : i \in [d]$ are also negatively associated. \square

$$\begin{aligned} \mathbb{E}[X_i^j] &= 1 - \Pr[S \cap V_i^j = \emptyset] = 1 - \left(1 - \frac{1}{d}\right)^{2k} \geq 1 - \left(\left(1 - \frac{1}{8 \cdot k/\epsilon}\right)^{8 \cdot k/\epsilon}\right)^{\epsilon/4} \\ &\geq 1 - e^{-\epsilon/4} \geq \frac{\epsilon \cdot (1 - \epsilon/8)}{4} \end{aligned}$$

Therefore, $\mathbb{E}[\bar{X}^j] \geq \left\lceil \frac{8 \cdot k}{\epsilon} \right\rceil \cdot \frac{\epsilon \cdot (1 - \epsilon/8)}{4} \geq 2k \cdot (1 - \epsilon/8)$. Now we apply Chernoff’s inequality for negatively associated random variables to get that:

$$\begin{aligned} \Pr[\bar{X}^j \leq 2k \cdot (1 - \epsilon/4)] &\leq \Pr[\bar{X}^j \leq \mathbb{E}[\bar{X}^j] \cdot (1 - \epsilon/8)] \leq \exp\left(-\frac{\mathbb{E}[\bar{X}^j] \cdot (\frac{\epsilon}{8})^2}{2}\right) \\ &\leq e^{-\frac{2k \cdot \epsilon^2}{128}} \end{aligned}$$

This implies that

$$\Pr\left[\min_{j \in [L]} \{\bar{X}^j\} \leq 2k \cdot (1 - \epsilon/4)\right] \leq e^{-4 \cdot \log(n) \cdot 2k}$$

Further applying a union bound over the $\binom{n}{2k}$ possible choices of S yields that regardless of the choice of S with probability $1 - e^{-2 \cdot \log(n) \cdot 2k} \geq 1 - 1/\text{poly}(n)$ there is a partitioning $\mathcal{V}^j = V_i^j : i \in [d]$ where at least $2k \cdot (1 - \epsilon/4)$ of the vertex sets of \mathcal{V}^j contain a vertex of S . This implies that there can be at most $2k \cdot (1 - \epsilon/2)$ vertices of S sharing a vertex set of \mathcal{V}^j with an other vertex of S . Furthermore, if S represents the endpoints of a matching of size k at least $k \cdot (1 - \epsilon)$ of it’s edges will have both their endpoints being assigned to unique vertex sets of \mathcal{V}^j with respect to S . This implies that the concatenation of G based on \mathcal{V}^j will preserve a $1 - \epsilon$ fraction of any matching of size k from G . Therefore, \mathcal{V} is a (k, ϵ) matching preserving set of partitionings for G .

Note that while we can simply sample the partitionings randomly in polynomial time, we could also consider all possible sets of partitionings and check weather any of them is (k, ϵ) matching preserving for all possible choice of $S \subseteq V$. From the fact that a random sampling based approach succeeds with positive probability we know

that there is a set of (k, ϵ) matching preserving partitionings therefore we will find one one eventually deterministically. \square

References

1. Abboud, A., Addanki, R., Grandoni, F., Panigrahi, D., Saha, B.: Dynamic set cover: improved algorithms and lower bounds. In: Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing, pp. 114–125 (2019)
2. Behnezhad, S., Derakhshan, M., Hajiaghayi, M., Stein, C., Sudan, M.: Fully dynamic maximal independent set with polylogarithmic update time. In: 2019 IEEE 60th Annual Symposium on Foundations of Computer Science (FOCS), pp. 382–405. IEEE (2019)
3. Bernstein, A., Forster, S., Henzinger, M.: A deamortization approach for dynamic spanner and dynamic maximal matching. In: Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms, pp. 1899–1918. SIAM (2019)
4. Bernstein, A., Gutenberg, M.P., Saranurak, T.: Deterministic decremental reachability, scc, and shortest paths via directed expanders and congestion balancing. In: 2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS), pp. 1123–1134. IEEE (2020)
5. Bhattacharya, S., Chakrabarty, D., Henzinger, M.: Deterministic fully dynamic approximate vertex cover and fractional matching in $o(1)$ amortized update time. In: International Conference on Integer Programming and Combinatorial Optimization, pp. 86–98. Springer (2017)
6. Gupta, A., Krishnaswamy, R., Kumar, A., Panigrahi, D.: Online and dynamic algorithms for set cover. In: Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, pp. 537–550 (2017)
7. Abboud, A., Williams, V.V.: Popular conjectures imply strong lower bounds for dynamic problems. In: 2014 IEEE 55th Annual Symposium on Foundations of Computer Science, pp. 434–443. IEEE (2014)
8. Charikar, M., Solomon, S.: Fully dynamic almost-maximal matching: breaking the polynomial barrier for worst-case time bounds. In: ICALP (2017)
9. Grandoni, F., Leonardi, S., Sankowski, P., Schwiegelshohn, C., Solomon, S.: $(1 + \epsilon)$ -approximate incremental matching in constant deterministic amortized time. In: Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms, pp. 1886–1898. SIAM (2019)
10. Gupta, M.: Maintaining approximate maximum matching in an incremental bipartite graph in polylogarithmic update time. In: 34th International Conference on Foundation of Software Technology and Theoretical Computer Science (FSTTCS 2014). Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik (2014)
11. Neiman, O., Solomon, S.: Simple deterministic algorithms for fully dynamic maximal matching. *ACM Trans Algorithms (TALG)* **12**(1), 1–15 (2015)
12. Bhattacharya, S., Henzinger, M., Italiano, G.F.: Deterministic fully dynamic data structures for vertex cover and matching. *SIAM J. Comput.* **47**(3), 859–887 (2018)
13. Bhattacharya, S., Henzinger, M., Nanongkai, D.: Fully dynamic approximate maximum matching and minimum vertex cover in $o(\log^3(n))$ worst case update time. In: Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms, pp. 470–489. SIAM (2017)
14. Bernstein, A., Stein, C.: Faster fully dynamic matchings with small approximation ratios. In: Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms, pp. 692–711. SIAM (2016)
15. Solomon, S.: Fully dynamic maximal matching in constant update time. In: 2016 IEEE 57th Annual Symposium on Foundations of Computer Science (FOCS), pp. 325–334. IEEE (2016)
16. Sankowski, P.: Faster dynamic matchings and vertex connectivity. In: Proceedings of the Eighteenth Annual ACM-SIAM Symposium on Discrete Algorithms, pp. 118–126 (2007)
17. Arar, M., Chechik, S., Cohen, S., Stein, C., Wajc, D.: Dynamic matching: reducing integral algorithms to approximately-maximal fractional algorithms. In: 45th International Colloquium on Automata, Languages, and Programming, ICALP 2018, 9–13 July, 2018, Prague, Czech Republic, pp. 7:1–7:16 (2018)
18. Wajc, D.: Rounding dynamic matchings against an adaptive adversary. In: Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing, STOC 2020, Chicago, IL, USA, 22–26 June, 2020, pp. 194–207. ACM (2020)

19. Baswana, S., Gupta, M., Sen, S.: Fully dynamic maximal matching in $o(\log(n))$ update time. *SIAM J. Comput.* **44**(1), 88–113 (2015)
20. Behnezhad, S., Lacki, J., Mirrokni, V.: Fully dynamic matching: Beating 2-approximation in δ^ϵ update time. In: Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms, pp. 2492–2508. SIAM (2020)
21. Henzinger, M.R., King, V.: Randomized dynamic graph algorithms with polylogarithmic time per operation. In: Proceedings of the Twenty-Seventh Annual ACM Symposium on Theory of Computing, pp. 519–527 (1995)
22. Henzinger, M., Krinninger, S., Nanongkai, D., Saranurak, T.: Unifying and strengthening hardness for dynamic problems via the online matrix-vector multiplication conjecture. In: Proceedings of the Forty-Seventh Annual ACM Symposium on Theory of Computing, pp. 21–30 (2015)
23. Amir, A., Pettie, S., Porat, E.: Higher lower bounds from the 3sum conjecture. In: Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms, pp. 1272–1287. SIAM (2016)
24. Larsen, K.G.: The cell probe complexity of dynamic range counting. In: Proceedings of the Forty-Fourth Annual ACM Symposium on Theory of Computing, pp. 85–94 (2012)
25. Patrascu, M.: Towards polynomial lower bounds for dynamic problems. In: Proceedings of the Forty-Second ACM Symposium on Theory of Computing, pp. 603–610 (2010)
26. Onak, K., Rubinfeld, R.: Maintaining a large matching and a small vertex cover. In: Proceedings of the Forty-Second ACM Symposium on Theory of Computing, pp. 457–464 (2010)
27. Bhattacharya, S., Henzinger, M., Nanongkai, D.: New deterministic approximation algorithms for fully dynamic matching. In: Proceedings of the Forty-Eighth Annual ACM Symposium on Theory of Computing, pp. 398–411 (2016)
28. Roghani, M., Saberi, A., Wajc, D.: Beating the folklore algorithm for dynamic matching (2021). arXiv preprint [arXiv:2106.10321](https://arxiv.org/abs/2106.10321)
29. Holm, J., De Lichtenberg, K., Thorup, M.: Poly-logarithmic deterministic fully-dynamic algorithms for connectivity, minimum spanning tree, 2-edge, and biconnectivity. *J. ACM (JACM)* **48**(4), 723–760 (2001)
30. Kapron, B.M., King, V., Mountjoy, B.: Dynamic graph connectivity in polylogarithmic worst case time. In: Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete algorithms, pp. 1131–1142. SIAM (2013)
31. Gupta, M., Peng, R.: Fully dynamic $1 + \epsilon$ -approximate matchings. In: 2013 IEEE 54th Annual Symposium on Foundations of Computer Science, pp. 548–557. IEEE (2013)
32. Kiss, P., Bhattacharya, S.: Deterministic rounding of dynamic fractional matchings. In: 48th International Colloquium on Automata, Languages, and Programming, ICALP 2021, 12–16 July, 2021, Glasgow (2021)
33. Bernstein, A., Brand, J.V.D., Gutenberg, M.P., Nanongkai, D., Saranurak, T., Sidford, A., Sun, H.: Fully-dynamic graph sparsifiers against an adaptive adversary. *CoRR*. [arXiv:2004.08432](https://arxiv.org/abs/2004.08432) (2020)
34. Nanongkai, D., Saranurak, T.: Dynamic spanning forest with worst-case update time: adaptive, las vegas, and $o(n^{1/2-\epsilon})$ -time. In: Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, pp. 1122–1129 (2017)
35. Bernstein, A., Stein, C.: Fully dynamic matching in bipartite graphs. In: International Colloquium on Automata, Languages, and Programming, pp. 167–179. Springer, Berlin (2015)
36. Grandoni, F., Schwiegelshohn, C., Solomon, S., Uzdad, A.: Maintaining an eds in general graphs: simpler, density-sensitive and with worst-case time bounds (2021). [arXiv:2108.08825](https://arxiv.org/abs/2108.08825)
37. Assadi, S., Bateni, M., Bernstein, A., Mirrokni, V., Stein, C.: Coresets meet eds: algorithms for matching and vertex cover on massive graphs. In: Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms, pp. 1616–1635. SIAM (2019)
38. Assadi, S., Behnezhad, S.: Towards a unified theory of sparsification for matching problems. In: 2nd Symposium on Simplicity in Algorithms, SOSA 2019, 8–9 Jan, 2019, San Diego, CA, USA (2019)
39. Bernstein, A.: Improved bounds for matching in random-order streams. In: 47th International Colloquium on Automata, Languages, and Programming, ICALP 2020, 8–11 July, 2020, Saarbrücken, Germany (Virtual Conference), pp. 12:1–12:13 (2020)
40. Bernstein, A., Dudeja, A., Langle, Z.: A framework for dynamic matching in weighted graphs. In: STOC'21: 53rd Annual ACM SIGACT Symposium on Theory of Computing, Virtual Event, Italy, 21–15 June, 2021, pp. 668–681. ACM (2021)

41. Assadi, S., Behnezhad, S.: Beating two-thirds for random-order streaming matching. In: 48th International Colloquium on Automata, Languages, and Programming, ICALP 2021, 12–16 July, 2021, Glasgow, Scotland (Virtual Conference), pp. 19:1–19:13 (2021)
42. Behnezhad, S., Derakhshan, M., Hajiaghayi, M.: Stochastic matching with few queries: $(1 - \epsilon)$ approximation. In: Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing, pp. 1111–1124 (2020)
43. Assadi, S., Khanna, S., Li, Y.: The stochastic matching problem with (very) few queries. *ACM Trans. Econ. Comput. (TEAC)* **7**(3), 1–19 (2019)
44. Solomon, N., Solomon, S.: A generalized matching reconfiguration problem. In: 12th Innovations in Theoretical Computer Science Conference, ITCS 2021, 6–8 Jan, 2021, Virtual Conference, LIPIcs (2021)
45. Dubhashi, D.P., Ranjan, D.: Balls and bins: a study in negative dependence. *BRICS Rep. Ser.* **3**(25) (1996)
46. Capalbo, M., Reingold, O., Vadhan, S., Wigderson, A.: Randomness conductors and constant-degree lossless expanders. In: Proceedings of the Thiry-Fourth Annual ACM Symposium on Theory of Computing, pp. 659–668 (2002)
47. Berinde, R., Gilbert, A.C., Indyk, P., Karloff, H., Strauss, M.J.: Combining geometry and combinatorics: a unified approach to sparse signal recovery. In: 2008 46th Annual Allerton Conference on Communication, Control, and Computing, pp. 798–805. IEEE (2008)
48. Gerasimov, M., Kruglov, V., Volodin, A.: On negatively associated random variables. *Lobachevskii J. Math.* **33**(1), 47–55 (2012)
49. Assadi, S., Khanna, S., Li, Y., Yaroslavtsev, G.: Maximum matchings in dynamic graph streams and the simultaneous communication model. In: Krauthgamer R. (ed.), Proceedings of the Twenty-Seventh Annual (ACM-SIAM) Symposium on Discrete Algorithms, SODA 2016, 10–12 Jan, Arlington, VA, USA, pp. 1345–1364. SIAM (2016)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.