

Simple Algorithms for Fully Dynamic Edge Connectivity

Yotam Kenneth-Mordoch Robert Krauthgamer*

Weizmann Institute of Science

{yotam.kenneth,robert.krauthgamer}@weizmann.ac.il

October 19, 2025

Abstract. In the fully dynamic edge connectivity problem, the input is a simple graph G undergoing edge insertions and deletions, and the goal is to maintain its edge connectivity, denoted λ_G . We present two simple randomized algorithms solving this problem. The first algorithm maintains the edge connectivity in worst-case update time $\tilde{O}(n)$ per edge update, matching the known bound but with simpler analysis. Our second algorithm achieves worst-case update time $\tilde{O}(n/\lambda_G)$ and worst-case query time $\tilde{O}(n^2/\lambda_G^2)$, which is the first algorithm with worst-case update and query time $o(n)$ for large edge connectivity, namely, $\lambda_G = \omega(\sqrt{n})$.

1 Introduction Finding the edge connectivity of a graph is a fundamental algorithmic problem in graph theory. The *edge connectivity* of a graph $G = (V, E)$, denoted λ_G , is defined as the minimum number of edges that need to be removed from G in order to disconnect it; this is equivalent to the value of a global minimum cut in G . It has long been known that it is possible to find the edge connectivity of a graph in $\tilde{O}(m)$ time, where $n = |V|, m = |E|$ and $\tilde{O}(\cdot)$ hides polylogarithmic factors in n [Kar00, KT19, HRW20, MN20, GMW20].

We consider the problem of maintaining the edge connectivity of a graph $G = (V, E)$ in a fully dynamic setting, where the input is a *dynamic graph* presented as a fixed set of vertices and a sequence of edge insertions and deletions, and the goal is to report the edge connectivity after each update. This problem has a rich history: it was originally studied for constant values of edge connectivity (namely $\lambda_G \in \{1, 2, 3, 4\}$) see for example [EGIN97, HdLT01, KKM13, CGL⁺20]; for the case $\lambda_G = 1$ there exists a randomized algorithm using worst-case update time $\text{polylog}(n)$ and a deterministic algorithm using amortized update time $\text{polylog}(n)$ [HdLT01, KKM13]. These results are complemented by a deterministic algorithm for the case $\lambda_G = \log^{o(1)} n$, which achieves worst-case update time $n^{o(1)}$ [JST24].

For the case of general λ_G , there exist deterministic algorithms that maintain the edge connectivity in worst-case time $\tilde{O}(\min(\lambda_{\max}^{5.5} \sqrt{n}, m^{1-1/12}, m^{11/13} n^{1/13}, n^{3/2}))$ per update, where λ_{\max} is the maximum edge connectivity of G throughout the algorithm's execution [Tho07, GHN⁺23, dVC25]. Allowing for randomization, there exists a randomized algorithm that maintains the edge connectivity in worst-case time $\tilde{O}(n)$ per update [GHN⁺23]. Finally, for approximate maintenance which asks to maintain a $(1 + \epsilon)$ -approximation of the edge connectivity [TK00, Tho07, EHL25], one can achieve a $(1 + o(1))$ -approximation with worst-case update time $n^{o(1)}$ [EHL25].

1.1 Contributions We focus on general edge connectivity, and present two new randomized algorithms that achieve worst-case update time $\tilde{O}(n)$ and $\tilde{O}(n^2/\lambda_G^2)$ respectively. Both algorithms are based on the same high-level idea, which is an efficient and simple implementation of a contraction algorithm that preserves the edge connectivity. Our first algorithm matches the result of [GHN⁺23] but has a simpler and more intuitive proof, while our second algorithm offers improved performance for dense graphs. Throughout, update time (of an algorithm) refers to the worst-case time per edge update unless stated otherwise. We say that a randomized algorithm succeeds with high probability if its probability of success is at least $1 - 1/n^c$ for any constant $c > 0$.

*The Harry Weinrebe Professorial Chair of Computer Science. Work partially supported by the Israel Science Foundation grant #1336/23, by the Israeli Council for Higher Education (CHE) via the Weizmann Data Science Research Center, and by a research grant from the Estate of Harry Schutzman.

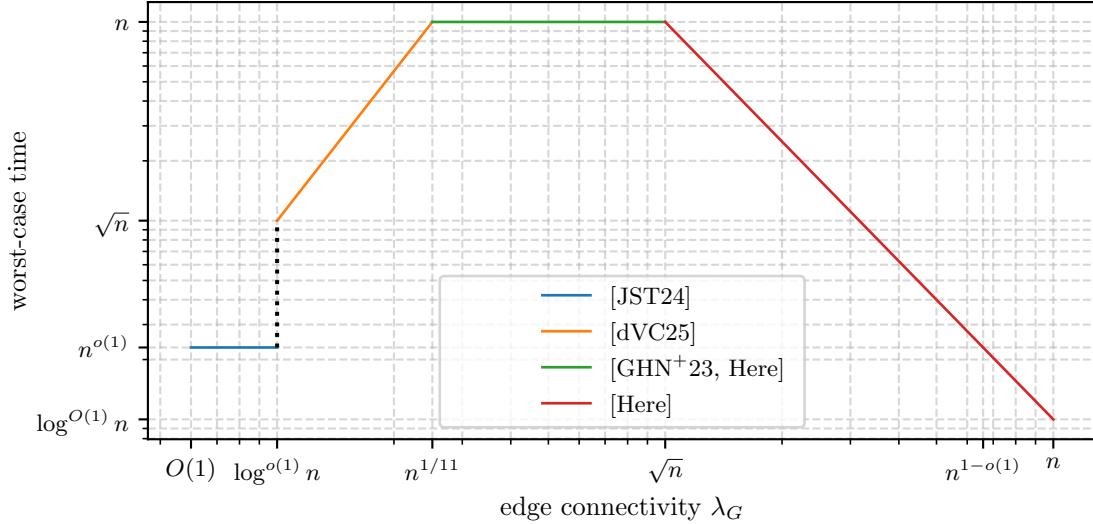


Figure 1.1: Schematic comparison of existing algorithms for dynamic edge connectivity.

THEOREM 1.1. *There exists a fully dynamic algorithm that, given an unweighted dynamic graph G on n vertices, maintains the edge connectivity of G with update time $\tilde{O}(n)$. The algorithm is randomized and succeeds with high probability.*

Our second algorithm improves on the result of [GHN⁺23] for large edge connectivity, namely when $\lambda_G = \tilde{\omega}(\sqrt{n})$. Our result is actually stronger, since it is parametrized by the minimum degree of the graph at the time of the update, δ_G , which is only larger than the edge connectivity λ_G and hence it implies an update time of $\tilde{O}(n^2/\lambda_G^2)$. Therefore, Theorem 1.2 improves on [GHN⁺23] for large edge connectivity, i.e. $\lambda_G \geq \tilde{\omega}(\sqrt{n})$, and is in fact the first algorithm to achieve worst-case update time $o(n)$ for large edge connectivity. Moreover, the improvement applies also to graphs with large minimum degree but small edge connectivity. The improvement is even larger when separating the update and query times, i.e. if the algorithm only needs to report the edge connectivity occasionally.

THEOREM 1.2. *There exists a fully dynamic algorithm that, given an unweighted dynamic graph G on n vertices, maintains the edge connectivity of G with update time $\tilde{O}(n/\delta_G)$ and query time $\tilde{O}(n^2/\delta_G^2)$. The algorithm is randomized and succeeds with high probability.*

We now provide several details about the above results. The first one is that both results can report not only the exact edge connectivity, but also the edges that form a minimum cut, which incurs an additional $\tilde{O}(\lambda_G)$ time per query. When reporting only the value of the edge connectivity, the algorithms are adversarially-robust, i.e. they work against an adaptive adversary. This is because the edge connectivity is a deterministic function of the graph and hence, as long as the algorithm returns the correct result, reporting it does not leak any information to the adversary. In contrast, when reporting the edges of the minimum cut, our algorithms only work against an oblivious adversary. Since in both cases the adversary lacks knowledge of the algorithms' internal state, our analysis only considers the case of an oblivious adversary.

The second detail is that it is possible to combine the results, to obtain an algorithm with update time $\tilde{O}(\min\{n, n^2/\delta_G^2\})$. A comparison of existing results with our work is given in Figure 1.1. As explained above we use the fact that the update time of Theorem 1.2 can be stated as $\tilde{O}(n^2/\lambda_G^2)$.

1.2 Future Work The main open question remaining is to achieve worst-case update time $o(n)$ for all edge connectivity values. Prior work achieves this for edge connectivity $\lambda_G = o(n^{1/11}/\text{polylog } n)$ [Tho07, JST24, dVC25]. Our results complement these results and achieve it for large edge connectivity, namely $\lambda_G = \tilde{\omega}(\sqrt{n})$. Perhaps surprisingly, the remaining case is that of intermediate connectivity, $\lambda_G \in [n^{1/11}, n^{1/2}]$.

2 Proof of Main Theorems A powerful tool for finding a minimum cut that has been used extensively in past work is a family of contraction procedures first introduced by Kawarabayashi and Thorup [KT19], which we call *KT-style contractions* [KT19, HRW20, GNT20, Sar21, AEG⁺22, KK25]. In general, KT-style contractions preserve the minimum cut value of an input graph G while reducing the number of vertices to $\tilde{O}(n/\delta_G)$, where δ_G is the minimum degree in G . The edge connectivity can then be more efficiently computed on a smaller contracted graph $G' = (V', E')$. This general approach has also been employed to design an efficient algorithm for dynamic edge connectivity [GHN⁺23]. Specifically, the algorithm leverages the 2-out contraction technique of [GNT20] to maintain a contracted multigraph G' of the input graph G [GHN⁺23], and constructs a maximal δ_G -packing of forests in G' .

DEFINITION 2.1 (Maximal k -Packing of Forests). *Given a (multi)graph G , a maximal k -packing of forests is a set of edge-disjoint forests T_1, \dots, T_k of G that are maximal, i.e. for every $i \in [k]$ and $e \in E \setminus \cup_j T_j$, the edge set $T_i \cup \{e\}$ contains a cycle. A maximal k -packing of forests in a multigraph includes at most one parallel edge in each forest.*

The main idea of [GHN⁺23] is that the edge connectivity of G is at most δ_G , and therefore if a global minimum cut of G is preserved in G' , then all its edges are contained in a maximal δ_G -packing of forests in the multigraph G' . Therefore, the edge connectivity of G' (and hence of G) is equal to the edge connectivity of $H = (V', T_1 \cup \dots \cup T_{\delta_G})$, where $\{T_i\}_{i=1}^{\delta_G}$ is a maximal δ_G -packing of forests in G' . Furthermore, the number of edges in H is at most $\tilde{O}(n)$ since each forest T_i has at most $O(|V'|) = \tilde{O}(n/\delta_G)$ edges, and hence the edge connectivity of H can be computed in time $\tilde{O}(n)$ using some static algorithm for edge connectivity.

From a simplicity perspective, the main drawback of [GHN⁺23] is that it does not maintain an explicit contraction of the graph, but rather only its connected components. It then recovers the forest packing using linear-sketching techniques, which complicates the algorithm. Another drawback of this approach is that linear sketching only recovers unweighted edges of the contracted multigraph G' . However, by reinterpreting the contracted multigraph as a weighted graph (merging parallel edges), we can take advantage of the fact that when $\delta_G \geq \omega(\sqrt{n})$ the graph G' becomes very small.¹ In fact, the number of edges in G' is $\tilde{O}(n^2/\delta_G^2)$, which enables solving the edge connectivity problem in $\tilde{o}(n)$ worst-case time per edge update when $\delta_G \geq \omega(\sqrt{n})$.

Our algorithms follow the general recipe of [GHN⁺23], with the main difference being that our contraction algorithm maintains an explicit KT-style contraction as a weighted graph.² This both simplifies the algorithm and improves the update time whenever $\delta_G \geq \omega(\sqrt{n})$. Specifically, we leverage the τ -star contraction procedure of [AEG⁺22, KK25], which was originally designed for finding a global minimum cut in the cut-query model. The procedure works as follows. Sample a set of *center vertices* $R \subseteq V$ by including each vertex $v \in V$ independently with probability $p = O(\log n/\tau)$. Let $H = \{v \in V \setminus R \mid d_G(v) \geq \tau\}$ be the set of non-center vertices with degree at least τ , where throughout $d_G(v)$ is the degree of v in G . Then, for every $v \in H$, uniformly sample a vertex $r \in N_G(v) \cap R$, where $N_G(v)$ is the neighborhood of v in G , and contract the corresponding edge (v, r) , keeping parallel edges, which yields a contracted multigraph G' . The main guarantee of this procedure is that if G has a minimum cut that is non-trivial, i.e., not composed of a single vertex, then $\lambda_G = \lambda_{G'}$ with constant probability. An illustration of the τ -star contraction is given in Figure 2. The following theorem states this formally.

THEOREM 2.2 ([AEG⁺22, KK25]). *Let $G = (V, E)$ be an unweighted graph on n vertices with at least one minimum cut that is non-trivial. Then, τ -star-contraction with $p = 800 \log n/\tau$ yields a contracted graph G' such that,*

1. *if $\tau \leq \delta_G$, then with probability at least $1 - 1/n^4$ all vertices in $V \setminus R$ are contracted and G' has at most $\tilde{O}(n/\tau)$ vertices, and*
2. *$\lambda_{G'} \geq \lambda_G$, and equality holds with probability at least $2 \cdot 3^{-13}$.*

The following technical lemma provides the guarantees of the τ -star contraction procedure in a fully dynamic setting. Using this lemma our algorithms maintain a KT-style contraction of the graph G as a weighted graph (by merging parallel edges) in a fully dynamic setting. We call a contraction G' of G *complete*, if it maps every

¹Such a reinterpretation is natural and was used before in the context of KT-style contractions, e.g. [KK25].

²To recover the edges of a minimum cut we can maintain a list of the edges of G that are mapped to each edge of G' and report the edges corresponding to the minimum cut in G' .

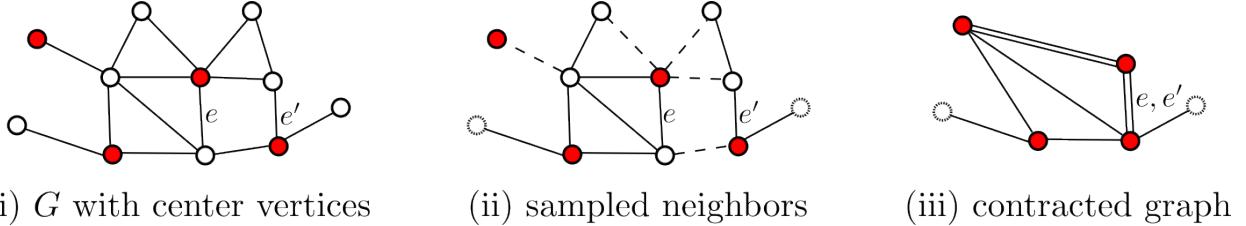


Figure 2.1: An illustration of the τ -star contraction procedure with $\tau = 2$. (i) A random set of center vertices R is sampled (indicated by red). (ii) Each non-center vertex in $H = \{v \in V \setminus R \mid d_G(v) \geq \tau\}$ chooses a random neighbor in R (indicated by dashed edges); vertices not in $H \cup R$ are indicated by a striped pattern. (iii) The corresponding edges are contracted (keeping parallel edges) to obtain a contracted graph G' .

edge of G to an edge in G' , and otherwise it is *incomplete*.³ Our contraction procedure only tracks edges that are internal to the center vertices R , hence if a vertex $v \in V \setminus R$ is not contracted, then its incident edges are not mapped in G' . Therefore, its contraction is complete only when all vertices in $V \setminus R$ are contracted. Fortunately, this happens with high probability whenever $\delta_G \geq \tau$ by Theorem 2.2.

LEMMA 2.3. *There is a fully-dynamic randomized algorithm that, given as input a dynamic graph G and a parameter $\tau > 0$, maintains a (possibly incomplete) contraction G' of G as a weighted graph. If $\delta_G \geq \tau$ then G' is complete and constitutes a τ -star contraction of G with high probability. The algorithm runs in amortized update time $\text{polylog}(n)$ and worst-case update time $\tilde{O}(n)$ (regardless of δ_G). Furthermore, the algorithm updates at most $O(1)$ edges in G' per edge update in G in expectation and $\tilde{O}(n/\tau)$ in the worst-case.*

We prove this lemma in Section 3 and proceed now to show that it implies our two main theorems.

2.1 Proof of Theorem 1.1 We begin by showing how to maintain a maximal k -packing of forests in a weighted graph under edge insertions and deletions. Note that since we interpret the contracted multigraph G' as a weighted graph we need the packing to maintain edge weights.⁴

LEMMA 2.4. *There exists a fully dynamic algorithm that maintains a maximal k -packing of forests of a dynamic weighted graph $G = (V, E)$ on n vertices with update time $\tilde{O}(k)$. The algorithm is randomized and succeeds with high probability against an oblivious adversary.*

To prove the lemma we follow the approach of [ADK⁺16], which requires the following result of [KKM13, GKKT15] on maintaining a single spanning forest.

THEOREM 2.5 ([KKM13, GKKT15]). *There exists a fully dynamic algorithm that maintains a weighted spanning forest T of a dynamic weighted graph G with n vertices, with worst-case update time $O(\log^4 n)$. Every edge insertion in G adds at most one edge to T . Every edge deletion in G removes at most one edge from T and may additionally add one edge to T . The algorithm is randomized and succeeds with high probability against an oblivious adversary.*

Proof of Theorem 2.4. The algorithm maintains a collection of k spanning trees $\{T_i\}_{i=1}^k$, where each T_i is spanning tree of the graph $G \setminus \cup_{j < i} T_j$, that is maintained using Theorem 2.5. To show that the update time is $\tilde{O}(k)$, examine some edge update e and level i . If the update is an insertion, then either e is added to T_i or the update is propagated to level $i + 1$, and in both cases at most one change is propagated to level $i + 1$. If the update is a deletion, then either $e \notin T_i$ and this update is propagated to level $i + 1$, or $e \in T_i$ and then removing it might cause a single other edge e' to be added to T_i from $G \setminus T_{j \leq i}$. In this case e' is from $G \setminus T_{j \leq i}$, and therefore we propagate a deletion of e' to level $i + 1$ (however, T_{i+1} is oblivious to the deletion of e). Therefore, an edge update in G causes at most $O(1)$ changes to each T_i and the overall update time is $\tilde{O}(k)$. \square

Proof of Theorem 1.1. It suffices to present an algorithm that returns a global minimum cut of G with constant probability, and returns some larger cut otherwise (if the algorithm is only required to return the edge

³Notice that an incomplete contraction is not a contraction in the strict sense, since it does not map every edge of G to an edge in G' .

⁴One can simulate changing the weight of an edge by deleting it and inserting it with the new weight.

connectivity then it only returns the value of this cut). Running $O(\log n)$ independent copies of the algorithm and taking the minimum cut found yields a global minimum cut with high probability. The desired algorithm works as follows. Begin by running in parallel $r = O(\log n)$ instances of the τ -star contraction algorithm of Theorem 2.3 to obtain contracted graphs $\{G'_i = (V'_i, E'_i)\}_{i=1}^r$, where the i -th instance uses $\tau = 2^i$. Furthermore, for each instance G'_i maintain a maximal 2^{i+1} -forest packing $\{T_j^i\}_{j=1}^{2^{i+1}}$ of the contracted graph G'_i using Theorem 2.4.⁵ In parallel, track the minimum degree δ_G of the input graph G .

Let i^* be the maximal such that $\delta_G \geq 2^{i^*}$, and let $H = (V'_{i^*}, \bigcup_{j=1}^{2^{i^*+1}} T_j^{i^*})$ be the graph obtained by taking the union of the maximal 2^{i^*+1} -forest packing of G'_{i^*} . The algorithm then finds a minimum cut of H using an offline algorithm for weighted graphs in time $\tilde{O}(n)$ and, returns the minimum between the edge connectivity found and δ_G (and the associated cut). The correctness guarantee follows from Theorem 2.2 and since the contraction is complete by Theorem 2.3 (and hence G'_i is a 2^i -star contraction of G) whenever $\delta_G \geq \tau$.

To analyze the time complexity, notice that maintaining each τ -star-contraction instance takes worst-case time $\tilde{O}(n)$ per update by Theorem 2.3. To analyze the time complexity of maintaining a maximal 2^{i+1} -forest packing, notice that each edge update in G updates at most $\tilde{O}(n/2^i)$ edges in G'_i by Theorem 2.3. Each update in G'_i then takes $\tilde{O}(2^{i+1})$ time to update the maximal 2^{i+1} -forest packing by Theorem 2.4, hence the total time per edge update in G is $\tilde{O}(n)$. Observe that the number of edges in each 2^{i+1} -forest packing is at most $\tilde{O}(n)$, since each forest has $O(|V'_i|) = \tilde{O}(n/2^i)$ edges and the packing has 2^{i+1} forests, and hence a minimum cut can be found in time $\tilde{O}(n)$. Finally, observe that tracking the minimum degree δ_G in G can be done using a heap data structure with worst-case update time $O(\log n)$ per edge update. This concludes the proof of Theorem 1.1. \square

2.2 Proof of Theorem 1.2 At a high level, the proof of Theorem 1.2 is similar to that of Theorem 1.1. Again the algorithm maintains in parallel $r = O(\log n)$ instances of the τ -star contraction algorithm of Theorem 2.3, where the i -th instance G'_i uses $\tau = 2^i$. However, instead of maintaining a maximal 2^{i+1} -forest packing of G'_i , the algorithm solves the edge-connectivity problem on G'_i directly.

Notice that using Theorem 2.3, the worst-case time per update in G is $\tilde{O}(n)$, which is larger than $\tilde{O}(n^2/\delta_G^2)$ when $\delta_G = \omega(n^{1/2})$. To circumvent this, we de-amortize Theorem 2.3; instead of processing immediately all $d_G(v)$ edges incident to v when its chosen center vertex changes, the contraction procedure places these edges in a queue and processes $\tilde{O}(n/\delta_G)$ of them after each update in G . This ensures that the worst-case time complexity per update in G is $\tilde{O}(n/\delta_G)$. However, now the contraction G' may be incomplete even when $\delta_G \geq \tau$, if the queue of the contraction procedure is non-empty. Fortunately, we can show that the probability that the queue is non-empty at any given moment in time is $O(1/\log^2 n)$ as stated next and proven in Section 3. Notice that the contraction procedure always knows whether its contraction is complete, by checking whether its queue is empty and if $\delta_G \geq \tau$, hence the algorithm can easily ignore incomplete instances.

LEMMA 2.6. *There is a fully-dynamic randomized algorithm that, given as input a dynamic graph G and a parameter $\tau > 0$, returns a (possibly incomplete) contraction G' of G as a weighted graph. The algorithm runs in amortized update time $\text{polylog}(n)$ and worst-case time $\tilde{O}(n)$ (regardless of δ_G). Finally, if $\delta_G \geq \tau$ then G' is complete and constitutes a τ -star contraction of G with probability $1 - O(1/\log^2 n)$.*

Proof of Theorem 1.2. Again it suffices to show an algorithm that succeeds with constant probability and returns a larger cut otherwise. The algorithm is as follows. Begin by running $r = O(\log n)$ independent instances of the τ -star contraction algorithm of Theorem 2.6 to obtain contracted graphs $\{G'_i\}_{i=1}^r$, where the i -th instance uses $\tau = 2^i$. In parallel, track the minimum degree δ_G of the input graph G . Let i^* be the maximal such that $\delta_G \geq 2^{i^*}$.

If the contraction G'_{i^*} is incomplete, then fail (or return an arbitrary cut). Otherwise, find a minimum cut using an offline algorithm for edge connectivity in weighted graphs in time $\tilde{O}(n^2/\delta_G^2)$. Finally, return the minimum between the edge connectivity of G'_{i^*} and the minimum degree δ_G (and the associated cut). Notice that by Theorem 2.6 if G'_{i^*} is a 2^{i^*} -star contraction, the edge connectivity of G'_{i^*} is equal to the edge connectivity of G with constant probability. Furthermore, G'_{i^*} is complete (and hence a τ -star contraction) with probability at least $1 - O(1/\log^2 n)$ by Theorem 2.6. Therefore, using a union bound over the two events the algorithm finds the edge connectivity with constant probability.

To analyze the time complexity, notice that maintaining a single τ -star contraction instance using Theorem 2.6

⁵An additional factor 2 in the size of the packing is needed since $\delta_G/\tau \in [1, 2)$.

takes amortized time $\text{polylog}(n)$ and worst-case time $\tilde{O}(n/\delta_G)$ per edge update in G . Maintaining the minimum degree δ_G in G can be done using a heap data structure with worst-case update time $O(\log n)$ per edge update. Therefore, the time per edge update in G is $\tilde{O}(n/\delta_G)$. Finally, the query time is determined by the edge-connectivity computation in G'_{i^*} which takes $\tilde{O}(n^2/2^{2i^*}) = \tilde{O}(n^2/\delta_G^2)$ time. \square

3 Dynamic Star Contraction In this section we present a fully dynamic algorithm that maintains a τ -star contraction of the graph G , proving Theorems 2.3 and 2.6. The algorithm for both results is the same, except for minor changes in the edge-update procedure which we detail in the proof of Theorem 2.6, and presented in Algorithm 1. To maintain the τ -star contraction, we use a stable dynamic uniform sampler. The construction of this sampler uses standard techniques and is described in Theorem 3.2, whose proof is provided in Subsection 3.2.

DEFINITION 3.1 (Stable Dynamic Uniform Sampling). *Given a dynamic set of elements (undergoing insertions and deletions) denoted S_t at time t , a stable uniform dynamic sampler maintains a sample $x_t \in S_t$ (or reports \emptyset if $S_t = \emptyset$) such that:*

1. *At every time t , each $x \in S_t$ is sampled with probability exactly $1/|S_t|$.*
2. *The probability that the sampled element changes after an update is at most $1/|S_t|$.*

LEMMA 3.2. *There exists a data structure for stable dynamic uniform sampling that uses $O(\log n)$ update time per operation under an oblivious adversary. The data structure is randomized and succeeds with high probability.*

Throughout this section denote each edge update by the pair $\{(u, v), s\}$, where $s \in \{\pm 1\}$ indicates whether the edge (u, v) is inserted or deleted. Note that Algorithm 1 excludes from G' all the vertices in $V \setminus R$ that are not contracted to any vertex in R . The proof of Theorem 2.3 follows immediately from the following two lemmas, which prove the correctness and time complexity of Algorithm 1 respectively.

Algorithm 3.1 Dynamic Star Contraction

```

1: Input: An unweighted multigraph  $G = (V, E)$ , parameter  $\tau > 0$ 
2: Output: A contracted graph  $G'$ 
3: procedure DYNAMIC-STAR-CONTRACTION( $G, \tau$ )
4:    $R \leftarrow$  sample every  $v \in V$  with probability  $p = 800 \log n / \tau$ 
5:   for each  $v \in V \setminus R$  let  $R_v \leftarrow$  initialize stable uniform dynamic sampler  $\triangleright$  the vertex  $r \in R$  to which  $v$  is
   contracted
6:   for each  $v \in R$  let  $R_v \leftarrow v$   $\triangleright$  by convention, since  $v \in R$  is never contracted
7:    $G' \leftarrow (R, \emptyset)$ 
8:   for each edge update  $\{(u, v), s\}$  do
9:     if  $s = 1$  then  $E \leftarrow E \cup \{(u, v)\}$  else  $E \leftarrow E \setminus \{(u, v)\}$ 
10:    if  $R_u, R_v$  are both not null then  $w(R_u, R_v) \leftarrow w(R_u, R_v) + s$ 
11:    if  $u, v \in R$  or  $u, v \notin R$  then continue
12:    if  $u \notin R$  then swap  $u, v$   $\triangleright |\{u, v\} \cap R| = 1$  by line 11
13:    add/remove  $u$  to/from  $R_v$   $\triangleright$  update the stable uniform dynamic sampler for  $v$ 
14:    if  $R_v$  has changed then
15:       $R_v^{\text{old}} \leftarrow$  old  $R_v$ ,  $R_v^{\text{new}} \leftarrow$  new  $R_v$ 
16:      for each  $w \in R$  do
17:         $E_{vw} \leftarrow \{x \in V \setminus R \mid (v, x) \in E, R_x = w\}$ 
18:        if  $R_v^{\text{old}}$  is not null then  $w(R_v^{\text{old}}, w) \leftarrow w(R_v^{\text{old}}, w) - |E_{vw}|$ 
19:        if  $R_v^{\text{new}}$  is not null then  $w(R_v^{\text{new}}, w) \leftarrow w(R_v^{\text{new}}, w) + |E_{vw}|$ 

```

LEMMA 3.3 (Correctness). *Whenever $\delta_G \geq \tau$, the output of Algorithm 1 with parameter τ is a complete contraction and constitutes a τ -star contraction of a given dynamic graph G .*

Proof. We begin by showing that if every edge of G is mapped to an edge in G' , then the output is a τ -star contraction. Notice that the vertices in R are sampled uniformly at random with probability $p = 800 \log n / \tau$. By the properties of the stable uniform dynamic sampling, we have that for every vertex $v \in V \setminus R$, the vertex R_v is sampled uniformly at random from the neighbors of v in R . Therefore, once every edge of G is mapped to

an edge in G' , the output is a τ -star contraction. We now show that the output of the algorithm is complete. Whenever $\tau \geq \delta_G$ by Theorem 2.2, we have that every vertex $v \in V \setminus R$ is contracted with high probability. In addition, notice that once every vertex $v \in V \setminus R$ is contracted, all edges of G have been mapped to edges in G' . Therefore the output of the algorithm is complete. \square

LEMMA 3.4 (Time Complexity). *Given an unweighted graph $G = (V, E)$ on n vertices, Algorithm 1 with parameter τ requires amortized time $\tilde{O}(1)$ (and worst case time $\tilde{O}(n)$) per edge. Furthermore, in expectation the algorithm updates at most $O(1)$ edges in G' per edge update in G and in the worst-case it updates $\tilde{O}(n/\tau)$ edges in G' per edge update in G .*

3.1 Amortized Time Complexity

The proof of Theorem 3.4 relies on the following claim.

CLAIM 3.5. *For each edge update $\{(u, v), s\}$, with probability at least $1 - 800 \log n / \max\{d_G(v), \tau\}$ both R_v, R_u remain the unchanged.*

Proof of Claim 3.5. To bound the probability R_v changes, we may assume without loss of generality that $v \notin R$ (since if $u, v \in R$ then R_v, R_u do not change). Notice that under this assumption, R_u does not change, since $v \notin R$ and so R_u is not updated. Observe that R_v can change only if $u \in R$, hence

$$(3.1) \quad \Pr[R_v \text{ changes}] = \Pr[u \in R] \cdot \Pr[R_v \text{ changes} \mid u \in R] = p \cdot \Pr[R_v \text{ changes} \mid u \in R],$$

where the last equality is since R is obtained by sampling each vertex independently with probability $p = 800 \log n / \tau$. Hence the probability that R_v changes is at most $800 \log n / \tau$. We now show a stronger bound on the probability that R_v changes when $d_G(v) \geq \tau$. Notice that $\Pr[R_v \text{ changes} \mid u \in R] = 1/d_R(v)$ by Theorem 3.2, where $d_R(v)$ is the number of neighbors of v in R . To bound $d_R(v)$, notice that $\mathbb{E}_R[d_R(v) \mid u \in R] \geq \mathbb{E}_R[d_R(v) \in R] = pd_G(v)$, where the inequality is since conditioning on the event that $u \in R$ increases the expected number of neighbors of v in R and the equality is since R is sampled uniformly at random from V . Therefore, by Chernoff's bound (Theorem A.1),

$$\Pr[d_R(v) \leq pd_G(v)/2] \leq 2 \exp\left(-\frac{pd_G(v)}{12}\right) \leq 2 \exp\left(-\frac{800 \log n}{12}\right) \leq n^{-10},$$

where the second inequality is from $d_G(v) \geq \tau$. Hence, with probability at least $1 - n^{-10}$ we have that $d_R(v) \geq pd_G(v)/2$. Plugging this back into Equation (1),

$$\Pr[R_v \text{ changes}] = \Pr[u \in R] \cdot \Pr[R_v \text{ changes} \mid u \in R] \leq p \cdot \left(\frac{2}{pd_G(v)} + n^{-10}\right) \leq \frac{800 \log n}{d_G(v)}.$$

Therefore, the probability that both R_v, R_u remain unchanged is at least $1 - 800 \log n / \max\{d_G(v), \tau\}$, as claimed. \square

Proof of Theorem 3.4. Observe that if R_v, R_u do not change, then updating the weight of the edge (R_u, R_v) in line 10 takes $O(1)$ time and updating the samplers for u, v in line 13 takes $O(\log n)$ time by Theorem 3.2. They cannot both change since that implies that both $u, v \notin R$, so assume exactly one of them changes and without loss of generality it is R_v . In this case, in line 16 the algorithm updates in G' the weight of all edges corresponding to edges incident to v in G .

To minimize the number of times we update the edges of G' , we go over all edges incident to v in G to find the size of the sets E_{vw} , for every $w \in R$.⁶ This requires $O(d_G(v))$ time, since we need to iterate over all edges incident to v . Then, we update the weights of the edges (R_v^{old}, w) and (R_v^{new}, w) in G' in $O(1)$ time for every $w \in R$, which requires $O(\min\{|R|, d_G(v)\}) \leq O(d_G(v))$ time overall, since at most $d_G(v)$ sets E_{uw} can be non-empty. Therefore, by Claim 3.5, the expected time spent on this operation is $O(\log n)$. In the worst case each update requires at most $O(d_G(v))$ time, which is $O(n)$ per edge.

To bound the number of changed edges in G' , notice that if R_v does not change then at most one edge is updated in G' . Otherwise, the number of edges updated in G' is at most $\max\{2d_G(v), \tilde{O}(n/\tau)\}$, since each vertex in G' has at most $\tilde{O}(n/\tau)$ neighbors by the size of R . Combining the above with Claim 3.5 we find that the

⁶Observe that it is simpler to iterate over all the edges incident to v in G and update the weights of the edges in G' during each iteration. However, this creates $d_G(v)$ edge updates in G' which then takes $O(d_G(v)n/\delta_G) = \omega(n)$ time (if $d_G(v) > \delta_G$) to update the tree packing in G' . In contrast, our approach only updates the edges of G' at most $\tilde{O}(n/\tau)$ times per update in G .

expected number of edges updated in G' per edge update in G is $O(1)$. By the same bound presented above, the worst-case number of edges whose weight is updated in line 16 per edge update in G is $\tilde{O}(n/\tau)$. \square

3.2 From Amortized to Worst-Case Time Complexity via Incremental Rebuilding In this section we prove Theorem 2.6, bounding the worst-case time complexity of Algorithm 1 with $(n \log^4 n / \delta_G(v))$ -incremental rebuilding. Notice that the incremental rebuilding completes the update of the star contraction G' after $\delta_G / \log^4 n$ edge updates in G . Therefore, the probability that the queue of the algorithm is not empty (and hence the contraction is incomplete) at any given time t is bounded by the probability that some vertex changes its representative R_v in the last $\delta_G / \log^3 n$ edge updates. The following lemma shows that this probability is $O(1 / \log^2 n)$ whenever $\delta_G \geq \log^4 n$. Otherwise, if $\delta_G < \log^4 n$, then the algorithm finishes processing its queue after a single edge update since $n \log^4 n / \delta_G \geq n$, and its output is complete (as long as $\delta_G \geq \tau$).

CLAIM 3.6. *If at a given time t , the minimum degree of G is $\delta > \log^4 n$, then given a subsequent sequence of updates of length $N := O(\delta / \log^3 n)$ the probability that at least one vertex $v \in V$ changes its center vertex R_v during the sequence is at most $O(1 / \log^2 n)$.*

Proof. Note that the minimum degree of G after the sequence of updates is at least $(1 - o(1)) \cdot \delta \geq \delta/2$. Let I_i be the indicator for the event that there exists $v \in V \setminus R$ such that R_v changes during the i -th edge update. Notice that $\mathbb{E}[I_i] \leq 800 \log n / (\delta/2)$ by Claim 3.5 and that $\delta_G \geq \delta/2$ throughout the sequence of updates. Therefore, $\mathbb{E}[\sum_{i=1}^N I_i] \leq 1600N \log n / \max\{\delta/2, \tau\} \leq O(1 / \log^2 n)$. To conclude, by Markov's inequality we have that

$$\Pr\left[\sum_{i=1}^N I_i \geq 1\right] \leq \frac{\mathbb{E}\left[\sum_{i=1}^N I_i\right]}{1} \leq O\left(\frac{1}{\log^2 n}\right).$$

3.3 Stable Dynamic Uniform Sampling In this section we prove Theorem 3.2, it is based on a standard techniques (see e.g. [Lub86, BDM02, BNR02, DLT07]) and we include it for completeness.

Proof of Theorem 3.2. As each element comes, assign it a priority by uniformly sampling an integer in $[1, n^{10}]$, if at any time two elements have the same priority, the algorithm fails. Then, maintain a heap of all the items in the sequence, ordered by their priority. When an element is deleted, remove it from the heap. The minimum element of the heap is the sampled item. It is known that a dynamic heap can be maintained in $O(\log n)$ worst-case time per operation, where n is the number of elements in the heap.

We now prove the two properties of the sampling. Notice that the minimum priority element is distributed uniformly, and hence each element has a probability of $1/|S_t|$ to be sampled at time t . To show the stability, notice that under a deletion the probability that the element deleted is the minimum element is $1/|S_t|$. If we are given an insertion, notice that again by symmetry we have that the probability that the new element is sampled is $1/|S_t|$. Therefore, we obtain the desired properties of the sampling. Finally, notice that the algorithm fails with probability at most n^{-10} . \square

References

- [ADK⁺16] Ittai Abraham, David Durfee, Ioannis Koutis, Sebastian Krinniger, and Richard Peng. On fully dynamic graph sparsifiers. In *IEEE 57th Annual Symposium on Foundations of Computer Science, FOCS 2016*, pages 335–344. IEEE Computer Society, 2016. doi:10.1109/FOCS.2016.44.
- [AEG⁺22] Simon Apers, Yuval Efron, Paweł Gawrychowski, Troy Lee, Sagnik Mukhopadhyay, and Danupon Nanongkai. Cut query algorithms with star contraction. In *63rd IEEE Annual Symposium on Foundations of Computer Science, FOCS 2022*, pages 507–518. IEEE, 2022. doi:10.1109/FOCS54457.2022.00055.
- [BDM02] Brian Babcock, Mayur Datar, and Rajeev Motwani. Sampling from a moving window over streaming data. In *Proceedings of the Thirteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '02, page 633–634. SIAM, 2002. URL: <http://dl.acm.org/citation.cfm?id=545381.545465>.
- [BNR02] Allan Borodin, Morten N. Nielsen, and Charles Rackoff. (incremental) priority algorithms. In David Eppstein, editor, *Proceedings of the Thirteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '02, pages 752–761. SIAM, 2002. URL: <http://dl.acm.org/citation.cfm?id=545381.545481>.
- [CGL⁺20] Julia Chuzhoy, Yu Gao, Jason Li, Danupon Nanongkai, Richard Peng, and Thatchaphol Saranurak. A deterministic algorithm for balanced cut with applications to dynamic connectivity, flows, and beyond. In *61st IEEE Annual Symposium on Foundations of Computer Science, FOCS 2020*, pages 1158–1167. IEEE, 2020. doi:10.1109/FOCS46700.2020.00111.

[DLT07] Nick G. Duffield, Carsten Lund, and Mikkel Thorup. Priority sampling for estimation of arbitrary subset sums. *J. ACM*, 54(6):32, 2007. doi:10.1145/1314690.1314696.

[dVC25] Tijn de Vos and Aleksander B. G. Christiansen. Tree-packing revisited: Faster fully dynamic min-cut and arboricity. In *Proceedings of the 2025 Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2025*, pages 700–749. SIAM, 2025. doi:10.1137/1.9781611978322.21.

[EGIN97] David Eppstein, Zvi Galil, Giuseppe F. Italiano, and Amnon Nissenzweig. Sparsification - a technique for speeding up dynamic graph algorithms. *J. ACM*, 44(5):669–696, 1997. doi:10.1145/265910.265914.

[EHL25] Antoine El-Hayek, Monika Henzinger, and Jason Li. Fully dynamic approximate minimum cut in subpolynomial time per operation. In *Proceedings of the 2025 Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2025*, pages 750–784. SIAM, 2025. doi:10.1137/1.9781611978322.22.

[GHN⁺23] Gramoz Goranci, Monika Henzinger, Danupon Nanongkai, Thatchaphol Saranurak, Mikkel Thorup, and Christian Wulff-Nilsen. Fully dynamic exact edge connectivity in sublinear time. In *Proceedings of the 2023 ACM-SIAM Symposium on Discrete Algorithms, SODA 2023*, pages 70–86. SIAM, 2023. doi:10.1137/1.9781611977554.CH3.

[GKKT15] David Gibb, Bruce M. Kapron, Valerie King, and Nolan Thorn. Dynamic graph connectivity with improved worst case update time and sublinear space. *CoRR*, abs/1509.06464, 2015. URL: <http://arxiv.org/abs/1509.06464>.

[GMW20] Paweł Gawrychowski, Shay Mozes, and Oren Weimann. Minimum cut in $o(m \log^2 n)$ time. In *47th International Colloquium on Automata, Languages, and Programming, ICALP 2020*, volume 168 of *LIPICS*, pages 57:1–57:15. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020. doi:10.4230/LIPICS.ICALP.2020.57.

[GNT20] Mohsen Ghaffari, Krzysztof Nowicki, and Mikkel Thorup. Faster algorithms for edge connectivity via random 2-out contractions. In *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1260–1279. SIAM, 2020.

[HdLT01] Jacob Holm, Kristian de Lichtenberg, and Mikkel Thorup. Poly-logarithmic deterministic fully-dynamic algorithms for connectivity, minimum spanning tree, 2-edge, and biconnectivity. *J. ACM*, 48(4):723–760, 2001. doi:10.1145/502090.502095.

[HRW20] Monika Henzinger, Satish Rao, and Di Wang. Local flow partitioning for faster edge connectivity. *SIAM J. Comput.*, 49(1):1–36, 2020. doi:10.1137/18M1180335.

[JST24] Wenyu Jin, Xiaorui Sun, and Mikkel Thorup. Fully dynamic min-cut of superconstant size in subpolynomial time. In *Proceedings of the 2024 ACM-SIAM Symposium on Discrete Algorithms, SODA 2024*, pages 2999–3026. SIAM, 2024. doi:10.1137/1.9781611977912.107.

[Kar00] David R. Karger. Minimum cuts in near-linear time. *J. ACM*, 47(1):46–76, 2000. doi:10.1145/331605.331608.

[KK25] Yotam Kenneth-Mordoch and Robert Krauthgamer. Cut-query algorithms with few rounds. *CoRR*, abs/2506.20412, 2025. arXiv:2506.20412, doi:10.48550/ARXIV.2506.20412.

[KKM13] Bruce M. Kapron, Valerie King, and Ben Mountjoy. Dynamic graph connectivity in polylogarithmic worst case time. In *Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2013*, pages 1131–1142. SIAM, 2013. doi:10.1137/1.9781611973105.81.

[KT19] Ken-ichi Kawarabayashi and Mikkel Thorup. Deterministic edge connectivity in near-linear time. *J. ACM*, 66(1):4:1–4:50, 2019. doi:10.1145/3274663.

[Lub86] Michael Luby. A simple parallel algorithm for the maximal independent set problem. *SIAM J. Comput.*, 15(4):1036–1053, 1986. doi:10.1137/0215074.

[MN20] Sagnik Mukhopadhyay and Danupon Nanongkai. Weighted min-cut: sequential, cut-query, and streaming algorithms. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing, STOC 2020*, pages 496–509. ACM, 2020. doi:10.1145/3357713.3384334.

[Sar21] Thatchaphol Saranurak. A simple deterministic algorithm for edge connectivity. In *4th Symposium on Simplicity in Algorithms, SOSA 2021*, pages 80–85. SIAM, 2021. doi:10.1137/1.9781611976496.9.

[Tho07] Mikkel Thorup. Fully-dynamic min-cut. *Comb.*, 27(1):91–127, 2007. doi:10.1007/S00493-007-0045-2.

[TK00] Mikkel Thorup and David R. Karger. Dynamic graph algorithms with applications. In *7th Scandinavian Workshop on Algorithm Theory (SWAT 2000)*, volume 1851 of *Lecture Notes in Computer Science*, pages 1–9. Springer, 2000. doi:10.1007/3-540-44985-X__1.

A Concentration Inequalities

THEOREM A.1. *Let $X_1, \dots, X_m \in [0, a]$ be independent random variables. For any $\delta \in [0, 1]$ and $\mu \geq \mathbb{E}[\sum_{i=1}^m X_i]$, we have*

$$\mathbb{P} \left[\left| \sum_{i=1}^m X_i - \mathbb{E} \left[\sum_{i=1}^m X_i \right] \right| \geq \delta \mu \right] \leq 2 \exp \left(-\frac{\delta^2 \mu}{3a} \right).$$