Near-Optimal Fault-Tolerant Strong Connectivity Preservers

Gary Hoppenworth* Thatchaphol Saranurak† Benyu Wang‡

Abstract

A *k-fault-tolerant connectivity preserver* of a directed n-vertex graph G is a subgraph H such that, for any edge set $F \subseteq E(G)$ of size $|F| \leq k$, the strongly connected components of $G - F$ and $H - F$ are the same. While some graphs require a preserver with $\Omega(2^k n)$ edges [BCR18], the best-known upper bound is $\tilde{O}(k2^k n^{2-1/k})$ edges [CC20], leaving a significant gap of $\Omega(n^{1-1/k})$. In contrast, there is no gap in *undirected* graphs; the optimal bound of $\Theta(kn)$ has been well-established since the 90s [NI92].

We nearly close the gap for directed graphs; we prove that there exists a k-fault-tolerant connectivity preserver with $O(k4^k n \log n)$ edges, and we can construct one with $O(8^k n \log^{5/2} n)$ edges in $\mathrm{poly}(2^k n)$ time.

Our results also improve the state-of-the-art for a closely related object; a *k-connectivity preserver* of G is a subgraph H where, for all $i \leq k$, the strongly i-connected components of G and H agree. By a known reduction, we obtain a k-connectivity preserver with $O(k4^k n \log n)$ edges, improving the previous best bound of $\tilde{O}(k2^k n^{2-1/(k-1)})$ [CC20]. Therefore, for any constant k, our results are optimal to a $\log n$ factor for both problems.

Lastly, we show that the exponential dependency on k is not inherent for k-connectivity preservers by presenting another construction with $O(n\sqrt{kn})$ edges.

* garytho@umich.edu. Supported in part by NSF grant NSF:AF 2153680.
† thsa@umich.edu. Supported by NSF Grant CCF-2238138.
‡ benyuw@umich.edu.

Contents

1 Introduction

A k-fault-tolerant (k-FT) connectivity preserver of a *directed* graph G is a subgraph H that maintains the strongly connected components of G even after k edge failures. Specifically, for any edge set F with size $|F| \leq k$, the strongly connected components of $G - F$ and $H - F$ are the same. In undirected graphs, a k-FT connectivity preserver simplifies to a subgraph that preserves connected components even after edge failures. We study the following question:

Does every graph admit a sparse k-FT connectivity preserver?

Undirected Graphs: Settled. This question is well understood in undirected graphs. Since the 90s, Nagamochi and Ibaraki [NI92] developed a linear-time algorithm to compute a k-FT connectivity preserver with $O(nk)$ edges, which is optimal.[1]

Since then, sparse k-FT connectivity preservers have been highly influential. They were first used to accelerate algorithms for minimum cuts and cut sparsifiers [BK96, BK15]. Subsequently, they have been studied and applied in various computational models, including dynamic algorithms [EGIN97, Tho07, ADK+16], parallel and distributed algorithms [Thu95, DHNS19, Par19], and streaming algorithms [AGM12, GMT15, AS23]. Another line of work strengthens the definition to address stronger types of failures, such as both edge and vertex failures [CKT93, FIN93, EIR95, Fül05, BJ05b, Nag06], bounded-degree edge failures [PT25], and color failures [PST24, PPST25]. Even in these generalized settings, near-optimal bounds of $\tilde{O}(nk)$ have been established.[2]

In the past decade, extensive research has also focused on other fault-tolerant subgraphs beyond connectivity, such as approximate distances [CLPR09, DK11, BDPW18, BDR22, Par22, BDN22, BHP24, PST24, PT25, PPST25, BDKW25], exact distances [PP16, BGPW17, GK17], diameter [BCC+23], and spectral properties [ZSL+19]. Unfortunately, despite success in undirected graphs, progress in directed graphs has been very limited.

State-of-the-Art in Directed Graphs. Directed graphs are fundamentally harder to compress. Even without failures, a standard information-theoretic argument shows that directed graphs cannot be compressed into any $o(n^2)$-sized data structure that answers reachability queries, i.e., whether vertex a has a path to b. This presents a significant barrier for any attempt to compress *asymmetric* connectivity information.

This barrier motivates the study of k-FT connectivity preservers because a strongly connected component is a symmetric notion. In a breakthrough result [BCR18], Baswana, Choudhary, and Roditty completely solved the special case of k-FT *single-source* connectivity preservers, which preserve the strongly connected component containing a *single* vertex s instead of all strongly connected components.[3] Their preserver contains $O(2^k n)$ edges, and this bound is tight even in this special case.

However, for the general case, a significant gap remains. The only non-trivial k-FT connectivity preservers were shown by Chakraborty and Choudhary [CC20] and require $\tilde{O}(k 2^k n^{2-1/k})$ edges. In

[1] Strictly speaking, [NI92] proved that their algorithm outputs a *k-connectivity preserver*, which is slightly weaker and will be defined soon in Section 1.2. However, it is well known in the community that their output is stronger. Benczúr and Karger [BK96] explicitly proved that it preserves all cuts of size at most k, implying it is a k-FT connectivity preserver. See also [Cho21].

[2] $\tilde{O}(\cdot)$ hides logarithmic factors

[3] This is equivalent, up to a constant factor, to the *k-FT single-source reachability preserver* as stated in [BCR18].

particular, even their 2-FT connectivity preservers require $\Omega(n\sqrt{n})$ edges. They posed an open question whether the large $\tilde{\Omega}(n^{1-1/k})$-factor gap between upper and lower bounds can be closed.

1.1 Main Result: Near-Optimal k-Fault-Tolerant Connectivity Preservers

We present k-FT connectivity preservers with $O(n \log n)$ edges for every constant k, reducing the gap from $\tilde{\Omega}(n^{1-1/k})$ to $O(\log n)$.

Theorem 1.1. *For every positive integer k, every n-vertex directed graph admits a k-fault-tolerant connectivity preserver with $O(k4^k n \log n)$ edges. Moreover, we can construct one with $O(8^k n \log^{5/2} n)$ edges in* $\mathrm{poly}(2^k n)$ *time.*

For constant k, the k-FT connectivity preservers of $O(n)$ size prior to Theorem 1.1 are known only for undirected graphs [NI92] or for the single-source version in directed graphs [BCR18]. Theorem 1.1 generalizes both results to directed graphs with only an $O(\log n)$-factor increase in the preserver size. Quantitatively, Theorem 1.1 also significantly improves the bound of $\tilde{O}(n^{2-1/k})$ [CC20] to $O(n \log n)$.

Theorem 1.1 directly implies a data structure of size $O(k4^k n \log n)$ that, given a query set F of size $|F| \leq k$, reports all strongly connected components of $G - F$ in $O(k4^k n \log n)$ time. This already significantly improves upon the space of the previous data structure in [BCR19] which required $O(2^k n^2)$ space and $O(2^k n \log^2 n)$ query time. It also generalizes the data structure of [GIP20] with optimal $O(n)$ space and $O(n)$ query time, but that works only for $k = 1$.

Construction Time. A straightforward construction of our preservers would take $n^{O(k)}$ time. In the second statement of Theorem 1.1, we show an FPT-time construction, at the cost of a slightly larger size bound. It is very interesting if near-linear time construction is possible. However, this remains open even for the single-source version; the state-of-the-art algorithm [BCR18] still requires $\Omega(mn)$ construction time for constant k.

Strong Failure Models. Theorem 1.1 can be extended to address *vertex failures* using the standard vertex-splitting reduction.

Recently, two stronger types of failures, namely *bounded-degree edge failures* [BHP24] and *color failures* [PST24], were introduced. In [PST24, PPST25, BHP24, PT25], fault-tolerant connectivity preservers (and even fault-tolerant spanners) in undirected graphs have been strengthened to handle these types of failures.

We show that this generalization is impossible for directed graphs. In particular, we show directed graphs undergoing just 1-bounded-degree edge failures or 1-color failures require a fault-tolerant connectivity preserver with $\tilde{\Omega}(n^2)$ edges. This marks a fundamental distinction between undirected and directed graphs. We formally define these fault models and prove these lower bounds in Section B.

1.2 Side Result: k-Connectivity Preservers

Theorem 1.1 also immediately advances the state-of-the-art of k-connectivity preservers. At a high level, k-connectivity preservers are subgraphs that preserve higher connectivity but are not fault-tolerant. Before providing a formal definition, let us recall some basic concepts. Let $\mathrm{flow}(G, s, t)$ denote the maximum number of edge-disjoint paths from s to t. We say that s and t are (strongly)

i-connected if $\mathrm{flow}(G,s,t) \geq i$ and $\mathrm{flow}(G,t,s) \geq i$. For any i, the i-connected components uniquely partition the vertex set into parts such that two vertices are in the same part iff they are i-connected. The graph G is i-connected if all pairs of vertices are i-connected.

A *k-connectivity preserver* of a directed graph G is a subgraph H where, for all $i \leq k$, the i-connected components of G and H are the same. Note that every $(k-1)$-FT connectivity preserver of G must be a k-connectivity preserver, but the reverse is false; there exists a directed graph G that admits a k-connectivity preserver with $O(n)$ edges, but requires $\Omega(n2^k)$ edges to preserve strongly connected components under k edge faults. See Theorem B.3 for the proof.

State-of-the-Art in Directed Graphs. In undirected graphs, the $O(nk)$ optimal bounds for k-connectivity preservers have long been established [NI92], but they are much less understood in directed graphs. Currently, tight bounds are known only in special cases. Let $\lambda(G)$ denote the minimum cut size of G. Since the '70s, [Dal77, Mad85] (cf. [BJ05a]) showed that every directed graph G contains a k-connected subgraph H with $O(nk)$ edges for all $k \leq \lambda(G)$, and this is tight. This implies optimal k-connectivity preservers of size $\Theta(nk)$ when $k \leq \lambda(G)$. However, this argument fails when $k > \lambda(G)$. Much later, in 2016, [GILP16] showed a 2-connectivity preserver of size $O(n)$. Very recently, [BCDW24] implies a *single-source* version of k-connectivity preservers that only preserves, for all $i \leq k$, the strongly i-connected component containing the fixed vertex s. Their preserver has an optimal size of $\Theta(nk)$.

In the general case, the best-known k-connectivity preservers prior to our work follow from the best-known construction of k-FT connectivity preservers by [CC20], which implies the bound of $\tilde{O}(k2^k n^{2-1/(k-1)})$ edges. Theorem 1.1 directly provides a significant improvement upon this bound.

Corollary 1.2. *For every positive integer k, every n-vertex directed graph admits a k-connectivity preserver with $O(k4^k n \log n)$ edges.*

Corollary 1.2 is optimal up to an $O(\log n)$ factor when k is constant. But, are there non-trivial k-connectivity preservers even when k is large? We give a strong affirmative answer by showing the first $o(n^2)$-sized k-connectivity preserver for any $k = o(n)$.

Theorem 1.3. *For every positive integer k, every n-vertex directed graph admits a k-connectivity preserver with $O(n\sqrt{nk})$ edges. The preserver can be constructed in polynomial time.*

Theorem 1.3 establishes a strong separation between k-FT connectivity preservers and k-connectivity preservers in directed graphs. The former suffers from an exponential $\Omega(2^k)$ factor (due to [BCR18]) and cannot admit an $o(n^2)$ upper bound when $k \geq \log_2 n$. In contrast, the latter can achieve the $o(n^2)$ upper bound for any $k = o(n)$ by Theorem 1.3. Notably, in undirected graphs, the optimal bounds of the two notions coincide at $\Theta(kn)$ [NI92].

1.3 Mapping the Landscape of Connectivity Preservers in Directed Graphs

The strong separation between k-connectivity and k-FT connectivity established by Theorem 1.3 suggests a rich landscape of connectivity preserver problems in directed graphs. Table 1 lists the best-known upper and lower bounds for all variants of k-FT connectivity preservers and k-connectivity preservers in the literature.

These variants include the s-t, global, single-source, and all-pairs versions. See Section A for the formal definitions and reductions between them. For each variant, $(k-1)$-FT connectivity

	k-fault connectivity		k-connectivity	
	Upper Bound	Lower Bound	Upper Bound	Lower Bound
s-t	$O(n2^k)$ [Thm A.4]	$\Omega(n2^{k/2})$ [Thm A.5]	$O(n\sqrt{k})$ [KL98]	$\Omega(n\sqrt{k})$ [KL98]
global	$O(n2^k)$ [Thm A.4]	$\Omega(n2^{k/2})$ [Thm A.4]	$O(nk)$ [Dal77, BJ05a]	$\Omega(nk)$ [Dal77, BJ05a]
single-source	$O(n2^k)$ [BCR18]	$\Omega(n2^k)$ [BCR18]	$O(nk)$ [BCDW24]	$\Omega(nk)$ [Thm A.4]
all-pairs	$\tilde{O}(nk4^k)$ [Thm 1.1]	$\Omega(n2^k)$ [CC20]	$O(n\sqrt{nk})$ [Thm 1.3] $\tilde{O}(nk4^k)$ [Thm 1.2]	$\Omega(nk)$ [Thm A.4]

Table 1: Upper and lower bounds for connectivity preserver variants in directed graphs.

preservers are also k-connectivity preservers. It is clear from the definitions that the s-t version is a special case of the single-source version, which in turn is a special case of the all-pairs version. Interestingly, we can show reductions that place the global version between the s-t and single-source versions, establishing a linear order of difficulty among all versions.

Our results (Theorems 1.1 and 1.3 and corollary 1.2) for the all-pairs version, along with the reduction in Theorem A.4, and a lower bound for k-fault-tolerant s-t connectivity preservers in Theorem A.5, collectively provide a comprehensive understanding of the landscape of connectivity preservers in directed graphs. Notably, our reduction implies that the global version of k-fault-tolerant connectivity must exhibit an exponential dependency on k in its size, which was unclear prior to our work.

Given our results, the bounds for all variants of k-FT connectivity preservers are nearly tight up to an $O(\log n)$ factor and a constant in the exponent of $2^{\Theta(k)}$. The bounds for all variants of k-connectivity preservers are tight except for all-pairs k-connectivity preservers, where our upper bounds of $O(n\sqrt{nk})$ and $O(k4^k n \log n)$ do not yet match the $\Omega(nk)$ lower bound.

2 Technical Overview

In this section, we give an overview of the techniques and proof ideas in our work. We begin by describing a simple greedy strategy for constructing connectivity preservers that is standard in the area of graph spanners and preservers.

Greedy connectivity preservers. Given any graph G, we can construct a connectivity preserver H of G using the following greedy algorithm. Initially, let $H \leftarrow G$. While there exists an edge $e \in E(H)$ such that $H - e$ is a valid connectivity preserver of H, we let $H \leftarrow H - e$. This algorithm terminates and outputs a subgraph $H \subseteq G$ when no such edge $e \in E(H)$ exists. By a transitivity argument, H is a valid connectivity preserver of G. This algorithm easily extends to k-FT connectivity preservers and k-connectivity preservers as well.

This greedy algorithm is useful for analyzing k-FT connectivity preservers (respectively, k-connectivity preservers) because it allows us to assume wlog that the input graph G is an edge-minimal k-FT connectivity preserver (respectively, k-connectivity preserver) of itself. For simplicity, throughout this section we will assume all graphs we consider have this edge minimality property.

2.1 Fault-Tolerant Connectivity Preservers

We now discuss some of the proof ideas for our construction of fault-tolerant connectivity preservers.

Unbreakability. A key technical tool used in our fault-tolerant connectivity preserver construction is a directed expander hierarchy (formally defined in Lemma 3.7). At a high level, using a directed expander hierarchy allows us to assume the graph has certain good expansion properties. The specific graph expansion property we use is called *unbreakability* [CLP+14, CKL+20].

Definition 2.1 (unbreakable)**.** Let G be a directed graph, and let $U \subseteq V$. Let $q \geq k$ be positive integers. We say that U is (q, k)-unbreakable in G if for every cut $(S, V - S)$ in G such that $|\delta^+_G(S)| \leq k$ or $|\delta^-_G(S)| \leq k$, we have either $|S \cap U| \leq q$ or $|U - S| \leq q$.

Here $\delta^+_G(S)$ denotes the outgoing edges from S, and $\delta^-_G(S)$ denotes the incoming edges to S. We say that a graph G is (q, k)-unbreakable if its vertex set $V(G)$ is (q, k)-unbreakable in G. To give some intuition for why unbreakability might be a useful concept here, we quickly show that unbreakable graphs have sparse k-FT connectivity preservers.

Sparse k-FT connectivity preserver for unbreakable graphs (#1). Let G be a $(q, k+1)$-unbreakable directed graph that is an edge-minimal k-FT connectivity preserver of itself. We will show that G has at most $|E(G)| = O((q + k)n)$ edges. In the existential proof of our k-FT connectivity preservers in Section 4, we let $q = O(k)$, so $|E(G)|$ will be appropriately sparse.

Since graph G is an edge-minimal k-FT connectivity preserver of itself, subgraph $G - e$ is not a k-FT connectivity preserver of G for any $e \in E(G)$. Then for each edge $e = (u, v) \in E(G)$, there exists vertices $s, t \in V(G)$ and a fault set $F \subseteq E(G)$ of size at most $|F| \leq k$ such that s and t are strongly connected in $G - F$ but not strongly connected in $(G - e) - F$. In particular, this implies that there exists a (u, v)-cut $(S, V - S)$ in G of size at most $|\delta^+(S)| \leq k + 1$.

Now we observe that, for any edge $(u, w) \in E(G)$, if $(u, w) \notin \delta^+(S)$, then $w \in S$. Therefore, if $|\delta^+(u)| > q + k + 1$, then there will be more than q out-neighbors of u in S since $|\delta^+(S)| \leq k + 1$. Similarly, if $|\delta^-(v)| > q+k+1$, then there will be more than q vertices in $V - S$. By unbreakability, we know either $|\delta^+(u)| \leq q+k+1$ or $|\delta^-(v)| \leq q+k+1$. Therefore, we can charge edge $e = (u, v)$ to u if $|\delta^+(u)| \leq q+k+1$ or to v if $|\delta^-(v)| \leq q+k+1$. We observe that under this scheme, each vertex in $V(G)$ is charged $O(q+k)$ times and thus the total number of edges in G is $|E(G)| = O((q+k)n)$.

This simple argument shows that unbreakable graphs have sparse k-FT connectivity preservers. Frequently in graph algorithms and data structures, if we can solve a problem on graphs with good expansion, then we can extend our solution to general graphs by decomposing the graph into subgraphs with good expansion (e.g., using expander decomposition). In our case, we use a specific decomposition which we call the *directed expander hierarchy* (formally defined in Lemma 3.7). This directed expander hierarchy is a directed version of the (undirected) expander hierarchy proved in [LPS25], and it follows from a similar proof. However, to successfully apply the directed expander hierarchy, we will actually need a different argument for the sparsity of k-FT connectivity preservers of (q, k)-unbreakable graphs that is more adaptable to our purposes.

Sparse k-FT connectivity preserver for unbreakable graphs (#2). Let G be a directed (q, k)-unbreakable graph that is an edge-minimal k-FT connectivity preserver of itself. We will give an informal argument that G has at most $|E(G)| = O(2^k qn)$ edges. We need the following property of (q,k)-unbreakable sets, which we prove in Section C.1.

Claim 2.2 (Giant component). *Let G be a directed graph, let $U \subseteq V(G)$ be a (q,k)-unbreakable set in G, and let $F \subseteq E(G)$ be a fault set of size at most $|F| \le k$. Then $G - F$ contains a strongly connected component C containing at least $|C \cap U| \ge |U| - 2q$ of the vertices in U. We refer to C as a giant component of U in $G - F$.*

Since graph G is (q,k)-unbreakable, Claim 2.2 implies that after any k edge faults F in graph G, at least $n - 2q$ vertices in G are contained in the same strongly connected component (SCC) S in $G - F$.

By the edge-minimality of G, for each edge $e = (u,v) \in E(G)$, there exists a fault set $F_e \subseteq E(G)$ of size at most $|F_e| \le k$ such that u and v are strongly connected in $G - F_e$ but not strongly connected in $(G - e) - F_e$.[4] Let $C_e \subseteq V(G)$ be the SCC in $G - F_e$ containing vertices u and v. We will bound the number of edges $e \in E(G)$ in two cases, based on whether or not set C_e is a giant component in $G - F_e$.

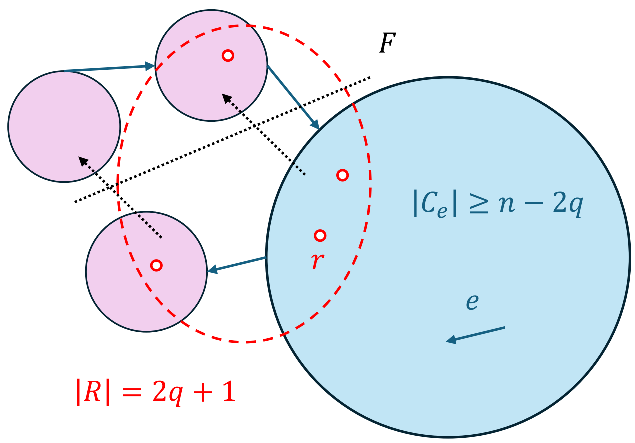

Figure 1: Illustration of the giant component analysis in Case 1.

Case 1: C_e is a giant component in $G - F_e$. In this case, we will directly use the k-FT single-source connectivity preservers of [BCR18]. Let $R \subseteq V(G)$ be an arbitrary set of vertices of size $|R| = 2q + 1$. Let $H \subseteq G$ be the graph obtained by including a k-FT single-source connectivity preserver of G rooted at r for all $r \in R$. Each of these $|R|$ preservers has size $O(2^k n)$ by [BCR18], so graph H has $|E(H)| = O(2^k qn)$ edges.

We claim that graph H contains every edge $e = (u,v)$ whose connected component C_e in $G - F_e$ is a giant component. Since C_e is a giant component in $G - F_e$, we know that $|C_e| \ge n - 2q$, so $C_e \cap R \ne \emptyset$. Then subgraph H contains a k-FT single-source connectivity preserver of G rooted at some vertex $r \in C_e$, so C_e is an SCC of $H - F_e$.

Finally, since u and v are strongly connected in $H - F_e$ but not strongly connected in $(H - e) - F_e$, we conclude that $e \in E(H)$. The argument in this case roughly corresponds to Claim 4.5 in Section 4.

Case 2: C_e is not a giant component in $G - F_e$. Let $E' \subseteq E(G)$ be the set of all edges $e \in E(G)$ such that C_e is not a giant component. We will give an informal argument that $|E'| = O(2^k qn)$.

[4] This corresponds to Claim 4.6 in Section 4.1.1.

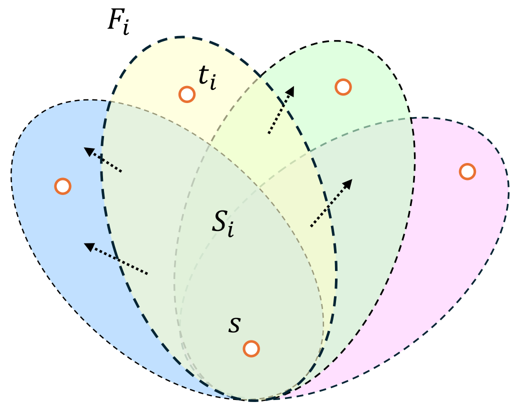

Figure 2: Illustration of the anti-isolation argument in Case 2.

By an averaging argument, there must exist vertices $s, t_1, \dots, t_r \in V(G)$ such that $(s, t_i) \in E'$ for $i \in [1, r]$ and $r = \Theta(|E'|/n)$. Let $e_i = (s, t_i) \in E'$ for $i \in [1, r]$. By Claim 2.2, the SCC C_{e_i} in $G - F_{e_i}$ containing vertices $\{s, t_i\}$ has size at most $|C_{e_i}| \le 2q$ for each $i \in [1, r]$. Since the sets C_{e_i} each have size $|C_{e_i}| = O(q)$, we can ensure that $t_j \notin C_{e_i}$ for all $j \neq i$, using a greedy argument that decreases the size of our set of vertices $\{t_1, \dots, t_r\}$ by at most a factor of q (see Claim 4.9 for details).

Next, we will apply the so-called *anti-isolation lemma*, which was previously proved in [PW18] with the weaker bound of $r \le k4^k$. Below, we state an improved bound using insights from [BCR18]. See Section 6.1 for the formal proof.

Lemma 2.3 (Improved anti-isolation)**.** *Let $s, t_1, \dots, t_r$ be vertices in a directed graph G, and let $F_1, \dots, F_r \subseteq E(G)$ be edge sets each of size at most k such that for all $i, j \in [1, r]$ there is an (s, t_j)-path in $G - F_i$ if and only if $i = j$. Then $r \le 2^k$.*

By our earlier discussion, we know that s and t_j are strongly connected in $G - F_{e_i}$ if and only if $i = j$ (see Figure 2 for reference). For simplicity, we will further assume that s can reach t_j in $G - F_{e_i}$ if and only if $i = j$. Then the anti-isolation lemma directly implies that $\Theta(|E'|/(qn)) = r \le 2^k$, so we conclude that $|E'| = O(2^k qn)$. This argument roughly corresponds to Claim 4.10.

Our two-case analysis shows that (q, k)-unbreakable graphs have k-FT connectivity preservers of size $O(2^k qn)$. Our proof in Section 4.1 roughly follows this argument. We now build some intuition for how we combine these ideas with the directed expander hierarchy.

Directed expander hierarchy with two levels. The two-level directed expander hierarchy has a particularly simple interpretation. It implies that given a directed graph G, there exists a set of vertices $U \subseteq V(G)$ such that U is (q, k)-unbreakable in G and every SCC C of $G - U$ is (q, k)-unbreakable in $G - U$, for an appropriate choice of q.

We briefly describe how we can use the ideas we have developed, along with this two-level hierarchy, to bound the number of edges we need to preserve in a k-FT connectivity preserver of G. As we have done throughout this section, we will assume that G is an edge-minimal k-FT connectivity preserver of itself; our goal then is to upper bound $|E(G)|$.

If C is an SCC in $G - U$, then since $G[C]$ is (q, k)-unbreakable, we can upper bound the number of edges in $G[C]$ using the sparse k-FT connectivity preservers for unbreakable graphs that we developed earlier; the same holds for $G[U]$ as well. The remaining edges in $E(G)$ that we need to upper bound fall into two categories: 1) edges between distinct SCC's C_1, C_2 in $G - U$, and 2) edges between U and $V(G) - U$.

Let $H_U \subseteq G$ be an edge-minimal subgraph of G that preserves the k-FT connectivity between all pairs of vertices (u, v) in $U \times V(G)$. We say that H_U is a k-FT *sourcewise* connectivity preserver of G with respect to set U.[5] It is not difficult to see that H_U must contain all edges in both of the two categories of edges in $E(G)$ identified above. Then in order to upper bound $|E(G)|$ it is sufficient to upper bound $|E(H_U)|$.

As it turns out, the second sparse preserver argument we saw in this section can be easily extended to solve this source-restricted connectivity problem using only $O(4^k qn)$ edges (Lemma 4.4).

Full directed expander hierarchy. In our full directed expander hierarchy in Lemma 3.7, we will have $\ell = O(\log n)$ levels, which allows us to set $q = O(k)$. The ideas discussed in the previous paragraph easily extend to this full expander hierarchy, leading to a final size bound of $O(\ell \cdot 4^k qn) = O(k4^k n \log n)$, as we prove in Section 4.1.2.

Computing a k-FT connectivity preserver in FPT-time. So far, all of the arguments in this section use a greedy algorithm to construct edge-minimal k-FT connectivity preservers. Naively implementing this algorithm requires $n^{O(k)}$ time to check every fault set $F \subseteq E(G)$ of size $|F| \leq k$. We implement a modified version of the greedy algorithm in poly$(2^k n)$ time by adapting some of the ideas in our existential bound. Our FPT-time algorithm uses an algorithmic lemma about important cuts (Lemma 4.1), which we prove in Section 6.

2.2 k-Connectivity Preservers

Next, we describe the high-level idea how to show a k-connectivity preserver with $O(n\sqrt{nk})$ edges. Again, we start by assuming the graph G is an edge-minimal k-connectivity preserver and argue that G must have $O(n\sqrt{nk})$ edges.

Decomposition. Our argument is based on the following two-level decomposition. Set $q = \sqrt{nk}$. We maintain a partition $\mathcal{S}$ of vertex set V and a collection of cuts $\mathcal{C}$. Initially, we let $\mathcal{S} = \{V\}$ and $\mathcal{C} = \emptyset$. While there exists a set $S \in \mathcal{S}$ and a cut (L, R) of size $|\delta^+(L)| \leq k$ such that $|S \cap L| \geq q$ and $|S \cap R| \geq q$, we replace S using $S \cap L$ and $S \cap R$ in $\mathcal{S}$ and add the cut (L, R) to $\mathcal{C}$. When the process terminates, we observe that:

- Every set $S \in \mathcal{S}$ has size $|S| \geq q$ and is (q, k)-unbreakable.
- All edges between components in $\mathcal{S}$ form a subset of $\cup_{(L,R)\in\mathcal{C}}(\delta^+(L) \cup \delta^-(L))$.

By these two observations, we bound the number of edges inside some final component S, or between two final components, respectively.

[5] We formally define k-FT sourcewise connectivity preservers in Definition 4.3.

Inside components. We bound the edges inside each component using the same strategy as we bound the size of a k-FT preserver in $(q, k+1)$-unbreakable graphs. By minimality, we can prove similarly that every edge $e \in G$ is in a small cut $(X, V-X)$ where $|\delta^+(X)| \leq k$. Let us consider any $S \in \mathcal{S}$ and any $e = (u, v) \in G[S]$. We use the same charging argument, with the only difference that, now we instead argue that either the out-degree of u that is *inside* $G[S]$ is at most $q+k$, or the in-degree of v that is *inside* $G[S]$ is at most $q+k$. Finally, we deduce that the total number of edges inside $G[S]$ is $O((q+k)|S|)$ for any final component $S \in \mathcal{S}$ and all components sum up to $O((q+k)n) = O(n\sqrt{nk})$.

Between components. Since all of the edges between components in $\mathcal{S}$ are contained in the set $\cup_{(L,R)\in\mathcal{C}}(\delta^+(L) \cup \delta^-(L))$, we aim to bound the number of edges $|\delta^+(L) \cup \delta^-(L)|$ crossing each cut (L, R) we operated on in the decomposition. By our decomposition, $|\delta^+(L)| \leq k$. We will now show that $|\delta^-(L)| \leq kn$.

Let us first think about a *path* from any vertex s to any vertex t. Since it will go through the edges $\delta^+(L)$ and $\delta^-(L)$ interleavingly, and $|\delta^+(L)| \leq k$, the number of edges in $\delta^-(L)$ it goes through can be at most $k+1$. Similarly, for a *flow* with at most k edge-disjoint paths from any vertex s to any vertex t, since every flow path will still go through the edges $\delta^+(L)$ and $\delta^-(L)$ interleavingly, and all of the paths go through $\delta^+(L)$ at most k times in total, the flow goes through $\delta^-(L)$ at most $k+k=2k$ times. Therefore, $2k$ edges in $\delta^-(L)$ will be sufficient to preserve the k-bounded connectivity between any vertex pair (s,t).

Another observation is, to preserve k-connectivity for *all pairs* of vertices (s,t), it suffices that we preserve k-connectivity for only $O(n)$ *demand pairs* of vertices (s,t). These demand pairs can be viewed as all edges in some analog of Gomory-Hu trees [GH61, CH91] for symmetric connectivity in directed graphs.[6] Therefore, to preserve k-connectivity for these $O(n)$ *demand pairs*, we only need $O(nk)$ edges in $\delta^-(L)$. This bound on the number of edges crossing cut (L,R) in $\mathcal{C}$ is sufficient to show that there are at most $O(n\sqrt{nk})$ edges between components.

3 Preliminaries

Let $G = (V, E)$ be a directed graph with n vertices and m edges. For any edge set $F \subseteq V \times V$, we use $G - F$ to denote the graph formed by removing edges in F from G, and we use $G + F$ to denote the graph formed by adding edges in F to G. When $F = \{e\}$ is a singleton set, we abbreviate this to $G - e$ or $G + e$. For a vertex set $S \subseteq V$, we use $G[S]$ to denote the induced subgraph of S in G.

For any vertex $v \in V$, we use $\delta_G^+(v)$ and $\delta_G^-(v)$ to denote the set of outgoing edges from v and incoming edges to v respectively. When $(u, v) \in E$, we say v is an out-neighbor of u, and u is an in-neighbor of v. We denote the set of out- and in- neighbors of $v \in V$ in G by $N_G^+(v)$ and $N_G^-(v)$.

We define a cut as a partition $(S, V-S)$ of the vertex set V into two disjoint subsets, where $S \subseteq V$. We use $\delta_G^+(S)$ to denote the set of outgoing edges from S, and $\delta_G^-(S)$ to denote the set of incoming edges to S. For two disjoint sets $X, Y \subseteq V$, we say that $(S, V-S)$ is an (X,Y)-cut if $X \subseteq S$ and $Y \subseteq V-S$. In particular, if $X = \{x\}$ and $Y = \{y\}$, we call $(S, V-S)$ an (x,y)-cut. We typically refer to the size of a cut $(S, V-S)$ as the number of outgoing edges from S, $|\delta_G^+(S)|$. In particular, we say that an (X,Y)-cut $(S, V-S)$ is a minimum (X,Y)-cut if $|\delta_G^+(S)| \leq |\delta_G^+(S')|$ for every (X,Y)-cut $(S', V-S')$.

[6] To be more specific, the trees lose some "cut-equivalency" but still satisfy that the connectivity between any pair of vertices is the minimum connectivity on the tree path. See Lemma 5.6 for details.

We define flow(G, s, t) to be the maximum number of pairwise edge-disjoint paths from s to t in G. Additionally, we define the *symmetric* connectivity $\lambda_G(s,t)$ between s and t to be $\lambda_G(s,t) = \min(\text{flow}(G,s,t), \text{flow}(G,t,s))$. We define the k-bounded connectivity to be $\lambda_G^k(s,t) = \min(\lambda_G(s,t), k)$. For terminal sets X, Y, we define flow(G, X, Y) to be the maximum number of pairwise edge-disjoint (X, Y)-paths.

3.1 Farthest minimum cuts and important cuts

We need to define a class of well-structured (X, Y)-cuts, which we call reachable cuts.

Definition 3.1 (Reachable Cuts)**.** We say that an (X, Y)-cut $(S, V - S)$ is *out-reachable* if for every $s \in S$ there exists an $x \in X$ such that x can reach s in $G - \delta_G^+(S)$. Likewise, we say that $(S, V - S)$ is *in-reachable* if for every $s \in S$ there exists an $x \in X$ such that s can reach x in $G - \delta_G^-(S)$.

In some sense, reachable cuts are without loss of generality.

Claim 3.2. *Let $(S, V - S)$ be an (X, Y)-cut. Then there exists an out-reachable (X, Y)-cut $(S', V - S')$ such that $\delta_G^+(S') \subseteq \delta_G^+(S)$. Likewise, there exists an in-reachable (X, Y)-cut $(S'', V - S'')$ such that $\delta_G^-(S'') \subseteq \delta_G^-(S)$.*

Proof. Consider the out-reachable case. For any (X, Y)-cut $(S, V - S)$, we let S' be the set that X can reach in $G - \delta_G^+(S)$. It follows that S' is out-reachable and $\delta_G^+(S') \subseteq \delta_G^+(S)$. The in-reachable case can be proved in the same way. □

A classic result of Ford-Fulkerson implies the existence of a unique min-cut that is "farther" than every other min-cut.

Lemma 3.3 (Farthest Minimum Cuts [FF62])**.** *Given a directed graph G, we say that an (X, Y)-cut $(S, V - S)$ is a farthest minimum (X, Y)-cut in G if $(S, V - S)$ is an out-reachable, minimum (X, Y)-cut, and there does not exist an out-reachable, minimum (X, Y)-cut $(S', V - S')$ such that $S \subset S'$. For disjoint vertex sets $X, Y \subseteq V(G)$, the unique farthest minimum (X, Y)-cut always exists and can be computed in max-flow time.*

We denote the farthest minimum (X, Y)-cut in G by FMC(G, X, Y). Important cuts generalize the notion of farthest minimum cuts to cuts that are not min-cuts.

Definition 3.4 (Important Cuts [Mar11])**.** We say that an (X, Y)-cut $(S, V - S)$ is an important (X, Y)-cut if $(S, V - S)$ is out-reachable, and there does not exist an out-reachable (X, Y)-cut $(S', V - S')$ with size at most $|\delta_G^+(S')| \le |\delta_G^+(S)|$ such that $S \subset S'$.

The definition of important cuts is well-known in the parametrized complexity community. We also state the following claim, which is clear from definition:

Claim 3.5. *If $(S, V - S)$ is an out-reachable (X, Y)-cut of size $|\delta_G^+(S)| \le k$, then there exists an important (X, Y)-cut $(S', V - S')$ of size $|\delta_G^+(S')| \le k$ such that $S \subseteq S'$.*

Proof. If $(S, V - S)$ is important, then we let $S' = S$ and the claim holds. Otherwise, since $(S, V - S)$ is not important, we can find $(S', V - S')$ that $|\delta_G^+(S')| \le |\delta_G^+(S)|$ and $S \subset S'$. We can set $S = S'$ and repeat this procedure until $(S, V - S)$ is important. □

3.2 Unbreakable sets and directed expander hierarchies

We begin this section by defining (q, k)-unbreakable sets.

Definition 2.1 (unbreakable)**.** Let G be a directed graph, and let $U \subseteq V$. Let $q \geq k$ be positive integers. We say that U is (q, k)-unbreakable in G if for every cut $(S, V - S)$ in G such that $|\delta_G^+(S)| \leq k$ or $|\delta_G^-(S)| \leq k$, we have either $|S \cap U| \leq q$ or $|U - S| \leq q$.

Unbreakable sets have the useful property that if we remove a small number of edges from the graph, then almost all of the vertices in the unbreakable set will lie in the same strongly connected component in the resulting graph. We defer its proof to Section C.1.

Claim 3.6 (Giant component)**.** *Let G be a directed graph, let $U \subseteq V(G)$ be a (q, k)-unbreakable terminal set in G, and let $F \subseteq E(G)$ be a set of edges with size at most k. Then $G - F$ contains a strongly connected component C containing at least $|C \cap U| \geq |U| - 2q$ of the nodes in U. We refer to C as a giant component of U in $G - F$.*

Our upper bounds for fault-tolerant symmetric connectivity preservers will rely on the following directed expander hierarchy decomposition for unbreakable sets.

Lemma 3.7 (Directed Expander Hierarchy for Unbreakable Sets)**.** *Let G be an n-node directed graph, and let q, k be positive integers with $q \geq k/\phi$, where $0 < \phi \leq 1$. Then there exists a partition $\{V_1, V_2, \ldots, V_\ell\}$ of $V(G)$ with the following properties:*

1. *Fix an $i \in [1, \ell]$. Let $V_{\leq i} = V_1 \cup \cdots \cup V_i$, and let $C \subseteq V(G)$ be a strong component in $G[V_{\leq i}]$. Then the set $V_i \cap C$ is (q, k)-unbreakable in $G[C]$.*

2. *The number of sets in the partition is at most $\ell = O(\log n)$.*

Moreover, when $\phi = 1/2$, this partition can be computed in exponential time, and when $\phi = O(1/\sqrt{\log n})$, this partition can be computed in polynomial time.

A version of expander hierarchies in *undirected* graphs was first introduced in the context of oblivious routing [Rac02, BKR03, HHR03] and is called *tree flow sparsifiers* or *Räcke tree*. It later finds many more applications for devising almost-linear time flow algorithms [RST14, ADK+16, GRST21, VDBCK+24, LRW25].

In 2007, a weaker variant of the expander hierarchy was introduced by Patrascu and Thorup [PT07] for connectivity oracles under edge faults. It turns out, in contrast to tree flow sparsifiers, that this weaker variant naturally generalizes from the usual edge expansion to vertex expansion [LS22, LPS25]. Recently, [BBST24] introduced a directed expander hierarchy as a further generalization to directed expansion. They also presented a fast construction algorithm that is highly complicated.

Here, to prove Lemma 3.7, we give a very simple construction of the directed expander hierarchy. The proof follows closely and generalizes the construction in undirected graphs of [LPS25]. For the sake of completeness, we defer its proof to Section C.2.

4 Fault-Tolerant Connectivity Preservers

We now present our new results for fault-tolerant connectivity preservers. In section 4.1, we prove our existential upper bound for fault-tolerant connectivity preservers (Theorem 4.2). In section 4.2, we show how to modify our construction to run in FPT-time (Theorem 4.12).

Throughout this section, we will repeatedly use the following lemma about important cuts, which is proved in full generality in Theorem 6.1 of Section 6.

Lemma 4.1 (Important cut container). *Let k be a positive integer, let G be an n-vertex directed graph, and let $X, Y \subseteq V(G)$ be disjoint sets. There exists an (X, Y)-cut $(S, V - S)$ of size at most $|\delta^+(S)| \leq 2^{k-1}$ such that every important (X, Y)-cut $(S', V - S')$ of size at most $|\delta^+(S')| \leq k$ satisfies $S' \subseteq S$. Cut $(S, V - S)$ can be computed in $O(k2^k m + k^2 4^k)$ time.*

We define some new notation that we will use when we apply Lemma 4.1.

- We will let $\mathrm{ImpCut}(G, X, Y, k, +)$ denote an (X, Y)-cut $(S, V - S)$ in G of size at most $|\delta^+(S)| \leq 2^{k-1}$ such that every out-reachable (X, Y)-cut $(S', V - S')$ of size $|\delta^+(S')| \leq k$ satisfies $S' \subseteq S$.
- We will let $\mathrm{ImpCut}(G, X, Y, k, -)$ denote an (X, Y)-cut $(S, V - S)$ in G of size at most $|\delta^-(S)| \leq 2^{k-1}$ such that every in-reachable (X, Y)-cut $(S', V - S')$ of size $|\delta^-(S')| \leq k$ satisfies $S' \subseteq S$.

By Lemma 4.1 and Claim 3.5, $\mathrm{ImpCut}(G, X, Y, k, +)$ and $\mathrm{ImpCut}(G, X, Y, k, -)$ are well-defined and can be computed in $O(k2^k m + k^2 4^k)$ time.

Additionally, throughout this section we will construct k-FT single-source connectivity preservers as a subroutine in our algorithms. As mentioned in Definition A.3, k-FT single-source connectivity preservers maintain the strong connectivity between a source node s and the rest of the vertex set $V(G)$ under k edge faults. We use $\mathrm{SSCP}(G, s, k)$ to denote a k-FT single-source connectivity preserver of G rooted at source vertex s with $O(2^k |V(G)|)$ edges. Subgraph $\mathrm{SSCP}(G, s, k)$ always exists by [BCR18].

4.1 Existential upper bound for fault-tolerant connectivity preservers

The goal of this section is to present our existential upper bound for fault-tolerant connectivity preservers.

Theorem 4.2. *For every positive integer k, every n-vertex directed graph admits a k-fault-tolerant connectivity preserver with $O(k4^k n \log n)$ edges.*

We will prove Theorem 4.2 in two steps. In the first step, we will show how to construct sparse fault-tolerant *sourcewise* connectivity preservers between a well-connected set of source vertices U and the rest of the graph (Lemma 4.4). In the second step, we combine our sourcewise preservers with a directed expander hierarchy (Lemma 3.7) to construct a sparse fault-tolerant connectivity preserver of the input graph.

4.1.1 Fault-tolerant sourcewise connectivity preservers for unbreakable sets

We begin with a formal definition of fault-tolerant sourcewise connectivity preservers.

Definition 4.3 (Fault-tolerant sourcewise connectivity preservers). Let G be an n-vertex directed graph, and let $S \subseteq V(G)$ be a set of source vertices. We say that a subgraph $H \subseteq G$ is a k-FT sourcewise connectivity preserver of G with respect to set S if for any pair of vertices $(s, v) \in S \times V(G)$ and any edge fault set $F \subseteq E(G)$ of size at most $|F| \leq k$, s and v are strongly connected in $G - F$ if and only if s and v are strongly connected in $H - F$.

Our fault-tolerant sourcewise connectivity preservers can be summarized with the following lemma.

Lemma 4.4. *Let G be an n-vertex directed graph, and let $U \subseteq V(G)$ be a (q,k)-unbreakable set in G. There exists a k-FT sourcewise connectivity preserver H of graph G with respect to source vertices U of size $|E(H)| = O(4^k qn)$.*

We will construct our fault-tolerant sourcewise connectivity preserver using a decremental greedy strategy that is standard in the area of graph spanners and preservers.

Greedy Algorithm. Let $G = (V, E)$ be an n-vertex directed graph, and let $U \subseteq V$ be a (q,k)-unbreakable set in G. We may assume G is strongly connected, without loss of generality. We will initialize our fault-tolerant sourcewise connectivity preserver H to be the input graph, so $H \leftarrow G$. While there exists an edge $e \in E(H)$ such that $H - e$ is a k-FT sourcewise connectivity preserver of G with respect to set U, we will remove edge e from H, so that $H \leftarrow H - e$. This algorithm terminates and outputs H when no such edge $e \in E(H)$ exists.

Correctness. By construction, H is a k-FT sourcewise connectivity preserver of G with respect to set U. Specifically, H is an *edge-minimal* k-FT sourcewise connectivity preserver of G with respect to U, by the terminating condition of the greedy algorithm.

Time Complexity. There are $n^{O(k)}$ fault sets $F \subseteq E(G)$ of size $|F| \leq k$. Consequently, we can check if a subgraph $H \subseteq G$ is a k-FT sourcewise connectivity preserver of G with respect to U in $n^{O(k)}$ time. We conclude that the greedy algorithm can be implemented in $n^{O(k)}$ time.

Size Analysis. Before we can bound the size of $E(H)$, we need to define a subgraph $J \subseteq H$ that is useful for our analysis. Let U' be an arbitrary subset $U' \subseteq U$ of size $|U'| = \min(2q+1, |U|)$. Let

$$J = \bigcup_{u \in U'} \mathrm{SSCP}(H, u, k)$$

be the subgraph of H obtained by unioning $|U'|$ many k-FT single-source connectivity preservers, each rooted at a distinct vertex in U'. The following claim states that subgraph J preserves strong connectivity between pairs of vertices $(u, v) \in U \times V(G)$ for a certain category of fault sets.

Claim 4.5 (Large SCC fault sets)**.** *Let $F \subseteq E(G)$ be a fault set of size $|F| \leq k$, and let $C \subseteq V(G)$ be an SCC of $G - F$. If $|C \cap U| > 2q$, then vertices in C are pairwise strongly connected in $J - F$.*

Proof. By Claim 3.6, if $|C \cap U| > 2q$, then $|C \cap U| \geq |U| - 2q$. Since $|U'| = \min(2q+1, |U|)$, it follows that $C \cap U' \neq \emptyset$. Let $x \in C \cap U'$ be a vertex in the nonempty intersection of U' and C.

Now we make the following sequence of observations. First, vertices in C are pairwise strongly connected in $G - F$. Second, since H is a k-FT sourcewise connectivity preserver of G with respect to U, and $C \cap U \neq \emptyset$, it follows that vertices in C are pairwise strongly connected in $H - F$. Finally, subgraph J contains $\mathrm{SSCP}(H, x, k)$, where $x \in C$, so we conclude that vertices in C are pairwise strongly connected in $J - F$ as well. $\square$

Since J is the union of $O(q)$ many k-FT single-source connectivity preservers, $|E(J)| = O(2^k qn)$ by Theorem 1.1 of [BCR18]. Then to prove $|E(H)| = O(4^k qn)$, it will be sufficient to prove that $|E(H) - E(J)| = O(4^k qn)$. Before proceeding with the proof, we will need the following structural claim about the edges in H.

Claim 4.6. *For each edge $(s,t) \in E(H)$, there exists a vertex $u \in U$ and a fault set $F \subseteq E(H)$, $|F| \leq k$, such that:*

1. *Vertices s, t, and u are strongly connected in $H - F$, and*
2. *$H - F$ contains either*
   - *an out-reachable (u,t)-cut $(S, V - S)$ such that $\delta^+_{H-F}(S) = \{(s,t)\}$, or*
   - *an in-reachable (u,s)-cut $(S, V - S)$ such that $\delta^-_{H-F}(S) = \{(s,t)\}$.*

Proof. If $(s,t) \in E(H)$, then there exists a fault set F, with $|F| \leq k$, such that $(H - (s,t)) - F$ does not preserve the strong connectivity of a pair of vertices u, v in $H - F$, where $u \in U$ and $v \in V$. Specifically, vertices u, v are strongly connected in $H - F$, but not strongly connected in $(H - (s,t)) - F$. This implies that edge (s,t) is inside the SCC containing u and v in $H - F$, so vertices s, t, and u are all strongly connected in $H - F$, as claimed.

Since u and v are strongly connected in $H - F$ and not strongly connected in $(H - (s,t)) - F$, there must exist a (u,v)-cut $(S, V - S)$ in H such that $\delta^+_{H-F}(S) = \{(s,t)\}$ or $\delta^-_{H-F}(S) = \{(s,t)\}$. If $\delta^+_{H-F}(S) = \{(s,t)\}$, then by Claim 3.2 there exists an out-reachable (u,t)-cut $(S', V - S')$ such that $\delta^+_{H-F}(S') = \{(s,t)\}$. Else if $\delta^-_{H-F}(S) = \{(s,t)\}$, then by Claim 3.2 there exists an in-reachable (u,s)-cut $(S', V - S')$ such that $\delta^-_{H-F}(S') = \{(s,t)\}$. □

For each edge $e = (s,t) \in E(H) - E(J)$ we will associate a label $\ell(e) \in U \times V$ as follows. By Claim 4.6, there exist a vertex $u \in U$ and a fault set $F \subseteq E(H)$, $|F| \leq k$, such that vertices s, t, and u are all strongly connected in $H - F$, but not all strongly connected in $(H - (s,t)) - F$.

- If H has an out-reachable (u,t)-cut $(S, V - S)$ such that $\delta^+_H(S) \subseteq F \cup \{(s,t)\}$, then we let $\ell(e) = (u,t)$.
- If H has an in-reachable (u,s)-cut $(S, V - S)$ such that $\delta^-_H(S) \subseteq F \cup \{(s,t)\}$, then we let $\ell(e) = (u,s)$.

By Claim 4.6, $\ell(e)$ is well-defined for all $e \in E(H)$.[7] Let

$$\mathcal{L} = \{\ell(e) \mid e \in E(H) - E(J)\}.$$

The set of labels $\mathcal{L}$ is useful because we can use it to upper bound the size of $E(H)$.

Claim 4.7. $|E(H)| \leq 2^{k+1} \cdot |\mathcal{L}| + O(2^k qn)$

Proof. Since $|E(J)| = O(2^k qn)$ by an earlier discussion, it will be sufficient to show that $|E(H) - E(J)| \leq 2^{k+1} \cdot |\mathcal{L}|$. Suppose towards contradiction that $|E(H) - E(J)| > 2^{k+1} \cdot |\mathcal{L}|$. Then there exists a vertex pair $(u,v) \in U \times V$ such that $|\{e \in E(H) - E(J) \mid \ell(e) = (u,v)\}| > 2^{k+1}$. Note that every edge $e \in E(H)$ with label $\ell(e) = (u,v)$ is of the form $e = (w,v)$ or $e = (v,w)$ for some vertex $w \in V$. Let $E_1 = \{(w,v) \in E(H) - E(J) \mid \ell((w,v)) = (u,v)\}$, and let $E_2 = \{(v,w) \in E(H) - E(J) \mid \ell((v,w)) = (u,v)\}$. Either $|E_1| > 2^k$ or $|E_2| > 2^k$.

Suppose that $|E_1| > 2^k$. By the definition of $\ell(e)$, for each edge $(w,v) \in E_1$, there exists an out-reachable (u,v)-cut $(S_w, V - S_w)$ of size $|\delta^+_H(S_w)| \leq k+1$ such that $(w,v) \in \delta^+_H(S_w)$. Let

[7] If there are multiple valid choices of a label $\ell(e)$, then we select one arbitrarily.

$(S', V - S') = \text{ImpCut}(H, u, v, k+1, +)$. Then by Lemma 4.1, $(S', V - S')$ is a (u, v)-cut such that $S_w \subseteq S'$, so $(w, v) \in \delta_H^+(S')$. This implies that $E_1 \subseteq \delta_H^+(S')$. However, by the important cut lemma (Lemma 4.1), $|\delta_H^+(S')| \leq 2^k$, a contradiction.

If instead $|E_2| > 2^k$, then we let $(S', V - S') = \text{ImpCut}(H, u, v, k+1, -)$. By a symmetric argument as before, we can argue that $E_2 \subseteq \delta_H^-(S')$. However, by Lemma 4.1, $|\delta_H^-(S')| \leq 2^k$, a contradiction. □

By Claim 4.7, it is sufficient to upper bound $|\mathcal{L}|$ in order to upper bound $|E(H)|$. We prove that $|\mathcal{L}| = O(2^k qn)$ in Claim 4.10 by appealing to our important cut lemma (Lemma 4.1). The proof of Claim 4.10 uses the following structural claim about the edges in $E(H) - E(J)$.

Claim 4.8. *Let $(s, t) \in E(H) - E(J)$. By Claim 4.6, there exists a vertex $u \in U$ and a fault set $F \subseteq E(H)$, $|F| \leq k$, such that vertices s, t, and u are strongly connected in $H - F$, but not strongly connected in $(H - (s, t)) - F$. Let C be the SCC in $H - F$ containing vertices s, t, and u. Then $|C \cap U| \leq 2q$.*

Proof. Suppose, towards a contradiction, that $|C \cap U| > 2q$. Let C' be the SCC in $G - F$ containing vertices s, t, and u. Since H is a k-FT sourcewise connectivity preserver of G with respect to U and $|F| \leq k$, it must follow that $C = C'$.

Then $|C' \cap U| > 2q$, and by Claim 4.5, vertices s, t, and u are strongly connected in $J - F$. Because $J \subseteq H$ and s and t are not strongly connected in $(H - (s, t)) - F$, it follows that s and t are not strongly connected in $(J - (s, t)) - F$ as well. Since s and t are strongly connected in $J - F$ but not strongly connected in $(J - (s, t)) - F$, we conclude that $(s, t) \in E(J)$. This contradicts our assumption that $(s, t) \in E(H) - E(J)$. □

Additionally, our proof of Claim 4.10 will need the following general claim about set systems.

Claim 4.9. *Let $\mathcal{S} = \{S_u\}_{u \in U}$ be a family of sets indexed by a set U such that $u \in S_u$ and $|S_u \cap U| \leq q$ for all $S_u \in \mathcal{S}$. Then there exists a subset $\mathcal{S}_1 \subseteq \mathcal{S}$ of size $|\mathcal{S}_1| \geq |\mathcal{S}|/(2q)$ such that for all distinct sets $S_u, S_{u'} \in \mathcal{S}_1$, $u' \notin S_u$ and $u \notin S_{u'}$.*

Proof. We define a directed graph $G_{\mathcal{S}}$ with vertex set $V(G_{\mathcal{S}}) = \mathcal{S}$. We add the directed edge $(S_u, S_w) \in \mathcal{S} \times \mathcal{S}$ to $G_{\mathcal{S}}$ if $w \in S_u$. Since $|S_u \cap U| \leq q$, each vertex $S_u \in \mathcal{S}$ has out degree at most $q - 1$ in graph $G_{\mathcal{S}}$.

We say that a (possibly directed) graph is α-degenerate if every subgraph of the graph contains a vertex of total degree at most α. Since every vertex of $G_{\mathcal{S}}$ has outdegree at most $q - 1$, every subgraph of $G_{\mathcal{S}}$ has average degree at most $2q - 2$. Consequently, graph $G_{\mathcal{S}}$ is $(2q-2)$-degenerate. By the standard greedy vertex coloring algorithm for bounded-degeneracy graphs, graph $G_{\mathcal{S}}$ has a proper vertex coloring with at most $2q - 1$ colors.

In particular, this implies that graph $G_{\mathcal{S}}$ has an independent set $\mathcal{S}_1 \subseteq \mathcal{S}$ of size at least $|\mathcal{S}|/(2q - 1) \geq |\mathcal{S}|/(2q)$. Since set $\mathcal{S}_1 \subseteq \mathcal{S}$ is an independent set in $G_{\mathcal{S}}$, it follows that for all distinct sets $S_u, S_{u'} \in \mathcal{S}_1$, $u' \notin S_u$ and $u \notin S_{u'}$, as claimed. □

We are now ready to prove Claim 4.10.

Claim 4.10. *$|\mathcal{L}| \leq 2^{k+4} qn$.*

Proof. Suppose towards contradiction that $|\mathcal{L}| > 2^{k+4}qn$. Then there exists a vertex $v \in V$ such that $|\mathcal{L} \cap (U \times \{v\})| > 2^{k+4}q$. Let $\mathcal{L}_v = \mathcal{L} \cap (U \times \{v\})$. For each $(u, v) \in \mathcal{L}_v$, let $e_u = (s_u, t_u) \in E(H) - E(J)$ be an edge such that $\ell(e_u) = (u, v)$. By Claim 4.6 and the definition of $\ell(e_u)$, there exists a fault set $F_u \subseteq E(H)$, $|F_u| \leq k$, such that vertices v and u are strongly connected in $H - F_u$, but not strongly connected in $(H - e_u) - F_u$. Let C_u denote the SCC in $H - F_u$ containing vertices v and u. By Claim 4.8, $|C_u \cap U| \leq 2q$.

We claim that for each $(u, v) \in \mathcal{L}_v$, there exists a cut $(W_u, V - W_u)$ in H such that

1. $C_u \subseteq W_u$,
2. $|W_u \cap U| \leq 4q$, and
3. either $|\delta^+_H(W_u)| \leq k$ or $|\delta^-_H(W_u)| \leq k$.

By Claim 3.6, there is an SCC C in $G - F_u$ such that $|C \cap U| \geq |U| - 2q$. Since H is a k-FT sourcewise connectivity preserver of G with respect to U, set C is an SCC in $H - F_u$ as well. If $C_u = C$, then $|U| \leq |C_u \cap U| + 2q \leq 4q$, so we can let $W_u = V$. Otherwise, C_u and C are distinct SCC's in $H - F_u$. Then there exists an (C_u, C)-cut $(W_u, V - W_u)$ in $H - F_u$ such that either $\delta^+_{H-F_u}(W_u) = \emptyset$ or $\delta^-_{H-F_u}(W_u) = \emptyset$, so $|\delta^+_H(W_u)| \leq k$ or $|\delta^-_H(W_u)| \leq k$. Moreover, $|W_u \cap U| \leq 2q$, since $|S \cap U| \geq |U| - 2q$. We conclude that cut $(W_u, V - W_u)$ satisfies the claim.

Let $\mathcal{W} = \{W_u \mid (u, v) \in \mathcal{L}_v\}$. Let $\mathcal{W}^+ = \{W_u \in \mathcal{W} \mid \delta^+_{H-F_u}(W_u) = \emptyset\}$, and let $\mathcal{W}^- = \{W_u \in \mathcal{W} \mid \delta^-_{H-F_u}(W_u) = \emptyset\}$. Observe that $|\mathcal{W}^+| + |\mathcal{W}^-| \geq |\mathcal{L}_v|$. We now split our analysis into two cases.

- **Case 1: $|\mathcal{W}^+| \geq |\mathcal{L}_v|/2 > 2^{k+3}q$.** By Claim 4.9, there exists a subset $\mathcal{W}^+_1 \subseteq \mathcal{W}^+$ of size $|\mathcal{W}^+_1| \geq |\mathcal{W}^+|/(8q) > 2^k$ such that for all distinct $W_u, W_{u'} \in \mathcal{W}^+_1$, $u \notin W_{u'}$ and $u' \notin W_u$. Let $U^+_1 = \{u \in U \mid W_u \in \mathcal{W}^+_1\}$, and note that $|U^+_1| = |\mathcal{W}^+_1| > 2^k$. For each $u \in U^+_1$, we observe the following:

    1. $|\delta^+_H(W_u)| \leq k$,
    2. v can reach u in $H - \delta^+_H(W_u)$, since $u, v \in C_u \subseteq W_u$, and
    3. v cannot reach any vertex $u' \in U^+_1 - \{u\}$ in $H - \delta^+_H(W_u)$, since $u' \notin W_u$ if $u \neq u'$.

    See Figure 2 for a visualization. We can apply the anti-isolation lemma (Lemma 2.3) with respect to source vertex v, sink vertices U^+_1, and edge sets $\{\delta^+_H(W_u) \mid u \in U^+_1\}$ to argue that $|U^+_1| \leq 2^k$. However, this contradicts our earlier claim that $|U^+_1| > 2^k$. We conclude that $|\mathcal{L}| \leq 2^{k+3}qn$.

- **Case 2: $|\mathcal{W}^-| \geq |\mathcal{L}_v|/2 > 2^{k+3}q$.** This case is symmetric to Case 1. Specifically, we can obtain a contradiction in this case by considering the reverse graph of H and applying an identical argument as in Case 1. □

We finish our size analysis of H with the following claim.

Claim 4.11. $|E(H)| = O(4^k qn)$.

Proof. $|E(H)| \leq 2^{k+1} \cdot |\mathcal{L}| + O(2^k qn) \leq 2^{k+1} \cdot 2^{k+4} qn + O(2^k qn) = O(4^k qn)$, where the first inequality follows from Claim 4.7 and the second inequality follows from Claim 4.10. □

4.1.2 Applying the directed expander hierarchy

We will combine our fault-tolerant sourcewise connectivity preservers from Lemma 4.4 with the existential directed expander hierarchy of Lemma 3.7 to complete the construction of our fault-tolerant connectivity preserver claimed in Theorem 4.2.

Theorem 4.2. *For every positive integer k, every n-vertex directed graph admits a k-fault-tolerant connectivity preserver with $O(k4^k n \log n)$ edges.*

Proof. Let G be an n-vertex directed graph, and let k be a positive integer. Let $\{V_1, \dots, V_\ell\}$ be the directed expander hierarchy for $(2k, k)$-unbreakable sets specified in Lemma 3.7. For each $i \in [1, \ell]$, let $V_{\leq i} = V_1 \cup \dots \cup V_i$, and let $C_i^1, \dots, C_i^{j_i} \subseteq V_{\leq i}$ be the SCCs of $G[V_{\leq i}]$, for some $j_i \in [1, n]$. Additionally, for each $i \in [1, \ell]$ and $j \in [1, j_i]$, we let $U_i^j = C_i^j \cap V_i$.

For each $i \in [1, \ell]$ and $j \in [1, j_i]$ we build a k-FT sourcewise connectivity preserver H_i^j of graph $G[C_i^j]$ with respect to source vertices U_i^j using Lemma 4.4. We let $H_i = \cup_{j \in [1, j_i]} H_i^j$, and let $H = \cup_{i \in [1, \ell]} H_i$. We claim that H is a k-FT connectivity preserver of G with size $|E(H)| = O(k4^k \cdot n \log n)$.

First, we verify that $|E(H)| = O(k4^k \cdot n \log n)$. By Lemma 3.7, set U_i^j is $(2k, k)$-unbreakable in $G[C_i^j]$ for all $i \in [1, \ell]$ and $j \in [1, j_i]$. Then by Lemma 4.4, $|E(H_i^j)| = O(4^k k |C_i^j|)$. Then we can upper bound the number of edges in H as follows:

$$|E(H_i)| \leq \sum_{j \in [1, j_i]} |E(H_i^j)| \leq O(k4^k) \cdot \sum_{[1, j_i]} |C_i^j| = O(k4^k n).$$

Since there are at most $\ell = O(\log n)$ levels in our directed expander hierarchy, we conclude that

$$|E(H)| \leq \sum_{i \in [1, \ell]} |E(H_i)| \leq O(k4^k n \log n).$$

We now verify that H is a k-FT connectivity preserver of G. Fix a pair of vertices $s, t \in V$ and a fault set $F \subseteq E(G)$ of size $|F| \leq k$ such that s and t are strongly connected in $G - F$. Let $C \subseteq V(G)$ be the SCC in $G - F$ containing s and t. Let $i \in [1, \ell]$ be the largest index such that $V_i \cap C \neq \emptyset$, and let $x \in V_i \cap C$. Since $C \subseteq V_{\leq i}$, set C is contained in some SCC C_i^j in $G[V_{\leq i}]$, where $j \in [1, j_i]$. It follows that $x \in U_i^j$, since $V_i \cap C \subseteq V_i \cap C_i^j = U_i^j$. Recall that subgraph H_i^j is a sourcewise connectivity preserver of $G[C_i^j]$ with respect to source vertices U_i^j. Then since $x \in U_i^j$ is strongly connected to $s, t \in C_i^j$ in $G[C_i^j] - F$, it follows that x is strongly connected to s and t in $H_i^j - F$. We conclude that s and t are strongly connected in $H_i^j - F \subseteq H - F$, as claimed. $\square$

4.2 Constructing fault-tolerant connectivity preservers in FPT time

The goal of this section is to present our FPT-time construction of fault-tolerant connectivity preservers.

Theorem 4.12. *For every positive integer k, every n-vertex directed graph admits a k-FT connectivity preserver with $O(8^k n \log^{5/2} n)$ edges that can be computed in $2^{O(k)}\text{poly}(n)$ time, with high probability.*

We use a decremental greedy strategy to construct our k-FT connectivity preserver H. Initially, we let $H \leftarrow G$. Then we repeatedly identify an edge $e \in E(H)$ such that $H - e$ is a k-FT connectivity preserver of G. Once we find such an edge e, we can update H so that $H \leftarrow H - e$. Our algorithm terminates when $|E(H)| = O(8^k n \log^{5/2} n)$.

Under this framework, the main algorithmic subroutine of our construction is to find an edge in H that is safe to remove. If an edge is not safe to remove from H, then we say it is *k-fault critical.*

Definition 4.13 (k-fault critical)**.** Let G be an n-vertex directed graph. We say that an edge $e \in E(G)$ is k-fault critical in G if there exists a pair of vertices $s, t \in V(G)$ and a fault set $F \subseteq E(G)$ of size $|F| \leq k$ such that s and t are strongly connected in $G - F$, but s and t are not strongly connected in $(G - e) - F$.

We are able to adapt the tools from our proof of Theorem 1.1 to efficiently compute a small set of edges $E' \subseteq E(G)$ that contains all k-fault critical edges in G.

Lemma 4.14. *Let G be an n-vertex directed graph. For every positive integer k, there exists an edge set $E' \subseteq E(G)$ of size $|E'| = O(8^k n \log^{5/2} n)$ that contains every edge in G that is k-fault critical. Moreover, E' can be computed with high probability in time $2^{O(k)} \cdot \mathrm{poly}(n)$.*

Lemma 4.14 directly implies Theorem 4.12.

Proof of Theorem 4.12. We initialize our k-FT connectivity preserver H of G to be $H \leftarrow G$. If $|E(H)| = O(8^k n \log^{5/2} n)$, then we are done. Otherwise, we can find an edge $e \in E(H)$ that is not k-fault critical in H, with high probability in time $2^{O(k)} \cdot \mathrm{poly}(n)$, by Lemma 4.14. Since $e \in E(H)$ is not k-fault critical in H, subgraph $H - e$ is a k-fault connectivity preserver of H. Moreover, since H is a k-FT connectivity preserver of G, it follows that $H - e$ is a k-FT connectivity preserver of G as well. We update H so that $H \leftarrow H - e$.

In general, given a subgraph H of G that is a k-FT connectivity preserver of G, either $|E(H)| = O(8^k n \log^{5/2} n)$ or we can apply Lemma 4.14 on H to find an edge $e \in E(H)$ such that $H - e$ is a k-FT connectivity preserver of G as well. Our algorithm terminates and outputs H when $|E(H)| = O(8^k n \log^{5/2} n)$. This algorithm outputs a k-FT connectivity preserver H of G of size $|E(H)| = O(8^k n \log^{5/2} n)$ with high probability in time $2^{O(k)} \cdot \mathrm{poly}(n)$. □

For the remainder of this section, we focus on proving Lemma 4.14.

4.2.1 k-fault critical edges for sourcewise connectivity

Similar to the proof of Theorem 4.2, we will begin our proof of Lemma 4.14 by first considering fault-tolerant *sourcewise* connectivity. We will need the following definition.

Definition 4.15 (k-fault critical with respect to P)**.** Let $G = (V, E)$ be an n-vertex directed graph. We say that an edge $e \in E$ is k-fault critical in G with respect to vertex pairs $P \subseteq V \times V$ if there exists $(s, t) \in P$ and a fault set $F \subseteq E$ of size $|F| \leq k$ such that s and t are strongly connected in $G - F$, but s and t are not strongly connected in $(G - e) - F$.

We will prove the following algorithmic lemma about k-fault critical edges.

Lemma 4.16. *Let G be an n-vertex, m-edge directed graph, and let $U \subseteq V(G)$ be a $(q, 2^k)$-unbreakable set in G. There exists an edge set $E' \subseteq E(G)$ of size $|E'| = O(2^k q^2 n + 2^k qn \log n)$ that contains every edge $e \in E(G)$ that is k-fault critical with respect to $U \times V(G)$. Moreover, E' can be computed with high probability in time $O(q^2 2^k mn + qn \log(n) \cdot (k2^k m + k^2 4^k))$.*

We will use Lemma 4.16, combined with a directed expander hierarchy, to prove Lemma 4.14 and complete the proof of Theorem 4.12. We first present the construction of E', and then we analyze the time complexity and correctness of our construction.

Construction of E'. Let $G = (V, E)$ be an n-vertex directed graph, and let $U \subseteq V$ be a $(q, 2^k)$-unbreakable set in G. Let $U' \subseteq U$ be an arbitrary subset of U of size $|U'| = \min(|U|, 5q^2)$. Let $J \subseteq G$ be the subgraph of G obtained by unioning $|U'|$ many k-FT single-source connectivity preservers of G, each rooted at a distinct vertex in U'. Formally, let $J = \bigcup_{u \in U'} \text{SSCP}(G, u, k)$. Subgraph $E(J)$ will be part of our final edge set E'.

We define a collection of $\lambda = 50 \log n$ subsets of U as follows. For $j \in [1, \lambda]$, let $Q_j \subseteq U$ be an element of $\binom{U}{q}$ sampled uniformly at random. We will use the collection of sets $\{Q_j\}_{j \in [1,\lambda]}$ to construct our final edge set E'.

Let $V = \{v_1, \dots, v_n\}$. For each $i \in [1, n]$, we will define a set U_i and a collection of edges $E_i \subseteq E$ as follows. For each $j \in [1, \lambda]$, if $v_i \notin Q_j$, then let $(U_i^j, V - U_i^j) = \text{ImpCut}(G, v_i, Q_j, k, +)$ and let $(W_i^j, V - W_i^j) = \text{ImpCut}(G, v_i, Q_j, k, -)$. Otherwise, if $v_i \in Q_j$, then let $U_i^j = W_i^j = \emptyset$. Let $S_i = \bigcup_{j \in [1,\lambda]} U_i^j \cup W_i^j$, and let $U_i = S_i \cap U$.

With set $U_i \subseteq U$ defined, we are ready to define edge set $E_i \subseteq E$. For each $u \in U_i$, let $(S_i^u, V - S_i^u) = \text{ImpCut}(G, u, v_i, k + 1, +)$ and let $(T_i^u, V - T_i^u) = \text{ImpCut}(G, u, v_i, k + 1, -)$. We define edge set E_i to be $E_i = \cup_{u \in U_i} \delta_G^+(S_i^u) \cup \delta_G^-(T_i^u)$.

We finish our construction of E' by letting $E' = E(J) \cup_{i \in [1,n]} E_i$. Our construction is summarized in Algorithm 1.

Time Complexity. We now verify that Algorithm 1 has the time complexity claimed in Lemma 4.16. We will first need the following result bounding the size of U_i.

Claim 4.17. *For each $i \in [1, n]$, $|U_i| \leq 2\lambda q$.*

Proof. We claim that for each $j \in [1, \lambda]$, $|U_i^j \cap U| < q$ and $|W_i^j \cap U| < q$. Note that by Lemma 4.1, $|\delta_G^+(U_i^j)| \leq 2^k$ and $|\delta_G^-(W_i^j)| \leq 2^k$. Additionally, $Q_j \subseteq V - U_i^j$ and $Q_j \subseteq V - W_i^j$, so $|(V - U_i^j) \cap U| \geq q$ and $|(V - W_i^j) \cap U| \geq q$. Then if $|U_i^j \cap U| \geq q$ or $|W_i^j \cap U| \geq q$, it would imply that U is $(q, 2^k)$-breakable, contradicting our assumption that U is $(q, 2^k)$-unbreakable in G. □

Using Claim 4.17, we can now bound the time complexity of our construction of H.

Claim 4.18. *Algorithm 1 runs in time $O(q^2 2^k mn + qn \log(n) \cdot (k2^k m + k^2 4^k))$.*

Proof. We can compute a k-FT single-source (symmetric) connectivity preserver in time $O(2^k mn)$ by Theorem 2 of [BCR18]. Consequently, subgraph J can be computed in $O(q^2 2^k mn)$ time. Sets $Q_1, \dots, Q_\lambda$ can be computed in $O(\lambda qn)$ time.

For each vertex $v_i \in V(G)$, we make $O(\lambda + |U_i|)$ calls to $\text{ImpCut}(G, \cdot, \cdot, k + 1, \cdot)$. Each call can be completed in time $O(k2^k m + k^2 4^k)$ by Lemma 4.1. By Claim 4.17, $|U_i| \leq 2\lambda q$. Then for each vertex $v_i \in V(G)$, we spend time $O((\lambda + |U_i|) \cdot (k2^k m + k^2 4^k)) = O(\lambda q \cdot (k2^k m + k^2 4^k))$. Then the total time complexity contributed by the for loop in our construction is $O(\lambda qn(k2^k m + k^2 4^k))$.

We conclude that our construction of E' runs in time $O(q^2 2^k mn + qn \log(n) \cdot (k2^k m + k^2 4^k))$. □

Algorithm 1 Construction of edge set E' from Lemma 4.16

Input: Directed graph $G = (V, E)$ and $(q, 2^k)$-unbreakable set $U \subseteq V$.

$J \leftarrow (V, \emptyset)$
Let $U' \subseteq U$ be a subset of U of size $|U'| = \min(5q^2, |U|)$.
$J \leftarrow \cup_{u \in U'} \mathrm{SSCP}(G, u, k)$
Let $\lambda = 50 \log n$
for $j \in [1, \lambda]$ **do**
  Let $Q_j \subseteq U$ be an element of $\binom{U}{q}$ sampled uniformly at random.
end for
Let $V = \{v_1, \ldots, v_n\}$.
for $i \in [1, n]$ **do**
  for $j \in [1, \lambda]$ **do**
    if $v_i \notin Q_j$ **then**
      $(U_i^j, V - U_i^j) \leftarrow \mathrm{ImpCut}(G, v_i, Q_j, k, +)$.
      $(W_i^j, V - W_i^j) \leftarrow \mathrm{ImpCut}(G, v_i, Q_j, k, -)$.
    else
      $U_i^j \leftarrow \emptyset$, $W_i^j \leftarrow \emptyset$
    end if
  end for
  $S_i \leftarrow \bigcup_{j \in [1, \lambda]} U_i^j \cup W_i^j$
  $U_i \leftarrow S_i \cap U$
  for $u \in U_i$ **do**
    $(S_i^u, V - S_i^u) \leftarrow \mathrm{ImpCut}(G, u, v_i, k+1, +)$.
    $(T_i^u, V - T_i^u) \leftarrow \mathrm{ImpCut}(G, u, v_i, k+1, -)$.
  end for
  $E_i \leftarrow \cup_{u \in U_i} \delta_G^+(S_i^u) \cup \delta_G^-(T_i^u)$
end for
$E' \leftarrow E(J) \cup_{i \in [1, n]} E_i$

Correctness. We now verify that our edge set E' contains all edges in $E(G)$ that cannot be safely removed from G. This property of E' is summarized in the following lemma.

Lemma 4.19. *For all $e \in E(G)$, if e is k-fault critical with respect to $U \times V(G)$, then $e \in E'$.*

Fix an edge $e = (v_i, v_j) \in E(G)$. We will show that if e is k-fault critical with respect to $U \times V(G)$, then $e \in E'$. We start with the following structural claim.

Claim 4.20. *If $e = (v_i, v_j) \in E(G)$ is k-fault critical with respect to $U \times V(G)$, then there exists a vertex $u_e \in U$ and a fault set $F_e \subseteq E(G)$ of size $|F_e| \leq k$ such that vertices v_i, v_j, and u_e are strongly connected in $G - F_e$, and $G - F_e$ contains either*

- *an out-reachable (u_e, v_j)-cut $(S, V - S)$ such that $\delta^+_{G-F_e}(S) = \{e\}$, or*
- *an in-reachable (u_e, v_i)-cut $(S, V - S)$ such that $\delta^-_{G-F_e}(S) = \{e\}$.*

Proof. Claim 4.20 follows from an argument identical to the proof of Claim 4.6. □

Let X_e be the SCC in $G - F_e$ containing vertices v_i, v_j, u_e.

Claim 4.21. *In the proof of Lemma 4.19, we may assume that $|X_e \cap U| \leq 2q$ and $|U| > 5q^2$, without loss of generality.*

Proof. By Claim 4.5, if $|X_e \cap U| > 2q$, then v_i, v_j, and u_e are strongly connected in $J - F_e$. However, v_i, v_j, and u_e are not all strongly connected in $(J - e) - F_e$, so $e \in E(J) \subseteq E(H)$. Then in the proof of Lemma 4.19, we may assume without loss of generality that $|X_e \cap U| \leq 2q$. Additionally, if $|U| \leq 5q^2$, then $u_e \in U'$, so v_i, v_j, and u_e are strongly connected in $J - F_e$. Again this implies that $e \in E(J) \subseteq E(H)$. Then we may assume without loss of generality that $|U| > 5q^2$. ☐

We may assume without loss of generality that $|X_e \cap U| \leq 2q$ and $|U| > 5q^2$ by Claim 4.21. Then by Claim 3.6, there is an SCC $Z_e \neq X_e$ in $G - F_e$ such that $|Z_e \cap U| \geq |U| - 2q$. Then there is an (X_e, Z_e)-cut $(Y_e, V - Y_e)$ in $G - F_e$ such that either $\delta^+_{G-F_e}(Y_e) = \emptyset$ or $\delta^-_{G-F_e}(Y_e) = \emptyset$, so $|\delta^+_G(Y_e)| \leq k$ or $|\delta^-_G(Y_e)| \leq k$. Moreover, $|Y_e \cap U| \leq 2q$, since $|Z_e \cap U| \geq |U| - 2q$.

Claim 4.22. *With high probability, at least one of the sets $Q_j \subseteq U$, $j \in [1, \lambda]$, sampled during the construction of H satisfies $Q_j \cap Y_e = \emptyset$.*

Proof. We first observe that for a fixed index $j \in [1, \lambda]$, $Q_j \cap Y_e = \emptyset$ with constant probability:

$$\Pr[Q_j \cap Y_e \neq \emptyset] \leq \sum_{u \in Y_e \cap U} \Pr[u \in Q_j] \leq |Y_e \cap U| \cdot \frac{\binom{|U|}{q-1}}{\binom{|U|}{q}} \leq 2q \cdot \frac{q}{|U| - q + 1} \leq \frac{4q^2}{|U|} \leq \frac{4}{5}.$$

Then the probability that none of our sampled sets $Q_j \subseteq U$, $j \in [1, \lambda]$ are disjoint from Y is at most

$$\left(\frac{4}{5}\right)^{\lambda} \leq e^{-\lambda/5} = e^{-10 \log n} \leq n^{-10}.$$

Then with high probability, there exists a set Q_j, $j \in [1, \lambda]$, such that $Q_j \cap Y_e = \emptyset$. ☐

With the above discussion in mind, we can now prove the following claim about edges in $E(G)$ that are k-fault critical with respect to $U \times V(G)$.

Claim 4.23. *If edge $e = (v_i, v_j) \in E(G)$ is k-fault critical with respect to vertex pair $(u_e, v) \in U \times V(G)$, then $u_e \in U_i \cap U_j$.*

Proof. By the earlier discussion, there exists an out-minimal (X_e, Z_e)-cut $(Y_e, V(G) - Y_e)$ in $G - F_e$ such that $|Y_e \cap U| \leq 2q$, and either $\delta^+_{G-F_e}(Y_e) = \emptyset$ or $\delta^-_{G-F_e}(Y_e) = \emptyset$. Moreover, by Claim 4.22, there exists a $\ell \in [1, \lambda]$ such that $Q_\ell \subseteq V(G) - Y_e$, with high probability.

We now split our analysis into two cases based on whether $\delta^+_{G-F_e}(Y_e) = \emptyset$ or $\delta^-_{G-F_e}(Y_e) = \emptyset$.

- If $\delta^+_{G-F_e}(Y_e) = \emptyset$, then $\delta^+_G(Y_e) \subseteq F_e$, so $|\delta^+_G(Y_e)| \leq k$. We observe that $(Y_e, V - Y_e)$ is a (v_i, Q_ℓ)-cut as well. Since v_i and u_e lie in the same SCC $X_e \subseteq Y_e$ in $G - F_e$, we conclude that there exists an *out-reachable* (v_i, Q_ℓ)-cut $(Y_e^*, V - Y_e^*)$ in G of size $|\delta^+_G(Y_e^*)| \leq k$ such that $u_e \in X_e \subseteq Y_e^*$. Then by Lemma 4.1, $Y_e^* \subseteq U_i^\ell$, so $u_e \in U_i^\ell \cap U \subseteq U_i$. By an identical argument, $u_e \in U_j$.

- If $\delta^-_{G-F_e}(Y_e) = \emptyset$, then $\delta^-_G(Y_e) \subseteq F_e$, so $|\delta^-_G(Y_e)| \leq k$. We observe that $(Y_e, V - Y_e)$ is a (v_i, Q_ℓ)-cut as well. Since v_i and u_e lie in the same SCC $X_e \subseteq Y_e$ in $G - F_e$, we conclude that there exists an *in-reachable* (v_i, Q_ℓ)-cut $(Y^*_e, V - Y^*_e)$ in G of size $|\delta^-_G(Y^*_e)| \leq k$ such that $u_e \in X_e \subseteq Y^*_e$. Then by Lemma 4.1, $Y^*_e \subseteq W^\ell_i$, so $u_e \in W^\ell_i \cap U \subseteq U_i$. By an identical argument, $u_e \in U_j$.

We conclude that $u_e \in U_i \cap U_j$, as claimed. □

We are ready to finish the proof of Lemma 4.19.

Lemma 4.19. *For all $e \in E(G)$, if e is k-fault critical with respect to $U \times V(G)$, then $e \in E'$.*

Proof of Lemma 4.19. If $e = (v_i, v_j)$ is k-fault critical with respect to $U \times V(G)$, then by Claim 4.20, there exists a vertex $u_e \in U$ and a fault set $F_e \subseteq E(G)$ of size $|F_e| \leq k$ such that vertices v_i, v_j, and u_e are strongly connected in $G - F_e$, and $G - F_e$ contains either

- an out-reachable (u_e, v_j)-cut $(S, V - S)$ such that $\delta^+_{G-F_e}(S) = \{e\}$, or
- an in-reachable (u_e, v_i)-cut $(S, V - S)$ such that $\delta^-_{G-F_e}(S) = \{e\}$.

Moreover, by Claim 4.23, $u_e \in U_i \cap U_j$. Now we split our analysis into two cases.

- **Case 1:** $G - F_e$ contains an out-reachable (u_e, v_j)-cut $(S, V - S)$ such that $\delta^+_{G-F_e}(S) = \{e\}$. Then $(S, V - S)$ is an out-reachable (u_e, v_j)-cut in G of size $|\delta^+_G(S)| \leq k + 1$. Recall that in our construction of set E_j in Algorithm 1, $(S^{u_e}_j, V - S^{u_e}_j) = \mathrm{ImpCut}(G, u_e, v_j, k + 1, +)$. We conclude that $S \subseteq S^{u_e}_j$ by Lemma 4.1, so $e \in \delta^+_G(S^{u_e}_j) \subseteq E_j \subseteq E'$.
- **Case 2:** $G - F_e$ contains an in-reachable (u_e, v_i)-cut $(S, V - S)$ such that $\delta^-_{G-F_e}(S) = \{e\}$. Then $(S, V - S)$ is an in-reachable (u_e, v_i)-cut in G of size $|\delta^-_G(S)| \leq k + 1$. Recall that in our construction of set E_i in Algorithm 1, $(T^{u_e}_i, V - T^{u_e}_i) = \mathrm{ImpCut}(G, u_e, v_i, k + 1, -)$. We conclude that $S \subseteq T^{u_e}_i$ by Lemma 4.1, so $e \in \delta^-_G(T^{u_e}_i) \subseteq E_i \subseteq E'$.

We conclude that $e \in E'$. □

Size Analysis. We now conclude the proof of Lemma 4.16 by bounding $|E'|$.

Claim 4.24. $|E'| = O(2^k q^2 n + 2^k qn \log n)$.

Proof. Since J is the union of $|U'| = O(q^2)$ many k-FT single-source connectivity preservers, $|E(J)| = O(2^k q^2 n)$ by Theorem 2 of [BCR18].

We conclude that

$$\begin{aligned}
|E'| &= |E(J)| + \sum_{i \in [1,n]} |E_i| \\
&\leq O(2^k q^2 n) + \sum_{i \in [1,n]} \sum_{u \in U_i} \left(|\delta^+_G(S^u_i)| + |\delta^-_G(T^u_i)|\right) \\
&\leq O(2^k q^2 n) + \sum_{i \in [1,n]} \sum_{u \in U_i} 2^{k+1} && \text{by Lemma 4.1} \\
&\leq O(2^k q^2 n) + 2^{k+1} \cdot 2\lambda q \cdot n && \text{by Claim 4.17} \\
&\leq O(2^k q^2 n + 2^k qn \log n),
\end{aligned}$$

as claimed. □

4.2.2 Applying the directed expander hierarchy

We will combine Lemma 4.16 with the polynomial directed expander hierarchy construction of Lemma 3.7 to complete the proof of Lemma 4.14.

Lemma 4.14. *Let G be an n-vertex directed graph. For every positive integer k, there exists an edge set $E' \subseteq E(G)$ of size $|E'| = O(8^k n \log^{5/2} n)$ that contains every edge in G that is k-fault critical. Moreover, E' can be computed with high probability in time $2^{O(k)} \cdot \mathrm{poly}(n)$.*

Proof. Let G be an n-vertex directed graph, and let k be a positive integer. Let $\{V_1, \dots, V_\ell\}$ be the directed expander hierarchy for $(\Theta(2^k\sqrt{\log n}), 2^k)$-unbreakable sets specified in Lemma 3.7. Since $q = \Omega(2^k\sqrt{\log n})$, this expander hierarchy can be computed in polynomial time. For each $i \in [1, \ell]$, let $V_{\leq i} = V_1 \cup \cdots \cup V_i$, and let $S_i^1, \dots, S_i^{j_i} \subseteq V_{\leq i}$ be the SCCs of $G[V_{\leq i}]$, for some $j_i \in [1, n]$. Additionally, for each $i \in [1, \ell]$ and $j \in [1, j_i]$, we let $U_i^j = S_i^j \cap V_i$.

For each $i \in [1, \ell]$ and $j \in [1, j_i]$ we apply Lemma 4.16 on graph $G[S_i^j]$ and $(\Theta(2^k\sqrt{\log n}), 2^k)$-unbreakable set U_i^j. Then in $2^{O(k)} \cdot \mathrm{poly}(n)$ time, with high probability, we can compute an edge set $E_i^j \subseteq G[S_i^j]$ that contains every edge in $G[S_i^j]$ that is k-fault critical in $G[S_i^j]$ with respect to vertex pairs $U_i^j \times S_i^j$.

We let $E_i = \cup_{j \in [1, j_i]} E_i^j$, and let $E' = \cup_{i \in [1,\ell]} E_i$. Set E' can be computed in $2^{O(k)} \cdot \mathrm{poly}(n)$ time. We claim that, with high probability, edge set E' contains every edge in G that is k-fault critical, and $|E'| = O(8^k n \log^{5/2} n)$.

First, we verify that $|E'| = O(8^k n \log^{5/2} n)$. By Lemma 3.7, set U_i^j is $(\Theta(2^k\sqrt{\log n}), 2^k)$-unbreakable in $G[S_i^j]$ for all $i \in [1, \ell]$ and $j \in [1, j_i]$. Then by Lemma 4.16, $|E_i^j| = O(8^k \log^{3/2} n |S_i^j|)$. Then we can upper bound the number of edges in H as follows:

$$|E_i| \leq \sum_{j \in [1, j_i]} |E_i^j| \leq O(8^k \log^{3/2} n) \cdot \sum_{[1, j_i]} |S_i^j| = O(8^k n \log^{3/2} n).$$

Since there are at most $\ell = O(\log n)$ levels in our directed expander hierarchy, we conclude that

$$|E'| \leq \sum_{i \in [1,\ell]} |E_i| \leq O(8^k n \log^{5/2} n).$$

We now verify that E' contains every edge in G that is k-fault critical, with high probability. Fix an edge $e \in E(G)$ that is k-fault critical in G. Then there exists a pair of vertices $s, t \in V$ and a fault set $F \subseteq E(G)$ of size $|F| \leq k$ such that s and t are strongly connected in $G - F$, but not strongly connected in $(G - e) - F$. Let $C \subseteq V(G)$ be the SCC in $G - F$ containing s and t. Let $i \in [1, \ell]$ be the largest index such that $V_i \cap C \neq \emptyset$, and let $x \in V_i \cap C$. Since $C \subseteq V_{\leq i}$, set C is contained in some SCC S_i^j in $G[V_{\leq i}]$, where $j \in [1, j_i]$. Then x is strongly connected to vertices s and t in $G[S_i^j] - F$. However, x cannot be strongly connected to both s and t in $(G[S_i^j] - e) - F$, since s and t are not strongly connected in $(G - e) - F$. We conclude that edge e is k-fault critical in $G[S_i^j]$ with respect to vertex pairs $\{x\} \times \{s, t\}$. Recall that edge set E_i^j contains every edge that is k-fault critical in $G[S_i^j]$ with respect to vertex pairs $U_i^j \times S_i^j$, with high probability. Since $x \in C \cap V_i \subseteq S_i^j \cap V_i = U_i^j$ and $s, t \in C \subseteq S_i^j$, we conclude that edge $e \in E_i^j \subseteq E'$ with high probability.

Then by the union bound, E' contains every edge in G that is k-fault critical, with high probability. $\square$

5 k-Connectivity Preservers

In this section, we present our upper bound on k-connectivity preservers in Theorem 1.3.

Theorem 1.3. *For every positive integer k, every n-vertex directed graph admits a k-connectivity preserver with $O(n\sqrt{nk})$ edges. The preserver can be constructed in polynomial time.*

Recall that a subgraph $H \subseteq G$ is a k-connectivity preserver of G if for all $s, t \in V(G)$, $\lambda_H^k(s,t) = \lambda_G^k(s,t)$, where $\lambda_G^k(s,t) = \min(\lambda_G(s,t), k)$ is the k-bounded connectivity between s, t in G.

Lemma 5.1 (transitivity)**.** *For three graphs $H' \subseteq H \subseteq G$, if H' is a k-connectivity preserver of H, and H is a k-connectivity preserver of G, then H' is a k-connectivity preserver of G.*

Proof. For each pair of vertices $s, t \in V(G)$, by the definition of k-connectivity preservers we have $\lambda_{H'}^k(s,t) = \lambda_H^k(s,t) = \lambda_G^k(s,t)$. Therefore, H' is a k-connectivity preserver of G. $\square$

Our proof of Theorem 1.3 requires the following definition.

Definition 5.2 (k-critical graphs)**.** We say that a graph H is *k-critical* if no strict edge subgraph $H' \subset H$ of H is a k-connectivity preserver of H.

We prove Theorem 1.3 by establishing the following upper bound on the number of edges in an n-vertex, k-critical directed graph.

Theorem 5.3. *Let H be an n-vertex, k-critical directed graph. Then $|E(H)| = O(k^{1/2}n^{3/2})$.*

We quickly verify that Theorem 5.3 implies Theorem 1.3.

Proof of Theorem 1.3. Let $G = (V, E)$ be an n-vertex directed graph, and let $H \subseteq G$ be an edge-minimal k-connectivity preserver of G. By Lemma 5.1 and the minimality of H, there does not exist a subgraph $H' \subsetneq H$ that is a k-connectivity preserver of H. Then H is an n-vertex, k-critical graph. By Theorem 5.3, $|E(H)| = O(k^{1/2}n^{3/2})$. $\square$

We will prove Theorem 5.3 using a two-level decomposition with regard to (q,k)-unbreakability. Let H be an n-vertex, k-critical directed graph. Let $q \in [1, n]$ be a parameter of the decomposition that we will optimize later. We maintain a partition $\mathcal{S}$ of vertex set V and a collection of cuts $\mathcal{C}$ as follows. Initially, we let $\mathcal{S} = \{V\}$ and $\mathcal{C} = \emptyset$. While there exists a set $S \in \mathcal{S}$ and a cut (L, R) in H of size at most $\min(|\delta_H^+(L)|, |\delta_H^-(L)|) \leq k$ such that $|S \cap L| \geq q$ and $|S \cap R| \geq q$, we remove set S from $\mathcal{S}$ and we add sets $S \cap L$ and $S \cap R$ to $\mathcal{S}$. Additionally, we add cut (L, R) to collection $\mathcal{C}$. Note that $\mathcal{S}$ remains a partition of V after this operation. We repeat this process until no such set $S \in \mathcal{S}$ and cut (L, R) exists.

Claim 5.4. *This process must terminate. After the above process terminates, every set $U \in \mathcal{S}$ is (q,k)-unbreakable in H. Moreover, $|\mathcal{C}| \leq n/q$.*

Proof. This process must terminate, since $|\mathcal{S}|$ increases by one after each iteration. By the terminating condition of this procedure, every set $U \in \mathcal{S}$ is (q,k)-unbreakable in H. Additionally, every set $U \in \mathcal{S}$ has size at least $|U| \geq q$. We conclude that $|\mathcal{S}| \leq n/q$, and therefore $|\mathcal{C}| \leq |\mathcal{S}| \leq n/q$. $\square$

We will bound the number of edges inside each set in $\mathcal{S}$ in Lemma 5.5, and we will bound the number of edges crossing sets in $\mathcal{S}$ using Lemma 5.8.

Lemma 5.5. *Consider a k-critical directed graph H. For every subset $U \subseteq V(H)$, if U is (q,k)-unbreakable in H, then the induced subgraph $H[U]$ has at most $O((q+k)|U|)$ edges.*

Proof. Since H is k-critical, for each edge $e \in E(H)$, graph $H - e$ is not a k-connectivity preserver of H. Therefore, there exists a pair of nodes $s,t \in V(H)$ such that $\lambda^k_{H-e}(s,t) < \lambda^k_H(s,t)$. Thus, for each edge $e \in E(H)$, there exists a minimum (s,t)-cut (L_e, R_e) of size $|\delta^+_H(L_e)| \leq k$ such that $e \in \delta^+_H(L_e)$.

Now we will show that for each edge $e = (u,v) \in H[U]$,

$$\min(\text{outdeg}_{H[U]}(u), \text{indeg}_{H[U]}(v)) \leq q + k.$$

Suppose towards contradiction that $\min(\text{outdeg}_{H[U]}(u), \text{indeg}_{H[U]}(v)) > q + k$. Then since (L_e, R_e) is a cut of size at most $|\delta^+_H(L_e)| \leq k$ and $e \in \delta^+_H(L_e)$, this implies that

$$|L_e \cap S| \geq |N^+_{H[U]}(u) \cap L_e| \geq |N^+_{H[U]}(u)| - k > q$$

and

$$|R_e \cap U| \geq |N^-_{H[U]}(v) \cap R_e| \geq |N^-_{H[U]}(v)| - k > q.$$

However, this contradicts our assumption that set U is (q,k)-unbreakable in H. We conclude that for each edge $(u,v) \in H[U]$, $\min(\text{outdeg}_{H[U]}(u), \text{indeg}_{H[U]}(v)) \leq q+k$, as claimed.

Using this inequality, it becomes straightforward to prove that the total number of edges in $H[U]$ is at most $|E(H[U])| \leq 2(q+k)|U|$. For each edge $(u,v) \in E(H[U])$, if $\text{outdeg}_{H[U]}(u) \leq q+k$, then we charge edge (u,v) to vertex $u \in U$. Otherwise, if $\text{indeg}_{H[U]}(v) \leq q+k$, then we charge edge (u,v) to vertex $v \in U$. Under this charging scheme each vertex in U is charged at most $2(q+k)$ times, so the total number of edges in $H[U]$ is at most $|E(H[U])| \leq 2(q+k)|U|$. ◻

We also need the next lemma, which reduces the task of constructing an all-pairs k-connectivity preserver to preserving k-connectivity between just $O(n)$ demand pairs. This lemma follows from well-known results about Gomory-Hu trees [GH61, CH91].

Lemma 5.6. *For any n-vertex directed graph $H = (V,E)$, there exists a collection of vertex demand pairs $P \subseteq V \times V$ of size $|P| = O(n)$ such that for any edge subgraph $H' \subseteq H$, H' is a k-connectivity preserver of H if and only if for every demand pair $(s,t) \in P$, $\lambda^k_{H'}(s,t) = \lambda^k_H(s,t)$.*

Proof. The "only if" direction is trivial, so we focus on proving the "if" direction below. Let C_H be a complete undirected weighted graph with vertex set V. We assign each edge $(u,v) \in E(C_H)$ weight $\lambda_H(u,v)$. Let T be a maximum spanning tree of C_H.

For any edge $(u,v) \in E(C_H)$, consider the (u,v)-path $Q = (u_0, u_1, \dots, u_\ell)$ in T, where $u = u_0$ and $v = u_\ell$. We claim that $\lambda_H(u,v) = \min_{(s,t)\in Q} \lambda_H(s,t)$. On one side, $\lambda_H(u,v) \geq \min_{(s,t)\in Q} \lambda_H(s,t)$, by the triangle inequality of symmetric connectivity. On the other side, $\lambda_H(u,v) \leq \min_{(s,t)\in Q} \lambda_H(s,t)$, since T is a maximum spanning tree in C_H.

We define our set of vertex demand pairs $P \subseteq V \times V$ to be the edges of tree T, so that $P = E(T)$. Now suppose that $H' \subseteq H$ is an edge subgraph of H that satisfies $\lambda^k_{H'}(s,t) = \lambda^k_H(s,t)$ for each demand pair $(s,t) \in P$. Then by the triangle inequality of symmetric connectivity, we have

$$\lambda^k_{H'}(u,v) \geq \min_{(s,t)\in Q} \lambda^k_{H'}(s,t) = \min_{(s,t)\in Q} \lambda^k_H(s,t) = \lambda^k_H(u,v).$$

Therefore, H' preserves the k-bounded connectivity for all vertex pairs (u,v) in H. ◻

We will need the following lemma in order to eventually bound the number of edges between components in $\mathcal{S}$.

Lemma 5.7. *Let H be an n-vertex, k-critical directed graph. For every cut (L, R) with size $|\delta^+(L)| \le k$, we have $|\delta^-(L)| = O(kn)$.*

Proof. Let $P \subseteq V(H) \times V(H)$ be the set of $|P| = O(n)$ demand pairs specified in Lemma 5.6. For each demand pair $(s,t) \in P$, there exists a collection $\mathcal{Q}_{s,t}$ of $\lambda^k_H(s,t)$ pairwise edge-disjoint paths from s to t in H, and there exists a collection $\mathcal{Q}_{t,s}$ of $\lambda^k_H(s,t)$ pairwise edge-disjoint paths from t to s in H. For each path $Q \in \mathcal{Q}_{s,t} \cup \mathcal{Q}_{t,s}$, it must pass through edges in $\delta^+_H(L)$ and edges in $\delta^-_H(L)$ interleavingly, which means the number of edges in $\delta^-_H(L)$ on the path is at most the number of edges in $\delta^+_H(L)$ on the path, plus one. Since the paths in $\mathcal{Q}_{s,t}$ are pairwise edge-disjoint, the paths in $\mathcal{Q}_{s,t}$ contain at most $|\delta^+_H(L)| + \lambda^k_H(s,t) \le 2k$ edges in $\delta^-_H(L)$. An identical argument implies that the paths in $\mathcal{Q}_{t,s}$ contain at most $2k$ edges in $\delta^-_H(L)$ as well.

Let $H' \subseteq H$ be the subgraph of H obtained by unioning all the paths in all collections $\mathcal{Q}_{s,t} \cup \mathcal{Q}_{t,s}$ for all $(s,t) \in P$, so that $H' = \cup_{\{Q \in \mathcal{Q}_{s,t} \cup \mathcal{Q}_{t,s} \mid (s,t) \in P\}} Q$. We observe that $\lambda^k_{H'}(s,t) = \lambda^k_H(s,t)$. Graph H' is the union of $2|P|$ collections of paths $\mathcal{Q}_{s,t}$, and the paths in $\mathcal{Q}_{s,t}$ contain at most $2k$ edges in $\delta^-_H(L)$, by the earlier discussion. We conclude that there are at most $|\delta^-_{H'}(L)| \le 2k \cdot 2|P| = O(kn)$ edges in $\delta^-_{H'}(L)$. Now observe that for any demand pair $(s,t) \in P$, $\lambda^k_{H'}(s,t) = \lambda^k_H(s,t)$, so H' is a k-connectivity preserver of H by Lemma 5.6. Finally, since H is k-critical, we know $H' = H$ and thus $|\delta^-_H(L)| = |\delta^-_{H'}(L)| = O(kn)$. ◻

We can now bound the number of cross-component edges in $H[S]$.

Lemma 5.8. *Let $E' = \{(u,v) \in E(H[S]) \mid (u,v) \in S_1 \times S_2, S_1 \ne S_2 \in \mathcal{S}\}$ be the set of edges in $H[S]$ between different components in $\mathcal{S}$. Then $|E'| = O(kn^2/q)$.*

Proof. Recall that $\mathcal{C}$ is the family of all cuts (L,R) that we considered in our construction of partition $\mathcal{S}$. By our construction of $\mathcal{S}$, every edge $e \in E'$ crosses a cut (L,R) in cut family $\mathcal{C}$ (i.e., $e \in \delta^+(L) \cup \delta^-(L)$ for some cut $(L,R) \in \mathcal{C}$).

By Claim 5.4, $|\mathcal{C}| \le n/q$, and by Lemma 5.7, we have that $|\delta^+(L) \cup \delta^-(L)| = O(kn)$ for each cut $(L,R) \in \mathcal{C}$. Therefore,

$$|E'| \le \sum_{(L,R) \in \mathcal{C}} |\delta^+(L) \cup \delta^-(L)| \le O((n/q) \cdot kn) = O(kn^2/q).$$

◻

Proof of Theorem 5.3. By Claim 5.4 and Lemma 5.5, we know that the total number of edges in $\cup_{S \in \mathcal{S}} H[S]$ is at most

$$\left| \bigcup_{S \in \mathcal{S}} H[S] \right| \le \sum_{S \in \mathcal{S}} O((q+k)|S|) = O((q+k)n).$$

Additionally, by Lemma 5.8, $|E'| = O(kn^2/q)$. We conclude that

$$|E(H)| \le |\cup_{S \in \mathcal{S}} H[S]| + |E'| = O((q+k)n + kn^2/q) = O(k^{1/2}n^{3/2}),$$

where the final inequality follows by choosing $q = \sqrt{nk}$ and observing that $k \le n$. ◻

6 A Small Cut Containing All Important Cuts

In this section, we develop a new theorem regarding important cuts that is useful for our fault-tolerant symmetric connectivity preservers. Roughly, we show that for any two disjoint sets of vertices X, Y, there exists a (small) (X, Y)-cut containing every (small) important (X, Y)-cut. We state our cut theorem below, in full generality.

Theorem 6.1. *Let G be an m-edge directed multigraph, and let $X, Y \subseteq V(G)$ be two disjoint sets of vertices with $\text{flow}(G, X, Y) = \lambda$. For every integer $k \geq 0$, there exists an (X, Y)-cut $(S, V - S)$ of size $|\delta^+(S)| \leq \lambda 2^k$ such that every important (X, Y)-cut $(S', V - S')$ of size $|\delta^+(S')| \leq \lambda + k$ satisfies $S' \subseteq S$. Moreover, cut $(S, V - S)$ can be computed in time $O(\lambda k 2^k m + \lambda^2 k^2 4^k)$.*

We use Theorem 6.1 frequently in Section 4 in the following (less general) form.

Lemma 4.1 (Important cut container)**.** *Let k be a positive integer, let G be an n-vertex directed graph, and let $X, Y \subseteq V(G)$ be disjoint sets. There exists an (X, Y)-cut $(S, V - S)$ of size at most $|\delta^+(S)| \leq 2^{k-1}$ such that every important (X, Y)-cut $(S', V - S')$ of size at most $|\delta^+(S')| \leq k$ satisfies $S' \subseteq S$. Cut $(S, V - S)$ can be computed in $O(k 2^k m + k^2 4^k)$ time.*

We quickly verify that Lemma 4.1 follows directly from Theorem 6.1.

Proof of Lemma 4.1. Let $\text{flow}(G, X, Y) = \lambda$. If $\lambda > k$, then there are no important (X, Y)-cut $(S, V - S)$ in G of size at most $|\delta^+(S)| \leq k$, and Lemma 4.1 is vacuously true.

Otherwise, $\lambda \leq k$. Let $k^* = k - \lambda$. We apply Theorem 6.1 with respect to vertex sets X, Y in graph G and integer parameter k^*. This implies that there exists an (X, Y)-cut $(S, V - S)$ of size $|\delta^+(S)| \leq \lambda 2^{k^*}$ such that every important (X, Y)-cut $(S', V - S')$ of size $|\delta^+(S')| \leq \lambda + k^*$ satisfies $S' \subseteq S$. Moreover, $(S, V - S)$ can be computed in the time $O(\lambda k^* 2^{k^*} m + \lambda^2 k^{*2} 4^{k^*})$.

Now we can use the fact that $k = \lambda + k^*$ to prove Lemma 4.1. Since $\lambda \leq 2^{\lambda - 1}$, we observe that $|\delta^+(S)| \leq \lambda 2^{k^*} \leq 2^{k-1}$. Likewise, $|\delta^+(S')| \leq \lambda + k^* = k$. Then the (X, Y)-cut $(S, V - S)$ has size at most $|\delta^+(S)| \leq 2^{k-1}$ and has the property that every important (X, Y)-cut $(S', V - S')$ of size at most $|\delta^+(S')| \leq k$ satisfies $S' \subseteq S$. Finally, using the fact that $\lambda \leq 2^\lambda$, cut $(S, V - S)$ can be computed in $O(k 2^k m + k^2 4^k)$ time. ☐

For the remainder of this section, let G be an n-vertex, m-edge directed multigraph, and let $X, Y \subseteq V(G)$ be two disjoint sets of vertices with $\text{flow}(G, X, Y) = \lambda$. We may assume without loss of generality that X and Y are singleton sets, so that $X = \{x\}$ and $Y = \{y\}$ (e.g., by creating an artificial source x with edges to X and an artificial sink y with edges from Y).

We first describe how to efficiently compute the cut claimed in Theorem 6.1, and then we prove that it has the desired properties. Our proof of Theorem 6.1 will follow [BCR18, BCDW24] and our results can be seen as a generalization of them.

Constructing the cut. We iteratively define a sequence of graphs $G_0, \dots, G_k$ over vertex set $V(G)$. Let graph $G_0 := G$. Given a graph G_i, where $i \in [0, k]$, let $(S_i, V - S_i) = \text{FMC}(G_i, x, y)$ be the farthest minimum (x, y)-cut in G_i. Given set S_i, where $i \in [0, k - 1]$, we define graph G_{i+1} as follows:

- Initially, let $G_{i+1} := G_i$.

- For each edge $(u, v) \in \delta^+_{G_i}(S_i)$, add an edge (x, v) to G_{i+1}. (If edge (x, v) already exists in G_i, then we create an additional parallel (x, v) edge.) This completes the construction of G_{i+1}.

Observe that $(S_i, V - S_i)$ is an (x, y)-cut in G for all $i \in [0, k]$. We return $(S_k, V - S_k)$ as the (x, y)-cut claimed in Theorem 6.1.

Time complexity. We now prove that cut $(S_k, V - S_k)$ can be computed in the time claimed in Theorem 6.1.

Lemma 6.2. *Cut $(S_k, V - S_k)$ can be computed in time $O(\lambda k 2^k m + \lambda^2 k^2 4^k)$.*

We will require the following claim bounding the size of cut $(S_k, V - S_k)$ in graph G_k.

Claim 6.3. *In graph G_k, the (x, y)-cut $(S_k, V - S_k)$ has size at most $|\delta^+_{G_k}(S_k)| \leq \lambda 2^k$.*

Proof. We will prove by induction that for every $i \in [0, k]$, $\text{flow}(G_i, x, y) \leq \lambda 2^i$. When $i = 0$, the claim is satisfied since we assumed $\text{flow}(G, x, y) = \lambda$. Next, we show that for all $i \in [0, k-1]$,

$$\text{flow}(G_{i+1}, x, y) \leq 2\text{flow}(G_i, x, y).$$

We will need several observations to prove this:

- For each $i \in [0, k]$, $\text{flow}(G_i, x, y) = |\delta^+_{G_i}(S_i)|$, since $(S_i, V - S_i)$ is the farthest minimum (x, y)-cut in G_i.
- For each $i \in [0, k-1]$, $\text{flow}(G_{i+1}, x, y) \leq |\delta^+_{G_{i+1}}(S_i)|$. This holds because $(S_i, V - S_i)$ is an (x, y)-cut in G_{i+1}.
- For each $i \in [0, k-1]$, $|\delta^+_{G_{i+1}}(S_i)| \leq 2|\delta^+_{G_i}(S_i)|$. We construct graph G_{i+1} from graph G_i by adding at most $|\delta^+_{G_i}(S_i)|$ edges to G_i. Then there can be at most $|\delta^+_{G_i}(S_i)|$ additional edges crossing cut $(S_i, V - S_i)$ in G_{i+1}. We conclude that $|\delta^+_{G_{i+1}}(S_i)| \leq 2|\delta^+_{G_i}(S_i)|$, as claimed.

Putting these observations together, we find that

$$\text{flow}(G_{i+1}, x, y) \leq |\delta^+_{G_{i+1}}(S_i)| \leq 2|\delta^+_{G_i}(S_i)| = 2 \cdot \text{flow}(G_i, x, y).$$

Then by induction, we conclude that $\text{flow}(G_k, x, y) \leq \lambda 2^k$. In particular, since $(S_k, V - S_k)$ is a minimum (x, y)-cut in G_k, $|\delta^+_{G_k}(S_k)| \leq \lambda 2^k$. □

We are now ready to prove Lemma 6.2.

Proof of Lemma 6.2. Cut $(S_k, V - S_k)$ can be obtained by computing an (x, y) maximum flow on the $k+1$ graphs $G_0, \ldots, G_k$. By Claim 6.3, for each $i \in [0, k-1]$,

$$|E(G_{i+1})| \leq |E(G_i)| + \lambda 2^i \leq m + k\lambda 2^k.$$

Likewise, by Claim 6.3, $\text{flow}(G_i, x, y) \leq \lambda 2^k$ for all $i \in [0, k]$. Then we can compute an (x, y) maximum flow on graph G_i in time $O(\lambda 2^k \cdot (m + k\lambda 2^k))$ using Ford-Fulkerson [FF62]. Then the overall time complexity becomes $O(\lambda k 2^k m + \lambda^2 k^2 4^k)$, as claimed. □

Completing the proof of Theorem 6.1. We have identified an (x,y)-cut $(S_k, V-S_k)$ in G that can be computed in time $O(\lambda k 2^k m + \lambda^2 k^2 4^k)$ time by Lemma 6.2. Moreover, $|\delta^+_G(S_k)| \le \lambda 2^k$ by Claim 6.3. All that remains to complete the proof of Theorem 6.1 is to verify that cut $(S_k, V-S_k)$ contains every important (x,y)-cut $(S, V-S)$ of size at most $|\delta^+_G(S)| \le \lambda + k$. We will need a technical lemma about farthest minimum cuts that was proved in [BCR18].

Lemma 6.4 (cf. Lemma 10 of [BCR18])**.** *Let G be a directed multigraph. Let $s,t \in V(G)$ be vertices in G, and let $(S, V-S) = \mathrm{FMC}(G,s,t)$. For each vertex $v \in V-S$, let G_v be the graph obtained by adding edge (s,v) to G, i.e., $G_v = G + (s,v)$. Then*

$$\mathrm{flow}(G_v, s, t) = \mathrm{flow}(G,s,t) + 1,$$

and $(S, V-S)$ is a minimum (s,t)-cut in G_v.

Lemma 6.4 has the following general implication about farthest minimum cuts in directed multigraphs.

Claim 6.5. *Let $G = (V,E)$ be a directed multigraph. Let $s,t \in V$ be vertices in G, and let $(S, V-S) = \mathrm{FMC}(G,s,t)$. Let $G' = (V, E \cup E')$, where $E' \subseteq \{s\} \times V$, and let $(S', V-S') = \mathrm{FMC}(G', s, t)$. Then $S \subseteq S'$.*

Proof. We can prove Claim 6.5 by repeatedly applying Lemma 6.4. Let $E' = \{e_1, \ldots, e_p\}$, Let $G_0 = G$, and let $G_i = G + e_1 + \cdots + e_i$ for $i \in [1,p]$. Let $(T_i, V-T_i) = \mathrm{FMC}(G_i, s, t)$ for $i \in [0,p]$. By Lemma 6.4, $T_i \subseteq T_{i+1}$ for $i \in [0, p-1]$. Then in particular, $S = T_0 \subseteq T_p = S'$ as claimed. ☐

In particular, Claim 6.5 implies that the sets $S_0, \ldots, S_k$ are nested.

Claim 6.6. *For each $i \in [0, k-1]$, $S_i \subseteq S_{i+1}$.*

Proof. Notice that for each $i \in [0, k-1]$, $G_{i+1} = (V, E(G_i) \cup E_{i+1})$, where edges E_{i+1} satisfy $E_{i+1} \subseteq \{x\} \times V$. Then by Claim 6.5, $S_i \subseteq S_{i+1}$ as claimed. ☐

We now finish our proof of Theorem 6.1 with the following lemma.

Lemma 6.7. *If $(S, V-S)$ is an important (x,y)-cut with size $|\delta^+_G(S)| \le \lambda + k$, then $S \subseteq S_k$.*

Proof. Suppose towards contradiction that there exists an important (x,y)-cut $(S, V-S)$ with size $|\delta^+_G(S)| \le \lambda + k$ such that $S \not\subseteq S_k$, and let $v \in S \cap (V - S_k)$. Since $(S, V-S)$ is an important (x,y)-cut, there exists an (x,v)-path P in $G - \delta^+_G(S)$. We will use the existence of P to obtain a contradiction. Since $v \notin S_k$ and $x \in S_i \subseteq S_k$ for each $i \in [0,k]$, path P contains an edge in $\delta^+_G(S_i)$ for each $i \in [0,k]$. We let $(u_i, v_i) \in E(P)$ denote an arbitrary edge in $E(P) \cap \delta^+_G(S_i)$ for each $i \in [0,k]$.

For analysis purposes, we define a sequence of multigraphs $H_0, H_1 \ldots, H_{k+1}$ over vertex set V. Let graph $H_0 := G$. For each $i \in [0,k]$, we let $H_{i+1} = H_0 + (x, v_0) + \cdots + (x, v_i)$. Notice that in particular, $H_{i+1} = H_i + (x, v_i)$.

Let $(U_i, V-U_i) = \mathrm{FMC}(H_i, x, y)$ for each $i \in [0,k]$. We observe that for each $i \in [0,k]$, $H_i \subseteq G_i$ and $E(G_i) - E(H_i) \subseteq \{x\} \times V$. Then by Claim 6.5, $U_i \subseteq S_i$ for each $i \in [0,k]$. Then since $H_{i+1} = H_i + (x, v_i)$, where $v_i \in V - S_i \subseteq V - U_i$, it follows that

$$\mathrm{flow}(H_{i+1}, x, y) \ge \mathrm{flow}(H_i, x, y) + 1$$

for each $i \in [0, k]$ by Lemma 6.4. In particular, since $\text{flow}(H_0, x, y) = \lambda$, we conclude that $\text{flow}(H_{k+1}, x, y) \geq \lambda + k + 1$.

We are ready to obtain our contradiction. Since $\text{flow}(H_{k+1}, x, y) \geq \lambda+k+1$ and $|\delta_G^+(S)| \leq \lambda+k$, there exists an (x, y)-path P^* in graph $H_{k+1} - \delta_G^+(S)$. If path P^* is contained in graph $G - \delta_G^+(S)$, then this immediately implies a contradiction because $(S, V - S)$ is an (x, y)-cut. Otherwise, path P^* is of the form $P^* = (x, v_i) \circ P^*[v_i, y]$ for some $i \in [0, k]$, where edge (x, v_i) is in graph H_{k+1} and path $P^*[v_i, y]$ is in graph G. However, in this case, path $P[x, v_i]$ is contained in graph $G - \delta_G^+(S)$, so we conclude that the (x, y)-path $P[x, v_i] \circ P^*[v_i, y]$ is contained in graph $G - \delta_G^+(S)$, contradicting our assumption that $(S, V - S)$ is an (x, y)-cut. □

6.1 Anti-isolation lemma

In this section, we will use Lemma 4.1 to show our improved anti-isolation lemma (Lemma 2.3).

Lemma 2.3 (Improved anti-isolation)**.** *Let $s, t_1, \ldots, t_r$ be vertices in a directed graph G, and let $F_1, \ldots, F_r \subseteq E(G)$ be edge sets each of size at most k such that for all $i, j \in [1, r]$ there is an (s, t_j)-path in $G - F_i$ if and only if $i = j$. Then $r \leq 2^k$.*

Proof. For each edge set F_i, where $i \in [1, r]$, there exists a cut $(S_i, V - S_i)$ such that

1. $\delta_G^+(S_i) \subseteq F_i$,
2. $t_i \in S_i$, and
3. $(S_i, V - S_i)$ is an out-reachable (s, t_j)-cut for each $j \neq i \in [1, r]$.

Now to obtain an upper bound on r using Lemma 4.1, we define a new graph H by adding an artificial sink vertex t^* to G and adding the edge (t_i, t^*) to H for all $i \in [1, r]$. Then graph H is defined as $H = G + t^* + (t_1, t^*) + \cdots + (t_r, t^*)$.

We observe that $(S_i, V \cup \{t^*\} - S_i)$ is an out-reachable (s, t^*)-cut in H of size $|\delta_H^+(S_i)| \leq k + 1$ for each $i \in [1, r]$. By Lemma 4.1 and Claim 3.5, there exists an (s, t^*)-cut $(S^*, V - S^*)$ in H of size $|\delta_H^+(S^*)| \leq 2^k$ such that $S_i \subseteq S^*$ for each $i \in [1, r]$. This implies that $(t_i, t^*) \in \delta_H^+(S^*)$ for each $i \in [1, r]$. We conclude that $r \leq |\delta_H^+(S^*)| \leq 2^k$. □

7 Conclusion and Open Problems

Theorem 1.1 and Corollary 1.2 present the first k-FT connectivity preservers and k-connectivity preservers in directed graphs with $O(n \log n)$ edges for every constant k, while Theorem 1.3 shows a k-connectivity preserver with a non-trivial size for any $k = o(n)$. Given the significance and wide-ranging applications of connectivity preservers in undirected graphs [NI92, BK96, BK15, EGIN97, Tho07, ADK+16, Thu95, DHNS19, Par19, AGM12, GMT15, AS23], we are optimistic that our generalization to directed graphs in Theorem 1.1 will also yield further applications. Below, we outline some intriguing open problems motivated by our results.

Optimal Size Bounds. Our approach for proving Theorem 1.1 incurs a factor of $O(\log n)$ due to the depth of the expander hierarchy. Eliminating this $O(\log n)$ factor from the k-FT connectivity preservers might require a completely different approach, making it an exciting challenge.

Additionally, it is important to determine the optimal constant in the factor $2^{\Theta(k)}$ across all variants of k-FT connectivity preservers presented in Table 1. Lastly, we believe that our upper bound of $O(n\sqrt{nk})$ for k-connectivity preservers can be improved to match the lower bound:

Conjecture 7.1. *For every positive integer k, every n-vertex directed graph admits a k-connectivity preserver with $O(kn)$ edges.*

Resolving this conjecture would be exciting.

Near-Linear Construction Time. For many applications including fast algorithms for computing minimum cuts, it is crucial that the preservers are constructed very efficiently. We consider it an important open problem whether one can construct k-fault-tolerant connectivity preservers or k-connectivity preservers in directed graphs in near-linear time, even for constant k. However, it is worth noting that this remains unresolved even in the single-source version; the algorithm in [BCR18] requires $O(2^k mn)$ time, or $O(kmn^{1+o(1)})$ when using the almost-linear time max flow algorithm by [CKL+22].

Acknowledgment We thank an anonymous reviewer for many constructive suggestions that helped us improve the writing.

References

[ACMM05] Amit Agarwal, Moses Charikar, Konstantin Makarychev, and Yury Makarychev. $o(\sqrt{\log n})$ approximation algorithms for min uncut, min 2cnf deletion, and directed cut problems. In *Proceedings of the thirty-seventh annual ACM symposium on Theory of computing*, pages 573–581, 2005. 45

[ADK+16] Ittai Abraham, David Durfee, Ioannis Koutis, Sebastian Krinninger, and Richard Peng. On fully dynamic graph sparsifiers. In *2016 IEEE 57th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 335–344. IEEE, 2016. 1, 11, 30

[AGM12] Kook Jin Ahn, Sudipto Guha, and Andrew McGregor. Graph sketches: sparsification, spanners, and subgraphs. In *Proceedings of the 31st ACM SIGMOD-SIGACT-SIGAI symposium on Principles of Database Systems*, pages 5–14, 2012. 1, 30

[AS23] Sepehr Assadi and Vihan Shah. Tight bounds for vertex connectivity in dynamic streams. In *Symposium on Simplicity in Algorithms (SOSA)*, pages 213–227. SIAM, 2023. 1, 30

[BBST24] Aaron Bernstein, Joakim Blikstad, Thatchaphol Saranurak, and Ta-Wei Tu. Maximum flow by augmenting paths in $n^{2+o(1)}$ time. In *2024 IEEE 65th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 2056–2077. IEEE, 2024. 11

[BCC+23] Davide Bilò, Keerti Choudhary, Sarel Cohen, Tobias Friedrich, Simon Krogmann, and Martin Schirneck. Fault-tolerant st-diameter oracles. In *50th International*

Colloquium on Automata, Languages, and Programming, ICALP 2023, page 24. Schloss Dagstuhl-Leibniz-Zentrum fur Informatik GmbH, Dagstuhl Publishing, 2023. 1

[BCDW24] Shivam Bansal, Keerti Choudhary, Harkirat Dhanoa, and Harsh Wardhan. Fault-tolerant bounded flow preservers, 2024. 3, 4, 27

[BCR18] Surender Baswana, Keerti Choudhary, and Liam Roditty. Fault-tolerant subgraph for single-source reachability: General and optimal. *SIAM Journal on Computing*, 47(1):80–95, 2018. 1, 2, 3, 4, 6, 7, 12, 13, 19, 22, 27, 29, 31, 42, 43

[BCR19] Surender Baswana, Keerti Choudhary, and Liam Roditty. An efficient strongly connected components algorithm in the fault tolerant model. *Algorithmica*, 81:967–985, 2019. 2

[BDKW25] Greg Bodwin, Michael Dinitz, Ama Koranteng, and Lily Wang. Light edge fault tolerant graph spanners. *arXiv preprint arXiv:2502.10890*, 2025. 1

[BDN22] Greg Bodwin, Michael Dinitz, and Yasamin Nazari. Vertex fault-tolerant emulators. In *13th Innovations in Theoretical Computer Science Conference, ITCS 2022*, page 25. Schloss Dagstuhl-Leibniz-Zentrum fur Informatik GmbH, Dagstuhl Publishing, 2022. 1

[BDPW18] Greg Bodwin, Michael Dinitz, Merav Parter, and Virginia Vassilevska Williams. Optimal vertex fault tolerant spanners (for fixed stretch). In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1884–1900. SIAM, 2018. 1

[BDR22] Greg Bodwin, Michael Dinitz, and Caleb Robelle. Partially optimal edge fault-tolerant spanners. In *Proceedings of the 2022 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 3272–3286. SIAM, 2022. 1

[BGPW17] Greg Bodwin, Fabrizio Grandoni, Merav Parter, and Virginia Vassilevska Williams. Preserving distances in very faulty graphs. In *44th International Colloquium on Automata, Languages, and Programming, ICALP 2017*, page 73. Schloss Dagstuhl-Leibniz-Zentrum fur Informatik GmbH, Dagstuhl Publishing, 2017. 1

[BHP24] Greg Bodwin, Bernhard Haeupler, and Merav Parter. Fault-tolerant spanners against bounded-degree edge failures: Linearly more faults, almost for free. In *Proceedings of the 2024 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 2609–2642. SIAM, 2024. 1, 2, 41

[BJ05a] Alex R Berg and Tibor Jordán. Minimally k-edge-connected directed graphs of maximal size. *Graphs and Combinatorics*, 21(1):39–50, 2005. 3, 4

[BJ05b] Alex R Berg and Tibor Jordán. Sparse certificates and removable cycles in l-mixed p-connected graphs. *Operations research letters*, 33(2):111–114, 2005. 1

[BK96] András A Benczúr and David R Karger. Approximating st minimum cuts in õ (n 2) time. In *Proceedings of the twenty-eighth annual ACM symposium on Theory of computing*, pages 47–55, 1996. 1, 30

[BK15] András A Benczúr and David R Karger. Randomized approximation schemes for cuts and flows in capacitated graphs. *SIAM Journal on Computing*, 44(2):290–319, 2015. 1, 30

[BKR03] Marcin Bienkowski, Miroslaw Korzeniowski, and Harald Räcke. A practical algorithm for constructing oblivious routing schemes. In *Proceedings of the fifteenth annual ACM symposium on Parallel algorithms and architectures*, pages 24–33, 2003. 11

[CC20] Diptarka Chakraborty and Keerti Choudhary. New extremal bounds for reachability and strong-connectivity preservers under failures. *ACM Transactions on Algorithms*, 2020. 1, 2, 3, 4

[CH91] Chung-Kuan Cheng and TC Hu. Ancestor tree for arbitrary multi-terminal cut functions. *Annals of Operations Research*, 33(3):199–213, 1991. 9, 25

[Cho21] Keerti Choudhary. k-fault-tolerant connectivity-preserver. https://web.iitd.ac.in/~keerti/Courses/COL866-Notes/kFT-connectivity-preserver.pdf, 2021. Lecture notes. 1

[CKL+20] Marek Cygan, Paweł Komosa, Daniel Lokshtanov, Marcin Pilipczuk, Michał Pilipczuk, Saket Saurabh, and Magnus Wahlström. Randomized contractions meet lean decompositions. *ACM Transactions on Algorithms (TALG)*, 17(1):1–30, 2020. 5

[CKL+22] Li Chen, Rasmus Kyng, Yang P Liu, Richard Peng, Maximilian Probst Gutenberg, and Sushant Sachdeva. Maximum flow and minimum-cost flow in almost-linear time. In *2022 IEEE 63rd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 612–623. IEEE, 2022. 31

[CKT93] Joseph Cheriyan, Ming-Yang Kao, and Ramakrishna Thurimella. Scan-first search and sparse certificates: an improved parallel algorithm for k-vertex connectivity. *SIAM Journal on Computing*, 22(1):157–174, 1993. 1

[CLP+14] Marek Cygan, Daniel Lokshtanov, Marcin Pilipczuk, Michał Pilipczuk, and Saket Saurabh. Minimum bisection is fixed parameter tractable. In *Proceedings of the forty-sixth annual ACM symposium on Theory of computing*, pages 323–332, 2014. 5

[CLPR09] Shiri Chechik, Michael Langberg, David Peleg, and Liam Roditty. Fault-tolerant spanners for general graphs. In *Proceedings of the forty-first annual ACM symposium on Theory of computing*, pages 435–444, 2009. 1

[Dal77] Michel Dalmazzo. Nombre d'arcs dans les graphes k-arc-fortement connexes minimaux., 1977. 3, 4

[DHNS19] Mohit Daga, Monika Henzinger, Danupon Nanongkai, and Thatchaphol Saranurak. Distributed edge connectivity in sublinear time. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, pages 343–354, 2019. 1, 30

[DK11] Michael Dinitz and Robert Krauthgamer. Fault-tolerant spanners: better and simpler. In *Proceedings of the 30th annual ACM SIGACT-SIGOPS symposium on Principles of distributed computing*, pages 169–178, 2011. 1

[EGIN97] David Eppstein, Zvi Galil, Giuseppe F Italiano, and Amnon Nissenzweig. Sparsification—a technique for speeding up dynamic graph algorithms. *Journal of the ACM (JACM)*, 44(5):669–696, 1997. 1, 30

[EIR95] Shimon Even, Gene Itkis, and Sergio Rajsbaum. On mixed connectivity certificates. In *Algorithms—ESA'95: Third Annual European Symposium Corfu, Greece, September 25–27, 1995 Proceedings 3*, pages 1–16. Springer, 1995. 1

[FF62] Lester Randolph Ford and Delbert Ray Fulkerson. Flows in networks. In *Flows in Networks*. Princeton university press, 1962. 10, 28

[FIN93] András Frank, Toshihide Ibaraki, and Hiroshi Nagamochi. On sparse subgraphs preserving connectivity properties. *Journal of graph theory*, 17(3):275–281, 1993. 1

[Fül05] Ottilia Fülöp. Sparse graph certificates for mixed connectivity. *Discrete mathematics*, 294(3):285–290, 2005. 1

[GH61] Ralph E Gomory and Tien Chung Hu. Multi-terminal network flows. *Journal of the Society for Industrial and Applied Mathematics*, 9(4):551–570, 1961. 9, 25

[GILP16] Loukas Georgiadis, Giuseppe F Italiano, Luigi Laura, and Nikos Parotsidis. 2-edge connectivity in directed graphs. *ACM Transactions on Algorithms (TALG)*, 13(1):1–24, 2016. 3

[GIP20] Loukas Georgiadis, Giuseppe F Italiano, and Nikos Parotsidis. Strong connectivity in directed graphs under failures, with applications. *SIAM Journal on Computing*, 49(5):865–926, 2020. 2

[GK17] Manoj Gupta and Shahbaz Khan. Multiple source dual fault tolerant bfs trees. In *44th International Colloquium on Automata, Languages, and Programming (ICALP 2017)*, pages 127–1. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2017. 1

[GMT15] Sudipto Guha, Andrew McGregor, and David Tench. Vertex and hyperedge connectivity in dynamic graph streams. In *Proceedings of the 34th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems*, pages 241–247, 2015. 1, 30

[GRST21] Gramoz Goranci, Harald Räcke, Thatchaphol Saranurak, and Zihan Tan. The expander hierarchy and its applications to dynamic graph algorithms. In *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 2212–2228. SIAM, 2021. 11

[HHR03] Chris Harrelson, Kirsten Hildrum, and Satish Rao. A polynomial-time tree decomposition to minimize congestion. In *Proceedings of the fifteenth annual ACM symposium on Parallel algorithms and architectures*, pages 34–43, 2003. 11

[KL98] David R Karger and Matthew S Levine. Finding maximum flows in undirected graphs seems easier than bipartite matching. In *Proceedings of the thirtieth annual ACM symposium on Theory of computing*, pages 69–78, 1998. 4

[LPS25] Yaowei Long, Seth Pettie, and Thatchaphol Saranurak. Connectivity labeling schemes for edge and vertex faults via expander hierarchies. In *Proceedings of the 2025 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1–47. SIAM, 2025. 5, 11, 45, 46

[LRW25] Jason Li, Satish Rao, and Di Wang. Congestion-approximators from the bottom up. In *Proceedings of the 2025 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 2111–2131. SIAM, 2025. 11

[LS22] Yaowei Long and Thatchaphol Saranurak. Near-optimal deterministic vertex-failure connectivity oracles. In *2022 IEEE 63rd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 1002–1010. IEEE, 2022. 11

[Mad85] Wolfgang Mader. Minimal n-fach zusammenhängende digraphen. *Journal of Combinatorial Theory, Series B*, 38(2):102–117, 1985. 3

[Mar11] Dániel Marx. Important separators and parameterized algorithms. In *International Workshop on Graph-Theoretic Concepts in Computer Science*, pages 5–10. Springer, 2011. 10

[Nag06] Hiroshi Nagamochi. Sparse connectivity certificates via ma orderings in graphs. *Discrete applied mathematics*, 154(16):2411–2417, 2006. 1

[NI92] Hiroshi Nagamochi and Toshihide Ibaraki. Linear time algorithms for finding k-edge connected and k-node connected spanning subgraphs. *Algorithmica*, 7(1992):583–596, 1992. 1, 2, 3, 30

[Par19] Merav Parter. Small cuts and connectivity certificates: A fault tolerant approach. In *33rd International Symposium on Distributed Computing, DISC 2019*, pages 30–1. Schloss Dagstuhl-Leibniz-Zentrum fur Informatik GmbH, Dagstuhl Publishing, 2019. 1, 30

[Par22] Merav Parter. Nearly optimal vertex fault-tolerant spanners in optimal time: sequential, distributed, and parallel. In *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing*, pages 1080–1092, 2022. 1

[PP16] Merav Parter and David Peleg. Sparse fault-tolerant bfs structures. *ACM Transactions on Algorithms (TALG)*, 13(1):1–24, 2016. 1

[PPST25] Merav Parter, Asaf Petruschka, Shay Sapir, and Elad Tzalik. Parks and recreation: Color fault-tolerant spanners made local. In *Proceedings of the 2025 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 4061–4094. SIAM, 2025. 1, 2

[PST24] Asaf Petruschka, Shay Sapir, and Elad Tzalik. Color fault-tolerant spanners. In *15th Innovations in Theoretical Computer Science Conference (ITCS 2024)*, pages 88–1. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2024. 1, 2, 41, 43

[PT07] Mihai Patrascu and Mikkel Thorup. Planning for fast connectivity updates. In *48th Annual IEEE Symposium on Foundations of Computer Science (FOCS'07)*, pages 263–271. IEEE, 2007. 11

[PT25] Merav Parter and Elad Tzalik. Connectivity certificate against bounded-degree faults: Simpler, better and supporting vertex faults. In *2025 Symposium on Simplicity in Algorithms (SOSA)*, pages 369–377. SIAM, 2025. 1, 2

[PW18] Marcin Pilipczuk and Magnus Wahlström. Directed multicut is w [1]-hard, even for four terminal pairs. *ACM Transactions on Computation Theory (TOCT)*, 10(3):1–18, 2018. 7

[Rac02] Harald Racke. Minimizing congestion in general networks. In *The 43rd Annual IEEE Symposium on Foundations of Computer Science, 2002. Proceedings.*, pages 43–52. IEEE, 2002. 11

[RST14] Harald Räcke, Chintan Shah, and Hanjo Täubig. Computing cut-based hierarchical decompositions in almost linear time. In *Proceedings of the twenty-fifth annual ACM-SIAM symposium on Discrete algorithms*, pages 227–238. SIAM, 2014. 11

[Tho07] Mikkel Thorup. Fully-dynamic min-cut. *Combinatorica*, 27(1):91–127, 2007. 1, 30

[Thu95] Ramakrishna Thurimella. Sub-linear distributed algorithms for sparse certificates and biconnected components. In *Proceedings of the fourteenth annual ACM symposium on Principles of distributed computing*, pages 28–37, 1995. 1, 30

[VDBCK+24] Jan Van Den Brand, Li Chen, Rasmus Kyng, Yang P Liu, Simon Meierhans, Maximilian Probst Gutenberg, and Sushant Sachdeva. Almost-linear time algorithms for decremental graphs: Min-cost flow and more via duality. In *2024 IEEE 65th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 2010–2032. IEEE, 2024. 11

[ZSL+19] Chunjiang Zhu, Sabine Storandt, Kam-Yiu Lam, Song Han, and Jinbo Bi. Improved dynamic graph learning through fault-tolerant sparsification. In *International Conference on Machine Learning*, pages 7624–7633. PMLR, 2019. 1

A A hierarchy of connectivity preserver problems

In this section, we aim to survey the landscape of connectivity preserver problems in directed graphs and develop an improved understanding of it. We begin by defining a small family of related connectivity preserver problems. We will need the following definition.

Definition A.1 (k-Bounded Connectivity)**.** For a graph G, we define $\lambda^k_G(s,t) = \min(\lambda_G(s,t), k)$ as the k-bounded connectivity between s and t in G. (Recall that we use $\lambda_G(s,t)$ to denote the symmetric edge-connectivity between s and t in G.) We also use $\lambda^k(G) = \min(\lambda(G), k)$ to denote the k-bounded connectivity of a graph.

With the definition of k-bounded connectivity, we can define two families of directed connectivity preservers as follows.

Definition A.2 (k-Connectivity Family)**.** Let $G = (V, E)$ be a directed graph. For a subgraph H of G, we say H is a

– ***s-t*** k-connectivity preserver if $\lambda^k_H(s,t) = \lambda^k_G(s,t)$ for a given source-sink pair $(s,t) \in V \times V$,

– **global** k-connectivity preserver if $\lambda^k(H) = \lambda^k(G)$,

– **single-source** k-connectivity preserver if $\lambda^k_H(s,t) = \lambda^k_G(s,t)$ for a given $s \in V$ and any $t \in V$,

– **all-pairs** k-connectivity preserver if $\lambda^k_H(s,t) = \lambda^k_G(s,t)$ for any $s, t \in V$.

Note that an all-pairs k-connectivity preserver is exactly the same as the k-connectivity preserver defined in Section 1.2. We can define an analogous family of connectivity preservers for the fault-tolerant setting.

Definition A.3 (k-FT Connectivity Family)**.** Let $G = (V, E)$ be a directed graph. For an subgraph H of G, we say H is a

– ***s-t*** k-FT connectivity preserver if $\lambda^1_{H-F}(s,t) = \lambda^1_{G-F}(s,t)$ for a given source-sink pair $(s,t) \in V \times V$ and any set of edges $F \subseteq E$ with $|F| \leq k$,

– **global** k-FT connectivity preserver if $\lambda^1(H - F) = \lambda^1(G - F)$ for any set of edges $F \subseteq E$ with $|F| \leq k$,

– **single-source** k-FT connectivity preserver if $\lambda^1_{H-F}(s,t) = \lambda^1_{G-F}(s,t)$ for a given source $s \in V$, any $t \in V$, and any set of edges $F \subseteq E$ with $|F| \leq k$,

– **all-pairs** k-FT connectivity preserver if $\lambda^1_{H-F}(s,t) = \lambda^1_{G-F}(s,t)$ for any $s, t \in V$ and any set of edges $F \subseteq E$ with $|F| \leq k$.

We again note that an all-pairs k-FT connectivity preserver is exactly the same as the k-FT connectivity preserver defined in the introduction.

As it turns out, the preservers in our connectivity preserver families can be ordered in a simple hierarchy according to how many edges they require. We formalize this hierarchy with the following theorem, which we prove in Section A.1.

Theorem A.4. *The sizes of the preservers in the k-FT preserver family can be related as follows:*

1. ***s-t*** $\Rightarrow$ ***global****. If there exists a global k-FT preserver with size $f(n,k)$, then there exists an s-t k-FT preserver with size $O(f(n,k))$.*

2. ***global*** $\Rightarrow$ ***single-source****. If there exists a single-source k-FT preserver with size $f(n,k)$, then there exists a global k-FT preserver with size $O(f(n,k))$.*

3. ***single-source*** $\Rightarrow$ ***all-pairs****. If there exists an all-pairs k-FT preserver with size $f(n,k)$, then there exists a single-source k-FT preserver with size $O(f(n,k))$.*

Analogous reductions also hold for the k-connectivity preserver family. Theorem A.4 can be visualized as below.

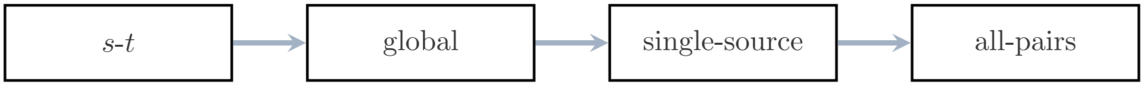

By Theorem 1.1 and our reductions in Theorem A.4, all preservers in the k-FT connectivity preserver family have size at most $O(k4^k n \log n)$. However, do all preservers in this family require size that depends exponentially on k? We confirm that this exponential dependence is necessary with the following theorem, which we prove in Section A.2.

Theorem A.5. *For any positive integers n, k with $n \geq 2^{k/2}$, there exists an n-vertex directed graph G and a pair of vertices $s, t \in V(G)$, such that any s-t k-FT connectivity preserver of G must have $\Omega(n2^{k/2})$ edges.*

A.1 Reductions between Preservers: Proof of Theorem A.4

Proof of Theorem A.4. We prove the three reductions in Theorem A.4 as follows:

1. ***s*-*t*** $\Rightarrow$ **global**. Suppose there exists a global k-FT connectivity preserver with size $f(n,k)$ for every n-vertex graph. We claim that for any n-vertex graph $G = (V, E)$ and every $s, t \in V$, we can build an s-t k-FT connectivity preserver with size $O(f(n,k))$

   We let $G_1 = G + \{(v,s) | v \in V\} + \{(t,v) | v \in V\}$ be the graph obtained by adding (v,s) for any $v \in V$ and (t,v) for any $v \in V$. Let H_1 be a global k-FT preserver of G_1. Suppose t is *reachable* from s in $G - F$ for some fault set $F \subseteq G$ with $|F| \leq k$, then t is *reachable* from s in $G_1 - F$. For all $v \in V$, both (v,s) and (t,v) are in $G_1 - F$. Given that t is *reachable* from s in $G_1 - F$, we know all vertices v will be in the same strong component with s, t.

   Thus, $G_1 - F$ is strongly connected. Since H_1 is a global k-FT preserver of G_1, $H_1 - F$ is also strongly connected, and t is reachable from s in $H_1 - F$. Finally, since any simple directed path from s to t cannot pass any edge (v,s) or (t,v), we deduce that t is *reachable* from s in $(H_1 \cap G) - F$.

   Let $G_2 = G \cup \{(v,t) | v \in V\} \cup \{(s,v) | v \in V\}$, and let H_2 be a global k-FT preserver of G_2. By the same argument we know if s is *reachable* from t in $G - F$, then s is *reachable* from t in $(H_2 \cap G) - F$.

   Finally, we define $H = (H_1 \cup H_2) \cap G$, and observe that $|E(H)| \leq 2f(n,k) + 4n = O(f(n,k))$. Now suppose s and t are *strongly connected* in $G - F$ for some fault set $F \subseteq G$ with $|F| \leq k$, then from above we get t is *reachable* from s in $(H_1 \cap G) - F$ and s is *reachable* from t in $(H_2 \cap G) - F$. Thus, they are still strongly connected in $H - F$. We conclude that H is an s-t k-FT preserver with size $O(f(n,k))$.

2. **global** $\Rightarrow$ **single-source**. Suppose there exists a single-source k-FT preserver with size $f(n,k)$ for every n-vertex graph. Now for any n-vertex graph G, we can pick any source $s \in V(G)$ and build a single-source k-FT preserver H with source s.

   H will have size $f(n,k)$, and we prove that H is a global k-FT preserver of G. Suppose $G - F$ is strongly connected for some fault set F with $|F| \leq k$, then in particular, every vertex v will be strongly connected to the source s in $G - F$. Therefore, since H is a single-source k-FT preserver, we know every vertex v will also be strongly connected to the source s in $H - F$, and thus $H - F$ is also connected.

3. **single-source $\Rightarrow$ all-pairs**. Since any all-pairs k-FT connectivity preserver is also a single-source k-FT connectivity preserver by definition, we immediately get this result.

$\square$

A.2 Lower Bound for s-t k-FT Connectivity Preservers: Proof of Theorem A.5

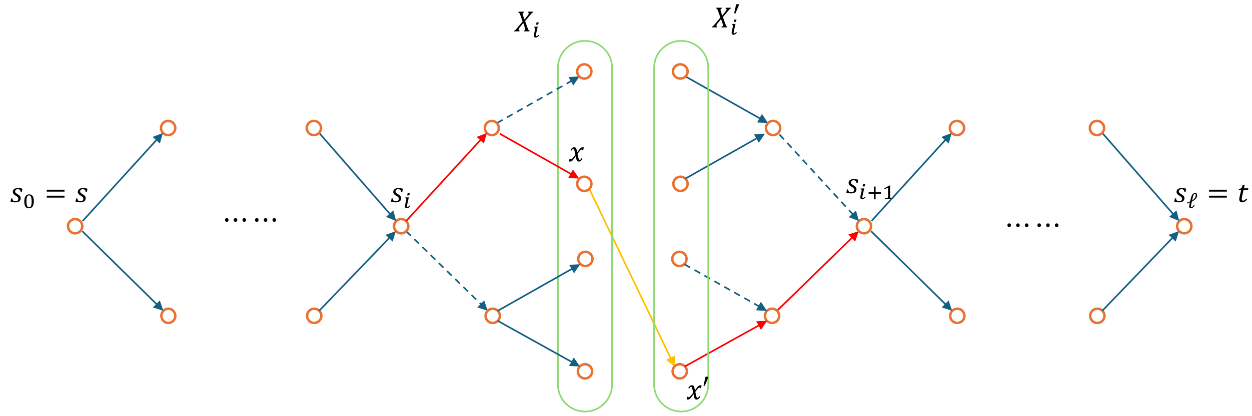

Figure 3: An $\Omega(n2^{k/2})$ Lower Bound for s-t k-FT Preserver

Proof of Theorem A.5. We consider the following construction, as in figure Figure 3. First, we define a set of $\ell = \Theta(n/2^{k/2})$ vertices $s_0, s_1, \ldots, s_\ell$. We will let $s = s_0$ and $t = s_\ell$. Between each pair of vertices s_i and s_{i+1}, $i \in [0, \ell - 1]$, we construct our lower bound graph as follows:

- Construct a full binary tree with depth $k/2$ rooted as s_i, with edges pointing away from root s_i. Denote the tree by T_i and the set of leaves by X_i.
- Construct a full binary tree with depth $k/2$ rooted as s_{i+1}, with edges pointing towards the root s_{i+1}. Denote the the tree by T_i' and set of leaves by X_i'.
- For every $x \in X_i$ and $x' \in X_i'$, add an edge from x to x'.

Additionally, we add an edge from t to s (not shown in Figure 3). Consider the graph G we constructed. We will prove that any s-t k-FT connectivity preserver $H \subseteq G$ must have $\Omega(n2^{k/2})$ edges.

Consider any edge (x, x') such that $x \in X_i$ and $x' \in X_i'$ for some $i \in [0, \ell - 1]$. Let P be an (s_i, x)-path in T_i. Let $F_1 \subseteq E(T_i) - E(P)$ be the set of $k/2$ edges in tree T_i that are incident to path P but not contained in it. We observe that x is the only leaf in X_i that is reachable from s_i in $G - F_1$. Similarly, there is a set $F_2 \subseteq E(T_i')$ of $k/2$ edges in T_i' such that x' is the only leaf in X_i' that can reach s_{i+1} in $G - F_2$.

Finally, we let $F = F_1 \cup F_2$. Now we can observe that any path from s to t in $G - F$ must use edge (x, x'). Since s, t are still strongly connected in $G - F$, it must be true that $(x, x') \in H$. Finally, since there are $\ell = \Omega(n/2^{k/2})$ layers, and each layer has $2^{k/2} \times 2^{k/2} = 2^k$ edges, we know that H must contain at least $\ell 2^k = \Omega(n2^{k/2})$ edges. Therefore, the statement holds. $\square$

A.3 Cut Characterizations

In this section, we state the cut characterizations of (all-pairs) k-connectivity preservers and k-FT connectivity preservers. We prove these characterizations are equivalent to the original flow characterizations of these connectivity preservers in Definition A.2 and Definition A.3. Our cut characterizations highlight the difference and relation between k-connectivity preservers and k-FT connectivity preservers.

Definition A.6 (Minimal symmetric cut)**.** We say that an (s,t)-cut $(S, V-S)$ is a minimal symmetric (s,t)-cut if there does not exist a cut $(S', V-S')$ such that $\delta^+(S') \subsetneq \delta^+(S)$ and $(S', V-S')$ is an (s,t)-cut or a (t,s)-cut.

We now prove our cut characterization of k-FT connectivity preservers.

Theorem A.7. *Given a directed graph G, a subgraph H is an (all-pairs) k-FT connectivity preserver of G if and only if the following flow and cut characterizations hold:*

- *(flow): for every $s,t \in V(G)$ and $F \subseteq E(G)$ where $|F| \leq k$, we have $\lambda^1_{H-F}(s,t) = \lambda^1_{G-F}(s,t)$.*
- *(cut): for every $s,t \in V(G)$ and every minimal symmetric (s,t)-cut $(S, V-S)$ in G, we have $\min(|\delta^+_H(S)|, k+1) = \min(|\delta^+_G(S)|, k+1)$.*

Proof. The flow characterization of k-FT connectivity preservers is identical to the definition of k-FT connectivity preservers in Definition A.3. What remains is to prove that H is a k-FT connectivity preserver of G if and only if the cut characterization holds.

(Only if direction.) Let H be a k-FT connectivity preserver of G. Suppose towards contradiction that H does not satisfy the cut characterization, and let $(S, V-S)$ be a minimal symmetric (s,t)-cut such that $\min(|\delta^+_H(S)|, k+1) < \min(|\delta^+_G(S)|, k+1)$. Let F be the edge fault set $F = \delta^+_H(S)$, and note that $|F| \leq k$. We observe that s,t are strongly connected in $G-F$ (by symmetric minimality and the fact that $\delta^+_H(S) \subsetneq \delta^+_G(S)$), but not in $H-F$, contradicting our assumption that H is a k-FT connectivity preserver. We conclude that H satisfies the cut characterization.

(If direction.) Let H be a subgraph of G that satisfies the cut characterization, and suppose towards contradiction that H is not a k-FT connectivity preserver. Then there exist a pair of vertices $s,t \in V(G)$ and a set of edge faults $F \subseteq E(G)$ with $|F| \leq k$ such that s and t are strongly connected in $G-F$, but not in $H-F$. Without loss of generality, we may assume that s cannot reach t in $H-F$, and we can let $(S, V-S)$ be an (s,t)-cut in $H-F$ of size $|\delta^+_{H-F}(S)| = 0$.

Let $(S', V-S')$ be a cut in G such that $\delta^+_G(S') \subseteq \delta^+_G(S)$, and $(S', V-S')$ is either a minimal symmetric (s,t)-cut or a minimal symmetric (t,s)-cut. We observe that $\delta^+_H(S') \subseteq F$, whereas $\delta^+_G(S') \not\subseteq F$ because s and t are strongly connected in $G-F$. This implies that $\delta^+_H(S') \subsetneq \delta^+_G(S')$. Since $|\delta^+_H(S')| \leq |F| \leq k$, we conclude that

$$\min(|\delta^+_H(S')|, k+1) < \min(|\delta^+_G(S')|, k+1),$$

contradicting our assumption that H satisfies the cut characterization. $\square$

Then, we prove our cut characterization for k-connectivity preservers.

Theorem A.8. *Given a directed graph G, a subgraph H is an (all-pairs) k-connectivity preserver of G if and only if the following flow and cut characterizations hold:*

- *(flow): for every $s, t \in V(G)$, we have $\lambda^k_H(s,t) = \lambda^k_G(s,t)$.*
- *(cut): for every $s, t \in V(G)$ and every minimal symmetric (s,t)-cut $(S, V-S)$ in G, we have $|\delta^+_H(S)| \geq \lambda^k_G(s,t)$.*

Proof. The flow characterization of k-connectivity preservers is identical to the definition of k-connectivity preservers in Definition A.2. What remains is to prove that H is a k-connectivity preserver of G if and only if the cut characterization holds.

(Only if direction.) Let H be a k-connectivity preserver of G. Suppose towards contradiction that H does not satisfy the cut characterization, and let $(S, V-S)$ be a minimal symmetric (s,t)-cut such that $|\delta^+_H(S)| < \lambda^k_G(s,t)$, then we immediately have the contradiction that $\lambda^k_H(s,t) \leq |\delta^+_H(S)| < \lambda^k_G(s,t)$. We conclude that H satisfies the cut characterization.

(If direction.) Let H be a subgraph of G that satisfies the cut characterization, and suppose towards contradiction that H is not a k-connectivity preserver. Then there exist a pair of vertices $s, t \in V(G)$ that $\lambda^k_H(s,t) < \lambda^k_G(s,t)$. Without loss of generality, we may assume there is a cut $(S, V-S)$ be an (s,t)-cut in H of size $|\delta^+_H(S)| < \lambda^k_G(s,t)$.

Let $(S', V-S')$ be a cut in G such that $\delta^+_G(S') \subseteq \delta^+_G(S)$, and $(S', V-S')$ is either a minimal symmetric (s,t)-cut or a minimal symmetric (t,s)-cut. We observe that $|\delta^+_H(S')| \leq |\delta^+_H(S)| < \lambda^k_G(s,t)$, contradicting our assumption that H satisfies the cut characterization.

□

Remark A.9. Theorem A.7 and Theorem A.8 highlight the similarities and differences between k-FT connectivity preservers and k-connectivity preservers. In particular, they both characterize k-connectivity preservers as preserving the minimum symmetric cuts: In k-FT connectivity preservers, we preserve the minimal symmetric cuts of size at most $k+1$, and in k-connectivity preservers, we preserve the minimal symmetric cuts of size at most $\lambda^k_G(s,t) \leq k$. These also imply any $k-1$-FT connectivity preserver is a k-connectivity preserver.

B Strong Connectivity Preservers in Other Fault Models

As mentioned in Section 1.1, we will prove that both the **bounded-degree fault model** (introduced in [BHP24]) and the **color fault model** (introduced in [PST24]) do not admit sparse fault-tolerant connectivity preservers in directed graphs, even in the easiest settings of 1-faulty-degree and 1-color fault, respectively. In this appendix, we will formally define these fault models and prove these lower bounds.

B.1 Bounded-Degree Faults

As in [BHP24], we define the bounded-degree fault model as follows:

Definition B.1 (Bounded-Degree Fault-Tolerant Connectivity Preservers)**.** Given a directed graph G, a subgraph H is called a k-faulty-degree-tolerant connectivity preserver if for any pair of vertices (u,v) and any set of failed edges F such that $|F \cap (\delta^+(x) \cup \delta^-(x))| \leq k$[8] for each vertex $x \in V(G)$, we have the guarantee that u, v are strongly connected in $G - F$ if and only if u, v are strongly connected in $H - F$.

[8] As a generalization to degree in undirected graphs, one may define this as $|F \cap \delta^+(x)| \leq k$ or $|F \cap \delta^-(x)| \leq k$, but these variants are even stronger than what we define here so we cannot hope for anything better.

We prove the following lower bound in the bounded-degree fault setting:

Theorem B.2. *There exists a directed graph G with n vertices such that any 1-faulty-degree-tolerant connectivity preserver $H \subseteq G$ of G requires $\Omega(n^2)$ edges.*

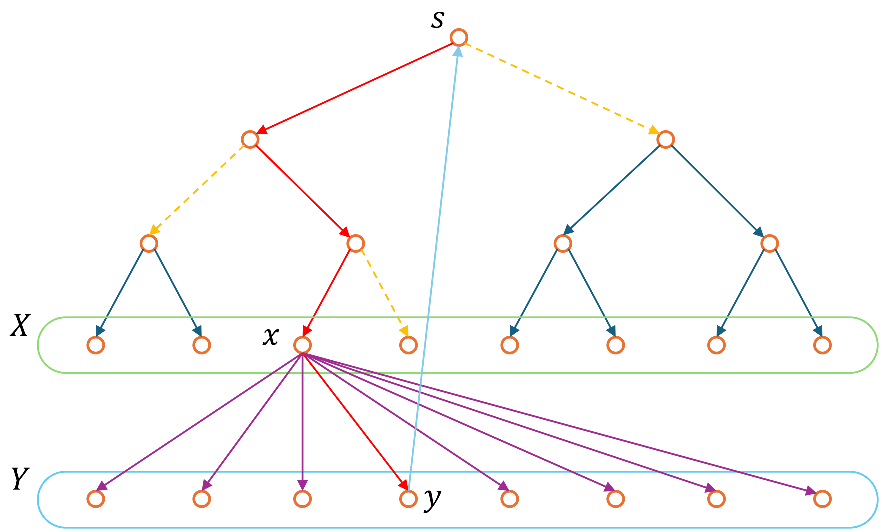

Figure 4: An $\Omega(n^2)$ lower bound for 1-faulty-degree preservers.

Proof. We show that a construction similar to the one in [BCR18] gives the desired lower bound. As in Figure 4, we construct our instance G:

- Construct a full binary tree T with root s and leaf set X of size $|X| = \Theta(n)$. The directed edges in T are oriented away from root s.
- Construct another vertex set Y with $|Y| = \Theta(n)$. Add a directed edge from every $x \in X$ to every $y \in Y$, and add a directed edge from every $y \in Y$ to s.

This final graph G has $\Theta(n)$ vertices and $\Theta(n^2)$ edges.

We prove that any 1-faulty-degree preserver $H \subseteq G$ must contain all edges in $X \times Y$. Consider any $x \in X$ and $y \in Y$. Let P be the unique (s, x)-path in tree T. We let the fault set F be all edges in $T - E(P)$ that are incident to a vertex in (s, x)-path P. This corresponds to all the dashed edges in Figure 4. We observe that fault set F is a 1-bounded-degree fault. Additionally, we observe that x is the only leaf in X that is reachable from s in $G - F$. Therefore, s, y are strongly connected in $G - F$ but not strongly connected in $G - (x, y) - F$. Therefore, to preserve the connectivity between (s, y), the edge (x, y) must be present in the preserver H. Finally, since there are $\Theta(n^2)$ edges in $X \times Y$, we conclude that there must be $\Theta(n^2)$ edges in H. □

As a brief aside, we also quickly point out that the graph constructed in the proof of Theorem B.2 implies a large separation between k-FT connectivity preservers and k-connectivity preservers.

Theorem B.3. *There exists a directed graph G with n vertices that admits a k-connectivity preserver of size $O(n)$, but any k-FT connectivity preserver of G requires $\Omega(n2^k)$ edges.*

Proof. The construction comes from [BCR18]. Consider the graph G in the proof of Theorem B.2 but instead of letting $|X| = \Theta(n)$, we let $|X| = 2^k$. Each pair of vertices in G is only 1-strongly-connected, so we can preserve the k-connectivity of G for any $k \in [1, n]$ using a strongly connected spanning subgraph of size $O(n)$. On the other hand, by cutting the same set of edges as in the proof of Theorem B.2 (which is with size at most k since the tree has k layers), we can show that every k-FT connectivity preserver of G, where $k \in [1, \log_2 n]$, requires $\Omega(2^k n)$ edges. □

B.2 Color Faults

As in [PST24], we will define the color fault model as follows:

Definition B.4 (Color-Fault-Tolerant Connectivity Preservers)**.** Given an edge-colored directed graph G with an associated partition $\mathcal{C}$ of $E(G)$ into $|\mathcal{C}|$ colors, a subgraph H is called a k-color-fault-tolerant connectivity preserver if for any pair of vertices (u, v) and any family of failed colors $\mathcal{F} \subseteq \mathcal{C}$ such that $|\mathcal{F}| \le k$, we have the guarantee that u, v are strongly connected in $G - \cup\mathcal{F}$ if and only if u, v are strongly connected in $H - \cup\mathcal{F}$, where $\cup\mathcal{F}$ denotes the union of all edge sets contained in $\mathcal{F} \subseteq \mathcal{C}$.

We prove the following lower bound in the color fault setting:

Theorem B.5. *There exists a directed graph G with n vertices such that any 1-color-fault-tolerant connectivity preserver $H \subseteq G$ of G requires $\tilde{\Omega}(n^2)$ edges.*

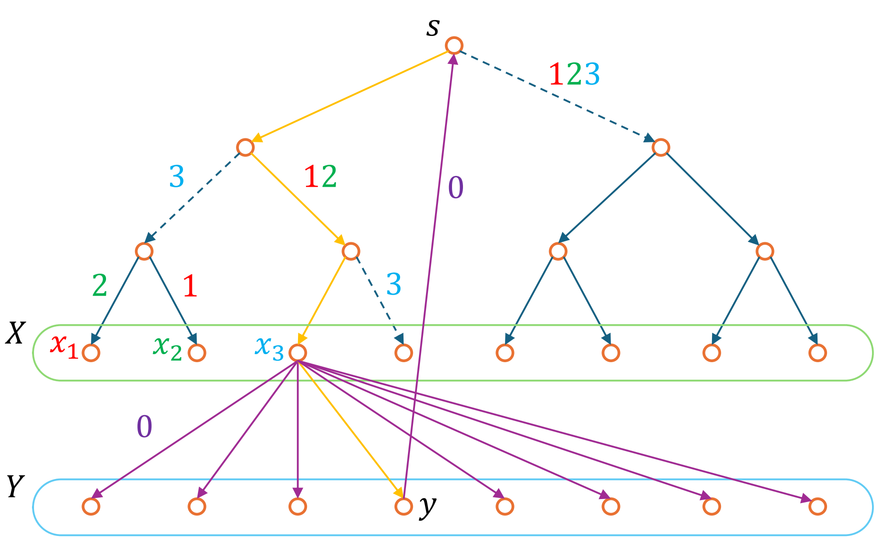

Figure 5: An $\tilde{\Omega}(n^2)$ lower bound for 1-color fault preservers. When the edges in color class 3 fail, every (s, y)-path must pass through vertex x_3.

Proof. As in Figure 5 (where a line with multiple colors denotes a path with an edge of each color), we construct our instance G:

- Construct a full binary tree T with root s and leaf set X of size $|X| = \Theta(n/\log n)$. The directed edges in T are oriented away from source s.

- For each vertex $x_i \in X$, where $i \in [1, |X|]$, we let P_i be the unique (s, x_i)-path in tree T. We color all edges in $T - E(P_i)$ that are incident to a vertex in path P_i using color i. Note that under this coloring, one edge may have been assigned multiple colors.
- For an edge e in tree T that is assigned a set $C_e \subseteq [1, |X|]$ of $|C_e| > 1$ distinct colors by the above coloring, we replace edge e with a path of length $|C_e|$, where each edge is colored with a different color in C_e. After applying this operation to all edges e in tree T, the resulting graph will have a valid edge coloring.

  Since each leaf $x_i \in X$ colors only $|P_i| = O(\log n)$ edges in T with color i, and $|X| = \Theta(n/\log n)$, the previous operation increases the number of vertices in G by at most $|X| \cdot O(\log n) = O(n)$.
- Construct another vertex set Y with $|Y| = \Theta(n)$. Add a directed edge from every $x \in X$ to every $y \in Y$, and add a directed edge from every $y \in Y$ to s. Color all these edges with a new color 0.

This final graph G has $\Theta(n)$ vertices and $\Theta(n^2/\log n)$ edges.

Now we prove that any 1-color-fault-tolerant connectivity preserver $H \subseteq G$ must contain all edges in $X \times Y$. Consider any $x_i \in X$ and $y \in Y$. We set the color fault set F to be the set of all edges in $E(G)$ with color i. We observe that x_i is the only leaf in X that is reachable from s in $G - F$. Therefore, s, y are strongly connected in $G - F$ but not strongly connected in $G - (x_i, y) - F$. Therefore, to preserve the connectivity between (s, y), the edge (x_i, y) must be present in the preserver H. Finally, since there are $\Theta(n^2/\log n)$ edges in $X \times Y$, we conclude that there must be $\Theta(n^2/\log n)$ edges in H. $\square$

C Omitted Proofs from Preliminaries

C.1 A Giant Component from Unbreakable Sets: Proof of Claim 3.6

Proof of Claim 3.6. Let $C_1, \ldots, C_r \subseteq V(G)$ be the strongly connected components in $G - F$. We may assume without loss of generality that for each edge $e \in E(G - F)$, there exist strongly connected components C_i, C_j of $G - F$ such that $e \in C_i \times C_j$, where $i \le j$. This follows from the fact that the condensation graph of $G - F$ is acyclic.

Let $C_{\le i} = C_1 \cup \cdots \cup C_i$ for each $i \in [1, r]$. Suppose towards contradiction that for each $i \in [1, r]$, SCC C_i in $G - F$ satisfies $|C_i \cap U| < |U| - 2q$. Let $i \in [1, r]$ be the largest index such that $|C_{\le i} \cap U| \le q$. Then $|C_{\le i+1} \cap U| > q$. Moreover, since $|C_{i+1} \cap U| < |U| - 2q$, it follows that

$$|(V - C_{\le i+1}) \cap U| = |U| - |C_{\le i} \cap U| - |C_{i+1} \cap U| > |U| - q - (|U| - 2q) = q.$$

Then the cut $(C_{\le i+1}, V - C_{\le i+1})$ will satisfy $|C_{\le i+1} \cap U| > q$ and $|(V - C_{\le i+1}) \cap U| > q$. Additionally, since $\delta^-_{G-F}(C_{\le i+1}) = \emptyset$ by our earlier discussion, it follows that $|\delta^-_G(C_{\le i+1})| \le |F| \le k$. However, this implies that set U is not (q, k)-unbreakable in G, a contradiction.

We conclude that there exists a strongly connected component C in $G - F$ such that $|C \cap U| \ge |U| - 2q$, as claimed. $\square$

C.2 Directed Expander Hierarchy: Proof of Lemma 3.7

In this appendix, we give a proof of the hierarchy in Lemma 3.7 based on approaches in [LPS25]. Although in Lemma 3.7 we state the unbreakability in each layer, in this appendix we will construct the hierarchy using *expansion*.

Definition C.1 (Expansion)**.** For any directed graph $G = (V, E)$, and any terminal set $U \subseteq V$, we say U is *ϕ-expanding* in G, if

$$|\delta^+(S)| \geq \phi \cdot \min(|S \cap U|, |U - S|)$$

for any cut $(S, V - S)$ in G that separates U. We say a graph G is ϕ-expanding if the whole vertex set V is ϕ-expanding.

By replacing S using $V - S$, we know $|\delta^-(S)| \geq \phi \cdot \min(|S \cap U|, |U - S|)$ for any cut $(S, V - S)$ in G that separates a ϕ-expanding set U as well.

Remark C.2. If a terminal set U is ϕ-expanding, then U is $(k/\phi, k)$-unbreakable for any k.

Proof. Since U is ϕ-expanding, we know $|\delta^+(S)|/\phi \geq \min(|S \cap U|, |U - S|)$. Therefore, for any cut $(S, V - S)$ with $|\delta^+(S)| \leq k$ that separates U, we know either $k/\phi \geq |S \cap U|$ or $k/\phi \geq |U - S|$. The same result holds when $|\delta^-(S)| \leq k$. Thus, by Definition 2.1, the above argument holds. □

Generally, to determine expansion is NP-hard and enumerating cuts requires exponential time. However, [ACMM05] shows that we can $O(\sqrt{\log n})$ approximate the expansion of any vertex set U in polynomial time, which we conclude in the following theorem:

Theorem C.3 ([ACMM05])**.** *Given any graph G and a terminal set U, there is a polynomial time algorithm that outputs a cut $(S, V - S)$ that approximates $\phi(U)$ within a factor of $O(\sqrt{\log n})$. i.e. $\phi_U(S) = O(\phi(U) \cdot \sqrt{\log n})$.*

Therefore, to build the hierarchy as in Lemma 3.7, we can consider an *expander hierarchy* with regard to expansion ϕ as follows:

Definition C.4 (Directed Expander Hierarchy)**.** Let G be an n-node directed graph, and let ϕ be any positive expansion. We define a ϕ-expander hierarchy to be a partition $\{V_1, V_2, \ldots, V_\ell\}$ of $V(G)$ into ℓ layers, with the following properties:

1. Fix an $i \in [1, \ell]$. Let $V_{\leq i} = V_1 \cup \cdots \cup V_i$, and let $C \subseteq V(G)$ be a SCC in $G[V_{\leq i}]$. Then the terminal set $V_i \cap C$ is ϕ-expanding in $G[C]$.
2. The number of layers ℓ in the partition is at most $\ell = O(\log n)$.

From Remark C.2, in the above expander hierarchy, the set $V_i \cap C$ is $(k/\phi, k)$-unbreakable in $G[C]$. Thus, it will satisfy Lemma 3.7.Below we show that such a hierarchy in Definition C.4 can be built by the following theorem:

Theorem C.5. *There exists an algorithm that, given a graph G, computes a directed expander hierarchy with $\phi = 1/2$ in exponential time*[9]*, or $\phi = O(1/\sqrt{\log n})$ in polynomial time.*

[9] Same as [LPS25], the factor $1/2$ here is artificial and can be 1 via the same proof.

Proof. WLOG we suppose G is strongly connected, otherwise we consider every strong component of G. As in [LPS25], we construct the whole hierarchy by a top-down approach, from V_ℓ to V_1. To construct the top level $U = V_\ell$, we hope that U satisfies the following two criteria:

1. U is ϕ-expanding in G.

2. Any strong component in $G[V - U]$ has at most $|V|/2$ vertices.

Initially, we let $U = V$ and finally we want to make U ϕ-expanding in G while satisfying the second criterion. Consider the following procedure. While U is not ϕ-expanding in G, there must exist a sparse cut $\delta^+(S)$ with $\phi_U(S) \leq 1/2$. Moreover, such cut $\delta^+(S)$ could be found in exponential time when $\phi(U) \leq 1/2$, or in polynomial time when $\phi(U) \leq O(1/\sqrt{\log n})$. Suppose $L = \{s \in S | \exists t \in V - S \text{ s.t. } (s,t) \in \delta^+(S)\}$, then:

- If $|S| \leq |V|/2$, then we update U as $U' = (U - S) \cup L$.

- Else, $|V - S| \leq |V|/2$, and we update as $U' = (U - (V - S)) \cup L$.

First, we show the procedure will terminate (and thus we have a final set U that is ϕ-expanding in G). In every update, we add at most $|L| \leq |\delta^+(S)|$ vertices to U, but we also delete at least $|U \cap S| \geq |\delta^+(S)|/(1/2)$ or $|U \cap (V - S)| \geq |\delta^+(S)|/(1/2)$ vertices from U. Thus, we know $|U|$ is strictly decreasing in every update and the procedure must terminate.

Then, we show that the second criterion will be invariant while we update U to U'. WLOG, we consider below the case when $|S| \leq |V|/2$, and the other case can be proved symmetrically. In one update, we let $U' = (U - S) \cup L$. Let's consider any strong component C in $G[V - U']$. We consider the following two cases:

1. If $C \cap S \neq \emptyset$, since C contains no vertex in $L \subseteq U'$, we know C must be fully contained in S, and with size $|C| \leq |S| \leq |V|/2$.

2. Otherwise, $C \cap S = \emptyset$. Therefore, C is a strong component in $G[V - U' - S]$. From the update, we know $(V - S) - U' = (V - S) - U$, thus, $G[V - U' - S] = G[V - U - S]$. We conclude that C must be also strongly connected in $G[V - U]$, which implies $|C| \leq |V|/2$.

Therefore, when the procedure terminates, all strong component C in $G[V - U]$ has at most $|V|/2$ vertices. We let $V_\ell = U$, so $V_{\leq \ell - 1} = V - U$ then recurse in every strong component C in $G[V_{\leq \ell-1}]$ to obtain $V_{\ell-1}$ through V_1. From the above construction, we know the size of any component C in $G[V_{\leq i}]$ will be halved compared to $G[V_{\leq (i+1)}]$ for any i. Since G contains one component with n vertices, there can be at most $\log n$ layers. $\square$

Finally, we have constructed a hierarchy that satisfies Lemma 3.7.