

Fault-Tolerant Spanners against Bounded-Degree Edge Failures: Linearly More Faults, Almost For Free

Greg Bodwin* Bernhard Haeupler† Merav Parter‡
 University of Michigan ETH Zurich & CMU Weizmann Institute of Science

Abstract

We study a new and stronger notion of fault-tolerant graph structures whose size bounds depend on the degree of the failing edge set, rather than the total number of faults. For a subset of faulty edges $F \subseteq G$, the *faulty-degree* $\deg(F)$ is the largest number of faults in F incident to any given vertex. For example, a matching F has $\deg(F) = 1$ while $|F|$ might be as large as $n/2$.

We design new fault-tolerant structures with size comparable to previous constructions, but which tolerate every fault set of small faulty-degree $\deg(F)$, rather than only fault sets of small size $|F|$. Thus, for example, our structures can tolerate a *linear* number of edge faults with almost the same size bounds currently known for handling a *single* edge failure, provided that the edge faults are arranged in a matching. Our main results are:

- **New FT-Certificates:** For every n -vertex graph G and degree threshold f , one can compute a connectivity certificate $H \subseteq G$ with $|E(H)| = \tilde{O}(fn)$ edges that has the following guarantee: for any edge set F with faulty-degree $\deg(F) \leq f$ and every vertex pair u, v , it holds that u and v are connected in $H \setminus F$ iff they are connected in $G \setminus F$. This bound on $|E(H)|$ is nearly tight. Since our certificates handle some fault sets of size up to $|F| = O(fn)$, prior work did not imply any nontrivial upper bound for this problem, even when $f = 1$.
- **New FT-Spanners:** We show that every n -vertex graph G admits a $(2k-1)$ -spanner H with $|E(H)| = O_k(f^{1-1/k}n^{1+1/k})$ edges, which tolerates any fault set F of faulty-degree at most f . This bound on $|E(H)|$ is optimal up to its hidden dependence on k , and it is close to the bound of $O_k(|F|^{1/2}n^{1+1/k} + |F|n)$ that is known for the case where the *total* number of faults is $|F|$ [Bodwin, Dinitz, Robelle SODA '22]. Our proof of this theorem is non-constructive, but by following a proof strategy of Dinitz and Robelle [PODC '20], we show that the runtime can be made polynomial by paying an additional $\text{polylog } n$ factor in spanner size.

Our techniques are based on an adaptation of the blocking set method used in previous work on fault tolerant spanners, as well as a new expander-based toolkit which translates the quality guarantees for expander routing into fault tolerance guarantees.

*bodwin@umich.edu. Supported by NSF:AF 2153680.

†bernhard.haeupler@inf.ethz.ch. Supported by the European Research Council (ERC grant 949272).

‡merav.parter@weizmann.ac.il. Supported by the European Research Council (ERC grant 949083) and the Israeli Science Foundation (ISF), grant 2084/18.

Contents

1	Introduction	1
1.1	Connectivity Certificates	1
1.2	Fault-Tolerant Spanners	2
1.3	Expanders and Rigid vs. Competitive Fault Tolerance	4
2	Technical Overview	6
2.1	Expander-Based Constructions	6
2.1.1	Sparse Certificates	6
2.1.2	Sparse Spanners	7
2.2	Blocking Set-Based Constructions	9
2.2.1	The FD Greedy Algorithm and Analysis	9
2.2.2	Speeding Up the Greedy Algorithm	10
3	Preliminaries	11
4	Sparse Connectivity Certificates against Bounded-Degree Faults	14
5	Sparse Spanners against Bounded-Degree Faults	16
5.1	Robustness of Length-Constrained Expanders against Bounded Degree Faults	16
5.2	Length-Constrained Expander Sparsification	17
5.3	Expander-Based Computation of FD-Spanners	18
6	Improved Existential Upper Bounds for FD Spanners	20
6.1	Proof Setup	20
6.2	MUCk Paths and Dispersion Lemma	22
6.3	Counting Lemma	23
6.4	Proof Wrapup	25
6.5	The MIN MAX LENGTH-BOUNDED CUT Problem	26
6.6	Improvement to Polynomial Time	29
7	Existential Lower Bounds for FD Spanners and Connectivity Certificates	30
7.1	FD Spanner Lower Bounds	30
7.2	FD Connectivity Certificate Lower Bounds	30
A	Missing Proofs	34
B	$G(n, p)$ is a good constant-hop router	36

1 Introduction

A basic problem in theoretical computer science and graph theory is to compress distance or connectivity information from an input graph G into much smaller space than G itself. In its most basic form, this corresponds to computation of a sparse spanning subgraph (which captures connectivity information), or to computation of a sparse spanner (which captures approximate distances). However, in some domains such as distributed computing and network design, it is common for a few nodes or edges of the input graph G to temporarily *fail* and be unusable while they await repair. This motivates a more constrained problem, in which the goal is to find a sparse subgraph H that still captures connectivity/distance information of G even after any “reasonable” failure event that might occur. This paper concerns the design of subgraphs of this kind.

1.1 Connectivity Certificates

Let us begin our discussion in the setting of connectivity. A common formalization of connectivity preservation under failures is known as *fault-tolerant connectivity certificates*:

Definition 1.1 (EFT Connectivity Certificates [NI92]). *Given a graph G , an edge-subgraph H is called an f -edge fault tolerant (EFT) connectivity certificate if, for any nodes u, v and for any set F of $|F| \leq f$ edges, there is a $u \rightsquigarrow v$ path in $G \setminus F$ iff there is a $u \rightsquigarrow v$ path in $H \setminus F$.*

This definition encodes a design choice: the idea of a “reasonable” failure event is formalized simply by choosing a size parameter f , and letting F be any edge set of size $|F| \leq f$. We will refer to this as **global** fault tolerance, since it counts faults equally no matter where they occur in the graph. Global fault tolerance is the dominant model studied in prior work. The tradeoff between certificate size and level of global fault tolerance was resolved in a classic work of Nagamochi and Ibaraki [NI92]:

Theorem 1.2 ([NI92]). *For any positive integers n, f , every n -node graph G has an f -EFT connectivity certificate H on $|E(H)| = O(fn)$ edges, and this bound is best possible.*

The main new ideas in Nagamochi and Ibaraki’s work lie on the upper bound side. The lower bound, that $|E(H)| = \Omega(fn)$ edges are sometimes necessary, is simple. It is achieved by letting G be any $(f + 1)$ -regular graph. Consider an edge $(u, v) \in E(G)$, and let the failure set F be the f additional edges incident to u besides (u, v) . Then (u, v) is the only edge still incident to u in $G \setminus F$, which implies that we must keep the edge (u, v) in the connectivity certificate H . Since this holds for an arbitrary edge, it follows that the only f -EFT connectivity certificate of G is G itself, which contains $\Theta(fn)$ edges.

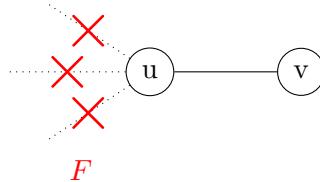


Figure 1: By failing all edges incident to u except for (u, v) , we see that any connectivity certificate must keep the edge (u, v) .

The $(f + 1)$ -regular graph lower bound is simple and optimal, but in a way it might also feel unsatisfying: is it really a “reasonable” failure event to see all f edge failures incident to the same

node u ? Although mass failures adjacent to u are certainly possible – failures can be correlated, and perhaps they all result from a common systemic problem with u – it is less reasonable to think that the rest of the graph would meanwhile remain completely failure-free. Failures in different parts of the graph tend to be independent or positively correlated, and so when u suffers failures, we may still expect to see *at least* the usual number of failures occurring elsewhere in the graph at the same time.

This motivates a much stronger kind of connectivity certificate, which can tolerate edge failures that occur elsewhere in the graph alongside the f failing edges incident to u . We will consider a new **local** model of fault tolerance, which only limits failures-per-node rather than total number of failures:

Definition 1.3 (Faulty-Degree Connectivity Certificates (New)). *Given a graph G , an edge-subgraph H is called an f -faulty-degree (FD) connectivity certificate if, for any nodes u, v and for any set of edges F such that $\deg_F(x) \leq f$ for all nodes x , there is a $u \rightsquigarrow v$ path in $G \setminus F$ iff there is a $u \rightsquigarrow v$ path in $H \setminus F$.*

An f -FD connectivity certificate is stronger than an f -EFT connectivity certificate, in the sense that every f -EFT connectivity certificate is also an f -FD connectivity certificate, but the converse is far from true. Indeed, since an f -FD connectivity certificate must tolerate some failure sets of size up to $|F| = \Theta(fn)$, the previous upper bounds for f -EFT connectivity certificates do not imply *any* nontrivial (subquadratic) upper bounds for f -FD connectivity certificates. Despite the stronger model, the $(f+1)$ -regular graph still only provides the same lower bound of $\Omega(fn)$ for f -FD connectivity certificates. Our main result is that this is not an accident: there is actually a nearly-matching upper bound.

Theorem 1.4 (First Main Result). *For all positive integers f, n , every n -node graph has an f -FD connectivity certificate on $\tilde{O}(fn)$ edges.*

This theorem is constructive, i.e., we also provide an algorithm to compute the certificate in polynomial time. The theorem proved in the paper gives a somewhat stronger guarantee; one can recover not just connectivity, but also shortest path distances up to a polylog n factor. Finally, we complement this theorem with a lower bound: although the hidden polylog n factors in this upper bound might be improvable, they cannot be removed entirely. This stands in contrast to the standard f -FT connectivity certificates of Theorem 1.2. Hence the cost of handling a faulty set with degree f (with possibly $\Theta(fn)$ edges) is at least logarithmic and at most polylogarithmic, compared to handling a total of f faults.

Theorem 1.5. *There are n -node graphs for which any f -FD connectivity certificate has $\Omega(fn \cdot \frac{\log n}{\log f})$ edges.*

Our lower bound construction is still relatively simple: when $f = 1$, the certificate size lower bound is $\Omega(n \log n)$, and this is achieved by analyzing the hypercube graph. For general f the construction is essentially the extension of the hypercube graph to an alphabet of size f .

1.2 Fault-Tolerant Spanners

A *spanner* is a subgraph that preserves approximate pairwise distances, rather than just connectivity. Spanners and their variants have many applications in algorithms and network design; see survey [ABS⁺20]. In this discussion, graphs can have arbitrary positive edge weights.

Definition 1.6 (Spanners [PS89]). *For an input graph G , an edge-subgraph H is called a t -spanner of G if it satisfies $\text{dist}_H(u, v) \leq t \cdot \text{dist}_G(u, v)$ for all nodes u, v .*

The parameter t of the spanner is often called its *stretch*. The existentially optimal size-stretch tradeoff was settled in a classic paper by Althöfer, Das, Dobkin, Joseph, and Soares:

Theorem 1.7 ([ADD⁺93]). *For all positive integers n, k , every n -node graph G has a $(2k-1)$ -spanner on $O(n^{1+1/k})$ edges.*

This upper bound is unconditional, and there is a matching lower bound assuming the girth conjecture [Erd64], i.e., this tradeoff is best possible up to the hidden constant. Much like connectivity certificates, spanners are often applied in settings with edge faults, in which case a notion of fault tolerance is needed. This was first considered in the special case of Euclidean input graphs by Levcopoulos, Narasimhan, and Smid [LNS98] and has since been studied intensively in this setting, and also in the more general setting of doubling metrics [CZ04, Luk99, NS07, Sol14, CLN15, CLNS15, LST23]. For general input graphs, the first results on fault-tolerant spanners were obtained by Chechik, Langberg, Peleg, and Roditty [CLPR10], in the following global model of fault-tolerance identical to the one used for connectivity certificates.

Definition 1.8 (Fault-Tolerant Spanners [CLPR10]). *Given a graph G , an edge-subgraph H is an f -edge fault-tolerant (f -EFT) t -spanner if for any set F of $|F| \leq f$ edges and any nodes u, v , we have $\text{dist}_{H \setminus F}(u, v) \leq t \cdot \text{dist}_{G \setminus F}(u, v)$.*

We recap the considerable prior work on EFT spanners in Table 1. For now, we note that the lower bound graphs from [BDPW18], which we describe in Section 7.1, partially share the same locality property as the lower bounds for connectivity certificates. In particular, we can consider the following model of local fault tolerance for spanners:

Definition 1.9 (Faulty-Degree Spanners (New)). *Given a graph G , an edge-subgraph H is called an f -faulty-degree (FD) t -spanner if, for any nodes u, v and for any set of edges F such that $\deg_F(x) \leq f$ for all nodes x , we have $\text{dist}_{H \setminus F}(u, v) \leq t \cdot \text{dist}_{G \setminus F}(u, v)$.*

Once again, every f -EFT spanner is necessarily also an f -FD spanner, but the converse is far from true, as f -FD spanners must handle some fault sets of size up to $\Theta(fn)$. The graph constructions from [BDPW18] imply a lower bound of $\Omega(f^{1-1/k}n^{1+1/k})$ for f -FD spanners. Our next main result is to show that this bound is essentially tight.

Theorem 1.10 (Second Main Result). *For all positive integers f, n, k , every n -node graph has an f -FD $(2k-1)$ spanner on $f^{1-1/k} \cdot n^{1+1/k} \cdot O(k)^k$ edges. Assuming the girth conjecture [Erd64], this tradeoff is best possible when k is a constant. Additionally, one can compute the spanner in polynomial time by paying an additional $\text{polylog } n$ factor in the spanner size.*

We obtain the result by analyzing a certain exponential-time greedy algorithm, analogous to the one used repeatedly in previous work on EFT spanners [BDPW18, BDR22]. In order to speed it up to polynomial time, we follow an approach of Dinitz and Robelle [DR20] to replace a key step in the greedy construction with an approximation algorithm. Thus costs roughly the approximation factor in spanner size. We overview the technical aspects of this approximation algorithm more in Section 2.2.2.

Although our results focus on spanners and connectivity certificates, we wish to emphasize here that the local (FD) model of fault tolerance is sensible for many other problems, e.g., distance oracles, flows, etc. It will also be interesting to provide analog local FD results for vertex faults. The conceptual message behind Theorems 1.4 and 1.10 is that, despite being considerably stronger, one can still prove meaningful upper bound results in the local model. A natural open question is therefore to investigate whether upper bounds can be achieved when considering these other problems in the local fault tolerant model, which have previously been considered only in the global fault tolerant setting.

Stretch	Size	Polytime?	Citation
Global Fault Tolerance (f-EFT spanners)			
$2k - 1$	$O(f \cdot n^{1+1/k})$	✓	[CLPR10]
$2k - 1$	$O_k(f^{1-1/k} \cdot n^{1+1/k})$		[BDPW18]
$2k - 1$	$O(f^{1-1/k} \cdot n^{1+1/k})$		[BP19]
$2k - 1$	$O_k(f^{1-1/k} \cdot n^{1+1/k})$	✓	[DR20]
$2k - 1$	$O(f^{1-1/k} \cdot n^{1+1/k})$	✓	[BDR21]
$2k - 1, k$ even	$O_k(f^{1/2} \cdot n^{1+1/k} + fn)$	✓	[BDR22]
$2k - 1, k$ odd	$O_k(f^{1/2-1/(2k)} \cdot n^{1+1/k} + fn)$	✓	[BDR22]
3	$\Omega(f^{1/2} \cdot n^{3/2})$	lower bound	[BDPW18]
$2k - 1, k > 2$	$\Omega(f^{1/2-1/(2k)} \cdot n^{1+1/k} + fn)$	lower bound*	[BDPW18]
Local Fault Tolerance (f-FD spanners)			
$\text{polylog } n$	$\tilde{O}(f \cdot n)$	✓	this paper
$2k - 1$	$\tilde{O}_k(f^{1-1/k} \cdot n^{1+1/k})$	✓	this paper
$2k - 1$	$O_k(f^{1-1/k} \cdot n^{1+1/k})$		this paper
$2k - 1$	$\Omega(f^{1-1/k} \cdot n^{1+1/k})$	lower bound*	[BDPW18]
∞ (connectivity)	$\Omega\left(\frac{f}{\log f} \cdot n \log n\right)$	lower bound	this paper

Table 1: Prior work on upper and lower bounds for edge fault-tolerant spanners, both in the global model (f -EFT spanners) and the local model (f -FD spanners). The lower bound entries marked with an asterisk (*) are conditional on the girth conjecture [Erd64]. We also show constructions of FD-spanners via expander decomposition methods, not listed in this table.

1.3 Expanders and Rigid vs. Competitive Fault Tolerance

An important theme in the literature on fault-tolerant graph structures is the distinction between *rigid* and *competitive* fault tolerance [CLPR10]. All of the previous discussion concerns the competitive notion of fault tolerance, because we compare distances or connectivity between the post-failure subgraph $H \setminus F$ and the *post-failure* subgraph $G \setminus F$. The rigid notion of fault tolerance is defined similarly, but we would compare distances or connectivity between $H \setminus F$ and the *pre-failure* subgraph G .

Rigid fault tolerance is a stronger notion than competitive fault tolerance, and hence it is preferable when it can be achieved. However, rigid fault tolerance is not available for all graphs. For example, when G is a tree, a single edge failure can disconnect the graph. Thus there is no 1-rigid-fault-tolerant subgraph of G , not even G itself. Despite this, rigid fault tolerance is available for certain restricted graph classes, e.g., Euclidean [LNS98]. An additional contribution of this work is a refined understanding of when rigid fault tolerance is available.

Rigid Fault Tolerance for Connectivity. A key technical ingredient in our construction of connectivity certificates is a new structural lemma, showing that every *expander* of high enough minimum degree admits a notion of rigid fault tolerance for connectivity (and, in fact, also for expansion).

Theorem 1.11 (Expanders are Rigid to Bounded-Degree Faults). *Let G be an n -vertex ϕ -expander with minimum degree $f' \geq 2f/\phi$. Then for every f -degree faulty-set F , the graph $G \setminus F$ remains connected. (In fact, $G \setminus F$ is itself a $\phi/2$ -expander.)*

For general graphs G , this implies a strategy to design FD connectivity certificates. Suppose that we compute an expander decomposition of G (or rather, a slight variant of expander decomposition that also enforces a high enough minimum degree of each expander). We then sparsify each expander in the decomposition, while maintaining its expansion properties and its minimum degree. Theorem 1.11 promises that every edge with both endpoints in the same expander is *automatically* protected from faults: these edge failures simply cannot disconnect the graph, so they can be ignored. Indeed, the only bad edges are the ones whose endpoints lie in two different expanders in the decomposition. These bad edges can be handled recursively, until very few of them remain, and the remaining edges can simply be added to the certificate. We overview this strategy in more detail in Section 2.1.1.

Rigid Fault Tolerance for Distances. Theorem 1.11 actually provides a notion of rigid fault tolerance for *distances*, rather than just connectivity. Since the post-failure graph $G \setminus F$ remains an expander, its distances can be changed only by a factor of $O(\log n)$ by the fault set F – in other words, $G \setminus F$ is an $O(\log n)$ spanner of G . Plugging this fact into the recursive expander decomposition framework outlined above, this ultimately implies that the FD connectivity certificate is in fact an FD spanner with stretch $\text{polylog } n$.

It is natural to ask whether this $O(\log n)$ factor in rigid fault tolerance for distance can be improved. We show that it can, by instead considering the recently-introduced class of length-constrained (LC) expanders [HRG22]. We will review these objects formally in Section 2.1.2, but briefly, an (h, s) -length expander is a graph where any h -length unit demand can be routed by a multi-commodity flow with congestion $\tilde{O}(1/\phi)$ and over paths of length at most hs . The following theorem shows that they provide an improved notion of rigid fault tolerance for distances:

Theorem 1.12 (Length-Constrained Expanders are Rigid to Bounded-Degree Faults). *Let G be an n -vertex (h, s) -length ϕ -expander with minimum degree $\tilde{\Omega}(f \cdot n^\epsilon/\phi)$. Then for every f -degree faulty-set F , the graph $G \setminus F$ is a $(hs)^{O(1/\epsilon)}$ -spanner of G .*

The analogous construction now works for FD spanners: one can take a LC expander decomposition (again enforcing high minimum degree), and then sparsify each expander while maintaining its expansion properties and its minimum degree. The precise arguments are more delicate in this setting and in particular, our sparsification does not provide an LC-expander but rather a sparse subgraph that has similar routing properties as those provided by LC expanders. But the effect is again that edge faults whose endpoints lie in the same expander can essentially be ignored, since Theorem 1.12 implies that these faults simply cannot change the pairwise distances by too much. The only bad edges are the ones whose endpoints lie in different expanders, which are again handled recursively, until only a few bad edges remain which can be taken into the spanner. Analyzing this construction leads to the following result:

Theorem 1.13 (Second Main Result, Part 2). *For all positive integers f, n, k , every n -node unweighted graph has an f -FD $k^{O(k)}$ -spanner $H \subseteq G$ with $|E(H)| = \tilde{O}(f \cdot n^{1+1/k})$ edges.*

Although the size-stretch tradeoff of the FD-spanner in Theorem 1.13 is worse than the one in Theorem 1.10 (which is achieved using an unrelated non-expander-based toolkit), we include it because we

think that the new connections between expanders, expander decomposition, and rigid fault tolerance illustrated by this theorem are worthwhile to explore, and could have applications in followup work, especially where computational aspects of the construction are emphasized. Additionally, the non-tightness of the size-stretch tradeoff in this theorem is mostly due to a lack of understanding of the distance rigidity properties of expanders, rather than a weakness in the expander-based construction framework itself. The following result states that the size-stretch tradeoff in this construction could improve from exponential to linear – hence tight, up to a constant factor in the stretch and a sublinear factor in the faulty-degree – if the distance rigidity bounds from Theorem 1.12 can be improved.

Theorem 1.14 (Informal). *Suppose that the distance rigidity bound in Theorem 1.12 can be improved from $O(hs)^{O(1/\varepsilon)}$ to ℓ , for some parameter ℓ . Then our construction yields an f -FD $O(\ell \cdot \text{poly}(s))$ -spanner $H \subseteq G$ with $|E(H)| = \tilde{O}(f \cdot n^{1+O(1/s)})$ edges.*

We think it is an interesting open problem to improve the distance rigidity bound ℓ in this theorem, both in its own right and to improve this particular application to FD spanners. We refer to Section 2.1.2 for more technical details on the construction and this possible extension.

Related Work on Spanners and Expanders. In a very recent work, Haeupler, Hershkowitz and Tan [HHT23b] provided a new implementation of the well-known greedy (non-FT) spanner algorithm of [ADD⁺93]. In particular, they show that one can add at a time a collection of matching edges M provided that each individual edge in M does not close a short cycle in the current spanner. On a high-level, their size analysis approach is based on exploiting the routing properties of the recently introduced notion of length-constrained expanders [HRG22]. While their paper does not address fault-tolerance aspects, their use of expanders and more specifically, expander routing, in the context of (standard) spanners highly inspired our expander-based constructions of FD spanners.

2 Technical Overview

2.1 Expander-Based Constructions

Our expander-based constructions use a two-step approach. We first show that, when the input graph is an expander, one can provide rigid fault-tolerant guarantees for connectivity and distances. By employing the expander-decomposition technique, we then translate the rigid fault-tolerant guarantees for expanders into competitive fault-tolerant guarantees for general input graphs. To the best of our knowledge, these tools have not been employed before in the context of fault-tolerant graph sparsification.

On a conceptual level, our expander-based sparsification approach allows us to pinpoint a critical set of edges in the input graph G that prevent one from providing rigid fault-tolerant guarantees for G . Our proof shows that the *only* problematic edges for rigid fault-tolerance are the inter-expander edges, that connect different expander subgraphs in the output of the expander decomposition procedure. Those edges can be handled recursively until their number becomes bounded by $\tilde{O}(f \cdot n)$. This set of small number of edges are the problematic edges for obtaining rigid fault-tolerance, and by adding them to the output subgraph, we provide competitive fault-tolerance for G .

2.1.1 Sparse Certificates

We first discuss our use of expanders for sparse connectivity certificates. Our starting observation is that any ϕ -expander with minimum degree $\Theta(f)$ can tolerate the removal of any faulty-set F with degree f , and maintain its expansion properties, up to a small loss in the conductance parameter. As formulated in Theorem 1.11.

To illustrate our ideas, assume for simplicity that the minimum degree of the input graph G is $\tilde{\Omega}(f)$. Our algorithm first employs an expander-decomposition procedure on G which partitions G into vertex-disjoint ϕ -expanders G_1, \dots, G_k for $\phi = 1/\log n$, and an additional set $H_0 \subseteq G$ of at most $m/2$ inter-expander edges. In the second step, the algorithm sparsifies each expander G_i into a $\Theta(\phi)$ -expander $H_i \subseteq G_i$ with only $\tilde{O}(f \cdot |V(G_i)|)$ edges. The final subgraph is then given by $H = \bigcup_{i=0}^k H_i$. Formally, our expander sparsification lemma can be stated as follows:

Lemma 2.1 (Expander Sparsification). *Given a ϕ -expander n -vertex G with minimum degree $\tilde{\Omega}(f/\phi)$, one can compute a subgraph $H \subseteq G$ such that (i) $|H| = \tilde{O}(f \cdot n/\phi^2)$, (ii) H is an n -vertex $\tilde{\Omega}(\phi^2)$ -expander (and $V(H) = V(G)$), and (iii) the minimum-degree of H is $\Omega(f)$.*

At a high level, the sparsification is based on the useful notion of graph embedding (see e.g., Section 3.16 in [LR99]). An embedding σ of a graph $\hat{G} = (V, \hat{E})$ into a graph $G = (V, E)$ is a function that maps each edge e' in \hat{G} to a path $\sigma(e')$ in G . We say that congestion of the embedding is at most c if each G -edge e appears on at most c paths in $\mathcal{P} = \{\sigma(e') \mid e' \in \hat{E}\}$. The dilation of the embedding is the length of the longest path in \mathcal{P} .

We then show that one can embed an f -regular ϕ -expander $\hat{G} = (V, \hat{E})$ into any given ϕ -expander $G = (V, E)$ with minimum degree f . The congestion and the dilation of the embedding are bounded by $\tilde{O}(1/\phi)$. The output subgraph H is taken as the union of the paths that embed \hat{G} into G .

Finally, we explain the intuition for the correctness of our algorithm. Fix a f -degree subset F and an edge $(u, v) \notin F$. We show that $\text{dist}_{H \setminus F}(u, v) = \tilde{O}(1)$. The interesting case is when $(u, v) \notin H$, and hence (u, v) is in some $\phi = 1/\log n$ expander G_i , obtained by the expander-decomposition. Since $H_i = (V(G_i), E_i)$ is a $\tilde{\Omega}(\phi^2)$ -expander with minimum degree $\Omega(f)$, by Theorem 1.11, $H_i \setminus F$ is also a $\tilde{\Omega}(\phi^2)$ -expander, hence of diameter $\tilde{O}(1)$, which provides the desired stretch guarantees.

2.1.2 Sparse Spanners

Warmup: FD 3-Spanners As a warm-up, we show a simple construction of Δ -FD 3-spanners that can be obtained by a mild adaptation of the classic Baswana-Sen spanner algorithm [BS07]. The following proof assumes familiarity with the Baswana-Sen algorithm, and proceeds quickly over some details used in this algorithm. We have:

Lemma 2.2. *For any graph $G = (V, E)$, one can compute a f -FD 3-spanner $H \subseteq G$ with $|E(H)| = \tilde{O}(fn^{3/2})$.*

Proof. Initially set $H = \emptyset$ and add to H all edges incident to vertices with degree at most $f\sqrt{n}$. We will call these *low-degree vertices*, and the others are *high-degree vertices*. Let S be a random sample of $O(\sqrt{n} \log n)$ vertices, which we call *cluster centers*. With high probability we sample at least $f+1$ cluster centers adjacent to each high-degree vertex. Connect each high-degree vertex to an arbitrary set of $(f+1)$ of its neighboring sampled centers in S . Finally, for each high-degree vertex v and each cluster center $c \in S$, we choose $3f$ arbitrary vertices that are neighbors of both v and c , and we add the edge connecting v to these vertices. (If there are fewer than $3f$ vertices that are neighbors of both v and c , add them all.) Overall, this procedure adds $\tilde{O}(fn^{3/2})$ edges, as desired.

To prove correctness, it suffices to consider an edge $(u, v) \notin F$ where both u, v are high-degree, where F is an arbitrary f -degree set of edge faults. Since u is adjacent to at least $f+1$ cluster centers, there exists a center $c(u) \in S$ such that the edge $(c(u), u)$ survives in $G \setminus F$. If $(u, v) \notin H$, then we have added $\geq 3f$ edges to H connecting v to neighbors in the cluster of $c(u)$. As $c(u), v$ are incident to at most $2f$ edges in F , there exists at least one such vertex x such that $(x, c(u)), (x, v) \notin F$. Thus the path $(u, c(u), x, v)$ survives in $G \setminus F$, and so we have $\text{dist}_{H \setminus F}(u, v) \leq 3$. \square

FD-Spanners with Constant Stretch. Unfortunately, directly extending the (FT) Baswana-Sen approach to the FD-setting, as in [Par22], seems to be quite nontrivial. The reason is that FT Baswana-Sen is based on f edge-disjoint clustering which can handle a total of f faults but cannot handle a f -degree subsets F of $f \cdot n$ failed edges.

Our approach is based on replacing the low-depth edge-disjoint trees that form the clustering in the FT Baswana-Sen algorithm by length-constrained expanders [HRG22] with a sufficiently large minimum degree. Our approach of connecting spanners to length-constrained expanders is inspired by the very recent work of [HHT23b]. They use expanders to provide an alternative, matching-based, implementation of the well-known greedy algorithm for (non fault-tolerant) spanners.

Step (I): From Expander Routing to Rigid Fault-Tolerant Properties. Our approach is based on translating the routing quality of LC-expanders into rigid FT properties against bounded-degree faults. Our starting observation is the following. Consider an (h, s) -length ϕ -expander with minimum degree $\tilde{\Omega}(f \cdot n^\epsilon / \phi)$. We use the fact that length-constrained expanders are good *routers* for h -hop pairs. Specifically, given an f -degree set F , one can route $d = \tilde{\Omega}(n^\epsilon / \phi)$ units of flow over each edge $(u, v) \in F$ along paths of length $O(h \cdot s)$ and with congestion¹ of $c = O(\log n / \phi)$.

We use these dilation and congestion bounds to deduce that $G \setminus F$ is an $(hs)^{O(1/\epsilon)}$ spanner of G , as follows: The solution to the routing instance consists of d many u - v paths of length at most hs for every $(u, v) \in F$. Since these paths might intersect F , we cannot provide an immediate guarantee on the u - v distance in $G \setminus F$ for any $(u, v) \in F$. However, by using an averaging argument, one can claim that for at least $(1 - c/d)$ fraction of the edges in F , their distance (between their endpoints) in $G \setminus F$ is at most hs . We then apply an iterative re-routing procedure that allows us to translate the original routing paths into those that provide at least one fault-free path (among the d many paths) for any $(u, v) \in F$. This re-routing procedure has $\log_{c/d} n = \Theta(1/\epsilon)$ iterations. In each iteration i , we can provide fault-free paths of length $(hs)^i$ for $(1 - c/d)$ fraction of the remaining pairs in F . This provides the intuition for Thm. 1.12.

Step (II): From Rigid-FT in Expanders into Competitive-FT in General Graphs. For simplicity, we present the high-level ideas when using the existential bounds of LC expander decomposition. Our actual algorithms can be implemented in polynomial time, and thus use expander decomposition with somewhat weaker bounds.

To translate the rigid-FT properties of Thm. 1.12 into a construction of FD spanners, we again employ the following two steps: expander decomposition and expander sparsification. Due to the nature of length-restricted expanders, the sparsification arguments are slightly more delicate. The output of the LC expander-decomposition provides us with a (h, s) LC-expander² G' with minimum degree $\tilde{\Omega}(f \cdot n^{O(1/s)})$ and a set of cut edges $C = G \setminus G'$ with $\tilde{O}(f \cdot n^{1+1/s})$ edges.

Our goal is then to sparsify G' into a subgraph $H' \subseteq G$ with a total number of $\tilde{O}(f \cdot n^{1+O(1/s)})$ edges, such that H' has approximately the same routing properties as G' in a way that allows us to apply the robustness lemma of Thm. 1.12 on H' . The sparsification is based on embedding a collection of virtual expanders $\hat{G}_1, \dots, \hat{G}_\ell$ into G' . The output H' is taken as the union of the embedding paths in G' . Each virtual graph \hat{G}_i is defined based on a subset $V_i \subseteq V(G)$ of pairwise distances at most h in G' . These subsets are obtained by using the well-known tool of neighborhood-cover of [ABCP98].

We believe that our approach for translating rigid-FT properties in expander into competitive-FT in general graphs might be also useful for other graph properties (beyond connectivity and distances) and under a wide collection of fault models. The high-level recipe is given some desired graph property

¹The congestion of a path collection \mathcal{P} is at most c if each edge e appears in at most c paths in \mathcal{P} .

²Unlike standard expanders, LC expanders are not required to be connected, hence the output of the expander decomposition is a possibly disconnected LC expander.

Π and a faulty model M and has the following flow: (i) showing that any expander G' approximates Π under any failing event F (where F is determined by the faulty model M). (ii) showing that one can sparsify the expander into $H' \subseteq G'$ such that H' is an expander that satisfies (i), and (iii) apply expander-decomposition (possibly recursively) accompanied with expander sparsification. At the point where the number of inter-expander edges is sufficiently small, add them to the desired output subgraph H . We also note that while we focus here on connectivity and distances, as a byproduct of our approach, we also get that our output subgraphs (e.g., FD certificates) also approximate the dilation and congestion bounds of routing instances in G .

2.2 Blocking Set-Based Constructions

2.2.1 The FD Greedy Algorithm and Analysis

The near-optimal size/stretch tradeoff for FD spanners in Theorem 1.10 is proved by analyzing the natural FD adaptation of the greedy spanner algorithm [ADD⁺93], also analogous to the EFT greedy algorithm studied in prior work [BDPW18, BDR22]. We state this algorithm formally in Section 6.1, but to quickly overview: the spanner H is initially empty, we consider the edges (u, v) of the input graph one at a time in order of increasing weight, and we add each edge (u, v) to the spanner iff there exists a possible fault set F (of max degree $\leq f$) under which $\text{dist}_{H \setminus F}(u, v) > (2k - 1) \cdot w(u, v)$.

The proof of correctness, i.e., that the algorithm returns a correct f -FD $(2k - 1)$ -spanner of the input graph G , is standard. The challenging part of the proof is to show that the output spanner H does not contain too many edges.

The Blocking Set Method. The analysis of the non-fault-tolerant greedy spanner algorithm [ADD⁺93] uses crucially that the output spanner has *high girth*; that is, one can prove that all cycles in H have $> 2k$ edges. It is not still true that the output spanner H from the FD greedy spanner algorithm must have high girth. Still: since the FD greedy algorithm is similar in spirit to the non-fault-tolerant greedy algorithm, it is intuitive to think that its output spanner might be “close to high girth” in some structural sense.

A blocking set formalizes the idea of a graph being *structurally close to high girth*, by asserting that its short cycles admit a particularly small or simple kind of hitting set. The following definition of blocking set is the relevant one for this work:

Definition 2.3 (FD Blocking Sets). *Let $H = (V, E)$ be a graph equipped with a total ordering of its edge set E . A f -fault-degree (FD) k -blocking set \mathcal{B} for H is a set of pairs of the form (e, F_e) , such that:*

- *Each edge e is the first edge of exactly one pair (e, F_e) , each F_e is a set of edges from $E(H)$ of degree $\deg(F_e) \leq f$, and each edge in F_e strictly precedes e in the edge-ordering of H .*
- *For each cycle C in H on $|C| \leq k$ edges, letting e be the latest edge in C in the edge-ordering of H , we have that $F_e \cap C$ is nonempty.*

It is fairly easy to prove that the output spanner H of the FD-greedy algorithm, with its edges ordered by their arrival in the algorithm, has a f -FD blocking set. The focus of the proof then shifts from bounding the number of edges in the particular output spanner from the FD-greedy algorithm to bounding the number of edges in *any* graph that admits a f -FD blocking set.

Path-Counting Methods. In order to explain how we limit the size of a graph with a f -FD blocking set, it may be helpful to first recall the proof of the *Moore bounds*, which are used to bound the maximum possible number of edges in an n -vertex graph H of girth $> 2k$. The Moore bounds are proved using a counting argument over the *edge-simple k -paths* of H . They include two steps:

- A *dispersion lemma*, which shows that no two edge-simple k -paths in H can share endpoints (or else they would form a short cycle), and
- A *counting lemma* which shows a lower bound on the number of edge-simple k -paths in H , where this lower bound is increasing with the number of edges $|E(H)|$.

By comparing the upper and lower bounds on the number of edge-simple k -paths that respectively arise from the dispersion and counting lemmas, and rearranging terms, one gets an upper bound on $|E(H)|$. The interested reader can refer to [Bod23], Section 2, for a recap of this proof of the Moore bounds in full detail.

Since a graph with a blocking set is conceptually interpreted as a graph that is “close to high girth,” it stands to reason that our upper bound on the size of a graph with a blocking set should follow the same basic strategy as the Moore bounds. Indeed, our strategy is to prove analogous dispersion and counting lemmas, and our final bound on $|E(H)|$ follows by comparing these bounds to each other and rearranging terms. However, in our arguments we do not consider *any* edge-simple k -paths in these lemmas: a graph with a f -FD-blocking set can unfortunately have an unbounded number of edge-simple k -paths that share endpoints, and so the dispersion lemma would be impossible. Instead, we restrict our attention to a subset of edge-simple k -paths satisfying a very particular set of properties, which are restrictive enough to enable a version of the dispersion lemma but not so restrictive as to break the counting lemma. We call these MUC k paths. This is an acronym, and the technical properties encoded by this acronym and their role in the proof are discussed more in Section 6.2).

2.2.2 Speeding Up the Greedy Algorithm

The main downside of the FD greedy algorithm is that, in a naive implementation, it takes exponential time. The part that takes exponential time is that in each round, we need to check whether or not there exists a bounded-degree fault set F for which $\text{dist}_{H \setminus F}(u, v) \leq (2k-1) \cdot w(u, v)$. There are $\exp(n)$ many fault sets to consider, and it is not clear how to avoid checking each of these potential fault sets by hand. In fact, there is some evidence that no substantial improvement on this brute force approach will be possible. For the EFT greedy algorithm (which is identical except that it searches only over fault sets of size $|F| \leq f$), the search for a valid fault set encodes an NP-hard problem called LENGTH-BOUNDED CUT (LBC) [BEH⁺06]:

LENGTH-BOUNDED CUT (LBC): Given a graph $G = (V, E)$, vertices s, t , and an integer k , find the least integer f for which there exists an edge set F of size $|F| \leq f$ with $\text{dist}_{G \setminus F}(s, t) > k$.

A nice paper by Dinitz and Robelle [DR20] provides a method to escape this NP-hardness barrier for the EFT greedy algorithm. Dinitz and Robelle essentially show an $O(k)$ -approximation algorithm for LBC in unweighted graphs. Using this approximation algorithm for the edge-test step of the greedy algorithm ultimately costs a factor of $O(k)$ in the size of the blocking set for the output spanner, which translates to an $O(k)$ factor in spanner size as the price to pay for polynomial runtime. The other important observation in [DR20] is that it suffices to solve LBC in an unweighted graph, even when the goal is to build a spanner of a weighted input graph.

It is natural to attempt the same method to improve the runtime of the FD greedy algorithm. The same high-level proof strategy works, but the catch is that we need to solve the following variant:

MIN MAX LBC: Given a graph $G = (V, E)$, vertices s, t , and an integer k , find the least integer f for which there exists an edge set F of degree $\deg(F) \leq f$ with $\text{dist}_{G \setminus F}(s, t) > k$.

As it happens, so-called “min max” cut problems of this type have been studied before through an independent motivating framework, in the research area of *correlation clustering*. Correlation clustering was introduced by Bansal, Blum, and Chawla [BBC04], and it concerns the following model. We are given a (possibly weighted) graph G , with some edges labeled $(+)$ and others labeled $(-)$. We cluster the vertices of G , and a $(+)$ edge is considered to be satisfied if its endpoints lie in the same cluster, while a $(-)$ edge is considered to be satisfied if its edges lie in different clusters. The general goal is to minimize the total weight of unsatisfied edges, although several different objective functions have been studied [Wir10]. Correlation clustering simultaneously generalizes several fundamental graph problems, including MIN S-T CUT.

More recently, researchers have considered *local* objective functions, which encode goals such as minimizing the maximum disagreement weight on any given vertex. This framework was first studied by Charikar, Gupta, and Schwarz [CGS17], who obtained a $O(\sqrt{n})$ approximation algorithm for the problem of local correlation clustering. They also pointed out that, in the same way that correlation clustering captures MAX S-T CUT, their algorithm implies a $O(\sqrt{n})$ approximation algorithm for MIN MAX S-T CUT in undirected graphs. This algorithm has been substantially extended in followup work [KMZ19, JKMM21].

A central message of the work of Charikar et al. [CGS17] is that graph problems of these type are fundamental, and warrant further study. We fully agree with this message, and indeed, this paper provides a concrete application for this direction of study. Our contribution is the following approximation algorithm, which when plugged into the framework of Dinitz and Robelle [DR20], implies our near-optimal FD spanners in polynomial time from Theorem 1.10.

Theorem 2.4. *The problem MIN MAX LBC on an n -vertex input graph with length parameter k has a $O(k \log n)$ approximation algorithm.*

3 Preliminaries

Throughout, we denote by n the number of vertices in the graph input graph. As usual the notions $\tilde{O}(\cdot)$ and $\tilde{\Omega}(\cdot)$ hide poly-logarithmic factors in n .

Vertex Weightings. A vertex weighting $W : V \mapsto \mathbb{Z}_{\geq 0}$ assigns each vertex v a positive integer weight $W(v)$. For a multi-set of undirected edges E' (potentially with (multiple) self-loops a vertices) the vertex weighting $W_{E'}$ assigns each vertex v the number of edges adjacent to v in E' . In this way any (multi-)graph $G = (V, E)$ comes with the so-called degree weighting $\deg_G = W_G = W_E$ which assigns each vertex v the weight $W(v) = \deg_G(v)$.

Our algorithms use the notion of vertex weighting only in the context of length-constrained expanders, nevertheless to make the framework general we provide all definitions based on a given vertex weighting W . If no W is provided then $W(u) = \deg_G(u)$ for every u .

Conductance, Volume, and Expanders. Consider a graph $G = (V, E)$. For a vertex subset $S \subseteq V$ and vertex weighting W , let $\text{Vol}_W(S) = \sum_{v \in S} W(v)$. We sometimes write Vol_G as a shorthand for Vol_{\deg_G} . Let $\partial_G(S) = E(S, V \setminus S)$ be the set of edges in G with one endpoint in S and the other endpoint in $V \setminus S$. The *conductance* of a cut $S \neq \emptyset$ with respect to vertex weighting W is defined by

$$\Phi_G(S, W) = \frac{|\partial_G(S)|}{\min\{\text{Vol}_W(S), \text{Vol}_W(V \setminus S)\}}.$$

For $S = \emptyset$, let $\Phi(S, W) = 0$. The *conductance* of a graph G is given by $\Phi(G, W) = \min_{S \subseteq V, S \neq \emptyset} \Phi_{G,W}(S)$. That is, $\Phi(G, W)$ is the minimum value of $\Phi_{G,W}(S)$ over all non-trivial cuts $S \subseteq V$.

If no vertex-weighting is specified W is assumed to be \deg_G , i.e., the volume $\text{Vol}_G(S)$ counts edges in $\partial_G(S)$ as one, edges induced by S as two, and all other edges as zero and $\Phi_G(S)$ is the conductance with respect to this volume.

Definition 3.1 (ϕ -expander (for W)). *[LS21] Let $G = (V, E)$ be a graph, W be a vertex weighting for V , and $\phi \geq 0$. G is a ϕ -expander for W if $\Phi(G, W) \geq \phi$. We also say W is ϕ -expanding in G . If G is ϕ -expanding for $W = \deg_G$ we simply say G is a ϕ -expander (leaving the "for \deg_G " implied).*

Demands, Routings, Congestion and Dilation. A demand $D : V \times V \rightarrow \mathbb{R}_{\geq 0}$ assigns each pair of vertices u, v a positive amount $D(u, v)$. A demand is integral if it only assigns integer amounts and symmetric if $D(u, v) = D(v, u)$ for all $u, v \in V$. All demands in this paper are symmetric and integral. The *load* of a vertex $v \in V$ in demand D is defined as $\text{load}(v, D) = \max\{\sum_{w \in V} D(v, w), \sum_{w \in V} D(w, v)\}$. Let (max) load of a demand is defined as $\text{load}(D) = \max_v \text{load}(v, D)$. We call a demand W -respecting (or W -unit) if $\text{load}(v, D) \leq W(v)$ and we say a demand D is a unit demand in G if it is \deg_G -respecting. We say that a demand D is supported on $V' \subseteq V$, if $D(u, v) = 0$ for any pair $\langle u, v \rangle \notin V' \times V'$.

Any multi-set of edges E' has an associated integral symmetric demand $D_{E'}$ with $D_{E'}(u, v)$ equal to the number of edges between u and v in E' . Similarly any integral symmetric demand maps to a multi-set of edges or a set of edges $E_D = \{(u, v) \mid D(u, v) > 0\}$.

Definition 3.2 (Dilation, Congestion of a Path Collection). *A routing is a collection of paths $\mathcal{P} = \{P_1, \dots, P_k\}$, the dilation of \mathcal{P} is the maximum length of the paths in \mathcal{P} . The congestion of \mathcal{P} is equal to $\max_{e \in G} |\{P \in \mathcal{P} \mid e \in P\}|$, i.e., the maximum number of times any edge appears in the total over all paths. The quality of \mathcal{P} is defined as $\text{congestion}(\mathcal{P}) + \text{dilation}(\mathcal{P})$. The demand $D_{\mathcal{P}}$ sets $D(u, v)$ to be the number of pathes from u to v in \mathcal{P} . (Iff $D = D_{\mathcal{P}}$ we say " \mathcal{P} routes D or \mathcal{P} is a routing for D .)*

We say a demand D is routable with congestion c and dilation ℓ in G if there exists a routing \mathcal{P} for D in G with congestion congestion c and dilation ℓ . Similarly we say a graph $H = (V', E')$ with $V' \subseteq V$ or a set of edges $F \subseteq V \times V$ is d -routable with congestion c and dilation ℓ in G if $d \cdot D_H$ or $d \cdot D_F$ are routable. A routing for a graph H in G is also called an embedding (of H into G).

We show in appendix that if a good routing or embedding exists than such a routing can also be found efficiently in polymomial time, up to a constant factor loss in congestion. This allows us to provide polynomial time implementation for our expander-based algorithms.

Lemma 3.3 (Length-Constrained Multi-Commodity Flow Routing [HHS23]). *There exists a randomized algorithm which given an n -vertex graph $G = (V, E)$ and a demand D which is routable with congestion $c = \omega(\log n)$ and dilation ℓ in G outputs a routing for D with congestion at most $1.1c$ and dilation ℓ in polynomial time, with high probability.*

The next well-known lemma summarizes the routing properties of ϕ -expanders.

Lemma 3.4. *If a graph G is a ϕ -expander, then any unit demand in G can be routed with congestion $O(\frac{\log n}{\phi})$ and dilation $O(\frac{\log n}{\phi})$.*

For a graph $G = (V, E)$ and a vertex subset $U \subseteq V$, let $G[U]$ be the graph G induced on U . The graph $G(U)$ consists of all edges of $G[U]$ and in addition, it includes (potentially multiple) self-loops on each vertex $u \in U$, such that $\deg_{G(U)}(u) = \deg_G(u)$. I.e., for each edge $(u, v) \in U \times (V \setminus U)$, the graph $G(U)$ includes a self-loop on u .

Theorem 3.5 ([LS21]). *For any any parameter $\phi > 0$, there is a randomized algorithm `ExpDecomp` that given an n -vertex graph $G = (V, E)$ partitions V into clusters $U = (U_1, \dots, U_k)$, such that w.h.p. the following holds:*

- $\sum_i |E(U_i, V \setminus U_i)| = O(\phi m \cdot \log^a n)$ for some constant $a \geq 2$, and
- Every $G(U_i)$ is a ϕ -expander.

The running time of the algorithm is $\tilde{O}(|E(G)|/\phi^2)$.

Length-Constrained Expanders. Length-constrained expanders are generalizations of expanders that give more and separate control over the length of routing paths. We say a demand is h -length if $D(u, v) > 0$ implies that $\text{dist}_G(u, v) \leq h$. The following description of length-constrained expanders follows the summary in [HHG22]: The general notion of (h, s) -length ϕ -expanders where h is called *length*, s is called *length slack*, and ϕ is called *conductance* is defined in terms of so-called “moving cuts” and can be found in [HRG22]. The exact details of their definition will not be important for us here.

All we will need is that up to a constant length factor and a poly-logarithmic congestion factor, (h, s) -hop ϕ -expanders for unit-demands are equivalent to the notion of routing any h -length unit-demand via $(s \cdot h)$ -length paths with congestion roughly $\frac{1}{\phi}$. This is made precise by the following theorem which is the equivalent of Lemma 3.4:

Theorem 3.6 (Flow Characterization of Length-Constrained Expanders (Lemma 3.16 of [HRG22])). *For any $h \geq 1$, $\phi < 1$, (vertex weighting W), and $s \geq 1$ we have that:*

1. *If G is an (h, s) -length ϕ -expander (for W), then every h -hop (W -)unit-demand can be routed in G along $(s \cdot h)$ long paths with congestion $O(\log(n)/\phi)$.*
2. *If G is not an (h, s) -hop ϕ -expander (for W), then some h -hop (W -)unit demand cannot be routed in G along $(\frac{s}{2} \cdot h)$ long paths with congestion $1/2\phi$.*

Theorem 3.7 ([HRG22, HHG22]). *For any graph $G = (V, E)$, conductance $\phi > 0$, length bound h , length slack $s > 2$ and vertex weighting W there exists a set of edges $C \subset E$ of size at most $O(h\phi|W|n^{s^{-0.1}} \log n)$ such that $G \setminus C$ is an (h, s) -length $\tilde{\Omega}(\phi)$ -expander with respect to W .*

Our use of LC-expanders w.r.t weighting W is mainly for the purpose of using the above mentioned properties of expander decomposition. That is, we enjoy the fact that $|C|$ depends on $|W|$, rather than on $|E(G)|$, to provide a cleaner algorithmic description in Sec. 5.3.

There also exists a randomized algorithm to compute the expander decomposition guaranteed by Theorem 3.5 with high probability and in polynomial time [HRG22, HHG22, HHT23a].

Neighborhood Cover. For a graph $G = (V, E)$ and a positive integer r , the r -neighborhood of a vertex v is defined by $N_{r,G}(u) = \{v \mid \text{dist}_G(u, v) \leq r\}$. Given a graph G and parameters β, r , an (r, β) -neighborhood cover is a collection of subgraphs G_1, \dots, G_k with the following properties: (i) the diameter of each G_i is $O(\beta \cdot r)$, (ii) for every $v \in V(G)$, there is a subgraph G_i that contains the entire r -neighborhood of v , i.e., $N_{r,G}(u) \subseteq G_i$ and (iii) each vertex appears in at most $\beta \cdot n^{1/\beta}$ subgraphs.

Lemma 3.8 ([ABCP98]). *Given a graph G and parameters r, β , there is a polynomial time algorithm NeighborhoodCover that computes the (r, β) -neighborhood cover of G .*

Roadmap. In Section 4, we provide the construction of FD certificates. In Section 5, we provide constructions of FD $O(1)$ -spanners that are based on LC-expanders. Then, in Section 6, we provide polynomial time algorithms for computing FD t -spanners with nearly optimal density (conditioned on the girth congestion). Finally, lower bounds arguments are provided in Section 7.

4 Sparse Connectivity Certificates against Bounded-Degree Faults

In this section we provide an algorithm for computing sparse FD-certificates and establish Theorem 1.4. In fact, our construction will provide an FD $\tilde{O}(1)$ -spanners. The structure of this section is as follows. First, we show that expanders are rigid to bounded-degree faults. Then we present an algorithm for computing rigid FD certificates for expander graphs, and finally we present an algorithm for computing competitive FD certificates for general graphs.

Rigidity of Expanders to Bounded Degree Faults. We show that ϕ -expanders with minimum degree $\tilde{\Omega}(f/\phi)$ remain good expanders under the failing on any f -degree edge set.

Proof of Theorem 1.11. Fix a cut $(S, V \setminus S)$ in G and let $\text{Vol}_G(S) \leq \text{Vol}_G(V \setminus S)$. Then, by the definition of expanders we have $|\partial_G(S)| \geq \phi \cdot |S| \cdot f'$. Since each vertex in S is incident to at most f edges in F , we have that $|\partial_{G \setminus F}(S)| \geq |\partial_G(S)| - |S|f \geq (\phi \cdot |S| \cdot f')/2$. Moreover, $\text{Vol}_G(V') \geq \text{Vol}_{G \setminus F}(V')$ for any subset $V' \subseteq V$. We therefore have that:

$$\begin{aligned} \Phi_{G \setminus F}(S) &= \frac{|\partial_{G \setminus F}(S)|}{\min\{\text{Vol}_{G \setminus F}(S), \text{Vol}_{G \setminus F}(V \setminus S)\}} \geq \frac{|\partial_G(S)| - |S|f}{\min\{\text{Vol}_G(S), \text{Vol}_G(V \setminus S)\}} \\ &= \frac{|\partial_G(S)|}{\min\{\text{Vol}_G(S), \text{Vol}_G(V \setminus S)\}} - \frac{|S|f}{\min\{\text{Vol}_G(S), \text{Vol}_G(V \setminus S)\}} \\ &\geq \Phi_G(S) - \frac{|S|f' \cdot \phi}{2 \min\{\text{Vol}_G(S), \text{Vol}_G(V \setminus S)\}} \geq \Phi_G(S) - \frac{|S|f' \cdot \phi}{2|S|f'} \\ &= \Phi_G(S) - \phi/2 \geq \phi/2. \end{aligned}$$

□

In Appendix A, we prove a more delicate robustness argument which holds for any graph G (even not expanders) provided that G admits *non-trivial* routing solutions for f -degree subsets of edges F , i.e., in which the output flow uses edges not in F .

Lemma 4.1. *[From Low-Congestion Routing to Fault-Tolerance] Consider an n -vertex graph G with minimum degree f' and let F be a f -degree subset of edges such that F is d -routable in G for $d = f'/f$ with congestion of $c < d$. Then, $G \setminus F$ is connected.*

Expander Sparsification. Our certificate computation is based on our ability to sparsify a possibly dense expanders while preserving their fault-tolerant properties w.r.t. their minimum degree. The following lemma summarizes the properties of our sparsification procedure:

Theorem 4.2 (Expander Sparsification). *For any $\phi \in [0, 1]$ and $f' \in [\Omega(\log n), n]$, there is an algorithm **Sparsify** that given an n -vertex ϕ -expander G with min-degree f' outputs a subgraph $H \subseteq G$ such that:*

- H is an $\tilde{\Omega}(\phi^2)$ -expander.
- $|E(H)| = O(f' \cdot n \cdot \log n/\phi)$, and
- the minimum degree of H is $\tilde{\Omega}(f' \cdot \phi)$.

We need the following auxiliary claim. We say that a graph $H \subseteq G$ is a *minimal* graph that embeds a graph \hat{G} into G if H is the union of the G -paths that embeds \hat{G} into G . The proof of the next lemma is provided in appendix.

Lemma 4.3. Let $H = (V, E_H) \subseteq G$ be a minimal graph that embeds a f -regular $(1/\log n)$ -expander $\widehat{G} = (V(H), \widehat{E})$ into G , such that the collection of embedded paths have congestion and dilation of $\tilde{O}(1/\phi)$. Then H is a $\tilde{\Omega}(\phi^2)$ -expander.

We are now ready to provide the proof of Thm. 4.2.

Proof of Thm. 4.2. Let $\widehat{G} = (V(G), \widehat{E})$ be an f' -regular $(1/\log n)$ -expander. We define a demand $D_{\widehat{E}}$ where $D_{\widehat{E}}((u, v)) = 1$ for every $(u, v) \in \widehat{E}$ and $D_{\widehat{E}}(u, v) = 0$, otherwise.

Since \widehat{G} is f' -regular, D is unit-demand in G , and can therefore be solved in G with congestion and dilation of $O(\log n/\phi)$ by Lemma 3.4. Moreover, the collection of routing paths \mathcal{P} can be computed in polynomial time using Lemma 3.3. The output subgraph H consists of the union of all paths $P \in \mathcal{P}$.

We next show that H satisfies the desired properties. Since H is a minimal graph for embedding \widehat{G} into G , by Lemma 4.3, we have that H is a $\tilde{\Omega}(\phi^2)$ -expander. By construction, we also have that $|E(H)| = O(f' \cdot n \cdot \log n/\phi)$, as H consists of $f'n$ paths, each of length at most $O(\log n/\phi)$. It remains to bound the minimum degree in H . Consider a vertex u and its collection of f neighbors in \widehat{G} . Since the congestion of the paths in \mathcal{P} is $O(\log n/\phi)$, we get that u is incident to at least $f = \tilde{\Omega}(f' \cdot \phi)$ distinct edges on the output u - v paths of the routing instance, which implies that u is incident to at least f neighbors in H . \square

The Algorithm for f -Faulty-Degree Certificates. We use the following procedure **MinDegree** that given a graph G and integer f' , outputs a subgraph $G' \subseteq G$ with minimum degree f' such that $|G \setminus G'| = O(f'n)$.

Lemma 4.4. Given an n -vertex graph G and integer f' , there is an algorithm **MinDegree** that computes a subgraph $G' \subseteq G$ with minimum degree f' and such that $|G \setminus G'| = O(f'n)$.

Proof. Letting $G' = G$, the algorithm iteratively omits a vertex with degree at most f' in the remaining graph G' until no such vertex exists. The algorithm omits at most $f'n$ edges from G which proves the desired bound. \square

We are now ready to present Algorithm **FDCertificate** which computes the desired f -FD certificate H with $\tilde{O}(f \cdot n)$ edges. The algorithm has $\ell = O(\log n)$ iterations. In each iteration i , the algorithm gets as input a subgraph $G_i \subseteq G$ with $\Omega(f'n)$ edges, where initially $G_1 = G$. The output of the iteration is given by a subgraph H_i and a subgraph $G_{i+1} \subseteq G_i$ to be provided as input to the next iteration $i+1$. The final output subgraph H is then defined by $H = \bigcup_{i=1}^{\ell} H_i$.

We now focus on iteration $i \geq 1$. The algorithm starts by applying Procedure **MinDegree** of Lemma 4.4 on G_i with degree threshold $f' = \tilde{O}(f)$. Let G'_i be the output subgraph of this procedure. Next, the algorithm applies expander decomposition on G'_i which partitions G into vertex-disjoint ϕ -expanders $G_{i,1}, \dots, G_{i,k_i}$ for $\phi = 1/(\log n)^{a+1}$ and a collection of at most $m/2$ inter-expander edges, where a is the constant in Theorem 3.5.

On each sufficiently dense ϕ -expander $G_{i,j}$ (i.e., (with $|G_{i,j}| \geq 2f' \cdot |V(G_{i,j})|$)), the algorithm applies Procedure **Sparsify** of Thm. 4.2 to obtain a $\tilde{\Omega}(\phi^2)$ -expander $H_{i,j} \subseteq G_{i,j}$. The output of the i^{th} phase is given by a subgraph $H_i = \bigcup_{j=1}^{k_i} H_{i,j}$ and the next graph G_{i+1} is the defined by the union of all inter-expander edges. As $|G_{i+1}| \leq |G_i|/2$ the process terminates within ℓ iterations. See a detailed description below.

Algorithm FDCertificate

Input: A graph G , integer f .

Output: An f -FD certificate $H \subseteq G$ with $\tilde{O}(f \cdot n)$ edges.

- $G_1 \leftarrow G$ and $\phi = 1/\log^{a+1} n$ (see Theorem 3.5).
- $f' \leftarrow \lceil f/\phi^5 \rceil$ and $f'' \leftarrow \lfloor f' \cdot \phi \rfloor$.
- For $i = \{1, \dots, \ell = 2 \log n\}$ do:
 1. $G'_i = \text{MinDegree}(G_i, f')$.
 2. $(U_{i,1}, \dots, U_{i,k_i}) = \text{ExpDecomp}(G'_i, \phi)$.
 3. For every $j \in \{1, \dots, k_i\}$ do:
 - Let $G_{i,j} = G'_i[U_{i,j}]$.
 - If $|G_{i,j}| \leq 2f' \cdot |U_{i,j}|$: $H_{i,j} \leftarrow G_{i,j}$.
 - Otherwise: $H_{i,j} \leftarrow \text{Sparsify}(G_{i,j}, f'')$.
 4. $G_{i+1} \leftarrow \bigcup_i G \setminus \bigcup_{j=1}^{k_i} G_{i,j}$.
 5. $H_i \leftarrow (G_i \setminus G'_i) \cup \bigcup_{j=1}^{k_i} H_{i,j}$.
 6. If $|G_{i+1}| \leq f'n$, set $H_{i+1} = G_{i+1}$, $\ell = (i+1)$ and QUIT.
- $H = \bigcup_{i=1}^{\ell} H_i$.

Lemma 4.5. H is an f -FD $\tilde{O}(1)$ -spanner (and hence also an f -FD certificate) with $|E(H)| = \tilde{O}(f \cdot n)$.

Proof. We start with the size bound by showing that $|H_i| = \tilde{O}(f \cdot n)$. By Lemma 4.4, we have $|G_i \setminus G'_i| = \tilde{O}(fn)$. By Thm. 4.2, $|H_{i,j}| = \tilde{O}(f \cdot |V(U_{i,j})|)$. The size bounds follows that $H_{i,1}, \dots, H_{i,k_i}$ are vertex-disjoint. We next claim that the algorithm quits within $2 \log n$ iterations. This holds as by the stopping criteria, and by the fact that $|G_{i+1}| \leq |G_i|/2$ (as guaranteed by Thm. 4.2).

Min-Degrees. We next claim that the minimum-degree of any ϕ -expander $G_{i,j} = G'_i[U_{i,j}]$ is at least f'' . By the properties of expander-decomposition, we have that the graph $\tilde{G} = G'_i(U_{i,j})$ is also a ϕ -expander. Recall that the graph \tilde{G} consists of all edges of $G_{i,j}$, and in addition, each vertex v has a self-loop for every edge $(u, v) \in G'_i \setminus G_{i,j}$. By adding these self-loops, we have that the degree of v in \tilde{G} is the same as its degree in G'_i . Since the minimum degree of G'_i is at least f' , we have that $\min\{\text{Vol}_{\tilde{G}}(\{v\}), \text{Vol}_{\tilde{G}}(V \setminus \{v\})\} \geq f'$ for every $v \in U_{i,j}$. By the expansion of \tilde{G} , for every $v \in U_{i,j}$, it then holds that $\deg_{G_{i,j}}(v) \geq \phi \cdot f' \geq f''$.

Stretch. Fix an f -degree set F and $(u, v) \in E \setminus F$. The interesting case is when $(u, v) \notin H$, and therefore in such a case there exist indices i, j such that $(u, v) \in G_{i,j} \setminus H_{i,j}$. By Theorem 4.2, $H_{i,j}$ is a $\tilde{\Omega}(\phi^2)$ -expander with minimum degree $\tilde{f} = \tilde{\Omega}(f'' \cdot \phi) = \tilde{\Omega}(f' \cdot \phi^2)$. Since $\tilde{f} = \tilde{\Omega}(f/\phi^3)$, by Theorem 1.11, we have that $H_{i,j} \setminus F$ is also $\tilde{\Omega}(\phi^2)$ expander. Concluding that $\text{dist}_{H_{i,j} \setminus F}(u, v) = \tilde{O}(1/\phi^2) = \tilde{O}(1)$. \square

5 Sparse Spanners against Bounded-Degree Faults

5.1 Robustness of Length-Constrained Expanders against Bounded Degree Faults

In this section, we show that dense LC-expanders G are robust to bounded-degree faults F , in the sense that $G \setminus F$ is a $O(1)$ -spanner for G . We show the following (a generalization of Thm. 1.12):

Theorem 5.1. Consider an (h, s) -length ϕ -expander G for vertex weighting W where $W(u) = \tilde{\Omega}(f \cdot n^{\epsilon}/\phi)$, for every $u \in V(G)$. Then, for every f -degree subset of edges F , it holds that $\text{dist}_{G \setminus F}(u, v) \leq (h \cdot s)^{O(1/\epsilon)}$ for every $(u, v) \in F$.

To prove the theorem, we use the routing properties of LC-expanders, and in particular the following corollary implied by Theorem 3.6:

Corollary 5.2 (Routing on Dense LC Expanders). *Let G be an n -vertex $O(h, s)$ -length ϕ -expander for vertex weighting W where $W(u) = \tilde{\Omega}(f \cdot n^\epsilon / \phi)$ for every $u \in V(G)$. Then every f -degree subset of edges F is $\Theta(n^\epsilon / \phi)$ -routable in G with dilation $h \cdot s$ and congestion $O(\log n / \phi)$.*

By Cor. 5.2, it is sufficient to show the following key lemma that translates routing properties of edge sets $F \subseteq G$ into fault-tolerant properties in $G \setminus F$. Recall that a given set of edges F is d -routable with congestion c and dilation ℓ , if one can solve a routing instance in which d units of flow are required to be sent over each edge $e \in F$ along paths with congestion c and dilation ℓ .

Theorem 5.3. *Let G be an n -vertex graph and let $F \subseteq E(G)$ be such that F is d -routable with congestion $c \leq d$ and dilation ℓ . Then, for every $(u, v) \in F$, it holds that $\text{dist}_{G \setminus F}(u, v) \leq \ell^k$ where $k = 2 \log n / \log(d/c)$.*

Proof. Set $k = 2 \log_{d/c} n$. We partition F into at most k subsets F_1, \dots, F_k , such that for every $(u, v) \in F_i$ it holds that $\text{dist}_{G \setminus F'_i}(u, v) \leq \ell$ where $F'_1 = F$ and $F'_i = F \setminus \bigcup_{j \leq i-1} F_j$ for $i \in \{2, \dots, k\}$. This is done in k iterations.

In iteration $i \in \{1, \dots, k\}$, we define a routing instance where the routing pairs are the edges in F'_i and with uniform demand $D(u, v) = d$ for every $(u, v) \in F'_i$ and $D(u, v) = 0$ otherwise. By solving this routing instance in G , we get the routing paths $\mathcal{P}_i = \{\mathcal{P}(u, v) \mid (u, v) \in F'_i\}$ where each $\mathcal{P}(u, v)$ is a multi-set of d paths connecting u and v in G . The set F_i is then defined by:

$$F_i = \{(u, v) \in F'_i \mid \exists P \in \mathcal{P}(u, v) \text{ such that } P \cap F'_i = \emptyset\}.$$

That is, F_i consists of all edges $(u, v) \in F'_i$ whose $\mathcal{P}(u, v)$ set contains at least one path that does not intersect F'_i . We next claim that F_1, \dots, F_k is a partitioning of F .

We show that $|F_i| \geq (1 - c/d)|F'_i|$ for every $i \in \{1, \dots, k\}$, which establishes the claim. This is can be shown by a simple averaging argument: The number of paths in \mathcal{P}_i that intersects F'_i is at most $c \cdot |F'_i|$. Hence $|F'_i \setminus F_i| \leq c \cdot |F'_i|/d$. Consequently, $F'_{k+1} = \emptyset$ and $\bigcup_{i=1}^k F_i = F$.

Finally, we show by induction on $i \in \{1, \dots, k\}$ that $\text{dist}_{G \setminus F}(u, v) \leq \ell^i$. The base case follows by the definition of F_1 , as $F'_1 = F$. Assume that the claim holds up to $(i-1)$ and consider i . Fix $(u, v) \in F_i$ and consider some u - v shortest path P in $G \setminus F'_i$. By the definition of F_i , we have that $|P| \leq \ell$. Since $P \cap F \subseteq \bigcup_{j \leq i-1} F_j$, we can replace each edge $(x, y) \in \bigcup_{j \leq i-1} F_j$ with a fault-free x - y path $P_{x,y} \subseteq G \setminus F$. By induction assumption we have that $|P_{x,y}| \leq \ell^{i-1}$. Altogether, we get that $\text{dist}_{G \setminus F}(u, v) \leq |P| \cdot \ell^{i-1} \leq \ell^i$. \square

Theorem 1.12 then follows by Lemma 5.3 and Lemma 3.6.

5.2 Length-Constrained Expander Sparsification

Our goal in this section is to sparsify an expander G into $H \subset G$ while maintaining the congestion and dilation parameters of routing instances. Our key sparsification result shows:

Theorem 5.4 (Key Sparsification Lemma). *Consider an n -vertex graph $G = (V, E)$ with and a subset $V' \subseteq V$ satisfying that any demand D supported on V' with $\text{load}(D) \leq f'$, for some $f' = \Omega(\log n)$, can be routed with dilation ℓ_G and congestion c_G . Then, there is a randomized polynomial-time procedure `LCExpanderSparsify` that computes a subgraph $H \subseteq G$ which satisfies the following properties w.h.p.:*

- $\min_{v \in V'} \deg_H(v) = \Omega(f'/c_G)$.
- The diameter of H is $O(\log_{f'} n \cdot \ell_G)$.
- Any demand function D supported on V' such that $\text{load}(D) \leq \Theta(f')$ is routable in H with congestion $\tilde{O}(c_G)$ and dilation $O(\ell_G \cdot \log_{f'} n)$.

- $|E(H)| \leq f' \cdot |V'| \cdot \ell_G$.

In the following we describe Algorithm **LCExpanderSparsify** and prove Thm. 5.4.

Algorithm LCExpanderSparsify. Algorithm **LCExpanderSparsify** that receives as input a graph $G = (V, E)$, a subset $V' \subseteq V$ and an integer parameter f' . The algorithm computes $H \subseteq G$ that preserves the routing quality of G on any demand D supported on V' with $\text{load}(D) \leq f'$.

Our approach is based on embedding a “good constant-hop expander” \widehat{G} on the vertex set V' in the graph G , and letting H be the union of the paths that embeds \widehat{G} in G . Here we take $\widehat{G} = (V', \widehat{E})$ to be the random graph $G(n, p)$ with $p = \Theta(f'/n)$ because random graphs with polynomially large degrees are excellent routers allowing any unit demand to be routed over constant length paths with low congestion. More precisely it is true that with high probability any unit-demand in \widehat{G} is be routable with dilation $\ell_{\widehat{G}} = O(\log_{f'} n)$ and congestion $c_{\widehat{G}} = O(\log n)$. We prove this in Lemma B.1. Also note by setting the constant hidden in the probability p to be small enough, w.h.p., all degrees in \widehat{G} are in $[f'/a, f']$ for some constant $a \geq 1$.

Next, define a routing instance in G with the following demands values: $D(u, v) = 1$ for every $(u, v) \in \widehat{E}$ and $D(u, v) = 0$, otherwise. Since D is supported on V' and $\text{load}(D) \leq f'$, by the properties of G , the demand is routable in G by a collection of paths $\mathcal{P} = \{P(x, y) \mid (x, y) \in \widehat{E}\}$ with dilation at most ℓ_G and congestion c_G . Moreover, these paths can be computed in polynomial time using Lemma 3.3. The output subgraph $H \subseteq G$ is then given by $H = \bigcup_{P \in \mathcal{P}} E(P)$.

Size and Diameter. The size bound is immediate as $|E(H)| \leq f' \cdot |V'| \cdot \ell_G$. We next consider the diameter bound. The diameter of \widehat{G} is $O(\log_{f'} n)$, and as each edge in \widehat{G} translates into a H -path of length ℓ_G , we have that the diameter of H is $O(\log_{f'} n \cdot \ell_G)$.

Minimum Degree. Since each G -edge appears on c_G paths in \mathcal{P} , and since each u is the source vertex of $\Theta(f')$ paths in \mathcal{P} , we have that u is incident to $\Omega(f'/c_G)$ edges in H . Note that some vertices in $V(H) \setminus V'$ might have arbitrarily small or large degrees.

Routing. Fix a demand D supported on V' with $\text{load}(D) \leq f'/a$ and recall that the minimum degree of \widehat{G} is at least f'/a . Define $I_H = \{\langle u, v \rangle \mid D(u, v) \neq \emptyset\}$, and as D is supported on V' , we have that all pairs $\langle u, v \rangle \in I_H$ are in $V' \times V'$.

Since D is a unit-demand instance in \widehat{G} , it can be routable in \widehat{G} with dilation at most $\ell_{\widehat{G}} = O(\log_{f'} n)$ and congestion $c_{\widehat{G}} = \widetilde{O}(1)$. Let $\mathcal{Q} = \{\mathcal{Q}(u, v) \mid \langle u, v \rangle \in I_H\}$ be the collection of output paths in \widehat{G} , where $\mathcal{Q}(u, v)$ is a multi-set of $D(u, v)$ paths connecting u and v in \widehat{G} . Since each edge (x, y) in \widehat{G} translates into a path $P(x, y)$ in H , the routing solution in \widehat{G} can be translated into a routing solution in H , as follows. Letting $Q = [u = x_1, \dots, x_k = v] \in \mathcal{Q}(u, v)$, then let

$$g(Q) = P(x_1, x_2) \circ \dots \circ P(x_{k-1}, x_k).$$

Then, define $\mathcal{P}'(u, v) = \{g(Q) \mid Q \in \mathcal{Q}(u, v)\}$ and $\mathcal{P}' = \bigcup_{\langle u, v \rangle \in I_H} \mathcal{P}'(u, v)$. We finally note that path collection \mathcal{P}' has dilation $\ell_G \cdot \ell_{\widehat{G}}$ and congestion is at most $c_{\widehat{G}} \cdot c_G$. The lemma follows.

5.3 Expander-Based Computation of FD-Spanners

We are now ready to present Alg. **FDSpanner**. The algorithm computes the FD spanner H by taking the following steps. First, it computes a subgraph $G' \subseteq G$ with minimum degree $f' = \widetilde{\Theta}(f \cdot n^{1/t})$ using Lemma 4.4. The edges in $G \setminus G'$ are added to H . Then, it applies the LC-expander decomposition of Theorem 3.7 with respect to expansion parameter $\phi = 1/\log n$, $h = O(t)$, $s = t^{20}$ and vertex weighting $W(u) = f \cdot n^{1/(2t)}$ for every $u \in V$. The output of this expander decomposition is a set C of cut edges such that $G'' = G' \setminus C$ is a (h, s) -length ϕ -expander w.r.t W . All cut edges C are then added to H . Next, the algorithm applies the neighborhood-cover procedure of Lemma 3.8 on G'' with radius parameter

$r = O(t)$ and overlap of $\beta = O(tn^{1/r})$. Formally, let $(G_1, \dots, G_\ell) = \text{NeighborhoodCover}(G'', r = 1, \beta = 3t)$. For every $j \in \{1, \dots, \ell\}$, it applies the expander-sparsification procedure of Thm. 5.4 to compute a subgraph H_j that preserves the routing quality of G'' with respect to pairs in $V[G_j]$. Each H_j is then added to H . This completes the description, and the pseudocode appears below.

Algorithm FDSpanner

Input: A graph G , integer f , a stretch parameter t .
Output: An f -FD $t^{O(t)}$ -spanner $H \subseteq G$ with $\tilde{O}(f \cdot n^{1+1/t})$ edges.

- $G' \leftarrow \text{MinDegree}(G, f')$ for $f' = \tilde{\Theta}(f \cdot n^{1/t})$.
- $\phi \leftarrow 1/\log n$, $W(u) = f \cdot n^{1/(2t)}$ for every $u \in V$.
- $C = \text{LCExpanderExpDecomp}(G', \phi, h = 3t, s = t^{20}, W)$.
- $G'' \leftarrow G' \setminus C$ and $H_0 \leftarrow C$.
- $(G_1, \dots, G_\ell) = \text{NeighborhoodCover}(G'', r = 1, \beta = 5t)$.
- For every $j \in \{1, \dots, \ell\}$ do:
 - $H_j \leftarrow \text{LCExpanderSparsify}(G'', V(G_j), f \cdot n^{1/(2t)})$.
- $H \leftarrow (G \setminus G') \cup \bigcup_{j=0}^{\ell} H_j$.

Size Analysis. By Lemma 3.7, we have that $|H_0| = \tilde{O}(h\phi \cdot n \cdot f \cdot n^{1/(2t)} \cdot n^{1/t^2}) = \tilde{O}(f \cdot n^{1+1/t})$.

By Theorem 5.4, $|H_j| = \tilde{O}(f \cdot n^{1/(2t)}(h \cdot s) \cdot |V(G_j)|)$. In addition, by Lemma 3.8, each vertex appears in $O(t \cdot n^{1/(5t)})$ subgraphs G_i . Therefore, as $t = O(\log n)$, we have that $\sum_j |H_j| = \tilde{O}(f \cdot n^{1+1/t})$. Finally, by Lemma 4.4, we have that $|G \setminus G'| = \tilde{O}(f \cdot n^{1/t})$.

The Stretch Argument. Let F be an f -degree faulty set. It is sufficient to show that for every $(u, v) \in E(G) \setminus F$, it holds that $\text{dist}_{H \setminus F}(u, v) = t^{O(t)}$. Fix $(u, v) \in E(G) \setminus F$. The intersetting case is clearly when $(u, v) \notin H$. This in particular implies that $(u, v) \in G''$, and thus by the properties of the neighborhood cover of Lemma 3.8, there exists j such that $(u, v) \in G_j$. Since G'' is a ϕ -expander w.r.t W and $W(u) = f \cdot n^{1/(2t)}$ for every u , we have:

Corollary 5.5. *For every $j \in \{1, \dots, \ell\}$, any demand D supported on $V[G_j]$ with $\text{load}(D) \leq f \cdot n^{1/(2t)}$ can be routed in G'' along paths of dilation at most $O(t^{21})$ and congestion $O(\log n/\phi)$.*

Our goal is to show that $\text{dist}_{H \setminus F}(x, y) \leq t^{O(t)}$ for every $(x, y) \in F$. To prove the claim, we show that (i) H_j is a good router for the edges in $F_j = F \cap G_j$ for every $j \in \{1, \dots, \ell\}$, and consequently that (ii) H is a good router for F .

Corollary 5.6. *$F_j = F \cap G_j$ is d -routable in H_j with dilation $O(t^{22})$ and congestion $c = O(\log n/\phi)$ for $d = n^{1/(3t)} \cdot c$.*

Proof. By Corollary 5.5 and Theorem 5.4, we have that any demand D supported on $V(G_j)$ of load at most $k = \Theta(f \cdot n^{1/(2t)})$ is routable in H_j with congestion of $\tilde{O}(1/\phi)$ and dilation $O(t^{22})$. In our setting, the demands are $D(u, v) = d$ for every $(u, v) \in F_j$ and $D(u, v) = 0$, otherwise. Hence, D is supported on $V(G_j)$ and $\text{load}(D) \leq d \cdot f \leq k$. Therefore, D is routable in H_j with congestion $\tilde{O}(1/\phi)$ and dilation $O(t^{22})$. \square

Observation 5.7. F is d' -routable in H with dilation $O(t^{22})$ and congestion $c' = \tilde{O}(n^{1/(5t)})$ for $d' = n^{1/(20t)} \cdot c'$.

Proof. By Corollary 5.6, each F_j is d -routable in H_j with dilation $O(t^{22})$ and congestion $c = O(\log n/\phi)$ for $d = n^{1/(3t)} \cdot c$.

By Lemma 3.8, each vertex appears in at most $q = O(tn^{1/(5t)})$ sets of G_1, \dots, G_ℓ , and that $\bigcup_j F_j = F$. We therefore get that F is d -routable in H with congestion of $c' = c \cdot q$ and dilation $O(t^{22})$. The claim follows as $d' \leq d$. \square

By combining Cor. 5.7 with Lemma 5.3, we get that $\text{dist}_{H \setminus F}(x, y) = t^{O(t)}$. Since $\text{dist}_H(u, v) \leq \text{dist}_{H_j}(u, v) = O(t)$, it holds that $\text{dist}_{H \setminus F}(u, v) = t^{O(t)}$.

6 Improved Existential Upper Bounds for FD Spanners

In this section, we provide near-tight bounds for FD spanners for any fixed stretch value. In contrast to the previous constructions, the approach provided in this section is structural and does not use expanders.

Theorem 6.1. *For all positive integers n, k, f , every n -node undirected weighted graph has a f -FD $(2k - 1)$ -spanner on at most $f^{1-1/k} n^{1+1/k} \cdot O(k)^k$ edges.*

By the lower bound in the following section, this bound is optimal for constant k , but may be suboptimal in its $O(k)^k$ dependence. We refer back to Section 2.2 for a technical overview of the proof strategy used in this section, although some of it will be repeated here.

6.1 Proof Setup

We construct our spanners using the following greedy algorithm, which is the natural adaptation to our setting of the standard greedy algorithm for non-faulty spanners [ADD⁺93], or more specifically, the greedy algorithm used in previous work on fault tolerant spanners [BDPW18, BP19, BDR22].

Algorithm GreedyFDSpanner

Input: A graph G and a positive integer stretch parameter k .

Output: An f -FD $(2k - 1)$ -spanner $H \subseteq G$ with at most $f^{1-1/k} \cdot n^{1+1/k} \cdot O(k)^k$ edges.

- Let $H \leftarrow (V, \emptyset, w)$ be the initially-empty spanner
- For each edge $(u, v) \in E$ in order of nondecreasing $w(u, v)$
 - If there exists an edge set $F_{(u,v)}$ of max degree f such that $\text{dist}_{H \setminus F}(u, v) > (2k - 1) \cdot w(u, v)$
 - * Add (u, v) to H
- Return H

We note that a naive implementation of this greedy algorithm runs in exponential time, since we would need to search over all possible choices of fault set $F_{(u,v)}$ in each round. We are not currently aware of a speedup in this step. The proof of correctness, that H is indeed a f -FD $(2k - 1)$ -spanner of the input graph G , is entirely standard. We give this next.

Lemma 6.2 (Correctness of the FD-Greedy Algorithm). *The output graph H from Algorithm GreedyFDSpanner is a f -FD $2k - 1$ spanner of the input graph G .*

Proof. Let F be any set of faulty edges of degree f , and so our goal is to show that $H \setminus F$ is a $(2k - 1)$ -spanner of $G \setminus F$. As is well known (e.g. [ADD⁺93]), it suffices to prove that for each edge $(u, v) \in E(G \setminus F)$, we have

$$\text{dist}_{H \setminus F}(u, v) \leq (2k - 1) \cdot w(u, v).$$

Indeed, this inequality follows straightforwardly from the greedy algorithm. If the greedy algorithm adds (u, v) to H , then we have $\text{dist}_{H \setminus F}(u, v) \leq w(u, v)$. If not, then by construction we have $\text{dist}_{H \setminus F}(u, v) \leq (2k - 1) \cdot w(u, v)$ at the moment (u, v) is considered by the greedy algorithm, and any further edges added to H after this point can only decrease $\text{dist}_{H \setminus F}(u, v)$. \square

The goal of the rest of the proof is to control the number of edges $|E(H)|$ in the output spanner. This is the more involved part of the proof. Our bound on $|E(H)|$ will use the blocking set method, as in [BP19, BDR22, BDN23]. The relevant kind of blocking set for this problem will be the following:

Definition 6.3 (FD Blocking Sets). *Let $H = (V, E)$ be a graph equipped with a total ordering of its edge set E . A f -fault-degree (FD) k -blocking set \mathcal{B} for H is a set of pairs of the form (e, F_e) , such that:*

- *Each edge e is the first edge of exactly one pair (e, F_e) , each F_e is a set of edges from $E(H)$ of maximum degree f , and each edge in F_e strictly precedes e in the edge-ordering of H .*
- *For each cycle C in H on $|C| \leq k$ edges, letting e be the latest edge in C in the edge-ordering of H , we have that $F_e \cap C$ is nonempty.*

Lemma 6.4 (Blocking Set Exists). *The output graph H from the greedy algorithm, with $E(H)$ ordered by the order its edges were added in the greedy algorithm, has a f -FD $2k$ -blocking set \mathcal{B} .*

Proof Sketch. Each time we add an edge e in the greedy algorithm, we do so because of a possible fault set F_e . (There might be many possible fault sets forcing us to add e , in which case we fix F_e to be any one of them.) We then add (e, F_e) as a pair to the blocking set \mathcal{B} .

Now letting \mathcal{B} be the final blocking set at the end of the greedy algorithm, it is immediate by construction that each edge $e \in E(H)$ is the first edge of exactly one pair in \mathcal{B} , and that each F_e has maximum degree f . For the last property, let $C \subseteq E(H)$ be a cycle on $\leq 2k$ edges, and let (u, v) be the latest edge in C considered by the greedy algorithm. When we add (u, v) , we must include at least one edge in C in the associated fault set, since otherwise there is a $u \rightsquigarrow v$ path through C of length $\leq (2k - 1) \cdot w(u, v)$. Thus $C \cap F_e$ is nonempty. \square

The focus of the proof now shifts. Instead of controlling the number of edges in the specific output graph H from Algorithm GreedyFDSpanner, instead we will control the number of edges in *any* n -node graph that admits a f -FD $2k$ -blocking set. That is, our goal is now to prove the following lemma:

Lemma 6.5 (Blocking Set Size Bounds). *For any positive integers n, k, f , any n -node graph H with a f -FD $2k$ -blocking set \mathcal{B} satisfies*

$$|E(H)| \leq f^{1-1/k} n^{1+1/k} \cdot O(k)^k.$$

In the following arguments, H is the input graph from this lemma, which thus has a blocking set as described.

6.2 MUCk Paths and Dispersion Lemma

Our bound on $|E(H)|$ will be a counting argument over a special type of path that we will call a MUCk path, defined as follows:

Definition 6.6 (MUCk Paths). *A path π in H is called a MUCk path if it satisfies the following properties:*

- (Monotonic) The edges in π occur in strictly increasing order in the edge-ordering of H .
- (Unblocked) π does not contain two edges e, f with $e \in F_f$.
- (Chain-Unblocked) For each edge $(u, v) \in \pi$, there are at most Cfk edges e with the property that $(u, v) \in F_e$, and also e is incident on a node x that strictly follows (u, v) on π (in particular $x \notin \{u, v\}$). Here $C \geq 1$ is a parameter that we will set later.
- (k -Path) π has exactly k edges. (The monotonic property implies that these edges are all distinct.)

In order to help motivate this definition, let us recall the discussion in Section 2.2. Analogizing the Moore bounds, which limit the maximum possible number of edges in a high-girth graph, our proof will include two steps:

- A *dispersion lemma*, which shows that not too many MUCk paths can share endpoints (or else they violate the existence of the FD blocking set), and
- A *counting lemma*, which shows a lower bound on the number of MUCk paths in H , where the bound scales with the number of edges in H .

Our edge bound will follow by comparing the upper and lower bounds on the number of MUCk paths that respectively arise from the dispersion and counting lemmas, and rearranging terms.

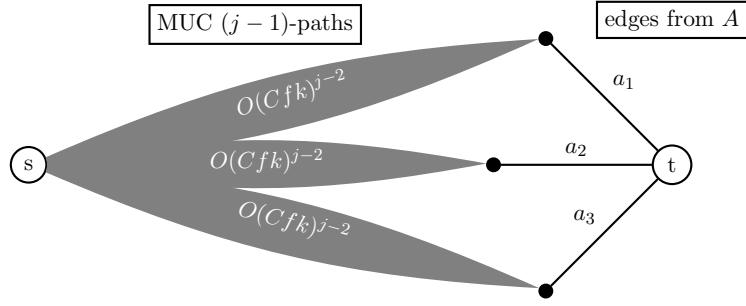
The main challenge in this proof method is to find a precise combination of properties that are specific enough to enable a dispersion lemma, but also general enough that many such paths are guaranteed to exist, so that the counting lemma applies. The MUCk definition is a cocktail of properties from prior work that does exactly that: monotonicity appears in [Bod23], a version of the chain-unblocked property is used in [BDN23], and the unblocked property has been used repeatedly, e.g., [BDR22]. Rather than saying more about the exact role played by each specific property, we will next show them in action by stating and proving our *dispersion lemma*. The properties can be best understood as outlining a maximally general class of paths over which the following dispersion lemma proof strategy applies. The lemma shows that the MUCk paths must be “dispersed” around the graph H , rather than having too many of them concentrated on the same pair of endpoints.

Lemma 6.7 (Dispersion Lemma). *For any two nodes s, t , H has at most $O(Cfk)^{k-1}$ MUCk paths with endpoints (s, t) .*

Proof. For the sake of an inductive proof, we will prove the following stronger claim: for each integer $1 \leq j \leq k$, H has at most $O(Cfk)^{j-1}$ MUC j -paths with endpoints (s, t) . The induction is on j , and the base case $j = 1$ holds since H has at most one 1-path (edge) between s and t .

For the inductive step, suppose $j \geq 2$, and let A be the set of edges that are used as the last edge of at least one MUC j -path with endpoints (s, t) (in particular, all edges in A are incident on t). If we have $|A| \leq O(Cfk)$, then the rest of the proof is a straightforward counting argument: every MUC j -path can be expressed by choosing an edge $a \in A$ and a MUC $j - 1$ path from s to the non- t endpoint of a , and so by applying the inductive hypothesis, the number of such paths is bounded by

$$|A| \cdot O(Cfk)^{j-2} \leq O(Cfk)^{j-1}.$$



So now it only remains to prove that $|A| \leq O(Cfk)$. Assume for contradiction that $|A| \geq 6Cfk$, let $A_{\text{late}} \subseteq A$ be the latest $4Cfk$ edges of A in the ordering, and let $A_{\text{early}} \subseteq A$ be the earliest $2Cfk$ edges of A in the ordering. For each edge $a \in A_{\text{late}}$, by the fault-degree constraint there are at most f edges in F_a incident to t , and therefore there are at most f edges in $F_a \cap A_{\text{early}}$. Thus (somewhat conservatively) there are at least Cfk edges in A_{early} that are *not* also in F_a . We can therefore count that there exists an edge $a' \in A_{\text{early}}$ with the property that there are more than Cfk edges $a \in A_{\text{late}}$ with $a' \notin F_a$.

Let us analyze this particular edge a' . Let π' be a MUC j -path with endpoints (s, t) that uses a' as its last edge. Let $a \in A_{\text{late}}$ be one of the edges described above with $a' \notin F_a$, and let π^a be a MUC j -path with endpoints (s, t) that uses a as its last edge. Notice that $\pi' \cup \pi^a$ contains a cycle C on $|C| \leq 2k$ edges, and by monotonicity, the latest edge in C is a . Thus, by definition of FD blocking sets, there exists an edge $e \in C \cap F_a$.

Let us now consider the placement of e in the paths $\pi' \cup \pi^a$. We know that $e \notin \pi^a$, since π^a is unblocked. We know that e is not the last edge of π' , since the last edge of π' is a' and we have been careful to consider an edge $a' \notin F_a$. Thus the only remaining possibility is that e is one of the first $j - 1$ edges in π' . That means that, for the edge $e \in \pi'$, the edge a counts towards the chain-blocked property: a is incident on a node (t) that strictly follows e along π' , and we have $e \in F_a$. Since we require that π' is chain-unblocked, there can be at most Cfk edges a with this property. However, we have shown that our analysis applies to at least $> Cfk$ different edges $a \in A_{\text{late}}$. This completes the contradiction and the proof. \square

6.3 Counting Lemma

Our next step is to prove lower bounds on the number of MU Ck paths in H . Our strategy is to first count the number of monotonic k -paths, and then separately count the number of blocked and chain-blocked k -paths, and argue that most monotonic k -paths must in fact be MU Ck paths (in the relevant parameter regime).

We start with the monotone counting lemma. Only the “full” monotone counting lemma in the following sequence is ultimately used; the previous two lemmas are used to help bootstrap the proof.

Lemma 6.8 (Weak Monotone Counting Lemma). *If $|E(H)| \geq kn/2$, then H contains a monotonic k -path.*

Proof. This follows from a folklore theorem in graph theory sometimes called the “hiker lemma” (see [Bod23] for context). We generate n different monotonic paths by the following process. Start by placing a hiker at each node of H . Then, consider the edges in H in increasing order. Each time we consider an edge (u, v) , we ask the two hikers who currently stand at u, v to walk across the edge, switching places with each other.

Once all edges have been considered, our n hikers have traversed $2|E(H)| \geq kn$ edges in total. Additionally, the path walked by each individual hiker is monotonic. Thus, there exists a hiker who walked a monotonic path of length $\geq k$. \square

Lemma 6.9 (Medium Monotone Counting Lemma). *If $|E(H)| \geq kn$, then H contains at least $kn/2$ monotonic k -paths.*

Proof. By the weak counting lemma, we would need to delete at least $kn/2$ edges from H in order to destroy all monotonic k -paths. This implies that H has at least $kn/2$ monotonic k -paths to begin with, since otherwise we could easily delete any one edge per path to destroy them all. \square

Lemma 6.10 (Full Monotone Counting Lemma). *Let d be the average degree in H . If $d \geq 2k$, then H contains at least $kn \cdot \Omega(d/k)^k$ monotonic k -paths.*

Proof. Let H' be a random edge-subgraph of H on exactly kn edges. Let x be the number of monotonic k -paths in H , and let x' be the number of monotonic k -paths that survive in H' . On one hand, by the medium counting lemma, we have $x' \geq kn/2$ (deterministically). On the other hand, for any monotonic k -path π in H , the probability that π survives in H' is:

$$\begin{aligned} & \left(\frac{kn}{|E(H)|} \right) \left(\frac{kn-1}{|E(H)|-1} \right) \cdots \left(\frac{kn-(k-1)}{|E(H)|-(k-1)} \right) \\ & \leq \left(\frac{kn}{|E(H)|} \right)^k \\ & = \Theta \left(\frac{k}{d} \right)^k. \end{aligned}$$

Thus we have

$$\frac{kn}{2} \leq \mathbb{E}[x'] \leq x \cdot \Theta \left(\frac{k}{d} \right)^k.$$

Rearranging terms, we get our desired inequality of

$$x \geq kn \cdot \Theta \left(\frac{d}{k} \right)^k.$$

\square

Finally, we count the number of blocked or chain-blocked k -paths in H , with the plan to argue that these are much fewer than the number of monotonic k -paths. The following arguments require a stronger assumption that the *max* degree is bounded; in the next part we will show how this assumption can be justified.

Lemma 6.11 (Chain-Blocked Path Upper Bound). *Let d be a parameter and suppose that the maximum node degree in H is $O(d)$. Then H has at most $kn \cdot O(d)^k/C$ k -paths that are edge-simple and chain-blocked.*

Proof. We can overcount our chain-blocked paths as follows:

- First, choose its ending node v_k (n choices).
- Then, choose an index $1 \leq i \leq k$ (k choices) such that the i^{th} edge causes π to be chain-blocked. That is, we commit to selecting the i^{th} edge $e_i \in \pi$ in such a way that there exist $> Cfk$ edges e incident to nodes in π that strictly follow e_i , with $e_i \in F_e$.

- Now build π iteratively. If currently we have determined a suffix (v_j, \dots, v_k) of π , then we choose the previous node v_{j-1} as follows.
 - If $j-1 \neq i$, then we choose any of the $O(d)$ neighbors of v_j besides v_{j+1} to serve as v_{j-1} .
 - If $j-1 = i$, then our choice is restricted, since we need this edge to cause π to be chain-blocked. Let us count the number of neighbors of v_j that we may choose in this step. There are $O(dk)$ edges e incident to any of the previously-selected nodes (v_{j+1}, \dots, v_k) . By the fault-degree constraint, each of these edges has at most f edges incident to v_j in its set F_e . We thus count that there are $O(d/C)$ edges incident to v_j that are in $> Cfk$ such sets F_e . We must choose one such edge in order for this edge to be chain-blocked, and so there are $O(d/C)$ choices of edge at this step.

Putting it together, we have $n \cdot k \cdot O(d)^{k-1} \cdot O(d/C) = kn \cdot O(d)^k/C$ edge-simple chain-blocked k -paths. \square

Lemma 6.12 (Blocked Path Upper Bound). *Let d be a parameter and suppose that the maximum node degree in H is $O(d)$. Then H has at most $k^2n \cdot O(d)^{k-1} \cdot f$ k -paths that are edge-simple and blocked.*

Proof. Recall that a k -path π is blocked if it contains edges $e, e' \in \pi$ with $e' \in F_e$. We may overcount our blocked k -paths as follows:

- Choose distinct indices $1 \leq i, i' \leq k$ where e, e' will occur ($O(k^2)$ choices). In the following we will assume w.l.o.g. that $i < i'$; the case where $i > i'$ is symmetric, but we build π from its last node towards the beginning instead of from its first node towards the end.
- Choose the first node v_1 of π (n choices).
- Now we build π iteratively. If currently we have determined a prefix (v_1, \dots, v_j) , then we choose the following node v_{j+1} as follows.
 - If $j+1 \neq i'$, then we choose any of the $O(d)$ neighbors of v_j besides v_{j-1} to serve as v_{j+1} .
 - If $j+1 = i'$, then our choice is restricted, since we need this edge to lie in F_e (since $i' > i$, note that we have already determined the edge e). By the fault-degree constraint, there are at most f edges incident on v_j in F_e , and hence there are at most f choices of edges at this step.

Putting it together, we have $O(k^2) \cdot n \cdot O(d)^{k-1} \cdot f$ blocked k -paths. \square

6.4 Proof Wrapup

We will first need to apply a standard ‘‘cleaning’’ step to H , in order to reduce to the case where the max degree and average degree are roughly equal, so that all of the above counting lemmas may be applied simultaneously. This step is standard and has appeared in some form in a lot of prior work on spanners, but we will outline the proof anyways.

Lemma 6.13 (Degree-Cleaning Lemma). *Let H be a graph with n nodes and a f -FD $2k$ -blocking set with average degree d . Then there exists a graph H' with $\Theta(n)$ nodes and a f -FD $2k$ -blocking set in which all nodes have degree $\Theta(d)$.*

Proof. Let H' be a graph obtained from H by the following process. Fix d as the initial average degree in H . While there exists a node of degree $< d/4$, delete that node and all its incident edges from H . Then, while there exists a node of degree $> d$, split that node into two copies, with its edges divided equitably between the copies. Define H' as the final graph at the end of this process.

It is immediate from the construction that all surviving nodes in H' have degree $\Theta(d)$; in particular this means that the average degree changes by at most a constant factor from H to H' . Additionally, notice that we delete at most $< nd/4 \leq |E(H)|/2$ edges from H to H' , and so the number of edges changes by at most a constant factor from H to H' . Since the average and total degrees both change by a constant factor, it follows that the number of nodes also changes by only a constant factor from H to H' . Thus H' has $\Theta(n)$ nodes, as desired.

The last detail is to argue that H' still has a f -FD $2k$ -blocking set. This follows from the fact that we create H' from H only by the operations of deleting and splitting nodes. These operations cannot create new $2k$ -cycles, and they can only reduce the fault degree of the blocking set. Thus, the remaining part of the f -FD $2k$ -blocking set for H is still a valid f -FD $2k$ -blocking set for H' . \square

Since our previous lemmas apply to *any* graph with a f -FD $2k$ -blocking set, we can apply them to H' . On one hand, by the dispersion lemma, H' has $n^2 \cdot O(Cfk)^{k-1}$ total MUck paths. On the other hand, by Lemmas 6.10, 6.11, and 6.12, the number of MUck paths in H' is at least

$$\underbrace{\left(kn \cdot \Omega\left(\frac{d}{k}\right)^k \right)}_{\text{l.b. on monotone } k\text{-paths}} - \underbrace{\left(\frac{kn \cdot O(d)^k}{C} \right)}_{\text{u.b. on chain-blocked } k\text{-paths}} - \underbrace{k^2nf \cdot O(d)^{k-1}}_{\text{u.b. on blocked } k\text{-paths}} .$$

If we set $C \geq \Omega(k)^k$ with a large enough implicit constant, then the monotone k -paths dominate the chain-blocked k -paths. If we have $d \geq \Omega(k)^k \cdot f$ with a large enough implicit constant, then the monotone k -paths dominate the blocked k -paths. So, under both these assumptions, we have

$$kn \cdot \Omega\left(\frac{d}{k}\right)^k \leq (\# \text{ MUck paths}) \leq n^2 \cdot O(Cfk)^{k-1}$$

and so

$$\frac{d}{k} \leq n^{1/k} \cdot O(Cfk)^{(k-1)/k}$$

and so

$$|E(H)| = \Theta(nd) \leq n^{1+1/k} \cdot f^{1-1/k} \cdot k^{O(k)}$$

where the last step uses our setting of C . Finally, in the case where $d \leq O(k)^k \cdot f$, then we immediately have $|E(H)| = \Theta(nd) \leq nf \cdot O(k)^k$, and so the bound still holds. In either case, Theorem 6.1 is proved.

6.5 The Min Max Length-Bounded Cut Problem

Our next goal is to develop a tweak on the FD-greedy algorithm that runs in polynomial time. Before explaining this, we will need to develop a subroutine to be used in the algorithm. We will discuss the following problem:

MIN MAX LENGTH-BOUNDED CUT (LBC): Given an n -node graph $G = (V, E)$, an integer k , and nodes u, v , find the least integer f such that there exists an edge set F of degree $\deg(F) \leq f$ with $\text{dist}_{G \setminus F}(u, v) > k$.

For motivation, notice that for the FD-greedy algorithm in the setting where the input graph G is unweighted, the test to keep or discard an edge (u, v) is precisely MIN MAX LBC. This problem is likely hard: it is a tweak on the NP-hard problem LENGTH-BOUNDED CUT (LBC) [BEH⁺06], in which the goal is to minimize $|F|$ subject to the same constraint. So, we will develop an approximation algorithm. This high-level approach to speeding up the greedy algorithm was first used by Dinitz and Robelle [DR20], although the min-max degree objective introduces some new challenges that mean our approximation algorithm works quite differently from theirs.

Theorem 6.14. *There is an $O(k \log n)$ approximation algorithm for MIN MAX LBC that runs in polynomial time. Specifically, on an instance of MIN MAX LBC with solution f^* , the algorithm returns a value \hat{f} satisfying $f^* \leq \hat{f} \leq Bf^* \cdot k \log n$, where $B > 0$ is some absolute constant.*

We will use an LP-rounding algorithm. The following LP encodes a fractional relaxation of MIN MAX LBC:

LP for Fractional Min Max LBC

Choose $c \in \mathbb{R}^{E(G)}$, $f \in \mathbb{R}$

To minimize f

Subject to:

- For all nodes v , $\sum_{e \text{ incident to } v} c_e \leq f$ // degree constraints
- For all simple $u \rightsquigarrow v$ paths π of length $|\pi| \leq k$, $\sum_{e \in \pi} c_e \geq 1$ // path cut constraints
- For all edges e , $0 \leq c_e \leq 1$. // cut value constraints

This LP could have exponentially many constraints – roughly n^k – since it has one constraint per short $u \rightsquigarrow v$ path. However, we next point out that it has a separation oracle, and thus can be efficiently solved.

Lemma 6.15. *There is a polynomial-time separation oracle for the above LP.*

Proof. Given a setting of c, f , it is trivial to check the degree constraints and the path cut constraints in $\text{poly } n$ time. To check the path constraints, let G' be the weighted graph with the same edge set as G and with the variables c_e interpreted as edge weights. Our goal is to check whether or not there exists a $u \rightsquigarrow v$ path π in G' that uses $|\pi| \leq k$ hops, and which has total weight $w(\pi) < 1$.

To check this, we design an auxiliary graph G'' by creating $k+1$ independent copies G'_0, G'_1, \dots, G'_k of the graph G' . We interpret these as $k+1$ layers, and for a node x , we write x_i to mean the copy of x in G'_i . For each edge $(x, y) \in E(G')$, for each $0 \leq i \leq k-1$ we add directed edges (x_i, y_{i+1}) and (y_i, x_{i+1}) to $E(G'')$, both with the same weight as (x, y) . Additionally, for each node x we add all directed edges of the form (x_i, x_{i+1}) with weight 0.

We now claim that a given choice of c, f satisfies the path cut constraints iff $\text{dist}_{G''}(u_0, v_k) \geq 1$, which is easy to test in polynomial time. Indeed, notice that any $u_0 \rightsquigarrow v_k$ path π in G'' corresponds to a $u \rightsquigarrow v$ path in G on $\leq k$ hops, and that its length $w(\pi)$ is the sum of cut variables c_e along π . \square

Now let c, f_{LP} be an optimal solution to this LP. Notice that the IP in which each $c_e \in \{0, 1\}$ encodes MIN MAX LBC exactly. Hence, letting f^* be the true solution to the (non-relaxed) LBC problem, we have

$$f_{LP} \leq f^*.$$

We will next use the entries of c to randomly generate an edge set F (this step is equivalent to rounding the entries of c to $\{0, 1\}$). Specifically: each edge e is included in F independently with probability $\min\{A \cdot c_e \cdot k \log n, 1\}$, where $A > 0$ is a large enough absolute constant that will control the probability of correctness. In the following analysis, for each edge e we will let X_e be the Bernoulli random variable indicating whether or not $e \in F$.

We can bound the degree of F as follows:

Lemma 6.16. *With high probability, $\deg(F) \leq O(f^* \cdot k \log n)$.*

Proof. Fix a node v . By the degree constraints, we have

$$\sum_{e \text{ incident to } v} c_e \leq f_{LP} \leq f^*.$$

It follows that

$$\mathbb{E}[\deg_F(v)] \leq \sum_{e \text{ incident to } v} A \cdot c_e \cdot k \log n \leq Af^* \cdot k \log n.$$

By the Chernoff bounds, we then have

$$\Pr[\deg_F(v) \geq 2Af^* \cdot k \log n] \leq \exp(-\Theta(Af^*k \log n)) \leq \frac{1}{n^{\Theta(Af^*k)}}.$$

By a union bound over the n nodes in the graph, we thus have

$$\Pr[\deg(F) \leq O_A(f^* \cdot k \log n)] \geq 1 - \frac{1}{n^{\Theta(Af^*k)-1}}.$$

By choice of large enough constant AA , this means $\deg(F) \leq O_A(f^* \cdot k \log n)$ with high probability, completing the proof. \square

Then we prove correctness:

Lemma 6.17. *With high probability, $\text{dist}_{G \setminus F}(u, v) > k$.*

Proof. Let π be a $u \rightsquigarrow v$ path in G with $|\pi| \leq k$ edges. We will measure the probability that F contains at least one edge in π . First, by the path cut constraints, we have

$$\sum_{e \in \pi} c_e \geq 1.$$

If any edge $e \in \pi$ has fractional cut value c_e satisfying $A \cdot c_e \cdot k \log n \geq 1$, then this edge is included in F deterministically. Otherwise, we have

$$\mathbb{E} \left[\sum_{e \in \pi} X_e \right] = \sum_{e \in \pi} A \cdot c_e \cdot k \log n \geq A \cdot k \log n.$$

By Hoeffding's Inequality, we can bound

$$\Pr[F \text{ does not contain an edge in } \pi] = \Pr \left[\sum_{e \in \pi} X_e < 1 \right] \leq \exp \left(-\Theta \left(\frac{A^2 k^2 \log^2 n}{k} \right) \right) \leq \frac{1}{n^{Ak}}$$

where the last step follows by choice of large enough constant A . There are at most n^k possible $u \rightsquigarrow v$ paths in G with $\leq k$ edges. By a union bound, the probability that there is *any* such path that does not intersect F is at most $\frac{1}{n^{(A-1)k}}$. So, with high probability, no $u \rightsquigarrow v$ path on $\leq k$ edges survives in $G \setminus F$. \square

6.6 Improvement to Polynomial Time

We now describe how to tweak the FD-greedy algorithm to run in polynomial time, at mild cost in spanner size. This part very directly follows the approach of Dinitz and Robelle [DR20]. We simply run the previous FD-greedy algorithm, but we use our approximation algorithm for MIN MAX LBC to test edges:

Algorithm Approximate GreedyFDSpanner

Input: A graph G and a positive integer stretch parameter k .

Output: An f -FD $(2k - 1)$ -spanner $H \subseteq G$ with at most $f^{1-1/k} \cdot n^{1+1/k} \cdot O(k)^k$ edges.

- Let $H \leftarrow (V, \emptyset, w)$ be the initially-empty spanner
- For each edge $(u, v) \in E$ in order of nondecreasing $w(u, v)$
 - If $\text{APPROXMINMAXLBC}(u, v, 2k - 1) \leq Bfk \log n$ // from Theorem 6.14
 - * Add (u, v) to H
- Return H

This algorithm now runs in polynomial time, but due to the changes we need to re-establish its correctness and size bound. This mostly follows by rehashing our early arguments used for the (non-approximate) GreedyFDSpanner.

Lemma 6.18. *The output spanner H of the Approximate GreedyFDSpanner algorithm is a f -FD $(2k - 1)$ -spanner of G .*

Proof. This is nearly the same argument as Lemma 6.2, with a couple minor tweaks. As before, let F be an arbitrary set of faulty edges with $\deg(F) \leq f$, and it suffices to verify the spanner stretch inequality for an arbitrary edge $(u, v) \in G \setminus F$. If (u, v) is added to the spanner H by the greedy algorithm, then we immediately have $\text{dist}_H(u, v) = w(u, v)$. Otherwise, if (u, v) is discarded from H by the algorithm, then it must be that $\text{APPROXMINMAXLBC}(u, v, 2k - 1) > Bfk \log n$. By Theorem 6.14, this means that at the time (u, v) is considered, there is no fault set of degree $\deg(F) \leq f$ that intersects all $u \rightsquigarrow v$ paths on $\leq 2k - 1$ edges. Hence at least one such path π survives in $G \setminus F$. Since the greedy algorithm considers edges in increasing order of weight, every edge in π has weight $\leq w(u, v)$. We therefore have

$$\text{dist}_{G \setminus F}(u, v) \leq |\pi| \cdot w(u, v) \leq (2k - 1) \cdot w(u, v)$$

and so the spanner inequality holds. \square

Lemma 6.19. *The output spanner H of the Approximate GreedyFDSpanner algorithm has size*

$$|E(H)| \leq \tilde{O} \left(f^{1-1/k} n^{1+1/k} \right) \cdot O(k)^k.$$

Proof. From Theorem 6.14, the algorithm APPROXMINMAXLBC overestimates the true value of MINMAXLBC on a given input instance. Thus, each time an edge (u, v) is added to the spanner H , there exists a length-bounded cut F of degree $\deg(F) \leq \tilde{O}_B(fk)$. Exactly as in Lemma 6.4, the output spanner H therefore has a $\tilde{O}(fk)$ -FD blocking set, and the edge bound then follows by plugging this into Lemma 6.5. \square

7 Existential Lower Bounds for FD Spanners and Connectivity Certificates

7.1 FD Spanner Lower Bounds

We first show that the upper bound in Theorem 6.1 is optimal, up to its dependence on k :

Theorem 7.1. *Assuming the girth conjecture [Erd64], for all k, f , there are n -node undirected unweighted graphs in which every f -FD $(2k - 1)$ -spanner has at least $\Omega(f^{1-1/k} n^{1+1/k})$ edges.*

The proof of this theorem is essentially from [BDPW18], where the following construction is used for a lower bound against vertex fault tolerant spanners.

The Construction. We will analyze the following graphs:

- The girth conjecture [Erd64] states that for all positive integers k , there exist n -node graphs on $\Theta(n^{1+1/k})$ edges with girth $> 2k$. We begin by letting $\Gamma = (V, E)$ be one such graph.
- Replace each node $v \in V$ with f copies, v_1, \dots, v_f . Each edge $(u, v) \in E$ is replaced by f^2 edges, arranged in a biclique on the vertex sets $\{u_1, \dots, u_f\} \times \{v_1, \dots, v_f\}$. Let G be the final graph after these replacements.

The graph G has $N := nf$ nodes, and its number of edges is

$$|E(G)| = \Theta\left(f^2 \cdot n^{1+1/k}\right) = \Theta\left(f^{1-1/k} \cdot N^{1+1/k}\right).$$

Thus, it suffices to argue that the only f -FD $(2k - 1)$ spanner of G is G itself.

Analysis. Let H be a f -FD $(2k - 1)$ spanner of G , and let $(u_i, v_j) \in E(H)$ (so our goal is to argue that $(u_i, v_j) \in E(H)$ as well). Let F be the fault set that contains every edge connecting any copy of u to any copy of v , except for the edge (u_i, v_j) . That is:

$$F := \{(u_a, v_b) \mid a \neq i \text{ or } b \neq j\}.$$

Notice that the degree of F on u_i, v_j is $f - 1$, and for each copy of u, v besides u_a, v_b , the degree of F on this node is f . Thus F is a f -FD fault set, which means that by definition of f -FD spanners, we require

$$\text{dist}_{H \setminus F}(u_i, v_j) \leq (2k - 1) \cdot \text{dist}_{G \setminus F}(u_i, v_j) = 2k - 1.$$

This inequality clearly holds if $(u_i, v_j) \in E(H)$ (in which case we have $\text{dist}_{H \setminus F}(u_i, v_j) = 1$). On the other hand, suppose for contradiction that $(u_i, v_j) \notin E(H)$. In particular this means there is no edge connecting any copy of u to any copy of v remaining in $H \setminus F$. Thus, any $u_i \rightsquigarrow v_j$ path in $H \setminus F$ corresponds to a $u \rightsquigarrow v$ path π in Γ that does not include the edge (u, v) . Since Γ has girth $> 2k$, and $\pi \cup \{(u, v)\}$ forms a cycle in Γ , there must be $> 2k - 1$ edges in π . So any $u_i \rightsquigarrow v_j$ path in $H \setminus F$ has length $> 2k - 1$, violating the above spanner inequality, and completing the contradiction.

We have now proved that every edge $(u_i, v_j) \in E(G)$ must lie in $E(H)$ as well, which means that $G = H$, i.e., G is the only f -FD $(2k - 1)$ -spanner of itself, which completes the proof.

7.2 FD Connectivity Certificate Lower Bounds

The naive lower bound for a f -FD connectivity certificate is $\Omega(nf)$ edges, from any f -regular graph. Here we point out that there is a minor improvement to this lower bound available:

Theorem 7.2. *For all f , there are n -node undirected unweighted graphs in which every f -FD connectivity certificate has $\Omega(fn \cdot \frac{\log n}{\log f})$ edges.*

The Construction. Let d be a parameter of the construction. Our lower bound graph G is defined as follows:

- The nodes of G are $[f]^d$, i.e., d -tuples of integers between 1 and f .
- Two nodes of G are joined by an edge iff their d -tuples are equal on all but one coordinate.

This graph has $n := f^d$ nodes, and each node has degree

$$(f-1)d = (f-1) \cdot \frac{\log n}{\log f}.$$

Thus it suffices to argue that G is the only f -FD connectivity certificate of itself.

The Analysis. Let $(u, v) \in E(G)$, where the d -tuples associated to nodes u, v agree on all but the i^{th} coordinate. Let H be a f -FD connectivity certificate of G , and suppose for contradiction that $(u, v) \notin E(H)$. Let F be the fault set that contains every edge in G that connects two nodes whose d -tuples agree on all but the i^{th} coordinate, except that we also omit (u, v) itself from F . Notice that the degree of F on any node x is at most $f-1$, since there are only $f-1$ other nodes whose d -tuples agree with x on all but the i^{th} coordinate.

We have that u, v are connected in $G \setminus F$, since $(u, v) \in E(G) \setminus F$. However, we claim that u, v are disconnected in $H \setminus F$. This follows by noticing that all nodes reachable from u in $H \setminus F$ must have the same i^{th} coordinate as u , but v has a different i^{th} coordinate from u . This completes the contradiction.

We have shown that, for all edges $(u, v) \in E(G)$, we have $(u, v) \in E(H)$. Thus $H = G$, and so G is the only f -FD connectivity certificate of itself, and so the theorem follows.

References

- [ABCP98] Baruch Awerbuch, Bonnie Berger, Lenore Cowen, and David Peleg. Near-linear time construction of sparse neighborhood covers. *SIAM J. Comput.*, 28(1):263–277, 1998.
- [ABS⁺20] Reyan Ahmed, Greg Bodwin, Faryad Darabi Sahneh, Keaton Hamm, Mohammad Javad Latifi Jebelli, Stephen Kobourov, and Richard Spence. Graph spanners: A tutorial review. *Computer Science Review*, 37:100253, 2020.
- [ADD⁺93] Ingo Althöfer, Gautam Das, David P. Dobkin, Deborah Joseph, and José Soares. On sparse spanners of weighted graphs. *Discrete & Computational Geometry*, 9:81–100, 1993.
- [BBC04] Nikhil Bansal, Avrim Blum, and Shuchi Chawla. Correlation clustering. *Machine learning*, 56:89–113, 2004.
- [BDN23] Greg Bodwin, Michael Dinitz, and Yasamin Nazari. Epic Fail: Emulators Can Tolerate Polynomially Many Edge Faults for Free. In Yael Tauman Kalai, editor, *14th Innovations in Theoretical Computer Science Conference (ITCS 2023)*, volume 251 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 20:1–20:22, 2023.
- [BDPW18] Greg Bodwin, Michael Dinitz, Merav Parter, and Virginia Vassilevska Williams. Optimal vertex fault tolerant spanners (for fixed stretch). In Artur Czumaj, editor, *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2018, New Orleans, LA, USA, January 7-10, 2018*, pages 1884–1900. SIAM, 2018.

[BDR21] Greg Bodwin, Michael Dinitz, and Caleb Robelle. Optimal vertex fault-tolerant spanners in polynomial time. In DánIEL Marx, editor, *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms, SODA 2021, Virtual Conference, January 10 - 13, 2021*, pages 2924–2938. SIAM, 2021.

[BDR22] Greg Bodwin, Michael Dinitz, and Caleb Robelle. Partially optimal edge fault-tolerant spanners. In Joseph (Seffi) Naor and Niv Buchbinder, editors, *Proceedings of the 2022 ACM-SIAM Symposium on Discrete Algorithms, SODA 2022, Virtual Conference / Alexandria, VA, USA, January 9 - 12, 2022*, pages 3272–3286. SIAM, 2022.

[BEH⁺06] Georg Baier, Thomas Erlebach, Alexander Hall, Ekkehard Köhler, Heiko Schilling, and Martin Skutella. Length-bounded cuts and flows. In Michele Bugliesi, Bart Preneel, Vladimiro Sassone, and Ingo Wegener, editors, *Automata, Languages and Programming*, pages 679–690, Berlin, Heidelberg, 2006. Springer Berlin Heidelberg.

[Bod23] Greg Bodwin. An alternate proof of near-optimal light spanners. *arXiv preprint arXiv:2305.18647*, 2023.

[BP19] Greg Bodwin and Shyamal Patel. A trivial yet optimal solution to vertex fault tolerant spanners. In *Proceedings of the 2019 ACM Symposium on Principles of Distributed Computing*, PODC ’19, page 541–543, New York, NY, USA, 2019. Association for Computing Machinery.

[BS07] Surender Baswana and Sandeep Sen. A simple and linear time randomized algorithm for computing sparse spanners in weighted graphs. *Random Struct. Algorithms*, 30(4):532–563, 2007.

[CGS17] Moses Charikar, Neha Gupta, and Roy Schwartz. Local guarantees in graph cuts and clustering. In *International Conference on Integer Programming and Combinatorial Optimization*, pages 136–147. Springer, 2017.

[CLN15] T-H Hubert Chan, Mingfei Li, and Li Ning. Sparse fault-tolerant spanners for doubling metrics with bounded hop-diameter or degree. *Algorithmica*, 71:53–65, 2015.

[CLNS15] T-H Hubert Chan, Mingfei Li, Li Ning, and Shay Solomon. New doubling spanners: Better and simpler. *SIAM Journal on Computing*, 44(1):37–53, 2015.

[CLPR10] Shiri Chechik, Michael Langberg, David Peleg, and Liam Roditty. Fault tolerant spanners for general graphs. *SIAM J. Comput.*, 39(7):3403–3423, 2010.

[CZ04] Artur Czumaj and Hairong Zhao. Fault-tolerant geometric spanners. *Discrete & Computational Geometry*, 32(2):207–230, 2004.

[DR20] Michael Dinitz and Caleb Robelle. Efficient and simple algorithms for fault-tolerant spanners. In *Proceedings of the 2020 ACM Symposium on Principles of Distributed Computing*, PODC ’20, 2020.

[Erd64] Paul Erdős. Extremal problems in graph theory. In *IN THEORY OF GRAPHS AND ITS APPLICATIONS, PROC. SYMPOS. SMOLENICE*. Citeseer, 1964.

[HHG22] Bernhard Haeupler, Jonas Hübotter, and Mohsen Ghaffari. A cut-matching game for constant-hop expanders. *CoRR*, abs/2211.11726, 2022.

[HHS23] Bernhard Haeupler, D. Ellis Hershkowitz, and Thatchaphol Saranurak. Maximum length-constrained flows and disjoint paths: Distributed, deterministic, and fast. In Barna Saha and Rocco A. Servedio, editors, *Proceedings of the 55th Annual ACM Symposium on Theory of Computing, STOC 2023, Orlando, FL, USA, June 20-23, 2023*, pages 1371–1383. ACM, 2023.

[HHT23a] Bernhard Haeupler, D. Ellis Hershkowitz, and Zihan Tan. Length-constrained expander decomposition: Deeper theory, simpler algorithm. *Under Submission*, 2023.

[HHT23b] Bernhard Haeupler, D. Ellis Hershkowitz, and Zihan Tan. Parallel greedy spanners. *CoRR*, abs/2304.08892, 2023.

[HRG22] Bernhard Haeupler, Harald Räcke, and Mohsen Ghaffari. Hop-constrained expander decompositions, oblivious routing, and distributed universal optimality. In Stefano Leonardi and Anupam Gupta, editors, *STOC '22: 54th Annual ACM SIGACT Symposium on Theory of Computing, Rome, Italy, June 20 - 24, 2022*, pages 1325–1338. ACM, 2022.

[JKMM21] Jafar Jafarov, Sanchit Kalhan, Konstantin Makarychev, and Yury Makarychev. Local correlation clustering with asymmetric classification errors. In *International Conference on Machine Learning*, pages 4677–4686. PMLR, 2021.

[KMZ19] Sanchit Kalhan, Konstantin Makarychev, and Timothy Zhou. Correlation clustering with local objectives. *Advances in Neural Information Processing Systems*, 32, 2019.

[LNS98] Christos Levcopoulos, Giri Narasimhan, and Michiel Smid. Efficient algorithms for constructing fault-tolerant geometric spanners. In *Proceedings of the thirtieth annual ACM symposium on Theory of computing*, pages 186–195. ACM, 1998.

[LR99] Frank Thomson Leighton and Satish Rao. Multicommodity max-flow min-cut theorems and their use in designing approximation algorithms. *J. ACM*, 46(6):787–832, 1999.

[LS21] Jason Li and Thatchaphol Saranurak. Deterministic weighted expander decomposition in almost-linear time. *CoRR*, abs/2106.01567, 2021.

[LST23] Hung Le, Shay Solomon, and Cuong Than. Optimal fault-tolerant spanners in euclidean and doubling metrics: Breaking the $\omega(\log n)$ lightness barrier. *arXiv preprint arXiv:2306.11226*, 2023.

[Luk99] Tamas Lukovszki. New results on fault tolerant geometric spanners. *Algorithms and Data Structures*, pages 774–774, 1999.

[NI92] Hiroshi Nagamochi and Toshihide Ibaraki. A linear-time algorithm for finding a sparse k-connected spanning subgraph of a k-connected graph. *Algorithmica*, 7(5&6):583–596, 1992.

[NS07] Giri Narasimhan and Michiel Smid. *Geometric Spanner Networks*. Cambridge University Press, 2007.

[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.

[PS89] David Peleg and Alejandro A. Schäffer. Graph spanners. *Journal of Graph Theory*, 13(1):99–116, 1989.

[Sol14] Shay Solomon. From hierarchical partitions to hierarchical covers: Optimal fault-tolerant spanners for doubling metrics. In *Proceedings of the forty-sixth annual ACM symposium on Theory of computing*, pages 363–372, 2014.

[Wir10] Anthony Wirth. Correlation clustering., 2010.

A Missing Proofs

Proof of Lemma 4.1. The solution for the routing instance (F, d) provides each $(u, v) \in F$ with a multi-set of d many u - v paths $\mathcal{P}(u, v)$. Since $c < d$, it holds that $\mathcal{P}(u, v)$ contains at least one u - v path P that does not contain e , for every $e = (u, v) \in F$. We use the collection of multi-paths $\mathcal{P} = \{\mathcal{P}(u, v)\}_{(u, v) \in F}$ to compute a directed multi-graph D which captures the dependencies between the edges in F . Specifically, $V(D) = F$ and the multiplicity of a directed edge $(e = (u, v), e')$ equals to the number of paths in $\mathcal{P}(u, v)$ that go through e' . In the final multi-graph D , we include only edges with positive multiplicity. We refer to the digraph D as the dependency graph of the output multicommodity solution. Note that the multiplicity of an edge is bounded by $\min\{c, d\} = c$. We start by observing the following:

Observation A.1. *The output solution $\mathcal{P} = \{\mathcal{P}(u, v)\}_{(u, v) \in F}$ with congestion c can be transformed into an alternative solution $\mathcal{P}_0 = \{\mathcal{P}_0(u, v)\}_{(u, v) \in F}$ with congestion $c(\phi)$ whose dependency (multi)graph D_0 is a DAG.*

Proof. The procedure eliminates directed cycles in an iterative manner. Consider some directed cycle $(e_1, e_2) \circ (e_2, e_3) \circ \dots \circ (e_{\ell-1}, e_\ell) \circ (e_\ell, e_1)$ in D . Let $e_i = (u_i, v_i)$. Each edge (e_i, e_{i+1}) is identified with a unique path $P_i \in \mathcal{P}(u_i, v_i)$ such that $e_{i+1} \in P_i$. We then modify the routing solution as follows: For every $i \in \{1, \dots, \ell\}$, modify $\mathcal{P}(u_i, v_i)$ by omitting P_i and adding $\{e_i\}$. It is easy to see that the congestion of each edge is preserved and moreover, all edges of the directed cycle are eliminated in the dependency graph of the resulting modified solution. This step can then be repeated until all directed cycles are removed. \square

We next claim the following lemma which shows that the routing solution $\mathcal{P}_0 = \{\mathcal{P}_0(u, v)\}_{(u, v) \in F}$ can be further modified into $\mathcal{P}^* = \{\mathcal{P}^*(u, v)\}_{(u, v) \in F}$ with an *empty* dependency graph, while maintaining the same congestion bound of $c(\phi)$.

Lemma A.2. *There is a routing solution $\mathcal{P}^* = \{\mathcal{P}^*(u, v)\}_{(u, v) \in F}$ with congestion at most c and such that its corresponding dependency graph is empty.*

Proof. We employ an iterative procedure where in each iteration $i \geq 1$, given is a routing solution $\mathcal{P}_{i-1} = \{\mathcal{P}_{i-1}(u, v)\}_{(u, v) \in F}$ that corresponds to a DAG D_{i-1} . By Obs. A.1, this holds for $i = 1$. The output of the i^{th} iteration is a routing solution $\mathcal{P}_i = \{\mathcal{P}_i(u, v)\}_{(u, v) \in F}$ with congestion of c and a dependency graph $D_i \subset D_{i-1}$. In particular, in each iteration i we eliminate one copy of some edge in D_{i-1} . Hence, within a finite number of steps ℓ , we get a routing solution $\mathcal{P}^* = \mathcal{P}_\ell$ with congestion $c(\phi)$ and an empty dependency graph D_ℓ .

We now describe the i^{th} iteration. Let $e' = (u', v')$ be some sink vertex in D_{i-1} and let $e = (u, v)$ be such that $(e, e') \in D_{i-1}$. Our goal is to output $\{\mathcal{P}_i(u, v)\}_{(u, v) \in F}$ with congestion $c(\phi)$ and $D_i = D_{i-1} \setminus \{(e, e')\}$. I.e., we will reduce the multiplicity of the edge (e, e') by exactly one in D_i . Let $P \in \mathcal{P}_{i-1}(u, v)$ be such that $e' \in P$. Since $(e, e') \in D_{i-1}$ such a path P must exist. In addition, let $P' \in \mathcal{P}_{i-1}(u', v')$ be such that $P' \cap F = \emptyset$. Since $c(\phi) < d$ and since e' is a sink, such a path P' exists, as well. We use P, P' to define the solution $\{\mathcal{P}_i(u, v)\}_{(u, v) \in F}$ as follows:

- Initially, let $\mathcal{P}_i(x, y) = \mathcal{P}_{i-1}(x, y)$ for every $(x, y) \in F$.

- Replace P from $\mathcal{P}_i(u, v)$ with the path $P'' = P[u, u'] \circ P' \circ P[v', v]$.
- Replace P' from $\mathcal{P}_i(u', v')$ with $\{(u', v')\}$.

Observe that the congestion bound is preserved, hence $\mathcal{P}_i = \{\mathcal{P}_i(u, v)\}_{(u, v) \in F}$ has congestion of c . Note that e' is still a sink in D_{i-1} . More specifically, we have that $D_i = D_{i-1} \setminus \{(e, e')\}$. Since each iteration omits reduces the multiplicity of one edge in the dependency graph, within a finite number of steps, ℓ , we obtain a routing solution $\{\mathcal{P}_\ell(u, v)\}_{(u, v) \in F}$ with congestion of c and an empty dependency graph, as required. The lemma follows. \square

We are now ready to complete the proof of Lemma 4.1. By Lemma A.2, there is an output solution \mathcal{P}^* with congestion of c and with an empty dependency graph. Since $c < d$, we get that every $e = (u, v)$, it holds that $\mathcal{P}^*(u, v) \in \mathcal{P}^*$ has a u - v path P such that $e = (u, v) \notin P$. Moreover, since the dependency graph is empty, it also holds that $P \cap F = \emptyset$. Therefore, every for every $(u, v) \in F$, we get an alternative u - v path in $G \setminus F$, and therefore $G \setminus F$ is connected.

Proof of Lemma 3.3. We use the multi-commodity $(1 + \varepsilon)$ -approximate length-constrained flow algorithm of [HHS23, Theorem A.1] together with the flow-path decomposition algorithm [HHS23, Theorem 10.1] followed by a simple independent randomized rounding. We set $\varepsilon = \frac{1}{100mnC}$ and define the capacity of all edges in G to be c . Since our flow size is always integral and of size at most mnc this guarantees that the fractional flow routed by the $(1 + \varepsilon)$ -approximation algorithm loses at most $1/100$ units of flow overall and therefore (fractionally) routes demand of every demand pair up to an additive $1/100$ (or multiplicative factor of 1.01). We do not use any batching structure of the multi-commodity demands and simply make each demand pair its own batch, i.e., $\kappa \leq n^2$. The running time is polynomial in ε , κ , ℓ , c and $|D|$ and therefore polynomial overall. The algorithm produces a fractional multi commodity flow in the form of an average of polynomially many integral flows. We compute the flow-path decomposition for each of these integral flows using [HHS23, Theorem 10.1]. Overall this gives for every demand pair u, v with $D(u, v) > 0$ an explicit distribution $P_{u, v}$ with polynomial support over flow paths between u and v of length at most ℓ . Without loss of generality all routing paths are simple. We now independently sample exactly $D(u, v)$ paths from $P_{u, v}$ for every u, v to obtain a routing for D with dilation at most ℓ . For any edge in G the congestion of these sampled paths is at most $1.01c$ in expectation. Since every (simple) path creates at most one unit of congestion per edge if sampled and since all paths are sampled independently a multiplicative Chernoff bound guarantees that the probability of any fixed edge e having congestion more than $1.1c$ is at most $\exp(-\Theta(c)) = n^{-\omega(1)}$. Therefore, with high probability, this randomized rounding of the fractional multi-commodity flow has congestion at most $1.1c$. \square

Proof of Lemma 4.3. In this proof and following [HRG22], we define demand between edges for convenience. For any edge $e \in E$ and any demand $D : E \times E \rightarrow \mathbb{R}_{\geq 0}$, the load of D on e is the maximum between the demand sent or received by e in D . I.e., $load(e, D) = \max\{\sum_{e' \in E} D(e, e'), \sum_{e' \in E} D(e', e)\}$. Let $load(D) = \max_{e \in E} load(D, e)$. The demand is called unit if $load(D) \leq 1$. By Lemma 3.16 of [HRG22] (Lemma 3.4 for the edge-demand setting), it is sufficient to show that any unit-demand instance $I \subseteq E_H \times E_H$ is routable in H with congestion and dilation of $O(\log n/\phi)$.

Consider a unit-demand edge-instance $I_H \subseteq E_H \times E_H$ where each pair $(e, e') \in I_H$ is weighted by a demand function $D_H : E(H) \rightarrow \mathbb{R}_{\geq 0}$. We then have that

$$load(e, D_H) = \sum_{e' \in E(H)} \max\{D_H(e, e'), D_H(e', e)\} \leq 1, \forall e \in E_H.$$

We translate I_H into a unit-demand instance $I_{\widehat{G}} \subseteq \widehat{E} \times \widehat{E}$, as follows. Let $\mathcal{P} = \{P(e) \subseteq H \mid e \in \widehat{E}\}$ be the collection of paths that embed \widehat{G} into H . Since H is minimal, we have that $E(H) = \bigcap_{P \in \mathcal{P}} E(P)$.

For every $e \in H$, let $M(e) = \{e' \in \widehat{E} \mid e \in P(e')\}$ be the subset of edges in \widehat{E} such that e appears on their corresponding paths in \mathcal{P} . By the definition of H , we have that $M(e) \neq \emptyset$ for every $e \in E_H$. For each edge $e \in I_H$, we choose one such edge $f(e) = e'$ for some arbitrary $e' \in M(e)$.

Note that each edge $e' \in \widehat{G}$ is associated with at most $\ell = \tilde{O}(1/\phi)$ edges in H . Define a routing instance $I_{\widehat{G}} = \{(f(e), f(e')) \mid (e, e') \in I_H\}$ and let $D_{\widehat{G}}(f(e), f(e')) = d_H(e, e')$. The load of every edge $e' \in \widehat{E}$ can be bounded by:

$$\text{load}(e', D_{\widehat{G}}) \leq \sum_{e \in P(e')} \text{load}(e, D_H) \leq \ell.$$

Let $\widehat{G}_\ell = (V, \widehat{E}, c)$ be a capacitated graph in which each edge e' of \widehat{G} has a capacity of $c(e') = \ell$. Clearly, \widehat{G}_ℓ is a $(1/\log n)$ -expander, as well. Note also that $I_{\widehat{G}}$ is a unit-demand instance in \widehat{G}_ℓ , and hence it is routable in \widehat{G}_ℓ with congestion and dilation $\tilde{O}(1)$. Since each edge in \widehat{G}_ℓ has capacity of ℓ , the output path collection \mathcal{P}' of this routing instance has congestion of $\tilde{O}(\ell)$ and dilation $\tilde{O}(1)$. Each path $P' = [e_1, \dots, e_k]$ in \mathcal{P}' then translates into an H -path $P'' = P(e_1) \circ \dots \circ P(e_k)$ that contains a path between the edge pair $e, e' \in E(H)$ for every $(e, e') \in I_H$ such that $f(e) = e_1$ and $f(e') = e_k$. The length of these paths is bounded by $\tilde{O}(1/\phi)$ and the congestion is bounded by $\tilde{O}(\ell/\phi) = \tilde{O}(1/\phi^2)$. \square

B $G(n, p)$ is a good constant-hop router

Lemma B.1. *For any $n, p = \omega(\frac{\log n}{n})$ let G be a randomly sampled $G(n, p)$ graph. With high probability, any unit demand D in G can be routed with dilation $\ell = O(\log_{np} n)$ and congestion $c = O(\log n)$.*

Proof. Define $\varepsilon > 0$ such that $np = n^\varepsilon$. It is well known that for any sufficiently large $t = \Theta(\frac{1}{\varepsilon})$ a t -step random walk in $G(n, p)$ from any starting vertex suffices to polynomially close to its stationary distribution. This means no matter what node one starts from a t -step random walk will, up to polynomially small deviations, end at every node with probability proportional to its degree. Degrees in $G(n, p)$ with $p = \omega(\frac{\log n}{n})$ and average degree $np = \omega(\log n)$ are furthermore highly concentrated.

We will show that these facts imply that, with high probability, any unit demand D in G can be routed with dilation at most $2t$ and congestion at most $O(\log n)$ as claimed.

In particular, suppose each node v starts with $\deg(v) \cdot 2m$ many pebbles and each pebble (simultaneously) performs a random walk for t steps. Let P be the routing (i.e., collection of paths) traced out by these pebbles. Because of the above mixing property it is the case that, with high probability and up to $(1 + o(1))$ -factors, this routing P ends up again with exactly $\deg(v) \cdot 2m = \deg(v) \cdot \sum_u \deg(u)$ many pebbles at node v of which $\deg(u)$ many started from node u . This means P integrally routes exactly the demand $m \cdot D_{\text{mix}}$ where $D_{\text{mix}}(u, v) = \frac{\deg(v) \cdot \deg(u)}{m}$ is the so-called uniform mixing demand. The dilation of this routing is at most t and the expected congestion of the routing is at most mt . We can obviously route any unit demand D in G by using P to mix and unmix. In particular, for each u, v we construct $D(u, v)$ many routing paths by independently sampling an intermediate node x in G (sampling nodes proportional to their degree) and then independently sampling a path from u to x and from x to v from P and concatenating these two paths into a routing path from u to v . The dilation of this routing is at most $2t$ and the expected congestion for this routing is at most $2t$ on any edge e . After making routing paths simple the multiplicative Chernoff bound guarantees that any edge has congestion at most $O(\log n)$ with high probability. \square