

Partially Optimal Edge Fault-Tolerant Spanners

Greg Bodwin
University of Michigan
bodwin@umich.edu

Michael Dinitz*
Johns Hopkins University
mdinitz@cs.jhu.edu

Caleb Robelle
MIT
robelle@mit.edu

Abstract

Recent work has established that, for every positive integer k , every n -node graph has a $(2k - 1)$ -spanner with $O(f^{1-1/k} n^{1+1/k})$ edges that is resilient to f edge or vertex faults. For *vertex* faults, this bound is tight. However, the case of *edge* faults is not as well understood: the best known lower bound for general k is $\Omega(f^{\frac{1}{2}-\frac{1}{2k}} n^{1+1/k} + fn)$. Our main result is to nearly close this gap with an improved upper bound, thus separating the cases of edge and vertex faults. For odd k , our new upper bound is $O_k(f^{\frac{1}{2}-\frac{1}{2k}} n^{1+1/k} + fn)$, which is tight up to hidden $\text{poly}(k)$ factors. For even k , our new upper bound is $O_k(f^{1/2} n^{1+1/k} + fn)$, which leaves a gap of $\text{poly}(k) f^{1/(2k)}$. Our proof is an analysis of the fault-tolerant greedy algorithm, which requires exponential time, but we also show that there is a polynomial-time algorithm which creates edge fault tolerant spanners that are larger only by factors of k .

1 Introduction

Let $G = (V, E)$ be a graph, possibly with edge lengths $w : E \rightarrow \mathbb{R}_{\geq 0}$. A t -spanner of G , for $t \geq 1$, is a subgraph $H = (V, E')$ that preserves all pairwise distances within a factor of t , i.e., $d_H(u, v) \leq t \cdot d_G(u, v)$ for all $u, v \in V$ (where d_X denotes the shortest-path distance in a graph X). Since H is a subgraph of G it is also true that $d_G(u, v) \leq d_H(u, v)$, and so distances in H are the same as in G up to a factor of t . The distance preservation factor t is called the *stretch* of the spanner. Spanners were introduced by Peleg and Ullman [24] and Peleg and Schäffer [23], and have a wide range of applications in routing [25], synchronizers [3], broadcasting [2, 22], distance oracles [27], graph sparsifiers [17], preconditioning of linear systems [14], etc.

The most common objective in spanners research is to achieve the best possible existential size-stretch trade-off. Most notably, a landmark result of Althöfer et al. [1] analyzed the following simple and natural greedy algorithm: given an n -node graph G and an integer $k \geq 1$, consider the edges of G in non-decreasing order of their weight and add an edge (u, v) to the current spanner H if and only if $d_H(u, v) > (2k - 1)w(u, v)$. They proved that this algorithm produces $(2k - 1)$ -spanners of *existentially optimal size*. In particular, the spanner produced has size $O(n^{1+1/k})$, and assuming the well-known *Erdős girth conjecture* [15] there are graphs in which every $(2k - 1)$ spanner (and in fact every $2k$ -spanner) has at least $\Omega(n^{1+1/k})$ edges.

A crucial aspect of real-life systems that is not captured by the standard notion of spanners is the possibility of *failure*. If some edges (e.g., communication links) or vertices (e.g., computer processors) fail, what remains of the spanner might not still approximate the distances in what remains of the original graph. This motivates the notion of *fault tolerant* spanners:

DEFINITION 1.1. (FAULT TOLERANT SPANNERS) *A subgraph H is an f -edge fault tolerant (f -EFT) t -spanner of $G = (V, E)$ if*

$$d_{H \setminus F}(u, v) \leq t \cdot d_{G \setminus F}(u, v)$$

for all $u, v \in V$ and $F \subseteq E$ with $|F| \leq f$.

In other words, an f -EFT spanner contains a spanner of $G \setminus F$ for every set of $|F| \leq f$ edges that could fail. The definition for vertex fault tolerance (VFT) is equivalent, with the only change being that $F \subseteq V \setminus \{u, v\}$.

Fault tolerant spanners were originally introduced in the setting of geometric graphs by Levkopoulou, Narasimhan, and Smid [18] and have since been studied extensively in the settings of geometric graphs and

*Supported in part by NSF award CCF-1909111.

(more generally) doubling metrics; see e.g., [8, 9, 11, 19–21, 26] and references within. Chechik, Langberg, Peleg and Roditty [10] were the first to study fault-tolerant spanners in general graphs, giving a construction of an f -VFT $(2k - 1)$ -spanner of size $O(f^2 k^{f+1} \cdot n^{1+1/k} \log^{1-1/k} n)$ and an f -EFT $(2k - 1)$ -spanner of size $O(f \cdot n^{1+1/k})$. So they showed that introducing tolerance to f edge faults costs us an extra factor of f in the size of the spanner, while introducing tolerance to f vertex faults costs us a factor of roughly $f^2 k^{f+1}$ in the size (compared to the size of a non-fault tolerant spanner of the same stretch). Dinitz and Krauthgamer [12] later improved the extra factor paid in the VFT setting to about f^3 . At this point, it appeared that the EFT setting might be substantially easier than the VFT setting, in the sense that it allowed for a smaller dependence on f in spanner size. However, a recent series of papers has developed a set of techniques that apply equally well to both settings, yielding the same improved bounds for each [5–7, 13]. This has culminated in the following theorem:

THEOREM 1.1. (FT UPPER BOUNDS [6]) *There is a polynomial time algorithm that, given any positive integers f, k and any n -node graph G , constructs an f -EFT or VFT $(2k - 1)$ -spanner of G with $O(f^{1-1/k} n^{1+1/k})$ edges.*

This upper bound is notable because it fully matches the known lower bound for VFT spanners:

THEOREM 1.2. (VFT LOWER BOUNDS [5]) *Assuming the girth conjecture [15], for any positive integers n, f, k , there exist n -node graphs in which every f -VFT $(2k - 1)$ -spanner has $\Omega(f^{1-1/k} n^{1+1/k})$ edges.*

Thus the VFT setting is now essentially completely understood. However, there are still significant gaps in the EFT setting. While the VFT lower bound from [5] also holds for the EFT setting when $k = 2$, for general k it drops off considerably. (We note that the fn term in the following theorem statement is not explicitly given in [5], but is a simple folklore result stating that any f -regular graph is the unique f -EFT spanner of itself):

THEOREM 1.3. (EFT LOWER BOUNDS [5]) *Assuming the girth conjecture [15], for any positive integers n, f, k , there exist n -node graphs in which every f -EFT $(2k - 1)$ -spanner H with*

$$|E(H)| = \begin{cases} \Omega(f^{1/2} n^{3/2}) & k = 2 \\ \Omega(f^{1/2-1/(2k)} n^{1+1/k} + fn) & k > 2 \end{cases}$$

Hence there is a substantial $\text{poly}(f)$ gap in the EFT setting when $k > 2$ (note that f can be polynomial in n , so this gap is significant). So the main open question left is how to close this gap: how much extra do we have to pay to achieve edge fault tolerance compared to a non-fault tolerant spanner? Is the right answer $f^{\frac{1}{2}-\frac{1}{2k}}$, or $f^{1-1/k}$, or somewhere in between?

1.1 Our Results

In this paper we nearly resolve this question by improving the upper bound:

THEOREM 1.4. (MAIN RESULT) *For any positive integers n, f, k , every n -node graph has an f -EFT $(2k - 1)$ -spanner H with*

$$|E(H)| = \begin{cases} O(k^2 f^{1/2-1/(2k)} n^{1+1/k} + kfn) & k \text{ is odd} \\ O(k^2 f^{1/2} n^{1+1/k} + kfn) & k \text{ is even.} \end{cases}$$

Hence Theorem 1.4 *entirely resolves* the question of edge fault tolerance for constant odd k , and is off from the lower bound (Theorem 1.3) by only quadratic factors of k for nonconstant odd k , and only $k^2 f^{1/(2k)}$ for even k . See Table 1 for the full context of our results, and see Figure 1 in Appendix A for a visualization.

While it is interesting and important to optimize the dependence on k (and close the gap between our new upper bounds and the known lower bounds), we remark that spanners are never used with k larger than $O(\log n)$, since the additional stretch no longer yields additional sparsity. Hence the factor of k^2 in our bounds is at most polylogarithmic in n . On the other hand, the fault parameter f can be significantly larger (for example, polynomial in n). Hence the more central question in the area is to optimize dependence on f , and we are more concerned with our gap of $f^{1/(2k)}$ for even k than our gap of k^2 . However, we conjecture that our upper bound is actually correct, and that our current construction and analysis are tight with respect to f :

CONJECTURE 1.1. *For any positive integers n, f and even integer k , there exist n -node graphs in which every f -EFT $(2k - 1)$ -spanner has $\Omega(f^{1/2} n^{1+1/k})$ edges.*

Spanner size	Tight?	Polytime?	Citation
$O(f \cdot n^{1+1/k})$		✓	[10]
$O(\exp(k)f^{1-1/k} \cdot n^{1+1/k})$	$k = 2$ only		[5]
$O(f^{1-1/k} \cdot n^{1+1/k})$	$k = 2$ only		[7]
$O(kf^{1-1/k} \cdot n^{1+1/k})$	$k = 2$ only	✓	[13]
$O(f^{1-1/k} \cdot n^{1+1/k})$	$k = 2$ only	✓	[6]
$O(k^2 f^{1/2} \cdot n^{1+1/k} + kfn)$ for even k	$k = 2$ only (*)		this paper
$O(k^2 f^{1/2-1/(2k)} \cdot n^{1+1/k} + kfn)$ for odd k	all fixed odd k		this paper
$O(k^{5/2} f^{1/2} \cdot n^{1+1/k} + k^2 fn)$ for even k	$k = 2$ only (*)	✓	this paper
$O(k^{5/2} f^{1/2-1/(2k)} \cdot n^{1+1/k} + k^2 fn)$ for odd k	all fixed odd k	✓	this paper

Table 1: Prior work on the size of f -EFT $2k - 1$ spanners, and our new bounds in this paper. Tightness is implied by a lower bound from [5] of $\Omega(f^{1/2-1/(2k)}n^{1+1/k} + fn)$, except for $k = 2$ where the lower bound improves to $\Omega(f^{1/2}n^{1+1/k})$. Except for $k \in \{2, 3, 5\}$, and for $k = 7$ and large enough f (see Theorem B.1 to follow), this lower bound and our claims of tightness are conditional on the girth conjecture [15]. The (*) indicates that we conjecture our bound for even k to be tight for all fixed even k .

Settling this conjecture, or generally making progress on the gap between the $f^{1/2}$ dependence in this work and the $f^{\frac{1}{2}-\frac{1}{2k}}$ lower bound on f -dependence from [5], is the main open question left by this work. We remark that there is some precedent in the literature for fundamentally different behavior for $(2k - 1)$ -spanners under even/odd k . This occurs, for example, in the *unbalanced Moore bounds* that control the maximum possible density of a bipartite graph with a very different number of nodes on its two sides [16]. We believe that a similar even/odd effect may occur here, where the imbalance is essentially caused by the structure of edge faults. The fact that the dependence $f^{1/2}$ is known to be tight for $k = 2$ partially validates this intuition, and it may be “simpler” for the correct bounds to distinguish between even/odd k rather than taking $k = 2$ as a special case.

Algorithmic Efficiency. The spanner construction algorithm that we use to prove Theorem 1.4 is the same greedy algorithm as in [5, 7] (adapted for edge fault tolerance), which requires exponential time. However, by combining our new analysis with the ideas used by [13], we can obtain polynomial time at the price of a slightly worse dependence on k :

THEOREM 1.5. *There is a polynomial time algorithm that, given positive integers f, k and an n -node input graph, outputs an f -EFT $(2k - 1)$ -spanner H with*

$$|E(H)| = \begin{cases} O(k^{5/2} f^{1/2-1/(2k)} n^{1+1/k} + k^2 fn) & k \text{ is odd} \\ O(k^{5/2} f^{1/2} n^{1+1/k} + k^2 fn) & k \text{ is even.} \end{cases}$$

2 Technical Overview

Before diving into the details, we first give some intuition and context for our techniques.

Background on Non-Faulty Spanners. In the non-faulty setting, the analysis of the greedy construction by Althöfer et al. [1] passes through the *girth* of the output spanner (recall that the girth of a graph is the smallest number of edges in a cycle). One proves that the output spanner always has girth $> 2k$, and conversely, any unweighted graph of girth $> 2k$ has no $(2k - 1)$ -spanner except for the graph itself. Phrased another way, let us write $\gamma(n, 2k)$ for the maximum number of edges in any n -node graph of girth $> 2k$. Then the greedy algorithm outputs spanners on $\leq \gamma(n, 2k)$ edges, and this bound is best possible.

The only known argument to upper bound the value of $\gamma(n, 2k)$ is the *Moore bounds*, which prove that $\gamma(n, 2k) = O(n^{1+1/k})$ (the girth conjecture [15] is that this is tight, i.e., $\gamma(n, 2k) = \Omega(n^{1+1/k})$). These use a

counting argument over the simple k -paths of the input graph, where a “simple” path is one that does not repeat nodes. The Moore bounds are proved in two steps. First, one proves a *counting lemma*: any n -node graph of average degree $d \gg 1$ and girth $> 2k$ has $n \cdot \Omega(d)^k$ total simple k -paths. Then, one proves a *dispersion lemma*: in a graph of girth $> 2k$, no two simple k -paths may share endpoints. Together these imply $n \cdot \Omega(d)^k = O(n^2)$, and the Moore bound are obtained by rearranging this equation.

Background on FT-Spanners. In the faulty setting, it is natural to consider the “FT-Greedy Algorithm,” which straightforwardly extends the non-faulty greedy algorithm. That is: consider each edge (u, v) in order of nondecreasing weight, and add (u, v) to the spanner if and only if there is a set of $|F| \leq f$ faults such that $d_{H \setminus F}(u, v) > (2k - 1) \cdot d_{G \setminus F}(u, v)$. Correctness is again easy to show, but the challenge is to control the number of edges in the output spanner. Informally speaking, the output spanner ought to be “close” to a high-girth graph – not in the sense that its girth is high (the girth of the output spanner could be 3), but in the sense that the techniques used to bound the density of high-girth graphs also apply to the output spanner. In the previous literature, there are two basic approaches to analyzing the density of the output spanner from the greedy algorithm:

1. The first analysis of the greedy algorithm [5] argued that greedy FT-spanners are similar to high-girth graphs in the sense that they are amenable to a generalized version of the Moore bound analysis. As discussed, the original Moore bounds work by arguing that (1) any two nodes in the output spanner of the (non-faulty) greedy algorithm have at most one simple k -path going between them, and (2) any graph satisfying this property must be sparse. Very roughly, the analysis of the FT-greedy algorithm in [5] works by limiting the number of “non-backtracking” k -paths going between any two nodes as a function of f and k , and then arguing that a bounded number of these non-backtracking k -paths again implies a (weaker) bound on the number of edges in the spanner. Quite a few technical details arise in this approach, and in [5] it led to a rather complicated analysis that took on $\exp(k)$ factors in spanner size, but which gave optimal dependence on n and f for the VFT setting, assuming the girth conjecture. Indeed, arguments of this type can only hope to achieve *conditional* optimality on the girth conjecture, since the Moore bounds themselves are only optimal if the girth conjecture holds.
2. The following analyses [6, 7, 13] took an alternate view that greedy FT-spanners are *structurally similar* to high-girth graphs; in particular, they certify sparsity by showing that the output spanners of the FT-greedy algorithm have high-girth subgraphs that keep most of the density of the original graph. This high-girth subgraph must be sparse, since the Moore bounds apply to this subgraph directly, and the way the subgraph is constructed ensures that its density cannot be too far away from that of the original output spanner. This approach has a number of benefits over the generalized Moore argument: it leads to a far simpler analysis, it removes all extra factors of k in the spanner size, and it lets one “black-box” the Moore bounds. That is, it gives bounds that depend directly on $\gamma(n, 2k)$ rather than the Moore bounds for this function, thus establishing optimal spanner size in the VFT setting even if the girth conjecture fails.

So the message of these papers seemed to be that the structural approach, which only uses the Moore bounds as a black-box to plug into $\gamma(n, 2k)$, dominates in every important way the approach that opens the black box of the Moore bounds.

Since a gap remained for EFT spanners, it was still technically unresolved whether Moore or structural arguments are to be preferred in this arena. This point was discussed in [7], which contained a technical barrier suggesting that it would be difficult to push the EFT upper bounds any further using the latter structural approach. Specifically, they suggested the following objects for EFT structural analysis:

DEFINITION 2.1. (EDGE BLOCKING SETS [7]) *Given a graph $G = (V, E)$, a t -edge blocking set for G is a set $B \subseteq \binom{E}{2}$ such that, for every cycle $C \in G$ with at most t edges, there exists $(e_1, e_2) \in B$ such that $e_1, e_2 \in C$.*

One proves the previous-best upper bounds on FT-greedy spanners by (1) observing that the output spanners H from the EFT-greedy algorithm have $2k$ -edge blocking sets of size $|B| \leq f|E(H)|$, and then (2) proving that *any* graph with a $2k$ -edge blocking set of size $|B| \leq f|E(H)|$ can only have $O(f^{1-1/k}n^{1+1/k})$ edges. The second lemma is the key to the structural approach, which appeared again in some form in followup work [6, 13]. However, it was further proved in [7] that this second lemma is tight: there exist graphs that have a blocking set of size $|B| \leq f|E(H)|$ and $\Omega(f^{1-1/k}n^{1+1/k})$ edges. So, if better EFT upper bounds are to be achieved, they at least need to use a different object from blocking sets, and perhaps depart from this style of argument altogether.

2.1 Our Approach

We show, perhaps surprisingly, that the original Moore-based approach seems to dominate the structural one in the context of EFT spanners. Our analysis is best viewed as a generalization of the Moore bounds, and moreover, as an auxiliary result we strengthen the evidence from [7] that structural arguments cannot possibly achieve the improved EFT upper bounds obtained in this paper.

First, to sidestep the barrier from [7], we need to change the focus of our analysis. We use *strong blocking sets*, a natural extension that explicitly takes edge weights into account. While the original definition of blocking sets is just about edges in cycles, we now require one of the edges to be the *heaviest* edge. More formally:

DEFINITION 2.2. (STRONG BLOCKING SETS) *A strong t -blocking set is some $B \subseteq \binom{E}{2}$ such that, for any cycle C on at most t edges, there is $(e_1, e_2) \in B$ with $e_1, e_2 \in C$ and such that one of e_1, e_2 is the heaviest edge in C . The elements $(e_1, e_2) \in B$ are called blocks.*

(We assume that ties between edge weights are broken in some canonical way, and so “the heaviest edge” in C is unambiguous.) It is again easy to prove that the FT-Greedy algorithm produces spanners that have strong blocking sets of size $\leq f|E(H)|$. Since the girth of a graph has nothing to do with its edge weights, we can view this definition as a step away from a purely structural approach, and towards an approach that cares about the specific interaction between the input graph and the FT-greedy algorithm (which considers edges in increasing weight order).

The main new technical idea goes into our generalized dispersion lemma. The goal is to bound the total number of simple k -paths in the graph. The core of the argument is the following lemma: up to some technical details that we will not discuss here (see Section 4), due to the existence of a small strong blocking set, it is nearly true that for any two nodes s, t in the output spanner, there is a set of $O(kf)$ edges such that every simple $s \rightsquigarrow t$ k -path uses one of these edges as its *heaviest* edge. We may then bound the number of simple $s \rightsquigarrow t$ k -paths by considering each edge (u, v) in this set and recursively bounding the number of $s \rightsquigarrow u$ j -paths and the number of $v \rightsquigarrow t$ $(k - j - 1)$ -paths (for different values of j).

This recursion turns out to give the strongest bounds when $1 \leq j \leq k - 1$; that is, (u, v) is a *middle* edge along a path, not the first or last one. This motivates a narrowing of focus in our counting argument from all simple k -paths to only those that are *middle-heavy* (have their heaviest edge in a middle position). In fact, we need our paths to stay middle-heavy over the recursion, as they are repeatedly split over their heaviest edge. So we focus even more narrowly on a class of *alternating paths*, for which each even-numbered edge is heavier than the two odd-numbered edges on either side, which have this favorable property. This induces some additional technical work in the counting lemma: standard arguments guarantee the existence of many simple k -paths, but we need to rework these arguments to show that a good fraction of these simple k -paths are indeed alternating.

Lower Bounds on Structural Analysis. Although strong blocking sets move away from structural analyses a bit, it is still reasonable to ask whether our argument can be replaced with a structural one somehow. Next, we explain our new evidence that the answer to this question is probably “no”: we show that a structural analysis would by necessity involve solving a long-standing open problem in extremal graph theory. As mentioned earlier, a key advantage of structural analyses is that they give spanner upper bounds directly in terms of the function γ . For example, [7] shows that all graphs have EFT or VFT $(2k - 1)$ spanners on $O(f^2 \cdot \gamma(n/f, 2k))$ edges. The Moore bounds are applied as a secondary step to bound γ , giving the upper bound of $O(f^{1-1/k} n^{1+1/k})$ from Theorem 1.1. But if the girth conjecture fails, and there is an even better upper bound for γ than the Moore bounds, then the upper bounds from these arguments automatically improve as well. Conversely, if one can prove *unconditionally* that f -VFT spanners need $\Omega(f^{1-1/k} n^{1+1/k})$ edges, this would imply a lower bound on γ that would be enough to prove the girth conjecture, which is currently unknown and is a famous and longstanding open problem.

Let us now entertain the hypothetical that one could prove our upper bounds for EFT spanners with a structural argument, arguing that greedy EFT-spanners can be reduced to high-girth graphs while sacrificing only a small amount of density. For example, analogous to [7], let us imagine that we have a bound of the form $O_k(f \cdot \gamma(n/\sqrt{f}, 2k) + nf)$ edges for odd k , which would in turn imply our upper bound of $O_k(f^{1/2-1/(2k)} n^{1+1/k} + nf)$ for odd k by applying the Moore bounds to γ . Then if one could prove *unconditionally* that f -EFT spanners need $\Omega(f^{1/2-1/(2k)} n^{1+1/k} + nf)$ edges for fixed odd k , it would imply a lower bound on γ that would be enough to prove the girth conjecture for fixed odd k , which would again be a major breakthrough.

In fact, we can prove a new unconditional bound of this type – not for all fixed odd k , but specifically for

$k = 7$. A proof of the following theorem appears in Appendix B.

THEOREM 2.1. *For $k = 7$ and $f = \Omega(n^{1/11})$, the lower bound on EFT spanners in Theorem 1.3 holds unconditionally.*

Thus, a structural argument as explained above would be enough to prove the girth conjecture for $k = 7$, which is currently unknown. Since no new setting of the girth conjecture has been settled in over 50 years, this is likely very hard, and would constitute a major breakthrough.

3 From Greedy to Strong Blocking

We begin our technical work by proving that the greedy FT-spanner algorithm analyzed in [5, 7] gives spanners with small strong blocking sets.

Algorithm 1 Greedy f -EFT $(2k - 1)$ -Spanner Algorithm

```

function FT-GREEDY( $G = (V, E, w)$ ,  $k, f$ )
 $H \leftarrow (V, \emptyset, w)$ 
for all  $(u, v) \in E$  in nondecreasing weight order do
    if there exists a set  $F$  of at most  $f$  edges such that  $d_{H \setminus F}(u, v) > (2k - 1)w(u, v)$  then
        add  $(u, v)$  to  $H$ 
return  $H$ 

```

This algorithm implicitly breaks ties between edges of equal weight, to consider one before the other in the main for-loop. This tiebreaking is arbitrary, but in the sequel it will be convenient to unambiguously refer to *the heaviest* edge among a set of edges, which means the heaviest edge under this tiebreaking (equivalently, the edge considered last by the greedy algorithm). In particular, in the definition of strong blocking sets, it is required that for each cycle C , there is a block (e_1, e_2) with $e_1, e_2 \in C$ and such that either e_1, e_2 is the heaviest edge in C under *this tiebreaking*.

The following lemma is a straightforward adaptation of Lemma 3 of [7], where the same statement was proved for blocking sets rather than strong blocking sets.

LEMMA 3.1. *Any graph H returned by Algorithm 1 is an f -EFT $(2k - 1)$ -spanner that has a strong $2k$ -blocking set of size at most $f|E(H)|$.*

Proof. See Appendix C. □

One downside of this algorithm is its running time. In each iteration of this algorithm, the algorithm needs to decide the following question: given an edge (u, v) , is there a fault set $F \subseteq E$ where $|F| \leq f$ and $d_{H \setminus F}(u, v) \leq (2k - 1) \cdot d_{G \setminus F}(u, v)$? Doing this in the obvious way (checking all sets of f or fewer edges) takes time exponential in f . And, unfortunately, it turns out that this question is equivalent to the Length Bounded Cut problem (LBC), which is known to be NP-hard [4].

In order to develop a polynomial time algorithm for constructing FT-spanners, [13] replaced this exponential time subroutine with one that ran in polynomial time. They designed a simple $(2k - 1)$ -approximation algorithm for the *unweighted* setting which is essentially the standard frequency approximation of Set Cover. In particular, when used for the decision variant it has the property that if there is a fault set of size at most f that intersects all $u - v$ paths with at most $2k - 1$ hops then the algorithm will return YES, while if all such fault sets have size larger than $(2k - 1)f$ then it will return NO (if neither is true then the algorithm can return either answer). We call this algorithm $\mathcal{A}(G, u, v, f, k)$.

A key insight of [13] was that, even though \mathcal{A} is designed for the unweighted setting, it can still be used in the weighted setting. In particular, if we use the weights to set the initial ordering of edges (as in Algorithm 1) but from then on pretend that the graph is unweighted, this modified greedy algorithm still returns a valid f -EFT $(2k - 1)$ -spanner. This gives Algorithm 2, a polynomial time algorithm for constructing EFT spanners.

In order to analyze the sparsity of spanners constructed by this algorithm, [13] showed that the blocking set analysis of [7] can be applied with an extra loss of $O(k)$ in the size (due to the approximation). We show the same, but for our new notion of strong blocking sets.

Algorithm 2 Modified Greedy EFT Spanner Algorithm

```

function FT-GREEDY( $G = (V, E), k, f$ )
 $H \leftarrow (V, \emptyset)$ 
for all  $(u, v) \in E$  in nondecreasing weight order do
    if  $\mathcal{A}(H, u, v, f, k)$  returns YES then
        Add  $(u, v)$  to  $H$ 
return  $H$ 

```

LEMMA 3.2. *The graph H returned by Algorithm 2 is an f -EFT $(2k - 1)$ -spanner that has a strong $2k$ -blocking set of size at most $O(kf) \cdot |E(H)|$.*

Proof. See Appendix C □

4 From Strong Blocking Sets to Extremal Density Bounds

In this section, we will prove:

THEOREM 4.1. *Let f be a positive integer, let k be a positive integer, and let H be an n -node graph that has a strong $2k$ -blocking set B of size $|B| \leq |E|f$. Then*

$$|E(H)| = \begin{cases} O(k^2 f^{1/2-1/(2k)} n^{1+1/k} + nfk) & k \text{ is odd} \\ O(k^2 f^{1/2} n^{1+1/k} + nfk) & k \text{ is even.} \end{cases}$$

This theorem together with Lemma 3.1 directly implies Theorem 1.4. For the polynomial-time version, note that Lemma 3.2 gives us a polytime algorithm with $|B| = O(kf|E(H)|)$, so we can reparameterize Theorem 4.1 by setting the f in the theorem equal to $O(kf)$. This implies Theorem 1.5.

4.1 Proof Preliminaries

Our proof will be a counting argument over a special type of path in the graph H :

DEFINITION 4.1. (ALTERNATING PATHS) *An alternating k -path in H is a path π with edge sequence (e_1, e_2, \dots, e_k) with the following property: every even-numbered edge is heavier than the two odd-numbered edges adjacent to it. (When k is even, the last edge e_k is only required to be heavier than e_{k-1} .)*

More specifically, we will count *simple unblocked* alternating paths. As usual, a *simple* path is one that does not repeat nodes. An *unblocked* path is defined as follows:

DEFINITION 4.2. (BLOCKED PATHS) *A path π is blocked by a strong blocking set B if there is $(e_1, e_2) \in B$ with $e_1, e_2 \in \pi$. Otherwise, π is unblocked. (Note: unlike for cycles, π is blocked even if neither e_1 nor e_2 is the heaviest edge in π .)*

In the remainder of this section we follow the outline discussed in Section 2, giving a counting lemma and a dispersion lemma for simple unblocked alternating k -paths. But we first start by making some useful simplifications, which are all without loss of generality:

1. We will use the stronger hypothesis that in the strong blocking set B , each edge appears in $\leq f$ blocks. This is without loss of generality for the following reason. Since every block contains two edges, the average edge participates in $\leq 2f$ blocks, and hence a simple counting argument (or Markov's inequality) implies that at most half of the edges participate in at least $4f$ blocks. Suppose that, while there is an edge that participates in $\geq 4f$ blocks, we delete that edge from the graph and the corresponding blocks from the blocking set. We thus delete at most half of the graph edges in this way, and we arrive at a graph with a blocking set in which every block participates in $< 4f$ blocks. We can then reparametrize $f \leftarrow 4f$ while changing our claimed upper bound on spanner size by only a constant factor.
2. We assume that H has average degree $d \geq 32fk$. In the regime where this does not hold, H has $nd/2 = O(nfk)$ total edges and so the claimed bounds hold already due to the trailing additive term.

3. Additionally, letting d be the average degree in H , we will assume that the maximum degree is $O(d)$. This is justified by the following operation: while there is a node v of degree $\geq 4d$, split it into two nodes v_1, v_2 , equitably partitioning the edges that use v between its two copies. Notice that this splitting operation does not create cycles or edges (though it may destroy some cycles), and thus B is still a strong blocking set for the modified graph. A quick counting argument, standard in prior work (for example, [5]), shows that one introduces only $O(n)$ new nodes in this way, and hence this again changes our claimed upper bounds on spanner size only by a constant factor.
4. We assume that all edges in H have distinct weights, so that we may unambiguously refer to *the heaviest* or *lightest* edge among a set of edges. This is justified by simply breaking weight ties in some arbitrary but consistent way, as in our discussion of the FT-greedy algorithm.

4.2 Counting Lemma for Alternating Paths

First, we prove a counting lemma for alternating paths. We will start with some weaker versions of the counting lemma we need, and then gradually bootstrap them into a full one. In the following lemma statements, an *edge-simple alternating path* is an alternating path that does not repeat edges (but which might repeat nodes).

LEMMA 4.1. (WEAK COUNTING LEMMA) *For any positive integer k , any n -node graph Γ with $\geq kn$ edges has an edge-simple alternating k -paths.*

Proof. The proof is by induction on k . For the base case, when $k = 0$, any individual node may be viewed as an edge-simple alternating 0-path.

Now we prove the inductive step. We will first assume that k is odd. Assume without loss of generality that H has no isolated nodes, and preprocess Γ by removing the lightest edge incident to each node. We remove at most n edges in this way, so at least $(k - 1)n$ edges remain. By the inductive hypothesis, there exists an edge-simple alternating $k - 1$ path π in the surviving graph. Let (u, v) be the last edge of π . Since $k - 1$ is even, we can extend π into an alternating k -path by appending any edge incident to v that is lighter than (u, v) . Letting (v, x) be the edge removed in the preprocessing, we know that (v, x) is lighter than (u, v) , since it was selected to be removed instead of (u, v) . Thus we can extend π into an alternating k -path by appending (v, x) , and since π was generated in a graph that did not contain (v, x) , it remains edge-simple under this extension.

The case where k is even is similar, except that we remove the *heaviest* edge incident to each node in the preprocessing. \square

We next bootstrap into a stronger counting lemma over the same kinds of paths:

LEMMA 4.2. (INTERMEDIATE COUNTING LEMMA) *For any positive integer k , any n -node graph Γ with $\geq 2kn$ edges has $\geq kn$ total edge-simple alternating k -paths.*

Proof. Consider the following process: pick an edge that is contained in at least one edge-simple alternating k -path, remove it, and repeat until there are no more edge-simple alternating k -paths. Since we remove the edge we find in each iteration, there are at least as many edge-simple alternating k -paths as there are iterations. And there are at least kn iterations, since as long as $\geq kn$ edges remain in the graph, Lemma 4.1 implies that there will still be an edge-simple alternating k -path. Hence there are at least kn edge-simple alternating k -paths in total. \square

Now we bootstrap this once again into the counting lemma we will actually use in our main argument. While the previous two lemmas hold for arbitrary graphs Γ , the following one holds specifically for the graph H under consideration, which has a small strong blocking set.

LEMMA 4.3. (FULL COUNTING LEMMA) *The graph H has $n \cdot \Omega(d/k)^k$ total unblocked simple alternating k -paths, where d is its average degree.*

Proof. In this proof, we will actually count the unblocked edge-simple alternating k -paths in H . We notice that, if an edge-simple k -path repeats a node, then it must have a cycle of length $\leq k$ as a subpath, and this cycle must be blocked by B . Thus every path π that is edge-simple, but not simple, is already blocked, and hence

the number of unblocked edge-simple alternating k -paths is equal to the number of unblocked simple alternating k -paths. In the following, let $\alpha(H)$ denote the number of unblocked (edge-)simple alternating k -paths in H (and we use similar notation for other graphs that will arise in the argument).

Let H' be a random subgraph of H obtained by including each edge with probability $8(d/k)^{-1}$. So we have

$$\mathbb{E}[|E(H')|] = |E(H)| \cdot 8(d/k)^{-1} = \frac{nd}{2} \cdot 8(d/k)^{-1} = 4nk.$$

Since each edge is sampled independently, a standard Chernoff bound implies that with high probability $|E(H')| \geq (15/4)nk$.

Let $B' \subseteq B$ be the strong blocking set for H' obtained by keeping the blocks $(e_1, e_2) \in B$ where both e_1, e_2 survive in H' . We have

$$\begin{aligned} \mathbb{E}[|B'|] &= |B| \cdot 64(d/k)^{-2} \leq |E(H)|f \cdot 64(d/k)^{-2} \\ &= \frac{nd}{2} \cdot f \cdot 64(d/k)^{-2} \\ &= 32nfk^2/d \\ &\leq nk \end{aligned}$$

where the last inequality uses our assumption that $d \geq 32fk$. Hence another standard Chernoff bound implies that $|B'| \leq (5/4)nk$ with high probability. Now, build a subgraph H'' from H' by considering each block $(e_1, e_2) \in B'$ and deleting e_1 . We thus have that with high probability,

$$|E(H'')| = |E(H')| - |B'| \geq (15/4)nk - (5/4)nk = (5/2)nk.$$

By Lemma 4.2, a graph on $\geq 2nk$ edges has $\geq kn$ total edge-simple alternating k -paths. By construction *no* k -paths in H'' are blocked (since we have removed at least one edge from every block originally in B). Thus, we have that $\alpha(H'') \geq kn$ with high probability, and hence $\mathbb{E}[\alpha(H'')] = \Omega(kn)$. Since $H'' \subseteq H$, this implies that $\mathbb{E}[\alpha(H')] = \Omega(kn)$.

Let us now compute $\mathbb{E}[\alpha(H')]$ in a different way. Recall that H' was obtained from H by including each edge with probability $O(d/k)^{-1}$. So each individual edge-simple k -path from H survives in H' with probability $O(d/k)^{-k}$. So we have

$$\mathbb{E}[\alpha(H')] = \alpha(H) \cdot O(d/k)^{-k}.$$

Putting together these two bounds on $\mathbb{E}[\alpha(H')]$, we get

$$\alpha(H) = \Omega(kn) \cdot \Omega(d/k)^k = n \cdot \Omega(d/k)^k,$$

as claimed. □

4.3 Dispersion Lemma for Alternating Paths

In this part, we bound the maximum number of alternating paths that can go between two nodes. First we prove a useful intermediate lemma which applies to all unblocked k -paths (not just alternating):

LEMMA 4.4. *For any nodes s, t , there is a set $F_{s,t}$ containing $|F_{s,t}| \leq kf + 1$ edges such that every simple unblocked $s \rightsquigarrow t$ path of length $\leq k$ in H uses an edge in $F_{s,t}$ as its heaviest edge.*

Proof. We build $F_{s,t}$ iteratively as follows. Initially $F_{s,t} := \emptyset$. In each round, pick the heaviest edge contained in any simple unblocked $s \rightsquigarrow t$ path of length $\leq k$ whose heaviest edge is not already in $F_{s,t}$, and add this edge to $F_{s,t}$. Halt once no such path exists.

To get the size bound on $F_{s,t}$, suppose for contradiction that after $kf + 1$ rounds there remains a simple unblocked $s \rightsquigarrow t$ path π of length $\leq k$ whose heaviest edge is not already in $F_{s,t}$. For each $f \in F_{s,t}$, by construction there is a simple unblocked $s \rightsquigarrow t$ path π' of length $\leq k$ that uses f as its heaviest edge, and f is heavier than any edge in π . Thus $\pi \cup \pi'$ contains a cycle C of length $\leq 2k$, which has f as its heaviest edge. This cycle C must be blocked by B , so there is $(f, e) \in B$ with $f, e \in C$. Since the path π' is unblocked, and $f \in \pi'$, it must be that $e \in \pi$.

But there are $\leq k$ edges in π , and there are $kf + 1$ edges in $F_{s,t}$, so by the pigeonhole principle one edge in π appears in $\geq f + 1$ blocks. This contradicts our initial assumption from Section 4.1 that each edge appears in $\leq f$ total blocks in B . \square

LEMMA 4.5. (DISPERSION LEMMA) *For any nodes s, t and positive integer $j \leq k$, the number of simple unblocked alternating $s \rightsquigarrow t$ paths of length j in H is*

$$\begin{cases} O(k^2 f)^{(j-1)/2} & j \text{ is odd} \\ O(k^2 f)^{j/2} & j \text{ is even} \end{cases}$$

Proof. The proof is by strong induction over j . The base cases are when $j = 0$, in which case there is at most one 0-path (when $s = t$), and also when $j = 1$, in which case only a single (s, t) edge can be a 1-path.

For the inductive step, we apply Lemma 4.4 to find a set of $|F_{s,t}| = O(kf)$ edges for which every simple unblocked alternating $s \rightsquigarrow t$ path of length k has its heaviest edge in $F_{s,t}$. The alternating $s \rightsquigarrow t$ paths may then be partitioned into $O(k^2 f)$ equivalence classes, by (1) which edge in $F_{s,t}$ is their heaviest, and (2) in which of the $j \leq k$ possible positions in the path the heaviest edge occurs.

We now bound the size of each individual equivalence class using the inductive hypothesis. Consider the equivalence class of simple unblocked alternating paths whose heaviest edge is $(u, v) \in F_{s,t}$, and which use (u, v) in the i^{th} position. Notice that each such path may be viewed as the concatenation of a simple unblocked $s \rightsquigarrow u$ alternating path of length $i - 1$, and then the edge (u, v) , and then a simple unblocked $v \rightsquigarrow t$ alternating path of length $j - i$. So the number of such paths may be upper bounded by the product of the number of simple unblocked $s \rightsquigarrow u$ alternating paths of length $i - 1$ and the number of simple unblocked alternating $v \rightsquigarrow t$ paths of length $j - i$.

We have that i is always even, since the heaviest edge in any alternating path is even-numbered. So $i - 1$ is odd, and using the strong inductive hypothesis, the number of $s \rightsquigarrow u$ simple unblocked alternating paths of length $i - 1$ is $O(k^2 f)^{(i-2)/2}$. The parity $j - i$ depends on the parity of j , so we need to split into two cases:

- if j is even, and so $j - i$ is even, by the strong inductive hypothesis the number of simple unblocked alternating $v \rightsquigarrow t$ paths is $O(k^2 f)^{(j-i)/2}$. In this case the product is bounded by

$$O(k^2 f)^{(i-2)/2} \cdot O(k^2 f)^{(j-i)/2} = O(k^2 f)^{(j-2)/2}.$$

Since there are $O(k^2 f)$ total equivalence classes, the total number of simple unblocked alternating $s \rightsquigarrow t$ j -paths is

$$O(k^2 f) \cdot O(k^2 f)^{(j-2)/2} = O(k^2 f)^{j/2},$$

as claimed.

- if j is odd, and so $j - i$ is odd, by the strong inductive hypothesis the number of simple unblocked alternating $v \rightsquigarrow t$ paths is $O(k^2 f)^{(j-i-1)/2}$. In this case the product is bounded by

$$O(k^2 f)^{(i-2)/2} \cdot O(k^2 f)^{(j-i-1)/2} = O(k^2 f)^{(j-3)/2}.$$

Since there are $O(k^2 f)$ total equivalence classes, the total number of simple unblocked alternating $s \rightsquigarrow t$ j -paths is

$$O(k^2 f) \cdot O(k^2 f)^{(j-3)/2} = O(k^2 f)^{(j-1)/2},$$

also as claimed. \square

We have stated this dispersion lemma for general j , for the sake of its inductive proof, but we will only need to use the setting $j = k$ in the following.

4.4 Putting it Together

Now we prove Theorem 4.1. Let d be the average degree of H . By our counting lemma, H has at least $n \cdot \Omega(d/k)^k$ total simple unblocked alternating k -paths. To apply our dispersion lemma, let us first consider the case where k is odd. Then there are $O(k^2 f)^{(k-1)/2}$ total simple unblocked alternating k -paths between any given pair of nodes in H , for an upper bound of $n^2 \cdot O(k^2 f)^{(k-1)/2}$ on the total number of simple unblocked alternating k -paths in the entire graph. We can put this upper and lower bound together and rearrange as follows:

$$\begin{aligned} n \cdot \Omega(d/k)^k &= n^2 \cdot O(k^2 f)^{(k-1)/2} \\ \implies d/k &= n^{1/k} O(k^2 f)^{1/2-1/(2k)} \\ \implies |E(H)| &= \frac{nd}{2} = O\left(k^2 n^{1+1/k} f^{1/2-1/(2k)}\right). \end{aligned}$$

In the case where k is even, we proceed similarly: the total number of simple unblocked alternating k -paths in the entire graph is at most $n^2 \cdot O(k^2 f)^{k/2}$, and so we put this upper bound together with the same lower bound as before to get:

$$\begin{aligned} n \cdot \Omega(d/k)^k &= n^2 \cdot O(k^2 f)^{k/2} \\ \implies d/k &= n^{1/k} O(k^2 f)^{1/2} \\ \implies |E(H)| &= \frac{nd}{2} = O\left(k^2 n^{1+1/k} f^{1/2}\right). \end{aligned}$$

The additional additive $+nfk$ term in Theorem 4.1 is needed to justify the assumption that $d = \Omega(fk)$, which was used to prove our counting lemma.

References

- [1] I. Althöfer, G. Das, D. P. Dobkin, D. Joseph, and J. Soares. On sparse spanners of weighted graphs. *Discrete & Computational Geometry*, 9:81–100, 1993.
- [2] B. Awerbuch, A. Baratz, and D. Peleg. Efficient broadcast and light-weight spanners. *Unpublished manuscript, November, 1991*.
- [3] B. Awerbuch and D. Peleg. Network synchronization with polylogarithmic overhead. In *Foundations of Computer Science, 1990. Proceedings., 31st Annual Symposium on*, pages 514–522. IEEE, 1990.
- [4] G. Baier, T. Erlebach, A. Hall, E. Köhler, H. Schilling, and M. Skutella. Length-bounded cuts and flows. In M. Bugliesi, B. Preneel, V. Sassone, and I. Wegener, editors, *Automata, Languages and Programming*, pages 679–690, Berlin, Heidelberg, 2006. Springer Berlin Heidelberg.
- [5] G. Bodwin, M. Dinitz, M. Parter, and V. V. Williams. Optimal vertex fault tolerant spanners (for fixed stretch). In A. 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.
- [6] G. Bodwin, M. Dinitz, and C. Robelle. Optimal vertex fault-tolerant spanners in polynomial time. In *Proceedings of the Thirty-Second Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2021, 2021*.
- [7] G. Bodwin and S. 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.
- [8] P. Bose, V. Dujmovic, P. Morin, and M. Smid. Robust geometric spanners. *SIAM Journal on Computing*, 42(4):1720–1736, 2013.
- [9] K. Buchin, S. Har-Peled, and D. Oláh. A spanner for the day after. *Discrete & Computational Geometry*, 64(4):1167–1191, 2020.
- [10] S. Chechik, M. Langberg, D. Peleg, and L. Roditty. Fault tolerant spanners for general graphs. *SIAM J. Comput.*, 39(7):3403–3423, 2010.
- [11] A. Czumaj and H. Zhao. Fault-tolerant geometric spanners. *Discrete & Computational Geometry*, 32(2):207–230, 2004.
- [12] M. Dinitz and R. Krauthgamer. Fault-tolerant spanners: better and simpler. In *Proceedings of the 30th Annual ACM Symposium on Principles of Distributed Computing, PODC 2011, San Jose, CA, USA, June 6-8, 2011*, pages 169–178, 2011.

[13] M. Dinitz and C. Robelle. Efficient and simple algorithms for fault-tolerant spanners. In *Proceedings of the 2020 ACM Symposium on Principles of Distributed Computing*, PODC '20, 2020.

[14] M. Elkin, Y. Emek, D. A. Spielman, and S.-H. Teng. Lower-stretch spanning trees. *SIAM Journal on Computing*, 38(2):608–628, 2008.

[15] P. Erdős. Extremal problems in graph theory. In *In Theory of Graphs and its Applications, Proc. Sympos. Smolenice*, 1964.

[16] S. Hoory. The size of bipartite graphs with a given girth. *Journal of Combinatorial Theory, Series B*, 86(2):215–220, 2002.

[17] M. Kapralov and R. Panigrahy. Spectral sparsification via random spanners. In *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference*, pages 393–398. ACM, 2012.

[18] C. Levcopoulos, G. Narasimhan, and M. 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.

[19] C. Levcopoulos, G. Narasimhan, and M. 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.

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

[21] G. Narasimhan and M. Smid. *Geometric Spanner Networks*. Cambridge University Press, 2007.

[22] D. Peleg. *Distributed computing: a locality-sensitive approach*. SIAM, 2000.

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

[24] D. Peleg and J. D. Ullman. An optimal synchronizer for the hypercube. *SIAM J. Comput.*, 18(4):740–747, 1989.

[25] D. Peleg and E. Upfal. A trade-off between space and efficiency for routing tables. *J. ACM*, 36(3):510–530, 1989.

[26] S. 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.

[27] M. Thorup and U. Zwick. Approximate distance oracles. *Journal of the ACM (JACM)*, 52(1):1–24, 2005.

[28] R. A. Wilson. A simple construction of the ree groups of type 2f4. *Journal of Algebra*, 323(5):1468–1481, 2010.

A Figures

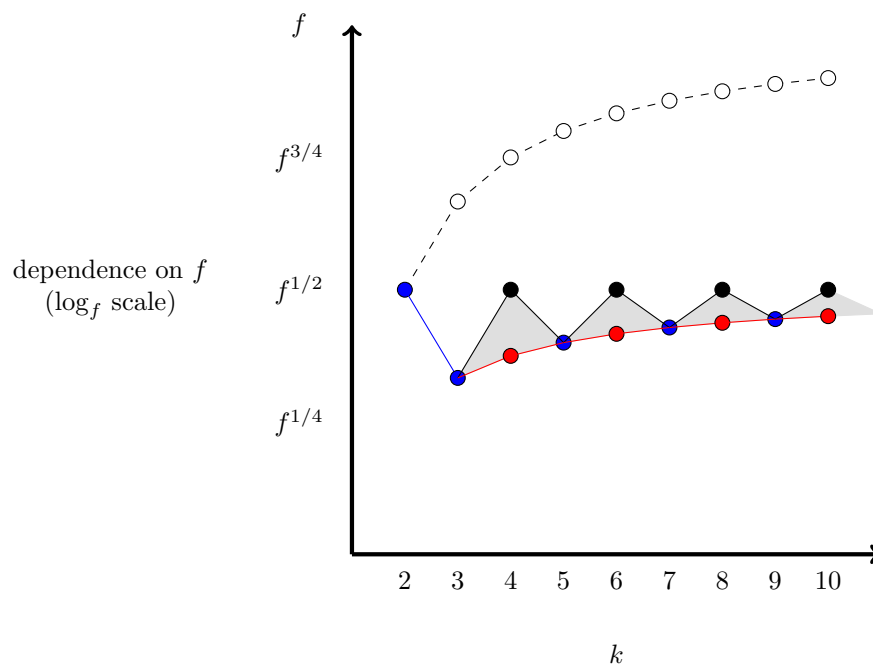


Figure 1: Dependence of our new EFT upper bounds on f : $f^{1/2}$ for even k and $f^{1/2-1/(2k)}$ for odd k . Blue points are where our f dependence is tight, and black/red points represent non-matching upper/lower bounds. The shaded regions represent knowledge gaps, and the white points are the previous best f dependence ($f^{1-1/k}$) that was known before this paper, in all of [5–7, 13].

B Unconditional Lower Bounds for EFT Spanners

In this section, we point out an extension to the lower bound of [5] for EFT spanners:

THEOREM B.1. *For $k = 7$ and $f = \Omega(n^{1/11})$, there is an infinite family of n -node graphs for which any f -EFT $(2k - 1 = 13)$ -spanner has $\Omega(f^{3/7} n^{8/7})$ edges.*

This is the same bound as in Theorem 1.3, but without the need to condition on the girth conjecture despite the fact that the girth conjecture is still unproved for the $k = 7$ case. Our construction is based on an object in incidence geometry called a Ree-Tits Octagon (see [28] for discussion and a relatively simple construction). This is a kind of incidence structure, meaning it is a collection of subsets (called “lines”) of the universe $[n]$ (called “points”). For integers $k \geq 2$, let us say that a k -gon is a circularly-ordered collection of k pairwise-distinct points $(v_0, v_1, \dots, v_k = v_0)$ such that, for any two adjacent points v_i, v_{i+1} , there exists a line that contains both points. The properties are:

THEOREM B.2. (REE-TITS OCTAGON, SEE E.G. [28]) *For infinitely many values of n , there is an incidence structure consisting of $\Theta(n^{10/11})$ lines over the points $[n]$ such that:*

- *The structure has no k -gon for any $2 \leq k \leq 7$,*
- *Every point is contained in $\Theta(n^{1/11})$ lines, and*
- *Every line contains $\Theta(n^{2/11})$ points.*

The structure has some other interesting properties as well, but these are the properties relevant to our proof. We can associate this structure to an incidence graph, which has the following properties:

COROLLARY B.1. (REE-TITS OCTAGON INCIDENCE GRAPH) *For infinitely many values of n , there is a bipartite graph with n nodes on one side, $\Theta(n^{10/11})$ nodes on the other side, $\Theta(n^{12/11})$ total edges, and girth > 15 .*

Proof. Build the incidence graph associated to the Ree-Tits Octagon. This is a bipartite graph with n nodes on the left side, corresponding to the points $[n]$, and $\Theta(n^{10/11})$ nodes on the right side corresponding to the lines. We put an edge from a point x on the left to a line ℓ on the right iff $x \in \ell$. The number of edges is immediate from the properties of the Ree-Tits Octagon, and we notice that a $2k$ -cycle in the incidence graph corresponds to a k -gon in the original incidence structure. Thus the graph has no cycles of length 4, 6, 8, 10, 12, or 14, and since it is bipartite it has no odd-length cycles either. So its girth is > 15 . \square

Now we can prove Theorem B.1; this is the method of [5], but applied to a different starting graph. Starting with a graph $G = (V, E)$ from Corollary B.1, copy each node v on the smaller side of the bipartite graph $n^{1/11} \cdot (f/n^{1/11})^{1/2}$ times, and replace each edge (u, v) with an edge (u, v_i) for every copy v_i . We similarly copy each node on the larger side of the bipartite graph $(f/n^{1/11})^{1/2}$ times in the same way. (We assume for convenience that these quantities are integers, which affects our bounds only by lower-order terms.) The number of nodes in the modified graph is

$$\Theta\left(n \cdot \left(\frac{f}{n^{1/11}}\right)^{1/2}\right) = \Theta\left(n^{21/22} \cdot f^{1/2}\right)$$

on the left side, and

$$\Theta\left(n^{10/11} \cdot n^{1/11} \cdot \left(\frac{f}{n^{1/11}}\right)^{1/2}\right) = \Theta\left(n^{21/22} \cdot f^{1/2}\right)$$

on the right side as well. The number of edges in the modified graph is

$$|E(G)| = \Theta\left(n^{12/11} \cdot n^{1/11} \cdot \frac{f}{n^{1/11}}\right) = \Theta\left(f \cdot n^{12/11}\right).$$

Writing $N := n^{21/22} \cdot f^{1/2}$, so that there are $\Theta(N)$ total nodes, we can then reparametrize the number of edges as

$$|E(G)| = \Theta\left(f^{3/7} \cdot N^{8/7}\right).$$

Notice that this is exactly our claimed lower bound. Next, we argue that G is the only f -EFT spanner of itself, and thus no edges can be removed. To see this, let us consider an edge (u_i, v_j) , which is a copy of an edge (u, v) from the original graph. The total number of copies of (u, v) is the product of the number of copies of u by the number of copies of v , which is

$$n^{1/11} \cdot \left(\frac{f}{n^{1/11}}\right)^{1/2} \cdot \left(\frac{f}{n^{1/11}}\right)^{1/2} = f.$$

So we may let F be all other copies of (u, v) , besides the edge (u_i, v_j) under consideration, and we have $|F| \leq f$. With these definitions of G, F , we have:

LEMMA B.1. *In $G \setminus F$, there is no cycle C of length ≤ 15 that includes the edge (u_i, v_j) .*

Proof. Suppose for contradiction that such a cycle C exists. Let ϕ be the homomorphism from $G \setminus F$ to the original Ree-Tits octagon incidence graph R , which maps each node x_i back to the original node x from which it was copied. Let $\phi(C)$ be the nodewise image of C , which is a closed walk in R . Notice that $\phi(C)$ only contains the edge (u, v) once, since we removed all copies of (u, v) in G except for (u_i, v_j) .

We then modify $\phi(C)$ as follows: while there exists a node x that occurs at least twice in C , replace $\phi(C)$ by its $x \rightsquigarrow x$ subwalk that contains the single copy of (u, v) . Once this process halts, we have a simple circularly-ordered walk that contains the edge (u, v) only once, which must therefore be a cycle in R of length ≥ 3 and ≤ 15 . This contradicts that R has girth > 15 , completing the proof. \square

It follows from this lemma that the shortest $u_i \rightsquigarrow v_j$ path in $G \setminus (F \cup \{(u_i, v_j)\})$ has length > 14 . Thus, we must keep (u_i, v_j) in any f -EFT 14-spanner. Since (u_i, v_j) was an arbitrary edge, this means G is the only f -EFT 14-spanner of itself, completing the proof.

C Missing Proofs from Section 3

LEMMA 3.1. *Any graph H returned by Algorithm 1 is an f -EFT $(2k - 1)$ -spanner that has a strong $2k$ -blocking set of size at most $f|E(H)|$.*

Proof. The fact that H is an f -EFT $(2k - 1)$ -spanner is straightforward from the for-loop; see [5] for a formal proof.

We now prove that the output spanner H has a small strong $2k$ -blocking set. For each edge $e = (u, v) \in E(H)$, let F_e denote the set of edges that forced the algorithm to add e ; i.e., the set with $|F_e| \leq f$ such that $d_{H \setminus F_e}(u, v) > (2k - 1) \cdot w(e)$ when e is added to H . Let

$$B := \{(x, e) \mid e \in E(H), x \in F_e\}.$$

Since $|F_e| \leq f$ for all e , we have that $|B| \leq f|E(H)|$. We can now show that B is a strong $2k$ -blocking set for H . Let C be any cycle on $\leq 2k$ edges in the final graph H and let $e = (u, v)$ be the last edge in C considered by the greedy algorithm. By construction there is a $u \rightsquigarrow v$ path (through C) of total weight $\leq (2k - 1) \cdot w(e)$ when e is added to H , and so some edge $x \in C \setminus \{(u, v)\}$ must be included in F_e . Thus $(x, e) \in B$. Moreover, since e is the last edge in C added to H by the algorithm it is the heaviest edge in C (under tiebreaking), as the algorithm considers edges in non-decreasing weight order. Hence B is a strong $2k$ -blocking set. \square

LEMMA 3.2. *The graph H returned by Algorithm 2 is an f -EFT $(2k - 1)$ -spanner that has a strong $2k$ -blocking set of size at most $O(kf) \cdot |E(H)|$.*

Proof. The fact that H is an f -EFT $(2k - 1)$ -spanner is again straightforward, and was proved formally in [13].

We now prove that the output spanner H has a small strong $2k$ -blocking set. For each edge $e = (u, v) \in E(H)$, let $H_e \subseteq H$ denote the subgraph of H containing only the edges considered before e by the algorithm. Since e was added to H by Algorithm 2, \mathcal{A} must have returned YES. Thus there is some set $F_e \subseteq E \setminus (u, v)$ with $|F_e| \leq f(2k - 1)$ such that $\hat{d}_{H_e \setminus F_e}(u, v) > 2k - 1$ (where we use \hat{d} to denote the *unweighted* distance). Let

$$B := \{(x, e) \mid e \in E(H), x \in F_e\}.$$

Since $|F_e| \leq f(2k - 1)$ we have that $|B| \leq f(2k - 1)|E(H)|$. We can now show that B is a strong $2k$ -blocking set for H . Let C be any cycle in H on $\leq 2k$ edges, and let $e = (u, v)$ be the last edge in this cycle added to H . By

construction there is a $u \rightsquigarrow v$ path (through C) with at most $2k - 1$ hops when e is added to H , and so some edge $x \in C \setminus \{(u, v)\}$ must be included in F_e . Thus $(x, e) \in B$. Moreover, since the algorithm considers the edges in nondecreasing weight order, we know that e is the heaviest edge in C (under tiebreaking). Thus B is a strong $2k$ -blocking set for H . \square