

Optimal Vertex Fault Tolerant Spanners (for fixed stretch)

Greg Bodwin¹, Michael Dinitz², Merav Parter¹, and Virginia Vassilevska Williams¹

¹MIT CSAIL

²Johns Hopkins University

Abstract

A k -spanner of a graph G is a sparse subgraph H whose shortest path distances match those of G up to a multiplicative error k . In this paper we study spanners that are resistant to faults. A subgraph $H \subseteq G$ is an f vertex fault tolerant (VFT) k -spanner if $H \setminus F$ is a k -spanner of $G \setminus F$ for any small set F of f vertices that might “fail.” One of the main questions in the area is: what is the minimum size of an f fault tolerant k -spanner that holds for all n node graphs (as a function of f , k and n)? This question was first studied in the context of geometric graphs [Levcopoulos et al. STOC '98, Czumaj and Zhao SoCG '03] and has more recently been considered in general undirected graphs [Chechik et al. STOC '09, Dinitz and Krauthgamer PODC '11].

In this paper, we settle the question of the optimal size of a VFT spanner, in the setting where the stretch factor k is fixed. Specifically, we prove that every (undirected, possibly weighted) n -node graph G has a $(2k-1)$ -spanner resilient to f vertex faults with $O_k(f^{1-1/k}n^{1+1/k})$ edges, and this is fully optimal (unless the famous Erdős Girth Conjecture is false). Our lower bound even generalizes to imply that no *data structure* capable of approximating $\text{dist}_{G \setminus F}(s, t)$ similarly can beat the space usage of our spanner in the worst case. We note that all previous upper bounds carried an almost *quadratic* (or worse) dependence on f , whereas the dependence in our tight bound is *sublinear*. To the best of our knowledge, this is the first instance in fault tolerant network design in which introducing fault tolerance to the structure increases the size of the (non-FT) structure by a sublinear factor in f . Another advantage of this result is that our spanners are constructed by a very natural and simple greedy algorithm, which is the obvious extension of the standard greedy algorithm used to build spanners in the non-faulty setting.

We also consider the *edge fault tolerant (EFT)* model, defined analogously with edge failures rather than vertex failures. We show that the same spanner upper bound applies in this setting. Our data structure lower bound extends to the case $k = 2$ (and hence we close the EFT problem for 3-approximations), but it falls to $\Omega(f^{1/2-1/(2k)} \cdot n^{1+1/k})$ for $k \geq 3$. We leave it as an open problem to close this gap.

1 Introduction

A *spanner* ([44, 45]) of a graph is a subgraph that approximately preserves its shortest path metric. More formally, a subgraph $H = (V, E' \subseteq E)$ is a t -spanner of a graph $G = (V, E)$ if

$$\text{dist}_H(u, v) \leq t \cdot \text{dist}_G(u, v) \quad \text{for all } u, v \in V$$

(t is called the *stretch* of the spanner). Spanners were introduced by Peleg and Ullman [45] and Peleg and Schäffer [44], and have a wide range of applications in routing [46], synchronizers [5], broadcasting [4, 42], distance oracles [48], graph sparsifiers [31], and even preconditioning of linear systems [27]. 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. [3] proved that for any integer $k \geq 1$, every graph $G = (V, E)$ has a $(2k - 1)$ -spanner $H \subseteq G$ with $O(n^{1+1/k})$ edges, and moreover, there exist graphs for which this size-stretch tradeoff cannot be improved (if we assume the girth conjecture of Erdős [28]). In fact, their existentially optimal upper bound was obtained via an extremely simple and natural greedy construction algorithm: 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 $\text{dist}_H(u, v) > (2k - 1)w(u, v)$. It is easy to verify that this algorithm never creates cycles of length $2k$ or less in H , and simple folklore upper bounds imply that any graph of girth $> 2k$ has $O(n^{1+1/k})$ 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 of what remains of the original graph. This motivates the notion of *fault-tolerance* for spanners. The canonical model was first introduced by Levcopoulos, Narasimhan, and Smid [33] (who happened to work in the geometric setting): a subgraph H is an f vertex (edge) fault tolerant t -spanner for G if

$$\text{dist}_{H \setminus F}(u, v) \leq t \cdot \text{dist}_{G \setminus F}(u, v) \quad \text{for every } u, v \in V \text{ and } F \subseteq V (F \subseteq E), |F| \leq f.$$

In other words, a fault tolerant spanner H contains a spanner for $G \setminus F$ for every set F of f nodes/edges that could fail.

The question of whether it is possible to construct a sparse fault tolerant spanner for an arbitrary undirected weighted graph (rather than a geometric graph) was raised by Czumaj and Zhao [22]. This was answered in the affirmative by Chechik, Langberg, Peleg and Roditty [21] who gave the first results on fault-tolerant spanners for general graphs. They presented constructions of an f -vertex fault tolerant $(2k - 1)$ -spanner of an n -node graph G of size $O(f^2 k^{f+1} \cdot n^{1+1/k} \log^{1-1/k} n)$, and an f -edge fault tolerant $(2k - 1)$ -spanner of G of size $O(f \cdot n^{1+1/k})$. Hence, Chechik et al. [21] 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 $f^2 k^{f+1}$ in the size (compared to the size of a non-fault tolerant spanner of the same stretch). Dinitz and Krauthgamer [24] later improved the edge bound for vertex faults to $O\left(f^{2-\frac{1}{k}} n^{1+\frac{1}{k}} \log n\right)$.

While this prior work provided a significant lead on the problem, it left behind two interesting knowledge gaps.

1. It is open whether these dependencies are “right;” no nontrivial lower bounds for the problem are yet published, and so improvements to these upper bounds are conceivable. For example, it is natural to wonder: *is it possible to pay only an extra factor of f and still achieve vertex fault-tolerance (as is possible for edge failures)? Even more fundamentally: can we pay even less than f and still achieve f -fault tolerance?*

2. One of the major pros of the textbook (non-faulty) greedy spanner construction of Althöfer et al. [3] is the simplicity and obvious correctness of the algorithm. The prior work on fault tolerant spanners exhibits new and interesting techniques, but cannot reasonably be viewed as an analog or extension of the classic greedy spanner. It is thus open whether a similar degree of algorithmic simplicity can be achieved in the fault tolerant model.

In this paper, we present new upper and lower bounds for fault tolerant spanners that directly address both of these issues.

Our Contribution. We answer the above questions affirmatively. First, we show that we can construct vertex and edge-fault tolerant spanners that cost only $o(f)$ more than their non-fault tolerant counterparts.

Theorem 1 (Main Result, Upper Bound). *Let $G = (V, E, w)$ be an undirected graph with real edge weights and no negative-weight cycles. Let $k \geq 1$ be a fixed integer. For any (possibly non-constant) positive integer f , G has an f -VFT $(2k - 1)$ -spanner on $O_k(f^{1-1/k}n^{1+1/k})$ edges. The same bounds can be achieved for f -EFT spanners.*

Interestingly, the construction algorithm behind Theorem 1 (Algorithm 1) is indeed the natural generalization of Althöfer et al. [3]: we simply consider the edges of the graph in non-decreasing order of weight, and we add an edge (u, v) to the current spanner H if and only if $\text{dist}_{H \setminus F}(u, v) > (2k - 1)w(u, v)$ for any possible fault set F . Correctness of the algorithm is once again trivial. However, this time it is highly nontrivial to prove an upper bound the density of the final spanner H , and it is the main endeavor of this paper to establish this.

The O_k in the sparsity bound in Theorem 1 hides a $2^{O(k)}$ factor (which is a constant for fixed k). We leave it as an open problem to determine whether this exponential dependence on k is necessary.

To the best of our knowledge, Theorem 1 is the first construction of any fault tolerant graph structure whose size has a *sublinear* dependence in f . Perhaps the most basic fault tolerant structures are those preserving connectivity, where the connected components of $H \setminus F$ are the same as for $G \setminus F$ for any fault set F of size at most f . Even for this much simpler requirement, the existing constructions of FT connected subgraphs [39] pay a factor of at least f in the size compared to the size of the non-fault tolerant structure. In fact, a common belief in fault tolerant network design is that paying a factor f in the size of the fault tolerant structure is the best one can achieve [39]. Showing that this belief is wrong is a core contribution of this paper. We hope that it will motivate finding other graph predicates where requiring fault tolerance only costs us a sublinear factor of f compared to the size of the non-faulty structure.

For the case of vertex faults, we complement Theorem 1 with a matching lower bound:

Theorem 2 (Lower Bound, Vertex Faults). *For any positive integers k and f , assuming the Erdős Girth Conjecture [28], there exist infinite families of undirected unweighted n -node graphs on $\Omega(f^{1-1/k} \cdot n^{1+1/k})$ edges for which any f VFT $(2k - 1)$ -spanner H must contain all the edges of G .*

The Girth Conjecture is widely believed and widely used as a basis for lower bounds in spanners research. It has been confirmed for $k \in \{1, 2, 3, 5\}$ [50] (and thus our lower bounds are unconditional for these values of k), and it is open for all other values of k .

For edge fault-tolerance we can prove the same lower bound for the special case of $k = 2$, but for larger stretch values we give a weaker lower bound. We leave it as an open question to close this gap for $k \geq 3$.

Theorem 3 (Lower Bound, Edge Faults). *For any positive integers k and f , assuming the Erdős Girth Conjecture [28], there exist infinite families of undirected unweighted n -node graphs on*

$$\begin{cases} \Omega(f^{1/2}n^{3/2}) & \text{when } k = 2 \\ \Omega(f^{1/2-1/(2k)}n^{1+1/k}) & \text{when } k \geq 3 \end{cases}$$

edges for which any f EFT $(2k - 1)$ -spanner H must contain all the edges of G .

In fact, by applying some standard tricks to our new lower bound constructions, we can generalize to prove *strong incompressibility* theorems:

Theorem 4 (Strong Incompressibility, Vertex Faults). *For any positive integers $k \geq 2$ and f , assuming the Erdős Girth Conjecture [28], there is no algorithm that can process n -node graphs $G = (V, E)$ into a data structure \mathcal{D}_G on $o(f^{1-1/k}n^{1+1/k})$ bits such that \mathcal{D}_G can answer queries (in any amount of time) of the form (s, t, F) , where $s, t \in V, F \subseteq V, |F| \leq f$, with a value $\widehat{\text{dist}}_{G \setminus F}(s, t)$ satisfying*

$$\text{dist}_{G \setminus F}(s, t) \leq \widehat{\text{dist}}_{G \setminus F}(s, t) \leq (2k - 1) \cdot \text{dist}_{G \setminus F}(s, t).$$

Theorem 5 (Strong Incompressibility, Edge Faults). *For any positive integers $k \geq 2$ and f , assuming the Erdős Girth Conjecture [28], there is no algorithm that can process n -node graphs $G = (V, E)$ into a data structure \mathcal{D}_G on*

$$\begin{cases} o(f^{1/2}n^{3/2}) & \text{when } k = 2 \\ o(f^{1/2-1/(2k)}n^{1+1/k}) & \text{when } k \geq 3 \end{cases}$$

bits such that \mathcal{D}_G can answer queries (in any amount of time) of the form (s, t, F) , where $s, t \in V, F \subseteq E, |F| \leq f$, with a value $\widehat{\text{dist}}_{G \setminus F}(s, t)$ satisfying

$$\text{dist}_{G \setminus F}(s, t) \leq \widehat{\text{dist}}_{G \setminus F}(s, t) \leq (2k - 1) \cdot \text{dist}_{G \setminus F}(s, t).$$

Note that the spanner $H \subseteq G$ promised by Theorem 1 functions as a data structure \mathcal{D}_G for these purposes, and thus these bounds are essentially optimal.¹

Discussion of our techniques. As mentioned previously, the fault-tolerant spanners in Theorem 1 are constructed by the natural generalization of the greedy spanner of Althöfer et al. [3] into the fault-tolerant setting. It is easy to see that the spanner H built by the greedy algorithm is an f -fault tolerant $(2k - 1)$ -spanner, both for vertex and for edge faults. Our main technical contribution in this paper is in the analysis of the number of edges m in H . We overview this here.

In the non-faulty greedy spanner algorithm, the output spanner H is shown to have *no* cycles of length at most $2k$. In our algorithm, we focus on counting the closely related notion of closed walks of length exactly $2k$. Our output spanner H might have many closed walks of length $2k$; let this number be C . At a very high level, we attempt to give both a lower bound $L(m, n, f)$ and an upper bound $U(m, n, f)$ on C , and then derive an upper bound on m in terms of n, f by rearranging the inequality $L(m, n, f) \leq U(m, n, f)$.

First, it is not too hard to obtain a lower bound $L(m, n, f)$ on the number of closed $2k$ -walks of length up to $2k$ using an argument based on the Cauchy-Schwartz inequality: we have $C =$

¹More precisely, these data structure lower bounds are optimal up to a $\log n$ factor, since the spanner on $|E|$ edges can take up to $|E| \cdot \log n$ bits to encode.

$\Omega((m/n)^{2k})$ as long as $m \geq 100kn^{1+1/k}$, say (this lower bound actually holds for $2k$ -cycles as well, as shown by Morris and Saxton [37]; a simple argument for $k = 2$ can be found in [17]). Based on this lower bound, it turns out that the desired bound on m follows if we can upper bound the number of short cycles in H by $O(m(fm/n)^{k-1})$, as we then would get $(m/n)^{2k} \leq O(f^{k-1}m^k/n^{k-1})$ which implies $m \leq O(f^{1-1/k}n^{1+1/k})$, as desired.

Unfortunately, the only thing that our greedy VFT spanner algorithm gives us is that when an edge $\{u, v\}$ is added to H , there is a set F of f nodes such that every short enough $u \rightsquigarrow v$ path includes a node in F . A major difficulty is to go from this disjoint paths condition to upper-bounding the number of short closed walks. To overcome this difficulty, we take several steps. First we present a regularization technique that allows us to assume that the output spanner H is roughly regular. This rough regularity allows us to say that the number of i -walks starting from any given vertex is roughly D^i , where $D = \Theta(m/n)$ is the average degree of H . For the special case $k = 2$, this makes our upper bound argument quite clean (see Section 3): for any pair of nodes u, v , when the edge $\{u, v\}$ is added, the number of 3-walks from u to v is $\Theta(fD)$, since all these paths must intersect one of the f nodes $x \in F$, and (by rough regularity) the remaining node in the walk must be one of the $\Theta(D)$ possible neighbors of x . Thus, by a union bound over the edge set of H , the total number of closed 4-walks is $O(fDm)$ and the desired upper bound on m follows.

For larger k , however, a difficulty arises in extending this argument. Let us generously suppose that all the $2k - 1$ walks connecting u and v go through one of f neighbors of v . Can one still argue that the total number of $2k - 1$ -walks connecting u and v (when $\{u, v\}$ is added to H) is at most $O(m(fm/n)^{k-1})$? A naive extension of the previous argument would only give an upper bound of $O(mfD^{2k-2})$ which is quite far from what we want.

To obtain a better argument, we introduce quite a bit of machinery. For instance, instead of bounding the number of closed $2k$ -walks, we show that it is sufficient to bound the number of pairs of k -walks that meet at the same endpoints. We also introduce the notion of *blockades*, which allows us to only count certain types of walks, and allows us to push through a delicate inductive argument that finally achieves the correct upper bound of $O(m(fm/n)^{k-1})$ that leads to our main upper bound theorem.

Additional Related Work. Constructing fault tolerant spanners for geometric graphs (or graphs from “simple” metric spaces such as doubling metrics) has been further studied in [35, 34, 1, 47, 18]. The study of *robust* geometric spanners, where removing few vertices from the graph harms only a small number of other vertices, was initiated in [15].

In general graphs, the construction of sparse fault tolerant subgraphs has received a significant amount of recent attention [43, 39]. Construction of purely additive fault tolerant spanners were studied in [16, 10, 14]. Concerning *exact* distances from a single (or few) sources, [40] introduced the notion of FT-BFS structures that contain a BFS tree from s in $G \setminus \{e\}$ for every failing edge $e \in E(G)$. They showed an upper bound of $O(n^{3/2})$ edges and provided a matching lower bound graph example. FT-BFS structures avoiding 2 faults with optimal size were given in [38]. Approximate versions of FT-BFS structures (where the structure is allowed to have a stretch on the $s \times V$ distances) were studied in [32, 41, 11, 13].

A natural data structure analog of fault-tolerant subgraphs are *distance sensitivity oracles*, which are small data structure that are used to answer queries of the form: “what is the distance between s and t in G when a set of F edges fail”? There is a long line of literature on constructing distance sensitivity oracles [23, 8, 9, 30, 7, 49, 26, 12] and related structures such a fault tolerant routing schemes [29, 19] and labeling schemes [2]. Recently, [20] provided an efficient construction of distance sensitivity oracles that support $f = O(\log n / \log \log n)$ many faults with polylogarithmic

query time. In another breakthrough, [25] showed a connectivity sensitivity oracle that supports $f \in [1, n]$ vertex failures with $O(fm \log n)$ space, update time $O(g^2)$ and query time $O(g)$ where $g \leq f$ is the number of actual faults.

Turning to reachability in directed graphs, [6] showed that for any number of faults $f \geq 1$, there is a subgraph $H \subseteq G$ with $O(2^f n)$ edges that preserves the reachability from s after the failure of any f edges. They also showed that this upper bound is existentially tight.

2 Preliminaries

We begin with a formal definition of fault tolerant spanners:

Definition 1 (Fault Tolerant Spanners). *Let $G = (V, E, w)$ be an undirected weighted graph without negative weight cycles. We say that a subgraph $H \subseteq G$ over the same vertex set is an f vertex (edge) fault tolerant t -spanner of G if, for any set F of f vertices (edges), we have*

$$\text{dist}_{H \setminus F}(u, v) \leq t \cdot \text{dist}_{G \setminus F}(u, v) \quad \text{for all } u, v \in V$$

where $G \setminus F$ and $H \setminus F$ denote these graphs with the set F of vertices (edges) removed.

We abbreviate *vertex fault tolerant* and *edge fault tolerant* by VFT and EFT, respectively. We will construct our spanners in Theorem 1 using the following natural construction algorithm:

Definition 2. *In a graph H , a pair of nodes (u, v) is (t, f) vertex (edge) protected if there is no set F of f vertices (edges), with $u, v \notin F$,² such that $\text{dist}_{H \setminus F}(u, v) > t \cdot w(u, v)$.*

Algorithm 1: Construction of f VFT (EFT) $(2k - 1)$ -Spanners
<p>Input : An undirected weighted graph $G = (V, E, w)$ and positive integers f, k.</p> <p>1 Initialize $H \leftarrow (V, \emptyset)$;</p> <p>2 foreach $(u, v) \in E$ in order of ascending edge weight do</p> <p>3 if (u, v) is not currently $(2k - 1, f)$ vertex (edge) protected in H then</p> <p>4 add (u, v) to H;</p> <p>5 end</p> <p>6 end</p> <p>7 return H;</p> <p>Output: H, an f VFT (EFT) $(2k - 1)$-Spanner of G.</p>

It is essentially trivial to see that the output graph H is indeed a $(2k - 1)$ spanner of the input graph G :

Theorem 6. *The graph H returned by Algorithm 1 is an f VFT (EFT) $(2k - 1)$ -spanner of the input graph $G = (V, E)$.*

Proof. Let F be any set of f node (edge) faults. Consider any $u, v \in V$ and consider a shortest path π between u and v in $G \setminus F$. Consider any edge (x, y) on π . We have that either (x, y) is in H (and hence $H \setminus F$ since $(x, y) \in G \setminus F$), or the algorithm chose not to add it to H . If the latter event occurs, it must have been that $\text{dist}_{H \setminus F}(x, y) \leq (2k - 1) \cdot w(x, y)$ (by the protection definition, this holds for all fault sets). Thus, there is a path in $H \setminus F$ of weight at most $\sum_{(x, y) \in \pi} \text{dist}_{H \setminus F}(x, y) \leq (2k - 1) \text{dist}_{G \setminus F}(u, v)$. The theorem follows. \square

²The restriction $u, v \notin F$ is only meaningful when vertex protection is considered; in the case of edge protection there is no analogous requirement.

The vast majority of this paper is devoted to showing the upper bound $n^{1+1/k} f^{1-1/k} \cdot 2^{O(k)}$ on the density of H as stated in Theorem 1.

Our lower bounds are conditional on the standard Erdős Girth Conjecture, which we recall:

Conjecture 1 (Erdős Girth Conjecture [28]). *For any positive integer k , there exist infinite families of n -node graphs with $\Omega(n^{1+1/k})$ edges and girth $2k + 2$.*

The paper is outlined as follows. In the main body of this paper, we will highlight our matching upper and lower bounds for $k = 2$ (i.e. 3-spanners) and any f , for both edge and vertex faults. The upper bound for $k = 2$ is far simpler to prove than the upper bound for $k \geq 3$, and in some sense its supporting arguments form the base case for an inductive attack on larger k . We will informally describe the extension of our argument to larger k , but due to space constraints, the formalities are deferred to the appendix.

3 Upper Bound for 3-Spanners

First, we have:

Lemma 1. *Suppose that Theorem 1 holds for $k = 2$ for all graphs G whose corresponding output graphs H have the property that their maximum degree is at most c times their minimum degree, for some universal constant c . Then Theorem 1 holds in general.*

The proof is deferred to Appendix A, as it is somewhat long but orthogonal to the main new ideas of this paper. We shall assume in the rest of this section that H (the graph output by Algorithm 1) is “approximately regular” in the sense of Lemma 1. More specifically, throughout this paper, we let D be a number such that all nodes in H have degree $\Theta_k(D)$.

As usual, a *walk* in H is a sequence of nodes (possibly with repeats) in which every pair of adjacent nodes is connected by an edge. A walk has length i (also called an i -walk) if it contains $i + 1$ nodes. A walk is *closed* if its first and last nodes are the same. We shall also say that an edge (u, v) belongs to a walk w if the nodes u, v appear in adjacent positions of w . Let c_{2k} denote the number of closed $2k$ -walks in H .

Definition 3. *If H is the spanner produced by Algorithm 1 on an input graph $G = (V, E)$, then for an edge $(u, v) \in E$, the graph $H^{(u,v)} = (V, E')$ is defined as the subgraph of H containing exactly the edges considered before (u, v) during the greedy algorithm (not including (u, v) itself).*


Lemma 2.

$$c_4 = O(|E_H| \cdot fD)$$

Proof. We give the argument for vertex faults here; the argument for edge faults is essentially identical. We shall argue that each edge added to H during Algorithm 1 completes $O(fD)$ closed 4-walks, which then implies the lemma by a simple union bound.

If we choose to add an edge (u, v) to H , then (u, v) is not $(3, f)$ protected in $H^{(u,v)}$. Thus, there is a set F of $|F| \leq f$ nodes such that every $u \rightsquigarrow v$ 3-walk contains a node in F . Additionally, for each node $x \in F$, there are $O(D)$ $u \rightsquigarrow v$ 3-walks including the node x , since 3 of the 4 nodes on this walk are specified (u, v , and x) and the fourth node must be one of the $O(D)$ neighbors of x . Applying a union bound, the total number of $u \rightsquigarrow v$ 3-walks in $H^{(u,v)}$ is $O(fD)$. See Figure 1 for a picture associated with this argument.

For each such $u \rightsquigarrow v$ 3-walk in $H^{(u,v)}$, we complete at most 8 closed 4-walks by adding (u, v) to $H^{(u,v)}$ (specifically, each such 3-walk corresponds to a single “circular 4-walk;” we may then choose any of the 4 nodes on this walk to serve as the start/end node of the appropriate closed walk, and it may be travelled in either of two directions). The lemma follows. \square



First layer $O(fD)$ walks
size is from u to
 $|F| = f$ second layer

Figure 1: In Lemma 2, we show that each edge (u, v) added to H completes at most $O(fD)$ $u \rightsquigarrow v$ 3-walks. In this picture we have drawn the separating node set F as coinciding with the first layer of the graph.

Lemma 3.

$$|E_H| = O\left(nc_4^{1/4}\right)$$

Proof. The total number of 2-walks in H is $\Theta(nD^2)$, since we have n choices for start node and $\Theta(D)$ neighbors for each node. Each ordered pair of (possibly identical) 2-walks (w_1, w_2) in H with the same start and end nodes corresponds uniquely to a closed 4-walk, obtained by walking w_1 and then w_2 in reverse. Thus, denoting by $k_{u,v}$ the number of $u \rightsquigarrow v$ 2-walks, we may calculate

$$\begin{aligned} c_4 &= \sum_{(u,v) \in V \times V} k_{u,v}^2 \\ &\geq \frac{\left(\sum_{(u,v) \in V \times V} k_{u,v}\right)^2}{n^2} && \text{Cauchy-Schwarz Inequality} \\ &= \frac{\Theta(n^2 D^4)}{n^2} = \Theta(D^4) \end{aligned}$$

and thus $nc_4^{1/4} \geq \Theta(nD) = |E_H|$. □

We are now ready to prove our upper bound for $k = 2$.

Proof of Theorem 1 for $k = 2$. Combining Lemmas 2 and 3, we compute

$$\begin{aligned} |E_H| &= O\left(n(|E_H| \cdot fD)^{1/4}\right) \\ |E_H|^3 &= O(n^4 fD) \\ |E_H|^2 &= O(n^3 f) && \text{since } |E_H| = \Theta(nD) \\ |E_H| &= O\left(n^{3/2} f^{1/2}\right). && \square \end{aligned}$$

4 Lower Bound for 3-Spanners

We prove Theorems 2 and 3 in the special case $k = 2$. Incompressibility arguments for Theorems 4 and 5 can be found in Appendix C. Our lower bounds for $k = 2$ are unconditional, since the girth conjecture has been proved in this special case:

Lemma 4 (e.g. Wenger [50]). *For all n , there exist n -node graphs on $\Omega(n^{3/2})$ edges without cycles of length 4 or less.*

We argue first for vertex faults:

Proof of Theorem 2, case $k = 2$. Start with a graph G from Lemma 4. We construct a new graph G' as follows. Let $t = \lceil f/2 \rceil$. We set $V_{G'} = V(G) \times [t]$, and let $E_{G'} = \{(u, i), (v, j) : \{u, v\} \in E(G) \wedge i, j \in [t]\}$. Let $G' = (V_{G'}, E_{G'})$. Intuitively, we can think of G' as being obtained by replacing each vertex of G by a set of t copies of the vertex, and each edge of G is replaced by a complete bipartite graph between the two sets of copies. We will prove Theorem 2 by proving that the only f VFT 3-spanner of G' is itself, and that it has the required number of edges.

We first claim that the only f VFT 3-spanner of G' is G' itself. To see this, suppose that H is a subgraph of G' which does not contain some edge $\{(u, i), (v, j)\} \in E_{G'}$. Let $F = \{(u, \ell) : \ell \neq i\} \cup \{(v, \ell) : \ell \neq j\}$, i.e., we let the fault set be all copies of u except for (u, i) and all copies of v except (v, j) . Note that $|F| \leq f$. Now consider the shortest path from (u, i) to (v, j) in $H \setminus F$. Let this path be $(u, i) = (x^0, i^0), (x^1, i^1), (x^2, i^2), \dots, (x^p, i^p) = (v, j)$. Note that for any $0 \leq a \leq p-1$, it cannot be the case that $x^a = u$ and $x^{a+1} = v$, since no such edges exist in $H \setminus F$. Thus $u = x^0, x^1, \dots, x^p = v$ is a walk from u to v in G which does not use the edge $\{u, v\}$. By adding $\{u, v\}$ to this path, we get a cycle of length at most $p+1$. Since G has girth at least 5, this implies that $p \geq 4$. Thus in $H \setminus F$ the distance between (u, i) and (v, j) is at least 4, while in $G' \setminus F$ they are at distance 1. So H is not an f VFT 3-spanner of G' , and hence the only f VFT 3-spanner of G' is G' itself.

So G' is the only f VFT 3-spanner of itself, and it remains only to analyze its size. Clearly $|V_{G'}| = t|V(G)|$, and by Lemma 4 we know that $|E(G)| \geq \Omega(|V(G)|^{3/2})$. Thus

$$\begin{aligned} |E_{G'}| &= \Omega(|E(G)|t^2) = \Omega(t^2|V(G)|^{3/2}) = \Omega\left(t^2 \cdot \left(\frac{|V_{G'}|}{t}\right)^{3/2}\right) \\ &= \Omega\left(f^{1/2}|V_{G'}|^{3/2}\right) \end{aligned}$$

as claimed. □

And next for edge faults:

Proof of Theorem 3, case $k = 2$. We construct a new graph G' exactly as in the VFT case. As before, let H be a subgraph of G' missing some edge $\{(u, i), (v, j)\}$. Let $F = \{(u, i), (v, \ell) : \ell \in [t] \setminus \{j\}\} \cup \{(u, \ell), (v, j) : \ell \in [t] \setminus \{i\}\}$ be the fault set, and note that $|F| \leq f$. In other words, we fail every edge from (u, i) to copies of v except for the edge to (v, j) , and similarly we fail all edges from (v, j) to copies of u except for (u, i) . So in $G' \setminus F$ the edge $\{(u, i), (v, j)\}$ is still present, but it is not in $H \setminus F$. Moreover, in $H \setminus F$ there are no paths of length at most 3 from (u, i) to (v, j) that use as an intermediate node any other copy of u or copy of v . Hence any path of length at most 3 must be of the form $(u, i), (x, a), (y, b), (v, j)$ where u, x, y, v are all distinct (possibly with either x or y missing). This implies that u, x, y, v form either a 3- or a 4-cycle in G , which contradicts Lemma 4. Thus $H \setminus F$ is not a 3-spanner of $G' \setminus F$, so H is not a f EFT 3-spanner of G' .

Thus G' is the only f EFT 3-spanner of itself. Using the same analysis as in the VFT case, we get that $|E(G')| \geq \Omega(f^{1/2}|V(G')|^{3/2})$, proving the theorem. □

5 Overview: Upper Bounds for Larger k

For simplicity, we will focus on the case $k = 3$ in this overview. The most natural attempt to extend our upper bounds to $k = 3$ goes as follows: first one generalizes Lemma 2, and then one generalizes Lemma 3 by considering 3-walks rather than 2-walks, and then one combines the two statements as before. Unfortunately, the upper bounds implied by this approach are quite weak. One easily generalizes Lemma 3 to show that

$$|E_H| = O\left(nc_6^{1/6}\right),$$

and one easily generalizes Lemma 2 to show

$$c_6 = O(|E_H| \cdot fD^3).$$

However, plugging these equations into the proof used before, we actually find *no improvement* in the previous upper bound: we still get

$$|E_H| = O\left(n^{3/2}f^{1/2}\right)$$

while we should hope for something much better.

Let us now informally sketch a method for improving the generalization of Lemma 2. Recall that, from Lemma 2 itself, we have $c_4 = O(|E_H| \cdot fD)$ (note that this proof holds even in the setting $k = 3$). Thus, *on average*, each node v participates in $O(fD^2)$ closed 4-walks. Additionally, we have $\Theta(D^2)$ 2-walks starting at v . Hence each of these two walks (on average) is responsible for creating at most $O(f)$ closed 4-walks in H , and so (on average) there are only $O(f)$ 2-walks between any two nodes.

Assume for a moment that we could move from an average to worst-case version of this statement, and assume a hard limit of $O(f)$ 2-walks between any two nodes. This fact could be used to improve our generalization of Lemma 2, reasoning as follows. Suppose we add the edge (u, v) to $H^{(u, v)}$ during Algorithm 1. Then, there is a set F of f nodes that lie on any $u \rightsquigarrow v$ 5-walk. Consider any such node x , and assume without loss of generality that x is in the first half of any $u \rightsquigarrow v$ 5-walk. We then count the number of $u \rightsquigarrow v$ 5-walks including x as follows: there are D^2 ways to choose the first four nodes on the walk (since u, x are two of them), and there are f ways to choose the last two nodes on the walk (since v is the last node on the walk and we have specified the fourth node on the walk, and we have assumed that there are only $O(f)$ possible 2-walks between these nodes). Hence there are $O(fD^2)$ $u \rightsquigarrow v$ 5-walks including x ; applying a union bound over F , we have $O(f^2D^2)$ $u \rightsquigarrow v$ 5-walks in total, and this leads to an improved bound

$$c_6 = O(|E_H| \cdot (fD)^2).$$

This is precisely the bound on c_6 needed to prove Theorem 1 for $k = 3$! Of course, it “only” remains to justify our assumption that there are no more than $O(f)$ 2-walks between any two nodes, while we have only proved this statement in the average case. This requires new machinery. At a high level, we accomplish this by carefully restricting our attention to a subset of the paths in H ; in particular, we “block” 2-paths that exceed the desired $O(f)$ bound, and then use something morally similar to the above proof to show that each edge (u, v) completes only $O(f^2D^2)$ “unblocked” closed 6-walks. It is a careful balancing act to define these terms in the right way while still having an analog of Lemma 3 that applies when only “unblocked” closed 6-walks are considered. However, with enough precision, it can be done.

Some additional technical work not mentioned here is required to avoid a dependence of the form e.g. $O(\text{poly}(k))$ in our upper bound. We defer a discussion of this point to the appendix.




Figure 2: Assuming a maximum of $O(f)$ 2-paths between any two nodes, each edge (u, v) added to H completes at most $O(f^2D^2)$ $u \rightsquigarrow v$ 3-walks. In this picture we have again drawn the separating node set F as coinciding with the first layer of the graph.

6 Overview: Lower Bounds for Larger k

Details of these proofs can be found in Appendix C. The proof of Theorem 2 in its full generality is essentially the obvious generalization of the $k = 2$ case. We simply have to start with a graph G which arises from the girth conjecture rather than from Lemma 4. We then use the same construction (making $f/2$ copies of each node, and turning each edge into a complete bipartite graph between the associated copies). For any edge $\{(u, i), (v, j)\}$ in this graph, if the spanner does not include this edge and if we fail all other copies of u and v , the shortest path from (u, i) to (v, j) in the spanner must correspond to a walk from u to v in G which does not use the edge $\{u, v\}$. But by construction G has large girth, so all such walks are too long. Hence $\{(u, i), (v, j)\}$ must be in the spanner.

In the case of edge faults (Theorem 3), when $k > 2$ the same construction unfortunately no longer works. To see why, consider as before some edge $\{(u, i), (v, j)\}$ which is in the graph but not the spanner. If we use the same fault set F as in the $k = 2$ case, there will still be a path in $G \setminus F$ of length 5 from (u, i) to (v, j) of the form $(u, i), (x, a), (u, b), (v, c), (y, d), (v, j)$. Here x is any neighbor of u in G and y is any neighbor of v and a, b, c, d are arbitrary indices. This will be the case even if the original graph G has large girth.

To get around this problem, we make even fewer copies of each vertex. In particular, we will make only $t = \lfloor \sqrt{f} \rfloor$ copies of each vertex. Now, instead of our fault set being all edges from (u, i) to copies of v and all edges from (v, j) to copies of u , we can simply remove all edges between copies of u and copies of v except for $\{(u, i), (v, j)\}$. Now the same logic as before implies that this edge must be in the spanner, but when we analyze the size of the graph we get only $\Omega(f^{1/2-1/(2k)}|V_{G'}|^{1+1/k})$.

References

- [1] Mohammad Ali Abam, Mark De Berg, Mohammad Farshi, and Joachim Gudmundsson. Region-fault tolerant geometric spanners. *Discrete & Computational Geometry*, 41(4):556–582, 2009.

- [2] Ittai Abraham, Shiri Chechik, Cyril Gavoille, and David Peleg. Forbidden-set distance labels for graphs of bounded doubling dimension. *ACM Transactions on Algorithms (TALG)*, 12(2):22, 2016.
- [3] 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.
- [4] Baruch Awerbuch, Alan Baratz, and David Peleg. Efficient broadcast and light-weight spanners. *Unpublished manuscript, November, 1991*.
- [5] Baruch Awerbuch and David Peleg. Network synchronization with polylogarithmic overhead. In *Foundations of Computer Science, 1990. Proceedings., 31st Annual Symposium on*, pages 514–522. IEEE, 1990.
- [6] Surender Baswana, Keerti Choudhary, and Liam Roditty. Fault tolerant subgraph for single source reachability: generic and optimal. In *Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2016, Cambridge, MA, USA, June 18-21, 2016*, pages 509–518, 2016.
- [7] Surender Baswana, Utkarsh Lath, and Anuradha S Mehta. Single source distance oracle for planar digraphs avoiding a failed node or link. In *Proceedings of the twenty-third annual ACM-SIAM symposium on Discrete Algorithms*, pages 223–232. SIAM, 2012.
- [8] Aaron Bernstein and David Karger. Improved distance sensitivity oracles via random sampling. In *SODA '08: Proceedings of the nineteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 34–43, Philadelphia, PA, USA, 2008. Society for Industrial and Applied Mathematics.
- [9] Aaron Bernstein and David Karger. A nearly optimal oracle for avoiding failed vertices and edges. In *Proceedings of the forty-first annual ACM symposium on Theory of computing*, pages 101–110. ACM, 2009.
- [10] Davide Bilò, Fabrizio Grandoni, Luciano Gualà, Stefano Leucci, and Guido Proietti. Improved purely additive fault-tolerant spanners. In *Algorithms - ESA 2015 - 23rd Annual European Symposium, Patras, Greece, September 14-16, 2015, Proceedings*, pages 167–178, 2015.
- [11] Davide Bilò, Luciano Gualà, Stefano Leucci, and Guido Proietti. Fault-tolerant approximate shortest-path trees. In *Algorithms - ESA 2014 - 22th Annual European Symposium, Wroclaw, Poland, September 8-10, 2014. Proceedings*, pages 137–148, 2014.
- [12] Davide Bilò, Luciano Gualà, Stefano Leucci, and Guido Proietti. Compact and fast sensitivity oracles for single-source distances. In *24th Annual European Symposium on Algorithms, ESA 2016, August 22-24, 2016, Aarhus, Denmark*, pages 13:1–13:14, 2016.
- [13] Davide Bilò, Luciano Gualà, Stefano Leucci, and Guido Proietti. Multiple-edge-fault-tolerant approximate shortest-path trees. In *33rd Symposium on Theoretical Aspects of Computer Science, STACS 2016, February 17-20, 2016, Orléans, France*, pages 18:1–18:14, 2016.
- [14] G. Bodwin, F. Grandoni, M. Parter, and V. Vassilevska Williams. Preserving distances in very faulty graphs. In *Proc. of 44th International Colloquium on Automata, Languages, and Programming (ICALP)*, to appear, 2017.

- [15] Prosenjit Bose, Vida Dujmovic, Pat Morin, and Michiel Smid. Robust geometric spanners. *SIAM Journal on Computing*, 42(4):1720–1736, 2013.
- [16] Gilad Braunschvig, Shiri Chechik, David Peleg, and Adam Sealton. Fault tolerant additive and (μ, α) -spanners. *Theor. Comput. Sci.*, 580:94–100, 2015.
- [17] Karl Bringmann, Fabrizio Grandoni, Barna Saha, and Virginia Vassilevska Williams. Truly sub-cubic algorithms for language edit distance and rna-folding via fast bounded-difference min-plus product. In *IEEE 57th Annual Symposium on Foundations of Computer Science, FOCS 2016, 9-11 October 2016, Hyatt Regency, New Brunswick, New Jersey, USA*, pages 375–384, 2016.
- [18] T-H Hubert Chan, Mingfei Li, and Li Ning. Sparse fault-tolerant spanners for doubling metrics with bounded hop-diameter or degree. *Algorithmica*, 71(1):53–65, 2015.
- [19] Shiri Chechik. Fault-tolerant compact routing schemes for general graphs. *Automata, Languages and Programming*, pages 101–112, 2011.
- [20] Shiri Chechik, Sarel Cohen, Amos Fiat, and Haim Kaplan. $1 + \epsilon$ -approximate f -sensitive distance oracles. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2017, Barcelona, Spain, Hotel Porta Fira, January 16-19*, pages 1479–1496, 2017.
- [21] Shiri Chechik, Michael Langberg, David Peleg, and Liam Roditty. Fault tolerant spanners for general graphs. *SIAM J. Comput.*, 39(7):3403–3423, 2010.
- [22] Artur Czumaj and Hairong Zhao. Fault-tolerant geometric spanners. *Discrete & Computational Geometry*, 32(2):207–230, 2004.
- [23] Camil Demetrescu, Mikkel Thorup, Rezaul Alam Chowdhury, and Vijaya Ramachandran. Oracles for distances avoiding a failed node or link. *SIAM Journal on Computing*, 37(5):1299–1318, 2008.
- [24] Michael Dinitz and Robert 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.
- [25] Ran Duan and Seth Pettie. Connectivity oracles for graphs subject to vertex failures. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 490–509. SIAM, 2017.
- [26] Ran Duan and Tianyi Zhang. Improved distance sensitivity oracles via tree partitioning. *arXiv preprint arXiv:1605.04491*, 2016.
- [27] Michael Elkin, Yuval Emek, Daniel A Spielman, and Shang-Hua Teng. Lower-stretch spanning trees. *SIAM Journal on Computing*, 38(2):608–628, 2008.
- [28] Paul Erdős. Extremal problems in graph theory. In *IN “THEORY OF GRAPHS AND ITS APPLICATIONS,” PROC. SYMPOS. SMOLENICE*. Citeseer, 1964.
- [29] Cyril Gavoille and Andrew Twigg. Compact forbidden-set routing on planar graphs. *Unpublished as yet*, 2008.

- [30] Fabrizio Grandoni and Virginia Vassilevska Williams. Improved distance sensitivity oracles via fast single-source replacement paths. In *Foundations of Computer Science (FOCS), 2012 IEEE 53rd Annual Symposium on*, pages 748–757. IEEE, 2012.
- [31] Michael Kapralov and Rina Panigrahy. Spectral sparsification via random spanners. In *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference*, pages 393–398. ACM, 2012.
- [32] Neelesh Khanna and Surender Baswana. Approximate shortest paths avoiding a failed vertex: Optimal size data structures for unweighted graph. In *27th International Symposium on Theoretical Aspects of Computer Science-STACS 2010*, pages 513–524, 2010.
- [33] 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.
- [34] Christos Levcopoulos, Giri Narasimhan, and Michiel Smid. Improved algorithms for constructing fault-tolerant spanners. *Algorithmica*, 32(1):144–156, 2002.
- [35] Tamas Lukovszki. New results on fault tolerant geometric spanners. *Algorithms and Data Structures*, pages 774–774, 1999.
- [36] Jiří Matoušek. On the distortion required for embedding finite metric spaces into normed spaces. *Israel Journal of Mathematics*, 93(1):333–344, 1996.
- [37] R. Morris and D. Saxton. The number of C_{2k} -free graphs. *Adv. Math.*, 298:534–580, 2016.
- [38] Merav Parter. Dual failure resilient BFS structure. In *Proceedings of the 2015 ACM Symposium on Principles of Distributed Computing, PODC 2015, Donostia-San Sebastián, Spain, July 21 - 23, 2015*, pages 481–490, 2015.
- [39] Merav Parter. Fault-tolerant logical network structures. In *The Distributed Computing Column*. Stefan Schmid, 2016.
- [40] Merav Parter and David Peleg. Sparse fault-tolerant BFS trees. In *Algorithms - ESA 2013 - 21st Annual European Symposium, Sophia Antipolis, France, September 2-4, 2013. Proceedings*, pages 779–790, 2013.
- [41] Merav Parter and David Peleg. Fault tolerant approximate BFS structures. In *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2014, Portland, Oregon, USA, January 5-7, 2014*, pages 1073–1092, 2014.
- [42] David Peleg. *Distributed computing: a locality-sensitive approach*. SIAM, 2000.
- [43] David Peleg. As good as it gets: Competitive fault tolerance in network structures. In *Symposium on Self-Stabilizing Systems*, pages 35–46. Springer, 2009.
- [44] David Peleg and Alejandro A. Schäffer. Graph spanners. *Journal of Graph Theory*, 13(1):99–116, 1989.
- [45] David Peleg and Jeffrey D. Ullman. An optimal synchronizer for the hypercube. *SIAM J. Comput.*, 18(4):740–747, 1989.

- [46] David Peleg and Eli Upfal. A trade-off between space and efficiency for routing tables. *J. ACM*, 36(3):510–530, 1989.
- [47] Shay Solomon. From hierarchical partitions to hierarchical covers: Optimal fault-tolerant spanners for doubling metrics. In *Proceedings of the 46th Annual ACM Symposium on Theory of Computing*, pages 363–372. ACM, 2014.
- [48] Mikkel Thorup and Uri Zwick. Approximate distance oracles. *Journal of the ACM (JACM)*, 52(1):1–24, 2005.
- [49] Oren Weimann and Raphael Yuster. Replacement paths and distance sensitivity oracles via fast matrix multiplication. *ACM Transactions on Algorithms*, 9(2):14, 2013.
- [50] Rephael Wenger. Extremal graphs with no c_4 's, c_6 's, or c_{10} 's. *Journal of Combinatorial Theory, Series B*, 52(1):113–116, 1991.

A Regularizing H

Before diving into our main analysis, it will be convenient as before to assume that H is “approximately regular.” We justify this assumption by the following argument.

Lemma 5. *Suppose that Theorem 1 holds for all graphs G whose corresponding output graphs H have the property that their maximum degree is at most c times their minimum degree, for some universal constant c . Then Theorem 1 holds in general.*

Proof. We shall prove this lemma by showing the contrapositive. We suppose that $G = (V, E)$ is an n -node graph that serves as a counterexample to Theorem 1 (i.e. its corresponding output graph H has $n^{1+1/k} f^{1-1/k} \omega(1)^k$ edges). Our goal is then to show that there is a subgraph $G' \subseteq G$ on at least $n^{1/(2k)}$ nodes that also serves as a counterexample to Theorem 1, and the maximum degree of G' is larger than its minimum degree by a factor of $O_k(1)$. Hence, in the proof of Theorem 1 we may assume this approximate regularity condition. We will then rule out the existence of a family of G' , and this rules out the general existence of G . Note that the constraint that G' has at least $n^{1/(2k)}$ nodes is used to ensure that an infinite family of G implies an infinite family of G' ; any super-constant function would work equally well for our application of this lemma.

First, we may assume that $G = H$; if our algorithm discards any edge in G , we may simply delete this edge from G and it is clear that this deletion will not change the corresponding output graph H . Let n be the number of nodes in $G = H$ and let D be its average degree. Partition the nodes of G into two sets: the set A of nodes that have degree cD or less, and the set B of nodes that have degree more than cD (for some constant $c = c_k$ that will be chosen later). By Markov’s inequality we have $|B| \leq n/c$. Partition the edges of G into three sets: E_A where both endpoints are in A , E_B where both endpoints are in B , and E_{AB} that has one endpoint in each. Compare the sizes of E_A, E_B, E_{AB} ; keep all edges in the largest set and discard all edges in the other two (ties may be broken arbitrarily). We now split into cases based on which set survived.

(a) E_A Survived. We now have a maximum degree of cD in the remaining graph, while its average degree is still $\Omega(D)$ (since we have only discarded a constant fraction of the edges). Let D' be the new average degree of E_A , and repeatedly delete all nodes from G that have degree $D'/4$ or less until all nodes have degree at least $D'/4$. Note that only $n \cdot D'/4$ edges (which is at most half the remaining edges in G) will be removed in this way. The remaining subgraph G' is now sparser

than the original graph G by only a constant factor and its maximum and minimum degrees are both $\Theta(D)$. Additionally, the number of nodes n' in G' is $n' = \Omega(|E|)/\Theta(D) = \Omega(n)$. Thus G' satisfies the lemma.

(b) E_B Survived. We now obtain a subgraph G' by deleting all nodes in A (which now have degree 0). The average degree in the remaining subgraph G' is at least $cD/3$, and it has $1/c$ as many nodes as before. It is not necessarily the case that G' is approximately regular, but we may now recurse the argument on G' to enforce this property (it remains to be proved that this recursion terminates before removing too many nodes from G). Note that G' has higher average degree and fewer nodes than G , so it still serves as a counterexample graph to our claimed bound.

(c) E_{AB} Survived. The average degree of the nodes in A is now at least $D/6$, since we have deleted at most $1/3$ of the edges since setting D and (since the graph is now bipartite) each edge is incident on one node in A . Our next step is to delete all nodes from A except for the $|B|$ nodes in $|A|$ with the highest remaining degree. The average degree in the remaining subgraph G' is then still at least $D/6$. We also have $2|B| \leq \frac{2}{c} \cdot n$ nodes remaining in G' . We then once again recurse the analysis on the remaining subgraph (and we will prove shortly that this recursion terminates before removing too many nodes).

In the latter two cases, we recurse the analysis on a subgraph G' of the original graph. We will now show that, if c is chosen to be suitably large, then the recursion must eventually terminate in case (a) while the graph still has $\text{poly}(n)$ nodes (where n is the number of nodes in the original graph G). Naturally, throughout the recursion, the average degree D^* in the subgraph being considered cannot exceed the number of nodes n^* in that subgraph. However, after k rounds of recursion that avoid case (a), we have

$$D^* \geq \frac{D}{6^k} \geq \frac{n^{1/k}}{6^k}$$

and

$$n^* \leq n \cdot (2/c)^k$$

We thus have

$$\begin{aligned} \frac{n^{1/k}}{6^k} &\leq n \cdot (2/c)^k \\ (c/12)^k &\leq n^{1-1/k} \\ k \log(c/12) &\leq (1 - 1/k) \log n \leq \log n \end{aligned}$$

Thus, by choosing c sufficiently large, we have that $k \leq \frac{\log n}{c'}$ (for some constant $c' = \log(c/12)$ that can be made arbitrarily large by choice of c). This means that the recursion bottoms out at depth $\frac{\log n}{c'}$, at which point the average degree of the graph is

$$D^* \geq \frac{n^{1/k}}{3^{\log n/c'}} \geq \frac{n^{1/k}}{n^{1/c''}} = n^{1/k-1/c''}$$

where c'' is another constant that can be made arbitrarily large by pushing c' arbitrarily large. Choosing $c'' \geq 2k$ we have $D^* \geq n^{1/(2k)}$ and so the recursion must terminate while the graph still has at least $n^{1/(2k)}$ nodes, as claimed.

Now if we set $c'' = 2k$, then we are setting $c = 12 \cdot 9^k$. If the recursion bottoms out after $i \leq k$ levels in case (a), then after applying case (a), the average degree is $D^* \geq D/(2 \cdot 6^i)$

and the number of nodes is $n^* \leq n(2/c)^i$. We will explicitly show that the subgraph G^* we have obtained is a counterexample, provided G was. We know that in G , the number of edges is $Dn/2 \geq n^{1+1/k} f^{1-1/k} Q$ where $Q \geq \omega(1)$. We will show that in G^* , the number of edges is $\geq (n^*)^{1+1/k} f^{1-1/k} Q$, and hence G^* is also a counterexample.

First, notice that $n^* \leq n(2/c)^i = n/(6 \cdot 9^k)^i \leq n/(2^k \cdot 6^{ki})$. Thus, $(n^*)^{1/k} \leq n^{1/k}/(2 \cdot 6^i)$. Now, since $D \geq 2n^{1/k} f^{1-1/k} Q$ (since G is a counter example), we get that $D^* \geq D/(2 \cdot 6^i) \geq 2f^{1-1/k} Q n^{1/k}/(2 \cdot 6^i) \geq 2f^{1-1/k} (n^*)^{1/k} Q$, and hence the number of edges in G^* is $D^* n^*/2 \geq f^{1-1/k} (n^*)^{1+1/k} Q$, and so G^* is also a counterexample of nontrivial size. \square

Hence, it suffices to refute the existence of “approximately regular” graphs G' that violate Theorem 1. By contrapositive of Lemma 5, this would imply that no *general* graph may violate Theorem 1. Note that this argument implies Lemma 1 by plugging in $k = 2$.

As we proceed, we will assume that the maximum and minimum degrees of the input graph G differ by a factor of $O_k(1)$. In fact, as before we may assume that $G = H$, so we may assume this same approximate regularity property for H . Moreover, we will assume that all nodes in H have degree in the interval $[D, \psi_k D]$ for some parameter D that we carry into the following proofs. From the argument in Lemma 5, we have that $\psi_k \leq 2^{O(k)}$.

B Upper Bounds for $k \geq 3$

We now prove Theorem 1 for larger k . For convenience, we will assume that all edges in H have unique weights, and are thus considered by our algorithm in a consistent order (or it suffices that ties between equally-weighted edges are broken in some consistent fashion).

B.1 Definitions, Notation, and some Intuition

First let us introduce some new notation related to walks in H . We write \bar{w} to denote the reverse of a walk w , we write w^i to denote the i^{th} node in the walk w (indexing from 0), and for two walks w_1, w_2 such that the last node of w_1 equals the first node of w_2 , we write $w_1 w_2$ to denote their concatenation in the natural way.

As mentioned in Section 5, it is easy to prove a suitable generalization of Lemma 3 but hard to prove a suitable generalization of Lemma 2. The difficulty in proving Lemma 2 for $k > 2$ is in handling pair of nodes u, v for which the number of $u \rightsquigarrow v$ k -walks is much larger than the average number of k -walks over all nodes pairs in the graph. Our solution is to “block” these irregular parts of the graph, essentially throwing away a constant fraction of the available walks in the graph in exchange for a guarantee of “regularity” on the ones that survive. Specifically, we will shortly define a set \mathcal{B} of walks in H , where each $b \in \mathcal{B}$ is called a *blockade*. A walk w in H is *blocked* by some \mathcal{B} if there is a blockade $b \in \mathcal{B}$ such that b or \bar{b} is a sub-walk of w (we use the notation $b \subseteq w$ or $\bar{b} \subseteq w$).

After appropriately designing \mathcal{B} , it is tempting to generalize Lemma 2 by counting unblocked closed walks. A technical detail of introducing blocks is that this is no longer quite the right object to count. Let us define:

Definition 4. *An i -walk meet is an ordered pair of (possibly identical) i -walks with the same start and end points.*

The proof of Lemma 2 works by observing that the number of 2-walk meets is the same as the number of closed 4-walks, up to constant factors. However, it is not necessarily the case that the




Figure 3: In this picture, the closed 4-walk w starting and ending at x is blocked by the path $(a, b, c) \in \mathcal{B}$, but both paths in the 2-walk meet $[(x, a, b), (x, c, b)]$ associated with it are still unblocked. This explains why the number of unblocked 2-walk meets can potentially be much larger than the number of unblocked closed 4-walks.

number of *unblocked* i -walk meets is within a constant factor of the number of *unblocked* closed $2i$ -walks. Indeed, the number of i -walk meets can be much larger. This happens because there are $2i$ different i -walk meets corresponding to each closed $2i$ -walk, and one can imagine that some of these are blocked but others are unblocked. With this in mind, our proof works by counting i -walk meets directly rather than passing through any attempt to count closed walks, as before.

We shall write $W^{\mathcal{B}}, X^{\mathcal{B}}$ to denote the set of walks in H that are unblocked and blocked by \mathcal{B} , respectively. We will frequently omit the superscript when it is simply \mathcal{B} (we never refer to W in the absence of a blockade set). We also use the following modifiers on the sets \mathcal{B}, W, X : we denote by $\mathcal{B}_i, \mathcal{B}_{\leq i}$ the subset of walks in \mathcal{B} of length i or length at most i , respectively, and we denote by $\mathcal{B}[s \rightsquigarrow t]$ to denote the subset of walks in \mathcal{B} with endpoints s, t (and similar notation is used on W, X). Note that $X_i^{\mathcal{B}}$ is not to be confused with \mathcal{B}_i .

Similarly, we denote by $M_i^{\mathcal{B}}$ is the set of all unblocked i -walk meets in H , and $M_i^{\mathcal{B}}[s \rightsquigarrow t]$ is the subset of these where both walks in the meet start at s and end at t . We will also frequently suppress the superscript, although there is always an implicit \mathcal{B} .

Our goal is to establish the following inequality:

$$\Theta(D)^{2k} \leq \left| M_k^{\mathcal{B}_{\leq k-1}} \right| \leq \Theta_k \left(|E_H| \cdot (fD)^{k-1} \right)$$

which will then imply an upper bound on $|E_H|$ by some straightforward algebra. The left-hand inequality is shown in Lemmas 6 and 7. The high level idea here is that, if we only block a small enough constant fraction (dependent on k) fraction of the i -walks for all $i \leq k$, then the number of unblocked k -walk meets is still nearly as large as one would expect without blockades (using the Cauchy-Schwartz Inequality). The main technical effort goes into the right-hand inequality. This is proved using a delicate interlacing inductive argument, in which two complementary counting arguments are proved for the case $i = 2$ and then used to boost each other up to $i = k$. Specifically, we show an interesting interplay between (1) bounding the number of unblocked i -walk meets

$M_i^{\mathcal{B}^{\leq i-1}}$ and (2) choosing i -length blockades to bound the number of unblocked i -walks $W_i^{\mathcal{B}^{\leq i}}[u \rightsquigarrow v]$ between any given pair of nodes u, v (Lemma 9). The latter bound is then used to bounding the number of unblocked $i + 1$ -walk meets $M_{i+1}^{\mathcal{B}^{\leq i}}$, and so on.

B.2 Lower Bound for $|M_k|$ (Generalization of Lemma 3)

The following lemma generalizes the fact used in the $k = 2$ setting that we have $\Theta(nD^2)$ 2-walks in H to consider.

Lemma 6. *There is some $\phi_k > 0$ (independent of n, f) such that the following property holds: if \mathcal{B}_i contains at most a ϕ_k fraction of all i -walks in H for all i , then*

$$|W_i| \geq n\Theta(D)^i \quad \text{for all } i \leq k.$$

Proof. First note that we have

$$nD^i \leq |W_i \cup X_i| \leq n(\psi_k D)^i$$

total walks in H (recall that by our regularization procedure, we can assume that all vertices have degree in $[D, \psi_k D]$). We now upper bound X_i . Each $x \in X_i$ has a subpath $x \supseteq b \in \mathcal{B}$. We may thus upper bound X_i by the following (somewhat loose) estimate:

$$|X_i| \leq \sum_{j=2}^i |\mathcal{B}_j| \cdot (2\psi_k D)^{i-j},$$

since each blocked path $x \in X_i$ may be obtained via $i - |b|$ extensions of some blockade $b \in \mathcal{B}_{\leq i}$ by adding an edge to its front or back, and there are up to $2\psi_k D$ ways to extend any given walk by 1 edge on either end. Since $|\mathcal{B}_j| \leq \phi_k n(\psi_k D)^j$ we have

$$|X_i| \leq \sum_{j=2}^i n\phi_k \cdot (2)^{i-j} (\psi_k D)^i \leq n\phi_k (2\psi_k D)^i.$$

Hence, choosing $\phi_k \leq \frac{1}{(4\psi_k)^k}$ (this is overkill), we have

$$|X_i| \leq n(D/2)^i$$

And so

$$|W_i| \geq |W_i \cup X_i| - |X_i| \geq nD^i - n(D/2)^i = n\Theta(D)^i$$

and the lemma follows. \square

With this, we can prove:

Lemma 7 (Generalization of Lemma 3). *If $|W_k| = n\Theta(D)^k$, then $|E_H| = O_k \left(n|M_k|^{1/(2k)} \right)$.*

Proof. The proof is essentially identical to that of Lemma 3. We have

$$\begin{aligned} |M_k| &= \sum_{(u,v) \in V \times V} |W_k[u \rightsquigarrow v]|^2 \\ &\geq \frac{\left(\sum_{(u,v) \in V \times V} |W_k[u \rightsquigarrow v]| \right)^2}{n^2} && \text{Cauchy-Schwarz Inequality} \\ &= \frac{(n\Theta(D)^k)^2}{n^2} \\ &= \Theta(D)^{2k} \end{aligned}$$

and so

$$\begin{aligned} |M_k|^{1/(2k)} &\geq D \\ n |M_k|^{1/(2k)} &\geq nD \\ n |M_k|^{1/(2k)} &= \Omega_k(|E_H|) \end{aligned}$$

which implies the lemma. \square

B.3 Upper Bound for $|M_k|$ (Generalization of Lemma 2)

We now begin to work towards a generalization of Lemma 2. The argument will be inductive in nature. Lemma 2 will essentially serve as the base case, and the inductive step is split across the next two lemmas.

Lemma 8. *Suppose that $|W_j[u \rightsquigarrow v]| = O_k(f^{j-1})$ for all u, v and all $j < i$ for some fixed $i \leq k$. Then*

$$|M_i| = O_k(|E_H| \cdot (fD)^{i-1}).$$

Proof. For a given i -walk meet (w_1, w_2) , we will say that the *shift* s_x of a node $x \in w_1 \cup w_2$ is its first position in $w_1 \overline{w_2}$ (so $0 \leq s_x < 2i$). Similarly, the shift $s_{(u,v)}$ of an edge (u, v) is the first position of this edge in $w_1 \overline{w_2}$ (where the first edge in $w_1 \overline{w_2}$ is indexed from 0).

As before, when we choose to add any (u, v) to H , it is not $(2k - 1, f)$ protected in $H^{(u,v)}$ and so there is a set F of $|F| \leq f$ nodes (with $u, v \notin F$) such that every $u \rightsquigarrow v$ walk of length $2i - 1$ intersects some node $x \in F$. Note that each $(w_1, w_2) \in M_i$ that is completed by the addition of the edge (u, v) to H corresponds to a $u \rightsquigarrow v$ walk of length $2i - 1$ for $i \leq k$ (obtained by joining the common start/endpoint of $w_1 \overline{w_2}$ to create a *circular* walk, and then removing the edge (u, v) from this walk). Thus $x \in w_1 \cup w_2$ for some $x \in F$. Our proof strategy is: we fix a shift $s_{(u,v)}$ for the edge (u, v) , we fix a node $x \in F$ and a shift value s_x for x , and we will show that the addition of the edge (u, v) to $H^{(u,v)}$ completes only $O_k(f^{i-2} D^{i-1})$ i -walk meets including the node x in which $(u, v), x$ have their prescribed shift values. Since there are $O_k(f)$ possible choices of $x \in F, 0 \leq s_x, s_{(u,v)} \leq 2k$, this implies

$$|M_i| = O_k(|E_H| \cdot (fD)^{i-1})$$

by a simple union bound over $(u, v) \in E_H$ and over the possible choices of $x, s_x, s_{(u,v)}$.

In the rest of this proof we will assume $0 \leq s_{(u,v)} \leq i - 1$, and so $(u, v) \in w_1$ (we also assume without loss of generality that u precedes v in w_1). The proof for the remaining case $i \leq s_{(u,v)} < 2i$ is identical with the roles of w_1, w_2 swapped. We will denote by M'_i the subset of M_i satisfying the criteria mentioned thus far (i.e. with (u, v) at shift $s_{(u,v)}$ and x at shift s_x). We split into two cases, depending on the value of s_x .

(a) Suppose $0 \leq s_x \leq i + 1$. Let P be the set of $i + 1$ -walks in $H^{(u,v)}$ that include node u in position $s_{(u,v)}$, node v in position $s_{(u,v)} + 1$, and node x in position s_x (P can even include blocked walks). We next count $|P|$. Supposing $s_x > s_{(u,v)}$, after fixing the node u in position $s_{(u,v)}$, there are $(\psi_k D)^{s_{(u,v)}}$ ways to choose a prefix for the walk, and there are $(\psi_k D)^{i-1-s_{(u,v)}}$ ways to choose a suffix for the walk, since v, x have fixed positions following u . Hence $|P| = (\psi_k D)^{s_{(u,v)}} \cdot (\psi_k D)^{i-1-s_{(u,v)}} = O_k(D^{i-1})$. If instead we have $s_x < s_{(u,v)}$ then the argument and bound are identical (with a factor of $\psi_k D$ shifted from the prefix to the suffix).

For each $p = (p^0, \dots, p^{i+1}) \in P$, let Q^p be the set of suffixes such that for each $q \in Q^p$, we have $qp = w_1 \overline{w_2}$ for some $(w_1, w_2) \in M'_i$. We now count $|Q^p|$. Note that each $q \in Q^p$ has the form $q = w_1 \overline{w_2} \setminus p$, and so \overline{q} is an $i-1$ -subwalk of w_2 ; since w_2 is unblocked by definition of i -walk meets, it follows that q is unblocked as well. Thus $Q^p \subseteq W_{i-1}[p^{i+1} \rightsquigarrow p^0]$, and so

$$|Q^p| \leq |W_{i-1}[p^{i+1} \rightsquigarrow p^0]| = O_k(f^{i-2})$$

We then have

$$|M'_i| = \sum_{p \in P} |Q^p| \leq |P| \cdot O_k(f^{i-2}) = O_k(f^{i-2} D^{i-1})$$

as claimed.

(b) Instead suppose $i+2 \leq s_x < 2i$. This implies that $x \notin w_1$ so $x \in w_2$. The proof is now only a slight tweak on the previous case. We define P to be the set of i -walks (not $i+1$ walks as before) in $H^{(u,v)}$ that have u in position $s_{(u,v)}$ and v in position $s_{(u,v)} + 1$. By an identical counting argument as in the previous case, we have $|P| = O_k(D^{i-1})$. As before, for each fixed $p = (p^0, \dots, p^i) \in P$ we define Q^p as the set of suffixes such that for each $q \in Q^p$, we have $qp = w_1 \overline{w_2}$ for some $(w_1, w_2) \in M_i$, and our goal is to count $|Q^p|$. Note that we then have $\overline{q} = w_2$ and so q is unblocked. Unlike before, we have $x \in q$ so we may split each $q \in Q^p$ into two subpaths over q . Specifically, let

$$Q_1 := \left\{ q_1 \subseteq q \mid q_1^0 = p^i, q_1^{s_x-i} = x, q \in Q^p \right\}$$

and

$$Q_2 := \left\{ q_2 \subseteq q \mid q_2^0 = x, q_2^{2i-s_x} = p^0, q \in Q^p \right\}$$

(in other words, for any $q \in Q^p$, we split q over the node x and add its prefix to Q_1 and its suffix to Q_2). Since each $q \in Q^p$ is unblocked, we also have that each $q_1 \in Q_1, q_2 \in Q_2$ is unblocked. Thus $Q_1 \subseteq W_{s_x-i}[p^i \rightsquigarrow x]$ and $Q_2 \subseteq W_{2i-s_x}[x \rightsquigarrow p^0]$. We then have

$$|Q^p| \leq |Q_1| \cdot |Q_2| \leq O_k(f^{s_x-i-1}) \cdot O_k(f^{2i-s_x-1}) = O_k(f^{i-2})$$

and so, as before,

$$|M'_i| = \sum_{p \in P} |Q^p| = |P| \cdot O_k(f^{i-2}) = O_k(f^{i-2} D^{i-1})$$

as claimed. □

We next show:

Lemma 9. *Let $2 \leq i \leq k$. Suppose that*

$$\left| M_i^{\mathcal{B}_{\leq i-1}} \right| = O_k(|E_H| \cdot (fD)^{i-1}).$$

Then there exists a set \mathcal{B}_i of i -length blockades such that

$$\left| W_i^{\mathcal{B}_{\leq i-1} \cup \mathcal{B}_i}[u \rightsquigarrow v] \right| = O_k(f^{i-1}) \quad \text{for all } u, v \in V$$

and \mathcal{B}_i contains at most a ϕ_k fraction of all i -walks in H , for any desired constant $\phi_k > 0$ depending only on k .

Proof. We construct \mathcal{B}_i iteratively as follows: initially $\mathcal{B}_i = \emptyset$; we then repeatedly choose the node pair u, v that maximizes $\left|W_i^{\mathcal{B}_{\leq i-1} \cup \mathcal{B}_i}[u \rightsquigarrow v]\right|$ (ties may be broken arbitrarily), choose any walk $b \in W_i^{\mathcal{B}_{\leq i-1} \cup \mathcal{B}_i}[u \rightsquigarrow v]$, and add b to \mathcal{B}_i (note that the sets $W_i^{\mathcal{B}_{\leq i-1} \cup \mathcal{B}_i}[u \rightsquigarrow v]$ shrink throughout this process as we grow \mathcal{B}_i). Repeat until no more walks may be added to \mathcal{B}_i without destroying the property that \mathcal{B}_i contains at most a ϕ_k fraction of all i -walks in H .

We now argue that \mathcal{B}_i satisfies the lemma. Observe that $|M_i[s \rightsquigarrow t]| = |W_i[s \rightsquigarrow t]|^2$ for any $s, t \in V$. Suppose towards a contradiction that \mathcal{B}_i does *not* satisfy the lemma; that is,

$$\left|W_i^{\mathcal{B}_{\leq i-1} \cup \mathcal{B}_i}[u \rightsquigarrow v]\right| = \omega_k(f^{i-1}) \quad \text{for some } u, v \in V.$$

Then while we iteratively build \mathcal{B}_i , each time we choose a node pair s, t that maximizes $\left|W_i^{\mathcal{B}_{\leq i-1} \cup \mathcal{B}_i}[s \rightsquigarrow t]\right|$, we must have

$$\left|W_i^{\mathcal{B}_{\leq i-1} \cup \mathcal{B}_i}[s \rightsquigarrow t]\right| = \omega_k(f^{i-1}).$$

Thus, when we add some $b \in W_i^{\mathcal{B}_{\leq i-1} \cup \mathcal{B}_i}[s \rightsquigarrow t]$ to \mathcal{B}_i , the size of $(|W_i[s \rightsquigarrow t]| - 1)^2$ falls from $|W_i[s \rightsquigarrow t]|^2$ to

$$(|W_i[s \rightsquigarrow t]| - 1)^2 = |W_i[s \rightsquigarrow t]|^2 - \Omega(|W_i[s \rightsquigarrow t]|) = |W_i[s \rightsquigarrow t]|^2 - \omega_k(f^{i-1})$$

and so the total number of i -walk meets in H falls by an additive $\omega_k(f^{i-1})$ term in each iteration of building \mathcal{B}_i . By Lemma 6 there are $n\Theta_k(D)^i$ total i -walks in H , so

$$|\mathcal{B}_i| = \phi_k \cdot n\Theta_k(D)^i = n\Theta_k(D)^i.$$

Hence the number of i -walk meets falls by

$$\omega_k(f^{i-1}) \cdot n\Theta_k(D)^i = \omega_k(|E_H| \cdot (fD)^{i-1})$$

from start to finish of the process of building \mathcal{B}_i (since $|E_H| = n\Theta_k(D)$). This is a contradiction, since we assumed that initially $\left|M_i^{\mathcal{B}_{\leq i-1}}\right| = O_k(|E_H| \cdot (fD)^{i-1})$. Thus \mathcal{B}_i satisfies the lemma statement. \square

Lemma 10 (Generalization of Lemma 2). *There is a blockade set \mathcal{B} for which*

$$|M_k| = O_k\left(|E_H| (fD)^{k-1}\right)$$

and for all $i \leq k$, the set \mathcal{B}_i contains at most a ϕ_k fraction of all i -walks in H , for any desired constant $\phi_k > 0$ depending only on k .

Proof. We will inductively show that two statements hold for all $2 \leq i \leq k$:

- (The lemma statement) There is a blockade set $\mathcal{B}_{\leq i-1}$ (with maximum blockade length $i-1$) such that (1) \mathcal{B}_j contains at most a ϕ_k fraction of all j -walks in H for all $j \leq i-1$, and (2) $\left|M_i^{\mathcal{B}_{\leq i-1}}\right| = O_k(|E_H| (fD)^{i-1})$.
- (An auxiliary statement) We have $\left|W_j^{\mathcal{B}}[u \rightsquigarrow v]\right| = O_k(f^{j-1})$ for all $u, v \in V$ and $j \leq i-1$.

We first argue the base case $i = 2$. For the first condition (lemma statement), we appeal to Lemma 2 to argue that the bound holds even when $\mathcal{B}_{\leq 1} = \emptyset$. Specifically, note that without blockades, each closed 4-walk in H corresponds in the natural way to $\Theta(1)$ i -walk meets. Thus

$$|M_2^\emptyset| = O(|E_H| \cdot fD).$$

We also note that the auxiliary statement holds trivially when $i = 2$; it simply states that there is at most 1 edge between any two nodes.

Now we argue the inductive step, assuming that both of the above properties hold for i . The first part of the induction (the lemma statement) coincides with the premise of Lemma 9, so we may find a set \mathcal{B}_i as in the conclusion of Lemma 9. That is, we have

$$|W_i^{\mathcal{B}_{\leq i-1} \cup \mathcal{B}_i}[u \rightsquigarrow v]| = O_k(f^{i-1}) \quad \text{for all } u, v \in V.$$

This statement (together with the auxiliary inductive hypothesis) proves the auxiliary statement for the case $i + 1$. The auxiliary statement for the case $i + 1$ also coincides with the premise of Lemma 8; applying this lemma, we have

$$|M_{i+1}^{\mathcal{B}_{\leq i-1} \cup \mathcal{B}_i}| = O_k(|E_H| \cdot (fD)^i).$$

This is the lemma statement for the case $i + 1$, which completes the inductive hypothesis. Hence the first part of our inductive hypothesis (the lemma statement) holds in the case $i = k$, which completes the proof. \square

B.4 Proof of Theorem 1

It is now a matter of algebra to complete the proof of Theorem 1.

Proof of Theorem 1. Let \mathcal{B} be a blockade set as in Lemma 10. By Lemma 10, we have

$$|M_k| = O_k(|E_H| (fD)^{k-1}).$$

Additionally, since \mathcal{B}_i has at most a ϕ_k fraction of all i -walks for each i , we have $|W_i| \geq n\Theta(D)^k$ (by Lemma 6) and so

$$|E_H| = O_k(n|M_k|^{1/(2k)})$$

by Lemma 7. Combining these, we compute

$$\begin{aligned} |E_H| &= O_k\left(n\left(|E_H|(fD)^{k-1}\right)^{1/(2k)}\right) \\ |E_H|^{2k-1} &= O_k\left(n^{2k}(fD)^{k-1}\right) \\ |E_H|^k &= O_k\left(n^{k+1}f^{k-1}\right) && \text{since } |E_H| = \Theta_k(nD) \\ |E_H| &= O_k\left(n^{1+1/k}f^{1-1/k}\right) \end{aligned}$$

as claimed. \square

C Lower Bounds for Larger k

We now prove our lower bounds in full generality.

C.1 Proof of Theorem 2

Let G be a graph from the girth conjecture (Conjecture 1), i.e., a graph with girth at least $2k + 2$ and $\Omega(n^{1+1/k})$ edges. We construct a new graph G' as follows. Let $t = \lceil f/2 \rceil$. We set $V_{G'} = V(G) \times [t]$, and let $E_{G'} = \{\{(u, i), (v, j)\} : \{u, v\} \in E(G) \wedge i, j \in [t]\}$. Let $G' = (V_{G'}, E_{G'})$. Intuitively, we can think of G' as being obtained by replacing each vertex of G by a set of t copies of the vertex, and each edge of G is replaced by a complete bipartite graph between the two sets of copies. We will prove Theorem 2 by proving that the only f VFT $(2k - 1)$ -spanner of G' is itself, and that it has the required number of edges.

We first claim that the only f VFT $(2k - 1)$ -spanner of G' is G' itself. To see this, suppose that H is a subgraph of G' which does not contain some edge $\{(u, i), (v, j)\} \in E_{G'}$. Let $F = \{(u, \ell) : \ell \neq i\} \cup \{(v, \ell) : \ell \neq j\}$, i.e., we let the fault set be all copies of u except for (u, i) and all copies of v except (v, j) . Note that $|F| \leq f$. Now consider the shortest path from (u, i) to (v, j) in $H \setminus F$. Let this path be $(u, i) = (x^0, i^0), (x^1, i^1), (x^2, i^2), \dots, (x^p, i^p) = (v, j)$. Note that for any $0 \leq a \leq p - 1$, it cannot be the case that $x^a = u$ and $x^{a+1} = v$, since no such edges exist in $H \setminus F$. Thus $\langle u = x^0, x^1, \dots, x^p = v \rangle$ is a (possibly non-simple) path from u to v in G which does not use the edge $\{u, v\}$. By adding $\{u, v\}$ to this path, we get a cycle of length at most $p + 1$. Since G has girth at least $2k + 2$, this implies that $p \geq 2k + 1$, and thus that in $H \setminus F$ the distance between (u, i) and (v, j) is at least $2k + 1$, while in $G' \setminus F$ they are at distance 1. Thus H is not an f VFT $(2k - 1)$ -spanner of G' , and hence the only f VFT $(2k - 1)$ -spanner of G' is G' itself.

So G' is the only f VFT $(2k - 1)$ -spanner of itself, and it remains only to analyze its size. Clearly $|V_{G'}| = t|V(G)|$, and by assumption $|E(G)| \geq \Omega(|V(G)|^{1+1/k})$. Thus

$$\begin{aligned} |E_{G'}| &= |E(G)|t^2 = \Omega(t^2|V(G)|^{1+1/k}) = \Omega\left(t^2 \cdot \left(\frac{|V_{G'}|}{t}\right)^{1+1/k}\right) \\ &= \Omega\left(f^{1-1/k}|V_{G'}|^{1+1/k}\right). \end{aligned}$$

C.2 Proof of Theorem 3

Our proof is very similar to the proof of Theorem 2. The case of $k = 2$ was proved in Section 4, so we only prove the theorem for $k \geq 3$, where we are forced to use a slightly different construction which will give only the weaker bound claimed in Theorem 3.

Let $t = \lfloor \sqrt{f} \rfloor$. We set $V_{G'} = V(G) \times [t]$, and as before set $E_{G'} = \{\{(u, i), (v, j)\} : \{u, v\} \in E(G) \wedge i, j \in [t]\}$ (where G is a graph from the girth conjecture). Let H be a subgraph of G' missing some edge $e = \{(u, i), (v, j)\}$. As in the other lower bounds, we will prove that H cannot be an f EFT $(2k - 1)$ spanner of G' , and thus the only f EFT $(2k - 1)$ -spanner of G' is G' itself.

Consider the fault set $F = \{\{(u, i'), (v, j')\} : i' \neq i \vee j' \neq j\}$. In other words, we remove all edges between copies of u and copies of v *except* for the edge e (which is therefore in $G' \setminus F$ but not in $H \setminus F$). Now let $P = \langle (u, i) = (x^0, i^0), (x^1, i^1), (x^2, i^2), \dots, (x^p, i^p) = (v, j) \rangle$ be the shortest path from (u, i) to (v, j) in $H \setminus F$. Note that for any $0 \leq a \leq p - 1$, it cannot be the case that $x^a = u$ and $x^{a+1} = v$, since no such edges exist in $H \setminus F$. Thus $\langle u = x^0, x^1, \dots, x^p = v \rangle$ is a (possibly non-simple) path from u to v in G which does not use the edge $\{u, v\}$. By adding $\{u, v\}$ to this path, we get a cycle of length at most $p + 1$. Since G has girth at least $2k + 2$, this implies that $p \geq 2k + 1$, and thus that in $H \setminus F$ the distance between (u, i) and (v, j) is at least $2k + 1$, while in $G' \setminus F$ they are at distance 1. Thus H is not an f EFT $(2k - 1)$ -spanner of G' , and hence the only f EFT $(2k - 1)$ -spanner of G' is G' itself.

It remains only to analyze the size of G' . By construction, $|V_{G'}| = t|V(G)|$, and by assumption

$|E(G)| \geq \Omega(|V(G)|^{1+1/k})$. Thus

$$\begin{aligned} |E_{G'}| &= |E(G)|t^2 = \Omega(|V(G)|^{1+1/k}t^2) = \Omega\left(t^2 \cdot \left(\frac{|V_{G'}|}{t}\right)^{1+1/k}\right) \\ &= \Omega\left(t^{1-1/k}|V_{G'}|^{1+1/k}\right) = \Omega\left(f^{1/2-1/(2k)}|V_{G'}|^{1+1/k}\right). \end{aligned}$$

C.3 Strong Incompressibility

We now show how to generalize the above proofs to obtain strong incompressibility (Theorems 4 and 5). This type of argument was first used by Matoušek [36].

The proof of Theorem 4 (incompressibility under vertex faults) goes as follows. Let $G^* = (V, E^*)$ be the graph described in the proof of Theorem 2. There are $2^{|E^*|}$ possible subgraphs of G^* on the same vertex set. Consider two distinct such subgraphs $G_1 = (V, E_1) \neq G_2 = (V, E_2)$, let \mathcal{D}_1 be the data structure created by processing G_1 , let \mathcal{D}_2 be the data structure created by processing G_2 , and let $((u, i), (v, j)) \in E_1 \setminus E_2$. As before, let

$$F := \{(u, \ell) : \ell \neq i\} \cup \{(v, \ell) : \ell \neq j\}.$$

By the previous argument, we then have

$$\text{dist}_{G_1 \setminus F}((u, i), (v, j)) = 1 \quad \text{and} \quad \text{dist}_{G_2 \setminus F}((u, i), (v, j)) \geq 2k + 1$$

and so

$$\widehat{\text{dist}}_{G_1 \setminus F}((u, i), (v, j)) \leq 2k - 1 \quad \text{and} \quad \widehat{\text{dist}}_{G_2 \setminus F}((u, i), (v, j)) \geq 2k + 1.$$

Thus $\mathcal{D}_1, \mathcal{D}_2$ produce different answers to the query $((u, i), (v, j), F)$, so they must have different representations. Since this holds for any two subgraphs $G_1, G_2 \subseteq G^*$, we have a family of $2^{|E^*|}$ graphs that all have different representations, and so by the pigeonhole principle the data structure \mathcal{D}_G for one of these subgraphs $G \subseteq G^*$ occupies at least $|E^*| = \Omega(f^{1-1/k}n^{1+1/k})$ bits of space, thus completing the proof.

The proof of Theorem 5 is identical, taking F to be a set of edge faults as described in the above proof of Theorem 3.