

Unconditional Lower Bounds for Degree Fault Tolerant Spanners

Greg Bodwin  

Department of EECS, University of Michigan, Ann Arbor, MI, USA

Aleksey Lopez¹  

Department of EECS, University of Michigan, Ann Arbor, MI, USA

Abstract

We study multiplicative graph spanners in the *f-degree fault tolerant (f-DFT)* model, in which the spanner must approximately preserve distances even after any subset of edges of maximum degree f temporarily “fails” and is removed from the graph. We prove that there are n -node lower bound graphs for which any f -DFT $(2k - 1)$ -stretch spanner H must have size

$$|E(H)| \geq \Omega\left(f^{1-1/k} n^{1+1/k}\right).$$

This matches a lower bound that was previously only known to hold conditionally, under the 1963 *girth conjecture* of Erdős. It also matches the current upper bounds, up to a factor of $\exp(k)$. Our proof is an analysis of the so-called *Wenger graphs* (J. Comb. Theory 1991), via their recent reinterpretation by Szabó and by Conlon (Am. Math. Monthly 2021).

2012 ACM Subject Classification Mathematics of computing → Paths and connectivity problems; Mathematics of computing → Extremal graph theory; Mathematics of computing → Graph algorithms

Keywords and phrases Spanners, Fault Tolerance, Girth Conjecture, Wenger Graphs

Digital Object Identifier 10.4230/LIPIcs.ESA.2026.31

Funding Greg Bodwin: Supported in part by NSF:AF 2153680.

1 Introduction

1.1 Spanners and the Girth Conjecture

A *multiplicative spanner* is a primitive in graph theory and algorithms, which sparsifies an input graph while approximately preserving its shortest path distances. Spanners were abstracted in the late '80s after being implicitly studied in networking, distributed algorithms, etc (see survey [1]).

► **Definition 1** (Spanners [29]). *Given a graph $G = (V, E, w)$, a subgraph $H = (V, E', w)$ is a k -spanner of G if for all nodes s, t , we have $\text{dist}_H(s, t) \leq k \cdot \text{dist}_G(s, t)$.*

Spanners and their variants are often studied from the *extremal* perspective, where the goal is to prove that one can always construct a spanner with a good tradeoff between size and error. The following classic theorem of Althöfer, Das, Dobkin, Joseph, and Soares [2] conditionally settles the extremal tradeoff available for multiplicative spanners:

► **Theorem 2** ([2]). *For all positive integers n, k , every n -node graph has a $(2k - 1)$ -spanner on $O(n^{1+1/k})$ edges. Moreover, assuming the girth conjecture [18], for all positive integers k there are families of n -node graphs on which this size bound cannot be improved to $o(n^{1+1/k})$.*

¹ Author order has been alphabetized by surname in accordance with conference convention, but Aleksey Lopez is the primary and corresponding author for this paper.



The *girth conjecture* referenced in this theorem, attributed to Erdős [18], says the following. Let us denote by $\text{ex}(n, C_{\leq 2k})$ the maximum possible number of edges in an n -node graph that does not have any cycles of length $\leq 2k$ as subgraphs. There is a simple folklore counting argument, called the *Moore bound*, that implies

$$\text{ex}(n, C_{\leq 2k}) \leq O\left(n^{1+1/k}\right).$$

The girth conjecture posits that this bound is asymptotically best possible for all k . The girth conjecture is an old and very difficult problem in extremal graph theory, likely beyond the reach of current techniques. It has been confirmed only for $k \in \{1, 2, 3, 5\}$ [34, 35], with no new cases settled since the 1950s. In fact, the community has more recently started to regard it with some skepticism. For example, in an article from 2021, Conlon [14] writes:

“It is now quite commonplace to believe that the true value of certain extremal numbers lie below what the classical arguments give. We suspect that this should already be the case for C_8 , that is, that $\text{ex}(n, C_8) = o(n^{5/4})$.”


Here $\text{ex}(n, C_8)$ is the maximum possible number of edges in an n -node graph with no C_8 subgraph. It follows from the definitions that $\text{ex}(n, C_8) \geq \text{ex}(n, C_{\leq 8})$, and so the suggested bound of $o(n^{5/4})$ would imply that the girth conjecture is false already for $k = 4$. Since this seems difficult to prove, Conlon goes on to give a more attainable conjecture, very roughly stating that the techniques from incidence geometry used to confirm the girth conjecture for $k \in \{1, 2, 3, 5\}$ cannot be extended to other values of k . This conjecture is related to the classical Feit-Higman Theorem [19], which is sometimes interpreted as further evidence against the girth conjecture.² For further technical discussion and references on the girth conjecture and its surrounding literature, we recommend the excellent set of lecture notes by Szabó [33].

1.2 Fault Tolerant Spanners

Spanners are often used in distributed systems and networking, where there is a possibility of nodes or edges in the input graph G temporarily *failing* and becoming unusable while they await repair. A desirable feature in these application domains is for a spanner H to retain its distance approximation guarantees even after any “reasonable” failure event occurs.

The first formal model of fault tolerant spanners was introduced by Levkopoulos, Narasimhan, and Smid [22]. They defined H to be an f -vertex (edge) fault tolerant k -spanner of G if, for any set F of $|F| \leq f$ vertices (edges), the subgraph $H \setminus F$ is still a k -spanner of $G \setminus F$. The initial results in this model were for geometric graphs, which has remained an area of intensive focus for the community [9–12, 15, 21–24, 32]. Results for general graphs were obtained later by Chechik, Langberg, Peleg, and Roditty [13], which has inspired a similarly long line of followup work [4–6, 8, 16, 17, 25]. By now, this model of fault tolerance is mostly well understood. We have optimal existential bounds [4, 8] and near-optimal algorithms [25]

² The constructions implying the girth conjecture for $k \in \{1, 2, 3, 5\}$ arise from highly symmetric incidence structures known as *finite generalized polygons*. The Feit-Higman Theorem constrains the existence of finite generalized polygons, and in particular, it implies that finite generalized polygons cannot be used to prove the girth conjecture for any further values of k .



■ **Figure 1** An input graph G (left) and a 1-DFT 3-spanner H (right). Note that after any degree-1 fault set is removed from both G and H (such as the pair of edges marked with red X's), the remaining part of H is still a 3-spanner of the remaining part of G .

for the setting of vertex faults. We also have optimal existential bounds for edge faults for infinitely many values of spanner error k [4, 5], although some knowledge gaps remain. As usual, the lower bounds in both settings are conditioned on the girth conjecture.³

More recently, the community has turned toward new and more expressive models of fault tolerance, that strengthen the original model without harming the spanner size too much. One such model is *color-fault tolerant spanners*, a compelling model of correlated failure introduced recently by Petruschka, Sapir, and Tzalik [31] (see also [26, 27, 30, 31]). Optimal or near-optimal bounds are known for this setting as well, again with the lower bounds conditioned on the girth conjecture.⁴ Another more expressive fault tolerance model, which will be a focus of this work, is *degree fault tolerant spanners*:

► **Definition 3** (Degree Fault Tolerant Spanners [7]). *Given a graph $G = (V, E, w)$, a subgraph $H = (V, E', w)$ is an f -degree fault tolerant (f -DFT) k -spanner of G if, for any edge subset $F \subseteq E$ such that F contains at most f edges incident on each node, we have*

$$\text{dist}_{H \setminus F}(s, t) \leq k \cdot \text{dist}_{G \setminus F}(s, t) \quad \text{for all nodes } s, t.$$

An f -DFT spanner can tolerate many more edge failures at once than in the original model of fault tolerance; for example, it can handle $O(nf)$ edge failures, so long as they are spread throughout the graph such that at most f failing edges are incident to any particular node. Despite this possibility of massive failure, Bodwin, Haeupler, and Parter [7] showed that fairly reasonable existential size bounds are available, not much worse than the ones known for the original model:

► **Theorem 4** ([7]). *For all positive integers k , every n -node graph G has an f -DFT $(2k - 1)$ -spanner H on*

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

edges. Moreover, assuming the girth conjecture, there is a family of n -node input graphs G for which any f -DFT $(2k - 1)$ -spanner H must have

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

³ In the setting of vertex faults (but not edge faults), we remark that the upper bounds and lower bounds both work *by reduction* to $\text{ex}(n, C_{\leq 2k})$, and hence proving an unconditional tight lower bound would be equivalent to proving the girth conjecture itself [4, 8]. In the setting of edge faults, for higher fault tolerance parameters f , the lower bounds can also be proved by conditioning on a certain bipartite extension of the girth conjecture; see [6], Section 2.1 for discussion.

⁴ Again, the upper and lower bounds in the color-fault model work by reduction to $\text{ex}(n, C_{\leq 2k})$, and so proving an unconditional lower bound is equivalent to proving the girth conjecture [31].

Thus conditionally nearly-tight bounds are known for this problem, up to a factor of $\exp(k)$, and assuming the girth conjecture. Followup work by Parter and Tzalik [28] obtained matching upper and lower bounds in the reachability ($k = \infty$) setting, and Biniaz, Carufel, Maheshwari, and Smid [3] obtained results in the DFT model for metric and geometric spanners.

1.3 Our Results

The contribution of this paper is to show that the additional strength of the DFT model is enough to obtain *unconditional* lower bounds. We show:

► **Theorem 5.** *For all positive integers k , there are n -node input graphs G for which any f -DFT $(2k - 1)$ -spanner H must have*

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

These lower bounds are both optimal up to $O_k(1)$ factors, and they hold via explicit constructions that do not assume the girth conjecture. This makes DFT spanners essentially the first variants of multiplicative spanners where one can obtain unconditionally near-optimal lower bounds. For essentially every other variant, the best unconditional lower bounds are inherited from the lower bounds for the function $\text{ex}(n, C_{\leq 2k})$, where the state-of-the-art is from a 1995 work of Lazebnik, Ustimenko, and Woldar [20]. These bounds are polynomially below the bound posited by the girth conjecture, meaning that our results represent polynomial improvements in the known unconditional bounds.

These theorems may also have independent interest from the standpoint of extremal combinatorics. In some sense, they state that the DFT model captures the gap between known algebraic techniques and the structure that would be needed to prove the girth conjecture, if possible. More specifically, our proofs analyze the *Wenger graphs* [35], which are an elegant algebraic construction that notably achieves optimal lower bounds on the function $\text{ex}(n, C_{2k})$ for $k \in \{1, 2, 3, 5\}$.⁵ Szabó [33] and Conlon [14] independently devised a reinterpretation of Wenger’s original construction with a more algebraic flavor. We adopt their interpretation, and make algebraic arguments to reason about the fine-grained structure of the short cycles that appear in these graphs. The main part of our proofs can be viewed as a formalization of the sense in which the Wenger graphs come close to proving the girth conjecture; that is, their short cycles are sparsely “spread out” throughout the graph in a way that lets one destroy the relevant ones with a low-degree fault set.


2 Proof of Main Result

2.1 Construction of Incidence System and Fault Sets

We begin by constructing an incidence system, consisting of n *points* and n *lines* (subsets of the points). This construction and its incidence graph are originally by Wenger [35], while our algebraic exposition follows Szabó [33] and Conlon [14].

- **(The parameters).** The construction is parametrized by an arbitrary prime p and an arbitrary integer $k \geq 1$. We will have $n := p^k$ points and lines in the construction. We will use the standard shorthand $[p] := \{0, 1, \dots, p - 1\}$.

⁵ The Wenger graphs also imply the girth conjecture for $k \in \{1, 2, 3\}$, but not for $k = 5$: although they do not have C_{10} subgraphs, they do have C_8 subgraphs.



■ **Figure 2** (Left) The 25 points and first 3 lines used in the construction with $k = 2$ and $p = 5$. (Right) The corresponding incidence graph. In the full system, there would be 25 lines and therefore 25 nodes on the right side of the incidence graph.

- **(The points).** The set of points A is the elements of the finite field $\mathbb{F}_p^k$. We interpret $\mathbb{F}_p^k$ as a vector space with scalars in $\mathbb{F}_p$, and so we will perform scalar multiplication with scalars from $[p]$.
- **(The lines).** For any $t \in [p]$, let $\sigma_t := (1, t, t^2, \dots, t^{k-1}) \in \mathbb{F}_p^k$; we will refer to σ_t as a *slope vector*. The set of lines L is exactly the lines of slope σ_t for any $t \in [p]$. In other words, a subset of points forms a line ℓ iff there exists $t \in [p]$ such that ℓ is a maximal subset of points with the property that any two points in ℓ differ by a multiple of σ_t . Formally:

$$L := \{ \{x + \lambda \sigma_t \mid \lambda \in [p]\} \mid x \in \mathbb{F}_p^k, t \in [p] \}.$$

- **(The graph).** We define $\mathcal{B}$ to be the bipartite *incidence graph* between the set of points A and the set of lines L . That is: the left side is A , the right side is L , and there is an edge $\{a, \ell\} \in E(\mathcal{B})$ iff $a \in \ell$.

The graph $\mathcal{B}$ has the following statistics:


► **Lemma 6 (Size of $\mathcal{B}$).** *The graph $\mathcal{B}$ has $2n$ nodes and $n^{1+1/k}$ edges.*

Proof. There are $n = p^k$ nodes in $\mathbb{F}_p^k$. The lines are indexed by p^k possible choices of x and p possible choices of t , but each line can be generated in p ways (by choosing any of its p points for x), so there are $p^k \cdot p/p = p^k = n$ lines. Thus $\mathcal{B}$ has n nodes on each side of its bipartition, for $2n$ nodes in total. Each line contains exactly p points, so it generates p edges in $\mathcal{B}$. So there are $p^{k+1} = n^{(k+1)/k} = n^{1+1/k}$ edges in $\mathcal{B}$. ◀

Our goal will be to argue that $\mathcal{B}$ is the only 1-DFT $(2k - 1)$ -spanner of itself. This gives our desired lower bound in the special case $f = 1$. Later, in Section 2.4, we will describe a standard technique used to boost the lower bound to arbitrary $f \geq 1$. The first step towards this goal is, for every edge e , to define a corresponding edge set F_e of maximum degree 1 (i.e., a matching) such that e does not participate in any cycles of length $\leq 2k$ in the graph $\mathcal{B} \setminus F_e$. The next step in our proof is to define these sets F_e . The following definitions will be useful:

► **Definition 7 (Line and Path Properties).** *Two lines $\ell, \ell' \in L$ are **parallel**, written $\ell \parallel \ell'$, if they are distinct and they have the same direction vector σ_t . A path π in $\mathcal{B}$ with vertex sequence*

$$\pi = (a = a_0, \ell_0, a_1, \ell_1, \dots, \ell_{r-1}, a_r = b)$$



■ **Figure 3** When considering the edge e corresponding to the point-line incidence $a \in \ell$, we add to its associated fault set F_e the edges corresponding to point-line incidences $b \in \ell'$, where $\ell \parallel \ell'$ and b can be reached from a by an ℓ -transverse path of small line-length.

is called ℓ -**transverse** if it is simple⁶ and none of the lines $\{\ell_0, \ell_1, \dots, \ell_{r-1}\}$ are equal to ℓ or parallel to ℓ . Its **line-length** is the number of lines in L corresponding to nodes in the path; for example, the above path π has $\text{line-length}(\pi) = r$.

► **Definition 8** (Failure Set F_e). For any edge $e = \{a, \ell\} \in E(\mathcal{B})$, we define $F_e \subseteq E(\mathcal{B}) \setminus \{e\}$ to be the set of all edges $\{b, \ell'\}$ such that $\ell \parallel \ell'$, and there exists an ℓ -transverse path from a to b of line-length $\leq \frac{k}{2} - 1$.⁷

2.2 Bound on Degree of F_e

Our next goal is to prove that, for each edge $e \in E(\mathcal{B})$, the associated fault set F_e has maximum degree 1. We will need two supporting structural lemmas before the main proof.

► **Lemma 9.** For any closed, non-backtracking⁸ walk W in $\mathcal{B}$ of length $|W| \leq 2k$, there is no $t \in [p]$ such that W contains exactly one line of slope σ_t . (This allows the case where W contains two instances of the same line ℓ .)

Proof. Let the vertex sequence of W be

$$W = (a_0, \ell_0, a_1, \ell_1, \dots, \ell_{r-1}, a_r = a_0),$$

where each a_j is a point and each ℓ_j is a line with slope vector σ_j , and $r \leq k$. Since $a_j, a_{j+1} \in \ell_j$, there exists a scalar $\lambda_j \in \mathbb{F}_p$ such that $a_{j+1} = a_j + \lambda_j \sigma_j$. Moreover, since the walk is non-backtracking, we have $a_j \neq a_{j+1}$ and so these scalars are all nonzero. Summing these equations around the walk W gives the identity


$$\sum_{j=0}^{r-1} \lambda_j \sigma_j = 0.$$

However, it follows from the usual analysis of the Vandermonde matrix [14] that any k slope vectors $\{\sigma_j\}$ are linearly independent. Thus, this identity is possible only if after grouping

⁶ Recall: a simple path is one that does not repeat nodes.

⁷ The expression $\frac{k}{2} - 1$ is not necessarily integral, but line-length is necessarily integral, so this is equivalent to an upper bound on line-length of $\leq \lfloor \frac{k}{2} - 1 \rfloor$. Since the floor symbols are not technically necessary, we will omit them here and for other possibly-non-integral expressions encountered as upper bounds in the rest of this paper.

⁸ Recall: a walk is *closed* if it starts and ends at the same node, and it is *non-backtracking* if it never uses the same edge twice in a row with opposite orientations.



■ **Figure 4** A depiction of the “path-straightening” technique used to prove Lemma 12.

equal slope vectors together, their associated coefficients must sum to 0. Since each individual coefficient λ_j is nonzero, this implies that each slope vector must have at least two associated coefficients, implying the lemma. ◀

► **Lemma 10** (Unique endpoints). *Let ℓ, ℓ' be parallel lines in L and let $a \in \ell$. Then there is a point $b \in \ell'$ such that every ℓ -transverse path of line-length $\leq \frac{k}{2} - 1$ that starts at a and ends at a point in ℓ' specifically ends at b .*

Proof. Suppose there are ℓ -transverse paths with node sequences

$$\pi_1 = (a, \ell_0, a_1, \ell_1, \dots, \ell_{r-1}, b) \quad \text{and} \quad \pi_2 = (a, m_0, b_1, m_1, \dots, m_{s-1}, b')$$

ending at points $b, b' \in \ell'$, of respective line-lengths r, s both $\leq \frac{k}{2} - 1$. Suppose for contradiction that $b \neq b'$. Concatenating π_1 and the reverse of π_2 with the node ℓ' placed between them produces a closed walk W in $\mathcal{B}$. Consider a minimal nonempty closed subwalk $W' \subseteq W$ that contains an edge of ℓ' , which is a simple cycle (and hence non-backtracking). We have

$$\text{line-length}(W') \leq r + 1 + s \leq 2 \cdot \left(\frac{k}{2} - 1 \right) + 1 \leq k - 1.$$

Hence the total length of W' in $\mathcal{B}$ is $|W'| \leq 2k$. Moreover, the only line in W' parallel to ℓ is ℓ' , since both π_1 and π_2 are ℓ -transverse. This contradicts Lemma 9. Thus we have $b = b'$, completing the proof. ◀

► **Lemma 11.** *For any edge e , the edge set $F_e \cup \{e\}$ has maximum degree 1.*

Proof. Let $e =: \{a, \ell\}$. We analyze $\deg_{F_e}(v)$ in cases:

- Suppose that v corresponds to a point (rather than a line) in the incidence system. Since v lies on exactly one line ℓ' parallel to ℓ (or v lies on ℓ itself), the only edge containing v that might be in F_e is (v, ℓ') , so $\deg_{F_e}(v) \leq 1$.
- Suppose that v corresponds to a line in the incidence system that is *not* parallel to ℓ . Then by construction, $\deg_{F_e}(v) = 0$.
- Suppose that v corresponds to a line in the incidence system that is parallel to ℓ . By Lemma 10, there is at most one point $b \in v$ that is the endpoint of an ℓ -transverse path from a of line-length $\leq \frac{k}{2} - 1$. Hence v is incident to at most one edge of F_e .
- Suppose that $v = \ell$. By construction, the only edge in $F_e \cup \{e\}$ incident to v is e itself. ◀

2.3 Correctness of Fault Set F_e

Our next goal is to prove that each edge e does not participate in any cycles of length $\leq 2k$ in the graph $\mathcal{B} \setminus F_e$.

► **Lemma 12** (Path straightening). *Let π be an ℓ -transverse $a \rightsquigarrow b$ path in $\mathcal{B}$ that uses lines with z distinct slope vectors. Then there is an ℓ -transverse path π' with the same endpoints $a \rightsquigarrow b$ that has line-length $\leq z$.*

Proof. Let π' be an $a \rightsquigarrow b$ path in $\mathcal{B}$ of minimum length, among the paths with the property that every slope vector used by a line in π' is also used by at least one line in π . Note that, since π is ℓ -transverse, this implies that π' is ℓ -transverse as well. Let the vertex sequence of π' be

$$\pi' = (a = a_0, \ell_0, a_1, \dots, \ell_{s-1}, a_s = b).$$

If all lines $\{\ell_i\}$ used by π' have pairwise distinct slope vectors, then since π' has at most z distinct slope vectors, we get $s \leq z$. So π' has line-length $\leq z$ and it satisfies the lemma.

So, for the remaining part of the proof, assume towards contradiction that not all lines $\{\ell_i\}$ have distinct slope vectors. For each index i , let $\Delta_i := a_{i+1} - a_i$. We may therefore describe the sequence of point-nodes along π' as

$$\left(a, a + \Delta_0, a + \Delta_0 + \Delta_1, \dots, a + \sum_{i=0}^{s-1} \Delta_i = b \right).$$

Next, let ϕ be an arbitrary permutation on $[s]$. We define π_ϕ as the walk in $\mathcal{B}$ whose sequence of point-nodes is

$$\left(a, a + \Delta_{\phi(0)}, a + \Delta_{\phi(0)} + \Delta_{\phi(1)}, \dots, a + \sum_{i=0}^{s-1} \Delta_{\phi(i)} = b \right).$$

Note that adjacent point-nodes in this sequence differ by some vector Δ_i , which is the slope of the line ℓ_i , and so by construction there is indeed a line that contains these adjacent points. This confirms that there is a unique walk π_ϕ with this sequence of point-nodes, so π_ϕ is well defined.

We assumed that a slope vector is used in π' at least twice, so by choice of ϕ , we may have


$$\Delta_{\phi(0)} = \lambda \cdot \Delta_{\phi(1)}$$

for some scalar $\lambda \in [p]$. Thus, we may delete the second point-node along π_ϕ , giving a shorter sequence

$$\left(a, a + \Delta_{\phi(0)} + \Delta_{\phi(1)}, \dots, a + \sum_{i=0}^{s-1} \Delta_{\phi(i)} = b \right).$$

The difference between the first two point nodes is $(\lambda + 1) \cdot \Delta_{\phi(1)}$, so these points are still connected by a line of a viable slope vector. If $\lambda = -1$ then we simply omit that step, which still results in a shorter walk. Thus there is a corresponding $a \rightsquigarrow b$ walk of length shorter than π' . Deleting closed subwalks yields a path that contradicts minimality of the length of π' , completing the proof. ◀

► **Lemma 13.** *Every cycle C in $\mathcal{B}$ of length at most $2k$ that contains an edge $e = \{a, \ell\}$ also contains an edge of F_e .*



■ **Figure 5** Geometric view of the cycle-hitting argument. The solid orange path is the chosen ℓ -transverse arc of the cycle, and the dashed orange path is the straightened path with the same endpoints.

Proof. Let the vertex sequence of C be

$$C = (a, \ell, a_1, \ell_1, a_2, \ell_2, \dots, a_{r-1}, \ell_{r-1}, a_r = a)$$

with $r \leq k$. By Lemma 9, at least one of the other lines in C is parallel to ℓ . Let ℓ_i be such a line, with the index i selected as large as possible. Our goal will now be to argue that $\{a_{i+1}, \ell_i\} \in F_e$.

There are r distinct lines in C . By Lemma 9, each slope vector that is used by any of these r lines is used by at least two of them, so there are at most $\frac{r}{2}$ distinct slope vectors used by lines in C . Now consider the subpath $\pi \subseteq C$ with vertex sequence

$$\pi = (a_r, \ell_{r-1}, a_{r-1}, \dots, \ell_{i+1}, a_{i+1}).$$

By maximal choice of i , this is an ℓ -transverse path, and hence does not contain any lines whose slope vector matches ℓ . So there are at most $\frac{r}{2} - 1$ distinct slope vectors used by the lines in π . Thus, by Lemma 12, there is an ℓ -transverse path π' with the same endpoints (a_r, a_{i+1}) that has line-length at most $\frac{r}{2} - 1 \leq \frac{k}{2} - 1$. Thus, by construction, the set F_e associated to $e = \{a = a_r, \ell\}$ will include the edge $(a_{i+1}, \ell_i) \in C$. ◀

2.4 Proof Wrapup

The following theorem certifies our lower bound for the special case $f = 1$:


► **Theorem 14.** *The only 1-DFT $(2k - 1)$ -spanner of $\mathcal{B}$ is $\mathcal{B}$ itself.*

Proof. Let H be a proper subgraph of $\mathcal{B}$, and we will show that H is not a 1-DFT $(2k - 1)$ -spanner of $\mathcal{B}$. Consider an edge

$$e = \{u, v\} \in E(\mathcal{B}) \setminus E(H).$$

Let F_e denote the set associated to this edge (as in Definition 8). By Lemma 11, the set F_e has maximum degree 1. Next, let π be a shortest $u \rightsquigarrow v$ path in the graph $H \setminus F_e$. Since $e \notin E(H)$, we have $e \notin \pi$, and so $\pi \cup \{e\}$ is a cycle in $\mathcal{B}$. This cycle contains e but it does not contain any edge of F_e . By Lemma 13, this means $\pi \cup \{e\}$ contains $> 2k$ edges. Thus π contains $> 2k - 1$ edges, so we have

$$|E(\pi)| = \text{dist}_{H \setminus F_e}(u, v) > 2k - 1.$$



■ **Figure 6** A lifted edge in the cloud blowup. Each base vertex is replaced by a cloud of f copies, and each base edge $e = \{u, v\}$ is replaced by the complete bipartite graph $\mathcal{L}(e)$ between $C(u)$ and $C(v)$.

On the other hand, $e \notin F_e$, so we have $\text{dist}_{\mathcal{B} \setminus F_e}(u, v) = 1$. This implies that H is not a 1-DFT $(2k - 1)$ -spanner of $\mathcal{B}$. ◀

We now upgrade the previous lower bound to an arbitrary degree parameter f using a standard cloud-blowup technique, which has appeared in many previous papers on fault tolerance [4, 7, 31].

► **Definition 15** (Cloud blowup). *For an integer $f \geq 1$, the f -cloud blowup of the base graph is*

$$\mathcal{B}^{(f)} := \mathcal{B} \times [f],$$

where $[f]$ denotes the edgeless graph on f vertices, and $\times$ denotes the f -cloud blowup operation defined below. In other words, it is the graph obtained from $\mathcal{B}$ by the following steps:

- For each vertex $x \in V(\mathcal{B})$, create a “cloud” of f nodes, $C(x) := \{(x, i) : i \in [f]\}$.
- For each base edge $\{x, y\} \in E(\mathcal{B})$, lift it to a complete bipartite graph between the clouds $C(x)$ and $C(y)$. That is,

$$\mathcal{L}(\{x, y\}) := \{\{(x, i), (y, j)\} \mid i, j \in [f]\} \quad \text{and} \quad E(\mathcal{B}^{(f)}) := \bigcup_{\{x, y\} \in E(\mathcal{B})} \mathcal{L}(\{x, y\}).$$

The following lemma computes how the cloud blowup affects the parameters of the graph:

► **Lemma 16** (Size of Blown-Up Graph). *Letting $2n$ be the number of vertices in $\mathcal{B}$, the corresponding blown-up graph $\mathcal{B}^{(f)}$ has $N := 2nf$ vertices and $\Theta(f^{1-1/k} N^{1+1/k})$ edges.*

Proof. Each node in $\mathcal{B}$ corresponds to f copies in $\mathcal{B}^{(f)}$, and each edge in $\mathcal{B}$ corresponds to f^2 edges in $\mathcal{B}^{(f)}$. Thus we have

$$\begin{aligned} |E(\mathcal{B}^{(f)})| &= f^2 \cdot |E(\mathcal{B})| \\ &= \Theta(f^2 n^{1+1/k}) && \text{Lemma 6} \\ &= \Theta(f^{1-1/k} (nf)^{1+1/k}) \\ &= \Theta(f^{1-1/k} N^{1+1/k}). \end{aligned}$$

From the previous lemma, it suffices to prove that $\mathcal{B}^{(f)}$ is the only f -DFT $(2k - 1)$ -spanner of itself. To do so, let us first define lifted fault sets:

► **Definition 17** (Lifted Fault Sets $\widehat{F}_e$). For an edge $e \in E(\mathcal{B})$ and an edge $\widehat{e} \in \mathcal{L}(e)$, we define $\widehat{F}_e$ to be the set of all lifted edges from $F_e \cup \{e\}$, except for $\widehat{e}$ itself. That is:

$$\widehat{F}_e := \left(\bigcup_{g \in F_e} \mathcal{L}(g) \right) \cup (\mathcal{L}(e) \setminus \{\widehat{e}\}).$$

The next lemma shows that the lifted fault sets have degree $\leq f$:

► **Lemma 18** (Lifted fault set has maximum degree f). For any edge $\widehat{e}$, the corresponding edge set $\widehat{F}_e$ has maximum degree $\leq f$ in $\mathcal{B}^{(f)}$.

Proof. Since $F_e \cup \{e\}$ has maximum degree 1 in the base graph, each base vertex $x \in V(\mathcal{B})$ is incident to at most one edge of $F_e \cup \{e\}$. Hence each lifted vertex in $\mathcal{B}^{(f)}$ is incident to at most one lifted biclique, and within that biclique it is incident to at most f edges. ◀

We are finally ready to wrap up the proof:

► **Theorem 19** (Cloud-blowup lower bound). For any positive integer f , the only f -DFT $(2k-1)$ -spanner of $\mathcal{B}^{(f)}$ is $\mathcal{B}^{(f)}$ itself.

Proof. Since the following cloud-lift is standard in the previous literature, we will sketch parts of the proof. Let H be a proper subgraph of $\mathcal{B}^{(f)}$. Choose a missing lifted edge

$$\widehat{e} = \{(u, i), (v, j)\} \in E(\mathcal{B}^{(f)}) \setminus E(H),$$

and let $e = \{u, v\} \in E(\mathcal{B})$ be the corresponding base edge; that is, $\widehat{e} \in \mathcal{L}(e)$. Consider the associated fault set $\widehat{F}_e$. By the previous lemma, $\widehat{F}_e$ has maximum degree at most f . Next, let π be a shortest $(u, i) \rightsquigarrow (v, j)$ path in the graph $H \setminus \widehat{F}_e$. Projecting $\pi \cup \{\widehat{e}\}$ to the base graph yields a closed walk in $\mathcal{B}$ containing e and avoiding F_e . Choose a closed subwalk of minimum positive length that contains e ; this is a cycle, still contains e , and still avoids F_e . By Lemma 13, this cycle has length greater than $2k$. Hence π has length greater than $2k-1$, and so

$$|E(\pi)| = \text{dist}_{H \setminus \widehat{F}_e}((u, i), (v, j)) > 2k-1.$$

On the other hand, $\widehat{e} \notin \widehat{F}_e$, so we have $\text{dist}_{\mathcal{B}^{(f)} \setminus \widehat{F}_e}((u, i), (v, j)) = 1$. This implies that H is not an f -DFT $(2k-1)$ -spanner of $\mathcal{B}^{(f)}$. ◀

References

- 1 Reyan Ahmed, Greg Bodwin, Faryad Darabi Sahneh, Keaton Hamm, Mohammad Javad Latifi Jebelli, Stephen Kobourov, and Richard Spence. Graph spanners: A tutorial review. *Computer Science Review*, 37:100253, 2020. doi:10.1016/J.COSREV.2020.100253.
- 2 Ingo Althöfer, Gautam Das, David Dobkin, Deborah Joseph, and José Soares. On sparse spanners of weighted graphs. *Discrete & Computational Geometry*, 9(1):81–100, 1993. doi:10.1007/BF02189308.
- 3 Ahmad Biniiaz, Jean-Lou De Carufel, Anil Maheshwari, and Michiel Smid. Metric and geometric spanners that are resilient to degree-bounded edge faults. *arXiv preprint arXiv:2405.18134*, 2024. doi:10.48550/arXiv.2405.18134.
- 4 Greg Bodwin, Michael Dinitz, Merav Parter, and Virginia Vassilevska Williams. Optimal vertex fault tolerant spanners (for fixed stretch). In *Proceedings of the 29th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1884–1900. Society for Industrial and Applied Mathematics, 2018. doi:10.1137/1.9781611975031.123.

- 5 Greg Bodwin, Michael Dinitz, and Caleb Robelle. Optimal vertex fault-tolerant spanners in polynomial time. In *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 2924–2938. SIAM, 2021. doi:10.1137/1.9781611976465.174.
- 6 Greg Bodwin, Michael Dinitz, and Caleb Robelle. Partially optimal edge fault-tolerant spanners. In *Proceedings of the 2022 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 3272–3286. SIAM, 2022. doi:10.1137/1.9781611977073.129.
- 7 Greg Bodwin, Bernhard Haeupler, and Merav Parter. Fault-tolerant spanners against bounded-degree edge failures: Linearly more faults, almost for free. In *Proceedings of the 2024 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 2609–2642. SIAM, 2024. doi:10.1137/1.9781611977912.93.
- 8 Greg Bodwin and Shyamal Patel. A trivial yet optimal solution to vertex fault tolerant spanners. In *Proceedings of the 2019 ACM Symposium on Principles of Distributed Computing*, pages 541–543, 2019. doi:10.1145/3293611.3331588.
- 9 Prosenjit Bose, Vida Dujmovic, Pat Morin, and Michiel Smid. Robust geometric spanners. *SIAM Journal on Computing*, 42(4):1720–1736, 2013. doi:10.1137/120874473.
- 10 Kevin Buchin, Sarel Har-Peled, and Dániel Oláh. A spanner for the day after. *Discrete & Computational Geometry*, 64(4):1167–1191, 2020. doi:10.1007/S00454-020-00228-6.
- 11 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. doi:10.1007/S00453-013-9779-Y.
- 12 T-H Hubert Chan, Mingfei Li, Li Ning, and Shay Solomon. New doubling spanners: Better and simpler. *SIAM Journal on Computing*, 44(1):37–53, 2015. doi:10.1137/130930984.
- 13 Shiri Chechik, Michael Langberg, David Peleg, and Liam Roditty. Fault tolerant spanners for general graphs. *SIAM Journal on Computing*, 39(7):3403–3423, 2010. doi:10.1137/090758039.
- 14 David Conlon. Extremal numbers of cycles revisited. *The American Mathematical Monthly*, 128(5):464–466, 2021. doi:10.1080/00029890.2021.1886845.
- 15 Artur Czumaj and Hairong Zhao. Fault-tolerant geometric spanners. *Discrete & Computational Geometry*, 32(2):207–230, 2004. doi:10.1007/S00454-004-1121-7.
- 16 Michael Dinitz and Robert Krauthgamer. Fault-tolerant spanners: better and simpler. In *Proceedings of the 30th annual ACM SIGACT-SIGOPS symposium on Principles of distributed computing*, pages 169–178. ACM, 2011. doi:10.1145/1993806.1993830.
- 17 Michael Dinitz and Caleb Robelle. Efficient and simple algorithms for fault-tolerant spanners. In *Proceedings of the 39th Symposium on Principles of Distributed Computing*, pages 493–500, 2020. doi:10.1145/3382734.3405735.
- 18 Paul Erdős. Extremal problems in graph theory. In *Proceedings of the Symposium on Theory of Graphs and its Applications*, pages 29–36, 1963.
- 19 Walter Feit and Graham Higman. The nonexistence of certain generalized polygons. *Journal of Algebra*, 1(2):114–131, 1964.
- 20 Felix Lazebnik, Vasily A Ustimenko, and Andrew J Woldar. A new series of dense graphs of high girth. *Bulletin of the American Mathematical Society*, 32(1):73–79, 1995.
- 21 Hung Le, Shay Solomon, and Cuong Than. Optimal fault-tolerant spanners in euclidean and doubling metrics: Breaking the $\Omega(\log n)$ lightness barrier. In *64th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2023, Santa Cruz, CA, USA, November 6-9, 2023*, pages 77–97. IEEE, 2023. doi:10.1109/FOCS57990.2023.00013.
- 22 Christos Levkopoulos, 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. doi:10.1145/276698.276734.
- 23 Tamás Lukovszki. New results on fault tolerant geometric spanners. In *Workshop on Algorithms and Data Structures*, pages 193–204. Springer, 1999. doi:10.1007/3-540-48447-7_20.
- 24 Giri Narasimhan and Michiel Smid. *Geometric Spanner Networks*. Cambridge University Press, 2007. doi:10.1017/CB09780511546884.

- 25 Merav Parter. Nearly optimal vertex fault-tolerant spanners in optimal time: Sequential, distributed and parallel. In *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing (STOC)*. ACM, 2022.
- 26 Merav Parter and Asaf Petruschka. Preserving distances in faulty colored graphs. *arXiv preprint arXiv:2504.17868*, 2025. doi:10.48550/arXiv.2504.17868.
- 27 Merav Parter, Asaf Petruschka, Shay Sapir, and Elad Tzalik. Parks and recreation: Color fault-tolerant spanners made local. In *Proceedings of the 2025 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 4061–4094. SIAM, 2025. doi:10.1137/1.9781611978322.139.
- 28 Merav Parter and Elad Tzalik. Connectivity certificate against bounded-degree faults: Simpler, better and supporting vertex faults. In *2025 Symposium on Simplicity in Algorithms (SOSA)*, pages 369–377. SIAM, 2025. doi:10.1137/1.9781611978315.28.
- 29 David Peleg and Alejandro A Schäffer. Graph spanners. *Journal of Graph Theory*, 13(1):99–116, 1989. doi:10.1002/JGT.3190130114.
- 30 Asaf Petruschka, Shay Sapir, and Elad Tzalik. Connectivity Labeling in Faulty Colored Graphs. In *38th International Symposium on Distributed Computing (DISC 2024)*, pages 36:1–36:22, 2024.
- 31 Asaf Petruschka, Shay Sapir, and Elad Tzalik. Color fault-tolerant spanners. *ACM Transactions on Algorithms*, 22(1):1–21, 2025. doi:10.1145/3750728.
- 32 Shay Solomon. From hierarchical partitions to hierarchical covers: Optimal fault-tolerant spanners for doubling metrics. In *Proceedings of the forty-sixth annual ACM symposium on Theory of computing*, pages 363–372. ACM, 2014. doi:10.1145/2591796.2591864.
- 33 Tibor Szabó. Even cycles, 2018. Lecture notes, FU Berlin. URL: <http://discretemath.imp.fu-berlin.de/DMIII-2018/notes-evencycles.pdf>.
- 34 Jacques Tits. Sur la trialité et certains groupes qui s’ en déduisent. *Publications Mathématiques de l’IHÉS*, 2:13–60, 1959.
- 35 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. doi:10.1016/0095-8956(91)90097-4.