

Color Fault-Tolerant Spanners

Asaf Petruschka ✉

Weizmann Institute of Science, Rehovot, Israel

Shay Sapir ✉

Weizmann Institute of Science, Rehovot, Israel

Elad Tzalik ✉

Weizmann Institute of Science, Rehovot, Israel

Abstract

We initiate the study of spanners in arbitrarily vertex- or edge-colored graphs (with no “legality” restrictions), that are resilient to failures of *entire color classes*. When a color fails, all vertices/edges of that color crash. An f -color fault-tolerant (f -CFT) t -spanner of an n -vertex colored graph G is a subgraph H that preserves distances up to factor t , even in the presence of at most f color faults. This notion generalizes the well-studied f -vertex/edge fault-tolerant (f -V/EFT) spanners. The size (number of edges) of an f -V/EFT spanner crucially depends on the number f of vertex/edge faults to be tolerated. In the colored variants, even a single color fault can correspond to an unbounded number of vertex/edge faults.

The key conceptual contribution of this work is in showing that the size required by an f -CFT spanner is in fact comparable to its uncolored counterpart, with no dependency on the size of color classes. We provide optimal bounds on the size required by f -CFT $(2k - 1)$ -spanners, as follows:

- When vertices have colors, we show an upper bound of $O(f^{1-1/k}n^{1+1/k})$ edges. This precisely matches the (tight) bounds for $(2k - 1)$ -spanners resilient to f *individual* vertex faults [Bodwin et al., SODA 2018; Bodwin and Patel, PODC 2019].
- For colored edges, we show that $O(fn^{1+1/k})$ edges are always sufficient. Further, we prove this is tight, i.e., we provide an $\Omega(fn^{1+1/k})$ (worst-case) lower bound. The state-of-the-art bounds known for the corresponding uncolored setting of edge faults are (roughly) $\Theta(f^{1/2}n^{1+1/k})$ [Bodwin et al., SODA 2018; Bodwin, Dinitz and Robelle, SODA 2022].
- We also consider a *mixed* model where both vertices and edges are colored. In this case, we show tight $\Theta(f^{2-1/k}n^{1+1/k})$ bounds.

Thus, CFT spanners exhibit an interesting phenomenon: while (individual) edge faults are “easier” than vertex faults, edge-color faults are “harder” than vertex-color faults.

Our upper bounds are based on a generalization of the blocking set technique of [Bodwin and Patel, PODC 2019] for analyzing the (exponential-time) greedy algorithm for FT spanners. We complement them by providing *efficient* constructions of CFT spanners with similar size guarantees, based on the algorithm of [Dinitz and Robelle, PODC 2020].

2012 ACM Subject Classification Theory of computation → Sparsification and spanners

Keywords and phrases Fault tolerance, Graph spanners

Digital Object Identifier 10.4230/LIPIcs.ITCS.2024.88

Related Version *Extended Version*: <http://arxiv.org/abs/2311.08868>

Funding *Asaf Petruschka*: Supported by the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme, grant agreement No. 949083, and by the Israeli Science Foundation (ISF), grant 2084/18.

Shay Sapir: Partially supported by the Israeli Council for Higher Education (CHE) via the Weizmann Data Science Research Center.

Acknowledgements We are grateful to Merav Parter for encouraging this collaboration, and for helpful guidance and discussions. We thank Nathan Wallheimer for useful discussions.



© Asaf Petruschka, Shay Sapir, and Elad Tzalik;
licensed under Creative Commons License CC-BY 4.0

15th Innovations in Theoretical Computer Science Conference (ITCS 2024).

Editor: Venkatesan Guruswami; Article No. 88; pp. 88:1–88:17

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

1 Introduction

Let G be an n -vertex weighted graph. A *spanner* of G is a subgraph that approximately preserves distances. Formally, a subgraph H is a t -spanner of G if $\text{dist}_H(u, v) \leq t \cdot \text{dist}_G(u, v)$ for all $u, v \in V$ (where dist_X denotes the shortest path distance in a graph X). Since their introduction by Peleg and Ullman [32] and Peleg and Schäffer [30], spanners have been found to be extremely useful for a wide variety of applications, including network tasks like routing and synchronization [31, 33, 11, 12, 32], distance estimation [34, 2], preconditioning linear systems [15], and many others.

In real-world scenarios, spanners are frequently employed in systems whose components are susceptible to occasional breakdowns. It is therefore desirable to have spanners possessing resilience to such failures, leading to the notion of *fault-tolerant (FT) spanners*. Two extensively studied types of faults are *edge faults* and *vertex faults*. Spanners in the edge/vertex fault-tolerant (E/VFT) settings are defined as follows:

► **Definition 1 (E/VFT Spanner).** *An f -EFT (VFT) t -spanner of G is a subgraph H such that for every set F of at most f edges (vertices) in G , it holds that $H - F$ is a t -spanner of $G - F$.*

E/VFT spanners have received major interest in recent years; see, e.g., [8, 13, 4, 7, 14, 5, 6, 29] and references therein.

In this work, we introduce and study *color fault tolerant* models. The framework is the following: the edges or vertices are arbitrarily partitioned into classes, or equivalently, associated with colors, and f such *color classes* might fail. Each of the E/VFT settings gets a colored counterpart setting, denoted E/V**C**FT (the **C** stands for “color”). The standard E/VFT settings are obtained as special cases, by giving each edge/vertex a different color.

Faulty color classes have been used to model Shared Risk Resource Groups (SRRG), which arise in many practical contexts such as optical telecommunication networks and multi-layered networks; see [10, 23, 35] and references therein for a detailed discussion. Previous work in this context mainly focused on colored variants of classical graph problems (and their hardness). A notable such problem is *diverse routing*, where the goal is to determine if two (or more) color disjoint paths between two vertices [21, 16, 26] exist (and if so, finding them). In such algorithmic contexts, replacing a dense colored graph by a sparse CFT spanner could potentially help to improve algorithmic performance, similarly to V/EFT spanners for uncolored graphs.

Sparse Connectivity Certificates. Connectivity certificates are subgraph notions that are closely related to FT spanners. Intuitively, a λ -edge (or λ -vertex) *connectivity certificate* serves as concise “proof” to establish the λ -edge/vertex connectivity of a graph G . Formally, it is a subgraph H of G such that every two vertices u, v are λ -edge/vertex connected in G if and only if they are λ -edge/vertex connected in H . The concept of connectivity certificates was presented by Nagamochi and Ibaraki [27], who showed the existence of λ -connectivity certificates with at most $\lambda(n - 1)$ edges for every n -vertex graph. These play a crucial role in various graph algorithms [25, 9, 22, 19, 18, 24]. Every λ -V/EFT $(2k - 1)$ -spanner, for every $1 \leq k < \infty$, is *by definition* also a λ connectivity certificate, as observed e.g. in [28, 29]. Such a spanner with $O(\lambda^{1-1/k} n^{1+1/k})$ edges always exists [7], so choosing $k = \Theta(\log n)$ yields sparse λ -connectivity certificates of $O(\lambda n)$ size.

Due to the min-cut max-flow theorem, edge/vertex-connectivity can be defined equivalently using cuts and flows. In the colored case, this relationship no longer holds: the minimum set of colors whose removal disconnects two vertices can be arbitrarily larger than the

maximum number of color-disjoint paths connecting them [10]. Hence, “color-connectivity certificates” may take two different meanings for cuts and flows. CFT Spanners can then be seen as color-connectivity certificates *in the cut sense*. To avoid confusion, we call the latter λ -CFT-connectivity certificates. Similarly to their uncolored counterparts, such CFT-connectivity certificates may provide useful tools when dealing with colored graphs. As this work focuses on the stronger notion of CFT spanners, (worst-case) optimal results on the size of CFT-connectivity certificates are implied as direct corollaries.

A Mixed Setting. We additionally consider another possible setting, which we call *mixed*, that allows for both edge faults and vertex faults. That is, in the uncolored setting, the faulty-set F is a subset of $V(G) \cup E(G)$ of size $\leq f$. In the colored setting, both edges and vertices have colors (and a color may appear on both vertices and edges). We use the shorthand MFT for the uncolored mixed setting, and MCFT for its colored counterpart. To the best of our knowledge, MFT spanners have not been explicitly defined, but many known results for the E/VFT settings implicitly apply to them.

1.1 Our Contribution

We initiate the study of FT spanners supporting f color faults in colored graphs. A naïve approach for handling color faults is to treat them as many edge/vertex faults. If each color class contains at most Δ edges/vertices, then a Δf -FT spanner for edge/vertex faults is clearly an f -CFT spanner. While this solution might be plausible for small Δ , it is useless for large Δ . As the coloring is arbitrary, it may even be that $\Delta = \Omega(n)$. Thus, the resulting spanner from this naïve approach might have $\Omega(n^2)$ edges even when $f = 1$, i.e., when only a single color fault may occur. Conceptually, the main contribution of our work is showing that in the realm of FT spanners, one can handle f color faults almost as efficiently as f *individual* edge/vertex faults, without any dependence on Δ .

Our main results regarding CFT spanners are:

► **Theorem 2.** *For every $k, f \geq 1$, and every colored n -vertex graph G , there exists an f -CFT $(2k - 1)$ -spanner with the following size bounds:*

1. $O(fn^{1+1/k})$ edges in the edge-colored setting,
2. $O(f^{1-1/k}n^{1+1/k})$ edges in the vertex-colored setting, and
3. $O(f^{2-1/k}n^{1+1/k})$ edges in the mixed setting of both edge and vertex colors.

Moreover, the above bounds are tight assuming the Erdős Girth Conjecture.

Table 1 presents the known bounds for non-colored FT-spanners, compared to the color settings. The bounds of Theorem 2 are based on the *exponential-time* FT greedy algorithm. We further provide polynomial-time constructions of CFT spanners with nearly optimal size guarantees (up to poly(k) factors). We obtain λ -CFT connectivity certificates as immediate corollaries:

► **Theorem 3.** *Every colored n -vertex graph G has a subgraph H such that, for every two vertices u, v , and every set C of at most λ colors, u, v are connected in $H - C$ iff they are connected in $G - C$, where $|E(H)| = O(\lambda n)$ if the graph has either colored edges or colored vertices (but not both), and $|E(H)| = O(\lambda^2 n)$ when both vertices and edges are colored. This is (worst-case) optimal.*

■ **Table 1** Bounds on the worst-case size required by f -FT $(2k - 1)$ -spanners. All lower bounds are conditioned on the Erdős Girth Conjecture [17]. The lower bounds in the ECFT and MCFT settings cannot exceed $\Omega(n^2)$ for simple graphs, but we omit the $\min\{\cdot, n^2\}$ expression for ease of reading. The bounds in the MFT settings are implicit by the cited works.

Setting	Upper bound	Ref.	Lower bound	Ref.
EFT	$O_k(f^{\frac{1}{2}-\frac{1}{2k}} \cdot n^{1+\frac{1}{k}} + fn)$, odd k $O_k(f^{\frac{1}{2}} \cdot n^{1+\frac{1}{k}} + fn)$, even k	[6]	$\Omega(f^{\frac{1}{2}-\frac{1}{2k}} \cdot n^{1+\frac{1}{k}} + fn)$	[4]
ECFT	$O(f \cdot n^{1+\frac{1}{k}})$	New	$\Omega(f \cdot n^{1+\frac{1}{k}})$	New
VFT		[7]		
VCFT	$O(f^{1-\frac{1}{k}} \cdot n^{1+\frac{1}{k}})$	New	$\Omega(f^{1-\frac{1}{k}} \cdot n^{1+\frac{1}{k}})$	[4]
MFT	$O(f^{1-\frac{1}{k}} \cdot n^{1+\frac{1}{k}})$	[7]	$\Omega(f^{1-\frac{1}{k}} \cdot n^{1+\frac{1}{k}})$	[4]
MCFT	$O(f^{2-\frac{1}{k}} \cdot n^{1+\frac{1}{k}})$	New	$\Omega(f^{2-\frac{1}{k}} \cdot n^{1+\frac{1}{k}})$	New

On FT Spanners for Multi-Graphs. Interestingly, the bounds we provide for *simple* colored graphs match the bounds required by E/V/MFT spanners for uncolored *multi-graphs*. All of our results apply both to simple and multi-graphs. This is in contrast to the uncolored setting, where “multi” FT-spanners are denser than “simple” ones. See discussion in Appendix A.

1.2 Technical Overview

Upper Bounds (Section 3). To prove size upper bounds for CFT spanners, we build upon the powerful yet strikingly simple *blocking set* technique of Bodwin and Patel [7]. This is a method for analyzing the spanner H given by the (exponential-time) *FT greedy algorithm*, used in [7] to obtain optimal bounds in the VFT setting. For clarity of exposition, we focus our discussion here only on the corresponding vertex-colored setting.

Bodwin and Patel define a blocking set as a collection of edge-vertex pairs, such that each short ($\leq 2k$) cycle in H contains one of the pairs. On an intuitive level, a graph that admits a small blocking set is close to having high girth, forcing it to be sparse. The FT greedy algorithm naturally induces a blocking set of size $f|E(H)|$, which ultimately yields the upper bound on the size of H (as explained next). The trivial adaption of this algorithm to the CFT settings then leads to the “right” extension of a blocking set: a collection of edge-color pairs such that for every short cycle Σ , at least one pair has both its edge and its color appearing on Σ .

The analysis of $|E(H)|$ is by considering a random subgraph of H induced on a fraction of $\approx 1/f$ vertices, and removing the edges of surviving blocking-set-pairs. The resulting subgraph, denoted \hat{H}_S , has $O(n/f)$ vertices and large girth, and therefore $O((n/f)^{1+1/k})$ edges according to the folklore *Moore Bounds*. On the other hand, direct calculations show that $\Omega(|E(H)|/f^2)$ edges are expected to survive, hence $|E(H)| = O(f^{1-1/k}n^{1+1/k})$. Again, there seems to be a natural adaptation for (vertex) colors: sample a random fraction of $\approx 1/f$ *color classes*. The problem that arises is that even though \hat{H}_S has $O(n/f)$ vertices *in expectation*, it is quite probable to have $\Omega(n)$ vertices (say, if a color class of size $\Omega(n)$ is sampled). More conceptually, the problem is that vertices cannot be sampled independently, preventing us from using concentration bounds. Naïvely carrying out the same analysis with a trivial $O(n)$ bound on the number of sampled vertices yields only an $O(f^2n^{1+1/k})$ bound on the size of the spanner H .

We show that even in the independence-deprived circumstances of the VCFT setting, the same bounds of the VFT setting can be obtained. Our analysis is carefully designed to harness the little independence we do have, between different color classes. It has three main ingredients.

- (1) We first show that effectively, we may assume that greedy spanner H is *color-bipartite*: a bipartite graph where each color class is fully contained in one of the sides.
- (2) To complement this, we use a “symmetric” version of the Moore Bounds: a bipartite graph with sides L, R and girth $> 2k$ has $O(|L||R|^{1/k} + |R||L|^{1/k})$ edges.¹
- (3) Combining the independence between the color classes *from different sides* with the Moore Bounds, we show that the random graph \hat{H}_S has $O((n/f)^{1+1/k})$ edges in expectation.

Lower Bounds (Section 4). Our lower bounds, as most spanner lower bounds, are conditioned on the famous and widely believed Erdős Girth Conjecture [17], implying the existence of an n -vertex graph G_{ER} with $\Omega(n^{1+1/k})$ edges that has no proper subgraph as a $(2k - 1)$ -spanner.

For edge-colored graphs, our strategy is to pack f copies of an n -vertex Erdős graph G_{ER} , denoted G_1, \dots, G_f , on the same set of vertices, with little overlap. We then consider the union graph $\bigcup_i G_i$ (ignoring overlaps), where each edge originating from G_i is given the color i . This graph has $\Omega(fn^{1+1/k})$ edges, and no proper f -ECFT $(2k - 1)$ -spanner.

In the mixed-colored setting, we “encode” f^2 disjoint f -ECFT lower bound instances in a graph with fn properly colored vertices, such that each instance can be isolated by failing $2f$ vertex-colors. Thus, this graph has no proper $3f$ -MCFT spanner.

Efficient Algorithms (Section 5). We show efficient CFT spanner constructions based on the technique of Dinitz and Robelle [14]. Their approach was to modify the greedy algorithm in the E/VFT settings to have polynomial running time, while ensuring it still induces a small blocking set, only larger by an $O(k)$ factor than that of the “vanilla” FT greedy algorithm. Thus, the size analysis based on blocking sets still applies. Our (existential) upper bounds generalize this last component to colored settings. We then augment the [14] algorithm to produce small *colored* blocking sets, and thus achieve efficient constructions.

Remark on Optimality Beyond the Girth Conjecture. Let $\gamma(n, 2k)$ denote the maximum number of edges in an n -vertex graph of girth $> 2k$. The Moore Bounds show that $\gamma(n, 2k) = O(n^{1+1/k})$, and Erdős’ Girth Conjecture is that this is tight up to a constant factor. Our lower bounds can be formulated purely in terms of the γ function: $\Omega(f \cdot \gamma(n, 2k))$ edges in the ECFT setting, $\Omega(f^2 \cdot \gamma(n/f, 2k))$ edges in VCFT setting, and $\Omega(f^3 \cdot \gamma(n/f, 2k))$ edges in MCFT setting. As for the upper bounds: In the ECFT case, the upper bound matches the lower bound also in terms of γ . In the VCFT and MCFT settings, the upper bounds are matching assuming only much weaker conditions on γ than the Girth Conjecture; it is enough to assume that $\gamma(n, 2k) = \Theta(n \cdot \beta(n, 2k))$, where $\beta(\cdot, 2k)$ is *concave* (the Girth Conjecture is that $\beta(n, 2k) = n^{1/k}$).

2 Preliminaries: A Unified FT Framework

We now present a unified formal framework for all six E/V/M(C)FT settings, to help avoid tedious repetitions later on.

¹ Stronger bipartite Moore Bounds exist [20], but this weaker bound suffices for our needs.

Henceforth, assume every edge/vertex $x \in E(G) \cup V(G)$ is associated with a color $c(x) \in \mathcal{C}$, where \mathcal{C} is some finite set called the *color palette*. The coloring is arbitrary, and there are no “legality” restrictions. For example, two neighboring vertices may share the same color. It is also allowed that the same color appears both on edges and on vertices. For $c \in \mathcal{C}$, we denote the sets of c -colored edges and vertices by $E_c(G) = \{e \in E(G) \mid c(e) = c\}$ and $V_c(G) = \{v \in V(G) \mid c(v) = c\}$. The edges of G are weighted, where $w(e) = w(u, v)$ denotes the weight of edge $e = \{u, v\}$.

The *fault universe* \mathcal{F} is defined separately for each of the fault-tolerant settings:

Setting	EFT	VFT	MFT	ECFT	VCFT	MCFT
$\mathcal{F} =$	$E(G)$	$V(G)$	$E(G) \cup V(G)$	$\{E_c(G)\}_{c \in \mathcal{C}}$	$\{V_c(G)\}_{c \in \mathcal{C}}$	$\{E_c(G) \cup V_c(G)\}_{c \in \mathcal{C}}$

We say that an edge $e = \{u, v\} \in E(G)$ is *damaged* by $x \in \mathcal{F}$ if the failure of x causes e to fail. Formally, x damages e if: $x = e$ in EFT, $x \in \{u, v\}$ in VFT, $x \in \{e, u, v\}$ in MFT, $e \in x$ in ECFT, $\{u, v\} \cap x \neq \emptyset$ in VCFT, and $\{e, u, v\} \cap x \neq \emptyset$ in MCFT. Note that in the colored settings, there is a one-to-one correspondence between the fault universe \mathcal{F} and the set of colors \mathcal{C} . So, when $S \subseteq \mathcal{F}$, we say that S contains the color c , and abuse notation by writing $c \in S$, when S contains the element corresponding to c from the fault universe \mathcal{F} (which depends on the setting). We denote by $G - S$ the subgraph of G obtained by deleting all edges damaged by some $x \in S$, and by $G[S] := G - (\mathcal{F} - S)$, which is the subgraph obtained by keeping all edges that can only be damaged by faults from S .

We now formally state the general FT spanner definition, applying to all settings:

► **Definition 4 (General FT Spanner).** *An f -FT t -spanner of G with respect to a fault universe \mathcal{F} (defined by the given setting) is a subgraph H such that for every $F \subseteq \mathcal{F}$ with $|F| \leq f$, it holds that $H - F$ is a t -spanner of $G - F$.*

3 Upper Bounds

The FT variant of the greedy spanner algorithm of Alth’ofer et al. [1], presented as Algorithm 1, has been central in the study of spanners in general, and FT spanners in particular.

■ **Algorithm 1** FTGreedySpanner(G, k, f).

```

1:  $H \leftarrow (V(G), \emptyset)$ 
2: for each  $e = \{u, v\} \in E(G)$  in increasing order of weight do
3:   if there is  $F \subseteq \mathcal{F}$ ,  $|F| \leq f$ , not damaging  $e$ , such that  $\text{dist}_{H-F}(u, v) > (2k - 1)w(e)$ 
   then
4:      $H \leftarrow H \cup \{e\}$ 
5: return  $H$ 

```

Aside from its immediate correctness, the greedy approach has been proven to produce strong and often optimal upper bounds on the (worst-case) size of E/VFT spanners (e.g., see [7, 6]), making it an “immediate suspect” to investigate for getting upper bounds on the size of CFT spanners as well. The major downside of the naïve FT greedy algorithm is its exponential-in- f running time, which makes it intractable for many algorithmic applications. However, it is often the case that modified polynomial-time greedy approaches (with suitably modified analysis) can be devised to produce very close size bounds [5, 14].

In this section, we provide existential size bounds on CFT spanners via the naive FT greedy algorithm, using the *blocking set* technique of Bodwin and Patel [7]. We show:

► **Theorem 5.** *The number of edges m in the graph H constructed by the FT greedy algorithm (Algorithm 1) can be bounded as follows:*

1. *In the ECFT setting, $m = O(fn^{1+1/k})$.*
2. *In the VCFT setting, $m = O(f^{1-1/k}n^{1+1/k})$.*
3. *In the MCFT setting, $m = O(f^{2-1/k}n^{1+1/k})$.*

3.1 Blocking Sets and Random Blocked Subgraphs

Denote by H the output of Algorithm 1, and by m its number of edges. The size analysis of H hinges on a generalized definition of blocking sets, capturing all six E/V/M(C)FT settings.

► **Definition 6** (Blocking Set). *A $2k$ -blocking set of H is a subset $B \subseteq E(H) \times \mathcal{F}$ that satisfies the following conditions:*

- (i) *If $(e, x) \in B$, then e is not damaged by x .*
- (ii) *For every cycle Σ in H on $\leq 2k$ edges, there exist $(e, x) \in B$ such that $e \in \Sigma$, and some (other) edge of Σ is damaged by x . Abusing notation, we denote this condition by $(e, x) \in B \cap \Sigma$.*

The same arguments as in [7] imply the existence of small blocking set for H :

► **Lemma 7** (Modification of [7, Lemma 3]). *There is a $2k$ -blocking set B of H such that $|B| \leq fm$. Moreover, for every edge $e \in E(H)$, there are at most f pairs of the form (e, x) in B .*

Proof. For every $e = \{u, v\} \in E(H)$, let $F_e \subseteq \mathcal{F}$, $|F_e| \leq f$ be the set of faults not damaging e that caused e 's insertion. I.e., at the time e was considered by Algorithm 1, it was the case that $\text{dist}_{H-F_e}(u, v) > (2k-1)w(e)$. Define

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

Every edge $e \in E(H)$ participates in at most f pairs $(e, x) \in B$, and condition (i) of Definition 6 holds. We show condition (ii). Let Σ be a cycle on $\leq 2k$ edges in H . Consider the heaviest edge $e = \{u, v\}$ on Σ .² Then $\Sigma - e$ was already present in H before e was added, forming a u - v path of length $\leq (2k-1)w(e)$. But, before adding e , we had $\text{dist}_{H-F_e}(u, v) > (2k-1)w(e)$, so some $x \in F_e$ must damage an edge of $\Sigma - e$. Thus, $(e, x) \in B \cap \Sigma$. ◀

► **Remark 8.** Note that Lemma 7 stays true also for subgraphs of H , by deleting the pairs from B that correspond to the edges missing in the subgraph.

To analyze the size of H , we consider the expected number of edges in a randomly-formed subgraph \hat{H}_S of H , which we call a *random blocked subgraph*. Informally, we bound this expectation from below by $|E(H)|/\text{poly}(f)$, and from above by a function of n, f alone. Rearranging yields the final bounds.

² Ties are broken as in Algorithm 1. That is, e is edge in Σ that appeared last in the ordering by which Algorithm 1 went over the edges of G .

► **Definition 9** (*p*-Random Blocked Subgraph). Let $S \subseteq \mathcal{F}$ be a random subset obtained by sampling each $x \in \mathcal{F}$ into S independently with probability p . Consider $H[S]$, the subgraph of H obtained by removing all edges that are damaged by $\mathcal{F} - S$. Essentially, the p -random blocked subgraph \hat{H}_S is obtained from $H[S]$ by removing the edges of surviving elements of the blocking set. Formally:

$$\begin{aligned} B_S &= \{(e, x) \in B \mid e \in E(H[S]), x \in S\}, \\ E(B_S) &= \{e \mid (e, x) \in B_S \text{ for some } x \in S\}, \\ \hat{H}_S &= H[S] - E(B_S). \end{aligned}$$

The number of edges in \hat{H}_S is denoted by \hat{m} .

The point of this process is the following:

► **Observation 10.** A p -random blocked subgraph \hat{H}_S has girth $\geq 2k + 1$ with probability 1.

Proof. Suppose Σ is a cycle of $\leq 2k$ edges in $H[S]$ (and thus also in H). As B is a blocking set, there is some $(e, x) \in B \cap \Sigma$. Since Σ is sampled into $H[S]$, then $e \in E(H[S])$ and $x \in S$, showing that $(e, x) \in B_S$. Thus, e is removed in \hat{H}_S , so Σ does not survive in \hat{H}_S . ◀

3.2 The ECFT Setting

We start by analyzing the ECFT setting, in which a trivial analysis yields optimal results. It serves as a good warm up for the other settings. The analysis uses the folklore *Moore Bounds*:

► **Lemma 11** (Moore Bounds). An n -vertex graph with girth $\geq 2k + 1$ has $O(n^{1+1/k})$ edges.

Now, let \hat{H}_S be a $1/(2f)$ -random blocked subgraph. Consider an edge $e \in E(H)$ of the FT greedy spanner H . Clearly, $e \in H[S]$ iff $c(e) \in S$, which happens with probability $1/(2f)$. Next, consider a pair $(e, c) \in B$. As c does not damage e , meaning $c(e) \neq c$, then $(e, c) \in B_S$ iff both c and $c(e)$ are sampled, which happens with probability $1/(2f)^2$. Thus,

$$\begin{aligned} \mathbb{E}[\hat{m}] &= \mathbb{E}\left[|E(H[S]) - E(B_S)|\right] \\ &\geq \mathbb{E}\left[|E(H[S])| - |B_S|\right] = \frac{m}{2f} - \frac{|B|}{4f^2} \\ &\geq \frac{m}{2f} - \frac{fm}{4f^2} = \frac{m}{4f}. \end{aligned} \quad (\text{Lemma 7})$$

On the other hand, by Observation 10, \hat{H}_S has girth $\geq 2k + 1$, and thus by the Moore bounds (Lemma 11), $\hat{m} = O(n^{1+1/k})$. Hence,

$$\frac{m}{4f} \leq \mathbb{E}[\hat{m}] = O(n^{1+1/k}) \implies m = O(fn^{1+1/k}).$$

This concludes the proof of Theorem 5(1).

► **Remark 12.** The proof works seamlessly for multi-graphs, and the same holds for all other proofs in this section.

3.3 The VCFT Setting

We now analyze the size of the spanner produced by the greedy algorithm (Algorithm 1) in the VCFT setting. As explained in Section 1.2, a naive adaptation of the proof in the ECFT setting gives an $O(f^2 n^{1+1/k})$ bound. To get our tight bound here, we go in a different route.

3.3.1 Bipartite Moore Bounds

Our analysis for the VCFT setting relies on Moore Bounds for bipartite graphs. For our purposes, the strongest known bounds of Hoory [20] are not required, and the following weaker but simpler-to-prove bound is sufficient.

► **Lemma 13** (Weak Bipartite Moore Bounds). *Let G be a bipartite graph with bipartition $V = L \cup R$, $|L| = x$, $|R| = y$. Suppose G has girth $\geq 2k + 1$. Then $|E(G)| = O(x^{1/k}y + xy^{1/k})$.*

Proof. Suppose $x \leq y$. We will prove that $|E(G)| = O(x^{1/k}y)$, which implies the lemma.

Let $S \subseteq R$ be a uniformly random subset with $|S| = x$. The expected number of edges in the induced subgraph $G_S = G[L \cup S]$ is

$$\sum_{e \in E(G)} \Pr[e \in L \times S] = |E(G)| \cdot \frac{x}{y}.$$

Choose some S exceeding the expectation, i.e., such that

$$|E(G_S)| \geq |E(G)| \cdot \frac{x}{y}. \quad (1)$$

As G_S also has girth $\geq 2k + 1$, the Moore Bounds (Lemma 11) imply

$$|E(G_S)| \leq O(|V(G_S)|^{1+1/k}) = O(x^{1+1/k}). \quad (2)$$

Combining Eqns. (1) and (2) yields $|E(G)| = O(x^{1/k}y)$, which concludes the proof. ◀

3.3.2 Setting Up the Stage

As before, let H denote the output of the greedy spanner (Algorithm 1), now in the VCFT setting. Again, denote $m = |E(H)|$. We now perform a “clean-up” phase for the sake of analysis, where we neglect some fraction of the edges of H , as follows.

First, we claim that H contains at most $O(n^{1+1/k})$ *monochromatic* edges. Let $E_{mono} := \{\{u, v\} \in E(H) \mid c(u) = c(v)\}$. Note that E_{mono} cannot contain cycles of $\leq 2k$ edges, as such cycles must contain some pair $(e = \{u, v\}, c)$ of the blocking set B , but $|\{c, c(u), c(v)\}| \geq 2$. Hence, by the Moore Bounds (Lemma 11), $|E_{mono}| = O(n^{1+1/k})$.

From now on, we assume that H has no monochromatic edges. Namely, H is henceforth $H - E_{mono}$. This is justified as we can delete the $O(n^{1+1/k})$ edges of E_{mono} (which is less than the upper bound we prove), while still maintaining a small blocking set (recall Remark 8).

Next, for our analysis, we would like to “make” H into a bipartite graph with the special property that all the vertices of any specific color class always lie on the same side. To obtain this, we observe the following:

► **Observation 14.** *There is a partition $V(H) = L \cup R$, such that:*

- (i) *For every color c , all vertices with color c are found only in one of L, R .*
- (ii) *The number of edges between L and R is at least $|E(H)|/2$.*

Proof. Independently for every color c , with probability $1/2$, all the c -colored vertices go to L , and otherwise they go to R . Each edge is split between L, R with probability $1/2$ (as it is not monochromatic). Thus, the *expected* number of L -to- R edges is $|E(H)|/2$, and some partition achieves $\geq |E(H)|/2$ many L -to- R edges. ◀

Let H' be the subgraph of H obtained by deleting all edges that violate the bipartition. Clearly, by (ii) of the last observation, it is enough to prove that $|E(H')| = O(f^{1-1/k}n^{1+1/k})$. So from now on (justified again by Remark 8), we assume that H has only L -to- R edges.

3.3.3 Bounding the Size of H

Again, denote $m = |E(H)|$. The following is the key lemma of the analysis:

► **Lemma 15.** *Let \hat{H}_S be a p -random blocked subgraph. Then $\mathbb{E}[\hat{m}] = O((pn)^{1+1/k})$.*

Proof. Let $L_S = \{v \in L \mid c(v) \in S\}$, i.e., the vertices from L whose colors were sampled. Define R_S similarly. The crux of the special color-bipartition is ensuring that $|L_S|$ and $|R_S|$ are *independent* random variables. Recall that by Observation 10, \hat{H}_S has girth $\geq 2k + 1$. Using the bipartite Moore Bounds (Lemma 13), we now obtain

$$\begin{aligned} \mathbb{E}[\hat{m}] &\leq O\left(\mathbb{E}\left[|L_S|^{1/k} \cdot |R_S| + |L_S| \cdot |R_S|^{1/k}\right]\right) \\ &= O\left(\mathbb{E}\left[|L_S|^{1/k}\right] \cdot \mathbb{E}[|R_S|] + \mathbb{E}[|L_S|] \cdot \mathbb{E}\left[|R_S|^{1/k}\right]\right) && \text{(Independence)} \\ &\leq O\left(\mathbb{E}[|L_S|]^{1/k} \cdot \mathbb{E}[|R_S|] + \mathbb{E}[|L_S|] \cdot \mathbb{E}[|R_S|]^{1/k}\right) && \text{(Jensen's inequality)} \\ &= O\left((p|L|)^{1/k} \cdot p|R| + p|L| \cdot (p|R|)^{1/k}\right) \\ &\leq O\left((pn)^{1+1/k}\right) \end{aligned}$$

as required. ◀

We are now ready to complete the analysis. Let \hat{H}_S be a p -random blocked subgraph with $p = 1/(2f)$. Then,

$$\mathbb{E}[\hat{m}] \geq \mathbb{E}[|E(H[S])| - |B_S|] = \frac{m}{4f^2} - \frac{|B|}{8f^3} \geq \frac{m}{8f^2},$$

where the last inequality is by Lemma 7. Combining with Lemma 15, we now obtain

$$\frac{m}{8f^2} \leq \mathbb{E}[\hat{m}] = O\left(\left(\frac{n}{f}\right)^{1+1/k}\right) \implies m = O(f^{1-1/k} n^{1+1/k}).$$

Thus, we have proved Theorem 5(2).

3.4 The MCFT Setting

For the mixed setting, we suitably mix arguments from the ECFT and VCFT settings. We use the same paradigm of the VCFT proof: “clean-up”, color-partition, and applying bipartite Moore Bounds. The only step of the paradigm that needs to be adapted is the “clean-up” phase of monochromatic edges, which we now prove:

► **Lemma 16.** *H has at most $O(fn^{1+1/k})$ edges $e = \{u, v\} \in E(H)$ with $c(u) = c(v)$.*

Proof. Let $E' = \{e = \{u, v\} \in E(H) \mid c(u) = c(v)\}$, and let $H' = (V(H), E')$. I.e., H' is the subgraph with only edges that have endpoints of the same color. Run the greedy Algorithm 1, for computing f -ECFT $(2k - 1)$ -spanner, with H' being the *input*. We make sure to use the same ordering of the edges as was used by the MCFT greedy algorithm that generated H (i.e., ties are broken in the same way). Let H'' be the output. We claim that $H'' = H'$. This ends the proof, as we can then apply Theorem 5(1) to get $|E'| = |E(H'')| = O(fn^{1+1/k})$.

Suppose towards a contradiction that some edge $e = \{u, v\}$ is present in H' but was discarded from H'' . Denote $c = c(u) = c(v)$.

Let H_e be the state of H right before the MCFT greedy spanner considers the edge e . Let F_e be the set of colors that caused it to insert e into H . That is, F_e is a set of at most f colors that do not damage e (in particular $c \notin F_e$), such that in H_e , every u - v path that avoids the colors in F_e has weight $> (2k - 1)w(e)$.

Now, let H_e'' be the state of H'' right before the $\underline{\text{ECFT}}$ greedy algorithm considered e . Since e was discarded, there is a u - v path P in H_e'' of weight $\leq (2k-1)w(e)$ that avoided edges with colors from F_e . As H'' has only edges with the same color on both endpoints, and this path only uses edges with c -colored endpoints, P avoids *vertices* with colors from F_e . But, as we insisted on identical ordering when running the algorithms, H_e'' is a subgraph of H_e , so P exists also in H_e – contradiction. \blacktriangleleft

By the above lemma (and Remark 8), we can assume H has no edges with endpoints of the same color. For such H , we can apply the bipartitioning procedure as in the VCFT setting (ignoring the edge colors), and restrict the analysis to the resulting graph. By applying Lemma 15 on the p -random blocked subgraph \hat{H}_S , we conclude that $\mathbb{E}[\hat{m}] = O((pn)^{1+1/k})$.

We set $p = 1/(2f)$. Our final step, as before, is to give a lower bound on $\mathbb{E}[\hat{m}]$. Let

$$E_2 = \{e = \{u, v\} \in E(H) \mid |\{c(e), c(u), c(v)\}| = 2\}$$

$$E_3 = \{e = \{u, v\} \in E(H) \mid |\{c(e), c(u), c(v)\}| = 3\}$$

As every edge has two different colors on its endpoints, $E(H) = E_2 \cup E_3$. Next, let $B_2 = \{(e, x) \in B \mid e \in E_2\}$ and $B_3 = B - B_2$. By Lemma 7, $|B_2| \leq f|E_2|$ and $|B_3| \leq f|E_3|$. We now calculate:

$$\begin{aligned} \mathbb{E}[\hat{m}] &\geq \mathbb{E}[|E(H[S])| - |B_S|] \\ &= \left(\frac{|E_2|}{4f^2} - \frac{|B_2|}{8f^3}\right) + \left(\frac{|E_3|}{8f^3} - \frac{|B_3|}{16f^4}\right) \\ &\geq \frac{|E_2|}{8f^2} + \frac{|E_3|}{16f^3} \geq \frac{m}{16f^3}. \end{aligned}$$

Concluding, we get

$$\frac{m}{16f^3} \leq \mathbb{E}[\hat{m}] = O\left(\left(\frac{n}{f}\right)^{1+1/k}\right) \implies m = O(f^{2-1/k}n^{1+1/k}),$$

which proves Theorem 5(3).

4 Lower Bounds

Our lower bounds are conditioned on the famous *Erdős Girth Conjecture* [17]:

► **Conjecture 17** (Erdős Girth Conjecture). *For every integers $n, k \geq 1$, there exists an n -vertex graph with girth at least $2k+2$ and $\Omega(n^{1+1/k})$ edges.*

The Girth Conjecture is widely believed and commonly used as a basis for lower bounds on spanners, and particularly for the lower bounds of [4] on V/EFT-spanners.

We start by proving a lower bound in the edge-colored case for *simple* graphs (see Appendix A for a discussion on multi-graphs).

► **Theorem 18.** *Let $n, k, f \geq 1$ be integers. Assuming the Girth Conjecture, there exists a simple graph G with n vertices and colored edges, such that every f -ECFT $(2k-1)$ -spanner of G must have $\Omega(\min\{fn^{1+1/k}, n^2\})$ edges.*

Proof. Without loss of generality, we assume $f \leq \frac{1}{8}n^{1-1/k}$. If f is larger, use the graph constructed for $\frac{1}{8}n^{1-1/k}$ color faults, which already gives the desired bound.

Let G_{ER} be the graph from the Girth Conjecture with n vertices, i.e., it has girth $\geq 2k+2$ and $\alpha n^{1+1/k}$ edges, where $\alpha > 0$ is a fixed constant. Without loss of generality, we may

88:12 Color Fault-Tolerant Spanners

assume $\alpha \leq 1$ (otherwise, delete edges from G_{ER}). We prove that for every $1 \leq i \leq f$, there exists a *simple* n -vertex edge-colored graph G_i , with color palette $[i]$, such that the following conditions hold:

- (i) The number of edges in G_i satisfies $i \cdot \frac{\alpha}{2} \cdot n^{1+1/k} \leq |E(G_i)| \leq i \cdot \alpha \cdot n^{1+1/k}$.
- (ii) For every color $c \in [i]$, there is no cycle of $\leq 2k + 1$ c -colored edges in G_i .

The desired graph for the theorem is G_f , as it has $\Omega(fn^{1+1/k})$ edges by (i), and no proper subgraph which is an f color fault-tolerant $(2k - 1)$ -spanner by (ii).

The proof is by induction on i . The base case is just $G_1 = G_{ER}$ with all edges having the same color 1. The induction step constructs G_{i+1} from G_i , as follows. Consider a uniformly random bijection (or permutation) $\pi : V(G_{ER}) \rightarrow V(G_i)$. For every edge $\{u, v\} \in E(G_{ER})$,

$$\Pr [\{\pi(u), \pi(v)\} \in E(G_i)] = \frac{|E(G_i)|}{\binom{n}{2}} \leq \frac{i \cdot \alpha n^{1+1/k}}{\frac{1}{4}n^2} \leq \frac{\frac{1}{8}n^{1-1/k} \cdot \alpha n^{1+1/k}}{\frac{1}{4}n^2} \leq \frac{1}{2}.$$

Thus, in expectation, at least half of the edges in G_{ER} are mapped by π to *non-edges* in G_i . Choose π such that this last property holds. Define G_{i+1} by adding these edges to G_i , and color them with the new color $i + 1$. The resulting graph G_{i+1} is indeed simple, and properties (i) and (ii) follow immediately by the construction and the induction hypothesis. \blacktriangleleft

For the sake of completeness, we note that the lower bound of Bodwin et al. [4] for VFT spanners also applies to the more general VCFT setting.

► **Theorem 19** ([4]). *Let $n, k, f \geq 1$ be integers, $f \leq n$. Assuming the Girth Conjecture, there exists vertex-colored simple n -vertex graph G , such that every f - $V(C)$ FT $(2k - 1)$ -spanner of G must have $\Omega(f^{1-1/k}n^{1+1/k})$ edges.*

Finally, we give the lower bound in the MCFT setting:

► **Theorem 20.** *Let $n, k, f \geq 1$ be integers. Assuming the Girth Conjecture, there exists a simple graph G with n vertices, where both edges and vertices are colored, such that every f -MCFT $(2k - 1)$ -spanner of G must have $\Omega(\min\{f^{2-1/k}n^{1+1/k}, n^2\})$ edges.*

Proof. Let G_{EC} be the edge-colored n -vertex graph of Theorem 18, with $\Omega(\min\{fn^{1+1/k}, n^2\})$ edges and no proper f -ECFT $(2k - 1)$ -spanner. Denote its color palette by \mathcal{C}_E . We may also assume that G_{EC} is bipartite, as every graph can be made bipartite by deleting at most half of its edges.³ Let $V(G_{EC}) = L \cup R$ be the partition to left and right sides.

We construct a graph G with both edge and vertex colors, as follows. The vertex set is

$$V(G) = (L \times [f]) \cup (R \times [f]),$$

i.e., there are f copies of each vertex from $V(G) = L \cup R$. There are two sets of new colors for the vertices (disjoint from each other and from \mathcal{C}_E):

$$\mathcal{C}_L = \{\ell_1, \dots, \ell_f\} \quad \text{and} \quad \mathcal{C}_R = \{r_1, \dots, r_f\}.$$

A vertex $(u, i) \in L \times [f]$ gets the color ℓ_i . Similarly, a vertex $(v, j) \in R \times [f]$ gets the color r_j . We then place a copy of G_{EC} on $(L \times \{i\}) \cup (R \times \{j\})$ for every $i, j \in [f]$. Formally, the edge set is defined by

$$E(G) = \{\{(u, i), (v, j)\} \mid i, j \in [f], u \in L, v \in R, \{u, v\} \in E(G_{EC})\}.$$

³ This can be seen by considering a random bipartition.

Every edge $\{(u, i), (v, j)\} \in E(G)$ is assigned the color that is given in G_{EC} to the corresponding edge $\{u, v\}$ (from the palette \mathcal{C}_E). Note that the copies are edge-disjoint, so the resulting G is simple.

We now show that G has no $3f$ -MCFT $(2k-1)$ -spanner but itself. Indeed, suppose H is such a spanner. Choose any $i, j \in [f]$. Let $F_L^i = \mathcal{C}_L - \{\ell_i\}$ and $F_R^j = \mathcal{C}_R - \{r_j\}$. Consider the graph $H - (F_L^i \cup F_R^j)$. As $|F_L^i \cup F_R^j| \leq 2f$, this graph must be an f -MCFT (and, as $(\mathcal{C}_L \cup \mathcal{C}_R) \cap \mathcal{C}_E = \emptyset$, also f -ECFT) $(2k-1)$ -spanner of $G - (F_L^i \cup F_R^j)$. However, $G - (F_L^i \cup F_R^j)$ is just the copy of G_{EC} on $(L \times \{i\}) \cup (R \times \{j\})$, which has no such spanner but itself. Therefore, each of the copies of G_{EC} must be contained in H , hence $H = G$ as claimed.

Finally, we have that $|V(G)| = f|V(G_{EC})| = fn$ and

$$|E(G)| = f^2|E(G_{EC})| = f^2 \cdot \Omega(\min\{fn^{1+1/k}, n^2\}) = \Omega(\min\{f^{2-1/k}(fn)^{1+1/k}, (fn)^2\}).$$

The result follows by rescaling the proof: $n \leftarrow fn$ and $f \leftarrow 3f$. \blacktriangleleft

5 Efficient Constructions

We now show an efficient (polynomial time) construction yielding E/V/MCFT spanners whose sizes match the bounds of Section 3 up to a $\text{poly}(k)$ factor:

► **Theorem 21.** *There is a polynomial time algorithm for constructing f -CFT $(2k-1)$ -spanners, yielding spanner with size s such that:*

1. $s = O(kfn^{1+1/k})$ in the ECFT setting,
2. $s = O(kf^{1-1/k}n^{1+1/k})$ in the VCFT setting, and
3. $s = O(k^2f^{2-1/k}n^{1+1/k})$ in the MCFT setting.

The running time is $O(mf \cdot s)$.

This is by modifying the “Modified FT Greedy Spanner” algorithm of Dinitz and Robelle [14] to the colored settings. As in [14], we start by considering the case where G is *unweighted*. Assume this from now on; we discuss how weights are handled at the end of this section.

The High-Level Idea. Consider the beginning of some iteration of the FT greedy algorithm (Algorithm 1), which examines $e = \{u, v\} \in E(G)$. The intractable part of the algorithm is trying to find a fault-set F_e of size $\leq f$ such that $\text{dist}_{H-F_e}(u, v) > 2k-1$. If no such set is found, e can clearly be discarded. In the complementary case, e is added to H , and the set F_e serves to account for $\leq f$ pairs in the $2k$ -blocking set constructed by Lemma 7, which ultimately yields the size bound on the final H .

Dinitz and Robelle’s idea is replacing this by tractable *approximate* subroutine, that settles for finding a set \tilde{F}_e with size $O(kf)$, and ensures that e can be discarded if such a set is not found. This subroutine is based on *another greedy approach*, of computing a maximal set of “disjoint” (in some sense) u - v paths of length $\leq 2k-1$ in H . The “blame set” \tilde{F}_e is found if the maximal set of paths contains at most f paths, in which case e is inserted to the spanner. If there are more than f “disjoint” paths, e can be safely discarded. Dinitz and Robelle [14] consider the VFT setting, where “disjoint” means internally vertex-disjoint, and the set \tilde{F}_e contains the internal vertices of those paths. This approach can easily be adapted to the colored setting, as we now explain.

Subroutine for Deciding on an Edge. The subroutine works in $f+1$ iterations. We preserve the invariant that if we completed iteration i without halting, then we have i u - v paths $\{P_j\}_{j=1}^i$ in H , each of length $\leq 2k-1$. Furthermore, let \mathcal{C}_j be the set of colors appearing

on P_j that do not damage e . Then the color sets $\{\mathcal{C}_j\}_{j=1}^i$ are mutually disjoint. To execute iteration i , we compute the shortest u - v path P_i in $H - \bigcup_{j=1}^{i-1} \mathcal{C}_j$. If P_i is of length $\geq 2k$, we halt by deciding to include e , and “blame” the set $\tilde{F}_e = \bigcup_{j=1}^{i-1} \mathcal{C}_j$ for this decision. Crucially, we can bound $|\tilde{F}_e| \leq 4kf$ (as a path of length $\leq 2k - 1$ can have at most $(2k - 1) + 2k$ vertices and edges, and thus at most $4k$ different colors). Otherwise, if P_i has length $\leq 2k - 1$, we add P_i and continue to the next iteration. If the last iteration $f + 1$ was completed, we have found $f + 1$ u - v paths of length $\leq 2k - 1$, of which every pair may share only the colors that damage e . This means that no set of at most $\leq f$ color faults (that does not damage e itself) can cause the distance between u, v in H to be $> 2k - 1$. In this case, we call e *replaceable*, and we safely discard it while ensuring the correctness of the final output spanner.

The pseudo-code for the modified FT greedy algorithm using this subroutine appears as Algorithm 2. The running time analysis is identical to [14].

■ **Algorithm 2** ModifiedCFTGreedySpanner(G, k, f).

(Assuming G is unweighted)

-
- 1: $H \leftarrow (V(G), \emptyset)$
 - 2: **for** each $e = \{u, v\} \in E(G)$ in arbitrary order **do**
 - 3: **if** e is not replaceable because of color set \tilde{F}_e **then**
 - 4: $H \leftarrow H \cup \{e\}$
 - 5: **return** H
-

Size Analysis. The size analysis requires only minor modifications compared to the one for the (exponential-time) FT greedy algorithm, presented in Section 3. We define

$$\tilde{B} = \{(e, c) \mid e \in E(H), c \in \tilde{F}_e\}.$$

Since right before an edge $e = \{u, v\}$ is inserted into H we have that $\text{dist}_{H - \tilde{F}_e}(u, v) > 2k - 1$, then by the same proof as in Lemma 7, \tilde{B} is a $2k$ -blocking set. We can now repeat the analysis via random blocked subgraphs, with the following modifications:

- (1) The definition of a p -random blocked subgraph (Definition 9) uses \tilde{B} instead of B .
- (2) We set $p = \frac{1}{8kf}$ instead of $\frac{1}{2f}$, to account for the fact that $|\tilde{B}| \leq 4kfm$, instead of $|B| \leq fm$.

The resulting bounds on the size of the spanner are then bigger than the ones in Theorem 5 by a factor of $O(k)$ in the ECFT and the VCFT settings, and by a factor of $O(k^2)$ in the MCFT setting, proving Theorem 21 for the unweighted case.

Handling Weights. The way [14] handle a weighted graph G works “as is” also in our case. It is surprisingly simple: Run the *unweighted* Algorithm 2 on G , (i.e., ignoring the weights), but with the “correct” increasing order of weight. The correctness is by the following observation: If $e = \{u, v\} \in G$ was found replaceable, then at that time, for every set F of at most $\leq f$ faults (not damaging e), $H - F$ contains a u - v path of at most $2k - 1$ edges. All these edges precede e in the ordering, so the total weight of the path is $\leq (2k - 1)w(e)$.

6 Conclusion

In this work, we introduced the concept of color fault-tolerant (CFT) spanners, which generalize vertex/edge fault-tolerant spanners. We give tight bounds on the size required by CFT spanners, and polynomial time constructions of such spanners with nearly optimal size guarantees. We suggest a few possible directions for future research:

- Can we provide faster algorithms for CFT spanners with size guarantees close to those of the FT greedy algorithm? See e.g. the construction by Parter [29] in the VFT setting, based on the Baswana-Sen algorithm [3].
- As noted, the min-cut max-flow theorem does not hold in the colored setting. Hence, one may consider investigating a different notion from the “color cuts” of CFT-connectivity certificates, which pertains to “colorful flow”. That is, a subgraph H of G such that every two vertices u, v have λ color-disjoint paths in G iff they have λ such paths in H .

References

- 1 Ingo Althöfer, Gautam Das, David P. Dobkin, Deborah Joseph, and José Soares. On sparse spanners of weighted graphs. *Discret. Comput. Geom.*, 9:81–100, 1993. doi:10.1007/BF02189308.
- 2 Surender Baswana and Sandeep Sen. Approximate distance oracles for unweighted graphs in $\tilde{O}(n^2)$ time. In *Proceedings of the Fifteenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA*, pages 271–280, 2004. URL: <http://dl.acm.org/citation.cfm?id=982792.982830>.
- 3 Surender Baswana and Sandeep Sen. A simple and linear time randomized algorithm for computing sparse spanners in weighted graphs. *Random Struct. Algorithms*, 30(4):532–563, 2007. doi:10.1002/rsa.20130.
- 4 Greg Bodwin, Michael Dinitz, Merav Parter, and Virginia Vassilevska Williams. Optimal vertex fault tolerant spanners (for fixed stretch). In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA*, pages 1884–1900, 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, 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 ACM-SIAM Symposium on Discrete Algorithms, SODA*, pages 3272–3286, 2022. doi:10.1137/1.9781611977073.129.
- 7 Greg Bodwin and Shyamal Patel. A trivial yet optimal solution to vertex fault tolerant spanners. In *Proceedings of the 2019 ACM Symposium on Principles of Distributed Computing, PODC*, pages 541–543, 2019. doi:10.1145/3293611.3331588.
- 8 Shiri Chechik, Michael Langberg, David Peleg, and Liam Roditty. Fault tolerant spanners for general graphs. *SIAM J. Comput.*, 39(7):3403–3423, 2010. doi:10.1137/090758039.
- 9 Joseph Cheriyan, Ming-Yang Kao, and Ramakrishna Thurimella. Scan-first search and sparse certificates: An improved parallel algorithms for k-vertex connectivity. *SIAM J. Comput.*, 22(1):157–174, 1993.
- 10 David Coudert, Pallab Datta, Stephane Perennes, Hervé Rivano, and Marie-Emilie Voge. Shared risk resource group complexity and approximability issues. *Parallel Process. Lett.*, 17(2):169–184, 2007. doi:10.1142/S0129626407002958.
- 11 Lenore Cowen. Compact routing with minimum stretch. *J. Algorithms*, 38(1):170–183, 2001. doi:10.1006/jagm.2000.1134.
- 12 Lenore Cowen and Christopher G. Wagner. Compact roundtrip routing in directed networks. *J. Algorithms*, 50(1):79–95, 2004. doi:10.1016/j.jalgor.2003.08.001.
- 13 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*, pages 169–178, 2011. doi:10.1145/1993806.1993830.
- 14 Michael Dinitz and Caleb Robelle. Efficient and simple algorithms for fault-tolerant spanners. In *PODC '20: ACM Symposium on Principles of Distributed Computing*, pages 493–500, 2020. doi:10.1145/3382734.3405735.
- 15 Michael Elkin, Yuval Emek, Daniel A. Spielman, and Shang-Hua Teng. Lower-stretch spanning trees. *SIAM J. Comput.*, 38(2):608–628, 2008. doi:10.1137/050641661.

- 16 Georgios Ellinas, Eric Bouillet, Ramu Ramamurthy, J.-F. Labourdette, Sid Chaudhuri, and Krishna Bala. Routing and restoration architectures in mesh optical networks. *Opt Networks Mag*, 4, January 2003.
- 17 Paul Erdős. Extremal problems in graph theory. *Theory of Graphs and its Applications*, 1963.
- 18 Sebastian Forster, Danupon Nanongkai, Liu Yang, Thatchaphol Saranurak, and Sorrachai Yingchareonthawornchai. Computing and testing small connectivity in near-linear time and queries via fast local cut algorithms. In *Proceedings of the 2020 ACM-SIAM Symposium on Discrete Algorithms, SODA*, pages 2046–2065, 2020.
- 19 Mohsen Ghaffari and Fabian Kuhn. Distributed minimum cut approximation. In *Distributed Computing – 27th International Symposium, DISC*, volume 8205 of *Lecture Notes in Computer Science*, pages 1–15. Springer, 2013.
- 20 Shlomo Hoory. The size of bipartite graphs with a given girth. *J. Comb. Theory, Ser. B*, 86(2):215–220, 2002. doi:10.1006/JCTB.2002.2123.
- 21 Jian Qiang Hu. Diverse routing in optical mesh networks. *IEEE Transactions on Communications*, 51(3):489–494, 2003. doi:10.1109/TCOMM.2003.809779.
- 22 David R. Karger and Rajeev Motwani. An NC algorithm for minimum cuts. *SIAM J. Comput.*, 26(1):255–272, 1997.
- 23 F. Kuipers. An overview of algorithms for network survivability. *ISRN Communications and Networking*, 2012, December 2012. doi:10.5402/2012/932456.
- 24 Jason Li, Danupon Nanongkai, Debmalya Panigrahi, Thatchaphol Saranurak, and Sorrachai Yingchareonthawornchai. Vertex connectivity in poly-logarithmic max-flows. In *STOC '21: 53rd Annual ACM SIGACT Symposium on Theory of Computing*, pages 317–329, 2021.
- 25 David W. Matula. A linear time $2 + \epsilon$ approximation algorithm for edge connectivity. In *Proceedings of the Fourth Annual ACM/SIGACT-SIAM Symposium on Discrete Algorithms, SODA*, pages 500–504, 1993.
- 26 Eytan H. Modiano and Aradhana Narula-Tam. Survivable lightpath routing: a new approach to the design of wdm-based networks. *IEEE J. Sel. Areas Commun.*, 20(4):800–809, 2002. doi:10.1109/JSAC.2002.1003045.
- 27 Hiroshi Nagamochi and Toshihide Ibaraki. A linear-time algorithm for finding a sparse k-connected spanning subgraph of a k-connected graph. *Algorithmica*, 7(5&6):583–596, 1992. doi:10.1007/BF01758778.
- 28 Merav Parter. Small cuts and connectivity certificates: A fault tolerant approach. In *33rd International Symposium on Distributed Computing, DISC*, volume 146 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 30:1–30:16. Schloss Dagstuhl – Leibniz-Zentrum fuer Informatik, 2019. doi:10.4230/LIPIcs.DISC.2019.30.
- 29 Merav Parter. Nearly optimal vertex fault-tolerant spanners in optimal time: sequential, distributed, and parallel. In *STOC '22: 54th Annual ACM SIGACT Symposium on Theory of Computing*, pages 1080–1092, 2022. doi:10.1145/3519935.3520047.
- 30 David Peleg and Alejandro A. Schäffer. Graph spanners. *J. Graph Theory*, 13(1):99–116, 1989. doi:10.1002/jgt.3190130114.
- 31 David Peleg and Jeffrey D. Ullman. An optimal synchronizer for the hypercube. *SIAM J. Comput.*, 18(4):740–747, 1989. doi:10.1137/0218050.
- 32 David Peleg and Eli Upfal. A tradeoff between space and efficiency for routing tables. In *Proceedings of the 20th Annual ACM Symposium on Theory of Computing, STOC*, pages 43–52, 1988. doi:10.1145/62212.62217.
- 33 Mikkel Thorup and Uri Zwick. Compact routing schemes. In *Proceedings of the Thirteenth Annual ACM Symposium on Parallel Algorithms and Architectures, SPAA*, pages 1–10, 2001. doi:10.1145/378580.378581.
- 34 Mikkel Thorup and Uri Zwick. Approximate distance oracles. *J. ACM*, 52(1):1–24, 2005. doi:10.1145/1044731.1044732.
- 35 Peng Zhang, Jin-Yi Cai, Lin-Qing Tang, and Wen-Bo Zhao. Approximation and hardness results for label cut and related problems. *Journal of Combinatorial Optimization*, 21:192–208, February 2011. doi:10.1007/s10878-009-9222-0.

A FT Spanners for Multi-Graphs

All of the results in this paper apply also when G is a colored *multi-graph*, i.e., when many parallel edges between two vertices are allowed. Especially in edge-colored settings, this seems like a rather natural assumption in light of the real-world motivation of communication networks. However, it is presumably less natural in non-colored settings, so explicit discussions on multi-graphs have mostly been neglected in previous work on FT spanners. We next provide such a short discussion, which shows an interesting phenomenon: the bounds for uncolored multi-graphs are the same as for colored simple graphs.

EFT: $\Theta(fn^{1+1/k})$ size bounds. Our upper bound for ECFT-spanners in Theorem 5(1) applies for multi-graphs, and EFT is just a special case. Additionally, we note that the EFT-spanner algorithm of Chechik et al. [8] also works on multi-graphs, and achieves this size bound. For the lower bound, one may take the Girth Conjecture graph with $\Omega(n^{1+1/k})$ edges, and give each of its edges multiplicity f . Clearly, such a graph cannot have a proper f -EFT $(2k - 1)$ -spanner.

VFT: $\Theta(f^{1-1/k}n^{1+1/k})$ size bounds. In the VFT setting, parallel edges can be discarded by keeping the edge of lowest weight among them. This does not change distances (regardless of which vertex faults occur). Hence, the bounds are the same as for simple graphs.

MFT: $\Theta(f^{2-1/k}n^{1+1/k})$ size bounds. Constructing a lower bound instance is easy: Take the VFT lower bound instance which is a simple graph with $\Omega(f^{1-1/k}n^{1+1/k})$ edges and replace each edge by f parallel edges. Our upper bound for MCFT spanners applies also here.