

Fault Tolerant Approximate BFS Structures

Merav Parter, The Weizmann Institute of Science, Rehovot, Israel.

David Peleg, The Weizmann Institute of Science, Rehovot, Israel.

A *fault-tolerant* structure for a network is required to continue functioning following the failure of some of the network's edges or vertices. This paper addresses the problem of designing a *fault-tolerant* (α, β) approximate BFS structure (or FT-ABFS structure for short), namely, a subgraph H of the network G such that subsequent to the failure of some subset F of edges or vertices, the surviving part of H (namely, $H \setminus F$) still contains an *approximate* BFS spanning tree for (the surviving part of) G , satisfying $\text{dist}(s, v, H \setminus F) \leq \alpha \cdot \text{dist}(s, v, G \setminus F) + \beta$ for every $v \in V$.

Our first result is an algorithm that given an n -vertex unweighted undirected graph G and a source s constructs a *multiplicative* $(3, 0)$ FT-ABFS structure rooted at s resilient to a failure of a *single* edge with at most $4n$ edges (improving by an $O(\log n)$ factor on the near-tight result of [Baswana and Khanna 2010] for the special case of edge failures). This was recently improved to $2n$ edges by [Bilò et al. 2014]. Next, we consider the multiple edge faults case, for a constant integer $f > 1$, we prove that there exists a (polynomial-time constructible) $(3f, f \log n)$ FT-ABFS structure with $O(fn)$ edges that is resilient against f faults. We also show the existence of a $(3f + 1, 0)$ FT-ABFS structure with $O(f \log^f n \cdot n)$ edges.

We then consider *additive* $(1, \beta)$ FT-ABFS structures and demonstrate an interesting dichotomy between multiplicative and additive spanners. In contrast to the linear size of $(\alpha, 0)$ FT-ABFS structures, we show that for every n , there exist $\delta, \epsilon > 0$, and n -vertex graphs G with a source s for which any $(1, n^\delta)$ FT-ABFS structure rooted at s has $\Omega(n^{7/6-\epsilon})$ edges. For the case of additive stretch 3, we show that $(1, 3)$ FT-ABFS structures admit a lower bound of $\Omega(n^{5/4})$ edges.

Categories and Subject Descriptors: C.2.1 [Computer-Communication Networks]: Network Architecture and Design

General Terms: Design, Algorithms, Reliability

Additional Key Words and Phrases: Replacement Paths, Fault Tolerance, Additive Spanners

1. INTRODUCTION

Background and Motivation. Fault-tolerant subgraphs are subgraphs designed to maintain a certain desirable property in the presence of edge or vertex failures. This paper focuses on the property of containing a BFS tree with respect to some source s . A *fault tolerant BFS structure* (or FT-BFS structure) resistant to a single edge failure is a subgraph $H \subseteq G$ satisfying that $\text{dist}(s, v, H \setminus \{e\}) = \text{dist}(s, v, G \setminus \{e\})$ for every vertex $v \in V$ and edge $e \in E$.

To motivate our interest in such structures, consider a situation where it is required to lease a subnetwork of a given network, which will provide short routes from a source s to all other vertices. In a failure-free environment one can simply lease a BFS tree rooted at s . However, if links might disconnect, then one must prepare by leasing a larger set of links, and specifically an

Author's addresses: Department of Computer Science and Applied Mathematics. The Weizmann Institute of Science, Rehovot, Israel. E-mail: {merav.parter,david.peleg}@weizmann.ac.il. This work is supported in part by the Israel Science Foundation (grant 894/09), the United States-Israel Binational Science Foundation (grant 2008348), the I-CORE program of the Israel PBC and ISF (grant 4/11), the Israel Ministry of Science and Technology (infrastructures grant), and the Citi Foundation. M. Parter is recipient of the Google European Fellowship in distributed computing; research is supported in part by this Fellowship. An extended abstract of this paper has appeared in the proceedings of the 2014 ACM-SIAM Symposium on Discrete Algorithms.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org.

© YYYY ACM 1539-9087/YYYY/01-ARTA \$15.00

DOI: http://dx.doi.org/10.1145/0000000.0000000

FT-BFS structure. Moreover, taking costs into account, this example also motivates our interest in constructing *sparse* FT-BFS structure.

This question has recently been studied by us in [Parter and Peleg 2013]. Formally, a spanning graph $H \subseteq G$ is an f edge (resp., vertex) *fault-tolerant BFS* (FT-BFS) *structure* for G with respect to the source $s \in V$ iff for every $v \in V$ and every set $F \subseteq E(G)$ (resp., $F \subseteq V$), $|F| \leq f$, it holds that $\text{dist}(s, v, H \setminus F) = \text{dist}(s, v, G \setminus F)$. It is shown in [Parter and Peleg 2013] that for every graph G and source s there exists a (polynomial-time constructible) 1-edge FT-BFS structure H with $O(n^{3/2})$ edges. This result is complemented by a matching lower bound showing that for every sufficiently large integer n , there exist an n -vertex graph G and a source $s \in V$, for which every 1-edge FT-BFS structure is of size $\Omega(n^{3/2})$. Hence *exact* FT-BFS structures may be rather expensive.

This last observation motivates the approach of resorting to *approximate* distances, in order to allow the design of a sparse subgraph with properties resembling those of an FT-BFS structure. The current paper aims at exploring this approach, focusing mainly on subgraphs that contain *approximate* BFS structures and are resistant to a single edge failure. Formally, given an unweighted undirected n -vertex graph $G = (V, E)$ and a source $s \in V$, the subgraph $H \subseteq G$ is an f -edge (resp., vertex) (α, β) FT-ABFS structure with respect to s if for every vertex $v \in V$ and every set $F \subseteq E(G)$ (resp., $F \subseteq V$), $|F| \leq f$,

$$\text{dist}(s, v, H \setminus F) \leq \alpha \cdot \text{dist}(s, v, G \setminus F) + \beta.$$

(An (α, β) FT-ABFS structure is a fault-tolerant BFS (FT-BFS) structure if $\alpha = 1$ and $\beta = 0$.) We show that this relaxed requirement allows structures that are sparser than their exact counterparts.

Approximate BFS tree structures can also be compared against a different type of structures, namely, fault-tolerant *spanners*. Given an n -vertex graph $G = (V, E)$, the subgraph $H \subseteq G$ is an f -edge *fault-tolerant* (α, β) -*spanner* of G if for every two vertices $v, w \in V$ and every set $F \subseteq E(G)$, $|F| \leq f$, we have $\text{dist}(v, w, H \setminus F) \leq \alpha \cdot \text{dist}(v, w, G \setminus F) + \beta$. Observe that the union of (α, β) FT-ABFS structures with respect to every source $s \in V$ forms an (all-pairs) fault tolerant (α, β) -spanner for G . In fact, FT-ABFS structures can be viewed as *single source spanners*. Algorithms for constructing an f -vertex fault tolerant $(2k - 1)$ -spanner of size $O(f^{2k+1} \cdot n^{1+1/k} \log^{1-1/k} n)$ and an f -edge fault tolerant $(2k - 1)$ -spanner of size $O(f \cdot n^{1+1/k})$ for a given n -vertex graph G were presented in [Chechik et al. 2010]. A randomized construction attaining an improved tradeoff for vertex fault-tolerant spanners was then presented in [Dinitz and Krauthgamer 2011], yielding (with high probability) for every graph $G = (V, E)$ and integers k, f , an f -vertex fault-tolerant $(2k - 1)$ -spanner with $O(f^{2-1/k} n^{1+1/k} \log n)$ edges.

For the case of f edge failures for a constant $f \geq 1$, we show (in Sec. 2) that there exists a polynomial-time algorithm that for every n -vertex graph constructs a $(3f, f \log n)$ FT-ABFS structure H with $O(fn)$ edges overcoming up to f edge faults. We also considered the special case of a single edge failure ($f = 1$), and have shown a somewhat stronger result, namely, that for every n -vertex graph G and source s , there is a (polynomial-time constructible) $(3, 0)$ FT-ABFS structure with at most $4n$ edges, thus improving on the near-tight construction of [Baswana and Khanna 2010] by a $O(\log n)$ factor for the special case of $\alpha = 3$ and edge failures. This result was recently improved by [Bilò et al. 2014], obtaining a 3 approximation with at most $2n$ edges. Note that the algorithm of [Bilò et al. 2014] is limited to a single edge failure and cannot be used to overcome multiple edge faults. Towards the end of Sec. 2, we show the construction of a $(3f + 1, 0)$ FT-ABFS structure with $O(f(\log n)^f \cdot n)$ edges.

Our algorithm for approximately handling f edge failures is to be contrasted with two different structures: the (single-source) fault tolerant exact FT-BFS structure of [Parter and Peleg 2013], and the (all-pairs) fault tolerant $(3, 0)$ spanner of [Chechik et al. 2010], which both contain $\Theta(n^{3/2})$ edges. This implies that using FT-ABFS structures is more efficient than using fault-tolerant spanners even if it is necessary to handle not a single source s but a set $S \subseteq V$ of sources where $|S| = \Omega(n^\epsilon)$ for $\epsilon < 1/2$; a collection of approximate (α, β) FT-ABFS structures rooted at each of the sources $s \in S$ will still be cheaper than a fault-tolerant spanner.

Additive fault tolerant $(1, \beta)$ spanners were recently defined and studied by [Braunschvig et al. 2015], establishing the following general result. For a given n -vertex graph G , let H_1 be an ordinary additive $(1, \beta)$ spanner for G and H_2 be a fault tolerant $(\alpha, 0)$ spanner for G resilient against up to f edge faults. Then $H = H_1 \cup H_2$ is a $(1, \beta(f))$ additive fault tolerant spanner for G (for up to f edge faults) for $\beta(f) = O(f(\alpha + \beta))$. In particular, fixing the number of H edges to be $O(n^{4/3})$ and the number of faults to $f = 1$ yields an additive stretch of 38 (See [Braunschvig et al. 2015]; Cor. 1).

When considering FT-BFS structures with an *additive* stretch, namely, $(1, \beta)$ FT-ABFS structures, the improvement is less dramatic compared to the size of the single-source exact or the all-pairs approximate variants. In Sec. 3, we use the recent lower bound construction of [Abboud and Bodwin 2016] to show that there exists a superlinear lower bound on the size of the FT-ABFS structure with additive stretch n^δ and $\Omega(n^{7/6-\epsilon})$ edges. For additive stretch 3, we show a different construction with $\Omega(n^{5/4})$ edges. These new lower bound constructions are independent of the correctness of Erdős conjecture. Importantly, our results reveal an interesting dichotomy between *multiplicative* $(\alpha, 0)$ FT-ABFS and *additive* $(1, \beta)$ FT-ABFS structures: whereas every graph G contains a (polynomial-time constructible) $(3, 0)$ FT-ABFS structure rooted at $s \in V$ of size $\Theta(n)$, there exist an n -vertex graph G and a source $s \in V$ for which every $(1, \beta)$ FT-ABFS structure contains a *super-linear* number of edges.

Related work. FT-BFS structures are closely related to the notion of *replacement paths*. For a source s , a target vertex v and an edge $e \in G$, a *replacement path* is the shortest $s - v$ path that does not go through e . (Intuitively, we think of this replacement path as “protecting” the structure from the failure of the edge e , as it provides an alternative route from s to v that does not use e , and hence will continue to function even if e fails.) An FT-BFS structure is composed of a collection consisting of a replacement path for every target $v \in V$ and edge $e \in E$. Analogously, the notion of FT-ABFS structures is closely related to the problem of constructing *approximate replacement paths* [Baswana and Kavitha 2006; Chechik et al. 2012; Bernstein 2010], and in particular to its *single source* variant studied in [Baswana and Khanna 2010]. That problem requires to compute a collection \mathcal{P}_s consisting of an approximate $s - v$ replacement path $rp(s, v, G \setminus \{e\})$ for every $v \in V$ and every failed edge e that appears on the $s - v$ shortest-path in G , such that $|rp(s, v, G \setminus \{e\})| \leq \alpha \cdot \text{dist}(s, v, G \setminus \{e\})$. In the resulting *fault tolerant distance oracle*, in response to a query (s, v, F) consisting of an $s - v$ pair and a set F of failed edges or vertices (or both), the oracle S must return the distance between s and v in $G' = G \setminus F$. Such a structure is sometimes called an *F-sensitivity distance oracle*. The focus is on both fast preprocessing time, fast query time and low space. An approximate *single source fault tolerant distance oracle* has been first studied at [Baswana and Khanna 2010], which proposed an $O(n \log n / \epsilon^3)$ space data structure that can report a $(1 + \epsilon)$ approximate shortest path for any $\epsilon > 0$ for the case of unweighted graphs. An additional by-product of the data structure of [Baswana and Khanna 2010] is the construction of an $(1 + \epsilon, 0)$ FT-ABFS structure with $O(n/\epsilon^3 + n \log n)$ edges. Recently, this was extended for the case of weighted graphs [Bilò et al. 2014]. Setting $\epsilon = 2$, this yields a $(3, 0)$ FT-ABFS structure with $O(n \log n)$ edges. Hence our $(3, 0)$ FT-ABFS structure construction with at most $4n$ edges and the recent improvement of [Bilò et al. 2014] improves the construction of [Baswana and Khanna 2010] by a factor of $O(\log n)$ for the case of single edge failure (the construction of [Baswana and Khanna 2010] supports the case of vertex failures as well). For a summary table of the bounds see Table I.

It is important to note that the literature on approximate replacement paths (cf. [Baswana and Kavitha 2006; Bernstein 2010]) mainly focuses on *time-efficient* computation of these paths, as well as their efficient maintenance within distance oracles. In contrast, the main concern in the current paper is with optimizing the *size* of the resulting fault tolerant structure that contains the collection of approximate replacement paths.

Moreover, this paper considers both multiplicative and additive stretch, whereas the long line of existing approximate distance oracles concerned mostly multiplicative (and not additive) stretch, with the exception of [Pătraşcu and Roditty 2010; Abraham and Gavoille 2011]. To illustrate the

dichotomy between the additive and multiplicative setting, consider the issue of lower bounds for additive FT-ABFS structures. In the all-pairs fault-free setting, the best known lower bound for multiplicative spanners is based on the *girth conjecture* of Erdős [Erdős 1964], stating that there exist n -vertex graphs with $\Omega(n^{1+1/k})$ edges and girth (minimum cycle length) $2k + 2$ for any integer k . Removing any edge in such a graph increases the distance between its endpoints from 1 to $2k + 1$, hence any $(1, \beta)$ spanner with $\beta \leq 2k - 1$ must have $\Omega(n^{1+1/k})$ edges. This conjecture is settled only for $k = 1, 2, 3, 5$ (see [Wenger 1991]). Turning to additive spanners, in [Woodruff 2006], Woodruff presented a lower bound for additive spanners matching the girth conjecture bounds but independent of the correctness of the conjecture. More precisely, he showed the existence of graphs for which any spanner of size $O(k^{-1}n^{1+1/k})$ has an additive stretch of at least $2k - 1$, hence establishing a lower bound of $\Omega(k^{-1}n^{1+1/k})$ on the size of additive spanners. Recently, Abboud and Bodwin [Abboud and Bodwin 2016] presented a surprising lower bound of $\Omega(n^{4/3-\epsilon})$ for any additive stretch at most $n^{o(1)}$.

Upper bounds for constant stretch (non-fault-tolerant) additive spanners are currently known for but a few stretch values. A $(1, 2)$ spanner with $\tilde{O}(n^{3/2})$ edges is presented in [Aingworth et al. 1999], an improved construction with $O(n^{3/2})$ edges was first provided by [Elkin and Peleg 2004], a $(1, 6)$ spanner with $O(n^{4/3})$ edges is presented in [Baswana et al. 2005], and a $(1, 4)$ spanner with $\tilde{O}(n^{7/5})$ edges is presented in [Chechik 2013]. Constructions of fault tolerant additive spanners resilient to *edge* failures are presented in [Braunschvig et al. 2015]. Recently, [Parter 2014] considers the single *vertex* failure case, providing construction of FT spanners with additive stretch 2 and 6, this was later improved by [Bilò et al. 2015]. Currently, no lower bounds are known for this setting.

Table I: Fault Tolerant Distance Preservers

Approximation	Single Source	All Pairs
Exact, $f = 1$	$O(n^{3/2})$ [Parter and Peleg 2013]	$O(n^2)$
$(1 + \epsilon, \beta), f = 1$	$\beta = 0, O(n/\epsilon^3 + n \log n)$ [Baswana and Khanna 2010]	$\tilde{O}(f\beta(\beta\epsilon^{-1})^{2f} \cdot n^{1+1/k})$ [Braunschvig et al. 2015]
$(3, 0), f = 1,$	$2n$ [Bilò et al. 2014]	$O(n^{3/2})$ [Chechik et al. 2010]
$(3f, f \log n), f \geq 1,$	$O(f \cdot n)$ New	$O(f \cdot n^{1+2/(3f+1)})$ [Chechik et al. 2010]
$(1, \beta)$	$\beta = 4, f = 1, O(n^{4/3})$ [Parter and Peleg 2014a]	$\beta = 6, f = 1, \tilde{O}(n^{3/2})$ [Parter 2014] $\beta = 18f + 6, O(fn^{4/3})$ [Braunschvig et al. 2015; Bilò et al. 2015]

Preliminaries. Throughout, we fix an unweighted graph $G = (V, E)$ and a source vertex s in it. For a subgraph $G' = (V', E') \subseteq G$ (where $V' \subseteq V$ and $E' \subseteq E$) and a pair of vertices $u, v \in V$, let $\text{dist}(u, v, G')$ denote the distance in edges (a.k.a. *hop-distance*) between u and v in G' . Let $T_0 \subseteq G$ be a shortest paths (or BFS) tree rooted at s . Let $E(v, G) = \{(u, v) \in E(G)\}$ be the set of edges incident to v in the graph G and let $\deg(v, G) = |E(v, G)|$ denote the degree of vertex v in G . When the graph G is clear from the context, we may omit it and simply write $\deg(v)$. Let $\text{depth}(s, v) = \text{dist}(s, v, G)$ denote the *depth* of v in the BFS tree T_0 . When the source s is clear from the context, we may omit it and simply write $\text{depth}(v)$. Let $\text{Depth}(s) = \max_{u \in V} \{\text{depth}(s, u)\}$ be the *depth* of T_0 . For a path $P = [u_1, \dots, u_k]$, let $\text{LastE}(P)$ denote the last edge of P , let $|P|$ denote the length of P and let

$P[u_i, u_j]$ be the subpath of P from u_i to u_j . For paths P_1 and P_2 where the last vertex of P_1 equals the first vertex of P_2 , let $P_1 \circ P_2$ denote the path obtained by concatenating P_2 to P_1 .

Following other papers in this field [Hershberger and Suri 2001; Bernstein and Karger 2009], to avoid complications due to shortest-paths of the same length, we assume all path lengths are computed using some standard weight assignment $\omega : E(G) \rightarrow \mathbb{R}^+$ that guarantees the uniqueness of the shortest-paths. This weight assignment can be constructed either randomly by adding random perturbation to the edge weight or deterministically using the method of lexicographic perturbations, see Sec. 7 of [Borradaile et al. 2015]. Hereafter, for any subgraph G' and vertices x, y , let $sp(x, y, G')$ denote the unique $x - y$ shortest-path in G' .

Throughout, the edges of these paths are considered to be directed away from the source s . Given an $s - t$ path P and an edge $e = (u, v) \in P$, let $\text{dist}(s, e, P)$ be the distance between s and e 's far endpoint v on P . In addition, for an edge $e = (u, v) \in T_0$, define $\text{dist}(s, e) = i$ if $\text{depth}(u) = i - 1$ and $\text{depth}(v) = i$. For a subset $V' \subseteq V$, let $G(V')$ be the induced subgraph on V' . Let $\text{LCA}(V')$ be the least common ancestor of all the vertices in V' with respect to the BFS tree T_0 .

For a vertex u and an edge set $F \subseteq E$, define the *replacement path* of u upon the failure of F in G as $sp(s, u, G \setminus F)$, the shortest $s - u$ path in the surviving graph $G \setminus F$. Note that if $F \cap sp(s, u, T_0) = \emptyset$, then the replacement path $sp(s, u, G \setminus F)$ is simply the shortest-path $sp(s, u, T_0)$. (We sometimes simplify notation by omitting some of the indices when clear by context.)

For an edge $e = (u, v) \in T_0$, denote the set of vertices in $T_0(v)$, the subtree of T_0 rooted at v , by $S(e) = V(T_0(v))$. Note that the vertices of $S(e)$ are precisely those sensitive to the failure of the edge e , i.e., the vertices w having e on $sp(s, w, T_0)$, their $s - w$ path in T_0 , hence also $S(e) = \{w \mid e \in sp(s, w, T_0)\}$.

2. MULTIPLICATIVE FT-ABFS STRUCTURES

In this section, we consider the case of f edge failures for a constant integer $f \geq 1$. Let us first state our result for $f = 1$.

THEOREM 2.1. *There exists a poly-time algorithm that for every n -vertex graph G and source s constructs a $(3, 0)$ FT-ABFS structure with at most $4n$ edges overcoming a single edge fault.*

This result was recently improved by [Bilò et al. 2014], obtaining a 3 approximation with at most $2n$ edges. We therefore omit the details of our algorithm and analysis. Readers interested in those details are referred to [Parter and Peleg 2014b; 2014a].

The main result of this section is the following theorem for unweighted undirected graphs.

THEOREM 2.2. *There exists an algorithm $\text{Cons-FT-Add-Span}(s, G, f)$ that for every n -vertex graph G , source s and integer $f \geq 1$ constructs, in time $O(n^{2f})$, a $(3f, f \log n)$ FT-ABFS structure with $O(n)$ edges, overcoming up to f edge faults.*

PROOF. We present an algorithm, $\text{Cons-FT-Add-Span}(s, G, f)$, for constructing the sought structure. Let us start with an overview of the algorithm, which consists of three phases. Phase (1) constructs a (possibly dense) f -edge FT-BFS structure H_1 with respect to s by using Alg. $\text{Cons-FT-BFS}(s, G, f)$ to be defined later. Phase (2) constructs an f -edge $(3f, 0)$ FT-ABFS structure $H_2 \subseteq H_1$ by carefully sparsifying the edges of H_1 . However, H_2 might still be dense. Finally, phase (3) invokes Algorithm $\text{Cons-Span}(H_2 \setminus T_0, \log n, f)$ of [Chechik et al. 2010] and obtains a *sparse* multiplicative $O(\log n)$ spanner \hat{H} of H_2 such that $H = T_0 \cup \hat{H}$ is a sparse $(3f, f \cdot \log n)$ FT-ABFS structure $H \subseteq H_2$, where $|H| = O(n)$. For this phase we rely on the following fact.

LEMMA 2.3. [Chechik et al. 2010] *There exists an algorithm $\text{Cons-Span}(G, \alpha, f)$ that given an n -vertex graph G constructs an f edge fault tolerant $(2\alpha - 1, 0)$ spanner $G' \subseteq G$ such that $|G'| \leq O(f \cdot n^{1+1/\alpha})$.*

We proceed with a more detailed description of the algorithm. Algorithm Cons-FT-BFS(s, G, f) of phase (1) operates in a “brute-force” manner. For every $u \in V$, and every subset $F \subseteq E$, $|F| \leq f$, it constructs an $s-u$ replacement path $sp(s, u, G \setminus F)$ of minimal length. The f -edge FT-BFS structure H_1 is then simply the union of all these replacement paths, $H_1 = \bigcup \{sp(s, u, G \setminus F) \mid F \subseteq E, |F| \leq f, u \in V\}$. For a replacement path $P = sp(s, u, G \setminus F)$, let $\text{New}(P) = E(P) \setminus E(T_0)$ be the set of P ’s edges that are *new*, namely, not in the BFS tree T_0 .

In phase (2), each of these replacement paths $P = sp(s, u, G \setminus F)$ is transformed into another path P' in which the number of new edges is bounded by f . This has the penalty of introducing a stretch, that is, P' is an $s-u$ path in $G \setminus F$ but is not the *shortest* such path. To compute the path P' from P , the algorithm first labels the vertices according to the connected component they belong to in the forest $T_0 \setminus F$. Two vertices are given the same label iff they appear in the same connected subtree in the forest $T_0 \setminus F$. Note that since two vertices $u', u'' \in P$ of the same label are connected in $T_0 \setminus F$, the subpath $P[u', u'']$ (that potentially has many new edges) can be replaced by a T_0 path $sp(u', u'', T_0)$ that survived the failing of F . This allows us to save the new edges that occur on $P[u', u'']$. The potential drawback is that these bypasses might be longer than their counterparts in P . In the analysis we argue that the incurred multiplicative stretch is bounded by $3f$. Hence phase (2) turns the *exact* FT-BFS structure H_1 to an *approximate* $(3f, 0)$ FT-BFS structure H_2 .

We now describe formally the construction of H_2 . We start with partitioning the graph vertices as follows. Let $\mathcal{C} = \{T_0^1, \dots, T_0^\ell\}$ be the set of connected components (subtrees) of the forest $T_0 \setminus F$. Then each vertex $u \in T_0^y$ is assigned the label $\text{cc}_F(u) = y$. For pictorial illustration, see Fig. 2.

The procedure for selecting at most f new edges of $P = [u_1, \dots, u_k]$ for P' is as follows. For every $1 \leq i \leq k-1$, let $\varphi(i) \in \{i, \dots, k\}$ be the index of the farthest node on P with the same label as i , namely, $\varphi(i) = \max\{j \mid \text{cc}_F(v_j) = \text{cc}_F(v_i)\}$. Set $j \leftarrow 1$ and $i_1 \leftarrow 1$. Repeat the following operations until $\varphi(i_j) = k$: Set $i_{j+1} = \varphi(i_j) + 1$. Set the new edge $e_j^{\text{new}} = (v_{\varphi(i_j)}, v_{\varphi(i_j)+1})$ (note that e_j^{new} connects two vertices that are not connected in $T_0 \setminus F$, hence it is new), and add it to H_2 . Finally set $j \leftarrow j + 1$. Note that this procedure shortcuts the path P so that in the resulting path P' , hereafter denoted $rp(u)$, each component is entered at most once and the new edges that are added to H_2 are those connecting vertices in different components of $T_0 \setminus F$. Overall, the new edges of P that are added to H_2 are $\text{New}(rp(u)) = \{e_y^{\text{new}} \mid y \in \{1, \dots, \ell-1\}\}$, and the resulting approximate path is

$$rp(u) = sp(v_{i_1}, v_{\varphi(i_1)}, T_0) \circ e_1^{\text{new}} \circ \dots \circ sp(v_{i_y}, v_{\varphi(i_y)}, T_0) \circ e_y^{\text{new}} \circ \dots \circ sp(v_{i_{\ell'}}, v_{\varphi(i_{\ell'})}, T_0)$$

for some $\ell' \leq \ell$. The path $rp(u)$ replaces the $P[v_{i_y}, v_{\varphi(i_y)}]$ subpath with the T_0 subpath $sp(v_{i_y}, v_{\varphi(i_y)}, T_0)$. For a pictorial illustration, see Fig. 1.

Finally, we describe phase (3) of the algorithm. Given the $(3f, 0)$ FT-ABFS structure H_2 , a subgraph $(3f, f \log n)$ FT-ABFS structure T_3 of $O(n)$ edges is constructed as follows. Let $G' = H_2 \setminus T_0$ be the subgraph obtained by removing all original BFS edges of T_0 from H_2 . Let $\hat{H} \leftarrow \text{Cons-Span}(G', \log n, f)$ be an f -edge fault tolerant $(\log n, 0)$ spanner for G' , constructed using Algorithm Cons-Span($H_2 \setminus T_0, \log n, f$) of [Chechik et al. 2010]. The resulting structure is $H = T_0 \cup \hat{H}$.

We now analyze the properties of the construction. We first provide a straightforward auxiliary claim regarding the labels $\text{cc}_F(u_i)$, for $F \subseteq E$. Recall that $sp(u, v, T_0)$ is *not necessarily* the shortest $u-v$ path in G .

OBSERVATION 2.4. (1) $sp(u_i, u_{i'}, T_0) \subseteq T_0 \setminus F$ for every $u_i, u_{i'}$ such that $\text{cc}_F(u_i) = \text{cc}_F(u_{i'})$.
(2) $|\bigcup \{\text{cc}_F(u') \mid u' \in V\}| = |\mathcal{C}| \leq |F| + 1$.

Fix a vertex $u \in V$, edge faults $F \subseteq E$, $|F| \leq f$, and a replacement path $P = sp(s, u, G \setminus F)$. For $y \in \{1, \dots, \ell\}$, let $(v_{i_y}, v_{\varphi(i_y)})$ be the pairs of vertices (connected in T_0) according to their appearance on P when traversed by the procedure above when constructing H_2 , hence $s = v_{i_1}$. The following observation is immediate by the structure of the algorithm.




Fig. 1: Illustration of the approximate replacement path. Dashed lines correspond to BFS edges. The exact replacement path $sp(s, u, G \setminus F)$ is shown in blue. The green BFS paths between v_{i_1} and $v_{\varphi(i_1)}$ and between v_{i_2} and $v_{\varphi(i_2)}$ are used in $rp(u)$ to bypass segments in $sp(s, u, G \setminus F)$. As a result, the approximated replacement path $rp(u)$ contains $f = 2$ new edges, the edge (p, v_{i_1}) and the edge $(v_{\varphi(i_1)}, v_{i_2})$.




Fig. 2: Upon the failure of the BFS edges F (dotted lines), the BFS tree T_0 is decomposed into $\ell = |F| + 1 \leq f + 1$ components, denoted T_0^1, \dots, T_0^ℓ .

OBSERVATION 2.5. (1) $rp(u) \subseteq H_2 \setminus F$. (2) $sp(s, u, G \setminus F)[s, v_{\varphi(i_1)}] = sp(s, v_{\varphi(i_1)}, T_0)$.

We proceed by bounding the length of $rp(u)$. Recall that by Obs. 2.5(2), $P[s, v_{\varphi(i_1)}] = rp(u)[s, v_{\varphi(i_1)}] = sp(s, v_{\varphi(i_1)}, T_0)$, hence the distance increase when replacing the subpath $P[v_{i_y}, v_{\varphi(i_y)}]$ with $sp(v_{i_y}, v_{\varphi(i_y)}, T_0)$ occurs only for $y \geq 2$. For every $y \in \{2, \dots, \ell'\}$, define $\chi_y = \text{LCA}(v_{i_y}, v_{\varphi(i_y)})$, $A_y = P[s, v_{i_y}]$, $B_y = sp(s, \chi_y, T_0)$, $C_y = sp(\chi_y, v_{i_y}, T_0)$, $D_y = sp(\chi_y, v_{\varphi(i_y)}, T_0)$ and $R_y = P[v_{i_y}, v_{\varphi(i_y)}]$. Let $\tilde{P}_y = A_y \circ C_y \circ D_y$ and $P_y = P[s, v_{\varphi(i_y)}]$. See Fig. 3.

CLAIM 2.6. $|\tilde{P}_y| \leq 3 \cdot |P_y|$ for every $y \in \{2, \dots, \ell\}$.

PROOF. Note that since $B_y \circ C_y = sp(s, v_{i_y}, T_0)$, the alternative $s - v_{i_y}$ path A_y satisfies $|A_y| \geq |B_y| + |C_y|$. In addition, since $C_y \circ R_y$ is an alternative $\chi_y - v_{\varphi(i_y)}$ path and $D_y = sp(\chi_y, v_{\varphi(i_y)}, T_0)$ is the $\chi_y - v_{\varphi(i_y)}$ shortest-path, it holds that $|D_y| \leq |C_y| + |R_y|$. Hence,

$$|\tilde{P}_y| \leq |A_y| + 2|C_y| + |R_y| \leq 3|A_y| + |R_y| \leq 3|P_y|. \quad \square$$

CLAIM 2.7. $|rp(u)[s, v_{i_y}]| \leq 3(y-1)|sp(s, u, G \setminus F)[s, v_{i_y}]|$ for every $y \in \{2, \dots, \ell'\}$.

PROOF. By induction on y . Let $P_y = sp(s, u, G \setminus F)[s, v_{\varphi(i_y)}]$ and $Q_y = rp(u)[s, v_{\varphi(i_y)}]$. Note that $Q_{y+1} = Q_y \circ e_y^{new}$.

For the induction base, $y = 2$, as $B_2 \circ C_2 = sp(v_{i_2}, v_{\varphi(i_2)}, T_0)$ and since $P_1 = Q_1$ by Obs. 2.5(2), we get that $Q_2 = sp(s, u, G \setminus F)[s, v_{i_2}] \circ sp(v_{i_2}, v_{\varphi(i_2)}, T_0)$. Hence, the claim follows by Cl. 2.6 and $|Q_2| \leq 3 \cdot |P_2|$. Now assume the claim holds up to $y-1$ and consider y .

$$\begin{aligned} |Q_y| &= |Q_{y-1}| + 1 + |sp(v_{i_y}, v_{\varphi(i_y)}, T_0)| \leq 3(y-2)(|P_{y-1}| + 1) + |C_y \circ D_y| \\ &\leq 3(y-2)|P[s, v_{i_y}]| + |C_y \circ D_y| = 3(y-2)|A_y| + |C_y \circ D_y| \leq 3(y-2)|A_y| + 3|A_y \circ R_y| \\ &\leq 3(y-1) \cdot |P_y|, \end{aligned}$$

where the first inequality follows by the induction assumption, as $|Q_{y-1}| \leq 3(y-2)|P_{y-1}|$, and the third inequality follows by Cl. 2.6. Since by Obs. 2.4(2), $\ell \leq f+1$, we get that $|Q_\ell| \leq 3f|P_\ell|$ and $|rp(u)| \leq 3f|sp(s, u, G \setminus F)|$. The lemma follows. \square




Fig. 3: Illustration of the approximate replacement paths. The dashed lines represent the optimal replacement path $P = sp(s, u, G \setminus F)$. The bold edges are the BFS edges. The approximate replacement path $rp(u)$ replaces the subpath $P[v_{i_y}, v_{\varphi(i_y)}]$ with the tree path $C_y \circ D_y = sp(v_{i_y}, \chi_y, T_0) \circ sp(\chi_y, v_{\varphi(i_y)}, T_0)$.

We therefore have the following.

COROLLARY 2.8. For every $u \in V$ and $F \subseteq E$, there exists a replacement path $rp(u) \in H_2 \setminus F$ such that (1) $|\text{New}(rp(u))| \leq f$, and (2) $|rp(u)| \leq 3f \cdot |sp(s, u, G \setminus F)|$. Hence H_2 is an f -edge $(3f, 0)$ FT-ABFS structure.

Finally, we prove the correctness of the last phase and show the following.

CLAIM 2.9. (1) H is an f -edge $(3f, f \log n)$ FT-ABFS structure for s . (2) $|E(H)| = O(fn)$.

PROOF. Let $P = sp(s, u, G \setminus F)$ be the optimal replacement path and let $Q = rp(u)$ be the corresponding replacement path in H_2 obtained by using at most f bypasses between vertices

of the same label on P . Then, by Cor. 2.8, $|\text{New}(Q)| \leq f$ and $|Q| \leq 3f \cdot |P|$. Let $G' = H_2 \setminus T_0$ and $G'' = \text{Cons-Span}(G', \log n, f)$ be the f -edge FT $(\log n, 0)$ spanner for G' (see Fact 2.3). Let $E' = \text{New}(Q) \setminus H$ be the set of new edges in Q that are missing in the final H . Hence, $E' \subseteq G' \setminus G''$. Since $E' \subseteq \text{New}(Q)$, it holds that $|E'| \leq |\text{New}(Q)| \leq f$. We now claim that for every missing edge $e = (x, y) \in E'$ there exists an $x - y$ path in the structure $G'' \setminus F$ of length at most $\log n$. By the fact that G'' is an f -edge $(\log n, 0)$ spanner for G' , it holds that $\text{dist}(x, y, G'' \setminus F) \leq \log n \cdot \text{dist}(x, y, G' \setminus F) = \log n$, where the last equality follows by the fact that $e = (x, y)$ appears on the replacement path Q , hence $e \notin F$, so $\text{dist}(x, y, G' \setminus F) = 1$. We therefore have that Q contains at most f missing edges, and for each there exists a path in G'' of length at most $\log n$. Overall, $\text{dist}(s, u, G \setminus F) \leq |E(Q) \setminus E'| + (\log n) \cdot |E'| \leq 3f|P| + f \log n$. Part (1) is established. As for part (2), by Fact 2.3, $|G''| = O(n)$, hence $H = T_0 \cup G''$ has $O(fn)$ edges as well. \square

This completes the proof of Thm. 2.2. \square

Note that one could use any $(2k - 1)$ -spanner instead of using $k = O(\log n)$, but that would result in a denser subgraph. Fig. 4 illustrates the near-tightness of the analysis by showing a graph example in which the constructed replacement path is longer by a factor of $2f$ than the optimal path. Thus, the dependency of the stretch on f is unavoidable in our algorithm.




Fig. 4: Illustrating the tightness of the analysis. Shown in blue is a BFS tree T_0 rooted at s . Dashed lines represent paths and non-dashed lines are edges. The length of the the path from s to the leaf nodes is $L + 1$. The failure set is the set of f edges incident to s . The replacement path $sp(s, u, G \setminus F)$ goes from s to v and then proceeds on f new edges (shown in red). Hence, $|sp(s, u, G \setminus F)| = L + f + 1$. The path given by the algorithm replace a single red edge with a tree path of length $2L$. Hence $|rp(u)| \geq 2f \cdot L$.

We next assert that it is possible to get rid of the $f \log n$ additive factor in Thm. 2.2, albeit at the expense of increasing the size of the structure by polylog factor. Let $\ell \geq 1$ be a given threshold on the length of paths. In what follows we refer to a replacement path as *short* iff its length is at most ℓ . We first prove the following lemma, establishing the existence of a sparse structure (whose size depends on ℓ) that contains all short replacement paths.

LEMMA 2.10. *There exists an algorithm that constructs a subgraph $H(\ell)$ containing all f -failures short replacement paths from s with $O(n \cdot \ell^{f+1})$ edges.*

PROOF. We construct the collection of all replacement paths $P_F^*(u_i) = sp(s, u_i, G \setminus F)$ for every $F \subseteq E$, $|F| \leq f$ and every $u_i \in V$. The subgraph $H(\ell)$ contains the last edges of all short replacement paths, i.e., $H(\ell) = \{\text{LastE}(P_F^*(u_i)) \mid |P_F^*(u_i)| \leq \ell\}$. We first show that $H(\ell)$ preserve all the shortest replacement paths distances.

CLAIM 2.11. *For every u and $F \subseteq E$, $|F| \leq f$ such that $|P_F^*(u)| \leq \ell$, it holds that $\text{dist}(s, u, G \setminus F) = \text{dist}(s, u, H(\ell) \setminus F)$.*

PROOF. Towards contradiction, assume otherwise, and let

$$BP = \{(i, F) \mid |sp(s, u_i, G \setminus F)| \leq \ell \text{ and } \text{dist}(s, u_i, H(\ell) \setminus F) > |sp(s, u_i, G \setminus F)|\}$$

be the set of “bad pairs,” namely, pairs (i, F) for which the length of the replacement $s - u_i$ path in $H(\ell) \setminus F$ is greater than $|sp(s, u_i, G \setminus F)| = \text{dist}(s, u_i, G \setminus F)$. (By the contradictory assumption, $BP \neq \emptyset$.) For each bad pair $(i, F) \in BP$, define $BE(i, F) = sp(s, u_i, G \setminus F) \setminus E(H(\ell))$ to be the set of “bad edges,” namely, the set of $sp(s, u_i, G \setminus F)$ edges that are missing in $H(\ell)$. By definition, $BE(i, F) \neq \emptyset$ for every bad pair $(i, F) \in BP$. Let $d(i, F) = \max_{e \in BE(i, F)} \{\text{dist}(s, e, sp(s, u_i, G \setminus F))\}$ be the maximal depth of a missing edge in $BE(i, F)$, and let $DM(i, F)$ denote that “deepest missing edge”, i.e., the edge e on $P_F^*(u_i)$ satisfying $d(i, F) = \text{dist}(s, e, sp(s, u_i, G \setminus F))$. Finally, let $(i', F') \in BP$ be the pair that minimizes $d(i, F)$, and let $e_1 = (v_{\ell_1}, u_{i_1}) \in BE(i', F')$ be the deepest missing edge on $sp(s, u_{i'}, G \setminus F')$, namely, $e_1 = DM(i', F')$. Note that e_1 is the *shallowest* “deepest missing edge” over all bad pairs $(i, F) \in BP$.

We now claim that (i_1, F') is also a bad-pair. Towards contradiction, assume otherwise, and let $P' \in sp(s, u_{i_1}, H(\ell) \setminus F')$. Consider the $s - u_{i'}$ path $P = P' \circ sp(s, u_{i'}, G \setminus F')[u_{i_1}, u_{i'}]$. By definition, $P \subseteq H(\ell) \setminus F'$. We have that

$$\begin{aligned} \text{dist}(s, u_{i'}, H(\ell) \setminus F') &\leq |P| = |P'| + |sp(s, u_{i'}, G \setminus F')[u_{i_1}, u_{i'}]| \\ &= \text{dist}(s, u_{i_1}, G \setminus F') + |sp(s, u_{i'}, G \setminus F')[u_{i_1}, u_{i'}]| \\ &= |sp(s, u_{i'}, G \setminus F')| = \text{dist}(s, u_{i'}, G \setminus F'), \end{aligned}$$

in contradiction to the fact that $(i', F') \in BP$.

In addition, note that $|sp(s, u_{i_1}, G \setminus F')| \leq |sp(s, u_{i'}, G \setminus F')| \leq \ell$, hence the last edge of $sp(s, u_{i_1}, G \setminus F')$ was added to $H(\ell)$. Since $(i_1, F') \in BP$, and its last missing edge in $H(\ell)$ is not its last edge, we get that $d(i_1, F') < d(i', F')$, contradicting the definition of (i', F') . The claim follows. \square

We next bound the size of the subgraph $H(\ell)$.

CLAIM 2.12. $|E(H(\ell))| = O(n \cdot \ell^{f+1})$

PROOF. Let $\mathcal{P}(u) = \{sp(s, u, G \setminus F) \mid F \subseteq E, |F| \leq f\}$ be the collection of $s - u$ replacement paths. Define $\mathcal{P}_0(u) = \{sp(s, u, T_0)\}$ and $\mathcal{P}_{f'}(u) = \{sp(s, u, G \setminus F) \mid |F| = f'\}$ for $f' \geq 1$ as the collection of $s - u_i$ replacement paths supporting a sequence of f' edge faults. Hence, $\mathcal{P}_i = \bigcup_{f'=1}^f \mathcal{P}_{f'}(u)$.

We prove that for every $f' \in \{1, \dots, f\}$ it holds that $|\mathcal{P}_{f'}(u)| \leq \ell \cdot |\mathcal{P}_{f'-1}(u)|$. To see this, observe that every replacement path $sp(s, u, G \setminus F) \in \mathcal{P}_{f'}(u) \setminus \mathcal{P}_{f'-1}(u)$ was included in order to protect against the failure of some edge e in some $sp(s, u, G \setminus F') \in \mathcal{P}_{f'-1}(u)$. Hence, $F = F' \cup \{e\}$. Since $\mathcal{P}_{f'-1}(u)$ contains only short replacement paths, it holds that each short $s - u$ replacement path $sp(s, u, G \setminus F') \in \mathcal{P}_{f'-1}(u)$ has at most ℓ replacement paths of the form $sp(s, u, G \setminus (F' \cup \{e\}))$ in $\mathcal{P}_{f'}(u)$ that were included in order to protect against the failure of some edge $e \in sp(s, u, G \setminus F')$. This implies that $|\mathcal{P}_{f'}(u)| \leq \ell \cdot |\mathcal{P}_{f'-1}(u)|$. Overall, $|\bigcup_{f'=1}^f \mathcal{P}_{f'}(u)| \leq \ell^{f+1}$ and $|H(\ell)| \leq n \cdot \ell^{f+1}$. \square

Lemma 2.10 follows by Claims 2.11 and 2.12. \square

Letting $\ell = f \cdot (2 \log n - 1)$, we now show the following.

THEOREM 2.13. *There exists an algorithm $\text{Cons-FT-Add-Span}^+(s, G, f)$ that for every n -vertex graph G , source s and integer $f \geq 1$ constructs, in time $O(n^{2f})$, a $(3f + 1, 0)$ FT-ABFS structure with $O((f(2 \log n - 1))^{f+1} \cdot n)$ edges, overcoming up to f edge faults.*

PROOF. We first construct the collection of all replacement paths $P_F^*(u_i) = sp(s, u_i, G \setminus F)$ for every $F \subseteq E$, $|F| \leq f$ and every $u_i \in V$. The $(3f, 0)$ FT-ABFS structure H is constructed in two steps. In the first step, a subgraph $H_3 \subseteq G$ is constructed containing the set of last edges of all short replacement paths, i.e., $H_3 = \{\text{LastE}(P_F^*(u_i)) \mid |P_F^*(u_i)| \leq \ell\}$. In the second step, a $(3f, \ell)$ FT-ABFS structure H_4 is constructed by employing the algorithm of Thm. 2.2, where in step (3) of the algorithm we construct a $(\log n, 0)$ spanner.

We first prove the correctness of the final structure $H = H_3 \cup H_4$. Consider a pair (i, F) corresponding to a vertex $u_i \in V$ and edge set $F \subseteq E$. There are two cases to examine. Case (a) is where the replacement path is *long*, i.e., $|sp(s, u, G \setminus F)| > \ell$. By Thm. 2.2,

$$\text{dist}(s, u, H_4 \setminus F) \leq 3f \cdot \text{dist}(s, u, G \setminus F) + \ell \leq (3f + 1) \cdot \text{dist}(s, u, G \setminus F).$$

The complementary case (b) where the replacement path is *short*, i.e., $|sp(s, u, G \setminus F)| \leq \ell$ follows by Cl. 2.11.

It remains to bound the size of H . By Lemma 2.10, $|E(H_3)| = O(n \cdot (f \cdot (2 \log n - 1))^{f+1})$. By Fact 2.3 and the proof of Cl. 2.9, it holds that $|E(H_4)| = O(f \cdot n)$. The theorem follows. \square

3. LOWER BOUNDS FOR ADDITIVE FT-ABFS STRUCTURES

In this section we provide lower bound constructions for $(1, \beta)$ FT-ABFS structures for various values of the additive stretch β . The starting point for these constructions is the lower bound construction for the exact FT-BFS structures in [Parter and Peleg 2013]. In the latter construction, the bulk of the edges was due to a complete bipartite graph $B = (L, R)$ where $|L| = \Theta(n)$ and $|R| = \Theta(\sqrt{n})$. When relaxing the quality requirement to additive stretch 2, it is no longer necessary to include B entirely in the spanner; in fact, one can show that it suffices to include in the $(1, 2)$ FT-ABFS structure only a subgraph $B' \subseteq B$ of $O(n)$ edges. This subgraph B' is obtained by connecting an arbitrary vertex $p^* \in L$ to every vertex $r \in R$ and connecting an arbitrary vertex $r^* \in R$ to every vertex $p \in L$. Our lower bound strategy is based on carefully replacing B with different graphs (e.g., of large girth).

The structure of the section is as follows. We first observe that the exact lower bound of [Parter and Peleg 2013] extends for additive stretch 1. We then sketch a construction for additive stretch 3 and $\Omega(n^{5/4})$ edges. This is done by carefully plugging *multiple* copies of dense graphs with girth at least 6. Finally, we describe in more detail the general case of stretch $\beta \leq n^\delta$ for small constant $\delta > 0$ and $\Omega(n^{7/6-\epsilon})$ edges. Here we replace the complete bipartite graph B with a copy of Abboud and Bodwin's lower bound graph.

Additive Stretch $\beta = 1$. By a proof very similar to that of the exact case (see [Parter and Peleg 2013]) one can show the following (details are omitted).

THEOREM 3.1. *There exists an n -vertex graph $G(V, E)$ and a source $s \in V$ such that any $(1, 1)$ FT-ABFS structure rooted at s has $\Omega(n^{3/2})$ edges.*

Additive stretch $\beta = 3$. For additive stretch 3, we replace the complete bipartite graph B by using $\Theta(\sqrt{n})$ copies of graphs $G_6(n')$ for $n' = \sqrt{n}$ where $G_6(x)$ is an x -vertex graph with girth 6. Lemma 5.1.1 of [Godsil and Royle 2013], there are x -vertex graph $G_6(x)$ with girth at least 6 and $\Omega(x^{3/2})$ edges. Hence, overall the lower bound graph contains $\Theta(\sqrt{n} \times n^{3/4}) = \Theta(n^{5/4})$ edges. For pictorial illustration of the construction see Fig. 5. The analysis is very similar to that of [Parter and Peleg 2013] and its details can be found in [Parter and Peleg 2014a].




Fig. 5: Lower bound construction for (1, 3) FT-ABFS structures. Dashed red lines correspond to 4-length path. When an edge e in one of the $G_6(d)$ graphs is removed, there exists a pair $z_{j,i} \in Z_j$ and $x \in X_i$ whose distance in $G_6(d) \setminus \{e\}$ is increased by more than 3. When $e_i \in P_0$ fails, the path from s to x in $G \setminus \{e_i\}$ must go through $z_{j,i}$. We then have that $\text{dist}(s, x, G \setminus \{e_i, e\}) \geq \text{dist}(s, x, G \setminus \{e_i\}) + 4$.

Additive stretch $\beta = n^{o(1)}$. Here, we replace the complete bipartite graph B with a copy of the lower bound graph construction by Abboud and Bodwin [Abboud and Bodwin 2016]. The copy replaces the complete $n \times \sqrt{n}$ bipartite graph in the exact FT-BFS lower bound of [Parter and Peleg 2013]. We use the version of Lemma B.8 from [Bodwin et al. 2017].

LEMMA 3.2 ([ABBOUD AND BODWIN 2016],[BODWIN ET AL. 2017]). *There are absolute constants¹ $\delta, \epsilon > 0$, a family of n -node graphs $G^*(n) = (V, E)$, node subsets $X, Z \subseteq V$ of size $|X| = \Theta(n^{2/3-\delta})$, $|Z| = \Theta(n^{1/2-\delta})$, and a set $W \subseteq X \times Z$ of size $|W| = |X||Z|n^{-\epsilon}$ with the following property: for each pair $(x_i, z_j) \in W$, we may assign a set of edges in G such that (1) no edge is assigned to two or more pairs, and (2) if all edges assigned to a pair (x_i, z_j) are removed from G , resulting in G' , then $\text{dist}(x_i, z_j, G')$ increases by n^δ . Moreover, $\text{dist}(x_i, z_j) = L$ for all $(x_i, z_j) \in W$, and $\text{dist}(x_i, z_j) \geq L$ for all $(x_i, z_j) \in X \times Z$.*

We show the following.

THEOREM 3.3. *For every integer n , there exist $\delta, \epsilon > 0$, n -vertex graph $G(n) = (V, E)$ and a source $s \in V$ such that any $(1, n^\delta)$ FT-ABFS structure with respect to s has $\Omega(n^{7/6-\epsilon})$ edges.*

¹for some absolute constant ϵ , which can be made arbitrarily small by choosing a sufficiently small δ

PROOF. The construction is based on plugging the graph G^* in the exact lower bound construction of [Parter and Peleg 2013]. Set $\beta = n^\delta$. The resulting graph consists of the following components.

- (1) A path $P_0 = [s = v_1, \dots, v_{d+1} = v^*]$ of length $d = \Theta(n^{1/2-\delta})$. Our focus in the analysis is on what happens when some edge on this path fails.
- (2) A copy of the graph $G^*(c \cdot n)$ for some constant $c < 1$ such that $|Z| = d$ and $|X| = \Theta(n^{2/3-\delta})$. We think of the bottom layer as consisting of the vertices Z in the graph, and the top layer as consisting of the vertices X . This graph contributes most of the edges in G (with all the other components containing only $O(n)$ edges). See Figure 6.
- (3) A set of $|X| = \Theta(n^{2/3-\delta})$ vertex disjoint paths of length β each, $U_i = [v^* = u_1^i, \dots, u_\beta^i = x_i]$, for every $i \in \{1, \dots, |X|\}$ connecting the vertex $v^* \in P_0$ to the vertices of $x_i \in X$.
- (4) A collection of $d = |Z|$ vertex disjoint paths of decreasing length, P_1, \dots, P_d , where for $j \in \{1, \dots, d\}$, $P_j = [v_j = p_1^j, \dots, p_{\ell_j}^j = z_j]$ connects v_j with $z_j \in Z$ and its length is $\ell_j = |P_j| = d + 4 + \beta \cdot (d - j + 1)$.

The main bulk of vertices is in the paths P_1, \dots, P_d , which contain $\Theta(n^\delta \cdot d^2) = \Theta(n^\delta \cdot n^{1-2\delta})$ vertices, hence G contains $\Theta(n)$ vertices. We now turn to prove the correctness of the construction.

CLAIM 3.4. *For every $n = o(1)$, there exist $\delta, \epsilon > 0$ and an n -vertex graph G , such that every $(1, n^\delta)$ FT-ABFS structure H for G with respect to s must contain $O(n^{7/6-2\delta-\epsilon})$ edges.*

PROOF. Assume, towards contradiction, that there exists a $(1, \beta)$ FT-ABFS structure H for G with $o(n^{7/6-2\delta-\epsilon})$ edges. Let H' be the subgraph H restricted to the edges of $G^*(c \cdot n)$ (i.e., $H' \subseteq H \cap G^*(c \cdot n)$). Let $W \subseteq X \times Z$ be the subset of $|X| \cdot |Y| \cdot n^{-\epsilon}$ pairs asserted in Lemma 3.2. By Lemma 3.2, if every pair in W has at least one of its essential edges in H' , then $|E(H')| \geq |W| = n^{7/6-2\delta-\epsilon}$. Therefore, there must be a pair $(z_i, x_j) \in W$ for which $\text{dist}(z_i, x_j, H') \geq \text{dist}(z_i, x_j, G^*(c \cdot n)) + \Omega(\beta)$. Consider the distance from s to x_j when the edge $e_i = (v_i, v_{i+1})$ on the main path P_0 fails. Since the X vertices are connected to v^* via paths of length β , the only remaining option is to use the z_ℓ vertices for $\ell \leq i$. This is because the bypasses below the failing point e_i are no longer accessible from s . Finally, since the length difference between P_i and P_ℓ is at least $\beta + 1$ and by Lemma 3.2, $\text{dist}(z_\ell, x_j, G^*(c \cdot n)) \geq L$, using any other z_ℓ incurs an additive stretch that is larger by β . \square

We conclude that the graph $G(V, E)$ must contain $\Omega(n^{7/6-\epsilon})$ edges. \square

Acknowledgement: We are grateful to the anonymous referees for their useful comments and suggestions, and in particular for the suggestion to improve the lower bound of Sect. 3 by employing the construction of [Bodwin et al. 2017].

REFERENCES

- Amir Abboud and Greg Bodwin. 2016. The 4/3 additive spanner exponent is tight. In *Proc. ACM Symp. on Theory of Computing (STOC)*. 351–361.
- Ittai Abraham and Cyril Gavoille. 2011. On approximate distance labels and routing schemes with affine stretch. In *Distributed Computing*. SV, 404–415.
- Donald Aingworth, Chandra Chekuri, Piotr Indyk, and Rajeev Motwani. 1999. Fast estimation of diameter and shortest paths (without matrix multiplication). *SIAM J. Computing* 28, 4 (1999), 1167–1181.
- Surender Baswana and Telikepalli Kavitha. 2006. Faster algorithms for approximate distance oracles and all-pairs small stretch paths. In *Proc. 47th IEEE Symp. on Foundations of Computer Science (FOCS)*. 591–602.
- Surender Baswana, Telikepalli Kavitha, Kurt Mehlhorn, and Seth Pettie. 2005. New constructions of (α, β) -spanners and purely additive spanners. In *Proc. 16th ACM-SIAM Symp. on Discrete algorithms*. 672–681.




Fig. 6: Lower bound construction for $(1, n^\delta)$ FT-ABFS structures. If less than $|X||Z| \cdot n^{-\epsilon}$ edges are taken into the structure, there exists a pair (x_j, z_j) whose distance in $G^*(cn)$ is increased by at least n^δ . Consequently, when $e_i \in P_0$ fails, the path from s to x_j is increased by at least n^δ compared to the real distance $\text{dist}(s, x_j, G \setminus \{e_i\})$.

- Surender Baswana and Neelesh Khanna. 2010. Approximate Shortest Paths Avoiding a Failed Vertex: Optimal Size Data Structures for Unweighted Graph. In *Proc. 27th Symp. on Theoret. Aspects of Computer Sci. (STACS)*. 513–524.
- Aaron Bernstein. 2010. A nearly optimal algorithm for approximating replacement paths and k shortest simple paths in general graphs. In *Proc. 21st ACM-SIAM Symp. on Discrete Algorithms (SODA)*. 742–755.
- Aaron Bernstein and David Karger. 2009. A nearly optimal oracle for avoiding failed vertices and edges. In *Proc. 41st ACM Symp. on Theory of computing (STOC)*. 101–110.
- Davide Bilò, Fabrizio Grandoni, Luciano Gualà, Stefano Leucci, and Guido Proietti. 2015. Improved purely additive fault-tolerant spanners. In *Proc. European Symp. on Algorithms (ESA)*. SV, 167–178.
- Davide Bilò, Luciano Gualà, Stefano Leucci, and Guido Proietti. 2014. Fault-tolerant approximate shortest-path trees. In *Proc. European Symp. on Algorithms (ESA)*. SV, 137–148.
- Greg Bodwin, Fabrizio Grandoni, Merav Parter, and Virginia Vassilevska Williams. 2017. Preserving distances in very faulty graphs. *arXiv preprint arXiv:1703.10293* (2017).
- Glencora Borradaile, Piotr Sankowski, and Christian Wulff-Nilsen. 2015. Min st-cut oracle for planar graphs with near-linear preprocessing time. *ACM Trans. Algorithms* 11, 3 (2015), 16.
- Gilad Braunschvig, Shiri Chechik, David Peleg, and Adam Sealfon. 2015. Fault tolerant additive and (μ, α) -spanners. *Theoretical Computer Science* 580 (2015), 94–100.
- Shiri Chechik. 2013. New additive spanners. In *Proc. 24th ACM-SIAM Symp. on Discrete algorithms (SODA)*. 498–512.
- Shiri Chechik, Michael Langberg, David Peleg, and Liam Roditty. 2010. Fault tolerant spanners for general graphs. *SIAM J. Computing* 39, 7 (2010), 3403–3423.
- Shiri Chechik, Michael Langberg, David Peleg, and Liam Roditty. 2012. f-sensitivity distance Oracles and routing schemes. *Algorithmica* 63, 4 (2012), 861–882.
- Michael Dinitz and Robert Krauthgamer. 2011. Fault-tolerant spanners: better and simpler. In *Proc. 30th ACM SIGACT-SIGOPS Symp. on Principles of distributed computing (PODC)*. 169–178.
- Michael Elkin and David Peleg. 2004. $(1 + \epsilon, \beta)$ -Spanner Constructions for General Graphs. *SIAM J. Computing* 33, 3 (2004), 608–631.
- Paul Erdős. 1964. Extremal problems in graph theory. In *Proc. Symp. on Theory of Graphs and its Applications*.
- Chris Godsil and Gordon F Royle. 2013. *Algebraic graph theory*. Vol. 207. Springer Science & Business Media.
- John Hershberger and Subhash Suri. 2001. Vickrey prices and shortest paths: What is an edge worth?. In *Proc 42nd IEEE Symp. on Foundations of Computer Science (FOCS)*. 252–259.
- Merav Parter. 2014. Vertex fault tolerant additive spanners. In *Distributed Computing*. SV, 167–181.

- Merav Parter and David Peleg. 2013. Sparse fault-tolerant BFS trees. In *Proc. European Symp. on Algorithms (ESA)*. SV, 779–790.
- Merav Parter and David Peleg. 2014a. Fault Tolerant Approximate BFS Structures. *arXiv preprint arXiv:1406.6169* (2014).
- Merav Parter and David Peleg. 2014b. Fault tolerant approximate BFS structures. In *Proc. 25th ACM-SIAM Symp. on Discrete Algorithms (SODA)*. 1073–1092.
- Mihai Pătraşcu and Liam Roditty. 2010. Distance oracles beyond the Thorup-Zwick bound. In *Proc. 51st IEEE Symp. on Foundations of Computer Science (FOCS)*. 815–823.
- Rephael Wenger. 1991. Extremal graphs with no C_4 's, C_6 's, or C_{10} 's. *J. Combinatorial Theory, Ser. B* 52, 1 (1991), 113–116.
- David P Woodruff. 2006. Lower bounds for additive spanners, emulators, and more. In *Proc. 47th IEEE Symp. on Foundations of Computer Science (FOCS)*. 389–398.