Preserving Distances in Very Faulty Graphs*

Greg Bodwin[1], Fabrizio Grandoni[2], Merav Parter[3], and Virginia Vassilevska Williams[4]

1 **CSAIL, MIT, USA**
`gbodwin@mit.edu`
2 **IDSIA, USI-SUPSI, Switzerland**
`fabrizio@idsia.ch`
3 **CSAIL, MIT, USA**
`parter@mit.edu`
4 **CSAIL, MIT, USA**
`virgi@mit.edu`

Abstract

Preservers and additive spanners are sparse (hence cheap to store) subgraphs that preserve the distances between given pairs of nodes exactly or with some small additive error, respectively. Since real-world networks are prone to failures, it makes sense to study fault-tolerant versions of the above structures. This turns out to be a surprisingly difficult task. For every small but arbitrary set of edge or vertex failures, the preservers and spanners need to contain *replacement paths* around the faulted set. Unfortunately, the complexity of the interaction between replacement paths blows up significantly, even from 1 to 2 faults, and the structure of optimal preservers and spanners is poorly understood. In particular, no nontrivial bounds for preservers and additive spanners are known when the number of faults is bigger than 2.

Even the answer to the following innocent question is completely unknown: what is the worst-case size of a preserver for a *single pair* of nodes in the presence of f edge faults? There are no super-linear lower bounds, nor subquadratic upper bounds for $f > 2$. In this paper we make substantial progress on this and other fundamental questions:

- We present the first truly sub-quadratic size fault-tolerant single-pair preserver in unweighted (possibly directed) graphs: for any n node graph and *any* fixed number f of faults, $\tilde{O}(fn^{2-1/2^f})$ size suffices. Our result also generalizes to the single-source (all targets) case, and can be used to build new fault-tolerant additive spanners (for all pairs).
- The size of the above single-pair preserver grows to $O(n^2)$ for increasing f. We show that this is necessary even in undirected unweighted graphs, and even if you allow for a small additive error: If you aim at size $O(n^{2-\varepsilon})$ for $\varepsilon > 0$, then the additive error has to be $\Omega(\varepsilon f)$. This surprisingly matches known upper bounds in the literature.
- For weighted graphs, we provide matching upper and lower bounds for the single pair case. Namely, the size of the preserver is $\Theta(n^2)$ for $f \geq 2$ in both directed and undirected graphs, while for $f = 1$ the size is $\Theta(n)$ in undirected graphs. For directed graphs, we have a superlinear upper bound and a matching lower bound.

Most of our lower bounds extend to the distance oracle setting, where rather than a subgraph we ask for any compact data structure.

1998 ACM Subject Classification G.2.2 Graph Theory – I.1.2 Analysis of algorithms – B.8.1 Reliability, Testing, and Fault-Tolerance

* A full version of this paper is available at https://arxiv.org/abs/1703.10293. The first and fourth author were partially supported by NSF Grants CCF-1417238, CCF-1528078 and CCF-1514339, and BSF Grant BSF:2012338. The second author was partially supported by the ERC Starting Grant NEWNET 279352 and the SNSF Grant APPROXNET 200021_159697/1.

Keywords and phrases Fault Tolerance, shortest paths, replacement paths

Digital Object Identifier 10.4230/LIPIcs...

1 Introduction

Distance preservers and additive spanners are (sparse) subgraphs that preserve, either exactly or with some small additive error, the distances between given critical pairs P of nodes. This has been a subject of intense research in the last two decades [18, 11, 4, 3, 15, 6, 1, 34].

However, real-world networks are prone to failures. For this reason, more recently (e.g. [16, 14, 17, 32, 30, 9, 33, 8, 21, 27, 19, 28]) researchers have devoted their attention to fault-tolerant versions of the above structures, where distances are (approximately) preserved even in the presence of a few edge (or vertex) faults. For the sake of simplicity we focus here on edge faults, but many of these results generalize to the case of vertex faults where $F \subseteq V$.

▶ **Definition 1.** Given an n-node graph $G = (V, E)$ and $P \subseteq V \times V$, a subgraph $H \subseteq G$ is an f-fault tolerant (f-FT) β-additive P-pairwise spanner if

$$dist_{H \setminus F}(s,t) \leq dist_{G \setminus F}(s,t) + \beta, \quad \forall (s,t) \in P, \forall F \subseteq E, |F| \leq f.$$

If $\beta = 0$, then H is an f-FT P-pairwise preserver.

Finding sparse FT spanners/preservers turned out to be an incredibly challenging task. Despite intensive research, many simple questions have remained open, the most striking of which arguably is the following:

▶ **Question 1.** *What is the worst-case size of a preserver for a single pair (s,t) and $f \geq 1$ faults?*

Prior work [31, 32] considered the single-source $P = \{s\} \times V$ unweighted case, providing super-linear lower bounds for any f and tight upper bounds for $f = 1, 2$. However, first, there is nothing known for $f > 2$, and second, the lower bounds for the $\{s\} \times V$ case do not apply to the single pair case where much sparser preservers might exist. Prior to this work, it was conceivable that in this case $O(n)$ edges suffice for arbitrary fixed f.

Our first result is a *complete* answer to Question 1 for weighted graphs. For $f = 1$ and undirected graphs, we show that a $O(n)$ size preserver exists. Our result is achieved by proving the following interesting fact: for any replacement path $P_{s,t,e}$ protecting against a single edge fault e, there is an edge $(x,y) \in P_{s,t,e}$ such that there is no shortest path from s to x in G that includes e, and there is no shortest path from t to y in G that includes e. Therefore it is sufficient to build shortest path trees from s and to t, and then add one extra edge per possible fault e along the shortest path from s to t. With a trivial union bound, we get that any set P of node pairs can be preserved using $O(\min(n|P|, n^2))$ edges. It is natural to wonder if one can improve this union bound by doing something smarter in the construction. Surprisingly, the answer is NO: we are able to provide a matching lower bound.

▶ **Theorem 2.** *For any undirected n-node weighted graph G and any set P of p pairs of nodes, there exists a P-pairwise 1-FT preserver of size $O(\min(np, n^2))$. Furthermore, for any integer $1 \leq p \leq \binom{n}{2}$, there exists an undirected weighted graph G and a set P of p node pairs such that every 1-FT P-pairwise preserver of G contains $\Omega(\min(n|P|, n^2))$ edges.*

The lower bound part obtained by adapting the lower bound of [32] to the weighted case; this allows us to a lower bound graph whose number of edges is a function of the number of pairs.

For $f = 1$ and directed graphs, we achieve the following. Let $\mathtt{DP}(n)$ denote a tight bound for the sparsity of a pairwise distance preserver in directed weighted graphs with n nodes and $O(n)$ pairs.

▶ **Theorem 3.** *For every directed weighted n-node graph $G = (V, E)$ and for every pair of nodes $s, t \in V$, there is a 1-FT (s, t) preserver with $O(\mathtt{DP}(n))$ edges. For every n, there exists a directed weighted n-node graph $G = (V, E)$ and a node pair $s, t \in V$ such that any 1-FT (s, t) preserver for G has $\Omega(\mathtt{DP}(n))$ edges.*

Coppersmith and Elkin [18] show that $\Omega(n^{4/3}) \leq \mathtt{DP}(n) \leq O(n^{3/2})$. It is a major open question to close this gap, and we show that the no-fault n-pair distance preserver question is equivalent to the 1-fault single pair preserver question, thereby fully answering the latter question, up to resolving the major open problem for n-pair preservers.

We show that the situation dramatically changes for $f \geq 2$.

▶ **Theorem 4.** *There exists an undirected weighted graph G and a single node pair (s, t) in this graph such that every 2-FT (s, t) preserver of G requires $\Omega(n^2)$ edges.*

For unweighted graphs, we achieve several non-trivial upper and lower bounds concerning the worst-case size of (s, t) preservers and spanners. First of all, we address the following question.

▶ **Question 2.** *In unweighted graphs, is the worst-case size of an f-FT (s, t) preserver subquadratic for every constant $f \geq 2$?*

Prior work showed that the answer is YES for $f = 1, 2$ [31, 33], but nothing is known for $f \geq 3$. We show that the answer is YES. Indeed, our result is more general. First, it extends to the single-source case (i.e., $P = \{s\} \times V$) and even to a small enough set of sources (i.e., $P = S \times V$ for small $|S|$). Second, the same result holds for *any* fixed number f of vertex faults. Prior work was only able to address the simple case $f = 1$ [30]. We also remark that our preserver can be computed very efficiently in $O(fmn)$ time, and its analysis is relatively simple (e.g., compared to the cumbersome case analysis in [31]).

▶ **Theorem 5.** *For every directed or undirected unweighted graph $G = (V, E)$, integer $f \geq 1$ and $S \subseteq V$, one can construct in time $O(f\,n\,m)$ an f-FT S-sourcewise (i.e. $P = S \times V$) preserver of size $\widetilde{O}(f \cdot |S|^{1/2^f} \cdot n^{2-1/2^f})$, both in the case of edge and vertex faults.*

By standard techniques, we can exploit our S-sourcewise preserver to build an additive spanner (for all pairs): Let L be an integer parameter to be fixed later on. A vertex u is *low-degree* if it has degree less than L, otherwise it is *high-degree*. Let S be a random sample of $\Theta(\frac{n}{L} \cdot f \log n)$ vertices. Our spanner H consists of the f-VFT S-sourcewise preserver from Theorem 5 plus all the edges incident to low-degree vertices. This way we achieve:

▶ **Theorem 6.** *For every undirected unweighted graph $G = (V, E)$ and integer $f \geq 1$, there exists a randomized $\widetilde{O}(fnm)$-time construction of a $+2$-additive f-FT spanner of G of size $\widetilde{O}(f \cdot n^{2-1/(2^f+1)})$ that succeeds w.h.p.*[1].

In the above result the size of the preserver grows quickly to $O(n^2)$ for increasing f. This raises the following new question:

[1] The term w.h.p. (with high probability) here indicates a probability exceeding $1 - 1/n^c$, for an arbitrary constant $c \geq 2$. Since randomization is only used to select hitting sets, the algorithm can be derandomized.

▶ **Question 3.** *Does there exist a universal constant $\varepsilon > 0$ such that all unweighted graphs have an f-FT (s,t) preserver of size $O_f(n^{2-\varepsilon})$? What if we allow a small additive error?*

The only result with *strongly sub-quadratic* size in the above sense is an $O(f \cdot n^{4/3})$ size spanner with additive error $\Theta(f)$ [14, 8]. Can we remove or reduce the dependence of the error on f? We show that the answer is NO:

▶ **Theorem 7.** *For any two integers $q, h > 0$ and a sufficiently large n, there exists an unweighted undirected n-node graph $G = (V, E)$ and a pair $s, t \in V$ such that any $2hq$-FT $(2q-1)$-additive spanner for G for the single pair (s,t) has size $\Omega((\frac{n}{hq})^{2-2/(h+1)})$.*

▶ **Corollary 8.** *For any fixed constants $\varepsilon > 0$ and f, there exists an unweighted undirected n-node graph $G = (V, E)$ and a pair $s, t \in V$ such that any f-FT additive spanner for G for the single pair (s,t) of size $O(n^{2-\varepsilon})$ must have additive error $\Omega(\varepsilon f)$.*

Proof. This follows from Theorem 7 by choosing proper $h = \Theta(1/\varepsilon)$ and $q = \Theta(\varepsilon f)$. ◀

Hence the linear dependence in f in the additive error in [14, 8] is indeed necessary. We found this very surprising.

In Section 3 we present other related lower bounds which exploit the same basic construction plus ideas in [1, 10]: see Theorems 18, 19, and 20. In particular, we are able to achieve super-linear lower bounds for any $f \geq 2$, even if we allow for a small enough polynomial additive error n^{δ}.

So far we have focused on sparse distance preserving *subgraphs*. However, suppose that the distance estimates can be stored in a different way in memory. Data structures that store the distance information of a graph in the presence of faults are called *distance sensitivity oracles*. Distance sensitivity oracles are also intensely studied [20, 7, 38, 26, 22, 23]. Our main goal here[2] is to keep the size of the data structure as small as possible, leading to the following question.

▶ **Question 4.** *How much space do we need to preserve (exactly or with a small additive error) the distances between a given pair of nodes in the presence of f faults?*

Clearly all our preserver/spanner upper bounds extend to the oracle case, however the lower bounds might not: in principle a distance oracle can use much less space than a preserver/spanner with the same accuracy. Our main contribution here are the following incompressibility results:

▶ **Theorem 9.** *There exists an undirected weighted graph G and a single node pair (s,t) in this graph such that every 2-FT distance sensitivity oracle for the single pair (s,t) in G requires $\Omega(n^2)$ bits of space.*

Note that the optimal size for $f = 1$ is $\Theta(n)$ by simple folklore arguments, so our result completes our understanding in this setting.

We are able to achieve a super-linear lower bound for 3 faults even in the case of a small enough *polynomial additive error*: see Theorem 21 in Section 3.

[2] Other typical goals are to minimize preprocessing and query time - we will not address these.

1.1 Related Work

Fault-tolerant spanners were introduced in the geometric setting [27] (see also [28, 19]). FT-spanners with multiplicative stretch are relatively well understood: the error/sparsity for f-FT and f-VFT multiplicative spanners is (up to a small polynomial factor in f) the same as in the nonfaulty case. For f edge faults, Chechik et al. [16] showed how to construct f-FT $(2k-1)$-multiplicative spanners with size $\tilde{O}(fn^{1+\frac{1}{k}})$ for any $f, k \geq 1$. They also construct an f-VFT spanner with the same stretch and larger size. This was later improved by Dinitz and Krauthgamer [21] who showed the construction of f-VFT spanners with $2k-1$ error and $\tilde{O}\left(f^{2-\frac{1}{k}}n^{1+\frac{1}{k}}\right)$ edges.

FT additive spanners were first considered by Braunschvig, Chechik and Peleg in [14] (see also [8] for slightly improved results). They showed that FT $\Theta(f)$-additive spanners can be constructed by combining FT multiplicative spanners with (non-faulty) additive spanners. This construction, however, supports only edge faults. Parter and Peleg showed in [33] a lower bound of $\Omega(n^{1+\varepsilon_\beta})$ edges for single-source FT β-additive spanners. They also provided a construction of single-source FT-spanner with additive stretch 4 and $O(n^{4/3})$ edges that is resilient to one edge fault. The first constructions of FT-additive spanners resilient against *one vertex* fault were given in [30] and later on in [8]. Prior to our work, no construction of FT-additive spanners was known for $f \geq 2$ vertex faults.

As mentioned earlier, the computation of preservers and spanners in the non-faulty case (i.e. when $f = 0$) has been the subject of intense research in the last few decades. The current-best preservers can be found in [18, 11, 12]. Spanners are also well understood, both for multiplicative stretch [4, 25] and for additive stretch [3, 15, 6, 39, 1, 11, 15, 34, 2]. There are also a few results on "mixed" spanners with both multiplicative and additive stretch [24, 36, 6]

Distance sensitivity oracles are data structures that can answer queries about the distances in a given graph in the presence of faults. The first nontrivial construction was given by Demetrescu et al. [20] and later improved by Bernstein and Karger [7] who showed how to construct $\widetilde{O}(n^2)$-space, constant query time oracles for a single edge fault for an m-edge n-node graph in $\widetilde{O}(mn)$ time. The first work that considered the case of two faults (hence making the first jump from one to two) is due to Duan and Pettie in [22]. Their distance oracle has nearly optimal size of $\widetilde{O}(n^2)$ and query time of $\widetilde{O}(1)$. The case of bounded edge weights, and possibly multiple faults, is addressed in [38, 26] exploiting fast matrix multiplication techniques. The size of their oracle is super-quadratic.

The notion of FT-preservers is also closely related to the problem of constructing *replacement paths*. For a pair of vertices s and t and an edge e, the replacement path $P_{s,t,e}$ is the s-t shortest-path that avoids e[3]. The efficient computation of replacement paths is addressed, among others, in [29, 35, 38, 37]. A single-source version of the problem is studied in [26]. Single-source FT structures that preserve strong connectivity have been studied in [5].

1.2 Preliminaries and Notation

Assume throughout that all shortest paths ties are broken in a consistent manner. For every $s, t \in V$ and a subgraph $G' \subseteq G$, let $\pi_{G'}(s, t)$ be the (unique) u-v shortest path in G' (i.e., it is unique under breaking ties). If there is no path between s and t in G', we define

[3] Replacement paths were originally defined for the single edge fault case, but later on extended to the case of multiple faults as well.

$\pi_{G'}(s,t) = \emptyset$. When $G' = G$, we simply write $\pi(u,v)$. For any path P containing nodes u, v, let $P[u \rightsquigarrow v]$ be the subpath of P between u and v. For $s, t \in V$ and $F \subseteq E$, we let $P_{s,t,F} = \pi_{G \setminus F}(s,t)$ be the s-t shortest-path in $G \setminus F$. We call such paths *replacement paths*. When $F = \{e\}$, we simply write $P_{s,t,e}$. By m we denote the number of edges in the graph currently being considered.

The structure of the paper is as follows. In Sec. 2, we describe an efficient construction for FT-preservers and additive spanners with a subquadratic number of edges. Then, in Sec. 3, we provide several lower bound constructions for a single s-t pair, both for the exact and for the additive stretch case. All the proofs which are omitted due to lack of space appear in the full version of the paper (see [13]).

2 Efficient Construction of FT-Preservers and Spanners

In this section we prove Theorem 5. We next focus on the directed case, the undirected one being analogous and simpler. We begin by recapping the currently-known approaches for handling many faults, and we explain why these approaches fail to achieve interesting space/construction time bounds for large f.

The limits of previous approaches: A known approach for handling many faults is by *random sampling* of subgraphs, as introduced by Weimann and Yuster [38] in the setting of distance sensitivity oracles, and later on applied by Dinitz and Kraughgamer [21] in the setting of fault tolerant spanners. The high level idea is to generate multiple subgraphs $G_1, \ldots, G_r$ by removing each edge/vertex independently with sufficiently large probability p; intuitively, each G_i simultaneously captures many possible fault sets of size f. One can show that, for a sufficiently small parameter L and for any given (*short*) replacement path $P_{s,t,F}$ of length at most L (avoiding faults F), w.h.p. in at least one G_i the path $P_{s,t,F}$ is still present while all edges/vertices in F are deleted. Thus, if we compute a (non-faulty) preserver $H_i \subseteq G_i$ for each i, then the graph $H = \bigcup_i H_i$ will contain every short replacement path. For the remaining (*long*) replacement paths, Weimann and Yuster use a random decomposition into short subpaths. Unfortunately, any combination of the parameters p, r, L leads to a quadratic (or larger) space usage.

Another way to handle multiple faults is by extending the approach in [32, 33, 30] that works for $f \in \{1, 2\}$. A useful trick used in those papers (inspired by prior work in [35, 37]) is as follows: suppose $f = 1$, and fix a target node t. Consider the shortest path $\pi(s,t)$. It is sufficient to take the *last* edge of each replacement path $P_{s,t,e}$ and charge it to the node t; the rest of the path is then charged to other nodes by an inductive argument. Hence, one only needs to bound the number of *new-ending* paths – those that end in an edge that is not already in $\pi(s,t)$. In the case $f = 1$, these new-ending paths have a nice structure: they diverge from $\pi(s,t)$ at some vertex b (*divergence point*) above the failing edge/vertex and collide again with $\pi(s,t)$ only at the terminal t; the subpath connecting b and t on the replacement path is called its *detour*. One can divide the s-t replacement paths into two groups: short (resp., long) paths are those whose detour has length at most (resp., at least) $\sqrt{n}$. It is then straightforward enough to show that each category of path contributes only $\widetilde{O}(n^{1/2})$ edges entering t, and so (collecting these last edges over all nodes in the graph) the output subgraph has $\widetilde{O}(n^{3/2})$ edges in total. Generalizing this to the case of multiple faults is non-trivial already for the case of $f = 2$. The main obstacle here stems from a lack of structural understanding of replacement paths for multiple faults: in particular, any given divergence point $b \in \pi(s,t)$ can now be associated with many new-ending paths and not only one! In the only known positive solution for $f = 2$ [31], the approach works only for edge faults and is based on an extensive case analysis whose extension to larger f is beyond reasonable reach. Thus, in the absence of new structural understanding, further progress seems very difficult.

A second source of difficulties is related to the *running time* of the construction. A priori, it seems that constructing a preserver H should require computing all replacement paths $P_{s,t,F}$, which leads to a construction time that scales *exponentially* in f. In particular, by deciding to omit an edge e from the preserver H, we must somehow check that this edge does not appear on *any* of the replacement paths $P_{s,t,F}$ (possibly, without computing these replacement paths explicitly).

Our basic approach: The basic idea behind our algorithm is as follows. Similar to [32, 33, 30], we focus on each target node t, and define a set E_t of edges incident to t to be added to our preserver. Intuitively, these are the last edges of new-ending paths as described before. The construction of E_t, however, deviates substantially from prior work. Let us focus on the simpler case of edge deletions. The set E_t is constructed recursively, according to parameter f. Initially we consider the shortest path tree T from the source set S to t, and add to E_t the edges of T incident to t (at most $|S|$ many). Consider any new-ending replacement path P for t. By the previous discussion, this path has to leave T at some node b and it meets T again only at t: let D be the subpath of P between b and t (the *detour* of P). Note that D is edge-disjoint from T, i.e. it is contained in the graph $G' = G \setminus E(T)$. Therefore, it would be sufficient to compute recursively the set E'_t of final edges of new-ending replacement paths for t in the graph G' with source set S' given by the possible divergence points b and w.r.t. $f-1$ faults (recall that one fault must be in $E(T)$, hence we avoid that anyway in G'). This set E'_t can then be added to E_t.

The problem with this approach is that S' can contain $\Omega(n)$ many divergence points (hence E_t $\Omega(n)$ many edges), leading to a trivial $\Omega(n^2)$ size preserver. In order to circumvent this problem, we classify the divergence points b in two categories. Consider first the nodes b at distance at most L from t along T, for some parameter L. There are only $O(|S|L)$ many such nodes S^{short}, which is sublinear for $|S|$ and L small enough. Therefore we can safely add S^{short} to S'. For the remaining divergence points b, we observe that the corresponding detour D must have length at least L: therefore by sampling $\tilde{O}(n/L)$ nodes S^{long} we hit all such detours w.h.p. Suppose that $\sigma \in S^{long}$ hits detour D. Then the portion of D from σ to t also contains the final edge of D to be added to E_t. In other terms, it is sufficient to add S^{long} (which has sublinear size for polynomially large L) to S' to cover all the detours of nodes b of the second type. Altogether, in the recursive call we need to handle one less fault w.r.t. a larger (but sublinear) set of sources S'. Our approach has several benefits:

- It leads to a subquadratic size for any f (for a proper choice of the parameters);
- It leads to a very fast algorithm. In fact, for each target t we only need to compute a BFS tree in f different graphs, leading to an $O(fnm)$ running time;
- Our analysis is very simple, much simpler than in [31] for the case $f = 2$;
- It can be easily extended to the case of vertex faults.

Algorithm for Edge Faults: Let us start with the edge faults case. The algorithm constructs a set E_t of edges incident to each target node $t \in V$. The final preserver is simply the union $H = \bigcup_{t \in V} E_t$ of these edges. We next describe the construction of each E_t (see also Alg. 1). The computation proceeds in rounds $i = 0, \ldots, f$. At the beginning of round i we are given a subgraph G_i (with $G_0 = G$) and a set of sources S_i (with $S_0 = S$).

Algorithm 1 Construction of E_t in our f-FT S-Sourcewise Preserver Algorithm.

procedure ComputeSourcewiseFT(t, S, f, G)
  Input: A graph G with a source set S and terminal t, number of faults f.
  Output: Edges E_t incident to t in an f-FT S-sourcewise preserver H.
Set $G_0 = G$, $S_0 = S$, $E_t = \emptyset$.
  for $i \in \{0, \ldots, f\}$ **do**
    Compute the partial BFS tree $T_i = \bigcup_{s \in S_i} \pi_{G_i}(s,t)$.
    $E_t = E_t \cup \{\texttt{LastE}(\pi_{T_i}(s,t)) \mid s \in S_i\}$.
    Set distance threshold $d_i = \sqrt{n/|S_i| \cdot f \log n}$.
    Let $S_i^{short} = \{v \in V(T_i) \mid \texttt{dist}_{T_i}(v,t) \leq d_i\}$.
    Sample a collection $S_i^{long} \subseteq V(G_i)$ of $\Theta(n/d_i \cdot f \log n)$ vertices.
    Set $S_{i+1} = S_i^{short} \cup S_i^{long}$ and $G_{i+1} = G_i \setminus E(T_i)$.

We compute a partial BFS tree $T_i = \bigcup_{s \in S_i} \pi_{G_i}(s,t)$[4] from S_i to t, and add to E_t (which is initially empty) the edges $\{\texttt{LastE}(\pi_{T_i}(s,t)) \mid s \in S_i\}$ of this tree incident to t. Here, for a path π where one endpoint is the considered target node t, we denote by $\texttt{LastE}(\pi)$ the edge of π incident to t. The source set S_{i+1} is given by $S_i^{short} \cup S_i^{long}$. Here $S_i^{short} = \{v \in V(T_i) \mid \texttt{dist}_{T_i}(v,t) \leq d_i\}$ is the set of nodes at distance at most $d_i = \sqrt{n/|S_i| \cdot f \log n}$ from t, while S_i^{long} is a random sample of $\Theta(n/d_i \cdot f \log n)$ vertices. The graph G_{i+1} is obtained from G_i be removing the edges $E(T_i)$[5].

Adaptation for Vertex Faults: The only change in the algorithm is in the definition of the graph G_i inside the procedure to compute E_t. We cannot allow ourselves to remove all the vertices of the tree T_i from G_i and hence a more subtle definition is required. To define G_{i+1}, we first remove from G_i: (1) all edges of $S_i^{short} \times S_i^{short}$, (2) the edges of $E(T_i)$, and (3) the vertices of $V(T_i) \setminus S_i^{short}$. Finally, we delete all remaining edges incident to S_i^{short} which are directed towards any one of these vertices (i.e., the incoming degree of the S_i^{short} vertices in G_{i+1} is zero).

Analysis: We now analyze our algorithm. Since for each vertex t, we compute f (partial) BFS trees, we get trivially:

▶ **Lemma 10** (Running Time). *The subgraph H is computed within $O(f\,n\,m)$ time.*

We proceed with bounding the size of H.

▶ **Lemma 11** (Size Analysis). $|E_t| = \widetilde{O}(|S|^{1/2^f} \cdot (fn)^{1-1/2^f})$ *for every* $t \in V$*, hence* $|E(H)| = \widetilde{O}(f|S|^{1/2^f} n^{2-1/2^f})$.

Proof. Since the number of edges collected at the end each round i is bounded by the number of sources S_i, it is sufficient to bound $|S_i|$ for all i. Observe that for every $i \in \{0, \ldots, f-1\}$,

$$|S_{i+1}| \leq |S_i^{long}| + |S_i^{short}| \leq d_i \cdot |S_i| + \Theta(n/d_i \cdot f \log n) = \Theta(d_i \cdot |S_i|).$$

By resolving this recurrence starting with $|S_0| = |S|$ one obtains

$$|S_i| = O(|S|^{1/2^i} (fn \log n)^{1-1/2^i}).$$

The claim follows by summing over $i \in \{0, \ldots, f\}$. ◀

[4] If $\pi_{G_i}(s,t)$ does not exist, recall that we define it as an empty set of edges.
[5] Note that for $f = 1$, the algorithm has some similarity to the replacement path computation of [35]. Yet, there was no prior extension of this idea for $f \geq 2$.

We next show that the algorithm is correct. We focus on the vertex fault case, the edge fault case being similar and simpler. Let us define, for $t \in V$ and $i \in \{0, \dots, f\}$,

$$\mathcal{P}_{t,i} = \{\pi_{G_i \setminus F}(s,t) \ \mid \ s \in S_i,\ F \subseteq V(G_i),\ |F| \leq f-i\}.$$

▶ **Lemma 12.** *For every $t \in V$ and $i \in \{0, \dots, f\}$, it holds that*

$$\texttt{LastE}(\pi) \in E_t \ \textit{ for every } \ \pi \in \mathcal{P}_{t,i}.$$

Proof. We prove the claim by decreasing induction on $i \in \{f, \dots, 0\}$. For the base case $i = f$, $\mathcal{P}_{t,f} = \{\pi_{G_f}(s,t) \ \mid \ s \in S_f\}$. Since we add precisely the last edges of these paths to the set E_t, the claim holds. Assume that the lemma holds for rounds $f, f-1, \dots, i+1$ and consider round i. For every $\pi_{G_i \setminus F}(s,t) \in \mathcal{P}_{t,i}$, let $P'_{s,t,F} = \pi_{G_i \setminus F}(s,t)$. [6] Consider the partial BFS tree $T_i = \bigcup_{s \in S_i} \pi_{G_i}(s,t)$ rooted at t. Note that all (interesting) replacement paths $P'_{s,t,F}$ in G_i have at least one failing vertex $v \in F \cap V(T_i)$ as otherwise $P'_{s,t,F} = \pi_{G_i}(s,t)$.

We next partition the replacement paths $\pi \in \mathcal{P}_{t,i}$ into two types depending on their last edge $\texttt{LastE}(\pi)$. The first class contains all paths whose last edge is in T_i. The second class contains the remaining replacement paths, which end with an edge that is not in T_i. We call this second class of paths *new-ending* replacement paths. Observe that the first class is taken care of, since we add all edges incident to t in T_i. Hence it remains to prove the lemma for the set of new-ending paths.

For every new-ending path $P'_{s,t,F}$, let $b_{s,t,F}$ be the last vertex on $P'_{s,t,F}$ that is in $V(T_i) \setminus \{t\}$. We call the vertex $b_{s,t,F}$ the *last divergence point* of the new-ending replacement path. Note that the detour $D_{s,t,F} = P'_{s,t,F}[b_{s,t,F} \rightsquigarrow t]$ is vertex disjoint with the tree T_i except for the vertices $b_{s,t,F}$ and t. From now on, since we only wish to collect last edges, we may restrict our attention to this detour subpath. That is, since $\texttt{LastE}(D_{s,t,F}) = \texttt{LastE}(P'_{s,t,F})$, it is sufficient to show that $\texttt{LastE}(D_{s,t,F}) \in E_t$.

Our approach is based on dividing the set of new-ending paths in $\mathcal{P}_{t,i}$ into two classes based on the position of their last divergence point $b_{s,t,F}$. The first class $\mathcal{P}_{short}$ consists of new-ending paths in $\mathcal{P}_{t,i}$ whose last divergence point is at distance at most $d_i = \sqrt{n/|S_i| \cdot f \log n}$ from t on T_i. In other words, this class contains all new-ending paths whose last divergence point is in the set S_i^{short}. We now claim the following.

▶ **Claim 13.** *For every $P'_{s,t,F} \in \mathcal{P}_{short}$, the detour $D_{s,t,F}$ is in $\mathcal{P}_{t,i+1}$.*

Proof. Since $D_{s,t,F}$ is a subpath of the replacement path $P'_{s,t,F}$, $D_{s,t,f}$ is the shortest path between $b_{s,t,F}$ and t in $G_i \setminus F$. Recall that $D_{s,t,F}$ is vertex disjoint with $V(T_i) \setminus \{b_{s,t,F}, t\}$.

Since $b_{s,t,F}$ is the last divergence point of $P'_{s,t,F}$ with T_i, the detour $D_{s,t,F}$ starts from a vertex $b_{s,t,F} \in S_i^{short}$ and does not pass through any other vertex in $V(T_i) \setminus \{t\}$. Recall that in the construction of G_{i+1} we delete from G_i the edges directed *towards* S_i^{short}. In particular, the outgoing edge connecting $b_{s,t,F}$ to its neighbor x on $D_{s,t,F}[b_{s,t,F} \rightsquigarrow t]$ remains (i.e., this vertex x is not in $V(T_i) \setminus \{t\}$), this implies that the detour $D_{s,t,F}$ exists in G_{i+1}. In particular, note that the vertex $b_{s,t,F}$ cannot be a neighbor of t in T_i. Indeed, if $(b_{s,t,F}, t)$ were an edge in T_i, then we can replace the portion of the detour path between $b_{s,t,F}$ and t by this edge, getting a contradiction to the fact that $P'_{s,t,F}$ is a new-ending path[7].

Next, observe that at least one of the failing vertices in F occurs on the subpath $\pi_{G_i}[b_{s,t,F}, t]$, let this vertex be $v \in F$. Since $v \in S_i^{short}$, all the edges incident to v are

[6] We denote these replacement paths as $P'_{s,t,F}$ as they are computed in G_i and not in G.

[7] For the edge fault case, the argument is much simpler: by removing $E(T_i)$ from G_i, we avoid at least one of the failing edges in G_{i+1}.

directed away from v in G_{i+1} and hence the paths going out from the source $b_{s,t,F}$ in G_{i+1} cannot pass through v. Letting $F' = F \setminus V(T_i)$, it holds that (1) $|F'| \leq f - i - 1$ and (2) since the shortest path ties are decided in a consistent manner and by definition of G_{i+1}, it holds that $D_{s,t,F} = \pi_{G_{i+1} \setminus F'}(b_{s,t,F}, t)$. As $b_{s,t,F} \in S_i^{short}$, it holds that $D_{s,t,F} \in \mathcal{P}_{t,i+1}$. ◀

Hence by the inductive hypothesis for $i+1$, $\texttt{LastE}(P'_{s,t,F})$ is in E_t for every $P'_{s,t,F} \in \mathcal{P}_{short}$. We now turn to consider the second class of paths $\mathcal{P}_{long}$ which contains all remaining new-ending paths; i.e., those paths whose last divergence point is at distance at least d_i from t on T_i. Note that the detour $D_{s,t,F} = P'_{s,t,F}[b_{s,t,F} \rightsquigarrow t]$ of these paths is long – i.e., its length is at least d_i. For convenience, we will consider the internal part $D'_{s,t,F} = D_{s,t,F} \setminus \{b_{s,t,F}, t\}$ of these detours, so that the first and last vertices of these detours are not on T_i.

We exploit the lengths of these detours $D'_{s,t,F}$ and claim that w.h.p, the set S_i^{long} is a hitting set for these detours. This indeed holds by simple union bound overall possible $O(n^{f+2})$ detours. For every $P'_{s,t,F} \in \mathcal{P}_{long}$, let $w_{s,t,F} \in V(D'_{s,t,F}) \cap S_i^{long}$. (By the hitting set property, w.h.p., $w_{s,t,F}$ is well defined for each long detour). Let $W_{s,t,F} = P'_{s,t,F}[w_{s,t,F}, t]$ be the suffix of the path $P'_{s,t,F}$ starting at a vertex from the hitting set $w_{s,t,F} \in S_i^{long}$. Since $\texttt{LastE}(P'_{s,t,F}) = \texttt{LastE}(W_{s,t,F})$, it is sufficient to show that $\texttt{LastE}(W_{s,t,F})$ is in E_t.

▶ **Claim 14.** *For every $P'_{s,t,F} \in \mathcal{P}_{long}$, it holds that $W_{s,t,F} \in \mathcal{P}_{t,i+1}$.*

Proof. Clearly, $W_{s,t,f}$ is the shortest path between $w_{s,t,F}$ and t in $G_i \setminus F$. Since $W_{s,t,F} \subseteq D'_{s,t,F}$ is vertex disjoint with $V(T_i)$, it holds that $W_{s,t,F} = \pi_{G_{i+1} \setminus F'}(w_{s,t,F}, t)$ for $F' = F \setminus V(T_i)$. Note that since at least one fault occurred on T_i, we have that $|F'| \leq f - i - 1$. As $w_{s,t,F} \in S_i^{long}$, it holds that $W_{s,t,F} \in \mathcal{P}_{t,i+1}$. The lemma follows. ◀

By applying the claim for $i = 0$, we get that $\texttt{LastE}(P'_{s,t,F})$ is in E_t as required for every $P'_{s,t,F} \in \mathcal{P}_{long}$. This completes the proof. ◀

▶ **Lemma 15.** *(Correctness) H is an f-FT S-sourcewise preserver.*

Proof. By using Lemma 12 with $i = 0$, we get that for every $t \in V$, $s \in S$ and $F \subseteq V$, $|F| \leq f$, $\texttt{LastE}(P_{s,t,F}) \in E_t$ (and hence also $\texttt{LastE}(P_{s,t,F}) \in H$). It remains to show that taking the last edge of each replacement path $P_{s,t,F}$ is sufficient. The base case is for paths of length 1, where we have clearly kept the entire path in our preserver. Then, assuming the hypothesis holds for paths up to length $k-1$, consider a path $P_{s,t,F}$ of length k. Let $\texttt{LastE}(P_{s,t,F}) = (u, t)$. Then since we break ties in a consistent manner, $P_{s,t,F} = P_{s,u,F} \circ \texttt{LastE}(P_{s,t,e})$. By the inductive hypothesis $P_{s,u,F}$ is in H, and since we included the last edge, $P_{s,t,F}$ is also in H. The claim follows. ◀

Theorem 5 now immediately follows from Lemmas 10, 11, and 15.

3 Lower Bounds for FT Preservers and Additive Spanners

In this section, we provide the first non-trivial lower bounds for preservers and additive spanners for a single pair s-t.

We start by proving Theorem 7. The main building block in our lower bound is the construction of an (undirected unweighted) tree T^h, where h is a positive integer parameter related to the desired number of faults f. Tree T^h is taken from [31] with mild technical adaptations. Let d be a *size* parameter which is used to obtain the desired number n of nodes. It is convenient to interpret this tree as rooted at a specific node (though edges in this construction are undirected). We next let $rt(T^h)$ and $L(T^h)$ be the root and leaf set of T^h, respectively. We also let $\ell(h)$ and $n(h)$ be the height and number of nodes of T^h, respectively.

Tree T^h is constructed recursively as follows. The base case is given by T^0 which consists of a single isolated root node $rt(T^0)$. Note that $\ell(0) = 0$ and $n(0) = 1$. In order to construct T^h, we first create d copies $T_0^{h-1}, \ldots, T_{d-1}^{h-1}$ of T^{h-1}. Then we add a path $v_0, \ldots, v_{d-1}$ of length $d-1$ (consisting of new nodes), and choose $rt(T^h) = v_0$. Finally, we connect v_j to $rt(T_j^{h-1})$ with a path (whose internal nodes are new) of length $(d-j)\cdot(\ell(h-1)+3)$. Next lemma illustrates the crucial properties of T^h.

▶ **Lemma 16.** *The tree T^h satisfies the following properties:*
1. $n(h) \leq \frac{3}{2}(h+1)(d+1)^{h+1}$
2. $|L(T^h)| = d^h$
3. *For every $\ell \in L(T^h)$, there exists $F_\ell \subseteq E(T)$, $|F_\ell| = h$, such that $\mathtt{dist}_{T^h \setminus F_\ell}(s, \ell) \leq \mathtt{dist}_{T^h \setminus F_\ell}(s, \ell') + 2$ for every $\ell' \in L(T^h) \setminus \{\ell'\}$.*

We next construct a graph S^h as follows. We create two copies T_s and T_t of T^h. We add to S^h the complete bipartite graph with sides $L(T_s)$ and $L(T_t)$, which we will call the *bipartite core* B of S^h. Observe that $|L(T_s)| = |L(T_t)| = d^h$, and hence B contains d^{2h} edges. We will call $s = sr(S^h) = rt(T_s)$ the source of S^h, and $t = tg(S^h) = rt(T_t)$ its target.

▶ **Lemma 17.** *Every $2h$-FT (s,t) preserver (and 1-additive (s,t) spanner) H for S^h must contain each edge $e = (\ell_s, \ell_t) \in B$.*

Proof. Assume that $e = (\ell_s, \ell_t) \notin H$ and consider the case where F_{ℓ_s} fails in T_s and F_{ℓ_t} fails in T_t. Let $G' := S^h \setminus (F_{\ell_s} \cup F_{\ell_t})$, and d_s (resp., d_t) be the distance from s to ℓ_s (resp., from ℓ_t to t) in G'. By Lemma 16.3 the shortest s-t path in G' passes through e and has length $d_s + 1 + d_t$. By the same lemma, any path in G', hence in $H' := H \setminus (F_{\ell_s} \cup F_{\ell_t})$, that does not pass through ℓ_s (resp., ℓ_t) must have length at least $(d_s+2)+1+d_t$ (resp., $d_s + 1 + (d_t + 2)$). On the other hand, any path in H' that passes through ℓ_s and ℓ_t must use at least 3 edges of B, hence having length at least $d_s + 3 + d_t$. ◀

Our lower bound graph S_q^h is obtained by taking q copies $S_1, \ldots, S_q$ of graph S^h with $d = (\frac{n}{3q(h+1)} - 1)^{\frac{1}{h+1}}$, and chaining them with edges $(tg(S_i), sr(S_{i+1}))$, for $i = 1, \ldots, q-1$. We let $s = sr(S_1)$ and $t = tg(S_q)$.

Proof of Theorem 7. Consider S_q^h. By Lemma 16.1-2 this graph contains at most n nodes, and the bipartite core of each S_i contains $d^{2h} = \Omega((\frac{n}{qh})^{2-2/(h+1)})$ edges.

Finally, we show that any $(2q-1)$-additive (s,t) spanner needs to contain all the edges of at least one such bipartite core. Let us assume this does not happen, and let e_i be a missing edge in the bipartite core of S_i for each i. Observe that each s-t shortest path has to cross $sr(S_i)$ and $tg(S_i)$ for all i. Therefore, it is sufficient to choose $2h$ faulty edges corresponding to each e_i as in Lemma 17. This introduces an additive stretch of 2 in the distance between s and t for each e_i, leading to a total additive stretch of at least $2q$. ◀

The same construction can also be extended to the setting of $(2h)$-FT $S \times T$ preservers. To do that, we make *parallel* copies of the S^h graph.

▶ **Theorem 18.** *For every positive integer f, there exists a graph $G = (V, E)$ and subsets $S, T \subseteq V$, such that every $(2f)$-FT 1-additive $S \times T$ spanner (hence $S \times T$ preserver) of G has size $\Omega(|S|^{1/(f+1)} \cdot |T|^{1/(f+1)} \cdot (n/f)^{2-2/(f+1)})$.*

Improving over the Bipartite Core: The proof above only gives the trivial lower bound of $\Omega(n)$ for the case of two faults (using $h = q = 1$). We can strengthen the proof in this special case to show instead that $\Omega(n^{1+\varepsilon})$ edges are needed, and indeed this even holds in the presence of a *polynomial additive stretch*:

▶ **Theorem 19.** *A 2-FT distance preserver of a single (s,t) pair in an undirected unweighted graph needs $\Omega(n^{11/10-o(1)})$ edges.*

▶ **Theorem 20.** *There are absolute constants $\varepsilon, \delta > 0$ such that any $+n^{\delta}$-additive 2-FT preserver for a single (s,t) pair in an undirected unweighted graph needs $\Omega(n^{1+\varepsilon})$ edges.*

Finally, by tolerating one additional fault, we can obtain a strong incompressibility result:

▶ **Theorem 21.** *There are absolute constants $\varepsilon, \delta > 0$ such that any $+n^{\delta}$-additive 3-FT distance sensitivity oracle for a single (s,t) pair in an undirected unweighted graph uses $\Omega(n^{1+\varepsilon})$ bits of space.*

The proofs of Theorems 19, 20 and 21 are similar in spirit. The key observation is that the structure of T_s, T_t allows us to use our faults to select leaves ℓ_s, ℓ_t and enforce that a shortest ℓ_s-ℓ_t path is kept in the graph. When we use a bipartite core between the leaves of T_s and T_t, this "shortest path" is simply an edge, so the quality of our lower bound is equal to the product of the leaves in T_s and T_t. However, sometimes a better graph can be used instead. In the case $h = 1$, we can use a nontrivial lower bound graph against (non-faulty) subset distance preservers (from [10]), which improves the cost per leaf pair from 1 edge to roughly $n^{11/10}$ edges, yielding Theorem 19. Alternatively, we can use a nontrivial lower bound graph against $+n^{\delta}$ spanners (from [1]), which implies Theorem 20. The proof of Theorem 21 is similar in spirit, but requires an additional trick in which *unbalanced trees* are used: we take T_s as a copy of T^1 and T_t as a copy of T^2, and this improved number of leaf-pairs is enough to push the incompressibility argument through.

4 Open Problems

There are lots of open ends to be closed. Perhaps the main open problem is to resolve the current gap for f-FT single-source preservers. Since the lower bound of $\Omega(n^{2-1/(f+1)})$ edges given in [31] has been shown to be tight for $f \in [1, 2]$, it is reasonable to believe that this is the right bound for $f \geq 3$. Another interesting open question involves lower bounds for FT additive spanners. Our lower-bounds are super linear only for $f \geq 2$. The following basic question is still open though: is there a lower bound of $\Omega(n^{3/2+\epsilon})$ edges for some $\epsilon \in (0, 1]$ for 2-additive spanners with *one* fault? Whereas our lower bound machinery can be adapted to provide non trivial bounds for different types of f-FT P-preservers (e.g., $P = \{s, t\}, P = S \times T$, etc.), our upper bounds technique for general $f \geq 2$ is still limited to the sourcewise setting. Specifically, it is not clear how to construct an f-FT $S \times S$ preserver other than taking a (perhaps wasteful) f-FT S-sourcewise preserver. As suggested by our lower bounds, these questions are interesting already for a single pair.

References

1 A. Abboud and G. Bodwin. The 4/3 additive spanner exponent is tight. In *STOC*, pages 351–361, 2016.

2 Amir Abboud, Greg Bodwin, and Seth Pettie. A hierarchy of lower bounds for sublinear additive spanners. In *SODA*, pages 568–576, 2017.

3 D. Aingworth, C. Chekuri, P. Indyk, and R. Motwani. Fast estimation of diameter and shortest paths (without matrix multiplication). *SIAM J. Comput.*, 28(4):1167–1181, 1999.

4 I. Althöfer, G. Das, D. Dobkin, D. Joseph, and J. Soares. On sparse spanners of weighted graphs. *Discrete & Computational Geometry*, 9(1):81–100, 1993.

5 S. Baswana, K. Choudhary, and L. Roditty. Fault tolerant subgraph for single source reachability: generic and optimal. In *STOC*, pages 509–518, 2016.

6 S. Baswana, K. Telikepalli, K. Mehlhorn, and S. Pettie. New constructions of (alpha, beta)-spanners and purely additive spanners. In *SODA*, pages 672–681, 2005.

7 Aaron Bernstein and David R. Karger. A nearly optimal oracle for avoiding failed vertices and edges. In *STOC*, pages 101–110, 2009.

8 D. Bilò, F. Grandoni, L. Gualà, S. Leucci, and G. Proietti. Improved purely additive fault-tolerant spanners. In *ESA*, pages 167–178. 2015.

9 D. Bilò, L. Gualà, S. Leucci, and G. Proietti. Fault-tolerant approximate shortest-path trees. In *ESA*, pages 137–148. 2014.

10 G. Bodwin. Linear size distance preservers. In *SODA*, 2017.

11 G. Bodwin and V. Vassilevska Williams. Better distance preservers and additive spanners. In *SODA*, pages 855–872, 2016.

12 Greg Bodwin. Linear size distance preservers. In *SODA*, pages 600–615, 2017.

13 Greg Bodwin, Fabrizio Grandoni, Merav Parter, and Virginia Vassilevska Williams. Preserving distances in very faulty graphs. *CoRR*, abs/1703.10293, 2017. URL: `http://arxiv.org/abs/1703.10293`.

14 G. Braunschvig, S. Chechik, D. Peleg, and A. Sealfon. Fault tolerant additive and (μ, α)-spanners. *Theor. Comput. Sci.*, 580:94–100, 2015.

15 S. Chechik. New additive spanners. In *SODA*, pages 498–512, 2013.

16 S. Chechik, M. Langberg, D. Peleg, and L. Roditty. Fault-tolerant spanners for general graphs. In *STOC*, pages 435–444, 2009.

17 S. Chechik and D. Peleg. Rigid and competitive fault tolerance for logical information structures in networks. In *Electrical and Electronics Engineers in Israel (IEEEI), 2010 IEEE 26th Convention of*, pages 000024–000025. IEEE, 2010.

18 D. Coppersmith and M. Elkin. Sparse sourcewise and pairwise distance preservers. *SIAM Journal on Discrete Mathematics*, 20(2):463–501, 2006.

19 A. Czumaj and H. Zhao. Fault-tolerant geometric spanners. *Discrete & Computational Geometry*, 32(2):207–230, 2004.

20 C. Demetrescu, M. Thorup, R. A. Chowdhury, and V. Ramachandran. Oracles for distances avoiding a failed node or link. *SIAM Journal on Computing*, 37(5):1299–1318, 2008.

21 M. Dinitz and R. Krauthgamer. Fault-tolerant spanners: better and simpler. In *PODC*, pages 169–178, 2011.

22 R. Duan and S. Pettie. Dual-failure distance and connectivity oracles. In *SODA*, pages 506–515, 2009.

23 Ran Duan and Seth Pettie. Connectivity oracles for graphs subject to vertex failures. In *SODA*, pages 490–509, 2017.

24 M. Elkin and D. Peleg. (1+epsilon, beta)-spanner constructions for general graphs. *SIAM J. Comput.*, 33(3):608–631, 2004.

25 P. Erdös. Extremal problems in graph theory. *Theory of Graphs and its Applications (Proc. Sympos. Smolenice, 1963)*, pages 29–36, 1963.

26 F. Grandoni and V. Vassilevska Williams. Improved distance sensitivity oracles via fast single-source replacement paths. In *FOCS*, pages 748–757, 2012.

27 C. Levcopoulos, G. Narasimhan, and M. Smid. Improved algorithms for constructing fault-tolerant spanners. *Algorithmica*, 32(1):144–156, 2002.

28 T. Lukovszki. New results on fault tolerant geometric spanners. In *Algorithms and Data Structures*, pages 193–204. Springer, 1999.

29 K. Malik, A. K. Mittal, and S. K. Gupta. The k most vital arcs in the shortest path problem. *Operations Research Letters*, 8(4):223–227, 1989.

30 M. Parter. Vertex fault tolerant additive spanners. In *Distributed Computing*, pages 167–181. Springer, 2014.

31 M. Parter. Dual failure resilient BFS structure. In *PODC*, pages 481–490, 2015.

32 M. Parter and D. Peleg. Sparse fault-tolerant BFS trees. In *ESA*, pages 779–790, 2013.

33 M. Parter and D. Peleg. Fault tolerant approximate BFS structures. In *SODA*, pages 1073–1092, 2014.

34 S. Pettie. Low distortion spanners. *ACM Transactions on Algorithms*, 6(1), 2009.

35 L. Roditty and U. Zwick. Replacement paths and k simple shortest paths in unweighted directed graphs. *ACM Transactions on Algorithms*, 8(4):33, 2012.

36 M. Thorup and U. Zwick. Spanners and emulators with sublinear distance errors. In *SODA*, pages 802–809, 2006.

37 V. Vassilevska Williams. Faster replacement paths. In *SODA*, pages 1337–1346. SIAM, 2011.

38 O. Weimann and R. Yuster. Replacement paths and distance sensitivity oracles via fast matrix multiplication. *ACM Transactions on Algorithms*, 9(2):14, 2013.

39 D. P. Woodruff. Additive spanners in nearly quadratic time. In *ICALP*, pages 463–474, 2010.