

Vertex Fault Tolerant Additive Spanners

Merav Parter *

June 9, 2015

Abstract

A *fault-tolerant* structure for a network is required to continue functioning following the failure of some of the network's edges or vertices. In this paper, we address the problem of designing a *fault-tolerant* additive spanner, namely, a subgraph H of the network G such that subsequent to the failure of a single vertex, the surviving part of H still contains an *additive* spanner for (the surviving part of) G , satisfying $\text{dist}(s, t, H \setminus \{v\}) \leq \text{dist}(s, t, G \setminus \{v\}) + \beta$ for every $s, t, v \in V$. Recently, the problem of constructing fault-tolerant additive spanners resilient to the failure of up to f edges has been considered [9]. The problem of handling *vertex* failures was left open therein. In this paper we develop new techniques for constructing additive FT-spanners overcoming the failure of a single vertex in the graph. Our first result is an FT-spanner with additive stretch 2 and $O(n^{5/3})$ edges. Our second result is an FT-spanner with additive stretch 6 and $O(n^{3/2})$ edges. The construction algorithm consists of two main components: (a) constructing an FT-clustering graph and (b) applying a modified path-buying procedure suitably adapted to failure prone settings. Finally, we also describe two constructions for *fault-tolerant multi-source additive spanners*, aiming to guarantee a bounded additive stretch following a vertex failure, for every pair of vertices in $S \times V$ for a given subset of sources $S \subseteq V$. The additive stretch bounds of our constructions are 4 and 8 (using a different number of edges).

1 Introduction

An (α, β) -spanner H of an unweighted undirected graph G is a spanning subgraph satisfying for every pair of vertices $s, t \in V$ that $\text{dist}(s, t, H) \leq \alpha \cdot \text{dist}(s, t, G) + \beta$. When $\beta = 0$, the spanner is termed a *multiplicative* spanner and when $\alpha = 1$ the spanner is *additive*. Clearly, additive spanners provide a much stronger guarantee than multiplicative ones, especially for long distances. Constructions of additive spanners with *small* number of edges are currently known for $\beta = 2, 4, 6$ with $O(n^{3/2})$, $\tilde{O}(n^{7/5})$ and $O(n^{4/3})$ edges respectively [1, 2, 5, 12, 15, 16]. This paper considers a network G that may suffer a single *vertex* failure event, and looks for fault tolerant additive spanners that maintain their additive stretch guarantee under failures. Formally, a subgraph $H \subseteq G$ is a β -additive FT-spanner iff for every $(s, t) \in V \times V$ and for every failing vertex $v \in V$, $\text{dist}(s, t, H \setminus \{v\}) \leq \text{dist}(s, t, G \setminus \{v\}) + \beta$. As a motivation for such structures, consider a situation where it is required to lease a subnetwork of a given network, which will provide short routes from every source s and every target t with additive stretch 2. In a failure-free environment one can simply lease a 2-additive spanner H_0 of the graph with $\Theta(n^{3/2})$ edges. However, if one of the

*Department of Computer Science and Applied Mathematics, The Weizmann Institute of Science, Rehovot, Israel. E-mail: {merav.parter}@weizmann.ac.il. Recipient of the Google European Fellowship in distributed computing; research is supported in part by this Fellowship. Supported in part by the Israel Science Foundation (grant 894/09).

vertices in the graph fails, some $s - t$ routes in $H_0 \setminus \{v\}$ might be significantly longer than the corresponding route in the surviving graph $G \setminus \{v\}$. Moreover, s and t are not even guaranteed to be connected in $H_0 \setminus \{v\}$. One natural approach towards preparing for such eventuality is to lease a larger set of links, i.e., an additive FT-spanner.

The notion of fault-tolerant spanners for general graphs was initiated by Chechik et al. [11] for the case of multiplicative stretch. Specifically, [11] presented algorithms for constructing an f -vertex fault tolerant spanner with multiplicative stretch $(2k - 1)$ and $O(f^2 k^{f+1} \cdot n^{1+1/k} \log^{1-1/k} n)$ edges. Dinitz and Krauthgamer presented in [14], a randomized construction attaining an improved tradeoff for *vertex* fault-tolerant spanners, namely, f -vertex fault tolerant k -spanner with $\tilde{O}(f^2 \cdot n^{1+2/(k+1)})$ edges. Constructions of fault-tolerant spanners with *additive* stretch resilient to *edge* failures were recently given by Braunschvig et al. [9]. They establish 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 $(\alpha, 0)$ fault tolerant 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 $\beta(f) = 2f(2\beta + \alpha - 1) + \beta$. Note that since in general the degree of a vertex in the graph might be as large as $\Theta(n)$, using the construction of [9] by removing all edges incident to the failing vertex, might result with an additive stretch of $\Theta(n)$. Fixing the number of *edge* faults to $f = 1$, yields constructions with an additive stretch of 14 with $O(n^{3/2})$ edges and an additive stretch of 38 with $O(n^{4/3})$ edges. When fixing the number of edges in H to be $O(n^{3/4})$ edges and the number of edge faults to $f = 1$ yields $\alpha = 5$ and $\beta = 6$. Hence, in particular, there is no construction for additive stretch < 14 and $o(n^2)$ edges. In addition, note that these structures are resilient only to *edge* failures as the techniques of [9] cannot be utilized to protect even against a single vertex failure event. Indeed, the problem of handling *vertex* failures was left open therein.

In this paper, we make a first step in this direction and provide additive FT-structures resilient to the failure of a single *vertex* (and hence also edge) event. Our constructions provide additive stretch 2 and 6 and hence provide an improved alternative also for the case of a single edge failure event, compared to the constructions of [9].

The presented algorithms are based upon two important notions, namely, *replacement paths* and the *path-buying procedure*, which have been studied extensively in the literature. For a source s , a target vertex t and a failing vertex $v \in V$, a *replacement path* is the shortest $s - t$ path $P_{s,t,v}$ that does not go through v . The vast literature on *replacement paths* (cf. [7, 17, 23, 25]) focuses on *time-efficient* computation of these paths as well as their efficient maintenance in data structures (a.k.a *distance oracles*).

Fault-resilient structures that preserve exact distances for a given subset of sources $S \subseteq V$ have been studied in [20], which defines the notion of an FT-MBFS structure $H \subseteq G$ containing the collection of all replacement paths $P_{s,t,v}$ for every pair $(s, t) \in S \times V$ for a given subset of sources S and a failing vertex $v \in V$. Hence, FT-MBFS structures preserve the exact $s - t$ distances in $G \setminus \{v\}$ for every failing vertex v , for every source $s \in S$.

It is shown in [20] that for every graph G and a subset S of sources, there exists a (poly-time constructible) 1-edge (or vertex) FT-MBFS structure H with $O(\sqrt{|S|} \cdot n^{3/2})$ edges. This result is complemented by a matching lower bound showing that for sufficiently large n , there exist an n -vertex graph G and a source-set $S \subseteq V$, for which every FT-MBFS structure is of size $\Omega(\sqrt{|S|} \cdot n^{3/2})$. Hence *exact* FT-MBFS 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-MBFS structure.

The problem of constructing *multiplicative approximation replacement paths* $\tilde{P}_{s,t,v}$ (i.e., such that $|\tilde{P}_{s,t,v}| \leq \alpha \cdot |P_{s,t,v}|$) has been studied in [3, 10, 6]. In particular its *single source* variant has been studied in [4, 21, 8]. In this paper, we further explore this approach. For a given subset of sources S , we focus on constructions of subgraphs that contain an *approximate* BFS structure with additive stretch β for every source $s \in S$ that are resistant to a single vertex failure.

Indeed, the construction of additive sourcewise FT-spanners provides a key building block of additive FT-spanner constructions (in which bounded stretch is guaranteed for all pairs). We present two constructions of sourcewise spanners with different stretch-size tradeoffs. The first construction ensures an additive stretch 4 with $O(n \cdot |S| + (n/|S|)^3)$ edges and the second construction guarantees additive stretch 8 with $O(n \cdot |S| + (n/|S|)^2)$. As a direct consequence of these constructions, we get an additive FT-spanner with stretch 6 and $O(n^{3/2})$ edges and an additive sourcewise FT-spanner with additive stretch 8 and $O(n^{4/3})$ for at most $O(n^{1/3})$ sources.

Additive spanners for specified pairs or sources where the objective is to construct a subgraph $H \subseteq G$ that satisfies the bounded additive stretch requirement only for a *subset* of pairs, are given in [22, 18, 13, 19].

Contributions. This paper provides the first constructions for additive spanners resilient upon a single vertex failure. In addition, it provides the first additive FT-structures with stretch guarantee as low as 2 or 6 and with $o(n^2)$ edges.

Our constructions employ a modification of the *path-buying* strategy, which was originally devised in [5] to provide 6-additive spanners with $O(n^{4/3})$ edges. Recently, the path-buying strategy was employed in the context of pairwise spanners [13].

The adaptation of the path-buying strategy to the vertex failure setting has been initiated in [21] for the case of a *single-source* s and a single *edge* failure event. In this paper, we extend this technique in two senses: (1) dealing with many sources and (2) dealing with *vertex* failures. In particular, [21] achieves a construction of single source additive spanner with $O(n^{4/3})$ edges resilient to a single *edge* failure. In this paper, we extend this construction to provide a multiple source additive spanners resilient to a single vertex failure, for $O(n^{1/3})$ sources, additive stretch 8 and $O(n^{4/3})$ edges. In summary, we show the following.

Theorem 1.1 (2-additive FT-spanner) *For every n -vertex graph $G = (V, E)$, there exists a (polynomially constructible) subgraph $H \subseteq G$ of size $O(n^{5/3})$ such that $\text{dist}(s, t, H \setminus \{v\}) \leq \text{dist}(s, t, G \setminus \{v\}) + 2$ for every $s, t, v \in V$.*

Theorem 1.2 (4-additive sourcewise FT-spanner) *For every n -vertex graph $G = (V, E)$ and a subset of sources $S \subset V$, there exists a (polynomially constructible) subgraph $H \subseteq G$ of size $O(|S| \cdot n + (n/|S|)^3)$ such that $\text{dist}(s, t, H \setminus \{v\}) \leq \text{dist}(s, t, G \setminus \{v\}) + 4$ for every $s \in S$ and $t, v \in V$.*

Taking $|S| = O(\sqrt{n})$, Thm. 1.2 can be shown to yield the following.

Theorem 1.3 (6-additive FT-spanner) *For every n -vertex graph $G = (V, E)$, there exists a (polynomially constructible) subgraph $H \subseteq G$ of size $O(n^{3/2})$ such that $\text{dist}(s, t, H \setminus \{v\}) \leq \text{dist}(s, t, G \setminus \{v\}) + 6$ for every $s, t, v \in V$.*

Theorem 1.4 (8-additive sourcewise FT-spanner) *For every n -vertex graph $G = (V, E)$ and a subset of sources $S \subset V$ where $|S| = O(n^{1/3})$, there exists a (polynomially constructible) subgraph $H \subseteq G$ of size $O(n^{4/3})$ such that $\text{dist}(s, t, H \setminus \{v\}) \leq \text{dist}(s, t, G \setminus \{v\}) + 8$ for every $s \in S$ and $t, v \in V$.*

Techniques. Our main technical contribution is in adapting the path-buying technique to the fault-tolerance setting. The high-level idea of the basic path-buying procedure (without faults) is as follows. In an initial clustering phase, a suitable clustering of the vertices is computed, and an associated subset of edges is added to the spanner. Then comes a path-buying phase, where they consider an appropriate sequence of paths, and decide whether or not to add each path into the spanner. Each path P has a *cost*, given by the number of edges of p not already contained in the spanner, and a *value*, measuring P 's help in satisfying the considered set of constraints on pairwise distances. The considered path P is added to the spanner iff its value is sufficiently larger than its cost. In our adaptation to the FT-setting, an FT-clustering graph is computed first, providing vertices with a sufficiently *high* degree *two* clusters to which it belongs. Every cluster consists of a center vertex v connected via a star to a subset of its heavy neighbors. In our design not all replacement paths are candidates to be bought in the path-buying procedure. Let $\pi(s, t)$ be an $s - t$ shortest-path between a source s and a clustered vertex t . We divide the failing events on $\pi(s, t)$ into two classes depending on the position of the failing vertex on $\pi(s, t)$ with respect to the least common ancestor (LCA) $\ell(s, t)$ of t 's cluster members in the BFS tree rooted at s . Specifically, a vertex fault $\pi(s, t)$ that occurs on $\ell(s, t)$ is handled *directly* by adding the last edge of the corresponding replacement path to the spanner. Vertex failures that occur strictly below the LCA, use the shortest-path $\pi(s, x)$ between s and some member x in the cluster of t such that the failing vertex v does not appear on the $\pi(s, x)$ path. The approximate replacement path will follow $\pi(s, x)$ and then use the intercluster path between x and t . The main technicality is when concerning the complementary case when that failing event occurs strictly above $\ell(s, t)$. These events are further divided into two classes depending on the structure of their replacement path. Some of these replacement paths would again be handled directly by collecting their last edges into the structure and only the second type paths would be candidates to be bought by the path-buying procedure.

2 Preliminaries

Notation. Throughout, we assume that the shortest-paths in the given graph $G = (V, E)$ are unique. Specifically, the shortest path ties are broken in a consistent manner so that the subpath of a shortest-path is a shortest path itself. For a subgraph $G' \subseteq G$, let $\pi(u, v, G')$ denote the $u - v$ shortest-path in G' .

Given a graph $G = (V, E)$, a vertex pair s, t and an edge weight function $W : E(G) \rightarrow \mathbb{R}^+$, let $SP(s, t, G, W)$ be the set of $s - t$ shortest-paths in G according to the edge weights of W . Throughout, we make use of (an arbitrarily specified) weight assignment W that guarantees the uniqueness of the shortest paths¹. Hence, $SP(s, t, G', W)$ contains a single path for every $s, t \in V$ and for every subgraph $G' \subseteq G$, we override notation and let $SP(s, t, G, W)$ be the unique $s - t$ path in G according to W . When the shortest-paths are computed in G , let $\pi(s, t) = SP(s, t, G, W)$. To avoid cumbersome notation, we may omit W and simply refer to $\pi(s, t) = SP(s, t, G, W)$. For a subgraph $G' \subseteq G$, let $V(G')$ (resp., $E(G')$) denote the vertex set (resp. edge set) in G' .

For a given source node s , let $T_0(s) = \bigcup_{t \in V} \pi(s, t)$ be a shortest paths (or BFS) tree rooted at s . For a set $S \subseteq V$ of source nodes, let $T_0(S) = \bigcup_{s \in S} T_0(s)$ be a union of the single source BFS trees. For a vertex $t \in V$ and a subset of vertices $V' \subseteq V$, let $T(t, V') = \bigcup_{u \in V'} \pi(u, t)$ be the union of all $\{t\} \times V'$ shortest-paths (by the uniqueness of W , $T(t, V')$ is a subtree of $T_0(t)$). Let $\Gamma(v, G)$

¹The role of the weights W is to perturb the edge weights by letting $W(e) = 1 + \epsilon$ for a random infinitesimal $\epsilon > 0$.

be the set of v 's neighbors in G . 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 node v in G . For a given graph $G = (V, E)$ and an integer $\Delta \leq n$, a vertex v is Δ -heavy if $\deg(v, G) \geq \Delta$, otherwise it is Δ -light. When Δ is clear from the context, we may omit it and simply refer to v as *heavy* or *light*. For a graph $G = (V, E)$ and a positive integer $\Delta \leq n$, let $V_\Delta = \{v \mid \deg(v, G) \geq \Delta\}$ be the set of Δ -heavy vertices in G . (Throughout, we sometimes simplify notation by omitting parameters which are clear from the context.) 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 shortest-path distance in edges between u and v in G' . For a path $P = [v_1, \dots, v_k]$, let $\text{LastE}(P)$ be the last edge of P , let $|P|$ denote its length and let $P[v_i, v_j]$ be the subpath of P from v_i to v_j . For paths P_1 and P_2 , denote by $P_1 \circ P_2$ the path obtained by concatenating P_2 to P_1 . For “visual” clarity, the edges of these paths are considered throughout, to be directed away from the source node s . Given an $s - t$ path P and an edge $e = (x, y) \in P$, let $\text{dist}(s, e, P)$ be the distance (in edges) between s and y on P . In addition, for an edge $e = (x, y) \in T_0(s)$, define $\text{dist}(s, e) = i$ if $\text{dist}(s, x, G) = i - 1$ and $\text{dist}(s, y, G) = i$. A vertex w is a *divergence point* of the s - v paths P_1 and P_2 if $w \in P_1 \cap P_2$ but the next vertex u after w (i.e., such that u is closer to v) in the path P_1 is not in P_2 .

Basic Tools. We consider the following graph structures.

Definition 2.1 ((α, β, S) **FT-spanner**) *A subgraph $H \subseteq G$ is an (α, β, S) FT-spanner structure with respect to S if for every $(s, t) \in S \times V$ and every $v \in V$, $\text{dist}(s, t, H \setminus \{v\}) \leq \alpha \cdot \text{dist}(s, t, G \setminus \{v\}) + \beta$.*

Definition 2.2 ((α, β) **FT-spanners**) *A subgraph $H \subseteq G$ is an (α, β) FT-spanner if it is an (α, β, V) FT-spanner for G with respect to V .*

Throughout, we restrict attention to the case of a single vertex fault. When $\alpha = 1$, H is termed (β, S) -additive FT-spanner. In addition, in case where $S = V$, H is an β -additive FT-spanner.

FT-Clustering Graph G_Δ . For a subset $Z \subseteq V$, a heavy vertex $v \in V_\Delta$ is said to be *clustered* by Z if $|\Gamma(v, G) \cap Z| \geq 2$. When Z is clear from the context, we simply say that v is clustered. A subset $Z \subseteq V$ is an *FT-center set* for a subset V' of Δ -heavy vertices, if every vertex $v \in V'$ is clustered by Z . The FT-center set Z is constructed in the following manner. Initially, set $Z = \emptyset$. Let V' be the subset of unclustered heavy vertices by Z . Hence, initially $V' = V_\Delta$. As long as there exists a vertex $u \in V \setminus Z$ that has at least Δ neighbors in V' , add u to Z and remove the clustered vertices of $\Gamma(u, G) \cap V'$ from V' . Let $V_C = V_\Delta \setminus V'$ be the set of clustered heavy vertices by Z . For every clustered vertex $v \in V_C$, let $Z(v) = \{z_1(v), z_2(v)\}$ be two arbitrary neighbors of v in Z .

We then add to the graph G_Δ , the set of all edges incident to the unclustered vertices in $V \setminus V_C$, and in-addition connect each clustered vertex $v \in V_C$ to $Z(v)$. Formally,

$$G_\Delta = \bigcup_{v \in V_C} \{(v, z_1(v)), (v, z_2(v))\} \cup \bigcup_{v \in V \setminus V_C} E(v, G).$$

Note that every edge in $G \setminus G_\Delta$ is incident to a clustered vertex. For every center vertex $z \in Z$, let C_z be the cluster consisting of z and all the clustered vertices it represents, i.e., $C_z = \{z\} \cup \{v \in V_C \mid z \in Z(v)\}$. Note that every center z is connected via a star to each of the vertices in its cluster C_z , hence the diameter of each cluster C_z in G_Δ is 2.

For a failing vertex v and a clustered vertex t , let $z_v(t) \in Z(t) \setminus \{v\}$ be a cluster center of t in $G \setminus \{v\}$. In particular, if $z_1(t) \neq v$, then $z_v(t) = z_1(t)$, else $z_v(t) = z_2(t)$. Let $C_v(t)$ be the cluster centered at $z_v(t)$. Note that since every clustered vertex has two cluster centers $z_1(t)$ and $z_2(t)$, we have the guarantee that at least one of them survives the single vertex fault event. The next observation summarizes some important properties of the clustering graph.

Observation 2.3 (1) *Every missing edge $e \in G \setminus G_\Delta$ is incident to a clustered vertex v .*

(2) *The diameter of every cluster C_z is 2.*

(3) *The FT-center set $Z \subseteq V$ is of size $|Z| = O(n/\Delta)$.*

(4) *$|E(G_\Delta)| = O(\Delta \cdot n)$.*

Proof: Claims (1-2) are immediate. Consider claim (3). Since every center in Z has at least Δ neighbors in V_Δ and since every vertex has two neighbors in Z before removing it from V' it holds that $|Z| \cdot \Delta \leq 2n$ hence $|Z| \leq 2n/\Delta$. Consider claim (4). The edges incident to light vertices is bounded by $\Delta \cdot n$. Hence, it remains to bound the edges incident to the subset of unclustered heavy vertices $V' = V_\Delta \setminus V_C$. Note that every vertex $u \in V \setminus Z$ has at most $\Delta - 1$ neighbors in V' . In addition, every vertex v in V' has at most one neighbor in Z , hence the total number of edges in $E(V', G)$ is bounded by $\Delta \cdot n$. The observation follows. ■

Replacement Paths. For a source s , a target vertex t and a vertex $v \in G$, a *replacement path* is the shortest $s - t$ path $P_{s,t,v} \in SP(s, t, G \setminus \{v\})$ that does not go through v .

Observation 2.4 *Every path $P_{s,t,v}$ contains at most $3n/\Delta$ Δ -heavy vertices.*

Proof: Note that

$$3n \geq 3 \cdot \left| \bigcup_{x \in P_{s,t,v} \cap V_\Delta} \Gamma(x, G \setminus \{v\}) \right| \geq \sum_{x \in P_{s,t,v} \cap V_\Delta} \deg(x, G \setminus \{v\}) \geq |P_{s,t,v} \cap V_\Delta| \cdot \Delta,$$

where the second inequality follows by the fact the every vertex $u \in V \setminus \{v\}$ has at most 3 neighbors on $P_{s,t,v}$. The observation follows. ■

New-ending replacement paths. A replacement path $P_{s,t,v}$ is called *new-ending* if its last edge is different from the last edge of the shortest path $\pi(s, t)$. Put another way, a new-ending replacement path $P_{s,t,v}$ has the property that once it diverges from the shortest-path $\pi(s, t)$, it joins $\pi(s, t)$ again only at the final vertex t . It is shown in [20] that for a given graph G and a set S of source vertices, a structure $H \subseteq G$ containing a BFS tree rooted at each $s \in S$ plus the last edge of each new-ending replacement path $P_{s,t,v}$ for every $(s, t) \in S \times V$ and every $v \in V$, is an **FT-MBFS** structure with respect to S . Our algorithms exploit the structure of new-ending replacement paths to construct (β, S) -additive FT-spanners. Essentially, a key ingredient in our analysis concerns with collecting the last edges from a subset of new-ending replacement paths as well as bounding the number of new-ending paths $P_{s,t,v}$ whose detour segments intersect with $\pi(s', t) \setminus \{t\}$ for some other source $s' \in S$.

The basic building block. Our constructions of β -additive FT-spanners, for $\beta \geq 2$, consist of the following two building blocks: (1) an FT-clustering graph G_Δ for some parameter Δ , and (2) an $(\beta - 2)$ -additive FT-spanner $H_{\beta-2}(Z)$ where Z is an FT-center set (i.e., cluster centers) for the vertices.

Lemma 2.5 *Let $\beta \geq 2$ and $H = G_\Delta \cup H_{\beta-2}(Z)$ where Z is an FT-center set for the clustered heavy vertices V_C . Then H is an β additive FT-spanner.*

Proof: Consider vertices $u_1, u_2, u_3 \in V$. Let $P \in SP(u_1, u_2, G \setminus \{u_3\})$ be the $u_1 - u_2$ replacement path in $G \setminus \{u_3\}$ and let (x, y) be the last missing edge on P that is not in H (i.e., closest to u_2). Since $G_\Delta \subseteq H$, by Obs. 2.3(1), y is a clustered vertex. Let $z = z_{u_3}(y)$ be the cluster center of y in $G \setminus \{u_3\}$, and consider the following $u_1 - u_2$ path $P_3 = P_1 \circ P_2$ where $P_1 \in SP(u_1, z, H \setminus \{u_3\})$ and $P_2 = (z, y) \circ P[y, u_2]$. Clearly, $P_3 \subseteq H \setminus \{u_3\}$, so it remains to bound its length. Since $H_{\beta-2}(Z) \subseteq H$, it holds that $|P_1| \leq \text{dist}(u_1, z, G \setminus \{u_3\}) + \beta - 2$. Hence,

$$\begin{aligned} \text{dist}(u_1, u_2, H \setminus \{u_3\}) &\leq |P_3| = |P_1| + |P_2| \\ &\leq \text{dist}(u_1, z, G \setminus \{u_3\}) + \beta - 2 + \text{dist}(y, u_2, G \setminus \{u_3\}) + 1 \\ &\leq \text{dist}(u_1, y, G \setminus \{u_3\}) + \text{dist}(y, u_2, G \setminus \{u_3\}) + \beta \\ &\leq |P| + \beta = \text{dist}(u_1, u_2, G \setminus \{u_3\}) + \beta, \end{aligned}$$

where the second inequality follows by the triangle inequality using the fact that the edge (z, y) exists in $H \setminus \{u_3\}$. The lemma follows. \blacksquare

3 Additive Stretch 2

We begin by considering the case of additive stretch 2. We make use of the construction of FT-MBFS structures presented in [20].

Fact 3.1 ([20]) *There exists a polynomial time algorithm that for every n -vertex graph $G = (V, E)$ and a source set $S \subseteq V$ constructs an FT-MBFS structure $H_0(S)$ from each source $s_i \in S$, tolerant to one edge or vertex failure, with a total number of $O(\sqrt{|S|} \cdot n^{3/2})$ edges.*

Set $\Delta = \lceil n^{2/3} \rceil$ and let Z be an FT-center set (see Obs. 2.3(3)). Let $H_0(Z)$ be an FT-MBFS structure with respect to the source set Z as given by Fact 3.1. Then, let $H = G_\Delta \cup H_0(Z)$. Thm. 1.1 follows by Lemma 2.5, Obs. 2.3 and Fact 3.1.

4 Sourcewise additive FT-spanners

In this section, we present two constructions of $(4, S)$ and $(8, S)$ additive FT-spanners with respect to a given source set $S \subseteq V$. The single source and single *edge* failure case (where $|S| = 1$) is considered in [21], which provides a construction of a single source FT-spanner resilient against single edge failure with $O(n^{4/3})$ edges. The current construction deals with single vertex failures and increases the stretch to 8 while providing a bounded stretch for $O(n^{1/3})$ sources with the same order of edges, $O(n^{4/3})$.

4.1 Sourcewise spanner with additive stretch 4

Theorem 4.1 *There exists a subgraph $H_4(S) \subseteq G$ with $O(|S| \cdot n + (n/|S|)^3)$ edges satisfying $\text{dist}(s, t, H_4(S) \setminus \{v\}) \leq \text{dist}(s, t, G \setminus \{v\}) + 4$ for every $(s, t) \in S \times V$ and $v \in V$.*

The following notation is useful in our context. Let $\mathcal{C} = \{C_z \mid z \in Z\}$ be the collection of clusters corresponding to the FT-centers Z . For a source $s \in S$ and a cluster $C_z \in \mathcal{C}$ rooted at FT-center $z \in Z$, let $\text{LCA}(s, C_z)$ be the least common ancestor (LCA) of the cluster vertices of C_z in the BFS tree $T_0(s)$ rooted at s . Let $\pi(s, C_z)$ be the path connecting s and $\text{LCA}(s, C_z)$ in $T_0(s)$.

4.1.1 Algorithm Cons4SWSpanner for constructing $H_4(S)$ spanner

Step (0): Replacement-path definition. For every $(s, t) \in S \times V$ and every $v \in V$, let $P_{s,t,v} = SP(s, t, G \setminus \{v\}, W)$.

Step (1): Clustering. Set $\Delta = |S|$ and let $Z \subseteq V$ be an FT-center set of size $O(n/\Delta)$ (see Obs. 2.3(3)). Let V_C be the subset of heavy clustered vertices. $\mathcal{C} = \{C_z \mid z \in Z\}$ be the collection of $|Z|$ clusters. For a clustered vertex t , let $C_1(t), C_2(t)$ be its two clusters in \mathcal{C} corresponding to the centers $z_1(t)$ and $z_2(t)$ respectively. The initial spanner is then given by $\tilde{H}_0 = T_0(S) \cup G_\Delta$.

Step (2): Shortest-path segmentation. For every $(s, t) \in S \times V_C$, the algorithm uses the first cluster of t , $C_1(t)$, to segment the path $\pi(s, t)$. Define

$$\pi^{\text{far}}(s, t) = \pi(s, \ell(s, t)) \setminus \{\ell(s, t)\} \quad \text{and} \quad \pi^{\text{near}}(s, t) = \pi(\ell(s, t), t) \setminus \{\ell(s, t)\},$$

where $\ell(s, t) = \text{LCA}(s, C_1(t))$ is the LCA of the cluster $C_1(t)$ in the tree $T_0(s)$. Hence, $\pi(s, t) = \pi^{\text{far}}(s, t) \circ \ell(s, t) \circ \pi^{\text{near}}(s, t)$. The algorithm handles separately vertex faults in the near and far segments. Let $V^{\text{near}}(s, t) = V(\pi^{\text{near}}(s, t))$ and $V^{\text{far}}(s, t) = V(\pi^{\text{far}}(s, t))$.

Step (3): Handling faults in the cluster center and the LCA. Let

$$E^{\text{local}}(t) = \{\text{LastE}(P_{s,t,v}) \mid s \in S, v \in \{z_1(t), \text{LCA}(s, C_1(t))\}\} \quad \text{and} \quad E^{\text{local}} = \bigcup_{t \in V_C} E^{\text{local}}(t),$$

be the last edges of replacement-paths protecting against the failure of the primary cluster center $z_1(t)$ and the least common ancestor $\text{LCA}(s, C_1(t))$.

Step (4): Handling far vertex faults $V^{\text{far}}(s, t)$. Consider the $s - t$ new-ending replacement-paths $P_{s,t,v}$ of a clustered vertex t . Let $b_{s,t,v}$ be the unique divergence point of $P_{s,t,v}$ from $\pi(s, t)$ (the uniqueness of the divergence point is guaranteed by the uniqueness of the shortest-paths, as shown in the analysis). Let $D_{s,t,v} = P_{s,t,v}[b_{s,t,v}, t]$ denote the detour segment and let $D_{s,t,v}^- = D_{s,t,v} \setminus \{b_{s,t,v}\}$ denote the detour segment excluding the divergence point. For every clustered vertex t , let $\mathcal{P}^{\text{far}}(t)$ be the collection of new-ending² $s - t$ paths protecting against vertex faults in the far segments, i.e., $\mathcal{P}^{\text{far}}(t) = \{P_{s,t,v} \mid s \in S, \text{LastE}(P_{s,t,v}) \notin T_0(S) \text{ and } v \in V^{\text{far}}(s, t)\}$.

²actually, it would be sufficient to consider only the paths whose last edge is missing in the current spanner.

The algorithm divides this set into two subsets $\mathcal{P}_{\text{dep}}^{\text{far}}(t)$ and $\mathcal{P}_{\text{indep}}^{\text{far}}(t)$ depending on the structure of the partial detour segment $D_{s,t,v}^-$. A new-ending path $P_{s,t,v}$ is *dependent* if $D_{s,t,v}^-$ intersects $\pi(s',t) \setminus \{t\}$ for some $s' \in S$, i.e., for a dependent path $P_{s,t,v}$, it holds that

$$V(D_{s,t,v}^-) \cap V(T(t,S)) \neq \{t\}. \quad (1)$$

Otherwise, it is *independent*. Let

$$\mathcal{P}_{\text{dep}}^{\text{far}}(t) = \{P_{s,t,v} \in \mathcal{P}^{\text{far}}(t) \mid V(D_{s,t,v}^-) \cap V(T(t,S)) \neq \{t\}\}$$

be the set of all $S \times \{t\}$ dependent paths and let $\mathcal{P}_{\text{indep}}^{\text{far}}(t) = \mathcal{P}^{\text{far}} \setminus \mathcal{P}_{\text{dep}}^{\text{far}}(t)$ be the set of independent paths.

Step (4.1): Handling *dependent* new-ending paths. The algorithm simply takes the last edges $E_{\text{dep}}^{\text{far}}(t)$ of all dependent replacement paths where $E_{\text{dep}}^{\text{far}}(t) = \{\text{LastE}(P) \mid P \in \mathcal{P}_{\text{dep}}^{\text{far}}(t)\}$.

Let $E_{\text{dep}}^{\text{far}} = \bigcup_{t \in V_C} E_{\text{dep}}^{\text{far}}(t)$.

In the analysis we show that dependant paths have a special structure which imposes a constraints on the cardinality of $E_{\text{dep}}^{\text{far}}(t)$.

Step (4.2): Handling *independent* new-ending paths. The algorithm employs a modified *path-buying* procedure on the collection $\mathcal{P}_{\text{indep}}^{\text{far}} = \bigcup_{t \in V_C} \mathcal{P}_{\text{indep}}^{\text{far}}(t)$ of new-ending independent paths. The paths of $\mathcal{P}_{\text{indep}}^{\text{far}}$ are considered in some arbitrary order. A path $P \in \mathcal{P}_{\text{indep}}^{\text{far}}$ is bought, if it improves the pairwise cluster distances in some sense. Starting with

$$G_0 = T_0(S) \cup G_\Delta \cup E^{\text{local}} \cup E_{\text{dep}}^{\text{far}}, \quad (2)$$

at step $\tau \geq 0$, the algorithm is given $G_\tau \subseteq G$ and considers the path $P_\tau = P_{s,t,v}$. Let $e = (x,y)$ be the first edge on P_τ which is not in $E(G_\tau)$ (where x is closer to s). Note that since $G_\Delta \subseteq G_0$, both x and t are clustered. Recall that for a clustered vertex u and a failing vertex v , $C_v(u)$ is the cluster of u centered at $z_v(u) \in Z(u) \setminus \{v\}$. For every cluster C , let $V_f(C)$ be the collection of vertices appearing on the paths $\pi(s,C) = \pi(s, \text{LCA}(s,C))$ for every $s \in S$ excluding the vertices of the cluster. That is,

$$V_f(C) = \bigcup_{s \in S} V(\pi(s,C)) \setminus C. \quad (3)$$

Let $C_{1,\tau} = C_v(x)$ be the cluster of x in $G_\Delta \setminus \{v\}$ and $C_{2,\tau} = C_v(t)$ be the cluster of t in $G_\Delta \setminus \{v\}$. The path P_τ is added to G_τ resulting in $G_{\tau+1} = G_\tau \cup P_\tau$, only if

$$\text{dist}(x,t,P_\tau) < \text{dist}(C_{1,\tau}, C_{2,\tau}, G_\tau \setminus V_f(C_v(t))). \quad (4)$$

Let $\tau' = |\mathcal{P}_{\text{indep}}^{\text{far}}|$ be the total number of independent paths considered to be bought by the algorithm. Then, the algorithm outputs $H_4(S) = G_{\tau'}$. This completes the description of the algorithm.

Analysis. Throughout, we restricted attention to $s - t$ replacement paths of clustered vertices $t \in V_C$. Let $b_{s,t,v}$ be the first divergence point of $P_{s,t,v}$ and $\pi(s, t)$.

Lemma 4.2 *For every vertex $u \in P_{s,t,v}$ such that $\text{LastE}(P_{s,t,v}[s, u]) \notin T_0(S)$, it holds that: (a) $v \in V(\pi(s, u))$. (b) $V(P_{s,t,v}[b_{s,t,v}, u]) \cap V(\pi(s, u)) = \{b_{s,t,v}, u\}$.*

Proof: Begin with (a). Assume towards contradiction otherwise. By the uniqueness of the weight assignment W , we get that $P_{s,t,v}[s, u] = SP(s, u, G \setminus \{v\}, W) = \pi(s, u)$. We therefore get a contradiction to the fact that $\text{LastE}(P_{s,t,v})$ not in $T_0(S)$. We next prove (b) and show that the divergence point $b_{s,t,v}$ is unique. By the definition of $b_{s,t,v}$, it occurs on $\pi(s, t)$ above the failing vertex v . Since by Lemma 4.2, $v \in V(\pi(s, u))$, it also holds that $b_{s,t,v} \in V(\pi(s, u))$. Assume towards contradiction that there exists an additional point

$$w \in (V(P_{s,t,v}[b_{s,t,v}, u]) \cap V(\pi(s, u))) \setminus \{b_{s,t,v}, u\}.$$

There are two cases to consider (b1) $v \in V(\pi(b_{s,t,v}, w))$, in such a case, $v \notin V(\pi(w, u))$ and hence $\pi(w, u) = SP(w, u, G \setminus \{v\}, W) = P_{s,t,v}[w, u]$, contradicting the fact that $\text{LastE}(P_{s,t,v}[s, u]) \notin T_0(S)$. (b2) $v \in V(\pi(w, u))$. In such a case, $v \notin V(\pi(b_{s,t,v}, w))$ and hence $\pi(b_{s,t,v}, w) = SP(b_{s,t,v}, w, G \setminus \{v\}, W) = P_{s,t,v}[b_{s,t,v}, w]$, contradicting the fact that $b_{s,t,v}$ is a divergence point from $\pi(s, t)$. The lemma holds. ■

The next lemma shows that a new-ending $P_{s,t,v}$ path whose last edge is not in G_0 (see Eq. (2)), protecting against faults in the near segment, has a good approximate replacement path $\tilde{P}_{s,t,v}$ in G_0 .

Lemma 4.3 *If $\text{LastE}(P_{s,t,v}) \notin G_0$ and $v \in V^{\text{near}}(s, t)$, then $\text{dist}(s, t, G_0 \setminus \{v\}) \leq \text{dist}(s, t, G \setminus \{v\}) + 4$.*

Proof: Since $v \in V^{\text{near}}(s, t)$, i.e., the failing vertex occurs strictly below $\text{LCA}(s, C_1(t))$ on $\pi(s, t)$, there exists a vertex $w \in C_1(t)$ such that $v \notin V(\pi(s, w))$ (hence in particular $w \neq v$). See Fig. 1. Since $\text{LastE}(P_{s,t,v}) \notin E^{\text{local}}$, it holds that $v \neq z_1(t)$. Consider the following $s - t$ path $P = \pi(s, w) \circ [w, z_1(t), t]$. Clearly, $P \subseteq (T_0(S) \cup G_\Delta) \setminus \{v\}$. By the triangle inequality, as the diameter of the cluster $C_1(t)$ is 2, it holds that

$$\text{dist}(s, t, (T_0(S) \cup G_\Delta) \setminus \{v\}) \leq \text{dist}(s, w, G) + 2 \leq \text{dist}(s, t, G) + 4 \leq \text{dist}(s, t, G \setminus \{v\}) + 4.$$

The lemma follows. ■

The following auxiliary lemma is useful. It considers a new-ending path $P_{s,t,v}$ where $D_{s,t,v}^- = D_{s,t,v} \setminus \{b_{s,t,v}\}$ and (x, y) is the first missing edge in $P_{s,t,v} \setminus G_0$ (where x is closer to s).

Lemma 4.4 *For every vertex $u \in P_{s,t,v}$ such that $\text{LastE}(P_{s,t,v}[s, u]) \notin G_0$, it holds that: (a) $C_v(u) = C_1(u)$. (b) $P_{s,t,v}[x, t] \subseteq D_{s,t,v}^-$.*

Proof: Begin with (a). By the definition of the weight assignment W , it holds that $P_{s,t,v}[s, u] = P_{s,u,v} = SP(s, u, G \setminus \{v\}, W)$. Since $\text{LastE}(P_{s,t,v}[s, u]) \notin E^{\text{local}}$, it holds that $v \neq z_1(u)$, concluding that $z_v(u) = z_1(u)$ and hence $C_v(u) = C_1(u)$. Claim (a) follows. Consider claim (b). Let $b = b_{s,t,v}$. We show that $x \neq b$, which implies the claim. By Lemma 4.2, $v \in V(\pi(s, y))$. Since b appears above v on $\pi(s, t)$, b is a common vertex in both $\pi(s, t)$ and $\pi(s, y)$. Hence, the $s - y$ shortest path has the following form: $\pi(s, y) = \pi(s, b) \circ \pi(b, v) \circ \pi(v, y)$. Since $b \neq v \neq y$, $\text{dist}(b, y, G) \geq 2$, and hence also $\text{dist}(b, y, G \setminus \{v\}) \geq 2$, concluding that $b \neq x$. The lemma follows. ■




Figure 1: Handling near vertex faults. Schematic illustration of an approximate replacement path in $(T_0 \cup G_\Delta) \setminus \{v\}$. Shown is an $\pi(s, t)$ whose failing vertex v occurs strictly below the least common ancestor $LCA(s, C_1(t))$. The alternative replacement path exploits the surviving $\pi(s, w) \subseteq T_0(S)$ path for $w \in C_1(t)$ and the intracluster path connecting w and v through $z = z_v(t)$.

Corollary 4.5 *Let $t \in V_C$. For every $P_{s,t,v} \in \mathcal{P}_{\text{indep}}^{\text{far}}(t)$, $P_{s,t,v}[x, t] \cap V_f(C_v(t)) = \emptyset$ where x is the first vertex of $D_{s,t,v}^-$.*

Proof: Since $P_{s,t,v}$ is independent, by Eq. 1, $D_{s,t,v}^- \cap T(t, S) = \{t\}$. Since $t \in C_v(t)$, by Eq. (3), $t \notin V_f(C_v(t))$ and hence $D_{s,t,v}^- \cap V_f(C_v(t)) = \emptyset$. The corollary follows Lemma 4.4(b). ■

Correctness analysis of $H_4(S)$. We now show that $H_4(S)$ is a $(4, S)$ FT-spanner.

Lemma 4.6 *$H_4(S)$ is a $(4, S)$ FT-spanner.*

Proof: Fix a source $s \in S$ and let $H = H_4(S)$. We then show that for every pair $(t, v) \in V \times V$ it holds that

$$\text{dist}(s, t, H \setminus \{v\}) \leq \text{dist}(s, t, G \setminus \{v\}) + 4. \quad (5)$$

To show this, define $e_{t,v}$ to be the last edge of $P_{s,t,v} = SP(s, t, G \setminus \{v\}, W)$ that is missing in H and let $d(t, v) = \text{dist}(s, e_{t,v}, P_{s,t,v})$ be the distance of $e_{t,v}$ from s on $P_{s,t,v}$, for every pair $(t, v) \in V \times V$. The proof is shown by induction on $d(t, v)$. For $d(t, v) = 0$, $P_{s,t,v} \subseteq H$ and Eq. (5) holds vacuously. Assume that the claim holds for all (t, v) pairs with $d(t, v) \leq d - 1$ and consider some pair (t_0, v_0) pair with $d(t_0, v_0) = d$. Let $e_{t_0, v_0} = (u, t')$ be the last missing edge on P_{s, t_0, v_0} . We distinguish between two cases.

Case (1) $d(t', v_0) \leq d - 1$. Let $P'' \in SP(s, t', H \setminus \{v_0\})$. By induction assumption, it holds that

$$\begin{aligned} |P''| &\leq \text{dist}(s, t', G \setminus \{v_0\}) + 4 \\ &= |P_{s, t_0, v_0}[s, t']| + 4. \end{aligned} \quad (6)$$

We now consider the following $s - t_0$ replacement path $Q = P'' \circ P_{s, t_0, v_0}[t', t_0]$. By definition of (t', v_0) (last missing edge on P_{s, t', v_0}), $Q \subseteq H \setminus \{v_0\}$. In addition,

$$\begin{aligned} |Q| &= |P''| + |P_{t_0, v_0}[t', t_0]| \leq |P_{s, t_0, v_0}[s, t']| + 4 + |P_{s, t_0, v_0}[t', t_0]| \\ &= |P_{s, t_0, v_0}| + 4 = \text{dist}(s, t_0, G \setminus \{v_0\}) + 4, \end{aligned}$$

where the inequality follows by Eq. (6), and Eq. (5) holds as required.

Case (2) $d(t', v_0) \geq d$. In this case, since $|P_{s, t', v_0}| = |P_{s, t_0, v_0}|$, $d(t', v_0) = d$. Hence, the last edge of P_{s, t', v_0} is not in H . Since $\pi(s, t') \subseteq H$, we have that the failing vertex v_0 occurs on the shortest-path $\pi(s, t')$. Since the last edge of $P_{s, t', v_0} = P_{s, t_0, v_0}[s, t']$ is missing in H , by the fact that the clustering graph G_Δ is in H , by Obs. 2.3(1), it holds that t' is a clustered vertex. By step (2), since $E^{\text{local}} \subseteq H$, it holds that $v_0 \notin \{z_1(t'), \text{LCA}(s, C_1(t'))\}$. Combining with Lemma 4.3, it holds that $v_0 \notin V^{\text{near}}(s, t')$. Hence, $v_0 \in V^{\text{far}}(s, t')$. We next claim that P_{s, t', v_0} is an *independent* path. This holds since the last edges of t_0 's dependant replacement paths $E_{\text{dep}}^{\text{far}}(t_0)$ were added to H in step (4.1). Thus P_{s, t', v_0} is an *independent* path and hence it was considered to be bought in the path-buying procedure of Step (4.2). If the algorithm bought P_{s, t', v_0} , Eq. (5) holds. So, it remains to consider the case where the algorithm did not buy this path. Let τ be the iteration at which $P_\tau = P_{s, t', v_0}$ was considered to be purchased in the path-buying procedure. Let G_τ be the current spanner in iteration τ . Let x be the vertex incident to the first missing edge on $P_\tau \setminus E(G_\tau)$.

Let $C_0 = C_{v_0}(t')$ be the cluster of t' in $G_\Delta \setminus \{v_0\}$. Since $\text{LastE}(P_\tau) \notin E^{\text{local}}$, by Lemma 4.4(a) $C_0 = C_1(t')$ and $z_{v_0}(t') = z_1(t')$. By definition,

$$v_0 \in V^{\text{far}}(s, t') = \pi(s, C_0) \setminus \{\text{LCA}(s, C_0)\}.$$

Hence, the failing vertex is *not* in the cluster C_0 , i.e., $v_0 \notin C_0$ and by Eq. (3),

$$v_0 \in V_f(C_0). \quad (7)$$

Since P_τ was not bought by the algorithm, by Eq. (4), we have that


$$\text{dist}(C_{v_0}(x), C_0, G_\tau \setminus V_f(C_0)) \leq \text{dist}(x, t', P_\tau). \quad (8)$$

Let $w_1 \in C_{v_0}(x)$ and $w_2 \in C_0$ be an arbitrary closest pair in $G_\tau \setminus V_f(C_0)$ from the clusters $C_{v_0}(x)$ and C_0 respectively satisfying that $\text{dist}(w_1, w_2, G_\tau \setminus V_f(C_0)) = \text{dist}(C_{v_0}(x), C_0, G_\tau \setminus V_f(C_0))$.

Let z_1 (resp., z_2) be the cluster center of $C_{v_0}(x)$ (resp., C_0). Consider the following $s - t'$ replacement path in $H \setminus \{v_0\}$, $P_5 = P_1 \circ P_2 \circ P_3 \circ P_4$ where $P_1 = P_\tau[s, x]$, $P_2 = [x, z_1, w_1]$ and $P_3 \in SP(w_1, w_2, G_\tau \setminus V_f(C_0))$ and $P_4 = [w_2, z_2, t']$. For an illustration see Fig. 2. We first claim that $P_5 \subseteq H \setminus \{v_0\}$. Since x is incident to the first missing edge on P_τ , P_1 is in $H \setminus \{v_0\}$. By Eq. (7), $v_0 \in V_f(C_0)$ and since $w_1, w_2 \subseteq G_\tau \setminus V_f(C_0)$ it also holds that $w_1, w_2 \neq v_0$. Finally note that $G_\tau, G_\Delta \subseteq H$, hence $P_2, P_4 \subseteq H \setminus \{v_0\}$. We next bound the length of P_5 .

$$\begin{aligned} \text{dist}(s, t', H \setminus \{v_0\}) &\leq |P_5| \leq \text{dist}(s, x, G \setminus \{v_0\}) + 2 + \text{dist}(w_1, w_2, G_\tau \setminus V_f(C_0)) + 2 \\ &= \text{dist}(s, x, G \setminus \{v_0\}) + \text{dist}(C_{v_0}(x), C_0, G_\tau \setminus V_f(C_0)) + 4 \\ &\leq \text{dist}(s, x, G \setminus \{v_0\}) + \text{dist}(x, t', P_\tau) + 4 \\ &= |P_{s, t', v}| + 4 = \text{dist}(s, t', G \setminus \{v_0\}) + 4, \end{aligned} \quad (9)$$

where Eq. (9) follows by Eq. (8). Eq. (5) holds for the pair (t', v_0) , and hence also for the pair (t', v_0) (as (u, t') is the last missing edge on P_{s, t_0, v_0} . The claim holds. \blacksquare



4-additive sourcewise FT-spanner

Figure 2: Schematic illustration of the path-buying procedure of Alg. **Cons4SWSpanner**. Shown is an $s - t$ path $P_\tau = P_{s, t_i, v_0}$ considered to be bought in time τ . The green paths correspond to shortest-paths in $T_0(s)$ and the red edges correspond to missing edges on $P_\tau \setminus E(G_\tau)$. The first missing edge on $P_\tau \setminus E(G_\tau)$ is incident to x . If P_τ was not bought, then there exists a short route between a pair of vertices w_1 and w_2 belonging to $C_{v_0}(x)$ and $C_{v_0}(t_i)$ (respectively) in $H \setminus \{v_0\}$.

Size analysis of $H_4(S)$. We proceed with the size analysis. Consider Step (3). Since $E^{\text{local}}(t)$ includes for every source $s \in S$ at most two last edges of the $s - t$ replacement-paths protecting against the failing of $z_1(t)$ and $\text{LCA}(s, C_1(t))$, it holds that

Observation 4.7 For every $t \in V_C$, $|E^{\text{local}}(t)| = O(|S|)$, hence $|E^{\text{local}}| = O(|S| \cdot n)$.

Bounding the number of last edges in $E_{\text{dep}}^{\text{far}}(t)$. We now turn to bound the number of edges added due to step (4.1), i.e., the last edges of new-ending *dependent* paths $P_{s, t, v}$ protecting against the faults in the far segment $\pi^{\text{far}}(s, t)$. To bound the number of edges in $E_{\text{dep}}^{\text{far}}(t)$, consider the partial BFS tree rooted at t , $T(t, S) \subseteq T_0(t)$, whose leaf set is contained in the vertex set S where $T(t, S) = \bigcup_{s \in S} \pi(s, t)$. It is convenient to view this tree as going from the leafs towards the root, where the root t is at the bottom and the leafs are on the top of the tree. Let $V^+ = S \cup \{u \in T(t, S) \mid \text{deg}(u, T(t, S)) \geq 3\}$, be the union of S and the vertices with degree at least 3 in the tree $T(t, S)$. We have that $|V^+| < 2|S|$. A pair of vertices $x, y \in V^+$ is *adjacent* if their shortest-path $\pi(x, y)$ is contained in the tree $T(t, S)$ and it is free from any other V^+ vertex, i.e., $\pi(x, y) \subseteq T(t, S)$ and $\pi(x, y) \cap V^+ = \{x, y\}$. Let $\Pi(V^+) = \{\pi(x, y) \mid x, y \in V^+ \text{ and } x, y \text{ are adjacent}\}$ be the collection of paths between adjacent pairs.

Observation 4.8 (1) $T(t, S) = \Pi(V^+)$. (2) $\Pi(V^+)$ consists of at most $2|S| + 1$ paths $\pi(x, y)$ (i.e., there are at most $2|S| + 1$ adjacent pairs).

To bound $|E_{\text{dep}}^{\text{far}}(t)|$, we first claim that every two dependent replacement paths with the same divergence point have the same last edge.

Lemma 4.9 For every two dependent paths $P_{s_1, t, v_1}, P_{s_2, t, v_2} \in \mathcal{P}_{\text{dep}}^{\text{far}}(t)$, if $b_{s_1, t, v_1} = b_{s_2, t, v_2}$ then $\text{LastE}(P_{s_1, t, v_1}) = \text{LastE}(P_{s_2, t, v_2})$.

Proof: Let $b = b_{s_1, t, v_1} = b_{s_2, t, v_2}$. Since $b \in V(\pi(s_1, t)) \cap V(\pi(s_2, t))$ it holds that $\pi(s_i, t) = \pi(s_i, b) \circ \pi(b, t)$ for $i \in \{1, 2\}$. In addition, since $P_{s_i, t, v_i}[s_i, b] = \pi(s_i, b)$ for $i \in \{1, 2\}$, it holds that both failing vertices v_1 and v_2 occur in the common segment $\pi(b, t)$. Recall that P_{s_i, t, v_i} is a new-ending path, hence by the definition of the divergence point b (see Lemma 4.2(b)), it holds that $V(P_{s_i, t, v_i}[b, t]) \cap V(\pi(b, t)) = \{b, t\}$ and hence both detours are free from the failing vertices. Hence, $P_{s_1, t, v_1}[b, t] = P_{s_2, t, v_2}[b, t] = SP(b, t, G \setminus \{v_1, v_2\}, W)$. We get that $\text{LastE}(P_{s_1, t, v_1}) = \text{LastE}(P_{s_2, t, v_2})$ as needed. ■

Since our goal is to bound the number of last edges of the new ending dependent paths $\mathcal{P}_{\text{dep}}^{\text{far}}(t)$, to avoid double counting, we now restrict attention to $\mathcal{Q}^{\text{far}}(t)$, a collection of representative paths in $\mathcal{P}_{\text{dep}}^{\text{far}}(t)$ each ending with a distinct new edge from $E_{\text{dep}}^{\text{far}}(t)$. Formally, for each new edge $e \in E_{\text{dep}}^{\text{far}}(t)$, let $P(e)$ be an arbitrary path in $\mathcal{P}_{\text{dep}}^{\text{far}}(t)$ satisfying that $\text{LastE}(P(e)) = e$. Let $\mathcal{Q}^{\text{far}}(t) = \{P(e), e \in E_{\text{dep}}^{\text{far}}(t)\}$ (hence $|\mathcal{Q}^{\text{far}}(t)| = |E_{\text{dep}}^{\text{far}}(t)|$). From now on, we aim towards bounding the cardinality of $\mathcal{Q}^{\text{far}}(t)$. Let $\text{DP} = \{b_{s, t, v} \mid P_{s, t, v} \in \mathcal{Q}^{\text{far}}(t)\}$ be the set of divergence points of the new ending paths in $\mathcal{Q}^{\text{far}}(t)$. By Lemma 4.9, it holds that in order to bound the cardinality of $\mathcal{P}_{\text{dep}}^{\text{far}}(t)$, it is sufficient to bound the number of distinct divergence points. To do that, we show that every path $\pi(x, y)$ of two adjacent vertices $x, y \in V^+$, contains at most one divergence point in $\text{DP} \setminus V^+$.

Lemma 4.10 $|\pi(x, y) \cap (\text{DP} \setminus V^+)| \leq 1$ for every $\pi(x, y) \in \Pi(V^+)$.

Proof: Assume, towards contradiction, that there are two divergence points b_{s_1, t, v_1} and b_{s_2, t, v_2} on some path $\pi(x, y)$ for two adjacent vertices $x, y \in V^+$. For ease of notation, let $P_i = P_{s_i, t, v_i}$, $b_i = b_{s_i, t, v_i}$, $D_i = D_{s_i, t, v_i}$ and $D_i^- = D_i \setminus \{b_i\}$ for $i \in \{1, 2\}$. Without loss of generality, assume the following: (1) y is closer to t than x and (2) b_2 is closer to t than b_1 . By construction, the vertices s_1 and s_2 are in the subtree $T(x) \subseteq T(t, S)$. For an illustration see Fig. 3. We now claim that the failing vertices v_1, v_2 occur in the interior of $\pi(y, t)$. Since D_1^- and D_2^- are vertex disjoint with $\pi(y, t) \setminus \{t\}$, it would imply that both detour segments D_1 and D_2 are free from the failing vertices and hence at least one of the two new edges $\text{LastE}(P_1), \text{LastE}(P_2)$ could have been avoided. We now focus on v_1 and show that $v_1 \in V(\pi(y, t))$, the exact same argumentation holds for v_2 . Since P_1 is a new-ending dependent path, by Eq. (1), there exists some source $s_3 \in S \setminus \{s_1\}$ satisfying that $(D_1^- \cap \pi(s_3, t)) \setminus \{t\} \neq \emptyset$. Let $w \in (V(D_1^-) \cap V(\pi(s_3, t))) \setminus \{t\}$ be the first intersection point (closest to s_1). See Fig. 3 for schematic illustration. We first claim that s_3 is not in $T(x)$ where $T(x)$ is the subtree of $T(t, S)$ rooted at x . To see why this holds, assume, towards contradiction, that $s_3 \in T(x)$. It then holds that the replacement path P_1 has the following form $P_1 = \pi[s_1, x] \circ \pi(x, b_1) \circ P_1[b_1, w] \circ P_1[w, t]$. Recall, that since $b_1 \in \text{DP} \setminus V^+$, $b_1 \neq x$ and also $b_1 \neq w$. Since $P_1[x, w]$ goes through b_1 , by the optimality of P_1 , it holds that

$$\text{dist}(x, w, G \setminus \{v_1\}) > \text{dist}(b_1, w, G \setminus \{v_1\}) . \quad (10)$$

On the other hand, the path $\pi(s_3, t)$ has the following form: $\pi(s_3, t) = \pi(s_3, w) \circ \pi(w, x) \circ \pi(x, b_1) \circ \pi(b_1, t)$. Hence, $\pi(w, b_1)$ goes through x . Since the failing vertex $v_1 \in V(\pi(b_1, t))$ is not in $\pi(w, b_1)$, by the optimality of $\pi(w, b_1)$, we get that $\text{dist}(w, b_1, G \setminus \{v_1\}) > \text{dist}(x, w, G \setminus \{v_1\})$, leading to contradiction with Ineq. (10). Hence, we conclude that $s_3 \notin T(x)$ (in particular this implies that $s_3 \neq s_2$). Note that $\pi(w, t)$ is a segment of $\pi(s_3, t)$ and hence it is contained in the tree $T(t, S)$. Since P_1 is a new-ending path (i.e., $\text{LastE}(P_1) \notin T(t, S)$), we have that $P_1[w, t] \neq \pi(w, t)$ are distinct $w-t$ paths. We next claim that the failing vertex v_1 must occur on $\pi(w, t)$ and hence also on $\pi(s_3, t)$. To see this, observe that if $\pi(w, t)$ would have been free from the failing vertex v_1 , then it implies that $\pi(w, t) = SP(w, t, G \setminus \{v_1\}, W) = P_1[w, t]$, contradiction as $\text{LastE}(P_1) \neq \text{LastE}(\pi(w, t))$. Finally, we show that $v_1 \in V(\pi(y, t))$. By the above, the failing vertex v_1 is common to both paths $\pi(s_1, t)$ and $\pi(s_3, t)$, i.e., $v_1 \in V(\pi(s_1, t)) \cap V(\pi(s_3, t))$. By the definition of the path $\pi(x, y)$, all its internal vertices u have degree 2 and hence $(V(\pi(x, y)) \cap V(\pi(s_3, t))) \setminus \{y\} = \emptyset$, concluding that $v_1 \in V(\pi(y, t))$. By the same argumentation, it also holds that v_2 is in $\pi(y, t)$. As the detours D_1 and D_2 are vertex disjoint with $\pi(y, t) \setminus \{t\}$, it holds that they are free from the two failing vertices, i.e., $v_1, v_2 \notin D_1 \cup D_2$. Since $P_1, P_2 \in \mathcal{Q}^{\text{far}}(t)$, it holds that $\text{LastE}(P_1) \neq \text{LastE}(P_2)$, and hence there are two b_1-t distinct shortest paths in $G \setminus \{v_1, v_2\}$, given by D_1 and $\pi(b_1, b_2) \circ D_2$. By optimality of these paths, they are of the same lengths. Again, we end with contradiction to the uniqueness of the weight assignment W . The claim follows. \blacksquare

We now are now ready to bound $|E_{\text{dep}}^{\text{far}}(t)|$.

Lemma 4.11 *For every $t \in V_C$, $|E_{\text{dep}}^{\text{far}}(t)| = O(|S|)$.*

Proof: By Lemma 4.9 there are at most $|V^+|$ replacement paths with divergence point in V^+ . By Lemma 4.10, there is at most one divergence point on each segment $\pi(x, y)$ of an adjacent pair (x, y) . Note that the divergence points DP are in the tree $T(t, S)$ and that the internal segments of $\pi(x, y), \pi(x', y')$ for $x, x', y, y' \in V^+$ are vertex disjoint. Combining with Obs. 4.8(2), we get $|E^{\text{far}}(t)| = |\mathcal{Q}^{\text{far}}(t)| = O(|S|)$. The lemma follows. \blacksquare

We complete the size analysis and proves Theorem 4.1, by bounding the number of edges added by the path-buying procedure of Step (4.2).

Bounding the number of edges added due to the path-buying procedure. Finally, it remains to bound the number of edges added due to the path-buying procedure of step (4.2). Let $\mathcal{B} \subseteq \mathcal{P}_{\text{indep}}^{\text{far}}(t)$ be the set of paths bought in the path-buying procedure of Step (4.2). For every ordered pair of clusters $C_1, C_2 \in \mathcal{C}$, let $\mathcal{B}(C_1, C_2) \subseteq \mathcal{B}$ be the set of paths that were added since they improved the distance of C_1 and C_2 , that is

$$\mathcal{B}(C_1, C_2) = \{P_\tau \in \mathcal{B} \mid C_{1,\tau} = C_1 \text{ and } C_{2,\tau} = C_2\}$$

Clearly, $\mathcal{B} = \bigcup_{C_1, C_2 \in \mathcal{C}} \mathcal{B}(C_1, C_2)$. We next use the fact that the diameter of each cluster $C \in \mathcal{C}$ is small, to bound the cardinality of the set $\mathcal{B}(C_1, C_2)$.

Lemma 4.12 $|\mathcal{B}(C_1, C_2)| \leq 5$ for every $C_1, C_2 \in \mathcal{C}$.

Proof: Fix $C_1, C_2 \in \mathcal{C}$ and order the paths of $\mathcal{B}(C_1, C_2)$ according to the time step they were added to the spanner $\mathcal{B}(C_1, C_2) = \{P_{\tau_1}, \dots, P_{\tau_N}\}$ and $\tau_1 < \tau_2 < \dots < \tau_N$ where $N = |\mathcal{B}(C_1, C_2)|$. Since $P_{\tau_k} \in \mathcal{P}_{\text{indep}}^{\text{far}}$, it is a new-ending path, i.e., $\text{LastE}(P_{\tau_k}) \notin G_0$. Let $P_{\tau_k} = P_{s_k, t_k, v_k}$ and




Figure 3: Schematic illustration of new-ending dependent paths. Shown is the tree $T(t, S)$ with the root t at the bottom and leaf set is contained in the set of sources S . (a) The two replacement paths have the same divergence point b , hence one of the new last edges is redundant. (b) A new-ending $s_1 - t$ dependant path P_{s_1,t,v_1} with a divergence point $b_1 \in V(\pi(x, y))$ intersects with $\pi(s_3, t)$ at the vertex $w \notin \{b_1, t\}$. Since $P_{s_1,t,v}$ is a new-ending path (i.e., its last edges is not on $T(t, S)$), the failing vertex v must occur on the path $\pi(w, t)$. Hence $v_1 \in V(\pi(s_1, t)) \cap V(\pi(s_3, t))$, implying that $v_1 \in V(\pi(y, t))$. Since this holds for any new-ending path with a divergence point in $\pi(x, y)$, we get that only one new edge from all these paths is needed.

$D_{\tau_k} = D_{s_k, t_k, v_k}$ denote the detour segment of this path. Hence, each P_{τ_k} protects against the failing of v_k . Let x_k be the vertex adjacent to the first missing edge on P_{τ_k} . Hence, $C_1 = C_{v_k}(x_k)$ and $C_2 = C_{v_k}(t_k)$ for every $k \in \{1, \dots, N\}$ and also $V_f(C_{v_k}(t_k)) = V_f(C_2)$ for every $k \in \{1, \dots, N\}$. Since $T_0(S) \subseteq G_0$, the missing edges of P_{τ_k} are restricted to the detour segment D_{τ_k} .

In addition, since $P_{\tau_k} \in \mathcal{P}_{\text{indep}}^{\text{far}}$, it holds that the failing vertex v_k occurs on the far segment $\pi^{\text{far}}(s_k, t_k)$ and in particular, $v_k \notin C_2$ (i.e., v_k occurs strictly above the least common ancestor $\text{LCA}(s_k, C_2)$ and since all cluster members appear on $T_0(s_k)$ in the subtree rooted at $\text{LCA}(s_k, C_2)$, the far segment $\pi^{\text{far}}(s_k, t_k)$ is free from cluster members). We therefore have

$$\{v_1, \dots, v_N\} \subseteq V_f(C_2) . \quad (11)$$

Note that each path P_{τ_k} protects against the failing of the single vertex v_k , however, since each P_{τ_k} belongs to $\mathcal{B}(C_1, C_2)$, Eq. (11) holds.

Note that $\pi(s_k, C_2) \subseteq \pi(s_k, t_k)$ and hence it is contained in $T(t_k, S)$ for every $k \in \{1, \dots, N\}$. By the definition of independent paths (see Eq. (1) for the definition of *dependent* paths), we have that

$$D_{s_k, t_k, v_k}^- \cap T(t_k, S) = \{t_k\} . \quad (12)$$

Consequently, by Lemma 4.4(b),

$$P_{\tau_k}[x_k, t_k] \subseteq D_{\tau_k}^- \subseteq G \setminus V_f(C_2) . \quad (13)$$

where the last inclusion holds by the fact that $t_k \notin V_f(C_2)$ and $V_f(C_2) \subseteq \bigcup_{s \in S} \pi(s, C_2) \subseteq \bigcup_{s \in S} \pi(s, t_k) = T(t_k, S)$. Let $z_i \in Z$ be the cluster center of C_i for $i \in \{1, 2\}$. We therefore have that $z_1 = z_{v_1}(x_1) = \dots = z_{v_N}(x_N)$ and $z_2 = z_{v_1}(t_1) = \dots = z_{v_N}(t_N)$ and hence $z_2 \neq v_k$ for every $k \in \{1, \dots, N\}$. Hence,

$$z_1, z_2 \notin \{v_1, \dots, v_N\} . \quad (14)$$

For every $k \in \{1, \dots, N\}$, denote

$$X_k = \text{dist}(x_k, t_k, G_{\tau_{k+1}} \setminus V_f(C_2)).$$

We now show that $X_k < X_{k-1}$ for every $k \in \{2, \dots, N\}$.

Since the path P_{τ_k} is purchased at time τ_k , we have that

$$X_k \leq \text{dist}(x_k, t_k, P_{\tau_k} \setminus V_f(C_2)) \quad (15)$$

$$= \text{dist}(x_k, t_k, P_{\tau_k}) \quad (16)$$

$$< \text{dist}(C_1, C_2, G_{\tau_k} \setminus V_f(C_2)) \quad (17)$$

$$\leq X_{k-1}, \quad (18)$$

where Eq. (15) follows by the fact that $P_{\tau_k} \subseteq G_{\tau_{k+1}}$, Eq. (16) follows by Eq. (13). Eq. (17) follows by the fact that P_{τ_k} was bought and by Eq. (4), and Eq. (18) follows by the fact that $x_{k-1} \in C_1$ and $t_{k-1} \in C_2$.

Therefore, we have that

$$X_N \leq X_1 - (N - 1). \quad (19)$$

Conversely, we have that

$$X_N \geq \text{dist}(x_N, t_N, G \setminus \{v_1, \dots, v_N\}) \quad (20)$$

$$\geq \text{dist}(x_1, t_1, G \setminus \{v_1, \dots, v_N\}) - 4 \quad (21)$$

$$= \text{dist}(x_1, t_1, P_{\tau_1}) - 4 = \text{dist}(x_1, t_1, P_{\tau_1} \setminus V_f(C_2)) - 4 \quad (22)$$

$$\geq X_1 - 4, \quad (23)$$

where Eq. (20) follows as $G_{\tau_{N+1}} \subseteq G$ and by Eq. (11), $\{v_1, \dots, v_N\} \subseteq V_f(C_2)$. To see Eq. (21), we need to prove the existence of the intracluster paths $R_1 = [x_1, z_1, x_N]$ and $R_2 = [t_1, z_2, t_N]$ in $G \setminus \{v_1, \dots, v_N\}$ where z_1 (resp., z_2) is the cluster center of C_1 (resp., C_2). By definition, $x_1, x_N \in C_1$ and $t_1, t_N \in C_2$. Hence, z_1 (resp., z_2) is a common neighbor of both x_1 and x_N (resp., t_1 and t_N). By (14), $z_1, z_2 \notin \{v_1, \dots, v_N\}$.

In addition, by Eq. (13), $x_k, t_k \in G \setminus V_f(C_2)$ for every $k \in \{1, \dots, N\}$ and by Eq. (11), it also holds that $x_k, t_k \notin \{v_1, \dots, v_N\}$ for every $k \in \{1, \dots, N\}$. Hence, R_1 and R_2 exist in $G \setminus \{v_1, \dots, v_N\}$ and Eq. (21) follows by the triangle inequality. Eq. (22) follows by Eq. (13) and Eq. (11). Finally, Eq. (23) follows by the fact that P_{τ_1} was added at step τ_1 , hence $P_{\tau_1} \subseteq G_{\tau_2}$. We get that $N \leq 5$. Lemma 4.12 follows. ■

We are now ready to bound the number of edges added in the path-buying phase.

Lemma 4.13 $|H_4(S) \setminus G_0| = O((n/|S|)^3)$.

Proof: By Obs. 2.3(2) and Obs. 2.4, every path P_{τ_k} contains at most $O(n/\Delta)$ edges in $G \setminus G_\Delta$. Hence,

$$|E(G_{\tau'} \setminus G_0)| = O(n/\Delta) \cdot |\mathcal{B}| \quad (24)$$

$$= O(n/\Delta) \cdot \sum_{C_1, C_2 \in \mathcal{C}} |\mathcal{B}(C_1, C_2)| \quad (25)$$

$$\leq O(n/\Delta) \cdot |\mathcal{C}|^2 = O((n/\Delta)^3). \quad (26)$$

where the last equality follows by the fact that $|\mathcal{C}| = O(n/\Delta)$. Taking $\Delta = |S|$ establishes the lemma. ■

4.2 Sourcewise spanner with additive stretch 8

In this section, we present Alg. `Cons8SWSpanner` for constructing a sourcewise additive FT-spanner with additive stretch 8. The size of the resulting spanner is smaller (in order) than the $H_4(S)$ spanner of Alg. `Cons4SWSpanner`, at the expense of larger stretch. The algorithm is similar in spirit to Alg. `Cons4SWSpanner` and the major distinction is in the path-buying procedure of step (4.2).

Lemma 4.14 *There exists a subgraph $H_8(S) \subseteq G$ with $O(|S| \cdot n + (n/|S|)^2)$ edges s.t. $\text{dist}(s, t, H_8(S) \setminus \{v\}) \leq \text{dist}(s, t, G \setminus \{v\}) + 8$ for every $(s, t) \in S \times V$ and every $v \in V$.*

4.2.1 Algorithm `Cons8SWSpanner` for constructing $H_8(S)$ spanner

Step (0-4.1): Same as in Alg. `Cons4SWSpanner`. Let $E^{\text{local}}, E_{\text{dep}}^{\text{far}}$ be the set of last edges obtained at the end of step (3) and set (4.1) respectively. Let $\mathcal{P}_{\text{indep}}^{\text{far}}$ be the set of new-ending independent paths.

Step (4.2): Handling independent new-ending paths. Starting with G_0 as in Eq. (2), the paths of $\mathcal{P}_{\text{indep}}^{\text{far}}$ are considered in an arbitrary order. At step τ , we are given $G_\tau \subseteq G$ and consider the path $P_\tau = P_{s,t,v}$. Let $D_\tau = P_\tau \setminus \pi(s, t)$ be the detour segment of P_τ (since $\pi(s, t) \subseteq T_0(S)$ is in G_0 , all missing edges of P_τ occur on its detour segment).

To decide whether P_τ should be added to G_τ , the number of pairwise cluster “distance improvements” is compared to the number of new edges added due to P_τ . To do that we compute the set $\text{ValSet}(P_\tau)$ containing all pairs of clusters that achieves a better distance if P_τ is bought. The value and cost of P_τ are computed as follows. Let $\text{Val}(P_\tau) = |\text{ValSet}(P_\tau)|$ be the number of distance improvements as formally defined later. We next define a key vertex $\phi_\tau \in V_C$ on the path P_τ .

Definition 4.15 *Let $\phi_{s,t,v}$ (or ϕ_τ for short) be the last vertex on P_τ (closest to t) satisfying that: (N1) $\text{LastE}(P_\tau[s, \phi_\tau]) \notin G_\tau$, and (N2) $v \in V^{\text{near}}(s, \phi_\tau) = \pi(\ell, \phi_\tau) \setminus \{\ell\}$ where $\ell = \text{LCA}(s, C_v(\phi_\tau))$. If there is no vertex on P_τ that satisfies both (N1) and (N2), then let ϕ_τ be the first vertex incident to the first missing edge on $P_\tau \setminus E(G_\tau)$ (i.e., such that $P_\tau[s, \phi_\tau]$ is the maximal prefix that is contained in G_τ).*

Let $Q_\tau = P_\tau[\phi_\tau, t]$ and define $\text{Cost}(P_\tau) = |E(Q_\tau) \setminus E(G_\tau)|$ be the number of edges of Q_τ that are missing in the current subgraph G_τ . Thus $\text{Cost}(P_\tau)$ represents the increase in the size of the spanner G_τ if the procedure adds Q_τ . Our algorithm attempts to buy only the suffix Q_τ of P_τ when considering P_τ . We now define the set $\text{ValSet}(P_\tau) \subseteq \mathcal{C} \times \mathcal{C}$ which contains a collection of ordered cluster pairs. Let $C_{1,\tau} = C_v(\phi_\tau)$ and $C_{2,\tau} = C_v(t)$ be the clusters of ϕ_τ and t in $G_\Delta \setminus \{v\}$. Let $\kappa = \text{Cost}(P_\tau)$. The candidate P_τ is said to be *cheap* if $\kappa \leq 4$, otherwise it is *costly*. The definition of $\text{ValSet}(P_\tau)$ depends on whether or not the path is cheap. In particular, if P_τ is cheap, then $\text{ValSet}(P_\tau) = \{(C_{1,\tau}, C_{2,\tau})\}$ only if

$$\text{dist}(\phi_\tau, t, P_\tau) < \text{dist}(C_{1,\tau}, C_{2,\tau}, G_\tau \setminus V_f(C_{2,\tau})) , \quad (27)$$

where $V_f(C_{2,\tau})$ is as given by Eq. (3), and let $\text{ValSet}(P_\tau) = \emptyset$ otherwise. Alternatively, if P_τ is costly, we do the following.

Definition 4.16 Let $U_{s,t,v} = \{u_{3\ell+1} \mid \ell \in \{0, \dots, \lfloor (\kappa - 1)/3 \rfloor\}\} \subseteq Q_\tau$ be some representative endpoints of missing edges on Q_τ satisfying that

$$\text{LastE}(Q_\tau[\phi_\tau, u_\ell]) \notin G_\tau \text{ for every } u_\ell \in U_{s,t,v} \text{ and } \text{dist}(u_\ell, u_{\ell'}, Q_\tau) \geq 3$$

for every $u_\ell, u_{\ell'} \in U_{s,t,v}$.

Define

$$\begin{aligned} \text{ValSet}_1(P_\tau) = \{ & (C_{1,\tau}, C_\ell) \mid C_\ell = C_v(u_\ell), u_\ell \in U_{s,t,v} \\ & \text{and } \text{dist}(\phi_\tau, u_\ell, P_\tau) < \text{dist}(C_{1,\tau}, C_\ell, G_\tau \setminus V_f(C_\ell))\} \end{aligned} \quad (28)$$

and

$$\begin{aligned} \text{ValSet}_2(P_\tau) = \{ & (C_\ell, C_{2,\tau}) \mid C_\ell = C_v(u_\ell), u_\ell \in U_{s,t,v} \\ & \text{and } \text{dist}(u_\ell, t, P_\tau) < \text{dist}(C_\ell, C_{2,\tau}, G_\tau \setminus V_f(C_{2,\tau}))\} \end{aligned} \quad (29)$$

Let $\text{ValSet}(P_\tau) = \text{ValSet}_1(P_\tau) \cup \text{ValSet}_2(P_\tau)$. The subpath Q_τ is added to G_τ resulting in $G_{\tau+1}$ only if

$$\text{Cost}(P_\tau) \leq 4 \cdot \text{Val}(P_\tau) , \quad (30)$$

where $\text{Val}(P_\tau) = |\text{ValSet}(P_\tau)|$. (Note that when P_τ is cheap, Eq. (30) holds iff Eq. (27) holds.) The output of Alg. `Cons8SWSpanner` is the subgraph $H_8(S) = G_{\tau'}$ where $\tau' = |\mathcal{P}_{\text{indep}}^{\text{far}}|$. This completes the description of the algorithm.

Analysis. Throughout the discussion, a path $P_{s,t,v}$ is a new-ending path, if $\text{LastE}(P_{s,t,v}) \notin G_0$ (see Eq. (2)). Hence, we consider only $P_{s,t,v} \in \mathcal{P}_{\text{indep}}^{\text{far}}(t)$ paths for clustered vertices $t \in V_C$.

For a new-ending path $P_{s,t,v}$, recall that $b_{s,t,v}$ is the unique divergence point of $P_{s,t,v}$ and $\pi(s, t)$ and let $D_{s,t,v}$ be the detour segment, i.e., $D_{s,t,v} = P_{s,t,v}[b_{s,t,v}, t]$ and $D_{s,t,v}^- = D_{s,t,v} \setminus \{b_{s,t,v}\}$. Let $Q_{s,t,v} = P_{s,t,v}[\phi_{s,t,v}, t]$ be the path segment that was considered to be bought in step (4.2) (see Def. 4.15).

Observation 4.17 $Q_{s,t,v} \subseteq D_{s,t,v}^-$.

Proof: Let x be the first vertex incident to a missing-edge on $P_{s,t,v}$ (such that $P_{s,t,v}[s, x]$ is the maximal prefix that is contained in G_0). Since $\phi_{s,t,v}$ occurs not before x on $P_{s,t,v}$ the observation follows by Lemma 4.4(b). ■

The main essence of the path-buying procedure (in the fault-free setting) is that the number of distance improvements between any fixed pair of clusters with bounded diameter is *bounded*. This essential argument fails to hold when the distances are measured in different subgraphs. Since in the FT setting the candidate path to be bought, $P_{s,t,v}$, should be compared against some alternative path in the current spanner $H' \setminus \{v\}$, the distances between clusters might be evaluated in distinct subgraphs. Hence, the main challenge in adapting the path-buying scheme to the FT setting is in showing that due to the special structure of the independent paths $P_{s,t,v}$, the distance improvements between any pair of clusters C_1 and C_2 that are incident on $P_{s,t,v}$ can be carried out in the “same” subgraph, i.e., a subgraph that depends only on the clusters C_1 and C_2 and *independent* of the source s and the failing vertex v . The independence of the source s is given by the fact that the

paths are independent and hence their internal detour segments does not intersects any $\pi(s', t)$ path, for $s' \in S$. The independence of the failing vertex v is due to the fact that all failing events occur above the least common ancestor of the cluster members in the BFS tree rooted at s . The next lemma formalizes some of the above intuition and provide the key properties that enables the definition of the graph in which the path $P_{s,t,v}$ would be evaluated in our path-buying scheme.

Lemma 4.18 *Let $P_{s,t,v} \in \mathcal{P}_{\text{indep}}^{\text{far}}(t)$ be a new-ending replacement path. Then for every $u_k \in U_{s,t,v} \cup \{t\}$ with $C_k = C_v(u_k)$ it holds that:*

- (a) $C_k = C_1(u_k)$.
- (b) $V(P_{s,t,v}[b_{s,t,v}, u_k]) \cap V(T(u_k, S)) = \{b_{s,t,v}, u_k\}$.
- (c) $Q_{s,t,v}[\phi_{s,t,v}, u_k] \cap V_f(C_k) = \emptyset$.
- (d) $v \in V_f(C_k)$.

Proof: We begin with (a). By the uniqueness of the weight assignment W , $P_{s,t,v}[s, u_k] = P_{s,u_k,v} = SP(s, u_k, G \setminus \{v\}, W)$. By the uniqueness of the divergence point $b_{s,t,v}$ and in particular by Lemma 4.2(b),

$$b_{s,t,v} = b_{s,u_k,v} . \quad (31)$$

Since $\text{LastE}(P_{s,u_k,v}) \notin E^{\text{local}}$, it follows that $u_k \in V_C$, $v \neq z_1(u_k)$. Hence $z_v(u_k) = z_1(u_k)$ and (a) holds.

Consider (b). By the definition of the set $U_{s,t,v}$ (see Def. 4.16), it holds $\text{LastE}(P_{s,u_k,v}) \notin G_0$. Since $u_k \in Q_{s,t,v}$ occurs strictly after $\phi_{s,t,v}$, by the Def. 4.15, it holds that u_k did not satisfy property (N2). Hence, since $\text{LastE}(P_{s,u_k,v}) \notin E^{\text{local}}$, $v \notin \{z_1(u_k), \text{LCA}(s, C_k)\}$ and hence $v \in V^{\text{far}}(s, u_k)$. As $\text{LastE}(P_{s,u_k,v}) \notin E_{\text{dep}}^{\text{far}}(u_k)$, we get that $P_{s,u_k,v}$ is a new-ending *independent* path. By Eq. (1), $V(P_{s,u_k,v}[b_{s,u_k,v}, u_k]) \cap V(T(u_k, S)) = \{b_{s,u_k,v}, u_k\}$. Hence (b) holds by Eq. (31).

We now turn to consider claim (c). By Eq. (3), $V_f(C_k) \subseteq T(u_k, S)$. Since $C_v(u_k) \cap V_f(C_k) = \emptyset$, it holds that $u_k \notin V_f(C_k)$, and hence by combining with claim (b), we get that $P_{s,u_k,v}[b_{s,u_k,v}, u_k] \cap V_f(C_k) = \{b_{s,u_k,v}\}$. Since by Obs. 4.17, $\phi_{s,t,v} \neq b_{s,u_k,v}$, hence $Q_{s,t,v}[\phi_{s,t,v}, u_k] \cap V_f(C_k) = \emptyset$. Claim (c) holds.

Consider claim (d). By the above, v occurs on the far segment $\pi(s, C_k) \setminus \{\text{LCA}(s, C_k)\}$, hence $v \notin C_k$. Since $(\pi(s, C_k) \setminus C_k) \subseteq V_f(C_k)$, (d) holds. \blacksquare

The next observation is useful in our analysis.

Observation 4.19 *If $\phi_{s,t,v}$ satisfies (N1) and (N2), then there exists a vertex $x \in C_v(\phi_{s,t,v})$ satisfying that $v \notin V(\pi(s, x))$.*

Proof: Let $P_\tau = P_{s,t,v}$ and $\phi_\tau = \phi_{s,t,v}$. By the uniqueness of the weight assignment W , $P_\tau[s, \phi_\tau] = P_{s,\phi_\tau,v} = SP(s, \phi_\tau, G \setminus \{v\}, W)$. Since ϕ_τ satisfies (N2), it holds that the failing vertex v occurs on $\pi^{\text{near}}(s, \phi_\tau)$, strictly *below* (i.e., closer to ϕ_τ) the least common ancestor $\text{LCA}(s, C_v(\phi_\tau))$ on $\pi(s, \phi_\tau)$. Hence, there must exist a vertex $x \in C_v(\phi_\tau)$ such that $v \notin V(\pi(s, x))$ (otherwise, if v is shared by $\pi(s, u)$ for all cluster members u , then we end with contradiction to the definition of the least common ancestor $\text{LCA}(s, C_v(\phi_\tau))$). \blacksquare

We proceed by showing correctness.

Theorem 4.20 *$H_8(S)$ is a $(8, S)$ FT-spanner.*

Proof: Let $H = H_8(S)$. It is required to show that $\text{dist}(s, t, H \setminus \{v\}) \leq \text{dist}(s, t, G \setminus \{v\}) + 8$ for every $(s, t) \in S \times V$ and $v \in V$. By the analysis of Alg. **Cons4SWSpanner** (Lemma 4.6), it remains to consider the case of independent new-ending paths where $P_{s,t,v} \in \mathcal{P}_{\text{indep}}^{\text{far}}(t)$ for $t \in V_C$.

Let τ be the iteration at which $P_\tau = P_{s,t,v}$ was considered to be added to the spanner at step (4.2), and let $\kappa = \text{Cost}(P_\tau)$ denote its cost. Let ϕ_τ be as defined in Def. 4.15 and recall that $Q_\tau = P_\tau[\phi_\tau, t]$ is the candidate suffix to be bought by the procedure. (In particular, $\text{Cost}(P_\tau)$ counts the number of edges on $Q_\tau \setminus E(G_\tau)$.)

Case (1): Q_τ was bought. If ϕ_τ did not satisfy either properties (N1) or (N2), then $P_\tau[s, \phi_\tau] \subseteq G_\tau$. Since $P_\tau = P_\tau[s, \phi_\tau] \circ Q_\tau$ and Q_τ was added to the spanner, we get that $P_\tau \subseteq H \setminus \{v\}$.

It remains to consider the complementary case where ϕ_τ satisfies both (N1) and (N2). By Obs. 4.19, we get that there exist $x \in C_v(\phi_\tau)$ satisfying that $v \notin V(\pi(s, x))$.

Consider the path $P = \pi(s, x) \circ (x, z_v(\phi_\tau), \phi_\tau)$. By definition, $P \subseteq H \setminus \{v\}$ and by the existence of the intracluster path connecting x and ϕ_τ in $G \setminus \{v\}$, it holds that $|P| = \text{dist}(s, x, G \setminus \{v\}) + 2 \leq \text{dist}(s, \phi_\tau, G \setminus \{v\}) + 4$. Hence, letting $P' = P \circ Q_\tau$ (where $Q_\tau = P_\tau[\phi_\tau, t]$), since $Q_\tau \subseteq H \setminus \{v\}$, it holds that $P' \subseteq H \setminus \{v\}$ and $|P'| \leq |P_\tau| + 4$, as required.

Case (2): Q_τ was not bought. Let $x \in C_v(\phi_\tau)$ be defined as follows. If ϕ_τ satisfies both properties (N1) and (N2) of Def. 4.15, then using Obs. 4.19, let $x \in C_v(\phi_\tau)$ be the vertex satisfying that $v \notin V(\pi(s, x))$. Otherwise, if ϕ_τ did not satisfy (N1) or (N2) (or both), let $x = \phi_\tau$. Note that in any case, it holds that $x, \phi_\tau \in C_v(\phi_\tau)$. We have the following.

Lemma 4.21 $P_{s,x,v} \subseteq H \setminus \{v\}$.

Proof: If $x = \phi_\tau$, then it implies that ϕ_τ did not satisfy both of the properties (N1, N2). By Def. 4.15, in such a case ϕ_τ is the vertex incident to the first missing edge on $P_{s,t,v} \setminus E(G_\tau)$ and hence $P_{s,t,v}[s, x] = P_{s,x,v} \subseteq G_\tau \setminus \{v\}$.

Otherwise, if $x \neq \phi_\tau$, then $x \in C_v(\phi_\tau)$ and by the selection of x , $v \notin V(\pi(s, x))$. Hence, $P_{s,x,v} = \pi(s, x) \subseteq H \setminus \{v\}$. ■

Recall that $C_{1,\tau} = C_v(\phi_\tau)$ and $C_{2,\tau} = C_v(t)$. In addition, since $v \in V^{\text{far}}(s, t)$, it holds that $v \in V_f(C_{2,\tau})$.

Case (2.1): P_τ is cheap. Since Q_τ was not added, Eq. (27) did not hold and hence

$$\text{dist}(\phi_\tau, t, P_\tau) \geq \text{dist}(C_{1,\tau}, C_{2,\tau}, G_\tau \setminus V_f(C_{2,\tau})) . \quad (32)$$

Let $w_1 \in C_{1,\tau}$ and $w_2 \in C_{2,\tau}$ be a closest pair satisfying that $\text{dist}(w_1, w_2, G_\tau \setminus V_f(C_{2,\tau})) = \text{dist}(C_{1,\tau}, C_{2,\tau}, G_\tau \setminus V_f(C_{2,\tau}))$. Since the failing vertex v is in $V_f(C_{2,\tau})$, both auxiliary vertices w_1 and w_2 are in $G \setminus \{v\}$. Consider the following $s-t$ path: $P = P_0 \circ P_1 \circ P_2 \circ P_3$ where $P_0 = P_{s,x,v}$, $P_1 = [x, z_v(\phi_\tau), w_1]$, $P_2 \in SP(w_1, w_2, G_\tau \setminus V_f(C_{2,\tau}))$, and $P_3 = [w_2, z_v(t), t]$. For an illustration see Fig. 4. By Lemma 4.21, $P_0 \subseteq H \setminus \{v\}$. Note that since $x, w_1 \in C_v(\phi_\tau)$, the path P_1 exists in $H \setminus \{v\}$. Combining with the definitions of the vertices $z_v(x), z_v(t), w_1, w_2$, it holds that $P \subseteq H \setminus \{v\}$. So, it

remains to bound the length of the path.

$$\begin{aligned}
\text{dist}(s, t, H \setminus \{v\}) &\leq |P_0| + |P_1| + |P_2| + |P_3| \\
&= \text{dist}(s, x, G \setminus \{v\}) + \text{dist}(w_1, w_2, G_\tau \setminus V_f(C_{2,\tau})) + 4 \\
&\leq \text{dist}(s, \phi_\tau, G \setminus \{v\}) \\
&\quad + \text{dist}(w_1, w_2, G_\tau \setminus V_f(C_{2,\tau})) + 6 \tag{33} \\
&\leq \text{dist}(s, \phi_\tau, G \setminus \{v\}) + \text{dist}(\phi_\tau, t, P_\tau) + 6 = |P_\tau| + 6, \tag{34}
\end{aligned}$$

where Eq. (33) follows by the fact that $x, \phi_\tau \in C_v(\phi_\tau)$ and since $G_\Delta \subseteq H$, it holds that the intracluster path $R = [x, z_v(\phi_\tau), \phi_\tau]$ exists in $G \setminus \{v\}$, Eq. (34) follows by Eq. (32).

Case (2.2): P_τ is costly. Let $U_{s,t,v} = \{u_1, \dots, u_{\kappa'}\} \subseteq Q_\tau$ for $\kappa' = \lfloor \kappa/3 \rfloor \geq 1$ be as defined by Def. 4.16. Since by Obs. 2.3, the diameter of each cluster is 2, each $u_k \in U_{s,t,v}$ belongs to a distinct cluster $C_k = C_v(u_k) \in \mathcal{C}$. Hence there are at least κ' distinct clusters on Q_τ .

A cluster $C_k = C_v(u_k)$ is a *contributor* if adding Q_τ to G_τ improves either the $C_{1,\tau} - C_k$ distance (i.e., $(C_{1,\tau}, C_k) \in \text{ValSet}_1(P_\tau)$) or the $C_{2,\tau} - C_k$ distance (i.e., $(C_k, C_{2,\tau}) \in \text{ValSet}_2(P_\tau)$) in the corresponding appropriate graph. Otherwise, C_k is *neutral*. There are two cases to consider. If all clusters are contributors (i.e., there is no neutral cluster) then all the κ' clusters contribute to $\text{Val}(P_\tau)$ (either with $C_{1,\tau}$ or with $C_{2,\tau}$ or both). It then holds that $\text{Val}(P_\tau) \geq \kappa' \geq \text{Cost}(P_\tau)/4$. Hence, by Eq. (30), we get a contradiction to the fact that the suffix Q_τ was not added to G_τ .

In the other case, there exists at least one neutral cluster C_ℓ such that

$$\begin{aligned}
\text{dist}(C_{1,\tau}, C_k, \widehat{H}_1) &\leq \text{dist}(\phi_\tau, u_k, P_\tau) \quad \text{and} \tag{35} \\
\text{dist}(C_k, C_{2,\tau}, \widehat{H}_2) &\leq \text{dist}(u_k, t, P_\tau),
\end{aligned}$$


where $\widehat{H}_1 = G_\tau \setminus V_f(C_k)$ and $\widehat{H}_2 = G_\tau \setminus V_f(C_{2,\tau})$. Let $w_1 \in C_{1,\tau}$ and $w_2 \in C_k$ be the pair of vertices satisfying $\text{dist}(w_1, w_2, \widehat{H}_1) = \text{dist}(C_{1,\tau}, C_k, \widehat{H}_1)$. In addition, let $y_1 \in C_k$ and $y_2 \in C_{2,\tau}$ be the pair satisfying $\text{dist}(y_1, y_2, \widehat{H}_2) = \text{dist}(C_k, C_{2,\tau}, \widehat{H}_2)$.

Let $Q_1 = [x, z_v(\phi_\tau), w_1]$, $Q_2 = [w_2, z_v(u_k), y_1]$ and $Q_3 = [y_2, z_v(t), t]$ be the intracluster paths in $C_{1,\tau}, C_k$ and $C_{2,\tau}$ respectively. Note that by definition $x, w_1 \in C_v(\phi_\tau)$.

Since by Lemma 4.18(d), $v \in V_f(C_k) \cap V_f(C_{2,\tau})$, it also holds that $Q_1, Q_2, Q_3 \subseteq H \setminus \{v\}$. Let $P' = P_0 \circ Q_1 \circ P_1 \circ Q_2 \circ P_2 \circ Q_3$ where $P_0 = P_{s,x,v}$, $P_1 \in SP(w_1, w_2, \widehat{H}_1)$ and $P_2 \in SP(y_1, y_2, \widehat{H}_2)$. By Lemma 4.21, $P_0 \subseteq H \setminus \{v\}$ and by the above explanation, $P' \subseteq H \setminus \{v\}$. So, it remains to bound the length of the $s - t$ path P' .

$$\begin{aligned}
\text{dist}(s, t, H \setminus \{v\}) &\leq |P'| = |P_0| + |P_1| + |P_2| + 6 \\
&= \text{dist}(s, x, G \setminus \{v\}) + \text{dist}(w_1, w_2, \widehat{H}_1) + \text{dist}(y_1, y_2, \widehat{H}_2) + 6 \\
&\leq \text{dist}(s, \phi_\tau, G \setminus \{v\}) + \text{dist}(w_1, w_2, \widehat{H}_1) + \text{dist}(y_1, y_2, \widehat{H}_2) + 8 \\
&= \text{dist}(s, \phi_\tau, G \setminus \{v\}) + \text{dist}(C_{1,\tau}, C_k, \widehat{H}_1) \\
&\quad + \text{dist}(C_k, C_{2,\tau}, \widehat{H}_2) + 8 \\
&\leq \text{dist}(s, \phi_\tau, G \setminus \{v\}) + \text{dist}(\phi_\tau, u_k, P_\tau) + \text{dist}(u_k, t, P_\tau) + 8 \\
&= |P_{s,t,v}| + 8,
\end{aligned}$$

where the first inequality follows by the fact that $x, \phi_\tau \in C_v(\phi_\tau)$ and hence the intracluster path $R = [x, z_v(\phi_\tau), \phi_\tau]$ exists in $G \setminus \{v\}$ and last inequality follows Eq. (35). The lemma follows. \blacksquare



8-additive sourcewise FT-spanner.

Figure 4: Schematic illustration of the path-buying procedure of Alg. `Cons8SWSpanner`. The horizontal path is $P_\tau = P_{s,t,v}$ whose segment $Q_{s,t,v} = P_{s,t,v}[\phi_{s,t,v}, t]$ was considered to be bought at time τ . The green paths correspond to the shortest paths in $T_0(s)$. Red edges correspond to missing edges on $P_{s,t,v} \setminus E(G_\tau)$. The vertex $\phi_{s,t,v}$ satisfies properties (N1) and (N2), hence it is incident to a missing edge and the failing vertex v occurs on $\pi(s, \phi_{s,t,v})$ strictly below the LCA vertex $\text{LCA}(\phi_{s,t,v}) = \text{LCA}(s, C_v(\phi_{s,t,v}))$. The vertex $x \in C_v(\phi_{s,t,v})$ satisfies that $v \notin V(\pi(s, x))$. The $s - t$ replacement path in $H \setminus \{v\}$ is given by traveling from s to x on $\pi(s, x)$ and then use the closest vertex pairs w_1, w_2 and y_1, y_2 .

Finally, we turn to bound the size of $H = H_8(S)$. By the size-analysis of Alg. `Cons4SWSpanner`, it remains to bound the number of edges added due to the path-buying procedure of step (4.2). Let $\mathcal{B} \subseteq \mathcal{P}_{\text{indep}}^{\text{far}}$ be the set of paths corresponding to the path segments that were bought in the path-buying phase. For every ordered pair of clusters, $C_1, C_2 \in \mathcal{C}$ let $\mathcal{B}(C_1, C_2) = \{P_\tau \in \mathcal{B} \mid (C_1, C_2) \in \text{ValSet}(P_\tau)\}$. Clearly, $\mathcal{B} = \bigcup_{C_1, C_2 \in \mathcal{C}} \mathcal{B}(C_1, C_2)$. We next claim that since the diameter of each cluster is small, it holds that the cardinality of each subset $\mathcal{B}(C_1, C_2)$ is small as well.

Lemma 4.22 $|\mathcal{B}(C_1, C_2)| \leq 5$ for every $C_1, C_2 \in \mathcal{C}$.

Proof: Fix $C_1, C_2 \in \mathcal{C}$ and let $\mathcal{B}(C_1, C_2) = \{P_{\tau_1}, \dots, P_{\tau_N}\}$ be sorted according to the time τ_k their segment Q_τ was added to the spanner, for every $k \in \{1, \dots, N\}$ where $N = |\mathcal{B}(C_1, C_2)|$. Let $P_{\tau_k} = P_{s_k, t_k, v_k}$. Let $p_k, q_k \in P_{\tau_k}$ be the endpoints of Q_τ such that p_k is closer to the source s_k , and $C_{v_k}(p_k) = C_1$ and $C_{v_k}(q_k) = C_2$.

Recall that ϕ_{τ_k} is the first vertex of Q_{τ_k} (see Def. 4.15). Let $C_\ell = C_{v_k}(u_\ell)$ be the cluster of u_ℓ for every $u_\ell \in U_{s_k, t_k, v_k}$ (see Def. 4.16). By Obs. 2.3, it holds that $C_\ell \neq C_{\ell'}$ for every $u_\ell, u_{\ell'} \in U_{s_k, t_k, v_k}$. Recall that for every $u_\ell \in U_{s_k, t_k, v_k}$, $P_{s, u_\ell, v} = P_{s, t, v}[s, u_\ell]$. Since $\text{LastE}(P_{s, u_\ell, v}) \notin E^{\text{local}}$, it holds that $v \notin \{z_1(u_\ell), \text{LCA}(s, C_1(u_\ell))\}$. Combining that with the fact that $u_\ell \in U_{s_k, t_k, v_k}$ did not satisfy property (N2) (see Def. 4.15 and Def. 4.16), we conclude that $v_k \in V^{\text{far}}(s_k, u_\ell)$. Since $q_k \in U_{s_k, t_k, v_k} \cup \{t_k\}$, using Lemma 4.18(d), it holds that

$$v_k \in V_f(C_2) \text{ for every } k \in \{1, \dots, N\}, \quad (36)$$

and by Lemma 4.18(c),

$$P_{\tau_k}[p_k, q_k] \subseteq Q_{\tau_k}[p_k, q_k] \subseteq G \setminus V_f(C_2). \quad (37)$$

Since $C_1 = C_{v_k}(p_k)$ and $C_2 = C_{v_k}(q_k)$, for every $k \in \{1, \dots, N\}$, it holds that $z_{v_1}(p_1) = \dots = z_{v_N}(p_N)$ and also $z_{v_1}(q_1) = \dots = z_{v_N}(q_N)$. Hence, letting $z_1 = z_{v_1}(p_1)$ and $z_2 = z_{v_1}(q_1)$, it holds that

$$z_1, z_2 \notin \{v_1, \dots, v_N\}. \quad (38)$$

Denote

$$X_k = \text{dist}(p_k, q_k, G_{\tau_{k+1}} \setminus V_f(C_2)).$$

We now show that $X_k < X_{k-1}$ for every $k \in \{2, \dots, N\}$.

Each time a path segment Q_{τ_k} is purchased at time τ_k , it implies that

$$X_k \leq \text{dist}(p_k, q_k, P_{\tau_k} \setminus V_f(C_2)) \quad (39)$$

$$= \text{dist}(p_k, q_k, P_{\tau_k}) \quad (40)$$

$$< \text{dist}(C_1, C_2, G_{\tau_k} \setminus V_f(C_2)) \quad (41)$$

$$\leq X_{k-1}, \quad (42)$$

where Eq. (39) follows by the fact that $P_{\tau_k}[p_k, q_k] \subseteq Q_{\tau_k} \subseteq G_{\tau_{k+1}}$, Eq. (40) follows by Eq. (37), Eq. (41) follows by the fact that Q_{τ_k} was bought and by Eqs. (28) and (29), and Eq. (42) follows by the fact that $p_{k-1} \in C_1$ and $q_{k-1} \in C_2$.

Therefore, we have that

$$X_N \leq X_1 - (N - 1). \quad (43)$$

Conversely, we have that

$$X_N \geq \text{dist}(p_N, q_N, G \setminus \{v_1, \dots, v_N\}) \quad (44)$$

$$\geq \text{dist}(p_1, q_1, G \setminus \{v_1, \dots, v_N\}) - 4 \quad (45)$$

$$= \text{dist}(p_1, q_1, P_{\tau_1}) - 4 = \text{dist}(p_1, q_1, P_{\tau_1} \setminus V_f(C_2)) - 4 \quad (46)$$

$$\geq X_1 - 4, \quad (47)$$

where Eq. (44) follows as $G_{\tau_{N+1}} \subseteq G$ and by Eq. (36), $\{v_1, \dots, v_N\} \subseteq V_f(C_2)$. To see Eq. (45), note that $p_1, p_N \in C_1$ and $q_1, q_N \in C_2$ and by Obs. 2.3(2) the diameter of the cluster is 2. It remains to show that the intracluster paths $R_1 = [p_1, z_1, p_N]$, $R_2 = [q_1, z_2, q_N]$ exist in the surviving graph $G \setminus \{v_1, \dots, v_N\}$. This holds since by Eq. (38), $z_1, z_2 \notin \{v_1, \dots, v_N\}$, and by Eq. (36) and (37). Eq. (46) follows by Eq. (37). Finally, Eq. (47) follows by the fact that $P_{\tau_1}[p_1, q_1] \subseteq Q_{\tau_1}$ was added at step τ_1 , hence $P_{\tau_1}[p_1, q_1] \subseteq G_{\tau_2}$. By combining with Eq. (43), we get that $N \leq 5$. The lemma follows. \blacksquare

claim the following.

Lemma 4.23 $|E(H_8(S))| = O(|S| \cdot n + n^2/|S|^2)$.

Finally, since for every path $P \in \mathcal{B}$, it holds that $\text{Cost}(P) \leq 4 \cdot \text{Val}(P)$, we get that

$$\begin{aligned} |E(G_{\tau'}) \setminus E(G_0)| &= \sum_{P \in \mathcal{B}} \text{Cost}(P) \leq 4 \sum_{P \in \mathcal{B}} \text{Val}(P) \leq 4 \sum_{C_1, C_2 \in \mathcal{C}} |\mathcal{B}(C_1, C_2)| \\ &\leq O(|\mathcal{C}|^2) = O((n/|S|)^2). \end{aligned}$$

where the last equality follows by the fact that there are $|\mathcal{C}| = |Z| = O(n/|S|)$ clusters. The claim follows. \blacksquare

Additive stretch 6 (for all pairs). Set $\Delta = \sqrt{n}$ and let $Z \subseteq S$ be a collection of \sqrt{n} cluster centers and G_Δ be the corresponding clustering graph. By Obs. 2.5, $H = H_4(Z) \cup G_\Delta$ is a 6 additive FT-spanner. By Theorem 4.1, $|E(H_4(Z))| = O(n^{3/2})$ and by Obs. 2.3(4), $|E(G_\Delta)| = O(n^{3/2})$, hence Thm. 1.3 follows.

The achieved bounds should be compared with the single source additive FT-spanner $H_4(\{s\})$ of [21] and the (all-pairs, non FT) 6-additive spanner, both with $O(n^{4/3})$ edges.

5 Discussion

The current work provides the first upper bounds for handling vertex faults in additive spanner constructions. Unfortunately, currently there is no lower bound for this setting. Specifically, even for the basic case of additive stretch 2, there is no lower bound of $\Omega(n^{3/2+\epsilon})$ for any $\epsilon > 0$. In [21], there is a lower bound construction for the *single source* FT-additive spanners. Extending this lower bound for (all pairs) FT-additive spanners remains open. Turning to the upper bound side, there are two main challenges. The first involves the construction of FT structures with *nonconstant* additive stretch (e.g., polylogarithmic or sublinear in the distances). The second involves the extension of the presented constructions to support the case of *multiple* vertex faults.

References

- [1] 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.
- [2] Mathias Bæk Tejs Knudsen. Additive Spanners: A Simple Construction. In Proc. *SWAT*, 277–281, 2014.
- [3] S. Baswana and S. Sen. Approximate distance oracles for unweighted graphs in expected $O(n^2)$ time. *ACM Trans. Algorithms*, 2(4):557–577, 2006.
- [4] S. Baswana and N. Khanna. Approximate Shortest Paths Avoiding a Failed Vertex: Optimal Size Data Structures for Unweighted Graph. In Proc. *STACS*, 513–524, 2010.
- [5] S. Baswana, T. Kavitha, K. Mehlhorn, and S. Pettie. Additive spanners and (α, β) -spanners. *ACM Trans. Algo.* 7, A.5, 2010.
- [6] A. Bernstein. A nearly optimal algorithm for approximating replacement paths and k shortest simple paths in general graphs. In Proc. *SODA*, 2010.
- [7] A. Bernstein and D. Karger. A nearly optimal oracle for avoiding failed vertices and edges. In Proc. *STOC*, 101–110, 2009.
- [8] D. Bilò, L. Gualà, S. Leucci, G. Proietti. Fault-Tolerant Approximate Shortest-Path Trees. In Proc. *ESA*, 137–148, 2014.
- [9] G. Braunschvig, S. Chechik and D. Peleg. Fault tolerant additive spanners. In Proc. *WG*, 206–214, 2012.
- [10] S. Chechik, M. Langberg, D. Peleg, and L. Roditty. f -sensitivity distance oracles and routing schemes. *Algorithmica* **63**, (2012), 861–882.
- [11] S. Chechik, M. Langberg, D. Peleg, and L. Roditty. Fault-tolerant spanners for general graphs. In Proc. *STOC*, 435–444, 2009.
- [12] S. Chechik. New Additive Spanners. In Proc. *SODA*, 2013.
- [13] M. Cygan, F. Grandoni and T. Kavitha. On Pairwise Spanners. In Proc. *STACS*, 209–220, 2013.
- [14] M. Dinitz and R. Krauthgamer. Fault-tolerant spanners: better and simpler. In Proc. *PODC*, 2011, 169–178.
- [15] D. Dor, S. Halperin, and U. Zwick. All-pairs almost shortest paths. *SIAM J. Computing*, 29(5):1740–1759, 2000.
- [16] M. Elkin and D. Peleg. $(1 + \varepsilon, \beta)$ -Spanner Constructions for General Graphs. *SIAM J. Computing*, 33(3): 608–631, 2004.
- [17] F. Grandoni and V.V Williams. Improved Distance Sensitivity Oracles via Fast Single-Source Replacement Paths. In Proc. *FOCS*, 2012.
- [18] T. Kavitha and N.M. Varma. Small Stretch Pairwise Spanners. In Proc. *ICALP*, 2013.
- [19] T. Kavitha. New Pairwise Spanners. In Proc. *STACS*, 2015.
- [20] M. Parter and D. Peleg. Sparse Fault-tolerant BFS trees. In Proc. *ESA*, 2013.
- [21] M. Parter and D. Peleg. Fault Tolerant Approximate BFS Structures. In Proc. *SODA*, 2014.
- [22] S. Pettie. Low Distortion Spanners. *TALG*, 6(1), 2009.
- [23] L. Roditty and U. Zwick. Replacement paths and k simple shortest paths in unweighted directed graphs. *ACM Trans. Algorithms*, 2012.

- [24] M. Thorup and U. Zwick. Spanners and emulators with sublinear distance errors. In Proc. *SODA*, 802–809, 2006.
- [25] O. Weimann and R. Yuster. Replacement paths via fast matrix multiplication. In Proc. *FOCS*, 2010.
- [26] D.P Woodruff. Lower bounds for additive spanners, emulators, and more. In Proc. *FOCS*, 389–398, 2006.
- [27] D.P Woodruff. Additive spanners in nearly quadratic time. In Proc. *ICALP*, 463–474, 2010.