# Vertex Fault Tolerant Additive Spanners 

Merav Parter *

November 19, 2018


#### 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 $\operatorname{dist}(s, t, H \backslash\{v\}) \leq \operatorname{dist}(s, t, G \backslash\{v\})+\beta$ 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 [8]. 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 $\widetilde{O}\left(n^{5 / 3}\right)$ edges. Our second result is an FT-spanner with additive stretch 6 and $\widetilde{O}\left(n^{3 / 2}\right)$ edges. The construction algorithm consists of two main components: (a) constructing an FT-clustering graph and (b) applying a modified path-buying procedure suitably adopted 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 $(\alpha, \beta)$-spanner $H$ of an unweighted undirected graph $G$ is a spanning subgraph satisfying for every pair of vertices $s, t \in V$ that $\operatorname{dist}(s, t, H) \leq \alpha \cdot \operatorname{dist}(s, t, G)+\beta$. 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 $O\left(n^{3 / 2}\right), \widetilde{O}\left(n^{7 / 5}\right)$ and $O\left(n^{4 / 3}\right)$ edges respectively [1, 2, 5, 11, 14, 15]. This paper considers a network $G$ that may suffer single vertex failure events, and looks for fault tolerant additive spanners that maintain their additive stretch guarantee under failures. Formally, a subgraph $H \subseteq$ $G$ is a $\beta$-additive FT-spanner iff for every $(s, t) \in V \times V$ and for every failing vertex $v \in V$, $\operatorname{dist}(s, t, H \backslash\{v\}) \leq \operatorname{dist}(s, t, G \backslash\{v\})+\beta$. 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

[^0]simply lease a 2-additive spanner $H_{0}$ of the graph with $\Theta\left(n^{3 / 2}\right)$ edges. However, if one of the vertices in the graph fails, some $s-t$ routes in $H_{0} \backslash\{v\}$ might be significantly longer than the corresponding route in the surviving graph $G \backslash\{v\}$. Moreover, $s$ and $t$ are not even guaranteed to be connected in $H_{0} \backslash\{v\}$. 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 at el. [10] for the case of multiplicative stretch. Specifically, [10] presented algorithms for constructing an $f$ vertex fault tolerant spanner with multiplicative stretch $(2 k-1)$ and $O\left(f^{2} k^{f+1} \cdot n^{1+1 / k} \log ^{1-1 / k} n\right)$ edges. Dinitz and Krauthgamer presented in [13, a randomized construction attaining an improved tradeoff for vertex fault-tolerant spanners, namely, $f$-vertex fault tolerant $k$-spanner with $\widetilde{O}\left(f^{2}\right.$. $\left.n^{1+2 /(k+1)}\right)$ edges. Constructions of fault-tolerant spanners with additive stretch resilient to edge failures were recently given by Braunschvig at el. [8]. 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 $\beta(f)=O(f(\alpha+\beta))$. In particular, fixing the number of $H$ edges to be $O\left(n^{3 / 2}\right)$ and the number of faults to $f=1$ yields an additive stretch of 14 . Hence, in particular, there is no construction for additive stretch $<14$ and $o\left(n^{2}\right)$ edges. In addition, note that these structures are resilient only to edge failures as the techniques of $[8]$ 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 [8].

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, 16, 19, 21]) focuses on time-efficient computation of the 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 [17], 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 $G \backslash\{v\}$ for every failing vertex $v$, for every source $s \in S$.

It is shown in [17] 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 $O\left(\sqrt{|S|} \cdot n^{3 / 2}\right)$ 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 $\Omega\left(\sqrt{|S|} \cdot n^{3 / 2}\right)$. 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 $\widetilde{P}_{s, t, v}$ (i.e., such that $\left.\left|\widetilde{P}_{s, t, v}\right| \leq \alpha \cdot\left|P_{s, t, v}\right|\right)$ has been studied in [3, 9, 6]. In particular its single source variant has been studied in [4, 18]. 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 $\beta$ 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 $\widetilde{O}\left(\max \left\{n \cdot|S|,(n /|S|)^{3}\right\}\right)$ edges and the second construction guarantees additive stretch 8 with $\widetilde{O}\left(\max \left\{n \cdot|S|,(n /|S|)^{2}\right\}\right)$. As a direct consequence of these constructions, we get an additive FT-spanner with stretch 6 and $\widetilde{O}\left(n^{3 / 2}\right)$ edges and an additive sourcewise FT-spanner with additive stretch 8 and $\widetilde{O}\left(n^{4 / 3}\right)$ for at most $\widetilde{O}\left(n^{1 / 3}\right)$ sources.

Our constructions employ a modification of the path-buying strategy, which was originally devised in [5] to provide 6 -additive spanners with $O\left(n^{4 / 3}\right)$ edges. Recently, the path-buying strategy was employed in the context of pairwise spanners, where the objective is to construct a subgraph $H \subseteq G$ that satisfies the bounded additive stretch requirement only for a subset of pairs [12]. The high-level idea of this procedure 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 every vertex with a sufficiently high degree (termed hereafter a heavy vertex) 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 heavy vertex $t$ (in our constructions, all heavy vertices are clustered). 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 $t$ whose failing vertex $v$ does not appear on its $\pi(s, x)$ 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 events occur 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 candidate to be bought by the path-buying procedure. Essentially, the structure of these paths and the cost and value functions assigned to them would guarantee that the resulting structure is sparse, and in addition, that paths that were not bought have an alternative safe path in the surviving part of the structure.

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

The main technical contribution of our algorithms is in adapting the path-buying strategy to the vertex failure setting. Such an adaptation has been initiated in [18] 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, [18] achieves a construction of single source additive spanner with $O\left(n^{4 / 3}\right)$ 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 $O\left(n^{1 / 3}\right)$ sources, additive stretch 8 and $\widetilde{O}\left(n^{4 / 3}\right)$ 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 $\widetilde{O}\left(n^{5 / 3}\right)$ such that $\operatorname{dist}(s, t, H \backslash\{v\}) \leq \operatorname{dist}(s, t, G \backslash$ $\{v\})+2$ for every $s, t, v \in V$.

Theorem 1.2 (6-additive FT-spanner) For every $n$-vertex graph $G=(V, E)$, there exists a (polynomially constructible) subgraph $H \subseteq G$ of size $\widetilde{O}\left(n^{3 / 2}\right)$ such that $\operatorname{dist}(s, t, H \backslash\{v\}) \leq \operatorname{dist}(s, t, G \backslash$ $\{v\})+6$ for every $s, t, v \in V$.

Theorem 1.3 (8-additive sourcewise FT-spanner) For every n-vertex graph $G=(V, E)$ and a subset of sources $S \subset V$ where $|S|=\widetilde{O}\left(n^{1 / 3}\right)$, there exists a (polynomially constructible) subgraph $H \subseteq G$ of size $\widetilde{O}\left(n^{4 / 3}\right)$ such that $\operatorname{dist}(s, t, H \backslash\{v\}) \leq \operatorname{dist}(s, t, G \backslash\{v\})+8$ for every $s \in S$ and $t, v \in V$.

## 2 Preliminaries

Notation. Given a graph $G=(V, E)$, a vertex pair $s, t$ and an edge weight function $W: E(G) \rightarrow$ $\mathbb{R}^{+}$, let $S P(s, t, G, W)$ 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 paths ${ }^{1}$. Hence, $S P\left(s, t, G^{\prime}, W\right)$ contains a single path for every $s, t \in V$ and for every subgraph $G^{\prime} \subseteq G$, we override notation and let $S P(s, t, G, W)$ be the unique $s-t$ path in $G$ according to $W$. When the shortest-path are computed in $G$, let $\pi(s, t)=S P(s, t, G, W)$. To avoid cumbersome notation, we may omit $W$ and simply refer to $\pi(s, t)=S P(s, t, G, W)$. For a subgraph $G^{\prime} \subseteq G$, let $V\left(G^{\prime}\right)$ (resp., $E\left(G^{\prime}\right)$ ) denote the vertex set (resp. edge set) in $G^{\prime}$.

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 $V^{\prime} \in V$, let $T\left(t, V^{\prime}\right)=\bigcup_{u \in V^{\prime}} \pi(u, t)$ be the union of all $\{t\} \times V^{\prime}$ shortest-paths (by the uniqueness of $W, T\left(t, V^{\prime}\right)$ is a subtree of $T_{0}(t)$ ). Let $\Gamma(v, G)$ 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 $\operatorname{deg}(v, G)=|E(v, G)|$ 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 $\Delta$-heavy if $\operatorname{deg}(v, G) \geq \Delta$, otherwise it is $\Delta$-light. When $\Delta$ 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 $V_{\Delta}=\{v \mid \operatorname{deg}(v, G) \geq \Delta\}$ be the set of $\Delta$-heavy vertices in $G$. (Throughout, we sometimes simplify notation by omitting parameters which are clear from the context.) For a subgraph $G^{\prime}=\left(V^{\prime}, E^{\prime}\right) \subseteq G$ (where $V^{\prime} \subseteq V$ and $E^{\prime} \subseteq E$ ) and a pair of vertices $u, v \in V$, let $\operatorname{dist}\left(u, v, G^{\prime}\right)$ denote the shortest-path distance in edges between $u$ and $v$ in $G^{\prime}$. For a path $P=\left[v_{1}, \ldots, v_{k}\right]$, let LastE $(P)$ be the last edge of $P$, let $|P|$ denote its length and let $P\left[v_{i}, v_{j}\right]$ be the subpath of $P$ from $v_{i}$ to $v_{j}$. For paths $P_{1}$ and $P_{2}$, denote by $P_{1} \circ P_{2}$

[^1]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 $\operatorname{dist}(s, e, P)$ be the distance (in edges) between $s$ and $y$ on $P$. In addition, for an edge $e=(x, y) \in T_{0}(s)$, define $\operatorname{dist}(s, e)=i$ if $\operatorname{dist}(s, x, G)=i-1$ and $\operatorname{dist}(s, y, G)=i$. 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 ( $(\alpha, \beta, S)$-AMBFS FT-spanners) A subgraph $H \subseteq G$ is an ( $\alpha, \beta, S$ )-FT-AMBFS (approximate multi-BFS) structure with respect to $S$ if for every $(s, t) \in S \times V$ and every $v \in V$, $\operatorname{dist}(s, t, H \backslash\{v\}) \leq \alpha \cdot \operatorname{dist}(s, t, G \backslash\{v\})+\beta$.

Definition $2.2((\alpha, \beta)$ FT-spanners) A subgraph $H \subseteq G$ is an $(\alpha, \beta)$ FT-spanner if it is an $(\alpha, \beta, V)$-FT-AMBFS structure for $G$ with respect to $V$.

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

FT-Clustering Graph. A subset $Z \subseteq V$ is an $F T$-center set for $V$ if every $\Delta$-heavy vertex $v$ has at least two neighbors in $Z$, i.e., $|\Gamma(v, G) \cap Z| \geq 2$. For every heavy vertex $v \in V_{\Delta}$, let $Z(v)=\left\{z_{1}(v), z_{2}(v)\right\}$ be two arbitrary neighbors of $v$ in $Z$. The clustering graph $G_{\Delta} \subseteq G$ consists of the edges connecting the $\Delta$-heavy vertices $v$ to their two representatives in $Z$ as well as all edges incident to the $\Delta$-light vertices. Formally,

$$
G_{\Delta}=\bigcup_{v \in V_{\Delta}}\left\{\left(v, z_{1}(v)\right),\left(v, z_{2}(v)\right)\right\} \cup \bigcup_{v \notin V_{\Delta}} E(v, G) .
$$

The $\Delta$-heavy vertices are referred hereafter as clustered, hence every missing edge in $G \backslash G_{\Delta}$ 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 $\Delta$-heavy vertices it represents, i.e., $C_{z}=\{z\} \cup\left\{v \in V_{\Delta} \mid z \in Z(v)\right\}$. 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 $G_{\Delta}$ is 2 .

For a failing vertex $v$ and a heavy vertex $t$, let $z_{v}(t) \in Z(t) \backslash\{v\}$ be a cluster center of $t$ in $G \backslash\{v\}$. 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 heavy 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) $\left|E\left(G_{\Delta}\right)\right|=O(\Delta \cdot n)$.
(2) Every missing edge is incident to a clustered vertex in $V_{\Delta}$.
(3) The diameter of every cluster $C_{z}$ is 2 .
(4) There exists an FT-center set $Z \subseteq V$ of size $|Z|=\widetilde{O}(n / \Delta)$.

Obs. 2.3(4) follows by a standard hitting set argument.

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 $P_{s, t, v} \in S P(s, t, G \backslash\{v\})$ that does not go through $v$.

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

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

where the second inequality follows by the fact the every vertex $u \in V \backslash\{v\}$ 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 $\pi(s, t)$ at the vertex $b$, it joins $\pi(s, t)$ again only at the final vertex $t$. It is shown in [17] 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 ( $\beta, S$ )-additive FT-spanners. Essentially, a key section 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 $\pi\left(s^{\prime}, t\right) \backslash\{t\}$ for some other source $s^{\prime} \in S$.

The basic building block. Our constructions of $\beta$-additive FT-spanners, for $\beta \geq 2$, consist of the following two building blocks: (1) an FT-clustering graph $G_{\Delta}$ for some parameter $\Delta$, and (2) an ( $\beta-2, Z$ )-additive FT-spanner 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 $F T$-center set for $V_{\Delta}$. Then $H$ is an $\beta$ additive FT-spanner.

Proof: Consider vertices $u_{1}, u_{2}, u_{3} \in V$. Let $P \in S P\left(u_{1}, u_{2}, G \backslash\left\{u_{3}\right\}\right)$ be the $u_{1}-u_{2}$ replacement path in $G \backslash\left\{u_{3}\right\}$ and let $(x, y)$ be the last missing edge on $P \backslash H$ (i.e., closest to $u_{2}$ ). Since $G_{\Delta} \subseteq H$, by Obs. 2.3(2), $y$ is a clustered vertex. Let $z=z_{u_{3}}(y)$ be the cluster center of $y$ in $G \backslash\left\{u_{3}\right\}$, and consider the following $u_{1}-u_{2}$ path $P_{3}=P_{1} \circ P_{2}$ where $P_{1} \in S P\left(u_{1}, z, H \backslash\left\{u_{3}\right\}\right)$ and $P_{2}=(z, y) \circ P\left[y, u_{2}\right]$. Clearly, $P_{3} \subseteq H \backslash\left\{u_{3}\right\}$, so it remains to bound its length. Since $H_{\beta-2}(Z) \subseteq H$, it holds that $\left|P_{1}\right| \leq \operatorname{dist}\left(u_{1}, z, G \backslash\left\{u_{3}\right\}\right)+\beta-2$. Hence,

$$
\begin{aligned}
\operatorname{dist}\left(u_{1}, u_{2}, H \backslash\left\{u_{3}\right\}\right) & \leq\left|P_{3}\right|=\left|P_{1}\right|+\left|P_{2}\right| \\
& \leq \operatorname{dist}\left(u_{1}, z, G \backslash\left\{u_{3}\right\}\right)+\beta-2+\operatorname{dist}\left(y, u_{2}, G \backslash\left\{u_{3}\right\}\right) \\
& \leq \operatorname{dist}\left(u_{1}, y, G \backslash\left\{u_{3}\right\}\right)+\operatorname{dist}\left(y, u_{2}, G \backslash\left\{u_{3}\right\}\right)+\beta \\
& \leq|P|+\beta=\operatorname{dist}\left(u_{1}, u_{2}, G \backslash\left\{u_{3}\right\}\right)+\beta,
\end{aligned}
$$

where the second inequality follows by the triangle inequality using the fact that the edge $(z, y)$ exists in $H \backslash\left\{u_{3}\right\}$. The lemma follows.

## 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 [17.

Fact 3.1 ([17]) There exists a polynomial time algorithm that for every n-vertex graph $G=(V, E)$ and 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 $O\left(\sqrt{|S|} \cdot n^{3 / 2}\right)$ edges.

Set $\Delta=\left\lceil n^{2 / 3}\right\rceil$ and let $Z$ be an FT-center set for $V_{\Delta}$ as given by Obs. 2.3(4). 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 case (where $|S|=1$ ) is considered in [18], which provides a construction of a single source FT-spanner ${ }^{2}$ with $O\left(n^{4 / 3}\right)$ edges and additive stretch 4. The current construction increases the stretch to 8 to provide a bounded stretch for $\widetilde{O}\left(n^{1 / 3}\right)$ sources with the same order of edges, $\widetilde{O}\left(n^{4 / 3}\right)$.

### 4.1 Sourcewise spanner with additive stretch 4

Lemma 4.1 There exists a subgraph $H_{4}(S) \subseteq G$ with $\widetilde{O}\left(\max \left\{|S| \cdot n,(n /|S|)^{3}\right\}\right)$ edges satisfying $\operatorname{dist}\left(s, t, H_{4}(S) \backslash\{v\}\right) \leq \operatorname{dist}(s, t, G \backslash\{v\})+4$ for every $(s, t) \in S \times V$ and $v \in V$.

The following notation is useful in our context. Let $\mathcal{C}=\left\{C_{z} \mid z \in Z\right\}$ 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 $\operatorname{LCA}\left(s, C_{z}\right)$ be the least common ancestor (LCA) of the cluster vertices of $C_{z}$ in the BFS tree $T_{0}(s)$ rooted at $s$. Let $\pi\left(s, C_{z}\right)$ be the path connecting $s$ and LCA $\left(s, C_{z}\right)$ 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 $P_{s, t, v}=S P(s, t, G \backslash\{v\}, W)$.

Step (1): Clustering. Set $\Delta=|S|$ and let $Z \subseteq V$ be an FT-center set of size $\widetilde{O}(n / \Delta)$ (by Obs. 2.3(4) such set exists). Let $\mathcal{C}=\left\{C_{z} \mid z \in Z\right\}$ be the collection of $|Z|$ clusters. For a heavy 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.

[^2]Step (2): Shortest-path segmentation. For every $(s, t) \in S \times V_{\Delta}$, the algorithm uses the first cluster of $t, C_{1}(t)$, to segment the path $\pi(s, t)$. Define

$$
\pi^{f a r}(s, t)=\pi(s, \ell(s, t)) \backslash\{\ell(s, t)\} \text { and } \pi^{\text {near }}(s, t)=\pi(\ell(s, t), t) \backslash\{\ell(s, t)\},
$$

where $\ell(s, t)=\operatorname{LCA}\left(s, C_{1}(t)\right)$ is the LCA of the cluster $C_{1}(t)$ in the tree $T_{0}(s)$. Hence, $\pi(s, t)=$ $\pi^{f a r}(s, t) \circ \ell(s, t) \circ \pi^{\text {near }}(s, t)$. The algorithm handles separately vertex faults in the near and far segments. Let $V^{\text {near }}(s, t)=V\left(\pi^{\text {near }}(s, t)\right)$ and $V^{f a r}(s, t)=V\left(\pi^{f a r}(s, t)\right)$.

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

$$
E^{l o c a l}(t)=\left\{\operatorname{LastE}\left(P_{s, t, v}\right) \mid s \in S, v \in\left\{z_{1}(t), \operatorname{LCA}\left(s, C_{1}(t)\right)\right\}\right\} \text { and } E^{l o c a l}=\bigcup_{t \in V_{\Delta}} E^{l o c a l}(t)
$$

be the last edges of replacement-paths protecting against the failure of the primary cluster center $z_{1}(t)$ and the least common ancestor LCA $\left(s, C_{1}(t)\right)$.

Step (4): Handling far vertex faults $V^{f a r}(s, t)$. A replacement path $P_{s, t, v}$ is new-ending if its last edge is not in $\left(T_{0}(S) \cup G_{\Delta}\right)$. For a new-ending path $P_{s, t, v}$, let $b_{s, t, v}$ be the unique divergence point of $P_{s, t, v}$ from $\pi(s, t)$ (in the analysis we show that such point exists). Let $D_{s, t, v}=P_{s, t, v}\left[b_{s, t, v}, t\right]$ denote the detour segment and let $D_{s, t, v}^{+}=D_{s, t, v} \backslash\left\{b_{s, t, v}\right\}$ denote the detour segment excluding the divergence point. For every clustered vertex $t$, let $\mathcal{P}^{f a r}(t)$ be the collection of new-ending $s-t$ paths protecting against vertex faults in the far segments, i.e., $\mathcal{P}^{f a r}(t)=\left\{P_{s, t, v} \mid s \in S\right.$, LastE $\left(P_{s, t, v}\right) \notin$ $T_{0}(S)$ and $\left.v \in V^{f a r}(s, t)\right\}$.

The algorithm divides this set into two subsets $\mathcal{P}_{\text {dep }}^{f a r}(t)$ and $\mathcal{P}_{\text {indep }}^{f a r}(t)$ 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 $\pi\left(s^{\prime}, t\right) \backslash\{t\}$ for some $s^{\prime} \in S$, i.e., for a dependent path $P_{s, t, v}$, it holds that

$$
\begin{equation*}
V\left(D_{s, t, v}^{+}\right) \cap V(T(t, S)) \neq\{t\} \tag{1}
\end{equation*}
$$

Otherwise, it is independent. Let

$$
\mathcal{P}_{d e p}^{f a r}(t)=\left\{P_{s, t, v} \in \mathcal{P}^{f a r}(t) \mid s \in S, v \in V^{f a r}(s, t) \text { and } V\left(D_{s, t, v}^{+}\right) \cap V(T(t, S)) \neq\{t\}\right\}
$$

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

Step (4.1): Handling dependent new-ending paths. The algorithm simply takes the last edges $E_{d e p}^{f a r}(t)$ of all dependent replacement paths where $E_{d e p}^{f a r}(t)=\left\{\operatorname{LastE}(P) \mid P \in \mathcal{P}_{d e p}^{f a r}(t)\right\}$. (In the analysis section, we show that the $E_{d e p}^{f a r}(t)$ sets are sparse.) Let $E_{d e p}^{f a r}=\bigcup_{t \in V_{\Delta}} E_{d e p}^{f a r}(t)$.

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

$$
\begin{equation*}
G_{0}=T_{0}(S) \cup G_{\Delta} \cup E^{l o c a l} \cup E_{d e p}^{f a r} \tag{2}
\end{equation*}
$$

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 missing edge on $P_{\tau} \backslash E\left(G_{\tau}\right)$ (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 $z_{v}(u) \in Z(u) \backslash\{v\}$. For every cluster $C$, let $V_{f}(C)$ be the collection of vertices appearing on the paths $\pi(s, C)=\pi(s, \operatorname{LCA}(s, C))$ for every $s \in S$ excluding the vertices of the clusters. That is,

$$
\begin{equation*}
V_{f}(C)=\bigcup_{s \in S} V(\pi(s, C)) \backslash C \tag{3}
\end{equation*}
$$

The path $P_{\tau}$ is added to $G_{\tau}$ resulting in $G_{\tau+1}=G_{\tau} \cup P_{\tau}$, only if

$$
\begin{equation*}
\operatorname{dist}\left(x, t, P_{\tau}\right)<\operatorname{dist}\left(C_{v}(x), C_{v}(t), G_{\tau} \backslash V_{f}\left(C_{v}(t)\right)\right) \tag{4}
\end{equation*}
$$

Let $\tau^{\prime}=\left|\mathcal{P}_{\text {indep }}^{f a r}\right|$ be the total number of independent paths considered to be bought by the algorithm. Then, the algorithm outputs $H_{4}(S)=G_{\tau^{\prime}}$. This completes the description of the algorithm.

Analysis. Throughout the discussion, we consider a $P_{s, t, v}$ paths of clustered vertices $t \in V_{\Delta}$. A path $P_{s, t, v}$ is a new-ending path, if $\operatorname{LastE}\left(P_{s, t, v}\right) \notin G_{0}$ (see Eq. (22). 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 $\operatorname{LastE}\left(P_{s, t, v}[s, u]\right) \notin T_{0}(S)$, it holds that: (a) $v \in \pi(s, u)$. (b) $V\left(P_{s, t, v}\left[b_{s, t, v}, u\right]\right) \cap V(\pi(s, u))=\left\{b_{s, t, v}, u\right\}$.
Proof: Begin with (a). Assume towards contradiction otherwise. By the uniqueness of the weight assignment $W$, we get that $P_{s, t, v}[s, u]=S P(s, u, G \backslash\{v\}, W)=\pi(s, u)$. Leading to contradiction to the fact that LastE $\left(P_{s, t, v}\right)$ 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, $v \in \pi(s, u)$, it also holds that $b_{s, t, v} \in \pi(s, u)$. Assume towards contradiction that there exists an additional point

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

There are two cases to consider (b1) $v \in \pi\left(b_{s, t, v}, w\right)$, in such a case, $v \notin \pi(w, u)$ and hence $\pi(w, u)=$ $S P(w, u, G \backslash\{v\})=P_{s, t, v}[w, u]$, contradiction that $\operatorname{LastE}\left(P_{s, t, v}[s, u]\right) \notin T_{0}(S)$. (b2) $v \in \pi(w, u)$. In such a case, $v \notin \pi\left(b_{s, t, v}, w\right)$ and hence $\pi\left(b_{s, t, v}, w\right)=S P\left(b_{s, t, v}, w, G \backslash\{v\}\right)=P_{s, t, v}\left[b_{s, t, v}, w\right]$, contradiction to the fact the $b_{s, t, v}$ is a divergence point from $\pi(s, t)$. The claim holds.

The next claim 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 $\widetilde{P}_{s, t, v}$ in $T_{0} \cup G_{\Delta}$.

Lemma 4.3 If LastE $\left(P_{s, t, v}\right) \notin G_{0}$ and $v \in \pi^{\text {near }}(s, t)$, then $\operatorname{dist}\left(s, t,\left(G_{0} \cup G_{\Delta}\right) \backslash\{v\}\right) \leq \operatorname{dist}(s, t, G \backslash$ $\{v\})+4$.
Proof: Since $v \in \pi^{\text {near }}(s, t)$, i.e., the failing vertex occurs strictly below $\operatorname{LCA}\left(s, C_{1}(t)\right)$ on $\pi(s, t)$, there exists a vertex $w \in C_{1}(t)$ such that $v \notin \pi(s, w)$ (hence in particular $w \neq v$ ). See Fig. 1 . Since $\operatorname{LastE}\left(P_{s, t, v}\right) \notin E^{\text {local }}$, it holds that $v \neq z_{1}(t)$. Consider the following $s-w$ path $P=$ $\pi(s, w) \circ\left[w, z_{1}(t), t\right]$. Clearly, $P \subseteq\left(T_{0}(S) \cup G_{\Delta}\right) \backslash\{v\}$. By the triangle inequality, as the diameter of the cluster $C_{1}(t)$ is 2 , it holds that

$$
\operatorname{dist}\left(s, t,\left(T_{0}(S) \cup G_{\Delta}\right) \backslash\{v\}\right) \leq \operatorname{dist}(s, w, G)+2 \leq \operatorname{dist}(s, t, G)+4 \leq \operatorname{dist}(s, t, G \backslash\{v\})+4
$$



Figure 1: Handling near vertex faults. Schematic illustration of an approximate replacement path in $\left(T_{0} \cup G_{\Delta}\right) \backslash\{v\}$. Shown is an $\pi(s, t)$ whose failing vertex $v$ occurs strictly below the least common ancestor $\operatorname{LCA}\left(s, C_{1}(t)\right)$. 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 $z=z_{v}(t)$.

The claim follows.
For every new-ending path $P_{s, t, v}$, recall that $D_{s, t, v}^{+}=D_{s, t, v} \backslash\left\{b_{s, t, v}\right\}$. Let $(x, y)$ be the first missing in $P_{s, t, v} \backslash G_{0}$ (where $x$ is closer to $y$ ). The following auxiliary claims are useful.

Lemma 4.4 For every vertex $u \in P_{s, t, v}$ such that $\operatorname{LastE}\left(P_{s, t, v}[s, u]\right) \notin G_{0}$, 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 $P_{s, t, v}[s, u]=$ $P_{s, u, v}=S P(s, u, G \backslash\{v\}, W)$. Since LastE $\left(P_{s, t, v}[s, u]\right) \notin E^{l o c a l}$, 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 4.2, $v \in \pi(s, y)$. Since $b$ appears above $v$ on $\pi(s, t), b$ is mutual to $\pi(s, t)$ 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 $b \neq v \neq y$, $\operatorname{dist}(y, b, G) \geq 2$, and hence also $\operatorname{dist}(y, b, G \backslash\{v\}) \geq 2$, concluding that $b \neq x$. The lemma follows.

Corollary 4.5 Let $t \in V_{\Delta}$. For every $P_{s, t, v} \in \mathcal{P}_{\text {indep }}^{f a r}(t), P_{s, t, v}[x, t] \cap V_{f}\left(C_{v}(t)\right)=\emptyset$ 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), $t \notin V_{f}\left(C_{v}(t)\right)$ and hence $D_{s, t, v}^{+} \cap V_{f}\left(C_{v}(t)\right)=\emptyset$. The corollary follows by combining with Lemma $4.4(\mathrm{~b})$.

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: Let $H=H_{4}(S)$. The proof is by contradiction. Assume that there exists a vertex $s \in S$, a target vertex $t$ and a failing vertex $v$ such that $\operatorname{dist}(s, t, H \backslash\{v\})>\operatorname{dist}(s, t, G \backslash\{v\})+4$. Let

$$
\begin{gathered}
B P=\{(t, v) \mid t \in V, v \in(V \cup\{\emptyset\}) \text { and } \\
\operatorname{dist}(s, t, H \backslash\{v\})>\operatorname{dist}(s, t, G \backslash\{v\}+4\}
\end{gathered}
$$

be the set of "bad pairs" namely, vertex pair $(t, v)$ whose additive stretch in $H$ is greater than 4 with respect to $s \in S$. By the contradictory assumption $B P \neq \emptyset$. Since the BFS tree $T_{0}(s)$ is in $H$, it holds that the failing vertex $v \in \pi(s, t)$ for every pair $(t, v) \in B P$.

For every bad pair $(t, v) \in B P$ define $e_{t, v}$ to be the last missing edge of $P_{s, t, v}=S P(s, t, G \backslash$ $\{v\}, W)$ in $H$. Let $d(t, v)=\operatorname{dist}\left(s, e_{t, v}, P_{s, t, v}\right)$ be the distance of the last missing edge from $s$ on $P_{s, t, v}$. Finally, let $\left(t_{0}, v_{0}\right) \in B P$ be the pair that minimizes $d(t, v)$, and let $e e_{t_{0}, v_{0}}=\left(u, t_{i}\right)$. Note that $e_{t_{0}, v_{0}}$ is the shallowest "deepest missing edge" over all bad pairs $(t, v) \in B P$.

Lemma 4.7 The pair $\left(t_{i}, v_{0}\right) \in B P$.
Proof: Assume towards contradiction that $\left(t_{i}, v_{0}\right) \notin B P$ and let $P^{\prime \prime} \in S P\left(s, t_{i}, H \backslash\left\{v_{0}\right\}\right)$. Hence, since $\left(t_{i}, v_{0}\right) \notin B P$, it holds that

$$
\begin{align*}
\left|P^{\prime \prime}\right| & \leq \operatorname{dist}\left(s, t_{i}, G \backslash\left\{v_{0}\right\}\right)+4  \tag{5}\\
& =\left|P_{s, t_{0}, v_{0}}\left[s, t_{i}\right]\right|+4 .
\end{align*}
$$

We now consider the following $s-t_{0}$ replacement path $Q=P^{\prime \prime} \circ P_{s, t_{0}, v_{0}}\left[t_{i}, t_{0}\right]$. By definition of $\left(t_{i}, v_{0}\right)$ (last missing edge on $\left.P_{s, t_{i}, v_{0}}\right), Q \subseteq H \backslash\left\{v_{0}\right\}$. In addition,

$$
\begin{aligned}
|Q| & =\left|P^{\prime \prime}\right|+\left|P_{t_{0}, v_{0}}\left[t_{i}, t_{0}\right]\right| \leq\left|P_{s, t_{0}, v_{0}}\left[s, t_{i}\right]\right|+4+\left|P_{s, t_{0}, v_{0}}\left(t_{i}, t_{0}\right)\right| \\
& =\left|P_{s, t_{0}, v_{0}}\right|+4=\operatorname{dist}\left(s, t_{0}, G \backslash\left\{v_{0}\right\}\right)+4,
\end{aligned}
$$

where the inequality follows by Eq. (5). This contradicts the fact that $\left(t_{0}, v_{0}\right) \in B P$.
Since $\pi\left(s, t_{i}\right) \subseteq H$, by the fact that $\left(t_{i}, v_{0}\right) \in B P$, we have that the failing vertex $v_{0}$ occurs on the shortest-path $\pi\left(s, t_{i}\right)$. Since the last edge of $P_{s, t_{i}, v_{0}}=P_{s, t_{0}, v_{0}}\left[s, t_{i}\right]$ is missing, by the fact that the clustering graph $G_{\Delta}$ is in $H$, by Obs. 2.3(2), it holds that $t_{i}$ is clustered (i.e., $t_{i} \in V_{\Delta}$ ). By step (2), since $E^{l o c a l} \subseteq H$, it holds that $v_{0} \notin\left\{z_{1}\left(t_{i}\right), \operatorname{LCA}\left(s, C_{1}\left(t_{i}\right)\right)\right\}$. Combining with Lemma 4.3, it holds that $v_{0} \notin V^{\text {near }}\left(s, t_{i}\right)$. Hence, $v_{0} \in V^{f a r}\left(s, t_{i}\right)$. There are further two cases. Case (1) $\bar{P}_{s, t_{i}, v_{0}}$ is dependant. By step (4.1), we then have that $\operatorname{LastE}\left(P_{s, t_{i}, v_{0}}\right) \in E_{d e p}^{f a r}(t)$, contradiction to the fact that LastE $\left(P_{s, t_{i}, v_{0}}\right)$ is missing in $H$.

Consider case(2) where $P_{s, t_{i}, v_{0}}$ is an independent path, i.e., $P_{s, t_{i}, v_{0}} \in \mathcal{P}_{\text {indep }}^{\text {far }}(t)$ and hence it was considered to be bought in the path-buying procedure of Step (4.2). By the fact that the pair $\left(t_{i}, v_{0}\right)$ is a bad pair (i.e., $\left(t_{i}, v_{0}\right) \in B P$ ), we conclude that the algorithm did not buy the path. Let $\tau$ be the iteration at which $P_{\tau}=P_{s, t_{i}, v_{0}}$ was considered to be purchased in the path-buying procedure. Let $G_{\tau}$ be the current spanner in iteration $\tau$. Let $x$ be the vertex incident to the first missing edge on $P_{\tau} \backslash E\left(G_{\tau}\right)$.

Let $C_{0}=C_{v_{0}}\left(t_{i}\right)$ be the cluster of $t_{i}$ in $G_{\Delta} \backslash\left\{v_{0}\right\}$. Since LastE $\left(P_{\tau}\right) \notin E^{\text {local }}$, by Lemma 4.4 (a) $C_{0}=C_{1}\left(t_{i}\right)$ and $z_{v_{0}}\left(t_{i}\right)=z_{1}\left(t_{i}\right)$. By definition,

$$
v_{0} \in \pi^{f a r}\left(s, t_{i}\right)=\pi\left(s, C_{0}\right) \backslash\left\{\operatorname{LCA}\left(s, C_{0}\right)\right\} .
$$

Hence, in particular failing vertex is not in the cluster $C_{0}$, i.e., $v_{0} \notin C_{0}$ and by Eq. (3),

$$
\begin{equation*}
v_{0} \in V_{f}\left(C_{0}\right) . \tag{6}
\end{equation*}
$$

Since $P_{\tau}$ was not bought by the algorithm, by Eq. (4), we have that

$$
\begin{equation*}
\operatorname{dist}\left(C_{v_{0}}(x), C_{0}, G_{\tau} \backslash V_{f}\left(C_{0}\right)\right) \leq \operatorname{dist}\left(x, t_{i}, P_{\tau}\right) \tag{7}
\end{equation*}
$$

Let $w_{1} \in C_{v_{0}}(x)$ and $w_{2} \in C_{0}$ be an arbitrary closest pair in $G_{\tau} \backslash V_{f}\left(C_{0}\right)$ from the clusters $C_{v_{0}}(x)$ and $C_{0}$ respectively satisfying that $\operatorname{dist}\left(w_{1}, w_{2}, G_{\tau} \backslash V_{f}\left(C_{0}\right)\right)=\operatorname{dist}\left(C_{v_{0}}(x), C_{0}, G_{\tau} \backslash V_{f}\left(C_{0}\right)\right)$.

Let $z_{1}$ (resp., $z_{2}$ ) be the cluster center of $C_{v_{0}}(x)$ (resp., $C_{0}$ ). Consider the following $s-t_{i}$ replacement path in $H \backslash\left\{v_{0}\right\}, P_{5}=P_{1} \circ P_{2} \circ P_{3} \circ P_{4}$ where $P_{1}=P_{\tau}[s, x], P_{2}=\left[x, z_{1}, w_{1}\right]$ and $P_{3} \in S P\left(w_{1}, w_{2}, G_{\tau} \backslash V_{f}\left(C_{0}\right)\right)$ and $P_{4}=\left[w_{2}, z_{2}, t_{i}\right]$. For an illustration see Fig. 2. We first claim that $P_{5} \subseteq H \backslash\left\{v_{0}\right\}$. Since $x$ incident to the first missing edge on $P_{\tau}, P_{1}$ is in $H \backslash\left\{v_{0}\right\}$. By Eq. (6), $v_{0} \in V_{f}\left(C_{0}\right)$ and since $w_{1}, w_{2} \subseteq G_{\tau} \backslash V_{f}\left(C_{0}\right)$ it also holds that $w_{1}, w_{2} \neq v_{0}$. Finally note that $G_{\tau}, G_{\Delta} \subseteq H$, hence $P_{2}, P_{4} \subseteq H \backslash\left\{v_{0}\right\}$. We next bound the length of $P_{5}$.

$$
\begin{align*}
\operatorname{dist}\left(s, t_{i}, H \backslash\left\{v_{0}\right\}\right) & \leq P_{5} \leq \operatorname{dist}\left(s, x, G \backslash\left\{v_{0}\right\}\right)+2+\operatorname{dist}\left(w_{1}, w_{2}, G_{\tau} \backslash V_{f}\left(C_{0}\right)\right)+2 \\
& =\operatorname{dist}\left(s, x, G \backslash\left\{v_{0}\right\}\right)+\operatorname{dist}\left(C_{v_{0}}(x), C_{0}, G_{\tau} \backslash V_{f}\left(C_{0}\right)\right)+4 \\
& \leq \operatorname{dist}\left(s, x, G \backslash\left\{v_{0}\right\}\right)+\operatorname{dist}\left(x, t_{i}, P_{\tau}\right)+4  \tag{8}\\
& =\left|P_{s, t_{i}, v}\right|+4=\operatorname{dist}\left(s, t_{i}, G \backslash\left\{v_{0}\right\}\right)+4,
\end{align*}
$$

where Eq. (8) follows by Eq. (7). We end with contradiction to the fact that the pair $\left(t_{i}, v_{0}\right)$ is a bad pair. The claim follows.

Size analysis of $H_{4}(S)$. We proceed with the size analysis.
Lemma 4.8 For every $t \in V_{\Delta},\left|E^{\text {local }}(t)\right|=O(|S|)$, hence $\left|E^{\text {local }}\right|=O(|S| \cdot n)$.
Bounding the number of last edges in $E_{d e p}^{f a r}(t)$. 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 $\pi^{f a r}(s, t)$. To bound the number of edges in $E_{d e p}^{f a r}(t)$, 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 $V^{+}=S \cup\{u \in$ $T(t, S) \mid \operatorname{deg}(u, T(t, S)) \geq 3\}$, be the union of $S$ and the vertices with degree at least 3 in the tree $T(t, S)$. We have that $\left|V^{+}\right|<2|S|$. 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 $\Pi\left(V^{+}\right)=\left\{\pi(x, y) \quad \mid \quad x, y \in V^{+}\right.$and $x, y$ are adjacent $\}$ be the collection of paths between adjacent pairs.

Observation 4.9 (1) $T(t, S)=\Pi\left(V^{+}\right)$. (2) $\Pi\left(V^{+}\right)$consists of at most $2|S|+1$ paths $\pi(x, y)$ (i.e., there are at most $2|S|$ adjacent pairs).


4-additive sourcewise FT-spanner.

Figure 2: Schematic illustration of the path-buying procedure of Alg. Cons4SWSpanner. Shown is an $s-t$ path $P_{\tau}=P_{s, t, v}$ considered to be bought in time $\tau$. The green paths correspond to shortest-paths in $T_{0}(s)$ and the red edges correspond to missing edges on $P_{\tau} \backslash E\left(G_{\tau}\right)$. The first missing edge on $P_{\tau} \backslash E\left(G_{\tau}\right)$ is incident to $x$. If $P_{\tau}$ was not bought, then there exists a short route between a pair of vertices $w_{1}$ and $w_{2}$ belonging to $C_{v}(x)$ and $C_{v}(t)$ (respectively) in $H \backslash\{v\}$.

We now show the following.
Lemma 4.10 For every $t \in V_{\Delta},\left|E_{\text {dep }}^{f a r}(t)\right|=O(|S|)$.
We first claim that every two dependent replacement paths with the same divergence point have the same last edge.

Lemma 4.11 For every two dependent paths $P_{s_{1}, t, v_{1}}, P_{s_{2}, t, v_{2}} \in \mathcal{P}_{\text {dep }}^{f a r}(t)$, if $b_{s_{1}, t, v_{1}}=b_{s_{2}, t, v_{2}}$ then $\operatorname{LastE}\left(P_{s_{1}, t, v_{1}}\right)=\operatorname{LastE}\left(P_{s_{1}, t, v_{2}}\right)$.

Proof: Let $b=b_{s_{1}, t, v_{1}}=b_{s_{2}, t, v_{2}}$. Since $b \in \pi\left(s_{1}, t\right) \cap \pi\left(s_{2}, t\right)$ it holds that $\pi\left(s_{i}, t\right)=\pi\left(s_{i}, b\right) \circ$ $\pi(b, t)$ for $i \in\{1,2\}$. In addition, since $P_{s_{i}, t, v_{i}}\left[s_{i}, b\right]=\pi\left(s_{i}, b\right)$ 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 newending path, hence by the definition of the divergence point $b$ (see Lemma $4.2(\mathrm{~b})$ ), it holds that $V\left(P_{s_{i}, t, v_{i}}[b, t]\right) \cap V(\pi(b, t))=\{b, t\}$ and hence both detours are free from the failing vertices. Hence, $P_{s_{1}, t, v_{1}}[b, t], P_{s_{1}, t, v_{1}}[b, t]=S P\left(b, t, G \backslash\left\{v_{1}, v_{2}\right\}\right)$. We get that $\operatorname{LastE}\left(P_{s_{1}, t, v_{1}}\right)=\operatorname{LastE}\left(P_{s_{2}, t, v_{2}}\right)$ as needed.

Since our goal is to bound the number of last edges of the new ending dependent paths $\mathcal{P}_{d e p}^{f a r}(t)$, to avoid double counting, we now restrict attention to $\mathcal{Q}^{\text {far }}(t)$, a collection of representative paths in $\mathcal{P}_{d e p}^{f a r}(t)$ each ending with a distinct new edge from $E_{d e p}^{f a r}(t)$. Formally, for each new edge $e \in E_{d e p}^{f a r}(t)$, let $P(e)$ be an arbitrary path in $\mathcal{P}_{d e p}^{f a r}(t)$ satisfying that $\operatorname{LastE}(P(e))=e$. Let $\mathcal{Q}^{f a r}(t)=\{P(e), e \in$ $\left.E_{d e p}^{f a r}(t)\right\}$ (hence $\left.\left|\mathcal{Q}^{f a r}(t)\right|=\left|E_{d e p}^{f a r}(t)\right|\right)$. From now on, we aim towards bounding the cardinality of
$\mathcal{Q}^{f a r}(t)$. Let $\mathrm{DP}=\left\{b_{s, t, v} \mid \quad P_{s, t, v} \in \mathcal{Q}^{f a r}(t)\right\}$ be the set of divergence points of the new ending paths in $\mathcal{Q}^{f a r}(t)$. By Lemma 4.11, it holds that in order to bound the cardinality of $\mathcal{P}_{d e p}^{f a r}(t)$, 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 $\mathrm{DP} \backslash V^{+}$.
Lemma $4.12\left|\pi(x, y) \cap\left(\mathrm{DP} \backslash V^{+}\right)\right| \leq 1$ for every $\pi(x, y) \in \Pi\left(V^{+}\right)$.
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 $D_{i}^{+}=D_{i} \backslash\left\{b_{i}\right\}$ 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 on $\pi(y, t)$. Since $D_{1}^{+}$and $D_{2}^{+}$are vertex disjoint with $\pi(y, t) \backslash\{t\}$, 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 $\operatorname{LastE}\left(P_{1}\right), \operatorname{LastE}\left(P_{2}\right)$ could have been avoided. We now focus on $v_{1}$ and show that $v_{1} \in \pi(y, t)$, the exact same argumentation holds for $v_{2}$. Since $P_{1}$ is a new-ending dependent path, by Eq. (1), there exists some source $s_{3} \in s \backslash\left\{s_{1}\right\}$ satisfying that $\left(D_{1}^{+} \cap \pi\left(s_{3}, t\right)\right) \backslash\{t\} \neq \emptyset$. Let $w \in\left(D_{1}^{+} \cap \pi\left(s_{3}, t\right)\right) \backslash\{t\}$ be the first intersection point (closest to $s_{1}$ ). See Fig. 3 for schematic illustration. We first claim that $s_{3}$ is not in the subtree $T(x) \subseteq T(t, S)$ 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 $P_{1}=\pi\left[s_{1}, x\right] \circ \pi\left(x, b_{1}\right) \circ P_{1}\left[b_{1}, w\right] \circ P_{1}[w, t]$. Recall, that since $b_{1} \in \mathrm{DP} \backslash V^{+}$, $b_{1} \neq x$ and also $b_{1} \neq w$. Since $P_{1}[x, w]$ goes through $b_{1}$, by the optimality of $P_{1}$, it holds that

$$
\begin{equation*}
\operatorname{dist}\left(x, w, G \backslash\left\{v_{1}\right\}\right)>\operatorname{dist}\left(b_{1}, w, G \backslash\left\{v_{1}\right\}\right) . \tag{9}
\end{equation*}
$$

On the other hand, the path $\pi\left(s_{3}, t\right)$ has the following form: $\pi\left(s_{3}, t\right)=\pi\left(s_{3}, w\right) \circ \pi(w, x) \circ \pi\left(x, b_{1}\right) \circ$ $\pi\left(b_{1}, t\right)$. Hence, $\pi\left(w, b_{1}\right)$ goes through $x$. Since the failing vertex $v_{1} \in \pi\left(b_{1}, t\right)$ is not in $\pi\left(w, b_{1}\right)$, by the optimality of $\pi\left(w, b_{1}\right)$, we get that $\operatorname{dist}\left(w, b_{1}, G \backslash\left\{v_{1}\right\}\right)>\operatorname{dist}\left(x, w, G \backslash\left\{v_{1}\right\}\right)$, leading to contradiction with Ineq. (9). Hence, we conclude that $s_{3} \notin T(x)$ (in particular this implies that $\left.s_{3} \neq s_{2}\right)$. Note that $\pi(w, t)$ is a segment of $\pi\left(s_{3}, t\right)$ and hence it is contained in the tree $T(t, S)$. Since $P_{1}$ is a new-ending path (i.e., LastE $\left(P_{1}\right) \notin T(t, S)$ ), 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 $\pi\left(s_{3}, t\right)$. To see this, observe that if $\pi(w, t)$ would have been free from the failing vertex $v_{1}$, then it implies that $\pi(w, t)=S P\left(w, t, G \backslash\left\{v_{1}\right\}\right)=P_{1}[w, t]$, contradiction as $\operatorname{LastE}\left(P_{1}\right) \neq \operatorname{LastE}(\pi(w, t))$. Finally, we show that $v_{1} \in \pi(y, t)$. By the above, the failing vertex $v_{1}$ is common to both paths $\pi\left(s_{1}, t\right)$ and $\pi\left(s_{3}, t\right)$, i.e., $v_{1} \in \pi\left(s_{1}, t\right) \cap \pi\left(s_{3}, t\right)$. By the definition of the path $\pi(x, y)$, all its internal vertices $u$ have degree 2 and hence $\left(\pi(x, y) \cap \pi\left(s_{3}, t\right)\right) \backslash\{y\}=\emptyset$, concluding that $v_{1} \in \pi(y, t)$. 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 $\pi(y, t) \backslash\{t\}$, 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 $\operatorname{LastE}\left(P_{1}\right) \neq \operatorname{LastE}\left(P_{2}\right)$, and hence there are two $b_{1}-t$ distinct shortest paths in $G \backslash\left\{v_{1}, v_{2}\right\}$, given by $D_{1}$ and $\pi\left(b_{1}, b_{2}\right) \circ D_{2}$. 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.

By Lemma 4.11 there are at most $\left|V^{+}\right|$replacement paths with divergence point in $V^{+}$. By Lemma 4.12 , there is at most one divergence point on each segment $\pi(x, y)$ of an adjacent pair $(x, y)$. Combining with Obs. 4.9(2), we get $\left|E^{f a r}(t)\right|=\left|\mathcal{Q}^{\text {far }}(t)\right|=O(|S|)$. The lemma follows.


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 newending $s_{1}-t$ dependant path $P_{s_{1}, t, v_{1}}$ with a divergence point $b_{1} \in \pi(x, y)$ intersects with $\pi\left(s_{3}, t\right)$ at the vertex $w \notin\left\{b_{1}, t\right\}$. 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 $v_{1} \in \pi\left(s_{1}, t\right) \cap \pi\left(s_{3}, t\right)$, implying that $v_{1} \in \pi(y, t)$. 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.

We complete the size analysis and proves Lemma 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).
Lemma $4.13\left|H_{4}(S) \backslash G_{0}\right|=\widetilde{O}\left((n /|S|)^{3}\right)$.
Proof: Let $\mathcal{B} \subseteq \mathcal{P}_{\text {indep }}^{f a r}(t)$ 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 $\mathcal{B}\left(C_{1}, C_{2}\right) \subseteq \mathcal{B}$ be the set of paths that were added since they improved the distance of $C_{1}$ and $C_{2}$, that is

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

Clearly, $\mathcal{B}=\bigcup_{C_{1}, C_{2} \in \mathcal{C}} \mathcal{B}\left(C_{1}, C_{2}\right)$. We next use the fact that the diameter of each cluster $C \in \mathcal{C}$ is small, to bound the cardinality of the set $\mathcal{B}\left(C_{1}, C_{2}\right)$.

Lemma $4.14\left|\mathcal{B}\left(C_{1}, C_{2}\right)\right| \leq 5$ for every $C_{1}, C_{2} \in \mathcal{C}$.
Proof: Fix $C_{1}, C_{2} \in \mathcal{C}$ and order the paths of $\mathcal{B}\left(C_{1}, C_{2}\right)$ according to the time step they were added to the spanner $\mathcal{B}\left(C_{1}, C_{2}\right)=\left\{P_{\tau_{1}}, \ldots, P_{\tau_{N}}\right\}$ where $\tau_{1}<\tau_{2}<\ldots<\tau_{N}$ where $N=\left|\mathcal{B}\left(C_{1}, C_{2}\right)\right|$. Since $P_{\tau_{k}} \in \mathcal{P}_{\text {indep }}^{f a r}$, it is a new-ending path, i.e., LastE $\left(P_{\tau_{k}}\right) \notin T_{0}(S)$. Let $P_{\tau_{k}}=P_{s_{k}, t_{k}, v_{k}}$ and $D_{\tau_{k}}=D_{s_{k}, t_{k}, v_{k}}$ denote the detour segment of this path. Let $x_{k}$ be the vertex adjacent to the first missing edge on $P_{\tau_{k}}$. Hence, $C_{1}=C_{v_{k}}\left(x_{k}\right)$ and $C_{2}=C_{v_{k}}\left(t_{k}\right)$ for every $k \in\{1, \ldots, N\}$ and also $V_{f}\left(C_{v_{k}}\left(t_{k}\right)\right)=V_{f}\left(C_{2}\right)$ for every $k \in\{1, \ldots, N\}$. Since $T_{0}(S) \subseteq G_{0}$, the missing edges of $P_{\tau_{k}}$ are restricted to the detour segment $D_{\tau_{k}}$.

In addition, since $P_{\tau_{k}} \in \mathcal{P}_{\text {indep }}^{f a r}$, it holds that the failing vertex $v_{k}$ occurs on the far segment $\pi^{f a r}\left(s_{k}, t_{k}\right)$ and in particular, $v_{k} \notin C_{2}$ (i.e., $v_{k}$ occurs strictly above the least common ancestor $\operatorname{LCA}\left(s_{k}, C_{2}\right)$ and since all cluster members appears on $T_{0}\left(s_{k}\right)$ in the subtree rooted at LCA $\left(s_{k}, C_{2}\right)$, the far segment $\pi^{f a r}\left(s_{k}, t_{k}\right)$ is free from cluster members). We therefore have

$$
\begin{equation*}
\left\{v_{1}, \ldots, v_{N}\right\} \subseteq V_{f}\left(C_{2}\right) \tag{10}
\end{equation*}
$$

Note that $\pi\left(s_{k}, C_{2}\right) \subseteq \pi\left(s_{k}, t_{k}\right)$ and hence it is contained in $T\left(t_{k}, S\right)$ for every $k \in\{1, \ldots, N\}$. By the definition of independent paths (see Eq. (11) for the definition of dependent paths), we have that

$$
\begin{equation*}
D_{s_{k}, t_{k}, v_{k}}^{+} \cap T\left(t_{k}, S\right)=\left\{t_{k}\right\} . \tag{11}
\end{equation*}
$$

Consequently, by Lemma 4.4(b),

$$
\begin{equation*}
P_{\tau_{k}}\left[x_{k}, t_{k}\right] \subseteq D_{\tau_{k}}^{+} \subseteq G \backslash V_{f}\left(C_{2}\right) \tag{12}
\end{equation*}
$$

where the last inclusion holds by the fact that $t_{k} \notin V_{f}\left(C_{2}\right)$ and $V_{f}\left(C_{2}\right) \subseteq \bigcup_{s \in S} \pi\left(s, C_{2}\right) \subseteq$ $\bigcup_{s \in S} \pi(s, t)=T(t, S)$. Let $z_{i} \in Z$ be the cluster center of $C_{i}$ for $i \in\{1,2\}$. We therefore have that $z_{1}=z_{v_{1}}\left(x_{1}\right)=\ldots=z_{v_{N}}\left(x_{N}\right)$ and $z_{2}=z_{v_{1}}\left(t_{1}\right)=\ldots=z_{v_{N}}\left(t_{N}\right)$ and hence $z_{2} \neq v_{k}$ for every $k \in\{1, \ldots, N\}$. Hence,

$$
\begin{equation*}
z_{1}, z_{2} \notin\left\{v_{1}, \ldots, v_{N}\right\} . \tag{13}
\end{equation*}
$$

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

$$
X_{k}=\operatorname{dist}\left(x_{k}, t_{k}, G_{\tau_{k+1}} \backslash V_{f}\left(C_{2}\right)\right)
$$

We now show that $X_{k}<X_{k-1}$ for every $k \in\{2, \ldots N\}$.
Since the path $P_{\tau_{k}}$ is purchased at time $\tau_{k}$, we have that

$$
\begin{align*}
X_{k} & \leq \operatorname{dist}\left(x_{k}, t_{k}, P_{\tau_{k}} \backslash V_{f}\left(C_{2}\right)\right)  \tag{14}\\
& =\operatorname{dist}\left(x_{k}, t_{k}, P_{\tau_{k}}\right)  \tag{15}\\
& <\operatorname{dist}\left(C_{1}, C_{2}, G_{\tau_{k}} \backslash V_{f}\left(C_{2}\right)\right)  \tag{16}\\
& \leq X_{k-1}, \tag{17}
\end{align*}
$$

where Eq. (14) follows by the fact that $P_{\tau_{k}} \subseteq G_{\tau_{k+1}}$, Eq. (15) follows by Eq. (12). Eq. (16) follows by the fact that $P_{\tau_{k}}$ was bought and by Eq. (4), and Eq. (41) follows by the fact that $x_{k-1} \in C_{1}$ and $t_{k-1} \in C_{2}$.

Therefore, we have that

$$
\begin{equation*}
X_{N} \leq X_{1}-(N-1) \tag{18}
\end{equation*}
$$

Conversely, we have that

$$
\begin{align*}
X_{N} & \geq \operatorname{dist}\left(x_{N}, t_{N}, G \backslash\left\{v_{1}, \ldots, v_{N}\right\}\right)  \tag{19}\\
& \geq \operatorname{dist}\left(x_{1}, t_{1}, G \backslash\left\{v_{1}, \ldots, v_{N}\right\}\right)-4  \tag{20}\\
& =\operatorname{dist}\left(x_{1}, t_{1}, P_{\tau_{1}}\right)-4=\operatorname{dist}\left(x_{1}, t_{1}, P_{\tau_{1}} \backslash V_{f}\left(C_{2}\right)\right)-4  \tag{21}\\
& \geq X_{1}-4, \tag{22}
\end{align*}
$$

where Eq. 19) follows as $G_{\tau_{N+1}} \subseteq G$ and by Eq. 10p, $\left\{v_{1}, \ldots, v_{N}\right\} \subseteq V_{f}\left(C_{2}\right)$. To see Eq. (20), we need to prove the existence of the intracluster paths $R_{1}=\left[x_{1}, z_{1}, x_{N}\right]$ and $R_{2}=\left[t_{1}, z_{2}, t_{N}\right]$ in
$G \backslash\left\{v_{1}, \ldots, v_{N}\right\}$ 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 (13), $z_{1}, z_{2} \notin\left\{v_{1}, \ldots, v_{N}\right\}$.

In addition, by Eq. (12), $x_{k}, t_{k} \in G \backslash V_{f}\left(C_{2}\right)$ for every $k \in\{1, \ldots, N\}$ and by Eq. (10), it also holds that $x_{k}, t_{k} \notin\left\{v_{1}, \ldots, v_{N}\right\}$ for every $k \in\{1, \ldots, N\}$. Hence, $R_{1}$ and $R_{2}$ exists in $G \backslash\left\{v_{1}, \ldots, v_{N}\right\}$ and Eq. (20) follows by the triangle inequality. Eq. (21) follows by Eq. 12) and Eq. 10). Finally, Eq. (22) follows by the fact that $P_{\tau_{1}}$ was added at step $\tau_{1}$, hence $P_{\tau_{1}} \subseteq G_{\tau_{2}}$. We get that $N \leq 5$. The claim follows.

By Obs. 2.3(2) and Obs. 2.4, every path $P_{\tau_{k}}$ contains at most $O(n /|\Delta|)=O(n /|S|)$ missing edges in $G \backslash G_{\Delta}$. Hence,

$$
\begin{align*}
\left|E\left(G_{\tau^{\prime}} \backslash G_{0}\right)\right| & =O(n /|S|) \cdot|\mathcal{B}|  \tag{23}\\
& =O(n /|S|) \cdot \sum_{C_{1}, C_{2} \in \mathcal{C}}\left|\mathcal{B}\left(C_{1}, C_{2}\right)\right|  \tag{24}\\
& \leq O(n /|S|) \cdot|\mathcal{C}|^{2}=\widetilde{O}\left((n /|S|)^{3}\right) . \tag{25}
\end{align*}
$$

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

### 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. Cons 4 SWSpanner, 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 4.15 There exists a subgraph $H_{8}(S) \subseteq G$ with $\widetilde{O}\left(\max \left\{|S| \cdot n,(n /|S|)^{2}\right\}\right)$ edges s.t. dist $\left(\right.$ s, $t, H_{8}(S) \backslash$ $\{v\}) \leq \operatorname{dist}(s, t, G \backslash\{v\})+8$ 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 $E^{\text {local }}, E_{d e p}^{f a r}$ be the set of last edges obtained at the end of step (3) and set (4.1) respectively. Let $\mathcal{P}_{\text {indep }}^{f a r}$ 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 $\tau$, we are given $G_{\tau} \subseteq G$ and consider the path $P_{\tau}=P_{s, t, v}$. Let $D_{\tau}=P_{\tau} \backslash \pi(s, t)$ be the detour segment of $P_{\tau}$ (since $\pi(s, t) \subseteq T_{0}(S)$ is in $G_{0}$, all missing edges of $P_{\tau}$ occur on its detour segment).

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

Definition 4.16 Let $\phi_{s, t, v}$ (or $\phi_{\tau}$ for short) be the last vertex on $P_{\tau}$ (closest to $t$ ) satisfying that: $(N 1) \operatorname{LastE}\left(P_{\tau}\left[s, \phi_{\tau}\right]\right) \notin G_{\tau}$, and (N2) $v \in \pi^{\text {near }}\left(s, \phi_{\tau}\right)=\pi\left(\ell, \phi_{\tau}\right) \backslash\{\ell\}$ where $\ell=\operatorname{LCA}\left(s, C_{v}\left(\phi_{\tau}\right)\right)$. If there is no vertex on $P_{\tau}$ that satisfies both (N1) and (N2), then let $\phi_{\tau}$ be the first vertex incident to the first missing edge on $P_{\tau} \backslash E\left(G_{\tau}\right)$ (i.e., such that $P_{\tau}\left[s, \phi_{\tau}\right]$ is the maximal prefix that is contained in $G_{\tau}$ ).

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

$$
\begin{equation*}
\operatorname{dist}\left(\phi_{\tau}, t, P_{\tau}\right)<\operatorname{dist}\left(C_{1}(\tau), C_{2}(\tau), G_{\tau} \backslash V_{f}\left(C_{2}(\tau)\right)\right) \tag{26}
\end{equation*}
$$

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

Definition 4.17 Let $U_{s, t, v}=\left\{u_{3 \ell+1} \mid \quad \ell \in\{0, \ldots,\lfloor(\kappa-1) / 3\rfloor\}\right\} \subseteq Q_{\tau}$ be some representative endpoints of missing edges on $Q_{\tau}$ satisfying that

$$
\operatorname{LastE}\left(Q_{\tau}\left[\phi_{\tau}, u_{\ell}\right]\right) \notin G_{\tau} \text { for every } u_{\ell} \in U_{s, t, v} \text { and } \operatorname{dist}\left(u_{\ell}, u_{\ell^{\prime}}, Q_{\tau}\right) \geq 3
$$

for every $u_{\ell}, u_{\ell^{\prime}} \in U_{s, t, v}$.
Define

$$
\begin{array}{r}
\operatorname{ValSet}_{1}\left(P_{\tau}\right)=\left\{\left(C_{1}(\tau), C_{\ell}\right) \mid C_{\ell}=C_{v}\left(u_{\ell}\right), u_{\ell} \in U_{s, t, v}\right.  \tag{27}\\
\left.\quad \operatorname{and} \operatorname{dist}\left(\phi_{\tau}, u_{\ell}, P_{\tau}\right)<\operatorname{dist}\left(C_{1}(\tau), C_{\ell}, G_{\tau} \backslash V_{f}\left(C_{\ell}\right)\right)\right\}
\end{array}
$$

and

$$
\begin{equation*}
\operatorname{ValSet}_{2}\left(P_{\tau}\right)=\left\{\left(C_{\ell}, C_{2}(\tau)\right) \mid C_{\ell}=C_{v}\left(u_{\ell}\right), u_{\ell} \in U_{s, t, v}\right. \tag{28}
\end{equation*}
$$

$$
\text { and } \left.\operatorname{dist}\left(u_{\ell}, t, P_{\tau}\right)<\operatorname{dist}\left(C_{\ell}, C_{2}(\tau), G_{\tau} \backslash V_{f}\left(C_{2}(\tau)\right)\right)\right\}
$$

Let $\operatorname{ValSet}\left(P_{\tau}\right)=\operatorname{ValSet}_{1}\left(P_{\tau}\right) \cup \operatorname{ValSet}_{2}\left(P_{\tau}\right)$. The subpath $Q_{\tau}$ is added to $G_{\tau}$ resulting in $G_{\tau+1}$ only if

$$
\begin{equation*}
\operatorname{Cost}\left(P_{\tau}\right) \leq 4 \cdot \operatorname{Val}\left(P_{\tau}\right) \tag{29}
\end{equation*}
$$

where $\operatorname{Val}\left(P_{\tau}\right)=\left|\operatorname{ValSet}\left(P_{\tau}\right)\right|$. (Note that when $P_{\tau}$ is cheap, Eq. (29) holds iff Eq. (26) holds.) The output of Alg. Cons8SWSpanner is the subgraph $H_{8}(S)=G_{\tau^{\prime}}$ where $\tau^{\prime}=\left|\mathcal{P}_{\text {indep }}^{f a r}\right|$. This completes the description of the algorithm.

Analysis. Throughout the discussion, a path $P_{s, t, v}$ is a new-ending path, if $\operatorname{LastE}\left(P_{s, t, v}\right) \notin G_{0}$ (see Eq. (2)). Hence, we consider only $P_{s, t, v} \in \mathcal{P}_{\text {indep }}^{f a r}(t)$ paths for clustered vertices $t \in V_{\Delta}$.

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., $D_{s, t, v}=P_{s, t, v}\left[b_{s, t, v}, t\right]$ and $D_{s, t, v}^{+}=D_{s, t, v} \backslash\left\{b_{s, t, v}\right\}$. Let $Q_{s, t, v}=P_{s, t, v}\left[\phi_{s, t, v}, t\right]$ be the path segment that was considered to be bought in step (4.2) (see Def. 4.16).

Observation $4.18 Q_{s, t, v} \subseteq D_{s, t, v}^{+}$.
Proof: Let $x$ be the first vertex incident to a new-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).

Lemma 4.19 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 $C_{k}=C_{v}\left(u_{k}\right)$ it holds that:
(a) $C_{k}=C_{1}\left(u_{k}\right)$.
(b) $V\left(P_{s, t_{i}, v}\left[b_{s, t, v}, u_{k}\right]\right) \cap V\left(T\left(u_{k}, S\right)\right)=\left\{b_{s, t, v}, u_{k}\right\}$.
(c) $Q_{s, t, v}\left[\phi_{s, t, v}, u_{k}\right] \cap V_{f}\left(C_{k}\right)=\{\emptyset\}$.
(d) $v \in V_{f}\left(C_{k}\right)$.

Proof: We begin with (a). By the uniqueness of the weight assignment $W, P_{s, t_{i}, v}\left[s, u_{k}\right]=P_{s, u_{k}, v}=$ $S P\left(s, u_{k}, G \backslash\{v\}, W\right)$. By the uniqueness of the divergence point $b_{s, t, v}$ and in particular by Lemma 4.2 (b),

$$
\begin{equation*}
b_{s, t, v}=b_{s, u_{k}, v} . \tag{30}
\end{equation*}
$$

Since LastE $\left(P_{s, u_{k}, v}\right) \notin E^{l o c a l}$, concluding that $v \neq z_{1}\left(u_{k}\right)$ and hence $z_{v}\left(u_{k}\right)=z_{1}\left(u_{k}\right)$ and (a) holds.
Consider (b). By the definition of the set $U_{s, t, v}$ (see Def. 4.17), it holds LastE $\left(P_{s, u_{k}, v}\right) \notin G_{0}$. Since $u_{k} \in Q_{s, t, v}$ occurs strictly after $\phi_{s, t, v}$, by the definition Def. 4.16 , it holds that $u_{k}$ did not satisfy property (N2). Hence, since LastE $\left(P_{s, u_{k}, v}\right) \notin E^{l o c a l}, v \notin\left\{z_{1}\left(u_{k}\right), \operatorname{LCA}\left(s, C_{1}\left(u_{k}\right)\right)\right\}$ and hence $v \in \pi^{f a r}\left(s, u_{k}\right)$. As LastE $\left(P_{s, u_{k}, v}\right) \notin E_{d e p}^{f a r}\left(u_{k}\right)$, we get that $P_{s, u_{k}, v}$ is a new-ending independent path. By Eq. (1), $V\left(P_{s, u_{k}, v}\left[b_{s, u_{k}, v}, u_{k}\right]\right) \cap V\left(T\left(u_{k}, S\right)\right)=\left\{b_{s, u_{k}, v}, u_{k}\right\}$. Hence (b) holds by Eq. (30).

We now turn to consider claim (c). By Eq. (3), $V_{f}\left(C_{v}\left(u_{k}\right)\right) \subseteq T\left(u_{k}, S\right)$. Since $C_{v}\left(u_{k}\right) \cap$ $V_{f}\left(C_{v}\left(u_{k}\right)\right)=\emptyset$, it holds that $u_{k} \notin V_{f}\left(C_{v}\left(u_{k}\right)\right)$, and hence by combining with claim (a), we get that $P_{s, u_{k}, v}\left[b_{s, u_{k}, v}, u_{k}\right] \cap V_{f}\left(C_{v}\left(u_{k}\right)\right)=\left\{b_{s, u_{k}, v}\right\}$. Since by the proof of Lemma 4.4(b), $\phi_{s, t, v} \neq b_{s, u_{k}, v}$, hence $Q_{s, t, v}\left[\phi_{s, t, v}, u_{k}\right] \cap V_{f}\left(C_{v}\left(u_{k}\right)\right)=\emptyset$.

Consider claim (d). By the above, $v$ occurs on the far segment $\pi\left(s, C_{k}\right) \backslash\left\{\operatorname{LCA}\left(s, C_{k}\right)\right\}$, hence $v \notin C_{k}$. Since $\left(\pi\left(s, C_{k}\right) \backslash C_{k}\right) \subseteq V_{f}\left(C_{k}\right)$, (d) holds.

The next observation is useful in our analysis.
Observation 4.20 If $\phi_{s, t, v}$ satisfies (N1) and (N2), then there exists a vertex $x \in C_{v}\left(\phi_{s, t, v}\right)$ satisfying that $v \notin \pi(s, x)$.

Proof: Let $P_{\tau}=P_{s, t, v}$ and $\phi_{\tau}=\phi_{s, t, v}$. By the uniqueness of the weight assignment $W, P_{\tau}\left[s, \phi_{\tau}\right]=$ $P_{s, \phi_{\tau}, v}=S P\left(s, \phi_{\tau}, G \backslash\{v\}, W\right)$. Since $\phi_{\tau}$ satisfies (N2), it holds that the failing vertex $v$ occurs on $\pi^{\text {near }}\left(s, \phi_{\tau}\right)$, strictly below (i.e., closer to $\phi_{\tau}$ ) the least common ancestor LCA $\left(s, C_{v}\left(\phi_{\tau}\right)\right)$ on $\pi\left(s, \phi_{\tau}\right)$. Hence, there must exist a vertex $x \in C_{v}\left(\phi_{\tau}\right)$ such that $v \notin \pi(s, x)$ (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 LCA $\left.\left(s, C_{v}\left(\phi_{\tau}\right)\right)\right)$.

We proceed by showing correctness.
Theorem $4.21 H_{8}(S)$ is a $(8, S)$ FT-spanner.
Proof: Let $H=H_{8}(S)$. It is required to show that $\operatorname{dist}(s, t, H \backslash\{v\}) \leq \operatorname{dist}(s, t, G \backslash\{v\})+8$ 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 $P_{s, t, v} \in \mathcal{P}_{\text {indep }}^{f a r}(t)$ for $t \in V_{\Delta}$.

Let $\tau$ 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 $\kappa=\operatorname{Cost}\left(P_{\tau}\right)$ denote its cost. Let $\phi_{\tau}$ be as defined in Def. 4.16 and recall that $Q_{\tau}=P_{\tau}\left[\phi_{\tau}, t\right]$ is the candidate suffix to be bought by the procedure. (In particular, $\operatorname{Cost}\left(P_{\tau}\right)$ counts the number of edges on $Q_{\tau} \backslash E\left(G_{\tau}\right)$.)

Case (1): $Q_{\tau}$ was bought. If $\phi_{\tau}$ did not satisfy neither properties (N1), (N2) (or both), then $P_{\tau}\left[s, \phi_{\tau}\right] \subseteq G_{\tau}$. Since $P_{\tau}=P_{\tau}\left[s, \phi_{\tau}\right] \circ Q_{\tau}$ and $Q_{\tau}$ was added to the spanner, we get that $P_{\tau} \subseteq H \backslash\{v\}$.

It remains to consider the complementary case where $\phi_{\tau}$ satisfies both (N1) and (N2). By Obs. 4.20, we get that there exist $x \in C_{v}\left(\phi_{\tau}\right)$ satisfying that $v \notin \pi(s, x)$.

Consider the path $P=\pi(s, x) \circ\left(x, z_{v}\left(\phi_{\tau}\right), \phi_{\tau}\right)$. By definition, $P \subseteq H \backslash\{v\}$ and by the existence of the intracluster path connecting $x$ and $\phi_{\tau}$ in $G \backslash\{v\}$, it holds that $|P|=\operatorname{dist}(s, x, G \backslash\{v\})+2 \leq$ $\operatorname{dist}\left(s, \phi_{\tau}, G \backslash\{v\}\right)+4$. Hence, letting $P^{\prime}=P \circ Q_{\tau}$ (where $\left.Q_{\tau}=P_{\tau}\left[\phi_{\tau}, t\right]\right)$, since $Q_{\tau} \subseteq H \backslash\{v\}$, it holds that $P^{\prime} \subseteq H \backslash\{v\}$ and $\left|P^{\prime}\right| \leq\left|P_{\tau}\right|+4$, as required.

Case (2): $Q_{\tau}$ was not bought. Let $x \in C_{v}\left(\phi_{\tau}\right)$ be defined as follows. If $\phi_{\tau}$ satisfies both properties (N1) and (N2) of Def. 4.16, then using Obs. 4.20, let $x \in C_{v}\left(\phi_{\tau}\right)$ be the vertex satisfying that $v \notin \pi(s, x)$. Otherwise, if $\phi_{\tau}$ did not satisfy (N1) or (N2) (or both), let $x=\phi_{\tau}$. Note that in any case, it holds that $x, \phi_{\tau} \in C_{v}\left(\phi_{\tau}\right)$. We have the following.

Lemma 4.22 $P_{s, x, v} \subseteq H \backslash\{v\}$.
Proof: If $x=\phi_{\tau}$, then it implies that $\phi_{\tau}$ did not satisfy both of the properties (N1,N2). By Def. 4.16, in such a case $\phi_{\tau}$ is the vertex incident to the first missing edge on $P_{s, t, v} \backslash E\left(G_{\tau}\right)$ and hence $P_{s, t, v}[s, x]=P_{s, x, v} \subseteq G_{\tau} \backslash\{v\}$.

Otherwise, if $x \neq \phi_{\tau}$, then $x \in C_{v}\left(\phi_{\tau}\right)$ and by the selection of $x, v \notin \pi(s, x)$. Hence, $P_{s, x, v}=$ $\pi(s, x) \subseteq H \backslash\{v\}$.

Recall that $C_{1}(\tau)=C_{v}\left(\phi_{\tau}\right)$ and $C_{2}(\tau)=C_{v}(t)$. In addition, since $v \in \pi^{f a r}(s, t)$, it holds that $v \in V_{f}\left(C_{2}(\tau)\right)$.

Case (2.1): $P_{\tau}$ is cheap. Since $Q_{\tau}$ was not added, Eq. (26) did not hold and hence

$$
\begin{equation*}
\operatorname{dist}\left(\phi_{\tau}, t, P_{\tau}\right) \geq \operatorname{dist}\left(C_{1}(\tau), C_{2}(\tau), G_{\tau} \backslash V_{f}\left(C_{2}(\tau)\right)\right) \tag{31}
\end{equation*}
$$

Let $w_{1} \in C_{1}(\tau)$ and $w_{2} \in C_{2}(\tau)$ be a closest pair satisfying that $\operatorname{dist}\left(w_{1}, w_{2}, G_{\tau} \backslash V_{f}\left(C_{2}(\tau)\right)\right)=$ $\operatorname{dist}\left(C_{1}(\tau), C_{2}(\tau), G_{\tau} \backslash V_{f}\left(C_{2}(\tau)\right)\right)$. Since the failing vertex $v$ is in $V_{f}\left(C_{2}(\tau)\right)$, both auxiliary vertices $w_{1}$ and $w_{2}$ are in $G \backslash\{v\}$. Consider the following $s-t$ path: $P=P_{0} \circ P_{1} \circ P_{2} \circ P_{3}$ where $P_{0}=P_{s, x, v}$, $P_{1}=\left[x, z_{v}\left(\phi_{\tau}\right), w_{1}\right], P_{2} \in S P\left(w_{1}, w_{2}, G_{\tau} \backslash V_{f}\left(C_{2}(\tau)\right)\right)$, and $P_{3}=\left[w_{2}, z_{v}(t), t\right]$. For an illustration see

Fig. 4. By Lemma 4.22, $P_{0} \subseteq H \backslash\{v\}$. Note that since $x, w_{1} \in C_{v}\left(\phi_{\tau}\right)$, the path $P_{1}$ exists in $H \backslash\{v\}$. Combining with the definitions of the vertices $z_{v}(x), z_{v}(t), w_{1}, w_{2}$, it holds that $P \subseteq H \backslash\{v\}$. So, it remains to bound the length of the path.

$$
\begin{align*}
\operatorname{dist}(s, t, H \backslash\{v\}) & \leq\left|P_{0}\right|+\left|P_{1}\right|+\left|P_{2}\right|+\left|P_{3}\right| \\
& =\operatorname{dist}(s, x, G \backslash\{v\})+\operatorname{dist}\left(w_{1}, w_{2}, G_{\tau} \backslash V_{f}\left(C_{2}(\tau)\right)\right)+4 \\
& \leq \operatorname{dist}\left(s, \phi_{\tau}, G \backslash\{v\}\right) \\
& +\operatorname{dist}\left(w_{1}, w_{2}, G_{\tau} \backslash V_{f}\left(C_{2}(\tau)\right)\right)+6  \tag{32}\\
& \leq \operatorname{dist}\left(s, \phi_{\tau}, G \backslash\{v\}\right)+\operatorname{dist}\left(\phi_{\tau}, t, P_{\tau}\right)+6=\left|P_{\tau}\right|+6, \tag{33}
\end{align*}
$$

where Eq. (32) follows by the fact that $x, \phi_{\tau} \in C_{v}\left(\phi_{\tau}\right)$ and since $G_{\Delta} \subseteq H$, it holds that the intracluster path $R=\left[x, z_{v}\left(\phi_{\tau}\right), \phi_{\tau}\right]$ exists in $G \backslash\{v\}$, Eq. (33) follows by Eq. (31).

Case (2.2): $P_{\tau}$ is costly. Let $U_{s, t, v}=\left\{u_{1}, \ldots, u_{\kappa^{\prime}}\right\} \subseteq Q_{\tau}$ for $\kappa^{\prime}=\lfloor\kappa / 3\rfloor \geq 1$ be as defined by Def. 4.17. 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 $C_{k}=C_{v}\left(u_{k}\right) \in \mathcal{C}$. Hence there are at least $\kappa^{\prime}$ distinct clusters on $Q_{\tau}$.

A cluster $C_{k}=C_{v}\left(u_{k}\right)$ is a contributor if adding $Q_{\tau}$ to $G_{\tau}$ improves either the $C_{1}(\tau)-C_{k}$ distance (i.e., $\left.\left(C_{1}(\tau), C_{k}\right) \in \operatorname{ValSet}_{1}\left(P_{\tau}\right)\right)$ or the $C_{2}(\tau)-C_{k}$ distance (i.e., $\left.\left(C_{k}, C_{2}(\tau)\right) \in \operatorname{ValSet}_{2}\left(P_{\tau}\right)\right)$ 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 $\kappa^{\prime}$ clusters contribute to $\operatorname{Val}\left(P_{\tau}\right)$ (either with $C_{1}(\tau)$ or with $C_{2}(\tau)$ or both). It then holds that $\operatorname{Val}\left(P_{\tau}\right) \geq \kappa^{\prime} \geq \operatorname{Cost}\left(P_{\tau}\right) / 4$. Hence, by Eq. 29], we get a contradiction to the fact that the suffix $Q_{\tau}$ was not added to $G_{\tau}$.

In the other case, there exists at least one neutral cluster $C_{\ell}$ such that

$$
\begin{align*}
\operatorname{dist}\left(C_{1}(\tau), C_{k}, \widehat{H}_{1}\right) & \leq \operatorname{dist}\left(\phi_{\tau}, u_{k}, P_{\tau}\right) \text { and }  \tag{34}\\
\operatorname{dist}\left(C_{k}, C_{2}(\tau), \widehat{H}_{2}\right) & \leq \operatorname{dist}\left(u_{k}, t, P_{\tau}\right)
\end{align*}
$$

where $\widehat{H}_{1}=G_{\tau} \backslash V_{f}\left(C_{k}\right)$ and $\widehat{H}_{2}=G_{\tau} \backslash V_{f}\left(C_{2}(\tau)\right)$. Let $w_{1} \in C_{1}(\tau)$ and $w_{2} \in C_{k}$ be the pair of vertices satisfying $\operatorname{dist}\left(w_{1}, w_{2}, \widehat{H}_{1}\right)=\operatorname{dist}\left(C_{1}(\tau), C_{k}, \widehat{H}_{1}\right)$. In addition, let $y_{1} \in C_{k}$ and $y_{2} \in C_{2}(\tau)$ be the pair satisfying $\operatorname{dist}\left(y_{1}, y_{2}, \widehat{H}_{2}\right)=\operatorname{dist}\left(C_{k}, C_{2}(\tau), \widehat{H}_{2}\right)$.

Let $Q_{1}=\left[x, z_{v}\left(\phi_{\tau}\right), w_{1}\right], Q_{2}=\left[w_{2}, z_{v}\left(u_{k}\right), y_{1}\right]$ and $Q_{3}=\left[y_{2}, z_{v}(t), t\right]$ be the intracluster paths in $C_{1}(\tau), C_{k}$ and $C_{2}(\tau)$ respectively. Note that by definition $x, w_{1} \in C_{v}\left(\phi_{\tau}\right)$.

Since by Lemma 4.19(c), $v \in V_{f}\left(C_{k}\right) \cap V_{f}\left(C_{2}(\tau)\right)$, it also holds that $Q_{1}, Q_{2}, Q_{3} \subseteq H \backslash\{v\}$. Let $P^{\prime}=P_{0} \circ Q_{1} \circ P_{1} \circ Q_{2} \circ P_{2} \circ Q_{3}$ where $P_{0}=P_{s, x, v}, P_{1} \in S P\left(w_{1}, w_{2}, \widehat{H}_{1}\right)$ and $P_{2} \in S P\left(y_{1}, y_{2}, \widehat{H}_{2}\right)$. By Lemma 4.22, $P_{0} \subseteq H \backslash\{v\}$ and by the above explanation, $P^{\prime} \subseteq H \backslash\{v\}$. So, it remains to bound the length of the $s-t$ path $P^{\prime}$.

$$
\begin{aligned}
\operatorname{dist}(s, t, H \backslash\{v\}) & \leq\left|P^{\prime}\right|=\left|P_{0}\right|+\left|P_{1}\right|+\left|P_{2}\right|+6 \\
& =\operatorname{dist}(s, x, G \backslash\{v\})+\operatorname{dist}\left(w_{1}, w_{2}, \widehat{H}_{1}\right)+\operatorname{dist}\left(y_{1}, y_{2}, \widehat{H}_{2}\right)+6 \\
& \leq \operatorname{dist}\left(s, \phi_{\tau}, G \backslash\{v\}\right)+\operatorname{dist}\left(w_{1}, w_{2}, \widehat{H}_{1}\right)+\operatorname{dist}\left(y_{1}, y_{2}, \widehat{H}_{2}\right)+8 \\
& =\operatorname{dist}\left(s, \phi_{\tau}, G \backslash\{v\}\right)+\operatorname{dist}\left(C_{1}(\tau), C_{k}, \widehat{H}_{1}\right) \\
& +\operatorname{dist}\left(C_{k}, C_{2}(\tau), \widehat{H}_{2}\right)+8 \\
& \leq \operatorname{dist}\left(s, \phi_{\tau}, G \backslash\{v\}\right)+\operatorname{dist}\left(\phi_{\tau}, u_{k}, P_{\tau}\right)+\operatorname{dist}\left(u_{k}, t, P_{\tau}\right)+8 \\
& =\left|P_{s, t, v}\right|+8,
\end{aligned}
$$

where the first inequality follows by the fact that $x, \phi_{\tau} \in C_{v}\left(\phi_{\tau}\right)$ and hence the intraclusrer path $R=\left[x, z_{v}\left[\phi_{\tau}\right], \phi_{\tau}\right]$ exists in $G \backslash\{v\}$ and last inequality follows Eq. (34). The claim follows.


## 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 $Q_{s, t, v}=P_{s, t, v}\left[\phi_{s, t, v}, t\right]$ was considered to be bought at time $\tau$. The green paths correspond to the shortest paths in $T_{0}(s)$. Red edges correspond to missing edges on $P_{s, t, v} \backslash E\left(G_{\tau}\right)$. 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 $\pi\left(s, \phi_{s, t, v}\right)$ strictly below the LCA vertex $\operatorname{LCA}\left(\phi_{s, t, v}\right)=\operatorname{LCA}\left(s, C_{v}\left(\phi_{s, t, v}\right)\right)$. The vertex $x \in C_{v}\left(\phi_{s, t, v}\right)$ satisfies that $v \notin \pi(s, x)$. The $s-t$ replacement path in $H \backslash\{v\}$ 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 $H_{8}(S)$ and claim the following.
Lemma $4.23\left|E\left(H_{8}(S)\right)\right|=\widetilde{O}\left(\max \left\{|S| \cdot n, n^{2} /|S|^{2}\right\}\right)$.
Proof: Let $H=H_{8}(S)$. 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 paths segments that were bought in the path-buying phase. For every ordered pair of clusters, $C_{1}, C_{2} \in \mathcal{C}$ let $\mathcal{B}\left(C_{1}, C_{2}\right)=\left\{P_{\tau} \in\right.$ $\left.\mathcal{B} \mid\left(C_{1}, C_{2}\right) \in \operatorname{ValSet}\left(P_{\tau}\right)\right\}$.

Clearly, $\mathcal{B}=\bigcup_{C_{1}, C_{2} \in \mathcal{C}} \mathcal{B}\left(C_{1}, C_{2}\right)$. We next claim that since the diameter of each cluster is small, it holds that the cardinality of each subset $\mathcal{B}\left(C_{1}, C_{2}\right)$ is small as well.

Lemma $4.24\left|\mathcal{B}\left(C_{1}, C_{2}\right)\right| \leq 5$ for every $C_{1}, C_{2} \in \mathcal{C}$.
Proof: Fix $C_{1}, C_{2} \in \mathcal{C}$ and let $\mathcal{B}\left(C_{1}, C_{2}\right)=\left\{P_{\tau_{1}}, \ldots, P_{\tau_{N}}\right\}$ be sorted according to the time $\tau_{k}$ their segment $Q_{\tau}$ was added to the spanner, for every $k \in\{1, \ldots, N\}$ where $N=\left|\mathcal{B}\left(C_{1}, C_{2}\right)\right|$. Let $P_{\tau_{k}}=P_{s_{k}, t_{k}, v_{k}}$. Let $p_{k}, q_{k} \in P_{\tau_{k}}$ be such that $p_{k}$ is closer to the source $s_{k}$, and $C_{v_{k}}\left(p_{k}\right)=C_{1}$ and $C_{v_{k}}\left(q_{k}\right)=C_{2}$.

Recall that $\phi_{\tau_{k}}$ is the first vertex of $Q_{\tau_{k}}$ (see Def. 4.16). Let $C_{\ell}=C_{v_{k}}\left(u_{\ell}\right)$ be the cluster of $u_{\ell}$ for every $u_{\ell} \in U_{s_{k}, t_{k}, v_{k}}$ (see Def. 4.17). By Obs. 2.3, it holds that $C_{\ell} \neq C_{\ell^{\prime}}$ for every $u_{\ell}, u_{\ell^{\prime}} \in$ $U_{s_{k}, t_{k}, v_{k}}$. Recall that for every $u_{\ell} \in U_{s_{k}, t_{k}, v_{k}}, P_{s, u_{\ell}, v}=P_{s, t, v}\left[s, u_{\ell}\right]$. Since LastE $\left(P_{s, u_{\ell}, v}\right) \notin E^{l o c a l}$, it holds that $v \notin\left\{z_{1}\left(u_{\ell}\right), \operatorname{LCA}\left(s, C_{1}\left(u_{\ell}\right)\right)\right\}$. Combining that with the fact that $u_{\ell} \in U_{s_{k}, t_{k}, v_{k}}$ did not satisfy property (N2) (see Def. 4.16 and Def. 4.17), we conclude that $v_{k} \in \pi^{f a r}\left(s_{k}, u_{\ell}\right)$. Since $q_{k} \in U_{s_{k}, t_{k}, v_{k}} \cup\left\{t_{k}\right\}$, using Lemma 4.19(c), it holds that

$$
\begin{equation*}
v_{k} \in V_{f}\left(C_{2}\right) \text { for every } k \in\{1, \ldots, N\} \tag{35}
\end{equation*}
$$

and by Lemma 4.19(b),

$$
\begin{equation*}
P_{\tau_{k}}\left[p_{k}, q_{k}\right] \subseteq Q_{\tau_{k}}\left[p_{k}, q_{k}\right] \subseteq G \backslash V_{f}\left(C_{2}\right) . \tag{36}
\end{equation*}
$$

Since $C_{1}=C_{v_{k}}\left(p_{k}\right)$ and $C_{2}=C_{v_{k}}\left(q_{k}\right)$, for every $k \in\{1, \ldots, N\}$, it holds that $z_{v_{1}}\left(p_{1}\right)=\ldots=$ $z_{v_{N}}\left(p_{N}\right)$ and also $z_{v_{1}}\left(q_{1}\right)=\ldots=z_{v_{N}}\left(q_{N}\right)$. Hence, letting $z_{1}=z_{v_{1}}\left(p_{1}\right)$ and $z_{2}=z_{v_{1}}\left(q_{1}\right)$, it holds that

$$
\begin{equation*}
z_{1}, z_{2} \notin\left\{v_{1}, \ldots, v_{N}\right\} . \tag{37}
\end{equation*}
$$

Denote

$$
X_{k}=\operatorname{dist}\left(p_{k}, q_{k}, G_{\tau_{k+1}} \backslash V_{f}\left(C_{2}\right)\right)
$$

We now show that $X_{k}<X_{k-1}$ for every $k \in\{2, \ldots N\}$.
Each time a path segment $Q_{\tau_{k}}$ is purchased at time $\tau_{k}$, it implies that

$$
\begin{align*}
X_{k} & \leq \operatorname{dist}\left(p_{k}, q_{k}, P_{\tau_{k}} \backslash V_{f}\left(C_{2}\right)\right)  \tag{38}\\
& =\operatorname{dist}\left(p_{k}, q_{k}, P_{\tau_{k}}\right)  \tag{39}\\
& <\operatorname{dist}\left(C_{1}, C_{2}, G_{\tau_{k}} \backslash V_{f}\left(C_{2}\right)\right)  \tag{40}\\
& \leq X_{k-1}, \tag{41}
\end{align*}
$$

where Eq. (38) follows by the fact that $P_{\tau_{k}}\left[p_{k}, q_{k}\right] \subseteq Q_{\tau_{k}} \subseteq G_{\tau_{k+1}}$, Eq. (39) follows by Eq. (36), Eq. (40) follows by the fact that $Q_{\tau_{k}}$ was bought and by Eqs. (27) and (28), and Eq. (41) follows by the fact that $p_{k-1} \in C_{1}$ and $q_{k-1} \in C_{2}$.

Therefore, we have that

$$
\begin{equation*}
X_{N} \leq X_{1}-(N-1) . \tag{42}
\end{equation*}
$$

Conversely, we have that

$$
\begin{align*}
X_{N} & \geq \operatorname{dist}\left(p_{N}, q_{N}, G \backslash\left\{v_{1}, \ldots, v_{N}\right\}\right)  \tag{43}\\
& \geq \operatorname{dist}\left(p_{1}, q_{1}, G \backslash\left\{v_{1}, \ldots, v_{N}\right\}\right)-4  \tag{44}\\
& =\operatorname{dist}\left(p_{1}, q_{1}, P_{\tau_{1}}\right)-4=\operatorname{dist}\left(p_{1}, q_{1}, P_{\tau_{1}} \backslash V_{f}\left(C_{2}\right)\right)-4  \tag{45}\\
& \geq X_{1}-4, \tag{46}
\end{align*}
$$

where Eq. 43) follows as $G_{\tau_{N+1}} \subseteq G$ and by Eq. (35), $\left\{v_{1}, \ldots, v_{N}\right\} \subseteq V_{f}\left(C_{2}\right)$. To see Eq. (44), note that $p_{1}, p_{N} \in C_{1}$ and $q_{1}, q_{N} \in C_{2}$ and by Obs. 2.3(3) the diameter of the cluster is 2 . It remains to show that the intracluster paths $R_{1}=\left[p_{1}, z_{1}, p_{N}\right], R_{2}=\left[q_{1}, z_{2}, q_{N}\right]$ exist in the surviving graph $G \backslash\left\{v_{1}, \ldots, v_{N}\right\}$. This holds since by Eq. (37), $z_{1}, z_{2} \notin\left\{v_{1}, \ldots, v_{N}\right\}$, and by Eq. (35) and (36). Eq. (45) follows by Eq. (36). Finally, Eq. (46) follows by the fact that $P_{\tau_{1}}\left[p_{1}, q_{1}\right] \subseteq Q_{\tau_{1}}$ was added at step $\tau_{1}$, hence $P_{\tau_{1}}\left[p_{1}, q_{1}\right] \subseteq G_{\tau_{2}}$. By combining with Eq. (42), we get that $N \leq 5$. The claim follows.

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

$$
\begin{aligned}
\left|E\left(G_{\tau^{\prime}}\right) \backslash E\left(G_{0}\right)\right| & =\sum_{P \in \mathcal{B}} \operatorname{Cost}(P) \leq 4 \sum_{P \in \mathcal{B}} \operatorname{Val}(P) \leq 4 \sum_{C_{1}, C_{2} \in \mathcal{C}}\left|\mathcal{B}\left(C_{1}, C_{2}\right)\right| \\
& \leq O\left(|\mathcal{C}|^{2}\right)=\widetilde{O}\left((n /|S|)^{2}\right)
\end{aligned}
$$

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

Additive stretch 6 (for all pairs). Thm. 1.2 follows immediately by Lemma 4.1. This should be compared with the single source additive FT-spanner $H_{4}(\{s\})$ of [18] and the (all-pairs, non FT) 6 -additive spanner, both with $\widetilde{O}\left(n^{4 / 3}\right)$ edges.

## 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 SWAT, 277-281, 2014.
[3] S. Baswana and S. Sen. Approximate distance oracles for unweighted graphs in expected $O\left(n^{2}\right)$ 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 STACS, 513-524, 2010.
[5] S. Baswana, T. Kavitha, K. Mehlhorn, and S. Pettie. Additive spanners and ( $\alpha, \beta$ )-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. 21st ACM-SIAM Symp. on Discrete Algorithms, 2010.
[7] A. Bernstein and D. Karger. A nearly optimal oracle for avoiding failed vertices and edges. In Proc. 41st ACM Symp. on Theory of Computing, 101-110, 2009.
[8] G. Braunschvig, S. Chechik and D. Peleg. Fault tolerant additive spanners. In WG, 206-214, 2012.
[9] S. Chechik, M. Langberg, D. Peleg, and L. Roditty. $f$-sensitivity distance oracles and routing schemes. Algorithmica 63, (2012), 861-882.
[10] S. Chechik, M. Langberg, D. Peleg, and L. Roditty. Fault-tolerant spanners for general graphs. In STOC, 435-444, 2009.
[11] S. Chechik. New Additive Spanners. In SODA, 2013.
[12] M. Cygan, F. Grandoni and T. Kavitha. On Pairwise Spanners. In STACS, 209-220, 2013.
[13] M. Dinitz and R. Krauthgamer. Fault-tolerant spanners: better and simpler. In PODC, 2011, 169-178.
[14] D. Dor, S. Halperin, and U. Zwick. All-pairs almost shortest paths. SIAM J. Computing, 29(5):17401759, 2000.
[15] M. Elkin and D. Peleg. ( $1+\varepsilon, \beta$ )-Spanner Constructions for General Graphs. SIAM Journal on Computing, 33(3): 608-631, 2004.
[16] F. Grandoni and V.V Williams. Improved Distance Sensitivity Oracles via Fast Single-Source Replacement Paths. In FOCS, 2012.
[17] M. Parter and D. Peleg. Sparse Fault-tolerant BFS trees. In ESA, 2013.
[18] M. Parter and D. Peleg. Fault Tolerant Approximate BFS Structures. In SODA, 2014.
[19] L. Roditty and U. Zwick. Replacement paths and k simple shortest paths in unweighted directed graphs. ACM Trans. Algorithms ,2012.
[20] M. Thorup and U. Zwick. Spanners and emulators with sublinear distance errors. In $S O D A, 802-809$, 2006.
[21] O. Weimann and R. Yuster. Replacement paths via fast matrix multiplication. FOCS, 2010.
[22] D.P Woodruff. Lower bounds for additive spanners, emulators, and more. In FOCS, 389-398, 2006.
[23] D.P Woodruff. Additive spanners in nearly quadratic time. In ICALP, 463-474, 2010.


[^0]:    *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).

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

[^2]:    ${ }^{2}$ The construction of 18 supports a single edge failure, yet, it can be modified to overcome a single vertex failure as well.

