Multipath Spanners via Fault-Tolerant Spanners

Shiri Chechik, Quentin Godfroy, David Peleg

Multipath Spanners via Fault-Tolerant Spanners

Shiri Chechik[1,⋆], Quentin Godfroy[2,⋆⋆], and David Peleg[3,⋆]

[1] Microsoft Research, Silicon Valley Center, USA
shiri.chechik@gmail.com
[2] LaBRI, Université Bordeaux-I, 351 cours de la Libération, Talence, France
quentin@godfroy.eu
[3] Department of Computer Science, The Weizmann Institute, Rehovot, Israel
david.peleg@weizmann.ac.il

Abstract. An s-spanner H of a graph G is a subgraph such that the distance between any two vertices u and v in H is greater by at most a multiplicative factor s than the distance in G. In this paper, we focus on an extension of the concept of spanners to *p-multipath distance*, defined as the smallest length of a collection of p pairwise (vertex or edge) disjoint paths. The notion of multipath spanners was introduced in [15, 16] for edge (respectively, vertex) disjoint paths. This paper significantly improves the stretch-size tradeoff result of the two previous papers, using the related concept of fault-tolerant s-spanners, introduced in [6] for general graphs. More precisely, we show that at the cost of increasing the number of edges by a polynomial factor in p and s, it is possible to obtain an s-multipath spanner, thereby improving on the large stretch obtained in [15, 16].

1 Introduction

Consider a graph $G = (V, E)$. An s-spanner H of a graph G is a spanning subgraph that preserves distances between all pairs of nodes within a factor of s, namely, such that $\delta(u, v, H) \leq s \cdot \delta(u, v, G)$ for every two nodes $u, v \in V$, where $\delta(u, v, G')$ for a graph G' is the distance from u to v in G'.

Graph spanners were introduced in [21, 22] in the context of distributed networks. It is well-known how to efficiently construct an s-spanner of size $O(n^{1+\frac{2}{s+1}})$ [1], for an odd integer s. Spanners have numerous applications, such as synchronizers [22], efficient routing [7, 8, 22, 25, 26], broadcasting [14], near-shortest path algorithms [11–13], covers [2], dominating sets [10], distance oracles [4, 27], or emulators and distance preservers [5]. For recent reviews on spanners see [23, 28].

This paper considers spanners for the *multipath* graph length. A *p-vertex (resp., p-edge) multipath* between two vertices u and v is a subgraph consisting

⋆ Supported in part by the Israel Science Foundation (grant 894/09), the United States-Israel Binational Science Foundation (grant 2008348), the Israel Ministry of Science and Technology (infrastructures grant), and the Citi Foundation.
⋆⋆ Supported by the Institut Universitaire de France.

G. Even and D. Rawitz (Eds.): MedAlg 2012, LNCS 7659, pp. 108–119, 2012.

of the union of p pairwise vertex-disjoint (resp., edge-disjoint) paths between u and v (except for the endpoints). The cost of a p-vertex (resp., p-edge) multipath between u and v is the sum of the weights of its edges.

For a weighted graph G and two vertices u and v, let $\delta_V^p(u,v,G)$ (resp., $\delta_E^p(u,v,G)$) be the minimum cost of a p-vertex (resp., p-edge) multipath in G between u and v if one exists, and ∞ otherwise. We say that a subgraph H is a *p-vertex (resp., p-edge) multipath s-spanner* of G if $\delta_V^p(u,v,H) \leq s \cdot \delta_V^p(u,v,G)$ (resp., $\delta_E^p(u,v,H) \leq s \cdot \delta_E^p(u,v,G)$) for every two vertices u and v such that $\delta_V^p(u,v,G) \neq \infty$ (resp., $\delta_E^p(u,v,G) \neq \infty$).

Our interest in the disjoint multipath graph length stems from the need for multipath routing in networks. Using multiple paths between a pair of nodes is an obvious way to aggregate bandwidth. Additionally, a classical approach to quickly overcome link failures consists in pre-computing alternate paths which are disjoint from the primary paths [18, 24, 20]. Multipath routing can be used for traffic load balancing and for minimizing delays. It has been extensively studied in ad hoc networks for load balancing, fault-tolerance, higher aggregate bandwidth, diversity coding, minimizing energy consumption (see [19] for a quick overview). Considering only a subset of links is a practical concern in link state routing in ad hoc networks [17]. This raises the problem of computing spanners for the multipath graph length, a first step towards constructing compact multipath routing schemes.

Edge-disjoint multipath spanners were first introduced in [15], where the authors provided construction examples. More specifically, they showed how to construct a p-edge multipath spanner with $O(pn^{1+\frac{2}{s+1}})$ edges and stretch $s \cdot p$. It was also shown that the factor p in the stretch can be discarded by an ad-hoc construction for $p = 2$ and $s = 3$. In [16] the authors proved that for every weighted graph one can efficiently construct a p-vertex multipath spanner with $O(p^2 n^{1+\frac{2}{s+1}})$ edges, but with a large stretch for large values of p (about $(1 + p/s)^s$).

In this paper, we significantly improve the stretch bounds from the two previous papers, using the related concept of fault-tolerant s-spanners, introduced in [6] for general graphs. We show that the constructions of edge-fault tolerant s-spanners from [6] and vertex-fault tolerant s-spanners from [9] can yield p-multipath s-spanners, by fixing the right number of faults. Specifically, for edge multipath spanners we reduce the stretch from $s \cdot p$ to s, thus deriving edge multipath spanners with stretch-size tradeoff close to the best known bounds for standard spanners. For vertex multipath spanners we reduce the stretch bound from a function of both s and p to just s in the case of unweighted graphs, and to $s \cdot p$ in the case of weighted graphs, s being the stretch of the underlying fault-tolerant spanner.

More precisely, we show the following theorems.

Theorem 1. *Given a weighted graph $G = (V, E)$ with maximal edge weight $\hat{\omega}$ and minimal edge weight 1, odd integer s and integer p, one can efficiently construct a p-edge-multipath s-spanner with $O\left(sp^2 \cdot \log \hat{\omega} \cdot n^{1+\frac{2}{s+1}}\right)$ edges.*

Theorem 2. *Given an unweighted graph $G = (V, E)$, odd integer s and integer p, one can efficiently construct a p-vertex-multipath s-spanner with $O\left(\left(s^{6\alpha}p^{2\alpha} + (s \cdot p)^{4\alpha}\right) n^{1+\frac{2}{s+1}} \log n\right)$ edges, where $\alpha = 1 - \frac{1}{s+1}$.*

Theorem 3. *Given a weighted graph $G = (V, E)$, odd integer s and integer p, one can efficiently construct a p-vertex-multipath $(s \cdot p)$-spanner with $O\left((p \cdot s)^{2\alpha} \cdot n^{1+\frac{2}{s+1}} \log n\right)$ edges, where $\alpha = 1 - \frac{1}{s+1}$.*

2 Preliminaries

For a graph H and two vertices u and v, let $\delta(u, v, H)$ denote the distance between u and v in H. A subgraph H of G is a *spanner* of *stretch* s (or, an s-spanner) of G if $\delta(u, v, H) \leq s \cdot \delta(u, v, G)$ for every $u, v \in V$. For a path P and two vertices x and y on it, let $P[x, y]$ denote the subpath of P from x to y.

A graph H is an *r-vertex (resp., edge) fault-tolerant s-spanner* of G if for any set $F \subseteq V$ (resp., $F \subseteq E$) of size at most r, the subgraph $H \setminus F$ is an s-spanner of $G \setminus F$, where $H' \setminus F'$ for a subgraph H' and set of vertices F' is the graph obtained by removing all vertices F' from H' together with their edges. Similarly $H' \setminus F'$ for a subgraph $H' = (V', E')$ and set of edges F' is the graph $(V', E' \setminus F')$.

For a subgraph H, let $\mathbf{cost}(H)$ denote the sum of weights of the edges in H.

For a path P, let $V(P)$ be the set of vertices on P and let $E(P)$ be the set of edges on P. Similarly for a set of paths $\mathcal{S}$, let $V(\mathcal{S}) = \bigcup_{P \in \mathcal{S}} V(P)$ and $E(\mathcal{S}) = \bigcup_{P \in \mathcal{S}} E(P)$. Consider two paths P_1 and P_2 between the same two vertices x and y. We say that P_1 and P_2 are *internal vertex-disjoint* if $V(P_1) \cap V(P_2) = \{x, y\}$. We say that a path is an *s-path* if it is of length at most s.

3 Edge Disjoint Multipath Spanners

In this section we show that every edge fault-tolerant spanner is an edge multipath spanner with the same stretch, by fixing the right number of faults.

We begin by considering the unweighted case, and show later how to proceed with weighted graphs.

A set of paths $\mathcal{B}$ is called a set of *edge-disjoint-s-bypasses* of u and v if the paths in $\mathcal{B}$ are edge-disjoint s-paths between u and v.

Consider two vertices x and y and let $\mathcal{S} = \{Q_1, ..., Q_p\}$ be the set of p vertex-disjoint paths from x to y in G of minimal cost. Consider an edge $e = (u, v)$ such that $e \in E(\mathcal{S})$. The next lemma shows that G does not contain "many" edge-disjoint-s-bypasses, namely, edge disjoint s-paths between u and v, that intersect with the other paths of $\mathcal{S}$.

Lemma 1. *Consider two vertices x and y and let $\mathcal{S} = \{Q_1, ..., Q_p\}$ be the set of p edge-disjoint paths from x to y in G of minimal cost. Consider an edge $e = (u, v) \in Q_i$ for some $1 \leq i \leq p$. There are at most $2sp + 2p$ edge-disjoint-s-bypasses between u to v, that intersect with $E(\mathcal{S}) \setminus E(Q_i)$.*

Proof: Assume towards contradiction that there are more than $2sp + 2p$ edge-disjoint-s-bypasses between u and v that intersect with $E(\mathcal{S}) \setminus E(Q_i)$. Let $\mathcal{B}$ be the set of all these edge-disjoint-s-bypasses that intersect with $E(\mathcal{S}) \setminus E(Q_i)$. Let J be the set of indices j such that $1 \leq j \leq p$, $j \neq i$ and $E(Q_j) \cap E(\mathcal{B}) \neq \emptyset$. For an index $j \in J$, let $e_h^j = (u_h^j, v_h^j)$ be the edge in $E(Q_j) \cap E(\mathcal{B})$ closest to x in Q_j, and let $e_l^j = (u_l^j, v_l^j)$ be the edge in $E(Q_j) \cap E(\mathcal{B})$ closest to y in Q_j. Let $M_j = [u_h^j, v_l^j]$ for $j \in J$ and let $M = \bigcup_{j \in J} M_j$. We show that it is possible to replace the set of edges $E(M)$ with a set of edges E' such that the resulting graph contains p edge disjoint paths from x to y and its cost is less than $\mathbf{cost}(\mathcal{S})$ and thus derive a contradiction to the optimality of $\mathcal{S}$.

We now explain how to gradually build the set of edges E'. Roughly speaking, the set of edges E' is the union of some prefixes and suffixes of the paths M_j together with some edge-disjoint-s-bypasses from $\mathcal{B}$.

Let $p_j = e_h^j$ (resp., $s_j = e_l^j$) be the prefix (resp., suffix) of M_j. Our construction process will gradually add edges to these prefixes and suffixes.

Let $\text{tip}(p_j)$ (resp., $\text{tip}(s_j)$) be the edge in $E(p_j) \cap E(\mathcal{B})$ (resp., $E(s_j) \cap E(\mathcal{B})$) closest to y on Q_j (resp., to x on Q_j). We examine the set of tips of the prefixes and suffixes, and the set of edge-disjoint-s-bypasses $\mathcal{B}'$ in which they appear. Let X be the set of prefixes and suffixes, namely, $X = \bigcup_{j \in J} \{p_j, s_j\}$. The set $\mathcal{B}'$ is the set of edge-disjoint-s-bypasses containing one of the edges $\{\text{tip}(p_j), \text{tip}(s_j) \mid j \in J\}$. For an edge $e \in E(\mathcal{B}')$, let $B(e)$ be the edge-disjoint-s-bypass such that $e \in E(B)$. Note that there is exactly one such edge-disjoint-s-bypass since the edge-disjoint-s-bypasses are disjoint. We say that a path $P \in X$ is *clean* if the sub-path $B(\text{tip}(e))[\text{tip}(P), u]$ does not contain other edges from $E(X)$. For an edge-disjoint-s-bypass $B \in \mathcal{B}'$, let $P_{clean}(B)$ be the path $P \in X$ such that $\text{tip}(P) \in E(B)$ and P is clean; note that there is exactly one such path.

We say that a prefix $p_j \in X$ (resp., suffix s_j) is *complete* if $p_j \circ s_j = M_j$. We apply the following process until all paths in X are clean. Choose an unclean incomplete path and add edges to it until it becomes clean. By adding an edge to a prefix p_j (resp., suffix s_j) we mean adding the edge on M_j adjacent to $\text{tip}(p_j)$ (resp., $\text{tip}(s_j)$) closest to y (resp., to x).

Note that it could happen that during this process some clean path becomes unclean. Note also that edge-disjoint-s-bypasses are only added to $\mathcal{B}'$ (but never removed). Namely, $\mathcal{B}'(t_1) \subseteq \mathcal{B}'(t_2)$ for $t_1 \leq t_2$, where $\mathcal{B}'(t)$ is the set $\mathcal{B}'$ in the t'th step of this process. To see this, note that the process does not add edges to $P_{clean}(B)$ for any $B \in \mathcal{B}'$. Thus in any stage of this process, B contains $\text{tip}(P)$ for $P = P_{clean}(B)$. Hence, by definition, $B \in \mathcal{B}'$. Notice that the path $P_{clean}(B)$ itself may change (since the process might add an edge to another path and this edge could belong to the path from $\text{tip}(P)$ to u).

We claim that $\mathcal{B}'$ contains at most $2p$ edge-disjoint-s-bypasses. This follows directly from the fact that each edge-disjoint-s-bypass $B \in \mathcal{B}'$ contains a different path $P_{clean}(B)$ and that there are $2p$ paths in X.

We now show that it possible to substitute the paths in M with "cheaper" paths and thus derive a contradiction to the optimality of $\mathcal{S}$. For every incomplete prefix $p_j \in X$, let p'_j be the clean subpath $B(p_j)[\text{tip}(p_j), u]$. Similarly, let s'_j be

the clean subpath $B(s_j)[u, \text{tip}(s_j)]$. For every index $j \in J$, if p_j is complete then set $Q'_j = Q_j$, otherwise set $Q'_j = Q_j[x, u_h^j] \circ p_j \circ p'_j \circ s'_j \circ s_j \circ Q_j[v_l^j, y]$.

Let $D' = \bigcup_{j \in J} E(p'_j \circ s'_j)$ and $D = \bigcup_{j \neq i} E(M_j) \setminus (E(p_j) \cup E(s_j))$. Let $\mathcal{S}' = \{Q'_1, ..., Q'_p\}$. Note that $|E(\mathcal{S}')| = |E(\mathcal{S})| + |D'| - |D|$. It is not hard to verify that the paths Q'_j are disjoint and each of them leads from x to y. Moreover, M intersects with at least $2sp+2p+1$ edge-disjoint-s-bypasses, and the set $E(p_j) \cup E(s_j)$ intersects with at most $2p$ edge-disjoint-s-bypasses, thus $|D| > 2sp + 1$. In addition, $|D'| \leq 2sp$. We get that $|D| > |D'|$ and thus $\mathbf{cost}(\mathcal{S}') < \mathbf{cost}(\mathcal{S})$, contradiction. ∎

We now show that edge fault-tolerant spanners constructed by the algorithm of [6] are also edge multipath spanners with the same stretch, by fixing the right number of faults. For completeness, we outline the algorithm and its properties that are essential for our needs. The algorithm operates in q iterations. Initially, set $H = (V, \emptyset)$. In each iteration i, consider the graph $G_i = (V, E \setminus E(H))$, and construct an s-spanner H_i for this graph, (say, using the algorithm of [1]), and add the edges of H_i to the subgraph H.

We summarize the required properties for our purposes in the following lemma.

Lemma 2. *[6] For every graph $G = (V, E)$, odd integer s and integer q, one can efficiently construct in polynomial time a collection of edge disjoint subgraphs $\{H_1, ..., H_q\}$ with the following properties. Let H be the union of the subgraphs $\{H_1, ..., H_q\}$.*

(1) The number of edges in H is at most $O(q \cdot n^{1+\frac{2}{s+1}})$.
(2) For every edge $e = (u, v) \in E$, either $e \in E(H)$ or each H_i contains a path from u to v of length at most s.

Theorem 4. *Given a graph $G = (V, E)$, odd integer s and integer p, one can efficiently construct a p-multipath s-spanner with $O\left(sp^2 \cdot n^{1+\frac{2}{s+1}}\right)$ edges.*

Proof: Construct the collection of subgraphs $\{H_1, ..., H_{2sp^2+2p^2+p}\}$ of Lemma 2 with parameters s and $q = 2sp^2 + 2p^2 + p$. Let H be the union of all subgraphs $\{H_1, ..., H_{2sp^2+2p^2+p}\}$. Consider two vertices x and y and let $\mathcal{S} = \{Q_1, ..., Q_p\}$ be the set of p edge-disjoint paths from x to y in G of minimal cost. We now show how to find a set of edge-disjoint paths $\mathcal{S}' = \{Q'_1, ..., Q'_p\}$ from x to y such that $E(\mathcal{S}') \subseteq E(H)$ and $\mathbf{cost}(Q'_i) \leq s \cdot \mathbf{cost}(Q_i)$. Let $T_i = \{H_j \mid (2sp + 2p + 1) \cdot (i - 1) + 1 \leq j \leq (2sp+2p+1) \cdot i\}$ for $1 \leq i \leq p$. Note that $E(T_i)$ contains $2sp+2p+1$ edge-disjoint paths from u to v for every edge $(u, v) \notin E(H)$. Moreover, the sets $E(T_i)$ are disjoint for $1 \leq i \leq p$.

The path Q'_i is constructed as follows. For every edge $e \in E(Q_i) \cap E(H)$ add e to Q'_i. For every edge $e = (u, v) \in E(Q_i) \setminus E(H)$, consider the set $\mathcal{B}_i$ with the maximum number of edge-disjoint-s-bypasses from u to v in T_i. By Lemma 1, there are at most $2sp + 2p$ edge-disjoint-s-bypasses in $\mathcal{B}_i$ that intersect with $E(\mathcal{S}) \setminus E(Q_i)$. Since $\mathcal{B}_i$ contains at least $2sp+2p+1$ edge-disjoint-s-bypasses, at least one edge-disjoint-s-bypass $B(e) \in \mathcal{B}_i$ does not intersect with $E(\mathcal{S}) \setminus E(Q_i)$. Add $B(e)$ to Q'_i instead of e.

We claim that (1) the paths Q'_i for $1 \leq i \leq p$ are edge-disjoint, and (2) $\mathbf{cost}(\mathcal{S}') \leq s \cdot \mathbf{cost}(\mathcal{S})$.

To see claim (1), consider an edge $e = (u, v)$ such that $e \in E(Q'_i)$ for some $1 \leq i \leq p$. We consider two cases. The first case is when $e \in E(Q_i)$. Note that e does not appear in any $E(Q_j)$ for $i \neq j$, since the paths in $\mathcal{S}$ are disjoint. Moreover, e does not appear in any $B(e')$ for $e' \in E(Q_j)$ for some $j \neq i$. To see this, recall that $B(e')$ does not intersect with $E(\mathcal{S}) \setminus E(Q_j)$. The second case is when $e \in B(\tilde{e})$ for some $\tilde{e} \in E(Q_i)$. The edge-disjoint-s-bypass $B(\tilde{e})$ does not intersect with $E(\mathcal{S}) \setminus E(Q_i)$. Moreover, the edge-disjoint-s-bypass $B(\tilde{e})$ does not intersect with any $B(e')$ for $e' \in E(Q_j)$ for some $j \neq i$. To see this, recall that $E(B(\tilde{e})) \subseteq E(T_i)$, $E(B(e')) \subseteq E(T_j)$, and $E(T_i) \cap E(T_j) = \emptyset$. It follows that the paths Q'_i are edge-disjoint for $1 \leq i \leq p$.

To see claim (2), note that for every edge $e \in E(Q_i)$, either e itself or an alternative path of length s is added to $E(Q'_i)$. We get that $\mathbf{cost}(Q'_i) \leq s \cdot \mathbf{cost}(Q_i)$. Claim (2) follows. ∎

Weighted Graphs. We now show the modifications needed for weighted graphs. Assume the minimal edge weight is 1 and let $\hat{\omega}$ be the maximal edge weight. We now describe the algorithm for constructing p-edge multipath s-spanner. Initially, set $H = (V, \emptyset)$. Consider the graphs $G_i = (V, E_i)$ such that $E_i = \{e \in E \mid 2^{i-1} \leq \omega(e) \leq 2^i\}$ for every $1 \leq i \leq \lceil \log \hat{\omega} \rceil$. Construct the collection of subgraphs $F_i = \{H_1, ..., H_{4sp^2+2p^2+p}\}$ of Lemma 2 for parameters s and $q = 4sp^2 + 2p^2 + p$ on the graph G_i. Add $E(F_i)$ to H.

We claim that H is a p-edge multipath s-spanner. The analysis is very similar to the unweighted case. We now outline the slight changes.

Here we call a set of paths $\mathcal{B}$ a set of *edge-disjoint-s-bypasses* of two nodes u and v that are connected by an edge if the paths in $\mathcal{B}$ are edge-disjoint paths between u and v of length at most $s \cdot \omega(u, v)$ each.

Consider two vertices x and y and let $\mathcal{S} = \{Q_1, ..., Q_p\}$ be the set of p edge-disjoint paths from x to y in G of minimal cost. Consider an edge $e = (u, v) \in Q_i$ for some $1 \leq i \leq p$. Let i be the index such $2^{i-1} \leq \omega(e) \leq 2^i$. In Lemma 1 we prove that the graph $(V, E(F_i))$ contains at most $4sp + 4p$ edge-disjoint-s-bypasses $\mathcal{B}$ from u to v. Note that since the weight of the edges in the edge-disjoint-s-bypasses $\mathcal{B}$ could be half the weight $\omega(e)$, we double the factor of sp^2 in the number of edge-disjoint-s-bypasses. The rest of the proof of Lemma 1 is similar to the unweighted case.

The proof of Theorem 4 is also similar to the unweighted case, where for each edge $e = (u, v) \in Q_j$ for some $1 \leq j \leq p$ such that $e \notin E(H)$, we pick an edge-disjoint-s-bypass from F_i that does not intersect $E(\mathcal{S}) \setminus E(Q_j)$, for i such that $2^{i-1} \leq \omega(e) \leq 2^i$.

We thus conclude with Theorem 1.

4 Vertex Disjoint Multipath Spanners

In this section we show that every vertex fault-tolerant s-spanner is a vertex multipath spanner with the same stretch, by fixing the right number of faults.

Note that it is unclear how to generalize the analysis from the previous section to vertex disjoint multipath spanners. To see this, recall that in the previous section we consider a set $\mathcal{S} = \{Q_1, ..., Q_p\}$ of p vertex-disjoint paths from x to y in G of minimal cost. We claim that every edge $e \in E(Q_i)$ does not contain too many bypasses that intersect with the other paths of $\mathcal{S}$. To prove this claim, we substitute sub-paths of each Q_j with some E-bypasses of the edge $e = (u, v)$. All these E-bypasses are edge disjoint but not vertex-disjoint. Specifically, all these E-bypasses contain the nodes u and v. Therefore, it is unclear how to use these E-bypasses to substitute multiple sub-paths and stay with vertex disjoint paths. We thus present a different analysis for vertex disjoint multipath spanners at the price of slightly increasing the size of the spanner. Moreover, our analysis for vertex disjoint multipath spanners works only for unweighted graphs. We later show a simple construction for weighted vertex multipath spanners with stretch sp (instead of s).

A subgraph H is *q-vertex-resilient with stretch s* if for every edge $e = (x, y) \in E$, either $e \in E(H)$ or H has at least q internal vertex-disjoint s-paths between x and y. For a path P between two nodes x and y and a vertex $v \in V(P)$, let $\mathbf{index}(v, P)$ be the distance (number of hops) between x and v in P. Two paths are said to intersect if they have at least one common vertex.

A set of paths $\mathcal{B}$ is called a set of *vertex-disjoint-s-bypasses* of u and v if the paths in $\mathcal{B}$ are internal vertex-disjoint s-paths between u and v. The next lemma shows that every vertex fault-tolerant s-spanner H has "many" vertex-disjoint-s-bypasses between u and v for every edge $e = (u, v)$ in $E \setminus E(H)$.

Lemma 3. *Every r-fault tolerant s-spanner is $\lfloor r/(s-1) \rfloor$-vertex-resilient with stretch s.*

Proof: Consider an r-fault tolerant s-spanner H. Consider an edge $e = (u, v) \in E \setminus E(H)$. We need to show that H contains $\lfloor r/(s-1) \rfloor$ vertex-disjoint-s-bypasses. Assume towards contradiction that H contains only k vertex-disjoint-s-bypasses between u and v such that $k < \lfloor r/(s-1) \rfloor$. Let $\mathcal{B}$ be the set of these k vertex-disjoint-s-bypasses. Note that $V(\mathcal{B}) \setminus \{u, v\}$ contains at most $(s-1) \cdot k < r$ vertices. Fix the set of vertices $F = V(\mathcal{B}) \setminus \{u, v\}$ to be faulty. Since the subgraph H is an r-fault tolerant s-spanner, by definition $H \setminus F$ contains an s-path between u and v. Therefore H contains more than k vertex-disjoint-s-bypasses between u and v, contradiction. ∎

Throughout, we consider a graph $G(V, E)$. Consider two vertices x and y and let $\mathcal{S} = \{Q_1, ..., Q_p\}$ be a set of p vertex-disjoint paths from x to y in G of minimal cost. Consider an edge $e = (u, v)$ in one of the paths of $\mathcal{S}$. The next lemma shows that G does not contain "many" vertex-disjoint-s-bypasses, namely, internal vertex disjoint s-paths between u and v, that intersect with the other paths of $\mathcal{S}$.

Lemma 4. *Consider two vertices x and y and let $\mathcal{S} = \{Q_1, ..., Q_p\}$ be a set of p vertex-disjoint paths from x to y in G of minimal cost. Consider an edge $e = (u, v) \in Q_i$ and $\notin H$ for some $1 \leq i \leq p$. There are at most $2sp(p-1) + 2p(p-1)$ vertex-disjoint-s-bypasses between u and v that intersect with $V(\mathcal{S}) \setminus V(Q_i)$.*

Proof: Assume towards contradiction that there are more than $2sp(p-1) + 2p(p-1)$ such vertex-disjoint-s-bypasses that intersect $V(\mathcal{S}) \setminus V(Q_i)$. By the Pigeonhole principle, there exists a path Q_j for some $j \neq i$ such that at least $2sp + 2p + 1$ of these vertex-disjoint-s-bypasses intersect with Q_j. Let $\mathcal{B}$ be the set of all these vertex-disjoint-s-bypasses that intersect Q_j. For every vertex-disjoint-s-bypass $A \in \mathcal{B}$, let $\mathbf{top}(A)$ be the earliest vertex of A on Q_j, i.e., the vertex in $V(A) \cap V(Q_j)$ with minimal $\mathbf{index}(\mathbf{top}(A), Q_j)$, and let $\mathbf{bottom}(A)$ be the last vertex of A on Q_j, i.e., the vertex in $V(A) \cap V(Q_j)$ with maximal $\mathbf{index}(\mathbf{top}(A), Q_j)$. Let $\mathcal{B}_h$ be the set of p vertex-disjoint-s-bypasses $A \in \mathcal{B}$ with minimal $\mathbf{index}(\mathbf{top}(A), Q_j)$ and let $\mathcal{B}_l$ be the set of p vertex-disjoint-s-bypasses $A \in \mathcal{B}$ with maximal $\mathbf{index}(\mathbf{bottom}(A), Q_j)$. Let $A_h \in \mathcal{B}_h$ be the vertex-disjoint-s-bypass with maximal $\mathbf{index}(\mathbf{top}(A_h), Q_j)$ and let $q_h = \mathbf{top}(A_h)$. Let $A_l \in \mathcal{B}_l$ be the vertex-disjoint-s-bypass with minimal $\mathbf{index}(\mathbf{bottom}(A_l), Q_j)$ and let $q_l = \mathbf{bottom}(A_l)$. Let $M = Q_j[q_h, q_l]$ (i.e., the subpath of Q_j from q_h to q_l). See Figure 1 for illustration.

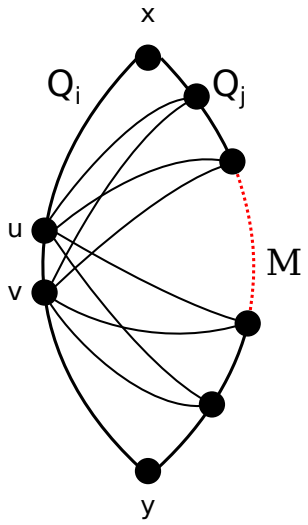

Fig. 1. Illustration of the sets $\mathcal{B}_h$, $\mathcal{B}_l$ ($p = 2$) and the path M

We claim that (1) the subgraph $H'(V, \mathcal{S}')$ for $\mathcal{S}' = (E(\mathcal{S}) \setminus E(M)) \cup E(\mathcal{B}_h) \cup E(\mathcal{B}_l)$ contains p vertex-disjoint paths from x to y and (2) $\mathbf{cost}(\mathcal{S}') \leq \mathbf{cost}(\mathcal{S})$.

To prove claim (1) we use Menger's theorem. We show that there is no set F of $p-1$ vertices such that x and y are disconnected in $\mathcal{S}' \setminus F$. Consider a set F of at most $p-1$ vertices. If F fails to intersect a path $Q_r \in \mathcal{S}$, for some $r \neq j$, then clearly x and y are connected in $\mathcal{S}' \setminus F$. So suppose the set F disconnects every path $Q_r \in \mathcal{S}$, for $r \neq j$, hence F contains exactly one vertex from each path $Q_r \in \mathcal{S}$ for every $r \neq j$. In particular, F contains only one vertex of Q_i. Therefore, one of u or v is not in F. Assume without loss of generality that $u \notin F$. Note that both sets $\mathcal{B}_h$ and $\mathcal{B}_l$ contain p vertex-disjoint-s-bypasses.

Since F contains at most $p-1$ vertices, there must be a vertex-disjoint-s-bypass $B_h \in \mathcal{B}_h$ and a vertex-disjoint-s-bypass $B_l \in \mathcal{B}_l$ whose internal vertices are not in F. Let $x_1 = \mathbf{top}(B_h)$ and $y_1 = \mathbf{bottom}(B_l)$. Note that the subpaths $Q_j[x, x_1]$ and $Q_j[y_1, y]$ do not contain any vertex from F, as $F \cap V(Q_j) = \emptyset$. Moreover, the vertex-disjoint-s-bypasses B_h and B_l contain subpaths $B_h[x_1, u]$ and $B_l[u, y_1]$ that do not intersect F. Concatenating all these paths together, we get a path $Q_j[x, x_1] \circ B_h[x_1, u] \circ B_l[u, y_1] \circ Q_j[y_1, y]$ from x to y. We thus conclude that H' contains p vertex-disjoint paths from x to y, establishing (1). Next, we show claim (2). Recall that $\mathcal{B}$ contains at least $2sp + 2p + 1$ vertex-disjoint-s-bypasses where each of which intersects Q_j. Moreover, each of the subpaths $Q_j[x, q_h]$ and $Q_j[q_l, y]$ intersect with exactly p vertex-disjoint-s-bypasses from $\mathcal{B}$. We get that the remaining part of Q_j, namely, the path $M = Q_j[q_h, q_l]$, intersects with at least $2sp + 2p + 1 - 2p = 2sp + 1$ vertex-disjoint-s-bypasses from $\mathcal{B}$. Thus, the length of M is at least $2sp + 1$. In contrast, the number of edges in the edge collection $E(\mathcal{B}_h) \cup E(\mathcal{B}_l)$ that replaced M in $\mathcal{S}$ is at most $2ps$. Hence, $\mathbf{cost}(\mathcal{S}') < \mathbf{cost}(\mathcal{S})$.

Finally parts (1) and (2) of the claim imply a contradiction to the optimality of $\mathcal{S}$. The lemma follows. ∎

Let $f = ((4s + 2)(p - 1)s + 1 + 2sp(p - 1) + 2p(p - 1))$.

Lemma 5. *Every f-vertex-resilient subgraph H is a p-vertex disjoint multipath spanner.*

Proof: Consider two vertices x and y and let $\mathcal{S} = \{Q_1, ..., Q_p\}$ be the set of p vertex-disjoint paths from x to y in G of minimal cost. Consider an edge $e = (u, v) \in E(Q_i)$ such that $e \notin E(H)$ for some $1 \leq i \leq p$. By definition of H, H contains f vertex-disjoint-s-bypasses between u to v. By Lemma 4, there are at most $2sp(p - 1) + 2p(p - 1)$ vertex-disjoint-s-bypasses between u to v, that intersect $V(\mathcal{S}) \setminus V(Q_i)$. We thus get that there exists a set $\mathbf{Bypasses}(e)$ of at least $f - 2sp(p - 1) - 2p(p - 1)$ vertex-disjoint-s-bypasses between u to v in H, that do not intersect $V(\mathcal{S}) \setminus V(Q_i)$. We now show how to select for each edge $e \in Q_i$ for some i a vertex-disjoint-s-bypass $B_e \in \mathbf{Bypasses}(e)$, such that B_e is vertex disjoint with any $B_{e'}$ for any $e' \in Q_j$ such that $e' \notin H$ and $j \neq i$.

Consider an edge $e = (u, v) \in Q_i$ such that $e \notin E(H)$. Let E_e be the set of edges e' such that $e' \in E(Q_j) \setminus E(H)$ and $V(B) \cap V(B') \neq \emptyset$ for some $j \neq i$, $B \in \mathbf{Bypasses}(e)$ and $B' \in \mathbf{Bypasses}(e')$. Towards proving Lemma 5, we first prove the next auxiliary lemma.

Lemma 6. *For every edge $e \in E(\mathcal{S}) \setminus E(H)$, the set E_e contains at most $(4s + 2)(p - 1)$ edges*

Proof: Assume, towards contradiction, that $|E_e| \geq (4s + 2)(p - 1) + 1$. By the Pigeonhole principle, there is a path Q_j (for $j \neq i$) such that $|E(Q_j) \cap E_e| > 4s+2$. Let $e_h = (u_h, v_h)$ be the edge in $E(Q_j) \cap E_e$ closest to x in Q_j, and let $e_l = (u_l, v_l)$ be the edge in $E(Q_j) \cap E_e$ closest to y in Q_j. Let $h_1, h_2 \in \mathbf{Bypasses}(e)$, $h_3 \in \mathbf{Bypasses}(e_h)$ and $h_4 \in \mathbf{Bypasses}(e_l)$ such that h_1 and h_3 intersect and h_2 and h_4 intersect (it could be that $h_1 = h_2$). Let M be the subpath $Q_j[v_h, v_l]$.

We now claim that (1) the subgraph $H''(V, \mathcal{S}'')$ for $\mathcal{S}'' = E(\mathcal{S}) \setminus E(M) \cup E(h_1) \cup E(h_2) \cup E(h_3) \cup E(h_4)$ contains p vertex disjoint paths from x to y and (2) $\mathbf{cost}(\mathcal{S}'') < \mathbf{cost}(\mathcal{S})$.

We show claim (1) by using Menger's theorem to establish that $\mathcal{S}''$ contains p vertex-disjoint paths from x to y. Consider a set F of at most $p-1$ vertices. If F fails to intersect a path $Q_r \in \mathcal{S}$, for some $r \neq j$, then clearly x and y are connected in $\mathcal{S}'' \setminus F$. So suppose the set F disconnects every path $Q_r \in \mathcal{S}$, for $r \neq j$, hence F contains exactly one vertex from each path $Q_r \in \mathcal{S}$ for every $r \neq j$. Note that all vertices in h_3 and h_4 are not in F as h_3 and h_4 are disjoint from all Q_r for $r \neq j$ and we assume that F contains only nodes from Q_r for some $r \neq j$. Since F contains exactly one vertex in Q_i then one of u and v are not in F, assume w.l.o.g. that $u \notin F$. Let r_3 be a vertex in $h_3 \cap h_1$ and let r_4 be a vertex in $h_4 \cap h_2$. Note that the subgraph $(E(Q_j) \setminus E(M)) \cup E(h_3)$ contains a path from x to r_3 that does not intersect F. Similarly, the subgraph $(E(Q_j) \setminus E(M)) \cup E(h_4)$ contains a path from r_4 to y that does not intersect F. The cycle $h_1 \cup \{(u, v)\}$ contains at most one vertex in F and thus there is a path from r_3 to u that does not contain vertices from F. Similarly, $h_1 \cup \{(u, v)\}$ contains a path from r_4 to u that does not contain vertices from F. Concatenating all these paths together we get a path from x to y. Claim (1) follows. Next, we show claim (2) and thus derive a contradiction to the optimality of $\mathcal{S}$. the path Q_j contains at least $4s+3$ edges from E_e and since the path M contains all these edges but two (e_h and e_l), the length of M is at least $4s+1$. In contrast, the number of edges in the edge collection $E(h_1) \cup E(h_2) \cup E(h_3) \cup E(h_4)$ that replaced M in $\mathcal{S}$ is at most $4s$. We thus get that $\mathbf{cost}(\mathcal{S}'') < \mathbf{cost}(\mathcal{S})$. This implies a contradiction to the optimality of $\mathcal{S}$. The lemma follows. ∎

Consider the edges $e \in E(\mathcal{S}) \setminus E(H)$ one by one. For each edge $e \in E(Q_i) \setminus E(H)$, choose a vertex-disjoint-s-bypass B_e that does not intersect with any $B_{e'}$ for an edge e' that was already considered by this process and such that $e' \in E(Q_r) \setminus E(H)$ for some $r \neq i$. We claim that this process never gets stuck, namely, each time we consider an edge $e \in E(Q_i) \setminus E(H)$, there is at least one vertex-disjoint-s-bypass in $\mathbf{Bypasses}(e)$ that does not intersect with the other vertex-disjoint-s-bypasses selected so far. Let $\tilde{E}_e \subseteq E_e$ be the set of edges that were considered before e by this process. Note that $|\tilde{E}_e| \leq (4s+2)(p-1)$ by Lemma 6. Moreover, note that each path $B_{e'}$ for some $e' \in \tilde{E}_e$ intersects with at most s vertex-disjoint-s-bypasses in $\mathbf{Bypasses}(e)$. Since there are more than $(4s+2)(p-1)s$ vertex-disjoint-s-bypasses in $\mathbf{Bypasses}(e)$, at least one of these vertex-disjoint-s-bypasses does not intersect with any of $B_{e'}$ for $e' \in \tilde{E}_e$.

For each path $Q_r \in \mathcal{S}$ for $1 \leq r \leq p$, construct a path $\tilde{Q}_r$ as follows. For every edge $e \in E(Q_r)$, if $e \in E(H)$ then add e to $\tilde{Q}_r$, otherwise add B_e to $\tilde{Q}_r$. It is not hard to verify that $V(Q_i) \cap V(Q_j) = \emptyset$ for any $i \neq j$ and that each Q_r is a path from x to y such that $\mathbf{cost}(\tilde{Q}_r) \leq s \cdot \mathbf{cost}(Q_r)$. The lemma follows. ∎

The following theorem was shown by Dinitz and Krauthgamer in [9].

Theorem 5. *[9] For every graph $G = (V, E)$, odd integer s and integer r, one can construct in polynomial time with high probability a r-vertex fault-tolerant s-spanner with $O\left(r^{2-\frac{2}{s+1}} n^{1+\frac{2}{s+1}} \log n\right)$ edges.*

Combining Theorem 5, Lemma 3 and Lemma 5, we get Theorem 2.

We note that for weighted graphs, every p-vertex-resilient H with stretch s is a p-vertex multipath $(s \cdot p)$-spanner. To see this, consider two vertices x and y and let $\mathcal{S} = \{Q_1, ..., Q_p\}$ be the set of p vertex-disjoint paths from x to y in G of minimal cost. We now show how to construct a subgraph $H'(V, \mathcal{S}')$ such that (1) H' contains p vertex-disjoint paths from x to y and (2) $\mathbf{cost}(\mathcal{S}') \leq s \cdot \mathbf{cost}(\mathcal{S})$.

Initially, set $\mathcal{S}' = \emptyset$. For every edge $e = (u, v) \in E(Q_r)$, if $e \in E(H)$ then add e to $\mathcal{S}'$, otherwise add to $\mathcal{S}'$ a set $\mathcal{B}(e)$ of V-bypasses such that $|\mathcal{B}(e)| = p$ and $E(\mathcal{B}(e)) \subseteq E(H)$. Note that such a set $\mathcal{B}(e)$ exists since H is p-vertex-resilient with stretch s.

To prove claim (1) we use Menger's theorem. We show that there is no set F of $p-1$ vertices such that x and y are disconnected in $\mathcal{S}' \setminus F$. Consider a set F of at most $p-1$ vertices. Note that, for every edge $e \in (E(Q_r) \setminus E(H))$, there exists a V-bypass $B(e) \in \mathcal{B}(e)$ that is internal vertex disjoint from F. To see this, recall that $\mathcal{B}(e)$ contains p V-bypasses. In addition, there exists at least one path $Q_r \in \mathcal{S}$ such that $V(Q_r) \cap F = \emptyset$ for some $1 \leq r \leq p$. Construct the path Q'_r as follows. For every edge $e = (u, v) \in E(Q_r)$, if $e \in E(H)$ then add e to Q'_r, otherwise add the alternative path $B(e)$ to Q'_r. It is not hard to verify that Q'_r leads from x to y, $E(Q'_r) \subseteq E(H)$, and $V(Q'_r) \cap F = \emptyset$. Claim (1) follows.

To see claim (2), note that for every edge $e \in E(\mathcal{S})$, either e itself or an alternative path of length s is added to $E(\mathcal{S}')$. We get that $\mathbf{cost}(\mathcal{S}') \leq s \cdot \mathbf{cost}(\mathcal{S})$.

We note that Lemma 3 can be generalized to weighted graphs. This can be done by invoking a construction for s-spanner, in the algorithm for constructing vertex fault-tolerant spanners of Dinitz and Krauthgamer in [9], that satisfies the following property, defined in [16]. Every edge e is either included in the spanner H or H contains an alternative path to e of length at most $s \cdot \omega(e)$ and with at most s hops. The details are omitted.

Combining with Theorem 5, we conclude Theorem 3.

References

1. Althöfer, I., Das, G., Dobkin, D., Joseph, D., Soares, J.: On sparse spanners of weighted graphs. Discrete & Computational Geometry 9, 81–100 (1993)
2. Awerbuch, B., Berger, B., Cowen, L., Peleg, D.: Near-linear cost sequential and distributed constructions of sparse neighborhood covers. In: Proc. 34th IEEE FOCS, pp. 638–647 (1993)
3. Barenboim, L., Elkin, M.: Deterministic distributed vertex coloring in polylogarithmic time. In: Proc. 29th ACM PODC, pp. 410–419 (2010)
4. Baswana, S., Kavitha, T.: Faster algorithms for approximate distance oracles and all-pairs small stretch paths. In: Proc. 47th IEEE FOCS, pp. 591–602 (2006)
5. Bollobás, B., Coppersmith, D., Elkin, M.: Sparse distance preservers and additive spanners. In: Proc. 14th ACM-SIAM SODA, pp. 414–423 (2003)

6. Chechik, S., Langberg, M., Peleg, D., Roditty, L.: Fault-tolerant spanners for general graphs. In: Proc. 41st ACM STOC, pp. 435–444 (2009)
7. Cowen, L.: Compact routing with minimum stretch. J. Algo. 38, 170–183 (2001)
8. Cowen, L., Wagner, C.: Compact roundtrip routing in directed networks. J. Algo. 50, 79–95 (2004)
9. Dinitz, M., Krauthgamer, R.: Fault-Tolerant Spanners: Better and Simpler. In: Proc. 30th ACM PODC, pp. 169–178 (2011)
10. Dubhashi, D., Mei, A., Panconesi, A., Radhakrishnan, J., Srinivasan, A.: Fast distributed algorithms for (weakly) connected dominating sets and linear-size skeletons. J. Computer and System Sciences 71, 467–479 (2005)
11. Elkin, M.: Computing almost shortest paths. ACM Tr. Algo. 1, 283–323 (2005)
12. Elkin, M.: A near-optimal distributed fully dynamic algorithm for maintaining sparse spanners. In: Proc. 26th ACM PODC, pp. 185–194 (2007)
13. Elkin, M., Zhang, J.: Efficient algorithms for constructing $(1 + \epsilon, \beta)$-spanners in the distributed and streaming models. In: Proc. 23rd ACM PODC, pp. 160–168 (2004)
14. Farley, A.M., Proskurowski, A., Zappala, D., Windisch, K.: Spanners and message distribution in networks. Discrete Applied Mathematics 137(2), 159–171 (2004)
15. Gavoille, C., Godfroy, Q., Viennot, L.: Multipath Spanners. In: Patt-Shamir, B., Ekim, T. (eds.) SIROCCO 2010. LNCS, vol. 6058, pp. 211–223. Springer, Heidelberg (2010)
16. Gavoille, C., Godfroy, Q., Viennot, L.: Node-Disjoint Multipath Spanners and Their Relationship with Fault-Tolerant Spanners. In: Fernàndez Anta, A., Lipari, G., Roy, M. (eds.) OPODIS 2011. LNCS, vol. 7109, pp. 143–158. Springer, Heidelberg (2011)
17. Jacquet, P., Viennot, L.: Remote spanners: what to know beyond neighbors. In: Proc. 23rd IEEE IPDPS (2009)
18. Kushman, N., Kandula, S., Katabi, D., Maggs, B.M.: R-BGP: Staying connected in a connected world. In: Proc. 4th NSDI (2007)
19. Mueller, S., Tsang, R.P., Ghosal, D.: Multipath Routing in Mobile Ad Hoc Networks: Issues and Challenges. In: Calzarossa, M.C., Gelenbe, E. (eds.) MASCOTS 2003. LNCS, vol. 2965, pp. 209–234. Springer, Heidelberg (2004)
20. Nasipuri, A., Castañeda, R., Das, S.R.: Performance of multipath routing for on-demand protocols in mobile ad hoc networks. Mobile Networks and Applications 6(4), 339–349 (2001)
21. Peleg, D., Schäffer, A.A.: Graph spanners. J. Graph Theory, 99–116 (1989)
22. Peleg, D., Ullman, J.D.: An optimal synchronizer for the hypercube. SIAM J. Computing 18(4), 740–747 (1989)
23. Pettie, S.: Low Distortion Spanners. In: Arge, L., Cachin, C., Jurdziński, T., Tarlecki, A. (eds.) ICALP 2007. LNCS, vol. 4596, pp. 78–89. Springer, Heidelberg (2007)
24. Pan, P., Swallow, G., Atlas, A.: Fast Reroute Extensions to RSVP-TE for LSP Tunnels. RFC 4090 (Proposed Standard) (2005)
25. Roditty, L., Thorup, M., Zwick, U.: Roundtrip spanners and roundtrip routing in directed graphs. ACM Trans. Algorithms 3(4), Article 29 (2008)
26. Thorup, M., Zwick, U.: Compact routing schemes. In: Proc. SPAA, pp. 1–10 (2001)
27. Thorup, M., Zwick, U.: Approximate distance oracles. JACM 52, 1–24 (2005)
28. Woodruff, D.P.: Lower bounds for additive spanners, emulators, and more. In: Proc. 47th IEEE FOCS, pp. 389–398 (2006)