

Multipath Spanners

Cyril Gavoille^{1*}, Quentin Godfroy^{1*}, and Laurent Viennot^{2**}

¹ University of Bordeaux, LaBRI

² INRIA, University Paris 7, LIAFA

Abstract. This paper concerns graph spanners that approximate multipaths between pair of vertices of an undirected graphs with n vertices. Classically, a spanner H of stretch s for a graph G is a spanning sub-graph such that the distance in H between any two vertices is at most s times the distance in G . We study in this paper spanners that approximate short cycles, and more generally p edge-disjoint paths with $p > 1$, between any pair of vertices.

For every unweighted graph G , we construct a 2-multipath 3-spanner of $O(n^{3/2})$ edges. In other words, for any two vertices u, v of G , the length of the shortest cycle (with no edge replication) traversing u, v in the spanner is at most thrice the length of the shortest one in G . This construction is shown to be optimal in term of stretch and of size. In a second construction, we produce a 2-multipath $(2, 8)$ -spanner of $O(n^{3/2})$ edges, i.e., the length of the shortest cycle traversing any two vertices have length at most twice the shortest length in G plus eight. For arbitrary p , we observe that, for each integer $k \geq 1$, every weighted graph has a p -multipath $p(2k - 1)$ -spanner with $O(pn^{1+1/k})$ edges, leaving open the question whether, with similar size, the stretch of the spanner can be reduced to $2k - 1$ for all $p > 1$.

Keywords: spanner, multipath

1 Introduction

This paper concerns the computation of sparse spanners for the multipath graph metric. We call *graph metric* a function δ that associates a metric δ_G with the vertex-set of a given graph G . A graph metric δ is *non-increasing* when distances can only decrease when adding edges. In other words, δ is non-increasing when $H \subseteq G$ implies $\delta_H \geq \delta_G$. (Here $H \subseteq G$ stands for H is a subgraph³ of G , and $\delta_H \geq \delta_G$ stands for $\delta_H(u, v) \geq \delta_G(u, v)$ for all⁴ u, v .) The *graph distance* d is the

* Supported by the ANR-project “ALADDIN”, the équipe-projet INRIA “CÉPAGE”, and the French-Israeli “Multi-Computing” project. The first author is Member of the “Insitut Universitaire de France”.

** Supported by the ANR-project “ALADDIN”, and the équipe-projet INRIA “GANG”.

³ I.e., $V(H) \subseteq V(G)$ and $E(H) \subseteq E(G)$.

⁴ For convenience, we set $\delta_H(u, v) = \infty$ if u or v is not in $V(H)$.

most classical graph metric: given an weighted undirected graph G , $d_G(u, v)$ is defined as the cost of a shortest path between u and v .

The notion of spanner is usually defined for the graph distance but it can be formulated for any non-increasing graph metric δ : an f -spanner (for δ) of a graph G is a subgraph $H \subseteq G$ such that $\delta_H(u, v) \leq f(\delta_G(u, v))$ for all $u, v \in V(G)$ where f is a stretch function satisfying $f(d) \geq d$ for $d \geq 0$. As δ is non-increasing, H must satisfy $\delta_G(u, v) \leq \delta_H(u, v) \leq f(\delta_G(u, v))$ for all u, v . Intuitively, an f -spanner H approximates the distances of G using possibly fewer edges. A rich literature studies the trade-off between the sparsity of H and the tightness of the stretch function f .

In this paper, we focus on the notion of spanner for the multipath graph metric. Given an integral value $p \geq 1$, the p -multipath graph metric d^p is defined as follows: given a weighted undirected graph G , $d_G^p(u, v)$ is the minimum weight sum of the edges of a set of p edge-disjoint paths joining u and v in G . The value $d_G^p(u, v)$ can be determined in polynomial time by computing a minimum-cost flow⁵ of value p between source u and sink v . For $p = 1$, we fall back on the graph distance: $d_G^1(u, v) = d_G(u, v)$. This introduces a new notion of spanner that we call *multipath spanner*.

1.1 Motivation

Our interest to the multipath graph metric 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 fail-over paths which are link-wise disjoint from primary paths [18,23,22]. Multipath routing can be used for traffic load balancing and for minimizing delays [32,15]. Multipath routing has been extensively studied in ad hoc networks for load balancing, fault-tolerance, higher aggregate bandwidth, diversity coding, minimizing energy consumption (see [21] for a quick overview). Heuristics have been proposed to provide disjoint routes [22,19] in on-demand protocols. There is a wide variety of optimization requirements when using several paths between pairs of nodes. However, using edge-disjoint or vertex-disjoint paths is a recurrent concern in optimizing routing in networks. Using disjoint paths is a dissertation subject in itself [17] and has many problem variants.

Considering only a subset of links is a practical concern in link state routing in ad hoc networks [16]. This raises the problem of computing spanners for the multipath graph metric. As node mobility results in link failures, having an edge-disjoint multipath between two nodes reduces the probability of disconnection. Additionally, spanners are a key ingredient in the design of compact routing schemes [25,30]. Designing multipath spanners is thus a first step toward multipath compact routing.

⁵ To avoid using an edge in both directions, we apply a standard reduction to digraphs: each undirected edge xy is replaced by the dipath $x \rightarrow x' \rightarrow y' \rightarrow y \rightarrow x' \rightarrow y' \rightarrow x$.

Another reason for considering edge-disjoint paths rather than vertex-disjoint paths is that the resulting distance is a metric. This is not the case with vertex-disjoint paths. More specifically, if we define $d_G^{*p}(u, v)$ as the minimum cost of a set of p internally vertex-disjoint paths joining u and v , the triangle inequality may be violated. For $p \geq 2$, one can easily build a graph G where $d_G^{*p}(u, v)$ and $d_G^{*p}(v, w)$ are finite whereas $d_G^{*p}(u, w)$ is not⁶. The proof that d_G^p satisfies the triangle inequality is given in Proposition 1 (see Section 2.1). Choosing the weight sum as the cost of a set of p edge-disjoint paths follows again the same rationale. d_G^p is the most natural metric generalizing d_G in the context of multipath routing.

1.2 Related work

The concept of spanner has been mostly studied for the graph distance as introduced by Peleg and Schäffer [24], and generalized to weighted graphs in [3]. There is an abundant literature on graph distance spanners which is surveyed, e.g., by Pettie in [26]. It is well-known that, for each integer $k \geq 1$, every weighted graph with n vertices has an f -spanner of $O(n^{1+1/k})$ edges with stretch function $f(d) = (2k - 1)d$. This can be proved using a greedy construction based on a simple modification of Kruskal's algorithm⁷, and this size-stretch tradeoff is conjectured to be tight. However, other tradeoffs exist, in particular for unweighted graphs and for stretch functions on the form $f(d) = \alpha d + \beta$. For $\alpha = 1$, the size is $O(n^{3/2})$ with $\beta = 2$ [1], and $O(n^{4/3})$ [4] with $\beta = 6$. There are also constructions for $\alpha = 1 + \varepsilon$ and arbitrary small $\varepsilon > 0$. The term $\beta = \beta(\varepsilon, k)$ depends on ε and k , and the size is $O(\beta n^{1+1/k})$ [11,31]. For $\alpha = O(1)$ and $\beta = \tilde{O}(1)$, spanners of size⁸ $O(n)$ exist [26]. Other stretch functions have been considered, e.g., $f(d) = d + O(d^{1-1/k})$ [26,31]. Distributed algorithms for constructing such spanners have been also developed in [7,8,10,12,27].

The concept of spanner has been extended to some other graph metrics. In weighted directed graphs, the one-way distance is not a graph metric. However, the *roundtrip distance*, defined as the one-way distance from u to v plus the one way distance from v to u , is a non-increasing graph metric. Roundtrip f -spanners of size $O(n^{3/2})$ exist for $f(d) = 3d$ [28], and the size is $\tilde{O}(n^{1+1/k})$ for $f(d) = (2^k - 1)d$ and for every $k > 1$ [6]. It is also proved that no such size-stretch tradeoff can exist if the usual one-way distance in directed graph is considered. Interestingly, these solutions lead to compact routing scheme in directed graphs.

Dragan and Yan [9] study the complexity of computing spanners approximating maximum flow between any two vertices. We observe that the inverse⁹ of the maximum flow is a graph metric (it is even an ultrametric), and thus is captured by our framework.

⁶ E.g., if v is a cut-vertex.

⁷ Visit all the edges of G by increasing order of their weights, and add the edge (x, y) to the current spanner H , initialized to the empty graph, only if $d_H(x, y) > (2k - 1) \cdot d_G(x, y)$.

⁸ Where $\tilde{O}(g(n))$ stands for $g(n) \cdot (\log n)^{O(1)}$.

⁹ The flow from u to u is considered to be ∞ .

A concept independent of spanner for a graph metric δ , is the one of fault-tolerant spanner. An f -spanner H is t -fault-tolerant if $\delta_{H \setminus F}(u, v) \leq f(\delta_{G \setminus F}(u, v))$ for all vertices u, v and for any set F of at most t faulty vertices and/or edges. Introduced in [20] for the graph distance and for geometric graphs, namely the d -dimensional Euclidean complete graph, it has been recently generalized to arbitrary weighted graphs in [5]. Surprisingly, size-stretch trade-offs are similar to the classic case where no-fault ($t = 0$) are tolerated. Note that t -fault-tolerant spanners preserve $(t+1)$ -connectivity: if there exists $t+1$ disjoint paths in the graph from u to v , then the spanner also contains $t+1$ disjoint paths from u to v . However, the stretch guarantee is different from ours. In particular, we can build examples where the stretch guarantee is low with respect to fault tolerance whereas it is high when considering the multipath graph metric.

1.3 Our contributions

Our results hold for undirected multi-edge graphs, and for the p -multipath graph metric d^p , for integer $p \geq 1$. Our contributions are the following:

1. We observe (Proposition 2) that every weighted graph has a p -multipath f -spanner of $O(pn^{1+1/k})$ edges, where $f(d) = p(2k - 1)d$. This is done by an iterative construction of standard graph distance spanners.
2. The analysis of the stretch can be refined for unweighted graphs, and we show that the previous construction for $p = k = 2$ actually leads to a 2-multipath f -spanner of $O(n^{3/2})$ edges, where $f(d) = 3d$.
3. We also show that all the lower bounds for $p = 1$, i.e., for the standard graph distance, generalize to p -multipath distance for any $p > 1$. In particular, the size-stretch tradeoff of our second result is optimal.
4. Using a quite different approach, we show that 2-multipath f -spanner of size $\Phi \cdot n^{3/2} + n$ edges exists for $f(d) = 2d + 8W$, where W is the largest edge-weight of the graph and $\Phi \approx 1.61$ is the golden ratio.

It may be worth to mention that, as long as δ is a non-increasing graph metric, the greedy Kruskal's algorithm can be naively applied to produce sparse approximate skeletons. It suffices to construct, from G , the weighted complete graph K defined by $V(K) = V(G)$ and the weight $\delta_G(u, v)$ assigned to the edge (u, v) . An f -spanner H of $O(n^{1+1/k})$ edges with $f(d) = (2k - 1)d$ can therefore be constructed. Unfortunately, H is not a spanner of the input graph G , it is only a subgraph of K . This general solution might be acceptable for emulator construction [31] where the output graph H is not required to be a subgraph.

A second observation is that, in spite of similarities between roundtrip and 2-multipath distances (in both cases a subgraph realizing the distance between any two vertices indeed consists of a shortest cycle), there is no reduction from 2-multipath to roundtrip spanners¹⁰.

¹⁰ By simply directing the edges of G and applying an efficient roundtrip spanner to G , we may obtain a roundtrip from u to v using twice the same arc, say $u \rightarrow a \rightarrow b \rightarrow v \rightarrow a \rightarrow b \rightarrow u$, which is not an acceptable solution for a 2-multipath spanner.

In the next section we give formal definitions and prove some important basic facts, including simple upper and lower bounds. In Section 3, we present the optimal 2-multipath f -spanner with $O(n^{3/2})$ edges with $f(d) = 3d$. In Section 4, we improve the stretch function to $f(d) = 2d + 8$. We give some open problems in Section 5.

Due to space limitation, some proofs have been moved to the Appendix.

2 Preliminaries

In this paper, we consider an undirected multi-edge weighted graph G with weight function ω . The *cost* of any subgraph H of G is the sum of the weights of its edges. It is denoted by $\omega(H) = \sum_{e \in E(H)} \omega(e)$. We set $\omega(H) = 0$ if $E(H) = \emptyset$.

2.1 Multipath distance and multipath spanner

A p -path from a vertex u to a vertex v is a subgraph of G composed of p edge-disjoint paths from u to v . We define the p -multipath distance between two vertices u and v , denoted by $d_G^p(u, v)$, as the minimum cost of a p -path from u to v . We set $d_G^p(u, v) = \infty$ if there are no p edge-disjoint paths from u to v .

Indeed, d_G^p is a distance. It clearly satisfies $d_G^p(u, v) = 0$ if and only if $u = v$. Symmetry follows from the fact that G is undirected, and the triangle inequality comes from Proposition 1.

Proposition 1. *Let u, v, w be any triple of vertices of a multi-edge graph G . If A is a p -path from u to v , and B a p -path from v to w , then $A \cup B$ contains a p -path from u to w . In particular, $d_G^p(u, w) \leq d_G^p(u, v) + d_G^p(v, w)$.*

Proof. Let (U, W) be a minimum cut between $u \in U$ and $w \in W$ in the graph induced by $H = A \cup B$. Let t be the number of edges. Consider a maximum flow on H with source u and sink v with unit capacity on each edge. From the min-cut max-flow theorem we derive that there are t edge-disjoint paths from u to w . If $t < p$, we conclude that u and v are both in U as there exists a p -path between them. The same argument can be used to show that $v, w \in W$. This is a contradiction, $U \cap W = \emptyset$.

The notions of p -path and p -multipath distance extend the usual notions of path and distance which correspond to the case $p = 1$. We write d_G a short for d_G^1 . We observe that the multipath distance does not extend to a vertex-disjoint version. Indeed, for each $p > 1$, the existence of p vertex-disjoint paths from u to v and of p vertex-disjoint paths from v to w does not imply there are p vertex-disjoint paths from u to w .

A subgraph H of G is a p -multipath s -spanner if $d_H^p(u, v) \leq s \cdot d_G^p(u, v)$ for all pairs of vertices u, v . The parameter s is called the p -multipath stretch of H . We also use the term *stretch*, instead of multipath stretch, when the context is clear. For $p = 1$, we fall back on the regular definition of s -spanner.

2.2 Iterative spanners

A p -iterative s -spanner of G is a subgraph $H = \bigcup_{i=1}^p H_i$, where H_i is any 1-multipath s -spanner of $G \setminus \bigcup_{j<i} H_j$. We observe that the union of p such 1-multipath spanners is actually a p -multipath spanner.

Proposition 2. *For all integers $k, p \geq 1$, every multi-edge weighted graph with n vertices has a p -multipath $p(2k-1)$ -spanner with less than $p \cdot n^{1+1/k}$ edges that can be constructed as a p -iterative $(2k-1)$ -spanner.*

Proof. Let $H = \bigcup_{i=1}^p H_i$ be a p -iterative $(2k-1)$ -spanner of G , where H_i is a $(2k-1)$ -spanner of G with less than $n^{1+1/k}$ edges. Each spanner H_i does exist (cf. [3]). Hence, H has less than $p \cdot n^{1+1/k}$ edges.

We now prove that H is a p -multipath $p(2k-1)$ -spanner. Let u, v be two vertices of G . If there is no p -path from u to v in G , then $d_G^p(u, v) = \infty$. In particular, $d_H^p(u, v) \leq p(2k-1) \cdot d_G^p(u, v)$. So, we assume there exists a p -path from u to v . Let P be any minimum-cost p -path from u to v in G . We have $\omega(P) = d_G^p(u, v)$. For an edge $e \notin H$, we denote by $H_i(e)$ the simple path which replaces the edge e of G in the i -th spanner member of H . Observe that, for each i , $\omega(H_i(e)) \leq (2k-1) \cdot \omega(e)$ because $e \in G \setminus H_i$ and H_i has stretch $2k-1$.

Given P and H , we define the subgraph F as follows:

$$F := (P \cap H) \cup \bigcup_{e \in P \setminus H} \bigcup_{i=1}^p H_i(e).$$

Clearly, $F \subseteq H$, since each edge $e \in P$ is either in H or is replaced by $H_i(e)$ for some i . Moreover, we have $\omega(F) \leq p(2k-1) \cdot \omega(P)$ because:

$$\begin{aligned} \omega(F) &\leq \omega(P \cap H) + \sum_{e \in P \setminus H} \sum_{i=1}^p \omega(H_i(e)) \\ &\leq \sum_{e \in P \cap H} \omega(e) + \sum_{e \in P \setminus H} p(2k-1) \cdot \omega(e) \\ &\leq \sum_{e \in P} p(2k-1) \cdot \omega(e) = p(2k-1) \cdot \omega(P). \end{aligned}$$

Therefore, the stretch of H is at most $p(2k-1)$ as claimed.

We now show that F contains a p -path from u to v , and for that we shall use the min-cut max-flow theorem. Consider a cut (X, \bar{X}) with $u \in X$ and $v \in \bar{X}$. Since P is a p -path from u to v , there is a subset C of the cut of at least p edges of P which have one endpoint in X and the other in \bar{X} . Two cases are possible:

1. Every edge of C belongs to F : the cut in F is already at least p .
2. One edge of C does not belong to F : p paths were added in F in replacement for this edge, so the minimum cut is at least p .

Therefore, the minimum cut in F is at least p . By the min-cut max-flow theorem there is p edge-disjoint paths from u to v in F . It follows that F contains a p -path from u to v . This completes the proof.

The rest of the paper studies how the $p(2k - 1)$ stretch bound can be improved.

2.3 Lower bounds

For all integers p, n and real $s > 1$, denote by $m_p(n, s)$ the largest integer such that there exists a multi-edge weighted graph with n vertices for which every p -multipath spanner of stretch $< s$ requires $m_p(n, s)$ edges.

The value of $m_p(n, s)$ provides a lower bound on the sparsity of p -multipath spanners of stretch $< s$. To illustrate this, consider for instance $p = 1$ and $s = 3$. It is known that $m_1(n, 3) = \Omega(n^2)$, by considering the complete bipartite graph $B = K_{\lfloor n/2 \rfloor, \lceil n/2 \rceil}$. Since all cycles of B have length at least 4, every proper subgraph H contains two vertices x and y which are neighbors in B but at distance at least 3 in the spanner. Thus H is an s -spanner with $s \geq 3$. In other words, every s -spanner of B , with $s < 3$ contains all the edges of B that is $\Omega(n^2)$ edges.

Unfortunately, this argument does not transfer to p -multipath spanners whenever $p > 1$. Indeed, with the same graph B , we get $d_B^p(x, y) = 1 + 3(p - 1)$. And, if (x, y) is removed from any candidate spanner H , we only get $d_H^p(x, y) = 3p$. Hence, the stretch for H so proved is only $d_H^p(x, y)/d_G^p(x, y) = 1 + O(1/p)$. The argument transfers however if multi-edges are allowed.

Proposition 3. *For all integers $n, p \geq 1$ and real $s > 1$, $m_p(n, s) \geq m_1(n, s)$.*

In particular, under the Erdős-Simonovits [13, 14] Conjecture¹¹ which implies $m_1(n, 2k + 1) = \Omega(n^{1+1/k})$ for every integer $k \geq 1$ and proved for $k \in \{1, 2, 3, 5\}$, there is a multi-edge unweighted graph with n vertices for which every p -multipath spanner with stretch $< 2k + 1$ has $\Omega(n^{1+1/k})$ edges.

Proof. Let G be an n -vertex graph with the minimum number of edges such that every spanner of stretch $< s$ has $m_1(n, s)$ edges. Let ω be the weight function of G . Clearly, G has $m_1(n, s)$ edges, since otherwise we could remove an edge of G . Observe also that any path between two neighbors x, y of G that does not use the edge (x, y) has length at least s , since otherwise we could remove it from G . In other words, $d_{G \setminus \{(x, y)\}}(x, y) \geq s \cdot \omega(x, y)$.

Let G_p be the graph obtained from G by adding, for each edge of G , $p - 1$ extra multi-edges with same weight. We have $G_1 = G$, and G_p has $p \cdot m_1(n, s)$ edges. Let H be any p -multipath spanner of G_p with $< m_1(n, s)$ edges. There must exist two vertices x, y adjacent in G_p that are not adjacent in H . We have

¹¹ It states that there are n -vertex graphs with $\Omega(n^{1+1/k})$ edges without cycles of length $\leq 2k$.

$d_{G_p}^p(x, y) = p \cdot \omega(x, y)$, and $d_H^p(x, y) \geq p \cdot d_{G \setminus \{(x, y)\}}(x, y) \geq p \cdot s \cdot \omega(x, y)$. We conclude that the p -multipath stretch of H is at least $d_{G_p}^p(x, y)/d_H^p(x, y) \geq s$.

In other words, every p -multipath spanner H of G_p with stretch $< s$ has $\geq m_1(n, s)$ edges, proving that $m_p(n, s) \geq m_1(n, s)$.

We leave open the question of determining whether the same lower bound of $2k - 1$ on the stretch applies if the graphs are restricted to be simple graphs only.

3 An unweighted 2-multipath 3-spanner

In this section, we focus on unweighted 2-multipath 3-spanners. The lower bound of Proposition 3 tells us that $\Theta(n^{3/2})$ is the required size of any 2-multipath 3-spanner. However, the p -iterative $(2k - 1)$ -spanner given by Proposition 2 (with $p = k = 2$) provides a 2-multipath spanner of stretch 6 only. In fact a finer analysis shows that the same construction yields a 2-multipath 3-spanner.

Theorem 1. *Every multi-edge unweighted graph with n vertices has a 2-multipath 3-spanner with less than $2n^{3/2}$ edges that can be constructed as a 2-iterative 3-spanner.*

It is obvious from the construction that a 2-iterative 3-spanner contains less than $2n^{3/2}$ edges. For the stretch, the long proof is divided in lemmas 2 to 6, that have been moved to the appendix.

4 A 2-multipath (2,8)-spanner

In this section, we construct a 2-multipath spanner with $O(n^{3/2})$ edges whose stretch is below 3 for (2-multipath) distances > 8 . In the remaining, (α, β) -spanner stands for f -spanner of stretch function $f(d) = \alpha d + \beta$.

4.1 Multipath spanning trees

To prove the main result of this section we extend the notion of spanning tree.

A p -multipath spanning tree of G is a subgraph T of G with a distinguished vertex u , called the *root* of T , such that, for every vertex v of G , T contains a p -path from u to v . Moreover, T is a p -shortest-path spanning tree if $d_T^p(u, v) = d_G^p(u, v)$ for every vertex v . For $p = 1$, we come back to the standard notions of spanning tree and shortest-path spanning tree. Observe that T may not exist, for instance, if G is not 2-edge-connected.

Lemma 1. *Every n -vertex 2-edge-connected graph with a given vertex u has a 2-shortest-path spanning tree rooted at u with at most $2(n - 1)$ edges constructible in polynomial time.*

Proof. We use the construction of [29] that can be extended to undirected graphs. The algorithm of [29] constructs 2-paths, each one of minimum cost, from every vertex to a fixed source of the graph. Roughly speaking, the construction results of two shortest-path spanning trees computed with Dijkstra’s algorithm. The 2-paths (from every vertex v to the source) are reconstructed via a specific procedure. This latter can be analyzed so that the number of edges used in the 2-multipath tree is at most $2(n-1)$: all vertices, but the source, have two parents.

The bound of $2(n-1)$ is tight because of the graph $K_{2,n-2}$. More generally, the number of edges in any p -multipath spanning tree must be, in the worst-case, at least $p(n-p)$, for every $p \leq n/2$. Indeed, every p -multipath spanning tree T must be p -edge-connected¹², and the graph $K_{p,n-p}$ is minimal for the p -edge-connectivity. Therefore, T contains all the edges of $K_{p,n-p}$, and there are $p(n-p)$ edges. Obviously, there are p -multipath spanning tree with less than $p(n-p)$ edges. Typically a subdivision of $K_{2,p}$ with n vertices has p -multipath spanning tree rooted at a degree- p vertex with a total of $n+p-2$ edges.

4.2 A stretch-(2,8W) spanner

Theorem 2. *Every multi-edge weighted graph with n vertices and largest edge-weight W has a 2-multipath (2,8W)-spanner with less than $\Phi n^{3/2} + n$ edges, where $\Phi \approx 1.618$ is the golden ratio.*

Proof. Let denote by $B_H^p(u, r) = \{v \in V(H) : d_H^p(u, v) \leq r\}$ the p -multipath ball of radius r in H centered at u , and denote by $N_H^p(u, r)$ the neighbors of u in H that are in $B_H^p(u, r)$. Note that for $p \geq 2$, some neighbor v of u might not be in $B_H^p(u, r)$ for every $r < \infty$: for instance if u and v are not in the same 2-edge-connected component. We denote by $\text{SPT}_H^p(u)$ a p -shortest-path tree rooted at u spanning the 2-edge-connected component of H containing u , and having at most $2(|E(H)| - 1)$ edges. According to Lemma 1, such p -shortest path tree can be constructed.

Let G be a multi-edge weighted graph with n vertices and largest edge-weight W . We denote by ω its edge-weight function. The 2-multipath spanner H of G is constructed thanks to the following algorithm (see Algorithm 1).

Size: Denote by G_3 and H_3 respectively the graphs G and H obtained after running the while-loop. Let b be the number of while-loops performed by the algorithm, and let $a = \sqrt{5} - 1$. Observe that $a^2 + 2a = 4$. From Lemma 1, the 2-shortest-path tree $\text{SPT}_G^2(u)$ has at most $2(n-1)$ edges. Hence, the size of H_3 is at most:

$$|E(H_3)| < 2b \cdot n .$$

The number of vertices of G_3 is at most $n - ab\sqrt{n}$, since at each loop, at least $a\sqrt{n}$ vertices are removed from G . Let G_3^1 be the graph induced by all

¹² By Proposition 1, there are two edge-disjoint paths between any two vertices of the p -multipath tree, through its root.

1. For each edge e of G : if there are in G two other edges between the endpoints of e with weight at most $\omega(e)$, then $G := G \setminus \{e\}$
2. $H := (V(G), \emptyset)$ and $W := \max \{\omega(e) : e \in E(G)\}$
3. While there exists $u \in V(G)$ such that $|N_G^2(u, 4W)| \geq (\sqrt{5} - 1)\sqrt{n}$:
 - (a) $H := H \cup \text{SPT}_G^2(u)$
 - (b) $G := G \setminus N_G^2(u, 4W)$
 - (c) $W := \max \{\omega(e) : e \in E(G)\}$
4. $H := H \cup G$

Algorithm 1: A 2-multipath $(2, 8W)$ -spanner algorithm.

the edges (u, v) of G_3 such that $v \in N_{G_3}^2(u, 4W_1)$, where W_1 is the maximum weight of an edge of the graph obtained after running Instruction 1. Let G_3^2 be the graph induced by the edges of $G_3 \setminus G_3^1$. The degree of each vertex u of G_3^1 is $|N_{G_3}^2(u, 4W_1)| - 1$ which is $< \lceil a\sqrt{n} \rceil - 1 \leq a\sqrt{n}$ because of the while-condition. Therefore, the size of G_3^1 is at most:

$$|E(G_3^1)| \leq \frac{1}{2} \sum_{u \in V(G_3)} a\sqrt{n} < \frac{1}{2} (n - ab\sqrt{n}) \cdot a\sqrt{n} < \frac{a}{2} \cdot n^{3/2} - \frac{a^2b}{2} \cdot n.$$

Let S_3 be the graph obtained from G_3^2 where each multi-edge is replaced by a single unweighted edge. More formally, vertices u and v are adjacent in S_3 if and only if there is at least one edge between u and v in G_3^2 . From Instruction 1, there is at most two edges between two adjacent vertices, so $|E(G_3^2)| \leq 2|E(S_3)|$.

Let us show that S_3 has no cycle of length ≤ 4 . Consider any edge (u, v) of S_3 . Observe that $v \notin N_{G_3}^2(u, 4W_3)$, where W_3 is the maximum weight of an edge of G_3 . Assume there is a path cycle of length at most 4 in S_3 going through (u, v) . Then in G_3^2 there is a 2-path from u to v of cost at most $4W_3$. Contradiction: $v \notin N_{G_3}^2(u, 4W_3)$ implies $d_{G_3^2}^2(u, v) > 4W_3$.

It has been proved in [2] that every simple η -vertex μ -edge graph without cycle of length $\leq 2k$, must verify the Moore bound:

$$\eta \geq 1 + \delta \sum_{i=0}^{k-1} (\delta - 1)^i > (\delta - 1)^k$$

where $\delta = 2\mu/\eta$ is the average degree of the graph. This implies that $\mu < \frac{1}{2}(\eta^{1+1/k} + \eta)$.

We observe that S_3 is simple. It follows, for $k = 2$ and $\eta \leq n - ab\sqrt{n}$, that the size of G_3^2 is at most (twice the one of S_3):

$$|E(G_3^2)| \leq (n - ab\sqrt{n})^{3/2} + n - ab\sqrt{n} < (n - ab\sqrt{n})\sqrt{n} + n = n^{3/2} + (1 - ab) \cdot n.$$

Overall, the total number of edges of the final spanner H is bounded by:

$$|E(H)| \leq |E(H_3)| + |E(G_3^1)| + |E(G_3^2)|$$

$$\begin{aligned}
&< \left(1 + \frac{a}{2}\right) \cdot n^{3/2} + \left(2b - \frac{a^2b}{2} + 1 - ab\right) \cdot n \\
&= \left(1 + \frac{a}{2}\right) \cdot n^{3/2} + n = \frac{1 + \sqrt{5}}{2} \cdot n^{3/2} + n = \Phi n^{3/2} + n
\end{aligned}$$

because the term $2b - a^2b/2 + 1 - ab = b/2 \cdot (4 - a^2 - 2a) + 1 = 1$. (Recall that, by the choice of a , $a^2 + 2a = 4$.)

Stretch: Let G_0 be the input graph G , before applying the algorithm. We first observe that we can restrict our attention to the stretch analysis of G_1 (instead of G_0), the graph obtained after applying Instruction 1.

Let H be a 2-multipath spanner for G_1 . Consider two vertices u, v of H , and let A be a minimum-cost 2-path between u and v in H . A is composed of two edge-disjoint paths and is of minimum cost in H , so A traverses (at most) two edges with same endpoints having the smallest weight. These (possibly) two edges exist in G_0 and in G_1 , therefore the 2-multipath stretch of H in G_0 or in G_1 is the same.

From the above observation, it suffices to prove that H is a 2-multipath $(2, 8W_1)$ -spanner of G_1 , where $W_1 \leq W_0$ is the maximum weight of an edge of G_1 .

Let x, y be any two vertices of G_1 , and A be a minimum-cost 2-path between x and y in G_1 . Let $d = d_{G_1}^2(x, y) = \omega(A)$. If all the edges of A are in H , then $d_H^2(x, y) = d_{G_1}^2(x, y) = d$, and the stretch is $(1, 0)$. So, assume that $A \not\subset H$. Let u be the first vertex selected in the while-loop such that $N_G^2(u, 4W)$ intersects A , so that Instruction 3(b) removes at least one edge of A . Let G, H be the graphs at the time u is selected, but before running Instruction 3(a) and 3(b). Let $v \in N_G^2(u, 4W) \cap A$, and B a minimum-cost 2-path from u and v in G . By definition of $N_G^2(u, 4W)$, $d_G^2(u, v) = \omega(B) \leq 4W$. Let $T = \text{SPT}_G^2(u)$.

An important observation is that u, x, y are in the same 2-edge-connected component of G . This comes from the fact that every 2-path is a 2-edge-connected subgraph¹³. So, A and B are 2-edge-connected, and $A \cup B$ as well, since A intersects B (in v).

Using the triangle inequality (Proposition 1) between x and y in H , we have $d_H^2(x, y) \leq d_H^2(u, x) + d_H^2(u, y)$. By construction of H and T , $d_H^2(u, x) = d_T^2(u, x) = d_G^2(u, x) \leq \omega(A \cup B)$ since we have seen that $u, x \in A \cup B$ that is 2-edge-connected. Thus, $d_H^2(u, x) \leq \omega(A) + \omega(B) \leq d + 4W_1$. Similarly, $d_H^2(u, y) \leq d + 4W_1$. Therefore, $d_H^2(x, y) \leq 2d + 4W_1$. The subgraph H is a 2-multipath $(2, 8W_1)$ -spanner, completing the proof.

5 Conclusion

We have introduced multipath spanners, some subgraphs that approximate the cost of p edge-disjoint paths between any two vertices of a graph. The usual

¹³ This observation becomes wrong whenever p -paths with $p > 2$ are considered.




Fig. 1. A weighted graph G with $d_G^2(x, y) = d$ showing that the stretch analysis in the proof of Theorem 2 is tight. The 2-shortest-path tree rooted at u spans all the edges but (x, y) . We have $d_T^2(u, x) = d_T^2(u, y) = d + 4$, and $d_T^2(x, y) = 2d + 8 - 2\varepsilon$.

notion of spanner is obtained for $p = 1$. This new notion leaves open several questions. We propose some of them:

- Is it true that every weighted graph has a p -multipath $(2k - 1)$ -spanner with $O(pn^{1+1/k})$ edges?
- Let $s < 2k + 1$. Does the bound $\Omega(n^{1+1/k})$ on the size of any p -multipath s -spanner hold for the class of n -vertex simple graphs (with no multi-edges)?
- Let $\varepsilon \in (0, 1]$. Is there a 2-multipath $(2 - \varepsilon, O(W))$ -spanner with $o(n^2)$ edges for every graph of largest edge-weight W ? and with $O(n^{3/2})$ edges? or even $O(n^{4/3})$ edges? And for $p > 2$?
- Prove, for $p > 2$, that every p -edge-connected graph has a p -shortest-path spanning tree with at most $p(n - 1)$ edges.

References

1. D. AINGWORTH, C. CHEKURI, P. INDYK, AND R. MOTWANI, *Fast estimation of diameter and shortest paths (without matrix multiplication)*, SIAM J. on Computing, 28 (1999), pp. 1167–1181.
2. N. ALON, S. HOORY, AND N. LINIAL, *The Moore bound for irregular graphs*, Graphs and Combinatorics, 18 (2002), pp. 53–57.
3. I. ALTHÖFER, G. DAS, D. DOBKIN, D. A. JOSEPH, AND J. SOARES, *On sparse spanners of weighted graphs*, Discrete & Computational Geometry, 9 (1993), pp. 81–100.
4. S. BASWANA, T. KAVITHA, K. MEHLHORN, AND S. PETTIE, *New constructions of (α, β) -spanners and purely additive spanners*, in 16th Symposium on Discrete Algorithms (SODA), ACM-SIAM, Jan. 2005, pp. 672–681.
5. S. CHECHIK, M. LANGBERG, D. PELEG, AND L. RODITTY, *Fault-tolerant spanners for general graphs*, in 41st Annual ACM Symposium on Theory of Computing (STOC), ACM Press, May 2009, pp. 435–444.

6. L. J. COWEN AND C. WAGNER, *Compact roundtrip routing in directed networks*, in 19th Annual ACM Symposium on Principles of Distributed Computing (PODC), ACM Press, July 2000, pp. 51–59.
7. B. DERBEL, C. GAVOILLE, D. PELEG, AND L. VIENNOT, *On the locality of distributed sparse spanner construction*, in 27th Annual ACM Symposium on Principles of Distributed Computing (PODC), ACM Press, Aug. 2008, pp. 273–282.
8. ———, *Local computation of nearly additive spanners*, in 23rd International Symposium on Distributed Computing (DISC), vol. 5805 of LNCS, Springer, Sept. 2009, pp. 176–190.
9. F. F. DRAGAN AND C. YAN, *Network flow spanners*, in 7th Latin American Symposium on Theoretical Informatics (LATIN), vol. 3887 of LNCS, Springer, Mar. 2006, pp. 410–422.
10. M. ELKIN, *Computing almost shortest paths*, ACM Transactions on Algorithms, 1 (2005), pp. 283–323.
11. M. ELKIN AND D. PELEG, $(1 + \epsilon, \beta)$ -spanner constructions for general graphs, SIAM J. on Computing, 33 (2004), pp. 608–631.
12. M. ELKIN AND J. ZHANG, *Efficient algorithms for constructing $(1 + \epsilon, \beta)$ -spanners in the distributed and streaming models*, Distributed Computing, 18 (2006), pp. 375–385.
13. P. ERDÖS, *Extremal problems in graph theory*, in Publ. House Czechoslovak Acad. Sci., Prague, 1964, pp. 29–36.
14. P. ERDÖS AND M. SIMONOVITS, *Compactness results in extremal graph theory*, Combinatorica, 2 (1982), pp. 275–288.
15. R. G. GALLAGER, *A minimum delay routing algorithm using distributed computation*, IEEE Transactions on Communications, (1977).
16. P. JACQUET AND L. VIENNOT, *Remote spanners: what to know beyond neighbors*, in 23rd IEEE International Parallel & Distributed Processing Symposium (IPDPS), IEEE Computer Society Press, May 2009.
17. J. KLEINBERG, *Approximation algorithms for disjoint paths problems*, PhD thesis, Massachusetts Institute of Technology, 1996.
18. N. KUSHMAN, S. KANDULA, D. KATABI, AND B. M. MAGGS, *R-bgp: Staying connected in a connected world*, in 4th Symposium on Networked Systems Design and Implementation (NSDI), 2007.
19. S. LEE AND M. GERLA, *Split multipath routing with maximally disjoint paths in ad hoc networks*, in IEEE International Conference on Communications (ICC), vol. 10, 2001, pp. 3201–3205.
20. C. LEVCOPOULOS, G. NARASIMHAN, AND M. SMID, *Efficient algorithms for constructing fault-tolerant geometric spanners*, in 30th Annual ACM Symposium on Theory of Computing (STOC), ACM Press, May 1998, pp. 186–195.
21. S. MUELLER, R. P. TSANG, AND D. GHOSAL, *Multipath routing in mobile ad hoc networks: Issues and challenges*, in Performance Tools and Applications to Networked Systems, Revised Tutorial Lectures [from MASCOTS 2003], 2003, pp. 209–234.
22. A. NASIPURI, R. CASTAÑEDA, AND S. R. DAS, *Performance of multipath routing for on-demand protocols in mobile ad hoc networks*, Mobile Networks and Applications, 6 (2001), pp. 339–349.
23. P. PAN, G. SWALLOW, AND A. ATLAS, *Fast Reroute Extensions to RSVP-TE for LSP Tunnels*. RFC 4090 (Proposed Standard), 2005.
24. D. PELEG AND A. A. SCHÄFFER, *Graph spanners*, J. of Graph Theory, 13 (1989), pp. 99–116.

25. D. PELEG AND J. D. ULLMAN, *An optimal synchronizer for the hypercube*, SIAM J. on Computing, 18 (1989), pp. 740–747.
26. S. PETTIE, *Low distortion spanners*, in 34th International Colloquium on Automata, Languages and Programming (ICALP), vol. 4596 of LNCS, Springer, July 2007, pp. 78–89.
27. ———, *Distributed algorithms for ultrasparse spanners and linear size skeletons*, in 27th Annual ACM Symposium on Principles of Distributed Computing (PODC), ACM Press, Aug. 2008, pp. 253–262.
28. L. RODITTY, M. THORUP, AND U. ZWICK, *Roundtrip spanners and roundtrip routing in directed graphs*, ACM Transactions on Algorithms, 3 (2008), p. Article 29.
29. J. W. SUURBALLE AND R. E. TARJAN, *A quick method for finding shortest pairs of disjoint paths*, Networks, 14 (1984), pp. 325–336.
30. M. THORUP AND U. ZWICK, *Approximate distance oracles*, J. of the ACM, 52 (2005), pp. 1–24.
31. ———, *Spanners and emulators with sublinear distance errors*, in 17th Symposium on Discrete Algorithms (SODA), ACM-SIAM, Jan. 2006, pp. 802–809.
32. S. VUTUKURY AND J. J. GARCIA-LUNA-ACEVES, *A simple approximation to minimum-delay routing*, in SIGCOMM, 1999, pp. 227–238.

A The proof of Theorem 1

Let G be an unweighted multi-edge graph. Since G is unweighted, for any subgraph $H \subseteq G$, we denote by $|H|$ the cost of H , i.e., its number of edges. Consider any 2-iterative 3-spanner H constructed from G , and call H_1 and H_2 the two consisting spanners. Let u, v be two vertices, and P be a minimal 2-path in G which connects u and v . In the following, we use a decomposition of P in two simple edge-disjoint paths P_1 and P_2 . If they share a common vertex w , it is called an *intersection point*. P is said to be simple if P_1 and P_2 are in fact vertex-disjoint, i.e., they do not share a single intersection point.

We first show that any minimal 2-path P can be decomposed in a sequence of simple 2-paths. Let $u_0 = u, u_1, \dots, u_t = v$ denote the intersection points as they are ordered on the path P_1 from u to v .

Lemma 2. *The intersection points u_0, \dots, u_t appear in the same order on P_2 .*

Proof. Suppose they are not in the same order on the two paths. Then there exists $i < j$ such that u_j appears before u_i on P_2 . The sequence on P_1 (resp. P_2) looks like : $u_0, \dots, u_i, \dots, u_j, \dots, u_t$ (resp. $u_0, \dots, u_j, \dots, u_i, \dots, u_t$). Let P'_1 and P'_2 be the paths realized by interleaving P_1 and P_2 at the crossing point u_i , P'_1 (resp. P'_2) having the same prefix as P_1 (resp. P_2). More precisely, $P'_1 = P_1[u_0, u_j] \circ P_2[u_j, u_i] = P_1[u_0, u_i] \circ P_1[u_i, u_j] \circ P_2[u_j, u_i] \circ P_2[u_i, u_t]$ and $P'_2 = P_2[u_0, u_j] \circ P_1[u_j, u_t]$. The transformation is shown in Fig. 2. The 2-path (P'_1, P'_2) has same cost as (P_1, P_2) . However, an explicit loop in P'_1 was introduced. It can be discarded, resulting in an improvement on the total cost which we supposed to be minimal. This is a contradiction.




Fig. 2. Proof of Lemma 2.

Among all 2-paths (P_1, P_2) satisfying $|P_1| + |P_2| = d_G^2(u, v)$, there is at least one which contains a maximum number of intersection points. We now suppose that $P = (P_1, P_2)$ is such a 2-path and that u_0, \dots, u_t now denote the intersections points of P_1 and P_2 as they are ordered on both paths.

Define P_1^i (resp. P_2^i) as the portion of P_1 (resp. P_2) from u_{i-1} to u_i . Let P^i denote the 2-path formed by the union of P_1^i and P_2^i . Note that P^i is simple. P is indeed the union of P^1, \dots, P^t .

We construct a replacement graph F with edges of H according to the rules bellow. The idea is to replace each edge $x - y \notin H$ of P with a shortest path from x to y in H_1 or H_2 , by using rules which guarantee the stretch property while ensuring the biconnectivity.

Consider an edge $e \notin H$. Let $H_1(e)$ (resp. $H_2(e)$) be a shortest path in H_1 (resp. H_2) between the extremities of e . As H_1 is a 3-spanner, we have $|H_1(e)| \leq 3$. As H_2 is a 3-spanner of $G \setminus H_1$ and $e \notin H_1$, we also have $|H_2(e)| \leq 3$.

We define a subgraph F^i of H by applying the following disjoint rules, for each 2-path P^i , in the order which they are presented and for which a schema is shown in Fig. 3:

- R0: If P_1^i or P_2^i is a single edge that belongs to $H_j, j \in \{1, 2\}$, add it to F^i . Then for each edge e of the other path: if it belongs to H add it to F^i , if it does not add the replacement path $H_{(j \bmod 2)+1}(e)$ to F^i .
- R1: If an edge of P^i belongs to H , then add it to F^i .
- R2: For all edges not concerned by R0 nor R1 do:
 - If $e \in P_1^i$ and $H_1(e) \cap P_2^i = \emptyset$, then $F^i = F^i \cup H_1(e)$
 - If $e \in P_2^i$ and $H_2(e) \cap P_1^i = \emptyset$, then $F^i = F^i \cup H_2(e)$
- R3 (disjointly from R2):
 - If $e \in P_1^i$ and $H_1(e) \cap P_2^i \neq \emptyset$ then $F^i = F^i \cup H_1(e) \cup H_2(e)$.
 - If $e \in P_2^i$ and $H_2(e) \cap P_1^i \neq \emptyset$ then $F^i = F^i \cup H_1(e) \cup H_2(e)$.

F is then the union of all $F^i, i \in \{0, \dots, t\}$.

The following lemmatas (3–6) show that the rules enforce that there is always:

- a 2-path between u and v
- the stretch is controlled.

Lemma 3. F contains a 2-path from u to v .

Proof. The proof will show that the capacity of any cut which separates u from v in F is at least two, which will allow us to conclude with a min-cut max-flow argument.

Suppose there is a cut $X \subset V$ such that $u \in X, v \in \overline{X}$ which contains u but not v . Then there exists i such that $u_{i-1} \in X$ and $u_i \in \overline{X}$. Let be e_1 (resp. e_2) an edge in P_1^i (resp. P_2^i) crossing X (i.e., having an extremity in X and the other in \overline{X}).

Several cases are possible:

- R0 has been applied to e_1 or e_2 : then the other edge has a replacement path in F^i , so the cut is at least two in F .




Fig. 3. The different cases for the rules are shown here for an edge from P_1^i . In blue are shown the edges and paths belonging to H_1 , and in red those from H_2 .

- R1 was applied to both e_1 and e_2 : then each of these edges belong to F^i , are disjoint, so the cut is at least two in F .
- R2 was applied to both e_1 and e_2 : then each of these edges have a disjoint replacement path in F^i , so the cut is at least two in F .
- R1 was applied to one of $\{e_1, e_2\}$, and R2 to the other: then as one of the edges is in F^i and the other has a replacement path in F^i , disjoint from the other (if it wasn't the case, R3 would have been applied) there are at least two edges from F^i crossing X .
- R3 was applied to either e_1 or e_2 : then there are two disjoint replacement paths for either e_1 or e_2 , which are cut by X , so it is at least two in F .

As the cut is at minimum two in F , by the min-cut max-flow theorem the flow between u and v is at least two in F .

Let be $|P^i| = x_0 + x_1 + x_2 + x_3$ the number of edges of P^i , x_j being the number of edges from P^i where rule R_j was applied.

Lemma 4. $|F^i| \leq 3x_0 + x_1 + 3x_2 + 5x_3$.

- Proof.* – For rule R0 it is easy to show that the number of edges added in F does not exceed $3x_0$ (one path is a single edge which is added, replacement paths of length at most 3 are added for the edges of the other path).
- For rule R1, the number of edges added in F is exactly x_1 , as these edges are in H .
 - For rule R2, the number of edges added is $3x_2$, because the spanners H_1 and H_2 have stretch 3.

- For rule R3, there are at most 5 edges added for each application of the rule because one is already part of F .

Lemma 5. $x_3 \leq x_1$.

Proof. We show that an edge added according to rule R1 can be generally in at most one replacement path of an edge considered under rule R3, thus implying $x_3 \leq x_1$. Suppose by contradiction that this is not the case. W.l.o.g. there exists $e_1, e'_1 \in P_1^i, e_2 \in P_2^i$ such that R1 applies to e_2 , and that R3 applies to e_1 and e'_1 . We prove that this may happen only in special cases where $x_3 \leq x_1$ is still satisfied. Let a, b, c, d be the paths such that $a - e_2 - b = H_1(e_1)$ and $c - e_2 - d = H_1(e'_1)$, $e_2 \in H_1(e) \cap H_1(e'_1)$. (The paths a, b, c, d have length 0, 1 or 2 since H_1 is a 3-spanner.). There are two possible cases as shown in Fig. 4:

1. either a and c are connected together on one endpoint of e_2 — and then b and d are connected to the other end;
2. or a and d to one endpoint of e_2 and c and b to the other.

Let e be the sub-path of P_1^i which lies in between e_1 and e'_1 .




Fig. 4. Two possible cases.

In Case 1, we first show $|e| > 0$. Suppose by contradiction that e is empty. Then $a-c$ (resp. $b-d$) is a replacement path for e_1 (resp. e'_1) in H_1 . As $H_1(e_1)$ and $H_1(e'_1)$ are shortest paths, we have $|a|+|c| \geq |a|+1+|b|$ and $|b|+|d| \geq |c|+1+|d|$, i.e., $|c| > |b|$ and $|b| > |c|$ which is a contradiction.

Now, we show that each of the paths a, b, c, d is composed of a single edge. As we are not in the case of rule R0, a and d cannot be both empty. Suppose w.l.o.g. $|a| > 0$. $|H_1(e_1)| \leq 3$ implies $|b| \leq 1$. Indeed, we have $|b| = 1$ as b cannot be empty since there are no intersection point of P_1 and P_2 in between e_1 and e'_1 . For the same reason, b cannot be composed of an edge of P_2 . As a consequence, d cannot be empty otherwise b would be a shortcut violating the minimality of the cost of P^i . We then similarly show that c is composed of a single edge which

is not in P_2 . As $|c| + 1 + |d| \leq 3$ and $|d| > 0$, we obtain $|d| = 1$. Similarly, we have $|a| = 1$.

We now prove that a is in P_2 . If this is not the case, then $a - c$ is a shortcut that can be substituted to $e_1 - e$ in P_1 . This does not increase the cost of the 2-path and it increases the number of intersection points, a contradiction with the choice of P . Similarly, d is in P_2 .

Let us recall what we have obtain so far for Case 1: a and d are edges of P_2 , b and c are edges not in P_2 . P_2^i is thus the path $a - e_2 - d$ and P_1^i is $e_1 - e - e'_1$. Note that $x_1 \geq 3$ since the three edges of P_2^i follow rule R1. We then show that at most one edge of e may fall under rule R3, yielding $x_3 \leq 3 \leq x_1$. Consider an edge e'_1 of e falling under rule R3. We write $e = e' - e'_1 - e''$ where e' and e'' are sub-paths of e . We can write $H_1(e'_1) = b' - e'_2 - c'$ where b' and c' are edges of H_1 and where e'_2 is either an empty path or an edge. As e'_1 follow rule R3, e'_2 must be an edge of P_2^i ($H_1(e'_1)$ must have length 3).

First consider the case where e'_2 is a . Then, e' must be empty (e'_1 must be the first edge after e_1). W.l.o.g. b' is e_1 and c' contains $a \cap e_2$. We must have $|e''| \leq 1$. Otherwise $c' - c$ is a shortcut violating the choice of P . If e'' is an edge, it cannot follow rule R3 as $c' - c$ is a replacement path in H_1 with length 2 and $H_1(e'')$ cannot have length 3. We thus have $x_3 \leq 3 \leq x_1$. The case where e'_2 is d can be treated similarly.

Now consider the case where e'_2 is e_2 . W.l.o.g. b' contains $b \cap e_2$ and c' contains $c \cap e_2$. As $b - b'$ cannot be a shortcut violating the choice of P , we have $|e'| \leq 1$. If e' is an edge, it cannot follow rule R3 as $b - b'$ is a replacement path in H_1 with length 2. Similarly $c - c'$ cannot be a shortcut violating the choice of P . We thus have $|e''| \leq 1$ and if e'' is an edge, it cannot follow rule R3. We again obtain $x_3 \leq 3 \leq x_1$.

In Case 2, we first prove that a, b, c, d are single edges. b and c cannot be empty as there are no intersection point in between e_1 and e'_1 . We thus have $|a| \leq 1$ and $|d| \leq 1$. If a is empty, then d is a shortcut violating the minimality of the cost of P . We thus have $|a| = 1$ which implies $|b| = 1$ as $|H_1(e_1)| \leq 3$. Similarly, we have $|d| = 1$ and $|c| = 1$. Note that b and c cannot be in P_2 as there are no intersection point in between e_1 and e'_1 .

We show that if we are not in Case 1 again, then we have two edges following rule R1 contained in at most two replacement paths of edges following rule R3. As $b - c$ cannot be a shortcut violating the choice of P , we have $|e| \leq 1$. If e is an edge, it admits $b - c$ as a replacement path in H_1 . Thus it cannot follow rule R3. The path $a - d$ cannot be a shortcut violating the choice of P . This implies that a or d is in P_2 . W.l.o.g., suppose that a is in P_2 ($e_1 \cap a$ is indeed u_{i-1}). d cannot be in P_2 as e_2 is in P_2^i . Consider another edge e'_1 falling under rule R3. $H_1(e'_1)$ cannot contain a as we would get $e'_1 = e$ which cannot follow rule R3. If $H_1(e'_1)$ contains e_2 , we fall back into Case 1.

Case 2 may occur both at u_{i-1} and u_i . All other edges falling under rule R1 are contained in at most one replacement path of an edge following R3. We thus have $x_3 \leq x_1$ in any case.

Lemma 6. F has at most $3d_G^2(u, v)$ edges.

Proof. As $|F^i| \leq 3x_0 + x_1 + 3x_2 + 5x_3$ and $x_3 \leq x_1$,

$$|F^i| \leq 3x_0 + x_1 + 3x_2 + 2x_3 + 3x_3 \leq 3x_0 + 3x_1 + 3x_2 + 3x_3 \leq 3|P^i|$$

As F is the union of all F^i , F has a maximum weight of $3d_G^2(u, v)$.