

# Greedy spanners are optimal in doubling metrics

Glencora Borradaile

Hung Le

Christian Wulff-Nilsen

## Abstract

We show that the greedy spanner algorithm constructs a  $(1+\epsilon)$ -spanner of weight  $\epsilon^{-O(d)}w(\text{MST})$  for a point set in metrics of doubling dimension  $d$ , resolving an open problem posed by Gottlieb [11]. Our result generalizes the result by Narasimhan and Smid [15] who showed that a point set in  $d$ -dimension Euclidean space has a  $(1+\epsilon)$ -spanner of weight at most  $\epsilon^{-O(d)}w(\text{MST})$ . Our proof only uses the packing property of doubling metrics and thus implies a much simpler proof for the same result in Euclidean space.

## 1 Introduction

For a real value  $t \geq 1$ , a  $t$ -spanner of an edge-weighted graph  $G$  is a subgraph  $S$  such that  $d_G(x, y) \leq d_S(x, y) \leq t \cdot d_G(x, y)$  where  $d_G$  and  $d_S$  denote the shortest path distance functions for vertex pairs in  $G$  and  $S$ , respectively. In this work, we study  $t$ -spanners of a metric space; more precisely, we study  $t$ -spanners of the corresponding *metric graph* which is the complete graph on the set of points in the metric where the weight of each edge  $pq$  denotes the metric distance between  $p$  and  $q$ . When we refer to metric spaces in the following, we assume them to be finite.

Spanners have been used in many applications including distributed systems, communication networks, robotics and more [1, 15]. In this work, we are interested in  $(1+\epsilon)$ -spanners (herein referred to simply as *spanners*) of geometric and metric graphs where  $\epsilon < 1$  denotes a fixed constant. One way to measure the quality of spanners is by their *lightness* which is the ratio of the weight of spanner's edges to the weight of a minimum spanning tree of  $G$ . One prominent application of spanners with constant lightness is in designing faster polynomial time approximation schemes<sup>1</sup> (PTAS) for the TSP problem [4, 5, 11, 13, 16]. A classical spanner algorithm [1] constructs a  $(1+\epsilon)$ -spanner of a graph  $G$  by considering edges in non-decreasing order of weight and adding the current edge  $pq$  to the spanner if there is not already a  $p$ -to- $q$  path of weight at most  $(1+\epsilon)$  times the weight of  $pq$ ; the resulting spanner is called the *greedy  $(1+\epsilon)$ -spanner* of  $G$ .

The doubling dimension [3, 12] of a metric space is the smallest  $d$  such that every ball of radius  $r$  is covered by  $2^d$  balls of radius most  $\frac{r}{2}$ . In this work, we show that:

**Theorem 1.** *The greedy  $(1+\epsilon)$ -spanner of a metric space of doubling dimension  $d$  has lightness  $\epsilon^{-O(d)}$ .*

Geometric spanners have a rich history. In 2-dimensional Euclidean space,  $O(1)$ -spanners of lightness  $O(1)$  have been known since the late 80s [1, 14]. Das, Heffernan and Narasimhan [7] sketched an intricate argument showing that  $t$ -spanners for any fixed  $t > 1$  in 3-dimensional Euclidean space have lightness  $O(1)$ . Their main contribution is an analysis of the *leap-frog property* of

---

<sup>1</sup>A polynomial-time approximation scheme is an algorithm which, for a fixed error parameter  $\epsilon$ , finds a solution whose value is within  $1 \pm \epsilon$  of optimal in polynomial time.

the spanner edges found by the greedy algorithm. Later, Das, Narasimhan and Salowe [8] sketched a generalization of the proof by Das, Heffernan and Narasimhan [7] to show that  $t$ -spanners for any fixed  $t > 1$  in  $d$ -dimensional Euclidean space have lightness  $O(1)$ . However, the dependency of the lightness on  $t$  and  $d$  was not explicitly computed. Rao and Smith [16] redid the analysis of Arya, Das, Mount, Salowe and Smid [2] to show that the constant in the work of Das, Heffernan and Narasimhan [7] is  $(\frac{d}{\epsilon})^{O(d)}$ . Narasimhan and Smid [15] devoted a 60-page chapter of *Geometric Spanner Networks* to give full details of the analysis of the lightness of greedy spanners. They show that greedy spanners for  $d$ -dimensional Euclidean space have lightness  $\epsilon^{-O(d)}$ ; their proof heavily relies on the geometry of Euclidean space. Our Theorem 1 immediately implies a simpler proof for the same lightness bound in Euclidean space; it is well-known that a point set in  $d$ -dimensional Euclidean metric has doubling dimension  $\Theta(d)$ . Instead of relying on the leap-frog property as in previous works, which is not easy to analyze in doubling metrics, we only use the simple packing property of the doubling metrics where Euclidean space is a special case.

Spanners in doubling metrics were first considered by Gao, Guibas and Nguyen [10] who showed that a  $n$ -point set in doubling dimension  $d$  has a spanner of  $\epsilon^{-O(d)}n$  edges. By analyzing the greedy algorithm, Smid [17] showed that greedy spanners have  $O(n)$  edges and  $O(\log n)$  lightness. Beating the  $O(\log n)$  lightness bound of Smid [17] had been an important open problem until the recent work by Gottlieb [11], who showed that a metric of doubling dimension  $d$  has a spanner of lightness  $(\frac{d}{\epsilon})^{O(d)}$ . We note that the construction of Gottlieb [11] is non-greedy, conceptually involved and takes  $O(n \log^2 n)$  time. However, two questions remain open. First, can we design a spanner of lightness  $\epsilon^{-O(d)}$  to match the bound in Euclidean space? Second, is there a more refined analysis of the greedy algorithm to achieve the bound  $\epsilon^{-O(d)}$ ? Gottlieb [11] asked the first question in his paper. The second question was partially addressed by Filtser and Solomon [9], who showed that greedy spanners in doubling metrics (as well as graph classes closed under edge removal) are *existentially optimal*: if there is a spanner construction of lightness bound  $l(\epsilon, d)$ , then greedy spanners have lightness  $O(l(\epsilon, d))$ . Combined with Gottlieb's results [11], the existential optimality implies that greedy spanners have  $(\frac{d}{\epsilon})^{O(d)}$  lightness. In this paper, we resolve both questions affirmatively by presenting a refined and comparatively simple analysis of the greedy algorithm. Our result, in combination with the result of Filtser and Solomon [9], implies an  $O(n \log n)$  time algorithm to find a spanner of lightness  $\epsilon^{-O(d)}$ .

## 1.1 Techniques

Our analysis is built primarily upon our result [5] in minor-free graphs which in turn is based on the techniques of Chechick and Wulff-Nilsen [6] for general graphs. We briefly review the  $H$ -minor-free techniques [5] here, highlighting new ideas required for doubling metrics. The first step is to reduce the problem on the input graph to the same problem on graphs that have unit-edge-weight minimum spanning trees (MSTs), by rounding small-weight edges and subdividing large-weight edges. Then the greedy algorithm is applied to this slightly modified graph making the construction non-greedy as a whole. However, in the setting of doubling metrics, we cannot use the same simplification since rounding changes the metric. We instead directly analyze the greedy spanner of the input.

To analyze the spanner of an  $H$ -minor free graph, [5] uses iterative clustering. Spanner edges are partitioned into  $\log \frac{1}{\epsilon}$  sets<sup>2</sup>, and then the total weight of each set is bounded separately; this induces the  $\log \frac{1}{\epsilon}$  factor in the lightness. Each set consists of spanner edges in an exponential scale

---

<sup>2</sup>Here,  $\log$  denotes the base 2 logarithm.

of many *levels*. First, a non-negative credit  $c(\epsilon)$  is assigned to each MST edge of unit weight;  $c(\epsilon) \log \frac{1}{\epsilon}$  is also the lightness of the spanner. In each level, clusters are constructed iteratively from clusters of the previous level; level-1 clusters are constructed directly from the MST. An invariant is maintained that each cluster must have some amount of credit to pay for spanner edges in their level. Credits of level-1 clusters are taken directly from MST edges. Credits of level- $i$  clusters are taken from credits of clusters of level  $i - 1$  and MST edges connecting those lower-level clusters. However, to pay for the spanner edges, level- $i$  clusters cannot take all credits from level- $(i - 1)$  clusters. Instead, it is guaranteed that on average, each level- $(i - 1)$  cluster has a non-trivial amount of credit left to pay for spanner edges. The minor-free property is then used to argue that each cluster on average must pay for only a constant number of spanner edges in each level.

Our new argument is also based on our earlier iterative cluster construction. However, we rely on the packing property of doubling metrics (defined below) to show that each cluster needs to pay for a constant number of spanner edges in the same level. This property of doubling metric spanners allows us in fact to simplify the cluster construction that was used for  $H$ -minor free graphs.

Let  $G(V, E)$  be the graph representing a metric of doubling dimension  $d$ . For each edge  $e \in E$ , we define the weight function  $w(e)$  to be the distance between its endpoints in the metric. Let  $n = |V(G)|$  and  $m = |E(G)|$ . We directly analyze the spanner produced by the greedy algorithm. For a review of the greedy spanner algorithm, see Appendix B. Let  $S$  be the greedy spanner of  $G$ . Smid [17] showed (in two pages) that:

**Lemma 2.**  $|E(S)| \leq \delta(\epsilon)n$  where  $\delta(\epsilon) = \epsilon^{O(-d)}$ .

The following packing property of doubling metric is well-known (see [17]):

**Lemma 3** (Packing property). *A point set  $X$  of a metric of doubling dimension  $d$  that is contained in a ball of radius  $R$  and for every  $x \neq y \in X$ ,  $d(x, y) > r$ , has  $|X| \leq \left(\frac{4R}{r}\right)^d$ .*

## 2 Assigning credits to MST edges

Let  $w_0 = \frac{w(\text{MST})}{n-1}$  be the average weight of an MST edge. We first bound the total weight of edges that have weight at most  $w_0$ .

**Claim 4.** *Let  $L_S$  be the set of edges of  $S$  of weight at most  $w_0$ . Then,  $w(L_S) \leq 2\delta(\epsilon)w(\text{MST})$ .*

*Proof.*  $w(L_S) \leq w_0|L_S| \stackrel{\text{Lemma 2}}{\leq} w_0\delta(\epsilon)n = \frac{w(\text{MST})}{n-1}\delta(\epsilon)n \leq 2\delta(\epsilon)w(\text{MST})$ . □

We now focus on bounding the total weight of edges of weight at least  $w_0$  in  $S$ . We subdivide and allocate credits to MST edges such that every MST edge has weight at most  $w_0$  and at least  $c(\epsilon)w_0$  credits where  $c(\epsilon)$  is a constant that only depends on  $\epsilon$  and will be specified later. We will guarantee that the total allocated credit is  $O(c(\epsilon))w(\text{MST})$  where  $O(c(\epsilon))$  is also the lightness of the spanner. First, we subdivide every MST edge  $e$  of weight more than  $w_0$  into  $\lceil \frac{w(e)}{w_0} \rceil$  new edges with equal weights summing up to  $w(e)$ ; note that each new edge has weight at most  $w_0$ . Letting  $S'$  be the new graph, we have  $w(\text{MST}(S')) = w(\text{MST})$ . We then allocate  $c(\epsilon)w_0$  credits to each MST edge of  $S'$ .

**Claim 5.** *The total credit allocated to the MST edges of  $S'$  is at most  $2c(\epsilon)w(\text{MST})$ .*

*Proof.* The total credits assigned to MST edges of  $S'$  is:

$$\begin{aligned} c(\epsilon)w_0|E(\text{MST}(S'))| &\leq c(\epsilon)w_0 \sum_{e \in \text{MST}} \lceil \frac{w(e)}{w_0} \rceil \leq c(\epsilon)w_0 \left( \sum_{e \in \text{MST}} \left( \frac{w(e)}{w_0} + 1 \right) \right) \\ &= c(\epsilon)w(\text{MST}) + c(\epsilon)w_0(n-1) = 2c(\epsilon)w(\text{MST}) \end{aligned} \quad (1)$$

□

### 3 Iterative Clustering

Let  $J_0 = \{e \in S', w_0 < w(e) \leq \frac{2w_0}{\epsilon}\}$ . We first bound the weight of  $J_0$  and pay for edges in  $J_0$  separately. The purpose is to simplify the base case in the inductive amortized argument that we present below.

**Claim 6.**  $w(J_0) \leq \frac{4\delta(\epsilon)}{\epsilon}w(\text{MST})$ .

*Proof.*

$$w(J_0) = \sum_{e \in J_0, w(e) > w_0} w(e) \stackrel{\text{Lemma 2}}{\leq} \delta(\epsilon)n \frac{2w_0}{\epsilon} \leq \frac{4\delta(\epsilon)}{\epsilon}w(\text{MST}) \quad \square$$

Let  $I_\epsilon = \lceil \log \frac{1}{\epsilon} \rceil$  and  $I_n = \lceil \log n \rceil$ . Note that the longest distance between any two vertices in  $S'$  is at most  $n \cdot w_0$ . We partition the spanner edges (of weight at least  $w_0$ ) of  $S'$  into  $I_n \cdot I_\epsilon$  sets  $\{\Pi_i^j, 0 \leq i \leq I_n - 1, 0 \leq j \leq I_\epsilon - 1\}$  where each edge  $e \in \Pi_i^j$  has weight in the range  $(\frac{2^j}{\epsilon^i}w_0, \frac{2^{j+1}}{\epsilon^i}w_0]$ . For each  $0 \leq j \leq I_\epsilon - 1$ , let

$$S_j = \bigcup_{i=0}^{I_n-1} \Pi_i^j \quad (2)$$

**Lemma 7.** For each  $0 \leq j \leq I_\epsilon - 1$ , there is a set of spanner edges  $B$  such that  $w(B) \leq \epsilon^{-O(d)} \cdot w(\text{MST})$  and  $w(S_j \setminus B) \leq \epsilon^{-O(d)}w(\text{MST})$ .

It is not hard to see that Lemma 7 directly implies Theorem 1. Thus, we only focus on proving Lemma 7 for a fixed  $j$ . We refer to edges of  $\Pi_i^j$  as edges in *level  $i$*  (Equation 2). Let  $\ell_i = \frac{2^{j+1}}{\epsilon^i}w_0$ . Similar to our analysis for  $H$ -minor free graphs [5], we construct a set of clusters, which are subgraphs of  $S'$ , for each level and guarantee inductively two diameter-credits invariants:

**DC1** A cluster of level  $i$  of diameter  $k$  has at least  $c(\epsilon) \cdot \max\{k, \frac{\ell_i}{2}\}$  credits.

**DC2** A cluster of level  $i$  has diameter at most  $g\ell_i$  for some constant  $g > 2$  (specified later).

A cluster of level  $i$ , say  $C_i$ , is the union of a subset of clusters in level  $i-1$  connected by MST and level- $i$  spanner edges. Clusters of level  $i-1$  are referred to as  $\epsilon$ -clusters. To satisfy DC1, we assign the credits from  $\epsilon$ -clusters in  $C_i$  and the MST edges connecting the  $\epsilon$ -clusters to  $C_i$ . However, we need to group  $\epsilon$ -clusters in such a way that there are some extra  $\epsilon$ -clusters whose credits are not needed to maintain DC1 for  $C_i$ . We will use credits of these extra  $\epsilon$ -clusters to pay for level- $i$  spanner edges incident to every  $\epsilon$ -cluster in  $C_i$ . The credit lower bound  $c\ell_i/2$  (DC1) helps us achieve the goal.

To guarantee the diameter-credit invariants for level 0, we greedily break the MST into components (level-0 clusters) of diameter at least  $\ell_0$  and at most  $4\ell_0$ . Recall  $\ell_0 = 2^{j+1}w_0 \leq \frac{2w_0}{\epsilon}$ . To guarantee DC1, we use the credits of MST edges in the longest path of each cluster. Since the credit of each MST edge is at least its length, DC1 is satisfied. Invariant DC2 follows directly from the construction. Note that we have already accounted for the weight of spanner edges of  $E_0$  in Claim 6.

### 3.1 Constructing higher level clusters

We construct clusters of level  $i$  from the  $\epsilon$ -clusters of level  $i-1$ . We assume that the stretch of the spanner is  $1+s\epsilon$  for some constant  $s$  (independent of  $\epsilon$ ) that we will pick sufficiently big to make our claims below hold. Furthermore, we assume that  $\epsilon$  is bounded from above by a sufficiently small positive constant. We call vertices of  $V(S') \setminus V(S)$  *virtual vertices*. We call a cluster *virtual* if it only contains virtual vertices and *non-virtual* otherwise. Let  $\mathcal{K}(\mathcal{C}_\epsilon, E_i)$  be the multigraph obtained by taking the subgraph of  $G$  consisting of  $\epsilon$ -clusters and spanner edges in  $E_i$  and contracting each  $\epsilon$ -cluster into a single vertex. Let  $\ell = \ell_i$ .

**Lemma 8.**  $\mathcal{K}(\mathcal{C}_\epsilon, E_i)$  is a simple graph of degree  $\epsilon^{-O(d)}$ .

*Proof.* We leave the details of the proof that  $\mathcal{K}(\mathcal{C}_\epsilon, E_i)$  is simple to Appendix C. To show the degree bound, first note that virtual clusters are isolated vertices in  $\mathcal{K}(\mathcal{C}_\epsilon, E_i)$  since virtual vertices are subdividing vertices incident to edges of weight at most  $w_0$ . For each non-virtual  $\epsilon$ -cluster, we designate a non-virtual vertex to be its center. Assuming w.l.o.g. that  $\epsilon \leq \frac{1}{4}$  and picking  $s \geq 12g+4$ , we get the following claim:

**Claim 9.** Let  $C_1, C_2, C_3$  be three  $\epsilon$ -clusters that have  $x_1, x_2, x_3$  as centers. Suppose that  $C_2, C_3$  are neighbors of  $C_1$  in  $\mathcal{K}$ . Then,  $d_G(x_i, x_j) > \epsilon\ell$  for any  $1 \leq i \neq j \leq 3$ .

*Proof.* Suppose  $y_1y_3$  and  $z_1z_2$  are two level- $i$  spanner edges such that  $y_1, z_1 \in C_1, z_2 \in C_2, y_3 \in C_3$ . We assume, w.l.o.g., that  $w(y_1y_3) \leq w(z_1z_2)$ . Recall  $\ell/2 < w(y_1y_3), w(z_1z_2) \leq \ell$ .

We only present the proof showing that  $d_G(x_2, x_3) > \epsilon\ell$  since a similar but simpler proof holds for  $d_G(x_1, x_2)$  and  $d_G(x_1, x_3)$ . We assume that  $d_G(x_2, x_3) \leq \epsilon\ell$ . Let  $Q$  be the  $z_1$ -to- $z_2$  path that consists of: (i) a shortest  $z_1$ -to- $y_1$  subpath in  $C_1$ , (ii) edge  $y_1y_3$ , (iii) a shortest  $y_3$ -to- $x_3$  subpath in  $C_3$ , (iv) a shortest  $x_3$ -to- $x_2$  path in  $S$  and (v) a shortest  $x_2$ -to- $z_2$  subpath in  $C_2$ . Since in the greedy spanner, edges are added by increasing weight, by the time  $z_1z_2$  is added,  $d_S(x_2, x_3) \leq (1+s\epsilon)d_G(x_2, x_3) < (1+s\epsilon)\epsilon\ell$ . Thus, we have:

$$\begin{aligned} w(Q) &\leq g\epsilon\ell + w(y_1y_3) + g\epsilon\ell + (1+s\epsilon)\epsilon\ell + g\epsilon\ell \\ &\leq (3g+1+s\epsilon)\epsilon\ell + w(y_1y_3) \\ &\leq 2(3g+1+s\epsilon)\epsilon w(z_1z_2) + w(z_1z_2) \\ &= ((6g+2+2s\epsilon)\epsilon+1)w(z_1z_2) \end{aligned}$$

Since  $\epsilon \leq \frac{1}{4}$ , we have  $s \geq 12g+4 \geq (6g+2)/(1-2\epsilon)$  and hence  $6g+2+2s\epsilon \leq s$ . But then  $w(Q) \leq (1+s\epsilon)w(z_1z_2)$ ; contradicting that  $z_1z_2$  is a spanner edge.  $\square$

Let  $C_0$  be an  $\epsilon$ -cluster with neighbors  $C_1, C_2, \dots, C_p$  in  $\mathcal{K}$ . Let  $X = \{x_0, x_1, \dots, x_p\}$  where  $x_i$  is the center of  $C_i, 0 \leq i \leq p$ . We show below that  $d_G(x_0, x_i) \leq 3\ell$  for every  $1 \leq i \leq p$  when  $\epsilon$  is

sufficiently small. Thus,  $X$  is contained in a ball centered at  $x_0$  of radius at most  $3\ell$ . By Claim 9,  $d_G(x_i, x_j) > \epsilon\ell$  for every  $0 \leq i \leq p$ . Thus, by Lemma 3,  $|X| \leq \epsilon^{-O(d)}$ .

We now show that  $d_G(x_0, x_i) \leq 3\ell$  for a fixed  $i$  with  $1 \leq i \leq p$ . Let  $y_0y_i$  be the spanner edge in  $E_i$  connecting  $C_0$  and  $C_i$  in  $\mathcal{K}$ . Then the  $x_0$ -to- $x_i$  path  $P$  consisting of an  $x_0$ -to- $y_0$  shortest path in  $C_0$ , edge  $y_0y_i$  and a  $y_i$ -to- $x_i$  shortest path in  $C_i$  has length at most  $2g\epsilon\ell + \ell \leq 3\ell$  when  $\epsilon$  is smaller than  $1/g$ .  $\square$

Let  $\mathcal{T}$  be a tree of  $\epsilon$ -clusters connected by MST edges. We say an  $\epsilon$ -cluster of  $\mathcal{T}$  is *branching* if it is incident to at least three MST edges in  $\mathcal{T}$ . Let  $\mathcal{P}$  be a path of  $\mathcal{T}$ . We define the diameter of  $\mathcal{P}$ , denoted by  $\text{diam}(\mathcal{P})$ , to be the diameter of the subgraph of  $S'$  formed by edges inside  $\epsilon$ -clusters and MST edges connecting  $\epsilon$ -clusters of  $\mathcal{P}$ . We define *effective diameter* of  $\mathcal{P}$ , denoted by  $\text{ediam}(\mathcal{P})$ , to be the diameters of  $\epsilon$ -clusters in  $\mathcal{P}$ . Since  $\epsilon$ -clusters have diameter at least  $w_0$  (by construction of the base case) which is at least the weight of edges connecting them in  $\mathcal{P}$ , we have:

**Observation 10.**  $\text{diam}(\mathcal{P}) \leq 2\text{ediam}(\mathcal{P})$ .

We define the effective diameter of a subtree  $\mathcal{T}'$  of  $\mathcal{T}$  to be the effective diameter of the diameter path of  $\mathcal{T}'$ . We construct clusters in four phases:

**Phase 1: Branching  $\epsilon$ -clusters.** We have two steps. Since our construction is recursive, we update the set branching vertices of  $\mathcal{T}$  after each recursive step. (Step 1) Let  $\mathcal{T}'$  be a subtree of  $\mathcal{T}$  of effective diameter at least  $\ell$  and at most  $2\ell$  that contains a branching vertex  $X$  and its neighbors so that  $X$  is still branching in  $\mathcal{T}'$ . We group  $\epsilon$ -clusters and MST edges of  $\mathcal{T}'$  as a new level- $i$  cluster. We remove  $\mathcal{T}'$  from  $\mathcal{T}$  and repeat until every component of  $\mathcal{T}$  either has effective diameter less than  $\ell$  or is a path of  $\epsilon$ -clusters, called a *cluster path*, of effective diameter at least  $\ell$ .

(Step 2) Let  $\mathcal{P}$  be a cluster path of diameter at least  $\ell$ . Let  $X$  be the set of internal  $\epsilon$ -clusters of  $\mathcal{P}$  such that  $X$  has at least one MST edge, say  $e$ , to a level- $i$  cluster, say  $C$ , formed in Step 1. Observe that  $X$  is branching before execution of Step 1 and the removal of subtrees of  $\mathcal{T}$  in Step 1 reduces degree of  $X$  to 2. We augment  $C$  with  $X$  and  $e$ . We then remove  $X$  from  $\mathcal{P}$  and repeat until every cluster path of effective diameter at least  $\ell$  only has MST edges to level- $i$  clusters incident to its endpoint  $\epsilon$ -clusters.

**Phase 2:  $\epsilon$ -clusters in high diameter paths.** Let  $e$  be a spanner edge in  $E_i$  whose endpoints,  $x$  and  $y$  are in high-diameter cluster paths,  $\mathcal{P}$  and  $\mathcal{Q}$ , respectively, where it may be that  $\mathcal{P} = \mathcal{Q}$ . Let  $C_x$  and  $C_y$  be the  $\epsilon$ -clusters containing  $x$  and  $y$ , respectively. We only proceed with this phase if the two affix cluster subpaths of  $\mathcal{P}$  ending at  $C_x$  have effective diameter at least  $\ell$  (likewise for  $\mathcal{Q}$ ). Let  $\mathcal{P}_1$  and  $\mathcal{P}_2$  be the two minimal subpaths of  $\mathcal{P}$  ending at  $C_x$  that have effective diameter at least  $\ell$ . Likewise define  $\mathcal{Q}_1$  and  $\mathcal{Q}_2$ . We group  $\epsilon$ -clusters and MST edges of  $\mathcal{P}_1 \cup \mathcal{Q}_1 \cup \mathcal{P}_2 \cup \mathcal{Q}_2$  and  $e$  as a new level- $i$  cluster. See Figure 1 for an illustration of the different forms this cluster can take.

**Phase 3: Low diameter components.** Let  $\mathcal{F}$  be the set of trees (and paths) remaining of effective diameter at most  $\ell$ . By construction, each component  $\mathcal{T}'$  of  $\mathcal{F}$  has a MST edge, say  $e$ , to a level- $i$  cluster constructed in previous phases, say  $C$ . We augment  $C$  by  $\mathcal{T}'$  and  $e$ .




Figure 1: Three different forms that a cluster (enclosed in the dotted red curves) in Phase 2 can take. The solid blue line is the spanner edge  $e$ . (a)  $e$  connects  $\epsilon$ -clusters in different clusters paths, (b)  $e$  connects  $\epsilon$ -cluster in the same path and  $\mathcal{P}_1$  and  $\mathcal{Q}_1$  are disjoint and (c)  $e$  connects  $\epsilon$ -cluster in the same path and  $\mathcal{P}_1$  and  $\mathcal{Q}_1$  are overlapped. In case (c), we redefine  $\mathcal{P}_1 = \mathcal{Q}_1 = \mathcal{P}_{xy}$ .

**Phase 4: Remaining high diameter paths.** Let  $\mathcal{P}$  be a cluster path of effective diameter at least  $\ell$ . We greedily break  $\mathcal{P}$  into subpaths of effective diameter at least  $\ell$  and at most  $2\ell$ . If any affix of  $\mathcal{P}$ , say  $\mathcal{P}'$ , has a MST edge, say  $e$ , to a level- $i$  cluster constructed in previous phases, say  $C$ , we augment  $C$  with  $\mathcal{P}'$  and  $e$ . We then make each remaining cluster subpath of  $\mathcal{P}$  into a new level- $i$  cluster.

This completes the cluster construction for level  $i$ .

### 3.2 Showing diameter-credit invariant DC2

By construction, each level- $i$  cluster constructed in Phase 4 is a cluster path of effective diameter at most  $2\ell$ . By Observation 10, we have:

**Claim 11.** *Level- $i$  clusters constructed in Phase 4 have diameter at most  $4\ell$ .*

**Claim 12.** *Level- $i$  clusters have diameter at most  $33\ell$  when  $\epsilon$  is smaller than  $\frac{1}{g}$ .*

*Proof.* Let  $C$  be a level- $i$  cluster that is initially formed in Phase 1 or 2. By construction,  $C$  may be augmented in Phases 3 and 4. Let  $C'$  and  $C''$  be the augmented clusters of  $C$  after Phase 3 and Phase 4, respectively. It could be that  $C = C' = C''$ .  $C'$  is obtained from  $C$  by attaching trees of effective diameter at most  $\ell$  via MST edges.  $C''$  is obtained from  $C'$  by attaching trees of effective diameter at most  $2\ell$  via MST edges. Recall each MST edge has length at most  $w_0$ . By Observation 10, we have:

$$\text{diam}(C') \leq \text{diam}(C) + 4\ell + 2w_0 \quad \text{and} \quad \text{diam}(C'') \leq \text{diam}(C') + 8\ell + 2w_0 \quad (3)$$

If  $C$  is constructed in Phase 1, by Observation 10, after Step 1,  $\text{diam}(C) \leq 4\ell$ . Since in Step 2,  $C$  is augmented by  $\epsilon$ -clusters via MST edges, after Step 2,  $\text{diam}(C) \leq 4\ell + 2w_0 + 2g\epsilon\ell \leq 8\ell$  ( $\ell \geq w_0$  by construction of the base case). If  $C$  is constructed in Phase 2, we have:

$$\text{diam}(C) \leq \text{diam}(\mathcal{P}_1) + \text{diam}(\mathcal{P}_2) + \text{diam}(\mathcal{Q}_1) + \text{diam}(\mathcal{Q}_2) + \ell(e)$$

Since  $\mathcal{P}_1, \mathcal{P}_2, \mathcal{Q}_1, \mathcal{Q}_2$  are minimal, each has effective diameter at most  $\ell + g\epsilon\ell$ . Thus,  $\text{diam}(C) \leq 4(2\ell + 2g\epsilon\ell) + \ell = 9\ell + 9g\epsilon\ell \leq 17\ell$ . Thus, in both cases,  $\text{diam}(C) \leq 17\ell$ . By Equation (3),  $\text{diam}(C'') \leq 29\ell + 4w_0 \leq 33\ell$ .  $\square$

Thus, by Claim 12, we can choose  $g = 33$ .

### 3.3 Showing diameter-credit invariant DC1

Let  $\Delta_{\mathcal{K}}$  be the maximum degree of the cluster graph  $\mathcal{K}$ . By Lemma 8,  $\Delta_{\mathcal{K}} = \epsilon^{-O(d)}$ . We define  $cr(\mathcal{X})$  to be the total credits of a set of  $\epsilon$ -clusters  $\mathcal{X}$ .

#### 3.3.1 Clusters originating in Phase 4

Let  $C$  be a level- $i$  cluster formed in Phase 4. We call  $C$  a *long cluster* if it has at least  $\frac{2g}{\epsilon} + 1$   $\epsilon$ -clusters and a *short cluster* otherwise. We have:

**Claim 13.** *A long cluster can both maintain invariant DC1 and pay for its incident spanner edges when  $c(\epsilon) = \epsilon^{-O(d)}$ .*

*Proof.* Let  $\mathcal{X}$  be a set of any  $\frac{2g}{\epsilon}$   $\epsilon$ -clusters of  $C$ . By invariant DC1 for level  $i - 1$ , we have:

$$cr(\mathcal{X}) \geq \frac{2g}{\epsilon} c(\epsilon) \ell / 2 = c(\epsilon) g \ell$$

which is at least  $c(\epsilon) \cdot \max(\text{diam}(C), \ell/2)$  since  $\text{diam}(C) \leq g\ell$  as shown in Claim 11 (since  $g = 33$ ). Thus, credits of  $\mathcal{X}$  are enough to maintain DC1 for  $C$ .

Since  $C$  is a long cluster, there is at least one  $\epsilon$ -cluster, say  $Y$ , not in  $\mathcal{X}$ . By DC1 for level  $i - 1$ ,  $Y$  has at least  $c(\epsilon)\ell/2$  credits. Since there are at most:

$$\Delta_{\mathcal{K}} \cdot \left( \frac{2g}{\epsilon} + 1 \right) = \epsilon^{-O(d)}$$

level- $i$  spanner edges incident to  $\epsilon$ -clusters in  $\mathcal{X} \cup \{Y\}$ ,  $Y$ 's credits are enough to pay for those spanner edges when  $c(\epsilon) = \epsilon^{-O(d)}$ .

For each  $\epsilon$ -cluster  $z \in C \setminus (\mathcal{X} \cup \{Y\})$ , we use  $z$ 's credit to pay for the spanner edges incident to  $z$ . By Lemma 8 and invariant DC1 for level  $i - 1$ , this amount of credit is sufficient when  $c(\epsilon) = \epsilon^{-O(d)}$ .  $\square$

**Claim 14.** *The credits of  $\epsilon$ -clusters and MST edges connecting  $\epsilon$ -clusters of each short cluster  $C$  are enough to maintain invariant DC1 for  $C$ .*

*Proof.* We abuse notation by letting  $\text{MST}(C)$  be the set of MST edges in  $C$  that connects its  $\epsilon$ -clusters. Since  $C$  is a cluster path, we have:

$$\text{diam}(C) \leq \sum_{X_\epsilon \in C} \text{diam}(X_\epsilon) + \sum_{e \in \text{MST}(C)} w(e)$$

By invariant DC1 for level  $i - 1$ ,  $cr(X_\epsilon) \geq c(\epsilon) \cdot \text{diam}(X_\epsilon)$  and since each MST edge has credit at least  $c(\epsilon)$  times its length, the claim follows.  $\square$

A short cluster may need to use all the credits of  $\epsilon$ -clusters and MST edges to maintain DC1, hence, it may not have extra credit to pay for any incident level- $i$  spanner edges. In this case, we need to use credits of other level- $i$  clusters to pay for those spanner edges. We call a short cluster *internal* if it is not an affix of a long path  $\mathcal{P}$  in Phase 4.

**Observation 15.** *There is no level- $i$  spanner edge  $e$  that has both endpoints in internally short clusters.*



*Proof.* If there is such an edge  $e$ , it would be grouped into a level- $i$  cluster in Phase 2.  $\square$

Thus, a level- $i$  spanner edge incident to an internally short cluster can be paid by the level- $i$  cluster that contains the other endpoint of  $e$ . However, if a short cluster is not internal, we must find a way to pay for its incident spanner edges. Recall after Phase 1, every cluster path of effective diameter at least  $\ell$  must have an MST edge from one of its endpoint  $\epsilon$ -clusters to a level- $i$  cluster. By construction in Phase 4, if a short cluster is an affix of  $\mathcal{P}$ , called a *short affix cluster*, the other affix of  $\mathcal{P}$ , called the *sibling affix*, must have an MST edge to a cluster originating in the first two phases and thus augments it. (The only exception is when there is no level- $i$  clusters after the first two phases and we will handle this case at the end of this paper.) Thus, we can use the credit of  $\epsilon$ -clusters of the sibling affices to pay for incident spanner edges of affix short clusters. To that end, we analyze clusters originally constructed in the first two phases.

### 3.3.2 Clusters originating in Phase 1+2

Let  $C$  be a level- $i$  cluster constructed in Phase 1 or 2. Let  $C'$  and  $C''$  be the augmentations of  $C$  in Phase 3 and 4, respectively. Let  $D$  be the diameter path of the spanner given by edges and vertices in  $C''$ . Let  $\mathcal{D}$  be the walk obtained from  $D$  by contracting each maximal subpath of  $D$  that is inside an  $\epsilon$ -cluster of  $C''$  to a single vertex.

**Definition 16** (Canonical pair). *Let  $\mathcal{S} \subseteq C \cup \mathcal{D}$  be a subset of  $\epsilon$ -clusters of  $C''$  such that  $|\mathcal{S}| \leq \frac{2g}{\epsilon}$  and the credits of  $\epsilon$ -clusters in  $\mathcal{S}$  and MST edges of  $C''$  are sufficient to maintain invariant DC1 for  $C''$ . Let  $Y$  be an  $\epsilon$ -cluster of  $C$  that is not in  $\mathcal{S}$ . We call  $(\mathcal{S}, Y)$  a canonical pair of  $C''$ .*

Note that we do not claim the existence of canonical pairs. Indeed, the main goal of this subsection is to prove that a canonical pair exists for  $C''$  since its existence implies that  $C'' \setminus \mathcal{S} \neq \emptyset$ . Thus, we can use credits of  $\epsilon$ -clusters in  $C'' \setminus \mathcal{S}$  to pay for level- $i$  spanner edges incident to  $\epsilon$ -clusters of  $C''$  and  $\epsilon$ -clusters of short affix clusters that have sibling affices in  $C''$ .

**Claim 17.** *If  $C''$  has a canonical pair  $(\mathcal{S}, Y)$ , we then can pay for every level- $i$  spanner edge that is incident to  $\epsilon$ -clusters of  $C''$  and  $\epsilon$ -clusters of short affix clusters that have sibling affices in  $C''$  using credits of  $\epsilon$ -clusters in  $C'' \setminus \mathcal{S}$  when  $c(\epsilon) = \epsilon^{-\Theta(d)}$ .*

*Proof.* Let  $\mathcal{R}$  be a set of  $\epsilon$ -clusters that contains every  $\epsilon$ -cluster in  $\mathcal{S} \cup \{Y\}$  and affix short clusters in Phase 4 whose sibling affices contain an  $\epsilon$ -cluster of  $\mathcal{D}$ . Recall that  $C'$  is augmented by attaching paths of  $\epsilon$ -clusters via MST edges. Thus,  $\mathcal{R}$  contains at most two short clusters as a result of Phase 4 (see Figure 2). Since  $|\mathcal{S}| \leq \frac{2g}{\epsilon}$  and each short cluster has at most  $\frac{2g}{\epsilon}$   $\epsilon$ -clusters,  $|\mathcal{R}| = O(\frac{g}{\epsilon})$ . Since each  $\epsilon$ -cluster is incident to at most  $\Delta_K$  level- $i$  spanner edges by Lemma 8,  $\epsilon$ -clusters in  $\mathcal{R}$  are incident to at most  $O(\frac{g\Delta_K}{\epsilon})$  level- $i$  spanner edges. Recall that each level- $i$  spanner edge has length at most  $\ell$ . By invariant DC1 for level  $i-1$ ,  $Y$  has at least  $c(\epsilon)\ell/2$  credits. Thus, by choosing  $c(\epsilon) = \Theta(\frac{g\Delta_K}{\epsilon^2}) = \epsilon^{-\Theta(d)}$ ,  $Y$ 's credit is sufficient to pay for every spanner edge incident to  $\epsilon$ -clusters in  $\mathcal{R}$ .

For each  $\epsilon$ -cluster  $z$  in  $C'' \setminus \mathcal{R}$ , the credit of  $z$  is sufficient to pay for incident level- $i$  spanner edges incident to  $z$ . However, we also need to pay for short affix clusters in Phase 4, whose siblings augment  $C'$  in Phase 4. To afford this, we use half the credit of each  $\epsilon$ -cluster in  $C \setminus \mathcal{R}$  (of value at least  $c(\epsilon)\ell/4$  by invariant DC1 for level  $i-1$ ) to pay for level- $i$  spanner edges incident to it. Since each  $\epsilon$ -cluster is incident to at most  $\Delta_K$  level- $i$  spanner edges by Lemma 8, this credit is sufficient when  $c(\epsilon) \geq \frac{4\Delta_K}{\epsilon} = \epsilon^{-O(d)}$ .



*Proof.* Let  $\mathcal{F} = \{\mathcal{P}_1, \mathcal{P}_2, \mathcal{Q}_1, \mathcal{Q}_2\}$ . By construction in Phase 2,  $C$  is an acyclic graph of  $\epsilon$ -clusters connected by MST edges and a single spanner edge  $e$ . Thus,  $\mathcal{D}$  is a simple path. That implies at most two paths, say  $\mathcal{P}'$  and  $\mathcal{Q}'$ , among four paths in  $\mathcal{F}$  share  $\epsilon$ -clusters with  $\mathcal{D}$ . Let other two paths of  $\mathcal{F}$  be  $\mathcal{P}''$  and  $\mathcal{Q}''$ . Let  $Y$  be an arbitrary  $\epsilon$ -cluster of  $\mathcal{Q}''$  and

$$\mathcal{S} = \lfloor \mathcal{D} \cup \mathcal{P}' \cup \mathcal{Q}' \cup \mathcal{P}'' \rfloor_{2g/\epsilon}$$

If  $\mathcal{S} = \frac{2g}{\epsilon}$ , then  $cr(\mathcal{S}) \geq c(\epsilon)g\ell$  by invariant DC1 for level  $i-1$ . Hence, credits of  $\epsilon$ -clusters in  $\mathcal{S}$  are sufficient to maintain DC1 for  $C''$  since  $\text{diam}(C'') \leq g\ell$  as shown in the previous section.

Thus, we can assume that  $\mathcal{S} < \frac{2g}{\epsilon}$ . In this case,  $\mathcal{S} = \mathcal{D} \cup \mathcal{P}' \cup \mathcal{Q}' \cup \mathcal{P}''$ . If  $\mathcal{D}$  does not contain the spanner edge  $e$ , using the same argument in Claim 18, we can show that credits of  $\epsilon$ -clusters and MST edges of  $\mathcal{D}$  are sufficient to maintain invariant DC1 for  $C''$ . Otherwise, we assign credits of  $\mathcal{P}''$  to  $e$ . Since  $\text{ediam}(\mathcal{P}'') \geq \ell \geq w(e)$ , by invariant DC1 for level  $i-1$ ,  $cr(\mathcal{P}'') \geq c(\epsilon)\text{ediam}(\mathcal{P}'') \geq c(\epsilon)w(e)$ . Thus  $e$  is assigned credit of at least  $c(\epsilon)$  times its length. We then use credits of  $\epsilon$ -clusters and edges of  $\mathcal{D}$  to maintain DC1. The rest of the proof is similar to Claim 18.  $\square$

We assume that  $\mathcal{P}_1 = \mathcal{Q}_1 = \mathcal{P}_{xy}$ . In this case,  $C$  contains a unique cycle, which is  $\{e\} \cup \mathcal{P}_{xy}$ . We first prove that  $\mathcal{D}$  is a path when  $\epsilon$  is sufficiently small.

**Claim 20.**  $\mathcal{D}$  is a path if  $\epsilon$  is smaller than  $\frac{1}{2g}$ .

*Proof.* If  $\mathcal{D}$  is not simple, it contains a cycle  $\mathcal{C}_{xy}$ . Let  $u$  and  $v$  be two vertices of the same  $\epsilon$ -cluster, say  $X_\epsilon$ , such that  $\mathcal{D}$  enters and leaves  $\mathcal{C}_{xy}$  at  $u$  and  $v$ , respectively. Then, the subpath  $D_{uv}$  between  $u$  and  $v$  of  $\mathcal{D}$  must contain edge  $e$  of length at least  $\ell/2$ . However, we can shortcut  $D_{uv}$  through  $X_\epsilon$  by a path of length at most  $\text{diam}(X_\epsilon) \leq g\ell$  by DC2. For  $\epsilon < \frac{1}{2g}$ , the shortcut has length smaller than  $w(D_{uv})$ , contradicting that  $\mathcal{D}$  is a shortest path.  $\square$

**Observation 21.**  $\mathcal{P}_{xy} \not\subseteq \mathcal{D}$ .

*Proof.* For otherwise,  $\mathcal{D}$  could be shortcut through  $e$  at a cost of

$$\begin{aligned} &\leq \underbrace{\text{diam}(C_x) + \text{diam}(C_y) + w(e)}_{\text{cost of shortcut}} - \underbrace{(\text{diam}(\mathcal{P}_{xy}) - \text{diam}(C_x) - \text{diam}(C_y))}_{\text{lower bound on diameter}} \\ &\leq w(e) + 4g\ell - (1 + s\epsilon)w(e) \quad (\text{by the stretch condition for } e) \\ &\leq 4g\ell - s\epsilon\ell/2 \quad (\text{since } w(e) \geq \ell/2) \end{aligned}$$

This change in cost is negative for  $s \geq 8g + 1$ .  $\square$

**Claim 22.**  $C''$  has a canonical pair.

*Proof.* Let  $Y$  be an  $\epsilon$ -cluster of  $\mathcal{P}_{xy} \setminus \mathcal{D}$ .  $Y$  exists by Observation 21. We define:

$$\mathcal{S} = \lfloor \mathcal{D} \cup \mathcal{P}_2 \cup \mathcal{Q}_2 \cup \mathcal{P}_{xy} \setminus \{Y\} \rfloor_{2g/\epsilon}$$

If  $|\mathcal{S}| = \frac{2g}{\epsilon}$ , then the total credit of  $\epsilon$ -clusters in  $\mathcal{S}$  is at least  $c(\epsilon)g\ell$  by invariant DC1 for level  $i-1$ . Thus credits of  $\epsilon$ -clusters in  $\mathcal{S}$  is sufficient to maintain invariant DC1 for  $C''$ . That implies  $C''$  has a canonical pair.

Otherwise,  $\mathcal{S} = \mathcal{D} \cup \mathcal{P}_2 \cup \mathcal{Q}_2 \cup \mathcal{P}_{xy} \setminus \{Y\}$ . If  $\mathcal{D}$  does not contain the spanner edge  $e$ , then by the same argument in Claim 18, we can argue that credits of  $\epsilon$ -clusters and MST edges in  $\mathcal{D}$  is enough to maintain invariant DC1 for  $C''$ . Thus, we can assume that  $\mathcal{D}$  contains  $e$ . We consider two cases:

1.  $\mathcal{D}$  contains an internal  $\epsilon$ -cluster of  $\mathcal{P}_{xy}$ . Since  $\mathcal{D}$  is a path by Claim 20, it does not contain any internal  $\epsilon$ -cluster of at least one of two paths  $\mathcal{P}_2, \mathcal{Q}_2$ , w.l.o.g., say  $\mathcal{P}_2$ . Since  $\text{ediam}(\mathcal{P}_2) \geq \ell$ , by invariant DC1 for level  $i - 1$ , the total credit of  $\epsilon$ -clusters in  $\mathcal{P}_2$  is at least  $c(\epsilon)\ell$  which is at least  $c(\epsilon)w(e)$ . Thus, by assigning credits of  $\mathcal{P}_2$  to  $e$ , every edge of  $\mathcal{D}$  has credit at least  $c(\epsilon)$  times its length. Thus, credits of  $\epsilon$ -clusters and edges of  $\mathcal{D}$  are enough to maintain DC1 for  $C''$ .
2.  $\mathcal{D}$  does not contain any internal  $\epsilon$ -cluster of  $\mathcal{P}_{xy}$ . We have:

$$\begin{aligned}
& \text{diam}(\mathcal{P}_{xy} \setminus \{C_x, C_y\}) \\
& \geq \text{diam}(\mathcal{P}_{xy}) - \text{diam}(C_x) - \text{diam}(C_y) \\
& \geq (1 + s\epsilon)w(e) - \text{diam}(C_x) - \text{diam}(C_y) \quad (\text{by the stretch condition}) \\
& \geq w(e) + s\epsilon\ell/2 - 2g\epsilon\ell \quad (\text{by bounds on } w(e) \text{ and DC2}) \\
& \geq w(e) + g\epsilon\ell \quad (\text{for } s \geq 8g + 1, \text{ as previously required})
\end{aligned} \tag{4}$$

The credit of the MST edges and  $\epsilon$ -clusters of  $\mathcal{P}_{xy} \setminus \{C_x, C_y\}$  is at least:

$$\begin{aligned}
& c(\epsilon) \cdot (\text{MST}(\mathcal{P}_{xy} \setminus \{C_x, C_y\}) + \text{ediam}(\mathcal{P}_{xy} \setminus \{C_x, C_y\})) \\
& \geq c(\epsilon) \cdot \text{diam}(\mathcal{P}_{xy} \setminus \{C_x, C_y\}) \\
& \geq c(\epsilon)(w(e) + g\epsilon\ell)
\end{aligned} \quad \square$$

Since  $\text{diam}(Y) \leq g\epsilon\ell$  by invariant DC2 for level  $i - 1$ , the total credit of  $\epsilon$ -clusters of  $\mathcal{P}_{xy} \setminus \{C_x, C_y, Y\}$  and MST edges of  $\mathcal{P}_{xy} \setminus \{C_x, C_y\}$  is at least  $c(\epsilon) \cdot w(e)$ . Thus, by assigning this credit to  $e$ , we can argue that credits of  $\epsilon$ -clusters and edges of  $\mathcal{D}$  are enough to maintain DC1 for  $C''$ .

### 3.3.3 No Phase 1 or 2 clusters

We now deal with the case when there are no level- $i$  clusters formed in Phase 1 and 2.

**Observation 23.** *There is no level- $i$  cluster formed in Phase 1 and 2 if and only if (i) the tree  $\mathcal{T}$  of  $\epsilon$ -clusters is a path and (ii) every spanner edge is incident to an  $\epsilon$ -cluster in an affix of  $\mathcal{T}$  having effective diameter at most  $2\ell$ .*

By Claim 13, we only need to pay for spanner edges incident to short affix clusters of  $\mathcal{T}$ . Since short clusters have at most  $\frac{2g}{\epsilon}$   $\epsilon$ -clusters, there are at most  $\frac{4g\Delta\kappa}{\epsilon}$  such spanner edges, that we assign to set  $B$  (Lemma 7). Below, we show that  $w(B) \leq \epsilon^{-O(d)} \cdot w(\text{MST})$  across all levels, implying Lemma 7.

**Claim 24.**  $w(B) \leq \epsilon^{-O(d)} \cdot w(\text{MST})$ .

*Proof.* We have:

$$\begin{aligned}
\frac{4g\Delta\kappa}{\epsilon} \sum_i \ell_i & \leq \frac{4g\Delta\kappa}{\epsilon} \ell_{\max} \sum_i \epsilon^i, \text{ where } \ell_{\max} = \max_{e \in S} \{w(e)\} \\
& \leq \frac{4g\Delta\kappa}{\epsilon} w(\text{MST}) \sum_i \epsilon^i \\
& \leq \frac{4g\Delta\kappa}{\epsilon} w(\text{MST}) \frac{1}{1 - \epsilon} = \epsilon^{-O(d)} \cdot w(\text{MST})
\end{aligned} \tag{5}$$

□

## References

- [1] I. Althöfer, G. Das, D. Dobkin, D. Joseph, and J. Soares. On sparse spanners of weighted graphs. *Discrete Computational Geometry*, 9(1):81–100, 1993.
- [2] S. Arya, G. Das, D. M. Mount, J. S. Salowe, and M. Smid. Euclidean spanners: Short, thin, and lanky. In *Proceedings of the Twenty-seventh Annual ACM Symposium on Theory of Computing*, STOC '95, pages 489–498, 1995.
- [3] P. Assouad. Plongements lipschitziens dans  $\mathbb{R}^n$ . *Société mathématique de France*, 111:429–448, 1983.
- [4] G. Borradaile, E. D. Demaine, and S. Tazari. Polynomial-time approximation schemes for subset-connectivity problems in bounded-genus graphs. *Algorithmica*, 68(2):287–311, 2014.
- [5] G. Borradaile, H. Le, and C. Wulff-Nilsen. Minor-free graphs have light spanners. In *2017 IEEE 58th Annual Symposium on Foundations of Computer Science*, 2017.
- [6] S. Chechik and C. Wulff-Nilsen. Near-optimal light spanners. In *Proceedings of the 27th Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA'16, pages 883–892, 2016.
- [7] G. Das, P. Heffernan, and G. Narasimhan. Optimally sparse spanners in 3-dimensional euclidean space. In *Proceedings of the 9th Annual Symposium on Computational Geometry*, SCG '93, pages 53–62, 1993.
- [8] G. Das, G. Narasimhan, and J. Salowe. A new way to weigh malnourished euclidean graphs. In *Proceedings of the 6th Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '95, pages 215–222, 1995.
- [9] A. Filtser and S. Solomon. The greedy spanner is existentially optimal. In *Proceedings of the 2016 ACM Symposium on Principles of Distributed Computing*, PODC '16, pages 9–17, 2016.
- [10] J. Gao, L. J. Guibas, and A. Nguyen. Deformable spanners and applications. *Computational Geometry*, 35(1):2–19, 2006.
- [11] L. A. Gottlieb. A light metric spanner. In *2015 IEEE 56th Annual Symposium on Foundations of Computer Science*, pages 759–772, 2015.
- [12] J. Heinonen. *Lectures on analysis on metric spaces*. Universitext. Springer New York, 2001.
- [13] P. N. Klein. A linear-time approximation scheme for planar weighted TSP. In *Proceedings of the 46th Annual IEEE Symposium on Foundations of Computer Science*, FOCS '05, pages 647–657, 2005.
- [14] C. Levcopoulos and A. Lingas. There are planar graphs almost as good as the complete graphs and as short as minimum spanning trees. In *International Symposium on Optimal Algorithms*, pages 9–13, 1989.
- [15] G. Narasimhan and M. Smid. *Geometric Spanner Networks*, chapter Geometric Analysis: The Leapfrog Property, pages 257–317. Cambridge University Press, 2007.

- [16] S. B. Rao and W. D. Smith. Approximating geometrical graphs via “spanners” and “banyans”. In *Proceedings of the 30th Annual ACM Symposium on Theory of Computing*, STOC ’98, pages 540–550, 1998.
- [17] M. Smid. The weak gap property in metric spaces of bounded doubling dimension. In Susanne Albers, Helmut Alt, and Stefan Näher, editors, *Efficient Algorithms*, pages 275–289. Springer-Verlag, 2009.

## A Notation and definitions

Let  $G(V(G), E(G))$  be a connected and undirected graph with a positive edge weight function  $w : E(G) \rightarrow \mathbb{R}^+ \setminus \{0\}$ . We denote  $|V(G)|$  and  $|E(G)|$  by  $n$  and  $m$ , respectively. Let  $\text{MST}(G)$  be a minimum spanning tree of  $G$ ; when the graph is clear from the context, we simply write  $\text{MST}$ . A walk of length  $p$  is a sequence of alternating vertices and edges  $\{v_0, e_0, v_1, e_1, \dots, e_{p-1}, v_p\}$  such that  $e_i = v_i v_{i+1}$  for every  $i$  such that  $0 \leq i \leq p-1$ . A path is a *simple walk* where every vertex appears exactly once in the walk. For two vertices  $x, y$  of  $G$ , we use  $d_G(x, y)$  to denote the shortest distance between  $x$  and  $y$ .

Let  $S$  be a subgraph of  $G$ . We define  $w(S) = \sum_{e \in E(S)} w(e)$ . Let  $X \subseteq V(G)$  be a set of vertices. We use  $G[X]$  to denote the subgraph of  $G$  induced by  $X$ . Let  $Y \subseteq E(G)$  be a subset of edges of  $G$ . We denote the graph with vertex set  $V(G)$  and edge set  $Y$  by  $G[Y]$ .

## B Greedy spanners

A subgraph  $S$  of  $G$  is a  $(1 + \epsilon)$ -spanner of  $G$  if  $V(S) = V(G)$  and  $d_S(x, y) \leq (1 + \epsilon)d_G(x, y)$  for all  $x, y \in V(G)$ . The following greedy algorithm by Althöfer et al. [1] finds a  $(1 + \epsilon)$ -spanner of  $G$ :

```

GREEDYSPANNER( $G(V, E), \epsilon$ )
   $S \leftarrow (V, \emptyset)$ .
  Sort edges of  $E$  in non-decreasing order of weights.
  For each edge  $xy \in E$  in sorted order
    if  $(1 + \epsilon)w(xy) < d_S(x, y)$ 
       $E(S) \leftarrow E(S) \cup \{e\}$ 
  return  $S$ 

```

Observe that as algorithm `GREEDYSPANNER` is a relaxation of Kruskal’s algorithm,  $\text{MST}(G) = \text{MST}(S)$ . Since we only consider  $(1 + \epsilon)$ -spanners in this work, we simply call an  $(1 + \epsilon)$ -spanner a *spanner*. We define the lightness of a spanner  $S$  to be the ratio  $\frac{w(S)}{w(\text{MST}(G))}$ . We call  $S$  *light* if its lightness is independent of the number of vertices or edges of  $G$ .

## C Omitted Proofs

*Proof of Lemma 8.* To show that  $\mathcal{K}$  is simple, we use the same argument as Borradaile, Le and Wulff-Nilsen [5] that we briefly sketch here. Recall  $\epsilon$ -clusters have diameter at most  $g\epsilon\ell$  by the diameter-credit invariants. Recall edges in  $E_i$  have weight in range  $(\ell/2, \ell]$ . Thus, when  $\epsilon$  is sufficiently small,  $\mathcal{K}$  has no self-loops. To show that  $\mathcal{K}$  has no parallel edges, we assume there are

such two  $xy$  and  $uv$  where  $w(xy) < w(uv)$ . Let  $C_u, C_v$  be two  $\epsilon$ -clusters that contain  $u$  and  $v$ , respectively. We further assume, w.l.o.g, that  $x \in C_u, y \in C_v$ . Then the  $u$ -to- $v$  path  $P_{uv}$  from  $u$  to  $x$  inside  $C_u$ , edge  $xy$  and then  $y$  to  $v$  inside  $C_v$  has length at most  $w(xy) + 2g\epsilon\ell$ , which is at most  $(1 + 4g\epsilon)w(uv)$  since  $w(uv) \geq \ell/2$ . Thus, by choosing  $s \geq 4g$ , edge  $uv$  is not added to the spanner by the greedy algorithm. Thus,  $\mathcal{K}$  is simple.  $\square$