



University of
Massachusetts
Amherst

Greedy Spanners in Euclidean Spaces Admit Sublinear Separators

Item Type	article;article
Authors	Le, Hung;Than, Cuong
DOI	https://doi.org/10.1145/3590771
Rights	UMass Amherst Open Access Policy
Download date	2024-08-14 23:17:16
Link to Item	https://hdl.handle.net/20.500.14394/9792

Greedy Spanners in Euclidean Spaces Admit Sublinear Separators

HUNG LE and CUONG THAN, University of Massachusetts Amherst, USA

The greedy spanner in a low dimensional Euclidean space is a fundamental geometric construction that has been extensively studied over three decades as it possesses the two most basic properties of a good spanner: constant maximum degree and constant lightness. Recently, Eppstein and Khodabandeh [28] showed that the greedy spanner in \mathbb{R}^2 admits a sublinear separator in a strong sense: any subgraph of k vertices of the greedy spanner in \mathbb{R}^2 has a separator of size $O(\sqrt{k})$. Their technique is inherently planar and is not extensible to higher dimensions. They left showing the existence of a small separator for the greedy spanner in \mathbb{R}^d for any constant $d \geq 3$ as an open problem.

In this paper, we resolve the problem of Eppstein and Khodabandeh [28] by showing that any subgraph of k vertices of the greedy spanner in \mathbb{R}^d has a separator of size $O(k^{1-1/d})$. We introduce a new technique that gives a simple criterion for any geometric graph to have a sublinear separator that we dub τ -lanky: a geometric graph is τ -lanky if any ball of radius r cuts at most τ edges of length at least r in the graph. We show that any τ -lanky geometric graph of n vertices in \mathbb{R}^d has a separator of size $O(\tau n^{1-1/d})$. We then derive our main result by showing that the greedy spanner is $O(1)$ -lanky. We indeed obtain a more general result that applies to unit ball graphs and point sets of low fractal dimensions in \mathbb{R}^d .

Our technique naturally extends to doubling metrics. We use the τ -lanky criterion to show that there exists a $(1+\epsilon)$ -spanner for doubling metrics of dimension d with a constant maximum degree and a separator of size $O(n^{1-\frac{1}{d}})$; this result resolves an open problem posed by Abam and Har-Peled [1] a decade ago. We then introduce another simple criterion for a graph in doubling metrics of dimension d to have a sublinear separator. We use the new criterion to show that the greedy spanner of an n -point metric space of doubling dimension d has a separator of size $O((n^{1-\frac{1}{d}}) + \log \Delta)$ where Δ is the spread of the metric; the factor $\log(\Delta)$ is tightly connected to the fact that, unlike its Euclidean counterpart, the greedy spanner in doubling metrics has *unbounded maximum degree*. Finally, we discuss algorithmic implications of our results.

CCS Concepts: • **Theory of computation** → **Computational geometry**.

Additional Key Words and Phrases: spanner, greedy spanner, sublinear separator, geometric graph

Authors' address: Hung Le, hungle@cs.umass.edu; Cuong Than, cthan@umass.edu, University of Massachusetts Amherst, 120 Tillson Farm Road, Amherst, MA, USA, 01003-9346.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.

1549-6325/2023/4-ART \$15.00

<https://doi.org/10.1145/3590771>

CONTENTS

Abstract	I
Contents	II
1 Introduction	1
1.1 Our Contributions	3
1.2 Algorithmic Implications	5
2 Preliminaries	5
3 Criteria of Graphs with Sublinear Separators	6
3.1 Bounded Degree Graphs	7
3.2 Graphs with High Vertex Degrees	9
4 Separators of Greedy Spanners for Graphs in Euclidean Spaces	12
4.1 Euclidean Spaces	12
4.2 Point Sets of Low Fractal Dimensions	14
4.3 Unit Ball Graphs	14
4.4 Derandomization	15
5 Bounded Degree Spanners with Sublinear Separators in Doubling Metrics	18
6 Separators of Greedy Spanners for Doubling Metrics	24
Acknowledgments	26
References	26

1 INTRODUCTION

A t -spanner of an edge-weighted graph $G = (V, E, w)$ is a spanning edge-weighted subgraph H such that $\delta_H(u, v) \leq t \cdot \delta_G(u, v)$ for every pair of vertices $u, v \in V$; t is called the *stretch* of the spanner. This definition naturally extends to a t -spanner of a point set P in a metric space (X, δ_X) where the graph $G = (P, \binom{P}{2}, \delta_X)$ in this context is the complete graph representing the metric induced by P . The study of spanners dated back to the work of Chew [15] in 1986 who constructed a $\sqrt{10}$ -spanner for a point set on the Euclidean plane, though the term *spanner* had not yet appeared until the work of Peleg and Schäffer [53].

Over more than 30 years, research on spanners has evolved into an independent field of study, and has found numerous applications, such as to VLSI circuit design [17, 18, 57], to distributed computing [4, 5, 52], to approximation algorithms [9, 16, 35, 43, 44, 55], to wireless and sensor networks [8, 58, 63], to machine learning [36], and to computational biology [56]. Spanners enjoy wide applicability since they possess many desirable properties such as *low sparsity*, implying, e.g., low storage cost; *small lightness*, implying, e.g., low construction cost; *low maximum degree*, implying, e.g., small routing tables in routing applications. Here, the *sparsity* of a spanner is the ratio of its number of edges to the number of edges of the minimum spanning tree (MST), which is $n - 1$, and the *lightness* of a spanner is the ratio of its total edge weight to the weight of the MST. Some spanners only have one property among three properties while others possess all of them and even more [3, 26].

The *path greedy algorithm*, or greedy algorithm for short, is perhaps the simplest and one of the most well-studied algorithms for constructing a t -spanner: consider edges of G in the non-decreasing order of the weights, and add to the spanner an edge e if the distance between its endpoints in the current spanner is larger than $t \cdot w(u, v)$. The output spanner is called the *greedy spanner*. The greedy algorithm was introduced by Althöfer *et al.* [2] and independently discovered by Bern¹. In addition to its great simplicity, the greedy algorithm has shown to be one-algorithm-fits-all. The greedy $(1 + \epsilon)$ -spanner has constant sparsity and constant lightness for point sets in Euclidean spaces² [14, 19, 45, 55], doubling metrics [10, 35, 61], and minor-free graphs [10]. For *edge-weighted* general graphs of n vertices, the greedy spanner with stretch $(2k - 1)$ has $O(n^{1/k})$ sparsity and with stretch $(2k - 1)(1 + \epsilon)$ has $O(n^{1/k})$ lightness for a constant ϵ , which are (nearly) optimal bounds assuming the Erdős' girth conjecture [30]. Filtser and Solmon [31] showed that greedy spanners are *existentially optimal* for both the size and the lightness for several graph families, which include general graphs and minor-free graphs. Roughly speaking, existential optimality for a graph family \mathcal{G} means that the worst-case sparsity and lightness of the greedy spanner over all graphs in \mathcal{G} are as good as the sparsity and lightness of an optimal spanner over all graphs in \mathcal{G} . In summary, understanding the greedy spanner is an important problem, and there have been many papers dedicated to this task.

In this work, we investigate greedy spanners for point sets from a different perspective: the existence of sublinear balanced separators³. For brevity, we simply refer to a balanced separator as a separator. While our work covers several settings, understanding the greedy spanner for point sets in Euclidean spaces is of special interest. This fundamental setting has been studied extensively; the chapter Geometric Analysis: The Leapfrog Property of the book by Narasimhan and Smid [51] is dedicated solely to Euclidean spanners. Many algorithmic ideas developed in the context of Euclidean spanners could be carried over different settings, such as doubling metrics and unit ball graphs (UBGs), which are intersection graphs of unit balls in \mathbb{R}^d . On the other hand, having a small separator is algorithmically significant. In their seminal work, Lipton and Tarjan [47] demonstrated the algorithmic power of sublinear separators in designing divide-and-conquer algorithms for solving many algorithmic tasks. (Though their results were stated for planar graphs, their techniques are applicable to any

¹Althöfer *et al.* [2] attributed the independent discovery of the greedy algorithm to an unpublished work of Marshall Bern.

²In this work, we assume that ϵ and d are fixed constants, unless specified otherwise.

³We say that a separator S of an n -vertex graph $G = (V, E)$ is balanced if every connected component of $G[V \setminus S]$ has at most $c \cdot n$ vertices for some constant c .

graph class that is closed under taking subgraphs and has sublinear separators.) Over four decades since the work of Lipton and Tarjan [47], sublinear separators have played a central role in the development of many algorithmic paradigms, such as subexponential time parameterized algorithms [20, 21, 32, 49], analyzing local search heuristics [11, 40], designing polynomial time approximation schemes (PTASes) for problems definable in first-order logic [6, 24], to name a few.

There have been some prior works that focus on (non-greedy) spanners with sublinear separators. Smith and Wormald [62] sketched a proof showing that the spanner of Arya *et al.* [3] has a separator of size $O(n^{1-1/d})$; their proof relies on the (κ, λ) -thick property⁴. Abam and Har-Peled [1] constructed a $(1 + \epsilon)$ -spanner with a separator of size $O(n^{1-1/d})$ for point sets in \mathbb{R}^d with maximum degree $O(\log^2 n)$ and constant sparsity using a semi-separated pair decomposition (SSPD), a concept introduced in the same paper. They left as an open problem of constructing a spanner with a sublinear separator and a *constant maximum degree* in metrics of constant doubling dimensions. Fürer and Kasiviswanathan [33] constructed a spanner of constant sparsity with a separator of size $O(n^{1-1/d} + \log(\Gamma))$ for ball graphs, which are intersection graphs of balls of arbitrary radii in \mathbb{R}^d ; here Γ is the ratio of the maximum radius to minimum radius over the balls, which could be exponentially large. However, for the important special case of UBGs ($\Gamma = 1$), the result of Fürer and Kasiviswanathan [33] implies a separator of size $O(n^{1-1/d})$. Sidiropoulos and Sridhar [59] devised a construction of a spanner with a separator of size $O(n^{1-1/d_f})$ for point sets in \mathbb{R}^d with a *fractal dimension* d_f (see Definition 4 for a formal definition). Both results of Fürer and Kasiviswanathan [33] and Sidiropoulos and Sridhar [59] imply $(1 + \epsilon)$ -spanners with sublinear separators for point sets in \mathbb{R}^d as their settings are more general⁵. However, both constructions are fairly complicated and based on the (κ, λ) -thick property of Smith and Wormald [62]. Furthermore, their spanners are not known to have desirable properties of spanners such as maximum degree and constant lightness; Abam and Har-Peled [1] showed that the max degree of spanner of Fürer and Kasiviswanathan [33] could be $\Omega(n)$. These results motivate the following question:

Question 1. *Does the greedy $(1 + \epsilon)$ -spanner for any point set in \mathbb{R}^d have a separator of size $O(n^{1-1/d})$ for fixed constants $d \geq 2, \epsilon \leq 1$? More generally, do greedy $(1 + \epsilon)$ -spanners for more general settings, such as ball graphs in \mathbb{R}^d and point sets in Euclidean spaces of small fractal dimension, have sublinear separators?*

We remark that the upper bound $O(n^{1-1/d})$ of the separator size is optimal; the greedy $(1 + \epsilon)$ -spanner of point sets in the d -dimensional grid has size $\Omega(n^{1-1/d})$ [59].

Recently, Eppstein and Khodabandeh [28] showed that the greedy spanner for point sets in \mathbb{R}^2 admits a separator of size $O(\sqrt{n})$, thereby resolving Question 1 for the case of Euclidean graphs and $d = 2$. They derived their result by studying the crossing patterns of the straight-line drawing of the greedy spanner on the Euclidean plane. Since edge crossings are not useful in dimensions at least 3—the edges of the greedy spanners for point sets in a general position generally have no crossing—their technique cannot be applied to higher dimensions. They left Question 1 for dimension $d \geq 3$ as an open problem.

Nevertheless, the technique based on edge crossings of Eppstein and Khodabandeh [28] gives a sublinear separator in a strong sense: any subgraph of k vertices of the greedy spanner in \mathbb{R}^2 has a separator of size $O(\sqrt{k})$. That is, having a square-root separator is a *monotone property*: it is closed under taking subgraphs. Being a monotone property is not a mere generalization; all aforementioned algorithmic applications of sublinear separators require monotonicity. Fürer and Kasiviswanathan [33] showed that their spanner has a monotone sublinear separator, and this is crucial for the algorithmic applications in their paper. It is not clear whether other non-greedy spanners mentioned above have the same property. On the other hand, all spanners in our paper have monotone sublinear separators.

⁴A set of objects (e.g., balls) in \mathbb{R}^d is (κ, λ) -thick if every point in the space belongs to at most κ objects and the maximum size (linear dimension, e.g., diameter) of an object is at most λ times the minimum size.

⁵The fractal dimension of any point set in \mathbb{R}^d is $d_f \leq d$.

While resolving Question 1 is an important goal, it is probably more important and challenging to have a simple tool to resolve it, as asked in the following question:

Question 2. *Is there a simple criterion of graphs with sublinear separators in Euclidean and doubling metrics of constant dimensions that could be easily applied to (greedy or non-greedy) spanners?*

A simple criterion asked in Question 2 would be particularly significant, as there are numerous spanner constructions in the literature that are faster than the greedy algorithm and give spanners with different properties than those guaranteed by greedy spanners. A prime example is greedy spanners for doubling metrics of constant dimensions. Smid [61] constructed a doubling metric (X, δ_X) of dimension 1 such that the greedy spanner for (X, δ_X) has maximum degree $n - 1$. However, there are other constructions [12, 37] that give a spanner with constant maximum degree for doubling metrics. Does any of the spanners given by these algorithms have a sublinear separator? A positive answer to this question would resolve the open problem by Abam and Har-Peled [1]. And in this work, we provide a positive answer (Theorem 3), by showing that the construction of [12] gives a spanner with a sublinear separator (and constant maximum degree); we obtain our result based on a simple criterion that we develop.

There have been several criteria of *geometric graphs* with sublinear separators: (κ, λ) -thick property [62], k -ply property of the intersection graphs of balls [50], k -th nearest neighbor graphs [50], intersection graphs of convex shapes that are tame [23]; we refer readers to the cited papers for precise definitions of these concepts. While these criteria are fairly simple and generally applicable, it is not clear how to use them, except for the (κ, λ) -thick property [62], in the context of spanners where distances between points play a central role. (Sidiropoulos and Sridhar [59] and Fürer and Kasiviswanathan [33] used the (κ, λ) -thick property [62] in their constructions, but the algorithms become complicated.) Furthermore, except k -th nearest neighbor graphs [66], these criteria have been used for Euclidean spaces only and it is unclear how they could be useful to find sublinear separators of spanners in the more general setting of doubling metrics.

1.1 Our Contributions

Our first contribution is a simple criterion of graphs with sublinear separators in Euclidean and doubling metrics, that we call τ -lanky. Roughly speaking, a graph is τ -lanky if for every ball of radius r , there are at most τ edges of length at least r is cut by the ball. Here, an edge is cut by a ball if exactly one endpoint of the edge is contained in the ball.

Definition 1 (τ -Lanky). *We say that a graph $G = (V, E, w)$ in a metric (X, δ_X) is τ -lanky if for any non-negative r , and for any ball $\mathbf{B}_X(x, r)$ of radius r centered at a vertex $x \in V$, there are at most τ edges of length at least r that are cut by $\mathbf{B}_X(x, r)$.*

We note that in Definition 1, we only count edges of length at least r . There could be $\Omega(n)$ edges of length less than r that are cut by a ball of radius r . We show in Section 3.1 that lankiness implies the existence of a balanced sublinear separator.

THEOREM 1. *Let (X, δ_X) be the Euclidean or a doubling metrics of constant dimension $d \geq 2$, and $G = (V, E, w)$ be an n -vertex graph in (X, δ_X) such that G is τ -lanky. Then, G has a balanced separator of size $O(\tau n^{1-1/d})$. Furthermore, the separator can be found in $O(\tau \cdot n)$ expected time.*

Since τ -lanky property is *closed under taking subgraphs*, i.e., it is a monotone property, Theorem 1 implies that any subgraph of k vertices of a τ -lanky graph has a separator of size $O(\tau k^{1-1/d})$. Thus, as a by-product, we obtain monotone sublinear separators.

Unlike criteria introduced earlier for Euclidean spaces which are mostly based on intersections of objects such as balls or convex shapes, our criterion in Theorem 1 directly related to the edge length of G via the notion of

τ -lanky. As a result, our criterion works better for spanners, where distances (and hence edge lengths) play a central role. Furthermore, as we will show later, for most spanners in this paper, it is relatively easy to see that they are τ -lanky for some constant τ (that depends on the dimension and ϵ). Thus, Theorem 1 can be seen as a positive answer to Question 2.

Our second contribution is to show that greedy spanners of Euclidean graphs, unit ball graphs, and point sets with small fractal dimensions are τ -lanky. Our proof is relatively simple; see Section 4 for details. Since τ -lanky is monotone, Theorem 1 implies that:

THEOREM 2. *Let H be a subgraph of k vertices of the greedy $(1 + \epsilon)$ -spanner for a graph G . Then,*

- (1) *H has a balanced separator of size $O(k^{1-1/d})$ if G is a complete Euclidean graph representing a set of points in \mathbb{R}^d .*
- (2) *H has a balanced separator of size $O(k^{1-1/d})$ if G is a unit ball graph in \mathbb{R}^d .*
- (3) *H has a balanced separator of size $O(k^{1-1/d_f})$ if G is a complete graph representing a set of points in \mathbb{R}^d that has a fractal dimension $d_f \geq 2$.*

Furthermore, the separator of H can be found in $O(k)$ expected time given H .

Theorem 2 provides a positive answer to Question 1, and completely resolves the open problem raised by Eppstein and Khodabandeh [28]. We also give a deterministic linear time algorithm finding a sublinear separator when the spread of the metric is polynomially large (Section 4.4).

Our third contribution is to show that the classical $(1 + \epsilon)$ -spanner of Chan *et al.* [12] which is the first spanner for doubling metrics of constant maximum degree, has a sublinear separator. This result resolves the open problem raised by Abam and Har-Peled [1] a decade ago. We obtain our result by showing that the spanner is τ -lanky for a constant τ . It then follows by Theorem 1 that:

THEOREM 3. *There exists a $(1 + \epsilon)$ -spanner in doubling metrics of dimension $d \geq 2$ with a constant maximum degree such that any subgraph H of k vertices of the spanner has a balanced separator of size $O(k^{1-1/d})$. When $d = 1$, H has a balanced separator of size $O(\log k)$.*

Furthermore, the separator of H can be found in $O(k)$ expected time given H .

Finally, we study the greedy $(1 + \epsilon)$ -spanner for doubling metrics. By taking $r = 0$, the τ -lanky property implies that the maximum degree of a τ -lanky graph G is at most τ . However, Smid [61] showed that the greedy spanner could have a maximum degree up to $n - 1$, and thus, the lankiness parameter τ of the greedy spanner is $n - 1$ in the worst case. On the other hand, the spread of the metric by Smid [61] is $\Delta = 2^{\Omega(n)}$. Therefore, it is natural to ask: Does the greedy spanner for a doubling metric with a *subexponential spread* have a sublinear separator? We answer this question by introducing another property in our criterion: κ -thinness. A graph G is κ -thin if for any two sets of diameter at most r whose distance is at least r has κ edges between them; see Definition 8 for a formal definition. We then show (Theorem 6) that any τ -lanky and κ -thin graph in doubling metric of dimension d has a separator of size $O(\kappa n^{1-1/d} + \tau)$. Indeed, our Theorem 6 is stronger as it holds for a weaker notion of lankiness which does not impose the maximum degree of the graph by τ . The size of the separator now depends on τ additively instead of multiplicatively as in Theorem 1. By showing that the greedy spanner for doubling metrics is τ -lanky and κ -thin for $\tau = O(\log \Delta)$ and $\kappa = O(1)$, we have that the greedy spanner of doubling metrics has a sublinear separator for any subexponential spread, as described in the following theorem.

THEOREM 4. *Let (X, δ_X) be a doubling metric of n points with a constant doubling dimension d . Let G be a $(1 + \epsilon)$ -greedy spanner of (X, δ_X) . G has a balanced separator of size $O(n^{1-1/d} + \log(\Delta))$ when $d \geq 2$ and of size $O(\log(n) + \log(\Delta))$ when $d = 1$.*

1.2 Algorithmic Implications

One of the most popular applications of the spanners is to serve as overlay networks in wireless networks [8, 58, 63]. In these applications, we are often interested in solving computational problems in the spanner, such as shortest path [34, 46], independent set [7, 48], dominating sets [48, 54], connected dominating set [65]. The existence of sublinear separators in spanners implies that we can design provably good algorithms for these problems.

Specifically, by the result of Henzinger *et al.* [41], Theorem 2 and Theorem 3 directly imply that Dijkstra's algorithm can be implemented in expected *linear time* on the spanners in these theorems. This improves upon the (deterministic) algorithm of Eppstein and Khodabandeh [28] that has running time $O(n \log^i(n))$ time for any constant i ; here $\log^{(i)}(x) = \log(\log^{(i-1)}(x))$ is the i -times iterated logarithm.

In graph-theoretic terms, the spanners in Theorem 2 and Theorem 3 have *polynomial expansion* [25]. Har-Peled and Quanrud [40] showed that many *unweighted* optimization problems such as independent set, vertex cover, dominating set, connected dominating set, packing problems, admit a polynomial-time approximation scheme (PTAS) in graphs with polynomial expansion. Thus, all of these problems admit PTAS in our spanners. However, if the graphs have weights on vertices, the algorithm of Har-Peled and Quanrud [40], which is based on local search, does not have any guarantee on the approximation ratio. Indeed, designing a PTAS for vertex-weighted NP-hard problems in graphs with polynomial expansion remains an open problem [24]. We remark that planar graphs and minor-free graphs on which these problems were extensively studied [6, 21, 22, 27] are special cases of graphs with polynomial expansion.

On the other hand, Dvořák [24] showed that if a graph has a polynomial expansion and *bounded maximum degree*, then vertex-weighted optimization problems considered above admit a PTAS. This implies that these problems admit a PTAS on the spanners considered in Theorem 2 and Theorem 3 even when we have weights on the vertices, as these spanners have bounded maximum degree. As approximating the vertex-weighted problems has been studied in many wireless network applications [7, 42, 60, 64], our results could be of interest in these settings.

2 PRELIMINARIES

We denote an edge-weighted graph by $G = (V, E, w)$ where V is the vertex set, E is the edge set, and $w : E(G) \rightarrow \mathbb{R}^+$ is the weight function on the edge set. We denote a minimum spanning tree of G by $\text{MST}(G)$. When the graph is clear from context, we use MST as a shorthand for $\text{MST}(G)$. The *distance* between two vertices p, q in G , denoted by $\delta_G(p, q)$, is the minimum weight of a path between them in G . The diameter of G , denoted by $\text{diam}(G)$, is the maximum pairwise distance in G .

For a subset of vertices $A \subseteq V$, we denote by $G[A]$ the subgraph induced by A . We denote by $G - \{e\}$ a subgraph of G obtained from G by removing the edge e . For a subgraph H of G , we define $w(H) = \sum_{e \in E(H)} w(e)$ to be the total weight of edges in H .

A *c-balanced separator* of a graph G is a subset S of V such that every connected component of $G[V \setminus S]$ has at most $c|V|$ vertices. If c is a constant, we simply refer to S as a separator.

In this paper, we study the spanner obtained by the (path) greedy algorithm. The following fact is well known for greedy spanners (see, e.g., Fact 6.1 in the full version of [45])

Fact 1. *Let H be a t -spanner of a graph $G = (V, E, w)$, then $t \cdot w(x, y) \leq \delta_{H-e}(x, y)$ for any edge $e = (x, y)$ in H .*

We denote a metric space X with a distance function δ_X by (X, δ_X) . The diameter of a point set is the maximum distance between points in the set. A ball $\mathbf{B}_X(p, r)$ centered at p of radius r is the set of points within distance at most r from the center p . We say that a ball $\mathbf{B}_X(p, r)$ cuts an edge if exactly one endpoint of the edge is in $\mathbf{B}_X(p, r)$. When the metric is clear from the context, we drop the subscript X in the notation $\mathbf{B}_X(p, r)$. We say that a set of points P is *r-separated* with $r > 0$ if the distance between any two points in P is at least r .

Given a set of points $P \in (X, \delta_X)$, we say that a subset $N \subseteq P$ is an r -net of P if N is r -separated and for every point $y \in P$, there exists a point $x \in N$ such that $\delta_X(x, y) \leq r$.

Given two sets of points A and B , we denote by $\delta_X(A, B) = \min_{a \in A, b \in B} \delta_X(a, b)$ the distance between A and B .

Definition 2 (c -Separated Pair). *A pair of subsets (A, B) in a metric (X, δ_X) is a c -separated pair for some $c > 0$ if the distance between A and B is at least c times the maximum diameter of A and B .*

If (X, δ_X) is the Euclidean metric, we use $|uv|$ to denote the distance between u and v , and use $|A, B|$ to denote the distance between two point sets A and B . Another metric studied in this paper is doubling metrics.

Definition 3 (Doubling Metric). *A metric space (X, δ_X) has doubling constant λ if any ball of radius R can be covered by at most λ balls of radius $\frac{R}{2}$. The number $\text{ddim} = \log_2 \lambda$ is called the doubling dimension of (X, δ_X) .*

It is well-known that the metric induced by any point set in \mathbb{R}^d has doubling dimension $O(d)$. Doubling metrics and the Euclidean metric of dimension d satisfy the following packing bound.

Lemma 1. *Let $B(p, r)$ be a ball of radius R in a Euclidean/doubling metric of dimension d , and $Y \subseteq B(p, x)$ be an r -separated subset for some $r \leq R$, then $|Y| = 2^{O(d)} \left(\frac{R}{r}\right)^d$.*

We also consider the notion of *fractal dimension* introduced by Sidiropoulos and Sridhard [59].

Definition 4 (Fractal Dimension). *Let P be a set of points in \mathbb{R}^d . P has a fractal dimension d_f if and only if for any two positive numbers $r > 0, R \geq 2r$, any point $p \in \mathbb{R}^d$, and any r -net N of P , $|N \cap B(p, R)| = O((R/r)^{d_f})$.*

3 CRITERIA OF GRAPHS WITH SUBLINEAR SEPARATORS

In this section, we provide two criteria of graphs with sublinear separators in low dimensional Euclidean and doubling metrics: one for bounded degree graphs (Section 3.1), and one for graphs with high vertex degrees (Section 3.2). Our proof uses the following lemma due to Har-Peled and Mendel [39].

Lemma 2 (Lemma 2.4 [39]). *Let P be a set of n points in a metric (X, δ_X) with doubling constant λ_X . There exists a point $v \in P$ and a radius $r \geq 0$ such that (a) $|B_X(v, r) \cap P| \geq \frac{n}{2\lambda_X}$ and (b) $|B_X(v, 2r) \cap P| \leq \frac{n}{2}$. Furthermore, v and r can be found in $O(\lambda_X^3 n)$ expected time.*

Before presenting these criteria in details, we introduce a notion of a *packable metric space* (Definition 5), which captures the packing bound (Lemma 1) in both Euclidean and doubling metrics. The notion of packable metric is very similar to the notion of fractal dimension in Definition 4; the main difference is that we do not restrict a packable metric to being a submetric of an Euclidean metric.

Definition 5 ((η, d) -Packable Metric Space). *A metric (X, δ_X) is (η, d) -packable if for any $r \in (0, 1]$ and any r -separated set $P \subseteq X$ contained in a unit ball, $|P| \leq \eta \left(\frac{1}{r}\right)^d$. We call d the packing dimension of the metric and η the packing constant of the metric.*

A folklore result is that the doubling dimension of the Euclidean metric of dimension d is $O(d)$. However, since the dimension of the metric will appear in the exponent of the separator, treating Euclidean metrics as a special case of doubling metrics would result in a *polynomial loss* in the size of the separator. A strength of our technique is that, we only need the packing bound (Lemma 1) in the construction of the separator. This property allows us to unify the construction of both doubling and Euclidean metrics via packable metrics.

By setting $R = 1$ in Lemma 1, the Euclidean metric and doubling metrics of dimension d are both (η, d) -packable for some $\eta = 2^{O(d)}$. We observe that the packing Lemma 1 also holds for (η, d) -packable metric space as well.

Observation 1. Let B be any ball of radius R in a (η, d) -packable metric space (X, δ_X) , and $P \subseteq B$ be any r -separated set for some r such that $0 < r \leq R$. Then $|P| \leq \eta \left(\frac{R}{r}\right)^d$.

PROOF. Scaling the metric by $\frac{1}{R}$, B has radius 1 and P is (r/R) -separated. The observation now follows from Definition 5. \square

We observe that the doubling dimension of a (η, d) -packable metric (X, δ_X) is also close to the packing dimension of X .

Observation 2. Let Y be a subset of points in an (η, d) -packable metric (X, δ_X) . Let δ_Y be the restriction of the distance function δ on $Y \times Y$. Then (Y, δ_Y) has doubling constant $\lambda_Y \leq \eta 2^d$, and hence doubling dimension $\text{ddim}_Y = d + \log(\eta)$.

PROOF. Let $\mathbf{B}_Y(v, R)$ be a ball of radius R in Y centered at a vertex $v \in Y$. Let $N \subseteq \mathbf{B}_Y(v, R)$ be a $R/2$ -net of $\mathbf{B}_Y(v, R)$. By Definition 5, $|N| \leq \eta 2^d$. Since $\mathbf{B}_Y(v, R)$ can be covered by balls of radius $R/2$ centered at points in N , it follows that $\lambda_Y \leq \eta 2^d$. \square

We say that an edge-weighted graph $G = (V, E, w)$ with n vertices is a graph in a metric space (X, δ_X) if $V \subseteq X$ and for any two vertices $u \neq v \in V$, $w(u, v) = \delta_X(u, v)$.

3.1 Bounded Degree Graphs

We now introduce the *lanky* property, and then we show that lanky graphs have sublinear separators.

Definition 6 (τ -Lanky). We say that a graph $G = (V, E, w)$ in a metric (X, δ_X) is τ -lanky if for any non-negative r , and for any ball $\mathbf{B}_X(x, r)$ of radius r centered at a vertex $x \in V$, there are at most τ edges of length at least r that are cut by $\mathbf{B}_X(x, r)$.

Intuitively, if G is τ -lanky for a constant τ , then for any ball B of radius r , there is only a constant number of edges of length at least r coming out from B . We note that there could be as many as $\Omega(|E|)$ short edges of G that are cut by B . We observe in the following that if G is lanky, it has small degree.

Observation 3. If G τ -lanky, then its maximum degree is at most τ .

PROOF. Let v be any vertex in G , and let $r = 0$. By Definition 1, there are at most τ edges in G of positive length that cut the ball $\mathbf{B}(v, r)$, which only contains v . Thus, the degree v is at most τ . \square

We now show the main theorem in this section: if G is thin and lanky, it has a sublinear separator.

Theorem 5. Let (X, δ_X) is an (η, d) -packable metric space and $G = (V, E, w)$ is an n -vertex graph in (X, δ_X) such that G is τ -lanky. Then, G has a $\left(1 - \frac{1}{\eta 2^{d+1}}\right)$ -balanced separator S such that $|S| = O(\tau \eta 8^d n^{1-1/d})$ when $d \geq 2$ and $|S| = O(\tau \eta 8^d \log n)$ when $d = 1$. Furthermore, S can be found in $O((\eta^3 8^d + \tau)n)$ expected time.

We first show the following lemma, which says that there exists a ball of radius r^* that contains a constant fraction of vertices of G and cut at most $O(n^{1-\frac{1}{d}})$ short edges of G . Our proof uses the random ball technique introduced by Miller, Teng, Thurston, and Vavasis [50]; the same technique was used by Har-Peled in [38] and Sidiropoulos and Sridhar in [59] to construct separators for graphs of various kinds.

Lemma 3. There exists a vertex $v \in V$ and a radius r^* such that:

- (1) $\frac{n}{\eta 2^{d+1}} \leq |\mathbf{B}_X(v, r^*) \cap V| \leq \frac{n}{2}$
- (2) $|E^*| = O(\tau \eta 8^d n^{1-1/d})$ when $d \geq 2$ and $|E^*| = O(\tau \eta 8^d \log n)$ when $d = 1$. Here E^* is the set of all edges in G of length at most r^* that are cut by $\mathbf{B}_X(v, r^*)$.

Furthermore, v and r^* can be found in $O((\eta^3 8^d + \tau)n)$ expected time.

PROOF. Let (V, δ_V) be the submetric of (X, δ_X) induced by V . By Observation 2, (V, δ_V) has doubling constant $\lambda_V = \eta 2^d$. By Lemma 2, we can find in $O(\lambda_V^3 n) = O(\eta^3 8^d n)$ time a vertex v and a radius r such that: (a) $|\mathbf{B}_V(v, r)| \geq \frac{n}{2\lambda_V} \geq \frac{n}{\eta 2^{d+1}}$ and (b) $|\mathbf{B}_V(v, 2r)| \leq \frac{n}{2}$.

Let $\sigma \in [0, 1]$ be chosen uniformly at random. Let $r^* = (1 + \sigma)r$; $r^* \leq 2r$ since $\sigma \leq 1$. By properties (a) and (b), we have that $|\mathbf{B}_X(v, r^*) \cap V| = |\mathbf{B}_V(v, r^*)| \geq |\mathbf{B}_V(v, r)| \geq \frac{n}{\eta 2^{d+1}}$ and that $|\mathbf{B}_X(v, r^*) \cap V| = |\mathbf{B}_V(v, r^*)| \leq |\mathbf{B}_V(v, 2r)| \leq \frac{n}{2}$. Thus, Item (1) is satisfied.

We now bound the expected size of E^* , the set of edges of length at most r^* that are cut by $\mathbf{B}_X(v, r^*)$. Let E' be the set of edges of length at most $2r$ that are cut by $\mathbf{B}_X(v, r^*)$. Then $E^* \subseteq E'$. We will bound the expected size of E' instead, which implies the same bound on the expected size of E^* . We partition E' into two sets M_1, M_2 where M_1 contains every edge of weight at most $rn^{-1/d}$ and $M_2 = E' \setminus M_1$.

Claim 1. Any edge of weight at most $rn^{-1/d}$ is cut by $\mathbf{B}_X(v, r^*)$ with probability at most $n^{-1/d}$.

PROOF. Let (u, v) be an edge of weight at most $rn^{-1/d}$. We observe that an edge is cut by a ball if exactly one of its endpoints is inside the ball. Thus, the probability that (u, v) is cut by the ball $\mathbf{B}_X(p, r^*)$ is equal to the probability of the radius r^* landing between $\delta_X(p, u)$ and $\delta_X(p, v)$. Assume without loss of generality that $\delta_X(p, u) \leq \delta_X(p, v)$. Recall that $r^* = (1 + \sigma)r$ with σ being a random variable chosen uniformly from $[0, 1]$. Therefore,

$$\begin{aligned} \Pr[\delta_X(p, u) \leq r^* \leq \delta_X(p, v)] &= \Pr[\delta_X(p, u)/r \leq (1 + \sigma) \leq \delta_X(p, v)/r] \\ &= \Pr[\delta_X(p, u)/r - 1 \leq \sigma \leq \delta_X(p, v)/r - 1]. \end{aligned}$$

Since σ is chosen uniformly random between 0 and 1, the above probability is at most:

$$(\delta_X(p, v)/r - 1) - (\delta_X(p, u)/r - 1) = (\delta_X(p, v) - \delta_X(p, u))/r \leq \delta_X(u, v)/r \leq n^{-1/d},$$

as desired. \square

By Observation 3, $|E| \leq \tau n/2$, and hence there are at most $\tau n/2$ edges of weight at most $rn^{-1/d}$. By Claim 1, it holds that:

$$\mathbb{E}(|M_1|) \leq n^{-1/d}(\tau n)/2 = \tau n^{1-1/d}/2 \quad (1)$$

We now bound the expected size of M_2 . For each $i \in [1, \lceil \log(2n^{-1/d}) \rceil]$, we define $r_i = 2^{i-1}n^{-1/d}r$ and a set of edges $M_2^i = \{(u, v) \in M_2 \mid r_i < w(u, v) \leq r_{i+1}\}$. Observe that $M_2 = \bigcup_{i=1}^{\lceil \log(2n^{-1/d}) \rceil} M_2^i$. Let \mathcal{N}_i be a $(r_i/2)$ -net of $\mathbf{B}_X(v, 2r)$. By Observation 1, we have that:

$$|\mathcal{N}_i| \leq \eta \left(\frac{2r}{r_i/2} \right)^d = \eta \left(\frac{4r}{2^{i-1}n^{-1/d}r} \right)^d = \eta \frac{8^d n}{2^{id}} \quad (2)$$

Let \mathcal{N}_i be set of balls with center in \mathcal{N}_i and radius $2^{i-2}n^{-1/d}r$. Since $r^* \leq 2r$, \mathcal{N}_i covers all points in $\mathbf{B}_X(v, r^*)$. Furthermore, since each ball in \mathcal{N}_i has diameter at most r_i , every edge of length at least r_i (including edges in M_2^i) will be cut by at least one ball (and at most two balls) in \mathcal{N}_i . Note that the construction of \mathcal{N}_i is deterministic.

Let E_i be the set of edges of length at least r_i and at most $2r_i$ that are cut by at least one ball in \mathcal{N}_i . Since G is τ -lanky, there are at most τ edges of length at least $r_i/2$ that are cut by a ball in \mathcal{N}_i . It follows that

$$|E_i| \leq \tau |\mathcal{N}_i| = \tau \eta \frac{8^d n}{2^{id}} \quad (3)$$

by Equation (2). Since every edge in e in E_i is cut by $\mathbf{B}_X(v, r^*)$ with probability at most $w(e)/r \leq r_{i+1}/r$. Thus, it follows that $\mathbb{E}(|M_2^i|) \leq |E_i|r_{i+1}/r$. By the linearity of expectation and Equation (3), we have:

$$\begin{aligned} \mathbb{E}(|M_2|) &= \sum_{i=1}^{\lceil \log(2n^{1/d}) \rceil} \mathbb{E}(|M_2^i|) = \sum_{i=1}^{O(\log n)} \tau\eta \frac{8^d n}{2^{id}} \cdot \frac{r_{i+1}}{r} \\ &= \tau\eta 8^d \sum_{i=1}^{O(\log n)} \frac{n}{2^{id}} \cdot \frac{2^i n^{1/d} r}{r} = \tau\eta 8^d n^{1-1/d} \sum_{i=1}^{O(\log n)} \left(\frac{1}{2^{d-1}}\right)^i \end{aligned} \quad (4)$$

We conclude that $\mathbb{E}(|M_2|) = O(\tau\eta 8^d \log n)$ when $d = 1$ and $\mathbb{E}(|M_2|) = O(\tau\eta 8^d n^{1-1/d})$ when $d \geq 2$. By Equation (1), we have that:

$$\mathbb{E}[|E^*|] = \tau n^{1-1/d} / 2 + O(\tau\eta 8^d n^{1-1/d}) = O(\tau\eta 8^d n^{1-1/d}) \quad (5)$$

Thus, by Markov's inequality, with a constant probability, E^* has size $O(\tau\eta 8^d n^{1-1/d})$. Since the running time to find E^* for each random choice of r^* is $O(|E|) = O(\tau n)$, by repeating many times until we find r^* such that E^* satisfies Item (2), the expected running time is still $O(\tau n)$. Recall that v can be found in $O(\eta^3 8^d n)$ expected time. Thus, the total expected running time is $O((\eta^3 8^d + \tau)n)$ as claimed. \square

We are now ready to prove Theorem 5.

PROOF OF THEOREM 5. Let v and r^* the center and the radius of the ball $\mathbf{B}_X(v, r^*)$ as in Lemma 3. We define the function $g(d, n)$ as follows $g(d, n) = O(\log(n))$ when $d = 1$ and $g(d, n) = O(n^{1-1/d})$ when $d \geq 2$. Let \tilde{E} be the set of edges that are cut by $\mathbf{B}_X(v, r^*)$. By Item (2) in Lemma 3, there are at most $O(\tau\eta 8^d g(d, n))$ edges in \tilde{E} of length at most r^* . Since G is τ -lanky, there are at most τ edges of length at least r^* in \tilde{E} . Thus, it holds that

$$|\tilde{E}| = O(\tau\eta 8^d g(d, n) + \tau) \quad (6)$$

Let S be the set of all endpoints of edges in \tilde{E} . Then removing S from G disconnects the set of points in $\mathbf{B}_X(v, r^*)$ from the set of points outside $\mathbf{B}_X(v, r^*)$. Thus, S is a $\left(1 - \frac{1}{\eta 2^{d+1}}\right)$ -balanced separator by Item (1) of Lemma 3. The running time to construct S is dominated by the running time to construct v and r^* , which is $O((\eta^3 8^d + \tau)n)$ by Lemma 3. \square

Observe by the definition that τ -lanky property is closed under taking subgraph: if G is τ -lanky, then any subgraph H of G is also τ -lanky. Thus, by Theorem 5, we have:

Corollary 1. *Let (X, δ_X) is an (η, d) -packable metric space and $G = (V, E, w)$ is a graph in (X, δ_X) such that G is τ -lanky. Let H be any subgraph of G with k vertices. Then, H has a $\left(1 - \frac{1}{\eta 2^{d+1}}\right)$ -balanced separator S such that $|S| = O(\tau\eta 8^d k^{1-1/d})$ when $d \geq 2$ and $|S| = O(\tau\eta 8^d \log k)$ when $d = 1$. Furthermore, S can be found in $O((\eta^3 8^d + \tau)k)$ expected time given H .*

3.2 Graphs with High Vertex Degrees

By Observation 3, the τ -lanky property implies that the maximum degree is bounded by τ . In this section, we introduce another criterion that could be used for graphs with high vertex degrees. Our criterion is based on two properties: weakly lanky and thin.

Definition 7 (Weakly τ -lanky). *A graph $G = (V, E, w)$ in a metric (X, δ_X) is weakly τ -lanky if for any non-negative r , and for any ball $\mathbf{B}_X(x, r)$ of radius r centered at a vertex $x \in V$, there are at most τ vertices inside $\mathbf{B}_X(x, r)$ that are incident to all edges of length at least r cut by $\mathbf{B}_X(x, r)$.*

Observe that τ -lanky implies weakly τ -lanky but the converse statement does not hold. There could be $\Omega(n)$ edges of length at least r incident to a single vertex that is cut by a ball of radius r . We observe that the same proof in Section 3.1 is applicable to yield a separator of size $O(\tau n^{1-1/d})$ for a weakly τ -lanky graph. However, in this section, we look for a separator of size $O(n^{1-1/d} + \tau)$. To this end, we introduce another property that we call κ -thin.

Definition 8 (κ -Thin). *A graph $G = (V, E, w)$ in a metric (X, δ_X) is κ -thin if for any 1-separated pair (A, B) of V , there are at most κ edges between A and B .*

Our goal in this section is to show the following theorem, which implies that if G is κ -thin and weakly τ -lanky, it has sublinear separator.

THEOREM 6. *Let (X, δ_X) is an (η, d) -packable metric space and $G = (V, E, w)$ is a graph in (X, δ_X) that is weakly τ -lanky and κ -thin. Then, G has a $(1 - \frac{1}{\eta^{2d+1}})$ -balanced separator S such that $|S| = O(\eta^{O(1)} 2^{O(d)} \kappa n^{1-1/d} + \tau)$ when $d \geq 2$ and $|S| = O(\eta^{O(1)} 2^{O(d)} \kappa \log n + \tau)$ when $d = 1$. Furthermore, S can be constructed in $O((\eta^3 8^d + \kappa)n)$ expected time.*

In the proof of Theorem 6, we follow the same construction presented in Section 3.1: take a ball $B_X(v, r^*)$ of random radius $r^* \in [r, 2r]$ centered at a specific vertex v , and construct the separator S by taking the endpoint inside $B_X(v, r^*)$ of every edge cut by the ball. We then show that in expectation, the size of S is small. There are three places in the proof of Theorem 5 where the τ -lanky property is used to bound the size of the separator; here we point out how the κ -thin property could be used to replace τ . First, the number of edges in E is bounded by $\tau n/2$, and this fact is used to bound the expected size of edges M_1 in Equation (1). We show in Lemma 4 below that κ -thin property, implies that $|E| = O(\kappa n)$. Second, the number of edges in E of size at most $2r_i$ for some radius r_i cut by some ball of radius at least r_i is bounded by τ , and this fact is used to bound the size of E_i in Equation (3). We show in Lemma 5 that the number of such edges is $O(\kappa)$ if G is κ -thin, thereby, removing the dependency on τ . Finally, the τ -lanky property is used to bound the number of edges of length at least $2r$ cut by $B_X(v, r^*)$, which contribute an additive τ in the size of \tilde{E} in Equation (6). Since we can tolerate the additive term τ in Theorem 6, we do not need to do anything.

Now we focus on showing that κ -thin property implies that G has $O(\kappa n)$ edges. Our proof uses well-separated pair decomposition. An s -well-separated pair decomposition (s -WSPD) for a point set P for some non-negative parameter s in a metric (X, δ_X) is the set of s -separated pairs $\mathcal{P} = \{(A_1, B_1), \dots, (A_m, B_m)\}$ such that (i) $A_i, B_i \subseteq P, A_i \cap B_i = \emptyset$ for all $i \in [1, m]$ and (ii) for every pair of points (p, q) , there exists a some $i \in [1, m]$ such that either $p \in A_i, q \in B_i$ or $q \in A_i$ and $p \in B_i$.

Lemma 4. *If an n -vertex graph $G = (V, E, w)$ in an (η, d) -packable metric (X, δ_X) is κ -thin, then $|E| = O(\eta^{O(1)} 2^{O(d)} \kappa n)$.*

PROOF. Let $\mathcal{D} = \{(A_1, B_1), (A_2, B_2), \dots, (A_m, B_m)\}$ be a 1-WSPD of V with minimum number of pairs. Har-Peled and Mendel (Lemma 5.1 [39]) show that for any n -point set P in a metric of doubling dimension ddim has an s -well-separated pair decomposition \mathcal{P} with $|\mathcal{P}| = 2^{O(\text{ddim})} s^{\text{ddim}} n$ for any $s \geq 1$. Since X has doubling dimension $d + \log(\eta)$ by Observation 2, $|\mathcal{D}| = \eta^{O(1)} 2^{O(d)} n$. For each $i \in [1, m]$, let H_i be the set of edges between A_i and B_i . Since (A_i, B_i) is 1-separated, $|H_i| = \kappa$ by Definition 8. By the definition of well-separated pair decomposition, for each edge $e = (u, v) \in E$, there exists an index i such that (u, v) or (v, u) is in $A_i \times B_i$. Thus, $\bigcup_{i=1}^m H_i = E$. It follows that:

$$|E| = \left| \bigcup_{i=1}^m H_i \right| \leq \sum_{i=1}^m |H_i| = O(\eta^{O(1)} 2^{O(d)} \kappa n), \quad (7)$$

as claimed. \square

Next, we show that, if G is κ -thin, then the number of edges of length at least r and at most $2r$ cut by a ball of radius r is $O(\kappa)$.

Lemma 5. *Let $\mathbf{B}_X(x, r)$ be any ball of radius r centered at some point $x \in X$. If G is κ -thin, there are at most $O(\eta^2 2^{O(d)} \kappa)$ edges of G of length at least r and at most $2r$ that are cut by $\mathbf{B}_X(x, r)$.*

PROOF. Let E_{short} be the set of edges of length in $[r, 2r]$. Observe by the triangle inequality that for every edge $(u, v) \in E_{\text{short}}$, both endpoints u and v are in $\mathbf{B}_X(x, 3r)$. Let $\mathcal{B}_{\text{in}} (\mathcal{B}_{\text{out}})$ be the set of balls obtained by taking balls of radius $\frac{r}{6}$ centered at points in a $\frac{r}{6}$ -net of $\mathbf{B}_X(x, r)$ ($\mathbf{B}_X(x, 3r)$). By Observation 1, $|\mathcal{B}_{\text{in}}| = O(\eta \left(\frac{r}{r/6}\right)^d) = O(\eta 6^d)$ and $|\mathcal{B}_{\text{out}}| = O(\eta \left(\frac{3r}{r/6}\right)^d) = O(\eta 18^d)$ that covers $\mathbf{B}_X(x, 3r)$. Let $(B_{\text{in}}, B_{\text{out}}) \in \mathcal{B}_{\text{in}} \times \mathcal{B}_{\text{out}}$ be a pair of balls such that there exists an edge $(u, v) \in E_{\text{short}}$ between them. By the triangle inequality, $\delta_X(B_{\text{in}}, B_{\text{out}}) \geq \delta_X(u, v) - \text{diam}(B_{\text{in}}) - \text{diam}(B_{\text{out}}) \geq r - 2r/6 - 2r/6 = r/3$. Hence, $\frac{\delta_X(B_{\text{in}}, B_{\text{out}})}{\max\{\text{diam}(B_{\text{in}}, B_{\text{out}})\}} \geq 1$, which implies that $(B_{\text{in}}, B_{\text{out}})$ is 1-separated. Thus, there are at most κ edges between B_{in} and B_{out} by the definition of κ -thin (Definition 8). It follows that $|E_{\text{short}}| = \kappa |\mathcal{B}_{\text{in}} \times \mathcal{B}_{\text{out}}| = O(\eta^2 \kappa 108^d) = O(\eta^2 2^{O(d)} \kappa)$, as claimed. \square

Next, we show the following lemma, which is analogous to Lemma 3. The key difference is that the size of E^* , the set of edges of length at most r^* cut by $\mathbf{B}_X(v, r^*)$.

Lemma 6. *Let $G = \{V, E, w\}$ is a weakly τ -lanky and κ -thin graph in an (η, d) -packable metric space (X, δ_X) . There exists a vertex $v \in V$ and a radius r^* such that:*

- (1) $\frac{n}{\eta^{2d+1}} \leq |\mathbf{B}_X(v, r^*) \cap V| \leq \frac{n}{2}$
- (2) $|E^*| = O(\eta^{O(1)} 2^{O(d)} \kappa n^{1-1/d})$ when $d \geq 2$ and $|E^*| = O(\eta^{O(1)} 2^{O(d)} \kappa \log n)$ when $d = 1$. Here E^* is the set of all edges in G of length at most r^* that are cut by $\mathbf{B}_X(v, r^*)$.

Furthermore, v and r^* can be found in $O(\eta^{O(1)} 2^{O(d)} \kappa n)$ expected time.

PROOF. We reuse the notation in the proof of Lemma 3. Specifically, we construct v, r and r^* as in Lemma 3; Item (1) follows directly from the construction. Next, we bound the size of E^* following the same strategy: partitioning E^* into two set M_1 and M_2 where M_1 is the set of edges with length at most $rn^{1-1/d}$ and M_2 contains the other edges. The same argument in Lemma 3, specifically Equation (1), yields:

$$\mathbb{E}(|M_1|) \leq n^{-1/d} |E| \leq O(\eta^{O(1)} 2^{O(d)} \kappa) n \quad (8)$$

by Lemma 4.

To bound the expected size of M_2 , we partition M_2 into $\{M_1, \dots, M_{\lceil \log(2n^{-1/d}) \rceil}\}$, where M_2^i be the set of edges with length in $(r_i, r_{i+1}]$; here $r_i = 2^{i-1} n^{1-1/d} r$. Following the same construction in Lemma 3, we construct a set of \mathcal{N}_i of balls of radius $2^{i-2} n^{-1/d} r = r_i/2$ covering $\mathbf{B}_X(v, 2r)$. By Equation (2), $|\mathcal{N}_i| = \eta \frac{8^d n}{2^{id}}$. Let E_i be the set of edges of length at least r_i and at most $2r_i$ cut by at least one ball in \mathcal{N}_i . By Lemma 5, it follows that:

$$|E_i| \leq \tau |\mathcal{N}_i| = O(\eta^2 2^{O(d)} \kappa) \eta \frac{8^d n}{2^{id}} = O\left(\eta^3 2^{O(d)} \frac{n}{2^{id}}\right) \quad (9)$$

The rest of the argument is exactly the same as the proof of Lemma 3 (with different constants), that yields $\mathbb{E}[|M_2|] = O(\kappa \eta^3 2^{O(d)} n^{1-1/d})$ when $d \geq 2$ and $\mathbb{E}[|M_2|] = O(\kappa \eta^3 2^{O(d)} \log(n))$ when $d = 1$. This in turn implies $|E^*| = O(\eta^{O(1)} 2^{O(d)} \kappa n^{1-1/d})$ when $d \geq 2$ and $|E^*| = O(\eta^{O(1)} 2^{O(d)} \kappa \log n)$ when $d = 1$ as desired. The expected running time bound follow the same line of reasoning in the proof of Lemma 3. \square

We are now ready to complete the proof of Theorem 6.

PROOF OF THEOREM 6. Let v and r^* be the center and radius of the ball $\mathbf{B}_X(v, r^*)$ in Lemma 6. Let $g(d, n) = O(\log n)$ if $d = 1$ and $g(d, n) = O(n^{1-1/d})$ if $d \geq 2$. Let S be the set of all endpoints in $\mathbf{B}_X(v, r^*)$ of edges in \tilde{E} . Recall that \tilde{E} is the set of edges that are cut by $\mathbf{B}_X(v, r^*)$. Removing S from G disconnects $\mathbf{B}_V(v, r^*)$ and $V \setminus \mathbf{B}_V(v, r^*)$. By Item (2) of Lemma 6, there are $O(\eta^{O(1)} 2^{O(d)} \kappa)g(d, n)$ edges of length at most r^* cut by $\mathbf{B}_X(v, r^*)$. By Definition 7, there are at most τ points in $\mathbf{B}_X(v, r^*)$ that are incident to all edges of length at least r^* in \tilde{E} . Thus, $|S| = O(\eta^{O(1)} 2^{O(d)} \kappa)g(d, n) + \tau$. By Item (1) of Lemma 6, S is a $\left(1 - \frac{1}{\eta^{2^{d+1}}}\right)$ -balanced separator of G . The expected time to find v and r^* is $O(\eta^{O(1)} 2^{O(d)} \kappa n)$ by Lemma 6. Since the time to find all endpoints in $\mathbf{B}_X(v, r^*)$ of edges in \tilde{E} is $O(|E|) = O(\eta^{O(1)} 2^{O(d)} \kappa n)$, the total time complexity to construct S is $O(\eta^{O(1)} 2^{O(d)} \kappa n)$. \square

Since weakly τ -lanky and κ -thin properties are closed under taking subgraphs, we have:

Corollary 2. *Let (X, δ_X) is an (η, d) -packable metric space and $G = (V, E, w)$ is a graph in (X, δ_X) such that G is weakly τ -lanky and κ -thin. Let H be any subgraph of G with k vertices. Then, H has a $\left(1 - \frac{1}{\eta^{2^{d+1}}}\right)$ -balanced separator S such that $|S| = O(\eta^{O(1)} 2^{O(d)} \kappa k^{1-1/d} + \tau)$ when $d \geq 2$ and $|S| = O(\eta^{O(1)} 2^{O(d)} \kappa \log k + \tau)$ when $d = 1$. Furthermore, S can be found in $O(\eta^{O(1)} 2^{O(d)} \kappa k)$ expected time given H .*

4 SEPARATORS OF GREEDY SPANNERS FOR GRAPHS IN EUCLIDEAN SPACES

In this section, we prove Theorem 2. Specifically, in Section 4.1, we focus on the greedy spanners of point sets in Euclidean spaces. In Section 4.2, we focus on the greedy spanners of point sets that have a low fractal dimension. In Section 4.3, we focus on the greedy spanners of unit ball graphs (UBGs).

4.1 Euclidean Spaces

First, we focus on proving Theorem 7 below, which implies Item (1) of Theorem 2.

THEOREM 7 (SEPARATORS FOR SUBGRAPHS OF GREEDY SPANNERS). *Let P be a given set of points in the d -dimensional Euclidean space and G be the greedy $(1 + \epsilon)$ -spanner of P for some $\epsilon \in (0, 1/2]$. Then any k -vertex subgraph H of G has a $\left(1 - \frac{1}{2^{O(d)}}\right)$ -balanced separator S of size $O(2^{O(d)} \epsilon^{1-2d} k^{1-1/d})$. Furthermore, S can be found in expected $O((2^{O(d)} + \epsilon^{1-2d})k)$ time given H .*

Our goal in the proof of Theorem 7 is to show that the greedy $(1 + \epsilon)$ -spanner G is τ -lanky for some τ that depends on ϵ and d only. To this end, we rely on the following lemma, which bounds the number of edges in the greedy $(1 + \epsilon)$ -spanner between subsets of points X and Y , where X has a small diameter and Y is sufficiently far from X .

Lemma 7. *Let P be a set of n points in \mathbb{R}^d for a constant d and G be a greedy $(1 + \epsilon)$ -spanner of P for $\epsilon \in (0, 1/2]$. Let X and Y be two subsets of P such that $\text{diam}(X) \leq \frac{\epsilon R}{12}$ and $|X, Y| \geq R$. Then G has $O(\epsilon^{1-d})$ edges between X and Y .*

Our proof of Lemma 7 uses the following lemma.

Lemma 8. *Given a point x in \mathbb{R}^d and an arbitrarily small number $\epsilon \in (0, 1/2]$. Let C be a cone with apex x and angle at most $\epsilon/8$. Then, for any two points y and z in C such that $|xy| \geq |xq|$, $|xy| \geq |yq| + (1 - \epsilon/4)|xq|$.*

PROOF. Consider the triangle xyq and let h be the projection of q on xy . Since $|xy| \geq |xq|$, h is in the segment between x and y . Let $\gamma = \angle yxq$; see Figure 1(b). Observe that $\gamma \leq \theta = \epsilon/8 < 1$ when $\epsilon \in (0, 1/2)$. We have:

$$\begin{aligned} |xy| - |yq| &\geq |xy| - |yh| - |hq| = |xh| - |hq| = (\cos \gamma - \sin \gamma)|xq| \\ &\geq (\cos \theta - \sin \theta)|xq| \quad (\cos(x) - \sin(x) \text{ is monotonically decreasing for } x \in (0, 1)) \\ &\geq (1 - 2\theta)|xq| = (1 - \epsilon/4)|xq|. \end{aligned}$$

The penultimate inequality is due to that $\cos(x) - \sin(x) \geq 1 - 2x$ every $x \in [0, 1]$. Thus, we conclude that $|xy| \geq |yq| + (1 - \epsilon/4)|xq|$ as claimed. \square

We are now ready to prove Lemma 7.

PROOF OF LEMMA 7. Let x be an (arbitrary) point in X . We cover the space \mathbb{R}^d by a minimum number of cones with apex x , each with an angle $\theta = \epsilon/8$; the number of cones in the cover is $O(d^{(d+1)/2}(\pi/\theta)^{d-1}) = O(\epsilon^{1-d})$ [51]. We show below that for each cone, there is at most *one edge* in G between X and the subset of points of Y in the cone. It then follows that the number of edges in G between X and Y is $O(\epsilon^{1-d})$.




Fig. 1. (a) Points y and q of Y are in the same cone bounded by two red dashed lines at apex x and (b) the triangle xqy with $\angle yxq = \gamma$.

Suppose for a contradiction that there two edges (s, q) and (t, y) in G , where $s, t \in X$ and $q, y \in Y$. W.l.o.g, we assume $|xy| \geq |xq|$; see Figure 1(a). Our goal is to show that:

$$(1 + \epsilon)|ty| > \delta_{G-(t,y)}(t, y) \quad (10)$$

which contradicts Fact 1. From Lemma 8, we obtain $|xy| \geq |yq| + (1 - \epsilon/4)|xq|$. Observe that:

$$\begin{aligned} (1 + \epsilon)|ty| &\geq (1 + \epsilon)(|xy| - |tx|) && \text{(by the triangle inequality)} \\ &\geq (1 + \epsilon)(|yq| + (1 - \epsilon/4)|xq|) - (1 + \epsilon)|tx| && \text{(by Lemma 8)} \\ &> (1 + \epsilon)|yq| + (1 + \epsilon/2)|xq| - (1 + \epsilon)|tx| && \text{(since } 0 < \epsilon < 1\text{)} \\ &\geq (1 + \epsilon)|yq| + (1 + \epsilon/2)(|sq| - |xs|) - (1 + \epsilon)|tx| && \text{(by the triangle inequality)} \\ &\geq (1 + \epsilon)|yq| + |sq| + \epsilon R/2 - (1 + \epsilon/2)|xs| - (1 + \epsilon)|tx| && \text{(since } |sq| \geq |X, Y| \geq R\text{)} \\ &\geq (1 + \epsilon)|yq| + |sq| + (1 + \epsilon)|ts| + \epsilon R/2 - \underbrace{((1 + \epsilon)|ts| + (1 + \epsilon/2)|xs| + (1 + \epsilon)|tx|)}_{\leq 3(1+\epsilon)\epsilon R/12 = \epsilon(1+\epsilon)R/4 \text{ since } \text{diam}(X) \leq \epsilon R/12} \\ &\geq (1 + \epsilon)|yq| + |sq| + (1 + \epsilon)|ts| + \epsilon R/2 - (1 + \epsilon)\epsilon R/4 \\ &> (1 + \epsilon)|yq| + |sq| + (1 + \epsilon)|ts| && \text{(since } \epsilon \leq 1/2\text{)} \end{aligned}$$

Since G is a $(1 + \epsilon)$ -spanner, $(1 + \epsilon)|yq| \geq \delta_G(y, q)$ and $(1 + \epsilon)|ts| \geq \delta_G(t, s)$. Thus, we have:

$$(1 + \epsilon)|ty| > \delta_G(y, q) + \delta_G(s, q) + \delta_G(t, s)$$

Observe that $\max\{|yq|, |st|\} \leq \text{diam}(X) \leq \epsilon R/12 \leq \epsilon|X, Y|/12 \leq \epsilon|ty|/12$. This implies $(1 + \epsilon) \max\{|yq|, |st|\} < |ty|$ when $\epsilon \in (0, 1/2]$; that is, edge (t, y) cannot be in the shortest paths from y to q and from s to t in G . Thus, $\delta_G(y, q) = \delta_{G-(t,y)}(y, q)$ and $\delta_G(t, s) = \delta_{G-(t,y)}(t, s)$. Furthermore, since (s, q) is an edge in G , $\delta_G(s, q) = \delta_{G-(t,y)}(s, q)$. Thus, it holds that:

$$(1 + \epsilon)|ty| > \delta_{G-(t,y)}(y, q) + \delta_{G-(t,y)}(s, q) + \delta_{G-(t,y)}(t, s) = \delta_{G-(t,y)}(t, y).$$

That is, Equation (10) holds as claimed. \square

Lemma 9. *Let P be a set of n points in \mathbb{R}^d for a constant d and G be a greedy $(1 + \epsilon)$ -spanner of P for $\epsilon \in (0, 1/2]$. Then G is $O(\epsilon^{1-2d})$ -lanky.*

PROOF. Let $\mathbf{B}(x, r)$ be a ball of radius r centered at a vertex $x \in P$. Let \tilde{E} be the set of edges of length at least r that are cut by $\mathbf{B}(x, r)$. Let \mathcal{B} be a minimum collection of balls of radius $\frac{\epsilon r}{48}$ that covers $\mathbf{B}(x, r)$. By Lemma 1, $|\mathcal{B}| = O(\epsilon^{-d})$. We observe that, since every edge in \tilde{E} has length at least the diameter of every ball in \mathcal{B} , each edge in \tilde{E} must be cut by some ball in \mathcal{B} .

Consider a ball $B \in \mathcal{B}$, and $e = (u, v) \in \tilde{E}$ be an edge that is cut by B . W.l.o.g., we assume that $u \in B$. Then by the triangle inequality, $|v, B| \geq w(e) - \text{diam}(B) \geq r - \frac{\epsilon r}{24} \geq r/2$ when $\epsilon \leq 1/2$. Since $\text{diam}(B) = \frac{\epsilon r}{24}$, by Lemma 7, there are at most $O(\epsilon^{1-d})$ such edges e in \tilde{E} that cut B . Thus, $|\tilde{E}| \leq |\mathcal{B}|O(\epsilon^{1-d}) = O(\epsilon^{1-2d})$; the lemma follows. \square

PROOF OF THEOREM 7. When $d = 1$, the theorem holds trivially. Thus, we assume that $d \geq 2$. By Lemma 9, G is τ -lanky with $\tau = O(\epsilon^{1-2d})$. Since the Euclidean space of dimension d is (η, d) -packable with $\eta = 2^{O(d)}$ by Lemma 1, by Corollary 1, H has a separator S of size $O(\epsilon^{1-2d} 2^{O(d)} k^{1-1/d})$. Also by Corollary 1, S can be found in $O((2^{O(d)} 8^d + \epsilon^{1-2d})k)$ time. \square

4.2 Point Sets of Low Fractal Dimensions

We now consider a point set P in \mathbb{R}^d that has a small fractal dimension. Since P is still a set of points in \mathbb{R}^d , Lemma 9 remains true for the greedy $(1 + \epsilon)$ -spanner G of P . The only difference is that, the metric induced by P , by Definition 4, is an $(O(1), d_f)$ -packable metric. This implies a separator of smaller size.

THEOREM 8. *Let P be a given set of points in \mathbb{R}^d that has fractal dimension d_f , and G be the greedy $(1 + \epsilon)$ -spanner of P . Then any k -vertex subgraph H of G has a $\left(1 - \frac{1}{2^{O(d_f)}}\right)$ -balanced separator S of size $O(2^{O(d_f)} \epsilon^{1-2d} k^{1-1/d_f})$. Furthermore, S can be found in expected $O((2^{O(d_f)} + \epsilon^{1-2d})k)$ time given H .*

PROOF. By Definition 4 and Definition 5, the metric induced by P is $(O(1), d_f)$ -packable. By Lemma 9, G is $O(\epsilon^{1-2d})$ -lanky. By Corollary 1, every k -vertex subgraph H of G has a $\left(1 - \frac{1}{2^{O(d_f)}}\right)$ -balanced separator of size $O(2^{O(d_f)} \epsilon^{1-2d} k^{1-1/d_f})$ that can be found in expected $O((2^{O(d_f)} + \epsilon^{1-2d})k)$ time. \square

We note that Theorem 8 implies Item (3) in Theorem 2.

4.3 Unit Ball Graphs

In this section, we show that the greedy spanners of UBGs have sublinear separators, as described in the following theorem. We note that Theorem 9 implies Item (2) of Theorem 2.

THEOREM 9. *Let G be the greedy $(1 + \epsilon)$ -spanner of a unit ball graph in \mathbb{R}^d . Then any k -vertex subgraph H of G has a $\left(1 - \frac{1}{2^{O(d)}}\right)$ -balanced separator S of size $O(2^{O(d)} \epsilon^{1-2d} k^{1-1/d})$. Furthermore, S can be found in expected $O((2^{O(d)} + \epsilon^{1-2d})k)$ time given H .*

Unit ball graphs are intersection graphs of unit balls. Since we scale the metric so that the minimum inter-point distance is 1, we assume that the unit ball graphs are the intersection graphs of balls of radius μ for some positive μ .

We consider the greedy spanners of unit ball graphs. We prove that a greedy spanner of a unit ball graph in a d -dimensional Euclidean space also admits τ -lanky property. We prove a lemma which is analogous to Lemma 7.

Lemma 10. *Let P be a set of n points in \mathbb{R}^d for a constant d . Let G_B be the unit ball graphs with the set of centers P and G be a greedy $(1 + \epsilon)$ -spanner of G_B for $\epsilon \in (0, 1/2]$. Let X and Y be two subsets of P such that $\text{diam}(X) \leq \frac{\epsilon R}{12}$ and $|X, Y| \geq R$. Then G has $O(\epsilon^{1-d})$ edges between X and Y .*

PROOF. Let x be an arbitrary point in X . We partition the space \mathbb{R}^d into $O(\epsilon^{1-d})$ cones with apex x and angle $\theta = \epsilon/8$. We prove that for each cone C , there is at most one edge in G from X to $Y \cap C$. We follow the same proof strategy of Lemma 7. Assume that there are two edges (s, q) and (t, y) in G with $s, t \in X$ and $q, y \in Y$ (see Figure 1). W.o.l.g., we suppose that $|xy| \geq |xq|$. Our goal is to show that $(1 + \epsilon)d\delta_{G_B}(t, y) > \delta_{G-(t,y)}(t, y)$, which contradicts Fact 1.

First, we prove that $(q, y) \in E(G_B)$. By Lemma 8, we obtain $|xy| - |yq| \geq (1 - \epsilon/4)|xq|$. Using the triangle inequality, we have that:

$$\begin{aligned} |yq| &\leq \underbrace{|xy|}_{\leq |ty|+|xt|} - (1 - \epsilon/4) \underbrace{|xq|}_{\geq |sq|-|sx|} \leq |ty| + \underbrace{|xt|}_{\leq \frac{\epsilon R}{12}} + (1 - \epsilon/4) \underbrace{|sx|}_{\leq \frac{\epsilon R}{12}} - (1 - \epsilon/4) \underbrace{|sq|}_{\geq R} \\ &\leq |ty| + \frac{\epsilon R}{12} + (1 - \epsilon/4) \frac{\epsilon R}{12} - (1 - \epsilon/4)R \leq |ty| + \frac{\epsilon R}{12} + \frac{\epsilon R}{12} - \frac{3R}{4} \leq |ty| \leq 2\mu \quad (\text{since } \epsilon \leq 1/2). \end{aligned} \quad (11)$$

The bound $|yq| \leq 2\mu$ implies that y is adjacent to q in G_B . For any pair of vertices (v_1, v_2) with $v_1, v_2 \in X, v_1 \neq v_2$, we observe that $|v_1 v_2| \leq \frac{\epsilon R}{12} \leq R \leq |ty| \leq 2\mu$. Hence, v_1 is adjacent to v_2 in G_B . Then, we have $(t, s) \in E(G_B)$. In Lemma 7, we have prove that $(1 + \epsilon)|ty| > (1 + \epsilon)|yq| + |sq| + (1 + \epsilon)|ts|$. Since G is a $(1 + \epsilon)$ -spanner of G_B , $\delta_G(y, q) \leq (1 + \epsilon)\delta_{G_B}(y, q) = (1 + \epsilon)|yq|$ and $\delta_G(t, s) \leq (1 + \epsilon)\delta_{G_B}(t, s) = (1 + \epsilon)|ts|$. Then, we have:

$$(1 + \epsilon)|ty| > \delta_G(y, q) + \delta_G(q, s) + \delta_G(s, t). \quad (12)$$

The edge (t, y) cannot be in the shortest path from s to t in G or the path from y to q because $|ty| > \max\{\delta_{G_B}(y, q), \delta_{G_B}(s, t)\}$. Since $(q, s) \in G$, we obtain $\delta_{G-(t,y)}(q, s) = \delta_G(q, s)$. Hence, we have:

$$(1 + \epsilon)|ty| > \delta_{G-(t,y)}(y, q) + \delta_{G-(t,y)}(q, s) + \delta_{G-(t,y)}(s, t) \geq \delta_{G-(t,y)}(t, y). \quad (13)$$

Thus, $(1 + \epsilon)\delta_{G_B}(t, y) > \delta_{G-(t,y)}(t, y)$, a contradiction. \square

The proof of Lemma 9 implies that G is $O(\epsilon^{1-2d})$ -lanky. The same proof of Theorem 7 yields Theorem 9.

4.4 Derandomization

In this section, we give a deterministic algorithm for finding a sublinear separator of a set of points in \mathbb{R}^d whose spread is bounded by a parameter $\Delta > 0$. The same argument could be applied to point sets of low fractal dimension and unit ball graphs.

THEOREM 10. *Let P be a given set of points in \mathbb{R}^d with spread Δ and G be the greedy $(1 + \epsilon)$ -spanner of P for some $\epsilon \in (0, 1/2]$. Then any k -vertex subgraph H of G has a $\left(1 - \frac{\epsilon^{O(d)}}{d}\right)$ -balanced separator S of size $O(\log \Delta^{1/d} \epsilon^{-O(1)} k^{1-1/d})$. Furthermore, S can be found in $O(k)$ time given H .*

Our algorithm uses a result of [29].

Definition 9. *A set of balls \mathcal{B} in \mathbb{R}^d is ψ -ply if for any point $p \in \mathbb{R}^d$, there are at most ψ balls in \mathcal{B} containing p .*

Given a set ψ -ply set \mathcal{B} of n balls, we say that a ball $\mathbf{B}(p, r)$ is an (α, β) -separating ball if (a) the number of balls entirely contained in $\mathbf{B}(p, r)$ and the number of balls outside and disjoint from $\mathbf{B}(p, r)$ are at most α each and (b) the number of balls intersecting $\mathbf{B}(p, r)$ is at most β . Eppstein, Miller and Teng[29] gave a deterministic linear time algorithm that finds a balanced separating ball for any given ψ -ply set of balls.

THEOREM 11 (THEOREM 4.2 IN [29], ADAPTED). *Given a ψ -ply set \mathcal{B} of n balls in \mathbb{R}^d , one can find a $\left(\frac{d+1}{d+2}n, O(\psi^{1/d}n^{1-1/d})\right)$ -separating ball for \mathcal{B} in (deterministic) linear time.*

Given a set of edges E , we define a set of balls \mathcal{B}_E as follows. For each edge $e \in E$, we construct a ball $\mathbf{B}_e = \mathbf{B}(m_e, w(e)/2)$ with diameter e where the center m_e is the midpoint of e . We then define $\mathcal{B}_E = \{\mathbf{B}_e : e \in E\}$. In Lemma 11 below, we show that, if a graph G is a greedy spanner, then $\mathcal{B}_{E(G)}$ is $(\epsilon^{-O(d)} \log(\Delta))$ -ply. We later use this result and Theorem 11 to find a separator of G deterministically.

Lemma 11. *Let P be a given set of points in \mathbb{R}^d with spread Δ and G be the greedy $(1 + \epsilon)$ -spanner of P for some $\epsilon \in (0, 1/2]$. Then $\mathcal{B}_{E(G)}$ is $(\epsilon^{-O(d)} \log \Delta)$ -ply.*

We now focus on proving Lemma 11. Let p be an arbitrary point in \mathbb{R}^d and \mathcal{B}_p be the set of balls in $\mathcal{B}_{E(G)}$ containing p . Our goal is to bound \mathcal{B}_p by a function of d and ϵ . Observe that:

Observation 4. *Given an edge $e = (u, v)$. If p is in \mathbf{B}_e , then $m_e \in \mathbf{B}(p, |uv|/2)$.*

PROOF. Observe that p is in \mathbf{B}_e if and only if $|pm_e| \leq \frac{|uv|}{2}$, which implies that m_e is in $\mathbf{B}(p, |uv|/2)$. \square

Let d_{min} and d_{max} respectively be the minimum and maximum distance between two points in P . By Observation 4, every center of the balls in \mathcal{B}_p must be inside $\mathbf{B}(p, d_{max}/2)$. Let $M_p = \{e \in E(G) : \mathbf{B}_e \in \mathcal{B}_p\}$. For each $i \in [1, \lfloor \log \Delta \rfloor + 1]$, let $M_p^i = \{e \in M_p, 2^{i-1}d_{min} \leq w(e) < 2^i d_{min}\}$. By definition, M_p is the union of all M_p^i . By Observation 4, we have:

Observation 5. *The centers of the balls in $\mathcal{B}_{M_p^i}$ are in $\mathbf{B}(p, 2^{i-1}d_{min})$.*

Next, we cover $\mathbf{B}(p, 2^{i-1}d_{min})$ by a set of $O(\epsilon^{-d})$ balls of radius $(2^{i-1}\epsilon d_{min})/24$, denoted by \mathcal{N}_i . Note that \mathcal{N}_i exists due to Lemma 1. In the following claim, we show that each ball in \mathcal{N}_i contains at most $\epsilon^{-O(d)}$ centers of the balls in $\mathcal{B}_{M_p^i}$.

Lemma 12. *For any point x and a positive real number D , $\mathbf{B}(x, \epsilon D/24)$ contains at most $\epsilon^{-O(d)}$ midpoints of the edges in $E(G)$ whose lengths are between D and $2D$.*

PROOF. Similar to the proof of Lemma 7, we cover the space \mathbb{R}^d by a set of $O(\epsilon^{-d})$ cones with apex x and angle $\epsilon/8$. Let $M_{x,D} = \{e \in E(G) : D \leq w(e) \leq 2D, m_e \in \mathbf{B}(x, \epsilon D/24)\}$ and $V(M_{x,D})$ is the set of endpoints of the edges in $M_{x,D}$. Our goal is to show that $|M_{x,D}| = \epsilon^{-O(d)}$. First, we show that each cone contains at most one point $V(M_{x,D})$ which is incident to exactly one edge in $M_{x,D}$. Assume that there is a cone containing two points y and q that are endpoints of two different edges in $M_{x,D}$; y and q might be the same point. Let t and s be the points such that $(t, y), (s, q) \in M_{x,D}$. See Figure 2 for an illustration. W.l.o.g, we assume that $|xy| \geq |xq|$. By Lemma 8, we have $|xy| \geq |yq| + (1 - \epsilon/4)|xq|$. To obtain a contradiction, we show below that $(1 + \epsilon)|ty| > \delta_{G-(t,y)}(t, y)$.

Let m_{ty} and m_{sq} be the midpoints of (t, y) and (s, q) , respectively. Because m_{ty} and m_{sq} are both in $\mathbf{B}(x, \epsilon D/24)$, we have $|xm_{ty}|, |xm_{sq}| \leq \epsilon D/24$.

Claim 2. $2\overrightarrow{m_{ty}m_{sq}} = \overrightarrow{ts} + \overrightarrow{yq}$.

PROOF. Since $\overrightarrow{ts} + \overrightarrow{sy} + \overrightarrow{yq} + \overrightarrow{qt} = 0$, it holds that $\overrightarrow{ts} + \overrightarrow{yq} = \overrightarrow{ys} + \overrightarrow{tq}$. We have:

$$2\overrightarrow{m_{ty}m_{sq}} = \overrightarrow{m_{ty}s} + \overrightarrow{m_{ty}q} = \frac{\overrightarrow{ts} + \overrightarrow{ys} + \overrightarrow{tq} + \overrightarrow{yq}}{2} = \overrightarrow{ts} + \overrightarrow{yq} \quad , \quad (14)$$




Fig. 2. A cone containing two points in $V(M_{x,D})$. We later prove that this case cannot happen.

as claimed. □

Since $\vec{ts} = \vec{yq} - 2\vec{m_{ty}m_{sq}}$, it follows from the triangle inequality that:

$$|ts| \leq |yq| + 2|m_{ty}m_{sq}| \leq |yq| + \epsilon D/6. \quad (15)$$

Moreover,

$$\begin{aligned} |ty| &= 2|ym_{ty}| \geq 2|xy| - 2|xm_{ty}| \geq 2|xy| - \epsilon D/12 && \text{(by the triangle inequality)} \\ &\geq 2|yq| + 2(1 - \epsilon/4)|xq| - \epsilon D/12 && \text{(since } |xy| \geq |yq| + (1 - \epsilon/4)|xq| \text{ by Lemma 8)} \\ &\geq |yq| + |ts| - \epsilon D/6 + 2(1 - \epsilon/4)|xq| - \epsilon D/12 && \text{(by Equation (15))} \\ &\geq |yq| + |ts| + 2(1 - \epsilon/4)(|m_{sq}q| - |m_{sq}x|) - \epsilon D/4 && \text{(by the triangle inequality)} \\ &\geq |yq| + |ts| + (1 - \epsilon/4)|sq| - \epsilon D/3 && \text{(since } |m_{sq}x| \leq \epsilon D/24 \text{ and } 2|m_{sq}q| = |sq|) \end{aligned} \quad (16)$$

Multiplying both sides of Equation (16) by $(1 + \epsilon)$, we obtain:

$$\begin{aligned} (1 + \epsilon)|ty| &\geq (1 + \epsilon)(|yq| + |ts| + (1 - \epsilon/4)|sq| - \epsilon D/3) \\ &= (1 + \epsilon)|yq| + (1 + \epsilon)|ts| + |sq| + \underbrace{(3\epsilon/4 - \epsilon^2/4)}_{\geq D} |sq| - (1 + \epsilon)\epsilon D/3 \\ &\geq (1 + \epsilon)|yq| + (1 + \epsilon)|ts| + |sq| + \epsilon D \underbrace{(3/4 - \epsilon/4 - (1 + \epsilon)/3)}_{\text{positive for all } \epsilon \in (0, 1/2]} \\ &> (1 + \epsilon)|yq| + (1 + \epsilon)|ts| + \delta_{G-(t,y)}(s, q) && \text{(since } G - (t, y) \text{ contains } (s, q)) \end{aligned} \quad (17)$$

To bound the value of $(1 + \epsilon)|ty|$, we show in the following claim that the distances from y to q and from t to s are unchanged after removing (t, y) from G .

Claim 3. $\delta_{G-(t,y)}(t, s) = \delta_G(t, s)$ and $\delta_{G-(t,y)}(y, q) = \delta_G(y, q)$.

PROOF. We show that $|yq|, |ts| < |ty|/(1 + \epsilon)$. By Lemma 8, we have:

$$\begin{aligned} |yq| &\leq |xy| - (1 - \epsilon/4)|xq| \leq |xy| \leq |m_{ty}y| + |xm_{ty}| && \text{(by the triangle inequality)} \\ &\leq |ty|/2 + \underbrace{\epsilon D/24}_{\leq \epsilon |ty|/24} \leq (1/2 + \epsilon/24)|ty| \\ &< |ty|/(1 + \epsilon) && \text{since } \epsilon \in (0, 1/2] \end{aligned} \quad (18)$$

By Equation (15), $|ts| \leq |yq| + \epsilon D/6$. Thus, by Equation (18), it holds that:

$$|ts| \leq |ty|/2 + \epsilon D/24 + \epsilon D/6 = |ty|/2 + 5\epsilon D/24 \leq (1/2 + 5\epsilon/24)|ty| < |ty|/(1 + \epsilon) \quad (19)$$

By Equations (18) and (19), $(1+\epsilon) |ts|$, $(1+\epsilon) |yq|$ are both less than $|ty|$. On the other hand, $\delta_G(t, s) \leq (1+\epsilon) |ts| < |ty|$ and $\delta_G(y, q) \leq (1+\epsilon) |yq| < |ty|$. It follows that edge (t, y) is not on the shortest path from t to s and from y to q . This means that $\delta_{G-(t,y)}(t, s) = \delta_G(t, s)$ and $\delta_{G-(t,y)}(y, q) = \delta_G(y, q)$ as desired. \square

Apply $\delta_{G-(t,y)}(y, q) = \delta_G(y, q) \leq (1+\epsilon) |yq|$ and $\delta_{G-(t,y)}(t, s) = \delta_G(t, s) \leq (1+\epsilon) |ts|$ to Equation (17), we obtain:

$$(1+\epsilon) |ty| > \delta_{G-(t,y)}(y, q) + \delta_{G-(t,y)}(t, s) + \delta_{G-(t,y)}(s, q) \geq \delta_{G-(t,y)}(t, y). \quad (20)$$

By Fact 1, we obtain a contradiction. Thus, there is at most one point in $V(M_{x,D})$ in each cone which is incident to exactly one edge in $M_{x,D}$. Since there are $\epsilon^{-O(d)}$ cones, $|V(M_{x,D})| = \epsilon^{-O(d)}$. Therefore, $M_{x,D}$ contains at most $\epsilon^{-O(d)}$ edges as claimed. \square

We now finish the proof of Lemma 11.

PROOF OF LEMMA 11. By Lemma 12, each ball in \mathcal{N}_i contains at most $\epsilon^{O(-d)}$ centers of the balls in $\mathcal{B}_{M_p^i}$. Because $\mathbf{B}(p, 2^{i-1}d_{min})$ is covered by \mathcal{N}_i , $\mathbf{B}(p, 2^{i-1}d_{min})$ contains at most $|\mathcal{N}_i|\epsilon^{-O(d)} = \epsilon^{-O(d)}$ centers of the balls in $\mathcal{B}_{M_p^i}$. By Observation 5, all centers of balls in $\mathcal{B}_{M_p^i}$ are in $\mathbf{B}(p, 2^{i-1}d_{min})$. Therefore, $|M_p^i| = |\mathcal{B}_{M_p^i}| = \epsilon^{-O(d)}$. Since \mathcal{B}_p is the union of $\mathcal{B}_{M_p^i}$ for $i \in [1, \lfloor \log \Delta \rfloor + 1]$, $|\mathcal{B}_p| \leq \sum_i |\mathcal{B}_{M_p^i}| = \epsilon^{-O(d)} \log \Delta$ as desired. \square

We now give a deterministic linear-time algorithm to find a separator of a greedy spanner in Euclidean spaces as claimed in Theorem 10.

PROOF OF THEOREM 10. We prove Theorem 10 for the special case $H = G$. The argument for any subgraph H follows the same line. Let n be the number of vertices in P and E be the edge set of G .

If $n = O(d^{2d} \log \Delta \epsilon^{-O(d)})$, then $\log \Delta^{1/d} \epsilon^{-O(1)} n^{1-1/d} = \Omega(n/d^2)$. Thus, if we let S be any set of vertices of size n/d^2 of G , then $|S| = O(\log \Delta^{1/d} \epsilon^{-O(1)} n^{1-1/d})$. Furthermore, S is a $(1 - \frac{1}{d^2})$ -balanced separator of G , Theorem 10 follows from the fact that $1 - 1/d^2 \leq 1 - \epsilon^d/d$ when $\epsilon \leq 1/2$.

Otherwise, $n = \Omega(d^{2d} \log \Delta \epsilon^{-O(d)})$. By Lemma 11, \mathcal{B}_E is $(\epsilon^{-O(d)} \log \Delta)$ -ply. By Theorem 11, we can find a $(\frac{d+1}{d+2}, O(\log \Delta^{1/d} \epsilon^{-O(1)} n^{1-1/d}))$ -separating ball $\mathbf{B}(p, r)$ of \mathcal{B}_E in deterministic linear time. Let \mathcal{B}_{in} , \mathcal{B}_{ext} and \mathcal{B}_{cut} be the set of balls strictly contained in the interior, exterior and cutting $\mathbf{B}(p, r)$. Similar to Theorem 5, we construct a separator S by picking all vertices in $\mathbf{B}(p, r)$ that are incident to the edges cutting $\mathbf{B}(p, r)$. By the construction of \mathcal{B}_E , we get $|S| \leq |\mathcal{B}_{cut}| = O(\log \Delta^{1/d} \epsilon^{-O(1)} n^{1-1/d})$. To complete the proof, we show that there are at least $\frac{\epsilon^{O(d)} n}{d}$ vertices outside $\mathbf{B}(p, r)$. The bound for the number of inside vertices follows by a symmetric argument. The number of balls in the exterior of $\mathbf{B}(p, r)$ is $|\mathcal{B}_{ext}| = |\mathcal{B}_E| - |\mathcal{B}_{in}| - |\mathcal{B}_{cut}| \geq |\mathcal{B}_E| - \frac{d+1}{d+2} |\mathcal{B}_E| - |\mathcal{B}_{cut}| = \frac{|\mathcal{B}_E|}{d+2} - |\mathcal{B}_{cut}|$. Because $|\mathcal{B}_E| = |E| \geq n-1$ and $n = \Omega(d^{2d} \log \Delta \epsilon^{-O(d)})$, $|\mathcal{B}_{cut}| = O(\log \Delta^{1/d} \epsilon^{-O(1)} n^{1-1/d}) \leq \frac{|\mathcal{B}_E|}{(d+2)(d+3)}$. It follows that $|\mathcal{B}_{ext}| \geq \frac{|\mathcal{B}_E|}{d+2} - \frac{|\mathcal{B}_E|}{(d+2)(d+3)} \geq \frac{|\mathcal{B}_E|}{d+3}$. That is, there are at least $\frac{|E|}{d+3}$ edges in the exterior of $\mathbf{B}(p, r)$. By Lemma 9, each vertex in G has degree at most ϵ^{1-2d} . Therefore, there are at least $\frac{|E|}{(d+3)\epsilon^{1-2d}}$ vertices in the exterior of $\mathbf{B}(p, r)$. Using the fact that $|E| \geq n-1$, we conclude the number of points outside $\mathbf{B}(p, r)$ is at least $\frac{n\epsilon^{O(d)}}{d}$. By the same argument, the number of points inside $\mathbf{B}(p, r)$ is at least $\frac{n\epsilon^{O(d)}}{d}$. Therefore, S is a $(1 - \frac{\epsilon^{O(d)}}{d})$ -balanced separator of G . \square

5 BOUNDED DEGREE SPANNERS WITH SUBLINEAR SEPARATORS IN DOUBLING METRICS

Chan, Gupta, Maggs and Zhou [13], hereafter CGMZ, constructed a $(1+\epsilon)$ -spanner with a maximum degree of $\epsilon^{-O(d)}$ for points in doubling metrics of dimension d . We show in this section that their spanner is τ -lanky for

$\tau = \epsilon^{-O(d)}$. This implies that CGMZ's spanner has sublinear separator, thereby resolves the question asked by Abam and Hal-Peled in [1].

We assume doubling metric (X, δ_X) has dimension d , minimum pairwise distance 1 and spread Δ . CGMZ's construction relies on a net-tree that we define below.

Definition 10 (Net Tree). *Let $r_0 = 1/4$ and $N_0 = X$. For each integer $i \in [1, \lceil \log(\Delta) \rceil + 2]$, we define $r_i = 2^i r_0$ and $N_i \subseteq N_{i-1}$ be the r_i -net of N_{i-1} . The hierarchy of nets $N_0 \supseteq N_1 \supseteq \dots \supseteq N_{\lceil \log(\Delta) \rceil + 2}$ induces a rooted tree T where:*

- (1) T has $\lceil \log(\Delta) \rceil + 3$ levels in which level i corresponds to the set of points N_i for every $i \in [0, \lceil \log(\Delta) \rceil + 2]$.
- (2) The parent of a vertex v at level $i \in [0, \lceil \log(\Delta) \rceil + 1]$ is the closest vertex in N_{i+1} ; ties are broken arbitrarily.

Note that since Δ is the maximum pairwise distance, $N_{\lceil \log(\Delta) \rceil + 2}$ contains a single point. Additionally, the starting radius $r_0 = 1/4$ is an artifact introduced by CGMZ [13] to handle some edge cases gracefully; one could construct the net tree with the starting radius $r_0 = 1$.

Since a point $x \in X$ may appear in many different nets, we sometimes write (x, i) to explicitly indicate the copy of x in the net N_i . We denote by $p(x, i)$ the parent of the point (x, i) , which is a point in N_{i+1} when (x, i) is not the root of the tree T . For each point $x \in X$, we denote by $i^*(x)$ the maximum level that x appears in the net tree T .

CGMZ Spanner Construction. The spanner construction of CGMZ has two steps. In the first step, they construct a $(1 + \epsilon)$ -spanner G_1 with $O(n)$ edges using the *cross edge rule*. Edges of G_1 are then oriented by the maximum level of the endpoints: an edge (u, v) is oriented as $(u \rightarrow v)$ if $i^*(u) < i^*(v)$; if both endpoints have the same maximum level, then the edge is oriented arbitrarily. CGMZ showed that the out-degree of every vertex of G_1 after the orientation is a constant. In the second step, they reduced the in-degree of every vertex by rerouting to its neighbors following a *rerouting rule*. The graph after rerouting edges has bounded degree (for constant ϵ and d). In the following, we formally describe two steps.

CGMZ Algorithm [13]: The construction has two steps.

- **Step 1:** Let $\gamma = 4 + \frac{32}{\epsilon}$. For each net N_i at level i of the net tree T , let E_i be the set of all pairs (u, v) of vertices in N_i such that $\delta_X(u, v) \leq \gamma r_i$ and (u, v) has not appeared in $\bigcup_{j=0}^{i-1} E_j$. We call E_i the set of *cross edges* at level i . The edge set of G_1 is the union of all E_i 's: $E(G_1) = \bigcup_{i=0}^{\lceil \log(\Delta) \rceil + 2} E_i$. Let \vec{G}_1 be obtained from G_1 by orienting every edge (u, v) as $(u \rightarrow v)$ if $i^*(u) < i^*(v)$, as $(v \rightarrow u)$ if $i^*(u) > i^*(v)$, and arbitrarily if $i^*(u) = i^*(v)$. Let \vec{E}_i be the set of oriented edges corresponding to E_i .
- **Step 2:** Let $l = \lceil (1/\epsilon) \rceil + 1$. For each point $w \in N_i$, $i \in [0, \lceil \log(\Delta) \rceil + 2]$, $N_i^{in}(w) = \{v \in N_i : (v \rightarrow w) \in E(\vec{G}_1)\}$ be the set of *in-neighbors at level i of w* . We construct another directed graph \vec{G}_2 with the same vertex set X as follows. Initially, \vec{G}_2 has no edge. Let $I_w = \{i_1, i_2, \dots, i_{m_w}\}$ be the list of levels sorted by increasing order such that w has *at least one* in-neighbor at each level in the list; $m_w \leq i^*(w)$. For each index $j \leq [1, m_w]$, if $j \leq l$, we add every directed edge in \vec{E}_{i_j} to \vec{G}_2 ; otherwise, $j > l$, we choose a point $u \in N_{j-l}$ and for every edge $(v \rightarrow w) \in \vec{E}_{i_j}$, we add an edge $(v \rightarrow u)$ to \vec{G}_2 . We say that $(v \rightarrow w)$ is *rerouted* to u . Let G_2 be obtained from \vec{G}_2 by ignoring the directions of the edges; we return G_2 as the output spanner.

Our goal is to show that G_2 is τ -lankly for $\tau = \epsilon^{-O(d)}$. To this end, we show several properties of G_1 and G_2 . CGMZ [12] showed the following lemmas.

Lemma 13. G_1 and \vec{G}_1 have following properties:

- (1) For every point $u \in N_i$, $i \in [0, \lceil \log(\Delta) \rceil + 2]$, $|N_i^{\text{in}}(u)| \leq (4\gamma)^d$; Claim 5.9 [12].
- (2) The out-degree of a vertex u in \vec{G}_1 is $\epsilon^{-O(d)}$; see Lemma 5.10 [12].

The proof of Item (1) in Lemma 13 is by using the packing bound (Lemma 1). For Item (2) in Lemma 13, CGMZ observed that that for every vertex u , there are $O(\log(1/\epsilon))$ levels where u could have edges oriented out from u ; the bound on the out-degree then follows from Item (1).

Lemma 14 (Lemma 5.12 [12]). G_2 is a $(1 + 4\epsilon)$ -spanner of (X, δ_X) with maximum degree $\epsilon^{-O(d)}$.

CGMZ used the following claim to show the stretch bound $(1 + 4\epsilon)$ of G_2 (Equation (4) [12]). Since we also use this claim here, we include its proof for completeness.

Claim 4. If $(v \rightarrow w)$ is rerouted to $(v \rightarrow u)$, then $\delta_X(u, w) \leq \epsilon \delta_X(v, w)$ and $\delta_X(v, w) \geq \delta_X(v, u)/(1 + \epsilon)$.

PROOF. First, we show that $\delta_X(u, w) \leq \epsilon \delta_X(v, w)$. Let i and j be two levels such that $(v, w) \in E_i$ and $(u, w) \in E_j$. It follows that $\delta_X(v, w) \geq \gamma r_{i-1}$ since (v, w) does not appear in E_{i-1} , and that $\delta_X(u, w) \leq \gamma r_j$ by the definition of E_j . By the construction in Step 2, $i \geq j + l$. We have that $\delta_X(v, w) \geq \gamma r_{i-1} \geq 2^l \gamma r_j \geq \frac{\delta_X(u, w)}{\epsilon}$. Thus, $\delta_X(u, w) \leq \epsilon \delta_X(v, w)$ as claimed.

By the triangle inequality, $\delta_X(v, u) \leq \delta_X(v, w) + \delta_X(w, u) < (1 + \epsilon)\delta_X(v, w)$; the second claim follows. \square

Next, we show that any ball of radius r only contains exactly one point that is incident to a long edge in G_1 . Recall that $\gamma = 4 + \frac{32}{\epsilon}$ is the parameter defined in CGMZ Algorithm.

Lemma 15. Let $\mathbf{B}(p, r)$ be any ball of radius r centered at some point p . There is at most one point in $\mathbf{B}(p, r)$ that is incident to an edge in G_1 with length at least $4\gamma r$.

More generally, for any parameter $\beta > 0$, there are at most $2^{O(d)} \left(\frac{\gamma}{\beta}\right)^d$ points in $\mathbf{B}(p, r)$ that are incident to edges in G_1 of length at least βr .

PROOF. By the construction of G_1 , two endpoints of an edge with length at least $4\gamma r$ must be in r_i -net N_i such that $r_i \geq 4r$. Since any two points in N_i has distance at least $r_i \geq 4r$, $|\mathbf{B}(p, r) \cap N_i| \leq 1$; this implies the first claim.

We now show the second claim. Let \mathcal{B} be a set of balls obtained by taking balls of radius $\frac{\beta r}{4\gamma}$ centered at points in a $\left(\frac{\beta r}{4\gamma}\right)$ -net of $\mathbf{B}(p, r)$. By Lemma 1, $|\mathcal{B}| = 2^{O(d)} \left(\frac{\gamma}{\beta}\right)^d$. Let E_{long} be the set of edges in G_1 of length at least βr . For each ball $B \in \mathcal{B}$, there is at most one point in $B \cap \mathbf{B}(p, r)$ that is incident to an edge in E_{long} by the first claim. It follows that the total number of points in $\mathbf{B}(p, r)$ incident to an edge in E_{long} is at most $|\mathcal{B}| = 2^{O(d)} \left(\frac{\gamma}{\beta}\right)^d$. \square

Next, we show that there is a small number of edges in G_2 between any well-separated pair. For any two sets A and B , we define $r_{AB} = \max\{\text{diam}(A), \text{diam}(B)\}$.

Lemma 16. Let (A, B) be a γ -separated pair in (X, δ_X) . Then, there are at most $\epsilon^{-O(d)}$ edges in G_2 between A and B .

PROOF. We count the number of directed edges with length at least γr_{AB} between A and B in \vec{G}_2 . (Recall that $\gamma = 4 + \frac{32}{\epsilon}$.) We consider the edges that are directed from A to B ; those directed from B to A can be counted by the same argument. By Claim 4, an edge, say $(v \rightarrow u)$, of length at least γr_{AB} from a point v in A must be rerouted from an edge $(v \rightarrow w) \in E(\vec{G}_1)$ of length at least $\frac{\gamma r_{AB}}{1+\epsilon} \geq \frac{\gamma \text{diam}(A)}{1+\epsilon}$. It follows that the set of the endpoints of the directed edges of length at least γr_{AB} from A to B is a subset of the set of the endpoints in A , denoted by U_A , of edges of length at least $\frac{\gamma r_{AB}}{1+\epsilon}$ in G_1 .

By applying Lemma 15 to a minimum ball enclosing A (of radius at most $\text{diam}(A)$), we have that $|U_A| = 2^{O(d)} \left(\frac{\gamma}{\gamma/(1+\epsilon)} \right)^d = 2^{O(d)} (1+\epsilon)^d = 2^{O(d)}$. By Lemma 14, there are at most $|U_A| \epsilon^{-O(d)} = 2^{O(d)} \epsilon^{-O(d)} = \epsilon^{-O(d)}$ edges in G_2 incident to points in U_A ; here we use the fact that $\epsilon \leq 1/2$. It follows that the number of directed edges from A to B in \vec{G}_2 of length at least γr_{AB} is $\epsilon^{-O(d)}$. By the same argument, there are $\epsilon^{-O(d)}$ edges from B to A in \vec{G}_2 of length at least γr_{AB} . The lemma now follows. \square

We obtain the following generalization of Lemma 16.

Corollary 3. *Let (A, B) be a β -separated pair in (X, δ_X) . Then there are at most $\left(\frac{\gamma}{\beta}\right)^{2d} \epsilon^{-O(d)}$ edges in G_2 between A and B .*

PROOF. We have that $\delta_X(A, B) \geq \beta r_{AB}$. Let \mathcal{A} and \mathcal{B} be minimum collections of balls with radius $\frac{\beta r_{AB}}{2\gamma}$ that cover A and B , respectively. By Lemma 1, we have that $|\mathcal{A}| = 2^{O(d)} \left(\frac{2\text{diam}(A)}{\beta r_{AB}/(2\gamma)} \right)^d \leq 2^{O(d)} \left(\frac{2r_{AB}}{\beta r_{AB}/(2\gamma)} \right)^d = 2^{O(d)} \left(\frac{\gamma}{\beta} \right)^d$. By the same argument, $|\mathcal{B}| = 2^{O(d)} \left(\frac{\gamma}{\beta} \right)^d$. Let $(A', B') \in \mathcal{A} \times \mathcal{B}$ be any pair of balls. Observe that $(A' \cap A, B' \cap B)$ is a γ -separated pair. By Lemma 16, there are $\epsilon^{-O(d)}$ edges between $A' \cap A$ and $B' \cap B$. It follows that the number of edges between A and B in G_2 is $|\mathcal{A}||\mathcal{B}| \epsilon^{-O(d)} = 2^{O(d)} \left(\frac{\gamma}{\beta} \right)^{2d} \epsilon^{-O(d)} = \left(\frac{\gamma}{\beta} \right)^{2d} \epsilon^{-O(d)}$ since $\epsilon \leq 1/2$. \square

Next, we show that the number of short edges, those of length at least r and at most γr , cut by a ball with radius r is small. This reduce the task of showing lankiness of G_2 to bounding the number of long edges of G .

Lemma 17. *Let $\mathbf{B}(p, r)$ be any ball of radius r centered at some point p . The number of edges with length from r to γr cut by $\mathbf{B}(p, r)$ is $\epsilon^{-O(d)}$.*

PROOF. Let $E_{\text{short}} \subseteq E(G_2)$ be the set of edges of length in $[r, \gamma r]$ cut by $\mathbf{B}(p, r)$. Observe by the triangle inequality that for every edge $(u, v) \in E_{\text{short}}$, both u and v are in $\mathbf{B}(p, (\gamma+1)r)$. Let \mathcal{B}_{in} (\mathcal{B}_{out} , resp.) be the minimum collection of balls with radius ϵr that covers $\mathbf{B}(p, r)$ ($\mathbf{B}(p, (\gamma+1)r)$, resp.). By the packing lemma (Lemma 1), $|\mathcal{B}_{\text{in}}^i| = 2^{O(d)} \left(\frac{r}{\epsilon r} \right)^d = \epsilon^{-O(d)}$, and $|\mathcal{B}_{\text{out}}^i| = 2^{O(d)} \left(\frac{\gamma+1}{\epsilon} \right)^d = \epsilon^{-O(d)}$.

Let $(B_{\text{in}}, B_{\text{out}}) \in \mathcal{B}_{\text{in}} \times \mathcal{B}_{\text{out}}$ be any pair in of ball such that there is an edge $e = (u, v)$ between B_{in} and B_{out} . By the triangle inequality, we have that $\delta_X(B_{\text{in}}, B_{\text{out}}) \geq \delta_X(u, v) - \text{diam}(B_{\text{in}}) - \text{diam}(B_{\text{out}}) \geq r - 4\epsilon r$. Thus, $\frac{\delta_X(B_{\text{in}}, B_{\text{out}})}{r_{B_{\text{in}}, B_{\text{out}}}} \geq \frac{r-4\epsilon r}{2\epsilon r} = \frac{1-4\epsilon}{2\epsilon}$. (Recall that $r_{B_{\text{in}}, B_{\text{out}}}$ by definition is the maximum diameter of B_{in} and B_{out} .) This implies that $(B_{\text{in}}, B_{\text{out}})$ is a $\left(\frac{1-4\epsilon}{2\epsilon}\right)$ -separated pair. By applying Corollary 3 with $\beta = \frac{1-4\epsilon}{2\epsilon}$, it follows that there are $\left(\frac{2\gamma\epsilon}{1-4\epsilon}\right)^{2d} \epsilon^{-O(d)} = 2^{O(d)} \epsilon^{-O(d)} = \epsilon^{-O(d)}$ edges in E_{short} between B_{in} and B_{out} . Thus, $|E_{\text{short}}| = \epsilon^{-O(d)} |\mathcal{B}_{\text{in}} \times \mathcal{B}_{\text{out}}| = \epsilon^{-O(d)}$, as desired. \square

By Lemma 17, to show that G_2 is lanky, it suffices to focus on edges of length at least γr . To that end, we show the following structure lemma, which bounds the number of directed edges in \vec{G}_1 of length at least γr that are incident the same vertex w in a ball of radius $c \cdot r$ for some parameter $c \geq 1$ and rerouted (from w) to other points in the same ball in the construction of \vec{G}_2 . See Figure 3 for an illustration.

Lemma 18. *Let $\mathbf{B}(p, c \cdot r)$ be any ball centered at some point p that has radius $c \cdot r > 0$ for some parameter $c \geq 1$. For every point $w \in \mathbf{B}(p, c \cdot r)$, there are $\epsilon^{-O(d)} (\lceil \log c \rceil + 1)$ edges $(v \rightarrow w)$ in \vec{G}_1 such that $(v \rightarrow w)$ is rerouted to a point u in $\mathbf{B}(p, c \cdot r)$ and $\delta_X(v, u) \geq \gamma r$.*




Fig. 3. An edge $(v \rightarrow w) \in E(\vec{G}_1)$ incident to a vertex $w \in \mathbf{B}(p, c \cdot r)$ is rerouted to some point $u \in \mathbf{B}(p, c \cdot r)$. That is, $(v \rightarrow u) \in E(\vec{G}_2)$.

PROOF. Let $I_w = \{i_1, i_2, \dots, i_{m_w}\}$ be the increasing sequence of indices such that $N_{i_t}^{in}(w) \neq \emptyset$ for every $1 \leq t \leq m_w$. Let $P_{out}(w)$ be the set of vertices v such that $(v \rightarrow w) \in E(\vec{G}_1)$, $(v \rightarrow w)$ is rerouted to a point, say u , in $\mathbf{B}(p, c \cdot r)$ and $\delta_X(v, u) \geq \gamma r$. Observe that $P_{out}(w)$ is a subset of $\bigcup_{t=1}^{m_w} N_{i_t}^{in}(w)$. Let i_j be the smallest index in I_w such that $N_{i_j}^{in}(w) \cap P_{out}(w) \neq \emptyset$. Our goal is to show that $P_{out}(w) \subseteq \bigcup_{t=0}^{l+\lceil \log c \rceil+2} N_{i_{j+t}}^{in}(w)$. Recall that $l = \lceil (1/\epsilon) \rceil + 1$, which is defined in Step 2 of CGMZ Algorithm.

Suppose that there exists an integer $h \geq l + \lceil \log c \rceil + 3$ such that $N_{i_{j+h}}^{in}(w) \cap P_{out}(w) \neq \emptyset$. Let $v_1 \neq v_2$ be any two points in $P_{out}(w)$ such that $v_1 \in N_{i_j}^{in}(w)$ and $v_2 \in N_{i_{j+h}}^{in}(w)$. Let u_a be the point such that $(v_a \rightarrow w)$ is rerouted to, $a \in \{1, 2\}$. (See Figure 4 for an illustration.) By the definition of $P_{out}(w)$, we have that $u_1, u_2 \in \mathbf{B}(p, c \cdot r)$ and that $\delta_X(v_1, u_1), \delta_X(v_2, u_2) \geq \gamma r$. By Claim 4, $\delta_X(v_1, w) \geq \frac{\delta_X(v_1, u_1)}{1+\epsilon} \geq \frac{\gamma r}{1+\epsilon}$. Hence, $r_{i_j} \geq \frac{r}{1+\epsilon}$ by the construction in Step 1 of CGMZ Algorithm. Since $v_2 \in N_{i_{j+h}}^{in}(w)$, by the construction in Step 2 of CGMZ Algorithm, $u_2 \in N_{i_{j+h-1}}^{in}(w)$. It follows that:

$$\delta_X(w, u_2) \geq r_{i_{j+h-1}} \geq r_{i_{j+h-l}} \geq 2^{\lceil \log c \rceil+3} r_{i_j} \geq \frac{8c \cdot r}{1+\epsilon} > 2c \cdot r, \quad (21)$$

since $\epsilon \leq 1$ and $c \cdot r > 0$. However, by the triangle inequality, $\delta_X(w, u_2) \leq \delta_X(w, p) + \delta_X(p, u_2) \leq 2cr$, which contradicts Equation (21). Thus, $P_{out}(w)$ only contains points in $\bigcup_{t=0}^{l+\lceil \log c \rceil+2} N_{i_{j+t}}^{in}(w)$.

By Item (1) of Lemma 13, there are at most $(4\gamma)^d$ points in each net that are adjacent to w in G_1 . Hence, $|P_{out}(w)| \leq (l + \lceil \log c \rceil + 3)O(4\gamma)^d = \epsilon^{-O(d)}(\lceil \log c \rceil + 1)$; this implies the lemma. \square

We now have all necessary tools to show that G_2 is lanky.

Lemma 19. G_2 is $\epsilon^{-O(d)}$ -lanky.

PROOF. Let $\mathbf{B}(p, r)$ be a ball of radius r . We assume that $r \geq 1/2$ since otherwise, we could enlarge r to $1/2$ without changing the number of edges that cut $\mathbf{B}(p, r)$. We show that the number of edges of length at least r cut by $\mathbf{B}(p, r)$ is bounded by $\epsilon^{-O(d)}$, which implies the lemma by the definition (Definition 1).

Let E_{short} be the set of edges of length in $[r, \gamma r]$ cut by $\mathbf{B}(p, r)$. Observe by Lemma 17 that $|E_{short}| = \epsilon^{-O(d)}$. Thus, it remains to focus on the set of edges in G_2 of length at least γr cut by $\mathbf{B}(p, r)$, which we denote by E_{long} . Let \vec{E}_{long} be the set of directed edges corresponding to E_{long} in \vec{G}_2 .




Fig. 4. (a): The positions of u_1, u_2, v_1, v_2 and w in the net tree and (b) the edges $(v_1 \rightarrow w)$ and $(v_2 \rightarrow w)$ are rerouted to u_1 and u_2 , respectively. We show in Lemma 18 that u_2 must be outside $\mathbf{B}(p, c \cdot r)$.

Recall that in the construction of Step 2 of CGMZ Algorithm, some edges in \vec{G}_1 are copied over \vec{G}_2 (the case where $j \geq l$) while other edges are rerouted to different endpoints (the case where $j > l$). Let \vec{E}_{long}^1 be the set of edges in \vec{G}_2 that are copied from \vec{G}_1 , and $\vec{E}_{\text{long}}^2 = \vec{E}_{\text{long}} \setminus \vec{E}_{\text{long}}^1$. Let E_{long}^1 and E_{long}^2 be the undirected counterparts of \vec{E}_{long}^1 and \vec{E}_{long}^2 , respectively.

We observe by the definition that if a point is incident to an edge in E_{long}^1 , it is also incident to an edge with length at least γr in G_1 . By applying Lemma 15 with $\beta = 1$, there are $2^{O(d)}$ points in $\mathbf{B}(p, r)$ incident to an edge with length at least γr in $E(G_1)$. By Lemma 14, each of these points is incident to at most $\epsilon^{-O(d)}$ edges in E_{long} since $E_{\text{long}} \subseteq E(\vec{G}_2)$. Thus, $|E_{\text{long}}^1| = 2^{O(d)} \epsilon^{-O(d)} = \epsilon^{-O(d)}$ since $\epsilon \leq 1/2$.

To bound the size of E_{long}^2 , we partition \vec{E}_{long}^2 into two sets: $\vec{E}_{\text{in}}^2 = \{(v \rightarrow u) \in \vec{E}_{\text{long}}^2 \mid u \in \mathbf{B}(p, r)\}$ is the set of edges directed into $\mathbf{B}(p, r)$ and $\vec{E}_{\text{out}}^2 = \{(u \rightarrow y) \in \vec{E}_{\text{long}}^2 \mid u \in \mathbf{B}(p, r)\}$ is the set of edges directed out of $\mathbf{B}(p, r)$. We bound the number of edges of \vec{E}_{in}^2 and \vec{E}_{out}^2 in Claim 5 and Claim 6, respectively.

Claim 5. $|\vec{E}_{\text{in}}^2| = \epsilon^{-O(d)}$.

PROOF. Let j be the index that $4r \geq r_j > 2r$. Recall that for each $i \in [0, \lceil \log(\Delta) \rceil + 2]$, E_i is the set of cross edge at level i defined in Step 1. Let \vec{M}_1 be the set of edges $(v \rightarrow u) \in \vec{E}_{\text{in}}^2$ such that (a) $(v \rightarrow u)$ is formed by rerouting $(v \rightarrow w)$ to u in Step 2 and (b) $(u, w) \in E_{j'}$ for some $j' \geq j$. Let $\vec{M}_2 = \vec{E}_{\text{in}}^2 \setminus \vec{M}_1$.

We first bound $|\vec{M}_1|$ by showing that every edge in \vec{M}_1 is directed towards the same endpoint, say u . If so, $|\vec{M}_1| = \epsilon^{-O(d)}$ since u has a degree at most $\epsilon^{-O(d)}$ by Lemma 14. Suppose for contradiction that there is another

edge $(v' \rightarrow u') \in \vec{M}_1$ such that $u' \neq u$. We observe by the definition of cross edges in Step 1 that u and u' belongs to nets of level at least j . That means they both belong to the net N_j at level j since nets are nested. Thus, $\delta_X(u, u') \geq r_j > 2r$ by the definition of j . It follows that u and u' could not both be in $\mathbf{B}(p, r)$; a contradiction.

It remains to bound $|\vec{M}_2|$. To this end, we define $W = \{w \mid \exists (v \rightarrow u) \in \vec{M}_2 \text{ s.t. } (v \rightarrow w) \text{ is rerouted to } u\}$. We claim that $W \subseteq \mathbf{B}(p, (4\gamma+1)r)$. Let w be an arbitrary point in W . Let u, v be the points that $(v \rightarrow w)$ is rerouted to u . Let j' be the index such that $(u, w) \in E_{j'}$. By the definition of \vec{M}_2 , $j' < j$. Since $(u, w) \in E_{j'}$, $\delta_X(u, w) \leq \gamma r_j \leq 4\gamma r$. By the triangle inequality, $\delta_X(p, w) \leq \delta_X(p, u) + \delta_X(u, w) \leq (4\gamma+1)r$. Thus, $w \in \mathbf{B}(p, (4\gamma+1)r)$, as claimed.

By Claim 4, $\delta_X(v, w) \geq \frac{\delta_X(v, u)}{1+\epsilon} \geq \frac{\gamma r}{1+\epsilon}$. It follows that $w \in N_h$ where h is the minimum index such that $r_h \geq \frac{r}{1+\epsilon}$. Since nets are nested, $W \subseteq N_h \cap \mathbf{B}(p, (4\gamma+1)r)$. By the packing lemma (Lemma 1), $|N_h \cap \mathbf{B}(p, (4\gamma+1)r)| = 2^{O(d)} \left(\frac{(4\gamma+1)r}{r/(1+\epsilon)}\right)^d = \epsilon^{-O(d)}$. By Lemma 18, for each point $w \in W$, there are at most $\epsilon^{-O(d)} (\lceil \log(4\gamma+1) \rceil + 1) = \epsilon^{-O(d)}$ edges in \vec{M}_2 . It follows that $|\vec{M}_2| = |W| \cdot \epsilon^{-O(d)} = \epsilon^{-O(d)}$. The lemma now follows from the fact that both \vec{M}_1 and \vec{M}_2 have $\epsilon^{-O(d)}$ edges. \square

Claim 6. $|\vec{E}_{out}^2| = \epsilon^{-O(d)}$.

PROOF. Let U be the set of endpoints in $\mathbf{B}(p, r)$ of edges in \vec{E}_{out}^2 . By Lemma 14, $|\vec{E}_{out}^2| = O(|U|\epsilon^{-O(d)})$. Let u be an arbitrary point in U . Let x, y be two points that there exists an edge $(u \rightarrow x) \in E(\vec{G}_1)$ that is rerouted to y . That is, $(u \rightarrow y) \in \vec{E}_{out}^2$. Since $\vec{E}_{out}^2 \subseteq \vec{E}_{long}^2$, $\delta_X(u, y) \geq \gamma r$. By Claim 4, we have $\delta_X(u, x) \geq \frac{\delta_X(u, y)}{1+\epsilon} \geq \frac{\gamma r}{1+\epsilon}$. Hence, $u \in N_i$ where i is the minimum index such that $r_i \geq \frac{r}{1+\epsilon}$. It follow that $u \in N_i \cap \mathbf{B}(p, r)$ and hence, $U \subseteq N_i \cap \mathbf{B}(p, r)$. By the packing lemma (Lemma 15), $|N_i \cap \mathbf{B}(p, r)| = 2^{O(d)} \left(\frac{r}{r/(1+\epsilon)}\right)^d = 2^{O(d)}$. We conclude that the total number of edges in \vec{M}_2 is bounded by $|U|\epsilon^{-O(d)} \leq |N_i \cap \mathbf{B}(p, r)|\epsilon^{-O(d)} = \epsilon^{-O(d)}$ as claimed. \square

Since $\vec{E}_{long}^2 = \vec{E}_{in}^2 \cup \vec{E}_{out}^2$, by Claim 5 and Claim 6, $|\vec{E}_{long}^2| = \epsilon^{-O(d)}$. As we showed above, $|E_{long}^1| = \epsilon^{-O(d)}$. Thus, $|E_{long}| = |E_{long}^1| + |\vec{E}_{long}^2| = \epsilon^{-O(d)}$. Since $|E_{short}| = \epsilon^{-O(d)}$, the number of edges of length at least r cut by $\mathbf{B}(p, r)$ is $|E_{short}| + |E_{long}| = \epsilon^{-O(d)}$, as desired. \square

6 SEPARATORS OF GREEDY SPANNERS FOR DOUBLING METRICS

In this section, we prove Theorem 4. We begin with the following lemma which bounds the number of edges in the greedy spanner between two well separated pair.

Lemma 20. *Let (A, B) be a $(4/\epsilon)$ -separated pair in (X, δ_X) . There is at most one edge in G between A and B .*

PROOF. Suppose that there are two edges (u, v) and (u', v') such that u and u' are in A and v and v' are in B . W.l.o.g, we assume that $\delta_X(u, v) \geq \delta_X(u', v')$. By the definition of separated pairs, the diameters of A and B are at most $(\epsilon/4)\delta_X(A, B)$. Thus, $\max\{\delta_X(u, u'), \delta_X(v, v')\} \leq (\epsilon/4)\delta_X(A, B)$. By Fact 1, $\delta_{G-(u,v)}(u, u') \leq (1+\epsilon)\delta_X(u, u') \leq (1+\epsilon)(\epsilon/4)\delta_X(A, B)$. Similarly, $\delta_{G-(u,v)}(v, v') \leq (1+\epsilon)\delta_X(v, v') \leq (1+\epsilon)(\epsilon/4)\delta_X(A, B)$. We have:

$$\begin{aligned} \delta_{G-(u,v)}(u, v) &\leq \underbrace{\delta_{G-(u,v)}(u, u')}_{\leq (1+\epsilon)\epsilon\delta_X(A,B)/4} + \underbrace{\delta_{G-(u,v)}(u', v')}_{\leq (1+\epsilon)\epsilon\delta_X(A,B)/4} + \underbrace{\delta_{G-(u,v)}(v', v)}_{\leq (1+\epsilon)\epsilon\delta_X(A,B)/4} \\ &\leq 2(1+\epsilon)\epsilon\delta_X(A, B)/4 + \delta_X(u', v') \leq \epsilon\delta_X(A, B) + \delta_X(u', v') \quad (\text{since } \epsilon < 1) \\ &< (1+\epsilon)\delta_X(u, v) \quad (\text{since } \delta_X(A, B) \leq \delta_X(u', v') \leq \delta_X(u, v)). \end{aligned} \tag{22}$$

Thus, $(1+\epsilon)\delta_X(u, v) > \delta_{G-(u,v)}(u, v)$, contradicting that G is a $(1+\epsilon)$ -spanner (Fact 1). \square

We obtain the following corollary of Lemma 20.

Corollary 4. *Let (A, B) be a β -separated pair in (X, δ_X) . Then, there are at most $O\left(\frac{1}{\epsilon\beta}\right)^{2d}$ edges in G between A and B .*

PROOF. Let $r_{AB} = \max\{\text{diam}(A), \text{diam}(B)\}$. By the definition of β -separated pairs, $r_{AB} \leq \frac{\delta_X(A, B)}{\beta}$. Let N_A (N_B) be a $\left(\frac{\epsilon\beta r_{AB}}{8}\right)$ -net of A (B). By Lemma 1, it holds that:

$$|N_A| \leq 2^{O(d)} \left(\frac{\text{diam}(A)}{\epsilon\beta r_{AB}/8}\right)^d \leq 2^{O(d)} \left(\frac{r_{AB}}{\epsilon\beta r_{AB}/8}\right)^d = O\left(\frac{1}{\epsilon\beta}\right)^d. \quad (23)$$

By the same argument $|N_B| = O\left(\frac{1}{\epsilon\beta}\right)^d$. Let \mathcal{A} (\mathcal{B}) be a collection of balls with radius $\frac{\epsilon\beta r_{AB}}{8}$ centered at points in A (B). Note that balls in \mathcal{A} and \mathcal{B} cover A and B , respectively. For any two pairs of balls $(A', B') \in \mathcal{A} \times \mathcal{B}$, we observe that $(A' \cap A, B' \cap B)$ is a $\left(\frac{1}{\epsilon}\right)$ -separated pair. Thus, by Lemma 20, there is at most one edge in G between $A' \cap A$ and $B' \cap B$. Therefore, the number of edges in G between A and B is at most:

$$|\mathcal{A}| \cdot |\mathcal{B}| = |N_A| \cdot |N_B| = O\left(\frac{1}{\epsilon\beta}\right)^{2d},$$

by Equation (23); the corollary follows. \square

Next, we show that G is lanky.

Lemma 21. *Let P be a set of n points in a metric (X, δ) of dimension d and spread Δ . Let G be a greedy $(1+\epsilon)$ -spanner for P with $\epsilon \in (0, 1/2]$. Then G is $O(\epsilon^{-d} \log(\Delta))$ -lanky.*

PROOF. Let $\mathbf{B}(p, r)$ be a ball of radius r centered at a point $p \in P$. We assume that the minimum distance is at least 1 and that $r \geq 1/2$; otherwise, we could increase r to $1/2$ without changing the number of edges that cut $\mathbf{B}(p, r)$.

Let E_{cut} be the set of edges with length at least 1 and cutting $\mathbf{B}(p, r)$. We partition the set E_{cut} into $O(\log \Delta)$ subsets: $E_{\text{cut}}^i = \{(u, v) \in E_{\text{cut}} \mid 2^i r < \delta_X(u, v) \leq 2^{i+1} r\}$ for $1 \leq i \leq \lceil \log \Delta \rceil$. Let (u, v) be an edge in E_{cut}^i . Since (u, v) is cut by $\mathbf{B}(p, r)$, at least one of the two endpoints, say u , of (u, v) is in $\mathbf{B}(p, r)$. By the triangle inequality, it follows that $\delta_X(p, v) \leq \delta_X(p, u) + \delta_X(u, v) \leq (2^{i+1} + 1)r$.

Let $t = \max\{\epsilon 2^i r, r\}/4$. Let \mathcal{B}_R for $R > 0$ be a minimal collection of balls of radius t that covered the ball $\mathbf{B}(p, R)$ of radius R centered at p ; \mathcal{B}_R can be constructed by taking balls of radius R centered at a R -net of $\mathbf{B}(p, R)$. By Lemma 1, when $R = r$, we have:

$$|\mathcal{B}_r| = 2^{O(d)} \left(\frac{r}{r/4}\right)^d = 2^{O(d)}$$

and when $R = (2^{i+1} + 1)r$, we have:

$$|\mathcal{B}_{(2^{i+1}+1)r}| = 2^{O(d)} \left(\frac{(2^{i+1} + 1)r}{\epsilon 2^i r/4}\right)^d = O(\epsilon)^{-d}$$

since $\epsilon \leq 1/2$.

Let A and B be two balls in \mathcal{B}_r and $\mathcal{B}_{(2^{i+1}+1)r}$, respectively, that contain two endpoints of an edge $(u, v) \in E_{\text{cut}}^i$. Observe by the triangle inequality that:

$$\delta_X(A, B) \geq w(u, v) - \text{diam}(A) - \text{diam}(B) \geq w(u, v) - 4(\epsilon 2^i r/4) = 2^i(1 - \epsilon)r \geq 2^{i-1}r$$

since $\epsilon \leq 1/2$. By Corollary 4 with $\beta = \frac{2^{i-1}r}{\epsilon 2^i r/4} = \frac{2}{\epsilon}$, it follows that there are $O(1)^{-d}$ edges between A and B . Thus, we have:

$$|E_{\text{cut}}^i| = O(1)^{-d} |\mathcal{B}_r| \cdot |\mathcal{B}_{(2^{i+1})r}| = O(\epsilon)^{-d}.$$

It follows that $|E_{\text{cut}}| = O(\epsilon)^{-d} \log(\Delta)$, which implies the lemma. \square

We are now ready to show Theorem 4.

THEOREM 4. *Let (X, δ_X) be a doubling metric of n points with a constant doubling dimension d . Let G be a $(1 + \epsilon)$ -greedy spanner of (X, δ_X) . G has a balanced separator of size $O(n^{1-1/d} + \log(\Delta))$ when $d \geq 2$ and of size $O(\log(n) + \log(\Delta))$ when $d = 1$.*

PROOF. By Lemma 21, G is τ -lanky with $\tau = O(\epsilon^{-d} \log(\Delta))$, and hence G is weakly τ -lanky with the same parameter. By applying Corollary 4 with $\beta = 1$, G is κ -thin with $\kappa = \epsilon^{-2d}$. Since a doubling metric of dimension d is (η, d) -packable with $\eta = 2^{O(d)}$ by Lemma 1, by Corollary 2, H has a separator S of size $O(2^{O(d)} \epsilon^{-d} (\epsilon^{1-d} k^{1-1/d} + \log(\Delta))) = O(k^{1-1/d})$ for fixed constants d and $\epsilon \in (0, 1/2]$. Also by Corollary 1, S can be found in $O(2^{O(d)} \epsilon^{1-2d} k) = O(k)$ expected time. \square

ACKNOWLEDGMENTS

The first author thanks Shay Solomon for helpful discussions. This material is based upon work supported by the National Science Foundation under Grant No. CCF-2121952. We thank an anonymous reviewer of SODA22 for suggesting derandomizing Theorem 7 and pointing us to [29], which leads to results in Section 4.4.

REFERENCES

- [1] M. A. Abam and S. Har-Peled. 2010. New Constructions of SSPDs and Their Applications. In *Proceedings of the 26th Annual Symposium on Computational Geometry (SoCG'10)*. 192–200. <https://doi.org/10.1145/1810959.1810993>
- [2] I. Althöfer, G. Das, D. Dobkin, D. Joseph, and J. Soares. 1993. On Sparse Spanners of Weighted Graphs. *Discrete Computational Geometry* 9, 1 (1993), 81–100. <https://doi.org/10.1007/BF02189308>
- [3] S. Arya, G. Das, D. M. Mount, J. S. Salowe, and M. Smid. 1995. Euclidean Spanners: Short, Thin, and Lanky. In *Proceedings of the Twenty-seventh Annual ACM Symposium on Theory of Computing (STOC '95)*. 489–498.
- [4] B. Awerbuch. 1985. Communication-Time Trade-Offs in Network Synchronization. In *Proc. of 4th PODC*. 272–276.
- [5] B. Awerbuch, A. Baratz, and D. Peleg. October, 1992. Efficient Broadcast and Light-Weight Spanners. *Technical Report CS92-22, Weizmann Institute* (October, 1992).
- [6] B. S. Baker. 1994. Approximation Algorithms for NP-complete Problems on Planar Graphs. *J. ACM* 41, 1 (1994), 153–180. <https://doi.org/10.1145/174644.174650> preliminary version published in FOCS 1983.
- [7] S. Basagni. 2001. Finding a Maximal Weighted Independent Set in Wireless Networks. *Telecommunication Systems* 18, 1/3 (2001), 155–168. <https://doi.org/10.1023/a:1016747704458>
- [8] Y. Ben-Shimol, A. Dvir, and M. Segal. 2004. SPLAST: a novel approach for multicasting in mobile wireless ad hoc networks. In *Proceedings of the IEEE 15th International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC'04)*. 1011–1015. <https://doi.org/10.1109/PIMRC.2004.1373851>
- [9] G. Borradaile, H. Le, and C. Wulff-Nilsen. 2017. Minor-free graphs have light spanners. In *2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS '17)*. 767–778. <https://doi.org/10.1109/FOCS.2017.76>
- [10] G. Borradaile, H. Le, and C. Wulff-Nilsen. 2019. Greedy spanners are optimal in doubling metrics. In *Proceedings of the 30th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA '19)*. 2371–2379. <https://doi.org/10.1137/1.9781611975482.145>
- [11] S. Cabello and D. Gajser. 2015. Simple PTAS's for Families of Graphs Excluding a Minor. *Discrete Applied Mathematics* 189, C (2015), 41–48. <https://doi.org/10.1016/j.dam.2015.03.004>
- [12] T.-H. H. Chan, A. Gupta, B. M. Maggs, and S. Zhou. 2016. On Hierarchical Routing in Doubling Metrics. *ACM Trans. Algorithms* 12, 4 (2016), 55:1–55:22. Preliminary version appeared in SODA 2005..
- [13] T-H. Hubert Chan, Anupam Gupta, Bruce M. Maggs, and Shuheng Zhou. 2005. On Hierarchical routing in Doubling Metrics. In *Proc. 16th Ann. ACM-SIAM Symposium on Discrete Algorithms (SODA)*.
- [14] B. Chandra, G. Das, G. Narasimhan, and J. Soares. 1992. New Sparseness Results on Graph Spanners. In *Proceedings of the Eighth Annual Symposium on Computational Geometry*.

- [15] P. Chew. 1986. There is a Planar Graph Almost as Good as the Complete Graph. In *Proceedings of the 2nd Annual Symposium on Computational Geometry* (Yorktown Heights, New York, USA) (SOCG '86). Association for Computing Machinery, New York, NY, USA, 169–177. <https://doi.org/10.1145/10515.10534>
- [16] V. Cohen-Addad, A. Filtser, P. N. Klein, and H. Le. 2020. On Light Spanners, Low-treewidth Embeddings and Efficient Traversing in Minor-free Graphs. In *2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE, 589–600.
- [17] J. Cong, A. B. Kahng, G. Robins, M. Sarrafzadeh, and C. K. Wong. 1991. Performance-Driven Global Routing for Cell Based ICs. In *Proceedings of the 9th International Conference on Computer Design: VLSI in Computers and Processors (ICCD'91)*. 170–173.
- [18] J. Cong, A. B. Kahng, G. Robins, M. Sarrafzadeh, and C. K. Wong. 1992. Provably good performance-driven global routing. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* 11, 6 (1992), 739–752.
- [19] G. Das, P. Heffernan, and G. Narasimhan. 1993. Optimally Sparse Spanners in 3-dimensional Euclidean Space. In *Proceedings of the 9th Annual Symposium on Computational Geometry (SCG '93)*. 53–62.
- [20] E. D. Demaine, F. V. Fomin, M. T. Hajiaghayi, and D. M. Thilikos. 2005. Fixed-parameter algorithms for (k, r) -center in planar graphs and map graphs. *ACM Trans. Algorithms* 1, 1 (2005), 33–47. <https://doi.org/10.1145/1077464.1077468> preliminary version published in IICALP 2003.
- [21] E. D. Demaine and M. Hajiaghayi. 2005. Bidimensionality: New Connections Between FPT Algorithms and PTASs. In *Proceedings of the Sixteenth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'05)*. 590–601.
- [22] E. D. Demaine, M. Hajiaghayi, and K. Kawarabayashi. 2005. Algorithmic Graph Minor Theory: Decomposition, Approximation, and Coloring. In *Proceedings of the 46th Annual IEEE Symposium on Foundations of Computer Science*. 637–646. <https://doi.org/10.1109/SFCS.2005.14>
- [23] Z. Dvořák, R. McCarty, and S. Norin. 2021. Sublinear Separators in Intersection Graphs of Convex Shapes. *SIAM Journal on Discrete Mathematics* 35, 2 (2021), 1149–1164. <https://doi.org/10.1137/20m1311156>
- [24] Z. Dvořák. 2018. Thin Graph Classes and Polynomial-Time Approximation Schemes. In *Proceedings of the 29th Annual ACM-SIAM Symposium on Discrete Algorithms (New Orleans, Louisiana) (SODA '18)*. 1685–1701. <https://doi.org/10.5555/3174304.3175415>
- [25] Z. Dvořák and S. Norin. 2016. Strongly Sublinear Separators and Polynomial Expansion. *SIAM Journal of Discrete Math* 30 (2016), 1095–1101. <https://doi.org/10.1137/15M1017569>
- [26] M. Elkin and S. Solomon. 2015. Optimal Euclidean Spanners: Really Short, Thin, and Lanky. *J. ACM* 62, 5 (2015), 35:1–35:45.
- [27] D. Eppstein. 2000. Diameter and Treewidth in Minor-Closed Graph Families. *Algorithmica* 27, 3 (2000), 275–291. <https://doi.org/10.1007/s004530010020>
- [28] D. Eppstein and H. Khodabandeh. 2021. On the Edge Crossings of the Greedy Spanner. In *37th International Symposium on Computational Geometry (SoCG' 2021)*. To appear.
- [29] D. Eppstein, G. L. Miller, and S.-H. Teng. 1995. A deterministic linear time algorithm for geometric separators and its applications. *Fundamenta Informaticae* 22, 4 (1995), 309–329.
- [30] P. Erdős. 1964. Extremal problems in graph theory. *Theory of Graphs and Its Applications (Proc. Sympos. Smolenice)* (1964), 29–36.
- [31] A. Filtser and S. Solomon. 2016. The Greedy Spanner is Existentially Optimal. In *Proceedings of the 2016 ACM Symposium on Principles of Distributed Computing (PODC '16)*. 9–17. <https://doi.org/10.1145/2933057.2933114> to appear in SICOMP 2020.
- [32] F. V. Fomin, D. Lokshtanov, D. Marx, M. Pilipczuk, M. Pilipczuk, and S. Saurabh. 2016. Subexponential Parameterized Algorithms for Planar and Apex-Minor-Free Graphs via Low Treewidth Pattern Covering. In *2016 IEEE 57th Annual Symposium on Foundations of Computer Science (FOCS' 16)*. <https://doi.org/10.1109/focs.2016.62>
- [33] M. Fürer and S. P. Kasiviswanathan. 2012. Spanners for geometric intersection graphs with applications. *Journal of Computational Geometry* 3, 1 (2012), 31–64. Announced at WADS' 07.
- [34] J. Gao and L. Zhang. 2005. Well-Separated Pair Decomposition for the Unit-Disk Graph Metric and Its Applications. *SIAM J. Comput.* 35, 1 (2005), 151–169. <https://doi.org/10.1137/s0097539703436357>
- [35] L.-A. Gottlieb. 2015. A Light Metric Spanner. In *IEEE 56th Annual Symposium on Foundations of Computer Science, FOCS 2015, Berkeley, CA, USA, 17-20 October, 2015*. 759–772. <https://doi.org/10.1109/FOCS.2015.52>
- [36] L.-A. Gottlieb, A. Kontorovich, and R. Krauthgamer. 2017. Efficient Regression in Metric Spaces via Approximate Lipschitz Extension. *IEEE Trans. Inf. Theory* 63, 8 (2017), 4838–4849. <https://doi.org/10.1109/TIT.2017.2713820> preliminary version published in SIMBAD 2013.
- [37] L.-A. Gottlieb and L. Roditty. 2008. An optimal dynamic spanner for doubling metric spaces. In *Proceedings of the 16th Annual European Symposium on Algorithms (ESA'08)*. 478–489. <https://doi.org/10.1137/20m1311156>
- [38] S. Har-Peled. 2011. A simple proof of the existence of a planar separator. *arXiv preprint arXiv:1105.0103* (2011). <https://arxiv.org/abs/1105.0103>.
- [39] S. Har-Peled and M. Mendel. 2006. Fast Construction of Nets in Low-Dimensional Metrics and Their Applications. *SIAM J. Comput.* 35, 5 (2006), 1148–1184. <https://doi.org/10.1137/s0097539704446281>
- [40] S. Har-Peled and K. Quanrud. 2017. Approximation Algorithms for Polynomial-Expansion and Low-Density Graphs. *SIAM J. Comput.* 46, 6 (2017), 1712–1744. <https://doi.org/10.1137/16m1079336>

- [41] M. R. Henzinger, P. Klein, S. Rao, and S. Subramanian. 1997. Faster Shortest-Path Algorithms for Planar Graphs. *J. Comput. System Sci.* 55, 1 (1997), 3–23. <https://doi.org/10.1006/jcss.1997.1493>
- [42] F. Huang. 2013. *On the maximum weighted independent set problem with applications in wireless sensor networks*. Boston University.
- [43] P. N. Klein. 2005. A Linear-time Approximation Scheme for Planar Weighted TSP. In *Proceedings of the 46th Annual IEEE Symposium on Foundations of Computer Science (FOCS '05)*. 647–657. <https://doi.org/10.1109/SFCS.2005.7>
- [44] P. N. Klein. 2006. Subset Spanner for Planar Graphs, with Application to Subset TSP. In *Proceedings of the 38th Annual ACM Symposium on Theory of Computing (STOC '06)*. 749–756. <https://doi.org/10.1145/1132516.1132620>
- [45] H. Le and S. Solomon. 2019. Truly Optimal Euclidean Spanners. In *60th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2019, Baltimore, Maryland, USA, November 9-12, 2019*. 1078–1100. <https://doi.org/10.1109/FOCS.2019.00069>
- [46] X. Li, P. Wan, and O. Frieder. 2003. Coverage in wireless ad hoc sensor networks. *IEEE Transactions on computers* 52, 6 (2003), 753–763.
- [47] R. J. Lipton and R. E. Tarjan. 1980. Applications of a Planar Separator Theorem. *SIAM J. Comput.* 9, 3 (1980), 615–627. <https://doi.org/10.1137/0209046>
- [48] D. Mahjoub and D. W. Matula. 2009. Experimental Study of Independent and Dominating Sets in Wireless Sensor Networks Using Graph Coloring Algorithms. In *Wireless Algorithms, Systems, and Applications*. 32–42. https://doi.org/10.1007/978-3-642-03417-6_4
- [49] D. Marx and M. Pilipczuk. 2017. Subexponential Parameterized Algorithms for Graphs of Polynomial Growth. In *25th Annual European Symposium on Algorithms (ESA' 17)*. 59:1–59:15. <https://doi.org/10.4230/LIPIcs.ESA.2017.59>
- [50] G. L. Miller, S. Teng, W. Thurston, and S. A. Vavasis. 1997. Separators for sphere-packings and nearest neighbor graphs. *Journal of the ACM* 44, 1 (1997), 1–29. <https://doi.org/10.1145/256292.256294>
- [51] G. Narasimhan and M. Smid. 2007. *Geometric Spanner Networks*. Cambridge University Press.
- [52] D. Peleg. 2000. *Distributed Computing: A Locality-Sensitive Approach*. SIAM, Philadelphia, PA.
- [53] D. Peleg and A. A. Schäffer. 1989. Graph spanners. *Journal of Graph Theory* 13, 1 (1989), 99–116. <https://doi.org/10.1002/jgt.3190130114>
- [54] T. Pino, S. Choudhury, and F. Al-Turjman. 2018. Dominating Set Algorithms for Wireless Sensor Networks Survivability. *IEEE Access* 6 (2018), 17527–17532. <https://doi.org/10.1109/access.2018.2819083>
- [55] S. B. Rao and W. D. Smith. 1998. Approximating Geometrical Graphs via “Spanners” and “Banyans”. In *Proceedings of the 30th Annual ACM Symposium on Theory of Computing (STOC '98)*. 540–550. <https://doi.org/10.1145/276698.276868> Full version at <http://graphics.stanford.edu/courses/cs468-06-winter/Papers/rs-tsp.pdf>.
- [56] D. Russel and L. Guibas. 2005. Exploring protein folding trajectories using geometric spanners. In *Pacific Symposium on Biocomputing 2005*. https://doi.org/10.1142/9789812702456_0005
- [57] F. Sibel Salman, Joseph Cheriyan, R. Ravi, and S. Subramanian. 2001. Approximating the Single-Sink Link-Installation Problem in Network Design. *SIAM Journal on Optimization* 11, 3 (2001), 595–610. <https://doi.org/10.1137/S1052623497321432>
- [58] H. Shpungin and M. Segal. 2010. Near-Optimal Multicriteria Spanner Constructions in Wireless Ad Hoc Networks. *IEEE/ACM Transactions on Networking* 18, 6 (2010), 1963–1976. <https://doi.org/10.1109/TNET.2010.2053381>
- [59] A. Sidiropoulos and V. Sridhar. 2017. Algorithmic Interpretations of Fractal Dimension. In *Proceedings of the 33rd Annual Symposium on Computational Geometry (SoCG'17)*. 58:1–58:16. <https://doi.org/10.4230/LIPIcs.SoCG.2017.58>
- [60] C. Singh, A. Nedic, and R. Srikant. 2014. LP-relaxation based distributed algorithms for scheduling in wireless networks. In *IEEE INFOCOM 2014 - IEEE Conference on Computer Communications*. <https://doi.org/10.1109/infocom.2014.6848130>
- [61] M. Smid. 2009. The Weak Gap Property in Metric Spaces of Bounded Doubling Dimension. In *Efficient Algorithms*, Susanne Albers, Helmut Alt, and Stefan Näher (Eds.). Springer-Verlag, 275–289.
- [62] W. D. Smith and N. C. Wormald. [n. d.]. Geometric separator theorems and applications. In *Proceedings 39th Annual Symposium on Foundations of Computer Science (FOCS'98)*. 232–243. <https://doi.org/10.1109/sfcs.1998.743449>
- [63] P. von Rickenbach and R. Wattenhofer. 2004. Gathering correlated data in sensor networks. In *DIALM-POMC Joint Workshop on Foundations of Mobile Computing*. 60–66. <https://doi.org/10.1145/1022630.1022640>
- [64] P. Wang and S. Bohacek. 2008. On the Practical Complexity of Solving the Maximum Weighted Independent Set Problem for Optimal Scheduling in Wireless Networks. In *Proceedings of the 4th Annual International Conference on Wireless Internet (WICON '08)*. Article 15, 9 pages.
- [65] J. Yu, N. Wang, G. Wang, and D. Yu. 2013. Connected dominating sets in wireless ad hoc and sensor networks – A comprehensive survey. *Computer Communications* 36, 2 (2013), 121–134. <https://doi.org/10.1016/j.comcom.2012.10.005>
- [66] Y. Zhao and S. Teng. 2009. Combinatorial and spectral aspects of nearest neighbor graphs in doubling dimensional and nearly-Euclidean spaces. *Theoretical Computer Science* 410, 11 (2009), 1081–1092. <https://doi.org/10.1016/j.tcs.2008.10.032>