

Light Euclidean Spanners with Steiner Points

Hung Le

University of Victoria, Canada
University of Massachusetts, Amherst, MA, USA
hungle@uvic.ca

Shay Solomon

Tel-Aviv University, Israel
solo.shay@gmail.com

Abstract

The FOCS'19 paper of Le and Solomon [59], culminating a long line of research on Euclidean spanners, proves that the lightness (normalized weight) of the greedy $(1 + \epsilon)$ -spanner in \mathbb{R}^d is $\tilde{O}(\epsilon^{-d})$ for any $d = O(1)$ and any $\epsilon = \Omega(n^{-\frac{1}{d-1}})$ (where \tilde{O} hides polylogarithmic factors of $\frac{1}{\epsilon}$), and also shows the existence of point sets in \mathbb{R}^d for which any $(1 + \epsilon)$ -spanner must have lightness $\Omega(\epsilon^{-d})$.¹ Given this tight bound on the lightness, a natural arising question is whether a better lightness bound can be achieved using *Steiner points*.

Our first result is a construction of Steiner spanners in \mathbb{R}^2 with lightness $O(\epsilon^{-1} \log \Delta)$, where Δ is the spread of the point set.² In the regime of $\Delta \ll 2^{1/\epsilon}$, this provides an improvement over the lightness bound of [59]; this regime of parameters is of practical interest, as point sets arising in real-life applications (e.g., for various random distributions) have polynomially bounded spread, while in spanner applications ϵ often controls the precision, and it sometimes needs to be much smaller than $O(1/\log n)$. Moreover, for spread polynomially bounded in $1/\epsilon$, this upper bound provides a quadratic improvement over the non-Steiner bound of [59]. We then demonstrate that such a light spanner can be constructed in $O_\epsilon(n)$ time for polynomially bounded spread, where O_ϵ hides a factor of $\text{poly}(\frac{1}{\epsilon})$. Finally, we extend the construction to higher dimensions, proving a lightness upper bound of $\tilde{O}(\epsilon^{-(d+1)/2} + \epsilon^{-2} \log \Delta)$ for any $3 \leq d = O(1)$ and any $\epsilon = \Omega(n^{-\frac{1}{d-1}})$.

2012 ACM Subject Classification Theory of computation \rightarrow Computational geometry

Keywords and phrases Euclidean spanners, Steiner spanners, light spanners

Digital Object Identifier 10.4230/LIPIcs.ESA.2020.67

Funding *Hung Le*: Supported by a PIMS postdoctoral fellowship and an NSERC grant.

Shay Solomon: Partially supported by the Israel Science Foundation grant No.1991/19 and by Len Blavatnik and the Blavatnik Family foundation.

Acknowledgements Shay Solomon is grateful to Michael Elkin, Ofer Neiman and Michiel Smid for fruitful discussions.

1 Introduction

A t -spanner for a set P of points in the d -dimensional Euclidean space \mathbb{R}^d is a *geometric graph* that preserves all the pairwise Euclidean distances between points in P to within a factor of t , called the *stretch factor*; by geometric graph we mean a weighted graph in which the vertices correspond to points in \mathbb{R}^d and the edge weights are the Euclidean distances between the corresponding points. The study of Euclidean spanners dates back to the seminal work of Chew [27, 28] from 1986, which presented a spanner with constant stretch and $O(n)$ edges for any set of n points in \mathbb{R}^2 . In the three following decades, Euclidean spanners

¹ The lightness of a spanner is the ratio of its weight and the MST weight.

² The spread $\Delta = \Delta(P)$ of a point set P in \mathbb{R}^d is the ratio of the largest to the smallest pairwise distance.



have evolved into an important subarea of Discrete and Computational Geometry, having found applications in many different areas, such as approximation algorithms [65], geometric distance oracles [42, 45, 44, 43], and network design [52, 60]; see the book by Narasimhan and Smid “Geometric Spanner Networks” [61] for an excellent account on Euclidean spanners and some of their applications. Numerous constructions of Euclidean spanners in two and higher dimensions were introduced over the years, such as Yao graphs [70], Θ -graphs [29, 54, 55, 66], the (path-)greedy spanner [4, 23, 61] and the gap-greedy spanner [67, 6] – unveiling an abundance of techniques, tools and insights along the way; refer to the book of [61] for more spanner constructions.

In addition to low stretch, many applications require that the spanner would be *sparse*, in the unweighted and/or weighted sense. The *sparsity* (respectively, *lightness*) of a spanner is the ratio of its size (respectively, weight) to the size (resp., weight) of a spanning tree (resp., MST), providing a normalized notion of size (resp., weight), which should ideally be $O(1)$. For any dimension $d = O(1)$, spanners with constant sparsity are known since the 80s [70, 29, 54, 55, 66, 4, 67, 6]; also, it is known since the early 90s that the greedy spanner has constant lightness [4]. The constant bounds on the sparsity and lightness depend on both ϵ and d . In some applications, ϵ must be a very small sub-constant parameter, so as to achieve the highest possible precision and minimize potential errors, and in some situations ϵ may be as small as n^{-c} for some constant $0 < c < 1$. Consequently, besides the theoretical appeal, achieving the precise dependencies on ϵ and d in the sparsity and lightness bounds is of practical importance.

Culminating a long line of work, in FOCS’19 Le and Solomon [59] showed that the precise dependencies on ϵ in the sparsity and lightness bounds are $\Theta(\epsilon^{1-d})$ and $\tilde{\Theta}(\epsilon^{-d})$, respectively, for any $d = O(1)$ and any $\epsilon = \Omega(n^{-\frac{1}{d-1}})$; throughout we shall use $\tilde{O}, \tilde{\Omega}, \tilde{\Theta}$ to hide polylogarithmic factors of $\frac{1}{\epsilon}$. The lower bounds of [59] are proved for the d -dimensional sphere, for $d = O(1)$. On the upper bound side, sparsity $O(\epsilon^{1-d})$ is achieved by a number of classic constructions such as Yao graphs [70], Θ -graphs [29, 54, 55, 66], the (path-)greedy spanner [4, 23, 61] and the gap-greedy spanner [67, 6], and the argument underlying all these upper bounds is basic and simple. On the other hand, constant lightness upper bound (regardless of the dependency on ϵ and d) is achieved only by the greedy algorithm [4, 31, 32, 65, 61, 41, 13, 59], and all known arguments for constant lightness are highly nontrivial, even for $d = 2$; in fact, the proofs in [4, 31, 32, 65] have missing details. The first complete proof was given in the book of [61], in a 60-page chapter, where it is shown that the greedy $(1 + \epsilon)$ -spanner has lightness $O(\epsilon^{-2d})$, which improved the dependencies on ϵ and d given in all previous work. In SODA’19, Borradaile, Le and Wulff-Nilsen [13] presented a shorter and arguably simpler alternative proof that applies to the wider family of *doubling metrics*, which improves the ϵ dependency provided in FOCS’15 by Gottlieb [41] for doubling metrics, but is inferior to the lightness bound of $O(\epsilon^{-2d})$ by [61].³ Finally, the lightness bound of $\tilde{O}(\epsilon^{-d})$ of [59] was proved via a tour-de-force argument; interestingly, the proof for $d = 2$ is much more intricate than for higher dimensions $d \geq 3$.

Le and Solomon [59] showed that, counter-intuitively, one can use Steiner points to bypass the sparsity lower bound of Euclidean spanners. Specifically, in this way they reduced the sparsity upper bound almost quadratically to $\tilde{O}(\epsilon^{(1-d)/2})$, and also provided a matching

³ The *doubling dimension* of a metric space (X, δ) is the smallest value d such that every ball B in the metric space can be covered by at most 2^d balls of half the radius of B . This notion generalizes the Euclidean dimension, since the doubling dimension of the Euclidean space \mathbb{R}^d is $\Theta(d)$. A metric space is called *doubling* if its doubling dimension is constant.

lower bound for $d = 2$. Their lower bound for sparsity is derived from a lightness lower bound; specifically, they first prove that, for a set of points evenly spaced on the boundary of the unit square, with distances $\Theta(\sqrt{\epsilon})$ between neighboring points, any Steiner $(1 + \epsilon)$ -spanner must incur lightness $\tilde{\Omega}(\frac{1}{\epsilon})$, and then they translated the lightness lower bound into a sparsity lower bound of $\tilde{\Omega}(\sqrt{1/\epsilon})$. Whether or not Steiner points can be used to reduce the lightness remained open in [59], but this should not come as a surprise.⁴ First, bounding the lightness of spanners is inherently more difficult than bounding their sparsity – this is true for both Euclidean spaces and doubling metrics, as well as other graph families including general weighted graphs [4, 31, 32, 65, 61, 41, 13, 59, 26, 39, 33]. Second, constructing Steiner trees and spanners *with asymptotically improved bounds* is inherently more difficult than constructing their non-Steiner counterparts [34, 35, 68, 57, 11, 58]. In our particular case, while the classic sparsity upper bound for non-Steiner Euclidean spanners is simple, its improved counterpart for Steiner spanners by [59] requires a number of nontrivial insights and is rather intricate. On the other hand, the lightness upper bound of [59] uses a tour-de-force argument, and as mentioned already this is true even for $d = 2$. Consequently, obtaining an improved lightness bound using Steiner points seems currently out of reach, at least until an inherently simpler proof to [59] for non-Steiner lightness is found (if one exists).

1.1 Our Contribution

In this paper, we explore the power of Steiner points in reducing lightness for Euclidean spaces of bounded spread Δ .⁵ Point sets of bounded spread have been studied extensively for Euclidean spanners and related geometric objects [9, 53, 30, 51, 17, 16, 36, 64, 1, 2, 5, 49, 50, 19, 62, 63, 18, 69, 21]. The motivation for studying point sets of bounded spread is three-fold.

1. Such point sets arise naturally in practice, and are thus important in their own right; indeed, for many random distributions, the spread is polynomial in the number of points – in expectation and with high probability. In particular, it is known that for n -point sets drawn uniformly at random from the unit square, the expected spread is $\Theta(n)$, and the expected spread in the unit d -dimensional hypercube is $n^{2/d}$ for any $d = O(1)$. Researchers have studied random distributions of point sets, in part to explain the success of solving various geometric optimization problems in practice [9, 53, 30], and there are many results on spanners for random point sets [24, 7, 37, 10, 38, 14, 61, 69, 15, 21, 8, 56]. Of course, the family of bounded spread point sets is much wider than that of random point sets.
2. Euclidean spanners can be constructed in (deterministic) $O(n)$ time in such point sets [20], while there is no $o(n \log n)$ -time algorithm for constructing Euclidean spanners in arbitrary point sets. (The result of [20] extends to other geometric objects, such as WSPD and compressed quadtrees, and some of the aforementioned references (e.g., [50, 17]) build on this result to achieve faster algorithms for other geometric problems for point sets of polynomially bounded spread.)
3. The case of bounded spread is sometimes used as a stepping stone towards the general case; see, e.g., [40, 47, 46, 21, 22, 3, 59].

⁴ The lightness of Steiner spanners can be defined with respect to the SMT (Steiner minimum tree) weight, but we can also stick to the original definition, since the SMT and MST weights differ by a constant factor smaller than 2.

⁵ The *spread* of a Euclidean space is the ratio of the maximum to minimum pairwise distances in it.

As mentioned above, Le and Solomon [59] showed that for a set of points evenly spaced on the boundary of the unit square, with distances $\Theta(\sqrt{\epsilon})$ between neighboring points, any Steiner $(1 + \epsilon)$ -spanner must incur lightness $\tilde{\Omega}(\frac{1}{\epsilon})$. Note that the spread of this point set is $\frac{1}{\sqrt{\epsilon}}$. The non-Steiner upper bound for $d = 2$ by [59] is $\tilde{O}(\epsilon^{-2})$. A natural arising question is whether one can improve the lightness upper bound of $\tilde{O}(\epsilon^{-2})$ using Steiner points, ideally quadratically, for point sets of bounded spread. We answer this question in the affirmative.

► **Theorem 1.** *Any point set P in \mathbb{R}^2 of spread Δ admits a Steiner $(1 + \epsilon)$ -spanner of lightness $O(\frac{\log \Delta}{\epsilon})$.*

Recalling that the lower bound of $\tilde{\Omega}(\frac{1}{\epsilon})$ by [59] applies to a point set of spread $\frac{1}{\sqrt{\epsilon}}$, the lightness upper bound provided by Theorem 1 is therefore tight (up to polylogarithmic factors in $\frac{1}{\epsilon}$) for point sets of spread $\text{poly}(\frac{1}{\epsilon})$. Moreover, this upper bound improves the general upper bound of $\tilde{O}(\epsilon^{-2})$ from [59] in the regime $\log \Delta \ll \epsilon^{-1}$, i.e., when $\Delta \ll 2^{1/\epsilon}$. For spread polynomial in n , we get an improvement over [59] as long as $\epsilon \ll \frac{1}{\log n}$. Of course, the improvement gets more significant as ϵ decays – in the most extreme situation ϵ is inverse polynomial in n , and then the improvement over [59] is polynomial in n even when Δ is exponential in n .

Our second result is that a Steiner spanner with near-optimal lightness can be constructed in linear time, when the spread Δ is polynomial in n .

► **Theorem 2.** *For any point set P in \mathbb{R}^2 of spread $\Delta = O(n^c)$, for any $c = O(1)$, a Steiner $(1 + \epsilon)$ -spanner of lightness $O(\frac{\log \Delta}{\epsilon})$ can be constructed in $O_\epsilon(n)$ time, where $O_\epsilon(\cdot)$ hides a factor of $\text{poly}(\frac{1}{\epsilon})$.*

Higher dimensions

The lower bound of [58] states that any (non-Steiner) $(1 + \epsilon)$ -spanner must incur lightness $\Omega(\epsilon^{-d})$, for any $d = O(1)$. We show that, similarly to the 2-dimensional case, one can improve the lightness almost quadratically using Steiner points, for point sets of bounded spread.

► **Theorem 3.** *For any $d \geq 3$, any point set P in \mathbb{R}^d of spread Δ admits a Steiner $(1 + \epsilon)$ -spanner of lightness $\tilde{O}(\epsilon^{-(d+1)/2} + \epsilon^{-2} \log \Delta)$.*

Interestingly, the dependence on the spread in the lightness bound provided by Theorem 3 does not grow with the dimension, hence the improvement over the non-Steiner bound gets more significant as the dimension grows, provided of course that the spread is not too large.

1.2 Proof Overview

Proof of Theorem 1

We partition the set of pairs of points into $m = O(\log \Delta)$ subsets $\{\mathcal{P}_1, \dots, \mathcal{P}_m\}$ where \mathcal{P}_i contains pairs of distances in $[2^{i-1}, 2^i]$. The objective is to show that one can preserve distances between all pairs in \mathcal{P}_i to within a factor of $(1 + \epsilon)$ using a Steiner spanner S_i of weight $O(\frac{w(\text{MST})}{\epsilon})$; by taking the union of all such spanners, we obtain a Steiner spanner with the required lightness.

Let $L_i = 2^i$. A natural idea to preserve distances in \mathcal{P}_i is to (a) find an (ϵL_i) -net N_i ⁶ of P and (b) add to the spanner edges between any two net points p and q such that $(u, v) \in \mathcal{P}_i$ and there is a pair (u, v) in \mathcal{P}_i such that u and v are covered by (i.e., within distance ϵL_i

⁶ A subset of points $N \subseteq P$ is an r -net if every point in P is within distance r from (or covered by) some point in N and pairwise distances between points in N are larger than r .

from) p and q , respectively. The stretch will be in check because u and v are at distance roughly $O(\epsilon)\|u, v\|$ from their net points and thus, the *additive* stretch between u and v is $O(\epsilon)\|u, v\|$. For lightness, we can show that the number of net points $|N_i| = O(\frac{w(\text{MST})}{\epsilon L_i})$, and using a (nontrivial) packing argument, there are about $\frac{|N_i|}{\epsilon}$ edges of length $O(L_i)$ added in step (b). Thus, the total weight of the spanner is $O(\frac{w(\text{MST})}{\epsilon^2})$, which is bigger than our aimed lightness bound by a factor of $\frac{1}{\epsilon}$.

To shave the factor of $\frac{1}{\epsilon}$, we employ two ideas. First, we take N_i to be a $\sqrt{\epsilon}L_i$ -net of P . In this way $|N_i| = O(\frac{w(\text{MST})}{\sqrt{\epsilon}})$. By applying the 2-dimensional Steiner spanner construction of [59] as a blackbox, we obtain a Steiner spanner with only $\frac{|N_i|}{\sqrt{\epsilon}}$ edges of length $O(L_i)$ that approximates distances between points in N_i ; we have thus reduced the weight bound to $O(\frac{w(\text{MST})}{\epsilon})$, as required. The problem now is with the stretch guarantee: Preserving distances between the points in N_i is no longer sufficient. Indeed, for every pair $(u, v) \in \mathcal{P}_i$, the distance between u and v to the nearest net points is $O(\sqrt{\epsilon})\|u, v\|$, hence the resulting stretch is $(1 + O(\sqrt{\epsilon}))\|u, v\|$. We overcome this hurdle by introducing a novel construction of *single-source spanners*, which generalize Steiner shallow-light trees of Solomon [68]. Specifically, we open the black-box of [59] and observe that every time we want to preserve the distance from (some) Steiner point s to a net-point $p \in N_i$, instead of connecting s to p by a straight line of weight $O(L_i)$, we can use a single-source spanner (rooted at s) of *weight* $O(L_i)$ to preserve distances (up to a $(1 + \epsilon)$ factor) from s to *every point* within distances $O(\sqrt{\epsilon}L_i)$ from p . As a result, our Steiner spanner can preserve distances between any two points $(u, v) \in \mathcal{P}_i$ to within a factor of $(1 + \epsilon)$, where $u \in B(p, O(\sqrt{\epsilon}L_i))$, $v \in B(q, O(\sqrt{\epsilon}L_i))$, and p, q are their nearest net points.

Proof of Theorem 3

By extending the construction in Theorem 1, we can construct a Steiner spanner with lightness $\tilde{O}(\epsilon^{-(d+1)/2} \log \Delta)$ as follows: for each \mathcal{P}_i , we construct an ϵL_i -net N_i and then apply the construction of [59] as a black box to obtain a Steiner spanner S_i for N_i with weight $\tilde{O}(\epsilon^{-(d-1)/2}|N_i|L_i)$. The stretch will be in check since N_i is an ϵL_i -net. Since $|N_i| = O(\frac{w(\text{MST})}{\epsilon L_i})$, $w(S_i) = \tilde{O}(\epsilon^{-(d+1)/2}w(\text{MST}))$. The union of Steiner spanners S_i for all $i \in [1, m]$ has weight $\tilde{O}(\epsilon^{-(d+1)/2} \log \Delta)w(\text{MST})$. Note when $d \geq 3$, we can not take N_i as a $\sqrt{\epsilon}L_i$ -net since the construction of single-source spanners with $O(L_i)$ weight in the proof of Theorem 1 only works when $d = 2$.

Most of our effort is to further refine the result in a way that $\log \Delta$ term is multiplied only by ϵ^{-2} and not by the term that depends on d . We first reduce to the problem of approximating distances between pairs (of endpoints) in a family of edge sets $\mathcal{E} = \{E_1, \dots, E_m\}$ with $m = O(\log \Delta)$, where edges of E_i have length (roughly) in the interval $(\frac{1}{2\epsilon^i}, \frac{1}{\epsilon^i}]$ and edges in E_1 have length in $[1, \frac{1}{\epsilon})$. Let $L_i = \frac{1}{\epsilon^i}$. For a technical reason, we will subdivide edges of MST by using a set of Steiner points K so that each new edge has length in $(1/2, 1]$.

We construct a Steiner spanner for edges in \mathcal{E} using a *charging cover tree* T : T has depth $m + 1$, level i of T is associated with an ϵL_i -cover⁷ of $P \cup K$, and leaves (at level 0) of T are points in $P \cup K$. For each cover N_i at level i of T , we construct a graph H_i where $V(H_i) = N_i$ and there is an edge (u, v) between $u, v \in N_i$ if there is a corresponding edge in E_i whose endpoints are covered by u and v , respectively. Graph H_i is used to distinguish between *low degree points*, whose degree in H_i is $O(\frac{1}{\epsilon})$, and *high degree points*, whose degree

⁷ A subset of points $N \subseteq P$ is an r -cover if every point in P is within distance r from some point in N .

in H_i is $\Omega(\frac{1}{\epsilon})$. T will have two *charging properties*: (1) every point $p \in N_i$ has at least ϵL_i *uncharged descendants* and (2) at every level i , one can charge up to $\frac{\epsilon L_i}{2}$ uncharged descendants of high degree points. Note that, once an uncharged point is *charged* at level i , it will be marked as charged at higher levels; initially at level 0, every point is *uncharged*.

We then use a charging cover tree T to guide the Steiner spanner construction. Specifically, at level i , we add all edges incident to low degree points to the spanner and we can show that the total weight of all these edges over all levels is at most $O(\frac{w(\text{MST}) \log \Delta}{\epsilon^2})$. For high degree points, we apply the construction of Le and Solomon [59] to obtain a Steiner spanner S_i , and we charge the weight of S_i to $\frac{L_i \epsilon}{2}$ uncharged descendants of each high degree point. This charging is possible by the charging property (2) of T . We then show that each point in $K \cup P$ is charged a weight at most $\tilde{O}(\epsilon^{-(d+1)/2})$. Thus, the total weight of the Steiner spanners (for high degree points) at all levels is $\tilde{O}(\epsilon^{-(d+1)/2})w(\text{MST})$.

Our construction of a charging cover tree is inspired by the construction of a hierarchy of clusters in the iterative clustering technique. The technique was initially developed by Chechick and Wulff-Nilsen [25] to construct light spanners for general graphs, and then was adapted to many other different settings [25, 12, 13, 59, 58]. Our construction is directly inspired by the construction of Borradaile et al. [13] in the doubling dimension setting. However, our construction is much simpler. Specifically, we are able to decouple the Steiner spanner construction from the charging cover tree construction. We refer readers to Section 5 for more details.

2 Preliminaries

Let P be a point set of n points in \mathbb{R}^d . We denote by $\|p, q\|$ the Euclidean distance between two points $p, q \in \mathbb{R}^d$. Let $B(p, r) = \{x \in \mathbb{R}^d, \|p, x\| \leq r\}$ be the ball of radius r centered at p . Given a point p and a set of point Q on the plane, we define the distance between p and Q , denoted by $d(p, Q)$, to be $\inf_{x \in Q} \|p, x\|$.

An r -cover of P is a subset of points $N \subseteq P$ such that for every point $x \in P$, there is at least one point $p \in N$ such that $\|p, x\| \leq r$; we say x is *covered* by p . When the value of r is clear from the context, we simply call N a cover of P . A subset of point $N \subseteq P$ is called an r -net if N is an r -cover of P and also an r -packing of P , i.e., for every two points $p \neq q \in N$, $\|p, q\| > r$.

Let G be a graph with weight function w on the edges. We denote the vertex set and edge set of G by $V(G)$ and $E(G)$, respectively. Let $d_G(p, q)$ be the distance between two vertices p, q of G . We denote by $G[X]$ the subgraph induced by a subset of vertices X .

G is *geometric* in \mathbb{R}^d if each vertex of G corresponds to a point $p \in \mathbb{R}^d$ and for every edge (p, q) , $w(p, q) = \|p, q\|$. In this case, we use points to refer to vertices of G . We say that a geometric graph G is a $(1 + \epsilon)$ -spanner of P if $V(G) = P$ and for every two points $p \neq q \in P$, $d_G(p, q) \leq (1 + \epsilon)\|p, q\|$. We say that G is a Steiner $(1 + \epsilon)$ -spanner for P if $P \subseteq V(G)$ and for every two points $p \neq q \in P$, $d_G(p, q) \leq (1 + \epsilon)\|p, q\|$. Points in $V(G) \setminus P$ are called *Steiner points*. Note that distances between Steiner points may not be preserved in a Steiner $(1 + \epsilon)$ -spanner.

3 Steiner Spanners on the Plane

We focus on constructing a Steiner spanner with good lightness; the fast construction is in Section 4. We will use the following geometric Steiner *shallow-light tree* (SLT) construction by Solomon [68].

► **Lemma 4.** *Let L be a line segment of length $\sqrt{\epsilon}$ and p be a point on the plane such that $d(p, L) = 1$. For any point set $X \in L$, there is a geometric graph H of weight $\Theta(1)$ such that $d_H(p, x) \leq (1 + \epsilon)\|p, x\|$ for any point $x \in X$.*

We will use *single-source spanners* (defined below) as a black box in our construction.

► **Definition 5** (Single-source spanners). *Given a point p (source), a set of points X on the plane and a connected geometric graph S_X spanning X , a single source $(1 + \epsilon)$ -spanner w.r.t. (p, X, S_X) is a graph H such that for every $x \in X$: $\|p, x\| \leq d_{H \cup S_X}(p, x) \leq (1 + \epsilon)\|p, x\|$.*

Our starting point is the construction of a single source spanner from a point p to point set X enclosed in a circle C of radius $\sqrt{\epsilon}$ such that $d(p, C) = 1$. We show that, if S_X approximately preserves the distances between pairs of points in X up to a $(1 + g\epsilon)$ factor for any constant g , it is possible to construct a single-source spanner with weight $O(1)$. It is not so hard to see that if Steiner points are not allowed, a lower bound of weight $\Omega(\frac{1}{\sqrt{\epsilon}})$ holds here.

► **Lemma 6.** *Let X be a set of points in a circle C of radius $\sqrt{\epsilon}$ on the plane and a point p of distance 1 from C . Let S_X be a $(1 + g\epsilon)$ -spanner of X for any constant g . Then there is a single-source $(1 + 13\epsilon)$ -spanner H w.r.t. (p, X, S_X) of weight $O(1)$ when $g \ll \frac{1}{\epsilon}$.*

Proof. Let c be a center of C . W.l.o.g, we assume that pc is parallel to y -axis. Let Q be the axis-aligned smallest square bounding C . Observe that the side length of Q is at mos $2\sqrt{\epsilon}$. Place a $\frac{2}{\sqrt{\epsilon}} \times \frac{2}{\sqrt{\epsilon}}$ grid W on Q , so that every cell of W is a square of side length ϵ . Observe that:

$$w(W) \leq \epsilon \cdot \frac{4}{\epsilon} = O(1)$$

We extend W to W_1 by connecting an (arbitrary) corner of each grid cell to an arbitrary point of X in the cell. Observe that: $w(W_1) = O(W) = O(1)$. Let P be the set of grid points on the side, say L , of Q that is closer to p (than the opposite side). We apply the construction in Lemma 4 to p and L to obtain a geometric graph K . Let $H = W_1 \cup K$. Since $w(K) = O(1)$ by Lemma 4, it holds that $w(H) = O(1)$.

It remains to show the stretch bound. Let x be any point of X and v be the point in the same cell with x that is connected to a corner, say z of grid W . We will show below that:


$$d_{W \cup K}(z, p) \leq (1 + 3\epsilon)\|p, z\| \tag{1}$$

If Equation 1 holds, it would imply:

$$\begin{aligned} d_{S_X \cup H}(x, p) &\leq d_{S_X}(x, v) + d_H(v, p) \leq (1 + g\epsilon)\sqrt{2}\epsilon + d_H(v, p) \\ &\leq (1 + g\epsilon)\sqrt{2}\epsilon + d_H(z, v) + d_H(z, p) \stackrel{g \ll 1/\epsilon}{\leq} 2\sqrt{2}\epsilon + \sqrt{2}\epsilon + d_{W \cup K}(z, p) \\ &\stackrel{\text{Eq. 1}}{\leq} (1 + 3\epsilon)\|p, z\| + 3\sqrt{2}\epsilon \stackrel{\|x, z\| \leq \sqrt{2}\epsilon}{\leq} (1 + 3\epsilon)(\|p, x\| + \sqrt{2}\epsilon) + 3\sqrt{2}\epsilon \\ &\leq (1 + 3\epsilon)\|p, x\| + 7\sqrt{2}\epsilon \leq (1 + 13\epsilon)\|p, x\| \quad \text{since } \|p, x\| \geq 1 \end{aligned}$$

Thus, it remains to prove Equation 1. To this end, let y be the projection of z on L . (Point y is also a grid point; see Figure 1) Let y', z' be projections of y and z on the line containing pc , respectively. Let u be the intersection of pz and yy' . Observe that $\|p, y\| \leq (1 + \epsilon) \leq (1 + \epsilon)\|p, u\|$. Thus,

$$\|p, y\| + \|y, z\| \leq (1 + \epsilon)\|p, u\| + \|y, z\| \leq (1 + \epsilon)\|p, u\| + \|u, z\| \leq (1 + \epsilon)\|p, z\| \tag{2}$$



■ **Figure 1** (Left) A single source spanner from p to a set of points enclosed by a circle of radius $\sqrt{\epsilon}$. One point in each non-empty cell is connected to a corner by a thick edge. (Right) An illustration for analyzing the stretch of pz .

Since $d_W(z, y) = \|z, y\|$ and $d_K(y, p) = (1 + \epsilon)\|p, y\|$ by Lemma 4, we have:

$$d_{W \cup K}(z, p) \leq (1 + \epsilon)(\|z, y\| + \|p, y\|) \stackrel{\text{Eq. 2}}{\leq} (1 + \epsilon)(1 + \epsilon)\|p, z\| \leq (1 + 3\epsilon)\|p, z\| \quad (3)$$

which implies Equation 1. ◀

We obtain the following corollary of Lemma 6.


► **Corollary 7.** Let X be a set of points in a circle C of radius $\sqrt{\epsilon}L$ on the plane and a point p of distance L/h from C for some constant $h \geq 1$. Let S_X be a $(1 + g\epsilon)$ -spanner of X for any constant g . Then there is a single-source $(1 + 13\epsilon)$ -spanner H w.r.t. (p, X, S_X) of weight $O(hL)$ when $g \ll \frac{1}{\epsilon}$.

Proof. We scale the space by L/h . In the scaled space, C has radius $h\sqrt{\epsilon}$ and $d(p, C) = 1$. Let C_1, C_2, \dots, C_m , where $m = O(h^2)$, be circles of radius $\sqrt{\epsilon}$ covering C ; such a set of circles can be constructed greedily. We apply Lemma 6 to p and each C_i to construct a single-source $(1 + 13\epsilon)$ -spanner H_i from p to each C_i . The final spanner is $H = \cup_{i=1}^m H_i$ that has total weight $O(h^2)$ in the scaled metric. Thus, in the original metric, $w(H) = O(h^2L/h) = O(hL)$. ◀

We are now ready to prove Theorem 1.

Proof of Theorem 1. Assume that the minimum pairwise distance is 1. Let $\mathcal{P} = \binom{P}{2}$ be all pairs of points. Partition \mathcal{P} into $O(\log \Delta)$ sets $\mathcal{P}_1, \mathcal{P}_2, \dots, \mathcal{P}_{\lceil \log \Delta \rceil}$ where \mathcal{P}_i is the set of pairs (x, y) such that $\|x, y\| \in [2^{i-1}, 2^i]$.

For a fixed i , we claim that there is a geometric graph H_i such that for every two distinct points $(x, y) \in \mathcal{P}_i$, $d_{H_1 \cup \dots \cup H_i}(x, y) \leq (1 + \epsilon)\|x, y\|$ and that $w(H_i) = O(\frac{1}{\epsilon})w(\text{MST})$. Thus, $H_1 \cup \dots \cup H_{\lceil \log \Delta \rceil}$ is a Steiner spanner with weight $O(\frac{\log \Delta}{\epsilon})w(\text{MST})$.



■ **Figure 2** (a) Square B is divided into $O(1)$ horizontal and vertical bands of width $L_i/8$ each. (b) The Steiner spanner construction for two non-adjacent horizontal bands H_a and H_b . The dashed cone represents a single-source spanner from r to circle $B(p, \sqrt{\epsilon}L_i)$.

We now focus on constructing H_i . Let $S_{i-1} = H_1 \cup H_2 \dots \cup H_{i-1}$. We will construct a spanner with stretch $(1 + c\epsilon)$ for some constant c . By induction, we can assume that:

$$d_{S_{i-1}}(p, q) \leq (1 + c\epsilon)\|p, q\| \tag{4}$$

for any pair $(p, q) \in \mathcal{P}_1 \cup \dots \cup \mathcal{P}_{i-1}$.

Let $L_i = 2^i$ and N_i be a $(\sqrt{\epsilon}L_i)$ -net of P . For each point x , let $N_i(x)$ be the net point that covers x : the distance from x to $N_i(x)$ is at most $\sqrt{\epsilon}L_i$.

▷ **Claim 8.** $|N_i| = O(\frac{w(\text{MST})}{L_i\sqrt{\epsilon}})$.

Proof. Consider the circle $B(p, \sqrt{\epsilon}L_i)$ centered at p ; $B(p, \sqrt{\epsilon}L_i)$ contains a segment of length $\Omega(\sqrt{\epsilon}L_i)$ of the MST, which is not contained in any other circle. Thus, the claim holds. ◁

Next, we consider the smallest axis-aligned square Q bounding the point set. In the following, we divide Q into a set of (overlapping) sub-squares \mathcal{B} of side length $\Theta(L_i)$ each. This way, for any pair $(x, y) \in \mathcal{P}_i$ (of distance at most L_i), there is a sub-square entirely containing $N_i(x), N_i(y)$, and the balls of radius $O(\sqrt{\epsilon}L_i)$ around the two net points.

Constructing \mathcal{B} . We first divide Q into subsquares of side length $5L_i$ each⁸. For each subsquare B , we extend its borders equally to four directions by an amount of $2L_i$ in each direction. After this extension, B has side length $9L_i$.

▷ **Claim 9.** Every point in Q belongs to at most 4 subsquares in \mathcal{B} . Furthermore, for each pair $(x, y) \in \mathcal{P}_i$, there is a subsquare $B \in \mathcal{B}$ such that $B(N_i(x), \sqrt{\epsilon}L_i), B(N_i(y), \sqrt{\epsilon}L_i)$ are entirely contained in B .

Proof. Let B be a subsquare in \mathcal{B} containing one of the endpoints of (x, y) , say x , before extension. Then, after extension, B will contain both x, y since $\|x, y\| \leq L_i$, and furthermore, x and y are at least L_i away from the boundary since we extended B by $2L_i$ in each direction. Thus, points in $B(N_i(x), \sqrt{\epsilon}L_i)$ ($B(N_i(y), \sqrt{\epsilon}L_i)$) will be at most $2\sqrt{\epsilon}L_i < L_i$ from x (y) when $\epsilon \ll 1$. ◁

⁸ We assume that the side length of Q is divisible by $5L_i$; otherwise, we can extend Q in such a way.

67:10 Light Euclidean Spanners with Steiner Points

Consider a subsquare $B \in \mathcal{B}$. Let $N_B = N_i \cap B$. By abusing notation, we denote by $\mathcal{P}_i \cap B$ all the pairs in B of \mathcal{P}_i . We will show that:

▷ **Claim 10.** There is a Steiner spanner S_B of weight at most $O(|N_B|L_i/\sqrt{\epsilon})$ such that for any pair of points $(x, y) \in \mathcal{P}_i \cap B$, it holds that:

$$d_{S_B \cup S_{i-1}}(x, y) \leq (1 + c\epsilon)\|x, y\|$$

for some big enough constant c .

Proof. We divide B into $O(1)$ horizontal (vertical) *bands* of length (width) $L_i/8$ so that for any two points $x, y \in \mathcal{P}_i \cap B$, $N_i(x)$ and $N_i(y)$ are in two non-adjacent horizontal bands and/or vertical bands (see Figure 2). These bands exist since

$$\|N_i(x), N_i(y)\| \geq \|x, y\| - 2\sqrt{\epsilon}L_i \geq L_i/2 - 2\sqrt{\epsilon}L_i \geq L_i/4$$

Now for each pair of non-adjacent horizontal bands H_a and H_b , $d(H_a, H_b) \geq \frac{L_i}{8}$. Draw a bisecting segment (touching two sides of B) between H_a and H_b and place $O(\frac{1}{\sqrt{\epsilon}})$ equally-spaced Steiner points, say R , on the bisecting line in a way that the distance between any two nearby Steiner points is $L_i\sqrt{\epsilon}$ (see Figure 2(a)). For each point $r \in R$, and each net point $p_i \in N_B \cap (H_a \cup H_b)$, we apply the construction of Corollary 7 to r , the set of endpoints of \mathcal{P}_i inclosed in circle $B(p_i, \sqrt{\epsilon}L_i)$ and S_{i-1} ; let $S_{a,b}(p_i)$ be the obtained geometric graph (see Figure 2(b)). Note that $d(r, B(p_i, \sqrt{\epsilon}L_i)) = \Omega(L_i)$. Thus, by Corollary 7, $w(S_{a,b}(p_i)) = O(L_i)$ and that:

$$d_{S_{i-1} \cup S_{a,b}(p_i)}(r, q) \leq (1 + 13\epsilon)\|r, q\| \quad (5)$$

for any $q \in B(p_i, \sqrt{\epsilon}L_i) \cap P$. Let $S_{a,b}(r) = \cup_{p_i \in N_B \cap (H_a \cup H_b)} S_{a,b}(p_i)$. It holds that

$$w(S_{a,b}(r)) \leq O(L_i)|N_B \cap (H_a \cup H_b)|.$$

Let $S_{a,b} = \cup_{r \in R} S_{a,b}(r)$. Then, we have:

$$w(S_{a,b}) \leq O(L_i \cdot |R| \cdot |N_B \cap (H_a \cup H_b)|) = O\left(\frac{L_i}{\sqrt{\epsilon}}|N_B \cap (H_a \cup H_b)|\right) \quad (6)$$

We apply the same construction for every pair of non-adjacent vertical bands. We then let S_B be the union of all $S_{a,b}$ for every pair of non-adjacent horizontal/vertical bands H_a, H_b . It holds that:

$$w(S_B) = O(L_i \cdot |R| \cdot |N_B|) = O\left(\frac{L_i|N_B|}{\sqrt{\epsilon}}\right) \quad (7)$$

since there are only $O(1)$ pairs of bands. To bound the stretch, let (x, y) be a pair in \mathcal{P}_i whose endpoints are in B . W.l.o.g, assume that H_a, H_b are two non-adjacent horizontal bands that contain $N_i(x)$ and $N_i(y)$, respectively. Let v be the intersection of segment xy and the bisecting line L of H_a, H_b (see Figure 2(b)). Let $r \in R$ be the closest Steiner point to v in L and z be the projection of r on xy . Observe that $\|z, x\|, \|z, y\| \geq \frac{L_i}{16} - \|z, v\| \geq L_i/16 - \sqrt{\epsilon}L_i \geq L_i/32$ when $\epsilon \ll 1$. We have:

$$\begin{aligned} \|r, x\| + \|r, y\| &= \sqrt{\|x, z\|^2 + \|r, z\|^2} + \sqrt{\|y, z\|^2 + \|r, z\|^2} \\ &\leq \|x, z\|\sqrt{1 + \frac{\epsilon L_i^2}{\|x, z\|^2}} + \|y, z\|\sqrt{1 + \frac{\epsilon L_i^2}{\|y, z\|^2}} \quad \text{since } \|r, z\| \leq \sqrt{\epsilon}L_i \\ &\leq \|x, z\|\sqrt{1 + 1024\epsilon} + \|y, z\|\sqrt{1 + 1024\epsilon} \quad \text{since } \|x, z\|, \|y, z\| \geq L_i/32 \\ &\leq \|x, z\|(1 + 512\epsilon) + \|y, z\|(1 + 512\epsilon) = (1 + 512\epsilon)\|x, y\| \end{aligned}$$

Thus, by Equation 5, we have:

$$d_{S_B \cup S_{i-1}}(x, y) \leq (1 + 13\epsilon)(\|r, x\| + \|r, y\|) = (1 + 13\epsilon)(1 + 512\epsilon)\|x, y\| = (1 + O(\epsilon))\|x, y\|$$

Thus, the stretch is $(1 + c\epsilon)$ for a sufficiently big constant c . ◁

Let $H_i = \cup_{B \in \mathcal{B}} S_B$. Since each net point belongs to at most 4 subsquares in \mathcal{B} by Claim 9, and $w(H_i) \leq O(\frac{|N_i|L_i}{\sqrt{\epsilon}})$ by Claim 10, it holds that:

$$w(H_i) = O\left(\frac{L_i}{\sqrt{\epsilon}} \frac{w(\text{MST})}{L_i \sqrt{\epsilon}}\right) = O\left(\frac{w(\text{MST})}{\epsilon}\right) \tag{8}$$

by Claim 8 as desired. ◀

4 A Linear Time Construction

In this section, we assume that $\Delta = O(n^c)$ for some constant c , and ϵ is a constant. We use the same model of computation used by Chan [20]: the real-RAM model with $\Theta(\log n)$ word size and floor function. We will use O_ϵ notation to hide a polynomial factor of $\frac{1}{\epsilon}$. Chan [20] showed that:

► **Theorem 11** (Step 4 in [20]). *Given a point set $P \in \mathbb{R}^d$ with spread $\Delta = O(n^c)$ for constant d and c , a $(1 + \epsilon)$ -spanner of P can be constructed in $O_\epsilon(P)$ time.*

We will use a construction of an r -net for a point set P for any r in time $O(n)$. Such a construction was implicit in the work of Har-Peled [48] which was made explicit by Har-Peled and Raichel (Lemma 2.3 [50]).

► **Lemma 12** (Lemma 2.3 and Corollary 2.4 [50]). *Given $r \geq 1$ and an n -point set P in \mathbb{R}^d , an r -net N of P can be constructed in $O(n)$ time. Furthermore, for each net point p , one can compute all the points covered by p in total $O(n)$ time.*

We first show that the spanner in Corollary 7 can be implemented in $O_\epsilon(|X|)$ time.

▷ **Claim 13.** The single-source spanner H in Corollary 7 can be found in $O_\epsilon(|X|)$ time.

Proof. First, we observe that the single source spanner in Lemma 6 can be constructed in time $O_\epsilon(|X|)$. This is because the grid W has size $O_\epsilon(1)$ and the SLT tree from p to (a set of $O(\frac{1}{\sqrt{\epsilon}})$ grid points on) L can be constructed in $O_\epsilon(1)$ time. The single-source spanner in Corollary 7 uses a constant number of constructions in Lemma 6. Thus, the total running time is $O_\epsilon(|X|)$. ◁

Proof of Theorem 2. Our implementation will follow the construction in Section 3 ; we will reuse notation in that section as well. Let K be a $(1 + \epsilon)$ -spanner H for P constructed in time $O(n)$ by Theorem 11. In our fast construction algorithm, instead of considering all pairs of points $\mathcal{P} = \binom{P}{2}$, we only consider the pairs corresponding to edges of H ; there are $O(n)$ such pairs. Our algorithm has four steps:

- **Step 1.** Partition pairs of endpoints in $E(H)$ into at most $O(\log(\Delta))$ sets $\mathcal{P}_1, \dots, \mathcal{P}_{\lceil \log \Delta \rceil}$ where \mathcal{P}_i is the set of pairs (x, y) such that $\|x, y\| \in [2^{i-1}, 2^i)$. Let $L_i = 2^i$. This step can be implemented in time $O(|E(H)|) = O_\epsilon(n)$. The following steps are applied to each $i \in [1, \lceil \log \Delta \rceil]$. Let P_i be the set of endpoints of \mathcal{P}_i . We observe that:

$$\sum_{i=1}^n |P_i| \leq 2 \sum_{i=1}^n |\mathcal{P}_i| = 2|E(H)| = O(n) \tag{9}$$

67:12 Light Euclidean Spanners with Steiner Points

- **Step 2.** Construct a $(\sqrt{\epsilon}L_i)$ -net N_i for P_i in $O(|P_i|)$ time using the algorithm in Lemma 12.
- **Step 3.** Compute a bounding square Q and divide it into (overlapping) subsquares of length $\Theta(L_i)$ each. Let \mathcal{B} be the set of subsquares that contain at least one point participating in the pairs \mathcal{P}_i . Since there are only $O(|P_i|)$ non-empty subsquares, \mathcal{B} can be computed in time $O(|P_i|)$ by iterating over each point and check (in $O(1)$ time) which subsquare the point falls into. Here we use the fact that each floor operation takes $O(1)$ time.
- **Step 4.** For each subsquare $B \in \mathcal{B}$, we divide it into $O(1)$ horizontal bands and vertical bands of length $L_i/8$. For each pair of non-adjacent (horizontal) bands H_a, H_b , construct a set of $O(\frac{1}{\sqrt{\epsilon}})$ Steiner points on the bisecting line between H_a and H_b as in Section 3. For each Steiner point r and each net point $p \in N_i \cap (H_a \cap H_b)$, we apply Corollary 7 to construct a single source spanner from r to a set of points $B(p, \sqrt{\epsilon}L_i) \cap P_i$; this step can be implemented in time $O_\epsilon(|B(p, \sqrt{\epsilon}L_i) \cap P_i|)$ by Claim 13. By Claim 9, the construction in this step can be implemented in $O_\epsilon(|P_i|)$ time. Our final spanner is the union of all single source spanners in all subsquares in \mathcal{B} .

The running time needed to implement Steps 2 to 4 is $O_\epsilon(\sum_{i=1} |P_i|) = O_\epsilon(n)$ by Equation 9. The same analysis in Section 3 gives $O(\frac{\log \Delta}{\epsilon})$ lightness. For stretch, we observe that the stretch of the spanner for each edge of H is $(1 + O(\epsilon))$. Thus, the stretch for every pair of points in P is $(1 + \epsilon)(1 + O(\epsilon)) = (1 + O(\epsilon))$. We can recover stretch $(1 + \epsilon')$ by setting $\epsilon' = \frac{\epsilon}{c}$ where c is the constant behind big- O . ◀

5 Steiner Spanners in High Dimension

In this section, we a light Steiner spanner for a point set $P \in \mathbb{R}^d$ with spread Δ as in Theorem 3. We rescale the metric so that every edge in $\binom{P}{2}$ has weight at least $\frac{1}{\epsilon}$. Let MST be the minimum spanning tree of P . We subdivide each MST edge of length > 1 , by placing Steiner points greedily, in a way that each new edge has length *at least* $1/2$ and *at most* 1 . Let K be the set of Steiner points. We observe that:

$$w(\text{MST}) = \Theta(|P| + |K|) \quad (10)$$

Let $\delta > 1$ be some parameter and $L_i = \frac{\delta}{\epsilon^i}$. Let $\mathcal{E}_\delta = \{E_1, \dots, E_m\}$ be the set of edges such that

$$E_i = \{e | e \in \binom{P}{2} \wedge w(e) \in (L_i/2, L_i]\} \quad (11)$$

where $m = \lceil \log_{\frac{1}{\epsilon}}(\Delta/(\epsilon\delta)) \rceil \leq \lceil \log_{\frac{1}{\epsilon}} \Delta \rceil + 1$. If an edge $e \in E_i$ for some $E_i \in \mathcal{E}_\delta$, we will abuse notation by saying that $e \in \mathcal{E}_\delta$. The main focus of this section is to show that:

► **Lemma 14.** *There is a Steiner spanner that preserves distances between the endpoints of edges in \mathcal{E}_δ with weight $\tilde{O}(\epsilon^{-(d+1)/2} + (\delta + \epsilon^{-2}) \log_{\frac{1}{\epsilon}} \Delta)w(\text{MST})$.*

We will show below that Lemma 14 implies Theorem 3.

Proof of Theorem 3. We assume that $\frac{1}{\epsilon}$ is a power of 2. We partition the interval $[1, \epsilon)$ into $J = \log_2(\frac{1}{\epsilon})$ intervals $[1, 2), \dots, [2^{J-1}, 2^J)$. For each fixed $j \in [1, J]$, let $\delta_j = 2^j$, and \mathcal{E}_{δ_j} be the set of edges with $\delta = \delta_j$ in the definition of \mathcal{E}_δ . Recall that we scale the metric so that every edge in $\binom{P}{2}$ has weight at least $\frac{1}{\epsilon}$. Thus, $\binom{P}{2} = \cup_{j=1}^J \mathcal{E}_{\delta_j}$

Observe that $\delta_j \leq \frac{1}{\epsilon}$ for all $j \in [1, J]$. By Lemma 14, there exists a Steiner spanner S_j with weight $\tilde{O}(\epsilon^{-(d+1)/2} + (\epsilon^{-2} \log_{\frac{1}{\epsilon}} \Delta)w(\text{MST}))$ preseving distances between endpoints of edges in \mathcal{E}_{δ_j} up to a $(1 + \epsilon)$ factor. Then $S = \cup_{j=1}^J S_j$ is a Steiner $(1 + \epsilon)$ -spanner with weight $\tilde{O}(\epsilon^{-(d+1)/2} + \epsilon^{-2} \log \Delta)w(\text{MST})$. ◀

We now focus on constructing a Steiner spanner in Lemma 14. We will use the following Steiner spanner construction as a black box.

► **Theorem 15** (Theorem 1.3 [59]). *For a given point set P , there is a Steiner $(1+\epsilon)$ -spanner, denoted by $\text{STP}(P)$, with $\tilde{O}(\epsilon^{-(d-1)/2})|P|$ edges that preserves pairwise distances of points in P up to a $(1+\epsilon)$ factor.*

We will rely on a *cover tree* to construct a Steiner spanner. Let c be a sufficiently big constant chosen later ($c = 20$).

► **Definition 16** (Cover tree). *A cover T for point set $P \cup K$ with $(m+1)$ levels has each node associated with a point of P such that (a) level-0 of T is the point set $P \cup K$, (b) level- i of T is associated with a $(c\epsilon L_i)$ -cover N_i of P and (c) $N_m \subseteq N_{m-1} \subseteq \dots \subseteq N_0$.*

A point p may appear in many levels of a cover tree T . To avoid confusion, we denote by (p, i) the copy of p at level i , and we still call (p, i) a *point* of $P \cup K$. For each point (p, i) , we denote by $\text{child}(p, i)$ and $\text{desc}(p, i)$ the set of children and descendants of (p, i) in T , respectively. Note that $\text{desc}(p, i)$ includes (p, i) .

We will construct the Steiner spanner level by level, starting from level 1. At every level, we will add a certain set of edges to E_{sp} . We then charge the weight of a subset of the edges to a subset of *uncharged points* of $P \cup K$; initially, every point of $P \cup K$ is *uncharged*. To decide which uncharged points we will charge to at level i , we consider a geometric graph H_i where $V(H_i) = N_i$ and there is an edge between two points $(p, i) \neq (q, i)$ (of weight $\|p, q\|$) in H_i if there exists at least one $e \in E_i$ between two descendants of (p, i) and (q, i) , respectively. We say a cover point (p, i) has *high degree* if its degree in H_i is at least $\frac{4c}{\epsilon}$. At level i , we only charge to uncharged points which are descendants of high degree cover points. The intuition is that high cover points have many descendants. This leads to a notion of a *charging cover tree*. We call a cover tree a *charging cover tree* for \mathcal{E}_δ if for all level $i \geq 1$,

- (1) Each point (p, i) has at least ϵL_i descendants that are uncharged at level less than i .
- (2) Up to $\epsilon L_i/2$ uncharged descendant of each high-degree cover point (p, i) can be charged at level i . No descendant of low-degree points is charged at level i .

We show how to construct a charging cover tree in Appendix 5.4. We now show that given a charging cover tree, we can construct a Steiner spanner with the lightness bound in Lemma 14. Let T be such a charging cover tree. We define a set of edges E_T as follows:

$$E_T = \{(p, q) \mid (p \in \text{child}(q, i) \vee q \in \text{child}(p, i)) \text{ for some } i\} \quad (12)$$

We abuse notation by denoting E_T the graph induced by the set of edges in E_T .

▷ **Claim 17.** $w(E_T) = O(\delta + \epsilon^{-1} \log_{\frac{1}{\epsilon}} \Delta)w(\text{MST})$ and for any p and every $x \neq y \in \text{desc}(p, i)$, $d_{E_T}(x, y) \leq 4c\epsilon L_i$.

Proof. Edges in E_T can be partitioned according to levels where an edge is at level i if it connects a point (p, i) and its parent. Observe that at level 0, the total edge weight is at most $c\delta$ times the number of points and hence, the total weight is $O(\delta|K \cup P|) = O(\delta w(\text{MST}))$ by Equation 10. At higher level, we observe that the total weight of edges of E_T at level $i \geq 1$ is at most $L_i|N_i|$, and that $N_i \leq \frac{|K \cup P|}{\epsilon L_i}$ since each point (p, i) has $|\text{desc}(p, i)| \geq \epsilon L_i$ by property (1) of the charging tree. Thus, the total weight of edges E_t at level at least 1 is at most:

$$\sum_{i=1}^m \frac{|K \cup P|}{\epsilon} = m \frac{|K \cup P|}{\epsilon} = O(\epsilon^{-1} \log_{\frac{1}{\epsilon}} \Delta |K \cup P|) = O(\epsilon^{-1} \log_{\frac{1}{\epsilon}} \Delta)w(\text{MST})$$

67:14 Light Euclidean Spanners with Steiner Points

This implies the weight bound of E_T . We now bound the distance between $x \neq y \in \text{desc}(p, i)$. Let $x = v_0, v_1, \dots, v_k = p$ be the (unique) path from x to p . By construction, $\|v_{i-1}, v_i\| \leq \frac{v_i v_{i+1}}{\epsilon}$ for $i \in [1, k-1]$. This implies:

$$d_{E_T}(x, p) = \|p, v_{k-1}\| \sum_{i=0}^{k-1} \epsilon^i \leq \frac{\|p, v_{k-1}\|}{1 - \epsilon} \leq 2\|p, v_{k-1}\| = 2c\epsilon L_i$$

when $\epsilon \leq \frac{1}{2}$. Similarly, $d_{E_T}(y, p) \leq 2c\epsilon L_i$ and hence $d_{E_T}(x, y) \leq 4c\epsilon L_i$. \triangleleft

Claim 17 implies that the descendants of any level i node in a charging cover tree form a subgraph of diameter at most $O(c\epsilon L_i)$. Let E_{sp} be the set of edges that will be our final spanner. Initially, $E_{sp} = E_T \cup \text{MST}$. We will abuse notation by denoting E_{sp} the graph induced by edge set E_{sp} .

5.1 Spanner construction at level i

Recall that H_i is a geometric graph where $V(H_i) = N_i$ and there is an edge between two points $(p, i) \neq (q, i)$ in H_i if there exists at least one $e \in E_i$ between two descendants of (p, i) and (q, i) , respectively. Recall that a high degree point (p, i) has at least $\frac{4c}{\epsilon}$ neighbors in H_i . We proceed in two steps.

- **Step 1.** For every low degree point (p, i) , we add all incident edges of (p, i) in H_i to E_{sp} .
- **Step 2.** Let Q be the set of high degree points in N_i . We add to E_{sp} the set of edges of $\text{STP}(Q)$. We take from each high degree cover point (p, i) exactly $\epsilon L_i/2$ uncharged descendants and let X be the set of these uncharged points. We charge the cost of $\text{STP}(Q)$ equally to all points in X and mark them *charged*. This charging is possible by property (2) of T .

5.2 Bounding the stretch

We will show that the stretch is $(1 + (48c + 1)\epsilon)$. We can recover stretch $(1 + \epsilon')$ by setting $\epsilon' = \frac{\epsilon}{48c+1}$.

Observe by construction that for every edge $e \in H_i$, the stretch of e in E_{sp} is at most $(1 + \epsilon)$. Recall that for every edge $(u, v) \in E_i$, there is an edge $(p, q) \in E(H_i)$ such that $u \in \text{desc}(p, i), v \in \text{desc}(q, i)$. By Claim 17, there is a path between u and v in E_{sp} of length at most $d_{E_{sp}}(p, q) + 8c\epsilon L_i$. By triangle inequality, $\|p, q\| - 2c\epsilon L_i \leq \|u, v\| \leq \|p, q\| + 2c\epsilon L_i$. Thus, we have:

$$\begin{aligned} \frac{d_{E_{sp}}(u, v)}{\|u, v\|} &\leq \frac{d_{E_{sp}}(p, q) + 8c\epsilon L_i}{\|p, q\| - 2c\epsilon L_i} = \frac{d_{E_{sp}}(p, q)}{\|p, q\|} + \frac{(d_{E_{sp}}(p, q) + 4\|p, q\|)2c\epsilon L_i}{\|p, q\|(\|p, q\| - 2c\epsilon L_i)} \\ &\leq \frac{d_{E_{sp}}(p, q)}{\|p, q\|} + \frac{12c\epsilon L_i}{\|p, q\| - 2c\epsilon L_i} \quad \text{since } d_{E_{sp}}(p, q) \leq 2\|p, q\| \\ &\leq \frac{d_{E_{sp}}(p, q)}{\|p, q\|} + 48c\epsilon \leq 1 + (48c + 1)\epsilon \end{aligned}$$

The penultimate inequality follows from the fact that

$$\|p, q\| - 2c\epsilon L_i \geq \|u, v\| - 4c\epsilon L_i \geq L_i/2 - 4c\epsilon L_i \geq L_i/4$$

when $\epsilon \ll \frac{1}{c}$.

5.3 Bounding $w(E_{sp})$

Observe that the total number of low degree points in N_i is at most $\frac{|K \cup P|}{\epsilon L_i}$ since each low degree point has at least ϵL_i (uncharged) descendants by property (1) of charging tree T . By triangle inequality, each edge of H_i has weight at most $L_i + 2c\epsilon L_i \leq 3L_i$ when $\epsilon < \frac{1}{c}$. Thus, the total weight of edges added to E_{sp} in Step 1 is bounded by:

$$\frac{3L_i \cdot 4c}{\epsilon} |\{p|p \in N_i \text{ and } p \text{ has low degree}\}| = O\left(\frac{L_i}{\epsilon} \frac{|K \cup P|}{\epsilon L_i}\right) = O\left(\frac{1}{\epsilon^2}\right)w(\text{MST})$$

by Equation 10. Thus, the total weight of the edges added to E_{sp} in Step 1 over m levels is $O(\epsilon^{-2} \log_{\frac{1}{\epsilon}} \Delta)w(\text{MST})$. Note here that $m \leq \lceil \log_{\frac{1}{\epsilon}} \Delta \rceil + 1$.

We now bound the total weight of the edges added to E_{sp} in Step 2 over m levels. Since each edge of H_i has weight at most $3L_i$, each edge of $\text{STP}(Q)$ has weight a most $(1 + \epsilon)3L_i \leq 6L_i$. By Theorem 15,

$$w(\text{STP}(Q)) = \tilde{O}(\epsilon^{-(d-1)/2})|Q|6L_i = \tilde{O}(\epsilon^{-(d-1)/2+o(1)})|Q|L_i$$

Thus, in Step 2, each uncharged point is charged at most:

$$\frac{\tilde{O}(\epsilon^{-(d-1)/2})|Q|L_i}{|Q|\epsilon L_i/2} = \tilde{O}(\epsilon^{-(d+1)/2}) \tag{13}$$

This implies the total weight of the edges added to E_{sp} in Step 2 over all levels is $\tilde{O}(\epsilon^{-(d+1)/2})(|P \cup K|) = \tilde{O}(\epsilon^{-(d+1)/2})w(\text{MST})$. Together with Claim 17, we conclude that:

$$w(E_{sp}) = \tilde{O}(\epsilon^{-(d+1)/2} + (\delta + \epsilon^{-2}) \log_{\frac{1}{\epsilon}} \Delta)w(\text{MST})$$

This completes the proof of Lemma 14.

5.4 Constructing a Charging Cover Tree

The main difficulty is to guarantee property (1) in constructing a charging cover tree; for property (2) at each level i , we simply charge to exactly $\epsilon L_i/2$ uncharged descendants of each high degree cover point.

A natural idea is to guarantee that each cover point has $\frac{1}{\epsilon}$ children. Then inductively, if each child of a cover point (p, i) has at least $\epsilon L_{i-1}/2$ uncharged descendants after the charging at level $i - 1$, we can hope that p has at least $\frac{1}{\epsilon}\epsilon L_{i-1} = \epsilon L_i$ uncharged descendants. There are two issues with this idea: (a) if at least one child, say $(q, i - 1)$, of (p, i) has high degree in the graph H_{i-1} , up to $\epsilon L_{i-1}/2$ uncharged points in $\text{desc}(q, i - 1)$ were charged at level $i - 1$ by property (2), and thus, $(q, i - 1)$ only contributes $\frac{\epsilon L_{i-1}}{2}$ uncharged descendants to (p, i) ; and (b) there may not be enough cover points at level $i - 1$ close to p as these points and their descendants must be within distance $c\epsilon L_i$ from p .

In our construction, we resolve both issues by picking a cover point in a way that the total number of uncharged descendants of its children is at least ϵL_i . We do so by having a more accurate way to track the number of uncharged descendants of a cover point, instead of simply relying on the lower bound ϵL_i of uncharged descendants. Specifically, denote by $D(X)$ the diameter of a point set X . We will construct a charging cover tree in a way that the following invariant is maintained at all levels.

Strong Charging Invariant: (SCI)

Each point (p, i) has at least $\max(\epsilon L_i, D(\text{desc}(p, i)) + 1)$ uncharged descendants (before the charging happened at level i).

Clearly, SCI implies property (1) of T . We begin by constructing the level-1 cover. Recall that MST edges have weight at most 1 and at least $\frac{1}{2}$, and that $\delta \geq 1$.

Level 1

We construct level-1 cover points N_1 by greedily breaking MST edges into subtrees of diameter at least δ and at most $3\delta + 2 \leq 5\delta$. Let X be such a subtree of MST with diameter d_X ; X will have at least δ points since MST edge has weight at most 1. We pick any point, say $p \in X$, to be a level-1 cover point, and make other points in X become p 's children; p will have at least δ children (uncharged at level 0). Recall that $\epsilon L_1 = \delta$. Since each MST edge has weight at most 1, the number of descendants of $(p, 1)$ is at least:

$$d_X + 1 \geq \max(\epsilon L_1, D(\text{desc}(p, 1)) + 1)$$

Thus, SCI holds for this level.

Level $i + 1$

Recall H_i is a graph with vertex set N_i . We construct a cover N_{i+1} in three steps A, B and C.

- **Step A.** For each high degree point p (with at least $\frac{4c}{\epsilon}$ unmarked neighbors in H_i), we pick p to N_{i+1} and make its unmarked neighbors become its children. We then mark p and all of its neighbors. For each remaining unmarked high degree point x in H_i , at least one of its neighbors, say q , must be marked before. We make x become a child of q 's parent.

The intuition of the construction in Step A is that $(p, i+1)$ picked at this step has at least $\frac{4c}{\epsilon}$ children. Since at least $(\epsilon L_i)/2$ descendants of each child of (p, i) remain uncharged after level i , the total number of uncharged descendants of $(p, i+1)$ is $\frac{4c}{\epsilon} \frac{\epsilon L_i}{2} = 2cL_i = 2c\epsilon L_{i+1} > \epsilon L_{i+1}$. Furthermore, since every high degree point is marked in this step, points in subsequent steps have low degree and hence no uncharged descendant of these points is charged at level i by property (2) of charging cover trees.

Let W be the set of remaining points in N_i . We construct a forest F from W as follows. The vertex set of F is W , and there is an edge between two vertices p, q of F if there is an MST edge, called the source of the edge, connecting a point in $\text{desc}(p)$ and a point in $\text{desc}(q)$. We set the weight of each edge in F to be the weight of the source edge. Note that F is not a geometric graph. We observe that:

► **Observation 18.** *For every connected component $C \in F$, there must be a point $p \in C$ and a point q marked in Step A such that there is an MST edge between a descendant of (p, i) and a descendant of (p, i) , except when there is no point marked in Step A (and F is a tree in this case).*

Proof. The observation follows from the fact that MST spans $P \cup K$. ◀

We define the weight on each vertex p of F to be $w(p) = D(\text{desc}(p, i))$, and the *vertex-weight* of a path P , denoted by $\text{vw}(P)$, of F to be the total weight of vertices on the path. We define the *absolute weight* of P , denoted by $\text{aw}(P)$, to be the total vertex and edge weight of P . Since each MST edge has weight at most 1 and each vertex has weight at least 1,

$$\text{aw}(P) \leq 2\text{vw}(P) \tag{14}$$

The *vertex-diameter* of a subtree C , denoted by $\text{VD}(C)$, of F is defined to be the vertex-weight of the path of maximum vertex-weight in the subtree. The *absolute diameter* of C , denoted by $\text{AD}(C)$, is defined similarly but w.r.t absolute weight.

- **Step B.** For each component C of F of vertex-diameter at least L_i , we greedily break C into sub-trees of vertex-diameter at least L_i and at most $3L_i$. For each subtree of C , choose an arbitrary point p to be a level- $(i + 1)$ cover point and make other points become p 's children.
- **Step C.** For each component C of F of vertex-diameter at most L_i , by Observation 18, there must be at least one MST connecting a point in $\text{desc}(u, i)$ for some $u \in C$ to a point in $\text{desc}(v, i)$ for some point v marked in Step 1. We make all points in C become children of v 's parent.

The following claim implies that N_{i+1} is a $(c\epsilon L_{i+1})$ -cover.

▷ **Claim 19.** For each point $(p, i + 1)$, $D(\text{desc}(p, i + 1)) \leq c\epsilon L_{i+1}$ for $c = 20$.

Proof. Note that for each point (q, i) , $1 \leq D(\text{desc}(q, i)) \leq 2c\epsilon L_i$ since every point in $\text{desc}(q, i)$ is within distance $c\epsilon L_i$ from q .

First, consider the case that $(p, i + 1)$ is chosen in Step B. Then p and its children belong to a subtree C of F of vertex-diameter at most $3L_i$. by Equation 14, $\text{AD}(C) \leq 2 \cdot 3L_i = 6L_i$. Thus, for a point $p \in N_{i+1}$ selected in Step B, $D(\text{desc}(p, i + 1)) \leq \text{AD}(C) \leq 6L_i$.

We now consider the case where $(p, i + 1)$ is chosen in Step A. (There is no cover point at level $(i + 1)$ selected in Step C.) Recall that each edge of H_i has length at most $L_i + 2c\epsilon L_i$ by the triangle inequality. Observe that after Step 1, for any $(q, i) \in \text{child}(p, i + 1)$, the hop distance in H_i between (p, i) and (q, i) is at most 2, hence $\|p, q\| \leq 2(L_i + 2c\epsilon L_i) = 2L_i + 4c\epsilon L_i$. That implies, after Step A,

$$\|p, x\| \leq (2L_i + 4c\epsilon L_i) + c\epsilon L_i = 2L_i + 5c\epsilon L_i \quad \text{for any } x \in \text{desc}(p, i + 1) \quad (15)$$

In Step C, we add more points belonging to subtrees of F to $\text{child}(p, i + 1)$. Let C be any of these subtrees. Since $\text{VD}(C) \leq L_i$, $\text{AD}(C) \leq 2L_i$. By construction, there exists a point $(v, i) \in C$ and a point $(u, i) \in \text{child}(p, i + 1)$ such that there is an MST edge e connecting a point in $\text{desc}(v, i)$ and a point in $\text{desc}(u, i)$. Thus, the augmentation in Step C increases the distance from p to the furthest point of $\text{desc}(p, i + 1)$ by at most $w(e) + \text{AD}(C) \leq 1 + 2L_i \leq 3L_i$. This implies:

$$D(\text{desc}(p, i + 1)) \leq 2(2L_i + 5c\epsilon L_i + 3L_i) \leq 20L_i$$

when $\epsilon < \frac{1}{c}$. ◁

To complete the proof of Lemma 14, it remains to show SCI for level $(i + 1)$.

▷ **Claim 20.** Each point $(p, i + 1)$ has at least $\max(\epsilon L_{i+1}, D(\text{desc}(p, i + 1)) + 1)$ uncharged descendants.

Proof. We first consider the case $(p, i + 1)$ is picked in Step B. Let X be the subtree of F that p belongs to. Let P be a path of X of maximum absolute weight. By definition on absolute weight, $\text{aw}(P) \geq D(\text{desc}(p, i + 1))$. Since MST has length at most 1, we have:

$$\sum_{q \in P} D(\text{desc}(q, i)) + |E(P)| \geq \text{aw}(P) \geq D(\text{desc}(p, i + 1))$$

By SCI for level i , we conclude that the number of uncharged descendants of $(p, i + 1)$ is at least:

$$\sum_{q \in P} (D(\text{desc}(q, i)) + 1) \geq \left(\sum_{q \in P} D(\text{desc}(q, i)) + |E(P)| \right) + 1 \geq D(\text{desc}(p, i + 1)) + 1$$

To show that X has at least ϵL_{i+1} uncharged descendants, we observe by construction that X has a path Q with $\text{vw}(Q) \geq L_i$. By definition of vertex-weight, $\text{vw}(Q) = \sum_{q \in Q} D(\text{desc}(q, i))$. Thus, the total number of uncharged descendants of all $q \in Q$ by SCI is at least:

$$\sum_{q \in Q} (D(\text{desc}(q, i)) + 1) > \sum_{q \in Q} D(\text{desc}(q, i)) \geq L_i = \epsilon L_{i+1}$$

Thus, $(p, i + 1)$ has at least $\max(\epsilon L_{i+1}, D(\text{desc}(p, i + 1)) + 1)$ uncharged descendants.

It remains to consider the case $(p, i + 1)$ is picked in Step A. By construction, $(p, i + 1)$ has at least $\frac{4c}{\epsilon}$ children, and each has at least $\epsilon L_{i-1}/2$ uncharged descendants by property (2) of a charging cover tree. Note that $D(\text{desc}(p, i + 1)) \leq c\epsilon L_{i+1} = cL_i$ by Claim 19. Thus, $(p, i + 1)$ has at least:

$$\frac{4c}{\epsilon} \frac{\epsilon L_i}{2} = 2cL_i \geq \max(\epsilon L_{i+1}, D(\text{desc}(p, i + 1))) + cL_i > \max(\epsilon L_{i+1}, D(\text{desc}(p, i + 1)) + 1)$$

uncharged descendants as desired. \triangleleft

References

- 1 M. A. Abam and S. Har-Peled. New constructions of sspds and their applications. In *Proceedings of the 26th Annual Symposium on Computational Geometry*, SoCG '10, page 192–200, 2010. doi:10.1145/1810959.1810993.
- 2 P. K. Agarwal, K. Fox, D. Panigrahi, K. R. Varadarajan, and A. Xiao. Faster algorithms for the geometric transportation problem. In *33rd International Symposium on Computational Geometry*, volume 77 of *SoCG '17*, pages 7:1–7:16, 2017. doi:10.4230/LIPIcs.SoCG.2017.7.
- 3 Stephen Alstrup, Søren Dahlgaard, Arnold Filtser, Morten Stöckel, and Christian Wulff-Nilsen. Constructing light spanners deterministically in near-linear time. In Michael A. Bender, Ola Svensson, and Grzegorz Herman, editors, *27th Annual European Symposium on Algorithms, ESA 2019, September 9–11, 2019, Munich/Garching, Germany*, volume 144 of *LIPIcs*, pages 4:1–4:15. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019.
- 4 I. Althöfer, G. Das, D. Dobkin, D. Joseph, and J. Soares. On sparse spanners of weighted graphs. *Discrete Computational Geometry*, 9(1):81–100, 1993. doi:10.1007/BF02189308.
- 5 B. Aronov, M. Berg, O. Cheong, J. Gudmundsson, H. Haverkort, and A. Vigneron. Sparse geometric graphs with small dilation. In *International Symposium on Algorithms and Computation*, ISAAC '05, pages 50–59, 2005. doi:10.1007/11602613_7.
- 6 S. Arya and M. H. M. Smid. Efficient construction of a bounded degree spanner with low weight. *Algorithmica*, 17(1):33–54, 1997.
- 7 Sunil Arya, David M. Mount, and Michiel H. M. Smid. Randomized and deterministic algorithms for geometric spanners of small diameter. In *35th Annual Symposium on Foundations of Computer Science, Santa Fe, New Mexico, USA, 20–22 November 1994*, pages 703–712. IEEE Computer Society, 1994.
- 8 Gali Bar-On and Paz Carmi. δ -greedy t-spanner. In Faith Ellen, Antonina Kolokolova, and Jörg-Rüdiger Sack, editors, *Algorithms and Data Structures - 15th International Symposium, WADS 2017, St. John's, NL, Canada, July 31 - August 2, 2017, Proceedings*, volume 10389 of *Lecture Notes in Computer Science*, pages 85–96. Springer, 2017.
- 9 J. Beardwood, J. H. Halton, and J. M. Hammersley. The shortest path through many points. *Mathematical Proceedings of the Cambridge Philosophical Society*, 55(4):299–327, 1959. doi:10.1017/s0305004100034095.
- 10 Marc Benkert, Alexander Wolff, Florian Widmann, and Takeshi Shirabe. The minimum manhattan network problem: Approximations and exact solutions. *Comput. Geom.*, 35(3):188–208, 2006.

- 11 G. Bodwin and V. V. Williams. Better distance preservers and additive spanners. In *Proceedings of the Twenty-seventh Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '16, pages 855–872, 2016.
- 12 G. Borradaile, H. Le, and C. Wulff-Nilsen. Minor-free graphs have light spanners. In *2017 IEEE 58th Annual Symposium on Foundations of Computer Science*, FOCS '17, pages 767–778, 2017. doi:10.1109/FOCS.2017.76.
- 13 G. Borradaile, H. Le, and C. Wulff-Nilsen. Greedy spanners are optimal in doubling metrics. In *Proceedings of the 30th Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '19, pages 2371–2379, 2019. doi:10.1137/1.9781611975482.145.
- 14 Prosenjit Bose, Paz Carmi, Mathieu Couture, Michiel H. M. Smid, and Daming Xu. On a family of strong geometric spanners that admit local routing strategies. In Frank K. H. A. Dehne, Jörg-Rüdiger Sack, and Norbert Zeh, editors, *Algorithms and Data Structures, 10th International Workshop, WADS 2007, Halifax, Canada, August 15-17, 2007, Proceedings*, volume 4619 of *Lecture Notes in Computer Science*, pages 300–311. Springer, 2007.
- 15 Prosenjit Bose, Luc Devroye, Maarten Löffler, Jack Snoeyink, and Vishal Verma. Almost all delaunay triangulations have stretch factor greater than $\pi/2$. *Comput. Geom.*, 44(2):121–127, 2011.
- 16 K. Buchin. Delaunay triangulations in linear time? (part I). arXiv preprint, 2008. arXiv:arXiv:0812.0387.
- 17 K. Buchin and W. Mulzer. Linear-time delaunay triangulations simplified. In *25th European Workshop on Computational Geometry*, EuroCG '09, pages 235–238, 2009.
- 18 M. Buneffeddou, J. Chuzhoy, P. Indyk, and A. Sidiropoulos. Low-distortion embeddings of general metrics into the line. In *Proceedings of the 37th Annual ACM Symposium on Theory of Computing*, STOC '05, page 225–233, 2005. doi:10.1145/1060590.1060624.
- 19 T. Carpenter, F. V. Fomin, D. Lokshtanov, S. Saurabh, and A. Sidiropoulos. Algorithms for low-distortion embeddings into arbitrary 1-dimensional spaces. In *34th International Symposium on Computational Geometry*, SoCG '2018, pages 21:1–21:14, 2018. doi:10.4230/LIPIcs.SoCG.2018.21.
- 20 T. Chan. Well-separated pair decomposition in linear time? *Information Processing Letters*, 107(5):138–141, 2008. doi:10.1016/j.ipl.2008.02.008.
- 21 T.-H. H. Chan, M. Li, L. Ning, and S. Solomon. New doubling spanners: Better and simpler. *SIAM Journal on Computing*, 44(1):37–53, 2015. doi:10.1137/130930984.
- 22 T.-H. Hubert Chan, Anupam Gupta, Bruce M. Maggs, and Shuheng Zhou. On hierarchical routing in doubling metrics. *ACM Trans. Algorithms*, 12(4):55:1–55:22, 2016. Preliminary version appeared in SODA 2005.
- 23 B. Chandra, G. Das, G. Narasimhan, and J. Soares. New sparseness results on graph spanners. In *Proceedings of the Eighth Annual Symposium on Computational Geometry*, 1992.
- 24 Barun Chandra. Constructing sparse spanners for most graphs in higher dimensions. *Inf. Process. Lett.*, 51(6):289–294, 1994.
- 25 S. Chechik and C. Wulff-Nilsen. Near-optimal light spanners. In *Proceedings of the 27th Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA'16, pages 883–892, 2016.
- 26 Shiri Chechik and Christian Wulff-Nilsen. Near-optimal light spanners. *ACM Trans. Algorithms*, 14(3):33:1–33:15, 2018. preliminary version published in SODA 2016. doi:10.1145/3199607.
- 27 L. P. Chew. There is a planar graph almost as good as the complete graph. In *Proceedings of the Second Annual Symposium on Computational Geometry*, SCG '86, pages 169–177, 1986.
- 28 L. P. Chew. There are planar graphs almost as good as the complete graph. *Journal of Computer and System Sciences*, 39(2):205–219, 1989.
- 29 K. Clarkson. Approximation algorithms for shortest path motion planning. In *Proceedings of the Nineteenth Annual ACM Symposium on Theory of Computing*, STOC '87, pages 56–65, 1987.

- 30 V. Cohen-Addad and C. Mathieu. Effectiveness of local search for geometric optimization. In *31st International Symposium on Computational Geometry*, volume 35 of *SoCG '2015*, pages 329–343, 2015. doi:10.4230/LIPIcs.SOCG.2015.329.
- 31 G. Das, P. Heffernan, and G. Narasimhan. Optimally sparse spanners in 3-dimensional euclidean space. In *Proceedings of the 9th Annual Symposium on Computational Geometry*, SCG '93, pages 53–62, 1993.
- 32 G. Das, G. Narasimhan, and J. Salowe. A new way to weigh malnourished euclidean graphs. In *Proceedings of the 6th Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '95, pages 215–222, 1995.
- 33 Michael Elkin, Ofer Neiman, and Shay Solomon. Light spanners. *SIAM J. Discret. Math.*, 29(3):1312–1321, 2015. preliminary version published in ICALP 2014.
- 34 Michael Elkin and Shay Solomon. Narrow-shallow-low-light trees with and without steiner points. *SIAM J. Discret. Math.*, 25(1):181–210, 2011. preliminary version published in ESA 2009.
- 35 Michael Elkin and Shay Solomon. Steiner shallow-light trees are exponentially lighter than spanning ones. *SIAM J. Comput.*, 44(4):996–1025, 2015. preliminary version published in FOCS 2011. doi:10.1137/13094791X.
- 36 J. Erickson. Dense point sets have sparse delaunay triangulations or “... but not too nasty”. *Discrete & Computational Geometry*, 33(1):83–115, 2004. doi:10.1007/s00454-004-1089-3.
- 37 Mohammad Farshi and Joachim Gudmundsson. Experimental study of geometric t -spanners. In Gerth Stølting Brodal and Stefano Leonardi, editors, *Algorithms - ESA 2005, 13th Annual European Symposium, Palma de Mallorca, Spain, October 3-6, 2005, Proceedings*, volume 3669 of *Lecture Notes in Computer Science*, pages 556–567. Springer, 2005.
- 38 Mohammad Farshi and Joachim Gudmundsson. Experimental study of geometric t -spanners: A running time comparison. In Camil Demetrescu, editor, *Experimental Algorithms, 6th International Workshop, WEA 2007, Rome, Italy, June 6-8, 2007, Proceedings*, volume 4525 of *Lecture Notes in Computer Science*, pages 270–284. Springer, 2007.
- 39 Arnold Filtser and Ofer Neiman. Light spanners for high dimensional norms via stochastic decompositions. In *26th Annual European Symposium on Algorithms, ESA 2018, August 20-22, 2018, Helsinki, Finland*, pages 29:1–29:15, 2018. doi:10.4230/LIPIcs.ESA.2018.29.
- 40 J. Gao, L. J. Guibas, and A. Nguyen. Deformable spanners and applications. *Computational Geometry*, 35(1):2–19, 2006. doi:10.1016/j.comgeo.2005.10.001.
- 41 Lee-Ad Gottlieb. A light metric spanner. In *IEEE 56th Annual Symposium on Foundations of Computer Science, FOCS 2015, Berkeley, CA, USA, 17-20 October, 2015*, pages 759–772, 2015. doi:10.1109/FOCS.2015.52.
- 42 J. Gudmundsson, C. Levcopoulos, G. Narasimhan, and M. H. M. Smid. Approximate distance oracles for geometric graphs. In *Proc. of 13th SODA*, pages 828–837, 2002.
- 43 J. Gudmundsson, C. Levcopoulos, G. Narasimhan, and M. H. M. Smid. Approximate distance oracles for geometric spanners. *ACM Transactions on Algorithms*, 4(1), 2008.
- 44 J. Gudmundsson, G. Narasimhan, and M. H. M. Smid. Fast pruning of geometric spanners. In *Proc. of 22nd STACS*, pages 508–520, 2005.
- 45 Joachim Gudmundsson, Christos Levcopoulos, Giri Narasimhan, and Michiel H. M. Smid. Approximate distance oracles revisited. In *Proc. of 13th ISAAC*, pages 357–368, 2002.
- 46 Joachim Gudmundsson, Christos Levcopoulos, Giri Narasimhan, and Michiel H. M. Smid. Approximate distance oracles for geometric spanners. *ACM Trans. Algorithms*, 4(1):10:1–10:34, 2008.
- 47 Joachim Gudmundsson, Giri Narasimhan, and Michiel H. M. Smid. Fast pruning of geometric spanners. In Volker Diekert and Bruno Durand, editors, *STACS 2005, 22nd Annual Symposium on Theoretical Aspects of Computer Science, Stuttgart, Germany, February 24-26, 2005, Proceedings*, volume 3404 of *Lecture Notes in Computer Science*, pages 508–520. Springer, 2005.

- 48 S. Har-Peled. Clustering motion. *Discrete and Computational Geometry*, 31(4):545–565, 2004. doi:10.1007/s00454-004-2822-7.
- 49 S. Har-peled. *Geometric Approximation Algorithms*. American Mathematical Society, 2011. doi:10.5555/2031416.
- 50 S. Har-Peled and B. Raichel. Net and prune: A linear time algorithm for euclidean distance problems. *Journal of the ACM*, 62(6):1–35, 2015. doi:10.1145/2831230.
- 51 S. Har-Peled and B. Sadri. How fast is the k -means method? In *Proceedings of the 16th Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '05, page 877–885, 2006.
- 52 Y. Hassin and D. Peleg. Sparse communication networks and efficient routing in the plane. In *Proc. of 19th PODC*, pages 41–50, 2000.
- 53 R. M. Karp. Probabilistic analysis of partitioning algorithms for the traveling-salesman problem in the plane. *Mathematics of Operations Research*, 2(3):209–224, 1977. doi:10.1287/moor.2.3.209.
- 54 J. M. Keil. Approximating the complete euclidean graph. In *Proceedings of the first Scandinavian Workshop on Algorithm Theory*, SWAT '88, pages 208–213, 1988.
- 55 J. M. Keil and C. A. Gutwin. Classes of graphs which approximate the complete Euclidean graph. *Discrete and Computational Geometry*, 7(1):13–28, 1992.
- 56 Michael Kerber and Arnur Nigmatov. Metric spaces with expensive distances. *CoRR*, abs/1901.08805, 2019. arXiv:1901.08805.
- 57 P. N. Klein. Subset spanner for planar graphs, with application to subset TSP. In *Proceedings of the 38th Annual ACM Symposium on Theory of Computing*, STOC '06, pages 749–756, 2006. doi:10.1145/1132516.1132620.
- 58 Hung Le. A PTAS for subset TSP in minor-free graphs. In *Proceedings of the 2020 ACM-SIAM Symposium on Discrete Algorithms, SODA 2020, Salt Lake City, UT, USA, January 5-8, 2020*, pages 2279–2298, 2020. Full version: arxiv:1804.01588. doi:10.1137/1.9781611975994.140.
- 59 Hung Le and Shay Solomon. Truly optimal euclidean spanners. In *60th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2019, Baltimore, Maryland, USA, November 9-12, 2019*, pages 1078–1100, 2019. Full version at arXiv:1904.12042. doi:10.1109/FOCS.2019.00069.
- 60 Y. Mansour and D. Peleg. An approximation algorithm for min-cost network design. *DIMACS Series in Discr. Math and TCS*, 53:97–106, 2000.
- 61 G. Narasimhan and M. Smid. *Geometric Spanner Networks*. Cambridge University Press, 2007.
- 62 A. Nayyeri and B. Raichel. Reality distortion: Exact and approximate algorithms for embedding into the line. In *Proceedings of the 56th Annual Symposium on Foundations of Computer Science*, FOCS '15, page 729–747, 2015. doi:10.1109/FOCS.2015.50.
- 63 A. Nayyeri and B. Raichel. A treehouse with custom windows: Minimum distortion embeddings into bounded treewidth graphs. In *Proceedings of the 28th Annual ACM-SIAM Symposium on Discrete Algorithm*, SODA '17, page 724–736, 2017. doi:10.1137/1.9781611974782.46.
- 64 A. Nayyeri and B. Raichel. Viewing the rings of a tree: Minimum distortion embeddings into trees. In *Proceedings of the 30th Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '19, pages 2380–2399, 2019. doi:10.1137/1.9781611975482.146.
- 65 S. B. Rao and W. D. Smith. Approximating geometrical graphs via “spanners” and “ban-yans”. In *Proceedings of the 30th Annual ACM Symposium on Theory of Computing*, STOC '98, pages 540–550, 1998. Full version at <http://graphics.stanford.edu/courses/cs468-06-winter/Papers/rs-tsp.pdf>. doi:10.1145/276698.276868.
- 66 J. Ruppert and R. Seidel. Approximating the d -dimensional complete Euclidean graph. In *Proceedings of the 3rd Canadian Conference on Computational Geometry*, CCCG '91, page 207–210, 1991.
- 67 Jeffrey S. Salowe. On euclidean spanner graphs with small degree. In *Proceedings of the Eighth Annual Symposium on Computational Geometry, Berlin, Germany, June 10-12, 1992*, pages 186–191, 1992.

67:22 Light Euclidean Spanners with Steiner Points

- 68 S. Solomon. Euclidean Steiner shallow-light trees. *Journal of Computational Geometry*, 6(2), 2015. preliminary version published in SoCG 2014. doi:10.20382/JOCG.V6I2A7.
- 69 S. Solomon and M. Elkin. Balancing degree, diameter and weight in euclidean spanners. In *Proceedings of the 18th Annual European Symposium on Algorithms, ESA '10*, pages 48–59, 2010. doi:10.1007/978-3-642-15775-2_5.
- 70 A. C. Yao. On constructing minimum spanning trees in k -dimensional spaces and related problems. *SIAM Journal on Computing*, 11(4):721–736, 1982.