


Light Euclidean Steiner Spanners in the Plane

Sujoy Bhore ✉ 

Université Libre de Bruxelles, Belgium

Csaba D. Tóth ✉ 

California State University Northridge, Los Angeles, CA, USA

Tufts University, Medford, MA, USA

Abstract

Lightness is a fundamental parameter for Euclidean spanners; it is the ratio of the spanner weight to the weight of the minimum spanning tree of a finite set of points in \mathbb{R}^d . In a recent breakthrough, Le and Solomon (2019) established the precise dependencies on $\varepsilon > 0$ and $d \in \mathbb{N}$ of the minimum lightness of a $(1 + \varepsilon)$ -spanner, and observed that additional Steiner points can substantially improve the lightness. Le and Solomon (2020) constructed Steiner $(1 + \varepsilon)$ -spanners of lightness $O(\varepsilon^{-1} \log \Delta)$ in the plane, where $\Delta \geq \Omega(\sqrt{n})$ is the *spread* of the point set, defined as the ratio between the maximum and minimum distance between a pair of points. They also constructed spanners of lightness $\tilde{O}(\varepsilon^{-(d+1)/2})$ in dimensions $d \geq 3$. Recently, Bhore and Tóth (2020) established a lower bound of $\Omega(\varepsilon^{-d/2})$ for the lightness of Steiner $(1 + \varepsilon)$ -spanners in \mathbb{R}^d , for $d \geq 2$. The central open problem in this area is to close the gap between the lower and upper bounds in all dimensions $d \geq 2$.

In this work, we show that for every finite set of points in the plane and every $\varepsilon > 0$, there exists a Euclidean Steiner $(1 + \varepsilon)$ -spanner of lightness $O(\varepsilon^{-1})$; this matches the lower bound for $d = 2$. We generalize the notion of shallow light trees, which may be of independent interest, and use directional spanners and a modified window partitioning scheme to achieve a tight weight analysis.

2012 ACM Subject Classification Mathematics of computing \rightarrow Approximation algorithms; Mathematics of computing \rightarrow Paths and connectivity problems; Theory of computation \rightarrow Computational geometry

Keywords and phrases Geometric spanner, lightness, minimum weight

Digital Object Identifier 10.4230/LIPIcs.SoCG.2021.15

Related Version *Full Version*: <https://arxiv.org/abs/2012.02216>

Funding *Sujoy Bhore*: Research on this paper was supported by the Fonds de la Recherche Scientifique-FNRS under Grant no MISU F 6001.

Csaba D. Tóth: Research on this paper was partially supported by the NSF award DMS-1800734.

1 Introduction

Given an edge-weighted graph G , a spanner is a subgraph H of G that preserves the length of the shortest paths in G up to some amount of multiplicative or additive distortion. Formally, a subgraph H of a given edge-weighted graph G is a t -spanner, for some $t \geq 1$, if for every $pq \in \binom{V(G)}{2}$ we have $d_H(p, q) \leq t \cdot d_G(p, q)$, where $d_G(p, q)$ denotes the length of the shortest path in G . The parameter t is called the *stretch factor* of the spanner. Graph spanners were introduced by Peleg and Schäffer [40], and since then it has turned out to be a fundamental graph structure with numerous applications in the field of distributed systems and communication, distributed queuing protocol, compact routing schemes, etc.; see [19, 29, 41, 42]. For edge-weighted graphs, a natural parameter is the *lightness* of a spanner, that is associated with the total weight of the spanner. The *lightness* of a spanner H of an input graph G , is the ratio $w(H)/w(\text{MST})$ between the total weight of H and the weight of a minimum spanning tree (MST) of G . Note that, since a spanner H is a connected graph, the trivial lower bound for lightness is 1.



© Sujoy Bhore and Csaba D. Tóth;

licensed under Creative Commons License CC-BY 4.0


37th International Symposium on Computational Geometry (SoCG 2021).

Editors: Kevin Buchin and Éric Colin de Verdière; Article No. 15; pp. 15:1–15:17

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany



In geometric settings, a t -spanner for a finite set S of points in \mathbb{R}^d is a subgraph of the underlying complete graph $G = (S, \binom{S}{2})$, that preserves the pairwise Euclidean distances between points in S to within a factor of t , that is the *stretch factor*. The edge weights of G are the Euclidean distances between the vertices. Chew [14, 15] initiated the study of Euclidean spanners in 1986, and showed that for a set of n points in \mathbb{R}^2 , there exists a spanner with $O(n)$ edges and constant stretch factor. Since then a large body of research has been devoted to Euclidean spanners due to its many applications across domains, such as, topology control in wireless networks [45], efficient regression in metric spaces [26], approximate distance oracles [28], and others. Moreover, Rao and Smith [43] showed the relevance of Euclidean spanners in the context of other fundamental geometric NP-hard problems, e.g., Euclidean traveling salesman problem and Euclidean minimum Steiner tree problem. Many different spanner construction approaches have been developed for Euclidean spanners over the years, that each found further applications in geometric optimization, such as spanners based on well-separated pair decomposition (WSPD) [11, 27], skip-lists [3], path-greedy and gap-greedy approaches [1, 4], locality-sensitive orderings [12], and more. We refer to the book by Narasimhan and Smid [39] and the survey of Bose and Smid [10] for a summary of results and techniques on Euclidean spanners up to 2013.

Lightness and sparsity are two natural parameters for Euclidean spanners. For a set S of points in \mathbb{R}^d , the lightness is the ratio of the spanner weight (i.e., the sum of all edge weights) to the weight of the Euclidean minimum spanning tree $\text{MST}(S)$. Then, the *sparsity* of a spanner on S is the ratio of its size to the size of a spanning tree. Classical results show that when the dimension $d \in \mathbb{N}$ and $\varepsilon > 0$ are constant, every set S of n points in d -space admits an $(1 + \varepsilon)$ -spanners with $O(n)$ edges and weight proportional to that of the Euclidean MST of S . We refer to a series of spanners constructions for a comprehensible understanding of sparse spanners [15, 16, 30, 31, 44, 49].

Of particular interest, we elaborate on the *lightness* aspect of Euclidean spanners. Das et al. [17] showed that the *greedy-spanner* (cf. [1]) has constant lightness in \mathbb{R}^3 . This was generalized later to \mathbb{R}^d , for all $d \in \mathbb{N}$, by Das et al. [18]. However the dependencies on ε and d have not been addressed. Rao and Smith [43] showed that the greedy spanner has lightness $\varepsilon^{-O(d)}$ in \mathbb{R}^d for every constant d , and asked what is the best possible constant in the exponent. A complete proof for the existence of a $(1 + \varepsilon)$ -spanner with lightness $O(\varepsilon^{-2d})$ is in the book on geometric spanners [39]. Gao et al. [24] considered the spanners in doubling metrics, and showed that every finite set of n points in doubling dimension d has a spanner of sparsity $\varepsilon^{-O(d)}$. In 2015, Gottlieb [25] showed that a metric of doubling dimension d has a spanner of lightness $(d/\varepsilon)^{O(d)}$, which improved the $O(\log n)$ lightness bound of Smid [46]. Recently, Borradaile et al. [9] showed that the greedy $(1 + \varepsilon)$ -spanner of a finite metric space of doubling dimension d has lightness $\varepsilon^{-O(d)}$. In [33], Le and Solomon established the precise dependencies of ε in the *lightness* and *sparsity* bounds of Euclidean $(1 + \varepsilon)$ -spanners. They constructed, for every $\varepsilon > 0$ and constant $d \in \mathbb{N}$, a set S of n points in \mathbb{R}^d for which any $(1 + \varepsilon)$ -spanner must have lightness $\Omega(\varepsilon^{-d})$ and sparsity $\Omega(\varepsilon^{-d+1})$, whenever $\varepsilon = \Omega(n^{-1/(d-1)})$. Moreover, they showed that the greedy $(1 + \varepsilon)$ -spanner in \mathbb{R}^d has lightness $O(\varepsilon^{-d} \log \varepsilon^{-1})$.

Steiner Spanners. *Steiner points* are additional vertices in a network that are not part of the input, and a t -spanner must achieve stretch factor t only between pairs of the input points in S . Le and Solomon [33] observed that it is possible to use Steiner points to bypass the lower bounds and substantially improve the bounds for *lightness* and *sparsity* of Euclidean $(1 + \varepsilon)$ -spanners. For minimum sparsity, they gave an upper bound of $O(\varepsilon^{(1-d)/2})$ for d -space

and a lower bound of $\Omega(\varepsilon^{-1/2}/\log \varepsilon^{-1})$. For minimum lightness, they gave a lower bound of $\Omega(\varepsilon^{-1}/\log \varepsilon^{-1})$, for points in the plane ($d = 2$) [33]. In a subsequent work [34], they have constructed Steiner $(1 + \varepsilon)$ -spanners of lightness $O(\varepsilon^{-1} \log \Delta)$ in the plane, where Δ is the *spread* of the point set, defined as the ratio between the maximum and minimum distance between a pair of points. In particular, $\log \Delta \in \Omega(\log n)$ in doubling metrics.

Recently, Bhore and Tóth [7] established a lower bound of $\Omega(\varepsilon^{-d/2})$ for the lightness of Steiner $(1 + \varepsilon)$ -spanners in Euclidean d -space for all $d \geq 2$. Moreover, for points in the plane, they established an upper bound of $O(\varepsilon^{-1} \log n)$. In [35], Le and Solomon constructed spanners of lightness $\tilde{O}(\varepsilon^{-(d+1)/2})$ in dimensions $d \geq 3$, nearly matching the lower bound $\Omega(\varepsilon^{-d/2})$, for $d \geq 3$. The central open problem in this area is to close the gap between the lower and upper bounds of lightness, in all dimensions $d \geq 2$.

► **Question 1.** *Do there exist Euclidean Steiner $(1 + \varepsilon)$ -spanners for a finite set of points in \mathbb{R}^d , of lightness $O(\varepsilon^{-d/2})$, for any $d \geq 2$?*

Bounding the *lightness* of Euclidean spanners is often harder than bounding the *sparsity*, as also noted by Le and Solomon [34]. Several works portrayed the difficulties of constructing light spanners in Euclidean spaces, doubling metrics, as well as on other weighted graphs; see [1, 9, 13, 22, 25, 33, 46, 18, 43]. A delicate aspect of the problem is to find suitable locations for *Steiner points*. Recent results on Steiner spanners [7, 33, 34, 35] suggest that highly nontrivial insights are required to argue the upper bounds for Steiner spanners, and they tend to be even more intricate than their non-Steiner counterpart.

Related Previous Work. Steiner points were used in several occasions to improve the overall weight of a network. Previously, Elkin and Solomon [23] and Solomon [47] showed that Steiner points can improve the weight of the network in the single-source setting. In particular, they introduced the so-called *shallow-light trees* (SLT), that is a single-source spanning tree that concurrently approximates a shortest-path tree (between the source and all other points) and a minimum spanning tree (for the total weight). They proved that Steiner points help to obtain exponential improvement on the lightness of SLTs in a general metric space [23], and quadratic improvement on the lightness in Euclidean spaces [47].

Our Contribution. In this work, we show that for every finite set of points in the plane and every $\varepsilon > 0$, there exists a Euclidean Steiner $(1 + \varepsilon)$ -spanner of lightness $O(\varepsilon^{-1})$ (Theorem 2). This matches the lower bound for $d = 2$, and thereby closes the gap between lower and upper bounds of lightness for Euclidean $(1 + \varepsilon)$ -spanners in \mathbb{R}^2 .

On the one hand, without Steiner points, the greedy spanner in Euclidean plane has lightness $\tilde{O}(\varepsilon^{-2})$, which is the best possible up to lower-order terms [33]. On the other hand, with Steiner points, recent constructions achieved linear dependence on ε^{-1} , while losing the independence from n ; see [7, 34]. Our result is the first that constructs Steiner spanners with sub-quadratic dependence on ε^{-1} without any dependence on n or any assumption on the point set, in fact our result achieves the optimal dependence on ε .

► **Theorem 2.** *For every finite point sets $S \subset \mathbb{R}^2$ and $\varepsilon > 0$, there exists a Euclidean Steiner $(1 + \varepsilon)$ -spanner of weight $O(\frac{1}{\varepsilon} \|MST(S)\|)$.*

Outline. We review previous results on angle-bounded paths, SLTs, and window partitions that we use in our construction (Section 2). The tight bound in Theorem 2 relies on three new ideas, which may be of independent interest: First, we generalize Solomon's SLTs to points on a staircase path (Section 3). Second, we reduce the proof of Theorem 2 to the

construction of “directional” spanners, in each of $\Theta(\varepsilon^{-1/2})$ directions, where it is enough to establish the stretch factor $1 + \varepsilon$ for point pairs $s, t \in S$ where $\text{dir}(st)$ is in an interval of size $\sqrt{\varepsilon}$ (Section 4). Combining the first two ideas, we show how to construct light directional spanners for points on a staircase path (Section 5). In each direction, we start with a rectilinear MST of S , and augment it into a directional spanner. We modify the classical window partition of a rectilinear polygon into histograms by replacing vertical edges with angle-bounded paths; this is the final piece of the puzzle. Near-vertical paths (with slopes $\pm \varepsilon^{-1/2}$) allow sufficient flexibility to reduce the weight of a histogram subdivision, and we can construct directional $(1 + \varepsilon)$ -spanners for each face of such a subdivision (Section 6).


2 Preliminaries

The *direction* of a line segment ab in the plane, denoted $\text{dir}(ab)$, is the minimum counterclockwise angle $\alpha \in [0, \pi)$ that rotates the x -axis to be parallel to ab . The set of possible directions $[0, \pi)$ is homeomorphic to the unit circle \mathbb{S}^1 , and an interval (α, β) of directions corresponds to the counterclockwise arc of \mathbb{S}^1 from $\alpha \pmod{\pi}$ to $\beta \pmod{\pi}$.

Angle-Bounded Paths. For $\delta \in (0, \pi/2]$, a polygonal path (v_0, \dots, v_m) is $(\theta \pm \delta)$ -*angle-bounded* if the direction of every segment $v_{i-1}v_i$ is in the interval $[\theta - \delta, \theta + \delta]$; see Fig. 1(a). Borradaile and Eppstein [8, Lemma 5] observed that the weight of a $(\theta \pm \delta)$ -angle-bounded st -path is at most $(1 + O(\delta^2))\|st\|$. We prove this observation in a more precise form. The quadratic growth rate in δ is due to the Taylor estimate $\sec(x) = \frac{1}{\cos(x)} \leq 1 + x^2$ for $x \leq \frac{\pi}{4}$.

► **Lemma 3.** *Let $a, b \in \mathbb{R}^2$ and let $P = v_0v_1 \dots v_m$ be an ab -path such that P is monotonic in direction \vec{ab} and $|\text{dir}(v_{i-1}v_i) - \text{dir}(ab)| \leq \delta \leq \frac{\pi}{4}$, for $i = 1, \dots, m$. Then $\|P\| \leq (1 + \delta^2)\|ab\|$.*

Proof. For $i = 0, \dots, m$, let u_i be the orthogonal projection of v_i to the line ab , and let $\alpha_i = \text{dir}(v_{i-1}v_i) - \text{dir}(ab)$. Then $\|ab\| = \sum_{i=1}^m \|u_{i-1}u_i\| = \sum_{i=1}^m \|v_{i-1}v_i\| \sec \alpha_i \leq \|P\| \sec \delta \leq (1 + \delta^2)\|P\|$, as claimed. ◀



■ **Figure 1** (a) A $(0 \pm \delta)$ -angle-bounded path. (b) A shallow-light tree between a source s and a horizontal line segment L . (c)–(d) An x - and a y -monotone histogram.

Shallow-Light Trees. Shallow-light trees (SLT) were introduced by Awerbuch et al. [5] and Khuller et al. [32]. Given a source s and a point set S in a metric space, an (α, β) -SLT is a Steiner tree rooted at s that contains a path of weight at most $\alpha \|ab\|$ between the *source* s and any point $t \in S$, and has weight at most $\beta \|MST(S)\|$. We build on the following basic variant of SLT between a source s and a set S of collinear points in the plane; see Fig. 1(b).

► **Lemma 4** (Solomon [47, Section 2.1]). *Let $0 < \varepsilon < 1$, let $s = (0, \varepsilon^{-1/2})$ be a point on the y -axis, and let S be a set of points in the line segment $L = [-\frac{1}{2}, \frac{1}{2}] \times \{0\}$ in the x -axis. Then there exists a geometric graph of weight $O(\varepsilon^{-1/2})$ that contains, for every point $t \in L$, an st -path P_{st} with $\|P_{st}\| \leq (1 + \varepsilon)\|st\|$.*

We note that the weight analysis of the st -path P_{st} in an SLT does not use angle-boundedness. In particular, an SLT may contain short edges of arbitrary directions close to t , but the directions of long edges are close to vertical. In Section 3 below, we generalize the shallow-light trees to obtain $(1 + \varepsilon)$ -spanners between points on two staircase paths.

Stretch Factor of $1 + \varepsilon$ Versus $1 + O(\varepsilon)$. In the geometric spanners we construct, an st -path may comprise $O(1)$ subpaths, each of which is angle-bounded or contained in an SLT. For the ease of presentation, we typically establish a stretch factor of $1 + O(\varepsilon)$ in our proofs. It is understood that $1 + \varepsilon$ can be achieved by a suitable scaling of the constant coefficients.

Histogram Decomposition. A path in the plane is x -monotone (resp., y -monotone) if its intersection with every vertical (resp., horizontal) line is connected. A *histogram* is a rectilinear simple polygon bounded by an axis-parallel line segment and an x - or y -monotone path; see Fig. 1(c–d). It is well known that every rectilinear simple polygon P can be subdivided into histograms (faces) such that every axis-parallel line segment in P intersects (*stabs*) at most three histograms [21, 36]; such a subdivision is also called a *window partition* [37, 48] of P , and can be computed in $O(n \log n)$ time if P has n vertices. The stabbing property implies that the total perimeter of the histograms in such a subdivision is $O(\text{per}(P))$.

Dumitrescu and Tóth [20] showed that for a finite point set $S \subset \mathbb{R}^2$, one can refine the window partition, while increasing the weight by a constant factor, to construct a graph with constant geometric dilation. The *geometric dilation* of a geometric graph G is $\sup_{a,b \in G} d_G(a,b)/\|ab\|$, where $d_G(a,b)$ denotes the Euclidean length of a shortest path in G , and the supremum is taken over all point pairs $\{a,b\}$ at vertices and along edges of G . We follow a similar approach here, but we construct a subdivision of “modified” histograms (defined in Section 6), where the vertical edges are replaced by angle-bounded paths.

3 Generalized Shallow Light Trees

In Section 3.1, we generalize Lemma 4, and construct shallow-light trees between a source s and points on an x - and y -monotone rectilinear path L , which is called a *staircase path*. In Section 3.2, we show how to combine two shallow-light trees to obtain a spanner between point pairs on two staircase paths.


3.1 Single Source and Staircase Chain

We present a new, slightly modified proof for Solomon’s result on SLTs between a single source s and a horizontal line segment, and then adapt the modified proof to obtain a SLT between s and an x - and y -monotone polygonal chain. In the proof below, we use the Taylor estimates $\cos x \geq 1 - x^2/2$ and $\sin x \geq x/2$ for $x \leq \pi/3$.

Alternative proof for Lemma 4. Assume w.l.o.g. that $\varepsilon = 2^{-k}$ for $k \in \mathbb{N}$. Let $T = \{t_i : i = 1, \dots, 2^{k+1}\}$ be 2^{k+1} points on the line segment $L = [-\frac{1}{2}, \frac{1}{2}] \times \{0\}$ with uniform $1/(2^{k+1} - 1) < \varepsilon$ spacing between consecutive points. Consider the standard binary partition of $\{1, \dots, 2^{k+1}\}$ into intervals, associated with a binary tree: At level 0, the root corresponds to the interval $[1, 2^{k+1}]$ of all 2^{k+1} integer. At level j , we have intervals $[i \cdot 2^{k-j} + 1, (i+1) \cdot 2^{k-j}]$ for $i = 0, \dots, 2^j$. Note that if a point q is the left (resp., right) endpoint of an interval at a level j , then it is a left (resp., right) endpoint of all descendant intervals that contains it.

For every $q \in \{1, \dots, 2^{k+1}\}$, we define a line segment ℓ_q with one endpoint at t_q : Let $j \geq 0$ be the smallest level such that q is an endpoint of some interval I_q at level j . If q is the left (resp., right) endpoint of I_q , then let ℓ_q be the line segment of direction $\frac{\pi}{2} - 2^{(j-k)/2}$

15:6 Light Euclidean Steiner Spanners in the Plane



■ **Figure 2** The segments added to graph G at level $j = 0, 1, 2$ for $m = 2^3 = 8$ points. The intervals $[t_a, t_b]$ at level j are indicated below the line L .

(resp., $\frac{\pi}{2} + 2^{(j-k)/2}$) such that its orthogonal projection to the x -axis is I_q ; see Fig. 2. Note that for $j = 0$, we use directions $\frac{\pi}{2} \pm 2^{-k/2} = \frac{\pi}{2} \pm \sqrt{\varepsilon}$. Let G be the union of segments ℓ_q for $q = 1, \dots, 2^{k+1}$, the horizontal segment L , and the vertical segment from s to the origin.


Lightness analysis. We show that $\|G\| = O(\varepsilon^{-1/2})$. We have $\|L\| = 1$, and the length of the vertical segment between s and the origin is $\varepsilon^{-1/2}$. At level j of the binary tree, we construct 2^j segments ℓ , each of length $\|\ell\| \leq 2^{-j} / \sin(2^{(j-k)/2}) \leq 2 \cdot 2^{(k-3j)/2}$. Summation over all levels yields $\sum_{j=0}^k 2^j \cdot 2 \cdot 2^{(k-3j)/2} = 2^{k/2} \cdot 2 \cdot \sum_{j=0}^k 2^{-j/2} = O(2^{k/2}) = O(\varepsilon^{-1/2})$.

Source-stretch analysis. We show that G contains an st_q -path of length $(1 + O(\varepsilon))\|st_q\|$ for all $q = 1, \dots, 2^{k+1}$. First note that $\|st_q\| \geq \varepsilon^{-1/2}$, as the distance between s and L is $\varepsilon^{-1/2}$. For each interval $[t_a, t_b]$ in the binary tree, ℓ_a and ℓ_b have positive and negative slopes, respectively, and so they cross above the interval $[t_a, t_b]$. Consequently, for every point t_q , the union of the $k+1$ segments corresponding to the intervals that contain t_q must contain a y -monotonically increasing path P_q from t_q to s . The y -projection of this path has length $\varepsilon^{-1/2}$. Consider one edge e of P_q along a segment ℓ at level j , which has direction $\frac{\pi}{2} \pm \alpha = \frac{\pi}{2} \pm 2^{(j-k)/2}$. Then the difference between the length of e and the y -projection of e is $\|e\|(1 - \cos \alpha) \leq \|\ell\|(1 - \cos \alpha) \leq 2^{-j} \frac{1 - \cos \alpha}{\sin \alpha} \leq 2^{-j} \frac{\alpha^2/2}{\alpha/2} = 2^{-j} \alpha = 2^{-j} \cdot 2^{(j-k)/2} = 2^{-(j+k)/2}$. Since P_q contains at most one edge in each level, summation over all edges of P_q yields $\sum_{j=0}^k 2^{-(j+k)/2} = 2^{-k/2} \sum_{j=0}^k 2^{-j/2} = O(\varepsilon^{1/2}) \leq \|st_q\| \cdot O(\varepsilon)$.

Finally, for an arbitrary point $t \in L$, we have $\|st\| \geq \varepsilon^{-1/2}$, and G contains an st -path that consists of an st_q -path from s to the point t_q closest to t , followed by the horizontal segment $t_q t$ of weight $\|t_q t\| \leq 1/2^k \leq \varepsilon$. The total weight of this path is $(1 + O(\varepsilon))\|st\|$. After suitable scaling of the constant coefficients, G contains a path of weight at most $(1 + \varepsilon)\|st\|$ for any $t \in L$, as required. ◀

► **Lemma 5.** Let $0 < \varepsilon < 1$, let $s = (0, \varepsilon^{-1/2})$ be a point on the y -axis, and let L be an x - and y -monotone increasing staircase path between the vertical lines $x = \pm \frac{1}{2}$, such that the right endpoint of L is $(\frac{1}{2}, 0)$ on the x -axis. Then there exists a geometric graph G comprised of L and additional edges of weight $O(\varepsilon^{-1/2})$ such that G contains, for every $t \in L$, an st -path P_{st} with $\|P_{st}\| \leq (1 + O(\varepsilon))\|st\|$.

We can adjust the construction above as follows; refer to Fig. 3.



■ **Figure 3** The paths γ_q added to graph G at level $j = 0, 1, 2$ for $m = 2^3 = 8$ points. The intervals $[t_a, t_b]$ at level j are indicated below the staircase path L .

Proof. Assume w.l.o.g. that $\varepsilon = 2^{-k}$ for some $k \in \mathbb{N}$. Let $T = \{t_i : i = 1, \dots, 2^{k+1}\}$ be 2^{k+1} points in L on equally spaced vertical lines, with spacing $1/(2^{k+1} - 1) < \varepsilon$. Consider the standard binary partition of $\{1, \dots, 2^{k+1}\}$ into intervals as in the previous proof.

For every $q \in \{1, \dots, 2^{k+1}\}$, we define a polygonal path γ_q with one endpoint at t_q ; see Fig. 3. Let $j \geq 0$ be the smallest level such that t_q is an endpoint of some interval I_q at level j . If t_q is the right endpoint of I_q , then let γ_q be the line segment of direction $\frac{\pi}{2} + 2^{(j-k)/2}$ such that its x -projection is I_q . If t_q is the left endpoint of I_q , then γ_q will be an x - and y -monotone path whose x -projection is I_q , and its edges will be vertical segments along L and segments of direction $\alpha_q = \frac{\pi}{2} - 2^{(j-k)/2}$. Specifically, γ_q starts from t_q with a line of direction α_q . Whenever γ_q encounters a vertical edge of L , it follows it upward until its upper endpoint, and then continues in direction α_q .

Let G be the union of all paths γ_q for $q = 1, \dots, 2^{k+1}$, as well as the path L , and the vertical segment from s to the origin. This completes the construction of G .

Lightness analysis. We show that $\|G\| = \|L\| + O(\varepsilon^{-1/2})$. The distance between s and L is $\varepsilon^{-1/2}$. For every $q \in \{1, \dots, 2^{k+1}\}$, the path γ_q is composed of vertical segments along L , and nonvertical segments whose total weight is the same as $\|\ell_q\|$ in the proof of Lemma 4, where we have seen that $\sum_{q=1}^{2^{k+1}} \|\ell_q\| = O(\varepsilon^{-1/2})$. Consequently, $\|G\| = \|L\| + O(\varepsilon^{-1/2})$.

Source-stretch analysis. We show that G contains an st_q -path of weight $(1 + (\varepsilon))\|st_q\|$ for all $q = 1, \dots, 2^{k+1}$. Denoting $y(t_q)$ the y -coordinate of point t_q , we have $\|st_q\| \geq \varepsilon^{-1/2} + |y(t_q)|$. For each interval $[t_a, t_b]$ in the binary tree, the paths γ_a and γ_b cross above the portion of L between t_a and t_b . Consequently, for every point t_q , the union of the $k + 1$ paths γ corresponding to the intervals that contain t_q must contain a y -monotonically increasing path P_q from t_q to s . The y -projection of this path has length $\varepsilon^{-1/2} + |y(t_q)|$. Some of the edges of this path are vertical. Consider the union of all nonvertical edges e of P_q along a path γ at level j , which all have direction $\frac{\pi}{2} \pm 2^{(j-k)/2}$. The difference between the length of e and the y -projection of e is bounded by the same analysis as in the proof of Lemma 4. Summation over all levels yields $O(\varepsilon^{1/2}) \leq \|st_q\| \cdot O(\varepsilon)$.

Finally, for an arbitrary point $t \in L$, we have $\|st\| \geq |y(s) - y(t)| = \varepsilon^{-1/2} + |y(t)|$, and G contains an st -path that consists of an st_q -path from s to the closest point t_q to the right of t , followed by an x - and y -monotone path along L in which the total length of the horizontal edges is bounded by $1/2^k \leq \varepsilon$ (and the length of vertical segment might be arbitrary). We use the lower bound $\|st\| \geq \varepsilon^{-1/2} + |y(t)|$. The vertical segments between t_q and t do not contribute to the error term $\|st\| - (\varepsilon^{-1/2} + |y(t)|)$. The analysis in the proof of Lemma 4 yields $\|st\| - (\varepsilon^{-1/2} + |y(t)|) \leq O(\sqrt{\varepsilon}) \leq O(\varepsilon)\|st\|$. ◀

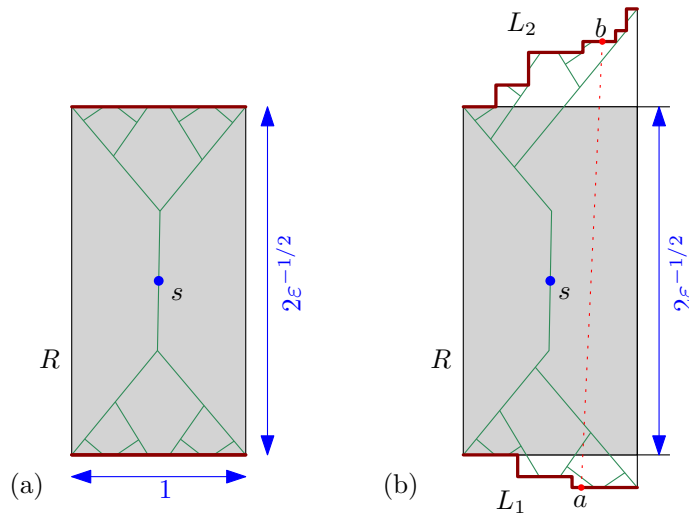
Note that the source-stretch analysis assumed that the vertical edges of an st -path (along the vertical edges of L) do not accumulate any error. Consequently, the same analysis carries over if we replace the vertical edges of L by $(\frac{\pi}{2} \pm \frac{\sqrt{\varepsilon}}{2})$ -angle-bounded paths. The key observation is that in the proof of Lemma 5, all nonvertical edges have directions that differ from vertical (i.e., from $\frac{\pi}{2}$) by $\sqrt{\varepsilon}$ or more.

► **Corollary 6.** *Let $0 < \varepsilon < 1$, let $s = (0, \varepsilon^{-1/2})$ be a point on the y -axis, and let L be a path between the vertical lines $x = \pm \frac{1}{2}$, obtained from an x - and y -monotone increasing staircase path with the right endpoint at $(\frac{1}{2}, 0)$ on the x -axis, by replacing the vertical edges with y -monotonically increasing $(\frac{\pi}{2} \pm \frac{\sqrt{\varepsilon}}{2})$ -angle-bounded paths. Then there exists a geometric graph G that contains L and additional edges of weight $O(\varepsilon^{-1/2})$ such that G contains, for every $t \in L$, an st -path P_{st} with $\|P_{st}\| \leq (1 + O(\varepsilon))\|st\|$.*

3.2 Combination of Shallow-Light Trees

The combination of two SLTs yields a light $(1 + \varepsilon)$ -spanner between points on two staircases.

► **Lemma 7.** *Let R be an axis-parallel rectangle of width 1 and height $2\varepsilon^{-1/2}$; and let L_1 (resp., L_2) be a staircase path from the lower-left (upper-left) corner of R to a point on the vertical line passing through the right side of R , lying below (above) R ; see Fig. 4. Then there exists a geometric graph comprised of $L_1 \cup L_2$ and additional edges of weight $O(\varepsilon^{-1/2})$ that contains an ab -path P_{ab} with $\|P_{ab}\| \leq (1 + O(\varepsilon))\|ab\|$ for any $a \in L_1$ and any $b \in L_2$.*



■ **Figure 4** (a) A combination of two SLTs between the two horizontal sides of R . (b) A combination of two SLTs between two staircases above and below R , respectively.

Proof. Let s be the center of the rectangle R . Let G be the geometric graph formed by the SLTs from the source s to L_1 and L_2 , resp., using Lemma 5. By construction, $\|G\| = \|L_1\| + \|L_2\| + O(\varepsilon^{-1/2})$. It remains to show that G has the desired spanning ratio. Let $a \in L_1$ and $b \in L_2$. Let h_a be the distance of a from bottom side of R , and h_b the distance of b from the top side of R . By Lemma 4, the two SLTs jointly contain an ab -path P_{ab} of length $\|P_{ab}\| \leq (1 + O(\varepsilon))(\|as\| + \|bs\|)$.

On the one hand, s is the center of R , and so $\|as\| + \|bs\| \leq \text{diam}(R) + h_a + h_b \leq (1 + \frac{\varepsilon}{8})2\varepsilon^{-1/2} + h_a + h_b$. On the other hand, $\|ab\| \geq \text{height}(R) + h_a + h_b = 2\varepsilon^{-1/2} + h_a + h_b$. Overall, $\|P_{ab}\| \leq (1 + O(\varepsilon))(1 + \frac{\varepsilon}{8})\|ab\| \leq (1 + O(\varepsilon))\|ab\|$. ◀

4 Reduction to Directional Spanners in Histograms

In this section, we present our general strategy for the proof of Theorem 2, and reduce the construction of a light $(1 + \varepsilon)$ -spanner for a point set S in the plane to a special case of *directional spanners* for a point set on the boundary of faces in a (modified) *window partition*.

Directional $(1 + \varepsilon)$ -Spanners. Our strategy to construct a $(1 + \varepsilon)$ -spanner for a point set S is to partition the interval of directions $[0, \pi)$ into $O(\varepsilon^{-1/2})$ intervals, each of length $O(\varepsilon^{1/2})$. For each interval $D \subset [0, \pi)$, we construct a geometric graph that serves point pairs $a, b \in S$ with $\text{dir}(ab) \in D$. Then the union of these graphs over all $O(\varepsilon^{-1/2})$ intervals will serve all point pairs $ab \in S$. The following definition formalizes this idea.

► **Definition 8.** Let S be a finite point set in \mathbb{R}^2 , and let $D \subset [0, \pi)$ be a set of directions. A geometric graph G is a directional $(1 + \varepsilon)$ -spanner for S and D if G contains an ab -path of weight at most $(1 + \varepsilon)\|ab\|$ for every $a, b \in S$ with $\text{dir}(ab) \in D$.

In Section 6, we modify the standard window partition algorithm and partition a bounding box of S into *fuzzy staircases* and *tame histograms* (defined below). We also construct directional spanners for point pairs $a, b \in S$, where ab is a chord of a face in this partition. A line segment ab is a *chord* of a simple polygon P if $a, b \in \partial P$, and $ab \subset P$. The *perimeter* of a simple polygon P , denoted $\text{per}(P)$, is the total weight of the edges of P ; and the *horizontal perimeter*, denoted $\text{hper}(P)$, is the total weight of the horizontal edges of P .

► **Lemma 9.** We can subdivide a simple rectilinear polygon P into a collection \mathcal{F} of fuzzy staircases and tame histograms of total perimeter $\sum_{F \in \mathcal{F}} \text{per}(F) \leq O(\varepsilon^{-1/2} \text{per}(P))$ and total horizontal perimeter $\sum_{F \in \mathcal{F}} \text{hper}(F) \leq O(\text{per}(P))$.

► **Lemma 10.** Let F be a fuzzy staircase or a tame histogram, $S \subset \partial F$ a finite point set, $\varepsilon > 0$, and $D = [\frac{\pi - \sqrt{\varepsilon}}{2}, \frac{\pi + \sqrt{\varepsilon}}{2}]$ an interval of nearly vertical directions. Then there exists a geometric graph of weight $O(\text{per}(F) + \varepsilon^{-1/2} \text{hper}(F))$ such that for all $a, b \in S$, if ab is a chord of F and $\text{dir}(ab) \in D$, then G contains an ab -path of weight at most $(1 + O(\varepsilon))\|ab\|$.

For the proof of Lemmas 9 and 10, refer to the full paper [6]. In the remainder of this section, we show that these lemmas imply Lemma 11, which in turn implies Theorem 2.

► **Lemma 11.** Let $S \subset \mathbb{R}^2$ be a finite point set, $\varepsilon > 0$, and $D \subset [0, \pi)$ an interval of length $\sqrt{\varepsilon}$. Then there exists a directional $(1 + \varepsilon)$ -spanner for S and D of weight $O(\varepsilon^{-1/2} \|MST(S)\|)$.

Proof. We may assume, by applying a suitable rotation, that $D = [\frac{\pi - \sqrt{\varepsilon}}{2}, \frac{\pi + \sqrt{\varepsilon}}{2}]$, that is, an interval of nearly vertical directions. We construct a directional $(1 + \varepsilon)$ -spanner for S and D of weight $O(\varepsilon^{-1/2} \cdot \|MST(S)\|)$.

15:10 Light Euclidean Steiner Spanners in the Plane

Assume w.l.o.g. that the unit square $U = [0, 1]^2$ is the minimum axis-parallel bounding square of S . In particular, S has two points on two opposite sides of U , and so $1 \leq \text{diam}(S) \leq \|\text{MST}(S)\|$. Our initial graph G_0 is composed of the boundary of U and a *rectilinear MST* of S , where $\|G_0\| = O(\|\text{MST}(S)\|)$. Since each edge of G_0 is on the boundary of at most two faces, the total perimeter of all faces of G_0 is also $O(\|\text{MST}(S)\|)$. Lemma 9 yields subdivisions of the faces of G_0 into a collection \mathcal{F} of fuzzy staircases and tame histograms of total perimeter $\sum_{F \in \mathcal{F}} \text{per}(F) = O(\varepsilon^{-1/2} \|\text{MST}(S)\|)$ and horizontal perimeter $\sum_{F \in \mathcal{F}} \text{hper}(F) = O(\|\text{MST}(S)\|)$,

Let $K(S)$ be the complete graph induced by S . For each face $F \in \mathcal{F}$, let S_F be the set of all intersection points between the boundary ∂F and the edges of $K(S)$. For each face F , Lemma 10 yields a geometric graph G_F of weight $O(\text{per}(F) + \varepsilon^{-1/2} \text{hper}(F))$ with respect to the finite point set $S_F \subset \partial F$.

We can now put the pieces back together. Let G be the union of G_0 and the graphs G_F for all $F \in \mathcal{F}$. The weight of G is bounded by $\|G\| = \|G_0\| + \sum_{F \in \mathcal{F}} \|G_F\| = O(\|\text{MST}(S)\| + \sum_{F \in \mathcal{F}} (\text{per}(F) + \varepsilon^{-1/2} \text{hper}(F))) = O(\varepsilon^{-1/2} \|\text{MST}(S)\|)$.

Let $a, b \in S$. The edges of G_0 subdivide the line segment ab into a path $v_0 v_1 \dots v_m$ of collinear segments, each of which is a chord of some face in \mathcal{F} . For $i = 1, \dots, m$, graph G contains a $v_{i-1} v_i$ -path of weight at most $(1 + \varepsilon) \|v_{i-1} v_i\|$. The concatenation of these paths is an ab -path of length at most $\sum_{i=1}^m (1 + \varepsilon) \|v_{i-1} v_i\| = (1 + \varepsilon) \|ab\|$, as required. \blacktriangleleft

Proof of Theorem 2. Let S be a finite set of points in the plane. Let $\varepsilon > 0$ be given. For $k = \lceil \pi \varepsilon^{-1/2} \rceil$, we partition the space of directions as $[0, \pi) = \bigcup_{i=1}^k D_i$, into k intervals of equal length. By Lemma 11, there exists a directional $(1 + \varepsilon)$ -spanner of weight $O(\varepsilon^{-1/2} \|\text{MST}(S)\|)$ for S and D_i for all i . Let $G = \bigcup_{i=1}^k G_i$. For every point pair $s, t \in S$, we have $\text{dir}(st) \in D_i$ for some $i \in \{1, \dots, k\}$, and $G_i \subset G$ contains an st -path of weight at most $(1 + \varepsilon) \|st\|$. Consequently, G is a Euclidean Steiner $(1 + \varepsilon)$ -spanner for S . The weight of G is $\|G\| = \sum_{i=1}^k \|G_i\| \leq \lceil \pi \varepsilon^{-1/2} \rceil \cdot O(\varepsilon^{-1/2} \|\text{MST}(S)\|) \leq O(\varepsilon^{-1} \|\text{MST}(S)\|)$, as required. \blacktriangleleft


5 Construction of Directional Spanners for Staircases

In this section, we handle the special case where the points are on a x - and y -monotone rectilinear path L , which is called a *staircase path*. Our recursive construction uses a type of polygons that we define now. A *shadow polygon* is bounded by a staircase path L and a single line segment of slope $\varepsilon^{-1/2}$; see Fig. 5(a) for examples.

► **Lemma 12.** *Let L be an x - and y -monotonically increasing staircase path, and let $S \subset L$ be a finite point set. Then there exists a geometric graph G comprised of L and additional edges of weight $O(\varepsilon^{-1/2} \text{width}(L))$ such that G contains a path P_{ab} of weight $\|P_{ab}\| \leq (1 + O(\varepsilon)) \|ab\|$ for any $a, b \in L$ where $\text{slope}(ab) \geq \varepsilon^{-1/2}$ and the line segment ab lies below L .*

Proof. If $a, b \in L$ and ab lies below L , then either both a and b are in the same edge of L (hence L contains a straight-line path ab), or one point in $\{a, b\}$ is on a vertical edge of L and the other is on a horizontal edge of L . We may assume w.l.o.g. that a is on a vertical edge and b is on a horizontal edge of L .

Let A be the set of all points p such that there exists $a \in L$ on some vertical edge of L such that $\text{slope}(ap) \geq \varepsilon^{-1/2}$ and ap is below L ; see Fig. 5(a). The set A is not necessarily connected, the connected components of A are shadow polygons for disjoint subpaths of L . Let \mathcal{U} be the set of these shadow polygons. Note that for every pair $a, b \in L$, if $\text{slope}(ab) \geq \varepsilon^{-1/2}$ and ab lies below the path L , then ab lies in some polygon in \mathcal{U} . For each polygon $U \in \mathcal{U}$, we construct a geometric graph $G(U)$ of weight $O(\varepsilon^{-1/2} \text{width}(U))$ such



■ **Figure 5** (a) A staircase path L ; the shadow of vertical edges of L is shaded light gray. (b) The shadow of the horizontal edges in the subpolygons is shaded dark gray. (c) Recursive subproblems generated in the proof of Lemma 12.

that $G(U) \cup L$ is a directional spanner for the point pairs in $S \cap U$. Then L together with $\bigcup_{U \in \mathcal{U}} G(U)$ is a directional spanner for all possible ab pairs. Since the shadow polygons in \mathcal{U} are adjacent to disjoint portions of L , we have $\sum_{U \in \mathcal{U}} \text{width}(U) \leq \text{width}(L)$, and so $\sum_{U \in \mathcal{U}} \|G(U)\| = O(\varepsilon^{-1/2} \text{width}(L))$, as required.

Recursive Construction. For all $U \in \mathcal{U}$, we construct $G(U)$ recursively as follows. Assume that $|S \cap U| \geq 2$. Let $B(U)$ be the set of all points $p \in U$ for which there exists a point b on some horizontal edge of U such that $bp \subset U$ and $\text{slope}(ab) \geq \frac{1}{2}\varepsilon^{-1/2}$; see Fig. 5(b). The set $B(U)$ may be disconnected, each component is a simple polygon bounded by a contiguous portion of L and a line segment of slope $\frac{1}{2}\varepsilon^{-1/2}$. Denote by \mathcal{V} the set of connected components of $B(U)$.

For every $V \in \mathcal{V}$, let $C(V)$ be the set of all points $p \in V$ for which there exists a point a on some vertical edge of V such that $ap \subset V$ and $\text{slope}(ap) \geq \varepsilon^{-1/2}$; see Fig. 5(b). Again, the set $C(V)$ may be disconnected, each component is a shadow polygon. Denote by \mathcal{W} the set of all connected components of $C(V)$ for all $V \in \mathcal{V}$.

Since $\text{height}(W)/\text{width}(W) = \varepsilon^{-1/2}$ for all $W \in \mathcal{W}$ and $\text{height}(V)/\text{width}(V) = \frac{1}{2}\varepsilon^{-1/2}$ for all $V \in \mathcal{V}$, we have

$$\begin{aligned} \sum_{W \in \mathcal{W}} \text{width}(W) &= \sqrt{\varepsilon} \cdot \sum_{W \in \mathcal{W}} \text{height}(W) = \sqrt{\varepsilon} \cdot \sum_{V \in \mathcal{V}} \text{height}(V) \\ &= \frac{1}{2} \sum_{V \in \mathcal{V}} \text{width}(V) = \frac{1}{2} \sum_{U \in \mathcal{U}} \text{width}(U). \end{aligned} \tag{1}$$

For every polygon $V \in \mathcal{V}$, let s_V be the bottom vertex of V . We construct a sequence of shallow-light trees from source s_V as follows. For every nonnegative integer $i \geq 0$, let h_i be a horizontal line at distance $\text{height}(V)/2^i$ above s_V . If there is any point in S between h_i and h_{i+1} , then we construct an SLT from s_V to the portion of L between h_i and h_{i+1} . By Lemma 12, the total weight of these trees is $O(\varepsilon^{-1/2} \text{width}(V))$. Over all $V \in \mathcal{V}$, the weight of these SLTs is $\sum_{V \in \mathcal{V}} O(\varepsilon^{-1/2} \text{width}(V)) = O(\varepsilon^{-1/2} \text{width}(U))$. For all $V \in \mathcal{V}$, we also add the boundary ∂V to our spanner, at a cost of $\sum_{V \in \mathcal{V}} \text{per}(V) = \sum_{V \in \mathcal{V}} O(\varepsilon^{-1/2} \text{width}(V)) = O(\varepsilon^{-1/2} \text{width}(U))$. This completes the description of one iteration. Recurse on all $W \in \mathcal{W}$ that contain any point in S .


Lightness analysis. Each iteration of the algorithm, for a shadow polygon U , constructs SLTs of total weight $O(\varepsilon^{-1/2}\text{width}(U))$, and produces subproblems whose combined width is at most $\frac{1}{2}\text{width}(U)$ by Equation (1). Consequently, summation over all levels of the recursion yields $\|G(U)\| = O(\varepsilon^{-1/2}\text{width}(U) \cdot \sum_{i \geq 0} 2^{-i}) = O(\varepsilon^{-1/2}\text{width}(U))$, as required.

Stretch analysis. Now consider a point pair $a, b \in S$ such that $\text{slope}(ab) \geq \varepsilon^{-1/2}$, a is in a vertical edge of L , and b is in a horizontal edge of L . Assume that U is the smallest shadow polygon in the recursive algorithm above that contains both a and b . Then $b \in V$ for some $V \in \mathcal{V}$, and a is at or below vertex s_V of V . Now we can find an ab -path P_{ab} as follows: First construct a y -monotonically increasing path from a to V_S along vertical edges of L and along edges of some polygons in \mathcal{V} ; all these edges have slope larger than $\frac{1}{2}\varepsilon^{-1/2}$. Then from s_V , follow an SLT to b . All edges of P_{ab} from a to s_V have slope at least $\frac{1}{2}\varepsilon^{-1/2}$, and so their directions differ from vertical by at most $\arctan(2\varepsilon^{1/2}) \leq 3\varepsilon^{1/2}$, using the Taylor expansion of $\tan(x)$ near 0. By Lemma 3 the stretch factor of the paths from a to s_V and the path $a s_V b$ are each at most $1 + O(\varepsilon)$. By Lemma 12, the SLT contains a path from s_V to b with stretch factor $1 + O(\varepsilon)$. Overall, $\|P_{ab}\| \leq (1 + O(\varepsilon))\|ab\|$. ◀

In the full paper [6], it is shown that Lemma 12 continues to hold if we replace the vertical edges of the staircase L with angle-bounded paths. Furthermore, the horizontal edges can also be replaced by x -monotone paths of approximately the same length.

6 Construction of Directional Spanners in Histograms


We would like to partition a simple rectilinear polygon P into a collection \mathcal{F} of simple polygons (faces), and then design a directional $(1 + \varepsilon)$ -spanner for each face $F \in \mathcal{F}$ such that the total weight of these spanners is under control. Lemma 12 tells us that we can handle staircase polygons efficiently. The standard window partition [37, 48] would partition P into histograms as indicated in Fig. 6(a). We would like to further reduce the problem to staircase polygons. However, the worst-case weight of a standard decomposition of a histogram H into staircases is $\Theta(\text{per}(H) \log n)$, where n is the number of vertices of H . We cannot afford a $\log n$ factor (or any function of the cardinality $|S|$). To overcome this technical difficulty, we replace the vertical edges by nearly vertical δ -angle-bounded paths (cf. Lemma 3). By setting $\delta = \Theta(\sqrt{\varepsilon})$, these paths provide enough flexibility to keep the weight of the subdivision under control; and our result on SLTs for these “modified” staircases carry over with only a constant increase in their total weight.



■ **Figure 6** (a) A standard window partition of a rectilinear polygon P into histograms, starting from a horizontal edge e_0 . (b) A decomposition of a y -monotone histogram into staircase polygons. (c) The modified window partition of a rectilinear polygon P into x -monotone Λ -histograms and y -monotone fuzzy histograms.

We introduce some terminology; refer to Fig. 7. Let $\Lambda \geq 8$ be a constant.

- A Λ -path is a y -monotone path in which every edge is vertical, or has slope $\pm \Lambda \varepsilon^{-1/2}$.
- A Λ -histogram is a simple polygon obtained from a histogram by replacing vertical edges with some Λ -paths. A Λ -histogram is x -monotone (resp., y -monotone) if it is obtained from an x -monotone (resp., y -monotone) histogram.
- A fuzzy staircase is a simple polygon bounded by a path pqr , where pq is horizontal and $\text{slope}(qr) = \pm \Lambda \varepsilon^{-1/2}$, and a pr -path obtained from an x - and y -monotone staircase by replacing vertical edges with some Λ -paths.
- A fuzzy histogram is a simple polygon bounded by a y -monotone rectilinear path L and a path γ of one or two edges of slopes $\pm \Lambda \varepsilon^{-1/2}$; if the latter path has two edges, then its interior vertex is a reflex vertex of the polygon.
- A tame histogram (Fig. 8(a)) is a simple polygon bounded by a horizontal line segment pq and an pq -path L that consists of ascending or descending Λ -paths and x -monotone increasing horizontal edges with the following properties: (i) there is no chord between interior points of any two ascending (resp., descending) Λ -paths; (ii) for every horizontal chord ab , with $a, b \in L$, the subpath L_{ab} of L between a and b satisfies $\|L_{ab}\| \leq 2\|ab\|$.
- A tame path is a subpath of the pq -path L of a tame histogram.

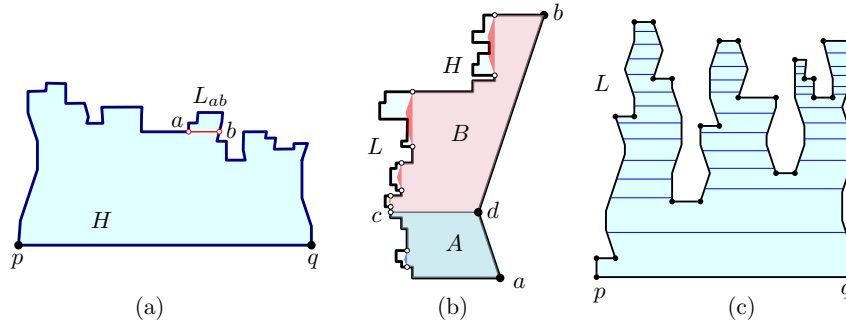


■ **Figure 7** (a) A Λ -path. (b) A y -monotone Λ -histogram. (c) An x -monotone Λ -histogram. (d) Three fuzzy staircases. (e) A fuzzy histogram.

In what follows, we describe our spanner constructions in five modules. However, due to space constraints we provide only an overview of these modules and refer to the full version of the paper [6] for the formal description and the proofs.

Fuzzy Window Decomposition. Let R be a rectilinear simple polygon. By modifying the standard window-partition, we show how to partition R into a collection \mathcal{H} of x -monotone Λ -histograms and y -monotone fuzzy histograms such that $\sum_{H \in \mathcal{H}} \text{per}(H) = O(\text{per}(P))$; see Fig. 6(b). Furthermore, we show that in the x -monotone Λ -histograms, there is no chord between interior points of two ascending (resp., two descending) Λ -paths.

y -Monotone Histograms. Let H be a (y -monotone) fuzzy histogram. We recursively subdivide H into a family \mathcal{F} of fuzzy staircases using subdivision edges of total weight $O(\varepsilon^{-1/2} \text{per}(H))$ such that $\sum_{F \in \mathcal{F}} \text{hper}(F) = O(\text{per}(H))$; see Fig. 8(b) for an illustration.



■ **Figure 8** (a) A tame histogram. (b) Recursive subdivision of a fuzzy histogram into fuzzy staircase polygons. (c) Recursive subdivision of an x -monotone Λ -histograms into tame histograms.

x -Monotone Λ -Histograms. Let H be an x -monotone Λ -histogram that does not have any chords between interior points of any two ascending (resp., two descending) Λ -paths. We use a sweepline algorithm to subdivide H into tame histograms; see Fig. 8(c) for an illustration. Specifically, we subdivide H into a collection \mathcal{T} of tame histograms such that $\sum_{T \in \mathcal{T}} \text{per}(T) = O(\text{per}(H))$. This module provides a proof for Lemma 9.

Directional Spanners for Tame Histograms. Given a tame histogram H and a finite set of points $S \subset \partial H$, we construct a directional spanner for S with respect to point pairs $a, b \in S$ with $|\text{slope}(ab)| \geq \varepsilon^{-1/2}$. First we adapt Lemma 12 to construct a directional spanner for points $a, b \in S$ on a tame path $L \subset \partial H$; and then generalize Lemma 7 to handle point pairs where a is in the horizontal base of H and $b \in L$.

Directional Spanners for Fuzzy Staircases. Given a fuzzy staircase polygon F and a finite point set $S \subset \partial F$, we construct a directional spanner of weight $\|G\| = O(\varepsilon^{-1/2} \text{hper}(F))$ for S with respect to point pairs $a, b \in S$ with $|\text{slope}(ab)| \geq \varepsilon^{-1/2}$. The last two modules jointly imply Lemma 10, and complete all components needed for Theorem 2.

7 Conclusion and Outlook

We have proved a tight upper bound of $O(\varepsilon^{-1})$ on the lightness of Euclidean Steiner $(1 + \varepsilon)$ -spanners in the plane. That is, for every finite set $S \subset \mathbb{R}^2$, there is a Euclidean Steiner $(1 + \varepsilon)$ -spanner of weight $O(\varepsilon^{-1} \|\text{MST}(S)\|)$. Our proof is constructive, but we do not control the number of Steiner points. This immediately raises the question about the optimum number of Steiner points: What is the minimum sparsity of a Euclidean Steiner $(1 + \varepsilon)$ -spanner of weight $O(\frac{1}{\varepsilon} \|\text{MST}(S)\|)$ that can be attained for all finite set of points in \mathbb{R}^2 ?

Planarity is an important aspect of any geometric networks. Therefore, it is desirable to construct Euclidean $(1 + \varepsilon)$ -spanners that are planar, i.e., no two edges of the spanner cross. Any Steiner spanner can be turned into a plane spanner (planarized), with the same weight and the same spanning ratio between the input points, by introducing Steiner points at all edge crossings. However, planarization may substantially increase the number of Steiner points. Bose and Smid [10, Sec. 4] note that Arikati et al. [2] constructed a Euclidean plane $(1 + \varepsilon)$ -spanner with $O(\varepsilon^{-4}n)$ Steiner points for any n points in \mathbb{R}^2 ; see also [38]. Borradaile and Eppstein [8] improved the bound to $O(\varepsilon^{-3}n \log \varepsilon^{-1})$ in certain special cases where all Delaunay faces are fat. It remains an open problem to find the optimum dependence of ε for plane Steiner $(1 + \varepsilon)$ -spanners; and for plane Steiner $(1 + \varepsilon)$ -spanners of lightness $O(\varepsilon^{-1})$.

References

- 1 Ingo Althöfer, Gautam Das, David Dobkin, Deborah Joseph, and José Soares. On sparse spanners of weighted graphs. *Discrete & Computational Geometry*, 9(1):81–100, 1993.
- 2 Srinivasa Rao Arikati, Danny Z. Chen, L. Paul Chew, Gautam Das, Michiel H. M. Smid, and Christos D. Zaroliagis. Planar spanners and approximate shortest path queries among obstacles in the plane. In *Proc. 4th European Symposium on Algorithms (ESA)*, volume 1136 of *LNCS*, pages 514–528. Springer, 1996.
- 3 Sunil Arya, David M Mount, and Michiel Smid. Randomized and deterministic algorithms for geometric spanners of small diameter. In *Proc. 35th IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 703–712, 1994.
- 4 Sunil Arya and Michiel Smid. Efficient construction of a bounded-degree spanner with low weight. *Algorithmica*, 17(1):33–54, 1997.
- 5 Baruch Awerbuch, Alan E. Baratz, and David Peleg. Cost-sensitive analysis of communication protocols. In *Proc. 9th ACM Symposium on Principles of Distributed Computing (PODC)*, pages 177–187, 1990.
- 6 Sujoy Bhore and Csaba D. Tóth. Light Euclidean Steiner spanners in the plane. *Preprint*, 2020. [arXiv:2012.02216](https://arxiv.org/abs/2012.02216).
- 7 Sujoy Bhore and Csaba D. Tóth. On Euclidean Steiner $(1+\varepsilon)$ -spanners. In *Proc. 38th Symposium on Theoretical Aspects of Computer Science (STACS)*, volume 187 of *LIPICs*, pages 13:1–13:16. Schloss Dagstuhl, 2021.
- 8 Glencora Borradaile and David Eppstein. Near-linear-time deterministic plane Steiner spanners for well-spaced point sets. *Comput. Geom.*, 49:8–16, 2015.
- 9 Glencora Borradaile, Hung Le, and Christian Wulff-Nilsen. Greedy spanners are optimal in doubling metrics. In *Proc. 13th ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 2371–2379, 2019.
- 10 Prosenjit Bose and Michiel H. M. Smid. On plane geometric spanners: A survey and open problems. *Comput. Geom.*, 46(7):818–830, 2013.
- 11 Paul B. Callahan. Optimal parallel all-nearest-neighbors using the well-separated pair decomposition. In *Proc. 34th IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 332–340, 1993.
- 12 Timothy M. Chan, Sarel Har-Peled, and Mitchell Jones. On locality-sensitive orderings and their applications. *SIAM J. Comput.*, 49(3):583–600, 2020.
- 13 Shiri Chechik and Christian Wulff-Nilsen. Near-optimal light spanners. *ACM Transactions on Algorithms (TALG)*, 14(3):1–15, 2018.
- 14 L. Paul Chew. There is a planar graph almost as good as the complete graph. In *Proc. 2nd Symposium on Computational Geometry*, pages 169–177. ACM Press, 1986.
- 15 L. Paul Chew. There are planar graphs almost as good as the complete graph. *J. Comput. Syst. Sci.*, 39(2):205–219, 1989.
- 16 Kenneth L. Clarkson. Approximation algorithms for shortest path motion planning. In *Proc. 19th ACM Symposium on Theory of Computing (STOC)*, pages 56–65, 1987.
- 17 Gautam Das, Paul Heffernan, and Giri Narasimhan. Optimally sparse spanners in 3-dimensional Euclidean space. In *Proc. 9th Symposium on Computational Geometry (SoCG)*, pages 53–62. ACM Press, 1993.
- 18 Gautam Das, Giri Narasimhan, and Jeffrey S. Salowe. A new way to weigh malnourished Euclidean graphs. In *Proc. 6th ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 215–222, 1995.
- 19 Michael J. Demmer and Maurice P. Herlihy. The arrow distributed directory protocol. In *Proc. 12th Symposium on Distributed Computing (DISC)*, volume 1499 of *LNCS*, pages 119–133. Springer, 1998.
- 20 Adrian Dumitrescu and Csaba D. Tóth. Light orthogonal networks with constant geometric dilation. *J. Discrete Algorithms*, 7(1):112–129, 2009.

- 21 Herbert Edelsbrunner, Joseph O'Rourke, and Emmerich Welzl. Stationing guards in rectilinear art galleries. *Computer Vision, Graphics, and Image Processing*, 27(2):167–176, 1984.
- 22 Michael Elkin, Ofer Neiman, and Shay Solomon. Light spanners. *SIAM Journal on Discrete Mathematics*, 29(3):1312–1321, 2015.
- 23 Michael Elkin and Shay Solomon. Steiner shallow-light trees are exponentially lighter than spanning ones. *SIAM Journal on Computing*, 44(4):996–1025, 2015.
- 24 Jie Gao, Leonidas J. Guibas, and An Nguyen. Deformable spanners and applications. *Comput. Geom.*, 35(1-2):2–19, 2006.
- 25 Lee-Ad Gottlieb. A light metric spanner. In *2015 IEEE 56th Symposium on Foundations of Computer Science*, pages 759–772, 2015.
- 26 Lee-Ad Gottlieb, Aryeh Kontorovich, and Robert Krauthgamer. Efficient regression in metric spaces via approximate Lipschitz extension. *IEEE Transactions on Information Theory*, 63(8):4838–4849, 2017.
- 27 Joachim Gudmundsson, Christos Levcopoulos, and Giri Narasimhan. Fast greedy algorithms for constructing sparse geometric spanners. *SIAM J. Comput.*, 31(5):1479–1500, 2002.
- 28 Joachim Gudmundsson, Christos Levcopoulos, Giri Narasimhan, and Michiel Smid. Approximate distance oracles for geometric spanners. *ACM Transactions on Algorithms (TALG)*, 4(1):1–34, 2008.
- 29 Maurice Herlihy, Srikanta Tirthapura, and Rogert Wattenhofer. Competitive concurrent distributed queuing. In *Proc. 20th ACM Symposium on Principles of Distributed Computing (PODC)*, pages 127–133, 2001.
- 30 J. Mark Keil. Approximating the complete Euclidean graph. In *Proc. 1st Scandinavian Workshop on Algorithm Theory (SWAT)*, volume 318 of *LNCS*, pages 208–213. Springer, 1988.
- 31 J. Mark Keil and Carl A. Gutwin. Classes of graphs which approximate the complete Euclidean graph. *Discrete & Computational Geometry*, 7:13–28, 1992.
- 32 Samir Khuller, Balaji Raghavachari, and Neal E. Young. Balancing minimum spanning and shortest path trees. In *Proc. 4th ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 243–250, 1993.
- 33 Hung Le and Shay Solomon. Truly optimal Euclidean spanners. In *Proc. 60th IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 1078–1100, 2019.
- 34 Hung Le and Shay Solomon. Light Euclidean spanners with Steiner points. In *Proc. 28th European Symposium on Algorithms (ESA)*, volume 173 of *LIPIcs*, pages 67:1–67:22. Schloss Dagstuhl, 2020.
- 35 Hung Le and Shay Solomon. A unified and fine-grained approach for light spanners. *CoRR*, abs/2008.10582, 2020. [arXiv:2008.10582](https://arxiv.org/abs/2008.10582).
- 36 Christos Levcopoulos. *Heuristics for Minimum Decompositions of Polygons*. PhD thesis, Linköping, 1987. No. 74 of Linköping Studies in Science and Technology.
- 37 Anil Maheshwari, Jörg-Rüdiger Sack, and Hristo N. Djidjev. Link distance problems. In Jörg-Rüdiger Sacks and Jorge Urutia, editors, *Handbook of Computational Geometry*, chapter 12, pages 519–558. North-Holland, 2000.
- 38 Anil Maheshwari, Michiel H. M. Smid, and Norbert Zeh. I/O-efficient algorithms for computing planar geometric spanners. *Comput. Geom.*, 40(3):252–271, 2008.
- 39 Giri Narasimhan and Michiel Smid. *Geometric Spanner Networks*. Cambridge University Press, 2007.
- 40 David Peleg and Alejandro A. Schäffer. Graph spanners. *Journal of Graph Theory*, 13(1):99–116, 1989.
- 41 David Peleg and Jeffrey D. Ullman. An optimal synchronizer for the hypercube. *SIAM J. Comput.*, 18(4):740–747, 1989.
- 42 David Peleg and Eli Upfal. A trade-off between space and efficiency for routing tables. *Journal of the ACM (JACM)*, 36(3):510–530, 1989.

- 43 Satish B. Rao and Warren D. Smith. Approximating geometrical graphs via “spanners” and “banyans”. In *Proc. 13th ACM Symposium on Theory of Computing (STOC)*, pages 540–550, 1998.
- 44 Jim Ruppert and Raimund Seidel. Approximating the d -dimensional complete Euclidean graph. In *Proc. 3rd Canadian Conference on Computational Geometry (CCCG)*, pages 207–210, 1991.
- 45 Christian Schindelhauer, Klaus Volbert, and Martin Ziegler. Geometric spanners with applications in wireless networks. *Comput. Geom.*, 36(3):197–214, 2007.
- 46 Michiel Smid. The weak gap property in metric spaces of bounded doubling dimension. In *Efficient Algorithms*, pages 275–289. Springer, 2009.
- 47 Shay Solomon. Euclidean Steiner shallow-light trees. *J. Comput. Geom.*, 6(2):113–139, 2015.
- 48 Subhash Suri. On some link distance problems in a simple polygon. *IEEE Trans. Robotics Autom.*, 6(1):108–113, 1990.
- 49 Andrew Chi-Chih Yao. On constructing minimum spanning trees in k -dimensional spaces and related problems. *SIAM J. Comput.*, 11(4):721–736, 1982.