Sparse Euclidean Spanners with Optimal Diameter: A General and Robust Lower Bound via a Concave Inverse-Ackermann Function

Hung Le
University of Massachusetts, Amherst, MA, USA

Lazar Milenković
Tel Aviv University, Israel

Shay Solomon
Tel Aviv University, Israel

Abstract

In STOC'95 [6] Arya et al. showed that any set of n points in $\mathbb{R}^d$ admits a $(1+\epsilon)$-spanner with hop-diameter at most 2 (respectively, 3) and $O(n \log n)$ edges (resp., $O(n \log\log n)$ edges). They also gave a general upper bound tradeoff of hop-diameter k with $O(n\alpha_k(n))$ edges, for any $k \geq 2$. The function α_k is the inverse of a certain Ackermann-style function, where $\alpha_0(n) = \lceil n/2 \rceil$, $\alpha_1(n) = \lceil \sqrt{n} \rceil$, $\alpha_2(n) = \lceil \log n \rceil$, $\alpha_3(n) = \lceil \log\log n \rceil$, $\alpha_4(n) = \log^* n$, $\alpha_5(n) = \lfloor \frac{1}{2} \log^* n \rfloor$, ... Roughly speaking, for $k \geq 2$ the function α_k is close to $\lfloor \frac{k-2}{2} \rfloor$-iterated log-star function, i.e., log with $\lfloor \frac{k-2}{2} \rfloor$ stars.

Despite a large body of work on spanners of bounded hop-diameter, the fundamental question of whether this tradeoff between size and hop-diameter of Euclidean $(1+\epsilon)$-spanners is optimal has remained open, even in one-dimensional spaces. Three lower bound tradeoffs are known:

- An optimal k versus $\Omega(n\alpha_k(n))$ by Alon and Schieber [4], but it applies to stretch 1 (not $1+\epsilon$).
- A suboptimal k versus $\Omega(n\alpha_{2k+6}(n))$ by Chan and Gupta [13].
- A suboptimal k versus $\Omega(\frac{n}{2^{6\lfloor k/2 \rfloor}}\alpha_k(n))$ by Le et al. [38].

This paper establishes the optimal k versus $\Omega(n\alpha_k(n))$ lower bound tradeoff for stretch $1+\epsilon$, for any $\epsilon > 0$, and for any k. An important conceptual contribution of this work is in achieving optimality by shaving off an extremely slowly growing term, namely $2^{6\lfloor k/2 \rfloor}$ for $k \leq O(\alpha(n))$; such a fine-grained optimization (that achieves optimality) is very rare in the literature.

To shave off the $2^{6\lfloor k/2 \rfloor}$ term from the previous bound of Le et al., our argument has to drill much deeper. In particular, we propose a new way of analyzing recurrences that involve inverse-Ackermann style functions, and our key technical contribution is in presenting the **first explicit construction of concave versions** of these functions. An important advantage of our approach over previous ones is its *robustness*: While all previous lower bounds are applicable only to restricted 1-dimensional point sets, ours applies even to *random* point sets in constant-dimensional spaces.

2012 ACM Subject Classification Theory of computation → Sparsification and spanners

Keywords and phrases Euclidean spanners, Ackermann functions, convex functions, hop-diameter

Digital Object Identifier 10.4230/LIPIcs.SoCG.2023.47

Funding Hung Le is supported by the National Science Foundation under Grants No. CCF-2121952 and No. CCF-2237288. Lazar Milenković and Shay Solomon are co-funded by the European Union (ERC, DynOpt, 101043159). Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Council. Neither the European Union nor the granting authority can be held responsible for them. This research was also supported by the Israel Science Foundation (ISF) grant No.1991/1, and by a grant from the United States-Israel Binational Science Foundation (BSF), Jerusalem, Israel, and the United States National Science Foundation (NSF).

1 Introduction

Let P be a set of n points in $\mathbb{R}^d$ and let $G_P = (P, \binom{P}{2}, \|\cdot\|)$ be the complete weighted graph induced by P, which contains an edge (p,q) of weight $w(p,q) = \|p-q\|$, for every $p,q \in S$. A subgraph graph $G = (P, E, \|\cdot\|)$ of G_P, $E \subseteq \binom{P}{2}$, is called a *geometric graph*. For a parameter $t \geq 1$, a geometric graph G is called a *t-spanner* for P if, for all $p,q \in S$, G contains a *t-spanner path* between p and q (i.e., a path of *weight* at most $t\|p-q\|$).

Euclidean spanners have been studied extensively [17, 35, 5, 12, 19, 6, 20, 8, 47, 2, 13, 21, 51, 53, 23, 39, 32, 38]. They are important in theory and practice, having found many applications, e.g., in geometric approximation algorithms, network topology design, and distributed computing [19, 40, 47, 26, 28, 27, 31, 41]; see also the book [42].

The most basic requirement of a spanner is to be *sparse*, while achieving small stretch. Cornerstone results settle the stretch-size tradoeff: for any d-dimensional n-point Euclidean space and for any $\epsilon > 0$, there exists a $(1+\epsilon)$-spanner with $O_{\epsilon,d}(n)$ edges [58, 18, 34, 48, 35, 5], where the $O_{\epsilon,d}$ suppresses the dependence on ϵ and d. (More precisely, the size upper bound is $n \cdot O(\epsilon^{-d+1})$, and it was shown to be tight [39].) In many applications, however, the spanner should have additional useful properties of the underlying metric. One such property is the *(hop-)diameter*: a t-spanner for P has (hop-)diameter of k if, for any $p,q \in S$, there is a t-spanner path between p and q with at most k edges (or hops). Having a small diameter is important for various applications (e.g., routing protocols) [7, 1, 2, 13, 21, 32].

While the stretch-size tradeoff is fully understood including the dependence on ϵ and d, the extended tradeoff of stretch-size-diameter is not fully understood yet even for fixed ϵ and d. Our goal is to achieve a full understanding of this tradeoff for fixed ϵ and d.

If the points are in general position, a 1-spanner must include basically all $\binom{n}{2}$ edges of the underlying metric. For points lying on a line, the simple path connecting them provides 1-spanner, but its diameter is worst-possible, $n-1$. Surprisingly perhaps, *all* previous lower bounds for the stretch-size-diameter tradeoff apply to line metrics. Understanding line metrics, and more generally tree metrics, is also important from the upper bounds front. In particular, the problem of constructing sparse 1-spanners with bounded diameter for line and tree metrics is closely related to several other fundamental problems. As an example, consider the extremely well-studied problem of partial sums, where we are given an array A of semigroup elements $A[1], \ldots, A[n]$ and are asked to construct a small-sized data structure, so that given a query i,j for $1 \leq i < j \leq n$, the partial sum $\sum_{i \leq k \leq j} A[k]$ can be computed efficiently. A 1-spanner for the corresponding set $A[1], \ldots, A[n]$ with bounded diameter is basically what we are looking for: A 1-spanner path between $A[i]$ and $A[j]$ that consists of at most k edges can be used for answering a query i,j within time $O(k)$. Other closely related problems include the tree product queries in semigroup problem (a generalization of partial sums) and its variants (see [55, 58, 4, 16, 46, 2], and the references therein), the MST verification problem [37, 36, 44], and the problem of shortcutting digraphs [56, 57, 10].

1.1 Previous Work on Spanners with Tiny diameter

1.1.1 Upper bounds

1-spanners for line and tree metrics. Let $T = (T, rt)$ be a (possibly weighted) n-vertex rooted tree, and let M_T be the tree metric induced by T. A spanning subgraph G of M_T is said to be a *1-spanner* for T, if for every pair of vertices, their distance in G is equal to their distance in T. One can define t-spanners for T, with $t \geq 1$, but essentially all previous work here concerned stretch 1. Alon and Schieber [4] showed that for any

n-point tree metric, a 1-spanner with diameter 2 (respectively, 3) and $O(n \log n)$ edges (resp., $O(n \log \log n)$ edges) can be built within time linear in its size; for $k \geq 4$, they showed that 1-spanners with diameter at most $2k$ and $O(n\alpha_k(n))$ edges can be built in $O(n\alpha_k(n))$ time. The function α_k is the inverse of a certain Ackermann-style function at the $\lfloor k/2 \rfloor$th level of the primitive recursive hierarchy, where $\alpha_0(n) = \lceil n/2 \rceil$, $\alpha_1(n) = \lceil \sqrt{n} \rceil$, $\alpha_2(n) = \lceil \log n \rceil$, $\alpha_3(n) = \lceil \log \log n \rceil$, $\alpha_4(n) = \log^* n$, $\alpha_5(n) = \left\lfloor \frac{1}{2} \log^* n \right\rfloor$, etc. Roughly speaking, for $k \geq 2$ the function α_k is close to $\lfloor \frac{k-2}{2} \rfloor$-iterated log-star function, i.e., log with $\lfloor \frac{k-2}{2} \rfloor$ stars. Also, $\alpha_{2\alpha(n)+2}(n) \leq 4$, where $\alpha(n)$ is the one-parameter inverse Ackermann function, which is an extremely slowly growing function. (See [38] for a formal definition.) Bodlaender et al. [11] constructed 1-spanners with diameter at most k and $O(n\alpha_k(n))$ edges, but for $k \geq 4$ their construction time is rather high ($\Omega(n^2)$). Solomon [52] gave a linear-time construction with the same diameter-size tradeoff k versus $O(n\alpha_k(n))$ as [11].

Alternative constructions, by Yao [58] for line metrics and by Chazelle [14] for general tree metrics, achieve a tradeoff of m edges versus diameter $\Theta(\alpha(m, n))$, where $\alpha(m, n)$ is the two-parameter inverse-Ackermann function (defined in [38]). However, these constructions provide 1-spanners with diameter $\Gamma' \cdot k$, only for constant $\Gamma' > 30$.

1.1.1.1 $(1 + \epsilon)$-spanners

The seminal STOC'95 of Arya et al. [6] established the "Dumbbell Theorem": For any d-dimensional Euclidean space, a $(1+\epsilon, O(\frac{\log(1/\epsilon)}{\epsilon^d}))$-tree cover can be constructed in $O(\frac{\log(1/\epsilon)}{\epsilon^d} \cdot n \log n + \frac{1}{\epsilon^{2d}} \cdot n) = O_{\epsilon,d}(n \log n)$ time. (For the definition of *tree cover*, see e.g. [32].) The consequence of the Dummbell Theorem is that any construction of 1-spanners for tree metrics can be tranformed into a construction of Euclidean $(1 + \epsilon)$-spanners, and the running time of the transformation is $O_{\epsilon,d}(n \log n)$ (plus a linear term in the size bound of the 1-spanner construction). The construction of 1-spanners for tree metrics from [52] thus yields an $O(n \log n)$-time construction of Euclidean $(1 + \epsilon)$-spanners with diameter k and $O(n\alpha_k(n))$ edges. Moreover, this result of [52] generalizes for the wider family of *doubling metrics* via the recent tree cover theorem of Bartal et al. [9].

1.1.2 Lower bounds

The celebrated work of Yao [58] provided the first lower bound on 1-spanners for tree metrics, where a tradeoff of m edges versus diameter of $\Omega(\alpha(m, n))$ was proved for the uniform line metric. A stronger lower bound on 1-spanners, still for the uniform line metric, was given in [4]: diameter k versus $\Omega(n\alpha_k(n))$ edges, for any k; as shown in [38], the lower bound of [4] implies that of [58], but the converse isn't true. These lower bounds apply only to 1-spanners.

Chan and Gupta [13] extended the lower bound of [58] to $(1 + \epsilon)$-spanners, still for line metrics, proving a lower bound tradeoff of m edges versus diameter of $\Omega(\alpha(m, n))$. This tradeoff only provides a meaningful lower bound for sufficiently large values of diameter (above say 30). Specifically, the result of [13] can be used to show that any $(1 + \epsilon)$-spanner for a certain line metric with diameter at most k must have $\Omega(n\alpha_{2k+6}(n))$ edges. When $k = 2$ (resp. $k = 3$), this gives $\Omega(n \log^{****} n)$ (resp. $\Omega(n \log^{*****} n)$) edges, which is far from the upper bound of $O(n \log n)$ (resp., $O(n \log \log n)$).

In SoCG'22 Le et al. [38] gave the following suboptimal lower bound tradeoff, for $(1 + \epsilon)$-spanners of the uniform line metric: k versus $\Omega(\frac{n}{2^{6\lfloor k/2 \rfloor}}\alpha_k(n))$. While the result of [38] is tight for constant k, the following question remains open for more than three decades:

▶ **Question 1.1.** *Is there a lower bound of k versus $\Omega(n\alpha_k(n))$ between the diameter and the number of edges,* for all k, *for Euclidean $(1 + \epsilon)$-spanners?*

1.1.2.1 Putting Question 1.1 into perspective.

Question 1.1 has been answered affirmatively by [38] for constant values of k. Recall that $\alpha_{2\alpha(n)+4}(n) \leq 4$, where $\alpha(n)$ is the one-parameter inverse Ackermann function. In other words, the gap underlying Question 1.1 holds only for $k = \omega(1), \ldots, O(\alpha(n))$, which is admittedly a very small regime. The gap itself is exponential in k, which is at most exponential in $\alpha(n)$, hence it is a very small gap.

One might wonder – why is Question 1.1 of any interest? Indeed, from a quantitative perspective, $\alpha(n)$ grows asymptotically even more slowly than $\log^* n$, which, in turn, is at most 5 for $n < 2^{65536}$. Thus a gap of $\exp(\alpha(n))$ is a constant factor gap for all practical purposes. However, we argue that Question 1.1 is important from a qualitative perspective. Indeed, there are numerous breakthrough works whose "only goal" was to shave off factors that grow as slowly as inverse-Ackermann type functions. For example, for the Union-Find data structure, efforts to achieve a linear time algorithm led to a lower bound showing that inverse-Ackermann function dependence is necessary [25], matching Tarjan's cornerstone upper bound [54]. Another prime example is in the context of the MST problem, where the inverse-Ackermann function dependence was shaved off from the upper bound of [15] to achieve a linear time algorithm by means of randomization [33] or under certain assumptions [24]; and it remains a major question whether there exists a linear time deterministic comparison-based MST algorithm. Yet another example is in the context of Davenport-Schinzel sequences, whose study involves optimizing inverse-Ackermann style functions – including the functions $\alpha(n)$ and $\alpha_k(n)$ – has led to important advances in discrete and computational Geometry. Indeed, Davenport and Schinzel [29] gave sharp bounds on sequences of order 1 and 2, namely $\lambda_1(n) = n$ and $\lambda_2(n) = 2n - 1$, and since then numerous applications of the sequences have been found, such as to geometric containment problems, computing shortest paths, and convex hulls. Achieving a tight bound for order-3 sequences spanned a long line of work [29, 22, 30, 43], and it is now understood that $\lambda_3(s) = 2n\alpha(n) + O(n\sqrt{\alpha(n)})$, i.e., the asymptotic behavior is known up to the leading constant. The case for $k \geq 4$ also spanned much work [22, 30, 49, 50, 3, 43, 45] and was settled up to leading constants in front of $\alpha(n)$ in the exponent, i.e., $\lambda_4(n) = \Theta(n2^{\alpha(n)})$, $\lambda_5(n) = \Theta(n\alpha(n)2^{\alpha(n)})$, $\lambda_6(n) = 2^{(1+o(1))\alpha^t(n)/t!}$.

We stress that in this work we are not merely shaving off an inverse-Ackermann function dependence slack from a previous upper bound (that of [38]) – we shave off such a slack *to achieve a tight bound.* This is a rare example where such a tiny slack is shaved to achieve optimality, and we believe that it is a significant evidence for the importance of our result, especially in light of our technical contribution.

1.1.2.2 A robust lower bound?

All previous lower bounds [58, 4, 13, 38] apply to very specific line metrics: either to the uniform line metric [58, 4, 38] or to one that is derived from hierarchically well-separated trees (HSTs) and is very far from being uniform [13].

A natural question is whether one can improve the longstanding construction of Euclidean $(1+\epsilon)$-spanners by Arya et al. [6] for "typical" point sets, which arise in real-life applications – such as *random points in low-dimensional spaces.* While random point sets are important from a practical perspective, none of the previous lower bounds [58, 4, 13, 38] precludes the existence of improved spanner constructions for such point sets.

▶ **Question 1.2.** *Can one improve the k versus $O(n\alpha_k(n))$ longstanding upper bound by Arya et al. [6] for random point sets in constant-dimensional Euclidean space?*

1.2 Our Contribution

1.2.1 The basic lower bound (settling Question 1.1 in the affirmative)

We prove that any $(1+\epsilon)$-spanner for the uniform line metric with diameter k has $\Omega(n\alpha_k(n))$ edges, for any k. We first prove the following general statement, which applies to subspaces of the uniform line metrics of any *density*.

▶ **Theorem 1.** *Let P be a set of p points in the interval $[0, L]$ such that every unit sub-interval $[i, i+1]$ for integer i, $1 \leq i \leq L-1$ contains at most 1 point of P. For any $\epsilon \in [0, 1/4]$ and integer $k \geq 1$, any $(1+\epsilon)$-spanner with diameter k for P contains $\Omega\left((p^2/L)\alpha_k(p)\right)$ edges.*

For technical reasons we prove a more general lower bound, stated in Lemmas 12, 14, and 16, which applies to *Steiner spanners*, namely, spanners that may contain additional Steiner points. The following direct corollary of Theorem 1 improves the previous lower bound by Le et al. [38] by a factor of $2^{\Omega(k)}$, and it settles Question 1.1 in the affirmative.

▶ **Corollary 2** (The longstanding upper bound is tight for all k). *Let $P = \{0, 1, \ldots, n-1\}$ be the set of n points on the uniform line metric contained on interval $[0, n)$. For any $\epsilon \in [0, 1/4]$ and integer $k \geq 1$, any $(1+\epsilon)$-spanner with diameter k for P contains $\Omega(n\alpha_k(n))$ edges.*

1.2.2 A robust lower bound (settling Question 1.2 in the negative)

Our lower bound of Theorem 1 applies to subspaces of the uniform line metric. We first demonstrate that this lower bound can be naturally extended to obtain analogs for *constant* dimensions. Second, we show that this lower bound carries over for random point sets in spaces of constant dimension, thereby settling Question 1.2 in the negative. We note that our approach seamlessly extends to higher constant dimensions.

The constant-dimensional hypercube and grid

The proof of the following theorem is omitted from this version due to space constraints.

▶ **Theorem 3.** *Let P be a set of p points in the hypercube $[0, L]^d$ for a constant $d \geq 2$ and some integer $L \geq 0$ such that every unit hypercube with integer vertices in $[0, L]^d$ contains at most one point of P. For any $\epsilon \in [0, 1/4]$ and any integer $k \geq 1$, any $(1+\epsilon)$-spanner with diameter k for P contains $\Omega\left((p^d/L^d)\alpha_k(p^d)\right)$ edges.*

Thus for $d = 2$ and $d = 3$, we get lower bounds $\Omega((p^2/L^2)\alpha_k(p^2))$ and $\Omega((p^3/L^3)\alpha_k(p^3))$.

▶ **Corollary 4.** *Let P be the set of n^d points on the d-dimensional grid $[0, n)^d$, for a constant $d \geq 2$. Then, for any $\epsilon \in [0, 1/4]$ and any integer $k \geq 1$, any $(1+\epsilon)$-spanner with diameter k for P contains $\Omega\left(n^d\alpha_k(n^d)\right)$ edges.*

Random point sets in the d-dimensional hypercube

We omit the proof of the following theorem from this version due to space constraints.

▶ **Theorem 5.** *Let P be a set of n points sampled uniformly at random on the hypercube $[0, 1]^d$ for any constant $d \geq 1$. For any $\epsilon \in [0, 1/4]$, and any integer $k \geq 1$, any $(1+\epsilon)$-spanner with diameter k for P contains $\Omega(n\alpha_k(n))$ edges.*

Remark. Theorem 5 applies to $d = 1$ as well, i.e., random points on the unit interval $[0, 1]$.

1.2.3 A concave inverse-Ackermann function

Our technique for proving Theorem 1 requires a significantly deeper understanding of inverse-Ackermann style functions than used in previous works [58, 4, 13, 38]. A key technical contribution in our work is an explicit construction of continuous versions of these functions. To our knowledge, this work is the first to introduce such functions for $\alpha_k(n)$ for $k > 4$. We then show that these functions are *concave*, which allows us to apply Jensen's inequality in our inductive proof, leading to a lower bound that is not only optimal for all values of k, but is also more robust, and in particular precludes the existence of better constructions for random point sets.

▶ **Theorem 6.** *Fix an arbitrary constant $\frac{1}{10000} \le \Delta \le \frac{1}{256}$. There exists a family of functions $\{f_k(x) : k \ge 2, k \in \mathbb{Z}\}$ such that each $f_k : \mathbb{R}^+ \to \mathbb{R}^+$ is twice differentiable in $(0, +\infty)$ and:*
1. *For $x > 1$, $f_2(x) = \log x$; $f_3(x) = \log\log x$; and $f_k(x) = \Delta + f_k(f_{k-2}(x))$ for every $k \ge 4$.*
2. *For all $x \in \mathbb{R}^{\ge 1}$ and $k \ge 4$, function $x^2 f_k(x)$ is convex.*
3. *For all $x \in \mathbb{R}^{\ge 0}$ and $k \ge 4$, it holds that $f_k(x) \ge \frac{\Delta}{5}\alpha_k(\lceil x \rceil) - 1$.*
4. *For all $x \in \mathbb{R}^{\ge 0}$ and $k \ge 2$, it holds that $f_k(\lceil x \rceil) \le \alpha_k(\lceil x \rceil)$.*
5. *For all $k \ge 2$, $k \in \mathbb{Z}$ and $x \ge 200$, $x \in \mathbb{R}$, it holds that $2\left\lfloor \frac{x}{f_k(x)} \right\rfloor f_k\left(\frac{2\lfloor x/f_k(x)\rfloor}{4}\right) \ge x/2$.*

Items 3 and 4 of Theorem 6 imply that $f_k(n) = \Theta(\alpha_k(n))$. Item 2 is a key property of our function $f_k(x)$, which does not hold for its discrete counterpart $\alpha_k(n)$.

2 One-dimensional instances

This section is dedicated to proving Theorem 1. The proof is by double induction on the number of points and the diameter of the spanner. There are two base cases in the proof: $k = 2$ and $k = 3$ presented in Section 2.2 and Section 2.3, respectively. The proof for $k \ge 4$ is given in Section 2.4. Together, they imply Theorem 1. We choose $\Delta = 1/256$.

For a constant d and given set of points P on the d-dimensional hypercube $[0, L]^d$, we require that every unit hypercube with integer vertices in $[0, L]^d$ contains at most one point in P. We call the condition *unit interval condition*.

2.1 Classification of cross edges

Given a point set P contained on an interval $[0, n]$ and given an $\epsilon \in [0, 1/4]$, let H be any $(1+\epsilon)$-spanner for P. Consider Algorithm 1 with parameter $\ell = 0$ being the *recursion level*, k being the diameter, and I being the interval containing P. This algorithm is used to classify the edges of H only. It divides I into a smaller set of b subintervals and defines a set of separators, which are the endpoints of the subintervals excluding the two endpoints of I. A cross edge of the interval I at level ℓ is an edge (1) needed to preserve the distance between two points in P and (2) crossing a separator.

Next we study properties of cross edges and classify them.

▶ **Lemma 7.** *Let e be a cross edge of some interval $I = [c, d]$ and let $L := |d - c|$ denote the interval length. Then, both endpoints of e are within $[c - L/4, d + L/4]$.*

Proof. Suppose toward the contradiction that there is an edge containing an endpoint outside of $[c - L/4, d + L/4]$. Without loss of generality we take the case where the right endpoint of e has coordinate larger than $d + L/4$. Let $x < y$ be two points in I for which e is on their

Algorithm 1 Procedure describing the terms used in the proof. It is initially invoked with a given set of p points P on interval I, and $\ell = 0$. Here, H is a $(1+\epsilon)$-spanner for P.

procedure CrossEdges($P, I = [c,d], k, H, \ell$)
  if $(k \le 3$ **and** $p \le 1)$ **or** $(k \ge 4$ **and** $f_k(p) < 1)$ **then return**
  Let $b \leftarrow 2$ if $k = 2$, $b \leftarrow \lceil \sqrt{p} \rceil$ if $k = 3$, and $b \leftarrow 2 \cdot \lfloor p/f_{k-2}(p) \rfloor$ otherwise.
  $M \leftarrow (d-c)/b$ ▷ dividing I into b subintervals
  for $1 \le j \le b-1$ **do**
    $I_j \leftarrow [c+(j-1)M, c+jM]$
    $P_j \leftarrow P \cap [c+(j-1)M, c+jM)$
  $P_b \leftarrow P \cap [c+(b-1)M, c+bM]$
  Let $\{c + jM \mid 1 \le j \le b-1\}$ be the set of *separators* of I.
  A *cross edge* of interval I is every edge $e = (x,y)$ of H such that: *(i)* e is on some $(1+\epsilon)$-spanner path between two points in P and *(ii)* there exists a separator s such that $x \le s \le y$.
  for $1 \le j \le b$ **do** CrossEdges($P_j, I_j, k, H, \ell+1$)

$(1+\epsilon)$ spanner path, say $\pi_{x,y}$, in H. Since $\pi_{x,y}$ is a $(1+\epsilon)$-spanner path, its length $\pi_{x,y}$ must be at most $(1+\epsilon)|y-x| \le 5|y-x|/4 \le |y-x| + L/4$. However, the length of $\pi_{x,y}$ is strictly greater than $|x-y| + L/4$ since the right endpoint of e is larger than $d + L/4$, a contradiction. ◀

We classify the cross edges as follows. We call a cross edge of some interval *interior* if it contains both endpoints inside of the interval. If both of its endpoints are outside of the interval, we call it *exterior*. Otherwise, we call it *mixed*. See Figure 1 for an illustration.

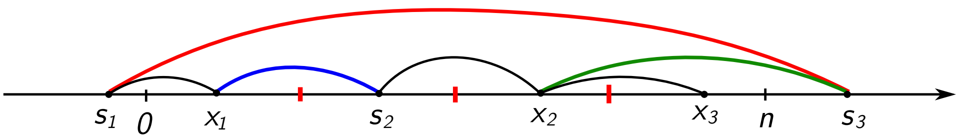

Figure 1 The separators are marked by short red lines. Here $P = \{x_1, x_2, x_3\}$. The spanner could use Steiner points which are points not in P; they are s_1, s_2, s_3 in this figure. The red edge (s_1, s_3) is an exterior cross edge of $[0,n]$; the blue edge (x_1, s_2) is an interior cross edge of $[0,n]$; and the green edge (x_2, s_3) is a mixed cross edge of $[0,n]$. Edge (s_1, x_1) is not a mixed edge since it does not cross any separator.

▶ **Lemma 8.** *Let e be an interior cross edge of some interval. Then, it cannot be an interior cross edge of any other interval.*

Proof. Let ℓ be the level at which e is an interior cross edge of some interval $I = [c,d]$. By definition, e cannot be an interior cross edge of any other interval at level ℓ, since the intervals at the same level are disjoint. Since the intervals of levels lower than ℓ contain no separators inside $[c,d]$, e cannot be a cross edge at these levels. Finally, after level ℓ, I is split at the separators into smaller intervals, and hence e cannot have two endpoints in the same interval at levels higher than ℓ. ◀

▶ **Lemma 9.** *Let e be an exterior cross edge of some interval. Then, it cannot be an exterior cross edge of any other interval.*

Proof. Suppose that $e = (u, v)$ is an exterior cross edge of more than one interval. Among such intervals, let $[c, d]$ be of the highest level, say ℓ. We have that $u < c$ and $d < v$ since e is exterior. Let $L = |d - c|$. The length of intervals at levels lower than ℓ are at least $2L$. From Lemma 7, we know that $c - L/4 \leq u$ and $v \leq d + L/2$, so the length of e is at most $3L/2$. This means that e cannot be an exterior edge at levels lower than ℓ. ◀

▶ **Lemma 10.** *Let e be a mixed cross edge of some interval. Then, it can be a mixed cross edge for at most one other interval, an exterior cross edge for at most one other interval and an interior cross edge for at most one other interval.*

Proof. Let ℓ be the level at which $e = (u, v)$ is a mixed cross edge of some interval $I = [c, d]$ of length $L := |d - c|$. Without loss of generality, we assume that $u \in [c, d]$ and $v \geq d$. By Lemma 7, $v < d + L/4$. Let $I' = [c', d']$ be another interval such that $I' \neq I$ and e is a cross edge of I'. We consider three cases.

If the level of I' is strictly smaller than ℓ. If d is not a separator of I', then by definition e cannot be a cross edge of I'. If d is a separator of I', then $d' \geq d + L$. On the other hand, $v < d + L/4$, so $e = (u, v)$ must be an interior cross edge of I'. By Lemma 8, it cannot be an interior cross edge of any other interval.

If the level of I' is exactly ℓ. I. Since $v < d + L/4$, the only case where e is a cross edge of I' is that d is the left endpoint of I' and e is a mixed cross edge of I'. Thus, e could not be a mixed cross edge of any other interval at level ℓ.

If the level of I' is strictly larger than ℓ. Then the length of I' is at most L/b. Since u is on the left of at least one separator, say s, of I and $v > d$, the distance between s and d is at least L/b. It follows that the length of e is at least L/b. Hence, the only possible way for e to be a cross edge of I' is that it is an exterior cross edge. By Lemma 9, e' will not be an exterior cross edge of any other interval. ◀

▶ **Corollary 11.** *Every cross edge considered in the process above is counted at most 4 times.*

In the sequel, we will be proving the lower bound on the number of cross edges. We say that a point of P in an interval I is *global* if it is incident on at least one cross edge of I. Otherwise, we say that it is *non-global*.

2.2 Hop-diameter 2

In this section, we show one of the two base cases of our inductive proof: a lower bound for diameter $k = 2$.

▶ **Lemma 12.** *Let P be a set of $p \geq 2$ points in the interval $[0, L]$ satisfying the unit interval condition. For any $\epsilon \in [0, 1/4]$, any Steiner $(1 + \epsilon)$-spanner for P with diameter 2 contains at least $T_2(p, L) \geq \frac{p^2 \log p}{16L}$ edges.*

Proof. Our proof is by induction on the number of points in P. Let H be any $(1 + \epsilon)$-spanner for P with diameter 2. We split the interval $[0, L]$ into two disjoint intervals $[0, L/2]$ and $[L/2, L]$. Let the number of points in the intervals be $p_1 := |P \cap [0, L/2]|$ and $p_2 := |P \cap (L/2, L]|$. We claim that the number of edges of H can be lower bounded by $T_2(p, L)$ which satisfies:

$$T_2(p, L) \geq T_2(p_1, L/2) + T_2(p_2, L/2) + \min(p_1, p_2)/4 \tag{1}$$

The base cases are $T_2(0, L_0) = T_2(1, L_0) = 0$, for any $L_0 > 0$. The terms $T_2(p_1, L/2)$ (resp., $T_2(p_2, L/2)$) come from the cross edges contributed by the intervals in $[0, L/2]$ (resp., $[L/2, L]$)

and their recursive divisions in Algorithm 1. We will show in Claim 13 that the number of cross edges of $[0, L]$ is at least $\min(p_1, p_2)$. By Corollary 11, each cross edge is counted at most 4 times. Thus, we use $\min(p_1, p_2)/4$ in Equation (1). This implies that the number of edges of H is bounded by $T_2(p, L)$.

▷ Claim 13. H contains at least $\min(p_1, p_2)$ cross edges of the interval $[0, L]$.

Proof. Without loss of generality, assume $p_1 \le p_2$. For contradiction, assume that the number of cross edges is less than p_1. This means that there is a non-global point a in $[0, L/2]$. (Recall that we call a point non-global if it is not incident on any cross edge of the interval $[0, L]$.) A path from a to any point b in $P \cap [L/2, L]$ is of the form (a, a_b, b), where a_b is a point on the left of $L/2$. Then (a_b, b) is a cross edge by definition. That is, for each point in $P \cap [L/2, L]$, there is a corresponding cross edge in the path to a. Thus, $[0, L]$ contains $p_2 \ge p_1$ different cross edges, which is a contradiction. ◁

We now solve the recurrence in Equation (1). We prove by induction that $T_2(p, L) \ge \frac{p^2 \log p}{16L}$. Note that $L \ge p$ by the unit interval conditoin in Lemma 12. Assume without loss of generality that $p_1 \le p_2$. First, we assume that $p_1 \ge p/4$.

$$T_2(p, L) \ge T_2(p_1, L/2) + T_2(p_2, L/2) + \frac{p_1}{4} \ge \frac{p_1^2 \log p_1}{8L} + \frac{p_2^2 \log p_2}{8L} + \frac{p_1}{4} \ge \frac{p^2 \log(p/2)}{16L} + \frac{p_1}{4}$$
$$= \frac{p^2(\log(p) - 1) + 4Lp_1}{16L} \ge \frac{p^2(\log(p) - 1) + 4pp_1}{16L} \ge \frac{p^2 \log p}{16L} (\text{since } p_1 \ge p/4)$$

The second inequality follows by induction hypothesis, third by Jensen's inequality, fourth by the unit interval condition, and the fifth since $p1 \ge p/4$. When $p_1 < p/4$, we have the following.

$$T_2(p, L) \ge T_2(p_1, L/2) + T_2(p_2, L/2) + \frac{p_1}{4} \ge T_2(p_2, L/2) \ge \frac{(3p/4)^2 \log(3p/4)}{8L} \ge \frac{p^2 \log p}{16L}$$

The penultimate inequality follows by using $p_2 \ge 3p/4$ and the induction hypothesis, whereas the last one holds for all $p \ge 14$. When $2 \le p \le 13$, we use $T_2(p_1, L/2) + T_2(p_2, L/2) + p_1/4 \ge \frac{p_1^2 \log p_1}{8L} + \frac{p_2^2 \log p_2}{8L} + \frac{p_1}{4} \ge \frac{p^2 \log p}{16L}$, where the last inequality can be manually verified. The lemma now follows. ◀

2.3 Hop-diameter 3

In this section, we show the remaining base case of our inductive proof: a lower bound for diameter $k = 3$.

▶ **Lemma 14.** *Let P be a set of $p \ge 2$ points in the interval $[0, L]$ satisfying the unit interval condition. For any $\epsilon \in [0, 1/4]$, any Steiner $(1 + \epsilon)$-spanner for P with diameter 3 contains at least $T_3(p, L) \ge \frac{p^2 \log \log p}{800L}$ edges.*

Proof. Let H be any $(1 + \epsilon)$-spanner for P with diameter 3. We split the interval $[0, L]$ into $b := \lceil \sqrt{p} \rceil$ disjoint intervals of length L/b: $[0, L/b], [L/b, 2(L/b)], \dots, [(b - 1)(L/b), L]$. Let $P_i = P \cap [(i - 1)(L/b), i(L/b))$ for $1 \le i < b$ and $P_b = P \cap [L - L/b, L]$. In other words, we divide the interval as in Algorithm 1. Let the number of points in the i-th interval be denoted by $p_i := |P_i|$. We claim that the number of edges of H can be lower bounded by $T_3(p, L)$ which satisfies:

$$T_3(p, L) \ge \sum_{i=1}^{b} T_3(p_i, L/b) + |E_C|/4 \tag{2}$$

Here E_C denotes the set of cross edges for the interval $[0, L]$ and the term $T_3(p_i, l_i)$, where $1 \le i \le b$, is the lower bound on the number of cross edges of H at higher levels restricted to preserving distances in P_i. By Corollary 11, each cross edge is counted at most 4 times. Thus, we use $|E_C|/4$ in Equation (2). Thus, $|E(H)| \ge T_3(p, L)$. The base cases are $T_3(0, L_0) = T_3(1, L_0) = 0$, for any $L_0 > 0$.

We now inductively show that $T_3(p, L) \ge \frac{p^2 \log\log p}{800L}$. Suppose first that there is a collection of $c \le \sqrt{p}/2$ intervals which in total contain at least $9p/10$ points. Without loss of generality, assume that these are the first c intervals; that is, $\sum_{i=1}^{c} b_i = 9p/10$. In this case, we show that the inequality holds even without the contribution of the cross edges.

$$\sum_{i=1}^{c} T_3(p_i, L/b) \ge \sum_{i=1}^{c} \frac{p_i^2 \log\log p_i}{800L/b} \ge c \cdot \frac{\left(\frac{9p}{10c}\right)^2 \log\log\left(\frac{9p}{10c}\right)}{800L/b}$$
$$\ge \frac{81}{50} \cdot \frac{p^2 \log\log\left(\frac{9\sqrt{p}}{5}\right)}{800L} \ge \frac{p^2 \log\log p}{800L}$$

The first inequality follows from the induction hypothesis, second by Jensen's inequality, and third using $b \ge \sqrt{p}$ and $c \le \sqrt{p}/2$. We next bound the number of cross edges in the complementary case.

▷ Claim 15. Assume that there is no collection of $c \le \sqrt{p}/2$ intervals that in total contain at least $9p/10$ points. Then, $|E_c| \ge p/100$.

Proof. Suppose first there are at least $p/10$ global points. The number of cross edges they contribute is at least $p/20$, since each edge can be counted at most twice. In the complementary regime, there are at least $9p/10$ non-global points. By the assumption of the claim, we know that they are contained in at least $\sqrt{p}/2$ intervals. Consider two non-global points x and y contained in two different intervals, X and Y, respectively. Since x and y are non-global, i.e., they are not incident on any cross edge, every 3-hop path between x and y must be of the form $\langle x, x', y', y\rangle$, where $x' \in X$ and $y' \in Y$. We conclude that every pair of different intervals containing non-global points induces a different cross edge. Hence, the number of cross edges can be lower bounded by $\binom{\sqrt{p}/2}{2} \ge \frac{p}{100}$ for $p \ge 5$. When $2 \le p \le 4$, there is at least one cross edge, and the bound holds as well. ◁

We now solve Equation (2) by induction. By Claim 15, we have:

$$T_3(p, L) \ge \sum_{i=1}^{b} T_3(p_i, L/b) + \frac{p}{400} \ge \sum_{i=1}^{b} \frac{p_i^2 \log\log p_i}{800L/b} + \frac{p}{400} \ge b \cdot \frac{(p/b)^2 \log\log(p/b)}{800L/b} + \frac{p}{400}$$
$$= \frac{p^2 \log\log(p/b)}{800L} + \frac{p}{400} = \frac{p^2 \log\log(p/b) + 2pL}{800L} \ge \frac{p^2 \log\log p}{800L}$$

The second inequality follows from the induction hypothesis, third by Jensen's inequality, and the last from the unit interval condition and the choice $b = \lceil \sqrt{p} \rceil$. The lemma now follows. ◀

2.4 Hop-diameter $k \ge 4$

In this section, we show a lower bound for $k \ge 4$, concluding the proof of Theorem 1. Our proof will use function $f_k(x)$ in Theorem 6 with $\Delta = 1/256$. In particular, we will show the lower bound $\Omega(\frac{p^2 f_k(p)}{L})$ on the number of edges. Since $f_k(p) = \Omega(\alpha_k(b))$ by Item 3 of Theorem 6, the number of edges of the spanner is $\Omega(\frac{p^2 \alpha_k(p)}{L})$ as claimed in Theorem 1.

▶ **Lemma 16.** *Let P be a set of $p \geq 2$ points in the interval $[0, L]$ satisfying the unit interval condition. For any $\epsilon \in [0, 1/4]$, any Steiner $(1+\epsilon)$-spanner for P with hop-diameter $k \geq 2$ contains at least $T_k(p, n) \geq \frac{p^2 f_k(p)}{800L}$ edges.*

Proof. The base cases $k = 2$ and $k = 3$ follow from the definition of $f_2(x) = \log x$ and $f_3(x) = \log\log x$ and Lemmas 12 and 14. The base case for p happens when $f_k(p) < 1$. Here, we use the fact that any spanner on p points must have at least $p-1$ edges and $p-1 \geq \frac{p^2 f_k(p)}{800L}$ so the claim follows.

Let H be any $(1+\epsilon)$-spanner for P with hop-diameter k. We split the interval $[0, L]$ into $b := 2 \cdot \lfloor p/f_{k-2}(p) \rfloor$ disjoint intervals of length L/b: $I_1 = [0, L/b), I_2 = [L/b, 2(L/b)), \ldots, I_{b-1} = [(b-2)(L/b), (b-1)(L/b)), I_b = [(b-1)(L/b), L]$. Let the number of points in the i-th interval be denoted by $p_i := |P \cap I_b|$. By the same proof of Lemma 14, the number of edges of H can be lower bounded by $T_k(p, n)$ which satisfies:

$$T_k(p, L) \geq \sum_{i=1}^{b} T_k(p_i, L/b) + |E_C|/4 \tag{3}$$

Here E_C denotes the set of cross edges for the interval $[0, L]$ and the term $T_k(p_i, L/b)$, where $1 \leq i \leq b$, come from the cross edges contributed by the i-th interval and its recursive subdivisions.

We now inductively show that $T_k(p, L) \geq \frac{p^2 f_k(p)}{800L}$ for $k \geq 4$. Suppose first that there is a collection of $c \leq b/4$ intervals that in total contain at least $3p/4$ points. Then the inequality holds even without considering $|E_C|$. Recall that by Item 2 in Theorem 6, $x^2 fk(x)$ is convex and hence we can apply the Jensen's inequality.

$$\begin{aligned}
T_k(p, L) &\geq \sum_{i=1}^{c} T_k(p_i, L/b) \geq \sum_{i=1}^{c} \frac{p_i^2 f_k(p_i)}{800L/b} \\
&\geq c \cdot \frac{\left(\frac{3p}{4c}\right)^2 f_k\left(\frac{3p}{4c}\right)}{800L/b} && \text{(Jensen's inequality)} \\
&\geq \frac{9}{4} \cdot \frac{p^2 f_k\left(\frac{3p}{b}\right)}{800L} \geq \frac{9}{4} \cdot \frac{p^2 f_k\left(f_{k-2}(p)\right)}{800L} && \text{(using } c \leq b/4 \text{ and } b := 2 \cdot \lfloor p/f_{k-2}(p) \rfloor) \\
&= \frac{9}{4} \cdot \frac{p^2 (f_k(p) - \Delta)}{800L} && \text{(by Item 1 in Theorem 6)} \\
&\geq \frac{p^2 f_k(p)}{800L} && \text{using that } f_k(p) \geq 1
\end{aligned}$$

Now we consider the complementary case where there is no collection of $c \leq b/4$ intervals that in total contain at least $3p/4$ points. For this case, we need to take the number of cross edges into account.

▷ Claim 17. Assume that there is no collection of $c \leq b/4$ intervals that in total contain at least $3p/4$ points. Then, $|E_C| \geq p/25600$.

Proof. If there is at least $p/4$ global points, then we have at least $p/8$ cross edges. In the complementary regime, there are at least $3p/4$ non-global points. By assumption, they are contained in at least $b/4$ non-global blocks. From each interval that contains non-global points we take exactly one non-global point and let the resulting set of points be denoted P'. We use the induction hypothesis with $k-2$ on P'. Note that $|P'| \geq b/4$. The following observation allows us to use the scaled version of the induction hypothesis.

▶ **Observation 18.** *Suppose that a set of points P' on interval $[0, L]$ satisfies that when we divide $[0, L]$ into consecutive intervals of length M, every such interval contains at most one point from P' and H' be any $(1+\epsilon)$ spanner of P' with hop-diameter k. Let Q' be a set of points in P' scaled down by a factor of L. Such a set of points is contained on interval $[0, L/M]$ and it satisfies the unit interval condition. Let H'' be the scaled version of H'. Then, H'' is a $(1+\epsilon)$-spanner for Q' with hop-diameter k.*

We proceed to lower bound the number of cross edges, using the observation.

$$T_{k-2}(|P'|, b) \geq T_{k-2}\left(\frac{b}{4}, b\right) \geq \frac{\frac{b^2}{16} f_{k-2}\left(\frac{b}{4}\right)}{800b} = \frac{2\lfloor p/f_{k-2}(p)\rfloor f_{k-2}\left(\frac{2\lfloor p/f_{k-2}(p)\rfloor}{4}\right)}{12800} \geq \frac{p}{25600}$$

The second inequality follows by the induction hypothesis for $k-2$, and the last by Item 5 in Theorem 6. This concludes the proof of Claim 17. ◁

We now solve Equation (3) by induction. Recall that we choose $\Delta = 1/256$. By Claim 17, we have:

$$\begin{aligned}
T_k(p, L) &\geq \sum_{i=1}^{b} T_k(p_i, L/b) + \frac{p}{102400} \\
&\geq \sum_{i=1}^{b} \frac{p_i^2 f_k(p_i)}{800L/b} + \frac{p}{102400} && \text{(induction hypothesis)} \\
&\geq b \cdot \frac{\left(\frac{p}{b}\right)^2 f_k\left(\frac{p}{b}\right)}{800L/b} + \frac{p}{102400} && \text{(Jensen's inequality)} \\
&= \frac{p^2 f_k\left(\frac{p}{2\lfloor p/f_{k-2}(p)\rfloor}\right)}{800L} + \frac{p}{102400} && \text{(replacing } b := 2\lfloor p/f_{k-2}(p)\rfloor\text{)} \\
&\geq \frac{p^2(f_k(p) - 3\Delta)}{800L} + \frac{p}{102400} \\
&\geq \frac{p^2 f_k(p)}{800L} && \text{(using } p \leq L \text{ and } \Delta = 1/256\text{)}
\end{aligned}$$

The lemma now follows. ◀

3 Concave Ackermann-type functions

In this section, we introduce the concave inverse-Ackermann function $f_k(x)$. We omit the details from this extended abstract due to space constraints. We fix a constant $\Delta < 1/256$.

▶ **Definition 19** ($f_k(n)$ for even k). *For all $x \in \mathbb{R}^{\geq 0}$ and even $k \geq 2$, we let $f_k(x)$ be:*

$$\begin{aligned}
&f_2(x) = \log x \\
&f_k(x) = a_k x^3 + b_k x^2 + c_k x - \Delta && \textit{for } 0 \leq x \leq 1, k \geq 4 \\
&f_k(x) = \Delta + f_k(f_{k-2}(x)) && \textit{for } x > 1, k \geq 4
\end{aligned}$$

Constants a_k, b_k, and c_k are chosen so that they satisfy the following relations.

$$a_k + b_k + c_k = \Delta \quad \forall k \geq 4 \tag{4}$$

$$3a_4 + 2b_4 + c_4 = \frac{c_4}{\ln 2} \tag{5}$$

$$6a_4 + 2b_4 = \frac{2b_4 - c_4 \ln 2}{\ln^2 2} \tag{6}$$

$$3a_k + 2b_k + c_k = c_k \cdot (3a_{k-2} + 2b_{k-2} + c_{k-2}) \tag{7}$$

$$6a_k + 2b_k = 2b_k \cdot (3a_{k-2} + 2b_{k-2} + c_{k-2})^2 + c_k \cdot (6a_{k-2} + 2b_{k-2}) \tag{8}$$

In this section, we solve the recurrence in Definition 19 for even k by giving estimates on the values of a_k, b_k and c_k. We will use these estimates in the proof of Theorem 6, which is omitted from this extended abstract due to space constraints.

For $k = 4$, by solving a linear system of equations defined by Equations (4), (5), and (8) we obtain the following estimates.

▶ **Lemma 20.** *a_4, b_4 and c_4 satisfy the following equation:*

$$\begin{aligned} -0.0819\Delta \leq a_4 &\leq -0.0818\Delta \\ 0.2966\Delta \leq b_4 &\leq 0.2967\Delta \\ 0.7852\Delta \leq c_4 &\leq 0.7853\Delta \end{aligned} \tag{9}$$

In estimating the values of a_k, b_k and c_k, we will use the following sequences:

$$\lambda_4 = 1.1328\Delta, \qquad \lambda_k = \frac{3\Delta\lambda_{k-2}}{1 + 4\lambda_{k-2}}$$

$$r_4 = 11.0439, \qquad r_k = \frac{\Lambda_{k-2}^3 + \Lambda_{k-2}}{2\Lambda_{k-2}^3 - 2\Lambda_{k-2}^2 + \frac{\Lambda_{k-2}}{r_{k-2}}}, \text{ where } \Lambda_k = 0.3777 \cdot (3\Delta)^{(k-2)/2} \tag{10}$$

▶ **Lemma 21.** *$\lambda_k \geq 0.3265(3\Delta)^{\frac{k-2}{2}}$ and $r_k < 25$ for all $k \geq 4$.*

Proof. Solving the recurrence we get

$$\begin{aligned} \lambda_k &= \frac{236 \cdot (1 - 3\Delta) \cdot 3^{(k-2)/2}}{(625 + 957\Delta)\left(\frac{1}{\Delta}\right)^{(k-2)/2} - 944 \cdot 3^{(k-2)/2}} \\ &\geq \frac{236 \cdot (1 - 3\Delta) \cdot (3\Delta)^{(k-2)/2}}{625 + 957\Delta} \geq 0.3265 \cdot (3\Delta)^{\frac{k-2}{2}}. \end{aligned}$$

The last inequality holds whenever $\Delta \leq 1/32$.

We use induction to show that $r_k < 25$; the base case holds by definition of r_4. Observe that $0 \leq \Lambda_k \leq \Lambda_4 \leq 0.3777 \cdot (3\Delta) \leq 0.3777 \cdot \frac{3}{256}$. By induction, $r_{k-2} < 25$. Thus, we have $r_k = \frac{\Lambda_{k-2}^3 + \Lambda_{k-2}}{2\Lambda_{k-2}^3 - 2\Lambda_{k-2}^2 + \frac{\Lambda_{k-2}}{r_{k-2}}} \leq r_{k-2} + 3000\Lambda_{k-2}$, where the last inequality follows since the left-hand side grows with Λ_{k-2} for all $0 \leq \Lambda_k \leq 0.3777 \cdot \frac{3}{256}$, when $11.0439 \leq r_{k-2} < 25$. It follows that: $r_k \leq r_4 + 3000 \sum_{i=1}^{(k-2)/2} \Lambda_k \leq r_4 + 3000 \sum_{i=1}^{\infty} \Lambda_k \leq 11.0439 + 3000 \cdot 0.3777 \cdot \frac{3\Delta}{1-3\Delta} < 25$, as desired. ◀

▶ **Lemma 22.** *Let $X_k = 2a_k + b_k + \Delta$ and $Y_k = 6a_k + 2b_k$. Then*

$$0.3265 \cdot (3\Delta)^{(k-2)/2} \leq \lambda_k \leq X_k \leq 0.3777 \cdot (3\Delta)^{(k-2)/2} \tag{11}$$

$$11.041 \leq \frac{X_k}{Y_k} \leq r_k < 25 \tag{12}$$

$$\Delta - X_k \leq a_k \leq \Delta - X_k + \frac{X_k}{22} \tag{13}$$

$$-3\Delta + 3X_k - \frac{X_k}{11} \leq b_k \leq -3\Delta + 3X_k \tag{14}$$

$$3\Delta - 2X_k \leq c_k \leq 3\Delta - 2X_k + \frac{X_k}{22} \tag{15}$$

Proof. Observe by Equation (4) that $X_k = 2a_k + b_k + \Delta = 3a_k + 2b_k + c_k$ and that $c_k = \frac{Y_k}{2} - 2X_k + 3\Delta$. Thus, the system from Definition 19 for $k \geq 6$ can be written as follows.

$$X_k = X_{k-2} \cdot (\frac{Y_k}{2} - 2X_k + 3\Delta)$$
$$Y_k = X_{k-2}^2 \cdot (6X_k - 2Y_k - 6\Delta) + Y_{k-2} \cdot (\frac{Y_k}{2} - 2X_k + 3\Delta)$$

Solving the above system of equations for X_k and Y_k, we get:

$$X_k = \frac{\begin{vmatrix} 6X_{k-2}\Delta & -X_{k-2} \\ (6Y_{k-2} - 12X_{k-2}^2)\Delta & (4X_{k-2}^2 + 2 - Y_{k-2}) \end{vmatrix}}{\begin{vmatrix} (4X_{k-2}+2) & -X_{k-2} \\ (4Y_{k-2} - 12X_{k-2}^2) & (4X_{k-2}^2 + 2 - Y_{k-2}) \end{vmatrix}} = \frac{6\Delta(X_{k-2}^3 + X_{k-2})}{2X_{k-2}^3 + 4X_{k-2}^2 + 4X_{k-2} - Y_{k-2} + 2}$$
$$Y_k = \frac{\begin{vmatrix} (4X_{k-2}+2) & 6X_{k-2}\Delta \\ (4Y_{k-2} - 12X_{k-2}^2) & (6Y_{k-2} - 12X_{k-2}^2)\Delta \end{vmatrix}}{\begin{vmatrix} (4X_{k-2}+2) & -X_{k-2} \\ (4Y_{k-2} - 12X_{k-2}^2) & (4X_{k-2}^2 + 2 - Y_{k-2}) \end{vmatrix}} = \frac{6\Delta(2X_{k-2}^3 - 2X_{k-2}^2 + Y_{k-2})}{2X_{k-2}^3 + 4X_{k-2}^2 + 4X_{k-2} - Y_{k-2} + 2} \tag{16}$$

For the base case, $X_4 = 2a_4 + b_4 + \Delta$ and $Y_4 = 6a_4 + 2b_4$. By Lemma 20, we have:

$$1.1328\Delta \leq X_4 \leq 1.1331\Delta \quad \text{and} \quad 0.1018\Delta \leq Y_4 \leq 0.1026\Delta \tag{17}$$

Next, we show both Equation (11) and Equation (12) by induction; the base case ($k = 4$) holds by Equation (17). By Equation (16), we have: $X_k \leq \frac{6\Delta(X_{k-2}^3+X_{k-2})}{2X_{k-2}^3+4X_{k-2}^2+4X_{k-2}-\frac{X_{k-2}}{11.041}+2} \leq 3\Delta X_{k-2} \leq 3\Delta \cdot 0.3777 \cdot (3\Delta)^{(k-4)/2} = 0.3777 \cdot (3\Delta)^{(k-2)/2}$ The lower bound on X_k follows also by induction: $X_k = \frac{6\Delta(X_{k-2}^3+X_{k-2})}{2X_{k-2}^3+4X_{k-2}^2+4X_{k-2}-Y_{k-2}+2} \geq \frac{3\Delta X_{k-2}}{1+4X_{k-2}} \geq \frac{3\Delta\lambda_{k-2}}{1+4\lambda_{k-2}} = \lambda_k$, by Equation (10). For the lower bound on $\frac{X_k}{Y_k}$, by Equation (16), we have: $\frac{X_k}{Y_k} = \frac{X_{k-2}^3+X_{k-2}}{2X_{k-2}^3-2X_{k-2}^2+Y_{k-2}} \geq \frac{X_{k-2}^3+X_{k-2}}{2X_{k-2}^3-2X_{k-2}^2+\frac{X_{k-2}}{11.041}} \geq 11.041$, where the last inequality holds since $X_{k-2} \leq 1.1331\Delta \leq \frac{1.1331}{256}$. Finally, we show an upper bound on $\frac{X_k}{Y_k} = \frac{X_{k-2}^3+X_{k-2}}{2X_{k-2}^3-2X_{k-2}^2+Y_{k-2}} \leq \frac{X_{k-2}^3+X_{k-2}}{2X_{k-2}^3-2X_{k-2}^2+\frac{X_{k-2}}{r_{k-2}}} \leq \frac{\Lambda_{k-2}^3+\Lambda_{k-2}}{2\Lambda_{k-2}^3-2\Lambda_{k-2}^2+\frac{\Lambda_{k-2}}{r_{k-2}}} = r_k < 25$, by Lemma 21. This concludes the inductive proof of Equation (11) and Equation (12). For Equations (13)–(15), we express a_k, b_k, and c_k in terms of X_k and Y_k as follows: $a_k = \Delta + \frac{Y_k}{2} - X_k$, $b_k = -3\Delta + 3X_k - Y_k$, and $c_k = 3\Delta + \frac{Y_k}{2} - 2X_k$. Eq. (13)-(15) follow. ◀

References

1 I. Abraham and D. Malkhi. Compact routing on Euclidean metrics. In *Proc. of 23rd PODC*, pages 141–149, 2004.

2 P. K. Agarwal, Y. Wang, and P. Yin. Lower bound for sparse Euclidean spanners. In *Proc. of 16th SODA*, pages 670–671, 2005.

3 Pankaj K. Agarwal, Micha Sharir, and Peter W. Shor. Sharp upper and lower bounds on the length of general davenport-schinzel sequences. *J. Comb. Theory, Ser. A*, 52(2):228–274, 1989.

4 N. Alon and B. Schieber. Optimal preprocessing for answering on-line product queries. *Manuscript*, 1987.

5 I. Althöfer, G. Das, D. P. Dobkin, D. Joseph, and J. Soares. On sparse spanners of weighted graphs. *Discrete & Computational Geometry*, 9:81–100, 1993.

6 S. Arya, G. Das, D. M. Mount, J. S. Salowe, and M. H. M. Smid. Euclidean spanners: short, thin, and lanky. In *Proc. of 27th STOC*, pages 489–498, 1995.

7 S. Arya, D. M. Mount, and M. H. M. Smid. Randomized and deterministic algorithms for geometric spanners of small diameter. In *Proc. of 35th FOCS*, pages 703–712, 1994.

8 S. Arya and M. H. M. Smid. Efficient construction of a bounded degree spanner with low weight. *Algorithmica*, 17(1):33–54, 1997.

9 Yair Bartal, Nova Fandina, and Ofer Neiman. Covering metric spaces by few trees. In *ICALP*, volume 132 of *LIPIcs*, pages 20:1–20:16. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019.

10 Arnab Bhattacharyya, Elena Grigorescu, Kyomin Jung, Sofya Raskhodnikova, and David P. Woodruff. Transitive-closure spanners. *SIAM J. Comput.*, 41(6):1380–1425, 2012.

11 H. L. Bodlaender, G. Tel, and N. Santoro. Trade-offs in non-reversing diameter. *Nord. J. Comput.*, 1(1):111–134, 1994.

12 P. B. Callahan and S. R. Kosaraju. Faster algorithms for some geometric graph problems in higher dimensions. In *Proc. of 4th SODA*, pages 291–300, 1993.

13 H. T.-H. Chan and A. Gupta. Small hop-diameter sparse spanners for doubling metrics. In *Proc. of 17th SODA*, pages 70–78, 2006.

14 B. Chazelle. Computing on a free tree via complexity-preserving mappings. *Algorithmica*, 2:337–361, 1987.

15 B. Chazelle. A minimum spanning tree algorithm with inverse-ackermann type complexity. *Journal of the ACM*, 47(6):1028–1047, 2000. doi:10.1145/355541.355562.

16 B. Chazelle and B. Rosenberg. The complexity of computing partial sums off-line. *Int. J. Comput. Geom. Appl.*, 1:33–45, 1991.

17 L. P. Chew. There is a planar graph almost as good as the complete graph. In *Proc. of 2nd SOCG*, pages 169–177, 1986.

18 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.

19 G. Das and G. Narasimhan. A fast algorithm for constructing sparse Euclidean spanners. In *Proc. of 10th SOCG*, pages 132–139, 1994.

20 G. Das, G. Narasimhan, and J. S. Salowe. A new way to weigh malnourished Euclidean graphs. In *Proc. of 6th SODA*, pages 215–222, 1995.

21 Y. Dinitz, M. Elkin, and S. Solomon. Low-light trees, and tight lower bounds for Euclidean spanners. *Discrete & Computational Geometry*, 43(4):736–783, 2010.

22 E. Szemerédi. On a problem of Davenport and Schinzel. *Acta Arith.*, 25, 1973.

23 Michael Elkin and Shay Solomon. Optimal euclidean spanners: Really short, thin, and lanky. *J. ACM*, 62(5):35:1–35:45, 2015.

24 M. L. Fredman and D. E. Willard. Trans-dichotomous algorithms for minimum spanning trees and shortest paths. *Journal of Computer and System Sciences*, 48(3):533–551, 1994. Announced at FOCS'90. doi:10.1016/s0022-0000(05)80064-9.

25 Michael L. Fredman and Michael E. Saks. The cell probe complexity of dynamic data structures. In *STOC*, pages 345–354. ACM, 1989.

26 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.

27 J. Gudmundsson, C. Levcopoulos, G. Narasimhan, and M. H. M. Smid. Approximate distance oracles for geometric spanners. *ACM Transactions on Algorithms*, 4(1), 2008.

28 J. Gudmundsson, G. Narasimhan, and M. H. M. Smid. Fast pruning of geometric spanners. In *Proc. of 22nd STACS*, pages 508–520, 2005.

29 H. Davenport and A. Schinzel. A combinatorial problem connected with differential equations. *Am. J. Math*, 1965.

30 Sergiu Hart and Micha Sharir. Nonlinearity of davenport - schinzel sequences and of generalized path compression schemes. *Comb.*, 6(2):151–178, 1986.

31 Y. Hassin and D. Peleg. Sparse communication networks and efficient routing in the plane. In *Proc. of 19th PODC*, pages 41–50, 2000.

32 Omri Kahalon, Hung Le, Lazar Milenkovic, and Shay Solomon. Can't see the forest for the trees: Navigating metric spaces by bounded hop-diameter spanners. In *PODC*, pages 151–162. ACM, 2022.

33 D. R. Karger, P. N. Klein, and R. E. Tarjan. A randomized linear-time algorithm to find minimum spanning trees. *Journal of the ACM*, 42(2):321–328, 1995. doi:10.1145/201019.201022.

34 J. M. Keil. Approximating the complete Euclidean graph. In *Proc. of 1st SWAT*, pages 208–213, 1988.

35 J. M. Keil and C. A. Gutwin. Classes of graphs which approximate the complete Euclidean graph. *Discrete & Computational Geometry*, 7:13–28, 1992.

36 Valerie King. A simpler minimum spanning tree verification algorithm. *Algorithmica*, 18(2):263–270, 1997.

37 János Komlós. Linear verification for spanning trees. *Comb.*, 5(1):57–65, 1985.

38 Hung Le, Lazar Milenkovic, and Shay Solomon. Sparse euclidean spanners with tiny diameter: A tight lower bound. In *SoCG*, volume 224 of *LIPIcs*, pages 54:1–54:15. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2022.

39 Hung Le and Shay Solomon. Truly optimal euclidean spanners. In *FOCS*, pages 1078–1100. IEEE Computer Society, 2019.

40 C. Levcopoulos, G. Narasimhan, and M. H. M. Smid. Efficient algorithms for constructing fault-tolerant geometric spanners. In *Proc. of 30th STOC*, pages 186–195, 1998.

41 Y. Mansour and D. Peleg. An approximation algorithm for min-cost network design. *DIMACS Series in Discr. Math and TCS*, 53:97–106, 2000.

42 G. Narasimhan and M. Smid. *Geometric Spanner Networks.* Cambridge University Press, 2007.

43 Gabriel Nivasch. Improved bounds and new techniques for davenport-schinzel sequences and their generalizations. *J. ACM*, 57(3):17:1–17:44, 2010.

44 Seth Pettie. An inverse-ackermann type lower bound for online minimum spanning tree verification. *Comb.*, 26(2):207–230, 2006.

45 Seth Pettie. Sharp bounds on davenport-schinzel sequences of every order. *J. ACM*, 62(5):36:1–36:40, 2015.

46 M. Pătraşcu and E. D. Demaine. Tight bounds for the partial-sums problem. In *Proc. of 15th SODA*, pages 20–29, 2004.

47 S. Rao and W. D. Smith. Approximating geometrical graphs via "spanners" and "banyans". In *Proc. of 30th STOC*, pages 540–550, 1998.

48 J. Ruppert and R. Seidel. Approximating the d-dimensional complete Euclidean graph. In *Proc. of 3rd CCCG*, pages 207–210, 1991.

49 Micha Sharir. Almost linear upper bounds on the length of general davenport-schinzel sequences. *Comb.*, 7(1):131–143, 1987.

50 Micha Sharir. Improved lower bounds on the length of davenport - schinzel sequences. *Comb.*, 8(1):117–124, 1988.

51 S. Solomon and M. Elkin. Balancing degree, diameter and weight in Euclidean spanners. In *Proc. of 18th ESA, Part 1*, pages 48–59, 2010.

52 Shay Solomon. Sparse euclidean spanners with tiny diameter. *ACM Trans. Algorithms*, 9(3):28:1–28:33, 2013.

53 Shay Solomon. From hierarchical partitions to hierarchical covers: optimal fault-tolerant spanners for doubling metrics. In *STOC*, pages 363–372. ACM, 2014.

54 R. E. Tarjan. Efficiency of a good but not linear set union algorithm. *J. ACM*, 22(2):215–225, 1975.

55 R. E. Tarjan. Applications of path compression on balanced trees. *J. ACM*, 26(4):690–715, 1979.

56 Mikkel Thorup. Shortcutting planar digraphs. *Comb. Probab. Comput.*, 4:287–315, 1995.

57 Mikkel Thorup. Parallel shortcutting of rooted trees. *J. Algorithms*, 23(1):139–159, 1997.

58 A. C. Yao. Space-time tradeoff for answering range queries. In *Proc. of 14th STOC*, pages 128–136, 1982.