Fractal Dimension and Lower Bounds for Geometric Problems

Anastasios Sidiropoulos
Department of Computer Science, University of Illinois at Chicago
Chicago IL, USA
sidiropo@uic.edu

Kritika Singhal
Department of Mathematics, The Ohio State University
Columbus OH, USA
singhal.53@osu.edu

Vijay Sridhar
Department of Computer Science and Engineering, The Ohio State University
Columbus OH, USA
sridhar.38@osu.edu

──── **Abstract** ────

We study the complexity of geometric problems on spaces of low *fractal dimension*. It was recently shown by [Sidiropoulos & Sridhar, SoCG 2017] that several problems admit improved solutions when the input is a pointset in Euclidean space with fractal dimension smaller than the ambient dimension. In this paper we prove nearly-matching lower bounds, thus establishing nearly-optimal bounds for various problems as a function of the fractal dimension.

More specifically, we show that for any set of n points in d-dimensional Euclidean space, of fractal dimension $\delta \in (1, d)$, for any $\varepsilon > 0$ and $c \geq 1$, any c-spanner must have treewidth at least $\Omega\left(\frac{n^{1-1/(\delta-\epsilon)}}{c^{d-1}}\right)$, matching the previous upper bound. The construction used to prove this lower bound, can also be used to derive lower bounds on the running time of algorithms for various problems, assuming the Exponential Time Hypothesis. We provide two prototypical results of this type:

- For any $\delta \in (1, d)$ and any $\varepsilon > 0$, d-dimensional Euclidean TSP on n points with fractal dimension at most δ cannot be solved in time $2^{O\left(n^{1-1/(\delta-\varepsilon)}\right)}$. The best-known upper bound is $2^{O(n^{1-1/\delta}\log n)}$.

- For any $\delta \in (1, d)$ and any $\varepsilon > 0$, the problem of finding k-pairwise non-intersecting d-dimensional unit balls/axis parallel unit cubes with centers having fractal dimension at most δ cannot be solved in time $f(k)n^{O\left(k^{1-1/(\delta-\varepsilon)}\right)}$ for any computable function f. The best-known upper bound is $n^{O(k^{1-1/\delta}\log n)}$.

The above results nearly match previously known upper bounds from [Sidiropoulos & Sridhar, SoCG 2017], and generalize analogous lower bounds for the case of ambient dimension due to [Marx & Sidiropoulos, SoCG 2014].

1 Introduction

The *curse of dimensionality* is a general phenomenon in computational geometry, asserting that the complexity of many problems increases rapidly with the dimension of the input. Sets of fractional dimension can be used to model various processes and phenomena in science and engineering [25]. Recently, the complexity of various geometric optimization problems was studied as a function of the fractal dimension of the input [23]. It was shown that, for several problems, improved algorithms can be obtained when the fractal dimension is smaller than the ambient dimension.

Interestingly, the algorithms obtained in [23] nearly match the best-possible algorithms for integral dimension. In this paper, we give nearly-matching lower bounds, assuming the Exponential Time Hypothesis (ETH). We remark that there are several different definitions of fractal dimension that can be considered. Our results indicate that, for the case of Euclidean pointsets, the definition of fractal dimension we consider is the "correct" one for certain computational problems. That is, it precisely generalizes the dependence of the running time on the ambient dimension.

1.1 Our contribution

We obtain nearly-optimal lower bounds for various prototypical geometric problems. Our results are obtained via a general method that could be applicable to other problems.

Spanners

We begin with a lower bound on the treewidth of spanners. It is known that any set of n points in $\mathbb{R}^d$ admits a $(1 + \varepsilon)$-spanner of size $n(1/\varepsilon)^{O(d)}$ [22, 27]. This result has been generalized for the case of fractal dimension. Specifically, it was shown in [23] that any n-point set in $O(1)$-dimensional Euclidean space, of fractal dimension $\delta > 1$, admits a $(1 + \varepsilon)$-spanner of size $n(1/\varepsilon)^{O(d)}$, and of pathwidth $O(n^{1-1/\delta} \log n)$. We show the following lower bound, which establishes that the upper bound from [23] is essentially best-possible.

▶ **Theorem 1.** *Let $d \geq 2$ be an integer. Then for all $\delta \in (1, d)$, for all $\varepsilon > 0$ and for all $n_0 \in \mathbb{N}$, there exists a set of $n \geq n_0$ points $P \subset \mathbb{R}^d$, of fractal dimension at most δ', where $|\delta - \delta'| \leq \varepsilon$, such that for any $c \geq 1$, any c-spanner G of P has $tw(G) = \Omega\left(\frac{n^{1-1/\delta'}}{c^{d-1}}\right)$.*

Independent Set of Unit Balls

We consider the k-Independent Set of Unit Balls in $\mathbb{R}^d$, which is a prototypical geometric optimization problem, parameterized by the optimum. In this problem given a set of n unit balls in $\mathbb{R}^d$, we seek to find a set of k pairwise non-intersecting balls. It is known that this problem can be solved in time $n^{O(k^{1-1/d})}$, for any $d \geq 2$ [1, 19], and that there is no algorithm with running time $f(k)n^{o(k^{1-1/d})}$, for any computable function f, assuming ETH [19] (see also [17]). The upper bound has been generalized for fractal dimension as follows: It has been shown that when the set of centers of the balls has fractal dimension δ, the problem can be solved in time $n^{O(k^{1-1/\delta} \log n)}$ [23]. We show the following lower bound on the running time, which nearly matches this upper bound, up to a logarithmic term.

▶ **Theorem 2.** *Let $d \geq 2$ be an integer, and let $1 < \delta' < \delta < d$. If for all k, and for some computable function f, there exists an $f(k)n^{o(k^{1-1/\delta'})}$ time algorithm for finding k pairwise non-intersecting open balls in a collection of n unit balls with the centers having fractal dimension at most δ, then ETH fails.*

Euclidean TSP

It is known that TSP on a set of n points in d-dimensional Euclidean space can be solved in time $2^{1-1/d}n^{O(1)}$ [24], and that there is no algorithm with running time $2^{O(n^{1-1/d-\varepsilon})}$, for any $\varepsilon > 0$, assuming ETH. The upper bound has been generalized to the case of fractal dimension as follows. It has been shown that for a set of fractal dimension $\delta > 1$, in $O(1)$-dimensional Euclidean space, TSP can be solved in time $2^{O(n^{1-1/\delta} \log n)}$. Here, we obtain the following nearly-tight lower bound.

▶ **Theorem 3.** *Let $d \geq 2$ be an integer. Then, for all $\delta \in (2, d)$, for all $\delta' < \delta$ and for all $n_0 \in \mathbb{N}$, if there exists $n \geq n_0$ such that Euclidean TSP in $\mathbb{R}^d$ on all pointsets of size n and fractal dimension at most δ can be solved in time $2^{O(n^{1-1/\delta'})}$, then ETH fails.*

1.2 Overview of techniques

We now briefly highlight the main technical tools used in the paper.

High-level idea

We derive our lower bounds by adapting a method from the case of graph problems. It is known that, for many problems on graphs, large treewidth implies large running time lower bounds (see, e.g. [18]). This is typically done by exploiting the duality between treewidth and grid minors. Specifically, it is known that, for any r, graphs of treewidth at least $f(r)$, for some function f, have the $(r \times r)$-grid as a minor [21]. In fact, the function f is known to be linear for planar graphs [20], and polynomial in general [5]. One can often obtain a lower bound on the running time by using the grid minor to embed a large hard instance in the input. We apply the above approach to the geometric setting by relating fractal dimension to treewidth. Specifically, we construct a pointset in Euclidean space, such that any $O(1)$-spanner must have large treewidth.

From fractal dimension to treewidth

A main technical ingredient for obtaining nearly-optimal lower bounds is constructing pointsets such that the treewidth of any $O(1)$-spanner is as large as possible. For the case of exposition, we will describe the construction in the continuous case. We construct some $X \subset \mathbb{R}^d$, and we discretize X by taking some $O(1)$-approximate[1] ε-net N_ε of X. Let us refer to the fractal dimension of the resulting infinite family of nets N_ε, as the fractal dimension of X (see Section 1.4 for precise definitions). Our goal is to construct some X, with some fixed fractal dimension $\delta \in (1, d)$, such that the treewidth of any $O(1)$-spanner of N_ε is as large as possible as a function of $1/\varepsilon$.

---

[1] Recall that a $O(1)$-approximate r-net in some metric space (X, ρ) is some $N \subseteq X$, such that for all $x \neq y \in N$, $\rho(x, y) > r$, and for all $z \notin N$, $\rho(z, N) = \inf_{y \in N} d(z, y) = O(r)$.
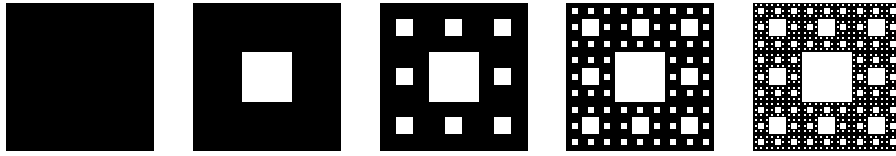
Figure 1 The first few iterations of the Sierpiński carpet.

Figure 2 The first few iterations of the Cantor ternary set.

### A first failed attempt: the Sierpiński carpet

Let us now briefly describe the construction and point out the main technical challenges. A natural first attempt in $\mathbb{R}^2$ is to let X be the Sierpiński carpet. This is a set obtained from the unit square by removing the central square of side length $1/3$, and by recursing on the remaining 8 sub-squares (see Figure 1). Unfortunately, this construction does not lead to a tight treewidth lower bound. Specifically, the resulting set has fractal dimension $\delta = \log 8 / \log 3$, while there exist $O(1)$-spanners of treewidth $O(n^{1-1/\gamma})$, where γ is a constant arbitrarily close to $\log 6 / \log 3$.

Intuitively, this happens for the following reason. Let S_ε be any $O(1)$-spanner for N_ε. Then, the largest possible grid minor in S_ε does not use most of the vertices in S_ε. Thus, roughly speaking, we can obtain a larger grid minor by constructing a set X so that as few vertices of S_ε as possible are being "wasted".

Constructing a treewidth-extremal fractal: the Cantor crossbar

Using the above observation, we define the set X as follows. We first recall that the Cantor set $\mathcal{C}$ is obtained from the unit interval by removing the central interval of length $1/3$, and recursing on the other two (see Figure 2). We define $\mathcal{C}'$ to be the Cartesian product of $\mathcal{C}$ with $[0, 1]$, and we set X to be the union of two copies of $\mathcal{C}'$, where one is rotated by $\pi/2$. We refer to the resulting set as the *Cantor crossbar* (see Figure 3). We can show that the resulting set achieves a nearly-optimal treewidth lower bound.

The above construction can be generalized to the case where the ambient dimension is $d \geq 2$ as follows. Recall that, for any $d' \geq 1$, the *Cantor dust* in $\mathbb{R}^{d'}$, denoted by $\mathcal{D}_{d'}$, is the Cartesian product of d' copies of the Cantor set (see Figure 4). Let $e_1, \ldots, e_d$ be the standard orthonormal basis in $\mathbb{R}^d$. For each $i \in \{1, \ldots, d\}$, we define $\mathcal{T}_i$ to be the Cartesian product of $\mathcal{D}_{d-1}$ with $[0, 1]$, rotated so that $[0, 1]$ is parallel to e_i. Finally, we set $X = \mathcal{T}_1 \cup \ldots \cup \mathcal{T}_d$ (see Figure 5).
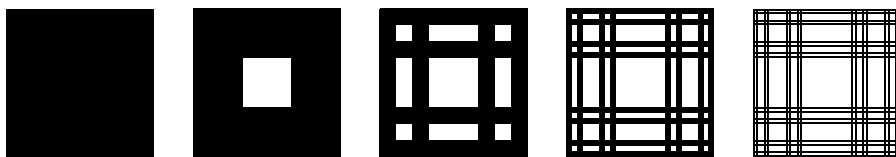
Figure 3 The first few iterations of the Cantor crossbar in $\mathbb{R}^2$.
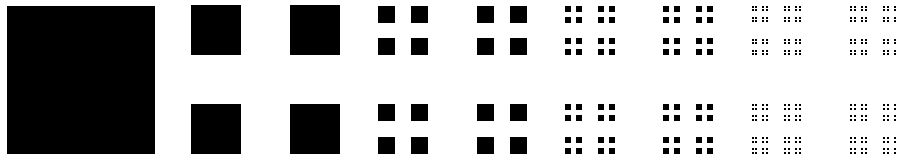
Figure 4 The first few iterations of the Cantor dust in $\mathbb{R}^2$.
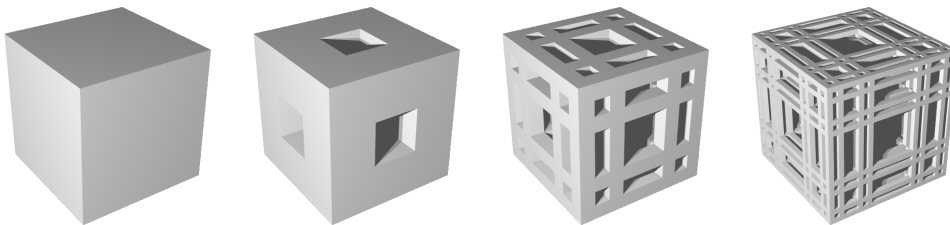
Figure 5 The first few iterations of the Cantor crossbar in $\mathbb{R}^3$.

The above construction gives a set with some fixed fractal dimension δ, for each fixed $d \geq 2$. We can generalize the construction so that δ attains any desired value in the range $(1, d]$. The only difference is that, when defining the Cantor dust, we start with a Cantor set of smaller dimension. This can be done by removing the central interval of length $\alpha \in (0, 1)$, instead of $1/3$, and recursing on the remaining two intervals of length $(1 - \alpha)/2$.

From spanner lower bounds to running time lower bounds

The above construction of the Cantor crossbar leads to a nearly-optimal lower bound for the treewidth of $O(1)$-spanners. We next use this construction to obtain running time lower bounds. Informally, a typical NP-hardness reduction for some geometric problem in the plane works as follows: One encodes some known computationally hard problem by constructing "gadgets" that are arranged in a grid-like fashion in $\mathbb{R}^2$ (see, e.g. [19]). More generally, for problems in $\mathbb{R}^d$, the gadgets are arranged along some d-dimensional grid. We follow a similar approach, with the main difference being that we arrange the gadgets along a Cantor crossbar.

1.3 Other related work

There has been a large body of work on determining the effect of *doubling* dimension on the complexity of various geometric problems [10, 2, 6, 14, 8, 16, 4, 3, 9, 26]. Other notions of dimension that have been considered include low-dimensional negatively curved spaces [15], growth-restricted metrics [12], as well as generalizations of doubling dimension to metrics of so-called bounded global growth [11]. In all of the above lines of research the goal is to extend tools and ideas from the Euclidean setting to more general geometries. In contrast, we study restricted classes of Euclidean instances, with the goal of obtaining better bounds than what is possible in the general case.

1.4 Preliminaries

We give some definitions that are used throughout the paper.

▶ **Definition 4** (ε-covering). Let X be a set and $A \subseteq X$. Then, A is an ε-covering of X if for every $x \in X$, there exists $a \in A$ such that $\text{dist}(x, a) \le \varepsilon$.

▶ **Definition 5** (ε-packing). A set A is called an ε-packing if for all $x, y \in A$, $\text{dist}(x, y) \ge \varepsilon$.

▶ **Definition 6** (ε-net). A set $A \subseteq X$ is called an ε-net of X if A is an ε-packing as well as an ε-covering of X.

▶ **Definition 7** (Fractal dimension). [23] Given $x \in \mathbb{R}^d$ and $r > 0$, let $B(x, r)$ denote the ball of radius r centered at x. The *fractal dimension* of some $P \subseteq \mathbb{R}^d$, denoted by $\dim_f(P)$, is defined as the infimum δ, such that for any $\varepsilon > 0$ and $r \ge 2\varepsilon$, for any ε-net N of P, and for any $x \in \mathbb{R}^d$, we have $|N \cap B(x, r)| = O((r/\varepsilon)^\delta)$.

We have the following lemmas showing invariance of fractal dimension under certain operations.

▶ **Lemma 8.** *Let $d \ge 1$ be an integer, let $0 < \delta \le d$ and let $c > 0$ be some constant. Let $P \subset \mathbb{R}^d$ be a pointset such that $\dim_f(P) = \delta$. Let P' be the pointset obtained by uniformly scaling the points of P about the origin by a factor of c. Then $\dim_f(P') = \delta$.*

▶ **Lemma 9.** *Let $d \ge 1$ be some integer, let $0 < \delta \le d$ and let $c > 0, k > 0$ be some constants. Let $P \subset \mathbb{R}^d$ be a pointset such that $\dim_f(P) = \delta$ and for all $u, v \in P$, $d(u, v) > 4c$. For all $p \in P$ let $S_p \subset \mathbb{R}^d$ be a set of points such that $|S_p| \le k$ and for all $x \in S_p$, $d(x, p) \le c$. Let $P' = \bigcup_{p \in P} S_p$. Then $\dim_f(P') = \delta$.*

▶ **Definition 10** (c-spanner). For any pointset $P \subset \mathbf{R}^d$, and for any $c \ge 1$, a *c-spanner* for P is a graph G with $V(G) = P$, such that for all $x, y \in P$, we have

$$||x - y||_2 \le d_G(x, y) \le c \cdot ||x - y||_2,$$

where d_G denotes the shortest path distance in G.

▶ **Definition 11** (Treewidth). [7] Let G be a graph, T a tree and $\mathcal{V} = \{V_t\}_{t \in T}$ be a family of vertex sets $V_t \subseteq V(G)$ indexed by the vertices t of T. The pair $(T, \mathcal{V})$ is called a tree-decomposition of G if it satisfies the following three conditions:
1. $V(G) = \cup_{t \in T} V_t$.
2. For every edge $e \in G$, there exists a $t \in T$ such that both ends of e lie in V_t.
3. $V_{t_1} \cap V_{t_3} \subseteq V_{t_2}$, whenever t_2 lies in the unique path joining t_1 and t_3 in T.
The width of $(T, \mathcal{V})$ is the number $\max\{|V_t| : t \in T\}$ and the *treewidth* of G is the least width of any tree-decomposition of G.

Exponential Time Hypothesis.   The Exponential Time Hypothesis states that there is no $2^{o(n)}$ algorithm to solve the n-variable 3SAT.

1.5   Organization

This paper is organized as follows. In section 2, we present lower bounds on the treewidth of spanners for arbitrary pointsets with integral dimension, and with fractal dimension. In section 3, we present running time lower bound on the Independent Set of Balls problem on pointsets with arbitrary fractal dimension in $\mathbf{R}^d$. The proof of Theorem 3 can be found in the full version of the paper. This theorem proves running time lower bound on the Euclidean TSP problem on pointsets with arbitrary fractal dimension in $\mathbf{R}^d$.

2 Lower bounds on the treewidth of spanners

In this section, we obtain lower bounds on the treewidth of spanners for arbitrary pointsets. In subsection 2.1, we consider pointsets with integral fractal dimension. In subsection 2.2, we consider a discretized version of the Sierpiński carpet whose fractal dimension is less than two but greater than 1. In subsection 2.3, we use a carefully chosen inductive construction to obtain a specific fractal pointset of fractal dimension $\frac{\log 6}{\log 3}$. This pointset gives us a nearly tight lower bound on the treewidth of a spanner. We finally generalize this construction in subsection 2.4 and present the proof of Theorem 1.

2.1 Treewidth and integral dimension

We obtain lower bounds on the treewidth of spanners for pointsets with integral fractal dimension. We will make use of the following theorem due to Kozawa *et al.*[13] for the proofs in this section.

▶ **Theorem 12** ([13]). *The treewidth of the d-dimensional grid on n vertices is $\Theta(n^{1-1/d})$.*

▶ **Theorem 13.** *For any integer $d \geq 1$, there exists a set of n points $P \subseteq \mathbb{R}^d$ such that for any $c \geq 1$, and for any c-spanner G of P, $\mathsf{tw}(G) = \Omega\left(\frac{n^{1-1/d}}{c^{d-1}}\right)$.*

Proof. Fix an integer $d \geq 1$. Let $n \in \mathbf{N}$ be such that $n^{\frac{1}{d}}$ is an integer. Let $P = \{1, 2, \cdots, n^{\frac{1}{d}}\}^d$ and let G be any c-spanner of P. Let $p = (p_1, \ldots, p_d)$ be a point in P. We define P' and X' as follows:

$P' = \{p \in P \mid \exists\ 1 \leq i \leq d \text{ such that } \forall\ j \neq i,\ (p_j \bmod (c+1)) = 1\}.$

$X = \{p \in P' \mid \forall\ 1 \leq i \leq d,\ (p_i \bmod (c+1)) = 1\}.$

Consider the points in P'. We call a row of points in P' that is parallel to one of the d axes a *full row* if this row consists of exactly $n^{\frac{1}{d}}$ points with adjacent points unit distance apart. Consider for any $1 \leq i \leq d$, the points of any pair of full rows R and T that are both parallel to the i^{th} axis. We have that for any pair of consecutive points x_1, x_2 in R and for any pair of consecutive points y_1, y_2 in T, no shortest path in G joining x_1 and x_2 can intersect any shortest path in G joining y_1 and y_2. Suppose not, then let z be a point of intersection between two such shortest paths. Since G is a c-spanner and consecutive points in any row are distance 1 apart, we have $d_G(x_1, z) + d_G(x_2, z) \leq c$ and similarly, $d_G(y_1, z) + d_G(y_2, z) \leq c$. But this implies that at least one of $d_G(x_1, y_1)$, $d_G(x_1, y_2)$, $d_G(x_2, y_1)$ and $d_G(x_2, y_2)$ is at most c due to triangle inequality. This is a contradiction because G is non-contracting and the distance between R and T is at least $c + 1$ by our choice of R and T.

Now if we consider a shortest path between every pair of consecutive points in the row R and concatenate these paths, remove all loops, then we can obtain a path from one end of this row to the other end. Doing the same for all full rows in P', we end up with a set of paths traversing the points in the full rows. Moreover from the earlier argument, it follows that any two such paths obtained from parallel full rows are vertex disjoint. Thus for all $1 \leq i \leq d$, we can obtain a set of vertex disjoint paths Q_i in G that traverse the points in the full rows of P' parallel to the i^{th} axis.

Finally we define a subgraph H of G as follows: H consists of the points in X. Furthermore, for any pair of points $p, q \in X$ such that p and q differ only along one coordinate, say the i^{th} coordinate, and differ by exactly $c + 1$, H also consists of the sub-path between p and q of the corresponding path in Q_i connecting p and q. Now contracting these paths in H between adjacent points in X results in a d-dimensional grid with $\frac{n}{\lceil (c+1) \rceil^d}$ points. Thus, we conclude that $\mathsf{tw}(G) = \Omega\left(\frac{n^{1-1/d}}{(c+1)^{d-1}}\right) = \Omega\left(\frac{n^{1-1/d}}{c^{d-1}}\right)$. ◀

2.2 A first attempt: The Sierpiński carpet

Consider a set of points X obtained by the following method: start with a $3^k \times 3^k$ integer grid for some $k \in \mathbb{N}$ and partition it into 9 subgrids of equal size. We delete all the points in the central subgrid and recurse on the remaining 8 subgrids. The recursion stops when we arrive at a subgrid containing a single point. This is a natural discrete variant of the Sierpiński carpet.

▶ **Theorem 14.** *Let δ be the fractal dimension of the set of points X obtained above. Then we have $\delta \geq \log_3 8$.*

Proof. We start by recalling the definition of fractal dimension of a set of points P. It is the infimum $\delta > 0$ such that for any $\varepsilon > 0$, for any ball B of radius $r \geq 2\varepsilon$ and for any ε-net N, we have $|B \cap N| = O((r/\varepsilon)^\delta)$.

As seen in the construction, the width of the grid reduces by $\frac{1}{3}$ at every step of recursion. Thus we may assume that the width is $\frac{1}{3^i}$ when we stop. Let $\varepsilon = \frac{1}{2} \cdot \frac{1}{3^i}$. Let $N = X$. We have $n = |X| = 8^i$ since every step of recursion is done on the remaining 8 subgrids. Let $r = \sqrt{2}$. Let x be the point on the top left corner of the grid. Then $|B(x, r) \cap N| = n = 8^i \leq m \cdot \left(\frac{r}{\varepsilon}\right)^\delta$ for some constant $m > 0$. This gives $8^i \leq m \cdot (\sqrt{2} \cdot 2 \cdot 3^i)^\delta$. Taking log on both sides, we get $\log 8 \leq \delta \log 3$. Thus, $\delta \geq \log_3 8$. ◀

▶ **Theorem 15.** *There exists a set of n points P with fractal dimension $\delta \in (1, 2)$, and some c-spanner G of P, where $c < 1 + \sqrt{2}$, such that $\mathsf{tw}(G) = \Theta(n^{1-\frac{1}{\delta}-\varepsilon})$, for some $\varepsilon > 0$.*

2.3 One treewidth-extremal fractal: A discretized Cantor crossbar

In this section we describe the construction of a pointset in $\mathbb{R}^2$ with a specific fractal dimension of $\frac{\log 6}{\log 3}$. This pointset is a discretized version of the *Cantor crossbar*. We generalize this construction in the next section to construct pointsets with arbitrary fractal dimension.

▶ **Theorem 16.** *Let $\delta = \frac{\log(6)}{\log(3)}$. Then for all $n_0 \in \mathbb{N}$, there exists a set $P \subset \mathbb{R}^2$ of $n \geq n_0$ points of fractal dimension at most δ, such that for any $c \geq 1$, any c-spanner G of P has $\mathsf{tw}(G) = \Omega\left(\frac{n^{1-\frac{1}{\delta}}}{c}\right).$*

Construction of the discrete Cantor crossbar.   To prove the above theorem, we consider the set of points obtained as follows. First we define $f(0)$ and $h(0)$ and $g(0)$ to be a single point. Then we inductively define $h(i)$, $g(i)$ and $f(i)$ as follows. To get $h(i)$ we start with a $3^i \times 3^i$ integer grid and subdivide it into nine $3^{i-1} \times 3^{i-1}$ integer grids. Then we remove all 3 sub-grids in the middle row and replace each of the 6 remaining sub-grids with copies of $h(i-1)$. To get $g(i)$ we again start with a $3^i \times 3^i$ integer grid and subdivide it into nine $3^{i-1} \times 3^{i-1}$ integer grids. Then we remove all 3 sub-grids in the middle column and replace each of the 6 remaining sub-grids with copies of $g(i-1)$. Finally to get $f(i)$ we start with a $3^i \times 3^i$ integer grid and subdivide it into nine $3^{i-1} \times 3^{i-1}$ integer grids. Then we remove the central sub-grid. We then replace the four corner sub-grids with copies of $f(i-1)$. We replace the middle sub-grid in the first and last rows with copies of $h(i-1)$ and we replace the middle sub-grids of the first and last columns with copies of $g(i-1)$ as depicted in Figure 6. The pointset we require is given by $f(k)$ where k is any positive integer. We have the following two lemmas regarding the pointset $f(k)$.

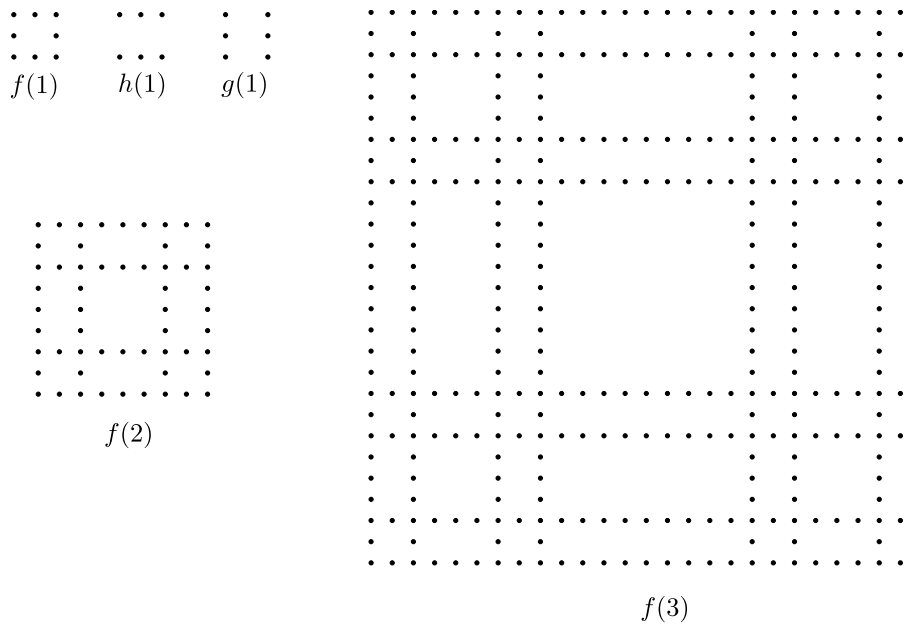▶ **Lemma 17.** $|f(k)| \leq 2 \cdot 6^k$.

Figure 6 The construction of the discrete Cantor crossbar.

▶ **Lemma 18.** *Let $P = f(k)$. Then,* $\dim_{\mathsf{f}}(P) \leq \frac{\log(6)}{\log(3)}$.

▶ **Lemma 19.** *Let G be a c-spanner of P, where $c \geq 1$. Then,* $\mathsf{tw}(G) = \Omega\left(\frac{2^k}{c}\right)$.

Proof of Theorem 16. Using Lemma 18 and Lemma 19, we get that

$$\mathsf{tw}(G) = \Omega\left(\frac{(3^{k\delta})^{1-\frac{1}{\delta}}}{c}\right) = \Omega\left(\frac{n^{1-\frac{1}{\delta}}}{c}\right) \qquad \blacktriangleleft$$

.

2.4 A family of treewidth-extremal fractals for all dimensions

We can now generalize the ideas from the previous section to obtain a family of pointsets that allow us to a get a lower bound on the treewidth of spanners for any given choice of fractal dimension.

Construction of treewidth-extremal fractal pointsets. Consider the family of pointsets defined as follows: For all integers $d > 0$, we name the dimensions from $\{1, 2, \ldots, d\}$ in an arbitrary manner. For all odd integers l and v such that $l > v$, we define each of $f^{l,v,d}(0), h_1^{l,v,d}(0), \ldots, h_d^{l,v,d}(0)$ to be a single point. We inductively define $f^{l,v,d}(i), h_1^{l,v,d}(i), \ldots$ $\ldots, h_d^{l,v,d}(i)$ as follows: For $h_1^{l,v,d}(i)$, we start with a d-dimensional l^i integer grid and subdivide it to get l^d identical l^{i-1} d-dimensional integer subgrids. Now, along every dimension $j \in \{2, \ldots, d\}$, we remove all $l^{d-1}v$ subgrids in the middle v rows of the subgrids. We then replace each of the remaining $l(l-v)^{d-1}$ subgrids with copies of $h_1^{l,v,d}(i-1)$. The pointset obtained is $h_1^{l,v,d}(i)$. In general for any $m \in [d]$, we construct $h_m^{l,v,d}(i)$ as follows: we start with a d-dimensional l^i integer grid and subdivide it to get l^d identical l^{i-1} d-dimensional integer subgrids. Then, along every dimension $j \neq m$, we remove all $l^{d-1}v$ subgrids in the middle v rows of the subgrids. We replace each of the remaining $l(l-v)^{d-1}$ subgrids

with copies of $h_m^{l,v,d}(i-1)$. The pointset thus obtained is $h_m^{l,v,d}(i)$. In order to construct $f^{l,v,d}(i)$, we start with a d-dimensional l^i integer grid and subdivide it into l^d identical l^{i-1} d-dimensional integer subgrids. Let S denote the set of the central v^d subgrids. Note that S is a d-dimensional grid with side length $v(l^{i-1})$. Now along every dimension $m \in [d]$, there are $\frac{(l-v)v^{d-1}}{2}$ subgrids on each side of S. We remove S as well as all such $\frac{(l-v)v^{d-1}}{2}$ subgrids lying on either side of S along every dimension. We replace each of the $\left(\frac{l-v}{2}\right)^d$ sub-grids in the 2^d corners with copies of $f^{l,v,d}(i-1)$. Then along every dimension m, we replace each of the remaining $\left(\frac{l-v}{2}\right)^{d-1} v2^{d-1}$ subgrids with copies of $h_m^{l,v,d}(i-1)$. The pointset thus obtained is $f(i)$. To generate the pointset P mentioned in the statement of Theorem 1 we pick l and v such that $\left|\delta - \frac{\log(l(l-v)^{d-1})}{\log l}\right| \leq \varepsilon$. Such a pair of odd numbers always exists. Let $\delta' = \frac{(\log l(l-v)^{d-1})}{\log l}$. Then, we set P to be $f^{l,v,d}(k)$, where k is any positive integer. We have the following lemmas regarding pointset P.

▶ **Lemma 20.** *For all odd integers l and v such that $l > v$, and for all positive integers $k > 0$, we have that $\dim_f(f^{l,v,d}(k)) \leq \frac{(\log l(l-v)^{d-1})}{\log l}$.*

▶ **Lemma 21.** $\dim_f(P) \leq \delta'$.

Proof. This follows from Lemma 20 when applied to our choice of l,v and P. ◀

▶ **Lemma 22.** *Let G be a c-spanner of P, where $c \geq 1$. Then, $\mathsf{tw}(G) = \Omega\left(\frac{(l-v)^{k(d-1)}}{c^{d-1}}\right)$.*

Proof of Theorem 1. Using Lemma 22 and Lemma 21, we get that

$$\mathsf{tw}(G) = \Omega\left(\frac{(l-v)^{k(d-1)}}{c^{d-1}}\right) = \Omega\left(\frac{l^{k\delta'\left(1-\frac{1}{\delta'}\right)}}{c^{d-1}}\right) = \Omega\left(\frac{n^{1-\frac{1}{\delta'}}}{c^{d-1}}\right).$$

This combined with the result of Lemma 21 proves the statement of the theorem. ◀

3 Running time lower bound for Independent Set of Unit Balls

In this section, we present the proof of Theorem 2. Our argument uses a reduction from a type of Constraint Satisfaction Problem called the Geometric Constraint Satisfaction Problem. The definitions in this section are taken from [19].

▶ **Definition 23** (The Constraint Satisfaction Problem). [19] The input instance I of a constraint satisfaction problem is a triple (V, D, C), where V is a set of variables that can take values in the domain D, and C is a set of constraints, with each constraint being a pair $\langle s_i, R_i \rangle$ such that:
- s_i is a tuple of variables of size m_i.
- R_i is an m_i-ary relation over D.
A valid solution to the problem is an assignment of values from D to each of the variables in V such that for all constraints $\langle s_i, R_i \rangle$, the assignment for each tuple s_i is in R_i.

For our purposes, we only need to consider the case where the constraints are binary, or in other words, for all i, we have that $m_i = 2$. We may assume that the input size $|I|$ of a binary CSP instance is a polynomial in $|V|$ and $|D|$.

For an instance I of the Constraint Satisfaction Problem, the *primal graph* is a graph G with vertex set V and an edge between $u, w \in V$ if and only if there exists a constraint $\langle s_i, R_i \rangle \in C$, such that $s_i = (u, w)$.

Let $R[n, d]$ denote the d-dimensional grid with vertex set $[n]^d$ and let $\mathcal{R}_d$ denote the set of graphs $R[n, d]$ for all $n \geq 1$.
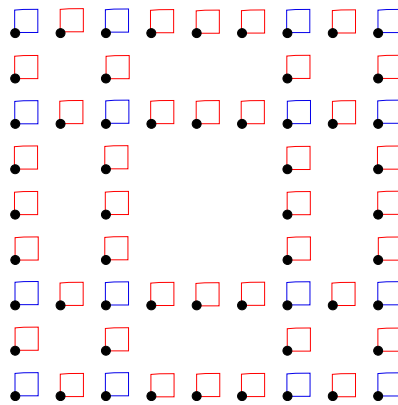
Figure 7 The case $d = 2$, $l = 3$, $v = 1$ and $m = 2$. Boxes contain centers of balls in B. The blue boxes are obtained from variables in I and the red boxes are obtained from variables in $I' \setminus I$.

▶ **Definition 24** (The $\leq$-CSP). [19] A d-dimensional geometric $\leq$-CSP is a constraint satisfaction problem of the following form: The set of variables V is a subset of vertices of $R[n, d]$ for some n and the primal graph is an induced subgraph of $R[n, d]$. The domain is $[\Delta]^d$ for some integer $\Delta \geq 1$. The instance can contain arbitrary unary constraints but the binary constraints are of a special form. A geometric constraint is a constraint $\langle (\mathbf{a}, \mathbf{a}'), R \rangle$ with $\mathbf{a}' = \mathbf{a} + \mathbf{e}_i$ such that

$$R = \{((x_1, \ldots, x_d), (y_1, \ldots, y_d)) \mid x_i \leq y_i\}$$

This means that, if variables $\mathbf{a}$ and $\mathbf{a}'$ are adjacent with $\mathbf{a}'$ being larger by one in the i-th coordinate, then the i-th coordinate of the value of $\mathbf{a}$ is at most as large as the i-th coordinate of the value of $\mathbf{a}'$.

We will use the following theorem from [19] in the proof of Theorem 2. We remark that the condition $|V| = \Theta(n^d)$ is implicit in [19].

▶ **Theorem 25** ([19, Theorem 2.20]). *If for some fixed $d \geq 1$, there is an $f(|V|)n^{o(|V|^{1-1/d})}$ time algorithm for d-dimensional geometric $\leq$-CSP for some function f, where $|V| = \Theta(n^d)$, then ETH fails.*

Construction of $\leq$-CSP I'. Given $\delta \in (1, d)$ and $\varepsilon' > 0$, we can find odd integers l and v such that $\left| \delta - \frac{\log l(l-v)^{d-1}}{\log l} \right| \leq \varepsilon'$. Let $\delta' = \frac{\log l(l-v)^{d-1}}{\log l}$. Let I be a d-dimensional $\leq$-CSP instance with variables V and domain $[\Delta]^d$, where Δ is any positive integer. Let the primal graph of I be $R[n, d]$. We now define a new d-dimensional $\leq$-CSP I', with variables V' and domain $[\Delta]^d$, such that $|V'| = O\left(|V|^{\frac{d-1}{d} \cdot \frac{\delta'}{\delta'-1}} \right)$, and the primal graph of I' is $R[n_{new}, d]$, where $n_{new} = n^{O(1)}$. Let $m > 0$ be the smallest integer such that $l^{\frac{(m-1)(\delta'-1)}{d-1}} < n \leq l^{\frac{m(\delta'-1)}{d-1}}$. We construct $f^{l,v,d}(m)$, as in the proof of Theorem 1. From Lemma 22, we know that $f^{l,v,d}(m)$ contains a subset of points that form a d-dimensional grid of side length $l^{\frac{m(\delta'-1)}{d-1}}$. Let M denote this grid contained in $f^{l,v,d}(m)$. Since the variables V lie on the grid $R[n, d]$ and $n^d \leq l^{\frac{md(\delta'-1)}{d-1}}$, we can place the variables V on grid M such that their position relative to each other in M is the same as in $R[n, d]$. By abuse of notation, we refer to the subset of M containing variables V as V. We refer to the pointset $f^{l,v,d}(m) \setminus M$ as C. We observe

that the points in C connect adjacent points of M. We now define a new set of variables $V' = V \cup C$. If $\mathbf{a} \in V$, then the set of unary constraints for $\mathbf{a}$ is $R_{\mathbf{a}}$. If $\mathbf{a} \in C$, then $\mathbf{a}$ belongs to a chain of points that connects two points of M along some dimension i, where $i \in [d]$. For an $\mathbf{a} \in C$ connecting two points of M along dimension i, we define $R_{\mathbf{a}}$ as follows:

$$R_{\mathbf{a}} = \{(a_1, \ldots, a_{i-1}, a_i, a_{i+1} \ldots, a_d) \mid a_j = 0 \; \forall \; j \neq i \text{ and } a_i \in [\Delta]\}.$$

Thus for every $\mathbf{a} \in C$, we have $|R_{\mathbf{a}}| = \Delta$. We define $R_{\mathbf{a}}$ for all variables $\mathbf{a} \in C$ in a similar manner. The binary constraints on variables V' are described as follows: a binary constraint is a constraint $\langle (\mathbf{a}, \mathbf{a}'), R \rangle$, with $\mathbf{a}' = \mathbf{a} + \mathbf{e}_i$ in $f^{l,v,d}(m)$ such that

$$R = \{((x_1, \ldots, x_d), (y_1, \ldots, y_d)) \mid x_i \leq y_i\}.$$

Let I' denote the new $\leq$-geometric CSP with variables V', and unary and binary constraints as defined above. From the choice of m, we have that $|V'| \leq l^{\delta'} \cdot n^{(d-1) \cdot \frac{\delta'}{\delta'-1}}$. Thus $|V'| = O\left(|V|^{\frac{d-1}{d} \cdot \frac{\delta'}{\delta'-1}}\right)$, where we use the fact that δ, δ', and l are fixed constants. Similarly, we get that $n_{new} = n^{O(1)}$.

▶ **Lemma 26.** *I' is satisfiable if and only if I is satisfiable.*

Proof of Theorem 2. This proof is similar to the proof of Theorem 3.1 in [19]. Let I be a d-dimensional $\leq$-CSP with variables V and domain $[\Delta]^d$, and with primal graph $R[n, d]$. As explained before, we construct a new $\leq$-CSP I' with variables V' and domain $[\Delta]^d$, such that $|V'| = O\left(|V|^{\frac{d-1}{d} \cdot \frac{\delta}{\delta-1}}\right)$ and, the primal graph of I' is $R[n_{new}, d]$, where $n_{new} = n^{O(1)}$. We construct a set B of n_{new} balls, of diameter 1 each, such that the centers of the balls have fractal dimension at most δ', and B contains a set of $|V'|$ pairwise non-intersecting balls if and only if I' has a satisfying assignment.

If there is an $f(|V'|)n_{new}^{o(|V'|^{1-1/\delta'})}$ time algorithm for finding $|V'|$ pairwise non-intersecting balls in B, then we can solve I' in $f(|V'|)n_{new}^{o(|V'|^{1-1/\delta'})}$ time. Since I' is satisfiable if and only if I is satisfiable, we obtain a solution for I in time $f(|V'|)n_{new}^{o(|V'|^{1-1/\delta'})}$. Let h be a computable function such that $h(|V|) = f(|V'|)$. Since $f(|V'|)n_{new}^{o(|V'|^{1-1/\delta'})} = h(|V|)n^{o(|V|^{1-1/d})}$, we obtain a solution for I in time $h(|V|)n^{o(|V|^{1-1/d})}$. By Theorem 25, this contradicts ETH. Please refer to the full version of the paper for a detailed proof. ◀

▶ **Remark 27.** We can similarly show that assuming the Exponential Time Hypothesis, for any $\delta \in (1, d)$ and any $\varepsilon > 0$, the problem of finding k-pairwise non-intersecting d-dimensional axis parallel unit cubes in a collection of n cubes with centers having fractal dimension at most δ cannot be solved in time $f(k)n^{O(k^{1-1/(\delta-\varepsilon)})}$, for any computable function f. Given a $\leq$-CSP instance I, we reduce I to another $\leq$-CSP instance I' as explained in the beginning of this section. We then use the construction and analysis as in the proof of Theorem 3.2 of [19], replacing I with I'. Using Theorem 25 and the proof of Theorem 2, we get the desired result.

References

1 Jochen Alber and Jiří Fiala. Geometric separation and exact solutions for the parameterized independent set problem on disk graphs. In *Foundations of Information Technology in the Era of Network and Mobile Computing*, pages 26–37. Springer, 2002.

2 Yair Bartal, Lee-Ad Gottlieb, and Robert Krauthgamer. The traveling salesman problem: low-dimensionality implies a polynomial time approximation scheme. In *Proceedings of the forty-fourth annual ACM symposium on Theory of computing*, pages 663–672. ACM, 2012.

3 Hubert TH Chan, Anupam Gupta, Bruce M Maggs, and Shuheng Zhou. On hierarchical routing in doubling metrics. In *Proceedings of the sixteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 762–771. Society for Industrial and Applied Mathematics, 2005.

4 T-H Hubert Chan and Anupam Gupta. Small hop-diameter sparse spanners for doubling metrics. *Discrete & Computational Geometry*, 41(1):28–44, 2009.

5 Chandra Chekuri and Julia Chuzhoy. Polynomial bounds for the grid-minor theorem. *Journal of the ACM (JACM)*, 63(5):40, 2016.

6 Richard Cole and Lee-Ad Gottlieb. Searching dynamic point sets in spaces with bounded doubling dimension. In *Proceedings of the thirty-eighth annual ACM symposium on Theory of computing*, pages 574–583. ACM, 2006.

7 Reinhard Diestel. *Graph Theory, 4th Edition*, volume 173 of *Graduate texts in mathematics*. Springer, 2012.

8 Lee-Ad Gottlieb and Liam Roditty. An optimal dynamic spanner for doubling metric spaces. In *European Symposium on Algorithms*, pages 478–489. Springer, 2008.

9 Anupam Gupta and Kevin Lewi. The online metric matching problem for doubling metrics. In *International Colloquium on Automata, Languages, and Programming*, pages 424–435. Springer, 2012.

10 Sariel Har-Peled and Manor Mendel. Fast construction of nets in low-dimensional metrics and their applications. *SIAM Journal on Computing*, 35(5):1148–1184, 2006.

11 T-H Hubert Chan and Anupam Gupta. Approximating tsp on metrics with bounded global growth. *SIAM Journal on Computing*, 41(3):587–617, 2012.

12 David R Karger and Matthias Ruhl. Finding nearest neighbors in growth-restricted metrics. In *Proceedings of the thiry-fourth annual ACM symposium on Theory of computing*, pages 741–750. ACM, 2002.

13 Kyohei Kozawa, Yota Otachi, and Koichi Yamazaki. The carving-width of generalized hypercubes. *Discrete Math.*, 310(21):2867–2876, 2010. `doi:10.1016/j.disc.2010.06.039`.

14 Robert Krauthgamer and James R Lee. The black-box complexity of nearest-neighbor search. *Theoretical Computer Science*, 348(2):262–276, 2005.

15 Robert Krauthgamer and James R Lee. Algorithms on negatively curved spaces. In *2006 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS'06)*, pages 119–132. IEEE, 2006.

16 Robert Krauthgamer, James R Lee, Manor Mendel, and Assaf Naor. Measured descent: A new embedding method for finite metrics. *Geometric & Functional Analysis GAFA*, 15(4):839–858, 2005.

17 Dániel Marx. Efficient approximation schemes for geometric problems. In *European Symposium on Algorithms*, pages 448–459. Springer, 2005.

18 Dániel Marx. Can you beat treewidth? *Theory OF Computing*, 6:85–112, 2010.

19 Dániel Marx and Anastasios Sidiropoulos. The limited blessing of low dimensionality: When 1-1/d is the best possible exponent for d-dimensional geometric problems. In *Proceedings of the Thirtieth Annual Symposium on Computational Geometry*, SOCG'14, pages 67:67–67:76, New York, NY, USA, 2014. ACM. `doi:10.1145/2582112.2582124`.

20 Neil Robertson, Paul Seymour, and Robin Thomas. Quickly excluding a planar graph. *Journal of Combinatorial Theory, Series B*, 62(2):323–348, 1994.

21 Neil Robertson and Paul D Seymour. Graph minors. v. excluding a planar graph. *Journal of Combinatorial Theory, Series B*, 41(1):92–114, 1986.

22 Jeffrey S Salowe. Construction of multidimensional spanner graphs, with applications to minimum spanning trees. In *Proceedings of the seventh annual symposium on Computational geometry*, pages 256–261. ACM, 1991.

23   Anastasios Sidiropoulos and Vijay Sridhar.   Algorithmic interpretations of fractal dimension.   In *33rd International Symposium on Computational Geometry, SoCG 2017, July 4-7, 2017, Brisbane, Australia*, volume 77 of *LIPIcs*, pages 58:1–58:16. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2017.   URL: `http://www.dagstuhl.de/dagpub/978-3-95977-038-5`, `doi:10.4230/LIPIcs.SoCG.2017.58`.

24   Warren D Smith and Nicholas C Wormald. Geometric separator theorems and applications. In *Foundations of Computer Science, 1998. Proceedings. 39th Annual Symposium on*, pages 232–243. IEEE, 1998.

25   Hideki Takayasu. *Fractals in the physical sciences.* Manchester University Press, 1990.

26   Kunal Talwar.   Bypassing the embedding: algorithms for low dimensional metrics.   In *Proceedings of the thirty-sixth annual ACM symposium on Theory of computing*, pages 281–290. ACM, 2004.

27   Pravin M Vaidya. A sparse graph almost as good as the complete graph on points ink dimensions. *Discrete & Computational Geometry*, 6(3):369–381, 1991.