

Algorithmic Interpretations of Fractal Dimension^{*†}

Anastasios Sidiropoulos¹ and Vijay Sridhar²

- 1 Dept. of Mathematics and Dept. of Computer Science & Engineering,
The Ohio State University, Columbus, OH, USA
sidiropoulos.1@osu.edu
- 2 Dept. of Computer Science & Engineering, The Ohio State University,
Columbus, OH, USA
sridhar.38@buckeyemail.osu.edu

Abstract

We study algorithmic problems on subsets of Euclidean space of low *fractal dimension*. These spaces are the subject of intensive study in various branches of mathematics, including geometry, topology, and measure theory. There are several well-studied notions of fractal dimension for sets and measures in Euclidean space. We consider a definition of fractal dimension for finite metric spaces which agrees with standard notions used to empirically estimate the fractal dimension of various sets. We define the fractal dimension of some metric space to be 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 (that is, for any maximal ε -packing), we have $|B \cap N| = O((r/\varepsilon)^\delta)$.

Using this definition we obtain faster algorithms for a plethora of classical problems on sets of low fractal dimension in Euclidean space. Our results apply to exact and fixed-parameter algorithms, approximation schemes, and spanner constructions. Interestingly, the dependence of the performance of these algorithms on the fractal dimension nearly matches the currently best-known dependence on the standard Euclidean dimension. Thus, when the fractal dimension is strictly smaller than the ambient dimension, our results yield improved solutions in all of these settings.

We remark that our definition of fractal definition is equivalent up to constant factors to the well-studied notion of doubling dimension. However, in the problems that we consider, the dimension appears in the exponent of the running time, and doubling dimension is not precise enough for capturing the best possible such exponent for subsets of Euclidean space. Thus our work is *orthogonal* to previous results on spaces of low doubling dimension; while algorithms on spaces of low doubling dimension seek to extend results from the case of low dimensional Euclidean spaces to more *general* metric spaces, our goal is to obtain faster algorithms for *special* pointsets in Euclidean space.

1998 ACM Subject Classification F.2.2 [Nonnumerical Algorithms and Problems] Geometrical Problems and Computations

Keywords and phrases fractal dimension, exact algorithms, fixed parameter tractability, approximation schemes, spanners

Digital Object Identifier 10.4230/LIPIcs.SoCG.2017.58

* A full version of the paper is available at <https://arxiv.org/abs/1703.09324>.

† This work was supported by the National Science Foundation (NSF) under grant CCF 1423230 and award CAREER 1453472.



© Anastasios Sidiropoulos and Vijay Sridhar;
licensed under Creative Commons License CC-BY

33rd International Symposium on Computational Geometry (SoCG 2017).

Editors: Boris Aronov and Matthew J. Katz; Article No. 58; pp. 58:1–58:16

Leibniz International Proceedings in Informatics



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

1 Introduction

Sets of non-integral dimension are ubiquitous in nature and can be used to model a plethora of processes and phenomena in science and engineering [26]. Sets and measures in Euclidean space of certain fractal dimension are the subject of study in several branches of mathematics, including geometry, topology, and measure theory.

In many problems in computational geometry, the dimension of the input set often determines the complexity of the best-possible algorithms. In this work we study the computational complexity of geometric problems on sets of bounded fractal dimension in low-dimensional Euclidean space. We observe the following interesting phenomenon: For many problems, it is possible to obtain algorithms with dependence on the fractal dimension similar to the best-possible dependence to the standard Euclidean dimension. This implies asymptotically faster algorithms when the fractal dimension of the input is smaller than the ambient dimension.

1.1 Definition of fractal dimension


Intuitively, some $X \subseteq \mathbb{R}^d$ has fractal dimension $\delta \in [0, d]$ if when scaling X by a factor of $\alpha > 0$, the “volume” of X is multiplied by a factor α^δ . There are many different ways this intuition can be formalized, such as Hausdorff dimension, Minkowski dimension, and so on. Unfortunately, some of these definitions are not directly applicable in the context of discrete computational problems. For example, the Hausdorff dimension of any countable set is 0.

Despite this, there are some natural methods that are used to estimate the fractal dimension of a set in practice. Let $X \subseteq \mathbb{R}^d$. Let Γ_ε be a d -dimensional grid where each cell has width $\varepsilon > 0$, and let $I_\varepsilon(X)$ be the number of cells in Γ_ε that intersect X . The *fractal box-counting dimension* of X is defined to be $\lim_{\varepsilon \rightarrow 0} \log(I_\varepsilon(X)) / \log(1/\varepsilon)$ [8]. This definition is often used experimentally as follows: Intersect X with a regular lattice $(\varepsilon\mathbb{Z})^d$, and estimate the rate by which the cardinality of the intersection grows when $\varepsilon \rightarrow 0$. In that context, X has fractal dimension δ when the size of the intersection grows as $(1/\varepsilon)^\delta$ [18].

We consider a definition that is closely related to box-counting dimension, but is more easily amenable to algorithmic analysis. Let $S \subseteq A$. We say that S is an ε -covering of A if for any $x \in A$ we have that $\text{dist}(\{x\}, S) \leq \varepsilon$. For any $x \in A$ and $y \in S$ we say that x is covered by y if $\rho(x, y) \leq \varepsilon$. S is an ε -packing if for any $x, y \in S$ we have $\rho(x, y) \geq \varepsilon$. If S is both an ε -covering and an ε -packing of A then we say that S is an ε -net of A . We define the *fractal dimension* of some family of pointsets $P \subseteq \mathbb{R}^d$, denoted by $\text{dim}_f(P)$, to be the infimum δ , such that for any $\varepsilon > 0$ and $r \geq 2\varepsilon$, for any ε -net¹ N of P , and for any $x \in \mathbb{R}^d$, we have $|N \cap \text{ball}(x, r)| = O((r/\varepsilon)^\delta)$. For the sake of notational simplicity, we will be referring to the fractal dimension of some family of pointsets P , as the fractal dimension of the pointset P , with the understanding that in the asymptotic notation $|P|$ is unbounded.

Figure 1 depicts an example of an infinite family of discrete pointsets P with non-integral fractal dimension constructed as follows: We begin with the $3^k \times 3^k$ integer grid, for some $k \in \mathbb{N}$, we partition it into 9 subgrids of equal size, we delete all the points in the central subgrid, and we 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. It can be shown that $\text{dim}_f(P) = \log_3 8$, which is equal to the Hausdorff dimension of the standard Sierpiński carpet.

¹ We arrive at an equivalent definition if we require N to be a ε -packing instead of a ε -net.



■ **Figure 1** A discrete variant of the Sierpiński carpet for $k = 3$.

1.2 Why yet another notion of dimension?

We now briefly compare the above notion of fractal dimension to previous definitions and motivate its importance. The most closely related notion that has been previously studied in the context of algorithm design is *doubling dimension* [2, 15, 10]. We recall that the doubling dimension of some metric space M , denoted by $\dim_d(M)$, is defined to be $\log \kappa$, where κ is the minimum integer such that for all $r > 0$, any ball in M of radius r can be covered by at most κ balls of radius $r/2$. It is easy to show that for any metric space M , we have² $\dim_d(M) = \dim_f(M) + O(1)$ and $\dim_f(M) = O(\dim_d(M))$. Thus our definition is equivalent to doubling dimension up to constant factors. However, in the problems we consider, the dimension appears in the *exponent* of the running time of the best-known algorithms; therefore, determining the best-possible constant is of importance. As we shall see, for several algorithmic problems, our definition yields nearly optimal bounds on this exponent, while doubling dimension is not precise enough for this task.

Let us illustrate this phenomenon on the problem of solving TSP exactly of a set of n points in the Euclidean plane. It is known that TSP admits an algorithm with running time $2^{O(\sqrt{n} \log n)} n^{O(1)}$ in this case [25]. Moreover, the exponent of $O(\sqrt{n} \log n)$ is known to be nearly optimal assuming the Exponential Time Hypothesis (ETH) [23] (see later in this Section for a more precise statement). We show that for sets of fractal dimension $\delta \in (1, 2]$, there exists an algorithm with running time $2^{O(n^{1-1/\delta} \log n)}$. Thus, for any fixed $\delta < 2$, we achieve an asymptotically faster algorithm than what is possible for general pointsets (assuming ETH). On the other hand, it is known that the unit disk cannot be covered with 6 disks of radius $1/2$ (see [29]). Thus $\dim_d(\mathbb{R}^2) \geq \log_2 7 > 2.807$, while $\dim_f(\mathbb{R}^2) = 2$. Therefore doubling dimension is not precise enough to capture the best-possible exponent in this setting.

In summary, while algorithms on spaces of low doubling dimension seek to extend results from the case of low dimensional Euclidean space to a more *general* setting, our goal is to obtain faster algorithms for *special* classes of pointsets in Euclidean space.

1.3 Our results

We obtain algorithms for various problems on sets of low fractal dimension in Euclidean space. We consider exact algorithms, fixed parameter algorithms, and approximation schemes. In each one of these settings, we pick classical representative problems. We believe that our techniques should be directly applicable to many other problems.

² Note that for a set X containing two distinct points we have $\dim_f(X) = 0$ while $\dim_d(X) = 1$ and thus it is not always the case that $\dim_d(X) = O(\dim_f(X))$.

Exact algorithms. We first consider exact algorithms in \mathbb{R}^d . It is known that for any fixed d , TSP on a set of n points in \mathbb{R}^d can be solved in time $2^{O(n^{1-1/d} \log n)}$ [25]. By adapting ideas from the Euclidean setting, we show that TSP on a set of n points of fractal dimension $\delta > 1$ in constant-dimensional Euclidean space, can be solved in time $2^{O(n^{1-1/\delta} \log n)}$. When $\delta = 1$ and $\delta < 1$, our algorithm has running time $n^{O(\log^2 n)}$ and $n^{O(\log n)}$ respectively. We remark that it has been shown by Marx and Sidiropoulos [23] that assuming ETH, there is no algorithm for TSP in \mathbb{R}^d with running time $2^{O(n^{1-1/d-\varepsilon})}$, for any $\varepsilon > 0$. Thus, our result bypasses this lower bound for sets of low fractal dimension. In particular, our result implies that, in a certain sense, the hardest instances for TSP in \mathbb{R}^d must be close to *full-dimensional*; that is, they must have fractal dimension close to d . Our technique also extends to the Minimum Rectilinear Steiner Tree problem in \mathbb{R}^2 . Due to lack of space, this extension is omitted.

Parameterized problems. We also consider algorithms for problems parameterized by the value of the optimum solution. A prototypical geometric problem in this setting is Independent Set of unit balls in \mathbb{R}^d . Formally, we show that given a set D of unit balls in \mathbb{R}^d , the k -Independent Set problem on D can be solved in time $n^{O(k^{1-1/\delta})}$, for any fixed d , where $\delta > 1$ is the fractal dimension of the set of centers of the disks in D . When $\delta \leq 1$, we get an algorithm with running time $n^{O(\log k)}$. Previously known algorithms for this problem in d -dimensional Euclidean space have running time $n^{O(k^{1-1/d})}$, for any $d \geq 2$ [1, 23]. Moreover, it has been shown that there is no algorithm with running time $f(k)n^{o(k^{1-1/d})}$, for any computable function f , assuming ETH [23] (see also [22]). Thus, our result implies that this lower bound can also be bypassed for sets of fractal dimension $\delta < d$.

Approximation schemes. We next consider approximation schemes. Let P be a set of n points of fractal dimension $\delta > 0$, in d -dimensional Euclidean space. We show that for any $R > 0$, for any $\ell > 0$, we can compute a $(1 + d/\ell)$ -approximate R -cover of P in time $\ell^{d+\delta} n^{O((\ell\sqrt{d})^\delta)}$. This matches the performance of the algorithm of Hochbaum and Maass [16] after replacing δ by d . We also obtain a similar algorithm for the R -packing problem.

Spanners and pathwidth. Recall that for any pointset in \mathbb{R}^d , and for any $c \geq 1$, a c -spanner for P is a graph G with $V(G) = P$, such that for all $x, y \in P$, we have $\|x - y\|_2 \leq d_G(x, y) \leq c \cdot \|x - y\|_2$, where d_G denotes the shortest path distance in G . The parameter c is called the *dilation* of G . It is known that for any $\varepsilon > 0$, any set of n points in \mathbb{R}^d admits a $(1 + \varepsilon)$ -spanner of size $n(1/\varepsilon)^{O(d)}$ [24, 28]. We strengthen this result in the following way. We show that for any $\varepsilon > 0$, any set of n points of fractal dimension δ in constant-dimensional Euclidean space admits a $(1 + \varepsilon)$ -spanner of size $n(1/\varepsilon)^{O(d)}$, and of pathwidth at most $O(n^{1-1/\delta} \log n)$ if $\delta > 1$, at most $O(\log^2 n)$ if $\delta = 1$, and at most $O(\log n)$ if $\delta < 1$. Our spanner is obtained via a modification of the construction due to Vaidya [28]. This provides a general polynomial-time reduction for geometric optimization problems on Euclidean instances of low fractal dimension to corresponding graph instances of low pathwidth. This result can be understood as justification for the fact that instances of low fractal dimension appear to be “easier” than arbitrary instances. We remark that our construction also implies, as a special case, that arbitrary n -pointsets in \mathbb{R}^d admit $(1 + \varepsilon)$ -spanners of size $n(1/\varepsilon)^{O(d)}$ and pathwidth $O(n^{1-1/d} \log n)$; this bound on the pathwidth appears to be new, even for the case $d = 2$.

1.4 Related work

There is a large body of work on various notions of dimensionality in computational geometry. Most notably, there has been a lot of effort on determining the effect of doubling dimension

on the complexity of many problems [14, 3, 7, 19, 9, 21, 4, 6, 11, 27]. Other notions that have been considered include low-dimensional negatively curved spaces [20], growth-restricted metrics [17], as well as generalizations of doubling dimension to metrics of bounded global growth [5].

A common goal in all of the above lines of research is to extend tools and ideas from the Euclidean setting to more general geometries. In contrast, as explained above, we study restricted classes of Euclidean instances, with the goal of obtaining faster algorithms than what is possible for arbitrary Euclidean pointsets.

1.5 Notation and definitions

Let (X, ρ) be some metric space. For any $x \in X$ and $r \geq 0$, we define $\text{ball}(x, r) = \{y \in X : \rho(x, y) \leq r\}$ and $\text{sphere}(x, r) = \{y \in X : \rho(x, y) = r\}$. For some $A, B \subseteq X$, we write $\text{dist}(A, B) = \inf_{x \in A, y \in B} \{\rho(x, y)\}$. For some $r \geq 0$, we write $N(A, r) = \{x \in X : \text{dist}(A, \{x\}) \leq r\}$. Let $S \subseteq A$. We say that S is an ε -covering of A if for any $x \in A$ we have that $\text{dist}(\{x\}, S) \leq \varepsilon$. For any $x \in A$ and $y \in S$ we say that x is covered by y if $\rho(x, y) \leq \varepsilon$. S is an ε -packing if for any $x, y \in S$ we have $\rho(x, y) \geq \varepsilon$. If S is both an ε -covering and an ε -packing of A then we say that S is an ε -net of A .

We recall the following definition from [25]. Let D be a collection of subsets of \mathbb{R}^d . D is said to be κ -thick if no point is covered by more than κ elements of D . Let D' be any subset of D such that the ratio between the diameters of any pair of elements in D' is at most λ . Then D' is said to be λ -related. D is said to be (λ, κ) -thick if no point is covered by more than κ elements of any λ -related subset of D .

The pathwidth of some graph G , denoted by $\text{pw}(G)$, is the minimum integer $k \geq 1$, such that there exists a sequence C_1, \dots, C_ℓ of subsets of $V(G)$ of cardinality at most $k + 1$, such that for all $\{u, v\} \in E(G)$, there exists $i \in \{1, \dots, \ell\}$ with $\{u, v\} \subseteq C_i$, and for all $w \in V(G)$, for all $i_1 < i_2 < i_3 \in \{1, \dots, \ell\}$, if $w \in C_{i_1} \cap C_{i_3}$ then $w \in C_{i_2}$.

1.6 Organization

The rest of the paper is organized as follows. In Section 2 we derive a separator Theorem for a set of balls whose set of centers has bounded fractal dimension. In Section 3 we present our exact algorithm for TSP. In Section 4 we give a fixed-parameter algorithm for Independent Set of unit balls. In Section 5 we give approximation schemes for packing and covering unit balls. Finally, in Section 6 we present our spanner construction. Due to lack of space, some of the proofs have been deferred to the Appendix.

2 A separator theorem

In this section we prove a separator theorem for a set of d -balls intersecting a set of points with bounded fractal dimension. Subsequently, this result will form the basis for some of our algorithms. The proof uses an argument due to Har-Peled [13].

► **Theorem 1.** *Let $d \geq 2$ be some integer, and let $\delta \in (0, d]$ be some real number. Let $P \subset \mathbb{R}^d$ such that $\dim_f(P) = \delta$. Let B be a (λ, κ) -thick set of d -balls in \mathbb{R}^d , with $|B| = n$, $\lambda \geq 2$ and such that for all $b \in B$ we have $b \cap P \neq \emptyset$. Then there exists a $(d - 1)$ -sphere C such that at most $(1 - 2^{-O(d)})n$ of the elements in B are entirely contained in the interior of C , at most*

$(1 - 2^{-O(d)})n$ of the elements in B are entirely outside C , and

$$|A| = \begin{cases} O\left(\kappa(5\lambda)^d 6^\delta \frac{\lambda}{1-\lambda^{(1-\delta)}} n^{1-1/\delta}\right) & \text{if } \delta > 1 \\ O\left(\kappa(5\lambda)^d 6^\delta \log n\right) & \text{if } \delta = 1 \\ O\left(\frac{\kappa(5\lambda)^d 6^\delta}{\lambda^{1-\delta}-1}\right) & \text{if } \delta < 1 \end{cases},$$

where $A = \{b \in B : \text{diam}(b) \leq \text{diam}(C) \text{ and } b \cap C \neq \emptyset\}$.

Proof. It is known that any ball in \mathbb{R}^d of radius r can be covered by at most $k(d) = 2^{O(d)}$ balls of radius $\frac{r}{2}$. Let C' be the d -ball of minimum radius that contains at least $\frac{1}{k(d)+1}n$ of the elements in B , breaking ties by choosing the ball that contains the maximum number of elements in B . Let \mathbf{o} denote the origin in \mathbb{R}^d . Without loss of generality we can scale and translate the elements of B and P until the radius of C' is 1 and it is centered at \mathbf{o} . Now, let B^* denote the set of d -balls in B of diameter less than or equal to 4 after scaling. We pick uniformly at random $r \in [1, 2]$ and let $C = \text{sphere}(\mathbf{o}, r)$. Now we are ready to obtain an upper bound on the number of elements of B^* that intersect $\text{sphere}(\mathbf{o}, r)$ in expectation.

Consider any d -ball $b \in B^*$ of diameter x . The probability that $\text{sphere}(\mathbf{o}, r)$ intersects b is at most x . Now let $M_1 = \{b \in B^* : \text{diam}(b) \leq n^{-\frac{1}{\delta}} \text{ and } b \cap \text{sphere}(\mathbf{o}, r) \neq \emptyset\}$ and $M_2 = \{b \in B^* : n^{-\frac{1}{\delta}} < \text{diam}(b) \leq 4 \text{ and } b \cap \text{sphere}(\mathbf{o}, r) \neq \emptyset\}$. $|M_1|$ in expectation is at most $O(n^{1-\frac{1}{\delta}})$ as $|B^*| \leq n$. It remains to bound the expected value of $|M_2|$.

Let $B_i = \{b \in B^* : \lambda^i n^{-\frac{1}{\delta}} < \text{diam}(b) \leq \min\{\lambda^{i+1} n^{-\frac{1}{\delta}}, 4\} \text{ and } b \cap \text{sphere}(\mathbf{o}, r) \neq \emptyset\}$. Let n_i denote $|B_i|$. We will construct a $\lambda^i n^{-\frac{1}{\delta}}$ -net of P as follows. Let $B'_i = B_i$. Let π be some arbitrary ordering of the elements of B'_i . In the sequence determined by π pick the next d -ball b from B'_i . Remove all d -balls from B'_i that are entirely within a ball of diameter $5 \cdot \lambda^{i+1} n^{-\frac{1}{\delta}}$ centered at the center of b . Repeat this procedure for the next element determined by π until all the remaining d -balls in B'_i have been visited. From the fact that B is (λ, κ) -thick we have that there can be at most $\kappa 5^d \lambda^d$ elements in B'_i that are contained within a ball of diameter $5 \cdot \lambda^{i+1} n^{-\frac{1}{\delta}}$. This implies that we retain at least a constant fraction of the elements of B_i in B'_i . Now from each $b \in B'_i$ pick a point p_b that also belongs to P and take the union of all such points to get a set of points N_i . From the choice of d -balls in the above argument $|N_i| \geq \frac{1}{\kappa 5^d \lambda^d} n_i$ and N_i is $\lambda^i n^{-\frac{1}{\delta}}$ -packing. We can add more points from P to N_i to obtain a $\lambda^i n^{-\frac{1}{\delta}}$ -net N'_i . We have that $|N_i| \leq |N'_i \cap \text{ball}(\mathbf{o}, 6)| \leq O\left(\left(\frac{6}{\lambda^i n^{-\frac{1}{\delta}}}\right)^\delta\right)$ since $\dim_f(P) = \delta$ and the points of N_i are contained within the ball of radius 6 centered at the origin. This implies that $|B_i| \leq O(\kappa(5\lambda)^d 6^\delta \lambda^{-i\delta} n)$. Since the d -balls in B_i are intersected by $\text{sphere}(\mathbf{o}, r)$ with probability at most $\lambda^{i+1} n^{-\frac{1}{\delta}}$ we have that the expected number of elements of B_i that are intersected by $\text{sphere}(\mathbf{o}, r)$ is $O(\kappa(5\lambda)^d 6^\delta \lambda^{i+1-i\delta} n^{1-\frac{1}{\delta}})$. We thus get $\mathbb{E}[|M_2|] \leq \sum_{i=0}^{\frac{\log n}{\delta}+2} |B_i| \lambda^{i+1} n^{-\frac{1}{\delta}} \leq \sum_{i=0}^{\frac{\log n}{\delta}+2} O(\kappa(5\lambda)^d 6^\delta \lambda^{i+1-i\delta} n^{1-\frac{1}{\delta}})$. When $\delta > 1$ this implies $\mathbb{E}[|M_2|] \leq O(\kappa(5\lambda)^d 6^\delta \left(\frac{\lambda}{1-\lambda^{(1-\delta)}}\right) n^{1-\frac{1}{\delta}})$. When $\delta = 1$ we have $\mathbb{E}[|M_2|] \leq O(\kappa(5\lambda)^d 6^\delta \lambda \left(\frac{\log n}{\delta} + 3\right) n^{1-\frac{1}{\delta}}) \leq O(\kappa(5\lambda)^d 6^\delta \log n)$. When $\delta < 1$ we have $\mathbb{E}[|M_2|] \leq O(\kappa(5\lambda)^d 6^\delta \lambda \left(\frac{\lambda^{(\frac{\log n}{\delta}+3)(1-\delta)}-1}{\lambda^{(1-\delta)}-1}\right) n^{1-\frac{1}{\delta}}) \leq O\left(\frac{\kappa(5\lambda)^d 6^\delta}{\lambda^{1-\delta}-1}\right)$.

For any $r \in [1, 2]$ we have that $A \subseteq B^*$. Thus $\mathbb{E}[|A|] = \mathbb{E}[|M_1|] + \mathbb{E}[|M_2|] \leq O(n^{1-\frac{1}{\delta}}) + \mathbb{E}[|M_2|]$, which implies that

$$\mathbb{E}[|A|] = \begin{cases} O(\kappa(5\lambda)^d 6^\delta \left(\frac{\lambda}{1-\lambda^{(1-\delta)}}\right) n^{1-\frac{1}{\delta}}) & \text{if } \delta > 1 \\ O(\kappa(5\lambda)^d 6^\delta \log n) & \text{if } \delta = 1 \\ O\left(\frac{\kappa(5\lambda)^d 6^\delta}{\lambda^{1-\delta}-1}\right) & \text{if } \delta < 1 \end{cases}$$

Finally we need to ensure that C separates a constant fraction of the elements of B . The choice of C' ensures that at least $\frac{1}{k(d)+1}n = \frac{1}{2^{O(d)}}n$ of the elements in B are entirely contained in the interior of C . This implies that at most $(1 - 2^{-O(d)})n$ of the elements of B are in the exterior of C . Since the $(d - 1)$ -ball of radius 2 is covered by the union of at most $k(d)$ $(d - 1)$ -balls of unit radius we have that there are at most $\frac{k(d)}{k(d)+1}n = (1 - 2^{-O(d)})n$ of the elements in B contained in the interior of C . We note that the upper bound on $\mathbb{E}[|A|]$ remains unaltered for any choice of C' . We further remark that using a more complicated argument similar to the one used by Smith and Wormald [25] a cube separator can be found that separates a constant fraction of d -balls where the constant is independent of d . ◀

3 Exact algorithms

In this Section we give an exact algorithm for TSP. We first use Theorem 1 with the following Lemmas due to Smith and Wormald [25] to obtain a separator for any optimal TSP solution.

▶ **Lemma 2** (Smith and Wormald [25]). *Let $d \geq 2$ be some integer, and let $P \subset \mathbb{R}^d$. Let W be the edge set of an optimal traveling salesman tour of the points of P . Let B be the set of circumballs of the edges of W . Then B is $(2, \kappa)$ -thick where $\kappa = 2^{O(d)}$.*

▶ **Lemma 3** (Smith and Wormald [25]). *Let $d \geq 2$ be some integer, and let $P \subset \mathbb{R}^d$. Let W be the edge set of an optimal traveling salesman tour of the points of P . For any $x \in \mathbb{R}^d$ let $W_x = \{w \in W : \text{diam}(w) \geq 1 \text{ and } w \cap \text{ball}(x, 1) \neq \emptyset\}$. Then $|W_x| \leq 2^{O(d)}$ for all $x \in \mathbb{R}^d$.*

▶ **Theorem 4.** *Let $d \geq 2$ be some integer, and let $\delta \in (0, d]$ be some real number. Let P be a set of n points in \mathbb{R}^d with $\dim_f(P) = \delta$. Let W be the set of edges of any optimal Euclidean TSP tour of P . Then there exists a $(d - 1)$ -sphere C such that at most $(1 - 2^{-O(d)})n$ points in P are contained in the interior of C , at most $(1 - 2^{-O(d)})n$ points in P are contained outside C , and*

$$|W_C| = \begin{cases} O(n^{1-1/\delta}) & \text{if } \delta > 1 \\ O(\log n) & \text{if } \delta = 1 \\ O\left(\frac{1}{2^{1-\delta}-1}\right) & \text{if } \delta < 1 \end{cases},$$

where $W_C = \{w \in W : w \cap C \neq \emptyset\}$.

Proof. Let B be the set of circumballs of the edges in W . From Lemma 2 we have that B is $(2, 2^{O(d)})$ -thick. Every ball in B contains an edge in W and therefore also two points in P . Therefore we can use Theorem 1 on B to find a separator C . It remains to bound the number of edges in W that are intersected by C . Let $W_1 = \{w \in W : \text{diam}(w) \leq \text{diam}(C) \text{ and } w \cap C \neq \emptyset\}$ and $W_2 = \{w \in W : \text{diam}(w) > \text{diam}(C) \text{ and } w \cap C \neq \emptyset\}$. Therefore $W_C = W_1 \cup W_2$. Let B_1 denote the circumballs of the edges in W_1 and B_2 denote the circumballs of the edges in W_2 . If an edge in W_1 is intersected by C then the corresponding circumball in B_1 is also intersected by C . From Theorem 1 we have that

$$|W_1| = \begin{cases} O(n^{1-1/\delta}) & \text{if } \delta > 1 \\ O(\log n) & \text{if } \delta = 1 \\ O\left(\frac{1}{2^{1-\delta}-1}\right) & \text{if } \delta < 1 \end{cases}$$

W.l.o.g. we can assume that C has unit radius and is centered at the origin by scaling and translation. Therefore any edge in W_2 also intersects the unit ball centered at the origin. Combining this with Lemma 3 we have that $|W_2| \leq O(1)$. Since $|W_C| \leq |W_1| + |W_2|$, this concludes the proof. ◀

We now use Theorem 4 to obtain an exact algorithm for TSP. We note that the O -notation hides a factor of $n^{O(1)^d}$.

► **Theorem 5.** *Let $d \geq 2$ be some fixed integer, and let $\delta \in (0, d]$ be some real number. Let P be a set of n points in \mathbb{R}^d with $\dim_f(P) = \delta$. Then for any fixed d an optimal Euclidean TSP tour for P can be found in time $T(n)$, where*

$$T(n) = \begin{cases} n^{O(n^{1-1/\delta})} & \text{if } \delta > 1 \\ n^{O(\log^2 n)} & \text{if } \delta = 1 \\ n^{O(\log n)} & \text{if } \delta < 1 \end{cases}$$

Proof. First we observe that the $(d - 1)$ -sphere separator C described in Theorem 4 can be assumed to intersect at least $d + 1$ points in P . This is because we can always decrease the radius of C without changing W_C until at least one point in P lies on it. We exhaustively consider all separating $(d - 1)$ -spheres to find the separator from Theorem 4. Since every relevant $(d - 1)$ -sphere is uniquely defined by at most $d + 1$ points of P intersecting it, there are at most $n^{O(d)}$ spheres to consider. Let $f(n, \delta)$ denote the number of edges intersected by the separator C . From Theorem 4 we have that

$$f(n, \delta) = \begin{cases} O(n^{1-1/\delta}) & \text{if } \delta > 1 \\ O(\log n) & \text{if } \delta = 1 \\ O(1) & \text{if } \delta < 1 \end{cases}$$

We guess a set E' of at most $f(n, \delta)$ edges in the optimal tour that intersect C . For each such guess E' , we also guess the permutation of E' defined by the order in which the optimal tour traverses the edges in E' . For each such permutation we solve the two sub-problems in the exterior and interior of the separator respecting the boundary conditions. The resulting running time is $T(n) \leq n^{O(d)} n^{O(f(n, \delta))} 2T((1 - 2^{-O(d)})n)$ which implies that for any fixed d implies the assertion. ◀

4 Parameterized problems

In this section we present an algorithm for the parameterized version of the Independent Set problem on a set of unit d -balls in \mathbb{R}^d , where set of centers of the d -balls has bounded fractal dimension. We first prove a separator theorem which will be used in the algorithm.

► **Theorem 6.** *Let $d \geq 2$ be an integer. Let $\delta \in (0, d]$ be a real number. Let P be a set of n points in \mathbb{R}^d with $\dim_f(P) = \delta$. Let $D = \{\text{ball}(x, 1) : x \in P\}$. Let $D' \subseteq D$ be a set of disjoint elements of D such that $|D'| = k$. Then there exists $c \in \mathbb{R}^d$ and $r > 0$ such that at most H d -balls in D' intersect $\text{sphere}(c, r)$ and at most $(1 - 2^{-O(d)})k$ d -balls in D' are contained on either side (interior and exterior) of $\text{sphere}(c, r)$ where*

$$H = \begin{cases} O(k^{1-\frac{1}{\delta}}) & \text{if } \delta > 1 \\ O(1) & \text{if } \delta \leq 1 \end{cases}$$

Proof. Let P' denote the set of centers of the d -balls in D' . We have $|P'| = |D'| = k$. Also since the d -balls in D' are disjoint we have that P' is a 2-packing of P . Consider any $c \in \mathbb{R}^d$ and any $r \geq 1$. Consider a random $(d - 1)$ -sphere $\text{sphere}(c, r')$ with radius $r' \in [r, 2r]$ chosen uniformly at random. Now we can bound the number of d -balls in D' that intersect the $\text{sphere}(c, r')$. First we note that the center of any d -ball in D' that potentially intersects $\text{sphere}(c, r')$ lies within $\text{ball}(c, 2r + 1)$. Therefore the number of d -balls that potentially intersect

$\text{sphere}(c, r')$ is at most $|P' \cap \text{ball}(c, 2r + 1)|$. Since P' is a 2-packing that can be augmented into a 2-net by only adding points, we have that $|P' \cap \text{ball}(c, 2r + 1)| \leq O((\frac{2r+1}{2})^\delta) = O(r^\delta)$. Therefore we have that the number of d -balls that potentially intersect $\text{sphere}(c, r')$ is at most $\min\{k, O(r^\delta)\}$. Any d -ball in D' intersects $\text{sphere}(c, r')$ with probability at most $\frac{2}{r}$ since r' is chosen uniformly at random from the interval $[r, 2r]$. So in expectation the number of d -balls in D' that intersect $\text{sphere}(c, r')$ is at most $\min\{k \cdot \frac{2}{r}, O(r^\delta) \cdot \frac{2}{r}\}$. When $r \leq k^{\frac{1}{\delta}}$ and $\delta > 1$ this is at most $O(r^\delta) \cdot \frac{2}{r} = O(k^{1-\frac{1}{\delta}})$, when $r \leq k^{\frac{1}{\delta}}$ and $\delta \leq 1$ this is at most $O(r^\delta) \cdot \frac{2}{r} = O(r^{\delta-1}) = O(1)$, and when $r > k^{\frac{1}{\delta}}$ this is again at most $k \cdot \frac{2}{r} = O(k^{1-\frac{1}{\delta}})$. This implies that there exists some specific $r' \in [r, 2r]$ such that the number of d -balls in D' that intersect $\text{sphere}(c, r')$ is at most $O(k^{1-\frac{1}{\delta}})$ when $\delta > 1$ and $O(1)$ when $\delta \leq 1$.

Now it remains to specify our choice of c and r so that $\text{sphere}(c, r')$ induces a balanced separator. We will use the fact that for any $r > 0$ any d -ball of radius $2r$ can be covered by at most $g(d) = 2^{O(d)}$ d -balls of radius r . Let c and r be chosen such that $\text{sphere}(c, r)$ is the $(d - 1)$ -sphere with minimum radius that also contains in its interior $\frac{1}{g(d)+1}k$ elements of D' . Since the d -balls have unit radius it follows that the $r \geq 1$. This ensures that there are at least $\frac{1}{2^{O(d)}}$ elements of D' in the interior of $\text{sphere}(c, r)$ and therefore at most $(1 - 2^{-O(d)})k$ elements of D' in the exterior of $\text{sphere}(c, r)$. We have that $\text{ball}(c, r')$ is contained within $\text{ball}(c, 2r)$. Since $\text{ball}(c, 2r)$ can be covered by at most $g(d)$ d -balls of radius r , by our choice of c and r we have that $\text{sphere}(c, r')$ encloses at most $\frac{g(d)}{g(d)+1}k = (1 - 2^{-O(d)})k$ d -balls in D' concluding the proof. ◀

► **Theorem 7.** *Let $d \geq 2$ be an integer. Let $\delta \in (0, d]$ be a real number. Let P be a set of n points in \mathbb{R}^d with $\dim_f(P) = \delta$. Let $D = \{\text{ball}(x, 1) : x \in P\}$. Then there exists an algorithm that computes an independent set in D of size k , if one exists, in time $T(n, k)$, where for any fixed d we have*

$$T(n, k) = \begin{cases} n^{O(k^{1-1/\delta})} & \text{if } \delta > 1 \\ n^{O(\log k)} & \text{if } \delta \leq 1 \end{cases}$$

Proof. Let $D' \subseteq D$ denote the set of k disjoint d -balls in any fixed optimal solution. Let P' denote the set of centers of the d -balls in D' . We have $|P'| = |D'| = k$. We use a divide and conquer approach using the separator from Theorem 6. First we guess the center c and radius r of the smallest $(d - 1)$ -sphere enclosing $\frac{1}{g(d)+1}$ of the d -balls in D' . W.l.o.g. we can assume that there exist a set of d -balls in D' that are tangential to $\text{sphere}(c, r)$ and are enclosed by $\text{sphere}(c, r)$, of cardinality $d+1$. Moreover $\text{sphere}(c, r)$ is uniquely defined by the d -balls that it is tangential to. This implies that $\text{sphere}(c, r - 1)$ intersects at least $d + 1$ points in P and can be uniquely defined by at most $d + 1$ points in P . We can exhaustively guess c by searching through all $(d - 1)$ -spheres uniquely defined by at most $d + 1$ points in P in time $n^{O(d)}$. Next we can assume w.l.o.g. that $\text{sphere}(c, r')$ from Theorem 6 is tangential to at least one d -ball in D' (otherwise r' can be increased or decreased until this condition is met without altering the set of d -balls in D' that are intersected by the separator). This means that given a fixed center c we need to search through at most $2n$ different radii to guess r' . We enumerate over all such separators. For each such separator we again enumerate over all ways to pick the d -balls in D' that are intersected. This can be done in time $n^{O(k^{1-1/\delta})}$ when $\delta > 1$ and $n^{O(1)}$ when $\delta \leq 1$. Therefore we have $T(n, k) = n^{O(d)} \cdot O(n) \cdot n^{O(k^{1-1/\delta})} \cdot 2 \cdot T(n, (1 - 2^{-O(d)})k)$ when $\delta > 1$ or $T(n, k) = n^{O(d)} \cdot O(n) \cdot n^{O(1)} \cdot 2 \cdot T(n, (1 - 2^{-O(d)})k)$ when $\delta \leq 1$, which solves to the desired bound. ◀

5 Approximation schemes

In this section we describe polynomial time approximation schemes for covering and packing problems. We use the approach of Hochbaum and Maass [16].

► **Theorem 8.** *Let $d \geq 2$ be some integer, and let $\delta \in (0, d]$ be some real number. Let P be a set of n points in \mathbb{R}^d with $\dim_f(P) = \delta$. Then there exists a polynomial time approximation scheme which given a natural number $l > 0$ and any $\varepsilon > 0$, computes a $(1 + \frac{d}{l})$ -approximation to the ε -cover of P , in time $l^{d+\delta} n^{O((l\sqrt{d})^\delta)}$.*

Proof. Let A be a d -rectangle that encloses the points in P . Consider a set of hyperplanes perpendicular to an axis of the ambient space that subdivide A into strips of width $2l\varepsilon$, which are left closed and right open. This gives a partition P_0 of A where each strip has width $2l\varepsilon$. Now for any integer i where $0 < i < l$ we shift the hyperplanes that define the partition P_0 by $2i\varepsilon$ to the right to get the partition P_i . Let $S = \{P_0, P_1, \dots, P_{l-1}\}$. Let OPT be the optimal ε -cover of P . Let D be the set of d -balls of radius ε centered at the points in OPT. Any d -ball in D intersects the hyperplanes from at most one partition in S . Therefore there exists a partition P_i such that at most $\frac{|D|}{l}$ d -balls in D are intersected by the hyperplanes defining P_i . In other words at most $\frac{|D|}{l}$ d -balls in D intersect more than one strip in P_i . Now we consider partitioning A similarly along each axis to get a grid of hypercubes of side length $2l\varepsilon$, which we call *cells*. Using the argument described above it follows that there exists a partition P' such that at most $\frac{d|D|}{l}$ d -balls in D intersect more than one cell in P' .

Now consider a cell C of side length $2l\varepsilon$. Since $\dim_f(P) = \delta$ and C is contained in a ball of radius $\sqrt{d}l\varepsilon$ we have that there exists an ε -cover of the points in C of cardinality at most $O(\frac{\sqrt{d}l\varepsilon}{\varepsilon})^\delta = O(\sqrt{d}l)^\delta$.

We combine the above observations to obtain our algorithm as follows. The algorithm enumerates all l^d partitions of P into cells of side length $2l\varepsilon$. Next it enumerates exhaustively all ε -covers of cardinality at most $O((\sqrt{d}l)^\delta)$ for each cell. Since verifying whether a set of points is a valid cover takes time $O(n(\sqrt{d}l)^\delta) = O(nl^\delta)$ this step overall takes time at most $n^{O((\sqrt{d}l)^\delta)} \cdot l^\delta$. Finally the algorithm takes the union of the ε -covers of all the cells to get an ε -cover of P and returns the best solution over all partitions. Since there exists at least one partition where at most $\frac{d|D|}{l}$ d -balls in D intersect more than one cell in the partition, we have that the size of the solution returned is at most $(1 + \frac{d}{l})|D| = (1 + \frac{d}{l})|\text{OPT}|$. The running time of the algorithm is $l^d \cdot n^{O((\sqrt{d}l)^\delta)} \cdot l^\delta = l^{d+\delta} n^{O((l\sqrt{d})^\delta)}$. ◀

► **Theorem 9.** *Let $d \geq 2$ be some integer, and let $\delta \in (0, d]$ be some real number. Let P be a set of n points in \mathbb{R}^d with $\dim_f(P) = \delta$. There exists a polynomial time approximation scheme which given a natural number $l > 0$ and any $\varepsilon > 0$, computes a $(1 + \frac{d}{l-d})$ -approximation to the ε -packing of P , in time $l^{d+\delta} n^{O((l\sqrt{d})^\delta)}$.*

Proof. We use the partitioning approach described in Theorem 8. We consider cells of side length $l\varepsilon$. Since any ε -packing can be augmented into an ε -net we have that any ε -packing of the points in a cell has cardinality at most $O((\frac{\sqrt{d}l}{2})^\delta)$. We consider ε -packings for each cell where the points in the packing are all at least distance $\frac{\varepsilon}{2}$ from the boundary of the cell; this ensures that the d -balls of radius $\frac{\varepsilon}{2}$ centered at these points do not intersect multiple cells. Then we take the union of these points over all cells and take the minimum cardinality set over all partitions. The running time is $l^{d+\delta} n^{O((l\sqrt{d})^\delta)}$ by the same reasoning used in Theorem 8. Let OPT be the optimal ε -packing of P . Since at most $\frac{d}{l}|\text{OPT}|$ d -balls in the optimal packing intersect more than one cell we have that the solution returned by the algorithm has cardinality at least $(1 - \frac{d}{l})|\text{OPT}|$, as required. ◀

6 Spanners and pathwidth

We remark that several other constructions of $(1 + \varepsilon)$ -spanners for finite subsets of d -dimensional Euclidean space are known. However, they do not yield graphs of small pathwidth. Here we use a construction that is a modified version of the spanner due to Vaidya [28]. Let P be a set of n points in \mathbb{R}^d . Let us first recall the construction from [28]. Let $\varepsilon > 0$. We will define a graph G with $V(G) = P$, that is a $(1 + \varepsilon)$ -spanner for P .

Let $I_1, \dots, I_d \subset \mathbb{R}$ be intervals, all having the same length, and such that each I_i is either closed, open, or half-open. Then we say that $b = I_1 \times \dots \times I_d$ is a *box*. We define $\text{size}(b)$ to be the length of the interval I_1 . For each $i \in \{1, \dots, d\}$, let $\psi_i(b)$ be the center of I_i , and define the half-spaces $L_i(b) = \{(x_1, \dots, x_d) \in \mathbb{R}^d : x_i < \psi_i(b)\}$ and $R_i(b) = \{(x_1, \dots, x_d) \in \mathbb{R}^d : x_i \geq \psi_i(b)\}$. Let $S(b)$ be the set of boxes such that $S(b) = \{b' : b' = b \cap (\bigcap_{i=1}^d f_i), \text{ where for all } i \in \{1, \dots, d\}, f_i = L_i(b) \text{ or } f_i = R_i(b)\}$. We also define $\text{shrunk}(b)$ to be some box satisfying the following conditions:

1. If $|b \cap P| \leq 1$ then $\text{shrunk}(b) = b \cap P$. Note that we allow $\text{shrunk}(b)$ to be empty.
2. If $|b \cap P| \geq 2$ then $\text{shrunk}(b)$ is some minimal box contained in b with $\text{shrunk}(b) \cap P = b \cap P$. Note that if there are multiple choices for $\text{shrunk}(b)$, then we choose one arbitrarily.

For some box b with $|b \cap P| \geq 2$, we define $S'(b)$ to be the set of boxes such that $S'(b) = \{b' : \text{there exists } b'' \in S(b) \text{ s.t. } b'' \cap P \neq \emptyset \text{ and } b' = \text{shrunk}(b'')\}$. If $|b \cap P| \leq 1$, then we define $S'(b) = \emptyset$.

The *box-tree* of P is defined to be a tree T where every node is some box. We set the root of T to be some minimal box b^* containing P . For each $b \in V(T)$, the set of children of b in T is $S'(b)$. Note that $|b \cap P| = 1$ if and only if b is a leaf of T . For each $b \in V(T) \setminus \{b^*\}$ we denote by $\text{father}(b)$ the father of b in T .

For each $b \in V(T)$ let

$$\text{Near}(b) = \left\{ b' \in V(T) \setminus \{b^*\} : \text{size}(b') < \text{size}(b) \leq \text{size}(\text{father}(b')) \text{ and } \text{dist}(b, b') \leq \frac{6\sqrt{d}}{\varepsilon} \text{size}(b) \right\}.$$

It follows by the construction that for each $b \in V(T)$, we have $b \cap P \neq \emptyset$. For each $b \in V(T)$ pick some arbitrary point $\text{rep}(b) \in b \cap P$. We say that $\text{rep}(b)$ is the *representative* of b . We further impose the constraint that for each non-leaf $b \in V(T)$, if b' is the unique child of b with $\text{rep}(b) \in b'$, then $\text{rep}(b') = \text{rep}(b)$. This implies that for every $b \in V(T)$, there exists a branch in T starting at b and terminating at some leaf, such that all the boxes in the branch have the same representative as b . We remark that this additional requirement is not necessary in the original construction of Vaidya [28].

We define $E(G) = E_1 \cup E_2$, where $E_1 = \{\{\text{rep}(b), \text{rep}(b')\} : b \in V(T), b' \in S'(b), \text{rep}(b) \neq \text{rep}(b')\}$ and $E_2 = \{\{\text{rep}(b), \text{rep}(b')\} : b \in V(T), b' \in \text{Near}(\text{father}(b))\}$. This completes the description of the spanner construction due to Vaidya [28]. His result is summarized in the following.

► **Theorem 10** (Vaidya [28]). *G is a $(1 + \varepsilon)$ -spanner for P . Moreover $|E(G)| = O(\varepsilon^{-d}n)$.*

For each $e = \{u, v\} \in E(G)$, let D_e be the circumscribed ball for the segment $u-v$. Let $\mathcal{D} = \bigcup_{e \in E(G)} \{D_e\}$. For each $i \in \{1, 2\}$ let $\mathcal{D}_i = \bigcup_{e \in E_i} \{D_e\}$.

► **Lemma 11.** *\mathcal{D}_1 is $(2, d^{O(d)})$ -thick.*

Proof. Let $r > 0$ and define $E_{1,r} = \{\{x, y\} \in E_1 : r \leq \|x - y\|_2 < 2r\}$. Let $\mathcal{D}_{1,r} = \{D_e \in \mathcal{D}_1 : e \in E_{1,r}\}$. It suffices to show that $\mathcal{D}_{1,r}$ is $d^{O(d)}$ -thick.

For each $e = \{x, y\} \in E_{1,r}$ we define some unordered pair of boxes $\gamma(e) = \{B(e), B'(e)\}$, as follows. By the definition of E_1 , there exists some $b \in V(T)$, $b' \in S'(b)$, with $\text{rep}(b) \neq \text{rep}(b')$, such that $\{x, y\} = \{\text{rep}(b), \text{rep}(b')\}$. Assume w.l.o.g. that $x = \text{rep}(b)$ and $y = \text{rep}(b')$. By the choice of the representatives, there exist some branch b_0, \dots, b_t of T , for some $t \geq 1$ with $b_0 = b$, that terminates at some leaf b_t , such that $x = \text{rep}(b) = \text{rep}(b_0) = \dots = \text{rep}(b_t)$. Since $x, y \in b$, it follows that $r \leq \|x - y\|_2 \leq \sqrt{d} \cdot \text{size}(b)$. Since b_t is a leaf, we have $\text{size}(b_t) = 0$. Let $t^* \in \{1, \dots, t\}$ be the maximum integer such that $\text{size}(b_{t^*-1}) \geq r/\sqrt{d}$. Let $A \in S(b_{t^*-1})$ such that $b_{t^*} \subseteq A$. Note that $\text{size}(A) \geq r/(2\sqrt{d})$, and $\text{size}(b_{t^*}) < r/\sqrt{d}$. Pick some box $B(e)$, with $b_{t^*} \subseteq B(e) \subseteq A$, such that

$$\text{size}(B(e)) \in \left[r/(2\sqrt{d}), r/\sqrt{d} \right] \quad (1)$$

in a consistent fashion (i.e. for a fixed choice of b_{t^*} and A we always pick the same box). Similarly, let b'_0, \dots, b'_s be a sequence of boxes such that $b'_0 \in S'(b)$, with $b' \subseteq b'_0$, and b'_1, \dots, b'_s is a branch of T starting at $b'_1 = b'$ and terminating at some leaf b'_s . Arguing as before, let $s^* \in \{1, \dots, s\}$ be the maximum integer such that $\text{size}(b'_{s^*-1}) \geq r/(2\sqrt{d})$. If $s^* = 1$ then let $A' \in S'(b'_{s^*-1})$, with $b'_{s^*} \subseteq A'$; pick some box $B'(e)$, with $b'_{s^*} \subseteq B'(e) \subseteq A'$, such that

$$\text{size}(B'(e)) \in \left[r/(4\sqrt{d}), r/(2\sqrt{d}) \right] \quad (2)$$

in a consistent fashion.

We say that e is *charged* to $\gamma(e)$. By construction, there exists at most one edge in $E_{1,r}$ that is charged to each pair of boxes.

By (1) and (2) we have that for each $e \in E_{1,r}$, the pair $\gamma(e)$ consists of two boxes, each of size $\Theta(r/\sqrt{d})$. Moreover by construction and our choice of boxes we have that for any $e, f \in E_{1,r}$ $B(e)$ and $B(f)$ are disjoint or equal. Similarly $B'(e)$ and $B'(f)$ are also disjoint or equal. Thus, each point in \mathbb{R}^d can be contained in at most $O(1)$ boxes in all the pairs $\gamma(e)$, for all $e \in E_{1,r}$. Moreover, $\text{dist}(B(e), B'(e)) \leq \|x - y\| < 2r$. Thus, each box participates in at most $(\sqrt{d})^{O(d)} = d^{O(d)}$ pairs. For each $e \in E_{1,r}$, let $A(e) = N(B(e), r) \cup N(B'(e), r)$, where $N(X, r)$ denotes the r -neighborhood of X in \mathbb{R}^d . It follows that $\{A(e)\}_{e \in E_{1,r}}$ is $d^{O(d)}$ -thick. Since for each $e \in E_{1,r}$, we have $D_e \in A(e)$, it follows that \mathcal{D}_1 is $d^{O(d)}$ -thick, as required. \blacktriangleleft

► **Lemma 12.** \mathcal{D}_2 is $(2, (d/\varepsilon)^{O(d)})$ -thick.

Proof. Let $r > 0$ and define $E_{2,r} = \{\{x, y\} \in E_2 : r \leq \|x - y\|_2 < 2r\}$. Let $\mathcal{D}_{2,r} = \{D_e \in \mathcal{D}_2 : e \in E_{2,r}\}$. It suffices to show that $\mathcal{D}_{2,r}$ is $d^{O(d)}$ -thick.

As in the proof of Lemma 11, for each $e \in \mathcal{D}_{1,r}$ we define some unordered pair of boxes $\gamma(e) = \{B(e), B'(e)\}$. By the definition of E_2 , there exists some $b \in V(T)$, $b' \in \text{Near}(\text{father}(b))$, such that $\{x, y\} = \{\text{rep}(b), \text{rep}(b')\}$. Assume w.l.o.g. that $x = \text{rep}(b)$ and $y = \text{rep}(b')$. Thus we have $\text{size}(b') < \text{size}(\text{father}(b)) \leq \text{size}(\text{father}(b'))$ and $\text{dist}(b', \text{father}(b)) \leq \frac{6\sqrt{d}}{\varepsilon} \text{size}(\text{father}(b))$. Thus $r \leq \|x - y\|_2 \leq \sqrt{d} \cdot \text{size}(b) + \sqrt{d} \cdot \text{size}(b') + \text{dist}(b, b') < \sqrt{d} \cdot \text{size}(\text{father}(b)) + \sqrt{d} \cdot \text{size}(\text{father}(b)) + \text{dist}(b', \text{father}(b)) + \sqrt{d} \cdot \text{size}(\text{father}(b)) \leq (3 + 6/\varepsilon)\sqrt{d} \cdot \text{size}(\text{father}(b))$. Thus $\text{size}(\text{father}(b)) > r\varepsilon/(9\sqrt{d})$. Let b_0, \dots, b_t be a branch in T with $b_0 = \text{father}(b)$, and $b_t = \{x\}$. Arguing as in Lemma 11, let $t^* \in \{0, \dots, t-1\}$ be the maximum integer such that $\text{size}(b_{t^*}) \geq r\varepsilon/(9\sqrt{d})$. Let $A \in S'(b_{t^*})$, with $b_{t^*+1} \subseteq A$, and pick some box $B(e) \subset A$, with

$$\text{size}(B(e)) \in \left[(r\varepsilon)/(18\sqrt{d}), (r\varepsilon)/(9\sqrt{d}) \right] \quad (3)$$

Similarly, let b'_0, \dots, b'_s be a branch of T with $b'_0 = \text{father}(b')$, and b'_s is a leaf with $b'_s = \{y\}$. Arguing as in Lemma 11, let $s^* \in \{0, \dots, s - 1\}$ be the maximum integer such that $\text{size}(b'_{s^*-1}) \geq (r\varepsilon)/(9\sqrt{d})$. Let $A' \in S'(b_{s^*-1})$ with $b_{s^*} \subseteq A$, and pick some box $B'(e)$, with $b_{s^*} \subseteq A \subseteq B'(e)$, such that

$$\text{size}(B'(e)) \in \left[(\varepsilon r)/(18\sqrt{d}), (\varepsilon r)/(9\sqrt{d}) \right] \tag{4}$$

We say that e is *charged* to $\gamma(e)$. By construction, there exists at most one edge in $E_{1,r}$ that is charged to each pair of boxes.

By (3) and (4) we have that for each $e \in E_{2,r}$, the pair $\gamma(e)$ consists of two boxes, each of size $\Theta((\varepsilon r)/\sqrt{d})$. Thus, each point in \mathbb{R}^d can be contained in at most $O(1)$ distinct boxes in all the pairs $\gamma(e)$, for all $e \in E_{2,r}$. Moreover, $\text{dist}(B(e), B'(e)) \leq \|x - y\| < 2r$. Thus, each box participates in at most $(\sqrt{d}/\varepsilon)^{O(d)} = (d/\varepsilon)^{O(d)}$ pairs. For each $e \in E_{2,r}$, let $A(e) = N(B(e), r) \cup N(B'(e), r)$. It follows that $\{A(e)\}_{e \in E_{2,r}}$ is $(d/\varepsilon)^{O(d)}$ -thick. Since for each $e \in E_{2,r}$, we have $D_e \in A(e)$, it follows that \mathcal{D}_2 is $(d/\varepsilon)^{O(d)}$ -thick, as required. ◀

► **Lemma 13.** \mathcal{D} is $(2, (d/\varepsilon)^{O(d)})$ -thick.

Proof of Lemma 13. By Lemma 11 we have that \mathcal{D}_1 is $(2, d^{O(d)})$ -thick, and by Lemma 12 we have that \mathcal{D}_2 is $(2, (d/\varepsilon)^{O(d)})$ -thick. Since $\mathcal{D} = \mathcal{D}_1 \cup \mathcal{D}_2$, we get that \mathcal{D} is $(2, \kappa)$ -thick, where $\kappa = d^{O(d)} + (d/\varepsilon)^{O(d)} = (d/\varepsilon)^{O(d)}$, as required. ◀

Let $x, y, z, w \in \mathbb{R}^d$. We say that zw is a *shortcut* for xy if the following conditions holds:

1. $\|x - z\|_2 \leq \varepsilon \|z - w\|_2 / 20$.
2. The angle formed by the segments $x-y$ and $x-(w - z + x)$ is at most $\varepsilon/20$.

We now proceed to modify G to obtain a graph G' . Initially, G' contains no edges. We consider all edges in G in increasing order of length. When considering an edge $e = \{x, y\}$, if there exists $\{z, w\} \in E(G')$ such that either zw is a shortcut for xy or zw is a shortcut for yx , then we do not add e to G' ; otherwise we add e to G' . This completes the construction of G' . We next argue that G' is a spanner with low dilation for P . The proof of the following is standard (see e.g. [12]).

► **Lemma 14.** G' is a $(1 + 2\varepsilon)$ -spanner for P .

Due to lack of space, the proof of Lemma 14 is deferred to the full version.

► **Lemma 15.** Let $c \in \mathbb{R}^d$ and let $r > 0$. Let $E^* = \{\{x, y\} \in E(G') : \|x - y\| > 2r \text{ and } x-y \cap \text{sphere}(c, r) \neq \emptyset\}$. Then $|E^*| \leq (d/\varepsilon)^{O(d)}$.

Proof of Lemma 15. Let $E_0^* = \{\{x, y\} \in E^* : \|x - y\| \leq 100r/\varepsilon\}$. We can partition E_0^* into $O(\log(1/\varepsilon))$ buckets, where the i -th bucket contains the balls with radius in $[r2^i, r2^{i+1})$. Since by Lemma 13, \mathcal{D} is $(2, (d/\varepsilon)^{O(d)})$ -thick, and all the balls in E^* are contained in a ball of radius $O(r/\varepsilon)$, it follows that each bucket can contain at most $(1/\varepsilon)^{O(d)} \cdot (d/\varepsilon)^{O(d)}$ balls. Thus $|E_0^*| = O(\log(1/\varepsilon)) \cdot (1/\varepsilon)^{O(d)} \cdot (d/\varepsilon)^{O(d)} = (d/\varepsilon)^{O(d)}$.

Let $E_1^* = E^* \setminus E_0^*$. Suppose that $|E_1^*| > (d/\varepsilon)^{Cd}$. Setting C to be a sufficiently large universal constant it follows that there exist distinct edges $\{x, y\}, \{z, w\} \in E_1^*$ that form an angle of less than $\varepsilon/20$. Assume w.l.o.g. that $\|x - y\|_2 \geq \|z - w\|_2$, $x \in \text{ball}(c, r)$, and $z \in \text{ball}(c, r)$. Then zw must be a shortcut for xy , which is a contradiction since $\{x, y\} \in E(G')$, concluding the proof. ◀

We now prove the main result of this section.

► **Theorem 16.** Let $d \geq 2$ be some fixed integer, and let $\delta \in (0, d]$ be some real number. Let $P \subset \mathbb{R}^d$ be some finite point set with $|P| = n$, such that $\dim_f(P) = \delta$. Then, for any fixed $\varepsilon \in (0, 1]$, there exists a $(1 + \varepsilon)$ -spanner, G' , for P , with a linear number of edges, and with

$$\text{pw}(G) = \begin{cases} O(n^{1-1/\delta} \log n) & \text{if } \delta > 1 \\ O(\log^2 n) & \text{if } \delta = 1 \\ O(\log n) & \text{if } \delta < 1 \end{cases}$$

Moreover, given P , the graph G' can be computed in polynomial time.

Proof. Let G' be the spanner constructed above. The bound on the number of edges of G' follows by Theorem 10 since $G' \subseteq G$. We will bound the pathwidth of G' . By Lemma 13 we have that \mathcal{D} is $(2, (d/\varepsilon)^{O(d)})$ -thick. By Theorem 1 there exists some $(d-1)$ -sphere C with radius r such that at most $(1 - 2^{-O(d)})n$ points of P are contained in either side of C , and $|A| \leq M$, where

$$M = \begin{cases} O(n^{1-1/\delta}) & \text{if } \delta > 1 \\ O(\log n) & \text{if } \delta = 1 \\ O(1) & \text{if } \delta < 1 \end{cases},$$

and $A = \{\{x, y\} \in E(G) : \|x - y\| \leq 2r \text{ and } x-y \cap C \neq \emptyset\}$.

Let $A' = \{\{x, y\} \in E(G') : \|x - y\| \leq 2r \text{ and } x-y \cap C \neq \emptyset\}$. We have $A' \subseteq A$, and thus $|A'| \leq |A|$. Let $A'' = \{\{x, y\} \in E(G') : \|x - y\| > 2r \text{ and } x-y \cap C \neq \emptyset\}$. By Lemma 15 we have $|A''| = O(1)$ (for fixed d and ε). Let S be the set of all endpoints of all the edges in $A' \cup A''$. We have $|S| \leq 2|A' \cup A''| = O(|A|)$. Let U (resp. U') be the set of points in P that are inside (resp. outside) C . Then S separates in G' every vertex in U from every vertex in U' . We may thus recurse on $G'[U \setminus (S \cup U')]$ and $G'[U' \setminus (S \cup U)]$ and obtain path decompositions X_1, \dots, X_t and Y_1, \dots, Y_s respectively. We now obtain the path decomposition $X_1 \cup S, \dots, X_t \cup S, Y_1 \cup S, \dots, Y_s \cup S$ for G' . The width of the resulting path decomposition is at most $O(M \log n)$, concluding the proof. ◀

Acknowledgements. The authors wish to thank Sariel Har-Peled, Dimitrios Thilikos, and Yusu Wang for fruitful discussions.

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 Patrice Assouad. Plongements lipschitziens dans \mathbb{R}^n . *Bulletin de la Société Mathématique de France*, 111:429–448, 1983.
- 3 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.
- 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 T. H. Hubert Chan and Anupam Gupta. Approximating tsp on metrics with bounded global growth. *SIAM Journal on Computing*, 41(3):587–617, 2012.
- 6 T. H. Hubert 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.

- 7 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.
- 8 Kenneth Falconer. *Fractal geometry: mathematical foundations and applications*. John Wiley & Sons, 2004.
- 9 Lee-Ad Gottlieb and Liam Roditty. An optimal dynamic spanner for doubling metric spaces. In *European Symposium on Algorithms*, pages 478–489. Springer, 2008.
- 10 Anupam Gupta, Robert Krauthgamer, and James R. Lee. Bounded geometries, fractals, and low-distortion embeddings. In *Foundations of Computer Science, 2003. Proceedings. 44th Annual IEEE Symposium on*, pages 534–543. IEEE, 2003.
- 11 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.
- 12 Sariel Har-Peled. *Geometric approximation algorithms*, volume 173. American Mathematical Society, Providence, 2011.
- 13 Sariel Har-Peled. A simple proof of the existence of a planar separator. *arXiv preprint arXiv:1105.0103*, 2011.
- 14 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.
- 15 Juha Heinonen. *Lectures on analysis on metric spaces*. Springer Science & Business Media, 2012.
- 16 Dorit S. Hochbaum and Wolfgang Maass. Approximation schemes for covering and packing problems in image processing and VLSI. *Journal of the ACM (JACM)*, 32(1):130–136, 1985.
- 17 David R. Karger and Matthias Ruhl. Finding nearest neighbors in growth-restricted metrics. In *Proceedings of the 34th Annual ACM Symposium on Theory of Computing*, pages 741–750. ACM, 2002.
- 18 Marc Khoury and Rephael Wenger. On the fractal dimension of isosurfaces. *IEEE Transactions on Visualization and Computer Graphics*, 16(6):1198–1205, 2010.
- 19 Robert Krauthgamer and James R. Lee. The black-box complexity of nearest-neighbor search. *Theoretical Computer Science*, 348(2):262–276, 2005.
- 20 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.
- 21 Robert Krauthgamer, James R. Lee, Manor Mendel, and Assaf Naor. Measured descent: A new embedding method for finite metrics. *Geometric & Functional Analysis GFA*, 15(4):839–858, 2005.
- 22 Dániel Marx. Efficient approximation schemes for geometric problems? In *European Symposium on Algorithms*, pages 448–459. Springer, 2005.
- 23 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*, page 67. ACM, 2014.
- 24 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.
- 25 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.
- 26 Hideki Takayasu. *Fractals in the physical sciences*. Manchester University Press, 1990.

58:16 Algorithmic Interpretations of Fractal Dimension

- 27 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.
- 28 Pravin M. Vaidya. A sparse graph almost as good as the complete graph on points in k dimensions. *Discrete & Computational Geometry*, 6(3):369–381, 1991.
- 29 C.T. Zahn. Black box maximization of circular coverage. *Journal of Research of the National Bureau of Standards B*, 66:181–216, 1962.