ON THE STRETCH FACTOR OF POLYGONAL CHAINS[*]

KE CHEN[†], ADRIAN DUMITRESCU[‡], WOLFGANG MULZER[§], AND CSABA D. TÓTH[¶]

Abstract. Let $P = (p_1, p_2, \dots, p_n)$ be a polygonal chain in $\mathbb{R}^d$. The *stretch factor* of P is the ratio between the total length of P and the distance of its endpoints, $\sum_{i=1}^{n-1} |p_i p_{i+1}|/|p_1 p_n|$. For a parameter $c \geq 1$, we call P a *c-chain* if $|p_i p_j| + |p_j p_k| \leq c|p_i p_k|$, for every triple (i, j, k), $1 \leq i < j < k \leq n$. The stretch factor is a global property: it measures how close P is to a straight line, and it involves all the vertices of P; being a c-chain, on the other hand, is a *fingerprint*-property: it only depends on subsets of $O(1)$ vertices of the chain.

We investigate how the c-chain property influences the stretch factor in the plane: (i) we show that for every $\varepsilon > 0$, there is a noncrossing c-chain that has stretch factor $\Omega(n^{1/2-\varepsilon})$, for sufficiently large constant $c = c(\varepsilon)$; (ii) on the other hand, the stretch factor of a c-chain P is $O\left(n^{1/2}\right)$, for every constant $c \geq 1$, regardless of whether P is crossing or noncrossing; and (iii) we give a randomized algorithm that can determine, for a polygonal chain P in $\mathbb{R}^2$ with n vertices, the minimum $c \geq 1$ for which P is a c-chain in $O\left(n^{2.5} \text{ polylog } n\right)$ expected time and $O(n \log n)$ space. These results generalize to $\mathbb{R}^d$. For every dimension $d \geq 2$ and every $\varepsilon > 0$, we construct a noncrossing c-chain that has stretch factor $\Omega\left(n^{(1-\varepsilon)(d-1)/d}\right)$; on the other hand, the stretch factor of any c-chain is $O\left((n-1)^{(d-1)/d}\right)$; for every $c > 1$, we can test whether an n-vertex chain in $\mathbb{R}^d$ is a c-chain in $O\left(n^{3-1/d} \text{ polylog } n\right)$ expected time and $O(n \log n)$ space.

Key words. polygonal chain, vertex dilation, Koch curve, recursive construction

AMS subject classifications. 05C38, 05C40, 68Q25, 68U05

1. Introduction.

Given a set S of n point sites in a Euclidean space $\mathbb{R}^d$, what is the best way to connect S into a *geometric network (graph)*? This question has motivated researchers for a long time, going back as far as the 1940s, and beyond [20, 36]. Numerous possible criteria for a good geometric network have been proposed, perhaps the most basic being the *length*. In 1955, Few [21] showed that for any set of n points in a unit square, there is a traveling salesman tour of length at most $\sqrt{2n} + 7/4$. This was improved to at most $0.984\sqrt{2n} + 11$ by Karloff [24]. Similar bounds hold for the shortest spanning tree and the shortest rectilinear spanning tree [14, 17, 22]. Besides length, two further key factors in the quality of a geometric network are the *vertex dilation* and the *geometric dilation* [32], both of which measure how closely shortest paths in a network approximate the Euclidean distances between their endpoints.

The *dilation* (also called *stretch factor* [30] or *detour* [2]) between two points p and q in a geometric graph G is defined as the ratio between the length of a shortest path from p to q and the Euclidean distance $|pq|$. The *dilation* of the graph G is the maximum dilation over all pairs of vertices in G. A graph in which the dilation is bounded above by $t \geq 1$ is also called a *t-spanner* (or simply a *spanner* if t is a constant). A complete graph in Euclidean space is clearly a 1-spanner. Therefore, researchers focused on the dilation of graphs with certain additional constraints, for

[*]A preliminary version of this paper appeared in the *Proceedings of the 44th International Symposium on Mathematical Foundations of Computer Science*, (MFCS 2019), Aachen, Germany, August 2019, Vol. 138 of LIPIcs, 56:1–56:14.

[†]Department of Computer Science, University of Wisconsin–Milwaukee, USA (kechen@uwm.edu).

[‡]adriandumitrescu.org (ad.dumitrescu@gmail.com).

[§]Institut für Informatik, Freie Universität Berlin, Germany (mulzer@inf.fu-berlin.de). Partially supported by ERC STG 757609.

[¶]Department of Mathematics, California State University Northridge, Los Angeles, CA; and Department of Computer Science, Tufts University, Medford, MA, USA (csaba.toth@csun.edu). Supported in part by NSF CCF-1422311, CCF-1423615, and DMS-1800734.

example, noncrossing (i.e., plane) graphs. In 1989, Das and Joseph [16] identified a large class of plane spanners (characterized by two simple local properties). Bose et al. [7] gave an algorithm that constructs for any set of planar sites a plane 11-spanner with bounded degree. On the other hand, Eppstein [19] analyzed a fractal construction showing that *β-skeletons*, a natural class of geometric networks, can have arbitrarily large dilation.

The study of dilation also raises algorithmic questions. Agarwal et al. [2] described randomized algorithms for computing the dilation of a given path (on n vertices) in $\mathbb{R}^2$ in $O(n \log n)$ expected time. They also presented randomized algorithms for computing the dilation of a given tree, or cycle, in $\mathbb{R}^2$ in $O(n \log^2 n)$ expected time. Previously, Narasimhan and Smid [31] showed that an $(1 + \varepsilon)$-approximation of the stretch factor of any path, cycle, or tree can be computed in $O(n \log n)$ time. Klein et al. [25] gave randomized algorithms for a path, tree, or cycle in $\mathbb{R}^2$ to count the number of vertex pairs whose dilation is below a given threshold in $O(n^{3/2+\varepsilon})$ expected time. Cheong et al. [13] showed that it is NP-hard to determine the existence of a spanning tree on a planar point set whose dilation is at most a given value. More results on plane spanners can be found in the monograph dedicated to this subject [32] or in several surveys [9, 18, 30].

We investigate a basic question about the dilation of polygonal chains. We ask how the dilation between the endpoints of a polygonal chain (which we will call the *stretch factor*, to distinguish it from the more general notion of dilation) is influenced by *fingerprint* properties of the chain, i.e., by properties that are defined on $O(1)$-size subsets of the vertex set. Such fingerprint properties play an important role in geometry; classic examples include the *Carathéodory property*[1] [27, Theorem 1.2.3] or the *Helly property*[2] [27, Theorem 1.3.2]. In general, determining the effect of a fingerprint property may prove elusive—given n points in the plane, consider the simple property that every 3 points determine 3 distinct distances. It is unknown [10, p. 203] whether this property implies that the total number of distinct distances grows superlinearly in n. Furthermore, fingerprint properties appear in the general study of *local versus global properties of metric spaces*, which is highly relevant to combinatorial approximation algorithms based on mathematical programming relaxations [6].

In the study of dilation, interesting fingerprint properties have also been found. For example, a (continuous) curve C is said to have the *increasing chord property* [15, 26] if for any points a, b, c, d that appear on C in this order, we have $|ad| \geq |bc|$. The increasing chord property implies that C has (geometric) dilation at most $2\pi/3$ [34]. A weaker property is the *self-approaching property*: a (continuous) curve C is self-approaching if for any points a, b, c that appear on C in this order, we have $|ac| \geq |bc|$. Self-approaching curves have dilation at most 5.332 [23] (see also [4]), and they have found interesting applications in the field of graph drawing [5, 8, 33].

We introduce a new natural fingerprint property and see that it can constrain the stretch factor of a polygonal chain, but only in a weaker sense than one may expect; we also provide algorithmic results on this property. Before providing details, we give a few basic definitions.

Definitions. A *polygonal chain* P in $\mathbb{R}^d$ is specified by a sequence of n points $(p_1, p_2, \ldots, p_n)$, called *vertices*. The chain P consists of $n - 1$ line segments between

[1] Given a finite set S of points in d dimensions, if every $d + 2$ points in S are in convex position, then S is in convex position.

[2] Given a finite collection of convex sets in d dimensions, if every $d + 1$ sets have nonempty intersection, then all sets have nonempty intersection.

consecutive vertices. We say P is *simple* if only consecutive line segments intersect and they only intersect at their endpoints. Given a polygonal chain P in $\mathbb{R}^d$ with n vertices and a parameter $c \geq 1$, we call P a *c-chain* if for all $1 \leq i < j < k \leq n$, we have

$$|p_i p_j| + |p_j p_k| \leq c|p_i p_k|. \tag{1}$$

Observe that the c-chain condition is a fingerprint condition that is not really a local dilation condition—it is more a combination between the local chain substructure and the distribution of the points in the subchains.

The *stretch factor* δ_P of P is defined as the dilation between the two end points p_1 and p_n of the chain:

$$\delta_P = \frac{\sum_{i=1}^{n-1} |p_i p_{i+1}|}{|p_1 p_n|}.$$

Note that this definition is different from the more general notion of dilation (also called *stretch factor* [30]) of a graph which is the maximum dilation over all pairs of vertices. Since there is no ambiguity in this paper, we will just call δ_P the stretch factor of P.

For example, the polygonal chain $P = ((0,0),(1,0),\ldots,(n,0))$ in $\mathbb{R}^2$ is a 1-chain with stretch factor 1; and $Q = ((0,0),(0,1),(1,1),(1,0))$ is a $(\sqrt{2}+1)$-chain with stretch factor 3.

Without affecting the results, the floor and ceiling functions are omitted in our calculations. For a positive integer t, let $[t] = \{1,2,\ldots,t\}$. For a point set S, let $\mathrm{conv}(S)$ denote the convex hull of S. All logarithms are in base 2, unless stated otherwise.

Our results. In the Euclidean plane $\mathbb{R}^2$, we deduce three upper bounds on the stretch factor of a c-chain P with n vertices (Section 2). In particular, we have (i) $\delta_P \leq c(n-1)^{\log c}$, (ii) $\delta_P \leq c(n-2)+1$, and (iii) $\delta_P = O\left(c^2\sqrt{n-1}\right)$.

From the other direction, we obtain the following lower bound in $\mathbb{R}^2$ (Section 3): For every $c \geq 4$, there is a family $\mathcal{P}_c = \{P^m\}_{m\in\mathbb{N}}$ of simple c-chains, so that P^m has $n = 4^m + 1$ vertices and stretch factor $(n-1)^{\frac{1+\log(c-2)-\log c}{2}}$, where the exponent converges to 1/2 as c tends to infinity. The lower bound construction does not extend to the case of $1 < c < 4$, which remains open.

Then we generalize the results to higher dimensional Euclidean spaces (Section 4): For all integers $d \geq 2$, we show that any c-chain P with n vertices in $\mathbb{R}^d$ has stretch factor $\delta_P = O\left(c^2(n-1)^{(d-1)/d}\right)$. On the other hand, for any constant $\varepsilon > 0$ and sufficiently large $c = \Omega(d)$, we construct a c-chain in $\mathbb{R}^d$ with n vertices and stretch factor at least $(n-1)^{(1-\varepsilon)(d-1)/d}$.

Finally, we present two algorithmic results (Section 5) for all fixed dimensions $d \geq 2$: (i) A randomized algorithm that decides, given a polygonal chain P in $\mathbb{R}^d$ with n vertices and a threshold $c > 1$, whether P is a c-chain in $O\left(n^{3-1/d} \operatorname{polylog} n\right)$ expected time and $O(n \log n)$ space. (ii) As a corollary, there is a randomized algorithm that finds, for a polygonal chain P with n vertices, the minimum $c \geq 1$ for which P is a c-chain in $O\left(n^{3-1/d} \operatorname{polylog} n\right)$ expected time and $O(n \log n)$ space.

2. Upper Bounds in the Plane. At first glance, one might expect the stretch factor of a c-chain, for $c \geq 1$, to be bounded by some function of c. For example, the stretch factor of a 1-chain is necessarily 1. We derive three upper bounds on the stretch factor of a c-chain with n vertices in terms of c and n (cf. Theorems 1–3); see Fig. 1 for a visual comparison between the bounds. For large n, the bound in

Theorem 1 is the best for $1 \leq c \leq 2^{1/2}$, while the bound in Theorem 3 is the best for $c > 2^{1/2}$. In particular, the bound in Theorem 1 is tight for $c = 1$. When n is comparable with c, more specifically, for $c \geq 2$ and $n \leq 64c^2 + 2$, the bound in Theorem 2 is the best.

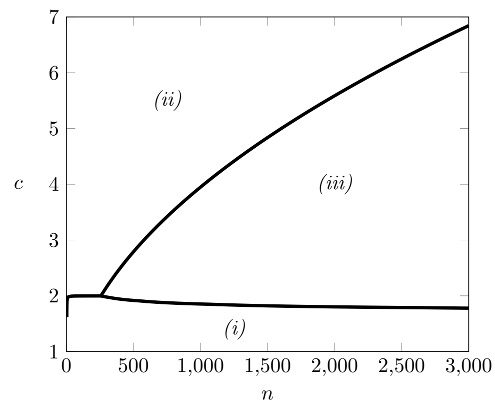

FIG. 1. *The values of n and c for which (i) Theorem 1: $\delta_P \leq c(n-1)^{\log c}$, (ii) Theorem 2: $\delta_P \leq c(n-2)+1$, and (iii) Theorem 3: $\delta_P \leq 8c^2\sqrt{n-1}$ give the current best upper bound.*

Our first upper bound is obtained by a recursive application of the c-chain property. It holds for any positive distance function that need not even satisfy the triangle inequality.

THEOREM 1. *For a c-chain P with n vertices, we have $\delta_P \leq c(n-1)^{\log c}$.*

Proof. We prove, by induction on n, that

$$\delta_P \leq c^{\lceil \log(n-1) \rceil}, \tag{2}$$

for every c-chain P with $n \geq 2$ vertices. In the base case, $n = 2$, we have $\delta_P = 1$ and $c^{\lceil \log(2-1) \rceil} = 1$. Now let $n \geq 3$, and assume that (2) holds for every c-chain with fewer than n vertices. Let $P = (p_1, \ldots, p_n)$ be a c-chain with n vertices. Then, applying (2) to the first and second half of P, followed by the c-chain property for the first, middle, and last vertex of P, we get

$$\begin{aligned}
\sum_{i=1}^{n-1} |p_i p_{i+1}| &\leq \sum_{i=1}^{\lceil n/2 \rceil - 1} |p_i p_{i+1}| + \sum_{i=\lceil n/2 \rceil}^{n-1} |p_i p_{i+1}| \\
&\leq c^{\lceil \log(\lceil n/2 \rceil - 1) \rceil} \left(|p_1 p_{\lceil n/2 \rceil}| + |p_{\lceil n/2 \rceil} p_n| \right) \\
&\leq c^{\lceil \log(\lceil n/2 \rceil - 1) \rceil} \cdot c |p_1 p_n| \\
&\leq c^{\lceil \log(n-1) \rceil} |p_1 p_n|,
\end{aligned}$$

so (2) holds also for P. Consequently,

$$\delta_P \leq c^{\lceil \log(n-1) \rceil} \leq c^{\log(n-1)+1} = c \cdot c^{\log(n-1)} = c\,(n-1)^{\log c},$$

as required. ∎

Our second upper bound combines the c-chain property with the triangle inequality, and it holds in any metric space.

THEOREM 2. *For a c-chain P with n vertices, we have $\delta_P \le c(n-2)+1$.*

Proof. Without loss of generality, assume that $|p_1p_n| = 1$. For every $1 < i < n$, the c-chain property implies $|p_1p_i| + |p_ip_n| \le c|p_1p_n| = c$, hence

$$|p_1p_i| \le c - |p_ip_n|. \tag{3}$$

The triangle inequality yields

$$|p_1p_i| \le |p_1p_n| + |p_np_i| = 1 + |p_ip_n|. \tag{4}$$

The combination of (3) and (4) gives $|p_1p_i| \le \frac{c+1}{2}$. Analogous argument for p_n (in place of p_1) yields $|p_ip_n| \le \frac{c+1}{2}$.

For every pair $1 < i < j < n$, the triangle inequality implies

$$2|p_ip_j| \le (|p_ip_1|+|p_1p_j|)+(|p_ip_n|+|p_np_j|) = (|p_1p_i|+|p_ip_n|)+(|p_1p_j|+|p_jp_n|) \le 2c,$$

hence $|p_ip_j| \le c$. Overall, the stretch factor of P is bounded above by

$$\begin{aligned}\delta_P = \frac{\sum_{j=1}^{n-1}|p_jp_{j+1}|}{|p_1p_n|} &= |p_1p_2| + |p_{n-1}p_n| + \sum_{j=2}^{n-2}|p_jp_{j+1}| \\ &\le \frac{c+1}{2} + \frac{c+1}{2} + c(n-3) = c(n-2)+1,\end{aligned}$$

as claimed. ∎

Our third upper bound uses properties of the Euclidean plane (specifically, a volume argument) to bound the number of long edges in P.

THEOREM 3. *For a c-chain P with n vertices, we have $\delta_P = O\left(c^2\sqrt{n-1}\right)$.*

Proof. Let $P = (p_1, \ldots, p_n)$ be a c-chain, for some constant $c \ge 1$, and let $L = \sum_{i=1}^{n-1}|p_ip_{i+1}|$ be its length. We may assume that p_1p_n is a horizontal segment of unit length. By the c-chain property, every point p_j, $1 < j < n$, lies in an ellipse E with foci p_1 and p_n; see FIG. 2. The diameter of E is its major axis, whose length is c. Let U be a disk of radius $c/2$ concentric with E, and note that $E \subset U$

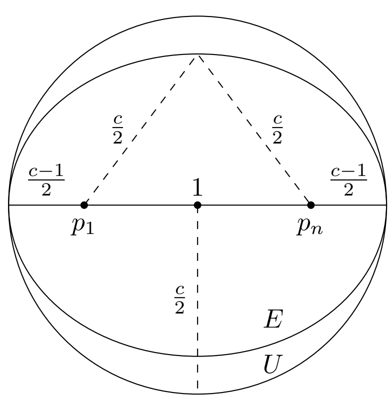

FIG. 2. *The entire chain P lies in an ellipse E with foci p_1 and p_n. E lies in a cocentric disk U of radius $c/2$.*

We set $x = 4c^2/\sqrt{n-1}$; and let L_0 and L_1 be the sum of lengths of all edges in P of length at most x and more than x, respectively. By definition, we have $L = L_0 + L_1$ and

$$L_0 \leq (n-1)x = (n-1) \cdot 4c^2/\sqrt{n-1} = 4c^2\sqrt{n-1}. \tag{5}$$

We shall prove that $L_1 \leq 4c^2\sqrt{n-1}$, implying $L \leq 8c^2\sqrt{n-1}$. For this, we further classify the edges in L_1 according to their lengths: For $\ell = 0, 1, \ldots, \infty$, let

$$P_\ell = \left\{p_i : 2^\ell x < |p_i p_{i+1}| \leq 2^{\ell+1} x\right\}. \tag{6}$$

Since all points lie in an ellipse of diameter c, we have $|p_i p_{i+1}| \leq c$, for all $i = 0, \ldots, n-1$. Consequently, $P_\ell = \emptyset$ when $c \leq 2^\ell x$, or equivalently $\log(c/x) \leq \ell$.

We use a volume argument to derive an upper bound on the cardinality of P_ℓ, for $\ell = 0, 1, \ldots, \lfloor \log(c/x) \rfloor$. Assume that $p_i, p_k \in P_\ell$, and w.l.o.g., $i < k$. If $k = i+1$, then by (6), $2^\ell x < |p_i p_k|$. Otherwise,

$$2^\ell x < |p_i p_{i+1}| < |p_i p_{i+1}| + |p_{i+1} p_k| \leq c|p_i p_k|, \text{ or } \frac{2^\ell x}{c} < |p_i p_k|.$$

Consequently, the disks of radius

$$R = \frac{2^\ell x}{2c} = \frac{2 \cdot 2^\ell c}{\sqrt{n-1}} \tag{7}$$

centered at the points in P_ℓ are interior-disjoint. The area of each disk is πR^2. Since $P_\ell \subset U$, these disks are contained in the R-neighborhood U_R of the disk U, which is a disk of radius $\frac{c}{2} + R$ concentric with U. For $\ell \leq \log(c/x)$, we have $2^\ell x \leq c$, hence $R = \frac{2^\ell x}{2c} \leq \frac{c}{2c} = \frac{1}{2} \leq \frac{c}{2}$. Thus the radius of U_R is at most c. Since U_R contains $|P_\ell|$ interior-disjoint disks of radius R, we obtain

$$|P_\ell| \leq \frac{\text{area}(U_R)}{\pi R^2} < \frac{\pi c^2}{\pi R^2} = \frac{4c^4}{2^{2\ell} x^2}. \tag{8}$$

For every segment $p_{i-1}p_i$ with length more than x, we have that $p_i \in P_\ell$, for some $\ell \in \{0, 1, \ldots, \lfloor \log(c/x) \rfloor\}$. The total length of these segments is

$$\begin{aligned} L_1 \leq \sum_{\ell=0}^{\lfloor \log(c/x) \rfloor} |P_\ell| \cdot 2^{\ell+1} x &< \sum_{\ell=0}^{\lfloor \log(c/x) \rfloor} \frac{4c^4}{2^{2\ell} x^2} \cdot 2^{\ell+1} x = \sum_{\ell=0}^{\lfloor \log(c/x) \rfloor} \frac{8c^4}{2^\ell x} \\ &< \frac{8c^4}{x} \sum_{\ell=0}^{\infty} \frac{1}{2^\ell} = \frac{16c^4}{x} = 4c^2 \cdot \sqrt{n-1}, \end{aligned}$$

as required. Together with (5), this yields $L \leq 8c^2 \cdot \sqrt{n-1}$. ∎

3. Lower Bounds in the Plane. We now present our lower bound construction, showing that the dependence on n for the stretch factor of a c-chain cannot be avoided.

THEOREM 4. *For every constant $c \geq 4$, there is a set $\mathcal{P}_c = \{P^m\}_{m \in \mathbb{N}}$ of simple c-chains, so that P^m has $n = 4^m + 1$ vertices and stretch factor $(n-1)^{\frac{1+\log(c-2)-\log c}{2}}$.*

By Theorem 3, the stretch factor of a c-chain in the plane is $O\left((n-1)^{1/2}\right)$ for every constant $c \geq 1$. Since

$$\lim_{c\to\infty} \frac{1+\log(c-2)-\log c}{2} = \frac{1}{2},$$

our lower bound construction shows that the limit of the exponent cannot be improved. Indeed, for every $\varepsilon > 0$, we can set $c = \frac{2^{2\varepsilon+1}}{2^{2\varepsilon}-1}$, and then the chains above have stretch factor

$$(n-1)^{\frac{1+\log(c-2)-\log c}{2}} = (n-1)^{1/2-\varepsilon} = \Omega(n^{1/2-\varepsilon}).$$

We first construct a family $\mathcal{P}_c = \{P^m\}_{m\in\mathbb{N}}$ of polygonal chains. Then we show, in Lemmata 5 and 7, that every chain in $\mathcal{P}_c$ is simple and indeed a c-chain. The theorem follows since the claimed stretch factor is a consequence of the construction.

Construction of $\mathcal{P}_c$. The construction here is a generalization of the iterative construction of the *Koch curve*; when $c = 6$, the result is the original Cesàro fractal (which is a variant of the Koch curve) [11]. We start with a unit line segment P^0, and for $m = 0, 1, \ldots$, we construct P^{m+1} by replacing each segment in P^m by four segments such that the middle three points achieve a stretch factor of $c_* = \frac{c-2}{2}$ (this choice will be justified in the proof of Lemma 7). Note that $c_* \geq 1$, since $c \geq 4$.

We continue with the details. Let P^0 be the unit line segment from $(0, 0)$ to $(1, 0)$; see FIG. 3 (left). Given the polygonal chain P^m ($m = 0, 1, \ldots$), we construct P^{m+1} by replacing each segment of P^m by four segments as follows. Consider a segment of P^m, and denote its length by ℓ. Subdivide this segment into three segments of lengths $(\frac{1}{2} - \frac{a}{c_*})\ell$, $\frac{2a}{c_*}\ell$, and $(\frac{1}{2} - \frac{a}{c_*})\ell$, respectively, where $0 < a < \frac{c_*}{2}$ is a parameter to be determined later. Replace the middle segment with the top part of an isosceles triangle of side length $a\ell$. The chains P^0, P^1, P^2, and P^4 are depicted in Figures 3 and 4.

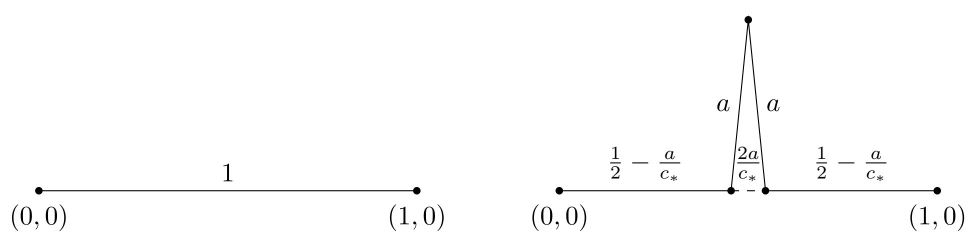

FIG. 3. *The chains P^0 (left) and P^1 (right).*

Note that each segment of length ℓ in P^m is replaced by four segments of total length $(1 + \frac{2a(c_*-1)}{c_*})\ell$. After m iterations, the chain P^m consists of 4^m line segments of total length $\left(1 + \frac{2a(c_*-1)}{c_*}\right)^m$.

By construction, the chain P^m (for $m \geq 1$) consists of four scaled copies of P^{m-1}. For $i = 1, 2, 3, 4$, let the *ith subchain of* P^m be the subchain of P^m consisting of 4^{m-1} segments starting from the $((i-1)4^{m-1}+1)$th segment. By construction, the ith subchain of P^m is similar to the chain P^{m-1}, for $i = 1, 2, 3, 4$.[3] The following functions allow us to refer to these subchains formally. For $i = 1, 2, 3, 4$, define a function $f_i^m : P^m \to P^m$ as the identity on the ith subchain of P^m that sends the

[3]Two geometric shapes are *similar* if one can be obtained from the other by translation, rotation, and scaling; and are *congruent* if one can be obtained from the other by translation and rotation.

remaining part(s) of P^m to the closest endpoint(s) along this subchain. So $f_i^m(P^m)$ is similar to P^{m-1}. Let $g_i : \mathcal{P}_c \setminus \{P^0\} \to \mathcal{P}_c$ be a piecewise defined function such that $g_i(C) = \sigma^{-1} \circ f_i^m \circ \sigma(C)$ if C is similar to P^m, where $\sigma : C \to P^m$ is a similarity transformation. Applying the function g_i on a chain P^m can be thought of as "cutting out" its ith subchain.

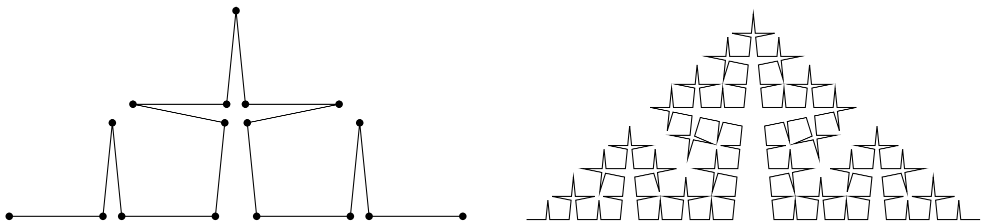

FIG. 4. *The chains P^2 (left) and P^4 (right).*

Clearly, the stretch factor of the chain monotonically increases with the parameter a. However, if a is too large, the chain is no longer simple. The following lemma gives a sufficient condition for the constructed chains to avoid self-crossings.

LEMMA 5. *For every constant $c \geq 4$, if $a \leq \frac{c-2}{2c}$, then every chain in $\mathcal{P}_c$ is simple.*

Proof. Let $T = \mathrm{conv}(P^1)$. Observe that T is an isosceles triangle; see FIG. 5 (left). We first show the following:

CLAIM 6. *If $a \leq \frac{c-2}{2c}$, then $\mathrm{conv}(P^m) = T$ for all $m \geq 1$.*

Proof. We prove the claim by induction on m. It holds for $m = 1$ by definition. For the induction step, assume that $m \geq 2$ and that the claim holds for $m - 1$. Consider the chain P^m. Since it contains all the vertices of P^1, $T \subset \mathrm{conv}(P^m)$. So we only need to show that $\mathrm{conv}(P^m) \subset T$.

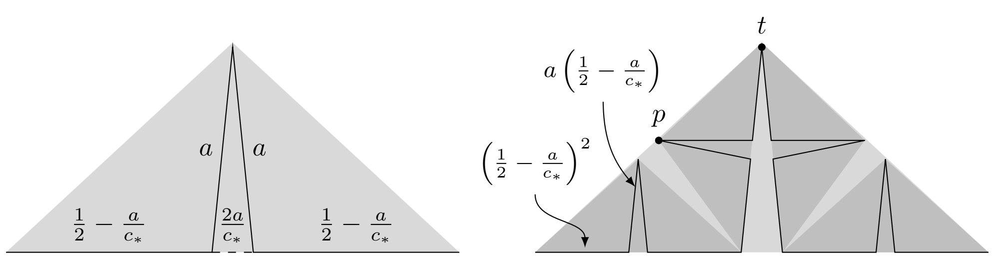

FIG. 5. *Left: Convex hull T of P^1 in light gray; Right: Convex hulls of $g_i(P^2)$, $i = 1, 2, 3, 4$, in dark gray, are contained in T.*

By construction, $P^m \subset \bigcup_{i=1}^{4} \mathrm{conv}(g_i(P^m))$; see FIG. 5 (right). By the inductive hypothesis, $\mathrm{conv}(g_i(P^m))$ is an isosceles triangle similar to T, for $i = 1, 2, 3, 4$. Since the bases of $\mathrm{conv}(g_1(P^m))$ and $\mathrm{conv}(g_4(P^m))$ are collinear with the base of T by construction, due to similarity, they are contained in T. The base of $\mathrm{conv}(g_2(P^m))$ is contained in T. In order to show $\mathrm{conv}(g_2(P^m)) \subset T$, by convexity, it suffices to ensure that its apex p is also in T. Note that the coordinates of the top point is

$t = \left(1/2, a\sqrt{c_*^2 - 1}/c_*\right)$, so the supporting line ℓ of the left side of T is

$$y = \frac{2a\sqrt{c_*^2-1}}{c_*}x, \text{ and}$$
$$p = \left(\frac{1}{2} - \frac{a}{2c_*} - \frac{a^2\left(c_*^2-1\right)}{c_*^2}, \left(\frac{a}{2c_*} + \frac{a^2}{c_*^2}\right)\sqrt{c_*^2-1}\right).$$

By the condition of $a \le \frac{c-2}{2c} = \frac{c_*}{2(c_*+1)}$ in the lemma, p lies on or below ℓ. Under the same condition, we have $\mathrm{conv}(g_3(P^m)) \subset T$ by symmetry. Then $P^m \subset \bigcup_{i=1}^4 \mathrm{conv}(g_i(P^m)) \subset T$. Since T is convex, $\mathrm{conv}(P^m) \subset T$. So $\mathrm{conv}(P^m) = T$, as claimed. ∎

We can now finish the proof of Lemma 5 by induction. Clearly, P^0 and P^1 are simple. Assume that $m \ge 2$, and P^{m-1} is simple. Consider the chain P^m. For $i = 1, 2, 3, 4$, $g_i(P^m)$ is similar to P^{m-1}, hence simple by the inductive hypothesis. Since $P^m = \bigcup_{i=1}^4 g_i(P^m)$, it is sufficient to show that for all $i, j \in \{1, 2, 3, 4\}$, where $i \ne j$, a segment in $g_i(P^m)$ does not intersect any segments in $g_j(P^m)$, unless they are consecutive in P^m and they intersect at a common endpoint. This follows from the above claim together with the observation that for $i \ne j$, the intersection $g_i(P^m) \cap g_j(P^m)$ is either empty or contains a single vertex which is the common endpoint of two consecutive segments in P^m. ∎

In the remainder of this section, we assume that

$$a = \frac{c-2}{2c} = \frac{c_*}{2(c_*+1)}. \tag{9}$$

Under this assumption, all segments in P^1 have the same length a. Therefore, by construction, all segments in P^m have the same length

$$a^m = \left(\frac{c_*}{2(c_*+1)}\right)^m.$$

There are 4^m segments in P^m, with $4^m + 1$ vertices, and its stretch factor is

$$\delta_{P^m} = 4^m\left(\frac{c_*}{2(c_*+1)}\right)^m = \left(\frac{2c_*}{c_*+1}\right)^m.$$

Consequently, $m = \log_4(n-1) = \frac{\log(n-1)}{2}$, and

$$\delta_{P^m} = \left(\frac{2c_*}{c_*+1}\right)^{\frac{\log(n-1)}{2}} = \left(\frac{2c-4}{c}\right)^{\frac{\log(n-1)}{2}} = (n-1)^{\frac{1+\log(c-2)-\log c}{2}},$$

as claimed. To finish the proof of Theorem 4, it remains to show the constructed polygonal chains are indeed c-chains.

Lemma 7. *For every constant $c \ge 4$, $\mathcal{P}_c$ is a family of c-chains.*

We first prove a couple of facts that will be useful in the proof of Lemma 7. We defer an intuitive explanation until after the formal statement of the following lemma.

Lemma 8. *Let $m \ge 1$ and let $P^m = (p_1, p_2, \ldots, p_n)$, where $n = 4^m + 1$. Then the following hold:*

(i) *There exists a sequence* $(q_1, q_2, \ldots, q_\ell)$ *of* $\ell = 2 \cdot 4^{m-1}$ *points in* $\mathbb{R}^2$ *such that the chain* $R^m = (p_1, q_1, p_2, q_2, \ldots, p_\ell, q_\ell, p_{\ell+1})$ *is similar to* P^m*.*

(ii) *For* $m \geq 2$*, define* $g_5 : \mathcal{P}_c \setminus \{P^0, P^1\} \to \mathcal{P}_c$ *by*

$$g_5(P^m) = (g_3 \circ g_2(P^m)) \cup (g_4 \circ g_2(P^m)) \cup (g_1 \circ g_3(P^m)) \cup (g_2 \circ g_3(P^m)).$$

Then $g_5(P^m)$ *is similar to* P^{m-1}*.*

Part (i) of Lemma 8 says that given P^m, we can construct a chain R^m similar to P^m by inserting one point between every two consecutive points of the left half of P^m, see Fig. 6 (left). Part (ii) says that the "top" subchain of P^m that consists of the right half of $g_2(P^m)$ and the left half of $g_3(P^m)$, see Fig. 6 (right), is similar to P^{m-1}.

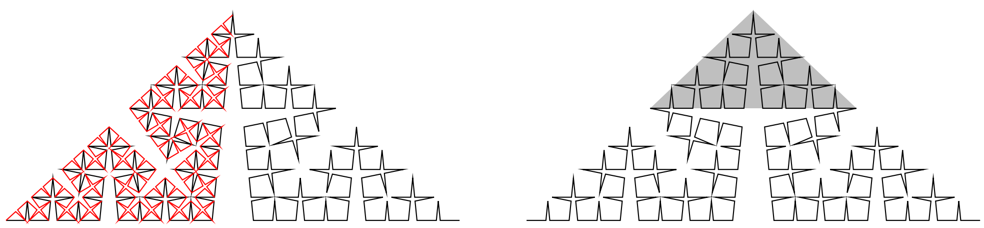

Fig. 6. *Left: Chain* P^m *with the scaled copy of itself* R^m *(in red); Right: Chain* P^m *with its subchain* $g_5(P^m)$ *marked by its convex hull.*

Proof of Lemma 8. For part (i), we review the construction of P^m, and show that R^m and P^m can be constructed in a coupled manner. In Fig. 7 (left), consider $P^1 = (p_1, p_2, p_3, p_4, p_5)$. Recall that all segments in P^1 are of the same length $a = \frac{c_*}{2(c_*+1)}$. The isosceles triangles $\Delta p_1p_2p_3$ and $\Delta p_1p_3p_5$ are similar. Let $\sigma : \Delta p_1p_3p_5 \to \Delta p_1p_2p_3$ be the similarity transformation. Let $q_1 = \sigma(p_2)$ and $q_2 = \sigma(p_4)$. By construction, the chain $R^1 = (p_1, q_1, p_2, q_2, p_3)$ is similar to P^1. In particular, all of its segments have the same length, and so the isosceles triangle $\Delta p_1q_1p_2$ is similar to $\Delta p_1p_3p_5$. Moreover, its base is the segment p_1p_2, so $\Delta p_1q_1p_2$ is precisely $\text{conv}(g_1(P^2))$, see Fig. 7 (right).

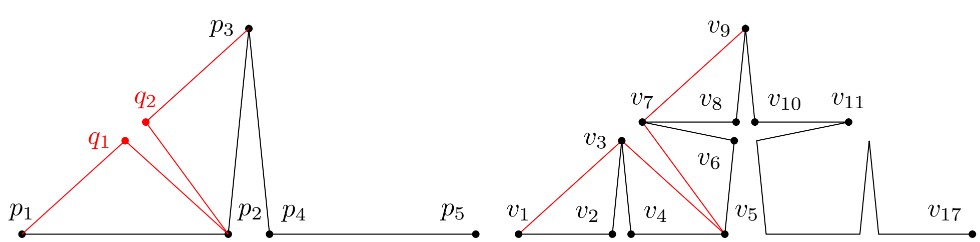

Fig. 7. *Left: the chains* P^1 *and* R^1 *(red); Right: the chains* P^2 *and* R^1 *(red).*

Write $P^2 = (v_1, v_2, \ldots, v_{17})$, then $v_3 = q_1$ by the above argument and $v_7 = q_2$ by symmetry. Now $\Delta v_1v_2v_3$, $\Delta v_3v_4v_5$, $\Delta v_5v_6v_7$, and $\Delta v_7v_8v_9$ are four congruent isosceles triangles, all of which are similar to $\Delta v_1v_9v_{17}$, since the angles are the same. Repeat the above procedure on each of them to obtain $R^2 = (v_1, u_1, v_2, u_2, \ldots, v_8, u_8, v_9)$, which is similar to P^2. Continue this construction inductively to get the desired chain R^m for any $m \geq 1$.

For part (ii), see FIG. 7 (right). By definition, $g_5(P^2)$ is the subchain $(v_7, v_8, v_9, v_{10}, v_{11})$. Observe that the segments v_7v_8 and $v_{10}v_{11}$ are collinear by symmetry. Moreover, they are parallel to v_1v_{17} since $\angle v_7v_8v_9 = \angle v_1v_5v_9$. So $g_5(P^2)$ is similar to P^1; see FIG. 7 (left). Then for $m \geq 2$, $g_5(P^m)$ is the subchain of P^m starting at vertex v_7, ending at vertex v_{11}. By the construction of P^m, $g_5(P^m)$ is similar to P^{m-1}. ◻

Proof of Lemma 7. We proceed by induction on m again. The claim is vacuously true for P^0. For P^1, among all ten choices of $1 \leq i < j < k \leq 5$, $\frac{|p_2p_3|+|p_3p_4|}{|p_2p_4|} = c_* = \frac{c-2}{2} < c$ is the largest, and so P^1 is also a c-chain. Assume that $m \geq 2$ and P^{m-1} is a c-chain. We need to show that P^m is also a c-chain. Consider a triplet of vertices $\{p_i, p_j, p_k\} \subset P^m$, where $1 \leq i < j < k \leq n = 4^m + 1$.

Recall that P^m consists of four copies of the subchain P^{m-1}, namely $g_1(P^m)$, $g_2(P^m)$, $g_3(P^m)$, and $g_4(P^m)$, see FIG. 8 (left). If $\{p_i, p_j, p_k\} \subset g_l(P^m)$ for any $l = 1, 2, 3, 4$, then by the induction hypothesis,

$$\frac{|p_ip_j| + |p_jp_k|}{|p_ip_k|} \leq c.$$

So we may assume that p_i and p_k belong to two different $g_l(P^m)$'s. There are four cases to consider up to symmetry:

Case 1. $p_i \in g_1(P^m)$ and $p_k \in g_2(P^m)$;
Case 2. $p_i \in g_1(P^m)$ and $p_k \in g_3(P^m)$;
Case 3. $p_i \in g_1(P^m)$ and $p_k \in g_4(P^m)$;
Case 4. $p_i \in g_2(P^m)$ and $p_k \in g_3(P^m)$.

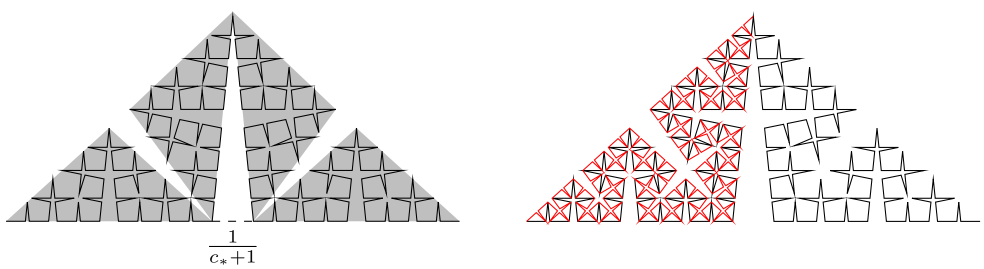

FIG. 8. *Left: Chain P^m with its four subchains of type P^{m-1} marked by their convex hulls; Right: Chain P^m with the scaled copy of itself R^m (in red) constructed in Lemma 8 (i).*

By Lemma 8 (i), the vertex set of $g_1(P^m) \cup g_2(P^m)$ is contained in the chain R^m shown in FIG. 8 (right). If we are in Case 1, i.e., $p_i \in g_1(P^m)$ and $p_k \in g_2(P^m)$, then p_i, p_j, p_k can be thought of as vertices of R^m. The similarity between R^m and P^m, maps points p_i, p_j, p_k to suitable points $p'_i, p'_j, p'_k \in P^m$ such that

$$\frac{|p'_ip'_j| + |p'_jp'_k|}{|p'_ip'_k|} = \frac{|p_ip_j| + |p_jp_k|}{|p_ip_k|}.$$

Since $p_i \in g_1(R^m) \cup g_2(R^m)$ while $p_k \in g_3(R^m) \cup g_4(R^m)$, the triplet (p'_i, p'_j, p'_k) does not belong to Case 1. In other words, Case 1 can be represented by other cases.

Recall that in Lemma 5, we showed that $\operatorname{conv}(P^m)$ is an isosceles triangle T of diameter 1. Observe that if $|p_ip_k| \geq \frac{1}{c_*+1}$, then

$$\frac{|p_ip_j| + |p_jp_k|}{|p_ip_k|} \leq \frac{1+1}{\frac{1}{c_*+1}} = 2c_* + 2 = c,$$

as required. So we may assume that $|p_i p_k| < \frac{1}{c_*+1}$, therefore only Case 4 remains, i.e., $p_i \in g_2(P^m)$ and $p_k \in g_3(P^m)$.

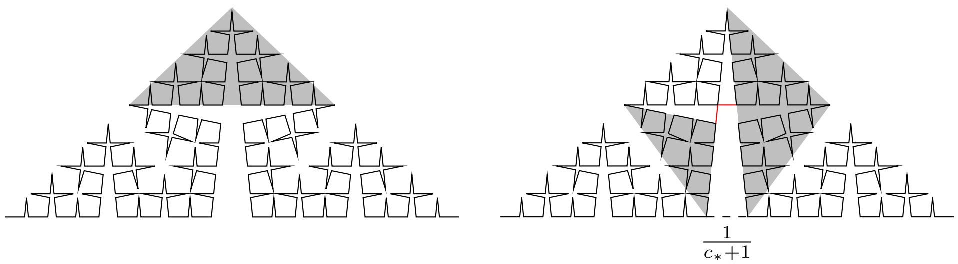

FIG. 9. *Left: Chain P^m with its subchain $g_5(P^m)$ marked by its convex hull; Right: The last case where p_i is in the left shaded subchain and p_k is in the right shaded subchain.*

By Lemma 8 (ii), the "top" subchain $g_5(P^m)$ of P^m is also similar to P^{m-1}, see FIG. 9 (left). If p_i and p_k are both in $g_5(P^m)$, i.e., $p_i \in (g_3 \circ g_2(P^m)) \cup (g_4 \circ g_2(P^m))$ and $p_k \in (g_1 \circ g_3(P^m)) \cup (g_2 \circ g_3(P^m))$, then so is p_j.

By the induction hypothesis, we have

$$\frac{|p_i p_j| + |p_j p_k|}{|p_i p_k|} \le c.$$

So we may assume that at least one of p_i and p_k is not in $g_5(P^m)$. Without loss of generality, let $p_i \in g_2(P^m) \setminus g_5(P^m)$. The similarity that maps P^{m-1} to $g_2(P^m)$ and $g_5(P^m)$, respectively, have the same scaling factor of $a = \frac{c_*}{2(c_*+1)}$, and they carry the bottom dashed segment in FIG. 9 (right), to the two red segments.

CLAIM 9. *If $p_i \in g_2(P^m) \setminus g_5(P^m)$ and $p_k \in g_3(P^m)$, then $|p_i p_k| > \frac{c_*}{2(c_*+1)^2}$.*

Proof. As noted above, we assume that p_i is in $\mathrm{conv}(g_2(P^m) \setminus g_5(P^m)) = \Delta q_1 q_2 q_3$ in FIG. 10. If $p_k \in g_5(P^m) \cap g_3(P^m) = \Delta q_7 q_6 q_5$, then the configuration is illustrated in FIG. 10 (left). Note that $\Delta q_1 q_2 q_3$ and $\Delta q_7 q_6 q_5$ are reflections of each other with respect to the bisector of $\angle q_3 q_4 q_5$. Hence the shortest distance between $\Delta q_1 q_2 q_3$ and $\Delta q_7 q_6 q_5$ is $\min\{|q_3 q_5|, |q_2 q_6|, |q_1 q_7|\}$. Since $c_* \ge 1$, we have

$$|q_1 q_7| > |q_7 q_9| = |q_3 q_5| = a^{3/2} = \left(\frac{c_*}{2(c_*+1)}\right)^{3/2} \ge \frac{c_*}{2(c_*+1)^2}.$$

Further note that $q_2 q_4 q_6 q_8$ is an isosceles trapezoid, so the length of its diagonal is bounded by $|q_2 q_6| > |q_2 q_4| = \frac{c_*}{2(c_*+1)^2}$. Therefore the claim holds when $p_k \in \Delta q_7 q_6 q_5$.

Otherwise $p_k \in g_3(P^m) \setminus g_5(P^m) = \Delta q_9 q_8 q_7$: see FIG. 10 (right). Note that $\Delta q_1 q_2 q_3$ and $\Delta q_9 q_8 q_7$ are reflections of each other with respect to the bisector of $\angle q_4 q_5 q_6$. So the shortest distance between the shaded triangles is the minimum between $|q_3 q_7|$, $|q_2 q_8|$, and $|q_1 q_9|$. However, all three candidates are strictly larger than $|q_4 q_6| = \frac{c_*}{2(c_*+1)^2}$. This completes the proof of the claim. ∎

Now the diameter of $g_2(P^m) \cup g_3(P^m)$ is $a = \frac{c_*}{2(c_*+1)}$ (note that there are three diameter pairs), so

$$\frac{|p_i p_j| + |p_j p_k|}{|p_i p_k|} < \frac{2 \cdot \frac{c_*}{2(c_*+1)}}{\frac{c_*}{2(c_*+1)^2}} = 2c_* + 2 = c,$$

as required. This concludes the proof of Lemma 7 and Theorem 4. ∎

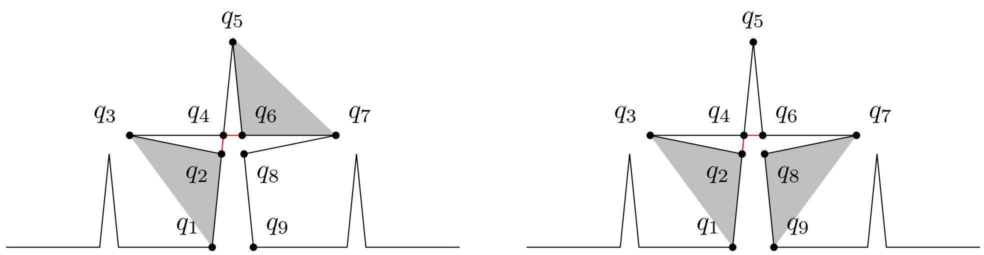

FIG. 10. $p_i \in \Delta q_1 q_2 q_3$, *Left:* $p_k \in \Delta q_7 q_6 q_5$*; Right:* $p_k \in \Delta q_9 q_8 q_7$.

4. Generalizations to Higher Dimensions.

A c-chain P with n vertices and its stretch factor δ_P can be defined in any metric space, not just the Euclidean plane. We now discuss how our results generalize to other metric spaces, with a particular focus on the high-dimensional Euclidean space $\mathbb{R}^d$. First, we examine the upper bounds from Section 2.

4.1. Upper bounds.

As already noted in Section 2, the upper bound $\delta_P \leq c(n-1)^{\log c}$ of Theorem 1 holds for any positive distance function that need not even satisfy the triangle inequality.

Theorem 2 uses only the triangle inequality, and the bound $\delta_P \leq c(n-2)+1$ holds in any metric space. This bound cannot be improved, in the following sense: For every $c \geq 2+\sqrt{5}$ and even n, we can define a finite metric space on the vertex set of P by $|p_1p_n| = 1$; for $1 < i < n$,

$$|p_1p_i| = \begin{cases} \frac{c+1}{2} & \text{if } i \text{ is even} \\ \frac{c-1}{2} & \text{if } i \text{ is odd} \end{cases} \text{ and } |p_ip_n| = \begin{cases} \frac{c-1}{2} & \text{if } i \text{ is even} \\ \frac{c+1}{2} & \text{if } i \text{ is odd} \end{cases};$$

and $|p_ip_j| = c$ for all $1 < i < j < n$. It is easy to verify that P is a c-chain (the case that puts the strongest constraint on c in (1) occurs if, e.g., $i = 1$, $1 < j < n$ is even, and $j < k < n$ is odd) and that P has stretch factor

$$\delta_P = \frac{\sum_{i=1}^{n-1} |p_ip_{i+1}|}{|p_1p_n|} = |p_1p_2| + |p_{n-1}p_n| + \sum_{i=2}^{n-2} |p_ip_{i+1}| = c(n-2)+1.$$

The proof of Theorem 3 uses a volume argument in the plane. The argument extends to $\mathbb{R}^d$, for all constant dimensions $d \geq 2$, and yields $\delta_P = O\left(c^2(n-1)^{(d-1)/d}\right)$.

THEOREM 10. *For a c-chain P with n vertices in $\mathbb{R}^d$, for some constant $d \geq 2$, we have*

$$\delta_P = O\left(c^2(n-1)^{(d-1)/d}\right).$$

Proof. Let $P = (p_1, \ldots, p_n)$ be a c-chain in $\mathbb{R}^d$, for some constants $c \geq 1$ and $d \in \mathbb{N}$. We may assume that $|p_1p_n| = 1$. By the c-chain property, all vertices of P lie in an ellipsoid E with foci at p_1 and p_n, with major axis of length c. Let U be a ball of radius $c/2$ concentric with E; and note that $E \subseteq U$.

We set $x = c^2/(n-1)^{1/d}$; and let L_0 and L_1 be the sum of lengths of all edges in P of length at most x and more than x, respectively. By definition, we have $L = L_0 + L_1$ and

$$L_0 \leq (n-1)x = c^2(n-1)^{(d-1)/d}. \tag{10}$$

We shall prove that $L_1 = O\left(c^2(n-1)^{(d-1)/d}\right)$. For this, we further classify the edges in L_1 according to their lengths: For $\ell = 0, 1, \ldots, \infty$, let

$$P_\ell = \left\{p_i : 2^\ell x < |p_i p_{i+1}| \le 2^{\ell+1} x\right\}. \tag{11}$$

As shown in the proof of Theorem 2, we have $|p_i p_{i+1}| \le c$, for all $i = 0, \ldots, n-1$. Consequently, $P_\ell = \emptyset$ when $c \le 2^\ell x$, or equivalently $\log(c/x) \le \ell$.

We use a volume argument to derive an upper bound on the cardinality of P_ℓ, for $\ell = 0, 1, \ldots, \lfloor \log(c/x) \rfloor$. Assume that $p_i, p_k \in P_\ell$, and w.l.o.g., $i < k$. If $k = i+1$, then $2^\ell x < |p_i p_k|$ by (11). Otherwise,

$$2^\ell x < |p_i p_{i+1}| < |p_i p_{i+1}| + |p_{i+1} p_k| \le c|p_i p_k|, \text{ or } \frac{2^\ell x}{c} < |p_i p_k|.$$

Consequently, the balls of radius

$$R = \frac{2^\ell x}{2c} = \frac{2^\ell c}{2(n-1)^{1/d}} \tag{12}$$

centered at the points in P_ℓ are interior-disjoint. The volume of each ball is $\alpha_d R^d$, where $\alpha_d > 0$ depends on d only. Since $P_\ell \subset E$, these balls are contained in the R-neighborhood of the ball U, which is a ball U_R of radius $\frac{c}{2} + R$ concentric with U. For $\ell \le \log(c/x)$, we have $2^\ell x \le c$, hence $R = \frac{2^\ell x}{2c} \le \frac{c}{2c} = \frac{1}{2}$. Consequently, the radius of U_R is at most c. Since U_R contains $|P_\ell|$ interior-disjoint balls of radius R, we obtain

$$|P_\ell| \le \frac{\alpha_d c^d}{\alpha_d R^d} = \left(\frac{c}{R}\right)^d = \left(\frac{2(n-1)^{1/d}}{2^\ell}\right)^d \le \frac{2^d}{2^{d\ell}}(n-1). \tag{13}$$

For every segment $p_i p_{i+1}$ with length more than x, we have that $p_i \in P_\ell$, for some $\ell \in \{0, 1, \ldots, \lfloor \log(c/x) \rfloor\}$. Using (13), the total length of these segments is

$$L_1 \le \sum_{\ell=0}^{\lfloor \log(c/x) \rfloor} |P_\ell| \cdot 2^{\ell+1} x < \sum_{\ell=0}^{\lfloor \log(c/x) \rfloor} \frac{2^d}{2^{d\ell}}(n-1) \cdot 2^{\ell+1} \cdot \frac{c^2}{(n-1)^{1/d}}$$

$$< 2^{d+1} c^2 (n-1)^{\frac{d-1}{d}} \sum_{\ell=0}^{\infty} \frac{1}{2^{(d-1)\ell}} \le 2^{d+2} c^2 (n-1)^{(d-1)/d},$$

as required. Together with (10), this yields $L = O\left(c^2(n-1)^{(d-1)/d}\right)$. □

4.2. Lower bounds in $\mathbb{R}^d$. We show that the exponent $(d-1)/d$ in Theorem 10 cannot be improved. More precisely, for every $\varepsilon > 0$, we construct a family of axis-parallel chains in $\mathbb{R}^d$ whose stretch factor is $n^{(1-\varepsilon)(d-1)/d}$ for sufficiently large $n(\varepsilon)$. For the higher-dimensional case, we focus on axis-parallel chains, as they are easier to analyze. In the plane ($d = 2$), this construction is also possible, but it yields weaker bounds than Theorem 4.

THEOREM 11. *Let $d \ge 2$ be an integer. For all constants $\varepsilon > 0$ and sufficiently large $c = \Omega(d)$, there is a positive integer n_0 such that for every $n \ge n_0$, there exists an axis-parallel c-chain in $\mathbb{R}^d$ with n vertices and stretch factor at least $(n-1)^{(1-\varepsilon)(d-1)/d}$.*

Proof. Let $d \ge 2$, $\varepsilon > 0$, and $c = \Omega(d)$ be given. We describe a recursive construction in terms of an even integer parameter

$$r > 3^{(1-\varepsilon)/(d\varepsilon)}. \tag{14}$$

We recursively define a family $\mathcal{Q}_c = \{Q^m\}_{m\in\mathbb{N}}$ of axis-parallel c-chains in $\mathbb{R}^d$, where each chain Q^m has $n_m \le 3^{m+1}r^{dm}$ vertices. Then, we show that the stretch factor of every Q^m is at least $(n_m-1)^{(1-\varepsilon)(d-1)/d}$ for sufficiently large $m \in \mathbb{N}$.

Construction of $\mathcal{Q}_c$. For each chain in $\mathcal{Q}_c$, we maintain a subset of *active* directed edges, which are disjoint, have the same length, and are parallel to the same coordinate axis. In a nutshell, the recursion works as follows. We start with a chain Q^0 that consists of a single segment that is labeled active; then for $m = 1, 2, \ldots$, we obtain Q^m by replacing each active edge in a fixed chain π by a homothetic copy of Q^{m-1}. The chain π is defined below; it consists of $6r^d + 1$ edges, $3r^d$ of which are active.

We define the chain π in four steps, see Fig. 11 for an illustration. Let $\mathbf{e_i}$, $i = 1, \ldots, d$, be the standard basis vectors in $\mathbb{R}^d$.

(1) Consider the $(d-1)$-dimensional hyperrectangle $A = [0,1] \times [0, r-1]^{d-2}$. Let γ_0 be an axis-parallel Hamiltonian cycle on the $2r^{d-2}$ integer points that lie in A such that the origin is incident to an edge parallel to the x_1-axis. We label the vertices of γ_0 by v_i, for $i = 1, \ldots, 2r^{d-2}$, in order, where v_1 is the origin.

(2) Let $a = (3r^2+1)/(3r) = r + 1/(3r)$, and consider the d-dimensional hyperrectangle $A \times [0,a] = [0,1] \times [0, r-1]^{d-2} \times [0,a]$. We construct a Hamiltonian cycle γ_1 on the $4r^{d-2}$ points in

$$\left\{v_i \times \{0,a\} \mid i = 1, \ldots, 2r^{d-2}\right\}$$

by replacing every edge (v_{2i-1}, v_{2i}) in γ_0 with three edges

$$((v_{2i-1},0),(v_{2i-1},a)),\ ((v_{2i-1},a),(v_{2i},a)),\ \text{and}\ ((v_{2i},a),(v_{2i},0)).$$

Note that γ_1 has $4r^{d-2}$ edges, such that $2r^{d-2}$ edges have length a and are parallel to the x_d-axis. Also note that the origin v_1 is incident to a unit edge parallel to the x_1-axis, and to an edge of length a parallel to the x_d-axis.

(3) Delete the edge of γ_1 that is incident to the origin v_1 and parallel to the x_1-axis. This turns γ_1 into a Hamiltonian chain γ_2 from the origin to the vertex $\mathbf{e}_1$ in the hyperrectangle $A \times [0,a] = [0,1] \times [0,r-1]^{d-2} \times [0,a]$.

(4) Consider the hyperrectangle $B(\pi) = \left[0, 3r^2+1\right] \times [0, r-1]^{d-2} \times [0,a]$. Let π be the chain from the origin to $(3r^2+1)\cdot\mathbf{e}_1$ that is obtained by the concatenation of $3r^2/2$ copies of γ_2, translated by vectors $(2j-1)\cdot\mathbf{e}_1$ for $j = 1, 2, \ldots, 3r^2/2$, interlaced with $3r^2/2+1$ unit segments parallel to $\mathbf{e}_1$. Note that π has $\left(3r^2/2\right)\cdot\left(4r^{d-2}-1\right) + 3r^2/2 + 1 = 6r^d + 1$ edges, $\left(3r^2/2\right)\cdot 2r^{d-2} = 3r^d$ of which have length a and are parallel to the x_d-axis. We label all these edges as active, so that π has $3r^d$ active edges. Observe that $B(\pi)$ is the minimum axis-parallel bounding box of π.

Lemma 12. *The chain π is a c'-chain for $c' = 8 + 2r\sqrt{d-1}$. Furthermore, if the points q_1, q_2, and q_3 are contained in active edges, in this order along π and not all in the same edge, then*

$$\frac{|q_1q_2| + |q_2q_3|}{|q_1q_3|} \le 8 + 2r\sqrt{d-1}.$$

Proof. We extend π to a chain π' by attaching a parallel copy of γ_2 to each end of π. We prove the lemma for π'. Then, the lemma also follows for π, as π is a subchain of π'. Write $\pi' = (p_1, \ldots, p_n)$. Since p_i, p_j, and p_k are endpoints of active edges, for any choice of $1 \le i < j < k \le n$, the second claim in the lemma implies that π' is a c'-chain.

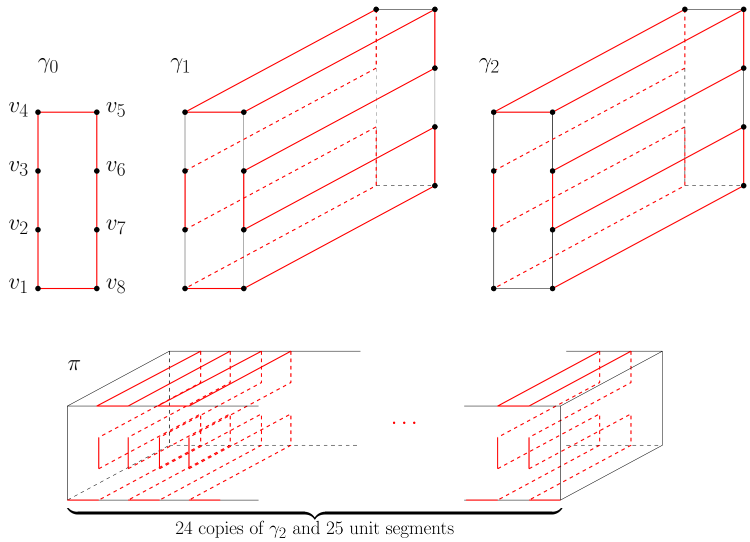

FIG. 11. *The cycles γ_0 (top left), γ_1 (top middle), and the chains γ_2 (top right), π (bottom) for $d = 3$ and $r = 4$. The cycles and chains are in red, their bounding boxes are outlined in black.*

We give an upper bound for the ratio $(|q_1q_2| + |q_2q_3|)/|q_1q_3|$. Recall that all the active edges in π' come from the $3r^2/2 + 2$ translated copies of the chain γ_2; and γ_2 has vertices in an axis-aligned bounding box $B = [0, 1] \times [0, r - 1]^{d-2} \times [0, a]$. Denote by $B_0, B_1, \ldots, B_{3r^2/2}, B_{3r^2/2+1}$ the minimum axis-aligned bounding boxes of the $3r^2/2 + 2$ translates of γ_2 in π'. Suppose that q_1, q_2, and q_3 are in B_{i_1}, B_{i_2}, and B_{i_3}, respectively. By assumption, $i_1 \le i_2 \le i_3$.

If $i_1 = i_3$, then q_1, q_2, and q_3 are in B_{i_1}. Since q_1 and q_3 are not on the same active edge, and since γ_0 has integer coordinates, we have $|q_1q_3| \ge 1$. Consequently,

$$\begin{aligned}
\frac{|q_1q_2| + |q_2q_3|}{|q_1q_3|} &\le \frac{2 \cdot \operatorname{diam}(B_{i_1})}{1} \\
&\le 2\sqrt{1^2 + (d-2)(r-1)^2 + a^2} \\
&= 2\sqrt{1 + (d-2)(r-1)^2 + (r + 1/(3r))^2} \\
&\le 2\sqrt{2 + (d-1)r^2} \\
&< 2\sqrt{2} + 2r\sqrt{d-1}.
\end{aligned}$$

Otherwise $i_1 < i_3$, and the first coordinates of q_1 and q_3 differ by at least $2(i_3 - i_1) - 1 \ge i_3 - i_1$, hence $|q_1q_3| \ge i_3 - i_1$. In this case,

$$\begin{aligned}
\frac{|q_1q_2| + |q_2q_3|}{|q_1q_3|} &\le \frac{2 \cdot \operatorname{diam}(B_{i_1} \cup B_{i_3})}{i_3 - i_1} \\
&\le \frac{2 \cdot \sqrt{(2(i_3 - i_1) + 1)^2 + (d-2)(r-1)^2 + a^2}}{i_3 - i_1} \\
&\le \frac{4(i_3 - i_1) + 4 + 2r\sqrt{d-1}}{i_3 - i_1} \\
&\le 8 + 2r\sqrt{d-1},
\end{aligned}$$

as claimed. This completes the proof of Lemma 12. ∎

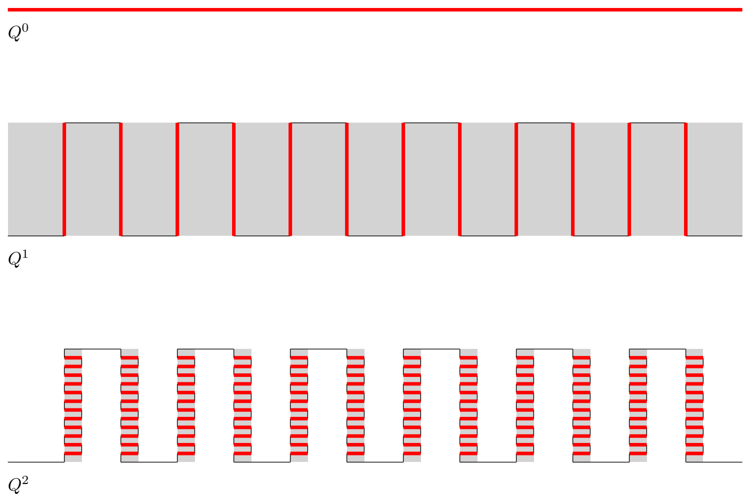

FIG. 12. *The chains Q^0 (top), Q^1 (middle), and Q_2 (bottom) for $d = r = 2$. The active edges are highlighted by red bold lines. The bounding box B of Q^1 and bounding boxes B' of homothetic copies of Q^1 in Q^2 are shaded.*

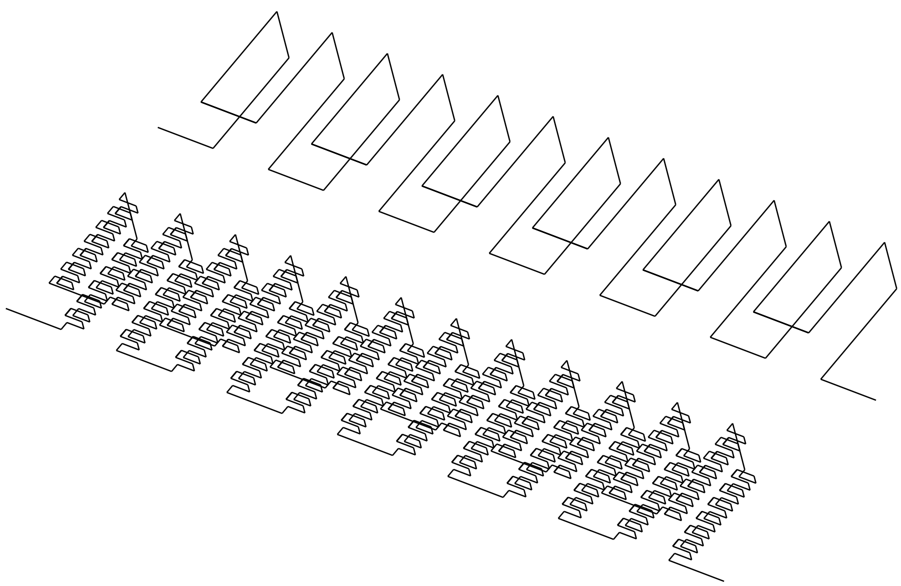

FIG. 13. *The chains Q^1 (top) and Q^2 (bottom) for $d = 3$ and $r = 2$.*

Now the axis-parallel chains Q^m can be defined recursively (see Fig. 12 for an illustration). Let Q^0 be a line segment of length $3r^2 + 1$, parallel to the x_1-axis, labeled active. Let Q^1 be π and let $B = B(\pi)$ be its minimum axis-parallel bounding box. Recall that $B = \big[0, 3r^2 + 1\big] \times [0, r - 1]^{d-2} \times [0, a]$.

We maintain the invariant that each chain Q^m ($m \in \mathbb{N}$) is contained in B. In order to do this, let B' be a hyperrectangle obtained from B by a rotation of 90 degrees in the $\langle \mathbf{e}_1, \mathbf{e}_d \rangle$ plane, and scaling by a factor of $a/(3r^2+1) = 1/(3r)$; i.e., $B' = [0, a/(3r)] \times [0, (r-1)/(3r)]^{d-2} \times [0, a]$. In particular, the longest edges of B' are parallel to the active edges in B, and they all have length a. Place a translate of B' along each active edge in Q^1 such that all such translates are contained in B. Note that the distance between any two translates is at least $1 - 2a/(3r) = 1/3 - 2/(9r^2) \geq 5/18$.

For all $m \geq 1$, we construct Q^{m+1} by replacing the active edges of Q^1 with a scaled (and rotated) copy of Q^m in each translate of B'; and we let the active edges of Q^{m+1} be the active edges in these new copies of Q^m.

Instead of keeping track of the total length of Q^m, we analyze the total length of the active edges of Q^m. In each iteration, the number of active edges increases by a factor of $3r^d$ and the length of an active edge decreases by a factor of $a/(3r^2+1) = 1/(3r)$. Overall the total length of active edges increases by a factor of r^{d-1}. It follows that for all $m \in \mathbb{N}$, the chain Q^m has $3^m r^{dm}$ active edges, and their total length is $(3r^2+1) \cdot r^{(d-1)m}$. Thus, we have

$$|Q^m| \geq (3r^2+1) \cdot r^{(d-1)m}, \tag{15}$$

for $m \in \mathbb{N}$. Next we estimate the number of vertices in Q^m. Recall that the recursive construction replaces each active edge with $3r^d$ active edges and $3r^d+1$ inactive edges (which are never replaced). Consequently, for $m \geq 1$, the number of inactive edges in Q^m is $(3r^d+1)\sum_{i=0}^{m-1} 3^i r^{di}$, and the total number of vertices is

$$n_m = 1 + 3^m r^{dm} + (3r^d+1)\sum_{i=0}^{m-1} 3^i r^{di} = 1 + 3^m r^{dm} + (3r^d+1)\frac{3^m r^{dm} - 1}{3r^d - 1}.$$

Note that

$$3^m r^{dm} < n_m \leq 3 \cdot 3^m r^{dm}. \tag{16}$$

Since the distance between the two endpoints of Q^m remains $3r^2+1$, we can use (15) and the upper bound in (16) to obtain

$$\frac{|Q^m|}{3r^2+1} \geq r^{(d-1)m} \geq \left(\frac{n_m}{3^{m+1}}\right)^{\frac{d-1}{d}}. \tag{17}$$

Now, (14) implies that $r = \beta \cdot 3^{(1-\varepsilon)/(d\varepsilon)}$, for a constant $\beta > 1$. Thus, using the lower bound in (16), we get that

$$n_m^{\varepsilon} > 3^{\varepsilon m} r^{\varepsilon dm} = 3^{\varepsilon m}\left(\beta \cdot 3^{\frac{(1-\varepsilon)}{\varepsilon d}}\right)^{\varepsilon dm} = \beta^{\varepsilon dm} \cdot 3^m \geq 3^{m+1},$$

for sufficiently large m. Hence, combining with (17), we can bound the stretch factor from below as

$$\frac{|Q^m|}{3r^2+1} \geq n_m^{(1-\varepsilon)\frac{d-1}{d}},$$

for sufficiently large m.

It remains to show that $\mathcal{Q}_c = \{Q^m : m \in \mathbb{N}\}$ is a family of c-chains, where $c = \Omega(d)$. We proceed by induction on m. The claim is trivial for $m = 0$, and it follows from Lemma 12 for $m = 1$.

Now, let $m \geq 2$. Write $Q^m = (p_1, \ldots, p_n)$, and let $1 \leq i < j < k \leq n$. We shall derive an upper bound for the ratio $(|p_i p_j| + |p_j p_k|)/|p_i p_k|$. Recall that Q^m is obtained by replacing each active edge of $Q^1 = \pi$ by a scaled copy of Q^{m-1}. If p_i and p_k are in the same copy of Q^{m-1}, then so is p_j and induction completes the proof.

Otherwise let B'_i, B'_j, and B'_k be the bounding boxes of the copies of Q^{m-1} that contain p_i, p_j, and p_k, respectively. Let a_i, a_j, and a_k be the active segments in Q^1 that are replaced by B'_i, B'_j, and B'_k; and let $q_i \in a_i$, $q_j \in a_j$, and $q_k \in a_k$ be the orthogonal projections of p_i, p_j, and p_k onto a_i, a_j, and a_k, respectively. (If $i = 1$, then let $q_i = p_1$; if $k = n$, then let $q_k = p_n$. Since the proof of Lemma 12 works on the extended chain π', it applies to q_i, q_j, and q_k regardless of this special condition.)

Since each projection happens within a hyperplane orthogonal to the x_d-axis onto an active edge in a translated copy of $[0, a/(3r)] \times [0, (r-1)/(3r)]^{d-2} \times [0, a]$, we have that $|p_iq_i|$, $|p_jq_j|$, and $|p_kq_k|$ are each bounded above by

$$\sqrt{\frac{a^2}{(3r)^2} + (d-2)\frac{(r-1)^2}{(3r)^2}} \leq \frac{\sqrt{d-1}}{3} + \frac{1}{3r} \leq \frac{\sqrt{d-1}}{3} + \frac{1}{6}.$$

As there are at least two distinct active edges among a_i, a_j, and a_k (and as the distance between p_1 or p_n and any active edge in π is at least 1), we have

$$|q_iq_j| + |q_jq_k| \geq \max\{|q_iq_j|, |q_jq_k|\} \geq 1.$$

Combining these two bounds with the triangle inequality, we get

$$\begin{aligned}
|p_ip_j| + |p_jp_k| &\leq (|p_iq_i| + |q_iq_j| + |q_jp_j|) + (|p_jq_j| + |q_jq_k| + |q_kp_k|) \\
&\leq |q_iq_j| + |q_jq_k| + \frac{4}{3}\sqrt{d-1} + \frac{2}{3} \\
&\leq \left(\frac{5}{3} + \frac{4}{3}\sqrt{d-1}\right)(|q_iq_j| + |q_jq_k|).
\end{aligned}$$

On the other hand, we have $|p_ip_k| \geq \frac{5}{18}|q_iq_k|$, as this lower bound holds for the projections of the edges to each coordinate axis. Now Lemma 12 yields

$$\begin{aligned}
\frac{|p_ip_j| + |p_jp_k|}{|p_ip_k|} &\leq \frac{5/3 + 4\sqrt{d-1}/3}{5/18} \cdot \frac{|q_iq_j| + |q_jq_k|}{|q_iq_k|} \\
&\leq (6 + 24\sqrt{d-1}/5) \cdot (8 + 2r\sqrt{d-1}) \\
&= O(r(d-1)).
\end{aligned}$$

This completes the proof of Theorem 11. ∎

5. Algorithm for Recognizing c-Chains. In this section, we design a randomized Las Vegas algorithm to recognize c-chains in d-dimensional Euclidean space. More precisely, given a polygonal chain $P = (p_1, \ldots, p_n)$ in $\mathbb{R}^d$, and a parameter $c \geq 1$, the algorithm decides whether P is a c-chain, in $O\left(n^{3-1/d} \operatorname{polylog} n\right)$ expected time. By definition, $P = (p_1, \ldots, p_n)$ is a c-chain if $|p_ip_j| + |p_jp_k| \leq c\,|p_ip_k|$ for all $1 \leq i < j < k \leq n$; equivalently, p_j lies in the ellipsoid of major axis c with foci p_i and p_k. Consequently, it suffices to test, for every pair $1 \leq i < k \leq n$, whether the ellipsoid of major axis $c|p_ip_k|$ with foci p_i and p_k contains p_j, for all j, $i < j < k$. For this, we can apply recent results from geometric range searching.

THEOREM 13. *For every integer $d \geq 2$, there are randomized algorithms that can decide, for a polygonal chain $P = (p_1, \ldots, p_n)$ in $\mathbb{R}^d$ and a threshold $c > 1$, whether P is a c-chain in $O\left(n^{3-1/d} \operatorname{polylog} n\right)$ expected time and $O(n \log n)$ space.*

Agarwal, Matoušek and Sharir [3, Theorem 1.4] constructed, for a set S of n points in $\mathbb{R}^d$, a data structure that can answer semi-algebraic range searching queries; in particular, it can report the number of points in S that are contained in a query ellipsoid. Specifically, they showed that, for every $d \geq 2$ and $\varepsilon > 0$, there is a constant B and a data structure with $O(n)$ space, $O\left(n^{1+\varepsilon}\right)$ expected preprocessing time, and $O\left(n^{1-1/d}\log^B n\right)$ query time. The construction was later simplified by Matoušek and Patáková [28]. Using this data structure, we can quickly decide whether a given polygonal chain is a c-chain.

Proof of Theorem 13. Subdivide the polygonal chain $P = (p_1, \ldots, p_n)$ into two equal-sized subchains (to within 1) $P_1 = (p_1, \ldots, p_{\lceil n/2 \rceil})$ and $P_2 = (p_{\lceil n/2 \rceil}, \ldots, p_n)$; and recursively subdivide P_1 and P_2 until reaching 1-vertex chains. Denote by T the recursion tree. Then, T is a binary tree of depth $\lceil \log n \rceil$. There are at most 2^i nodes at level i; the nodes at level i correspond to edge-disjoint subchains of P, each of which has at most $n/2^i$ edges. Let W_i be the set of subchains on level i of T; and let $W = \bigcup_{i \geq 0} W_i$. We have $|W| \leq 2n$.

For each polygonal chain $Q \in W$, construct an ellipsoid range searching data structure $\mathrm{DS}(Q)$ described above [3] for the vertices of Q, with a suitable parameter $\varepsilon > 0$. Their overall expected preprocessing time is

$$\sum_{i=0}^{\lceil \log n \rceil} 2^i \cdot O\left(\left(\frac{n}{2^i}\right)^{1+\varepsilon}\right) = O\left(n^{1+\varepsilon} \sum_{i=0}^{\lceil \log n \rceil} \left(\frac{1}{2^i}\right)^{\varepsilon}\right) = O\left(n^{1+\varepsilon}\right),$$

and their space requirement is $\sum_{i=0}^{\lceil \log n \rceil} 2^i \cdot O\left(n/2^i\right) = O(n \log n)$. The query time of each chain in W_i is $O\left(\left(n/2^i\right)^{1-1/d} \text{ polylog } \left(n/2^i\right)\right)$.

For each pair of indices $1 \leq i < k \leq n$, we do the following. Let $E_{i,k}$ denote the ellipsoid of major axis $c|p_i p_k|$ with foci p_i and p_k. The chain $(p_{i+1}, \ldots, p_{k-1})$ is subdivided into $O(\log n)$ maximal subchains in W, using at most two subchains from each set W_i, $i = 0, \ldots, \lceil \log n \rceil$. For each of these subchains $Q \in W$, query the data structure $\mathrm{DS}(Q)$ with the ellipsoid $E_{i,k}$. If all queries are positive (i.e., the count returned is $|Q|$ in *all* queries), then P is a c-chain; otherwise there exists j, $i < j < k$, such that $p_j \notin E_{i,k}$, hence $|p_i p_j| + |p_j p_k| > c|p_i p_k|$, witnessing that P is not a c-chain.

The query time over all pairs $1 \leq i < k \leq n$ is bounded above by

$$\begin{aligned} \binom{n}{2} \sum_{i=0}^{\lceil \log n \rceil} 2 \cdot O\left(\left(\frac{n}{2^i}\right)^{1-1/d} \text{ polylog } \left(\frac{n}{2^i}\right)\right) &= \binom{n}{2} \cdot O\left(n^{1-1/d} \text{ polylog } n\right) \\ &= O\left(n^{3-1/d} \text{ polylog } n\right). \end{aligned}$$

This subsumes the expected time needed for constructing the structures $\mathrm{DS}(Q)$, for all $Q \in W$. So the overall running time of the algorithm is $O\left(n^{3-1/d} \text{ polylog } n\right)$, as claimed. ∎

In the decision algorithm in the proof of Theorem 13, only the construction of the data structures $\mathrm{DS}(Q)$, $Q \in W$, uses randomization, which is independent of the value of c. The parameter c is used for defining the ellipsoid $E_{i,k}$, and the queries to the data structures; this part is deterministic. Hence, we can find the optimal value of c by Meggido's parametric search [29] in the second part of the algorithm.

Meggido's technique reduces an optimization problem to a corresponding decision problem at a polylogarithmic factor increase in the running time. An optimization problem is amenable to this technique if the following three conditions are met [35]: (1) the objective function is monotone in the given parameter; (2) the decision problem can be solved by evaluating bounded-degree polynomials, and (3) the decision problem admits an efficient parallel algorithm (with polylogarithmic running time using a polynomial number of processors). All three conditions hold in our case: The area of each ellipsoid with foci in S monotonically increases with c; the data structure of [28] answers ellipsoid range counting queries by evaluating polynomials of bounded degree; and the $\binom{n}{2}$ queries can be performed in parallel. Alternatively, Chan's randomized optimization technique [12] is also applicable. Both techniques yield the following result.

COROLLARY 14. *There are randomized algorithms that can find, for a polygonal chain* $P = (p_1, \ldots, p_n)$ *in* $\mathbb{R}^d$*, the minimum* $c \geq 1$ *for which* P *is a* c*-chain in* $O\left(n^{3-1/d} \text{ polylog } n\right)$ *expected time and* $O(n \log n)$ *space.*

We note that, for $c = 1$, the test takes $O(n)$ time: it suffices to check whether points $p_3, \ldots, p_n$ lie on the line spanned by p_1p_2, in that order.

Remark. Recently, Agarwal et al. [1, Theorem 13] designed a data structure for semi-algebraic range searching queries that supports $O(\log n)$ query time, at the expense of higher space and preprocessing time. The size and preprocessing time depend on the number of free parameters that describe the semi-algebraic set. An ellipsoid in $\mathbb{R}^d$ is defined by $2d+1$ parameters: the coordinates of its foci and the length of its major axis. Specifically, they showed that, for every $d \geq 2$ and $\varepsilon > 0$, there is a data structure with $O(n^{2d+1+\varepsilon})$ space and $O(n^{2d+1+\varepsilon})$ expected preprocessing time that can report the number of points in S contained in a query ellipsoid in $O(\log n)$ time. This data structure allows for a tradeoff between preprocessing time and overall query time in the algorithm above. However the resulting tradeoff does not seem to yield an improvement over the expected running time in Theorem 13 for any $d \geq 2$.

6. Conclusion.

We conclude with some remarks and open problems.

1. The lower bound construction in the plane can be slightly improved as follows. For $m \geq 1$, let $P_*^m = g_2(P^m) \cup g_3(P^m)$, see FIG. 14 (right). Observe that P_*^m is a c-chain with $n = 4^m/2 + 1$ vertices and stretch factor

$$\sqrt{c(c-2)/8}(n-1)^{\frac{1+\log(c-2)-\log c}{2}}.$$

Since $\sqrt{c(c-2)/8} \geq 1$ for $c \geq 4$, this improves the result of Theorem 4 by a constant factor. Since this construction does not improve the exponent, and the analysis would be longer (requiring a case analysis without new insights), we omit the details.

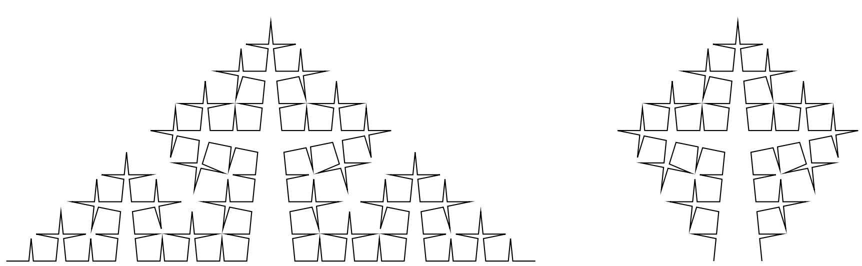

FIG. 14. *The chains* P^4 *(left) and* P_*^4 *(right).*

2. The lower bound construction in the plane depends on a parameter $c_* = (c-2)/2$. If c were used instead, the condition $c \geq 4$ in Theorem 4 could be replaced by $c \geq 1$, and the bound could be improved from

$$(n-1)^{\frac{1+\log(c-2)-\log c}{2}} \quad \text{to} \quad (n-1)^{\frac{1+\log c-\log(c+1)}{2}}.$$

   Although we were unable to prove that the resulting P^m's, $m \in \mathbb{N}$, are c-chains, a computer program has verified that the first few generations of them are indeed c-chains.
3. The upper bounds in Theorem 1–3 (and their generalizations to higher dimensions, e.g., Theorem 10) are valid regardless of whether the chain is crossing or not. On the other hand, the lower bounds in Theorem 4 and Theorem 11 are given by noncrossing chains. A natural question is whether sharper upper bounds hold if the chains are required to be noncrossing. Specifically, can the exponent of n in the upper bound for $\mathbb{R}^d$ be reduced to $\frac{d-1}{d} - \varepsilon$, where $\varepsilon > 0$ depends on c?
4. The running time of the algorithm in Theorem 13 is sub-cubic, but super-quadratic. Is this necessary, or is it possible to decide the c-chain property in time $O(n^2)$ or better?

Acknowledgments. This work was initiated at the Fields Workshop on Discrete and Computational Geometry, held July 31–August 4, 2017, at Carleton University. The authors thank the organizers and all participants of the workshop for inspiring discussions and for providing a great research atmosphere. This problem was initially posed by Rolf Klein in 2005. We would like to thank Rolf Klein and Christian Knauer for interesting discussions on the stretch factor and related topics.

REFERENCES

[1] P. K. Agarwal, B. Aronov, E. Ezra, and J. Zahl, *An efficient algorithm for generalized polynomial partitioning and its applications*, in 35th International Symposium on Computational Geometry (SoCG), vol. 129 of LIPIcs, Schloss Dagstuhl, 2019, pp. 5:1–5:14, https://doi.org/10.4230/LIPIcs.SoCG.2019.5.

[2] P. K. Agarwal, R. Klein, C. Knauer, S. Langerman, P. Morin, M. Sharir, and M. A. Soss, *Computing the detour and spanning ratio of paths, trees, and cycles in 2D and 3D*, Discrete & Computational Geometry, 39 (2008), pp. 17–37, https://doi.org/10.1007/s00454-007-9019-9.

[3] P. K. Agarwal, J. Matoušek, and M. Sharir, *On range searching with semialgebraic sets. II*, SIAM J. Computing, 42 (2013), pp. 2039–2062, https://doi.org/10.1137/120890855.

[4] O. Aichholzer, F. Aurenhammer, C. Icking, R. Klein, E. Langetepe, and G. Rote, *Generalized self-approaching curves*, Discrete Applied Mathematics, 109 (2001), pp. 3–24, https://doi.org/10.1016/S0166-218X(00)00233-X.

[5] S. Alamdari, T. M. Chan, E. Grant, A. Lubiw, and V. Pathak, *Self-approaching graphs*, in Proc. 20th Symposium on Graph Drawing (GD), W. Didimo and M. Patrignani, eds., vol. 7704 of LNCS, Berlin, 2012, Springer, pp. 260–271, https://doi.org/10.1007/978-3-642-36763-2_23.

[6] S. Arora, L. Lovász, I. Newman, Y. Rabani, Y. Rabinovich, and S. Vempala, *Local versus global properties of metric spaces*, SIAM J. Computing, 41 (2012), pp. 250–271, https://doi.org/10.1137/090780304.

[7] P. Bose, J. Gudmundsson, and M. H. M. Smid, *Constructing plane spanners of bounded degree and low weight*, Algorithmica, 42 (2005), pp. 249–264, https://doi.org/10.1007/s00453-005-1168-8.

[8] P. Bose, I. Kostitsyna, and S. Langerman, *Self-approaching paths in simple polygons*, Computational Geometry: Theory and Applications, 87 (2020), p. 101595, https://doi.org/10.1016/j.comgeo.2019.101595.

[9] P. Bose and M. H. M. Smid, *On plane geometric spanners: A survey and open problems*, Computational Geometry: Theory and Applications, 46 (2013), pp. 818–830, https://doi.org/10.1016/j.comgeo.2013.04.002, https://doi.org/10.1016/j.comgeo.2013.04.002.

[10] P. Brass, W. O. J. Moser, and J. Pach, *Research Problems in Discrete Geometry*, Springer, New York, 2005, https://doi.org/10.1007/0-387-29929-7.

[11] E. Cesàro, *Remarques sur la courbe de von Koch*, Atti della R. Accad. della Scienze fisiche e matem. Napoli, 12 (1905). Reprinted as §228 in *Opere scelte, a cura dell'Unione matematica italiana e col contributo del Consiglio nazionale delle ricerche*, Vol. 2: Geometria, analisi, fisica matematica, Rome, dizioni Cremonese, pp. 464–479, 1964.

[12] T. M. Chan, *Geometric applications of a randomized optimization technique*, Discrete & Computational Geometry, 22 (1999), pp. 547–567, https://doi.org/10.1007/PL00009478.

[13] O. Cheong, H. J. Haverkort, and M. Lee, *Computing a minimum-dilation spanning tree is NP-hard*, Computational Geometry: Theory and Applications, 41 (2008), pp. 188–205, https://doi.org/10.1016/j.comgeo.2007.12.001.

[14] F. R. K. Chung and R. L. Graham, *On Steiner trees for bounded point sets*, Geometriae Dedicata, 11 (1981), pp. 353–361, https://doi.org/10.1007/BF00149359.

[15] H. T. Croft, K. J. Falconer, and R. K. Guy, *Unsolved Problems in Geometry*, vol. 2 of Unsolved Problems in Intuitive Mathematics, Springer, New York, 1991, https://doi.org/10.1007/978-1-4612-0963-8.

[16] G. Das and D. Joseph, *Which triangulations approximate the complete graph?*, in Proc. International Symposium on Optimal Algorithms, H. Djidjev, ed., vol. 401 of LNCS, Berlin, 1989, Springer, pp. 168–192, https://doi.org/10.1007/3-540-51859-2_15.

[17] A. Dumitrescu and M. Jiang, *Minimum rectilinear Steiner tree of n points in the unit square*, Computational Geometry: Theory and Applications, 68 (2018), pp. 253–261, https://doi.org/10.1016/j.comgeo.2017.06.007.

[18] D. Eppstein, *Spanning trees and spanners*, in Handbook of Computational Geometry, J.-R. Sack and J. Urrutia, eds., Elsevier, Amsterdam, 2000, ch. 9, pp. 425–461, https://doi.org/10.1016/B978-044482537-7/50010-3.

[19] D. Eppstein, *Beta-skeletons have unbounded dilation*, Computational Geometry: Theory and Applications, 23 (2002), pp. 43–52, https://doi.org/10.1016/S0925-7721(01)00055-4.

[20] L. Fejes Tóth, *Über einen geometrischen Satz*, Mathematische Zeitschrift, 46 (1940), pp. 83–85, https://doi.org/10.1007/BF01181430.

[21] L. Few, *The shortest path and the shortest road through n points*, Mathematika, 2 (1955), pp. 141–144, https://doi.org/10.1112/S0025579300000784.

[22] E. N. Gilbert and H. O. Pollak, *Steiner minimal trees*, SIAM Journal on Applied Mathematics, 16 (1968), pp. 1–29, https://doi.org/10.1137/0116001.

[23] C. Icking, R. Klein, and E. Langetepe, *Self-approaching curves*, Mathematical Proceedings of the Cambridge Philosophical Society, 125 (1999), pp. 441–453, https://doi.org/10.1017/S0305004198003016.

[24] H. J. Karloff, *How long can a Euclidean traveling salesman tour be?*, SIAM Journal on Discrete Mathematics, 2 (1989), pp. 91–99, https://doi.org/10.1137/0402010.

[25] R. Klein, C. Knauer, G. Narasimhan, and M. H. M. Smid, *On the dilation spectrum of paths, cycles, and trees*, Computational Geometry: Theory and Applications, 42 (2009), pp. 923–933, https://doi.org/10.1016/j.comgeo.2009.03.004.

[26] D. G. Larman and P. McMullen, *Arcs with increasing chords*, Mathematical Proceedings of the Cambridge Philosophical Society, 72 (1972), pp. 205–207, https://doi.org/10.1017/S0305004100047022.

[27] J. Matoušek, *Lectures on Discrete Geometry*, vol. 212 of Graduate Texts in Mathematics, Springer-Verlag, New York, 2002, https://doi.org/10.1007/978-1-4613-0039-7.

[28] J. Matoušek and Z. Patáková, *Multilevel polynomial partitions and simplified range searching*, Discrete & Computational Geometry, 54 (2015), pp. 22–41, https://doi.org/10.1007/s00454-015-9701-2.

[29] N. Megiddo, *Linear-time algorithms for linear programming in $\mathbb{R}^3$ and related problems*, SIAM J. Computing, 12 (1983), pp. 759–776, https://doi.org/10.1137/0212052.

[30] J. S. B. Mitchell and W. Mulzer, *Proximity algorithms*, in Handbook of Discrete and Computational Geometry, J. E. Goodman, J. O'Rourke, and C. D. Tóth, eds., CRC Press, Boca Raton, 3rd ed., 2017, ch. 32, pp. 849–874, https://doi.org/10.1201/9781315119601.

[31] G. Narasimhan and M. H. M. Smid, *Approximating the stretch factor of Euclidean graphs*, SIAM J. Comput., 30 (2000), pp. 978–989, https://doi.org/10.1137/S0097539799361671.

[32] G. Narasimhan and M. H. M. Smid, *Geometric Spanner Networks*, Cambridge University Press, 2007, https://doi.org/10.1017/CBO9780511546884.

[33] M. Nöllenburg, R. Prutkin, and I. Rutter, *On self-approaching and increasing-chord draw-*

ings of 3-connected planar graphs, Journal of Computational Geometry, 7 (2016), pp. 47–69, http://jocg.org/index.php/jocg/article/view/223.

[34] G. Rote, *Curves with increasing chords*, Mathematical Proceedings of the Cambridge Philosophical Society, 115 (1994), pp. 1–12, https://doi.org/10.1017/S0305004100071875.

[35] J. S. Salowe, *Parametric search*, in Handbook of Discrete and Computational Geometry, J. E. Goodman and J. O'Rourke, eds., CRC Press, Boca Raton, 2nd ed., 2004, ch. 43, pp. 969–982, https://doi.org/10.1201/9781420035315.

[36] S. Verblunsky, *On the shortest path through a number of points*, Proceedings of the American Mathematical Society, 2 (1951), pp. 904–913, https://doi.org/10.1090/S0002-9939-1951-0045403-1.