THE TRAVELING SALESMAN THEOREM FOR JORDAN CURVES

CHRISTOPHER J. BISHOP

Abstract. If $\Gamma \subset \mathbb{R}^n$, $n \geq 2$, is a Jordan arc with endpoints z and w, we show that the arclength of Γ satisfies

$$\ell(\Gamma) - |z-w| \simeq \sum_Q \beta_\Gamma^2(Q)\mathrm{diam}(Q),$$

where the sum is over all dyadic cubes in $\mathbb{R}^n$ and $\beta_\Gamma(Q)$ is Peter Jones's β-number that measures the deviation of Γ from a straight line inside $3Q$. This estimate sharpens previously known results by replacing an $O(\mathrm{diam}(\Gamma))$ term by $|z-w|$. Applications of this improvement to the study of Weil-Petersson curves are described, and a new proof of Jones's traveling salesman theorem is given.

Date: April 15, 2020; revised August 17, 2020.

1991 *Mathematics Subject Classification.* Primary: 28A75 Secondary: 26B15, 30L05, 42C99, 49Q15 .

Key words and phrases. arclength, traveling salesman theorem, beta-numbers, rectifiable sets, Weil-Petersson class, integral geometry, convex sets .

The author is partially supported by NSF Grant DMS 1906259.

1. Introduction

Given a finite set $E \subset \mathbb{R}^n$, the traveling salesman problem asks for the shortest curve γ that contains E. This is one of the most famous intractable problems of combinatorial optimization and its study has had a profound impact on the development of computational geometry and discrete geometry. The "analyst's traveling salesman problem" asks whether an infinite set $E \subset \mathbb{R}^n$ is contained in any curve of finite length, i.e., it asks to characterize subsets of rectifiable curves. This problem has had a powerful influence on the development of harmonic analysis and geometric measure theory over the last three decades. For sets in $\mathbb{R}^2$ it was solved by Peter Jones in [16]; this is known as Jones's "traveling salesman theorem" (TST): he gave an infinite series whose sum estimates the length of the optimal curve containing E up to a bounded factor; thus E lies on a rectifiable curve if and only if the series converges. Jones's TST was extended to higher (finite) dimensions by Kate Okikiolu [27], but with constants that grow exponentially with the dimension, and later Raanan Schul [31] proved a version that holds for sets in Hilbert space, and thus in $\mathbb{R}^n$ with constants that are independent of n. This is one of only a handful of problems in Euclidean analysis where dimension independent bounds are known. Extensions to curves in other metric spaces are given in [10], [13], [21], [22]. There has also been much work in extending Jones' result from curves to higher dimensional objects in $\mathbb{R}^n$, e.g., what is the "smallest" surface containing a given set. This problem has proved extremely subtly, and is central to recent developments in harmonic analysis, geometric measure theory and rectifiability. For a sampling of applications of Jones TST and related work, see [3], [4], [6], [9], [20], [22], [28], [34].

The purpose of this paper is to return to the original setting of curves in $\mathbb{R}^n$, and prove a sharper version of Jones's and Okikiolu's theorems. In order to state their results precisely and explain the proposed improvement, we need a few definitions.

A dyadic interval I in $\mathbb{R}$ is one of the form $(2^{-n}j, 2^{-n}(j+1)]$ for $j, n \in \mathbb{Z}$. A dyadic cube Q in $\mathbb{R}^n$ is the product of n dyadic intervals of the same length. This common length is called the side length of Q and is denoted $\ell(Q)$. Note that $\text{diam}(Q) = \sqrt{n}\ell(Q)$. For a positive number $\lambda > 0$, we let λQ denote the cube concentric with Q but with diameter $\lambda \text{diam}(Q)$, e.g., $3Q$ is the "triple" of Q, a union of Q and $3^n - 1$

adjacent copies of itself. Given a set $E \subset \mathbb{R}^n$, $\lambda > 0$ and a dyadic cube Q, define

$$\beta(E, \lambda, Q) = \frac{1}{\mathrm{diam}(Q)} \inf_L \sup\{\mathrm{dist}(z, L) : z \in \lambda Q \cap E\},$$

where the infimum is over all lines L that hit Q. In most cases we take $\lambda = 3$ and for brevity we set $\beta_E(Q) = \beta(E, 3, Q)$. Note that $0 \le \beta_E(Q) \le 2$, and equals 0 if and only if E is a subset of a line. See Figure 1. There are several other versions of the β-numbers that can be used to state equivalent versions of Jones's TST; see Appendix B for a few of these.

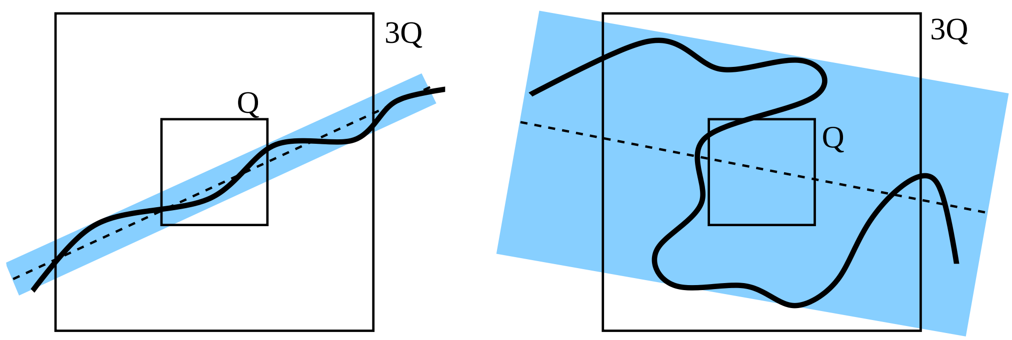

FIGURE 1. The definition of the β-numbers. The left shows a situation where $\beta_E(Q)$ is small and the right a situation where it is large. The shaded region represents the thinest strip containing $E \cap 3Q$ and the dashed line is its axis, the line L in the definition of $\beta_E(Q)$.

Jones's theorem in [16] says that the shortest curve Γ containing $E \subset \mathbb{R}^2$ has length

$$\ell(\Gamma) \simeq \mathrm{diam}(E) + \sum_Q \beta_E^2(Q)\mathrm{diam}(Q). \tag{1.1}$$

In this paper $A \lesssim B$ means the same as $A = O(B)$, i.e., A, B both depend on some parameter and $A \le C \cdot B$ where C is independent of the parameter. If $A \lesssim B$ and $B \lesssim A$, we write $A \simeq B$, and say that A and B are comparable. Thus Jones's β-sum estimates the length of the optimal curve up to a bounded factor.

Actually, [16] states that for any $\delta > 0$ and $E \subset \mathbb{R}^n$, $2 \le n < \infty$,

$$\ell(\Gamma) \le (1+\delta)\mathrm{diam}(E) + C(\delta) \sum_Q \beta_E^2(Q)\mathrm{diam}(Q). \tag{1.2}$$

For general sets E, this does not hold for $\delta = 0$. For example, if $E = \{0, 1, i\beta\} \subset \mathbb{R}^2$ with $0 < \beta << 1$, then the shortest curve Γ containing E satisfies $\ell(\Gamma) = 1 + \beta$, but $\mathrm{diam}(\Gamma) = \sqrt{1+\beta^2} = 1 + O(\beta^2)$. It is not hard to check that the β^2-sum for E is

$O(\beta^2) \ll \beta$. Thus the term $C(\delta)$ must tend to ∞ as $\delta \searrow 0$. However, we will show that (1.2) does hold for $\delta = 0$ when $E = \Gamma$ a Jordan curve:

Theorem 1.1. *For any Jordan arc in* $\mathbb{R}^n$*,*

$$\ell(\Gamma) - \operatorname{diam}(\Gamma) \simeq \sum_Q \beta_\Gamma^2(Q)\operatorname{diam}(Q), \tag{1.3}$$

where the sum is over all dyadic cubes.

In fact, we can do even better than this. Let $\operatorname{crd}(\Gamma) = |z-w|$ where $\{z, w\}$ are the endpoints of Γ; this is the "chord length" of Γ. We always have $\operatorname{crd}(\Gamma) \leq \operatorname{diam}(\gamma)$, so Theorem 1.1 implies

$$\ell(\Gamma) - \operatorname{crd}(\Gamma) \gtrsim \sum_Q \beta_\Gamma^2(Q)\operatorname{diam}(Q). \tag{1.4}$$

The opposite direction is less obvious, but also holds:

Theorem 1.2. *For any Jordan arc* $\Gamma \subset \mathbb{R}^n$*,*

$$\ell(\Gamma) - \operatorname{crd}(\Gamma) \simeq \sum_Q \beta_\Gamma^2(Q)\operatorname{diam}(Q), \tag{1.5}$$

where the sum is over all dyadic cubes.

One obvious consequence is that for any Jordan arc Γ,

$$\operatorname{diam}(\Gamma) - \operatorname{crd}(\Gamma) \lesssim \sum_Q \beta_\Gamma^2(Q)\operatorname{diam}(Q),$$

and another is:

Corollary 1.3. *If* Γ *is a closed Jordan curve, then*

$$\ell(\Gamma) \simeq \sum_Q \beta_\Gamma^2(Q)\operatorname{diam}(Q), \tag{1.6}$$

where the sum is over all dyadic cubes.

In Sections 2 and 3 we prove Theorem 1.1, and in Section 4 we use it prove Theorem 1.2. In Appendix A we show how to adapt our proof of Theorem 1.1 to give a new proof of (1.2) for general sets $E \subset \mathbb{R}^n$, and in Appendix B we discusses equivalent formulations of the β-numbers and Jones's theorem.

We end the introduction by describing one motivation for wanting this improved version of Jones's theorem. Although changing $O(\operatorname{diam}(\Gamma))$ to $\operatorname{crd}(\Gamma)$ may seem minor, it is crucial for proving the following result in [5]. Recall that a closed curve Γ is called chord-arc if any two points $z, w \in \Gamma$ are joined by a sub-arc γ with length $\ell(\gamma) = O(|z-w|)$.

Theorem 1.4. *The following are equivalent for a closed Jordan curve in* $\mathbb{R}^n$, $n \geq 2$*:*

(1) Γ *satisfies*

$$\sum_Q \beta_\Gamma^2(Q) < \infty.$$

(2) Γ *is chord-arc and for any dyadic decomposition of* Γ*, the inscribed polygons* $\{\Gamma_n\}$ *defined by the* n*th generation points satisfy*

$$\sum_{n=1}^{\infty} 2^n \left[\ell(\Gamma) - \ell(\Gamma_n)\right] < \infty,$$

with a bound that is independent of the choice of the decomposition.

(3) Γ *has finite Möbius energy, i.e.,*

$$\text{Möb}(\Gamma) = \int_\Gamma \int_\Gamma \left(\frac{1}{|x-y|^2} - \frac{1}{\ell(x,y)^2} \right) dx dy < \infty,$$

where dx, dy *denotes integration with respect to arclength measure.*

Note that (1) is Jones's sum without the $\text{diam}(Q)$ factor; thus this condition represents something stronger than rectifiability. The β's represent a measurement of local curvature of Γ, so (1) makes precise the idea that the that curvature of Γ is square integrable over all locations and scales.

If a closed Jordan curve Γ has finite length $\ell(\Gamma)$, choose a base point $z_1^0 \in \Gamma$ and for each $n \geq 1$, let $\{z_j^n\}$, $j = 1, \dots, 2^n$ be the unique set of ordered points with $z_1^n = z_1^0$ that divides Γ into 2^n equal length intervals (called the nth generation dyadic subintervals of Γ). Let Γ_n be the inscribed 2^n-gon with these vertices. Clearly $\ell(\Gamma_n) \nearrow \ell(\Gamma)$ and condition (2) measures the rate of convergence.

In (3), the Möbius energy of a curve is one of several "knot energies" on curves introduced by O'Hara [24], [25], [26], that blows up when the curve is close to self-intersecting, so continuously deforming a curve in $\mathbb{R}^3$ to minimize it should lead to a canonical "nice" representative of each knot type. This was proven by Freedman, He and Wang [15] for irreducible knots. They also showed that $\text{Möb}(\Gamma)$ is Möbius invariant (hence the name) and that $\text{Möb}(\Gamma)$ attains its minimal value 4 only for circles. Theorem 1.4 provides a geometric characterization of the curves for which this energy is finite. Function theoretic arguments suffice to prove Theorem 1.4 in the plane, but for $n \geq 3$, Theorem 1.2 is used in [5] to prove (1) $\Rightarrow$ (2); the other implications (2) $\Rightarrow$ (3) $\Rightarrow$ (1) follow by more elementary arguments.

In the special case $n = 2$, the class of closed curves described by Theorem 1.4 is known as the Weil-Petersson class. This is the closure of the smooth curves in the Weil-Petersson metric on universal Teichmüller space defined by Takhtajan and Teo [33]; their work was motivated by problems arising in string theory. Before [5] it had been an open problem to give a geometrical characterization of these curves, a question that also arose in the work of David Mumford and his students on computer vision and pattern recognition, e.g., [11], [12], and [32]. The Weil-Petersson class is also connected to the study of Schramm-Loewner evolutions (random Jordan paths) and the Brownian "loop soup" of Lawler and Werner. See [29], [35], [36], [37]. In addition to the conditions in Theorem 1.4, there are numerous other characterizations of the Weil-Petersson class involving conformal maps, Schwarzian derivatives, quasiconformal mappings, Sobolev spaces and minimal surfaces in hyperbolic 3-space with asymptotic boundary Γ. The results of this paper allow many of these characterizations to be extended to higher dimensions and proven equivalent there. They should also prove useful in a number of other constructions involving β-numbers in higher dimensions.

I thank Jack Burkart, María González, Joe Mitchell, David Mumford, and Raanan Schul for reading early drafts of this paper and for numerous helpful comments and suggestions.

2. Proof of the upper bound in Theorem 1.1

In this section we prove the inequality

$$\ell(\Gamma) - \mathrm{diam}(\Gamma) \lesssim \sum_{Q} \beta_\Gamma^2(Q)\mathrm{diam}(Q), \tag{2.1}$$

and we will prove the opposite direction in the next section.

Proof of (2.1). This direction closely follows the proof of Theorem 10.5.1 in [7] in the planar case, which itself is inspired by the argument in Section X.2 of [14] (but that proof contains a minor gap, fixed in [7]). However, several facts that are easy in the plane require more intricate proofs in higher dimensions.

We will define a sequence of nested, compact sets $\{\Gamma_n\}_0^\infty$ that shrinks down to Γ. Γ_0 is the convex hull of Γ. In general, suppose that Γ_n is the union of a collection $\mathcal{R}_n$ of compact, convex sets that cover Γ and that each set $R \in \mathcal{R}_n$ is the convex hull of $R \cap \Gamma$. For each such set R, choose a diameter segment I of R and divide I into two equal halves. Let R_1, R_2 be the convex hulls of the parts of $R \cap \Gamma$ than project orthogonally onto each of these segments. See Figure 2. We call this process splitting R. The collection $\mathcal{R}_{n+1}$ is obtained by splitting every element of $\mathcal{R}_n$ in this way. Thus $\mathcal{R}_{n+1}$ has twice as many elements as $\mathcal{R}_n$ and we will think of these elements as the nth generation of a binary tree whose root is $R_0 = \Gamma_0$. Below we will show that the diameters of these sets tend to zero uniformly in n and that the sets are well dispersed in space (only a bounded number with diameter $\simeq r$ can be within distance r of each other).

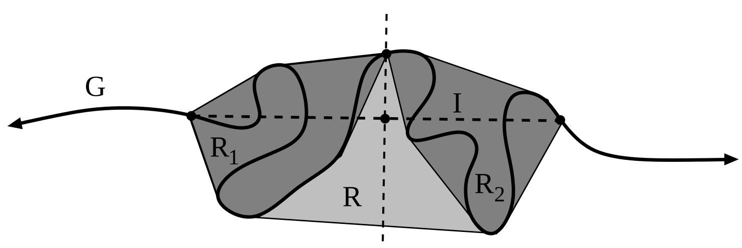

FIGURE 2. The convex set R is split into two smaller convex sets.

For a convex set R we define

$$\beta(R) = \inf_I \sup_{z \in R} \frac{\operatorname{dist}(z, I)}{\operatorname{diam}(R)},$$

and the infimum over all diameters I of R. (diameters are segments connecting pairs of points $z, w \in \partial R$ with $|z - w| = \operatorname{diam}(R)$).

Lemma 2.1. *If R is split into R_1, R_2 as above, then*

$$\operatorname{diam}(R_1) + \operatorname{diam}(R_2) \le \operatorname{diam}(R) + O\left(\beta^2(R)\operatorname{diam}(R)\right).$$

Proof. The subset $R_1 \subset R$ is contained in a cylinder W with axis length $\operatorname{diam}(R)/2$ and radius $\beta(R)\operatorname{diam}(R)$, so $\operatorname{diam}(R_1) \le \operatorname{diam}(W) = \frac{1}{2}\operatorname{diam}(R) + O(\beta^2(R)\operatorname{diam}(R))$. Similarly for R_2, and adding the estimates proves the lemma. $\square$

Lemma 2.2. *There is a constant $M = M(n)$, so that if the splitting operation is performed M times, then each of the 2^M resulting sets has diameter at most $\frac{3}{4}\mathrm{diam}(R)$.*

Proof. Suppose not, that is, suppose there is an R with $\mathrm{diam}(R) = 1$ and a large integer M, so that after M splittings some subset still has diameter $> 3/4$. After the first subdivision the projection onto the direction of the first diameter segment has length $1/2$, so the second diameter segment (or any of the next M diameter segments) can't point in the same direction. Indeed, since all the next M diameters are $> 3/4$ then can't lie within angle $\theta = \cos^{-1}(2/3)$ of the first direction. Similarly, the third direction can't be within θ of either the first or second directions, and so on. Since the $(n-1)$-sphere is compact, it contains at most a bounded number $C(n)$ of disjoint spherical caps of this size and so $M \le C(n) + 1$, as desired. □

By considering an n-dimensional ball, we see that n splittings may have to occur before the diameter drops at all. As a side remark, Borsuk's conjecture [8] asked if any bounded set in $\mathbb{R}^n$ could be partitioned into $n+1$ subsets of strictly smaller diameter, but this was disproven by Kahn and Kalai [17] who gave examples of sets requiring $\ge (1.1)^{\sqrt{n}}$ subsets when n is large. Schramm had earlier shown that $(1.3)^n$ subsets always suffice. See also Chapter 18 of [1] for some history and related results.

Using Lemma 2.1, induction and $\mathrm{diam}(\Gamma_0) = \mathrm{diam}(\Gamma)$, we get

$$\begin{aligned}\sum_{R\in\mathcal{R}_{n+1}} \mathrm{diam}(R) &\le \sum_{R\in\mathcal{R}_n} \mathrm{diam}(R) + O\left(\sum_{R\in\mathcal{R}_n} \beta^2(R)\mathrm{diam}(R)\right)\\ &\le \mathrm{diam}(\Gamma) + O\left(\sum_{k=1}^{n}\sum_{R\in\mathcal{R}_k} \beta^2(R)\mathrm{diam}(R)\right)\end{aligned}$$

The following is a standard fact.

Lemma 2.3. *For a Jordan curve, the definition of $\ell(\Gamma)$ via the supremum of lengths of inscribed polygons agrees with the definition of 1-dimensional Hausdorff measure $\mathcal{H}^1(\Gamma)$ as the limit $\lim_{\delta\searrow 0}\inf\sum_j \mathrm{diam}(X_J)$, where the infimum is over all coverings of Γ by set of diameter less than δ.*

Proof. For any arc σ, we have $\mathcal{H}^1(\sigma) \ge \mathrm{crd}(\sigma)$ so for any polygon P inscribed in γ, we have $\mathcal{H}^1(\gamma) \le \ell(P)$. By taking limits we get $\mathcal{H}^1(\gamma) \le \ell(\gamma)$. On the other hand, we can cover γ by finitely many disjoint (except for endpoints) subarcs $\{\gamma_k\}$ each of length $\leq \delta$. Hence $\sum_k \operatorname{diam}(\gamma_k) \leq \sum_l \ell(\gamma_k) \leq \ell(\gamma)$. Taking the limit as $\delta \searrow 0$ we get $\mathcal{H}^{-1}(\gamma) \leq \ell(\gamma)$. □

By Lemma 2.2 our collections $\mathcal{R}_n$ form such coverings and hence

$$\ell(\Gamma) = \mathcal{H}^1(\Gamma) \leq \limsup_{n\to\infty} \sum_{R\in\mathcal{R}_n} \operatorname{diam}(R). \tag{2.2}$$

Therefore,

$$\ell(\Gamma) \quad \leq \quad \operatorname{diam}(\Gamma) + O\left(\sum_{n=0}^{\infty} \sum_{R\in\mathcal{R}_n} \beta^2(R)\operatorname{diam}(R)\right). \tag{2.3}$$

So all that remains to do is to show that the β^2-sum over all the convex sets in the tree T is dominated by the usual β^2-sum over dyadic cubes. Given a set R in some $\mathcal{R}_n$ there is a dyadic cube Q that intersects R and satisfies $\operatorname{diam}(Q) \leq \operatorname{diam}(R) \leq 2 \cdot \operatorname{diam}(Q)$. Then $R \subset C(n) \cdot Q$ and $\beta(R) = O(\beta(Q))$. We will be done once we know that only a uniformly bounded number of R's can be associated to the same Q. This is implied by:

Lemma 2.4. *Suppose R is the convex hull of $\Gamma \subset \mathbb{R}^n$, $n \geq 2$. Consider the binary tree of subsets obtained by the subdivision rule described above. Given $0 < \epsilon < \operatorname{diam}(\Gamma)$ and a point $x \in \mathbb{R}^n$, the number of descendants of R that hit the ball $B(x, \epsilon)$ and have diameter between $\epsilon/2$ and ϵ is bounded depending only on the dimension n.*

Proof. Rescale so $\epsilon = 1/1000$ and $x = 0$. Let $\mathcal{C}$ be the collection of sets described in the lemma. Choose a large integer N and remove all the sets that are within tree distance N of the root; there are at most 2^N of these, so it suffices to bound the number of remaining sets. Replace each remaining set by its smallest (in terms of containment) ancestor to have diameter larger than 4. By Lemma 2.2 there must be such an ancestor, if N is large enough (depending only on n), and at most 2^N sets in $\mathcal{C}$ have the same replacement. Thus it suffices to bound the number of minimal sets R' in T so that $\operatorname{diam}(R') \geq 4$ and $R' \cap B_\epsilon \neq \emptyset$. We call these sets R' the admissible descendents of R and denote them by $\mathcal{A}$.

We say a set R' in $\mathcal{A}$ has rank k if it contains a k-dimensional ball of radius 10^{-k} centered on the unit n-sphere. We will call the center of this ball the center of R'. Since every admissible descendent hits B_ϵ and has diameter ≥ 4, it contains a segment that connects $\{|x| = 1/2\}$ to $\{|x| = 3/2\}$ and hence has rank at least 1. The maximum possible rank is n, and there are only a bounded number of such sets in $\mathcal{A}$ since they contain disjoint balls of fixed volume centered on the unit sphere. When considering the tree T, we will say a vertex has rank k if the corresponding set does.

The key observation is the following. Suppose that $\delta = 10^{-n-4}$ and that R_1, R_2 are two descendent sets whose center points are within δ of each other. Suppose also that R_2 has rank less than or equal to k, the rank of R_1. Let R_0 be the smallest common ancestor of R_1 and R_2 (on the tree T, this is the vertex where the paths from R_1 and R_2 to root first meet). Then R_1 and R_2 are on opposite sides of the hyperplane H (possibly each intersecting H) bisecting some diameter I of R_0, and hence H must come within δ of the center of R_1. Thus the k-ball B_k in R_1 is very close to parallel to H. The segment I is perpendicular to H and hits H at a point at most distance $\mathrm{diam}(R_0)$ from the center of B_k. By definition, R_0 contains the convex hull of its endpoints and the k-ball B_k. Since $\mathrm{diam}(I) = \mathrm{diam}(R_0) \geq \mathrm{diam}(R_1)$, R_0 contains a $(k+1)$-ball B_{k+1} with the same center as B_k and with radius at least $1/10$ as big, in particular, bigger than 10^{-k-1}. Thus any common ancestor of two sets whose center points are δ-close has strictly higher rank than either of them. See Figure 3.

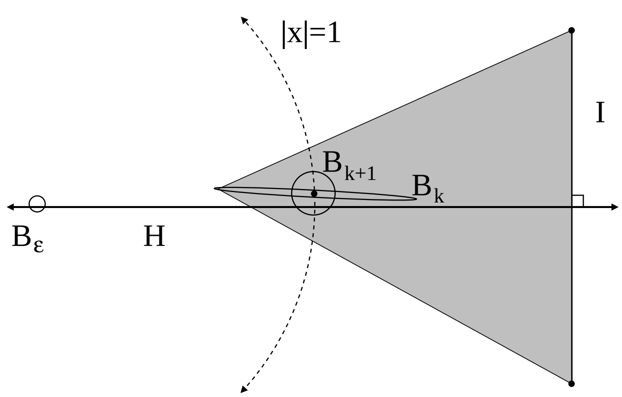

FIGURE 3. If the hyperplane H passes very close to the center of a k-dimensional ball B_k that lies on one side of H, then B_k is nearly parallel to H and a concentric ball of comparable radius and dimension $k+1$ lies in the convex hull of B_k and any segment I that is bisected and perpendicular to H, and whose length is comparable to its distance from B_k. This implies the common ancestor of two disjoint sets with very close centers must have strictly larger rank than either of them.

Now choose a point x on the unit sphere in $\mathbb{R}^n$ and consider all the admissible descendents whose centers are within δ of this point. These sets form the leaves of a finite subtree of T, where the only vertices of degree 3 are smallest common ancestors of some subcollection of the sets. By our remarks above, each vertex of degree three has strictly larger rank than any of the degree three vertices below it (closer to the leaves). Thus each leaf is connected to the root by a path that has at most n degree three vertices on it and so the tree is homeomorphic to a rooted binary tree with depth $\leq n$. Thus there are at most 2^n leaves.

Since the unit sphere in $\mathbb{R}^n$ is compact, we can partition the set of all admissible descendents into $N(n, \delta)$ collections, each of which has all their centers contained inside some ball of radius δ. By our previous argument, each such collection has at most 2^n elements, and this proves the lemma. $\square$

As noted earlier, this completes the proof of (2.1). $\square$

3. Proof of the lower bound in Theorem 1.1

Next we consider the opposite direction:

$$\ell(\Gamma) - \mathrm{diam}(\Gamma) \gtrsim \sum_Q \beta_\Gamma^2(Q)\mathrm{diam}(Q). \tag{3.1}$$

In the case $n = 2$ we will actually prove that

$$\ell(\Gamma) - \frac{1}{2}\mathrm{prm}(\Gamma) \gtrsim \sum_Q \beta_\Gamma^2(Q)\mathrm{diam}(Q), \tag{3.2}$$

where $\mathrm{prm}(\Gamma) = \ell(\partial(\mathrm{ch}(\Gamma)))$ denotes the perimeter of Γ, i.e., the length of the boundary of its planar convex hull (twice the length of Γ if it is a line segment). By noting that the orthogonal projection of a closed curve onto a diameter segment is 1-Lipschitz and at least 2-to-1, we see that the perimeter of Γ is at least twice its diameter. Hence (3.2) implies (3.1). In higher dimensions, the perimeter is replaced by a quantity called the "mean width" of Γ, defined below.

Proof of (3.1). Estimate (3.1) is proven using ideas from integral geometry. For the following facts, see [30].

There is a standard measure μ on the space of $(n-1)$-hyperplanes in $\mathbb{R}^n$, that is invariant under rigid motions of $\mathbb{R}^n$. In this proof "hyperplane" will always mean a $(n-1)$-dimensional affine space, and we will drop the explicit mention of the dimension. Each hyperplane $H \subset \mathbb{R}^n$ (except those passing though the origin; a set of μ measure zero) is determined by the point $p \in H$ closest to the origin. If $p = rx$ with $r > 0$, $x \in \mathbb{S}^{n-1} = \{x \in \mathbb{R}^n : |x| = 1\}$, the measure μ on hyperplanes is given by dr times $(n-1)$-measure on the unit sphere $\mathbb{S}^{n-1} \subset \mathbb{R}^n$.

Crofton's formula says there is a constant $c_n > 0$ so that

$$\ell(\Gamma) = c_n \int n(H, \Gamma) d\mu(H),$$

where $n(H, \Gamma)$ is number of points in $H \cap \Gamma$. See [30]. As a special case the measure of the set of hyperplanes hitting a line segment I is $c_n \ell(I)$ (almost every hyperplane hits a given segment at most once). The value of c_n is explicitly known, but not important to us; indeed, from this point on we normalize μ so that $c_n = 1$. Note that if S is the chord of Γ then

$$\mu(\{H : H \cap \Gamma \neq \emptyset\}) \geq \mu(H : H \cap S \neq \emptyset\}) = \mathrm{crd}(\Gamma). \tag{3.3}$$

The leftmost quantity in (3.3) is called the mean width of Γ; it is the average over S^{n-1} of the length of the projection of Γ onto the line every direction. For $n = 2$, this is a multiple of the perimeter of the convex hull K of Γ, and for $n = 3$ it is a multiple of the integral of the mean curvature over the surface of K (and this is often easier to compute). More generally, it is the coefficient V_1 of ϵ in Steiner's formula $\mathrm{vol}(K_\epsilon) = \mathrm{vol}(K) + V_1\epsilon + V_2\epsilon^2 + \cdots + V_n\epsilon^n$, where K_ϵ is the ϵ-neighborhood of K. The V_k's are the intrinsic volumes of K and every rigid motion invariant on convex sets is a combination of these. See, e.g., [18], [30].

Let $\mathcal{D}$ denote the collection of dyadic cubes in $\mathbb{R}^n$ and let $\mathcal{D}^*$ be the union of all possible translates of $\mathcal{D}$ by $\{-\frac{1}{3}, 0, \frac{1}{3}\}$ along any subset of the n coordinates. The family $\mathcal{D}^*$ has the property that any bounded subset of $\mathbb{R}^n$ is contained in some member of $\mathcal{D}^*$ of comparable size (the $\frac{1}{3}$-trick, [27]). Also, the translate of any dyadic cube by $\ell(Q)/3$ along any subset of the coordinates is in $\mathcal{D}^*$ (to see this, note that one of $2^n \pm 1$ is divisible by 3 hence $\frac{1}{3}2^{-n} = \pm\frac{1}{3} + k2^{-n}$ has a solution). The set of all such translates of Q is denoted $\mathcal{D}^*(Q)$.

For $Q \in \mathcal{D}^*$ let $S(Q, \Gamma)$ be the set of hyperplanes that intersect both $\frac{5}{3}Q \cap \Gamma$ and $(3Q \setminus 2Q) \cap \Gamma$. For a hyperplane H, let $N(H, \Gamma)$ be the number of cubes $Q \in \mathcal{D}^*$ such that $H \in S(Q, \Gamma)$.

Lemma 3.1. $n(H, \Gamma) - 1 \gtrsim N(H, \Gamma)$ *whenever* $n(H, \Gamma) > 0$.

Proof. Assume $n(H,\Gamma)$ is finite (otherwise there is nothing to do). Since $\mathcal{D}^*$ consists of a finite union of families of translations of the dyadic cubes $\mathcal{D}$, its suffices to bound the number of cubes belonging to each family. Since the argument is the same for each family, we just consider the dyadic cubes. By breaking the dyadic family into a finite number of sub-families, we may also assume the cubes are "M-sparse", i.e., there is a large constant M so that any two cubes Q, Q' of the same size satisfy $\mathrm{dist}(Q,Q') \geq M\mathrm{diam}(Q)$ and cubes Q, Q' of different sizes satisfy either $\mathrm{diam}(Q) \geq M\mathrm{diam}(Q')$ or the reverse inequality. Let $\mathcal{Q}$ be one such a collection of sparse dyadic cubes.

We define a graph $G = (V, E)$ with vertices $V = H \cap \Gamma$ and an edge between $z, w \in V$ if there is some cube $Q \in \mathcal{Q}$ with $z \in \frac{5}{3}Q$ and $w \in 3Q \setminus 2Q$. Let $\mathcal{C} \subset \mathcal{Q}$ be the cubes Q for which such a pair (z, w) exists. Note that $\mathrm{diam}(Q) \simeq |z-w|$, so the number of dyadic cubes associated to this pair is uniformly bounded. Moreover, by sparseness there is at most one $Q \in \mathcal{Q}$ associated to any pair (z,w). Thus $\#(\mathcal{C}) \leq \#(E)$ (where $\#$ is cardinality). In fact, if Q is associated to more than one edge in G, we remove all but one of these edges from the graph (that we still call G), so that $\#(\mathcal{C}) = \#(E)$. If G has no cycles, i.e., all its connected components are trees, then $\#(\mathcal{C}) = \#(E) \leq \#(V) - 1 = n(H,\Gamma) - 1$ and we are done.

Note that if M is chosen large enough, and $e_1, e_2 \in E$ are adjacent edges of G whose lengths are comparable to within a factor of 100, then by sparseness, e_1 and e_2 must correspond to the same cube Q, and hence $e_1 = e_2$. Thus adjacent edges in G have lengths differing by a factor of at least 100.

It suffices to prove G has no cycles, so suppose $e_1, \dots e_N$ are the ordered edges of cycle in G. We want to show this is impossible. We will call a pair of edges e_j, e_k in the cycle a "good pair" if they have sizes that are comparable within a factor of 100, but are connected by a non-empty path of edges that are all smaller by a factor of 100. We may assume e_1 is the longest edge. Thus it is one element of a good pair; the other element is either itself (if it is the only edge with comparable length in the cycle) or another edge e_k; in the latter case e_1 and e_k can't be adjacent by our earlier remarks. We claim there must be another good pair on the path connecting e_1 to e_k.

Let f be the largest edge between e_1 and e_k. Since e_1, e_k is a good pair, $\ell(f) \leq min(\ell(e_1), \ell(e_k))/100$. If there is a second edge with length comparable to f (within a factor of 100), then since the path from e_1, e_2 is finite, then some pair of these comparable edges gives another good pair. Otherwise there is no other edge with length comparable to f. Now consider the largest edge f' between e_1 and f. As above, there is either another edge between e_1 and f of comparable size (and hence another good pair), or just a single such edge f' with $\ell(f') \leq \ell(f')/100$ (since we have assumed there are no edges with length that is within a factor of 100 of f's). The same argument applies to the path from f to e_k. Thus either we find a good pair or there are at most two edges (besides f) that are longer than $\ell(e_1)/100^2$. Continuing in the same way, we either find a good pair or we prove that there are at most $2^n - 1$ edges of length $\geq 100^{-n-1}\ell(e_1)$. Thus the total length of all the edges in the path from e_1 to e_k is at most $\sum_{n=1}^{\infty} 2^n \cdot 100^{-n}\ell(e_1) \ll \ell(e_1)$. However, the path either connects the endpoints of e_1, or connects e_1 to an edge e_k whose distance from e_1 is $\geq 100 \cdot \ell(e_1)$. In either case, the path is too short, so the assumption that there are no good pairs must be wrong.

So a good pair of edges e_j, e_i with diameters $\leq 100 \cdot \mathrm{diam}(e_1)$ must exist. But these edges are separated by distance at least $100 \cdot \mathrm{diam}(e_j)$ and the same argument as above implies the path between them contains another good pair smaller by at least a factor of 100. Continuing in this way, we see the proposed cycle in G contains arbitrarily many edges, and this contradicts the fact that G is a finite graph. Thus G is a forest, and the lemma is proven. □

Continuing with the proof of (3.1), note that by (3.3)

$$\begin{aligned} \ell(\Gamma) - \mathrm{crd}(\Gamma) &\geq \int n(H,\Gamma)d\mu(H) - \int_{H\cap\Gamma\neq\emptyset} 1 d\mu(H) \\ &\gtrsim \int N(H,\Gamma)d\mu(H) \\ &= \sum_{Q^*} \mu(S(Q^*,\Gamma)). \end{aligned}$$

To complete the proof we need

Lemma 3.2. *For every dyadic cube Q there is a intersecting $Q^* \in \mathcal{D}^*$ of comparable size so that $\beta_\Gamma^2(Q)\mathrm{diam}(Q) \lesssim \mu(S(Q^*,\Gamma))$.*

Given the lemma, we deduce that

$$\sum_{Q\in\mathcal{D}} \beta^2(Q)\mathrm{diam}(Q) \lesssim \sum_{Q^*\in\mathcal{D}^*} \mu(S(Q^*,\Gamma)),$$

and hence

$$\ell(\Gamma) - \mathrm{crd}(\Gamma) \gtrsim \sum_{Q\in\mathcal{D}} \beta^2(Q)\mathrm{diam}(Q),$$

as desired. To prove Lemma 3.2 we will need some preliminary facts.

Lemma 3.3. *Suppose I is the unit segment between 0 and 1 on the x_1-axis, and suppose J is another unit length segment in $\mathbb{R}^n$ with $\mathrm{dist}(I,J) \geq 1$ and with at least one endpoint inside $\{|x| < 100\}$ that is distance $\beta > 0$ from the x_1-axis. Then the μ-measure of the hyperplanes hitting both I and J is $\gtrsim \beta^2$.*

Proof. Think of I as fixed and J as variable. The measure of the hyperplanes hitting both I and J is a continuous function of J and is non-zero as long as J is not a subset of X_1, the x_1-axis. Thus it is bounded away from zero as long as J has one endpoint outside the cylinder of radius $1/100$ around X_1.

Since the lemma is true if $\beta \geq 1/100$, now suppose $\beta < 1/100$. Then the hyperplanes that hit both I and J have unit perpendicular vectors that within $O(\beta)$ of the $(n-1)$-unit-sphere S^{n-1} in $X_1^\perp$, the orthogonal complement of X_1. Consider the linear map that is the identity on X_1 and expands by a factor of $b = 1/(100\beta)$ on $X_1^\perp$. In the $p = rx$ parameterization of hyperplanes, this map can change r by a factor of $O(b)$ and changes the x_1 coordinate of x by at most a factor of b. Therefore the μ measure of hyperplanes hitting I and J is increased by at most a factor of $O(b^2) = O(\beta^{-2})$ and the new measure is bounded uniformly away from zero. This proves the lemma. $\square$

Lemma 3.4. *Suppose Q is a dyadic cube in $\mathbb{R}^n$ and $Q^* \in \mathcal{D}^*$ is cube of the same size and is a translation of Q by at most $\ell(Q)/3$ in each coordinate direction. Then the distance from $2Q \cup 2Q^*$ to $\mathbb{R}^n \setminus (3Q \cap 3Q^*)$ is at least $\ell(Q)/6$.*

Proof. If $z \in 2Q \cup 2Q^*$ then $P_k(z) \in 2I \cup 2I^* = P_k(2Q) \cup P_k(2Q^*)$ for every $k = 1, \dots, n$, where P_k is the orthogonal projection onto the kth coordinate axis. Similarly, if $w \notin (3Q \cap 3Q^*)$ then $P_k(z) \notin 3I \cap 3I^* = P_k(3Q) \cap P_k(3Q^*)$ for some choice of k. But in dimension 1, the distance between $2I \cup 2I^*$ and $\mathbb{R} \setminus (3I \cap 3I^*)$ is easily computed to be $\ell(I)/6$, e.g., if $I=[0,1]$ and $I^*=[\frac{1}{3},\frac{4}{3}]$, then $2I\cup 2I^*=[-\frac{1}{2},\frac{11}{6}]$ and $3I\cap 3I^*=[-\frac{2}{3},2]$. Hence $|z-w|$ is at least this big. $\square$

Proof of Lemma 3.2. The idea is simple: by Lemma 3.3 and the fact that any hyperplane hitting a chord S of a sub-arc $\gamma\subset\Gamma$ also hits γ, the estimate reduces to finding two sub-arcs γ_1,γ_2 of $\Gamma\cap 3Q$ with length $\geq ell(Q)/100$ and $\leq \ell(Q)/2$ and so that $\mathrm{dist}(\gamma_1,\gamma_2)\geq \ell(Q)/2$, and also whose chords are β far from lying one the same line. The separation property will hold if one chord lies in $\frac{3}{2}Q$ and the other in $3Q\setminus 2Q$.

By rescaling we may assume Q has side length 1, and that Γ hits both Q and $3Q^c$. Set $\beta=\beta_\Gamma(2\frac{23}{24},Q)$. For any child Q' of Q, $3Q'\subset 2Q\subset 2\frac{23}{24}Q$, so $\beta_\Gamma(Q')\leq\beta=\beta_\Gamma(2\frac{23}{24},Q)$. Thus the β-numbers for the children of Q are all bounded by $O(\beta)$, and we will show that this is either $O(\mu(S(Q,\Gamma)))$ or $O(\mu(S(Q^*,\Gamma)))$ for some $Q^*\in\mathcal{D}^*(Q)$. Thus finding the two chords in $\Gamma\cap Q$ will actually bound the β numbers for the children of Q.

Choose $z\in\Gamma\cap\partial Q$ and define the ball $B_1=B(z,\frac{1}{24})$. Choose $w\in\Gamma\cap\partial B_1$ so that z and w are connected by a sub-arc $\gamma_1\subset\Gamma\cap B_1$. Let $S_1=[z,w]$ be the segment connecting them. Then $S_1\subset 1\frac{2}{3}Q$ and any hyperplane that hits S_1 must also hit $\gamma_1\subset\Gamma\cap 1\frac{2}{3}Q$. Let L_1 be the line that contains S_1 and let $W_0\subset W_1$ be the cylinders of radius $\beta/1000$ and $\beta/2$ respectively, both with axis L_1. Since these radii are less than β, we know $\Gamma\cap 2\frac{23}{24}Q$ contains a point outside the cylinder W_1. There are two cases to consider, depending on where this point is.

Case 1: Suppose there is a point

$$v\in\Gamma\cap\left(2\frac{23}{24}Q\setminus 2\frac{1}{24}Q\right)\setminus W_0.$$

See left side of Figure 4. Since Γ is path connected and has diameter $\geq\mathrm{diam}(Q)$, v can be connected to a point $u\in\Gamma$ with $|u-v|=\frac{1}{24}$ (we may assume u and v are connected by a sub-arc $\gamma_2\subset\Gamma$ that stays inside the ball $B_2=B(v,\frac{1}{24})$ by starting at v and following γ until it first leaves B_2). Let $S_2=[u,v]$ be the segment connecting these two points and note that any line that hits S_2 also hits $\gamma_2\subset\Gamma\cap D_2\subset 3Q\setminus 2Q$. Since S_2 has an endpoint outside W_0, the measure of the set of lines that hits both S_1 and S_2 is $\gtrsim\beta^2$ by Lemma 3.3. Thus $\mu(S(Q,\Gamma))\gtrsim\beta^2$ as well and the lemma is satisfied with $Q^*=Q$.

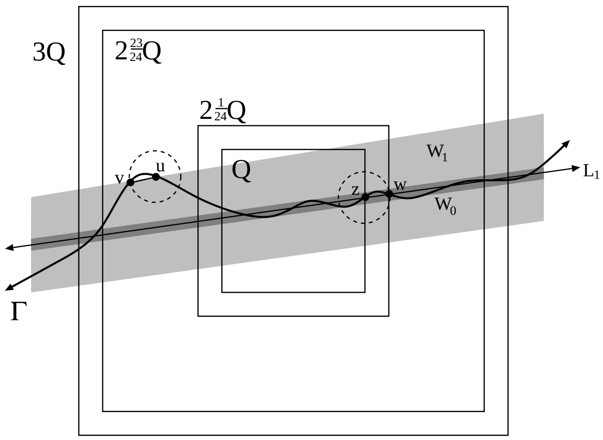

FIGURE 4. In both cases, we can find two segments inside $3Q$ that have lengths comparable to $\mathrm{diam}(Q)$ and are not both close to common line. Lemma 3.3 them implies at least measure $\simeq \beta^2$ of hyperplanes hit both segments. For clarity, the pictures are not quite to scale.

Case 2: Suppose Case 1 does not hold. Then there must be a point $p \in \Gamma \cap (2\frac{1}{24}Q \setminus W_1)$. Choose q with $|p-q| \leq \frac{1}{12}$ that is connected to p by a sub-arc $\gamma_3 \subset \Gamma \cap B_3$, where $B_3 = B(p, \frac{1}{12}) \subset 2\frac{1}{3}Q$. See right side of Figure 4. Let $S_3 = [p,q]$. As before, any line that hits S_3 must also hit γ_3. Choose an element $Q^* \in \mathcal{D}^*(Q)$ that is the same size as Q and translated by at most $\frac{1}{3}\ell(Q)$ in each coordinate direction (possibly Q itself) and so that $B_3 \subset 1\frac{2}{3}Q^*$ (it is easy to check there is at least one such cube Q^* by considering the projections onto each coordinate).

By Lemma 3.4, Γ must contain a point $u \in (3Q \cap 3Q^*) \setminus (2Q \cup 2Q^*)$ that is at least distance $\geq \frac{1}{12}$ from the boundaries of both $2Q$ and $3Q$. Therefore the ball $B_4 = B(u, \frac{1}{24})$ is inside of $2\frac{23}{24}Q \setminus 2\frac{1}{24}Q$. As before, we can find a radius S_4 of B_4 so that any line that hits S_4 also hits a sub-arc $\gamma_4 \subset \Gamma \cap B_4$. Since γ_4 lies inside the very thin cylinder W_0, the line L_4 containing S_4 is almost parallel to L_1 (the axis of W_0) and so $L_4 \cap 3Q$ lies inside a cylinder of radius $\beta/4$ around L_1. Since p is outside the larger cylinder W_1 we see that S_3 and S_4 satisfy Lemma 3.3, hence the measure of the set of lines that hit both S_4 and S_3 is $\gtrsim \beta^2$ by Lemma 3.3. Thus $\mu(S(Q^*, \Gamma)) \gtrsim \beta^2$, as desired. $\square$

This completes the proof of (3.1), and hence of Theorem 1.1. $\square$

4. Proof of Theorem 1.2

Proof. As noted earlier, "$\gtrsim$" is immediate from Theorem 1.1 since $\mathrm{crd}(\Gamma) \le \mathrm{diam}(\Gamma)$.

To prove the other direction, we may assume the β^2-sum in (1.4) is finite, for otherwise there is nothing to prove. Thus we may assume γ is rectifiable. We may also assume $\mathrm{diam}(\Gamma) = 1$. Let Q_0 be a dyadic cube hitting Γ with $1 \le \mathrm{diam}(Q_0) \le 2$, hence $\Gamma \subset 3Q$. Suppose β_0 is a small positive number (chosen to satisfy various conditions described below). If $\beta_\Gamma(Q_0) > \beta_0$, then the result is trivially true since then

$$\mathrm{crd}(\Gamma) \le \mathrm{diam}(\Gamma) = 1 \le \frac{1}{\beta_0^2}\beta_\Gamma^2(Q_0)\mathrm{diam}(Q_0) \lesssim \beta_\Gamma^2(Q_0)\mathrm{diam}(Q_0),$$

(with constant depending on β_0) and hence the $\mathrm{crd}(\Gamma)$ term in (1.4) can be absorbed into the β^2-sum term.

Therefore we may assume $\beta_\Gamma(Q_0) \le \beta_0$. Let $S = [x, y]$ be a diameter segment of Γ and let γ_0 be the open subarc of Γ connecting x and y. Then $\Gamma \setminus \gamma_0$ consists of two arcs, γ_1 connecting x to an endpoint z (possibly $z = x$) and γ_2 connecting y to the other endpoint w (possibly $w = y$).

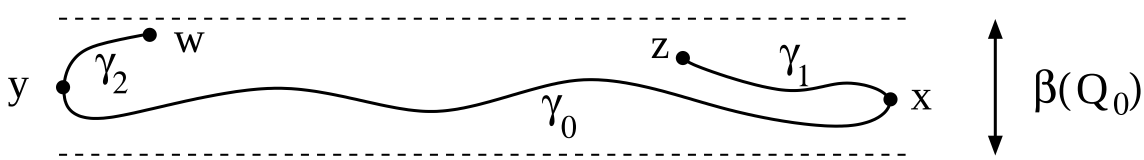

Figure 5. Definitions for the proof of Theorem 1.2.

By rotating and rescaling, we may assume that $x = 1, y = -1$ on the x_1-axis. See Figure 5. Note that

$$\mathrm{crd}(\Gamma) \ge \mathrm{diam}(\Gamma) - \ell(\gamma_1) - \ell(\gamma_2)$$

and hence using Theorem 1.1 (in particular (2.1)) we get

$$\begin{aligned} \ell(\Gamma) - \mathrm{crd}(\Gamma) &\le \ell(\Gamma) - \mathrm{diam}(\Gamma) + \ell(\gamma_1) + \ell(\gamma_2) \\ &\le O\left(\sum_Q \beta_\Gamma^2(Q)\mathrm{diam}(Q)\right) + \ell(\gamma_1) + \ell(\gamma_2). \end{aligned}$$

Thus Theorem 1.2 will follow if we can show that both $\ell(\gamma_1)$ and $\ell(\gamma_2)$ are bounded by a multiple of the β^2-sum for Γ. Because of (1.2), i.e., the usual form of the traveling salesman theorem, and the fact that $\beta_{\gamma_1}(Q), \beta_{\gamma_2}(Q)$ are both at most $\beta_\Gamma(Q)$, it is enough to bound the diameters of these arcs by the β_Γ^2-sum; then the diameters can be absorbed into the sum by making the comparability constant larger. The arguments for both arcs are the same, so we only discuss γ_1.

Let $\epsilon = \mathrm{diam}(\gamma_1)$. Assume $\epsilon > 0$ (otherwise there is nothing to do). Let $Q_1, \dots, Q_k$ be the nested dyadic cubes containing x with diameters going from $\mathrm{diam}(Q_0)$ to ϵ. Note that $k \simeq \log(\mathrm{diam}(Q_0)/\epsilon)$. If any one of these cubes satisfies $\beta_\Gamma(Q_j) \geq \beta_0$, then

$$\mathrm{diam}(\gamma_1) \leq \frac{\beta^2_{\Gamma(Q_j)}}{\beta_0^2}\mathrm{diam}(\gamma_1) \lesssim \beta^2_{\Gamma(Q_j)}\mathrm{diam}(Q_j),$$

and hence $\mathrm{diam}(\gamma_1)$ is dominated by the β^2-sum, as desired. For the remainder of the proof we may therefore assume that $\beta_\Gamma(Q_j) \leq \beta_0$ for all $j \in \{0, 1, \dots, k\}$. Let L_j be a best line in the definition of $\beta_\Gamma(Q_j)$.

Case 1: Assume that for some $j \in \{1, \dots, k\}$, the line L_j makes an angle larger than $10\beta_0$ with the x_1-axis. Since the angle between L_0 and L_j is bounded by $O\left(\sum_{i=0}^j \beta(Q_i)\right)$, and we have normalized so that the best line for Q_0 is within β_0 of the x_1-axis, we must have $\sum_{j=1}^k \beta_\Gamma(Q_j) \gtrsim \beta_0 \gtrsim 1$. The Cauchy-Schwarz inequality then implies

$$1 \quad \lesssim \quad \left(\sum_{j=1}^k \beta_\Gamma(Q_j)\right)^2 \leq \left(\sum_{j=1}^k \beta_\Gamma^2(Q_j)2^{-j}\right) \cdot \left(\sum_{j=1}^k 2^j\right) \simeq 2^k \sum_{j=1}^k \beta_\Gamma^2(Q_j)2^{-j}$$

so $\sum_{j=1}^k \beta_\Gamma^2(Q_j)2^{-j} \gtrsim 2^{-k} \gtrsim \epsilon$, and hence

$$\epsilon = \mathrm{diam}(\gamma_1) \lesssim \sum_{j=1}^k \beta_\Gamma^2(Q_j)\mathrm{diam}(Q_j),$$

as desired.

Case 2: Next we assume that all the lines L_j, $j = 0, \dots, k$ make angle $\leq 10\beta_0$ with the x_1-axis. Consider a subarc $\gamma_1' \subset \gamma_1$ that is contained in, and connects the boundary components of, the annulus

$$\{p \in \mathbb{R}^2 : \frac{1}{10}\mathrm{diam}(\gamma_1) \leq |p - x| \leq \frac{1}{5}\mathrm{diam}(\gamma_1)\}.$$

Since γ_1 and γ_1' have comparable diameters, it is enough to bound $\mathrm{diam}(\gamma_1')$.

For each $p \in \gamma_1'$ a dichotomy holds: either every dyadic cube Q containing p with $\mathrm{diam}(Q) \leq \mathrm{diam}(\gamma_1)/10$ satisfies $\beta_\Gamma(Q) \leq \beta_0$ or there is a cube Q_p of this form such that $\beta_\Gamma(Q_p) > \beta_0$. Let $E \subset \gamma_1'$ be the set of points p where such a Q_p exists. Since we can assume γ_1 is rectifiable, almost every point of γ_1 is a tangent point.

Lemma 4.1. *If $p \in \gamma_1' \setminus E$ and p is a tangent point of Γ, then p has the following "crossing property": if Q is a dyadic cube containing p with* $\operatorname{diam}(Q) \le \operatorname{diam}(\gamma_1)/10$ *then γ_1 must "cross" Q in the sense that the orthogonal projection of $\gamma_1 \cap 3Q$ onto L_Q covers $L_Q \cap Q$, where L_Q is a best approximating line for the definition of $\beta_\Gamma(Q)$. In other words, γ_1 must connect the two components of $W \cap \partial 3Q$ where W is a cylinder of radius $1/10$ with axis L_Q passing through p.*

Proof. Note that because p is not in E, that γ has small β-number for Q and for every dyadic subcube of Q that contains p. Using this, we claim we can construct a $(n-1)$-surface σ so that

(1) σ cuts $3Q$ into two pieces,
(2) σ separates the endpoints of $\gamma_1 \cap 3Q$,
(3) σ contains p, but no other points of γ_1, and
(4) $\sigma \cap 3Q' \setminus Q'$ is nearly orthogonal to $L_{Q'}$ for each dyadic Q' with $p \in Q' \subset Q$.

To do this, choose a $(n-1)$-sphere of the n-sphere of radius $t = \operatorname{diam}(Q)2^{-n}$ around p that is nearly orthogonal to the optimal line L_n passing through p for the definition of $\beta(p, 2^{-n})$ and then connecting these $(n-1)$-spheres by a surface (e.g., project both onto the $(n-1)$-plane $L_n^\perp$ orthogonal to L^n and connect two points if the projections are on the same ray in $L_n^\perp$). See Figure 6 for the planar picture.

If p is a tangent point of γ_1, then σ has a tangent plane at p that is perpendicular to γ_1's tangent direction. From this we see that γ_1 crosses σ, i.e., it hits both components of $3Q \setminus \sigma$. Since γ_1 only hits σ once, it must leave $3Q$ through a different component of $\partial(3Q) \cap W$ than it entered through. This implies Lemma 4.1. $\square$

Note in the previous proof, that if Γ contains two sub-arcs that connect different ends of the cylinder W, then both sub-arcs must cross through p and hence Γ is not a Jordan curve. This observation will be used later to prove Corollary 5.1.

Lemma 4.2. $\ell(E) = \ell(\gamma_1')$.

Proof. If not, then we can choose a non-empty subset $F \subset \gamma_1' \setminus E$ that consists entirely of tangent points of γ. Suppose $p \in F$ and define d to be the distance

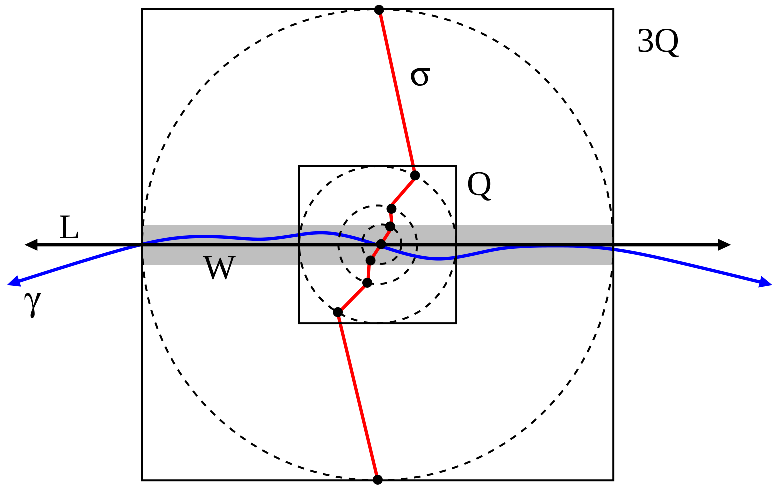

FIGURE 6. If all the β's are small at p, and p is a tangent point of γ, then γ must cross Q in the sense that the orthogonal projection of $\gamma \cap 3Q$ on the line L must cover $L \cap Q$, i.e., γ can't "double back" and leave using the same end of W that it entered.

from p to γ_0. By the assumption that every L_j is close to horizontal, we know $d = O(\beta_0 \operatorname{diam}(\gamma_1)) < \operatorname{diam}(\gamma_1)/10$. Also note that d is positive since p is not on γ_0. Let Q_p be the dyadic square containing p with diameter $2d < \operatorname{diam}(Q_p) \leq 4d$. Because $\operatorname{diam}(Q_p) \leq \operatorname{diam}(\gamma_1)/10$, the argument in the previous paragraph applies, and γ_1 must cross Q_p inside a cylinder S of width $\beta_0 \operatorname{diam}(Q_p)$. Moreover, since $\operatorname{diam}(Q_p) > 2d$, the curve γ_0 also hits $3Q$ and hence contains a point q in the same cylinder S, and hence γ_0 is at most distance $\beta_0 \operatorname{diam}(Q_p)$ from γ_1. For β_0 small, this value is much smaller than d, giving a contradiction. Thus no such p exists, and hence $\ell(E) = \ell(\gamma_1')$, so Lemma 4.2 holds. $\square$

By the nested property of dyadic cubes, we can find a collection $\{Q_p^j\}$ of cubes as in the definition of E that have disjoint interiors and that covers E. Hence

$$\ell(\gamma_1') \simeq \ell(E) \leq \sum_j \ell(Q_p^j \cap E) \lesssim \sum_j \left[\operatorname{diam}(Q_p^j) + \sum_{Q \subset Q_p^j} \beta_E^2(Q) \operatorname{diam}(Q) \right]$$

where we have applied (1.2), say with $\delta = 1$, to each set $Q_p^j \cap E$. Note that usual formulation of the traveling salesman theorem is to sum over all dyadic cubes in $\mathbb{R}^n$, but if $E \subset Q$, then it suffices to sum over all cubes contained in Q (including Q itself) since the β^2-sum over all larger cubes that hit E form a geometric series whose sum is $O(\beta^2(Q)\mathrm{diam}(Q))$. See also Lemma B.2.

Now we use (1.2) and the fact that $\beta_E \le \beta_\Gamma$, to show

$$\ell(\gamma_1') \;\lesssim\; \sum_j \sum_{Q\subset Q_p^j} \beta_\Gamma^2(Q)\mathrm{diam}(Q),$$

where we have also used $\beta_\Gamma(Q_p^j) \simeq 1$ to absorb the $\mathrm{diam}(Q_p^j)$ terms into the β^2-sums. Since this is a β^2-sum over disjoint collections of dyadic cubes, it is dominated by the full β^2-sum, and this completes the proof of Theorem 1.2. $\square$

For the proof of (1.2) we can refer to [16], or note that it is also proven in Appendix A of this paper. This proof uses a slight modification of the arguments in Section 2 and is independent of the current argument.

5. Corollaries and questions

Next we derive some consequences of our arguments that are used to derive the characterizations of Weil-Petersson curves given in [5]. As noted in the introduction, such curves satisfy the condition $\sum_Q \beta^2(Q) < \infty$, where the "diam$(Q)$" has been dropped from Jones's characterization of rectifiable curves.

Corollary 5.1. *If Γ is a closed Jordan curve and $S = \sum_Q \beta_\Gamma^2(Q) < \infty$, then Γ has bounded turning, i.e., there is an $M < \infty$ so that any pair $z, w \in \Gamma$ is connected by a sub-arc γ with* $\mathrm{diam}(\gamma) \le M|z-w|$. *We may take $M = O(e^{O(S)})$.*

Proof. Suppose not. Then given any $M < \infty$ there are $z, w \in \Gamma$, so that both subarcs connecting them have diameter $\ge M|z-w|$. Rescale so that $|z-w| = 1$. Then we can find disjoint arcs $z \in \gamma_1, w \in \gamma_2$ with endpoints z', z'' and w', w'' respectively, so that all four of these points are at least distance M from $v = (z+w)/2$. Choose a positive integer N so that $M/2 < 2^N \le M$. See Figure 7. Consider the annuli $A_n = \{y : 2^n \le |y-v| < 2^{n+1}\}$ for $n = 1, \dots, N$ and let $\mathcal{Q}_n$ be the collection of dyadic cubes that hit A_n and have diameter $\le 2^{n+1}$. These collections have bounded overlap, so $\sum_n \sum_{\mathcal{Q}_n} \beta_\Gamma^2(Q) \le C \cdot S$, and hence there is some n so that $\sum_{\mathcal{Q}_n} \beta_\Gamma^2(Q) \le 1/10$, if $\log M \ge 10 \cdot C \cdot S$. Both components of $\gamma_1 \setminus \{z\}$ cross A_n, as do both components of $\gamma_2 \setminus \{w\}$. Thus there is a radial cylinder of radius $2^n/100$ and length 2^n that connects the two boundary components of A_n and contains two disjoint sub-arcs of Γ that also cross A_n. The proof of Lemma 4.1, however, shows that there is a point p on in A_n that both arcs must pass through, a contradiction. Thus Γ has the bounded turning property for some M with $\log M \lesssim S = \sum_Q \beta^2_\gamma(Q)$. $\square$

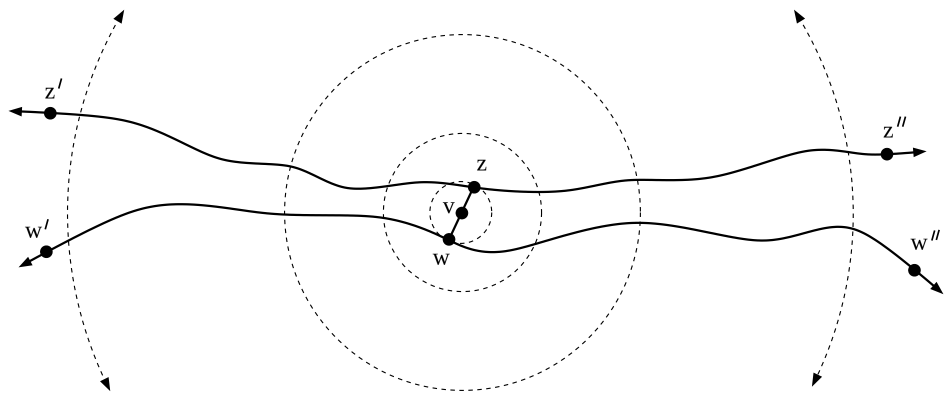

FIGURE 7. If Γ does not have bounded turning, then there are points z, w that cannot be connect by a subarc of diameter $O(|z-w|)$. This means these points lie on disjoint subarcs whose four endpoints are all distance $\gg |z-w|$ from z and w.

Corollary 5.2. *If Γ is a closed Jordan curve and $S = \sum_Q \beta^2_\Gamma(Q) < \infty$, then Γ is chord-arc, i.e., any pair of points $z, w \in \Gamma$ are connected by a sub-arc γ with $\ell(\gamma) \lesssim |z-w|$.*

Proof. Suppose $z, w \in \Gamma$ and $\gamma \subset \Gamma$ is a sub-arc with endpoints z, w and $\mathrm{diam}(\gamma) \leq M|z-w| = M\mathrm{crd}(\gamma)$, with $M = O(\exp(O(S))$ as in the previous corollary. Then by Theorem 1.2 and the fact that it suffices to sum over cubes Q with $\mathrm{diam}(Q) \lesssim \mathrm{diam}(\gamma)$ we get

$$\begin{aligned}
\ell(\gamma) &\leq \mathrm{crd}(\Gamma) + O\left(\sum_Q \beta^2_\gamma(Q)\mathrm{diam}(Q)\right) \\
&\leq \mathrm{crd}(\gamma) + O\left(\mathrm{diam}(\gamma)\sum_Q \beta^2_\gamma(Q)\frac{\mathrm{diam}(Q)}{\mathrm{diam}(\gamma)}\right) \\
&\leq \mathrm{crd}(\gamma) + M\mathrm{crd}(\gamma)\cdot\sum_Q \beta^2_\gamma(Q)\cdot O(1) \\
&= \mathrm{crd}(\gamma)\cdot(1+O(MS)) = \mathrm{crd}(\gamma)\cdot(1+O(S\exp(O(S)))). \quad \square
\end{aligned}$$

The corollaries fail if Γ is not a closed arc, e.g., take a circle with an ϵ-long arc removed. Is it true that if $\sum_Q \beta_\Gamma^2 < \infty$ for a Jordan arc Γ, then Γ is a subarc of a chord-arc curve Γ'? It seems likely that an argument similar to the proof of Lemma 4.1 can be used to extend Γ past an endpoint on scales where β is small, but something more clever is needed on scales where β is large. More generally, if E is a general set and $\sum_Q \beta_Q^2(E) < \infty$, is E a subset of a curve Γ with $\sum_\Gamma \beta_\Gamma^2(E) < \infty$? Jones's traveling salesman theorem shows this is true for sums of the form $\sum_Q \beta_Q^2(E)\mathrm{diam}(Q)$. What about sets satisfying $\sum_Q \beta_Q^2(E)\mathrm{diam}^s(Q)$ for $0 < s < 1$?

An earlier draft of this paper asked if Theorem 1.2 is true in Hilbert space. This has since been verified by Jared Krandel in [19]. Does this theorem hold in other metric spaces where the usual TST is known to hold? Does (1.2) hold for general sets in a metric space if (1.1) does? What about (1.3)? These questions seem analogous to the fact, proven independently by Arora [2] and Mitchell [23], that one can compute $(1+\epsilon)$-approximations to the classical traveling salesman problem in polynomial time for finite sets in Euclidean space, but this is unknown for $\epsilon < 1/2$ in metric spaces, and computing any bounded approximation is NP hard for general weighted graphs. What is the proper "analytic" version of this?

Appendix A. The TST for general sets

We have done most of the work needed to prove the traveling salesman theorem for general sets $E \subset \mathbb{R}^n$, not just for Jordan curves. For the convenience of the reader, we explain how to prove (1.2) by modifying the argument in Section 2..

The only change is in the splitting procedure. As before, we start with Γ_0 the convex hull of E. In general, we will have a collection of convex sets $\mathcal{R}_n$ and line segments $\mathcal{S}_n$ whose union is a closed, connected set Γ_n that contains E. As before, each convex, compact set $R \in \mathcal{R}_n$, will be convex hull of $R \cap E$. The intersection $\Gamma = \cap \Gamma_n$ is a compact, connected set containing E and we wish to bound its 1-dimensional Hausdorff measure.

Given $R \in \mathcal{R}_n$, we take a diameter segment I and split it into three equal thirds: the middle segment J_0 and the two ends J_1, J_2. If the orthogonal projection of $R \cap E$ contains a point $v \in J_0$, then cut I into two pieces I_1, I_2 using this point and replace R by two pieces R_1, R_2 that are the convex hulls of the parts of $R \cap E$ that project onto I_1, I_2 respectively. If the projection of onto J_0 is empty, then define R_1, R_2 as the convex hulls of the parts of $R \cap E$ that project onto J_1, J_2 respectively and a shortest possible line segment S connecting R_1 and R_2. The union of the nth generation sets and segments is clearly a connected compact covering of E, so the sum of the diameters of these sets and segments is an upper bound for the shortest connected set containing E (shortest in the sense of 1-dimensional Hausdorff measure).

The only change needed in the earlier proof is to Lemma 2.1. It becomes

Lemma A.1. *If R is split into R_1, R_2 sets and a segment S as above, then*

$$\mathrm{diam}(R_1) + \mathrm{diam}(R_2) + (1-\delta)\ell(S) \le \mathrm{diam}(R) + O(\frac{1}{\delta})\beta^2(R)\mathrm{diam}(R).$$

Proof. For the first case of the new splitting procedure, there is no segment S and each subset has diameter comparable to R, and the proof of Lemma 2.1 is the same as before. For the second case, first note that if $\beta(R) \ge \delta/20$, then

$$\begin{aligned} \mathrm{diam}(R_1) + \mathrm{diam}(R_2) + (1-\delta)\ell(S) &\le \ell(J_1) + \ell(J_2) + \ell(J_0) + 6\beta(R)\mathrm{diam}(R) \\ &\le \mathrm{diam}(R) + \frac{120}{\delta}\beta^2(R)\mathrm{diam}(R) \end{aligned}$$

If $\beta(R) < \delta/6$, then because $\ell(J_0) = \mathrm{diam}(R)/3$,

$$\begin{aligned} \mathrm{diam}(R_1) + \mathrm{diam}(R_2) + (1-\delta)\ell(S) &\le \ell(J_1) + \ell(J_2) + 8\beta^2(R)\mathrm{diam}(R) \\ &\qquad +(1-\delta)(\ell(J_0) + 4\beta^2(R)|R|) \\ &\le (1-\frac{\delta}{3})\mathrm{diam}(R) + 12\beta^2(R)\mathrm{diam}(R) \\ &\le (1-\frac{\delta}{3})\mathrm{diam}(R) + 12(\delta/6)^2\mathrm{diam}(R) \\ &\le \mathrm{diam}(R). \end{aligned}$$

This proves the lemma. □

The rest of the proof now proceeds as before, except that since

$$\sum_n \sum_{S \in \mathcal{S}_n} \ell(S) \le \ell(\Gamma),$$

we can replace (2.3) by

$$\begin{aligned}(1-\delta)\ell(\Gamma) &\leq \operatorname{diam}(\Gamma) + \limsup_n \sum_{R\in\mathcal{R}_n} \operatorname{diam}(R) + \sum_{k=1}^{n} \sum_{S\in\mathcal{S}_n} (1-\delta)\ell(S) \\ &\leq \operatorname{diam}(\Gamma) + O(\frac{1}{\delta}) \sum_n \sum_{R\in\mathcal{R}_n} \beta^2(R)\operatorname{diam}(R).\end{aligned}$$

Dividing both sides by $(1-\delta)$ proves (1.2). This implies the smallest connected set Γ containing a set E satisfies

$$C_1 \sum_Q \beta_E^2(Q)\operatorname{diam}(Q) \leq \ell(\Gamma) - \operatorname{diam}(E) \leq (1+\delta)\operatorname{diam}(E) + \frac{C_2}{\delta} \sum_Q \beta_E^2(Q)\operatorname{diam}(Q),$$

for some constants $0 < C_1, C_2 < \infty$, since the proof of the lower bound in Section 3 applies to any curve Γ containing E.

Appendix B. Equivalent formulations of TST

There are several formulations of Peter Jones's traveling salesman theorem and it is folklore that they are all equivalent to one another. Responding to requests from readers of an earlier draft of this paper, I give a precise formulation and proof of this "well known" fact.

A multi-resolution family in a metric space X is a collection of bounded sets $\{X_j\}$ in X such that there are $N, M < \infty$ so that

1. For each $r > 0$, the sets with diameter between r and Mr cover X,
2. each bounded subset of X hits at most N of the sets X_k with $\operatorname{diam}(X)/M \leq \operatorname{diam}(X_k) \leq M\operatorname{diam}(X)$.
3. any subset of X with positive, finite diameter is contained in at least one X_j with $\operatorname{diam}(X_j) \leq M\operatorname{diam}(X)$.

Dyadic intervals do not form a multi-resolution family, e.g., $X = [-1, 1] \subset \mathbb{R}$ is not contained in any dyadic interval, violating (3) above. However, the family of triples of all dyadic intervals (or cubes) do form a multi-resolution family. Similarly, if we "triple" the collection of dyadic intervals by adding all translates by $\pm 1/3$, we get a multi-resolution family (this is sometimes called the "$\frac{1}{3}$-trick", [27]). The analogous construction for dyadic cubes in $\mathbb{R}^n$ is to take all translates by elements of $\{-\frac{1}{3}, 0, \frac{1}{3}\}$ on each of the coordinates; this 3^n families of translates of $\mathcal{D}$ (including $\mathcal{D}$ itself). The union of these families is denoted $\mathcal{D}^*$.

We often deal with functions α that map a collection of sets into the non-negative reals (the β-numbers are an example), and will wish to decide if the sum $\sum_j \alpha(X_j)$ over some multi-resolution family converges or diverges. The following observation allows us to switch between various multi-resolution families without comment.

Lemma B.1. *Suppose $\{X_j\}$, $\{Y_k\}$ are two multi-resolution families on a space X and that α is a function mapping subsets of X to $[0,\infty)$ that satisfies $\alpha(E) \lesssim \alpha(F)$, whenever $E \subset F$ and $\mathrm{diam}(F) \lesssim \mathrm{diam}(E)$. Then*

$$\sum_j \alpha(X_j) \simeq \sum_k \alpha(Y_k).$$

Proof. By Condition (3) in the definition of a multi-resolution family, each X_j is contained in some set $Y_{k(j)}$ of comparable diameter. Hence $\alpha(X_j) \lesssim \alpha(Y_{k(j)})$ by assumption. By Condition (2), each Y_k can only contain a bounded number of X_j's of comparable size, so each Y_k is only chosen a bounded number of times as a $Y_{k(j)}$. Thus $\sum_j \alpha(X_j) \lesssim \sum_k \alpha(Y_k)$. The opposite direction follows by reversing the roles of the two families. $\square$

It is often convenient to consider several different formulations of the β-numbers. For $x \in \mathbb{R}^n$ and $t > 0$, define

$$\beta_\Gamma(x,t) = \frac{1}{t} \inf_L \max\{\mathrm{dist}(z,L) : z \in \Gamma, |x-z| \le t\},$$

where the infimum is over all lines hitting the ball $B = B(x,t)$ and let $\widetilde{\beta}_\Gamma(x,t)$ be the same, but where the infimum is only taken over lines L hitting x. Since this is a smaller collection, clearly $\beta(x,t) \le \widetilde{\beta}(x,t)$ and it is not hard to prove that $\widetilde{\beta}(x,t) \le 2\beta(x,t)$ if $x \in \Gamma$. See the center picture in Figure 8.

Given a Jordan arc γ with endpoints z, w we let

$$\beta(\gamma) = \frac{\max\{\mathrm{dist}(z,L) : z \in \gamma\}}{|z-w|},$$

where L is the line passing through z and w. See the right side of Figure 8.

Lemma B.2. *Suppose $-1 < s < 2$ and $\Gamma \subset \mathbb{R}^n$ is bounded Jordan curve (either closed or an arc). Then the following are equivalent:*

$$\sum_{Q \in \mathcal{D}} \beta_\Gamma^2(Q)\mathrm{diam}(Q)^s < \infty, \tag{B.1}$$

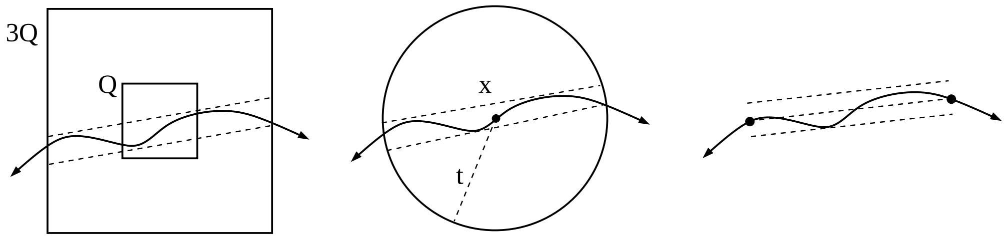

FIGURE 8. Three equivalent versions of the β-numbers.

$$\int_0^\infty \iint_{\mathbb{R}^n} \beta^2(x,t)\frac{dxdt}{t^{n+1-s}} < \infty, \tag{B.2}$$

If Γ is chord-arc, then (B.1) and (B.2) are also equivalent to

$$\int_0^\infty \int_\Gamma \widetilde{\beta}^2(x,t)\frac{dsdt}{t^{2-s}} < \infty, \tag{B.3}$$

$$\sum_j \beta^2(\Gamma_j)\mathrm{diam}(\Gamma_j)^s < \infty, \tag{B.4}$$

where dx is volume measure on $\mathbb{R}^n$, ds is arclength measure on Γ, and the sum in (B.4) is over a multi-resolution family $\{\Gamma_j\}$ for Γ. All four quantities are comparable with constants that depend only on the dimension n. Moreover, convergence or divergence in (B.2) and (B.3) is not changed if $\int_0^\infty$ is replaced by $\int_0^M$ (for any $M > 0$) and the values are all comparable for $M \geq \mathrm{diam}(\gamma)$. The convergence of the sum in (B.1) is unchanged if we only sum over cubes of diameter $\leq M$, for any $M > 0$ and are comparable for all values $M \geq \mathrm{diam}(\Gamma)$.

Since $\beta(x,t) \simeq \widetilde{\beta}(x,t)$ if $x \in \Gamma$, the integral in (B.3) is finite iff it is finite with β replacing $\widetilde{\beta}$. However, putting $\widetilde{\beta}$ into (B.2) gives a divergent integral for every closed Jordan curve Γ. The case $s = 1$ in the lemma corresponds to Peter Jones's traveling salesman theorem characterizing rectifiable curves, and $s = 0$ corresponds to the characterization of Weil-Petersson curves in [5]. Do other values of s correspond to interesting curve families?

Proof of Lemma B.2. Without loss of generality we may assume $\mathrm{diam}(\Gamma) = 1$.
(B.1) $\Leftrightarrow$ (B.2): If Q is a dyadic cube and $x \in Q$, then $\mathrm{diam}(Q) \geq \sqrt{n}t$ implies then $B(x,t) \subset 3Q$. In this case $\beta(x,t) \leq \sqrt{n}\beta(Q)$. Therefore

$$\int_{\mathrm{diam}(Q)/4}^{\mathrm{diam}(Q)/2} \iint_Q \beta^2(x,t)\frac{dxdt}{t^{n+1-s}} \lesssim \beta^2(Q)\frac{\mathrm{vol}(Q)}{\mathrm{diam}^{n-s}(Q)} \lesssim \beta^2(Q)\mathrm{diam}(Q)^s.$$

Since the domains of integration on the left are disjoint for distinct (but not necessarily disjoint) dyadic cubes Q, we see that

$$\int_0^\infty \iint_{\mathbb{R}^n} \beta^2(x,t)\frac{dxdt}{t^{n+1-s}} \lesssim \sum_Q \beta^2(Q)\mathrm{diam}(Q)^s.$$

Conversely, if $x \in Q$ and $t \geq 2\mathrm{diam}(Q)$, then $3Q \subset D(x,t)$, so $\beta(x,t) \geq \frac{1}{2}\beta(Q)$. This shows the β^2-integral is also bounded below by a multiple of the β^2-sum.

(B.1) $\Rightarrow$ (B.3): Assume (B.1) holds and that Γ is chord-arc. If $x \in \Gamma$, $0 < t \leq \mathrm{diam}(\Gamma)$, then $\ell(\Gamma \cap D(x,t)) \simeq t$. If $\gamma \subset \Gamma$ is a subarc of length t, then its diameter is at most t and we can choose a dyadic cube Q containing x and so that $2t \leq \mathrm{diam}(Q) \leq 4t$. Then $B(x,t) \subset 3Q$ and so $\widetilde{\beta}(x,t) \lesssim \beta(Q)$, and hence (if ds denotes arclength measure on Γ),

$$\int_{\mathrm{diam}(Q)/4}^{\mathrm{diam}(Q)/2} \int_\gamma \widetilde{\beta}^2(x,t)\frac{dsdt}{t^{2-s}} \lesssim \frac{\mathrm{diam}(Q)\beta^2(Q)}{\mathrm{diam}(Q)^{1-s}} \lesssim \beta^2(Q)\mathrm{diam}(Q)^s.$$

Now divide Γ into dyadic subintervals, $\{\gamma_j\}$ and let Q_j be the dyadic cube associated to γ_j as above. Then

$$\begin{aligned}
\int_0^\infty \iint_\Gamma \widetilde{\beta}^2(x,t)\frac{dsdt}{t^{2-s}} &= \sum_j \int_{\ell(\gamma_j)/2}^{\ell(\gamma_j)/2} \int_{\gamma_j} \widetilde{\beta}^2(x,t)\frac{dsdt}{t^{2-s}} \\
&\simeq \sum_j \beta^2(Q_j)\mathrm{diam}(Q_j)^s \\
&\lesssim \sum_Q \beta^2(Q)\mathrm{diam}(Q)^s,
\end{aligned}$$

where the last line holds if we know that each dyadic Q is only chosen a bounded number of times as a Q_j. But if Q is chosen for γ_j then γ_j hits Q and has length comparable to $\mathrm{diam}(Q)$. By the chord-arc condition,only a bounded number of such arcs can hit Q, for otherwise the arclength of $\Gamma \cap 3Q$ would be too large. This proves the arclength integral is bounded by the sum.

(B.3) $\Rightarrow$ (B.4): For each element Γ_j of the multi-resolution family, choose a dyadic arc $\gamma_k \subset \Gamma$ that hits Γ_j and and has comparable length. Since $\beta(\Gamma_j) \lesssim \widetilde{\beta}(x, 2\mathrm{diam}(\Gamma_j))$,

$$\beta^2(\Gamma_j)\mathrm{diam}(\Gamma_j)^s \lesssim \int_{\mathrm{diam}(\Gamma_j)}^{2\mathrm{diam}(\Gamma_j)} \int_{\gamma_k} \widetilde{\beta}_\Gamma(x,t)\frac{dxdt}{t^{2-s}}.$$

Since each γ_k can be associated to at most $O(1)$ arcs Γ_k (Γ_j can only hit a bounded number of dyadic arcs of comparable size), the multi-resolution sum over the whole family is bounded by the $\widetilde{\beta}$-integral over all Γ and all scales.

(B.4) $\Rightarrow$ (B.2): Suppose Γ is chord-arc, that Q is a dyadic cube, $x \in Q$, and that $\operatorname{diam}(Q) \leq t \leq 2\operatorname{diam}(Q)$. Let $X = \Gamma \cap 9Q$ and let Γ_j be a member of the multi-resolution family containing X and having comparable diameter. Then

$$\beta(x,t) \leq \beta(\Gamma_j) \cdot \frac{\operatorname{crd}(\Gamma_j)}{t} \lesssim \beta(\Gamma_j),$$

since $\operatorname{crd}(\Gamma_j) \simeq t$ by the chord-arc condition. The integral in (B.2) is obtained by summing all the integrals over product sets of the form $Q \times [\operatorname{diam}(Q), 2\operatorname{diam}(Q)]$, and each such integral is bounded by $\beta^2(\Gamma_j)\operatorname{diam}(\Gamma_j)^s$ for the corresponding Γ_j. Since Γ_j can only hit bounded number of dyadic cubes with $\operatorname{diam}(Q) \simeq \operatorname{diam}(\Gamma_j)$, we see that each Γ_j is used only a bounded number of times, hence the sum bounds the integral.

Changing limits of integration: Recall that we have assumed $\operatorname{diam}(\Gamma) = 1$. To see that (B.2) is equivalent to

$$\int_0^{\operatorname{diam}(\Gamma)} \iint_{\mathbb{R}^n} \beta^2(x,t) \frac{dxdt}{t^{n+1-s}} < \infty, \tag{B.5}$$

we simply note that for any $x \in \Gamma$ and $t > \operatorname{diam}(\Gamma)$, that $\beta(x,t) \leq \frac{1}{t}\beta(x,1)$ and that $\beta(x,1) \simeq \beta(y,1)$ for any $x, y \in \Gamma$. Hence

$$\begin{aligned}
\int_1^\infty \iint_{\mathbb{R}^n} \beta^2(x,t) \frac{dxdt}{t^{n+1-s}} &\lesssim \beta^2(x,1) \int_1^\infty \iint_{D(x,2t)} \frac{dxdt}{t^{n+3-s}} \\
&\lesssim \beta^2(x,1) \int_1^\infty t^{-3+s} dt \\
&\lesssim \int_{1/2}^1 \iint_{\mathbb{R}^n} \beta^2(x,t) \frac{dxdt}{t^{n+1-s}} \\
&\lesssim \int_0^1 \iint_{\mathbb{R}^n} \beta^2(x,t) \frac{dxdt}{t^{n+1-s}}
\end{aligned}$$

This is where we use the assumption $s < 2$, so that $-3 + s < -1$ and the integral above converges. Thus truncating the integral cannot convert it from divergent to convergent. A similar argument works for truncating the sum in (B.1) or the integral in (B.3). $\square$

References

[1] Martin Aigner and Günter M. Ziegler. *Proofs from The Book*. Springer, Berlin, sixth edition, 2018. See corrected reprint of the 1998 original [MR1723092], Including illustrations by Karl H. Hofmann.

[2] Sanjeev Arora. Polynomial time approximation schemes for Euclidean traveling salesman and other geometric problems. *J. ACM*, 45(5):753–782, 1998.

[3] Jonas Azzam and Raanan Schul. An analyst's traveling salesman theorem for sets of dimension larger than one. *Math. Ann.*, 370(3-4):1389–1476, 2018.

[4] Matthew Badger, Lisa Naples, and Vyron Vellis. Hölder curves and parameterizations in the Analyst's Traveling Salesman theorem. *Adv. Math.*, 349:564–647, 2019.

[5] Christopher J. Bishop. Weil-Petersson curves, conformal energies, β-numbers, and minimal surfaces. 2020. preprint.

[6] Christopher J. Bishop and Peter W. Jones. Harmonic measure, L^2 estimates and the Schwarzian derivative. *J. Anal. Math.*, 62:77–113, 1994.

[7] Christopher J. Bishop and Yuval Peres. *Fractals in probability and analysis*, volume 162 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, Cambridge, 2017.

[8] K. Borsuk. Drei sätze über die n-dimensionale euklidische sphäre. *Fundamenta Math*, 20:177–190, 1933.

[9] G. David and S. Semmes. Singular integrals and rectifiable sets in $\mathbf{R}^n$: Beyond Lipschitz graphs. *Astérisque*, (193):152, 1991.

[10] Guy C. David and Raanan Schul. The analyst's traveling salesman theorem in graph inverse limits. *Ann. Acad. Sci. Fenn. Math.*, 42(2):649–692, 2017.

[11] Matt Feiszli, Sergey Kushnarev, and Kathryn Leonard. Metric spaces of shapes and applications: compression, curve matching and low-dimensional representation. *Geom. Imaging Comput.*, 1(2):173–221, 2014.

[12] Matt Feiszli and Akil Narayan. Numerical computation of Weil-Peterson geodesics in the universal Teichmüller space. *SIAM J. Imaging Sci.*, 10(3):1322–1345, 2017.

[13] Fausto Ferrari, Bruno Franchi, and Hervé Pajot. The geometric traveling salesman problem in the Heisenberg group. *Rev. Mat. Iberoam.*, 23(2):437–480, 2007.

[14] John B. Garnett and Donald E. Marshall. *Harmonic measure*, volume 2 of *New Mathematical Monographs*. Cambridge University Press, Cambridge, 2008. Reprint of the 2005 original.

[15] Zheng-Xu He. The Euler-Lagrange equation and heat flow for the Möbius energy. *Comm. Pure Appl. Math.*, 53(4):399–431, 2000.

[16] Peter W. Jones. Rectifiable sets and the traveling salesman problem. *Invent. Math.*, 102(1):1–15, 1990.

[17] Jeff Kahn and Gil Kalai. A counterexample to Borsuk's conjecture. *Bull. Amer. Math. Soc. (N.S.)*, 29(1):60–62, 1993.

[18] Daniel A. Klain and Gian-Carlo Rota. *Introduction to geometric probability*. Lezioni Lincee. [Lincei Lectures]. Cambridge University Press, Cambridge, 1997.

[19] Jared Krandel. The traveling salesman theorem for Jordan curves in Hilbert space. 2022. preprint, arXiv:2107.07017v2 [math.CA].

[20] Gilad Lerman. Quantifying curvelike structures of measures by using L_2 Jones quantities. *Comm. Pure Appl. Math.*, 56(9):1294–1365, 2003.

[21] Sean Li and Raanan Schul. The traveling salesman problem in the Heisenberg group: upper bounding curvature. *Trans. Amer. Math. Soc.*, 368(7):4585–4620, 2016.

[22] Sean Li and Raanan Schul. An upper bound for the length of a traveling salesman path in the Heisenberg group. *Rev. Mat. Iberoam.*, 32(2):391–417, 2016.

[23] Joseph S. B. Mitchell. Guillotine subdivisions approximate polygonal subdivisions: a simple polynomial-time approximation scheme for geometric TSP, k-MST, and related problems. *SIAM J. Comput.*, 28(4):1298–1309, 1999.

[24] Jun O'Hara. Energy of a knot. *Topology*, 30(2):241–247, 1991.

[25] Jun O'Hara. Energy functionals of knots. In *Topology Hawaii (Honolulu, HI, 1990)*, pages 201–214. World Sci. Publ., River Edge, NJ, 1992.

[26] Jun O'Hara. Family of energy functionals of knots. *Topology Appl.*, 48(2):147–161, 1992.

[27] Kate Okikiolu. Characterization of subsets of rectifiable curves in $\mathbf{R}^n$. *J. London Math. Soc. (2)*, 46(2):336–348, 1992.

[28] Hervé Pajot. *Analytic capacity, rectifiability, Menger curvature and the Cauchy integral*, volume 1799 of *Lecture Notes in Mathematics*. Springer-Verlag, Berlin, 2002.

[29] Steffen Rohde and Yilin Wang. The Loewner energy of loops and regularity of driving functions. *Inter. Math. Res. Notes.*, 04 201p. arXiv:1601.05297v2 [math.CV].

[30] Luis A. Santaló. *Integral geometry and geometric probability.* Addison-Wesley Publishing Co., Reading, Mass.-London-Amsterdam, 1976. With a foreword by Mark Kac, Encyclopedia of Mathematics and its Applications, Vol. 1.

[31] Raanan Schul. Subsets of rectifiable curves in Hilbert space—the analyst's TSP. *J. Anal. Math.*, 103:331–375, 2007.

[32] E. Sharon and D. Mumford. 2D-Shape analysis using conformal mapping. *Int. J. Comput. Vis.*, 70:55–75, 2006.

[33] Leon A. Takhtajan and Lee-Peng Teo. Weil-Petersson metric on the universal Teichmüller space. *Mem. Amer. Math. Soc.*, 183(861):viii+119, 2006.

[34] Xavier Tolsa. Uniform rectifiability, Calderón-Zygmund operators with odd kernel, and quasiorthogonality. *Proc. Lond. Math. Soc. (3)*, 98(2):393–426, 2009.

[35] Yilin Wang. A note on Loewner energy, conformal restriction and Werner's measure on self-avoiding loops. 2018. preprint, arXiv:1810.04578v1 [math.CV], to appear, Annales de l'Institut Fourier.

[36] Yilin Wang. The energy of a deterministic Loewner chain: reversibility and interpretation via SLE_{0+}. *J. Eur. Math. Soc. (JEMS)*, 21(7):1915–1941, 2019.

[37] Yilin Wang. Equivalent descriptions of the Loewner energy. *Invent. Math.*, 218(2):573–621, 2019.

C.J. Bishop, Mathematics Department, Stony Brook University, Stony Brook, NY 11794-3651

Email address: `bishop@math.stonybrook.edu`