MAT 638, FALL 2020, Stony Brook University

TOPICS IN REAL ANALYSIS

THE TRAVELING SALESMAN THEOREM FOR JORDAN CURVES

Christopher Bishop, Stony Brook

1. Introduction and overview

Jones's theorem in [3] says that the shortest curve Γ containing $E \subset \mathbb{R}^2$ has length

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

Actually, [3] states that for any $\delta > 0$ and $E \subset \mathbb{R}^n$, $2 \leq n < \infty$,

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

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 $\operatorname{diam}(\Gamma) = \sqrt{1+\beta^2} = 1 + O(\beta^2)$. 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* $\mathbb{R}^n$*,*

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

where the sum is over all dyadic cubes.

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

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

The opposite direction is less obvious, but also holds:

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

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

where the sum is over all dyadic cubes.

Plan: For E a set and Γ a Jordan curve we will prove:

$$\ell(\Gamma) - \mathrm{diam}(\Gamma) \lesssim \sum_Q \beta_\Gamma^2(Q)\mathrm{diam}(Q), \tag{1.6}$$

$$\ell(E) \leq (1+\delta)\mathrm{diam}(E) + C\sum_Q \beta_\Gamma^2(Q)\mathrm{diam}(Q), \tag{1.7}$$

$$\ell(\Gamma) - \mathrm{diam}(\Gamma) \gtrsim \sum_Q \beta_\Gamma^2(Q)\mathrm{diam}(Q), \tag{1.8}$$

$$\ell(\Gamma) - \mathrm{crd}(\Gamma) \lesssim \sum_Q \beta_\Gamma^2(Q)\mathrm{diam}(Q), \tag{1.9}$$

First we prove the inequality

$$\ell(\Gamma) - \mathrm{diam}(\Gamma) \lesssim \sum_{Q} \beta_{\Gamma}^{2}(Q)\mathrm{diam}(Q), \tag{1.10}$$

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.

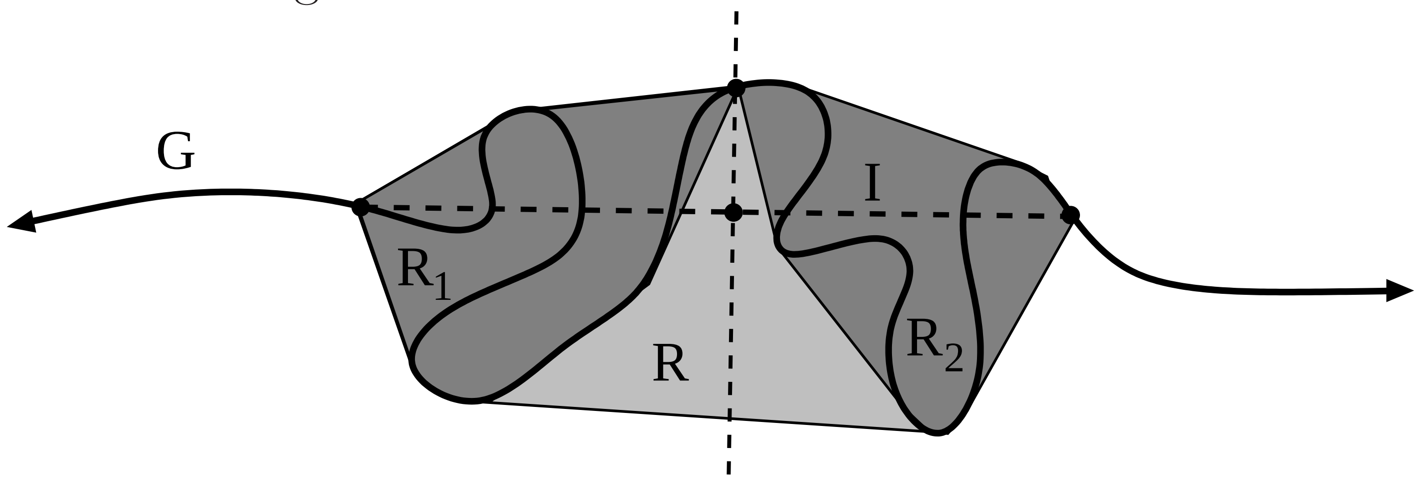

We call this process splitting R. The collection $\mathcal{R}_{n+1}$ is obtained by splitting very element of $\mathcal{R}_n$ in this way. Thus $\mathcal{R}_{n+1}$ has twice as many elements as $\mathcal{R}_n$.

We will think of these elements as the nth generation of a binary tree whose root is R_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).

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 first supremum is 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 1.3. *If R is split into R_1, R_2 as above, then*

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

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 $\operatorname{diam}(R_1) \leq \operatorname{diam}(W) = \frac{1}{2}\operatorname{diam}(R) + O(\beta^2(R)\operatorname{diam}(R))$. Similarly for R_2, and adding the estimates proves the lemma. $\square$

Lemma 1.4. *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 $M \leq C(n) + 1$, as desired. $\square$

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 [2] 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 [4] who gave examples of sets requiring $\geq (1.1)^{\sqrt{n}}$ 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 1.3, induction and $\mathrm{diam}(\Gamma_0) = \mathrm{diam}(\Gamma)$, we get

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

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 $\lim_{\delta\searrow 0}\inf\sum_j \operatorname{diam}(X_J)$, where the infimum is over all coverings of Γ by set of diameter less than δ.

By Lemma 1.4 our collections $\mathcal{R}_n$ form such coverings and hence

$$\ell(\Gamma) = \mathcal{H}^1(\Gamma) \le \limsup_{n\to\infty} \sum_{R\in\mathcal{R}_n} \operatorname{diam}(R). \tag{1.11}$$

Therefore,

$$\ell(\Gamma) \;\leq\; \operatorname{diam}(\Gamma) + O\left(\sum_{n=0}^{\infty} \sum_{R\in\mathcal{R}_n} \beta^2(R)\operatorname{diam}(R)\right). \tag{1.12}$$

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 $\operatorname{diam}(R) \leq \operatorname{diam}(Q) \leq 2\cdot\operatorname{diam}(Q)$. Then $R \subset 3Q$ 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 1.5. *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 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 1.4 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 $\mathrm{diam}(R) \geq 4$ and $R \cap B_\epsilon \neq \emptyset$. We call these sets 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 $\operatorname{diam}(R_0)$ from the center of B_k. By definition, R_0 contains the convex hull of its endpoints and the k-ball B_k. Since $\operatorname{diam}(I) = \operatorname{diam}(R_0) \geq \operatorname{diam}(R_1)$, 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.

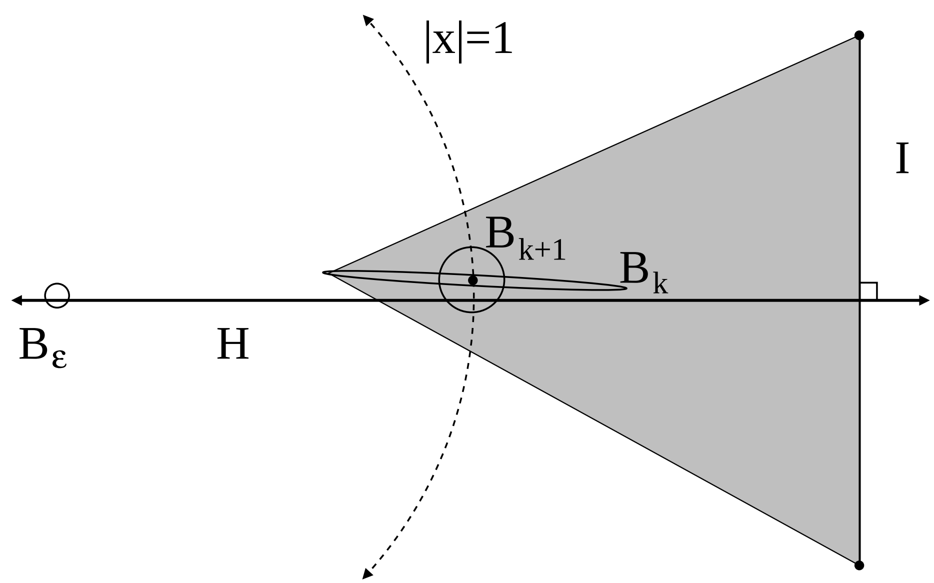

Now choose a point x on the unit sphere in $\mathbb{R}^n$ and consider all the admissible descendents whose centers are within δ of this point. These sets form the leaves of a finite subtree of T, where the only vertices of degree 3 are smallest common ancestors of some subcollection of the sets.

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. $\square$

This completes the proof of the upper bound.

TST upper bound for general sets: We have done most of the work needed to prove the

$$\ell(\Gamma) \leq (1+\delta)\mathrm{diam}(E) + C(\delta) \sum_Q \beta_E^2(Q)\mathrm{diam}(Q) \tag{1.13}$$

for a general set $E \subset \mathbb{R}^n$.

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 diagonal curve 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 to to Lemma 1.3. It becomes

Lemma 1.6. *If R is split into R_1, R_2 sets and a segment S as above, then*
$$\mathrm{diam}(R_1) + \mathrm{diam}(R_2) + (1-\delta)\ell(S) \leq \mathrm{diam}(R) + O(\frac{1}{\delta})\beta^2(R)\mathrm{diam}(R).$$

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 1.3 is the same as before.

For the second case, first note that if $\beta(R) \geq \delta/20$, then

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

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

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

This proves the lemma. $\square$

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 (1.12) by

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

Dividing both sides by $(1-\delta)$ proves (1.2).

This implies the smallest connected set Γ containing a set E satisfies

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

Later we prove that the lower bound

$$C_1 \sum_Q \beta_E^2(Q)\mathrm{diam}(Q) \leq\leq \ell(\Gamma) - \mathrm{diam}(E)$$

holds for any curve Γ containing E.

Next we consider the lower bound:

$$\ell(\Gamma) - \operatorname{diam}(\Gamma) \gtrsim \sum_Q \beta_\Gamma^2(Q)\operatorname{diam}(Q). \tag{1.14}$$

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

$$\ell(\Gamma) - \frac{1}{2}\operatorname{prm}(\Gamma) \gtrsim \sum_Q \beta_\Gamma^2(Q)\operatorname{diam}(Q), \tag{1.15}$$

where $\operatorname{prm}(\Gamma) = \ell(\partial(\operatorname{ch}(\Gamma)))$ 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 (1.15) implies (1.8). In higher dimensions, the perimeter is replaced by a quantity called the "mean width" of Γ, defined below.

Estimate (1.8) uses ideas from integral geometry. For the following facts, see [7].

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 [7]. 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

$$\mu(\{H : H \cap \Gamma \neq \emptyset\}) \geq \mu(H : H \cap S \neq \emptyset\}) = \mathrm{crd}(\Gamma). \tag{1.16}$$

The leftmost quantity in (1.16) 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., [5], [7].

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, [6]). 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 1.7. $n(H,\Gamma) - 1 \gtrsim N(H,\Gamma)$ *whenever* $n(L,\Gamma) > 0$.

Lemma 1.8. *1.7 a: Command not found.* $n(H,\Gamma)-1 \gtrsim N(H,\Gamma)$ *whenever* $n(L,\Gamma) > 0$.

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.

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 $\operatorname{diam}(Q) \simeq |z - w|$, 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 $\#(\mathcal{C}) = \#(E)$. We claim that G is a forest, i.e., all its connected components are trees, e.g., G has no cycles. If this holds, 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. Now 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. If not, then look at the connecting edges in groups of comparable size in decreasing order: there is one element of the first group, at most two in the second group, at most four in the third group, and so on.

Thus the total length of all the edges in the path 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. $\square$

Continuing with the proof of (1.8), note that by (1.16)

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

To complete the proof we need

Lemma 1.9. *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 1.9 we will need some preliminary facts.

Lemma 1.10. *Suppose I is the unit segment between 0 and 1 on the x_1-axis, and suppose J is another unit length segment with at least one endpoint inside $\{|x| < 100\}$ and 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 are close to the $(n-1)$-unit-sphere $\mathbb{S}^{n-1}$ 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 parameter by increasing its distance from $X_1^\perp$ by factor of at most $O(b)$, and thus changes the $(n-1)$-volume by a factor of $O(b)$ near S^{n-1}. 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 1.11. *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 $k = 1, \ldots, n$, 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 1.9. The idea is simple: by Lemma 1.10 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 of $\Gamma \cap 3Q$ whose chords are β far from lying one the same line. Moreover, we need these chords to have lengths comparable to $\mathrm{diam}(Q)$ and for one to lie 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 $O(\beta)$, that we shall show 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 $v \in \Gamma \cap \left(2\frac{23}{24}Q \setminus 2\frac{1}{24}Q\right) \setminus W_0$.

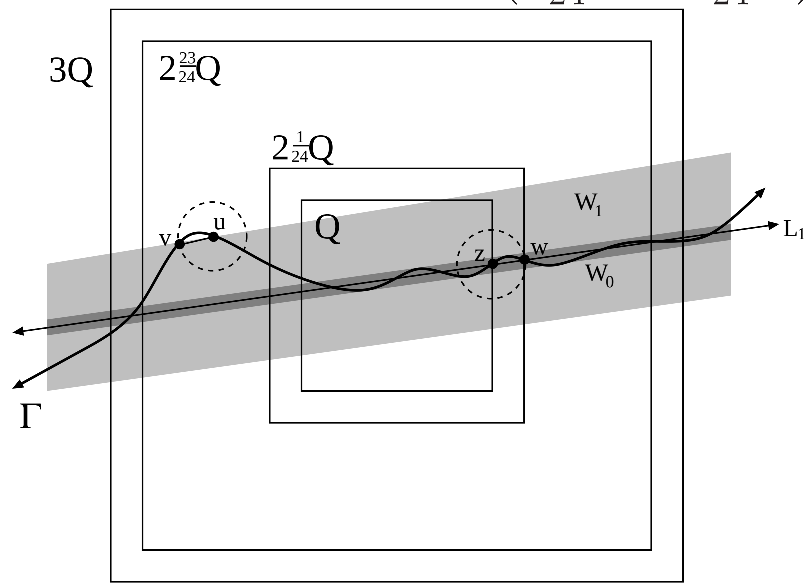

Since Γ is path connected and has diameter $\geq \operatorname{diam}(Q)$, v can be connected to a point $u \in \Gamma$ with $|u-v| = \frac{1}{24}$ by a sub-arc $\gamma_2 \subset \Gamma$ that stays inside the ball $B_2 = B(v, \frac{1}{24})$. Let $S_2 = [u, v]$ be the segment connecting these two points 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 1.10. Thus $\mu(S(Q,\Gamma)) \gtrsim$ as well and the lemma is satisfied with $Q^* = Q$.

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 $B_3 = B(p, \frac{1}{12}) \subset 2\frac{1}{3}Q$.

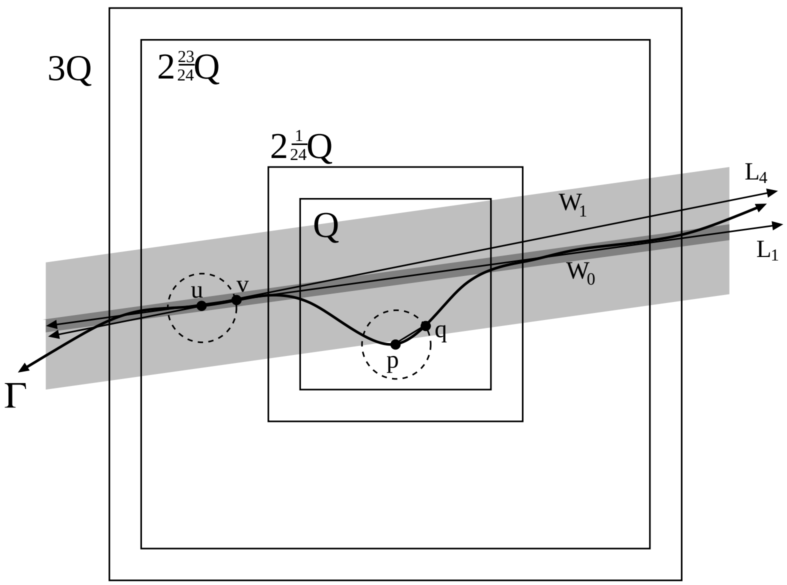

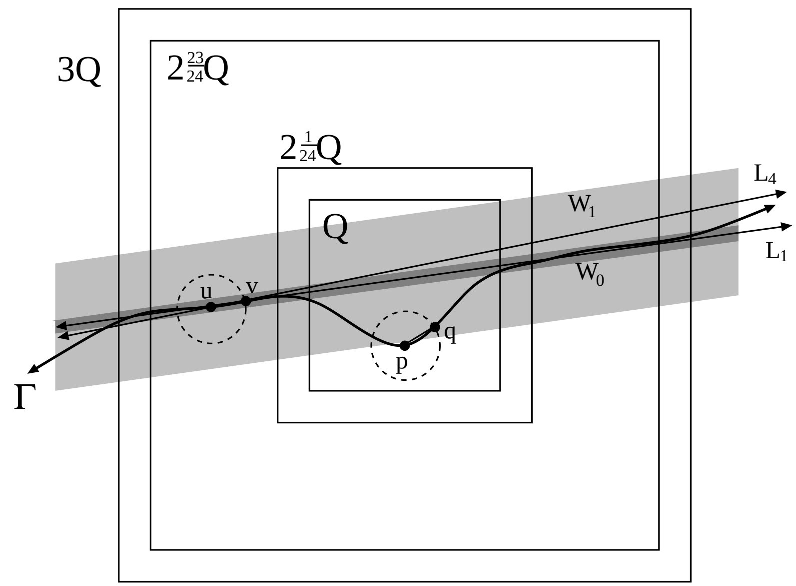

Let $S_3 = [p, q]$. As before, any line that hits S_3 must also hit γ_3. Choose an element $Q^* \in \mathcal{D}^*(Q)$ that is the same size as Q and translated by at most $\frac{1}{3}\ell(Q)$ in each coordinate direction (possibly Q itself) and so that $B_3 \subset 1\frac{2}{3}Q^*$ (it is easy to check there is at least one such cube Q^* by considering the projections onto each coordinate).

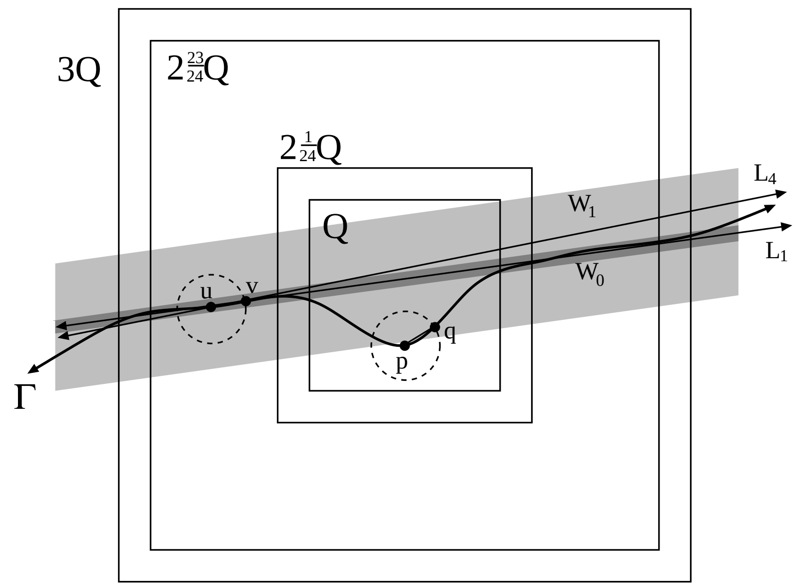

By Lemma 1.11, Γ must contain a point $u \in (3Q \cap 3Q^*) \setminus (2Q \cup 2Q^*)$ that is at least distance $\geq \frac{1}{12}$ from the boundaries of both $2Q$ and $3Q$. Therefore the ball $B_4 = B(u, \frac{1}{24})$ is inside of $2\frac{23}{24}Q \setminus 2\frac{1}{24}Q$. As before, we can find a radius S_4 of B_4 so that any line that hits S_4 also hits a sub-arc $\gamma_4 \subset \Gamma \cap B_4$.

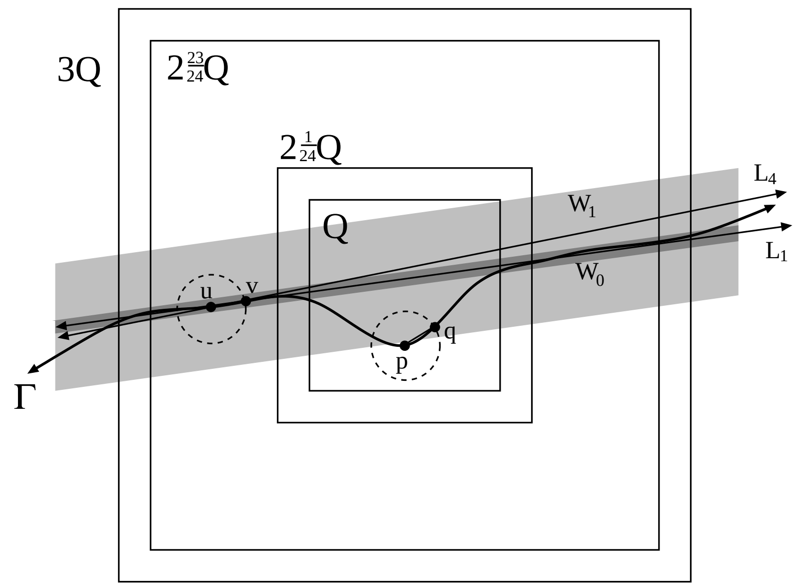

Since γ_4 lies inside the very thin cylinder w_0, the line L_4 containing S_4 is almost parallel to L_1 (the axis of W_0) and so $L_4 \cap 3Q$ lies inside a cylinder of radius $\beta/4$ around L_1. Since p is outside the larger cylinder W_1 we see that S_3 and S_4 satisfy Lemma 1.10, hence the measure of the set of lines that hit both S_4 and S_3 is $\gtrsim \beta^2$ by Lemma 1.10. Thus $\mu(S(Q^*, \Gamma)) \gtrsim \beta^2$, as desired. $\square$

This completes the proof of (1.8), and hence of Theorem 1.1:

$$\ell(\Gamma) - \mathrm{diam}(\Gamma) \simeq \sum_{Q} \beta^2(Q)\mathrm{diam}(Q).$$

Next we prove

Theorem 1.12. *For any Jordan arc* $\Gamma \subset \mathbb{R}^n$,

$$\ell(\Gamma) - \mathrm{crd}(\Gamma) \simeq \sum_Q \beta_\Gamma^2(Q)\mathrm{diam}(Q), \tag{1.17}$$

where the sum is over all dyadic cubes.

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

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 $1 \leq \mathrm{diam}(Q_0) \leq 2$, 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

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

(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 $w = y$). By rotating and rescaling, we may assume that $x = 1, y = -1$ on the x_1-axis.

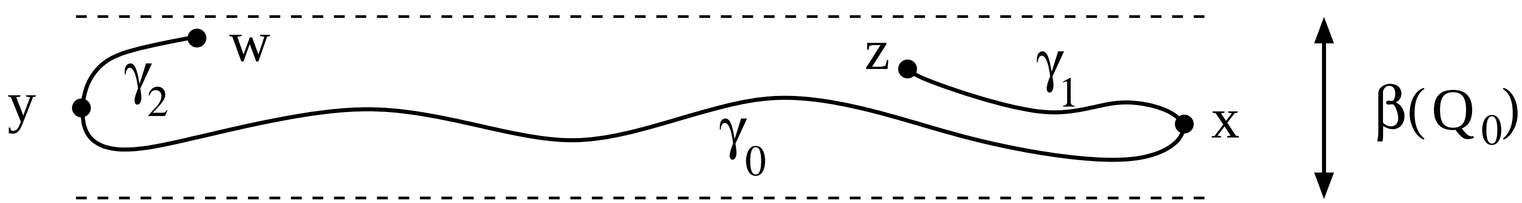

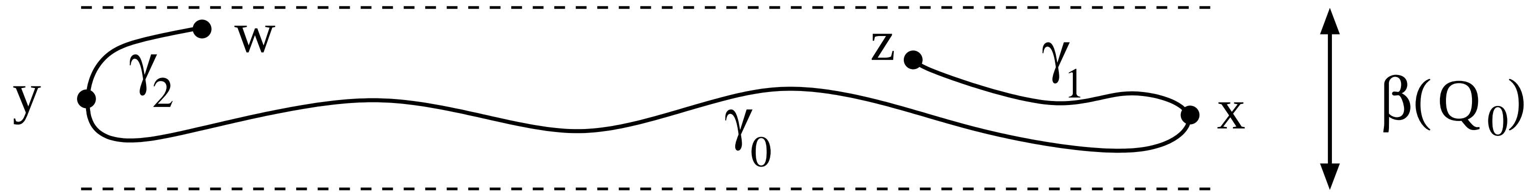

Note that

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

and hence using Theorem 1.1 we get

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

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 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 $Q_1, \ldots, Q_k$ 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 $j \in \{0, 1, \ldots, k\}$. Let L_j be a best line in the definition of $\beta_\Gamma(Q_j)$.

Case 1: Assume that for some $j \in \{1, \ldots, k\}$, 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

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

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

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

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 1.13. *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* $\mathrm{diam}(Q) \leq \mathrm{diam}(\gamma_1)/10$ *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* $1/10$ *passing through p.*

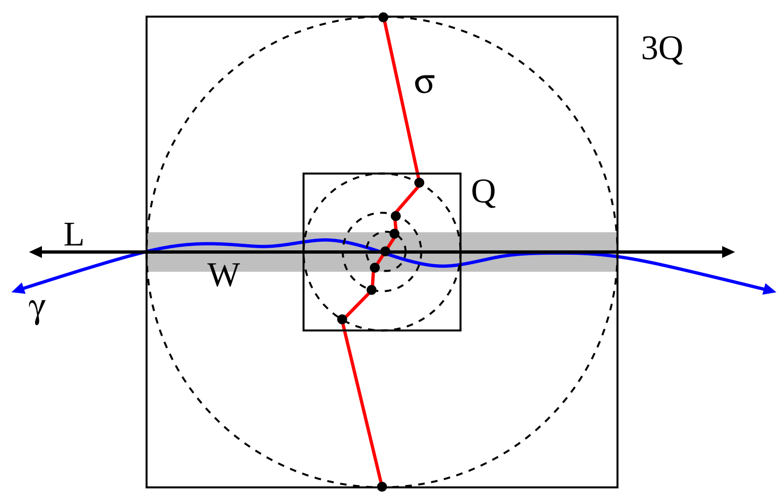

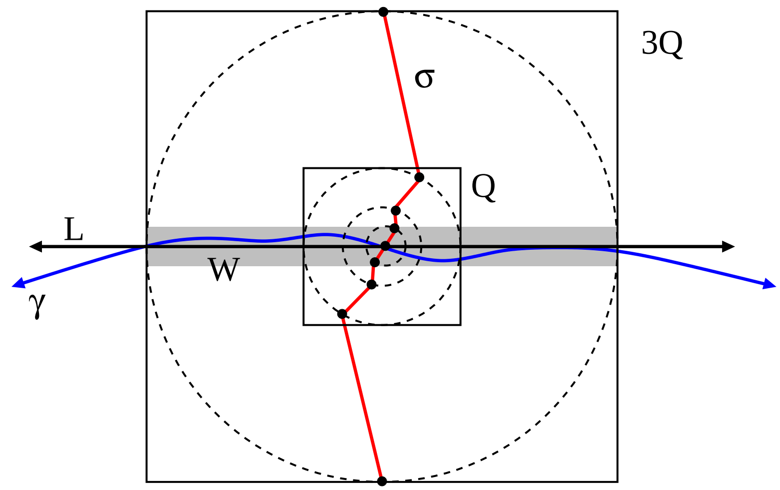

Proof. Note that because p is not in E, that γ has small β-number for Q and for every dyadic subcube of Q that contains p. Using this, we claim we can construct a surface σ so that

(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 $p \in Q' \subset Q$.

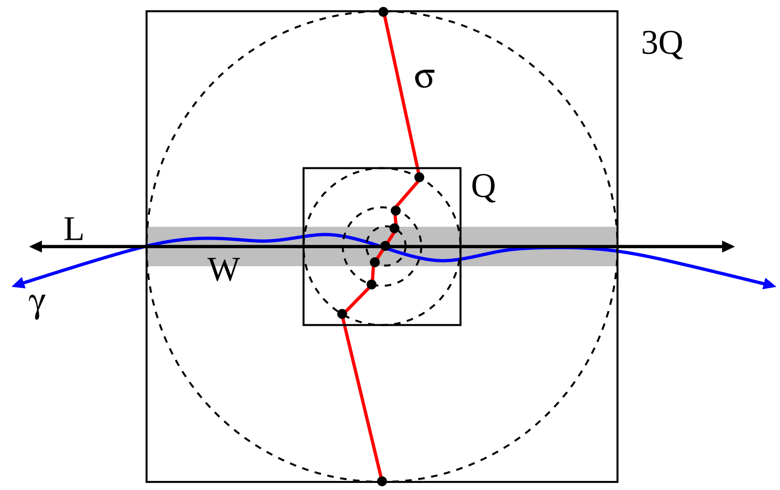

To do this, choose a $(n-1)$-sphere of the n-sphere of radius $t = \mathrm{diam}(Q)2^{-n}$ around p that is nearly orthogonal to the optimal line L_n passing through p for the definition of $\beta(p, 2^{-n})$ and then connecting these $(n-1)$-spheres by a surface (e.g., project both onto the $(n-1)$-plane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

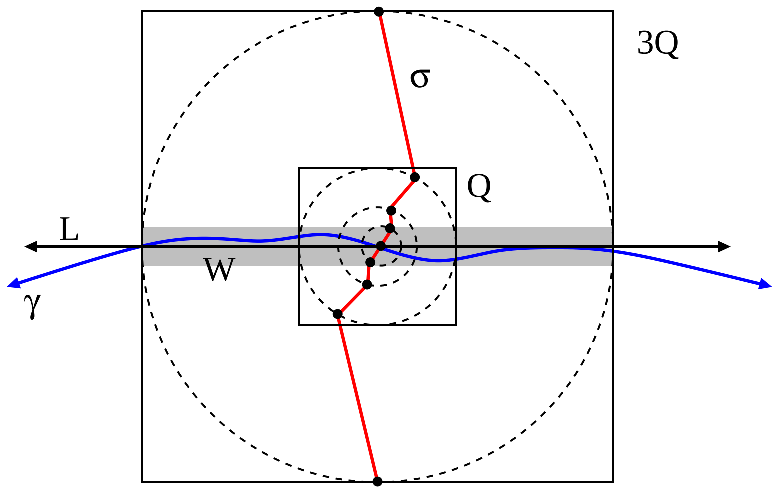

If p is a tangent point of γ_1, then σ has a tangent plane at p that is perpendicular to γ_1's tangent direction. From this we see that γ_1 crosses σ, i.e., it hits both components of $3Q \setminus \sigma$. Since γ_1 only hits σ once, it must leave $3Q$ through a different component of $\partial(3Q) \cap W$ than it entered through. This implies Lemma 1.13. $\square$

Lemma 1.14. $\ell(E) = \ell(\gamma_1')$.

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 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 $2d < \mathrm{diam}(Q_p) \le 4d$. Because $\mathrm{diam}(Q_p) \le \mathrm{diam}(\gamma_1)/10$, the argument in the previous paragraph applies, and γ_1 must cross Q_p inside a cylinder S of width $\beta_0 \mathrm{diam}(Q_p)$.

Moreover, since $\mathrm{diam}(Q_p) > 2d$, 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 $\beta_0 \, \mathrm{diam}(Q_p)$ 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 1.14 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 the plane, 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 $O(\beta^2(Q)\operatorname{diam}(Q))$.

Now we use the traveling salesman theorem 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.

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

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

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

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

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

[7] 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.