Catching sets with quasicircles

Paul MacManus

Abstract. We show how certain geometric conditions on a planar set imply that the set must lie on a quasicircle, and we give a geometric characterization of all subsets of the plane that are quasiconformally equivalent to the usual Cantor middle-third set.

0. Introduction.

Theorem 1. *For a subset E of $\mathbb{C}$ the following are equivalent*:

   i) *E has empty interior and uniform complement.*

   ii) *E is uniformly disconnected.*

   iii) *E is quasiconformally equivalent to a porous subset of $\mathbb{R}$.*

   The various constants depend only on each other.

One immediate consequence of this theorem is that any set E satisfying either i) or ii) lies on a quasicircle. Indeed, the main part of the proof consists of demonstrating this fact.

An NUD set is a compact set having no interior and whose complement is a uniform domain. Väisälä considered the family of NUD sets in [V], where he showed that if such a set is removed from a uniform domain, then the domain that remains is still uniform. Hence the name: NUD stands for nullsets for uniform domains. As a corollary of Theorem 1, we obtain the following characterisation of NUD sets in the plane.

Corollary 2. *E is an* NUD *set in* $\mathbb{C}$ *if and only if E is quasiconformally equivalent to a compact, porous subset of* $\mathbb{R}$.

We can also characterize those sets that are quasiconformally equivalent to the usual Cantor middle-third set.

Theorem 3. *For a compact set K in* $\mathbb{C}$ *whose interior is empty the following are equivalent*:

i) *K is uniformly perfect and has uniform complement.*

ii) *K is both uniformly perfect and uniformly disconnected.*

iii) *K is quasiconformally equivalent to the usual Cantor middle-third set.*

The various constants depend only on each other.

A result of David and Semmes [DS] says basically that any uniformly disconnected, uniformly perfect, compact metric space can be mapped quasisymmetrically to the middle-third set. In particular, if K satisfies ii) above then there is a quasisymmetric map from K to the Cantor set. The preceding theorem extends this by showing that the map can actually be taken to be a quasiconformal map of $\mathbb{C}$.

1. Preliminaries.

By a quasiconformal map of $\mathbb{C}$ we mean a quasiconformal map from $\mathbb{C}$ onto itself. Such maps can also be viewed as maps of $\overline{\mathbb{C}}$ onto itself that fix infinity. Two subsets of $\mathbb{C}$ are said to be quasiconformally equivalent if there is a quasiconformal map of $\mathbb{C}$ that maps one onto the other. A quasiconformal arc is the image of a closed sub-interval of $\mathbb{R}$ under a quasiconformal map of $\mathbb{C}$. If this map is L-quasiconformal, then we will say that the quasiconformal arc is an L-quasiconformal arc.

We call a set A *c-uniform* $(1 \leq c)$ under the following condition: A contains at least two points and for each pair of distinct points a and b of A, there exists a continuum F containing a and b such that $\operatorname{diam} F \leq c\,|a - b|$ and such that

$$\bigcup_{z \in F \setminus \{a,b\}} B(z, c^{-1}\, r(z)) \subset A\,,$$

where $r(z) = \min\left\{|z-a|, |z-b|\right\}$. A set is c-uniform if and only if there is a c-uniform domain D for which $D \subseteq A \subseteq \overline{D}$. See [V] for a proof of this last remark and for alternative definitions of uniformity.

Let $1 \leq m$. A subset E of $\mathbb{C}$ that contains at least two points is said to be m-*uniformly perfect* if at each point z_0 of E the closed annulus
$$A = \left\{ z : \ \frac{r}{m} \leq |z - z_0| \leq r \right\}$$
has non-empty intersection with E whenever $0 < r < \operatorname{diam} E$.

A subset E of $\mathbb{R}$ is said to be λ-*porous* $(1 \leq \lambda)$ if every interval I centred on E contains an interval of length $|I|/\lambda$ that lies in $\mathbb{R} \setminus E$.

We say that E is τ-*uniformly disconnected* $(\tau \geq 1)$ if for each $x \in E$ and each $r > 0$ we can find a subset A of E containing x, of diamter no more than r, and for which $d(A, E \setminus A) \geq \tau^{-1} r$. This concept was introduced recently in [DS].

$M, M_0, \ldots$ and $\varepsilon, \varepsilon_0, \ldots$ will denote constants that depend only on the relevant data (*e.g.* uniformity constants) associated to the set in question; the former are used for constants that are at least 1, and the latter are used for constants that are less than 1. The same symbol may be used to denote different constants. When we write $A \sim B$, we mean that the ratio of A to B is bounded above and below by a constant that depends, once again, only on the relevant data.

We are going to use a result from [M] on building quasiconformal arcs. In order to state this result we need to introduce the idea of a *chain*. A standard rectangle R is a closed rectangle whose major axis lies on the real line. Let $\{R_i\}_{i=1}^N$ be a family of disjoint standard rectangles of height h with each R_i at least a distance h to the left of R_{i+1}. Take I to be a closed interval in $\mathbb{R}$ whose left endpoint is at least a distance h to the left of R_1 and whose right endpoint is at least a distance h to the right of R_N. The union of I with the R_i is called a standard h-chain. Each R_i is referred to as a rectangle of the chain. The closed intervals joining R_i to R_{i+1}, along with the two closed intervals joining the endpoints of I to the nearest rectangle, are called the links of the chain, and the points where the links meet the rectangles are called the weld points of the chain. An (M, h)-chain is any M bi-Lipschitz image of a standard h-chain. The various parts of an (M, h)-chain are given the same names as their pre-images in the standard chain.

These chains can be used to build quasiconformal arcs. Assume that for each n we have a family $\mathcal{F}_n$ of (M, h_n)-chains with the following properties:

i) $\mathcal{F}_1$ contains only one chain.

ii) Each chain in $\mathcal{F}_{n+1}$ is contained in one of the rectangles of some chain in $\mathcal{F}_n$, and the endpoints of a chain in $\mathcal{F}_{n+1}$ are the same as the weld points of the associated rectangle.

iii) Each rectangle of each chain in $\mathcal{F}_n$ contains exactly one chain from $\mathcal{F}_{n+1}$.

iv) The h_n converge to zero.

Take T_n to be the union of every chain in $\mathcal{F}_n$ with the links of every chain in $\mathcal{F}_k$ for $1 \leq k \leq n-1$. Each T_n is a continuum and $T_{n+1} \subseteq T_n$.

Theorem. *Under the assumptions* i)-iv), $\Gamma_0 = \bigcap_{n=1}^{\infty} T_n$ *is a* $C(M)$ *quasiconformal arc.*

This is [M, Theorem 5.1] and it can be read independently of the rest of that paper.

2. Proofs of the Theorems.

We start with Theorem 1. It is easily confirmed that any quasiconformal map of $\mathbb{C}$ preserves both uniform domains and the property of being uniformly disconnected. It is also easy to check that a subset of the line is porous (as a subset of the line) if and only if its complement in the plane is uniform, and if and only if it is uniformly disconnected. Thus, iii) implies both i) and ii). Furthermore, in order to prove the opposite implications it suffices to show that E lies on a quasiconformal arc.

$\mathcal{G}_\delta$ is the square grid whose vertices are the points $(m\,\delta, n\,\delta)$ where m and n are any integers, and Σ_δ is the associated family of (closed) squares. A useful way of thickening up a set is the following. Let W be a bounded subset of the plane. W^δ is the union of the elements of Σ_δ that intersect W. Let $\mathcal{T}_\delta(W) = (W^{4\delta})^\delta$. The following facts are easily confirmed.

Lemma 2.1. *If W is a bounded subset of the plane, then the boundary of $\mathcal{T}_\delta(W)$ is a finite, disjoint union of Jordan curves, each of which is a subset of $\mathcal{G}_\delta$. Furthermore, the distance from any boundary point of $\mathcal{T}_\delta(W)$ to W is less than $8\,\delta$ and greater than δ.*

The next lemma will allow us to break the set up into manageable chunks at all scales.

Lemma 2.2. *Suppose that E satisfies* i) *or* ii) *in Theorem 1. Then for any $x \in E$ and any positive λ, there is a loop γ separating x from infinity for which* $\operatorname{diam} \gamma \sim d(E, \gamma) \sim \lambda$.

PROOF. Let us first assume that E satisfies i). Fix $x \in E$ and $\lambda > 0$. Set $r = \lambda/4$. Choose points $\{y_1, y_2, \ldots, y_n\}$ that are equally spaced with distance $\varepsilon\, r$ on the circle C_r of radius r centred at x. We will see as we go along how small an ε we need to choose. We will use the convention that $y_{n+1} = y_1$.

For each y_i there is $z_i \in E^c$ whose distance to E is at least $M^{-1}\varepsilon\, r$ and whose distance to y_i is at most $\varepsilon\, r$. There is a path γ_i joining z_i to z_{i+1} whose diameter is at most $M\varepsilon\, r$ and whose distance to E is at least $M^{-1}\varepsilon\, r$. Let γ denote the loop (*i.e.*, closed curve) obtained by joining up the γ_i in the obvious way. The next few statements hold for small enough ε. The diameter of γ is at most $3\, r$, which is less than λ, and the distance from γ to E is least $M^{-1}\varepsilon\, r = M_1^{-1}\lambda$. Furthermore, each γ_i is freely homotopic to the segment $[y_i, y_{i+1}]$ in $\mathbb{C} \setminus \{x\}$, and so γ is freely homotopic in $\mathbb{C} \setminus \{x\}$ to the circle C_r. As a result, γ must separate x from infinity.

Now let us assume that E satisfies ii). Once again fix $x \in E$ and $\lambda > 0$. There is a subset A of E containing x, of diameter no more than λ, and for which $d(A, E \setminus A) \geq \tau^{-1}\lambda$. Let $\delta = (20\,\tau)^{-1}\lambda$, and set $X = \mathcal{T}_\delta(A)$. Then X contains x, has diameter comparable to λ, and all of its boundary points are at least a distance δ, and no more than $8\,\delta$, from A. This information about the boundary combined with the estimate on the distance between A and $E \setminus A$ implies that the distance to E of every point on the boundary is comparable to λ. The boundary of X is a finite, disjoint, collection of Jordan curves. One of these must enclose the point x and this is the loop we seek.

It is clear that if E satisfies ii), then any subset of E also satisfies ii) with the same constant. This also holds for i), and depends on the simple observation that the complement of E is uniform and dense in $\mathbb{C}$. These facts combined with a standard limiting argument imply that it suffices to prove Theorem 1 for sets consisting of a finite number of points. We will assume from here on that that E contains only a finite number of points. This is by no means necessary but it means that we

do not have to worry about possible technical difficulties.

Set $\mathcal{D}_\delta$ to be the set of closures of all Jordan domains whose boundaries are both subsets of $\mathcal{G}_\delta$ and at least a distance δ from E.

Lemma 2.3. *Suppose that E satisfies* i) *or* ii) *in Theorem 1. Then for any positive δ there is a finite, disjoint collection $\{\Delta_1, \Delta_2, \dots\}$ of elements of $\mathcal{D}_\delta$ with the properties listed below.*

   i) *$E \cap \Delta_k \neq \varnothing$ for all k, and $E \subseteq \bigcup \Delta_k$.*

   ii) *$\operatorname{diam} \Delta_k \leq M_0 \, \delta$, for all k.*

PROOF. Fix $\delta > 0$. It follows from Lemma 2.2 that for each $x \in E$ there is a loop γ_x of diameter comparable to δ that separates x from infinity and whose distance to E is at least $10 \, \delta$. The diameter of C_x, the component of γ_x^c that contains x, is no more than $M \, \delta$.

Let $G = \bigcup_{x \in E} \gamma_x$. The distance from G to E is at least $10 \, \delta$. The boundary of $\mathcal{T}_\delta(G)$ is a finite, disjoint union of Jordan curves, each of which lies on $\mathcal{G}_\delta$ and is at least a distance δ from E. As a result, the set

$$\mathcal{C} = \{\overline{U} : \ U \text{ is the bounded domain determined}$$

$$\text{by some component of } \partial \mathcal{T}_\delta(G)\}$$

is contained in $\mathcal{D}_\delta$, and any two elements of $\mathcal{C}$ are either disjoint or one is contained in the other.

Now, the boundary of any bounded component V of $(\mathcal{T}_\delta(G))^c$ is a union of components of $\partial \mathcal{T}_\delta(G)$ and so there is a unique element of $\mathcal{C}$ that contains V and whose boundary is contained in the boundary of V.

For each $x \in E$, define V_x to be the component of $(\mathcal{T}_\delta(G))^c$ that contains x, and Δ_x to be the element of $\mathcal{C}$ that corresponds to V_x. Let $\{\Delta_1, \Delta_2, \dots\}$ be the maximal elements among the Δ_x. It is clear that they have the required properties, except perhaps ii). Each V_x must be a subset of C_x, as $\gamma_x \subseteq \mathcal{T}_\delta(G)$. Thus the diameter of V_x is less than $M\delta$. Consequently, we have that $\operatorname{diam} \Delta_x = \operatorname{diam} V_x \leq M \, \delta$, which is ii).

We are now ready to build the quasiconformal arc containing the set E. For convenience, we will assume that E lies in the unit disc. Let ε be a suitably small constant. Abbreviate $\mathcal{D}_{\varepsilon^n}$ to $\mathcal{D}_n$. For $n \geq 1$,

the subset of $\mathcal{D}_n$ obtained in Lemma 2.3 by setting $\delta = \varepsilon^n$ is called $\mathcal{S}_n$. We define $\mathcal{S}_0$ to be the square of sidelength 2 that is centred at the origin. Every element of $\mathcal{S}_{n+1}$ is a subset of some element of $\mathcal{S}_n$, and every element of $\mathcal{S}_n$ contains at least one element of $\mathcal{S}_{n+1}$. For any $S \in \mathcal{S}_n$, we have $\operatorname{diam} S \sim \varepsilon^n$. Furthermore, if $\widehat{S} \in \mathcal{S}_{n+1}$ lies in S, then $d(\partial S, \widehat{S}) \sim \varepsilon^n$. If we set E_n to be the union of the elements of $\mathcal{S}_n$, then $E = \bigcap_n E_n$. We now have nicely nested coverings of E, whose infinite intersection is E. All that remains in the proof of Theorem 1 is to use the $\mathcal{S}_n$ to construct families of chains satisfying conditions i)-iv) on pages 3-4. For all n, and for each $S \in \mathcal{S}_n$, choose two boundary points whose distance apart is the diameter of S. We will refer to these as the distinguished points of S.

Lemma 2.4. *Suppose that $S \in \mathcal{S}_{n-1}$. Denote the elements of $\mathcal{S}_n$ that lie in S by $\mathcal{C}(S)$. Then S contains an (M, ε^n)-chain whose endpoints are the distinguished points of S, whose rectangles are precisely $\mathcal{C}(S)$, and whose weld points are the distinguished points of the elements of $\mathcal{C}(S)$.*

This result is just a special case of [M, Corollary 4.2], but since this case is particularly simple we briefly outline the proof here.

SKETCH OF PROOF. S and the elements of $\mathcal{C}(S)$ are all elements of $\mathcal{D}_n$. Their boundaries are disjoint and each has at most M_1 edges. Take Q to be either S^c or one of the elements of $\mathcal{C}(S)$ and define $\hat{Q}$ to be the $\varepsilon^n/8$ neighbourhood of Q. There is a bi-Lipschitz map that is the identity outside $\hat{Q}$, that sends ∂Q to the boundary of a square, and that sends the distinguished points to the midpoints of opposite sides. The bi-Lipschitz constant will depend only on M_1. The regions $\hat{Q}$ are disjoint, so the composition of all the maps just described does not increase the bi-Lipschitz constant and it yields a bi-Lipschitz map F that sends S and all the elements of $\mathcal{C}(S)$ to squares and that also sends the distinguished points to the midpoints of opposite sides. It is easy to find a suitable chain for $F(S)$ and $F(\mathcal{C}(S))$, and the pullback of this chain by F^{-1} is the chain we seek.

Now let $\mathcal{F}_n$ be the family of chains obtained by applying the previous lemma to each element of $\mathcal{S}_{n-1}$. These satisfy conditions i)-iv). Consequently, E is contained in an M-quasiconformal arc.

We turn next to the proof of Theorem 3. Quasiconformal maps of $\mathbb{C}$ preserve uniform domains, uniformly disconnected sets, and uniformly perfect sets. As the middle-third set is uniformly perfect, uniformly disconnected, and has uniform complement, part iii) implies both parts i) and ii). Suppose now that K is compact and satisfies either i) or ii). Theorem 1 implies that we can assume that K is a porous and uniformly perfect subset of the real line. We need to find a quasiconformal map of $\mathbb{C}$ that maps K onto the middle-third Cantor set. We will only sketch the proof as the details are quite routine.

Set $\mathcal{O}$ to be the collection of disjoint open intervals that make up $\mathbb{R} \setminus K$. We will say that we *split* an interval I when we remove the largest subinterval that is an element of $\mathcal{O}$. Define K_0 to be the smallest closed interval containing K. Let K_1 be the union of the two intervals obtained by splitting K_0. Next split each of these to obtain another set, K_2, that is the union of four closed intervals. Continue indefinitely in this way. We summarize the properties of the sets K_n:

i) Each K_n is a finite union of disjoint closed intervals with endpoints in K, and K_0 consists of just one interval.

ii) $K_{n+1} \subseteq K_n$.

iii) Each of the intervals I that make up K_n contains exactly two of the intervals, I_l and I_r, that make up K_{n+1}, and $|I| \sim |I_l| \sim |I_r| \sim d(I_l, I_r)$.

iv) $K = \bigcap_{n=0}^{\infty} K_n$.

The key property is iii), and this is a consequence of K being porous and uniformly perfect. Suppose that I is a component of K_n and that J is a component of K_m. Let us say that $I < J$ if either $n < m$ or $n = m$ and J lies to the right of I. Label the collection of all components of all the K_n as $I_1, I_2, I_3, \ldots$ where $I_k < I_{k+1}$ for all k. Then $I_1 = K_0$, $I_2 \cup I_3 = K_1$, and so on.

From iii) above we deduce that there is a small constant ε, which we fix now, that ensures the validity of the statements that follow. To any $I_k = [x - r, x + r]$ we associate the rectangle $[x - (1 + \varepsilon) r, x + (1 + \varepsilon) r] \times [-r, r]$. Rectangles from the same level are disjoint, and the families consisting of the rectangles from each level are nested and nest nicely down to K. The next lemma is an easy consequence of iii) above. Here I, I_l, and I_r are as in iii), and R, R_l, and R_r are the corresponding rectangles.

Lemma 2.5. *There is an M_1-quasiconformal map of $\mathbb{C}$, which is the identity in R^c, is a similarity on R_l and on R_r, and which maps I_l to the left third of I and I_r to the right third of I.*

Denote by g_i the map we get from the lemma when $I = I_i$. Let τ be a similarity mapping K_0 onto $[0, 1]$. Set $G = \tau \circ g_1 \circ g_2 \circ \cdots$. Then G is an M_1-quasiconformal map of $\mathbb{C}$ that maps K onto the Cantor set. The composition does not increase the dilatation because the maps g_i only have non-trivial dilatation in the doubly connected region between the three corresponding rectangles and these regions are disjoint.

A natural question now is: do Theorems 1 and 3 hold in higher dimensions? The first point to note is that sets with uniform complement are no longer the same as uniformly disconnected sets. As an example, consider the compact set in $\mathbb{R}^3$ consisting of the line segment joining $(0, 0, 0)$ and $(0, 0, 1)$ and the family of line segments joining $(2^{-n}, 0, 0)$ and $(2^{-n}, 0, 1)$ for $n \geq 1$. This set is obviously not uniformly disconnected, yet its complement is a uniform domain. However, it is always true that any uniformly disconnected set has a uniform complement. This can be shown by the compactness method of Väisälä; see [V, Theorem 3.6]. Thus we have that ii) always implies i) in Theorem 1.

It is not reasonable to ask for quasiconformal equivalence in higher dimensions as topological issues complicate and cloud the issue. We saw in the planar case that the key to the whole problem is showing that the given set lies on a quasiconformal arc. A concept that makes sense in all dimensions and that agrees with that of a quasiconformal arc in the plane is that of an arc of bounded turning, *i.e.*, an arc with the property that the diameter of every sub-arc is comparable to the distance between its endpoints. Such arcs are precisely the quasisymmetric images of line-segments (see [TV, Section 4]). Thus a better formulation of the problem is: do uniformly disconnected sets or sets with uniform complement lie on an arc of bounded turning? The example in the previous paragraph shows that there are sets with uniform complement that do not lie on *any* Jordan arc. In contrast, it turns out that the answer is yes for uniformly disconnected sets. It follows immediately that a subset E of $\mathbb{R}^n$ is uniformly disconnected if and only if there is a quasisymmetric map from $[0, 1]$ into $\mathbb{R}^n$ that maps a porous subset of $[0, 1]$ onto E. Once we have this theorem, we find, following the planar case, that a subset of $\mathbb{R}^n$ is uniformly disconnected and uniformly perfect if and only if there is quasisymmetric map from

$[0, 1]$ into $\mathbb{R}^n$ that maps the middle-third Cantor set onto E.

The proof that a uniformly disconnected set in any Euclidean space lies on an arc of bounded turning follows essentially the same scheme as that of the planar case. We give a brief justification of this result. Lemma 2.3 will hold with $\mathcal{D}_\delta$ being those sets which are finite unions of cubes from the grid of sidelength δ, and which have connected interiors and complements. This part is straightforward. Define the families $\mathcal{S}_n$ as in the discussion preceding Lemma 2.4. We can uniformly bound the number of cubes in each element of any $\mathcal{S}_n$. Using the notation of Lemma 2.4, we define A to be

$$ S \setminus \bigcup_{\Delta \in \mathcal{C}(S)} \Delta. $$

The uniform bound on the number of cubes that make up both S and the elements of $\mathcal{C}(S)$ and the fact that these latter have connected interiors and complements allow us to show that S contains an arc γ of bounded turning with the following properties: the endpoints of γ are the distinguished points of S, the intersection of γ with any element Δ of $\mathcal{C}(S)$ consists of a sub-arc of γ whose end points are the distinguished points of Δ, and for any point on any component of $\gamma \cap A$ the distance to the nearest endpoint (of the component) and the distance to the boundary of A are comparable. We will refer to the union of γ and $\mathcal{C}(S)$ as a chain, although it does not fit our former definition. We now define $\mathcal{F}_n$ as before. The proof of [M, Theorem 5.1] shows that when such chains are nested (as described earlier in Section 1) they converge, as in the planar case, to an arc of bounded turning.

References.

[DS] David, G., Semmes, S., *Fractured Fractals and Broken Dreams: Self-Similar Geometry through Metric and Measure.* Oxford Lecture Series in Mathematics and its Applications. Oxford University Press, 1997.

[M] MacManus, P., Bi-Lipschitz Extensions in the Plane. *J. d'Analyse* **66** (1995), 85-115.

[TV] Tukia, P., Väisälä, J., Quasisymmetric embeddings of metric spaces. *Ann. Acad. Sci. Fenn. Ser. A I Math.* **5** (1980), 97-114.

[V] Väisälä, J., Uniform domains. *Tôhoku Math. J.* **40** (1988), 101-118.

Paul MacManus*
Department of Mathematics
NUI Maynooth
Maynooth
Co. Kildare, IRELAND
pmm@maths.may.ie