CONFORMAL DIMENSION AND THE QUASISYMMETRIC GEOMETRY OF METRIC SPACES

by

John M. Mackay

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Mathematics)
in The University of Michigan
2008

Doctoral Committee:

Professor Mario Bonk, Co-Chair
Professor Bruce A. Kleiner, Yale University, Co-Chair
Professor G. Peter Scott
Associate Professor James P. Tappenden
Assistant Professor Daniel Meyer

ACKNOWLEDGEMENTS

It is appropriate, but impossible, to thank all those who have helped me through my studies, however, I would like to mention some people in particular.

First, I am very grateful to my advisor Prof. Bruce Kleiner for all his guidance and encouragement, and for showing me some beautiful mathematics.

I learned a great deal too from Prof. Mario Bonk and thank him for this. I would also like to thank Prof. Yair Minsky for freely giving his time to talk. Prof. Peter Scott helped me navigate the treacherous waters of graduate school bureaucracy as I spent time at Yale, as well as teaching me about group splittings. I would like to thank the administrative staff at the mathematics departments of both Michigan and Yale for their help.

One of the most influential figures in my graduate education was Prof. Juha Heinonen. He taught me much about mathematics, and demonstrated how to live as a mathematician, teacher and mentor; he is greatly missed.

I would also like to thank my friends for their help through graduate school. There are too many to list here, but I would particularly like to mention Enrico, Amir, Diane, Kevin, Jasun, Marie, Dave and Glenn.

Finally, I gratefully thank my family for their love and support over many years and thousands of miles, particularly my sister, Fiona, and my parents, Mary and John.

TABLE OF CONTENTS

LIST OF FIGURES

Figure

CHAPTER I

Introduction

Studying a geometric object by means of its group of symmetries has been a central theme in mathematics ever since Klein proposed his Erlangen Program in 1872. In contrast, only in the last couple of decades has the converse approach been widely applied: using geometric methods to study infinite discrete groups, or Geometric Group Theory.

In a different direction, there has been a recent flourishing of research in the field of analysis on metric spaces; that is, the study of the behavior of functions and spaces with not necessarily smooth metric structure. (See [Hei07] for a survey of this area.)

At first glance these two fields may seem very far apart, since the ordinary (local) tools of topology or analysis are seemingly of little use in studying a discrete group. In fact, these fields have long had productive interactions. The classic example of applying low regularity analysis (of quasiconformal mappings) to the theory of discrete groups is given by Mostow's rigidity theorem [Mos73]. One application of this theorem is the surprising result that if two closed hyperbolic three dimensional manifolds have isomorphic fundamental groups, then the manifolds are isometric. This geometric and group theoretical result was proven by studying the behavior of quasiconformal maps on the sphere at infinity of $\mathbb{H}^3$.

In more recent developments, Gromov introduced the concept of a (Gromov) hyperbolic group, capturing the large scale geometry of the fundamental groups of compact negatively curved manifolds, but applicable in much wider contexts [Gro87]. Analogous to the sphere at infinity for hyperbolic space, hyperbolic groups have a boundary at infinity. This boundary has a canonical topological structure, and it carries a canonical *family* of quasisymmetrically equivalent metrics: an analytically rich object. Quasisymmetric maps between boundaries are in correspondence with quasi-isometric maps of the underlying groups [Pau96], and so any quasisymmetric invariant of a metric space will give a quasi-isometric invariant of a hyperbolic group.

The structure of the boundary at infinity has been studied to gain insight into the structure of the underlying group. For example, work of Bowditch shows that in a one ended hyperbolic group, the boundary has local cut points if and only if the group virtually splits over a virtually cyclic subgroup [Bow98]. Boundaries of hyperbolic groups also provide new examples of spaces with interesting analytical structure, such as those found in a recent work of Bourdon and Pajot [BP03]. These matters and more are discussed in [Kle06].

One of the most natural quasisymmetry invariants of a metric space is Pansu's conformal dimension [Pan89]: it is defined as the infimum of the Hausdorff dimensions of all quasisymmetrically equivalent metric spaces. In a sense, this invariant measures the dimension of the 'best shape' of the metric space. It was introduced by Pansu in his study of the boundary of rank one symmetric spaces which extended the work of Mostow. Although natural, it is often difficult to calculate the conformal dimension, or even to estimate it, particularly from below.

The conformal dimension of the boundary of a hyperbolic group is of particular interest in the light of recent work of Bonk and Kleiner on Cannon's conjecture. This conjecture states that if the boundary of a hyperbolic group is a topological 2-sphere, then the group acts properly, cocompactly and isometrically on $\mathbb{H}^3$. As a corollary to a theorem of Sullivan [Sul81], Cannon's conjecture can be stated as follows:

Conjecture 1.0.1. *If the boundary at infinity of a hyperbolic group is a topological 2-sphere, then the boundary is quasisymmetric to the standard 2-sphere.*

Bonk and Kleiner give a partial resolution of Cannon's conjecture [BK02a, BK05a]:

Theorem 1.0.2. *If the boundary at infinity of a hyperbolic group is a topological 2-sphere, and the (Ahlfors regular) conformal dimension of the boundary is realized, then the boundary is quasisymmetric to the standard 2-sphere.*

Consequently, it is pertinent to ask: Under what circumstances is the conformal dimension of the boundary of a hyperbolic group realized? Before dealing with boundaries that have topological dimension two, it is natural to study lower dimensional cases. The case of boundaries with topological dimension zero is trivial, since every such infinite hyperbolic group is virtually free, and so the boundary is either two points or a Cantor set. In both these cases the conformal dimension is zero; in the former it is realized, in the latter it is not.

When the topological dimension of the boundary is one, the situation is less clear. Any Fuchsian group has the standard circle as boundary, and so the conformal dimension is realized. However, there are examples of Pansu where the conformal dimension is one but it is not realized. We shall discuss this more in Chapters V and VI. In [BK05a] Bonk and Kleiner ask whether groups whose boundaries have no local cut points have conformal dimension greater than one. We answer this question affirmatively.

Theorem 5.1.1. *Suppose G is a non-elementary hyperbolic group whose boundary* *has no local cut points. Then the conformal dimension of $\partial_\infty G$ is strictly greater than one.*

This result follows from a study of the quasisymmetric geometry of metric spaces that are locally connected and have no local cut points in a uniform way. In this more general context we can again bound the conformal dimension.

Theorem 4.0.1. *Suppose (X, d) is a complete metric space which is doubling and annulus linearly connected. Then the conformal dimension $\dim_{\mathcal{C}}(X)$ is at least $C > 1$, where C depends only on the data of X (i.e., the constants associated to the two conditions above).*

(Chapter II provides general background on conformal dimension and quasisymmetric maps; the statement of this theorem is explained in more detail in Chapter IV.)

A good example of a space satisfying the hypotheses of Theorem 4.0.1 is the standard square Sierpiński carpet, especially as it has topological dimension one and so the trivial lower bound of one is no help in proving Theorem 4.0.1.

A standard way of finding an obstruction to lowering the conformal dimension of a metric space is to exhibit a good family of curves inside the space and then use a modulus-type argument. This will be the approach that we use to prove Theorem 4.0.1; the modulus argument is contained in a lemma of Pansu (see Section 2.3).

In order to construct our family of curves we will need a theorem of Tukia (Theorem 3.1.4) that allows us to construct arcs that are quasisymmetric to the unit interval ('quasi-arcs') inside a large class of metric spaces. We give a new and improved proof of this theorem in Chapter III.

The process of finding a good family of curves involves a somewhat delicate limiting argument whereby we begin with one arc, 'split' it into two, and then four, and so on. In between these steps we use Tukia's theorem to straighten out our arcs, and thereby control the limit of this process. Chapter IV contains the details of this construction; see Theorem 4.0.2 in particular.

After applying Theorem 4.0.1 to the boundaries of hyperbolic groups in Chapter V, we conclude with discussing natural questions that lead out of this work and provide future avenues of research.

It should be noted that the new results described here have previously appeared in the papers [Mac07a] and [Mac07b].

CHAPTER II

Metric space dimension and quasisymmetric geometry

In this chapter we recall the definition and basic properties of the Hausdorff dimension of a metric space. We will also describe its lesser-known counterpart in the world of quasisymmetric geometry: the conformal dimension of a metric space. As mentioned in Chapter I, a lemma of Pansu provides a key tool in bounding the conformal dimension from below. This lemma is described and proven in Section 2.3.

Finally, we define the Hausdorff and Gromov-Hausdorff metrics in Section 2.4, and state some of their properties.

We will use the notation (X, d) for a metric space X with metric $d : X \times X \to \mathbb{R}^+$; if the metric is understood we may refer simply to X. An open ball about $x \in X$ of radius r is denoted by $B(x, r) = \{y \in X : d(x, y) < r\}$. The corresponding closed ball is denoted by $\overline{B}(x, r) = \{y \in X : d(x, y) \leq r\}$.

2.1 Hausdorff dimension

For smooth manifolds such as curves and surfaces, the standard notion of topological dimension measures the number of degrees of freedom present. For example, a curve is one dimensional, a surface two dimensional, and so on. In the case of certain 'fractal' metric spaces the topological dimension, while still defined, often proves inadequate.

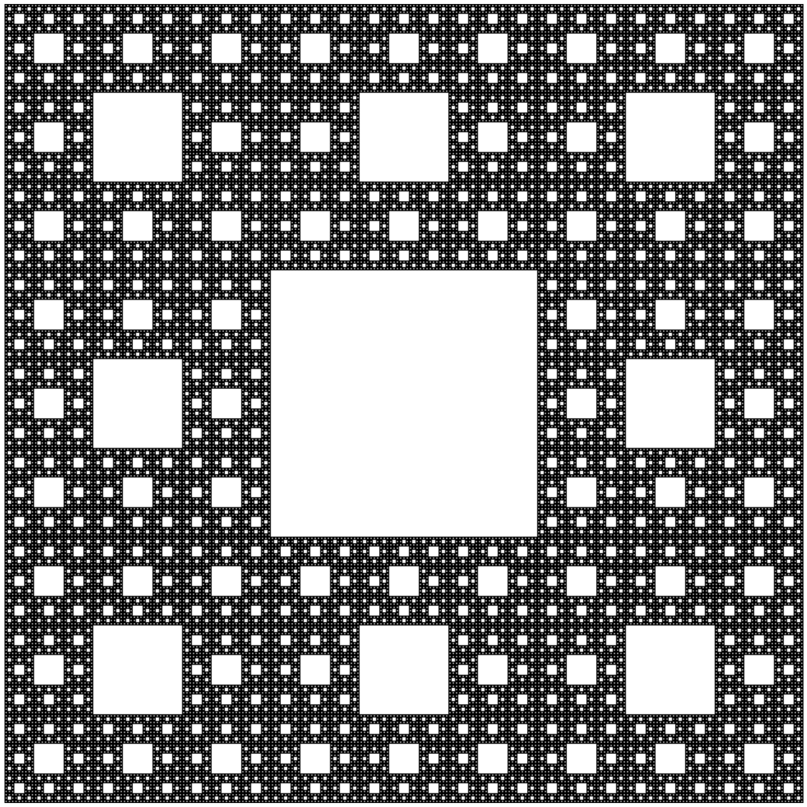

Figure 2.1: The Sierpiński Carpet

To see this, let us consider the example of the Sierpiński carpet. This is constructed from the unit square in the plane, S_0. We divide S_0 into nine squares of side length one third and remove the middle square to get the set S_1. This process is then repeated for the remaining eight squares to get the set S_2, and so on. The limit object $S = \cap_{i=0}^{\infty} S_i$ is called the Sierpiński carpet, see Figure 2.1.

It is a simple observation that the number of small balls of radius r required to cover a compact n-dimensional Riemannian manifold is proportional to r^{-n}. However, since S_i is made up of 8^i squares of side length 3^{-i} we can show that the number of balls of radius $r = 3^{-i}$ required to cover S is proportional to $8^i = (3^{-i})^{-(\log 8)/(\log 3)} = r^{-(\log 8)/(\log 3)}$: it seems that the Sierpiński carpet has 'dimension' $n = \frac{\log 8}{\log 3} \approx 1.893$.

One way to make this discussion precise is with the concept of the *Hausdorff dimension* of a metric space (X, d). First, given $d > 0$ we define the d-dimensional Hausdorff measure of a set $A \subset X$:

$$\mathcal{H}^d(A) = \lim_{\delta \to 0} \inf \left\{ \sum (r_i)^d : \{B(x_i, r_i)\} \text{ covers } X, r_i \leq \delta \right\} \in [0, \infty]. \tag{2.1}$$

In the infimum above we only consider covers by countable collections of sets.

(The standard definition of Hausdorff measure considers covers by arbitrary open sets and uses the diameter of a set instead of the radius of a ball in the equation above. However, definition (2.1) is comparable and will be more convenient to use.)

As is well known (for example, see [Hei01, Section 8.3]), the Hausdorff measure satisfies some basic dimension comparison results.

$$\begin{aligned} d < d', \mathcal{H}^d(A) < \infty &\implies \mathcal{H}^{d'}(A) = 0, \text{ and} \\ d > d', \mathcal{H}^d(A) > 0 &\implies \mathcal{H}^{d'}(A) = \infty. \end{aligned} \tag{2.2}$$

We can now define the Hausdorff dimension of $A \subset X$ as

$$\dim_{\mathcal{H}}(A) = \inf\{d : \mathcal{H}^d(X) = 0\} \in [0, \infty]. \tag{2.3}$$

Example 2.1.1. If (X, d) is an n-dimensional Riemannian manifold, then we have $\dim_{\mathcal{H}}(X) = n$.

Example 2.1.2. As expected, the Sierpiński carpet S has Hausdorff dimension $\dim_{\mathcal{H}}(S) = \frac{\log 8}{\log 3}$.

We also have the important inequality [Hei01, Theorem 8.14]

$$\dim_{\mathrm{top}}(X) \leq \dim_{\mathcal{H}}(X). \tag{2.4}$$

For proofs of these statements and others, see the reference [Hei01].

2.2 Quasisymmetric maps and conformal dimension

In the category of metric spaces, the most natural maps are isometries: maps that preserve distances. Isometries are much too rigid and rare in most contexts, where we consider bi-Lipschitz maps instead. These maps preserve distances up to a constant multiplicative factor. To be precise, we say that a map $f : (X, d) \to (X', d')$ is a L-bi-Lipschitz map, for some constant $L \geq 1$, if for all $x, y \in X$,

$$\frac{1}{L} d(x, y) \leq d'(f(x), f(y)) \leq L d(x, y). \tag{2.5}$$

It is straightforward to see that under a L-bi-Lipschitz map the d-dimensional Hausdorff content of a space varies by a multiplicative factor of at most L^d. Therefore, the Hausdorff dimension is preserved by bi-Lipschitz homeomorphisms.

We will be working in the category of quasisymmetric maps, where metrics can be distorted in a wilder way. Quasisymmetric maps were first considered in the context of complex analysis, where quasiconformal mappings have been a key tool since Grötzsch and Ahlfors. A quasiconformal map from $\mathbb{R}^n$ to $\mathbb{R}^n$ is a homeomorphism which sends infinitesimal circles to infinitesimal ellipses of uniformly controlled eccentricity. This infinitesimal control actually implies global control on such maps: they are quasisymmetric.

(In the context of general metric spaces the passage from infinitesimal to global control is not always possible. However, in recent work Heinonen and Koskela prove this implication is true provided the metric space is Loewner [HK98].)

Quasisymmetric maps are defined using a three point condition which says that, in effect, balls are sent to 'quasi-balls' of uniformly controlled eccentricity.

Definition 2.2.1. A topological embedding $f : (X, d) \to (X', d')$ is *quasisymmetric* if there exists a homeomorphism $\eta : [0, \infty) \to [0, \infty)$ such that for every triple of points $x, y, z \in X$, $x \neq z$, we have

$$\frac{d'(f(x), f(y))}{d'(f(x), f(z))} \leq \eta\left(\frac{d(x,y)}{d(x,z)}\right). \tag{2.6}$$

If we wish to specify the distortion function used, we say that f is an η-quasisymmetric embedding.

If f is also surjective then we call f a quasisymmetric homeomorphism and say (X, d) and (X', d') are quasisymmetrically equivalent, written as $(X, d) \overset{\text{qs}}{\simeq} (X', d')$.

Example 2.2.2. If f is a L-bi-Lipschitz embedding, then it is also a quasisymmetric embedding with distortion function $\eta(t) = L^2 t$.

Example 2.2.3. For any metric space (X, d) and $\epsilon \in (0, 1]$, Hölder's inequality implies that d^ϵ satisfies the triangle inequality and so (X, d^ϵ) is also a metric space. We can further see that $\dim_{\mathcal{H}}(X, d^\epsilon) = \frac{1}{\epsilon}\dim_{\mathcal{H}}(X, d)$.

The process of raising the metric by a power smaller than one is called 'snowflaking' because in the case of the unit interval, choosing $\epsilon = \frac{\log 3}{\log 4}$ gives a metric space that is bi-Lipschitz to the von Koch snowflake.

We have shown that the Hausdorff dimension of a metric space is not a quasisymmetric invariant because, using snowflake maps, it is possible to raise the dimension arbitrarily high by 'crinkling up' the space. On the other hand, it is a much more difficult, and interesting, challenge to 'straighten out' a space, i.e., to lower its Hausdorff dimension.

With this in mind, we can define a natural notion of dimension that is quasisymmetrically invariant.

Definition 2.2.4. The *conformal dimension* of a metric space (X, d) is the infimal Hausdorff dimension among all quasisymmetrically equivalent metric spaces, and is denoted by $\dim_{\mathcal{C}}(X,d)$:

$$\dim_{\mathcal{C}}(X,d) := \inf\{\dim_{\mathcal{H}}(X',d') : (X,d) \overset{\text{qs}}{\simeq} (X',d')\}. \tag{2.7}$$

This definition was introduced by Pansu in his study of the conformal structure possessed by the boundary at infinity of negatively curved manifolds [Pan89]. We will discuss its relevance to hyperbolic groups in Section 5.1.

Notice that for a metric space (X,d), by equations (2.4) and (2.7) we have that

$$\dim_{\text{top}}(X,d) \leq \dim_{\mathcal{C}}(X,d) \leq \dim_{\mathcal{H}}(X,d). \tag{2.8}$$

We have already stated that the image of a ball under a quasisymmetric homeomorphism is a 'quasi-ball' of controlled distortion. In the next section we will need a stronger statement concerning the image of an annulus, which we now prove. Recall that a proper metric space is one where all closed balls are compact.

Lemma 2.2.5. *If $f : (X,d) \to (X',d')$ is an η-quasisymmetric homeomorphism between proper, non-trivial metric spaces, then for all $x \in X$, $r > 0$ and $C \geq 1$, there exists some $R > 0$ such that*

$$B(f(x),R) \subset f(B(x,r)) \subset f(B(x,Cr)) \subset B(f(x), 2\eta(C)R). \tag{2.9}$$

Proof. First, we assume that $B(x,r) \subsetneq X$.

Set $x' = f(x)$, let y' be the furthest point from x' in $f(\overline{B(x,Cr)})$ and let z' be the closest point to x' in $X' \setminus f(B(x,r))$. If we set $R = d'(x',z')$, then we have

$$B(x',R) \subset f(B(x,r)) \subset f(B(x,Cr)) \subset B\left(x', 2\frac{d'(x',y')}{d'(x',z')}R\right).$$

However, if $y = f^{-1}(y')$, we know that y lies in $\overline{B(x,Cr)}$, and so $d(x,y) \leq Cr$. Similarly, if $z = f^{-1}(z')$ then z lies outside $B(x,r)$, and $d(x,z) \geq r$. Thus

$$\frac{d'(x',y')}{d'(x',z')} \leq \eta\left(\frac{d(x,y)}{d(x,z)}\right) \leq \eta(C).$$

As a corollary to the proof so far, note that if X is bounded then so is X'. To see this, suppose $\operatorname{diam}(X) \leq \infty$. Then there exist $u, v \in X$ such that $d(u,v) \geq \frac{3}{4}\operatorname{diam}(X)$, thus

$$B(u, \tfrac{1}{2}\operatorname{diam}(X)) \subsetneq X = B(u, 3\tfrac{1}{2}\operatorname{diam}(X)),$$

and so

$$X' = f\left(B(u, \tfrac{3}{2}\operatorname{diam}(X))\right) \subset B(x', 2\eta(3)R)$$

for some R, in particular, X' is bounded.

So, finally, if $B(x,r) = X$ then (2.9) is satisfied for $R = \frac{3}{2}\operatorname{diam}(X')$. (Strictly speaking, we are using that $\eta(1) \geq 1$; this follows from the definition of a quasisymmetric map by taking $y = z$ in (2.6).) $\square$

As we have seen, a quasisymmetric map $f : (X,d) \to (X',d')$ need not be bi-Lipschitz; however, on good spaces it will be Hölder, i.e., there exist constants $C > 0$, and $\alpha, \beta \in (0,1]$ such that

$$\frac{1}{C} d(x,y)^{\alpha} \leq d(f(x), f(y)) \leq C d(x,y)^{\beta}.$$

Theorem 2.2.6 (Corollary 11.5 of [Hei01])**.** *Quasisymmetric embeddings of uniformly perfect spaces are Hölder continuous on bounded sets.*

A metric space X is *uniformly perfect* if there exists a constant $C > 1$ such that for all $x \in X$ and $0 < r \leq \operatorname{diam}(X)$, then $B(x,r) \setminus B(x, r/C) \neq \emptyset$. For example, connected spaces are uniformly perfect.

2.3 Pansu's lower bound for conformal dimension

Finding a lower bound for conformal dimension that is greater than the topological dimension is a difficult challenge. A key tool that we will use is a lemma of Pansu that gives a nontrivial lower bound in the presence of certain families of curves [Pan89, Lemma 6.3].

We will use one further notion of dimension in this proof: the *packing dimension* of a metric space. This is defined in a way analogous to the definition of Hausdorff dimension. First, we define the d-dimensional packing premeasure of a set $A \subset X$.

(2.10)
$$\mathcal{P}^d(A) = \lim_{\delta \to 0} \sup \left\{ \sum (r_i)^d : \{B(x_i, r_i)\} \text{ is a packing}, x_i \in A, r_i \leq \delta \right\} \in [0, \infty].$$

By a 'packing' we mean a pairwise disjoint countable collection of balls.

As before, we have the following inequalities [Cut95, Theorem 3.7(f)]:

$$(2.11) \qquad \begin{aligned} d < d', \mathcal{P}^d(A) < \infty \quad &\Longrightarrow \quad \mathcal{P}^{d'}(A) = 0, \text{ and} \\ d > d', \mathcal{P}^d(A) > 0 \quad &\Longrightarrow \quad \mathcal{P}^{d'}(A) = \infty, \end{aligned}$$

and we define the packing dimension of $A \subset X$ as

$$(2.12) \qquad \dim_{\mathcal{P}}(A) = \inf\{d : \mathcal{P}^d(X) = 0\} \in [0, \infty].$$

The packing dimension, as defined here, is the same as the upper Minkowski dimension [Cut95, Theorem 3.7(g)], and so for all $A \subset X$,

$$(2.13) \qquad \dim_{\text{top}}(A) \leq \dim_{\mathcal{H}}(A) \leq \dim_{\mathcal{P}}(A).$$

The version of Pansu's lemma that we shall prove is due to Bourdon [Bou95, Lemma 1.6]. We include it below because of its importance to our work and for the reader's convenience; I have not found a proof in the literature written in English.

Lemma 2.3.1. *Suppose (X, d) is a uniformly perfect, compact metric space containing a collection of arcs $\mathcal{C} = \{\gamma_i | i \in I\}$ whose diameters are bounded away from zero. Suppose further that we have a Borel probability measure μ on $\mathcal{C}$ and constants*

$A > 0$, $\sigma \geq 0$ such that for all balls $B(x,r)$ in X the set $\{\gamma \in \mathcal{C} | \gamma \cap B(x,r) \neq \emptyset\}$ is μ-measurable with measure at most Ar^{σ}.

Then the packing dimension $\tau = \dim_{\mathcal{P}}(X)$ satisfies $\tau - \sigma \geq 1$, and the conformal dimension of X is at least $1 + \frac{\sigma}{\tau - \sigma} > 1$.

Proof. Suppose (X, d) carries a quasisymmetrically equivalent metric d'. We wish to find a lower bound for the Hausdorff dimension of (X, d') that depends only on the data described in the lemma.

Note that X is bounded since it is compact, and since (X, d) is uniformly perfect Theorem 2.2.6 implies that the metrics d and d' are Hölder equivalent. We will consider balls and distances in both the d and d' metrics; those in the latter will be distinguished by the use of a prime symbol.

Fix $r' > 0$, and let $\{B'_k : k \in K\}$ be a countable cover of X by (d'-)balls, where each B'_k has radius r'_k less than r'. Since X is compact we may assume that the index set K is finite.

Using the '5B'-lemma for ball coverings in metric spaces [Hei01, Theorem 1.2], we may extract a disjoint collection of balls $\{B'_j : j \in J\}$, where the collection $\{5B'_j : j \in J\}$ also covers X'. (Here, if $B = B(x,r)$ is a ball and $C > 0$ a constant, then CB denotes the ball $CB = B(x, Cr)$.)

By applying Lemma 2.2.5 to the identity map and the ball pairs $B' \subset 5B'$, there exists a constant $C \geq 1$ and balls B_j, for each $j \in J$, with the same center as B'_j and (d-)radii r_j that satisfy

$$B_j \subset B'_j \subset 5B'_j \subset CB_j.$$

We will use the following characteristic function to detect when a curve $\gamma \in \mathcal{C}$

meets $5B'_j$.

$$\chi_j(\gamma) = \begin{cases} 1 & \text{if } \gamma \cap 5B'_j \neq \emptyset, \\ 0 & \text{otherwise.} \end{cases}$$

Since we know that the curves $\gamma \in \mathcal{C}$ have uniformly large d-diameter, and the metrics d and d' are Hölder equivalent, each $\gamma \in \mathcal{C}$ has d'-diameter greater than some constant $D > 0$.

The collection of balls $\{5B'_j : j \in J\}$ covers X, so it also covers each $\gamma \in \mathcal{C}$, and thus for each $\gamma \in \mathcal{C}$ we have

$$D \leq \sum_{j \in J} 10 r'_j \chi_j(\gamma).$$

We will average this inequality over $(\mathcal{C}, \mu)$, and use Hölder's inequality to show the following bound:

$$\begin{aligned}
\frac{1}{10} D &\leq \int_{\mathcal{C}} \sum_{j \in J} r'_j \chi_j(\gamma) d\mu(\gamma) \\
&= \sum_{j \in J} r'_j \int_{\mathcal{C}} \chi_j(\gamma) d\mu(\gamma) \\
&\leq \sum_{j \in J} r'_j \mu\{\gamma \in \mathcal{C} : \gamma \cap CB_j \neq \emptyset\} \\
&\leq AC^\sigma \sum_{j \in J} r'_j r_j^\sigma \qquad (2.14) \\
&\leq AC^\sigma \left(\sum_{j \in J} {r'_j}^{\frac{\beta}{\beta-1}} \right)^{\frac{\beta-1}{\beta}} \left(\sum_{j \in J} {r_j}^{\sigma\beta} \right)^{\frac{1}{\beta}}. \qquad (2.15)
\end{aligned}$$

Now, $\{B_j : j \in J\}$ is a packing of (X, d) by balls of radius less than r. Since d and d' are Hölder equivalent, we have that $r \to 0$ as $r' \to 0$. Consequently, if $\sigma\beta > \tau = \dim_{\mathcal{P}}(X, d)$ then for r' (and hence r) sufficiently small we have

$$\sum_{j \in J} {r_j}^{\sigma\beta} \leq 1.$$

So we obtain

$$\sum_{j\in J} r_j'^{\frac{\beta}{\beta-1}} \geq \left(A^{-1}C^{-\sigma}\frac{D}{10}\right)^{\frac{\beta}{\beta-1}} = \tilde{C} > 0,$$

where $\tilde{C}$ is independent of the choice of cover of (X, d'). Thus $\mathcal{H}^{\frac{\beta}{\beta-1}}(X, d) \geq \tilde{C}$ and

$$\dim_{\mathcal{H}}(X, d') \geq \frac{\beta}{\beta - 1}. \tag{2.16}$$

We can take any $\beta > \frac{\tau}{\sigma}$, and so we wish to bound the packing dimension τ of (X, d). If we apply (2.14) to the case $d = d'$, $r_j' = r_j$, then we get that the Hausdorff dimension $\dim_{\mathcal{H}}(X, d)$ satisfies $\dim_{\mathcal{H}}(X, d) \geq \sigma + 1$, and so by (2.13) $\tau = \dim_{\mathcal{P}}(X, d) \geq \sigma + 1$.

Finally, taking $\beta \searrow \frac{\tau}{\sigma}$ in (2.16) we get that

$$\dim_{\mathcal{H}}(X, d') \geq 1 + \frac{\sigma}{\tau - \sigma}. \qquad \square$$

Remark 2.3.2. It is clear that we can apply the lemma when we only assume that X is proper, rather than compact, but require the collection $\mathcal{C}$ to be contained in some bounded set.

Example 2.3.3. Let (X, d) be an Ahlfors regular compact metric space of Hausdorff dimension Q, that is to say, there exists some constant C such that for all $x \in X$ and $r \leq \operatorname{diam}(X)$,

$$\frac{1}{C} r^Q \leq \mathcal{H}^Q(B(x, r)) \leq C r^Q.$$

In this case, $X \times [0, 1]$ is a compact, uniformly perfect metric space of Hausdorff, and packing, dimension $1 + Q$. We can take the collection of curves $\mathcal{C} = \{\gamma_x : [0, 1] \to X \times [0, 1]\}_{x\in X}$, where $\gamma_x(t) = (x, t)$, and the measure on $\mathcal{C}$ is $\mathcal{H}^Q$.

Lemma 2.3.1 together with (2.8) now demonstrates that

$$\begin{aligned} 1 + Q = \dim_{\mathcal{H}}(X \times [0, 1]) &\geq \dim_{\mathcal{C}}(X \times [0, 1]) \\ &\geq 1 + \frac{Q}{(Q + 1) - Q} = 1 + Q, \end{aligned}$$

and so the conformal dimension of $X \times [0,1]$ equals $1 + Q$.

Example 2.1.2, continued. Since the Sierpiński carpet S contains a copy of $C \times [0,1]$, where C is the standard $\frac{1}{3}$-Cantor set, the previous example gives the bound $\dim_{\mathcal{C}}(S) \geq 1 + \frac{\log 2}{\log 3} \approx 1.631$.

2.4 Hausdorff distance and convergence

In a metric space (X, d) there are different ways to measure the distance between two sets. The most obvious is the (infimal) distance between two subsets U, $V \subset X$. This is defined as

$$d(U, V) = \inf\{d(u, v) : u \in U, v \in V\}. \tag{2.17}$$

The problem with this definition is that two sets can be close in the infimal distance but look very different. For example, the x and y axes in the plane have distance zero, despite being very different sets.

The Hausdorff distance between U and V, $d_{\mathcal{H}}(U, V)$, is another measurement of when two sets are close that better avoids this problem. First, given a point $u \in X$, we set $d(u, V) = d(\{u\}, V)$, and say that the r-neighborhood of U is the set $N(U, r) = \{x : d(x, U) < r\}$. We can now define

$$d_{\mathcal{H}}(U, V) = \inf\{r : U \subset N(V, r), V \subset N(U, r)\} \in [0, \infty]. \tag{2.18}$$

Example 2.4.1. If we consider $\mathbb{Z} \subset \mathbb{R}$, then $d_{\mathcal{H}}(\mathbb{Z}, \mathbb{R}) = \frac{1}{2}$.

Example 2.4.2. If we let U and V denote the x and y axes in $\mathbb{R}^2$ respectively, then $d_{\mathcal{H}}(U, V) = \infty$.

The Hausdorff distance has many useful properties [BBI01, Section 7.3]; here are some that we will use.

Theorem 2.4.3. *Let $\mathcal{M}(X)$ denote the set of all closed subsets of a metric space X. Then $(\mathcal{M}(X), d_{\mathcal{H}})$ is a metric space. Furthermore,*

- *If X is complete, then $(\mathcal{M}(X), d_{\mathcal{H}})$ is complete.*
- *If X is compact, then $(\mathcal{M}(X), d_{\mathcal{H}})$ is compact.*

We denote the Hausdorff limit of a sequence of sets S_n by $\lim_{\mathcal{H}} S_n$. For example, if $\mathbb{Z}/n$ denotes the rescaling of the integers by $\frac{1}{n}$, then $\lim_{\mathcal{H}} \mathbb{Z}/n = \mathbb{R}$.

We can further develop the Hausdorff metric into a measurement of distance between *different* metric spaces. (See [BBI01, Chapter 7].) We define the *Gromov-Hausdorff distance* between metric spaces X and Y by considering all metric spaces Z containing subsets X' and Y' isometric to X and Y respectively, and let $d_{GH}(X, Y)$ be the infimum of the distances $d_{\mathcal{H}}(X', Y')$.

A more intuitive interpretation of this metric is the following: If X and Y are metric spaces with $d_{GH}(X, Y) \leq \epsilon$, then there exists a function $f : X \to Y$, not necessarily continuous, so that

- f distorts distances by at most an additive error of 2ϵ, and
- every point of Y is within 2ϵ of the image of f.

Conversely, the existence of a map f satisfying these two conditions implies that $d_{GH}(X, Y) \leq 4\epsilon$. (See [BBI01, Cor. 7.3.28].)

If the spaces X_i satisfy a condition that can be reasonably expressed in terms of distances of finitely many points, then the limit space $X = \lim_{GH} X_i$ will satisfy this condition also.

Example 2.4.4. A metric space is a length space if the distance between two points equals the infimum of the lengths of paths between the points. For a complete metric space X, this is equivalent to the existence of ϵ-midpoints for every $\epsilon > 0$. A point z is an ϵ-midpoint for x and y if $d(x,z)$ and $d(y,z)$ are both greater than or equal to $\frac{1}{2}d(x,y) - \epsilon$. Consequently, if X_i are length spaces, and $\lim_{GH} X_i = X$, where X is a complete metric space, then X will be a length space also.

Example 2.4.5. A compact metric space X is connected if and only if for every pair of points $x, y \in X$ and $\epsilon > 0$ there exists a finite chain of points joining x to y by jumps of distance less than ϵ.

Therefore, if X_i are compact, connected metric spaces and $X = \lim_{GH} X_i$ is a compact metric space, then X will also be connected.

The following definition provides a convenient way of ensuring that a metric space is finite dimensional on every scale. It is equivalent to the condition of having finite Assouad dimension.

Definition 2.4.6. A metric space (X,d) is N-*doubling* if every closed ball of radius R can be covered using at most N closed balls of radius $\frac{R}{2}$.

Example 2.4.7. The doubling condition is equivalent to the existence of N points $x_1, \ldots, x_N$ in each closed ball $\overline{B}(R)$, so that $\overline{B}(R) \subset \bigcup \overline{B}(x_i, \frac{R}{2})$.

This formulation makes it easy to prove that if each X_i is N-doubling, for the same N, and $X = \lim_{GH} X_i$ is complete, then X will be N-doubling as well.

The importance of the doubling condition is provided by the following precompactness result.

Theorem 2.4.8 (Theorem 7.4.15 of [BBI01])**.** *The class of N-doubling, complete metric spaces of diameter at most D is precompact: any sequence of such spaces has a convergent subsequence in the Gromov-Hausdorff metric.*

Sketch of proof. Suppose X_i is a sequence of such spaces. The doubling condition implies that there exist constants $0 \leq N_1 \leq N_2 \leq \cdots \leq N_m \leq \cdots$, independent of i, and countable sets $S_i = \{x^i_j \in X_i : j \in \mathbb{N}\}$ in each X_i, so that for every i and m, the $\frac{1}{m}$ neighborhood of $\{x^i_j : 1 \leq j \leq N_m\}$ equals X_i.

Now, for $i, j, k \in \mathbb{N}$ we have that $d(x^i_j, x^i_k)$ lies in $[0, D]$. By choosing subsequences using an Arzelà-Ascoli style of argument – we re-label our indices for convenience – we can ensure that, for each j and k, the limit

$$\lim_{i\to\infty} d(x^i_j, x^i_k) \tag{2.19}$$

exists.

Let S be an abstract countably infinite set $S = \{y_j : j \in \mathbb{N}\}$. It is natural to define a metric on S by setting $d(y_j, y_k)$ to be the limit given by (2.19). However, d may be only a pseudometric: two different points in S may be zero distance from each other. This is easily remedied by considering the quotient space $S/(d = 0)$. Finally, we set X to be the completion of $S/(d = 0)$.

By construction, X is complete, and for each m the $\frac{2}{m}$ neighborhood of the set $T_m = \{y_j : 1 \leq j \leq N_m\}$ covers X, so X is totally bounded and therefore compact. The fact that X is the Gromov-Hausdorff limit of (our subsequence of) the spaces X_i follows from the fact that X is well approximated by the finite sets T_m, which are themselves well approximated by their counterparts in $S_i \subset X_i$. □

This proof immediately adapts to more complicated situations. For example, suppose we have a sequence of configurations (X_i, A_i), where each A_i is a closed subset of a compact, N-doubling metric space X_i of diameter at most D. Using the same proof, we can choose an appropriate subsequence so that the configurations (X_i, A_i) Gromov-Hausdorff converge to a configuration (X_∞, A_∞), where $A_\infty \subset X_\infty$

is a closed subset.

Gromov-Hausdorff convergence for such a sequence of configurations means that for any $\epsilon > 0$ there exists $N > 0$ and maps $f_i : X_i \to X_\infty$ so that for all $i \geq N$, we have the following properties:

- f_i distorts distances by at most ϵ.
- Every point of X_∞ is within ϵ of $f_i(X_i)$.
- $d_{\mathcal{H}}(f_i(A_i), A_\infty) \leq \epsilon$.

To see how this convergence is achieved, fix m and consider the following sets for each i:

$$I^i_m := \left\{ j : d(x^i_j, A_i) \leq \frac{1}{m}, 1 \leq j \leq N_m \right\}. \tag{2.20}$$

We can choose a subsequence so that I^i_m is the same for every i. Doing this for each m and then choosing a diagonal subsequence gives the desired convergent subsequence.

We will use the precompactness of configurations of sets in Chapter IV.

CHAPTER III

Existence of quasi-arcs

3.1 Introduction and definitions

It is a standard topological fact that a complete metric space which is locally connected, connected and locally compact is arc-wise connected. Tukia [Tuk96] showed that an analogous geometric statement is true: if a complete metric space is linearly connected and doubling, then it is connected by quasi-arcs, quantitatively. In fact, he proved a stronger result: any arc in such a space may be approximated by a local quasi-arc in a uniform way. In this chapter we give a new and more direct proof of this fact.

Besides its intrinsic interest, this result is useful in studying the quasisymmetric geometry of metric spaces. For example, Tukia's result was used in the context of boundaries of hyperbolic groups by Bonk and Kleiner [BK05b], and it will be a key tool in our proof of Theorem 4.0.1. Bonk and Kleiner use Assouad's embedding theorem to translate Tukia's result from its original context of subsets of $\mathbb{R}^n$ into our setting of doubling and linearly connected metric spaces.

Before stating the theorem precisely, we recall some definitions. As we saw in Definition 2.4.6, a metric space (X, d) is called doubling if there exists a constant N such that every ball can be covered by at most N balls of half the radius. Note that any complete, doubling metric space is proper: all closed balls are compact.

Definition 3.1.1. We say (X, d) is *L-linearly connected* for some $L \geq 1$ if for all x, $y \in X$ there exists a compact, connected set $J \ni x$, y of diameter less than or equal to $Ld(x, y)$. (This is also known as bounded turning or LLC(1).)

We can actually assume that J is an arc, at the cost of increasing L by an arbitrarily small amount. To see this, note that X is locally connected, and so the connected components of an open set are open. Thus, for any open neighborhood U of J, the connected component of U that contains J is an open set. We can replace J inside U by an arc with the same endpoints, since any open, connected subset of a locally compact, locally connected metric space is arc-wise connected [Cul68, Corollary 32.36].

It is straightforward to adapt Example 2.4.5 to show that if X_n are L-linearly connected compact metric spaces, and $X = \lim_{\mathcal{H}} X_n$ is a compact metric space, then X is also L-linearly connected. We will need to use this fact in Chapter IV.

For any x and y in an embedded arc A, we denote by $A[x, y]$ the closed, possibly trivial, subarc of A that lies between them.

Definition 3.1.2. We say that an arc A in a doubling and complete metric space is an *ϵ-local λ-quasi-arc* if $\operatorname{diam}(A[x, y]) \leq \lambda d(x, y)$ for all x, $y \in A$ such that $d(x, y) \leq \epsilon$.

This terminology is explained by Tukia and Väisälä's characterization of quasisymmetric images of the unit interval as those metric arcs that are doubling and bounded turning [TV80].

One non-standard definition will be useful in our exposition.

Definition 3.1.3. We say that an arc B *ϵ-follows* an arc A if there exists a coarse map $p : B \to A$, sending endpoints to endpoints, such that for all x, $y \in B$, $B[x, y]$

is in the ϵ-neighborhood of $A[p(x), p(y)]$; in particular, p displaces points at most ϵ. (We call the map p *coarse* to emphasize that it is not necessarily continuous.)

The condition that B ϵ-follows A is stronger than the condition that B is contained in the ϵ-neighborhood of A. It says that, coarsely, the arc B can be obtained from the arc A by cutting out 'loops.' (Of course, A contains no actual loops, but it may have subarcs of large diameter whose endpoints are 2ϵ-close.)

We can now state the stronger version of Tukia's theorem precisely, and as an immediate corollary our initial statement [Tuk96, Theorem 1B, Theorem 1A]:

Theorem 3.1.4 (Tukia). *Suppose (X, d) is a L-linearly connected, N-doubling, complete metric space. For every arc A in X and every $\epsilon > 0$, there is an arc J that ϵ-follows A, has the same endpoints as A, and is an $\alpha\epsilon$-local λ-quasi-arc, where $\lambda = \lambda(L, N) \geq 1$ and $\alpha = \alpha(L, N) > 0$.*

Corollary 3.1.5 (Tukia). *Every pair of points in a L-linearly connected, N-doubling, complete metric space is connected by a λ-quasi-arc, where $\lambda = \lambda(L, N) \geq 1$.*

Our strategy for proving Theorem 3.1.4 is straightforward: find a method of straightening an arc on a given scale (Proposition 3.2.1), then apply this result on a geometrically decreasing sequence of scales to get the desired local quasi-arc as a limiting object. The statement of this proposition and the resulting proof of the theorem essentially follow Tukia [Tuk96], but the proof of the proposition is new and much shorter.

3.2 Finding quasi-arcs as a limit of arcs

Given any arc A and $\iota > 0$, the following proposition allows us to straighten A on a scale ι inside the ι-neighborhood of A.

Proposition 3.2.1. *Given a complete metric space X that is L-linearly connected and N-doubling, there exist constants $s = s(L, N) > 0$ and $S = S(L, N) > 0$ with the following property: for each $\iota > 0$ and each arc $A \subset X$, there exists an arc J that ι-follows A, has the same endpoints as A, and satisfies*

$$(*) \qquad \forall x, y \in J,\ d(x, y) < s\iota \implies \operatorname{diam}(J[x, y]) < S\iota.$$

We will apply this proposition on a decreasing sequence of scales to get a local quasi-arc in the limit. The key step in proving this is given by the following lemma.

Lemma 3.2.2. *Suppose (X, d) is a L-linearly connected, N-doubling, complete metric space, and let s, S, ϵ and δ be fixed positive constants satisfying $\delta \le \min\{\frac{s}{4+2S}, \frac{1}{10}\}$. Now, if we have a sequence of arcs $J_1, J_2, \ldots, J_n, \ldots$ in X, such that for every $n \ge 1$*

- *J_{n+1} $\epsilon\delta^n$-follows J_n, and*
- *J_{n+1} satisfies $(*)$ with $\iota = \epsilon\delta^n$ and s, S as fixed above,*

then the Hausdorff limit $J = \lim_{\mathcal{H}} J_n$ exists, and is an $\epsilon\delta^2$-local $\frac{4S+3\delta}{\delta^2}$-quasi-arc.

Moreover, the endpoints of J_n converge to the endpoints of J, and J ϵ-follows J_1.

We will now prove Theorem 3.1.4.

Proof of Theorem 3.1.4. Let s and S be given by Proposition 3.2.1, and set $\delta = \min\{\frac{s}{4+2S}, \frac{1}{10}\}$.

Let $J_1 = A$ and apply Proposition 3.2.1 to J_1 and $\iota = \epsilon\delta$ to get an arc J_2 that $\epsilon\delta$-follows J_1. Repeat, applying the lemma to J_n and $\iota = \epsilon\delta^n$, to get a sequence of arcs J_n, where each J_{n+1} $\epsilon\delta^n$-follows J_n, and satisfies $(*)$ with $\iota = \epsilon\delta^n$.

We can now apply Lemma 3.2.2 to find an $\alpha\epsilon$-local λ-quasi-arc J that ϵ-follows A, where $\alpha = \delta^2$ and $\lambda = \frac{4S+3\delta}{\delta^2}$. Every J_n has the same endpoints as A, so J will also have the same endpoints. $\square$

The proof of Lemma 3.2.2 relies on some fairly straightforward estimates and a classical characterization of an arc.

Proof of Lemma 3.2.2. For every $n \geq 1$, J_{n+1} $\epsilon\delta^n$-follows J_n. We denote the associated coarse map by $p_{n+1} : J_{n+1} \to J_n$.

In the following, we will make frequent use of the inequality $\sum_{n=0}^{\infty} \delta^n < \frac{11}{9}$.

We begin by showing that the Hausdorff limit $J = \lim_{\mathcal{H}} J_n$ exists. The collection of all compact subsets of a compact metric space, given the Hausdorff metric, is itself a compact metric space (Theorem 2.4.3). Since $\{J_n\}$ is a sequence of compact sets in a bounded region of a proper metric space, to show that the sequence converges with respect to the Hausdorff metric, it suffices to show that the sequence is Cauchy.

One bound follows by construction: $J_{n+m} \subset N(J_n, \frac{11}{9}\epsilon\delta^n)$ for all $m \geq 0$. For the second bound, take J_{n+m} and split it into subarcs of diameter at most $\epsilon\delta^n$, and write this as $J_{n+m} = J_{n+m}[z_0, z_1] \cup \cdots \cup J_{n+m}[z_{k-1}, z_k]$ for some $z_0, \ldots, z_k$ and some $k > 0$. Our coarse maps compose to give $p : J_{n+m} \to J_n$, showing that J_{n+m} $\frac{11}{9}\epsilon\delta^n$-follows J_n. Furthermore, since $d(z_i, z_{i+1}) \leq \epsilon\delta^n$, we have $d(p(z_i), p(z_{i+1})) \leq 4\epsilon\delta^n \leq s\epsilon\delta^{n-1}$. Combining this with the fact that p maps endpoints to endpoints, for $n \geq 2$ we have

$$\begin{aligned} J_n &= J_n[p(z_0), p(z_1)] \cup \cdots \cup J_n[p(z_{k-1}), p(z_k)] \subset N(\{p(z_0), \ldots, p(z_k)\}, S\epsilon\delta^{n-1}) \\ &\subset N\left(J_{n+m}, \frac{11}{9}\epsilon\delta^n + S\epsilon\delta^{n-1}\right). \end{aligned}$$

Taken together, these bounds give $d_{\mathcal{H}}(J_n, J_{n+m}) \leq \frac{11}{9}\epsilon\delta^n + S\epsilon\delta^{n-1}$, so $\{J_n\}$ is Cauchy and the limit $J = \lim_{\mathcal{H}} J_n$ exists. Moreover, J is compact (by definition) and connected (because each J_n is connected).

Now we let a_n, b_n denote the endpoints of J_n. Since $p_n(a_n) = a_{n-1}$, and p_n displaces points at most $\epsilon\delta^n$, the sequence $\{a_n\}$ is Cauchy and hence converges to some point $a \in J$. Similarly, $\{b_n\}$ converges to a point $b \in J$.

There are two cases to consider. If $a = b$, then $d(a_n, b_n) \leq 2\frac{11}{9}\epsilon\delta^n \leq s\epsilon\delta^{n-1}$. Consequently, $\text{diam}(J_n) \leq S\epsilon\delta^{n-1}$, $J = \lim_{\mathcal{H}} J_n$ has diameter zero, and thus $J = \{a\}$. Otherwise, $a \neq b$ and so J is non-trivial. We claim that in this case J is a local quasi-arc.

To show J is an arc with endpoints a and b it suffices to demonstrate that every point $x \in J \setminus \{a, b\}$ is a cut point [HY61, Theorems 2-18 and 2-27]. The topology of J_n induces an order on J_n with least element a_n and greatest b_n. Given $x \in J$, we define three points $h_n(x)$, x_n and $t_n(x)$ that satisfy $a_n < h_n(x) < x_n < t_n(x) < b_n$, where x_n is chosen such that $d(x, x_n) \leq \frac{11}{9}\epsilon\delta^n$, and $h_n(x)$ and $t_n(x)$ are the first and last elements of J_n at distance $(S+1)\epsilon\delta^{n-1}$ from x. As long as x is not equal to a or b, for n greater than some n_0 these points will exist and this definition will be valid.

We shall denote the $\frac{11}{9}\epsilon\delta^n$-neighborhoods of $J_n[a_n, h_n(x)]$ and $J_n[t_n(x), b_n]$ by $H_n(x)$ and $T_n(x)$ respectively, and define $H(x) = \cup\{H_n(x) : n \geq n_0\}$ (the Head) and $T(x) = \cup\{T_n(x) : n \geq n_0\}$ (the Tail). By definition, $H(x)$ and $T(x)$ are open. We claim that, in addition, they are disjoint and cover $J \setminus \{x\}$, and so x is a cut point.

Fix $y \in J$, and suppose $y \notin H(x) \cup T(x)$. We want to show that $y = x$. To this end, we assume that $n \geq 3$ and bound the diameter of $J_n[h_n(x), t_n(x)]$ using J_{n-1}. Because $d(p_n(h_n(x)), p_n(t_n(x))) \leq 2\epsilon\delta^{n-1} + 2(S+1)\epsilon\delta^{n-1} \leq s\epsilon\delta^{n-2}$, we know that the diameter of $J_{n-1}[p_n(h_n(x)), p_n(t_n(x))]$ must be less than $S\epsilon\delta^{n-2}$. Thus the diameter of $J_n[h_n(x), t_n(x)]$ is less than $S\epsilon\delta^{n-2} + 2\epsilon\delta^{n-1}$, as J_n $\epsilon\delta^{n-1}$-follows J_{n-1}.

For every $n \geq n_0$, y is $\frac{11}{9}\epsilon\delta^n$ close to some $y_n \in J_n$. Since $y \notin H(x) \cup T(x)$, y_n

must lie in $J_n[h_n(x), t_n(x)]$, so

$$\begin{aligned} d(x,y) &\le d(x, J_n[h_n(x), t_n(x)]) + \operatorname{diam}(J_n[h_n(x), t_n(x)]) + d(y_n, y) \\ &\le 2\frac{11}{9}\epsilon\delta^n + (S+2\delta)\epsilon\delta^{n-2} = \left(2\frac{11}{9}\delta^2 + S + 2\delta\right)\epsilon\delta^{n-2}, \end{aligned}$$

therefore $d(x,y) = 0$ and $J \setminus (H(x) \cup T(x)) = \{x\}$.

We now show that $H(x)$ and $T(x)$ are disjoint. If not, then $H_n(x) \cap T_m(x) \neq \emptyset$ for some n and m. It suffices to assume $n \le m$. Now $T_m(x) \subset N(J_m[x_m, b_m], \frac{11}{9}\epsilon\delta^m)$ by definition. We send J_m to J_n using $f = p_{n+1} \circ \cdots \circ p_m : J_m \to J_n$, to get that $T_m(x) \subset N(J_n[f(x_m), b_n], 3\epsilon\delta^n)$. Since

$$d(f(x_m), x_n) \le d(f(x_m), x_m) + d(x_m, x) + d(x, x_n) < 4\epsilon\delta^n < s\epsilon\delta^{n-1},$$

we have, even for $n = m$,

$$T_m(x) \subset N(J_n[x_n, b_n], 3\epsilon\delta^n) \cup B(x_n, (S+3\delta)\epsilon\delta^{n-1}).$$

Since $(S+3\delta)\epsilon\delta^{n-1} + \frac{11}{9}\epsilon\delta^n < (S+\frac{1}{2})\epsilon\delta^{n-1}$, $H_n(x)$ cannot meet $T_m(x)$ in the ball $B(x_n, (S+3\delta)\epsilon\delta^{n-1})$. Thus $H_n(x) \cap T_m(x) \neq \emptyset$ implies that there exist points p and q in J_n such that $a_n \le p \le h_n(x) < x_n \le q \le b_n$ and $d(p,q) < 3\epsilon\delta^n < s\epsilon\delta^{n-1}$. But then we know that $J_n[p,q]$ has diameter less than $S\epsilon\delta^{n-1}$, while containing both $h_n(x)$ and x_n. This contradicts the definition of $h_n(x)$, so $H(x) \cap T(x) = \emptyset$.

We have shown that J is an arc with endpoints a and b; it remains to show that J is a local quasi-arc with the required constants.

Say we are given x and y in J, with x_n and y_n as before. Our argument shows that the segments $J_n[x_n, y_n]$ converge to some arc $\tilde{J}[x,y]$, because $J_{n+1}[x_{n+1}, y_{n+1}]$ $(\epsilon\delta^n + S\epsilon\delta^{n-1})$-follows $J_n[x_n, y_n]$ for all $n \ge 2$. This arc $\tilde{J}[x,y]$ must lie in J, therefore $\tilde{J}[x,y]$ must equal $J[x,y]$. Now, suppose that $d(x,y) \in (\epsilon\delta^{n+1}, \epsilon\delta^n]$ holds for some $n \ge 2$. Then $d(x_n, y_n) \le 3\epsilon\delta^n + \epsilon\delta^n < s\epsilon\delta^{n-1}$, and so the subarc $J[x,y]$, which lies in $N(J_n[x_n, y_n], \frac{11}{9}\epsilon(\delta^n + S\delta^{n-1}))$, has diameter less than $S\epsilon\delta^{n-1} + 3\epsilon(\delta^n + S\delta^{n-1}) \leq \frac{4S+3\delta}{\delta^2} d(x, y)$, as desired.

Furthermore, this same argument gives that, for all $n \geq 2$, J $\frac{11}{9}\epsilon(\delta^n + S\delta^{n-1})$-follows J_n, which itself $\frac{11}{9}\epsilon\delta$-follows $J_1 = A$. Taking n sufficiently large, we have that J ϵ-follows A. □

Let us remark that a similar, but simpler, proof gives the following statement:

Proposition 3.2.3. *If X_n, for $n = 1, 2, \ldots$, are each ϵ-local λ-quasi-arcs, and $X = \lim_{GH} X_n$ is a compact metric space, then X is an ϵ'-local λ-quasi-arc, for all $\epsilon' < \epsilon$.*

Proof. First, by Example 2.4.5 and the remark following Definition 3.1.1, we have that X is ϵ'-locally λ-linearly connected as a metric space, for all $\epsilon' < \epsilon$.

Second, to show X is an arc, it suffices to show that all points except two are cut points. This is proven in a similar way as in Lemma 3.2.2.

Let $f_n : X_n \to X$ be maps distorting distance by at most δ_n, and with image that is δ_n-dense, where $\lim_{n\to\infty} \delta_n = 0$. Since X is compact, we can (after possibly choosing subsequences and adjusting δ_n) assume that the endpoints a_n and b_n of X_n have images under f_n at a distance of at most δ_n to points a and b in X.

Assume $x \in X$ is not equal to a or b. To show that x is a cut point, it suffices to show that there exists a constant $D > 0$ so that for all sufficiently small δ, every δ-chain from a to b passes through $B(x, D\delta)$. Recall that a δ-chain from a to b is a sequence of finitely many points, beginning with a and ending with b, where the distance between subsequent points is at most δ.

Now, x is δ_n close to the image of some $x_n \in X_n$. Since X_n is a ϵ-local λ-quasi-arc, if $\eta < \epsilon$ then any η-chain from a_n to b_n will come within a distance of $\lambda\eta$ from x_n. If we have a δ-chain $\mathcal{C}$ in X that joins a to b, then f_n lifts this to a $(\delta + 3\delta_n)$-chain joining a_n to b_n, therefore this chain must be at most $(\delta + 3\delta_n)\lambda$ from x_n, and so $\mathcal{C}$ is within $(\delta + 3\delta_n)\lambda + \delta_n$ of x.

Choosing n large enough so that $\delta_n \leq \delta$, and taking $D = 4\lambda + 1$, we are done. $\square$

3.3 Discrete paths and straightened arcs

The proof of Proposition 3.2.1 is based on a quantitative version of a simple geometric result. Before we state this result, recall that a maximal r-separated set $\mathcal{N}$ is a subset of X such that for all distinct x, $y \in \mathcal{N}$ we have $d(x, y) \geq r$, and for all $z \in X$ there exists some $x \in \mathcal{N}$ with $d(z, x) < r$.

Now suppose that we are given a maximal r-separated set $\mathcal{N}$ in an L-linearly connected, N-doubling, complete metric space X. Then we can find a collection of sets $\{V_x\}_{x \in \mathcal{N}}$ so that each V_x is a connected union of finitely many arcs in X, and for all x, $y \in \mathcal{N}$:

(1) $d(x, y) \leq 2r \implies y \in V_x$.

(2) $\operatorname{diam}(V_x) \leq 5Lr$.

(3) $V_x \cap V_y = \emptyset \implies d(V_x, V_y) > 0$.

For $x \in \mathcal{N}$, we can construct each V_x by defining it to be the union of finitely many arcs joining x to each $y \in \mathcal{N}$ with $d(x, y) \leq 2r$. By linear connectedness, we can ensure that $\operatorname{diam}(V_x) \leq 4Lr$. Condition (3) is trivially satisfied for compact subsets of a metric space, but we will strengthen it to the following:

(3′) $V_x \cap V_y = \emptyset \implies d(V_x, V_y) > \delta r$.

Lemma 3.3.1. *We can construct the sets V_x satisfying (1), (2) and (3′) for $\delta = \delta(L, N)$.*

Proof. Without loss of generality, we can rescale the metric to set $r = 1$.

Since X is doubling, the maximum number of 1-separated points in a $20L$-ball is bounded by a constant $M = M(20L, N)$. We can label every point of $\mathcal{N}$ with an integer between 1 and M, such that no two points have the same label if they are separated by a distance less than $20L$.

To find this labeling, consider the collection of all such labelings on subsets of $\mathcal{N}$ under the natural partial order. A Zorn's Lemma argument gives the existence of a maximal element: our desired labeling. So $\mathcal{N}$ is the disjoint union of $20L$-separated sets $\mathcal{N}_1, \ldots, \mathcal{N}_M$.

Now let $\mathcal{N}_{\leq n} = \cup_{k=1}^{n} \mathcal{N}_k$, and consider the following

Claim $\Delta(n)$. *We can find V_x for all $x \in \mathcal{N}_{\leq n}$, such that for all $x, y \in \mathcal{N}_{\leq n}$ (1), (2) and (3′) are satisfied with $\delta = \frac{1}{2}(5L)^{-n}$.*

$\Delta(0)$ holds trivially, and Lemma 3.3.1 immediately follows from $\Delta(M)$, with $\delta = \delta(L, N) = \frac{1}{2}(5L)^{-M}$. So we are done, modulo the statement that $\Delta(n) \implies \Delta(n+1)$ for $n < M$. $\square$

Proof that $\Delta(n) \implies \Delta(n+1)$, for $n < M$. By $\Delta(n)$, we have sets V_x for all x in $\mathcal{N}_{\leq n}$.

As $\mathcal{N}_{n+1}$ is $20L$-separated we can treat the constructions of V_x for each $x \in \mathcal{N}_{n+1}$ independently. We begin by creating a set $V_x^{(0)}$ that is the union of finitely many arcs joining x to its 2-neighbors in $\mathcal{N}$. We can ensure that $\operatorname{diam}(V_x^{(0)}) \leq 4L$.

Now construct $V_x^{(i)}$ inductively, for $1 \leq i \leq n$. $V_x^{(i-1)}$ can be $5L$-close to at most one $y \in \mathcal{N}_i$. If $d(V_x^{(i-1)}, V_y) \in (0, \frac{1}{2}(5L)^{-i})$, then define $V_x^{(i)}$ by adding to $V_x^{(i-1)}$ an arc of diameter at most $L(5L)^{-i}$ joining $V_x^{(i)}$ to V_y. Otherwise, let $V_x^{(i)} = V_x^{(i-1)}$. Continue until $i = n$ and set $V_x = V_x^{(n)}$.

Note that V_x satisfies (1) and (2) by construction. The only non-trivial case we need to check for (3′) is whether $d(V_x, V_y) \in (0, \frac{1}{2}(5L)^{-n})$ for some $y \in \mathcal{N}_i$, $i < n$. (The $i = n$ case follows from the last step of the construction.) Then, since $V_x = V_x^{(n)} \supset V_x^{(i)}$, $V_x^{(i)} \cap V_y \neq \emptyset$, and $d(V_x^{(i)}, V_y) \geq \frac{1}{2}(5L)^{-i}$. The construction then implies that

$$\begin{aligned} d(V_x, V_y) &\geq \frac{1}{2}(5L)^{-i}(1 - (2L)(5L)^{-1} - (2L)(5L)^{-2} - \cdots - (2L)(5L)^{-(n-i)}) \\ &> \frac{1}{2}(5L)^{-n}(5L)\left(1 - \frac{2/5}{1-(1/(5L))}\right) \geq \frac{5}{2}\left(\frac{1}{2}(5L)^{-n}\right), \end{aligned}$$

contradicting our assumption, so $\Delta(n+1)$ holds. □

We now finish by using this construction to prove our proposition.

Proof of Proposition 3.2.1. By rescaling the metric, we may assume that $\iota = 20L$. If the endpoints a and b of A satisfy $d(a, b) \leq 20 = \frac{\iota}{L}$, then join a to b by an arc of diameter less than ι. This arc will, trivially, satisfy our conclusion for any $S \geq 1$.

Otherwise, $d(a, b) > 20$. In the hypotheses for Lemma 3.3.1, let $r = 1$ and let $\mathcal{N}$ be a maximal 1-separated set in X that contains both a and b. Now apply Lemma 3.3.1 to get $\{V_x\}_{x\in\mathcal{N}}$ satisfying (1), (2) and (3′) for $\delta = \delta(L, N) > 0$.

We want to 'discretize' A by finding a corresponding sequence of points in $\mathcal{N}$. Now, every open ball in X meets the arc A in a collection of disjoint, relatively open intervals. Since $\mathcal{N}$ is a maximal 1-separated set, the collection of open balls $\{B(x, 1) : x \in \mathcal{N}\}$ covers X; in particular, it covers A. By the compactness of A, we can find a finite cover of A by connected, relatively open intervals, each lying in some $B(x, 1)$, $x \in \mathcal{N}$.

Using this finite cover, we can find points $x_i \in \mathcal{N}$ and $y_i \in A$ for $0 \leq i \leq n$, such that $a = y_0 < \cdots < y_n = b$ in the order on A, and $A[y_i, y_{i+1}] \subset B(x_i, 1)$ for each $0 \leq i < n$. Since $a, b \in \mathcal{N}$, we have that $x_0 = a$ and $x_n = b$. The sequence $(x_0, \ldots, x_n)$ is a discrete path in $\mathcal{N}$ that corresponds naturally to A.

We now find a subsequence (x_{r_j}) of (x_i) such that the corresponding sequence of sets $(V_{x_{r_j}})$ forms a 'path' without repeats. Let $r_0 = 0$, and for $j \in \mathbb{N}^+$ define r_j inductively as $r_j = \max\{k : V_{x_k} \cap V_{x_{r_{j-1}}} \neq \emptyset\}$, until $r_m = n$ for some $m \leq n$. The integer r_j is well defined since $d(y_{(r_{j-1}+1)}, x_k) \leq 1$ for $k = r_{j-1}$ and $k = r_{j-1} + 1$, so $V_{x_{(r_{j-1}+1)}} \cap V_{x_{r_{j-1}}} \neq \emptyset$. Note that if $i + 2 \leq j$ then $V_{x_{r_i}} \cap V_{x_{r_j}} = \emptyset$, and thus $d(V_{x_{r_i}}, V_{x_{r_j}}) > \delta$.

Let us construct our arc J in segments. First, let $z_0 = x_{r_0}$. Second, for each i from 0 to $m - 1$, let $J_i = J_i[z_i, z_{i+1}]$ be an arc in $V_{x_{r_i}}$ that joins $z_i \in V_{x_{r_i}}$ to some $z_{i+1} \in V_{x_{r_{i+1}}}$, where z_{i+1} is the first point of J_i to meet $V_{x_{r_{i+1}}}$. (In the case $i = m-1$, join z_{m-1} to $x_{r_m} = z_m$.) Set $J = J_0 \cup \cdots \cup J_m$.

This path J is an arc since each $J_i \subset V_{x_{r_i}}$ is an arc, and if there exists a point $p \in J_i \cap J_j$ for some $i < j$, then $j = i + 1$ and $p = z_{i+1} = z_j$. This is true because $V_{x_{r_i}} \cap V_{x_{r_j}} \neq \emptyset$ implies that $j = i + 1$, and the definition of z_{i+1} implies that $J_i \cap V_{x_{r_{i+1}}} = \{z_{i+1}\}$. Any finite sequence of arcs that meet only at consecutive endpoints is also an arc, so we have that J is an arc.

In fact, J satisfies $(*)$. For any y, $y' \in J$, $y < y'$, we can find $i \leq j$ such that $z_i \leq y < z_{i+1}$, $z_j \leq y' < z_{j+1}$. (If $y = z_m$, set $i = m$; likewise for y'.) If $d(y, y') \leq \delta$ then, because $y \in V_{x_{r_i}}$ and $y' \in V_{x_{r_j}}$, we have $d(V_{x_{r_i}}, V_{x_{r_j}}) \leq \delta$, so either $j = i$ or $j = i + 1$. This gives that $J[y, y'] \subset V_{x_{r_i}} \cup V_{x_{r_j}}$, and so $\operatorname{diam}(J[y, y'])$ is bounded above by $10L$.

Furthermore, J ι-follows A. There is a coarse map $f : J \to A$ defined by the following composition: first map J to $\mathcal{N}$ by sending $y \in J[z_i, z_{i+1}) \subset J$ to $x_{r_i} \in \mathcal{N}$, and sending x_{r_m} to itself. Second, map each x_{r_i} to the corresponding y_{r_i} in A. Taking arbitrary $y < y'$ in J as before, we see that

$$J[y,y'] \subset J[z_i, z_{j+1}] \subset N(\{x_{r_i}, \ldots, x_{r_j}\}, 5L) \subset N(\{y_{r_i}, \ldots, y_{r_j}\}, 5L+1)$$
$$\subset N(A[y_{r_i}, y_{r_j}], 5L+1) \subset N(A[f(y), f(y')], \iota).$$

We let $s = \frac{1}{20L}\delta$ and $S = \frac{1}{20L}10L$, and have proven the Proposition. □

3.4 Quantitative estimates

Unlike Tukia's original proof, our proof of Theorem 3.1.4 allows us to give explicit estimates for the constants α and λ. (Recall that we found an $\alpha\epsilon$-local λ-quasi-arc in the ϵ neighborhood of any arc A.) These estimates are certainly not optimal, but since they may be of interest we collect them in this section.

First, if X is N-doubling, then the maximum number of r-separated points in a ball B of radius R is at most N. For such a maximal set of points S, the collection of balls of radius $\frac{R}{2}$ centered at S forms a cover of B. Iterating this argument allows us to estimate that the maximum number of $\frac{R}{2^n}$ separated points in a R ball is less than N^{1+n}.

We use this to see that the constant $M(20L, N)$, which describes the maximum number of 1-separated points in a $20L$ ball, may be taken to equal

$$M(20L, N) = N^{1+\lceil \log(20L)/\log(2) \rceil}.$$

Therefore, in Lemma 3.3.1 we take

$$\delta(L,N) = \frac{1}{2}(5L)^{-M} = \frac{1}{2}(5L)^{-N^{1+\lceil \log(20L)/\log(2)\rceil}}.$$

Finally, since $S = \frac{1}{2}$ and $s = \frac{\delta}{20L}$, we have that

$$\alpha = \delta^2 \quad \text{and} \quad \lambda = \frac{2+3\delta}{\delta^2}.$$

CHAPTER IV

Conformal dimension bounds

In this chapter we shall prove the following lower bound on conformal dimension:

Theorem 4.0.1. *Suppose* (X, d) *is a complete metric space which is doubling and annulus linearly connected. Then the conformal dimension* $\dim_{\mathcal{C}}(X)$ *is at least* $C > 1$*, where* C *depends only on the data of* X*, i.e., the constants associated to the two conditions above.*

The annulus linearly connected condition is a quantitative analogue of the topological conditions of being locally connected and having no local cut points; see Definition 4.1.2 and the subsequent discussion. For now, a good example of such a space is the standard square Sierpiński carpet.

It is easy to find a family of curves in the Sierpiński carpet that allow us to apply Pansu's lemma on conformal dimension, as we saw in Example 2.1.2. We will use the lemma to prove Theorem 4.0.1, but to do so we have to construct a family of arcs in X akin to the product of an interval and a regular Cantor set (of controlled dimension).

To see how this may be done, let us consider a topological analogue. Let X be a connected, locally connected, complete metric space. It is well known that such a space is arc-wise connected. If we now further assume that X has no local cut points then a topological argument shows that the product of a Cantor set and the unit interval embeds homeomorphically into X. A weaker statement is that there exists a collection of arcs $\{J_\sigma\}$ in X such that, under the topology induced by the Hausdorff metric, the set $\{J_\sigma\}$ is a topological Cantor set.

We will show a quantitatively controlled analogue of this weaker statement. First, recall that $(\mathcal{M}(X), d_{\mathcal{H}})$ is the set of all closed subsets of X endowed with the Hausdorff metric, and it is complete (Theorem 2.4.3). For each $\sigma > 0$, Z_σ denotes a standard Ahlfors regular Cantor set of Hausdorff dimension σ. (Details are given in Section 4.1.)

Theorem 4.0.2. *For all $L \geq 1$ and $N \geq 1$, there exist $C' \geq 1$, $\sigma > 0$ and $\lambda' \geq 1$ such that if X is an L-annulus linearly connected, N-doubling, complete metric space of diameter at least one, then there exists a C'-bi-Lipschitz embedding of Z_σ into $\mathcal{M}(X)$, where each point in the image is a λ'-quasi-arc of diameter at least $\frac{1}{C'}$.*

So, how do we create such a good collection of arcs? First, use the topological properties of the space to split one arc into two arcs and apply Tukia's theorem (Theorem 3.1.4) to straighten these arcs into uniformly local quasi-arcs. Second, repeat this procedure in a controlled way by using the compactness properties of the quasi-arcs and spaces. This process gives four arcs, then eight, and so on, limiting to a collection of arcs indexed by a Cantor set. We describe this in detail in Section 4.1.

Once we have proven Theorem 4.0.2 it is straightforward to prove Theorem 4.0.1 using Pansu's lemma. We do so in Section 4.2.

For more discussion on conformal dimension we refer the reader to the Bonk and Kleiner paper [BK05a]. Incidentally, they work with the Ahlfors regular conformal dimension, where the Hausdorff dimension is infimized over all quasisymmetrically equivalent metric spaces that are *Ahlfors regular*. However, this dimension is bounded below by the conformal dimension so our lower bounds apply to the Ahlfors regular conformal dimension as well.

4.1 Unzipping arcs

Consider a complete locally connected metric space with no local cut points, i.e., no connected open set is disconnected by removing a point. In such a space it is straightforward to 'unzip' a given arc A into two disjoint arcs J_1 and J_2 lying in a specified neighborhood of A. Repeating this procedure to get four arcs, then eight, and so on, it is possible, with some care, to get a limiting set homeomorphic to the product of a Cantor set and the interval. Such a limit set is useless for our purposes because there is no control on the minimum distance between two unzipped arcs, and so no way to get a lower bound on conformal dimension that is greater than one. We will use compactness type arguments to overcome this problem.

We begin by proving the topological unzipping result.

Lemma 4.1.1. *Given an arc A in a complete, locally connected metric space with no local cut points, and $\epsilon > 0$, it is possible to find two disjoint arcs J_1 and J_2 in $N(A, \epsilon)$ such that the endpoints of J_i are ϵ-close to the endpoints of A. Furthermore, the arcs J_i ϵ-follow the arc A.*

Proof. Here, $B_0(x, r)$ will denote the connected component of an open ball $B(x, r) \subset X$ that contains its center x. As X is locally connected, $B_0(x, r)$ is always open and connected, and, moreover, $B_0(x, r) \setminus \{x\}$ is also open and connected because x is not a local cut point. Any open and connected subset of X is arcwise connected.

Let a and b be the initial and final points of A respectively (in a fixed order given by the topology). We are going to define J_1 and J_2 inductively. There exists $w \in B_0(a, \frac{1}{2}\epsilon) \setminus A$, otherwise there would be a open set in X homeomorphic to an

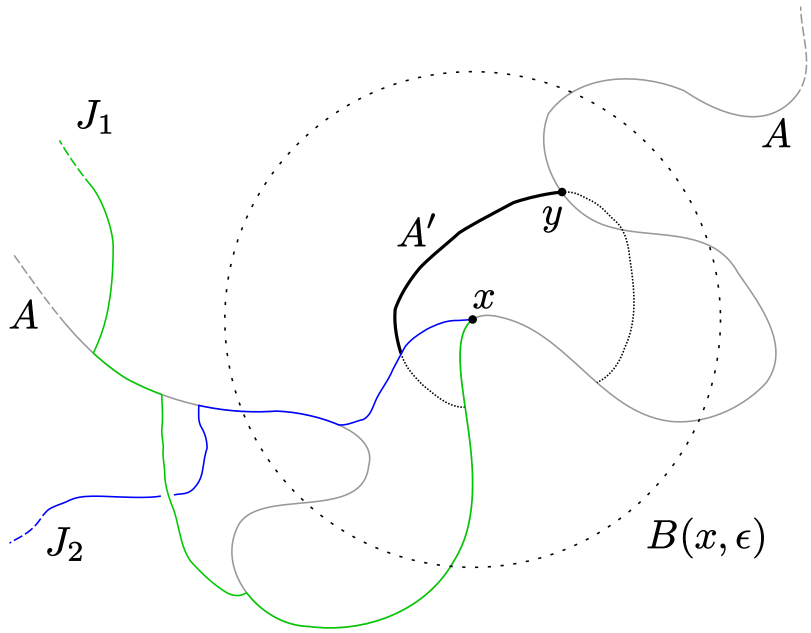

Figure 4.1: Unzipping an arc

arc segment, violating the no local cut point condition. Now join w to a by an arc in $B_0(a, \frac{1}{2}\epsilon)$. Stop this arc at x, the first time it meets A, and call it $J_1 = J_1[w, x]$. Set $J_2 = A[a, x]$. (Perhaps $x = a$, but this is not a problem).

Now we have two head segments for J_1 and J_2 meeting only at $x \in A$, and we want to 'unzip' this configuration further along A. This is possible since in $B_0(x, \epsilon)$ there is a tripod type configuration with two incoming arcs J_1 and J_2 and one outgoing arc A. As noted above, $B_0(x, \epsilon) \setminus \{x\}$ is arcwise connected, and so we can find an arc in this set that joins some point in J_1 (not x) to a point in A (not x). The arc may meet J_1, J_2 and A in many places but there must be some sub-arc A' joining some point in J_1 or J_2 to some point y in A with interior disjoint from them all. (See Figure 4.1, where A' is emphasized.) Use A' to detour one of J_1 and J_2 around x to the new unzipping point y, and extend the other J_i to y using $A[x, y]$.

What if this unzipping process approaches a limit before we are ϵ-close to the final point b in A? This cannot happen. Suppose it is not possible to unzip past $z \in A$.

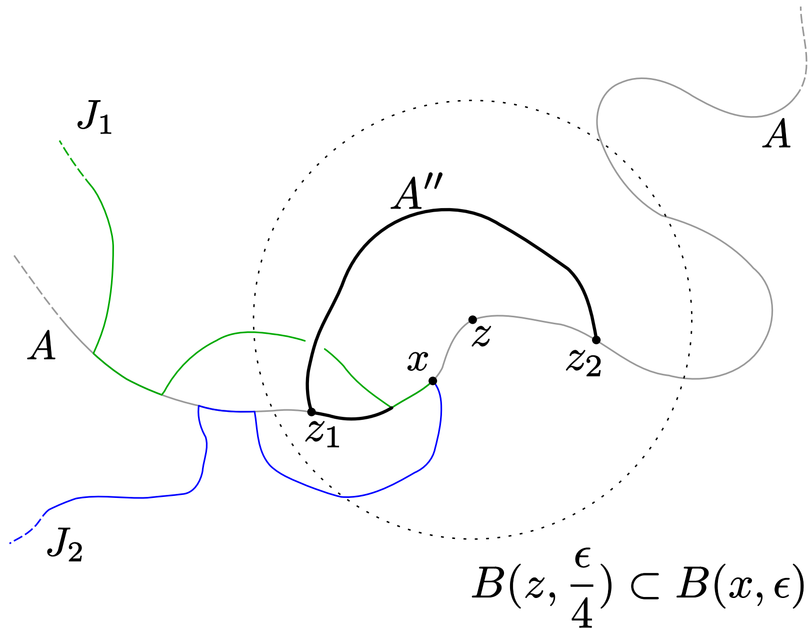

Figure 4.2: Avoiding a limit point

Since $B_0(z, \frac{\epsilon}{4}) \setminus \{z\}$ is arcwise connected, inside this set we can construct an arc A'' that detours around z, from $z_1 \in A$ to $z_2 \in A$, where $z_1 < z < z_2$ in the order on A.

Now by the limit point hypothesis, we can unzip J_1 and J_2 past z_1 to x, where $z_1 < x < z$. To continue the construction of J_1 and J_2 past z, find the arc given by following z_2 to z_1 along A'', stopping if one of J_1 or J_2 is met. If we reach z_1 without intersecting J_1 or J_2, as is the case in Figure 4.2, then continue to follow A from z_1 towards z. By the construction of J_1 and J_2, this arc will meet J_1 or J_2 before reaching z. In either case, this arc can be used as a legitimate detour around x and z, contradicting the assumption on z. Thus it is possible to continue unzipping until $x \in B(b, \frac{\epsilon}{2})$.

Label each point of J_1 and J_2 by the point $x \in A$ they were used to detour round. Call the resulting labeling a map $f_i : J_i \to A$ and this will coarsely preserve order as desired. $\square$

We would like to give a lower bound for the distance between the two split arcs.

To do this we need a quantitative metric version of the no local cut points condition as having no local cut points is not preserved under Gromov-Hausdorff convergence. (Neither is the standard LLC(2) condition.)

Definition 4.1.2. We say a metric space X is *(L-)annulus linearly connected* for some $L \geq 1$ if it is L-linearly connected and, in addition, given $r > 0$ and three distinct points p, x and y in X such that x and y lie in the annulus $A(p, r, 2r)$, then there is an arc J joining x to y that lies in the annulus $A(p, \frac{r}{L}, 2Lr)$.

There are many possible equivalent formulations of this condition. Its key features are that if X_i are L-annulus linearly connected and $X_i \to X_\infty$ in the Gromov-Hausdorff topology then X_∞ is also L-annulus linearly connected. Furthermore, annulus linearly connected implies no local cut points.

We do need a stronger condition than no local cut points as a hypothesis for Theorem 4.0.1: it is straightforward to modify the Sierpiński carpet construction to get a doubling, linearly connected, complete metric space with no local cut points whose *Hausdorff* dimension is one, therefore by (2.8) the conformal dimension is also one.

Now for the remainder of this section we will assume that L and N are fixed constants, and $\lambda \geq 1$, $\alpha \in (0, 1]$ are as given by Theorem 3.1.4. Consider the collection $\mathcal{C}$ of all λ-quasi-arcs A in any complete metric space X that is L-annulus linearly connected and N-doubling, and whose endpoints a and b satisfy $d(a, b) \in [\frac{1}{R}, R]$ for some $R \geq 1$. Fix $\epsilon > 0$, and consider the supremum of possible separations of two arcs split from A by the topological lemma above. Call this δ_A ($\delta_A > 0$).

Lemma 4.1.3. *There exists $\delta^\star = \delta^\star(\lambda, L, N, \epsilon, R) > 0$ such that for all $A \in \mathcal{C}$, $\delta_A > \delta^\star$.*

Proof. If not, then we can find a sequence of arcs $A_i \subset X_i$ such that $\delta_{A_i} < \frac{1}{i}$. Let a_i and b_i denote the endpoints of A_i. We are only interested in what happens inside the ball $B_i := B(a_i, 2R(2L+\epsilon))$. As the sequence of configurations $(B_i \subset X_i, A_i, a_i, b_i)$ is precompact in the Gromov-Hausdorff topology, we can take a subsequence so that $B_i \to B_\infty$, and (inside B_∞) $A_i \to A_\infty$; this will also be a λ-quasi-arc. This means that there exist constants $C_i \to 0$ and maps $f_i : B_\infty \to B_i$ such that f_i distorts distances by an additive error of at most C_i, and every point of B_i is within C_i of $f_i(B_\infty)$. Furthermore, $f_i(A_\infty) \subset A_i$, $f_i(a_\infty) = a_i$ and $f_i(b_\infty) = b_i$.

Since B_∞ will be L-annulus linearly connected (away from the edge of the ball), it will have no local cut points in its interior. Consequently, we can split A_∞ into two arcs J_1 and J_2 using Lemma 4.1.1 inside an $\frac{\epsilon}{3}$-neighborhood of A_∞. These arcs are disjoint so they are separated by some distance $0 < \delta' \le \frac{\epsilon}{3}$. The remainder of the proof consists of showing that this contradicts the assumption on $A_i \subset B_i$ for some large i.

For sufficiently large i, $C_i \le \frac{\delta'}{8L}$ because $C_i \to 0$ as $i \to \infty$. For $j = 1, 2$, the arc J_j in B_∞ contains a discrete path D_j with C_i sized jumps that corresponds to a discrete path $D'_j = f_i(D_j)$ in X_i with $2C_i \le \frac{\delta'}{4L}$ jumps. The L-linearly connected condition can then be used to join each D'_j up into a continuous arc J'_j.

To be precise, if $D'_j = \{p_1, \ldots, p_M\}$, join p_1 to p_2 by an arc J'_j of diameter at most $2C_iL \le \frac{\delta'}{4}$. Assume that at a stage k we have an arc J'_j from p_1 to p_k. There is an arc I of diameter at most $\frac{\delta'}{4}$ joining p_{k+1} to p_k. We extend J'_j to p_{k+1} by following I from p_{k+1} to p_k, stopping at x, the first time it meets J'_j, and gluing together $J'_j[p_1, x]$ and $I[x, p_{k+1}]$ to make a new arc J'_j, and repeat until $k = M$. Define a map $h_j : J'_j \to D'_j$ that sends each of the points added at stage k to the point p_k. Note that for all x, $y \in J'_j$, $J'_j[x,y] \subset N(D'_j[h_j(x), h_j(y)], \frac{\delta'}{4})$; in a coarse sense, J'_j $\frac{\delta'}{4}$-follows D'_j.

By construction, for sufficiently large i, J_1' and J_2' are $\frac{\delta'}{4}$-separated and ϵ-close to A_i, but to get a contradiction we need that they ϵ-follow A_i.

Since A_∞ and A_i are both λ-quasi-arcs, Fact 4.1.4 below implies that for all $x,\ y \in A_\infty$, $f_i(A_\infty[x,y])$ is contained in the $((2C_i\lambda + C_i)\lambda + C_i)$-neighborhood of $A_i[f_i(x), f_i(y)]$. For each j, we can lift the map $h_j : J_j' \to D_j'$ to a map $h_j' : J_j' \to D_j \subset B_\infty$. By Lemma 4.1.1, D_j $\frac{\epsilon}{3}$-follows A_∞, so further compose with the associated map $D_j \to A_\infty$. Finally, compose with $f_i : A_\infty \to A_i$.

The composed maps $J_j' \to D_j \to A_\infty \to A_i$, for each j, show that each J_j' follows A_i with constant $\left(\frac{\delta'}{4} + \frac{\epsilon}{3} + C_i + (2C_i\lambda + C_i)\lambda + C_i\right)$. This is smaller than ϵ for sufficiently large i because $C_i \to 0$ as $i \to \infty$. We have contradicted our initial assumption, so the proof is complete. $\square$

We used the following fact in the proof:

Fact 4.1.4. *If A and A' are λ-quasi-arcs, and $f : A \to A'$ is a map distorting distances by at most C, then for all x and y in A,*

$$f(A[x,y]) \subset N(A'[f(x), f(y)], (2C\lambda + C)\lambda + C). \tag{4.1}$$

Proof. Let $x = p_0 < p_1 < \cdots < p_n = y$ be a chain of points in A so that the diameter of $A[p_{i-1}, p_i]$ is less than C, for $i = 1, \ldots, n$.

Let $x' = f(x)$, $y' = f(y)$, and $p_i' = f(p_i)$. For some l and some m we have $p_{l-1}' \leq x' < p_l'$ and $p_m' < y' \leq p_{m+1}'$. Assume l is the greatest such and m is the least such number. Since $d(p_i', p_{i+1}') \leq 2C$, we have $d(p_l', x')$ and $d(p_m', y')$ are both less than or equal to $2C\lambda$.

This lifts, by f, to give that $d(p_l, x)$ and $d(p_m, y)$ are both less than or equal to $2C\lambda + C$, and so

$$\operatorname{diam}(A[x, p_l]) \leq (2C\lambda + C)\lambda \quad \text{and} \quad \operatorname{diam}(A[p_m, y]) \leq (2C\lambda + C)\lambda. \tag{4.2}$$

Therefore,

$$
\begin{aligned}
f(A[x,y]) &= f(A[x,p_l] \cup A[p_l,p_m] \cup A[p_m,y]) \\
&\subset N(\{x',y'\},(2C\lambda+C)\lambda+C) \cup N(A'[x',y'],2C\lambda) \\
&\subset N(A'[f(x),f(y)],(2C\lambda+C)\lambda+C). \qquad \square
\end{aligned}
$$

The important point to note in Lemma 4.1.3 was the presence of the diameter constraint R allowing us to use a compactness type technique. Without this constraint we have various problems: our sequence of counterexamples still converges in some sense, but would likely give an unbounded arc. Topological unzipping still works but the resulting arcs would not necessarily have a positive lower bound on separation.

We can deal with the problem of no diameter bounds by dividing the problem into two collections of non-interacting smaller problems. To be precise, given a λ-quasi-arc A, or even just a local λ-quasi-arc, we can use Lemma 4.1.3 on uniformly spaced out small subarcs of A (that are genuine λ-quasi-arcs) with a sufficiently small ϵ value – this is the first collection of problems.

Now the second collection of independent problems is how to join together two of these small splittings with two disjoint arcs having uniform bound on their separation – but this a problem with bounded diameter! So compactness arguments allow us to fix this and to remove the dependence of $\delta^\star$ on R in Lemma 4.1.3.

Lemma 4.1.5. *Given* $0 < \epsilon \leq \operatorname{diam}(X)$ *and an* $\alpha\epsilon$*-local* λ*-quasi-arc* A *in* X*, where* $\alpha \in (0,1]$ *is a constant, there exists* $\delta^\star = \delta^\star(\lambda, L, N, \alpha) > 0$ *such that for all* $\delta < \delta^\star$ *we can split* A *into two arcs that* ϵ*-follow* A *and that are* $\delta\epsilon$ *separated.*

Proof. Without loss of generality we can rescale to $\epsilon = 1$. As before, choose a linear order on A compatible with its topology. Let x_0 be the first point in A, and y_0 be the first point at distance $D_1 = \frac{\alpha}{5\lambda}$ from x_0. (If there is no such point, $\mathrm{diam}(A) \leq \frac{1}{5} = \frac{\epsilon}{5}$ and so we can split A into two points that are $\frac{\epsilon}{2}$ separated.) Label the next point at distance D_1 from y_0 by x_1. Continue in this manner with all jumps D_1 until the last label y_n, with $d(x_n, y_n) \in [D_1, 3D_1)$.

Let $D_2 = \frac{1}{4}D_1 = \frac{\alpha}{20\lambda}$, and $D_3 = \frac{1}{2\lambda(L\lambda+2)}D_2$. We can control the interactions of the collection of sub-arcs of types $A[x_i, y_i]$ and $A[y_i, x_{i+1}]$: the D_3 neighborhoods of two different such sub-arcs are disjoint outside the collection of balls $\{B(x_i, D_2)\} \cup \{B(y_i, D_2)\}$. This is because otherwise there are points z and z' in two different sub-arcs that satisfy $d(z, z') \leq 2D_3 < \alpha$; so the diameter of $A[z, z']$ is less than $2\lambda D_3 < \frac{1}{2}D_2$ – but $A[z, z']$ has to pass through the center of a D_2-ball that does not contain z or z', contradiction.

Now $A[x_i, y_i]$ is a λ-quasi-arc and we use Lemma 4.1.3 to create J_i and J_i' in a $\frac{1}{2}D_3$ neighborhood of $A[x_i, y_i]$ that are $\frac{1}{2}\delta_0$ separated for some $\delta_0 = \delta_0(\lambda, L, N, D_3) > 0$. By applying Theorem 3.1.4 to straighten the arcs we may assume that they are λ'-quasi-arcs in a D_3 neighborhood of $A[x_i, y_i]$ that are $\frac{1}{4}\delta_0$ separated, where $\lambda' = \lambda'(L, N, \delta_0, D_3)$.

We want to join up the pair of arcs J_i and J_i' ending in $B(y_i, D_3)$ to the arcs J_{i+1} and J_{i+1}' starting in $B(x_{i+1}, D_3)$, without altering the setup outside the set $\mathrm{Join}(i) = B(y_i, D_2) \cup N(A[y_i, x_{i+1}], D_3) \cup B(x_{i+1}, D_2)$. Figure 4.3 shows this configuration. We will do this joining in two stages; first a topological joining that keeps the arcs disjoint, and second a quantitative version that controls the separation of the arcs in the joining.

Topological joining: Join the endpoints of J_i and J_i' to the arc A in the ball $B(y_i, LD_3)$ and the endpoints of J_{i+1} and J_{i+1}' to A in the ball $B(x_{i+1}, LD_3)$. Use the topological unzipping argument of Lemma 4.1.1 to unzip A along this segment

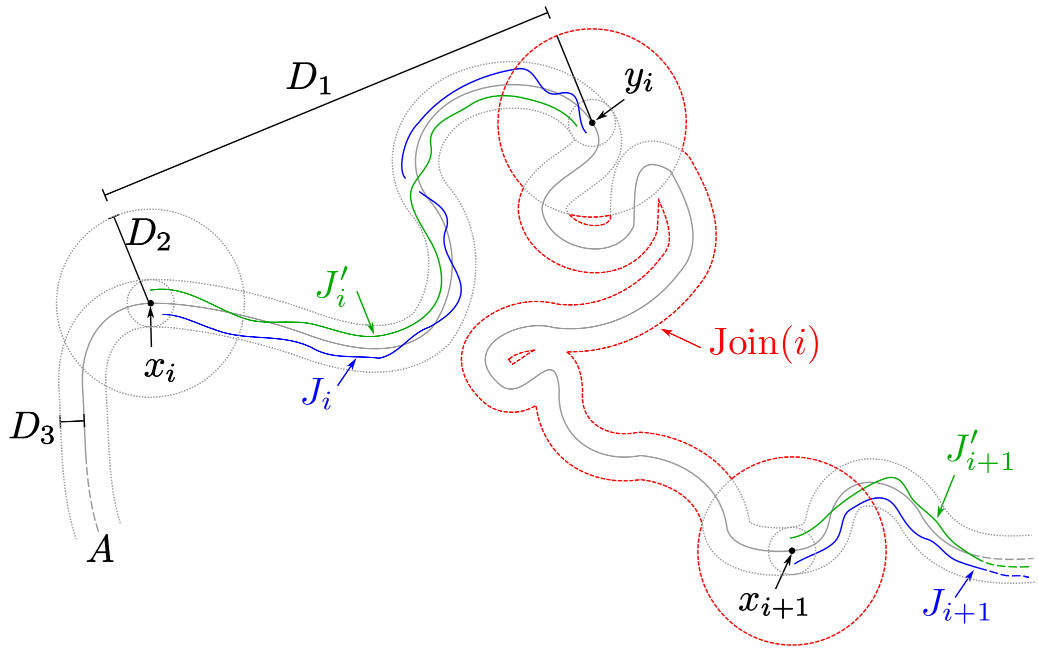

Figure 4.3: Joining unzipped arcs

resulting in 'wiring' the pair (J_i, J_i') to the pair (J_{i+1}, J_{i+1}') (not necessarily in that order) inside Join(i). These arcs are disjoint, and so separated by some distance $\delta > 0$.

Quantitative bound on δ: If there is no quantitative lower bound on δ then there are configurations (relabeling for convenience our joining arcs)

$$\mathcal{C}^n = (X^n, A^n, J_1^n, J_1'^n, J_2^n, J_2'^n),$$

where the best joining of the pair J_1^n and $J_1'^n$ to the pair J_2^n, $J_2'^n$ is at most $\frac{1}{n}$ separated.

But this configuration is precompact in the Gromov-Hausdorff topology as the X^n are all N-doubling, and the A^n and J^n arcs are uniform quasi-arcs. (This is the importance of Tukia's theorem.) So we can take a subsequence converging to a configuration $\mathcal{C}^\infty = (X^\infty, A^\infty, J_1^\infty, J_1'^\infty, J_2^\infty, J_2'^\infty)$ in a suitable ball, and join the arcs using the topological method above, giving some valid rewiring with some positive separation $\delta^\infty > 0$. Following the proof of Lemma 4.1.3 we can lift this to $\mathcal{C}^n$ for sufficiently large n retaining a separation of $\frac{1}{2}\delta^\infty > 0$: contradiction for large n.

Now since we have some $\delta^\star > 0$ to use when joining together our wirings in the disjoint collection of all Join(i), we can apply this procedure for all i to create two arcs along A that are $\delta^\star$-separated, for $\delta^\star$ depending only on λ, L, N, and α as desired. We assumed $\epsilon = 1$, but rescaling to any ϵ gives the same conclusion with our resulting arcs $\delta^\star\epsilon$ separated. □

Now we will use this lemma to create a 'Cantor set' of arcs. Recall that the space $Z = \{0,1\}^{\mathbb{N}}$ can be given the metric

$$d_\sigma((a_1, a_2, \ldots), (b_1, b_2, \ldots)) = \exp(-(\log(2)/\sigma)\inf\{n | a_n \neq b_n\}),$$

where $\sigma > 0$ is a constant, and the infimum of the empty set is positive infinity. The space (Z, d_σ) has Hausdorff dimension σ, and is Ahlfors regular since there is a Borel probability measure ν_σ on Z that satisfies $r^\sigma \leq \nu_\sigma(B(z,r)) \leq 2r^\sigma$, for all $z \in Z$ and $r < \operatorname{diam}(Z)$.

4.2 Collections of arcs and conformal dimension bounds

We now have the tools to prove Theorem 4.0.2.

Proof of Theorem 4.0.2. Begin with any arc J', assume it has endpoints 1 unit apart and apply Theorem 3.1.4 to J' and $\epsilon = \frac{1}{10}$ to get $J_\emptyset$, a λ-quasi-arc on scales below $\frac{\alpha}{10}$. Let our scaling factor be $\beta = \frac{\alpha\delta^\star}{32\lambda} \leq \frac{1}{32}$.

We can assume that for a given n we have a collection of λ-quasi-arcs on scales below β^n, written as $\{J_{a_1 a_2 \ldots a_n} | a_i \in \{0,1\}, 1 \leq i \leq n\}$, and that these arcs are β^n separated.

Now for each $J_{a_1a_2\ldots a_n}$ we split it into two arcs using Lemma 4.1.5 applied to $\epsilon = \frac{1}{8}\beta^n$, then straighten each arc using Theorem 3.1.4 with $\epsilon = \frac{\delta^\star}{32}\beta^n$ to get two new arcs $J_{a_1a_2\ldots a_n0}$ and $J_{a_1a_2\ldots a_n1}$ that are λ-quasi-arcs on scales below $\frac{\alpha\delta^\star}{32}\beta^n \geq \beta^{n+1}$, and are $\frac{\delta^\star}{16}\beta^n \geq \beta^{n+1}$ separated. In fact, all the arcs created at this stage are β^{n+1} separated as the new arcs arising from different arcs in the previous generation can only get $2\left(\frac{1}{8}\beta^n + \frac{\delta^\star}{32}\beta^n\right) < \frac{1}{2}\beta^n$ closer, still at least β^{n+1} separated.

At this point it is useful to record the following simple

Fact 4.2.1. *If J is a λ-quasi-arc on scales below ϵ, and we have an arc $J' \subset N(J, \frac{\epsilon}{4})$, whose endpoints are $\frac{\epsilon}{4}$ close to those of J, then we must have $J \subset N(J', \lambda\epsilon)$. In particular, $d_{\mathcal{H}}(J, J') \leq \lambda\epsilon$.*

Proof. To see this, fix an order on J' compatible with its topology. Take $p_0 < p_1 < \cdots < p_N$ in J' such that p_0 and p_N are the endpoints of J', and for $i = 0, \ldots, N-1$, $\mathrm{diam}(J'[p_i, p_{i+1}]) \leq \frac{\epsilon}{12}$. Associate a point $q_i \in J$ to each p_i so that $d(p_i, q_i) \leq \frac{\epsilon}{4}$, and $J = J[q_0, q_N]$. By construction,

$$d(q_i, q_{i+1}) \leq d(q_i, p_i) + d(p_i, p_{i+1}) + d(p_{i+1}, q_{i+1}) \leq \frac{7}{12}\epsilon.$$

Therefore, $\mathrm{diam}(J[q_i, q_{i+1}]) \leq \frac{7}{12}\epsilon\lambda$, and we have

$$\begin{aligned} J &= J[q_0, q_1] \cup \cdots \cup J[q_{N-1}, q_N] \\ &\subset \overline{N}(\{q_0, \ldots, q_N\}, \tfrac{7}{12}\epsilon\lambda) \\ &\subset \overline{N}(\{p_0, \ldots, p_N\}, \tfrac{10}{12}\epsilon\lambda) \\ &\subset N(J, \epsilon\lambda). \end{aligned}$$

∎

Given a sequence $a = (a_1, a_2, \ldots) \in \{0,1\}^{\mathbb{N}}$ the sequence of arcs $J_\emptyset$, J_{a_1}, $J_{a_1a_2}, \ldots$ is Cauchy in the Hausdorff metric (using Fact 4.2.1), and hence convergent to $J_{a_1a_2\ldots} = J_a$, a set of diameter at least $\frac{1}{2}$. A priori, this set need not be an arc, but only compact and connected: this is actually enough to prove Pansu's lemma. On the other hand, for each n we know that $J_{a_1 a_2 \ldots a_n}$ is a β^n-local λ-quasi-arc that β^n-follows $J_{a_1 a_2 \ldots a_{n-1}}$, and we know that $\beta < \frac{1}{10\lambda}$. Using these facts, Lemma 3.2.2 gives us that $J_{a_1 a_2 \ldots}$ is a λ'-quasi-arc, with $\lambda' = \lambda'(\beta, L, N) = \lambda'(L, N)$, that β^n-follows $J_{a_1 a_2 \ldots a_n}$ for each n.

(Finding quasi-arcs in the limit is not unexpected since on each scale the limit set will look like the quasi-arc approximation on the same scale.)

If we set $\mathcal{M}(X)$ to be the set of all closed sets in X, we can define a map $F : Z \to \mathcal{M}(X)$ by $F(a) = J_a$. Let $\mathcal{J} = F(Z)$ be the image of this map and choose the metric d_σ for Z, $\sigma = \frac{-\log(2)}{\log(\beta)} > 0$. It remains to show that $F : (Z, d_\sigma) \to (\mathcal{M}(X), d_{\mathcal{H}})$ is a bi-Lipschitz embedding.

Take $a = (a_1, a_2, \ldots), b = (b_1, b_2, \ldots) \in Z$. Then $d_\sigma(a,b) \in (\beta^{n+1}, \beta^n]$ if and only if $a_i = b_i$ for $1 \leq i < n$ and $a_n \neq b_n$. By construction, and a geometric series, $J_a \subset N(J_{a_1 \ldots a_n}, \frac{1}{4}\beta^n)$, and so as n stage arcs are β^n separated we get that

$$d_{\mathcal{H}}(J_a, J_b) \geq d(J_a, J_b) \geq \frac{1}{2}\beta^n \geq \frac{1}{2} d_\sigma(a,b). \tag{4.3}$$

Conversely, applying the triangle inequality and Fact 4.2.1 we get that

$$d_{\mathcal{H}}(J_a, J_b) \leq d_{\mathcal{H}}(J_a, J_{a_1 \ldots a_{n-1}}) + d_{\mathcal{H}}(J_{b_1 \ldots b_{n-1}}, J_b) \leq 2\lambda\beta^{n-1} \leq \frac{2\lambda}{\beta^2} d_\sigma(a,b), \tag{4.4}$$

so F is bi-Lipschitz, quantitatively.

As a final remark, note that there is a natural measure $\mu_\sigma = F_*(\nu_\sigma)$ on $\mathcal{J}$. The estimates (4.3) and (4.4) imply that for any ball $B(x,r) \subset X$ the set $\{J_a \in \mathcal{J} | J_a \cap B(x,r) \neq \emptyset\}$ is measurable (in fact open), and if two arcs J_a and J_b both meet this ball we have that $2r \geq d(J_a, J_b) \geq \frac{1}{2} d_\sigma(a,b)$, and so

$$\mu_\sigma\{J_a \in \mathcal{J} | J_a \cap B(x,r) \neq \emptyset\} \leq 4^\sigma r^\sigma.$$

□

Proof of Theorem 4.0.1. The construction of Theorem 4.0.2 gives a lower bound for conformal dimension by virtue of Lemma 2.3.1.

Two remarks are required. First, in Theorem 4.0.1, X may be non-compact, but we know that it is complete and doubling, therefore it is proper and, as noted in Remark 2.3.2, Lemma 2.3.1 applies, provided all arcs $\gamma \in \mathcal{C}$ lie in some fixed ball in X. Second, the packing dimension of X is finite and bounded above by a constant derived from the doubling constant N.

Therefore, following Theorem 4.0.2, we can apply Lemma 2.3.1 with $\mathcal{C} = \mathcal{J}$, $\mu = \mu_\sigma$ and $A = 4^\sigma$, where σ depends only on L and N, to find a lower bound for the conformal dimension of $C = C(L, N) > 1$. $\square$

CHAPTER V

One dimensional boundaries of hyperbolic groups

The original motivation for proving Theorem 4.0.1 was in the context of boundaries of hyperbolic groups. In this chapter we apply Theorem 4.0.1 to this case in order to answer a question of Bonk and Kleiner.

5.1 Hyperbolic groups and their boundaries

Given a finitely generated group G and a choice of finite symmetric generating set $S \subset G$ (i.e., $s \in S$ implies $s^{-1} \in S$), then we can define a word norm for $g \in G$ by

$$\|g\|_S = \inf\{n : g = s_1 s_2 \cdots s_n, s_i \in S\},$$

and word metric

$$d_S(g,h) = \|h^{-1}g\|.$$

Another way to view this word metric is as the path metric on the Cayley graph $\Gamma(G,S)$. Recall that $\Gamma(G,S)$ is the graph with vertex set G and one edge (of length one) between pairs of vertices g and gs, for all $g \in G$ and $s \in S$. The action of G on itself by left multiplication induces an (isometric) action on $\Gamma(G,S)$.

The standard observation in geometric group theory is that different choices of finite generating set give quasi-isometric Cayley graphs. Therefore, quasi-isometrically invariant properties of a metric space, applied to a Cayley graph, will be well defined for the underlying group.

An exceptionally useful property of a metric space was introduced by Gromov in his study of negatively curved groups and spaces. A geodesic metric space is called *(Gromov) hyperbolic* if every geodesic triangle is δ-thin for some $\delta > 0$: every point on each edge of the triangle is within a distance of δ of one of the other two sides. The name 'hyperbolic' comes from the fact that, as an a easy consequence of the Gauss-Bonnet theorem, $\mathbb{H}^n$ is δ-hyperbolic for $\delta = \log(1+\sqrt{2})$.

The property of being hyperbolic is a quasi-isometric invariant of a metric space, and so we can call a group hyperbolic if one (and hence every) Cayley graph of the group is hyperbolic. Hyperbolic groups are precisely those groups that satisfy a sub-quadratic isoperimetric inequality, and have many good properties, such as having solvable word problem [Gro87].

The fundamental group of a compact hyperbolic manifold M^n with totally geodesic boundary is a hyperbolic group. As is well known, the universal cover of M is a convex subset of $\mathbb{H}^n$, quasi-isometric to $\pi_1(M)$, and we can find the limit set in the sphere at infinity $\partial_\infty \pi_1(M) = \partial_\infty M \subset \mathbb{S}^{n-1}$. We can do a similar construction for any hyperbolic group G and define its boundary at infinity $\partial_\infty G$. The quasi-isometric ambiguity in the metric on G leads to a quasisymmetric ambiguity in the choice of metric on $\partial_\infty G$.

Therefore, a quasisymmetric invariant of a metric space will give a quasi-isometric invariant of a (hyperbolic) group. In our case, we can speak of the conformal dimension of the boundary of a hyperbolic group G, and denote this by $\dim_{\mathcal{C}}(\partial_\infty G)$.

From the basic properties of conformal dimension we have that

$$\dim_{\text{top}}(\partial_\infty G) \leq \dim_{\mathcal{C}}(\partial_\infty G).$$

If the boundary has topological dimension zero it is totally disconnected. A simple argument using Stallings' theorem on ends [Sta68] and Dunwoody's accessibility theorem for finitely presented groups [Dun85] gives that in this case the group is virtually free and the boundary has conformal dimension zero.

When the boundary has topological dimension one the situation is more complicated. There are three basic cases.

Case 1: The boundary is a circle, thus by [Gab92, CJ94] the group is virtually Fuchsian. This happens if and only if the (Ahlfors regular) conformal dimension of the boundary is one, and it is realized for a particular choice of metric [BK02b].

Case 2: The boundary is not a circle, but it does have local cut points. In this case the group virtually splits over a virtually cyclic subgroup [Bow99]. There are examples in this case which have conformal dimension one, and examples with conformal dimension greater than one. For more discussion, see Chapter VI.

Case 3: The boundary has no local cut points, and so by [KK00] is homeomorphic to the Sierpiński carpet or the Menger curve. Bonk and Kleiner asked if in this case the conformal dimension was greater than one [BK05a, Problem 6.2]. This is the content of the following theorem.

Theorem 5.1.1. *Suppose G is a non-elementary hyperbolic group which does not virtually split over any elementary group. Then the conformal dimension of $\partial_\infty G$ is strictly greater than one.*

Work of Bowditch and Swarup [Bow99, Swa96] on the boundaries of hyperbolic groups shows that the algebraic criterion in the statement of the theorem is equivalent to the boundary being connected with no local cut points. Essentially, the self-similarity of the boundary then allows us to promote these conditions to the annulus linearly connected condition, using a short dynamical argument similar to one given by Bonk and Kleiner in [BK05b]. We prove this in the next section.

5.2 A new conformal dimension bound.

Proof of Theorem 5.1.1. By work of Bowditch and Swarup [Bow99, Swa96] a hyperbolic group G does not (virtually) split over an elementary subgroup if and only if the conformal boundary has no local cut points.

Proposition 4 in [BK05b] shows that $\partial_\infty G$ with some visual metric d is compact, doubling and linearly connected. It remains to show that $(X, d) = (\partial_\infty G, d)$ is annulus linearly connected, but this follows by a proof similar to that of Bonk and Kleiner's proposition.

Suppose (X, d) is not annulus linearly connected. Then there is a sequence of annuli $A_n = A(z_n, r_n, 2r_n)$ containing points x_n and y_n such that there is no arc joining x_n to y_n inside $A(z_n, \frac{1}{n}r_n, 2nr_n)$. As X is compact we have that $r_n \to 0$, otherwise there would be a subsequence $n_j \to \infty$ as $j \to \infty$ with $r_{n_j} > \epsilon > 0$ for some ϵ. In this case take further subsequences so that $r_{n_j} \to r_\infty \in [\epsilon, \mathrm{diam}(X)]$, $z_{n_j} \to z_\infty$, $x_{n_j} \to x_\infty$ and $y_{n_j} \to y_\infty$. Then a contradiction follows from the fact that z_∞ is not a local cut point, so we must have $r_n \to 0$.

Now we can consider the rescaled sequence $(X, \frac{1}{r_n}d, z_n)$. By doubling, this subconverges to a limit (W, d_W, z_∞) with respect to pointed Gromov-Hausdorff convergence. By Lemma 5.2 of [BK02b], W is homeomorphic to $\partial_\infty G \setminus \{p\}$ for some p, and so z_∞ cannot be a local cut point in W. So we can connect the components of $A(z_\infty, 0.9, 2.1)$ in $W \setminus z_\infty$ by finitely many compact sets, and these must lie in some $A(z_\infty, 1/M, 2M)$ for some $1 \leq M < \infty$. For sufficiently large n we can lift these connecting sets to $A(z_n, \frac{1}{2M}r_n, 4Mr_n)$, contradicting the assumption.

In conclusion, $\partial_\infty G$ is annulus linearly connected, doubling and complete and so

Theorem 4.0.1 gives that the conformal dimension of $\partial_\infty G$ is strictly greater than one. $\square$

CHAPTER VI

Conclusion

So what can we say about the conformal dimension of hyperbolic groups with connected one dimensional boundary? We have noted that Fuchsian groups are precisely those groups whose Ahlfors regular conformal dimension is one and is realized. Theorem 5.1.1 showed that if the group does not virtually split over a virtually cyclic subgroup then the conformal dimension is greater than one. However, we do not have an answer to the following

Question 6.0.1. *Which hyperbolic groups have a boundary at infinity of (Ahlfors regular) conformal dimension one?*

The answer is not yet clear. For example, if we amalgamate a group with carpet boundary and a Fuchsian group along a cyclic subgroup, then the resulting group will have local cut points in the boundary, and the conformal dimension will be greater than one. This is because the carpet group, which has conformal dimension greater than one (Theorem 5.1.1), is a quasi-convex subgroup of the amalgam [Bow99], and so the boundary of the carpet group is embedded quasi-symmetrically in the boundary of the amalgam.

On the other hand, Pansu remarked that if two copies of a closed hyperbolic surface are glued along a simple closed geodesic then the conformal dimension of the fundamental group is one. This is because we can deform the surface groups to make the simple closed geodesic smaller and smaller; this has the consequence of lowering the Hausdorff dimension of the boundary of the universal cover arbitrarily close to one.

One would like to know under what circumstances the amalgam of two Fuchsian groups along a cyclic subgroup has a boundary at infinity of conformal dimension one; this is a current research goal. For more background and discussion on this, see [BK05a]. A more general question is:

Question 6.0.2. *When is the conformal dimension of the boundary of a hyperbolic group realized?*

Another corollary of Theorem 4.0.1 is worth noting. If we have a carpet contained in $\mathbb{S}^2$, i.e., a compact subset homeomorphic to the Sierpiński carpet, where the boundary circles are uniform quasi-circles with a uniform bound on their relative distance, then the conformal dimension is greater than one. This is because the annulus linearly connected condition is satisfied for such spaces.

Example 6.0.3. As pointed out to me in a discussion with Juha Heinonen, there are 'round' carpets of conformal dimension arbitrarily close to one. A round carpet is a carpet in $\mathbb{S}^2$ with round circles for its boundary circles.

To construct a sequence of examples X_m, modify the standard square Sierpiński carpet construction by removing a square of side $\frac{3^m-2}{3^m}$ from the middle of each square at each step. This gives a space of Hausdorff dimension

$$\dim_{\mathcal{H}}(X_m) = \frac{\log\left(4 \cdot (3^m - 1)\right)}{\log(3^m)} \to 1 \text{ as } m \to \infty.$$

Each X_m is quasisymmetric to a round carpet Y_m by work of Bonk [Bon06], and therefore Y_m, $m = 1, 2, \ldots$, is a sequence of round carpets whose conformal dimension converges to one. (By Theorem 4.0.1, we have $\dim_{\mathcal{C}}(Y_m) = \dim_{\mathcal{C}}(X_m) > 1$.)

However, for boundaries of hyperbolic groups we do not have such examples.

Question 6.0.4. *Are there hyperbolic groups with Sierpiński carpet boundary and conformal dimension arbitrarily close to one?*

For hyperbolic groups with Menger curve boundary, this question has been answered by examples of Bourdon and Pajot [BP99]; in fact, for these groups the conformal dimension takes values in a dense subset of $(1, \infty)$.

Finally, it would be nice to know the answer to the following question.

Question 6.0.5. *What is the conformal dimension of the square Sierpiński carpet?*

BIBLIOGRAPHY

BIBLIOGRAPHY

[BBI01] D. Burago, Y. Burago, and S. Ivanov, *A course in metric geometry*, Graduate Studies in Mathematics, vol. 33, American Mathematical Society, Providence, RI, 2001. MR 1835418 (2002e:53053)

[BK02a] M. Bonk and B. Kleiner, *Quasisymmetric parametrizations of two-dimensional metric spheres*, Invent. Math. **150** (2002), no. 1, 127–183.

[BK02b] ———, *Rigidity for quasi-Möbius group actions*, J. Differential Geom. **61** (2002), no. 1, 81–106.

[BK05a] ———, *Conformal dimension and Gromov hyperbolic groups with 2-sphere boundary*, Geom. Topol. **9** (2005), 219–246 (electronic).

[BK05b] ———, *Quasi-hyperbolic planes in hyperbolic groups*, Proc. Amer. Math. Soc. **133** (2005), no. 9, 2491–2494 (electronic).

[Bon06] M. Bonk, *Quasiconformal geometry of fractals*, International Congress of Mathematicians. Vol. II, Eur. Math. Soc., Zürich, 2006, pp. 1349–1373.

[Bou95] M. Bourdon, *Au bord de certains polyèdres hyperboliques*, Ann. Inst. Fourier (Grenoble) **45** (1995), no. 1, 119–141.

[Bow98] B. H. Bowditch, *Cut points and canonical splittings of hyperbolic groups*, Acta Math. **180** (1998), no. 2, 145–186.

[Bow99] ———, *Connectedness properties of limit sets*, Trans. Amer. Math. Soc. **351** (1999), no. 9, 3673–3686.

[BP99] M. Bourdon and H. Pajot, *Poincaré inequalities and quasiconformal structure on the boundary of some hyperbolic buildings*, Proc. Amer. Math. Soc. **127** (1999), no. 8, 2315–2324.

[BP03] ———, *Cohomologie l_p et espaces de Besov*, J. Reine Angew. Math. **558** (2003), 85–108.

[CJ94] A. Casson and D. Jungreis, *Convergence groups and Seifert fibered 3-manifolds*, Invent. Math. **118** (1994), no. 3, 441–456.

[Cul68] H. F. Cullen, *Introduction to general topology*, D. C. Heath and Co., Boston, Mass., 1968. MR 0221455 (36 #4507)

[Cut95] C. D. Cutler, *The density theorem and Hausdorff inequality for packing measure in general metric spaces*, Illinois J. Math. **39** (1995), no. 4, 676–694.

[Dun85] M. J. Dunwoody, *The accessibility of finitely presented groups*, Invent. Math. **81** (1985), no. 3, 449–457.

[Gab92] D. Gabai, *Convergence groups are Fuchsian groups*, Ann. of Math. (2) **136** (1992), no. 3, 447–510.

[Gro87] M. Gromov, *Hyperbolic groups*, Essays in group theory, Math. Sci. Res. Inst. Publ., vol. 8, Springer, New York, 1987, pp. 75–263.

[Hei01] J. Heinonen, *Lectures on analysis on metric spaces*, Universitext, Springer-Verlag, New York, 2001.

[Hei07] ______, *Nonsmooth calculus*, Bull. Amer. Math. Soc. (N.S.) **44** (2007), no. 2, 163–232 (electronic).

[HK98] J. Heinonen and P. Koskela, *Quasiconformal maps in metric spaces with controlled geometry*, Acta Math. **181** (1998), no. 1, 1–61.

[HY61] J. G. Hocking and G. S. Young, *Topology*, Addison-Wesley, Reading, Mass., 1961.

[KK00] M. Kapovich and B. Kleiner, *Hyperbolic groups with low-dimensional boundary*, Ann. Sci. École Norm. Sup. (4) **33** (2000), no. 5, 647–669.

[Kle06] B. Kleiner, *The asymptotic geometry of negatively curved spaces: uniformization, geometrization and rigidity*, International Congress of Mathematicians. Vol. II, Eur. Math. Soc., Zürich, 2006, pp. 743–768.

[Mac07a] J. M. Mackay, *Existence of quasi-arcs*, Accepted by Proc. Amer. Math. Soc. (2007), `arXiv:0708.0824`.

[Mac07b] ______, *Spaces with conformal dimension greater than one*, Preprint (2007), `arXiv:0711.0417`.

[Mos73] G. D. Mostow, *Strong rigidity of locally symmetric spaces*, Princeton University Press, Princeton, N.J., 1973, Annals of Mathematics Studies, No. 78.

[Pan89] P. Pansu, *Métriques de Carnot-Carathéodory et quasiisométries des espaces symétriques de rang un*, Ann. of Math. (2) **129** (1989), no. 1, 1–60.

[Pau96] F. Paulin, *Un groupe hyperbolique est déterminé par son bord*, J. London Math. Soc. (2) **54** (1996), no. 1, 50–74.

[Sta68] J. R. Stallings, *On torsion-free groups with infinitely many ends*, Ann. of Math. (2) **88** (1968), 312–334.

[Sul81] D. Sullivan, *On the ergodic theory at infinity of an arbitrary discrete group of hyperbolic motions*, Riemann surfaces and related topics: Proceedings of the 1978 Stony Brook Conference (State Univ. New York, Stony Brook, N.Y., 1978) (Princeton, N.J.), Ann. of Math. Stud., vol. 97, Princeton Univ. Press, 1981, pp. 465–496.

[Swa96] G. A. Swarup, *On the cut point conjecture*, Electron. Res. Announc. Amer. Math. Soc. **2** (1996), no. 2, 98–100 (electronic).

[Tuk96] P. Tukia, *Spaces and arcs of bounded turning*, Michigan Math. J. **43** (1996), no. 3, 559–584.

[TV80] P. Tukia and J. Väisälä, *Quasisymmetric embeddings of metric spaces*, Ann. Acad. Sci. Fenn. Ser. A I Math. **5** (1980), no. 1, 97–114.