A rigidity theorem for the solvable Baumslag-Solitar groups

Benson Farb and Lee Mosher
with an appendix by
Daryl Cooper

September 20, 1996

1 Introduction

In [Gr1], [Gr2] Gromov proposes and studies the problem of classifying finitely generated groups up to quasi-isometry. On the one hand this has motivated an industry of producing quasi-isometry invariants. On the other hand there are rigidity results which give a complete quasi-isometry classification within certain classes of groups. So far most of the progress in rigidity has been made for groups which arise in geometry. For example, in a series of papers by several authors (see, e.g. [Tu], [CC], [Me], [Pa], [Sch2], [FS], [Ch], [Sch1], [KlL], [EF], [Es], or [Fa] for a summary), a battery of new techniques related to the proof of the Mostow Rigidity Theorem has led to a nearly complete quasi-isometry classification of lattices in semisimple Lie groups. Quasi-isometric rigidity results for other groups arising in geometry and topology include [Ri], [BG], [KaL].

In this paper we take the first steps towards applying some of these ideas to proving rigidity results for groups that arise most naturally not in geometry but in combinatorial group theory.

The solvable Baumslag-Solitar groups. The solvable Baumslag-Solitar groups $\mathrm{BS}(1,n)$ are given by the presentation

$$\mathrm{BS}(1,n) = \langle a, b \mid aba^{-1} = b^n \rangle.$$

*The first author is supported in part by an NSF postdoctoral fellowship, the second author by NSF grant DMS 95–04946, and the third author by NSF grant DMS 95–10505. The work of the authors at MSRI was supported by NSF grant DMS–9022140.

These groups have served as a proving ground for many new ideas in combinatorial and geometric group theory. For example, they are the simplest groups to have an exponential isoperimetric function, and they were the first known groups to be asynchronously automatic but not automatic (see [ECH+]).

The rigidity machinery for lattices in Lie groups does not directly apply to the group $\mathrm{BS}(1,n)$ when $n \geq 2$, for although $\mathrm{BS}(1,n)$ has a faithful, *indiscrete* representation into $PSL_2(\boldsymbol{R})$ given by

$$a \mapsto \begin{pmatrix} n^{1/2} & 0 \\ 0 & n^{-1/2} \end{pmatrix} \qquad b \mapsto \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}$$

it does not seem to be a lattice in any connected Lie group (see §9). Furthermore, $\mathrm{BS}(1,n), n \geq 2$ is not a 3-manifold group ([JS], [Kr]), and is not a nonpositively curved group in any sense of the word (see §9).

Our main result is the following (for definitions see below):

Theorem 7.1 (Quasi-isometric iff commensurable). *Given integers $m, n \geq 2$, the groups $BS(1,m)$ and $BS(1,n)$ are quasi-isometric if and only if they are commensurable. This happens if and only if there exist integers $r, j, k > 0$ such that $m = r^j$ and $n = r^k$.*

We heard about this question from S. Weinberger [We] and S. Gersten [Ge]. As far as we know, Theorem 7.1 is the first quasi-isometric rigidity result for non-nilpotent solvable groups. Note that $\mathrm{BS}(1,1) \approx \boldsymbol{Z} \oplus \boldsymbol{Z}$ is not quasi-isometric to $\mathrm{BS}(1,n)$, $n \geq 2$, because they have different isoperimetric functions (see [ECH+]).

Geometry, boundaries, methods. As does any finitely presented group, $\mathrm{BS}(1,n)$ acts properly discontinuously and cocompactly by isometries on a certain metric 2-complex, in this case denoted X_n. In Section 3 we construct X_n with an explicit metric (see Figure 1). The space X_n admits a beautiful geometry which has aspects of both positive and negative curvature. By a basic observation of Svarc and Milnor (see §2), X_n is quasi-isometric to $\mathrm{BS}(1,n)$, so we can use X_n to study quasi-isometries of $\mathrm{BS}(1,n)$.

A key object of study in our proof of Theorem 7.1 is the *upper boundary* $\partial^u X_n$ of X_n (see §6), which is the *space of hyperbolic planes* in X_n. This is to be compared/contrasted to the space of flats, or *Furstenberg boundary*, studied in Mostow's proof [Mo], for example the boundary at infinity of a hyperbolic space. It turns out that the upper boundary $\partial^u X_n$ is a certain

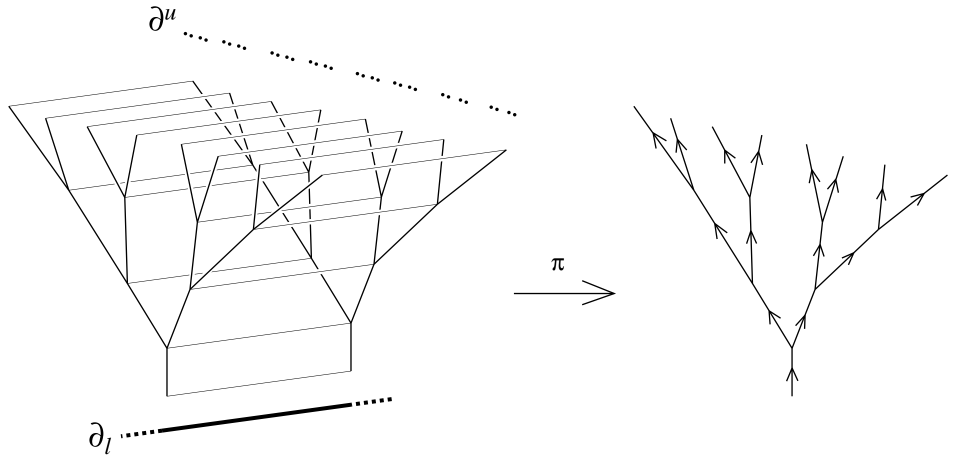

Figure 1: A piece of the 2-complex X_2 associated to BS$(1, 2)$, together with the fibration $\pi\colon X_2 \to T_2$ (see §3). For each edge $E \subset T_2$, the inverse image $\pi^{-1}(E)$ is isometric to the "horostrip" $1 \leq y \leq 2$ in the upper half plane model of $\boldsymbol{H}^2$. The geodesic curvature vector on a horocycle induces an orientation on each edge of T_2. For each coherently oriented proper line $L \subset T_2$, the set $H_L = \pi^{-1}(L)$ is an isometrically embedded hyperbolic plane in X_2. The point inverse images form a family of horocycles based at a point $\infty_L \in \partial H_L$. The figure shows the "upper boundary" $\partial^u X_2$ (see §5), the set of hyperbolic planes in X_2. Also shown is the "lower boundary" $\partial_\ell X_n$ (see §6), a line obtained by identifying all of the lines $\partial H_L - \{\infty_L\}$.

self-similar Cantor set, which can be thought of as the n-adic rationals $\mathbf{Q}_n$ with the usual metric.

In the core of this paper, Sections 4, 5, and 6, we use topological and geometric arguments to show that a quasi-isometry $f : X_n \to X_m$ induces a *bilipschitz homeomorphism* $f^u : \partial^u X_n \to \partial^u X_m$. Homeomorphisms of self-similar Cantor sets have been studied by Cooper and Pignataro [CP]. In an appendix to our paper, Cooper provides an extension of their work which we use to finish the proof of Theorem 7.1. The idea is that bilipschitz homeomorphisms of self-similar Cantor sets are locally measure linear (for Hausdorff measure in the appropriate dimension). It is the local similarity groups of the self-similar Cantor set $\partial^u X_n$ which are quasi-isometry invariants of X_n, and which allow detection of the prime factors (with multiplicities) of n.

The quasi-isometry group. For any metric space X, one can form the group $\mathrm{QI}(X)$ of all self quasi-isometries of X modulo those that lie a bounded distance (in the sup norm) from the identity. Modding-out by this equivalence relation makes $\mathrm{QI}(X)$ into a group. In proving Theorem 7.1, we actually compute the quasi-isometry group of $\mathrm{BS}(1,n)$, namely:

$$\mathrm{QI}(\mathrm{BS}(1,n)) \approx \mathrm{Bilip}(\boldsymbol{R}) \times \mathrm{Bilip}(\mathbf{Q}_n)$$

where $\mathrm{Bilip}(Y)$ denotes the group of bilipschitz homeomorphisms of a metric space Y.

Acknowledgements

We'd like to thank Alex Eskin, Curt McMullen, and Bill Thurston for useful conversations, and Jen Taback for corrections. We especially thank Curt McMullen and Bill Thurston for critical help on Proposition 5.3. Thanks also to Steve Gersten and to Shmuel Weinberger, both of whom asked about quasi-isometries of Baumslag-Solitar groups. Finally, we'd like to express our thanks and appreciation to Daryl Cooper for telling us about homeomorphisms of self-similar Cantor sets, and for writing the appendix to this paper.

2 Preliminaries on quasi-isometries

A (coarse) *quasi-isometry* between metric spaces is a map $f : X \to Y$ such that, for some constants $K, C, C' > 0$:

1. $K^{-1}d_X(x_1, x_2) - C \leq d_Y(f(x_1), f(x_2)) \leq Kd_X(x_1, x_2) + C$ for all $x_1, x_2 \in X$.
2. The C'-neighborhood of $f(X)$ is all of Y.

The map f is called a (K, C)-quasi-isometry, or a K-quasi-isometry for short. There is always a *coarse inverse* of f, a quasi-isometry $g: Y \to X$ such that $f \circ g$ and $g \circ f$ are a bounded distance from the identity maps in the sup norm; these bounds, and the quasi-isometry constants for g, depend only on the quasi-isometry constants of f.

A map satisfying (1) but not necessarily (2) is called a *quasi-isometric embedding* of X into Y.

It is a fundamental observation of Svarc and Milnor (see [Mi], and also [Ca]) that the fundamental group $\pi_1(M)$ (endowed with the word metric) of a compact Riemannian manifold M is quasi-isometric to the universal cover $\widetilde{M}$ of M. More generally, let X be a proper metric space and let Γ be a group of isometries acting properly discontinuously on X. If $\Gamma \backslash X$ is compact, then Γ is finitely generated and is quasi-isometric to X.

Two groups are *(abstractly) commensurable* if they have finite index subgroups which are isomorphic. Clearly commensurable groups are quasi-isometric.

Furthermore, by the now-standard "connect-the-dots" technique (see, e.g. [FS], 5.4), we can change a quasi-isometry (and it's coarse inverse) by a bounded amount in the sup norm so that it is a lipschitz map, while still being a quasi-isometry. This may require enlarging the constants, but their new values depend only on the old values.

By Rademacher's Theorem (see e.g. [EG], 3.1.2), lipschitz maps are differentiable almost everywhere, enabling one to define the length of a lipschitz path $\gamma : [0, 1] \to X$ when X is a piecewise Riemannian simplicial complex. More precisely, if the derivative $D_x\gamma$ exists at a point $x \in [0, 1]$, it sends a unit vector at x to a vector in $T_{\gamma(x)}(X)$. We can compute the length of this vector using the metric on X. This defines a function $L(x)$ almost everywhere on $[0, 1]$, and we can then define the *length* of γ to be $\ell(\gamma) = \int_{[0,1]} L(x)dx$. This integral exists because $L(x)$ is a bounded measurable function defined almost everywhere on $[0, 1]$, as $||D_x\gamma||$ is bounded by the lipschitz constant of γ.

Hence it makes sense to talk about how a Lipschitz quasi-isometry between piecewise Riemannian complexes distorts path length ℓ_X. We also could have defined and used an obvious coarse notion of path length, which agrees with the above notion up to constant multiplicative and additive factors.

Lemma 2.1 (bounded stretch). *Let $q : X \to Y$ be a lipschitz K-quasi-isometric embedding between piecewise Riemannian complexes. Then there exists K', depending only on K and not on q, so that if γ is lipschitz path in X, then*

$$\frac{1}{K'}\,\ell_X(\gamma) - K' \leq \ell_Y(q \circ \gamma) \leq K'\,\ell_X(\gamma) + K'$$

Proof. The second inequality follows easily from the fact that q is lipschitz. The first inequality follows from applying the same reasoning to a lipschitz map which is a bounded distance from a coarse inverse of q. $\diamond$

3 The 2-complex X_n

For the rest of the paper let $\boldsymbol{H}^2 = \{(x,y) \mid x \in \boldsymbol{R},\, y > 0\}$ denote the hyperbolic plane in the upper half plane model, with metric $(dx^2 + dy^2)/y^2$, and with $\partial \boldsymbol{H}^2 = \boldsymbol{R} \cup \{\infty\}$.

In this section we construct a metric 2-complex X_n on which $\mathrm{BS}(1,n)$ acts properly discontinuously and cocompactly by isometries. This 2-complex is a well-known object in combinatorial group theory (see e.g. [ECH$^+$]). We shall put an equivariant metric on this complex, which is designed so that X_n contains many isometric copies of $\boldsymbol{H}^2$.

To describe X_n topologically, consider the "horobrick" $H_n \subset \boldsymbol{H}^2$ defined by $0 \leq x \leq n$, $1 \leq y \leq n$. The left and right sides of H_n are both geodesic segments of hyperbolic length $\log(n)$; orient these segments upward and label each with a. The top of H_n is a horocyclic segment of hyperbolic length 1; orient this segment to the right and label it with b. The bottom of H_n is a horocyclic segment of hyperbolic length n; divide the bottom into n equal subsegments of length 1, orient each subsegment to the right, and label with b. Now form a piecewise hyperbolic complex C_n by identifying the a segments isometrically, and identifying the b segments isometrically. The subdivision of ∂H_n induces a cell structure on C_n, from which one easily reads off the presentation $\pi_1(C_n) = \mathrm{BS}(1,n) = \langle a, b \mid aba^{-1} = b^n \rangle$. Let X_n be the universal cover of C_n, and lift the metric on C_n to give a piecewise hyperbolic metric on X_n. The group $\mathrm{BS}(1,n)$ acts on X_n as the deck transformation group of the universal covering map $X_n \to C_n$.

To describe further features of X_n, let s be the "horostrip" in $\boldsymbol{H}^2$ given by $1 \leq y \leq n$. The parabolic isometry $\tau\colon (x,y) \mapsto (x+n,y)$ acts on s, and the quotient s/τ is a metric cylinder having one boundary circle of length 1 and the other boundary circle of length n. We alternatively construct C_n by gluing the long boundary circle of s/τ to the short boundary circle, via a degree n locally isometric convering map, induced by the path isometry taking the horocycle $y=1$ to the horocycle $y=n$ via the formula $(x,1) \mapsto (nx,n)$.

Define a *horocycle in* C_n to be the image, under the map $s \to s/\tau \to C_n$, of a horocycle $\{y=t\}$ in s. Define a *horocycle in* X_n to be a connected lift of a horocycle in C_n. Collapsing each horocycle of X_n to a point defines a fibration $\pi\colon X_n \to T_n$, where T_n is an infinite, regular, $(n+1)$-valent tree (see Figure 1). The fibration of X_n by horocycles is preserved by the action of $\mathrm{BS}(1,n)$ on X_n, and hence there is an induced action of $\mathrm{BS}(1,n)$ on the tree T_n. The tree T_n is the usual Bass-Serre tree (see [Se]) associated to the HNN extension $\mathrm{BS}(1,n) = \boldsymbol{Z}*_\phi$ where $\phi(k) = nk$.

The horocycles on the horostrip s have a transverse orientation pointing towards ∞. This induces a transverse orientation on each horocycle of X_n, which in turn induces an orientation on each edge of T_n. Each vertex of T_n has one incoming and n outgoing edges (see Figure 1).

A *proper line* in T_n is the image of a proper embedding $\boldsymbol{R} \to T_n$. There are two types of proper lines L:

- L is coherently oriented, i.e. the orientation on L induced by T_n agrees with a global orientation on L.
- L is not coherently oriented. Since each vertex has exactly one incoming edge, there is a unique point $p \in L$ with the property that each component of $L - p$ is coherently oriented, and these two orientations point away from p.

The inverse image $P = \pi^{-1}(L)$ of a proper line L is a *proper plane* in X_n. There are two types of proper planes, depending on the orientation type of L:

- If L is coherently oriented then P is an isometrically embedded *hyperbolic plane* in X_n.
- If L is not coherently oriented then P is an isometrically embedded *doubled horoball* in X_n, that is, the union $S_1 \cup S_2$ of horoballs with common horocircle boundary $\partial S_1 = \partial S_2$ called the *seam*. The horocircles are identified via an isometry in the path metric.

Given a hyperbolic plane $P \subset X_n$, when we speak of a horocycle in P we shall *always* mean one of the form $\pi^{-1}(x)$ for some point $x \in T_n$. These form a concentric family of horocycles, based at a point in ∂P which we denote ∞. The *line at infinity* of P will mean $\partial P - \infty$.

Note that a doubled horoball has concentrated *positive curvature* along its seam. In particular X_n is *not* a nonpositively curved space. Indeed it is easy to see that $\mathrm{BS}(1, n)$ has no properly discontinuous cocompact isometric action on any CAT(0) space, because $\mathrm{BS}(1, n)$ has exponential isoperimetric function.

4 Quasi-hyperbolic planes in X_n

Our first "rigidity" result is a study of subsets of X_n which are quasi-isometric to hyperbolic planes. Let d_H denote Hausdorff distance between two subsets of a metric space, i.e. $d_H(C, C')$ is the infimum of numbers $\delta \in [0, \infty]$ such that $C \subset N_\delta(C')$ and $C' \subset N_\delta(C)$.

Proposition 4.1 (Quasi-hyperbolic planes). *For all $K \geq 1, C \geq 0$ there exists $A \geq 0$, such that if $\phi : \boldsymbol{H}^2 \to X_n$ is a (K, C)-quasi-isometric embedding, then there exists a hyperbolic plane $Q \subset X_n$ such that $d_H(\phi(\boldsymbol{H}^2), Q) \leq A$.*

Proposition 4.1 may be thought of as an analogy to X_n of the fact that quasi-geodesics in a hyperbolic space lie close to geodesics, although Proposition 4.1 seems much harder to prove.

Applying the "connect-the-dots" argument discussed above, for the rest of the proof we may replace ϕ by a lipschitz quasi-isometric embedding which is close to ϕ in the sup norm.

Coarse topology

In proving Proposition 4.1 we will use two coarse topological theorems. The first is due to R. Schwartz, and both were first applied to quasi-isometric rigidity problems in [FS].

Recall that a metric space (M, d) is *uniformly contractible* if there is a function $\alpha : \boldsymbol{R}^+ \to \boldsymbol{R}^+$ having the following property: If a continuous map of a finite simplicial complex $\Delta \to M$ is contained in an r-ball, then it is contractible in an $\alpha(r)$-ball. The function α is not supposed to depend on the dimension of Δ. Clearly any contractible space admitting a cocompact group of isometries is uniformly contractible.

Proposition 4.2 (Coarse Separation Theorem). *Suppose $\phi\colon \boldsymbol{H}^2 \to Y$ is a lipschitz K-quasi-isometric embedding of $\boldsymbol{H}^2$ into a uniformly contractible Riemannian manifold Y diffeomorphic to $\boldsymbol{R}^3$. Then $Y - \phi(\boldsymbol{H}^2)$ has at least two components.*

Remark. This is an immediate consequence of the general Coarse Separation Theorem given as Theorem 5.2 of [FS], or more precisely from the lipschitz version of that theorem given in Section 5 of [FS]. To apply Theorem 5.2 of [FS], one uses the fact that $\boldsymbol{H}^2$ is uniformly contractible, and has an "expanding sphere", namely the r-spheres for $r > 0$.

The following is a metrical version of invariance of domain:

Proposition 4.3 (Packing Theorem). *Suppose that J is diffeomorphic to $\boldsymbol{R}^n$ and is uniformly contractible, and $\phi\colon J \to J$ is a K-quasi-isometric embedding. Then $nbhd_{K'}(\phi(J)) = J$ for some K' depending only on the pair (K, J).*

A short proof of the packing theorem is given as Lemma 8.2 in [EF]. A different proof (assuming more hypotheses) is given in Section 5 of [FS].

Proof of Proposition 4.1. Let $X = X_n$, and let $\pi : X \to T$ be the natural projection to the tree $T = T_n$. The proof of Proposition 4.1 is divided into several steps.

Step 1. *There exists a proper line $L \subset T$ such that the plane $Q = \pi^{-1}(L)$ is contained in $\phi(\boldsymbol{H}^2)$.*

We will adjust the given setup to a situation in which Coarse Separation (Proposition 4.2) can be applied.

To start with, let $\alpha\colon T \to \boldsymbol{R}^2$ be any proper embedding. For example, if S is a cusped hyperbolic surface of sufficiently large topological type, then S has a spine Σ which is a rose of valence at least as large as the valence of all vertices in T, and then there is a proper embedding of T into $\widetilde{\Sigma} \subset \widetilde{S} \approx \boldsymbol{H}^2 \approx \boldsymbol{R}^2$.

Choose a homeomorphism $\beta\colon X \to T \times \boldsymbol{R}$. The map $(\alpha \times \mathrm{Id}) \circ \beta$ gives a proper topological embedding

$$X \xrightarrow{\beta} T \times \boldsymbol{R} \xrightarrow{\alpha \times Id} \boldsymbol{R}^2 \times \boldsymbol{R} = \boldsymbol{R}^3$$

Regarding X now as a subset of $\boldsymbol{R}^3$, for each component C of $\boldsymbol{R}^3 - X$ the frontier ∂C is a topological plane in X. Metrically the plane ∂C is either a hyperbolic plane or a doubled horoball in X, and in either case ∂C is totally geodesic in X, in that the induced metric from X and path metric on ∂C are the same. In fact, the embedding $\alpha\colon T \to \boldsymbol{R}^2$ can be chosen so that each ∂C is a doubled horoball.

Choose a homeomorphism from $\overline{C} = C \cup \partial C$ to $\partial C \times [0, \infty)$. Using this product decomposition of $\overline{C}$, we may define the metric on $\overline{C}$ to be the product of the metric on ∂C (as a totally geodesic subset of X) and the usual metric on $[0, \infty)$. This defines a path metric on $\overline{C}$ for each component C of $\boldsymbol{R}^3 - X$. These path metrics agree where they meet along X, so pasting them together gives a geodesic metric d on all of $\boldsymbol{R}^3$. In this metric, the embedding $X \hookrightarrow (\boldsymbol{R}^3, d)$ is totally geodesic, meaning that for any geodesic in $(\boldsymbol{R}^3, d)$ whose endpoints lie in X, the entire geodesic lies in X.

The metric space $(\boldsymbol{R}^3, d)$ is uniformly contractible: any ball $B \subset (\boldsymbol{R}^3, d)$ with $B \cap X \neq \emptyset$ admits a deformation retraction $B \times [0,1] \to B$ such that $B \times 1 \to B \cap X$; and furthermore X is uniformly contractible because it admits a cocompact isometric group action. If $B \cap X = \emptyset$ then B admits such a deformation into a (uniformly contractible) copy $\partial C \times \{x\}$, $x \in (0, \infty)$ of ∂C.

Consider the lipschitz quasi-isometric embedding $\phi\colon \boldsymbol{H}^2 \to X$. Composing this with the isometric embedding $X \to (\boldsymbol{R}^3, d)$, we may regard ϕ as a lipschitz quasi-isometric embedding $\boldsymbol{H}^2 \xrightarrow{\phi} (\boldsymbol{R}^3, d)$. Applying Coarse Separation (Proposition 4.2), we know $\boldsymbol{R}^3 - \phi(\boldsymbol{H}^2)$ has at least two components. Since $\phi(\boldsymbol{H}^2) \subset X$, it follows that there are at least two components C, C' of $\boldsymbol{R}^3 - X$ which are separated by $\phi(\boldsymbol{H}^2)$. There exists a finite sequence of components $C = C_0, \ldots, C_J = C'$ of $\boldsymbol{R}^3 - X$ such that C_{j-1}, C_j are *adjacent*, meaning that $\overline{C}_{j-1} \cap \overline{C}_j = \pi^{-1}(E_j)$ for some closed edge E_j of T. Without loss of generality, we may therefore choose C, C' to be adjacent, with $\overline{C} \cap \overline{C}' = \pi^{-1}(E_0)$.

Let U be the union of all edges $E \subset T$ such that $\pi^{-1}(E) \subset \phi(\boldsymbol{H}^2)$. Note that $E_0 \subset U$, for otherwise there would be a path from C to C' piercing X in a single point lying in $\pi^{-1}(E_0) - \phi(\boldsymbol{H}^2)$, contradicting that C, C' are separated by $\phi(\boldsymbol{H}^2)$.

Let E_n be the union of all edge paths in T that start with E_0, are contained in U, and consist of at most $n+1$ edges. We assume by induction that E_n contains an edge path consisting of $2n+1$ edges centered on E_0. Let γ be the boundary of a regular neighborhood of E_n in $\boldsymbol{R}^2$, and write $\gamma = \gamma_1 * \gamma_2$, where γ_1, γ_2 meet at a transversal to the midpoint of E_0 (see Figure 2).

The next step of the induction says that E_{n+1} contains an edge path consisting of $2n+3$ edges centered on E_0. If this were not true, then one

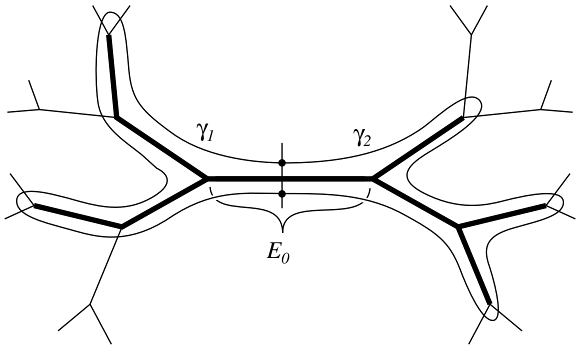

Figure 2: An example of E_n with $n = 2$. If E_{n+1} does not contain an edge path of length $2n+3$ centered on E_0, then one of γ_1 or γ_2 misses U, because they both miss any edge E' which intersects E_n at a vertex lying at distance less than n from E_0, and at least one of them misses any edge which intersects E_n at a vertex lying at distance exactly n from E_0.

of the two paths γ_1, γ_2 would be disjoint from U (see Figure 2). Suppose γ_1 is disjoint from U. For each edge E' that γ_1 pierces, there exists a point $v \in E'$ such that $\pi^{-1}(v) \not\subset \phi(\boldsymbol{H}^2)$, and we may assume that γ_1 pierces E' at v. We may then lift γ_1 to a path in $\boldsymbol{R}_3$ that misses $\phi(\boldsymbol{H}^2)$ and connects C to C', a contradiction.

By induction there is a bi-infinite path $L \subset \bigcup_{i=0}^{\infty} E_i \subset U$, and so $\pi^{-1}(L) \subset \phi(\boldsymbol{H}^2)$, finishing the proof of step 1.

Step 2. *$\phi(\boldsymbol{H}^2)$ is Hausdorff close to the plane $Q = \pi^{-1}(L)$.*

The (K, C) quasi-isometric embedding ϕ has a coarse inverse $\psi\colon \phi(\boldsymbol{H}^2) \to \boldsymbol{H}^2$. The map ψ is a quasi-isometry with constants depending only on K, C. Applying the Packing Theorem (Proposition 4.3) to the restriction of ψ to Q gives a constant $K' = K'(K, C)$ so that $nbhd_{K'}(\psi(Q)) \supset \boldsymbol{H}^2$. It follows that $\phi(\boldsymbol{H}^2) \subset nbhd_{KK'+C'}(Q)$. Step 1 showed that $Q \subset \phi(\boldsymbol{H}^2)$, and hence the Hausdorff distance between Q and $\phi(\boldsymbol{H}^2)$ is at most $KK' + C$.

Step 3. *Q is a hyperbolic plane.*

Since $Q = \pi^{-1}(L)$ for some line L in T, it follows that the plane Q is metrically either a hyperbolic plane or a doubled horoball. By step 2, Q is quasi-isometric to $\boldsymbol{H}^2$. But it is well-known that a doubled horoball is not quasi-isometric to $\boldsymbol{H}^2$ (see e.g. [ECH+] Figure 7.8). For example, the hyperbolic plane satisfies the "fellow traveller property", a quasi-isometry invariant that says: for each $K \geq 0$ there exists $D \geq 0$ such that two K-quasigeodesics with the same endpoints are D-Hausdorff close. But a doubled horoball $B_1 \cup B_2$ does not satisfy the fellow traveller property: by choosing two points $x, y \in B_1 \cap B_2$ sufficiently far apart, the two geodesics connecting x, y in B_1 and in B_2 can be made to have arbitrarily large Hausdorff distance in $B_1 \cup B_2$.

This completes the proof of Proposition 4.1. $\diamond$

5 Horocycles in X_n

The results of Proposition 4.1 motivate the following definitions. We define the *upper boundary* of X_n, denoted $\partial^u X_n$, to be the set of hyperbolic planes in X_n. Note that two distinct hyperbolic planes in X_n are an infinite Hausdorff distance from each other. Combining this fact with Proposition 4.1, it follows that any quasi-isometry $f\colon X_m \to X_n$ induces a map $f^u\colon \partial^u X_m \to \partial^u X_n$, characterized as follows: for each $Q \in \partial^u X_m$ we have $d_H(f(Q), f^u(Q)) < \infty$, where the bound depends only on the quasi-isometry constants of f.

In this section we use the map f^u to investigate how horocycles behave under a quasi-isometry $f\colon X_m \to X_n$. This information is used in turn to further pin down the structure of f.

Quasi-isometries preserve horocycles.

Given a quasi-isometry $f\colon X_m \to X_n$, define an induced map θ_f from the set of horocycles of X_m to the set of horocycles of X_n, as follows. Consider first a *branching horocycle* in X_m, one of the form $\sigma = \pi^{-1}(\text{vertex})$ for some vertex in T_n. There exist hyperbolic planes $Q_1, Q_2 \subset X_m$ such that $\sigma = \partial(Q_1 \cap Q_2)$. Let $Q_i' = f^u(Q_i)$ for $i = 1, 2$. Let $\sigma' = \partial(Q_1' \cap Q_2')$, a branching horocycle in X_n. Then we define $\theta_f(\sigma) = \sigma'$. For any nonbranching horocycle $\sigma \subset X_n$, define $\theta_f(\sigma)$ to be the same as $\theta_f(\sigma_1)$ where σ_1 is the closest branching horocycle above σ. Note that if $\overline{f}\colon X_m \to X_n$ is a coarse inverse for f then $(\theta_f)^{-1} = \theta_{\overline{f}}$ when restricted to the set of branching horocycles.

Although θ_f is only a set map, the next lemma explains the geometric significance of θ_f.

Lemma 5.1 (Horocycles are coarsely preserved). *Given $K \geq 1, C \geq 0$, there exists a constant $\lambda \geq 0$ so that if $f\colon X_m \to X_n$ is a (K, C) quasi-isometry, then for each horocycle $h \subset X_m$ we have*

$$d_H(f(h), \theta_f(h)) \leq \lambda$$

Proof. Clearly we may assume that h is a branching horocycle. The idea of the proof is that h is the boundary of the closeness set of two hyperbolic planes, but closeness sets and their boundaries are coarsely preserved by quasi-isometries.

Let $Q_1, Q_2 \subset X_m$ be the hyperbolic planes such that $h = \partial(Q_1 \cap Q_2)$. By Proposition 4.1, we have hyperbolic planes $Q_i' = f^u(Q_i)$, $i = 1, 2$ in X_n, satisfying $d_H(f(Q_i), Q_i') \le A$ for some $A = A(K, C)$. We also have a branching horocycle $h' = \theta_f(h) = \partial(Q_1' \cap Q_2')$ in X_n.

Choose $x \in h$. We must show that $x' = f(x)$ is close to h'. The set $X_n - h'$ has $n+1$ components. The component intersecting $Q_1' \cap Q_2'$ is said to be *below* h', and the remaining n components are *above* h'.

Case 1: x' is above h'

It follows that h' separates x' from $Q_1' \cap Q_2'$, and hence $d(x', h') \le d(x', Q_1' \cap Q_2') \le A$.

Case 2: x' is below h'

Choose a point $y \in Q_1 - Q_2$ such that $d(y, x) = d(y, Q_2) = D = 2K(A+C)$. The point $y' = f(y)$ satisfies:

$$d(y', Q_1') \le A$$

and

$$d(y', Q_2') \ge \frac{D}{K} - C$$

If y' is below h', then the point on Q_1' closest to y' is also in Q_2', and so $d(y', Q_2') = d(y', Q_1') \le A = D/2K - C < D/K - C$, a contradiction. It follows that y' is above h', and so h' separates x' and y'. Therefore $d(x', h') \le d(x', y') \le K d(x, y) + C = 2K^2(A+C) + C$.

This shows that $f(h)$ lies in a bounded neighborhood of h'. Applying a similar argument to the coarse inverse of f, we conclude that h' lies in a bounded neighborhood of $f(h)$. The bounds clearly depend only on K, C and on A, which in turn depends only on K, C by Proposition 4.1. $\diamond$

Consider a (K, C) quasi-isometry $f \colon X_m \to X_n$, a hyperbolic plane $Q \subset X_m$, and the hyperbolic plane $Q' = f^u(Q) \subset X_n$. Let $f_Q \colon Q \to Q'$ be the composition of f with the closest point projection to Q'. Since $d_H(f(Q), Q') \le A$ it follows that the closest point projection stretches distances by at most a factor depending on A. Applying the preceding lemma we have:

Corollary 5.2. *Given K, C there exist constants K', C', A such that for each (K,C)-quasi-isometry $f\colon X_m \to X_n$ and each $Q \in \partial^u(X_m), Q' = f^u(Q) \in \partial^u(X_n)$, the map $f_Q\colon Q \to Q'$ is a (K', C')-quasi-isometry of hyperbolic planes, and for each horocycle $\sigma \subset Q$ there is a horocycle $\sigma' \subset Q'$ such that $d_H(f_Q(\sigma), \sigma') \le A$.*

Horocycle preserving quasi-isometries of the upper half-plane

Consider a quasi-isometry $f\colon \boldsymbol{H}^2 \to \boldsymbol{H}^2$ which fixes ∞ in the upper half-space model. The induced boundary map $f_\ell\colon \boldsymbol{R} \to \boldsymbol{R}$ is quasi-symmetric, but typically this boundary map will not be bilipschitz. The following proposition explains under what circumstances f_ℓ is bilipschitz. In the next section we will see that this kind of improvement in the quality of f_ℓ imposes strong restrictions on quasi-isometries $X_m \to X_n$.

Recall that for $\boldsymbol{H}^2$ we use the upper half plane model with metric $(dx^2 + dy^2)/y^2$.

Proposition 5.3 ((bilipschitz)$\times$(translation)). *For each $K \ge 1, C \ge 0, A \ge 0$ there exists $D \ge 0$ satisfying the following. Given a (K,C) quasi-isometry $f\colon \boldsymbol{H}^2 \to \boldsymbol{H}^2$ which preserves the point ∞, suppose that for each horocycle σ based at ∞, there is a horocycle $\sigma' = \theta_f(\sigma)$ based at ∞ such that $d_H(f(\sigma), \sigma') \le A$. Then the induced boundary map $f_\ell\colon \boldsymbol{R} \to \boldsymbol{R}$ is bilipschitz. Moreover there exists a real number T such $d(f(x,y), (f_\ell(x), ye^T)) \le D$ for all $(x,y) \in \boldsymbol{H}^2$.*

In other words, the map f is a bounded distance in the sup norm from a map of the form (bilipschitz)$\times$(translation), where the bilipschitz map acts on the line at infinity and by *translation* we mean a translation of the "hyperbolic height" parameter $\log(y)$.

Remark. The sup norm bound D depends only on K, C, A, but the bilipschitz constant for f_ℓ and the translation constant T can depend on f in addition to K, C, A. For example, take f to be any homothety $(x,y) \to (sx, sy)$, $s > 0$. However, the proof of Proposition 5.3 will exhibit the following relation: if $[a,b] \subset \boldsymbol{R}$ is the "bilipschitz norm" of f_ℓ, i.e. the smallest interval such that $a|x-y| \le |f_\ell(x) - f_\ell(y)| \le b|x-y|$ for all $x, y \in \boldsymbol{R}$, then the interval of ratios $[a/e^T, b/e^T]$ has bounded size depending only on K, C, A. In other words, f is a "quasi-homothety".

Proof of Proposition 5.3. By changing the values of K, C we may assume that K, C are Lipschitz quasi-isometry constants for both f and a coarse inverse of f. Also, by changing the value of A we may assume that there is a bijective "Euclidean height transformation" $s \leftrightarrow s'$ such that for each horocycle $\sigma = \{y = s\}$, setting $\sigma' = \theta_f(\sigma) = \{y = s'\}$, we have $d_H(f(\sigma), \sigma') \leq A$ and similarly for a coarse inverse of f. The new values of K, C, A needed to acheive these effects depend only on the old values.

We shall also need the fact that $\lim_{s \to 0} s' = 0$.

Consider a rectangle of the form $R = [p, q] \times [0, s]$. Define the *eccentricity* of R to be $\text{ecc}(R) = \text{width}(R)/\text{height}(R) = |p - q|/s$. The next lemma describes how eccentricity is distorted by f. The *quasi-image* of R under f is defined to be the rectangle $R' = [p', q'] \times [0, s']$ where $p' = f_\ell(p), q' = f_\ell(q)$.

Lemma 5.4 (Eccentricity is quasi-preserved). *There exist constants $D \geq 0$ depending only on K, C, A and $E \geq 1$ depending on K, A, both independent of f, such that for any rectangle $R = [p, q] \times [0, s]$ with quasi-image $R' = [p', q'] \times [0, s]$, if $ecc(R) \geq D$ then*

$$ecc(R') \leq E \cdot ecc(R)$$

and if $ecc(R') \geq D$ then

$$\frac{1}{E} ecc(R) \leq ecc(R')$$

Proof. Since f preserves ∞, each vertical geodesic $x = a$ maps to a neighborhood of another vertical geodesic $x = a'$, where the size of the neighborhood depends only on K and C. It follows that there is a constant $B = B(K, C)$ such that each point on the image of $x = a$ is connected to a point of $x = a'$ by a horocyclic segment of length at most B.

Let $l = |p - q|$, let $\gamma = \text{Top}(R) = [p, q] \times s$, and note that $\text{Length}(\gamma) = l/s = \text{ecc}(R)$. Similarly, $\text{Length}(\gamma') = l'/s' = \text{ecc}(R')$, where $l' = |p' - q'|$ and $\gamma' = \text{Top}(R') = [p', q'] \times s'$. We need to compare the lengths of γ and γ' (see Figure 3).

Consider the path $f(\gamma)$, which lies below the horocycle $y = s' e^A$, and whose endpoints are connected to the vertical geodesics $x = p'$ and $x = q'$ by horocyclic segments of length at most B. Concatenating $f(\gamma)$ with these two horocyclic segments we obtain a path α. From Lemma 2.1 we have

$$\frac{1}{K} \cdot \text{Length}(\gamma) - C \leq \text{Length}(f(\gamma)) \leq K \cdot \text{Length}(\gamma) + C$$

and so

$$\frac{1}{K} \cdot \text{ecc}(R) - C \leq \text{Length}(f(\gamma)) \leq K \cdot \text{ecc}(R) + C.$$

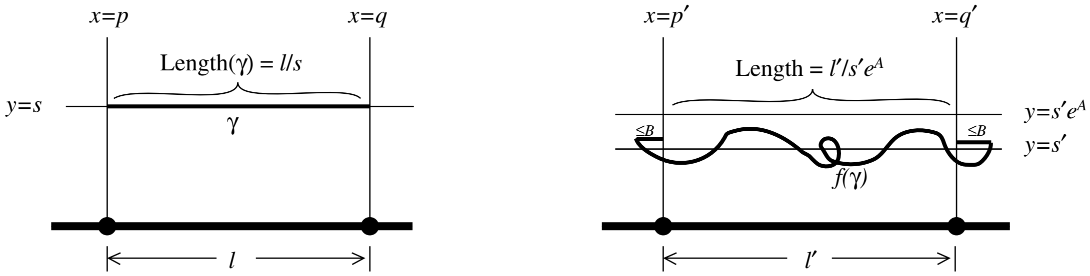

Figure 3: Connect $f(\gamma)$ to the vertical geodesics $x = p'$, $x = q'$ by short horocyclic segments, then project upward to the horocycle $y = s'e^A$.

We also have

$$\text{Length}(\alpha) \leq \text{Length}(f(\gamma)) + 2B$$

and by choosing $\text{ecc}(R) \geq K(2B + C)$ we have

$$\text{Length}(\alpha) \leq 2 \cdot \text{Length}(f(\gamma))$$

and therefore

$$\frac{1}{2}\,\text{Length}(\alpha) \leq K \cdot \text{ecc}(R) + C.$$

By further choosing $\text{ecc}(R) \geq C$ we obtain

$$\text{Length}(\alpha) \leq 2(K+1)\,\text{ecc}(R).$$

By projecting α straight upward to the horocycle $y = s'e^A$ we obtain

$$\text{Length}(\alpha) \geq \frac{l'}{s'e^A} = \text{ecc}(R')e^{-A}$$

and thus

$$\text{ecc}(R') \leq 2(K+1)e^A\,\text{ecc}(R).$$

The same argument using the coarse inverse of f shows that if $\text{ecc}(R')$ has the same lower bounds as $\text{ecc}(R)$ then

$$\text{ecc}(R) \leq 2(K+1)e^A\,\text{ecc}(R')$$

proving the lemma.

$\diamond$

As a first application of the lemma we show that f_ℓ is lipschitz in the large. Given a horocycle $y = s_0$, choose l_0 sufficiently large so that $l_0/s_0 \geq D$. For any $p_0 < q_0 \in \boldsymbol{R}$ such that $|p_0 - q_0| \geq l_0$, the first inequality of Lemma 5.4 applies to the rectangle $R = [p_0, q_0] \times [0, s_0]$. Setting $s_1 = s_0'$, $p_1 = f_\ell(p_0)$, $q_1 = f_\ell(q_0)$, and $l_1 = |p_1 - q_1|$ we therefore have

$$\frac{l_1}{l_0} \leq E\frac{s_1}{s_0}$$

Fixing the horocycle $y = s_0$ and its coarse image $y = s_1$, we have shown that f_ℓ is lipschitz in the large: if $|p-q| \geq l_0$ then $|f_\ell(p) - f_\ell(q)| \leq (Es_1/s_0)|p-q|$.

To show that f_ℓ is lipschitz, consider $\xi < \eta \in \boldsymbol{R}$ with $m = |\xi - \eta| < l_0$. Choose the points p_0, q_0 so that $p_0 < \xi, \eta < q_0$ and $l_0 = q_0 - p_0$. Choose a horocycle $y = r$ so close to the line at infinity that the rectangle $R = [\xi, \eta] \times [0, r]$ has eccentricity $m/r \geq D$, and hence the first inequality of Lemma 5.4 applies. Setting $m' = |f_\ell(\xi) - f_\ell(\eta)|$ we therefore have

$$\frac{m'}{m} \leq E\frac{r'}{r}$$

Now we find an upper bound for r'/r. We know that $r' \to 0$ as $r \to 0$, and so we may choose r so close to zero that $l_1/r' \geq D$. The second inequality of Lemma 5.4 therefore applies with $p = p_0$, $q = q_0$, and $s = r$, and we have

$$\frac{1}{E}\frac{r'}{r} \leq \frac{l_1}{l_0}$$

and so

$$\frac{m'}{m} \leq E^2\frac{l_1}{l_0} \leq E^3\frac{s_1}{s_0}$$

proving that f_ℓ is lipschitz.

A similar argument using the coarse inverse of f gives

$$\frac{m}{m'} \leq E^3\frac{s_0}{s_1}$$

and so

$$\frac{1}{E^3}s_1/s_0 \leq \frac{m'}{m} \leq E^3 s_1/s_0$$

Fixing, say, $s_0 = 1$ and letting $T = \log(s_1)$, we have

$$\frac{1}{E^3}e^T \leq \frac{m'}{m} \leq E^3 e^T$$

proving that f_ℓ is bilipschitz.

The above inequalities also show that for *any* horocycle $y = s_0$, setting $s_1 = s_0'$ we have

$$\frac{1}{E^6}e^T \le \frac{s_1}{s_0} \le E^6 e^T$$

and so

$$\log(s_0) + T - \log(E^6) \le \log(s_1) \le \log(s_0) + T + \log(E^6).$$

For any point $P = (x_0, y_0)$ in the upper half plane, we know that the point $f(P)$ lies in the A-neighborhood of the horocycle $y = y_0'$. The point $f(P)$ also lies in a certain bounded neighborhood of the geodesic $x = f_\ell(x_0)$ where the bound depends only on K, C. It follows that $d(f(P), (f_\ell(x_0), y_0'))$ is bounded above by a constant depending only on K, C, A. From the above inequality with $s_0 = y_0$, $s_1 = y_0'$ it also follows that

$$d((f_\ell(x_0), y_0'), (f_\ell(x_0), y_0 e^T)) \le |\log(E^6)|$$

completing the proof of the proposition. $\diamond$

Remark. It is easy to see that for any bilipschitz homeomorphism $\phi\colon \boldsymbol{R} \to \boldsymbol{R}$ and any $T \in \boldsymbol{R}$, the map $(x, y) \to (\phi(x), ye^T)$ is a quasi-isometry of $\boldsymbol{H}^2$. Proposition 5.3 therefore *characterizes* quasi-isometries that take horocycles based at ∞ to bounded neighborhoods of horocycles based at ∞.

6 The upper and lower boundaries of X_n

We have defined the upper boundary $\partial^u X_n$ and the map on upper boundaries induced by a quasi-isometry. We now define a metric on $\partial^u X_n$, and we study how the induced map acts on that metric.

Height function

Choose once and for all a *height function* on T_n, that is a continuous function $h\colon T_n \to \boldsymbol{R}$ which maps each oriented edge $E \subset T_n$ homeomorphically to an oriented interval in $\boldsymbol{R}$ of length $\log(n)$. We obtain by composition a height function $X_n \xrightarrow{\pi} T_n \xrightarrow{h} \boldsymbol{R}$ also denoted h. The choices in the definition of h may be fixed by requiring that the base point of X_n map to the origin of $\boldsymbol{R}$, and by requiring that for each edge $E \subset T_n$, the map $h \mid \pi^{-1}(E)$ is an isometry on each vertical geodesic segment. We assume that the base point of X_n lies on a branching horocycle, and hence for each branching horocycle $\sigma \subset X_n$ we have $h(\sigma) = k\log(n)$ for some integer k called the *combinatorial height* of σ.

Metric on $\partial^u X_n$

Given hyperbolic planes $Q_1, Q_2 \in \partial^u X_n$, define the distance between them to be $d(Q_1, Q_2) = e^{-h(\sigma)} = n^{-k}$, where the horocycle $\sigma = \partial(Q_1 \cap Q_2)$ has combinatorial height k. We leave the reader to check that this defines a locally compact metric on $\partial^u X_n$. In fact the metric d is an *ultrametric*, also called a *nonarchimedean metric*, in other words a metric satisfying $d(x,z) \leq \max\{d(x,y), d(y,z)\}$ for any $x, y, z \in \partial^u X_n$.

Note that $\partial^u X_n$ is isometric to the n-adic rational numbers $\mathbf{Q}_n$ with the usual n-adic metric. This is the space of all bi-infinite sequences of the form (a_i) where $a_i \in \{0, \ldots, n-1\}$ for $i \in \boldsymbol{Z}$, and $a_i = 0$ for i sufficiently close to $-\infty$. In number theory one thinks of (a_i) as the infinite series $\sum_i a_i n^i$ (although this is usually only studied when n is a prime). The distance between two distinct n-adic rational numbers $(a_i), (b_i) \in \mathbf{Q}_n$ is equal to n^{-I}, where I is the largest integer such that $a_i = b_i$ for all $i \leq I$. The isometry between $\partial^u X_n$ and $\mathbf{Q}_n$ can be exhibited by choosing, for each vertex $v \in T_n$, a bijective labelling of the set of outgoing edges at v by elements of the set $\{0, \ldots, n-1\}$, thereby assigning a label sequence (a_i) to each coherently oriented line in T_n. The labelling must be chosen so that some coherent line has a label sequence of all zeroes.

Lower boundary

We define the *lower boundary* of X_n, denoted $\partial_\ell X_n$, to be the common line at infinity of any two hyperbolic planes in X_n. That is, given hyperbolic planes $Q_1, Q_2 \subset X_n$, the inclusions $Q_1 \cap Q_2 \subset Q_i$, $i = 1, 2$ induce a bijection $\partial_\ell Q_1 \approx \partial_\ell Q_2$. Moreover, for any three hyperbolic planes we obtain a commutative triangle

$$\partial_\ell Q_1 \approx \partial_\ell Q_2 \approx \partial_\ell Q_3 \approx \partial_\ell Q_1.$$

We may therefore define $\partial_\ell X_n$ by pairwise identifying $\partial_\ell Q$ for all hyperbolic planes $Q \subset X_n$. The choice of a height function on X_n determines a specific isometry between Q and the upper half plane, taking the height zero horocycle on Q to the horocycle $y = 1$. This identification determines a specific Euclidean metric on $\partial_\ell Q$, and these metrics all agree, thereby defining a Euclidean metric on $\partial_\ell X_n$. A quasi-isometry $f \colon X_m \to X_n$ induces a well-defined map $f_\ell \colon \partial_\ell X_m \to \partial_\ell X_n$, as the reader may easily check.

Theorem 6.1 (Boundary maps are bilipschitz). *If $f \colon X_m \to X_n$ is a quasi-isometry, then the induced maps $f_\ell \colon \partial_\ell X_m \to \partial_\ell X_n$ and $f^u \colon \partial^u X_m \to \partial^u X_n$ are bilipschitz.*

Proof. We have already proved this for f_ℓ. Let K, C be quasi-isometry constants for f. There exists a unique map $\rho_j\colon X_j \to \boldsymbol{H}^2$ satisfying the following properties:

- ρ_j takes horocycles to horocycles.
- ρ_j is an isometry when restricted to each hyperbolic plane in X_j.
- ρ_j is normalized to take the base point of X_j to the point $(x, y) = (0, 1)$.

For each hyperbolic plane $Q \subset X_m$ define the map $\overline{f}_Q\colon \boldsymbol{H}^2 \to \boldsymbol{H}^2$ to be the composition

$$\boldsymbol{H}^2 \xrightarrow{\rho_m^{-1}} Q \xrightarrow{f} X_n \xrightarrow{\rho_n} \boldsymbol{H}^2.$$

Claim 6.2. *The map $\overline{f}_Q$ is at most a distance $R_1 = R_1(K, C)$ in the sup norm from a map of the form $(x, y) \to (f_\ell(x), ye^{T_Q})$, where the translation constant T_Q depends ostensibly on Q.*

To prove the claim, define $\hat{f}_Q$ to be the composition

$$\boldsymbol{H}^2 \xrightarrow{\rho_m^{-1}} Q \xrightarrow{f} X_n \to Q' = f^u(Q) \xrightarrow{\rho_n} \boldsymbol{H}^2$$

where the map $X_n \to Q'$ is the closest point projection. By Proposition 4.1 the maps $\overline{f}_Q$ and $\hat{f}_Q$ are at bounded distance in the sup norm. By Lemma 5.1 the map $\hat{f}_Q$ coarsely preserves horocycles. By Proposition 5.3 the map $\hat{f}_Q$ is a bounded distance in the sup norm from a map of the form $(x, y) \to (f_\ell(x), ye^{T_Q})$. All the distances in this argument depend only on K, C, proving the claim.

We can remove the dependence of the translation constant on Q, using the fact that any two hyperbolic planes intersect below some horocycle, and on that region of intersection the translation constants must nearly agree. More formally:

Claim 6.3. *There exists $T \in \boldsymbol{R}$ (depending on f) and $R_2 \geq 0$ (depending only on K, C) such that for each hyperbolic plane $Q \subset X_m$, the map $\overline{f}_Q$ is at most a distance R_2 in the sup norm from the map*

$$(x, y) \to (f_\ell(x), ye^T)$$

To prove this claim, set $\overline{f}_i = \overline{f}_{Q_i}$ and $T_i = T_{Q_i}$, and it suffices to bound $|T_1 - T_2|$ independent of the hyperbolic planes Q_1, Q_2. Choose a point $P \in Q_1 \cap Q_2$, and let $(x, y) = \rho_m(P)$. By the previous claim, for $i = 1, 2$ we have

$$d(\overline{f}_i(x, y), (f_\ell(x), ye^{T_i})) \leq R_1$$

However, $\overline{f}_1(x, y) = \overline{f}_2(x, y)$ and so

$$|T_1 - T_2| = d((f_\ell(x), ye^{T_1}), (f_\ell(x), ye^{T_2})) \le 2R_1$$

proving the claim.

To prove Theorem 6.1, consider hyperbolic planes $Q_1, Q_2 \in \partial^u(X_m)$ and $Q'_i = f^u(Q_i) \in \partial^u(X_n)$. Let $\sigma = \partial(Q_1 \cap Q_2)$ and let $\sigma' = \partial(Q'_1 \cap Q'_2)$. We must show that

$$\log(d(Q'_1, Q'_2)) - \log(d(Q_1, Q_2)) = -h(\sigma') + h(\sigma)$$

is bounded above and below by constants that depend only on f, not on Q_1 and Q_2. For each $p \in \sigma$, it follows from Lemma 5.1 that $|h(\sigma') - h(f(p))|$ is bounded depending only on K, C. Also, the second claim above shows that $|h(f(p)) - h(\sigma) - T|$ is bounded depending only on K, C. It follows that $|h(\sigma') - h(\sigma) - T|$ is bounded, finishing the proof of the theorem. $\diamond$

7 Proof of main theorem

Recall the statement:

Theorem 7.1 (Quasi-isometric iff commensurable). *For integers $m, n \ge 2$, the groups $BS(1, m)$ and $BS(1, n)$ are quasi-isometric if and only if they are commensurable. This happens if and only if there exist integers $r, j, k > 0$ such that $m = r^j$ and $n = r^k$.*

Proof. The group $\mathrm{BS}(1, m) = \langle a, b \mid aba^{-1} = b^m\rangle$ has a finite index subgroup isomorphic to $\mathrm{BS}(1, m^I)$, generated by a^I, b. This subgroup is the kernel of the homomorphism $\mathrm{BS}(1, m) \to \boldsymbol{Z}/I\boldsymbol{Z}$ taking $a \to 1, b \to 0$, and so the index equals I. The groups $\mathrm{BS}(1, r^j)$ and $\mathrm{BS}(1, r^k)$ are therefore commensurable, and hence quasi-isometric.

If $\mathrm{BS}(1, m)$ and $\mathrm{BS}(1, n)$ are quasi-isometric, then by the observation of Svarc-Milnor there is a quasi-isometry $f \colon X_m \to X_n$. By Theorem 6.1, f induces a map $f^u \colon \mathbf{Q}_m \approx \partial^u X_m \to \partial^u X_n \approx \mathbf{Q}_n$ which is a bilipschitz homeomorphism.

The *n-adic integers* $\boldsymbol{Z}_n$ are defined to be the subset of $\mathbf{Q}_n$ consisting of all $(a_i) \in \mathbf{Q}_n$ where $a_i = 0$ for $i < 0$. Given $a \in \{0, \ldots, n-1\}$ and $I \in \boldsymbol{Z}$, the set

$$C(a, I) = \{(b_i) \in \mathbf{Q}_n \mid b_i = 0 \quad \text{if} \quad i < I,\ b_I = a\}$$

is called a *clone of $\boldsymbol{Z}_n$*. Noting that $\boldsymbol{Z}_m$ has diameter 1 in $\mathbf{Q}_m$, it follows that the set $f^u(\boldsymbol{Z}_m)$ has bounded diameter in $\mathbf{Q}_n$. But every bounded diameter subset of $\mathbf{Q}_n$ is contained in some clone $C(a, I)$ of $\boldsymbol{Z}_n$. The clone $C(a, I)$ is isometric to $\boldsymbol{Z}_n$ with its metric scaled by n^{-I}. Thus, multiplying the metric on $f^u(\boldsymbol{Z}_m)$ by n^I, we obtain a bilipschitz embedding $\boldsymbol{Z}_m \to \boldsymbol{Z}_n$.

The proof of Theorem 7.1 is completed by applying the following theorem, whose proof is found in the appendix as Corollary 10.11:

Theorem 7.2 (D. Cooper). *For integers $m, n \geq 2$, if there is a bilipschitz embedding $\boldsymbol{Z}_m \to \boldsymbol{Z}_n$ then there exist integers $r, j, k > 0$ such that $m = r^j, n = r^k$.*

$\diamond$

8 The quasi-isometry group

Given a metric space X, the *quasi-isometry group* $\mathrm{QI}(X)$ is defined as follows. Define an equivalence relation on quasi-isometries where $f \equiv g$ if f and g have bounded distance in the sup norm. Let $\mathrm{QI}(X)$ be the set of equivalence classes of quasi-isometries from X to itself. Composition of quasi-isometries gives a well-defined group structure on $\mathrm{QI}(X)$. A quasi-isometry between two metric spaces induces an isomorphism between quasi-isometry groups. In particular, $\mathrm{QI}(\mathrm{BS}(1, n)) \approx \mathrm{QI}(X_n)$.

Let $\mathrm{Bilip}(Y)$ be the group of bilipschitz homeomorphisms of a metric space Y.

Theorem 8.1 (Quasi-isometry group). *$QI(BS(1, n))$ is isomorphic to*

$$Bilip(\boldsymbol{R}) \times Bilip(\mathbf{Q}_n).$$

Before proving this theorem we need to state a basic principle, dubbed the "rubber band principle" by Rich Schwartz. The trivial but useful idea is that if one takes any size chain of rubber bands connecting two distant points, and if each rubber band is stretched by a factor of at most K, then the whole chain is stretched by at most a factor of K.

Lemma 8.2 (Rubber Band Principle). *For all $L, M > 0$, there exists $C > 0$ so that the following holds: Suppose X and Y are path metric spaces and $f : X \to Y$ is a map. Suppose that there are collections of isometrically embedded subspaces $\mathcal{C}_X$ of X and $\mathcal{C}_Y$ of Y satisfying:*

- *Any two points in X can be connected by an M-quasigeodesic which is a concatenation of a finite number of subpaths, each lying in an element of $\mathcal{C}_X$; similarly for Y.*

- *f induces a one-to-one correspondence between elements of $\mathcal{C}_X$ and $\mathcal{C}_Y$.*

- *f restricts to a L-quasi-isometry between corresponding elements of $\mathcal{C}_X$ and $\mathcal{C}_Y$.*

Then $f : X \to Y$ is a C-quasi-isometry.

We leave the proof of the Rubber Band Principle as an easy exercise.

Proof of Theorem 8.1. Choosing a base point of X_n induces maps

$$\begin{aligned} \mathrm{QI}(X_n) &\mapsto \mathrm{Bilip}(\partial_\ell X_n) \approx \mathrm{Bilip}(\boldsymbol{R}) \\ f &\mapsto f_\ell \end{aligned}$$

and

$$\begin{aligned} \mathrm{QI}(X_n) &\mapsto \mathrm{Bilip}(\partial^u X_n) \approx \mathrm{Bilip}(\mathbf{Q}_n) \\ f &\mapsto f^u \end{aligned}$$

These two maps are obviously homomorphisms, and so we obtain a homomorphism $\mathrm{QI}(X_n) \to \mathrm{Bilip}(\partial_\ell X_n) \times \mathrm{Bilip}(\partial^u X_n)$.

To show that the map is onto, consider $\alpha \in \mathrm{Bilip}(\partial_\ell X_n), \beta \in \mathrm{Bilip}(\partial^u X_n)$. Choose a constant $T \in \boldsymbol{R}$ so that $d(\beta(Q), \beta(Q')) \leq e^T d(Q, Q')$ for all $Q, Q' \in \partial^u X_n$. Given $Q \in \partial^u X_n$, define the map $f_Q \colon Q \to Q' = \beta(Q)$ to be the composition

$$Q \xrightarrow{\rho_n} \boldsymbol{H}^2 \xrightarrow{\alpha \times \tau_T} \boldsymbol{H}^2 \xrightarrow{\rho_n^{-1}} Q'$$

where $(\alpha \times \tau_T)(x, y) = (\alpha(x), \tau_T(y)) = (\alpha(x), ye^{-T})$. Note that f_Q is a quasi-isometry from Q to Q', with constants that are independent of Q. Given $Q_1, Q_2 \in \partial^u(X_n)$, we shall show that the maps $f_i = f_{Q_i}$ agree on the overlap $Q_1 \cap Q_2$. Let $\sigma = \partial(Q_1 \cap Q_2)$ and let $\sigma' = \partial(Q_1' \cap Q_2')$. By the choice of T it follows that $h(\sigma') - h(\sigma) = -\log(d(\beta(Q_1'), \beta(Q_2'))) + \log(d(Q_1, Q_2)) \geq -T$, and so $h(\sigma') \geq h(\sigma) - T$. It follows that if $x \in Q_1 \cap Q_2$ then $h(x) \leq h(\sigma)$ and hence $h(f_1(x)) = h(f_2(x)) \leq h(\sigma')$, so $f_1(x)$ and $f_2(x)$ both lie in $Q_1' \cap Q_2'$. But the points $f_1(x), f_2(x)$ are both taken by ρ_n to the same point of $\boldsymbol{H}^2$, and ρ_n is $1-1$ on $Q_1' \cap Q_2'$, and so it follows that $f_1(x) = f_2(x)$.

We may now paste together the maps f_Q, as Q ranges over $\partial^u(X_n)$, to get a map $f \colon X_n \to X_n$. This map is a quasi-isometry on each hyperbolic plane Q, with constants independent of Q. By the Rubber Band Principle applied to the collection of isometrically embedded hyperbolic planes in X_n, we have that f is a quasi-isometry. It is clear that $f_\ell = \alpha$ and $f^u = \beta$, proving that $\mathrm{QI}(\mathrm{BS}(1,n))$ maps onto $\mathrm{Bilip}(\boldsymbol{R}) \times \mathrm{Bilip}(\mathbf{Q}_n)$.

To show that the map is injective, given a quasi-isometry $f\colon X_n \to X_n$, suppose that f_ℓ is the identity on $\partial_\ell X_n$ and f^u is the identity on $\partial^u X_n$. It follows that for each hyperbolic plane $Q \subset X_n$, $f(Q)$ is a bounded Hausdorff distance from Q, and so we obtain a horocycle preserving quasi-isometry from $\boldsymbol{H}^2$ to itself defined as the composition

$$\boldsymbol{H}^2 \xrightarrow{\rho_n^{-1}} Q \xrightarrow{f} X_n \to Q \xrightarrow{\rho_n} \boldsymbol{H}^2$$

where $X_n \to Q$ is the closest point projection. The quasi-isometry constants of this map depend only on those of f. This quasi-isometry of $\boldsymbol{H}^2$ is therefore a bounded distance from a quasi-isometry of the form $(x,y) \to (f_\ell(x), ye^T) = (x, ye^T)$ for some constant $T = T(Q)$. As remarked earlier, if $[a,b]$ is the smallest interval containing the bilipschitz constants for f_ℓ, then the interval $[a/e^T, b/e^T]$ is bounded above and below by positive constants depending only on the quasi-isometry constants of f. In the present situation where $a = b = 1$, it follows that $T = T(Q)$ is bounded independent of Q. The distance between (x, ye^T) and (x,y) is therefore bounded, and it follows that f is a bounded distance from the identity map in the sup norm. $\diamond$

Remark. The quasi-isometry group $\mathrm{QI}(G)$ has a natural topological group structure, and it can be proved that the isomorphism

$$\mathrm{QI}(\mathrm{BS}(1,n)) \approx \mathrm{Bilip}(\boldsymbol{R}) \times \mathrm{Bilip}(\mathbf{Q}_n)$$

is an isomorphism of topological groups.

9 Is BS$(1,n)$ a lattice?

Here is an easy argument to show that $\mathrm{BS}(1,n)$ is not a lattice in any 1-connected solvable real Lie group S. If it were then it must be cocompact since all lattices in a solvable Lie group are cocompact. Since S is simply connected it is actually contractible (the matrix exponential map is a diffeomorphism), hence $\Gamma\backslash S$ would be a $K(\Gamma,1)$ space. But $\mathrm{BS}(1,n)$ has (virtual) cohomological dimension 2, hence S would have to be two-dimensional. But the only connected two-dimensional (real) solvable Lie group admitting a lattice is $\boldsymbol{R}^2$, which gives a contradiction since $\mathrm{BS}(1,n), n \geq 2$ is not abelian.

We also know that $\mathrm{BS}(1,n)$ is not a subgroup of any group G acting properly discontinuously and cocompactly on any space X which is nonpositively curved in the CAT(0) sense. In particular $\mathrm{BS}(1,n)$ is not a subgroup of a cocompact lattice in any (real, p-adic, or product of real and p-adic) semisimple Lie group. The reason is that $\mathrm{BS}(1,n)$ contains an infinitely generated rank 1 abelian group Q, namely the kernel of the homomorphism $\mathrm{BS}(1,n) \to \boldsymbol{Z}$ taking $a \to 1$, $b \to 0$. The *translation number* of any $g \in G$ is defined to be $\tau_g = \inf_{x\in X} d(x, g(x))$. The number τ_g is bounded away from zero as g ranges over the infinite order elements of G, and $\tau_{g^n} = n\tau_g$. If Q were a subgroup of G, then Q would have elements with nonzero translation numbers arbitrarily close to zero, a contradiction.

We should note that, at least for prime p, $\mathrm{BS}(1,p)$ in some sense exhibits a mixture of real and p-adic behaviour as follows: $\mathrm{BS}(1,p)$ is an infinite index subgroup of $GL(2, \boldsymbol{Z}[1/p])$, which is a discrete subgroup of $GL(2,\boldsymbol{R}) \times GL(2,\mathbf{Q}_p)$ via the representation $M \mapsto (M, M)$. The group $PGL(2,\boldsymbol{R}) \times PGL(2,\mathbf{Q}_p)$ acts transitively by isometries on $\boldsymbol{H}^2 \times T$, where T is the Bruhat-Tits tree for $PGL(2,\mathbf{Q}_p)$ (see [Se]). Furthermore, $PGL(2,\boldsymbol{Z}[1/p])$ acts properly discontinuously and isometrically with cofinite volume (but not cocompactly) on $\boldsymbol{H}^2 \times T$ [Se]. By restriction we then obtain an isometric, properly discontinuous action of $\mathrm{BS}(1,p)$ on $\boldsymbol{H}^2 \times T$. This action has infinite covolume (as it must by the above comments).

10 Appendix: Bilipschitz homeomorphisms of Self-Similar Cantor Sets by D. Cooper

Introduction

The study of properties of a metric space which are preserved by bilipschitz homeomorphism occurs in the study of groups via a word metric. It was also studied in [CP] for a certain type of Cantor set embedded in the real line. The Cantor sets concerned are mild generalizations of the original middle-third Cantor set. This Cantor set has the basic property that it is the union of two exact copies of itself each scaled down in size by a factor of 1/3. The generalization allows finitely many linear scale factors. It is easy to see that the Hausdorff dimension of this type of Cantor set depends only on these scale factors. Now a bilipschitz homeomorphism preserves Hausdorff dimension, and so a natural question is what, if any, further invariants other than Hausdorff dimension are there. An almost complete answer was given in [CP], using invariants derived from the Hausdorff measure. This work was generalized to analogous Cantor sets in Euclidean space of dimension n by H. Vuong in his thesis, [Vu1], [Vu2]. In this section, we generalize in a different direction to abstract metric Cantor sets which possess a certain linear self-similarity structure. There is a further generalization to a much wider class of Cantor set, where the self-similarity structure is smooth rather than linear. This will not be dealt with here.

The main result is Theorem (10.6) which states that every bilipschitz homeomorphism between two Cantor sets of the type we consider has constant Radon-Nikodym derivative on some clopen. The reason is, roughly speaking, that a bilipschitz map only permits a limited amount of metric distortion. The self-similar bumpiness of a Cantor set forces non-linearity to come in discrete amounts. Thus a finite number of discrete distortions exhausts the amount of metric distortion possible. Actually, this argument should be made at the level of Hausdorff measure rather than at the level of the metric.

The main ideas are already in [CP], except that there the topological properties of the real line are heavily used. This obstacle was overcome in [Vu1] with the introduction of *separation* defined below.

Cantor Sets

A map $f : X \longrightarrow Y$ between metric spaces is called a *similarity map* if there is a constant $K > 0$ such that $d_Y(fx_1, fx_2) = K\ d_X(x_1, x_2)$ for all x_1, x_2 in X. Consider a Cantor set C equipped with a metric d_C then a *clone structure* on C is the following data. Let $A_1, A_2, \cdots, A_n$ be a partition of C into clopens and for each A_i let τ_i be a similarity map of A_i onto C. We will call such a Cantor set a *clone Cantor set.*

Any composition of the maps τ_i is called a *clone map,* (we include the identity map as a clone map) thus the set of clone maps forms a semi-group. A *clone* of C is the pre-image of C under a clone map. Thus C and each of the A_i are clones. More generally any pre-image of a clone under a clone map is also a clone. For every clone there is a clone map sending the clone onto C. The clone structure is not unique, since one may replace any clone by finitely many clones which partition it. The *level* of a clone is defined inductively. The clone C is level 0, and a clone has level n if the smallest clone properly containing it has level $n - 1$.

Two clones are either disjoint or else one is a subset of the other. Thus one may construct a tree where the vertices are clones and there is an edge between a clone of level n and any clone of level $n + 1$ that it meets (hence contains.)

Lemma 10.1. *Every clopen is a finite union of clones.*

Given a subset A of a Cantor set C define the *separation* of A to be the minimum distance between A and its complement,

$$sep(A) = \inf \{d_C(x,y) \ : \ x \in A, \ \ y \in C - A \ \}.$$

By convention, $sep(C) = \infty$. Recall that the *diameter* of a metric space is the supremum of the distance between points in the metric space. Now define the *relative separation* of A to be $rel(A) = sep(A)/diam(A)$. We will make frequent use of the minimum separation between level-1 clones:

$$\alpha(C) = \min\{ \ sep(A) \ : A \text{ is a level-1 clone of } C \ \}.$$

An *ϵ-cover* of a metric space, X, is a cover by sets each of which has diameter at most ϵ. Given a cover $\mathcal{U} = \{ \ U_i \ \}$ of X and $\delta > 0$ define

$$\mathcal{H}_\delta(\mathcal{U}) = \sum_i \ [diam(U_i)]^\delta.$$

The *Hausdorff-δ outer measure* of X is the limit as $\epsilon \to 0$ of the $inf \ \mathcal{H}_\delta(\mathcal{U}_\epsilon)$ where the infimum is taken over all ϵ-covers, $\mathcal{U}_\epsilon$, of X. We will assume that the Hausdorff measure μ_C of C in its Hausdorff dimension is finite and not zero.

A map $f : (X, \mu_X) \longrightarrow (Y, \mu_Y)$ between measure spaces is called *measure linear* if there is a constant $K > 0$ such that for every measurable set A in X we have $\mu_Y(fA) = K\mu_X(A)$. In other words, the Radon-Nikodym derivative of f is K everywhere. We also describe this by saying that the map is linear at the level of measure theory. A similarity map between metric spaces is measure linear for Hausdorff measures, thus in particular clone maps are measure linear. Similarly the notion of bilipschitz at the level of measure theory may be defined. Let f be a bilipschitz homeomorphism between two such Cantor sets with bilipschitz constant K. Then f is K^δ-bilipschitz for δ-dimensional Hausdorff measure. It follows that the two Cantor sets have the same Hausdorff dimension.

We will use the term *approximates* to mean that the ratio of two numbers is bounded away from 0 and ∞ by numbers depending only on certain constants.

The following says that the separation of a clone approximates its diameter.

Lemma 10.2. *Given a clone Cantor set C there is a constant $\xi(C) > 1$ such that for every clone A of C with $A \neq C$ we have*

$$\xi(C)^{-1} \ \leq \ rel(A) \ \leq \ \xi(C).$$

Proof. Let x, y be points in A and $C - A$ of minimal distance apart and let B be the smallest clone containing containing both of them. Let τ be the clone map taking B onto C, thus $rel(\tau A) = rel(A)$. After replacing A by τA we may assume that $B = C$ and τ is the identity.

Now x, y are in different level-1 clones by choice of B so $d_C(x, y) \geq \alpha(C)$. Let D be the smallest clone properly containing A. Suppose that $diam(D) < \alpha$. Then there every point in $D - A$ is closer to A than x is to y, a contradiction. Thus $diam(D) \geq \alpha$ and there are only finitely many clones this large, so that there is a bound on $level(D)$. By choice of D we have $level(A) = level(D) + 1$ and so there are only finitely many possibilities for A. Hence $\xi(C)$ is the maximum of a finite set. $\diamond$

The next result says that the smallest clone containing a given set has approximately the same diameter as that set.

Lemma 10.3. *Given a clone Cantor set C there is a constant $\mu(C)$ with the following property. Let X be any subset of C of positive diameter and let A be the smallest clone of C containing X. Then*

$$diam(X) \leq diam(A) \leq diam(X)\mu(C).$$

Proof. Let τ be the clone map sending A onto C. Then $\tau(X)$ is not contained in any level-1 clone of C by definition of A. Thus $\tau(X)$ contains points in two different level-1 clones of C, and so $diam(\tau X) \geq \alpha(C)$. Hence

$$\frac{diam(A)}{diam(X)} = \frac{diam(\tau A)}{diam(\tau X)} \leq \frac{diam(C)}{\alpha(C)}.$$

Taking $\mu(C) = diam(C)\alpha(C)^{-1}$ gives the result. $\diamond$

The clone structure of a clone Cantor set provides a natural way to magnify phenomena, in particular a way to magnify subsets. Given a subset, A, of a clone Cantor set C there is some smallest clone, B, which contains A. Now B is a copy of C but linearly scaled down. The clone structure provides an identification, $\tau : B \longrightarrow C$, of B with C. Then τA is a magnified copy of A in C, in particular the diameter of τA is at least α. This technique is used to enlarge phenomena which happen on a very small scale.

Lemma 10.4. *Given $K > 1$, $\epsilon > 0$ and clone Cantor sets C, C' there is a constant M with the following property. Let f be any K-bilipschitz map of C onto a clopen in C', and suppose that $rel(fC) \geq \epsilon$. Let A be any clone of C. If P' is the smallest clone of C' containing fA and τ' is the clone map taking P' onto C' then there are at most $M = M(C, C', K, \epsilon)$ possibilities for the image $\tau'(fA)$ as a subset of C'.*

Proof. Since f is K-bilipschitz $diam(fA) \leq K\ diam(A)$. Choose x in fA and y in $C' - fA$ minimizing the distance between x and y. If y is in $C' - fC$ then $sep(fA) \geq sep(fC)$ and $diam(fA) \leq diam(fC)$ so $rel(fA) \geq \epsilon$. Otherwise if y is in fC then,

$$sep(fA) = d_{C'}(x, y) \geq K^{-1} d_C(f^{-1}x, f^{-1}y) \geq K^{-1} sep(A).$$

Using lemma (10.2) we get that $sep(fA) \geq K^{-1} diam(A)\xi(C)^{-1}$. Hence in either case

$$rel(fA) \geq min\left(\frac{K^{-1} diam(A)\xi(C)^{-1}}{K\ diam(A)}, \epsilon\right) \equiv \gamma.$$

Since P' is the smallest clone containing fA it follows that C' is the smallest clone containing $\tau' fA$ thus $diam(\tau'(fA)) \geq \alpha(C')$. Since τ' is a similarity, $rel(\tau'(fA)) \geq rel(fA) \geq \gamma$, thus

$$sep(\tau'(fA)) = rel(\tau'(fA)) diam(\tau'(fA)) \geq \gamma\alpha(C') \equiv \beta.$$

Consider a clone B of C' contained in $\tau'(fA)$ and not a proper subset of any other clone contained in $\tau'(fA)$. We will call such a clone a *maximal clone* in $\tau'(fA)$. Let D be a clone containing B with $level(D) = level(B) - 1$. There is a point $y \in D$ which is not in $\tau'(fA)$ for otherwise B would not be maximal. Hence $sep(\tau'(fA)) \leq diam(D)$. Combining this with the previous inequality gives $diam(D) \geq \beta$. The argument of the last paragraph of lemma (10.2) again shows that $level(D)$ is bounded and so there only a bounded number of possibilities for B. This bound depends only on β and C'.

Thus there is a bounded number of choices for the collection of maximal clones. By lemma (10.1) the clopen $\tau'(fA)$ is a finite union of (maximal) clones, hence there are a bounded number of possibilities for $\tau'(fA)$. $\diamond$

Corollary 10.5. *Given $K > 1$, $\epsilon > 0$ and clone Cantor sets C, C' there is a finite set S of positive real numbers with the following propery. Suppose that f is any K-bilipschitz map of C onto a clopen in C', and suppose that $rel(fC) \geq \epsilon$. If B is any clone of level n containing a clone D of level $n+1$ in C then $\mu_{C'}(fD)/\mu_{C'}(fB)$ is in S.*

Proof. Let P' be the smallest clone of C' containing fB and τ' the clone map of P' onto C'. Since τ' is a similarity

$$\frac{\mu_{C'}(fD)}{\mu_{C'}(fB)} = \frac{\mu_{C'}(\tau'(fD))}{\mu_{C'}(\tau'(fB))}.$$

Then there are finitely many possible images $\tau'(fB)$ by lemma (10.4). We will show that there are only finitely many possibilities for the set $\tau'(fD)$, then the result follows.

The ratio $diam(D)/diam(B)$ is one of the finitely many numbers $diam(A)/diam(C)$ for A a level-1 clone of C, and is thus bounded away from 0. By lemma (10.3) $diam(P')$ approximates $diam(fB)$ which approximates $diam(B)$ since f is bilipschitz. Let Q' be the smallest clone of C' containing fD. Then $diam(Q')$ approximates $diam(D)$ and so the ratio $diam(Q')/diam(P')$ approximates $diam(D)/diam(B)$ and so is bounded away from 0. Now

$$\frac{diam(Q')}{diam(P')} = \frac{diam(\tau'(Q'))}{diam(\tau'(P'))}$$

and $\tau'(P') = C'$ hence $diam(\tau'Q')$ is bounded away from 0.

There are only finitely many clones $\tau'Q'$ in C' which are this large. If we map Q' onto C' by a clone map ϕ, then $\phi\tau'(fD)$ is one of finitely many possibilities by lemma (10.4), and since there are only finitely many possible choices for ϕ there are only finitely many possible sets $\tau'(fD)$. $\diamond$

The following is the basic result in this subject.

Theorem 10.6. *Suppose that C, C' are clone Cantor sets and that f is a bilipschitz map of C onto a clopen subset of C'. Then there is a clopen A in C such that the restriction $f|A$ of f to A is measure linear.*

Proof. Given a measurable subset A of C of positive measure define the *mass ratio* of A to be

$$MR(A) = \frac{\mu_{C'}(fA)}{\mu_C(A)}.$$

This is the multiplicative factor by which f changes the mass of A. Since f is bilipschitz the mass ratio function is bounded above, let L be the supremum. We remark that

$$MR(A \cup B) \leq \max\{MR(A), MR(B)\}$$

and the inequality is strict unless the mass ratios are equal. Thus given $\epsilon > 0$ there is a clone A of C with $MR(A) > L - \epsilon$. We will show that for ϵ small enough, that $f|A$ is measure linear.

Assertion. There is a finite set S of real numbers such that for every pair of clones $E \subset D$ of C with $level(E) = level(D)+1$ then $MR(E)/MR(D)$ is in S.

Assuming this, suppose that $f|A$ is not measure linear. Then there is some clone E in A such that $MR(A) < MR(E)$. Otherwise for every clone E in A we would have $MR(E) \leq MR(A)$ but this would imply, by the remark, there is always equality and this implies that $f|A$ is measure linear.

Thus if $f|A$ is not measure linear we may choose a clone E in A with $MR(A) < MR(E)$ and with $n = level(E)$ minimal. Note that if D is any clone in A of level less than n then $MR(D) = MR(A)$. Let D be a clone in A containing E with $level(D) = n - 1$. Then $MR(E)/MR(D)$ is in S. Define λ to be the minimum of the elements of S which are larger than 1. Thus $MR(E)/MR(D) \geq \lambda$. If we choose ϵ such that $(L-\epsilon)\lambda > L$ then

$$MR(E) > \lambda MR(D) = \lambda MR(A) > \lambda(L-\epsilon) > L$$

which contradicts the definition of L, proving that $f|A$ is measure linear.

To prove the assertion observe that

$$\frac{MR(E)}{MR(D)} = \frac{\mu_{C'}(fE)/\mu_{C'}(fD)}{\mu_C(E)/\mu_C(D)}.$$

Since $level(E) = level(D) + 1$, as discussed in the proof of (10.5), there are only finitely many values for the ratio $\mu_C(E)/\mu_C(D)$. By (10.5), applied with $\epsilon = rel(fC)$, there are only finitely many values for $\mu_{C'}(fE)/\mu_{C'}(fD)$, and this proves the assertion. $\diamond$

One may now define various invariants of the bilipschitz type of clone Cantor sets. Let C be a clone Cantor set with Hausdorff measure μ_C in its Hausdorff dimension. We will call two sets of real numbers A and B *similar* if there are non-zero scalars α, β such that

$$\alpha A \subset B \qquad and \qquad \beta B \subset A.$$

The *clopen invariant* of C is the similarity class of the countable set of real numbers $\{\mu_C(A) : \ A \text{ a clopen in } C\ \}$.

Corollary 10.7. *If C, C' are clone Cantor sets and there is a bilipschitz map of C onto a clopen in C', then C and C' have the same clopen invariant.*

Proof. By the theorem, there is a clopen, hence a clone, A in C and a measure linear map of A onto a clopen in C'. The set A is a clone Cantor set and the clone map of A onto C shows that it has the same clopen invariant as C. Let α be the Radon-Nikodym derivative of the measure linear map of A into C'. The image of a clopen B in A is a clopen in C' with α times the measure of B. Using that the inverse of a bilipschitz map is bilipschitz, one obtains the reverse relationship. $\diamond$

There is another invariant which measures the expansions of the similarity self maps of a clone Cantor set. A clone map is an example of such a map. In order to obtain a notion which is an invariant of bilipschitz homeomorphism, one must allow for conjugating a map which is a metric similarity by a bilipschitz map which is measure preserving, and this leads to the following.

Definition. Given $K > 1$ we will call a map $\psi : (X, d_X) \longrightarrow (Y, d_Y)$ between metric spaces *K-quasi-conformal* if

$$K^{-1} \leq \frac{d(\psi x, \psi y)/d(x,y)}{d(\psi x, \psi z)/d(x,z)} \leq K$$

for all distinct points x, y, z. Thus a similarity is 1-quasi-conformal. A K-bilipschitz map is K^2-quasi-conformal. The composition of K and K' quasi-conformal maps is (KK')-quasi-conformal. Suppose that x, y, w are distinct then

$$K^{-1} \leq \frac{d(\psi y, \psi w)/d(y,w)}{d(\psi y, \psi x)/d(y,x)} \leq K$$

thus

$$K^{-2} \leq \frac{d(\psi y, \psi w)/d(y,w)}{d(\psi x, \psi z)/d(x,z)} \leq K^2.$$

Regard x and z as fixed, then we see that ψ is M-bilipschitz with $M = K^2 \max(L, 1/L)$ where $L = d(\psi x, \psi z)/d(x,z)$.

Lemma 10.8. *Suppose that $\psi : X \longrightarrow Y$ is a K-quasi-conformal bijection. Then ψ is M-bilipschitz with $M = K^4 \max(N, 1/N)$ where $N = diam(Y)/diam(X)$.*

Proof. Choose x, y, z, w in X with $d_X(y,w) = diam(X)$ and $d_Y(\psi x, \psi z) = diam(Y)$. Then

$$K^{-2} \leq \frac{d_Y(\psi y, \psi w)/d_X(y,w)}{d_Y(\psi x, \psi z)/d_X(x,z)} = \frac{d_Y(\psi y, \psi w)}{d_X(y,w)} \frac{d_X(x,z)}{diam(Y)}.$$

Now $d_X(x,z) \leq diam(X)$ thus

$$(K^2 N^{-1})^{-1} = K^{-2} \frac{diam(Y)}{diam(X)} \leq \frac{d_Y(\psi y, \psi w)}{d_X(y,w)} \leq \frac{diam(Y)}{diam(X)} = N \leq K^2 N.$$

It follows that ψ is M-bilipschitz with $M = K^2 \max(K^2 N, K^2 N^{-1})$. $\diamond$

Definition. A map $\sigma : (X, d_X) \longrightarrow (X, d_X)$ of a metric space (X, d_X) to itself has *bounded distortion* or is *BD* if there is a constant $K > 1$ such that for every integer $n > 0$ the map σ^n is K-quasi-conformal. Informally, the iterates of σ are almost similarities. For example the conjugate of a contracting similarity by a bilipschitz map is BD. This definition is not quite equivalent to the one in [CP], but is better for present purposes.

Definition. A *local similarity* of a clone Cantor set C is a continuous map $\sigma : U \longrightarrow U$ defined on a non-empty clopen U of C which is a homeomorphism onto $V = \sigma U$. We require that V is a clopen contained in U, and the following properties:
(1) σ has a unique fixed point x called a *similarity point* of C.
(2) σ is measure linear with respect to Hausdorff measure. The Radon-Nikodym derivative of σ is called the *scale factor* of σ.
(3) σ is BD.

If x is a similarity point of C we define the *similarity group* $S_x(C)$ of C at x to be the subgroup of the multiplicative group $\boldsymbol{R}^+$ generated by the scale factors of the local similarities of C fixing x.

Corollary 10.9. *Suppose that C and C' are clone Cantor sets and there is a bilipschitz homeomorphism of C onto a clopen subset of C'. If x is a similarity point of C there is a similarity point y in C' such that $S_x(C) = S_y(C')$.*

Proof. By theorem (10.6) there is a clopen A of C on which the given bilipschitz homeomorphism

$$\phi : C \longrightarrow C'$$

is measure linear. The orbit of x under the semi-group of clone maps of C contains a point in every clone and hence contains a point z in A. Clearly $S_x(C) = S_z(C)$. We claim that the conjugate of a local similarity by a measure-linear bilipschitz homeomorphism is a local similarity. To verify this, let $\sigma : A \longrightarrow A$ be a local similarity defined on the clopen A of C and with fixed point z. Let $\sigma' = \phi \circ \sigma \circ \phi^{-1}$ then σ' has a unique fixed point $\phi(z)$. Since ϕ is a homeomorphism, ϕA is a clopen of C and so the domain of σ' is the clopen ϕA. Since $\phi|A$ and σ are measure linear, σ' is measure linear. Since σ is BD then the iterates σ^n are all K-quasi-conformal. Also ϕ is bilipschitz and thus M-quasi-conformal for some M. Thus

$$\sigma'^n = \phi \circ \sigma^n \circ \phi^{-1}$$

is (M^2K)-quasi-conformal. Hence σ' is BD. ◇

An Application to n-adic Cantor sets

Given $n > 1$ let C_n, or just C, be the Cantor set of sequences $(x_0, x_1, \cdots)$ where x_i is an element of $\{0, \cdots, n-1\}$. The metric on C_n is defined by $d_n(x, y) = n^{-p}$ if p is minimal subject to $x_p \neq y_p$. Thus C has diameter 1. The shift map σ mapping C onto itself is defined by $\sigma(x_0, x_1, \cdots) = (x_1, x_2, \cdots)$. Define clones A_i for $0 \leq i \leq n-1$ to be the subsets of C with first digit i. Define a clone map $\tau_i = \sigma|A_i$ which linearly maps A_i onto C.

Proposition 10.10. *The Hausdorff dimension of the Cantor set C is 1 and it has measure 1 in this dimension.*

Proof. There is a map ϕ of C onto the unit interval $[0, 1]$ given by regarding a sequence $(x_0, x_1, \cdots)$ as an expansion $0.x_0x_1x_2\cdots$ in base-n of a real number. Consider a subset A in C with diameter R. Choose p so that $n^{-p} \leq R < n^{1-p}$. Then all points in A consist of sequences which have the same digits up to and including x_{p-1} and therefore ϕA has Euclidean diameter at most n^{-p}. Thus

$$diameter_C(A) \geq Euclidean\ diameter(\phi A).$$

Thus the sum of the diameters of the sets in any cover of C is at least 1 and so the Hausdorff 1-measure of C is at least 1. Given $\epsilon > 0$ there is a cover of C by level-k clones each of which has diameter in the d_C-metric $n^{-k} < \epsilon$. There are n^k such clones, so the sum of their diameters is 1. Therefore the Hausdorff 1-measure is at most 1. $\diamond$

Since clone maps are measure linear it follows that the measure of a clone of C equals the Euclidean length of its image under ϕ. Thus ϕ is measure preserving. Hence the clopen invariant of C_n is the similarity class of $[0, 1] \cap \boldsymbol{Z}[1/n]$.

Corollary 10.11. *There is a bilipschitz homeomorphism of the clone Cantor set C_n onto a clopen subset of C_m if and only if there are integers $i, j > 0$ such that $n^i = m^j$.*

Proof. If $n^i = m^j$ then it is clear that the Cantor sets are bilipschitz homeomorphic. For the converse, suppose that σ is a local similarity of C_n with scale factor $s < 1$ for Hausdorff measure. We first show that s is a rational power of n.

Let A be a clone in the domain of σ which contains the fixed point of σ. Then some iterate, σ^m, of σ maps A into A. We replace σ with $\sigma^m|A$. The scale factor of σ^m is s^m and it suffices to prove that s^m is a rational power of n. By this means we may suppose that the domain of σ is a clone A.

Since σ is BD there is $L > 1$ such that for all $p > 0$ we have that σ^p is L-quasi-conformal. Define $\alpha \equiv \alpha(C_n)$. The lemma (10.4) provides a constant, $M = M(C_n, C_n, K, \epsilon)$, for the case that C and C' are both C_n, and $K = L^6\alpha^{-1}$, and $\epsilon = K^{-1}\alpha$. Given an integer $a > 0$, define a new metric, d_{new} on C_n which is n^a times the n-adic metric thus $d_{new} = n^a d_n$. There is a unique clone, D, of diameter 1 in the d_{new}-metric which contains the unique fixed point of σ. We may choose a large enough that D is contained in $\sigma^{2M}A$. Then D with the metric d_{new} is a clone Cantor set isometric to C_n with the n-adic metric. Also $\sigma|D$ is a local similarity of D with scale factor s. Also for all $p > 0$ we have σ^p is L-quasi-conformal. We now replace (C_n, d_n) and σ by (D, d_{new}) and $\sigma|D$. By means of this we may assume our original local similarity, σ, admits an extension to a local similarity, $\overline{\sigma}$, defined on all of a clone Cantor set, E, of large diameter containing C_n as a clone. Furthermore we may assume that the image of $\overline{\sigma}^{2M}$ contains C_n.

Suppose that A_p is the smallest clone containing $\sigma^p C_n$. If τ_p is the clone map taking A_p onto C_n, define $\psi_p \equiv \tau_p\sigma^p$ then $\psi_p C_n$ is not contained in any level-1 clone of C_n thus $diam(\psi_p C_n) \geq \alpha$. The clone map τ_p has domain the clone A_p which it maps bijectively to C_n. Define $B_p = (\overline{\sigma})^{-p}A_p$ and define a bijection

$$\overline{\psi}_p \equiv \tau_p\overline{\sigma}^p : B_p \longrightarrow C_n.$$

Since $\overline{\sigma}^p$ is L-quasi-conformal and τ_p is linear $\overline{\psi}_p$ is also L-quasi-conformal. Thus, in particular, ψ_p is L-quasi-conformal. Apply lemma (10.8) to

$$\psi_p : X \equiv C_n \longrightarrow Y \equiv \psi_p C_n.$$

Then $diam(X) = 1$ and $diam(Y) = diam(\psi_p C_n) \geq \alpha$ so $\alpha \leq N \leq 1$ thus ψ_p is $(L^4\alpha^{-1})$-bilipschitz. Since $\overline{\psi}_p$ is L-quasi-conformal and $(L^4\alpha^{-1})$-bilipschitz on C_n it follows that $\overline{\psi}_p$ is K-bilipschitz with $K = L^6\alpha^{-1}$.

We wish to apply the lemma (10.4) to the maps $\psi_p : C_n \longrightarrow C_n$ with $1 \leq p \leq 2M$. These maps are all K-bilipshitz. We claim that $sep(\psi_p C_n) \geq \epsilon$. Since $p \leq 2M$ we have $\overline{\sigma}^p E$ contains C_n thus $\overline{\psi}_p(B_p) = C_n$. Since $\overline{\psi}_p$ is K-bilipschitz $sep(\psi_p C_n)$ is at least $K^{-1}d_E(C_n, B_p - C_n)$. Now $d_E(C_n, B_p - C_n) \geq \alpha$ because C_n is a clone in E and the separation of C_n in E is therefore bigger than the separation of any clone contained in C_n. Thus $sep(\psi_p A) \geq K^{-1}\alpha = \epsilon$, which proves the claim.

By lemma (10.4) there are $M = M(C_n, C_n, K, \epsilon)$ possibilities for $\psi_p C_n$. So for some $1 \leq k < j \leq 2M$ we have that $\psi_j C_n = \psi_k C_n$. The metric scale factor of every clone map of C_n is an integral power of n. Since the Hausdorff dimension of C_n is 1 this means that the scale factor for Hausdorff measure of a clone map is also an integral power of n. Thus τ_j, τ_k are measure linear with scale factors for Hausdorff measure which are powers of n. Thus the scale factor for Hausdorff measure for ψ_j and ψ_k is $n^a s^j = n^b s^k$ and so s is a rational power of n.

By (10.9) we have that s is also the scale factor of a local similarity of C_m. The above then implies that s is also a rational power of m. $\diamond$

References

[BG] M. Bridson and S. Gersten, *The optimal isoperimetric inequality for torus bundles over the circle*, Quart. J. Math. Oxford Ser. (2) **47** (1996), 1–23.

[Ca] J. Cannon, *The theory of negatively curved spaces and groups*, Ergodic theory, symbolic dynamics, and hyperbolic spaces (C. Series T. Bedford, M. Keane, ed.), Oxford Univ. Press, 1991.

[CC] J. Cannon and D. Cooper, *A characterization of cocompact hyperbolic and finite-volume hyperbolic groups in dimension three*, Trans. AMS **330** (1992), 419–431.

[Ch] R. Chow, *Groups quasi-isometric to complex hyperbolic space*, Trans. AMS **348** (1996), no. 5, 1757–1769.

[CP] D. Cooper and T. Pignataro, *On the shape of Cantor sets*, J. Diff. Geom. **28** (1988), 203–221.

[ECH+] D.B.A. Epstein, J. Cannon, D.F. Holt, S. Levy, M.S. Paterson, and W.P. Thurston, *Word processing and group theory*, Jones and Bartlett, 1992.

[EF] A. Eskin and B. Farb, *Quasi-flats and rigidity in higher rank symmetric spaces*, Univ. of Chicago, October 1995, preprint.

[EG] L. Evans and R. Gariepy, *Measure theory and fine properties of functions*, Studies in Advanced Mathematics, CRC press, 1992.

[Es] A. Eskin, *Quasi-isometric rigidity for nonuniform lattices in higher rank*, (1996), preprint.

[Fa] B. Farb, *The quasi-isometry classification of lattices in semisimple Lie groups*, in preparation.

[FS] B. Farb and R. Schwartz, *The large-scale geometry of Hilbert modular groups*, J. Diff. Geom. (1996), to appear.

[Ge] S. Gersten, *Problems in geometric group theory*, Max dehn seminar preprint, Univ. of Utah, 1994.

[Gr1] M. Gromov, *Infinite groups as geometric objects*, Proceedings of the International Congress of Mathematicians (Warsaw, 1983) 385–392, PWN, Warsaw, 1984.

[Gr2] M. Gromov, *Asymptotic invariants of infinite groups*, Geometric Group Theory (G. Niblo and M. Roller, eds.), LMS Lecture Notes, vol. 182, Cambridge Univ. Press, 1993.

[JS] W. Jaco and P. Shalen, *Seifert fibered spaces in 3-manifolds*, Mem. AMS **21** (1979), no. 220.

[KaL] M. Kapovich and B. Leeb, *On asymptotic cones and quasi-isometry classes of fundamental groups of* 3*-manifolds*, Geom. Funct. Anal. **5** (1995), no. 3, 582–603.

[KlL] B. Kleiner and B. Leeb, *Rigidity of quasi-isometries for symmetric spaces of higher rank*, preprint, January 1995.

[Kr] P.H. Kropholler, *A note on centrality in 3-manifold groups*, Math. Proc. Camb. Philos. Soc. **107** (1990), 261–266.

[Me] G. Mess, *The Seifert conjecture and groups which are coarse quasi-isometric to planes*, preprint.

[Mi] J. Milnor, *A note on curvature and fundamental group*, J. Diff. Geom. **2** (1968), 1–7.

[Mo] G.D. Mostow, *Strong rigidity of locally symmetric spaces*, Annals of Math. Studies, no. 78, Princeton Univ. Press, 1973.

[Pa] P. Pansu, *Metriques de Carnot-Caratheodory et quasiisometries des espaces symmetriques de rang un*, Ann. of Math. **129** (1989), 1–60.

[Ri] E. Rieffel, *Groups coarse quasi-isometric to* $\boldsymbol{H}^2 \times \boldsymbol{R}$, Ph.D. thesis, UCLA, 1993.

[Sch1] R. Schwartz, *Quasi-isometric rigidity and diophantine approximation*, preprint, 1995.

[Sch2] R. Schwartz, *The quasi-isometry classification of rank one lattices*, IHES Sci. Publ. Math. (1996), to appear.

[Se] J. P. Serre, *Trees*, Springer-Verlag, New York, 1980.

[Tu] P. Tukia, *On quasi-conformal groups*, J. D'Anal. Math. **46** (1986), 318–346.

[Vu1] H. Vuong, *Classification of self-similar Cantor sets of* $\boldsymbol{R}^n$ *up to bi-Lipschitz equivalence*, Ph.D. thesis, UCSB, 1994.

[Vu2] H. Vuong, *Measure linearity of bi-Lipschitz maps on self-similar Cantor sets*, The Real Analysis Exchange (1996), to appear.

[We] S. Weinberger, personal communication.

Benson Farb:
Dept. of Math., Univ. of Chicago
5734 University Ave.
Chicago, Il 60637
E-mail: farb@math.uchicago.edu

Lee Mosher:
Dept. of Math. and C.S., Rutgers Univ., Newark
Newark, NJ 07102
E-mail: mosher@andromeda.rutgers.edu

Daryl Cooper:
Dept. of Math., Univ. of California, Santa Barbara
Santa Barbara, CA 93106
E-mail: cooper@math.ucsb.edu