Quasi-flats in $\mathbf{H}^2 \times \mathbf{H}^2$

Alex Eskin and Benson Farb *
University of Chicago

February 4, 1996

Abstract

The purpose of this paper is to exhibit most of the main ideas of [EF] in the familiar special case of $\mathbf{H}^2 \times \mathbf{H}^2$. We prove here that any (coarse) quasi-isometric embedding of $\mathbf{R}^2$ into $\mathbf{H}^2 \times \mathbf{H}^2$ is close to a finite union of flats. The paper [EF] deals with the general case of a quasi-isometric embedding $\mathbf{R}^n \to X$ for an arbitrary rank n symmetric space X of noncompact type.

1 Introduction

The main goal of this paper is to prove the following :

Theorem 1.1 (Quasi-flats are close to a finite union of flats). *Let $\phi : \mathbf{R}^2 \to \mathbf{H}^2 \times \mathbf{H}^2$ be a (κ, C)-quasi-isometric embedding. Then there are finite constants $F = F(\kappa), N = N(\kappa, C)$ such that $\phi(\mathbf{R}^2)$ lies in the N-neighborhood of a union of F flats in $\mathbf{H}^2 \times \mathbf{H}^2$.*

In the paper [EF] we prove the analogous result for quasi-isometric embeddings $\mathbf{R}^n \to X$ into an arbitrary nonpositively curved rank n symmetric space X with no Euclidean direct factors.

The purpose of the present paper is to exhibit most of the ideas used in the proof of the general result, but in a context that is easier to understand and to visualize. We refer the reader to [EF] for an introduction and motivation for the study of quasi-flats and quasi-isometries in symmetric spaces, as well as further applications.

*Both authors are supported in part by N.S.F. Postdoctoral Fellowships. The work of the second author at MSRI was supported by NSF grant DMS-9022140.

1.1 Quasi-flats which are not close to flats

In this subsection we give an example, suggested to us by Bruce Kleiner, of a quasi-isometric embedding of $\mathbf{R}^2$ in $\mathbf{H}^2 \times \mathbf{H}^2$ which does not lie in a bounded neighborhood of a single flat in $\mathbf{H}^2 \times \mathbf{H}^2$. This type of example can be easily generalized to any symmetric space X with $\mathrm{rank}(X) \geq 2$.

Key Example: For each factor of $\mathbf{H}^2 \times \mathbf{H}^2$, pick three distinct geodesic rays in $\mathbf{H}^2$, call them A, B, C and X, Y, Z. By taking metric products, each pair of rays gives rise to an isometrically embedded copy of a Euclidean quadrant in $\mathbf{H}^2 \times \mathbf{H}^2$. We denote these quarter-flats by AX, XB, etc. Divide the Euclidean plane $\mathbf{R}^2$ into six congruent sectors meeting at the origin. Let ϕ be the map which maps each sector, in clockwise order, to the quarter-flats AX, XB, BY, YC, CZ, ZA by the quasi-isometry which maps a sector to a quadrant (see Figure 1).

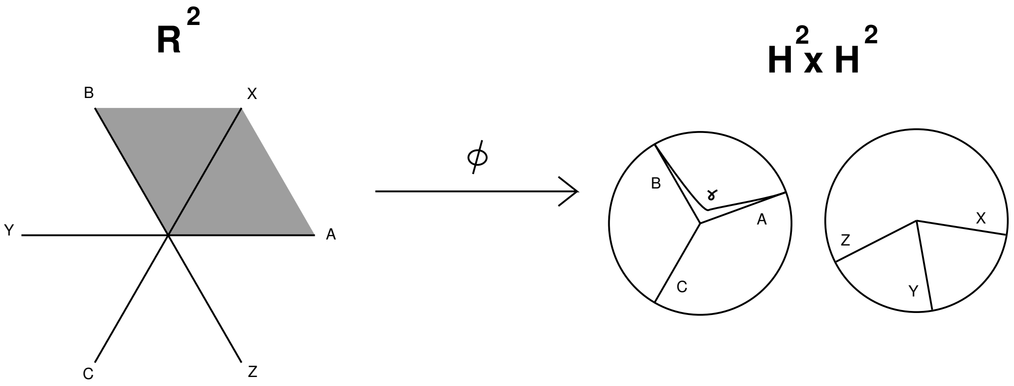

Figure 1. The map $\phi : \mathbf{R}^2 \to \mathbf{H}^2 \times \mathbf{H}^2$ maps each pair of adjacent sectors close to a half-flat: for example, the image of the shaded sectors under ϕ lies in a 4-neighborhood of the Euclidean half-flat $\gamma \times X$ in $\mathbf{H}^2 \times \mathbf{H}^2$, where γ is the geodesic connecting the endpoints of the rays A and B. It follows that ϕ is a quasi-isometry.

Note that the image of any adjacent pair of sectors lies uniformly close to a half-flat in $\mathbf{H}^2 \times \mathbf{H}^2$ since triangles in $\mathbf{H}^2$ are uniformly thin (see Figure 1). It follows that the restriction of ϕ to any pair of adjacent sectors is a 4-quasi-isometry, hence that ϕ is a K-quasi-isometry for some K. While $\phi(\mathbf{R}^2)$ lies in a bounded neighborhood of a union of three flats, it clearly does not lie in a bounded neighborhood of a single flat.

Remarks.

1. The above construction can be carried out using any number of geodesic rays instead of just three. In this case the constant K of the quasi-isometry ϕ will get larger as the number of geodesic rays used increases.

2. Note that we may precompose the map $\phi : \mathbf{R}^2 \to \mathbf{H}^2 \times \mathbf{H}^2$ with a self-quasi-isometry of $\mathbf{R}^2$, such as a quasi-isometry taking rays to logarithmic spirals. The point is that the six sectors may be geometrically quite complicated. For this reason we will not be able to use the geometry of $\mathbf{R}^2$ in the proof of Theorem 1.1

In what follows we always keep in mind the above examples.

1.2 Method of proof

In this section we give an outline of the proof of Theorem 1.1.

1.2.1 Nondegeneracy in X

In $X = \mathbf{H}^2 \times \mathbf{H}^2$ there is a "degenerate subset", which is simply the set of $x \in X$ which project to the origin in one of the factors. For a fixed, very small tolerance $\delta > 0$, we define the "nondegenerate subset" X^δ of X to be the set of $x \in X$ with $d(\pi_i(x), e) \geq \delta d(x, e)$ for $i = 1, 2$, where $\pi_i : X \to \mathbf{H}^2$ is the natural projection onto a factor and $e \in X$ is the origin.

1.2.2 Coarse topology and pinched subsets

It is a well-known fact that there is no (κ, C)-quasi-isometric embedding of a large, equilateral, Euclidean triangle into $\mathbf{H}^2$: if there were such a map, then the image would lie in a neighborhood of a geodesic triangle in $\mathbf{H}^2$ (since quasi-geodesics are close to geodesics), and triangles in $\mathbf{H}^2$ are uniformly thin (unlike equilateral triangles in $\mathbf{R}^2$).

Intuitively, any subset of $\mathbf{R}^2$ which admits a (κ, C)-quasi-isometric embedding into $\mathbf{H}^2$ should not contain any fat loops. We formalize this in the Main Lemma (Lemma 4.6): If W is any subset of $\mathbf{R}^2$ which admits a (κ, C)-quasi-isometric embedding to a hyperbolic space, then for a certain (small) neighborhood $\tilde{W}$ of W, the natural homomorphism $H_1(W) \to H_1(\tilde{W})$ induced by inclusion is the zero map. The size of the neighborhood $\tilde{W}$ of course depends on κ. We apply the Main Lemma to subsets of $\phi^{-1}(X^\delta)^c$.

The coarse topology we use is in the sprint of [FS].

1.2.3 Outline of the proof of Theorem 1.1

Associated to X is its *Furstenberg boundary* $\hat{X} = S^1 \times S^1$, which is the product of the boundaries at infinity of the two $\mathbf{H}^2$ factors. Pairs of points in $\hat{X}$ parameterize the set of (maximal) flats in X, just as pairs of points of the boundary at infinity of $\mathbf{H}^2$ parameterize the set of geodesics in $\mathbf{H}^2$. We identify $\hat{X}$ with the "visual sphere" $S^1 \times S^1$ at the origin of X[1].

We consider the set of points S_R in $\mathbf{R}^2$ lying in the annulus $A(R, 2R)$ which map into the nondegenerate subset X^δ, i.e. $S_R = A(R, 2R) \cap \phi^{-1}(X^\delta)$. We also consider the *visual image* of S_R, that is the set of points in the visual sphere $S1 \times S^1$ for which the infinite ray in that direction hits $\phi(S_R)$. In §3 we use basic geometry of X to show that the size of the visual image of $\phi(S_R)$ is small, in fact it can be covered by a bounded number (independent of R) of balls in $S^1 \times S^1$ of size roughly e^{-R}.

Proving that the visual image of ϕ actually *converges* to a finite set of points $(\alpha_1, \beta_1), \ldots, (\alpha_n, \beta_n)$ in $\hat{X}$ is a major step towards proving the theorem. Once this is known, we connect each α_i to each α_j by geodesics, and each β_i to each β_j by geodesics, and take all possible products of these geodesics to obtain a finite collection of flats. A final geometric argument, outlined in §6, proves that $\phi(\mathbf{R}^2)$ lies close to this finite union of flats.

The main difficulty is that the visual images of the sets $\phi(S_R)$ for different R have no *a priori* relationship, for example the visual image of $\phi(S_{R_j})$ may be a different finite collection of points for each $R_j = 2^j$. Hence there is no reason for the visual images of the sets $\phi(S_R)$ to converge to a fixed finite set of points in $\hat{X}$.

We rule out this phenomenon by showing that visual images *persist* (cf. Lemma 4.8 and Proposition 4.9): Suppose that, for some sufficiently large R_0, some point $x \in S_{R_0}$ maps into a "very nondegenerate subset" of X, i.e. $\phi(x) \in X^b$, $b \gg \delta$, and denote by $\beta \in S^1 \times S^1$ the visual position of $\phi(x)$. Then for *every* $R \geq R_0$ some point of $\phi(S_R)$ is seen in a direction near β. This fact together with the fact about the visual size of $\phi(S_R)$ noted above implies that there are finitely many limit points.

Proving persistence. A simple geometric observation (Lemma 3.3) shows that the visual size of any connected component of $\phi^{-1}(X^\delta)$ is exponentially small. Thus, to prove visual images persist, it is enough to prove (Lemma 4.8) that every connected component of $\phi^{-1}(X^\delta)$ which contains a "very nondegenerate point" x as above is unbounded.

The proof is simple in outline. First off, since x maps into a very nondegenerate

[1] Actually this is the product of the visual circles in each factor, as opposed to the visual sphere in X.

direction, the connected component of $\phi^{-1}(X^\delta)$ containing x contains some large ball around x. If the lemma were not true then this component would be bounded, so could be completely surrounded by a set $T \subset \mathbf{R}^2$, thought of as a circle S^1, which maps completely into one hyperbolic factor of X. By the coarse topology discussed above, this set T must be pinched, which contradicts the fact that it surrounds a large, round ball.

2 Preliminaries

2.1 Some definitions

Throughout the paper we let $X = \mathbf{H}^2 \times \mathbf{H}^2$. Both this and the rest of the notation are meant to match up with the notation in the general case of [EF], when X is any symmetric space. Let $0 \in \mathbf{R}^2$ and $e \in X$ be basepoints. Let $B(0,r)$ denote the ball of radius r centered at 0, and let $A(r_1, r_2)$ denote the annulus (centered at 0 unless otherwise specified) with inner radius r_1 and outer radius r_2. Let ρ be a non-negative real number, and let $d_\rho(x,0) = \max(d(x,0), \rho)$. So d_ρ is the usual distance on $\mathbf{R}^2$, except inside the ρ-ball all distances equal ρ.

For convenience we will assume that ϕ is continuous. It is possible to make this assumption by using the standard "connect the dots" argument, which changes ϕ into a continuous map while moving any image point by only a bounded amount. Connect-the-dots works roughly as follows: first triangulate the domain $\mathbf{R}^2$ by triangles of a uniformly bounded size; then build a quasi-isometry defined inductively on the skeleta of this triangulation, starting with the map being ϕ on the 0-skeleton, and extend to higher skeleta by connecting points (and then edges) by geodesics. This new quasi-isometry is continuous, and lies a bounded distance (in the sup-norm) from ϕ.

Let $K, \epsilon > 0$ be fixed. Unless otherwise stated, we will henceforth assume the following two conditions on the map ϕ, the point $0 \in \mathbf{R}^2$, and the point $e \in X$:

Condition I: ϕ is a "very coarse quasi-isometry" with additive coarseness constant depending on the distance from 0:

$$\frac{1}{K} d(x,y) - \epsilon d_\rho(x,0) \leq d(\phi(x), \phi(y)) \leq K d(x,y) + \epsilon d_\rho(x,0) \tag{1}$$

where for simplicity of notation we assume here and throughout that $d_\rho(x,0) > d_\rho(y,0)$.

Condition II: ϕ satisfies a radial condition:

$$\frac{1}{2K}d_\rho(x,0) \leq d(\phi(x),e) \leq 2K d_\rho(x,0) \tag{2}$$

Note that if ϕ is a (κ, C) quasi-isometry in the usual sense, then ϕ satisfies (1) with $\rho = C/\epsilon$ and $K = \kappa$. The constant ϵ can be made arbitrarily small by taking ρ sufficiently large. If $\phi(0) = e$ and ϵ is sufficiently small (depending on K), then (1) implies (2) outside some fixed ball.

The point of separating out the two conditions is that it will be necessary to consider the map ϕ with arbitrary points of $\mathbf{R}^2$ viewed as basepoints. The other reason to use Conditions I and II is to allow for maps ϕ that behave like quasi-isometries, but may remain undefined on many "holes" $H \subset \mathbf{R}^2$ of diameter roughly $d(H, 0)$. Such maps occur naturally when considering noncocompact lattices in higher rank, such as $SL(n, \mathbf{Z}), n \geq 3$. The higher-rank analog of Theorem 1.1, or more precisely its extension to quasi-flats with holes (Theorem 7.1), plays a major role in the proof of quasi-isometric rigidity for such lattices (see [E]).

In all that follows we allow the conditions (1) and (2) to be violated inside some fixed ball B centered at the origin, as long as all points and paths we are considering are outside B. However, we never mention this point explicitly until §6.

For any set $U \subset \mathbf{R}^2$, we let $U[c]$ be a neighborhood of U whose size is proportional to the distance from the origin, namely:

$$U[c] = \{x \in \mathbf{R}^2 \ : \ \exists u \in U \text{ with } d(x,u) < cd_\rho(x,0).\}$$

We use the same convention for subsets of X, with the origin replaced by the point e; namely, if $U \subset X$, then

$$U[c] = \{x \in X \ : \ \exists u \in U \text{ with } d(x,u) < cd(x,e).\}$$

Let $Y_1 = \mathbf{H}^2 \times \{e\}$ and $Y_2 = \{e\} \times \mathbf{H}^2$ be the direct factors of X, and denote the composition of ϕ with the natural projection $\pi_i : X \to Y_i$ by $\phi_i, i = 1, 2$. For any $\delta > 0$, we define

$$X^\delta = X \setminus (Y_1[\delta] \cup Y_2[\delta])$$

So the elements $x \in X^\delta$ are precisely those elements whose distance $d(0, x)$ is comparable to the distance to the origin of each of its projections to the factors Y_1, Y_2 (see Figure 2). More explicitly,

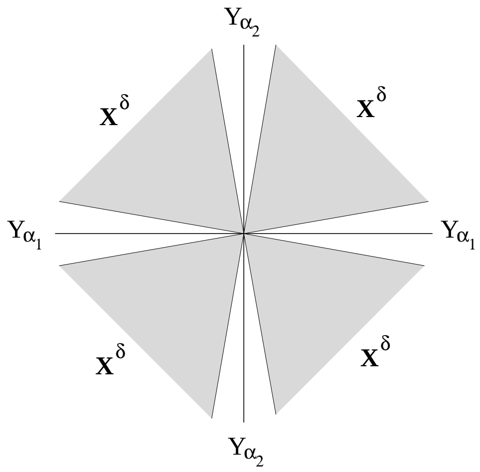

Figure 2. What the intersections of Y_i and X^δ with a flat through $e \in X$ look like in $\mathbf{H}^2 \times \mathbf{H}^2$.

$$X^\delta = \{x \in X : d(\pi_i(x), e) \geq \delta d(x, e) \text{ for each } i = 1, 2\}$$

Notational convention. All constants in this paper will depend on the quasi-isometry constant K. Hence constants introduced in statements of lemmas and theorems without further mention should thought of as numerical constants which could, in principal, be computed from K and basic constants depending on X. We refrain from such explicit computations of constants only for clarity of presentation.

If a constant η further depends on constants a and b, we will simply write $\eta = \eta(a, b)$. Unless otherwise specified, we denote by η_j "large" constants > 1, and by ν_j and δ_j "small constants", > 0 and < 1. We use the notation $O(b)$ to mean a quantity bounded above by λb for some implied constant $\lambda = \lambda(K, X)$. We also write $a \ll b$ to mean $a < \nu b$ for some small implied constant $\nu = \nu(K, X) < 1$, and $a \gg b$ to mean $a > \eta b$ for some large implied constant $\eta = \eta(K, X) > 1$. Finally we use the notations $a \prec b$ and $a \succ b$ to mean the same as $a \ll b$ and $a \gg b$ respectively, but without specifying the size of the implied constant.

We will always choose δ so that δ is much smaller then 1 but much larger then ϵ, i.e. $\epsilon \ll \delta \ll 1$, where ϵ is defined in Condition I.

2.2 Approximations

In this subsection we prove a few simple technical statements used throughout the paper.

2.2.1 Pullbacks and neighborhoods

Lemma 2.1. *Let $U \subset X$ be any subset. Then*

(a) If $\epsilon \ll \delta \ll 1$ then $\phi^{-1}(U)[\delta] \subset \phi^{-1}(U[4K^2\delta])$.

(b) $U[\alpha][\beta] \subset U[\alpha+\beta+\alpha\beta]$. If $\alpha < 1$ and $\beta < 1$, then $U[\alpha][\beta] \subset U[\alpha+2\beta]$.

(c) If α and β, $\epsilon \ll \alpha, \beta \ll 1$ then

$$\phi^{-1}\left(U[\alpha]\right)[\beta] \subset \phi^{-1}\left(U[\alpha+4K^2\beta+4K^2\beta\alpha]\right) \subset \phi^{-1}\left(U[\alpha+8K^2\beta]\right).$$

In particular, if $\epsilon \ll \delta \ll 1$ then for $i = 1, 2$

$$\phi^{-1}(Y_i)[\delta])[\epsilon] \subset \phi^{-1}(Y_i[2\delta]).$$

Proof. Part (a) is an easy consequence of (1), (2), and the triangle inequality. Part (b) follows immediately from the triangle inequality. Part (c) follows immediately from combining Parts (a) and (b). The last statement is the case $U = Y_i$. □

2.2.2 Approximating paths by edge-paths

For any $\epsilon > 0$, we divide $\mathbf{R}^2$ into a *grid* $\mathcal{G}$, that is a cellular decomposition of $\mathbf{R}^2$ into rectangular cells so that each point $x \in \mathbf{R}^2$ is contained in a cell each of whose edges has length between $\epsilon d_\rho(x, 0)/2$ and $2\epsilon d_\rho(x, 0)$. We emphasize that the size of the cells are not uniform: the length of an edge of a cell C is proportional to $d_\rho(C, 0)$.

It will be necessary for us to approximate paths $\phi(\gamma) \subset X$, γ an arbitrary path in $\mathbf{R}^2$, by paths $\hat{\gamma}$ which stay "close" to $\phi(\gamma)$, and so that there is some control on the length of $\hat{\gamma}$.

Lemma 2.2 (Approximating paths by edge-paths). *Let $R > 0$ be given, and let $R_j = 2^j R$ for $j = 1, 2 \ldots$. Let $\hat{A}_j$ denote the annulus in X centered at e with inner radius R_j and outer radius R_{j+1}. Then for any path $\gamma : [a, b] \to \mathbf{R}^2$, there is a path $\hat{\gamma} : [a, b] \to X$ connecting $(\phi \circ \gamma)(a)$ and $(\phi \circ \gamma)(b)$ such that:*

1. *$\hat{\gamma} \subset (\phi \circ \gamma)[\lambda\epsilon]$, for some constant λ.*

2. $\ell(\hat{\gamma} \cap \hat{A}_j)) \prec \frac{1}{\epsilon} R_j$ *for every* $j = 1, 2, \ldots$.

Proof. Let $\mathcal{C}$ denote the union of cells C with $C \cap \gamma \neq \emptyset$. Then $\mathcal{C}$ is connected subcomplex of the grid $\mathcal{G}$. Let a' be the closest grid vertex to $\gamma(a)$, b' the closest grid vertex to $\gamma(b)$. We define $\tilde{\gamma}^{(\epsilon)}$ to be the shortest edge-path in the 1-skeleton of $\mathcal{C}$ connecting a' to b', and let $\gamma^{(\epsilon)} = \overline{\gamma(a)a'} \cup \tilde{\gamma}^{(\epsilon)} \cup \overline{b'\gamma(b)}$, where $\overline{\gamma(a)a'}$ and $\overline{b'\gamma(b)}$ are segments in $\mathbf{R}^2$. Then $\hat{\gamma}$ is the path in X connecting $\phi \circ \gamma(a)$ and $\phi \circ \gamma(b)$ obtained by connecting the images under ϕ of the vertices of $\gamma^{(\epsilon)}$ by geodesics.

By definition of $\gamma[2\epsilon]$, and since every point of $\gamma^{(\epsilon)}$ lies in a cell intersecting γ and every cell C has edges of length at most $2\epsilon d_\rho(C, 0)$, we have $\gamma^{(\epsilon)} \subset \gamma[2\epsilon]$.

Now for any $z \in \gamma^{(\epsilon)}$, $d(z, \gamma(t)) \leq 2\epsilon d_\rho(z, 0)$ for some t. Hence

$$d(\phi(z), \phi(\gamma(t))) \leq (2K+1)\epsilon d_\rho(z, 0) \leq (4K^2 + 2K)\epsilon d(\phi(z), 0)$$

the second inequality following from condition II. Hence $\phi(\gamma^{(\epsilon)}) \subset \phi(\gamma)[(4K^2+2K)\epsilon]$. By construction of $\hat{\gamma}$ we now know $\hat{\gamma} \subset \phi(\gamma^{(\epsilon)})[(2K+1)\epsilon]$. Therefore by Lemma 2.1 part (b), $\hat{\gamma} \subset \phi(\gamma)[\lambda\epsilon]$, with e.g. $\lambda = 12K^3$.

Now since $\gamma^{(\epsilon)}$ is shortest, it traverses any edge in $\mathcal{C}$ at most once. Hence

$$\begin{aligned} \ell(\gamma^{(\epsilon)} \cap A(R_{j-1}, R_j)) &\prec (\text{\# of edges in a cell}) \cdot (\text{max. length of edge}) \cdot (\text{\# of cells}) \\ &\prec 2\epsilon R_j \cdot \frac{(R_{j+1}^2 - R_j^2)}{(\epsilon R_j/2)^2} \\ &\prec \tfrac{R_j}{\epsilon} \end{aligned}$$

By Condition II, $\phi^{-1}(\hat{A}_j)$ is contained in a bounded number of annuli $A(R_{k-1}, R_k)$. Thus property 2 holds, by the definition of $\hat{\gamma}$ and by Condition I. Note that $\gamma^{(\epsilon)}$ and $\hat{\gamma}$ consist of geodesic segments, so that the inequality on the arc lengths follows from Condition I. $\square$

3 Visual Angles

3.1 A fact about $\mathbf{H}^2$

For $x, y \in \mathbf{H}^2$, let $\Theta_e(x, y)$ denote the visual angle between the points x and y when viewed from e, where e denotes the origin in $\mathbf{H}^2$. We begin with an elementary but crucial lemma about the hyperbolic plane.

Lemma 3.1 (visual angle vs. hyperbolic distance in $\mathbf{H}^2$). *Let $x, y \in \mathbf{H}^2$ with $d(x,e) = d(y,e) = r$. If r is sufficiently large then :*

(a) Any path from x to y staying outside the r-ball has length at least $\Theta_e(x,y)e^r/2$.

(b) For any constant $0 < \alpha < 1$, if $d(x,y) \geq 2\alpha r$ then $\Theta_e(x,y) \geq e^{-(1-\alpha)r}$. (Taking $\alpha = 1/2$ gives $\Theta_e(x,y) \geq e^{-r/2}$ if $d(x,y) \geq r$).

(c) For any constant $0 < \alpha < 1$, if $\Theta_e(x,y) \geq e^{-\alpha r}$ then $d(x,y) \geq 2(1-\alpha)r$. (Taking $\alpha = 1/2$ gives $d(x,y) \geq r$ if $\Theta_e(x,y) \geq e^{-r/2}$).

Proof. Now, the length of an arc on a circle of radius r is $\theta \sinh r$, where θ is the angle of the arc; part (a) follows by projecting the given path to the circle of radius r about the origin.

Parts (b) and (c) follows easily from the hyperbolic law of cosines. □

3.2 Bounding the visual size of $\phi(\mathbf{R}^2)$.

Let $\Theta(y) \in S^1$ denote the point such that $\exp(r\Theta(y)) = y$ for some $r > 0$, where $\exp : T\mathbf{H}^2_e \to \mathbf{H}^2$ denotes the exponential map.

The following lemma says that the part of the ϕ-image of big annular regions in $\mathbf{R}^2$ which map into nondegenerate directions have small visual size when viewed from the origin in X.

Lemma 3.2 (bounded number of visual angles). *There exists a constant $\nu_1 = \nu_1(\delta)$ so that for each $i = 1, 2$, the set $\Theta(\phi_i(A(R, 2R)) \cap X^\delta)$ can be covered by $c_0 = c_0(\delta)$ balls of radius $e^{-\nu_1 R}$ as long as R is sufficiently large (depending on δ, ρ).*

In the Key Example (§1.1), for any $R > 0$ the set $\Theta(\phi_1(A(R, 2R) \cap X^\delta)$ consists of three points, corresponding to the directions of the three geodesic rays A,B, and C.

Proof of Lemma 3.2. For any $R > 0$, we decompose the unit circle $S^1 \subset T\mathbf{H}^2_e$ into intervals of length at least $e^{-\nu_1 R}$ and at most $e^{-\nu_1 R/4}$. We number the these intervals as $I_p, p = 1, 2, \ldots$.

Suppose $\Theta(\phi_i(x)) \in I_p$ and $\Theta(\phi_i(y)) \in I_{p+2m}$ for m a nonzero integer. Then if $\phi(x)$ and $\phi(y)$ belong to X^δ,

$$d(\phi_i(x), e)) > \delta d(\phi(x), e) > \frac{\delta R}{2K}$$

Similarly, $d(\phi_i(y), e) > \delta R/(2K)$. Since $\Theta_e(\phi_i(x), \phi_i(y)) \geq e^{-\nu_1 R}$ by our choice of x and y, part (c) of Lemma 3.1 gives

$$d(\phi_i(x), \phi_i(y)) \succ \delta R$$

for R sufficiently large. Then by Condition I, $d(x, y) > \lambda_2 \delta R/K$ for some constant λ_2.

From this it follows that the number of distinct m for which I_{p+2m} intersects $\Theta(\phi_i(X^\delta) \cap A(R, 2R))$ is bounded by $(2R)^2/(\lambda_2 \delta R/2K)^2 = O(1/\delta^2)$. Hence the the number of odd p (and even p) for which $I_p \cap \Theta(\phi_i(X^\delta) \cap A(R, 2R))$ is non-empty is $O(1/\delta^2)$. □

We emphasize that the size of the balls in Lemma 3.2 decreases exponentially with R, while the number of balls is bounded independent of R. However, Lemma 3.2 gives no information about the relative positions of the balls at different radii R.

The following lemma says that the visual size of the image of any connected component of $\phi^{-1}(X^\delta)$ is small. In the lemma, and in the rest of the paper, connected means path-connected.

Lemma 3.3 (No Shifting Lemma). *For every δ, $\epsilon \ll \delta \ll 1$, there exists $\nu_2 = \nu_2(\delta)$ so that if x, y are any two points in the same connected component of $\phi^{-1}(X^\delta) \cap B(0, r)^c$, then*

$$\Theta_e(\phi_i(x), \phi_i(y)) \leq e^{-\nu_2 \max(r, \rho)}$$

for each $i = 1, 2$ as long as $\max(r, \rho)$ is sufficiently large (depending on ϵ, δ).

In the Key Example (§1.1), a connected component of $\phi^{-1}(X^\delta)$ consists of a Euclidean sector S of angle $< \pi/3$. The set $\Theta(\phi_i(S)) \subset S^1$ is a single point, so in this case Lemma 3.3 is obviously true.

Proof of Lemma 3.3. Let $R = \max(r, \rho)/(2K)$. By Condition II, if $d(x, 0) \geq r$ and $d(y, 0) \geq r$ then $d(\phi(x), e) \geq R$ and $d(\phi(y), e) \geq R$.

Let $\gamma \subset \phi^{-1}(X^\delta)$ be a path such that $\gamma(0) = x$ and $\gamma(1) = y$. Let $\hat{\gamma}$ be the path constructed from $\phi \circ \gamma$ as in Lemma 2.2. Since $\phi \circ \gamma \subset X^\delta$, Property 1 of Lemma 2.2 gives $\hat{\gamma} \subset X^{\delta/2}$, as long as $\epsilon \ll \delta$. Let $\hat{A}_j$ denote the annulus $A(2^j R, 2^{j+1} R)$ in X. By Property 2 of Lemma 2.2 we have

$$\ell(\hat{\gamma} \cap \hat{A}_j) \prec 2^j R/\epsilon, \tag{3}$$

If $\phi(x) \in \hat{A}_j$ then $d(\phi(x), e) \geq 2^j R$. So if $\phi(x) \in \hat{A}_j \cap X^{\delta/2}$ then $d(\phi_i(x), e) \geq \delta 2^{j-1} R$ for each i. Applying part (a) of Lemma 3.1 to the path $\pi_i(\hat{\gamma})$ and with $r = \delta 2^{j-1} R$ gives

$$\begin{aligned} |\Theta(\pi_i(\hat{\gamma} \cap \hat{A}_j))| &\leq 2\ell(\pi_i(\hat{\gamma} \cap \hat{A}_j))e^{-\delta 2^{j-1} R} && \text{by Lemma 3.1} \\ &\prec (2^{j+1} R/\epsilon)e^{-\delta 2^{j-1} R} && \text{by (3)} \end{aligned}$$

Summing over j then gives

$$|\Theta(\pi_i(\hat{\gamma}))| \prec (R/\epsilon) \sum_{j=0}^{\infty} 2^{j+1} e^{-\delta 2^{j+1} R} \tag{4}$$

For sufficiently large r (and hence R), (4) is bounded by $e^{-\lambda \delta r}$ for some $\lambda > 0$, hence we can let $\nu_2 = \lambda\delta$. $\square$

4 Limit Sets

Let $\hat{X} = S^1 \times S^1$ be the product of the boundaries at infinity of the hyperbolic planes Y_1 and Y_2 in X. We call $\hat{X}$ the *boundary* of X.

4.1 Definition and finiteness

Basic to our study is the definition of limit set of ϕ: the collection of points in the boundary $\hat{X}$ on which $\phi(\mathbf{R}^2)$ limits in a strong sense.

Definition 4.1 (Limit Point). A point $\beta = (\beta_1, \beta_2) \in \hat{X}$ is called a *δ-limit point* of ϕ if there exists a path γ in $\phi^{-1}(X^\delta)$ such that $\gamma(t)$ leaves every compact set and the limit $\lim_{t\to\infty} \Theta(\phi_i(\gamma(t)))$ exists for each i and equals β_i. The collection of all δ-limit points of ϕ is called the *δ-limit set* of ϕ.

We will eventually see that, for δ sufficiently small, the limit set is independent of δ.

In the Key Example (§1.1), the limit set of ϕ consists of six pairs of points. For example, the point $(a, x) \in \hat{X}$, where a (resp. x) is the endpoint at infinity of the ray A (resp. X) in $\mathbf{H}^2$, is a limit point of ϕ.

Proposition 4.2 (The limit set is finite). *The cardinality of the δ-limit set of ϕ is at most c_0 (defined in Lemma 3.2).*

Proof of Proposition 4.2. Suppose there exist c_0+1 limit points $\beta_1 = (\beta_1^1, \beta_1^2), \ldots, \beta_{c_0+1} = (\beta_{c_0+1}^1, \beta_{c_0+1}^2)$. Since the β_j are limit points there exist paths $\gamma_j : [0, \infty) \to \phi^{-1}(X^\delta)$ which leave every compact set and $\lim_{t\to\infty} \Theta(\phi_i(\gamma_j(t))) = \beta_j^i$ for $i = 1, 2$. Choose r large so that Lemma 3.3 holds, so that for each $i = 1, 2$ the intervals of radius $2e^{-\nu_1 r} + 2e^{-\nu_2 r}$ centered at the limit points $\beta_1^i, \ldots, \beta_{c_0+1}^i$ in S^1 are disjoint, and so that $r > d(\gamma_j(0), 0)$ for all j. Here ν_1 and ν_2 are the same as in Lemma 3.2 and Lemma 3.3. For each j let t_j denote the greatest t such that $d(\gamma_j(t_j), 0) = r$, and let $x_j = \gamma_j(t_j)$. Then by Lemma 3.3, $\Theta_e(\phi_i(x_j)), \beta_j^i) \leq e^{-\nu_2 r}$ for $i = 1, 2$. Hence, by the triangle inequality, for $p \neq q$, $\Theta_e(\phi_i(x_p), \phi_i(x_q)) \geq 2e^{-\nu_1 r}$ for $i = 1, 2$. Since there are $c_0 + 1$ points x_j, this contradicts Lemma 3.2. □

4.2 Independence of basepoint

In this subsection show that the δ-limit set of ϕ is independent of various choices of parameters.

For any point $e' \in X$, we define $X^\delta(e')$ the same as we defined X^δ, except now e' is used as basepoint. More explicitly

$$X^\delta(e') = \{x \in X : d(\pi_i(x), e') \geq \delta d(x, e')) \text{ for each } i = 1, 2\}$$

Lemma 4.3. *Suppose $x \in \phi^{-1}(X^\delta(e))$, $d(x, 0) \geq R$, and $d(e, e') \leq r$. Then $x \in \phi^{-1}(X^{\delta/2}(e'))$ if $R \gg r/\delta$.*

Proof. This follows from Condition II and the triangle inequality. □

The unit tangent circle $T\mathbf{H}_z^2$ at a point $z \in \mathbf{H}^2$ has a natural identification $h_z : T\mathbf{H}_z^2 \to S_\infty^1$, where a vector $\vec{v} \in T\mathbf{H}_z^2$ maps to the endpoint of the unique geodesic ray through z in the $\vec{v}$ direction. The map h_z is a homeomorphism, indeed a conformal map. In this way the unit tangent circles at any two points $z, y \in \mathbf{H}^2$ may be identified in a natural way by the conformal homeomorphism $h_y^{-1} \circ h_z$.

Lemma 4.4 (limit set is independent of basepoint). *Suppose the map ϕ satisfies Condition I and Condition II with $e \in X$ replaced by some other point e' and $0 \in \mathbf{R}^2$ replaced by $0'$. Then the δ-limit set of ϕ defined at any e' is contained the $\delta/2$-limit set of ϕ defined at the origin e.*

Proof. Let $e' \in X$ be any point, let $r_0 = d(e', e)$ and let $\mathcal{L}'_\delta(\phi)$ denote the δ-limit set defined using e' as basepoint. By Lemma 4.3, if $\beta = (\beta_1, \beta_2) \in \mathcal{L}_\delta(\phi)$ and γ is the path in $\phi^{-1}(X^\delta)$ so that $\lim_{t\to\infty} \Theta(\phi_i(\gamma(t))) = \beta_i$, then the part of γ lying outside a ball centered at e with radius $\gg r/\delta$ is inside $X^{\delta/2}(e')$. Also because hyperbolic triangles are thin, $\lim_{t\to\infty} \Theta(g\gamma(t)) = g \cdot \beta$, where g is an isometry taking e to e'. Hence $\beta \in \mathcal{L}'_{\delta/2}(\phi)$ (recall that the isometry g^{-1} sending e' to e also sends β to $g^{-1} \cdot \beta$). $\square$

We also note the following:

Lemma 4.5 (invariance under linear growth perturbations). *Suppose ϕ and ψ are two functions satisfying Condition I and Condition II, and $d(\phi(x), \psi(x)) \leq \epsilon d_\rho(x, 0)$ for all $x \in X$. Then for $\epsilon \ll \delta \ll 1$, the δ-limit set of ϕ is contained in the $(\delta/2)$-limit set of ψ and vice versa.*

Proof. From part (c) of Lemma 2.1 $\phi^{-1}(X^\delta) \subset \psi^{-1}(X^{\delta/2})$. Also by Lemma 3.1 part (c), $\Theta_e(\phi_i(x), \psi_i(x)) \to 0$ as $x \to \infty$ in $\phi^{-1}(X^\delta)$ for $i = 1, 2$. Thus any δ limit point of ϕ is also a $\delta/2$ limit point of ψ. $\square$

4.3 Existence of limit points

Our proof that limit points exist follows from the main lemma of this paper, whose proof we will leave for a later section. The idea is that the set of points in $\mathbf{R}^2$ which map (nearly) into one of the $\mathbf{H}^2$ factors of X must be a pinched subset of $\mathbf{R}^2$. As stated in the introduction, our notion of W being pinched is the inclusion of W into some neighborhood inducing the zero map on first homology. Here we take the neighborhood of W to be $W[\eta\delta]$, for some constant η. Since we wish to ignore effects inside a large (but fixed) ball, we union the sets in question with such a ball.

Lemma 4.6 (Main Lemma). *Suppose $\epsilon \ll \delta \ll 1$ and $\max(r, \rho)$ is sufficiently large (depending on ϵ, δ). Then for any subset W of $\phi^{-1}(X^\delta)^c$, the homomorphism $i_* : H_1(W \cup B(0, r)) \to H_1(W[\eta\delta] \cup B(0, r))$ induced by the inclusion is the zero map, for some $\eta > 1$ independent of δ.*

In the Key Example (§1.1), the set $\phi^{-1}(X^\delta)^c$ consists of a neighborhood $L[\lambda]$ (λ some constant depending on δ) of the union L of the six rays emanating from the origin in $\mathbf{R}^2$ (see Figure 2); in fact $L = \phi^{-1}(Y_1 \cup Y_2)$. Hence in this special example Lemma 4.6 is clearly true.

In the rest of this subsection we present consequences of the main lemma.

Lemma 4.7 (Non-Degenerate Points exist). *If $\epsilon \ll \delta \ll 1$, and if R is sufficiently large (depending on ϵ, δ), then the sphere of radius R around 0 contains a point in $\phi^{-1}(X^\delta)$.*

Proof. Suppose r is so large that Lemma 4.6 holds. Let $B = B(0,r)$ and let W denote the circle of radius R. If $R > r$, the sets W and B do not intersect. Also $W[\eta\delta] \subset A((1+\eta\delta)^{-1}R, (1-\eta\delta)^{-1}R)$. Thus if $R \gg r$, the sets $W[\eta\delta]$ and B do not intersect. If the lemma is false, then $W \subset \phi^{-1}(X^\delta)^c$. By Lemma 4.6, the inclusion induced homomorphism $H_1(W \cup B) \to H_1(W[\eta\delta] \cup B)$ is the zero map. Since $W[\eta\delta] \cap B = \emptyset$, this implies that $H_1(W) \to H_1(W[\eta\delta])$ is the zero map, and so $H_1(W) \to H_1(A((1+\eta\delta)^{-1}R, (1-\eta\delta)^{-1}R))$ is also the zero map. This is clearly a contradiction since the inclusion $W \to A((1+\eta\delta)^{-1}R, (1-\eta\delta)^{-1}R)$ is that of a circle into an annulus. $\square$

The following lemma says that you can see forever after you see a deep point.

Lemma 4.8 (Connected components of $\phi^{-1}(X^b)$ are unbounded). *Suppose $\epsilon \ll b \ll 1$, and z is a point of $\phi^{-1}(X^b)$ with $d_\rho(z,0)$ sufficiently large (depending on b, ϵ). Then the connected component of $\phi^{-1}(X^{b/\eta_1}) \cap B(0, d(z,0))^c$ containing z is unbounded, for some $\eta_1 \gg 1$ independent of b.*

Proof. All homologies and cohomologies in this proof are reduced with real coefficients. If $\epsilon \ll \delta \ll 1$, by Lemma 4.7 $\phi^{-1}(X^b)$ is unbounded. Pick δ so that $\epsilon \ll \delta$ and $\eta\delta < b$ (here η is from Lemma 4.6). Let B denote the open ball $B(0, d(z,0))$. Suppose the lemma is false, so that the component of $\phi^{-1}(X^\delta) \cap B^c$ containing z is bounded. Choose R so big that $B(0, R-1)$ contains this bounded component. Let $W_\delta = \phi^{-1}(\text{interior}(X^\delta)) \cap B(0, R+1)$ and let $W_b = \phi^{-1}(\text{interior}(X^b)) \cap B(0,R)$. Since $W_b \subset W_\delta$, $W_b \cap B^c$ has at least two connected components (one containing z and at least one intersecting the sphere of radius $R-1$). Hence $H_0(W_b \cap B^c) \neq 0$. Furthermore the image of the map $j_* : H_0(W_b \cap B^c) \to H_0(W_\delta \cap B^c)$ induced by the inclusion is not zero. Note that $W_\delta^c = \bigcup_{i=1,2} \phi^{-1}(\overline{Y_i[\delta]}) \cup B(0, R+1)^c$. Therefore, by part (a) of Lemma 2.1, $W_\delta^c[\eta\delta] \subset W_b^c$, since $\delta = b/\eta_1$ and $\eta_1 \gg 1$.

The universal coefficient theorem then implies that the corresponding map on cohomology $j^* : H^0(W_\delta \cap B^c) \to H^0(W_b \cap B^c)$ is not zero. Since W_δ and W_b are open, and $W_\delta^c \cup B$ is orientable, we have the following commutative diagram:

$$\begin{array}{ccc} H_1(W_\delta^c \cup B) & \longrightarrow & H_1(W_b^c \cup B) \\ \cong\big\downarrow & & \cong\big\downarrow \\ H^0(W_\delta \cap B^c) & \longrightarrow & H^0(W_b \cap B^c) \end{array}$$

The vertical maps are Alexander duality isomorphisms. Hence the top horizontal map is non-zero. But by the Main Lemma (Lemma 4.6) the top horizontal map is zero. This is a contradiction. $\square$

Proposition 4.9 (Limit Points Exist). *Suppose $\epsilon \ll b \ll 1$, and z is a point of $\phi^{-1}(X^b)$ with $d_\rho(z,0)$ sufficiently large (depending on b, ϵ). Then there exists a b/η_1-limit point $\beta = (\beta_1, \beta_2)$ in $\hat{X}$ such that*

$$\Theta_e(\beta_i, \phi_i(z)) < e^{-\nu_2 d_\rho(z,0)} \quad \textit{for } i = 1, 2$$

The constant η_1 is the same as in Lemma 4.8 and depends only on K; the constant ν_2 depends on K and δ.

Proof. Let $s = d_\rho(z,0)$. By Lemma 4.8, there exists a path $\gamma : [0,\infty) \to \phi^{-1}(X^{b/\eta_1}) \subset \mathbf{R}^2$ such that $\gamma(0) = z$, $d(\gamma(t),0) \geq s$ for all t and such γ leaves every ball centered at 0. Let $r = \max(\rho, s)$. For a non-negative integer j let t_j denote the greatest value of t for which $d(\gamma(t_j),0) = 2^j r$. Then by Lemma 3.3, for all $t \geq t_j$, $\Theta_e(\phi_i \circ \gamma(t), \phi_i \circ \gamma(t_j)) \leq e^{-\nu_2 2^j r}$ for $i = 1, 2$. This implies that the limit as $t \to \infty$ of $\Theta(\phi_i \circ \gamma(t))$ exists for each i: call it β_i. The above estimate for $j = 0$ shows that $\Theta_e(\phi(z), \beta_i) \leq e^{-\nu_2 r}$. $\square$

5 Proof of the main lemma (Lemma 4.6)

5.1 The expanding annulus lemma

The following lemma is another statement of what it means for movement in $\mathbf{H}^2$ transverse to the radial direction to have "exponential cost".

Lemma 5.1 (Expanding annulus lemma in $\mathbf{H}^2$). *Let $c_1 \ll 1$ be given, and let ψ be a path in $\mathbf{H}^2$ connecting points a_0 and b_0. Suppose that:*

1. *$a_0, b_0 \in A(r, (1 + c_1/4)r)$.*
2. *$d(\psi(t), e) \geq r$ for all t.*
3. *$\ell(\psi \cap A(2^i r, 2^{i+1} r)) \prec 2^i r$.*

Then $d(a_0, b_0) < c_1 r$ if r is sufficiently large (depending on c_1).

Proof. Let a_0' (resp. b_0') denote the radial projection of a_0 (resp. b_0) on a circle of radius r. Suppose $d(a_0, b_0) \geq c_1 r$. Then $d(a_0', b_0') \geq (c_1/2)r$. Hence by part (b) of Lemma 3.1:

$$|\Theta(\psi)| \geq \Theta_e(a_0, b_0) = \Theta_e(a_0', b_0') \succ e^{-r} \tag{5}$$

Part (a) of Lemma 3.1 gives

$$|\Theta(\psi \cap \hat{A}_i)| \prec \ell(\psi \cap A(2^i r, 2^{i+1} r)) e^{-2^j r} \prec (2^j r) e^{-2^j r}$$

Summing over j gives

$$|\Theta(\psi)| \prec r \sum_{j=0}^{\infty} 2^j e^{-2^j r}$$

which contradicts (5) if r is large enough. $\square$

5.2 Coarse topology and pinched subsets of $\mathbf{R}^2$

The following definition is motivated by Lemma 5.1.

Definition 5.2 (Pinching function). Let W be a subset of $\mathbf{R}^2$. A non-negative, proper, continuous $f : W \to \mathbf{R}$ is a *pinching function* on W, with *pinching constants* (r_0, η, β) if for all $s \geq r_0$ and all $x, y \in W$ satisfying $s \leq f(x) \leq f(y) \leq \eta s$ the following holds: if there exists a path $\gamma : [0,1] \to W$ connecting x and y with $f(\gamma(t)) \geq s$ for all t, then $d(x,y) < \beta s$.

Note. If f is a pinching function for W then f is a pinching function for any subset of W, with the same pinching constants.

The goal of this subsection is to show that subsets $W \subset \mathbf{R}^2$ with pinching functions are *pinched*, by which we mean the homology $H_q(W), q \geq 1$ vanishes under the inclusion map into some neighborhood of W.

We begin with a fact about $\mathbf{R}^2$.

Lemma 5.3 (Trivial $\mathbf{R}^2$ fact). *If $r > 0$ and if $T \subset \mathbf{R}^2$ is a set of diameter at most r, then for any function $r(x)$ on T satisfying $2r < r(x) < 4r$ for all $x \in T$, the set $\bigcup_{x \in T} B(x, r(x))$ is contractible.*

Proof. Let u be any point of T. Then $T \subset B(u, r)$. Hence for each $x \in T$, the ball $B(x, r(x))$ contains u. Thus $\bigcup_{x \in T} B(x, r(x))$ is star-shaped with respect to u, hence contractible. $\square$

Lemma 5.4 (Pinching lemma). *Suppose $W \subset \mathbf{R}^2$, and suppose f is pinching function on W with pinching constants (r_0, η, β). Let $\tilde{F}$ be any ball containing $f^{-1}([0, r_0])$, and let $\tilde{W}$ be the following the neighborhood of W:*

$$\tilde{W} = \tilde{F} \cup \{x \in \mathbf{R}^2 : \exists w \in W \, with \; d(x, w) < 4\beta f(w)\}$$

Then the homomorphism $i_* : H_1(W) \to H_1(\tilde{W})$ induced by the inclusion $i : W \to \tilde{W}$ is the zero homomorphism.

Proof. Let $F = W \cap \tilde{F}$, so $\tilde{F}$ is a neighborhood of F and $\tilde{W}$ is a neighborhood of W. From the commutative diagram

$$\begin{array}{ccccc} H_1(F) & \longrightarrow & H_1(W) & \longrightarrow & H_1(W, F) \\ \downarrow & & \downarrow & & \downarrow \\ H_1(\tilde{F}) & \longrightarrow & H_1(\tilde{W}) & \longrightarrow & H_1(\tilde{W}, \tilde{F}) \end{array} \tag{6}$$

and the contractibility of $\tilde{F}$, it is enough to show that the right vertical homomorphism in (6) is the zero map.

In the proof of this lemma, for any subset Q of W, we denote by $\tilde{Q}$ the union of the set $\tilde{F}$ with the set of points $x \in \mathbf{R}^2$ such that there exists $w \in Q$ with $d(x, w) < 4\beta f(w)$.

Clearly, without loss of generality we may assume that W is bounded. (Indeed, if c is a cycle such that $i_*(c) \neq 0$, we may replace W by a suitable metric neighborhood of c.) Pick a decreasing sequence $s_0, \ldots, s_N$ so that $s_0 = \sup_{w \in W} f(w)$, $s_{k-1}/s_k < \eta^{1/2}$ for all k, and $s_N = r_0$. Let $W_k = \{w \in W \;:\; f(w) \leq s_k\} \cup F$, so that W_k is a decreasing family of sets with $W_0 = W$, and $W_N = F$.

We claim that for each k, and any connected component Q of $W \setminus W_{k+1}$,

$$j_* : H_1(Q, Q \cap W_k) \to H_1(\tilde{Q}, \widetilde{Q \cap W_k}) \quad \text{is the zero map} \qquad (*)_k$$

where j_* is induced by the inclusion $j : Q \to \tilde{Q}$. The point of considering the connected components as opposed to all of $W \setminus W_k$ at once is that this allows us to get the additional property of connectedness, which will be crucial.

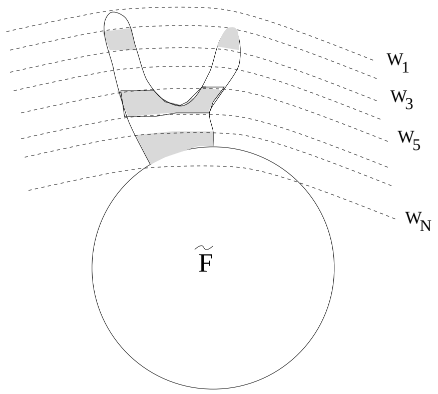

Figure 3. Applying the definition of pinching function shows that each of the regions like the shaded regions must have small diameter. Hence $W = \cup_{i=1}^{N} W_i$ must be pinched.

We prove $(*)_k$ by induction on k. If $k = 0$ then $H_1(Q, W_0) = 0$ because $Q \subset W = W_0$. Now suppose $(*)_{k-1}$ is true. Consider the commutative diagram:

$$\begin{tikzcd}
H_1(Q \cap W_{k-1}, Q \cap W_k) \arrow[r] \arrow[d] & H_1(Q, Q \cap W_k) \arrow[r] \arrow[d, "j_*"'] & H_1(Q, Q \cap W_{k-1}) \arrow[d, "j'_*"'] \\
H_1(\widetilde{Q \cap W_{k-1}}, \widetilde{Q \cap W_k}) \arrow[r] & H_1(\tilde{Q}, \widetilde{Q \cap W_k}) \arrow[r, "\Phi_*"] & H_1(\tilde{Q}, \widetilde{Q \cap W_{k-1}})
\end{tikzcd} \tag{7}$$

where the top row is the exact sequence of the triple $Q \cap W_k \subset Q \cap W_{k-1} \subset Q$, the bottom row is the exact sequence of the triple $\widetilde{Q \cap W_k} \subset \widetilde{Q \cap W_{k-1}} \subset \tilde{Q}$, and the vertical maps are induced by the inclusion $Q \to \tilde{Q}$.

The fact that j'_* in (7) is zero does not formally follow from the inductive hypothesis $(*)_{k-1}$ because Q is a connected component of $W \setminus W_{k+1}$, not $W \setminus W_k$. However, we may argue as follows: since $W_k \subset W_{k-1}$ for all k,

$$H_1(Q, Q \cap W_{k-1}) = H_1(Q \setminus W_k, (Q \setminus W_k) \cap W_{k-1}) = \oplus_{i \in I} H_1(S_i, S_i \cap W_{k-1})$$

where the S_i are the connected components of $Q \setminus W_k$. The inductive hypothesis $(*)_{k-1}$ is the statement that for each i, the homomorphism $j_* : H_1(S_i, S_i \cap W_{k-1}) \to H_1(\widetilde{S}_i, \widetilde{S_i \cap W_{k-1}})$ induced by the inclusion $S_i \to \tilde{S}_i$ is the zero map. Since for each i the inclusion $S_i \to \tilde{Q}$ factors through the inclusion $S_i \to \tilde{S}_i$, this implies that the map j'_* in (7) is the zero map.

Denote $(Q \cap W_{k-1}) \setminus (Q \cap W_k)$ by T. Now suppose x and y are any two points in T. Then $s_k \leq f(x) \leq s_{k-1}$, and $s_k \leq f(y) \leq s_{k-1}$. By the definition of T, and since Q is *connected*, x and y can be connected by a path $\gamma : [0,1] \to W$ such that $f(\gamma(t)) \geq s_{k+1}$ for all t. Note that $s_{k-1}/s_{k+1} < \eta$. Since f is a pinching function, this implies that $d(x,y) < \beta s_{k+1}$. Now Lemma 5.3 implies that $\tilde{T}$ is contractible. Hence $H_1(\widetilde{Q \cap W_{k-1}}, \widetilde{Q \cap W_k}) = 0$, so the map Φ_* in (7) is injective. This, together with the fact that $j'_* = 0$, implies that the map j_* in (7) is the zero map, which proves $(*)_k$. Since $W_N = F$, this proves that the right vertical homomorphism in (6) is indeed zero. □

5.3 Constructing a pinching function on $\phi^{-1}(Y_i[\delta])$

5.3.1 An approximation

Lemma 5.5 (ϕ_i satisfies Conditions I and II). *Let $\epsilon \ll \delta \ll 1$. For each $i = 1, 2$, the following two facts hold for any $x, y \in \phi^{-1}(Y_i[\delta])$:*

1.

$$d(x,y) - O(\delta d_\rho(x,0)) \prec d(\phi_i(x), \phi_i(y)) \prec d(x,y) + O(\delta d_\rho(x,0)) \tag{8}$$

2.

$$\frac{1}{4K} d_\rho(x,0) \leq d(\phi_i(x), e) \leq 4K d_\rho(x,0) \tag{9}$$

Proof. If $x, y \in \phi^{-1}(Y_i[\delta])$ then $\phi(x), \phi(y) \subset Y_i[\delta]$. Note that on $Y_i[\delta]$,

$$d(x, \pi_i(x)) = O(\delta d(x,e)). \tag{10}$$

Hence

$$d(\phi(x), \phi_i(x)) = O(\delta d(\phi(x), e)) = O(\delta d_\rho(x,0)) \tag{11}$$

by Condition II. Applying the triangle inequality together with (11) and (1) then gives (8).

To prove the second fact, we know from (11) and Condition II that, if $\epsilon \ll \delta$,

$$\frac{1}{2K} d_\rho(x,0) - O(\delta d_\rho(x,0)) \leq d(\phi_i(x), e) \leq 2K d_\rho(x,0) + O(\delta d_\rho(x,0)) \tag{12}$$

Now (9) follows from (12) if $\delta \ll 1$. □

5.3.2 Constructing a pinching function

The motivation for the definition of pinching function came from the behavior of functions such as $d(g(x), e)$, where g is a reasonably efficient map of $\mathbf{R}^2$ into $\mathbf{H}^2$, such as ϕ_i restricted to $\phi^{-1}((X^\delta)^c)$. The property of $\mathbf{H}^2$ used was that "travel transverse to the radial direction in $\mathbf{H}^2$ is exponentially expensive", as given precisely by the Expanding Annulus Lemma (Lemma 5.1).

The pinching functions we need to construct, however, use instead the neighborhood $Y_i[\delta]$ of a hyperbolic plane. This neighborhood only introduces a small linear error term $O(\delta r)$ in the functions involved, so proving that $d(\phi_i(x), e)$ is a pinching function on $\phi^{-1}(Y_i[\delta])$ should clearly follow from the Expanding Annulus Lemma. However, one must keep track of the error term and show that the axioms of a pinching function are satisfied.

Lemma 5.6 (Constructing a pinching function on $\phi^{-1}(Y_i[\delta])$). *If $\epsilon \ll \delta \ll 1$, then for each $i = 1, 2$ the set $\phi^{-1}(Y_i[\delta])$ has a pinching function f with pinching constants $(r_0, \omega = 1 + O(\delta), \beta = O(\delta))$, where $r_0 = r_0(\delta)$. Furthermore, f also satisfies:*

$$\frac{1}{2K} d_\rho(x, 0) \leq f(x) \leq 2K d_\rho(x, 0) \quad \text{for all } x \in \phi^{-1}(Y_i[\delta]) \tag{13}$$

Proof. We claim that $f(x) = d(\phi_i(x), e)$ is a pinching function on $\phi^{-1}(Y_i[\delta])$. Equation (13) follows directly from Condition II and the fact that the domain of f lies in $\phi^{-1}(Y_i[\delta])$. This also shows that f is proper.

Now suppose $x, y \in \phi^{-1}(Y_i[\delta])$, $s \geq r_0$ and $s \leq f(x) \leq f(y) \leq \omega s$. Let $\gamma : [0, 1] \to \phi^{-1}(Y_i[\delta])$ be a path from x to y satisfying $f(\gamma(t)) \geq s$ for all t. We must show $d(x, y) \prec \delta s$.

Let $\hat{\gamma}$ be the path constructed from $\phi \circ \gamma$ as in Lemma 2.2. Property 1 of Lemma 2.2 gives $\hat{\gamma} \subset \phi \circ \gamma[\lambda\epsilon]$. Hence $\hat{\gamma} \subset Y_i[2\delta]$ since $\epsilon \ll \delta$. Also the fact that $\hat{\gamma} \subset \phi \circ \gamma[\lambda\epsilon]$ implies that for every point $\hat{x} \in \hat{\gamma}$ there is a point $\hat{y} \in \phi(\gamma)$ so that $d(\hat{x}, \hat{y}) = O(\epsilon d(\hat{y}, e))$. Hence, if $\epsilon \ll \delta \ll 1$,

$$d(\pi_i(\hat{y}), \pi_i(\hat{x})) \prec \delta d(\hat{y}, e) \prec \delta d(\pi_i(\hat{y}), e)$$

Thus

$$d(\pi_i(\hat{x}), e) \geq d(\pi_i(\hat{y}), e) - d(\pi_i(\hat{y}), \pi_i(\hat{x})) = (1 - O(\delta)) d(\pi_i(\hat{y}), e).$$

Since $\hat{y} = \phi \circ \gamma(t)$ for some $t \in [0,1]$, $d(\pi_i(\hat{y}), e) = f(\gamma(t)) \geq s$. Thus for all $\hat{x}$ in the image of $\hat{\gamma}$,

$$(1 - \lambda\delta)s \leq d(\pi_i(\hat{x}), e) \leq d(\hat{x}, e) \tag{14}$$

for some constant λ.

We want to apply Lemma 5.1 to the path $\psi = \pi_i \circ \hat{\gamma}$. Let $r = (1 - \lambda\delta)s$. Then by (14), $d(\psi(t), e) \geq r$ for all t in $[0,1]$, hence the first hypothesis of Lemma 5.1 holds.

Let $\hat{A}_j$ denote the annulus $A(2^j r, 2^{j+1} r)$ in $\mathbf{H}^2$. From (14) it follows that $\pi_i^{-1}(\hat{A}_j) \cap Y_i[2\delta]$ lies in an annulus $A(2^{j-2} r, 2^{j+2} r)$ in X. By Property (2) of Lemma 2.2, we then get the bound

$$\ell(\pi_i \circ \hat{\gamma}) \prec 2^{j+3} R/\epsilon$$

so that hypothesis (3) in Lemma 5.1 holds.

Let $c_1 = 8\lambda\delta$, $\omega = 1 + (\lambda/2)\delta$, $a_0 = \phi_i(x)$, and $b_0 = \phi_i(y)$. Since $d(a_0, e), d(b_0, e) \in [s, \omega s]$, it easy to check that our choice of c_1 and ω gives $d(a_0, e), d(b_0, e) \in [r, (1 + c_1/4)r]$, as long as $\delta \ll 1$. So we may apply Lemma 5.1, which says that if r is sufficiently large then

$$d(\phi_i(x), \phi_i(y)) = d(a_0, b_0) < c_1 r < c_1 s = O(\delta s)$$

Hence

$$\begin{aligned} d(x,y) &\prec c_1 s + O(\delta d_\rho(x, 0)) && \text{by (8)} \\ &\prec c_1 s + O(\delta d(\phi_i(x), e) && \text{by (9)} \\ &\prec c_1 s + O(\delta s) && \text{since } d(\phi_i(x), e) \leq \omega s \leq 2s \\ &\prec \delta s \end{aligned}$$

$\square$

5.4 Proof of the Main Lemma (Lemma 4.6)

Proof of Lemma 4.6. Since if f is a pinching function on a set S, then f is a pinching function on any subset of S, and the first assertion follows immediately from Lemma 5.6 and Lemma 5.4, and the observation that by (13) and the definition of the pinching constant β, $\tilde{W} \subset W[\eta\delta]$ for $\eta = O(1)$. $\square$

6 Proof of Theorem 1.1

We will need the following lemma about triangles in hyperbolic space.

Lemma 6.1 (lemma about triangles in $\mathbf{H}^2$). *For every $\nu > 0$ there exists $\Lambda > 0$ so that if γ is a geodesic in $\mathbf{H}^2$ with endpoints x_1, x_2 and $\Theta_p(x_1, x_2) > \nu$. then $d(p, \gamma) < \Lambda$.*

Proof. This follows easily from the hyperbolic law of cosines. $\square$

There are simpler proofs of Theorem 1.1 than the one we are about to give. We give the following proof because it is the one that can be generalized to other higher rank symmetric spaces X (see [EF]).

We are given a (κ, C)-quasi-isometric embedding $\phi : \mathbf{R}^2 \to X$. The idea of the proof is roughly as follows.

First off, suppose $z \in X$ is a point which sees two far-away points $\phi(x_1), \phi(x_2) \in X^\delta(z)$ with both visual angles $\Theta_{\pi_i(z)}(\phi_i(x_1), \phi_i(x_2))$ bounded below. By Lemma 4.8, for each $j = 1, 2$, there is a path γ_j from x_j to infinity in $\phi^{-1}(X^\delta(z))$ with $\phi \circ \gamma_i$ limiting to a limit point $\beta_j \in S^1 \times S^1$. If we knew that the paths $\phi \circ \gamma_j$ stayed outside a large ball containing z, then by the Lemma 3.3 (no shifting lemma) $\beta_1 \neq \beta_2$. Since the visual distances $\Theta_{\pi_i(z)}(\phi_i(x_1), \phi_i(x_2))$ are bounded below, the visual distance $\Theta_{\pi_i(z)}(\pi_i(\beta_1), \pi_i(\beta_2))$ is bounded below (again by no shifting). Then by Lemma 6.1 $\pi_i(z)$ must be a bounded distance from the geodesic connecting $\pi_i(\beta_1)$ to $\pi_i(\beta_2)$. Let $F(\beta_1, \beta_2)$ be the flat which is the product of these geodesics for $i = 1, 2$. It is then clear that z is a bounded distance from the flat $F(\beta_1, \beta_2)$.

Why do the paths $\phi \circ \gamma_j$ stay outside a large ball containing z? There are two cases: either z is far from $\phi(\mathbf{R}^2)$ (and we are obviously done), or $\phi(\mathbf{R}^2)$ comes close to z at a point $p \in \phi(\mathbf{R}^2)$, in which case we can (Lemma 4.8) find paths γ_i as above which stay far from $\phi^{-1}(p)$. Then $\phi \circ \gamma_j$ would stay far from p since ϕ is a quasi-isometry. Since p and z are close, this implies that the paths stay far away from z as well.

Now given an arbitrary $x_1 \in \mathbf{R}^2$, we can use Lemma 4.7 to find $x_2 \in \mathbf{R}^2$ which is a large (but fixed) distance R from x_1, and so that each $\phi(x_i) \in X^\delta(z)$, where z is the midpoint of the geodesic from $\phi(x_1)$ to $\phi(x_2)$. We then apply the above argument to show that z is close to a flat connecting two limit points of ϕ, and so $\phi(x_1)$ is close plus $\kappa R/2$ to that flat.

Notice that, in the above argument, we really had to apply all of the previous machinery that gave existence of limit points two different times. We write out the step which will be repeated in the following:

Proposition 6.2. *Suppose $0, x_1, x_2 \in \mathbf{R}^2$, $e \in X$ and $\epsilon, \delta, K, \lambda_0, r > 0$ are such that*

(a) *ϕ satisfies (1) and (2) outside some ball B centered at 0.*

(b) *$\epsilon \ll \delta \ll 1$ (the implied constant depending on K).*

(c) *$x_1, x_2 \notin B$.*

(d) *$x_1, x_2 \in \phi^{-1}(X^\delta)$.*

(e) *For $i = 1, 2$, $\Theta_e(\phi_i(x_1), \phi_i(x_2)) \geq \lambda_0 > 0$.*

(f) *$d(\phi(x_2), e) \geq d(\phi(x_1), e) \geq r$ with r sufficiently large (depending on $K, \delta, \epsilon, \lambda_0$).*

Then there exist constants δ' (depending on K, δ) and Λ (depending on λ_0) such that

$$d(e, F(\beta, \gamma)) \leq \Lambda, \qquad \text{with } \beta, \gamma \in \mathcal{L}_{\delta'}(\phi).$$

Proof. Let $\delta' = \delta/\eta_1$ where η_1 is as in Proposition 4.9. Then by Proposition 4.9 there exist δ'-limit points β_1, β_2 so that for $i = 1, 2$, $j = 1, 2$,

$$\Theta_e(\phi_i(x_j), \pi_i(\beta_j)) \leq e^{-\nu_2 d_\rho(x_j, 0)}.$$

As r increases, so does $d_\rho(x_j, 0)$ by (2). Then for sufficiently large r, $\Theta_e(\phi_i(x_j), \pi_i(\beta_j)) \leq \lambda_0/4$, so by the triangle inequality, $\Theta_e(\pi_i(\beta_1), \pi_i(\beta_2)) > \lambda_0/2$. Then by Lemma 6.1, e must be a bounded distance from the geodesic connecting $\pi_i(\beta_1)$ to $\pi_i(\beta_2)$. Let $F(\beta_1, \beta_2)$ be the flat which is the product of these geodesics for $i = 1, 2$. It is then clear that e is a bounded distance from the flat $F(\beta_1, \beta_2)$. $\square$

Choosing constants. Choose ν smaller than a large negative power of κ (say $\nu = 1/(32\kappa^3) < 1/2$). Choose ϵ, δ so that Lemma 4.7 is satisfied with $K = \kappa$, and so that condition (b) of Proposition 6.2 is satisfied both with $K = \kappa$ and $K = \kappa^2/\nu$. Note that ϵ and δ depend only on κ. Now choose R so large and that condition (f) of Proposition 6.2 is satisfied with $r = R/(4\kappa)$, with ϵ, δ as chosen, and both with $K = \kappa$ and with $K = \kappa^2/\nu$. We also require that $R > 16\kappa^2 C/\epsilon$. Now let $\bar{\delta}$ denote one half of the smaller of the two constants δ' produced by the two applications ($K = \kappa$ and $K = \kappa^2/\nu$) of Proposition 6.2.

Method of proof. We will show the following: for every point $x_1 \in \mathbf{R}^2$, $\phi(x_1)$ is within a bounded distance of a flat of the form $F(\beta_1, \beta_2)$ where β_1, β_2 belong to the

$\bar{\delta}$-limit set $\mathcal{L}_{\bar{\delta}}(\phi)$, and $\pi_i(\beta_1) \neq \pi_i(\beta_2)$, $i = 1, 2$. This is enough since the cardinality of the limit set is bounded by Proposition 4.2 (limit set is finite), and the limit set is independent of the point x_1 by Lemma 4.4 and the choice of $\bar{\delta}$.

Separation into cases. Let $x_1 \in \mathbf{R}^2$ be an arbitrary point. By Lemma 4.7 there exists a point $x_2 \in \phi^{-1}(X^\delta(\phi(x_1)))$ with $d(x_1, x_2) = R$. Let z denote the midpoint of the geodesic in X between $\phi(x_1)$ and $\phi(x_2)$. We choose coordinates on X so that the origin e is z. An easy similar triangles argument shows that $x_1, x_2 \in \phi^{-1}(X^\delta(e))$, hence condition (d) of Proposition 6.2 holds. Since e is the midpoint of the geodesic connecting $\phi(x_1)$ and $\phi(x_2)$, $\Theta_e(\phi_i(x_1), \phi_i(x_2)) = \pi$. Thus condition (e) of Proposition 6.2 also holds. The conditions (b) and (f) hold by the choice of constants. Hence we need to verify the conditions (a) and (c) of Proposition 6.2. There are two cases to consider:

Case 1: $d(\phi(\mathbf{R}^2), z) > \nu R$. In this case choose coordinates so that $x_1 = 0 \in \mathbf{R}^2$. We claim that the map ϕ satisfies (1) and (2) on all of $\mathbf{R}^2$ with $\rho = \kappa^2 R$ and $K = \kappa^2/\nu$. The equation (1) holds because ϕ is a κ-quasi-isometry, $\rho > R > 4C/\epsilon$. and $K > \kappa$. The lower bound in (2) follows like this: For any $u \in \mathbf{R}^2$,

$$\begin{aligned}
d(\phi(u), e) &\geq \max(d(\phi(u), \phi(0)) - d(\phi(0), e), \nu R) \\
&\geq \max(\frac{1}{\kappa} d(u, 0) - \kappa R/2 - 3C/2, \nu R) \\
&\geq \max(\frac{1}{\kappa} d(u, 0) - 2\kappa R/3, \nu R) \\
&\geq \frac{\nu}{\kappa^2} \max\{\frac{\kappa}{\nu}(d(u, 0) - 2\kappa^2 R/3), \kappa^2 R)\} \geq \frac{\nu}{\kappa^2} \max\{d(u, 0), \kappa^2 R\} \\
&= \frac{1}{K} d_\rho(u, 0)
\end{aligned}$$

where the second line follows from $d(\phi(0), e) = \frac{1}{2} d(\phi(x_1), \phi(x_2)) \leq \frac{1}{2}(\kappa R + C)$, and the second to last line following since: If $d(u, 0) \leq \kappa^2 R$ then it is clear; if $d(u, 0) > \kappa^2 R$ then $\frac{\kappa}{\nu} \cdot \frac{1}{3} d(u, 0) > d(u, 0)$ since $\kappa > 3\nu$.

The upper bound in (2) follows from

$$\begin{aligned}
d(\phi(u), e) \leq d(\phi(u), \phi(0)) + d(\phi(0), e) &\leq \kappa d(u, 0) + \frac{\kappa R}{2} + \frac{3}{2} C \\
&\leq \frac{\kappa^2}{\nu} \max(d(u, 0), \kappa^2 R) = K d_\rho(u, 0) \quad (15)
\end{aligned}$$

since $\nu < 1$, $\kappa > 1$ and $R > 4C/\epsilon > 4C$.

Thus condition (a) of Proposition 6.2 is verified in case 1. Condition (c) of Proposition 6.2 is vacuous since B is empty.

Case 2: There exists $p \in \phi(\mathbf{R}^2)$ such that $d(p,e) \leq \nu R$.

Choose coordinates on $\mathbf{R}^2$ so that $\phi^{-1}(p)$ is the origin 0. We claim that ϕ satisfies (1) and (2) with $K = \kappa$ and $\rho = 0$ outside the ball $B = B(0, r')$ where $r' = \max(2C/\epsilon, 2\kappa(\nu R + C))$. Indeed (1) is satisfied outside B because ϕ is a (κ, C) quasi-isometry, and for points u outside B, $\epsilon d(u,0) \geq 2C$. The lower bound of (2) is satisfied outside B because if $u \notin B$ then $d(u,0) \geq 2\kappa(\nu R + C)$, hence

$$\begin{aligned} d(\phi(u),e) \geq d(\phi(u),p) - d(p,e) &\geq \frac{1}{K} d(u,0) - d(p,e) - C \\ &\geq \frac{1}{K} d(u,0) - (\nu R + C) \geq \frac{1}{2K} d(u,0) \end{aligned} \tag{16}$$

and the upper bound following similarly. Thus condition (a) of Proposition 6.2 is satisfied.

Note that for $i = 1, 2$,

$$d(x_i, \phi^{-1}(p)) \geq \frac{1}{\kappa} d(\phi(x_i), p) - C \geq \frac{1}{\kappa} d(\phi(x_i), e) - \frac{1}{\kappa} d(e,p) - C \geq \left[\frac{1}{4\kappa^2} - \frac{2\nu}{\kappa} \right] R. \tag{17}$$

In the last inequality we used $d(\phi(x_i), e) = \frac{1}{2} d(\phi(x_1), \phi(x_2)) \geq \frac{1}{\kappa} R - C$. By (17) and the choices of ν and R, we know $x_1, x_2 \notin B$, so condition (c) of Proposition 6.2 is also satisfied.

Completing the proof. By Proposition 6.2, $d(e, F(\beta,\gamma)) \leq \Lambda$. But then

$$d(\phi(x_1), F(\beta,\gamma)) \leq d(\phi(x_1), z) + d(z, F(\beta,\gamma)) \leq \frac{1}{2}(\kappa R + C) + \Lambda.$$

This proves the theorem. $\square$

7 Consequences

7.1 Quasi-flats with holes

Suppose $\Omega \subset \mathbf{R}^2$. Let $\Omega_{\epsilon,R}$ be the set of all $x \in \Omega$ such that for any $y \in \mathbf{R}^n \setminus B(x,R)$ the closed ball $B(y, \epsilon d(x,y))$ contains points of Ω, i.e there exists $z \in \Omega$ with $d(z,y) \leq \epsilon d(x,y)$. If $\Omega_{\epsilon,R}$ is non-empty, we think of Ω as a "quasi-flat with holes".

As a corollary of the proof of Theorem 1.1 we have the following theorem:

Theorem 7.1 (Quasi-flats with holes). *Suppose ϕ is a (κ, C) quasi-isometric embedding from $\Omega \subset \mathbf{R}^2$ into $X = \mathbf{H}^2 \times \mathbf{H}^2$. Then there exist constants $F = F(\kappa)$ and $\epsilon_0 = \epsilon_0(\kappa)$ such that if $\epsilon < \epsilon_0$ then, if $\Omega_{\epsilon,R}$ is non-empty, $\phi(\Omega_{\epsilon,R})$ lies in the N-neighborhood of F flats in X, where N depends on κ, C, and R.*

The version of this theorem for arbitrary higher rank symmetric spaces X ([EF]) is is used in [E] to deduce quasi-isometric rigidity for non uniform lattices in semisimple groups without rank one factors.

Proof. Suppose $x \in \Omega_{\epsilon,R}$. We build a grid as in §2.2, so that each point y is contained in a cell of size between $(\epsilon/2)d(x, y)$, and $\epsilon d(x, y)$. Then, by the definition of $\Omega_{\epsilon,R}$, for every cell C with $d(C, x) \geq R$, $\Omega \cap C$ is not empty; we pick a point y_C in the intersection.

We can now "connecting the dots" (see §2) to construct a continuous function ϕ_x defined on all of $\mathbf{R}^2$ which agrees with ϕ on every y_C, and satisfying for all $y \in \Omega - B(x, R)$

$$d(\phi(y), \phi_x(y)) \leq 4\kappa\epsilon d(y, x). \tag{18}$$

Choose ϵ_0 to be $1/(8\kappa)$ times the ϵ chosen in the proof of Theorem 1.1. We then repeat the proof of Theorem 1.1, except that immediately replace ϕ by ϕ_x. The function ϕ_x satisfies (1) and (2) because of (18) combined with the fact that ϕ does. Thus the proof of Theorem 1.1[2] shows that x is within a bounded distance from one of the flats $F(\beta, \gamma)$ where β and γ belong to the limit set $\mathcal{L}_{\bar{\delta}}(\phi_x)$. But the this limit set is (up to replacing $\bar{\delta}$ by $\bar{\delta}/2$) independent of the choice of basepoint $x \in \Omega_{\epsilon,R}$ because of (18) and an argument analogous to the proof of Lemma 4.5. $\square$

7.2 Concluding Remarks

It is possible to use Theorem 1.1 to prove that any quasi-isometry of $X = \mathbf{H}^2 \times \mathbf{H}^2$ lies a bounded distance from a product of quasi-isometries of $\mathbf{H}^2$. The idea is that Theorem 1.1 implies that, under a quasi-isometry $\phi : X \to X$, the image of a flat lies close to a finite union of flats. Since the domain and range contain many flats, it is possible to make arguments using the geometry of X to show that this finite union actually consists of a single flat. Once one knows that ϕ take flats to flats, the result (in the case $X = \mathbf{H}^2 \times \mathbf{H}^2$) is an easy exercise.

For more details we refer the reader to [EF], where the argument is worked out for more general X.

[2]To handle case 2 in the proof of Theorem 1.1 we must remark that if $x \in \Omega_{\epsilon,R}$ then so is any point p as long as R is sufficiently large compared to ϵ and $d(x, p)$.

References

[E] A. Eskin, Quasi-isometric rigidity of nonuniform higher rank lattices, in preparation.

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

[FS] B. Farb and R. Schwartz, The large-scale geometry of Hilbert modular groups, 1994 preprint (submitted).