

QUASI-FLATS AND RIGIDITY IN HIGHER RANK SYMMETRIC SPACES

ALEX ESKIN AND BENSON FARB

1. INTRODUCTION

In this paper we use elementary geometrical and topological methods to study some questions about the coarse geometry of symmetric spaces. Our results are powerful enough to apply to noncocompact lattices in higher rank symmetric spaces, such as $SL(n, \mathbb{Z})$, $n \geq 3$: Theorem 8.1 is a major step towards the proof of quasi-isometric rigidity of such lattices ([E]). We also give a different, and effective, proof of the theorem of Kleiner-Leeb on the quasi-isometric rigidity of higher rank symmetric spaces ([KL]).

Symmetric spaces of noncompact type. A symmetric space of noncompact type is a nonpositively curved symmetric space with no Euclidean (de Rham) factors, or, what is the same thing, the quotient G/K of a semisimple¹ Lie group G by a maximal compact subgroup K . These spaces are the most classical and important examples of nonpositively curved spaces. The *rank* of a symmetric space X is the dimension of a (maximal) *flat* in X , i.e. the maximal dimension of an isometrically embedded Euclidean space in X . A flat in rank one is just a geodesic. Examples of symmetric spaces of noncompact type include the hyperbolic spaces (rank 1), the spaces $SL(n, \mathbb{R})/SO(n, \mathbb{R})$ (rank $n - 1$, where a flat is given by the subgroup of diagonal matrices), and their products (rank is additive).

Quasi-flats in symmetric spaces. A (coarse) *quasi-isometry* between metric spaces is a map $f : X \rightarrow Y$ such that, for some constants $\kappa, C, C' > 0$:

- (1) $\frac{1}{\kappa}d_X(x_1, x_2) - C \leq d_Y(f(x_1), f(x_2)) \leq \kappa d_X(x_1, x_2) + C$ 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 (κ, C) -quasi-isometry. A map satisfying (1) but not necessarily (2) is called a *quasi-isometric embedding* of X into Y . A basic example of a quasi-isometry is the following: 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 .

A central step in the proof of Mostow Rigidity ([Mo]) involves showing that, under a Γ -equivariant quasi-isometry of a symmetric space X (Γ a cocompact lattice

Received by the editors March 8, 1996 and, in revised form, March 10, 1997.

1991 *Mathematics Subject Classification.* Primary 22E40, 20F32.

Key words and phrases. Lie groups, discrete subgroups, geometric group theory.

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.

¹By *semisimple* we mean connected semisimple with finite center and no nontrivial compact factors.

in X), a flat is taken to within a uniformly bounded neighborhood of a flat. In this paper we consider the problem of determining the quasi-flats in X , i.e. the quasi-isometric embeddings $\mathbb{R}^n \rightarrow X$, $n = \text{rank}(X)$.

For the rank one symmetric spaces X of noncompact type, i.e. for the (real, complex, quaternionic, Cayley) hyperbolic spaces, any quasi-isometric embedding of \mathbb{R} , i.e. any quasi-geodesic, is contained in a neighborhood of a single geodesic. This fact is due to Mostow, and in fact goes back to a 1924 paper of Morse in dimension two. One would hope that, as in the rank one case, any quasi-isometric embedding of \mathbb{R}^n into a rank n symmetric space X of noncompact type would lie in a bounded neighborhood of a single flat in X . Unfortunately this is not true for any X with $\text{rank}(X) > 1$, as we will see in §1.1.

The main goal of this paper is to prove the following:²

Theorem 1.1 (Quasi-flats are close to a finite union of flats). *Let X be a symmetric space of noncompact type of rank $n \geq 2$, and let $\phi : \mathbb{R}^n \rightarrow X$ be a (κ, C) -quasi-isometric embedding. Then there are constants $F = F(\kappa), N = N(\kappa, C) < \infty$ such that $\phi(\mathbb{R}^n)$ lies in the N -neighborhood of a union of F flats in X .*

A proof of Theorem 1.1 in the more familiar, easy to visualize special case of $X = \mathbb{H}^2 \times \mathbb{H}^2$ (i.e. the product of two hyperbolic planes) can be found in [EF].

Application of the proof. In proving Theorem 1.1, we actually consider maps $\phi : \mathbb{R}^n \rightarrow X$ that are weaker than quasi-isometries (see Conditions I and II on page 661): the additive coarseness constant is allowed to depend on the distance to the origin. The point of this is to allow for maps ϕ that behave like quasi-isometries, but may remain undefined on many “holes” $H \subset \mathbb{R}^n$ of diameter roughly $d(H, 0)$. Such maps occur naturally when considering noncocompact lattices in higher rank, such as $SL(n, \mathbb{Z}), n \geq 3$. Theorem 1.1, or more precisely its extension to quasi-flats with holes (Theorem 8.1), plays a major role in the proof of quasi-isometric rigidity for such lattices (see [E]).

Quasi-isometries of higher rank symmetric spaces. For any metric space X , one can form the group $QI(X)$ 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 $QI(X)$ into a group.

When X has rank one, i.e. when X is a hyperbolic space, then $QI(X)$ is the group of (Carnot) quasi-conformal transformations of the sphere at infinity. In the cases when X is \mathbb{H}^n or $\mathbb{C}\mathbb{H}^n$ this group is infinite dimensional. Pansu ([Pa]) showed that, when X is quaternionic hyperbolic space (in dimension ≥ 2) or the Cayley hyperbolic plane, then $QI(X)$ is isomorphic to $Isom(X)$. We use Theorem 1.1 to prove that this strengthening of Mostow rigidity also holds for most higher rank symmetric spaces. This result, first proved by B. Kleiner and B. Leeb ([KL]) using different methods, was conjectured by G. Margulis nearly twenty years ago.

Theorem 1.2 (Kleiner-Leeb [KL]). *Let X be a symmetric space of noncompact type with no rank one (de Rham) factors. Then any quasi-isometry of X is a bounded distance from an isometry of X . Hence the natural map*

$$Isom(X) \rightarrow QI(X)$$

is an isomorphism.

²After completion of work on this manuscript, we learned that Kleiner-Leeb have also obtained Theorem 1.1 (but not its extension Theorem 8.1) by different methods.

Roughly, Theorem 1.2 can be derived from Theorem 1.1 as follows: apply Theorem 1.1 to flats in X , then use the geometry of X (as in the proof of Mostow rigidity in [Mo]) to show that any global quasi-isometry q of X must actually take a flat to a *single* flat. Hence q induces an automorphism of the Tits building of X , so by Tits' theorem this boundary map is induced by an isometry, which one shows is a bounded distance from q . We give the details not present in [Mo] in §8.3.

An immediate corollary of Theorem 1.2 is the following:

Corollary 1.3 (Quasi-isometric iff isometric). *Let X and Y be two symmetric spaces of noncompact type with no rank one factors. Then X is quasi-isometric to Y if and only if X is isometric to Y (after multiplying the metrics on the direct factors of Y by constants).*

Theorem 1.2 can also be used to characterize cocompact lattices in $Isom(X)$ among all finitely generated groups. Recall that any finitely generated group Γ with chosen generating set may be endowed with the *word metric*, where the distance $d(g, h)$ is the minimal number of generators needed to represent gh^{-1} . The word metric on Γ is unique up to quasi-isometry, so for example is independent of choice of generating set.

Corollary 1.4 (QI rigidity). *Let X be a symmetric space of noncompact type with no rank one (de Rham) factors. If Γ is any finitely generated group which is quasi-isometric to X , then Γ is a finite extension of a cocompact lattice in $Isom(X)$.*

Reader's guide. To roughly understand at a glance the main issues and ideas of this paper, we suggest understanding the example in §1.1, followed by the method of proof in §1.2.

For the reader who is not comfortable with the geometry of symmetric spaces, many of the ideas of this paper can be gleaned from the special case when $X = \mathbb{H}^2 \times \mathbb{H}^2$. This special case is worked out in detail in [EF].

In §2, we state the generalizations of the geometric facts about $\mathbb{H}^2 \times \mathbb{H}^2$ used in [EF]. Some of the proofs of these facts are relegated to the Appendix. The proof of Theorem 1.1 beginning in §1.2 follows closely the proof of the $\mathbb{H}^2 \times \mathbb{H}^2$ case, with the exception of §6 where some extra ideas (an induction involving higher-dimensional pinching) are needed.

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 \mathbb{R}^2 in $\mathbb{H}^2 \times \mathbb{H}^2$ which does not lie in a bounded neighborhood of a single flat in $\mathbb{H}^2 \times \mathbb{H}^2$. This type of example can be easily generalized to any symmetric space X with $\text{rank}(X) \geq 2$.

Key Example. For each factor of $\mathbb{H}^2 \times \mathbb{H}^2$, pick three distinct geodesic rays in \mathbb{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 $\mathbb{H}^2 \times \mathbb{H}^2$. We denote these quarter-flats by AX, XB , etc. Divide the Euclidean plane \mathbb{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).

Note that the image of any adjacent pair of sectors lies uniformly close to a half-flat in $\mathbb{H}^2 \times \mathbb{H}^2$ since triangles in \mathbb{H}^2 are uniformly thin (see Figure 1). It follows that

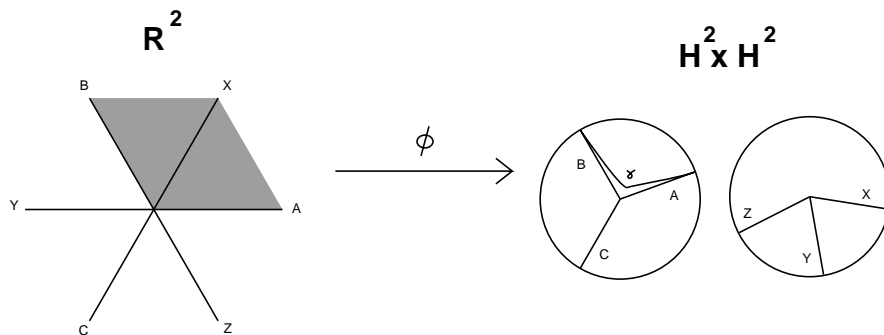


FIGURE 1. The map $\phi : \mathbb{R}^2 \rightarrow \mathbb{H}^2 \times \mathbb{H}^2$ maps each pair of adjacent sectors close to a half-flat: for example, the image of the shaded sectors under ϕ lies in a 4-neighborhood of the Euclidean half-flat $\gamma \times X$ in $\mathbb{H}^2 \times \mathbb{H}^2$,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

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 $\phi(\mathbb{R}^2)$ 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 $\phi : \mathbb{R}^2 \rightarrow \mathbb{H}^2 \times \mathbb{H}^2$ with a self-quasi-isometry of \mathbb{R}^2 , such as a quasi-isometry taking rays based at 0 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 \mathbb{R}^n 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 a symmetric space X there is a certain “degenerate subset”; when $X = \mathbb{H}^2 \times \mathbb{H}^2$ this 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 (page 661) the “nondegenerate subset” X^δ of X , where δ is a lower bound to the nondegeneracy. For example, when $X = \mathbb{H}^2 \times \mathbb{H}^2$, the set X^δ is the set of $x \in X$ with $d(\pi_i(x), e) \geq \delta d(x, e)$ for $i = 1, 2$, where $\pi_i : X \rightarrow \mathbb{H}^2$ 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 triangle³ into \mathbb{H}^2 : if there were such a map, then the image would lie in a neighborhood of a geodesic triangle in \mathbb{H}^2 (since quasi-geodesics are close to geodesics), and triangles in \mathbb{H}^2 are uniformly thin (unlike equilateral triangles in \mathbb{R}^2).

³By triangle we mean the sides only, not the interior.

Intuitively, any subset of \mathbb{R}^2 which admits a (κ, C) -quasi-isometric embedding into \mathbb{H}^2 should not contain any fat loops. We formalize this in Corollary 6.9, which states roughly: If W is any subset of \mathbb{R}^n which admits a (κ, C) -quasi-isometric embedding to a hyperbolic-like space, then for a certain metric neighborhood \tilde{W} of W , the natural homomorphism $H_p(W) \rightarrow H_p(\tilde{W}), p \geq 1$, induced by inclusion is the zero map. The size of the neighborhood \tilde{W} of course depends on κ .

The Main Lemma (Lemma 5.6) involves some variations on this, the most important being that the target hyperbolic space is replaced by the degenerate subset of X . The flat directions of X make this statement true only for $p = n - 1$. We apply the Main Lemma to subsets of $\phi^{-1}(X^\delta)^c$.

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

1.2.3. *Outline of the proof of Theorem 1.1.* Associated to a symmetric space X is its *Furstenberg boundary* \hat{X} . 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 \mathbb{H}^n parameterize the set of geodesics in \mathbb{H}^n . We identify \hat{X} with the “visual sphere” K/M at the origin of X .⁴

We consider the set of points S_R in \mathbb{R}^n 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 K/M for which the infinite ray in that direction hits $\phi(S_R)$. In §4 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 K/M of size roughly e^{-R} .

Proving that the visual image of ϕ actually *converges* to a finite set of points in \hat{X} is a major step towards proving the theorem. Once this is known, we connect each pair of these *limit points* $\alpha, \beta \in \hat{X}$ by a (unique) flat. A final geometric argument, outlined in §7, proves that $\phi(\mathbb{R}^n)$ 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 5.8 and Proposition 5.9): 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 $\beta \in K/M$ 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 4.3) shows that the visual size of any connected component of $\phi^{-1}(X^\delta)$ is exponentially small, i.e. the visual size of the intersection of $\phi^{-1}(X^\delta)$ with the complement of an r -ball has size $O(e^{-r})$. Thus, to prove visual images persist, it is enough to prove (Lemma 5.8) 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 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,

⁴Actually K/M is a quotient space of the usual visual sphere, although for simplicity we blur this distinction in this outline.

so could be completely surrounded by a set $T \subset \mathbb{R}^n$, thought of as a sphere S^{n-1} , which maps completely into the *degenerate* directions 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. THE METRIC GEOMETRY OF X

2.1. Polar coordinates and the metric. We collect here some basic facts about symmetric spaces and set some notation. The proofs of the facts stated here without reference may be found in [He] or [BGS]. Let $X = G/K$ be a Riemannian symmetric space, where G is a semisimple Lie group, and K a maximal compact subgroup. The Killing form $\langle \cdot, \cdot \rangle$ is a nondegenerate quadratic form on the Lie algebra \mathfrak{g} of G defined by $\langle X, Y \rangle = \text{tr ad } X \text{ ad } Y$. The Lie algebra \mathfrak{g} can be decomposed as $\mathfrak{k} \oplus \mathfrak{p}$ where \mathfrak{k} is the Lie algebra of K and \mathfrak{p} is the orthogonal complement relative to the Killing form. We choose a maximal Abelian subalgebra \mathfrak{a} in \mathfrak{p} , and denote $\exp \mathfrak{a}$ by A . The dimension of \mathfrak{a} is called the *rank* of the symmetric space X .

Simultaneous diagonalization of the linear transformations $\text{ad } H$ where H varies over \mathfrak{a} yields a decomposition of \mathfrak{g} into eigenspaces \mathfrak{g}^α , where each root space \mathfrak{g}^α is associated with a linear functional α on \mathfrak{a} . These are defined by the equation

$$(1) \quad [H, X_\alpha] = \alpha(H)X_\alpha$$

for $H \in \mathfrak{a}$ and $X_\alpha \in \mathfrak{g}^\alpha$. The linear functionals α are called the *roots*, and the \mathfrak{g}^α are called *root spaces*. We denote the collection of roots by Σ .

The root system Σ always has a lot of extra structure: in particular Σ has a basis Δ , such that every root $\beta \in \Sigma$ can be expressed as a linear combination $\beta = \sum_{\alpha \in \Delta} c_\alpha \alpha$, where the coefficients c_α are either all nonnegative integers or all nonpositive integers. Δ is called a *simple system*, and the elements of Δ *simple roots*. The cardinality of Δ is equal to the dimension of \mathfrak{a} , i.e. the rank of the symmetric space. We fix a simple system Δ . Let $\Sigma^+ \subset \Sigma$ denote the roots which are expressible as a linear combination of simple roots with nonnegative coefficients, and $\Sigma^- \subset \Sigma$ denote the roots which are expressible as a linear combination of simple roots with nonpositive coefficients. Then $\Sigma = \Sigma^+ \sqcup \Sigma^-$. Let \mathfrak{a}_+ denote the subset $\{H \in \mathfrak{a} : \alpha(H) > 0 \text{ for all } \alpha \in \Delta\}$; the set \mathfrak{a}_+ is called the *positive Weyl chamber*. Let $A_+ = \exp \mathfrak{a}_+ \subset A$, and let M denote the centralizer of A in K . We also call A_+ the (canonical) Weyl chamber.

Proposition 2.1 (Polar coordinates on X). *The map $\Phi : K/M \times \overline{A_+} \rightarrow G/K$ given by $\Phi(kM, a) = kaK$ is onto. The map Φ when restricted to $K/M \times A_+$ is a diffeomorphism, and the image is an open dense subset of X , called the set of nondegenerate elements.*

Proof. See [He, Chapter IX, Theorem 1.1]. □

Remark. If $G = SO(2, 1)$, then the “polar coordinates” defined above coincide with the polar coordinates on the hyperbolic plane $\mathbb{H}^2 = G/K$.

Definition 2.2. Proposition 2.1 allows us to consider the roots $\beta \in \Sigma$ as functions on X , via the formula $\beta(kaK) = \beta(\log a)$. We define the set of *nondegenerate elements* of X to be those $x \in X$ with $\beta(x) > 0$ for all $\beta \in \Delta$.

A good example to keep in mind is $X = \mathbb{H}^2 \times \mathbb{H}^2$, which has two roots α_1, α_2 . In this case $\alpha_i(x) = d(\pi_i(x), e)$, where $\pi_i : X \rightarrow \mathbb{H}^2$ is the natural projection onto the i th factor.

For a nondegenerate $y \in X$, let $\Theta(y) \in K/M$ be the unique element such that $y \in \Theta(y)A_+K$. For nondegenerate elements x, y of X , we can define the “angle” $\Theta_e(x, y)$ as $d_K(\Theta(x), \Theta(y))$, where d_K denotes the K -invariant metric on K/M . We think of $\Theta_e(x, y)$ as the “angle” between x and y as viewed from e . This interpretation is indeed correct when X has rank 1. In the higher rank case, Θ_e is not an angle at all; it is not always defined, and for example if $x, y \in A_+$ then $\Theta_e(x, y) = 0$. However Θ_e often behaves like an angle in hyperbolic space, see for example Lemma 2.4 below. We use the Θ -angles extensively in our analysis. We will also consider Θ with a different basepoint by defining for any $p \in X$

$$\Theta_p(x, y) = \Theta_e(g^{-1}x, g^{-1}y)$$

where g is any isometry taking e to p ; the right hand side is independent of the choice of g since d_K is left K -invariant. We think of $\Theta_p(x, y)$ as the “angle” between x and y as viewed from p .

We now write out the metric on X in polar coordinates. Let $\mathfrak{k}' \subset \mathfrak{k}$ denote the complement in \mathfrak{k} of the Lie algebra \mathfrak{m} of M . We identify the tangent spaces of K/M and A at kM and a , respectively, with \mathfrak{k}' and \mathfrak{a} via the differentials of the maps $\mathfrak{k}' \ni X \rightarrow k(\exp X)M \in K/M$ and $\mathfrak{a} \ni Y \rightarrow a \exp Y \in A$.

Proposition 2.3 (Symmetric space metric). *For a certain orthonormal basis $\{X'_\alpha\}$ of \mathfrak{k}' , the metric on X at the point kaK is given by*

$$ds^2 = da^2 + \sum_{\alpha \in \Sigma^+} \sinh^2 \alpha(a) (dX'_\alpha)^2$$

where da^2 is the standard Euclidean metric on \mathfrak{a} .

Proof. See the Appendix for the proof, and the definition of the basis $\{X'_\alpha\}$. □

In analogy with the ball model of hyperbolic space, we think of the \mathfrak{a} directions as radial and the X'_α directions as transverse to radial, with movement in the X'_α directions having exponential cost (cf. Lemma 4.1).

2.2. The boundary. A *Weyl chamber* in X is a translate gA_+K of the canonical Weyl chamber A_+K .⁵ The *boundary* of the symmetric space X , denoted \hat{X} , is defined as the set of Weyl chambers in X under the equivalence relation that two chambers are equivalent if the Hausdorff distance between them is finite. Each equivalence class contains a unique Weyl chamber of the form kA_+ where $k \in K/M$. This identifies K/M with \hat{X} . Note that the notion of boundary defined above differs from the Tits boundary of X in that the geodesic rays belonging to the same Weyl chamber are collapsed to one point. \hat{X} is often called the *Furstenberg boundary*, or the *maximal boundary* of X .

The identification of \hat{X} with K/M allows us to make sense of expressions like $\Theta_e(\beta, \gamma)$ where β, γ are elements of $X \cup \hat{X}$.

We say that a canonical Weyl chamber A_+ is *based* at e , and that a Weyl chamber gA_+ is based at the point $ge \in X$. For any point p of X we can define the visual map: $\text{Vis}_p : X \rightarrow \hat{X}$ sending x to the (equivalence class of) the unique Weyl chamber based at p containing x . This map is only defined if x is nondegenerate with respect to p , i.e. $\alpha(g^{-1}x) \neq 0$ for all $\alpha \in \Delta$, where g is any isometry taking e

⁵For simplicity of notation we will often identify $gA_+ \subset G$ with $gA_+K \subset X$.

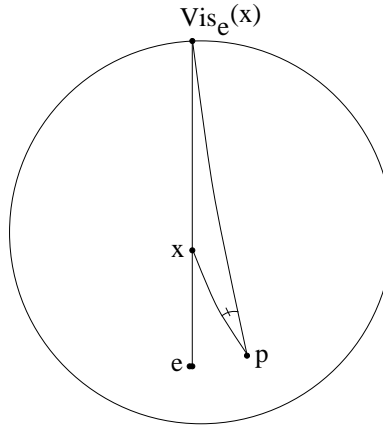


FIGURE 2. Illustration of the statement of Lemma 2.4. The marked angle is $\Theta_p(x, \text{Vis}_e(x))$.

to p . Note that under the identification of \hat{X} with K/M , $\text{Vis}_e(x)$ is identified with $\Theta(x)$.

G acts on \hat{X} in a way that is compatible with all the the visual maps $\text{Vis}_p, p \in X$:

$$(2) \quad \text{Vis}_{gp}(gx) = g \text{Vis}_p(x) \quad \text{for all } g \in G.$$

The definition $\Theta_p(x, y)$ extends naturally to \hat{X} and

$$\Theta_{gp}(gx, gy) = \Theta_p(x, y).$$

The following lemma illustrates why Θ can be thought of as “hyperbolic angle”.

Lemma 2.4. *For sufficiently small ν and sufficiently large r (depending on ν), the following holds: Suppose $p, x \in X$, $\alpha(x) \geq r$ for all $\alpha \in \Delta$, and $d(p, e) \leq \nu r$. Then*

$$\Theta_p(x, \text{Vis}_e(x)) \leq e^{-\lambda' r}$$

where λ' depends on X only, i.e. is independent of ν and r . See Figure 2.

Proof. See Appendix. □

3. PRELIMINARIES

3.1. Some definitions. Let $0 \in \mathbb{R}^n$ 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 annular region (centered at 0 unless otherwise specified) with inner radius r_1 and outer radius r_2 . Let ρ be a nonnegative real number, and let $d_\rho(x, 0) = \max(d(x, 0), \rho)$. One should think of d_ρ as the usual distance on \mathbb{R}^n , 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 \mathbb{R}^n by simplices 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, etc.)

by geodesics. This new quasi-isometry is continuous, and lies a bounded distance (in the sup-norm) from ϕ .

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

$$(3) \quad \frac{1}{\mathcal{K}}d(x, y) - \epsilon d_\rho(x, 0) \leq d(\phi(x), \phi(y)) \leq \mathcal{K}d(x, y) + \epsilon d_\rho(x, 0)$$

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:

$$(4) \quad \frac{1}{2\mathcal{K}}d_\rho(x, 0) \leq d(\phi(x), e) \leq 2\mathcal{K}d_\rho(x, 0).$$

Note that if ϕ is a (κ, C) -quasi-isometry in the usual sense, then ϕ satisfies (3) with $\rho = C/\epsilon$ and $\mathcal{K} = \kappa$. The constant ϵ can be made arbitrarily small by taking ρ sufficiently large.

If $\phi(0) = e$ and ϵ is sufficiently small (depending on \mathcal{K}), then (3) implies (4) outside a ball of radius ρ . If $\phi(0) \neq e$, then $d(\phi(x), \phi(0)) - d(\phi(0), e) \leq d(\phi(x), \phi(0)) \leq d(\phi(x), \phi(0)) + d(\phi(0), e)$, and (4) holds outside of a ball of radius $\max(\rho, 4\kappa d(\phi(0), e))$. The point of separating out the two conditions is that it will be necessary to consider the map ϕ with arbitrary points of \mathbb{R}^n viewed as basepoints. The more general form of (3) is needed in order to prove Theorem 8.1.

In all that follows we allow the conditions (3) and (4) 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 §7.

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

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

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

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

For any subset σ of Δ (where Δ is defined in §2), let

$$Y_\sigma = K \exp(\{H \in \mathfrak{a} : \alpha(H) = 0 \ \forall \alpha \notin \sigma\})K \subset X.$$

The union of the Y_σ is the set of degenerate elements of X . For any $\delta > 0$, let

$$(5) \quad X^\delta = X \setminus \left(\bigcup_{\alpha \in \Delta} Y_{\Delta - \{\alpha\}}[\delta] \right).$$

X^δ then consists of nondegenerate elements which are bounded away from the degenerate elements (see Figure 3).

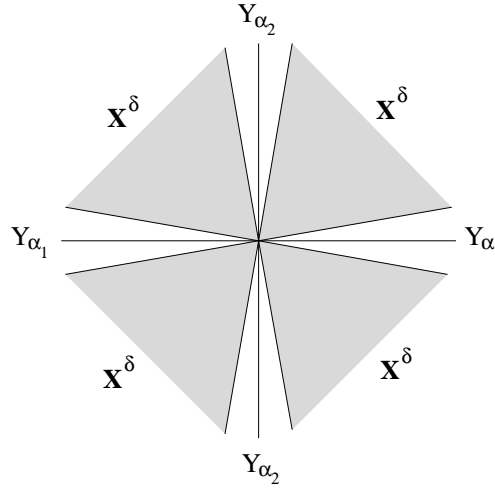


FIGURE 3. What the intersections of Y_σ and X^δ with a flat through $e \in X$ look like in the case $X = \mathbb{H}^2 \times \mathbb{H}^2$. In this case there are two roots α_1, α_2 .

Notational convention. All constants will depend on the symmetric space X as well as the quasi-isometry constant \mathcal{K} . Hence constants introduced in statements of lemmas and theorems without further mention should be thought of as numerical constants which could, in principal, be computed from \mathcal{K} and basic constants depending on X , such as the dimension and rank of X . We refrain from such explicit computations of constants 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 $\lambda = \lambda(\mathcal{K}, X)$. We also write $a \ll b$ to mean $a < \nu b$ for some small implied constant $\nu = \nu(\mathcal{K}, X) < 1$, and $a \gg b$ to mean $a > \eta b$ for some large implied constant $\eta = \eta(\mathcal{K}, X) > 1$. 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 than 1 but much larger than ϵ , i.e. $\epsilon \ll \delta \ll 1$, where ϵ is defined in (3).

Lemma 3.1. *For any $\alpha \in \Delta$, $|\alpha(x) - \alpha(y)| \prec d(x, y)$.*

Proof. This is clear if x and y lie in a common flat. The general case follows since Proposition 2.3 implies that the map $\pi_A : X \rightarrow \overline{A}_+$ given by $\pi_A(kaK) = a$ is distance nonincreasing. \square

It follows from the definitions that for all $x \in X^\delta$ and all $\alpha \in \Delta$,

$$(6) \quad \alpha(x) \succ \delta d(x, e).$$

Indeed $x \in X^\delta$ if and only if $d(x, y) \geq \delta d(x, e)$ for all degenerate $y \in X$. For $x = kaK$ and $\alpha \in \Delta$ take $y = y_\alpha = ka_\alpha K$ s.t. $d(x, y) = d(aK, a_\alpha K) \leq \lambda^{-1}\alpha(x)$, so that $\alpha(x) \geq \lambda\delta d(x, e)$ for all $\alpha \in \Delta$.

3.2. Approximations. In this section we prove a few simple technical statements used throughout the paper.

3.2.1. *Pullbacks and neighborhoods.*

Lemma 3.2. *Let $U \subset X$ be any subset. Then*

- (a) *If $\epsilon \ll \delta \ll 1$, then $\phi^{-1}(U)[\delta] \subset \phi^{-1}(U[4\mathcal{K}^2\delta])$.*
- (b) *$U[\alpha][\beta] \subset U[\alpha + \beta + \alpha\beta]$. Hence if $\alpha, \beta \ll 1$, then $U[\alpha][\beta] \subset U[\alpha + 2\beta]$.*
- (c) *If α and β , $\epsilon \ll \alpha, \beta \ll 1$, then*

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

In particular, for any $\sigma \subset \Delta$, if $\epsilon \ll \delta \ll 1$, then

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

Proof. Part (a) is an easy consequence of Condition I, Condition II, 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 $U = Y_\sigma$. □

3.2.2. *Approximating paths by edge-paths.* It will be necessary for us to approximate paths $\phi(\gamma) \subset X$, γ an arbitrary path in \mathbb{R}^n , by paths $\hat{\gamma}$ which stay “close” to $\phi(\gamma)$, and so that there is some control on the length of $\hat{\gamma}$.

For any $\epsilon > 0$, we divide \mathbb{R}^n into a *grid* \mathcal{G} , that is, a cellular decomposition of \mathbb{R}^n into rectangular cells so that each point $x \in \mathbb{R}^n$ 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 $d_\rho(C, 0)$.

Lemma 3.3 (Approximating paths by edge-paths). *Let $R > 0$ be given, and let $R_j = 2^j R$ for $j = 1, 2, \dots$. 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 $\gamma : [a, b] \rightarrow \mathbb{R}^n$, there is a path $\hat{\gamma} : [a, b] \rightarrow X$ 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) $\ell(\hat{\gamma} \cap \hat{A}_j) \prec \frac{1}{\epsilon^{n-1}} R_j$ for every $j = 1, 2, \dots$

Proof. Let \mathcal{C} denote the union of cells C with $C \cap \gamma \neq \emptyset$. Then \mathcal{C} is a 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 \mathbb{R}^n . 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

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

with the second inequality following from condition II. Hence $\phi(\gamma^{(\epsilon)}) \subset \phi(\gamma)[(4\mathcal{K}^2 + 2\mathcal{K})\epsilon]$. By construction of $\hat{\gamma}$ we now know $\hat{\gamma} \subset \phi(\gamma^{(\epsilon)})[(4\mathcal{K}^2 + 2\mathcal{K})\epsilon]$.

Therefore, by Lemma 3.2 part (b), $\hat{\gamma} \subset \phi(\gamma)[\lambda\epsilon]$, with e.g. $\lambda = 12\mathcal{K}^2 + 6\mathcal{K}$. Now since $\gamma^{(\epsilon)}$ is shortest, it traverses any edge in \mathcal{C} at most once. Hence

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

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

4. VISUAL ANGLES

4.1. Travel transverse to flats is expensive. Recall that for a nondegenerate $y \in X$, $\Theta(y) \in K/M$ denotes the unique element such that $y \in \Theta(y)A_+K$, and for nondegenerate $x, y \in X$, $\Theta_e(x, y) = d_K(\Theta(x), \Theta(y))$ where d_K denotes the K -invariant metric on K/M .

Lemma 4.1 (Travel transverse to flats is expensive). *For sufficiently large r and for x, y satisfying $\alpha(x) > r$, $\alpha(y) > r$ for all $\alpha \in \Delta$,*

- (i) *If $\Theta_e(x, y) > e^{-\nu r}$, then $d(x, y) > \lambda' r$.*
- (ii) *Any path $\hat{\gamma}$ connecting x, y and staying in the set $\{x : \alpha(x) > r \text{ for all } \alpha \in \Delta\}$ has length at least $\ell_K(\Theta(\hat{\gamma}))e^{\lambda'' r}$, where ℓ_K denotes lengths in the K -invariant metric on K/M .*

Proof. Part (ii) of Lemma 4.1 follows immediately from Proposition 2.3. For a proof of part (i) see the Appendix. \square

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

Lemma 4.2 (Bounded number of visual angles). *There exists a constant $\nu_1 = \nu_1(\delta)$ so that the set $\Theta(\phi(A(R, 2R)) \cap X^\delta) \subset K/M$ 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 $\Theta(\phi(A(R, 2R)) \cap X^\delta)$ consists of six points on $S^1 \times S^1$, corresponding to the six quarter-flats AX, XB, BY, YC, CZ, ZA in $\mathbb{H}^2 \times \mathbb{H}^2$.

Proof of Lemma 4.2. We use the shorthand $A_R = A(R, 2R)$. By (6), for all $x \in A_R \cap \phi^{-1}(X^\delta)$ and all $\alpha \in \Delta$, $\alpha(\phi(x)) \succ \delta d(\phi(x), e) \succ \delta R$, where the last estimate used Condition II. Hence Lemma 4.1 part (i) applies, giving a constant λ_1 such that the following holds: if $x, y \in A_R \cap \phi^{-1}(X^\delta)$ and $\Theta_e(\phi(x), \phi(y)) \geq e^{-\lambda_1 \delta R}$, then $d(\phi(x), \phi(y)) \succ \delta R$. For R sufficiently large (depending on ϵ, ρ), this combined with Condition I implies $d(x, y) \succ \delta R$, i.e. $d(x, y) \geq \lambda_2 \delta R$ for some constant λ_2 .

Decompose K/M into a disjoint union of cells $\bigsqcup_{p=1}^P I_p$, where each I_p is a cell with edge-lengths at most $e^{-\lambda_1 \delta R/4}$ and at least $e^{-\lambda_1 \delta R}$. Since K/M is finite dimensional, we can organize the I_p into $O(1)$ families, each family consisting of pairwise nonadjacent cells. Given any $x, y \in \phi^{-1}(X^\delta) \cap A_R$, suppose that $\Theta(\phi(x)) \in I_p$ and

$\Theta(\phi(y)) \in I_q$, where I_p and I_q are distinct and belong to the same family. Then since I_p and I_q are nonadjacent, $d_K(I_p, I_q) \geq e^{-\lambda_1 \delta R}$, hence $d(x, y) \geq \lambda_2 \delta R$.

From this it follows that the number of cells from within one family for which $I_p \cap \Theta(X^\delta \cap \phi(A_R))$ is nonempty is bounded by $R^n / (\lambda_2 \delta R)^n = O(1/\delta^n)$. Since the number of families is $O(1)$, this shows that the number of cells I_p for which $I_p \cap \Theta(X^\delta \cap \phi(A_R))$ is nonempty is $O(1/\delta^n)$. \square

We emphasize that the size of the balls in Lemma 4.2 decreases exponentially with R , while the number of balls is bounded independent of R . However, Lemma 4.2 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 4.3 (No Shifting Lemma). *For every $\delta, \epsilon \ll \delta \ll 1$, 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*

$$\Theta_\epsilon(\phi(x), \phi(y)) \leq e^{-\nu_2 \max(r, \rho)}$$

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 wedge S . The set $\Theta(\phi(S)) \subset K/M$ is a single point, so in this case Lemma 4.3 is obviously true.

Proof of Lemma 4.3. Let $R = \max(r, \rho)/(2K)$. By Condition II, if $d(x, 0), d(y, 0) \geq r$, then $d(\phi(x), e), d(\phi(y), e) \geq R$.

Let $\gamma \subset \phi^{-1}(X^\delta)$ be a path such that $\gamma(0) = x$ and $\gamma(1) = y$. Let \hat{A}_j denote the annular region $A(2^j R, 2^{j+1} R)$ in X , and let $\hat{\gamma}$ be the path constructed from $\phi \circ \gamma$ as in Lemma 3.3. Property (1) of Lemma 3.3 combined with part (c) of Lemma 3.2 and (5) gives $\hat{\gamma} \subset X^{\delta/2}$, as long as $\epsilon \ll \delta$.

By property (2) of Lemma 3.3 we have

$$(7) \quad \ell(\hat{\gamma} \cap \hat{A}_j) \prec 2^j R / \epsilon^{n-1}.$$

If $z \in \hat{A}_j$, then $d(z, e) \geq 2^j R$, thus if $z \in \hat{A}_j \cap X^{\delta/2}$, then by (6), $\alpha(z) \succ \delta 2^j R$ for all $\alpha \in \Delta$. Hence by part (ii) of Lemma 4.1 and (7),

$$\ell_K(\Theta(\hat{\gamma} \cap \hat{A}_j)) \leq \ell(\hat{\gamma} \cap \hat{A}_j) e^{-\lambda'' 2^j \delta R} \prec (2^j R / \epsilon^{n-1}) e^{-\lambda'' 2^j \delta R}$$

where ℓ_K denotes length in the K -invariant metric on K/M . Summing over j we get

$$(8) \quad \ell_K(\Theta(\hat{\gamma})) \prec (R / \epsilon^{n-1}) \sum_{j=0}^{\infty} 2^j e^{-\lambda'' 2^j \delta R}.$$

For sufficiently large r (and hence R), (8) is bounded by $e^{-\lambda \delta r}$ for some $\lambda > 0$, hence we can let $\nu_2 = \lambda \delta$. \square

5. LIMIT SETS

5.1. Definition and finiteness. Recall that, under the identification of the boundary \hat{X} with K/M , $\Theta(x) \in K/M$ is identified with $\text{Vis}_e(x) \in \hat{X}$.

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

Definition 5.1 (Limit point). A point $\beta \in \hat{X}$ is called a δ -limit point of ϕ if there exists a path in γ in $\phi^{-1}(X^\delta)$ such that $\gamma(t)$ leaves every compact set and the limit $\lim_{t \rightarrow \infty} \text{Vis}_e(\phi(\gamma(t)))$ exists and equals β . 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, denoted $\mathcal{L}_\delta(\phi)$, 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 $(a, x) \in \hat{X} = S^1 \times S^1$, where a (resp. x) corresponds to the direction of the ray A (resp. X) in \mathbb{H}^2 , is a limit point of ϕ .

Proposition 5.2 (The limit set is finite). *The cardinality of the δ -limit set of ϕ is at most c_0 (defined in Lemma 4.2).*

Proof of Proposition 5.2. Suppose there exist c_0+1 limit points $\beta_1, \dots, \beta_{c_0+1}$. Since the β_j are limit points there exist paths $\gamma_j : [0, \infty) \rightarrow \phi^{-1}(X^\delta)$ which leave every compact set and $\lim_{t \rightarrow \infty} \Theta(\phi(\gamma_j(t))) = \beta_j$. (Here we have identified Vis_e with Θ .) Choose r large so that Lemma 4.3 holds, so that the balls of radius $2e^{-\nu_1 r} + 2e^{-\nu_2 r}$ centered at the limit points $\beta_1, \dots, \beta_{c_0+1}$ 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 4.2 and Lemma 4.3. 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 4.3, $\Theta_e(\phi(x_j), \beta_j) \leq e^{-\nu_2 r}$. Hence, by the triangle inequality, for $i \neq j$, $\Theta_e(\phi(x_i), \phi(x_j)) \geq 2e^{-\nu_1 r}$. Since there are $c_0 + 1$ points x_j , this contradicts Lemma 4.2. \square

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

For any point $e' \in X$, we define $X^\delta(e')$ to be the translate gX^δ where g is any isometry such that $ge = e'$.

Lemma 5.3. *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 part (c) of Lemma 3.2 and the fact that if $x \in X^\delta(y)$, then $y \in X^\delta(x)$. \square

Lemma 5.4 (Limit set is independent of basepoint). *Suppose the map ϕ satisfies Conditions I and II with $e \in X$ replaced by some other point e' and $0 \in \mathbb{R}^n$ 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 $e' = gK \in X$ 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 5.3, if $\beta \in \mathcal{L}'_\delta(\phi)$ and γ is the path in $\phi^{-1}(X^\delta)$ so that

$$(9) \quad \lim_{t \rightarrow \infty} \text{Vis}_e(\phi \circ \gamma(t)) = \beta,$$

then the part of γ lying outside a sufficiently large ball centered at e is inside $X^{\delta/2}(e')$. By Lemma 2.4 with $r_0 = \nu r$,

$$(10) \quad \lim_{t \rightarrow \infty} \Theta_{e'}(\phi \circ \gamma(t), \text{Vis}_e(\phi \circ \gamma(t))) = 0.$$

Combining (9) with (10) we get

$$\lim_{t \rightarrow \infty} \text{Vis}_{e'}(\phi \circ \gamma(t)) = \beta.$$

Hence $\beta \in \mathcal{L}'_{\delta/2}(\phi)$. □

We also note the following, used in the proof of Theorem 8.1:

Lemma 5.5 (Invariance under linear growth perturbations). *Suppose ϕ and ψ are two functions satisfying Conditions I and 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 3.2 and (5), $\phi^{-1}(X^\delta) \subset \psi^{-1}(X^{\delta/2})$. Also by Lemma 4.1 part (i), $\Theta_e(\phi(x), \psi(x)) \rightarrow 0$ as $x \rightarrow \infty$ in $\phi^{-1}(X^\delta)$. Thus any δ -limit point of ϕ is also a $(\delta/2)$ -limit point of ψ . □

5.3. 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 of \mathbb{R}^n which map into the degenerate subset of X , that is, $\phi^{-1}(X^\delta)^c$, must be pinched. As stated in the introduction, our notion of W being pinched is the inclusion of W into some neighborhood inducing the zero map on $(n - 1)$ st 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 5.6 (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 $i_* : H_{n-1}(W \cup B(0, r)) \rightarrow H_{n-1}(W[\eta\delta] \cup B(0, r))$ induced by the inclusion is the zero map, for some $\eta > 1$ independent of δ .*

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

Lemma 5.7 (Nondegenerate 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 5.6 holds. Let $B = B(0, r)$ and let W denote the sphere 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 5.6, the inclusion induced homomorphism $H_{n-1}(W \cup B) \rightarrow H_{n-1}(W[\eta\delta] \cup B)$ is the zero map. Since $W[\eta\delta] \cap B = \emptyset$, this implies that $H_{n-1}(W) \rightarrow H_{n-1}(W[\eta\delta])$ is the zero map, and so $H_{n-1}(W) \rightarrow H_{n-1}(A((1 + \eta\delta)^{-1}R, (1 - \eta\delta)^{-1}R))$ is also the zero map. This is clearly a contradiction since the inclusion $W \rightarrow A((1 + \eta\delta)^{-1}R, (1 - \eta\delta)^{-1}R)$ is that of a sphere into a surrounding annular region. □

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

Lemma 5.8 (Connected components of $\phi^{-1}(X^\delta)$ 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. We choose $\eta_1 \gg \eta$, where η is from Lemma 5.6. Let $\delta = b/\eta_1$. We may assume $\epsilon \ll \delta \ll 1$; then by Lemma 5.7, $\phi^{-1}(X^\delta)$ is unbounded. 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 $R-1$). Without loss of generality one can assume that $z \in \phi^{-1}(\text{interior}(X^b))$, otherwise take any $b' < b$ and then $z \in \phi^{-1}(\text{interior}(X^{b'}))$. Hence $H_0(W_b \cap B^c) \neq 0$. Furthermore the image of the map $j_* : H_0(W_b \cap B^c) \rightarrow H_0(W_\delta \cap B^c)$ induced by the inclusion is not zero. Note that $W_\delta^c = \bigcup_{\alpha \in \Delta} \phi^{-1}(\overline{Y}_{\Delta-\{\alpha\}}[\delta]) \cup B(0, R+1)^c$. Therefore, by part (a) of Lemma 3.2, $W_\delta^c[\eta\delta] \subset W_b^c$, since $\delta = b/\eta_1$ and $\eta_1 \gg 1$.

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

$$\begin{array}{ccc} H_{n-1}(W_\delta^c \cup B) & \longrightarrow & H_{n-1}(W_b^c \cup B) \\ \cong \downarrow & & \cong \downarrow \\ H^0(W_\delta \cap B^c) & \longrightarrow & H^0(W_b \cap B^c) \end{array}$$

The vertical maps are Alexander duality isomorphisms. Hence the top horizontal map is nonzero. But by Lemma 5.6 the top horizontal map is zero. This is a contradiction. \square

Proposition 5.9 (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 β in \hat{X} such that*

$$\Theta_\epsilon(\beta, \phi(z)) < e^{-\nu_2 d_\rho(z, 0)}.$$

The constant η_1 is the same as in Lemma 5.8 and depends only on \mathcal{K} ; the constant ν_2 depends on \mathcal{K} and δ .

Proof. Let $s = d(z, 0)$. By Lemma 5.8, there exists a path $\gamma : [0, \infty) \rightarrow \phi^{-1}(X^{b/\eta_1}) \subset \mathbb{R}^n$ such that $\gamma(0) = z$, $d(\gamma(t), 0) \geq s$ for all t and such γ leaves every ball centered at 0. Let $r = \max(\rho, s)$, and for a nonnegative integer j let t_j denote the greatest value of t for which $d_\rho(\gamma(t_j), 0) = 2^j r$. Then by Lemma 4.3, for all $t \geq t_j$, $\Theta_\epsilon(\phi \circ \gamma(t), \phi \circ \gamma(t_j)) \leq e^{-\nu_2 2^j r}$. This implies that the limit as $t \rightarrow \infty$ of $\Theta(\phi \circ \gamma(t))$ exists: call it β . The above estimate for $j = 0$ shows that $\Theta_\epsilon(\phi(z), \beta) \leq e^{-\nu_2 r}$. \square

6. PROOF OF THE MAIN LEMMA (LEMMA 5.6)

6.1. The Expanding Annulus Lemma. Let v denote a unit vector in $\overline{\mathfrak{a}}_+$ (typically v will be taken on the walls of $\overline{\mathfrak{a}}_+$). We define a subset $Y_v = K \exp(\mathbb{R}^+ v) K$ of X . We think of the subset Y_v as playing the role of a hyperbolic space in X . The

v -direction is the “radial direction” of Y_v , with the K -movement (rotation) transverse to this radial direction and having exponential cost, just as in the hyperbolic plane. The first of the two geometric facts we use in the the proof of the Main Lemma is the following.

Lemma 6.1 (Expanding Annulus Lemma in Y_v). *Let $c_1 \ll 1$ and $c_2 > 0$ be given. Let a_0, b_0 be two points in $Y_v \cap A(r, (1 + c_1/4)r)$. Suppose $\hat{\psi}$ is a path in Y_v from a_0 to b_0 for which*

- $d(\hat{\psi}(t), e) \geq r$ for all t .
- $\ell(\hat{\psi} \cap A(2^i r, 2^{i+1} r)) \leq c_2 2^i r$.

Then $d(a_0, b_0) < c_1 r$ if r is sufficiently large (depending on c_1 and c_2).

Proof. See the Appendix. □

Note on uniformity. From the proof, we see that minimal radius r_0 for which the lemma holds, depends on v only via the the quantity $\alpha_{\min}(v) = \min \alpha(v)$ where the minimum is taken over the $\alpha \in \Delta$ for which $\alpha(v) \neq 0$.

6.2. Coarse topology and pinched subsets of \mathbb{R}^n . The following definition is motivated by Lemma 6.1.

Definition 6.2 (Pinching function). Let W be a subset of \mathbb{R}^n . A nonnegative, proper, continuous function $f : W \rightarrow \mathbb{R}$ 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 $\gamma : [0, 1] \rightarrow W$ 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 $W \subset \mathbb{R}^n$ 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 \mathbb{R}^n .

Lemma 6.3 (Trivial \mathbb{R}^n fact). *If $r > 0$ and if $T \subset \mathbb{R}^n$ 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. □

Lemma 6.4 (Pinching lemma). *Suppose $W \subset \mathbb{R}^n$, and suppose f is a 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 neighborhood of W :*

$$\tilde{W} = \tilde{F} \cup \{x \in \mathbb{R}^n : \exists w \in W \text{ with } d(x, w) < 4\beta f(w)\}.$$

Then for any $p \geq 1$, the homomorphism $i_ : H_p(W) \rightarrow H_p(\tilde{W})$ induced by the inclusion $i : W \rightarrow \tilde{W}$ is the zero homomorphism.*

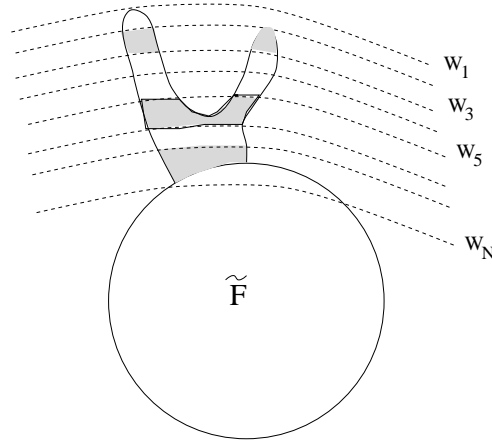


FIGURE 4. Applying the definition of pinching function shows that each of the regions like the shaded regions must have small diameter. Hence $W = \bigcup_{i=1}^N W_i$ must be pinched.

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

$$(11) \quad \begin{array}{ccccc} H_p(F) & \longrightarrow & H_p(W) & \longrightarrow & H_p(W, F) \\ \downarrow & & \downarrow & & \downarrow \\ H_p(\tilde{F}) & \longrightarrow & H_p(\tilde{W}) & \longrightarrow & H_p(\tilde{W}, \tilde{F}) \end{array}$$

and the contractibility of \tilde{F} , it is enough to show that the right vertical homomorphism in (11) 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 $x \in \mathbb{R}^n$ 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. Pick a decreasing sequence s_0, \dots, 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 $W_N = F$. The idea of the proof is illustrated in Figure 4.

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

$$(*)_k \quad j_* : H_p(Q, Q \cap W_k) \rightarrow H_p(\tilde{Q}, \tilde{Q} \cap \tilde{W}_k) \quad \text{is the zero map}$$

where j_* is induced by the inclusion $j : Q \rightarrow \tilde{Q}$. 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.

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

$$(12) \quad \begin{array}{ccccc} H_p(Q \cap W_{k-1}, Q \cap W_k) & \longrightarrow & H_p(Q, Q \cap W_k) & \longrightarrow & H_p(Q, Q \cap W_{k-1}) \\ \downarrow & & \downarrow j_* & & \downarrow j'_* \\ H_p(\tilde{Q} \cap \tilde{W}_{k-1}, \tilde{Q} \cap \tilde{W}_k) & \longrightarrow & H_p(\tilde{Q}, \tilde{Q} \cap \tilde{W}_k) & \xrightarrow{\Phi_*} & H_p(\tilde{Q}, \tilde{Q} \cap \tilde{W}_{k-1}) \end{array}$$

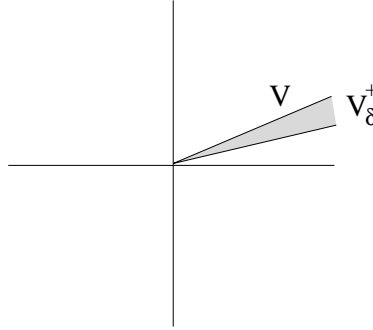


FIGURE 5. A schematic of V_δ^+ sitting inside the Abelian Lie sub-algebra \mathfrak{a} .

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 $\widetilde{Q \cap W_k} \subset \widetilde{Q \cap W_{k-1}} \subset \widetilde{Q}$, and the vertical maps are induced by the inclusion $Q \rightarrow \widetilde{Q}$.

The fact that j'_* in (12) 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 ,

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

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 $j_* : H_p(S_i, S_i \cap W_{k-1}) \rightarrow H_p(\widetilde{S}_i, \widetilde{S_i \cap W_{k-1}})$ induced by the inclusion $S_i \rightarrow \widetilde{S}_i$ is the zero map. Since for each i the inclusion $S_i \rightarrow \widetilde{Q}$ factors through the inclusion $S_i \rightarrow \widetilde{S}_i$, this implies that the map j'_* in (12) 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 $\gamma : [0, 1] \rightarrow W$ 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 6.3 implies that \widetilde{T} is contractible. Hence $H_p(\widetilde{Q \cap W_{k-1}}, \widetilde{Q \cap W_k}) = 0$, so the map Φ_* in (12) is injective. This, together with the fact that $j'_* = 0$, implies that the map j_* in (12) is the zero map, which proves $(*)_k$. Since $W_N = F$, this proves that the right vertical homomorphism in (11) is indeed zero. \square

6.3. Constructing a pinching function on $\phi^{-1}(Y_v[\delta])$. This section contains the second of the two geometric facts we use in the the proof of the Main Lemma.

6.3.1. A contraction. Let v be a unit vector in $\overline{\mathfrak{a}_+}$. Let $\alpha \in \Delta$ be such that $\alpha(v)$ is maximal. Then there are numbers $\tau_\beta \in [0, 1]$ for all simple roots $\beta \neq \alpha$ so that

$$(13) \quad V = \mathbb{R}^+ v = \{H \in \overline{\mathfrak{a}_+} : \beta(H) = \tau_\beta \alpha(H) \text{ for all } \beta \in \Delta - \{\alpha\}\}.$$

Let

$$V_\delta^+ = \{H \in \overline{\mathfrak{a}_+} : \tau_\beta \alpha(H) \leq \beta(H) \leq (\tau_\beta + \delta)\alpha(H) \text{ for all } \beta \in \Delta - \{\alpha\}\}$$

and let $Y_{V_\delta^+} = K \exp V_\delta^+ K \subset X$ (see Figure 5).

We will need to construct a contraction of $Y_{V_\delta^+}$ to Y_v . We begin by defining a map $\pi_V : V_\delta^+ \rightarrow V$ as follows: if V is given in the form (13), then $\pi_V(H)$ is the unique point in V such that $\alpha(\pi_V(H)) = \alpha(H)$. Then there exists a constant λ such that for $x, y \in V_\delta^+$,

$$(14) \quad d(\pi_V(x), \pi_V(y)) \leq \lambda d(x, y).$$

Now define π_V on $\exp V_\delta^+$ by $\pi_V(a) := \exp \pi_V(\log a)$, and extend π_V to all of $Y_{V_\delta^+}$ by $\pi_V(kaK) = k\pi_V(a)K$.

Lemma 6.5 (A contraction). *The map $\pi_V : Y_{V_\delta^+} \rightarrow Y_v$ is a contraction:*

$$d(\pi_V(x), \pi_V(y)) \leq \lambda d(x, y)$$

for all $x, y \in Y_{V_\delta^+}$.

Proof. Note that by construction, $\beta(\pi_V(x)) \leq \beta(x)$ for all $x \in Y_{V_\delta^+}$ and all $\beta \in \Sigma^+$. It is enough to show that for any tangent vector T , $\|\pi_V^*(T)\| \leq \lambda' \|T\|$, where $\|\cdot\|$ denotes the length in the Riemannian metric. But this follows immediately from (14), Proposition 2.3, and the observation that $\beta(\pi_V(x)) \leq \beta(x)$ for all $\beta \in \Sigma^+$. \square

6.3.2. More approximations.

Lemma 6.6. *For any vector v and any $\delta > 0$ there exists a one-dimensional linear subspace U so that $Y_v[\delta] \subset Y_{U_{2\lambda_1\delta}^+}$.*

Proof. Assume V is given in the form (13). Then there exists λ_1 so that

$$V[\delta] \cap \overline{\mathfrak{a}_+} \subset \{H \in \overline{\mathfrak{a}_+} : (\tau_\beta - \lambda_1\delta)\alpha(H) \leq \beta(H) \leq (\tau_\beta + \lambda_1\delta)\alpha(H), \\ \text{for all } \beta \in \Delta - \{\alpha\}\}.$$

Let $\tau'_\beta = \max(\tau_\beta - \lambda_1\delta, 0)$, and let

$$U = \{H \in \overline{\mathfrak{a}_+} : \beta(H) = \tau'_\beta\alpha(H) \text{ for all } \beta \in \Delta - \{\alpha\}\}.$$

Then $Y_v[\delta] \subset Y_{U_{2\lambda_1\delta}^+}$. Clearly we may ensure that $U = \mathbb{R}^+u$ for some unit vector u , and for every $\alpha \in \Delta$, either $\alpha(u) = 0$ or $\alpha(u) \succ \delta$. \square

For $x \in \phi^{-1}(Y_{V_\delta^+})$, define $\phi_V(x) = \pi_V \circ \phi(x)$.

Lemma 6.7 (ϕ_U satisfies Conditions I and II). *Let U, v, δ be such that $\epsilon \ll \delta \ll 1$, and $Y_v[2\delta] \subset Y_{U_{2\lambda_1\delta}^+}$. Then for any $x, y \in \phi^{-1}(Y_v[\delta])$ the following two facts hold:*

$$(15) \quad d(x, y) - O(\delta d_\rho(x, 0)) \prec d(\phi_U(x), \phi_U(y)) \prec d(x, y) + O(\delta d_\rho(x, 0)),$$

$$(16) \quad \frac{1}{4\mathcal{K}} d_\rho(x, 0) \leq d(\phi_U(x), e) \leq 4\mathcal{K} d_\rho(x, 0).$$

Proof. If $x, y \in \phi^{-1}(Y_v[\delta])$, then $\phi(x), \phi(y) \subset Y_v[2\delta] \subset Y_{U_{2\lambda_1\delta}^+}$. Note that on $Y_{U_{2\lambda_1\delta}^+}$,

$$(17) \quad d(x, \pi_U(x)) = O(\delta d(x, e)).$$

Hence

$$(18) \quad d(\phi(x), \phi_U(x)) = O(\delta d(\phi(x), e)) = O(\delta d_\rho(x, 0))$$

by Condition II. Applying the triangle inequality together with (18) and Condition I then gives (15).

To prove the second fact, we know from (18) and Condition II that, if $\epsilon \ll \delta$,

$$(19) \quad \frac{1}{2\mathcal{K}}d_\rho(x, 0) - O(\delta d_\rho(x, 0)) \leq d(\phi_U(x), e) \leq 2\mathcal{K}d_\rho(x, 0) + O(\delta d_\rho(x, 0)).$$

Now (16) follows from (19) if $\delta \ll 1$. □

6.3.3. Proof of 1-dimensional pinching. In this subsection we prove a “1-dimensional” version of Lemma 5.6. This is the base case of the inductive proof in §6.4.

The motivation for the definition of pinching function came from functions such as $d(g(x), e)$, where g is a reasonably efficient map of \mathbb{R}^n into a hyperbolic space. For example a quasi-isometric embedding $\phi : \mathbb{R}^2 \rightarrow \mathbb{H}^2 \times \mathbb{H}^2$ composed with projection onto an \mathbb{H}^2 factor, restricted to $\phi^{-1}(X^\delta)$.

For general X , the role of a hyperbolic plane factor is played by Y_v as in the Expanding Annulus Lemma (Lemma 6.1). The property of Y_v used is that “travel transverse to the radial direction in Y_v is exponentially expensive”, as given precisely by the Expanding Annulus Lemma.

The pinching functions we need to construct, however, use instead the *neighborhood* $Y_v[\delta]$ of Y_v in X . This neighborhood only introduces a linear error term $O(\delta r)$ in the functions involved, so proving that $d(\pi_U \circ \phi(x), e)$ is a pinching function on $\phi^{-1}(Y_v[\delta])$ follows easily 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 6.8 (Constructing a pinching function on $\phi^{-1}(Y_v[\delta])$). *If $\epsilon \ll \delta \ll 1$, then for any unit vector v in \mathfrak{a}_+ , the set $\phi^{-1}(Y_v[\delta])$ 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:*

$$(20) \quad \frac{1}{4\mathcal{K}}d_\rho(x, 0) \leq f(x) \leq 4\mathcal{K}d_\rho(x, 0) \quad \text{for all } x \in \phi^{-1}(Y_v[\delta]).$$

Proof. By Lemma 6.6 we can choose U so that $Y_v[2\delta] \subset Y_{U_{2\lambda_1\delta}^+}$. Let $\phi_U = \pi_U \circ \phi$. We claim that $f(x) = d(\phi_U(x), e)$ is the required pinching function. The equation (20) is the same as (16); this also shows that f is proper.

Now suppose $x, y \in \phi^{-1}(Y_v[\delta])$, $s \geq r_0$ and $s \leq f(x) \leq f(y) \leq \omega s$. Let $\gamma : [0, 1] \rightarrow \phi^{-1}(Y_v[\delta])$ 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 3.3. Property (1) of Lemma 3.3 gives $\hat{\gamma} \subset \phi \circ \gamma[\lambda\epsilon]$. Hence by part (b) of Lemma 3.2, $\hat{\gamma} \subset Y_v[2\delta]$, 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 $\epsilon \ll \delta \ll 1$, by Lemma 6.5,

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

where the last estimate follows from (17) and the triangle inequality. Thus

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

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

$$(21) \quad d(\pi_U(\hat{x}), e) \geq (1 - \lambda\delta)s,$$

for some constant λ .

We want to apply Lemma 6.1 to the path $\psi = \pi_U \circ \hat{\gamma}$. Let $r = (1 - \lambda\delta)s$. Then by (21), $d(\psi(t), e) \geq r$ for all t in $[0, 1]$, hence the first hypothesis of Lemma 6.1 holds.

Let $c_1 = 8\lambda\delta$, $\omega = 1 + (\lambda/2)\delta$, $a_0 = \phi_U(x)$, and $b_0 = \phi_U(y)$. Since $d(a_0, e), d(b_0, e) \in [s, \omega s]$, it is 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 $\delta \ll 1$. Property (2) of Lemma 3.3 and Lemma 6.5 imply that the second hypothesis of Lemma 6.1 holds with $c_2 = O(1/\epsilon^{n-1})$. Hence Lemma 6.1 gives a constant r_0 so that if $r \geq r_0$

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

Hence

$$\begin{aligned} d(x, y) &< c_1 s + O(\delta d_p(x, 0)) && \text{by (15)} \\ &< c_1 s + O(\delta d(\phi_U(x), e)) && \text{by (16)} \\ &< c_1 s + O(\delta s) && \text{since } d(\phi_U(x), e) \leq \omega s \leq 2s \\ &< \delta s. \end{aligned}$$

By the construction of $U = \mathbb{R}^+u$, for every $\alpha \in \Delta$ either $\alpha(u) = 0$ or $\alpha(u) > \delta$. Thus by the “remark on uniformity” following the statement of Lemma 6.1, r_0 depends only on ϵ, δ , and not on v . \square

Corollary 6.9 (The set $\phi^{-1}(Y_v[\delta])$ is pinched). *Suppose $\epsilon \ll \delta \ll 1$ and $\max(r, \rho)$ is sufficiently large (depending on δ, ϵ). Then for every subset $W \subset \phi^{-1}(Y_v[\delta])$ and all $p \geq 1$, the homomorphism $i_* : H_p(W \cup B(0, r)) \rightarrow H_p(W[\eta\delta] \cup B(0, r))$ induced by the inclusion is the zero map. Here $\eta = \eta(\mathcal{K}) > 1$ is independent of δ .*

Proof. 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 6.8 and Lemma 6.4, and the observation that by (20) and the definition of the pinching constant β , $\tilde{W} \subset W[\eta\delta]$ for $\eta = O(1)$. \square

6.4. Higher dimensional pinching and Mayer-Vietoris. We have a finite group W (the Weyl group) acting properly discontinuously on S^{n-1} . Let \mathcal{C} be a linear complex on $S^{n-1} \cap \overline{\mathfrak{a}_+}$, i.e. \mathcal{C} is a finite cell complex whose cells are intersections with S^{n-1} of linear subspaces. For a linear complex \mathcal{C} , we define $Y_{\mathcal{C}}$ to be $K \exp(\mathbb{R}^+\mathcal{C})K \subset X$.

We now wish to generalize the pinching phenomenon to a setting where the target space has flat directions, and so does not admit a pinching function like Y_v . Specifically we want to generalize the conclusion of the pinching lemma, with Y_v replaced by $Y_{\mathcal{C}}$.⁶

Although $Y_{\mathcal{C}}$ does not have an Expanding Annulus Lemma or any other hyperbolic properties, it is foliated by copies of Y_v . We use this together with Mayer-Vietoris to prove pinching.

For any linear complex \mathcal{D} , we define the *inner size* (resp. *outer size*) of \mathcal{D} to be the maximal (resp. minimal) number s so that every cell in \mathcal{D} contains (resp. is contained in) a ball of radius s . Note that every complex can be completed to a (W -invariant) triangulation of S^{n-1} with “bounded geometry”; i.e. there exists a $\lambda = \lambda(n) \gg 1$ such that every linear complex \mathcal{D} has a refinement \mathcal{D}' (which may be

⁶We note that this generalization is not necessary in the case $X = \mathbb{H}^2 \times \mathbb{H}^2$, and so is not covered in [EF].

completed to a W -invariant triangulation of S^{n-1}) or which the ratio of the outer size of \mathcal{D}' to the inner size of \mathcal{D}' is less than λ .

Now define the *feature size* of \mathcal{C} , denoted $\text{f.s.}(\mathcal{C})$, to be the maximum over $s \in \mathbb{R}^+$ such that \mathcal{C} can be completed to a W -invariant triangulation \mathcal{C}' of S^{n-1} with inner size at least s and outer size at most λs . This number is finite and nonzero for any linear complex.

We note the following trivial fact: there is a constant $\lambda_1 = \lambda_1(n)$ so that if \mathcal{C} is a linear complex and $\mu \leq \text{f.s.}(\mathcal{C})$, then there exists a refinement \mathcal{C}'' of \mathcal{C} such that $\lambda_1^{-1}\mu \leq \text{f.s.}(\mathcal{C}'') \leq \lambda_1\mu$.

We also note here that ‘bracketing with $[\]$ ’ commutes with translating by K ; so, for example, $Y_{\mathcal{C}}[\delta] = K((\exp \mathbb{R}^+ \mathcal{C})[\delta])K$.

Lemma 6.10 (The sets $\phi^{-1}(Y_{\mathcal{C}}[\delta])$ are pinched for linear complexes \mathcal{C}). *For all $1 \leq q \leq n - 1$, there exist constants $\beta_q = \beta_q(\mathcal{K}), \delta_q = \delta_q(\mathcal{K})$ such that for every $(q - 1)$ -dimensional linear complex \mathcal{C} in $S^{n-1} \cap \overline{\mathfrak{a}_+}$, for all $\epsilon \ll \delta < \delta_q \text{f.s.}(\mathcal{C})$, and for every $W \subset \phi^{-1}(Y_{\mathcal{C}}[\delta])$, the map $i_* : H_p(W \cup F) \rightarrow H_p(W[\beta_q\delta] \cup F)$ is zero for all $p \geq q$, and every ball F centered at 0, with $d_p(\partial F, 0)$ sufficiently large (depending on ϵ, δ).*

Proof. For simplicity we provide a proof ignoring the ball F . The proof which includes F is quite similar, as in the proof of Lemma 6.4.

We use induction on q . Suppose $q = 1$. Then \mathcal{C} is a 0-complex, i.e. a collection of points $\{v_i\}$. Its feature size is at most $\lambda_2 \min_{i \neq j} d(v_i, v_j)$, where $\lambda_2 = \lambda_2(n)$. Hence if $\delta < 1/(3\lambda_2)\text{f.s.}(\mathcal{C})$, the sets $Y_{L_i}[\delta]$ are disjoint. Hence the case $q = 1$ is Corollary 6.9, with $\delta_1 = 1/(3\lambda_2)$.

Suppose $q > 1$ and that the lemma holds for $q - 1$ with constants $\beta_{q-1}, \delta_{q-1}$. Without loss of generality, $\beta_{q-1} \geq 1$, and $\delta_{q-1} \ll 1$. We choose $1 \ll \eta_1 \ll \eta_2 \ll \eta_3 \ll \eta_4$, with sufficiently large implied constants (depending only on \mathcal{K} and n). Let \mathcal{C} be a linear complex of dimension $q - 1$. Choose $\delta_q = \delta_{q-1}/(\eta_4\lambda_1\beta_{q-1})$, and suppose $\delta \leq \delta_q \text{f.s.}(\mathcal{C})$. By the ‘trivial fact’ above with $\mu = (\delta/\delta_q)(\eta_3/\eta_4) = \delta\eta_3\lambda_1\beta_{q-1}/\delta_{q-1}$, we may refine \mathcal{C} so that $\text{f.s.}(\mathcal{C}) \in [\lambda_1^{-1}\mu, \lambda_1\mu]$, i.e.

$$(22) \quad \frac{\eta_3\beta_{q-1}}{\delta_{q-1}}\delta \leq \text{f.s.}(\mathcal{C}) \leq \frac{\eta_3\lambda_1^2\beta_{q-1}}{\delta_{q-1}}\delta.$$

We choose

$$\beta_q = \max(4\eta_3\lambda_1^3\beta_{q-1}^2, (3\eta_3/\delta_{q-1} + 8\eta_1)\beta_{q-1}\beta_1).$$

We will show that for $W \subset Y_{\mathcal{C}}[\delta]$, the map $i_* : H_p(W) \rightarrow H_p(W[\beta_q\delta])$ induced by inclusion is zero.

Let \mathcal{S} be the $(q - 2)$ -skeleton of \mathcal{C} . Let $\hat{B} = Y_{\mathcal{S}}[\eta_2\beta_{q-1}\delta]$, and let $\hat{A} = (Y_{\mathcal{C}}[\delta] \setminus \hat{B})[\delta]$. So $Y_{\mathcal{C}}[\delta] \subset \hat{A} \cup \hat{B}$. Let $A = \phi^{-1}(\hat{A}) \cap W$, $B = \phi^{-1}(\hat{B}) \cap W$. Then $W = A \cup B$. Note that $W[\eta] = A[\eta] \cup B[\eta]$ for all $\eta > 0$. See Figure 6.

For $p \geq q$, consider the commutative diagram

$$\begin{array}{ccccc} H_p(A) \oplus H_p(B) & \longrightarrow & H_p(W) & \longrightarrow & H_{p-1}(A \cap B) \\ \downarrow & & j_* \downarrow & & k_* \downarrow \\ H_p(A[4\beta_{q-1}\delta]) \oplus H_p(B[4\beta_{q-1}\delta]) & \xrightarrow{\Psi} & H_p(W[4\beta_{q-1}\delta]) & \longrightarrow & H_{p-1}(A[4\beta_{q-1}\delta] \cap B[4\beta_{q-1}\delta]) \\ \downarrow \iota_* & & m_* \downarrow & & n_* \downarrow \\ H_p(A[\beta_q\delta]) \oplus H_p(B[\beta_q\delta]) & \xrightarrow{\Psi'} & H_p(W[\beta_q\delta]) & \longrightarrow & H_{p-1}(A[\beta_q\delta] \cap B[\beta_q\delta]) \end{array}$$

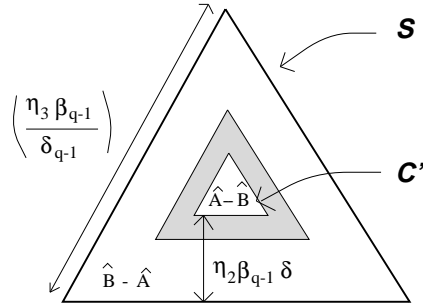


FIGURE 6. How Mayer-Vietoris is applied. The entire picture is a single cell of \mathcal{C} . The shaded region is $\hat{A} \cap \hat{B}$, which has width δ . We emphasize that the size of any cell is much bigger than the thickness of \hat{B} , which in turn is much bigger than the thickness of $\hat{A} \cap \hat{B}$. In other words $(\eta_3 \beta_{q-1}) / (\delta_{q-1}) \gg \eta_2 \beta_{q-1} \delta \gg \delta$. One reason these sizes are chosen this way is so that the sets \hat{A} in different cells remain disjoint after thickening. This allows us to apply the inductive hypothesis.

where the rows are Mayer-Vietoris exact sequences, and the vertical maps are inclusions. From the definition of \hat{A} and \hat{B} and the continuity of ϕ , it is easy to check that the hypotheses of Mayer-Vietoris are satisfied.

We claim that the map k_* is zero. Let C' denote the intersection of \mathcal{C} with $\partial \hat{B}$. Since $\partial \hat{B}$ is defined by linear equations, C' is also a linear complex of dimension $q-2$. C' has one component C'_i in each cell of \mathcal{C} . By construction, $A \cap B \subset \phi^{-1}(\hat{A} \cap \hat{B})$, and $\hat{A} \cap \hat{B}$ is a disjoint union of sets C_i where each C_i is contained in $Y_{C'_i}[4\delta]$. Since $\delta_{q-1} \ll 1$ and $\eta_2 \ll \eta_3$, the distance from C'_i to \mathcal{S} equals $\eta_2 \beta_{q-1} \delta \ll \text{f.s.}(\mathcal{C})$ by (22). Hence $\text{f.s.}(C'_i) \geq (1/2)\text{f.s.}(\mathcal{C})$. Thus, by (22), $\delta_{q-1} \text{f.s.}(C'_i) \geq (1/2)\eta_3 \beta_{q-1} \delta \geq 4\delta$, since $\beta_{q-1} > 1$ and $\eta_3 \gg 1$. Thus we can apply the inductive hypothesis to each C_i to conclude that $k_* = 0$.

We now claim that the map $l_* = l_*^{(1)} \oplus l_*^{(2)}$ is also zero. We first consider $l_*^{(2)} : H_p(B[4\beta_{q-1}\delta]) \rightarrow H_p(B[\beta_q\delta])$. Note that

$$B[4\beta_{q-1}\delta] \subset \phi^{-1}(Y_{\mathcal{S}}[\eta_2\beta_{q-1}\delta])[4\beta_{q-1}\delta] \subset \phi^{-1}(Y_{\mathcal{S}}[\eta_3\beta_{q-1}\delta])$$

with the last inclusion by part (c) of Lemma 3.2 and $\eta_2 \ll \eta_3$. Also $\eta_3\beta_{q-1}\delta \leq \delta_{q-1} \text{f.s.}(\mathcal{C})$ by (22). Thus, by the induction hypothesis, the map $H_p(B[4\beta_{q-1}\delta]) \rightarrow H_p(B[4\beta_{q-1}\delta][\eta_3\beta_{q-1}^2\delta])$ is the zero map. But

$$B[4\beta_{q-1}\delta][\eta_3\beta_{q-1}^2\delta] \subset B[4\eta_3\beta_{q-1}^2\delta] \subset B[\beta_q\delta]$$

by Lemma 3.2 part (b) and $\beta_q \geq 4\eta_3\beta_{q-1}^2$, so that $l_*^{(2)}$ is indeed the zero map.

We now consider $l_*^{(1)} : H_p(A[4\beta_{q-1}\delta]) \rightarrow H_p(A[\beta_q\delta])$. Note that $A[4\beta_{q-1}\delta] \subset \phi^{-1}(\hat{A})[4\beta_{q-1}\delta] \subset \phi^{-1}(\hat{A}[\eta_1\beta_{q-1}\delta])$, since $\eta_1 \gg 1$ by part (c) of Lemma 3.2. Since $\eta_1 \ll \eta_2$, the set $\hat{A}[\eta_1\beta_{q-1}\delta]$ is a disjoint union of sets \hat{A}_i , where each \hat{A}_i is contained in some $D_i[\eta_1\beta_{q-1}\delta]$, where D_i is a single cell of \mathcal{C} . Then $A[4\beta_{q-1}\delta]$ is a disjoint union of sets $A_i = \phi^{-1}(\hat{A}_i) \cap A[4\beta_{q-1}\delta]$. If v_i is the barycenter of D_i , then $D_i \subset$

$Y_{\{v_i\}}[\lambda_1 \text{f.s.}(\mathcal{C})]$ and

$$\begin{aligned} A_i \subset \phi^{-1} (Y_{\{v_i\}}[\lambda_1 \text{f.s.}(\mathcal{C})][\eta_1 \beta_{q-1} \delta]) &\subset \phi^{-1} (Y_{\{v_i\}}[2\lambda_1 \text{f.s.}(\mathcal{C}) + \eta_1 \beta_{q-1} \delta]) \\ &\subset \phi^{-1} (Y_{\{v_i\}}[(2\eta_3 \lambda_1^3 / \delta_{q-1} + \eta_1) \beta_{q-1} \delta]) \end{aligned}$$

with the second inclusion by part (b) of Lemma 3.2 and the third inclusion by formula (22). Note that if $\delta < \delta_q \text{f.s.}(\mathcal{C})$, then $(2\eta_3 \lambda_1^3 / \delta_{q-1} + \eta_1) \beta_{q-1} \delta < \delta_1 \text{f.s.}(\mathcal{C})$, because $\eta_4 \gg \eta_3$. Hence by the induction hypothesis applied in turn to each 0-complex $\{v_i\}$, and by the disjointness of the A_i 's, the map $H_p(A[4\beta_{q-1} \delta]) \rightarrow H_p(A[(3\eta_3 \lambda_1^3 / \delta_{q-1} + \eta_1) \beta_{q-1} \delta])$ is the zero map. Since $(3\eta_3 \lambda_1^3 / \delta_{q-1} + \eta_1) \beta_{q-1} \delta < \beta_q$, the map $l_*^{(1)} : H_p(A[4\beta_{q-1} \delta]) \rightarrow H_p(A[\beta_q \delta])$ is the zero map, hence l_* is the zero map.

The rest of the argument is a diagram chase: if $x \in H_p(W)$, then $j_*(x) = \Psi(y)$ for some y since $k_* = 0$. But then $m_*(j_*(x)) = \Psi'(l_*(y)) = 0$ since $l_* = 0$. Hence $i_* = m_* \circ j_* = 0$, which is what we needed to show. \square

Proof of the Main Lemma (Lemma 5.6). Recall that $(X^\delta)^c = \bigcup_{\alpha \in \Delta} Y_{\Delta \setminus \{\alpha\}}[\delta]$. Apply Lemma 6.10 to the “wall complex” whose cells are the intersections of the unit sphere with subspaces of the form $L_\sigma := \{H \in \overline{\mathfrak{a}}_+ : \alpha(H) = 0 \text{ for all } \alpha \notin \sigma\}$ (so that $Y_\sigma = K \exp L_\sigma K$). The dimension of this complex is $n - 2$. \square

7. PROOF OF THEOREM 1.1

7.1. Flats and the boundary. The hyperplanes $\beta = 0$ for $\beta \in \Sigma$ divide \mathfrak{a} into finitely many regions called *Weyl chambers*. The *Weyl group*, denoted by W , is defined to be $N_K(\mathfrak{a})/M$, where $N_K(\mathfrak{a})$ is the normalizer of \mathfrak{a} in K under the adjoint action. As above, M is the intersection with K of the centralizer of \mathfrak{a} . The Weyl group is finite, and acts transitively on the Weyl chambers.

We first write the action on the boundary of X explicitly. Let $\mathfrak{n} = \bigoplus_{\alpha \in \Sigma^+} \mathfrak{g}^\alpha$. Then \mathfrak{n} is a nilpotent Lie subalgebra of \mathfrak{g} , and its exponential is a nilpotent Lie group N . From the definition it is clear that MA normalizes N , hence $B = MAN$ is a subgroup of G , called a *Borel subgroup*. In view of the Iwasawa decomposition $G = KAN$ and the fact that M is a normal subgroup of B , we have a diffeomorphism of K/M onto G/B given by

$$kM \rightarrow kB.$$

The relation $K/M = G/B$ shows in particular that G acts as a transformation group on the boundary \tilde{X} ; this action agrees with the action defined by (2). The action in the K/M picture is given by

$$g \cdot kM = \mathfrak{K}(gk)M$$

if for $x \in G$, $\mathfrak{K}(x) \in K$ is given by $x \in \mathfrak{K}(x)AN$. Thus the stabilizer of the point $o = eM$ in K/M is B .

The Bruhat decomposition decomposes G into double cosets of B . In fact

$$G = \bigsqcup_{w \in W} BwB$$

where the union is disjoint. For exactly one $w_0 \in W$ the set Bw_0B is open; w_0 is called the *longest element* of the Weyl group. The set Bw_0B/B is an open and dense subset of the boundary G/B , on which B acts transitively. The complement of Bw_0B in G is called the *singular set*; the intersection of the singular set with K

is the *singular subset*, denoted by \mathcal{S} . The complement of \mathcal{S} in K is open and dense, and $M\mathcal{S}M = \mathcal{S}$.

Now for each flat F in X we can associate a finite set $F(\infty) \subset \hat{X}$ of (equivalence classes) of Weyl chambers in F . We call $F(\infty)$ the limit set of F . For every pair of points (β, γ) of \hat{X} , there exists a flat F so that $\beta, \gamma \in F(\infty)$. If (under the identification of \hat{X} with K/M) $\beta^{-1}\gamma \notin \mathcal{S}$, then F is unique; we denote it by $F(\beta, \gamma)$.

Definition 7.1. We define a function $\Theta'_e(\beta, \gamma)$ to be $d_K(\beta, \gamma\mathcal{S})$, and then extend it to all $p \in X$, $\beta, \gamma \in \hat{X}$ so that for any $g \in G$, $\Theta'_{ge}(g\beta, g\gamma) = \Theta'_e(\beta, \gamma)$.

The following geometric fact generalizes the fact that, in the hyperbolic plane, if the angle $\angle ABC$ of a triangle is bounded below, then the distance from B to \overline{AC} is bounded above. The proof is given in the Appendix.

Lemma 7.2. *For every $\epsilon > 0$ there exists a number $R > 0$ such that for $p \in X$ and $\beta, \gamma \in \hat{X}$, if $\Theta'_p(\beta, \gamma) > \epsilon$, then $d(p, F(\beta, \gamma)) < R$.*

7.2. Proof of Theorem 1.1. We are given a (κ, C) -quasi-isometric embedding $\phi: \mathbb{R}^n \rightarrow X$. The idea of the proof is roughly as follows.

First off, suppose $z \in X$ is a point which sees two far-away points $\phi(x_1), \phi(x_2) \in X^\delta(z)$ at a visual angle bounded below (actually we need a lower bound on $\Theta'_z(\phi(x_1), \phi(x_2))$). By Lemma 5.8, for each $i = 1, 2$, there is a path γ_i from x_i to infinity in $\phi^{-1}(X^\delta(z))$ with $\phi \circ \gamma_i$ limiting to a limit point β_i . If we knew that the paths $\phi \circ \gamma_i$ stayed outside a large ball containing z , then by Lemma 4.3 (no shifting lemma) $\beta_1 \neq \beta_2$. Since the visual distance $\Theta'_z(\phi(x_1), \phi(x_2))$ is bounded below, the visual distance $\Theta'_z(\beta_1, \beta_2)$ is bounded below (again by no shifting). Then by an analog of thin triangles (Lemma 7.2), z must be a bounded distance from the flat $F(\beta_1, \beta_2)$.

Why do the paths $\phi \circ \gamma_i$ stay outside a large ball containing z ? There are two cases: either z is far from $\phi(\mathbb{R}^n)$ (and we are obviously done), or $\phi(\mathbb{R}^n)$ comes close to z at a point $p \in \phi(\mathbb{R}^n)$, in which case we can (Lemma 5.8) find paths γ_i as above which stay far from $\phi^{-1}(p)$. Then $\phi \circ \gamma_i$ 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 $x_1 \in \mathbb{R}^n$, we can use Lemma 5.7 to find $x_2 \in \mathbb{R}^n$ 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 7.3. *Suppose $0, x_1, x_2 \in \mathbb{R}^n$, $e \in X$ and $\epsilon, \delta, \mathcal{K}, \lambda_0, r > 0$ are such that*

- (a) ϕ satisfies Condition I and Condition II outside some ball B centered at 0.
- (b) $\epsilon \ll \delta \ll 1$ (the implied constant depending on \mathcal{K}).
- (c) $x_1, x_2 \notin B$.
- (d) $x_1, x_2 \in \phi^{-1}(X^\delta)$.
- (e) $\Theta'_e(\phi(x_1), \phi(x_2)) \geq \lambda_0 > 0$.

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

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

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

Proof. Let $\delta' = \delta/\eta_1$ where η_1 is as in Proposition 5.9. Then by Proposition 5.9 there exist δ' -limit points β_1, β_2 so that for $i = 1, 2$,

$$\Theta_e(\phi(x_i), \beta_i) \leq e^{-\nu_2 d_\rho(x_i, 0)}.$$

As r increases, so does $d_\rho(x_i, 0)$ by Condition II. Then for sufficiently large r , $\Theta_e(\phi(x_i), \beta_i) \leq \lambda_0/4$, so by the triangle inequality, $\Theta'_e(\beta_1, \beta_2) > \lambda_0/2$. Then by Lemma 7.2, $d(e, F(\beta_1, \beta_2))$ is bounded by some constant Λ depending only on λ_0 and X . □

Choosing constants. Choose ν smaller than a large negative power of κ ($\nu = 1/(32\kappa^3)$, where $\kappa > 1$). Choose ϵ, δ so that $\epsilon < 1/(12\kappa)$, Lemma 5.7 is satisfied with $\mathcal{K} = \kappa$, and so that condition ?? of Proposition 7.3 is satisfied both with $\mathcal{K} = \kappa$ and $\mathcal{K} = \kappa^2/\nu$. Note that ϵ and δ depend only on κ . Now choose R so large that condition ?? of Proposition 7.3 is satisfied with $r = R/(4\mathcal{K})$, with ϵ, δ as chosen, and both with $\mathcal{K} = \kappa$ and with $\mathcal{K} = \kappa^2/\nu$. Now let δ denote one half of the smaller of the two constants δ' produced by the two applications ($\mathcal{K} = \kappa$ and $\mathcal{K} = \kappa^2/\nu$) of Proposition 7.3.

Method of proof. We will show the following: for every point $x_1 \in \mathbb{R}^n$, $\phi(x_1)$ is within a bounded distance of a flat of the form $F(\beta_1, \beta_2)$ where β_1, β_2 belong to the δ -limit set $\mathcal{L}_\delta(\phi)$, and $\beta_2^{-1}\beta_1 \notin \mathcal{S}$. This is enough since the cardinality of the limit set is bounded by Proposition 5.2 (limit set is finite), and the limit set is independent of the point x_1 by Lemma 5.4 and the choice of δ .

Separation into cases. Let $x_1 \in \mathbb{R}^n$ be an arbitrary point. Since ϕ is a quasi-isometry, we may assume, without loss of generality, that⁷

$$(23) \quad \phi \text{ satisfies Condition I with base point } x_1 \text{ and some } \rho_1 = \rho_1(\epsilon, C).$$

Choose $R > 16\kappa^2\rho_1$. By Lemma 5.7 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 7.3 holds. Since $\phi(x_1)$ and $\phi(x_2)$ are points in opposite Weyl chambers of a flat passing through $z = e$, $\Theta(\phi(x_1)) = \Theta(\phi(x_2))w_0$, where w_0 is the longest element of the Weyl group W . Hence

$$\Theta'_e(\phi(x_1), \phi(x_2)) = d_K(w_0, \mathcal{S}) = \lambda_0$$

depending only on X . Thus condition (e) of Proposition 7.3 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 7.3. There are two cases to consider:

⁷We will use only this more general condition in the rest of the proof. This is done in order to simplify the proof of Theorem 8.1.

Case 1. $d(\phi(\mathbb{R}^n), z) > \nu R$. In this case choose coordinates so that $x_1 = 0 \in \mathbb{R}^n$. We claim that the map ϕ satisfies Condition I and Condition II on all of \mathbb{R}^n with $\rho = \kappa^2 R$ and $\mathcal{K} = \kappa^2/\nu$. Condition I holds because $\rho > R > \rho_1$ and $\mathcal{K} > \kappa$. The lower bound in Condition II follows like this: For any $u \in \mathbb{R}^n$,

$$\begin{aligned} d(\phi(u), e) &\geq \max(d(\phi(u), \phi(0)) - d(\phi(0), e), \nu R) \\ &\geq \max\left(\frac{1}{\kappa}d(u, 0) - (\kappa + \epsilon)R/2 - \epsilon d_{\rho_1}(u, 0), \nu R\right) \\ &\geq \max\left(\frac{1}{\kappa}d(u, 0) - 2\kappa R/3 - \epsilon d_{\rho_1}(u, 0), \nu R\right) \\ &\geq \frac{\nu}{\kappa^2} \max\left\{\frac{\kappa}{\nu}((d(u, 0) - 2\kappa^2 R/3 - \epsilon \kappa d_{\rho_1}(u, 0)), \kappa^2 R)\right\} \\ &\geq \frac{\nu}{\kappa^2} \max\{d(u, 0), \kappa^2 R\} \\ &= \frac{1}{\mathcal{K}}d_\rho(u, 0) \end{aligned}$$

where the second line follows from

$$d(\phi(0), e) = \frac{1}{2}d(\phi(x_1), \phi(x_2)) \leq \frac{1}{2}(\kappa R + \epsilon d_{\rho_1}(x_1, x_2)) \leq \frac{1}{2}(\kappa + \epsilon)R,$$

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 $\frac{\kappa}{\nu} \cdot (\frac{1}{3} - \kappa\epsilon)d(u, 0) > d(u, 0)$ since $\kappa > 4\nu$, $\epsilon\kappa < 1/12$ and $\kappa^2 R = \rho > \rho_1$.

The upper bound in Condition II follows from

$$\begin{aligned} (24) \quad d(\phi(u), e) &\leq d(\phi(u), \phi(0)) + d(\phi(0), e) \leq \kappa d(u, 0) + \frac{(\kappa + \epsilon)R}{2} + \epsilon d_{\rho_1}(u, 0) \\ &\leq \frac{\kappa^2}{\nu} \max(d(u, 0), \kappa^2 R) = \mathcal{K}d_\rho(u, 0) \end{aligned}$$

since $\nu < 1/2$ and $\kappa > 1$.

Thus condition (a) of Proposition 7.3 is verified in Case 1. Condition (c) of Proposition 7.3 is vacuous since B is empty.

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

Choose coordinates on \mathbb{R}^n so that $\phi^{-1}(p)$ is the origin 0. We claim that ϕ satisfies the conditions of Proposition 7.3 with $\mathcal{K} = \kappa$, $\rho = 0$ and $B = B(0, r')$, where $r' = \frac{R}{8\kappa^2}$.

We first check condition (c) of Proposition 7.3. It follows from (23) that

$$\begin{aligned} d(x_1, 0) &\leq \kappa d(\phi(x_1), p) + \kappa \epsilon d_{\rho_1}(x_1, 0) \\ &\leq \kappa d(\phi(x_1), e) + \kappa d(p, e) + \kappa \epsilon d_{\rho_1}(x_1, 0) \\ &\leq \frac{1}{2}\kappa d(\phi(x_1), \phi(x_2)) + \kappa \nu R + \kappa \epsilon d_{\rho_1}(x_1, 0) \\ &\leq \frac{\kappa}{2}(\kappa + \epsilon)R + \kappa \nu R + \kappa \epsilon d_{\rho_1}(x_1, 0) \leq \kappa^2 R + \kappa \epsilon d_{\rho_1}(x_1, 0). \end{aligned}$$

Also,

$$\begin{aligned} d(x_1, 0) &\geq \frac{1}{\kappa}d(\phi(x_1), p) - \frac{1}{\kappa}\epsilon d_{\rho_1}(x_1, 0) \\ &\geq \frac{1}{\kappa}d(\phi(x_1), e) - \frac{1}{\kappa}d(p, e) - \frac{1}{\kappa}\epsilon d_{\rho_1}(x_1, 0) \\ &\geq \frac{1}{2\kappa^2}(\kappa - \epsilon)R - \frac{1}{\kappa}\nu R - \frac{1}{\kappa}\epsilon d_{\rho_1}(x_1, 0) \\ &\geq \frac{1}{4\kappa^2}R - \frac{\epsilon}{\kappa}d_{\rho_1}(x_1, 0) \end{aligned}$$

where the first line used (23) and to go from the second line to the third we used $d(\phi(x_1), e) \geq ((1/\kappa) - \epsilon)R$ and $d(p, e) \leq \nu R$. Hence

$$(25) \quad \frac{1}{8\kappa^2}R \leq d(x_1, 0) \leq 2\kappa^2R.$$

Analogously

$$\begin{aligned} d(x_2, 0) &\geq \frac{1}{\kappa}d(\phi(x_2), p) - \frac{\epsilon}{\kappa}d_{\rho_1}(x_1, x_2) - \frac{\epsilon}{\kappa}d_{\rho}(x_1, 0) \\ &\geq \frac{1}{4\kappa^2}R - \frac{\epsilon}{\kappa}d_{\rho_1}(x_1, 0) \geq \frac{1}{8\kappa^2}R. \end{aligned}$$

Therefore $x_1, x_2 \notin B$, so condition (c) of Proposition 7.3 is satisfied.

We now check condition (a) of Proposition 7.3. We have by (23)

$$d(\phi(u), \phi(u')) \leq \kappa d(u, u') + \epsilon d(u, x_1) + \epsilon d(u', x_1).$$

If $d(u, 0) \geq 2\kappa^2R$, then

$$d(u, 0) \geq 2\kappa^2R \geq d(x_1, 0).$$

If $d(u, 0) \leq 2\kappa^2R$, then

$$d(u, x_1) \leq 4\kappa^2R \leq 32\kappa^4d(u, 0)$$

since $d(u, 0) > \frac{1}{8\kappa^2}R$. Therefore $d_{\rho_1}(u, x_1) \leq 32\kappa^4d(u, 0)$ and $\phi(x)$ satisfies Condition I outside B with $\mathcal{K} = \kappa$, $\rho = 0$ and ϵ replaced by $32\kappa^4\epsilon = \epsilon'$.

Finally we check Condition II.

$$\begin{aligned} d(\phi(u), e) &\geq d(\phi(u), p) - d(p, e) \\ &\geq \frac{1}{\mathcal{K}}d(u, 0) - \epsilon d(u, 0) - \nu R \\ (26) \quad &\geq \frac{1}{2\mathcal{K}}d(u, 0) \end{aligned}$$

since $d(u, 0) > \frac{1}{8\kappa^2}R$ and $\nu < \frac{1}{32\kappa^3}$. Analogously

$$d(\phi(u), p) \leq \mathcal{K}d(u, 0) + \epsilon d(u, 0) \leq 2\mathcal{K}d(u, 0).$$

Note that for the simplicity of notation we write in (26) ϵ instead of $\epsilon' = 32\kappa^4\epsilon$. Thus condition (a) of Proposition 7.3 is satisfied.

Completing the proof. By Proposition 7.3, $d(e, F(\beta, \gamma)) \leq \Lambda$. 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

Corollary 7.4 (Single flat). *Suppose $\phi : \mathbb{R}^n \rightarrow X$ is a (κ, C) -quasi-isometric embedding, and for some sufficiently small $\delta > 0$ (depending on κ), $\mathcal{L}_\delta(\phi) = F(\infty)$ for some flat $F \subset X$. Then the image of ϕ is in an N -neighborhood of F , where N depends only on (κ, C) and X .*

Proof. This is clear from the proof since in this case, the union of flats constructed coincides with F . \square

8. CONSEQUENCES

8.1. Quasi-flats with holes. Suppose $\Omega \subset \mathbb{R}^n$. Let $\Omega_{\epsilon, R}$ be the set of all $x \in \Omega$ such that for any $y \in \mathbb{R}^n \setminus B(x, R)$ 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 for some ϵ and R , $\Omega_{\epsilon, R}$ is nonempty, 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 8.1 (Quasi-flats with holes). *Suppose ϕ is a (κ, C) -quasi-isometric embedding from $\Omega \subset \mathbb{R}^n$ into X . Then there exist constants $F = F(\kappa)$ and $\epsilon_0 = \epsilon_0(\kappa)$ such that if $\epsilon < \epsilon_0$, $R \gg 1$ and $\Omega_{\epsilon, R}$ is not empty, then $\phi(\Omega_{\epsilon, R})$ lies in the N -neighborhood of F flats in X , where N depends on κ, C , and R .*

This theorem is used in [E] to deduce quasi-isometric rigidity for nonuniform lattices in semisimple groups without rank one factors.

Proof. Suppose $x \in \Omega_{\epsilon, R}$. We build a grid as in §3.2, 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 also consider the ball of radius R as a cell, with the point x as the grid point.

We can now use the “connecting the dots” argument (see §3) to construct a continuous function ϕ_x defined on all of \mathbb{R}^n which agrees with ϕ on every y_C . We may also ensure that for all $y \in \Omega - B(x, R)$,

$$(27) \quad d(\phi(y), \phi_x(y)) \leq c_n \kappa \epsilon d(y, x)$$

where c_n depends only on the dimension.

Choose ϵ_0 to be $1/(2c_n \kappa)$ times the ϵ chosen in the proof of Theorem 1.1. The function ϕ_x satisfies Condition I with base point x_1 because of (27) combined with the fact that ϕ does. We then repeat the proof of Theorem 1.1, except that we immediately replace ϕ by ϕ_x and choose $x_1 = x$. Thus the proof of Theorem 1.1 shows that $\phi(x)$ is within a bounded distance from one of the flats $F(\beta, \gamma)$ where β and γ belong to the limit set $\mathcal{L}_\delta(\phi_x)$. We must show that this limit set is independent of $x \in \Omega_{\epsilon, R}$.

Let \bar{x} be another point in $\Omega_{\epsilon, R}$, and let $\phi_{\bar{x}}$ be a continuous function constructed analogously to ϕ_x . We have

$$(28) \quad d(\phi_x(y), \phi_{\bar{x}}(y)) = O(\epsilon d(y, x) + \epsilon d(y, \bar{x})), \quad y \in \mathbb{R}^n.$$

By Lemma 5.4 $\mathcal{L}_{\bar{\delta}}(\phi_x)$ is independent of the choice of basepoint (up to replacing $\bar{\delta}$ by $\bar{\delta}/2$). Therefore by choosing \bar{x} as basepoint for ϕ_x and using (28) and Lemma 5.5 we get $\mathcal{L}_{\bar{\delta}}(\phi_x) \subset \mathcal{L}_{\bar{\delta}/4}(\phi_{\bar{x}})$. The opposite containment $\mathcal{L}_{\bar{\delta}}(\phi_{\bar{x}}) \subset \mathcal{L}_{\bar{\delta}/4}(\phi_x)$ is proved identically. Therefore the limit set is independent of the base point x , up to replacing $\bar{\delta}$ by $\bar{\delta}/4$. \square

8.2. More on quasi-flats. The following lemma is a coarse version of invariance of domain. The proof is a variation of an argument due to Geoff Mess.⁸ See also [FS, Corollary 5.3].

Lemma 8.2 (Local packing). *Let $f : B \rightarrow D$ be a continuous map which is a (κ, C) -quasi-isometric embedding from the ball $B = B(x, r_1) \subset \mathbb{R}^n$ into a ball $D \subset \mathbb{R}^n$ around $f(x)$. Then for any $N > (2\kappa + C)$, the image $f(B)$ contains the ball $B(f(x), N)$ if r_1 is sufficiently large (depending on N). We emphasize that r_1 does not depend on κ or f .*

Note. A variant of the statement of Lemma 8.2 is true without the assumption that f is continuous, as one can see by using the “connect-the-dots” construction (§3).

Proof. Let $C = B(f(x), N) \subset D$. Let $U \subset B$ be a ball of radius $r_2 < r_1$ centered at x . Choose r_1 and r_2 so large that $f(B - U) \cap C$ is empty. Hence we have a continuous map of pairs

$$f : (B, B - U) \rightarrow (f(B), f(B) \cap C^c).$$

We define a continuous map $g : f(B) \rightarrow B$ by choosing a sufficiently sparse net of points in $f(B)$, and sending a point y of this net to any point in $f^{-1}(y)$, and extending the map to all of $f(B)$ by connecting-the-dots. Since $N > 2\kappa + C$, we get a continuous map of pairs

$$g : (f(B), f(B) \cap C^c) \rightarrow (B, B - \{x\}) \rightarrow (B, B - U)$$

where the second map is the obvious deformation retraction along rays emanating from x . The map $g \circ f$ is pair-preserving homotopic to the identity, so the induced map

$$f_* : H_n(B, B - U) \rightarrow H_n(f(B), f(B) \cap C^c)$$

is injective (here we use homology with \mathbb{Z} coefficients). Clearly $H_n(B, B - U) \approx \mathbb{Z}$, so that $H_n(f(B), f(B) \cap C^c) \neq 0$. Now suppose the proposition were false, so that there exists $y \in C$ with $y \notin f(B)$. Since we are in the top dimension, it follows easily from the definitions that the inclusion of pairs $H_n(f(B), f(B) \cap C^c) \rightarrow H_n(\mathbb{R}^n - \{y\}, C^c)$ is an injection, so the second group is nonzero. This is clearly a contradiction. \square

We note that Lemma 8.2 implies that any quasi-isometric embedding of \mathbb{R}^n (with a bounded geometry condition on the metric) into itself is in fact a quasi-isometry, in particular some neighborhood of the image is all of \mathbb{R}^n .

If ϕ and ϕ' are two maps, we write $d(\phi, \phi')$ for $\sup_x d(\phi(x), \phi'(x))$.

Lemma 8.3 ($\phi(\mathbb{R}^n)$ is close to cone over limit set). *Let $\phi : \mathbb{R}^n \rightarrow X$ be a (κ, C) -quasi-isometric embedding.*

- For δ sufficiently small the limit set $\mathcal{L}_\delta(\phi)$ does not depend on δ (we denote it by $\mathcal{L}(\phi)$).

⁸Thanks to Shmuel Weinberger for telling us about this argument.

- There exists a (κ, C') -quasi-isometric embedding $\phi' : \mathbb{R}^n \rightarrow X$ so that $d(\phi, \phi') < \infty$ and

$$\phi'(\mathbb{R}^n) \subset V \quad \text{and} \quad V \subset \text{Nbhd}_R(\phi'(\mathbb{R}^n))$$

where $R > 0$ and V is the union of Weyl chambers passing through the origin:

$$V = \bigcup_{k \in \mathcal{L}(\phi)} kA_+K \subset X.$$

The functions ϕ and ϕ' have the same limit set.

Note. We do not claim that C' or $d(\phi, \phi')$ are bounded depending only on κ and C .

Proof. In the proof we abuse notation by identifying sets $S \subset G$ with their projections to the symmetric space $SK \subset X$.

We first prove the second assertion. By Theorem 1.1, the image of ϕ is within a bounded distance of a finite set \mathcal{F} of flats. Let $\mathcal{L}' = \bigcup_{F \in \mathcal{F}} F(\infty)$. Then \mathcal{L}' is a finite subset of \hat{X} , which contains $\mathcal{L}_\delta(\phi)$ for δ sufficiently small. Note that in the Key Example (§1.1), $\mathcal{L}_\delta(\phi)$ is a proper subset of \mathcal{L}' for every value of δ .

We decompose each flat in \mathcal{F} as a finite union of Weyl chambers. Every Weyl chamber \mathfrak{C} thus obtained is equivalent to a Weyl chamber of the form kA_+ , where (with \hat{X} identified with K/M) $kM \in \mathcal{L}'$ (see §2). Let $\mathcal{U} = \bigcup_{kM \in \mathcal{L}'} k\overline{A}_+$. Since \mathcal{L}' is finite, some neighborhood of \mathcal{U} contains $\phi(\mathbb{R}^n)$. Thus we may construct a (κ, C') -quasi-isometry $\phi' : \mathbb{R}^n \rightarrow \mathcal{U}$ with $d(\phi, \phi') < \infty$ by composing ϕ with the nearest-point projection from the image of ϕ to \mathcal{U} . By connecting the dots (see §3.2) we may assume that ϕ' is continuous.

Let $\mathcal{I} \subset \mathcal{U}$ be the union of the pairwise intersections of the Weyl chambers $k\overline{A}_+$, $kM \in \mathcal{L}'$. Pick $R > C'$ sufficiently large, and let \mathcal{U}' denote the subset of \mathcal{U} which is at least $2R$ away from \mathcal{I} . It is clear \mathcal{U}' is not empty, and is a disjoint union of connected components, where each component is in the interior of a single Weyl chamber. Then for $u \in \mathcal{U}'$, the connected component of $B(u, R) \cap \mathcal{U}$ containing u is homeomorphic to a ball in \mathbb{R}^n . Take $u \in \mathcal{U}' \cap \phi'(\mathbb{R}^n)$. Then for $r < R/\kappa$, $\phi'(B(\phi'^{-1}(u), r))$ is contained in a single Weyl chamber in \mathcal{U} . Hence by Lemma 8.2, $B(u, r'') \cap \mathcal{U} \subset \phi'(\mathbb{R}^n)$. Since r'' does not depend on u , each connected component of \mathcal{U}' is either completely contained in $\phi'(\mathbb{R}^n)$ or is disjoint from $\phi'(\mathbb{R}^n)$. Now by Proposition 5.9 it is clear that the connected components of \mathcal{U}' which are contained in $\phi'(\mathbb{R}^n)$ are precisely the interiors of the Weyl chambers associated to limit points in $\mathcal{L}_\delta(\phi)$.

The first assertion now follows immediately from the second. \square

8.3. Quasi-isometries of higher rank symmetric spaces. The derivation of Theorem 1.2 from Theorem 1.1 is a variation of the last part of Mostow's proof of his rigidity theorem. The difference is that we only know that the image of a flat is close to a finite union of flats, as opposed to a single flat. In this subsection we outline the relevant parts of Mostow's proof with the modifications we need, and refer the reader to [Mo] for details.

We denote the Hausdorff distance between two sets A, B as $hd(A, B)$.

Lemma 8.4 (Weyl chambers go to Weyl chambers). *Let $q : X \rightarrow X$ be a quasi-isometry. Then for every Weyl chamber \mathfrak{C} of X there exists a Weyl chamber \mathfrak{C}' so that $hd(q(\mathfrak{C}), \mathfrak{C}') < \infty$.*

Proof. For $v \in \mathbb{R}^+$ and $S \subset X$, let $T_v(S)$ denote the v -neighborhood of S in X . Consider the restriction of q to a flat F in X . Let V be as in Lemma 8.3, and let \mathfrak{C} denote a Weyl chamber in F . By Lemma 8.3, we may assume that $q(F) \subset V$. Let $B_j \subset \mathfrak{C}$ be a sequence of balls with radii tending to ∞ . Since V is a finite union of Weyl chambers, we may pass to a subsequence so that, after possibly shrinking the original balls B_j (with radii still tending to ∞), the images $q(B_j)$ are contained in the interior of some fixed Weyl chamber \mathfrak{C}' of V .⁹ [Mo, Lemma 15.1], combined with [Mo, Theorem 7.8] states that every Weyl chamber \mathfrak{C}' is within a finite Hausdorff distance from a set of the form $T_v(F_1) \cap F_2$ where F_1 and F_2 are flats with $F_2 \supset \mathfrak{C}'$. The proof of [Mo, Lemma 15.1] shows that there exists a flat $F' \in X$ such that for w sufficiently large, $hd(T_w(F') \cap V, \mathfrak{C}') < \infty$.

Now by the definition any quasi-isometry q has a coarse inverse, which is a quasi-isometry $q' : X \rightarrow X$, with $q' \circ q$ a bounded distance from the identity.

By Theorem 1.1, $q'(F')$ is contained in a neighborhood of a finite union of flats $\bigcup_{j=1}^{\ell} F_j$. [Mo, Theorem 7.8] states that for v sufficiently large, and for some (possibly unbounded) convex polyhedra \mathcal{P}_j in F with singular faces,¹⁰ $hd(T_v(F_j) \cap F, \mathcal{P}_j) < \infty$. Hence, for sufficiently large v , $hd(T_v(q'(F')) \cap F, \bigcup_{j=1}^{\ell} \mathcal{P}_j) < \infty$. Since for sufficiently large v , $T_v(q'(F'))$ contains the balls $B_j \subset \mathfrak{C}$, there is a sequence of balls $B'_j \subset \mathfrak{C} \cap \bigcup_{j=1}^{\ell} \mathcal{P}_j$ with radii increasing to infinity. An easy argument¹¹ shows that, by passing to a subsequence if needed, we may assume that all the B_j are contained in $\mathfrak{C} \cap \mathcal{P}_i$ for some i , say $i = 1$. Clearly the boundary of \mathcal{P}_1 cannot contain two parallel singular faces, hence \mathcal{P}_1 must contain a Weyl chamber \mathfrak{C}'' containing the B_j . Since the B_j belong to both \mathfrak{C} and \mathfrak{C}'' , and the radii of the B_j tends to infinity, $hd(\mathfrak{C}, \mathfrak{C}'') < \infty$. But for sufficiently large $u \gg w \gg v \gg 1$, $q(\mathfrak{C}'') \subset q(\mathcal{P}_1) \subset q(T_v(q'(F')) \cap F) \subset T_w(F') \cap V \subset T_u(\mathfrak{C}')$. Thus $q(\mathfrak{C}) \subset T_v(\mathfrak{C}')$ for sufficiently large v .

The same argument shows that $q'(\mathfrak{C}') \subset T_v(\mathfrak{C}'')$. Hence for sufficiently large $v \gg w \gg 1$, $\mathfrak{C} \subset T_w(q'(\mathfrak{C}')) \subset T_v(\mathfrak{C}'')$. Since \mathfrak{C} and \mathfrak{C}'' are both Weyl chambers, this implies that $hd(\mathfrak{C}, \mathfrak{C}'') < \infty$. Hence for large enough v , $q'(\mathfrak{C}') \subset T_v(\mathfrak{C})$. Applying q to both sides we get $\mathfrak{C}' \subset T_w(q(\mathfrak{C}))$ for large enough w . Hence $hd(q(\mathfrak{C}), \mathfrak{C}') < \infty$. \square

Remark 8.5. Quasi-isometries have the basic property that

$$(29) \quad hd(q(\text{Nbhd}(A) \cap \text{Nbhd}(B)), \text{Nbhd}(q(A)) \cap \text{Nbhd}(q(B))) < \infty$$

for sufficiently large metric neighborhoods. It is another basic fact that $hd(A, B) < \infty$ iff $hd(q(A), q(B)) < \infty$. Now given any chamber wall S , it is the intersection of two Weyl chambers \mathfrak{C} and \mathfrak{D} . Since q takes the Weyl chambers $\mathfrak{C}, \mathfrak{D}$ to within a finite Hausdorff distance of Weyl chambers $\mathfrak{C}', \mathfrak{D}'$, equation (29) shows that $q(S)$ lies a finite Hausdorff distance from the intersection of (metric neighborhoods of) \mathfrak{C}' and \mathfrak{D}' . It is implicit in [Mo, §7] that this intersection is within a finite Hausdorff distance of a chamber wall S' of \mathfrak{C}' : by successive applications of [Mo, Theorem 7.8], the intersection is within a finite Hausdorff distance from a convex polyhedron with singular faces which is contained in the Weyl chamber \mathfrak{C}' , hence it lies within a finite Hausdorff distance from a chamber wall.

⁹ $q(\mathfrak{C})$ is not contained in a tubular neighborhood of a union of faces of Weyl chambers since these sets have volume growth $\sim R^n$ and $\sim R^{n-1}$, respectively.

¹⁰A *singular face* is a hyperplane of the form $\alpha = c$ for $\alpha \in \Delta$ and $c \in \mathbb{R}$.

¹¹This argument is simply the fact that, if a union of ℓ convex sets in \mathbb{R}^n contains a ball of radius R , then at least one of the sets contains a ball of radius ϵR , where ϵ depends only on ℓ .

This argument shows that, up to finite Hausdorff distance, the quasi-isometry q takes chamber walls to chamber walls. From this it is easy to see that q induces an order-preserving bijection of the partially ordered (by inclusion) set of equivalence classes of Weyl chambers and their walls, where two sets are equivalent if they have finite Hausdorff distance.

Lemma 8.6 (Flats go to flats). *Let $q : X \rightarrow X$ be a (κ, C) -quasi-isometry. Then there exists a number $N = N(\kappa, C)$ such that for every flat $F \in X$ there exists a flat $F' \in X$ with $hd(q(F), F') < N$.*

Proof. Another way to state the conclusion drawn in Remark 8.5 is that q induces an order-preserving bijection q^* of the Tits building $\mathcal{T}(X)$ associated to X (see [Ti]). In particular q^* takes apartments to apartments. By [Ti, Theorem 5.2(i)], apartments \mathcal{A} in $\mathcal{T}(X)$ are in one-to-one correspondence with (maximal) flats F in X . Under this correspondence, the set of chambers of an apartment \mathcal{A} is precisely the set of points $F(\infty) \subset \hat{X}$ for some flat F in X . Hence the limit set $\mathcal{L}(q|_F)$ in \hat{X} is of the form $F'(\infty)$ for some flat F' in X . Now it follows from Corollary 7.4 that $q(F) \subset T_N(F')$. To get the opposite inclusion $F' \subset T_N(q(F))$ one can use Lemma 8.2 (or the remark immediately following its proof). \square

We remark that Lemma 8.6 implies a uniform version of Lemma 8.4 (see [Mo, (15.2.6)]). It is in fact this uniform version by which Mostow obtains the continuity of the induced map on the Furstenberg boundary.

Proof of Theorem 1.2. By Remark 8.5, q induces an isomorphism $q^* : \mathcal{T}(X) \rightarrow \mathcal{T}(X)$. Since q is a quasi-isometry which takes flats uniformly close to flats (Lemma 8.6), it can be shown (see [Mo, Ch. 15]) that q induces a homeomorphism on the Furstenberg maximal boundary \hat{X} . Now Tits' Theorem ([Ti]) states that, if X has no Euclidean or rank one factors, then every such isomorphism of $\mathcal{T}(X)$ is induced by an isometry of X . Hence there is an isometry ψ of X with $\psi^* = q^*$. Then $\psi^{-1} \circ q$ is a (κ, C) -quasi-isometry of X which induces the identity on $\mathcal{T}(X)$. Hence for some $N > 0$ and for all flats $F \subset X$,

$$(30) \quad \psi^{-1} \circ q(F) \subset \text{Nbhd}_N(F).$$

Now it is not hard to check that there is a constant N' , depending only on N (hence only on κ, C), so that for any $x \in X$, there exist flats F_1, F_2 in X with $F_1 \cap F_2 = x$, and with

$$(31) \quad \text{Nbhd}_N(F_1) \cap \text{Nbhd}_N(F_2) \subset \text{Nbhd}_{N'}(F_1 \cap F_2).$$

It follows from (30) and (31) that $d(\psi^{-1} \circ q(x), x) \leq N'$, and the theorem follows easily. \square

Proof of Corollary 1.4. By now this type of proof has become standard; we provide a sketch here for completeness. Recall that $QI(X)$ is defined to be the group of quasi-isometries of X modulo those quasi-isometries which are a bounded distance from the identity (indeed this modding-out makes $QI(X)$ into a group). Suppose $\phi : \Gamma \rightarrow X$ is a quasi-isometry, with coarse inverse $\psi : X \rightarrow \Gamma$. We denote by $g \mapsto L_g$ the isometric action of Γ on itself by left multiplication. We “conjugate” this action by defining $\rho(g) : X \rightarrow X$ to be $\rho(g) = \phi \circ L_g \circ \psi$, which is easily seen

to be a quasi-isometry of X with q.i. constants independent of g . It is also clear that there exists a constant $C > 0$ so that

$$d(\rho(g) \circ \rho(h), \rho(gh)) \leq C.$$

Since elements of $QI(X)$ which are a bounded distance from each other are identified, this equation shows that $\rho : \Gamma \rightarrow QI(X)$ is a homomorphism. Theorem 1.2 gives an isomorphism from $QI(X)$ to $Isom(X)$, so by composing we get a homomorphism $\rho : \Gamma \rightarrow Isom(X)$. What we need to show is that ρ has finite kernel, and that $\rho(\Gamma)$ is discrete and cocompact.

Pick a basepoint $x \in X$. Since $\rho(\Gamma)$ is a *uniform* family of quasi-isometries (i.e. the quasi-isometry constants for $\rho(g)$ are uniformly bounded), it follows that the set of $g \in \Gamma$ which take x into any fixed ball B in X must be finite. From this it easily follows that ρ has finite kernel and $\rho(\Gamma)$ is discrete. Finally, $\rho(\Gamma)$ acts cocompactly on X by definition of the action and by the fact that the C'' -neighborhood of $\phi(\Gamma)$ is all of X , for some constant $C'' > 0$. □

A. PROOFS OF SOME GEOMETRIC ESTIMATES

In this Appendix, we occasionally abuse notation and write $d(g_1, g_2)$ as a shorthand for $d(g_1K, g_2K)$, for elements g_1, g_2 of G .

A.1. Proof of Proposition 2.3. For $g \in G$, let ℓ_g denote left-translation by g , and ℓ_g^* the differential of ℓ_g . Since

$$\Phi(k \exp XM, a \exp Y) = ka \exp(\text{Ad } a^{-1}X)(\exp Y)K$$

it follows that $\ell_{(ka)^{-1}}^* \circ d\Phi_{(kM,a)}(X, Y)$ is given by the projection to \mathfrak{p} along \mathfrak{k} of the vector $\text{Ad } a^{-1}X + Y$ in \mathfrak{g} . We will now determine this explicitly in terms of bases for \mathfrak{k}' , \mathfrak{a} and \mathfrak{k} . For simplicity of notation we will use the convention that Σ^+ consists of the roots β in Σ each repeated according to its multiplicity p_β . For each β in Σ^+ we pick $X_\beta \in \mathfrak{g}^\beta$ so that the various X_β corresponding to the same element of Σ^+ form a basis for the root space \mathfrak{g}^β . We also require that the basis X_β is orthonormal with respect to the Killing form. Let $X_{-\beta} = \theta(X_\beta)$ where θ denotes the Cartan involution, defined by $\theta(X) = X$ if $X \in \mathfrak{k}$, and $\theta(X) = -X$ if $X \in \mathfrak{p}$. Let $X'_\beta = X_\beta - \theta(X_\beta)$, and let $X''_\beta = X_\beta + \theta(X_\beta)$. Then $X'_\beta \in \mathfrak{k}$ and $X''_\beta \in \mathfrak{p}$. The elements X'_β form a basis for \mathfrak{k}' , and the elements X''_β together with a basis for \mathfrak{a} form a basis for \mathfrak{p} . Then the basis X'_β is orthonormal, and the elements X''_β are also orthonormal with respect to the Killing form.

Lemma A.1. *Let $d\Phi$ denote the differential of Φ . Then*

$$\ell_{(ka)^{-1}}^* \circ d\Phi_{(kM,a)}(X'_\beta) = -\sinh \beta(a)X''_\beta$$

and

$$\ell_{(ka)^{-1}}^* \circ d\Phi_{(kM,a)}(Z) = Z \quad \text{if } Z \in \mathfrak{a}.$$

Proof. See [Sc, Lemma 8.1.2]. □

Proof of Proposition 2.3. Pick an orthonormal basis H_1, \dots, H_n for \mathfrak{a} . Since the bases $\{X'_\alpha\}$ for \mathfrak{k}' and $\{X''_\alpha\} \cup \{H_i\}$ for \mathfrak{p} are orthonormal, and $\ell_{(ka)^{-1}}$ is an isometry, the theorem follows immediately. □

A.2. Proof of (i) of Lemma 4.1.

Lemma A.2. *Suppose $g = (g_{ij}) \in SL(n, \mathbb{R})$ is sufficiently far from the origin. Then*

$$\lambda_1 \log \operatorname{tr} gg^t \leq d(gK, e) \leq \lambda_2 \log \operatorname{tr} gg^t$$

where g^t is the transposed matrix. Note that $\operatorname{tr} gg^t = \sum |g_{ij}|^2$.

Proof. Note that if $g = k_1 a k_2$, with $k_1, k_2 \in K$ and $a \in \overline{A}_+$, then $d(gK, e) = d(aK, e)$, and $\operatorname{tr} gg^t = \operatorname{tr} aa^t = \sum_i |a_i|^2$ where $a = \operatorname{diag}(a_1, \dots, a_n)$. Thus we may assume $g = a$. Now $d(aK, e) = (\sum |\log a_i|^2)^{1/2}$ and $\log \operatorname{tr} aa^t = \log \sum |a_i|^2$. Since $a_1 \geq a_2 \geq \dots \geq a_n$, $|\log a_1| \leq d(aK, e) \leq \sqrt{n} \max(|\log a_1|, |\log a_n|)$. Since $a_1 \cdots a_n = 1$, $a_1 \geq a_n = 1/(a_1 \cdots a_{n-1}) \geq 1/a_1^{n-1}$, so that $|\log a_n| \leq (n-1)|\log a_1|$. Hence,

$$|\log a_1| \leq d(aK, e) \leq \sqrt{n}(n-1)|\log a_1|.$$

Also

$$2|\log a_1| \leq \log aa^t = \log \left(\sum |a_i|^2 \right) \leq \log(n|a_1|^2) \leq 3|\log a_1|$$

if $d(aK, e)$ is large enough. The lemma follows. □

Lemma A.3. *Suppose $G = SL(n, \mathbb{R})$, and A is the diagonal subgroup. If n is upper triangular and is sufficiently far from the origin e , then $d(nK, AK) \geq \lambda d(nK, e)$, where λ depends only on the dimension.*

Proof. Let a be the element of A so that aK is closest to nK . Note that $d(nK, aK) = d(a^{-1}nK, e)$ and $d(aK, e) = d(e, a^{-1}K)$. By Lemma A.2,

$$\begin{aligned} d(nK, aK) &= d(a^{-1}nK, e) \geq \lambda_1 \operatorname{tr}(a^{-1}n)(a^{-1}n)^t \geq \lambda_1 \operatorname{tr} a^{-1}(a^{-1})^t \\ &\geq \lambda_1 \lambda_2^{-1} d(a^{-1}K, e) = \lambda_1 \lambda_2^{-1} d(aK, e). \end{aligned}$$

Hence $d(nK, e) \leq d(nK, aK) + d(aK, e) \leq (1 + \lambda_2 \lambda_1^{-1})d(nK, aK)$. □

To prove part (i) of Lemma 4.1, we formulate a more general statement for the case $G = SL(n, \mathbb{R})$.

Lemma A.4. *Suppose $x, y \in SL(n, \mathbb{R})/SO(n, \mathbb{R})$ satisfy $\alpha(x) \geq r$, $\alpha(y) \geq r$ for all $\alpha \in \sigma$, where $\sigma \subseteq \Delta$ is nonempty. If $d_K(\Theta(x), \Theta(y)M_\sigma) \geq e^{-\nu r}$, then $d(x, y) \geq \lambda_2 r$. Here, $M_\sigma = Z_K(H_\sigma)$ is the subgroup of K consisting of elements fixing H_σ under the adjoint action, where $H_\sigma \in \mathfrak{a}$ is defined by $\alpha(H_\sigma) = 1$ if $\alpha \in \sigma$, $\alpha(H_\sigma) = 0$ if $\alpha \notin \sigma$.*

Proof of Lemma A.4. The simple roots are $\alpha_i(H) = H_i - H_{i+1}$ for $1 \leq i \leq n-1$, where $H = \operatorname{diag}(H_1, \dots, H_n)$. Thus $\sigma \subseteq \Delta$ can be thought of as a subset of $1, \dots, n-1$. Since $(k_{ij}) \in M_\sigma$ if $k_{ij}(H_i - H_j) = 0$, we have

$$M_\sigma = \{k = (k_{ij}) \in K : k_{ij} = 0 \text{ if } [j, i-1] \not\subseteq \Delta \setminus \sigma\}.$$

We may write $x = k_1 a K = k_1 \operatorname{diag}(a_1, \dots, a_n) K$ and $y = k_2 a' K$, where $a' = \operatorname{diag}(a'_1, \dots, a'_n)$. Then $d(x, y) = d(k_1 a K, k_2 a' K) = d(aK, k a' K)$, where $k = k_1^{-1} k_2$. Since the map $X \rightarrow \overline{A}_+ K$ sending $gK = \hat{k} \hat{a} K$ to $\hat{a} K$ is a contraction, $d(x, y) \geq$

$d(aK, a'K) \geq \lambda_3 |\log a_j/a'_j|$, for any $j \in [1, n - 1]$. Also for any pair (i, j) , by Lemma A.2,

$$\begin{aligned} d(x, y) &\geq \lambda_4 \log \operatorname{tr} x^{-1}y(x^{-1}y)^t \\ &\geq \lambda_4 |\log(x^{-1}y)_{ij}| \\ &= \lambda_4 |\log(a_i k_{ij} a_j^{-1}) + \log(a_j/a'_j)| \\ &\geq \lambda_4 |\log(a_i k_{ij} a_j^{-1})| - \lambda_4 |\log(a_j/a'_j)| \\ &\geq \lambda_4 |\log(a_i k_{ij} a_j^{-1})| - \lambda_3 \lambda_4 d(x, y). \end{aligned}$$

Hence, for any pair (i, j) ,

$$(32) \quad d(x, y) \geq \lambda_5 |\log(a_i k_{ij} a_j^{-1})|.$$

By compactness, for any $k \in K$ there exists a pair (i, j) so that

$$(33) \quad |k_{ij}| > \lambda d_K(kM_\sigma, e)$$

where $\lambda = \lambda(X)$. Indeed, $\max_{k \in K} d_K(kM_\sigma, e) \leq d_1$, and

$$k : d_K(k, M_\sigma) \geq d_2 > 0 \quad \max_{[j, i-1] \notin \Delta - \sigma} |k_{ij}| = d_3 > 0,$$

since if $d_3 = 0$ then there exists $k_1 \in M_\sigma$ with $d_K(k_1 M_\sigma, e) > 0$. Thus for any $k \in K$ such that $d_K(kM_\sigma, e) \geq d_2$, there exists $[j, i - 1] \notin \Delta \setminus \sigma$ such that $|k_{ij}| \geq d_3/d_1 d_K(kM_\sigma, e)$.

We may now choose (i, j) so that (33) holds, and $[i, j] \notin \Delta \setminus \sigma$. Since exchanging x and y replaces k by k^{-1} and k is orthogonal, we may assume $i > j$. Then $a_i/a_j \geq e^r$ since for some $l, l + 1 \in [j, i]$, $\alpha_l \in \sigma$ and thus $a_l/a_{l+1} \geq e^r$. Combining this estimate with (32) and (33) proves the lemma. \square

Proof of (i) of Lemma 4.1. Let $H_0 \in \mathfrak{a}_+$ be defined by $\alpha(H_0) = 1$ for all $\alpha \in \Delta$. Let ρ be a faithful linear representation of G , i.e. an injective homomorphism $G \rightarrow G' = \operatorname{SL}(n, \mathbb{R})$. We may choose $K', \mathfrak{a}', \Delta'$ etc. so that $\rho(G) \cap K' = \rho(K)$ and $\rho(\mathfrak{a}_+) \subset \mathfrak{a}'_+$ etc. Let $\tau \subset \Delta'$ be defined by

$$\tau = \{\alpha' \in \Delta' : \alpha'(\rho(H_0)) > 0\}.$$

Then $M'_\tau = Z_{K'}(H_\tau) = Z_{K'}(\rho(H_0))$. Hence, $\rho(K) \cap M'_\tau = \rho(M)$, and so $\rho : K/M \rightarrow K'/M'_\tau$ is a totally geodesic embedding. Also ρ when viewed as a map from G/K to G'/K' is a totally geodesic embedding. Now let $\lambda' = \min_{\alpha' \in \tau} \alpha'(\rho(H))$. By construction, $\lambda' > 0$. Finally, if $H \in \mathfrak{a}_+$ satisfies $\alpha(H) \geq r$ for all $\alpha \in \Delta$, then $H = rH_0 + \hat{H}$, where $\hat{H} \in \overline{\mathfrak{a}_+}$. Hence, for all $\alpha' \in \tau$, $\alpha'(\rho(H)) = r\alpha'(\rho(H_0)) + \alpha'(\rho(\hat{H})) \geq r\alpha'(\rho(H_0)) \geq \lambda'r$. This shows that if x, y satisfy the hypothesis of the lemma, i.e. $\alpha(x) > r, \alpha(y) \geq r$, then $\rho(x), \rho(y)$ satisfy the hypothesis of Lemma A.4, i.e. $\alpha'(\rho(x)) > \lambda'r, \alpha'(\rho(y)) \geq \lambda'r$. Hence we can apply Lemma A.4 with r replaced by $\lambda'r$. \square

A.3. Proof of Lemma 7.2. Since W normalizes MA , for any $w \in W, BwB = NwB$. The longest element w_0 also has the property that $w_0 N w_0^{-1} \cap N = \{e\}$. This implies that for $n_1, n_2 \in N$, if $n_1 w_0 B = n_2 w_0 B$, then $n_1 = n_2$. Thus we can define a map $\mathcal{N} : (K \setminus \mathcal{S})/M \rightarrow N$ by requiring $k = \mathcal{N}(k)man$ for some $m \in M, a \in A, n \in N$. This map is in fact a diffeomorphism onto; an inverse is obtained by sending $n \in N$ to the K part of the Iwasawa decomposition of nw_0 .

Proof of Lemma 7.2. We denote $k_1 k_2^{-1}$ by k . Since $k \notin \mathcal{S}$, we may write

$$(34) \quad k = n w_0 a n',$$

where $n = \mathcal{N}(k)$. Then, since MAN stabilizes o ,

$$F = F(o, k \cdot o) = F(o, n w_0 a n' \cdot o) = F(o, n w_0 \cdot o).$$

Since n stabilizes o ,

$$d(e, F) = d(e, F(o, n w_0 \cdot o)) = d(n^{-1}, F(o, w_0 \cdot o)) = d(n^{-1}, A).$$

Since $n^{-1} = \mathcal{N}(k)^{-1}$, the theorem follows immediately from the fact that \mathcal{N} is a diffeomorphism onto $(K - \mathcal{S})/M$. \square

A.4. Proof of Lemma 2.4.

Proof of Lemma 2.4. By the assumption, x is nondegenerate. We write $x = kaK$; then $kM = \Theta(x)$. Then if we write $g^{-1}k = k''a''n''$, where $k'' \in K$, $a'' \in A$, and $n'' \in N$, then $g^{-1} \cdot \Theta(x) = g^{-1} \cdot k = k''M$, by the definition of the action on the boundary. But $g^{-1}x = g^{-1}kaK = k''a''n''aK = k''aa''(a^{-1}n''a)K$. Denote $y = k''aa''K$ so that $\Theta(y) = k''M$. Since $Ad(a)$ is a contraction on N , $d(a^{-1}n''aK, e) \leq d(n''K, e)$. By Lemma A.3, $d(n''K, e) \leq \lambda^{-1}d(gK, e)$. Thus $d(a^{-1}n''aK, e) \leq (\nu/\lambda)r$. But then $d(g^{-1}x, y) = d(a^{-1}n''aK, e) \leq (\nu/\lambda)r$, thus, if r is sufficiently large, $d_K(\Theta(g^{-1}x), \Theta(y)) \leq e^{-\lambda_1 r}$, since if $d_K(\Theta(g^{-1}x), \Theta(y)) \geq e^{-\lambda_1 r}$ then by Lemma 4.1 $d(g^{-1}x, y) \geq \lambda_2 r$ and ν can be chosen to be arbitrarily small. \square

A.5. **Proof of Lemma 6.1.** The following lemma is well known:

Lemma A.5 (Nilpotent orbits are exponentially distorted). *Let x be a point in X , and let y be a point in the orbit Nx . Then if x and y are sufficiently far apart, $d_N(x, y) \geq \lambda_1 e^{\lambda_2 d(x, y)}$, where d_N denotes the path metric along the orbit Nx .*

Proof. We may assume that $G = SL(n, \mathbb{R})$. Write $x = anK$. Then $Nx = aNK$, so that $yK = an'K$. Hence $d(x, y) = d(nK, n'K) \leq \lambda \log \operatorname{tr} n^{-1}n'(n^{-1}n')^t$. From the elementary properties of nilpotent groups, $d_N(nK, n'K)$ is bounded below by some polynomial in $\operatorname{tr} n^{-1}n'(n^{-1}n')^t$. \square

Lemma A.6. *Suppose $\sigma \subset \Delta$, x and y are points in X , and γ a path connecting x, y with $\alpha(\gamma(t)) \geq r$ for all $\alpha \in \sigma$ and $t \in [0, 1]$. Then $\ell(\gamma) \geq e^{\lambda r} d_K(xM_\sigma, yM_\sigma)$.*

Proof. This is clear from Proposition 2.3 and the observation that the only tangent vectors in \mathfrak{k} which are not multiplied by positive exponentials in the expression for the metric are those tangent to M_σ . \square

Lemma A.7 (Spheres in Y_v are exponentially distorted). *Let x and y be points in Y_v satisfying $d(x, e) = d(y, e) = r$. Then there exist constants $\lambda_1, \lambda_2, \lambda_3$ so that for sufficiently large r , if $d(x, y) > \lambda_3$, then any path in Y_v connecting x and y which stays outside the ball of radius r centered at the origin has length at least $\lambda_1 e^{\lambda_2 d(x, y)}$.*

Proof. Without loss of generality we may assume $x = e$. Let σ denote the set of $\alpha \in \Delta$ such that $\alpha(v) = 0$. Let $N_\sigma = N/(N \cap M_\sigma)$. Let \mathcal{U} be a compact set containing e in the interior on which the $N_\sigma M_\sigma AN$ decomposition is defined. Suppose $y \notin \mathcal{U}$. Let z be the point on which the projection to the sphere of the geodesic \overline{xy} leaves \mathcal{U} . Then $\ell(\gamma) \geq \ell(\overline{xy}) \geq \ell(\overline{xz})$. By Lemma A.6, $\ell(\overline{xz}) \geq e^{\lambda r}$,

since $d_K(xM_\sigma, zM_\sigma) > \lambda' > 0$. Since $d(x, y) \leq 2r$, $\ell(\gamma) \geq e^{(\lambda/2)d(x,y)}$ and the lemma follows in this case. Thus we may assume that the projection to K of the path γ stays in \mathcal{U} .

Suppose x and y are connected by a path $\gamma(t)$. Since radial projection on the sphere $\{x : d(x, e) = r\}$ in Y_v is clearly a contraction, we may assume that $\gamma(t)$ is contained in the sphere, i.e. $\gamma(t) = k(t)a_rK$ where a_r is the unique element in $\exp \mathfrak{a}_+$ such that $d(a_r, e) = r$. We may write $k(t) = \bar{n}(t)a(t)n(t)$, where $\bar{n}(t) \in \bar{N}_\sigma = w_0N_\sigma w_0^{-1}$. Then $\gamma(t) = k(t)a_rK = \bar{n}(t)a(t)a_r(a_r^{-1}n(t)a_r)K$. Let $\gamma_1(t) = \bar{n}(t)a(t)a_rK$. Since $Ad(a_r)$ is a contraction on n , $d(\gamma(t)K, \gamma_1(t)) \leq d(n(t)K, e) = O(1)$ by compactness. Let $x' = \gamma_1(0)$, $y' = \gamma_1(1)$. Then $d(x, x') = O(1)$, and $d(y, y') = O(1)$. Here and below in the proof of Lemma A.7, $O(1)$ means a constant independent of r .

Let $u \in A$ be such that for all $\alpha \in \Delta$, $\alpha(u) = \inf_{0 \leq t \leq 1} \alpha(a(t))$. Then $d(uK, e) = O(1)$, hence $d(a_ruK, a_rK) = O(1)$. Let p be the map $G/K \rightarrow \bar{N}a_ruK$ sending $\bar{n}aa_rK$ to $\bar{n}a_ruK$. By the definition of u , p is a contraction when restricted to the image of $\gamma_1(t)$. Hence $\ell(\gamma_1(t)) \geq d_{\bar{N}}(p(x'), p(y'))$. Now $d_{\bar{N}}(p(x), p(y)) \geq e^{\lambda_1 d(p(x'), p(y'))}$ by Lemma A.5. Note that

$$d(p(\gamma_1(t)), \gamma_1(t)) = d(\bar{n}ta_ruK, \bar{n}(t)a(t)a_rK) = d(uK, a(t)K) = O(1).$$

In particular, $d(x', p(x')) = O(1)$, and $d(y', p(y')) = O(1)$. Thus $d(x, p(x')) = O(1)$ and $d(y, p(y')) = O(1)$. Therefore by the triangle inequality, $d(p(x'), p(y')) \geq d(x, y) - O(1) \geq (1/2)d(x, y)$ if r is sufficiently large. This implies the lemma. \square

Proof of Lemma 6.1. Let a'_0 (resp. b'_0) denote the radial projection of a_0 (resp. b_0) on a circle of radius r . Suppose $d(a_0, b_0) \geq c_1r$. Then $d(a'_0, b'_0) \geq (c_1/2)r$.

As above let σ denote the set of $\alpha \in \Delta$ such that $\alpha(v) \neq 0$. Since Y_v need not consist of nondegenerate elements, for $x \in Y_v$, $\Theta(x)$ is only defined as a map to K/M_σ , where M_σ is the centralizer of $\exp v$ in K . (M_σ depends only on σ , not on v .) The space K/M_σ has a left K -invariant metric, denoted d_σ .

For $R > 0$ let $ds_\sigma^{(R)}$ denote the induced metric on the sphere of radius R in Y_v . Let $d_\sigma^{(R)}$ denote the metric on K/M_σ which is the pullback of $ds_\sigma^{(R)}$ via the map $K/M_\sigma \ni k \rightarrow k(\exp Rv)K \in Y_v$. For a path γ in K/M_σ , let $\ell_\sigma^{(R)}(\gamma)$ denote the length of γ with respect to the metric $d_\sigma^{(R)}$.

If r is large enough, by Lemma A.7 and the fact that projection on the ball of radius r is a contraction,

$$(35) \quad \ell_\sigma^{(r)}(\Theta(\hat{\psi})) \geq d_\sigma^{(R)}(a'_0, b'_0) > \lambda_1 e^{\lambda_2 d(a'_0, b'_0)} \geq \lambda_1 e^{(\lambda_2 c_1/2)r}.$$

Since projection on the ball of radius $2^j r$ is a contraction,

$$(36) \quad \ell_\sigma^{(2^j r)}(\Theta(\hat{\psi} \cap \hat{A}_j)) \leq \ell(\hat{\psi} \cap \hat{A}_j) \leq c_2 2^j r.$$

By Proposition 2.3, for any path γ in K/M_σ ,

$$(37) \quad \ell_\sigma^{(2^j r)}(\gamma) \geq e^{2^j \alpha_{\min}(a_r)} \ell_\sigma^{(r)}(\gamma)$$

where $\alpha_{\min}(a_r) = \min_{\alpha \in \sigma} \alpha(a_r) > 0$ and we have used $\alpha(a_{2^j r}) = 2^j \alpha(a_r)$. Hence combining (36) with (37) we get

$$\ell_\sigma^{(r)}(\Theta(\hat{\psi} \cap \hat{A}_j)) \leq c_2 (2^j r) e^{-2^j \alpha_{\min}(a_r)}.$$

Summing over j we get

$$\ell_{\sigma}^{(r)}(\Theta(\hat{\psi})) \leq c_2 r \sum_j 2^j e^{-2^j \alpha_{\min}(a_r)}$$

which contradicts (35) if r is large enough. \square

ACKNOWLEDGMENTS

Thanks to Bruce Kleiner for showing us that there are quasi-flats not close to a single flat, to Geoff Mess and Shmuel Weinberger for useful correspondence, and to Bob Zimmer for his support and encouragement. We also thank David Fisher, Martin Pergler and Jennifer Taback for thoroughly reading and correcting our manuscript.

REFERENCES

- [BGS] W. Ballmann, M. Gromov, and V. Schroeder, *Manifolds of Nonpositive Curvature*, Progress in Math. Vol. 61, Birkhauser, 1985. MR **87h**:53050
- [E] A. Eskin, *Quasi-isometric rigidity of higher rank nonuniform lattices*, preprint.
- [EF] A. Eskin and B. Farb, *Quasi-flats in $\mathbb{H}^2 \times \mathbb{H}^2$* , to appear in *Proceedings of the Colloquium on Lie Groups and Ergodic Theory, Tata Institute (Jan 1996)*.
- [FS] B. Farb and R. Schwartz, *The large-scale geometry of Hilbert modular groups*, J. Diff. Geom. **44**, No. 3, (1996) pp 435–478.
- [KL] B. Kleiner and B. Leeb, *Rigidity of quasi-isometries for symmetric spaces and Euclidean buildings*, to appear in Publ. IHES.
- [He] S. Helgason, *Differential Geometry, Lie Groups, and Symmetric Spaces*, Academic Press, 1978. MR **80k**:53081
- [Mo] G.D. Mostow, *Strong Rigidity of Locally Symmetric Spaces*, Annals of Math. Studies, No. 78, Princeton Univ. Press, 1973. MR **52**:5874
- [Pa] P. Pansu, *Metriques de Carnot-Caratheodory et quasiisometries des espaces symetriques de rang un*, Annals of Math. **129** (1989), pp. 1–60. MR **90e**:53058
- [Sc] H. Schlichtkrull, *Hyperfunctions and Harmonic Analysis on Symmetric Spaces*, Birkhauser, 1984. MR **86g**:22021
- [Ti] J. Tits, *Buildings of spherical type and finite BN-pairs*, Lecture Notes in Math., vol. 386, Springer-Verlag, 1974. MR **57**:9866

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF CHICAGO, CHICAGO, ILLINOIS 60637

E-mail address: eskin@math.uchicago.edu

E-mail address: farb@math.uchicago.edu