RICHARD EVAN SCHWARTZ

The quasi-isometry classification of rank one lattices

Publications mathématiques de l'I.H.É.S., tome 82 (1995), p. 133-168

http://www.numdam.org/item?id=PMIHES_1995__82__133_0

© Publications mathématiques de l'I.H.É.S., 1995, tous droits réservés.

L'accès aux archives de la revue « Publications mathématiques de l'I.H.É.S. » (<http://www.ihes.fr/IHES/Publications/Publications.html>) implique l'accord avec les conditions générales d'utilisation (<http://www.numdam.org/conditions>). Toute utilisation commerciale ou impression systématique est constitutive d'une infraction pénale. Toute copie ou impression de ce fichier doit contenir la présente mention de copyright.

NUMDAM

Article numérisé dans le cadre du programme
Numérisation de documents anciens mathématiques
<http://www.numdam.org/>

THE QUASI-ISOMETRY CLASSIFICATION OF RANK ONE LATTICES

by RICHARD EVAN SCHWARTZ*

ABSTRACT. Let X be a symmetric space—other than the hyperbolic plane—of strictly negative sectional curvature. Let G be the isometry group of X . We show that any quasi-isometry between non-uniform lattices in G is equivalent to (the restriction of) a group element of G which commensurates one lattice to the other. This result has the following corollaries:

1. Two non-uniform lattices in G are quasi-isometric if and only if they are commensurable.
2. Let Γ be a finitely generated group which is quasi-isometric to a non-uniform lattice in G . Then Γ is a finite extension of a non-uniform lattice in G .
3. A non-uniform lattice in G is arithmetic if and only if it has infinite index in its quasi-isometry group.

1. Introduction

A *quasi-isometry* between metric spaces is a map which distorts distances by a uniformly bounded factor, above a given scale. (See § 2.1 for a precise definition.) Quasi-isometries ignore the local structure of metric spaces, but capture a great deal of their large scale geometry.

A finitely generated group G has a natural *word metric*, which makes it into a path metric space. Different finite generating sets produce quasi-isometric spaces. In other words, quasi-isometric properties of the metric space associated to the group only depend on the group itself. There has been much interest recently in understanding these quasi-isometric properties. (See [Gr1] for a detailed survey.)

Lattices in Lie groups provide a concrete and interesting family of finitely generated groups. A uniform (i.e. co-compact) lattice in a Lie group is always quasi-isometric to the group itself (equipped with a left-invariant metric). In particular, two uniform lattices in the same Lie group are quasi-isometric to each other.

In contrast, much less is known about quasi-isometries between non-uniform lattices. S. Gersten posed the natural question:

When are two non-uniform lattices quasi-isometric to each other?

The purpose of this paper is to answer Gersten's question, completely, in the case of *rank one semi-simple* Lie groups. Such groups agree, up to index 2, with isometry groups

* Supported by an NSF Postdoctoral Fellowship.

of negatively curved symmetric spaces. (See § 2.3 for a list.) To save words, we shall enlarge our Lie groups so that they precisely coincide with isometry groups of negatively curved symmetric spaces. We will simply call these groups *rank one* Lie groups. We will call their lattices *rank one lattices*.

The most familiar examples of non-uniform rank one lattices are fundamental groups of finite volume, non-compact, hyperbolic Riemann surfaces. Such lattices are finitely generated free groups. Every two such are quasi-isometric to each other. In all other cases, we uncover a rigidity phenomenon which is related to, and in some sense stronger than, Mostow rigidity.

1.1. Statement of Results

Let G be a Lie group. Given two lattices $\Lambda_1, \Lambda_2 \subset G$, we say that an element $\gamma \in G$ *commensurates* Λ_1 to Λ_2 if $\gamma \circ \Lambda_1 \circ \gamma^{-1} \cap \Lambda_2$ has finite index in Λ_2 . In particular, the group of elements which commensurate Λ to itself is called the *commensurator* of Λ . Given two lattices Λ_1 and Λ_2 , an isometry which commensurates Λ_1 to Λ_2 induces, by restriction, a quasi-isometry between Λ_1 and Λ_2 . We shall let $\text{Isom}(\mathbf{H}^2)$ be the isometry group of the hyperbolic plane.

Theorem 1.1 (Main Theorem). — *Let $G \neq \text{Isom}(\mathbf{H}^2)$ be a rank one Lie group, and let Λ_1 and Λ_2 be two non-uniform lattices in G . Any quasi-isometry between Λ_1 and Λ_2 is equivalent to (the restriction of) an element of G which commensurates Λ_1 to Λ_2 .*

Roughly speaking, two quasi-isometries between metric spaces are said to be *equivalent* if one can be obtained from the other by a uniformly bounded modification. (See § 2.1 for a precise definition.) The group of quasi-isometries, modulo equivalence, of a metric space M to itself is called the *quasi-isometry group* of M . The Main Theorem immediately implies:

Corollary 1.2. — *Let $\Lambda \subset G$ be a non-uniform lattice in a rank one Lie group $G \neq \text{Isom}(\mathbf{H}^2)$. Then the commensurator of Λ and the quasi-isometry group of Λ are canonically isomorphic.*

Here is the complete quasi-isometry classification of rank one lattices.

Corollary 1.3. — *Suppose, for $j = 1, 2$, that Λ_j is a lattice in the rank one Lie group G_j . Then Λ_1 and Λ_2 are quasi-isometric if and only if $G_1 = G_2$ and exactly one of the following statements holds:*

1. Λ_j is a non-uniform lattice in $\text{Isom}(\mathbf{H}^2)$.
2. Λ_j is a uniform lattice.
3. Λ_j is a non-uniform lattice in $G_j \neq \text{Isom}(\mathbf{H}^2)$. Furthermore Λ_1 and Λ_2 are commensurable.

In a rather tautological way, the Main Theorem gives the following extremely general rigidity result.

Corollary 1.4. — *Let $G \neq \text{Isom}(\mathbf{H}^2)$ be a rank one Lie group. Suppose that Γ is an arbitrary finitely generated group, quasi-isometric to a non-uniform lattice Λ of G . Then Γ is a finite extension of a non-uniform lattice Λ' of G . Furthermore Λ and Λ' are commensurable.*

Finally, combining the Main Theorem with Margulis' well-known characterization of arithmeticity we obtain:

Corollary 1.5. — *Let $G \neq \text{Isom}(\mathbf{H}^2)$ be a rank one Lie group, and let Λ be a non-uniform lattice in G . Then Λ is arithmetic if and only if it has infinite index in its quasi-isometry group.*

This last corollary is philosophical in nature. It says that arithmeticity, which is a number theoretic concept, is actually implicit in the group structure.

1.2. Outline of the Proof

Let $X \neq \mathbf{H}^2$ be a negatively curved symmetric space. (See § 2.3 for a list.) A *neutered space* is defined to be the closure, in X , of the complement of a disjoint union of horoballs. The neutered space is equipped with the path metric induced from the Riemannian metric on X . This path metric is called the *neutered metric*. Let Λ denote the isometry group of Ω . We say that Ω is *equivariant* if the quotient space Ω/Λ is compact ⁽¹⁾.

Step 1. — Introducing Neutered Spaces

We will show in § 2 (the standard fact) that an equivariant neutered space is quasi-isometric to its isometry group. Since any non-uniform lattice can be realized—up to finite index—as the group of isometries of an equivariant neutered space, we can ignore the groups themselves, and work with neutered spaces.

Step 2. — Horospheres Quasi-Preserved

We will show that the image of a quasi-isometric embedding of a horosphere into a neutered space must stay close to some (unique) boundary horosphere of that neutered space. In particular, any quasi-isometry between neutered spaces must take boundary horospheres to boundary horospheres. This is done in § 3 and § 4.

Step 3. — Ambient Extension

Let $\Omega_1, \Omega_2 \subset X$ be neutered spaces, and let $q : \Omega_1 \rightarrow \Omega_2$ be a quasi-isometry (relative to the two neutered metrics). From Step 2, q pairs up the boundary components of Ω_1 with those of Ω_2 . In § 5, we extend q to a map $\bar{q} : X \rightarrow X$. It turns out that this extension is a quasi-isometry of X , which “remembers” the horoballs used to define Ω_j . We will abbreviate this by saying that $\bar{q}$ is *adapted* to the pair (Ω_1, Ω_2) .

⁽¹⁾ *Equivariant neutered spaces* are called *invariant cores* in [Gr1]. Technically speaking, we are only concerned with results about equivariant neutered spaces, since these arise naturally in connection with lattices. However, many of the steps in our argument do not require the assumption of equivariance, and the logic of the argument is clarified by the use of more general terminology.

Step 4. — Geometric Limits

We now introduce the assumption that Ω_j is an equivariant neutered space. Let $h = \partial\bar{q}$ be the boundary extension of $\bar{q}$ to ∂X . It is known that h is quasiconformal, and almost everywhere differentiable. (In the real hyperbolic case, this is due to Mostow [M2]; the general formulation we need is due to Pansu [P].) We will work in *stereographic coordinates*, in which $\partial X - \infty$ is a Heisenberg group. (Euclidean space in the real-hyperbolic case.) By “zooming-in” towards a generic point x of differentiability of h , we produce a new quasi-isometry $q' : X \rightarrow X$ having the following properties:

1. q' is adapted to a new pair (Ω'_1, Ω'_2) of equivariant neutered spaces.
2. $h' = \partial q'$ is a nilpotent group automorphism of $\partial X - \infty$. (Linear transformation in the real case.)
3. $h' = dh(x)$, the linear differential at x .

In § 6, we make this construction in the real hyperbolic case. In § 8, we work out the general case.

Step 5. — Inverted Linear Maps

Suppose q' is the quasi-isometry of X produced in Step 5. Using a trick involving inversion, we show that h' is in fact a Heisenberg similarity. (Ordinary similarity in the real-hyperbolic case.) This is to say that $dh(x)$ is a similarity. Since x is generic, the original map h is 1-quasi-conformal, and hence is the restriction of an isometry of X . We work out the real hyperbolic case in § 7, and the complex hyperbolic case in § 8. The quaternionic (and Cayley) cases have similar proofs, and also follow directly from [P, Th. 1].

Step 6. — The Commensurator

Let $\bar{q} : X \rightarrow X$ be a quasi-isometry adapted to the pair (Ω_1, Ω_2) of equivariant neutered spaces. We know from Steps 4-5 that $\bar{q}$ is equivalent to an isometry q_* . In § 9, we use a trick, similar to the one developed in Step 5, to show that q_* commensurates the isometry group of Ω_1 to that of Ω_2 . In § 10, we recall the correspondence between Ω_j and Λ_j , and see that q_* commensurates Λ_1 to Λ_2 .

Remark. — Our quasi-isometry from one lattice to another is not *a priori* assumed to (virtually) conjugate one lattice to the other. This situation is in marked contrast to the situation in Mostow rigidity. Accordingly, the details of Steps § 4-6 are quite different (in places) from those usually associated with Mostow rigidity [M1].

1.3. Suggested Itinerary

It should be possible for the reader to read only portions of the paper and still come away with an understanding of the main ideas. Here are some suggestions for the order in which to read the paper:

1. To see the basic skeleton (and prettiest part) of the paper, without getting bogged down in the details of Step 2 and Step 3, read § 2.1-2.4, § 5.1, § 6, § 7, § 9 and § 10 1-10.2.

2. To understand Step 3, use Step 2 as a black box. Read § 2, § 3.1-3.2 and § 5.
3. To understand Step 2, read § 2 and § 3, using § 4 as a black box.
4. Read § 4 last. It helps to draw a lot of pictures here.

In general—and especially for Step 2—the reader should first restrict his attention to real hyperbolic space. The general case is conceptually the same as this case, but requires tedious background results (§ 2.5-2.7 and § 4.1) on rank one geometry.

1.4. Acknowledgements

It almost goes without saying that this paper would have been impossible without the beautiful ideas of Mostow. I would also like to thank:

1. Benson Farb, for originally telling me about Gersten's question, for numerous conversations about geometric group theory, and for his enthusiasm about this work.
2. Peter Doyle, for a great deal of moral support.
3. Steve Gersten, for posing such an interesting question.
4. Pierre Pansu, for his theory of Carnot differentiability.
5. Mladen Bestvina, Martin Bridgeman, Jeremy Kahn, Misha Kapovich, Geoff Mess, Bob Miner, and Pierre Pansu, for helpful and interesting conversations on a variety of subjects related to this work.

1.5. Dedication

I dedicate this paper to my wife, Brienne Elisabeth, on the occasion of our wedding, May 13, 1995.

2. Background

2.1. Quasi-isometries

Let (M, d) be a metric space, with metric d . A subset $N \subset M$ is said to be a *K-net* if every point of M is within K of some point of N . A *K-quasi-isometric embedding* of (M, d) into (M', d') is a map $q : N \rightarrow M'$ such that:

1. N is a K -net in M .
2. $d'(q(x), q(y)) \in [d(x, y)/K - K, Kd(x, y) + K]$, for $x, y \in N$.

The map q is said to be a *K-quasi-isometry* if the set $N' = q(N)$ is a K -net in M' . In this case, the two metric spaces (M, d) and (M', d') are said to be *K-quasi-isometric*. When the choice of K is not important, we will drop K from the terminology.

Two quasi-isometric embeddings (or quasi-isometries) $q_1, q_2 : M \rightarrow M'$ are said to be *equivalent* if there are constants C_1 and C_2 having the following properties:

1. Every point of N_1 is within C_1 of N_2 , and vice versa.
2. If $x_j \in N_j$ are such that $d(x_1, x_2) \leq C_1$, then $d'(q_1(x_1), q_2(x_2)) \leq C_2$.

It is routine to verify that the above relation is an equivalence relation, and that, modulo this relation, the quasi-isometries of M form a group. We call this group the *quasi-isometry group* of M .

A *K-quasi-geodesic* (segment) in M is a K -quasi-isometric embedding of (a segment of) $\mathbf{R}$ into M . If M happens to be a path metric space, then every such segment is equivalent to a K -bi-Lipschitz segment. The uniformity of the equivalence only depends on K .

For more information about quasi-isometries, see [Gr1] and [E1].

2.2. The Word Metric

Let G be a finitely generated group. A finite generating set $S \subset G$ is said to be *symmetric* provided that $g \in S$ if and only if $g^{-1} \in S$. The word metric on G (resp. S) is defined as follows: The distance from $g_1, g_2 \in G$ is defined to be the minimum number of generators needed to generate the element $g_1 g_2^{-1}$. It is easy to see that this makes G into a path metric space. Two different finite generating sets S_1 and S_2 produce Lipschitz equivalent (and in particular, quasi-isometric) metric spaces.

2.3. Rank One Symmetric Spaces

Here is the list of symmetric spaces which have (strictly) negative sectional curvature:

1. Real hyperbolic n -space, $\mathbf{H}^n$.
2. Complex hyperbolic n -space, $\mathbf{CH}^n$.
3. Quaternionic hyperbolic n -space, $\mathbf{QH}^n$.
4. The Cayley plane.

These spaces are also known as *rank one symmetric spaces*. For basic information about these spaces, see [C], [Gr1], [Go], [E2], or [T].

We will let X denote a rank one symmetric space. We will never take $X = \mathbf{H}^2 = \mathbf{CH}^1$, unless we say so explicitly. Also, to simplify the exposition, we will omit the description of the Cayley plane. Readers who are familiar with (and interested in) this one exceptional case can easily adapt all our arguments to fit it. Let g_X be the symmetric Riemannian metric on X . In the usual way, g_X induces a path metric on X . We will denote the path metric by d_X .

2.4. Neutered Spaces

Let X be a rank one symmetric space. A *horoball* of X is defined to be the limit of unboundedly large metric balls, provided that this limit exists, and is not all of X . The isometry group of X transitively permutes the horoballs. A *horosphere* is the boundary of a horoball.

We define a *neutered space* Ω to be the closure, in X , of the complement of a non-empty disjoint union V of horoballs, equipped with the path metric. We will call this

metric the *neutered metric*, and denote it by d_Ω . We will say that the horoballs of V are *horoballs* of Ω . This is a slight abuse of language, because only the bounding horospheres actually belong to Ω .

The metrics $d_X|_\Omega$ and d_Ω are not Lipschitz equivalent. However, they are Lipschitz equivalent below any given scale. The Lipschitz constant only depends on the scale, and not on the neutered space. We say that Ω is *equivariant* if it admits a co-compact group of isometries. These isometries are necessarily restrictions of isometries of X .

Lemma 2.1. — *Suppose Λ is a non-uniform lattice of isometries of X . Then there is an equivariant neutered space $\Omega \subset X$ such that:*

1. Λ has finite index in the isometry group of Ω .
2. Λ is canonically quasi-isometric to Ω .

Proof. — We remove disjoint horoball neighborhoods of the quotient X/Λ , and then lift to X . These neighborhoods are covered by a disjoint union of horoballs. The closure, $\Omega \subset X$, of the complement of these horoballs, is an equivariant neutered space. A well-known criterion of Milnor-Svarc says that, since Λ acts virtually freely, with compact quotient, on the path space Ω , the two spaces Λ and Ω are quasi-isometric. The canonical quasi-isometry is given by mapping $\lambda \in \Lambda$ to the point $\lambda(x)$, for some pre-chosen point $x \in \Omega$. The equivalence class of this quasi-isometry is independent of the choice of x . $\square$

To avoid certain trivialities, we will assume that any neutered space we consider has at least 3 horospheres. Certainly, this condition is fulfilled for equivariant neutered spaces.

2.5. Rank One Geometry: Horospheres

Let $\sigma \subset X$ be a horosphere. We will let d_σ denote the path metric on σ induced from the Riemannian metric $g_X|_\sigma$. As is well known, (σ, d_σ) is isometric to a Euclidean space, when $X = \mathbf{H}^n$. Below, we will describe the geometry of σ , when $X = \mathbf{FH}^n$. Here, we will take $\mathbf{F}$ to be either the complex numbers $\mathbf{C}$, or the quaternions $\mathbf{Q}$.

Let $T(\sigma)$ denote the tangent bundle of σ , considered as a sub-bundle of the tangent bundle to $\mathbf{FH}^n$. There is a canonical codimension $\dim(\mathbf{F}) - 1$ distribution

$$D(\sigma) \subset T(\sigma)$$

defined by the maximal $\mathbf{F}$ -linear subspaces in $T(\sigma)$. This distribution is totally non-integrable, in the sense that any two points $p, q \in \sigma$ can be joined by a curve which is integral to $D(\sigma)$. Furthermore, $D(\sigma)$ is totally symmetric, in the sense that the stabilizer of σ in $\text{Isom}(X)$ acts transitively on pairs (p, v) , where $p \in \sigma$, and $v \in D_p(\sigma)$.

There are explicit “coordinates” for both σ and $D(\sigma)$, which we now describe. Let $\bar{\xi}$ denote the componentwise conjugate of ξ , in $\mathbf{F}^n$. Let $\text{Im}(\xi)$ denote the imaginary

part of ξ ; $\text{Im}(\xi)$ has exactly $\dim(\mathbf{F}) - 1$ components. The product $\xi_1 \xi_2$ will be the usual componentwise multiplication.

$G(\mathbf{F}, n)$ is defined to be the semidirect product of $\mathbf{F}^n$ and $\text{Im}(\mathbf{F})$. The multiplication law is:

$$(\xi_1, v_1) + (\xi_2, v_2) = (\xi_1 + \xi_2, v_1 + v_2 + \text{Im}(\xi_1 \bar{\xi}_2)).$$

The horosphere (σ, d_σ) is isometric to $G(\mathbf{F}, n)$, equipped with a left invariant Riemannian metric d_σ . The restriction of d_σ to the tangent space to $\mathbf{F}^n \times \{0\}$ is the same as the Euclidean metric $\langle \xi_1, \xi_2 \rangle = \text{Re}(\xi_1 \bar{\xi}_2)$. The precise choice of d_σ depends on the normalization of the metric on $\mathbf{F}^{n+1}$.

There is a codimension $\dim(\mathbf{F}) - 1$ distribution:

$$D(\mathbf{F}, n) \subset T(G(\mathbf{F}, n)).$$

It is defined to be left invariant, and to agree with the tangent space to $\mathbf{F}^n \times \{0\}$ at the point $(0, 0)$. Any isometry which identifies σ with $G(\mathbf{F}, n)$ identifies the distribution $D(\sigma)$ with $D(\mathbf{F}, n)$.

2.6. Rank One Geometry: C-C Metric

The *Carnot-Caratheodory* distance between two points $p, q \in \sigma$ is defined to be the infimal d_σ -arc-length of any (piecewise smooth) path which joins p to q and which remains integral to the distribution $D(\sigma)$. We will denote this metric by d'_σ , and call it the C-C metric for brevity. The C-C metric is defined on $G(\mathbf{F}, n)$ in an entirely analogous way. Any isometry which identifies σ with $G(\mathbf{F}, n)$ is also an isometry in the respective C-C metrics.

Lemma 2.2. — *Let $\varepsilon > 0$ be fixed. Then there is a constant K_ε having the following property:*

$$d_\sigma(p, q) \geq \varepsilon \Rightarrow d_\sigma(p, q) \leq d'_\sigma(p, q) \leq K_\varepsilon d_\sigma(p, q);$$

K_ε only depends on ε and X .

Proof. — The first inequality is true by definition, and independent of ε . Consider the second inequality. If $d_\sigma(p, q) \in [\varepsilon, 2\varepsilon]$, then by compactness and non-integrability, there is a constant K_ε such that p and q can be joined by a (piecewise) smooth integral path having arc length at most K_ε times $d_\sigma(p, q)$. Now, suppose that $d_\sigma(p, q) > 2\varepsilon$. Let γ be the shortest path—not necessarily integral—which connects them. We can subdivide γ into intervals having length between ε and 2ε , and then replace each of these intervals by piecewise smooth integral paths having arc-length at most K_ε times as long. The concatenation of these paths is integral and has the desired arc-length. $\square$

For more details on the C-C metric, see [Gr1], [Gr2], or [P].

2.7. Rank One Geometry: Projection

Let $\sigma \subset X$ be a horosphere. We let h_σ denote the horoball which is bounded by σ , we let $b_\sigma \subset \partial X$ denote the accumulation point of σ and call b_σ the *basepoint* of σ .

Given any point $x \in X - h_\sigma$ we define

$$\pi_\sigma(x) = \overline{xb_\sigma} \cap \sigma.$$

(Here $\overline{xb_\sigma}$ is the geodesic connecting x and b_σ .) On σ , we define π_σ to be the identity map. We call π the *horospherical projection* onto σ . We will persistently use the convention that x is disjoint from h_σ . Under this convention, π_σ is distance non-increasing.

We say that two horospheres are *parallel* if they share a common base-point. Let σ^t denote the horosphere of X which is parallel to σ , contained in h_σ , and exactly t units from σ . Let $\pi_\sigma^t : \sigma \rightarrow \sigma^t$ be the restriction of π_σ to σ .

For $j = 1, 2$, let (M_j, d_j) be metric spaces. A bijection $f : M_1 \rightarrow M_2$ is called a *K-similarity* if $d_2(f(x), f(y)) = Kd_1(x, y)$ for all $x, y \in M_1$.

Lemma 2.3. — *The projection $\pi_\sigma^t : \sigma \rightarrow \sigma_t$ is an $\exp(-kt)$ -similarity, relative to the two G-C metrics. The constant $k > 0$ only depends on X .*

Proof. — Let G denote the stabilizer of σ . Then G is also the stabilizer of σ_t . Furthermore, π commutes with G , and takes $D(\sigma)$ to $D(\sigma_t)$. Since G acts transitively on $D(\sigma)$ and $D(\sigma_t)$, we see that π is an $\eta(t)$ -similarity when restricted to any linear subspace of $D(\sigma)$. It now follows from the definitions that π is an $\eta(t)$ -similarity relative to the two G-C metrics. The equality $\eta(t) = \exp(-kt)$ follows from the fact that π_σ^t is part of a one-parameter group of maps. $\square$

If σ and σ_t are both identified with $G(\mathbf{F}, n)$, then $\pi_\sigma^t = D^{-kt}$, where

$$D^r(\xi, v) = (\exp(r)\xi, \exp(2r)v).$$

The map $(\xi, v) \rightarrow \xi$ gives a natural fibration $G(\mathbf{F}, n) \rightarrow \mathbf{F}^n$. The fiber is a Euclidean space of dimension $\dim(\mathbf{F}) - 1$. The subgroup consisting of elements of the form $(0, v)$ is central in $G(\mathbf{F}, n)$, and acts by translations along the fibers. The map D^r obviously preserves this fibration structure, and induces a similarity of $\mathbf{F}^n$ relative to the metric $\langle \xi_1, \xi_2 \rangle = \operatorname{Re}(\xi_1 \bar{\xi}_2)$.

3. Detecting Boundary Components

3.1. Metric Space Axioms

In this chapter, we shall be concerned with a metric space M which satisfies the following two axioms:

Axiom 1. — There is a constant K_0 having the following property: For every point $p \in M$, there is a K_0 -quasi-geodesic in M which contains p .

Axiom 2. — There is a constant K_1 and functions $\alpha, \beta : \mathbf{R}^+ \rightarrow \mathbf{R}^+$ having the following properties: If $q_1, q_2 \in M$ are points that avoid $p \in M$ by at least $\alpha(n)$ units, then there are points q'_1 and q'_2 , and a K_1 -quasi-geodesic segment $\gamma' \subset M$, such that:

1. γ' connects q'_1 and q'_2 .
2. γ' avoids the ball of radius n about p .
3. q'_j is within $\beta(n)$ of q_j .

Here are some examples:

1. Above dimension one, Euclidean spaces satisfy both axioms.
2. In § 3.7 we will see that a horosphere of a rank one symmetric space (other than $\mathbf{H}^2$) satisfies both axioms.
3. Axiom 2 fails for all simply connected manifolds with pinched negative curvature.

3.2. Quasi-Flat Lemma

Given a subset $S \subset X$, let $T_r(S)$ denote the r -tubular neighborhood of S in X . The goal of this chapter is to prove:

Lemma 3.1. (Quasi-Flat Lemma). — Suppose that:

1. $X \neq \mathbf{H}^2$ is a rank one symmetric space.
2. $\Omega \subset X$ is a neutered space.
3. M is a metric space satisfying Axioms 1 and 2.
4. $q : M \rightarrow \Omega$ is a K -quasi-isometric embedding.

Then there is a constant K' and a (unique) horosphere $\sigma \subset \partial\Omega$ such that $q(M) \subset T_{K'}(\sigma)$. The constant K' only depends on K and on M .

Here is an overview of the proof ⁽²⁾ of the Quasi-Flat Lemma. Let $q : N \rightarrow \Omega$ be a K -quasi-isometric embedding, defined relative to a net $N \subset M$. From Axiom 1, the set $q(N)$ has at least one accumulation point on ∂X . Let L denote the set of these accumulation points. We will show, successively, that:

1. L must be a single point.
2. L must be the basepoint of a horosphere of Ω .
3. $q(N)$ must stay within K' of the horosphere based at L .

The main technical tool for our analysis is the Rising Lemma, which we will state here, but prove in § 4.

3.3. The Rising Lemma

We will use the notation established in § 2.7. Given any horosphere $\sigma \subset X$, and a subset $S \subset X - h_\sigma$, we let $|S|_\sigma$ denote the d_σ -diameter of $\pi_\sigma(S) \subset \sigma$.

⁽²⁾ Steve Gersten has independently given a proof of the Quasi-Flat Lemma for the case where M is Euclidean space, and $\Omega \subset \mathbf{H}^n$ is an equivariant neutered space. Gersten's proof involves the asymptotic cone construction.

Suppose that Ω is a neutered space. We say that S is *visible* on σ , with respect to Ω , if

$$\pi_\sigma(S) \cap \Omega \neq \emptyset.$$

Here, σ is not assumed to be a horosphere of Ω .

Lemma 3.2 (Rising Lemma). — *Let $\Omega \subset X$ be a neutered space. Let σ be an arbitrary horosphere, not necessarily belonging to Ω . Let constants $\eta, K > 0$ be given. Then there is a constant $d = d(\eta, K)$ having the following property. Suppose that*

1. γ is a K -bi-Lipschitz segment.
2. $|\gamma|_\sigma \geq \eta$.
3. γ is visible on σ .

Then $d_X(p, \sigma) \leq d$ for some $p \in \gamma$.

3.4. Multiple Limit Points

Suppose that L contains (at least) two limit points $\bar{x}$ and $\bar{y}$. Since we are assuming that Ω has at least 3 horospheres, we can choose one of them, σ , whose basepoint is neither $\bar{x}$ nor $\bar{y}$.

Recall that q is defined relative to a net $N \subset M$. Let $\{x_n\}, \{y_n\} \in N$ be sequences of points such that $q(x_n) \rightarrow \bar{x}$ and $q(y_n) \rightarrow \bar{y}$. We will choose an extremely large number d_0 , whose value is, as yet, undetermined. Let $0 \in N$ be any chosen origin. Since M satisfies Axiom 2, we can find points x'_n and y'_n such that

1. There is a uniform bound $\beta(d_0)$ from x'_n to x_n .
2. There is a uniform bound $\beta(d_0)$ from y'_n to y_n .
3. There is a K_1 -quasi-geodesic segment η_n connecting x'_n to y'_n .
4. η_n avoids 0 by d_0 units, independent of n .

Clearly $d_\Omega(q(x_n), q(x'_n))$ is uniformly bounded. Since $d_X \leq d_\Omega$, we have that $d_X(q(x_n), q(x'_n))$ is uniformly bounded. It follows that $q(x'_n) \rightarrow \bar{x}$. Likewise, $q(y'_n) \rightarrow \bar{y}$. Let $\gamma_n = q(\eta_n)$. Since Ω is a path metric space, we can assume that γ_n is (uniformly) bi-Lipschitz.

Let $B \subset \sigma$ denote the ball of d_σ -radius 1 about the point $\pi_\sigma(\bar{x})$. Let $Y = \pi_\sigma^{-1}(B)$. Let $\delta_n = \gamma_n \cap Y$. It is easy to see that $\pi_\sigma(q(x'_n)) \rightarrow \pi_\sigma(\bar{x})$. Likewise, $\pi_\sigma(q(y'_n)) \rightarrow \pi_\sigma(\bar{y})$. Hence, there is a positive constant ε such that, once n is large enough, $|\delta_n|_\sigma \geq \varepsilon$.

Let $p = q(0) \in \Omega$. Let $\Delta_r(p)$ denote the ball of d_Ω -radius r about p . For any fixed value of r , we can guarantee that γ_n does not intersect $\Delta_r(p)$ by initially choosing d_0 large enough. The sets $\{\Delta_r(p) \cap Y\}$ exhaust Y . Therefore, the distance from δ_n to σ can be made arbitrarily large, by choosing the initial constant d_0 large enough. Since δ_n is visible on σ , we get a contradiction to the Rising Lemma.

3.5. Location of the Limit Point

From the previous section, we know that L cannot be more than one point. We will show in this section that L is the basepoint of a horosphere of Ω . Since M satisfies Axiom 1, it contains at least one quasi-geodesic η . The image $\gamma = q(\eta) \subset X$ is a quasi-geodesic which limits, on both ends, to L . We may assume that γ is a bi-Lipschitz curve, since Ω is a path metric space.

Choose any point $L' \subset \partial X$ distinct from L . Let l be the geodesic in X whose two endpoints are L' and L . Assume now that L is *not* the basepoint of a horosphere of Ω . Then we can find a sequence of points $p_1, p_2, \dots \subset l$ such that

1. p_n belongs to the interior of Ω .
2. $p_n \rightarrow L$.

Let σ_n be the horosphere based at L' and containing p_n . Note that σ_n need not belong to Ω . We will let d_n be the induced path metric on σ_n . Let $B_n \subset \sigma_n$ denote the ball of d_n -radius 1 about p_n . Let $Y_n = \pi_n^{-1}(B_n)$. (Here $\pi_n = \pi_{\sigma_n}$.)

It is easy to see that the sets Y_n are nested, and

$$(*) \quad \bigcap_{n=1}^{\infty} Y_n = \emptyset.$$

Let x be any point of γ . Let γ^1 and γ^2 be the two infinite half-rays of γ divided by x . From $(*)$, both γ^1 and γ^2 must intersect ∂Y_n , for all $n > n_0$. Note, also, that $\gamma^j \cap Y_n$ is visible from σ_n , by choice of p_n . Applying the Rising Lemma, we conclude there is a uniform constant d having the following property: There is a point $x_n^j \subset \gamma^j \cap Y_n$ which is at distance at most d from σ_n .

Since the points x_n^1 and x_n^2 belong to Y_n , and are both close to σ_n , they are uniformly close to each other, independent of n . However, the length of the segment of γ connecting x_n^1 to x_n^2 tends to ∞ with n . This contradicts the fact that γ is bi-Lipschitz.

3.6. Distance to a Horosphere

We know that L is the basepoint of a horosphere σ of Ω . We will now show that $q(N)$ remains within a small tubular neighborhood of σ . Since M satisfies Axiom 1, we just have to show that any (say) K'' -bi-Lipschitz curve in Ω , which limits to L on both ends, remains within K' of σ .

Let γ be such a curve. Let $p \in \gamma$ be any point, which divides γ into two infinite rays, γ^1 and γ^2 . Similar to the previous section, the Rising Lemma says that there are points $x^j \subset \gamma^j$ which are close to each other, and close to σ . The bi-Lipschitz nature of γ bounds the length of the portion of γ connecting x^1 to x^2 , independent of p . Since p lies on this (short) segment joining x_1 to x_2 , we see that p is also close to σ .

3.7. Examples

Let $X \neq \mathbf{H}^2$ be a symmetric space. Let $\sigma \subset X$ be a horosphere. Recall from § 2.5 that σ is isometric to the nilpotent group $G(\mathbf{F}, n)$. In this section, we will show that

$G(\mathbf{F}, n)$ satisfies Axioms 1 and 2. Axiom 1 is obvious, by homogeneity. We concentrate on Axiom 2. Since the left invariant metric on $G(\mathbf{F}, n)$ is quasi-isometric to the C-C metric, we will work with the C-C metric, which is more symmetric.

Recall that there is a fibration $\rho : G(\mathbf{F}, n) \rightarrow \mathbf{F}^n$, and that C-C similarities of $G(\mathbf{F}, n)$ cover Euclidean similarities of $\mathbf{F}^n$. From this, it is easy to see:

Lemma 3.3. — *Let $\varepsilon > 0$ be fixed. If two points $q_1, q_2 \in G(\mathbf{F}, n)$ are sufficiently far away, and γ is a length minimizing geodesic segment which connects them, then the curvature of $\rho(\gamma)$ is bounded above by ε .*

Axiom 2 essentially follows from the existence of a central direction. Let $p \in G(\mathbf{F}, n)$. Let B denote the ball of radius N about p . Given any real number r , let $g_r = (0, r) \in G(\mathbf{F}, n)$. We choose r sufficiently large so that $g_r(B)$ is disjoint from B . By homogeneity, this choice does not depend on p , but only on N . Since g_r is central, $d(x, g_r(x)) \leq r'$, for all $x \in G(\mathbf{F}, n)$. Here r' only depends on r , and r in turn only depends on N .

Suppose that two points q_1, q_2 are at distance r'' from p . Let γ be a length-minimizing curve connecting q_1 to q_2 . We choose r'' so large that the projection $\rho(\gamma)$ has very small curvature. More precisely, we choose r'' so that $\gamma \cap \rho^{-1}(\rho(B))$ consists of disjoint segments $\gamma_1, \gamma_2, \dots$ having the following properties:

1. At most one γ_j intersects B .
2. The distance from γ_i to γ_j is at least $100r$.

(We picture γ as a portion of a helix, and $\rho^{-1}(\rho(B))$ as an infinite solid tube, intersecting this helix transversally.) Our two conditions imply that $g_s(\gamma)$ is disjoint from B for some $s \leq r$. Also, $d(q_j, g_r(q_j)) \leq r'$. We get Axiom 2 if we set

$$q'_j = g_s(q_j); \quad \alpha(N) = r''; \quad \beta(N) = r'; \quad K_1 = 1.$$

4. The Rising Lemma

In this chapter, we will prove the Rising Lemma, which was stated in § 3.3.

4.1. Geodesics and Horospheres

Let X be a rank-one symmetric space. In this section, we give some estimates concerning the interaction of geodesics and horoballs in X . A theme implicit in our discussion below, and worth making explicit here, is that horoballs in X are convex with respect to d_X .

Here is a piece of notation we will use often: Let $|\gamma|'_\sigma$ denote the d'_σ diameter of γ . Here, d'_σ is the C-C metric on σ . We say that γ and σ are *tangent* if they intersect

in exactly one point. In this case, all other points of γ are disjoint from the horoball h_σ . The following lemma is well known:

Lemma 4.1. — *Let γ be a geodesic and σ a horosphere. Suppose that γ is tangent to σ . Then*

$$K^{-1} \leq |\gamma|'_\sigma \leq K,$$

where K only depends on X .

Proof. — Compactness and equivariance. $\square$

Lemma 4.2. — *Let γ be a geodesic and σ a horosphere. Suppose that*

$$d_X(\gamma, \sigma) \geq n.$$

Then

$$|\gamma|'_\sigma \leq K \exp(-n/K),$$

where K only depends on X .

Proof. — Combine Lemma 4.1 and Lemma 2.3. $\square$

Lemma 4.3. — *Let σ be a horosphere. Let γ' be a geodesic segment joining two points $x_1, x_2 \in \sigma$. Let $\gamma'_s \subset \gamma'$ denote the points which are at least s units away from both x_1 and x_2 . Then*

$$d_X(\gamma'_s, \sigma) \geq s - K,$$

provided that γ'_s is nonempty. Here, the constant K only depends on X .

Proof. — Let σ' be the horosphere parallel to σ and tangent to γ' at a single point, ξ . Let β'_j denote the geodesic segment connecting x_j to $\pi_{\sigma'}(x_j)$. Let γ'_j be the portion of γ' connecting x_j to ξ . Using Lemma 4.1, and the usual comparison theorems, γ'_j remains within the K -tubular neighborhood of β'_j . Hence, γ'_j moves away from σ at the same linear rate that β'_j does, up to the constant K . $\square$

The following Lemma is obvious for real hyperbolic space, and actually a bit surprising in the general case.

Lemma 4.4. — *Let σ_1 and σ_2 be two horospheres. Suppose that $x_j, y_j \in \sigma_j$ satisfy*

$$d_X(x_1, y_1) \leq W + d_X(x_2, y_2).$$

Then

$$d_{\sigma_1}(x_1, y_1) \leq W'' d_{\sigma_2}(x_2, y_2),$$

where W'' only depends W and on X .

Proof. — To avoid trivialities, we will assume that $d_X(x_j, y_j) \geq 1$. Below, the positive constants $K_1, K_2, \dots$ have the desired independence. Let γ_j be the geodesic segment

in X which connects x_j to y_j . Let σ'_j denote the horosphere parallel to σ_j and tangent to γ'_j . From Lemma 4.1,

$$1/K_1 \leq |\gamma_j|'_{\sigma_j} \leq K_1.$$

(The first inequality follows from the assumption that $d_X(x_j, y_j) \geq 1$.) From this, it follows that

$$2d_X(\sigma_j, \sigma'_j) \leq d_X(x_j, y_j) \leq 2d_X(\sigma_j, \sigma'_j) + K_1.$$

Therefore

$$d_X(\sigma_1, \sigma'_1) \leq d_X(\sigma_2, \sigma'_2) + K_2.$$

From Lemma 2.3, we see that

$$(*) \quad d'_\sigma(x_1, y_1) \leq K_3 d'_\sigma(x_2, y_2).$$

By assumption, $d_\sigma(x_j, y_j) \geq 1$. Now apply Lemma 2.2 to (*). $\square$

Let σ be a horosphere. Let h_σ be the closed horoball bounded by σ .

Lemma 4.5. — *Let σ be a horosphere. Suppose that $S \subset X - h_\sigma$ is closed, and $\pi_\sigma(S)$ is compact. Let Z be a horoball which intersects S but does not contain $\pi_\sigma(S)$. Then $\partial Z \cap \sigma$ has diameter at most K_1 , independent of Z and σ .*

Proof. — If this was false, let $Z_1, Z_2, \dots$ be a sequence forming a counterexample. Since none of these horoballs contains all of $\pi_\sigma(S)$, and since this set is compact, no subsequence of these horoballs can converge to all of X . Furthermore, since Z_n intersects σ and S , no subsequence can converge to the empty set. Hence, some subsequence converges to a horoball Z_∞ of X . The intersection $Z_\infty \cap \sigma$ has infinite diameter. This is impossible unless $Z_\infty = \sigma$. This contradicts the fact that Z_n always intersects S . $\square$

Let σ_t denote the horosphere parallel to σ , disjoint from h_σ , and exactly t units away from σ . The following result is obvious in the real hyperbolic case, but requires work in general:

Lemma 4.6. — *Let Z be a horoball. Then*

$$\text{diam}_X(\partial Z \cap \sigma) \leq K_1 \Rightarrow |\partial Z \cap \sigma_t|'_\sigma \leq K_2 \exp(-t/K_2),$$

where K_2 only depend on X and on K_1 .

Proof. — Let $\tau = \partial Z$. By moving σ parallel to itself by at most K_1 units, we can assume that h_σ and h_τ are disjoint. This move only changes the constants in the estimates by a uniformly bounded factor. Let γ be a geodesic joining two points $x_1, x_2 \in \tau \cap \sigma_t$.

Sub-Lemma 4.7. — We have $d_{\mathbf{x}}(\gamma, \sigma) \geq t/2 - K$.

Proof. — The endpoints of γ are t units away from σ . If our sub-lemma is false, then there is a point $p \in \gamma$ for which $d_{\mathbf{x}}(p, \sigma) \leq t/2 - R_t$, for unboundedly large R_t . By convexity, $\gamma \subset h_\tau$. Since h_τ and σ are disjoint,

$$(*) \quad d_{\mathbf{x}}(p, \tau) \leq t/2 - R_t.$$

On the other hand, since γ comes within $t/2 - R_t$ of σ , its length must be at least t . (Recall the location of the endpoints). This says—in the notation of Lemma 4.3—that $\gamma_{t/2}$ is nonempty. Lemma 4.3 now contradicts Equation (*) for large R_t . $\square$

It follows from Lemma 4.2 that

$$d'_\sigma(x, y) \leq K_2 \exp(-t/K_2).$$

But x_1 and x_2 were arbitrary points in $\tau \cap \sigma_t$. $\square$

4.2. Shading

Let φ be a horosphere. We say that a point $x \notin h_\varphi$ is *s-shaded* with respect to φ if $\pi_\varphi(x) \notin \Omega$, and $d_\varphi(\pi_\varphi(x), \partial\Omega \cap \varphi) \geq s$.

Lemma 4.8. — Suppose that

1. $x \in \Omega$.
2. $d_{\mathbf{x}}(x, \varphi) \leq W$.
3. x is not *s-shaded* with respect to φ .

Then $d_\Omega(x, \varphi) \leq W'(x + 1)$, where the constant W' only depends on X and on W .

Proof. — The case where $\pi_\varphi(x) \in \Omega$ is trivial. So, we will assume that there is a horoball Z of Ω which contains $\pi_\varphi(x)$. By hypothesis, there is a point $y \in \partial Z \cap \varphi$ such that $d_\varphi(y, \pi_\varphi(x)) \leq s$. Since $x \notin Z$, there is a point $z \in \partial Z$ such that $\pi_\varphi(z) = \pi_\varphi(x)$, and $d_{\mathbf{x}}(x, z) + d_{\mathbf{x}}(z, \varphi) = W$.

By the triangle inequality, we have

$$d_{\mathbf{x}}(y, z) \leq d_{\mathbf{x}}(y, \pi_\varphi(x)) + W.$$

From Lemma 4.4, there is a path on ∂Z which connects y to z , having arc-length at most $W''(s + 1)$. Hence, $d_\Omega(y, z) \leq W''(s + 1)$. Also, $d_\Omega(x, z)$ is uniformly bounded, in terms of W . The result now follows from the triangle inequality, and an appropriate choice of W' . $\square$

Let α be any (compact) curve in Ω . Let $\partial_1 \alpha$ and $\partial_2 \alpha$ denote the two endpoints of α . We say that α is *W-controlled* by the horosphere φ provided that

1. α is disjoint from h_φ .
2. For any $p \in \alpha$, $d_X(p, \varphi) \geq W$.
3. $d_X(\partial_j \alpha, \varphi) \leq 2W$.

We say that α is *λ -fractionally shaded* with respect to φ provided that at least one endpoint of α is $\lambda |\alpha|_\varphi$ shaded with respect to φ .

Lemma 4.9. — *Suppose α is a K-bi-Lipschitz segment. Then there are positive constants λ, W, L having the following property: If α is W-controlled by φ and $|\alpha|_\varphi \geq L$, then α is λ -fractionally shaded with respect to φ . These three constants only depend on K and on X .*

Proof. — For the present, the three constants W, L and λ will be undetermined. We will let $A(*)$ denote the arc-length.

Let $\Lambda = |\alpha|_\varphi \geq L$. For points at least W away from φ , the projection π_φ decreases distances exponentially. This is to say that

$$A(\alpha) \geq \exp(k' W) \Lambda.$$

The constant k' only depends on the symmetric space X .

We set $\partial_j = \partial_j \alpha$. Suppose that α is not λ -fractionally shaded with respect to φ . Then, according to Lemma 4.8, there are paths $\beta_j \subset \Omega$ connecting ∂_j to points $p_j \subset \varphi$ such that

$$A(\beta_j) \leq W'(\lambda \Lambda + 1).$$

Since the projection onto φ is distance non-increasing, we have, by the triangle inequality,

$$d_\varphi(p_1, p_2) \leq W'(\lambda \Lambda + 1) + \Lambda + W'(\lambda \Lambda + 1).$$

Let γ be the geodesic segment in X connecting p_1 to p_2 . Every time γ intersects a horoball Z of Ω , we replace $\gamma \cap Z$ by the shortest path on ∂Z having the same endpoints as $\gamma \cap \partial Z$. Call the resulting path δ . By Lemma 4.4,

$$A(\delta) \leq K_0 d_\varphi(p_1, p_2).$$

Here K_0 is a universal constant, only depending on X .

The path $\eta = \beta_1 \cup \delta \cup \beta_2 \subset \Omega$ is a rectifiable curve connecting the endpoints of α . Combining the previous equations, we get

$$A(\eta) \leq K_1 W' \lambda \Lambda + K_1 W' + K_1 \Lambda.$$

The constant K_1 is a universal constant, depending only on X . For whatever value of W we choose, we will take

$$L > W', \quad \lambda < \frac{1}{W'}.$$

With these choices, we get

$$A(\eta) \leq 3K_1 \Lambda \leq 3K_1 \exp(-k' W) A(\alpha).$$

For sufficiently large W , this contradicts the bi-Lipschitz constant of α . $\square$

4.3. Main Construction

We will use the notation from the Rising Lemma. Let $p \in \gamma$ be a point such that $\pi_\sigma(p) \in \Omega$. Let B denote the ball of d_σ -radius η about $\pi_\sigma(p)$. Let $Y = \pi_\sigma^{-1}(B)$. Let δ be the component of $\gamma \cap Y$ which contains p . Choose two points $x_0, y \in \delta$ such that

1. $|\{x_0, y\}|_\sigma \geq \eta/2$.
2. x_0 is a point of δ which minimizes the d_X -distance to σ .

We delete portions of δ so that x_0 and y are the endpoints. We orient δ so that it progresses from x_0 to y . By construction, x_0 is a point of δ which is d_X -closest to σ . Let N denote this distance.

Choose the constants W, L and λ as in Lemma 4.9. For $j = 1, 2, \dots$, we inductively define x_j to be the last point along δ which is at most $jW + N$ units away from σ . (It is possible that $x_1 = y$.) Let δ_j denote the segment of δ connecting x_{j-1} to x_j . By construction, δ_j is W -controlled by the horosphere

$$\varphi_j = \sigma_{(j-2)W+N}.$$

We insist that $N \geq 2W + 1$. This guarantees that φ_j is disjoint from h_σ , and at least one unit away from σ .

Below, the constants $K_1, K_2, \dots$ have the desired dependence. Recall that d'_σ is the C-C metric on σ . Let $|S|'_\sigma$ denote the d'_σ diameter of $\pi_\sigma(S)$. Note that $|S|_\sigma \leq |S|'_\sigma$.

Sub-Lemma 4.10. — *We have*

$$|\delta_j|'_\sigma \leq K_1 \exp(-N/K_1) \exp(-j/K_1).$$

Proof. — If $|\delta_j|_{\varphi_j} \leq L$, then, by compactness, $|\delta_j|'_{\varphi_j} \leq K_2$. The bound now follows immediately from Lemma 2.3. Suppose that $|\delta_j|_{\varphi_j} \geq L$. Then δ_j is λ -fractionally shaded, from Lemma 4.9. Let Z be the horoball of Ω with respect to which δ_j is λ -fractionally shaded. The shading condition says:

$$L \leq |\delta_j|_{\varphi_j} \leq \lambda^{-1} |\partial Z \cap \varphi_j|_{\varphi_j}.$$

Lemma 2.2 says that

$$(*) \quad |\delta_j|'_{\varphi_j} \leq K_3 |\partial Z \cap \varphi_j|'_{\varphi_j}.$$

Let $S \subset Y$ denote the set of points which are at least one unit away from h_σ . Since Z shades α_j with respect to φ_j , we see that Z intersects S . Also, Z cannot contain $B = \pi_\sigma(S)$, because $p \in B \cap \Omega$. Lemma 4.5 says therefore that

$$\text{diam}(\partial Z \cap \sigma) \leq K_4.$$

From Lemma 4.6 we conclude that

$$|\partial Z \cap \varphi_j|'_\sigma \leq K_5 \exp(-N/K_5) \exp(-j/K_5).$$

The desired bound now follows from (*) and Lemma 2.3. $\square$

Summing terms from both cases over all j , we see that

$$|\delta|_\sigma \leq |\delta|'_\sigma \leq K_6 \exp(-N/K_1).$$

For sufficiently large N , this contradicts the choice of x_0 and y .

5. Ambient Extension

5.1. Overview

Suppose that $\bar{q} : X \rightarrow X$ is a quasi-isometry, defined relative to nets $\bar{N}_1$ and $\bar{N}_2$. We say that $\bar{q}$ is *adapted* to the pair (Ω_1, Ω_2) provided that:

1. No two distinct points of $\bar{N}_j$ are within one unit of each other.
2. $\bar{q}$ is a bi-Lipschitz bijection between $\bar{N}_1$ and $\bar{N}_2$.
3. If $x \in \bar{N}_j \cap \Omega_j$, then $\bar{q}(x) \in \bar{N}_{j+1} \cap \Omega_{j+1}$.
4. Let $\sigma_j \subset \partial\Omega_j$ be a horosphere. Then $\bar{N}_j \cap \sigma_j$ is a K -net of σ_j , where the constant K does not depend on σ_j .
5. For each horosphere $\sigma_j \subset \partial\Omega_j$, there is a horosphere $\sigma_{j+1} \subset \partial\Omega_{j+1}$ such that $\bar{q}(\bar{N}_j \cap \sigma_j) = \bar{N}_{j+1} \cap \sigma_{j+1}$.

To simplify notation, we have taken the indices mod 2 above, and also blurred the distinction between $\bar{q}$ and $\bar{q}^{-1}$. Hopefully, this does not cause any confusion.

The goal of this chapter is to prove the following:

Lemma 5.1 (Ambient Extension). — *Suppose that $q : \Omega_1 \rightarrow \Omega_2$ is a quasi-isometry. Then there is a quasi-isometry $\bar{q} : X \rightarrow X$ such that*

1. $\bar{q}$ is adapted to (Ω_1, Ω_2) .
2. $\bar{q}|_{\Omega_j}$ is equivalent to q , relative to d_{Ω_j} .

Remark. — The map $\bar{q}$ is a quasi-isometry relative to d_X . The restriction $\bar{q}|_{\Omega_j}$ is being considered as a quasi-isometry relative to d_{Ω_j} . In fact, since $\bar{q}$ is adapted to (Ω_1, Ω_2) , this restricted map is a quasi-isometry relative to *both* d_{Ω_j} and $d_X|_{\Omega_j}$.

5.2. Cleaning Up

We can add points to N_1 so that $N_1 \cap \sigma_1$ is a uniform net of σ_1 , for every horosphere $\sigma_1 \subset \Omega_1$. Given a horosphere $\sigma_1 \subset \Omega_1$, there is, by the Quasi-Flat Lemma, a unique horosphere σ_2 such that $q(N_1 \cap \sigma_1)$ remains within a uniformly thin tubular neighborhood of σ_2 . We define a new quasi-isometry q' , as follows: For each point $x \in \sigma_1$, let

$$q'(x) = \pi_{\sigma_2}(q(x)).$$

Doing this for every horosphere of Ω_1 , we obtain a quasi-isometry q' which is equivalent to q , and which has the property that $q'(\sigma_1 \cap N_1) \subset \sigma_2$. We set $q = q'$.

Lemma 5.2. — *The image $q(\sigma_1 \cap N_1)$ is a uniform net of σ_2 .*

Proof. — The map q^{-1} is well defined, up to bounded modification. Applying the Quasi-Flat Lemma to q^{-1} , we see that every point near σ_2 is near a point of the form $y = q(x)$, where x is near σ_1 . Our lemma now follows from the triangle inequality. $\square$

Since q' is a quasi-isometry, there is a constant λ_0 having the following property: If $x, y \in N_1$ are at least λ_0 apart, then $q'(x)$ and $q'(y)$ are at least one unit apart. By thinning the net N_1 as necessary, we obtain nets N'_1 and $N'_2 = q'(N'_1)$ having the following properties:

1. No two distinct points of N'_j are within 1 unit of each other.
2. q' is a bi-Lipschitz bijection from N'_1 to N'_2 .
3. $q'(N'_j - \partial\Omega_j) = N'_{j+1} - \partial\Omega_{j+1}$.
4. Let $\sigma_j \subset \partial\Omega_j$ be a horosphere. Then $N'_j \cap \sigma_j$ is a uniform net of σ_j .
5. For each horosphere $\sigma_j \subset \partial\Omega_j$, there is a horosphere $\sigma_{j+1} \subset \partial\Omega_{j+1}$ such that $q'(N'_j \cap \sigma_j) = N'_{j+1} \cap \sigma_{j+1}$.

5.3. Main Construction

For the rest of § 5, we adopt the notation of § 2.7. We now construct certain special nets which extend N'_1 and N'_2 . For simplicity, we will drop subscripts.

Let $x \in N'$. If $x \in N' - \partial\Omega$, let $S_x = \{x\}$. If $x \in \partial\Omega$, then x belongs to a horosphere $\sigma \subset \partial\Omega$. Let S_x be the infinite ray joining x to the basepoint $b_o \in \partial X$. Define

$$\hat{N} = \bigcup_{x \in N'} S_x.$$

Clearly $\hat{N}$ is a net of X .

For each point $x \in N'_1$, there is an obvious isometric bijection from S_x to $S_{q'(x)}$. The union of these isometric bijections gives a bijection

$$\hat{q} : \hat{N}_1 \rightarrow \hat{N}_2.$$

Lemma 5.3. — *Let $C > 0$ be given. Then there exists a constant C' having the following property: If $x, y \in \hat{N}_1$ satisfy $d_X(x, y) < C$, then $d_X(\hat{q}(x), \hat{q}(y)) < C'$.*

Proof. — Below, the constants $C_1, C_2, \dots$ depend on (C, q', N'_1, N'_2) . For each point $p \in X$, let $\delta(p)$ denote the minimum distance from p to Ω . For positive k , the connected components of the level set $\delta^{-1}(k)$ are horospheres. Say that two points $p, q \in \hat{N}_1$

are *horizontal* if they belong to the same connected level set of δ . Say that two points of $\hat{N}_1$ are *vertical* if they both belong to the same set S_x . Now suppose that $d_X(x, y) \leq C$. It is clear from the construction of $\hat{q}$ that there are points $x = p_0, p_1, p_2, p_3 = y \in \hat{N}_1$ such that

1. p_0 and p_1 are vertical.
2. p_1 and p_2 are horizontal.
3. p_2 and p_3 are vertical.
4. $d_X(p_j, p_{j+1}) \leq C_1$.

By construction,

$$d_X(\hat{q}(p_0), \hat{q}(p_1)) \leq C_1, \quad d_X(\hat{q}(p_2), \hat{q}(p_3)) \leq C_1.$$

For the points p_1, p_2 , there are two cases:

Case 1. — If $p_1, p_2 \in \Omega_1$, then $d_{\Omega_1}(p_1, p_2) \leq C_2$, since d_X and d_{Ω} are Lipschitz equivalent below any given scale. Since q is a quasi-isometry relative to d_{Ω_j} , we have that $d_{\Omega_2}(\hat{q}(p_1), \hat{q}(p_2)) \leq C_3$. Since $d_X|_{\Omega_2} \leq d_{\Omega_2}$, the same bound holds for $d_X(\hat{q}(p_1), \hat{q}(p_2))$.

Case 2. — Suppose that $p_1, p_2 \notin \Omega_1$. Then there is a horosphere $\sigma_1 \subset \partial\Omega_1$ and a number d such that $p_1, p_2 \in \sigma_1^d$. Let $\sigma_2 \subset \partial\Omega_2$ be the horosphere which is paired to σ_1 , via q . From Lemma 2.2, the map

$$q'|_{N'_1 \cap \sigma_1} : N'_1 \cap \sigma_1 \rightarrow N'_2 \cap \sigma_2$$

is uniformly bi-Lipschitz relative to the C-C metrics on these two horospheres. From Lemma 2.3, therefore,

$$d'_2(\hat{q}(p_1), \hat{q}(p_2)) \leq C_4 d'_1(p_1, p_2).$$

Here d'_j is the C-C metric on σ_j^d . (Recall that σ^d is the horosphere parallel to σ , contained in h_σ , and d units away from σ .) Since $d_X(p_1, p_2) \leq C_1$, it follows from compactness that

$$d'_1(p_1, p_2) \leq C_5.$$

Finally,

$$d_X(\hat{q}(p_1), \hat{q}(p_2)) \leq d'_2(\hat{q}(p_1), \hat{q}(p_2)).$$

Putting everything together gives $d_X(\hat{q}(p_1), \hat{q}(p_2)) \leq C_6$.

The triangle inequality completes the proof. $\square$

Lemma 5.3 also applies to the inverse map $\hat{q}^{-1}$. Since X is a path metric space, it follows that $\hat{q}$ is a quasi-isometry. To finish the proof of the Ambient Extension Lemma, we thin out the nets $\hat{N}_j$ appropriately.

6. Geometric Limits: Real Case

Let $q : X \rightarrow X$ be a quasi-isometry, adapted to the pair (Ω_1, Ω_2) of neutered spaces. From this point onward, we will assume that Ω_1 and Ω_2 are equivariant neutered spaces. It is the goal of § 6-§ 8 to prove

Lemma 6.1 (Rigidity Lemma). — Suppose that q is a quasi-isometry of X which is adapted to the pair (Ω_1, Ω_2) of equivariant neutered spaces. Then q is equivalent to an isometry of X .

In this chapter, and the next, we will work out the real hyperbolic case. In § 8, we will make the modifications needed for the complex case. The other cases follow immediately from [P, Th. 1].

6.1. Quasiconformal Extension

Let $\mathbf{H} = \mathbf{H}^n$ be real hyperbolic space, for some $n \geq 3$. We will use the upper half-space model for $\mathbf{H}$, and set $\mathbf{E} = \partial\mathbf{H} - \infty$. Finally, we will let $T_x(\mathbf{E})$ denote the tangent space to $\mathbf{E}$ at x . It is well known that q has a quasi-conformal extension $h = \partial q$. (See [M2], or [T, Ch. 5].) We normalize so that $h(\infty) = \infty$. It is also known that

1. h is a.e. differentiable ⁽³⁾ on $\mathbf{E}$, and this differential is a.e. nonsingular [M2, Th. 9.1].

Let $dh(x)$ denote the linear differential at x .

2. If $dh(x)$ is a similarity for almost all x , then h is a conformal map. This is to say that q is equivalent to an isometry of $\mathbf{H}$ [M2, Lemma 12.2].

This chapter is devoted to proving:

Lemma 6.2 (Real Case). — Let q be a quasi-isometry of $\mathbf{H}$ which is adapted to the pair (Ω_1, Ω_2) of equivariant neutered spaces. Suppose that $h = \partial q$ fixes ∞ . Let $x \in \mathbf{E}$ be a generic point of differentiability for h . Then there are isometric copies Ω'_j of Ω_j , and a quasi-isometry $q' : X \rightarrow X$ such that

1. q' is adapted to the pair (Ω'_1, Ω'_2) .
2. $h' = \partial q'$ is a real linear transformation of $\mathbf{E}$.
3. $h' = dh(x)$, under the canonical identification of $\mathbf{E}$ and $T_x(\mathbf{E})$.

6.2. Hausdorff Topology

Let M be a metric space. The *Hausdorff distance* between two compact subsets $K_1, K_2 \subset M$ is defined to be the minimum value $\delta = \delta(K_1, K_2)$ such that every point of K_j is within δ of a point of K_{j+1} . (Indices are taken mod 2.) A sequence of closed

⁽³⁾ For the purist, an additional trick can make our proof work under the easier assumption that h is just ACL.

subsets $S_1, S_2, \dots \subset M$ is said to converge to $S \subset M$ in the *Hausdorff topology* if, for every compact set $K \subset M$, the sequence $\{\delta(S_n \cap K, S \cap K)\}$ converges to 0.

We say that a net $N \subset \mathbf{H}$ is *sparse* provided that no two distinct points of N are within 1 unit of each other. We say that a quasi-isometry $q : \mathbf{H} \rightarrow \mathbf{H}$ is *sparse* if it is a bi-Lipschitz bijection between sparse nets N_1 and N_2 . Let $\Gamma(q) \subset \mathbf{H} \times \mathbf{H}$ denote the graph of q .

Let q^n be a sparse quasi-isometry defined relative to sparse nets N_1^n and N_2^n . We say that $q^1, q^2, \dots$ converges to a sparse quasi-isometry q , defined relative to nets N_j , provided that

1. N_j^n converges to N_j in the Hausdorff topology.
2. $\Gamma(q^n)$ converges to $\Gamma(q)$ in the Hausdorff topology.

We will make use of several compactness results:

Lemma 6.3. — *Let $\{q^n\}$ be a sequence of sparse K -quasi-isometries of $\mathbf{H}$. Let $h^n = \partial q^n$ be the extension of q^n . Suppose that*

1. $h^n(0) = 0$.
2. $h^n(\mathbf{E}) = \mathbf{E}$.
3. h^n converges uniformly on compacta to a homeomorphism $h : \mathbf{E} \rightarrow \mathbf{E}$.

Then the maps q^n converge on a subsequence to a sparse quasi-isometry q . Furthermore $\partial q = h$.

Proof. — Let 0 be any chosen origin of hyperbolic space. We will first show that the set $\{q^n(0)\}$ is bounded. Let γ_1 and γ_2 be two distinct geodesics through 0. Then the quasi-geodesics $q^n(\gamma_j)$ remain within uniformly thin tubular neighborhoods of geodesics δ_1^n and δ_2^n . (This is a standard fact of hyperbolic geometry.)

The endpoints of δ_1^n and δ_2^n converge to four distinct points of $\partial \mathbf{H}$. Furthermore, the point $q^n(0)$ must lie close to both δ_1^n and δ_2^n . This implies that $\{q^n(0)\}$ is bounded. Statement 1 now follows from a routine diagonalization argument.

By thinning out the sequence, we can assume that q^n converges to a quasi-isometry q^∞ . Let $h^\infty = \partial q^\infty$. Let p be any point in $\mathbf{E}$. Let γ be any geodesic, one of whose endpoints is p . Let δ^n denote the geodesic whose tubular neighborhood contains $q^n(\gamma)$. Then δ^n converges to some geodesic δ^∞ , in the Hausdorff topology. Hence the endpoints of δ^n converge to those of δ^∞ . This means that $h^n(p)$ converges to $h^\infty(p)$. Hence $h(p) = h^\infty(p)$. Since p is arbitrary, we get Statement 2. $\square$

Lemma 6.4. — *Let I^n be a sequence of hyperbolic isometries. Let Ω be an equivariant neutered space. Let $\Omega^n = I^n(\Omega)$. Suppose that $\cap \Omega^n$ is nonempty. Then, on a subsequence, these neutered spaces converge to an isometric copy Ω' of Ω .*

Proof. — Let $K^n \subset \Omega^n$ denote a compact fundamental domain for Ω^n , modulo its isometry group. Since $\cap \Omega^n$ is nonempty, we can choose K^n so that $\cap K^n$ is nonempty.

Note that K^n has uniformly bounded diameter. Hence, we can choose isometries $J^1, J^2, \dots$ such that

1. $J^n(\Omega) = I^n(\Omega)$.
2. The sequence $\{J^j\}$ lies in a compact subset of the isometry group of $\mathbf{H}$.

From these two statements, the claims of the Lemma are obvious. $\square$

6.3. Taking a Derivative

Suppose that $f: \mathbf{E} \rightarrow \mathbf{E}$ is a homeomorphism which is differentiable at the origin. Suppose also that the differential $df(0)$ is a nonsingular linear transformation of the tangent space $T_0(\mathbf{E})$. Let D^n denote the dilation

$$v \mapsto \exp(n) v.$$

Consider the sequence of maps

$$f^n = D^n \circ f \circ D^{-n}.$$

It is a standard fact from several variable calculus that f^n converges, uniformly on compacta, to a linear transformation f^∞ , and that f^∞ equals $df(0)$, under the canonical identification of $T_0(\mathbf{E})$ with $\mathbf{E}$. We will call this the *differentiability principle*, and will use it below.

6.4. Zooming In

We will use the notation established above. Suppose that q is a hyperbolic K -quasi-isometry satisfying the hypothesis of Lemma 6.2. By translation we can assume that the point x is the origin, $0 \in \mathbf{E}$. By further translation, we can assume that $h(0) = 0$. Since x is generic, we can assume that 0 is *not* the basepoint of a horosphere of Ω_1 . Since q is adapted to the pair (Ω_1, Ω_2) , we see that 0 is not the basepoint of a horosphere of Ω_2 either.

Recall that D^r is the dilation by $\exp(r)$ discussed in the preceding section. Let T^r denote the hyperbolic isometry which extends D^r . Consider the following sequence of objects:

1. $q^r = T^r \circ q \circ T^{-r}$.
2. $\Omega_j^r = T^r(\Omega_j)$.
3. $h^r = \partial q^r$.

From the differentiability principle, the maps h^r converges to the nonsingular linear map $h' = dh(0)$, as $r \rightarrow \infty$. Let l be the geodesic connecting 0 to ∞ . Since 0 is not the basepoint of a horosphere of Ω_1 , we can find points $p_1, p_2, \dots \in l \cap \Omega_1$ which converge to 0 . By choosing r appropriately, we extract a subsequence $r_1, r_2, \dots$ such

that $T^{r_n}(p_n) = p_1$. By Lemma 6.4, $T^{r_n}(\Omega_1)$ converges to a neutered space Ω'_1 isometric to Ω_1 .

Every quasi-isometry in sight is sparse. By Lemma 6.3, the maps q^r converge, on a thinner subsequence, to a quasi-isometry q' with $\partial q' = h'$. Since the maps q^{r_n} converge, the points $q^{r_n}(p_1)$ remain in a compact subset. Hence, for some thinner subsequence (labelled the same way) there is some point p_1^* which belongs to every neutered space $\Omega_2^{r_n} = q^{r_n}(\Omega_1^{r_n})$. By Lemma 6.4 there is a thinner subsequence on which these neutered spaces converge to a limit Ω'_2 which is isometric to Ω_2 .

Note that q' is adapted to the pair (Ω'_1, Ω'_2) . Since everything in sight converges, q' is adapted to the pair (Ω'_1, Ω'_2) . This establishes Lemma 6.2.

7. Inversion Trick: Real Case

The goal of this chapter is to prove the Rigidity Lemma stated in § 6, in the real hyperbolic case. Using the facts about quasi-conformal maps of the Euclidean space listed in § 7.1, and Lemma 6.2, we just have to prove:

Lemma 7.1 (Real Case). — Suppose that q is a quasi-isometry of $\mathbf{H}$ which is adapted to a pair (Ω_1, Ω_2) of equivariant neutered spaces. Suppose also that $h = \partial q$ is a real linear transformation when restricted to $\mathbf{E}$. Then h is a similarity on $\mathbf{E}$.

The remainder of this chapter is devoted to proving this result. The technique is somewhat roundabout, since the map q is not assumed to conjugate (or virtually conjugate) the isometry group of Ω_1 to that of Ω_2 .

7.1. Inverted Linear Maps

Let $T : \mathbf{E} \rightarrow \mathbf{E}$ be a real linear transformation. Let I denote inversion in the unit sphere of $\mathbf{E}$. Technically, I is well-defined only on the one-point compactification $\mathbf{E} \cup \infty$. Alternatively, I is well-defined and conformal on $\mathbf{E} - \{0\}$. We will call the map $I \circ T \circ I$ an *inverted linear map*.

There are two possibilities. If T is a similarity, then so is $I \circ T \circ I$. However, if T is not a similarity, then $I \circ T \circ I$ is quite strange. The key to our proof of the Lemma 7.1 is a careful analysis of the map $I \circ T \circ I$, when T is *not* a similarity. For notational convenience, we will set $\underline{T} = I \circ T \circ I$.

We will say that a *dilation* of $\mathbf{E}$ is any map of the form $v \mapsto \lambda v$. Here, λ is a scalar, and v is a vector.

Lemma 7.2. — The transformation $\underline{T}$ commutes with the one-parameter subgroup of dilations. Furthermore, $\underline{T}$ is bi-Lipschitz on $\mathbf{E} - \{0\}$.

Proof. — Let D^n denote the dilation by $\exp(n)$. Then, we have $D^n \circ I = I \circ D^{-n}$. Also, D^n commutes with T . These two facts imply that $\underline{T}$ commutes with dilations.

Now, $\underline{T}$ is clearly (say) K -bi-Lipschitz when restricted to a thin annulus containing the unit sphere. Since $\underline{T}$ commutes with dilations, it must in fact be K -bi-Lipschitz on the image of this annulus under an arbitrary dilation. $\square$

7.2. Images of Orbits

In this section, we will use the vector space structure of $\mathbf{E}$. A *translation* of $\mathbf{E}$ is a map of the form $f_v(x) = x + v$, where $v \in \mathbf{E}$ is a vector. Let Λ be a co-compact lattice of translations of $\mathbf{E}$. Once and for all, we fix a basis $\{v_1, \dots, v_d\}$ for Λ . Here d is the dimension of $\mathbf{E}$.

Let G be an orbit of Λ . We will say that a *cell* of G is a collection of 2^d vertices made up of the form:

$$g + \varepsilon_1 v_1 + \dots + \varepsilon_d v_d.$$

Here $g \in G$, and ε_j can either be 0 or 1. The orbit G is made up of a countable union of cells, all of which are translation equivalent. Each cell consists of the vertices of a parallelepiped. Say that two cells of G are *adjacent* if they have nonempty intersection. Clearly, two cells can be joined by a sequence of adjacent cells.

We say that two compact sets $S_1, S_2 \subset \mathbf{E}$ are ε -translation equivalent if there is a translation f_v such that $f_v(S_1)$ and S_2 are ε -close in the Hausdorff topology. Let $B_n \subset \mathbf{E}$ denote the n -ball about 0.

Lemma 7.3. — *Let $\varepsilon > 0$ be fixed. Then for sufficiently large n , the following is true: If C_1 and C_2 are adjacent cells of G , and disjoint from B_n , then $\underline{T}(C_1)$ and $\underline{T}(C_2)$ are ε -translation equivalent. The constant n only depends on T , and on ε .*

Proof. — Let γ be any ray emanating from 0. Let ρ be a number significantly larger than the diameter of a cell. Let Δ_r denote the ball of radius ρ centered on the point of γ which is $\exp(r)$ units away from 0. Note that the ball $\Delta'_r = D^{-r}(\Delta_r)$ is centered on a point of the unit sphere, and has radius tending to zero as r tends to ∞ .

From Lemma 7.2, we know what $\underline{T}$ commutes with dilations. Hence, in particular

$$\underline{T}|_{\Delta_r} = D^r \circ \underline{T}|_{\Delta'_r} \circ D^{-r}.$$

Note that $\underline{T}$ is differentiable on the unit sphere. Hence, using the differentiability principle of § 6.4, we see that, pointwise, $\underline{T}$ is within ε of an affine map, when restricted to Δ_r , provided that r is sufficiently large. (It is worth emphasizing that the diameter of Δ_r is large, and independent of r .) By compactness, the choice of r can be made independent of the ray γ . $\square$

Lemma 7.4. — *Suppose that T is not a similarity. Then, for any n the following is true: There are cells C_1 and C_2 of G , which are disjoint from B_n , such that $\underline{T}(C_1)$ and $\underline{T}(C_2)$ are not translation equivalent.*

Proof. — If this was false, then there would be a value of n having the following property: The images of all cells avoiding B_n would be translation equivalent. Let L denote the union of these cells. (Note that L is a countable set of points.) From the reasoning in Lemma 7.3, it follows that $\underline{T}$ is affine when restricted to L . Since $\underline{T}$ commutes with dilations, $\underline{T}$ is affine when restricted to the set $L^n = D^{-n}(L)$. The set L^n becomes arbitrarily dense as $n \rightarrow \infty$. Taking a limit, we see that $\underline{T}$ is affine. This is only possible if T is a similarity. $\square$

We now come to the main fact about the map $\underline{T}$.

Lemma 7.5. — *Suppose that $\underline{T}$ is not a similarity. Suppose that Λ_1 and Λ_2 are two co-compact lattices of translations of $\mathbf{E}$. Let G be an orbit of Λ_1 . Then $\underline{T}(G)$ cannot be contained in a finite union of orbits of Λ_2 .*

Proof. — We will assume the contrary, and derive a contradiction. By Lemma 7.2, $\underline{T}$ is bi-Lipschitz. Hence, there is a bound, above and below, on the size of $\underline{T}(C)$, where C is a cell of G . Hence, if $\underline{T}(G)$ was contained in a finite union of Λ_2 -orbits, there would be only a finite set of possible shapes for the image $\underline{T}(C)$. But this contradicts Lemma 7.3 and Lemma 7.4. $\square$

7.3. Packing Contradiction

We will now apply the above theory to prove Lemma 7.1. Suppose that q and $h = \partial q$ satisfy the hypotheses of Lemma 7.1. By translating, we can assume that 0 is the basepoint of a horosphere of Ω_1 . Since q is adapted to the pair (Ω_1, Ω_2) , and h is real linear, 0 is also the basepoint of a horosphere of Ω_2 .

Let J be the isometry of $\mathbf{H}$ which extends inversion. Consider the following objects:

1. $\underline{\Omega}_j = J(\Omega_j)$.
2. $\underline{q} = J \circ q \circ J$.
3. $\underline{h} = J \circ h \circ J = I \circ h \circ I$.

Note that $\underline{q}$ is a quasi-isometry adapted to the pair $(\underline{\Omega}_1, \underline{\Omega}_2)$. Note also that there is a horosphere σ_j of $\underline{\Omega}_j$ based at ∞ .

Since $\underline{\Omega}_j$ is an equivariant neutered space, there is a co-compact lattice of translations Λ_j of $\mathbf{E}$ having the following property: The hyperbolic extension of any element of Λ_j is an isometry of $\underline{\Omega}_j$ which preserves σ_j .

For each point $x \in \mathbf{E}$, we define $f_j(x)$ to be the hyperbolic distance from the horosphere of $\underline{\Omega}_j$ based at x to σ_j . If x is not the basepoint of a horosphere of $\underline{\Omega}_j$, we define $f_j(x) = \infty$. Note that f_j is Λ_j -equivariant. It follows easily from packing considerations that $(*)$: the set

$$S_{j,r} = f_j^{-1}[0, r]$$

is contained in a finite union of Λ_j -orbits.

Recall that $\underline{h} = \partial \underline{q}$, where $\underline{q}$ is a hyperbolic quasi-isometry adapted to $(\underline{\Omega}_1, \underline{\Omega}_2)$. This implies (**): for each $r > 0$,

$$\underline{h}(S_{1,r}) \subset S_{2,s}$$

for some s which depends on r and the quasi-isometry constant of $\underline{q}$.

Now, let x be any basepoint of a horosphere of $\underline{\Omega}_1$, and let G be the orbit of x under Λ_1 . From (*) and (**) we see that $\underline{h}(G)$ is contained in a finite union of Λ_2 -orbits. This contradicts Lemma 7.5, unless h is a similarity.

8. Rigidity Lemma: Complex Case

The purpose of this chapter is to prove the Rigidity Lemma in the case of complex hyperbolic space $\mathbf{CH}$. The technique is exactly the same as that for the real case. The only difference is that the analytic underpinnings are less well known. Our source for information about quasi-conformal maps on the boundary of complex hyperbolic space is [P]. Another reference is [KR].

8.1. Three Kinds of Automorphisms

A horosphere of $\mathbf{CH}^{n+1}$ has the geometry of the *Heisenberg group* $G(\mathbf{C}, n)$, described in § 2. In this section, we will describe some of the automorphisms of the Heisenberg group. We distinguish three types, listed in order of generality.

Heisenberg Dilations. — A Heisenberg dilation is a map of the form D^r , where

$$D^r(\xi, v) = (\exp(r) \xi, \exp(2r) v).$$

Here r is a real number.

Heisenberg Similarities. — A Heisenberg similarity is a map of the form

$$(\xi, v) \rightarrow (T(\xi), \det(T)^{1/n} v).$$

Here T is a similarity of $\mathbf{C}^n$ relative to the inner product $\langle \xi_1, \xi_2 \rangle = \operatorname{Re}(\xi_1 \bar{\xi}_2)$.

Linear Contact Automorphisms. — An LCA of $G(\mathbf{C}, n)$ is a smooth group automorphism which preserves the (contact) distribution $D(\mathbf{C}, n)$. Such transformations have the form:

$$(\xi, v) \rightarrow (T(\xi), \det(T)^{1/n} v).$$

Here $T = S_1 S_2$, where S_1 is a similarity of $\mathbf{C}^n$, as described above, and S_2 is a symplectic transformation of $\mathbf{C}^n$. In other words, S_2 preserves the symplectic form $(\xi_1, \xi_2) \rightarrow \operatorname{Im}(\xi_1 \bar{\xi}_2)$.

8.2. Stereographic Projection

Let $X = \mathbf{CH}^{n+1}$. The sphere at infinity, ∂X , can be considered the one point compactification of $G(\mathbf{C}, n)$, as follows: Let $\infty \in \partial X$ be any point, and let σ be a horosphere based at ∞ . For any point $x \in \partial X - \infty$, we define $\rho^\infty(x) = \pi_\sigma(x)$.

Suppose that $0 \neq \infty$ is another point of ∂X . We identify σ to $G(\mathbf{C}, n)$ by an isometry I_0 which takes $\rho^\infty(0)$ to the identity element $(0, 0) \in G(\mathbf{C}, n)$. We define the *stereographic projection*

$$\varphi_0^\infty = I_0 \circ \rho^\infty : \partial X - \infty \rightarrow G(\mathbf{C}, n)$$

φ_0^∞ is well defined up to post-composition with a Heisenberg similarity.

For map $h : \partial X \rightarrow \partial X$ which fixes ∞ , we define

$$h_\varphi = \varphi \circ h \circ \varphi^{-1}.$$

If T is an isometry of X which fixes 0 and ∞ , and $h = \partial T$, then h_φ is a Heisenberg similarity. If T is a pure translation along the geodesic connecting 0 to ∞ , then h_φ is a Heisenberg dilation.

8.3. Heisenberg Differentiability

Let D^r be the Heisenberg dilation defined above. Let $f : G(\mathbf{C}, n) \rightarrow G(\mathbf{C}, n)$ be a homeomorphism such that $f(0, 0) = (0, 0)$. We say that f is *Heisenberg differentiable* at $(0, 0)$ if the sequence

$$D^r \circ f \circ D^{-r}; \quad r \rightarrow \infty$$

converges uniformly on compacta to an LCA. More generally, we say that f is Heisenberg differentiable at some other point of $g \in G(\mathbf{C}, n)$ if the map

$$f(g)^{-1} \circ f \circ g$$

is Heisenberg differentiable at $(0, 0)$. We will denote this differential by $df(g)$. By hypothesis, $df(g)$ is an LCA.

Let $q : X \rightarrow X$ be a quasi-isometry. Composing by isometries of $\mathbf{CH}^{n+1}$, we can assume that $h = \partial q$ fixes ∞ . Let φ be the stereographic projection described above. We now list two facts about h .

1. h_φ is a.e. (resp. Haar measure) Heisenberg differentiable [P, Th. 5].
2. If dh_φ is a.e. a Heisenberg similarity, then h is the boundary extension of an isometry of X [P, Prop. 11.5].

8.4. Zooming In

Suppose that $T : X \rightarrow X$ is an isometry which is pure translation along the geodesic whose endpoints are 0 and ∞ . Let $h : \partial X \rightarrow \partial X$ be a map which fixes both 0 and ∞ . Then the stereographic projection φ_0^∞ conjugates

$$\partial T \circ h \circ \partial T^{-1}$$

to

$$D^r \circ h_\varphi \circ D^{-r}.$$

The constant r is essentially the translation length of T .

All of the limiting/compactness arguments of § 6 work in the complex hyperbolic setting, *mutatis mutandis*. The conjugation above, together with these arguments, gives:

Lemma 8.1 (Complex Case). — Suppose that $q : \mathbf{CH}^n \rightarrow \mathbf{CH}^n$ is a quasi-isometry, adapted to a pair (Ω_1, Ω_2) of equivariant neutered spaces. Suppose $h = \partial q$ fixes ∞ . Let $x \in G(\mathbf{C}, n)$ be a generic point of Heisenberg differentiability for h_φ . Then there are isometric copies Ω'_j of Ω_j , and a quasi-isometry $q' : X \rightarrow X$ such that

1. q' is adapted to the pair (Ω'_1, Ω'_2) .
2. h'_φ is an LCA.
3. $h'_\varphi = dh_\varphi(x)$.

8.5. Inversion

Let $J : X \rightarrow X$ be an isometric involution. (Unlike the real hyperbolic case, such an involution cannot have a codimension-one fixed point set.) Choose any point $0 \in \partial X$, such that $J(0) \neq 0$, and let $\infty = \partial J(0)$. We define *Heisenberg Inversion*

$$I : G(\mathbf{C}, n) - (0, 0) \rightarrow G(\mathbf{C}, n) - (0, 0)$$

to be the composition

$$I = \varphi_0^\infty \circ J \circ (\varphi_0^\infty)^{-1}.$$

Lemma 8.2. — If T is a Heisenberg similarity, then so is $\underline{T} = I \circ T \circ I$.

Proof. — Clearly, if T is an isometry of X which fixes 0 and ∞ , then $J \circ T \circ J$ is also an isometry fixing 0 and ∞ . Our lemma now follows from the fact such isometries, under stereographic projection, induce Heisenberg similarities. $\square$

Lemma 8.3. — Suppose T is an LCA. If T is not Heisenberg similarity, then $\underline{T}$ is not an LCA.

Proof. — It follows from symmetry that I preserves $V = \mathbf{C}^n \times \{0\}$. It follows from Lemma 8.2 that I_V is the composition of a Euclidean similarity and an inversion. Also, T preserves V . The map $T|_V$ is a linear transformation which is not a similarity of V . Hence $\underline{T}|_V$ cannot be a linear transformation. This implies that $\underline{T}$ is not an LCA. $\square$

Apology. — For the reader who is unwilling to use “symmetry” to see that I preserves V , here is a more *ad-hoc* line of reasoning. The manifold $G(\mathbf{F}, n) - (0, 0)$ is foliated by codimension-one hypersurfaces which are invariant under Heisenberg similarities. Call this foliation $\mathcal{F}$. Lemma 8.2 implies that I preserves the leaves of $\mathcal{F}$. Every leaf of $\mathcal{F}$, except V , is a punctured rotationally symmetric paraboloid. It follows that T cannot take a leaf of $\mathcal{F}$ to a leaf of $\mathcal{F}$, unless this leaf is V . If $\underline{T}$ is an LCA, then it is certainly not a similarity. Hence $\underline{T}$ preserves V . But then I preserves V as well.

The following results have proofs exactly analogous to those in the real case.

Lemma 8.4. — Suppose that T is an LCA. Then

1. $\underline{T}$ commutes with Heisenberg dilations.
2. $\underline{T}$ is Heisenberg differentiable away from $(0, 0)$.
3. $\underline{T}$ is bi-Lipschitz in the C - C metric.

Exactly as in § 7, Lemma 8.3 and Lemma 8.4 imply:

Lemma 8.5. — Let $\Lambda_1, \Lambda_2 \subset G(\mathbf{C}, n)$ be two co-compact discrete subgroups. If T is an LCA, but not a Heisenberg similarity, then $\underline{T}$ cannot take an orbit of Λ_1 into a finite union of orbits of Λ_2 .

8.6. Packing Contradiction

Let J be the isometric involution of X defined above. Let $\varphi = \varphi_0^\infty$ be stereographic projection. Let I be Heisenberg inversion. We normalize so that 0 is the basepoint of a horosphere of Ω_j .

We define $\underline{\Omega}_j$, q' and h' as in § 7. Note, in particular, that

$$h'_\varphi = I \circ h'_\varphi \circ I.$$

There are co-compact lattices $\Lambda_j \subset G(\mathbf{C}, n)$ whose elements have the form ∂T_φ , where T is an isometry of $\underline{\Omega}_j$ preserving the relevant horosphere based at ∞ . Let $x \in G(\mathbf{C}, n)$ be a point such that $\varphi^{-1}(x)$ is the basepoint of a horosphere of $\underline{\Omega}_1$. Let G_x denote the Λ_1 -orbit of x . The same argument as in § 7 says that h'_φ takes G_x into a finite union of Λ_2 -orbits. Lemma 8.3 therefore implies that h'_φ is a Heisenberg similarity. This suffices to prove the Rigidity Lemma in the complex case.

9. The Commensurator

Let $X \neq \mathbf{H}^2$ be any rank one symmetric space. Let q be a quasi-isometry of X adapted to the pair (Ω_1, Ω_2) , where Ω_j is an equivariant neutered space. We know from the Rigidity Lemma of § 6 that q is equivalent to an isometry

$$q_* : X \rightarrow X.$$

The goal of this chapter is to show that q_* commensurates the isometry group of Ω_1 to that of Ω_2 . For the sake of exposition, we first sketch a proof in the simplest arithmetic case. Afterwards, we turn to the general case.

9.1. Special Case

Suppose that $X = \mathbf{H}^3$. Nonuniform arithmetic lattices in X are all commensurable with $\mathrm{PSL}_2(\mathcal{O})$, where $\mathcal{O}$ is the ring of integers in an imaginary quadratic field F . We will consider exactly these lattices.

Let Ω_j and q_* be as above. If Ω_j is suitably normalized, then the union of basepoints of horoballs of Ω_j coincides with $F_j \cup \infty$. Hence, ∂q_* is a Möbius transformation which induces a bijection between $F_1 \cup \infty$ and $F_2 \cup \infty$. From this it is easy to see that $F_1 = F_2$, and $\partial q_* \in \text{PGL}_2(F_j)$. This last group is isomorphic to the commensurator of $\text{PSL}_2(\mathcal{O}_j)$.

9.2. General Case

We now give a general argument, which works in all cases. We begin by isolating the key feature of q_* .

Lemma 9.1. (Bounded Distance). — *Let σ_1 be a horosphere of Ω_1 . Then the horosphere $q_*(\sigma_1)$ is parallel to some horosphere σ_2 of Ω_2 . Furthermore,*

$$d_X(q_*(\sigma_1), \sigma_2) \leq K,$$

where K does not depend on the choice of σ_1 .

Proof. — Since q is adapted to (Ω_1, Ω_2) , the extension $h = \partial q$ induces a bijection between basepoints of horospheres of Ω_1 and basepoints of horospheres of Ω_2 . Hence $q_*(\sigma_1)$ is parallel to a horosphere σ_2 of Ω_2 . Since q and q_* are equivalent, $q_*(\sigma_1)$ and σ_2 are uniformly close. $\square$

To say that q_* commensurates the isometry group of Ω_1 to that of Ω_2 is to say that the common isometry group of Ω_2 and $q_*(\Omega_1)$ has finite index in the isometry group of Ω_2 . We will suppose that this is false, and derive a contradiction.

By assumption, there is an infinite family of isometries I_n having the following properties:

1. $I_n(\Omega_2) = \Omega_2$.
2. $I_m(q_*(\Omega_1)) \neq I_n(q_*(\Omega_1))$ if $m \neq n$.

Let ∞ be any point of ∂X . We normalize so that ∞ is the basepoint of a horosphere ω_j of Ω_j . Let $\Phi_0 = q_*(\Omega_1)$, and let $\Phi_n = I_n(\Phi_0)$. Modulo the isometry group of Φ_0 , there are only finitely many equivalent horospheres. Let φ_n denote the horosphere of Φ_n based at ∞ . By taking a subsequence, we can assume that the horospheres $I_n^{-1}(\varphi_n)$ of Φ_0 are all equivalent. By replacing Φ_0 by Φ_1 if necessary, we can assume that the horospheres $I_n^{-1}(\varphi_n)$ are in fact all equivalent to φ_0 . This is to say that there is an isometry J_n of X such that

1. $J_n(\Phi_0) = \Phi_n$.
2. J_n fixes ∞ .

Let Σ denote the group of isometries of X which fix ∞ . Let Λ_0 denote the isometry subgroup of Φ_0 that fixes ∞ . Clearly $\Lambda_0 \subset \Sigma$. We form the coset space $\Sigma/\Lambda_0 = \mathcal{E}$ as follows: We identify elements of the form S and $S \circ \lambda$, where $\lambda \in \Lambda_0$. The space $\mathcal{E}$ is

topologically the product of a line with a compact nilpotent orbifold. Let $\mathcal{J} \subset \mathcal{E}$ denote the image of $\{J_n\}$ in the quotient space. Since $\Phi_n \neq \Phi_m$, the set $\mathcal{J}$ is infinite.

We now contradict this. Observe that

$$d_X(\varphi_n, \omega_2) = d_X(I_n^{-1}(\varphi_n), I_n^{-1}(\omega_2)).$$

These last two horospheres are parallel horospheres, belonging respectively to $q_*(\Omega_1)$ and Ω_2 . It follows from Lemma 9.1 that

$$(1) \quad d_X(\varphi_n, \omega_2) \leq K_1.$$

From (1) it follows that $\mathcal{J}$ is precompact in $\mathcal{E}$. To show that $\mathcal{J}$ is finite, then, we just have to show it is a discrete set.

Suppose $\mathcal{J}$ is not a discrete set. Let $x \in \partial X$ be the basepoint of a horosphere φ_0^x of Φ_0 . Let G_x denote the orbit of x under Λ_0 . Define

$$Y = \bigcup_1^\infty J_n(G_x).$$

If J_n is indiscrete then so is $Y \subset \partial X$.

Let $y \in Y$ be a point. For some n , there is a horosphere φ_n^y of Φ_n based at y . Clearly, φ_n^y is parallel to a horosphere ω_2^y of Ω_2 . The same reasoning as above shows that

$$(2) \quad d_X(\varphi_n^y, \omega_2^y) \leq K_2.$$

By construction,

$$(3) \quad d_X(\varphi_n^y, \varphi_n) = d_X(\varphi_0^x, \varphi_0) = K_3.$$

Putting (1), (2) and (3) together, we see that

$$d_X(\omega_2^y, \omega_2) \leq K_4.$$

Since y was arbitrary, we see that every point of Y is the basepoint of a horosphere of Ω_2 which is within K_4 of the horosphere ω_2 . Packing considerations say that Y is therefore a discrete set.

10. Main Theorem and Corollaries

10.1. The Main Theorem

Suppose that $\Lambda_1, \Lambda_2 \subset G$ are two non-uniform rank one lattices. Let X be the rank one symmetric space associated to G . As usual, we assume $X \neq \mathbf{H}^2$. Suppose that

$$q_0 : \Lambda_1 \rightarrow \Lambda_2$$

is a quasi-isometry.

From Lemma 2.1, Λ_j is canonically quasi-isometric to an equivariant neutered space Ω_j . Furthermore, Λ_j has finite index in the isometry group of Ω_j . Under the identification of Λ_j and Ω_j , the quasi-isometry q_0 induces a quasi-isometry

$$q_1 : \Omega_1 \rightarrow \Omega_2.$$

From the Ambient Extension Lemma, q_1 is equivalent to (the restriction of) a quasi-isometry

$$q_3 : X \rightarrow X$$

which is adapted to the pair (Ω_1, Ω_2) .

From the Rigidity Lemma of § 7, q_3 is equivalent to an isometry $q_4 \in G$. From § 9, q_4 commensurates the isometry group of Ω_1 to that of Ω_2 . Hence q_4 commensurates Λ_1 to Λ_2 . All of this proves that q_0 is equivalent to (the restriction of) an element of G which commensurates Λ_1 to Λ_2 .

10.2. Canonical Isomorphism

Let $\Lambda \subset G$ be a non-uniform rank one lattice. Let X be the associated rank one symmetric space. Every commensurator of Λ induces, by restriction and bounded modification, a self-quasi-isometry of Λ . Thus there is a canonical homomorphism from the commensurator of Λ into the quasi-isometry group of Λ . Since different commensurators have different actions on ∂X , this homomorphism is injective. Our Main Theorem says that this injection is in fact a surjection. Hence, the commensurator of Λ is canonically isomorphic to the quasi-isometry group of Λ .

10.3. Classification

In this section, we will give the quasi-isometry classification of rank one lattices. The quasi-isometry classification of uniform lattices is well known. We will concentrate on the non-uniform case.

Suppose that Λ_j is a non-uniform rank one lattice, acting on the rank one symmetric space X_j . Suppose that $q : \Lambda_1 \rightarrow \Lambda_2$ is a quasi-isometry. If $X_1 = \mathbf{H}^2$, then $X_2 = \mathbf{H}^2$. This follows from the fact that non-uniform lattices in X_j have infinitely many ends iff $X_j = \mathbf{H}^2$.

Suppose that $X_j \neq \mathbf{H}^2$. We will now give three separate, but sketchy proofs, that $X_1 = X_2$.

1. If Λ_1 and Λ_2 are quasi-isometric, then they have isomorphic commensurators. Standard Lie group theory now implies that $X_1 = X_2$.
2. The Ambient Extension Lemma can be generalized (a bit) to show that q extends to a quasi-isometry between X_1 and X_2 . It is (fairly) well known that this implies $X_1 = X_2$.
3. By the Quasi-Flat Lemma of § 3, q induces a quasi-isometry between a horosphere of X_1 and a horosphere of X_2 . It is (fairly) well known that this implies $X_1 = X_2$.

We now reach the interesting part of the classification. If Λ_1 and Λ_2 are two non-uniform lattices acting on X , then any quasi-isometry between them commensurates Λ_1 to Λ_2 , by the Main Theorem. This implies that the two lattices are commensurable.

10.4. Quasi-Isometric Rigidity

Suppose that Γ is a finitely generated group quasi-isometric to a non-uniform lattice Λ . We will show that Γ is the finite extension of a non-uniform lattice Λ' , and that Λ' is commensurable with Λ .

Let Ω be an equivariant neutered space quasi-isometric to Λ . Let

$$q : \Gamma \rightarrow \Omega$$

be a quasi-isometry. Let q^{-1} be a (near) inverse of q .

Note that each element $\gamma \in \Gamma$ acts isometrically on Γ . Let I_γ be this isometry. Let

$$\Phi(\gamma) = q \circ I_\gamma \circ q^{-1}.$$

Then $\Phi(\gamma)$ is a uniform quasi-isometry of Ω . Let $\Psi(\gamma)$ denote the isometry of X whose restriction to Ω is equivalent to $\Phi(\gamma)$. The constants in all of our arguments only depend on the pair $(K(q), \Omega)$ where $K(q)$ is the quasi-isometry constant of q . Hence, there is a constant K' such that $\Phi(\gamma)$ and $\Psi(\gamma)|_\Omega$ differ pointwise by at most K' , independent of γ .

Now, Ψ gives a representation of Γ into the commensurator of Λ . In steps, we will characterize Ψ .

Finite Kernel. — Let e be any element of Γ . All but finitely many elements of Γ move the point $e \in \Gamma$ more than N units away from itself. Hence, all but finitely many $\Phi(\gamma)$ move the point $q(e)$ more than N units away from itself. From the uniformity of the distance between Φ and Ψ , we see that only finitely many $\Psi(\gamma)$ can fix $q(e)$. This is to say that Ψ has finite kernel ⁽⁴⁾.

Discreteness. — If the image of Ψ was not discrete, then there would be an infinite sequence of elements $\Psi(\gamma_1), \Psi(\gamma_2) \dots$ converging to the identity. This is ruled out by an argument just like the one given for the finiteness of the kernel.

Cofinite Volume. — Suppose that $g = \Psi(\gamma)$. From Lemma 9.1, and the uniformity mentioned above, there is a uniform constant K'' having the following property: Every horosphere of $g(\Omega)$ is within K'' of the corresponding horosphere of Ω . It follows that

$$\Omega' = \bigcup_{\gamma \in \Gamma} \Psi(\gamma)(\Omega)$$

is a neutered space on which $\Psi(\Gamma)$ acts isometrically. Since Γ acts transitively on itself, the quotient Ω'/Ψ has finite diameter, and hence is compact. This says that Ψ acts on X with finite volume quotient.

⁽⁴⁾ I would like to thank B. Farb for supplying this argument.

Thus, we see that Γ is the finite extension of a non-uniform lattice Λ' . From the classification above, we see that Λ' is commensurable with Λ .

10.5. Arithmeticity

Let $\Lambda \subset G \neq \mathrm{PSL}_2(\mathbf{R})$ be a non-uniform rank one lattice. It is a result of Margulis that Λ is arithmetic if and only if it has infinite index in its commensurator. (See [Z, Ch. 6] for details.) By our Main Theorem, the commensurator of Λ is isomorphic to the quasi-isometry group of Λ . Thus, Λ is arithmetic if and only if Λ has infinite index in its quasi-isometry group.

REFERENCES

- [C] K. CORLETTE, Hausdorff Dimension of Limits Sets I, *Invent. Math.*, **102** (1990), 521-541.
- [E1] D. B. A. Epstein *et al.*, *Word Processing in Groups*, Jones and Bartlett, 1992.
- [E2] D. B. A. Epstein, *Analytical and Geometric Aspects of Hyperbolic Space*, LMS Lecture Notes, series 111, Cambridge University Press, 1984.
- [Go] W. GOLDMAN, *Complex Hyperbolic Space*, notes available from the author.
- [Gr1] M. GROMOV, *Asymptotic Invariants of Infinite Groups*, LMS Lecture Notes Series, 1994.
- [Gr2] M. GROMOV, *Carnot-Caratheodory Spaces Seen from Within*, IHES, preprint, 1994.
- [KR] A. KORANYI and H. M. REIMANN, *Foundations for the Theory of Quasi-Conformal Mappings of the Heisenberg Group*, Preprint, 1991.
- [M1] G. D. MOSTOW, *Strong Rigidity of Locally Symmetric Spaces*, Annals of Math. Studies, No. 78, Princeton University Press, 1973.
- [M2] G. D. MOSTOW, Quasiconformal Mappings in n -Space and the Rigidity of Hyperbolic Space Forms, *Publ. Math. IHES*, **34** (1968), 53-104.
- [P] P. PANSU, Métriques de Carnot-Caratheodory et Quasi-Isométries des Espaces Symétriques de Rang Un., *Annals of Math.*, **129** (1989), 1-60.
- [T] W. THURSTON, *The Geometry and Topology of Three-Manifolds*, Princeton University Lecture Notes, 1978.
- [Z] R. ZIMMER, *Ergodic Theory and Semi-Simple Lie Groups*, Birkhauser, Boston, 1984.

Dept. of Mathematics
University of Maryland
College Park, MD 20742-0001
USA

Manuscrit reçu le 19 juillet 1994.