Geometrical finiteness for hyperbolic groups

B. H. Bowditch
Faculty of Mathematical Studies, University of Southampton,
Highfield, Southampton SO9 5NH, Great Britain

1. General remarks about finiteness.

1.1. Hyperbolic space.

We begin with a general discussion of hyperbolic geometry in order to introduce our terminology and notation. More details may be found in [Bea, Chapter 7].

We shall write S^{n} for the unit n-sphere in euclidean space. We write $\mathbf{E}^{\mathbf{n}}$ for euclidean n-space, and $\mathbf{H}^{\mathbf{n}}$ for hyperbolic n-space. We shall denote the metrics on these spaces by d_{sph}, d_{euc} and d_{hyp} respectively. We shall drop the subscripts where there can be no confusion. In each case, we write $\mathrm{Isom}\,\mathrm{X}$ for the group of all isometries of X.

If Q is a subset of our space X, we write $\mathrm{int}\, Q$ for the topological interior of Q, and $\bar{Q}$ for the closure of Q. If $r \geq 0$, we write $N_r(Q) = \{x \in X \mid d(x,Q) \leq r\}$ for the uniform r-neighbourhood of Q. If Γ is a subgroup of $\mathrm{Isom}\, X$, we write $\mathrm{stab}_\Gamma\, Q = \{\gamma \in \Gamma \,|\, \gamma Q = Q\}$ for the setwise stabliser of Q. We write $\Gamma Q = \{\gamma Q \,|\, \gamma \in \Gamma\}$ for the set of images of Q under Γ.

In the cases of euclidean and hyperbolic spaces, we write $[x,y]$ for the geodesic segment joining the points x and y. If $z \in X \setminus \{x\}$, we write $\overrightarrow{zx}$ for the unit tangent vector based at z in the direction of x. If also $z \neq y$, we write $x\hat{z}y$ for the angle between $\overrightarrow{zx}$ and $\overrightarrow{zy}$.

We can represent $\mathbf{H}^{\mathbf{n}}$ conformally as the open unit ball in $\mathbf{R}^{\mathbf{n}}$ with infinitessimal metric $d_{hyp} = \frac{2}{1-r^2} d_{euc}$, where r is the euclidean distance from the centre. This is the *Poincaré model*. The closed unit ball gives a canonical compactification of $\mathbf{H}^{\mathbf{n}}$, which we denote by $\mathbf{H}^{\mathbf{n}}_{\mathbf{C}}$. We write $\mathbf{H}^{\mathbf{n}}_{\mathbf{I}}$ for the $(n-1)$-sphere of ideal points, so that $\mathbf{H}^{\mathbf{n}}_{\mathbf{C}} = \mathbf{H}^{\mathbf{n}} \cup \mathbf{H}^{\mathbf{n}}_{\mathbf{I}}$. Any isometry $\gamma \in \mathrm{Isom}\, \mathbf{H}^{\mathbf{n}}$ can be extended to act conformally on $\mathbf{H}^{\mathbf{n}}_{\mathbf{I}}$.

Another conformal representation of $\mathbf{H}^{\mathbf{n}}$ is as the *upper half-space* in $\mathbf{R}^{\mathbf{n}}$; that is, $\mathbf{R}^{\mathbf{n}}_+ = \{\underline{\mathbf{x}} \in \mathbf{R}^{\mathbf{n}} \mid \mathbf{x_n} > \mathbf{0}\}$, where x_n is the last coordinate of $\underline{x}$. The metric is given infinitessimally by $d_{hyp} = \frac{1}{x_n} d_{euc}$. Writing $\partial \mathbf{R}^{\mathbf{n}}_+ = \{\underline{\mathbf{x}} \in \mathbf{R}^{\mathbf{n}} \mid \mathbf{x_n} = \mathbf{0}\}$, we may identify $\mathbf{H}^{\mathbf{n}}_{\mathbf{I}}$ as $\partial \mathbf{R}^{\mathbf{n}}_+ \cup \{\infty\}$, where the ideal point ∞ compactifies $\mathbf{R}^{\mathbf{n}}_+ \cup \partial \mathbf{R}^{\mathbf{n}}_+$ into a ball. We shall refer to the nth coordinate in $\mathbf{R}^{\mathbf{n}}_+$ as the *vertical coordinate*. Note that if $\gamma \in \mathrm{Isom}\, \mathbf{H}^{\mathbf{n}}$ fixes ∞, then it acts as a euclidean similarity on $\partial \mathbf{R}^{\mathbf{n}}_+$.

A third model for hyperbolic space we shall use is the *Klein model*. This consists of the open unit ball with a (non-conformal) Riemannian metric, such that all hyperbolic geodesics correspond to euclidean line segments (see [Bea, Chapter 7]).

In hyperbolic space, we may extend the notations $[x,y]$, $\overrightarrow{zx}$ and $x\hat{z}y$ to the case where x or y or both lie in $\mathbf{H}^{\mathbf{n}}_{\mathbf{I}}$. If $x, y \in \mathbf{H}^{\mathbf{n}}$ we call $[x,y]$ a *geodesic segment*. If $x \in \mathbf{H}^{\mathbf{n}}$ and $y \in \mathbf{H}^{\mathbf{n}}_{\mathbf{I}}$, we call $[x,y] \cap \mathbf{H}^{\mathbf{n}}$ a *geodesic ray based at* x and *tending to* y. If $x, y \in \mathbf{H}^{\mathbf{n}}_{\mathbf{I}}$, we call $[x,y] \cap \mathbf{H}^{\mathbf{n}}$ a *bi-infinite geodesic*.

By a *subspace* of $\mathbf{H^n}$, we mean a complete totally geodesic subset. By a *compactified subspace* of $\mathbf{H^n_C}$, we mean the closure in $\mathbf{H^n_C}$ of such a subspace.

We may classify non-trivial isometries $\mathbf{H^n}$ into three types, namely elliptic, parabolic and loxodromic as follows.

Let γ be an isometry of $\mathbf{H^n}$. We write $\operatorname{fix}\gamma$ for the set of fixed points of γ in $\mathbf{H^n_C}$. Brouwer's fixed point theorem tells us that $\operatorname{fix}\gamma$ must be non-empty.

Suppose that there is some point x in $\operatorname{fix}\gamma \cap \mathbf{H^n}$. We may take x to be the centre of the ball in the Poincaré model. Then, γ acts as a euclidean rotation on the ball, and we see that $\operatorname{fix}\gamma$ is a (possibly 0-dimensional) compactified subspace in $\mathbf{H^n_C}$. We call this case *elliptic.*

If γ is not elliptic, then $\operatorname{fix}\gamma$ is a subset of $\mathbf{H^n_I}$. Suppose that $\operatorname{fix}\gamma$ consists of just a single point in $\mathbf{H^n_I}$. We may take this point to be ∞ in the upper half space model, $\mathbf{R^n_+}$. Now, since γ has no fixed point in $\partial\mathbf{R^n_+}$, it must act as a euclidean isometry of $\partial\mathbf{R^n_+}$. Moreover, it must preserve setwise each horosphere about ∞. We call this case *parabolic.*

Suppose that γ fixes precisely two points, x and y, in $\mathbf{H^n_I}$. Let β be the geodesic joining x to y. In this case, γ acts as a translation on β, and (in general) has a rotational component in the orthogonal direction. We call this case *loxodromic*, and we call β the *loxodromic axis.* If γ translates β in the direction of x, we call x the *attracting* fixed point of γ, and y the *repelling* fixed point.

Finally, note that if γ has three (or more) fixed points in $\mathbf{H^n_I}$, then these must determine a fixed point in $\mathbf{H^n}$, so we are back in the elliptic case.

Given a subgroup $G \subseteq \operatorname{Isom}\mathbf{H^n}$, we write $\operatorname{fix} G = \bigcap_{\gamma\in G}\operatorname{fix}\gamma$ for the set of all fixed points of G in $\mathbf{H^n_C}$.

1.2. Groups of isometries.

Let Γ be a subgroup of $\operatorname{Isom}\mathbf{H^n}$. It is an elementary result that Γ is a discrete subgroup if and only if it acts properly discontinuously on $\mathbf{H^n}$, that is to say, each compact subset of $\mathbf{H^n}$ meets only finitely many images of itself under Γ.

In such a discrete group, the finite-order elements are precisely the elliptic isometries. Thus, Γ acts freely if and only if it is torsion-free. If Γ acts freely, we may form the quotient manifold $M = \mathbf{H^n}/\mathbf{\Gamma}$ which inherits a complete hyperbolic structure.

More generally, if Γ has torsion, the quotient $M = \mathbf{H^n}/\mathbf{\Gamma}$ is a complete hyperbolic "orbifold", as defined by Thurston [Th1, chapter 13]. That is to say, there is a closed cell complex Σ in M, such that $M \setminus \Sigma$ is an (incomplete) hyperbolic manifold. The set Σ can be defined as the projection of the set of all fixed points of elliptic elements of Γ, i.e., $\Sigma = \bigcup_{\gamma\in\Gamma}(\operatorname{fix}\gamma \cap \mathbf{H^n})/\mathbf{\Gamma}$. A neighbourhood of a point of Σ may or may not be topologically singular, but it will always be geometrically singular. In an orientable 2-orbifold, for example, Σ consists of a discrete set of cone singularities, which may be thought of as points of concentrated positive curvature. We shall call Σ the *singular set* of M.

Let $\Gamma \subseteq \operatorname{Isom}\mathbf{H^n}$ be discrete. The action of Γ may be extended to $\mathbf{H^n_C}$, and we may define the *limit set* $\Lambda \subseteq \mathbf{H^n_I}$ as the set of accumulation points of some Γ-orbit in $\mathbf{H^n}$, i.e.

$$\Lambda = \{y \in \mathbf{H^n_I} \mid \text{there exist } \gamma_\mathbf{i} \in \mathbf{\Gamma} \text{ and } \mathbf{x} \in \mathbf{H^n} \text{ with } \gamma_\mathbf{i}\mathbf{x} \to \mathbf{y}\}.$$

It turns out that this definition is independent of our choice of x. Moreover, Λ is a minimal closed Γ-invariant set, and Γ acts properly discontinuously on its complement Ω in $\mathbf{H^n_I}$. The set $\Omega = \mathbf{H^n_I} \setminus \mathbf{\Lambda}$ is called the *discontinuity domain.* (It is possible for Ω to be empty.) We may form the quotient orbifold $M_I = \Omega/\Gamma$ of Ω. Since Γ acts conformally on $\mathbf{H^n_I}$, we see that M_I inherits a (singular) conformal structure from Ω. In fact, Γ acts properly discontinuously on $\mathbf{H^n} \cup \mathbf{\Omega}$, so we may write

$$M_C = (\mathbf{H^n} \cup \mathbf{\Omega})/\mathbf{\Gamma} = \mathbf{M} \cup \mathbf{M_I}.$$

Note that when $n = 3$, M_I is a Riemann surface (in general not connected). This fact gives rise to a rich analytical theory in this dimension.

One direction of research in discrete hyperbolic groups, is to study the relationship of various types of "finiteness" — group theoretic, topological and geometric. The simplest group theoretic restriction is to demand that Γ be finitely generated. We can then ask what this tells us about the topology and geometry of M.

The first result is purely algebraic.

Selberg Lemma [Sel]. *Let k be a field of characteristic 0. Then, any finitely-generated subgroup of* $\mathrm{GL_n(k)}$ *is virtually torsion-free, (i.e. contains a torsion-free subgroup of finite index).*

For a simpler proof, see [Cas].

Since $\mathrm{Isom}\,\mathbf{H^n}$ can be represented as a subgroup of $\mathrm{GL_{n+1}}(\mathbf{R})$, the Selberg Lemma can be applied to finitely-generated subgroups of $\mathrm{Isom}\,\mathbf{H^n}$.

Beyond the Selberg Lemma, little seems to be known in general. The main thrust of research is in dimension 3, and we shall give a summary of 3-dimensional results in Section 1.3. First, we describe how the 2-dimensional case is trivial from the point of view of finiteness.

For simplicity, we restrict attention to orientable surfaces. Let M be a complete, orientable, hyperbolic surface with finitely-generated fundamental group. Then, it turns out that M consists of a compact surface with boundary, together with a finite number of "cusps" and "funnels". A *cusp* is (isometric to) a horoball in $\mathbf{H^2}$, quotiented out by a cyclic parabolic group (Figure 1a). A *funnel* consists of a hyperbolic half-space quotiented out by a loxodromic element (Figure 1b). We see that M_I is a disjoint union of finitely many circles, which serve to compactify the funnels in $M_C = M \cup M_I$. Thus the topological ends of M_C correspond precisely to the cusps (Figure 1c). We see that, in any meaningful sense, the geometry of M is only finitely complicated. This is about the strongest assertion of finiteness one could make.

1.3. Some 3-dimensional finiteness results.

In this section, we shall give a summary of some finiteness results in 3 dimensions. It is not meant to reflect the historical development of the subject.

Because of the Selberg Lemma, we shall not leave much essential out of the story if we restrict attention to torsion-free groups. It will also simplify the exposition a little if we assume that our group is orientation preserving.

Let Γ be a discrete, torsion-free, orientation-preserving subgroup of $\mathrm{Isom}\,\mathbf{H^n}$. Much of the technical complication of the subject arises from having to deal with parabolic subgroups of Γ. Suppose that $\gamma \in \Gamma$ is parabolic with fixed point p. Let $G = \mathrm{stab}_\Gamma\, p$ be the stabiliser of p in Γ. In a discrete group, a parabolic and a loxodromic cannot share a common fixed point (Lemma 3.1.2). Thus, G consists entirely of parabolics. We call p a *parabolic fixed point.* We can let p be the point ∞ in the upper half-space model. Now, G acts freely as a group of isometries on $\partial \mathbf{R}^{\mathbf{3}}_{+} \equiv \mathbf{E^2}$. This dimension is special in that such a group must act by translation. We see that G is isomorphic to either S or $S \oplus S$. Taking B to be any horoball about p, we may form the quotient B/G. If $G \cong S$, then $\partial B/G$ is a bi-infinite euclidean cylinder, and we call B/G a *S-cusp* (Figure 1d). (Note that a S-cusp is not quite the same as a "cusp cylinder" or "standard cusp region" which will be described later in this section. See Chapter 2 for details.) If $G \cong S \oplus S$ then $\partial B/G$ is a euclidean torus, and we call B/G a *$S \oplus S$-cusp* (Figure 1e). We may define such cusps to correspond to each orbit of parabolic fixed points. In general, one would expect these cusps to project to a collection of immersed submanifolds in M. However the Margulis Lemma (see Sections 2.2, and 3.3) tells us that (in dimension 3), by taking our horoballs small enough, we can arrange that the cusps be disjoint and embedded in M. We shall write cusp(M) for the disjoint union of all the cusps.

The construction of this set of disjoint cusps is valid for infinitely-generated groups. From now on, however, we shall insist that Γ be finitely-generated. We first use a purely topological result.

Theorem (Scott [Sc]). *Let M be a 3-manifold with finitely-generated fundamental group. Then, there is a compact submanifold M_T of M, such that the inclusion $M_T \hookrightarrow M$ induces an isomorphism of fundamental groups.*

We call M_T a *topological core* for M. With $M = \mathbf{H^n}/\mathbf{\Gamma}$, we deduce immediately that Γ is finitely presented.

In our case, $M = \mathbf{H^n}/\mathbf{\Gamma}$ is an irreducible 3-manifold, that is each embedded 2-sphere in M bounds a 3-ball. Because of this, we can arrange that ∂M_T contains no 2-spheres, and then the inclusion of M_T into M is a homotopy equivalence. Moreover, there is a bijective correspondence between the boundary components of M_T and the topological ends of M. We deduce that M has only finitely many ends. In particular, it contains only finitely many $S \oplus S$-cusps.

In fact (provided that Γ is not cyclic loxodromic), the $S \oplus S$-cusps correspond precisely to the toroidal components of ∂M_T. The remaining ends correspond to components of genus at least 2. The aim now is to understand something of the geometry of these remaining ends, which we shall call "non-cuspidal ends".

Now, a S-cusp is topologically just a product. Thus, we can assume that each S-cusp lies entirely within some non-cuspidal end. The effect of removing the S-cusps would (in general) be to subdivide each such end into smaller pieces, on which we may see qualitatively different behaviour. It is therefore necessary to take account of these S-cusps before going on to consider the geometry. We can do this by applying a relative version of Scott's theorem to the closure, $\mathrm{noncusp}(M)$, of $M \setminus \mathrm{cusp}(M)$ in M. In this section, we write $M' = \mathrm{noncusp}(M)$.

Theorem [Mc]. *Let N be a 3-manifold with boundary, whose fundamental group is finitely generated. Let S be a compact submanifold of ∂N. Then, we can find a topological core, N_T, for N such that $N_T \cap \partial N = S$.*

By using this result, together with an Euler characteristic argument, one may deduce [FMa] that there are only a finite number of S-cusps — a result due originally to Sullivan [Sul1]. We may now take a core M'_T of M' which meets each $S \oplus S$-cusp in the bounding torus, and each S-cusp in a compact annular core of its boundary cylinder. Again, we may take the inclusion to be a (relative) homotopy equivalence, so that the topological ends of M' correspond to the frontier components of M'_T in M'. We now look for geometric information about the ends of M'.

We have already remarked that, for $n = 3$, $M_I = \Omega/\Gamma$ is a Riemann surface. A fundamental result about M_I is the following.

Ahlfors' Finiteness Theorem [Ah1,Sul1]. *Let Γ be a finitely-generated discrete subgroup of* $\mathrm{Isom}\,\mathbf{H^3}$. *Then $M_I = \Omega/\Gamma$ is a Riemann surface of "finite type". That is to say, M_I is conformally equivalent to a compact surface with finitely many punctures.*

For a proof using deformation theory, see [Sul4]. One needs a special argument to rule out the possibility of there being an infinite number of (rigid) thrice punctured spheres, but this is taken care of if we know that there are only finitely many S-cusps [Sul1].

Moreover one may show that the punctures of M_I arise only from parabolic elements of Γ; that is, a small loop around a puncture represents a conjugacy class of parabolics in Γ.

We want to give Ahlfors' Finiteness Theorem a more geometric interpretation. We can do this by using the convex hull of the limit set — a generalisation of the Nielsen convex region in dimension 2. Let Y be the smallest convex set in $\mathbf{H^3}$ whose closure, Y_C, in $\mathbf{H^3_C}$ contains the limit set Λ. Then, Y_C meets $\mathbf{H^3_I}$ precisely in Λ. Since the construction is equivariant, we may form the quotient $\mathrm{core}(M) = Y/\Gamma \subseteq M$, which we call the *convex core* of M. The nearest point retraction of $\mathbf{H^3}$ onto Y extends continuously to all of $\mathbf{H^3_C}$, and therefore gives rise to a map from M_C to $\mathrm{core}(M)$ (see for example [Th1]). We shall denote by q, the restriction of this map to M_I. Note that $q(M_I) = \partial\mathrm{core}(M)$.

It is possible for $\mathrm{core}(M)$ to have empty interior, but if so, then Γ is either abelian or "fuchsian" (i.e. preserves some 2-plane in $\mathbf{H^3}$). Both these cases are completely understood, so we shall assume that the interior of $\mathrm{core}(M)$ is non-empty. In this case one may show that $\partial\mathrm{core}(M)$ has the structure of a complete hyperbolic surface in the induced path metric [Th1]. Moreover q is a homotopy equivalence from M_I to $\partial\mathrm{core}(M)$. In fact, by applying some kind of smoothing to the nearest point retraction, one may show that q is homotopic to a quasiconformal homeomorphism. ([EM] includes details of this in the case when Λ is connected.) We deduce that the surface $\partial\mathrm{core}(M)$ also has finite conformal type and thus finite hyperbolic area. In other words, we can restate Ahlfors' Finiteness Theorem to say that $\partial\mathrm{core}(M)$ must have finite 2-dimensional area. (In fact the discussion applies equally well if Γ has torsion, and then $\partial\mathrm{core}(M)$ becomes a finite-area orbifold.)

The parabolic cusps of the hyperbolic surface $\partial\mathrm{core}(M)$ are essentially the connected components of $\partial\mathrm{core}(M)\cap\mathrm{cusp}(\mathrm{M})$. In fact the cusps of $\partial\mathrm{core}(M)$ must lie inside S-cusps of M. The remainder of $\partial\mathrm{core}(M)$, namely $\partial\mathrm{core}(M)\cap M'$, is compact. Thus, each component of $\partial\mathrm{core}(M)$ corresponds to an end of M'. Such an end is topologically a product, being foliated by components of $\partial N_r(\mathrm{core}(M))$ for $r > 0$, where $N_r(\mathrm{core}(M))$ is the uniform r-neighbourhood of $\mathrm{core}(M)$. We call such ends *geometrically finite*. We see that the geometrically finite ends of M correspond bijectively to components of $\partial\mathrm{core}(M)$, and thus to components of M_I. (We may think of M_I as the limit of the surfaces $\partial N_r(\mathrm{core}(M))$ as r tends to ∞.) If we fix some $\eta > 0$, we can modify the topological core M_T', so that $\partial N_\eta(\mathrm{core}(M)) \cap M'$ becomes a subset of the frontier of M_T. That is, those frontier components of M_T' in M' which correspond to geometrically finite ends, coincide with frontier components of $N_\eta(\mathrm{core}(M)) \cap M'$.

The geometrically finite ends, however, might not account for all the ends of M'. It may be that an end makes no impression on the discontinuity domain Ω, so that Ahlfors' Finiteness Theorem tells us nothing. Such ends were shown to exist by Bers and Maskit [Ber] [Mas], their geometrically infinite nature being made explicit by Greenberg [Gr]. Jørgensen later described more concrete examples [J]. Thurston [Th3] gives a more general method of construction.

All the non-geometrically finite ends constructed so far have been “simply degenerate” as defined by Thurston [Th1, Chapter 9]. A simply degenerate end turns out to be just a product topologically (i.e. homeomorphic to a surface times a half-open interval), but its geometry is infinite. For example, every neighbourhood of the end will contain infinitely many closed geodesics. Bonahon and Otal construct an example of an end containing closed geodesics of arbitrarily small length [BonO]. There are also examples where lengths of closed geodesics have a positive lower bound. In the latter case the end has bounded diameter as one tends to infinity. In general, one may say that the volume of a simply-degenerate end grows at most linearly. This explains why such an end makes no impression on the discontinuity domain — geometrically finite ends have exponential growth.

If, as in all the examples constructed so far, each (non-cuspidal) end is either geometrically finite or simply degenerate, we call M *geometrically tame*. In this case, M is topologically finite, i.e. homeomorphic to the interior of a compact manifold with boundary. Moreover, one can show that the limit set of such a group has either zero or full 2-dimensional Lebesgue measure (see [Th1] or [Bon]) — a property conjectured, by Ahlfors, for all finitely-generated discrete groups. There are examples, however, where the limit set has Hausdorff dimension equal to 2, while still having zero 2-dimensional Lebesgue measure [Sul2].

It has been conjectured that all finitely-generated discrete groups are geometrically tame. Bonahon [Bon] has proven this under the hypothesis that for any free-product decomposition $\Gamma \cong A * B$, there is some parabolic in Γ not conjugate to any element of A or B. Also, Otal has shown that if Γ is isomorphic to the free product of two compact surface groups, then Γ is geometrically tame. (A partial account is given in [O].) In fact, it seems that the case of free groups is the most difficult to handle.

We now restrict attention to the case where all the ends of $M' = \mathrm{noncusp}(M)$ are geometrically finite. Then, we call M "geometrically finite". In this case, we can assume that each end of M' is bounded by a component of $\partial N_\eta(\mathrm{core}(M))$, which means that we can take the topological core M'_T to be equal to $N_\eta(\mathrm{core}(M)) \cap M' = N_\eta(\mathrm{core}(M)) \cap \mathrm{noncusp}(M)$. In other words, geometric finiteness says that $N_\eta(\mathrm{core}(M)) \cap \mathrm{noncusp}(M)$ is compact. This is more or less the definition of geometric finiteness (GF4) due to Thurston [Th1, Chapter 8] (see Section 3.4). (Taking the η-neighbourhood of the convex core allows us to include Fuchsian groups and cyclic loxodromic groups in the discussion, without making special qualifications.)

Clearly, $N_\eta(\mathrm{core}(M))$ meets the boundary of any S-cusp in a compact set. From this we see that the intersection of $N_\eta(\mathrm{core}(M))$ with any S-cusp has finite volume. (A S-cusp admits a totally geodesic embedding of a 2-dimensional cusp. We see that $N_\eta(\mathrm{core}(M))$ meets the S-cusp is some uniform neighbourhood, $N_r(c)$, of such a 2-dimensional cusp, c (Figure 1f).) Since each $S \oplus S$-cusp has finite volume, we arrive at Thurston's second definition of geometric finiteness (GF5), namely that $N_\eta(\mathrm{core}(M))$ should have finite volume. (For the definition GF4, it is enough to insist that $\mathrm{core}(M) \cap \mathrm{noncusp}(M)$ be compact. For GF5, however, it is essential to take some uniform neighbourhood of $\mathrm{core}(M)$, as the example of an infinitely generated Fuchsian group shows.)

If M had no cusps, one sees that $M_I = \Omega/\Gamma$ would give a compactification of M to M_C. In the general case, the topological ends of M_C correspond precisely to the cusps. In fact, each end of M_C has a neighbourhood isometric to one of two standard types — "cusp tori" and "cusp cylinders". Cusp tori are the same as $S \oplus S$-cusps, whereas a cusp cylinder is an enlargement of a S-cusp to include a portion of M_I (Figure 1g). This description of geometric finiteness (GF1) is due to Marden [Mar].

A fourth description (GF2), due to Beardon and Maskit [BeaM], demands that the limit set should be a union of (what we call here) "conical limit points" and "bounded parabolic fixed points". These will be defined in Sections 3.2 and 3.1 respectively. The notion of a conical limit point (also called a "radial limit point" or "approximation point") originates in [H], and has proven useful to the study of dynamics on limit sets.

Finally, the original and simplest definition of geometric finiteness (GF3) demands that Γ should possess a finite-sided convex fundamental polyhedron. This hypothesis was introduced by Ahlfors [Ah2], where he showed that the limit set of such a group must have either zero or full Lebesgue measure in $\mathbf{H}^3_\mathbf{I}$.

It has been known for some time, from the references already cited, that these five definitions are all equivalent in dimension 3. Geometrically finite groups occur frequently as the simplest examples of 3-dimensional hyperbolic groups. It is conjectured that they contain an open dense subset of the space of all finitely-generated discrete groups, given the appropriate topology (see [Sul5]). The hypothesis of geometrical finiteness has often been used in the study of the dynamics on limit sets. Sullivan, for example showed that the limit set of a geometrically finite group is either the whole sphere $\mathbf{H}^3_\mathbf{I}$, or else has Hausdorff dimension strictly less than 2 [Sul3].

1.4. Higher dimensions.

The study of discrete hyperbolic groups in dimensions greater than 3 is much less well developed. Restricting attention to finitely generated groups does not seem to help very much. In 3 dimensions, this hypothesis gives us topological information (Scott's theorem), as well as analytical information (Ahlfors' Finiteness Theorem). Both these results fail in higher dimensions. For example, Kapovich and Potyagailo [KP] constructed an example of a finitely generated discrete subgroup of $\mathrm{Isom}\,\mathbf{H^4}$ which is not finitely presented, and for which the quotient of the discontinuity domain, $M_I = \Omega/\Gamma$ is topologically infinite. In fact one can find such a group which is a subgroup of a discrete cocompact group acting on $\mathbf{H^4}$ [BowM]. Kapovich also has an example [K] of a finitely generated discrete subgroup of $\mathrm{Isom}\,\mathbf{H^4}$ with infinitely many conjugacy classes of parabolic subgroups (in contrast to Sullivan's cusp-finiteness theorem [Sul1] in dimension 3). In the same paper, he constructs another such group containing infinitely many conjugacy classes of elliptic elements. This is in contrast to a result of Feighn and Mess which says that any finitely generated discrete subgroup of $\mathrm{Isom}\,\mathbf{H^3}$ contains only finitely many conjugacy classes of finite subgroups [FMe].

A natural question to ask is how one should define geometric finiteness in dimensions greater than 3. Most authors have taken geometrical finiteness in this case to mean that the group should possess a finite-sided convex fundamental polyhedron — a direct generalisation of the original definition. However, in dimension 4 and higher, this definition becomes more restrictive than the obvious generalisations of the other four definitions. It seems that these other definitions give rise to a more natural notion of geometrical finiteness which we aim to elucidate in this work. All the applications of the traditional geometrical finiteness hypothesis seem to be valid for this more general notion.

The question of defining geometric finiteness in higher dimensions has also been considered by Apanasov [Ap1,Ap2], as well as by Weilenberg [We] and Tukia [Tu1]. In [Tu2], Tukia generalises, to dimension n, Sullivan's result about the Hausdorff dimension of the limit set. Thus, the limit set of a geometrically finite group is either equal to $\mathbf{H^n_I}$, or else has Hausdorff dimension less than $n-1$.

1.5. Variable curvature.

We remark that the definitions GF1, GF2, GF4 and GF5, as described in this paper, can easily be interpreted in the case where $\mathbf{H^n}$ is replaced by a simply-connected manifold, X, of pinched negative curvature (i.e. all the sectional curvatures of X are bounded between two negative constants). It turns out that these give rise to a well-defined notion of "geometrical finiteness" for discrete group actions on X. Of course, here, the term "finiteness" is less appropriate, though the context seems more natural, in that it removes the dependence on certain special features of constant curvature, such as the existence of half-spaces. These matters are discussed in my other paper [Bow].

2. The Margulis Lemma and Bieberbach Theorem.

In this section we shall be discussing results related to the Margulis Lemma and Bieberbach Theorems. One form of the Margulis Lemma says the following. Given any positive integer n, we can find some $\epsilon(n) > 0$ with the following property. Let (X, d) be any simply connected Riemannian n-manifold, all of whose sectional curvatures lie in the closed interval $[-1, 0]$. Let Γ be any discrete group of isometries acting on X, and $x \in X$ be any point. Let $\Gamma_\epsilon(x)$ be the group generated by those elements of $\gamma \in \Gamma$ such that $d(\gamma x, x) \leq \epsilon(n)$. In symbols, $\Gamma_\epsilon(x) = \langle \gamma \in \Gamma \mid d(\gamma x, x) \leq \epsilon(n) \rangle$. Then, $\Gamma_\epsilon(x)$ is virtually nilpotent (i.e. it contains a nilpotent subgroup of finite index). Moreover, the index of the nilpotent subgroup in $\Gamma_\epsilon(x)$ can be bounded by some $\nu(n)$ depending only on n. We say that groups of the form $\Gamma_\epsilon(x)$ are *uniformly virtually nilpotent.*

A proof of this result may be found in [BaGS]. In this paper, we shall restrict attention to the constant curvature cases, namely $\mathbf{E^n}$ and $\mathbf{H^n}$, where we can give a simple proof of the Margulis Lemma. Also, in these cases we may identify the nilpotent subgroup as being generated by elements of small rotational part, and it turns out always to be abelian. This final observation is a consequence of nilpotency, rather than discreteness, so we begin with a discussion of nilpotent groups of isometries in the geometries $\mathrm{S^n}$, $\mathbf{E^n}$ and $\mathbf{H^n}$. We shall prove that nilpotent subgroups of $\mathrm{Isom}\,\mathrm{S^n}$, $\mathrm{Sim}\,\mathbf{E^n}$, and $\mathrm{Isom}\,\mathbf{H^n}$ are uniformly virtually abelian (where $\mathrm{Sim}\,\mathbf{E^n}$ is the group of all euclidean similarities of $\mathbf{E^n}$). This fact seems to be well known, though I know of no explicit reference. However all the essential ingredients may be found in [Th2]. We shall go on to show how nilpotent groups arise out of discrete isometry groups. In the course of the discussion we deduce some of the classical Bieberbach Theorems. These results are also described in [Th2], [CarD] and [Wo].

2.1. Nilpotent implies virtually abelian.

Let $\mathrm{S^n}$, $\mathbf{E^n}$ and $\mathbf{H^n}$ denote the unit n-sphere, euclidean n-space and hyperbolic n-space respectively, with metrics d_{sph}, d_{euc}, and d_{hyp}. We shall omit the subscripts where there can be no confusion. Let $\mathrm{Isom\,X}$ denote the entire group of isometries of X, and $\mathrm{Sim}\,\mathbf{E^n}$ be the group of euclidean similarities. Throughout, we use the convention on commutators that $[x, y] = xyx^{-1}y^{-1}$.

We shall deal with the three geometries in turn.

2.1(i). Spherical geometry.

Let

$$U(\mathrm{S^n}) = \{\gamma \in \mathrm{Isom\,S^n} \mid \mathrm{d}(\gamma \mathrm{x}, \mathrm{x}) < \pi/2 \text{ for all } \mathrm{x} \in \mathrm{S^n}\}.$$

We may think of S^n as the unit sphere in $\mathbf{R^{n+1}}$, with the standard inner-product $\langle , \rangle$. We see that γ lies in $U(\mathrm{S}^n)$ if it moves each vector in $\mathbf{R^{n+1}}$ through an angle of less than $\pi/2$, in other words $\langle \gamma v, v \rangle > 0$ for each $v \in \mathbf{R^{n+1}} \setminus \{\underline{\mathbf{0}}\}$.

Let $\gamma \in \mathrm{Isom}\, \mathrm{S}^{\mathrm{n}}$. By complexifying, we can extend γ to act on $\mathbf{C}^{\mathbf{n+1}}$. Now, γ preserves the standard hermitian form on $\mathbf{C}^{\mathbf{n+1}}$, i.e. the form that restricts to the inner product on $\mathbf{R}^{\mathbf{n+1}}$. We also use $\langle\, ,\, \rangle$ to denote this hermitian form.

Now, let $v \in \mathbf{C}^{\mathbf{n+1}}$ be any non-trivial complex vector. Write $v = x + iy$, with $x, y \in \mathbf{R}^{\mathbf{n+1}}$. Then,

$$\mathrm{Re}\langle \gamma \mathrm{v}, \mathrm{v}\rangle = \langle \gamma \mathrm{x}, \mathrm{x}\rangle + \langle \gamma \mathrm{y}, \mathrm{y}\rangle.$$

If $\gamma \in U(\mathrm{S}^{\mathrm{n}})$, both the terms on the right hand side are non-negative, and at least one is strictly positive. It follows that γ lies in $U(\mathrm{S}^{\mathrm{n}})$ if and only if $\mathrm{Re}\langle \gamma \mathrm{v}, \mathrm{v}\rangle > 0$ for each non-trivial $v \in \mathbf{C}^{\mathbf{n+1}}$.

We can now prove:

Lemma 2.1.1 : *Let $\beta \in U(\mathrm{S}^{\mathrm{n}})$ and $\alpha \in \mathrm{Isom}\, \mathrm{S}^{\mathrm{n}}$. If α commutes with $[\alpha, \beta]$, then α commutes with β.*

Proof : Complexifying, we imagine α and β acting on $\mathbf{C}^{\mathbf{n+1}}$. We see that α commutes with $\beta^{-1}\alpha\beta$, so that they are simultaneously diagonalisable. Let V be an eigenspace of α. Then βV is an eigenspace of $\beta\alpha\beta^{-1}$. If $V \neq \beta V$, then V must intersect non-trivially some other eigenspace V' of $\beta\alpha\beta^{-1}$, orthogonal to βV. Let $v \in V \cap V'$ be non-zero. Then βv lies in βV, so that $\langle \beta v, v\rangle = 0$. However, since β lies in $U(\mathrm{S}^{\mathrm{n}})$, the discussion immediately prior to the lemma tells us that $\mathrm{Re}\langle \beta \mathrm{v}, \mathrm{v}\rangle > 0$. This contradiction means that $\beta V = V$. Since V was an arbitrary eigenspace of α, we deduce that α and β are simultaneously diagonalisable, and hence commute. $\diamondsuit$

Corollary 2.1.2 : *If $\Gamma \subseteq \mathrm{Isom}\, \mathrm{S}^{\mathrm{n}}$ is nilpotent, then $\langle \Gamma \cap U(\mathrm{S}^{\mathrm{n}})\rangle$ is abelian.*

Proof : Let a and b lie in $\Gamma \cap U(\mathrm{S}^{\mathrm{n}})$. By a "nested chain of commutators" in a and b, we mean an expression of the form $d = [c_1, [c_2, \ldots [c_n, c_{n+1}] \ldots]]$, where each c_i is either a or b. We take d to be of maximal length, n, such that $d \neq 1$. This means that d commutes with both a and b. It follows that $[c_2, \ldots [c_n, c_{n+1}] \ldots]$ commutes with d. Applying Lemma 2.1.1, with $\alpha = [c_2, \ldots [c_n, c_{n+1}] \ldots]$ and $\beta = c_1$, we deduce that α and β commute, so that $d = 1$. We have contradicted the assumption that $n \geq 1$, and so a must commute with b. $\diamondsuit$

Let V be an open symmetric neighbourhood of the identity in $\mathrm{Isom}\, \mathrm{S}^{\mathrm{n}}$ such that $V^2 \subseteq U(\mathrm{S}^{\mathrm{n}})$. There is an upper bound $N(n)$ on the number of disjoint translates of V by $\mathrm{Isom}\, \mathrm{S}^{\mathrm{n}}$ that we can embed in $\mathrm{Isom}\, \mathrm{S}^{\mathrm{n}}$. We deduce that $[\Gamma : \langle \Gamma \cap U(\mathrm{S}^{\mathrm{n}})\rangle] < \mathrm{N(n)}$, and so,

Corollary 2.1.3 : *Nilpotent subgroups of $\mathrm{Isom}\, \mathrm{S}^{\mathrm{n}}$ are uniformly virtually abelian.*

2.1(ii). Euclidean geometry.

To prepare for the hyperbolic case, it will be useful to consider the group $\mathrm{Sim}\, \mathbf{E}^{\mathbf{n}}$ of euclidean similarities. Let $S(\mathbf{E}^{\mathbf{n}})$ be the set of parallel classes of (semi-infinite) geodesic rays in $\mathbf{E}^{\mathbf{n}}$. We shall embed $S(\mathbf{E}^{\mathbf{n}})$ as the unit $(n-1)$-sphere in an inner-product space

$V(\mathbf{E^n})$, which we can imagine as euclidean space with a preferred basepoint. There is an obvious bijective correspondence between r-dimensional subspaces of $V(\mathbf{E^n})$, and foliations of $\mathbf{E^n}$ by parallel r-planes.

The group $\mathrm{Sim}\,\mathbf{E^n}$ acts isometrically on $S(\mathbf{E^n})$, so identifying $S(\mathbf{E^n})$ with $\mathrm{S^{n-1}}$ gives us a homomorphism

$$\mathrm{rot}\ :\ \mathrm{Sim}\,\mathbf{E^n} \longrightarrow \mathrm{Isom}\,\mathrm{S^{n-1}}.$$

We call $\mathrm{rot}\,\gamma$ the *rotational part* of γ. We define

$$U(\mathbf{E^n}) = \{\gamma \in \mathrm{Sim}\,\mathbf{E^n} \mid \mathrm{rot}\,\gamma \in \mathrm{U}(\mathrm{S^{n-1}})\}.$$

Note that if we embed $\mathbf{E^m}$ as a plane in $\mathbf{E^n}$, then $U(\mathbf{E^m})$ may be obtained by intersecting $U(\mathbf{E^n})$ with the stabiliser of this plane. This observation will allow us to use induction over dimension. Given $\gamma \in \mathrm{Sim}\,\mathbf{E^n}$, we shall write

$$\min\gamma = \{\mathrm{x} \in \mathbf{E^n} \mid \mathbf{d}(\mathbf{x}, \gamma\mathbf{x}) \text{ is minimal}\}.$$

Then, $\min\gamma$ is a plane in $\mathbf{E^n}$ on which γ acts either trivially or by translation. Of course, $\min\gamma$ may consist of just a single fixed point.

Theorem 2.1.4 : *If $\Gamma \subseteq \mathrm{Sim}\,\mathbf{E^n}$ is nilpotent, then $\langle \Gamma \cap U(\mathbf{E^n})\rangle$ is abelian.*

We shall begin with a lemma.

Lemma 2.1.5 : *Let Γ be an abelian subgroup of $\mathrm{Sim}\,\mathbf{E^n}$. Let $\sigma(\Gamma) = \bigcap_{\gamma\in\Gamma} \min\gamma$. Then, $\sigma(\Gamma)$ is a non-empty, Γ-invariant plane, on which Γ acts by translations.*

Proof : If Γ is already a translation group, then $\sigma(\Gamma) = \mathbf{E^n}$, and we are done. Otherwise, choose any $\gamma \in \Gamma$ which is not a translation. Then, $\min\gamma$ is a proper subspace, and since Γ is abelian, it is Γ-invariant. The result now follows by induction on dimension. $\diamondsuit$

In fact, our plane $\sigma(\Gamma)$ has a natural foliation by (in general) smaller Γ-invariant planes, namely the set of minimal Γ-invariant planes. That is to say, each leaf is obtained as the affine span of some Γ-orbit. This foliation determines a subspace W_1 of $V(\mathbf{E^n})$, by taking the set of geodesic rays lying in any one leaf. Now, W_1 lies in a larger subspace W' of $V(\mathbf{E^n})$, determined by $\sigma(\Gamma)$ itself. Let W_2 be the orthogonal complement of W_1 in W', and W_3 be the orthogonal complement of W' in $V(\mathbf{E^n})$. This gives us a canonical decomposition $V(\mathbf{E^n}) = \mathbf{W_1} \oplus \mathbf{W_2} \oplus \mathbf{W_3}$. Let m_i be the dimension of W_i. We shall say that the decomposition is trivial if $m_i = n$ for some i.

If $m_1 = n$, then Γ is a pure translation group, and the directions of translations span $\mathbf{E^n}$. If $m_2 = n$, then each point of $\mathbf{E^n}$ is a fixed point of Γ, thus Γ is trivial. If $m_3 = n$, then Γ has a unique fixed point in $\mathbf{E^n}$. We are now ready for:

Proof of Theorem 2.1.4 : Let Γ be a nilpotent subgroup of $\mathrm{Sim}\,\mathbf{E^n}$. We shall assume that Γ is generated by elements of $U(\mathbf{E^n})$, i.e. that $\Gamma = \langle \Gamma \cap U(\mathbf{E^n}) \rangle$. We want to show that Γ is abelian.

Let $Z(\Gamma)$ be the centre of Γ. From the preceding discussion, $Z(\Gamma)$ determines a decomposition $W_1 \oplus W_2 \oplus W_3$ of $V(\mathbf{E^n})$. Since this is canonical, it is respected by the whole group Γ. Thus Γ splits as a subgroup of $\mathrm{Sim}\,\mathbf{E^{m_1}} \times \mathrm{Sim}\,\mathbf{E^{m_2}} \times \mathrm{Sim}\,\mathbf{E^{m_3}}$, and the projection of Γ onto each component is nilpotent. If the decomposition is non-trivial, we may suppose, by induction on dimension, that each projection of Γ is abelian. It then follows that Γ itself is abelian. We need therefore deal only with the cases when the decomposition is trivial.

Suppose $m_1 = n$. This means that $Z(\Gamma)$ is a translation group with no non-empty proper invariant plane in $\mathbf{E^n}$. Consider any $\gamma \in \Gamma$. Since γ commutes with everything in $Z(\Gamma)$, $\min \gamma$ is $Z(\Gamma)$-invariant, and hence equal to $\mathbf{E^n}$. It follows that γ is a translation of $\mathbf{E^n}$. Since translations commute, Γ is abelian.

Suppose $m_2 = n$. Now $Z(\Gamma)$ is trivial. Since Γ is nilpotent, it is also trivial.

Finally, suppose $m_3 = n$. In this case, $Z(\Gamma)$ has a unique fixed point in $\mathbf{E^n}$. This point must be fixed by Γ, so Γ can be regarded as a subgroup of $\mathbf{R}_+ \times \mathrm{Isom}\,\mathrm{S^n}$, where the first component measures the magnification, and the second, the rotational part of an element. The projection into $\mathrm{Isom}\,\mathrm{S^n}$ is nilpotent and generated by elements of $U(\mathrm{S^n})$. By Corollary 2.1.2, this projection is abelian. We deduce that Γ is abelian. $\diamondsuit$

As in the spherical case, for any group Γ, the index of $\langle \Gamma \cap U(\mathbf{E^n}) \rangle$ in Γ is finite, and has a bound dependent only on n. Thus,

Corollary 2.1.6 : *Nilpotent subgroups of* $\mathrm{Sim}\,\mathbf{E^n}$ *are uniformly virtually abelian.*

2.1(iii). Hyperbolic geometry.

We shall write $\mathbf{H^n_I}$ for the ideal $(n-1)$-sphere at infinity of hyperbolic space $\mathbf{H^n}$, and write $\mathbf{H^n_C}$ for the compactification of hyperbolic space as $\mathbf{H^n} \cup \mathbf{H^n_I}$. By a *Möbius transformation* on the sphere $\mathrm{S^n}$, we mean any map which can be represented as a composition of inversions in $(n-1)$-spheres. (We are allowing Möbius transformations that reverse orientation.)

We may represent $\mathbf{H^n}$, conformally, as a hemisphere Σ of $\mathrm{S^n}$. $\mathrm{Isom}\,\mathbf{H^n}$ then consists of those Möbius transformations which preserve Σ. Let γ be a Möbius transformation of $\mathrm{S^n}$, with some fixed point y. Since γ acts conformally, it induces (after scaling) an isometry of the unit tangent space $(T_1\mathrm{S^n})_\mathrm{y}$ at y. Moreover, we may check that if z is any other fixed point of γ, then the induced isometries on $(T_1\mathrm{S^n})_\mathrm{y}$ and $(T_1\mathrm{S^n})_\mathrm{z}$ are conjugate. Thus, γ determines a conjugacy class in $\mathrm{Isom}\,\mathrm{S^n}$, which we shall call $\mathrm{rot}\,\gamma$. Since our subset $U(\mathrm{S^n})$ of $\mathrm{Isom}\,\mathrm{S^n}$ is invariant under conjugacy, it makes sense to demand that $\mathrm{rot}\,\gamma$ should lie in $U(\mathrm{S^n})$. Restricting to $\mathrm{Isom}\,\mathbf{H^n}$, where all Möbius transformations have fixed points, we may define

$$U(\mathbf{H^n}) = \{\gamma \in \mathrm{Isom}\,\mathbf{H^n} \mid \mathrm{rot}\,\gamma \subseteq \mathrm{U(S^n)}\}.$$

Theorem 2.1.7 : *If* $\Gamma \subseteq \mathrm{Isom}\,\mathbf{H^n}$ *is nilpotent, then* $\langle \Gamma \cap U(\mathbf{H^n}) \rangle$ *is abelian.*

We begin with two lemmas.

Lemma 2.1.8 : *If* $\Gamma \subseteq \mathrm{Isom}\,\mathbf{H^n}$ *is abelian, then* $\mathrm{fix}\,\Gamma$*, the set of points fixed by* Γ*, consists of either one or two points in* $\mathbf{H^n_I}$*, or else is a subspace of* $\mathbf{H^n_C}$ *(i.e. the closure, in* $\mathbf{H^n_C}$*, of a plane in* $\mathbf{H^n}$*).*

Proof : Let γ be any non-trivial element of Γ. If γ is parabolic, then its fixed point is preserved by Γ, so that Γ has a unique fixed point. If γ is elliptic, then $\mathrm{fix}\,\gamma$ is a proper Γ-invariant subspace, and we use induction on dimension. For this, we need to check the 1-dimensional case. But it is easily seen that an abelian group of isometries of the real line must either act trivially, or by translation (thus respecting the two "ideal" points), or else consist of an involution with a single fixed point. Finally, if γ is loxodromic, then its axis is Γ-invariant, and we are immediately reduced to the 1-dimensional case. $\diamondsuit$

Lemma 2.1.9 : *Suppose* $\Gamma \subseteq \mathrm{Isom}\,\mathbf{H^n}$ *is nilpotent, then* Γ *has a fixed point in* $\mathbf{H^n_C}$*.*

Proof : Let σ be the set of points fixed by the centre $Z(\Gamma)$. Let $\Gamma' \supseteq Z(\Gamma)$ be the subgroup that fixes σ pointwise. Since σ is canonical with respect to Γ, Γ' is normal in Γ. Thus Γ/Γ' is nilpotent, and acts effectively on σ.

From Lemma 2.1.8, we distinguish three possibilities for σ. Firstly, if σ is a single point of $\mathbf{H^n_I}$, this point is fixed by Γ, and we are done. Secondly, if σ is a proper subspace of $\mathbf{H^n_C}$, we use induction on dimension. Thus, we may assume that we are in the third case, namely that σ consists of precisely two points, x and y, in $\mathbf{H^n_I}$. If Γ/Γ' is trivial, we are done. Therefore we may suppose that Γ/Γ' is an involution. This means that there is some $\gamma \in \Gamma$ that swaps x and y. Now, each element of $Z(\Gamma)$ fixes x and y, and commutes with γ. We see that $Z(\Gamma)$ must fix pointwise the geodesic joining x and y. This contradicts the definition of σ as $\mathrm{fix}\,\mathrm{Z}(\Gamma)$. $\diamondsuit$

Proof of Theorem 2.1.7 : By Lemma 2.1.9, $\langle \Gamma \cap U(\mathbf{H^n}) \rangle$ fixes some point, x, of $\mathbf{H^n_C}$. If $x \in \mathbf{H^n}$, we are reduced to the spherical case, and if $x \in \mathbf{H^n_I}$, we are reduced to the case of euclidean similarities. We observe that our definitions of the rotational part of an isometry (or similarity) are in agreement, so that the theorem follows from Corollary 2.1.2, and Theorem 2.1.4. $\diamondsuit$

For completeness, we state:

Corollary 2.1.10 : *Nilpotent subgroups of* $\mathrm{Isom}\,\mathbf{H^n}$ *are uniformly virtually abelian.*

Proof : If $\Gamma \subseteq \mathrm{Isom}\,\mathbf{H^n}$ is nilpotent, we need that $[\Gamma : \langle \Gamma \cap U(\mathbf{H^n}) \rangle]$ is uniformly bounded. But by Lemma 2.1.8, Γ has a fixed point in $\mathbf{H^n_C}$, so the result follows from the spherical and euclidean (similarity) cases. $\diamondsuit$

Note that all the abelian subgroups constructed in this section are normal, since the neighbourhoods $U(\mathrm{S^n})$, $U(\mathbf{E^n})$ and $U(\mathbf{H^n})$ are all conjugacy invariant.

2.2. Discrete subgroups.

In this section, we describe how nilpotent groups occur naturally when considering discrete group actions.

Let g be a Lie group, and let $|\ |$ be any smooth norm on G, for example, distance from the identity in some Riemannian metric. For any g, $h \in G$, sufficiently near the identity, we will have $|[g,h]| < C|g||h|$, for some constant C. Thus, we can find a bounded symmetric neighbourhood, $O(G)$ of the identity in G such that whenever g, $h \in O(G)$, we have $[g,h] \in O(G)$ and $|[g,h]| < |g|/2$.

Lemma 2.2.1 : *If Γ is a discrete subgroup of G, then $\langle \Gamma \cap O(G) \rangle$ is nilpotent.*

Proof : The elements of Γ have norms bounded below by some number $c > 0$, and the elements of $O(G)$ have norms bounded above by some number k. If m is any integer greater than $\log_2(k/c)$, we see that any m-fold commutator in elements of $\Gamma \cap O(G)$ will be trivial. By repeated application of the identity $[xy,z] = [x,[y,z]][y,z][x,z]$, we deduce that any m-fold commutator in $\langle \Gamma \cap O(G) \rangle$ is trivial. Thus, $\langle \Gamma \cap O(G) \rangle$ is nilpotent. $\diamondsuit$

The following lemma is a modified version of one to be found in [Th2]. It is relevant to our discussion of the Margulis Lemma. First, we introduce some notation. Given a subset X of the a group G, we write X^r for those $g \in G$ expressible as words of length at most r, in elements of X together with their inverses (in X^{-1}), i.e. inductively, $X^1 = X \cup \{1\} \cup X^{-1}$, $X^r = X^{r-1}X^1$. If Γ is a subgroup of G, we write Γ_X for $\langle \Gamma \cap X \rangle$.

Lemma 2.2.2 : *Let G be a (locally compact) topological group, with W a neighbourhood of the identity. Let $K_i, i \in \mathbf{N}$ be a sequence of symmetric neighbourhoods of the identity. Suppose K_1 is compact, and $(K_i)^i \subseteq K_1$ for each i. Then, there exists some $N \in \mathbf{N}$ such that for any discrete group $\Gamma \leq G$, $[\Gamma_{K_N} : \langle \Gamma_{K_N} \cap W \rangle] \leq N$.*

Proof : Let V be a neighbourhood of 1 with $V^{-1}V \subseteq W$. Since K_1 is compact, there is an upper bound, k, on the number of right translates $Vg, g \in K_1$, of V, that we can pack disjointly into G. Let $N = k+1$.

Suppose that $\Gamma \leq G$ is discrete. Let $\{Va_i | i = 1, \ldots, p\}$ be a disjoint packing with $a_i \in \Gamma_{K_N} \cap K_1$, and p maximal. Note that $p \leq k$. Write $\Gamma_N = \langle \Gamma_{K_N} \cap W \rangle$. We claim that $\{\Gamma_N a_i | i = 1, \ldots, p\}$ includes a complete set of cosets for Γ_N in Γ_{K_N}, so that $[\Gamma_{K_N} : \Gamma_N] \leq N$, as required.

To see this, consider $\Gamma_N h$ with $h \in \Gamma_{K_N}$. Write $h = \prod_{i=1}^{l} g_i$, with $g_i \in \Gamma \cap K_N$. If $l \geq k+1$, consider the collection $\{Vh_j | j = 1, \ldots, k+1\}$, where $h_j = \prod_{i=1}^{j} g_i$, so that $h_j \in (K_N)^N \subseteq K_1$. These sets cannot all be disjoint. Thus, we can write $h = \alpha\beta\gamma$, with $\alpha\beta \in K_1$ and $V\alpha \cap V\alpha\beta \neq \emptyset$. Now, $\alpha\beta\alpha^{-1} \in V^{-1}V \subseteq W$, so $\alpha\beta\alpha^{-1} \in \Gamma_N$. Thus, $\Gamma_N h = \Gamma_N(\alpha\beta\alpha^{-1})\alpha\gamma = \Gamma_N h'$, where $h' = \alpha\gamma$. We have reduced the word-length of h, so, by induction, $\Gamma_N h = \Gamma_N h''$, with $h'' \in K_1$. But then, $Vh'' \cap Va_i \neq \emptyset$, for some a_i, so that $h'' a_i^{-1} \in W$, and $\Gamma_N h'' = \Gamma_N a_i$. Hence, $\Gamma_N h = \Gamma_N a_i$.

$\diamondsuit$

We again consider the three geometries in turn.

2.2(i). Spherical geometry.

We write $U_0(\mathrm{S^n})$ for $O(\mathrm{Isom}\,\mathrm{S^n})$, the neighbourhood of the identity defined at the beginning of Section 2.2. Since this set may be chosen to be arbitrarily small, we may suppose that $U_0(\mathrm{S^n}) \subseteq \mathrm{U}(\mathrm{S^n})$. We may also suppose that $U_0(\mathrm{S^n})$ is conjugacy invariant. Now if Γ is a discrete subgroup of $\mathrm{Isom}\,\mathrm{S^n}$, then $\langle \Gamma \cap U_0(\mathrm{S^n})\rangle$ is nilpotent by Lemma 2.2.1, and thus abelian by Corollary 2.1.2. It is easily checked that $\langle \Gamma \cap U_0(\mathrm{S^n})\rangle$ has a finite index in Γ, which is bounded as Γ varies. Thus we have:

Lemma 2.2.3 (Jordan Lemma) : *Discrete subgroups of* $\mathrm{Isom}\,\mathrm{S^n}$ *are uniformly virtually abelian.*

2.2(ii). Euclidean geometry.

We can assume that $O(\mathrm{Isom}\,\mathbf{E^n})$ has the form

$$O(\mathrm{Isom}\,\mathbf{E^n}) = \{\gamma \in \mathrm{Isom}\,\mathbf{E^n} \mid \mathbf{d}(\gamma\mathbf{a}, \mathbf{a}) < \epsilon \text{ and } \mathrm{rot}\,\gamma \in \mathrm{U}_1\}$$

where $\epsilon > 0$, a is some point of $\mathbf{E^n}$, and U_1 is some neighbourhood of the identity in $\mathrm{Isom}\,\mathrm{S^{n-1}}$ that is contained in $U(\mathrm{S^{n-1}})$. For notational convenience we shall identify U_1 with the set $U_0(\mathrm{S^{n-1}})$ of the Jordan Lemma. We set

$$U_0(\mathbf{E^n}) = \{\mathbf{\Gamma} \in \mathrm{Isom}\,\mathbf{E^n} \mid \mathrm{rot}\,\gamma \in \mathrm{U}_0(\mathrm{S^n})\}.$$

Proposition 2.2.4 : *Suppose that* Γ *is a discrete subgroup of* $\mathrm{Isom}\,\mathbf{E^n}$*; then* $\langle \Gamma \cap U_0(\mathbf{E^n})\rangle$ *is abelian.*

Proof : To begin with, we do not know that $\langle \Gamma \cap U_0(\mathbf{E^n})\rangle$ is finitely generated, so we proceed as follows. Let $D_r = \{\gamma \in \mathrm{Isom}\,\mathbf{E^n} \mid \mathbf{d}(\gamma\mathbf{a}, \mathbf{a}) < \mathbf{r}\epsilon\}$. Let g_r be the dilation of magnification r about a. Considering $\Gamma_r = \langle \Gamma \cap U_0(\mathbf{E^n}) \cap \mathbf{D_r}\rangle$, we see that

$$\begin{aligned} g_r^{-1}\Gamma_r g_r &= \langle g_r^{-1}\Gamma g_r \cap U_0(\mathbf{E^n}) \cap \mathbf{D_1}\rangle \\ &= \langle g_r^{-1}\Gamma g_r \cap O(\mathrm{Isom}\,\mathbf{E^n})\rangle \end{aligned}$$

which is nilpotent (Lemma 2.2.1) and hence abelian (Theorem 2.1.4). Thus, Γ_r is abelian for all r, and so $\langle \Gamma \cap U_0(\mathbf{E^n})\rangle = \bigcup_{\mathbf{r}} \mathbf{\Gamma_r}$ is abelian. ◊

Again, it is easily seen that $\langle \Gamma \cap U_0(\mathbf{E^n})\rangle$ has bounded index in Γ, so we have:

Theorem 2.2.5 (Bieberbach) : *Discrete subgroups of* $\mathrm{Isom}\,\mathbf{E^n}$ *are uniformly virtually abelian.* $\diamondsuit$

Note that since $U_0(\mathrm{S^n})$ is conjugacy invariant in $\mathrm{Isom}\,\mathrm{S^n}$, the abelian subgroups we produce in this way will be normal. We shall write $\nu(n)$ for the bound on their index.

We can say a little more about the structure of discrete euclidean groups. Clearly, if two subspaces, $\tau_1, \tau_2 \subseteq \mathbf{E^n}$ are Γ-invariant, then $\tau_1 \cap \tau_2$ is Γ-invariant. We may thus speak of minimal (non-empty) Γ-invariant subspaces of $\mathbf{E^n}$. Clearly such minimal subspaces must always exist.

Proposition 2.2.6 : *Suppose* Γ *acts properly discontinuously on* $\mathbf{E^n}$. *Then, a* Γ-*invariant subspace,* $\mu \subseteq \mathbf{E^n}$, *is minimal if and only if* μ/Γ *is compact. Moreover, any two such subspaces are parallel, and the action of* Γ *commutes with the perpendicular translation between them.*

Proof : Clearly, if μ/Γ is compact, then μ is minimal.

Let μ_1 and μ_2 be two minimal subspaces. Let $\lambda(\mu_i, \mu_j) = \{x \in \mu_i | d(x, \mu_i) = d(\mu_i, \mu_j)\} \subseteq \mu_i$. Γ preserves $\lambda(\mu_i, \mu_j)$. Hence, by minimality, $\lambda(\mu_i, \mu_j) = \mu_i$. It follows easily that μ_1 and μ_2 must be parallel.

Given any two parallel subspaces in $\mathbf{E^n}$, there is a unique perpendicular translation mapping one to the other. Any isometry that preserves these two planes must commute with this translation. It follows that the action of Γ on $\mathbf{E^n}$ must commute with the perpendicular translation sending μ_1 to μ_2.

It now remains to show that if Γ acts minimally on $\mathbf{E^n}$, then it is cocompact. From the Bieberbach theorem, and the discussion of abelian groups in Section 2.1(ii), we can find a normal abelian subgroup Γ', of finite index in Γ, and a plane $\tau \leq \mathbf{E^n}$, on which Γ' acts as a cocompact translation group. There are finitely many images, $\{\tau_1, \ldots, \tau_k\}$, of τ under Γ, each preserved by Γ'. Since a cocompact action is minimal, it follows that the τ_i are all parallel. We may now find τ', parallel to τ, which represents the centre of mass of the τ_i in any transverse plane. Now, Γ preserves τ', so, by minimality, $\tau' = \mathbf{E^n} = \tau$. Hence, $\mathbf{E^n/\Gamma}$ is compact.

$\diamondsuit$

As in the earlier discussion of the abelian case (Section 2.1(ii)), it is easily seen that the set of minimal planes in $\mathbf{E^n}$ form a foliation of a larger, canonical subspace.

2.2(iii). Hyperbolic geometry.

Given $x \in \mathbf{H^n}$, we write

$$\mathcal{I}_\epsilon(x) = \{\gamma \in \mathrm{Isom}\,\mathbf{H^n} \mid \mathbf{d}(\gamma\mathbf{x}, \mathbf{x}) \leq \epsilon\}.$$

Let d_1 be any Riemannian metric on the unit tangent bundle $T_1\mathbf{H^n}$ of $\mathbf{H^n}$, invariant under the action of $\mathrm{Isom}\,\mathbf{H^n}$. Given $x \in \mathbf{H^n}$, we write

$$\mathcal{I}'_\epsilon(x) = \{\gamma \in \mathrm{Isom}\,\mathbf{H^n} \mid \mathbf{d_1}(\gamma\vec{\mathbf{v}}, \vec{\mathbf{v}}) < \epsilon \text{ for each unit vector } \vec{\mathbf{v}} \text{ based at } \mathbf{x}\}.$$

If Γ is a subgroup of $\mathrm{Isom}\,\mathbf{H^n}$, we write

$$\Gamma_\epsilon(x) = \langle \Gamma \cap \mathcal{I}_\epsilon(x) \rangle$$

and

$$\Gamma'_\epsilon(x) = \langle \Gamma \cap \mathcal{I}'_\epsilon(x) \rangle.$$

Now we may suppose that $O(\mathrm{Isom}\,\mathbf{H^n})$ has the form $\mathcal{I}'_{\epsilon_1}(x)$ for some $\epsilon_1 > 0$ and $x \in \mathbf{H^n}$. We also assume that $\mathcal{I}'_{\epsilon_1}(x) \subseteq U(\mathbf{H^n})$. We now have:

Proposition 2.2.7 : *If Γ is a discrete subgroup of* $\mathrm{Isom}\,\mathbf{H^n}$, *then* $\Gamma'_{\epsilon_1}(x) = \langle \Gamma \cap \mathcal{I}'_{\epsilon_1}(x) \rangle$ *is abelian.*

Note that, by homogeneity, this remains true if we fix ϵ_1, and choose x arbitrarily.

We next show that for small ϵ, $\Gamma_\epsilon(x)$ is virtually abelian. To this end, we take $\mathcal{I}'_{\epsilon_1}(x)$ to be the set W of Lemma 2.2.2, and the sets K_r to be $\mathcal{I}_{1/r}(x)$. The lemma now tells us that, for some $N > 0$,

$$[\Gamma_{\epsilon(n)}(x) : \langle \Gamma_{\epsilon(n)}(x) \cap \mathcal{I}'_{\epsilon_1}(x) \rangle] \leq N,$$

where $\epsilon(n) = 1/N$. Thus,

$$[\Gamma_{\epsilon(n)}(x) : \Gamma_{\epsilon(n)}(x) \cap \Gamma'_{\epsilon_1}(x)] \leq N.$$

For notational convenience, we shall assume that $N \leq \nu(n)$, the constant of the Bieberbach Theorem, and that we have chosen the metric on $T_1\mathbf{H^n}$ so that $\epsilon_1 = \epsilon(n)$. We call $\epsilon(n)$ the *Margulis constant.* In summary, we have:

Theorem 2.2.8 (Margulis Lemma) : *For all n, there exist $\epsilon(n) > 0$ and $\nu(n) \in \mathbf{N}$ such that if Γ is any discrete subgroup of* $\mathrm{Isom}\,\mathbf{H^n}$, *and $x \in \mathbf{H^n}$, then $\Gamma_{\epsilon(n)}(x)$ has an abelian subgroup (namely $\Gamma'_{\epsilon(n)}(x) \cap \Gamma_{\epsilon(n)}(x)$) of index at most $\nu(n)$.*

Note that if $0 < \epsilon \leq \epsilon(n)$, then $\Gamma_\epsilon(x) \cap \Gamma'_{\epsilon(n)}(x)$ has index at most $\nu(n)$ in $\Gamma_\epsilon(x)$. By intersecting all conjugate subgroups to $\Gamma_\epsilon(x) \cap \Gamma'_{\epsilon(n)}(x)$, we see that $\Gamma_\epsilon(x)$ contains a normal abelian subgroup of bounded index, where the bound is independent of the choice of discrete subgroup Γ.

Chapter 3 : Five definitions of geometrical finiteness.

3.1. Parabolic groups and definition GF1.

Definition : A subgroup G of $\mathrm{Isom}\,\mathbf{H^n}$ is *parabolic* if $\mathrm{fix}\,G$ consists of a single point $p \in \mathbf{H^n_I}$, and if G preserves setwise some horosphere about p.

It follows that G contains a parabolic element and preserves setwise every horosphere about p. If we represent $\mathbf{H^n}$ by the upper half-space model $\mathbf{R^n_+}$, with the fixed point p at ∞, then G acts by euclidean isometry on $\mathbf{R^n_+} \cup \partial\mathbf{R^n_+}$.

Suppose now that G is a discrete parabolic subgroup. Then, the limit set $\Lambda(G) = \{p\}$. Thus, $M_C(G) = (\mathbf{H^n} \cup \mathbf{\Omega(G)})/\mathbf{G} = (\mathbf{H^n_C} \setminus \{\mathbf{p}\})/\mathbf{G}$. By Proposition 2.2.6, there is a G-invariant euclidean subspace $\mu \subseteq \partial\mathbf{R^n_+}$ with μ/G compact. The subspace μ may not be unique, but any two such will be euclidean-parallel. Note that there is a unique closed hyperbolic subspace, σ of $\mathbf{H^n}$, with $\sigma \cap \mathbf{H^n} \setminus \{\mathbf{p}\} = \mu$. In the upper half-space model, we may write $\sigma \setminus \{p\} = \mu \times [0,\infty) \subseteq \mathbf{R^{n-1}} \times [\mathbf{0},\infty) \equiv \mathbf{R^n_+} \cup \partial\mathbf{R^n_+}$. Thus, σ is G-invariant, contains the point p, and $(\sigma \cap \partial B)/G$ is compact for any horoball B about p.

Given any $r > 0$, we write

$$C(\mu,r) = \{x \in \mathbf{R^n_+} \cup \partial\mathbf{R^n_+} \mid \mathbf{d_{euc}(x},\mu) \geq \mathbf{r}\}.$$

Thus $C(\mu,r)$ is G-invariant, and hyperbolically convex (Figure 3a). We have $C(\mu,r)/G \subseteq M_C(G)$ (Figure 3b). The complement of $C(\mu,r)/G$ in $M_C(G)$ is relatively compact. Clearly $\bigcap_{r\in[0,\infty)} C(\mu,r) = \emptyset$, and so we have:

Lemma 3.1.1 : *If $G \subseteq \mathrm{Isom}\,\mathbf{H^n}$ is discrete parabolic, then $M_C(G)$ has precisely one topological end. Moreover, the collection $\{C(\mu,r)/G \mid r \geq 0\}$, as described above, forms a base of neighbourhoods for that end.*

Definition : We call a set of the form $C(\mu,r)$, for some r and μ, a *standard parabolic region*.

Let $\rho_\sigma : \mathbf{H^n_C} \longrightarrow \sigma$ be the nearest point retraction (as described in Chapter 4). Let B be the horoball about p of euclidean height r in $\mathbf{R^n_+}$, i.e. such that $d_{euc}(\partial B, \partial\mathbf{R^n_+}) = \mathbf{r}$. Thus, $\sigma \cap \partial B = \sigma \cap \partial C(\mu,r) = \partial B \cap \partial C(\mu,r)$. It is not hard to see that $C(\mu,r) = \rho_\sigma^{-1}(\sigma \cap \partial B)$. This gives us an alternative way of defining standard parabolic regions without explicit reference to the upper half-space model.

We are primarily interested in parabolic subgroups of more general discrete isometry groups. The following is an essential lemma.

Lemma 3.1.2 : *Suppose the elements $g, h \in \mathrm{Isom}\,\mathbf{H^n}$ have a common fixed point $p \in \mathbf{H^n_I}$. Suppose that g is parabolic, and that h is loxodromic. Then, the subgroup $\langle g,h\rangle$ generated by g and h is not discrete.*

Proof : Without loss of generality, we can suppose that p is the attracting fixed point of h. Then the sequence $(h^i g h^{-i})_{i\in\mathbf{N}}$ of elements of $\langle g,h\rangle$ tends to the identity. $\diamondsuit$

By a similar argument, we have:

Lemma 3.1.3 : *If $g, h \in \mathrm{Isom}\,\mathbf{H^n}$ ar loxodromic, and have pricisely one common fixed point, then $\langle g, h \rangle$ is not discrete.*

Now, suppose that Γ is any discrete subgroup of $\mathrm{Isom}\,\mathbf{H^n}$. Suppose that $G \subseteq \Gamma$ is a parabolic subgroup with fixed point p. By Lemma 3.1.2, we see that $\mathrm{stab}_\Gamma\, p$ cannot contain any loxodromic element. Thus $\mathrm{stab}_\Gamma\, p$ is also a parabolic subgroup. In fact, it is clear that $\mathrm{stab}_\Gamma\, p$ is maximal parabolic. This shows:

Lemma 3.1.4 : *Suppose $\Gamma \subseteq \mathrm{Isom}\,\mathbf{H^n}$ is discrete. Then every parabolic subgroup of Γ is contained in a unique maximal parabolic subgroup.*

Definition : We say that $p \in \mathbf{H^n_I}$ is a *parabolic fixed point* of Γ if $\mathrm{stab}_\Gamma\, p$ is parabolic.

We see that p is a parabolic fixed point if and only if it is fixed by some parabolic element of Γ. Note that every parabolic fixed point lies in the limit set, $\Lambda(\Gamma)$, of Γ.

If $\gamma \in \Gamma$, then clearly $\mathrm{stab}_\Gamma(\gamma p) = \gamma(\mathrm{stab}_\Gamma\, p)\gamma^{-1}$. Thus there is a bijective correspondence between Γ-orbits of parabolic fixed points, and conjugacy classes of maximal parabolic subgroups of Γ.

Suppose p is a parabolic fixed point of Γ. Suppose it happens that we can find a standard parabolic region, C, for the group $G = \mathrm{stab}_\Gamma\, p$, with the property that $C \subseteq \mathbf{H^n} \cup \mathbf{\Omega(\Gamma)}$, and that $\gamma C \cap C = \emptyset$ if $\gamma \in \Gamma \setminus G$. Then C descends to a set $E = (\bigcup \Gamma C)/\Gamma \subseteq M_C(\Gamma)$, which is isometric to C/G. (Strictly speaking, we mean that $E \cap M(\Gamma)$ and $(C/G) \cap M(G)$ are isometric.) We see that E is a neighbourhood of an end of $M_C(\Gamma)$. We call such a set, E, a *standard cusp region* in $M_C(\Gamma)$. In such a case, we say that the parabolic fixed point p (or more accurately, the orbit Γp) is *associated* to the standard cusp region E.

Suppose that p is associated to the region E as above. We may represent $\mathbf{H^n}$ by the upper half-space model $\mathbf{R^n_+}$ with p at ∞. By by hypothesis, $C \subseteq \mathbf{H^n} \cup \mathbf{\Omega}$, and so if $C = C(\mu, r)$, we see that $\Lambda \setminus \{p\}$ lies inside the euclidean r-neighbourhood of μ in $\partial \mathbf{R^n_+}$, i.e. $d_{euc}(x, \mu) \leq r$ for all $x \in \Lambda \setminus \{p\}$. Since μ/G is compact, and Λ is closed, we see that $(\Lambda \setminus \{p\})/G$ is compact.

Definition : A parabolic fixed point $p \in \mathbf{H^n_I}$ is *bounded* if $(\Lambda \setminus \{p\})/\,\mathrm{stab}_\Gamma\, p$ is compact.

In other words, p is bounded if and only if $d_{euc}(x, \mu)$ is bounded for $x \in \Lambda \setminus \{p\}$, where μ is some (or every) minimal $(\mathrm{stab}_\Gamma\, p)$-invariant subspace of $\partial \mathbf{R^n_+} = \mathbf{H^n_I} \setminus \{\mathbf{p}\}$.

We have seen:

Lemma 3.1.5 : *If the parabolic fixed point p is associated to a standard cusp region of $M_C(\Gamma)$, then p is bounded.*

In fact, we shall see in Chapter 4 that the converse of of Lemma 3.1.5 is also true (Corollary 4.5).

Definition (GF1) : Suppose Γ is a discrete subgroup of $\mathrm{Isom}\,\mathbf{H^n}$. We say that Γ is

"GF1" if we may write $M_C(\Gamma)$ as the union of a compact set and a finite number of disjoint standard cusp regions.

Suppose that Γ is GF1, and that E_i for $i \in \{1, 2, \dots, k\}$ are the standard cusp regions given by the hypothesis. We can take the compact set K to be the closure, in $M_C(\Gamma)$, of $M_C(\Gamma) \setminus \bigcup_{i=1}^k E_i$. In this way, K is a suborbifold of $M_C(\Gamma)$ (Figure 3c). In fact, $M_C(\Gamma)$ retracts onto K. As a simple consequence, we have:

Proposition 3.1.6 : *If Γ is GF1, then Γ is finitely generated.*

We also have:

Proposition 3.1.7 : *If Γ is GF1, then there is a bound on the orders of finite subgroups of Γ.*

Proof : We can just apply the Selberg Lemma (Chapter 1). Alternatively, one can give a direct proof by noting that if $G \subseteq \Gamma$ is finite, then fix$G$ meets the lift of the compact set K to $\mathbf{H^n} \cup \mathbf{\Omega}$. $\diamondsuit$

In fact, if Γ is GF1, then it has finitely many conjugacy classes of finite subgroups.

We note that if Γ is GF1, then the quotient, $M(\Gamma)$ is topologically finite, i.e. it is orbifold-homeomorphic to the interior of a compact orbifold with boundary.

Another way to say that Γ is GF1 is to say that $M_C(\Gamma)$ has only finitely many ends, and that each such end has a neighbourhood "isometric" to a neighbourhood of the end of $M_C(G)$, where G is a discrete parabolic group. When we say that two such sets are "isometric", we really mean that their metric (non-ideal) parts are isometric. We remark that this definition also makes sense for manifolds of variable negative curvature.

3.2. Conical limit points and definition GF2.

Suppose $\Gamma \subseteq \mathrm{Isom}\,\mathbf{H^n}$ is discrete. We write $\pi : \mathbf{H^n} \cup \mathbf{\Omega(\Gamma)} \longrightarrow \mathbf{M_C(\Gamma)}$ for the projection.

Definition : We say $y \in \mathbf{H^n_I}$ is a *conical limit point* if for some (and hence every) geodesic ray, β, tending to y, and some (and hence every) point $x \in \mathbf{H^n}$, there is a sequence (γ_i) of elements of Γ with $\gamma_i x \to y$ and $d(\gamma_i x, \beta)$ bounded. (Elsewhere, conical limit points have been called "approximation points" or "radial limit points")

It is not hard to see that this is equivalent to the following statement. Suppose β is a geodesic ray tending to y. Then y is a conical limit point if and only if $\pi\beta$ accumulates somewhere in $M = M(\Gamma)$, i.e. there is a sequence of points, $x_i \in \beta$, with $x_i \to y$ and (πx_i) convergent in M.

Note that saying that $\pi\beta$ accumulates somewhere in M is the same as saying that the collection of images, $\Gamma\beta$, of β under Γ is not locally finite in $\mathbf{H^n}$. The following tells us that if $\pi\beta$ has an accumulation point in $M_C(\Gamma)$, then it has an accumulation point in $M(\Gamma)$.

Lemma 3.2.1 : *Suppose $\Gamma \subseteq \mathrm{Isom}\,\mathbf{H^n}$ is discrete, and that $Q \subseteq \mathbf{H^n}$ is convex. Suppose that the set of images, ΓQ, of Q under Γ is locally finite in $\mathbf{H^n}$. (We take this to imply that $\mathrm{stab}_\Gamma Q$ is finite.) Then ΓQ is locally finite on $\mathbf{H^n} \cup \mathbf{\Omega(\Gamma)}$.*

Proof : Suppose $y \in \Omega$. Let H_1, H_2, H_3 be three closed half spaces with $y \in \mathrm{int}\, H_1$, $H_1 \subseteq \mathrm{int}\, H_2$, $H_2 \subseteq \mathrm{int}\, H_3$ and $H_3 \subseteq \mathbf{H^n} \cup \mathbf{\Omega}$. Thus, $\partial H_2 \cap \mathrm{hull}(\partial H_1 \cup \partial H_3)$ is a compact subset of $\mathbf{H^n}$ (Figure 3d). Clearly only finitely many images of Q can lie inside H_3. If γQ meets H_1, but is not contained in H_3, then γQ must meet $\partial H_2 \cap \mathrm{hull}(\partial H_1 \cup \partial H_3)$. Since ΓQ is locally finite on $\mathbf{H^n}$, this can happen for only finitely many γ. Thus, only finitely many images of Q can meet H_1. $\diamondsuit$

(Lemma 3.2.1, will also be used in relation to convex cell complexes, Section 3.5.)

Definition (GF2) : Suppose that Γ is a discrete subgroup of $\mathrm{Isom}\,\mathbf{H^n}$. We say that Γ is "GF2" if the limit set, $\Lambda(\Gamma)$, consists entirely of conical limit points and bounded parabolic fixed points.

It is easily seen that, for any discrete group, the set of bounded parabolic fixed points and the set of conical limit points are disjoint. In fact no parabolic fixed point can also be a conical limit point [SusS]. We shall see (Lemma 4.6) that in a geometrically finite group, every parabolic fixed point is bounded.

Beardon and Maskit [BeaM] give several equivalent definitions of conical limit point, including one that makes sense in $\mathbf{H^n_I}$. Thus, we have a definitions of geometrical finiteness intrinsic to the action of Γ on $\mathbf{H^n_I}$. Indeed, one can give a definition intrinsic to the limit set Λ (see [Bow]).

We remark that if Γ is geometrically finite, then the convergence to conical limit points is "uniform", as defined in [BeaM]. One way to say this is that there is a compact set $K \subseteq \mathbf{H^n}$, such that if y is conical limit point, and β is any geodesic tending to y, then $\Gamma\beta$ accumulates in K. The argument in [BeaM] generalises to any dimension to show that the limit set of any geometrically finite group has either zero or full spherical Lebesgue measure [Ap1]. Another simple proof of this fact, uses the definition GF5 (see [Th1, Chapter 8] for the 3-dimensional case).

3.3. The thick-thin decomposition.

In this section, we describe how to divide a hyperbolic orbifold into a "thick" and a "thin" part. Our construction agrees with the usual definition in the case of manifolds (Γ torsion-free).

We have already talked about parabolic subgroups of $\mathrm{Isom}\,\mathbf{H^n}$. Another important class are what we shall call "loxodromic" subgroups.

Definition : A subgroup, G, of $\mathrm{Isom}\,\mathbf{H^n}$ is *loxodromic* if G preserves setwise some bi-infinite geodesic, and contains a loxodromic element.

The bi-infinite geodesic, β, referred to must be unique. (If α were another G-invariant geodesic, then the set of endpoints of α and β would give us a G-invariant subset of $\mathbf{H^n_I}$ with at least three elements, and so give us a fixed point for G in $\mathbf{H^n}$. This contradicts the existence of a loxodromic element of G.) We shall call the geodesic β the *loxodromic axis*.

Suppose that Γ is any discrete subgroup of $\mathbf{H^n}$. Suppose that $G \subseteq \Gamma$ is loxodromic with axis β. Then, $\mathrm{stab}_\Gamma\,\beta$ is also loxodromic, in fact, a maximal loxodromic subgroup of Γ. We see (c.f. Lemma 3.1.3) that:

Lemma 3.3.1 : *If $\Gamma \subseteq \mathrm{Isom}\,\mathbf{H^n}$ is discrete, then every loxodromic subgroup of Γ is contained in a unique maximal loxodromic subgroup.*

We shall need:

Lemma 3.3.2 : *If $G \subseteq \mathrm{Isom}\,\mathbf{H^n}$ is discrete and virtually abelian, then G is finite, parabolic or loxodromic.*

Proof : Let $N \subseteq G$ be a normal abelian finite-index subgroup. We can suppose that N is infinite, and thus contains an element of infinite order. By Lemma 2.1.9, we know that $\mathrm{fix}\,N$ consists of either one or two points of $\mathbf{H^n_I}$. Thus, $X = \bigcup \Gamma(\mathrm{fix}\,N)$ is a non-empty finite G-invariant subset of $\mathbf{H^n_I}$. Since G is infinite, X has at most two elements. If X has precisely two points, then G preserves the bi-infinite geodesic joining them, and is thus loxodromic. Thus, we suppose that $X = \{p\}$. If G contains a parabolic element, then by Lemma 3.1.2, G contains no loxodromic element, and must therefore be a parabolic group. On the other hand, if G contained no parabolic element, then it must contain a loxodromic element g, and an element h which does not preserve the axis of g. Thus, hgh^{-1} is another loxodromic, which shares precisely one fixed point with g. Lemma 3.1.3 now gives us a contradiction. $\diamondsuit$

We shall now describe the thick-thin decomposition. A more detailed account is given in [Bow].

Suppose $\Gamma \subseteq \mathrm{Isom}\,\mathbf{H^n}$ is discrete. Suppose $\epsilon > 0$, and $x \in \mathbf{H^n}$. In Section 2.2(iii), we defined $\Gamma_\epsilon(x)$ to be the subgroup of Γ generated by those elements that move the point x a distance at most ϵ. We set

$$T_\epsilon(\Gamma) = \{x \in \mathbf{H^n} \mid \mathbf{\Gamma_\epsilon(x)} \text{ is infinite}\}.$$

Thus, $T_\epsilon(\Gamma)$ is a closed Γ-invariant subset of $\mathbf{H^n}$. We first consider the cases of parabolic and loxodromic groups.

Suppose $G \subseteq \mathrm{Isom}\,\mathbf{H^n}$ is discrete parabolic with fixed point p. It is not hard to see that $T_\epsilon(G)$ is connected, and meets every bi-infinite geodesic through p in a geodesic ray tending to p. It follows that $T_\epsilon(G)$ is closed in $\mathbf{H^n_C} \setminus \{\mathbf{p}\}$. Note that $T_\epsilon(G)/G$ is connected.

Suppose $G \subseteq \mathrm{Isom}\,\mathbf{H^n}$ is discrete loxodromic with axis β. Then $T_\epsilon(G)$ retracts onto $\beta \cap T_\epsilon(G)$. If $G_0 \triangleleft G$ is the subgroup which fixes β pointwise, then G/G_0 acts faithfully on β and is either infinite cyclic or infinite dihedral. It is thus not hard to describe $\beta \cap T_\epsilon(G)$ explicitly. In fact $\beta \cap T_\epsilon(G)$ may be empty, equal to β, or consist of a countable disjoint union of closed intervals or of points. In each case $(\beta \cap T_\epsilon(G))/G$ is connected. It follows that $T_\epsilon(G)$ is connected.

Now, let $\Gamma \subseteq \mathrm{Isom}\,\mathbf{H^n}$ be any discrete group. Suppose that $\epsilon < \epsilon(n)$ where $\epsilon(n)$ is the Margulis constant (Section 2.2(iii)). Suppose that $0 < \delta \leq (\epsilon(n) - \epsilon)/2$. By the Margulis Lemma (2.2.8), we know that for any $x \in \mathbf{H^n}$, the group $\Gamma_\epsilon(x)$ is virtually abelian, and hence, by Lemma 3.3.2, is finite, parabolic or loxodromic. We have observed (Lemmas 3.1.3 and 3.3.1), that any parabolic or loxodromic subgroup of Γ lies in a maximal such. We see that if $x \in T_\epsilon(\Gamma)$, then $x \in T_\epsilon(G)$ for some maximal parabolic or loxodromic subgroup G of Γ. Note that $\Gamma_{\epsilon(n)}(x)$ is also virtually abelian, and so either parabolic or loxodromic. Now $G \cap \Gamma_{\epsilon(n)}(x)$ contains $\Gamma_{\epsilon(n)}(x)$ and is thus infinite. It follows that $\Gamma_{\epsilon(n)}(x) \subseteq G$. If $y \in \mathbf{H^n}$ with $d(x,y) \leq \delta$, then $\Gamma_\epsilon(y) \subseteq \Gamma_{\epsilon(n)}(x) \subseteq G$. Thus, if $y \in T_\epsilon(\Gamma)$, then $y \in T_\epsilon(G)$. We have shown:

Proposition 3.3.3 : *Suppose $\Gamma \subseteq \mathrm{Isom}\,\mathbf{H^n}$ is discrete, and $\epsilon < \epsilon(n)$. Then $T_\epsilon(\Gamma)$ is a disjoint union of the sets $T_\epsilon(G)$, as G ranges over all maximal parabolic and maximal loxodromic subgroups of Γ. Moreover, if G and G' are two distinct such subgroups, then $d(T_\epsilon(G), T_\epsilon(G')) \geq (\epsilon(n)-\epsilon)/2$. (Possibly one or both of $T_\epsilon(G)$ and $T_\epsilon(G')$ may be empty.)*

We write

$$\mathrm{thin}_\epsilon(M) = T_\epsilon(\Gamma)/\Gamma \subseteq M.$$

We call $\mathrm{thin}_\epsilon(M)$ the *thin part* of M. Thus, $\mathrm{thin}_\epsilon(M)$ is, topologically, a disjoint union of its connected components. Each such component has the form $T_\epsilon(G)/G$, where $G \subseteq \Gamma$ is either maximal parabolic or maximal loxodromic. We call these components *Margulis cusps* and *Margulis tubes* respectively. We write $\mathrm{cusp}_\epsilon(M)$ for the union of all the Margulis cusps, and call $\mathrm{cusp}_\epsilon(M)$ the *cuspidal part* of M. We write $\mathrm{thick}_\epsilon(M)$ and $\mathrm{noncusp}_\epsilon(M)$ for the closures in M of, respectively, $M \setminus \mathrm{thin}_\epsilon(M)$ and $M \setminus \mathrm{cusp}_\epsilon(M)$. We call these sets the *thick part* and the *noncuspidal part* of M.

From the description given of $T_\epsilon(G)$, it is not hard to see that the thick and thin parts of M meet precisely in their topological boundaries in M. Note that if Γ is torsion-free, so that M is a manifold, then

$$\mathrm{thin}_\epsilon(M) = \{x \in M \mid \mathrm{inj}(x, M) \leq \epsilon/2\},$$

where $\mathrm{inj}(x, M)$ is the injectivity radius of M at the point x. In this case, the definition has become standard.

We remark that there is an alternative one might attempt to define the thin part of an orbifold. Suppose $\Gamma \subseteq \mathrm{Isom}\,\mathbf{H^n}$ is discrete, and $\epsilon > 0$. We set $T'_\epsilon(\Gamma)$ to be the set of all those $x \in \mathbf{H^n}$ for which there is some infinite-order $\gamma \in \Gamma$ with $d(x, \gamma x) \leq \epsilon$. This gives us an alternative thick-thin decomposition with $\mathrm{thin}'_\epsilon(M) = T'_\epsilon(\Gamma)/\Gamma$, and $\mathrm{thick}'_\epsilon(M)$ being the closure of $M \setminus \mathrm{thin}'_\epsilon(M)$. Qualitatively, $T'_\epsilon(\Gamma)$ behaves like $T_\epsilon(\Gamma)$, except that $T'_\epsilon(G)$ is always connected if G is loxodromic. Note that clearly $T'_\epsilon(\Gamma) \subseteq T_\epsilon(\Gamma)$. We also have that if $\epsilon \leq \epsilon(n)$, then $T_{\epsilon/N}(\Gamma) \subseteq T'_\epsilon(\Gamma)$, where the constant N depends only on the dimension n. We leave this last statement as an exercise. We have no explicit use for this alternative thick-thin decomposition in this paper.

3.4. The convex core and definitions GF4 and GF5.

Suppose $\Gamma \subseteq \mathrm{Isom}\,\mathbf{H^n}$ is discrete. In Chapter 1, we defined the convex hull, $\mathrm{hull}\,\Lambda$ of the limit set Λ of Γ. Clearly, $\mathrm{hull}\,\Lambda$ is Γ-invariant, so we may define

$$\mathrm{core}(M) = (\mathbf{H^n} \cap \mathrm{hull}\,\mathbf{\Lambda})/\mathbf{\Gamma} \subseteq \mathbf{M}.$$

We call $\mathrm{core}(M)$ the *convex core* of M (Figure 3e).

Definition (GF4) : Suppose Γ is a discrete subgroup of $\mathrm{Isom}\,\mathbf{H^n}$. We say that Γ is "GF4" if $\mathrm{core}(M) \cap \mathrm{thick}_\epsilon(M)$ is compact for some $\epsilon \in (0, \epsilon(n))$, where $\epsilon(n)$ is the Margulis constant.

We could equally well say "for all $\epsilon \in (0, \epsilon(n))$" or replace $\mathrm{thick}_\epsilon(M)$ by $\mathrm{noncusp}_\epsilon(M)$. It should be apparent from the proofs of equivalence that these variations give rise to the same notion of geometrical finiteness. We could also use the alternative version of the thick-thin decomposition as described above.

We remark that the thick part of the convex core, $\mathrm{core}(M) \cap \mathrm{thick}_\epsilon(M)$ is defined intrinsically to $\mathrm{core}(M)$.

Definition GF5 : Suppose Γ is a discrete subgroup of $\mathrm{Isom}\,\mathbf{H^n}$. We say that Γ is "GF5" if there is a bound on the orders of finite subgroups of Γ, and if for some $\eta > 0$, $N_\eta(\mathrm{core}(M))$ has finite volume.

We could equally well say "for all $\eta > 0$".

I suspect that the bound on the orders of finite subgroups of Γ is superfluous, i.e. that it is implied by the statement that $N_\eta \mathrm{core}(M)$ has finite volume. We shall show that this is indeed the case if:

(i) M, itself, has finite volume, or if

(ii) the dimension $n \leq 3$.

Note that, by the Selberg Lemma (Chapter 1), there is always a bound on the orders of finite subgroups if:

(iii) Γ is finitely generated.

Case (i) therefore tells us, in particular, that a finite volume hyperbolic orbifold is topologically finite.

Note that any finite subgroup of $\mathrm{Isom}\,\mathbf{H^n}$ is conjugate to a subgroup of $\mathrm{Isom}\,\mathrm{S}^{n-1}$, and thus, by Lemma 2.2.3, contains an abelian subgroup of bounded finite index. We therefore see that bounding the orders of finite subgroups is the same as bounding the orders of elliptic (torsion) elements.

3.5. Polyhedral complexes and definition GF3.

Definition GF3 is a generalisation of the existence of a finite sided fundamental domain. One way to formulate this definition is in terms of what we shall call "polyhedral complexes". An example of a polyhedral complex is a locally finite CW-complex, equipped with a complete, locally compact path-metric such that each open r-cell of the CW-complex is isometric to an open r-dimensional polyhedron. However, a polyhedral complex is more general than this, in that we allow for unbounded polyhedral cells. In the complex, the closure of such a cell is the metric completion of such a polyhedron, or perhaps the quotient of this completion after performing certain identifications around the boundary. At any rate, this closure need not be compact. Thus, topologically, each open cell is attached only along some subset of the compactifying sphere.

To say that a group Γ is GF3 means that the quotient $\mathbf{H^n}/\mathbf{\Gamma}$ may be represented as a finite polyhedral complex. The idea is thus intuitively simple, though it will take some work to get the subject properly off the ground.

We can speak equally well of euclidean as well as hyperbolic complexes, and we shall begin with reference to euclidean space, since it is easier to describe examples in this case. We start with some general remarks about convex sets.

Given any subset $A \subseteq \mathbf{E^n}$, we write $\langle A \rangle$ for the affine span of A, i.e. the smallest subspace which contains A.

It is a simple observation that if convex subset of euclidean space has empty interior, then it must lie inside some lower-dimensional subspace. As a consequence, if $A \subseteq \mathbf{E^n}$ is convex, then A must have non-empty interior in $\langle A \rangle$.

By an *open interval* in the real line, $\mathbf{E^1}$, we mean a connected open subset (possibly empty). We note that a subset $A \in \mathbf{E^n}$ is both open and convex in $\mathbf{E^n}$ if and only if it meets each 1-dimensional subspace in an open interval. More generally, we give the following definition:

Definition : A subset $A \subseteq \mathbf{E^n}$ is *convex-open* if it is convex, and meets each 1-dimensional subspace either in an open interval, or in a single point.

Thus, a convex-open set is certainly convex, but not necessarily open. For example, any subspace of $\mathbf{E^n}$ is convex-open. Note that the intersection of two convex-open sets is convex-open. Suppose that μ is a subspace of $\mathbf{E^n}$, and that $A \subseteq \mathbf{E^n}$. If A is a convex-open set, then $A \cap \mu$ is convex-open, both as a subset of $\mathbf{E^n}$, and as a subset of μ. Conversely, if $A \subseteq \mu$, and A is convex-open in μ, then A is convex-open in $\mathbf{E^n}$. Also, the orthogonal projection of a convex-open set onto a subspace has convex-open image.

It is not hard to see that a convex-open set is open if and only if it has non-empty interior. We conclude that a convex set $A \subseteq \mathbf{E^n}$ is convex-open if and only if it is open in $\langle A \rangle$.

We may give the following intrinsic definition of a convex-open set. Given any $x, y \in \mathbf{E^n}$, write $[x, y]$ for the geodesic segment joining x and y. Write $\mathrm{join}(x, y) = [x, y] \setminus \{x, y\}$ if $x \neq y$, and $\mathrm{join}(x, x) = \{x\}$. Then, $A \subseteq \mathbf{E^n}$ is convex-open if and only if for all $x, y \in A$, there exist $x', y' \in A$ such that $x, y \in \mathrm{join}(x', y')$.

More generally, given subsets $A, B \subseteq \mathbf{E^n}$, we define the *join* of A and B as

$$\mathrm{join}(A,B) = \bigcup_{\substack{x \in A \\ y \in B}} \mathrm{join}(x,y).$$

Clearly, if A is non-empty, then B lies in the closure of $\mathrm{join}(A,B)$.

Lemma 3.5.1 : *If $A, B \subseteq \mathbf{E^n}$ are convex-open, then $\mathrm{join}(A,B)$ is a convex-open.*

Proof : Using the intrinsic definition, we are reduced to considering the join of two open intervals or points in $\mathbf{E^3}$. Generically, we have two skew intervals whose join is the interior of the tetrahedral convex hull. There are a few degenerate cases, each of which is intuitively clear. We leave the formal verification to the reader. $\diamondsuit$

Lemma 3.5.2 : *Suppose $A \subseteq \mathbf{E^n}$ is convex-open, and that $x \in \bar{A}$. Then, $\mathrm{join}(A,x) = A$.*

Proof : This is easily verified if A is open. However, this case suffices, since A is always open in $\langle A \rangle$. $\diamondsuit$

As an obvious corollary, we have:

Lemma 3.5.3 : *If A, B are convex-open, and $B \cap \bar{A} \neq \emptyset$, then $A \subseteq \mathrm{join}(A,B)$.*

Corollary 3.5.4 : *If A, B are convex-open and $A \cap B \neq \emptyset$, then $\mathrm{join}(A,B) = \mathrm{hull}(A \cup B)$.*

Proof : By Lemma 3.5.3, we have $A \cup B \subseteq \mathrm{join}(A,B)$. $\diamondsuit$

Lemma 3.5.5 : *If A, B are convex-open, and $A \cap B \neq \emptyset$, then $\bar{A} \cap \bar{B} = \overline{A \cap B}$.*

Proof : Clearly, $\overline{A \cap B} \subseteq \bar{A} \cap \bar{B}$. Suppose $x \in \bar{A} \cap \bar{B}$. Choose some $y \in A \cap B$. We have $\mathrm{join}(x,y) \subseteq \mathrm{join}(A,x) \cap \mathrm{join}(B,x) = A \cap B$, by Lemma 3.5.2. Thus $x \in \overline{A \cap B}$. $\diamondsuit$

Given a subset $S \subseteq \mathbf{E^n}$, and a collection $\mathcal{A}$ of subsets of S, we say that $\mathcal{A}$ is *convex-open* decomposition of S if the following are satisfied.

(1) Each element of $\mathcal{A}$ is a convex-open set.

(2) The elements of $\mathcal{A}$ are all disjoint.

(3) $S = \bigcup \mathcal{A}$.

(4) If $A, B \in \mathcal{A}$, and $B \cap \bar{A} \neq \emptyset$, then $B \subseteq \bar{A}$.

Given two convex-open decompositions, $\mathcal{A}$ and $\mathcal{B}$, of S, we say that $\mathcal{B}$ is a *subdivision* of $\mathcal{A}$ if each element of $\mathcal{B}$ is subset of some element of $\mathcal{A}$. We write $\mathcal{B} \leq \mathcal{A}$. Thus, the set of all convex-open decompositions of S is partially ordered by subdivision.

Given two convex-open decompositions, $\mathcal{A}_1$ and $\mathcal{A}_2$, of a set S, there is a unique maximal common subdivision, namely

$$\langle \mathcal{A}_1, \mathcal{A}_2 \rangle = \{ A_1 \cap A_2 \mid A_1 \in \mathcal{A}_1, A_2 \in \mathcal{A}_2 \}.$$

Property (4) may be verified using Lemma 3.5.5. By "maximal", we mean that if $\mathcal{B} \leq \mathcal{A}_1$ and $\mathcal{B} \leq \mathcal{A}_2$, then $\mathcal{B} \leq \langle \mathcal{A}_1, \mathcal{A}_2 \rangle$.

Proposition 3.5.6 : *Suppose $S \subseteq \mathbf{E^n}$ is convex, then S admits a maximal convex cell decomposition $\mathcal{S} = \mathcal{S}(S)$ (i.e. if $\mathcal{A}$ is any convex cell decomposition of S, then $\mathcal{A} \leq \mathcal{S}$).*

Proof : We take $\mathcal{S}$ to be the set of all maximal convex-open subsets of S, i.e. if $A \in \mathcal{S}$ and B is a convex-open set with $A \subseteq B \subseteq S$, then $A = B$.

Suppose $x \in S$. Let $\mathcal{C}(x)$ be the set of all convex-open subsets of S containing the point x. $\mathcal{C}(x)$ is partially ordered by set inclusion. Note that $\mathcal{C}(x)$ is non-empty, since $\{x\} \in \mathcal{C}(x)$. By Lemma 3.5.1 and Corollary 3.5.4, we see that if $A, B \in \mathcal{C}(x)$, then $A \cup B \subseteq \mathrm{join}(A, B) \in \mathcal{C}(x)$. Also, any increasing union of convex-open sets must be convex-open. We conclude that $\mathcal{C}(x)$ must have a maximal element, namely $A(x) = \bigcup \mathcal{C}(x) \in \mathcal{S}$. We see that the elements of $\mathcal{S}$ cover S. Clearly, if $A \in \mathcal{S}$ and $x \in A$, then $A = A(x)$. Thus the elements of $\mathcal{S}$ are disjoint. Finally, Property (4) may be verified using Lemma 3.5.3 as follows. If $A, B \in \mathcal{S}$ with $B \cap \bar{A} \neq \emptyset$, then $A \subseteq \mathrm{join}(A, B) \subseteq S$. Thus, $A = \mathrm{join}(A, B)$ and so $B \subseteq \bar{A}$. We see that $\mathcal{S}$ is a convex-open decomposition of S.

Note that if $B \subseteq S$ is a convex-open subset, then $B \subseteq A(x) \in \mathcal{S}$ for any $x \in B$. It follows that $\mathcal{S}$ is maximal with respect to subdivision. $\diamondsuit$

Given a convex subset $S \subseteq \mathbf{E^n}$, we write $\mathrm{ri}(S)$ for the interior of S in $\langle S \rangle$. We call $\mathrm{ri}(S)$ is *relative interior* of S. Thus, $\mathrm{ri}(S)$ is convex-open, and dense in S. We write $\mathrm{rb}(S) = \bar{S} \setminus \mathrm{ri}(S)$ for the *relative boundary* of S. By the *dimension*, dim(S), of S, we mean the dimension of $\langle S \rangle$. It is easily seen that $\mathrm{ri}(S) \in \mathcal{S}(S)$. In fact, $\mathrm{ri}(S)$ is the unique element of $\mathcal{S}(S)$ of maximal dimension.

Definition : A closed convex subset $P \subseteq \mathbf{E^n}$ is a *polyhedron* if $\mathcal{S}(P)$ is locally finite.

In this case, $\mathcal{S}(P)$ is an example of a "convex cell complex". Suppose F is a closed subset of $\mathbf{E^n}$, then:

Definition : A *convex cell complex*, $\mathcal{A}$, representing F is a locally finite convex-open decomposition of F.

In other words, $\mathcal{A}$ satisfies Properties (1)–(4) above, together with:

(5) $\mathcal{A}$ is locally finite.

We refer to the elements of a convex cell complex as *cells*.

Note that it is an immediate consequence of Properties (4) and (5) that every cell of a convex cell complex meets the closures of only finitely many other cells.

Suppose $P \subseteq \mathbf{E^n}$ is closed and convex, and admits a representation as a convex cell complex, $\mathcal{A}$. Then $\mathcal{A}$ is a subdivision of $\mathcal{S}(P)$ and so $\mathcal{S}(P)$ is locally finite. Thus:

Lemma 3.5.7 : *A closed set $P \subseteq \mathbf{E^n}$ is a polyhedron if and only if it is convex, and admits a representation as a convex cell complex.*

Later, we shall relate this to a more obvious description of polyhedra.

Suppose that $\mathcal{A}$ is any convex cell complex in $\mathbf{E^n}$ (representing the closed set $\bigcup \mathcal{A}$). Suppose $A \in \mathcal{A}$. We say that $B \in \mathcal{A}$ is a *face* of A if $B \subseteq \bar{A}$. We write $\mathcal{F}_{\mathcal{A}}(A)$, or just $\mathcal{F}(A)$, for the set of faces of A. Property (4) tells us that $\mathcal{F}(A)$ is itself a convex cell complex, representing $\bar{A}$. We conclude that $\bar{A}$ is a polyhedron. Note that $\mathcal{F}(A) \leq \mathcal{S}(\bar{A})$.

If P is a polyhedron, we shall refer to the elements of $\mathcal{S}(P)$ as *sides* of P. Thus the notion of a side is intrinsic to P, and is distinct from the notion of a face in the case where P is the closure of a cell of a convex cell complex. We say that P is *finite-sided* if $\mathcal{S}(P)$ is finite.

Definition : If $\mathcal{A}$ is convex cell complex, and for each $A \in \mathcal{A}$ we have $\mathcal{F}(A) = \mathcal{S}(\bar{A})$, then we say that $\mathcal{A}$ is *strictly convex*.

In other words $\mathcal{A}$ is strictly convex if faces and sides coincide.

The following discussion will give rise to several ways of constructing convex cell complexes.

Suppose $F \subseteq \mathbf{E^n}$ is closed. Suppose we have some set $\mathcal{J}$, and a means of assigning to each point of $x \in F$ a finite subset $J(x) \subseteq \mathcal{J}$. Given any finite subset $J \subseteq \mathcal{J}$, we write $A(J) = \{x \in F \mid J(x) = J\}$ and $C(J) = \{x \in F \mid J(x) \supseteq J\}$. We make the following assumptions about the map $[J \mapsto J(x)]$. Firstly, we assume that $[x \mapsto J(x)]$ is "upper-semicontinuous", by which we mean that, given any $x \in F$, there is some neighbourhood, V, about x, such that for all $y \in V \cap F$, we have $J(y) \subseteq J(x)$. We also suppose that $C(\{j\})$ is convex for each $j \in \mathcal{J}$. Thus, $C(J)$ is closed and convex for each finite $J \subseteq \mathcal{J}$. Finally, we suppose that $A(J)$ is convex-open for each finite $J \subseteq \mathcal{J}$. Now let $\mathcal{A} = \{A(J) \mid J \subseteq \mathcal{J} \text{ is finite}\}$. We claim that $\mathcal{A}$ is a convex cell complex representing F. Properties (1), (2), (3) and (5) are immediate. We leave the verification of (4) as an exercise. In fact, in the cases which interest us, we will have that each $C(J)$ is the closure of $A(J)$, provided $A(J) \neq \emptyset$, and so property (4) becomes trivial.

For an example of this construction, suppose that $X \subseteq \mathbf{E^n}$ is a discrete set of points. Set $\mathcal{J} = X$. Given any $x \in \mathbf{E^n}$, set

$$J(x) = \{a \in X \mid d(a,x) \leq d(b,x) \text{ for all } b \in X\}.$$

In other words, $J(x)$ is the set of all nearest points to x in X. It is not hard to see that $[x \mapsto J(x)]$ satisfies the hypotheses given above. Upper-semicontinuity is more or less immediate. To see that $A(J)$ is convex-open for each finite $J \subseteq X$, suppose that $x \neq y \in \mathbf{E^n}$ with $J(x) = J(y) = J$. Then all the points of J lie in a codimension-1 subspace orthogonal to the line β through x and y. If $a \in J$ and $b \in X \setminus J$, then $d(x,a) < d(x,b)$ and $d(y,a) < d(y,b)$. Thus if $z \in [x,y]$, simple plane trigonometry shows that $d(z,a) < d(z,b)$. Thus $J(z) \subseteq J$. By upper-semicontinuity, we see that $J(w) \subseteq J$ for all w in some neighbourhood V of $[x,y]$. Since each point on β is equidistant from each of the points of J, we see that $J(w) = J$ for all $w \in \beta \cap V$, and so $\beta \cap V \subseteq A(J)$. The remaining property of $[x \mapsto J(x)]$ is similar. We write $\mathcal{D}(X)$ for the convex cell complex representing $\mathbf{E^n}$ arising in this way.

Definition : We call $\mathcal{D}(X)$ the *Voronoi tesselation* of $\mathbf{E^n}$ corresponding to X.

Note that any Voronoi tesselation is strictly convex in the sense described above.

Suppose that $\mathcal{A}$ is any representation of $\mathbf{E^n}$ as a convex cell complex. Let $\mathcal{A}^n$ be the set of top-dimensional cell (i.e. those cells that are open in $\mathbf{E^n}$). It is easy to see that $\bigcup \mathcal{A}^n$ is dense in $\mathbf{E^n}$. (For example note that lower dimensional cells have zero Lebesgue measure.) This observation gives rise to the following characterisation of the set of top-dimensional cells of a convex cell complex representing $\mathbf{E^n}$.

Lemma 3.5.8 : *Suppose that $\mathcal{U}$ is disjoint, locally finite collection of open convex subsets of $\mathbf{E^n}$ whose closures cover $\mathbf{E^n}$. Then there is a convex cell complex $\mathcal{A} = \mathcal{A}(\mathcal{U})$ representing $\mathbf{E^n}$ such that $\mathcal{U}$ is the collection of top-dimensional cells of $\mathbf{E^n}$.*

Proof : Given such a collection $\mathcal{U}$ and a point $x \in \mathbf{E^n}$, let $J(x) = \{U \in \mathcal{U} \mid x \in \bar{U}\}$. Let $\mathcal{A}(\mathcal{U}) = \{A(J) \mid J \subseteq \mathcal{U} \text{ is finite}\}$, where $A(J) = \{y \in \mathbf{E^n} \mid \mathbf{J(y) = J}\}$. We leave as an exercise the statement that $[x \mapsto J(x)]$ satisfies the properties given above. It follows that $\mathcal{A}(\mathcal{U})$ is a convex cell complex. Note that if $U \in \mathcal{U}$ then $A(\{U\}) = U$. Thus $\mathcal{U} \subseteq \mathcal{A}(\mathcal{U})$. Now any open set must meet some element of $\mathcal{U}$. It follows that $\mathcal{U}$ is precisely the set of top-dimensional cells of $\mathcal{A}(\mathcal{U})$. $\diamondsuit$

Note that $\mathcal{A}(\mathcal{U})$ is maximal with respect to subdivision, i.e. if $\mathcal{B}$ is a convex cell complex with $\mathcal{B}^n = \mathcal{U}$, then $\mathcal{B} \subseteq \mathcal{A}(\mathcal{U})$.

We can give now a couple of examples of convex cell complexes by describing their top-dimensional cells. For example, consider a tiling of $\mathbf{E^2}$ in the pattern of a "brick wall" (Figure 3f). In this picture, each 2-dimensional cell has six 0-dimensional faces, six 1-dimensional faces and one 2-dimensional face (itself). However each such cell is intrinsically a rectangle, and so has four 0-dimensional sides, four 1-dimensional sides and one 2-dimensional side.

As another example, we consider a tesselation of $\mathbf{E^3}$ by "planks", i.e. bi-infinite square prisms (Figure 3g). The tesselation is made up of successive layers. Each layer is one plank thick, and consists of an infinite number of of planks laid parallelly. In one layer, the planks are laid north-south, in the next they are laid east-west, and so on alternately. In this example, each top-dimensional cell has infinitely many faces, but only finitely many sides.

We remark that a Voronoi tesselation may be defined in terms of its top-dimensional cells. Suppose $X \subseteq \mathbf{E^n}$ is discrete. Given $a \in X$, write

$$U(a) = \{x \in \mathbf{E^n} \mid \mathbf{d(x,a) < d(x,b)} \text{ for all } \mathbf{b \in X \setminus \{a\}}\}.$$

In other words, $U(a)$ is the set of points nearer to a than to any other point of X. One can check that $\mathcal{U}(X) = \{U(a) \mid a \in X\}$ satisfies the hypotheses of Lemma 3.5.8, and that $\mathcal{D}(X) = \mathcal{A}(\mathcal{U}(X))$.

We now give another description polyhedra as promised above. We can restrict attention to polyhedra having non-empty interior. By a *half-space* in $\mathbf{E^n}$, we mean closed subset of $\mathbf{E^n}$ of non-empty interior, whose boundary is a codimension-1 subspace.

Lemma 3.5.9 : *Suppose $P \subseteq \mathbf{E^n}$ is a closed set of non-empty interior. Then P is a polyhedron if and only if it may be written in the form $P = \bigcap_{i \in I} H_i$, where I is an indexing set, where each H_i is as a half-space, and where the boundaries ∂H_i are locally finite on P.*

Proof : Suppose P has the form $\bigcap_{i \in I} H_i$ as described above. (We can suppose that $H_i \neq H_j$ if $i \neq j$.) It is easily seen that $\mathrm{rb}(P) = \partial P = \bigcup_{i \in I}(P \cap \partial H_i)$. Given $x \in P$, let $J(x) = \{i \in I \mid x \in P \cap H_i\}$. Thus each $J(x) \subseteq I$ is finite. It is an exercise that the map $[x \mapsto J(x)]$ satisfies the hypotheses stated above. Thus, we get a representation of ∂P as a convex cell complex. It follows that $\{\mathrm{ri}(P)\} \cup \mathcal{A}$ gives a representation of P as a convex cell complex. (In fact this gives the set of sides, $\mathcal{S}(P)$, of P.)

Conversely, suppose that the set of sides, $\mathcal{S}(P)$, of P is locally finite. To each codimension-1 side, $A \in \mathcal{S}(P)$, we may associate a half-space H such that $\partial H = \langle A \rangle$ (and so $P \cap \partial H = \bar{A}$), and such that $\mathrm{ri}(P) \cap H \neq \emptyset$. We leave as an exercise that P is the intersection of all such half-spaces. $\diamondsuit$

Note that it follows from the argument that P is finite-sided if and only if it is a finite intersection of half-spaces. It also follows that P is finite-sided if and only if it has finitely many codimension-1 sides.

We intend to relate the the notion of a convex cell complex to that of a CW-complex. Suppose $\mathcal{A}$ is a convex cell complex representing some closed subset of $\mathbf{E^n}$. We write $\mathcal{A}^i$ for the set of all *i-cells* of $\mathcal{A}$, i.e. all cells of dimension i. We write $K^i = K^i(\mathcal{A}) = \bigcup \bigcup_{j \leq i} \mathcal{A}^j$. We call K^i the *i-skeleton* of $\mathcal{A}$. Suppose $A \in \mathcal{A}^i$, and $B \neq A$ is a face of A. It is clear that B must have empty interior in $\langle A \rangle$, and so must have dimension less than that of A. Since $\mathrm{rb}(A)$ is the union of all faces of A, other than A itself, it follows that $\mathrm{rb}(A) \subseteq K^{i-1}$. Note also that it follows that the i-skeleton is closed in $\mathbf{E^n}$.

In fact, we claim that, in the definition of a convex cell complex, we could replace Property (4) by the following:

(4′) If $A \in \mathcal{A}^r$, then $\mathrm{rb}(A) \subseteq K^{r-1}$.

We have already seen that, given axioms (1), (2), (3) and (5), then (4) implies (4′). The converse is a little more complicated. Suppose that $\mathcal{A}$ satisfies (1), (2), (3), (4′) and (5), and that $A \in \mathcal{A}^r$. Using (for example) a measure-theoretic argument, we see that $(\bigcup \mathcal{A}^{r-1}) \cap \mathrm{rb\,A}$ is a dense subset of ∂A. Suppose that $B \in \mathcal{A}$ intersects $\mathrm{rb}(A)$. By axiom (4′), B has dimension at most $r-1$. Suppose, for the moment, that $B \in \mathcal{A}^{r-1}$. If B is not a subset of $\mathrm{rb\,A}$, then there is some point x in the relative boundary of $\bar{A} \cap B$ in B. By considering a neighbourhood of x in $\mathrm{rb\,A}$, we see that $x \in \mathrm{rb\,C}$ for some $C \in \mathcal{A}^{r-1}$, different from B. But by (4′), we have that $\mathrm{rb\,C} \subseteq \mathrm{K^{r-2}}$. This contradiction tells us that $B \subseteq \mathrm{rb\,A}$. In other words, $\mathrm{rb\,A}$ is a union of closures of $(r-1)$-cells. We have deduced property (4) in the case where $\mathrm{dim\,A} - \mathrm{dim\,B} = 1$. We now use induction over $\mathrm{dim\,A} - \mathrm{dim\,B}$. Suppose then, that $B \in \mathcal{A}$ is an i-cell intersecting $\mathrm{rb\,A}$, with $\dim B < r-1$. Now, B intersects

$\mathrm{rb}\,\mathrm{D}$ for some $(r-1)$-cell $D \subseteq \mathrm{rb}\,\mathrm{A}$. By the induction hypothesis, $B \subseteq \mathrm{rb}\,\mathrm{D}$. We see that $B \subseteq \mathrm{rb}\,\mathrm{A}$. Thus we have shown the equivalence of the two descriptions of convex cell complexes.

We shall want the following generalisation of Lemma 3.5.8.

Lemma 3.5.10 : *Suppose $\mathcal{A}$ is a convex cell complex representing $\mathbf{E^n}$. Suppose that $\mathcal{U}$ is a locally finite collection of r-dimensional convex-open sets, whose closures cover the r-skeleton of $\mathcal{A}$. Suppose that each $U \in \mathcal{U}$ is a subset of some $A \in \mathcal{A}^r$. Then there is a natural subdivision, $\mathcal{B}$ of $\mathcal{A}$ such that $\mathcal{B}^r = \mathcal{U}$, and $\mathcal{B}^i = \mathcal{A}^i$ for all $i \in \{r+1, \ldots, n\}$.*

Thus, Lemma 2.5.8 is the case where $\mathcal{A} = \{\mathbf{E^n}\}$.

Proof : Given a point x in the the r-skeleton, K^r, let $J(x) = \{U \in \mathcal{U} \mid x \in \bar{U}\}$. This gives rise to a representation of K^r as a convex cell complex, $\mathcal{B}_0$, with $\mathcal{B}_0^r = \mathcal{U}$ (c.f. Lemma 3.5.8). Thus, $\mathcal{B}_1 = \mathcal{B}_0 \cup \bigcup_{i=r+1}^n \mathcal{A}^i$ is a convex cell complex representing $\mathbf{E^n}$. Let $\mathcal{B}$ be the common subdivision $\mathcal{B} = \langle \mathcal{A}, \mathcal{B}_1 \rangle$. $\diamondsuit$

Note that the convex cell complex, $\mathcal{B}$ is maximal with respect to subdivision.

The entire discussion we have given so far in this section (except the examples!) works equally well with euclidean space, $\mathbf{E^n}$, replaced by hyperbolic space, $\mathbf{H^n}$. (For example, consider the Klein model for hyperbolic space, where the notions of euclidean and hyperbolic convexity coincide.)

We want to consider convex cell complexes invariant under a discrete group action. It will be better to do this with reference to hyperbolic space. Our definition GF3 will be in terms of such equivariant cell complexes. To give the definitions more intuitive meaning, we shall relate this to "polyhedral complexes" representing the quotient orbifold. The discussion of polyhedral complexes is not logically essential for the rest of the paper.

Suppose that Γ is a discrete subgroup of $\mathrm{Isom}\,\mathbf{H^n}$. We say that a convex cell complex representing $\mathbf{H^n}$ is *associated* to Γ if $\mathcal{A}$ is Γ-invariant, and if $\mathrm{stab}_\Gamma A$ is finite for each $A \in \mathcal{A}$.

For an example, choose any $a \in X$. The Voronoi tesselation, $\mathcal{D}(\Gamma a)$, is associated to Γ. We call $\mathcal{D}(\Gamma a)$ a *Dirichlet tesselation* of $\mathbf{H^n}$ for the group Γ. (It is often assumed elsewhere that $\mathrm{stab}_\Gamma\, a$ is trivial, though we shall have no need for this hypothesis.) More generally, if $X \subseteq \mathbf{H^n}$ is discrete and Γ-invariant, and if the set of orbits, X/Γ is finite, we shall call $\mathcal{D}(X)$ a *generalised Dirichlet tesselation*.

Lemma 3.5.11 : *Suppose $X \subseteq \mathbf{H^n}$ is Γ-invariant, with X/Γ finite. Then the generalised Dirichlet tesselation, $\mathcal{D}(X)$ is locally finite on $\mathbf{H^n} \cup \mathbf{\Omega(\Gamma)}$.*

Proof : If $x \in \Omega(\Gamma)$, then there is a horoball B about x with $X \cap B = X \cap \partial B$ consisting of a non-empty finite set of points. If $y \in \mathbf{H^n}$ is sufficiently close to x, then the set of nearest points to y in x is a subset of $X \cap B$. $\diamondsuit$

Given $a \in \mathbf{H^n}$, write

$$D_\Gamma(a) = \{x \in \mathbf{H^n} \mid \mathbf{d(x, a)} \leq \mathbf{d(x, \Gamma a)}\}.$$

Definition : We call $D_\Gamma(a)$ the *Dirichlet domain* for Γ about a.

From the discussion of Voronoi tesselations, we see that $D_\Gamma(a)$ is the closure, in $\mathbf{H^n}$, of the cell of the Dirichlet tesselation $\mathcal{D}(\Gamma a)$ containing a. (This cell is necessarily top-dimensional). We see that $D_\Gamma(a)$ is a polyhedron. Clearly, the set of images $\Gamma D_\Gamma(a)$ of $D_\Gamma(a)$ is locally finite, and covers $\mathbf{H^n}$. By Lemma 3.2.1, (or Lemma 3.5.11) we see that in fact $\Gamma D_\Gamma(a)$ is locally finite on $\mathbf{H^n} \cup \mathbf{\Omega(\Gamma)}$. Let D be the closure, in $\mathbf{H^n_C}$, of $D_\Gamma(a)$. It follows that $\bigcup \Gamma D \setminus \Lambda(\Gamma)$ is closed in $\mathbf{H^n} \cup \mathbf{\Omega(\Gamma)}$. Thus $\mathbf{H^n} \cup \mathbf{\Omega(\Gamma)} \subseteq \bigcup \mathbf{\Gamma D}$.

We now want to describe polyhedral complexes. Suppose, for the moment, that Γ is torsion-free. Let $\pi : \mathbf{H^n} \longrightarrow \mathbf{M}$ be the projection to the quotient manifold M. Suppose that $\mathcal{A}$ is a convex cell complex associated to Γ. In this case, each cell $A \subseteq \mathcal{A}$ descends to a subset πA of M. Since $\pi|A$ is injective, πA is itself a convex open set in the in the induced path-metric. The collection $\mathcal{B} = \{\pi A \subseteq M \mid A \in \mathcal{A}\}$ has the structure of a "polyhedral complex" which we shall axiomatise as follows.

Definition : A *polyhedral complex*, $\mathcal{B}$, representing a complete locally compact metric space, M, is a partition of M into a locally-finite collection of subsets (*cells*) satisfying the following. For each $B \in \mathcal{B}$, there is a polyhedron $P \subseteq \mathbf{H^n}$, together with a representation of P as a convex-cell complex, $\mathcal{F}$, containing $\mathrm{ri}(P)$, and a map $f : P \longrightarrow M$ with $f(\mathrm{ri}(P)) = B$ and satisfying the following property. For each $C \in \mathcal{F}$ we have that $f(C) \in \mathcal{B}$, and that $f|C$ is an isometry from C onto $f(C)$ in the path-metric induced from M.

In particular, each element of $\mathcal{B}$ is isometric to a convex-open set in the path-metric induced from M.

Comparing with the usual definition of a CW-complex, we see that P is analogous to a closed cell, and that $f|\mathrm{rb}(P)$ is the attaching map.

Note that if $F \subseteq \mathbf{H^n}$ is closed, then a representation of F as a convex cell complex is the same as a representation as a polyhedral complex by this definition.

Suppose now, that Γ is any discrete subgroup of $\mathrm{Isom}\,\mathbf{H^n}$, and that $\mathcal{A}$ is a Γ-invariant convex cell complex representing $\mathbf{H^n}$. Given $A \in \mathcal{A}$, write $\mathrm{stab}^0_\Gamma(A)$ for the pointwise stabliser of A in Γ. Thus $\mathrm{stab}^0_\Gamma(A)$ is a normal subgroup of $\mathrm{stab}_\Gamma(A)$. The property we need for $\mathcal{A}$ to descend to a polyhedral complex representing the quotient orbifold $M = \mathbf{H^n/\Gamma}$ is that $\mathrm{stab}^0_\Gamma(A) = \mathrm{stab}_\Gamma(A)$ for all $A \in \mathcal{A}$. Note that this is stronger than the property of being associated to Γ.

Proposition 3.5.12 : *If $\mathcal{A}$ is a Γ-invariant convex cell complex representing $\mathbf{H^n}$, then there is a subdivision, $\mathcal{B}$ of $\mathcal{A}$, such that $\mathrm{stab}^0_\Gamma(B) = \mathrm{stab}_\Gamma(B)$ for all $B \in \mathcal{B}$. Moreover, if $\mathcal{A}$ is associated to Γ (i.e. $\mathrm{stab}_\Gamma(A)$ is finite for all $A \in \mathcal{A}$), then we can arrange that each cell of $\mathcal{A}$ contains only finitely many cells on $\mathcal{B}$.*

Proof : Suppose $\mathcal{A}$ is Γ-invariant. We begin by subdividing the top-dimensional cells. Choose $A \in \mathcal{A}$, and let $G = \mathrm{stab}_\Gamma(A)$ and $G_0 = \mathrm{stab}^0_\Gamma(A)$. Thus G/G_0 acts faithfully on A. We can find some point $a \in A$ such that $\mathrm{stab}_\Gamma(a) = G_0$. Let

$$U = \{x \in A \mid d(x,a) < d(x,b) \text{ for all } b \in Ga \setminus \{a\}\}.$$

Thus, the images of $\bar{U}$ under G cover $\bar{A}$. Let $\mathcal{U}(A) = \Gamma A$ be the set of images of U under Γ. We perform such a construction for some representative, A, of each orbit of $\mathcal{A}^n$ under Γ, and let $\mathcal{U}$ be the union of all the $\mathcal{U}(A)$. Note that if $U \in \mathcal{U}$, then $\mathrm{stab}_\Gamma(U) = \mathrm{stab}^0_\Gamma(U)$. Moreover, we see that $\mathcal{U}$ satisfies the hypotheses of Lemma 2.5.10 (with $r = n$). This gives us a Γ-invariant subdivision, $\mathcal{A}_1$ of $\mathcal{A}$, with $\mathcal{A}_1^n = \mathcal{U}$. We now get to work on the $(n-1)$-dimensional cell of $\mathcal{A}$, and, in a similar way, arrive at a subdivision, $\mathcal{A}_2$ of $\mathcal{A}_1$ with $\mathcal{A}_2^n = \mathcal{A}_1^n$ and with the property that for all $U \in \mathcal{A}_2^{n-1}$, we have $\mathrm{stab}_\Gamma(U) = \mathrm{stab}^0_\Gamma(U)$. We continue inductively, and after n steps set $\mathcal{B} = \mathcal{A}_n$.

If $\mathcal{A}$ is associated to Γ it is easy to see that, at each stage, each cell gets subdivided into only finitely many subsets. $\diamondsuit$

Given a Γ-invariant convex cell complex $\mathcal{A}$, we write $\mathcal{A}/\Gamma$ for the set of orbits of $\mathcal{A}$ under Γ.

We can now give our last definition of geometrical finiteness.

Definition (GF3) : Suppose Γ is a discrete subgroup of $\mathrm{Isom}\,\mathbf{H^n}$. We say that Γ is "GF3" if there is a Γ-invariant convex cell complex, $\mathcal{A}$, representing $\mathbf{H^n}$, with $\mathrm{stab}_\Gamma(A)$ finite for each $A \in \mathcal{A}$, and with the set of orbits, $\mathcal{A}/\Gamma$, finite.

By Lemma 3.5.12, we see that Γ is GF3 if and only if the quotient orbifold, $M = \mathbf{H^n/\Gamma}$, admits a representation as a finite polyhedral complex.

Definition GF3 is roughly equivalent to saying that Γ has a fundamental domain which is a finite union of polyhedra, each with a finite number of faces (Lemma 3.5.13). In this context, it has often been assumed elsewhere that the convex cell complex $\mathcal{A}$ should be strictly convex (c.f. the axiom of "side-pairing" [BeaM]). However, we shall have no need for this hypothesis.

We shall need the following observation.

Lemma 3.5.13 : *Suppose $\Gamma \subseteq \mathrm{Isom}\,\mathbf{H^n}$ is discrete. Suppose that $\mathcal{A}$ is convex cell complex associated to Γ, and that $\mathcal{A}/\Gamma$ is finite. Then, for each $A \in \mathcal{A}$, the set of faces, $\mathcal{F}(A)$ of $\mathcal{A}$ is finite.*

In particular, we see that each cell of $\mathcal{A}$ is the relative interior of a finite-sided polyhedron.

Proof : We have already remarked that, in a convex cell complex, each cell meets the closures of only a finite number of others. Since $\mathcal{A}/\Gamma$ is finite, these numbers are bounded by some constant k_1. Moreover, the stabliser of each cell is finite. Thus, there is some constant, k_2, such that $|\,\mathrm{stab}_\Gamma(A)| \leq k_2$ for all $A \in \mathcal{A}$. If $A, B \in \mathcal{A}$, then

$$|\mathcal{F}(A) \cap \Gamma B| \leq |\{\gamma \in \Gamma \mid \gamma B \subseteq \bar{A}\}| = |\{\gamma \in \Gamma \mid B \subseteq \gamma \bar{A}\}| \leq k_1 |\,\mathrm{stab}_\Gamma(A)| \leq k_1 k_2.$$

In other words, at most $k_1 k_2$ faces of A can lie in any given orbit under Γ. Since there are only finitely many such orbits, it follows that A has finitely many faces. $\diamondsuit$

We conclude this section by introducing a notion that will be used in the proof of GF3$\Rightarrow$GF1, as well as in Chapter 5.

In Section 2.1(ii), we defined $S(\mathbf{E^n})$ to be the $(n-1)$-sphere of parallel classes of geodesic rays in $\mathbf{E^n}$. Given a closed convex subset, $P \subseteq \mathbf{E^n}$, we write $\Theta(P) \subseteq S(\mathbf{E^n})$ to be the set of those parallel classes which contain some representative ray lying entirely in P. If λ is the spherical Lebesgue measure on $S(\mathbf{E^n})$, we may define $\theta(P) \in [0,1]$ by

$$\theta(P) = \lambda(\Theta(P))/\lambda(S(\mathbf{E^n})).$$

Thus, if $a \in P$, the quantity $\theta(P)$ measures the proportion of geodesic rays based at a, which lie in P. The following is a simple observation.

Lemma 3.5.14 : *If $\mathcal{P}$ is a collection of closed convex sets in $\mathbf{E^n}$ having disjoint interiors, then $\sum_{P\in\mathcal{P}} \theta(P) \leq 1$.*

4. Proofs of equivalence.

In this chapter, we prove the equivalence of the various definitions of geometrical finiteness we have given. We show the following implications:

We include GF1⇒GF2 and GF1⇒GF4 since they are much simpler than following the cycle. More detailed arguments for the equivalence of GF1, GF2, GF4 and GF5, in a more general context, are given in [Bow]. For this reason, we shall try to keep the proofs given here brief. We also include a proof that any finite volume hyperbolic orbifold (complete, without boundary) is geometrically finite.

We shall begin with a general discussion of bounded parabolic fixed points.

Suppose $\Gamma \subseteq \operatorname{Isom} \mathbf{H^n}$ is discrete, and p is a bounded parabolic fixed point. Let $G = \operatorname{stab}_\Gamma p$. Put $p = \infty$ in the upper-half space model, $\mathbf{R^n_+} \cup \partial \mathbf{R^n_+}$. Let $v : \mathbf{R^n_+} \cup \partial \mathbf{R^n_+} \longrightarrow \partial \mathbf{R^n_+}$ be the vertical projection. Let $\mu \subseteq \partial \mathbf{R^n_+}$ be a minimal G-invariant subspace. Now, p is bounded, and so $\Lambda \setminus \{p\} \subseteq R = \{x \in \mathbf{R^n_+} \mid \mathbf{d_{euc}(x}, \mu) \leq \mathbf{r}\}$ for some $r > 0$. Since $v^{-1}R$ is convex, we have $\operatorname{hull} \Lambda \setminus \{p\} \subseteq v^{-1}R$ (Figure 4a).

Let $T = T_\epsilon(G)$ be as described in Section 3.3. From the discussion given there, it is not hard to see that we can find a horoball B about p such that $B \cap v^{-1}R \subseteq T$. Also, given any horoball B about p, we can find a standard parabolic region C about p such that $C \cap v^{-1}R \subseteq B$. In particular, we conclude the following (Figure 4b):

Lemma 4.1 : *Suppose* $\Gamma \subseteq \mathrm{Isom}\,\mathbf{H^n}$ *is discrete, and* p *is a bounded parabolic fixed point. Let* $G = \mathrm{stab}_\Gamma\, p$*. If* $\epsilon > 0$*, then we can find a horoball* B*, and a standard parabolic region* C*, about* p *such that*

$$C \cap Y \subseteq B \cap Y \subseteq T_\epsilon(G),$$

where $Y = \mathrm{hull}\,\Lambda$.

Note that we have not shown yet that we can choose C to be strictly invariant, i.e. that $\gamma C \cap C = \emptyset$ if $\gamma \in \Gamma \setminus G$. This will be Corollary 4.5.

Proof of GF1⇒GF2 : Suppose Γ is GF1. Let $\pi : \mathbf{H^n} \cup \mathbf{\Omega} \longrightarrow \mathbf{M_C}$ be the projection. We have $M_C = K \cup \bigcup_{i=1}^k E_i$, where K is compact, and each E_i is a standard cusp region. Let Π be the set of parabolic fixed points associated to one of the standard cusp regions. (Thus $\Pi = \bigcup_{i=1}^k \Gamma p_i$ where p_i is the fixed point of the stabliser of a component C_i of $\pi^{-1} E_i$.) From Lemma 3.1.5, we know that each element of Π is a bounded parabolic fixed point.

Suppose $y \in \Lambda \setminus \Pi$. Choose any geodesic ray β tending to y. It is clear that $\beta \cap \pi^{-1} K$ must be unbounded. Thus, the projection $\pi\beta$ must accumulate somewhere in M_C. By Lemma 3.2.1, it follows that y is a conical limit point. $\diamondsuit$

Note that we can find a compact set $K_0 \subseteq \mathbf{H^n}$ such that $\Gamma K_0 = \{\gamma K_0 \mid \gamma \in \Gamma\}$ covers $\pi^{-1} K \cap \mathrm{hull}\,\Lambda$. It is not hard to see that the geodesic β described in the proof, must accumulate somewhere in K_0. This gives the "uniformity of convergence" to conical limit points mentioned in Section 3.2.

We next want to give a proof of GF2⇒GF1. We begin with some general remarks.

Suppose that $Q \subseteq \mathbf{H^n_C}$ is a closed subset such that $Q \cap \mathbf{H^n}$ is dense in Q. We write $N_r(Q)$ for the closure, in $\mathbf{H^n_C}$, of the uniform r-neighbourhood $N_r(Q \cap \mathbf{H^n})$. Given two such subsets, Q_1 and Q_2, we write $\mathrm{hd}(Q_1, Q_2)$ for the minimal $r \in [0, \infty]$ such that both $Q_1 \subseteq N_r(Q_2)$ and $Q_2 \subseteq N_r(Q_1)$. We call $\mathrm{hd}(Q_1, Q_2)$ the *Hausdorff distance* between Q_1 and Q_2. Note that if $\mathrm{hd}(Q_1, Q_2) < \infty$ then $Q_1 \cap \mathbf{H^n_I} = \mathbf{Q_2} \cap \mathbf{H^n_I}$.

We will need the notion of the *nearest point retraction* to convex sets. (See for example [EM].) Suppose that $Q \subseteq \mathbf{H^n_C}$ is closed and convex, and meets $\mathbf{H^n}$ (so that $Q \cap \mathbf{H^n}$ is dense in Q). There is a natural retaction $\rho_Q : \mathbf{H^n_C} \longrightarrow \mathbf{Q}$ which may be defined as follows. If $x \in \mathbf{H^n}$, then $\rho_Q(x)$ is the (unique) nearest point of Q to x. If $x \in Q \cap \mathbf{H^n_I}$, then $\rho_Q(x) = x$. If $x \in \mathbf{H^n_I} \setminus \mathbf{Q}$, then there is a horoball B about x such that $B \cap Q = \partial B \cap Q = \{\rho_Q(x)\}$. One has to check that ρ_Q is continuous. (This is done in [EM] and in [Bow].)

Lemma 4.2 : *Given any* $\lambda > 0$*, there is some* $L = L(\lambda) > 0$ *such that if* $Q_1, Q_2 \subseteq \mathbf{H^n_C}$ *are closed and convex, with* $\mathrm{hd}(Q_1, Q_2) \leq \lambda$*, then for all* $x \in \mathbf{H^n_C} \setminus \mathbf{Q}$*, we* $d(\rho_1(x), \rho_2(x)) < L$*, where* ρ_i *is the retraction* ρ_{Q_i}.

Proof : We first make two observations from hyperbolic trigonometry. First, note that there is a fixed constant L_1 such that if $x, y, z \in \mathbf{H^2}$, with $x\hat{y}z \geq \pi/4$ and $x\hat{z}y \geq \pi/4$, then $d(x, y) \leq L_1$. Second, given any $\lambda > 0$, there is some constant $L_2 = L_2(\lambda)$ such that if $x, y, z \in \mathbf{H^2}$ with $d(y, z) \leq \lambda$ and $d(x, y) \geq L_2$, then $y\hat{x}z < \pi/4$. Let $L = \max(L_1, L_2)$. Note that $L > \lambda$.

Suppose, for contradiction, that there exists some $x \in \mathbf{H^n} \setminus \mathbf{Q}$ with $d(y_1, y_2) \geq L$, where $y_i = \rho_i(x)$. Now $y_2 \in Q_2 \subseteq N_\lambda(Q_1)$, and so there is some $w \in Q_1$ with $d(w, y_2) \leq \lambda$. Thus $w\hat{y}_1 y_2 \leq \pi/4$. By convexity, $[y_1, w] \subseteq Q_1$. Since y_1 is the projection of x to Q_1, we see that $x\hat{y}_1 w \geq \pi/2$, and so $x\hat{y}_1 y_2 \geq \pi/4$. Similarly $x\hat{y}_2 y_1 \geq \pi/4$. We conclude that $d(y_1, y_2) < L_1 \leq L$. This contradiction shows that no such x exists. The case where $x \in \mathbf{H^n_I} \setminus \mathbf{Q}$, follows by continuity. $\diamondsuit$

Corollary 4.3 : *Suppose $X, Q_1, Q_2 \subseteq \mathbf{H^n_C}$ are closed and convex. Suppose* $\mathrm{hd}(Q_1 \cap N_L(X), Q_2 \cap N_L(X)) \leq \lambda$*, where $L = L(\lambda)$ comes from Lemma 4.2. Then, $\rho_1^{-1}(Q_1 \cap X) \subseteq \rho_2^{-1}(Q_2 \cap N_L(X))$, where ρ_i is the nearest point retraction to Q_i.*

Proof : Let ρ'_i be the retraction onto $Q_i \cap N_L(X)$. (Note that $N_L(X)$ is closed and convex.) Suppose that $\rho'_1(x) \in X$. Certainly then $\rho_1(x) = \rho'_1(x)$, and so, by Lemma 4.2, we have $d(\rho'_1(x), \rho'_2(x)) < L$. Thus, $\rho'_2(x)$ lies in the interior of $N_L(X)$. If $x \in \mathbf{H^n}$, we see that $\rho'_2(x)$ is locally the nearest point to x in Q_2. It follows easily that it must, in fact, be globally distance-minimising, and so $\rho_2(x) = \rho'_2(x)$. The case where $x \in \mathbf{H^n_I}$ is similar. $\diamondsuit$

Suppose now that $\Gamma \subseteq \mathrm{Isom}\,\mathbf{H^n}$ is discrete, and let $\Pi \subseteq \Lambda$ be a Γ-invariant set of parabolic fixed points. Suppose that to each $p \in \Pi$ we associate some subset $Q(p) \subseteq \mathbf{H^n_C}$.

Definition : We say that the collection $Q(p)$ is *strictly invariant* if $Q(\gamma p) = \gamma Q(p)$ for all $\gamma \in \Gamma$, and if $Q(p) \cap Q(q) = \emptyset$ if $p \neq q$.

Proposition 4.4 : *Suppose $\Gamma \subseteq \mathrm{Isom}\,\mathbf{H^n}$ is discrete. Let $\Pi \subseteq \Lambda$ be the set of all parabolic fixed points. Then, to each $p \in \Pi$ we may associate a standard parabolic region $C(p) \subseteq \mathbf{H^n} \cup \mathbf{\Omega}$, so that the collection $\{C(p) \,|\, p \in \Pi\}$ is strictly invariant. Moreover, given any $r \in [0, \infty)$, we can arrange that $d(C(p), C(q)) \geq r$ whenever $p \neq q$.*

Proof : If Γ is finite or parabolic, the result is trivial, so we can assume that $Y = \mathrm{hull}\,\Lambda$ meets $\mathbf{H^n}$. The retraction $\rho : \mathbf{H^n_C} \longrightarrow \mathbf{Y}$ is clearly Γ-equivariant. Choose some $\epsilon \in (0, \epsilon(n))$ where $\epsilon(n)$ is the Margulis constant. Then, associated to each $p \in \Pi$, we have the Margulis region $T(p) = T_\epsilon(\mathrm{stab}_\Gamma\, p)$. From the discussion of Section 3.3, we see that the collection $\{T(p) \mid p \in \Pi\}$ is strictly invariant. It follows that the regions $S(p) = \rho^{-1}(T(p) \cap Y)$ are also strictly invariant. We show that each $S(p)$ contains a standard parabolic region $C(p)$.

We fix some $p \in \Pi$, and set $G = \mathrm{stab}_\Gamma\, p$ and $T = T_\epsilon(G)$. We put $p = \infty$ in the upper half-space model. From Lemma 4.1, we can find a horoball B about p such that $Y \cap B \subseteq T$. Let $\mu \subseteq \partial\mathbf{R^n_+}$ be a minimal G-invariant subspace, and let σ be the vertical subspace based on μ (i.e. σ is spanned by $\mu \cup \{p\}$). Now since $\Lambda \neq \infty$, we see that every point of μ lies within a certain bounded euclidean distance from $\Lambda \setminus \{p\} \subseteq \partial\mathbf{R^n_+}$. It follows easily that $\lambda < \infty$, where λ is the hyperbolic hausdorff distance $\mathrm{hd}(Y \cap B, \sigma \cap B)$.

Let $L = L(\lambda)$ be the constant of Lemma 4.2, and let B' be the horoball with $\partial B'$ a hyperbolic distance L above ∂B (so that $B = N_L(B')$). From Corollary 4.3, we see that $\rho_\sigma^{-1}(\sigma \cap B') \subseteq \rho^{-1}(Y \cap B) \subseteq \rho^{-1}(Y \cap T) = S(p)$ (Figure 4c). From Section 3.1, we know that $C(p) = \rho_\sigma^{-1}(\sigma \cap B')$ is a standard parabolic region about p.

Given any $\gamma \in \Gamma$, we may set $C(\gamma p) = \gamma C(p)$. Performing this construction for each orbit of parabolic fixed point in Π, we arrive at a strictly invariant set of standard parabolic regions.

Given any $r > 0$, we may find another set of standard parabolic regions $\{C'(p) | p \in \Pi\}$, so that $C'(p) \subseteq C(p)$, and $d(\partial C'(p), \partial C(p)) \geq r/2$. Then, $d(C'(p), C'(q)) \geq r$ if $p \neq q$. $\diamondsuit$

Corollary 4.5 : *A parabolic fixed point p is associated to a standard cusp region of $M_C(\Gamma)$ if and only if p is bounded.*

Proof : Lemma 3.1.5 and Proposition 4.4. $\diamondsuit$

Proof of GF2⇒GF1 : Suppose Γ is GF2. Let $\Pi \subseteq \Lambda$ be the set of all bounded parabolic fixed points. Lemma 4.3 gives us a collection $\{C(p) \mid p \in \Pi\}$ of standard parabolic regions such that $d(C(p), C(q)) \geq 1$ if $p \neq q$. These project to a disjoint collection $\{E_i \mid i \in I\}$ of standard cusp regions in $M_C(\Gamma)$ where I is some indexing set. Let $K = M_C(\Gamma) \setminus \bigcup_{i \in I} \mathrm{int} E_i$. We claim that K is compact. It then follows that I is finite, and so Γ is GF1.

To prove the claim, choose any $a \in \mathbf{H^n}$, and let $D_\Gamma(a)$ be the Dirichlet domain about a, as described in Section 2.5. Let D be the closure of $D_\Gamma(a)$ in $\mathbf{H^n_C}$. We saw in Section 2.5 that $\mathbf{H^n} \cup \mathbf{\Omega} \subseteq \mathbf{\Gamma D}$. It follows that K is the image under projection to M_C of the set $D' = D \setminus (\Lambda \cup \bigcup_{p \in \Pi} \mathrm{int}\, C(p))$. It thus suffices to see that D' is closed in $\mathbf{H^n_C}$ and hence compact.

Now since D is convex, and ΓD is locally finite in $\mathbf{H^n}$, it is clear that D cannot contain any conical limit point. Since Γ is GF2, we thus have $D \cap \Lambda \subseteq \Pi$, and so $D' = D \setminus (\{p\} \cup \mathrm{int}\, C(p))$. We thus need to see that, for any $p \in D \cap \Pi$, the set $D \cap C(p)$ is a neighbourhood p in D. This most easily seen in the upper half-space model. We have that $(\mathbf{H^n_C} \setminus (\{\mathbf{p}\} \cup \mathrm{int}\, \mathbf{C(p)}))/\, \mathrm{stab}_{\mathbf{\Gamma}}\, \mathbf{p}$ is compact, and the images of D under $\mathrm{stab}_\Gamma\, p$ are locally finite. Thus, $D \setminus (\{p\} \cup C(p))$ is bounded in the euclidean metric. $\diamondsuit$

Proof of GF1⇒GF4 : Suppose Γ is GF1. Write $M_C = K \cup \bigcup_{i=1}^k E_i$, where K is compact, and each E_i is a standard cusp region. Let $\mathbf{H^n} \cup \mathbf{\Omega} \longrightarrow \mathbf{M_C}$ be the projection. Choose a component C_i of $\pi^{-1} E_i$. Let $G_i = \mathrm{stab}_\Gamma\, C_i$, and let p_i be the fixed point of G_i. Thus, p_i is a bounded parabolic fixed point. Suppose $\epsilon \in (0, \epsilon(n))$, where $\epsilon(n)$ is the Margulis constant. By Lemma 4.1, we can find a standard parabolic region, $C_i' \subseteq C_i$, with $C_i' \cap \mathrm{hull}\, \Lambda \subseteq T_\epsilon(G_i)$. Let $E_i' = \pi(C_i') \subseteq E_i \subseteq M_C$. The closure, K', of $M_C \setminus \bigcup_{i=1}^k E_i'$ is compact. Now, $\mathrm{core}(M) \cap \mathrm{thick}_\epsilon(M)$ is a closed subset of K' and hence compact. $\diamondsuit$

Lemma 4.6 : *If Γ is GF4, then every parabolic fixed point is bounded.*

Proof : Suppose Γ is GF4. Let $Y = \mathrm{hull}\,\Lambda$. Let $\epsilon \in (0, \epsilon(n))$, where $\epsilon(n)$ is the Margulis constant. Suppose p is a parabolic fixed point of Γ. Let $G = \mathrm{stab}_\Gamma p$. Put $p = \infty$ in the upper half-space model, $\mathbf{R^n_+}$. Let $v : \mathbf{R^n_+} : \mathbf{R^n_+} \cup \partial\mathbf{R^n_+} \longrightarrow \partial\mathbf{R^n_+}$ be vertical projection. From the description of $T_\epsilon(G)$ given in Section 3.3, it is clear that $v(\partial T_\epsilon(G)) = \partial\mathbf{R^n_+}$. Thus $\Lambda \setminus \{p\} \subseteq v(Y \cap \partial T_\epsilon(G))$. Now $(Y \cap \partial T_\epsilon(G))/G$ may be identified with a component of the boundary of $\mathrm{thick}_\epsilon(M) \cap \mathrm{core}(M)$ in $\mathrm{core}(M)$, and so $(Y \cap T_\epsilon(G))/G$ is compact. It follows that $(\Lambda \setminus \{p\})/G$ is compact. $\diamondsuit$

Proof of GF4⇒GF2 : Suppose that Γ is GF4. We can suppose that Γ is not finite or parabolic, so that $Y = \mathrm{hull}\,\Lambda$ meets $\mathbf{H^n}$. Let $\pi : \mathbf{H^n} \longrightarrow \mathbf{M}$ be the projection.

Suppose $y \in \Lambda$. Let $\beta \subseteq Y$ be a geodesic ray tending to y. If $\beta \subseteq T_\epsilon(\Gamma)$, then $\beta \subseteq T_\epsilon(G)$, where $G \subseteq \Gamma$ is maximal loxodromic or parabolic. It follows that y is either a parabolic fixed point and hence bounded (by Lemma 4.6) or else a loxodromic fixed point and hence a conical limit point. We may thus suppose that $\beta \setminus T_\epsilon(\Gamma)$ is unbounded. It follows that $\pi\beta$ must accumulate somewhere in $\mathrm{core}(M) \cap \mathrm{thick}_\epsilon(M)$. Thus y is a conical limit point. $\diamondsuit$

We next prove GF1⇒GF5.

Suppose $\tau \subseteq \mathbf{H^n_C}$ is a r-dimensional subspace, and that $\rho_\tau : \mathbf{H^n_C} \longrightarrow \tau$ is the nearest point retraction. If $X \subseteq \tau$ is measurable, and $h > 0$, we have that $\mathrm{vol}_n(N_h(\tau) \cap \rho_\tau^{-1}(X)) = k\mathrm{vol}_r(X)$ where vol_i is the i-dimensional volume, and k is some constant depending on h, n and r. In particular, we see that $N_h(\tau) \cap \rho_\tau^{-1}(X)$ has finite n-volume if and only if X has finite r-volume.

Suppose now that G is a discrete parabolic subgroup of $\mathrm{Isom}\,\mathbf{H^n}$ with fixed point p. Let C be a standard parabolic region about p. Thus, C has the form $\rho_\sigma^{-1}(\sigma \cap B)$ where B is a horoball about p, and σ is a (compactified) G-invariant subspace of $\mathbf{H^n_C}$ with $((\sigma \cap \mathbf{H^n_I}) \setminus \{\mathbf{p}\})/\mathbf{G}$ is compact. We must have that $r = \dim\sigma \geq 2$. Now G acts as a cocompact group on $\partial B \cap \sigma = \partial C \cap \sigma$. Thus we can find a compact $K \subseteq \partial B \cap \sigma$ with $\partial B \cap \sigma \subseteq \bigcup GK$. Now, $X = \mathrm{hull}(K \cup p) \setminus \{p\} \subseteq \sigma$ has finite r-volume, and so for any $h > 0$, the set $N_h(\sigma) \cap \rho_\sigma^{-1}(X)$ has finite n-volume. The images of $N_h(\sigma) \cap \rho_\sigma^{-1}(X)$ under G cover $N_h(\sigma) \cap C$, and so $(N_h(\sigma) \cap C)/G$ has finite n-volume.

Proof of GF1⇒GF5 : Suppose Γ is GF1, and choose any $\eta > 0$. By Proposition 2.1.6, there is a bound on the orders of finite subgroups of Γ. Let $M_C = K \cup \bigcup_{i=1}^k E_i$. Let C_i be a component of $\pi^{-1}E_i$. Thus $\mathrm{stab}_\Gamma C_i = \mathrm{stab}_\Gamma p_i$, where p_i is a bounded parabolic fixed point. Using the upper half-space model, it is not hard to see that there is some standard parabolic region $C_i' = \rho_{\sigma_i}^{-1}(\sigma_i \cap B_i) \subseteq C_i$ about p_i, such that $C_i' \cap N_\eta(Y) \subseteq C_i' \cap N_{2\eta}(\sigma)$, where $Y = \mathrm{hull}\,\Lambda$ (Figure 4d). From the discussion immediately before the proof, we see that $(C_i' \cap N_\eta(Y))/G_i$ has finite volume. Each C_i' projects to a standard cusp region $E_i' \subseteq E_i \subseteq M_C$. For each i, therefore, $N_\eta(\mathrm{core}(M)) \cap E_i'$ has finite volume. Since Γ is GF1, $M_C \setminus \bigcup_{i=1}^k E_i$ is relatively compact in M_C. Thus $N_\eta(\mathrm{core}(M)) \setminus \bigcup_{i=1}^k E_i$ is relatively compact in M and thus has finite volume. Thus $N_\eta(\mathrm{core}(M))$ has finite volume. $\diamondsuit$

Proof of GF5⇒GF4 : Suppose Γ is GF5. Then $N_\eta(\mathrm{core}(M))$ has finite volume, and there is some bound, k, on the orders of finite subgroups of Γ. Choose $\epsilon \in (0, \epsilon(n))$ where $\epsilon(n)$ is the Margulis constant. Let $\delta = \min(\eta, \epsilon/2)$. If $x \in \mathrm{thick}_\epsilon(M)$, then we see that $N_\delta(x)$ has volume at most V/k, where V is the volume of a hyperbolic δ-ball. It follows that a disjoint packing of δ-balls in M, each centred on some point of $\mathrm{core}(M) \cap \mathrm{thick}_\epsilon(M)$, must be finite (having at most $k\mathrm{vol}_n(N_\eta(\mathrm{core}(M)))/V$ elements). If we take a maximal such packing, then the corresponding (2δ)-balls must cover $\mathrm{core}(M) \cap \mathrm{thick}_\epsilon(M)$. It follows that $\mathrm{core}(M) \cap \mathrm{thick}_\epsilon(M)$ is compact. $\diamondsuit$

In Section 3.5, we mentioned three situations where the bound on the orders of finite subgroups can be shown to be superfluous. The first case is when $M = \mathbf{H^n}/\mathbf{\Gamma}$ itself has finite volume:

Proposition 4.7 : *A finite volume hyperbolic orbifold (complete, without boundary) is geometrically finite.*

In fact, Proposition 4.7, is a consequence of the work of Garland and Ragunathan [GaR], who construct fundamental domains for lattices in rank-1 semisimple Lie groups. Here we shall offer an alternative proof, which can easily be generalised to the context of orbifolds of pinched negative curvature, as we shall observe at the end. (This fact is quoted in [Bow]). We give the proof as a series of lemmas, aiming to establish that $\mathrm{thick}_\epsilon(M)$ is compact, and so Γ is GF4. The proof of one of the lemmas (4.8) is unfortunately a bit of a mess, so we shall leave it till last.

Let Γ be any discrete subgroup of $\mathrm{Isom}\,\mathbf{H^n}$. Suppose $\epsilon \in (0, \epsilon(n))$. Recall the definitions of $\mathcal{I}_\epsilon(x)$, $\Gamma_\epsilon(x)$ and $T_\epsilon(\Gamma)$ from Sections 2.2(iii) and 3.3. We know from the Margulis Lemma (Theorem 2.2.8) and Lemma 3.3.2, that for any $x \in \mathbf{H^n}$ the group $\Gamma_\epsilon(x)$ is finite, parabolic or loxodromic. Thus, $\mathrm{fix}\,\Gamma_\epsilon(x)$ is either a non-empty (compactified) subspace of $\mathbf{H^n_C}$, or else consists of one or two points of $\mathbf{H^n_I}$. We write

$$f(x) = \mathrm{fix}\,\Gamma_\epsilon(x).$$

We see that f is lower-semicontinuous with respect to set-inclusion. That is to say, given any $x \in \mathbf{H^n}$, there is some neighbourhood, U of x, such that for all $y \in U$, we have $f(x) \subseteq f(y)$.

We define an equivalence relation $\sim$ on $\mathbf{H^n}$ by $x \sim y$ if $f(x) = f(y)$. We write $\mathcal{F}$ for the set of equivalence classes. Thus, $\mathcal{F}$ is a partition of $\mathbf{H^n}$. If $F \in \mathcal{F}$, we write $f(F) = f(x)$, where $x \in F$.

The collection $\mathcal{F}$ is locally finite on $\mathbf{H^n}$. To see this, suppose $x \in \mathbf{H^n}$, and $r > 0$. If $y \in N_r(x)$, then $\mathcal{I}_\epsilon(y) \subseteq \mathcal{I}_{\epsilon+2r}(x)$. Now $\Gamma \cap \mathcal{I}_{\epsilon+2r}(x)$ is finite, and so there are only finitely many possibilities for $\Gamma_\epsilon(y) = \langle \Gamma \cap \mathcal{I}_\epsilon(y) \rangle$.

Let $\mathcal{F}_\infty = \{F \in \mathcal{F} \mid f(F) \subseteq \mathbf{H^n_I}\}$. In other words, if $x \in F \in \mathcal{F}_\infty$, then $\Gamma_\epsilon(x)$ is infinite. Thus, by definition, $\bigcup \mathcal{F}_\infty = T_\epsilon(\Gamma)$. We write $\mathcal{F}_0 = \mathcal{F} \setminus \mathcal{F}_\infty$. From the description of $T_\epsilon(\Gamma)$ in Section 3.3, it is clear that $T_\epsilon(\Gamma) \neq \mathbf{H^n}$, and so $\mathcal{F}_0 \neq \emptyset$. For this, we need to assume that $n \geq 2$. Given $F \in \mathcal{F}_0$, we shall write $d(F)$ for the dimension of the subspace $f(F)$.

Let $F_0 = \{x \in \mathbf{H^n} \mid \mathbf{\Gamma_\epsilon(x)}$ is trivial$\}$. If $F_0 \neq \emptyset$ (as we shall see must be the case), then $F_0 \in \mathcal{F}_0$. In fact, F_0 is the unique element of $\mathcal{F}_0$ with $d(F_0) = n$.

Given $F_1, F_2 \in \mathcal{F}_0$, we write $F_1 \to F_2$ to mean that $F_1 \cap \bar{F}_2 \neq \emptyset$. By the lower semicontinuity of f, this implies that $f(F_1)$ is strictly included in $f(F_2)$ and so $d(F_1) < d(F_2)$.

Lemma 4.8 : *If $F \in \mathcal{F}_0 \setminus \{F_0\}$, then F is not open in $\mathbf{H^n}$.*

We postpone the proof for the moment.

Corollary 4.9 : *If $F \in \mathcal{F}_0 \setminus \{F_0\}$, then there is some $F' \in \mathcal{F}_0$ such that $F \to F'$.*

Proof : By Lemma 4.8, there is some point $x \in F$ which is not an interior point. Since $\mathcal{F}$ is a locally finite partition of $\mathbf{H^n}$, there must be some $F' \in \mathcal{F}$ with $x \in \bar{F}'$. By lower semicontinuity of f, we have $f(F') \supseteq f(F)$, and so $F' \in \mathcal{F}_0$. $\diamondsuit$

We know that $\mathcal{F}_0 \neq \emptyset$. We choose any $F_1 \in \mathcal{F}_0$. Applying Corollary 4.9, we get a sequence $F_i \in \mathcal{F}$ with $F_1 \to F_2 \to F_3 \to \ldots$. This continues until we arrive at some $F_k = F_0$. This must happen after at most n steps, since at each stage we have $d(F_{i+1}) > d(F_i)$. This argument shows that $F_0 \neq \emptyset$. We remark that we have proven:

Proposition 4.10 : *If $\epsilon < \epsilon(n)$, and M is a hyperbolic orbifold (complete, without boundary) of dimension at least 2, then M contains an embedded hyperbolic $(\epsilon/2)$-ball.*

Now, the partition $\mathcal{F}$ of $\mathbf{H^n}$ is Γ-invariant, and so projects to partition, $\mathcal{E}$, of $M = \mathbf{H^n}/\mathbf{\Gamma}$. Thus, each element of $\mathcal{E}$ is the image of an element of $\mathcal{F}$ under the projection $\pi : \mathbf{H^n} \longrightarrow \mathbf{M}$. Let $\mathcal{E}_0 = \{\pi(F) \in \mathcal{E} \mid F \in \mathcal{F}_0\}$. Let $E_0 = \pi(F_0) \in \mathcal{E}_0$. Given $E_1, E_2 \in \mathcal{E}_0$ we write $E_1 \to E_2$ to mean that $F_1 \to F_2$ for some $F_1, F_2 \in \mathcal{F}_0$ with $\pi(F_i) = E_i$. We may think of $\mathcal{E}$ as the vertex set of a directed graph, where we join E_1 to E_2 by an edge if $E_1 \to E_2$. As we saw above, every element of $\mathcal{E}$ can be joined to the element E_0 by a directed path in the graph which contains at most n edges. It follows that the graph has diameter at most $2n$.

Proof of Proposition 4.7 : We suppose now that $M = \mathbf{H^n}/\mathbf{\Gamma}$ has finite volume. Since every 1-dimensional orbifold is topologically finite, we can assume that $n \geq 2$. We claim that $\bigcup \mathcal{E}_0$ is bounded (has finite diameter). Since the closure of $\bigcup \mathcal{E}_0$ in M is, by definition, equal to $\mathrm{thick}_\epsilon(M)$, it then follows that Γ is GF4.

Consider any $E \in \mathcal{E}_0$. Choose $F \in \mathcal{F}_0$ with $\pi(F) = E$. Let $G \subseteq \Gamma$ be the pointwise stabliser of the subspace $f(F)$. Thus, if $y \in F$, we have $\Gamma_\epsilon(y) \subseteq G$, and so $|\Gamma_\epsilon(y)| \leq k = |G|$. It follows that if $x \in E$, then the $(\epsilon/2)$-ball $N_{\epsilon/2}(x) \subseteq M$ has volume at least V/k, where V is the volume of an $(\epsilon/2)$-ball in $\mathbf{H^n}$. Let $X \subseteq E$ be a maximal subset such that the balls $N_{\epsilon/2}(x)$ for $x \in X$ are all disjoint. Now X must be finite (it has at most $k\mathrm{vol}(M)/V$ elements), and $E \subseteq N_\epsilon(X)$. Thus E has finite diameter, and so $\bar{E}$ is compact.

Since the collection $\mathcal{E}_0$ is locally finite in M, it follows that $\bar{E}$ meets the closures of only finitely many other elements of $\mathcal{E}_0$. We conclude that, in the graph described above, each vertex has finite degree. Since the graph also has finite diameter, it must be finite. Thus, $\mathcal{E}_0$ is finite, and so $\bigcup \mathcal{E}_0$ has finite diameter. $\diamondsuit$

It remains to give a proof of Lemma 4.8.

Lemma 4.11 : *Suppose G is a discrete group acting by isometry on S^n. Then there is some $\eta > 0$ such that for all $x \in \mathrm{S}^n$, the group $G_\eta(x)$ has a fixed point in S^n.*

Proof : Given $x \in \mathrm{S}^n$, let $\delta(x) = \min\{d_{sph}(x, gx) \mid g \in G \setminus \mathrm{stab}_G\, x\} \in (0, 2\pi]$. Let $0 < \eta(x) < \delta(x)/3$. If $y \in N_{\eta(x)}(x)$, then $G_{\eta(x)}(x) \subseteq \mathrm{stab}_G\, x$. The result follows by compactness. $\diamondsuit$

Lemma 4.12 : *Suppose $\Gamma \subseteq \mathrm{Isom}\,\mathbf{H^n}$ is discrete, and $\mathcal{F}$ be the partition described above. Suppose $F \in \mathcal{F}_0$ so that $\sigma = f(F)$ is a subspace of $\mathbf{H^n_C}$. Then, there is some $r > 0$ such that $F \subseteq N_r(\sigma)$.*

Proof : Let $G \subseteq \Gamma$ be the pointwise stabliser of σ. If $x \in \sigma \cap \mathbf{H^n}$, we identify the unit tangent space, $T^1_x \mathrm{H}^n$, at x with the unit tangent sphere S^{n-1}. Thus, G acts faithfully by isometry on S^{n-1}. Any two actions arising in this way are conjugate in $\mathrm{Isom}\,\mathrm{S}^{n-1}$, and so we can find some $\eta > 0$ (depending on σ) satisfying the conclusion of Lemma 4.11. We can find $r = r(\eta, \epsilon) > 0$ such that if x, y, z are three points of $\mathbf{H^n}$ with $d(x, y) \geq r$, $d(x, z) \geq r$ and $d(y, z) \leq \epsilon$, then $y\hat{x}z < \eta$. Now suppose $y \in F \setminus N_r(\sigma)$. Let x be the nearest point on σ to y. We have $\Gamma_\epsilon(y) \subseteq G_\eta(\xi)$, where ξ is the unit vector $\overrightarrow{xy}$ based at x. Applying Lemma 4.11, we find that $\Gamma_\epsilon(y)$ fixes some unit vector in $T^1_x \mathrm{H}^n$. Thus, $f(y) = \mathrm{fix}\,\Gamma_\epsilon(y)$ must be strictly larger than σ, contradicting the assumption that $y \in F$. We conclude that $F \subseteq N_r(\sigma)$. $\diamondsuit$

Lemma 4.13 : *Suppose that $Q \subseteq \mathbf{H^n_C}$ is a closed convex set, and that $Z \subseteq \mathbf{H^n}$ is closed in $\mathbf{H^n}$. Suppose $Z \not\subseteq Q$, but that $Z \subseteq N_r(Q)$ for some $r > 0$. Then there is some $x \in Z$ with the following property. If $y \in Z \setminus \{x\}$, and β is the geodesic ray based at y through x, then $\beta \cap Q \subseteq [x, y]$*

We are interested in the case where Q is a (compactified) subspace (Figure 4e).

Proof : We define a relation $<$ on $Z \setminus Q$ by writing $x < y$ if $x \neq y$ and the geodesic ray from y through x meets Q in some point outside the segment $[x, y]$. One verifies that the relation $<$ is transitive and antisymmetric. (This reduces to the two dimensional case). We thus have a partial order on $Z \setminus Q$, and we need to find a $<$-maximal element. We can do this by applying Zorn's lemma, so we need to verify that any (infinite) chain $C \subseteq Z \setminus Q$ has an upper bound. However, since $Z \subseteq N_r(Q)$ we see easily that, for any $x \in Z \setminus Q$, the set $\{y \in Z \mid x \leq y\}$ is compact. We may thus take as an upper bound for C any (in fact the unique) accumulation point of C. $\diamondsuit$

Proof of Lemma 4.8 : We want to show that if $F \in \mathcal{F}_0 \setminus \{F_0\}$, then $F \cap \partial F \neq \emptyset$.

Let $\sigma = f(F)$. Lemma 4.12 gives us some $r > 0$ such that $F \subseteq N_r(\sigma)$. Since $F \neq F_0$, we can suppose that F is not a subset of σ. Let $x \in \bar{F}$ be a point given by Lemma 4.13, with $Z = \bar{F}$ and $Q = \sigma$. Let $\tau = f(x)$. By the lower semicontinuity of f, we have that $\tau \subseteq \sigma$. Again by lower semicontinuity, we can find some neighbourhood U of x such that $f(y) \supseteq \tau$ for all $y \in U$. We can suppose that U is convex, and disjoint from σ.

Given any $u \in \mathbf{H^n}$, define $h(u) \in \sigma$ as follows. If τ is a subspace of σ, we take $h(u)$ to be the nearest point on τ to u. If $\tau = \{p\}$ with $p \in \sigma \cap \mathbf{H^n_I}$, we take $h(u) = p$. If $\tau = \{p, q\}$ with $p \neq q \in \mathbf{H^n_I}$, we take $h(u)$ to be the nearest point to u on the geodesic $[p, q] \subseteq \sigma$. In any case, we see that $h(u)$ varies continuously with u.

Let $a = h(x)$. Let α be the geodesic ray (or bi-infinite geodesic) based at a through the point x. Choose any $y \in U \cap \alpha \setminus [a, x]$. From the choice of x (Lemma 4.13), we know that $y \neq \bar{F}$. Since $x \in \bar{F}$, using the continuity of the map h, we can find $y' \in U \setminus \bar{F}$ (near y) and $x' \in U \cap \bar{F}$ (near x) and $a' \in \sigma$ (near a) so that $a' = h(x')$ and $x' \in [a', y']$ (Figure 4f). Now choose any $z \in [x', y'] \cap \partial F$. Since U is convex, we have $z \in U$, and so $f(z) \supseteq \tau$. Thus, $\Gamma_\epsilon(z)$ fixes τ pointwise. Since $x' \in [a', z]$ and $a' = h(x')$ we clearly must have $\Gamma_\epsilon(z) \subseteq \Gamma_\epsilon(x')$. Thus $f(z) \supseteq f(x') = \sigma$. However, since $z \in \bar{F}$, we have, by lower semicontinuity, that $f(z) \subseteq \sigma$. Thus $f(z) = \sigma$ and so $z \in F$. We have shown that $F \cap \partial F \neq \emptyset$. $\diamondsuit$

This concludes the proof of Proposition 4.7.

We remarked earlier that our argument generalises to the case of orbifolds of pinched negative curvature. That is, if X is a complete simply-connected manifold of pinched negative curvature, and Γ acts properly discontinuously on X with $M = X/\Gamma$ having finite volume, then M is topologically finite (as an orbifold). All the ingredients, notably the thick-thin decomposition, and a lower bound on the volumes of uniform balls are present in this context (see [Bow]), and so the proof proceeds in the same way. However, we should comment on the proof of Lemma 4.13 in this situation. We need to replace $<$ by a different equivalence relation. Given any $\lambda > 0$, and $x, y \in Z \setminus Q$, write $x <_\lambda y$ if $x \neq y$ and $d(y, Q) \geq d(x, Q) + \lambda d(x, y)$. For any $\lambda > 0$, this relation is transitive and antisymetric. By hypothesis, there is some $z \in Z \setminus Q$. Let $\mu = d(z, Q) > 0$. Applying Zorn's lemma, we arrive at a $<_\lambda$-maximal element, x, with $z <_\lambda x$ or $z = x$. Thus $d(z, Q) \geq \mu$. By choosing λ sufficiently small in relation to μ, and using standard comparison theorems, we see that x has the required property.

In Section 3.4, we mentioned another situation in which the bounds on the orders of finite subgroups in the definition GF5 can be shown to be superfluous. This is when the dimension $n \leq 3$. The 2-dimemsional case is simple, so we sketch a proof in the case where $n = 3$.

Proposition 4.14 : *If $M = \mathbf{H^3}/\mathbf{\Gamma}$ is a hyperbolic 3-orbifold, and if for some $\eta > 0$, $N_\eta(\mathrm{core}(M))$ has finite volume, then there is a bound on the orders of finite subgroups of Γ.*

Proof : Note that any finite subgroup, G, of Γ is conjugate to a subgroup of $\mathrm{Isom}\,\mathrm{S}^2$. It follows that, if $|G| > 60$, then G contains a cyclic subgroup, G', of index at most 4 in G. Also, $\beta(G) = \mathrm{fix}\, G'$ is a bi-infinite geodesic in $\mathbf{H^3_C}$, which depends only on G. Also note that any parabolic subgroup of Γ is conjugate to a discrete subgroup of $\mathrm{Isom}\,\mathbf{E^2}$, and so cannot contain any element of order greater than 6. (These particular numbers, 60, 4 and 6, come from the classification of discrete groups in dimension 2. However, the existence of such numbers may be deduced easily from the results of Chapter 2.)

Now, let $\epsilon = \min(\eta, \epsilon(n)/3)$. Let V be the set of all $x \in \mathbf{H^3}$ such that $\Gamma_\epsilon(x)$ is either loxodromic, or else finite of order greater than 60. Given $x \in V$, we set $\beta(x)$ to be the loxodromic axis, in the first case, or to be the geodesic $\beta(\Gamma_\epsilon(x))$ in the second case.

If $x, y \in V$, and $d(x,y) < \epsilon$, then $\Gamma_\epsilon(x)$ and $\Gamma_\epsilon(y)$ are both subgroups of $\Gamma_{3\epsilon}(x)$ which must be finite, parabolic or loxodromic. We conclude that $\beta(x) = \beta(y)$. We thus have a decomposition of V into subsets $V(\beta) = \{x \in V \mid \beta(x) = \beta\}$, so that $d(V(\alpha), V(\beta)) \geq \epsilon$ if $\alpha \neq \beta$. Note that if $G \subseteq \Gamma$ is finite of order greater than 60, and $\beta = \beta(G)$, then $\beta \subseteq V(\beta)$.

Now suppose, for contradiction, that there are finite subgroups $G_i \subseteq \Gamma$ with $|G_i| \to \infty$. We assume that $|G_i| > 60$ for all i. Let $\beta_i = \beta(G_i)$. We can assume that the geodesics β_i are all inequivalent under Γ. Since $|\Lambda| > 2$, we can find $z_i \in \Lambda \setminus \beta_i$. Let y_i be the nearest point on β_i to z_i (in the sense of the nearest point retraction). It is easily seen that $y_i \in \mathrm{hull}\,(G_i z_i) \subseteq \mathrm{hull}\,\Lambda$. Choose $x_i \in [y_i, z_i] \in \partial V(\beta_i)$. Then, $N_{\epsilon/3}(x_i)$ meets at most 60 images of itself under Γ. The projection of the balls $N_{\epsilon/3}(x_i)$ to M give disjoint subsets of $N_\eta(\mathrm{core}(M))$ whose volumes are bounded below. This means that $N_\eta(\mathrm{core}(M))$ has infinite volume. $\diamondsuit$

It remains to show that GF3 is equivalent to the other definitions. We shall use the following notation. Given a collection $\mathcal{B}$ of subsets of $\mathbf{H^n}$, and another fixed subset $A \subseteq \mathbf{H^n}$, we write

$$A \wedge \mathcal{B} = \{A \cap B \mid B \in \mathcal{B}\}.$$

Proof of GF1⇒GF3 : Suppose Γ is GF1. We write $M_C(\Gamma) = K \cup \bigcup_{i=1}^k E_i$, where K is compact, and each E_i is a standard cusp region. Let $\pi : \mathbf{H^n} \cup \mathbf{\Omega} \longrightarrow \mathbf{M_C(\Gamma)}$ be the projection. We shall construct the desired cell complex as a generalised Dirichlet tesselation.

For $i \in \{1, 2, \ldots, k\}$, choose some parabolic fixed point $p_i \in \Lambda$ associated to the end E_i. Let $G_i = \mathrm{stab}_\Gamma\, p_i$. Put $p_i = \infty$ in the upper half-space model, so that E_i has the form C_i/G_i, where $C_i = C(\mu_i, r_i) = \{x \in \mathbf{R^n_+} \mid \mathbf{d_{euc}(x, \mu_i) \geq r_i}\}$ with $r_i > 0$ and μ_i a minimal G_i-invariant subspace of $\partial\mathbf{R^n_+}$. Let $\sigma_i = \langle \mu_i, p_i \rangle$ be the vertical subspace meeting $\partial\mathbf{R^n_+}$ in μ_i. Now choose any $a_i \in \sigma_i \cap \mathrm{int}\, C_i$.

We perform this construction for each $i \in \{1, 2, \ldots, k\}$, and choose an arbitrary $a_0 \in \pi^{-1}K$ (just in case $k = 0$). Set $X = \bigcup_{i=0}^k \Gamma a_i$ and let $\mathcal{D} = \mathcal{D}(X)$ be the generalised Dirichlet tesselation defined by X (Figure 4g). Thus $\mathcal{D}$ is associated to Γ (i.e. $\mathrm{stab}_\Gamma\, D$ is finite for all $D \in \mathcal{D}$). We claim that $\mathcal{D}/\Gamma$ is finite.

Fix some $i \in \{1, 2, \ldots, k\}$. Note that $X \cap C_i = G_i a_i$. In fact, $G_i a_i \subseteq \sigma_i \cap \partial B_i$ where B_i is a horoball contained in $\mathrm{int}\, C_i$. We have that $(\sigma_i \cap \partial B_i)/G_i$ is compact, and so every point of $\sigma_i \cap \partial B_i$ lies within a bounded distance of $G_i a_i$. Simple hyperbolic geometry shows that we can find another standard parabolic region $C'_i = C(\mu_i, r'_i) \subseteq C_i$ such that if $x \in \mathbf{H^n} \cap \mathbf{C'_i}$, then $d_{hyp}(x, G_i a_i) < d_{hyp}(x, \partial C_i)$. Thus the set of nearest points of X to x is a subset of $G_i a_i$. It follows that $C'_i \wedge \mathcal{D}(X) = C'_i \wedge \mathcal{D}(G_i a_i)$. Now in the upper half-space model, we see that the structure of $\mathcal{D}(G_i a_i)$ is independent of the vertical coordinate. So, we restrict attention to $\partial B_i \wedge \mathcal{D}(G_i a_i)$. Note that the euclidean distance in ∂B_i is a certain fixed monotonic function of the hyperbolic distance. Thus, if $y, z, y', z' \in \partial B_i$, we have $d_{euc}(y,z) < d_{euc}(y',z')$ if and only if $d_{hyp}(y,z) < d_{hyp}(y',z')$. It follows that $\partial B_i \wedge \mathcal{D}(G_i a_i)$ is just the euclidean Dirichlet tesselation for the action of G_i on ∂B_i. Since $a_i \in \sigma_i$, this in turn is a euclidean product of $(\sigma_i \cap \partial B_i) \wedge \mathcal{D}(G_i a_i)$ with an orthogonal subspace. Now $(\sigma_i \cap \partial B_i)/G_i$ is compact. We conclude that $((\sigma_i \cap \partial B_i) \wedge \mathcal{D}(G_i a_i))/G_i$, and so also $(\mathcal{D}(G_i a_i))/G_i$ are finite. We have shown that only finitely many orbits of $\mathcal{D}(X)$ under Γ meet C'_i.

We find such a standard parabolic region, C'_i, for each $i \in \{1, 2, \dots, k\}$. These regions project to standard cusp regions $E'_i \subseteq M_C(\Gamma)$. Since Γ is GF1, we have that $K' = M_C(\Gamma) \setminus \bigcup_{i=1}^k \operatorname{int} E'_i$ is compact. From the local finiteness of $\mathcal{D}(X)$, it follows that only finitely many orbits of $\mathcal{D}(X)$ under Γ meet $\pi^{-1}K'$. We conclude that $\mathcal{D}(X)/\Gamma$ is finite, as claimed. $\diamondsuit$

We next aim to prove the converse, GF3$\Rightarrow$GF1. First, we give a few definitions.

In Section 3.5, we saw that a “finite sided polyhedron”, with non-empty interior in $\mathbf{H^n}$, could be defined as a finite intersection of half-spaces in $\mathbf{H^n}$. By a *compactified finite-sided polyhedron* we mean the closure in $\mathbf{H^n_C}$ of such a polyhedron. It is easy to see that a compactified finite-sided polyhedron is a finite intersection of *compactified half-spaces*, where a compactified half-space is the closure in $\mathbf{H^n_C}$ of a half-space in $\mathbf{H^n}$. For the rest of this section, we shall drop the word “compactified”, and assume that our polyhedra and half-spaces are closed in $\mathbf{H^n_C}$.

If $Q \subseteq \mathbf{H^n_C}$ is a closed convex set, one may define a lower-semicontinuous function, $\omega(Q, .) : Q \longrightarrow [0,1]$ which measures the proportion (in the sense of spherical Lebesgue measure) of unit tangent vectors based at x which “point inside” Q. Thus, $\omega(Q,x)$ is the Lebesgue density of Q at x. If Q happens to be a finite-sided polyhedron, then this function has a natural extension to $Q \cap \mathbf{H^n_I}$. In fact, it is only this extension which interests us here, so we give a more detailed treatment in this context.

Suppose that $P \subseteq \mathbf{H^n_C}$ is a (compactified) finite-sided polyhedron with non-empty interior. We may write $P = \bigcap_{i \in I} H_i$ where I is a finite indexing set, and each H_i is a half-space. We want to define $\omega(P,y)$ for $y \in P \cap \mathbf{H^n_I}$.

We first deal with the case where $\mathbf{H^n_I} \cap \bigcap_{\mathbf{i} \in \mathbf{I}} \partial \mathbf{H_i} \neq \emptyset$, and where $y \in \bigcap_{i \in I} \partial H_i$. Put $y = \infty$ in the upper half-space model. Then for each i, $H_i \cap \partial \mathbf{R^n_+}$ is a euclidean half-space, and so $P \cap \partial \mathbf{R^n_+}$ is a finite-sided euclidean polyhedron. We set $\omega(P,y) = \theta(P \cap \partial \mathbf{R^n_+})$, where θ is as defined at the end of Section 3.5. If $\bigcap_{i \in I} \partial H_i = \{y\}$, then $\omega(P,y)$ may take the value 0. However, if $\bigcap_{i \in I} \partial H_i \neq \{y\}$, then we will have $\omega(P,y) > 0$. Moreover, for any $x \in \partial \mathbf{R^n_+} \cap \bigcap_{\mathbf{i} \in \mathbf{I}} \partial \mathbf{H_i}$, we have $\omega(P,x) = \omega(P,y)$. The latter statement follows from the observation that P is symmetric under reflection in any codimension-1 hyperbolic subspace orthogonal to the bi-infinite geodesic $[x,y]$. To see that $\omega(P,y) > 0$, choose any $x \in \partial \mathbf{R^n_+} \cap \bigcap_{\mathbf{i} \in \mathbf{I}} \partial \mathbf{H_i}$. Then $P \cap \partial \mathbf{R^n_+}$ is a euclidean cone about x, in the sense that if a euclidean geodesic ray, based at x, meets $P \cap \partial \mathbf{R}^{\mathbf{n}}_{+}$ at some point other than x, then it lies entirely in $P \cap \partial \mathbf{R}^{\mathbf{n}}_{+}$. Since $P \cap \partial \mathbf{R}^{\mathbf{n}}_{+}$ has non-empty interior in $\partial \mathbf{R}^{\mathbf{n}}_{+}$, it follows easily that $\theta(P \cap \partial \mathbf{R}^{\mathbf{n}}_{+}) > \mathbf{0}$.

Suppose now that $P = \bigcap_{i \in I} H_i$ is any finite sided polyhedron. Suppose that $y \in P \cap \mathbf{H}^{\mathbf{n}}_{\mathbf{I}}$. Let $P(y)$ be the intersection of all those half-spaces H_i for which $y \in \partial H_i$. (Thus, for example, $P(y) = \mathbf{H}^{\mathbf{n}}_{\mathbf{C}}$ if y lies in the topological interior of $P \cap \mathbf{H}^{\mathbf{n}}_{\mathbf{I}}$ in $\mathbf{H}^{\mathbf{n}}_{\mathbf{I}}$.) We set $\omega(P, y) = \omega(P(y), y)$ as defined in the previous paragraph (with the convention that $\omega(\mathbf{H}^{\mathbf{n}}_{\mathbf{C}}, \mathbf{y}) = \mathbf{1}$). One needs to check that $P(y)$ is well-defined however we care to write P as an intersection of half-spaces. (In fact there is a unique minimal such representation.)

If $\omega(P, y) = 0$, we call y a *cusp point* of P. Thus each cusp point is the unique intersection point of the boundaries of some subset of the half-spaces defining P. We write $\kappa(P)$ for the set of all cusp points of P.

By considering all possible intersections of half-spaces, we arrive at:

Lemma 4.15 : *Suppose $P \subseteq \mathbf{H}^{\mathbf{n}}_{\mathbf{C}}$ is a (compactified) finite-sided polyhedron. Then $\kappa(P)$ is finite, and there is some $\delta(P) > 0$ such that if $y \in P \cap \mathbf{H}^{\mathbf{n}}_{\mathbf{I}} \setminus \kappa(\mathbf{P})$, then $\omega(P, y) \geq \delta(P)$.* ♢

Lemma 4.16 : *Suppose that $\mathcal{P}$ is a collection of (compactified) finite-sided polyhedra in $\mathbf{H}^{\mathbf{n}}_{\mathbf{C}}$, with disjoint topological interiors in $\mathbf{H}^{\mathbf{n}}$. Suppose that $y \in \mathbf{H}^{\mathbf{n}}_{\mathbf{I}} \cap \bigcap \mathcal{P}$. Then $\sum_{P \in \mathcal{P}} \omega(P, y) \leq 1$.*

Proof : One checks that if P_1 and P_2 are distinct elements of $\mathcal{P}$, then $P_1(y)$ and $P_2(y)$ (as in the definition of ω) have disjoint interiors. Thus $P_1(y) \cap \mathbf{H}^{\mathbf{n}}_{\mathbf{I}}$ and $P_2(y) \cap \mathbf{H}^{\mathbf{n}}_{\mathbf{I}}$ have disjoint interiors in $\mathbf{H}^{\mathbf{n}}_{\mathbf{I}}$. Now apply Lemma 3.5.14. ♢

Proof of GF3⇒GF1 : Suppose Γ is GF3. Let $\mathcal{A}$ be a convex cell complex associated to Γ with $\mathcal{A}/\Gamma$ finite. Let $\mathcal{P}$ be the set of (compactified) polyhedra formed by taking the closures, in $\mathbf{H}^{\mathbf{n}}_{\mathbf{C}}$, of all the top-dimensional cells of $\mathcal{A}$. Lemma 3.5.13 tells us that each element of $\mathcal{P}$ is, in fact, a finite-sided polyhedron. Also, by Lemma 3.2.1, $\mathcal{P}$ is locally finite on $\mathbf{H}^{\mathbf{n}} \cup \mathbf{\Omega}$, and so $\mathbf{H}^{\mathbf{n}} \cup \mathbf{\Omega} \subseteq \bigcup \mathcal{P}$.

Given any $P \in \mathcal{P}$, we claim that $P \cap \Lambda$ is a subset of $\kappa(P)$, and hence finite. Moreover, each $p \in P \cap \Lambda$ is a bounded parabolic fixed point, and we can find a base of neighbourhoods of p in P consisting of sets of the form $C(p) \cap P$, where $C(p)$ is a standard parabolic region. We can clearly choose these regions so that $C(p) \cap C(q) \cap P = \emptyset$ if $p \neq q \in P \cap \Lambda$.

Given the claim of the last paragraph, the proof that Γ is GF1 may be completed as follows. We choose a set of orbit representatives, $\{P_1, P_2, \ldots, P_m\}$ for $\mathcal{P}$ under Γ. Let $\Pi_0 = \Lambda \cap \bigcup_{j=1}^{m} P_j$. For each $p \in \Pi_0$, we choose a standard parabolic region $C(p)$, in such a way that $C(\gamma p) = \gamma C(p)$ if $\gamma p \in \Pi_0$, and such that $C(p) \cap C(q) \cap P_j = \emptyset$ for each j if $p \neq q \in \Pi_0$. Since the images of the P_j under Γ cover $\mathbf{H}^{\mathbf{n}} \cup \mathbf{\Omega}$, we see that $C(p) \cap C(q) = \emptyset$ whenever p and q are distinct points of Π_0. It follows that the sets $C(p)$ for $p \in \Pi_0$ project to a set of disjoint cusp regions, $\{E_1, E_2, \ldots, E_k\}$ in $M_C(\Gamma)$. Also, for each j, the set $P_j \setminus \bigcup_{p \in \Pi_0} (\{p\} \cup \operatorname{int} C(p))$ is closed in P, and hence compact. Let $\pi : \mathbf{H}^{\mathbf{n}} \cup \mathbf{\Omega} \longrightarrow \mathbf{M}_{\mathbf{C}}(\mathbf{\Gamma})$

be projection. We see that $M_C(\Gamma) \setminus \bigcup_{i=1}^{k} \operatorname{int} E_i = \bigcup_{j=1}^{m} \pi(P_j \setminus \bigcup_{p \in \Pi_0}(\{p\} \cup \operatorname{int} C(p)))$ is compact. Thus Γ is GF1.

We now prove the claims. Fix any $P_0 \in \mathcal{P}$, and suppose $y \in P_0 \cap \mathbf{H^n_I}$. We show that either $y \in \Omega$, or else $y \in \kappa(P_0)$ and is a bounded parabolic fixed point.

Let $G = \operatorname{stab}_\Gamma p$. Certainly G cannot contain any loxodromic element, otherwise we could contradict the local finiteness of $\mathcal{P}$ by iterating the loxodomic with y as repelling fixed point. Thus G is either finite or parabolic.

Let $\mathcal{P}(y) = \{P \in \mathcal{P} \mid y \in P$. Thus $P_0 \in \mathcal{P}(y)$. We put $p = \infty$ in the upper half-space model. we distinguish two cases.

Case (1) : y is a cusp point of each polyhedron of $\mathcal{P}(y)$ (i.e. $\omega(P, y) = 0$ for all $P \in \mathcal{P}(y)$).

By Lemma 4.15, each polyhedron P has only finitely many cusp points, and so $P \cap \Gamma y$ is finite. Since $\mathcal{P}/\Gamma$ is finite, it follows that $\mathcal{P}(y)/G$ is finite. Moreover, we can find a horoball B about y such that each $P \in \mathcal{P}(y)$ meets B in a vertical prism in the upper half-space model. (By this we mean that $P \cap B$ is euclidean-isometric to a euclidean product $(P \cap \partial B) \times [0, \infty)$.) By Lemma 3.5.13, each polyhedron $P \in \mathcal{P}(y)$ has only finitely many faces. Thus, by raising the height of ∂B if necessary, we can arrange that each face of each polyhedron of $\mathcal{P}(y)$ also meets B either in a vertical prism, or not at all. It now follows that $B \subseteq \bigcup \mathcal{P}(y)$. For if not, there must be a polyhedron $P \in \mathcal{P}(y)$ with a codimension-1 face $A \subseteq P$ which meets the boundary of $B \cap \bigcup \mathcal{P}(y)$ in B. However, since A is a vertical prism, we must have $y \in P'$, where P' is the polyhedron of $\mathcal{P}$ on the other side of A. Thus $P' \in \mathcal{P}(y)$, and so A lies in the interior of $\bigcap \mathcal{P}(y)$. This contradiction shows that $B \subseteq \bigcup \mathcal{P}(y)$ as claimed.

We next want to arrange that the horoball B be strictly invariant, i.e. that $\gamma B \cap B = \emptyset$ if $\gamma \in \Gamma \setminus G$. Suppose we have $\gamma B \cap B \neq \emptyset$ for some $\gamma \in \Gamma \setminus G$. Since $B \subseteq \bigcup \mathcal{P}(y)$, it follows that γB meets some polyhedron $P \in \mathcal{P}(y)$. Now, $\gamma^{-1} P \cap B \neq \emptyset$ and so $\gamma^{-1} P \in \mathcal{P}(y)$. From the finiteness of $\mathcal{P}(y)/G$, and of $\operatorname{stab}_\Gamma P$ for each $P \in \mathcal{P}(y)$, we conclude that γ must lie in one of a certain finite number of double cosets of G in Γ. This puts an upper bound on the euclidean height of the highest point of γB in $\mathbf{R^n_+}$. Thus by raising the height of ∂B, we can arrange that B is strictly invariant.

We next want to construct a strictly invariant region C, which will be either a standard parabolic region about y, or else a half-space depending on whether G is parabolic or finite. (In fact, it turns out that the latter case cannot arise in Case (1).)

Since $B \subseteq \bigcup \mathcal{P}(y)$, we see that the boundary of $\bigcup \mathcal{P}(y)$ is a union of lower-dimensional faces of polyhedra in $\mathcal{P}(y)$, none of which meet $y = \infty$. Since $\mathcal{P}(y)/G$ is finite, there is a bound on the euclidean diameters of such faces. Thus, this boundary lies inside some uniform euclidean neighbourhood of a minimal G-invariant subspace, μ, of $\partial \mathbf{R^n_+}$. In other words, for some $r > 0$, we have that $C = \{x \in \mathbf{R^n_+} \cup \partial \mathbf{R^n_+} \mid \mathbf{d_{euc}(x}, \mu) \geq \mathbf{r}\} \subseteq \bigcup \mathcal{P}(\mathbf{y})$. Now, the structure of $\mathcal{P}(y)$ restricted to C is independent of the vertical coordinate. Since $\mathcal{P}(y)$ is locally finite on $\mathbf{R^n_+}$, it must be locally finite on $C \cap \partial \mathbf{R^n_+}$. Thus, $C \cap \partial \mathbf{R^n_+} \subseteq \mathbf{\Omega}$. If G is parabolic, it follows that y is a bounded parabolic fixed point, and that C is a standard parabolic region. Moreover, C can be chosen so that $C \cap P_0$ is an arbitrarily small neighbourhood of y in P_0. If G is finite, then C is a half space, and it follows that $y \in \Omega$. (In this case we may go on to get a contradiction to the hypothesis of Case (1), though logically we do not need this.)

Case (2) : There is some $P_1 \in \mathcal{P}(y)$ with $\omega(P_1, y) > 0$.

Applying Lemma 4.16 to GP_1, and using the fact that $\mathrm{stab}_G P_1$ is finite, we see that G must be finite. Since $\mathcal{P}/\Gamma$ is finite, we may set $\delta = \min\{\delta(P) \mid P \in \mathcal{P}\}$, where $\delta(P)$ is as defined by Lemma 4.15. By Lemma 4.16, again, the set $\mathcal{P}_+(y) = \{P \in \mathcal{P}(y) \mid \omega(P, y) > 0\}$ is finite (having at most $1/\delta$ elements). Also, from the fact that G is finite, and that each polyhedron has a finite number of cusp points, we conclude that $\mathcal{P}(y) \backslash \mathcal{P}_+(y)$ is finite. Thus $\mathcal{P}(y)$ is finite, and using an argument similar to that of Case (1), we find some half-space, containing y, and contained in $\bigcup \mathcal{P}(y)$. Thus $y \in \Omega$. $\diamondsuit$

5. Convex fundamental polyhedra.

The principal objective of this chapter is to give an account of which hyperbolic groups admit finite-sided fundamental polyhedra. We give a complete description of when Dirichlet domains are finite-sided.

Suppose $\Gamma \subseteq \mathrm{Isom}\, \mathbf{H^n}$ is discrete, and suppose $P \subseteq \mathbf{H^n}$ is closed and convex.

Definition : We say that P is a *convex fundamental domain* for Γ if ΓP is locally finite on $\mathbf{H^n}$, if $\bigcup \Gamma P = \mathbf{H^n}$, and if $\mathrm{int}\, P \cap \gamma \mathrm{int}\, P = \emptyset$ for all $\gamma \in \Gamma \setminus \{1\}$.

Note that $\Gamma(\mathrm{int}\, P)$ satisfies the hypotheses of Lemma 3.5.8 (with $\mathbf{H^n}$ replacing $\mathbf{E^n}$). It follows that P is a polyhedron. In Section 3.5, we made a distinction between the "sides" and "faces" of P. To say that P is finite-sided is thus an intrinsic property of P, meaning that P is a finite intersection of half-spaces. To say that P has finitely many faces means that P meets only finitely many images of itself under Γ, and is thus a stronger assertion.

Given any $a \in \mathbf{H^n}$, the Dirichlet domain, $D(a, \Gamma a)$, about a may be defined by

$$D(a, \Gamma a) = \bigcap \{H_\gamma(a) \mid \gamma \in \Gamma \setminus \mathrm{stab}_\Gamma a\},$$

where $H_\gamma(a)$ is the half-space $\{x \in \mathbf{H^n} \mid \mathbf{d}(\mathbf{x}, \mathbf{a}) \leq \mathbf{d}(\mathbf{x}, \gamma\mathbf{a})\}$. If a is not fixed by any element of Γ, then $D(a, \Gamma a)$ is a fundamental domain. For a Dirichlet domain, the notions of sides and faces coincide.

In dimension $n = 3$, we have the following equivalent formulations of geometrical finiteness:

1a (1b) : Some (each) convex fundamental polyhedron has finitely many faces.

2a (2b) : Some (each) Dirichlet domain is finite-sided.

The equivalence of these four notions to GF1 was shown by Marden [Mar], and will follow also from the results of this chapter. However, these definitions diverge in higher dimensions as we will demonstrate below. The problem, of course, comes from the parabolic cusps, so we need to say something about euclidean Dirichlet tesselations.

As in Section 2.1(ii), we define the $(n-1)$-sphere at infinity, $S(\mathbf{E^n})$, of euclidean space to be the set of parallel classes of geodesic rays. Thus the set of all geodesic rays in $\mathbf{E^n}$ may be identified with $\mathbf{E^n} \times \mathbf{S}(\mathbf{E^n})$ where the first coordinate is the basepoint of the ray, and the second coordinate is the "direction" of the ray. We shall write $[\beta]$ for the direction of the ray β. Note that $S(\mathbf{E^n})$ comes equipped with a natural spherical metric, and that $\operatorname{Isom}\mathbf{E^n}$ acts by isometry on $S(\mathbf{E^n})$.

Given a closed convex set $P \subseteq \mathbf{E^n}$, we set (as in Section 3.5) $\Theta(P)$ to be the set of directions of all those rays lying in P. Thus, $\Theta(P)$ is closed in $S(\mathbf{E^n})$, and we write $\operatorname{int}\Theta(P)$ for its interior. One can check that if P and Q are convex, then $\operatorname{int}\Theta(P\cap Q) = \operatorname{int}\Theta(P) \cap \operatorname{int}\Theta(Q)$.

We now go in to consider a specific example. Let $\Gamma \subseteq \operatorname{Isom}\mathbf{E^3}$ be an infinite cyclic group generated by an "irrational screw motion" of $\mathbf{E^3}$. By this we mean a translation parallel to τ composed with an irrational rotation of $\mathbf{E^3}$ with axis τ (Figure 5a). We claim that if $a \in \mathbf{E^3} \setminus \tau$, then the Dirichlet domain $D(a, \Gamma a)$ is infinite-sided.

To see this, let β be the ray through a meeting τ orthogonally at its basepoint. Now, Γ acts by irrational rotation on $S(\mathbf{E^3})$, and so $[\gamma\beta] \neq [\beta]$ for all $\gamma \in \Gamma \setminus \{1\}$. We see that, for all $\gamma \in \Gamma \setminus \{1\}$, we have $[\beta] \in \operatorname{int}\Theta(H_\gamma(a))$, where $H_\gamma(a)$ is the half-space $\{x \in \mathbf{E^3} \mid \mathbf{d(x,a)} \leq \mathbf{d(x,\gamma a)}\}$. If $D = D(a,\Gamma a)$ were finite sided, we could write it as an intersection $D = \bigcap_{\gamma\in G} H_\gamma(a)$, where $G \subseteq \Gamma$ is finite. Thus, $[\beta] \in \operatorname{int}\Theta(D) = \bigcap_{\gamma\in G} \operatorname{int}\Theta(H_\gamma(a))$. Since Γ acts non-discretely on $S(\mathbf{E^3})$, there must be some $\gamma \in \Gamma$ with $\operatorname{int}\Theta(D) \cap \gamma\operatorname{int}\Theta(D) \neq \emptyset$. Thus $D \cap \gamma D$ has non-empty interior. This contradiction shows that D is infinite-sided.

The above argument does not give much insight into what the Dirichlet tesselation actually looks like, so we shall give an informal qualitative description of this. We describe only the tesselation a long way from the axis τ. To this end, we imagine intersecting the tesselation with cylinders $S_r = \{x \in \mathbf{E^3} \mid \mathbf{d}(\mathbf{x},\tau) = \mathbf{r}\}$, and see how the picture changes as $r \to \infty$. Let $\tilde{S}_r$ be the universal cover of S_r. In the induced Riemannian metric, $\tilde{S}_r$ is isometric to the euclidean plane, $\mathbf{E^2}$. There is a $S \oplus S$-action on this plane generated by Γ and the covering transformations of $\tilde{S}_r$ over S_r. The Dirichlet tesselation gives us a representation of $\mathbf{E^2}$ as a CW-complex invariant under this action. In the generic situation, this decomposition is combinatorially equivalent to a regular hexagon tesselation of the plane. As r tends to infinity, the pattern of hexagons changes by an infinite sequence of "Whitehead moves". This process is best described with reference to the quotient torus, $S_r/\Gamma \equiv \mathbf{E^2}/(\mathbf{S}\oplus\mathbf{S})$. For a generic r, this torus is decomposed into two 0-cells, three 1-cells and one 2-cell. As r becomes critical, one of the 1-cells collapses to a single point, giving rise (combinatorially) to a square tesselation of $\mathbf{E^2}$. The 4-valent vertex then splits again into two 3-valent vertices to give another hexagon tesselation (Figure 5b). The combinatorial structure of the Dirichlet tesselation far away from τ, is thus determined by the sequence of 1-cells which get contracted by Whitehead moves. This sequence is, in turn, determined by the continued fraction expansion of the rotation angle θ, measured as a fraction of a full rotation. (The situation is analogous to following a geodesic in the moduli space of euclidean tori — see [Ser].) The geometry of the Dirichlet domain, D, is also related to rational approximation of θ. Clearly, the area of the cross section $D \cap S_r$ grows linearly with r, but its diameter grows much more quickly in the radial direction than in the longitudinal direction (parallel to τ). The relative rates depend on rational approximations to θ — the better θ is approximated, the quicker the cross sections flatten out radially. For a quadratic surd, the radial diameter grows asymptotically like $r^{3/4}$, while the longitudinal diameter grows like $r^{1/4}$.

Now, we may extend our cyclic group, Γ, to act on $\mathbf{H^4}$ as a parabolic group, with $\mathbf{E^3} \subseteq \mathbf{H^4}$ a horosphere about the fixed point p. Let ρ be the 2-dimensional subspace spanned by τ and p. If $a \in \mathbf{H^4} \setminus \rho$, the hyperbolic Dirichlet domain $D(a, \Gamma a)$ will be infinite-sided. (This is best seen by putting $p = \infty$ in the upper half-space model, so that $D(a, \Gamma a)$ is a vertical euclidean prism, with base a euclidean Dirirchlet domain.) However, Γ is geometrically finite with any of the definitions of Chapter 3.

We may now find a half-space, in $\mathbf{H^4}$, disjoint from all its images under Γ, and disjoint from ρ. This set projects to an embedded half space in the quotient manifold M. By removing this half space, and doubling M across the boundary, we get a new manifold M', with fundamental group $S * S$. This gives us a geometrically finite action of $S * S$ on $\mathbf{H^4}$ with no finite sided Dirichlet domain. This example was constructed by Apanasov.

We remark that we may carry this construction further to produce an example of a discrete group acting on $\mathbf{H^4}$ which has a parabolic fixed point with no strictly invariant horoball. To do this, put $p = \infty$ in the upper half-space model, and note that we can find an infinite sequence of half-spaces in $\mathbf{H^4}$, whose euclidean diameters tend to ∞, and which project to disjoint embedded half-spaces in the quotient manifold. We now cut out all these half spaces, and double the resulting manifold in its boundary. This gives us an action of an infinitely generated free group on $\mathbf{H^4}$. There is no strictly invariant horoball about p. It seems to be an open question as to whether a finitely generated disctrete group must have a strictly invariant set of horoballs for each orbit of parabolic fixed point. In dimensions 2 and 3, this is a consequence of the Margulis Lemma (even for infinitely generated groups). For geometrically finite groups, it is a consequence of Lemma 4.6 and Corollary 4.5.

We now move on to a general account of Dirichlet domains.

We say that an isometry, γ, of $\mathbf{E^n}$ is *rational* if some power of γ is a translation. We saw in Chapter 2 that any discrete subgroup of $\mathbf{E^n}$ is virtually abelian. Thus, if $\Gamma \subseteq \mathbf{E^n}$ is discrete and consists entirely of rational isometries, then Γ contains a finite index translation group.

If $\gamma \in \operatorname{Isom} \mathbf{H^n}$ is parabolic, then we say that γ is *rational* if γ restricted to a horosphere about the fixed point is a rational euclidean isometry. Otherwise, we say that γ is *irrational*.

Suppose that $\Gamma \subseteq \operatorname{Isom} \mathbf{H^n}$ is geometrically finite. We shall see that if every parabolic element of Γ is rational, then every Dirichlet domain is finite-sided. If, however, Γ contains an irrational parabolic, then almost every Dirichlet domain is infinite-sided, i.e. for at least an open dense set of basepoints $a \in \mathbf{H^n}$, $D(a, \Gamma a)$ is infinite-sided.

Note that if $n \leq 3$, then every parabolic element is rational. This shows the equivalence of properties (2a) and (2b) with the the other notions of geometrical finiteness in these dimensions.

We first discuss the euclidean case.

Suppose $\Gamma \subseteq \operatorname{Isom} \mathbf{E^n}$ is discrete. In Chapter 2, we saw that if μ is a minimal Γ-

invariant subspace, then Γ contains a finite-index subgroup which acts by translation on μ. Suppose that τ_1 and τ_2 are subspaces of $\mathbf{E^n}$, and that Γ_1 and Γ_2 are subgroups of Γ. Suppose that τ_i is Γ_i-invariant, and that Γ_i acts by translation of τ_i. Then it's not hard to see that $\Gamma_1 \cap \Gamma_2$ acts by translation on the subspace $\langle \tau_1, \tau_2 \rangle$ spanned by τ_1 and τ_2. We may thus define $\tau = \tau(\Gamma)$ to be the maximal subspace of $\mathbf{E^n}$ on which some finite-index subgroup acts by translation. Let $\Gamma' \subseteq \Gamma$ be the subgroup of all elements acting by translation on τ. If $g \in \Gamma$, then $g\Gamma' g^{-1}$ acts by translation on $g\tau$. Thus, $g\tau = \tau$ and $g\Gamma' g^{-1} = \Gamma'$. In other words, τ is Γ-invariant, and Γ' is normal in Γ. Note that the set of all minimal Γ-invariant subspaces form a foliation of some subspace of τ.

Proposition 5.1 : *Suppose that $\Gamma \subseteq \operatorname{Isom} \mathbf{E^n}$ is discrete, and that $a \in \mathbf{E^n}$. Then the Dirichlet domain $D(a, \Gamma a)$ is finite sided if and only if $a \in \tau(\Gamma)$.*

Proof : Write $\tau = \tau(\Gamma)$ and $D = D(a, \Gamma a)$.

Suppose first that $a \in \tau$. Now $D \cap \tau$ is the Dirichlet domain for the action of Γ restricted to τ, and D is a product of $D \cap \tau$ with an orthogonal subspace. Since Γ has a finite-index subgroup acting by translation on τ, one sees easily that $D \cap \tau$ is finite-sided. (This also follows from Lemma 5.4 below.)

Suppose now that $a \in \mathbf{E^n} \setminus \tau$. Let b be the nearest point on τ to a, and let β be the geodesic ray based at b through a. Thus β is orthogonal to τ, and it is easily verified that β lies in the interior of D.

Now, as described above, Γ acts by isometry on the $(n-1)$-sphere, $S(\mathbf{E^n})$. Let $\Gamma_0 \subseteq \Gamma$ be the subgroup of elements which preserve the direction $[\beta] \in S(\mathbf{E^n})$ of the ray β. By the maximality of τ, we see that Γ_0 must have infinite index in Γ. It follows that the action of Γ on $S(\mathbf{E^n})$ is not discrete. Note that $\operatorname{stab}_\Gamma D = \operatorname{stab}_\Gamma a \subseteq \Gamma_0$.

Now suppose, for contradiction, that D is finite-sided. Thus we may write D as a finite intersection of half-spaces $D = \bigcap_{\gamma \in G} H_\gamma(a)$, where $G \subseteq \Gamma$ is finite.

If Γ_0 were trivial, the proof could proceed as in the case of an irrational screw motion on $\mathbf{E^n}$ as described above. We would have $[\beta] \in \operatorname{int} \Theta(D)$, so we could find $\gamma \in \Gamma$ with $[\gamma\beta] \in \operatorname{int} \Theta(D)$, and arrive at the contradiction that $D \cap \gamma D$ has non-empty interior.

To deal with the general case, we write $D = D_1 \cap D_2$, where $D_1 = \bigcap\{H_\gamma(a) | \gamma \in G \setminus \Gamma_0\}$ and $D_2 = \bigcap\{H_\gamma(a) \mid \gamma \in \Gamma_0\}$.

For all $\gamma \in G \setminus \Gamma_0$, we have $[\gamma\beta] \neq [\beta]$ and so $[\beta] \in \operatorname{int} \Theta(H_\gamma(a))$. Thus $[\beta] \in \operatorname{int} \Theta(D_1)$. It follows that there is a ball of some radius, $\epsilon > 0$, about $[\beta]$, in the spherical metric on $S(\mathbf{E^n})$ which is contained in $\Theta(D_1)$. Since Γ acts non-discretely on $S(\mathbf{E^n})$, we can find some $\gamma_1 \in \Gamma \setminus \Gamma_0$ with $d_{sph}([\beta], [\gamma_1\beta]) < \epsilon$.

Now, D_2 is the Dirichlet domain for the action of Γ_0 on $\mathbf{E^n}$. Since Γ_0 preserves the subspace τ and the direction $[\beta]$, it follows that D_2 is a euclidean product of $D_2 \cap \tau$ with an orthogonal subspace, and that $D_2 \cap \tau$ is the Dirichlet domain about the point b, for the action of Γ_0 restricted to τ. Thus there is some $\gamma_2 \in \Gamma_0$, such that $\gamma_2\gamma_1 b \in D_2 \cap \tau$. Since β is orthogonal to τ, we have $\gamma_2\gamma_1\beta \subseteq D_2$. Now, $d_{sph}([\gamma_2\gamma_1\beta], [\beta]) = d_{sph}([\gamma_2\gamma_1\beta], [\gamma_2\beta]) = d_{sph}([\gamma_1\beta], [\beta]) \leq \epsilon$. Thus $[\gamma_2\gamma_1\beta] \in \operatorname{int} \Theta(D_1)$. It follows that $\gamma_2\gamma_1\beta$ meets D_1, and so meets $D_1 \cap D_2 = D$. Thus $\beta \cap (\gamma_2\gamma_1)^{-1} D \neq \emptyset$. Since $\gamma_2\gamma_1 \notin \operatorname{stab}_\Gamma a$ and since β lies in the interior of D, we arrive at a contradiction. ◊

Using a similar argument, we may deduce the following generalisation:

Proposition 5.2 : *Suppose* $\Gamma \subseteq \mathrm{Isom}\,\mathbf{E^n}$ *is discrete, and* $X \subseteq \mathbf{E^n}$ *is a finite union of* Γ*-orbits (i.e.* X/Γ *is finite). Then, the generalised Dirichlet domains* $D(a,X)$ *are all finite-sided if and only if* $X \subseteq \tau(\Gamma)$.

Proof (sketch) : The "if" part, as with Proposition 5.1, is easy. So, suppose that $X \not\subseteq \tau = \tau(\Gamma)$. Let $Y \subseteq X$ be the set of points a maximal distance from τ. Thus Y is Γ-invariant, and Y/Γ is finite. It is easily seen that, sufficiently far away from τ, the generalised Dirichlet tesselations, $\mathcal{D}(X)$ and $\mathcal{D}(Y)$ agree. It thus suffices to see that some element of $\mathcal{D}(Y)$ is infinite-sided.

Let $a_1, a_2, \dots, a_m$ be a complete set of orbit representatives of Y under Γ. Let b_i be the nearest point on τ to a_i, and let β_i be the geodesic ray based at b_i through a_i. Let $I \subseteq \{1, 2, \dots, m\}$ be the set of all those i for which there is some $\gamma \in \Gamma$ with $[\gamma\beta_i] = [\beta_1]$. Thus $I \neq \emptyset$, and we may as well assume that $[\beta_i] = [\beta_1]$ for all $i \in I$. Note that the orbits Γb_i must be disjoint. Let $\Gamma_0 \subseteq \Gamma$ be the subgroup fixing the direction $[\beta_1]$, so that $[\Gamma : \Gamma_0] = \infty$.

Suppose now that, for each $i \in I$, the generalised Dirichlet domain, $D(a_i, Y)$, is finite sided. Following the proof of Proposition 5.1, for each $i \in I$, we may write $D(a_i, Y) = D_1^i \cap D_2^i$ with $[\beta_1] \in \mathrm{int}\,\Theta(D_1^i)$ and $D_2^i = D(a_i, \bigcup_{i \in I} \Gamma_0 a_i)$. We see that D_2^i is a euclidean product of $D_2^i \cap \tau$ with an orthogonal subspace, and that $D_2^i \cap \tau = D_\tau(b_i, \bigcup_{i \in I} \Gamma_0 b_i)$ where D_τ denotes the generalised Dirichlet domain restricted to τ. We now proceed to a contradiction as in Proposition 5.1. $\diamondsuit$

We now want to apply these results to hyperbolic groups. Suppose that $\Gamma \subseteq \mathrm{Isom}\,\mathbf{H^n}$ is discrete. Given $a \in \mathbf{H^n}$, we write $D_C(a, \Gamma a)$ for the closure of $D(a, \Gamma a)$ in $\mathbf{H^n_C}$.

Suppose that $\Pi_0 \subseteq \mathbf{H^n_I}$ is a an orbit of parabolic fixed points, which admits a strictly invariant collection of horoballs $\{B(p) \mid p \in \Pi_0\}$ (so that $B(p) \cap B(q) = \emptyset$ if $p \neq q$). Given $x \in \mathbf{H^n}$, write $J(x)$ for the set of $p \in \Pi_0$ which minimise $d(x, B(p))$. We check that the hypotheses of the construction of Section 3.5 are satisfied. Thus we arrive at a Γ-invariant convex cell complex

$$\mathcal{A}(\Pi_0, \Gamma) = \{A(Q, \Pi_0) \mid Q \subseteq \Pi_0 \text{ is finite}\},$$

where $A(Q, \Pi_0) = \{x \in \mathbf{H^n} \mid \mathbf{J(x) = Q}\}$ (so that $A(\emptyset, \Pi_0) = \emptyset$). We see that $x \in A(Q, \Pi_0)$ if and only if $\Pi_0 \cap D_C(x, \Gamma x) = Q$. Note that $\mathcal{A}(\Pi_0, \Gamma)$ is independent of the choice of strictly invariant horoballs.

Now suppose that Π_i, for $i = 1, 2, \dots, k$, are a finite set of distinct orbits of parabolic fixed points, each admitting a collection of strictly invariant horoballs. Set $\Pi = \bigcup_{i=1}^k \Pi_i$. The construction of the previous paragraph gives a set of k convex cell complexes, $\mathcal{A}(\Pi_i, \Gamma)$. These have a common subdivision $\mathcal{A}(\Pi, \Gamma) = \bigwedge_{i=1}^k \mathcal{A}(\Pi_i, \Gamma)$ obtained by intersecting all the cells. Thus, each element of $\mathcal{A}(\Pi, \Gamma)$ has the form $A(Q, \Pi) = \bigcap_{i=1}^k A(Q \cap \Pi_i, \Pi_i)$, where $Q \subseteq \Pi$ is finite. Again, x lies in $A(Q, \Pi)$ if and only if $\Pi \cap D_C(x, \Gamma x) = Q$. Note that, generically, $D_C(x, \Gamma x)$ meets Π in an orbit-transversal.

Suppose now that Γ is geometrically finite, and let Π be the set of all (bounded) parabolic fixed points. Write Π as a disjoint union, $\Pi = \bigsqcup_{i=1}^k \Pi_i$, of orbits under Γ. Let

$\mathcal{A} = \mathcal{A}(\Pi, \Gamma) = \{A(Q, \Pi) \mid Q \subseteq \Pi \text{ is finite}\}$. Given $p \in \Pi$, we may put $p = \infty$ in the upper half-space model, and set $\tau = \tau(G) \subseteq \partial\mathbf{R}^{\mathbf{n}}_{+}$ where $G = \mathrm{stab}_{\Gamma}\, p$ acts on $\partial\mathbf{R}^{\mathbf{n}}_{+}$. We set $\rho(p) \subseteq \mathbf{H}^{\mathbf{n}}$ to be the hyperbolic subspace which is compactified by $\{p\} \cup \tau$. The construction of $\rho(p)$ is natural, so we have $\rho(\gamma p) = \gamma\rho(p)$ for all $p \in \Pi$ and $\gamma \in \Gamma$. Given a finite subset $Q \subseteq \Pi$, we set $\rho(Q) = \bigcap_{p \in Q} \rho(p)$. Let $F(\Gamma) = \bigcup\{A(Q, \Pi) \cap \rho(Q) \mid Q \subseteq \Pi \text{ is finite}\}$. Note that if every parabolic element is rational, then $F(\Gamma) = \mathbf{H}^{\mathbf{n}}$. Otherwise $\mathbf{H}^{\mathbf{n}} \setminus \mathbf{F}(\mathbf{\Gamma})$ contains an open dense subset of $\mathbf{H}^{\mathbf{n}}$.

Proposition 5.3 : *Suppose $\Gamma \subseteq \mathrm{Isom}\, \mathbf{H}^{\mathbf{n}}$ is geometrically finite. Then $D(a, \Gamma a)$ is finite-sided if and only if $a \in F(\Gamma)$.*

Proof : Suppose first that a lies in a top-dimensional cell of $\mathcal{A}$, so that $D_C(a, \Gamma a)$ meets a single representative p_i, from each orbit of parabolic fixed points, Π_i. Let $B(p_i)$ be a strictly invariant horoball about p_i. We see that the set of nearest points to $B(p_i)$ in Γa is precisely $G_i a$, where $G_i = \mathrm{stab}_{\Gamma}\, p_i$. It is not hard to see that we can find another horoball $B'(p_i) \subseteq B(p_i)$ such that $D(a, \Gamma a) \cap \partial B'(p_i)$ is precisely the euclidean Dirichlet domain for the action of G_i on $\partial B'(p_i)$ based at the intersection point of $\partial B'(p_i)$ and $[a_i, p_i]$ (c.f. the proof of GF1$\Rightarrow$GF3). This euclidean domain is finite-sided if an only if $a \in \rho(p_i)$. Moreover, one can find a standard parabolic region $C(p_i)$ about p_i so that that the intersection of the Dirichlet tesselation $\mathcal{D}(\Gamma a)$ with $C(p_i)$ is independent of the vertical coordinate in the upper half-space model. We see that $C(p_i)$ meets only finitely many faces of $D(a, \Gamma a)$ if and only if $a \in \rho(p_i)$. The result in this case now follows from the description GF1 of geometrical finiteness.

The general case can be dealt with using Proposition 5.2 in place of Proposition 5.1. Suppose $a \in A(Q, \Pi)$, and that $Q \cap \Pi_i = \{p_i, \gamma_1 p_i, \ldots, \gamma_r p_i\}$. In this case, we can find a horoball $B'(p_i)$ about p_i so that $D(a, \Gamma a) \cap \partial B'(p_i) = D_{\partial B'(p_i)}(b_i, G_i b_i \cup \bigcup_{j=1}^{r} G_i \gamma_j^{-1} b_i)$, where $[a_i, p_i] \cap \partial B'(p_i) = \{b_i\}$ and $D_{\partial B'(p_i)}$ denotes the euclidean generalised Dirichlet domain in $\partial B'(p_i)$. $\diamondsuit$

We now say a few words about convex fundamental domains in general.

Lemma 5.4 : *Suppose $\Gamma \subseteq \mathrm{Isom}\, \mathbf{E}^{\mathbf{n}}$ is discrete and acts by translation. Suppose $P_1, P_2, \ldots, P_k \subseteq \mathbf{E}^{\mathbf{n}}$ are closed convex subsets, with non-empty interiors, such that $\mathbf{E}^{\mathbf{n}} = \bigcup_{\mathbf{i=1}}^{\mathbf{k}} \bigcup \mathbf{\Gamma P_i}$ and such that $\mathrm{int}\, P_i \cap \gamma \mathrm{int}\, P_j = \emptyset$ unless $i = j$ and $\gamma = 1$. Then, each P_i is a finite sided polyhedron.*

Proof : We leave as an exercise that each ΓP_i must be locally finite. Thus the set $\mathcal{U} = \bigcup_{i=1}^{k} \Gamma(\mathrm{int}\, P_i)$ satisfies the hypotheses of Lemma 3.5.8, and so gives rise to a convex cell complex $\mathcal{A} = \mathcal{A}(\mathcal{U})$. It follows that each P_i is a polyhedron. Write $\mathcal{F}^{n-1}(P_i)$ for the set of codimension-1 faces of P_i. From the construction of $\mathcal{A}$, we see that if $A \in \mathcal{F}^{n-1}(P_i)$, then $\bar{A} = P_i \cap \gamma P_j$ for a unique $\gamma \in \Gamma$ and $j \in \{1, \ldots, k\}$. Now Γ is free abelian, so there is an isomorphism $\phi : \Gamma \longrightarrow S^m$ for some $m \leq n$. Let $\rho : S^m \longrightarrow S_2^m$ be reduction mod 2. We define $\lambda : \mathcal{F}^{n-1}(P_i) \longrightarrow \{1, \ldots, k\} \times S_2^m$ by setting $\lambda(A) = (j, \rho \circ \phi(\gamma))$ where $\bar{A} = P_i \cap \gamma P_j$. Thus we can think of λ as labelling the codimension-1 faces of P_i with at most $2^m k$ labels.

We claim that if $\lambda(A_1) = \lambda(A_2)$ then A_1 and A_2 lie in the same codimension-1 subspace. From this it follows easily that P_i is a finite intersection of half-spaces.

Suppose then that $\lambda(A_1) = \lambda(A_2)$. Thus $\bar{A}_1 = P_i \cap \gamma_1 P_j$ and $\bar{A}_2 = P_i \cap \gamma_2 P_j$, with $\gamma_1^{-1}\gamma_2 = g^2$ for some $g \in \Gamma$. Suppose that A_1 and A_2 do not lie in the same codimension-1 subspace. Choose $a_1 \in A_1$ and $a_2 \in A_2$. Let b be the midpoint of $[a_1, a_2]$ so that $b \in \operatorname{int}\operatorname{hull}(A_1 \cup A_2) \subseteq \operatorname{int} P_i$. Now b is also the midpoint of $[ga_1, g^{-1}a_2]$. We have $ga_1 \in g\gamma_1 P_j$ and $g^{-1}a_2 \in g^{-1}\gamma_2 P_j = g\gamma_1 P_j$. Thus $b \in g\gamma_1 P_j$, and so $\operatorname{int} P_i \cap g\gamma_1 P_j \neq \emptyset$. Since P_j has non-empty interior, this gives a contradiction. $\diamondsuit$

It is possible that the polyhedra P_i may have infinitely many faces. Consider the tesselation of $\mathbf{E^3}$ by planks described in Section 3.5 (Figure 3g). This tesselation is invariant under a $S \oplus S$-action, where one generator translates vertically through two layers, and the other translates in a north-easterly direction. The planks satisfy the hypotheses of Lemma 5.4, but each has infinitely many faces. This phenomenon, however, cannot occur in dimensions 1 and 2:

Lemma 5.5 : *With the same hypotheses as Lemma 5.4, if $n \leq 2$, then each polyhedron has finitely many faces.*

Proof : The only case that requires verification is that of an infinite cyclic action on $\mathbf{E^2}$. Figure 5c shows the typical situations that can arise. We leave the reader to work out the details. $\diamondsuit$

Proposition 5.6 : *Suppose that $\Gamma \subseteq \operatorname{Isom}\mathbf{H^n}$ is geometrically finite, and that each parabolic element is rational. Then every convex fundamental domain for Γ is finite-sided.*

Proof : Let P be a convex fundamental domain for Γ. Suppose p is a (bounded) parabolic fixed point of Γ. Let $G = \operatorname{stab}_\Gamma p$ and let C_3 be a standard parabolic region about p. Let $C_2 \subseteq \operatorname{int} C_3$ and $C_1 \subseteq \operatorname{int} C_1$ be strictly smaller standard parabolic regions. Now $(\partial C_2 \cap \operatorname{hull}(C_1 \cup C_3))/G$ is compact. From the local finiteness of ΓP on $\mathbf{H^n}$, we see that $\{\gamma \in \Gamma \mid \gamma P \cap C_3 \neq \emptyset\}$ consists of finitely many right cosets of G in Γ (c.f. Lemma 3.2.1). We can now find another standard parabolic region $C \subseteq C_3$ such that $\gamma P \cap C \neq \emptyset$ if and only p lies in the closure of γP in $\mathbf{H^n_C}$. Thus, in the upper half-space model, the intersection of γP with C is independent of the vertical coordinate, and so it is enough to show that $\gamma P \cap \partial B$ is a finite-sided euclidean polyhedron for some horoball $B \subseteq C$. Now G contains some finite index normal subgroup G' which acts by translation on ∂B. Writing $\{\gamma \in \Gamma \mid \gamma P \cap C \neq \emptyset\} = \bigsqcup_{i=1}^m G'\gamma_i$, we see that $\partial B \cap \gamma_1 P, \ldots, \partial B \cap \gamma_m P$ satisfy the hypotheses of Lemma 5.4, and the result follows. $\diamondsuit$

In the 3-dimensional case, we may apply Lemma 5.5 in place of Lemma 5.4 to get:

Proposition 5.7 : *If $\Gamma \subseteq \operatorname{Isom}\mathbf{H^3}$ is geometrically finite, then every convex fundamental domain has finitely many faces.* $\diamondsuit$

We have thus shown the equivalence of Properties (1a) and (1b) with geometrical finiteness in this dimension.

The question remains whether, in each dimension, every geometrically finite group admits some finite sided convex fundamental domain. This seems very improbable. We have already seen examples which do not admit any finite-sided Dirichlet domains, and there are many variations on this construction which ought to yield counterexamples to the more general question. However I cannot claim to have a definitive proof.

References.

[Ah1] L.V.Ahlfors, *Finitely generated Kleinian groups* : Amer. J. Math. **86** (1964) 413–429, and **87** (1965) 759.

[Ah2] L.V.Ahlfors, *Fundamental polyhedrons and limit point sets for Kleinian groups* : Proc. Nat. Acad. Sci. **55** (1966) 251–254.

[Ap1] B.N.Apanasov, *Geometrically finite groups of transformations of space* : Siberian Math. J. **23** No.6 (1982) 771–780 (English translation).

[Ap2] B.N.Apanasov, *Geometrically finite hyperbolic structures on manifolds* : Ann. Glob. Analysis & Geometry **1** No.3 (1983) 1–22 (English translation).

[BaGS] W.Ballman, M.Gromov, V.Schroeder, *Manifolds of nonpositive curvature* : Progress in Maths. 61, Birkhäuser (1985).

[Bea] A.Beardon, *The geometry of discrete groups* : Graduate Texts in Maths. 91, Springer-Verlag (1983).

[BeaM] A.Beardon, B.Maskit, *Limit sets of Kleinian groups and finite sided fundamental polyhedra* : Acta Math. **132** (1974) 1–2.

[Ber] L.Bers, *On boundaries of Teichmüller spaces and on Kleinian groups I* : Ann. Math. **91** (1970) 570–600.

[Bon] F.Bonahon, *Bouts des variétés hyperboliques de dimension 3* : Ann. Math. **124** (1986) 71–158.

[BonO] F.Bonahon, J.-P.Otal, *Variétés hyperboliques à géodésiques arbritairement courtes* : Bull. London Math. Soc. **20** (1988) 255–261.

[Bow] B.H.Bowditch, *Geometrical finiteness with variable negative curvature* : preprint, IHES (1990).

[BowM] B.H.Bowditch, G.Mess, : in preparation.

[CanEG] R.D.Canary, D.B.A.Epstein, P.Green, *Notes on notes of Thurston* : "Analytic and geometric aspects of hyperbolic space", L.M.S. Lecture Notes Series No.111, ed.

D.B.A.Epstein, Cambridge U.P. (1987) 3–92.

[CarD] Y.Carrière, F.Dal'bo, *Généralisations du première théorème de Bieberbach sur les groupes crystallographiques* : Enseign. Math. **35** (1989) 245–262.

[Cas] J.W.S.Cassels, *An embedding theorem for fields* : Bull. Austral. Math. Soc. **14** (1976) 193–198 and 479–480.

[EM] D.B.A.Epstein, A.Marden, *Convex hulls in hyperbolic space, a theorem of Sullivan, and measured pleated surfaces* : "Analytic and geometric aspects of hyperbolic space", L.M.S. Lecture Notes Series No.111, ed. D.B.A.Epstein, Cambridge U.P. (1987) 113–253.

[FMc] M.Feign, D.McCullough, *Finiteness conditions for 3-manifolds with boundary* : preprint.

[FMe] M.Feign, G.Mess, *Conjugacy classes of finite subgroups of Kleinian groups* : to appear in Amer. J. Math.

[GaR] H.Garland, M.S.Raghunathan, *Fundamental domains for lattices in (R-)rank 1 semisimple Lie groups* : Ann. Math. **92** (1970) 279–326.

[Gr] L.Greenberg, *Fundamental polyhedra for Kleinian groups* : Ann. Math. **84** (1966) 433–441.

[H] G.A.Heglund, *Fuchsian groups and transitive horocycles* : Duke Math. J. **2** (1936) 530–542.

[J] T.Jørgensen, *Compact 3-manifolds of constant negative curvature fibring over the circle* : Ann. Math. **106** (1977) 61–72.

[K] M.Kapovich, *On absence of Sullivan's cusp finiteness theorem in higher dimensions* : preprint (1990).

[KP] M.Kapovich, L.Potyagailo, *On absence of Ahlfors' finiteness theorem for Kleinian groups in dimension 3* : preprint.

[Mar] A.Marden, *The geometry of finitely generated Kleinian groups* : Ann. Math. **99** (1974) 383–462.

[Mas] B.Maskit, *On boundaries of Teichmüller spaces and on Kleinian groups II* : Ann. Math. **91** (1970) 607–639.

[Mc] D.McCullough, *Compact submanifolds of 3-manifolds with boundary* : Quart. J. Math. **37** (1986) 299–307.

[O] J.-P.Otal, *Les surfaces plissées dans les bretzels creux* : Thèse, Orsay.

[Sc] P.Scott, *Compact submanifolds of 3-manifolds* J. London Math. Soc. (2) **7** (1973) 246–250.

[Sel] A.Selberg, *On discontinuous groups in higher dimensional symmetric spaces* : "Contributions to function theory", Bombay (1960) 147–164.

[Ser] C.Series, *The geometry of Markoff mumbers* : Math. Intelligencer, Vol 7, No. 3 (1985) 20–29.

[Sul1] D.Sullivan, *A finiteness theorem for cusps* : Acta Math. **147** (1981) 289–299.

[Sul2] D.Sullivan, *Growth of positive harmonic functions and Kleinian group limit sets of zero planar measure and Hausdorf dimension two* : "Geometry Symposium, Utrecht 1980", Springer Lecture Notes 894 (1981) ed. E.Looijenga, D.Seisma, F.Taken, 127–144.

[Sul3] D.Sullivan, *Entropy, Hausdorff measures old and new, and the limit sets of geometrically finite Möbius groups* : Acta Math. **153** (1984) 259–277.

[Sul4] D.Sullivan, *Quasiconformal homeomorphisms and dynamics I — Solution to the Fatau-Julia problem for wandering domains* : Ann. Math. **122** (1985) 401–418.

[Sul5] D.Sullivan, *Quasiconformal homeomorphisms and dynamics II — Structural stability implies hyperbolicity for Kleinian groups* : Acta Math. **155** (1985) 243–260.

[SusS] P.D.Susskind, G.A.Swarup, *Limit sets of geometrically finite hyperbolic groups* : to appear in Amer. J. Math.

[Th1] W.P.Thurston, *The geometry and topology of 3-manifolds* : notes, Princeton Univ. Maths. Department (1979).

[Th2] W.P.Thurston, *The geometry and topology of 3-manifolds,* Chapter 4 : notes (1982).

[Th3] W.P.Thurston, *Hyperbolic structures on 3-manifolds II — Surface groups and 3-manifolds which fiber over the circle* : preprint.

[Tu1] P.Tukia, *On isomorphisms of geometrically finite Möbius groups* : I.H.E.S. Publ. Math. **61** (1985) 171–214.

[Tu2] P.Tukia, *The Hausdorff dimension of the limit set of geometrically finite Kleinian groups* : Acta Math. **152** (1984) 127–140.

[We] N.J.Weilenberg, *On the limit set of discrete Moebius groups with finite sided fundamental polyhedra* : unpublished.

[Wo] J.A.Wolf, *Spaces of constant curvature* : Publish or Perish (1974).