Geometrical Finiteness for Hyperbolic Groups

Brian Hayward Bowditch

Thesis submitted for the degree of Doctor of Philosophy at the University of Warwick.

Mathematics Institute,
University of Warwick,
Coventry, CV4 7AL,
United Kingdom.

June 1988

Declaration

Chapter 0, and to a large extent, Chapter 1, are of an expository nature. Chapters 2, 3, and 4 represent original work except where otherwise acknowledged. No part of this thesis has been previously submitted for any degree.

Summary

In this paper, we describe various definitions of geometrical finiteness for discrete hyperbolic groups in any dimension, and prove their equivalence. This generalises what has been worked out in two and three dimensions by Marden, Beardon, Maskit, Thurston and others. We also discuss the nature of convex fundamental domains for such groups. We begin the paper with a discussion of results related to the Margulis Lemma and Bieberbach Theorems.

Acknowledgements

This work was originally an offshoot from the M.Sc. dissertation of Dick Canary and Paul Green, which has been published, in augmented form, as [CEG]. Their paper focusses on other aspects of Thurston's notes [T1], but the start they made on geometrical finiteness was very helpful to me. My greatest debt is to David Epstein, for introducing me to the subject, and for his many suggestions and comments. I would also like to thank the S.E.R.C. for their financial support.

Geometrical Finiteness for Hyperbolic Groups

Brian H. Bowditch
University of Warwick,
Coventry, CV4 7AL, U.K.

Abstract.

Let Γ be a group acting properly discontinuously on hyperbolic space $\mathbf{H}^n$. The aim of this work is to clarify the meaning of "geometrical finiteness" for such groups. In dimension 3, with Γ acting freely, the principal definition is that Γ should possess a "finite-sided fundamental domain". Various definitions in this dimension have been worked out. Marden [Mar] shows that it is equivalent to the statement that the quotient manifold, including its ideal points, can be decomposed into a compact part and standard neighbourhoods of its cusps ("cusp cylinders" and "cusp tori"). Thurston [Th1] introduced two new definitions: that the thick part of the convex core should be compact, or that an ϵ-neighbourhood of the convex core should have finite volume. Finally, Beardon and Maskit [BeaM] say that Γ is geometrically finite if and only if its limit set consists of (what we call here) conical limit points and bounded parabolic fixed points.

We shall investigate here how these definitions generalise to arbitrary discrete actions in n dimensions. This matter has also been considered by Apanasov [Ap1,Ap2], and to some extent by Weilenberg [We].

Our central definition (GF1) will be essentially that of Marden, with appropriate definitions of cusp regions. The Beardon and Maskit description (as they point out in their paper) generalises unchanged (GF2). The use of finite-sided fundamental domains runs into problems when $n \geq 4$. The natural generalisation seems to be in terms of what we call "convex cell complexes" (GF3). For the first of Thurston's definitions, we need to clarify what we mean by the "thick part" of an orbifold. The definition chosen here (GF4) does not seem particularly natural, but it proves useful in discussing the final definition (GF5). For this, we impose the additional condition that Γ be finitely generated. We suspect that this is unnecessary, and show it to be unnecessary for manifolds, finite-volume orbifolds, or when $n \leq 3$. In the course of this discussion, we give a proof of the existence of an embedded ball in a hyperbolic n-orbifold, of uniform radius, depending on n, but not on the particular orbifold.

Finally, in Ch.4, we discuss the existence of finite-sided fundamental domains. We give (in principle) a complete description of when a Dirichlet region is finite-sided, and show that in certain special cases, all convex fundamental domains are finite-sided. We give an example (due to Apanasov) of a geometrically finite (henceforth abbreviated to GF) manifold with no finite sided Dirichlet domain.

Acknowledgements.
This work was originally an offshoot from the Warwick M.Sc. dissertation of Dick Canary and Paul Green, which has been published, in augmented form, as [CanEG]. Their paper focuses on other aspects of Thurston's notes [Th1], but the start they made on geometrical finiteness was very helpful to me. My greatest debt is to David Epstein, for introducing me to the subject, and for his many suggestions and comments. I would also like to thank the S.E.R.C. for their financial support.

0. Introduction.

0.1. Hyperbolic Space.

We begin with a general discussion of hyperbolic geometry in order to induce 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 $\mathbf{E}^n$ for euclidean n-space, and $\mathbf{H}^n$ 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 suffices where there can be no confusion. In each case, we write $\operatorname{Isom} X$ for the group of all isometries of X.

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

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

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]).

We may classify non-trivial isometries $\mathbf{H}^n$ into three types, namely elliptic, parabolic and hyperbolic 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) plane 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 l be the geodesic joining x to y. In this case, γ acts as a translation on l, and (in general) has a rotational component in the orthogonal direction. We call this case *loxodromic*, and we call l the *loxodromic axis*.

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.

0.2. 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 $M = \mathbf{H}^n/\Gamma$ which inherits a complete hyperbolic structure.

More generally, if Γ has torsion, the quotient $M = \mathbf{H}^n/\Gamma$ 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., $\Sigma = \bigcup_{\gamma \in \Gamma} (\operatorname{fix}\gamma \cap \mathbf{H}^n)/\Gamma$. 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 $\Gamma \subset \operatorname{Isom} \mathbf{H}^n$ 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, i.e.,

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

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 $\Omega = \mathbf{H}^n_I \setminus \Lambda$ 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 $\mathbf{H}^n \cup \Omega$, so we may write

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

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 study to 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, and ask what this tells us about the topology and geometry of M.

The first result is pure 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 $\mathrm{Isom}\,\mathbf{H}^n$. Geometrically, this tells us that we can restrict attention to the case where Γ acts freely on $\mathbf{H}^n$. Given this, we may as well assume also that Γ preserves orientation. This latter restriction is solely to simplify the exposition. Thus, for the rest of the introduction, unless otherwise stated, we shall be taking Γ to be a finitely-generated, discrete, torsion-free group of orientation preserving isometries of $\mathbf{H}^n$.

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 0.3. First, we describe how the 2-dimensional case is trivial from the point of view of finiteness.

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 (FIG 0.1). A *funnel* consists of a hyberbolic half-space quotiented out by a loxodromic element (FIG 0.2). 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 (FIG 0.3). 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.

0.3. 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.

Let Γ be a discrete, torsion-free, orientation-preserving subgroup of $\mathrm{Isom}\,\mathbf{H}^n$. Much of the technical complication of the subject arrises from having to deal with parabolic subgroups of Γ. Suppose that $\gamma \in \Gamma$ is parabolic with fixed point p. Let Γ_p be the stabiliser of p in Γ. In a discrete group, a parabolic and a loxodromic cannot share a common fixed point. Thus, Γ_p consists entirely of parabolics. We call p a *parabolic fixed point*, which we abbreviate to *p.f.p.* We can let p be the point ∞ in the upper half-space model. Now, Γ_p acts freely as a group of isometries on $\partial\mathbf{R}^3_+ \equiv \mathbf{E}^2$. This dimension is special in that such a group must act by translation. We see that Γ_p is isomorphic to either $\mathbf{Z}$ or $\mathbf{Z} \oplus \mathbf{Z}$. Taking B to be any horoball about p, we may form the quotient B/Γ_p. If $\Gamma_p \cong \mathbf{Z}$, then $\partial B/\Gamma_p$ is a bi-infinite euclidean cylinder, and we call B/Γ_p a $\mathbf{Z}$*-cusp* (FIG 0.4). (Note that a $\mathbf{Z}$-cusp is not quite the same as a "cusp cylinder" or "standard cusp region" which will be described later in this section. See Section 2 for details.) If $\Gamma_p \cong \mathbf{Z} \oplus \mathbf{Z}$ then $\partial B/\Gamma_p$ is a euclidean torus, and we call B/Γ_p a $\mathbf{Z} \oplus \mathbf{Z}$-cusp (FIG 0.5). 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 1.2, and 2.(GF4)) 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 $\mathrm{cusp}(M)$ 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 $M = \mathbf{H}^n/\Gamma$, we deduce immediately that Γ is finitely presented.

In our case, $M = \mathbf{H}^n/\Gamma$ 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 $\mathbf{Z} \oplus \mathbf{Z}$-cusps.

In fact (provided that Γ is not cyclic loxodromic), the $\mathbf{Z} \oplus \mathbf{Z}$-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 $\mathbf{Z}$-cusp is topologically just a product. Thus, we can assume that each $\mathbf{Z}$-cusp lies entirely within some non-cuspidal end. The effect of removing the $\mathbf{Z}$-cusps would (in general) be to subdivide each each such end into smaller pieces, on which we may see qualitatively different behaviour. It is therefore necessary to take account of these $\mathbf{Z}$-cusps before going on to consider the geometry. We can do this by applying a relative version of Scott's theorem to $M' = M \setminus \mathrm{cusp}(M)$.

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 [FM] that there are only a finite number of $\mathbf{Z}$-cusps — a result due originally to Sullivan [Sul1]. We may now take a core M'_T of M' which meets each $\mathbf{Z} \oplus \mathbf{Z}$-cusp in the bounding torus, and each $\mathbf{Z}$-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 non-cuspidal ends of M' (i.e. ends other that $\mathbf{Z} \oplus \mathbf{Z}$-cusps).

We have already said 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].)

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 $\hat{Y} = Y/\Gamma \subseteq M$, 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 $\hat{Y}$ (see for example [Th1]). We shall denote by q, the restriction of this map to M_I. Note that $q(M_I) = \partial\hat{Y}$.

It is possible for $\hat{Y}$ 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 $\hat{Y}$ is non-empty. In this case one may show that $\partial\hat{Y}$ has the structure of a complete hyperbolic surface in the induced path metric [Th1]. Moreover q is a homotopy equivalence from M_I to $\partial\hat{Y}$. 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 $\partial\hat{Y}$ also has finite conformal type and thus finite hyperbolic area. In other words, we can restate Ahlfors'

Finiteness Theorem to say that $\partial\hat{Y}$ should have finite 2-dimensional area. (In fact the discussion applies equally well if Γ has torsion, and then $\partial\hat{Y}$ becomes a finite-area orbifold.)

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

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 [G]. Jørgensen later described more concrete examples [J]. Thurston [Th2] gives a more general method of construction.

All the non-GF 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 [BoO]. 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 — GF 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 [Bo]) — 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 [Bo] 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. Otal has recently claimed the result for $\Gamma \cong \mathbb{Z} * \mathbb{Z}$.

We now restrict attention to the case where all the ends of $M' = M \backslash \mathrm{cusp}(M)$ 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 $\partial N_\eta(\hat{Y})$, which means that we can take the topological core M_T^+ to be equal to $N_\eta(\hat{Y}) \cap M' = N_\eta(\hat{Y}) \backslash \mathrm{cusp}(M)$. In other words, geometric finiteness says that $N_\eta(\hat{Y}) \backslash \mathrm{cusp}(M)$ is compact. This is more or less the definition of geometric finiteness (GF4) due to Thurston [Th1 Chapter 8] (see Section 2(GF4)). (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, $N_\eta(\hat{Y})$ meets the boundary of any $\mathbb{Z}$-cusp in a compact set. From this we see that the intersection of $N_\eta(\hat{Y})$ with any $\mathbb{Z}$-cusp has finite volume. (In fact the intersection will be contained in some r-neighbourhood of a totally geodesic 2-dimensional cusp — see FIG 0.6.) Since each $\mathbb{Z} \oplus \mathbb{Z}$-cusp has finite volume, we arrive at Thurston's second definition of geometric finiteness (GF5), namely that $N_\eta(\hat{Y})$ should have finite volume. (For the definition GF4, it is enough to insist that $\hat{Y} \backslash \mathrm{cusp}(M)$ be compact. For GF5, however, it is essential to take some uniform neighbourhood of $\hat{Y}$, 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 $\mathbb{Z} \oplus \mathbb{Z}$-cusps, whereas a cusp cylinder is an enlargement of a $\mathbb{Z}$-cusp to include a portion of M_I (FIG 0.7). 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 Section 2 (GF2). 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 $\mathbf{H}^3_I$.

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 set 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 examples showed that the limit set of a geometrically finite group is either the whole sphere $\mathbf{H}^2_I$, or else has Hausdorff dimension strictly less than 2 [Sul3].

0.4. Higher Dimensions.

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 Ahlfors' 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 geometric finiteness which we aim to elucidate in this work. All the applications of the traditional geometrical finiteness hypothesis seem to be valid for this slightly 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 GF group is either equal to $\mathbf{H}^n_I$, or else has Hausdorff dimension less than $n-1$.

1. 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 $[0,1]$. 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 $d(\gamma(x), x) < \epsilon(n)$. In symbols, $\Gamma_\epsilon(x) = \langle \gamma \in \Gamma \mid d(\gamma(x), x) < \epsilon(n) \rangle$. 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 section, 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 $\mathbf{S}^n$, $\mathbf{E}^n$ and $\mathbf{H}^n$. We shall prove that nilpotent subgroups of $\operatorname{Isom}\mathbf{S}^n$, $\operatorname{Sim}\mathbf{E}^n$, and $\operatorname{Isom}\mathbf{H}^n$ are uniformly virtually abelian. This fact seems to be well known, though I know of no explicit reference. However all the essential ingredients may be found in [Th2, Chapter 4]. 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, Chapter 4] and [Wo].

1.1. Nilpotent Implies Virtually Abelian.

Let $\mathbf{S}^n$, $\mathbf{E}^n$ 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 suffices 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.

1.1(i). Spherical Geometry.

Let

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

If we think of $\mathrm{Isom}\,\mathbf{S}^n$ acting on $\mathbf{E}^{n+1}$, this says that γ lies in $U(\mathbf{S}^n)$ if it moves each vector through an acute angle. In other words, $\langle \gamma v, v\rangle > 0$ for each non-trivial vector $v \in \mathbf{E}^{n+1}$, where $\langle\, ,\, \rangle$ is the inner product defining the metric on $\mathbf{E}^{n+1}$.

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

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

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

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

We can now prove:

Lemma 1.1.1 : *Let $\beta \in U(\mathbf{S}^n)$ and $\alpha \in \mathrm{Isom}\,\mathbf{S}^n$. If α commutes with $[\alpha, \beta]$, then α commutes with β.*

Proof : Complexifying, we imagine α and β acting on $\mathbf{C}^{n+1}$. 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 $U(\mathbf{S}^n)$, the discussion immediately prior to the lemma tells us that $\mathrm{Re}\langle \beta v, v\rangle > 0$. 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.

◇

Corollary 1.1.2 : *If $\Gamma \subset \mathrm{Isom}\,\mathbf{S}^n$ is nilpotent, then $\langle \Gamma \cap U(\mathbf{S}^n)\rangle$ is abelian.*

Proof : Let a and b lie in $\Gamma \cap U(\mathbf{S}^n)$. By a "nested chain of commutators" in a and b, we mean an expression of the form $d = [c_1, [c_2, \ldots [c_m, c_{m+1}] \ldots]]$, where each c_i is either a or b. We take d to be of maximal length, m, such that $d \neq 1$. This means that d commutes with both a and b. It follows that $[c_2, \ldots [c_m, c_{m+1}] \ldots]$ commutes with d. Applying Lemma 1.1.1, with $\alpha = [c_2, \ldots [c_m, c_{m+1}] \ldots]$ and $\beta = c_1$, we deduce that α and β commute, so that $d = 1$. We have contradicted the assumption that $m \geq 1$, and so a must commute with b.

◇

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

Corollary 1.1.3 : *Nilpotent subgroups of* $\operatorname{Isom} \mathbb{S}^n$ *are uniformly virtually abelian.*

1.1(ii). Euclidean Geometry.

To prepare for the hyperbolic case, it will be useful to consider the group $\operatorname{Sim} \mathbb{E}^n$ of euclidean similarities. Let Φ be the set of parallel classes of (semi-infinite) geodesic rays in $\mathbb{E}^n$. We shall embed Φ as the unit $(n-1)$-sphere in an inner-product space $V(\mathbb{E}^n)$, which we can imagine as euclidean space with a preferred basepoint. There is an obvious bijective correspondence between r-dimensional subspaces of $V(\mathbb{E}^n)$, and foliations of $\mathbb{E}^n$ by parallel r-planes.

The group $\operatorname{Sim} \mathbb{E}^n$ acts isometrically on Φ, so identifying Φ with $\mathbb{S}^{n-1}$ gives us a homomorphism

$$\operatorname{rot} : \operatorname{Sim} \mathbb{E}^n \longrightarrow \operatorname{Isom} \mathbb{S}^{n-1}.$$

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

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

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

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

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

Theorem 1.1.4 : *If* $\Gamma \subset \operatorname{Sim} \mathbb{E}^n$ *is nilpotent, then* $\langle \Gamma \cap U(\mathbb{E}^n) \rangle$ *is abelian.*

We shall begin with a lemma.

Lemma 1.1.5 : *Let* Γ *be an abelian subgroup of* $\operatorname{Sim} \mathbb{E}^n$. *Let* $\tau(\Gamma) = \bigcap_{\gamma \in \Gamma} \min \gamma$. *Then,* $\tau(\Gamma)$ *is a non-empty,* Γ*-invariant plane, on which* Γ *acts by translations.*

Proof : If Γ is already a translation group, then $\tau(\Gamma) = \mathbb{E}^n$, 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.

◇

In fact, our plane $\tau(\Gamma)$ 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 $V(\mathbb{E}^n)$, by taking the set of geodesic rays lying in any one leaf. Now, W_1 lies in a larger subspace W' of $V(\mathbb{E}^n)$, determined by $\tau(\Gamma)$ itself. Let W_2 be the orthogonal complement of W_1 in W', and W_3 be the orthogonal complement of W' in $V(\mathbb{E}^n)$. This gives us a canonical decomposition $V(\mathbb{E}^n) = W_1 \oplus W_2 \oplus W_3$. 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 $\mathbb{E}^n$. If $m_2 = n$, then each point of $\mathbb{E}^n$ is a fixed point of Γ, thus Γ is trivial. If $m_3 = n$, then Γ has a unique fixed point in $\mathbb{E}^n$. We are now ready for:

Proof of Theorem 1.1.4 : Let Γ be a nilpotent subgroup of $\operatorname{Sim} \mathbb{E}^n$. We shall assume that Γ is generated by elements of $U(\mathbb{E}^n)$, i.e. that $\Gamma = \langle \Gamma \cap U(\mathbb{E}^n) \rangle$. 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 $V(\mathbb{E}^n)$. Since this is canonical, it is respected by the whole group Γ. Thus Γ splits as a subgroup of $\operatorname{Sim} \mathbb{E}^{m_1} \times \operatorname{Sim} \mathbb{E}^{m_2} \times \operatorname{Sim} \mathbb{E}^{m_3}$, 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 $\mathbb{R}_+ \times \operatorname{Isom} \mathbf{S}^n$, where the first component measures the magnification, and the second, the rotational part of an element. The projection into $\operatorname{Isom}\mathbf{S}^n$ is nilpotent and generated by elements of $U(\mathbf{S}^n)$. By Corollary 1.1.2, this projection is abelian. We deduce that Γ is abelian.

◇

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

Corollary 1.1.6 : *Nilpotent subgroups of* $\operatorname{Sim}\mathbf{E}^n$ *are uniformly virtually abelian.*

1.1(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 $\mathbf{S}^n$, 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 $\mathbf{S}^n$. $\operatorname{Isom}\mathbf{H}^n$ then consists of those Möbius transformations which preserve Σ. Let γ be a Möbius transformation of $\mathbf{S}^n$, with some fixed point y. Since γ acts conformally, it induces (after scaling) an isometry of the unit tangent space $(\mathrm{T}_1\mathbf{S}^n)_y$ at y. Moreover, we may check that if z is any other fixed point of γ, then the induced isometries on $(\mathrm{T}_1\mathbf{S}^n)_y$ and $(\mathrm{T}_1\mathbf{S}^n)_z$ are conjugate. Thus, γ determines a conjugacy class in $\operatorname{Isom}\mathbf{S}^n$, which we shall call $\operatorname{rot}\gamma$. Since our subset $U(\mathbf{S}^n)$ of $\operatorname{Isom}\mathbf{S}^n$ is invariant under conjugacy, it makes sense to demand that $\operatorname{rot}\gamma$ should lie in $U(\mathbf{S}^n)$. Restricting to $\operatorname{Isom}\mathbf{H}^n$, where all Möbius transformations have fixed points, we may define

$$U(\mathbf{H}^n) = \{\gamma \in \operatorname{Isom}\mathbf{H}^n \mid \operatorname{rot}\gamma \subset U(\mathbf{S}^n)\}.$$

Theorem 1.1.7 : *If* $\Gamma \subset \operatorname{Isom}\mathbf{H}^n$ *is nilpotent, then* $(\Gamma \cap U(\mathbf{H}^n))$ *is abelian.*

We begin with two lemmas.

Lemma 1.1.8 : *If* $\Gamma \subset \operatorname{Isom}\mathbf{H}^n$ *is abelian, then* $\operatorname{fix}\Gamma$*, 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 $\operatorname{fix}\gamma$ 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.

◇

Lemma 1.1.9 : *Suppose* $\Gamma \subset \operatorname{Isom}\mathbf{H}^n$ *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 1.1.8, we distinguish three possibilities for σ. Firstly, if σ is a single point of $\mathbf{H}^n_?$, this point is fixed by Γ, and we are done. Secondly, if σ is a proper subspace of $\mathbf{H}^n_?$, 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 $\mathbf{H}^n_?$. 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 $\operatorname{fix} Z(\Gamma)$.

$\diamond$

Proof of Theorem 1.1.7 : By Lemma 1.1.9, $\langle \Gamma \cap U(\mathbf{H}^n)\rangle$ 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 1.1.2, and Theorem 1.1.4.

$\diamond$

For completeness, we state:

Corollary 1.1.10 : *Nilpotent subgroups of* $\operatorname{Isom} \mathbf{H}^n$ *are uniformly virtually abelian.*

Proof : If $\Gamma \subseteq \operatorname{Isom} \mathbf{H}^n$ is nilpotent, we need that $[\Gamma : \langle \Gamma \cap U(\mathbf{H}^n)\rangle]$ is uniformly bounded. But by Lemma 1.1.8, Γ has a fixed point in $\mathbf{H}^n_C$, so the result follows from the spherical and euclidean (similarity) cases.

$\diamond$

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

1.2. 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 1.2.1 : *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 $\epsilon > 0$, and the elements of $O(G)$ have norms bounded above by some number k. If m is any integer greater than $\log_2(k/\epsilon)$, 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.

$\diamond$

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 Lie group G, we write X^r for those $g \in G$ expressible as words of length 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 1.2.2 : *Let* G *be a Lie 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* $(K_i)^4 \subseteq K_1$ *for each* i. *Then, there exists some* $N \in \mathbf{N}$ *such that for any discrete group* $\Gamma \leq G$, $|\Gamma_{K_N} : \langle \Gamma_{K_N} \cap W\rangle| < N$.

Proof : Let V be a neighbourhood of 1 with $V^{-1}V \subset W$. 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 $h_j \in (K_N)^N \subset K_1$. 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, $\alpha\beta\alpha^{-1} \in V^{-1}V \subset W$, 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, $\Gamma_N h = \Gamma_N a_i$.

◇

We again consider the three geometries in turn.

1.2(i). Spherical Geometry.

We write $U_0(\mathbb{S}^n)$ for $O(\operatorname{Isom} \mathbb{S}^n)$, the neighbourhood of the identity defined at the beginning of Section 1.2. Since this set may be chosen to be arbitrarily small, we may suppose that $U_0(\mathbb{S}^n) \subset U(\mathbb{S}^n)$. We may also suppose that $U_0(\mathbb{S}^n)$ is conjugacy invariant. Now if Γ is a discrete subgroup of $\operatorname{Isom} \mathbb{S}^n$, then $\langle \Gamma \cap U_0(\mathbb{S}^n) \rangle$ is nilpotent by Lemma 1.2.1, and thus abelian by Corollary 1.1.2. It is easily checked that $\langle \Gamma \cap U_0(\mathbb{S}^n) \rangle$ has a finite index in Γ, which is bounded as Γ varies. Thus we have:

Lemma 1.2.3 (Jordan Lemma) : *Discrete subgroups of* $\operatorname{Isom} \mathbb{S}^n$ *are uniformly virtually abelian.*

1.2(ii). Euclidean Geometry.

We can assume that $O(\operatorname{Isom} \mathbb{E}^n)$ has the form

$$O(\operatorname{Isom} \mathbb{E}^n) = \{\gamma \in \operatorname{Isom} \mathbb{E}^n \mid d(\gamma a, a) < \epsilon \text{ and } \operatorname{rot} \gamma \in U_1\}$$

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

$$U_0(\mathbb{E}^n) = \{\Gamma \in \operatorname{Isom} \mathbb{E}^n \mid \operatorname{rot} \gamma \in U_0(\mathbb{S}^n)\}.$$

Proposition 1.2.4 : *Suppose that* Γ *is a discrete subgroup of* $\operatorname{Isom} \mathbb{E}^n$*; then* $\langle \Gamma \cap U_0(\mathbb{E}^n) \rangle$ *is abelian.*

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

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

which is nilpotent (Lemma 1.2.1) and hence abelian (Theorem 1.1.4). Thus, Γ_r is abelian for all r, and so $\langle \Gamma \cap U_0(\mathbb{E}^n) \rangle = \bigcup_r \Gamma_r$ is abelian.

◇

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

Theorem 1.2.5 (Bieberbach) : *Discrete subgroups of* $\operatorname{Isom} \mathbb{E}^n$ *are uniformly virtually abelian.*

◇

Note that since $U_0(\mathbf{S}^n)$ is conjugacy invariant in $\mathrm{Isom}\,\mathbf{S}^n$, 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:

Proposition 1.2.6 : *Suppose Γ acts properly discontinuously on $\mathbf{E}^n$. Then, there is a plane $\mu \subseteq \mathbf{E}^n$, preserved by Γ, with μ/Γ compact. Moreover, any two such subspaces are parallel, and the action of Γ commutes with the perpendicular translation between them.*

Proof : If Γ preserves each of two planes τ_1 and τ_2, then it preserves $\tau_1 \cap \tau_2$. It therefore makes sense to speak of a Γ-invariant plane $\mu \neq \emptyset$ being minimal.

Let μ_1 and μ_2 be two such minimal planes. Let $\lambda(\mu_i,\mu_j) = \{x \in \mu_i \mid d(x,\mu_j) = d(\mu_i,\mu_j)\} \subseteq \mu_i$. 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 planes 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 1.1(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. Γ preserves τ', so, by minimality, $\tau' = \mathbf{E}^n = \tau$. Hence, $\mathbf{E}^n/\Gamma$ is compact.
◇

As in the earlier discussion of the abelian case (Section 1.1(ii)), it is easily seen that the set of minimal planes in $\mathbf{E}^n$ form a foliation of a larger, canonical subspace.

1.2(iii). Hyperbolic Geometry.

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

$$I_\epsilon(x) = \{\gamma \in \mathrm{Isom}\,\mathbf{H}^n \mid d(\gamma x, x) < \epsilon\}.$$

Let d_1 be any Riemannian metric on the unit tangent bundle $\mathrm{T}_1\mathbf{H}^n$ of $\mathbf{H}^n$, invariant under the action of $\mathrm{Isom}\,\mathbf{H}^n$. Given $x \in \mathbf{H}^n$, we write

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

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

$$\Gamma_\epsilon(x) = \langle \Gamma \cap I_\epsilon(x)\rangle$$

and

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

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

Proposition 1.2.7 : *If Γ is a discrete subgroup of $\mathrm{Isom}\,\mathbf{H}^n$, then $\Gamma'_{\epsilon_1}(x) = \langle \Gamma \cap I'_{\epsilon_1}(x)\rangle$ 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 $I'_{\epsilon_1}(x)$ to be the set W of Lemma 1.2.2, and the sets K_r to be $I_{1/r}(x)$. The lemma now tells us that, for some $N > 0$,

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

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 $N \le \nu(n)$, 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 1.2.8 (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* Isom$\mathbf{H}^n$, *and $x \in \mathbf{H}^n$, then $\Gamma_{\epsilon(n)}(x)$ has an abelian subgroup ($= \Gamma'_{\epsilon(n)}(x) \cap \Gamma_{\epsilon(n)}(x)$) of index at most $\nu(n)$.*

Note that if $0 < \epsilon \le \epsilon(n)$, 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.

2. Five Definitions of Geometrical Finiteness.

In this section, we shall give details of the five definitions of geometrical finiteness that we intend to use. First, we clarify a few points of terminology, and notation.

By a (discrete) *parabolic group* of hyperbolic isometries, G, we mean a discrete group which fixes a unique point in $\mathbf{H}^n_+$, i.e. $\bigcap_{\gamma \in G} \operatorname{fix}\gamma = \{p\}$, where $p \in \mathbf{H}^n_+$. In this case, G must contain at least one parabolic with fixed point p. Since no loxodomic can share a fixed point with a parabolic in any discrete group, we see that G consists entirely of parabolics and elliptics. We may represent $\mathbf{H}^n$ using the upper half-space model $\mathbf{R}^n_+$ with $p = \infty$. It then follows that G acts as a group of euclidean isometries of $\partial\mathbf{R}^n_+ = \mathbf{E}^{n-1}$. From Proposition 1.2.6, we know that G preserves some plane in $\partial\mathbf{R}^n_+$ whose quotient by G is compact. Moreover, any two such planes are parallel. We shall write σ_I for some choice of such plane, and write σ for the vertical euclidean half-space with $\sigma \cap \partial\mathbf{R}^n_+ = \sigma_I$. Thus, σ is the hyperbolic subspace spanned by σ_I and p. Now, if Γ is any discrete subgroup of $\mathbf{H}^n_+$, and $\gamma \in \Gamma$ is any parabolic, with fixed point p, then the stabiliser of p is Γ will be a parabolic subgroup. We call p a *parabolic fixed point*, which we abbreviate to p.f.p.

GF1 Let Γ be a discrete group. In Section 0.2, we defined M_C as the quotient, by Γ, of hyperbolic space together with the discontinuity domain, that is $M_C = (\mathbf{H}^n \cup \Omega)/\Gamma$. Thus $M_C = M \cup M_I$, where $M = \mathbf{H}^n/\Gamma$ is a complete hyperbolic manifold, and $M_I = \Omega/\Gamma$ consists of "ideal points" of M. Where there is more than one group in question, we shall be specific by writing $M(\Gamma)$, $M_I(\Gamma)$ and $M_C(\Gamma)$.

Suppose that Γ and Γ' are two discrete groups. Suppose e and e' are (topological) ends of $M_C(\Gamma)$ and $M'_C(\Gamma')$ respectively. We say that e and e' are *equivalent* if they admit isometric neighbouroods. (Here, we use the term "isometric" loosely, in that the orbifolds in question may contain ideal points. In saying that two such orbifolds are isometric, we mean that there is an isometry of the metric parts which extends to a homeomorphism on the ideal points.) Note that if Γ' is a parabolic group, then $M_C(\Gamma')$ has precisely one end (see the discussion below).

Definition 1 : *Γ is GF1 if $M_C(\Gamma)$ has finitely many ends, and each such end is equivalent to the end of the quotient of a parabolic group.*

It will be convenient to give this definition a mor · concrete formulation in terms of the structure of parabolic groups. Let Γ_p be such a group, fixing $p = \infty$ in the upper half-space model, $\mathbf{R}^n_+$. Let σ be a Γ-invariant vertical plane with σ_I/Γ_p compact. Let $C(r) = \{x \in \mathbf{R}^n_+ \cup \partial\mathbf{R}^n_+ \mid d_{euc}(x, \sigma_I) \ge r\}$. Thus, $\mathbf{R}^n_+ \cup \partial\mathbf{R}^n_+ \setminus C(r)$ is an open uniform r neighbourhood of σ_I in the euclidean metric on $\mathbf{R}^n_+ \cup \partial\mathbf{R}^n_+$. Since Γ_p preserves the euclidean metric, the constuction is Γ_p-equivariant, so we may form the quotient $\hat{C}(r) = C(r)/\Gamma_p$. Since σ_I is compact, $(\mathbf{R}^n_+ \cup \partial\mathbf{R}^n_+ \setminus C(r))/\Gamma_p = M_C(\Gamma_p) \setminus \hat{C}(r)$ is relatively compact in $M_C(\Gamma_p)$. If we take a sequence (r_n) tending to infinity, then the sets $M_C(\Gamma_p) \setminus \hat{C}(r_n)$ give a relatively compact exhaustion of $M_C(\Gamma_p)$. Since each $\hat{C}(r_n)$ is connected, we see that $M_C(\Gamma_p)$ has precisely one end. The sets $\hat{C}(r_n)$ give a neighbourhood base for the end. Given any r, we call $C(r)$ a *standard parabolic region* (FIG 2.1), and the quotient $\hat{C}(r)$, a *standard cusp* (FIG 2.2).

We may now say that Γ is GF1 if and only if we can write $M_O = \hat{N} \cup (\bigcup \hat{\mathcal{C}})$, with $\hat{N}$ compact, $\hat{\mathcal{C}}$ finite, and each $\hat{C} \in \hat{\mathcal{C}}$ (isometric to) a standard cusp (FIG 2.3). We shall see (Chapter 3, GF4 $\Rightarrow$ GF2) that the cusps $\hat{C}$ are in bijective correspondence to the orbits of parabolic fixed points of Γ.

We write N for the lift of $\hat{N}$ to $\mathbf{H}^n_C$.

GF2

Let p be a parabolic fixed point (p.f.p.) of the discrete group Γ. Let $\Gamma_p = \mathrm{stab}_\Gamma p$ — the stabiliser of p. We say that P is a *bounded p.f.p. (b.p.f.p.)* if $(\Lambda \setminus \{p\})/\Gamma_p$ is compact. Let $p = \infty$ in the upper half-space model, and let σ_I be a minimal Γ_p invariant plane. Then it is not difficult to see that $(\Lambda \setminus \{\infty\})/\Gamma_p$ is compact if and only if $d_{euc}(y, \sigma_I)$ is bounded as y varies in $\Lambda \setminus \{\infty\}$. In other words, p is a b.p.f.p. if and only if $\Lambda(\Gamma) \setminus \{\infty\} \subseteq Q_r = \{x \in \partial \mathbf{R}^n_+ \mid d_{euc}(x, \sigma_I) \leq r\}$ for some r (FIG 2.4).

Let $y \in \Lambda(\Gamma)$. We say that y is a *conical limit point (c.l.p.)* if for some (and hence every) geodesic ray l joining some point of $\mathbf{H}^n$ to y, the orbit Γl of l accumulates somewhere in $\mathbf{H}^n$, i.e., $|\{\gamma \in \Gamma | \gamma l \cap K \neq \emptyset\}| = \infty$ for some compact $K \subset \mathbf{H}^n$. (The term derives from an alternative desription, namely that there should exist a sequence $\gamma_n \in \Gamma$, and a point $x \in \mathbf{H}^n$, with $\gamma_n x$ tending to y, while remaining a bounded distance from some geodesic ray — see FIG 2.5.)

Definition 2 : *Γ is GF2 if every point of Λ is either a c.l.p. or a b.p.f.p.*

We shall see in Ch.3 that these two classes are, in any discrete group, mutually exclusive. In fact, it is shown in [SuaS] that, in any discrete group, no p.f.p. can also be a c.l.p. It will follow from the discussion in Chapter 3 GF4 $\Rightarrow$ GF2 that in the special case of a geometrically finite group, any p.f.p. is necesserily a b.p.f.p., and thus not a c.l.p.

Beardon and Maskit [BeaM] give several equivalent definitions of c.l.p., including one that makes sense in $\mathbf{H}^n_T$. This gives GF2 as a definition of geometric finiteness intrinsic to the action of Γ on $\mathbf{H}^n_T$.

Finally, we remark that, for a GF group, the convergence of orbits under Γ to c.l.p.s can be chosen to be "uniform". For us, this means that the set K, in the definition, can be chosen independently of the point y and the ray l. Together with a certain convergence property for the radii of isometric spheres, this implies that the limit set of a GF group has either zero or full spherical Lebesgue measure (see [BeaM,Ap1]). A more geometric proof of this fact is based on the definition GF5 (see [Th1]).

GF3

Let Γ be a discrete group. We have said that the hypothesis that Γ should possess a finite-sided fundamental polyhedron is more restrictive than we would like in dimension 4 onwards. In Section 4, we give an example to illustrate this point (at least for the case of Dirichlet domains). However it is possible to modify the criterion so that it works in all dimensions. The idea is that we should allow ourselves more than one polyhedron to constitute a fundamental domain for Γ. The definition is most clearly expressed in terms of what we shall call "convex cell complexes". A convex cell complex is cell complex in which all the cells are convex, and hence necessarily polyhedra. It need not quite be a CW-complex since we only attach cells along their relative boundaries in hyperbolic space. Thus a finite complex is complete, but not in general compact. We give a more formal description below.

Let A be a subset of $\mathbf{E}^n$. We call A an *open (convex) cell* if any two distinct points of A lie in the interior of some geodesic segment contained entirely in A. Note that by demanding that the two points be distinct, we allow any one-point set as an open cell. We see that the property of being an open cell is closed under taking finite intersections.

Definition : *A collection $\mathcal{A}$ of subsets of $\mathbf{E}^n$ is "convex cell complex" if:*

(1) each element of $\mathcal{A}$ is an open cell,

(2) the sets of $\mathcal{A}$ are all disjoint,

(3) the collection $\mathcal{A}$ is locally finite,

(4) $\bigcup \mathcal{A} = \mathbf{E}^n$,

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

Let $\mathcal{A}$ be such a cell complex, and let $B \in \mathcal{A}$. Suppose that x and y are two points of B, and suppose that $x \in \bar{A}$ for some cell $A \in \mathcal{A}$. Then, from (5), we see that $y \in \bar{A}$. Thus, $\{A \in \mathcal{A} \mid x \in \bar{A}\}$ is independent of the choice of $x \in B$. In particular, from the local finiteness of $\mathcal{A}$, (2), we see that any cell of $\mathcal{A}$ meets the closures of only finitely many other cells.

Now, given two cell complexes $\mathcal{A}$ and $\mathcal{B}$, we call $\mathcal{B}$ a *subdivision* of $\mathcal{A}$ if each $B \in \mathcal{B}$ is a subset of some $A \in \mathcal{A}$. Any two cell complexes $\mathcal{A}_1$ and $\mathcal{A}_2$, have a natural common subdivision, namely $\langle \mathcal{A}_1, \mathcal{A}_2 \rangle = \{A_1 \cap A_2 \mid A_1 \in \mathcal{A}_1, A_2 \in \mathcal{A}_2\}$. In fact $\langle \mathcal{A}_1, \mathcal{A}_2 \rangle$ is minimal with respect to subdivision — if $\mathcal{B}$ is a subdivision of both $\mathcal{A}_1$ and $\mathcal{A}_2$, then it is a subdivision of $\langle \mathcal{A}_1, \mathcal{A}_2 \rangle$. We also remark that intersecting a cell complex with an affine subspace of $\mathbf{E}^n$ gives a cell complex in that subspace.

All the above properties are easily verified from the definition above. However, to make the analogy with CW-complexes more explicit, we offer a slightly different description of convex cell complexes as follows.

Suppose that $A \subseteq \mathbf{E}^n$ is a convex set. We write $\langle A \rangle$ for the affine span of A, i.e. the smallest subspace of $\mathbf{E}^n$ containing A. We may define the dimension of A, $\dim A$, to be equal to the dimension of $\langle A \rangle$. We also define $\operatorname{ri} A$ and $\operatorname{rb} A$ to be, respectively, the relative interior and the relative boundary of A in $\langle A \rangle$. Note that $\operatorname{ri} A$ is always nonempty, provided A is nonempty. Moreover, it is not difficult to see that A is an open cell if and only if $\operatorname{ri} A = A$. By an *open* i*-cell*, we mean an open cell of dimension i.

Let $\mathcal{A}$ be a collection of convex cells of $\mathbf{E}^n$. We write $\mathcal{A}^i$ for set of all i-cells in $\mathcal{A}$. The r*-skeleton*, K^r, of $\mathcal{A}$ is the union of all i-cells with $i \leq r$, i.e. $K^r = \bigcup_1^r (\bigcup \mathcal{A}^i)$. We now claim that, if we know that $\mathcal{A}$ satisfies properties (1)–(4), then property (5) is equivalent to the following:

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

To see (5′) $\Rightarrow$ (5), it is enough to not that if one open cell lies in the relative boundary of another, then its dimension must be strictly less. To see (5) $\Rightarrow$ (5′) is a little more complicated. Suppose that $\mathcal{A}$ satisfies (1)–(4) and (5), and let $A \in \mathcal{A}^r$. Using (for example) a measure-theoretic argument, we see that $(\bigcup \mathcal{A}^{r-1}) \cap \operatorname{rb} A$ is a dense subset of ∂A. Suppose that $B \in \mathcal{A}^{r-1}$ intersects $\operatorname{rb} A$. If B is not a subset of $\operatorname{rb} A$, then there is some point x in the relative boundary of $\bar{A} \cap B$ in B. By considering a neighbourhood of x in $\operatorname{rb} A$, we see that $x \in \operatorname{rb} C$ for some $C \in \mathcal{A}^{r-1}$, different from B. But by (5′), we have that $\operatorname{rb} C \subseteq K^{r-2}$. This contradiction tells us that $B \subseteq \operatorname{rb} A$. In other words, $\operatorname{rb} A$ is a union of closures of $(r-1)$-cells. We have deduced property (5) in the case where $\dim A - \dim B = 1$. We now use induction over $\dim A - \dim B$. Let $D \in \mathcal{A}$ be an i-cell intersecting $\operatorname{rb} A$. From (5′), we know that $i \leq r-1$. If $i = r-1$, we are done. If $i < r-1$, then D intersects $\operatorname{rb} E$ for some $(r-1)$-cell $E \subseteq \operatorname{rb} A$. By the induction hypothesis, $D \subseteq \operatorname{rb} E$. We see that $D \subseteq \operatorname{rb} A$. Thus we have shown the equivalence of the two descriptions of convex cell complexes.

Now, let $\mathcal{A}$ be a convex cell complex, and let $\mathcal{U} = \mathcal{A}^n$ be the collection of top-dimensional cells. We claim that $\mathcal{U}$ is characterised by the following properties.

(a) $\mathcal{U}$ is a collection of open convex subsets of $\mathbf{E}^n$.

(b) $\mathcal{U}$ is locally finite.

(c) The closures of all the sets in $\mathcal{U}$ cover $\mathbf{E}^n$.

(d) The elements of $\mathcal{U}$ are disjoint.

In fact, if we are given such a collection, we may recover a convex cell complex as follows. Given $x \in \mathbf{E}^n$, we write $P(x) = \{U \in \mathcal{U} \mid x \in \bar{U}\}$. Let $A(x) = \{y \in \mathbf{E}^n \mid P(y) = P(x)\}$, and let $\mathcal{A}(\mathcal{U}) = \{A(x) \mid x \in \mathbf{E}^n\}$. Then, we claim that $\mathcal{A}(\mathcal{U})$ is a convex cell complex with $\mathcal{A}(\mathcal{U})^n = \mathcal{U}$. The only property that is not immediate is property (1), namely that each $A(x)$ is an open cell. For this, it is enough to check that if y and z are distinct points in $\mathbf{E}^n$ with $P(y) = P(y)$, then y and z lie in the interior of a line segment l, with $P(u) = P(y)$ for all $u \in l$. In fact, if T is the translation of $\mathbf{E}^n$ sending y to z, then it is not difficult to see that, for some $\epsilon > 0$, we will have $N_\epsilon(z) \cap U = T(N_\epsilon(y) \cap U)$, for any $U \in \mathcal{U}$. (N_ϵ denotes a uniform ϵ-neighbourhood.) This means that we can find such a line segment with a neighbourhood on which the sets of $\mathcal{U}$ are a cartesian product. We deduce that $\mathcal{A}(\mathcal{U})$ is a cell complex, which has $\mathcal{U}$ as its collection of top-dimensional cells. Moreover, $\mathcal{A}(\mathcal{U})$ is minimal with respect to subdivision. Thus, our original $\mathcal{A}$ is a subdivision of $\mathcal{A}(\mathcal{U})$.

We state a refinement of the above result.

Proposition 2.1 : *Let $\mathcal{A}$ be a convex cell complex. Let $\mathcal{U}$ be a locally finite collection of open r-cells, 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 $i = r+1, \ldots, n$.*

The proposition may be proven by similar arguments to those given above. In fact, our discussion dealt with the special case when $\mathcal{A} = \{\mathbf{E}^n\}$ and $r = n$. We shall not give details here, since it is not central to the paper, but used only to relate the notion of cell complexes with fundamental domains.

One natural way in which convex cell complexes arrise is as follows. Let X be a discrete subset of $\mathbf{E}^n$. Given $x \in X$, we define $D_X(x)$ to be the set of points in $\mathbf{E}^n$, nearer to x than to any other point of X, i.e.

$$D_X(x) = \{y \in \mathbf{E}^n \mid d(y,x) < d(y,z) \text{ for all } z \in X \setminus \{x\}\}.$$

It is easily checked that the collection of sets $\{D_X(x) \mid x \in X\}$ satisfies all the conditions of being the set of top-dimensional cells of some convex cell complex, namely properties (a)–(d) listed above. Let $\mathcal{A}_X$ be the cell complex, minimal with respect to subdivision with $\mathcal{A}_X^n$.

Another description of $\mathcal{A}_X$ as follows. Given any finite subset Y of X, we write $D_X(Y)$ to be the set of points y for which the minimal value of $d(z,y)$ with $z \in X$ is attained equally at each point $z \in Y$. Then $\mathcal{A}$ is the set of all $D_X(Y)$ as Y ranges over all finite subsets of X.

Let $\mathcal{A}$ a convex cell complex, with $A \in \mathcal{A}$. We call $B \in \mathcal{A}$ a *face* of A if $B \subset \bar{A}$. We write $\mathcal{F}(A)$ for the set of all faces of A. We call a subset $\mathcal{B}$ of $\mathcal{A}$ a *full subcomplex* if a face of any element of $\mathcal{B}$ also lies in $\mathcal{B}$. In this case, we write $|\mathcal{B}|$ for the union of all the cells of $\mathcal{B}$. We see that $|\mathcal{B}|$ is a closed subset of $\mathbf{E}^n$.

We can make sense of the notion of convex cell complex on certain closed subsets of $\mathbf{E}^n$ by replacing property (4) in the definition by the hypothesis that $\bigcup \mathcal{B} = F$. Examples are thus full subcomplexes of a given complex $\mathcal{A}$.

Suppose that $A \in \mathcal{A}^r$ for some complex $\mathcal{A}$. Let $\langle A \rangle$ be the affine span of A. It is not difficult to see that we may represent A as an intersection of half-spaces, in $\langle A \rangle$, determined by the codimension-one faces of A. Thus each cell of a complex is necessarily the relative interior of a polyhedron, according to the following definition.

Definition : *An "r-dimensional polyhedron" is a (countable) intersection of closed half-spaces of* $\mathbf{E}^r$, $P = \bigcap_{\alpha \in A} H_\alpha$, *where the sets* $P \cap \partial H_\alpha$ *are locally finite. We insist that* P *have non-empty interior in* $\mathbf{E}^r$.

Given such a polyhedron, we may reconstruct a convex cell complex $\mathcal{S}(P)$ on P by taking, as lower-dimensional faces, the relative interiors of the intersections with P of the supporting hyperplanes. We call such faces the *sides* of P. If P is obtained as the closure of a cell in a convex cell complex, then $\mathcal{F}(P)$ is a subdivision of $\mathcal{S}(P)$.

As an example, consider the tesselation of $\mathbf{E}^n$ by bi-infinite square prisms (each isometric to $[0,1]^2 \times \mathbf{R}$), obtained by stacking the tiles in horizontal layers (FIG 2.6). First, the tiles are laid parallelly north-south, then east-west, and so on alternately. Each tile has infinitely many codimension-1 faces (in the associated cell complex), but only finitely many codimension-1 sides (in fact, four).

So far, we have talked only about cell complexes in euclidean space. However, all the above discussion is valid with $\mathbf{E}^n$ replaced by $\mathbf{H}^n$. To see this, we note that in the Klein model for hyperbolic space, hyperbolic and euclidean convexity coincide.

Now, let Γ be a discrete group acting on $\mathbf{H}^n$. Let X be a discrete Γ-invariant set, and let $\mathcal{A}_X$ be the complex derived from X, as described above. The complex $\mathcal{A}_X$ has the following properties.

(i) It is Γ-invariant.

(ii) The setwise stabiliser of any cell is finite.

Suppose in particular, that $X = \bigcup_{i=1}^{k} \Gamma a_i$, where the orbits Γa_i are disjoint, and each point a_i has trivial stabiliser in Γ. Then, we call the top-dimensional cells of $\mathcal{A}_X$ *(generalised) Dirichlet domains*. We write $\mathcal{A}_\Gamma(\underline{a})$ for the complex $\mathcal{A}_X$, and write $D_i(\underline{a})$ for $D_X(a_i)$ — the Dirichlet domain about a_i. Here, $\underline{a}$ represents the finite set $\{a_1, \ldots, a_k\}$.

More generally, we say that a convex complex $\mathcal{A}$ is *associated* to Γ if it satisfies the two properties (i)–(ii) above. If we are given such a complex $\mathcal{A}$, we may find a Γ-invariant subdivision $\mathcal{A}_0$ of $\mathcal{A}$ with the property that if any $\gamma \in \Gamma$ preserves, setwise, a cell $A \in \mathcal{A}_0$, then it fixes A pointwise. This means that $\mathcal{A}_0$ projects to a cell complex in $M = \mathbf{H}^n/\Gamma$.

One may obtain $\mathcal{A}_0$ as follows. For $A \in \mathcal{A}^n$, let $\mathrm{stab}_\Gamma A$ be the (finite) stabiliser of A in Γ, and let $\mathcal{U}(A)$ be the set of intersections of A with a collection of Dirichlet domains for $\mathrm{stab}_\Gamma A$. Given any $\gamma \in \Gamma \setminus \mathrm{stab}_\Gamma A$, we define $\mathcal{U}(\gamma A) = \gamma \mathcal{U}(A)$. Performing this construction for each orbit of top-dimensional cell gives us a Γ-invariant collection of convex sets $\mathcal{U} = \bigcup_{A \in \mathcal{A}^n} \mathcal{U}(A)$, satisfying the hypotheses Proposition 2.1. This gives us a subdivision $\mathcal{A}_1$ of $\mathcal{A}$. Let $\langle \mathcal{A}, \mathcal{A}_1 \rangle$ be the common subdivision of $\mathcal{A}$ and $\mathcal{A}_1$ as defined above. We can now cut up all the codimension-1 cells in $(\mathcal{A}, \mathcal{A}_1)$ in a similar way. Applying Proposition 2.1 again gives us a further subdivision $\mathcal{A}_2$. Continuing this process inductively gives us, after n steps our required subdivision $\mathcal{A}_0$. Note that each cell of $\mathcal{A}_0$ is divided into only finitely many pieces in $\mathcal{A}_0$.

We can relate the complex $\mathcal{A}_0$ to convex fundamental domains. Suppose, for a moment, that Γ is orientation-preserving, so that the singular set lies in $K^{n-2}(\mathcal{A}_0)$. Let $\mathcal{B}^n$ be a set of orbit representatives of $\mathcal{A}_0^n$ under Γ. Let $\mathcal{B}^{n-1} = (\mathcal{F}(\mathcal{B}^n))^{n-1}$, the set of codimension-1 faces. Each face in $\mathcal{B}^{n-1}$ meets either one or two cells in $\mathcal{B}^n$. Those that meet only one are paired under Γ. The set of *face-pairing* isometries generate Γ. If $\mathcal{B}^n$ has only one element, say B, then B is called a (convex) fundamental polyhedron. In defining geometrical finiteness, it has been usual to demand that the codimension-1 sides and faces of B coincide (the axiom of side-pairings - see [BM]). However, from our point of view, this restriction does not seem particularly natural, and we shall not use it. Note that, if Γ is not orientation-preserving, we may have to allow for reflections in codimension-1 faces.

Definition 3 : *Γ is GF3 if there exists a convex cell complex $\mathcal{A}$ on $\mathbf{H}^n$, preserved under Γ, with $\{\gamma \in \Gamma \mid \gamma A = A\}$ finite for all $A \in \mathcal{A}$, and with $\mathcal{A}/\Gamma$ (the set of orbits under Γ) finite.*

In such a case, if we subdivide $\mathcal{A}$ to $\mathcal{A}_0$ as described above, then $\mathcal{A}_0/\Gamma$ will also be finite. Thus, $\mathcal{A}_0$ projects to a finite complex in M. We may thus rephrase GF3 by saying that M may represented by a finite complex in which cell is isometric to an open convex set in $\mathbf{H}^n$. As stated at the beginning, each cell is attached only along its relative boundary in $\mathbf{H}^n$.

GF4

Let Γ be a discrete group of isometries of $\mathbf{H}^n$. We define $\mathrm{free}(\Gamma) = \{\gamma \in \Gamma \mid \mathrm{fix}(\gamma) \cap \mathbf{H}^n = \emptyset\}$ to be the subset of elements acting freely, i.e. without fixed points in $\mathbf{H}^n$. Let $0 < \epsilon < \epsilon(n)$, where $\epsilon(n)$ is the Margulis constant. The set $\mathrm{thin}_\epsilon^\sim(M) = \{x \mid d(x, \gamma x) \leq \epsilon$ for some $\gamma \in \mathrm{free}(\Gamma)\}$ projects to what we shall call the *thin part* of the quotient orbifold M, denoted by $\mathrm{thin}_\epsilon(M)$.

We claim that the connected components of $\mathrm{thin}_\epsilon(M)$ have the form $\mathrm{thin}_\epsilon(M(G))$, where G is either a parabolic group, i.e. fixes a unique point in $\mathbf{H}^n_I$, or else is, what we shall call here, a "loxodromic group", i.e. it preserves, setwise, a geodesic, whose quotient under G is compact. This is well known in the case where Γ is torsion-free, and we can use essentially the same reasoning for our more general situation. For completeness, we give the argument below.

Let T be a component of $\mathrm{thin}_\epsilon^\sim(M)$, and let G be the setwise stabiliser of $\tilde{T}$ in Γ, so that $T = \tilde{T}/G$ is a component of $\mathrm{thin}_\epsilon(M)$. We first show that $\tilde{T} \subseteq \mathrm{thin}_\epsilon^\sim(M(G))$. We then show that G is either parabolic or loxodromic, from which it follows that $\mathrm{thin}_\epsilon^\sim(M(G))$ is connected, and thus equal to $\tilde{T}$. We can then deduce that $T = \mathrm{thin}_\epsilon(M(G))$.

For the first part, consider $x \in \tilde{T}$. There is some $\gamma \in \mathrm{free}\,\Gamma$, with $d(x, \gamma x) < \epsilon$. Let l be the geodesic segment joining x to γx. For any $y \in l$, we have $d(y, \gamma y) \leq d(x, \gamma x) < \epsilon$. Thus $l \subset \mathrm{thin}_\epsilon^\sim(M)$, and so $\gamma x \in \tilde{T}$. Now any element of Γ must either preserve $\tilde{T}$ setwise, or map it onto a disjoint component. We deduce that $\gamma \in G$, and so $x \in \mathrm{thin}_\epsilon^\sim(M(G))$.

For the second part, we fix $x \in \tilde{T}$. Now, $\Gamma_\epsilon(x) = \langle \gamma \in \Gamma \mid d(x, \gamma x) < \epsilon \rangle$ contains an element of infinite order. From the discussion of the Margulis Lemma in Section 1, we may deduce that $\Gamma_\epsilon(x)$ is either a parabolic or a loxodromic group. Now, let $\gamma \in G$. We join x to γx by a path $\lambda \subset \tilde{T}$. Consider the groups $\Gamma_\epsilon(\lambda(t))$, as the parameter t varies. Suppose at some time t_0, $\Gamma_\epsilon(\lambda(t))$ changes from one subgroup G_1 of Γ to another, $G_2 \subset \Gamma$. Choose any ϵ' lying strictly between ϵ and the Margulis constant $\epsilon(n)$. Then G_1 and G_2 are both subgroups of $\Gamma_{\epsilon'}(\lambda(t_0))$. Again from the Margulis lemma, we see that G_1 and G_2 are either both parabolic with the same fixed point, or loxodromic with the same axis. Thus, if $\Gamma_\epsilon(x)$ is parabolic with fixed point p, then $\Gamma_\epsilon(\gamma x)$ is also parabolic with fixed point p. But, $\Gamma_\epsilon(\gamma x) = \gamma \Gamma_\epsilon(x) \gamma^{-1}$, and so $\gamma p = p$. It is clear from the first part of the proof that G contains elements of infinite order, and so we see that G is a parabolic group. In this case, note that if $y \in \mathrm{thin}_\epsilon^\sim(M(G))$, then the geodesic joining y to p lies within $\mathrm{thin}_\epsilon^\sim(M(G))$. From this, it is easy to deduce that $\mathrm{thin}_\epsilon^\sim(M(G))$ is connected. Similarly, if $\Gamma_\epsilon(x)$ is a loxodromic group, then so is G. Again, we may see that $\mathrm{thin}_\epsilon^\sim(M(G))$ is connected, since the shortest path from any point of $\mathrm{thin}_\epsilon^\sim(M(G))$ to the axis lies within $\mathrm{thin}_\epsilon^\sim(M(G))$. This completes the proof of the claim.

If G is parabolic, we call $T = \mathrm{thin}_\epsilon(M(G))$ a *Margulis cusp*. If G is loxodromic, we call T a *Margulis tube*. In the latter case, the quotient of the loxodromic axis is either a short arc, or a short closed geodesic, which we call the *core* of the tube. The tube is a regular neighbourhood of the core in the quotient orbifold. A cross-section of the tube is starlike about its intersection with the axis. In fact it is a finite union of convex sets, since the tube is a finite union of sets of the form $\{x \in \mathbf{H}^n \mid d(x, \gamma x) < \epsilon\}/G$, each of which has convex cross-section.

We shall denote by $\mathrm{thick}_\epsilon(M)$, the closure of the complement of $\mathrm{thin}_\epsilon(M)$, in M. We call $\mathrm{thick}_\epsilon(M)$ the *thick part* of M.

These definitions are most natural when Γ acts freely. Then, $\mathrm{thin}_\epsilon(M)$ is the set of points with injectivity radius at most $\epsilon/2$. The definitions for the orbifold case are not standard, but are convenient for our purposes.

To give the fourth definition of geometric finiteness, we need to define the "convex core" of a hyperbolic orbifold. The definition is the same as that given for a hyperbolic 3-manifold in the introduction. Let Λ be the limit set of Γ. The "convex hull", Y, of Λ is the minimal closed convex subset of $\mathbf{H}^n$ whose closure Y_C in $\mathbf{H}^n_C$ contains Λ. The construction of Y is best seen in the Klein model for hypebolic space (see [Th1]). From this picture, it is clear that $Y_C \cap \mathbf{H}^n_I = \Lambda$. Since the construction is Γ-equivariant, we may project Y to a subset, $\hat{Y}$, in the quotient orbifold, M. We call $\hat{Y}$ the *convex core of* M (FIG 2.7).

Definition 4 : *Γ is GF4 if, for some $\epsilon < \epsilon(n)$, $\epsilon > 0$, $\hat{Y}(\Gamma) \cap \mathrm{thick}_\epsilon(M(\Gamma))$ is compact.*

We will describe below an alternative way of defining a thick-thin decomposition for orbifolds. The resulting decomposition is identical for manifolds, and qualitatively similar for other orbifolds. The definition is suggested by the following proposition, which we also use in discussing GF5 in Ch.3.

Proposition 2.1 : *For each n there is some $N = N(n)$, such that if $x \in \mathbf{H}^n$ lies in the interior of $\mathrm{thick}_\epsilon^\sim(M)$ (the lift of $\mathrm{thick}_\epsilon(M)$ to $\mathbf{H}^n_C$), then $\Gamma_{\epsilon/N}(x)$ is finite.*

We begin the proof of Proposition 2.1 with the following lemma.

Lemma 2.2 : *Let G be any group, and $H \leq G$, a subgroup with $[G : H] = k$. If $G = \langle A \rangle$, $A \subset G$, then $H = \langle H \cap (A^{2k+1}) \rangle$.*

(As in Lemma 1.1, if $X \subset G$, we denote by X^r the set of those $g \in G$ expressible as words of length r in elements of $X \cup \{1\} \cup X^{-1}$.)

Proof of Lemma : The proof will be similar to that of Lemma 1.1.

Given any $h \in H$, we can write $h = \prod_1^p g_i$ with $g_i \in A$. If $p > 2k + 1$, consider the collection $\{Hh_j \mid j = 1, \ldots, k+1\}$, where $h_j = \prod_1^j g_i$. These cosets cannot all be distinct. Thus, $h = \alpha\beta\gamma$ with $H\alpha\beta = H\alpha$, $\alpha \in A^k$, $\beta \neq 1$, and $\alpha\beta \in A^{k+1}$. We can write $h = (\alpha\beta\alpha^{-1})h'$ where $h' = \alpha\gamma$. But $\alpha\beta\alpha^{-1} \in H \cap A^{2k+1}$, and h' has shorter word-length than h, so the result follows by induction.

◇

Proof of Proposition : Let $N = 2\nu(n) + 1$, where $\nu(n)$ is the bound on index in the Margulis Lemma. We fix $x \in \mathrm{int\,thick}_\epsilon^\sim(M)$, and consider $\Gamma_\epsilon(x) = \langle \Gamma \cap I_\epsilon(x) \rangle$. From Chapter 1, we know that $\Gamma_\epsilon(x)$ contains an abelian subgroup K, with $|\Gamma_\epsilon(x) : K| \leq \nu(n)$. Let T be the torsion subgroup of K. Since $x \in \mathrm{int\,thick}_\epsilon^\sim(M)$, we have $I_\epsilon(x) \cap \mathrm{free}\,\Gamma = \emptyset$, so that $K \cap I_\epsilon(x) \subseteq T$. We write I_ϵ for $I_\epsilon(x)$, etc.

Let $\eta = \epsilon/N$, so that $I_\eta^N \subset I_\epsilon$. Then $[\Gamma_\eta : \Gamma_\eta \cap K] \leq \nu(n)$, where $\Gamma_\eta = \langle \Gamma \cap I_\eta \rangle = \langle \Gamma_\epsilon \cap I_\eta \rangle$. From the lemma,

$$\begin{aligned} \Gamma_\eta \cap K &= \langle (\Gamma_\epsilon \cap I_\eta)^N \cap K \rangle \\ &\subseteq \langle \Gamma_\epsilon \cap I_\eta^N \cap K \rangle \\ &\subseteq \langle \Gamma_\epsilon \cap (I_\epsilon \cap K) \rangle. \end{aligned}$$

But $I_\epsilon \cap K \subset T$, so $|\Gamma_\eta| \leq \nu(n)|T| < \infty$.

The case when Γ_ϵ is loxodromic is similar.

◇

Suppose that $\eta \leq \epsilon(n)/N(n)$, and let $F_\eta(\Gamma) = \{x \in \mathbf{H}^n \mid \Gamma_\eta(x) \text{ is infinite }\}$. $F_\eta(\Gamma)$ is closed in $\mathbf{H}^n$, since we defined our sets $I_\eta(x)$ to be closed. It projects to a set which we denote by $\mathrm{thin}'_\eta(M)$ in M. We write $\mathrm{thick}'_\eta(M)$ for the closure of its complement in $\mathbf{H}^n_C$. For $\epsilon \leq \epsilon(n)$, we have the inclusions

$$\mathrm{thin}_{\epsilon/N}(M) \subset \mathrm{thin}'_{\epsilon/N}(M) \subset \mathrm{thin}_\epsilon(M).$$

Again, the connected components of $\operatorname{thin}'_\eta(M)$ are of two types - tubes and cusps. In GF4 $\Rightarrow$ GF1, we shall see that if M is GF, then $Y \cap \operatorname{thick}_\epsilon(M)$ is compact for any $\epsilon > 0$, arbitrarily small. This fact means that we can reformulate GF4 by demanding that $\hat{Y} \cap \operatorname{thick}'_\eta(M)$ be compact for some $\eta < \epsilon(n)/N(n)$.

GF5

Definition 5 : *Γ is GF5 if it is finitely generated and, for some $\eta > 0$, the η-neighbourhood, $N_\eta(\hat{Y}(\Gamma))$ of $\hat{Y}(\Gamma)$ has finite volume.*

We suspect that the assumption of finite-generation is unnecessary. We show this to be the case:

(i) if $|\operatorname{stab}_\Gamma(x)|$ is bounded for $x \in \mathbf{H}^n$ (for example, if Γ acts freely); or

(ii) if $M(\Gamma)$, itself, has finite volume; or

(iii) if $n \leq 3$.

3. Proofs of Equivalence.

The main cycle of proofs will be:

$$\begin{tikzcd} 3 \arrow[r, leftrightarrow] & 1 \arrow[r] \arrow[d, leftrightarrow] \arrow[rd] & 5 \arrow[d] \\ & 2 & 4 \arrow[l] \end{tikzcd}$$

We use GF1 as our central definition, since most the facts about geometrically finite groups are most easily deduced from this. We include proofs of $1 \Rightarrow 2$ and $1 \Rightarrow 4$ since they are very much shorter than following the cycle. The only non-geometric input is an appeal to the Selberg Lemma (Chapter 0) which overcomes a technical difficulty in the proof of $5 \Rightarrow 4$.

GF1 $\Rightarrow$ GF2

We have $M_C = \hat{N} \cup (\bigcup \hat{\mathcal{C}})$, where $\hat{N}$ is the projection of a compact set $N_0 \subseteq \mathbf{H}^n \cup \Omega$, and $\hat{\mathcal{C}}$ is a finite set of cusp regions. Let $y \in \Lambda(\Gamma)$.

Suppose that y is the is the fixed point of a parabolic group Γ_y, which stabilises some cusp region C. Then $(\Lambda \setminus \{\infty\})/\Gamma_y$ is a closed subset of the relatively compact set $(\mathbf{H}^n_C \setminus C)/\Gamma_y$, and is thus compact. We see that, in this case, y is a b.p.f.p.

So, suppose that y does not correspond to a cusp region in the way described above. (It is still conceivable, at the stage, that y may be a p.f.p., though this does not affect the argument.) We must have $|\Lambda| \geq 2$, so that the convex hull Y meets $\mathbf{H}^n$. We join y to a point $x \in Y \cap \mathbf{H}^n$ by a geodesic ray l. Note that $\bigcup(\Gamma l) \subseteq Y$. Clearly, l must leave any cusp region it enters, so the quotient $\hat{l}$ must accumulate somewhere in $\hat{N}$. Hence, Γl accumulates somewhere in $N_0 \cap Y$, and so, in this case y is a c.l.p.

GF2 $\Rightarrow$ GF1

First some general remarks.

Let $K \subseteq \mathbf{H}^n$ be a closed convex set. We may define the *nearest point retraction* $\rho_K : \mathbf{H}^n \to K$, where $\rho_K(x)$ is the nearest point of K to x. This map extends continuously to ideal points, $\rho_K : \mathbf{H}^n_C \to K_C$, where K_C is the closure of K in $\mathbf{H}^n_C$. We may describe the extension as follows. For $x \in K_C \cap \mathbf{H}^n_I$, take $\rho_K(x) = x$, and for $x \in \mathbf{H}^n_I \setminus K_C$, take $\rho_K(x)$ to be the unique point such that $K_C \cap B = \{\rho_K(x)\}$ for some horoball B about y. Notice that if, for a pair K, L of convex sets, $\rho_{K \cap L}(x) = y \in \operatorname{int} K$, then $\rho_L(x) = y$.

Given a set $X \subseteq \mathbf{H}^n$, we shall denote by $N_r(X)$ the uniform r-neighbourhood of X, i.e. $\{x \in \mathbf{H}^n \mid d(x, X) \leq r\}$. We shall say that two closed convex sets, K_1 and K_2, are λ-*near* (for some $\lambda > 0$) if $K_1 \subseteq N_\lambda(K_2)$ and $K_2 \subseteq N_\lambda(K_1)$. We show:

Lemma 3.1 : *Given $\lambda > 0$, there exists $L = L(\lambda) > 0$ such that if K_1 and K_2 are λ-near, then $d(\rho_1(x), \rho_2(x)) < L$ for all $x \in \mathbf{H}^n_C$, where $\rho_i = \rho_{K_i}$.*

Proof : Given a triangle xyz in $\mathbf{H}^2$, possibly with x an ideal point, if the angles at y and z are both at least $\pi/4$, then $d(y,z) \leq L_1$, where L_1 is some fixed constant. Also, given λ, we may find L_2 so that any two points, a distance no more than λ apart, subtend an angle of less than $\pi/4$ at any third point, distant at least L_2 from one of them. Let $L = \max(L_1, L_2)$. Then $L > \lambda$.

Suppose now, $x \in \mathbf{H}^n_C$, with $y_i = \rho_i(x)$ and $d(y_1, y_2) \geq L$. This means that x, y_1, y_2 are all distinct. Since $y_2 \in K_2$, there is a point $y_1' \in N_\lambda(y_2) \cap K_1$. Similarly, there is some $y_2' \in N_\lambda(y_1) \cap K_2$ (FIG 3.1). By convexity, the line segment $y_1 y_1'$ lies in K_1. Since $d(x, y_1)$ is minimal, the angle $x\hat{y}_1 y_1'$ is at least $\pi/2$. So the angle $x\hat{y}_1 y_2$ is at least $\pi/2 - \pi/4 = \pi/4$. Similarly, $x\hat{y}_2 y_1$ is at least $\pi/4$. But $d(y_1, y_2) \geq L_1$, contradicting the fact that x, y_1, y_2 form a triangle.

$\diamond$

Suppose, more generally, we have a closed convex set X, with $K_1 \cap N_L(X)$ and $K_2 \cap N_L(X)$ λ-near. (Note that $N_L(X)$ is also closed and convex.) Let ρ_i' be the retraction onto $K_i \cap N_L(X)$. Let $x \in \mathbf{H}^n_C$. By the lemma, $d(\rho_1'(x), \rho_2'(x)) < L$. Suppose $\rho_1(x) \in X$. We must have $\rho_1'(x) = \rho_1(x)$, so $\rho_2'(x) \in \operatorname{int} N_L(X)$. Hence, $\rho_2(x) = \rho_2'(x)$. In other words, $\rho_1(x) \in X$ implies that $\rho_2(x) \in N_L(X)$. We have shown:

Corollary 3.2 : *Let $X, K_1, K_2 \subseteq \mathbf{H}^n$ be closed convex subsets. If $K_1 \cap N_L(X)$, $K_2 \cap N_L(X)$ are λ-near, then $\rho_1^{-1}(K_1 \cap X) \subseteq \rho_2^{-1}(K_2 \cap N_L(X))$, where ρ_i is the retraction onto K_i.*

In proving GF2 $\Rightarrow$ GF1, the first step will be to construct standard parabolic regions about each b.p.f.p. We shall arrange that these regions are *strictly invariant* under Γ, i.e. they are disjoint, collectively invariant under Γ, and the stabiliser of each region is equal to the stabiliser of the corresponding p.f.p. They therefore project to disjoint cusp regions in M. A priori, there may be infinitely many of these. However, in the second part of the proof, we go on to show that their complement, in M_C, is relatively compact, so that there could only have been finitely many.

Since the construction of standard parabolic regions about b.p.f.p.s is valid for any discrete group, we stae it as a separate proposition.

Proposition 3.3 : *Let Γ be a discrete group, and let $P \subset \Lambda$ be a Γ-invariant collection of b.p.f.p.s. Then, there exists a collection of cusp regions $\{C(p) \mid p \in P\}$, which is strictly invariant under Γ, i.e. the regions are mutually disjoint, and $C(\gamma p) = \gamma C(p)$ for all $\gamma \in \Gamma$ and $p \in P$.*

Proof : If Γ is parabolic, the result is trivial. Hence we shall assume that $|\Lambda| \geq 2$ so that the convex hull, Y, of Λ meets $\mathbf{H}^n$. The retraction ρ_Y onto Y is clearly equivariant under the action of Γ. Let $p \in P$, and let $T(p) \subset \mathbf{H}^n$ be a Margulis region about p, as defined in Chapter 2 (GF4), i.e. $T(p) = \operatorname{thin}_\epsilon(\operatorname{stab}_\Gamma p)$. The regions $\{T(p) \mid p \in P\}$ are strictly invariant in the sense defined above. It follows that this is true of the regions $T(p) \cap Y$ and hence $S(p) = \rho_Y^{-1}(T(p) \cap Y)$ also. We need therefore only to show that each $S(p)$ contains a standard parabolic region $C(p)$.

Focusing on one such $p = \infty$ in $\mathbf{R}^n_+$, with stabliser Γ_p, we know that $\Lambda(\Gamma) \setminus \{\infty\} \subset Q_k = \{x \mid d_{euc}(x, \sigma_I) \leq k\} \subset \partial\mathbf{R}^n_+$, where σ_I is a minimal Γ_p-invariant plane. Let $v : \mathbf{R}^n_+ \to \mathbf{R}^n_+$ be vertical euclidean projection. It is not difficult to see that $v(T(p)) = \partial\mathbf{R}^n_+$. Moreover, we can choose a horoball, $B(p)$, about p, so that $B(p) \cap v^{-1}Q_k \subset T(p)$. Since $Y \subset v^{-1}Q_k$, we have that $Y \cap B(p) \subset T(p)$.

We have assumed that $\Lambda \neq \{\infty\}$. Hence each point of σ_I lies within some bounded euclidean distance of Λ. Since $v^{-1}\Lambda \subset Y$, we see that $Y \cap B(p)$ and $\sigma \cap B(p)$ are λ-near for some $\lambda > 0$ (recalling the notation $\sigma = v^{-1}\sigma_I$). Let $B'(p)$ be the horoball with $\partial B'(p)$ a hyperbolic distance $L(\lambda)$ above $\partial B(p)$, i.e. $B(p) = N_L(B'(p))$. From the corollary to our lemma, we have the inclusions $\rho_\sigma^{-1}(\sigma \cap B'(p)) \subset \rho_Y^{-1}(Y \cap B(p)) \subset \rho_Y^{-1}(Y \cap T(p)) = S(p)$ (FIG 3.2). Then, $\rho_\sigma^{-1}(\sigma \cap B'(p))$ is a standard parabolic region $C(p)$.

$\diamond$

Proof of 2 $\Rightarrow$ 1 : Let Γ be GF2, and let $P \subset \mathbf{H}^n_I$ be the set of all b.p.f.p.s. Let $\mathcal{C} = \{C(p) \mid p \in P\}$ be the collection of standard parabolic regions constructed as in Proposition 3.3. Let N be the closure, in $\mathbf{H}^n \cup \Omega$, of the complement $(\mathbf{H}^n \cup \Omega) \setminus \bigcup_{p \in P} C(p)$. In the quotient, we may write $M_C = \hat{N} \cup \bigcup \hat{\mathcal{C}}$, where $\hat{N}$ is the projection of N, and $\hat{\mathcal{C}}$ is a collection of standard cusps. We want to show that $\hat{N}$ is compact.

Let a lie in $N \cap \mathbf{H}^n \setminus \Sigma$, where N is the lift of $\hat{N}$ to $\mathbf{H}^n \cup \Omega$, and Σ is the singular set. Let $P = D(a)$ be the Dirichlet region about a. The set P is convex, and its images under Γ are locally finite in $\mathbf{H}^n$. Hence,

P_C, its closure in $\mathbf{H}^n_C$, can contain no c.l.p..

Suppose P_C meets a b.p.f.p., $p = \infty$ in $\mathbf{R}^n_+$. $(\mathbf{H}^n_C \backslash \mathrm{int}C(p))/\Gamma_p$ is compact. Since the images of $P \backslash \mathrm{int}C(p)$ under Γ_p are disjoint and locally finite, $P \backslash \mathrm{int}C(p)$ must have finite euclidean diameter. We see that $\mathrm{int}C(p)$ is an open neighbourhood of p in P_C. Thus, $N \cap P_C - P_C \backslash (\bigcup_p \mathrm{int}C(p))$ is a closed subset of P_C, and hence a compact subset of $\mathbf{H}^n \cup \Omega$. $P_C \backslash A$, being the closure of a Dirich'et region in $\mathbf{H}^n \cup \Omega$, is a fundamental domain for Γ acting on this set. It follows that $\hat{N}$ is a quotient of $P_C \cap N$, and is therefore compact.

Finally, suppose there were an infinite sequence $(\hat{C}_n)$ of distinct cusp regions. We take $x_n \in \partial C_n \cap P_C \cap \mathbf{H}^n$, with $d(x_n, \gamma_n x_n) < \epsilon(n)$ for some parabolic γ_n stabilising C_n. (Here $\epsilon(n)$ is the Margulis constant.) Margulis constant by some corresponding parabolic element. Taking a subsequence, we have $x_n \to y \in \mathbf{H}^n \cup \Omega$. If $y \in \mathbf{H}^n$, then $\Gamma_\epsilon(y)$ contains parabolics with different fixed points, contradicting what we know about the structure of $\Gamma_\epsilon(y)$ from Ch.1. If $y \in \Omega$, then $\min\{d(x_n, \gamma x_n) | \gamma \in \mathrm{free}(\Gamma)\} \to \infty$, contradicting the choice of x_n.

$\diamond$ 2 $\Rightarrow$ 1

GF1 $\Rightarrow$ GF3

We have $M_C = \hat{N} \cup (\bigcup \hat{\mathcal{C}})$, where $\hat{\mathcal{C}} = \{\hat{C}_1, \ldots, \hat{C}_k\}$. We may write $\hat{C}_i = C_i/\Gamma$, where $C_i = \{x | d_{euc}(x, (\sigma_i)_I) \geq r_i\} \subset \mathbf{R}^n_+ \cup \partial \mathbf{R}^n_+$, and Γ_i is the stabiliser of ∞. Choose $a_0 \in N$ and $a_i \in \mathrm{int}C_i \cap \sigma_i$, where σ_i is the vertical plane above $(\sigma_i)_I$. Let $\underline{a} = \{a_0, a_1, \ldots, a_k\}$. Let $A_\Gamma(\underline{a})$ be the complex defined in Chapter 2, GF3 (FIG 3.3). We fix our attention on some C_i. It is clear that $C_i \cap \Gamma \underline{a} = \Gamma_i a_i$. Since $\Gamma_i a_i \subset C_i \cap \sigma_i$, we see that the highest points of $\Gamma \underline{a}$ (those with largest nth euclidean coordinate) are precicely the points of $\Gamma_i a_i$. Thus, if $d_{euc}(x, (\sigma_i)_I) \geq r'_i$ for some fixed $r'_i > r_i$, the nearest points of $\Gamma \underline{a}$ to x lie in $\Gamma_i a_i$. Let C'_i be the standard region with radius r'_i. Within C'_i, the complex $A_\Gamma(\underline{a})$ is identical to that obtained from a_i for the group Γ_i, i.e. $C'_i \cap A_\Gamma(\underline{a}) = C'_i \cap A_{\Gamma_i}(a_i)$. Write $A(i) = A_{\Gamma_i}(a_i)$. Since g_i fixes σ_i, $A(i)$ is a euclidean product in the directions orthogonal to σ_i. Since $(\sigma_i)_I/\Gamma_I$ is compact, $A(i)/\Gamma_i$ must be finite. Rewriting $M_C = \hat{N}' \cup (\bigcup_{i=1}^k \hat{C}'_i)$, with $\hat{N}'$ compact, we see that A/Γ is finite.

$\diamond$ 1 $\Rightarrow$ 3

GF3 $\Rightarrow$ GF1

First we make a few remarks about general convex sets in $\mathbf{H}^n$.

Let $K \subseteq \mathbf{H}^n$ be a convex set with non-empty interior. Let $a \in K$, and let $T_1(a)$ be the unit tangent space to $\mathbf{H}^n$ at a. Clearly, K determines a cone in the tangent space at a, which intersects the the unit tangent space in a subset $T_1^K(a)$. We define $\omega(K, a)$ to equal $\mu(T_1^K(a))/\mu(T_1(a))$, where μ is spherical Lebesgue measure (FIG 3.4). (Alternatively, $\omega(K, a)$ may defined as the Lebesgue density of K at a.) The function $\omega(K, -)$ is strictly positive and lower semicontinuous on K. Let K_C be the closure of K in $\mathbf{H}^n_C$, and let $K_I = K_C \cap \mathbf{H}^n_I$. We may extend ω to a funtion on K_C as follows. If $y \in K_I = K_C \cap \mathbf{H}^n_I$, we consider the tangent space to $\mathbf{H}^n_I$ at y, and measure the proportion of unit tangent vectors to $\mathbf{H}^n_I$ lying in K_I. Note that on K_I it is possible (for example if $K_I = \{y\}$) to have $\omega(K, y) = 0$. If this is so, we call y a *cusp point* of K.

We shall restrict our attention to the case where K is a finite intersection of closed half-spaces, and $\mathrm{int}K \neq \emptyset$. We may write $K_C = \bigcap_{i=1}^k H_i$, where each H_i is a closed half-space in $\mathbf{H}^n_C$. By ∂H_i, we shall mean the closure, in $\mathbf{H}^n_C$, of the boundary of $H_i \cap \mathbf{H}^n$ in $\mathbf{H}^n$.

Suppose first, that $\bigcap_i \partial H_i \neq \emptyset$. If $\bigcap_i \partial H_i$ contains a point of $\mathbf{H}^n$, then it is a (possibly 0-dimensional) plane $\pi \subset \mathbf{H}^n_C$, and in this case, $\omega(K, x)$ takes a fixed, strictly positive value for all $x \in \pi$. If, on the other hand, $\bigcap_i \partial H_i$ contains only ideal points, it must consist of a single point $y \in \mathbf{H}^n_I$, with $\omega(K, y) \geq 0$.

More generally, if K is a finite-sided convex polyhedron, by considering all possible intersections of half-spaces H_i, we may deduce the following.

Lemma 3.4 : *Let K_C be a finite-sided convex polyhedron in $\mathbf{H}^n_C$. Then, there exists a finite set $\kappa(K_C)$ of cusp points in K_I, and $\delta(K_C) > 0$ such that for all $x \in K_C \backslash \kappa(K_C)$, we have $\omega(K, x) > \delta(K_C)$.*

Proof of 3 $\Rightarrow$ 1 : Let Γ be GF3. Let A be a Γ-invariant convex cell complex so that A/Γ is finite, and such that for each $A \in A$, $\{\gamma | \gamma A = A\}$ is finite. We stated in Chapter 2 (GF3) that any cell of a convex cell complex meets the closures of only finitely many other cells. Since $\mathcal{A}/\Gamma$ is finite, we may find a fixed constant, k_1, such that each point of $\mathbf{H}^n$ meets the closures of at most k_1 cells of $\mathcal{A}$. Also, the orders of $\mathrm{stab}_\Gamma A$ are all finite and therefore bounded by some k_2. Thus, for each $x \in \mathbf{H}^n$ and $A \in \mathcal{A}$, we have $|\{\gamma | x \in \gamma \bar{A}\}| \leq k_1 k_2$, and so at most $k_1 k_2$ faces of any A can be equivalent under Γ. It follows that each cell of $\mathcal{A}$ has only finitely many faces, and in particular, is (the interior of) a finite-sided polyhedron.

By hypothesis, $\mathcal{A}$ is locally finite in $\mathbf{H}^n$. Let $y \in \Omega$, and $H_1 \subset \mathbf{H}^n \cup \Omega$ be a half-space, containing y in its interior. We can insist that H_1 is invariant under $\mathrm{stab}_\Gamma y$, and $H_1 \cap \gamma H_1 = \emptyset$ if $\gamma \notin \mathrm{stab}_\Gamma y$. Let H_2 and H_3 be successively smaller ($\mathrm{stab}_\Gamma y$)-invariant half-spaces containing y. Only finitely many cells of $\mathcal{A}$ lie entirely within H_1. Any cell that meets both ∂H_1 and ∂H_3 also meets $\partial H_2 \cap \mathrm{hull}(\partial H_1 \cup \partial H_3)$, which is a compact subset of $\mathbf{H}^n$ (FIG 3.5). It follows that only finitely many cells meet H_3. We have shown that $\mathcal{A}$ is locally finite on $\mathbf{H}^n \cup \Omega$. Hence, every point of $\mathbf{H}^n \cup \Omega$ lies in the closure of some top-dimensional cell.

The set of closures of top-dimensional cells consists of the image under Γ of a finite set of polyhedra $\{P^1, \dots, P^k\}$. M_C is thus a quotient of the set $\bigcup_1^k P_C^i \setminus \Lambda(\Gamma)$.

We shall show below that each P_C^i can meet Λ only in a finite set of points (a subset of the set of cusp points $\kappa(P_C^i)$ of P_C^i), and that each of these points is a b.p.f.p. When we have done this, the proof of GF1 can be completed as follows. If $C(p)$ is a standard parabolic region about $p \in P_C^i \cap \Lambda$, then $\mathrm{int}C(p) \cap P_C^i$ is an open neighbourhood of p in P_C^i, which we shall call a *cusp*. We choose standard cusp regions for each conjugacy class of b.p.f.p. meeting some polyhedron P_C^i. By taking these regions small enough, we can ensure that the cusps they form in each polyhedron are disjoint. The parabolic regions C are then themselves disjoint. The closure of each $P_C^i \setminus (\bigcup C)$ is compact, and so therefore is the quotient $N = \mathrm{closure}(M_C \setminus (\bigcup \mathcal{C}))$, as required.

We now investigate $P_C^i \cap \Lambda$. Let $y \in \mathbf{H}_I^n$ lie in the closure of some polyhedron, P_C. We want to show that either $y \in \Omega$ or y is a b.p.f.p.. The stabliser Γ_y of y cannot contain any loxodromic element, since otherwise, by applying a loxodromic element with y as repelling fixed point, we would get a contradiction to the local finiteness of $\mathcal{A}$ along the axis of the element. Therefore, Γ_y preserves setwise each horosphere about y, and so Γ_y is either a finite or a parabolic group. Let $y = \infty$ in $\mathbf{R}_+^n$. Let $\mathcal{P}$ ($\neq \emptyset$) be the set of closures of top-dimensional cells containing y. We distinguish two cases.

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

By lemma 3.4, there are only finitely many such points in each polyhedron. It follows that there is a horoball B for which $(B \cap \mathcal{P})/\Gamma_y$ is finite. (By $B \cap \mathcal{P}$, we mean $\{B \cap P \,|\, p \in \mathcal{P}\}$.) Thus, if we take ∂B to be high enough, we can ensure that each $B \cap P$ is a vertical prism on $\partial B \cap P$ (i.e. isometric to $(\partial B \cap P) \times [0, \infty)$). It is still possible that the boundary of such a prism may be subdivided into many cells of $\mathcal{A}$. However, we know from the first paragraph of the proof that each polyhedron has only finitely many faces in $\mathcal{A}$. Thus, by raising the level of ∂B if necessary, we ensure that if A is a face of some $P \in \mathcal{P}$, then $A \cap B$ is a vertical prism (possibly empty). It now follows that $\bigcup(B \cap \mathcal{P}) = B$. For if not, consider a codimension-1 face A that bounds $\bigcup(B \cap \mathcal{P})$ in B. We know that A is a vertical prism, and so the (top-dimensional) polyhedron on each side of A has ∞ in its closure. This contradicts the definition of $\mathcal{P}$.

Suppose now that, for some $\gamma \in \Gamma \setminus \Gamma_y$, we have $\gamma B \cap B \neq \emptyset$. Since $B \subset \bigcup \mathcal{P}$, we see that γB meets some polyhedron $P \in \mathcal{P}$. Thus, P is the image under γ of some $Q \in \mathcal{P}$. From the finiteness of $\mathcal{P}/\Gamma_y$, and of the setwise stabiliser of each polyhedron, we see that γ must lie in one of a finite number of right cosets of Γ_y in Γ. This gives us an upper bound on the height of the highest point of γB in $\mathbf{R}_+^n$. Thus by raising the level of ∂B still further, we can arrange that $\gamma B \cap B = \emptyset$ for all $\gamma \in \Gamma \setminus \Gamma_y$. Thus, we have found a strictly invariant horoball $B \subset \bigcup \mathcal{P}$. Now, using B, we want to construct another strictly invariant region, C, which will be either a standard parabolic region or else a half space, depending on whether Γ_y is parabolic or finite. (In fact, it turns out that the latter case cannot arise in Case 1.)

Let $\sigma_I \subset \partial \mathbf{R}_+^n$ be a minimal Γ_y-invariant subspace. By the finiteness of $\mathcal{P}/\Gamma_y$, there is a bound on the euclidean diameters of the lower-dimensional faces, of the polyhedra in $\mathcal{P}$, which do not contain ∞. Hence, for some r, we see that $C = \{x \,|\, d_{euc}(x, \sigma_I) \geq r\}$ is contained in $\bigcup \mathcal{P}$. The structure of $C \cap \mathcal{P}$ is independent of the vertical coordinate. Since $\mathcal{P}$ is locally finite on $\mathbf{R}_+^n$, we see that $\mathcal{P}$ must also be locally finite on $C \cap \partial \mathbf{R}_+^n$. Therefore the limit set does not meet $C \cap \partial \mathbf{R}_+^n$. If Γ_y is parabolic, then y is a b.p.f.p., and C is a standard parabolic region. Note that C meets P_C^i in an (arbitrarily small) neighbourood of a cusp point. If Γ_y is finite, then C is a half space, with $C \cap \mathcal{P}$ finite. Thus $y \in \Omega$. (Now, it is fairly easy to see that we get a contadiction to the hypothesis of Case 1, though logically we do not need this.)

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

Let $\delta = \min_i \delta(P_C^i) > 0$ (see Lemma 3.4). We know that $\omega(P, y) > \delta$, so that Γ_y must be finite (its order bounded by $|\mathrm{stab}_\Gamma(P)|/\delta$). Since $|\Gamma_y| < \infty$, we can have $\omega(Q, y) = 0$ for only finitely many $Q \in \mathcal{P}$. Otherwise, $\omega(Q, y) > \delta$. It follows that $\mathcal{P}$ is finite. By an argument similar to that in Case 1, we show that some half-space, containing y, lies in $\bigcup \mathcal{P}$. Thus, $y \in \Omega$.

Conclusion : We chose an arbitrary $y \in P_C^i$. If y lies in the limit set, we must be in Case 1. Then, y is a b.p.f.p., and the standard parabolic regions about y define a base of neighbourhoods for y in P_C^i. The proof may now be completed as indicated above.

$\diamond$ 3 $\Rightarrow$ 1

GF1 $\Rightarrow$ GF4

Let $\epsilon(n)$ be the Margulis constant, and $\epsilon > 0$ be any number less than $\epsilon(n)$.

Suppose Γ is GF1. Let $Y(\Gamma) = \mathrm{hull}(\Lambda(\Gamma))$. If p is a b.p.f.p., we may take the corresponding parabolic region C so that $C \cap Y(\Gamma)$ is contained in the corresponding Margulis region. Writing $M_C = \hat{N} \cup (\bigcup \hat{C})$, we have that $\hat{Y}(\Gamma) \cap \mathrm{thick}_\epsilon(M)$ is contained in $\hat{N}$, and is thus compact.

$\diamond$ 1 $\Rightarrow$ 4

GF4 $\Rightarrow$ GF2

We have $\mathrm{thick}_\epsilon(M) \cap \hat{Y}(\Gamma)$ compact for some $\epsilon < \epsilon(n)$.

Let R be a component of $\mathrm{thin}_\epsilon^\sim(M)$, as defined in Ch2, GF4. If R is the lift of a Margulis tube, then R lies within a uniform neighbourhood of a loxodromic axis. If R is the lift of a Margulis cusp, then it lies in some horoball about the p.f.p. Thus if l is a geodesic ray lying entirely within $\mathrm{thin}_\epsilon^\sim(M)$, its ideal endpoint is either a loxodromic fixed point, hence a c.l.p., or a p.f.p.

Suppose $y \in \Lambda(\Gamma)$ is neither a parabolic nor a loxodromic fixed point. Join y to $x \in Y(\Gamma)$ by a geodesic ray l. The union of those parts of l lying outside $\mathrm{thin}_\epsilon^\sim(M)$ is unbounded. Thus, its projection $\hat{l}$ must accumulate in $\mathrm{thick}_\epsilon(M) \cap \hat{Y}(\Gamma)$, and so y is a c.l.p.

It remains to show that any p.f.p. of Γ is bounded.

Let $\delta > 0$ be such that $\epsilon + 4\delta < \epsilon(n)$. Let $p = \infty$ in $\mathbf{R}_+^n$ be a p.f.p. Let R_ϵ, $R_{\epsilon+2\delta}$, $R_{\epsilon+4\delta}$ be the corresponding Margulis regions. The shell $\hat{S} = (R_{\epsilon+4\delta} \setminus \mathrm{int} R_\epsilon)/\Gamma_p$ embeds as a closed subset of $\mathrm{thick}_\epsilon(M)$. We show that if p is not bounded, then $\hat{S} \cap \hat{Y}$ is not compact.

Let $\sigma_I \subset \mathbf{H}_I^n$ be a minimal Γ_p-invariant plane. Suppose we have a sequence (x_n) in $\Lambda(\Gamma)$ with $d_{euc}(x_n, \sigma_I) \to \infty$. Let $y_n \in \partial R_{\epsilon+2\delta}$ lie vertically above x_n (FIG 3.6). The hyperbolic ball $N_\delta(y_n)$ is a subset of $S = R_{\epsilon+4\delta} \setminus \mathrm{int} R_\epsilon$. Taking a subsequence, we can assume that no such ball meets the image of a different ball under Γ_p. Since Γ_p is virtually abelian, it has a torsion-free subgroup of index k (say). Thus, each $(\delta/2)$-ball meets at most k images of itself. It follows that the quotient sequence $\hat{y}_n \in \hat{Y} \cap \hat{S}$ has no convergent subsequence.

$\diamond$ 4 $\Rightarrow$ 2

GF1 $\Rightarrow$ GF5

Suppose τ is an r-plane in $\mathbf{H}^n$. Let $\rho_\tau : \mathbf{H}^n \to \tau$ be the nearest point retraction. For $X \subset \tau$ and $h > 0$, let $X_{\tau,h} = N_h(\tau) \cap \rho_\tau^{-1} X$. ($N_h(\tau)$ is the uniform h-neighbourhood of τ.) There is a function $f : \mathbf{R}_+ \to \mathbf{R}_+$, for which $\mathrm{vol}_n(X_{\tau,h}) = f(h)\mathrm{vol}_r(X)$, where vol_i is the i-dimensional volume. In particular, $X_{\tau,h}$ has finite n-volume if and only if X has finite r-volume.

Let $\Gamma_p = \mathrm{stab}_\Gamma p$ be an infinite parabolic group, and σ a minimal Γ_p-invariant plane containing p. Let $r = \dim\sigma$. We know that $r \geq 2$. Let C be a standard parabolic region. Γ_p acts as a cocompact group on $\partial C \cap \sigma$, so we can find a compact $K \subseteq \partial C \cap \sigma \subset \mathbf{H}^n$ with $\bigcup(\Gamma_p K) = \partial C \cap \sigma$. $K' = \mathrm{hull}(K, p)$ has finite

r-volume, so $K'_{\sigma,h}$ has finite n-volume. Note that $C = \rho_\sigma^{-1}(C \cap \sigma)$, and so the images of $K'_{\sigma,h}$ under Γ_p cover $C \cap N_h(\sigma)$. We see that $\mathrm{vol}_n((C \cap N_h(\sigma))/\Gamma_p) < \infty$.

Now, suppose that Γ is GF1, $\eta > 0$. For each b.p.f.p. p, we can find a region C about p so that $C \cap N_\eta Y \subseteq C \cap N_{2\eta}(\sigma)$. Thus, $\mathrm{vol}_n(C \cap N_\eta Y)/\Gamma_p < \infty$. The remainder of $N_\eta \hat{Y}$ is compact, so $\mathrm{vol}_n(N_\eta \hat{Y}) < \infty$.

Finally, to see that Γ is finitely generated, note that each cusp is topologically a product orbifold $\hat{C} \cong \partial\hat{C} \times [0,1)$. Hence, $\Gamma = \pi_1 \hat{N}$.

◇ 1 ⇒ 5

GF5 ⇒ GF4

Let Γ be GF5. For some $\eta > 0$, $\mathrm{vol}_n N_\eta \hat{Y}(\Gamma) < \infty$. Let $\epsilon < \epsilon(n)$.

Suppose, first, that Γ acts freely. Let $\epsilon_1 < \epsilon, \eta$. Take a maximal packing of $N_\eta(Y)$ with $\epsilon_1/2$-balls centred in $\mathrm{thick}_\epsilon(M) \cap \hat{Y}$. Since each of these balls is isometric to a standard ball in $\mathbf{H}^n$, this packing is finite. By maximality, the corresponding ϵ_1-balls cover $\mathrm{thick}_\epsilon(M) \cap \hat{Y}$. Thus, $\mathrm{thick}_\epsilon(M) \cap \hat{Y}$ is closed, and covered by finitely many compact sets, and hence is itself compact.

Suppose, now, that Γ is any GF5 group. By the Selberg Lemma, Γ has a torsion-free subgroup of finite index. This must act freely on $\mathbf{H}^n$, so that the volume of an $\epsilon_1/2$-ball centred on $\mathrm{thick}_\epsilon(M)$ is bounded away from zero. The proof now works as before.

◇ 5 ⇒ 4

Let Γ be a discrete group of isometries of $\mathbf{H}^n$.

In Ch.2 GF5, we stated three cases in which finite generation is automatically implied by $\mathrm{vol}_n N_\eta(Y) < \infty$.

Case (i) : *If there is some number k, such that for each $x \in \mathbf{H}^n$, $|stab_\Gamma(x)| \leq k$, then Γ is GF.*

Proof : Let $\delta < \min(\eta, \epsilon(n)/N(n))$, where $N(n)$ is as defined in Ch.2 GF4. Let $y \in \mathrm{thick}_\delta^{\sim}(M) \subseteq \mathbf{H}^n$. We know from Proposition 2.2 that $\Gamma_\delta(y)$ fixes some point of $\mathbf{H}^n$. Hence, $|\Gamma_\delta(y)| \leq k$. $N_{\delta/4}(y)$ meets at most k images of itself under Γ. Thus, in the quotient, $\mathrm{vol}_n N_{\delta/4}(\bar{y})$ is bounded away from zero, for $\bar{y} \in \mathrm{thick}_\delta(M)$. The argument is now as for free actions on $\mathbf{H}^n$.

◇ Case(i)

Case (ii) :

Theorem 3.5 : *A finite-volume complete hyperbolic orbifold (without boundary) is geometrically finite.*

Let $\eta < \epsilon = \epsilon(n)/N(n)$, where $\epsilon(n)$ is the Margulis constant, and $N(n)$ is as defined in Proposition 2.2. In Ch.2 GF4, we defined $\mathrm{thin}'_\eta(M)$ to be the projection of the set $\{x \in \mathbf{H}^n \mid |\Gamma_\eta(x)| = \infty\}$. We saw that GF4 was equivalent to the statement that $\hat{Y}(\Gamma) \cap \mathrm{thick}'_\eta(M(\Gamma))$ be compact, for any $\eta < \epsilon$. We aim to show here, that if M has finite volume, then $\mathrm{thick}'_\delta(M)$ is compact for a certain $\delta < \epsilon$. We begin by giving a proof of the following proposition about general hyperbolic orbifolds. We shall then build upon our argument to deduce the main theorem (3.5).

Proposition 3.6 : *Given n, there is some universal $\delta > 0$, such that for any discrete subgroup Γ of $\mathrm{Isom}\mathbf{H}^n$, there is some $\tilde{x} \in \mathbf{H}^n$ for which $\Gamma_\delta(\tilde{x})$ is trivial, (so that $N_{\delta/2}(x)$ is an embedded hyperbolic $\delta/2$-ball in the quotient).*

Proof : Let Γ be a discrete group, with $M = \mathbf{H}^n/\Gamma$. Given $x \in M$, we define the *injectivity radius*, $\mathrm{inj}(x) = \frac{1}{2}\min\{d(\tilde{x}, \gamma\tilde{x}) \mid \gamma \in \Gamma\}$. (This is usually defined only for manifolds.) We write $\tilde{\Sigma}$ for the singular set of Γ in $\mathbf{H}^n$, so that $\Sigma = \tilde{\Sigma}/\Gamma = \{x \in M \mid \mathrm{inj}(x) = 0\}$. Let L be the union of all the loxodromic axes with translation distance less than $\epsilon(n)$. By the Margulis lemma, the collection of such axes is locally finite, so that $\tilde{L}$ is closed. It projection, $L \subset M$ is a disjoint union of arcs and simple closed curves.

Let $\eta \leq \epsilon = \epsilon(n)/N(n)$. We define a decomposition of M into disjoint pieces as follows. Given $X \subset \mathbf{H}^n_C$, let $\tilde{T}_\eta(X) = \{x \in \mathbf{H}^n \mid \mathrm{fix}\Gamma_\eta(x) = X\}$. Let $\tilde{\mathcal{T}}_\eta = \{\tilde{T}_\eta(X) \mid X \subset \mathbf{H}^n_C \text{ and } T_\eta(X) \neq \emptyset\}$. Clearly, $\tilde{\mathcal{T}}_\eta$ is invariant under Γ, so it projects to a decomposition $\mathcal{T}_\eta$ of M. We shall call the elements of the decomposition *η-compartments*. Let $T = T_\eta(X)$ be one such η-compartment. T is non-empty, so from Ch1, we know that X must either consist of one or two points of $\mathbf{H}^n_I$, or be a plane in $\mathbf{H}^n_C$. We may thus define $d(T)$ to be the dimension of X, with the convention that $d(T) = -1$ if $X \subset \mathbf{H}^n_I$. (This is well-defined since, if $T_\eta(X_1) = T_\eta(X_2)$, then X_1 and X_2 are equivalent under Γ.) Suppose $x \in T$ lifts to $\tilde{x} \in \mathbf{H}^n$. Then, $\mathrm{fix}\Gamma_\eta(\tilde{x})$ is a subset of $\mathbf{H}^n_I$ if and only if $\Gamma_\eta(\tilde{x})$ is infinite. Thus $d(T) = -1$ if and only if $x \in \mathrm{thin}'_\eta(M)$. We see that $\mathrm{thin}'_\eta(M) = \bigcup\{T \in \mathcal{T}_\eta \mid d(T) = -1\}$. We write $\mathcal{F}_\eta = \{T \in \mathcal{T}_\eta \mid d(T) \neq -1\}$, and $P_\eta = T_\eta(\mathbf{H}^n_C) = \{x \in \mathbf{H}^n \mid \mathrm{inj}(x) > \eta/2\}$.

If $T = T_\eta(X) \in \mathcal{T}_\eta \setminus \{P_\eta\}$, (i.e. $d(T) < n$), we define a smooth unit vector field v_η on $T \setminus (\Sigma \cup L)$ as follows. For $x \in T \setminus (\Sigma \cup L)$, we define $v_\eta(x)$ to be the unit tangent vector pointing directly away from $\mathrm{hull}(X)$, i.e. $v_\eta = -\alpha'(0)$, where α is the geodesic arc from $\tilde{x} = \alpha(0)$ to the nearest point on $\mathrm{hull}(X)$. (Note that $\mathrm{hull}(X) = X$, unless $\Gamma_\eta(\tilde{x})$ is an (infinite) loxodromic group, in which case $\mathrm{hull}(X)$ is the loxodromic axis.) This gives us a well-defined vector $v_\eta(x)$ at $x \in T$. Performing this construction for each $T \in \mathcal{T}_\eta \setminus \{P_\eta\}$, we get a (usually discontinuous) piecewise analytic vector field on $M \setminus (P_\eta \cup \Sigma \cup L)$. The integral curves are piecewise geodesic.

Now, we fix $\epsilon = \epsilon(n)/N(n)$, and choose any $\delta < \epsilon$. Suppose $x \in M \setminus (\mathrm{thin}'_\delta(M) \cup P_\delta \cup \Sigma \cup L)$, so that $\Gamma_\delta(\tilde{x})$ is finite and non-trivial. Let β be the integral curve through x for the vector field v_ϵ. We take $\beta(0) = x$, and write Γ_0 for $\Gamma_\delta(\tilde{x})$.

Imagine following the integral curve β in $\mathbf{H}^n$. At time t, we have $\Gamma_\delta(\tilde{\beta}(t)) \subseteq \Gamma_\epsilon(\tilde{\beta}(t))$. Let $h_\delta = \mathrm{hull}\,\mathrm{fix}\Gamma_\delta(\tilde{\beta}(t))$, and $h_\epsilon = \mathrm{hull}\,\mathrm{fix}\Gamma_\epsilon(\tilde{\beta}(t))$. Now, h_δ, and h_ϵ are both subspaces of $\mathbf{H}^n$, and $h_\delta \subset h_\epsilon$. Thus, it is easily checked that $\tilde{v}_\delta(\beta(t))$ always makes an acute angle with $\tilde{v}_\epsilon(\beta(t))$. This mean that the distance of $\beta(t)$ from h_δ increases at least linearly with t. Now, while $\Gamma_\delta(\tilde{\beta}(t))$ remains constant and equal to Γ_0, we see that the injectivity radius $\mathrm{inj}(\beta(t))$ increases steadily, and the derivative of the injectivity radius with respect to t is non decreasing. Thus, after a finite distance, at $\tilde{\beta}(t_1)$, say, $\Gamma_\delta(\tilde{\beta}(t))$ must change to a new group Γ_1. Now Γ_0 and Γ_1 are both subgroups of $\Gamma_\epsilon(\tilde{\beta}(t_1))$. Since we are moving away from $\mathrm{fix}\Gamma_\epsilon(\tilde{\beta}(t_1))$, it is easy to see that $\Gamma_1 < \Gamma_0$. Again, after another finite distance, $\Gamma_\epsilon(\tilde{\beta}(t))$ changes to a third group $\Gamma_2 < \Gamma_1$. Since Γ_0 is finite, it follows that, after a finite number of steps, we shall arrive at a point y, with $\Gamma_\delta(y) = \{1\}$. (Note that we can never run into $\Sigma \cup L$. Wherever the $\tilde{v}_\epsilon$ is discontinuous along $\Sigma \cup L$, the vector field radiates away.)

It is easy to see, from the form of the components of $\mathrm{thin}'_\delta(M)$, that $\mathrm{thin}'_\delta(M)$ cannot occupy all of M. $\Sigma \cup L$ is a lower-dimensional object, so either $\mathrm{int}\,\mathrm{thick}'_\delta(M) \subset P_\delta$, or we can find some point x as above. Either way, there is some $y \in \mathbf{H}^n$ with $\Gamma_\delta(y)$ trivial.

◇ Prop.3.6

Now, think of M as cut into δ-compartments. We have shown that $P_\delta \neq \emptyset$. In fact, we have shown that any point $x \in \mathrm{thick}'_\delta(M) \setminus (\Sigma \cup L)$ can be joined to some $y \in P_\delta$, by a path β with $\Gamma_\delta(\tilde{\beta}(t))$ monotonically decreasing (i.e. if $t' > t$, then $\Gamma_\delta(\tilde{\beta}(t')) < \Gamma_\delta(\tilde{\beta}(t))$). Consequently, $\mathrm{fix}\Gamma_\delta(\tilde{\beta}(t))$ is monotonically increasing with respect to set inclusion. Each time $\mathrm{fix}\Gamma_\delta(\tilde{\beta}(t))$ changes, its dimension must strictly increase. It follows that $\beta(t)$ passes through at most $n+1$ δ-compartments of M. Hence, any point of $\mathrm{thick}'_\delta(M)$ can be joined to P_δ by a path that passes through at most $k = n+1$ δ-compartments.

Now, the collection $\mathcal{T}_\delta$ is locally finite. To see this, note that, for any point $x \in \mathbf{H}^n$, the set $\{\gamma \in \Gamma \mid d(\gamma x, x) < 3\delta\}$ is finite, by the discreteness of Γ. So, there are only finitely many candidates for the generating set of any group $\Gamma_\delta(y)$, with $y \in N_\delta(x)$.

Let $\mathcal{F}_\delta = \{T \in \mathcal{T}_\delta \mid \mathrm{d}(T) \neq -1\}$. We know that this covers $\mathrm{int}\,\mathrm{thick}'_\delta(M)$.

Now, suppose that M has finite volume. We aim to show that $\mathrm{thick}'_\delta(M)$ is compact. Consider the δ-compartment $T \in \mathcal{F}_\delta$. We have $T = \tilde{T}_\delta(\sigma)$ for some plane $\sigma \subset \mathbf{H}^n_C$. Let $\mathrm{stab}''_\Gamma(\sigma)$ be the subgroup of Γ that fixes σ pointwise. Let $p = |\mathrm{stab}''_\Gamma(\sigma)|$. For any $x \in \tilde{T}$, $\Gamma_\delta(x) \subset \mathrm{stab}''_\Gamma(\sigma)$, so $N_{\delta/4}(x)$ meets at most p images of itself under Γ. This gives a positive lower bound on the volume of any metric $\delta/4$-ball in M, centred on a point of T, namely $1/p$ times the volume of a $\delta/4$-ball in $\mathbf{H}^n$. Since M has finite volume, we see that T has finite diameter, and is thus relatively compact.

We now think of $\mathrm{thick}_\delta^*(M)$ just as a topological space W. We summarise what we know about W.

W has a locally finite cover K by compact subsets (the closures of the δ-compartments). Also, there is a constant k $(= n+1)$, and a fixed element K_0 of the cover such that any point of W can be joined to K_0 by a path which is covered by at most k sets from K.

It follows from this, that K must itself be compact. To see this, think of the elements of the cover as vertices in an abstract graph. Two vertices are joined by an edge if the corresponding sets intersect. The graph has finite diameter (path condition), and each vertex has finite degree (compactness and local finiteness). Thus the graph is finite.

We have shown that $\mathrm{thick}_\delta^*(M)$ is compact. The discussion of GF4 shows us that M is GF.

◇ Thm.3.5

Case (iii) : *Let* Γ *be a discrete subgroup of* $\mathrm{Isom}\mathbf{H}^n$*, with* $n = 2$ *or* 3*. If* $\mathrm{vol}_n N_\eta \hat{\Upsilon}(\Gamma) < \infty$*, for some* $\eta > 0$*, then* Γ *is GF.*

Proof : We deal with the case $n = 3$ ($n = 2$ is similar). We aim to reduce this to Case (i), by showing that $|\mathrm{stab}_\Gamma(x)|$ is bounded for $x \in \mathbf{H}^n$.

Note that we can assume (by taking an index-2 subgroup if necessary) that Γ is orientation-preserving. We shall also suppose that $|\Lambda(\Gamma)| > 2$, otherwise Γ is trivially GF. We can take η to be less than $\epsilon(n)$, the Margulis constant.

Suppose then, that $|\mathrm{stab}_\Gamma(x)|$ is unbounded. From the Jordan Lemma, we can find a sequence (G_i) of finite abelian subgroups of Γ, with $|G_i| \to \infty$. From Lemma 1.3, using the fact that $n = 3$, we see that the fixed-point set of each G_i is a geodesic l_i in $\mathbf{H}^3$. (Thus, each G_i can be assumed to be cyclic of large order.) We can assume that these geodesics are all inequivalent under Γ.

Let $G = G_i$ be one such group, with fixed-point set l. Since $|\Lambda| > 2$, we can choose some $t \in \Lambda \setminus l$. Let m be the perpendicular from t, meeting l at a. Note that $m \subset \hat{\Upsilon}(\Gamma)$, and that $G \leq \Gamma''_\eta(a)$, where $\Gamma''_\eta(x) = \Gamma_\eta(x) \cap \Gamma'_{\epsilon(n)}(x)$. $\Gamma''_\eta(x)$ has index at most $\nu(n)$ in in $\Gamma_\eta(x)$. (see the discussion of the Margulis Lemma Ch.1). It is not difficult to see that we can find a point $y \in m$, with $\Gamma''_{\eta/2}(y)$ trivial, but with $\Gamma''_\eta(y)$ a non-trivial subgroup of $\Gamma''_\eta(a)$. Thus, $l = \mathrm{fix}\Gamma''_\eta(y)$, and $N_{\eta/8}(y)$ meets at most $\nu(n)$ images of itself under Γ. Note that $N_{\eta/8}(y) \subset N_\eta \hat{\Upsilon}(\Gamma)$.

Performing this construction for each group G_i, we obtain a sequence of points $y_i \in \mathbf{H}^n$. In the quotient, the $(\eta/8)$-balls $N_{\eta/8}(y_i)$ are disjoint (since $\mathrm{fix}\Gamma''_\eta(y_i) = l_i$), and their volumes are bounded below. This means that $N_\eta \hat{\Upsilon}(\Gamma)$ must have infinite volume.

◇ Case(iii)

4. Convex Fundamental Polyhedra.

In dimension 3, the central definitions of geometrical finiteness have traditionally been in terms of finite-sided fundamental polyhedra. In particular, the following statements are all equivalent to geometrical finiteness.

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

2a (2b) : Some (each) Dirichlet polyhedron has finitely many faces.

Here we use "face" in the sense of Chapter 2, GF3. This means that each polyhedron meets only finitely many images of itself under the group.

We can interpret our definition GF3 as being equivalent to the statement 1a without the assumption of convexity. The remaining definitions, 1b, 2a and 2b, no longer work in higher dimensions, as the following discussion shows.

First, consider on $\mathbf{E}^3$, an infinite cyclic group Γ of irrational screw motions with axis r, i.e. Γ is generated by a translation parallel to r composed with an irrational rotation with r as axis (FIG 4.1). If $a \notin r$, then the Dirichlet domain $D(a)$ about a is infinite-sided.

To see this, let Φ be the $(n-1)$-sphere of parallel classes of rays in $\mathbf{E}^n$. Suppose that $a \notin r$, and let l be the ray through a, perpendicular to r. Suppose $D(a)$ is finite-sided. Then, $D(a) = \bigcap_{\gamma \in G} H_\gamma$, where $G \subset \Gamma$ is finite, and H_γ is the half-space $\{x \in \mathbf{H}^n \mid d(x, a) < d(x, \gamma a)\}$. Note that l is never parallel to ∂H_γ for $\gamma \neq 1$. It follows that $l \in \operatorname{int}\Theta$, where $\Theta \subseteq \Phi$ is the set of rays lying in $D(a)$. (We may identify Φ with the set of rays emanating from a.) Now, Γ acts on Φ as a non-discrete rotation group fixing r. Thus, for some $\gamma \in \Gamma$, we have $\operatorname{int}\Theta \cap \gamma \operatorname{int}\Theta \neq \emptyset$, so that $D(a) \cap \gamma D(a) \neq \emptyset$, contradicting the assumption that $D(a)$ is a Dirichlet domain. This proves that $D(a)$ is infinite-sided.

We may get a picture of how the domains $\gamma D(a)$, for $\gamma \in \Gamma$, tesselate $\mathbf{E}^3$, as follows. Let r be some large positive number, and let $S_r = \{x \in \mathbf{E}^3 \mid d(x, r) = r\}$ be the surface of a cylinder of radius r about r. Let $\tilde{S}_r$ be the universal cover of S_r. In the induced Riemannian metric, $\tilde{S}_r$ is isometric to $\mathbf{E}^2$. Thus, the tesselation of $\mathbf{E}^3$ determines a CW-decomposition of $\mathbf{E}^2$, invariant under a $\mathbf{Z} \oplus \mathbf{Z}$ action. In the generic situation, this decomposition is combinatorially equivalent to a regular hexagon tesselation of the plane. As we 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, $S_r/\Gamma \approx \mathbf{E}^2/\mathbf{Z} \oplus \mathbf{Z}$. 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 (FIG 4.2). The combinatorial structure of the tesselation $\{\gamma D(a) \mid \gamma \in \Gamma$, far away from r, 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 metric structure of each domain $D(a)$ is also related to rational approximation of θ. Clearly, the area of the cross section $D(a) \cap S_r$ grows linearly with r, but its diameter grows much more quickly in the radial direction than in the direction parallel to r. 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 diameter parallel to r grows like $r^{1/4}$.

Now, we may extend our cyclic group, Γ, to act on $\mathbf{H}^4$ as a parabolic group, with $\mathbf{E}^3 \subset \mathbf{H}^4$ a horosphere about the fixed point p. Let ρ be the 2-plane spanned by r and p. If $a \in \mathbf{H}^4 \setminus \rho$, the Dirichlet domain $D(a)$ will be infinite-sided. (With $p = \infty$ in the upper half-space model, $D(a)$ is a vertical prism on the euclidean Dirichlet domain, D', i.e., $D(a)$ is euclidean-isometric to $D' \times (0, \infty)$.) However, Γ is GF with any of the definitions of Chapter 2.

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 $\mathbf{Z} * \mathbf{Z}$. This gives us a geometrically finite action of $\mathbf{Z} * \mathbf{Z}$ on $\mathbf{H}^4$ with no finite sided Dirichlet domain. This example was constructed by Apanasov.

Remark : In the upper half-space model, we may find a sequence (H_i) of such half-spaces, disjoint in the quotient, M, with $\operatorname{diam}_{euc}(H_i) \to \infty$. We replace each half-space H_i with a copy of $M \setminus H_1$ to give a new manifold M'. On $\mathbf{H}^n$, this gives us a discrete, infinitely generated free group, with no standard horoball about p. I do not know of any finitely generated group with this property. In contrast, we know that the GF groups have standard horoballs.

Let Γ be a GF group. We have seen that there is no reason to expect a Dirichlet domain for Γ to be finite-sided. We say that a p.f.p., p of Γ is *rational* if $\operatorname{stab}_\Gamma(p)$, acting on a standard horosphere, contains a finite-index translation group. Otherwise, we say that p is *irrational*. We shall show that if Γ contains an irrational p.f.p., then $D(a)$ will be infinite-sided if we choose a anywhere on a certain open dense subset of $\mathbf{H}^n$. However, if there are no such p.f.p.s, we show that any convex fundamental domain, P, for Γ will be a finite-sided polyhedron. Here, we use the term "finite-sided" in the sense of Chapter 2, GF3; namely that P should be a finite intersection of half-spaces. It is still possible that P may meet infinitely many images of itself under Γ. If P happens to be a Dirichlet domain, however, it is fairly easy to see that for any $\gamma \in \Gamma$, we must have $P \cap \gamma P = P \cap \sigma$, for some $(n-1)$-plane σ. In other words, the "faces" and "sides" of P coincide. Thus, if Γ contains no irrational p.f.p.s, then each Dirichlet domain has only finitely many faces. In fact, we shall see that in $\mathbf{H}^2$ and $\mathbf{H}^3$, there can be no irrational p.f.p.s, and that each convex fundamental domain has finitely many faces. This will prove the equivalence of 1a, 1b, 2a and 2b in these dimensions.

Below we give (in principle) a complete description of when a Dirichlet domain is finite-sided. We begin by discussing the euclidean case.

Let Γ act discontinuously by isometries on $\mathbf{E}^n$. Suppose the subgroup $\Gamma_i \leq \Gamma$ acts on the plane $\tau_i \leq \mathbf{E}^n$ as a translation group, $i = 1,2$. The group $\Gamma_1 \cap \Gamma_2$ acts as a translation group on $\langle \tau_1, \tau_2 \rangle$. (If γ acts by translation on τ_1, and by translation on τ_2, then the two translations are parallel and have the same translation distance. Hence, γ acts by translation on $\langle \tau_1, \tau_2 \rangle$.) It therefore makes sense to define τ to be the largest plane on which some finite-index subgroup of Γ acts as a translation group. Let $\Gamma_0 \leq \Gamma$ be the subgroup of all elements acting by translation on τ. If $g \in \Gamma$, then $g\Gamma_0 g^{-1}$ is a translation group on $g\tau$. Thus, $g\tau = \tau$, and $g\Gamma_0 g^{-1} = \Gamma_0$, i.e. Γ_0 is a normal subgroup of Γ, and τ is fixed setwise by Γ.

Proposition 4.1 : *Suppose $a \in \mathbf{E}^n$ is not fixed by any element of Γ. Then, the Dirichlet domain $D(a)$ is finite-sided if $a \in \tau$, and infinite-sided if $a \notin \tau$.*

Proof : Suppose $a \in \tau$. Then, $D(a)$ is a euclidean product, with an orthogonal plane, of the Dirichlet domain $D(a) \cap \tau$ of Γ restricted to τ. On this subspace, Γ has a finite-index translation group, so that any convex fundamental domain is finite-sided (see Lemma 4.2 below).

Suppose $a \notin \tau$. Let μ be a minimal Γ-invariant affine subspace (see Chapter 1). Note that μ is a subspace of τ. Let b be the nearest point in μ to a. Let l be the ray from b through a, and let σ be the plane $\langle \mu, a \rangle$ (FIG 4.3). We have $l \subseteq D(a)$. To see this, take any $c \in l$. The images of a under Γ all lie a fixed distance from μ. It follows that the nearest image to c must be a itself, i.e. $c \in D(a)$.

Suppose $D(a)$ were finite-sided, $D(a) = \bigcap_{\gamma \in G} H_\gamma$, where G is a finite subset of Γ. Let $\Gamma_1 \leq \Gamma$ be the subgroup of Γ that fixes the plane σ setwise, and preserves the direction of l. (One can see that Γ_1 is defined independently of the choice of μ, though this is not important for our discussion.) By maximality of τ, we must have $|\Gamma : \Gamma_1| = \infty$.

If Γ_1 were trivial, the proof could proceed as in the example of an irrational screw motion described above. $D(a)$ would contain a cone about l, and, since the action of Γ on the sphere of rays is not discrete ($|\Gamma : \Gamma_1| = \infty$), we could find $\gamma \in \Gamma$ with γl arbitrarily close (in direction) to l, so that $\gamma D(a) \cap D(a) \neq \emptyset$.

To deal with the general case, we write $D(a) = D^1(a) \cap D^2(a)$ with $D^1(a) = \bigcap_{\gamma \in (\Gamma_1 \cap G)} H_\gamma$ and $D^2(a) = \bigcap_{\gamma \in (G \setminus \Gamma_1)} H_\gamma$. As before, $D^2(a)$ contains a cone about l, which we can take to be an open spherical cone C centred on l. Now, $D^1(a)$ contains $\bigcap_{\gamma \in \Gamma_1} H_\gamma = D'$, where D' is the Dirichlet domain about a for the group Γ_1.

Since σ is preserved by Γ_1, D' is a euclidean product, with an orthogonal plane, of the Dirichlet domain $D' \cap \sigma$, of $\Gamma_1 | \sigma$ about a. Since Γ_1 fixes the direction l, we see that $D' \cap \sigma$, in turn, is a euclidean product, with the line $\langle l \rangle$, of the Dirichlet domain D'' about b for $\Gamma_1 | \mu$.

As before, Γ acts as a non-discrete group on the sphere of rays. So, we can find $\gamma \in \Gamma \setminus \Gamma_1$ with $l \cap \gamma C \neq \emptyset$ (so that $l \cap \gamma C$ is an infinite ray). For some $g \in \Gamma_1$, we have $\gamma^{-1} b \in gD''$. Now, $g^{-1}\gamma^{-1} l$ is a ray, orthogonal to μ, emanating from the point $g^{-1}\gamma^{-1} b$ of D''. From the previous paragraph, we see that $g^{-1}\gamma^{-1} l$ lies entirely within D'. Hence, we see that $l \subseteq \gamma g D' \subseteq \gamma g D^1(a)$.

Now $\gamma^{-1} l \cap C \neq \emptyset$. Since C is a spherical cone about l, and gl is parallel to l, we see that gC is just a translate of C. Thus $\gamma^{-1} l \cap gC \neq \emptyset$, and so $l \cap \gamma g D^2(a) \neq \emptyset$.

We have shown that l intersects both $\gamma g D^1(a)$ and $\gamma g D^2(a)$. So, $\gamma g D(a)$ meets the interior of $D(a)$. But $\gamma g \neq 1$, contradicting the assumption that $D(a)$ is a Dirichlet domain.

◇

One may easily generalise the above proposition to the following.

A collection of generalised Dirichlet domains $\{D_i(a)\}$ contains at least one infinite-sided member if and only if some basepoint $a_i \notin \tau$.

We now want to give a description of when Dirichlet domains for hyperbolic groups are finite-sided. Given a Dirichlet domain, $D(a)$, we need to describe which of the p.f.p.s are contained in the closure, $(D(a))_C$, in $\mathbf{H}^n_C$. To do this, we introduce a construction analogous to that of Dirichlet domains, for p.f.p.s.

Let Γ be a discrete group, and let $P_0 \subset \Lambda$ be an orbit of p.f.p.s. We choose a horoball $B(p)$ about each $p \in P_0$, so that the collection $\{B(p) | p \in P_0\}$ is strictly invariant. Given $p \in P_0$, we define $U(p)$ to be the set of points nearer to $B(p)$ than to any other horoball; i.e. $U(p) = \{x \in \mathbf{H}^n | d(x, B(p)) < d(x, B(q)) \text{ for all } q \neq p\}$. Let $\mathcal{U} = \{U(p) \mid p \in P_0\}$. Since $B(\gamma p) = \gamma B(p)$, the collection $\{B(p) \mid p \in P_0\}$ is determined by the choice of just one horoball $B(p)$. It is easily seen that any choice of $B(p)$ will give rise to the same collection $\mathcal{U}$. We also see that $\mathcal{U}$ is locally finite, and that $\bigcup \mathcal{U}$ is dense in $\mathbf{H}^n$.

Suppose that $a \in \mathbf{H}^n$ is not fixed by any $\gamma \in \Gamma$. Then, we see that $(D(a))_C$ contains the point p.f.p. p, if and only if a lies in the closure $\bar{U}(p)$ of $U(p)$. This means that $(D(a))_C \cap P_0 = \{p \in P_0 \mid a \in \bar{U}(p)\}$. In particular, we see that $(D(a))_C \cap P_0 = \{p\}$ if and only if $a \in U(p)$. From this, we see that $(D(a))_C$ contains only finitely many p.f.p.s in a given orbit, and generically contains only one from each orbit.

Now, the collection $\mathcal{U}$ satisfies all the conditions (Chapter 2, GF3, (a)–(d)) to be the set of top-dimensional cells for a convex cell complex. We write $\mathcal{B}_\Gamma(P_0)$ for unique such complex which is minimal with respect to subdivision (see the construction in Chapter 2, GF3). More generally, suppose that $P \subset \Lambda$ consists of finitely many orbits of p.f.p.s, $P = \bigcup_{i=1}^k P_i$. From Chapter 2, GF3, we see that the complexes $\mathcal{B}_\Gamma(P_i)$ have a minimal common subdivision, $\mathcal{B}_\Gamma(P) = (\mathcal{B}_\Gamma(P_i) \mid i = 1, \dots, k)$, obtained by intersecting cells from each complex.

If Γ acts freely, we may describe the cells of $\mathcal{B}_\Gamma(P)$ as follows. Let $Q \subset P$ be finite, and let $A(Q) = \{a \in \mathbf{H}^n \mid (D(a))_C \cap P = Q\}$. Then $\mathcal{B}_\Gamma(P)$ is the set of all $A(Q)$ as Q ranges over all finite subsets of P.

Suppose now that Γ is GF. Then there are only finitely many orbits of p.f.p., so we may let P be the set of all parabolic fixed points. In this case, we write $\mathcal{B}_\Gamma$ for $\mathcal{B}_\Gamma(P)$. Note that if for any Dirichlet domain, we have $(D(a))_C \cap \Lambda = (D(a))_C \cap P$.

We are now in a position to describe when a Dirichlet domain for a hyperbolic group is finite-sided. This is only possible when the group is GF. Let Γ be a GF group, and let $P \subseteq \Lambda$ be the set of all p.f.p.s. To each $p \in P$, we may associate a unique plane $\rho(p)$ through p, which is maximal with the property that some finite-index subgroup of $\mathrm{stab}_\Gamma(p)$ act as a translation group on $\rho(p) \cap \partial B$, for some (and hence each) horoball B about p.

Suppose now that the point a is not fixed by any $\gamma \in \Gamma$. Suppose that $a \in U(p)$, so that $(D(a))_C$ contains the p.f.p. p, but no other point in the orbit of p. Let $p = \infty$ in $\mathbf{H}^n_C$, and let B be a horoball about p. Let l be the ray joining a to p. If we choose ∂B high enough, we see that $D(a) \cap \partial B$ is the Dirichlet domain, about the point $l \cap \partial B$, for the action of $\mathrm{stab}_\Gamma(p)$ on ∂B. In fact, $D(a) \cap B$ is then a vertical prism on $D(a) \cap \partial B$; i.e. it is euclidean-isometric to $(D(a) \cap \partial B) \times [0, \infty)$. Moreover, it is fairly easy to see that the images of $D(a)$ under $\mathrm{stab}_\Gamma(p)$ cover some standard parabolic region $C(p)$ about p. From Proposition 4.1, we see that $C(p)$ meets $D(a)$ in only finitely many sides if and only if $a \in \rho(p)$.

Suppose now, that a lies in a top dimensional cell of the complex $\mathcal{B}_\Gamma$. This means that $(D(a))_C$ meets Λ in an orbit-transversal, Q, of p.f.p.s. If we take a standard parabolic regions $C(p)$ about each $p \in Q$, as above, we see that $D(a) \setminus \bigcup_{p \in Q} C(p)$ is relatively compact. We deduce that $D(a)$ is finite-sided if and only if $a \in \bigcap_{p \in Q} \rho(p)$. If Γ contains at least one irrational p.f.p., then $\bigcap_{p \in Q} \rho(p)$ is a proper subspace of $\mathbf{H}^n$. So, in this case, we see that the set of a for which $D(a)$ is infinite-sided contains an open dense subset of $\mathbf{H}^n$. (In fact we may find a convex cell complex so that the set of a with $D(a)$ finite-sided lies in the $(n-1)$-skeleton.)

Using the generalisation of Proposition 4.1 stated above, we see that if a lies in a lower-dimensional cell of $\mathcal{B}_\Gamma$, we see that $D(a)$ is finite-sided if and only if $a \in \bigcap \{\rho(p) \mid p \in (D(a))_C \cap \Lambda\}$. Note that the set $(D(a))_C \cap \Lambda$ is determined by the cell of $\mathcal{B}_\Gamma$ in which a lies.

Finally, we say a few things about general convex fundamental domains.

Let Γ act discontinuously on $\mathbf{H}^n$. Let $X^1, \dots, X^k$ be a collection of disjoint open convex subsets of $\mathbf{H}^n$. Suppose that, as $\gamma \in \Gamma$ and $i = 1, \dots, k$ vary, the sets γX^i are disjoint, locally finite, and their closures cover $\mathbf{H}^n$. This means that $\mathcal{U} = \{\gamma X^i \mid \gamma \in \Gamma,\ i = 1, \dots, k\}$ satisfies all the criteria (Chapter 2, GF3, (a)–(d)) to be the set of top-dimensional cells of some cell complex $\mathcal{A}$, which we take to be minimal with respect to subdivision. The argument in Chapter 3, GF3 $\Rightarrow$ GF1, (applied to top-dimensional cells) shows that $\mathcal{A}$ is locally finite on $\mathbf{H}^n \cup \Omega$. This means that, if we write X^i_C for the closure of X^i in $\mathbf{H}^n_C$, then $\bigcup X^i_C \setminus \Lambda$ is a fundamental domain for the action of Γ on $\mathbf{H}^n \cup \Omega$.

Suppose now, that Γ is GF. By local finiteness, none of the sets X_C^i can meet a c.l.p. So, each $X_C^i \cap \Lambda$ consists of only (bounded) p.f.p.s. We show that each $X_C^i \cap \Lambda$ is finite, and, moreover, that we can assume that the only standard parabolic regions $C(p)$ that meet X_C^i are those corresponding to $p \in X_C^i \cap \Lambda$.

The argument is similar to that for local finiteness on $\mathbf{H}^n \cup \Omega$. Let C_1, C_2, C_3 be three successively smaller standard parabolic regions about p. Any set X^i that meets both ∂C_1 and ∂C_3, meets also $\partial C_2 \cap \operatorname{hull}(\partial C_1 \cup \partial C_3)$, which has compact quotient under $\operatorname{stab}_\Gamma p$. Thus $\{\gamma \in \Gamma \mid \gamma X^i \cap C_3 \neq \emptyset\}$ represents only finitely many cosets of the form $(\operatorname{stab}_\Gamma p)\gamma$, i.e. X_C^i meets only finitely many elements in the orbit of p. Shrinking C_3 further to C, we can assume that any γX^j with $\gamma X_C^j \cap C \neq \emptyset$ has $p \in \gamma X_C^j$.

Let $p = \infty$ in $\mathbf{R}^n_+$, and B, a horoball contained in C. We must have that each $X^i \cap B$ is a vertical prism on $X^i \cap \partial B$. We show below that if p is rational, $X_C^i \cap \partial B$ is finite-sided. (It is possible however that $X_C^i \cap \partial B$ meets infinitely many other $\gamma X_C^j \cap \partial B$ - recall the distinction between "faces" and "sides" made in Chapter 2, GF3.) It follows that if Γ is *rational*, (i.e. every p.f.p. is rational), then each X^i is finite-sided. From this we shall be able to deduce the equivalence, for $n \leq 3$, of definitions 1a, 1b, 2a and 2b, mentioned at the start of this chapter.

Lemma 4.2 : *Let Γ be a discrete group of translations acting on $\mathbf{E}^n$. Suppose that the open convex sets $X_1, \dots, X_k$ together constitute a fundamental domain for Γ. Then, each X_i is (the interior of) a finite sided polyhedron.*

(Note that the orbit of a convex set under a discrete euclidean group is necessarily locally finite, if the sets in the orbit are all disjoint.)

Proof : We know that Γ is a free abelian group. Let $\{g_1, \dots, g_r\}$ be a free set of generators. Let $\Gamma' < \Gamma$ be the subgroup $< 2g_1, \dots, 2g_r >$, so that $|\Gamma : \Gamma'| = 2^r$. The construction of the convex cell complex from $\{\gamma X_i\}$, enables us to define the set $\mathcal{F}^{n-1}$ of codimension-1 faces of X_i. Each $A \in \mathcal{F}^{n-1}$ corresponds to some γX_j with $\bar{A} = \gamma \bar{X}_j \cap \bar{X}_i$. We label A by the pair $(j, [\gamma])$, where $[\gamma]$ is the coset of Γ' containing γ. We claim that if A and B in $\mathcal{F}^{n-1}$ have the same label, then they lie in the same codimension-1 plane of $\mathbf{E}^n$. We have $\bar{A} = \bar{X}_i \cap \gamma_1 \bar{X}_j$, $\bar{B} = \bar{X}_i \cap \gamma_2 \bar{X}_j$, with γ_1 and γ_2 differing by twice some translation $g \in \Gamma$, that is, $\gamma_2 - \gamma_1 = 2g$. Let $a \in A$, $b \in B$. If A and B do not lie in the same plane, then the midpoint $c = (a+b)/2$ lies in $\operatorname{int}\operatorname{hull}(A \cup B) \subset X_i$. However, for some $u, v \in \bar{X}_j$, we also have $c = (\gamma_1 u + \gamma_2 v)/2 = \gamma_1((u+v)/2) + gv \in (\gamma_1 + g)\bar{X}_j$, by convexity of X_j. This gives us the contradiction $X_i \cap (\gamma_1 + g)\bar{X}_j \neq \emptyset$.

◇

The tesselation of $\mathbf{E}^3$ with square prisms described in Chapter 2, GF3, gives us an example where the X_i do not have a finite number of faces. Note that the tesselation is invariant under a $\mathbf{Z} \oplus \mathbf{Z}$ action, acting vertically and in the NW-SE direction. The problem arises because each tile meets infinitely many images of itself under the $\mathbf{Z} \oplus \mathbf{Z}$ action. However, this phenomenon cannot occur in euclidean space of dimension less than 3. The only case where we get a non-compact quotient is for an infinite cyclic action on $\mathbf{E}^2$. In this case, it is fairly easy to see that if we have a tesselation with finite quotient, then each tile meets only finitely many other tiles. In fact we may classify such tiles according to whether they are compact, or have one or two topological ends (FIG 4.4).

Now, any isometry of $\mathbf{E}^1$, or $\mathbf{E}^2$, with no fixed point, must be a translation. Thus, any discrete group action on these spaces must have a finite-index translation subgroup. We see that any discrete subgroup of $\operatorname{Isom}\mathbf{H}^2$, or of $\operatorname{Isom}\mathbf{H}^3$, can have only rational p.f.p.s. From this, we deduce the equivalence, in these dimensions, of the four descripions of geometric finiteness stated at the start of the chapter.

The question remains of whether or not a GF group necessarily has a single, convex finite-sided (or finite-faced) fundamental polyhedron. I suspect not, but I do not have a counterexample.

References.

[Ahl] 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, Birkhauser (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.

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

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

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

[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.

[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.

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

[G] 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.

[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.

[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 numbers* : 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.Seiama, 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 Fatou-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.

[SuS] P.D.Susskind, G.A.Swarup, *Limit sets of geometrically finite hyperbolic groups* : preprint.

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

[Th2] 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).

Figure 0.1 :

2-dimensional cusp

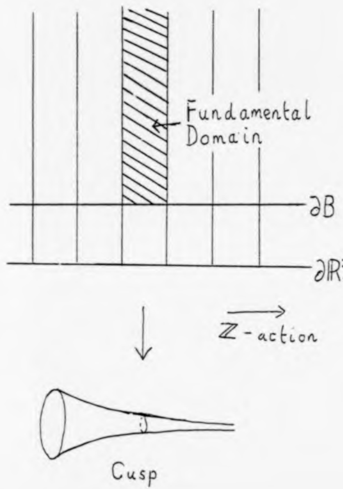

Figure 0.2 :

2-dimensional funnel

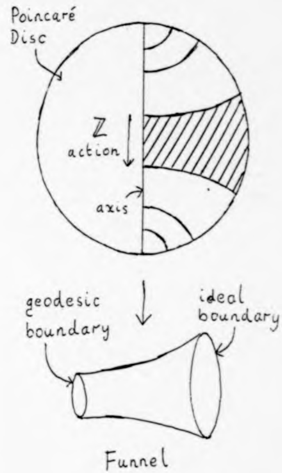

Figure 0.3 :
Geometrically finite
2-manifold

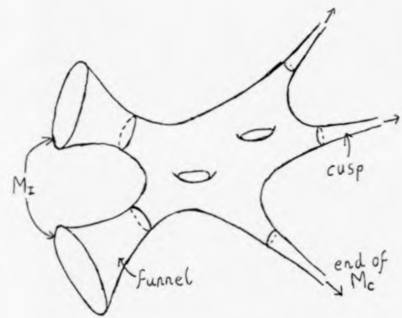

Figure 0.4 :
$\mathbb{Z}$-cusp

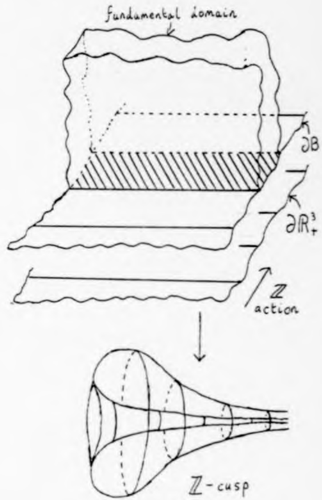

Figure 0.5 :

$\mathbb{Z}\oplus\mathbb{Z}$ - cusp

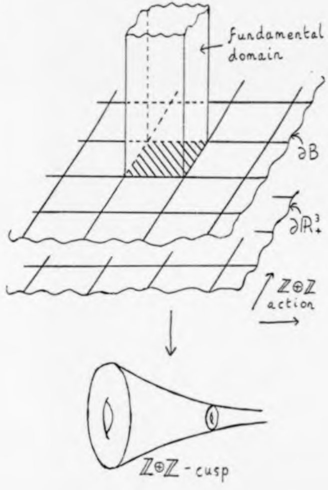

Figure 0.6 :

Cross-section of $\mathbb{Z}$-cusp showing intersection with $N_\eta \hat{Y}$.
Upper half-space model.

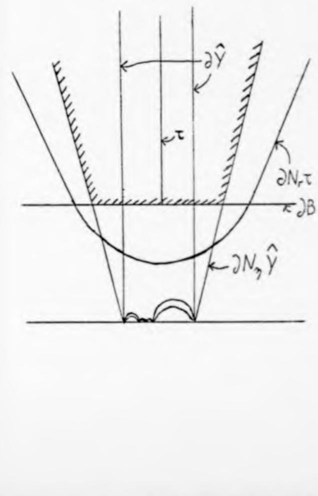

Figure 0.7 :
Standard cusp regions

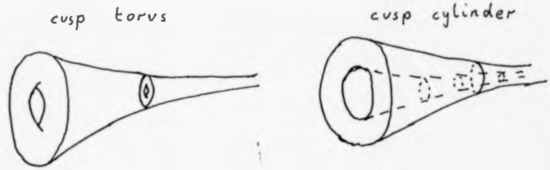

Figure 2.1 :
A standard parabolic region in the upper half-space model.
$n=3$, $\Gamma_p \cong \mathbb{Z}$

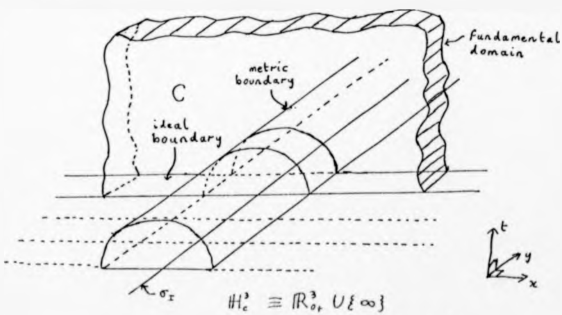

$\mathbb{H}^3_c \equiv \mathbb{R}^3_{0+} \cup \{\infty\}$

$\Gamma_p = \langle (x,y,t) \mapsto (x,y+1,t) \rangle$

Figure 2.2:
A standard cusp – the quotient from
Figure 2.

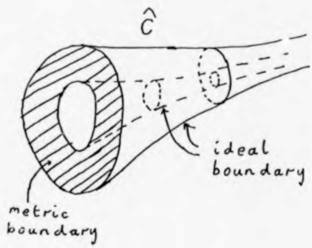

Figure 2.3:
A GFI 2-manifold.

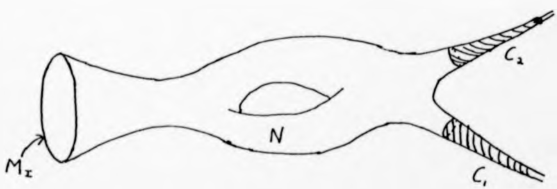

$M_c = N \cup C_1 \cup C_2$

Figure 2.4:

A bounded parabolic fixed point in the boundary of the upper half-space model – $n = 4$, $\Gamma_p \cong \mathbb{Z}$

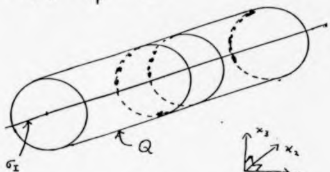

$$\mathbb{R}^3 \equiv \partial \mathbb{R}^4_+$$

$$\Gamma_p = \langle (x_1, x_2, x_3, x_4) \mapsto (x_1, x_2+1, x_3, x_4) \rangle$$

$$\Lambda \subseteq Q \cup \{\infty\}$$

Figure 2.5:

A conical limit point – upper half-space model.

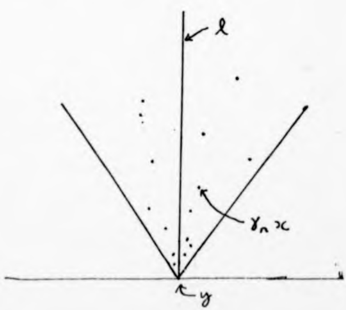

$$\gamma_n x \to y$$

Figure 2.6 :
A convex cell complex in $\mathbb{E}^3$.

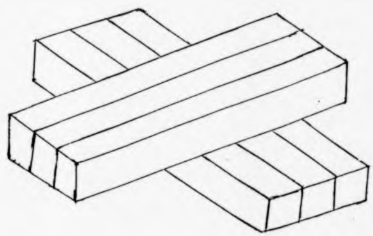

Figure 2.7 :
The convex core of a 2-manifold.

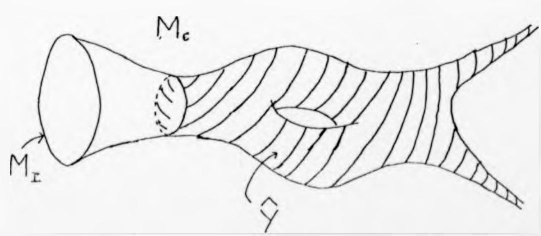

Figure 3.1 :
Schematic.

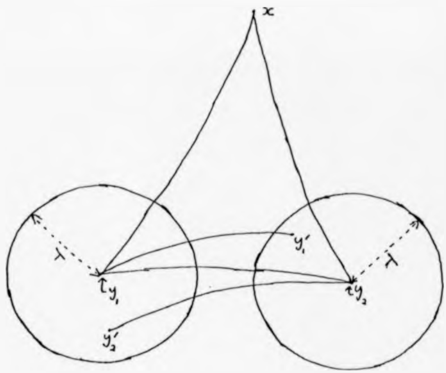

Figure 3.2 :
Upper half space model.

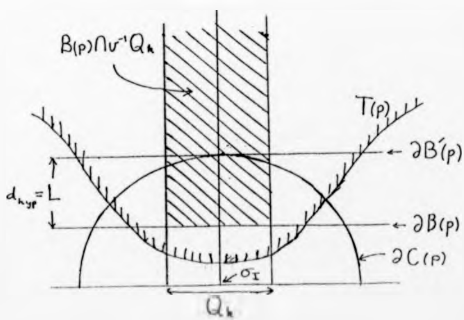

Figure 3,3 :
Generalised dirichlet domains on a 2-manifold.

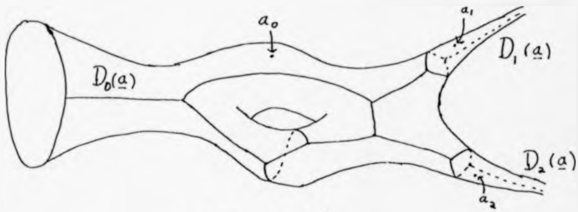

$\underline{a} = (a_0, a_1, a_2)$

Figure 3.4 :
The function $\omega(K,-)$ for $K \subseteq \mathbb{E}^2$.

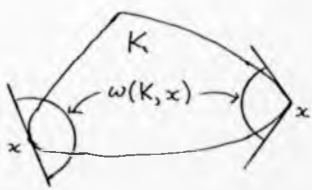

Figure 3.5 :
Upper half space model —
$n = 2$.

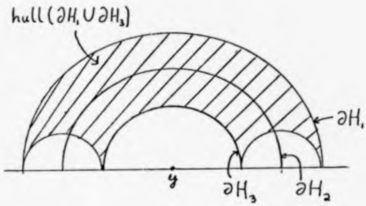

Figure 3.6 :
Schematic.

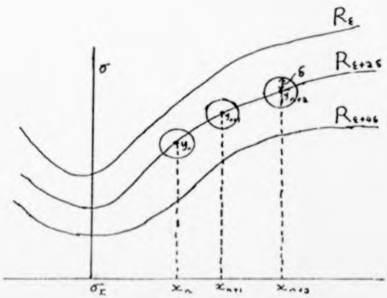

Figure 4.1 :
Irrational screw motion on $\mathbb{E}^3$

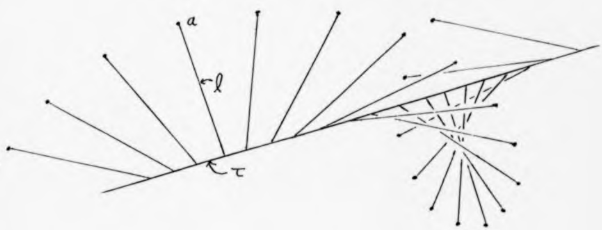

Figure 4.2 :
Whitehead move on a torus

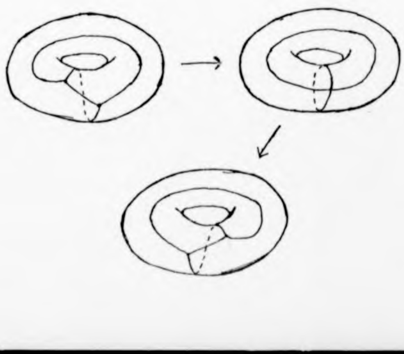

Figure 4.3 :
(Schematic)

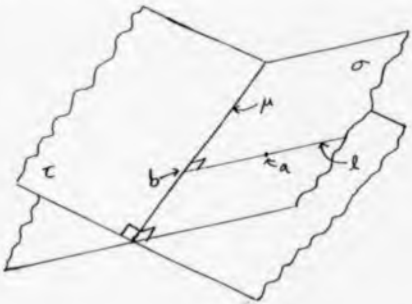

Figure 4.4 :
Two singly-periodic tilings of $\mathbb{E}^2$

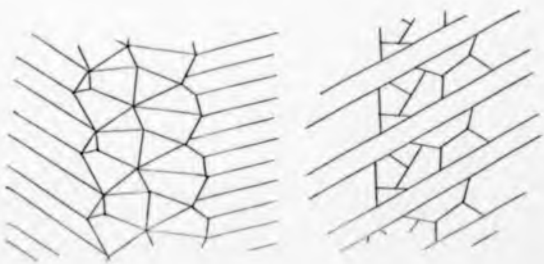

TITLE Geometrical Finiteness for Hyperbolic Groups

AUTHOR Brian Hayward Bowditch

INSTITUTION and DATE University of Warwick, 1988

D87395