William P. Thurston

The Geometry and Topology of Three-Manifolds

Electronic version 1.1 - March 2002
http://www.msri.org/publications/books/gt3m/

This is an electronic edition of the 1980 notes distributed by Princeton University. The text was typed in TEX by Sheila Newbery, who also scanned the figures. Typos have been corrected (and probably others introduced), but otherwise no attempt has been made to update the contents. Genevieve Walsh compiled the index.

Numbers on the right margin correspond to the original edition's page numbers.

Thurston's *Three-Dimensional Geometry and Topology*, Vol. 1 (Princeton University Press, 1997) is a considerable expansion of the first few chapters of these notes. Later chapters have not yet appeared in book form.

Please send corrections to Silvio Levy at levy@msri.org.

CHAPTER 8

Kleinian groups

Our discussion so far has centered on hyperbolic manifolds which are closed, or at least complete with finite volume. The theory of complete hyperbolic manifolds with infinite volume takes on a somewhat different character. Such manifolds occur very naturally as covering spaces of closed manifolds. They also arise in the study of hyperbolic structures on compact three-manifolds whose boundary has negative Euler characteristic. We will study such manifolds by passing back and forth between the manifold and the action of its fundamental group on the disk. 8.1

8.1. The limit set

Let Γ be any discrete group of orientation-preserving isometries of H^n. If $x \in H^n$ is any point, the *limit set* $L_\Gamma \subset S_\infty^{n-1}$ is defined to be the set of accumulation points of the orbit Γ_x of x. One readily sees that L_Γ is independent of the choice of x by picturing the Poincaré disk model. If $y \in H^n$ is any other point and if $\{\gamma_i\}$ is a sequence of elements of Γ such that $\{\gamma_i x\}$ converges to a point on S_∞^{n-1}, the hyperbolic distance $d(\gamma_i x, \gamma_i y)$ is constant so the Euclidean distance goes to 0; hence $\lim \gamma_i y = \lim \gamma_i x$.

The group Γ is called *elementary* if the limit set consists of $0, 1$ or 2 points.

PROPOSITION 8.1.1. *Γ is elementary if and only if Γ has an abelian subgroup of finite index.* $\square$

When Γ is not elementary, then L_Γ is also the limit set of any orbit on the sphere at infinity. Another way to put it is this:

PROPOSITION 8.1.2. *If Γ is not elementary, then every non-empty closed subset of S_∞ invariant by Γ contains L_Γ.*

PROOF. Let $K \subset S_\infty$ be any closed set invariant by Γ. Since Γ is not elementary, K contains more than one element. Consider the projective (Klein) model for H^n, and let $H(K)$ denote the convex hull of K. $H(K)$ may be regarded either as the Euclidean convex hull, or equivalently, as the hyperbolic convex hull in the sense that it is the intersection of all hyperbolic half-spaces whose "intersection" with S_∞ contains K. Clearly $H(K) \cap S_\infty = K$.

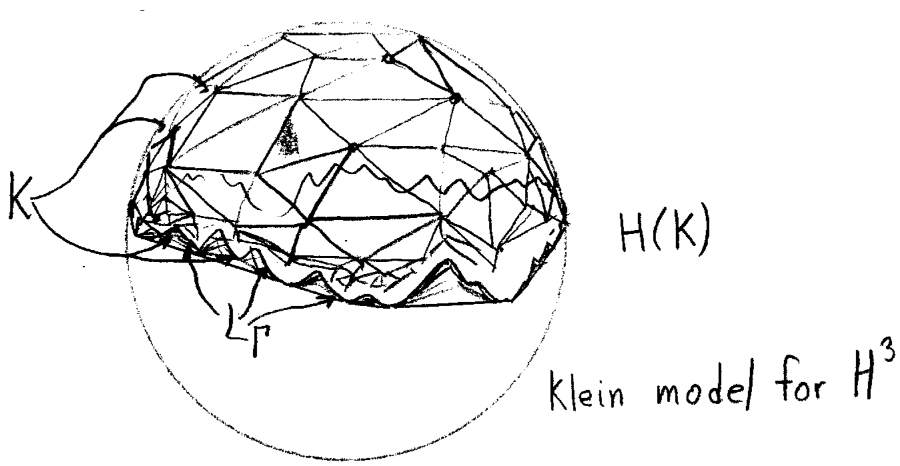

Since K is invariant by Γ, $H(K)$ is also invariant by Γ. If x is any point in $H^n \cap H(K)$, the limit set of the orbit Γ_x must be contained in the closed set $H(K)$. Therefore $L_\Gamma \subset K$. $\square$

8.3

A closed set K invariant by a group Γ which contains no smaller closed invariant set is called a minimal set. It is easy to show, by Zorn's lemma, that a closed invariant set always contains at least one minimal set. It is remarkable that in the present situation, L_Γ is the *unique* minimal set for Γ.

Corollary 8.1.3. *If Γ is a non-elementary group and $1 \neq \Gamma' \triangleleft \Gamma$ is a normal subgroup, then $L_{\Gamma'} = L_\Gamma$.*

Proof. An element of Γ conjugates Γ' to itself, hence it takes $L_{\Gamma'}$ to $L_{\Gamma'}$. Γ' must be infinite, otherwise Γ' would have a fixed point in H^n which would be invariant by Γ so Γ would be finite. It follows from 8.1.2 that $L_{\Gamma'} \supset L_\Gamma$. The opposite inclusion is immediate. $\square$

Examples. If M^2 is a hyperbolic surface, we may regular $\pi_1(M)$ as a group of isometries of a hyperbolic plane in H^3. The limit set is a circle. A group with limit set contained in a geometric circle is called a *Fuchsian group*.

The limit set for a closed hyperbolic manifold is the entire sphere S^{n-1}_∞.

If M^3 is a closed hyperbolic three-manifold which fibers over the circle, then the fundamental group of the fiber is a normal subgroup, hence its limit set is the entire sphere. For instance, the figure eight knot complement has fundamental group $\langle A, B : ABA^{-1}BA = BAB^{-1}AB \rangle$. 8.4

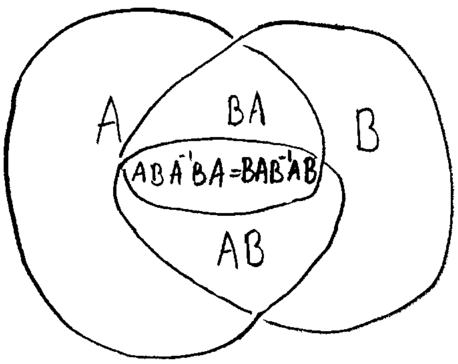

It fibers over S^1 with fiber F a punctured torus. The fundamental group $\pi_1(F)$ is the commutator subgroup, generated by AB^{-1} and $A^{-1}B$. Thus, the limit set of a finitely generated group may be all of S^2 even when the quotient space does not have finite volume.

A more typical example of a free group action is a Schottky group, whose limit set is a Cantor set. Examples of Schottky groups may be obtained by considering H^n minus $2k$ disjoint half-spaces, bounded by hyperplanes. If we choose isometric identifications between pairs of the bounding hyperplanes, we obtain a complete hyperbolic manifold with fundamental group the free group on k generators.

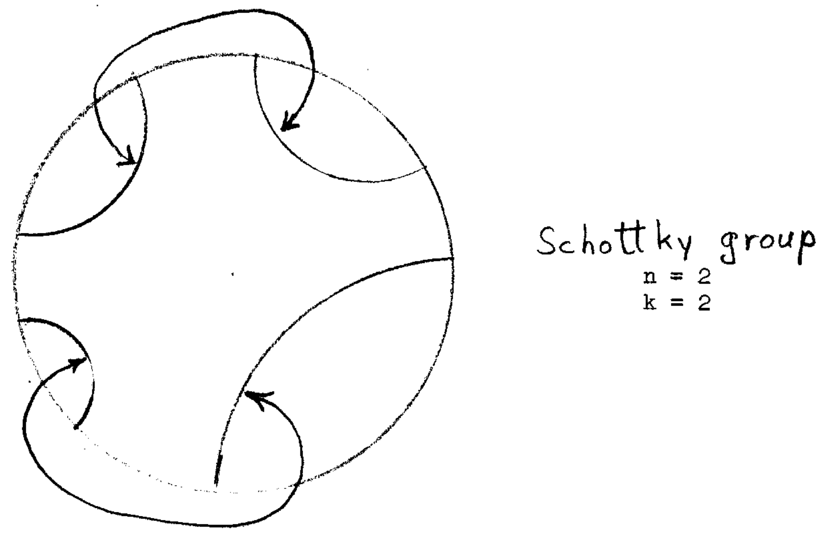

8.5

It is easy to see that the limit set for the group of covering transformations is a Cantor set.

8.2. The domain of discontinuity

The domain of discontinuity for a discrete group Γ is defined to be $D_\Gamma = S^{n-1}_\infty - L_\Gamma$. A discrete subgroup of $\mathrm{PSL}(2,\mathbb{C})$ whose domain of discontinuity is non-empty is called a *Kleinian group*. (There are actually two ways in which the term Kleinian group is generally used. Some people refer to any discrete subgroup of $\mathrm{PSL}(2,\mathbb{C})$ as a Kleinian group, and then distinguish between a type I group, for which $L_\Gamma = S^2_\infty$, and a type II group, where $D_\Gamma \neq \emptyset$. As a field of mathematics, it makes sense for Kleinian groups to cover both cases, but as mathematical objects it seems useful to have a word to distinguish between these cases $D_\Gamma \neq \emptyset$ and $D_\Gamma = \emptyset$.)

We have seen that the action of Γ on L_Γ is minimal—it mixes up L_Γ as much as possible. In contrast, the action of Γ on D_Γ is as discrete as possible.

Definition 8.2.1. If Γ is a group acting on a locally compact space X, the action is *properly discontinuous* if for every compact set $K \subset X$, there are only finitely many $\gamma \in \Gamma$ such that $\gamma K \cap K \neq \emptyset$.

Another way to put this is to say that for any compact set K, the map $\Gamma \times K \to X$ given by the action is a proper map, where Γ has the discrete topology. (Otherwise 8.6
there would be a compact set K' such that the preimage of K' is non-compact. Then infinitely many elements of Γ would carry $K \cup K'$ to itself.)

Proposition 8.2.2. *If Γ acts properly discontinuously on the locally compact Hausdorff space X, then the quotient space X is Hausdorff. If the action is free, the quotient map $X \to X/\Gamma$ is a covering projection.*

Proof. Let $x_1, x_2 \in X$ be points on distinct orbits of Γ. Let N_1 be a compact neighborhood of x_1. Finitely many translates of x_2 intersect N_1, so we may assume N_1 is disjoint from the orbit of x_2. Then $\bigcup_{\gamma\in\Gamma} \gamma N_1$ gives an invariant neighborhood of x_1 disjoint from x_2. Similarly, x_2 has an invariant neighborhood N_2 disjoint from N_1; this shows that X/Γ is Hausdorff. If the action of Γ is free, we may find, again by a similar argument, a neighborhood of any point x which is disjoint from all its translates. This neighborhood projects homeomorphically to X/Γ. Since Γ acts transitively on the sheets of X over X/Γ, it is immediate that the projection $X \to X/\Gamma$ is an even covering, hence a covering space. $\square$

Proposition 8.2.3. *If Γ is a discrete group of isometries of H^n, the action of Γ on D_Γ (and in fact on $H^n \cup D^\Gamma$) is properly discontinuous.*

PROOF. Consider the convex hull $H(L_\Gamma)$. There is a retraction r of the ball $H^n \cup S_\infty$ to $H(L_\Gamma)$ defined as follows. 8.7

If $x \in H(L_\Gamma)$, $r(x) = x$. Otherwise, map x to the nearest point of $H(L_\Gamma)$. If x is an infinite point in D_Γ, the nearest point is interpreted to be the first point of $H(L_\Gamma)$ where a horosphere "centered" about x touches L_Γ. This point $r(x)$ is always uniquely defined

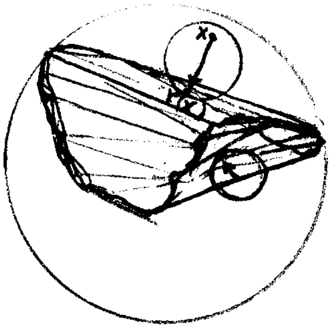

because $H(L_\Gamma)$ is convex, and spheres or horospheres about a point in the ball are strictly convex. Clearly r is a proper map of $H^n \cup D_\Gamma$ to $H(L_\Gamma) - L_\Gamma$. The action of Γ on $H(L_\Gamma) - L_\Gamma$ is obviously properly discontinuous, since Γ is a discrete group of isometries of $H(L_\Gamma) - L_\Gamma$; the property of $H^n \cup D_\Gamma$ follows immediately. □

REMARK. This proof doesn't work for certain elementary groups; we will ignore such technicalities.

It is both easy and common to confuse the definition of properly discontinuous with other similar properties. To give two examples, one might make these definitions:

DEFINITION 8.2.4. The action of Γ is *wandering* if every point has a neighborhood N such that only finitely many translates of N intersect N.

DEFINITION 8.2.5. The action of Γ has *discrete orbits* if every orbit of Γ has an empty limit set.

PROPOSITION 8.2.6. *If Γ is a free, wandering action on a Hausdorff space X, the projection $X \to X/\Gamma$ is a covering projection.*

PROOF. An exercise. □

WARNING. Even when X is a manifold, X/Γ may not be Hausdorff. For instance, consider the map

$$L : \mathbb{R}^2 - 0 \to \mathbb{R}^2 - 0$$

$$L\,(x, y) = (2x, \tfrac{1}{2}y).$$

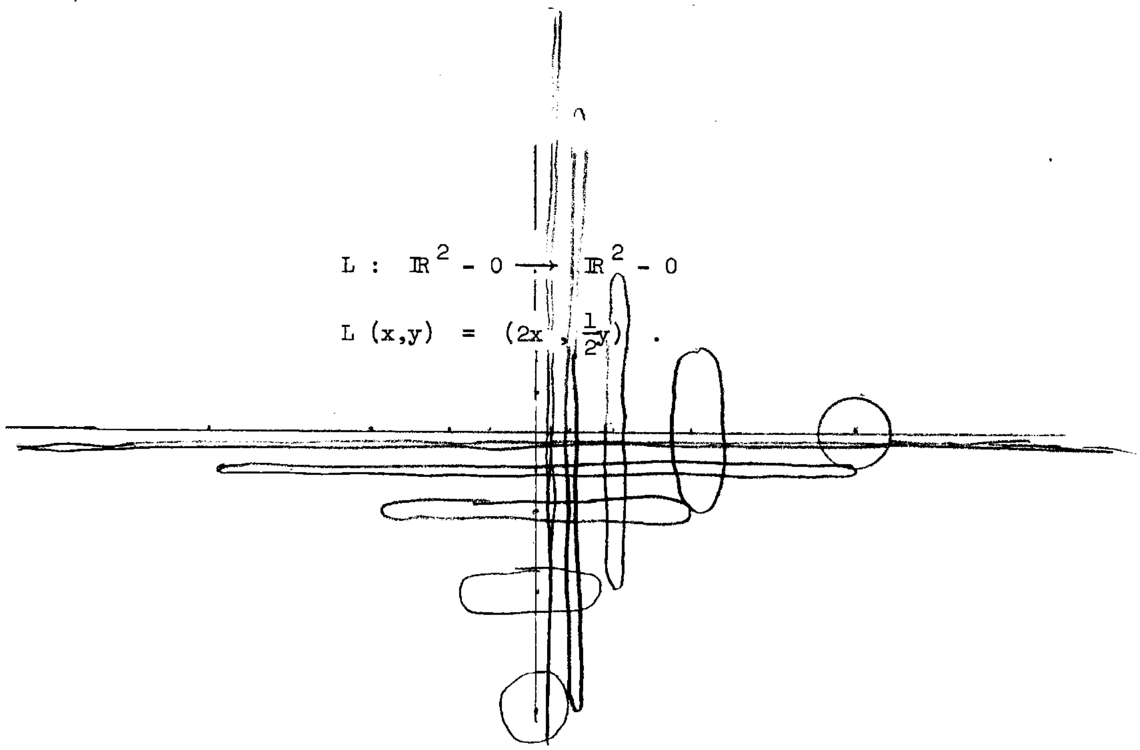

It is easy to see this is a wandering action. The quotient space is a surface with fundamental group $\mathbb{Z} \oplus \mathbb{Z}$. The surface is non-Hausdorff, however, since points such 8.9
as $(1, 0)$ and $(0, 1)$ do not have disjoint neighborhoods.

Such examples arise commonly and naturally; it is wise to be aware of this phenomenon.

The property that Γ has discrete orbits simply means that for every pair of points x, y in the quotient space X/Γ, x has a neighborhood disjoint from y. This can occur, for instance, in a l-parameter family of Kleinian groups Γ_t, $t \in [0, 1]$. There are examples where $\Gamma_t = \mathbb{Z}$, and the family defines the action of $\mathbb{Z}$ on $[0, 1] \times H^3$ with discrete orbits which is not a wandering action. See § . It is remarkable that the action of a Kleinian group on the set of all points with discrete orbits is properly discontinuous.

8.3. Convex hyperbolic manifolds

The limit set of a group action is determined by a limiting process, so that it is often hard to "know" the limit set directly. The condition that a given group action is discrete involves infinitely many group elements, so it is difficult to verify directly. Thus it is important to have a concrete object, satisfying concrete conditions, corresponding to a discrete group action.

We consider for the present only groups acting freely.

DEFINITION 8.3.1. A complete hyperbolic manifold M with boundary is *convex* if every path in M is homotopic (rel endpoints) to a geodesic arc. (The degenerate 8.10
case of an arc which is a single point may occur.)

PROPOSITION 8.3.2. *A complete hyperbolic manifold M is convex if and only if the developing map $D : \tilde{M} \to H^n$ is a homeomorphism to a convex subset of H^n.*

PROOF. If $\tilde{M}$ is a convex subset S of H^n, then it is clear that M is convex, since any path in M lifts to a path in S, which is homotopic to a geodesic arc in S, hence in M.

If M is convex, then D is $1-1$, since any two points in $\tilde{M}$ may be joined by a path, which is homotopic in M and hence in $\tilde{M}$ to a geodesic arc. D must take the endpoints of a geodesic arc to distinct points. $D(\tilde{M})$ is clearly convex. □

We need also a local criterion for M to be convex. We can define M to be *locally convex* if each point

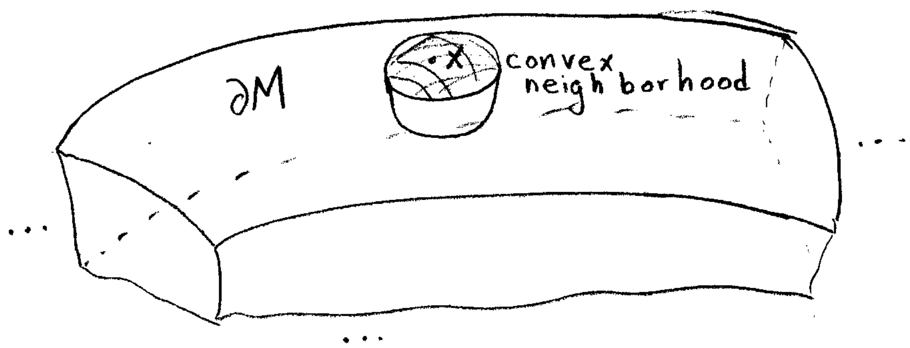

$x \in M$ has a neighborhood isometric to a convex subset of H^n. If $x \in \partial M$, then x will be on the boundary of this set. It is easy to convince oneself that local convexity implies convexity: picture a bath and imagine straightening it out. Because of local 8.11
convexity, one never needs to push it out of ∂M. To make this a rigorous argument, given a path p of length l there is an ϵ such that any path of length $\leq \epsilon$ intersecting $\mathcal{N}_l(p_0)$ is homotopic to a geodesic arc. Subdivide p into subintervals of length between $\epsilon/4$ and $\epsilon/2$. Straighten out adjacent pairs of intervals in turn, putting a new division point in the middle of the resulting arc unless it has length $\leq \epsilon/2$. Any time an interval becomes too small, change the subdivision. This process converges, giving a homotopy of p to a geodesic arc, since any time there are angles not close to π, the homotopy significantly shortens the path.

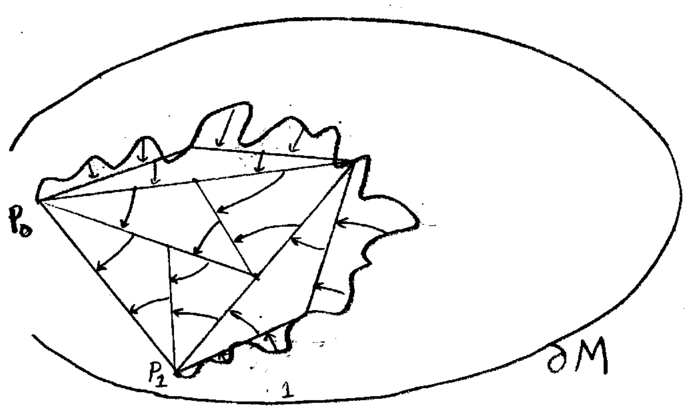

This give us a very concrete object corresponding to a Kleinian group: a complete convex hyperbolic three-manifold M with non-empty boundary. Given a convex manifold M, we can define $H(M)$ to be the intersection of all convex submanifolds M' of M such that $\pi_1 M' \to \pi_1 M$ is an isomorphism. $H(M)$ is clearly the same as $HL_{\pi_1}(M)/\pi_1(M)$. $H(M)$ is a convex manifold, with the same dimension as M except in degenerate cases. 8.12

Proposition 8.3.3. *If M is a compact convex hyperbolic manifold, then any small deformation of the hyperbolic structure on M can be enlarged slightly to give a new convex hyperbolic manifold homeomorphic to M.*

Proof. A convex manifold is *strictly convex* if every geodesic arc in M has interior in the interior of M. If M is not already strictly convex, it can be enlarged slightly to make it strictly convex. (This follows from the fact that a neighborhood of radius ϵ about a hyperplane is strictly convex.)

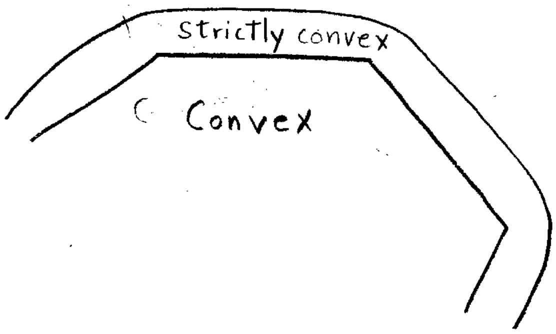

Thus we may assume that M' is a hyperbolic structure that is a slight deformation of a strictly convex manifold M. We may assume that our deformation M' is small

enough that it can be enlarged to a hyperbolic manifold M'' which contains a 2ϵ-neighborhood of M'. Every arc of length l greater than ϵ in M has the middle $(l-\epsilon)$ some uniform distance δ from ∂M; we may take our deformation M' of M small 8.13
enough that such intervals in M' have the middle $l-\epsilon$ still in the interior of M'. This implies that the union of the convex hulls of intersections of balls of radius 3ϵ with M' is locally convex, hence convex. $\square$

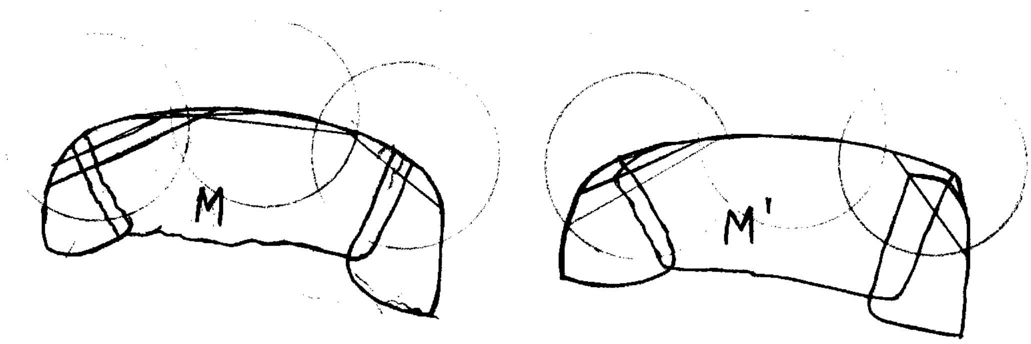

The convex hull of a uniformly small deformation of a uniformly convex manifold is locally determined.

Remark. When M is non-compact, the proof of 8.3.3 applies provided that M has a uniformly convex neighborhood and we consider only uniformly small deformations. We will study deformations in more generality in § .

Proposition 8.3.4. *Suppose M_1^n and M_2^n are strictly convex, compact hyperbolic manifolds and suppose $\phi : M_1^n \to M_2^n$ is a homotopy equivalence which is a diffeomorphism on ∂M_1. Then there is a quasi-conformal homeomorphism $f : B^n \to B^n$ of the Poincaré disk to itself conjugating $\pi_1 M_1$ to $\pi_1 M_2$. f is a pseudo-isometry on H^n.*

8.14

Proof. Let $\tilde{\phi}$ be a lift of ϕ to a map from $\tilde{M}_1$ to $\tilde{M}_2$. We may assume that $\tilde{\phi}$ is already a pseudo-isometry between the developing images of M_1 and M_2. Each point p on $\partial\tilde{M}_1$ and $\partial\tilde{M}_2$ has a unique normal ray γ_p; if $x \in \gamma_p$ has distance t from $\partial\tilde{M}_1$ let $f(x)$ be the point on $\gamma_{\tilde{\phi}(p)}$ a distance t from $\partial\tilde{M}_2$. The distance between points at a distance of t along two normal rays γ_{p_1} and γ_{p_2} at nearby points is approximately $\cosh t + \alpha \sinh t$, where d is the distance and θ is the angle between the normals of p_1 and p_2. From this it is evident that f is a pseudo-isometry extending to $\bar{\phi}$. $\square$

Associated with a discrete group Γ of isometries of H^n, there are at least four distinct and interesting quotient spaces (which are manifolds when Γ acts freely). Let us name them:

Definition 8.3.5.

$M_\Gamma = H(L_\Gamma)/\Gamma$, the convex hull quotient.
$N_\Gamma = H^n/\Gamma$, the complete hyperbolic manifold without boundary.
$O_\Gamma = (H^n \cup D_\Gamma)/\Gamma$, the Kleinian manifold.
$P_\Gamma = (H^n \cup D_\Gamma \cup W_\Gamma)/\Gamma$. Here $W_\Gamma \subset \mathbb{P}^n$ is the set of points in the projective model dual to planes in H^n whose intersection with S_∞ is contained in D_Γ.

We have inclusions $H(N_\Gamma) = M_\Gamma \subset N_\Gamma \subset O_\Gamma \subset P_\Gamma$. It is easy to derive the fact that Γ acts properly discontinuously on $H^n \cup D_\Gamma \cup W_\Gamma$ from the proper discontinuity on $H^n \cup D_\Gamma$. M_Γ, N_Γ and O_Γ have the same homotopy type. M_Γ and O_Γ are homeomorphic except in degenerate cases, and $N_\Gamma = \operatorname{int}(O_\Gamma)$ P_Γ is not always connected 8.15
when L_Γ is not connected.

8.4. Geometrically finite groups

Definition 8.4.1. Γ is geometrically finite if $\mathcal{N}_\epsilon(M_\Gamma)$ has finite volume.

The reason that $\mathcal{N}_\epsilon(M_\Gamma)$ is required to have finite volume, and not just M_Γ, is to rule out the case that Γ is an arbitary discrete group of isometries of $H^{n-1} \subset H^n$. We shall soon prove that geometrically finite *means* geometrically finite (8.4.3).

Theorem 8.4.2 (Ahlfors' Theorem). *If Γ is geometrically finite, then $L_\Gamma \subset S_\infty$ has full measure or 0 measure. If L_Γ has full measure, the action of Γ on S_∞ is ergodic.*

Proof. This statement is equivalent to the assertion that every bounded measurable function f supported on L_Γ and invariant by Γ is constant a.e. (with respect to Lebesque measure on S_∞). Following Ahlfors, we consider the function h_f on H^n determined by f as follows. If $x \in H^n$, the points on S_∞ correspond to rays through x; these rays have a natural "visual" measure V_x. Define $h_f(x)$ to be the average of f with respect to the visual measure V_x. This function h_f is harmonic, i.e., the gradient flow of h_f preserves volume,

$$\operatorname{div}\operatorname{grad} h_f = 0.$$

For this reason, the measure $\frac{1}{V_x(S_\infty)}V_x$ is called *harmonic measure*. To prove this, 8.16
consider the contribution to h_f coming from an infinitesimal area A centered at $p \in S^{n-1}$ (i.e., a Green's function). As x moves a distance d in the direction of p, the visual measure of A goes up exponentially, in proportion to $e^{(n-1)d}$. The gradient of any multiple of the characteristic function of A is in the direction of p, and also proportional in size to $e^{(n-1)d}$. The flow lines of the gradient are orthogonal trajectories to horospheres; this flow contracts linear dimensions along the horosphere in proportion to e^{-d}, so it preserves volume.

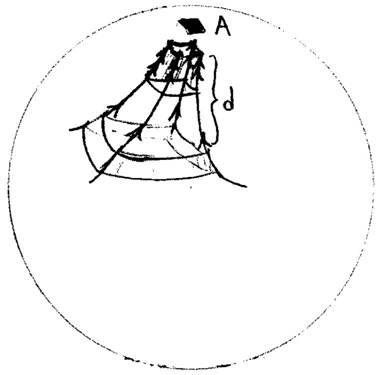

The average h_f of contributions from all the infinitesimal areas is therefore harmonic. We may suppose that f takes only the values of 0 and 1. Since f is invariant by Γ, so is h_f, and h_f goes over to a harmonic function, also h_f, on N_Γ. To complete the proof, observe that $h_f < \frac{1}{2}$ in $N_\Gamma - M_\Gamma$, since each point x in $H^n - H(L_\Gamma)$ lies in a half-space whose intersection with infinity does not meet L_Γ, which means that f is 0 on more than half the sphere, with respect to V_x. The set $\{x \in N_\Gamma | h_f(x) = \frac{1}{2}\}$ 8.17
must be empty, since it bounds the set $\{x \in N_\Gamma | h_f(x) \geq \frac{1}{2}\}$ of finite volume which flows into itself by the volume preserving flow generated by $\operatorname{grad} h_f$. (Observe that $\operatorname{grad} h_f$ has bounded length, so it generates a flow defined everywhere for all time.) But if $\{p | f(p) = 1\}$ has any points of density, then there are $x \in H^{n-1}$ near p with $h_f(x)$ near 1. It follows that f is a.e. 0 or a.e. 1. □

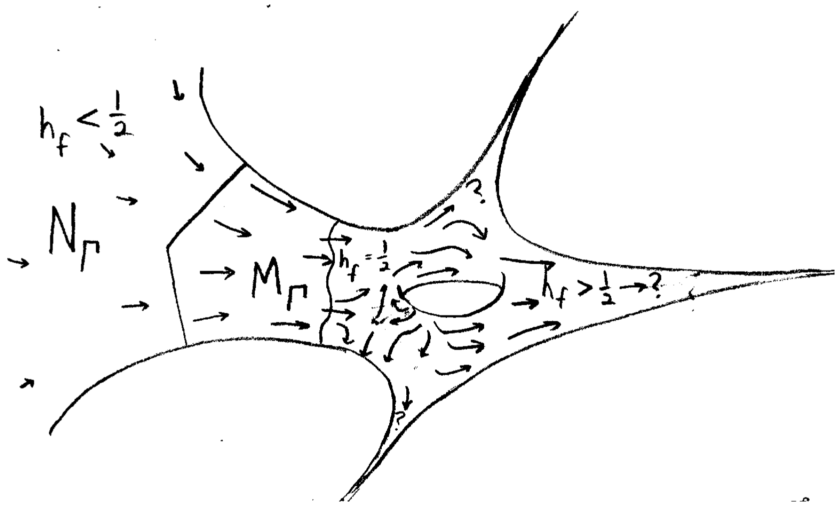

Let us now relate definition 8.4.1 to other possible notions of geometric finiteness. The usual definition is in terms of a fundamental polyhedron for the action of Γ. For concreteness, let us consider only the case $n = 3$. For the present discussion, a finite-sided polyhedron means a region P in H^3 bounded by finitely many planes. P

is a fundamental polyhedron for Γ if its translates by Γ cover H^3, and the translates of its interior are pairwise disjoint. P intersects S_∞ in a polygon which unfortunately 8.18
may be somewhat bizarre, since tangencies between sides of $P \cap S_\infty$ may occur.

Sometimes these tangencies are forced by the existence of parabolic fixed points for Γ. Suppose that $p \in S_\infty$ is a parabolic fixed point for some element of Γ, and let π be the subgroup of Γ fixing p. Let B be a horoball centered at p and sufficiently small that the projection of B/P to N_Γ is an embedding. (Compare §5.10.) If $\pi \supset \mathbb{Z} \oplus \mathbb{Z}$, for any point $x \in B \cap H(L_\Gamma)$, the convex hull of πx contains a horoball B', so in particular there is a horoball $B' \subset H(L_\Gamma) \cap B$. Otherwise, $\mathbb{Z}$ is a maximal torsion-free subgroup of π. Coordinates can be chosen so that p is the point at ∞ in the upper half-space model, and $\mathbb{Z}$ acts as translations by real integers. There is some minimal strip $S \subseteq \mathbb{C}$ containing $L_\Gamma \cap \mathbb{C}$; S may interesect the imaginary axis in a finite, half-infinite, or doubly infinite interval. In any case, $H(L_\Gamma)$ is contained in the region R of upper half-space above S, and the part of ∂R of height ≥ 1 lies on ∂H_Γ. 8.19

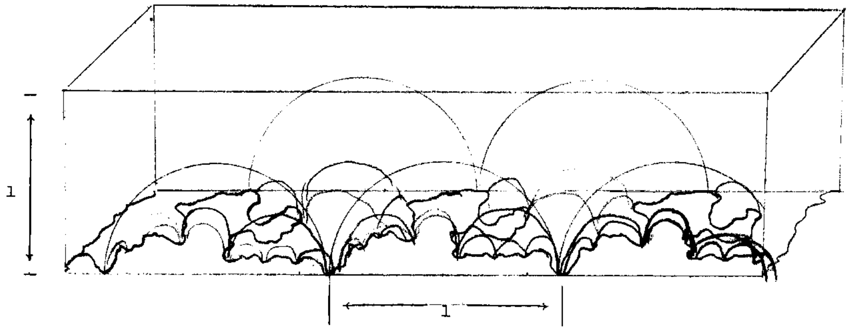

It may happen that there are wide substrips $S' \subset S$ in the complement of L_Γ. If S' is sufficiently wide, then the plane above its center line intersects $H(L_\Gamma)$ in B, so it gives a half-open annulus in $B/\mathbb{Z}$. If Γ is torsion-free, then maximal, sufficiently wide strips in $S - L_\Gamma$ give disjoint non-parallel half-open annuli in M_Γ; an easy argument shows they must be finite in number if Γ is finitely generated. (This also follows from Ahlfors's finiteness theorem.) Therefore, there is some horoball B' centered at p so that $H(L_\Gamma) \cap B' = R \cap B'$. This holds even if Γ has torsion.

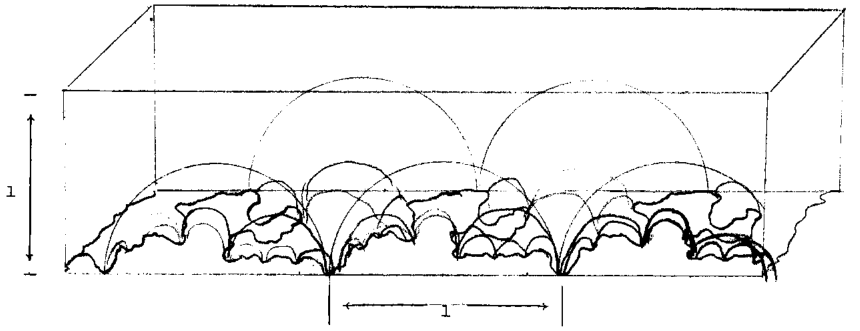

8.20

With an understanding of this picture of the behaviour of M_Γ near a cusp, it is not hard to relate various notions of geometric finiteness. For convenience suppose Γ is torsion-free. (This is not an essential restriction in view of Selberg's theorem—see § .) When the context is clear, we abbreviate $M_\Gamma = M$, $N_\Gamma = N$, etc.

PROPOSITION 8.4.3. *Let $\Gamma \subset \mathrm{PSL}(2, \mathbb{C})$ be a discrete, torsion-free group. The following conditions are equivalent:*

(a) Γ *is geometrically finite* (see dfn. 8.4.1).
(b) $M_{[\epsilon,\infty)}$ *is compact.*
(c) Γ *admits a finite-sided fundamental polyhedron.*

PROOF. (a) $\Rightarrow$ (b).

Each point in $M_{[\epsilon,\infty)}$ has an embedded $\epsilon/2$ ball in $\mathcal{N}_{\epsilon/2}(M_\Gamma)$, by definition. If (a) holds, $\mathcal{N}_{\epsilon/n}(M_\Gamma)$ has finite volume, so only finitely many of these balls can be disjoint and $M_{\Gamma[\epsilon,\infty)}$ is compact.

(b) $\Rightarrow$ (c). First, find fundamental polyhedra near the non-equivalent parabolic fixed points. To do this, observe that if p is a $\mathbb{Z}$-cusp, then in the upper half-space model, when $p = \infty$, $L_\Gamma \cap \mathbb{C}$ lies in a strip S of finite width. Let R denote the region above S. Let B' be a horoball centered at ∞ such that $R \cap B' = H(L_\Gamma) \cap B'$. Let $r : H^3 \cup D_\Gamma \to H(L_\Gamma)$ be the canonical retraction. If Q is any fundamental polyhedra for the action of $\mathbb{Z}$ in some neighborhood of p in $H(L_\Gamma)$ then $r^{-1}(Q)$ is a fundamental polyhedron in some neighborhood of p in $H^3 \cup D_\Gamma$. 8.21

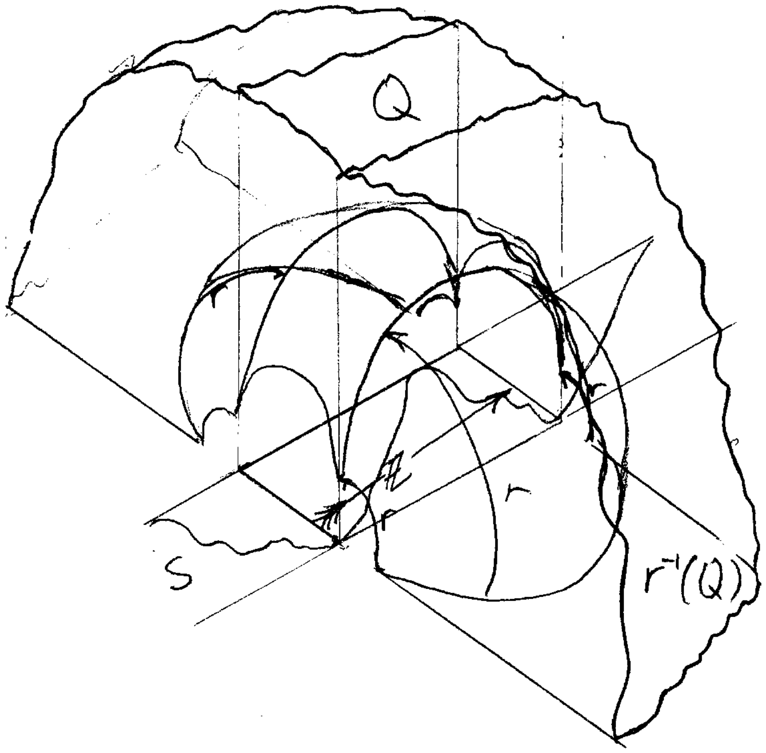

A fundamental polyhedron near the cusps is easily extended to a global fundamental polyhedron, since O_Γ-(neighborhoods of the cusps) is compact.

(c) $\Rightarrow$ (a). Suppose that Γ has a finite-sided fundamental polyhedron P.

A point $x \in P \cap S_\infty$ is a regular point ($\in D_\Gamma$) if it is in the interior of $P \cap S_\infty$ or of some finite union of translates of P. Thus, the only way x can be a limit point is for x to be a point of tangency of sides of infinitely many translates of P. Since P can have only finitely many points of tangency of sides, infinitely many $\gamma\Gamma$ must identify one of these points to x, so x is a fixed point for some element $\gamma\Gamma$. γ must be parabolic, otherwise the translates of P by powers of γ would limit on the axis of γ. If x is arranged to be ∞ in upper half-space, it is easy to see that $L_\Gamma \mathbb{C}$ must be contained in a strip of finite width. (Finitely many translates of P must form a fundamental domain for $\{\gamma^n\}$, acting on some horoball centered at ∞, since $\{\gamma^n\}$ has finite index in the group fixing ∞. Th faces of these translates of P which do not pass through ∞ lie on hemispheres. Every point in $\mathbb{C}$ outside this finite collection of hemispheres and their translates by $\{\gamma^n\}$ lies in D_Γ.) 8.22

It follows that $v(\mathcal{N}_\epsilon(M)) = v(\mathcal{N}_\epsilon(H(L_\Gamma)) \cap P)$ if finite, since the contribution near any point of $L_\Gamma \cap P$ is finite and the rest of $\mathcal{N}_\epsilon(H(L_\Gamma)) \cap P$ is compact. $\square$

8.5. The geometry of the boundary of the convex hull

Consider a closed curve σ in Euclidean space, and its convex hull $H(\sigma)$. The boundary of a convex body always has non-negative Gaussian curvature. On the other hand, each point p in $\partial H(\sigma) - \sigma$ lies in the interior of some line segment or triangle with vertices on σ. Thus, there is some line segment on $\partial H(\sigma)$ through p, so that $\partial H(\sigma)$ has non-positive curvature at p. It follows that $\partial H(\sigma) - \sigma$ has zero curvature, i.e., it is "developable". If you are not familiar with this idea, you can see it by bending a curve out of a piece of stiff wire (like a coathanger). Now roll the wire around on a big piece of paper, tracing out a curve where the wire touches. Sometimes, the wire may touch at three or more points; this gives alternate ways to roll, and you should carefully follow all of them. Cut out the region in the plane bounded by this curve (piecing if necessary). By taping the paper together, you can envelope the wire in a nice paper model of its convex hull. The physical process of unrolling a developable surface onto the plane is the origin of the notion of the developing map. 8.23

The same physical notion applies in hyperbolic three-space. If K is any closed set on S_∞, then $H(K)$ is convex, yet each point on $\partial H(K)$ lies on a line segment in $\partial H(K)$. Thus, $\partial H(K)$ can be developed to a hyperbolic plane. (In terms of Riemannian geometry, $\partial H(K)$ has extrinsic curvature 0, so its intrinsic curvature is the ambient sectional curvature, -1. Note however that $\partial H(K)$ is not usually differentiable). Thus $\partial H(K)$ has the natural structure of a complete hyperbolic surface.

Proposition 8.5.1. *If Γ is a torsion-free Kleinian group, the ∂M_Γ is a hyperbolic surface.* □

The boundary of M_Γ is of course not generally flat—it is bent in some pattern. Let $\gamma \subset \partial M_\Gamma$ consist of those points which are not in the interior of a flat region of 8.24
∂M_Γ. Through each point x in γ, there is a unique geodesic g_x on ∂M_Γ. g_x is also a geodesic in the hyperbolic structure of ∂M_Γ. γ is a closed set. If ∂M_Γ has finite area, then γ is compact, since a neighborhood of each cusp of ∂M_Γ is flat. (See §8.4.)

Definition 8.5.2. A *lamination* L on a manifold M^n is a closed subset $A \subset M$ (the *support* of L) with a local product structure for A. More precisely, there is a covering of a neighborhood of A in M with coordinate neighborhoods $U_i \xrightarrow{\phi_i} \mathbb{R}^{n-k} \times \mathbb{R}^k$ so that $\phi_i(A \cap U_i)$ is of the form $\mathbb{R}^{n-k} \times B$, $B \subset \mathbb{R}^k$. The coordinate changes ϕ_{ij} must be of the form $\phi_{ij}(x, y) = (f_{ij}(x, y), g_{ij}(y))$ when $y \in B$. A lamination is like a foliation of a closed subset of M. Leaves of the lamination are defined just as for a foliation.

EXAMPLES. If $\mathcal{F}$ is a foliation of M and $S \subset M$ is any set, the closure of the union of leaves which meet S is a lamination.

Any submanifold of a manifold M is a lamination, with a single leaf. Clearly, the bending locus γ for ∂M_Γ has the structure of a lamination: whenever two points of γ are nearby, the directions of bending must be nearly parallel in order that the lines of bending do not intersect. A lamination whose leaves are geodesics we will call a *geodesic lamination.* 8.25

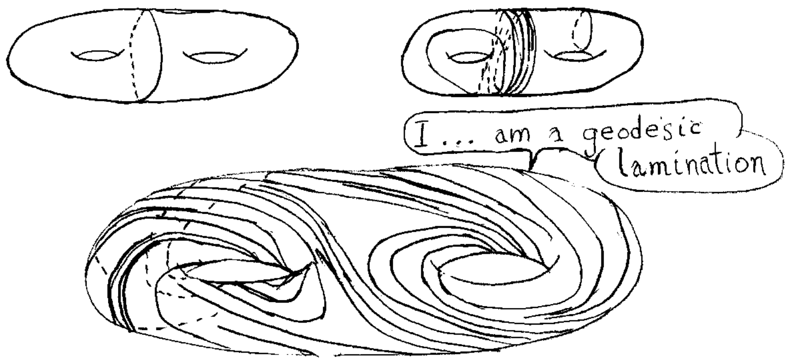

By consideration of Euler characteristic, the lamination γ cannot have all of ∂M as its support, or in other words it cannot be a foliation. The complement $\partial M - \gamma$ consists of regions bounded by closed geodesics and infinite geodesics. Each of these regions can be doubled along its boundary to give a complete hyperbolic surface, which of course has finite area. There

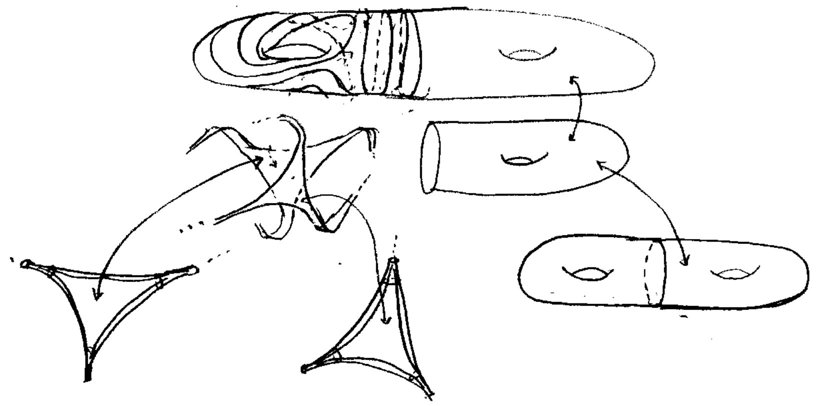

8.26

is a lower bound for π for the area of such a region, hence an upper bound of $2|\chi(\partial M)|$ for the number of components of $\partial M - \gamma$. Every geodesic lamination γ on a hyperbolic surface S can be extended to a foliation with isolated singularities on the complement. There

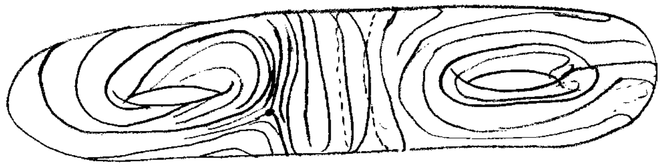

is an index formula for the Euler characteristic of S in terms of these singularities. Here are some values for the index.

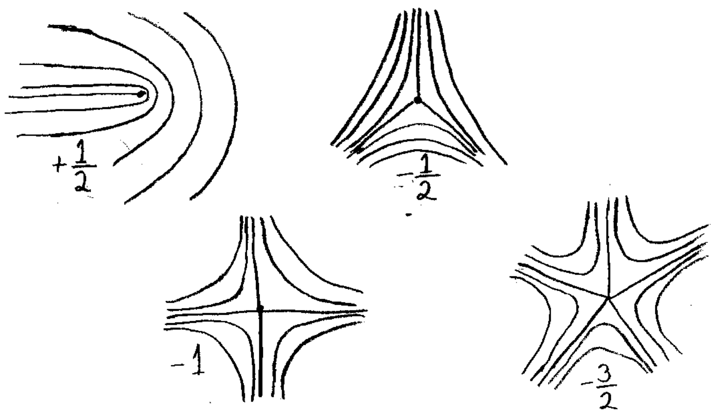

8.27

From the *existence* of an index formula, one concludes that the Euler characteristic of S is half the Euler characteristic of the double of $S - \gamma$. By the Gauss-Bonnet theorem,

$$\text{Area}(S - \gamma) = \text{Area}(S)$$

or in other words, γ has measure 0. To give an idea of the range of possibilities for geodesic laminations, one can consider an arbitrary sequence $\{\gamma_i\}$ of geodesic laminations: simple closed curves, for instance. Let us say that $\{\gamma_i\}$ converges geometrically to γ if for each $x \in \text{support}\,\gamma$, and for each ϵ, for all great enough i the support of γ_i intersects $\mathcal{N}_\epsilon(x)$ and the leaves of $\gamma_i \cap \mathcal{N}_\epsilon(x)$ are within ϵ of the direction of the leaf of γ through x. Note that the support of γ may be smaller than the limiting support of γ_i, so the limit of a sequence may not be unique. See §8.10. An easy diagonal argument shows that every sequence $\{\gamma_i\}$ has a subsequence which converges geometrically. From limits of sequences of simple closed geodesics, uncountably many geodesic laminations are obtained.

Geodesic laminations on two homeomorphic hyperbolic surfaces may be compared by passing to the circle at ∞. A directed geodesic is determined by a pair of points $(x_1, x_2) \in S^1_\infty \times S^1_\infty - \Delta$, where Δ is the diagonal $\{(x, x)\}$. A geodesic without direction is a point on $J = (S^1_\infty \times S^1_\infty - \Delta/\mathbb{Z}_2)$, where $\mathbb{Z}_2$ acts by interchanging coordinates. Topologically, J is an open Moebius band. It is geometrically realized in the Klein (projective) model for H^2 as the region outside H^2. A geodesic g projects 8.28

to a simple geodesic on the surface S if and only if the covering translates of its pairs of end points never strictly separate each other.

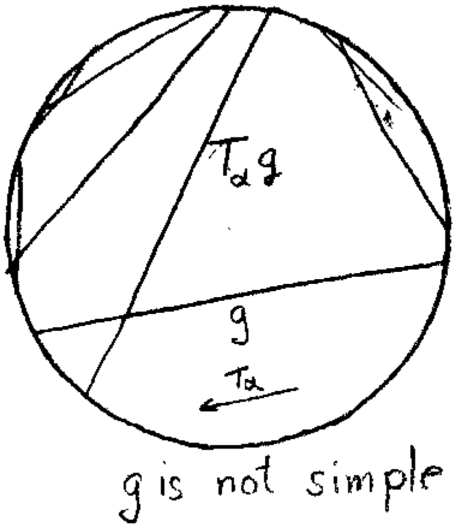

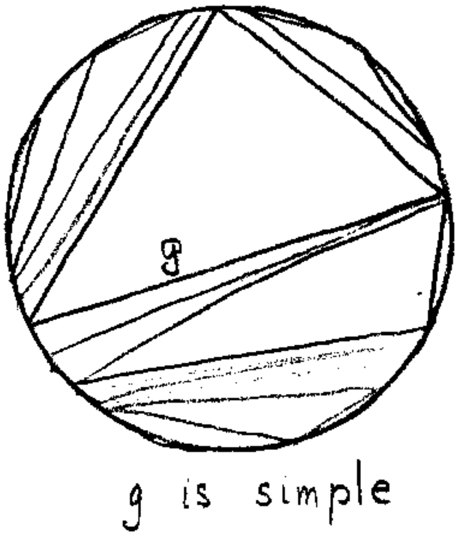

Geometrically, J has an indefinite metric of type $(1,1)$, invariant by covering translates. (See §2.6.) The light-like geodesics, of zero length, are lines tangent to S^1_∞; lines which meet H^2 when extended have imaginary arc length. A point $g \in J$ projects to a simple geodesic in S if and only if no covering translate $T_\alpha(g)$ has a positive real distance from g.

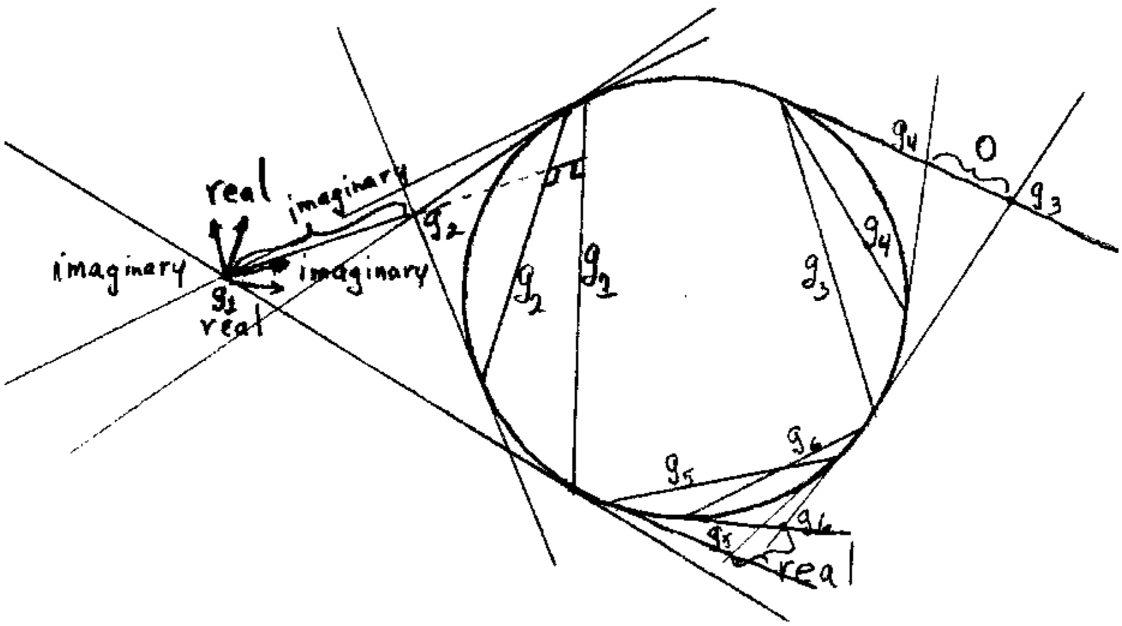

8.29

Let $\mathcal{S} \subset J$ consist of all elements g projecting to simple geodesics on S. Any geodesic $\subset H^2$ which has a translate intersecting itself has a neighborhood with the same property, hence $\mathcal{S}$ is closed.

If γ is any geodesic lamination on S, Let $\mathcal{S}_\gamma \subset J$ be the set of lifts of leaves of γ to H^2. $\mathcal{S}_\gamma$ is a closed invariant subset of $\mathcal{S}$. A closed invariant subset of $C \subset J$ gives rise to a geodesic lamination if and only if all pairs of points of C are separated by an imaginary (or 0) distance. If $g \in \mathcal{S}$, then the closure of its orbit, $\overline{\pi_1(S)g}$ is such a set, corresponding to the geodesic lamination $\bar{g}$ of S. Every homeomorphism between surfaces when lifted to H^2 extends to S^1_∞ (by 5.9.5). This determines an extension to J. Geodesic laminations are transferred from one surface to another via this correspondence.

8.6. Measuring laminations

Let L be a lamination, so that it has local homeomorphisms $\phi_i : L \cap U_i \approx \mathbb{R}^{n-k} \times B_i$. A *transverse measure* μ for L means a measure μ_i defined on each local leaf space B_i, in such a way that the coordinate changes are measure preserving. Alternatively one may think of μ as a measure defined on every k-dimensional submanifold transverse to L, supported on $T^k \cap L$ and invariant under local projections along leaves of L. We will always suppose that μ is finite on compact transversals. The simplest example of a transverse measure arises when L is a closed submanifold; in this case, one can take μ to count the number of intersections of a transversal with L. 8.30

We know that for a torsion-free Kleinian group Γ, ∂M_Γ is a hyperbolic surface bent along some geodesic lamination γ. In order to complete the picture of ∂M_Γ, we need a quantitative description of the bending. When two planes in H^3 meet along a line, the angle they form is constant along that line. The flat pieces of ∂M_Γ meet each other along the geodesic lamination γ; the angle of meeting of two planes generalizes to a transverse "bending" measure, β, for γ. The measure β applied to an arc α on ∂M_Γ transverse to γ is the total angle of turning of the normal to ∂M_Γ along α (appropriately interpreted when γ has isolated geodesics with sharp bending). In order to prove that β is well-defined, and that it determines the local isometric embedding in H^3, one can use local polyhedral approximations to ∂M_Γ. Local outer approximations to ∂M_Γ can be obtained by extending the planes of local flat regions. Observe that when three planes have pairwise intersections in H^3 but no triple intersection, the dihedral angles satisfy the inequality

$$\alpha + \beta \leq \gamma.$$

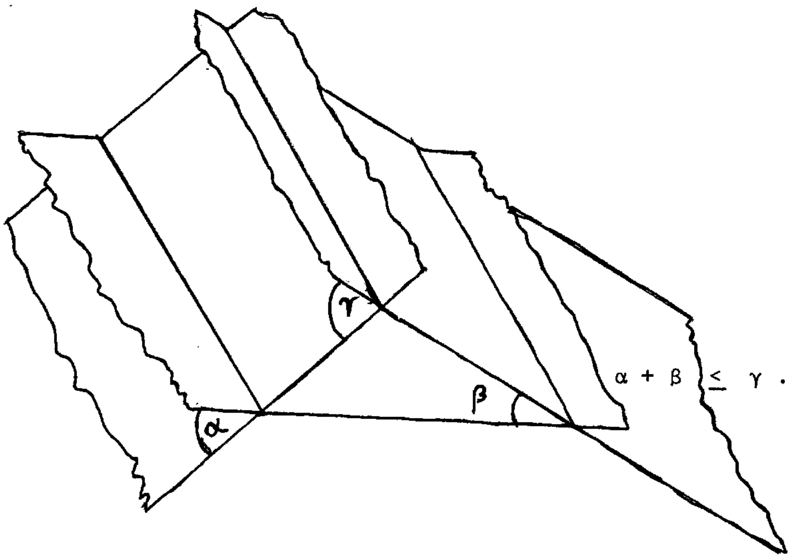

8.31

(The difference $\gamma - (\alpha + \beta)$ is the area of a triangle on the common perpendicular plane.) From this it follows that as outer polyhedral approximations shrink toward M_Γ, the angle sum corresponding to some path α on ∂M_Γ is a monotone sequence, converging to a value $\beta(\alpha)$. Also from the monotonicity, it is easy to see that for short paths α_t, $[0 \le t \le 1]$, $\beta(\alpha)$ is a close approximation to the angle between the tangent planes at α_0 and α_1. This implies that the hyperbolic structure on ∂M_Γ, together with the geodesic lamination γ and the transverse measure β, completely determines the hyperbolic structure of N_Γ in a neighborhood of ∂M_Γ.

The bending measure β has for its support all of γ. This puts a restriction on the structure of γ: every isolated leaf L of γ must be a closed geodesic on ∂M_Γ. (Otherwise, a transverse arc through any limit point of L would have infinite measure.) This limits the possibilities for the intersection of a transverse arc with γ to a Cantor set and/or a finite set of points.

When γ contains more than one closed geodesic, there is obviously a whole family of possibilities for transverse measures. There are (probably atypical) examples of families of distinct transverse measures which are not multiples of each other even for certain geodesic laminations such that every leaf is dense. There are many other examples which possess unique transverse measures, up to constant multiples. Compare Katok. 8.32

Here is a geometric interpretation for the bending measure β in the Klein model. Let P_0 be the component of P_Γ containing N_Γ (recall definition 8.3.5). Each point in $\tilde{P}_0$ outside S_∞ is dual to a plane which bounds a half-space whose intersection with S_∞ is contained in D_Γ. $\partial \tilde{P}_0$ consists of points dual to planes which meet L_Γ in at least one point. In particular, each plane meeting $\tilde{M}_\Gamma$ in a line or flat of $\partial \tilde{M}_\Gamma$ is dual