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. 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 $\lim \gamma_i y = \lim \gamma_i x$.

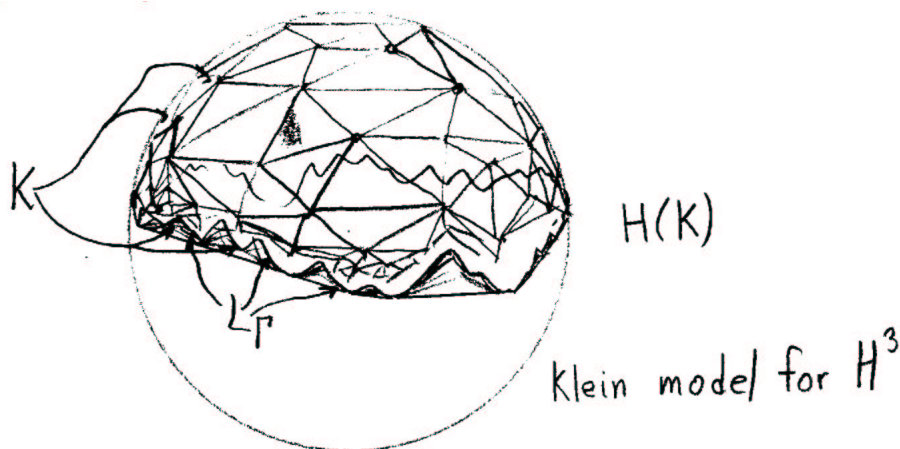The group Γ is called *elementary* if the limit set consists of $0, 1$ or 2 points.

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 $L_\Gamma \subset K$. $\qquad \square$

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 $1 \neq \Gamma' \lhd \Gamma$ 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. $\qquad \square$
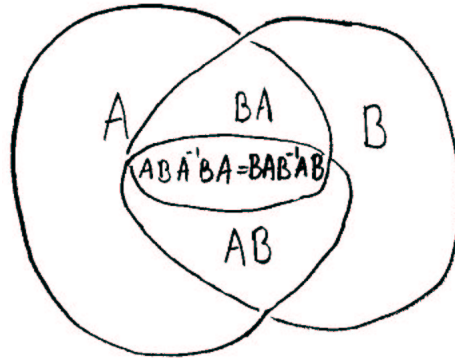
EXAMPLES. If M^2 is a hyperbolic surface, we may regular $\pi_1(M)$ as a group of isometries of a hyperbolic plane in H^3. The limit set is a circle. A group with limit set contained in a geometric circle is called a *Fuchsian group*.

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 $\langle A, B : ABA^{-1}BA = BAB^{-1}AB \rangle$.
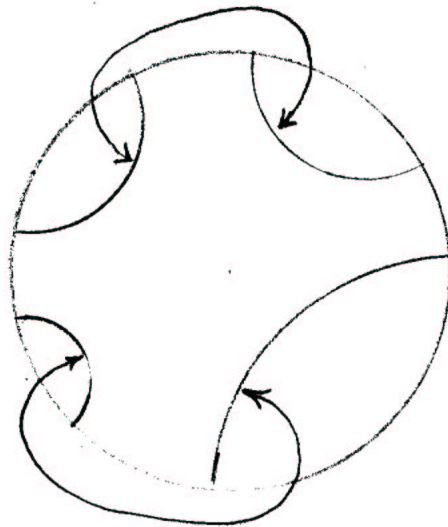
8.4

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.

Schottky group
n = 2
k = 2

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 $D_\Gamma = S_\infty^{n-1} - L_\Gamma$. 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 $L_\Gamma = S_\infty^2$, 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 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.*
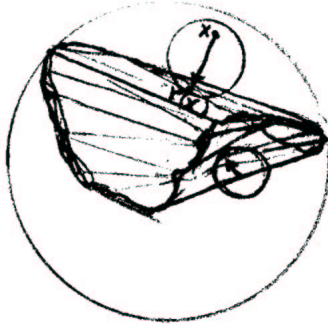
PROOF. Let $x_1, x_2 \in X$ be points on distinct orbits of Γ. Let N_1 be a compact neighborhood of x_1. Finitely many translates of x_2 intersect N_1, so we may assume N_1 is disjoint from the orbit of x_2. Then $\bigcup_{\gamma \in \Gamma} \gamma N_1$ gives an invariant neighborhood of x_1 disjoint 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.

<div style="text-align: right">8.7</div>

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. $\quad\square$

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.

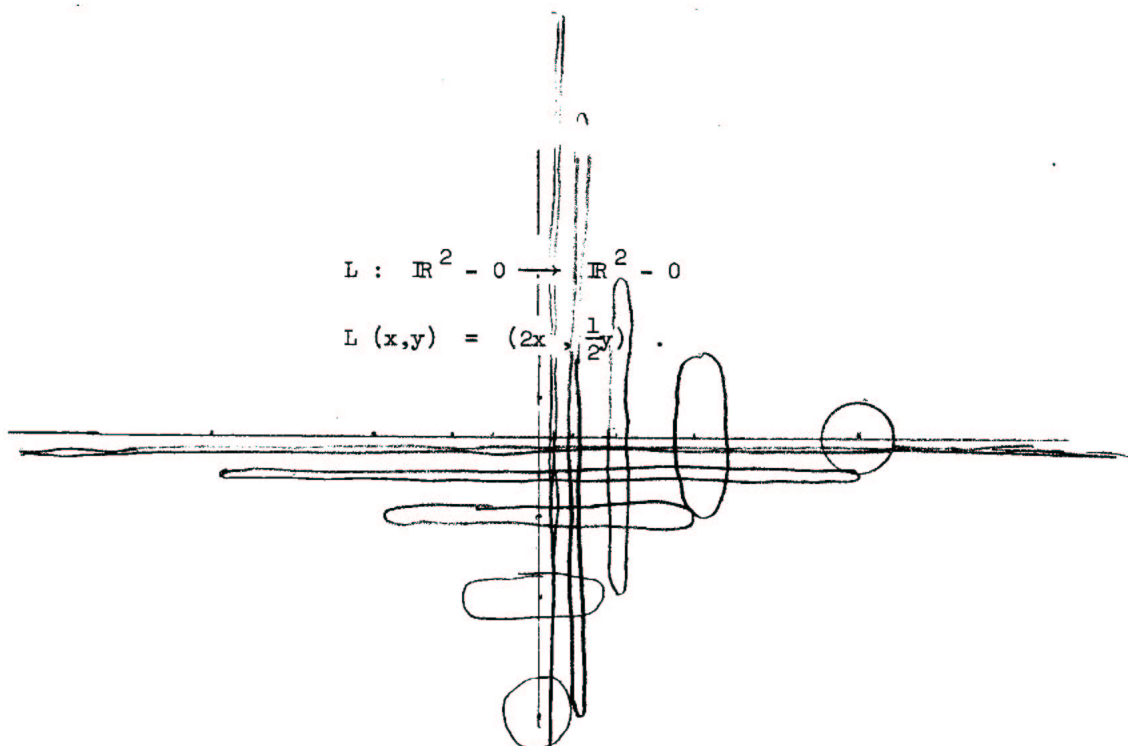DEFINITION 8.2.5. The action of Γ has *discrete orbits* if every orbit of Γ has an empty limit set.

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. $\quad\square$

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

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

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

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

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

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 as $(1,0)$ and $(0,1)$ do not have disjoint neighborhoods.

8.9

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.

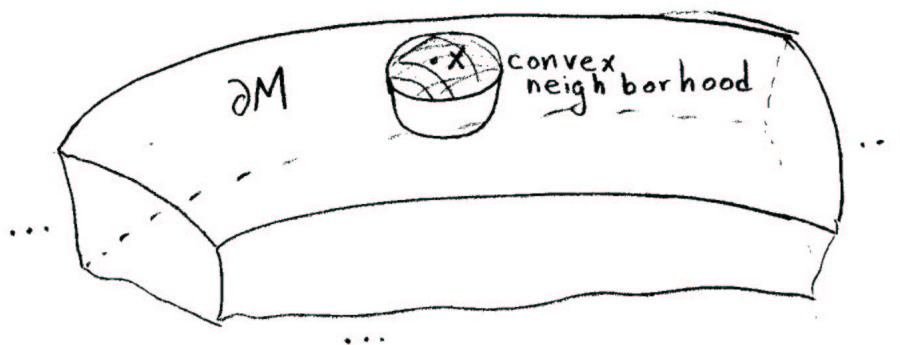We consider for the present only groups acting freely.

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 case of an arc which is a single point may occur.)

8.10

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.

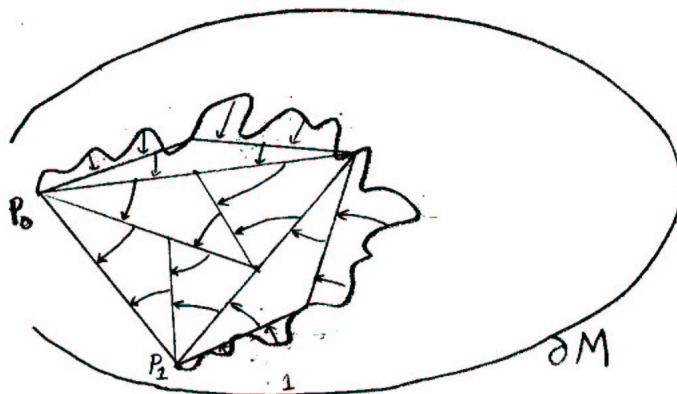If M is convex, then D is $1 - 1$, since any two points in $\tilde{M}$ may be joined by a path, which is homotopic in M and hence in $\tilde{M}$ to a geodesic arc. D must take the endpoints of a geodesic arc to distinct points. $D(\tilde{M})$ is clearly convex. □

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 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.
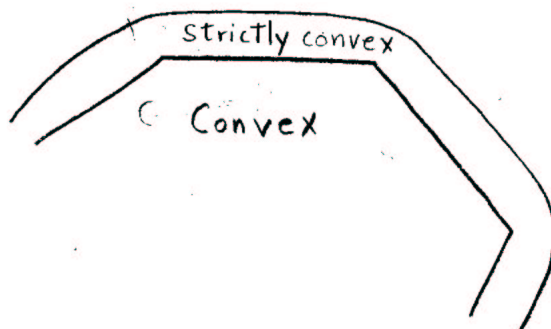
8.11

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.
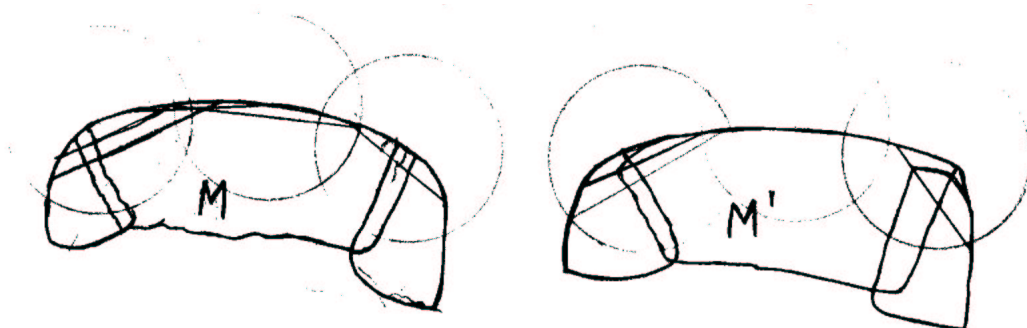
8.12

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.   □

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 $\bar{\phi}$.   □

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 $N_\Gamma = \mathrm{int}(O_\Gamma)$ P_Γ is not always connected when L_Γ is not connected.

8.15

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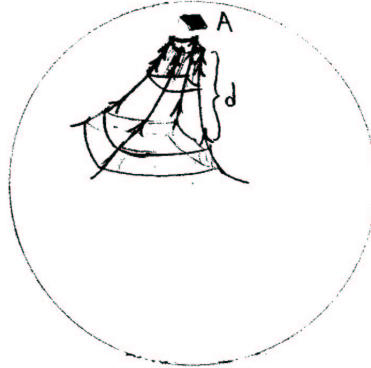
THEOREM 8.4.2 (Ahlfors' Theorem). *If Γ is geometrically finite, then $L_\Gamma \subset S_\infty$ has full measure or 0 measure. If L_Γ has full measure, the action of Γ on S_∞ is ergodic.*

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,

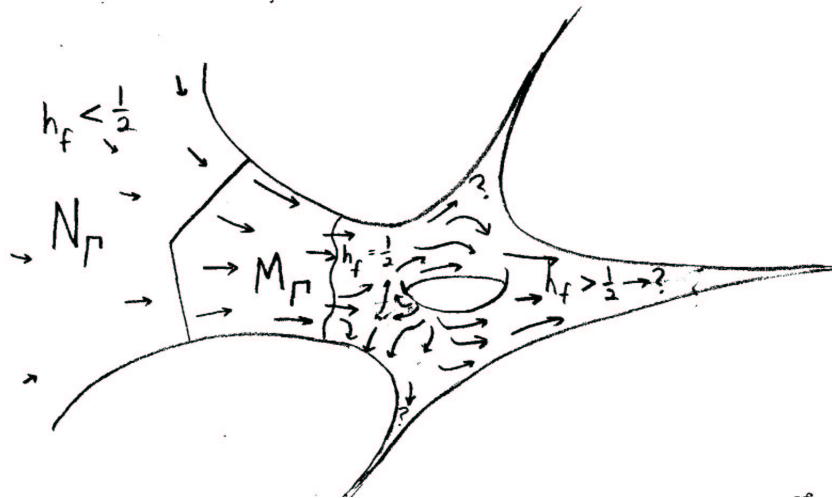$$\mathrm{div}\,\mathrm{grad}\,h_f = 0.$$

For this reason, the measure $\frac{1}{V_x(S_\infty)}V_x$ is called *harmonic measure*. To prove this, 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.

8.16

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. $\square$
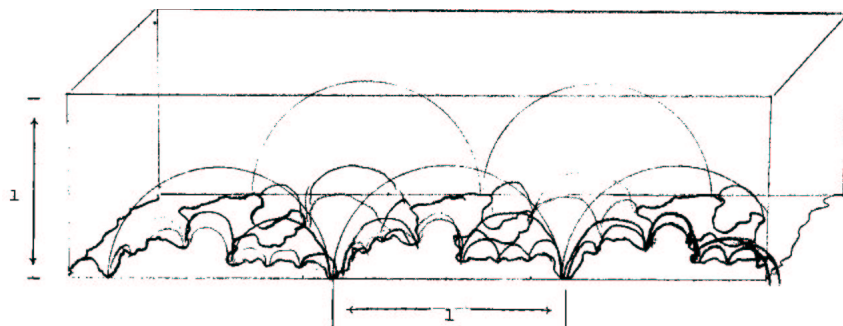

Let us now relate definition 8.4.1 to other possible notions of geometric finiteness. The usual definition is in terms of a fundamental polyhedron for the action of Γ. For concreteness, let us consider only the case $n = 3$. For the present discussion, a finite-sided polyhedron means a region P in H^3 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
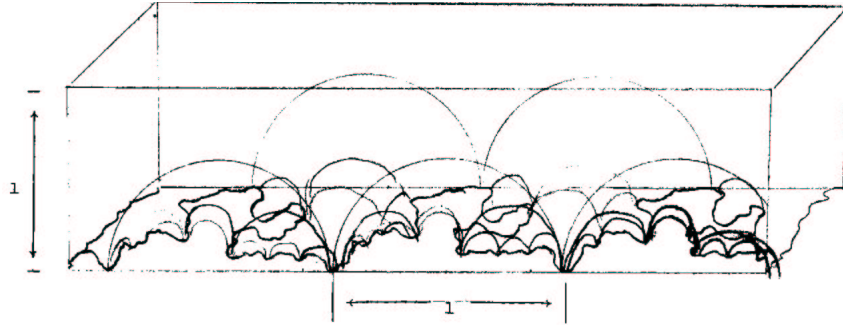may be somewhat bizarre, since tangencies between sides of $P \cap S_\infty$ may occur.

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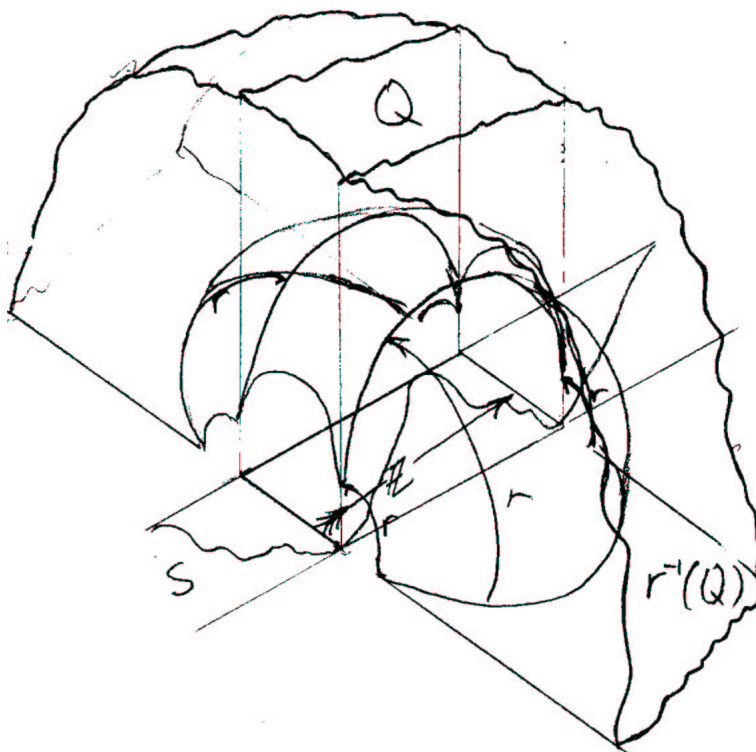
    (a) Γ *is geometrically finite* (see dfn. 8.4.1).
    (b) $M_{[\epsilon, \infty)}$ *is compact.*
    (c) Γ *admits a finite-sided fundamental polyhedron.*

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.)

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. $\square$

8.22

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.* $\square$

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 ∂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.)

8.24

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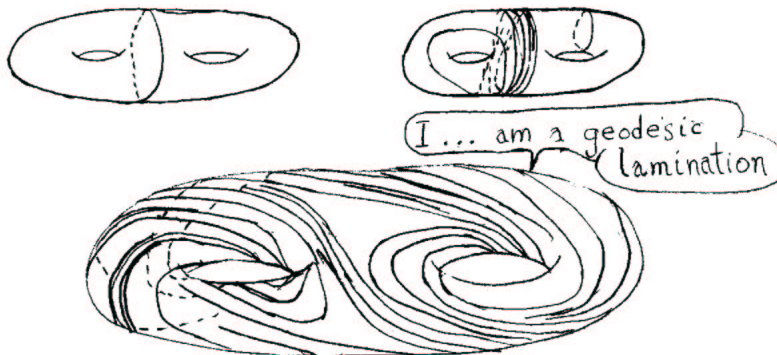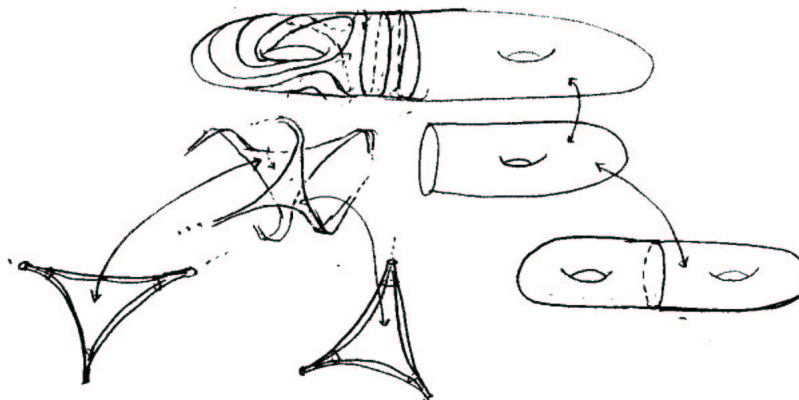
EXAMPLES. If $\mathcal{F}$ is a foliation of M and $S \subset M$ is any set, the closure of the union of leaves which meet S is a lamination.

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*.
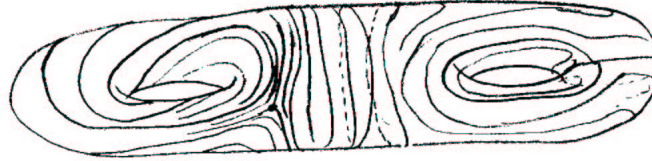
8.25

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
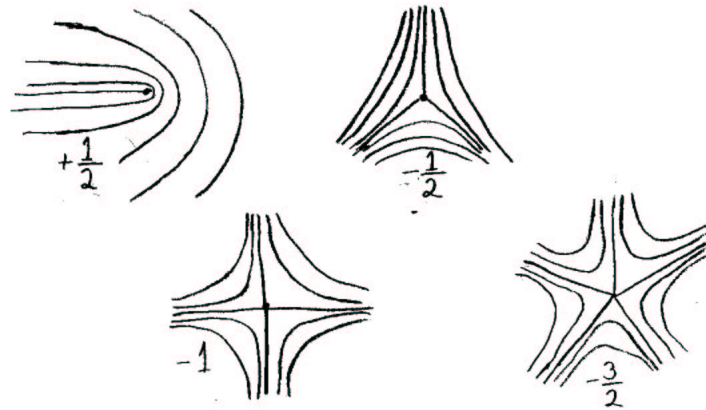

8.26

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,

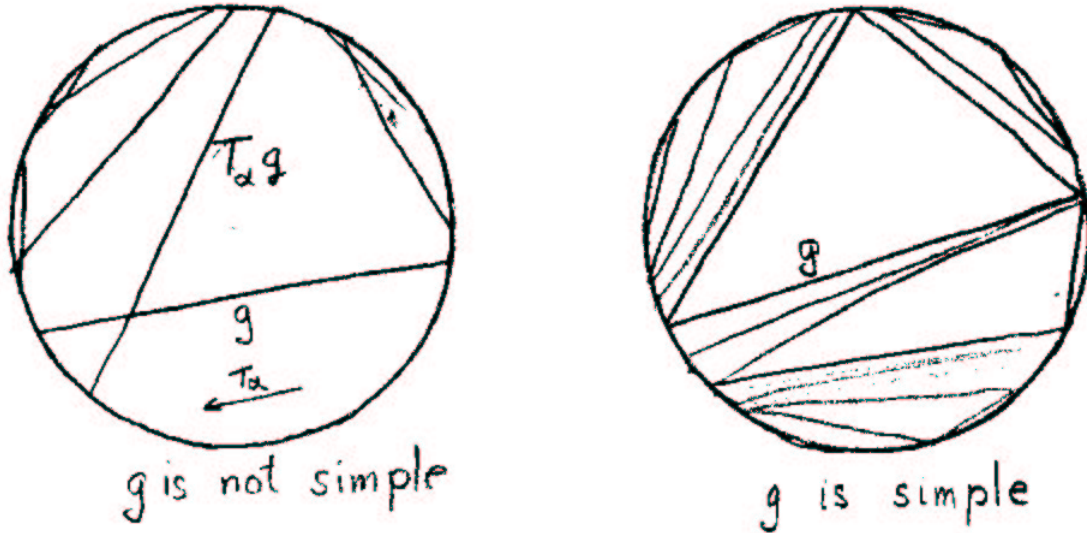$$\text{Area}(S - \gamma) = \text{Area}(S)$$

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
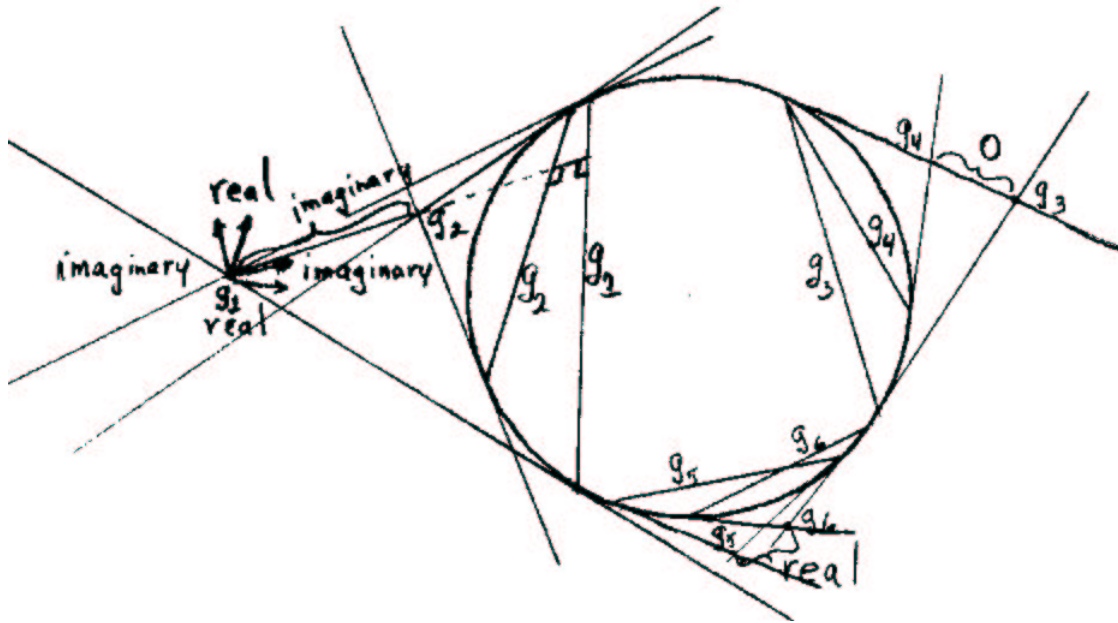
8.28

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.

g is not simple

g is simple

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.
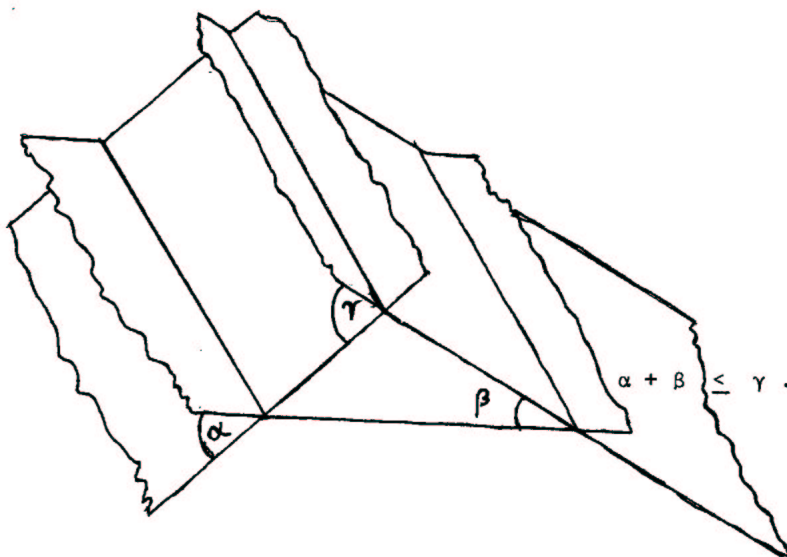
8.6. Measuring laminations

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

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

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

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, $[0 \leq t \leq 1]$, $\beta(\alpha)$ 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.
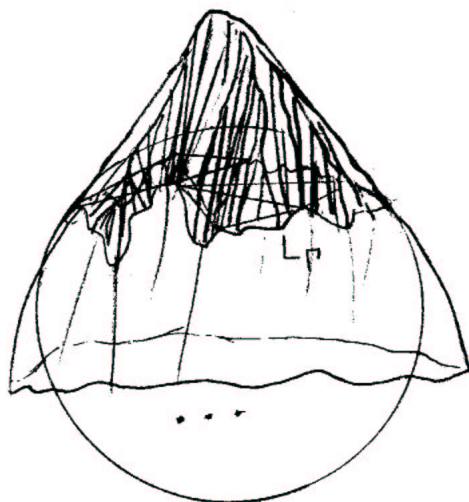
When γ contains more than one closed geodesic, there is obviously a whole family of possibilities for transverse measures. There are (probably atypical) examples of families of distinct transverse measures which are not multiples of each other even for certain geodesic laminations such that every leaf 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

to a point on $\partial\tilde{P}_0$. If $\bar{\pi} \in \partial\tilde{P}_0$ is dual to a plane π touching L_Γ at x, then one of the line segments $\overline{\overline{\pi}x}$ is also on $\partial\tilde{P}_0$. This line segments consists of points dual to planes touching L_Γ at x and contained in a half-space bounded by π. The reader may check that $\tilde{P}_0$ is convex. The natural metric of type $(2,1)$ in the exterior of S_∞ is degenerate on $\partial\tilde{P}_0$, since it vanishes on all line segments corresponding to a family of planes tangent at S_∞. Given a path α on $\partial\tilde{M}_\Gamma$, there is a dual path $\bar{\alpha}$ consisting of points dual to planes just skimming M_Γ along α. The length of $\bar{\alpha}$ is the same as $\beta(\alpha)$.

8.33

REMARK. The interested reader may verify that when N is a component of ∂M_Γ such that every leaf of $\gamma \cap N$ is dense in $\gamma \cap N$, then the action of $\pi_1 n$ on the appropriate component of $\partial\tilde{P}_0 - L_\Gamma$ is minimal (i.e., every orbit is dense). This action is approximated by actions of $\pi_1 N$ as covering transformations on surfaces just inside $\partial\tilde{P}_0$.
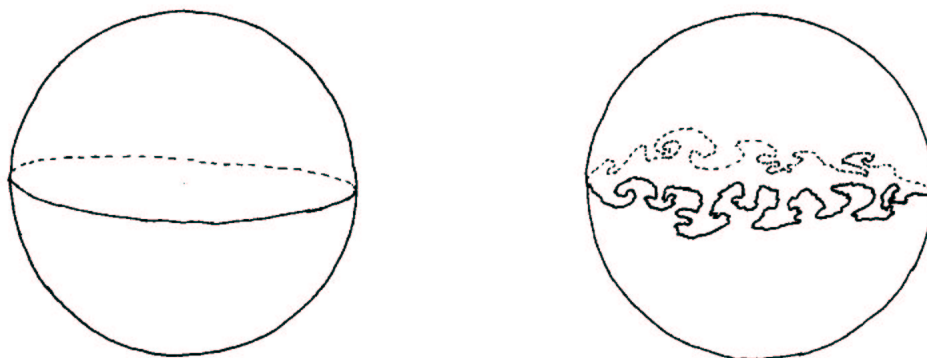
8.7. Quasi-Fuchsian groups

Recall that a Fuchsian group (of type I) is a Kleinian group Γ whose limit set L_Γ is a geometric circle. Examples are the fundamental groups of closed, hyperbolic surfaces. In fact, if the Fuchsian group Γ is torsion-free and has no parabolic elements, then Γ is the group of covering transformations of a hyperbolic surface. Furthermore, the Kleinian manifold $O_\Gamma = (H^3 \cup D_\Gamma)/\Gamma$ has a totally geodesic surface as a spine.

NOTE. The type of a Fuchsian group should not be confused with its type as a Kleinian group. To say that Γ is a Fuchsian group of type I means that $L_\Gamma = S^1$, but it is a Kleinian group of type II since $D_\Gamma \neq \emptyset$.

Suppose $M = N^2 \times I$ is a convex hyperbolic manifold, where N^2 is a closed surface. Let Γ' be the group of covering transformations of M, and let Γ be a Fuchsian group coming from a hyperbolic structure on N. Γ and Γ' are isomorphic as groups; we want to show that their actions on the closed ball B^3 are topologically conjugate.

Let M_Γ and $M_{\Gamma'}$ be the convex hull quotients ($M_\Gamma \approx N^2$ and $M_{\Gamma'} \approx N^2 \times I$). Thicken M_Γ and $M_{\Gamma'}$ to strictly convex manifolds. The thickened manifolds are diffeomorphic, so by Proposition 8.3.4 there is a quasi-conformal homeomorphism of B^3 conjugating Γ to Γ'. In particular, $L_{\Gamma'}$ is homeomorphic to a circle. Γ', which has convex hull manifold homeomorphic to $N^2 \times I$ and limit set $\approx S^1$, is an example of a quasi-Fuchsian group. 8.34

DEFINITION 8.7.1. The Kleinian group Γ is called a *quasi-Fuchsian group* if L_Γ is topologically S^1.

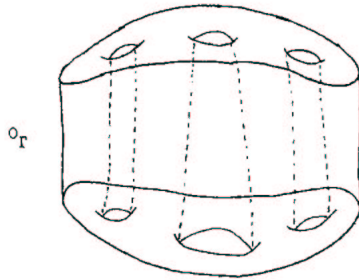PROPOSITION 8.7.2 (Marden). *For a torsion-free Kleinian group Γ, the following conditions are equivalent.*

(i) Γ *is quasi-Fuchsian.*
(ii) D_Γ *has precisely two components.*
(iii) Γ *is quasi-conformally conjugate to a Fuchsian group.*

PROOF. Clearly (iii) $\Longrightarrow$ (i) $\Longrightarrow$ (ii). To show (ii) $\Longrightarrow$ (iii), consider
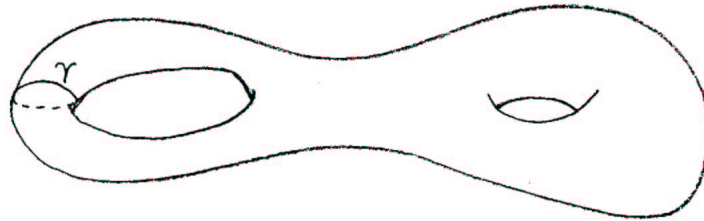
$$O_\Gamma = (H^3 \cup D_\Gamma)/\Gamma.$$

Suppose that no element of Γ interchanges the two components of D_Γ. Then O_Γ is a three-manifold with two boundary components (labelled, for example, N_1 and N_2), and 8.35

$\Gamma = \pi_1(O_\Gamma) \approx \pi_1(N_1) \approx \pi_1(N_2)$. By a well-known theorem about three-manifolds (see Hempel for a proof), this implies that O_Γ is homeomorphic to $N_1 \times I$. By the above discussion, this implies that Γ' is quasi-conformally conjugate to a Fuchsian group. A similar argument applies if O_Γ has one boundary component; in that case, O_Γ is the orientable interval bundle over a non-orientable surface. The reverse implication is clear. $\square$
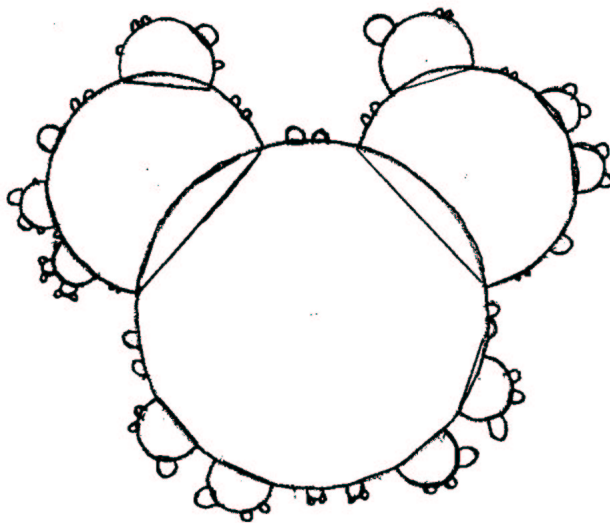
EXAMPLE 8.7.3 (Mickey mouse). Consider a hyperbolic structure on a surface of genus two. Let us construct a deformation of the corresponding Fuchsian group by bending along a single closed geodesci γ by an angle of $\pi/2$. This

8.36

will give rise to a quasi-Fuchsian group if the geodesic is short enough. We may visualize the limit set by imagining bending a hyperbolic plane along the lifts of γ, one by one.
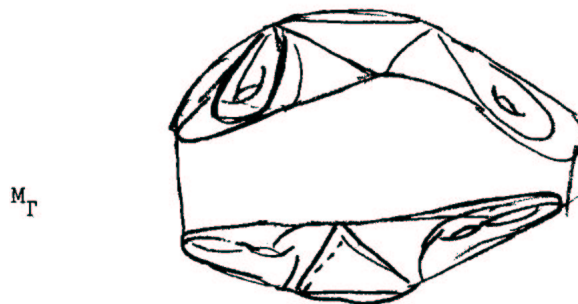
We want to understand how the geometry changes as we deform quasi-Fuchsian groups. Even though the topology doesn't change, geometrically things can become very complicated. For example, given any $\epsilon > 0$, there is a quasi-Fuchsian group Γ whose limit set L_Γ is ϵ-dense in S^2, and there are limits of quasi-Fuchsian groups with $L_\Gamma = S^2$.

Our goal here is to try to get a grasp of the geometry of the convex hull quotient $M = M_\Gamma$ of a quasi-Fuchsian group Γ. M_Γ is a convex hyperbolic manifold which is homeomorphic to $N^2 \times I$, and the two boundary components are hyperbolic surfaces bent along geodesic laminations.
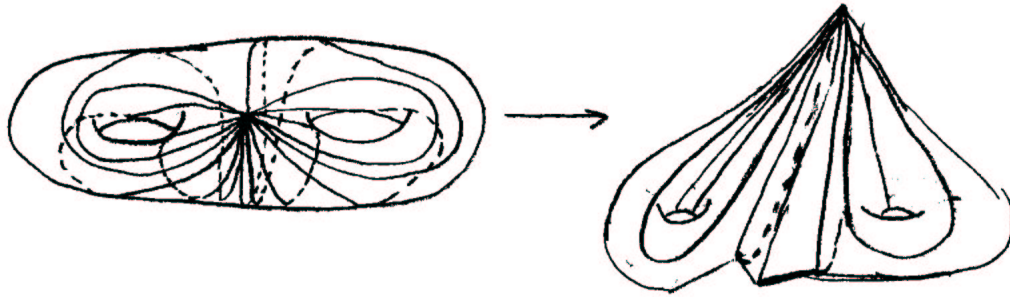
8.37

We also need to analyze intermediate surfaces in M_Γ. For example, what kinds of nice surfaces are embedded (or immersed) in M_Γ? Are there isometrically embedded cross sections? Are there cross sections of bounded area near any point in M_Γ?

Here are some ways to map in surfaces.

(a) Take the abstract surface N^2, and choose a "triangulation" of N with one vertex. Choose an arbitrary map of N into M. Then straighten the map (see §6.1).
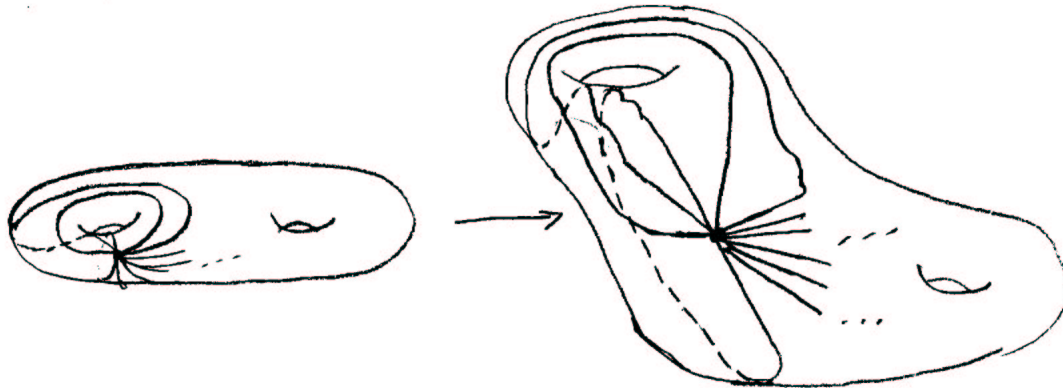
This is a fairly good way to map in a surface, since the surface is hyperbolic away from the vertex. There may be positive curvature concentrated at the vertex, however, since the sum of the angles around the vertex may be quite small. This map can be changed by moving the image of the vertex in M or by changing the triangulation on N.

8.38

(b) Here is another method, which insures that the map is not too bad near the vertex. First pick a closed loop in N, and then choose a vertex on the loop. Now extend this to a triangulation of N with one vertex. To map in N, first map

in the loop to the unique geodesic in M in its free homotopy class (this uses a homeomorphism of M to $N \times I$). Now extend this as in (a) to a piecewise straight map $f : N \to M$. The sum of the angles around the vertex is at least 2π, since there is a straight line segment going through the vertex (so the vertex cannot be spiked). It is possible to have the sum of the angles $> 2\pi$, in which case there is negative curvature concentrated near the vertex.

(c) Here is a way to map in a surface with constant negative curvature. Pick an example, as in (b), of a triangulation of N coming from a closed geodesic, and map N as in (b). Consider the isotopy obtained by moving the vertex around the loop more and more. The loop stays the same, but the other line segments start spiraling
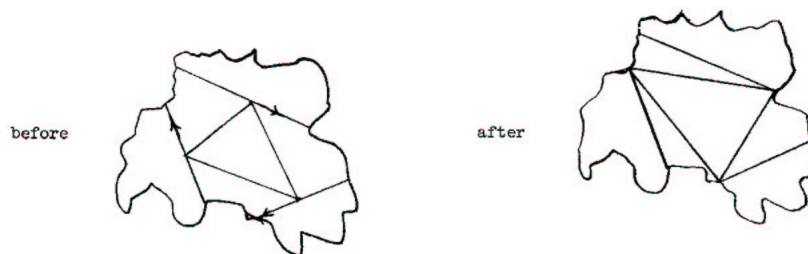
8.39

around the loop, more and more, converging, in the limit, to a geodesic laminated set. The surface N maps into M at each finite stage, and this carries over in the limit to an isometric embedding of a hyperbolic surface. The triangles with an edge on the fixed loop have disappeared in the limit. Compare 3.9.

One can picture what is going on by looking upstairs at the convex hull $H(L_\Gamma)$. The lift $\tilde{f} : \tilde{N} \to H(L_\Gamma)$ of the map from the original triangulation (before isotoping the vertex) is defined as follows. First the geodesic (coming from the loop) and its conjugates are mapped in (these are in the convex hull since their

8.40

endpoints are in L_Γ). The line segments connect different conjugates of the geodesic, and the triangles either connect three distinct conjugates or two conjugates (when the original loop is an edge of the triangle). As we isotope the vertex around the loop, the image vertices slide along the geodesic (and its conjugates), and in the limit the triangles become asymptotic (and the triangles connecting only two conjugates disappear).
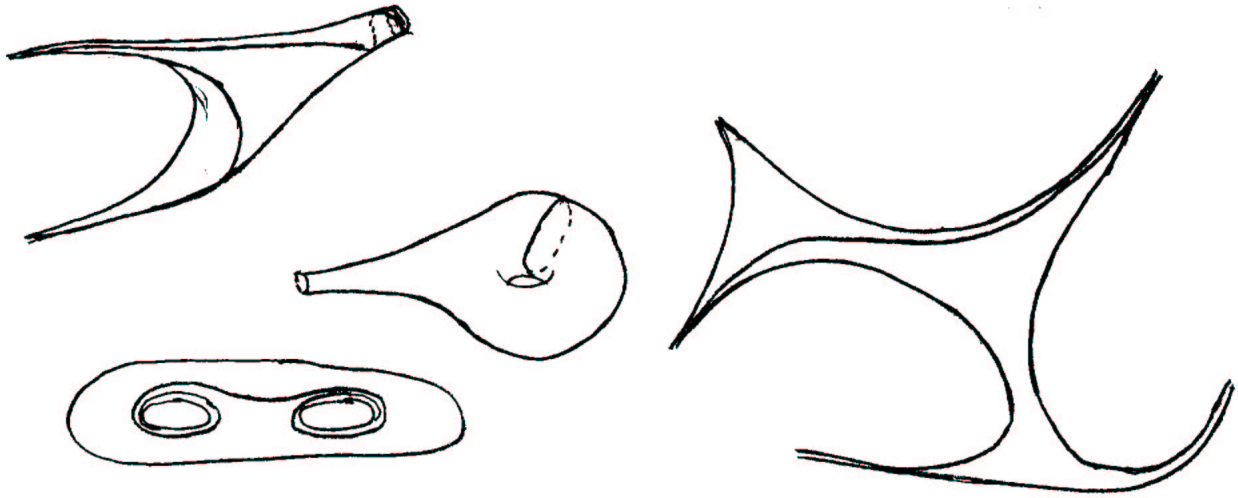
The above method works because the complement of the geodesic lamination (obtained by spinning the triangulation) consists solely of asymptotic triangles. Here is a more general method of mapping in a surface N by using geodesic laminations.

DEFINITION 8.7.5. A geodesic lamination γ on hyperbolic surface S is *complete* if the complementary regions in $S - \gamma$ are all asymptotic triangles.

PROPOSITION 8.7.6. *Any geodesic lamination γ on a hyperbolic surface S can be completed, i.e., γ can be extended to a complete geodesic lamination $\gamma' \supset \gamma$ on S.*

PROOF. Suppose γ is not complete, and pick a complementary region A which is not an asymptotic triangle. If A is simply connected, then it is a finite-sided asymptotic polygon, and it is easy to divide A into asymptotic triangles by adding simple geodesics. If A is not simply connected, extend γ to a larger geodesic lamination by adding a simple geodesic α in A
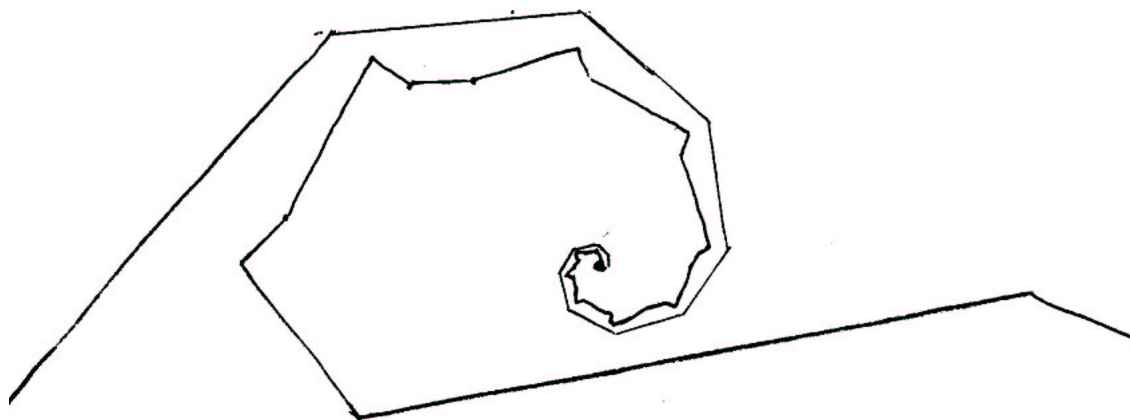
8.41

(being careful to add a *simple* geodesic). Either α separates A into two pieces (each of which has less area) or α does not separate A (in which case, cutting along α reduces the rank of the homology. Continuing inductively, after a finite number of steps A separates into asymptotic triangles. $\qquad\square$

Completeness is exactly the property we need to map in surfaces by using geodesic laminations.

PROPOSITION 8.7.7. *Let S be an oriented hyperbolic surface, and Γ a quasi-Fuchsian group isomorphic to $\pi_1 S$. For every complete geodesic lamination γ on S, there is a unique hyperbolic surface $S' \approx S$ and an isometric map $f : S' \to M_\Gamma$ which is straight (totally geodesic) in the complement of γ. (γ here denotes the corresponding geodesic lamination on any hyperbolic surface homeomorphic to S.)*
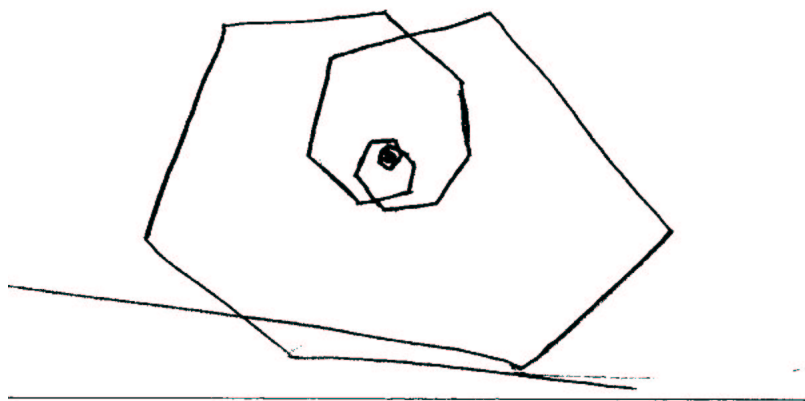
REMARK. By an isometric map $f : M_1 \to M_2$ from one Riemannian manifold to another, we mean that for every rectifiable path α_t in M_1, $f \circ \alpha_t$ is rectifiable and has the same length as α_t. When f is differentiable, this means that df preserves lengths of tangent vectors. We shall be dealing with maps which are not usually differentiable, however. Our maps are likely not even to be local embeddings. A cross-section of the image of a surface mapped in by method (c) has two polygonal spiral branches, if the closed geodesic corresponds to a covering transformation which is not a pure translation:

8.42

(This picture is obtained by considering triangles in H^3 asymptotic to a loxodromic axis, together with their translates.)

If the triangulation is spun in opposite directions on opposite sides of the geodesic, the polygonal spiral have opposite senses, so there are infinitely many self-intersections.
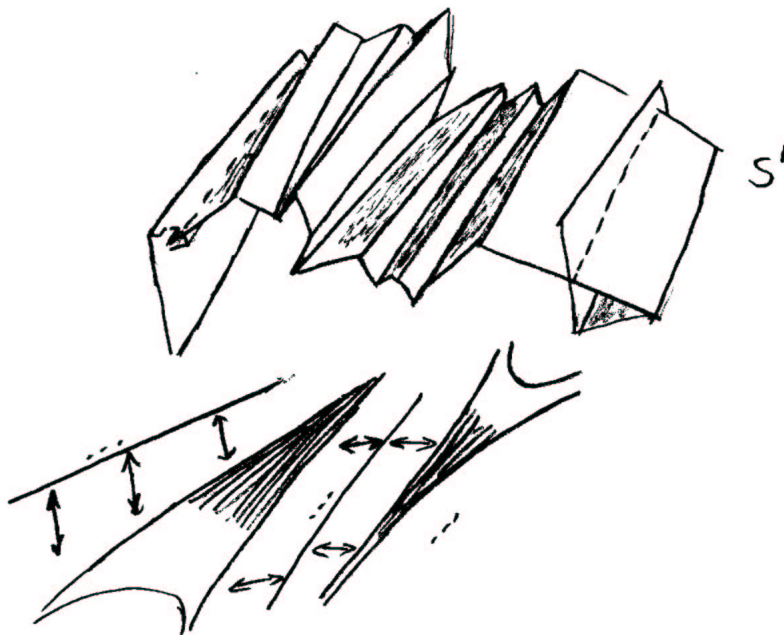

8.43

PROOF. The hyperbolic surface $\tilde{S}'$ is constructed out of pieces. The asymptotic triangles in $\tilde{S} - \tilde{\gamma}$ are determined by triples of points on S^1_∞. We have a canonical identification of S^1_∞ with L_Γ; the corresponding triple of points in L_Γ spans a triangle in H^3, which will be a piece of $\tilde{S}'$. Similarly, corresponding to each leaf of $\tilde{\gamma}$ there is a canonical line in H^3. These triangles and lines fit together just as on $\tilde{S}$; from this the picture of $\tilde{S}'$ should be clear. Here is a formal definition. Let P_γ be the set of all "pieces" of $\tilde{\gamma}$, i.e., P_γ consists of all leaves of $\tilde{\gamma}$, together with all components of $\tilde{S} - \tilde{\gamma}$. Let P_γ have the (non-Hausdorff) quotient topology. The universal cover $\tilde{S}'$

is defined first, to consist of ordered pairs (x, p), where $p \in P_\gamma$ and x is an element of the piece of H^3 corresponding to p. Γ acts on this space $\tilde{S}'$ in an obvious way; the quotient space is defined to be S'. It is not hard to find local coordinates for S', showing that it is a (Hausdorff) surface.
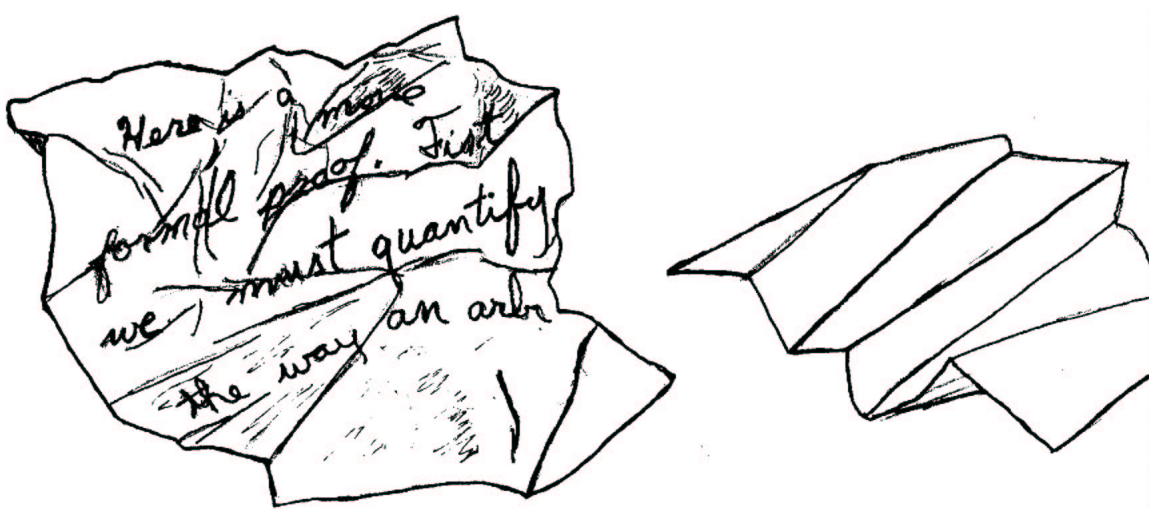

8.44

An appeal to geometric intuition demonstrates that S' is a hyperbolic surface, mapped isometrically to M_Γ, straight in the complement of γ. Uniqueness is evident from consideration of the circle at ∞. $\qquad\square$

REMARK. There are two approaches which a reader who prefers more formal proofs may wish to check. The first approach is to verify 8.7.7 first for laminations all of whose leaves are either isolated or simple limits of other leaves (as in (c)), and then extend to all laminations by passing to limits, using compactness properties of uncrumpled surfaces (§8.8). Alternatively, he can construct the hyperbolic structure on S' directly by describing the local developing map, as a limit of maps obtained by considering only finitely many local flat pieces. Convergence is a consequence of the finite total area of the flat pieces of S'.

8.8. Uncrumpled surfaces

There is a large qualitative difference between a crumpled sheet of paper and one which is only wrinkled or crinkled. Crumpled paper has fold lines or bending lines going any which way, often converging in bad points.
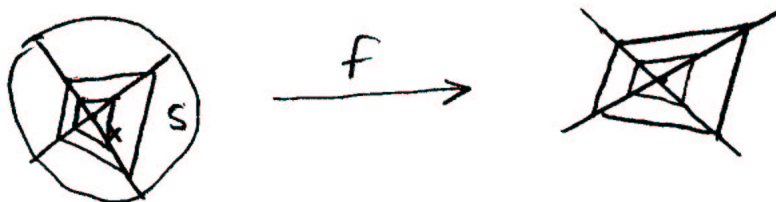
DEFINITION 8.8.1. An *uncrumpled surface* in a hyperbolic three-manifold N is a complete hyperbolic surface S of finite area, together with an isometric map $f : S \to N$ such that every $x \in S$ is in the interior of some straight line segment which is mapped by f to a straight line segment. Also, f must take every cusp of S to a cusp of N.

The set of uncrumpled surfaces in N has a well-behaved topology, in which two surfaces $f_1 : S_1 \to N$ and $f_2 : S_2 \to N$ are close if there is an approximate isometry $\phi : S_1 \to S_2$ making f_1 uniformly close to $f_2 \circ \phi$. Note that the surfaces have no preferred coordinate systems.

Let $\gamma \subset S$ consist of those points in the uncrumpled surfaces which are in the interior of unique line segments mapped to line segments.

PROPOSITION 8.8.2. *γ is a geodesic lamination. The map f is totally geodesic in the complement of γ.*

PROOF. If $x \in S - \gamma$, then there are two transverse line segments through x mapped to line segments. Consider any quadrilateral about x with vertices on these segments; since f does not increase distances, the quadrilateral must be mapped to a plane. Hence, a neighborhood of x is mapped to a plane.
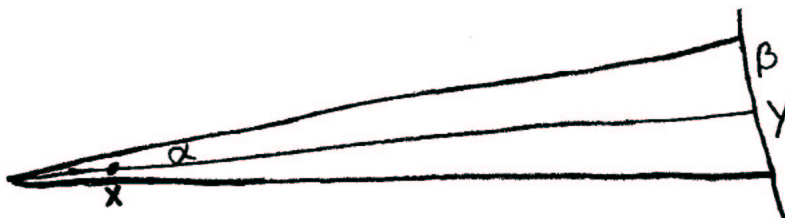
Consider now any point $x \in \gamma$, and let α be the unique line segment through x which is mapped straight. Let α be extended indefinitely on S. Suppose there were some point y on α in the interior of some line segment $\beta \not\subset \alpha$ which is mapped straight. One may assume that the segment $\overline{xy}$ of α is mapped straight. Then, by considering long skinny triangles with two vertices on β and one vertex on α, it would follow that a neighborhood of x is mapped to a plane—a contradiction.

Thus, the line segments in γ can be extended indefinitely without crossings, so γ must be a geodesic lamination. $\qquad\square$

If $U = S \xrightarrow{f} N$ is an uncrumpled surface, then this geodesic lamination $\gamma \subset S$ (which consists of points where U is not locally flat) is the *wrinkling locus* $\omega(U)$.

The *modular space* $\mathcal{M}(S)$ of a surface S of negative Euler characteristic is the space of hyperbolic surfaces with finite area which are homeomorphic to S. In other words, $\mathcal{M}(S)$ is the Teichmüller space $\mathcal{T}(S)$ modulo the action of the group of homeomorphisms of S.

8.47

PROPOSITION 8.8.3 (Mumford). *For a surface S, the set $A_\epsilon \subset \mathcal{M}(S)$ consisting of surfaces with no geodesic shorter than ϵ is compact.*
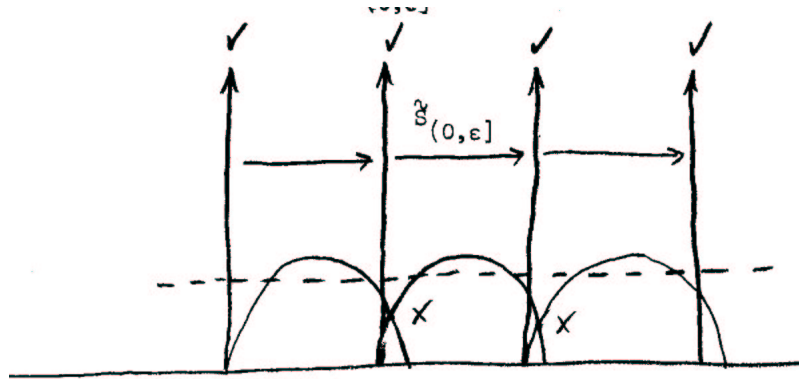
PROOF. By the Gauss–Bonnet theorem, all surfaces in $\mathcal{M}(S)$ have the same area. Every non-compact component of $S_{(0,\epsilon]}$ is isometric to a standard model, so the result follows as the two-dimensional version of a part of 5.12. (It is also not hard to give a more direct specifically two-dimensional geometric argument.) $\qquad\square$

Denote by $\mathcal{U}(S, N)$ the space of uncrumpled surfaces in N homeomorphic to S with $\pi_1(S) \to \pi_1(N)$ injective. There is a continuous map $\mathcal{U}(S, N) \to \mathcal{M}(S)$ which forgets the isometric map to N.

The behavior of an uncrumpled surface near a cusp is completely determined by its behavior on some compact subset. To see this, first let us prove

PROPOSITION 8.8.4. *There is some ϵ such that for every hyperbolic surface S and every geodesic lamination γ on S, the intersection of γ with every non-compact component of $S_{(0,\epsilon]}$ consists of lines tending toward that cusp.*

PROOF. Thus there are uniform horoball neighborhoods of the cusps of uncrumpled surfaces which are always mapped as cones to the cusp point. Uniform convergence of a sequence of uncrumpled surfaces away from the cusp points implies uniform convergence elsewhere. □

PROPOSITION 8.8.5. *Let $K \subset N$ be a compact subset of a complete hyperbolic manifold N. For any surface S_0, let $W \subset \mathcal{U}(S_0, N)$ be the subset of uncrumpled surfaces $S \xrightarrow{f} N$ such that $f(S)$ intersects K, and satisfying the condition*

(np) $\pi_1(f)$ *takes non-parabolic elements of $\pi_1 S$ to non-parabolic elements of $\pi_1 N$.*
Then W is compact.

PROOF. The first step is to bound the image of an uncrumpled surface, away from its cusps.

Let ϵ be small enough that for every complete hyperbolic three-manifold M, components of $M_{(0,\epsilon]}$ are separated by a distance of at least (say) 1. Since the area of surfaces in $\mathcal{U}(S_0, N)$ is constant, there is some number d such that any two points in an uncrumpled surface S can be connected (on S) by a path p such that $p \cap S_{[\epsilon,\infty)}$ has length $\leq d$.

If neither point lies in a non-compact component of $S_{(0,\epsilon]}$, one can assume, furthermore, that p does not intersect these components. Let $K' \subset N$ be the set of points which are connected to K by paths whose total length outside compact components of $N_{(0,\epsilon]}$ is bounded by d. Clearly K' is compact and an uncrumpled surface of W must have image in K', except for horoball neighborhoods of its cusps.

Consider now any sequence $S_1, S_2, \ldots$ in W. Since each short closed geodesic in S_i is mapped into K', there is a lower bound ϵ' to the length of such a geodesic, so by 8.8.3 we can pass to a subsequence such that the underlying hyperbolic surfaces converge in $\mathcal{M}(S)$. There are approximate isometries $\phi_i : S \to S_i$. Then the compositions $f_i \circ \phi_i : S \to N$ are equicontinuous, hence there is a subsequence converging uniformly on $S_{[\epsilon,\infty)}$. The limit is obviously an uncrumpled surface. [To make the picture

clear, one can always pass to a further subsequence to make sure that the wrinkling laminations γ_i of S_i converge geometrically.] □

COROLLARY 8.8.6. (a) *Let S be any closed hyperbolic surface, and N any closed hyperbolic manifold. There are only finitely many conjugacy classes of subgroups $G \subset \pi_1 N$ isomorphic to $\pi_1 S$.*
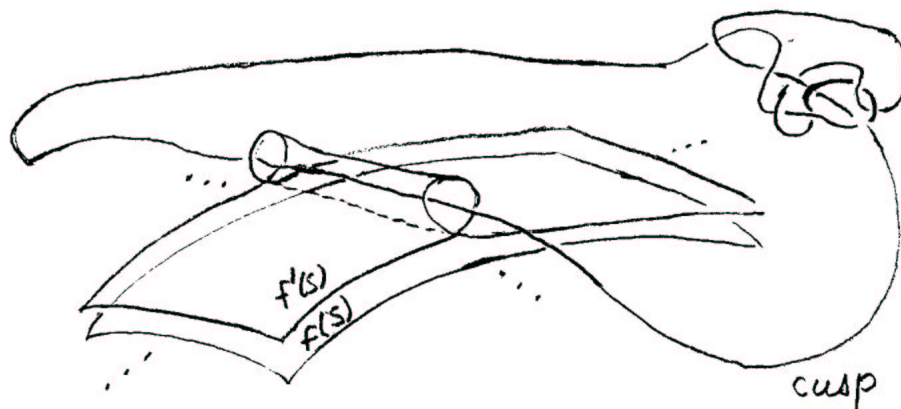
(b) *Let S be any surface of finite area and N any geometrically finite hyperbolic three-manifold. There are only finitely many conjugacy classes of subgroups $G \subset \pi_1 N$ isomorphic to $\pi_1 S$ by an isomorphism which preserves parabolicity (in both directions).*

PROOF. Statement (a) is contained in statement (b). The conjugacy class of every subgroup G is represented by a homotopy class of maps of S into N, which is homotopic to an uncrumpled surface (say, by method (c) of §8.7). Nearby uncrumpled surfaces represent the same conjugacy class of subgroups. Thus we have an open cover of the space W by surfaces with conjugate subgroups; by 8.8.5, this is a finite subcover. □

8.50

REMARK. If non-parabolic elements of $\pi_1 S$ are allowed to correspond to parabolic elements of $\pi_1 N$, then this statement is no longer true.

In fact, if $S \xrightarrow{f} N$ is any surface mapped into a hyperbolic manifold N of finite volume such that a non-peripheral simple closed curve γ in S is homotopic to a cusp of N, one can modify f in a small neighborhood of γ to wrap this annulus a number of times around the cusp. This is likely to give infinitely many homotopy classes of surfaces in N.

In place of 8.8.5, there is a compactness statement in the topology of geometric convergence provided each component of $S_{[\epsilon,\infty)}$ is required to intersect K. One would allow S to converge to a surface where a simple closed geodesic is pinched to yield a pair of cusps. From this, one deduces that there are finitely many classes of groups G
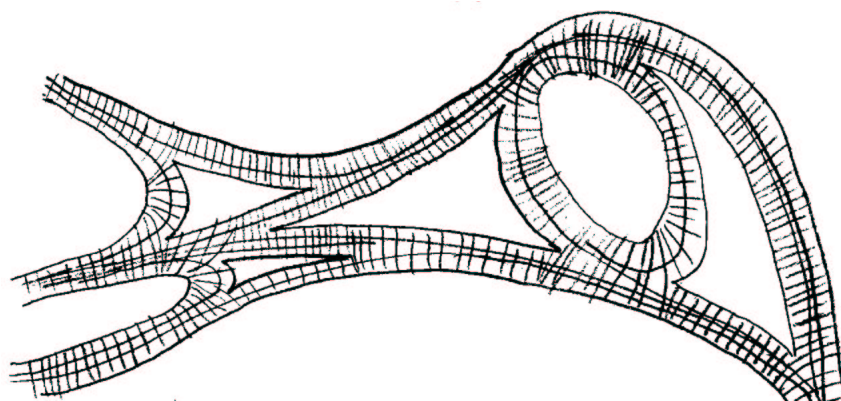
isomorphic to S up to the operations of conjugacy, and wrapping a surface carrying G around cusps.

Haken proved a finiteness statement analogous to 8.8.6 for embedded incompressible surfaces in atoroidal Haken manifolds.
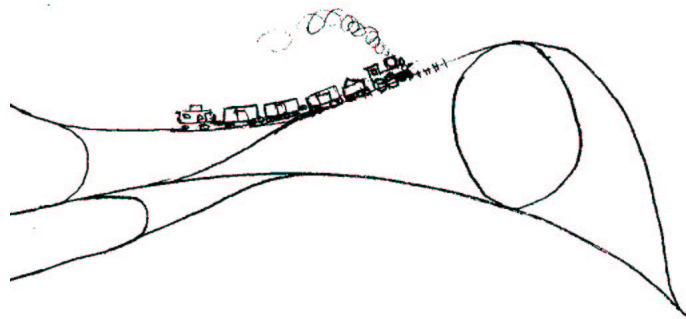
8.51

8.9. The structure of geodesic laminations: train tracks

Since a geodesic lamination γ on a hyperbolic surface S has measure zero, one can picture γ as consisting of many parallel strands in thin, branching corridors of S which have small total area.

Imagine squeezing the nearly parallel strands of γ in each corridor to a single strand. One obtains a train track τ (with switches) which approximates γ. Each leaf of γ may be imagined as the path of a train running around along τ.

8.52

Here is a construction which gives a precise and nice sequence of train track approximations of γ. Consider a complementary region R in $S - \gamma$. The double dR is a hyperbolic surface of finite area, so $(dR)_{(0,2\epsilon]}$ has a simple structure: it consists of neighborhoods of geodesics shorter than 2ϵ and of cusps. In each such neighborhood there is a canonical foliation by curves of constant curvature: horocycles about a cusp or equidistant curves about a short geodesic. Transfer this foliation to R, and then

to S. This yields a foliation $\mathcal{F}$ in the subset of S where leaves of γ are not farther than 2ϵ apart. (A local vector field tangent to $\mathcal{F}$ is Lipschitz, so it is integrable; this is why $\mathcal{F}$ exists. If γ has no leaves tending toward a cusp, then we can make all the leaves of $\mathcal{F}$ be arbitrarily short arcs by making ϵ sufficiently small. If γ has leaves tending toward a cusp, then there can be only finitely many such leaves, since there is an upper bound to the total number of cusps of the complementary regions. Erase all parts of $\mathcal{F}$ in a cusp of a region tending toward a cusp of S; again, when ϵ is sufficiently small all leaves of $\mathcal{F}$ will be short arcs. The space obtained by collapsing all arcs of $\mathcal{F}$ to a point is a surface S' homeomorphic to S, and the image of γ is a train track τ_ϵ on S'. Observe that each switch of τ_ϵ comes from a boundary component of some $dR_{(0,2\epsilon]}$. In particular, there is a uniform bound to the number of switches. From this it is easy to see that there are only finitely many possible types of τ_ϵ, up to homeomorphisms of S' (not necessarily homotopic to the identity).

8.53

In working with actual geodesic laminations, it is better to use more arbitrary train track approximations, and simply sketch pictures; the train tracks are analogous to decimal approximations of real numbers.

Here is a definition of a useful class of train tracks.

DEFINITIONS 8.9.1. A *train track* on a differentiable surface S is an embedded graph τ on S. The edges (branch lines) of τ must be C^1, and all edges at a given vertex (switch) must be tangent. If S has "cusps", τ may have open edges tending toward the cusps. Dead ends are not permitted. (Each vertex v must be in the interior of a C^1 interval on τ through v.) Furthermore, for each component R of $S - \tau$, the double dR of R along the interiors of edges of ∂R must have negative Euler characteristic. A lamination γ on S is *carried* by τ if there is a differentiable map $f : S \to S$ homotopic to the identity taking γ to τ and non-singular on the

tangent spaces of the leaves of γ. (In other words, the leaves of γ are trains running around on τ.) The lamination γ is *compatible* with τ if τ can be enlarged to a train track τ' which carries γ.

8.54

PROPOSITION 8.9.2. *Let S be a hyperbolic surface, and let $\delta > 0$ be arbitrary. There is some $\epsilon > 0$ such that for all geodesic laminations γ of S, the train track approximation τ_ϵ can be realized on S in such a way that all branch lines τ_ϵ are C^2 curves with curvature $< \delta$.*

PROOF. Note first that by making ϵ sufficiently small, one can make the leaves of the foliation $\mathcal{F}$ very short, uniformly for all γ: otherwise there would be a sequence of γ's converging to a geodesic lamination containing an open set. [One can also see this directly from area considerations.] When all branches of τ_ϵ are reasonably long, one can simply choose the tangent vectors to the switches to be tangent to any geodesic of γ where it crosses the corresponding leaf of $\mathcal{F}$; the branches can be filled in by curves of small curvature. When some of the branch lines are short, group each set of switches connected by very short branch lines together. First map each of these sets into S, then extend over the reasonably long branches. $\square$

COROLLARY 8.9.3. *Every geodesic lamination which is carried by a close train track approximation τ_ϵ to a geodesic lamination γ has all leaves close to leaves of γ.*

PROOF. This follows from the elementary geometrical fact that a curve in hyperbolic space with uniformly small curvature is uniformly close to a unique geodesic. (One way to see this is by considering the planes perpendicular to the curve—they always advance at a uniform rate, so in particular the curve crosses each one only once.)

8.55

$\square$

PROPOSITION 8.9.4. *A lamination λ of a surface S is isotopic to a geodesic lamination if and only if*
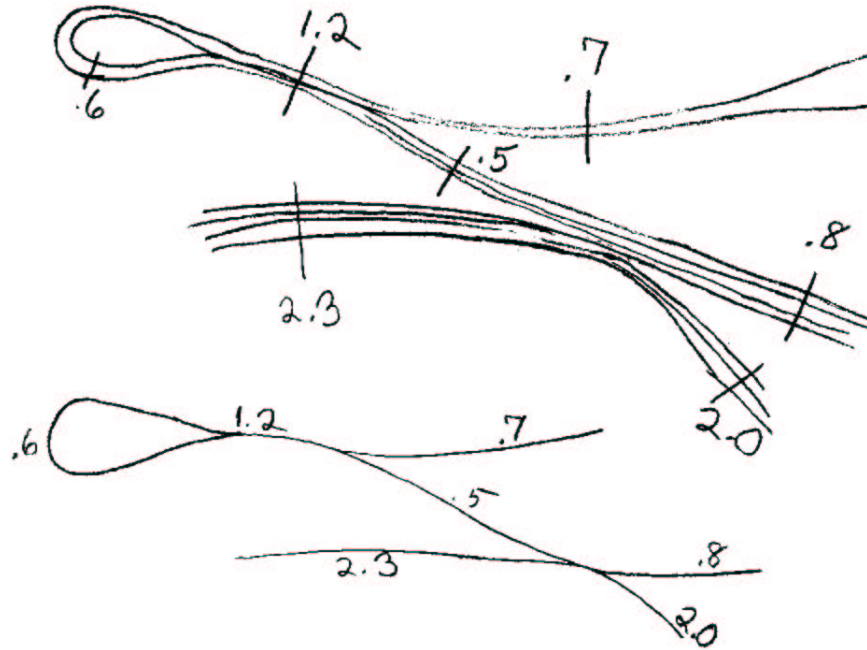
(a) *λ is carried by some train track τ, and*
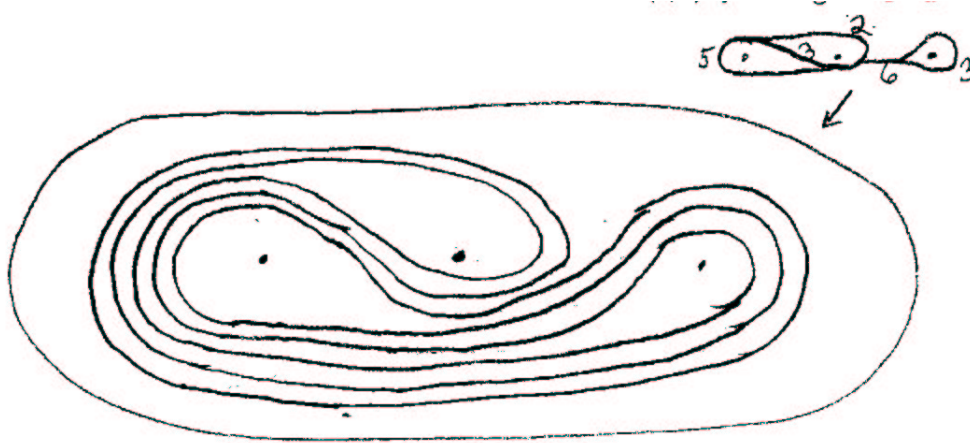(b) *no two leaves of λ take the same (bi-infinite) path on τ.*

PROOF. Given an arbitrary train track τ, it is easy to construct *some* hyperbolic structure for S on which τ is realized by lines with small curvature. The leaves of λ then correspond to a set of geodesics on S, near τ. These geodesics do not cross, since the leaves of λ do not. Condition (b) means that distinct leaves of λ determine distinct geodesics. When leaves of λ are close, they must follow the same path for a long finite interval, which implies the corresponding geodesics are close. Thus, we obtain a geodesic lamination γ which is isotopic to λ. (To have an isotopy, it suffices to construct a homeomorphism homotopic to the identity. This homeomorphism is constructed first in a neighborhood of τ, then on the rest of S.) □

REMARK. From this, one sees that as the hyperbolic structure on S varies, the corresponding geodesic laminations are all isotopic. This issue was quietly skirted in §8.5.

When a lamination λ has an invariant measure μ, this gives a way to associate a number $\mu(b)$ to each branch line b of any train track which dominates γ: $\mu(b)$ is just the transverse measure of the leaves of λ collapsed to a point on b. At a switch, the sum of the "entering" numbers equals the sum of the "exiting" numbers.

8.56

Conversely, any assignment of numbers satisfying the switch condition determines a unique geodesic lamination with transverse measure: first widen each branch line b of τ to a corridor of constant width $\mu(b)$, and give it a foliation $\mathcal{G}$ by equally spaced lines.

8.57

As in 8.9.4, this determines a lamination γ; possibly there are many leaves of $\mathcal{G}$ collapsed to a single leaf of γ, if these leaves of $\mathcal{G}$ all have the same infinite path. $\mathcal{G}$ has a transverse measure, defined by the distance between leaves; this goes over to a transverse measure for γ.

8.10. Realizing laminations in three-manifolds

For a quasi-Fuchsian group Γ, it was relatively easy to "realize" a geodesic lamination of the corresponding surface in M_Γ, by using the circle at infinity. However, not every complete hyperbolic three-manifold whose fundamental group is isomorphic to a surface group is quasi-Fuchsian, so we must make a more systematic study of realizability of geodesic laminations.

DEFINITION 8.10.1. Let $f : S \to N$ be a map of a hyperbolic surface to a hyperbolic three-manifold which sends cusps to cusps. A geodesic lamination γ on S is *realizable* in the homotopy class of f if f is homotopic (by a cusp-preserving homotopy) to a map sending each leaf of γ to a geodesic.

PROPOSITION 8.10.2. *If γ is realizable in the homotopy class of f, the realization is (essentially) unique: that is, the image of each leaf of γ is uniquely determined.*

PROOF. Consider a lift $\tilde{h}$ of a homotopy connecting two maps of S into N. If S is closed, $\tilde{h}$ moves every point a bounded distance, so it can't move a geodesic to a different geodesic. If S has cusps, the homotopy can be modified near the cusps of S so $\tilde{h}$ again is bounded. $\square$

8.58

In Section 8.5, we touched on the notion of geometric convergence of geodesic laminations. The *geometric topology* on geodesic laminations is the topology of geometric convergence, that is, a neighborhood of γ consists of laminations γ' which

have leaves near every point of γ, and nearly parallel to the leaves of γ. If γ_1 and γ_2 are disjoint simple closed curves, then $\gamma_1 \cup \gamma_2$ is in every neighborhood of γ_1 as well as in every neighborhood of γ_2. The space of geodesic laminations on S with the geometric topology we shall denote $\mathcal{GL}$. The geodesic laminations compatible with train track approximations of γ give a neighborhood basis for γ.

The *measure topology* on geodesic laminations with transverse measures (of full support) is the topology induced from the weak topology on measures in the Möbius band J outside S_∞ in the Klein model. That is, a neighborhood of (γ, μ) consists of (γ', μ') such that for a finite set $f_1, \ldots, f_k$ of continuous functions with compact support in J,

$$\left| \int f_i \, d\mu - \int f_i \, d\mu' \right| < \epsilon.$$

This can also be interpreted in terms of integrating finitely many continuous functions on finitely many *transverse* arcs. Let $\mathcal{ML}(S)$ be the space of (γ, μ) on S with the measure topology. Let $\mathcal{PL}(S)$ be $\mathcal{ML}(S)$ modulo the relation $(\gamma, \mu) \sim (\gamma, a\mu)$ where $a > 0$ is a real number.

PROPOSITION 8.10.3. *The natural map* $\mathcal{ML} \to \mathcal{GL}$ *is continuous.*

8.59

PROPOSITION 8.10.4. *The map* $w : \mathcal{U}(S, N) \to \mathcal{GL}(S)$ *which assigns to each uncrumpled surface its wrinkling locus is continuous.*

PROOF OF 8.10.3. For any point x in the support of a measure μ and any neighborhood UU of x, the support of a measure close enough to μ must intersect U. $\square$

PROOF OF 8.10.4. An interval which is bent cannot suddenly become straight. Away from any cusps, there is a positive infimum to the "amount" of bending of an interval of length ϵ which intersects the wrinkling locus $w(S)$ in its middle third, and makes an angle of at least ϵ with $w(S)$. (The "amount" of bending can be measure, say, by the different between the length of α and the distance between the image endpoints.) All such arcs must still cross $w(S')$ for any nearby uncrumpled surface S'. $\square$
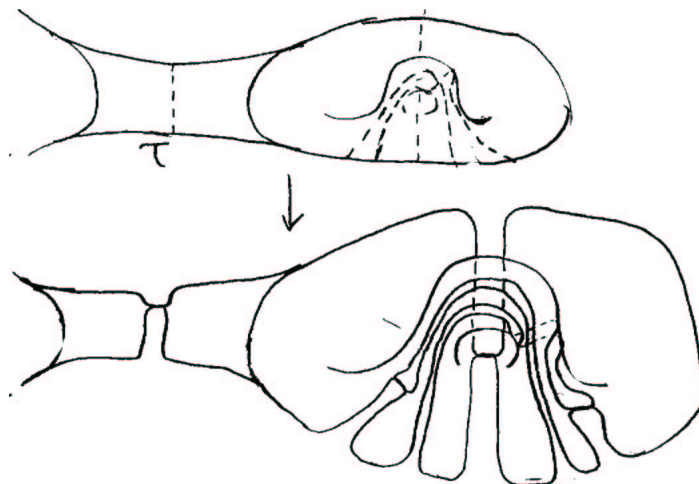
When S has cusps, we are also interested in measures supported on compact geodesic laminations. We denote this space by $\mathcal{ML}_0(S)$. If (τ, μ) is a train track *description* for (γ, μ), *where* $\mu(b) \neq 0$ *for any branch of* τ, then neighborhoods for (γ, μ) are described by $\{(\tau', \mu')\}$, where $\tau \subset \tau'$ and $\left| \mu(b) - \mu'(b) \right| < \epsilon$. (If b is a branch of τ' not in τ, then $\mu(b) = 0$ by definition.)

In fact, one can always choose a hyperbolic structure on S so that τ is a good *approximation* to γ. If S is closed, it is always possible to squeeze branches of τ

together along non-trivial arcs in the complementary regions to obtain a new train track which cannot be enlarged.

This implies that a neighborhood of (γ, μ) is parametrized by a finite number of real parameters. Thus, $\mathcal{ML}(S)$ is a manifold. Similarly, when S has cusps, $\mathcal{ML}(S)$ is a manifold with boundary $\mathcal{ML}_0(S)$.

PROPOSITION 8.10.5. $\mathcal{GL}(S)$ *is compact, and* $\mathcal{PL}(S)$ *is a compact manifold with boundary* $\mathcal{PL}_0(S)$ *if* S *is not compact.*

PROOF. There is a finite set of train tracks $\tau_1, \ldots, \tau_k$ carrying every possible geodesic lamination. (There is an upper bound to the length of a compact branch of τ_ϵ, when S and ϵ are fixed.) The set of projective classes of measures on any particular τ is obviously compact, so this implies $\mathcal{PL}(S)$ is compact. That $\mathcal{PL}(S)$ is a manifold follows from the preceding remarks. Later we shall see that in fact it is the simplest of possible manifolds.

In 8.5, we indicated one proof of the compactness of $\mathcal{GL}(S)$. Another proof goes as follows. First, note that

PROPOSITION 8.10.6. *Every geodesic lamination* γ *admits some transverse measure* μ *(possibly with smaller support).*

PROOF. Choose a finite set of transversals $\alpha_1, \ldots, \alpha_k$ which meet every leaf of γ. Suppose there is a sequence $\{l_i\}$ of intervals on leaves of γ such that the total number N_1 of intersection of l_i with the α_j's goes to infinity. Let μ_i be the measure on $\bigcup \alpha_j$ which is $1/N_i$ times the counting measure on $l_i \cap \alpha_j$. The sequence $\{\mu_i\}$ has a subsequence converging (in the weak topology) to a measure μ. It is easy to see that μ is invariant under local projections along leaves of γ, so that it determines a transverse measure.
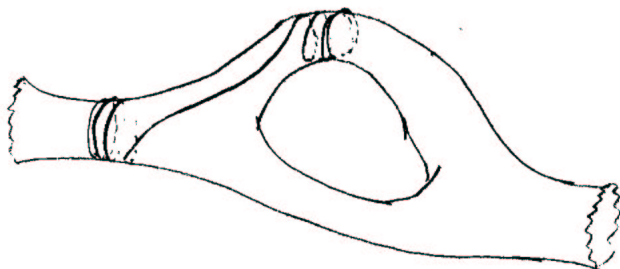
If there is no such sequence $\{l_i\}$ then every leaf is proper, so the counting measure for any leaf will do. □

We continue with the proof of 8.10.5. Because of the preceding result, the image $\mathfrak{I}$ of $\mathcal{PL}(S)$ in $\mathcal{GL}(S)$ intersects the closure of every point of $\mathcal{GL}(S)$. Any collection of open sets which covers $\mathcal{GL}(S)$ has a finite subcollection which covers the compact set $\mathfrak{I}$; therefore, it covers all of $\mathcal{GL}(S)$. □
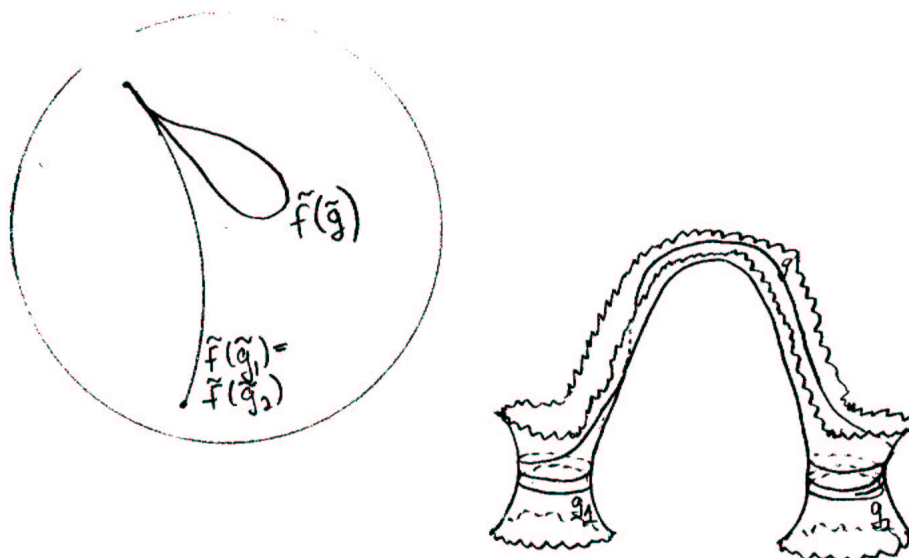
Armed with topology, we return to the question of realizing geodesic laminations. Let $\mathcal{R}_f \subset \mathcal{GL}(S)$ consist of the laminations realizable in the homotopy class of f.

First, if γ consists of finitely many simple closed geodesics, then γ is realizable   8.62
provided $\pi_1(f)$ maps each of these simple closed curves to non-trivial, non-parabolic elements.

If we add finitely many geodesics whose ends spiral around these closed geodesics or converge toward cusps the resulting lamination is also realizable except in the degenerate case that f restricted to an appropriate non-trivial pair of pants on S factors through a map to S^1.

To see this, consider for instance the case of a geodesic g on S whose ends spiral around closed geodesics g_1 and g_2. Lifting f to H^3, we see that the two ends of $\tilde{f}(\tilde{g})$ are asymptotic to geodesics $\tilde{f}(\tilde{g}_1)$ and $\tilde{f}(\tilde{g}_2)$. Then f is homotopic to a map taking g to a geodesic unless $\tilde{f}(\tilde{g}_1)$ and $\tilde{f}(\tilde{g}_2)$ converge to the same point on S_∞, which can only happen if $\tilde{f}(\tilde{g}_1) = \tilde{f}(\tilde{g}_2)$ (by 5.3.2). In this case, f is homotopic to a map taking a neighborhood of $g \cup g_1 \cup g_2$ to $f(g_1) = f(g_2)$.   8.63
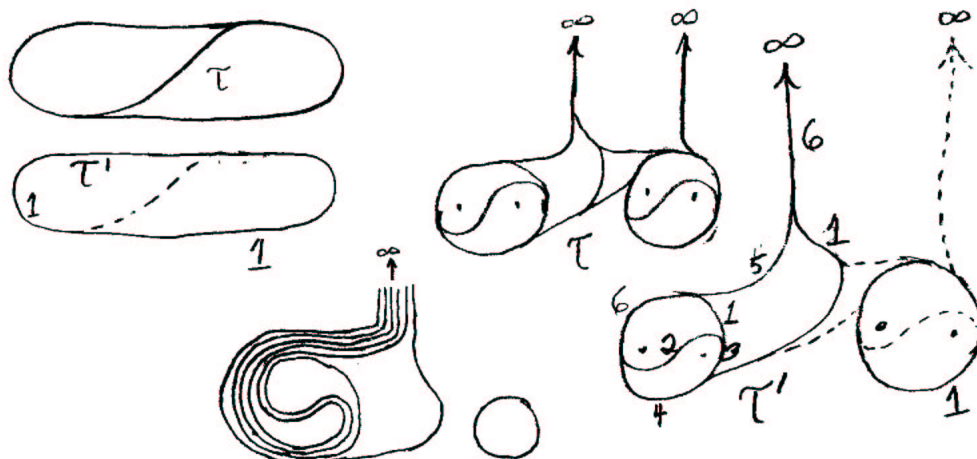
The situation is similar when the ends of g tend toward cusps.

These realizations of laminations with finitely many leaves take on significance in view of the next result:

PROPOSITION 8.10.7.   (a) *Measures supported on finitely many compact or proper geodesics are dense in $\mathcal{ML}$.*
  (b) *Geodesic laminations with finitely many leaves are dense in $\mathcal{GL}$.*
  (c) *Each end of a non-compact leaf of a geodesic lamination with only finitely many leaves spirals around some closed geodesic, or tends toward a cusp.*
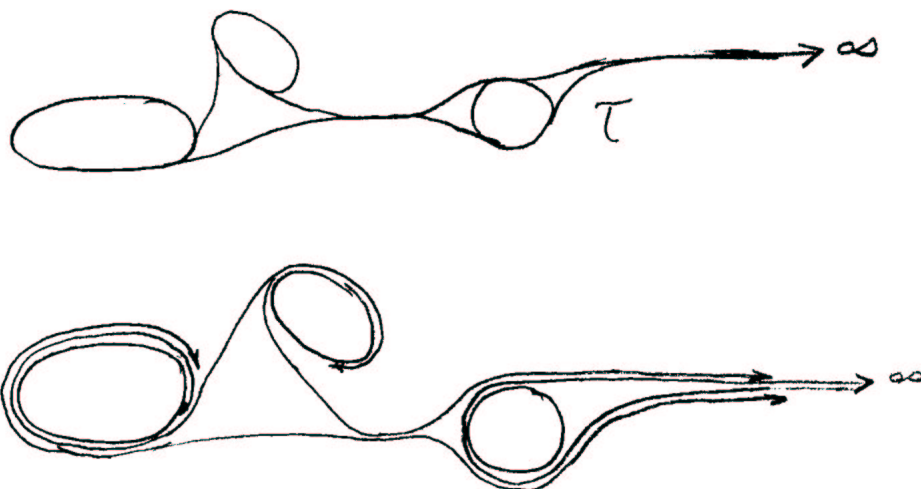
PROOF. If τ is any train track and μ is any measure which is positive on each branch, μ can be approximated by measures μ' which are rational on each branch, since μ is subject only to linear equations with integer coefficients. μ' gives rise to geodesic laminations with only finitely many leaves, all compact or proper. This proves (a).

If γ is an arbitrary geodesic lamination, let τ be a close train track approximation of γ and proceed as follows. Let $\tau' \subset \tau$ consist of all branches b of τ such that there exists either a cyclic (repeating) train route or a proper train route through b.
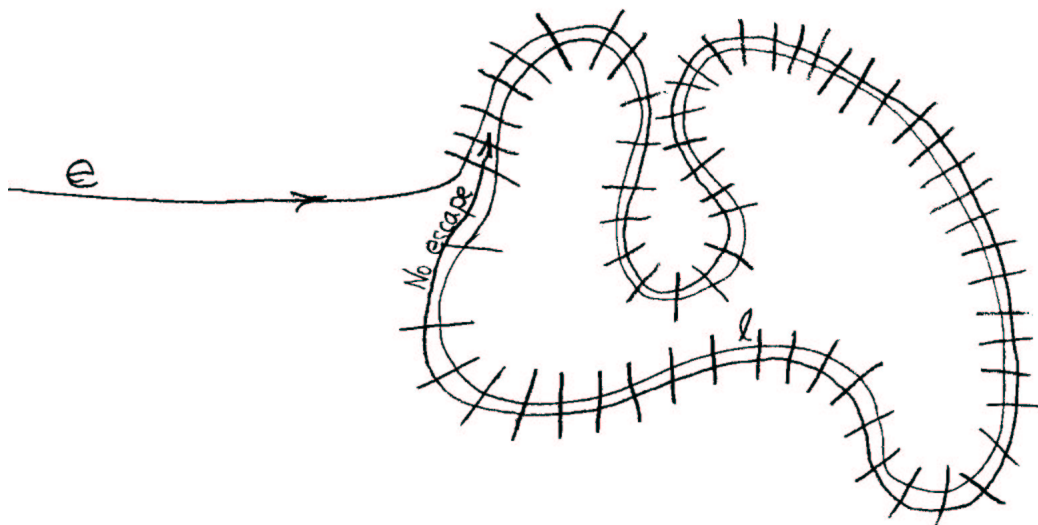
8.64

(The reader experienced with toy trains is aware of the subtlety of this question.) There is a measure supported on τ', obtained by choosing a finite set of cyclic and proper paths covering τ' and assigning to a branch b the total number of times these paths traverse. Thus there is a lamination λ' consisting of finitely many compact or proper leaves supported in a narrow corridor about τ'. Now let b be any branch of $\tau - \tau'$. A train starting on b can continue indefinitely, so it must eventually come to τ', in each direction. Add a leaf to λ' representing a shortest path from b to τ' in each direction; if the two ends meet, make them run along side by side (to avoid crossings). When the ends approach τ, make them "merge"—either spiral around a closed leaf, or follow along close to a proper leaf. Continue inductively in this way, adding leaves one by one until you obtain a lamination λ dominated by τ and filling out all the branches. This proves (b).

8.65

If γ is any geodesic lamination with finitely many (or even countably many) leaves, then the only possible minimal sets are closed leaves; thus each end e of a non-compact must either be a proper end or come arbitrarily near some compact leaf l. By tracing the leaves near l once around l, it is easy to see that this means e spirals around l. □
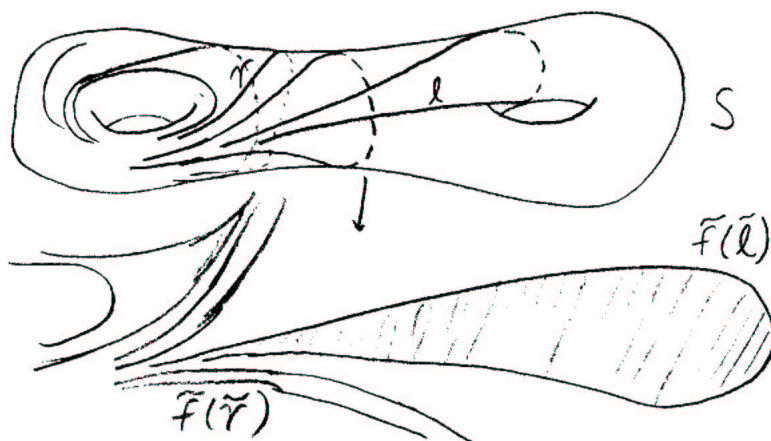
8.66

Thus, if f is non-degenerate, $\mathcal{R}_f$ is dense. Furthermore,

THEOREM 8.10.8. *If $\pi_1 f$ is injective, and f satisfies* (np) (*that is, if $\pi_1 f$ preserves non-parabolicity), then $\mathcal{R}_f$ is an open dense subset of $\mathcal{GL}(S)$.*

PROOF. If γ is any complete geodesic lamination which is realizable, then a train track approximation τ can be constructed for the image of γ in N^3, in such a way that all branch lines have curvature close to 0. Then all laminations carried by τ are also realizable; they form a neighborhood of γ. Next we will show that any enlargement $\gamma' \supset \gamma$ of a realizable geodesic lamination γ is also realizable. First note that if γ' is obtained by adding a single leaf l to γ, then γ' is also realizable. This is proved in the same way as in the case of a lamination with finitely many leaves: note that each end of l is asymptotic to a leaf of γ. (You can see this by considering $S - \gamma$.) If f is homotoped so that $f(\gamma)$ consists of geodesics, then both ends of $\tilde{f}(l)$ are asymptotic to geodesics in $\tilde{f}(\gamma)$. If the two endpoints were not distinct on S_∞, this would imply the existence of some non-trivial identification of γ by f so that $\pi_1 f$ could not be injective.

8.67

By adding finitely many leaves to any geodesic lamination γ' we can complete it. This implies that γ' is contained in the wrinkling locus of some uncrumpled surface. By 8.8.5 and 8.10.1, the set of uncrumpled surfaces whose wrinkling locus contains γ is compact. Since the wrinkling locus depends continuously on an uncrumpled surface, the set of $\gamma' \in \mathcal{R}_f$ which contains γ is compact. But any $\gamma' \supset \gamma$ can be approximated by laminations such that $\gamma' - \gamma$ consists of a finite number of leaves. This actually follows from 8.10.7, applied to $d(S - \gamma)$. Therefore, every enlargement $\gamma' \supset \gamma$ is in $\mathcal{R}_f$.

8.68

Since the set of uncrumpled surfaces whose wrinkling locus contains γ is compact, there is a finite set of train tracks $\tau_1, \ldots, \tau_k$ such that for any such surface, $w(S)$ is closely approximated by one of $\tau_1, \ldots, \tau_k$. The set of all laminations carried by at least one of the τ_i is a neighborhood of γ contained in $\mathcal{R}_f$. $\qquad\square$

COROLLARY 8.10.9. *Let Γ be a geometrically finite group, and let $f : S \to N_\Gamma$ be a map as in 8.10.8. Then either $\mathcal{R}_f = \mathcal{GL}(S)$ (that is, all geodesic laminations are realizable in the homotopy class of f), or Γ has a subgroup Γ' of finite index such that $N_{\Gamma'}$ is a three-manifold with finite volume which fibers over the circle.*

CONJECTURE 8.10.10. *If $f : S \to N$ is any map from a hyperbolic surface to a complete hyperbolic three-manifold taking cusps to cusps, then the image $\pi_1(f)(\pi_1(S))$ is quasi-Fuchsian if and only if $\mathcal{R}_f = \mathcal{GL}(S)$.*

PROOF OF 8.10.9. Under the hypotheses, the set of uncrumpled surfaces homotopic to $f(S)$ is compact. If each such surface has an essentially unique homotopy to $f(S)$, so that the wrinkling locus on S is well-defined, then the set of wrinkling loci of uncrumpled surfaces homotopic to f is compact, so by 8.10.8 it is all of $\mathcal{GL}(S)$.

Otherwise, there is some non-trivial $h_t : S \to M$ such that $h_1 = h_0 \circ \phi$, where $\phi : S \to S$ is a homotopically non-trivial diffeomorphism. It may happen that ϕ has

finite order up to isotopy, as when S is a finite regular covering of another surface in M. The set of all isotopy classes of diffeomorphisms ϕ which arise in this way form a group. If the group is finite, then as in the previous case, $\mathcal{R}_F = \mathcal{GL}(S)$. Otherwise, there is a torsion-free subgroup of finite index (see), so there is an element ϕ of infinite order. The maps f and $\phi \circ f$ are conjugate in Γ, by some element $\beta \in \Gamma$. The group generated by β and $f(\pi_1 S)$ is the fundamental group of a three-manifold which fibers over S^1. $\square$ 8.69

We shall see some justification for the conjecture in the remaining sections of chapter 8 and in chapter 9: we will prove it under certain circumstances.

8.11. The structure of cusps

Consider a hyperbolic manifold N which admits a map $f : S \to N$, taking cusps to cusps such that $\pi_1(f)$ is an isomorphism, where S is a hyperbolic surface. Let $B \subset N$ be the union of the components of $N_{(0,\epsilon]}$ corresponding to cusps of S. f is a relative homotopy equivalence from $(S, S_{(0,\epsilon)})$ to (N, B), so there is a homotopy inverse $g : (N, B) \to (S, S_{(0,\epsilon)})$. If $X \in S_{(\epsilon,\infty)}$ is a regular value for g, then $g^{-1}(x)$ is a one-manifold having intersection number one with $f(S)$, so it has at least one component homeomorphic to R, going out toward infinity in $N - B$ on opposite sides of $f(S)$. Therefore there is a proper function $h : (N - B) \to \mathbb{R}$ with h restricted to $g^{-1}(x)$ a surjective map. One can modify h so that $h^{-1}(0)$ is an incompressible surface. Since g restricted to $h^{-1}(0)$ is a degree one map to S, it must map the fundamental group surjectively as well as injectively, so $h^{-1}(0)$ is homeomorphic to S. $h^{-1}(0)$ divides $N - B$ into two components N_+ and N_- with $\pi_1 N = \pi_1 N_+ = \pi_1 N_- = \pi_1 S$. We can assume that $h^{-1}(0)$ does not intersect $N_{(0,\epsilon]}$ except in B (say, by shrinking ϵ).

Suppose that N has parabolic elements that are not parabolic on S. The structure   8.71 of the extra cusps of N is described by the following:
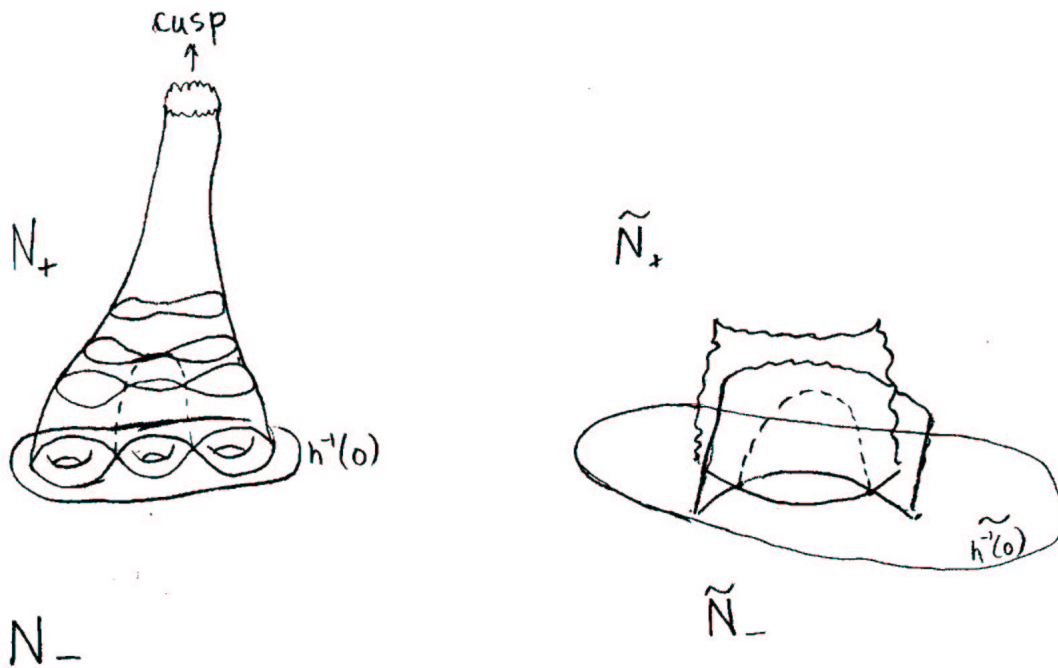
PROPOSITION 8.11.1. *There are geodesic laminations γ_+ and γ_- on S with all leaves compact (i.e., they are finite collections of disjoint simple closed curves) such that the extra cusps in N_e correspond one-to-one with leaves of $\gamma_e (e = +, -)$. In particular, for any element $\alpha \in \pi_1(S)$, $\pi_1(f)(\alpha)$ is parabolic if and only if α is freely homotopic to a cusp of S or to a simple closed curve in γ_+ or γ_-.*

PROOF. We need consider only one half, say N_+. For each extra cusp of N_+, there is a half-open cylinder mapped into N_+, with one end on $h^{-1}(0)$ and the other end tending toward the cusp. Furthermore, we can assume that the union of these cylinders is embedded outside a compact set, since we understand the picture in a neighborhood of the cusps. Homotope the ends of the cylinders which lie on $h^{-1}(0)$
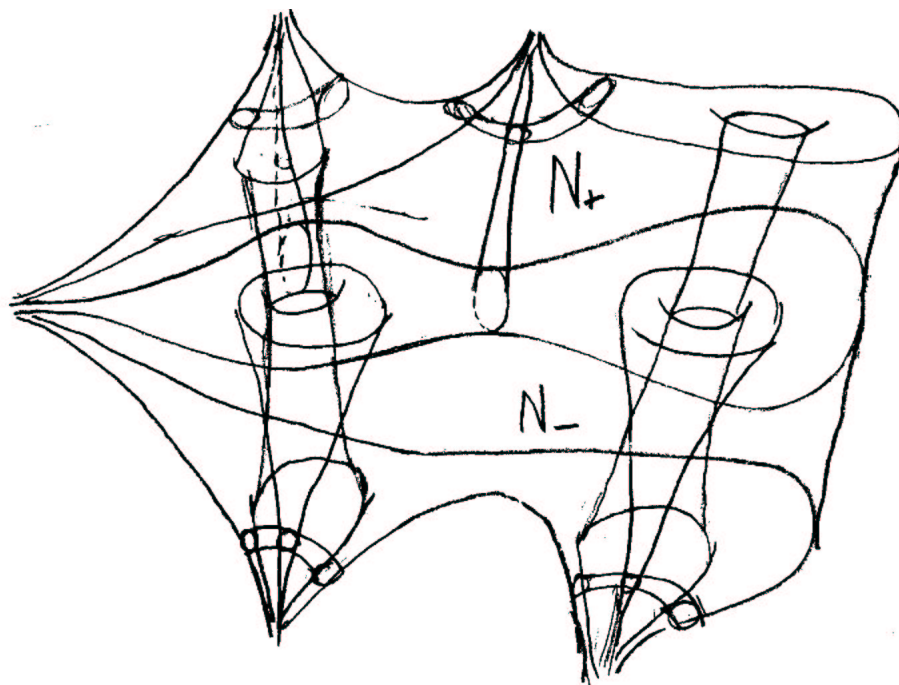
so they are geodesics in some hyperbolic structure on $h^{-1}(0)$. One can assume the cylinders are immersed (since maps of surfaces into three-manifolds are appoximable by immersions) and that they are transverse to themselves and to one another. If there are any self-intersections of the cylinders on $h^{-1}(0)$, there must be a double line which begins and ends on $h^{-1}(0)$. Consider the picture in $\tilde{N}$: there are two translates of universal covers of cylinders which meet in a double line, so that in particular their bounding lines meet twice on $h^{-1}(0)$. This contradicts the fact that they are geodesics in some hyperbolic structure. □

8.71

It actually follows that the collection of cylinders joining simple closed curves to the cusps can be embedded: we can modify g so that it takes each of the extra cusps to a neighborhood of the appropriate simple closed curve $\alpha \subset \gamma_\epsilon$, and then do surgery to make $g^{-1}(\alpha)$ incompressible.

8.72

To study N, we can replace S by various surfaces obtained by cutting S along curves in γ_+ or γ_-. Let P be the union of open horoball neighborhoods of all the cusps of N. Let $\{S_i\}$ be the set of all components of S cut by γ_+ together with those of S cut by γ_-. The union of the S_i can be embedded in $N - P$, with boundary on ∂P, within the convex hull M of N, so that they cut off a compact piece $N_0 \subset N - P$ homotopy equivalent to N, and non-compact ends E_i of $N - P$, with $\partial E_i \subset P \cup S_i$.

Let N now be an arbitrary hyperbolic manifold, and let P be the union of open horoball neighborhoods of its cusps. The picture of the structure of the cusps readily generalizes provided $N - P$ is homotopy equivalent to a compact submanifold N_O, obtained by cutting $N - P$ along finitely many incompressible surfaces $\{S_i\}$ with boundary ∂P.

Applying 8.11.1 to covering spaces of N corresponding to the S_i (or applying its proof directly), one can modify the S_i until no non-peripheral element of one of the S_i is homotopic, outside N_O, to a cusp. When this is done, the ends $\{E_i\}$ of $N - P$ are in one-to-one correspondence with the S_i.

According to a theorem of Peter Scott, every three-manifold with finitely generated fundamental group is homotopy equivalent to a compact submanifold. In general, such a submanifold will not have incompressible boundary, so it is not as well behaved. We will leave this case for future consideration.

8.73

DEFINITION 8.11.2. Let N be a complete hyperbolic manifold, P the union of open horoball neighborhoods of its cusps, and M the convex hull of N. Suppose

E is an end of $N - P$, with $\partial E - \partial P$ an incompressible surface $S \subset M$ homotopy equivalent to E. Then E is a *geometrically tame end* if either

(a) $E \cap M$ is compact, or
(b) the set of uncrumpled surfaces S' homotopic to S and with $S'_{[\epsilon, \infty)}$ contained in E is not compact.

If N has a compact submanifold N_O of $N - P$ homotopy equivalent to N such that $N - P - N_O$ is a disjoint union of geometrically tame ends, then N and $\pi_1 N$ are *geometrically tame*. (These definitions will be extended in § .) We shall justify this definition by showing geometric tameness implies that N is analytically, topologically and geometrically well-behaved.

8.12. Harmonic functions and ergodicity

Let N be a complete Riemannian manifold, and h a positive function on N. Let ϕ_t be the flow generated by $-(\operatorname{grad} h)$. The integral of the velocity of ϕ_t is bounded along any flow line:

$$\int_x^{\phi_T(x)} \| \operatorname{grad} h \| \, ds = h(x) - h(\phi_T(x))$$

$$\leq h(x) \quad (\text{ for } T > 0).$$

If A is a subset of a flow line $\{\phi_t(x)\}_{t \geq 0}$ of finite length $l(A)$, then by the Schwarz inequality

8.12.1.
$$T(A) = \int_A \frac{1}{\| \operatorname{grad} h \|} \, ds \geq \frac{l(A)^2}{\int_A \| \operatorname{grad} h \| \, ds} \geq \frac{l(A)^2}{h(x)}$$

where $T(A)$ is the total time the flow line spends in A. Note in particular that this implies $\phi_t(x)$ is defined for all positive time t (although ϕ_t may not be surjective). The flow lines of ϕ_t are moving very slowly for most of their length. If h is harmonic, then the flow ϕ_t preserves volume: this means that if it is not altogether stagnant, it must flow along a channel that grows very wide. A river, with elevation h, is a good image. It is scaled so $\operatorname{grad} h$ is small.

Suppose that N is a hyperbolic manifold, and $S \xrightarrow{f} N$ is an uncrumpled surface in N, so that it has area $-2\pi\chi(S)$. Let a be a fixed constant, suppose also that S has no loops of length $\leq a$ which are null-homotopic in N.

PROPOSITION 8.12.1. *There is a constant C depending only on a such that the volume of $\mathcal{N}_1(f(S))$ is not greater than $-C \cdot \chi(S)$. ($\mathcal{N}_1$ denotes the neighborhood of radius 1.)*

PROOF. For each point $x \in S$, let c_x be the "characteristic function" of an immersed hyperbolic ball of radius $1 + a/2$ centered at $f(x)$. In other words, $c_x(y)$ is the number of distinct homotopy classes of paths from x to y of length $\leq 1 + a/2$. Let g be defined by integrating c_x over S; in other words, for $y \in N$,

8.75

$$g(y) = \int_S c_x(y) \, dA.$$

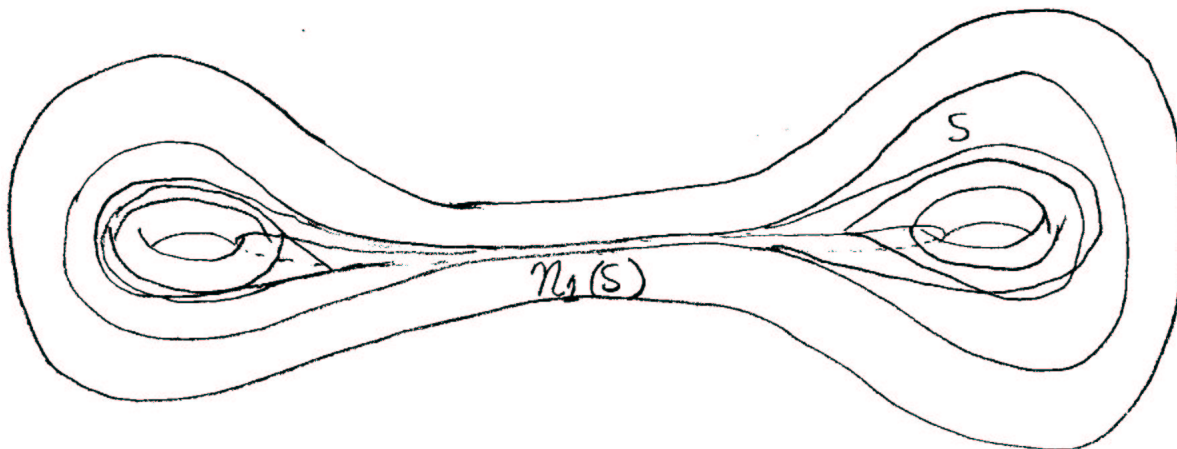If $v(B_r)$ is the volume of a ball of radius r in H^3, then

$$\int_N g \, dV = -2\pi\chi(S) \, v(B_{1+a/2}).$$

For each point $y \in \mathcal{N}_1(f(S))$, there is a point x with $d(fx, y) \leq 1$, so that there is a contribution to $g(y)$ for every point z on S with $d(z, y) \leq a/2$, and for each homotopy class of paths on S between z and x of length $\leq a/2$. Thus $g(y)$ is at least as great as the area $A(B_{a/2})$ of a ball in H^2 of radius $a/2$, so that

$$v(\mathcal{N}_1(f(S))) \leq \frac{1}{A(B_{a/2})} \int_N g \, dV \leq -C \cdot \chi(S).$$

$\square$

As $a \to 0$, the best constant C goes to ∞, since one can construct uncrumpled surfaces with long thin waists, whose neighborhoods have very large volume.

8.76

THEOREM 8.12.3. *If N is geometrically tame, then for every non-constant positive harmonic function h on its convex hull M,*

$$\inf_M h = \inf_{\partial M} h.$$

This inequality still holds if h is only a positive superharmonic function, i.e., if $\Delta h = \operatorname{div} \operatorname{grad} h \leq 0$.

COROLLARY 8.12.4. *If $\Gamma = \pi_1 N$, where N is geometrically tame, then L_Γ has measure 0 or 1. In the latter case, Γ acts ergodically on S^2.*
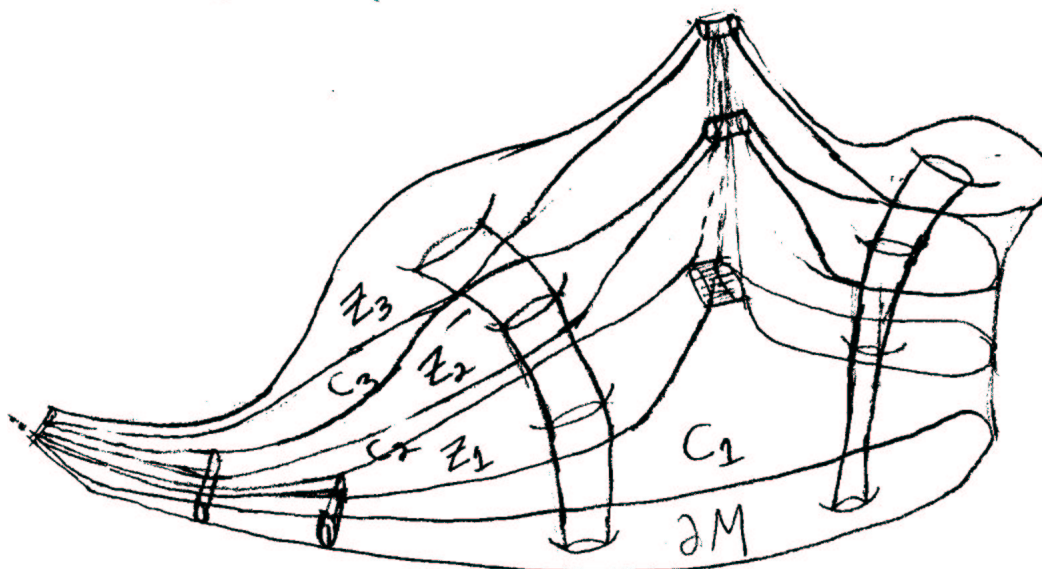
PROOF OF COROLLARY FROM THEOREM. This is similar to 8.4.2. Consider any invariant measurable set $A \subset L_\Gamma$, and let h be the harmonic extension of the characteristic function of A. Since A is invariant, h defines a harmonic function, also h, on N. If $L_\Gamma = S^2$, then by 8.12.3 h is constant, so A has measure 0 to 1. If $L_\Gamma \neq S^2$ then the infimum of $(1 - h)$ is the infimum on ∂M, so it is $\geq \frac{1}{2}$. This implies A has measure 0. This completes the proof of 8.12.4. ☐

Theorem 8.12.3 also implies that when $L_\Gamma = S^2$, the geodesic flow for N is ergodic. We shall give this proof in §, since the ergodicity of the geodesic flow is useful for the proof of Mostow's theorem and generalizations.

PROOF OF 8.12.3. The idea is that all the uncrumpled surfaces in M are narrows, which allow a high flow rate only at high velocities. In view of 8.12.1, most of the water is forced off M—in other words, ∂M is low. 8.77

Let P be the union of horoball neighborhoods of the cusps of N, and $\{S_i\}$ incompressible surfaces cutting $N - P$ into a compact piece N_0 and ends $\{E_i\}$. Observe that each component of P has two boundary components of $\cup S_i$. In each end E_i which does not have a compact intersection with M, there is a sequence of uncrumpled maps $f_{i,j} : S_i \to E_i \cup P$ moving out of all compact sets in $E_i \cup P$, by 8.8.5. Combine these maps into one sequence of maps $f_j : \cup S_i \to M$. Note that f_j maps $\sum[S_i]$ to a cycle which bounds a (unique) chain C_j of finite volume, and that the supports of the C_j's eventually exhaust M.

If there are no cusps, then there is a subsequence of the f_i whose images are disjoint, separated by distances of at least 2. If there are cusps, modify the cycles $f_j(\sum[S_i])$ by cutting them along horospherical cylinders in the cusps, and replacing the cusps of surfaces by cycles on these horospherical cylinders.

8.78

If the horospherical cylinders are sufficiently close to ∞, the resulting cycle Z_j will have area close to that of $f_j \sum [S_i]$, less than, say, $2\pi \sum |\chi(S_i)| + 1$. Z_j bounds a chain C_j with compact support. We may assume that the support of Z_{j+1} does not intersect $\mathcal{N}_2$ (support C_j). From 8.3.2, it follows that there is a constant K such that for all j,

$$v(\mathcal{N}_1(\text{support } Z_j)) \leq K.$$

If $x \in M$ is any regular point for h, then a small enough ball B about x is disjoint from $\phi_1(B)$. To prove the theorem, it suffices to show that almost every flow line through B eventually leaves M. Note that all the images $\{\phi_i(B)\}_{i \in N}$ are disjoint. Since ϕ_t does not decrease volume, almost all flow lines through B eventually leave the supports of all the C_j. If such a flow line does not cross ∂M, it must cross Z_j, hence it intersects $\mathcal{N}_1(\text{support } Z_j)$ with length at least two. By 8.12.1, the total length of time such a flow line spends in

$$\bigcup_{j=1}^{J} \mathcal{N}_1(\text{support } Z_j)$$

grows as J^2. Since the volume of

$$\bigcup_{j=1}^{J} \mathcal{N}_1(\text{support } Z_j)$$

grows only as $K \cdot J$, no set of positive measure of flow lines through B will fit—most have to run off the edge of M. $\square$

REMARK. The fact that the area of Z_j is constant is stronger than necessary to obtain the conclusion of 8.3.3. It would suffice for the sum of reciprocals of the areas to form a divergent series. Thus, $\mathbb{R}^2$ has no non-constant positive superharmonic function, although $\mathbb{R}^3$ has.