Notes on notes of Thurston

R. D. Canary

D. B. A. Epstein

P. Green

Introduction

This paper is based on our study of Bill Thurston's notes [Thurston, 1979], which consist of mimeographed notes produced by Princeton University Mathematics Department as a result of the the course given by Thurston in 1978/79. We shall refer to these notes as [T]. Thurston plans to expand parts of his notes into a book [Thurston, 1979]. There is very little overlap between the projected book and this paper. The basis of this paper was the joint M.Sc. dissertation written by two of us and supervised by the third. Thanks are due to Thurston who gave us help and encouragement and also to Francis Bonahon for additional help.

A useful reference for background information on hyperbolic geometry is [Epstein, 1983] or [Beardon, 1983].

Our work should be regarded as exposition of results of Thurston. There is not much genuinely original material. Nevertheless the effort of production has been considerable and we hope that readers will find our paper helpful. One way to use this paper would be to read it at the same time as reading Thurston's notes. Certainly Thurston's notes cover ground we do not cover, even in those areas to which we pay particular attention. There is some overlap between our work and that contained in [Lok, 1981]. Two good expositions of related work of Thurston are [] and [Scott, 1983].

Chapter 1. (G, X)-structures

1.1. (G, X)-structures on a manifold

The material of this section is discussed in Chapter 3 of [T].

Let X be a real analytic manifold and let G be a Lie group acting on X faithfully and analytically. Let N be a compact manifold, possibly with boundary, having the same dimension as X.

1.1.1 Definition. A (G, X)-*atlas* for N is a collection of charts $\{\phi_\lambda : U_\lambda \to X\}_{\lambda \in \Lambda}$ satisfying the following conditions.

1) The $\{U_\lambda\}$ form an open covering of N.
2) Each ϕ_λ is a homeomorphism onto its image. The image of the boundary, $\phi_\lambda(U_\lambda \cap \partial N)$, is locally flat in X.
3) For each $x \in U_\lambda \cap U_\mu$, there is a neighbourhood $N(x)$ of x in $U_\lambda \cap U_\mu$ and an element $g \in G$, such that

$$\phi_\lambda | N(x) = g \circ \phi_\mu | N(x).$$

We call g a *transition function*.

The last condition gives us a locally constant map

$$g_{\lambda\mu} = g : U_\lambda \cap U_\mu \to G.$$

Notice that $g_{\lambda\mu}$ is determined by $x \in U_\lambda \cap U_\mu$, λ and μ. To see this, note that $N(x)$ has a non-empty intersection with the interior of N. This means that g is equal to $\phi_\lambda \phi_\mu^{-1}$ on some open subset of X. Since the action of G is faithful, g is determined. Notice also that it does not work to insist that $g_{\lambda\mu}$ should be constant on all of $U_\lambda \cap U_\mu$. An example is given by taking $S^1 = \mathbf{R}/\mathbf{Z}$, with the standard $(\mathbf{R}, \mathbf{R})$-structure, where $\mathbf{R}$ acts on $\mathbf{R}$ by addition. If we cover S^1 by two open intervals U_1 and U_2, then $U_1 \cap U_2$ is the disjoint union of two open intervals, and g_{12} is not a constant element of $\mathbf{R}$.

Any (G, X)-atlas determines a unique maximal (G, X)-atlas.

1.1.2 Definition. We define a (G, X)-structure on N to be a maximal (G, X)-atlas.

We usually think in terms of atlases which are not maximal. The above definition refers to a C^0-structure, though the differentiability class only really depends on what happens at the boundary. If N is a C^r-manifold $(r \geq 1)$, we can insist that each ϕ_λ is a C^r-embedding. In that case it is automatic that the boundary is locally flat.

Given any open covering of N by coordinate charts $\{U_\lambda\}$, one may choose a refinement $\{V_j\}$, such that the same element $g \in G$ can be used as a transition function throughout the intersection of any two coordinate charts (see [Godement, 1958, p. 158]). We shall assume from now on that each $g_{\lambda\mu}$ is a transition function which works throughout $U_\lambda \cap U_\mu$. (As we have already pointed out, this is not possible if we insist on a maximal atlas.)

When X is a complete Riemannian manifold, G is the group of isometries of X, and N is a manifold without boundary, we say that N has a Riemannian (G, X)-structure. Under these circumstances, N has an induced Riemannian metric. In this paper, we shall mainly be interested in the case where X is $\mathbf{H}^n$ and G is the group of all isometries of X.

1.2. Developing Map and Holonomy

The material of this section is discussed in Section 3.5 of [T].

Let M be a manifold with a (G, X)-structure and $\gamma : I \to M$ a path (here I is the unit interval $[0,1]$). Holonomy results from the attempt to define a single chart in a coherent way, over the whole of γ and is a measure of the failure of that attempt. We shall give a short sketch of the construction. For more details we refer the reader to [Epstein, 1984].

It works like this: cover γ with a finite number of charts $\{Ui\}_{i=1}^{k}$ where $Ui \cap U_{i+1} \neq \varnothing$ and each transition function is constant as opposed to locally constant. Consider $U_1 \cap U_2$. There exists some $g_1 \in G$ such that $\phi_1 | U_1 \cap U_2 = g_1 \circ \phi_2 | \ U_1 \cap U_2$. If we replace ϕ_2 with $g_1 \circ \phi_2$ we still have a chart, which "extends" ϕ_1 to U_2. Similarly, associated with $U_2 \cap U_3$, is some $g_2 \in G$. Replacing ϕ_3 by $g_2 \circ \phi_3$ extends ϕ_2. So, replacing ϕ_3 by $g_1 \circ g_2 \circ \phi_3$ extends ϕ_1. Continuing along γ we arrive at $g_1 \circ g_2 \cdots \circ g_{k-1} \circ \phi_k$ which replaces ϕ_k.

1.2.1 Definition. We now define the *holonomy* of γ, denoted by $H(\gamma)$, to be $g_1 \circ g_2 \circ \cdots \circ g_{k-1}$. It can be shown that the holonomy depends only upon the homotopy class of γ, keeping the endpoints fixed, and upon the germs of the endpoints of γ. If we choose some basepoint x_0 and a germ about this point, then, by just considering closed loops from x_0, the holonomy gives a map $H: \pi_1(M, x_0) \to G$ which is a homomorphism. Changing the germ conjugates the image by some $g \in G$. In order that H be a homomorphism, rather than an anti-homomorphism, we need to take the correct definition of multiplication in $\pi_1(M, x_0)$. Here $g_1 g_2$ means traversing first g_1 and then g_2.

Now we are in a position to define the developing map. This map can be thought of as the result of analytically continuing the germ of some chart at the basepoint along all possible paths in M. Let $\tilde{M}$ be the universal cover of M. Fix a germ of a chart at x_0, say ϕ_0; then the *developing map* $D: \tilde{M} \to X$ is defined as follows. Take some point $[\omega]$ in $\tilde{M}$ i.e. $[\omega]$ is a homotopy class of paths represented by $\omega: I \to M$ with $\omega(0) = x_0$; let ϕ_1 be a chart at $\omega(1)$. Then we define $D([\omega]) = H(\omega) \circ \phi_1(\omega(1))$. D is independent of ϕ_1 because ϕ_1 is used in the definition of $H(\omega)$. D does depend on the germ of ϕ_0, but changing ϕ_0 merely composes D in X with some $g \in G$.

1.2.2. Theorem: Equivalent definitions of completeness. *Let M be a manifold with a Riemannian (G, X)-structure. Then if M has no boundary the following statements are equivalent:*

1) *The developing map $D: \tilde{M} \to X$ is the universal cover of X.*
2) *M is metrically complete.*
3) *For each $r \in \mathbf{R}^+$, and $m \in M$ the closed ball $B(m, r)$ of radius r about m is compact.*

4) *There exists a family of non-empty compact subsets $\{S_t\}_{t \in R^+}$ with $N_a(S_t) \subset S_{t+a}$ (where $N_a(S_t) = \{x \in M : d(x, S_t) < a\}$).*

For M with boundary, (2),(3) and (4) are all equivalent.

This result is proved in Section 3.6 of [T].

If M is a manifold with boundary, we cannot expect $D : \tilde{M} \to X$ to be a covering map onto its image. Take, for example, the structure on the closed disk D^2 determined by an immersion into the plane as (see Figure 1.3.1). Since D^2 is simply-connected the developing map is simply this immersion and is not a covering of its image.

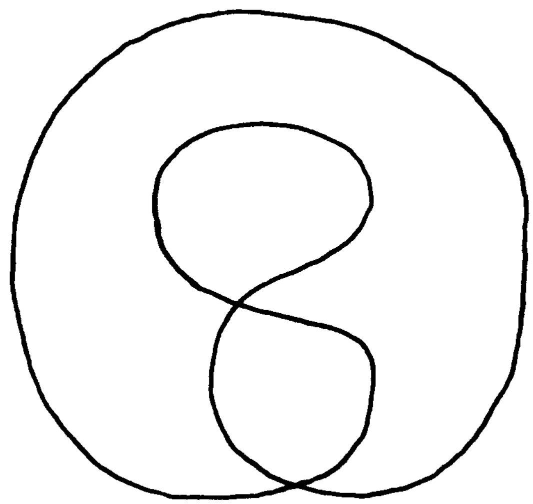

1.2.3 Figure.

1.3. Convexity

See Section 8.3 of [T].

1.3.1 Definition. Let M be a Riemannian manifold possibly with boundary. Then a *geodesic* in M is a path satisfying the geodesic differential equation.

For example let M be the annulus, as shown in Figure 1.3.4 , with a metric structure induced from its embedding in $\mathbf{R}^2$. Then γ as drawn is not a geodesic, even though it minimizes the arc length between its endpoints.

1.3.2 Definition. M is said to be *convex*, if, given any two points of M, each homotopy class of paths between them contains a geodesic arc. We say M is ***strictly convex*** if the interior of this arc is contained in the interior of M. A ***rectifiable*** path in M is a path "whose length makes sense".

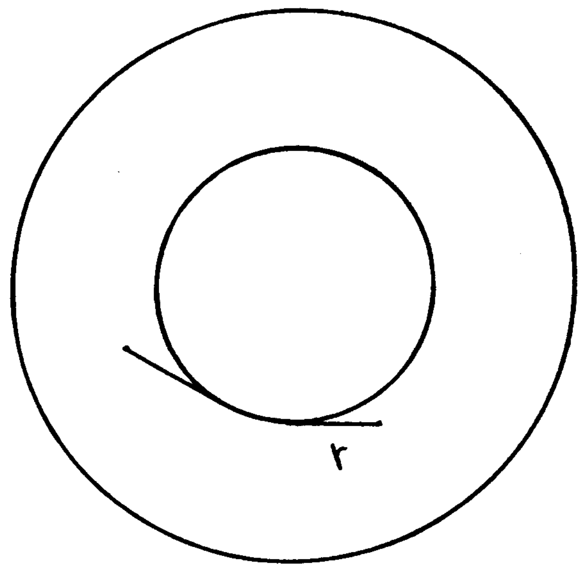

1.3.3 Figure.

That is, let $\rho: I \to M$ be a path. For a subdivision $\tau = \{0 = t_0 < t_1 < \ldots < t_n = 1\}$ (where n is an arbitrary integer), define $l_\tau(\rho)$ by

$$l_\tau(\rho) = \sum_{i=0}^{n-1} d(\rho(t_i), \rho(t_{i+1})).$$

If $l(\rho) = \sup_\tau\{l_\tau(\rho)\}$ (over all subdivisions τ) exists (i.e. if this set of real numbers is bounded), then we say ρ is rectifiable and that its length is l.

1.3.4 Definition. A map $f: M_1 \to M_2$ between Riemannian manifolds is an *isometric map* if it takes rectifiable paths to rectifiable paths of the same length. A *path space* is a path connected metric space in which the metric is determined by path length.

In a path space, there is not necessarily a path between a and b whose length is $d(a,b)$. For example the punctured disk $D^2 \setminus \{0\}$ is a path space, in which there exist such a and b. The annulus (see Figure 1.3.4), with the subspace metric induced from $\mathbf{R}^2$, is not a path space.

It is usual to consider only rectifiable paths parametrized proportionally to arc length. We leave to the reader the simple task of seeing how any rectifiable path γ can be converted, in a canonical way, into a path parametrized proportionally to arc length, having the same length as γ.

We shall develop two useful criteria for convexity. (see Section 8.3 of [T]). First we prove a generalization of a theorem which is well-known for Riemannian manifolds without boundary. See [Gromov, 1981b].

1.3.5. Theorem: Hopf Rinow. *Let (X,d) be a complete, locally compact path space. Then*

1) *$B(x,L)=\{y:d(x,y)\leq L\}$, the ball of radius L about x, is compact for all $x\in X$ and any $L>0$;*

2) *Given any two points, there is a distance minimizing path between them.*

Proof. (1) Suppose there exists $x\in X$ and $L>0$ such that $B(x,L)$ is non-compact. Then we claim that $B(x_1,L/2)$ is non-compact for some x_1 in $B(x,L)$. To prove this claim we argue by contradiction. So suppose not. Let M be the least upper bound of those $r\leq L$ for which $B(x,r)$ is compact. Then $M\geq L/2$ and $B(x,M-L/16)$ is compact. We may cover $B(x,M-L/16)$ by a finite number of balls of radius $L/8$ with centres $y_1, \ldots, y_n$ for some integer n. Enlarge each ball to an $L/2$-closed ball (each of which is compact by assumption). Then $\bigcup_{i=1}^{i=n}B(y_i,L/2)$ is compact.

Let y be an arbitrary point in $B(x,M+L/16)$. We choose a path from x to y of length $\leq M+L/8$. Let z be the point which is a distance $M-L/16$ along this path from x. Since $z\in B(x,M-L/16)$, we have $z\in B(y_j,L/8)$ for some j. Since $d(z,y)\leq L/8$, we have $y\in B(y_j,L/2)$. This argument shows that $B(x,M+L/16)\subset\bigcup_{i=1}^{i=n}B(y_i, L/2)$ is compact. Since M is maximal, this implies that $M=L$ and that $B(x,L)$ is compact. The resulting contradiction establishes our claim that $B(x_1,L/2)$ is non-compact for some $x_1\in B(x,L)$.

Continue this process inductively to obtain a sequence $\{x_n\}$ such that $B(x_n, L/2^n)$ is non-compact and $x_n\in B(x_{n-1}, L/2^{n-1})$. Since this sequence is Cauchy, x_n converges to some $v\in X$. Since X is locally compact, there exists $\delta>0$ such that $B(v,\delta)$ is compact. Choose N such that $\sum_{n=N}^{n=\infty}L/2^n<\delta/2$. Then $B(x,L/2^N)$ would be a non-compact closed subset of $B(v,\delta)$, which is a contradiction. This completes the proof that $B(x,L)$ is compact for each $x\in X$ and each $L>0$.

To prove 2), note that we may choose a sequence $\{\omega_i\}:I\to X$ such that $\omega_i(0)=x$ and $\omega_i(1)=y$, and

1) ω_i is rectifiable for all i;

2) $L(\omega_i)$ converges to $d(x,y)$ as i tends to infinity;

3) ω_i is parametrized proportional to arc length for all i.

Let $F=\{\omega_i\}$ be considered as a subset of $C(I,X)$ equipped with the compact-open topology. There exists L such that $L(\omega_i)\leq L$ for all i. F is equicontinuous since $|s-t|<\delta$ implies that $d(\omega_i(s),\omega_i(t))<\delta L$. Since

$B(x,L)$ is compact, we can apply Ascoli's theorem and deduce that ω_i converges (after taking a subsequence) to $\omega: I \to X$ in the compact-open topology on $C(I,X)$.

We claim that ω is a distance minimizing path. To prove this, first note that $\omega(0) = x$ and $\omega(1) = y$. For any partition $0 = t_0 < \ldots < t_k = 1$, we have

$$\sum_{i=0}^{i=k} d(\omega(t_i),\omega(t_{i+1})) = \lim_{j\to\infty} \sum_{i=0}^{i=k} d(\omega_j(t_i),\omega_j(t_{i+1}))$$
$$\leq \lim_{j\to\infty} L(\omega_j) = d(x,y).$$

This shows that ω is rectifiable and that $L(\omega) \leq d(x,y)$. Hence $L(\omega) = d(x,y)$. This completes the proof of the claim.

□

1.3.6 Definition. A metrically complete Riemannian manifold with boundary is *locally convex* if every point has a convex neighbourhood.

Remark: Whitehead's theorem (see [Kobayashi-Nomizu, 1963]), tells us that any point in the interior of M has a convex neighbourhood, so local convexity is really only a boundary condition. Furthermore, even for points x on the boundary, if y is near x and there is a geodesic from x to y, then this is a distance minimizing path.

1.3.7. Corollary: Local convexity implies convexity. *If M is a complete, locally convex, Riemannian manifold with boundary, then M is convex. In particular, any complete Riemannian manifold without boundary is convex.*

Proof. We may assume that M is simply connected, since we may always work in the universal cover.

By Theorem *(Hopf Rinow)*, there is a distance minimizing path between any two points in M, which we may parametrize proportional to arc length; denote this path by $\omega: I \to M$. We claim that ω is a geodesic. If not, there exists $t \in [0,1]$ such that $\omega[t-\eta, t-\eta]$ is not a geodesic for any $\eta > 0$. Choose η such that $\omega[t-\eta, t+\eta]$ is contained in a convex neighbourhood of $\omega(t)$. Now replace $\omega[t-\eta, t+\eta]$ by the geodesic from $\omega(t-\eta)$ to $\omega(t+\eta)$ to obtain $\overline{\omega}: I \to M$ with $L(\overline{\omega}) < L(\omega) = d(x,y)$, which is a contradiction.

□

Remarks: One may similarly prove that local strict convexity (i.e. every point on the boundary has a strictly convex neighbourhood), implies strict convexity.

1.4. The Developing Map and Convexity

(See Section 8.3 of [T].)

In this section we shall require M to be a (G, X)-manifold possibly with boundary, where G acts on X, a simply connected Riemannian manifold, as a group of isometries. We shall assume that if $\pi: M \to M$ is a covering map, then π is a local isometry and the covering translations are isometries. The following lemma is an immediate consequence of the covering homotopy property.

1.4.1. Lemma: Universal cover convex. *M is convex if and only if $\tilde{M}$ is.*

The next result is a natural generalization of Theorem 1.2.2 *(Equivalent definitions of completeness)* and Proposition 8.3.2 in [T]

1.4.2. Proposition: Coverings and convexity. *Suppose M is a (G, X)-manifold possibly with boundary, where X is a simply connected Riemannian manifold of non-positive curvature and G is the group of isometries of X. Then M is convex and metrically complete if and only if the developing map $D : M \to X$ is a homeomorphism onto a convex complete submanifold of X. In this case D is an isometry onto DM.*

Proof. Suppose M is convex and metrically complete. We have seen that M is convex and it is clear that M is metrically complete. Since the curvature is non-positive, no geodesic in X intersects itself. Since D takes geodesics in M to geodesics in X, D is injective on any geodesic of M. But M is convex, so any two points of M can be joined by a geodesic. Hence D is injective. Thus D is an isometric homeomorphism onto its image in X. Since M is convex, so is DM.

To prove the converse, suppose D is a homeomorphism of $\tilde{M}$ onto a convex complete submanifold of X. By Lemma 1.4.1 *(Universal cover convex)*, M is convex. To show M is metrically complete, let $\{x_i\}$ be a Cauchy sequence in M. Taking a subsequence, we may assume that $d(x_i, x_{i+1}) < 2^{-i-1}$. Let $\omega : [0,1] \to M$ be a path such that $\omega(2^{-i}) = x_i$. Let $\omega : (0,1] \to M$ be a lift of ω. Since M is complete, $\tilde{\omega}(t)$ converges to a limit as t tends to zero. Hence $\omega(t)$ has a limit as t tends to zero. This completes the proof of the proposition.

□

Remark: One may similarly prove the analogous result for strictly convex manifolds with boundary.

1.5. The Deformation Space

See Chapter 5 of [T].

1.5.1 Definition. We wish to consider the space of all possible (G, X)-structures on a fixed manifold N possibly with boundary. This is called the ***deformation space*** of N and is denoted $\Omega(N)$.

Suppose we have a fixed $M_0 \in \Omega(N)$, and a fixed covering by charts $\{\phi_i\}: U_i \to X$ (locally finite with one element of G acting as transition function for the whole of the intersection of any two of them) and a shrinking $\{U_i'\}$. A sub-basis for the topology of $\Omega(N)$ is given by sets of the following form:

$$N_j(M_0) = \{M \in \Omega(N) | M \text{ is defined by } \{\psi_i\}: U_i' \to X \text{ and } \psi_j \in V_j\}$$

where V_j is an open neighbourhood of ϕ_j in the compact-open topology on $C(U_j', X)$. If N is a C^r-manifold, we can restrict to C^r-charts and take V_j to be a neighbourhood of ϕ_j in the compact C^r-topology. We have chosen a shrinking $\{U_i'\}$ so that we may use $\{U_i'\}$ as coordinate charts for all "nearby" structures, i.e. we may deform ϕ_i a small amount in any direction without causing any self-intersections (see Figure 1.5.2). $\Omega(N)$ is infinite-dimensional if it is non-empty.

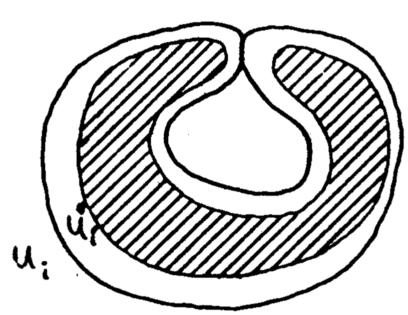

1.5.2 Figure.

Suppose $\{M_i\}$ converges to M in $\Omega(N)$. Then it is clear that we may choose associated developing maps $\{D_i\}$ and D such that D_i converges to D in the compact-C^r topology on $C(N,X)$ $(0 \leq r \leq \infty)$. Intuitively, any compact set in $\tilde{N}$ is covered by a finite number of lifts of coordinate charts, and one may control the behaviour on each coordinate chart.

1.5.3. Theorem: Limit is an embedding. *Suppose M_i converges to M in $\Omega(N)$. Let D_i and D be the associated developing maps such that D_i converges to D, and suppose that $D|K$ is an embedding, for some compact subset K of $\tilde{N}$. Then $D_i|K$ is also an embedding for sufficiently large i.*

Proof. We need only prove this in the C^0-case. Suppose that the result is false. Then there exists $\{x_i\}$, $\{y_i\} \subset K$ such that $D_i(x_i) = D_i(y_i)$ and $x_i \neq y_i$. Since K is compact we may assume that x_i converges to some x, and y_i converges to some y, both in K. Since D_i converges to D, $D(x) = D(y)$ and so $x = y$. Let $\pi : \tilde{N} \to N$ be the universal cover of N. It follows that, for sufficiently large i, $\pi(x_i)$, $\pi(y_i)$, $\pi(x)$, and $\pi(y)$, all lie in U_j' for some fixed j. We may assume that U_j' is contractible in N, so that there exists $\alpha : U_j' \to \tilde{N}$, a well-defined (G, X)-homeomorphism onto its image with $\pi \circ \alpha : U_j' \to N$ equal to the inclusion map and $\alpha(\pi(x_i)) = x_i$, and $\alpha(\pi(y_i)) = y_i$ for large i. Since ψ_i and $D_i\alpha$ are (G, X)-embeddings of U_j' in X, $D_i \circ \alpha | U_j' = g \circ \psi_i | u_j'$ for some $g \in G$. Then since $g \circ \psi_i$ is a homeomorphism onto its image and $\pi x_i \neq \pi y_i$, we have $D_i x_i \neq D_i y_i$, which is a contradiction.

□

1.5.4 Space of developing maps. It is often easier to regard $\Omega(N)$ as a function space. We can characterize a developing map $D : \tilde{N} \to X$ as a local C^r-diffeomorphism (homeomorphism if $r = 0$) such that for covering translation γ of $\tilde{N}$, there is an element $H(\gamma)$ of G with $D \circ \gamma = H(\gamma) \circ D$. We topologize the set of developing maps by means of a subbasis consisting of sets of the following form

1) U where U is open in the compact-C^r topology on $C(\tilde{N}, X)$,
2) $N(K) = \{D \mid D|K$ is an embedding$\}$ where K is a compact subset of $\tilde{N}$.

Then G acts continuously on the set of developing maps by composition on the left. We can identify $\Omega(N)$ with the quotient of the space of developing maps by G.

1.5.5 Remarks.

1) If we restrict our attention to C^r manifolds for some fixed r ($1 \leq r \leq \infty$), this topology on the space of developing maps is simply the compact-C^r topology (i.e. sub-basis sets of the second type are unnecessary.) It is a standard fact in differential topology that the

space of C^r embeddings is open in the space of all C^r maps.

2) Our topology is strictly finer than the compact-open topology — see for example Figure 1.5.6 .

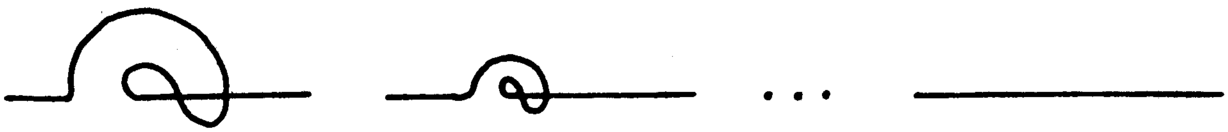

1.5.6 Figure. *The picture shows the images of various developing maps which are immersions not embeddings. This sequence converges in the compact-open topology, but not in ours.*

3) We shall use the notation $d_M(x,y)$ to denote the distance between x and y in N as measured in the path metric induced by M. If M_i converges to M, d_{M_i} does not necessarily converge to d_M. In fact, d_M may be infinite on a compact hyperbolic manifold with a boundary which is not smooth. In other words, it is possible to have two points on the boundary of M, with no rectifiable path joining them — see Figure 1.5.7 . To be sure that d_{M_i} converges to d_M, one needs to be working with the compact-C^r topology ($r \geq 1$)

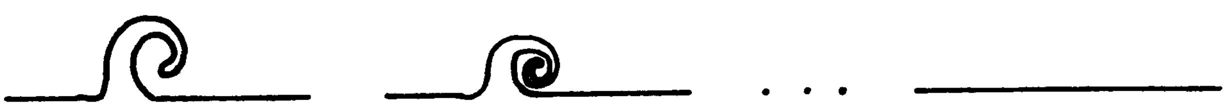

1.5.7 Figure. *d_{M_n} does not converge to d_M in this example, as we may choose the spiral to be arbitrarily long, and distances are measured by paths. In fact d_M is infinite for some pairs of points.*

4) If we consider a closed surface of negative Euler characteristic, Teichmüller space (see Section 3.1 *(The Geometric Topology)*) is a quotient of $\Omega(N)$. More precisely, $T(N)=\Omega(N)/H$, where H is the group of isotopies of N to the identity. An isotopy h acts on a developing map D by lifting h to an isotopy $\tilde{h}$, which starts at the identity, of $\tilde{N}=\mathbf{H}^2$ and taking $D\circ\tilde{h}_1$. For a general (G, X)-structure, it is also quite usual in the literature to define the space of structures as $\Omega(N)$ modulo isotopy.

1.6. Thickenings

The subject we are about to discuss is capable of considerable generalization. A very general treatment is given in [Haefliger, 1958]. Our discussion applies to a number of other situations (almost) verbatim. For example one can obtain a proof of the Whitney-Bruhat Theorem [Whitney-Bruhat, 1959], which states that every real analytic manifold can be thickened to a complex analytic manifold.

Let N be a manifold of dimension n, possibly with boundary. Let M be a fixed (G, X)-structure on N.

1.6.1 Definition. A *thickening* of M is a (G, X)-structure on a manifold without boundary N_T, of dimension n, containing N as a submanifold, which induces the given (G, X)-structure on N.

1.6.2 Definition. If M_1 and M_2 are Riemannian (G, X)-manifolds and M_1 is (G, X)-embedded in M_2, we say M_2 is an ϵ-*thickening* of M_1 if, for each point $x \in M_1$, there is an ϵ-neighbourhood in M_2 which is isometric to an ϵ-ball in X.

1.6.3. Theorem: Thickenings exist. *M has a thickening, and the germ of the thickening is unique in the following sense. Let $M \subset M_1^*$ and $M \subset M_2^*$ be two thickenings. Then we can find U_1^* and U_2^* with $M \subset U_i^* \subset M_i^*$, such that each U_i^* is a thickening of M and there is a (G, X)-isomorphism between U_1^* and U_2^* which extends the identity on M. Moreover this isomorphism is uniquely determined if each component of U_1^* meets M.*

Proof. Thickenings exist:

First we prove uniqueness.

... **1.6.4. Lemma: Thickening unique.** *Given two thickenings M_1^* and M_2^* of M, there exist isomorphic open neighbourhoods U_1^* of M in M_1^* and U_2^* of M in M_2^*. Moreover the isomorphism is unique.*

Proof. Thickening unique: For each point $x \in M$, we have a neighbourhood U, which is open in M_1^*, and a (G, X)-embedding $f: U \to M_2^*$, which is the identity on $U \cap M$. We may suppose that we have a family $\{fi: Ui \to M_2^*\}$ of such embeddings, such that the $\{Ui\}$ form a locally finite family and cover M. Let $\{Vi\}$ be a shrinking of $\{Ui\}$.

Let W be the set of $w \in \cup V_i$ such that if $w \in \overline{V_i} \cap \overline{V_j}$ then $f_i = f_j$ in some neighbourhood of w. From local finiteness, it is easy to deduce that W is open. Analytic continuation shows that $M \subset W$. We get a well-defined (G, X)-immersion of W in M_2^*. By restricting to a smaller neighbourhood of M we obtain an embedding, as required.

The uniqueness of the embedding follows by analytic continuation.

Thickening unique

Continuation, proof of Thickenings exist: We now prove the existence of thickenings. The clearest form of the proof is given in Figure 1.6.5 .

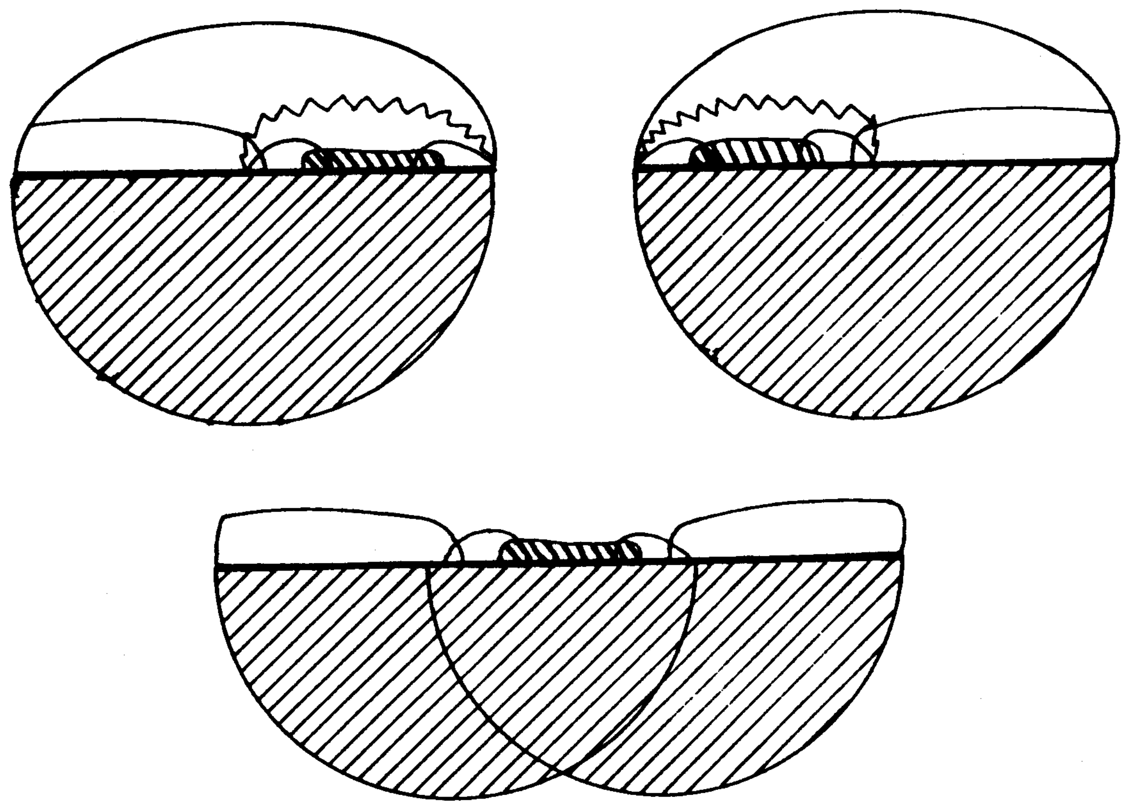

1.6.5 Figure.

Let $\{\phi_i : U_i \to X\}$ be a finite atlas of coordinate charts for N. First note that U_i itself may be thickened. To see this, note that we can identify a single U_i with a subspace of X. U_i is open in X if and only if $\partial U_i (= U_i \cap \partial N) = \emptyset$. By adding small open neighbourhoods in X of each point in ∂U_i, we obtain an open subset V_i of X such that U_i is closed in V_i.

We now find thickenings of $U_1 \cup \ldots \cup U_i$ by induction on i. What we have to show is that if a (G, X)-manifold M (with boundary) is the union of two open submanifolds U_1 and U_2, and if U_1 has a thickening U_1^* and U_2 has a thickening U_2^*, then M also has a thickening. To see this, we write $X_1 = M\setminus U_1$ and $X = M\setminus U_2$. Then X_1 and X_2 are disjoint closed subsets of M. Let $X_1 \subset V_1$ and $X_2 \subset V_2$, where V_1 and V_2 are open in M, $V_1 \cap V_2 = \emptyset$, $V_1 \subset U_1$ and $V_2 \subset U_2$. Let $P_i^* \subset U_i^*$ be an open neighbourhood of X_i, such that $\bar{P}_i^* \cap (M\setminus V_i) = \emptyset$. $M\setminus(V_1 \cup V_2) = P$ is closed subspace and $P \cap \bar{P}_i^* = \emptyset$ for $i = 1,2$. Note that $P \subset U_1 \cap U_2$. Let U^* be a thickening of $U_1 \cap U_2$. By the Lemma 1.6.4 *(Thickening unique)*, we may assume that $U^* \subset U_1^*$ and $U^* \subset U_2^*$. Let P^* be an open neighbourhood of P in U^* such that $\bar{P}^* \cap \bar{P}_i^* = \emptyset$ for $i = 1,2$. $\bar{V}_1 \setminus P^* \cup P_1^*$ and $\bar{V}_2 \setminus P^* \cup P_2^*$ are disjoint closed subsets of M, each contained in $U_1 \cap U_2$. Let W_1^* and W_2^* be disjoint open neighbourhoods of these sets in U^*. We may assume that $\bar{P}_2^* \cap W_1^* = \emptyset$, since

$$\bar{P}_2^* \cap \bar{V}_1 \subset \bar{P}_2^* \cap (M\setminus V_2) = \emptyset,$$

and similarly we may assume that $\bar{P}_1^* \cap W_2^* = \emptyset$.

Now we take the open subspace $P_1^* \cup W_1^* \cup P^* \cup W_2^*$ of U_1^* and the open subspace $P_2^* \cup W_2^* \cup P^* \cup W_1^*$ of U_2^* and glue them together by identifying $W_1^* \cup P^* \cup W_2^*$ in these two sets. We have

$$P_1^* \cap (W_1^* \cup P^* \cup W_2^*) = P_1^* \cap W_1^*$$

and

$$P_2^* \cap (W_1^* \cup P^* \cup W_2^*) = P_2^* \cap W_2^*$$

and these two sets are disjoint. It follows that the identification space is Hausdorff. (Note that P_1^* and P_2^* are disjoint open subsets.) Gluing manifolds together along an open subset automatically gives a manifold, usually a non-Hausdorff manifold.

This completes the proof of the theorem.

Thickenings exist

1.6.6 Remark. The same proof works for a non-compact manifold. Any (G, X)-manifold of dimension n with a countable basis can be covered with a finite number of charts. (These charts are not, of course, connected.) In fact $(n + 1)$ charts will do.

1.7. Varying the Structure

Let N be a compact C^r-manifold $(0 \leq r \leq \infty)$ with boundary. Let N_{Th} be the union of N with a collar $\partial N \times I$, where $\partial N \subset N$ is identified with $\partial N \times 0 \subset \partial N \times I$. We fix a (G, X)-structure M_{Th} on N_{Th}. By fixing a basepoint in N, the universal cover of N, we can identify the group of covering translations of N with $\pi_1(N) = \pi_1(N_{Th})$. The same group of covering translations acts on the universal cover of N_{Th}.

A developing map $D: \tilde{N} \to X$ induces a homomorphism from $\pi_1 N$, the group of covering translations, to G. Let $\mathfrak{D}(N)$ be the space of developing maps with the topology described in 1.5.4 *(Space of developing maps)*. Topologize $\mathrm{Hom}(\pi_1 N, G)$ with the compact-open topology.

The holonomy gives us a map

$$H: \mathfrak{D}(N) \to \mathrm{Hom}(\pi_1 N, G)$$

which is easily seen to be continuous. The induced map

$$\Omega(N) \to R(\pi_1 N, G)$$

of the space of (G, X)-structures on N into the space of conjugacy classes of homomorphisms of $\pi_1 N$ into G, is also continuous. In general $R(\pi_1 N, G)$ is not Hausdorff. Let M be a (G, X)-structure on N and let M_{Th} be an extension of M to N_{Th}. (From Theorem 1.6.3 *(Thickenings exist)* and the Collaring Theorem [Connelly, 1971] we can see that such an extension exists.)

The next theorem may be viewed as one of the ways to make the discussion in Section 5.1 of [T] more formal.

1.7.1. Theorem: Neighbourhood is a product. *Let $D_{Th}: \tilde{N}_{Th} \to X$ be a fixed developing map for M_{Th}. Then a small neighbourhood of $D_{Th}|\tilde{N}$ in the space of developing maps of N is homeomorphic to $\mathcal{A} \times \mathcal{B}$, where $\mathcal{A}$ is a small neighbourhood of the obvious embedding $N \subset N_{Th}$ in the space of locally flat embeddings, and $\mathcal{B}$ is small neighbourhood of the holonomy $h_M: \pi_1 N \to G$ in the space of all homomorphisms of $\pi_1 N$ into G. The projection of the neighbourhood of $D_{Th}|M$ to $\mathcal{B}$ is given by the holonomy $H: \mathfrak{D}(N) \to \mathrm{Hom}(\pi_1, N, G)$ defined above.*

The action of G on the space of developing maps corresponds to conjugation in $\mathcal{B}$.

Proof. Neighbourhood is a product: Let N_{TT} be the union of N_{Th} with a collar (i.e. a thickening of a thickening) and let M_{TT} be a (G, X)-thickening of M_{Th} with underlying manifold N_{TT}. We now show how to map a small neighbourhood $\mathcal{B}$ of $h_M:\pi_1 N \to G$, the holonomy of M, continuously into the space of structures of N_{Th}.

... 1.7.2. Lemma: Holonomy induces structure. *There is a continuous map*

$$D:\mathcal{B} \to \mathfrak{D}(N_{\mathrm{T}}) \ .$$

Proof. Holonomy induces structure: Let $\{Ui\}_{0\leq i\leq k}$ be a finite open covering of N_{TT} such that $U_0 = N_{\mathrm{TT}} \backslash N_{\mathrm{Th}}$ and such that $Ui \subset \mathrm{int} N_{\mathrm{TT}}$ for $1 \leq i \leq k$. Let $U_i^1 = Ui$ $(1 \leq i \leq k)$. For each $r > 0$ let $\{U_i^{r+1}\}_{0\leq i\leq k}$ be a shrinking of $\{U_i^r\}_{0\leq i\leq k}$. We may assume that U_i^r is simply connected for $i > 0$.

Let $D_M:\tilde{N}_{\mathrm{TT}} \to X$ be a developing map for M_{TT} and let $\pi:\tilde{N}_{\mathrm{T}} \to N_{\mathrm{T}}$ be the universal cover.

Choose an h near h_M. We show how to construct a (G, X)-structure on N_{Th}. The method is to construct a developing map $N_{\mathrm{Th}} \to X$. We define $D_1|\pi^{-1}U_1 = D_M|\pi^{-1}U_1$ and, inductively, an equivariant local homeomorphism $D_S : \pi^{-1}U_1^s \cup \ldots \cup \pi^{-1}U_s^s \to X$ which is equal to D_{s-1} on $\pi^{-1}U_1^s \cup \ldots \cup \pi^{-1}U_{s-1}^s$.

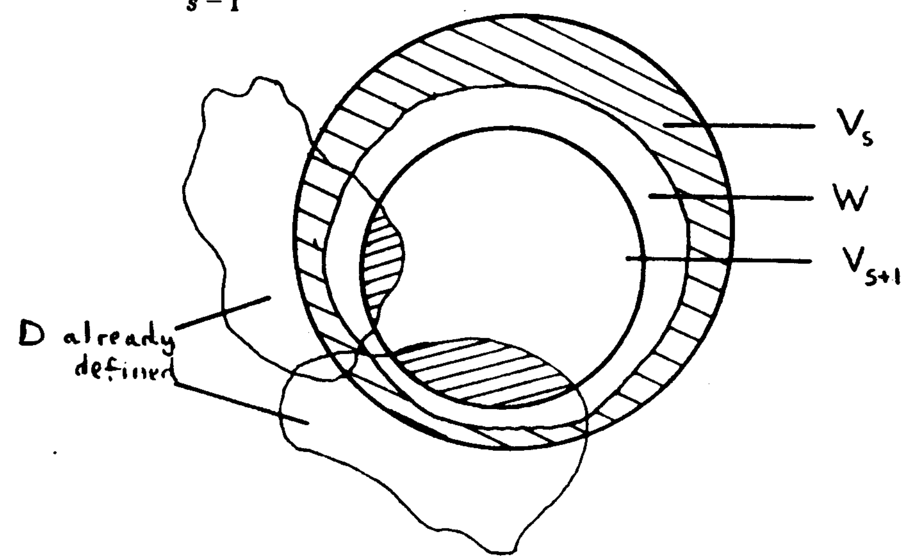

1.7.3 Figure.

To carry out the induction step we need to define D_{s+1} on $\pi^{-1}U_{s+1}^{s+1}$. We need to do this on only one component, because we can extend by h-equivariance. Let π map $V_s \subset \tilde{N}$ homeomorphically onto U_{s+1}^s and let

$V_{s+1} \subset V_s$ correspond to U^{s+1}_{s+1}. Let W be an open neighbourhood of V_{s+1} in N whose closure is contained in V_s. We define $f: V_s \to X$ as follows. On $V_s \setminus W$, f is equal to D_M. On $V_{s+1} \cap (\pi^{-1} U^{s+1}_1 \cup \ldots \cup \pi^{-1} U^{s+1}_s)$ f is equal to D_s. On the remainder of V_s, f is given using standard theorems about C^r-manifolds. If $r = 0$ we use a result of [Edwards-Kirby, 1971], in the form explained in [Siebenmann, 1972]. If $r > 0$ one can use standard bump function techniques and the openness of the space of embeddings in the space of all maps. We define $D_{s+1} | V_{s+1}$ to be equal to $f | V_{s+1}$. Note that whether $r = 0$ or $r > 0$, the extension f depends continuously in the C^r-topology, on D_s and on h. Hence D_{s+1} depends continuously on D_s and upon h. It follows that the developing map depends continuously on h. When $s = k$, the induction is done and we have defined a (G, X)-structure on N_{Th} which depends continuously on h.

Holonomy induces structure

Continuation, proof of Neighbourhood is a product: Given an element $(i: N \to N_{\mathrm{Th}}, h)$ of $\mathcal{A} \times \mathcal{B}$, we define an element of $\mathfrak{D}$ as follows. We obtain a developing map $D(h): \tilde{N}_{\mathrm{Th}} \to X$ from Lemma 1.7.2 *(Holonomy induces structure)*. By lifting i to the universal cover, choosing a lift $\tilde{i}$ which is near the identity, we obtain a composite developing map

$$\tilde{N} \xrightarrow{\tilde{i}} \tilde{N}_{\mathrm{Th}} \xrightarrow{D(h)} X \ .$$

This composition is near $D_M | \tilde{N}$.

Conversely, given a developing map $D: \tilde{N} \to X$, near to $D_M | \tilde{N}$, we have the holonomy $h = h(D)$ which is near to h_M, and hence the developing map $D(h): \tilde{N}_{\mathrm{Th}} \to X$ of 1.7.2 *(Holonomy induces structure)*. To complete the proof of Theorem 1.7.1 *(Neighbourhood is a product)* we need to construct an embedding $i: N \to N_{\mathrm{Th}}$, near the identity, such that $D(h) \circ \tilde{i} = D$. This will follow once we have proved the next lemma. (The proof that the maps between $\mathcal{A} \times \mathcal{B}$ and the neighbourhood of $D_T | N$ are inverse to each other is left to the reader.)

1.7.4. Lemma: Embedding exists. *There is a unique equivariant map $\tilde{i}: \tilde{N} \to \tilde{N}_{\mathrm{Th}}$ such that $D(h) \circ \tilde{i} = D$.*

Proof. Embedding exists: To carry out the construction we take a covering $\{U_i\}$ of N_{Th} by simply connected coordinate charts, and a shrinking $\{U_i'\}$, also consisting of simply connected open sets. We assume that $U_1' \cup \ldots \cup U_r'$ is connected for $1 \le r \le k$.

Let V_i be a lifting to $\tilde{N}_{\mathrm{Th}}$ of U_i and let $V'_i \subset V_i$ be the corresponding lift of U'_i. We choose the lifts so that $V'_1 \cup \ldots \cup V'_r$ is connected for $1 \leq r \leq k$. The map $D(h)|V_i$ is a homeomorphism onto its image. We may assume $D(V'_i \cap N) \subset D(h)(V_i)$ since D and $D(h)$ are both near D_M. The equation $D(h) \circ \tilde{i} = D$ then determines $\tilde{i}$ uniquely on $V'_i \cap N$. We shall show that $\tilde{i}$ is well-defined on N and equivariant at the same time. Let $i_1 : V'_i \cap N \to X$ satisfy $D(h)i_1 = D$ on $V'_i \cap N$ and let $i_2 : V'_j \cap N \to X$ satisfy $D(h)i_2 = D$ on $V'_j \cap N$. Let γ be a covering translation of N. Then, on $V'_i \cap \gamma^{-1} V'_j$ we have

$$\begin{aligned} D(h)i_1 = D &= h(\gamma^{-1})D\gamma \\ &= h(\gamma^{-1})D(h)i_2\gamma \\ &= D(h)\gamma^{-1}i_2\gamma \end{aligned}$$

Therefore $i_1 = \gamma^{-1}i_2\gamma : V'_i \cap \gamma^{-1} V'_j \cap \tilde{N} \to V_i \cap \gamma^{-1} V_j$.

Taking $\gamma = \mathrm{id}$ we see that $\tilde{i} = i_1 = i_2$ is well-defined on $V'_1 \cup \ldots \cup V'_k$. It also follows that we have a well-defined and unique equivariant extension to N.

Embedding exists

This completes the proof of the theorem

Neighbourhood is a product

1.7.5. Corollary: Epsilon thickenings exist. *If N is a manifold with a Riemannian (G, X)-structure then there is a neighbourhood of M in $\Omega(N)$ in which ϵ-thickenings exist for some value of $\epsilon > 0$.* (See Definition 1.6.2 *(Epsilon thickening)* for the definition of an ϵ-thickening.)

1.7.6 Weil's Theorem. In the particular case where N has no boundary, Theorem 1.7.1 *(Neighbourhood is a product)* implies the famous theorem of [Weil, 1960], that, up to isotopy, deformations in holonomy corresponding to deformations of the (G, X)-structure. Our theorem gives a very precise version of that result and is also a generalization to manifolds with boundary.

Chapter 2. Hyperbolic structures

2.1. Möbius Groups

The material in this section is covered in Chapter 8 of [T].

Suppose that M is a complete hyperbolic manifold without boundary, i.e M has an $(\mathbf{H}^n, \mathrm{Isom}(\mathbf{H}^n))$-structure. Then the developing map is a homeomorphism of M with $\mathbf{H}^n$ and we will consider M to be equal to $\mathbf{H}^n$. The covering transformations (which form a group isomorphic to $\pi_1 M$) are a discrete subgroup Γ of $\mathrm{Isom}(\mathbf{H}^n)$ and $M = \mathbf{H}^n/\Gamma$.

2.1.1 Definition. We define a *Möbius group* to be a discrete subgroup of $\mathrm{Isom}(\mathbf{H}^n)$. When the group consists of orientation preserving isometries and $n=3$, it is called a *Kleinian group* and when $n=2$ a *Fuchsian group*. If Γ is torsion-free, $\mathbf{H}^n/\Gamma$ is a complete hyperbolic manifold. From now on, Γ will denote a Möbius group, which will be assumed to be torsion free unless stated otherwise. $L\Gamma$ will denote the limit set of Γ (i.e. $\Gamma(x) - \Gamma(x)$ where $x \in \mathbf{H}^n$ and the closure is taken in the disk $\mathbf{H}^n \cup S^{n-1}_\infty$), and $D\Gamma = S^{n-1}_\infty \setminus L\Gamma$ the ordinary set. (We note that Γ acts properly discontinuously on $\mathbf{H}^n \cup D\Gamma$.)

2.1.2 Definition. The *convex hull* $\mathcal{C}(L\Gamma)$ of a Möbius group Γ is defined to be the convex hull of its limit set, i.e. the intersection of all closed hyperbolic half-spaces of $\mathbf{H}^n \cup S^{n-1}_\infty$ containing $L\Gamma$. We define three manifolds associated with Γ:

$C\Gamma = (\mathcal{C}(L\Gamma) \setminus L\Gamma)/\Gamma$ the *convex core*

$M\Gamma = \mathbf{H}^n/\Gamma$ the *Kleinian manifold*

$O\Gamma = \mathbf{H}^n \cup D\Gamma/\Gamma$ the *Kleinian manifold with boundary.*

The convex core carries all the essential information about Γ. $O\Gamma$ is often a compact manifold, even when $M\Gamma$ is non-compact.

Further information on and references for Möbius groups can be found in [Beardon, 1983] and [Harvey, 1977].

2.2. The Thick-Thin decomposition

See Section 5.10 of [T].

2.2.1 Definition. Let M be a complete hyperbolic n-manifold. We define inj(x), the ***injectivity radius*** at x, by

$$\text{inj}(x) = \inf_{\gamma}\{\text{length}(\gamma)\}/2$$

where γ varies over homotopically non-trivial loops through x. Given $\epsilon>0$ let $M_{[\epsilon,\infty)}=\{x \in M \mid \text{inj}(x)\geq\epsilon\}$ (this is often called the *thick* part of the manifold), and let $M_{(0,\epsilon]}=\{x \in M \mid \text{inj}(x) \leq \epsilon\}$ (known as the *thin* part). Note that given $x \in M_{(\epsilon,\infty)}$, the ϵ-neighbourhood of x is isometric to an ϵ-ball in $\mathbf{H}^n$.

We shall need to know the structure of the thin part of a hyperbolic n-manifold. The following result is due to Margulis [Kazdan-Margulis, 1968] although the following formulation is due to Thurston [Thurston, 1984]

2.2.2. Theorem: Thick-thin decomposition. *There is a universal constant ϵ, (called the Margulis constant), depending only on the dimension n, such that, given any complete hyperbolic n-manifold M, the thin part $M_{(0,\epsilon]}$ consists of a disjoint union of pieces of the following diffeomorphism types:*

1) $N^{n-1}\times[0,\infty)$ *where N^{n-1} is an Euclidean manifold. (These non-compact components of $M_{(0,\epsilon]}$ are neighbourhoods of cusps of M.)*

2) *A neighbourhood of a closed geodesic.*

2.2.3 Remarks:.

1) Neighbourhoods of cusps of M have finite volume if and only if N^{n-1} is compact. Thus M has finite volume if and only if $M_{[\epsilon,\infty)}$ is compact.

2) The neighbourhood of a geodesic can be non-orientable. (For example a Möbius band in two dimensions.)

When $n=2$ or when $n=3$ and M is orientable, we know, even more specifically, that the non-compact components of $M_{(0,\epsilon]}$ are of the form:

$$H_1/\Gamma_p \subset M_{(0,\epsilon]} \subseteq H_2/\Gamma_p$$

where H_1 and H_2 are horoballs centred at p (a parabolic fixed point of Γ) and Γ_p is the stabiliser of p in Γ. We also have the following corollary which will be used implicitly throughout the section on geodesic laminations.

2.2.4. Corollary: Simple geodesics go up a cusp. *If M is a hyperbolic surface or 3-manifold, there exists an ϵ such that if a simple geodesic enters a non-compact component of $M_{(0,\epsilon]}$ it must continue straight up the cusp (i.e. it must have a lift with endpoint p.)*

2.2.5 Remarks:.

1) In the orientable case, each non-compact component of $M_{(0,\epsilon]}$ is of the form H_1/Γ_p, where H_1 is a horoball centred on p.

2) A *uniform horoball* is a horoball whose images under the group Γ are disjoint or equal. If $n \leq 3$, then each cusp gives rise to a uniform horoball. There are examples in higher dimensions ([Apanasov, 1985]) where no uniform horoball exists, but the matter is still not entirely cleared up because all known counter- examples require an infinite number of generators for Γ.

The sort of decomposition we get in the case of a surface is illustrated in Figure 2.2.6 (*Thick-thin decomposition for surfaces).*

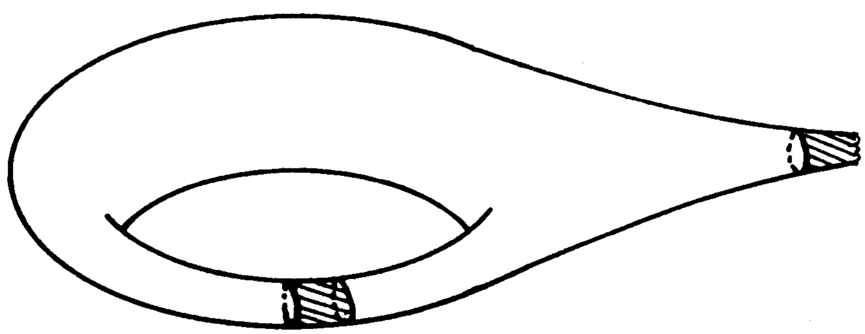

2.2.6 Figure.

2.3. The Nearest Point Retraction

See Section 8.4 of [T].

This section discusses properties of the nearest point retraction from hyperbolic space (with the sphere at infinity included) onto a convex subset. In general we will be considering it as a retraction of hyperbolic space onto the convex hull of the limit set of a Kleinian group. In this form it will induce a retraction from the Kleinian manifold with boundary onto the convex core. See also [Epstein-Marden].

Given a closed convex subset C of $\mathbf{H}^n \cup S^{n-1}_\infty$, there is a canonical retraction $r:\mathbf{H}^n \cup S^{n-1}_\infty \to C$. If $x \in C$ then $r(x)=x$, and if $x \in \mathbf{H}^n \setminus C$ then, since C is closed, there is a ball B_h of radius h about x and disjoint from C, for some h. We increase the radius of this ball continuously until it first touches C. This point of first contact we define to be $r(x)$. If $x \in S^{n-1}_\infty / C$, we do the same construction but with horoballs centred at x. Again we define the point of first contact to be $r(x)$.

2.3.1. Proposition: r continuous. *r is well-defined and continuous.*

For the proof of this result see [Epstein-Marden].

Now let Γ be a Möbius group and let $O\Gamma$ and $C\Gamma$ be the Kleinian manifold with boundary and the convex core respectively (see Section 2.1 *(Moebius groups)*). Then we can use r to define a map $\bar{r}:O\Gamma \to C\Gamma$ as follows. Given $x \in O\Gamma$ we define $\bar{r}(\pi(x))=\pi r(x)$ where $x \in \mathbf{H}^n \cup D\Gamma$. (Here $\pi:\mathbf{H}^n \cup D\Gamma \to O\Gamma$.) Clearly $\bar{r}$ is well-defined.

2.3.2. Proposition: r bar proper. *$\bar{r}$ is proper.*

Proof. Note that $r:\mathbf{H}^n \cup S^{n-1} \to \mathcal{C}(L\Gamma)$ is proper. Therefore the induced map $\mathbf{H}^n \cup D\Gamma \to \mathcal{C}(L\Gamma)\setminus L\Gamma$ is proper. Let K be a compact subset of $C\Gamma$. For each $x \in K$, we choose $\bar{x} \in \mathbf{H}^n \cup D\Gamma$ such that $\pi\bar{x}=x$ and we choose a compact neighbourhood $N(\bar{x})$ of $\bar{x}$ in $\mathcal{C}(L\Gamma)$. Let $x_1, \ldots, x_k$ be chosen so that $\pi\circ N(\bar{x}_1), \ldots, \pi\circ N(\bar{x}_k)$ cover K. Since r is proper, $L_i=r^{-1}N(\bar{x}_i)$ is compact. Clearly $\bar{r}^{-1}K$ is contained in $\pi L_1 \cup \ldots \cup \pi L_K$ and is therefore compact. So $\bar{r}$ is proper.

□

2.4. Neighbourhoods of Convex Hyperbolic Manifolds

See Section 8.3 of [T].

2.4.1. Theorem: Embedding convex manifolds. *Any convex complete hyperbolic manifold M with boundary can be embedded in a unique complete hyperbolic manifold with the same fundamental group as M.*

Proof. Since M is convex and complete, we can apply Theorem 1.4.2 *(Coverings and convexity)* and think of M as a closed convex subset C of $\mathbf{H}^n$ together with a group $\Gamma \subset \mathrm{Isom}\,\mathbf{H}^n$ which acts properly discontinuously on C. Since r is proper, Γ also acts properly discontinuously on $r^{-1}C = \mathbf{H}^n \cup D\Gamma$. Therefore it acts properly discontinuously on $\mathbf{H}^n$.

The only thing that might stop $\mathbf{H}^n/\Gamma$ being a manifold is that Γ could have torsion. For example, take the annulus in $\mathbf{H}^2$ as shown in Figure 2.4.2 . Then the quotient by a rotation of π about the centre gives a manifold. But $\mathbf{H}^2/\Gamma$ is a cone.

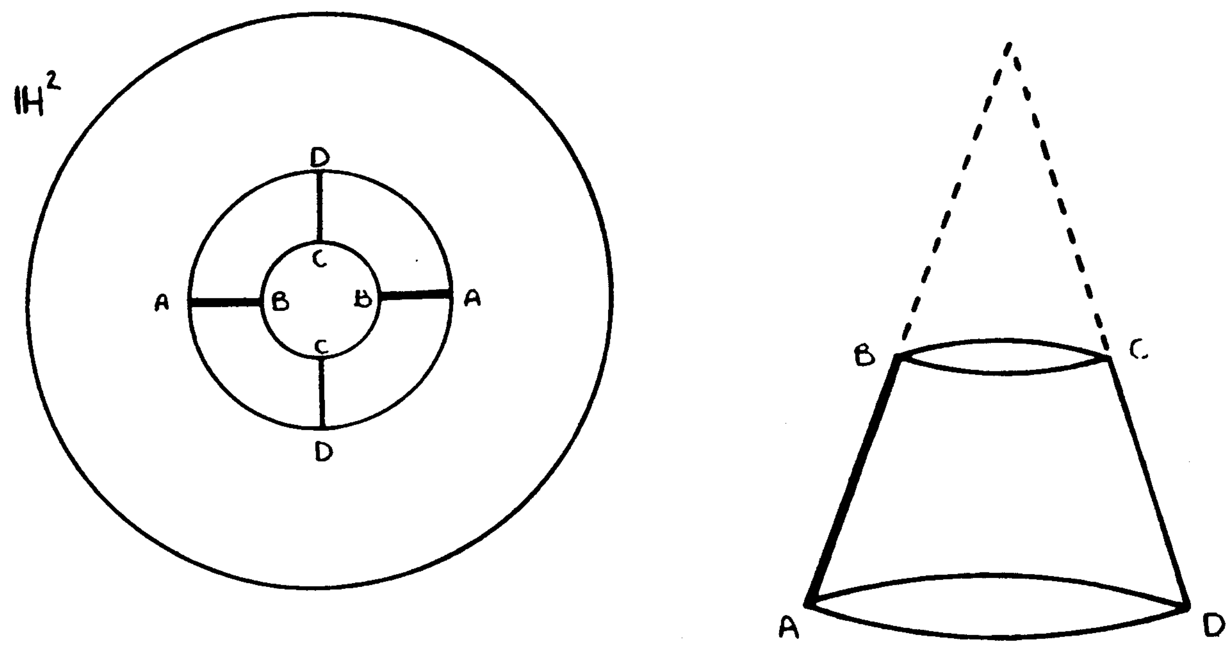

2.4.2 Figure.

We shall prove by contradiction that Γ is torsion free. Suppose Γ has an element g with order n. Take any point $x \in C$ and consider the set $\{x, g(x), \ldots, g^{n-1}(x)\} \subset C$. The convex hull $\mathcal{C}_1$, of these points is contained in C and also $\mathcal{C}_1$ is invariant under g. But then existence of a fixed point $p \in \mathcal{C}_1 \subset C$ for g contradicts the fact that C/Γ is a manifold. This contradiction shows that $\mathbf{H}^n/\Gamma$ is a

manifold.

□

2.4.3 Definition. Given a convex, metrically complete hyperbolic manifold M, with boundary, embedded in a metrically complete hyperbolic manifold, and an $\epsilon>0$, we define M_ϵ to be the ϵ-neighbourhood of M. ϵ does not need to be small. Unless otherwise stated we shall always assume that M is embedded in the manifold given by Theorem 2.4.1 *(Embedding convex manifolds)*.

2.4.4. Proposition: Epsilon neighbourhood smooth. *Let C be a closed, convex subset of $\mathbf{H}^n$. Then the distance function $\Delta:\mathbf{H}^n\setminus C\to\mathbf{R}$ given by $\Delta(x)=d(x,C)$ is C^1.*

For an elegant proof of this, due to Brian Bowditch, see [Epstein-Marden].

We shall assume the following well-known result which is actually a consequence of negative curvature. (See [Douady, 1979].)

2.4.5. Lemma: Strict convexity. *If ω_1 and ω_2 are two geodesic arcs in $\mathbf{H}^n$ parametrized proportional to arc length, then $f(x,y)=d(\omega_1(x),\omega_2(y))$ is a convex function defined on $\mathbf{R}\times\mathbf{R}$. Furthermore, f is strictly convex unless ω_1 and ω_2 are contained within the same geodesic.*

The next result is a vital step in the proof of Theorem 2.5.1 *(Nearby structures have convex thickenings)*.

2.4.6. Theorem: Geodesics dip. *Let M be a convex hyperbolic manifold and let ω be a geodesic, parametrized by arc length, of length $L\geq 2\eta$ in M_ϵ. There is a continuous function $\delta(\epsilon,\eta)>0$ such that $d(\omega(t),M)\leq\epsilon-\delta$ for all $t\in[\eta,L-\eta]$. δ is given by*

$$\delta(\epsilon,\eta) = \epsilon - \operatorname{arcsinh}\left(\frac{\sinh\epsilon}{\cosh\eta}\right).$$

Note that $\delta(\epsilon,\eta)<\eta$ and that $\delta(\epsilon,\eta)$ is a monotonic increasing function of η.

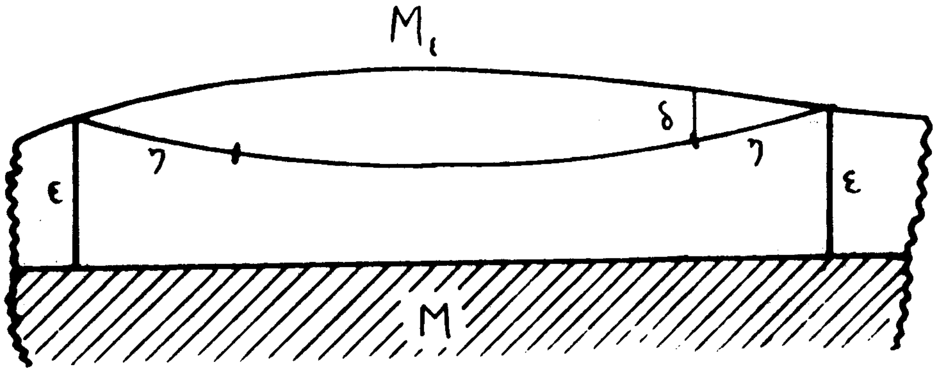

2.4.8 Figure.

Before we prove this, we prove two lemmas.

2.4.7. Lemma: Estimate for distance. *Let λ be a geodesic in $\mathbf{H}^2$. Let ω be a geodesic, and set $u = d(\omega(0),\lambda)$. Let the angle between $\omega'(0)$ and the perpendicular to λ through $\omega(0)$, oriented away from λ, be θ. Then*

$$\sinh d(\omega(t),\lambda) = \sinh u \cosh t + \sinh t \cosh u \cos\theta.$$

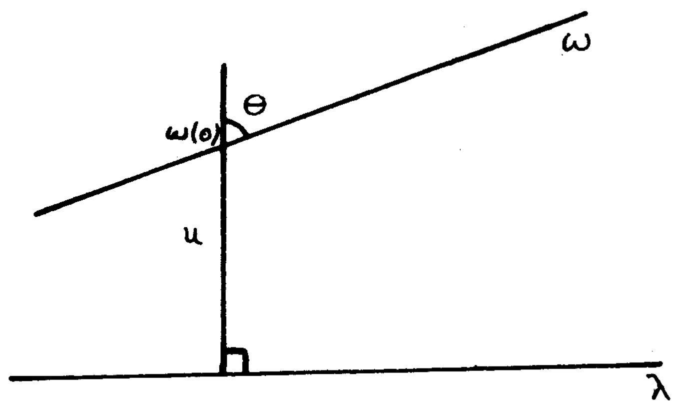

2.4.9 Figure.

Proof. Left to reader. □

2.4.10. Lemma: Distance function convex. *If C is a closed convex subset of $\mathbf{H}^n$, then $d(\omega(t),C)$ is convex.*

Proof. Suppose, for a contradiction, that

$$d(\omega(\alpha),C) > \alpha d(\omega(0),C) + (1-\alpha)d(\omega(1),C)$$

for some $0 < \alpha < 1$. Let $\gamma : I \to C$ be the geodesic with $\gamma(0) = r\omega(0)$ and $\gamma(1) = r\omega(1)$. We have

$$d(\omega(0),C) = d(\omega(0),\gamma(0)) \text{ and } d(\omega(1),C) = d(\omega(1),\gamma(1)).$$

Then

$$d(\omega(\alpha),\gamma(\alpha)) \geq d(\omega(\alpha),C) > \alpha d(\omega(0),\gamma(0)) + (1-\alpha)d(\omega(1),\gamma(1))$$

which is a contradiction.

□

We can know proceed with the proof of the theorem.

Proof. Geodesics dip: It is sufficient to prove the result for $M \subset \mathbf{H}^n$, since we can lift to the universal cover. In fact we shall prove it for $M = C$, an arbitrary closed convex subset. By Lemma 2.4.9 *(Distance function convex)*, $d(\omega(t),C) \leq \epsilon$ for $0 \leq t \leq L$. We must show that for $\eta \leq t_1 \leq L-\eta$,

$$d(\omega(t_1),C) \leq \epsilon - \delta(\epsilon,\eta).$$

Defining $\omega_1(t) = \omega(t+t_1)$, we see that what we have to show is that if $d(\omega_1(t),C) \leq \epsilon$ for $-\eta \leq t \leq \eta$ then

$$d(\omega_1(0),C) \leq \epsilon - \delta(\epsilon,\eta).$$

By the definition of $\delta(\epsilon, \eta)$ this is equivalent to showing that

$$\cosh\eta . \sinh d(\omega_1(0),C) \leq \sinh\epsilon.$$

This is obvious if $d(\omega_1(0),C) = 0$. So we may assume that $d(\omega_1(0),r\omega_1(0)) > 0$, where r is the nearest point retraction onto C. Let P be the $(n-1)$-dimensional subspace through $r\omega_1(0)$. orthogonal to $[\omega_1(0),r\omega_1(0)]$ and let H be the halfspace with boundary P, not containing $\omega_1(0)$. Then $C \subset H$. Hence $\epsilon \geq d(\omega_1(t),C) \geq d(\omega_1(t),H)$. Also $d(\omega_1(0),H) = d(\omega_1(0),r\omega_1(0)) = d(\omega_1(0),C)$.

After possibly changing the direction of ω, we may assume that the angle θ between $[r\omega_1(0),\omega_1(0)]$ and the tangent vector at $t = 0$, $\omega_1'(0)$, satisfies $0 \leq \theta \leq \pi/2$. Taking λ to be the line containing the orthogonal projection of $|\omega|$ in P, $t = \eta$ and $u = d(\omega_1(0),r\omega_1(0))$, Lemma 2.4.8 *(Estimate for distance)* gives $\sinh\epsilon \geq \sinh d(\omega_1(\eta),P) \geq \sinh d(\omega_1(0),C)\cosh\eta$ as required.

Geodesics dip

2.4.11. Corollary: Convex thickens to strictly convex. *If M is a convex hyperbolic manifold, M_ϵ (see Definition 2.4.3* (Epsilon neighbourhood of manifold)*) is strictly convex for all $\epsilon > 0$.*

2.5. Convex Thickenings

See Proposition 8.3.3 of [T].

2.5.1. Theorem: Nearby structures have convex thickenings. *Let N be a compact manifold with boundary. let M be a convex hyperbolic structure on N and let M_{Th} be a thickening of M. Then there is a neighbourhood U of M in $\Omega(N)$ such that, if $M' \in U$, then M' can be thickened to a compact convex manifold M''.*

Proof. By Theorem 1.7.1 *(Neighbourhood is a product)*, we may assume that all the manifolds in $\mathcal{U}$ are embedded in N_{Th} by an embedding near to the identity $i: N \to N_{Th}$ and that the hyperbolic structure on N_{Th} inducing the given structure M' is near to that of M_{Th}. It follows that we can restrict our attention to the varying hyperbolic structure on N_{Th}, and not worry about the embedding of N in N_{Th}. We shall use metrics which come from the path metric on even larger thickenings. (We have to be careful here because of examples such as that shown in Figure 1.5.7 .)

We choose $\epsilon > 0$ so that M_{Th} is a 10ϵ-thickening of M. We choose $\mathcal{U}$ sufficiently small so that, if $M' \in \mathcal{U}$, then the corresponding hyperbolic structure on N_{Th} makes N_{Th} a 9ϵ-thickening of N. This is possible by Corollary 1.7.5 *(Epsilon thickenings exist)*. We define $\delta = \delta(\epsilon, \epsilon/2)$, where δ is the function defined in Theorem 2.4.6 *(Geodesics dip)*. Recall that $\delta < \epsilon/2$. Let $K_0 \subset \operatorname{int} K_1 \subset K_1 \subset N_{Th}$ where K_0 and K_1 are compact connected subspaces, and the translates of K_0 by the covering translations cover N_{Th}. We assume that K_1 contains an r-neighbourhood of K_0 where $(r-1)$ is larger than the diameter of M_{Th}. We also assume that $\mathcal{U}$ is small enough so that $d(Dx, D'x) < \delta/4$, for all $x \in K_1$, where $D, D' : N_{Th} \to \mathbf{H}^n$ are developing maps for our fixed convex hyperbolic manifold M and for M' respectively. We shall also assume that $D'|K_1$ is an embedding (see 1.5.4 *(Space of developing maps)*). Then K_1 is an $(r-\delta)$-neighbourhood of K_0 in the metric d' related to the structure M'.

To define the convex thickening claimed in the statement, we fix $M' \in \mathcal{U}$ and let $\mathcal{A}$ be the collection of $(n+1)$-tuples of points $(x_0, \ldots, x_n)$ with $x_i \in M'_\epsilon$ and such that $d'(x_i, x_j) \leq \epsilon$. For each n-tuple $(x_0, \ldots,$

$x_n) \in \mathcal{A}$, we take the convex hull $\mathcal{C}(x_0, \ldots, x_n)$ and set $U = \bigcup_{\mathcal{A}} \mathcal{C}(x_0, \ldots, x_n)$. (The convex hull is defined inside an ϵ-ball centred on x_0.) We claim that U is the compact convex thickening we seek. The fact that U is compact is an immediate consequence of the compactness of $\mathcal{A}$. The fact that U is convex will be deduced from its local convexity (see Corollary 1.3.8 *(Local convexity implies convexity)*). Given $u,v \in U$, such that $d'(u,v) \leq \epsilon$, we show that the geodesic interval $[u,v]$ is in U. Let $u \in \mathcal{C}(u_0, \ldots, u_n)$ and $v \in \mathcal{C}(v_0, \ldots, v_n)$, where u_i, $v_i \in M'_\epsilon$ for $0 \leq i \leq n$, and $d'(u_i,u_j) \leq \epsilon$ and $d'(v_i, v_j) \leq \epsilon$ for $0 \leq i < j \leq n$. Then

$$d(u_i,v_j) \leq d(u_i,u) + d(u,v) + d(v,v_j) \leq 3\epsilon .$$

It follows that we need only establish the following claim. Let $\{w_0, \ldots, w_n\}$ be contained in an ϵ-neighbourhood of $D'(K_1 \cap N) \subset \mathbf{H}^n$, and in a 4ϵ-neighbourhood of some point $x \in D'(K_0 \cap N)$, and let $d(w_i, w_j) \leq 3\epsilon$. Then $\mathcal{C}(w_0, \ldots, w_n) \subset D'(K_1 \cap U)$, where U is the inverse image of U in the universal cover.

There is no loss of generality in supposing that $d(w_0,w_1)$ maximizes $d(w_i,w_j)$ $(0 \leq i < j \leq n)$. If $d(w_0, w_1) \leq \epsilon$ there is nothing to do. So we suppose that $d(w_0,w_1) > \epsilon$. Let z be the midpoint of $[w_0,w_1]$. We divide $\mathcal{C}(w_0, \ldots, w_n)$ into two pieces, $\mathcal{C}(z,w_1,w_2, \ldots, w_n)$ and $\mathcal{C}(z,w_0, w_2, \ldots, w_n)$. We want to show that each of these pieces satisfies the hypotheses for the claim. By the definition of $\mathcal{U}$, $d(w_i,D(K_1 \cap N)) \leq \epsilon + \delta/4$ for $0 \leq i \leq n$. Since DN is convex, Theorem 2.4.6 *(Geodesics dip)* implies that

$$\begin{aligned} d(z,D(\tilde{N})) &\leq \epsilon + \delta/4 - \delta(\epsilon + \delta/4, \epsilon/2) \\ &\leq \epsilon + \delta/4 - \delta(\epsilon, \delta/2) \\ &\leq \epsilon - 3\delta/4 . \end{aligned}$$

The second inequality is a result of the monotonicity of the function $\delta(\epsilon, \eta)$ in the variable ϵ, pointed out just after Theorem 2.4.6 *(Geodesics dip)*. It follows that

$$d(z,D'(K_1 \cap \tilde{N})) \leq \epsilon - \delta/2 < \epsilon .$$

The final point to check is that this chopping in half process gives a figure with all sides having length less than ϵ after a finite number of steps. This is not obvious, because, although the longest side is divided in two, the other sides will, in general, become longer than they were. Let $m = d(w_0,w_1)$, and fix m for the next few steps. At each step the number of edges of length greater than $0.9m$ decreases. Therefore, after a finite number of steps, all sides have length less than $0.9m$. (This can be seen by doing a computation in euclidean geometry. Because we are dealing with a small region, euclidean estimates give rise to hyperbolic estimates, though the hyperbolic estimates are slightly worse.) Continuing in this manner, the length of each side will eventually become less than ϵ.

Nearby structures have convex thickenings

2.5.2. Corollary: Nearby structures strongly complete. *With the same hypotheses as in the theorem above, if $M' \in U$, $D_{M'}$ is a homeomorphism, and M' can be embedded in a unique complete hyperbolic manifold with the same fundamental group.*

2.5.3 Remarks. We can see that convexity of the original manifold with boundary is a necessary condition for the above corollary to hold, by considering $M = S^3 \backslash N(K)$ (where K is a knot whose complement admits a complete hyperbolic structure.) By considering Dehn surgery space, we see that nearby hyperbolic structures on M, do not necessarily admit extensions to complete hyperbolic manifolds. See [T] for more details.

Chapter 3. Spaces of Hyperbolic Manifolds

3.1. The Geometric Topology

We shall define a topology on the set of closed subsets of a topological space. We shall thereby derive topologies for both the space of complete hyperbolic manifolds and the space of geodesic laminations (see Section 4.1 *(Geodesic Laminations)*). This topology was first considered by Chabauty [Chabauty, 1950] as a topology on the space of closed subgroups of a locally compact topological group, and later by Harvey [Harvey, 1977] with specific reference to Fuchsian groups. See also [Michael, 1951].

3.1.1 Definition. Given a topological space X, the *Chabauty topology* on $C(X)$ (the set of all closed subsets of X) has a sub-basis given by sets of the following form:

1) $O_1(K) = \{A \mid A \cap K = \varnothing\}$ where K is compact.

2) $O_2(U) = \{A \mid A \cap U \neq \varnothing\}$ where U is open.

If X is compact and metrizable, the Chabauty topology agrees with the topology induced by the Hausdorff metric. The Chabauty topology has the following nice topological properties.

3.1.2. Proposition: Properties of Chabauty topology. *Let X be an arbitrary topological space (no particular assumptions), then*

1) *$C(X)$ the set of closed subsets of X with the Chabauty topology is compact.*

2) *If X is Hausdorff, locally compact and second countable, $C(X)$ is separable and metrizable.*

Proof. 1) By Alexander's Sub-base Theorem [Rudin, 1973, page 368], we need only show that every covering by sub-basis elements has a finite sub-covering. Let the covering consist of:

$$\{O_1(K_i)\}_{i \in I} \text{ and } \{O_2(U_j)\}_{j \in J} \quad .$$

Let $C = X \setminus \bigcup_{j \in J} U_j$. C is closed and thus $C \in C(X)$. C is not in $O_2(U_j)$ for any j, therefore $C \in O_1(K_i)$ for some i. $\{U_j\}_{j \in J}$ is a covering for K_i, so there exists a finite sub-covering $\{U_{j(1)}, \ldots, U_{j(n)}\}$. Given a closed subset L, either $L \cap U_{j(k)} \neq \varnothing$ for some $k = 1, \ldots, n$ (i.e. $L \in O_2(U_{j(k)})$), or $L \cap K_i = \varnothing$ (i.e. $L \in O_1(K_i)$). Thus, $C(X) = O_1(K_i) \cup \bigcup_{k=1}^{k=n} O_2(U_{j(k)})$; which is the desired finite sub-covering.

2) Since X is Hausdorff and locally compact, it is also regular. Suppose $K, L \in C(X)$ and $x \in K \setminus L$. Because X is regular and locally compact, there exists an open set U such that $x \in U \subset \overline{U} \subset X - L$ where $\overline{U}$ is compact. So $K \in O_2(U)$, $L \in O_1(\overline{U})$, and $O_2(U) \cap O_1(\overline{U}) = \varnothing$. It follows that $C(X)$ is Hausdorff.

Let $\{B_1, \ldots, B_n, \ldots\}$ be a countable basis for X, such that $\overline{B_i}$ is compact for each i. We claim that $\{O_1(\overline{B_i})\} \cup \{O_2(B_i)\}$ is a sub-basis for the Chabauty topology on $C(X)$. Let $O_1(K)$ be a sub-basis element (as in our Definition 3.1.1 *(Chabauty topology)*.), $C \in O_1(K)$, and $\{B_{n(1)}, \ldots, B_{n(k)}\}$ a covering for K such that $C \cap \bar{B}_{n(i)} = \emptyset$ for $1 \leq i \leq k$. Then

$$C \in O_1(\bar{B}_{n(1)}) \cap \ldots \cap O_1(\bar{B}_{n(k)}) \subset O_1(K).$$

Now suppose $C \in O_2(U)$ and $x \in C \cap U$. Choose B, a neighbourhood of x such that $B \subset U$. Then $C \in O_2(B) \subset O_2(U)$. It follows that we have a countable subbasis consisting of sets of the form $O_1(B_i)$ and $O_2(B_i)$. Then, by the Urysohn Metrization Theorem $C(X)$ is separable and metrizable. □

The following easily proved lemma exposes the essential geometric nature of our topology.

3.1.3. Lemma: Geometric convergence. *Suppose X is a locally compact metric space. A sequence $\{A_n\}$ of closed subsets of X converges in $C(X)$ to the closed subset A if and only if:*

1) *If $\{x_{n(k)}\} \in \{A_{n(k)}\}$ converges to $x \in X$ then $x \in A$.*
2) *If $x \in A$, then there exists a sequence $\{x_n\}$, where each x_n is an element of A_n, converging to x.*

The proof is left to the reader.

We now restrict ourselves to the case of closed subgroups of Lie groups, and prove Chabauty's original theorem [Chabauty, 1950] (see also [Harvey, 1977].) The set of closed subgroups of a Lie group L is closed in the Chabauty topology and is thus compact and metrizable. There is a right invariant Haar measure on any Lie group, which induces a measure on $\Gamma \backslash L$ when Γ is discrete. Denote the total volume (which may be infinite), by $\mu\Gamma$. When L is $\mathrm{Isom}(\mathbb{H}^n)$, we shall consider instead $\mu\Gamma = \mathrm{vol}(\mathbb{H}^n/\Gamma)$. The two versions of $\mu\Gamma$ are equal if we normalize correctly.

3.1.4. Theorem: A(U) compact. *Let $\mathcal{G}(L)$ be the space of closed subgroups of a Lie group L with the Chabauty topology (so that $\mathcal{G}(L)$ is compact and metrizable). Let U be an open neighbourhood of $\{e\}$ in L. then*

1) $A(U) = \{G \in \mathcal{G}(L) \mid G \cap U = \{e\}\}$ *is compact.*
2) $B(U) = \{G \in A(U) \mid G$ is torsion-free$\}$ *is compact.*
3) *The set of discrete subgroups of L is open in the space of all closed subgroups. It is the union of the interiors in $\mathcal{G}(L)$ of the compact spaces $A(U)$, as U varies over open neighbourhoods of $\{e\}$ in L.*
4) *If $\{\Gamma(n)\}$ is a sequence of discrete subgroups converging to the discrete subgroup Γ, then*

$$\mu_\Gamma \leq \liminf_{n\to\infty} \mu_{\Gamma(n)} .$$

In particular, $A(U,M) = \{G \in A(U) | \mu(G) \leq M\}$, *and* $B(U,M) = \{G \in B(U) | \mu(G) \leq M\}$ *are compact.*

Proof. 1) We observe that $A(U) = \mathcal{G}(L) \setminus O_1(U \setminus \{e\})$ is a closed subset of the compact space $\mathcal{G}(L)$.

2) Suppose $\Gamma_i \in B(U)$ converges to Γ (which is in $A(U)$ by (1)). If Γ is not torsion-free, there is some $\gamma \in \Gamma$ and $n > 0$ such that $\gamma^n = e$. By Lemma 3.1.3 *(Geometric convergence)* there exist $\gamma_i \in \Gamma_i$ such that $\{\gamma_i\}$ converges to γ. Thus, $\{\gamma_i^n\}$ converges to e, but this would imply that $\gamma_i^n = e$ for large values of i, which is a contradiction.

3) Let V be a compact neighbourhood of e in L small enough not to contain a non-trivial subgroup and let U be a smaller open neighbourhood of e in L such that $U^2 \subset V$. Let $K = V \setminus U$. Then any closed subgroup Γ such that $\Gamma \cap K = \emptyset$ also satisfies $\Gamma \cap V = \{e\}$. To see this, let $g \in \Gamma \cap V$. Then $g \in \Gamma \cap U$. Let $n > 0$ be the smallest integer, if any, such that $g^n \notin U$. Then $g^n = g^{n-1} g \in U^2 \subset V$. This is impossible. Hence $g^n \in U \subset V$ for all $n > 0$ and so $g^n \in V$ for all $n \in \mathbf{Z}$. But V contains no subgroup, so $g = e$. What we have shown is that $O_1(K)$ consists of discrete subgroups. This proves that any discrete subgroup has an open neighbourhood consisting of discrete subgroups.

We caution the reader that a current important conjecture is that the set of discrete faithful hyperbolic representations of an abstract finitely generated abstract group in $PSL(2,\mathbf{C})$ is closed. This is with respect to a different topology from the Chabauty topology, namely the compact-open topology.

4) Choose a non-empty open set $W \subset \mathbf{H}^n$ (or, in general, $W \subset L$) such that $W \cap T(W) = \emptyset$ for all $T \in (\Gamma \setminus \{e\})$, and a non-empty compact subset K' of W. Let $C = \{T \in L : T(K') \cap K' \neq \emptyset\}$. C is a compact subset of L. Also let V be an open neighbourhood of e in L such that $\Gamma \cap V = \{e\}$. Then $\Gamma \in O_1(C \setminus V)$. Therefore $\Gamma(n) \in O_1(C \setminus V)$ for $n \geq N$. We can also assume that $\Gamma(n) \cap V = \{e\}$ for $n \geq N$ by the proof of 3). Hence, if $n \geq N$ and $T_n \in \Gamma(n)$, then $T_n = e$ or $T_n(K') \cap K' = \emptyset$.

Now W may be chosen to have the same measure as μ_Γ and K may be chosen to have measure arbitrarily close to μ_Γ. Therefore,

$$\mu_\Gamma \leq \liminf_{n\to\infty} \mu_{\Gamma(n)} .$$

□

3.1.5 Remarks. For those familiar with the theory of Kleinian groups, the Chabauty topology on the space of discrete subgroups of $PSL(2,\mathbb{C})$ is equivalent to the topology induced by convergence of Poincaré (Dirichlet) fundamental polyhedra (with a fixed origin) and their associated face-pairings.

We now obtain a few useful results of Thurston's (see Section 8.8 of [T] and [Thurstona]) as corollaries of Chabauty's Theorem.

3.1.6 We can think of a complete hyperbolic manifold provided with a frame (M,e) as a discrete torsion free subgroup Γ of Isom($\mathbf{H}^n$). To do this, we fix a point $s_0 \in \mathbf{H}^n$ and an orthonormal frame $(s_1, \ldots, s_n)$ for the tangent space to $\mathbf{H}^n$ at s_0. We refer to this fixed choice as the *standard frame* in $\mathbf{H}^n$. Then we choose the developing map for M which takes some lift of e to s, our standard frame in $\mathbf{H}^n$. The holonomy of this developing map gives an unique Möbius group Γ. Then $(\mathbf{H}^n \backslash \Gamma, s)$ is a complete hyperbolic manifold with baseframe which is isometric (in a frame-preserving way) to (M,e). In this way we can topologize the space $\mathfrak{MF}^n$ of complete hyperbolic $\mathfrak{M}$anifolds of dimension n with base $\mathfrak{F}$rame, using the Chabauty topology on the set of subgroups Γ of Isom($\mathbf{H}^n$). We call this the *geometric* topology on $\mathfrak{MF}^n$. We topologize $\mathfrak{MB}^n$, the space of complete hyperbolic $\mathfrak{M}$anifolds of dimension n with $\mathfrak{B}$asepoint as the quotient of $\mathfrak{MF}^n$, and $\mathfrak{MW}^n$, the space of complete hyperbolic n-$\mathfrak{M}$anifolds $\mathfrak{W}$ithout basepoint as a further quotient. All of these topologies are known as the geometric topology.

By Part 3) of Theorem 3.1.4 *(A(U) compact)*, $\mathfrak{MF}^n$ is a locally compact Hausdorff space. $\mathfrak{MB}^n$ is the quotient of $\mathfrak{MF}^n$ by the compact group $O(n)$, and is therefore Hausdorff. $\mathfrak{MW}^n$ is not Hausdorff.

3.1.7. Corollary: Set of hyperbolic manifolds compact.

1) $\mathfrak{MF}^n(\epsilon)$, *the space of complete hyperbolic manifolds with frame, having injectivity radius bounded below by* ϵ *at the basepoint is compact for any* $\epsilon > 0$,

2) $\mathfrak{MB}^n(\epsilon)$, *the space of complete hyperbolic* n*-manifolds with basepoint having injectivity radius not less than* ϵ *at the basepoint, is compact.*

3) *The space of complete hyperbolic* n*-manifolds,* $\mathfrak{MW}^n$, *is compact.*

Proof. Let x_0 be a fixed basepoint for $\mathbf{H}^n$ and let Γ be a torsion free discrete group of isometries for $\mathbf{H}^n$. To prove the first statement, we simply observe that the injectivity radius at the basepoint is equal to $\inf\{d(x_0,\gamma x_0) | \gamma \in \Gamma\}/2$. By the previous theorem the set of such Γ is compact. The second statement follows since $\mathfrak{MB}^n(\epsilon)$ is the image of $\mathfrak{MF}^n(\epsilon)$

under the obvious projection map. To prove 3) we recall that every hyperbolic n-manifold has a point with injectivity radius not less than ϵ (the Margulis constant for $\mathbf{H}^n$), so $\mathcal{MW}^n = Y(\mathcal{MB}^n(\epsilon))$, where $Y:\mathcal{MB}^n \to \mathcal{MW}^n$ is the map which forgets the basepoint.

□

3.1.8. Corollary: Compact with bounded volume.

1) *The space of complete hyperbolic n-manifolds with basepoint having injectivity radius at the basepoint not less than ϵ and volume not more than V, $\mathcal{MB}^n(\epsilon, V)$, is compact for any $\epsilon>0$ and any $V>0$.*

2) *The set of complete hyperbolic n-manifolds with volume not more than V, $\mathcal{MW}^n(V)$ is compact for any $V>0$.*

3.1.9 Definition. A *marked hyperbolic surface* of finite type is a topological surface T of finite type together with an isotopy class of homeomorphisms $h : T \to S$, where S is a complete hyperbolic surface. Two marked surfaces $h_1: T \to S_1$ and $h_2: T \to S_2$ are said to be *equivalent* (or, sometimes, *equal)* if there is an isometry $\phi : S_1 \to S_2$ such that ϕh_1 is isotopic to h_2.

3.1.10 The space of all equivalence classes of marked hyperbolic surfaces of a given homeomorphism type (where we specify also whether a puncture is to be a funnel or a cusp) is called Teichmüller space and we denote it $\mathbf{T}(T)$. An equivalence class of marked hyperbolic surfaces clearly determines and is determined by its holonomy map $H_{[h]}:\pi_1(T) \to \text{Isom}(\mathbf{H}^2)$ which is well-defined up to conjugacy. We may thus topologize Teichmüller space as a subspace of $R(\pi_1(T),\text{Isom}(\mathbf{H}^2))$(the set of conjugacy classes of representations of $\pi_1(T)$ into $\text{Isom}(\mathbf{H}^2)$) with the compact open topology (i.e. $\{\rho_i\}$ converges if $\{\rho_i(x)\}$ converges for every $x \in \pi_1(T)$.) Teichmüller space is an open subspace of $R(\pi_1(T),\text{Isom}(\mathbf{H}^2))$ and is in fact homeomorphic to $\mathbf{R}^{6g-6+2p+3f}$ where g is the genus, p is the number of punctures and f is the number of funnels. For further information on Teichmüller space see [Abikoff, 1980] or

To discuss the connection with the geometric topology, we generalize to any number of dimensions. We first remove the annoyance of dealing with conjugacy classes, by taking manifolds with baseframe. We can adapt the geometric topology to this situation, to include the marking. We topologize the space of injective homomorphisms $\pi_1(T) \to \text{Isom}(\mathbf{H}^n)$ onto a discrete subgroup of $\text{Isom}(\mathbf{H}^n)$, by using the Chabauty topology on the image, with a further refinement. We are allowed to specify a finite number (depending on the neighbourhood we are trying to describe) of elements of $\pi_1(T)$ and demand that the images of each of these lies in a certain open subset (depending on the element) of Isom($\mathbf{H}^n$). We call this the *marked geometric topology*. The compact-open topology on the space of homomorphisms is known as the *algebraic topology*.

There is an obvious continuous map from the marked geometric topology to the algebraic topology. This map is not a homeomorphism in general, though we will prove below that it is a homeomorphism for surfaces of finite type and also for manifolds of finite volume.

A counterexample to the map being a homeomorphism can be made for a surface of infinite type. We give a quick sketch. The basic building block is an infinite strip with one handle. The strip has two boundary components. We specify a hyperbolic structure on the building block, by insisting that the boundary components are infinite geodesics, which are asymptotic at each end, and that there is an orientation reversing isometry, interchanging the two ends of the strip, whose fixed point set consists of a geodesic arc joining the two boundary components and a geodesic simple closed curve going around the handle. There are two parameters for the hyperbolic structure, namely the length of the arc and the length of the circle. We glue a countable number of building blocks together in the simplest possible way. In the algebraic limit, the sum of the lengths of the arcs is convergent. The nth surface has all except a finite number of these lengths equal to 1. Details are left to the reader.

The finitely generated case is more important. Here a counterexample can only be given in dimensions greater than two. For example, let $\rho_n : \mathbf{Z} \to PSL(2,\mathbf{C})$ be generated by

$$\rho_n(1) = \begin{bmatrix} e^{w_n} & n \sinh w_n \\ 0 & e^{-w_n} \end{bmatrix} \quad \text{where } w_n = \frac{1}{n^2} + \frac{\pi i}{n}.$$

Then ρ_n converges to $\rho : \mathbf{Z} \to PSL(2,\mathbf{C})$ where

$$\rho(1) = \begin{bmatrix} 1 & \pi i \\ 0 & 1 \end{bmatrix}$$

in the algebraic topology. But $\{\rho_n(n)\}$ converges to

$$\begin{bmatrix} -1 & -1 \\ 0 & -1 \end{bmatrix} = \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix}.$$

Since $\{\rho_n(\mathbf{Z})\} \subset B(U)$ for some open neighbourhood U of e, the geometric limit of $\{\rho_n(\mathbf{Z})\}$ must be discrete and torsion-free, and thus a parabolic subgroup of rank 2. For a more in-depth discussion of algebraic and geometric convergence see Chapter 9 of [T]. Here we confine ourselves to the following case where the geometric and algebraic topologies coincide.

3.1.11. Proposition: Algebraic equals geometric. *We restrict our attention to finite volume complete hyperbolic manifolds (with baseframe). The map from the space of such manifolds with the geometric topology to the same space with the algebraic topology is a homeomorphism.*

Proof. Suppose $\{\rho_i:\pi_1 M \to \mathrm{Isom}(\mathbf{H}^2)\}$ converges to $\rho:\pi_1 M \to \mathrm{Isom}(\mathbf{H}^2)$ in the algebraic topology. We must show that it converges also in the geometric topology. Let $\Gamma_i=\rho_i(\pi_1 M)$ and $\Gamma=\rho(\pi_1 M)$. We choose $\gamma\in\Gamma$ to be a hyperbolic element which is not divisible. Let $\alpha\in\pi_1 M$ be defined by $\rho(\alpha)=\gamma$. Then α is indivisible. Let $\gamma_i=\rho_i(\alpha)\in\Gamma_i$. By passing to a subsequence, we may assume that each γ_i is hyperbolic.

We claim that for some neighbourhood U of the identity, $\Gamma_i\cap U=\{e\}$, for i large. For otherwise (taking a subsequence if necessary) there is a sequence of non-trivial indivisible elements $\beta_i\in\Gamma_i$, such that β_i converges to the identity. Then $[\gamma_i,\beta_i]$ converges to the identity. But then $[\gamma_i,\beta_i]$ must commute with β_i for i large, by the Margulis lemma.

If β_i is hyperbolic, this means that $[\gamma_i,\beta_i]$ is a power of β_i, so that γ_i normalizes the subgroup generated by β_i. Therefore γ_i is a power of α_i, which is impossible, since γ_i is indivisible and is not small. If β_i is parabolic, then $[\gamma_i,\beta_i]$ is either trivial or must be parabolic with the same unique fixed point p. But then $[\gamma_i,\alpha_i]\alpha_i^{-1}=\gamma_i^{-1}\alpha_i\gamma_i$ fixes p, and so γ_i fixes p. Since γ_i is hyperbolic, this is also impossible. This proves the existence of the claimed neighbourhood U.

By passing to a subsequence, we may assume that Γ_i converges to $G\subset\mathrm{Isom}(\mathbf{H}^n)$ in the Chabauty topology. Then $G\cap U=\{e\}$ and so G is discrete. Clearly $\Gamma\subset G$.

We want to prove $\Gamma=G$. So suppose $g\in G\setminus\Gamma$. Replacing g by $g\gamma^k$ for a large value of k, we see that there is no loss of generality in assuming that g is hyperbolic. Since Γ has finite co-volume, the index of Γ in G must be finite. It follows that $g^k=\rho(h)$ for some $k>0$ and some $h\in\pi_1 M$.

Let $g_i\in\Gamma_i$ converge to g and let $h_i=\rho_i(h)$. Then $g_i^{-k}h_i$ converges to the identity. But we have already shown that this implies $h_i=g_i^k$ for large i. This means that g_i centralizes h_i. Now the centralizer of a hyperbolic element in a discrete group is infinite cyclic, and so kth roots are unique. Hence h has a unique kth root α' in $\pi_1 M$. Then $g_i=\rho_i(\alpha')$ and so $g=\rho(\alpha')$. Hence $g\in\Gamma$, a contradiction.

□

Remark. In the case of a surface M of finite type, the algebraic topology and the geometric topology coincide, even when we allow surfaces with infinite area. This can be seen by breaking the surface up into a finite number of pairs of pants. The geometry of each pair of pants, and the way they are glued together, is determined by a finite number of elements of $\rho(\pi_1 T) \subset \mathrm{Isom}(\mathbf{H}^2)$.

3.2. ϵ-relations and Approximate Isometries

We now consider a generalization of the concept of Hausdorff metric which was developed by M. Gromov. Intuitively, compact metric spaces can be approximated very well by finite subsets of points and locally compact path spaces can be approximated very well by large compact subsets of themselves. This simple idea developed into the more formal notion of (ϵ, r)-relations, which provide us with a topology on the space of all complete locally compact path spaces. One of the first applications of this was also one of the most amazing; Gromov used ϵ-relations to show that if a finitely generated group has polynomial growth then it contains a nilpotent subgroup of finite index. ([Gromov, 1981a], [Gromov, 1981b]) contains an extensive investigation into the space of Riemannian manifolds with the geometric topology. His definitions are related to, but slightly different from ours. In this section we will establish that the topology induced by (ϵ, r)-relations when restricted to the space of hyperbolic manifolds with basepoint agrees with the geometric topology.

3.2.1 Definition. Let (X, x_0) and (Y, y_0) be two compact metric spaces with basepoint. An ϵ-relation between (X, x_0) and (Y, y_0) is a relation R with the following properties.

1) $x_0 \, R \, y_0$;
2) for each $x \in X$, there exists $y \in Y$ such that $x \, R \, y$;
3) for each $y \in Y$ there exists $x \in X$ such that $x \, R \, y$;
4) if $x \, R \, y$ and $x' \, R \, y'$, then $|d_X(x, x') - d_Y(y, y')| \leq \epsilon$.

3.2.2. Lemma: Metric on metric spaces. *If we define $d((X,x_0),(Y,y_0))$ to be the infimum of all values of ϵ for which there is an ϵ-relation between (X,x_0) and (Y,y_0) we obtain a metric on the set of isometry classes of compact metric spaces with basepoint.*

Remark: Of course, the class of all compact metric spaces is not a set in standard set theory. However, one of the usual tricks can be used to get around this objection. For example, every compact metric space has a countable dense subset, so we can consider all completions of all metrics on any subset of the natural numbers.

Proof. The only thing that needs proof is that if $d((X, x_0),(Y,y_0)) = 0$, then (X,x_0) is isometric to (Y,y_0). To prove this we fix a countable dense subset $\{x_i\}$ in X. Given a sequence R_n of ϵ_n-relations between (X,x_0) and (Y,y_0), such that $\epsilon_n \to 0$, we choose points $y_{i,n} \in Y$ such that $x_i \, R_n \, y_{i,n}$. By using the Cantor diagonalization process, we can assume that $\lim_{n\to\infty} y_{i,n}$ exists for each i. We denote the limit by y_i. Then

$$d(x_i,x_j) = \lim_{n\to\infty} d(y_{i,n},y_{j,n}) = d(y_i,y_j).$$

So the map which sends $\{x_i\}$ to $\{y_i\}$ is an isometry on this countable set. It is easy to see that $\{y_i\}$ is dense in Y. It now follows that this map extends to an isometry between (X,x_0) and (Y,y_0).

$\square$

3.2.3 Definition. Two metric spaces with basepoint (X,x_0) and (Y,y_0) are (ϵ,r)-*related* if there is an ϵ-relation between compact subspaces (X_1, x_0) and (Y_1, y_0) of (X,x_0) and (Y,y_0) respectively, where $B_X(x_0, r) \subset X_1$, and $B_Y(y_0, r) \subset Y_1$. (Recall from Theorem *(Hopf Rinow)* that balls of radius r are compact in a complete, locally compact path space.)

We can use this notion to topologize the space of (isometry classes of) complete locally compact path spaces with basepoint, as follows. Let (X,x_0) be a complete locally compact path space with basepoint and let $r > 0$ and $\epsilon > 0$. We define the neighbourhood $\mathfrak{N}(X,x_0,r,\epsilon)$ to be the set of complete locally compact path spaces with basepoint (Y,y_0) such that (Y,y_0) is (ϵ,r)-related to (X,x_0).

3.2.4. Lemma: Space of path spaces Hausdorff. *The space of complete locally compact path spaces with basepoint is Hausdorff.*

Proof. Let (X,x_0) and (Y,y_0) have no disjoint neighbourhoods. We must show that they are isometric. We fix r. By the method of Lemma 3.2.2 *(Metric on metric spaces)*, we construct an isometry of $(B_X(x_0,r),x_0)$ into (Y,y_0) from a sequence of ϵ_n-relations. From the method of construction and using the fact that Y is a path space, it is easy to see that $B_X(x_0,r)$ maps onto $B_Y(y_0,r)$. Let $\phi_r:(B_X(x_0,r),x_0) \to (B_Y(y_0,r),y_0)$ be the isometry. Using Ascoli's Theorem and the Cantor diagonalization process, we can find $\phi:(X,x_0) \to (Y,y_0)$, such that, for any fixed compact subset K of X, $\phi|K$ is the limit of maps of the form $\phi_r|K$. Hence ϕ is an isometry. □

In this way we get another Hausdorff topology on $\mathfrak{MB}^n$, the space of (isometry classes of) complete hyperbolic n-$\mathfrak{M}$anifolds with $\mathfrak{B}$asepoint in addition to the geometric topology. We shall also use $\mathfrak{MW}^n$, the space of (isometry classes of) complete hyperbolic n-$\mathfrak{M}$anifolds $\mathfrak{W}$ithout basepoint, which is given the quotient topology. Note that $\mathfrak{MW}^n$ is not Hausdorff. As an example, take a compact hyperbolic surface of genus three, in which a separating simple closed geodesic becomes shorter and shorter. If the basepoint is chosen on one side of the geodesic, the limit is a punctured torus. If the basepoint is chosen on the other side, the limit is a punctured surface of genus two. If the basepoint is chosen on the geodesic, the limit is the real line. Thus the sequence has two different limit points in $\mathfrak{MW}^2$. We also see that $\mathfrak{MB}^2$ is not closed in the space of all complete locally compact path spaces.

We shall also consider $\mathfrak{MF}^n$, the space of all (isometry classes of) complete hyperbolic n-$\mathfrak{M}$anifolds with base$\mathfrak{F}$rame. Such a manifold is a pair (M,e), where M is a complete hyperbolic n-manifold and $e = (e_1, \ldots, e_n)$ is an orthonormal frame of the tangent space at a point $e_0 \in M$. To topologize this space using ϵ-relations we introduce another definition.

3.2.5 Definition. Let (M_1,e_1) and (M_2,e_2) be two Riemannian manifolds with baseframe. Then a *framed* (ϵ,r)*-relation* between (M_1,e_1) and $(M_2,e_2)$is an (ϵ,r)-relation R between (M_1,e_1) and (M_2,e_2) together with the additional requirement that $(\exp v)\,R\,(\exp v')$ if $v = \sum v_i e_i$ is a tangent vector to M at e_0, $v' = \sum v_i e'_i$ is a tangent vector to M' at e'_0 and $\sum v_i^2 \le r^2$.

A neighbourhood $\mathfrak{N}(M,e,r,\epsilon)$ of (M,e) in $\mathfrak{MF}^n$ consists of all pairs (M',e') such that there exists a framed (ϵ,r)-relation R between (M,e) and (M',e'). Note that $\mathfrak{MF}^n$ has a countable basis of neighbourhoods of each point. Until a few lemmas are established the ϵ-relation topology is annoyingly difficult to handle. The problem is that the ϵ-relation controls only

C^0 behaviour whereas we need to control derivatives as well.

3.2.6. Lemma: Injectivity radius continuous. *Let $g_r(M,e_0)$ be the infimum of the injectivity radius at x, as x varies over $B(e_0,r)$. Then for any fixed value of r, the function $g_r : \mathcal{MB}^n \to (0,\infty)$ is continuous.*

Proof. Injectivity radius continuous: Let $x_0 \in B(e_0,r)$ be a point for which the injectivity radius is minimal. Let γ be a geodesic loop in M of length $2g_r(M,e_0)$, starting and ending at x_0. Let (x_0,x_1,x_2,x_3,x_4) be equally spaced points with $x_4 = x_0$ along γ. Then, since the interior of the ball in M with centre x_0 and radius $g_r(M,e_0)$ is isometric to an open round ball in $\mathbf{H}^n$, we have

$$d(x_0,x_1) = d(x_1,x_2) = d(x_2,x_3) = d(x_3,x_0) = \frac{g_r(M,e_0)}{2}.$$

Let (M', e_0') be near (M,e_0), and let x_0', x_1', x_2', x_3' be points such that $x_i \, R \, x_0'$ for an ϵ-relation R, where ϵ is small. We write $x_4' = x_0'$. We choose geodesic paths β_1, β_2, β_3, β_4 of minimal length, such that β_i joins x_{i-1}' to x_i'. We lift in turn β_1, β_2, β_3, β_4 to geodesics γ_1, γ_2, γ_3, γ_4 in $\mathbf{H}^n$, such that the end of γ_i in the beginning of γ_{i+1} ($i = 1,2,3$). We denote the endpoints of γ_i by y_{i-1} and y_i. Now, for $i = 0,1,2$,

$$d(y_i,y_{i+2}) \geq d(x_i',x_{i+2}') \geq g_r(M,e_0) - \epsilon$$

and $d(y_i,y_{i+1}) \leq g_r(M,e_0)/2 + \epsilon$. It follows from Lemma 4.2.10 *(Curve near geodesic)* that, by taking ϵ very small, we can ensure that the piecewise geodesic $\gamma_1\gamma_2\gamma_3\gamma_4$ is very close to a geodesic. In particular the endpoints are distinct. This means that $\beta_1\beta_2\beta_3\beta_4$ is an essential loop of length at most $2g_r(M,e_0) + 4\epsilon$. Therefore the injectivity radius of (M', e_0') at x_0' is bounded above by $g_r(M,e_0) + 2\epsilon$. Since $d(e_0',x_0') \leq r + \epsilon$, we have

$$g_r(M',e_0') \leq g_{r+\epsilon}(M',e_0') + \epsilon \leq g_r(M,e_0) + 3\epsilon\ .$$

We now prove that $\liminf g_r(M', e_0') \geq g_r(M,e_0)$, where the limit is taken over based manifolds (M',e_0') converging to (M,e_0), using the topology given by ϵ-relations. So suppose this is false. We take a sequence of based hyperbolic n-manifolds of the form (M',e_0') converging to (M,e_0) and a fixed $r > 0$, such that $g_r(M',e_0') < g_r(M,e_0) - 2\epsilon_0$, where $\epsilon_0 > 0$ is small and satisfies $\epsilon_0 < g_r(M,e_0)/4$.

Let $x_0' \in B(e_0',r) \subset M'$ and let $\gamma : [0,t_0] \to M'$ be a non-constant geodesic loop based at x_0', parametrized according to path length.

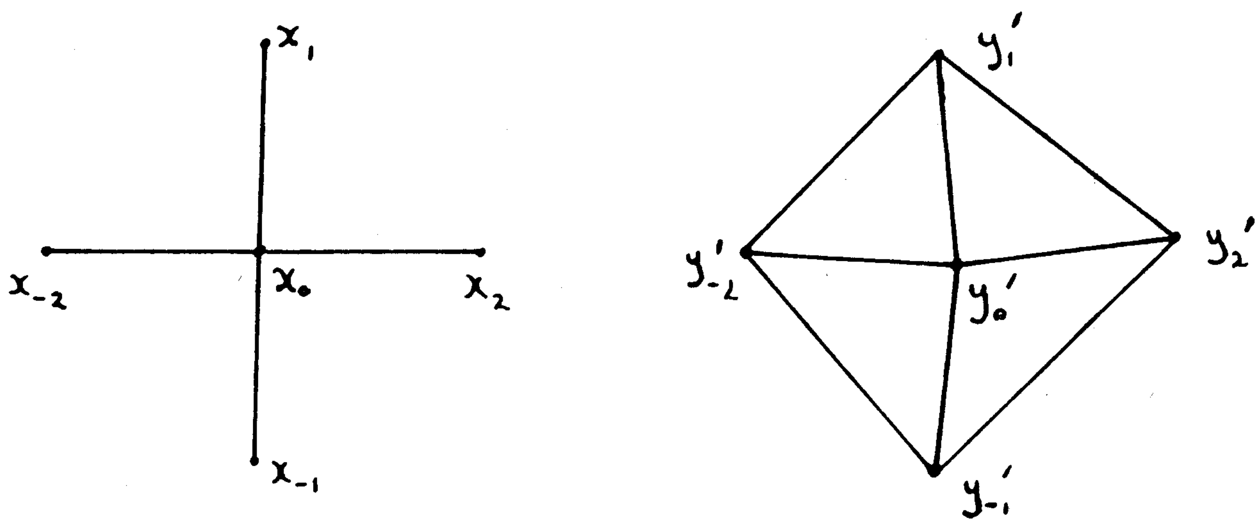

3.2.7 Figure.

We know that there is a correspondence between (M, e_0) and (M', e'_0) which gives an ϵ-relation between large neighbourhoods of e_0 and e'_0, with ϵ small. Let $x_0 \in B(e_0, r)$ correspond to x'_0 and let $(u_1, \ldots, u_n)$ be an orthonormal frame at x_0. Let $\delta = g_r(M, e_0)/2$. For $1 \leq i \leq n$, we set $x_{\pm i} = \exp(\pm \delta u_i)$. In M' let the corresponding points be $x'_{\pm i}$. We know approximately their distances apart in M'. By lifting shortest paths from x'_0 to x'_i $(1 \leq |i| \leq n)$, we obtain points y'_0, $y'_{\pm i}$ in $\mathbf{H}^n$. For each i, $y'_{-i} y'_0 y'_i$ is a once broken geodesic with $d(y'_0, y'_i) = \delta$ $(1 \leq |i| \leq n)$ and

$$d(y'_{-i}, y'_i) \geq d(x'_{-i}, x'_i) \geq 2\delta - \epsilon.$$

Therefore $y'_{-i} y'_0 y'_i$ is very near to a geodesic of length 2δ. Also, if $|i| \neq |j|$, then

$$d(y'_i, y'_j) \geq d(x'_i, x'_j) \geq d(x_i, x_j) - \epsilon = \operatorname{arccosh}(\cosh^2 \delta) - \epsilon.$$

It follows that $\angle y'_i y'_0 y'_j$ cannot be much less than $\pi/2$. The same applies to $\angle y'_i y'_0 y'_{-j}$. Since $y'_j y'_0 y'_{-j}$ is almost straight, we can deduce that $\angle y'_i y'_0 y'_j$ is approximately $\pi/2$. Therefore we can find an orthonormal basis $v_1, \ldots, v_n$ for the tangent space to $\mathbf{H}^n$ at y_0 such that $\exp(\pm \delta v_i)$ is approximately equal to $y'_{\pm j}$ $(i \leq j \leq n)$. Now let γ be a lifting of γ to $\mathbf{H}^n$, with $\gamma(0) = y'_0$. Then $\gamma(t) = \exp(t \sum \alpha_i v_i)$ for some $(\alpha_1, \ldots, \alpha_n) \in \mathbf{R}^n$ with $\sum \alpha_i^2 = 1$. Let $x \in M$ be defined by $x = \exp((t_0/2 + \epsilon_0) \sum \alpha_i u_i)$ and let $x' \in M'$ be a corresponding point. We take a shortest geodesic in M' from

x_0' to x' and lift it to a geodesic in $\mathbf{H}^n$ from y_0' to a point which we call y'. Then

$$d(y_0',y') = d(x_0',x') \leq d(x_0,x) + \epsilon = \frac{t_0}{2} + \epsilon_0 + \epsilon$$

and

$$d(y_i',y') \geq d(x_i',x') \geq d(x_i,x) - \epsilon \quad \text{for } 1 \leq |i| \leq n.$$

These inequalities specify the position of y' reasonably precisely in relation to the points $y_0', y'_{\pm i}$. Its position is almost the same as that of x relative to x_0, $x_{\pm i}$. This means that $\tilde{\gamma}(t_0/2 + \epsilon_0)$ is very near to y'. In fact their distance apart converges to zero as ϵ converges to zero. Now

$$d(x,x_0) \leq d(x',x_0') + \epsilon \leq d(y',\tilde{\gamma}(t_0/2 - \epsilon_0)) + d(\gamma(t_0/2 + \epsilon_0),\gamma(t_0)) + \epsilon.$$

The first and third terms on the right tend to zero with ϵ, and the second term is bounded by $t_0/2 - \epsilon_0$, so we see that $d(x,x_0) \leq t_0/2 - \epsilon_0$. But this contradicts $d(x,x_0) = t_0/2 + \epsilon_0$. This contradiction completes the proof.

Injectivity radius continuous

3.2.8. Lemma: Same as quotient. *The topology given by ϵ-relations on hyperbolic manifolds with basepoint is the quotient of the topology given by framed ϵ-relations on manifolds with baseframe.*

Proof. Clearly there is a continuous map $\mathcal{MF}^n \to \mathcal{MB}^n$. The injectivity radius is greater than some $\delta > 0$ throughout a neighbourhood of (M,e_0). Let $(e_1, \ldots, e_n)$ be a fixed orthonormal frame for the tangent space to M at e_0, and let $x_i = \exp(\delta e_i) \in M$. If (M',e_0') is near (M,e_0), we have points $x_1', \ldots, x_n'$ corresponding to $x_1, \ldots, x_n$. From $e_0', x_1', \ldots, x_n'$ we can form an orthonormal basis $(e_1', \ldots, e_n')$ in a canonical way as tangent vectors at e_0' as follows. We define u_i by the equation $x_i' = \exp(u_i)$, where the length of u_i is almost equal to δ, and then we obtain $(e_1', \ldots, e_n')$ by applying the Gram-Schmidt process to $(u_1, \ldots, u_n)$. If (M',e_0') is extremely near to (M,e_0), then it is easy to show that the manifold with baseframe (M',e') is near to (M,e). This shows that the map $\mathcal{MF}^n \to \mathcal{MB}^n$ is open and the lemma is proved.

□

There is an alternative way of defining a topology on $\mathcal{MF}^n$, which we have already encountered. We fix a standard orthonormal frame $(s_1, \ldots, s_n)$ for the tangent space at a fixed point $s_0 \in \mathbf{H}^n$. There is exactly one covering map $\mathbf{H}^n, s \to M, e$ which is a local isometry. The fundamental group of M acts on $\mathbf{H}^n$ as a group of covering translations. Thus we can write $\pi_1(M,e)$, and we get a well-defined subgroup of Isom($\mathbf{H}^n$). Recall that subgroups of Isom($\mathbf{H}^n$) are topologized with the Chabauty topology.

3.2.9. Theorem: Map to torsion free subgroups is a homeomorphism. *The map* $\pi_1 : \mathfrak{MJ}^n \to \mathfrak{DJ}^n$, *which sends* (M,e) *to the discrete torsion free subgroup* $\pi_1(M,e)$ *of* Isom($\mathbf{H}^n$), *is a homeomorphism.*

Proof. Map to torsion free subgroups is a homeomorphism: To prove the theorem, first note that the map $\pi_1 : \mathfrak{MJ}^n \to \mathfrak{DJ}^n$ is bijective: given a discrete torsion free subgroup Γ of Isom($\mathbf{H}^n$), we obtain a complete hyperbolic manifold by taking $M = \mathbf{H}^n/\Gamma$ and e to be the image of s.

We now prove that $\pi_1 : \mathfrak{MJ}^n \to \mathfrak{DJ}^n$ is continuous. We want to show that if (M',e') is near (M,e), then $\pi_1(M', e')$ is near $\pi_1(M,e)$ in the Chabauty topology on the space of subgroups of Isom($\mathbf{H}^n$). By Lemma 3.2.6 *(Injectivity radius continuous)*, we may assume that the injectivity radius is greater than some $\delta > 0$ throughout the regions of interest in the proof we are about to present.

Suppose $\gamma \in \pi_1(M,e)$. Let $B_1, B_2, \ldots, B_k$ be a circular chain of balls in M, where $B_i = B(x_i, \delta)$, $x_1 = e_0$, $x_i \in \gamma$, and the balls cover γ. We assume that the interior of B_i meets the interior of B_{i+1}. Let $x_i \; R \; x'_i$. We obtain a chain $B'(x'_i, \delta)$. The holonomy corresponding to this chain is nearly equal to γ. We see this by taking $(n+1)$ generic points in int$B_i \cap$ intB_{i+1} and keeping track of the corresponding points in int$B'_i \cap$ intB'_{i+1}. So if $\gamma \in \pi_1(M, e)$, then γ is the limit of elements $\gamma' \in \pi_1(M', e')$.

We also need to show that if $\gamma(i) \in \pi_1(M(i), e(i))$, and $\gamma(i) \to \gamma$, then $\gamma \in \pi_1(M, e)$. Let s_0 be the standard basepoint in $\mathbf{H}^n$. Let $d(s_0, \gamma s_0) = r$. Then the geodesic from s_0 to $\gamma(i)s_0$ has length approximately equal to r, for large i. This gives an essential loop in $M(i)$, based at $e(i)_0$, of length approximately r. It follows that there is a geodesic loop in M, based at e_0, of length approximately equal to r, and the holonomy of this loop is very near to that of γ. Therefore $\gamma \in \pi_1(M,e)$.

This shows that if (M',e') is near (M,e) in the ϵ-relation topology, then, (by Lemma 3.1.3 *(Geometric convergence)*), $\pi_1(M',e')$ is near $\pi_1(M,e)$ as a subgroup of Isom($\mathbf{H}^n$) in the Chabauty topology.

Conversely, suppose Γ_i converges to Γ in the Chabauty topology. We must prove that $(\mathbf{H}^n/\Gamma_i, e(i))$ converges to $(\mathbf{H}^n/\Gamma, e)$ where $e(i)$ and e are the images of the standard point and frame s. Let $A_r =$

$\{g \in \mathrm{Isom}(\mathbf{H}^n) : B(s_0,r) \cap gB(s_0,r) \neq \emptyset\}$. Then A_r is compact. We fix r and let id, $\gamma_1, \ldots, \gamma_k$ be the elements of Γ in A_r. Then, for large values of i, there are elements id, $\gamma(i)_1, \ldots, \gamma(i)_k$ in Γ_i with $\gamma(i)_j$ near to γ_j. Moreover, for fixed $\epsilon > 0$ (where ϵ is significantly smaller than δ), each element in $\Gamma_i \cap A_{r-\epsilon}$ appears in this list. We can now deduce that a large compact region of $\mathbf{H}^n/\Gamma$ is almost isometric to a large compact region of $\mathbf{H}^n/\Gamma_i$. In fact the isometry is induced by a map which is C^∞-near the identity for large values of i. The proof is the same as that presented in 1.7.2 *(Holonomy induces structure)*.

Map to torsion free subgroups is a homeomorphism

3.2.10 Definition. Let (M_1, e_1) and (M_2, e_2) be two Riemannian manifolds with baseframe. Then a *framed (K,r)-approximate isometry* between (M_1, e_1) and (M_2, e_2) is a diffeomorphism $f:(X_1, e_1) \to (X_2, e_2)$ such that $B_{M_1}(x_1, r) \subseteq (X_1, x_1) \subseteq (M_1, x_1)$, $B_{M_2}(x_2, r) \subseteq (X_2, x_2) \subseteq (M_2, x_2)$, $Df(e_1) = e_2$ and

$$\frac{d(x,y)}{K} \leq d(f(x),f(y)) \leq K\,d(x,y) \text{ for all } x,y \in X_1 .$$

We may similarly define (K,r)-approximate isometries, and K-approximate isometries

There are many possible definitions of an approximate isometry. We have chosen a relatively strong one (as we require differentiability) and (ϵ, r)-relations may be thought of as the weakest possible notion of an approximate isometry. However, the proof of the above result tells us that these two definitions (and thus many other definitions "between" the two) are equivalent for complete hyperbolic manifolds.

3.2.11. Corollary: Approximate isometries. *The topology on $\mathcal{MJ}^n$ induced by framed (K, r)- approximate isometries is equivalent to the topology induced by framed (ϵ, r)-relations (and thus to the Chabauty topology).*

We now return (hopefully with new insight) to two topics which we discussed at the end of the last section. We can extend Corollary 3.1.8 *(Compact with bounded volume)* by verifying that the volume map is continuous for hyperbolic n-manifolds of finite volume($n \geq 3$). Recall that Theorem 3.1.4 *(A(U) compact)* only guarantees that this map is lower semicontinuous; recall also that the statement for two-manifolds, corresponding to the next theorem, is actually false (for example, a sequence of compact surfaces of genus two may converge to a punctured torus).

3.2.12. Theorem (Jorgensen): Volume is continuous. *The map* vol *from the space of hyperbolic n-manifolds($n \geq 3$) of finite volume with baseframe to* **R** *which takes each manifold to its volume is continuous.*

Proof. Suppose (M_i, e_i) converges to (M,e). Notice that for hyperbolic n-manifolds ($n \geq 3$) the thick part is connected, since the boundary of each component of the thin part is connected. Also, we may assume that $\text{inj}_{M_i}(e_i) > \epsilon$ for some $\epsilon > 0$ and that $\text{vol}(M_i) < V$ for some V and all i. There exists a bound on the diameter of $M_{[\epsilon,\infty)}$ in terms of V and ϵ for all $\epsilon > 0$, so $M_{i,[\epsilon,\infty)}$ converges to $M_{[\epsilon,\infty)}$ using only the topology induced by K-approximate isometries (since they are all compact and of uniformly bounded diameter). Thus, $\{\text{vol}(M_{i,[\epsilon,\infty)}\}$ converges to $\text{vol}(M_{[\epsilon,\infty)})$ for all $\epsilon > 0$. But since the volume of the thin part converges to 0 as ϵ converges to 0, we see that $\{\text{vol}(M_i)\}$ converges to $\text{vol}(M)$

□

When we apply the above analysis to the two-dimensional case we obtain:

3.2.13. Proposition: Mumford's Lemma . *The subset* $\mathfrak{M}(T)_\delta$ *of* $\mathfrak{MW}^n$ *consisting of all finite area surfaces homeomorphic to a given surface T with no closed geodesics shorter than some* $\delta > 0$ *is compact, for any* $\delta > 0$.

Proof. Simply notice that if $S \in \mathfrak{M}(T)_\delta$, $S_{[\delta',\infty)}$ is connected for all $\delta' < \delta$. Then the argument above proves that $\mathfrak{M}(T)_\delta$ is a closed subset of $\mathfrak{MW}^2(A)$ where $A = \text{area}(S)$. But since $\mathfrak{MW}^2(A)$ is compact by Corollary 3.1.8 *(Compact with bounded volume)* so is $\mathfrak{M}(T)_\delta$.

□

3.2.14 Remarks. Thurston has further proved that the set of volumes of complete hyperbolic 3-manifolds form a closed, non-discrete set in $\mathbf{R}^+_\infty$. This set is well-ordered and has ordinal type ω^ω. For a detailed discussion of this result see 6.6 of [T] or [Gromov, 1980]. Thurston (see Section 6.6 of [T]), has also shown that there are finitely many complete hyperbolic manifolds of any fixed volume, but in [Wielenberg, 1981] it is shown that there is no bound on the number of complete hyperbolic manifolds of a given volume, by proving that there exist fundamental polyhedra with arbitrarily many associated non-conjugate Kleinian groups. In [Wang, 1972] it is shown that there are only finitely many hyperbolic manifolds with volume less than any given real number.

3.2.15 We finish this section by returning briefly to the subject of marked hyperbolic surfaces. We may also topologize the Teichmüller space of a surface of finite type using approximate isometries. We may define the K-neighbourhood of a marked surface $[h:T \to S]$ to be the set of all marked surfaces $[h':T \to S']$ such that there exist representatives h and h' for $[h]$ and $[h']$ and a diffeomorphism $\phi:S \to S'$ which is a K-approximate isometry when restricted to the respective convex cores and such that $\phi \circ h$ is isotopic to h'. Such neighbourhoods form a basis for the topology of Teichmüller space; it is left to the reader to to satisfy himself that this topology agrees with the topology defined in the last section. Equivalently, one may also topologize Teichmüller space using ϵ-relations, but approximate isometries are more frequently used.

Chapter 4. Laminations

4.1. Geodesic Laminations

For a more detailed treatment see [Casson, 1983], or [Harer-Penner, 1986], or Chapter 8 of [T].

4.1.1 Definition. Let S be a connected complete hyperbolic surface. Then a *geodesic lamination* on S is a closed subset λ of S which is a disjoint union of simple geodesics of S (which are called *leaves* of the lamination)

Remark: We allow the empty set as a geodesic lamination.

We denote the set of all geodesic laminations on S by $\mathcal{GL}(S)$.

4.1.2 Definition. If λ is a lamination on S, then a component of $S - \lambda$ is called a *flat piece* or a *complementary region*.

In general these need not be simply connected and may have a finite or infinite number of sides. The leaves of the geodesic lamination which form the boundary of some complementary region are called *boundary leaves.* For a surface of finite area there is an upper bound on the number of complementary regions, since each has area $n\pi$ for some positive integer n. Moreover each complementary region has finite type. On a surface with finite area, the set of boundary leaves is dense in the geodesic lamination.

4.1.3 Definition. A lamination such that each complementary region is isometric to an ideal triangle is said to be *maximal.*

We shall see, from Theorem 4.2.8 *(Structure of lamination)* that any lamination on a surface of finite area can be extended by adding a finite number of new leaves to obtain a maximal lamination.

On a surface of finite area, any geodesic lamination has measure zero in the surface. Equivalently any C^1-curve C transverse to the geodesic lamination intersects it in a set of measure zero in C. Since S is a complete hyperbolic surface, its universal cover is $\mathbf{H}^2$. We may lift a lamination λ to $\mathbf{H}^2$ to obtain a lamination $\tilde{\lambda}$ on $\mathbf{H}^2$, which is invariant under the action of the covering transformations (which are elements of some Fuchsian group Γ).

Consider the closed unit disk $\mathbf{B}^2$ as the compactification of $\mathbf{H}^2$. The space of geodesics is homeomorphic to an open Möbius band M. (See [Epstein-Marden].)

4.1.4 Transferring laminations. Given a geodesic lamination λ_1 on a complete hyperbolic surface S_1 of finite area, and a homeomorphism ϕ onto a complete hyperbolic surface S_2 of finite area, the lamination λ_1 can be canonically transferred to a geodesic lamination λ_2 of S_2. The reason is that the homeomorphism can be lifted to an equivariant map from $\mathbf{H}^2$ to $\mathbf{H}^2$ which extends to a homeomorphism between boundary circles. Since a geodesic is an unordered pair of elements of $\partial\mathbf{H}^2$, we see how to transfer a geodesic in S_1 to a geodesic in S_2. This transfer induces a canonical homeomorphism $\phi\#: GL(S_1) \to GL(S_2)$, which only depends on the isotopy class of ϕ.

4.1.5 Definition. The Chabauty Topology is the topology induced on $\mathcal{GL}(S)$ as a subspace of $C(M)$, the set of closed subsets of the open Möbius band M with the Chabauty topology (see Section 3.1 *(The Geometric Topology)* for the definition of the Chabauty topology).

If S has finite area, a lamination on S is determined by its underlying point set (see Proposition 4.1.6 *(Lamination determined by pointset)*). We can therefore also topologize $\mathcal{GL}(S)$ using the Chabauty topology on subsets of $\mathbf{H}^2$. We shall also see that these two topologies are homeomorphic.

If ϵ is small, $S_{(0,\epsilon)}$ consists entirely of cusps. All simple geodesics entering a cusp are asymptotic to each other; they are orthogonal to the horocycles. Therefore a lamination on S is completely determined by its intersection with $S_{[\epsilon,\infty)}$.

4.1.6. Proposition: Lamination determined by pointset. *Let S be a complete hyperbolic surface of finite area. Let $C(S)$ be the space of closed subsets of S with the Chabauty topology and let $L \subset C(S)$ be defined by*

$$L = \{X : X = |\lambda| \text{ for some lamination } \lambda \subset S\}.$$

Then L is a closed subset of $C(S)$ and the map $\mathcal{GL}(S) \to L$ defined by $\lambda \mapsto |\lambda|$ is a homeomorphism. In particular, the topologies induced by regarding GL(S) as a subset of C(M) and C(S) agree.

Proof. We need only show that $\mathcal{GL}(S) \to L$ is continuous and injective, since a continuous injective map onto a Hausdorff space is a homeomorphism. Geodesic laminations on a surface of finite area are nowhere dense, so no two laminations can share the same underlying pointset. (We note that this is not true for surfaces of infinite area.) Thus, our map is injective.

So suppose that $\{\lambda_i\}$ converges to λ in $\mathcal{GL}(S)$. Then, given a point $x \in |\lambda|$, x lies on some geodesic l, and l is the limit of geodesics $l_i \in \lambda_i$ as we see by looking in the universal cover of S. We choose $x_i \in l_i$, such that x_i converges to x. This is the first condition for the convergence of $|\lambda_i|$ to $|\lambda|$. To prove the second condition, we suppose that $x_i \in |\lambda_i|$ and that x_i converges to some point x. We must show that $x \in |\lambda|$. We have $x_i \in l_i \in \lambda_i$. By lifting to the universal cover, we see that we may assume l_i converges to $l \in \lambda$, and that $x \in l$. Hence $x \in |\lambda|$. This completes the proof of the proposition.

∎

4.1.7. Proposition: GL(S) compact. *$\mathcal{GL}(S)$ is compact, metrizable and separable in the Chabauty topology.*

Note that this is true even if S is not compact.

Proof. By Proposition 3.1.2 *(Properties of Chabauty topology)* it is sufficient to show that $\mathcal{GL}(S)$ is a closed subset of $C(M)$, where M is the open Möbius band.

Suppose we have a sequence of geodesic laminations $\{\lambda_i\} \subset \mathcal{GL}(S)$ with limit $\lambda \in C(M)$. By Lemma 3.1.3 *(Geometric convergence)* this means that every geodesic of λ is the limit of a sequence of geodesics, one in each λ_i, and that every convergent sequence of geodesics, one in each λ_i, converges to a geodesic of λ.

Let $S = \mathbf{H}^2/\Gamma$. We need to check two things: that λ is a *disjoint* union of geodesics, and that λ is invariant under Γ. Suppose l and k are two leaves in λ and that l_i converges to l and k_i converges to k where $l_i, k_i \in \lambda_i$. Since l_i and k_i are disjoint or equal for each i, the same is true of l and k. To check that λ is Γ-invariant, recall that if a group acts on a space, then, for any subset X which is pointwise fixed under Γ, its closure $\bar{X}$ is also fixed under Γ. Here the space is $C(M)$ and the subset X is the countable set $\{\lambda_i\} \subset C(M)$. So λ is a lamination and thus $\mathcal{GL}(S)$ is closed. ∎

The following lemma results immediately by applying Lemma 3.1.3 *(Geometric convergence)*.

4.1.8. Lemma: Geometric convergence for laminations. *If (l_i) and (k_i) are two sequences of geodesic laminations converging to l and k respectively, with $l_i \subset k_i$ for all i, then $l \subset k$.*

We set $\mathcal{GLM}(S)$ to be the set of maximal laminations on S. Then

4.1.9. Lemma: GLM closed in GL. *Let S have finite area. Then $\mathcal{GLM}(S)$ is a closed subset of $\mathcal{GL}(S)$.*

Proof. We shall work in $C(S)$, the set of closed subsets of S.

Suppose λ is not maximal, and $\lambda_i \to \lambda$, with each λ_i maximal. Let P be a complementary region for λ which is not an ideal triangle. First we show that P can have no simple closed geodesics in its interior. For if K is the underlying pointset of such a geodesic, then $\lambda \in O_1(K)$ and so $\lambda_i \in O_1(K)$ for i large. But this means that there is a simple closed geodesic in a complementary region of λ_i which is forbidden. Hence P is a finite sided polygon.

It is easy to construct a connected compact subset K in the interior of P such that each geodesic in P meets the interior of K and such that $\lambda \in O_1(K)$. If i is large, then λ_i contains geodesics near the boundary geodesics of P, and $\lambda_i \in O_1(K)$. Moreover none of the geodesics of λ_i meet K. Let P_i be the complementary region of λ_i in which K lies. We can see that among the boundary geodesics of P_i, there are geodesics very near to each of the boundary geodesics of P. Therefore some leaf of λ_i must pass through K since λ_i is maximal. This contradiction proves the result.

$\square$

We now discuss another topology on $\mathcal{GL}(S)$, which we call the *Thurston topology*. The reference for this is Section 8.10 of [T], where it is referred to as the geometric topology.

4.1.10 Definition. The *Thurston topology* on $\mathcal{GL}(S)$ is the topology induced by once again treating $\mathcal{GL}(S)$ as a subset of $C(S)$, but with the topology generated by sub-basis elements of the form $O_2(V)=\{A \in C(S) | A \cap V \neq \emptyset\}$, where V is open in S.

So we see that a neighbourhood of a geodesic lamination λ contains all geodesic laminations λ' such that $\lambda \supset \lambda'$.

Note that this topology is strictly weaker than the Chabauty topology and is non-Hausdorff, but it is more closely related to the consideration of pleated surfaces. For example, a sequence of surfaces, bent along a single geodesic converges to a geodesic surface if the bending angle converges to zero. The sequence of laminations corresponding to the surface converge to the lamination corresponding to the limit surface in the Thurston topology, but not in the Chabauty topology.

The analogue of Lemma 3.1.3 *(Geometric convergence)* is

4.1.11. Lemma: Geometric convergence in Thurston topology. *If (λ_i) converges to λ in the geometric topology then, given any geodesic $l \in \lambda$, there is a sequence of geodesics (l_i), where $l_i \in \lambda_i$, which converges to l.*

4.2. Minimal Laminations

See Section 8.10 of [T].

4.2.1 Definition. A non-empty geodesic lamination λ is said to be *minimal* if no proper subset of λ is a geodesic lamination.

A single geodesic is a minimal lamination if and only if it is closed subspace of the surface (either a simple closed geodesic, or an infinite geodesic, with each end converging to a cusp).

4.2.2. Lemma: One or uncountable. *Let λ be a minimal lamination on S. Then either*

1) *λ consists of a single geodesic, or*

2) *λ is uncountable.*

Proof. Let l be a leaf of λ. If l is isolated, then $\lambda \setminus l$ is again a lamination. Since λ is minimal, $\lambda \setminus l = \varnothing$, so that Condition 1) holds. So we assume that l is not isolated. Take some point x on l and a small transverse line L through x. Since l is not isolated, x must be an accumulation point of $L \cap \lambda$ in L. But since $L \cap \lambda$ is closed and each point is an accumulation point, it is a perfect set and is thus uncountable. Since each leaf of λ intersects L at most a countable number of times, λ must have uncountably many leaves.

□

We now discuss hyperbolic surfaces of finite area in order to discover precisely the structure of any lamination on such a surface

Let S be a complete hyperbolic surface without boundary, and let λ be a lamination on S. We cut S along λ. The formal definition of this process is to take $S \setminus \lambda$ with its Riemannian metric, and complete it. We obtain a complete (possibly disconnected) hyperbolic surface with geodesic boundary.

4.2.3. Lemma: Building a surface. *Let $\Delta_1, \ldots, \Delta_k$ be a finite set of ideal triangles. Let $(A_1, B_1), \ldots, (A_r, B_r)$ be r pairs made up from $2r$ distinct edges, chosen from the $3k$ edges of $\Delta_1, \ldots, \Delta_k$. Let $h_i : A_i \to B_i$ $(1 \le i \le r)$ be an isometry. Identifying using the h_i, we obtain a hyperbolic surface, possibly with boundary. The completion of this surface consists of adding a finite, possibly empty, collection of simple closed boundary geodesics.*

Proof. Each end E of an identified edge gives rise to two other ends E' and E'' of identified edges, namely those that occur on either side of E, see Figure 4.2.4 . Possibly $E' = E''$, and possibly $E = E' = E''$. (The last case occurs when, say, two edges of Δ_1 are identified with each other with a certain choice of orientation.) Here E and E' have a common ideal vertex in one of the ideal triangles $\Delta_1, \ldots, \Delta_k$, and similarly for E and E''. Thus, the ends of identified edges can be arranged in cycles $(E_1, \ldots, E_s)$, where E_i and E_{i+1} have a common ideal vertex (interpreting E_{s+1} as E_1).

We can explicitly work out the local geometry of the piece of the surface arising from such a cycle. In the upper half plane model, it is obtained

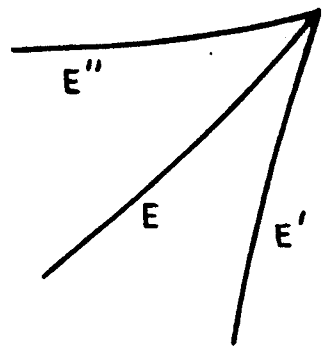

4.2.4 Figure.

from the strip

$$\{(x,y) \mid 0 \leq x \leq 1\,,\ y \geq k\}$$

modulo the gluing map $z \mapsto z+1$, or from the strip

$$\{(x,y) \mid 1 \leq x \leq a\,,\ y \geq kx\}$$

modulo the gluing map $z \mapsto az$. The first case gives us a cusp. The second case is isometric to the sector $\{(x,y) \mid y \geq kx\,,\ x > 0\}$ modulo $z \mapsto az$. The completion gives us the geodesic $\{x=0\,,\ 1 \leq y \leq a\}/\{z \mapsto az\}$, and our piece of surface is a nice neighbourhood of this geodesic.

∎

4.2.5. Corollary: Bound on boundary components. *Using the same notation as in the preceding lemma, the number of boundary components is at most* $3k-r$.

Proof. The $2r$ edges $\{A_1, B_1, \ldots, A_r, B_r\}$ give rise to r edges in the surface. So there are at most r cycles of edges of the type described in Lemma 4.2.3 *(Building a surface)*. Hence there are at most r new boundary components in the completed surface. It follows that the glued up and completed surface has at most $3k-2r+r = 3k-r$ boundary components.

∎

4.2.6. Lemma: Non-compact surfaces obtainable. *Every finite area complete hyperbolic surface S with geodesic boundary can be constructed (non-uniquely) by the above process, except for a compact surface without boundary. In the case of a compact surface without boundary we start by cutting the surface along a simple closed geodesic, and then cut what remains into triangles. If S is non-compact and ∂S has only non-compact* *components, then the completion step, after gluing the triangles together, can be avoided by cutting up into triangles correctly.*

Proof. By doubling S we obtain a complete finite area surface without boundary. It follows that S has only a finite number of boundary components. Also it is easy to see what S must look like topologically.

If we have a cusp or an ideal vertex x in S, we may cut a finite number of times along geodesics from x to x. We eventually get a disjoint union of surfaces, each of which is an ideal polygon or an annulus. An ideal polygon may be cut into triangles. In the annulus case, we may assume that one boundary is a geodesic from x to x and the other a geodesic circle or a cycle of non-compact oriented geodesics $B_1, \ldots, B_k$, with the positive end of one geodesic asymptotic to the negative end of the next. In the annulus case we can make the surface simply connected by cutting along a geodesic from x to an end point of one of the B_i, or by cutting along a geodesic from x which spirals around the simple closed geodesic at the other end of the annulus. This deals with the non-compact case. If S is compact with boundary, we reduce to the previous case by cutting along a geodesic which spirals to the boundary at each end. If S has no boundary, we cut first along a simple closed geodesic.

$\square$

4.2.7. Corollary: Another bound on boundary components. *Let S be a complete hyperbolic surface of finite area A with geodesic boundary. Let b be the number of boundary components. Then $\pi b \leq 3A$.*

Proof. Let S be obtained by gluing together k ideal triangles and completing. Then

$$\pi b = \pi(3k - r) = 3A - r\pi$$

using Corollary 4.2.5 *(Bound on boundary components).*

$\square$

4.2.8. Theorem: Structure of lamination. *Let λ be a lamination on a complete hyperbolic surface of finite area with geodesic boundary. Then λ consists of the disjoint union of a finite set of minimal sublaminations of λ together with a finite set of additional geodesics, each end of which either "spirals" onto a minimal lamination or goes up a cusp. Each of the additional geodesics is isolated – it is contained in an open subset which is disjoint from the rest of the lamination. Each cusp contains only a finite* *number of geodesics of* λ.

Proof. Structure of lamination: First note that minimal sublaminations exist. On a compact surface, this is a consequence of compactness in the usual way. On a non-compact surface, we take the intersection with the thick part of the surface and argue there. Recall that if we cut along any lamination we obtain a finite number of components, since each component has area $n\pi$ for some integer n.

If we cut S along a sublamination λ_1 of λ we obtain a new surface S' with a new lamination λ' obtained from λ in the obvious way. Let λ_1 be a minimal sublamination of λ. We claim that the minimal sublaminations of λ correspond one-to-one with the minimal sublaminations of λ', except that λ_1 itself disappears and is replaced by one or more new boundary leaves of S', each of which is minimal in λ'. (These boundary leaves form the set of points added in the process of completion.) Each boundary leaf is of course minimal. If $\lambda_2 \neq \lambda_1$ and λ_2 is minimal in λ, it is clearly minimal in λ'. A minimal sublamination λ_3 of λ' is either a boundary component of S' or is disjoint from $\partial S'$. In the second case λ_3 is a minimal sublamination of λ. In the first case, λ_3 is either a component of ∂S, in which case it is minimal in λ, or a new component of $\partial S'$. So the claim is established.

To prove the theorem, we cut successively along minimal sublaminations which are not boundary components. By our claim, each such minimal lamination is minimal in λ. Each cut increases the number of boundary components. Since this number is bounded by Corollary 4.2.5 *(Bound on boundary components)*, the process ends. This shows there are only a finite number of minimal sublaminations.

To complete the proof of the theorem, we need a lemma.

... **4.2.9. Lemma: L finite.** *If all minimal sublaminations are contained in the boundary,* λ *is finite.*

Proof. L finite: We define a *corner* of S to be the cusp-like region lying between two asymptotic boundary components of S. A *corner* of a complementary region is defined similarly, by completing the complementary region. If we double S, then the double of a corner is a cusp. First note that each corner of S and each cusp of S can contain only a finite number of geodesics of λ, since the area bounds both the number of complementary regions and the number of corners in each complementary region. Let K be the compact subspace of S obtained by cutting off the corners and cusps. (K is the intersection of S with the thick part of the double of S.) Every simple geodesic in S meets K.

Suppose that λ is infinite. Then there is a sequence of disjoint geodesics in λ converging to a a geodesic $l \in \lambda$. By what has already been shown, neither end of l can enter a cusp or corner. Hence l lies in K. So the closure of l is compact.

We claim that $\bar{l} \cap \partial S = \emptyset$. To see this note that if a component B of ∂S is contained in $\bar{l}$, then B is a circle since $\bar{l}$ is compact. Each leaf of λ which is near B must spiral onto B. Since there is a bounded number of corners in each complementary region of λ, and each spiral gives rise to a corner, there are only a finite number of leaves spiraling onto B. Since l is a limit of an infinite sequence of distinct geodesics of λ, l cannot spiral onto B. Therefore $\bar{l}$ contains a minimal sublamination disjoint from ∂S. This contradicts our hypothesis.

L finite

This completes the proof of the theorem

Structure of lamination

4.2.10. Theorem: Curve near geodesic. *Let $\epsilon > 0$ be fixed. Then there is a $\delta > 0$ with the following property. Let α be a piecewise geodesic curve parametrized by arc length in $\mathbf{H}^n$, whose pieces have length at least ϵ and such that pieces with a common endpoint meet at an angle greater than $\pi - \delta$. Let $\beta(t)$ be the geodesic parametrized by arc length, joining the endpoints of α. Then $d(\alpha(t), \beta(t)) < \epsilon$ for all t.*

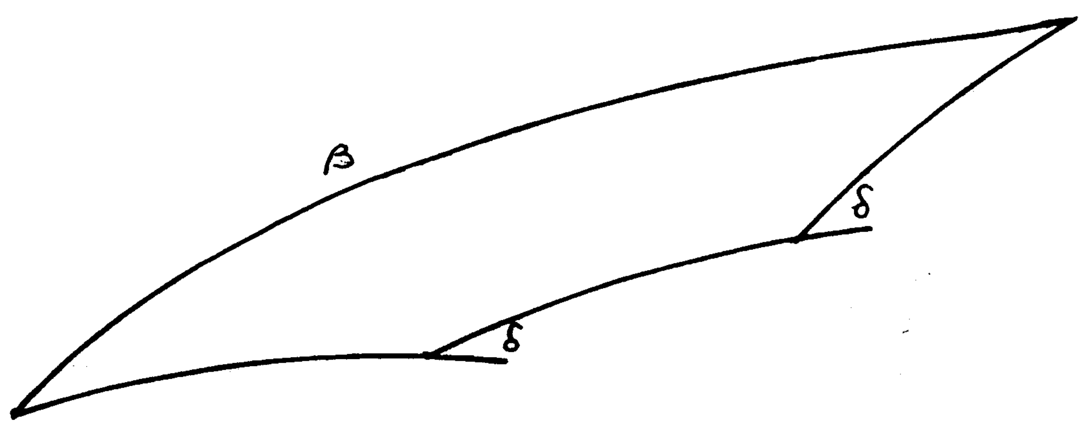

4.2.11 Figure.

Proof. Curve near geodesic: First we need a lemma which determines the situation in a triangle.

... **4.2.12. Lemma: Angle derivative.** *Let T be a fixed point and let Y move along a fixed geodesic containing a fixed point X. Set $t=d(X,Y)$ and $x=d(Y,T)$. Set $\theta=\angle XYT$. We regard x and θ as functions of t. Then*

$$\theta' = \frac{-\sin\theta}{\tanh x} .$$

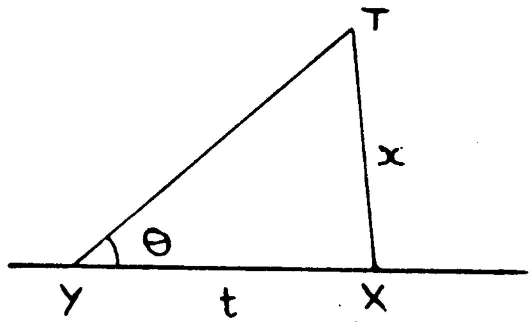

4.2.13 Figure.

Proof. Angle derivative: From the hyperbolic law of sines we see that $\sinh x \sin\theta$ is constant. Also $x' = \cos\theta$. The result follows.

Angle derivative

Continuation, proof of Curve near geodesic: Let $\theta(t)$ be the angle between the geodesic from $\alpha(0)$ to $\alpha(t)$ and the geodesic subarc on which $\alpha(t)$ lies. At the bends in α, θ is discontinuous, jumping by less than δ. Over the first segment of α, $\theta=\theta'=0$. Let $0=t_0<t_1<\ldots<t_k$ be the points at which α is not geodesic. Then $t_{i+1}-t_i > \epsilon$ for each i. We claim that $\theta(t)\leq\epsilon$ for all t. If not, let (t_i, t_{i+1}) be the first interval on which we do not have $\theta(t)\leq\epsilon$. On (t_0,t_1), $\theta=0$ so that $i\geq 1$. Since θ is monotonic decreasing on each subinterval, the limit of $\theta(t)$ as t decreases to t_i is greater than ϵ, and $\theta(t)>\epsilon-\delta$ on (t_{i-1},t_i). From Lemma 4.2.12 *(Angle derivative)* we see that, on (t_{i-1}, t_i),

$$\theta' < -\sin(\epsilon-\delta)$$

so that

$$\theta(t_{i-1}) > (\epsilon-\delta) + \epsilon\sin(\epsilon-\delta) .$$

If we choose δ very small, then the right hand side is bigger than ϵ, which is a contradiction, proving the claim.

It now follows that, by taking δ small, we can make $\theta(t)$ uniformly small. This ensures that dx/dt is near 1, where $x=d(\alpha(0), \alpha(t))$, so that $d(\alpha(0), \alpha(t))$ is uniformly near t.

Curve near geodesic

The following is a result which is frequently used.

4.2.14. Theorem: Finite laminations dense. *On any fixed hyperbolic surface of finite area, the finite laminations are dense in the space of all geodesic laminations, with the Chabauty topology. Hence the same is true with the Thurston topology.*

Proof. Finite laminations dense: We recall (from Theorem 4.2.8 *(Structure of lamination)*), that any lamination is the union of a finite set of disjoint minimal laminations $\mu_1, \ldots, \mu_k$, and a finite set of isolated leaves. We first approximate the minimal laminations.

... 4.2.15. Lemma: Approximating minimal laminations. *Let μ be a minimal lamination on a complete hyperbolic surface. Then μ can be approximated in the Chabauty topology by a simple closed geodesic.*

Proof. Approximating minimal laminations: We fix a point $p \in |\mu|$ and consider a geodesic arc in μ starting at p. Given $\epsilon > 0$, we can get within ϵ of each point of $|\mu|$, by taking the length of the arc long enough. We then expand this sideways to a strip P, of parallel geodesic arcs in μ of the same length. We assume that this strip has width less than ϵ at any point along its length.

Let γ be the complete geodesic of μ through p. We orient γ and this orients each arc of γ which lies in the strip P. There are two possible situations. The first is that there are two arcs α_1 and α_2 of the strip, lying in γ,oriented in the same direction along the strip, which are not separated in the strip by an arc α_3 of γ, lying between α_1 and α_2 as we travel along γ. In that case we get a simple closed curve, as shown in Figure 4.2.16 . The second possibility is that the first possibility does not occur. In that case the first three arcs of γ in which γ meets the strip must be arranged as in Figure 4.2.17 . In each case the simple closed curve is essential and the homotopic geodesic approximates μ. To see that the homotopic geodesic approximates μ, we change the transverse arc of the construction to an arc which is almost parallel to the arcs of the strip (see Figure 4.2.18). We then apply Theorem 4.2.10 *(Curve near geodesic)*.

Approximating minimal laminations

Continuation, proof of Finite laminations dense: To complete the proof that any lamination can be approximated by a finite lamination, we approximate each minimal sublamination μ_j by a simple closed geodesic C_j $(1 \leq j \leq k)$. Then $S \setminus C_1 \cup \ldots \cup C_k$ is nearly the same as $S \setminus \mu_1 \cup \ldots \cup \mu_k$. The difference is that complementary regions of $S \setminus \mu_1 \cup \ldots \cup \mu_k$ become

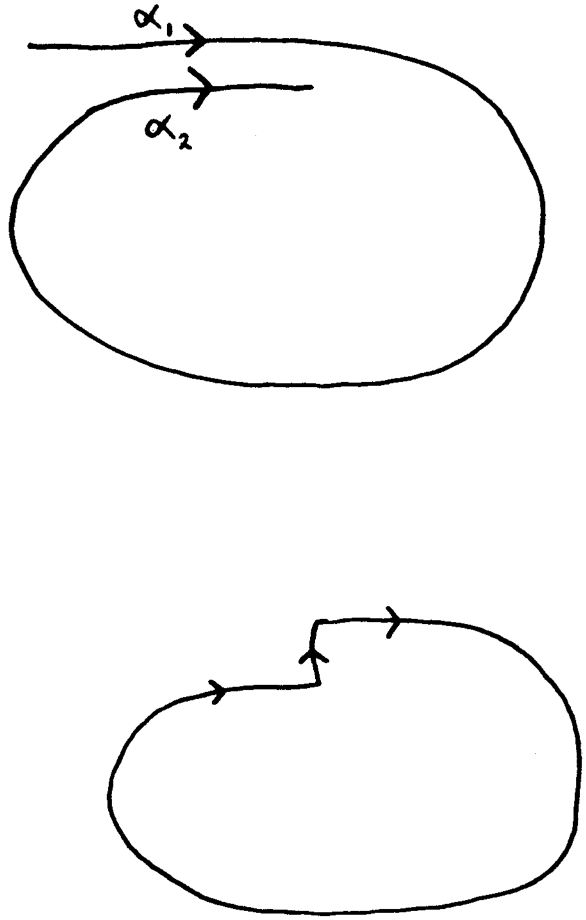

4.2.16 Figure. *Construction of the simple closed curve when the orientations agree.*

joined up in $S\backslash C_1 \cup \ldots \cup C_k$ through thin gaps between geodesics which were closed off in $S\backslash \mu_1 \cup \ldots \cup \mu_k$. It is easy to make sure that the geodesics C_i do not intersect each other since each minimal lamination intersects the convex core in a compact set, and these compact sets are then separated by fixed distances.

Geodesics in $S\backslash C_1 \cup \ldots \cup C_k$ corresponding to the isolated geodesics of $S\backslash \mu_1 \cup \ldots \cup \mu_k$ can now be drawn in. In the complement of a finite lamination they will extend through the new gaps and spiral around one of the geodesics C_j.

Finite laminations dense

We have the following corollary;

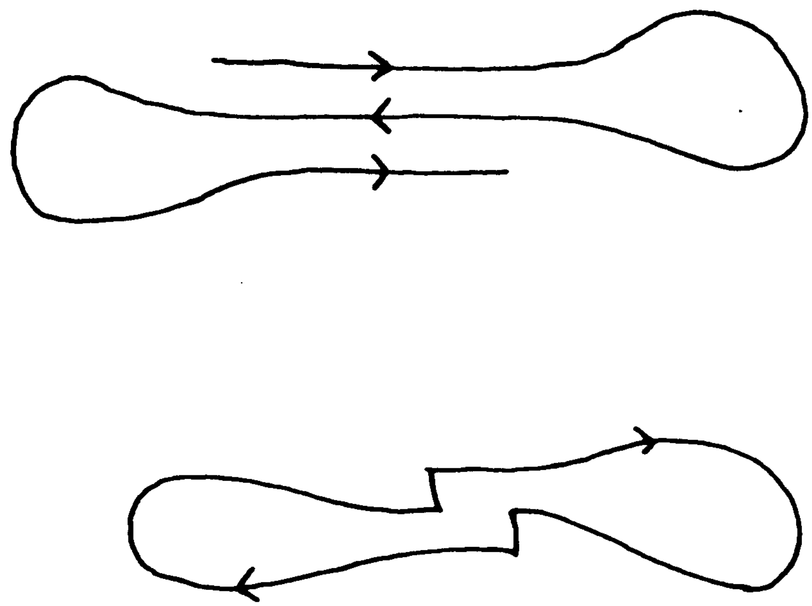

4.2.17 Figure. *Construction of the simple closed curve when the orientations do not agree.*

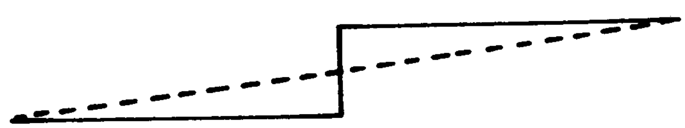

4.2.18 Figure.

4.2.19. Corollary: Finite laminations dense in GLM. *The set of finite maximal laminations in $\mathcal{GLM}(S)$ is dense*

Proof. Suppose that λ is maximal, then by the theorem above there is a sequence (λ_i) converging to λ, where each of the λ_i finite. Each λ_i can be extended to a finite maximal lamination $\hat{\lambda}_i$. Then $(\hat{\lambda}_i)$ converges to a lamination $\hat{\lambda}$ which contains λ. But since λ is maximal, $\hat{\lambda} = \lambda$ and so we have a sequence $(\hat{\lambda}_i)$ converging to λ as required. ∎

Chapter 5. Pleated Surfaces

5.1. Introduction

We now discuss pleated surfaces, which are a basic tool in Thurston's analysis of hyperbolic structures on three-manifolds. See Section 8.8 of [T]. In [T], pleated surfaces are called uncrumpled surfaces.

Recall from Definition 5.2.8 *(Isometric map)* that an isometric map takes rectifiable paths to rectifiable paths of the same length.

5.1.1 Definition. A map $f : M \to N$ from a manifold M to a second manifold N is said to be *homotopically incompressible* if the induced map $f_*:\pi_1(S) \to \pi_1(M)$ is injective.

5.1.2 Definition. A *pleated surface* in a hyperbolic three-manifold M is a complete hyperbolic surface S together with an isometric map $f:S \to M$ such that every point $s \in S$ is in the interior of some geodesic arc which is mapped by f to a geodesic arc in M. We shall also require that f be homotopically incompressible.

Note that this definition implies that a pleated surface f maps cusps to cusps since horocyclic loops on S are arbitrarily short and f is isometric and homotopically incompressible.

5.1.3 Definition. If (S,f) is a pleated surface, then we define its *pleating locus* to be those points of S contained in the interior of one and only one geodesic arc which is mapped by f to a geodesic arc.

An example of a pleated surface is the boundary of the convex core. (See [Epstein-Marden].)

5.1.4. Lemma: Pleating locus is a lamination. *Let (S,f) be a pleated surface. Then the pleating locus of (S,f) is a geodesic lamination and the map f is totally geodesic in the complement of the pleating locus.*

Proof. We need only consider pleated maps from $\mathbf{H}^2$ to $\mathbf{H}^3$ since we can always work in the universal covers. Let γ be the pleating locus of a pleated surface $f:\mathbf{H}^2 \to \mathbf{H}^3$. If $x \notin \gamma$, then there are two transverse geodesic arcs through x. Let $[a\ x\ b]$ and $[c\ x\ d]$ be geodesic arcs in $\mathbf{H}^2$, mapped by f to geodesic arcs $[a'x'b']$ and $[c'x'd']$ in $\mathbf{H}^3$. Let $\angle a\ x\ c = \theta$.

Then $\angle b\,x\,c = \pi - \theta$. Since f is isometric, $d(a', c') \leq d(a,c)$. Hence $\theta' = \angle a'\,x'\,c' \leq \theta$. Similarly $\angle b'\,x'\,c' \leq \angle b\,x\,c$, so that $\pi - \theta' \leq \pi - \theta$. Hence $\theta' = \theta$. It now follows that $d(a',c') = d(a,c)$ and that f maps the geodesic $[a,c]$ to the geodesic $[a',c']$. This implies that f is totally geodesic in a neighbourhood of x. In particular we see that the pleating locus is a closed set.

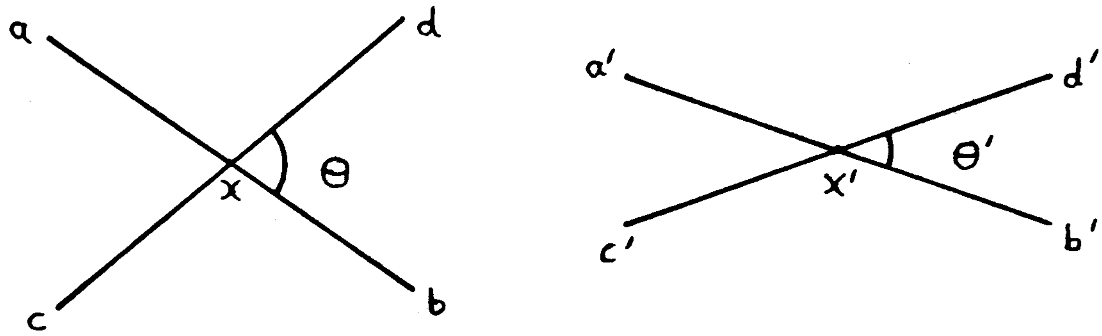

5.1.5 Figure. *Diagram showing two transverse arcs which are mapped by f to geodesics.*

Now let $x \in \gamma$, the pleating locus of f, and let x lie in the interior of the open arc α, such that $f|\alpha$ is geodesic. We take α maximal with this property and we claim that α is a complete geodesic. For if not, let y be a finite endpoint of α and let z be a point on α, on the other side of x from y. Let $[a\ y\ b]$ be a geodesic arc (a,y,b distinct), on which f is geodesic. Then $[a\ y\ b]$ meets the geodesic containing α only in y, by the maximality of α.

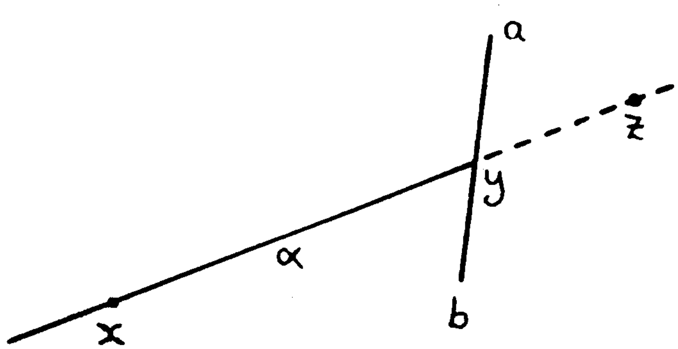

5.1.6 Figure.

Since f is isometric, $d(fa,\ fz) \leq d(a,z)$ and so $\angle fa\ fy\ fz \leq \angle a\ y\ z$. Similarly $\angle fb\ fy\ fz \leq \angle b\ y\ z$. Since these angles add up to π, we must have equality in each case. Hence f maps the geodesic $[a,z]$ to the geodesic $[fa,fz]$. We deduce that f is geodesic on the triangle $z\ a\ b$. But this contradicts the fact that $x \in \gamma$. The contradiction shows that α is a complete geodesic.

The same argument, with y defined as any point of α which is not on the pleating locus, shows that $\alpha \subset \gamma$.

$\square$

5.2. Compactness Properties of Pleated Surfaces

See Section 8.8 of [T].

We now wish to consider the space of all possible pleated surfaces (in all possible 3-manifolds), and derive the various compactness results which will be used later in our paper and are also used elsewhere in Thurston's work.

Let $\mathcal{PSF}$ be the set of *P*leated *S*urfaces with base*F*rame. More precisely, $\mathcal{PSF}$ is the set of triples (Γ_1, Γ_2, f), where Γ_1 is a torsion free discrete subgroup of $\mathrm{Isom}^+\mathbf{H}^2$, Γ_2 is a torsion free discrete subgroup of $\mathrm{Isom}^+\mathbf{H}^3$, and f is a pleated map such that there exists a homomorphism $h: \Gamma_1 \to \Gamma_2$ with

1) $f(s_0) = s_0$ (recall that s_0 is a fixed basepoint in $\mathbf{H}^2$ and we are thinking of $\mathbf{H}^2$ as embedded in $\mathbf{H}^3$);
2) $f \circ T = h(T) \circ f$ for $T \in \Gamma_1$;

One often assumes also that h is injective, but if we assume this, we will state it explicitly.

We topologize $\mathcal{PSF}$ by using the Chabauty topology for the groups Γ_1 and Γ_2 and the compact open topology for f.

Let S be a stratum for the pleating locus. Then $f|S = A(S)|S$ for some $A(S) \in \mathrm{Isom}(\mathbf{H}^3)$. (Since $\mathbf{H}^2$ is embedded in $\mathbf{H}^3$, the isometries of $\mathbf{H}^3$ map $\mathbf{H}^2$ into $\mathbf{H}^3$.) Let $T \in \Gamma_1$. Then

$$\begin{aligned} f \,|\, TS &= (fT)|S \circ T^{-1}|TS \\ &= (h(T)f)|S \circ T^{-1}|TS \\ &= h(T)A(S)T^{-1}|TS \,. \end{aligned}$$

Hence TS is contained in some stratum of the pleating locus. Applying the same argument to T^{-1}, we see that Γ_1 preserves the pleating locus.

Note that there is an induced basepoint preserving map $f: \mathbf{H}^2/\Gamma_1 \to \mathbf{H}^2/\Gamma_2$ and that h is the induced map of fundamental groups.

5.2.1 Definition. We define $\mathcal{PSJ}(A,\epsilon)$ to be the set of triples $(\Gamma_1, \Gamma_2, f) \in \mathcal{PSJ}$, such that Area$(\mathbf{H}^2/\Gamma_1) \leq A$, the injectivity radius of $\mathbf{H}^3/\Gamma_2$ at s_0 is greater than or equal to ϵ and the injectivity radius of $\mathbf{H}^2/\Gamma_1$ is greater than or equal to ϵ at s_0.

Let $|S|$ be the topological type of a surface S. Note that $|S|$ determines Area(S) via the Euler characteristic. Note also that if we assume that h is injective, the assumption on the injectivity radius in the domain in the preceding definition follows from the assumption in the range, because f is isometric.

5.2.2. Theorem: Compactness of pleated surfaces. *$\mathcal{PSJ}(A,\epsilon)$ is compact.*

In our experience facts about pleated surfaces are accepted far too readily by readers of Thurston's notes. Here are some examples which show how naive intuition can go wrong and which explain why the hypotheses in the theorem are necessary.

5.2.3 Example 1. Let ABC be a small equilateral hyperbolic triangle in $\mathbf{H}^3$. Let α, β and γ be the geodesics orthogonal to the plane of ABC, through A, B and C respectively. On $\mathbf{H}^2$ we fix a geodesic l and mark out lengths along l equal to the sidelength of ABC. We denote these marks by

$$\ldots A_{-1}, B_{-1}, C_{-1}, A_0, B_0, C_0, A_1, B_1, C_1 \ldots \quad .$$

Let α_n, β_n and γ_n $(n \in \mathbf{Z})$ be the geodesics orthogonal to l in $\mathbf{H}^2$. Let f be the obvious pleated map which sends A_n to A, B_n to B, C_n to C, α_n to α, β_n to β and γ_n to γ. We are taking $\Gamma_1 = \Gamma_2 =$ id here, so f is homotopically incompressible.

If we take the limit as the triangle shrinks to a point, we get a limit map which is orthogonal projection of $\mathbf{H}^2$ onto a line. In particular f is not isometric.

This example shows that it is essential for the area of $\mathbf{H}^2/\Gamma_1$ to be bounded.

5.2.4 Example 2. Here is another example, due to Thurston. Let P be a fixed pair of pants. We fix a hyperbolic structure P_0 on P, with geodesic boundary components. We define a map $f: P \to P$ which sends the generators of $\pi_1 P$, which we call x and y, to xyx and yx respectively. The map on homology is the familiar matrix

$$\begin{bmatrix} 2 & 1 \\ 1 & 1 \end{bmatrix} .$$

The boundary components of P correspond to the elements x, y and xy of $\pi_1 P$. They are sent to xyx, yx and $xyxyx$ respectively. The map $f^n : P \to P$ sends x and y to loops whose geodesic representatives have exponentially increasing length. Using Theorem 5.3.6 *(Finite laminations realizable)* we can find a hyperbolic structure P_n on P, with geodesic boundary components and a pleated map f_n representing f^n. Let α_n be a geodesic arc of minimal length joining two distinct boundary curves of P_n. Then the length of α_n tends to zero. We choose as a basepoint for P_n the midpoint of α_n and we choose the image under f_n of this basepoint as a basepoint for P_0. The pleating locus consists of the boundary components plus three spirals which cut P_n up into two ideal triangles.

Now $f_n : P_n \to P_1$ does not converge to a pleated surface. In fact, in the appropriate sense, P_n converges to a straight line. We get an example between closed surfaces of genus 2, by doubling P_n and P_1. However, when extending f_n to the double of P_n, we are obliged to send both copies of P_n to the same copy of P_1. The hypothesis of the theorem which then fails is the incompressibility.

Proof. Compactness of pleated surfaces: Consider $\{P_i\} = \{(\Gamma_{1,i}, \Gamma_{2,i}, f_i)\}$. By Corollary 3.1.7 *(Set of hyperbolic manifolds compact)*, and Corollary 3.1.8 *(Compact with bounded volume)*, we may choose a subsequence $\{P_j\}$ such that $\Gamma_{1,j}$ converges to Γ_1, a torsion free discrete subgroup of $\mathrm{Isom}(\mathbf{H}^2)$, $\Gamma_{2,j}$ converges to Γ_2, a torsion free subgroup of $\mathrm{Isom}(\mathbf{H}^3)$, and, by Ascoli's Theorem, f_j converges to some continuous map $f : \mathbf{H}^2 \to \mathbf{H}^3$. Let the pleating locus λ_j of f_j converge to λ. By Theorem 3.1.4 *(A(U) compact)*, $\mathbf{H}^2/\Gamma_1$ has area bounded by A. It is easy to see that λ is Γ_1 equivariant, and it therefore induces a lamination on a finite area surface. It follows that λ has measure zero. It is easy to see that f is geodesic on each stratum of λ.

... **5.2.5. Lemma: Injective homomorphism.** *The injective homomorphisms $h_j : \Gamma_{1,j} \to \Gamma_{2,j}$ converge to an injective homomorphism $h : \Gamma_1 \to \Gamma_2$. (possibly after taking a subsequence if Γ_2 does not consist entirely of orientation preserving isometries of $\mathbf{H}^3$).*

Proof. Injective homomorphism: Let S be a complementary region of the lamination λ. Given $T \in \Gamma_1$, we let $\{T_j \in \Gamma_1(j)\}$ be a sequence such that $\{T_j\}$ converges to T. Since f_j converges to f and $f_j T_j(x)$ converges to $fT(x)$ for each $x \in \mathbf{H}^2$, $h_j(T_j)$ converges on $f(S)$, which is an open subset of a hyperbolic plane. But since $\{h_j(T_j)\}$ is a sequence of isometries, it converges to some $U \in \text{Isom}(\mathbf{H}^3)$. By Lemma 3.1.3 *(Geometric convergence)*, $U \in \Gamma_2$. U will be called $h(T)$. Actually U is only determined if we know its orientation. So we have to pick a subsequence such that $h_j T_j$ converges whenever we have a sequence $\{T_j\}$ converging to T. It is easy to check that h is a homomorphism. Furthermore, h is injective, since if $h(T) = h(T')$ there would exist two sequences T_j and T'_j converging to T and T' respectively such that $h_j(T_j)$ and $h_j(T'_j)$ both converge to $h(T)$. But this implies that $h_j(T_j^{-1}T'_j)$ is equal to the identity for large values of j. Then we have $T_j = T'_j$ so $T = T'$.

Injective homomorphism

Continuation, proof of Compactness of pleated surfaces: To complete the proof of the theorem, we need to check that f is an isometric map. Before doing this we prove the following technical lemma.

... **5.2.6. Lemma: Technical lemma.** *Suppose $r > 0$ is given. Then there is a number K_r with the following property. Let X, Y, Z, U, V, W be points in $\mathbf{H}^2$ such that the distance between any two of them is at most r. Suppose that we have disjoint geodesics α, β, γ with $X, U \in \alpha$, $Y, V \in \beta$ and $W, Z \in \gamma$. Suppose further that X, Y, Z lie on a geodesic which is orthogonal to α, and that U, V and W are also collinear. (See Figure 5.2.7 .) Then $d(Y,Z) \leq K_r d(V,W)$.*

Proof. Technical lemma: We may regard Y and β as fixed. We assume the result is false and take a sequence of situations such that $d(Y,Z_i)/d(V_i,W_i)$ is unbounded. By taking subsequences, we assume that all sequences of points, geodesics or real numbers which occur in the proof converge (possibly to infinity).

Since $d(Y,Z_i) \leq r$, $d(V_i,W_i)$ tends to zero. Hence γ_i converges to β. Let V_i converge to $V \in \beta$. Then $d(Y,V) \leq r$. Using hyperbolic trigonometry, it is not difficult to show, using the fact that β and γ_i are disjoint, that

$$e^{-r} \leq \lim_{i \to \infty} \frac{d(Y,\gamma_i)}{d(V_i,\gamma_i)} \leq e^r .$$

Since γ_i and α_i are disjoint, we have

$$|\cos(\angle X_i Z_i W_i)| \leq \tanh d(X_i,Z_i)$$

which tends to zero. By the sine rule for hyperbolic triangles,

$$\lim \frac{d(Y,\gamma_i)}{d(Y,Z_i)} = 1 .$$

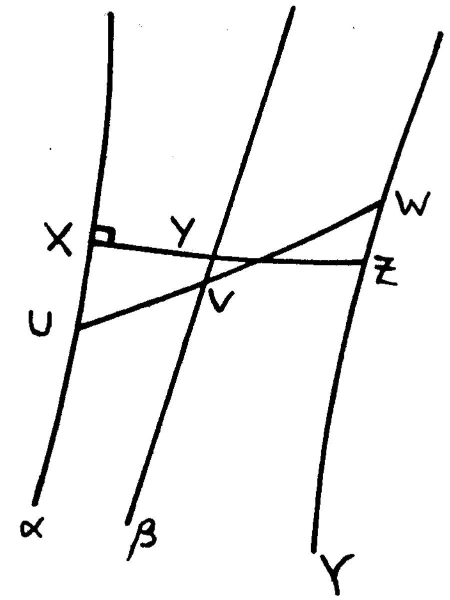

5.2.7 Figure.

Also

$$1 \geq \lim_{i \to \infty} \frac{d(V_i, \gamma_i)}{d(V_i, W_i)} .$$

It follows that

$$\lim_{i \to \infty} \frac{d(Y, Z_i)}{d(V_i, W_i)} = \lim_{i \to \infty} \frac{d(Y, Z_i)/d(Y, \gamma_i)}{d(V_i, W_i)/d(V_i, \gamma_i)} \frac{d(Y_i, \gamma_i)}{d(V_i, \gamma_i)} \leq e^r .$$

But this contradicts our assumption that $d(Y, Z_i)/d(V_i, W_i)$ is unbounded. This proves the lemma.

Technical lemma

Remark: Recently Thurston has circulated a preprint [Thurstona] in which the definition of a pleated surface is changed in order to avoid the necessity for the above technical lemma. In the new definition, a pleated map sends every geodesic to a rectifiable path of the same length, whereas in the old definition, every rectifiable path is sent to a rectifiable path of the same length. Our lemma can be regarded as showing that the new weaker hypothesis implies the old stronger hypothesis, or that the new definition is equivalent to the old one.

Continuation, proof of Compactness of pleated surfaces: Let l be a geodesic in λ. Then there exists $l_j \in \lambda_j$ such that l_j converges to l. We may assume that l and l_j are parametrized geodesics. Then l_j converges uniformly to l on compact subintervals of $\mathbf{R}$. Therefore $f_j \circ l_j$ converges uniformly to $f \circ l$ on compact subintervals of R. Therefore $f \circ l : \mathbf{R} \to \mathbf{H}^3$ is an isometry. Notice that f is lipschitz. In fact $d(fx,fy) \leq d(x,y)$ for all $x,y \in \mathbf{H}^2$.

... **5.2.8. Lemma: Isometric map.** *Let $\lambda \in \mathcal{GL}(\mathbf{H}^2)$ have measure zero, and suppose that $f:\mathbf{H}^2 \to \mathbf{H}^3$ is lipschitz with lipschitz constant less than 1, and is an isometry on each stratum of λ. Then f is an isometric map.*

Proof. Isometric map: Let $p:[0,L] \to \mathbf{H}^2$ be a rectifiable path parametrized proportional to arc length and choose $\epsilon > 0$. Since $p^{-1}(\mathbf{H}^2 \setminus \lambda) \subset [0,1]$, it is composed of a countable union of open intervals which we call the *complementary intervals.* Note that the sum of the lengths of the complementary intervals may be less than L. Pick a finite number of these intervals, $J_1, \ldots, J_r$, such that the total length of the remaining intervals is less than ϵ. Given any partition of $[0,1]$ we enlarge it so that

1) every endpoint of $J_1, \ldots, J_r$ is in the partition, and
2) if a point of the partition lies in a complementary interval, then the endpoints of the interval are in the partition.

Now take an open interval (t_j,t_{j+1}) of the partition which meets λ and let the total length of the complementary intervals contained within (t_j,t_{j+1}) be ϵ_j. Then $\sum \epsilon_j \leq \epsilon$. Let $C = p(t_j)$ and $B = p(t_{j+1})$. By our choice of partition, $B,C \in \lambda$. Let C lie on the geodesic $g \in \lambda$. Let A be the intersection of g with the geodesic orthogonal to g through B. See Figure 5.2.9 .

Let $\{(a_i, b_i)\}$ be the set of complementary intervals of $[A,B] \setminus \lambda$. Since λ has measure zero, $\sum d(a_i,b_i)$ is equal to $d(A,B)$. Let g_i' be the geodesic of λ containing a_i and g_i'' the geodesic of λ containing b_i. Let δ_i be the greatest length of any subpath of $\rho|[t_j,t_{j+1}]$ connecting g_i' to g_i''. By Lemma 5.2.6 *(Technical lemma)*, there is a constant $K = K_L$, such that

$$d(A,B) = \sum d(a_i,b_i) \leq K \textstyle\sum_i \delta_i \leq K\epsilon_j .$$

Since f is lipschitz with constant not greater than one, $d(f(A),f(B)) \leq d(A,B)$. By the triangle inequality

$$|d(f(A),f(C)) - d(f(B),f(C))| \leq d(f(A),f(B)) \leq K\epsilon_j.$$

Also by the triangle inequality we obtain

$$|d(A,C) - d(B,C)| \leq d(A,B) \leq K\epsilon_j.$$

Since f is an isometry on each stratum and both A and C lie on g, $d(f(A),f(C)) = d(A,C)$. So we may combine the two inequalities above to obtain

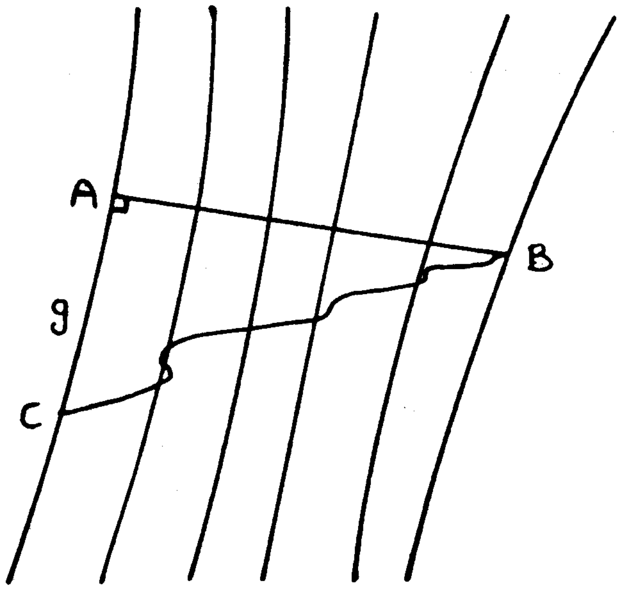

5.2.9 Figure.

$$|d(f(B),f(C)) - d(B,C)| \leq 2K\epsilon_j.$$

Also note that on each of the intervals (t_j,t_{j+1}) of the partition, whose interiors do not meet λ, $d(fp(t_j),fp(t_{j+1})) = d(p(t_j),p(t_{j+1}))$ since f is an isometry on each stratum.

We have thus proved that, given any partition $0 = t_0 < \ldots < t_n = L$ of $[0,L]$ and any $\epsilon > 0$, then (after possibly refining the partition)

$$\left|\sum_{j=0}^{n-1} d(f(p(t_j)),f(p(t_{j+1}))) - \sum_{j=0}^{n-1} d(p(t_j),p(t_{j+1}))\right| \leq 2K\sum\epsilon_j \leq 2K\epsilon\ .$$

This shows that $f \circ p$ is a rectifiable path with the same length as p, completing the proof.

Isometric map

This completes the proof of the theorem

Compactness of pleated surfaces

We can now deduce various consequences of Theorem 5.2.2 *(Compactness of pleated surfaces)*.

In the first version let $\mathcal{PSB}(A,\epsilon)$ ($\mathcal{P}$leated $\mathcal{S}$urfaces with $\mathcal{B}$asepoint) be the set of quintuplets (S,p,M,q,f), where S is a hyperbolic manifold of dimension two with basepoint p, M is a hyperbolic manifold of dimension three with basepoint q, and $f : S,p \to M,q$ is a pleated surface, with $\text{Area}(S) \leq A$ and $\text{inj}_q M \geq \epsilon$.

We topologize $\mathcal{PSB}(A,\epsilon)$ as a quotient of $\mathcal{PSI}(A,\epsilon)$.

5.2.10. Corollary: PSB compact. *$\mathcal{PSB}(A,\epsilon)$ is compact.*

This follows since $\mathcal{PSB}(A,\epsilon)$ is a quotient of $\mathcal{PSJ}(A,\epsilon)$.

Let K be a compact subset of a fixed complete hyperbolic three-manifold N. Let $\mathcal{PS}(A,K,N)$ be the space of pleated surfaces $f : S \to N$, without basepoint, such that $K \cap fS \neq \varnothing$ and with Area$(S) \leq A$. We give this space the quotient topology from the space of pleated surfaces with basepoint $f : S,p \to N,q$, where $q \in K$. The topology can be defined directly. Given $k > 1$, $\epsilon > 0$ and a compact subset K_1 of S, a neighbourhood of $f : S \to N$ consists of all pleated surfaces $f' : S' \to N$ such that there is a k-approximate isometry ϕ between K_1 and a compact subset $K'_1 \subset S'$, such that $d(fx, f' \circ \phi(x)) < \epsilon$ for all $x \in K_1$. The details of checking that this is a correct description of the quotient topology is left to the reader.

5.2.11. Corollary: Unmarked pleated surfaces compact. *Let K be a compact subset of a fixed hyperbolic three-manifold N. Let $A > 0$. Then $\mathcal{PS}(A,K,N)$, the set of pleated surfaces without basepoint which meet K, is compact.*

Proof. Let $\delta < \text{inj}_K(N)$. The space of all orthonormal frames over K forms a compact space K_1. There is a continuous map from K_1 into the space of discrete subgroups of Isom$(\mathbf{H}^3)$, defined by lifting a frame to the standard frame in $\mathbf{H}^3$ and taking the corresponding group of covering translations. This map is clearly continuous. It follows that the relevant set of triples (Γ_1,Γ_2,f) is a closed subset of $\mathcal{PS}(A,\delta)$ and is therefore compact.

□

If we wish to consider only a single homeomorphism type of surface, we have to introduce a condition which will prevent a simple closed geodesic pinching down into two cusps. Since a pleated surface is isometric, it will automatically map cusps of S to cusps of N. This condition means that parabolic elements of $\pi_1 S$ map to parabolic elements of $\pi_1 N$, but another condition is also needed. This condition is that hyperbolic elements in $\pi_1 S$ map to hyperbolic elements in $\pi_1 N$.

5.2.12 Definition. We say a homotopically incompressible map $f : S \to N$ is *non-parabolic,* (abbreviated to *np*) if the induced map $f_* : \pi_1(S) \to \pi_1(N)$ is injective and takes hyperbolic elements to hyperbolic elements and parabolic elements to parabolic elements.

5.2.13. Corollary: Compactness of pleated surfaces of fixed topological type. *Let T be a fixed topological surface of finite type and N a fixed complete hyperbolic three-manifold. Let $K \subset N$ be a fixed compact subset. Then $\mathcal{PS}^{(np)}(T,K,N)$ the space of non-parabolic pleated surfaces of finite area without basepoint $f:S \to N$ which meet K, and such that S homeomorphic to T, is compact.*

Proof. Since $\mathcal{PS}^{(np)}(T,K,N) \subset \mathcal{PS}(A,K,N)$ (where $A = \text{area}(T)$) which is compact, we need only prove that $\mathcal{PS}^{(np)}(T,K,N)$ is closed in $\mathcal{PS}(A,K,N)$. Suppose $\{[f_i:S_i \to N]\} \subset \mathcal{PS}^{(np)}(T,K,N)$ converges to $[f:S \to N] \in \mathcal{PS}(A,K,N)$. We need only prove that S has the same (np)-topological type as S_i. We do this by proving that, for some $\alpha > 0$, no closed geodesic in S_i has length less than α and then applying Proposition 3.2.13 *(Mumford's Lemma)* to obtain our result.

We define the δ-length of a path γ, $l_\delta(\gamma)$ in S_i to be $l(\gamma \cap S_{[\delta,\infty)})$. Correspondingly, we define δ-distance and the δ-diameter of a compact subset of the surface. Clearly, the δ-diameter of S_i is less than $\frac{4\text{area}(S_i)}{\delta} = \frac{4A}{\delta} = B_\delta$. Choose $\delta < \min(\frac{\epsilon_0}{2}, \min_{x \in K}(inj(x)))$ where ϵ_0 is the Margulis constant. Then since loxodromic components of $N_{[\delta/2,\infty)}$ are separated by a δ-distance of at least $\cosh^{-1}\left(\frac{\cosh\delta + 1}{\cosh\delta}\right)$, the set K_δ of all points of δ-distance less than or equal to B_δ from K, contains only finitely many such components. Thus K_δ contains no closed geodesics of length less than α for some $\alpha > 0$. But $f(S_i) \subset K_\delta$, since f decreases δ-distance, completing our proof.

$\square$

We now prove the compactness of the space of marked pleated surfaces, under appropriate conditions.

5.2.14 Definition. A *marked pleated surface* is a pair $([h], f)$, where $h: T \to S$ is a homeomorphism from a topological surface T to a complete hyperbolic 2-manifold S of finite area; $[h]$ is the isotopy class of h; and $f: S \to N$ is a pleated map

More directly, the pair (h_1, f_1) is equivalent to the pair (h_2, f_2) (where $h_i: T \to S_i$ are homeomorphisms and $f_i: S_i \to N$ are pleated surfaces) if there is an isometry $\phi: S_1 \to S_2$ such that the left hand side of the diagram in 5.2.15 commutes up to isotopy and the right hand side commutes precisely.

Our discussion of pleated surfaces will relate to a fixed choice of T and N. We shall denote the space of $\mathcal{M}$arked $\mathcal{P}$leated $\mathcal{S}$urfaces by

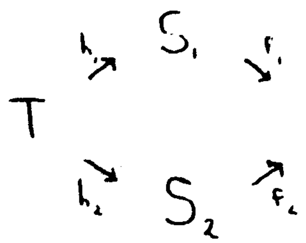

5.2.15 Figure.

$\mathcal{MPS}(T,N)$.

The space of marked pleated surfaces is topologized by saying that $([h_2]: T \to S_2, f_2: S_2 \to N)$ is near to $([h_1]: T \to S_1, f_1: S_2 \to N)$ if $[h_2]$ is near to $[h_1]$ in $\mathbf{T}(T)$ and there exist representatives h'_1 and h'_2 such that $f_1 \circ h'_1$ and $f_2 \circ h'_2$ are nearby in $C(T,N)$ in the compact-open topology. Thus we may consider $\mathcal{MPS}(T,N) \subset \mathbf{T}(T) \times C(T,N)$ where $C(T,N)$ is the quotient space of $C(T,N)$ induced by the equivalence relation of precomposition by homeomorphisms of T isotopic to the identity. Note that this is really only a condition on the thick part of S_1 and S_2, because the values of a pleated surface map on a cusp are determined by its values on the boundary of the thick part.

5.2.16. Lemma: Marked pleated surfaces Hausdorff. *The space of marked pleated surfaces is Hausdorff and has a countable basis.*

Proof. Suppose

$$T \xrightarrow{h_1} S_1 \xrightarrow{f_1} M \text{ and } T \xrightarrow{h_2} S_2 \xrightarrow{f_2} M$$

do not have disjoint neighbourhoods in $\mathcal{MPS}$. Then h_1 and h_2 must represent the same element of Teichmüller space. This means that we can take the hyperbolic surfaces $S_1 = S_2 = S$. We may also take h_1 to be the identity, and then $h = h_2$ is isotopic to the identity.

Let L_1 and L_2 be the pleating loci of f_1 and f_2 respectively. We claim that $L_1 = L_2$. For suppose l is a parametrized geodesic in L_1. By our hypothesis, we may assume that $f_2 h$ is arbitrarily close to f_1 in the compact open topology, if we change h by an isotopy. Let l be a lift of l to $\mathbf{H}^2$. Then $f_1 l$ is a geodesic in $\mathbf{H}^3$. Now $f_2 h(l)$ is close to $f_1 l$, and, by changing h by an isotopy, we can make it arbitrarily close to $f_1 l$ as a parametrized path. Since f_2 is an isometric map, hl is close to some parametrized geodesic. Since hl has the same endpoints as l, hl is close to the parametrized geodesic l. It follows that $f_2 l$ is equal to the parametrized geodesic $f_1 l$, except for a change of origin. Hence l is in the pleating locus of f_2. This means we have proved that $L_1 = L_2$.

But now f_1 and f_2 agree on the boundary of each complementary piece of the lamination (or, more precisely, their images agree). It follows that $f_1 = f_2$. This completes the proof that the space of measured laminations is Hausdorff.

$\mathcal{MPS}(T,N)$ has a countable basis since both $\mathbf{T}(T)$ and $C(T,N)$ do. ∎

Suppose $M \to S^1$ is a fibre bundle and M has a complete hyperbolic structure of finite volume. We know that such structures exist, firstly by the work of Jørgensen [J*orgensen, 1977] in special cases, and later, in great generality, by Thurston [Thurstonb]. The fibre is a surface, denoted by S. The monodromy $\phi : S \to S$ has infinite order in the mapping class group. We shall see later 5.3.6 *(Finite laminations realizable)* that the fibre can be represented by a pleated surface.

Since the monodromy has infinite order, we get an infinite sequence of homotopic marked pleated surfaces with no convergent subsequence. In fact there is not even a convergent subsequence in Teichmüller space. However, up to finite coverings, this is the only way in which the compactness of the space of homotopic marked pleated surfaces fails, as we now proceed to prove.

5.2.17 Definition. Let $f : S \to N$ be a map between manifolds. We say f is a *virtual fibre* of a fibre bundle over N, if N has a finite cover N which is a fibre bundle over the circle, and f lifts to a map into N which is homotopic to the inclusion of the fibre.

5.2.18. Theorem: Compactness of marked pleated surfaces. *Let N be a complete hyperbolic 3-manifold, and let K be a compact subspace of N. Let T be a topological surface of finite type. Then $\mathcal{MPS}^{np}(N,T,K)_{(\bar{f},\bar{h})}$, the space of marked np-pleated surfaces $(h : T \to S, f : S \to N)$ meeting K such that $f \circ h$ is homotopic to $f \circ h$ (by a cusp-preserving homotopy) is compact, unless $f \circ h$ is the virtual fibre of a fibre bundle over N.*

Proof. Compactness of marked pleated surfaces: Suppose (h_i, f_i) is a sequence of homotopic marked pleated surfaces meeting K, with no convergent subsequence. We may assume that $f_i : S_i \to N$ converges in the space of unmarked pleated surfaces, to a pleated surface $f : S \to N$ where S is homeomorphic to T. Let $\phi_i : S \to S_i$ be an approximate isometry, which is an isometry on cusps, such that $f_i \phi_i$ approximates f. Let $L \subset S$ be the complement of the cusps. Since L is compact, the injectivity radius is bounded below on a neighbourhood of L. Therefore $f_i\phi_i$ is homotopic to f by a linear homotopy along shortest geodesics on L and this homotopy extends in a unique standard way to a cusp preserving homotopy on all of S. Suppose there are only finitely many distinct homotopy classes among the maps $\{\phi_i^{-1}h_i : T \to S\}$. In this case we may suppose that this homotopy class is constant and is represented by $h : T \to S$. Now $f_i\phi_i$ is approximately equal to f. Therefore $fh \simeq f\phi_i^{-1}h_i \simeq f_ih_i$. It follows that (f,h) is the limit of the sequence $\{(f_i,h_i)\}$. This is a contradiction.

It follows that there are infinitely many distinct homotopy classes $\{\phi_i^{-1}h_i : T \to S\}$.

... **5.2.19. Lemma: Torsion free kernel.** *Let A be a free abelian group of finite rank, and let $n > 2$ be an integer. Then the kernel of* $\mathrm{Aut}A \to \mathrm{Aut}(A\otimes\mathbf{Z}_n)$ *contains no elements of finite order.*

Proof. Torsion free kernel: Suppose $g : A \to A$ is an automorphism of finite order which induces the identity on $A\otimes\mathbf{Z}_n$. By averaging, we construct a g-invariant bilinear positive definite inner product on A. Let N be the set of nearest elements to 0. This is a finite set (as we see by comparing with the associated positive definite g-invariant inner product on the associated vector space $A\otimes\mathbf{R}$), and it is g-invariant. If x_1 and x_2 are distinct elements of N, then x_1 and x_2 are distinct modulo n (for otherwise one of the $n-1$ intermediate points of A would be non-zero and nearer to 0).

It follows that g is fixed on N. Let B be the subgroup of A generated by N, and let

$$C = \{x \in A : kx \in B \text{ for some } k \neq 0\} .$$

Then g is fixed on B and hence fixed on C. Now A/C is a free abelian group of smaller rank than A, and g induces an automorphism of finite order on A/C, which becomes the identity on $(A/C)\otimes\mathbf{Z}_n$. It follows by induction that this automorphism is the identity on (A/C). Taking a basis for C and extending to a basis of A, we see that g has matrix

$$\begin{bmatrix} I & X \\ 0 & I \end{bmatrix}$$

and, since g has finite order, $X = 0$.

Torsion free kernel

... **5.2.20. Lemma: Non-trivial on homology.** *Let T be a surface of finite type and let $\chi(T) < 0$. Let $g : T \to T$ have finite order in the mapping class group, and suppose $g \not\simeq \mathrm{id}$. Then the induced map on $H_1(T;\mathbf{R})$ is non-trivial.*

The following proof was explained to us by Peter Scott.

Proof. Non-trivial on homology: Let $n > 0$ be the order of g. We may suppose that n is prime. According to [Kerckhoff, 1983] (in fact this was known earlier), we may impose on T a hyperbolic structure preserved by an isometry g of finite order. Let $Q = T/g$. Then Q is an orbifold. A component of the singular set of Q is either a point or a circle. (The latter case happens if g fixes a circle and interchanges the two sides of the circle.) Let E be the group which is the finite extension of $\pi_1 T$ by g. Then E is the fundamental group of the orbifold Q. It can be computed from van Kampen's Theorem, where the pieces to be glued together are nice neighbourhoods of the components of the singular set and Q_0, the complement of these neighbourhoods. It is easy to see that

$$H_1(|Q|;\mathbf{R}) \cong \frac{E}{[E,E]} \otimes \mathbf{R} \cong \frac{H_1(T;\mathbf{R})}{<g>}$$

by applying the Mayer-Vietoris Theorem to the same pieces of $|Q|$. Therefore g induces a non-trivial map on $H_1(T;\mathbf{R})$ if and only if $H_1(|Q|;\mathbf{R})$ has smaller dimension than $H_1(T;\mathbf{R})$.

Suppose, for a contradiction, that these two dimensions are equal and denote the common value by r. Since $\chi(T) < 0$, we must have $r \geq 2$. Let Q have q singular points, and let the inverse image of Q_0 in T be T_0. Amongst the boundary components of T_0, suppose q_1 of these bound disks in T. Then $q_1 \leq nq$. We have

$$n\chi|Q| = n\chi Q_0 + nq = \chi T_0 + nq \geq \chi T_0 + q_1 = \chi T \ .$$

If $|Q|$ is not a closed orientable surface, then $\chi|Q| = -r+1$, so that $n(r-1) = -n\chi|Q| \leq -\chi T = r-1$. Therefore $n \leq 1$, a contradiction. If $|Q|$ is a closed orientable surface, then $r \geq 4$ and $n(r-2) = -n\chi|Q| \leq -\chi T \leq r-1$, so that $n \leq (r-1)/(r-2) \leq 3/2$. Therefore $n = 1$, again a contradiction.

Non-trivial on homology

The following is a well-known result, which we do not in fact use in the sequel, but which is included for its general interest.

... **5.2.21. Corollary: Finite order implies G finite.** *If G is a subgroup of the mapping class group such that every element has finite order, then G is finite.*

Proof. Finite order implies G finite: Let A be the quotient of $H_1(T;\mathbf{Z})$ by its torsion subgroup. The map $G \to \mathrm{Aut}A$ is injective by Lemma *(Non-trivial on homology)*. The composite $G \to \mathrm{Aut}A \to \mathrm{Aut}(A \otimes \mathbf{Z}_3)$ is injective by Lemma 5.2.19 *(Torsion free kernel)*. Therefore G is finite.

Finite order implies G finite

Continuation, proof of Compactness of marked pleated surfaces: There must exist $r \neq s$ such that $h_1^{-1}\phi_1\phi_r^{-1}h_r$ and $h_1^{-1}\phi_1\phi_s^{-1}h_s$ have the same image in $\mathrm{Aut}(H_1(T;\mathbf{Z}_3))$. Then $h_r^{-1}\phi_r\phi_s^{-1}h_s$ is trivial in $\mathrm{Aut}(H_1(T;\mathbf{Z}_3))$ and therefore is of infinite order in the mapping class group. Now h_r, ϕ_r, ϕ_s, h_s can all be assumed to be standard on the cusps. So we can remove the cusps from T to obtain a manifold T_b with boundary and $\psi = h_r^{-1}\phi_r\phi_s^{-1}h_s$ is a homeomorphism of T_b.

Let M be the mapping torus of ψ. We shall define $F : M \to N$. We have already seen that $f_s\phi_s \simeq f_r\phi_r$ by a cusp preserving homotopy. Therefore $f_r h_r \psi = f_r\phi_r\phi_s^{-1}h_s \simeq f_s h_s$ by a cusp preserving homotopy. We can use this to construct a map of M into N. Note that the boundary components of M (all of which are tori or Klein bottles) are sent to cusps of N.

Now $F_\star : \pi_1 M \to \pi_1 N$ is injective. To see this note that $\pi_1 M$ is the extension of $\pi_1 T$ by $\mathbf{Z}$. Let α be the generator of $\mathbf{Z}$, and suppose that $F_\star(\gamma\alpha^n)$ is trivial in $\pi_1 N$ for some $\gamma \in \pi_1 T$. Then $(F_\star\alpha)^n$ lies in the image of $\pi_1 T$. This means $n = 0$ or that conjugation by $F_\star\alpha$ has finite order in the outer automorphism group of $\pi_1 T$. But the effect of this conjugation is equal to $\psi_\star$ which has infinite order. Therefore $n = 0$. Since the pleated surface is incompressible, γ is trivial.

Let $\tilde{N}$ be the covering space of N such that the lift $\tilde{F} : M \to \tilde{N}$ induces an isomorphism of fundamental groups.

First we dispose of the case where M is a closed manifold. Since the higher homotopy groups of M and $\tilde{N}$ are zero, we know that M and $\tilde{N}$ are homotopy equivalent. Hence $\tilde{N}$ is also a closed manifold. From a result of Stallings [Hempel, 1976, Theorem 11.6] we deduce that M is homeomorphic to $\tilde{N}$.

From now on we assume that M has a non-empty boundary. We remove from N uniform horoballs to make a manifold with boundary, which, by abuse of notation, we continue to denote by N. and we have

$\tilde{F}: M, \partial M \to \tilde{N}, \partial\tilde{N}$. A boundary component of M has fundamental group which is maximal solvable in $\pi_1 M$. Therefore its image is maximal solvable in $\pi_1 N$. Therefore each component of ∂N in the image of ∂M is a torus or a Klein bottle and the map has degree one on each component of ∂M. Since fundamental groups of different components of ∂M are not conjugate in $\pi_1 M$, these components are mapped to distinct components of ∂N. Now consider the following diagram

$$\begin{array}{ccc} H_3(M,\partial M;\mathbf{Z}_2) & \to & H_3(\tilde{N},\partial\tilde{N};\mathbf{Z}_2) \\ \downarrow & & \downarrow \\ H_2(\partial M;\mathbf{Z}_2) & \to & H_2(\partial\tilde{N};\mathbf{Z}_2) \end{array}$$

The map on H_2 is injective, and so is $\partial: H_3(M,\partial M;\mathbf{Z}_2) \to H_2(\partial M;\mathbf{Z}_2)$ (assuming for the moment that $\partial M \neq \varnothing$). Hence F_* is injective on H_3. It follows that $H_3(N,\partial N;\mathbf{Z}_2) = \mathbf{Z}_2$. Hence N is a compact manifold and the covering $N \to N$ is finite sheeted. It also follows that each boundary component of ∂N is in the image of ∂M.

M and $\tilde{N}$ are clearly sufficiently large, since they have non-empty boundary. Therefore a result of Waldhausen [Hempel, 1976, Theorem 13.7] can be applied. The twisted I-bundle situation described in the statement of that theorem is ruled out, since there is a bijection between the set of boundary components in our case, but not in the I-bundle situation. Once again, M is homeomorphic to N (with the cusps removed). The theorem follows.

Compactness of marked pleated surfaces

5.3. Realizations

See Section 8.10 of [T].

Given a pleated surface, we have its pleating locus. Rather surprisingly, it is often possible to reverse the direction of this construction. We shall show how, given a topological surface T of finite type, a homotopy class of maps $f: T \to N$ into a complete hyperbolic 3-manifold, and a maximal lamination λ of T, there often exists a complete hyperbolic structure on T and a pleated map $f: T \to N$ in the homotopy class, whose pleating locus is contained in λ.

First we make a definition

5.3.1 Definition. Let $f : S \to N$ be a pleated surface and let α be a finite geodesic arc in S. We define a non-negative real number which we call the *pleating* of α as follows. We lift α to $\tilde{\alpha}$ in $\mathbf{H}^3$. The pleating is the difference of the distance between the endpoints of $\tilde{\alpha}$ in $\mathbf{H}^3$ and the length of α.

If α has zero pleating, then it must lie in a single stratum of the pleating locus (except possibly for its endpoints).

We now prove that the pleating locus is a continuous function of the pleated surface.

Let N be a fixed complete hyperbolic 3-manifold. For each marked pleated surface $(h : T \to S, f : S \to N)$ the pleating locus of f is a geodesic lamination on S. By fixing a complete hyperbolic structure on T we obtain a geodesic lamination on T. (See 4.1.4 *(Transferring laminations)*.)

5.3.2. Lemma: Pleating locus continuous. *The map* $\mathcal{MPS}(T,N) \to \mathcal{GL}(T)$ *is continuous if we use the Thurston topology on* $\mathcal{GL}(T)$.

5.3.3 Remark. If we use the Chabauty topology on $\mathcal{GL}(T)$, the map is not continuous. For example, we can take a Fuchsian group with empty pleating locus and approximate it by a quasifuchsian group with any preassigned lamination of compact support as its pleating locus.

Proof. Let $(h_i : T \to S_i, f_i : S_i \to N)$ converge to (h,f). Let $\phi_i : S \to S_i$ be an approximate isometry such that $f_i\phi_i$ is approximately equal to f and h_i is homotopic to $h\phi$. We may assume that h and h_i are standard on the cusps and that ϕ_i is an isometry on the cusps. There is no loss of generality in taking $T = S$, h = id, and $h_i = \phi_i$. Let λ and λ_i be the pleating loci of f and f_i. Let α be a geodesic in $\mathbf{H}^2$ in the lamination $\tilde{\lambda}$. We must show that there is a geodesic α_i in $\tilde{\lambda}_1$ such that $\tilde{\phi}_i^{-1}\alpha_i$ has endpoints near to those of α.

By choosing appropriate coverings and lifts, we may assume that $\{f_i : \mathbf{H}^2 \to \mathbf{H}^3\}$ converges uniformly to $f : \mathbf{H}^2 \to \mathbf{H}^3$ on compact sets. Let β be a short arc crossing α. Then the pleating of (β, f) is non-zero. It follows that the pleating of (β, f_i) is non-zero for sufficiently large i. Therefore there exists a geodesic $\alpha_i \in \tilde{\lambda}_i$ which crosses β. We claim that $\{\alpha_i\}$ converges to α. For suppose not. Then by choosing a subsequence, we may assume $\{\alpha_i\}$ converges to $\alpha_\infty \neq \alpha$. Therefore α_∞ intersects α at a point x. But, by uniform convergence, $f|\alpha_\infty$ must be an isometry, and this would imply that x is not in the pleating locus of f, which is a contradiction. □

5.3.4 Definition. Suppose λ is a lamination on a complete hyperbolic surface T, and that $f : T \to N$ is a given *np* incompressible surface. We say (λ, f) is ***realizable*** and write $\lambda \in R_f$ if there is a map f' homotopic to f, by a cusp preserving homotopy which maps each geodesic of λ homeomorphically to a geodesic of N.

We do not require f' to preserve arclength (not even up to change of scale) on the geodesics of λ. Nor do we require the image of a geodesic of λ under f' to be a simple geodesic in N. It follows that the question of whether a lamination is realizable or not does not depend on the underlying hyperbolic structure of the surface. By an easily constructed homotopy, we may assume that every realization f is isometric on each geodesic ray of S, which lies in a cusp of S and travels straight towards the cusp point. We may also assume that f sends every horocyclic curve in a cusp of S to a piecewise linear curve in a horosphere of N (using the euclidean metric in a horosphere of N to give the meaning of "piecewise linear"). The map f is linear on each horocyclic interval disjoint from λ.

5.3.5. Lemma: Image well-defined. *If (λ, f) is realizable, then the image of λ in N under a realization is well-defined.*

Proof. Image well-defined: More precisely, we can lift any two realizations f_1, f_2 to two maps $\tilde{f}_1, \tilde{f}_2 : \mathbf{H}^2 \to \mathbf{H}^3$ such that $\tilde{f}_1(\alpha)$ and $\tilde{f}_2(\alpha)$ are equal as directed geodesics in $\mathbf{H}^3$ (though the parametrizations induced by f_1 and f_2 may be different) for every $\alpha \in \tilde{\lambda}$. To see this, we fix a complete hyperbolic structure on T. Let $h : T \times I \to N$ be a cusp-preserving homotopy between f_1 and f_2. We adjust h, f_1 and f_2 so that on each cusp they are "linear", i.e. isometric on each geodesic straight up the cusp and horocyclic intervals disjoint from λ are sent linearly to straight lines in a horosphere of N. The thick part of T is compact. It follows that by lifting h to $\tilde{h} : I \times \mathbf{H}^2 \to \mathbf{H}^3$ we obtain a homotopy such that the distance in $\mathbf{H}^3$ from $h_0(x)$ to $h_1(x)$ is bounded as $x \in \mathbf{H}^2$ varies. But then h_0 and h_1 have the same effect on endpoints of geodesics of $\tilde{\lambda}$.

Image well-defined

We spend the rest of this section proving that if f preserves parabolicity, then R_f is open and dense. The proof is by a series of lemmas

5.3.6. Theorem: Finite laminations realizable. *Any finite lamination is realizable. If μ is a geodesic lamination which is realizable and λ is any finite extension of μ, then λ is also realizable.*

Proof. Finite laminations realizable: We first do the case when λ is finite (i.e. $\mu = \emptyset$) first. Start by arranging for f to be a geodesic immersion on each cusp of T. Each simple closed geodesic in λ is sent to a well-defined conjugacy class in $\pi_1 N$ and hence to a well-defined closed geodesic. Fix a map on such simple geodesics which is parametrized proportional to arclength. By the Homotopy Extension Theorem we change f so that it is such a map on the closed geodesics of λ. We lift f to $\tilde{f} : \mathbf{H}^2 \to \mathbf{H}^3$. Now f realizes the sublamination μ of λ consisting of closed geodesics. Therefore $\tilde{f}$ defines a map on the endpoints of geodesics of $\tilde{\mu}$. $\tilde{f}$ also gives a map on the endpoints of geodesics of $\tilde{\lambda}$ which run up cusps of T. If α is any geodesic of $\tilde{\lambda}$, then each end either spirals onto a geodesic in μ or ends in a cusp. Therefore we know where the endpoints of α should be mapped to. This determines the image of α under the realization.

Now $|\lambda|$ is not a nice subspace of T and the Homotopy Extension Theorem cannot be applied to it, so we need to define the realization on a subspace which contains $|\lambda|$ and to which the Homotopy Extension Theorem can be applied.

Let N be the set of points x in T such that there are distinct homotopy classes of paths of length less than ϵ, with one endpoint at x and the other on λ There is a canonical foliation of N (see [Epstein-Marden]). We map each leaf of this foliation to the appropriate horocyclic curve in N, by a map which is proportional to arclength on each horocyclic interval complementary to λ. It is easy to see that this gives a definite map on N. We now extend to the rest of $|\lambda|$ and then use the Homotopy Extension Theorem to extend to the rest of T.

The same method is used to extend a realization from any geodesic lamination μ to a finite extension λ of μ. The only point we need to check on is that, when trying to define the image of $\alpha \in \tilde{\lambda} \setminus \tilde{\mu}$, we do not find that the two putative endpoints of $\tilde{f}\alpha$ are equal in $\mathbf{H}^3$. This is clear if μ is a finite lamination, but in general a little work is needed. Let $f : T \to N$ be a realization on μ and let $\tilde{f}$ be a lifting of f. Let λ be a geodesic lamination and let α be a geodesic in $\tilde{\lambda}$ with endpoints at the endpoints of $\tilde{\mu}$. Let $\lambda \setminus \mu$ consist of a finite number of geodesics. We now prove the following lemma.

... 5.3.7. Lemma: Image points not too close. *Let $f : S \to N$ be a cusp preserving incompressible map which is "linear" on the complement of λ in any cusp and let $\tilde{f} : \mathbf{H}^2 \to \mathbf{H}^3$ be a fixed lifting. Given $r > 0$, there exists $s > 0$ such that if $x, y \in \mathbf{H}^2$ and $d(x,y) > s$, then $d(\tilde{f}x, \tilde{f}y) > r$. In particular, $\tilde{f}$ is proper.*

Proof. Image points not too close: First we show $\tilde{f}$ is proper. Let K be a compact subspace of $\mathbf{H}^3$. Let $\delta > 0$ be chosen small enough so that $f(\pi^{-1}S_{(0,\delta]}) \cap K = \varnothing$. Let $F \subset \mathbf{H}^2$ be a fundamental region for S and let $F_\delta \subset F$ be a fundamental region for $S_{[\delta,\infty)}$. Then F_δ is compact. Since $\pi_1 N$ is a discrete subgroup of Isom($\mathbf{H}^3$), there are only a finite number of covering translations T of $\mathbf{H}^3$ such that $T f F_\delta \cap K \neq \varnothing$. It follows that $f^{-1}K$ is contained in a finite union of translates of F_δ, and is therefore compact. So f is proper.

Now suppose that $x_n, y_n \in \mathbf{H}^2$ and $d(x_n, y_n) \to \infty$. We have to show that $d(fx_n, fy_n) \to \infty$. We may suppose that $x_n \in F$ for each n. We may also suppose that neither x_n nor y_n converges to the end of a cusp, for the result is then obvious by our special assumptions on f. Therefore we may assume that x_n converges to a point $x \in F$. The result then follows since f is proper.

Image points not too close

... **5.3.8. Lemma: Endpoints not equal.** *The endpoints of* $\tilde{f}\alpha$ *are distinct in* $\mathbf{H}^3$.

Proof. Endpoints not equal: Suppose that the two endpoints of $\tilde{f}\alpha$ are equal in $\mathbf{H}^3$. Let the two ends of α be asymptotic to the oriented geodesics α_1 and α_2 at the positive ends of α_1 and α_2 where $\alpha_1, \alpha_2 \in \tilde{\mu}$. These ends are distinct in $\mathbf{H}^2$. Therefore, by Lemma 5.3.7 *(Image points not too close)*, $f\alpha_1$ and $f\alpha_2$ are distinct in $\mathbf{H}^3$. Therefore the endpoints of $f\alpha$ are distinct.

Endpoints not equal

This completes the proof of the theorem

Finite laminations realizable

Recall from Theorem 4.2.14 *(Finite laminations dense)* that finite laminations are dense in $\mathcal{GL}(S)$ in both the Chabauty topology and the Thurston topology.

5.3.9. Theorem: Existence of realizing structure. *Let* λ *be a maximal lamination in* S *and let* $f : S \to N$ *be a realization of* λ. *Then there is a unique hyperbolic structure* S' *on* S *and a pleated surface* $f' : S' \to N$ *homotopic to* f, *with pleating locus contained in* λ.

Proof. Existence of realizing structure: We prove it first for finite laminations. Each ideal triangle in $S \backslash \lambda$ is mapped by a unique isometry to a well-defined ideal triangle in N.

Let $S_1, \ldots, S_k$ be the components of the surface obtained by removing from S all the simple closed geodesics of λ. We get (new) well-defined incomplete hyperbolic structures on each S_i, defined in such a way that we obtain pleated maps on the completion S_i. S_i is constructed by starting with one ideal triangle, gluing on the next in the unique way so that we get a pleated map into N, and so on. Let α be a simple closed geodesic in λ. Then α corresponds to exactly two boundary curves α_1 and α_2 of the S_i, and we have to glue α_1 to α_2. Now $f : S_i \to \mathbf{H}^3$ maps horocycles centred at one end of α_i to horospheres centred at the corresponding end of $f\alpha$ in $\mathbf{H}^3$ and similarly for α_2. (Which end depends on the direction of spiralling of the lamination around α.) Extending by continuity, we get an isometric map $S_i \to \mathbf{H}^3$. This induces a well-defined isometric map $S_i \to N$. Now α_1 and α_2 map isometrically to the same closed geodesic of N. We identify α_1 to α_2 in such a way that we obtain a well-defined pleated surface $f' : S' \to N$.

For a general maximal lamination, the result follows from the compactness of pleated surfaces. By lifting to a cover, we may assume that $f : S \to N$ induces an isomorphism of fundamental groups. Let μ be a minimal sublamination of S, and let $U \subset N$ be a small closed ball meeting $f\mu$. By Theorem 4.2.10 *(Curve near geodesic)* and Theorem 4.2.14 *(Finite laminations dense)*, we may assume that we have a sequence λ_i of finite laminations, converging to λ. Since f is uniformly continuous, we may assume that there is a simple closed geodesic C_i in λ_i such that the simple closed geodesic corresponding to fC_i meets U. (See Theorem 4.2.10 *(Curve near geodesic)*.) Let f_i be the unique marked pleated surface with pleating locus contained in λ_i. By Theorem 5.2.18 *(Compactness of marked pleated surfaces)*, we may assume that f_i converges to a pleated surface $f' : S' \to N$(if f is a virtual fibre we must lift to a cover M such that $\pi_1(M) = f_*(\pi_1(T))$), and by Corollary 5.2.13 *(Compactness of pleated surfaces of fixed topological type)* S' is of the same type as S. By Lemma 5.3.2 *(Pleating locus continuous)*, the pleating locus λ_0 of f_0 is contained in λ. This shows the existence of $f' : S' \to N$. To prove the uniqueness of the pleated surface realizing λ, we give a more abstract version of S'. Let $\mathbf{H}^2 \to S$ be the universal cover of S. We lift λ to a lamination $\tilde{\lambda}$ of $\mathbf{H}^2$. We identify each geodesic of $\tilde{\lambda}$ to a point and each component of $\mathbf{H}^2 \backslash \tilde{\lambda}$ to a point. (Each such component is an open ideal triangle.) The resulting quotient space of $\mathbf{H}^2$ is called P_λ. Geodesics of $\tilde{\lambda}$ give closed points of P_λ and triangular strata give rise to open points. P_λ is not Hausdorff, because a point corresponding to a triangle is in every neighbourhood of each of its three sides. $\pi_1 S$ acts on P_λ via covering translations of $\mathbf{H}^2$.

If S' is a complete hyperbolic surface of finite area and we have a homeomorphism from S to S', we obtain a lifted homeomorphism of $\mathbf{H}^2$ with itself. This homeomorphism extends to the boundary circle, and is unaltered by an isotopy of homeomorphisms between S and S'. (The lift to $\mathbf{H}^2 \cup \partial \mathbf{H}^2 \to \mathbf{H}^2 \cup \partial \mathbf{H}^2$ can be changed by composition with an element of $\pi_1 S$ on the right, or, equivalently, with an element of $\pi_1 S'$ on the left.) A geodesic is represented by an unordered pair of distinct elements of S'. A geodesic of S is, in this way, sent to a geodesic of S' and so λ can be transferred from S to a lamination λ' of S'. Since P_λ can be defined entirely in terms of pairs and triples of points in $\partial \mathbf{H}^2$, we see that P_λ is homeomorphic to $P_{\lambda'}$.

Let $\tilde{f} : \mathbf{H}^2 \to \mathbf{H}^3$ be a lift of f, a realization of λ. If α is a geodesic of $\tilde{\lambda}$, then $f\alpha$ is a geodesic in $\mathbf{H}^3$ and we write $f(\alpha)$ for this geodesic. If s is a triangular component of $\mathbf{H}^2 \setminus \lambda$, then its boundary is mapped to an ideal triangle in $\mathbf{H}^3$, which we denote by $f(S)$.

Now consider the subset A of $P_\lambda \times \mathbf{H}^3$, consisting of pairs (y,x) such that $x \in fy$. It is easy to check that $A/\pi_1 S$ is Hausdorff. Let $f' : S' \to N$ be a marked pleated surface homotopic to $f : S \to N$, with pleating locus λ' corresponding to $\lambda \subset S$. Then the lift $f' : \mathbf{H}^2 \to \mathbf{H}^3$ of f' induces a map of S' into A and hence a map $S' \to A/\pi_1 S$. Each cusp of S only contains a finite number of geodesics of λ. From this it is easy to see that the map $S' \to A/\pi_1 S$ is a homeomorphism on the cusps. So we can apply the theorem that a bijective continuous map of a compact space to a Hausdorff space is a homeomorphism, to deduce that $S' \to A/\pi_1 S$ is a homeomorphism. The map $f' : S' \to N$ factors as

$$S' \to A/\pi_1 S \to N$$

where the second map is induced from the projection $A \to \mathbf{H}^3$. It follows that $f' : S' \to N$ is unique, up to composition of an isometry of another complete hyperbolic surface S'' with S'.

Existence of realizing structure

We have shown that R_f is dense (see Theorem 5.3.6 *(Finite laminations realizable)*); we now wish to show it is open. We need to make use of the theory of train tracks. We shall provide a few of the definitions, but the interested reader is encouraged to seek further detail in [Casson, 1983] or [Harer-Penner, 1986].

A *train track* τ on a hyperbolic surface S, is a finite graph embedded in S, such that all edges of the graph are C^1-embedded in S, all edges are tangent at any vertex (these vertices are called *switches*), and if you double any component of $S \setminus \tau$ along its (open) edges the resulting surface has negative Euler characteristic. A *train route* is a C^1-immersion $\rho: \mathbf{R} \to \tau$. A neighbourhood U of τ which is foliated by arcs (called *ties)* transverse to τ, such that each transverse arc meets τ in just one point (except near the switches) is called a *standard neighbourhood* of τ. A geodesic lamination λ is said to be *carried* by a train track τ if there is some standard neighbourhood U of τ such that $\lambda \subset U$ and each leaf of λ is transverse to every tie. So every leaf of the lamination is homotopic to an unique train route. (One way of obtaining such a train track is to take an ϵ-neighbourhood of λ, and "squeeze" it down into a train track.) Now, we simply note that the set $N(\tau)$ of all laminations carried by a given train track τ, is an open subset of $\mathcal{GL}(S)$ in the Thurston topology, and that every geodesic lamination is carried by some train track.

Suppose λ can be realized by a pleated map $g: \bar{S} \to N$ and suppose that λ is carried by a train track τ where $\tau \subset N_\epsilon(\lambda)$. We choose τ so that an ϵ-neighbourhood of τ contains λ. By taking ϵ small enough we can ensure that any train track path can be represented by a sequence of long (i.e of length greater than some $\alpha > 0$) geodesic arcs, joining each other at angles almost equal to π (we refer to points where these angles occur as bends.) Notice that all switches are bends, but that we may have to insert some bends at points which are not switches. Similarly, we may construct a "train track" τ' in N which is "near" to the $g(\tau)$, in particular the bends of τ' are the images of τ's bends, which also has the above properties. The image of every train route of τ is associated to a well-defined geodesic of N, since if we examine a lift of this train route to $\mathbf{H}^3$ it has well defined and distinct endpoints by Lemma 4.2.10 *(Curve near geodesic)*

Now suppose λ' is carried by τ, by the above λ' has a well-defined image in N which is a collection of geodesics (although it isn't necessarily a geodesic lamination as the geodesics may intersect). Now as in Theorem 5.3.9 *(Existence of realizing structure)* we may choose a sequence of finite laminations λ_i converging to λ' and prove the existence of a realizing map $f': S' \to N$. Hence any lamination carried by τ is realizable, thus proving R_f is open. We summarize the above results in a theorem.

5.3.10. Theorem: Realizable laminations open and dense. *If f preserves parabolicity, R_f is open and dense.*

5.3.11. Theorem: Lamination realizable. *Given a surface T, a geometrically finite complete hyperbolic 3-manifold N, and $f : T \to N$, an incompressible map taking cusps to cusps, satisfying np and which is not a virtual fibre, then every $\lambda \in \mathcal{GL}(S)$ may be realized by a pleated surface in the homotopy class of f (i.e. $R_f = \mathcal{GL}(S)$).*

Proof. Lamination realizable: First note that if f realizes λ, then $f\lambda \subset C_N$, the convex core of N. The reason for this is that it is true for finite laminations by construction, and these are dense. Let K be the thick part of the convex core of N. Then K is compact (see [] for example) and every pleated surface meets K. By Theorem 5.2.18 *(Compactness of marked pleated surfaces)*, the space of marked pleated surfaces is compact. The set of finite-leaved maximal laminations is dense in the set of maximal laminations, and every maximal geodesic lamination is realizable by a pleated surface in the homotopy class of f. Since every geodesic lamination is contained in a maximal lamination, every geodesic lamination is realizable.

$\square$

As an example of the power of the techniques in the last few sections, we state a trivial corollary of Corollary 5.2.13 *(Compactness of pleated surfaces of fixed topological type)* and the realizability of any (*np*)-surface subgroup of a hyperbolic 3-manifold by a pleated surface. (Of course, the main applications of the above material is in Thurston's proof of his uniformization theorem.)

5.3.12. Corollary: Finiteness theorem. *Let S be any surface of finite area and N any geometrically hyperbolic 3-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.*

References

Proceedings of the Smith Conjecture Symposium 1979.

Abikoff, 1980.

W. Abikoff, "The real analytic theory of Teichmüller space," *Springer Lecture Notes*, **820**, 1980.

Apanasov, 1985.
B. Apanasov, "Cusp Ends of Hyperbolic Manifolds," *Annals of Global Analysis and Geometry*, **3**, pp. 1-11, 1985.

Beardon, 1983.
A. F. Beardon, *The Geometry of Discrete Groups,* Springer, New York, 1983.

Casson, 1983.
A. Casson, *Automorphisms of Surfaces after Nielsen and Thurston,* University of Texas at Austin, 1983. Lecture Notes by S.Bleiler

Chabauty, 1950.
C. Chabauty, "Limites d'ensembles et géométrie des nombres," *Bull. Soc. Math. de France*, **78**, pp. 143-51, 1950.

Connelly, 1971.
R. Connelly, "A New Proof of Brown's Collaring Theorem," *Proceedings of the A.M.S.*, **27**, pp. 180-2, 1971.

Douady, 1979.
A. Douady, "L'espace de Teichmüller," *Astérisque*, **66-67**, pp. 127-137, Société Mathématique de France, 1979.

Edwards-Kirby, 1971.
R.D. Edwards and R.C. Kirby, "Deformations of spaces of imbeddings," *Annals of Maths.*, **93**, pp. 63-68, 1971.

Epstein, 1983.
D.B.A. Epstein, *Isomorfismi Conformi e Geometria Iperbolica,* University of Cagliari, 1983. Lecture notes, written by G. D'Ambra

Epstein, 1984.
D.B.A. Epstein, "Transversely hyperbolic 1-dimensional foliations," *Astérisque*, **116**, pp. 53-69, Société Mathématique de France, 1984.

Epstein-Marden.
D.B.A. Epstein and A. Marden, *Convex Hulls in Hyperbolic Space, a Theorem of Sullivan, and Measured Pleated Surfaces,* Cambridge University Press. to appear

Godement, 1958.
R. Godement, *Topologie algébrique et théorie des faisceaux,* Her-

mann, Paris, 1958.

Gromov, 1980.

M. Gromov, "Hyperbolic Manifolds according to Thurston & Jørgensen," in *Séminaire Bourbaki 1979/80*, Springer-Verlag, 1980.

Gromov, 1981a.

M. Gromov, "Groups of polynomial growth and expanding maps," *Publ. Math. I.H.E.S.*, **53**, pp. 53-78, 1981.

Gromov, 1981b.

M. Gromov, *Structures Metriques pour les Varietes Riemanniennes*, CEDIC/Fernand Nathan, Paris, 1981.

Haefliger, 1958.

A. Haefliger, "Structure Feuilletées et Cohomologie à Valeur dans un Faisceau de Groupoïdes," *Comm. Math. Helv.*, **32**, pp. 248-329, 1958.

Harer-Penner, 1986.

J. Harer and R. Penner, "Combinatorics of Train Tracks," *Annals of Math. Studies*, Princeton University Press, 1986. to appear

Harvey, 1977.

W. J. Harvey, "Spaces of Discrete Groups," in *Discrete Groups and Automorphic Functions*, ed. W. J. Harvey, Academic Press, London, 1977.

Hempel, 1976.

J. Hempel, *3-Manifolds,* Annals of Mathematics Studies, Princeton University Press, Princeton N.J., 1976.

Jørgensen, 1977.

T. Jørgensen, "Compact 3-manifolds of constant negative curvature fibering over the circle," *Annals of Maths.*, **106**, pp. 61-72, 1977.

Kazdan-Margulis, 1968.

D.A. Kazdan and G.A. Margulis, "A Proof of Selberg's Conjecture," *Math. U.S.S.R. Sbornik*, **4**, pp. 147-152, 1968.

Kerckhoff, 1983.

S. P. Kerckhoff, "The Nielsen realization problem." *Annals of Math.*, **117**, pp. 235-265, March, 1983.

Kobayashi-Nomizu, 1963.

S. Kobayashi and K. Nomizu. *Foundations of Differential Geometry*

I, Interscience, 1963.

Lok, 1984.

W.L. Lok, *Deformations of locally homogeneous spaces and Kleinian groups,* Columbia University, 1984. Ph.D. Thesis

Michael, 1951.

E. Michael, "Topologies on Spaces of Subsets," *Transactions of the A.M.S.*, **71**, pp. 152-182, 1951.

Rudin, 1973.

Rudin, *Functional Analysis,* Mc.Graw-Hill, New York, 1973.

Scott, 1983.

P. Scott, "The Geometries of 3-manifolds," *Bull. London Math. Soc.*, **15**, pp. 401-487, 1983.

Siebenmann, 1972.

L. Siebenmann, "Deformations of Homeomorphisms on Stratified Sets," *Comm. Math. Helv.*, **47**, pp. 123-136, 1972.

Thurstona.

W. P. Thurston, *Hyperbolic Structures on 3-Manifolds I: Deformations of Acylindrical Manifolds.* to appear

Thurstonb.

W. P. Thurston, "Surface Groups Which Fiber," *Hyperbolic Structures on 3-Manifolds,* **II**. To Appear

Thurston, 1979.

W. P. Thurston, *The Geometry and Topology of 3-Manifolds,* Princeton University Mathematics Department, 1979. Part of this material – plus additional material – will be published in book form by Princeton University Press

Thurston, 1984.

W. P. Thurston, "Three-dimensional geometry and topology," *Princeton University Press*, 1984. to appear

Wang, 1972.

H. C. Wang, "Topics in Totally Discontinuous Groups," in *Symmetric Spaces*, ed. Boothby-Weiss, New York, 1972.

Weil, 1960.

A. Weil, "On Discrete Subgroups of Lie Groups," *Annals of*

Mathematics, **72**, pp. 369-84, 1960.

Whitney-Bruhat, 1959.

H. Whitney and F. Bruhat, "Quelques propriétés fondamentales des ensembles analytiques-réels," *Comm. Math. Helv.*, **33**, pp. 132-160, 1959.

Wielenberg, 1981.

N.J. Wielenberg, "Hyperbolic 3-Manifolds which share a Fundamental Polyhedron," in *Riemann Surfaces and Related Topics: Proceedings of the 1978 Stony Brook Conference*, Princeton University Press, 1981.

R.D. Canary,
Fine Hall,
Princeton,
N.J. 08544,
U.S.A.

D.B.A. Epstein, University of Warwick, Coventry, CV4 7AL, ENGLAND.

P. Green, University of Warwick, Coventry, CV4 7AL, ENGLAND.