THEOREM 5.8.2. *If $M = M_{\infty,\dots,\infty}$ admits a hyperbolic structure then there is a neighborhood U of $(\infty, \dots, \infty)$ in $S^2 \times S^2 \times \cdots \times S^2$ such that for all $(d_1, \dots, dk) \in U$, $M_{d_1,\dots,d_k}$ admits a hyperbolic structure.*

PROOF. Consider the compact submanifold $M_0 \subset M$ gotten by truncating each end. M_0 has boundary a union of k tori and is homeomorphic to the manifold $\bar{M}$ such that $M =$ interior $\bar{M}$. By theorem 5.6, M_0 has a k complex parameter family of non-trivial deformations, one for each torus. From the lemma above, each small deformation gives a hyperbolic structure on some $M_{d_1,\dots,d_k}$. It remains to show that the d_i vary over a neighborhood of $(\infty, \dots, \infty)$.

Consider the function

$$\mathrm{Tr} : \mathrm{Def}(M) \to (\mathrm{Tr}(H(\alpha_1)), \dots, \mathrm{Tr}(H(\alpha_k)))$$

which sends a point in the deformation space to the k-tuple of traces of the holonomy of $\alpha_1, \alpha_2, \dots, \alpha_k$, where α_i, β_i generate the fundamental group of the i-th torus. Tr is a holomorphic (in fact, algebraic) function on the algebraic variety 5.36
$\mathrm{Def}(M)$. $\mathrm{Tr}(M_{\infty,\dots,\infty}) = (\pm 2, \dots, \pm 2)$ for some fixed choice of signs. Note that $\mathrm{Tr}(H(\alpha_i)) = \pm 2$ if and only if $H(\alpha_i)$ is parabolic and $H(\alpha_i)$ is parabolic if and only if the i-th surgery coefficient d_i equals ∞. By Mostow's Theorem the hyperbolic structure on $M_{\infty,\dots,\infty}$ is unique. Therefore $d_i = \infty$ for $i = l, \dots, k$ only in the original case and $\mathrm{Tr}^{-1}(\pm 2, \dots, \pm 2)$ consists of exactly one point. Since $\dim(\mathrm{Def}(M)) \geq k$ it follows from [ ] that the image under Tr of a small open neighborhood of $M_{\infty,\dots,\infty}$ is an open neighborhood of $(\pm 2, \dots, \pm 2)$.

Since the surgery coefficients of the i-th torus depend on the trace of both $H(\alpha_i)$ and $H(\beta_i)$, it is necessary to estimate $H(\beta_i)$ in terms of $H(\alpha_i)$ in order to see how the surgery coefficients vary. Restrict attention to one torus T and conjugate the original developing image of $M_{\infty,\dots,\infty}$ so that the parabolic fixed point of the holonomy, $H_0, (\pi_1 T)$, is the point at infinity. By further conjugation it is possible to put the holonomy matrices of the generators α, β of $\pi_1 T$ in the following form:

$$H_0(\alpha) = \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix} \quad H_0(\beta) = \begin{bmatrix} 1 & c \\ 0 & 1 \end{bmatrix}.$$

Note that since $H_0(\alpha), H_0(\beta)$ act on the horospheres about ∞ as a two-dimensional lattice of Euclidean translations, c and l are linearly independent over $\mathbb{R}$. Since 5.37
$H_0(\alpha), H_0(\beta)$ have $\left[\begin{smallmatrix} 1 \\ 0 \end{smallmatrix}\right]$ as an eigenvector, the perturbed holonomy matrices

$$H(\alpha), H(\beta)$$

will have common eigenvectors near $\left[\begin{smallmatrix}1\\0\end{smallmatrix}\right]$, say $\left[\begin{smallmatrix}1\\\epsilon_1\end{smallmatrix}\right]$ and $\left[\begin{smallmatrix}1\\\epsilon_2\end{smallmatrix}\right]$. Let the eigenvalues of $H(\alpha)$ and $H(\beta)$ be (λ, λ^{-1}) and (μ, μ^{-1}) respectively. Since $H(\alpha)$ is near $H_0(\alpha)$,

$$H(\alpha)\begin{bmatrix}0\\1\end{bmatrix} \approx \begin{bmatrix}1\\1\end{bmatrix}.$$

However

$$H(\alpha)\begin{bmatrix}0\\1\end{bmatrix} = \frac{1}{\epsilon_1-\epsilon_2}H(\alpha)\left(\begin{bmatrix}1\\\epsilon_1\end{bmatrix}-\begin{bmatrix}1\\\epsilon_2\end{bmatrix}\right) = \frac{1}{\epsilon_1-\epsilon_2}\left(\begin{bmatrix}\lambda\\\lambda\epsilon_1\end{bmatrix}-\begin{bmatrix}\lambda^{-1}\\\lambda^{-1}\epsilon_2\end{bmatrix}\right).$$

Therefore

$$\frac{\lambda-\lambda^{-1}}{\epsilon_1-\epsilon_2} \approx 1.$$

Similarly,

$$\frac{\mu-\mu^{-1}}{\epsilon_1-\epsilon_2} \approx c.$$

For λ, μ near l,

$$\frac{\log(\lambda)}{\log(\mu)} \approx \frac{\lambda-1}{\mu-1} \approx \frac{\lambda-\lambda^{-1}}{\mu-\mu^{-1}} \approx \frac{1}{c}.$$

Since $\tilde{H}(\alpha) = \log\lambda$ and $\tilde{H}(\beta) = \log\mu$ this is the desired relationship between $\tilde{H}(\alpha)$ and $\tilde{H}(\beta)$.

The surgery coefficients (a, b) are determined by the formula 5.38

$$a\tilde{H}(\alpha) + b\tilde{H}(\beta) = \pm 2\pi i.$$

From the above estimates this implies that

$$(a + bc) \approx \frac{\pm 2\pi i}{\log\lambda}.$$

(Note that the choice of sign corresponds to a choice of λ or λ^{-1}.) Since 1 and c are linearly independent over $\mathbb{R}$, the values of (a, b) vary over an open neighborhood of ∞ as λ varies over a neighborhood of 1. Since $\mathrm{Tr}(H(\alpha)) = \lambda + \lambda^{-1}$ varies over a neighborhood of 2 (up to sign) in the image of $\mathrm{Tr} : \mathrm{Def}(M) \to \mathbb{C}^k$, we have shown that the surgery coefficients for the $M_{d_1,\dots,d_k}$ possessing hyperbolic structures vary over an open neighborhood of ∞ in each component. □

Example. The complement of the Borromean rings has a complete hyperbolic structure. However, if the trivial surgery with coefficients $(1, 0)$ is performed on one component, the others are unlinked. (In other words, $M_{(1,0),\infty,\infty}$ is S^3 minus two unlinked circles.) The manifold $M_{(1,0),x,y}$ (where M is S^3 minus the Borromean rings) is then a connected sum of lens spaces if x, y are primitive elements of $H_1(T_i^2, \mathbb{Z})$ so it cannot have a hyperbolic structure. Thus it may often happen that an infinite number of non-hyperbolic manifolds can be obtained by surgery from a hyperbolic one. However, the theorem does imply that if a finite number of integral pairs of coefficients is excluded from *each boundary component*, then all remaining three-manifolds obtained by Dehn surgery on M are also hyperbolic.

5.39

5.9. A Proof of Mostow's Theorem.

This section is devoted to a proof of Mostow's Theorem for closed hyperbolic n-manifolds, $n \geq 3$. The proof will be sketchy where it seems to require analysis. With a knowledge of the structure of the ends in the noncompact, complete case, this proof extends to the case of a manifold of finite total volume; we omit details. The outline of this proof is Mostow's.

Given two closed hyperbolic manifolds M_1 and M_2, together with an isomorphism of their fundamental groups, there is a homotopy equivalence inducing the isomorphism since M_1 and M_2 are $K(\pi, 1)$'s. In other words, there are maps $f_1 : M_1 \to M_2$ and $f_2 : M_2 \to M_1$ such that $f_1 \circ f_2$ and $f_2 \circ f_1$ are homotopic to the identity. Denote lifts of f_1, f_2 to the universal cover H^n by $\tilde{f}_1, \tilde{f}_2$ and assume $\tilde{f}_1 \circ \tilde{f}_2$ and $\tilde{f}_2 \circ \tilde{f}_1$ are equivariantly homotopic to the identity.

The first step in the proof is to construct a correspondence between the spheres at infinity of H^n which extends $\tilde{f}_1$ and $\tilde{f}_2$.

Definition. A map $g : X \to Y$ between metric spaces is a *pseudo-isometry* if there are constants c_1, c_2 such that $c_1^{-1} d(x_1, x_2) - c_2 \leq d(gx_1, gx_2) \leq c_1 d(x_1, x_2)$ for all $x_1, x_2 \in X$.

Lemma 5.9.1. *$\tilde{f}_1, \tilde{f}_2$ can be chosen to be pseudo-isometries.*

Proof. Make f_1 and f_2 simplicial. Then since M_1 and M_2 are compact, f_1 and 5.40
f_2 are Lipschitz and lift to $\tilde{f}_1$ and $\tilde{f}_2$ which are Lipschitz with the same coefficient. It follows immediately that there is a constant c_1 so that $d(\tilde{f}_i x_1, \tilde{f}_i x_2) \leq c_1 d(x_1, x_2)$ for $i = 1, 2$ and all $x_1, x_2 \in H^n$.

If $x_i = \tilde{f}_1 y_i$, then this inequality implies that

$$d(\tilde{f}_2 \circ \tilde{f}_1(y_1),\ \tilde{f}_2 \circ \tilde{f}_1(y_2)) \leq c_1 d(\tilde{f}_1 y_1, \tilde{f}_1 y_2).$$

However, since M_1 is compact, $\tilde{f}_2 \circ \tilde{f}_1$ is homotopic to the identity by a homotopy that moves every point a distance less than some constant b. It follows that

$$d(y_1, y_2) - 2b \leq d(\tilde{f}_2 \circ \tilde{f}_1 y_1, \tilde{f}_2 \circ \tilde{f}_1 y_2),$$

from which the lower bound $c_1^{-1} d(y_1, y_2) - c_2 \leq d(\tilde{f}_1 y_1, \tilde{f}_1 y_2)$ follows. □

Using this lemma it is possible to associate a unique geodesic with the image of a geodesic.

Proposition 5.9.2. *For any geodesic $g \subset H^n$ there is a unique geodesic h such that $f_1(g)$ stays in a bounded neighborhood of h.*

Proof. If j is any geodesic in H^n, let $N_s(j)$ be the neighborhood of radius s about j. We will see first that if s is large enough there is an upper bound to the length of any bounded component of $g - \left(\tilde{f}_1^{-1}(N_s(j))\right)$, for any j. In fact, the perpendicular projection from $H^n - N_s(j)$ to j decreases every distance by at least a factor of $1/\cosh s$, so any long path in $H^n - N_s(j)$ with endpoints on $\partial N_s(j)$ can be replaced by a much shorter path consisting of two segments perpendicular to j, together with a segment of j. 5.41

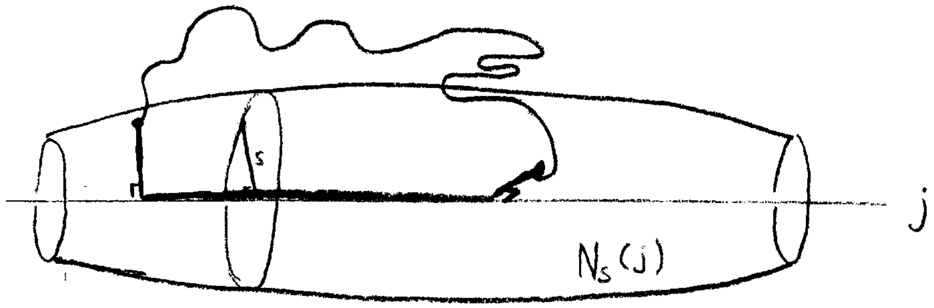

When this fact is applied to a line j joining distant points p_1 and p_2 on $\tilde{f}_1(g)$, it follows that the segment of g between p_1 and p_2 must intersect each plane perpendicular to j a bounded distance from j. It follows immediately that there is a limit line h to such lines j as p_1 and p_2 go to $+\infty$ and $-\infty$ on $\tilde{f}_1(g)$, and that $\tilde{f}_1(g)$ remains a bounded distance from h. Since no two lines in H^n remain a bounded distance apart, h is unique. $\square$

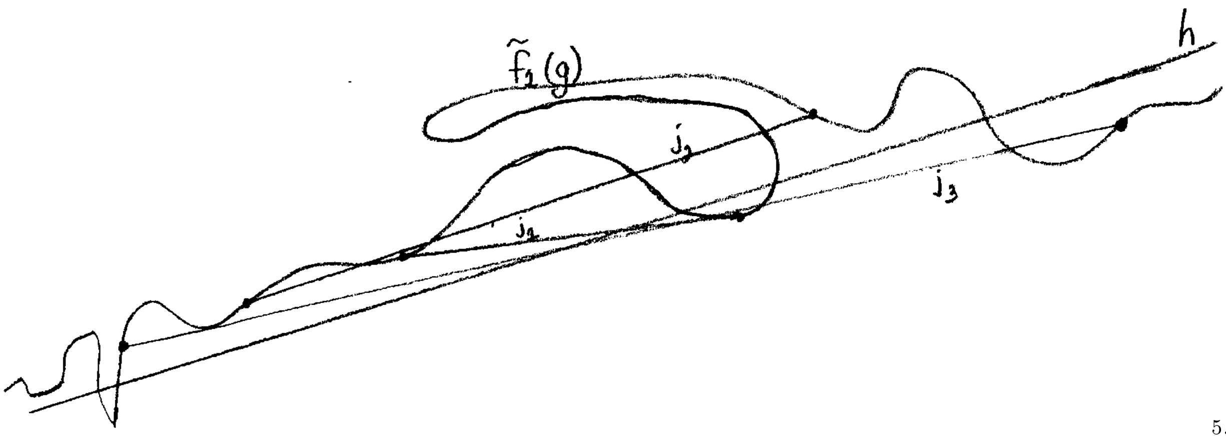

5.42

Corollary 5.9.3. *$\tilde{f}_1 : H^n \to H^n$ induces a one-to-one correspondence between the spheres at infinity.*

Proof. There is a one-to-one correspondence between points on the sphere at infinity and equivalence classes of directed geodesics, two geodesics being equivalent if they are parallel, or asymptotic in positive time. The correspondence of 5.9.2 between geodesics in $\tilde{M}_1$ and geodesics in $\tilde{M}_2$ obviously preserves this relation of parallelism, so it induces a map on the sphere at infinity. This map is one-to-one since any two distinct points in the sphere at infinity are joined by a geodesic, hence must be taken to the two ends of a geodesic. □

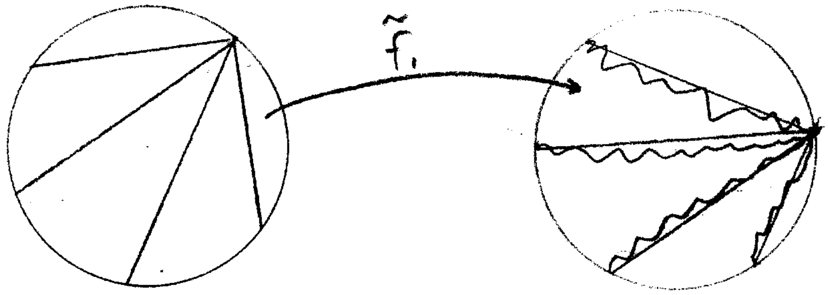

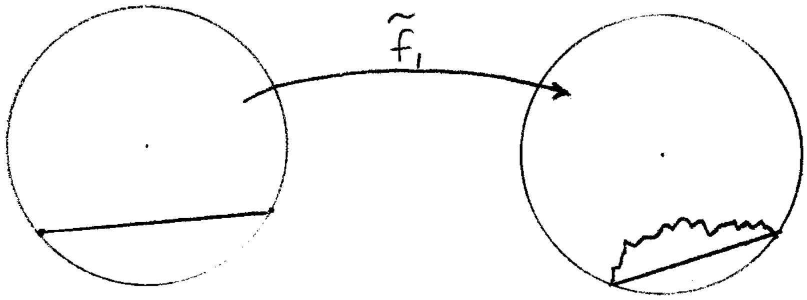

5.43

The next step in the proof of Mostow's Theorem is to show that the extension of $\tilde{f}_1$ to the sphere at infinity S^{n-1}_∞ is continuous. One way to prove this is by citing Brouwer's Theorem that every function is continuous. Since this way of thinking is not universally accepted (though it is valid in the current situation), we will give another proof, which will also show that f is quasi-conformal at S^{n-1}_∞. A basis for the neighborhoods of a point $x \in S^{n-1}_\infty$ is the set of disks with center x. The boundaries of the disks are $(n-2)$-spheres which correspond to hyperplanes in H^2 (i.e., to $(n-1)$-spheres perpendicular to S^{n-1}_∞ whose intersections with S^{n-1}_∞ are the $(n-2)$-spheres).

For any geodesic g in $\tilde{M}_1$, let $\phi(g)$ be the geodesic in $\tilde{M}_2$ which remains a bounded distance from $\tilde{f}_1(g)$.

LEMMA 5.9.4. *There is a constant c such that, for any hyperplane P in H^n and any geodesic g perpendicular to P, the projection of $\tilde{f}_1(P)$ onto $\phi(g)$ has diameter $\leq c$.*

PROOF. Let x be the point of intersection of P and g and let l be a geodesic ray based at x. Then there is a unique geodesic l_1 which is parallel to l in one direction and to a fixed end of g in the other. Let A denote the shortest arc between x and l_1. It has length d, where d is a fixed contrast $(= \operatorname{arc}\cosh\sqrt{2})$. 5.44

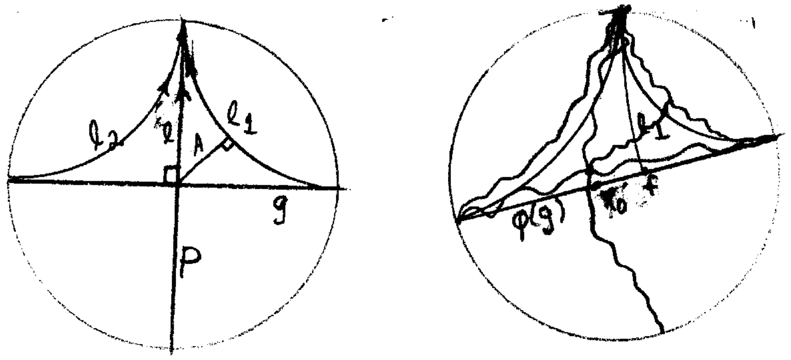

The idea of the proof is to consider the image of this picture under $\tilde{f}_l$. Let $\phi(l), \phi(l_1), \phi(g)$ denote the geodesics that remain a bounded distance from l, l_1 and g respectively. Since ϕ preserves parallelism $\phi(l)$ and $\phi(l_1)$ are parallel. Let $l^{\perp}$ denote the geodesic from the endpoint on S^{n-1}_{∞} of $\phi(l)$ which is perpendicular to $\phi(g)$. Also let x_0 be the point on $\phi(g)$ nearest to $\tilde{f}_l(x)$.

Since $\tilde{f}_l(x)$ is a pseudo-isometry the length of $\tilde{f}_l(A)$ is at most $c_1 d$ where c_1 is a fixed constant. Since $\phi(l_1)$ and $\phi(g)$ are less than distance s (for a fixed constant s) from $\tilde{f}_l(l_1)$ and $\tilde{f}_l(g)$ respectively, it follows that x_0 is distance less than $C_1 d + 2s = \bar{d}$ from $\phi(l_1)$. This implies that the foot of $l^{\perp}$ (i.e., $l^{\perp} \cap \phi(g)$) lies distance less than $\bar{d}$ to one side of x_0. By considering the geodesic l_2 which is parallel to l and to the other end of g, it follows that f lies a distance less than $\bar{d}$ from x_0. 5.45

Now consider any point $y \in P$. Let m be any line through y. The endpoints of $\phi(m)$ project to points on $\phi(g)$ within a distance $\bar{d}$ of x_0; since $\tilde{f}_l(y)$ is within a distance s of $\phi(m)$, it follows that y projects to a point not farther than $\bar{d} + s$ from x_0. $\square$

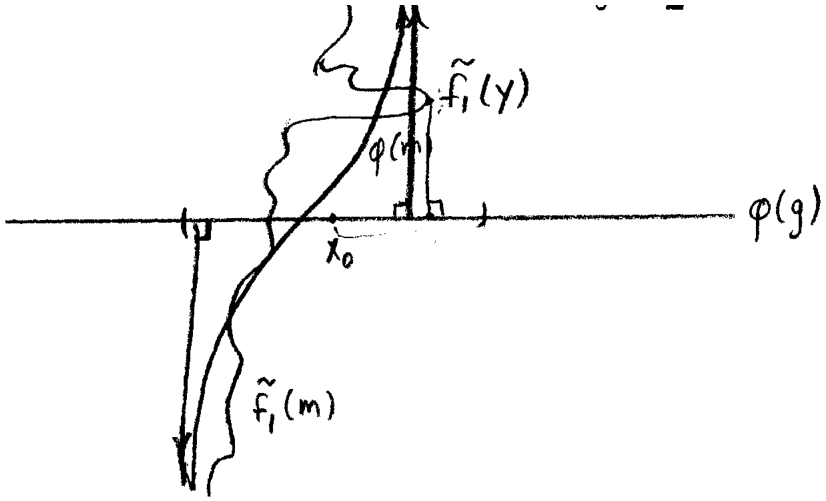

COROLLARY 5.9.5. *The extension of $\tilde{f}_l$ to S_∞^{n-1} is continuous.*

PROOF. For any point $y \in S_\infty^{n-1}$, consider a directed geodesic g bending toward y, and define $\tilde{f}_l(y)$ to be the endpoint of $\phi(g)$. The half-spaces bounded by hyperplanes perpendicular to $\phi(g)$ form a neighborhood basis for $\tilde{f}_l(y)$. For any such half-space H, there is a point $x \in g$ such that the projection of $\tilde{f}_l(x)$ to $\phi(g)$ is a distance $> C$ from ∂H. Then the neighborhood of y bounded by the hyperplane through x perpendicular to g is mapped within H. $\square$

Below it will be necessary to use the concept of quasi-conformality. If f is a homeomorphism of a metric space X to itself, f is K-*quasi-conformal* if and only if for all $z \in X$ 5.46

$$\lim_{r\to 0} \frac{\sup_{x,y\in S_r(z)} d\,(f(x), f(y))}{\inf_{x,y\in S_r(z)} d\,(f(x), f(y))} \le K$$

where $S_r(Z)$ is the sphere of radius r around Z. K measures the deviation of f from conformality, is equal to 1 if f is conformal, and is unchanged under composition with a conformal map. f is called *quasi-conformal* if it is K-quasi-conformal for some K.

COROLLARY 5.9.6. *$\tilde{f}_l$ is quasi-conformal at S_∞^{n-1}.*

PROOF. Use the upper half-space model for H^n since it is conformally equivalent to the ball model and suppose x and $\tilde{f}_l x$ are the origin since translation to the origin is also conformal. Then consider any hyperplane P perpendicular to the geodesic g from 0 to the point at infinity. By Lemma 5.9.4 there is a bound, depending only on $\tilde{f}_l$, to the diameter of the projection of $\tilde{f}_l(P)$ onto $\phi(g) = g$. Therefore, there are hyperplanes P_1, P_2 perpendicular to g contained in and containing $\tilde{f}_l(P)$ and the distance (along g) between P_1 and P_2 is uniformly bounded for all planes P.

But this distance equals $\log r, r > 1$, where r is the ratio of the radii of the $n-2$ spheres

$$S^{n-2}_{p_1}, \quad S^{n-2}_{p_2}$$

in S^{n-1}_∞ corresponding to P_1 and P_2. The image of the $n-2$ sphere S^{n-2}_P corresponding 5.47
to P lies between $S^{n-2}_{p_2}$ and $S^{n-2}_{p_1}$ so that r is an upper bound for the ratio of maximum to minimum distances on

$$\tilde{f}_l(S^{n-2}_p).$$

Since $\log r$ is uniformly bounded above, so is r and $\tilde{f}_l$ is quasi-conformal. $\square$

Corollary 5.9.6 was first proved by Gehring for dimension $n = 3$, and generalized to higher dimensions by Mostow.

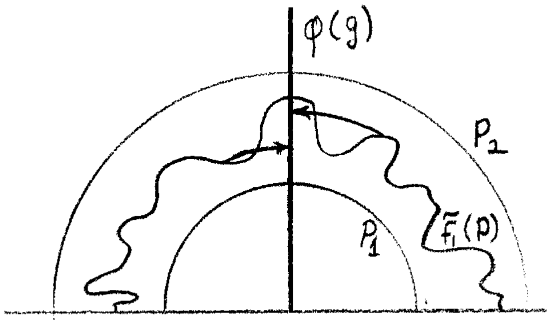

At this point, it is necessary to invoke a theorem from analysis (see Bers).

THEOREM 5.9.7. *A quasi-conformal map of an $n-1$-manifold, $n > 2$, has a derivative almost everywhere (= a.e.).*

REMARK. It is at this stage that the proof of Mostow's Theorem fails for $n = 2$. 5.48
The proof works to show that $\tilde{f}_l$ extends quasi-conformally to the sphere at infinity, S^1_∞, but for a one-manifold this does not imply much.

Consider $\tilde{f}_l : S^{n-1}_\infty \to S^{n-1}_\infty$; by theorem 5.9.7 $d\tilde{f}_l$ exists a.e. At any point x where the derivative exists, the linear map $d\tilde{f}_l(x)$ takes a sphere around the origin to an ellipsoid. Let $\lambda_1, \dots, \lambda_{n-1}$ be the lengths of the axes of the ellipsoid. If we normalize so that $\lambda_1 \cdot \lambda_2 \cdots \lambda_{n-1} = 1$, then the λ_i are conformal invariants. In particular denote the maximum ratio of the λ_i's at x by $e(x)$, the eccentricity of $\tilde{f}_l$ at x. Note that if $\tilde{f}_l$ is K-quasi-conformal, the supremum of $e(x)$, $x \in S^{n-1}_\infty$, is K. Since $\pi_1 M_1$ acts on S^{n-1}_∞ conformally and e is invariant under conformal maps, e is a measurable, $\pi_1 M_1$ invariant function on S^{n-1}_∞. However, such functions are very simple because of the following theorem:

THEOREM 5.9.8. *For a closed, hyperbolic n-manifold M, $\pi_1 M$ acts ergodically on S^{n-1}_∞, i.e., every measurable, invariant set has zero measure or full measure.*

Corollary 5.9.9. *e is constant a.e.*

Proof. Any level set of e is a measurable, invariant set so precisely one has full measure. □

In fact more is true: 5.49

Theorem 5.9.10. *$\pi_1(M)$ acts ergodically on $S^{n-1}_\infty \times S^{n-1}_\infty$.*

Remark. This theorem is equivalent to the fact that the geodesic flow of M is ergodic since pairs of distinct points on S^{n-1}_∞ are in a one-to-one correspondence to geodesics in H^n (whose endpoints are those points).

From Corollary 5.9.9 e is equal a.e. to a constant K, and if the derivative of $\tilde{f}_l$ is not conformal, $K \neq 1$.

Consider the case $n = 3$. The direction of maximum "stretch" of df defines a measurable line field l on S^2_∞. Then for any two points $x, y \in S^2_\infty$ it is possible to parallel translate the line $l(x)$ along the geodesic between x and y to y and compute the angle between the translation of $l(x)$ and $l(y)$. This defines a measurable $\pi_1 M$-invariant function on $S^2_\infty \times S^2_\infty$. By theorem 5.9.10 it must be constant a.e. In other words l is determined by its "value" at one point. It is not hard to see that this is impossible.

For example, the line field determined by a line at x agrees with the line field below a.e. However, any line field determined by its "value" at y will have the same form and will be incompatible. 5.50

The precise argument is easy, but slightly more subtle, since l is defined only a.e.

The case $n > 3$ is similar.

Now one must again invoke the theorem, from analysis, that a quasi-conformal map whose derivative is conformal a.e. is conformal in the usual sense; it is a sphere-preserving map of S^{n-1}_∞, so it extends to an isometry I of H^n. The isometry I conjugates the action of $\pi_1 M_1$ to the action of $\pi_1 M_2$, completing the proof of Mostow's Theorem. □

5.10. A decomposition of complete hyperbolic manifolds.

5.51

Let M be any complete hyperbolic manifold (possibly with infinite volume). For $\epsilon > 0$, we will study the decomposition $M = M_{(0,\epsilon]} \cup M_{[\epsilon,\infty)'}$ where $M_{(0,\epsilon]}$ consists of those points in M through which there is a non-trivial closed loop of length $\leq \epsilon$, and $M_{[\epsilon,\infty)}$ consists of those points through which every non-trivial loop has length $\geq \epsilon$.

In order to understand the geometry of $M_{(0,\epsilon]}$, we pass to the universal cover $\tilde{M} = H^n$. For any discrete group Γ of isometries of H^n and any $x \in H^n$ let $\Gamma_\epsilon(x)$ be the subgroup generated by all elements of Γ which move x a distance $\leq \epsilon$, and let

$\Gamma'_\epsilon(x) \subset \Gamma_\epsilon(x)$ be the subgroup consisting of elements whose derivative is also ϵ-close to the identity.

LEMMA 5.10.1 (The Margulis Lemma). *For every dimension n there is an $\epsilon > 0$ such that for every discrete group Γ of isometries of H^n and for every $x \in H^n$, $\Gamma'_\epsilon(x)$ is abelian and $\Gamma_\epsilon(x)$ has an abelian subgroup of finite index.*

REMARK. This proposition is much more general than stated; if "abelian" is replaced by "nilpotent," it applies in general to discrete groups of isometries of Riemannian manifolds with bounded curvature. The proof of the general statement is essentially the same.

PROOF. In any Lie group G, since the commutator map $[\ ,\ ] : G \times G \to G$ has derivative 0 at $(1, 1)$, it follows that the size of the commutator of two small elements 5.52
is bounded above by some constant times the product of their sizes. Hence, if Γ'_ϵ is any discrete subgroup of G generated by small elements, it follows immediately that the lower central series $\Gamma'_\epsilon \supset [\Gamma'_\epsilon, \Gamma'_\epsilon] \supset [\Gamma'_\epsilon, [\Gamma'_\epsilon, \Gamma'_\epsilon]], \ldots$ is finite (since there is a lower bound to the size of elements of Γ'_ϵ). In other words, Γ'_ϵ is nilpotent. When G is the group of isometries of hyperbolic space, it is not hard to see (by considering, for instance, the geometric classification of isometries) that this implies Γ'_ϵ is actually abelian.

To guarantee that $\Gamma_\epsilon(x)$ has an abelian subgroup of finite index, the idea is first to find an ϵ_1 such that $\Gamma'_{\epsilon_1}(x)$ is always abelian, and then choose ϵ many times smaller than ϵ_1, so the product of generators of $\Gamma_\epsilon(x)$ will lie in $\Gamma'_{\epsilon_1}(x)$. Here is a precise recipe:

Let N be large enough that any collection of elements of $O(n)$ with cardinality $\geq N$ contains at least one pair separated by a distance not more than $\epsilon_{1/3}$.

Choose $\epsilon_2 \leq \epsilon_{1/3}$ so that for any pair of isometries ϕ_1 and ϕ_2 of H^n which translate a point x a distance $\leq \epsilon_2$, the derivative at x of $\phi_1 \circ \phi_2$ (parallel translated back to x) is estimated within $\epsilon_{1/6}$ by the product of the derivatives at x of ϕ_1 and ϕ_2 (parallel translated back to x).

Now let $\epsilon = \epsilon_{2/2N}$, so that a product of $2N$ isometries, each translating x a distance $\leq \epsilon$, translates x a distance $\leq \epsilon_2$. Let $g_1, \ldots, g_k$ be the set of elements of Γ 5.53
which move x a distance $\leq \epsilon$; they generate $\Gamma_\epsilon(x)$. Consider the cosets $\gamma\,\Gamma'_{\epsilon_1}(x)$, where $\gamma \in \Gamma_\epsilon(x)$; the claim is that they are all represented by γ's which are words of length $< N$ in the generators $g_1, \ldots, g_k$. In fact, if $\gamma = g_{i_1} \cdot \ldots \cdot g_{i_l}$ is any word of length $\geq N$ in the g_i's, it can be written $\gamma = \alpha \cdot \epsilon' \cdot \beta, (\alpha, \epsilon', \beta \neq 1)$ where $\epsilon' \cdot \beta$ has length $\leq N$, and the derivative of ϵ' is within $\epsilon_{1/3}$ of 1. It follows that $(\alpha\beta)^{-1} \cdot (\alpha\epsilon'\beta) = \beta^{-1}\epsilon'\beta$ is in $\Gamma'_{\epsilon_1}(x)$; hence the coset $\gamma\Gamma'_{\epsilon_1}(x) = (\alpha\beta)\Gamma'_{\epsilon_1}(x)$. By induction, the claim is verified. Thus, the abelian group $\Gamma'_{\epsilon_1}(x)$ has finite index in the group generated by $\Gamma_\epsilon(x)$ and $\Gamma'_{\epsilon_1}(x)$, so $\Gamma'_{\epsilon_1}(x) \cap \Gamma_\epsilon(x)$ with finite index. $\square$

Examples. When $n = 3$, the only possibilities for discrete abelian groups are $\mathbb{Z}$ (acting hyperbolically or parabolically), $\mathbb{Z} \times \mathbb{Z}$ (acting parabolically, conjugate to a group of Euclidean translations of the upper half-space model), $\mathbb{Z} \times \mathbb{Z}_n$ (acting as a group of translations and rotations of some axis), and $\mathbb{Z}_2 \times \mathbb{Z}_2$ (acting by 180° rotations about three orthogonal axes). The last example of course cannot occur as $\Gamma'_\epsilon(x)$. Similarly, when ϵ is small compared to $1/n$, $\mathbb{Z} \times \mathbb{Z}_n$ cannot occur as $\Gamma'_\epsilon(x)$.

Any discrete group Γ of isometries of Euclidean space E^{n-1} acts as a group of isometries of H^n, via the upper half-space model. 5.54

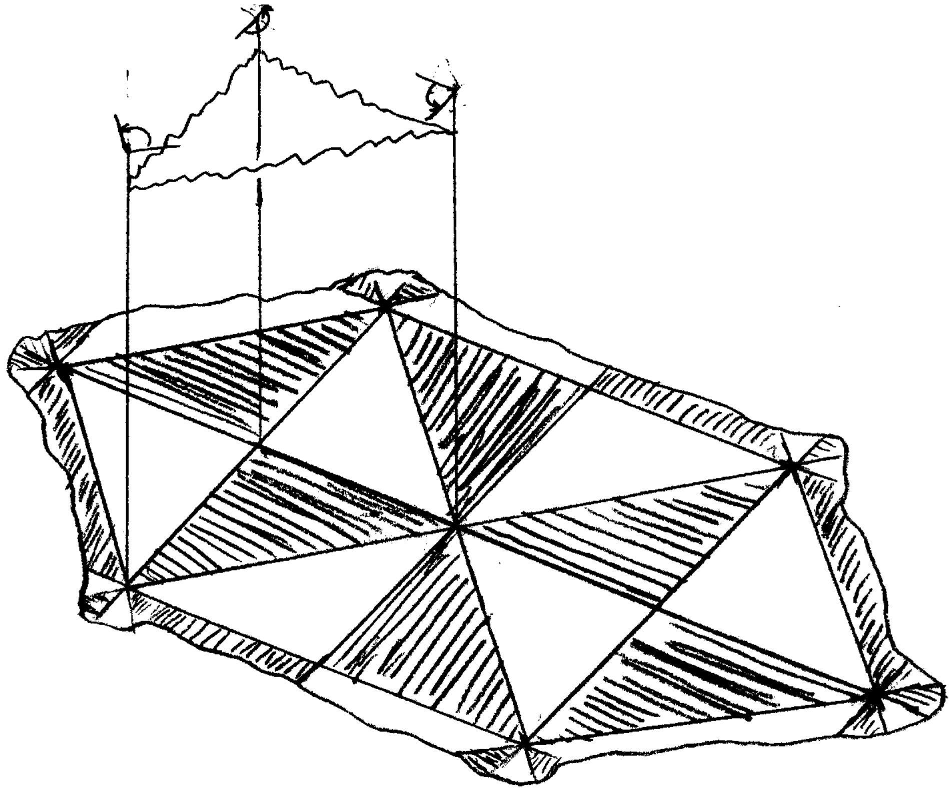

For any x sufficiently high (in the upper half space model), $\Gamma_\epsilon(x) = \Gamma$. Thus, 5.10.1 contains as a special case one of the Bieberbach theorems, that Γ contains an abelian subgroup of finite index. Conversely, when $\Gamma_\epsilon(x) \cap \Gamma'_{\epsilon_1}(x)$ is parabolic, $\Gamma_\epsilon(x)$ must be a Bieberbach group. To see this, note that if $\Gamma_\epsilon(x)$ contained any hyperbolic element γ, no power of γ could lie in $\Gamma'_{\epsilon_1}(x)$, a contradiction. Hence, $\Gamma_\epsilon(x)$ must consist of parabolic and elliptic elements with a common fixed point p at ∞, so it acts as a group of isometries on any horosphere centered at p.

If $\Gamma_\epsilon(x) \cap \Gamma'_{\epsilon_1}(x)$ is not parabolic, it must act as a group of translations and rotations of some axis a. Since it is discrete, it contains $\mathbb{Z}$ with finite index (provided $\Gamma_\epsilon(x)$ is infinite). It easily follows that $\Gamma_\epsilon(x)$ is either the product of some finite

FIGURE 1. The infinite dihedral group acting on H^3.

subgroup F of $O(n-1)$ (acting as rotations about a) with $\mathbb{Z}$, or it is the semidirect 5.55
product of such an F with the infinite dihedral group, $\mathbb{Z}/2 * \mathbb{Z}/2$.

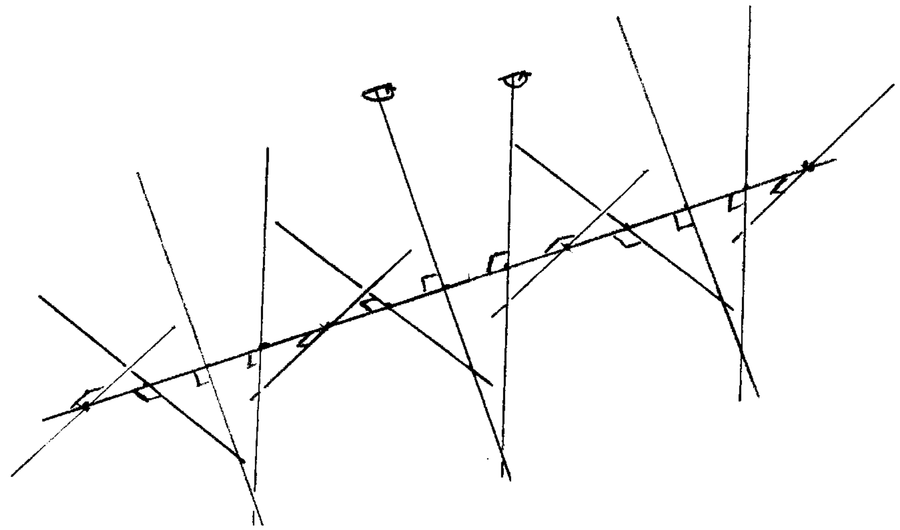

For any set $S \subset H^n$, let $B_r(S) = \{x \in H^n \,|\, d(x, S) \leq r\}$.

COROLLARY 5.10.2. *There is an $\epsilon > 0$ such that for any complete oriented hyperbolic three-manifold M, each component of $M_{(0,\epsilon]}$ is either*

(1) *a horoball modulo $\mathbb{Z}$ or $\mathbb{Z} \oplus \mathbb{Z}$, or*
(2) *$B_r(g)$ modulo $\mathbb{Z}$, where g is a geodesic.*

The degenerate case $r = 0$ may occur.

PROOF. Suppose $x \in M_{(0,\epsilon]}$. Let $\tilde{x} \in H^3$ be any point which projects to x. There is some covering translation γ which moves x a distance $\leq \epsilon$. If γ is hyperbolic, let a be its axis. All rotations around a, translations along a, and uniform contractions of hyperbolic space along orthogonals to a commute with γ. It follows that $\tilde{M}_{(0,\epsilon]}$ 5.56
contains $B_r(a)$, where $r = d(a, x)$, since γ moves any point in $B_r(a)$ a distance $\leq \epsilon$. Similarly, if γ is parabolic with fixed point p at ∞, $\tilde{M}_{(0,\epsilon]}$ contains a horoball about p passing through x. Hence $M_{(0,\epsilon]}$ is a union of horoballs and solid cylinders $B_r(a)$. Whenever two of these are not disjoint, they correspond to two covering transformations γ_1 and γ_2 which move some point x a distance $\leq \epsilon$; γ_1 and γ_2 must commute (using 5.10.1), so the corresponding horoballs or solid cylinders must be concentric, and 5.10.2 follows. $\square$

5.11. Complete hyperbolic manifolds with bounded volume.

It is easy now to describe the structure of a complete hyperbolic manifold with finite volume; for simplicity we stick to the case $n = 3$.

PROPOSITION 5.11.1. *A complete oriented hyperbolic three-manifold with finite volume is the union of a compact submanifold* (*bounded by tori*) *and a finite collection of horoballs modulo* $\mathbb{Z} \oplus \mathbb{Z}$ *actions.*

PROOF. $M_{[\epsilon,\infty)}$ must be compact, for otherwise there would be an infinite sequence of points in $M_{[\epsilon,\infty)}$ pairwise separated by at least ϵ. This would give a sequence of hyperbolic $\epsilon/2$ balls disjointly embedded in $M_{[\epsilon,\infty)}$, which has finite volume. $M_{(0,\epsilon]}$ must have finitely many components (since its boundary is compact). The proposition is obtained by lumping all compact components of $M_{(0,\epsilon]}$ with $M_{[\epsilon,\infty)}$. □

5.57

With somewhat more effort, we obtain Jørgensen's theorem, which beautifully describes the structure of the set of all complete hyperbolic three-manifolds with volume bounded by a constant C:

THEOREM 5.11.2 (Jørgensen's theorem [first version]). *Let* $C > 0$ *be any constant. Among all complete hyperbolic three-manifolds with volume* $\leq C$*, there are only finitely many homeomorphism types of* $M_{[\epsilon,\infty)}$*. In other words, there is a link* L_c *in* S^3 *such that every complete hyperbolic manifold with volume* $\leq C$ *is obtained by Dehn surgery along* L_C*.* (*The limiting case of deleting components of* L_C *to obtain a non-compact manifold is permitted.*)

PROOF. Let V be any maximal subset of $M_{[\epsilon,\infty)}$ having the property that no two elements of V have distance $\leq \epsilon/2$. The balls of radius $\epsilon/4$ about elements of V are embedded; since their total volume is $\leq C$, this gives an upper bound to the cardinality of V. The maximality of V is equivalent to the property that the balls of radius $\epsilon/2$ about V cover.

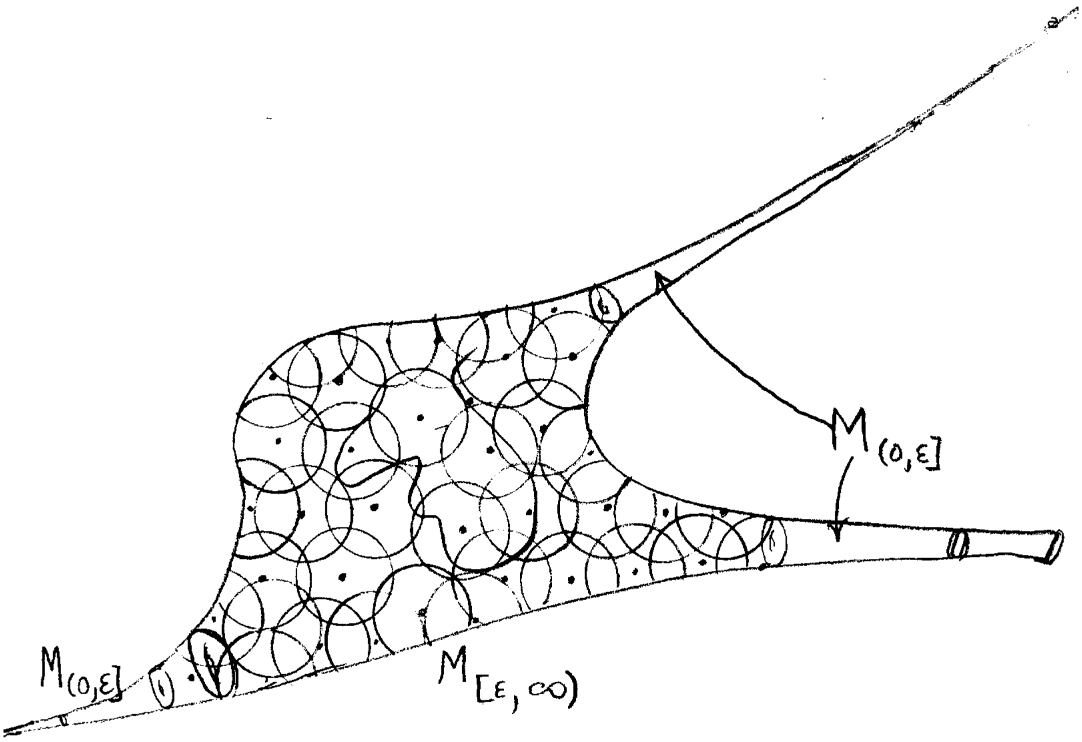

The combinatorial pattern of intersections of this set of $\epsilon/2$-balls determines $M_{[\epsilon,\infty)}$ up to diffeomorphism. There are only finitely many possibilities. (Alternatively a triangulation of $M_{[\epsilon,\infty)}$ with vertex set V can be constructed as follows. First, form a cell division of $M_{[\epsilon,\infty)}$ whose cells are indexed by V, associating to each $v \in V$ the subset of $M_{[\epsilon,\infty)}$ consisting of $x \in M_{[\epsilon,\infty)}$ such that $d(x,v) < d(x,v')$ for all $v' \in V$.

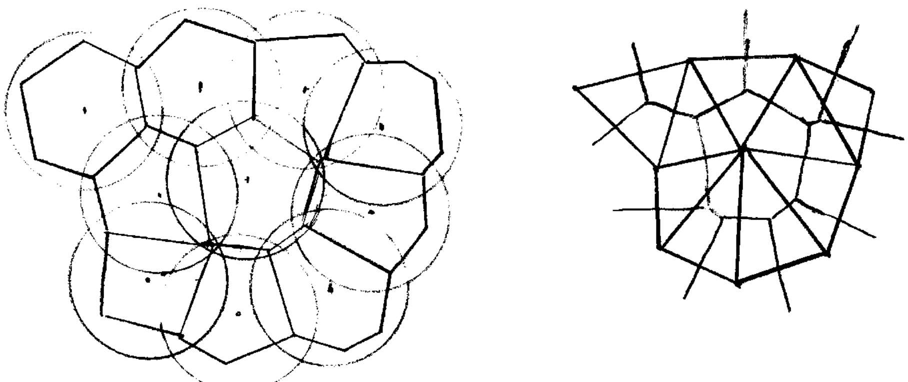

If

If V is in general position, faces of the cells meet at most four at a time. (The dual cell division is a triangulation.)

Any two hyperbolic manifolds M and N such that $M_{[\epsilon,\infty)} = N_{[\epsilon,\infty)}$ can be obtained from one another by Dehn surgery. All manifolds with volume $\leq C$ can therefore be obtained from a fixed finite set of manifolds by Dehn surgery on a fixed link in each manifold. Each member of this set can be obtained by Dehn surgery on some link in S^3, so all manifolds with volume $\leq C$ can be obtained from S^3 by Dehn surgery on the disjoint union of all the relevant links. □ 5.59

The full version of Jørgensen's Theorem involves the geometry as well as the topology of hyperbolic manifolds. The geometry of the manifold $M_{[\epsilon,\infty)}$ completely determines the geometry and topology of M itself, so an interesting statement comparing the geometry of $M_{[\epsilon,\infty)}$'s must involve the approximate geometric structure. Thus, if M and N are complete hyperbolic manifolds of finite volume, Jørgensen defines M to be *geometrically near* N if for some small ϵ, there is a diffeomorphism which is approximately an isometry from the hyperbolic manifold $M_{[\epsilon,\infty)}$ to $N_{[\epsilon,\infty)}$. It would suffice to keep ϵ fixed in this definition, except for the exceptional cases when M and N have closed geodesics with lengths near ϵ. This notion of geometric nearness gives a topology to the set $\mathcal{H}$ of isometry classes of complete hyperbolic manifolds of finite volume. Note that neither coordinate systems nor systems of generators for the fundamental groups have been chosen for these hyperbolic manifolds; the homotopy class of an approximate isometry is arbitrary, in contrast with the definition for Teichmüller space. Mostow's Theorem implies that every closed manifold M in $\mathcal{H}$ is an isolated point, since $M_{[\epsilon,\infty)} = M$ when ϵ is small enough. On the other hand, a manifold in $\mathcal{H}$ with one end or *cusp* is a limit point, by the hyperbolic Dehn surgery theorem 5.9. A manifold with two ends is a limit point of limit points and a manifold with k ends is a k-fold limit point. 5.60

Mostow's Theorem implies more generally that the number of cusps of a geometric limit M of a sequence $\{M_i\}$ of manifolds distinct from M must strictly exceed the lim sup of the number of cusps of M_i. In fact, if ϵ is small enough, $M_{(0,\epsilon]}$ consists only of cusps. The cusps of M_i are contained in $M_{i_{(0,\epsilon]}}$; if all its components are cusps, and if $M_{i_{[\epsilon,\infty)}}$ is diffeomorphic with $M_{[\epsilon,\infty)}$ then M_i is diffeomorphic with M so M_i is isometric with M.

The volume of a hyperbolic manifold gives a function $v : \mathcal{H} \to \mathbb{R}_+$. If two manifolds M and N are geometrically near, then the volumes of $M_{[\epsilon,\infty)}$ and $N_{[\epsilon,\infty)}$ are approximately equal. The volume of a hyperbolic solid torus r_0 centered around a geodesic of length l may be computed as

$$\text{volume (solid torus)} = \int_0^{r_0} \int_0^{2\pi} \int_0^{l} \sinh r \cosh r \, dt \, d\theta \, dr = \pi l \sinh^2 r_0$$

while the area of its boundary is

$$\text{area (torus)} = 2\pi l \sinh r_0 \cosh r_0.$$

Thus we obtain the inequality

$$\frac{\text{area } (\partial \text{ solid torus})}{\text{volume (solid torus)}} = \frac{1}{2}\frac{\sinh r_0}{\cosh r_0} < \frac{1}{2}.$$

The limiting case as $r_0 \to \infty$ can be computed similarly; the ratio is $1/2$. Applying 5.61
this to M, we have

5.11.2. $$\text{volume } (M) \leq \text{ volume } (M_{[\epsilon,\infty)}) + \tfrac{1}{2} \text{ area } (\partial M_{[\epsilon,\infty)}).$$

It follows easily that v is a continuous function on $\mathcal{H}$. Changed this label to 5.11.2a.

5.12. Jørgensen's Theorem.

THEOREM 5.12.1. *The function $v : \mathcal{H} \to \mathbb{R}_+$ is proper. In other words, every sequence in $\mathcal{H}$ with bounded volume has a convergent subsequence. For every C, there is a finite set $M_1, \ldots, M_k$ of complete hyperbolic manifolds with volume $\leq C$ such that all other complete hyperbolic manifolds with volume $\leq C$ are obtained from this set by the process of hyperbolic Dehn surgery* (as in 5.9).

PROOF. Consider a maximal subset of V of $M_{[\epsilon,\infty)}$ having the property that no two elements of V have distance $\leq \epsilon/2$ (as in 5.11.1). Choose a set of isometries of the $\epsilon/2$ balls centered at elements of V with a standard $\epsilon/2$-ball in hyperbolic space. The set of possible gluing maps ranges over a compact subset of $\text{Isom}(H^3)$, so any sequence of gluing maps (where the underlying sequence of manifolds has volume $\leq C$) has a convergent subsequence. It is clear that in the limit, the gluing maps still give a hyperbolic structure on $M_{[\epsilon,\infty)}$, approximately isometric to the limiting $M_{[\epsilon,\infty)}$'s. We must verify that $M_{[\epsilon,\infty)}$ extends to a complete hyperbolic manifold. To see this, note that whenever a complete hyperbolic manifold N has a geodesic which is very short compared to ϵ, the radius of the corresponding solid torus in $N_{(0,\epsilon]}$
becomes large. (Otherwise there would be a short non-trivial curve on $\partial N_{(0,\epsilon]}$—but 5.62
such a curve has length $\geq \epsilon$). Thus, when a sequence $\{M_{i_{[\epsilon,\infty)}}\}$ converges, there are approximate isometries between arbitrarily large balls $B_r(M_{i_{[\epsilon,\infty)}})$ for large i, so in the limit one obtains a complete hyperbolic manifold. This proves that v is a proper function. The rest of §5.12 is merely a restatement of this fact. □

REMARK. Our discussion in §5.10, 5.11 and 5.12 has made no attempt to be numerically efficient. For instance, the proof that there is an ϵ such that $\Gamma_\epsilon(x)$ has an abelian subgroup of finite index gives the impression that ϵ is microscopic. In fact, ϵ can be rather large; see Jørgensen, for a more efficient approach. It would be extremely interesting to have a good estimate for the number of distinct $M_{[\epsilon,\infty)}$'s

Figure eight knot

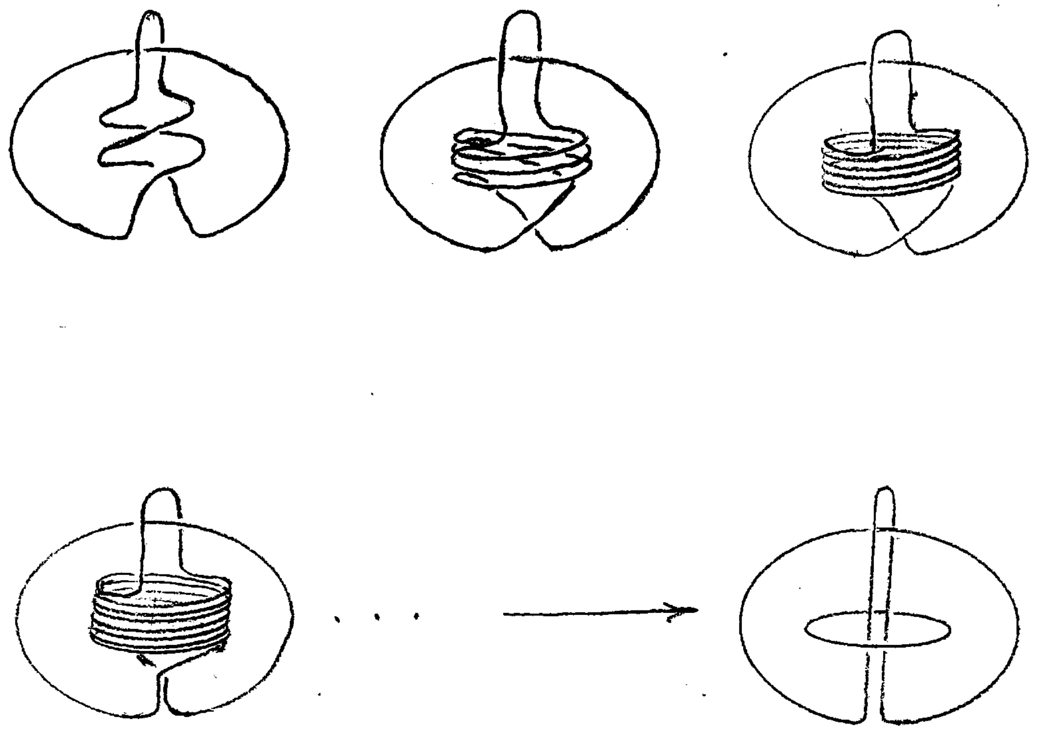

Whitehead Link

where M has volume $\leq C$, and it would be quite exciting to find a practical way of computing them. The development in 5.10, 5.11, and 5.12 is completely inefficient in this regard. Jørgensen's approach is much more explicit and efficient.

EXAMPLE. The sequence of knot complements below are all obtained by Dehn surgery on the Whitehead link, so 5.8.2 implies that all but a finite number possess complete hyperbolic structures. (A computation similar to that of Theorem 4.7 shows that in fact they all possess hyperbolic structures.) This sequence converges, in $\mathcal{H}$, to the Whitehead link complement: 5.63

NOTE. Gromov proved that in dimensions $n \neq 3$, there is only a finite number of complete hyperbolic manifolds with volume less than a given constant. He proved this more generally for negatively curved Riemannian manifolds with curvature varying between two negative constants. His basic method of analysis was to study the injectivity radius

$$\begin{aligned} \operatorname{inj}(x) &= \tfrac{1}{2} \inf\{\text{lengths of non-trivial closed loops through } x\} \\ &= \sup\{r \mid \text{the exponential map is injective on the ball of radius } r \text{ in } T(x)\}. \end{aligned}$$

Basically, in dimensions $n \neq 3$, little can happen in the region M^n_ϵ of M^n where $\operatorname{inj}(x)$ is small. This was the motivation for the approach taken in 5.10, 5.11 and 5.12. Gromov also gave a weaker version of hyperbolic Dehn surgery, 5.8.2: he showed that many of the manifolds obtained by Dehn surgery can be given metrics of negative curvature close to -1.