together along non-trivial arcs in the complementary regions to obtain a new train track which cannot be enlarged. 8.60

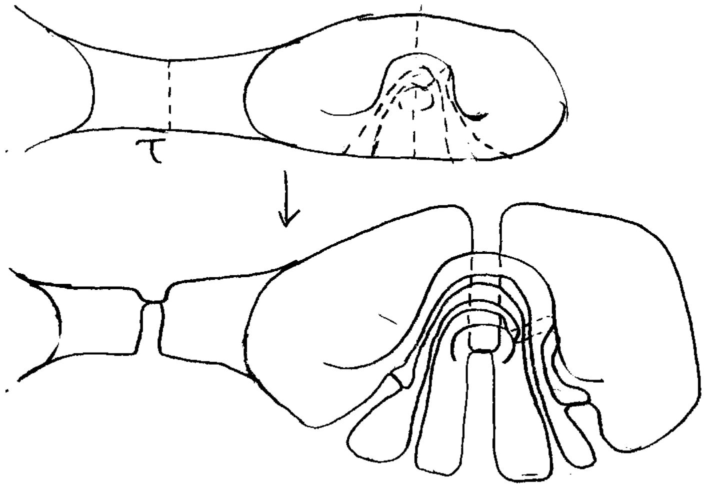

This implies that a neighborhood of (γ, μ) is parametrized by a finite number of real parameters. Thus, $\mathcal{ML}(S)$ is a manifold. Similarly, when S has cusps, $\mathcal{ML}(S)$ is a manifold with boundary $\mathcal{ML}_0(S)$.

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

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

In 8.5, we indicated one proof of the compactness of $\mathcal{GL}(S)$. Another proof goes as follows. First, note that 8.61

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

PROOF. Choose a finite set of transversals $\alpha_1, \ldots, \alpha_k$ which meet every leaf of γ. Suppose there is a sequence $\{l_i\}$ of intervals on leaves of γ such that the total number N_1 of intersection of l_i with the α_j's goes to infinity. Let μ_i be the measure on $\bigcup \alpha_j$ which is $1/N_i$ times the counting measure on $l_i \cap \alpha_j$. The sequence $\{\mu_i\}$ has a subsequence converging (in the weak topology) to a measure μ. It is easy to see that μ is invariant under local projections along leaves of γ, so that it determines a transverse measure.

If there is no such sequence $\{l_i\}$ then every leaf is proper, so the counting measure for any leaf will do. $\square$

We continue with the proof of 8.10.5. Because of the preceding result, the image $\mathcal{I}$ of $\mathcal{PL}(S)$ in $\mathcal{GL}(S)$ intersects the closure of every point of $\mathcal{GL}(S)$. Any collection of open sets which covers $\mathcal{GL}(S)$ has a finite subcollection which covers the compact set $\mathcal{I}$; therefore, it covers all of $\mathcal{GL}(S)$. $\square$

Armed with topology, we return to the question of realizing geodesic laminations. Let $\mathcal{R}_f \subset \mathcal{GL}(S)$ consist of the laminations realizable in the homotopy class of f.

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

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

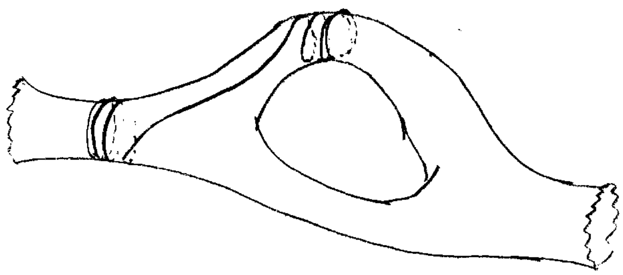

To see this, consider for instance the case of a geodesic g on S whose ends spiral around closed geodesics g_1 and g_2. Lifting f to H^3, we see that the two ends of $\tilde{f}(\tilde{g})$ are asymptotic to geodesics $\tilde{f}(\tilde{g}_1)$ and $\tilde{f}(\tilde{g}_2)$. Then f is homotopic to a map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

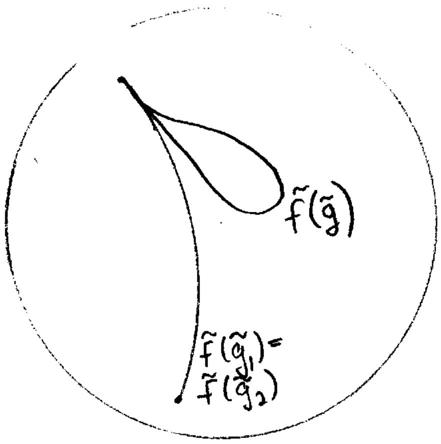

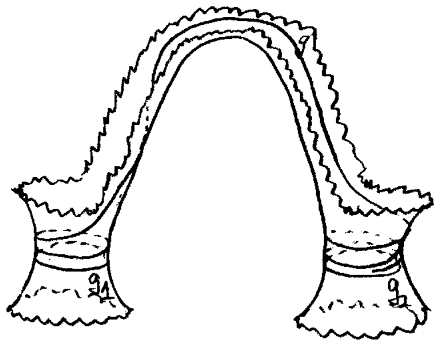

The situation is similar when the ends of g tend toward cusps.

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

Proposition 8.10.7. (a) *Measures supported on finitely many compact or proper geodesics are dense in $\mathcal{ML}$.*

(b) *Geodesic laminations with finitely many leaves are dense in $\mathcal{GL}$.*

(c) *Each end of a non-compact leaf of a geodesic lamination with only finitely many leaves spirals around some closed geodesic, or tends toward a cusp.*

Proof. If τ is any train track and μ is any measure which is positive on each branch, μ can be approximated by measures μ' which are rational on each branch, since μ is subject only to linear equations with integer coefficients. μ' gives rise to geodesic laminations with only finitely many leaves, all compact or proper. This 8.64
proves (a).

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

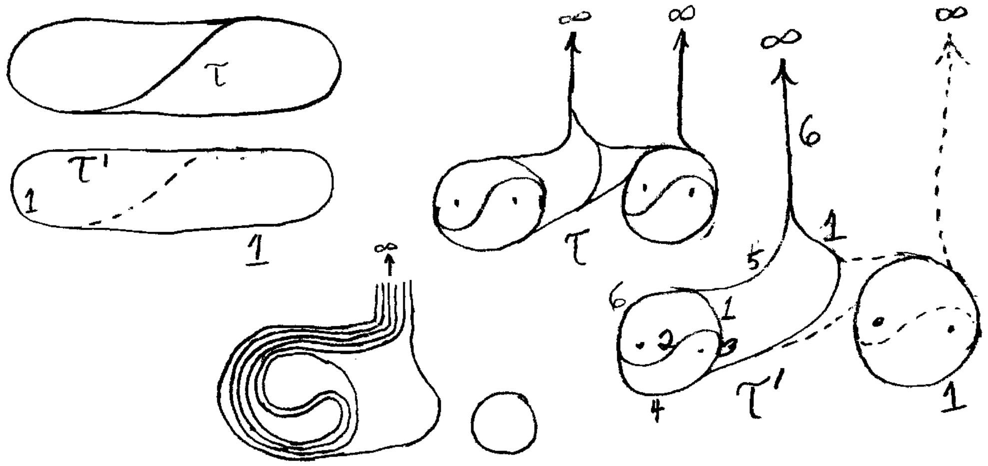

(The reader experienced with toy trains is aware of the subtlety of this question.) There is a measure supported on τ', obtained by choosing a finite set of cyclic and proper paths covering τ' and assigning to a branch b the total number of times these paths traverse. Thus there is a lamination λ' consisting of finitely many compact or proper leaves supported in a narrow corridor about τ'. Now let b be any branch of $\tau - \tau'$. A train starting on b can continue indefinitely, so it must eventually come to τ', in each direction. Add a leaf to λ' representing a shortest path from b to τ'
in each direction; if the two ends meet, make them run along side by side (to avoid 8.65
crossings). When the ends approach τ, make them "merge"—either spiral around a closed leaf, or follow along close to a proper leaf. Continue inductively in this way, adding leaves one by one until you obtain a lamination λ dominated by τ and filling out all the branches. This proves (b).

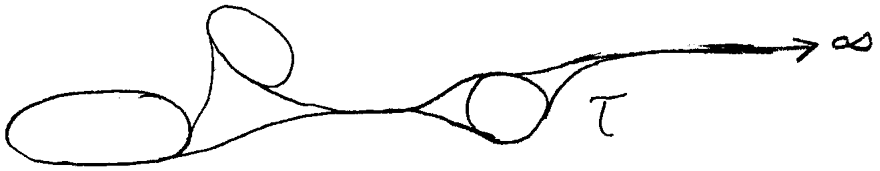

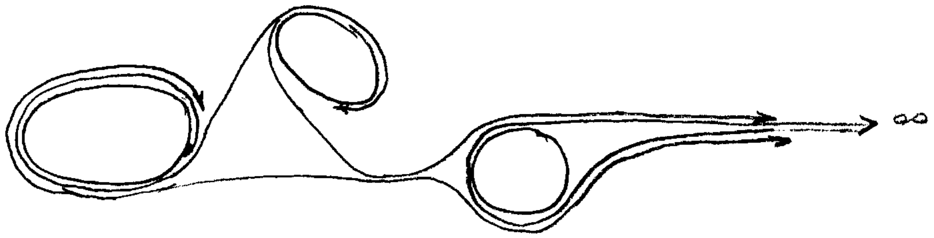

If γ is any geodesic lamination with finitely many (or even countably many) leaves, then the only possible minimal sets are closed leaves; thus each end e of a non-compact must either be a proper end or come arbitrarily near some compact leaf l. By tracing the leaves near l once around l, it is easy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

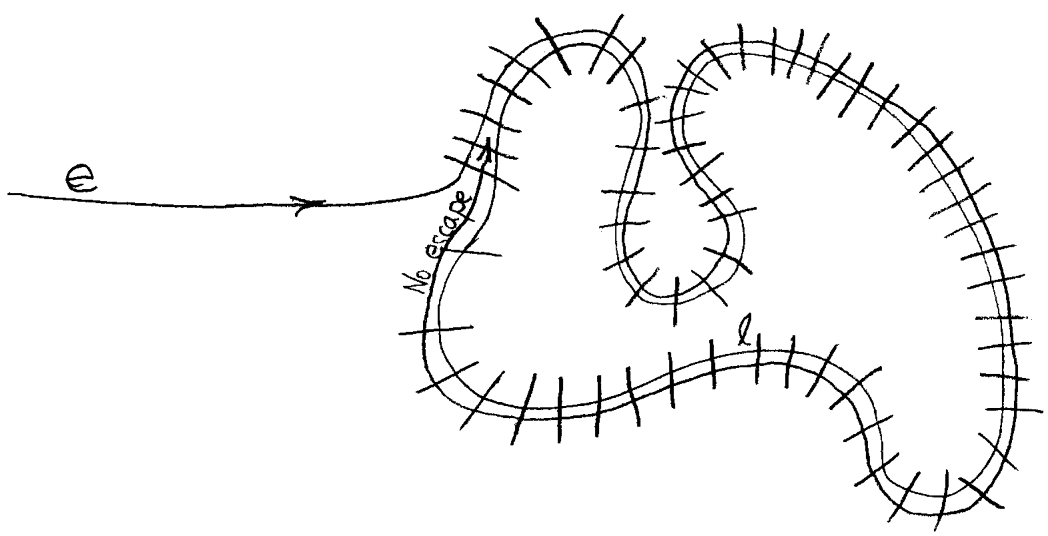

Thus, if f is non-degenerate, $\mathcal{R}_f$ is dense. Furthermore,

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

PROOF. If γ is any complete geodesic lamination which is realizable, then a train track approximation τ can be constructed for the image of γ in N^3, in such a way that all branch lines have curvature close to 0. Then all laminations carried by τ are also realizable; they form a neighborhood of γ. Next we will show that any enlargement $\gamma' \supset \gamma$ of a realizable geodesic lamination γ is also realizable. First note that if γ' is obtained by adding a single leaf l to γ, then γ' is also realizable. This is proved in the same way as in the case of a lamination with finitely many leaves: note that each end of l is asymptotic to a leaf of γ. (You can see this by considering $S - \gamma$.) If f is 8.67
homotoped so that $f(\gamma)$ consists of geodesics, then both ends of $\tilde{f}(l)$ are asymptotic to geodesics in $\tilde{f}(\gamma)$. If the two endpoints were not distinct on S_∞, this would imply the existence of some non-trivial identification of γ by f so that $\pi_1 f$ could not be injective.

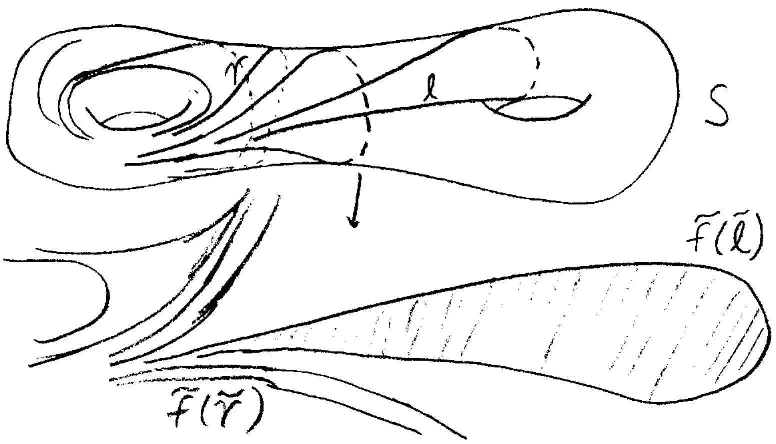

By adding finitely many leaves to any geodesic lamination γ' we can complete it. This implies that γ' is contained in the wrinkling locus of some uncrumpled surface. By 8.8.5 and 8.10.1, the set of uncrumpled surfaces whose wrinkling locus contains γ is compact. Since the wrinkling locus depends continuously on an uncrumpled surface, the set of $\gamma' \in \mathcal{R}_f$ which contains γ is compact. But any $\gamma' \supset \gamma$ can be approximated by laminations such that $\gamma' - \gamma$ consists of a finite number of leaves. This actually follows from 8.10.7, applied to $d(S - \gamma)$. Therefore, every enlargement $\gamma' \supset \gamma$ is in $\mathcal{R}_f$. 8.68

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

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

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

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

Otherwise, there is some non-trivial $h_t : S \to M$ such that $h_1 = h_0 \circ \phi$, where $\phi : S \to S$ is a homotopically non-trivial diffeomorphism. It may happen that ϕ has finite order up to isotopy, as when S is a finite regular covering of another surface in M. The set of all isotopy classes of diffeomorphisms ϕ which arise in this way form a group. If the group is finite, then as in the previous case, $\mathcal{R}_F = \mathcal{GL}(S)$. Otherwise, there is a torsion-free subgroup of finite index (see), so there is an element ϕ of infinite order. The maps f and $\phi \circ f$ are conjugate in Γ, by some element $\beta \in \Gamma$. The group generated by β and $f(\pi_1 S)$ is the fundamental group of a three-manifold which fibers over S^1. □ 8.69

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

8.11. The structure of cusps

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

Suppose that N has parabolic elements that are not parabolic on S. The structure 8.71
of the extra cusps of N is described by the following:

Proposition 8.11.1. *There are geodesic laminations γ_+ and γ_- on S with all leaves compact (i.e., they are finite collections of disjoint simple closed curves) such that the extra cusps in N_e correspond one-to-one with leaves of $\gamma_e (e = +, -)$. In particular, for any element $\alpha \in \pi_1(S)$, $\pi_1(f)(\alpha)$ is parabolic if and only if α is freely homotopic to a cusp of S or to a simple closed curve in γ_+ or γ_-.*

Proof. We need consider only one half, say N_+. For each extra cusp of N_+, there is a half-open cylinder mapped into N_+, with one end on $h^{-1}(0)$ and the other end tending toward the cusp. Furthermore, we can assume that the union of these cylinders is embedded outside a compact set, since we understand the picture in a neighborhood of the cusps. Homotope the ends of the cylinders which lie on $h^{-1}(0)$

so they are geodesics in some hyperbolic structure on $h^{-1}(0)$. One can assume the cylinders are immersed (since maps of surfaces into three-manifolds are appoximable by immersions) and that they are transverse to themselves and to one another. If there are any self-intersections of the cylinders on $h^{-1}(0)$, there must be a double line which begins and ends on $h^{-1}(0)$. Consider the picture in $\tilde{N}$: there are two translates of universal covers of cylinders which meet in a double line, so that in particular their bounding lines meet twice on $h^{-1}(0)$. This contradicts the fact that they are geodesics in some hyperbolic structure. □

8.71

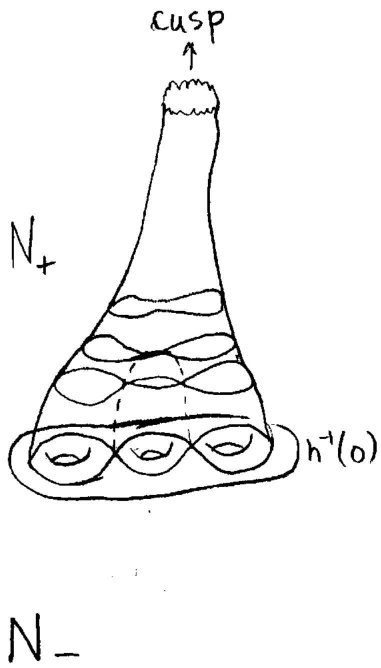

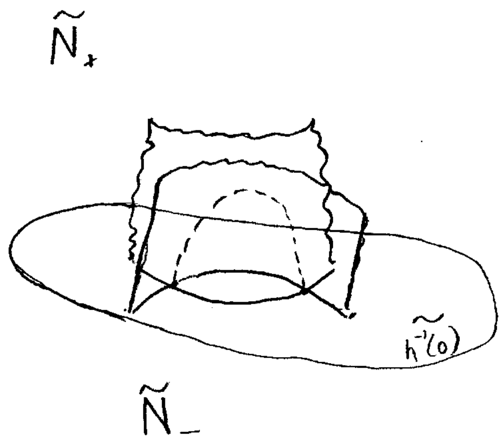

It actually follows that the collection of cylinders joining simple closed curves to the cusps can be embedded: we can modify g so that it takes each of the extra cusps to a neighborhood of the appropriate simple closed curve $\alpha \subset \gamma_\epsilon$, and then do surgery to make $g^{-1}(\alpha)$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

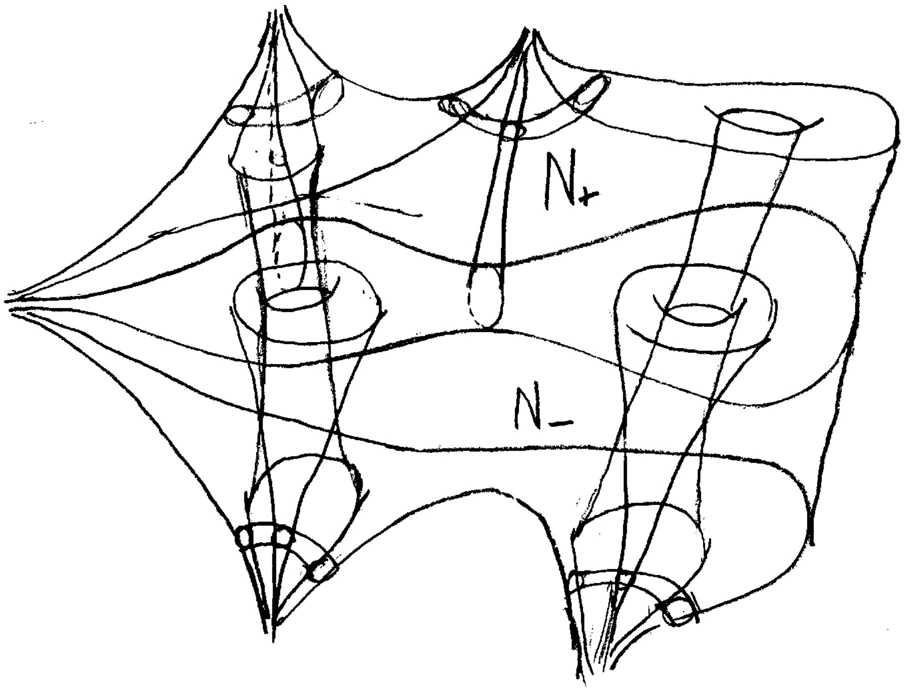

8.72

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

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

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

According to a theorem of Peter Scott, every three-manifold with finitely generated fundamental group is homotopy equivalent to a compact submanifold. In general, such a submanifold will not have incompressible boundary, so it is not as well behaved. We will leave this case for future consideration. 8.73

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

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

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

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

8.12. Harmonic functions and ergodicity

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

$$\begin{aligned} \int_x^{\phi_T(x)} \|\operatorname{grad} h\|\, ds &= h(x) - h(\phi_T(x)) \\ &\leq h(x) \quad (\text{ for } T > 0). \end{aligned}$$

If A is a subset of a flow line $\{\phi_t(x)\}_{t \geq 0}$ of finite length $l(A)$, then by the Schwarz inequality 8.74

Labelled this 8.12.1.eq

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

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

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

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

PROOF. For each point $x \in S$, let c_x be the "characteristic function" of an immersed hyperbolic ball of radius $1 + a/2$ centered at $f(x)$. In other words, $c_x(y)$ is the number of distinct homotopy classes of paths from x to y of length $\leq 1 + a/2$. Let g be defined by integrating c_x over S; in other words, for $y \in N$, 8.75

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

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

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

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

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

$\square$

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

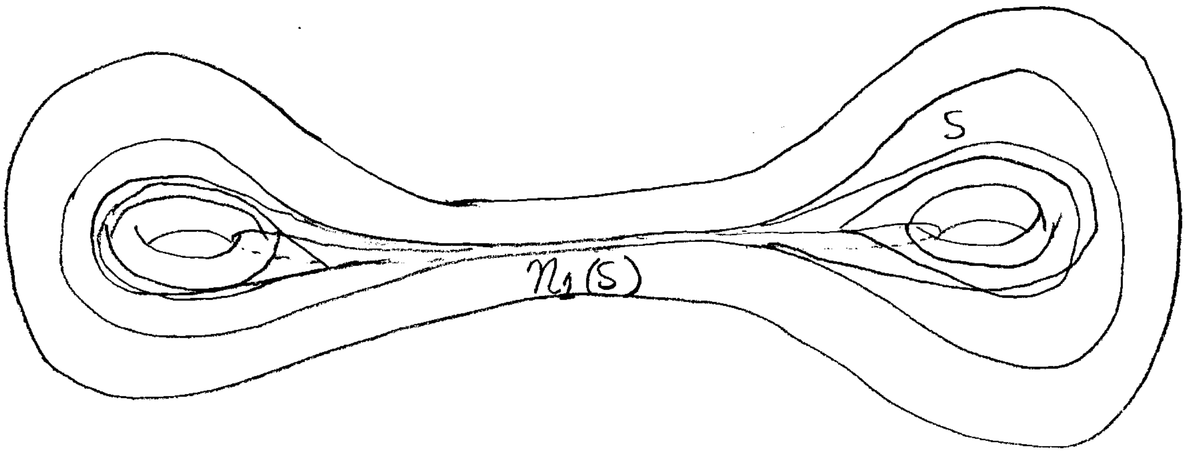

8.76

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

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

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

Corollary 8.12.4. *If $\Gamma = \pi_1 N$, where N is geometrically tame, then L_Γ has measure* 0 *or* 1. *In the latter case, Γ acts ergodically on S^2.*

Proof of Corollary from theorem. This is similar to 8.4.2. Consider any invariant measurable set $A \subset L_\Gamma$, and let h be the harmonic extension of the characteristic function of A. Since A is invariant, h defines a harmonic function, also h, on N. If $L_\Gamma = S^2$, then by 8.12.3 h is constant, so A has measure 0 to 1. If $L_\Gamma \neq S^2$ then the infimum of $(1 - h)$ is the infimum on ∂M, so it is $\geq \frac{1}{2}$. This implies A has measure 0. This completes the proof of 8.12.4. □

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

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

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

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

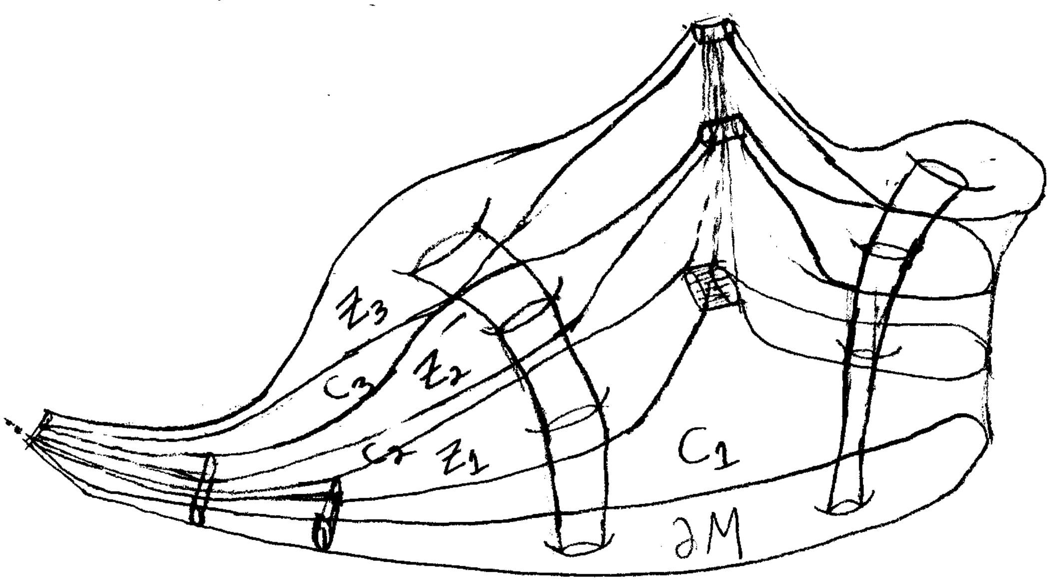

8.78

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

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

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

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

grows as J^2. Since the volume of

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

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

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