William P. Thurston

The Geometry and Topology of Three-Manifolds

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

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

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

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

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

CHAPTER 9

Algebraic convergence

9.1. Limits of discrete groups

It is important for us to develop an understanding of the geometry of deformations of a given discrete group. A qualitative understanding can be attained most concretely by considering limits of sequences of groups. The situation is complicated by the fact that there is more than one reasonable sense in which a group can be the limit of a sequence of discrete groups.

Definition 9.1.1. A sequence $\{\Gamma_i\}$ of closed subgroups of a Lie group G *converges geometrically* to a group Γ if

(i) each $\gamma \in \Gamma$ is the limit of a sequence $\{\gamma_i\}$, with $\gamma_i \in \Gamma_i$, and
(ii) the limit of every convergent sequence $\{\gamma_{i_j}\}$, with $\gamma_{i_j} \in \Gamma_{i_j}$, is in Γ.

Note that the geometric limit Γ is automatically closed. The definition means that Γ_i's look more and more like Γ, at least through a microscope with limited resolution. We shall be mainly interested in the case that the Γ_i's and Γ are discrete. The *geometric topology* on closed subgroups of G is the topology of geometric convergence.

The notion of geometric convergence of a sequence of discrete groups is closely related to geometric convergence of a sequence of complete hyperbolic manifolds of bounded volume, as discussed in 5.11. A hyperbolic three-manifold M determines a subgroup of $\mathrm{PSL}(2,\mathbb{C})$ well-defined *up to conjugacy.* A specific representative of this conjugacy class of discrete groups corresponds to a choice of a base frame: a base 9.2
point p in M together with an orthogonal frame for the tangent space of M at p. This gives a specific way to identify $\tilde{M}$ with H^3. Let $O(\mathcal{H}_{[\epsilon,\infty)})$ consist of all base frames contained in $M_{[\epsilon,\infty)}$, where M ranges over $\mathcal{H}$ (the space of hyperbolic three-manifolds with finite volume). $O(\mathcal{H}_{[\epsilon,\infty)})$ has a topology defined by geometric convergence of groups. The topology on $\mathcal{H}$ is the quotient topology by the equivalence relation of conjugacy of subgroups of $\mathrm{PSL}(2,\mathbb{C})$. This quotient topology is not well-behaved for groups which are not geometrically finite.

Definition 9.1.2. Let Γ be an abstract group, and $\rho_i : \Gamma \to G$ be a sequence of representations of Γ into G. The sequence $\{\rho_i\}$ *converges algebraically* if for every $\gamma \in \Gamma$, $\{\rho_i(\gamma)\}$ converges. The limit $\rho : \Gamma \to G$ is called the algebraic limit of $\{\rho_i\}$.

DEFINITION 9.1.3. Let Γ be a countable group, $\{\rho_i\}$ a sequence of representations of Γ in G with $\rho_i(\Gamma)$ discrete. $\{\rho_i\}$ *converges strongly* to a representation ρ if ρ is the algebraic limit of $\{\rho_i\}$ *and* $\rho\Gamma$ is the geometric limit of $\{\rho_i\Gamma\}$.

EXAMPLE 9.1.4 (Basic example). There is often a tremendous difference between algebraic limits and geometric limits, growing from the following phenomenon in a sequence of cyclic groups.

Pick a point x in H^3, a "horizontal" geodesic ray l starting at x, and a "vertical" plane through x containing the geodesic ray. Define a sequence of representations 9.3
$\rho_i : Z \to \mathrm{PSL}(2, \mathbb{C})$ as follows. Let x_i be

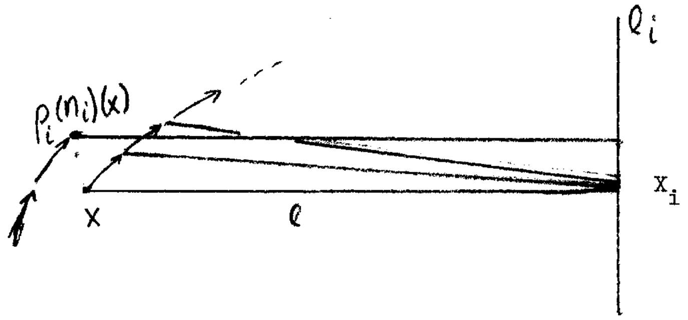

the point on l at distance i from x, and let l_i be the "vertical" geodesic through x_i: perpendicular to l and in the chosen plane. Now define ρ_i on the generator 1 by letting $\rho_i(1)$ be a screw motion around l_i with fine pitched thread so that $\rho_i(1)$ takes x to a point at approximately a horizontal distance of 1 from x and some high power $\rho_i(n_i)$ takes x to a point in the vertical plane a distance of 1 from x. The sequence $\{\rho_i\}$ converges algebraically to a parabolic representation $\rho : \mathbb{Z} \to \mathrm{PSL}(2, \mathbb{C})$, while $\{\rho_i\mathbb{Z}\}$ converges geometrically to a parabolic subgroup of rank 2, generated by $\rho(\mathbb{Z})$ plus an additional generator which moves x a distance of 1 in the vertical plane.

9.4

This example can be described in matrix form as follows. We make use of one-complex parameter subgroups of $\mathrm{PSL}(2, \mathbb{C})$ of the form

$$\begin{bmatrix} \exp w & a \sinh w \\ 0 & \exp -w \end{bmatrix},$$

with $w \in \mathbb{C}$. Define ρ_n by

$$\rho_n(1) = \begin{bmatrix} \exp w_n & n \sinh w_n \\ 0 & \exp -w_n \end{bmatrix}$$

where $w_n = 1/n^2 + \pi i/n$.

Thus $\{\rho_n(1)\}$ converges to

$$\begin{bmatrix} 1 & \pi i \\ 0 & 1 \end{bmatrix}$$

while $\{\rho_n(n)\}$ converges to

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

This example can be easily modified without changing the algebraic limit so that $\{\rho_i(\mathbb{Z})\}$ has no geometric limit, or so that its geometric limit is a one-complex-parameter parabolic subgroup, or so that the geometric limit is isomorphic to $\mathbb{Z} \times \mathbb{R}$.

9.5

This example can also be combined with more general groups: here is a simple case. Let Γ be a Fuchsian group, with M_Γ a punctured torus. Thus Γ is a free group on generators a and b, such that $[a, b]$ is parabolic. Let $\rho : \Gamma \to \mathrm{PSL}(2, \mathbb{C})$ be the identity representation. It is easy to see that $\operatorname{Tr} \rho'[a, b]$ ranges over a neighborhood of 2 as ρ' ranges over a neighborhood of ρ. Any nearby representation determines a nearby hyperbolic structure for $M_{[\epsilon,\infty)}$, which can be thickened to be locally convex except near $M_{(0,\epsilon]}$. Consider representations ρ_n with an eigenvalue for

$$\rho_n[a, b] \sim 1 + C/n^2 + \pi i/n.$$

$\rho_n[a, b]$ translates along its axis a distance of approximately $2\operatorname{Re}(C)/n^2$, while rotating an angle of approximately

$$\frac{2\pi}{n} + \frac{2\operatorname{Im}(C)}{n^2}.$$

Thus the n-th power translates by a distance of approximately $2\operatorname{Re}(C)/n$, and rotates approximately

$$2\pi + \frac{2\operatorname{Im}(C)}{n}.$$

The axis moves out toward infinity as $n \to \infty$. For C sufficiently large, the image of ρ_n will be a geometrically finite group (a Schottky group); a compact convex manifold with $\pi_1 = \rho_n(\Gamma)$ can be constructed by piecing together a neighborhood of $M_{[\epsilon,\infty)}$ with (the convex hull of a helix)/$\mathbb{Z}$. The algebraic limit of $\{\rho_n\}$ is ρ, while the geometric limit is the group generated by $\rho(\Gamma) = \Gamma$ together with an extra parabolic generator commuting with $[a, b]$.

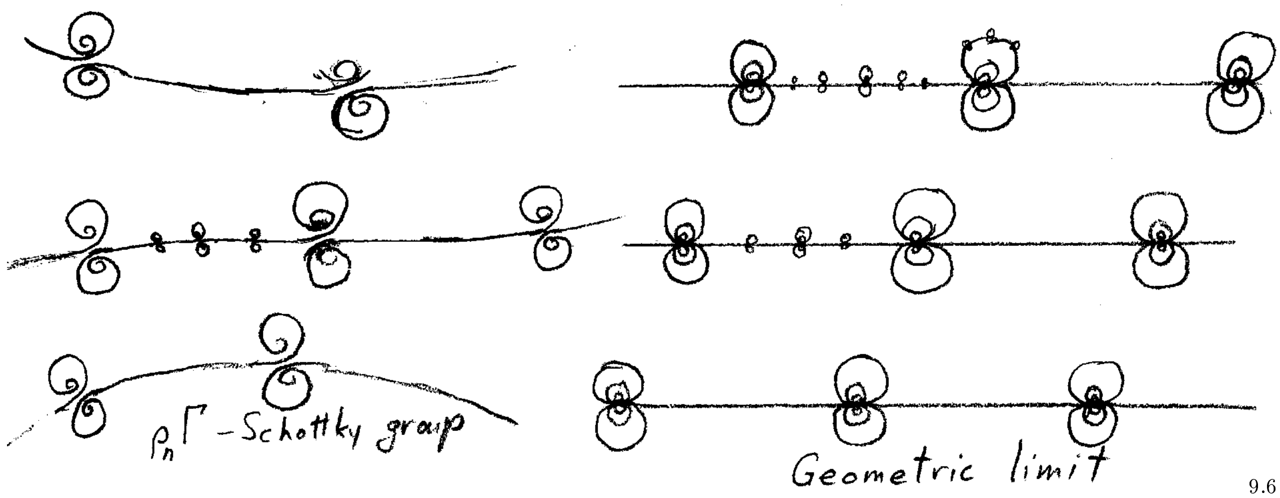

9.6

Troels Jørgensen was the first to analyze and understand this phenomenon. He showed that it is possible to iterate this construction and produce examples as above where the algebraic limit is the fundamental group of a punctured torus, but the geometric limit is not even finitely generated. See § .

Here are some basic properties of convergence of sequences of discrete groups.

PROPOSITION 9.1.5. *If $\{\rho_i\}$ converges algebraically to ρ and $\{\rho_i\Gamma\}$ converges geometrically to Γ', then $\Gamma' \supset \rho\Gamma$.*

PROOF. Obvious. □

PROPOSITION 9.1.6. *For any Lie group G, the space of closed subgroups of G (with the geometric topology) is compact.*

PROOF. Let $\{\Gamma_i\}$ be any sequence of closed subgroups. First consider the case that there is a lower bound to the "size" $d(e, \gamma)$ of elements of $\gamma \in \Gamma_i$. Then there is an upper bound to the number of elements of Γ_i in the ball of radius γ about e, for every γ. The Tychonoff product theorem easily implies the existence of a subsequence converging geometrically to a discrete group.

Now let S be a maximal subspace of $T_e(G)$, the tangent space of G at the identity element e, with the property that for any $\epsilon > 0$ there is a Γ_i whose ϵ-small elements fill out all directions in S, within an angle of ϵ. It is easy to see that S is closed under Lie brackets. Furthermore, a subsequence $\{\Gamma_{i_j}\}$ whose small elements fill out S has 9.7
the property that all small elements are in directions near S. It follows, just as in the previous case, that there is a subsequence converging to a closed subgroup whose tangent space at e is S.

COROLLARY 9.1.7. *The set of complete hyperbolic manifolds N together with base frames in $N_{[\epsilon,\infty)}$ is compact in the geometric topology.*

□

COROLLARY 9.1.8. *Let Γ be any countable group and $\{\rho_i\}$ a sequence of discrete representations of Γ in $\mathrm{PSL}(2,\mathbb{C})$ converging algebraically to a representation ρ. If $\rho\Gamma$ does not have an abelian subgroup of finite index then $\{\rho_i\}$ has a subsequence converging geometrically to a discrete group $\Gamma' \supset \circ\Gamma$. In particular, $\rho\Gamma$ is discrete.*

PROOF. By 9.1.7, there is a subsequence converging geometrically to *some* closed group Γ'. By 5.10.1, the identity component of Γ' must be abelian; since $\rho\Gamma \subset \Gamma'$, the identity component is trivial. □

Note that if the ρ_i are all faithful, then their algebraic limit is also faithful, since there is a lower bound to $d(\rho_i\gamma x, x)$. These basic facts were first proved in ????

Here is a simple example negating the converse of 9.1.8. Consider any discrete group $\Gamma \subset \mathrm{PSL}(2,\mathbb{C})$ which admits an automorphism ϕ of infinite order: for instance, Γ might be a fundamental group of a surface. The sequence of representations ϕ^i has no algebraically convergent subsequence, yet $\{\phi^i\Gamma\}$ converges geometrically to Γ. 9.8

There are some simple statements about the behavior of limit sets when passing to a limit. First, if Γ is the geometric limit of a sequence $\{\Gamma_i\}$, then each point $x \in L_\Gamma$ is the limit of a sequence $x_i \in L_{\Gamma_i}$. In fact, fixed points x (eigenvectors) of non-trivial elements of $\gamma \in \Gamma$ are dense in L_Γ; for high i, Γ_i must have an element near γ, with a fixed point near x. A similar statement follows for the algebraic limit ρ of a sequence of representations ρ_i. Thus, the limit set cannot suddenly increase in the limit. It may suddenly decrease, however. For instance, let $\Gamma \subset \mathrm{PSL}(2,\mathbb{C})$ be any finitely generated group. Γ is residually finite (see §), or in other words, it has a sequence $\{\Gamma_i\}$ of subgroups of finite index converging geometrically to the trivial group (e). $L_{\Gamma_i} = L_\Gamma$ is constant, but $L_{(e)}$ is empty. It is plausible that every finitely generated discrete group $\Gamma \subset \mathrm{PSL}(2,\mathbb{C})$ be a geometric limit of groups with compact quotient.

We have already seen (in 9.1.4) examples where the limit set suddenly decreases in an algebraic limit.

Let Γ be the fundamental group of a surface S with finite area and $\{\rho_i\}$ a sequence of faithful quasi-Fuchsian representations of Γ, preserving parabolicity. Suppose $\{\rho_i\}$ converges algebraically to a representation ρ as a group without any additional parabolic elements. Let N denote $N_{\rho(\Gamma)}$, N_i denote $N_{\rho_i(\Gamma)}$, etc. 9.9

THEOREM 9.2. *N is geometrically tame, and $\{\rho_i\}$ converges strongly to ρ.*

PROOF. If the set of uncrumpled maps of S into N homotopic to the standard map is compact, then using a finite cover of $\mathcal{GL}(S)$ carried by nearly straight train tracks, one sees that for any discrete representation ρ' near ρ, every geodesic lamination γ of S is realizable in N' near its realizations in N. (Logically, one can think of uncrumpled surfaces as equivariant uncrumpled maps of M^2 into H^3, with the compact-open topology, so that "nearness" makes sense.) Choose any subsequence of the ρ_i's so that the bending loci for the two boundary components of M_i converge in $\mathcal{GL}(S)$. Then the two boundary components must converge to locally convex disjoint embeddings of S in N (unless the limit is Fuchsian). These two surfaces are homotopic, hence they bound a convex submanifold M of N, so $\rho(\Gamma)$ is geometrically finite.

Since $M_{[\epsilon,\infty)}$ is compact, strong convergence of $\{\rho_i\}$ follows form 8.3.3: no unexpected identifications of N can be created by a small perturbation of ρ which preserves parabolicity.

If the set of uncrumpled maps of S homotopic to the standard map is not compact, then it follows immediately from the definition that N has at least one geometrically infinite tame end. We must show that both ends are geometrically tame. The possible phenomenon to be wary of is that the bending loci β_i^+ and β_i^- of the two boundary components of M_i might converge, for instance, to a single point λ in $\mathcal{GL}(S)$. (This would be conceivable if the "simplest" homotopy of one of the two boundary compo- 9.10
nents to a reference surface which persisted in the limit first carried it to the vicinity of the other boundary component.) To help in understanding the picture, we will first find a restriction for the way in which a hyperbolic manifold with a geometrically tame end can be a covering space.

Definition 9.2.1. Let N be a hyperbolic manifold, P a union of horoball neighborhoods of its cusps, E' an end of $N - P$. E' is *almost geometrically tame* if some finite-sheeted cover of E' is (up to a compact set) a geometrically tame end. (Later we shall prove that if E is almost geometrically tame it is geometrically tame.)

Theorem 9.2.2. *Let N be a hyperbolic manifold, and $\tilde N$ a covering space of N such that $\tilde N - \tilde P$ has a geometrically infinite tame end E bounded by a surface $S_{[\epsilon,\infty)}$. Then either N has finite volume and some finite cover of N fibers over S^1 with fiber S, or the image of E in $N - P$, up to a compact set, is an almost geometrically tame end of N.*

Proof. Consider first the case that all points of E identified with $S_{[\epsilon,\infty)}$ in the projection to N lie in a compact subset of E. Then the local degree of the projection of E to N is finite in a neighborhood of the image of S. Since the local degree is constant except at the image of S, it is everywhere finite.

Let $G \subset \pi_1 N$ be the set of covering transformations of H^3 over N consisting of 9.11
elements g such that $g\tilde E \cap \tilde E$ is all of $\tilde E$ except for a bounded neighborhood of $\tilde S$. G

is obviously a group, and it contains $\pi_1 S$ with finite index. Thus the image of E, up to compact sets, is an almost geometrically tame end of N.

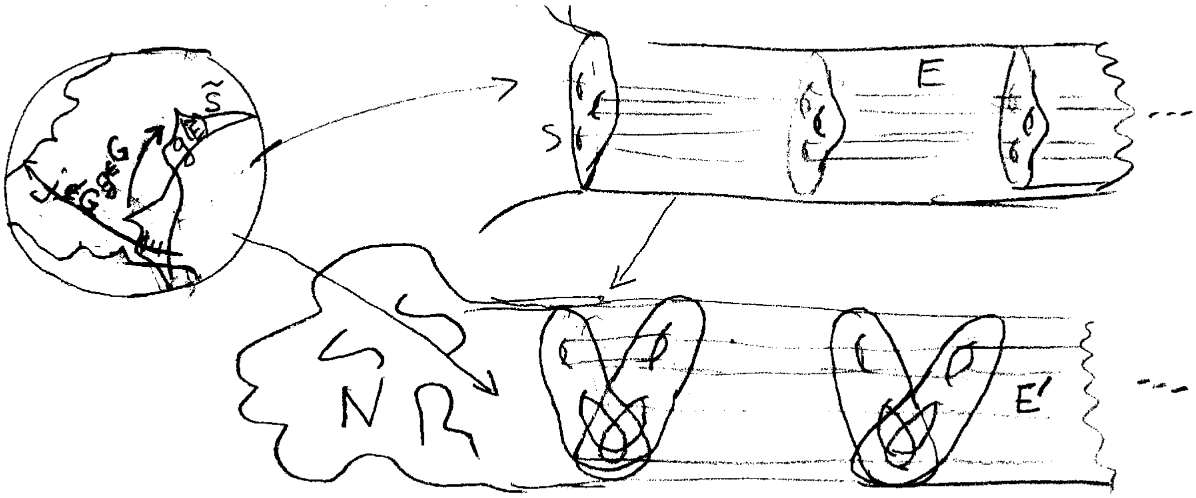

The other case is that $S_{[\epsilon,\infty)}$ is identified with a non-compact subset of E by projection to N. Consider the set I of all uncrumpled surfaces in E whose images intersect the image of $S_{[\epsilon,\infty)}$. Any short closed geodesic on an uncrumpled surface of E is homotopic to a short geodesic of E (not a cusp), since E contains no cusps other than the cusps of S. Therefore, by the proof of 8.8.5, the set of images of I in N is precompact (has a compact closure). If I itself is not compact, then N has a finite cover which fibers over S^1, by the proof of 8.10.9. If I is compact, then (since uncrumpled surfaces cut E into compact pieces), infinitely many components of the set of points identified with $S_{[\epsilon,\infty)}$ are compact and disjoint from S. 9.12

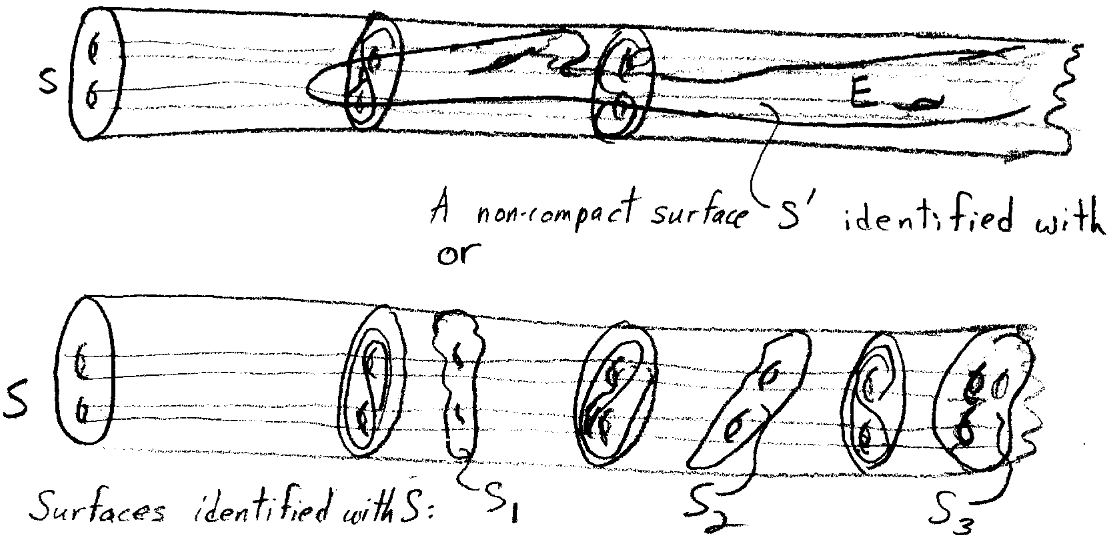

These components consist of immersions of k-sheeted covering spaces of S injective on π_1, which must be homologous to $\pm k\,[S]$. Pick two disjoint immersions with the same sign, homologous say to $-k\,[S]$ and $-l\,[S]$. Appropriate multiples of these cycles are homologous by a compactly supported three-chain which maps to a three-cycle in $N - P$, hence N has finite volume. Theorem 9.2.2 now follows from 8.10.9. $\square$

We continue the proof of Theorem 9.2. We may, without loss of generality, pass to a subsequence of representations ρ_i such that the sequences of bending loci $\{\beta_i^+\}$ and $\{\beta_i^-\}$ converge, in $\mathcal{PL}_0(S)$, to laminations β^+ and β^-. If β^+, say, is realizable for the limit representation ρ, then any uncrumpled surface whose wrinkling locus contains β^+ is embedded and locally convex—hence it gives a geometrically finite end of N. The only missing case for which we must prove geometric tameness is that neither 9.13
β^+ nor β^- is realizable. Let $\lambda_i^\epsilon \in \mathcal{PL}_0(S)$ (where $\epsilon = +, -$) be a sequence of geodesic laminations with finitely many leaves and with transverse measures approximating β_i^ϵ closely enough that the realization of λ_i^ϵ in N_i is near the realization of β_i^ϵ. Also suppose that $\lim \lambda_i^\epsilon = \beta^\epsilon$ in $\mathcal{PL}_0(S)$. The laminations λ_i^ϵ are all realized in N. They must tend toward ∞ in N, since their limit is not realized. We will show that they tend toward ∞ in the ϵ-direction. Imagine the contrary—for definiteness, suppose that the realizations of $\{\lambda_i^+\}$ in N go to ∞ in the $-$ direction. The realization of each λ_i^+ in N_j must be near the realization in N, for high enough j. Connect λ_j^+ to λ_i^+ by a short path $\lambda_{i,j,t}$ in $\mathcal{PL}_0(S)$. A family of uncrumpled surfaces $S_{i,j,t}$ realizing the $\lambda_{i,j,t}$ is not continuous, but has the property that for t near t_0, $S_{i,j,t}$ and S_{i,j,t_0} have points away from their cusps which are close in N. Therefore, for every uncrumpled surface U between $S_{i,j,0}$ and $S_{i,j,1}$ (in a homological sense), there is some t such that $S_{i,j,t} \cap U \cap (N - P)$ is non-void.

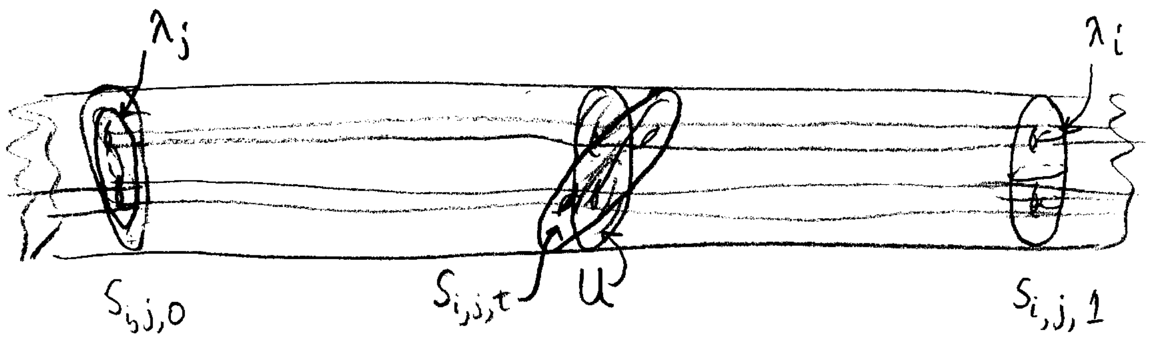

9.14

Let γ be any lamination realized in N, and U_j be a sequence of uncrumpled surfaces realizing γ in N_j, and converging to a surface in N. There is a sequence $S_{i(j),j,t(j)}$ of uncrumpled surfaces in N_j intersecting U_j whose wrinkling loci tend toward β^+.

Without loss of generality we may pass to a geometrically convergent subsequence, with geometric limit Q. Q is covered by N. It cannot have finite volume (from the analysis in Chapter 5, for instance), so by 8.14.2, it has an almost geometrically tame end E which is the image of the $-$ end E_- of N. Each element α of $\pi_1 E$ has a finite power $\alpha^k \in \pi_1 E_-$. Then a sequence $\{\alpha_i\}$ approximating α in $\pi_1(N_i)$ has the property that the α_i^k have bounded length in the generators of $\pi_1 S$, this implies that the α_i have bounded length, so α is in fact in $\pi_1 E_-$, and $E_- = E$ (up to compact sets). Using this, we may pass to a subsequence of $S_{i(j),j,t}$'s which converge to an uncrumpled surface R in E. R is incompressible, so it is in the standard homotopy class. It realizes β^+, which is absurd. there is no 8.14.2

We may conclude that N has two geometrically tame ends, each of which is mapped homeomorphically to the geometric limit Q. (This holds whether or not they are geometrically infinite.) This implies the local degree of $N \to Q$ is finite one or two (in case the two ends are identified in Q). But any covering transformation α of N over Q has a power (its square) in $\pi_1 N$, which implies, as before, that $\alpha \in \pi_1 N$, so that $N = Q$. This concludes the proof of 9.2. □

9.3. The ending of an end

In the interest of avoiding circumlocution, as well as developing our image of a geometrically tame end, we will analyze the possibilities for non-realizable laminations in a geometrically tame end.

We will need an estimate for the area of a cylinder in a hyperbolic three-manifold. Given any map $f : S^1 \times [0,1] \to N$, where N is a convex hyperbolic manifold, we may straighten each line $\theta \times [0,1]$ to a geodesic, obtaining a ruled cylinder with the same boundary.

THEOREM 9.3.1. *The area of a ruled cylinder (as above) is less than the length of its boundary.*

PROOF. The cylinder can be C^0-approximated by a union of small quadrilaterals each subdivided into two triangles. The area of a triangle is less than the minimum of the lengths of its sides (see p. 6.5). $\square$

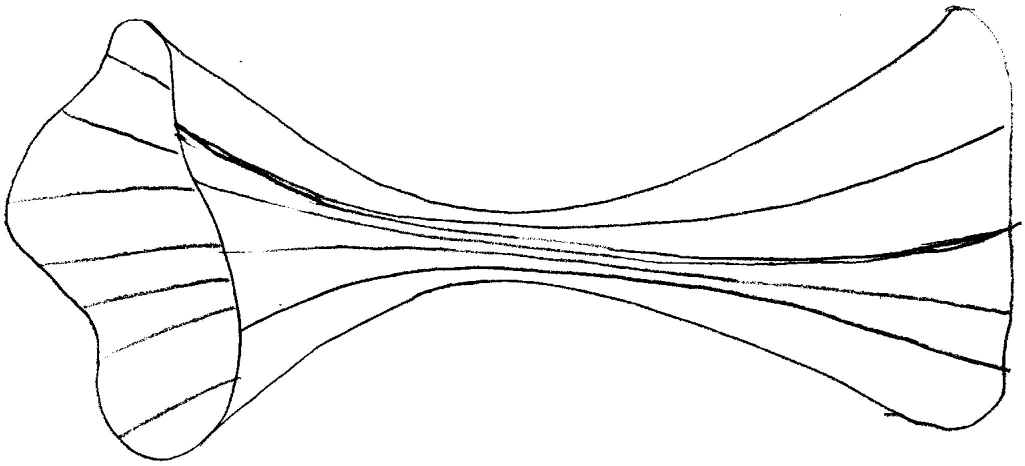

9.16

If the two boundary components of the cylinder C are far apart, then most of the area is concentrated near its boundary. Let γ_1 and γ_2 denote the two components of ∂C.

THEOREM 9.3.2. *Area* $(C - \mathcal{N}_r\gamma_1) \leq e^{-r}\, l(\gamma_1) + l(\gamma_2)$ *where* $r \geq 0$ *and* l *denotes length.*

This is derived by integrating the area of a triangle in polar coordinates from any vertex:

$$A = \int \int_0^{T(\theta)} \sinh t\, dt\, d\theta = \int (\cosh T(\theta) - 1)\, d\theta$$

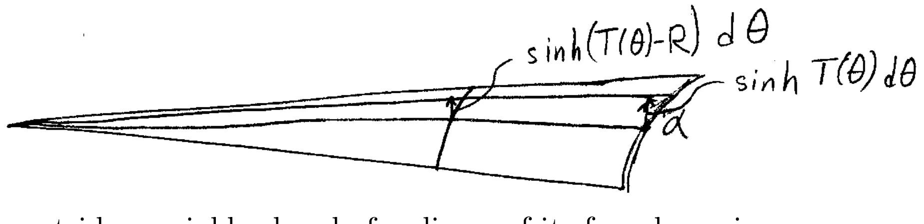

The area outside a neighborhood of radius r of its far edge α is

$$\int \cosh\,(T(\theta) - r) - 1\, d\theta < e^{-r} \int \sinh T(\theta)\, d\theta < e^{-r}\, l(\alpha).$$

This easily implies 9.3.2

Let E be a geometrically tame end, cut off by a surface $S_{[\epsilon,\infty)}$ in $N - P$, as usual. A curve α in E homotopic to a simple closed curve α' on S gives rise to a ruled cylinder $C_\alpha : S^1 \times [0,1] \to N$.

Now consider two curves α and β homotopic to simple closed curves α' and β' 9.17
on S. One would expect that if α' and β' are forced to intersect, then either α must intersect C_β or β must intersect C_α, as in 8.11.1

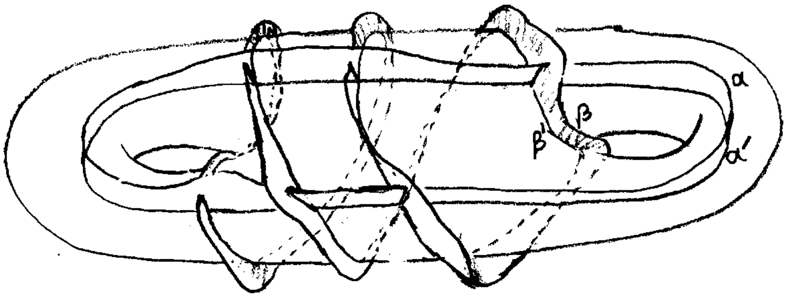

We will make this more precise by attaching an invariant to each intersection. Let us assume, for simplicity, that α' and β' are geodesics with respect to some hyperbolic structure on S. Choose one of the intersection points, p_0, of α' and β' as a base point for N. For each other intersection point p_i, let α_i and β_i be paths on α' and β' from p_0 to p_i. Then $\alpha_i * \beta_i^{-1}$ is a closed loop, which is non-trivial in $\pi_1(S)$ when $i \neq 0$ since two geodesics in $\tilde{S}$ have at most one intersection.

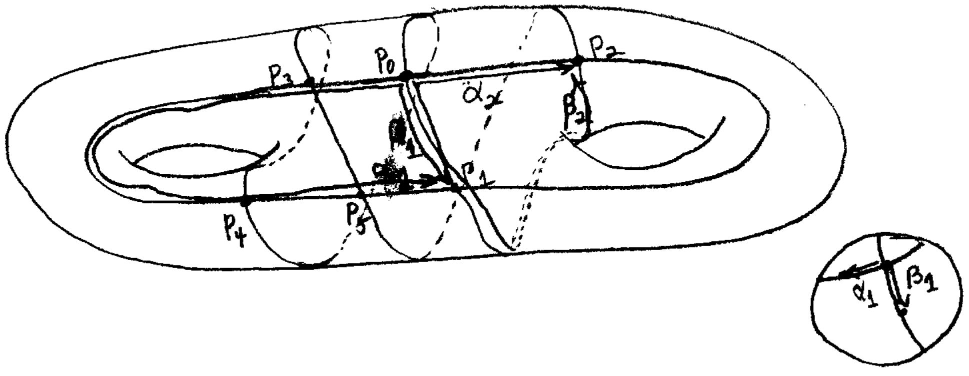

There is some ambiguity, since there is more than one path from α_0 to α_i on 9.18
α'; in fact, α_i is well-defined up to a power of α'. Let $\langle g \rangle$ denote the cyclic group generated by an element g. Then $\alpha_i \cdot \beta_i^{-1}$ gives a well-defined element of the double coset space $\langle \alpha' \rangle \backslash \pi_1(S) / \langle \beta' \rangle$. [The double coset $H_1 g H_2 \in H_1 \backslash G / H_2$ of an element $g \in G$ is the set of all elements $h_1 g h_2$, where $h_i \in H_i$.] The double cosets associated to two different intersections p_i and p_j are distinct: if $\langle \alpha' \rangle \alpha_i \beta_i^{-1} \langle \beta' \rangle = \langle \alpha' \rangle a_j \beta_j^{-1} \langle \beta' \rangle$, then there is some loop $\alpha_j^{-1} \alpha'^k \alpha_i \beta_i^{-1} \beta'^l \beta_j$ made up of a path on α' and a path on

β' which is homotopically trivial—a contradiction. In the same way, a double coset $D_{x,y}$ is attached to each intersection of the cylinders C_α and C_β. Formally, these intersection points should be parametrized by the domain: thus, an intersection point means a pair $(x, y) \in (S^1 \times I) \times (S' \times I)$ such that $C_\alpha x = C_\beta y$.

Let $i(\gamma, \delta)$ denote the number of intersections of any two simple geodesics γ and δ on S. Let $D(\gamma, \delta)$ be the set of double cosets attached to intersection points of γ and δ (including p_0). Thus $i(\gamma, \delta) = |D(\gamma, \delta)|$. $D(\alpha, C_\beta)$ and $D(C_\alpha, \beta)$ are defined similarly.

Proposition 9.3.3. $|\alpha \cap C_\beta| + |C_\alpha \cap \beta| \geq i(\alpha', \beta')$. *In fact*

$$D(a, C_\beta) \cup D(C_\alpha, \beta) \supset D(\alpha', \beta').$$

9.19

Proof. First consider cylinders C'_α and C'_β which are contained in E, and which are nicely collared near S. Make C'_α and C'_β transverse to each other, so that the double locus $L \subset (S^1 \times I) \times (S^1 \times I)$ is a one-manifold, with boundary mapped to $\alpha \cup \beta \cup \alpha' \cup \beta'$. The invariant $D_{(x,y)}$ is locally constant on L, so each invariant occurring for $\alpha' \cap \beta'$ occurs for the entire length of interval in L, which must end on α or β. In fact, each element of $D(\alpha', \beta')$ occurs as an invariant of an odd number of points $\alpha \cup \beta$.

Now consider a homotopy h_t of C'_β to C_β, fixing $\beta \cup \beta'$. The homotopy can be perturbed slightly to make it transverse to α, although this may necessitate a slight movement of C_β to a cylinder C''_β. Any invariant which occurs an odd number of times for $a \cap C'_\beta$ occurs also an odd number of times for $\alpha \cap C''_\beta$. This implies that the invariant must also occur for $a \cap C_\beta$. □

Remark. By choosing orientations, we could of course associate signs to intersection points, thereby obtaining an algebraic invariant $\mathcal{D}(\alpha', \beta') \in \mathbb{Z}^{\langle \alpha' \rangle \backslash \pi_1 S / \langle \beta' \rangle}$. Then 9.3.3 would become an equation,

$$\mathcal{D}(\alpha', \beta') = \mathcal{D}(\alpha, C_\beta) + \mathcal{D}(C_\alpha, \beta).$$

Since $\pi_1(S)$ is a discrete group, there is a restriction on how closely intersection points can be clustered, hence a restriction on $|D(\alpha, c_\beta)|$ in terms of the length of α times the area of C_β.

9.20

Proposition 9.3.4. *There is a constant K such that for every curve α in E with distance R from S homotopic to a simple closed curve α' on S and every curve β in E not intersecting C_α and homotopic to a simple curve β' on S,*

$$i(\alpha', \beta') \leq K\big[l(\alpha) + \big(l(\alpha) + 1\big)\big(l(\beta) + e^{-R} + l(\beta')\big)\big].$$

Proof. Consider intersection points $(x, y) \in S^1 \times (S^1 \times I)$ of α and C_β. Whenever two of them, (x, y) and (x', y'), are close in the product of the metrics induced from N, there is a short loop in N which is non-trivial if $D_{(x,y)} \neq D_{(x',y')}$.

Case (i). α is a short loop. Then there can be no short non-trivial loop on C_β near an intersection point with α. The disks of radius ϵ on C_β about intersection points with α have area greater than some constant, except in special cases when they are near ∂C_β. If necessary, extend the edges of C_β slightly, without substantially changing the area. The disks of radius ϵ must be disjoint, so this case follows from 9.3.2 and 9.3.3.

Case (ii). α is not short. Let $E \subset C_\beta$ consist of points through which there is a short loop homotopic to β. If (x, y) and (x', y') are intersection points with $D_{x,y} \neq D_{x',y'}$ and with y, y' in E, then x and x' cannot be close together—otherwise two distinct conjugates of β would be represented by short loops through the same point. The number of such intersections is thus estimated by some constant times $l(\alpha)$. 9.21

Three intersections of α with $C_\beta - E$ cannot occur close together. $S^1 \times (C_\beta - E)$ contains the balls of radius ϵ, with multiplicity at most 2, and each ball has a definite volume. This yields 9.3.4. □

Let us generalize 9.3.4 to a statement about measured geodesic laminations. Such a lamination (γ, μ) on a hyperbolic surface S has a well-defined "average length" $l_S(\gamma, \mu)$. This can be defined as the total mass of the measure which is locally the product of the transverse measure μ with one-dimensional Lebesgue measure on the leaves of γ. Similarly, a realization of γ in a homotopy class $f : S \to N$ has a length $l_f(\gamma, \mu)$. The length $l_S(\gamma, \mu)$ is a continuous function on $\mathcal{ML}_0(S)$, and $l_f(\gamma)$ is a continuous function where defined. If γ is realized a distance of R from an uncrumpled surface S, then $l_f(\gamma, \mu) \leq (1/\cosh R) l_S(\gamma, \mu)$. This implies that if f preserves non-parabolicity, l_f extends continuously over all of $\mathcal{ML}_0$ so that its zero set is the set of non-realizable laminations.

The intersection number $i\big((\gamma_1, \mu_1), (\gamma_2, \mu_2)\big)$ of two measured geodesic laminations is defined similarly, as the total mass of the measure $\mu_1 \times \mu_2$ which is locally the product of μ_1 and μ_2. (This measure $\mu_1 \times \mu_2$ is interpreted to be zero on any common leaves of γ_1 and γ_2.)

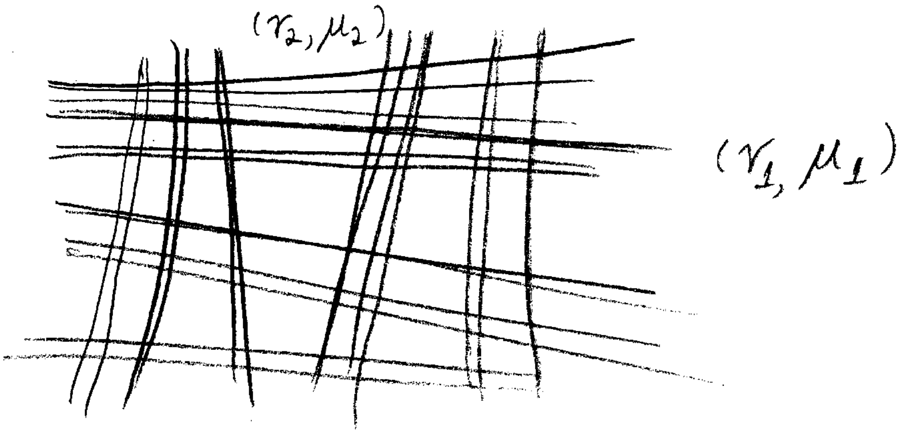

9.22

Given a geodesic lamination γ realized in E, let d_γ be the miniaml distance of an uncrumpled surface through γ from $S_{[\epsilon,\infty)}$.

THEOREM 9.3.5. *There is a constant K such that for any two measured geodesic laminations (γ_1, μ_1) and $(\gamma_2, \mu_2) \in \mathcal{ML}_0(S)$ realized in E,*

$$i\big((\gamma_1, \mu_1), (\gamma_2, \mu_2)\big) \leq K \cdot e^{-2R} l_S(\gamma_1, \mu_1) \cdot l_S(\gamma_2, \mu_2)$$

where $R = \inf(d_{\gamma_1}, d_{\gamma_2})$.

PROOF. First consider the case that γ_1 and γ_2 are simple closed geodesics which are not short. Apply the proof of 9.3.4 first to intersections of γ_1 with C_{γ_2}, then to intersections of C_{γ_1} with γ_2. Note that $l_S(\gamma_i)$ is estimated from below by $e^R l(\gamma_i)$, so the terms involving $l(\gamma_i)$ can be replaced by $Ce^{-R} l(\gamma_i)$. Since γ_1 and γ_2 are not short, one obtains

$$i(\gamma_1, \gamma_2) \leq K \cdot e^{-2R}\, l_S(\gamma_1)\, l_S(\gamma_2),$$

for some constant K. Since both sides of the inequality are homogeneous of degree one in γ_1 and γ_2, it extends by continuity to all of $\mathcal{ML}_0(S)$. □

Consider any sequence $\{(\gamma_i, \mu_i)\}$ of measured geodesic laminations in $\mathcal{ML}_0(S)$ whose realizations go to ∞ in E. If (λ_1, μ_1) and (λ_2, μ_2) are any two limit points of this sequence, 9.3.5 implies that $i(\lambda_1, \lambda_2) = 0$: in other words, the leaves do not cross. The union $\lambda_1 \cup \lambda_2$ is still a lamination. 9.23

DEFINITION 9.3.6. The *ending* lamination $\epsilon(E) \in \mathcal{GL}(S)$ is the union of all limit points λ_i, as above.

Clearly, $\epsilon(E)$ is compactly supported and it admits a measure with full support. The set $\Delta(E) \subset \mathcal{PL}_0(S)$ of all such measures on $\epsilon(E)$ is closed under convex combinations, hence its intersection with a local coordinate system (see p. 8.59) is convex.

In fact, a maximal train track carrying $\epsilon(E)$ defines a single coordinate system containing $\Delta(E)$.

The idea that the realization of a lamination depends continuously on the lamination can be generalized to the ending lamination $\epsilon(E)$, which can be regarded as being realized at ∞.

PROPOSITION 9.3.7. *For every compact subset K of E, there is a neighborhood U of $\Delta(E)$ in $\mathcal{PL}_0(S)$ such that every lamination in $U - \Delta(E)$ is realized in $E - K$.*

PROOF. It is convenient to pass to the covering of N corresponding to $\pi_1 S$. Let S' be an uncrumpled surface such that K is "below" S' (in a homological sense). Let $\{V_i\}$ be a neighborhood basis for $\Delta(E)$ such that $V_i - \Delta(E)$ is path-connected, and let $\lambda_i \in V_i - \Delta(E)$ be a sequence whose realizations go to ∞ in E. If there is any point $\pi_i \in V_i - \Delta(E)$ which is a non-realizable lamination or whose realization is not "above" S', connect λ_i to π_i by a path in V_i. There must be some element of this path whose realization intersects $S'_{[\epsilon,\infty)}$ (since the realizations cannot go to ∞ while in E.) Even if certain non-peripheral elements of S are parabolic, excess pinching of non-peripheral curves on uncrumpled surfaces intersecting S' can be avoided if S' is far from S, since there are no extra cusps in E. Therefore, only finitely many such π_i's can occur, or else there would be a limiting uncrumpled surface through S realizing the unrealizable. □

PROPOSITION 9.3.8. *Every leaf of $\epsilon(E)$ is dense in $\epsilon(E)$, and every non-trivial simple curve in the complement of $\epsilon(E)$ is peripheral.*

PROOF. The second statement follows easily from 8.10.8, suitably modified if there are extra cusps. The first statement then follows from the next result:

PROPOSITION 9.3.9. *If γ is a geodesic lamination of compact support which admits a nowhere zero transverse measure, then either every leaf of γ is dense, or there is a non-peripheral non-trivial simple closed curve in $S - \gamma$.*

PROOF. Suppose $\delta \subset \gamma$ is the closure of any leaf. Then δ is also an open subset of γ: all leaves of γ near δ are trapped forever in a neighborhood of δ. This is seen by considering the surface $S - \delta$.

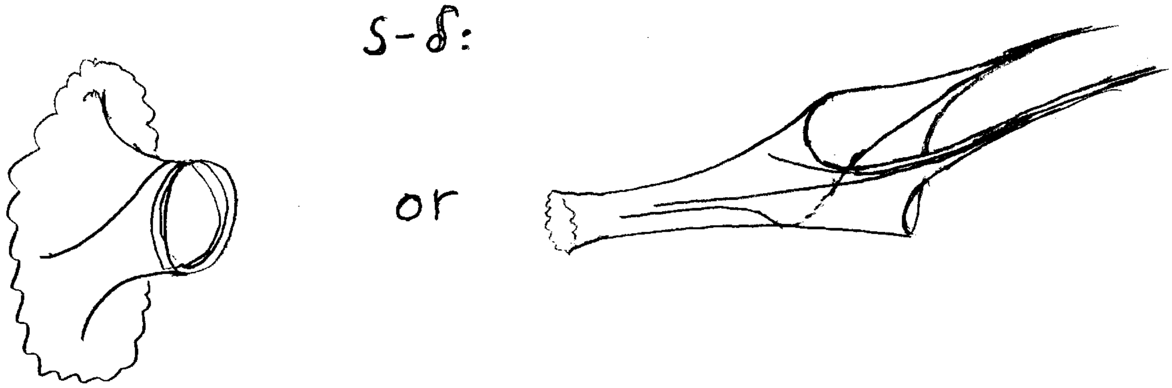

9.25

An arc transverse to these leaves would have positive measure, which would imply that a transverse arc intersecting these leaves infinitely often would have infinite measure. (In general, a closed union of leaves $\delta \subset \gamma$ in a general geodesic lamination has only a finite set of leaves of γ intersecting a small neighborhood.)

If $\delta \neq \gamma$, then δ has two components, which are separated by some homotopically non-trivial curve in $S - \gamma$. $\square$

$\square$

Corollary 9.3.10. *For any homotopy class of injective maps $f : S \to N$ from a hyperbolic surface of finite area to a complete hyperbolic manifold, if f preserves parabolicity and non-parabolicity, there are $n = 0$, 1 or 2 non-realizable laminations ϵ_i $[1 \leq i \leq n]$ such that a general lamination γ on S is non-realizable if and only if the union of its non-isolated leaves is an ϵ_i.*

9.4. Taming the topology of an end

We will develop further our image of a geometrically tame end, once again to avoid circumlocution.

Theorem 9.4.1. *A geometrically tame end $E \subset N - P$ is topologically tame. In other words, E is homeomorphic to the product $S_{[\epsilon,\infty)} \times [0, \infty)$.*

Theorem 9.4.1 will be proved in §§9.4 and 9.5.

Corollary 9.4.2. *Almost geometrically tame ends are geometrically tame.*

Proof that 9.4.1 implies 9.4.2. Let E' be an almost geometrically tame end, finitely covered (up to compact sets) by a geometrically tame end $E = S_{[\epsilon,\infty)} \times [0, \epsilon)$, 9.26
with projection $p : E \to E'$. Let $f : E' \to [0, \epsilon)$ be a proper map. The first step is to find an incompressible surface $S' \subset E'$ which cuts it off (except for compact sets).

Choose t_0 high enough that $p : E \to E'$ is defined on $S_{[\epsilon,\infty)} \times [t_0, \infty)$, and choose $t_1 > t_0$ so that $p\big(S_{[\epsilon,\infty)} \times [t_1, \infty)\big)$ does not intersect $p(S_{[\epsilon,\infty)} \times t_0)$.

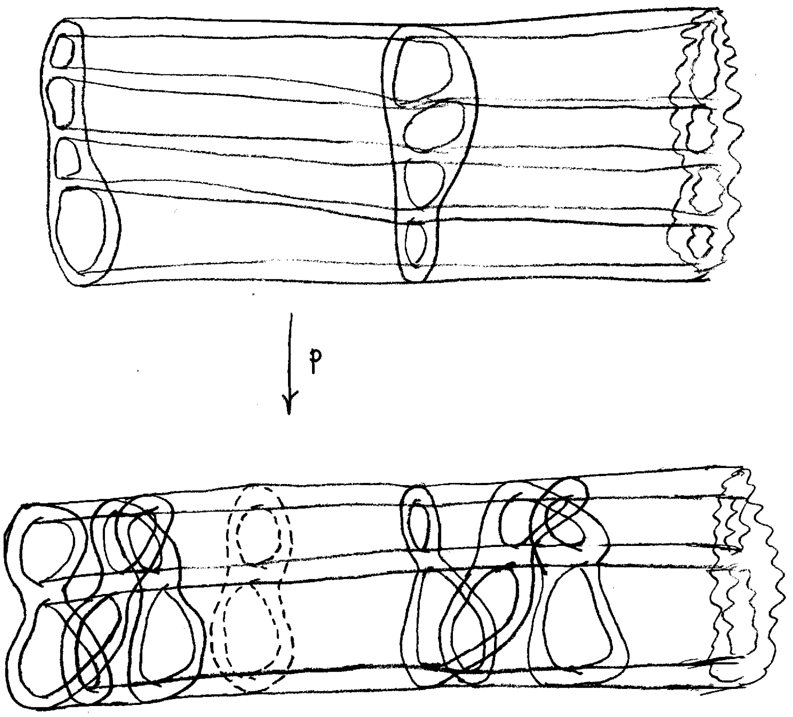

9.27

Let $r \in [0, \infty)$ be any regular value for f greater than the supremum of $f \circ p$ on $S_{[\epsilon,\infty)} \times [0, t_1)$. Perform surgery (that is, cut along circles and add pairs of disks) to $f^{-1}(r)$, to obtain a not necessarily connected surface S' in the same homology class which is incompressible in

$$E' - p\big(S_{[\epsilon,\infty)} \times [0, t_0)\big).$$

The fundamental group of S' is still generated by loops on the level set $f = r$. S' is covered by a surface $\tilde{S}'$ in E. $\tilde{S}'$ must be incompressible in $E-$ otherwise there would be a non-trivial disk D mapped into $S_{[\epsilon,\infty)} \times [t_1, \infty)$ with boundary on $\tilde{S}$; $p \circ D$ would be contained in

$$E' - p(S_{[\epsilon,\infty)} \times [0, t_0])$$

so S' would not be incompressible (by the loop theorem). One deduces that $\tilde{S}'$ is homotopic to $S_{[\epsilon,\infty)}$ and S' is incompressible in $N - P$.

If E is geometrically finite, there is essentially nothing to prove—E corresponds to a component of $\partial \tilde{M}$, which gives a convex embedded surface in E'. If E is geometrically infinite, then pass to a finite sheeted cover E'' of E which is a regular cover of E'. The ending lamination $\epsilon(E'')$ is invariant under all diffeomorphisms (up to compact sets) of E''. Therefore it projects to a non-realizable geodesic lamination $\epsilon(E')$ on S'. □

Proof of 9.4.1. We have made use of one-parameter families of uncrumpled surfaces in the last two sections. Unfortunately, these surfaces do not vary continuously. To prove 9.4.1, we will show, in §9.5, how to interpolate with more general surfaces, to obtain a (continuous) proper map $F\colon S_{[\epsilon,\infty)} \times [0,\infty) \to E$. The theorem will follow fairly easily once F is constructed: 9.28

Proposition 9.4.3. *Suppose there is a proper map $F : S_{[\epsilon,\infty)} \times [0,\infty) \to E$ with $F(S_{[\epsilon,\infty)} \times 0)$ standard and with $F\big(\partial S_{[\epsilon,\infty)} \times [0,\infty)\big) \subset \partial(N-P)$. Then E is homeomorphic to $S_{[\epsilon,\infty)} \times [0,\infty)$.*

Proof of 9.4.3. This is similar to 9.4.2. Let $f : E \to [0,\infty)$ be a proper map. For any compact set $K \subset E$, we can find a $t_1 > 0$ so that $F\big(S_{[\epsilon,\infty)} \times [t_1,\infty)\big)$ is disjoint from K. Let r be a regular value for f greater than the supremum of $f \circ F$ on $S_{[\epsilon,\infty)} \times [0,t_1]$. Let $S' = f^{-1}(r)$ and $S'' = (f \circ F)^{-1}(r)$. $F : S'' \to S'$ is a map of degree one, so it is surjective on π_1 (or else it would factor through a non-trivial covering space on S', hence have higher degree). Perform surgery on S' to make it incompressible in the complement of K, without changing the homology class. Now S' must be incompressible in E; otherwise there would be some element α of $\pi_1 S'$ which is null-homotopic in E. But α comes from an element β on S'' which is null-homotopic in $S_{[\epsilon,\infty)} \times [t_1,\infty)$, so its image α is null-homotopic in the complement of K. It follows that S' is homotopic to $S_{[\epsilon,\infty)}$, and that the compact region of E cut off by S' is homeomorphic to $S_{[\epsilon,\infty)} \times I$. By constructing a sequence of such disjoint surfaces going outside of every compact set, we obtain a homeomorphism with $S_{[\epsilon,\infty)} \times [0,\infty)$. □

□

9.5. Interpolating negatively curved surfaces

Now we turn to the task of constructing a continuous family of surfaces moving out 9.29
to a geometrically infinite tame end. The existence of this family, besides completing the proof of 9.4.1, will show that a geometrically tame end has uniform geometry, and it will lead us to a better understanding of $\mathcal{ML}_0(S)$.

We will work with surfaces which are totally geodesic near their cusps, on esthetic grounds. Our basic parameter will be a family of compactly supported geodesic laminations in $\mathcal{ML}_0(S)$. The first step is to understand when a family of uncrumpled surfaces realizing these laminations is continuous and when discontinuous.

DEFINITION 9.5.1. For a lamination $\gamma \in \mathcal{ML}_0(S)$, let $\mathcal{T}_\gamma$ be the limit set in $\mathcal{GL}(S)$ of a neighborhood system for γ in $\mathcal{ML}_0(S)$. ($\mathcal{T}_\gamma$ is the "qualitative tangent space" of $\mathcal{ML}_0(S)$ at γ).

Let $\overline{\mathcal{ML}}_0(S)$ denote the closure of the image of $\mathcal{ML}_0(S)$ in $\mathcal{GL}(S)$. Clearly $\overline{\mathcal{ML}}_0(S)$ consists of laminations with compact support, but not every lamination with compact support is in $\overline{\mathcal{ML}}_0(S)$:

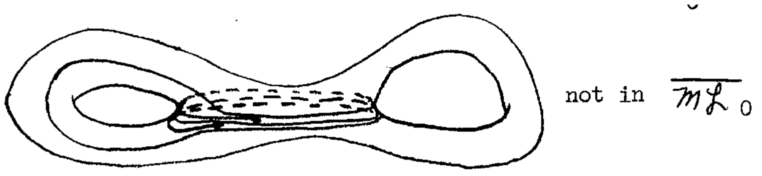

Every element of $\overline{\mathcal{ML}}_0$ is in $\mathcal{T}_\gamma$ for some $\gamma \in \mathcal{ML}_0$. Let us say that an element $\gamma \in \overline{\mathcal{ML}}_0$ is *essentially complete* if γ is a maximal element of $\overline{\mathcal{ML}}_0$. If $\gamma \in \mathcal{ML}_0$, then γ is essentially complete if and only if $\mathcal{T}_\gamma = \gamma$. A lamination γ is maximal among all compactly supported laminations if and only if each region of $S - \gamma$ is an asymptotic 9.30
triangle or a neighborhood of a cusp of S with one cusp on its boundary—a punctured monogon.

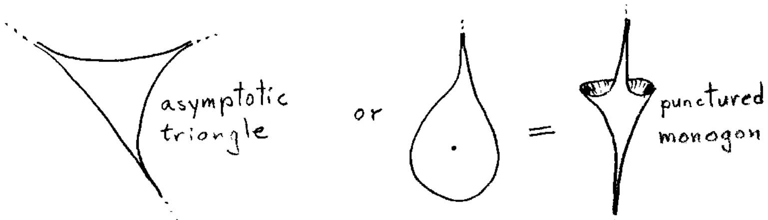

(These are the only possible regions with area π which are simply connected or whose fundamental group is peripheral.) Clearly, if $S - \gamma$ consists of such regions, then γ is essentially complete. There is one special case when essentially complete laminations are not of this form; we shall analyze this case first.

PROPOSITION 9.5.2. *Let $T - p$ denote the punctured torus. An element*

$$\gamma \in \overline{\mathcal{ML}}_0(T - p)$$

is essentially complete if and only if $(T - p) - \gamma$ is a punctured bigon.

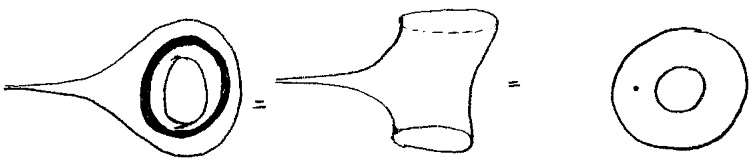

If $\gamma \in \mathcal{ML}_0(T - p)$, then either γ has a single leaf (which is closed), or every leaf of γ is non-compact and dense, in which case γ is essentially complete. If γ has a single closed leaf, then $\mathcal{T}_\gamma$ consists of γ and two other laminations: