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

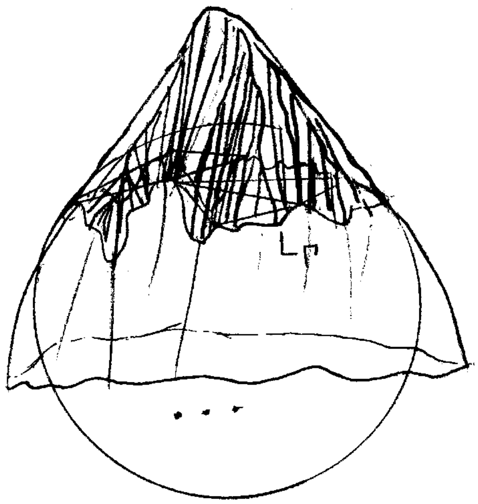

8.33

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

8.7. Quasi-Fuchsian groups

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

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

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

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

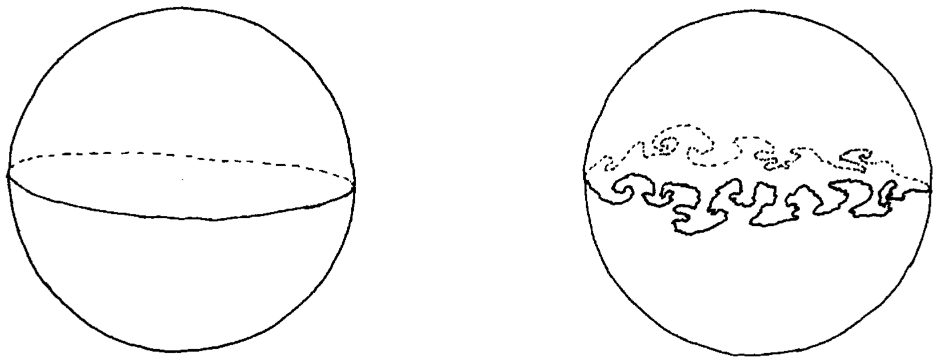

Definition 8.7.1. The Kleinian group Γ is called a *quasi-Fuchsian group* if L_Γ is topologically S^1.

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

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

Proof. Clearly (iii) $\Longrightarrow$ (i) $\Longrightarrow$ (ii). To show (ii) $\Longrightarrow$ (iii), consider

$$O_\Gamma = (H^3 \cup D_\Gamma)/\Gamma.$$

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

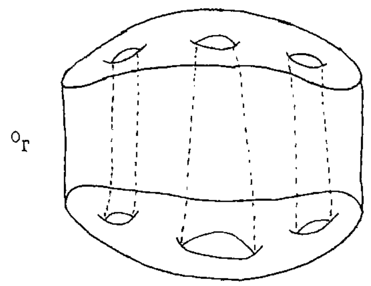

$\Gamma = \pi_1(O_\Gamma) \approx \pi_1(N_1) \approx \pi_1(N_2)$. By a well-known theorem about three-manifolds (see Hempel for a proof), this implies that O_Γ is homeomorphic to $N_1 \times I$. By the above discussion, this implies that Γ' is quasi-conformally conjugate to a Fuchsian group. A similar argument applies if O_Γ has one boundary component; in that case, O_Γ is the orientable interval bundle over a non-orientable surface. The reverse implication is clear. □

EXAMPLE 8.7.3 (Mickey mouse). Consider a hyperbolic structure on a surface of genus two. Let us construct a deformation of the corresponding Fuchsian group by bending along a single closed geodesci γ by an angle of $\pi/2$. This

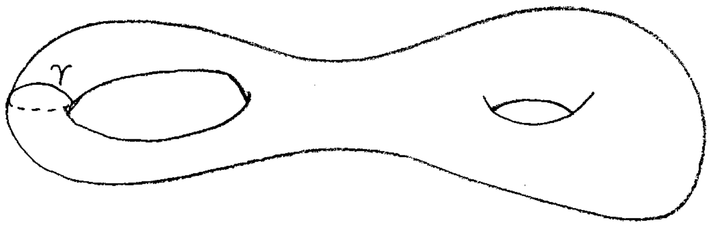

8.36

will give rise to a quasi-Fuchsian group if the geodesic is short enough. We may visualize the limit set by imagining bending a hyperbolic plane along the lifts of γ, one by one.

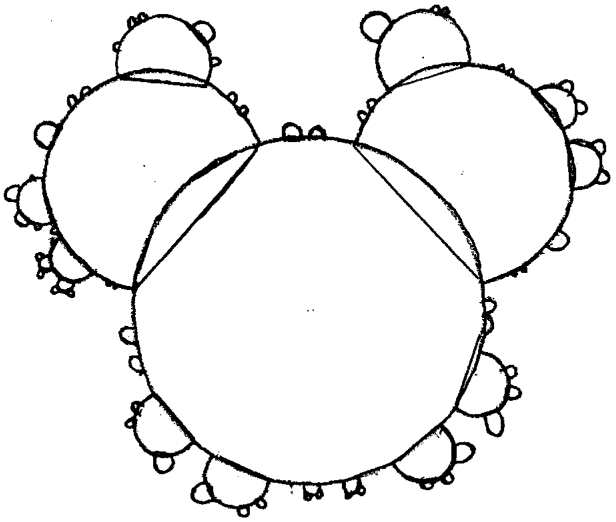

We want to understand how the geometry changes as we deform quasi-Fuchsian groups. Even though the topology doesn't change, geometrically things can become very complicated. For example, given any $\epsilon > 0$, there is a quasi-Fuchsian group Γ whose limit set L_Γ is ϵ-dense in S^2, and there are limits of quasi-Fuchsian groups with $L_\Gamma = S^2$.

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

M_Γ

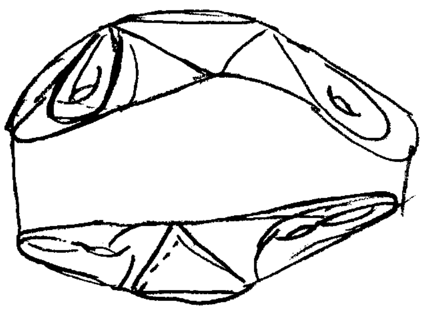

We also need to analyze intermediate surfaces in M_Γ. For example, what kinds of nice surfaces are embedded (or immersed) in M_Γ? Are there isometrically embedded cross sections? Are there cross sections of bounded area near any point in M_Γ?

Here are some ways to map in surfaces.

(a) Take the abstract surface N^2, and choose a "triangulation" of N with one vertex. Choose an arbitrary map of N into M. Then straighten the map (see §6.1).

This is a fairly good way to map in a surface, since the surface is hyperbolic away from the vertex. There may be positive curvature concentrated at the vertex, however, since the sum of the angles around the vertex may be quite small. This map can be changed by moving the image of the vertex in M or by changing the triangulation on N.

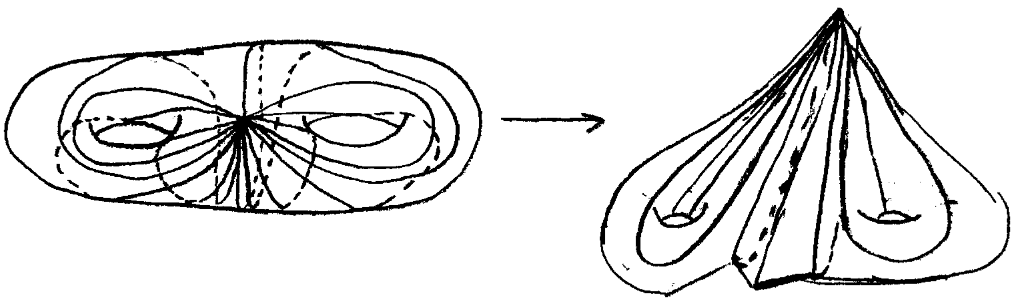

8.38

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

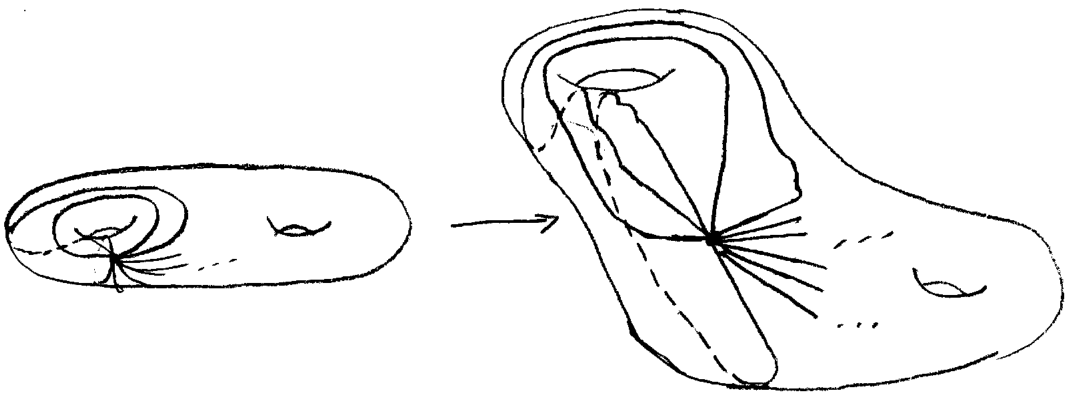

in the loop to the unique geodesic in M in its free homotopy class (this uses a homeomorphism of M to $N \times I$). Now extend this as in (a) to a piecewise straight map $f : N \to M$. The sum of the angles around the vertex is at least 2π, since there is a straight line segment going through the vertex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

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

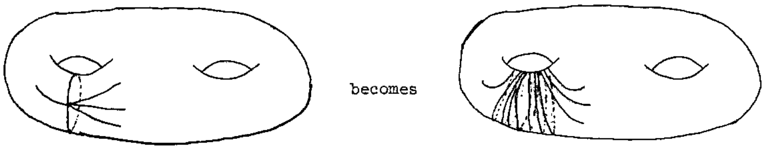

around the loop, more and more, converging, in the limit, to a geodesic laminated set. The surface N maps into M at each finite stage, and this carries over in the limit to an isometric embedding of a hyperbolic surface. The triangles with an edge on the fixed loop have disappeared in the limit. Compare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

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

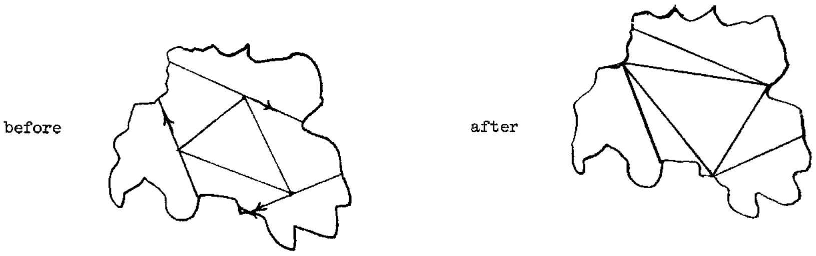

8.40

endpoints are in L_Γ). The line segments connect different conjugates of the geodesic, and the triangles either connect three distinct conjugates or two conjugates (when the original loop is an edge of the triangle). As we isotope the vertex around the loop, the image vertices slide along the geodesic (and its conjugates), and in the limit the triangles become asymptotic (and the triangles connecting only two conjugates disappear).

The above method works because the complement of the geodesic lamination (obtained by spinning the triangulation) consists solely of asymptotic triangles. Here is a more general method of mapping in a surface N by using geodesic laminations.

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

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

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

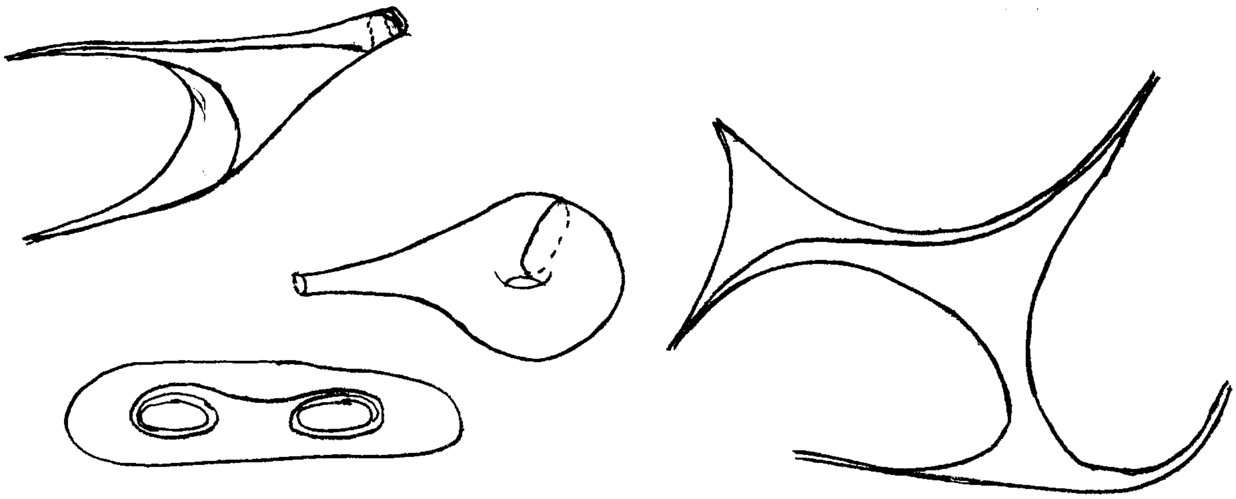

(being careful to add a *simple* geodesic). Either α separates A into two pieces (each of which has less area) or α does not separate A (in which case, cutting along α reduces the rank of the homology. Continuing inductively, after a finite number of steps A separates into asymptotic triangles. $\square$

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

Proposition 8.7.7. *Let S be an oriented hyperbolic surface, and Γ a quasi-Fuchsian group isomorphic to $\pi_1 S$. For every complete geodesic lamination γ on S, there is a unique hyperbolic surface $S' \approx S$ and an isometric map $f : S' \to M_\Gamma$ which is straight (totally geodesic) in the complement of γ.* (γ here denotes the corresponding geodesic lamination on any hyperbolic surface homeomorphic to S.)

Remark. By an isometric map $f : M_1 \to M_2$ from one Riemannian manifold to
another, we mean that for every rectifiable path α_t in M_1, $f \circ \alpha_t$ is rectifiable and 8.42
has the same length as α_t. When f is differentiable, this means that df preserves lengths of tangent vectors. We shall be dealing with maps which are not usually differentiable, however. Our maps are likely not even to be local embeddings. A cross-section of the image of a surface mapped in by method (c) has two polygonal spiral branches, if the closed geodesic corresponds to a covering transformation which is not a pure translation:

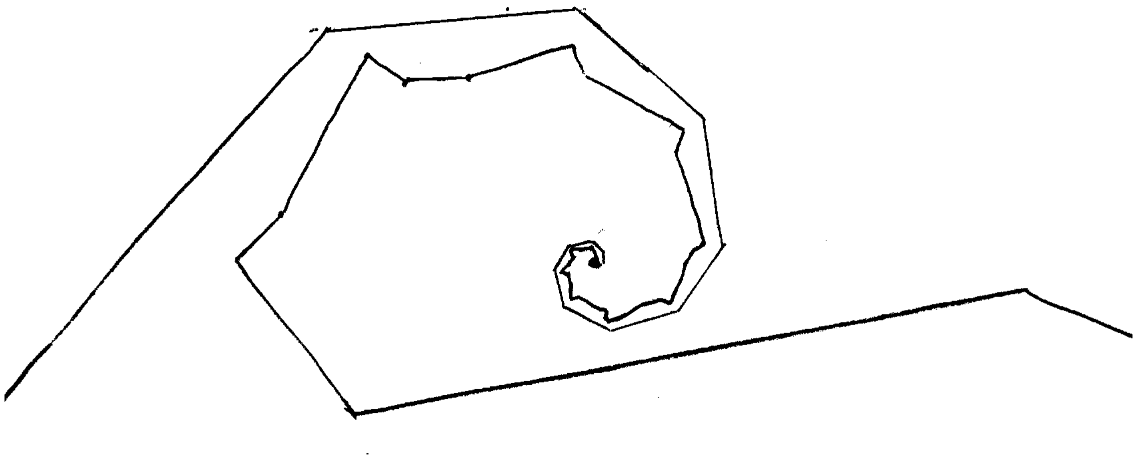

(This picture is obtained by considering triangles in H^3 asymptotic to a loxodromic axis, together with their translates.)

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

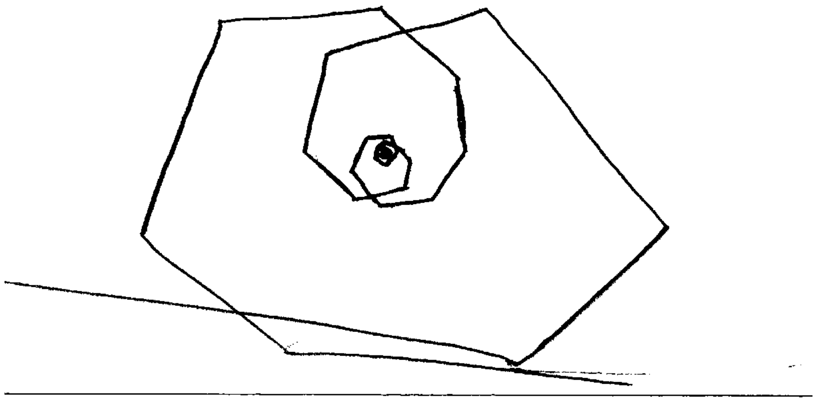

8.43

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

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

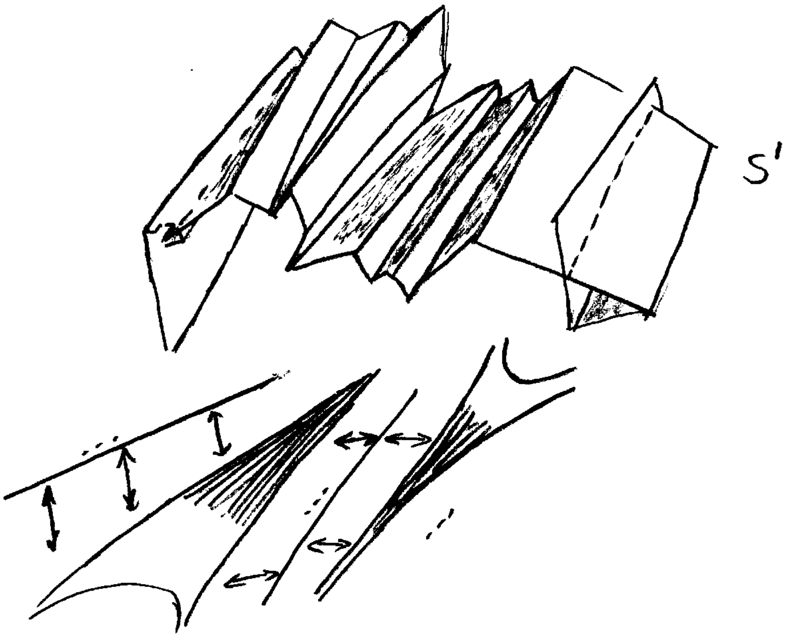

8.44

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

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

8.8. Uncrumpled surfaces

There is a large qualitative difference between a crumpled sheet of paper and one which is only wrinkled or crinkled. Crumpled paper has fold lines or bending lines going any which way, often converging in bad points.

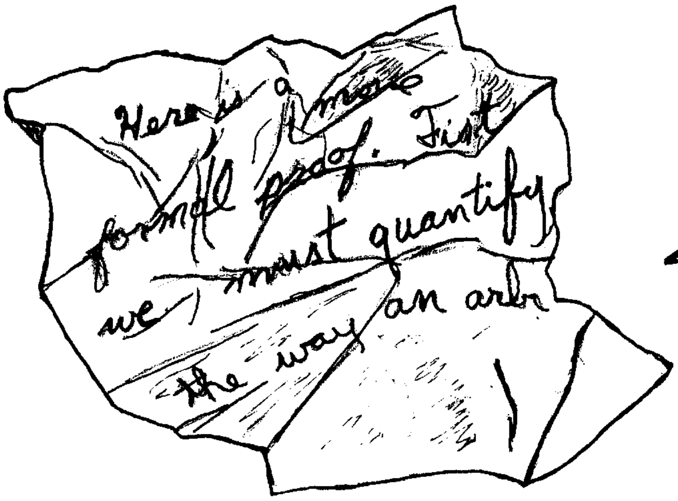

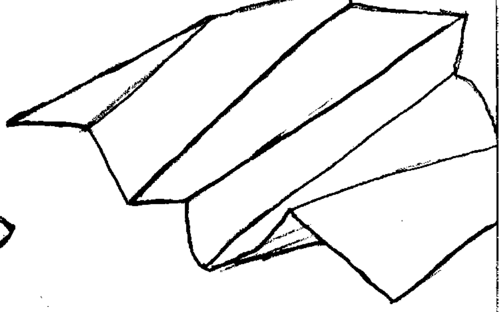

8.45

Definition 8.8.1. An *uncrumpled surface* in a hyperbolic three-manifold N is a complete hyperbolic surface S of finite area, together with an isometric map $f : S \to N$ such that every $x \in S$ is in the interior of some straight line segment which is mapped by f to a straight line segment. Also, f must take every cusp of S to a cusp of N.

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

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

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

Proof. If $x \in S - \gamma$, then there are two transverse line segments through x mapped to line segments. Consider any quadrilateral about x with vertices on these segments; since f does not increase distances, the quadrilateral must be mapped to a plane. Hence, a neighborhood of x is mapped to a plane.

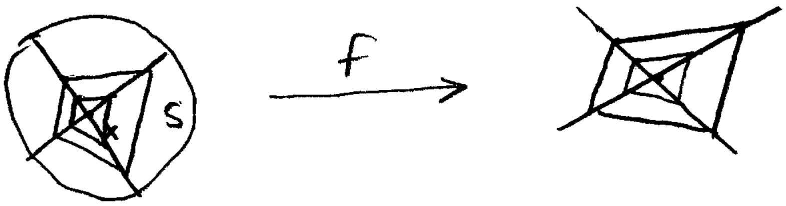

8.46

Consider now any point $x \in \gamma$, and let α be the unique line segment through x which is mapped straight. Let α be extended indefinitely on S. Suppose there were some point y on α in the interior of some line segment $\beta \not\subset \alpha$ which is mapped straight. One may assume that the segment $\overline{xy}$ of α is mapped straight. Then, by considering long skinny triangles with two vertices on β and one vertex on α, it would follow that a neighborhood of x is mapped to a plane—a contradiction.

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

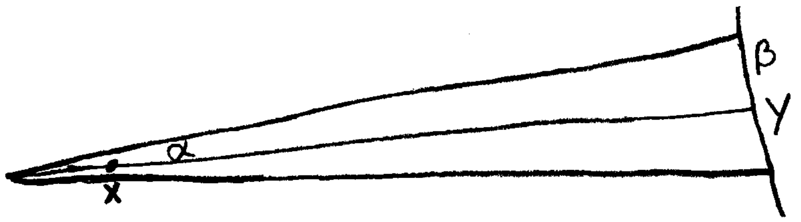

If $U = S \xrightarrow{f} N$ is an uncrumpled surface, then this geodesic lamination $\gamma \subset S$ (which consists of points where U is not locally flat) is the *wrinkling locus* $\omega(U)$.

The *modular space* $\mathcal{M}(S)$ of a surface S of negative Euler characteristic is the space of hyperbolic surfaces with finite area which are homeomorphic to S. In other words, $\mathcal{M}(S)$ is the Teichmüller space $\mathcal{T}(S)$ modulo the action of the group of home- 8.47
omorphisms of S.

PROPOSITION 8.8.3 (Mumford). *For a surface S, the set $A_\epsilon \subset \mathcal{M}(S)$ consisting of surfaces with no geodesic shorter than ϵ is compact.*

PROOF. By the Gauss–Bonnet theorem, all surfaces in $\mathcal{M}(S)$ have the same area. Every non-compact component of $S_{(0,\epsilon]}$ is isometric to a standard model, so the result follows as the two-dimensional version of a part of 5.12. (It is also not hard to give a more direct specifically two-dimensional geometric argument.) $\square$

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

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

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

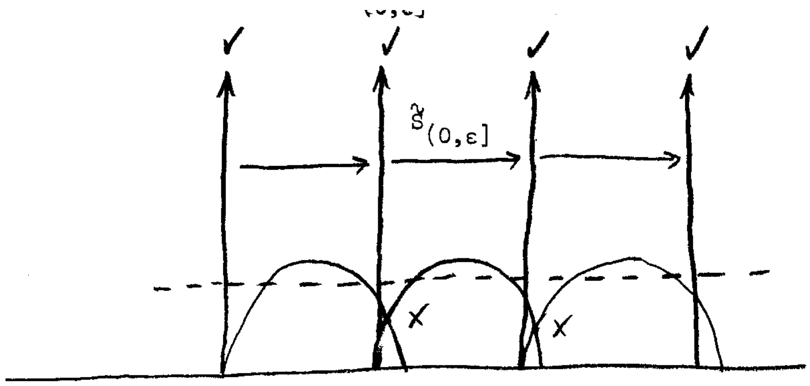

8.48

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

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

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

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

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

If neither point lies in a non-compact component of $S_{(0,\epsilon]}$, one can assume, furthermore, that p does not intersect these components. Let $K' \subset N$ be the set of points which are connected to K by paths whose total length outside compact components of $N_{(0,\epsilon]}$ is bounded by d. Clearly K' is compact and an uncrumpled surface of W must have image in K', except for horoball neighborhoods of its cusps. 8.49

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

Corollary 8.8.6. (a) *Let S be any closed hyperbolic surface, and N any closed hyperbolic manifold. There are only finitely many conjugacy classes of subgroups $G \subset \pi_1 N$ isomorphic to $\pi_1 S$.*

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

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

8.50

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

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

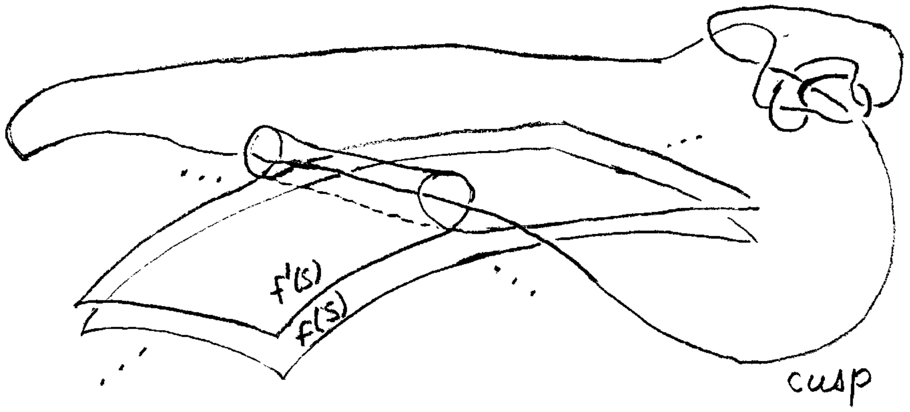

In place of 8.8.5, there is a compactness statement in the topology of geometric convergence provided each component of $S_{[\epsilon,\infty)}$ is required to intersect K. One would allow S to converge to a surface where a simple closed geodesic is pinched to yield a pair of cusps. From this, one deduces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

isomorphic to S up to the operations of conjugacy, and wrapping a surface carrying G around cusps.

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

8.9. The structure of geodesic laminations: train tracks

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

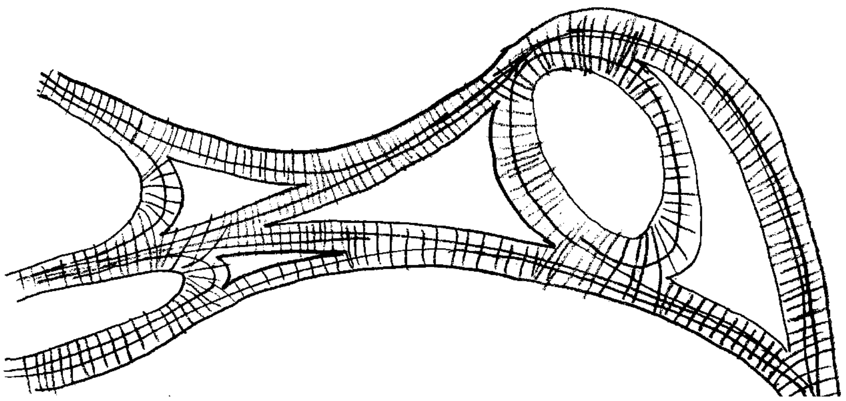

Imagine squeezing the nearly parallel strands of γ in each corridor to a single strand. One obtains a train track τ (with switches) which approximates γ. Each leaf of γ may be imagined as the path of a train running around along τ.

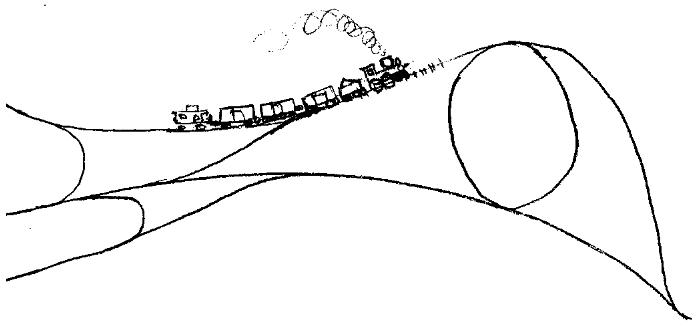

8.52

Here is a construction which gives a precise and nice sequence of train track approximations of γ. Consider a complementary region R in $S - \gamma$. The double dR is a hyperbolic surface of finite area, so $(dR)_{(0,2\epsilon]}$ has a simple structure: it consists of neighborhoods of geodesics shorter than 2ϵ and of cusps. In each such neighborhood there is a canonical foliation by curves of constant curvature: horocycles about a cusp or equidistant curves about a short geodesic. Transfer this foliation to R, and then to S. This yields a foliation $\mathcal{F}$ in the subset of S where leaves of γ are not farther than 2ϵ apart. (A local vector field tangent to $\mathcal{F}$ is Lipschitz, so it is integrable; this is why $\mathcal{F}$ exists. If γ has no leaves tending toward a cusp, then we can make all the leaves of $\mathcal{F}$ be arbitrarily short arcs by making ϵ sufficiently small. If γ has leaves tending toward a cusp, then there can be only finitely many such leaves, since there is an upper bound to the total number of cusps of the complementary regions. Erase all parts of $\mathcal{F}$ in a cusp of a region tending toward a cusp of S; again, when ϵ is sufficiently small all leaves of $\mathcal{F}$ will be short arcs. The space obtained by collapsing all arcs of $\mathcal{F}$ to a point is a surface S' homeomorphic to S, and the image of γ is a train track τ_ϵ on S'. Observe that each switch of τ_ϵ comes from a boundary component of some $dR_{(0,2\epsilon]}$. In particular, there is a uniform bound to the number of switches. From this it is easy to see that there are only finitely many possible types of τ_ϵ, up to homeomorphisms of S' (not necessarily homotopic to the identity). 8.53

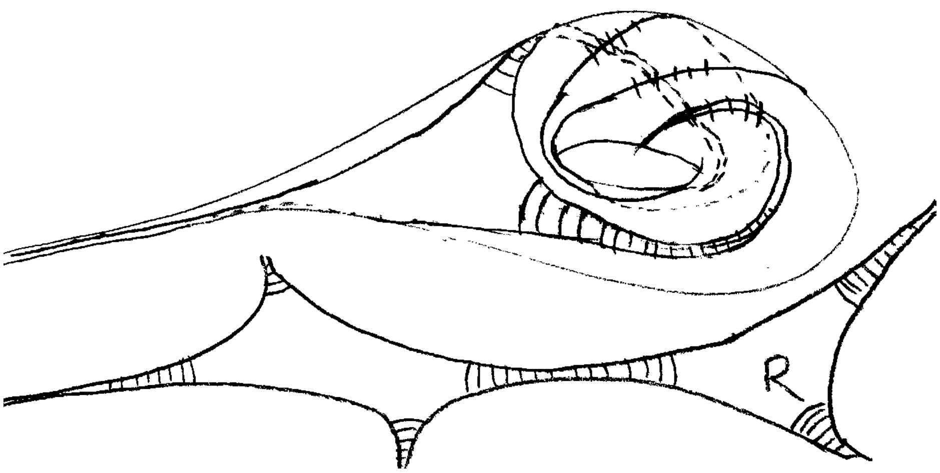

In working with actual geodesic laminations, it is better to use more arbitrary train track approximations, and simply sketch pictures; the train tracks are analogous to decimal approximations of real numbers.

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

DEFINITIONS 8.9.1. A *train track* on a differentiable surface S is an embedded graph τ on S. The edges (branch lines) of τ must be C^1, and all edges at a given vertex (switch) must be tangent. If S has "cusps", τ may have open edges tending toward the cusps. Dead ends are not permitted. (Each vertex v must be in the interior of a C^1 interval on τ through v.) Furthermore, for each component R of $S - \tau$, the double dR of R along the interiors of edges of ∂R must have negative Euler characteristic. A lamination γ on S is *carried* by τ if there is a differentiable map $f : S \to S$ homotopic to the identity taking γ to τ and non-singular on the tangent spaces of the leaves of γ. (In other words, the leaves of γ are trains running around on τ.) The lamination γ is *compatible* with τ if τ can be enlarged to a train track τ' which carries γ. 8.54

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

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

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

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

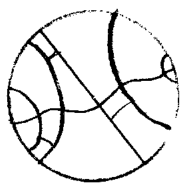

8.55

□

PROPOSITION 8.9.4. *A lamination λ of a surface S is isotopic to a geodesic lamination if and only if*

(a) *λ is carried by some train track τ, and*
(b) *no two leaves of λ take the same (bi-infinite) path on τ.*

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

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

When a lamination λ has an invariant measure μ, this gives a way to associate a number $\mu(b)$ to each branch line b of any train track which dominates γ: $\mu(b)$ is just 8.56
the transverse measure of the leaves of λ collapsed to a point on b. At a switch, the sum of the "entering" numbers equals the sum of the "exiting" numbers.

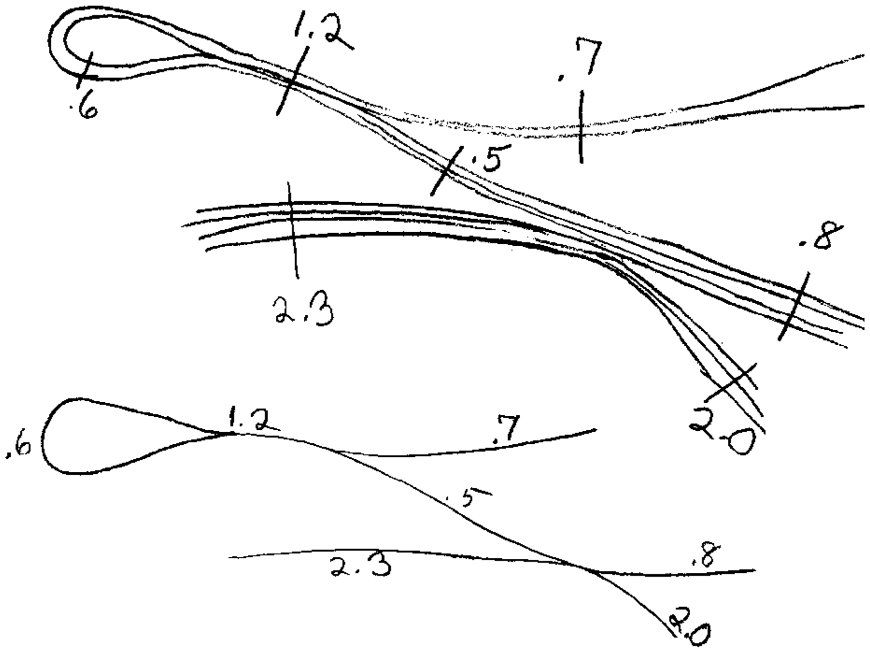

Conversely, any assignment of numbers satisfying the switch condition determines a unique geodesic lamination with transverse measure: first widen each branch line b of τ to a corridor of constant width $\mu(b)$, and give it a foliation $\mathcal{G}$ by equally spaced lines.

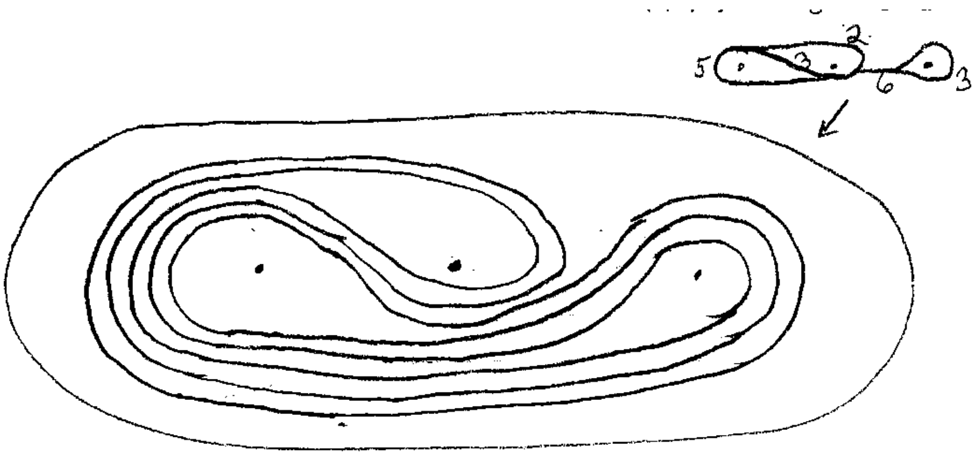

8.57

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

8.10. Realizing laminations in three-manifolds

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

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

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

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

8.58

In Section 8.5, we touched on the notion of geometric convergence of geodesic laminations. The *geometric topology* on geodesic laminations is the topology of geometric convergence, that is, a neighborhood of γ consists of laminations γ' which have leaves near every point of γ, and nearly parallel to the leaves of γ. If γ_1 and γ_2 are disjoint simple closed curves, then $\gamma_1 \cup \gamma_2$ is in every neighborhood of γ_1 as well as in every neighborhood of γ_2. The space of geodesic laminations on S with the geometric topology we shall denote $\mathcal{GL}$. The geodesic laminations compatible with train track approximations of γ give a neighborhood basis for γ.

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

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

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

PROPOSITION 8.10.3. *The natural map $\mathcal{ML} \to \mathcal{GL}$ is continuous.* 8.59

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

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

PROOF OF 8.10.4. An interval which is bent cannot suddenly become straight. Away from any cusps, there is a positive infimum to the "amount" of bending of an interval of length ϵ which intersects the wrinkling locus $w(S)$ in its middle third, and makes an angle of at least ϵ with $w(S)$. (The "amount" of bending can be measure, say, by the different between the length of α and the distance between the image endpoints.) All such arcs must still cross $w(S')$ for any nearby uncrumpled surface S'. $\square$

When S has cusps, we are also interested in measures supported on compact geodesic laminations. We denote this space by $\mathcal{ML}_0(S)$. If (τ, μ) is a train track *description* for (γ, μ), *where $\mu(b) \neq 0$ for any branch of* τ, then neighborhoods for (γ, μ) are described by $\{(\tau', \mu')\}$, where $\tau \subset \tau'$ and $|\mu(b) - \mu'(b)| < \epsilon$. (If b is a branch of τ' not in τ, then $\mu(b) = 0$ by definition.)

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

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

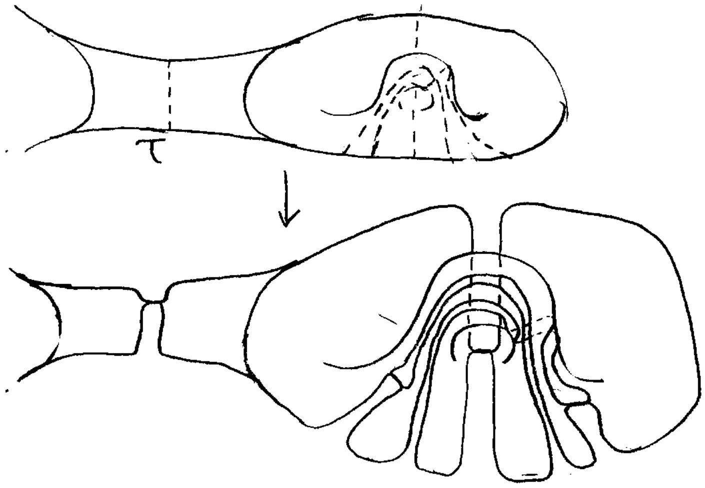

This implies that a neighborhood of (γ, μ) is parametrized by a finite number of real parameters. Thus, $\mathcal{ML}(S)$ is a manifold. Similarly, when S has cusps, $\mathcal{ML}(S)$ is a manifold with boundary $\mathcal{ML}_0(S)$.

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 8.61
as follows. First, note that

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

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