Measured Lamination Spaces for 3-Manifolds

Allen Hatcher

One of Thurston's important contributions to surface theory is his construction of measured lamination spaces. These spaces originally arose because their projectivizations form natural boundaries for Teichmüller spaces, but they are also of interest purely topologically. For example, they provide a nice global framework in which to view all the isotopy classes of simple closed curves in a surface. In view of the many parallels between surface theory and 3-manifold theory, it is natural to ask whether 3-manifolds also have measured lamination spaces. The idea would be that, after projectivization, the rational points of the measured lamination space of a given 3-manifold M would be the isotopy classes of incompressible surfaces in M, and the remaining points would be isotopy classes of measured laminations in M satisfying some sort of incompressibility conditions.

There is a 1988 paper of Oertel [O1] which takes some first steps in this direction. At about the same time I also put some effort into a project of working out the technical details of this theory. Unfortunately these details turned out to be much more complicated than I would have liked, and there were no striking applications on the horizon, so the project was eventually abandoned. Still, there are some nice ideas here which may some day prove useful, so it may be worthwhile to make some of this material available, even without full proofs of all the stated results.

The measured lamination space of a surface can be developed either geometrically, using geodesic laminations on a surface of constant curvature, or purely topologically. For 3-manifolds there is less geometric structure available to study incompressible surfaces, so it seems necessary to take a more topological viewpoint. Therefore we take as our model the more topological approach to measured lamination spaces of surfaces, as developed for example in [H1].

There are two or three purely topological constructions of the measured lamination space $ML(M)$ of a compact orientable surface M. Once one has defined precisely what $ML(M)$ is as a set (isotopy classes of certain measured laminations in M), one can topologize this set in the following ways.

— *Via length functions*: Each closed loop γ in M determines a function $\ell_\gamma : ML(M) \to [0,\infty)$ assigning to a measured lamination L the minimum length, with respect to the given measure on L, of all loops homotopic to γ. Letting γ vary over the set $\mathcal{C}$ of all conjugacy classes in $\pi_1(M)$, one has a total length function $\ell : ML(M) \to [0,\infty)^{\mathcal{C}}$.

After one proves that ℓ is injective, one identifies $ML(M)$ with its image under ℓ and gives it the subspace topology from the product topology on $[0,\infty)^{\mathcal{C}}$.

— *Via* $\mathbb{R}$*-trees*: If $\widetilde{L}$ is the preimage of $L \in ML(M)$ in the universal cover $\widetilde{M}$, then there is an $\mathbb{R}$-tree T_L 'dual' to $\widetilde{L}$ in $\widetilde{M}$, with an action of $\pi_1(M)$ on T_L by isometries, induced by the deck transformations of $\widetilde{M}$. The map $L \mapsto T_L$ embeds $ML(M)$ in the space of actions of $\pi_1(M)$ on $\mathbb{R}$-trees, and one gives it the induced topology. This approach is in fact not essentially different from the length function viewpoint since the space of actions of $\pi_1(M)$ on $\mathbb{R}$-trees is topologized via length functions associated to such actions.

— *Via train tracks*: Given a train track τ in M satisfying certain nontriviality conditions, then each assignment of positive weights to the various sectors of τ, subject to compatibility conditions at the branching loci, defines a measured lamination. The weights vary over a convex cone, and these cones give coordinate charts for a manifold structure on $ML(M)$, at least if one deletes the empty lamination.

It is then a theorem that the different topologies agree on $ML(M)$. To describe the global topological structure of $ML(M)$ it is convenient to consider the projectivization $PL(M)$ of $ML(M)$ obtained by deleting the empy lamination and factoring out by scalar multiplication of the measure. Then $ML(M)$ is the quotient of $PL(M) \times [0,\infty)$ obtained by collapsing $PL(M) \times \{0\}$ to a point. Using the train track approach one proves that when M is closed and of genus $g > 1$, the space $PL(M)$ is a sphere of dimension $6g-7$. For the torus a direct argument shows that $PL(M)$ is a circle. When M is not closed, $PL(M)$ is the join $PL_0(M) * \Delta^{b-1}$ where b is the number of boundary components of M and $PL_0(M) \approx S^{6g+2b-7}$ is the subspace of $PL(M)$ consisting of laminations disjoint from the boundary of M.

When M is a compact orientable 3-manifold, the constructions of $ML(M)$ via length function and dual $\mathbb{R}$-trees go through as one would hope, though the proof that the total length function $\ell\colon ML(M) \to [0,\infty)^{\mathcal{C}}$ is injective takes a fair amount of work. The analog of the train track approach in one lower dimension is to use branched surfaces, for which a well-developed theory exists. The global structure of the projectivization $PL(M)$ is however quite a bit more complicated than in one lower dimension. Branched surfaces give coordinate charts for a decomposition of $PL(M)$ into finitely many strata which are piecewise-linear manifolds of various dimensions. The frontier of a stratum is contained in strata of lower dimension, but the novel feature is that for some manifolds M, parts of the boundaries of some strata are missing. Thus $ML(M)$ need not be a closed subspace of $[0,\infty)^{\mathcal{C}}$, and $PL(M)$ may not be compact. Put another way, one can have a convergent sequence of $\mathbb{R}$-trees dual to measured laminations whose limit is an $\mathbb{R}$-tree not dual to any measured lamination. This does not happen in one lower dimension. What we expect to be true is that $PL(M)$ has a natural compactification $\overline{PL}(M)$ in terms of projective classes of actions on $\mathbb{R}$-trees, such that $\overline{PL}(M)$ is a finite polyhedron containing a subpolyhedron $\partial PL(M)$ with $PL(M) = \overline{PL}(M) - \partial PL(M)$.

To illustrate, let us consider the case that M is the product $F \times S^1$ with F a closed orientable surface of genus $g \geq 1$. When $g = 1$ everything is nice: M is the 3-torus, so incompressible surfaces are subtori embedded linearly, and more generally, measured incompressible laminations are just standard linear foliations of M. Such a foliation is uniquely determined by its orthogonal line field, so $PL(M)$ can be identified with the space of lines through the origin in $\mathbb{R}^3$, which is $\mathbb{R}\mathrm{P}^2$. We can also think of this as the projectivization of $H_2(M;\mathbb{R}) = \mathbb{R}^3$, corresponding to the fact that linear subtori of M are uniquely determined by their homology class, modulo orientations. The corresponding dual $\mathbb{R}$-trees are all isomorphic to $\mathbb{R}$.

When M is a product $F \times S^1$ with F of genus $g > 1$, things are more complicated. It is classical that all incompressible surfaces can be isotoped to be either horizontal or vertical, i.e., either transverse to the circle fibers of the product structure on M, or a union of this fibers. The same is true of incompressible measured laminations. The horizontal laminations can be taken to be measured foliations transverse to fibers, with dual $\mathbb{R}$-trees isomorphic to $\mathbb{R}$. These foliations are determined by their homology classes, modulo sign. We have $H_2(M;\mathbb{R}) = \mathbb{R}^{2g+1}$, and the vertical homology classes form a linear subspace $\mathbb{R}^{2g}$. The remaining classes are horizontal, so after projectivization we see that the horizontal elements of $PL(M)$ form an open cell $e^{2g} = \mathbb{R}\mathrm{P}^{2g} - \mathbb{R}\mathrm{P}^{2g-1}$. The vertical measured laminations, on the other hand, form a copy of $ML(F)$, so after projectivization these give a subspace $S^{6g-7} \subset PL(M)$. Set-theoretically we thus have $ML(M)$ as the disjoint union $S^{6g-7} \amalg e^{2g}$, and it is also topologically the disjoint union since the closure of e^{2g} in the space of actions of $\pi_1(M)$ on $\mathbb{R}$-trees is $\mathbb{R}\mathrm{P}^{2g}$, the projectivization of $H_2(M;\mathbb{R}) = \mathbb{R}^{2g+1}$, actions on $\mathbb{R}$, and the points of S^{6g-7} are actions on more complicated trees, the dual trees of measured laminations in F.

A similar analysis holds more generally for Seifert-fibered manifolds. In all these examples incompressible tori play an essential role, so one might wonder whether $PL(M)$ is compact when M is atoroidal. Unfortunately this is not true. Examples are provided by 2-bridge link exteriors; see [FH].

1. Incompressible Measured Laminations

Throughout, M will be a compact, orientable, irreducible, ∂-irreducible 3-manifold. Our first task is to give a precise definition of the set $ML(M)$, whose points will be isotopy classes of incompressible measured laminations in M.

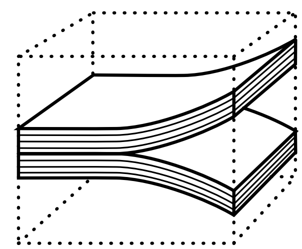

A *prelamination* in M is a compact 3-dimensional submanifold $P \subset M$ with a 2-dimensional singular foliation having local structure as shown in the figure. The singularities of the foliation form a 1-dimensional submanifold $\Sigma \subset \partial P$ along which bifurcation of a leaf occurs, and $\partial P - \Sigma$ consists of two parts: $\partial_M P = P \cap \partial M$, where leaves meet ∂M transversely; and the "horizontal" boundary $\partial_h P$, each component of which is contained in a leaf.

We can eliminate the singularities of the prelamination P by splitting it open along the leaves through the singular points Σ, taking care to damp down the magnitude of the splitting fast enough so that the process converges. The result is a lamination L. In this paper we shall consider only laminations constructed by this procedure. In particular, the laminations we consider have only finitely many complementary components, meeting the lamination in finitely many leaves, called *boundary leaves.* We remark also that laminations in our restricted class have no isolated leaves, as is clear from the construction.

Different (i.e., non-isotopic) prelaminations P can split open to the same lamination L. Namely, if P' is obtained from P by splitting open along compact subsurfaces of the singular leaves of P, then P' splits open to the same lamination as P does. Conversely, it is clear that two prelaminations P splitting open to the same L have a common compactly split prelamination P'. Thus laminations correspond to equivalence classes of prelaminations under the equivalence relation generated by compact splitting along singular leaves.

One could also consider splitting a prelamination P along a compact subsurface of a nonsingular leaf which is not a component of $\partial_h P$. The effect of this on the associated lamination L is to split L along a non-boundary leaf, creating a complementary I-bundle component for the resulting lamination L'. Conversely, a complementary I-bundle component of a lamination L' (with ends of fibers on L', and meeting ∂M in a union of fibers) can be collapsed fiberwise to create a new lamination L. Laminations related by splitting along non-boundary leaves will be regarded as equivalent. A convenient way to factor out by this equivalence relation is to restrict attention to laminations without complementary I-bundle components.

Next we give a definition of incompressibility for a lamination $L \subset M$. Let V be a complementary component of L with the abutting boundary leaves of L adjoined; V is given the weak topology, as a possibly non-compact 3-manifold. The boundary ∂V consists of certain boundary leaves, which we call $\partial_h V$, together with a part of ∂M, $\partial_M V$. Define L to be *incompressible* if for each complementary component V the following conditions hold:

— In V, $\partial_h V$ is incompressible, ∂-incompressible toward $\partial_M V$, and end-incompressible (see below).

— No component of $\partial_h V$ is a sphere or disk, or a compact surface isotopic across V to a component of ∂M.

Here $\partial_h V$ is *end-incompressible* in V if it contains no end-compressing disk: a closed half-plane H embedded properly in V (inverse images of compact sets are compact) with $H \cap \partial_h V = \partial H$ a proper arc in $\partial_h V$ not cutting off a half-plane from $\partial_h V$.

Since the ends of V are I-bundles by the construction of L, it is not hard to see that an end-compressing disk H can be isotoped so that either (a) H is obtained from a compressing disk D for $\partial_h V$ in V by pushing a point of ∂D out to infinity along an arc of $\partial_h V$, or (b) the end of H is a union of fibers in an I-bundle end of V. Note that condition (b) implies automatically that ∂H does not cut off a half-plane from $\partial_h V$. Type (b) end-compressing disks correspond to "infinitely long folds" in leaves of L.

Our definition of an incompressible lamination involves only conditions on the complement of the lamination, so all foliations of M are automatically incompressible according to this definition. One might want a more stringent definition which rules out Reeb components. This is the notion of an essential lamination, as in [GO] or [H2]. For measured laminations, incompressibility can be shown to be equivalent to essentiality.

By a *measure* on a prelamination P we mean a transverse Euclidean structure on P. This assigns a "length" to curves γ in P transverse to leaves, and this length is invariant under homotopy of γ through curves transverse to leaves, with $\partial\gamma$ not moving across leaves during the homotopy. The lamination L obtained from P by splitting open singular leaves inherits a measure from P, assigning lengths to curves in M transverse to L (portions of curves in the complement of L have length zero).

Definition. *$ML(M)$ is the set of isotopy classes of incompressible measured laminations $L \subset M$ without complementary I-bundle components.*

For example, if all leaves of a lamination $L \in ML(M)$ are compact, then each component of L is either (a) an I-bundle, (b) a bundle over S^1, or (c) the union of two nontrivial I-bundles meeting along their ∂I-subbundles. In the latter two cases, the component of L is all of M. In all three cases the component of L can be assigned a total measure: the length of a fiber in (a), the length of a circle cross-section in (b), the sum of the length of fibers in the two I-bundles in (c). If these total measures are all integers, we say L is *integral*. If $\mathcal{S}(M)$ denotes the set of isotopy classes of incompressible surfaces in M (not necessarily connected, but without boundary-parallel components), there is a natural map $\mathcal{S}(M) \to ML(M)$ onto the integral laminations, obtained by first thickening each component of a given $S \in \mathcal{S}(M)$ to a laminated neighborhood of itself with total measure one, then collapsing out complementary I-bundles. This map is very nearly injective, the only exception arising from the fact that in (c) the cross-sections of the two I-bundles are not isotopic in M (they represent different elements of $H_2(M, \partial M; \mathbb{Z}_2)$, in fact), but they yield the same measured lamination.

The conditions for incompressibility of a lamination L can be reformulated as conditions on the associated prelaminations P. Reversing the process of splitting P along singular leaves to form L, we can obtain P from L by collapsing the fibers of an I-bundle W in the complement of L. The boundary ∂W decomposes into three parts: the endpoints of the fibers, $\partial_h W$, lying in boundary leaves of L; $\partial_M W$, the part in ∂M, a union of fibers; and the rest, $\partial_v W$, also a union of fibers. We call W *essential* if it includes the ends of the complement of L, and if, in the complement of L, the components of $\partial_v W$ (annuli and rectangles) are incompressible, ∂-incompressible toward both L and ∂M, and also not ∂M-parallel.

If L is incompressible, it has an essential complementary I-bundle W. Namely, begin by letting W be the ends of the complement of L. Thus each component of W is noncompact. If a component of $\partial_v W$ is compressible in $M - L$, it must be an annulus whose boundary circles bound disks in their leaves of L (since L is incompressible), so the annulus bounds a $D^2 \times I$ in the complement of L, and we enlarge W by adjoining this $D^2 \times I$. Similarly, components of $\partial_v W$ which are ∂M-compressible or ∂M-parallel can be eliminated by enlarging W. Finally, a ∂-compressing disk for $\partial_v W$ toward L would give rise to an end-compressing disk for L since the components of W are noncompact, so such ∂-compressing disks cannot exist.

If P is obtained from the incompressible lamination L by collapsing the essential I-bundle W, then:

— In the complement of P, $\partial_h P$ is incompressible, ∂-incompressible toward ∂M, and without components which are parallel to ∂M, or which are spheres or disks with boundary contained in ∂M.

— P has no complementary monogon, i.e., disk $D \subset M$ with $D \cap P = \partial D$, the circle ∂D having one cusp point where it crosses Σ.

— The components of Σ do not bound disks in their leaves in P.

We call P *incompressible* if it satisfies these conditions.

Conversely, if an incompressible P splits open to a lamination L, then L is incompressible. Namely, P is obtained from L by collapsing an essential I-bundle W. Take a compressing, ∂-compressing, or end-compressing disk for L and simplify its intersections with $\partial_v W$ by the standard arguments, to produce a compressing or ∂-compressing disk for $\partial_h P$, or a monogon. (Details left to the reader.)

Branched Surfaces

For a prelamination P, consider a 1-dimensional foliation on the underlying set of P transverse to the 2-dimensional leaves of P. For convenience we call the 1-dimensional leaves "fibers:" If all these fibers in P are closed intervals, we can collapse them to points to produce a branched surface $B \subset M$ with fibered neighborhood $N = N(B)$ equal to P as a subset of M. In general however, not all fibers in P will be closed intervals, but we can achieve this by splitting P open along finitely many disjoint disks contained in leaves of P and lying either in the interior of P or meeting ∂P in an arc in ∂M. Such disks, and also the holes they create in the resulting P', we call *slits.* Such slits prevent P' from being incompressible, but only in a rather mild way.

A branched surface $B \subset M$ with fibered neighborhood N is *incompressible* if:

— In the complement of N, $\partial_h N$ is incompressible, ∂-incompressible toward ∂M, and without components which are parallel to ∂M, or which are spheres or disks with boundary contained in ∂M.

— N has no complementary monogon, i.e., disk $D \subset M$ with $D \cap N = \partial D$, the circle ∂D having one cusp point where it crosses Σ.

— B has no disk or half-disk of contact. (A disk of contact is a disk $D \subset N$ transverse to fibers, with $\partial D \subset \Sigma$, D lying on the side of Σ away from the abutting sheets of $\partial_h N$. A half-disk of contact is similar, ∂D consisting of an arc in Σ and an arc in ∂M.)

This third condition is a version of the third condition in the definition of incompressibility for prelaminations P which does not use the 2-dimensional leaf structure of P, but rather the transverse fiber structure. It also allows slits. This definition is the same as that used in [FO] and [O2], except for our mild additional requirement that $\partial_h N$ contain no components which are spheres, or disks with boundary in ∂M, or ∂M-parallel.

If B is incompressible and P is a prelamination having underlying set $N(B)$ with slits collapsed, and having leaves transverse to fibers, then clearly P is incompressible. Conversely:

Lemma 1.1. *If B is obtained from an incompressible measured prelamination P by inserting slits and collapsing transverse interval fibers, then B is incompressible.*

Proof: Take a disk or half-disk of contact D for $N(B)$. This lives in P after slits are collapsed. It suffices to isotope D (rel ∂D), staying transverse to fibers, until it lies in a leaf of P. To do this, perturb D to have general position intersection with the leaves of P, so P induces a foliation on D with center and saddle singularities in $int(D)$, with ∂D a leaf. By index considerations, there must be some center singularities. Around a center singularity the leaves form concentric circles. This growing family of circle leaves limits on a leaf which either contains a saddle or is ∂D. In the latter case we can push D vertically to lie in the leaf of ∂D; this is possible if P is measured, though not in general otherwise, e.g., in a Reeb component. If the limit leaf contains a saddle there are the two subcases shown in the next figure. In the first case we can push D vertically to cancel the center

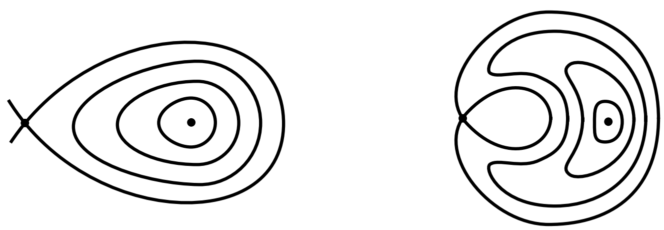

and the saddle. In the second case the limit leaf bounds a disk disjoint from the given center, and we transfer our attention to a center singularity in this disk. In both cases an induction finishes the proof. $\square$

Let $C(B)$ be the set of weight vectors $\alpha = (\alpha_1, \cdots, \alpha_n)$ on B, α_i being the weight assigned to the i^{th} nonsingular sector of B, satisfying the branch equations $\alpha_i + \alpha_j = \alpha_k$ along the branching locus of B, and with each $\alpha_i > 0$. $C(B)$ is thus an open convex polyhedral cone in $\mathbb{R}^n$, whose closure $\overline{C}(B)$ is obtained by allowing coordinate weights $\alpha_i = 0$. To $\alpha \in C(B)$ we associate a measured prelamination $P'_\alpha \subset M$ by first taking a family of parallel sheets of "thickness" α_i over the i^{th} nonsingular sector of B, then glueing these families together over the branching locus of B via the branch equations $\alpha_i + \alpha_j = \alpha_k$. Clearly every measured prelamination whose underlying set is $N(B)$ and with leaves transverse to fibers of $N(B)$ has this form P'_α for some $\alpha \in C(B)$. Let P_α be the prelamination obtained from P'_α by collapsing complementary I-bundle components, and let L_α be the lamination obtained by splitting open the singular leaves of P_α. (If B is incompressible this L_α could just as well be constructed by first splitting P'_α along its singular leaves to produce a lamination L'_α, then collapsing complementary I-bundle components L'_α.)

If $N(B)$ has complementary I-bundle components which are products, then different weights $\alpha \in C(B)$ can give the same lamination L_α since at the stage of constructing P'_α we could push a layer of sheets of P'_α across such a product yielding a prelamination P'_β with $P_\alpha = P_\beta$ and hence $L_\alpha = L_\beta$. Such moves generate a linear equivalence relation on $C(B)$, whose quotient $c(B)$ is therefore also an open convex linear cone in some quotient $\mathbb{R}^k$. If B is incompressible, we thus obtain a map $\varphi_B : c(B) \to ML(M)$. We will show in Proposition 4.1 that φ_B is injective.

We call an incompressible branched surface B *maximal* if it has no essential complementary annuli, rectangles, digons, or half-digons.

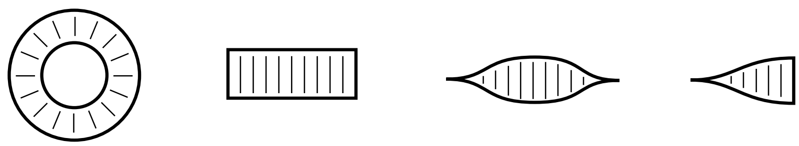

In terms of the associated essential complementary I-bundle W_α for the lamination $L_\alpha \in ML(M)$ with $\alpha \in C(B)$, maximality of B is simply the condition that W_α be maximal up to isotopy among essential I-bundles for L_α. (This condition is independent of α.)

Maximality can easily be achieved by finitely many pinching operations, modifying B by collapsing the fibers of neighborhoods of essential complementary surfaces of the types in the preceding figure. For this collapsing to yield a branched surface, a little care is needed to avoid for example collapsing a fiber whose two endpoints coincide on B. General position suffices for this: first put the boundary curves of an essential complementary surface in general position on B, to have only isolated intersections, then perturb the fiber structure near these isolated intersections to break any cycle of fibers.

Note that laminations L_α carried by B with positive weights before pinching are still carried with positive weights after pinching, and in fact pinching induces a linear map of the cones $C(B)$ projecting to a linear inclusion of the quotient cones $c(B)$.

Pinching essential annuli, rectangles, digons, or half-digons preserves incompressibility of B since it amounts to enlarging the essential complementary I-bundles W_α so that they stay essential (modulo slits).

2. Length Functions

Fix now an incompressible branched surface B and a weight vector $\alpha \in C(B)$. Consider loops in M which meet the associated prelamination P_α in finitely many paths, each a union of finitely many subpaths which are either vertical (in fibers of P_α) or horizontal (in leaves of P_α); we call such loops PVH — piecewise vertical or horizontal. PVH loops have a length, the total length of the vertical segments, as specified by α. For an arbitrary loop γ in M define $\ell_\gamma(P_\alpha)$ to be the infimum of the lengths of all PVH loops homotopic to γ. We will show in Proposition 2.1 below that this infimum is always achieved, if B is incompressible. We may assume B is maximal if we like, since pinching it to make it maximal clearly has no effect on $\ell_\gamma(P_\alpha)$.

The notion of "vertical" depends on the fiber structure of P_α, which is not really part of the data of P_α. To avoid this we could replace "vertical" by "transverse to leaves of P_α." This leads to the same $\ell_\gamma(P_\alpha)$, clearly. "Piecewise transverse or horizontal" loops are also defined for the associated lamination L_α, and also give the same $\ell_\gamma(L_\alpha) = \ell_\gamma(P_\alpha)$. In particular, $\ell_\gamma(L_\alpha)$ is independent of the choice of the incompressible branched surface B used to construct L_α. If α is integral it is easy to see that $\ell_\gamma(P_\alpha)$ equals the minimum number of points of intersection of the incompressible surface associated to α with loops homotopic to γ. Another elementary observation is that $\ell_\gamma(P_\alpha)$ is linear with respect to scalar multiplication of α.

Shortest Curves

Loops γ which are PVH with respect to P_α are partitioned into segments which are either vertical, horizontal, or outside P_α. We may assume the horizontal segments cross the singular locus (cusp curves) Σ of P_α transversely. In particular the horizontal segments do not switch branches at a cusp. The PVH loop γ is called *taut* if it cannot be deformed to decrease the number of segments outside P_α and if neither of the following moves is possible:

(1) shortening γ by pushing vertically a horizontal segment

(2) decreasing the number of intersections of γ with Σ by taking a horizontal subarc of γ with endpoints on Σ, and deforming this subarc either (a) into Σ, and a little beyond, to a horizontal subarc disjoint from Σ, or (b) into a horizontal path in ∂P_α meeting Σ only in its endpoints.

In (2) the subarc can go outside P_α during the deformation. Examples of type 2 deformations are shown in the next figure.

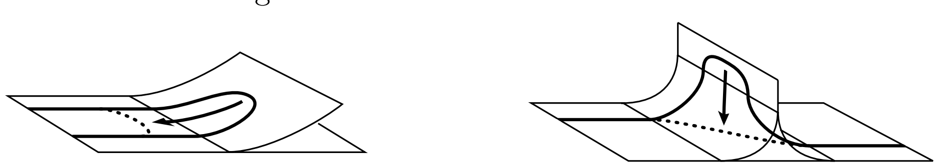

Any PVH γ can be made taut by a finite sequence of type 1 and 2 deformations, after first pushing as many segments outside P_α into P_α as possible. To see this, first note that a sequence of consecutive type 2 moves must be finite since each such move decreases intersections with Σ. So consider a horizontal segment s for which a type 1 move is possible, say in the downward direction. Let δ be the minimum length of the two vertical segments at the ends of s. We could hope to push s down a distance δ, shortening γ by 2δ. This motion may be obstructed by arcs of ∂P_α, however. If ε is the minimum weight α_i of α, then the obstructing arcs within vertical distance ε of s are non-overlapping when projected back up to s. Push s down to the first of these obstructing arcs, which thus becomes a contact arc of γ with ∂P_α. If possible, perform type 2 moves on subarcs of s to eliminate such contact arcs in s. The newly positioned subarcs can have no obstructing arcs within distance ε below them. Then push s down to the next obstructing arc and repeat the preceding step. Continuing in this way, after finitely many steps either we eliminate a vertical segment at one end of γ, or we produce an essential contact arc which cannot be eliminated by type 2 moves, or we shorten γ by at least 2ε. In the last case we repeat the argument. Eventually either we have eliminated a vertical segment of γ or we have a horizontal segment s for which no type 1 or 2 moves are possible. Then do the same with other horizontal segments of γ. After a finite number of steps, γ will be taut.

Proposition 2.1. *Given a maximal incompressible branched surface B and $\alpha \in C(B)$, a loop γ which is taut with respect to P_α achieves the minimum length of loops within its homotopy class.*

Proof: The first step is to reduce to the case that α is rational. Consider an n-parameter variation $\alpha(t) = \alpha_0 + t_1\alpha_1 + \cdots + t_n\alpha_n$ of the given $\alpha = \alpha_0 = \alpha(0)$, $t = (t_1, \cdots, t_n)$, $t_i \geq 0$, where $\alpha_1, \cdots, \alpha_n$ are integral points spanning the closure $\overline{C}(B)$. Corresponding to $\alpha_1, \cdots, \alpha_n$ are surfaces $S_1, \cdots, S_n$ in the interior of $P_\alpha(0)$ transverse to fibers and in general position with respect to each other and to γ, which is taut for $P_\alpha(0)$. The lamination $P_\alpha(t)$ may be obtained from $P_\alpha(0)$ by slitting it open along S_1, inserting a neighborhood of S_1 of vertical thickness t_1, extending the foliation over this neighborhood in the obvious way, then repeating this operation for $S_2, \cdots, S_n$. This converts γ into an n-parameter family of PVH loops $\gamma(t)$ by inserting vertical pieces of length t_i at the points where γ meets S_i. (These points may be in either horizontal or vertical segments of γ.) For $t = 0$, $\gamma(t)$ is taut by assumption. For small t, $\gamma(t)$ can be pulled to a taut $\widehat{\gamma}(t)$ using only small type 1 moves, since no contact arcs can be introduced by small vertical deformations. Thus we may choose the taut $\widehat{\gamma}(t)$ varying continuously with t near 0.

Suppose γ' is a PVH loop homotopic to γ and shorter than γ, with respect to the given P_α. We may extend γ', like γ, to a family $\gamma'(t)$ of PVH loops varying continuously with t. Choosing $\alpha(t)$ rational with t small enough, $\gamma'(t)$ remains shorter than the taut $\widehat{\gamma}(t)$. This reduces us to the case that α is rational. Rescaling, we may assume α is integral, corresponding to an incompressible surface S.

We may obtain P_α from a thickening $N(S)$ of S by collapsing the fibers of an I-bundle A in the complement of $N(S)$. Let γ_1 be a loop projecting to γ under the map which collapses fibers of A, γ_1 being obtained by inserting certain fibers of A. These fibers contribute nothing to the length of γ_1, which is the same as the length of γ, namely the number of points of intersection with S (we may assume γ_1 meets S transversely). Since γ_1 is not shortest in its homotopy class, a well-known elementary argument shows there is a homotopy of a subarc of γ_1 which eliminates two consecutive points of intersection of γ_1 with S. Thus there is a map $f: D^2 \to M$ taking one arc $\partial_+ D^2$ of ∂D^2 to γ_1, the remainder of ∂D^2 being $\partial_- D^2 = f^{-1}(S)$. By deforming f near $\partial_- D^2$, pushing it vertically away from S, we may assume that $f(D^2) \subset M - int(N(S))$. After this deformation $f(\partial_+ D^2)$ has length zero, since originally all the length of $f(\partial_+ D^2)$ came from its endpoints intersecting S. Note that $f(\partial_+ D^2)$ is contained in $\partial N(S) \cup A$ since otherwise we could use f to deform γ to decrease the number of segments outside P_α, contradicting tautness.

We may assume f is transverse to $\partial N(S)$ and to $\partial_v A$. Consider a component arc a of $f^{-1}(\partial_v A)$. If both endpoints of a lie on $\partial_- D^2$, a can be eliminated by rechoosing f near the disk in D^2 cut off by a. The next case is that a has only one endpoint on $\partial_- D^2$. Since the endpoints of $\partial_- D^2$ map into A, there must be another such arc a', the endpoints of a and a' separated on $\partial_- D^2$ by an interval I_- mapping to $\partial N(S) - A$. The interval I_+ in $\partial_+ D^2$ bounded by the other endpoints of a and a' corresponds to a horizontal subarc of γ for which a type 2(b) move is possible using the rectangle in D^2 cut off by a and a' (contrary to tautness) unless $int(I_+)$ is disjoint from $f^{-1}(\partial_v A)$. In the latter case a type 2(a) move will be possible for γ (using maximality of B) unless f maps I_+ to the same component of $\partial N(S)$ as I_-, a and likewise a' having both its endpoints mapped to the same end of an annulus or rectangle component of $\partial_v A$.

Consider next an arc a with both endpoints on $\partial_+ D^2$ which is "nearest" $\partial_- D^2$ among such arcs, so f maps points near a on the side away from $\partial_- D^2$ to points outside A. Then the arc of $\partial_+ D^2$ cut off by a corresponds to a horizontal subarc of γ on which a type 2(a) move is possible, using the subdisk of D^2 cut off by a, again contrary to tautness.

The upshot of all this is that $f(\partial_+ D^2)$ is contained in the union of A with the component of $\partial N(S)$ containing $f(\partial_- D^2)$. Collapsing the fibers of A, this means that a type 1 move was possible on γ. ◻

The same proof shows that a taut PVH path achieves the minimum length within a homotopy class of PVH paths with the same endpoints.

Trees Associated to Measured Laminations

When studying measured laminations it is sometimes convenient to factor out what is happening within leaves and within complementary regions of the lamination. Trees provide a nice way of doing this, just as in one lower dimension.

By a *tree*, or more precisely an $\mathbb{R}$-tree, we mean a metric space (T,d) such that:

(1) Any two points $x, y \in T$ are the endpoints of a unique segment $[x,y]$, i.e., a subset isometric to a closed interval in $\mathbb{R}$.

(2) $[x,y] \cap [x,z] = [x,w]$ for some w.

(3) $[x,y] \cup [y,z] = [x,z]$ if $[x,y] \cap [y,z] = y$.

This is the definition in the foundational paper [AB]. (Axiom (2) actually follows from (1) since d takes values in $\mathbb{R}$, rather than in a more general ordered abelian group.)

Verifying the uniqueness of segments $[x,y]$ with given endpoints can be difficult in practice. This problem is avoided with the following alternative characterization.

Lemma 2.2. *A metric space (T,d) is a tree if there is a function assigning to each unordered pair $x, y \in T$ a segment $[x,y]$ in T, such that (2) and (3) hold.*

Proof: Choose a basepoint $x_0 \in T$. According to Theorem 3.17 of [AB], the following two axioms characterize trees:

(4) Segments $[x_0, x]$ exist for all $x \in T$.

(5) Letting $x \wedge y = \big(d(x,x_0) + d(y,x_0) - d(x,y)\big)/2$, then $x \wedge z \geq min(x \wedge y, y \wedge z)$ for all $x, y, z \in T$.

Under the hypotheses of the Lemma, $x \wedge y$ is $d(x_0, w)$, where $[x_0, x] \cap [x_0, y] = [x_0, w]$. Then (5) follows by simply listing the possible configurations for the segments joining x_0 to x, y, and z. $\square$

Given a lamination $L_\alpha \in ML(M)$, let $\widetilde{L}_\alpha$ be its preimage in the universal cover $\widetilde{M}$ of M. For $x, y \in \widetilde{M}$ define $\widetilde{d}(x,y)$ to be the infimum of the lengths of all PVH paths joining x and y in $\widetilde{M}$. (Strictly speaking, we mean PVH with respect to the lift of the prelamination N_α here, instead of $\widetilde{L}_\alpha$.) As we observed earlier in this section, this infimum is always realized by some PVH path. Clearly $\widetilde{d}$ defines a pseudo-metric on $\widetilde{M}$, being symmetric and satisfying the triangle inequality. Let (T,d) be the associated metric space, obtained by identifying points of $\widetilde{M}$ of $\widetilde{d}$-distance zero apart, namely points in the same leaf or in the closure of the same complementary component of $\widetilde{L}_\alpha$ (since a PVH path of zero length cannot have nontrivial vertical segments).

Proposition 2.3. *(T, d) is a tree.*

Proof: We use the preceding lemma. For the segment $[x, y]$ we take the image in T of a shortest path in $\widetilde{M}$. To verify that this segment is well-defined, and at the same time check axiom (2), consider two shortest paths in $\widetilde{M}$ with the same initial point. The intersection of their images in T is closed, being the intersection of two closed sets. The intersection is also connected. For otherwise we would have an embedded circle in T. This would mean that for some leaf of $\widetilde{L}_\alpha$ there was a closed (PVH) loop in $\widetilde{M}$ meeting this leaf in only one point, transversely. But leaves of $\widetilde{L}_\alpha$ are properly embedded in $\widetilde{M}$. (If a leaf in $\widetilde{L}_\alpha$ were not properly embedded, it would meet some vertical segment in $\widetilde{M}$ infinitely often; but we have seen that vertical segments are length-minimizing.) This is a contradiction.

For (3), let segments $[x, y]$ and $[y, z]$ be given, intersecting only in y. If $[x, z]$ differed from $[x, y] \cup [y, z]$ we would have the configuration shown in the figure, and again T would contain an embedded circle. □

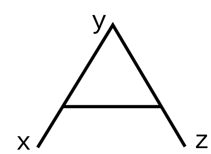

The tree T associated to L_α has an action of $\pi_1 M$ by isometries, coming from deck transformations in $\widetilde{M}$.

Proposition 2.4. *No proper subtree of T is invariant under the action of $\pi_1 M$.*

Proof: Given a point in T, there is a vertical segment s in $\widetilde{M}$ whose image in T contains the given point. Being vertical, s is taut. We claim that s can be extended to a taut path of infinite length in both directions. For this we may as well work with the prelamination $P_\alpha \subset M$ corresponding to L_α. Choose an end of s and extend this vertically until stopped by a component C of $\partial_h P_\alpha$. (If the vertical extension can be carried on infinitely far, we are already done.) There are two cases now. If C contains singular points (cusps) of P_α, extend s by a horizontal segment until it crosses such a cusp point, then continue s vertically in the same direction until it meets another component of $\partial_h P_\alpha$. This extended s can then be perturbed to miss the cusp locus, staying taut, by pushing the new horizontal segment slightly into the interior of P_α. The other case is that C has no cusp points. In this case we first extend s by a segment outside P_α, going to another component of $\partial_h P_\alpha$ if possible, and otherwise returning to C by a path not homotopic (rel endpoints) to a path in C; such a path must exist since the component of $M - P_\alpha$ cut off by C is not a product, in view of the maximality of the branched surface B used to define P_α. Then we extend s by a vertical segment until it meets $\partial_h P_\alpha$ again.

This is the inductive step in extending s. It can be repeated infinitely often if necessary to produce the desired taut bi-infinite line, which we shall still call s. This has infinite length since at each step we are able to lengthen it by at least the minimum weight of N_α.

Since M is compact, each end of s must eventually have vertical segments accumulating somewhere. Then we can rechoose the ends of s so they are eventually periodic, cycling around two taut loops. Lifting to $\widetilde{M}$ and projecting to T, s becomes an embedded $\mathbb{R} \subset T$ whose two ends are part of the invariant axes for the actions on T of the elements of $\pi_1 M$ represented by the two ending loops of s. Any invariant subtree of T must contain these two axes, hence also the embedded $\mathbb{R} \subset T$, which contains the arbitrarily given point of T we started with. □

Piecewise Linearity

We have seen in the proof of Proposition 2.1 that $\ell_\gamma(P_\alpha(t))$ is continuous on the open cone $C(B)$. A considerably stronger result is:

Proposition 2.5. *For B incompressible, the function $\alpha \mapsto \ell_\gamma(P_\alpha)$ is the restriction to the open cone $C(B)$ of a piecewise linear function on the closed cone $\overline{C}(B)$.*

Proof: We may assume B is maximal, since the pinching operation for achieving maximality corresponds to a linear change of coordinates in $C(B)$.

To prove the proposition, we shall deform the family $\gamma(t)$ constructed in the proof of Proposition 2.1 to a family of PVH loops $\widehat{\gamma}(t)$ which are taut for all t, such that:

(i) The segments of $\widehat{\gamma}(t)$ outside $P_\alpha(t)$ are independent of t.

(ii) The vertical segments of $\widehat{\gamma}(t)$ are subsegments of the vertical segments of $\gamma(t)$, varying finitely piecewise linearly with t, i.e., "piecewise" with respect to a partition of the parameter domain $[0,\infty)^n$ by finitely many hyperplanes; and having lengths which are convex functions of t.

(iii) The horizontal segments of $\widehat{\gamma}(t)$ vary piecewise continuously with t, "piecewise" with respect to a partition of the parameter domain by hyperplanes, perhaps infinite in number.

Since $\ell_\gamma(P_\alpha(t))$ is the sum of the lengths of the vertical segments of $\gamma(t)$, the Proposition will follow from (ii).

As in the unparametrized case, pulling $\gamma(t)$ taut is done inductively, one horizontal segment at a time. Since no new vertical segments are introduced in the process, horizontal segments can always be regarded as unions of horizontal segments of the original $\gamma(t)$ ($t \neq 0$). At the inductive step of pulling a segment $s(t)$ taut, we assume all "subsegments" of $s(t)$ have already been pulled taut, and that $s(t)$ lives over some convex subpolyhedron of the parameter domain $[0,\infty)^n$. Pulling $s(t)$ taut involves a number of type 2 moves eliminating inessential contact arcs. The number of such moves is finite for fixed t, as we have seen, and bounded for nearby t by the same argument (though it may be unbounded as t approaches ∞). A given sequence of moves pulling $s(t)$ taut can be chosen to vary continuously with t provided that the relative heights of the inessential contact arcs encountered do not change with t. These relative heights of contact arcs can be measured by comparing their vertical distances from points of $\partial P'_\alpha(t)$, where the prime on P indicates that we have inserted slits corresponding to the slits of B; we call such vertical distances displacements. Displacements of contact arcs are linear functions of t, with constant term a sum of coordinates α_i of $\alpha(0)$ and with coefficient of t_i the number of intersections of Si with the vertical arc from the contact arc to the reference point in $\partial P'\alpha(t)$. Without loss of generality we may assume the weight $\alpha(0)$ is integral, so displacements of contact arcs are $\mathbb{N}$-linear functions. Where two given contact arcs are on the same level thus defines a hyperplane in the parameter domain (if it is not all of the parameter domain). These hyperplanes define a nice polyhedral stratification of the support of $s(t)$, on each stratum of which the process of pulling $s(t)$ taut can be taken to vary continuously, using a single sequence of moves.

A further inductive hypothesis which we make is that displacements of points on horizontal segments of $\gamma(t)$ are finitely piecewise $\mathbb{N}$-linear functions of t. Since $\alpha(0)$ is integral, we may assume the displacements of horizontal segments of $\gamma(0)$ are integral, so before we pull $\gamma(t)$ taut, its horizontal segments have $\mathbb{N}$-linear displacements. We must check that after pulling $s(t)$ taut, its displacement is still finitely piecewise $\mathbb{N}$-linear. At the moment, we can say that after $s(t)$ is pulled taut its displacement varies at least continuously with t. This is because if we then extend it by the vertical segments at either end it becomes a family of taut paths with continuously varying endpoints, and, as we have already observed for families of taut loops, such families of paths have continuously varying lengths.

Toward strengthening this to finite piecewise $\mathbb{N}$-linearity, we first verify that there is a well-defined collection of components of Σ, independent of stratum, such that if for some t, $s(t)$ is pulled taut using type 1 moves, and possibly also type 2 moves, and meets such a component of Σ, then the resulting contact arc is essential. (Strictly speaking, we should be working in the universal cover of M at this point.) We may restrict to rational weights $\alpha(t)$ for this, since within strata everything varies continuously and rational points are dense in strata, the strata being defined by $\mathbb{N}$-linear equations.

During a type 1 move, $s(t)$ at each instant lies in a leaf of $P_\alpha(t)$ and can moreover be lifted to lie in a leaf of the lamination $L_\alpha(t)$. This holds also in the limit, at the end of the type 1 move, before any type 2 moves are performed. If type 2 moves are now to be done, we may assume they each replace an arc in the given leaf of $L_\alpha(t)$ with another arc in the same leaf. For otherwise the move would be a type 2(b) move, with the new arc lying in an adjacent leaf across a complementary component of $L_\alpha(t)$. Then (compare the proof of Proposition 2.1) the move could be realized by some type 2(a) moves eliminating all intersections of the interior of the arc with Σ, followed by a homotopy across a complementary component of $N_{\alpha(t)}$. Since B is maximal, the existence of the latter homotopy would mean that the two endpoint intersections of the arc with Σ could also be eliminated by a type 2(a) move. Thus type 2(b) moves are not really necessary at any time following type 1 moves.

Suppose that for t in a certain stratum we have performed on $s(t)$ some sequence of type 1 and 2(a) moves ending with a type 1 move, and that we also have another such sequence of moves arising as a limit from a sequence of moves for nearby t's in an abutting stratum. Suppose also that these two sequences produce horizontal segments with the same endpoints. These two horizontal segments can then be viewed as lying in the same leaf of $L_\alpha(t)$. The existence of essential contact arcs on either of these horizontal segments depends on whether or not the segment can be homotoped (staying in the leaf, since the leaf is incompressible) to be disjoint from those components of Σ in the leaf which lie on the side of the leaf to impede further vertical motion of the segment. Hence the two segments have the same status in this regard. Thus the essential contact components of Σ are well-defined locally. Global well-definedness then follows by considering a line segment in the parameter domain joining two points where a component is essential.

The displacement of the taut $s(t)$ is the infimum of the displacements, with respect to a suitable reference point, of:

(a) the opposite endpoints of the two vertical segments adjacent to $s(t)$, and
(b) any essential contact components of Σ which $s(t)$ meets.

The endpoints in (a) have displacements which are finitely piecewise $\mathbb{N}$-linear by induction. By Lemma 2.6 below, the infimum is effectively over a finite set of $\mathbb{N}$-linear functions. So the infimum is finitely piecewise $\mathbb{N}$-linear. This completes the induction step except for the convexity property of lengths of vertical segments. For this, the inductive assumption is that upper endpoints of vertical segments have convex displacements from reference points below, as do lower endpoints from reference points above. This is clearly preserved when $s(t)$ is pulled taut (an infimum of concave functions is again concave). □

Lemma 2.6. *The infimum of a collection of* $\mathbb{N}$*-linear functions on* $[0,\infty)^n$ *is always achieved with some finite subcollection.*

Proof: First, the inhomogeneous case reduces to the homogeneous case by adjoining a new variable $x_{n+1} = 1$. In the homogeneous case, choose one function φ in the collection. We need only consider functions taking on a smaller value than φ on one of the standard basis vectors for $[0,\infty)^n$, finitely many possibilities. For each of these possible values, the problem becomes one in $n-1$ variables, so by induction on n we are through. $\square$

3. Injectivity of Length Functions

The various length functions $\ell_\gamma : ML(M) \to [0,\infty)$, as γ ranges over the homotopy classes of loops in M (i.e., conjugacy classes in $\pi_1 M$), form the coordinates of a map $\ell : ML(M) \to [0,\infty)^\infty$. The goal of this section is to prove:

Theorem 3.1. *The function $\ell : ML(M) \to [0,\infty)^\infty$ is injective.*

It is not too difficult to show that the restriction of ℓ to the rational points of $ML(M)$ is injective. For this, it suffices to find, for a given pair of non-isotopic incompressible 2-sided surfaces S and S', a loop γ which is disjoint from S but which cannot be homotoped to be disjoint from S'. Finding such a loop is not hard after one first isotopes S to intersect S' in a minimal number of circles. The proof for irrational points of $ML(M)$ is quite a bit more difficult. For a start, it is not evident that one can isotope two incompressible measured laminations to intersect in a "minimal" fashion. Instead, it seems better to start by putting the laminations in "normal form" with respect to a triangulation of M. This takes some work, but it is easier than trying to put non-measured laminations in normal form; see [B]. Then the strategy is to find a loop having a much smaller length with respect to one lamination than the other. This is the best one could hope for, as one can see from the case of linear foliations of the 3-torus.

4. Piecewise-Linear Strata

At the end of §1 we associated to an incompressible branched surface B in M a map $\varphi_B : c(B) \to ML(M)$ where $c(B)$ is an open cone of equivalence classes of positive weights on B. One can think of φ_B as something like a coordinate chart for $ML(M)$. The following result bolsters this viewpoint.

Proposition 4.1. *If $B \subset M$ is a maximal incompressible branched surface, then the map $\varphi_B : c(B) \to ML(M)$ is injective. If $B' \subset M$ is another maximal incompressible branched surface, then the "coordinate change" map $\varphi_B^{-1} \circ \varphi_{B'}$ is a piecewise linear homeomorphism defined on an open subset (perhaps empty) of $c(B')$.*

It follows that $ML(M)$ inherits the structure of a piecewise linear manifold with components of various dimensions. These components we think of as strata, and the main problem, addressed in the remaining sections of the paper, is to see how the strata fit

together to give $ML(M)$ the structure of a stratified polyhedron in a natural way. We will also show the strata are Hausdorff manifolds, which does not follow from Proposition 4.1.
Proof: The main point will be to examine how different maximal B's carrying a given $L_\alpha \in ML(M)$ are related. There are three choices made in passing from L_α to B: the maximal essential I-bundle W_α, which determines the prelamination P_α; the transverse fiber structure on P_α; and the slits in P_α, which determine the fibered neighborhood $N(B)$.

First we consider the choice of slits. Given two collections of slits in P_α we can find a third collection disjoint from both collections, just by a local construction of inserting a multitude of small slits. So rechoosing slits can be achieved by a sequence of operations in which a single disjoint slit is inserted or deleted. If B' is obtained from B by inserting a slit $D \subset N(B)$, then pinching B' back to B induces a linear map $C(B') \to C(B)$ which passes down to a quotient map $c(B') \to c(B)$ which is a linear injection. The image of this map contains a neighborhood of α since if we consider P_β for β near α, the slit D which lies in a leaf of P_α lies near a disk D_β contained in a leaf of P_β and vertically (fiberwise) isotopic to D (since D is simply-connected), so realizing this vertical isotopy by an ambient isotopy takes P_β to a prelamination carried by B'. Thus the φ_B-image of the linear germ of α in $c(B)$ is independent of the choice of slits.

Next we consider the choice of transverse fiber structure on P_α. Any two such are isotopic, since they correspond to line fields on P_α transverse to leaves and tangent to ∂M, and the space of such line fields is contractible. So consider an isotopy of transverse fiber structures τ_t. For fixed $t = t_0$ choose a system of slits in leaves of P_α cutting all fibers of τ_{t_0} into compact intervals. These slits do the same thing for fibers of nearby τ_t as well. By compactness of the t-interval, we may then assume a single set of slits works for all t, producing branched surfaces B_t with fibered neighborhoods N_t.

Consider the collection C_t of annuli and rectangles in N_t formed by the fibers through the singular points of N_t. We may assume the components of C_t have general position intersections with each other for all t, namely isolated fibers in $int(M)$ where exactly two components meet transversely, except for three kinds of isolated "catastrophes:" a fiber of non-transverse intersection, occurring when a pair of nearby transverse intersection fibers is introduced or cancelled; a triple transverse intersection fiber, when one component of C_t moves across an intersection fiber of two other components; and a transverse intersection fiber in ∂M, when a transverse intersection fiber "moves across" ∂M. Away from these catastrophes B_t varies only by isotopy. Near one of these catastrophes we can insert a slit cutting the critical fiber in a leaf of P_α between two of the cusp points of N_t in the components of C_t involved. The non-isotopic change in B_t produced by the catastrophe can then be achieved just as well by slitting and regluing in a different way. This reduces us to a situation considered previously.

Now we consider the choice of the maximal essential I-bundle W_α. Varying this by isotopy only changes P_α by isotopy, so it is non-isotopic variation which we must consider. We consider four kinds of *elementary moves* which change one maximal essential I-bundle to another. In the first move we take the product of the disk shown in (a) of the figure below with S^1, obtaining a submanifold $S^1 \times D^2$ of a complementary component of L_α.

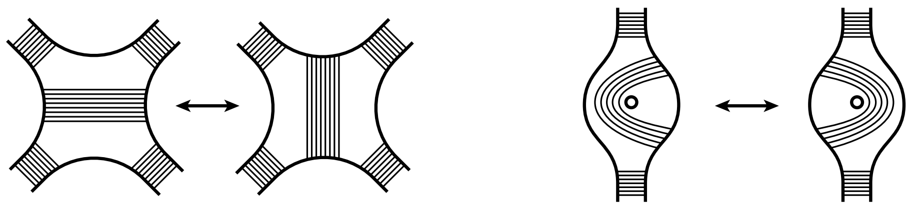

(a) (b)

The fibers of W_α are drawn as families of parallel lines. Thus in this move a component of W_α which is a thickened annulus is deleted and replaced by another such component "in the transverse direction." The second move is similar: Take the product of the disk shown in (a) with I instead of S^1. For the third move, indicated in (b), take the product of the disk shown with S^1, then delete a smaller concentric solid torus, indicated by the small central disk in (b), producing a torus component of ∂M. For the fourth move, reinsert this smaller deleted solid torus so as to produce a Seifert fibering of a solid torus with one multiple fiber, and then the picture in (b) represents the base surface of this Seifert fibering, with the small central disk a neighborhood of the multiple fiber. Both the third and fourth moves replace a component of W_α which is a thickened annulus by another such component.

Lemma 4.2. *Any two maximal essential I-bundles W_α are related by a finite sequence of elementary moves (and isotopies).*

Proof: Consider finite collections $\mathcal{A}$ of disjoint annuli and rectangles in $M - L_\alpha$, the annuli having boundary in L_α, the edges of the rectangles lying alternately in L_α and ∂M. We assume these annuli and rectangles are essential, as defined earlier. We shall use the following fact: Up to isotopy there is a unique minimal collection $\mathcal{A}$ splitting $M - L_\alpha$ into components which are either:

— I-bundles meeting $\mathcal{A}$ and ∂M in unions of fibers, and L_α in the ∂I-subbundle
— I-bundles meeting $\mathcal{A}$ and L_α in unions of fibers, and ∂M in the ∂I-subbundle
— Seifert fiberings meeting $\mathcal{A}$, L_α, and ∂M in unions of fibers
— acylindrical, containing no essential annuli or rectangles except parallel copies of com-

ponents of $\mathcal{A}$.

This is a consequence of the Jaco-Shalen, Johannson machinery, but can also be proved by an elementary geometric argument.

To obtain a maximal essential I-bundle W_α from such a minimal collection $\mathcal{A}$ we take the union of:

(1) thickenings of the the annuli and rectangles in $\mathcal{A}$ (For the rectangles there is no ambiguity about which direction fibers of W_α go since these fibers are to meet L_α in their endpoints.)
(2) the complementary components of $\mathcal{A}$ which are I-bundles of the first type
(3) thickenings of maximal collections of non-parallel essential rectangles which are unions of fibers in the second type of complementary I-bundle component of $\mathcal{A}$
(4) thickenings of maximal collections of non-parallel essential annuli which are unions of fibers in the Seifert fibered complementary components of $\mathcal{A}$.

It is not hard to see that, conversely, every maximal W_α arises in this way.

The non-isotopy variation in the choice of W_α comes from the choices of annuli and rectangles in (3) and (4). These project to arcs in the base surface S of the I-bundle or Seifert fibering. We may take S to be compact since the ends of noncompact leaves of L_α lie in the boundaries of complementary regions of type (2). Over certain arcs and circles in ∂S lie the points in L_α in the closure of the complementary region. Collapsing these arcs and circles in ∂S to points yields a new surface T with $V \subset T$ the image of the collapsed arcs and circles. In the Seifert fibered case delete from T small open disk neighborhoods of the images of the multiple fibers. Let E be the projection to T of the annuli and rectangles in (3) and (4). These are "edges" with endpoints on the "vertex" set V. By maximality of E, the complementary components of E in T are either triangles or once-punctured monogons, the punctures being components of ∂T disjoint from V; such punctures occur only in the Seifert fibered case, coming either from tori in ∂M or tori bounding neighborhoods of the multiple fibers. We call such an edge system E in (T,V) a *triangulation* of (T,V).

The four elementary moves changing W_α correspond to the two types of elementary moves on triangulations of (T,V) shown in the next figure. It is a classical fact, at least

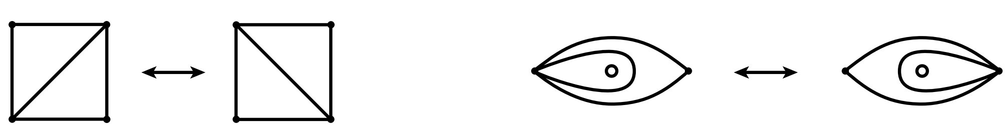

in the case with no punctures, that any two triangulations are related by a finite sequence of elementary moves. A proof of this in the general case can be found in [H3]. □

Continuing with the proof of 5.1, the branched surface versions of the four elementary moves on maximal I-bundles W_α are shown in (a) and (b) of the next figure, with the same pictorial conventions as before. Consider a move of the first or second type, changing

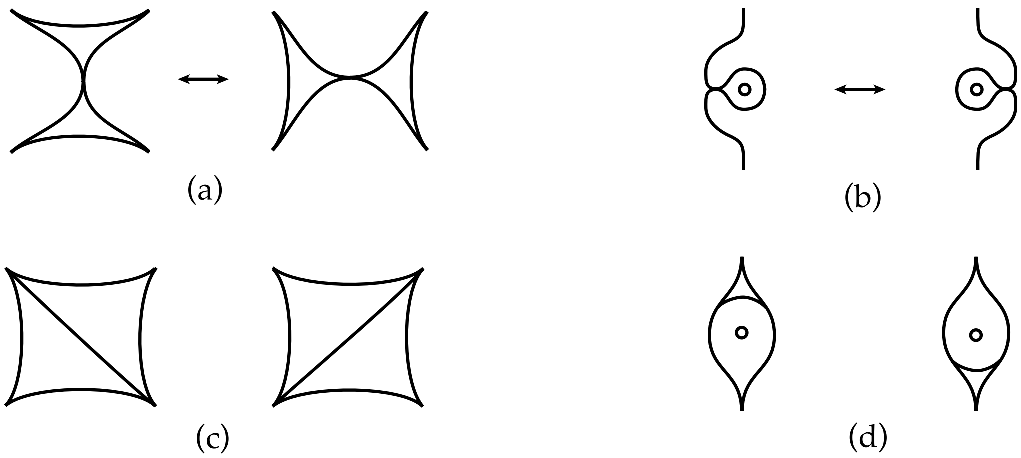

B to B'. Laminations L_β with β near the given α in $c(B)$ are carried by one of the two branched surfaces B_+, B_- indicated in (c) of the figure. To see this, one can insert slits in (a) of the first figure in this section, parallel to each of the four sheets of L_α shown, just outside the region shown. Laminations near L_α can be made disjoint from such slits, as we have seen, so these laminations can use the central component of W_α only to go diagonally from one corner of the figure to the opposite corner. Thus these laminations near L_α are carried by B_+ or B_-.

Define a function d from a neighborhood of α in $C(B)$ to $\mathbb{R}$ measuring the thickness of the layer of diagonal leaves, with a plus sign for leaves going from lower left to upper right and a minus sign for leaves going from upper left to lower right. The function d is linear, as can be seen as follows. In terms of B_+ and B_- (which pinch to B) d is just the weight on the central sector of B_+ or B_-, with the appropriate sign. For the common branched subsurface B_0 of B_+ and B_- where the central sector is deleted, d is zero. So d lifts to a continuous function $c(B_+) \cup c(B_0) \cup c(B_-) \to \mathbb{R}$, where we put the natural topology on $c(B_+) \cup c(B_0) \cup c(B_-)$, identifying $c(B_0)$ as a face of $c(B_\pm)$ (if it is not equal to $c(B_\pm)$). The projection $c(B_+) \cup c(B_0) \cup c(B_-) \to c(B)$ being continuous (linear on each piece, in fact), d is continuous. Therefore it is linear if it is linear on rational points. Since it preserves scalar multiplication, it is linear on rational points if it is linear on integer points. But linearity on integer points is clear if we interpret integer points as surfaces embedded in $N(B)$ transverse to fibers, with "sum" as the classical cut-and-paste operation (staying transverse to fibers).

Since d is linear (near α) it either takes on both positive and negative values or is identically zero. In the former case the projection $c(B_+) \cup c(B_0) \cup c(B_-) \to c(B)$ is linear and injective on each piece, hence a piecewise linear homeomorphism near L_α. (Points in different pieces cannot have the same image since they have distinct d-values.) This holds equally well with B' in place of B, with the same B_+, B_-, and B_0, so the elementary move replacing B by B' corresponds to a piecewise linear change of coordinates near α. In the case that d is identically zero, $c(B_0) \to c(B)$ and $c(B_0) \to c(B')$ are linear homeomorphisms, and the coordinate change is linear.

Elementary moves of the other two types are treated in entirely similar fashion; details are left to the reader.

What we have shown so far is the following. If L_α is isotopic to L'_α for $\alpha \in c(B)$ and $\alpha' \in c(B')$, B and B' being maximal incompressible branched surfaces, then there is a piecewise linear homeomorphism ψ of a neighborhood of α onto a neighborhood of α' such that L_β is isotopic to $L_{\psi(\beta)}$ for β near α. To finish the proof it therefore suffices to show that each φ_B is injective.

Suppose φ_B is not injective, so that there exist isotopic laminations L_α and L_β with distinct $\alpha, \beta \in c(B)$. By the preceding, whole neighborhoods of α and β would then correspond to isotopic laminations, so we may take α and β to be rational, and in fact integral by clearing denominators. Let S_α and S_β be the corresponding incompressible surfaces carried by B. Doubling α and β if necessary, we may assume S_α and S_β are orientable. Now we appeal to a result of Oertel [O2] which says that since S_α and S_β are isotopic in M, we can pass from S_α to S_β by a sequence of moves in which sheets of S_α are transferred across slits of B (and vertical isotopy in $N(B)$). Hence $\alpha = \beta$ in $c(B)$, a contradiction. ☐

5. Limits of Length Functions

Given an incompressible branched surface B, let $L_t \in ML(M)$ be a path of laminations determined by a linear path of weights $\alpha(t)$ on B, $t \geq 0$, which are all strictly positive for $t > 0$ but not for $t = 0$. By Proposition 2.5 we know that for any loop γ, $\ell_\gamma(L_t)$ approaches a limiting value as t goes to 0. The question we consider in this section is, when is this limit equal to $\ell_\gamma(L_0)$ for all γ?

Zero Weights

For a branched surface B, the branched subsurfaces $B' \subset B$ correspond to the faces $C(B')$ of $C(B)$ obtained by setting the weights α_i for the sectors in $B - B'$ equal to zero. Assuming B is incompressible, it can happen that B' is compressible, due to the presence of disks or half-disks of contact for B'. Let B'' be obtained from B' by splitting to eliminate disks and half-disks of contact; that is, we split $N(B')$ along a complete collection of disjoint disks and half-disks of contact to form $N(B'')$ for a branched surface B''. Any incompressible L_α with $\alpha \in C(B')$ can be isotoped, by vertical motion in $N(B')$ and shifting sheets across slits, so as to be disjoint from the chosen disks and half-disks. So the natural linear map $C(B'') \to C(B') \to c(B')$ is onto the subset of incompressible L_α's, which therefore form a nice subcone of $c(B')$. Assuming this subcone is non-empty (hence contains rational points), it follows as in [O2] that B'' is incompressible. [Details?] We call B'' an incompressible branched surface for the given face of $C(B)$.

Recall from [O2] that an incompressible branched surface B is said to have a Reeb component if it carries a torus or annulus bounding a solid torus which meets some surface carried by B with positive weights in a collection of meridian disks. In particular, B is without Reeb components if it carries no compressible tori or ∂-compressible annuli. By [O2], if B has no Reeb components then every surface carried by B (without restriction on weights) is incompressible and ∂-incompressible. A similar question, not considered in [O2], is whether B can carry closed surfaces parallel to ∂M. To handle this, we enlarge the notion of Reeb component to include the case that B carries a ∂-parallel torus T cutting off a $T \times I$ from M which meets some surface carried by B with positive weights in a collection of essential annuli. By methods as in [O2] one can easily show that when this broader class of Reeb components is excluded, B carries no ∂-parallel surfaces. Following [G-O] we call an incompressible branched surface without Reeb components *essential*.

Let $\mathcal{B}$ be the collection of essential branched surfaces $B \subset M$ for which $C(B)$ is non-empty. If $B \in \mathcal{B}$ and B' is an incompressible branched surface for a face C' of $C(B)$, then $B' \in \mathcal{B}$. Consider the following diagram of linear maps:

$$\begin{array}{ccc} C(B') & \xrightarrow{J} & \overline{C}(B) \\ \big\downarrow q' & & \big\downarrow \overline{q} \\ c(B') & \xrightarrow{j} & \overline{c}(B) \end{array}$$

Here J comes from the definition of B' as an incompressible branched surface for the face C' (hence $Im(J) \subset C'$), q' is the canonical quotient map, q is the continuous extension of the quotient map $q\colon C(B) \to c(B)$ to the closures of these cones, and j is defined to make the diagram commute. To see that j is well-defined, consider a product complementary region P of B'. The product structure on P can be isotoped so that its I-fibers are transverse to any surface carried by B with positive weights. [[[...DETAILS...]]]

The image of j is the whole face $c' = q(C')$ of $c(B)$. To show this it suffices to show that $Im(j)$ contains all the rational points in c'. These points correspond to surfaces carried by B with weights in C'. Since $B \in \mathcal{B}$ these surfaces are incompressible and ∂-incompressible, hence as earlier they can be isotoped off the disks and half-disks of contact along which the branched subsurface of B corresponding to C' is split to form B'. So these surfaces are carried by B', with positive weights.

The maps j are not in general injective, however.

Reeb bundles

We return now to the question raised at the beginning of this section: If B is an incompressible branched surface and $L_t \in ML(M)$ is a path of laminations determined by a linear path of weights $\alpha(t)$ on B, $t \geq 0$, which are positive for $t > 0$ but not for $t = 0$, when is it true that the lengths $\ell_\gamma(L_t)$ approach $\ell_\gamma(L_0)$ for all loops γ? We are assuming L_0 is incompressible, so it is carried with positive weights by some incompressible branched surface B' for the face of $C(B)$ containing $\alpha(0)$ in its interior. As described above, B' is obtained from the branched subsurface of B corresponding to this face by splitting disks and half-disks of contact.

A *Reeb bundle* for the pair (B, B') is a codimension zero compact submanifold R of the closure of $M - N(B')$, meeting $\partial N(B')$ in a disjoint union $\partial_h R$ of tori and annuli in $\partial_h N(B')$, such that:

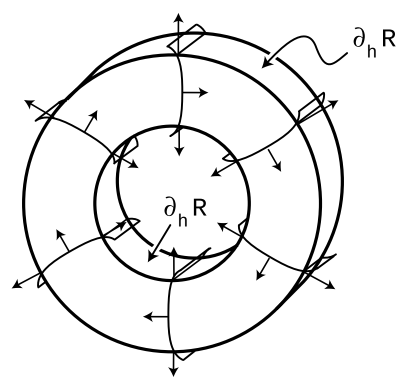

(a) For some surface S carried by B with positive weights, the leaves of the associated prelamination P meet R in the surface fibers of a fiber bundle $R \to S^1$, with $\partial_h R$ a subbundle.

(b) There is an oriented line field on R, transverse to the surface fibers, which on $\partial_h R$ is tangent to the interval fibers of $N(B')$, pointing inside $N(B')$.

The Reeb bundle R is *essential* if it contains a path with endpoints in $\partial_h R$ which cannot be deformed in M (rel endpoints) to a path in $\partial_h N(B')$. We shall see that if (B, B') has an essential Reeb bundle for one surface S carried by B with positive weights, then it has an essential Reeb bundle for each such S.

Proposition 5.1. $\lim \ell_\gamma(L_t) = \ell_\gamma(L_0)$ *for all* γ *if and only if there are no essential Reeb bundles for* (B, B').

Proof: First we show how to construct an essential Reeb bundle if $\lim \ell_\gamma(L_t) \neq \ell_\gamma(L_0)$ for some γ. For this it suffices by continuity to assume that the linear path L_t joins laminations L_0 and L_1 coming from rational, and in fact integral, weights $\alpha(0)$ and $\alpha(1)$ corresponding to surfaces S_0 and S_1.

Consider S_0 as embedded in the prelamination P_1 associated to S_1, meeting ∂P_1 only in ∂M. By vertical isotopy of S_0 we can arrange that S_0 and ∂S_0 are transverse to leaves of P_1 except for the following phenomena:

— saddle tangencies in $int(S_0)$ and half saddle tangencies on ∂S_0

— round saddles (a circle of tangency in $int(S_0)$, with a local maximum or minimum in cross-section) and half round saddles (a proper arc of tangency with a similar cross-section).

It follows that for the induced singular foliation on S_0 all the nonsingular leaves are essential, not cutting off a disk from S_0. We may assume all the various saddles of S_0 lie in distinct nonsingular leaves of P_1, and then we can choose for S_1 a union of nonsingular leaves of P_1 not containing saddles, such that in each region of P_1 between adjacent components of S_1 there is at most one saddle or singular leaf.

Let ν be the vertical linefield on P_1, modified near S_0 (staying transverse to leaves of P_1 and tangent to ∂M) to be tangent to S_0 except near the various saddles of S_0, where ν is given a standard form.

Let T_0 be S_1 split along S_0. For a fixed m (to be chosen later) let T_m be the subsurface of T_0 consisting of points from which the two trajectories of ν, above and below, flow across at least m sheets of T_0 without meeting $\partial_h P_1$ or saddles of S_0.

Lemma 5.2. *If $\lim \ell_\gamma(L_t) \neq \ell_\gamma(L_0)$ for some γ, then ν can be deformed so that there exists a component U_0 of T_m which is essential, containing a path with endpoints on S_0 which cannot be homotoped in M (rel endpoints) to a path in S_0.*

Proof: First we define a collection C_0 of curves in T_0. Each cusp circle of P_1 contributes two circles to C_0, obtained by flowing up and down via ν to the adjacent sheet of T_0. Similarly each cusp arc of P_1 contributes two arcs to C_0. The rest of C_0 comes from the various saddles. An ordinary saddle of S_0 contributes two core arcs to C_0 with boundary on S_0, one in the sheet of T_0 above the saddle, the other in the sheet below; these can be thought of as the core and co-core of the associated 1-handle. A half saddle yields only one core arc in C_0, with one end on S_0 and the other on ∂M. A round saddle contributes a circle to C_0 obtained from the core of the round saddle by flowing to the adjacent sheet of T_0 on the side away from S_0. Similarly a half round saddle gives an arc in C_0 with endpoints on ∂M. Each curve of C_0 is embedded and essential in T_0.

A curve of C_0 arises from trajectories of ν flowing into T_0 from one side or the other. Continue these trajectories until they meet either $\partial_h P_1$, a saddle (of some sort) in S_0, or the next sheet of T_0. Let C_1 be the union of C_0 with the points of T_0 at the ends of these enlarged trajectories. If ν is in general position, C_1 is obtained from C_0 by adding finitely many embedded circles and arcs with endpoints either on ∂T_0 or on curves of C_0 coming from cusp points of P_1 or round or half round saddles. Inductively, we construct curve collections C_n in this way, $C_n - C_{n-1}$ consisting of curves with boundary on $\partial T_0 \cup C_{n-1}$.

We wish now to modify ν so that the curves of C_n for a fixed $n \geq 2m$ have minimal intersections with each other, in the following sense: If c and c' are two curves of C_n then there is no disk $D \subset T_0$ with ∂D consisting of an arc of c, an arc of c', and, possibly, an arc of S_0. The procedure is inductive. Consider adjoining one curve c of $C_n - C_{n-1}$, assuming all previous curves, including those in C_{n-1}, have minimal intersections with each other. Let $c' \in C_i$ be a previous curve which c intersects non-minimally, with i minimal, and let D be a disk as above for c and c'. We may assume D is a minimal such disk. The curves c and c' arise from trajectories which arrive at c and c' from opposite sides of T_0, say from above at c and from below at c'. (Otherwise we would have a non-minimal intersection among curves of C_{n-1}.) In a neighborhood of D there can be no other points of c or previous curves coming from above, nor can there be other points of previous curves in C_i coming from below. So the obvious isotopy of c across D decreases the number of intersections of c with previous curves in C_i coming from below. This isotopy can be realized by an isotopy of ν in the region just above a neighborhood of D, finishing the induction step in our modification of ν.

To see how the hypothesis $\lim \ell_\gamma(L_t) \neq \ell_\gamma(L_0)$ implies the existence of a component U_0 of T_m containing an essential path, we consider as in the proof of Proposition 3.2 the prelaminations P_t corresponding to L_t (with maximal I-bundles W_t collapsed) and the family of loops $\gamma(t)$ which are pulled to loops $\widehat{\gamma}(t)$ taut with respect to P_t for $t > 0$. (Note the change in notation: now $\alpha(t) = t\alpha(1) + (1-t)\alpha(0)$ for $t \in [0,1]$, whereas formerly $\alpha(t) = \alpha(0) + t\alpha(1)$ for $t \in [0,\infty)$; projectively this is the same path, but now it has the limiting value $\alpha(0)$.)

Prior to pulling $\gamma(t)$ taut, we may first deform $\gamma(1)$ to minimize the number of intersections with S_0, cancelling pairs of unnecessary intersections in the usual way. Then we pull $\gamma(1)$ taut and construct $\gamma(t)$, with vertical segments created by the intersections of $\gamma(1)$ with S_0. Some of these vertical segments may be eliminated when $\gamma(t)$ is pulled taut. As t goes to 0, if every intersection point of the taut $\widehat{\gamma}(t)$ with S_0 lies in a vertical segment of $\widehat{\gamma}(t)$, then clearly $\ell_\gamma(L_t)$ approaches $\ell_\gamma(L_0)$, the number of points of intersection of $\gamma(t)$ with S_0. And otherwise the limiting value of $\ell_\gamma(L_t)$ is less than $\ell_\gamma(L_0)$. The latter case is the one we are considering.

Since some vertical segment is eliminated completely for t approaching 0, there must be a pair of such segments at the ends of a horizontal segment s of $\gamma(t)$ upon which an unbounded number of type 1 and 2 moves is performed, s meeting S_0 only in its endpoints.

(Since we earlier deformed $\gamma(1)$ to minimize intersections with S_0, s is essential with respect to S_0.) As shown in Figure 5.2, we see that for t near 0, s is being deformed via a sequence of type 1 and 2 moves through a one-parameter family s_s of horizontal segments with endpoints on S_0 and crossing S_1 arbitrarily often, say at the parameter values $s = 0, 1, \cdots, n$, where n is chosen as earlier in the proof.

Figure 5.2

For each s the path s_s may be viewed as a horizontal path in P_1, and we may perturb the homotopy s_s so that it is transverse to the leaves of P_1.

The claim is now that we can deform the homotopy s_s for $s \in [0, n]$ so that it is obtained by simply flowing along trajectories of ν, meeting components T_{0i} of T_0 for integer parameter values $s = i$. Suppose inductively that this has been done for $s \in [0, j]$. This means that s_j is disjoint from the collection C of curves of Cj lying in T_{0j} and arising from trajectories flowing to T_{0j} from the same side that s_j came from. To do the induction step we need to homotope s_j in T_{0j}, staying disjoint from C, so that it also misses the collection C' of curves of C_0 in T_{0j} flowing to T_{0j} from the opposite side. (The curves of C' come from cusp curves and cores of saddles obstructing the flow of s_j along trajectories of ν from T_{0j} to $T_{0,j+1}$.)

Suppose s_j meets C'. The existence of s_s for $s \in [j, j+1]$ implies that s_j can be homotoped in T_{0j} to be disjoint from C', so there is a subsegment s' of s_j, meeting C' only in one or both of its endpoints, which can be homotoped (rel its intersection with C') to a path s'' in C'. See Figure 5.3 for the case that both endpoints of s' lie on C'; we describe this case and leave it for the reader to make the minor modifications needed for the other case.

Figure 5.3

Let T' be the subsurface of T_{0j} consisting of a neighborhood of $C \cup C'$ with complementary disk components meeting S_0 in at most an arc filled in. One possibility is that s' lies in such a disk. In this simple case we may assume that s' is embedded, hence also that s'' is embedded. So $s' \cup s''$ bounds an disk $D \subset T_{0j}$, disjoint from C since we aranged that curves of C_n intersected minimally. So homotoping s' across D decreases intersections of s_j with C' without introducing intersections with C.

The other possibility is that s' lies outside T' except near its endpoints. Since s' can be homotoped into T', the part of s' outside T' can be homotoped to a path in $\partial T'$, without affecting intersections with $C \cup C'$; we assume now that this has been done. The part of s' in $\partial T'$ we may take to be a monotone (i.e., immersed) path, which might conceivably wrap around this component of $\partial T'$ more than once. If it does not, then s' is embedded and the previous argument applies. If s' does wrap around $\partial T'$ more than once, then by an easy argument in surface theory, s' must lie in an annulus bounded by the components of $\partial T'$ and C' which it meets. Then C must be disjoint from annulus since curves of C_n intersect minimally, so again we can deform s' to decrease intersections of s_j with C' without introducing intersections with C. This finishes the induction step in the improvement of the homotopy s_s to be a flow along ν.

We have chosen $n \geq 2m$, and therefore the essential path s_m can flow along ν crossing m sheets in both directions. Let U_0 be the component of T_m containing s_m. □

Returning to the proof of Proposition 5.1, let U_j for $j \leq m$ be the set of points in T_0 reachable from U_0 by flowing along trajectories of ν, crossing T_0 at most j times. Thus $U_0 \subset U_1 \subset \cdots \subset U_m$. Each component of U_j is essential since U_0 is. In particular, each component of U_j contains a path which cannot be homotoped in T_0 (rel endpoints) to a path in S_0. Let U_j' be obtained from U_j by filling in any disk components of $T_0 - U_j$ which are either disjoint from ∂T_0 or meet ∂T_0 in an arc in S_0. These surfaces also form an expanding sequence $U_0' \subset U_1' \subset \cdots \subset U_m'$.

If m is larger than a certain bound depending only on T_0 then for some $k \leq m$ either:

(1) Two components of U_k' are parallel rectangles in T_0 meeting S_0 in a pair of opposite sides (parallel via an isotopy through such rectangles).

(2) The surfaces U_{k-1}' and U_k' are isotopic via an expanding one-parameter family of surfaces $U_t' \subset T_0$, with $U_t' \cap S_0$ also an isotopy.

In case 1 we can take an essential arc in one rectangle and flow along trajectories until we obtain an essential arc in the other rectangle. From this we can then construct a bundle $V \to S^1$ whose fibers are rectangles in leaves of P_1, rectangles meeting S_0 in a pair of opposite edges and containing an essential arc.

In case 2 we also construct a bundle $V \to S^1$, in several steps. First let V_1 be the union of the trajectories of ν starting from U_0 (in both directions) and continuing until they cross T_0 k times, stopping a short distance ε after the *kth* crossing. Next fill in the disk components of $int(T_0) - U_k$ and enlarge V_1 to V_2 by letting these disks flow along trajectories as long as their boundaries stay in V_1. Enlarge V_2 to V_3 by adding the trajectories going a distance ε up and down from the disk components of $T_0 - U_k$ which meet S_0 in single arcs. (Unfortunately these disks must be treated differently from the disks of $int(T_0) - U_k$ due to the presence of saddles of S_0.) Now form V_4 by truncating ends of trajectories in V_3, tapering via the isotopy U_t' in the interval $[0, \varepsilon/2]$ on the appropriate side of U_k'. Similarly, form V_5 by tapering V_4 in the interval $[\varepsilon/2, \varepsilon]$ on each side of U_k' so as to delete by isotopy the disk components of $T_0 - U_k$ meeting S_0 in single arcs. See Figure 5.4 for what V_5 might look like along S_0.

Figure 5.4

V_5 is foliated by compact subsurfaces of leaves of P_1 which are transverse to ∂V_5, and is therefore a bundle $V_5 \to S^1$.

To get the desired $V \to S^1$ it remains to peel part of ∂V_5 away from S_0 so that V intersects S_0 in a subbundle, while preserving the property that fibers of $V \to S^1$ contain essential paths. To do this, take an essential path in $U'k$ and let it flow through the fibers of the bundle V. We may do this arbitrarily far keeping both endpoints of the path on S_0. One endpoint must eventually return to the same component of $U'k \cap S_0$. Taking the first return, we may assume this endpoint then follows periodically an embedded loop in S_0. And similarly for the other endpoint. If the two loops followed by the endpoints meet, we may assume they coincide. Peel all of ∂V_5 away from S_0 except for a neighborhood of these loops.

Now we show the converse. For this we may assume L_1 is carried by B with integral weights, corresponding to the surface S_1, while L_0 is carried with arbitrary positive weights by B'. Assume we have an essential Reeb bundle R. By definition, we then have a path γ' in R which cannot be deformed into $\partial_h N(B')$. We may deform this to lie in a fiber $S_1 \cap R$ of R. Extend γ' at its ends by vertical arcs in P_0 of length ε, say. This extended γ'' is a shortest path between its endpoints, length being measured with respect to P_0. This is because γ'' is taut, or can be made into a taut γ''' of the same length by deforming arcs of γ' outside $N(B')$ into $\partial N(B')$ if possible and then eliminating extra intersections of γ' with Σ by type 2 moves; since γ' cannot be deformed into $\partial_h N(B')$, the resulting γ''' will be taut. By the fact mentioned at the end of §2, γ'' can be extended to a PVH loop γ_0 minimizing length for P_0. We can extend this to a continuous family of loops γ_t for t near 0 which are PVH for P_t, with the part of γ_t corresponding to γ'' consisting of two vertical segments of length approximately ε joined by a horizontal segment. Now using the Reeb bundle structure on R this part of γ_t can be pushed vertically to shorten it by 2ε, provided $t > 0$. Hence $\lim \ell_\gamma(L_t) \neq \ell_\gamma(L_0)$. □

Proposition 5.3. *If laminations* $L_t \in ML(M)$, $t \in [0,1]$, *arise from a linear path of weights* $\alpha(t) \in \overline{C}(B)$ *for some incompressible branched surface* B, *with* $\alpha(t) \in C(B)$ *for* $t > 0$, *and if* $\lim \ell(L_t) \neq \ell(L_0)$, *then* $\lim \ell(L_t) \notin \ell(ML(M))$.

6. Assembling the Strata

In §4 we saw how the set $ML(M)$ is partitioned into disjoint strata which have the structure of piecewise linear manifolds (without boundary). Our task now is to fit these strata together in a way which reflects how the associated length functions converge.

Recall from §5 the collection $\mathcal{B}$ of essential branched surfaces. If $B \in \mathcal{B}$ and C' is an open face of the cone $C(B)$, let B' be an incompressible branched surface for this face (obtained as in §1 by slitting disks and half-disks of contact in the branched subsurface of B corresponding to C'). We call the corresponding face c' of $c(B)$ *good* if the pair (B, B') has no essential Reeb bundles. From Proposition 5.1 this is equivalent to continuity of all length functions ℓ_γ at the face C', and is independent of which face C' we choose projecting to c' by Proposition 4.?. From Proposition 4.1 it follows that the natural linear map $j: c(B') \to c'$, which is surjective since B and B' are without Reeb components, is also injective. (In these applications of Propositions 4.1 and 4.? it suffices to have the result just for rational points.) Let $c(B, B') = c(B) \cup c(B')$ where $c(B')$ is identified with the face c' of $c(B)$ via j. The natural map $\varphi: c(B, B') \to ML(M)$ is injective. Namely, it is injective on $c(B)$ and $c(B')$ separately by Proposition 5.1, since we can pinch B and B' to be maximal, inducing linear inclusions of the cones $c(-)$. Also by 5.1, if the φ-images of $c(B)$ and $c(B')$ intersected each other, they would do so at rational points. Injectivity of φ on all rational points of $c(B, B')$ was proved in [O2].

We define a topology on the set $ML(M)$ by specifying that $U \subset ML(M)$ is open iff $\varphi^{-1}(U)$ is open in $c(B, B')$ for all $B \in \mathcal{B}$ and all good faces c' of $c(B)$.

On each stratum of $ML(M)$ this topology is the same as that underlying the piecewise linear manifold structure on that stratum.

Proposition 6.1. *The map $\ell: ML(M) \to [0, \infty)^\infty$ having coordinates the length functions ℓ_γ is a homeomorphism onto its image.*

Proof: We have shown that ℓ is a continuous injection, so it remains to show that a sequence $L_n \in ML(M)$ converges to $L \in ML(M)$ whenever all the length functions $\ell_\gamma(L_n)$ converge to $\ell_\gamma(L)$. All the laminations L_n are carried with positive weights by a finite number of branched surfaces $B \in \mathcal{B}$, so by passing to a subsequence we may assume all the L_n's are carried with positive weights by one B, say L_n is carried with weight vector $\alpha_n \in C(B)$.

If the α_n's are unbounded, then after passing to a further subsequence we may choose positive scalars $t_n \to 0$ so that the sequence $t_n \alpha_n$ approaches a non-zero limit $\alpha \in \overline{C}(B)$. Since $\lim \ell(t_n L_n) = 0 \in \ell(ML(M))$, Proposition 4.? implies that $\lim \ell(t_n L_n) = \ell(L_\alpha)$. But $\ell(L_\alpha) \neq 0$ since ℓ is injective and $\alpha \neq 0$. This contradiction shows the α_n's are bounded.

Passing to a subsequence, we may then assume the α_n's converge to $\alpha \in \overline{C}(B)$. As in the preceding paragraph, $\lim \ell(L_n) \in \ell(ML(M))$ implies $\lim \ell(L_n) = \ell(L_\alpha)$, so $L_n \to L_\alpha$ in $ML(M)$. And $L_\alpha = L$ since ℓ is injective. ☐

Proposition 6.2. *The frontier of each stratum of $ML(M)$ is a union of strata of lower dimension.*

Proof: Let L_0 be a lamination in the frontier of a given stratum. This means there is a pair (B, B') where $B \in \mathcal{B}$ has $\varphi(c(B))$ contained in the given stratum and B' is an incompressible branched surface for a good face of $c(B)$, B' carrying L_0 with positive weights. Thus there is linear path L_t, $0 \leq t \leq 1$, approaching L_0, with $L_t \in \varphi(c(B))$ for $t > 0$. Our first goal is to modify B and B' so that B' becomes maximal, still carrying L_0 with positive weights, with the new B still carrying the linear path L_t at least for t near 0.

We take L_1 to have integer weights, corresponding to a surface $S \in \mathcal{S}(M)$. The prelamination P_0 determining L_0 lies in $N(B)$ transverse to I-fibers, as does S. Perturb S to be in general position with respect to $\partial_h P_0$. Split $N(B)$ outside $P_0 \cup S$ along surfaces transverse to fibers of $N(B)$ so that $N(B)$ becomes $P_0 \cup N(S)$. Inserting slits in P_0 and collapsing the resulting I-fibers in $P_0 \cup N(S)$ yields a new branched surface B. The branched surface B' we take to be obtained from P_0 by inserting slits and collapsing resulting I-fibers. So B is obtained from B' by adjoining $S - P_0$.

The incompressible branched surface B' can be made maximal by pinching P_0 near a collection A_i of complementary annuli, rectangles, Mobius bands, digons, and half-digons. We need to put these A_i's into a good position with respect to S. To simplify the discussion let us assume no A_i's are rectangles or half-digons; the additional arguments needed to handle such A_i's are entirely similar, and no more difficult. We may assume each A_i meets S transversely, with the curves of $S \cap A_i$ meeting ∂A_i tangentially since we may assume S meets $\partial_h P_0$ tangentially. There can be no monogon components of $A_i - S$ since the original B was incompressible and splitting cannot produce monogons. Disk components of $A_i - S$ (with smooth boundary) can be eliminated as follows. Consider the prelamination P_t associated to L_t (with respect to the branched surface B). The part of P_t outside P_0 consists of a slight thickening of $S - P_0$, foliated trivially. Let D be a disk component of $A_i - S$. We view D as a disk in the complement of P_t with $\partial D \subset \partial_h P_t$. We can begin to enlarge D by adding on a collar so that ∂D stays in a leaf of P_t at each time. Choosing a fixed $t = t_1$ small enough, this enlargement of D can be continued until ∂D meets another component of $A_i - P_t$. See Figure 6.1.

Figure 6.1

For the enlarged D, ∂D bounds a disk D' in a leaf of P_t (possibly meeting cusp points of P_t if P_t has disks of contact) since L_t is incompressible. Clearly, if $t = t_1$ is small enough, D' will meet P_0 only near $\partial_h P_0$.

Now we split P_t along the disk D' (slightly enlarged). After this splitting, we can eliminate the disk component D from $A_i - P_t$ by isotoping A_i so that D moves to D'. If $int(D')$ meets $\cup A_j$ we push disks in other A_j's along too so that $\cup A_j$ stays embedded. This splitting and isotopy decreases the number of components of $\cup A_j - P_t$. So by repeating this process finitely often (with smaller and smaller $t = t_1, t_2, \cdots$) we eventually reach the case that there are no disk components of $\cup A_j - P_t$. Each time we split P_t we also split $N(B)$, producing a new B. Any disks or half-disks of contact which this B may have can be eliminated by further splitting of the same sort. So we may assume B is incompressible.

There exist I-fiberings of the A_i's which restrict to I-fiberings of the components of $A_i - P_t$. For clearly such I-fiberings exist except for A_i's which are annuli or Mobius bands meeting P_t in "Reeb configurations" (Figure 6.2). But these would give rise to essential Reeb bundles for (B, B'), which are ruled out by Proposition 3.1.

Figure 6.2

Thus we can pinch B' to be maximal, simultaneously pinching B. The new B might have monogons, however, arising when an annulus component of $A_i - P_t$ is pinched. (Pinching a Mobius band component of $A_i - P_t$ cannot produce a monogon. For such a monogon would come from a ∂-compressing disk for this Mobius band relative to P_t; ∂-surgering the Mobius band via this ∂-compressing disk would produce a nonseparating compressing disk for $\partial_h P_t$). An annulus component of $A_i - P_t$ which pinches to create a monogon would be ∂-parallel relative to P_t, isotopic to an annulus in $\partial_h P_t$. Just as disk components of $A_i - P_t$ were eliminated before, we can eliminate this annulus component of $A_i - P_t$ by splitting P_t and isotoping A_i, decreasing the number of components of $A_i - P_t$. Repeating this process, we eventually reach a position where pinching the A_i's produces no monogons for B, so the new B is incompressible (pinching the A_i's cannot create disks or half-disks of contact). This gives the desired pair (B, B'), which still carries the linear path L_t for t small.

The branched surface B' still corresponds to a good face c' of $c(B)$ since $\lim \ell_\gamma(L_t) = \ell_\gamma(L_0)$ for all γ. This is a proper face since we assume L_0 is not in the same stratum as L_t for $t > 0$. By Proposition 3.1 we have continuity of all length funcions ℓ_γ at c', so $\varphi(c') = \varphi(c(B'))$ is in the closure of $\varphi(c(B))$ by Proposition 6.1. Since B' is maximal, this means a neighborhood of L_0 in its stratum is contained in the closure of the stratum of L_t, $t > 0$. Since c' is a proper face of $c(B)$, the stratum of L_0 has dimension less than the dimension of the stratum of L_t.

Strata being connected, it remains only to verify that the closure of the stratum of L_t meets the stratum of L_0 in a closed set. After Proposition 6.1 this is just simple point-set topology (the set of accumulation points of a set is closed). □

7. The Linear Substratification in the Atoroidal Case

In this section we assume M is atoroidal and acylindrical, i.e., the set of incompressible surfaces $\mathcal{S}(M)$ defined in §1 contains no surfaces of Euler characteristic zero. We will show that the piecewise linear strata of $ML(M)$ then have, themselves, a natural stratification into linear substrata. The key ingredient is the following:

Lemma 7.3. *If M is atoroidal and acylindrical, then a maximal incompressible branched surface $B \subset M$ has only a finite number of annuli and rectangles of contact, modulo vertical isotopy in $N(B)$.*

Proof: Suppose there are infinitely many annuli of contact, hence an infinite number all having the same boundary consisting of one or two branching circles of B. For convenience in describing this situation, consider the enlarged branched surface B' in the manifold M' obtained by deleting a thin solid torus in $M - B$ running parallel to each of the branching circles in question, as shown in cross-section in Figure 7.1.

Figure 7.1

We have an infinite sequence of annuli A_i carried by B' with integer weight vectors $\alpha_i \in \overline{C}(B')$. Since all the α_i's are distinct by assumption, some coordinate α_{j_i} of the α_i's must be unbounded. Normalizing the α_i's by dividing by the sum of their coordinates, we get an infinite sequence of weight vectors in the intersection of $\overline{C}(B')$ with the hyperplane $\sum_j \alpha_j = 1$. This intersection is compact, so some subsequence of the normalized α_i's converges to a weight $\alpha \in \overline{C}(B')$. The coordinate of α along $B' - B$ is zero since the α_i's have bounded weight (1 or 2) there. The coordinate weights of α along ∂M are also zero, since this was the case for the α_i's.

The Euler characteristic function χ on $\overline{C}(B')$ is linear with rational coefficients. The α_i's lie in the linear subspace where $\chi = 0$, hence α does also. Rational points are dense in this subspace. These correspond to collections of tori carried by B. Thus we can find a torus collection arbitrarily close (in terms of normalized weights) to an annulus A_i. Let T be a torus in such a collection.

Given a maximal incompressible branched surface B, consider P_α associated to $\alpha \in C(B)$. Let $\mathcal{A}(P_\alpha)$ be the collection of annuli and rectangles of contact for B which lie in leaves of P_α, the annuli meeting the singular locus of P_α only in their boundary circles, the rectangles only in pairs of opposite sides. For a fixed set $\mathcal{A}$ of annuli and rectangles of contact for B, let $c(\mathcal{A}) \subset c(B)$ be the set of P_α's with $\mathcal{A}(P_\alpha) = \mathcal{A}$. These subsets $c(\mathcal{A})$ partition $c(B)$ into disjoint subsets, finite in number by Lemma 7.1.

Each $c(\mathcal{A})$ is defined by a finite number of linear equations and strict inequalities in $c(B)$. This can be seen as follows. The condition that a given annulus or rectangle of contact A have its two opposite boundary curves on the same height in P_α is given by a linear equation on α, since the difference in heights between these two curves is a linear function of α. If A is in $\mathcal{A}(P_\alpha)$ for some α then α must satisfy this equation, but also α must satisfy some strict linear inequalities expressing the condition that no other cusp curves of P_α move up or down into A (regarding A as contained in a leaf of P_α). When such a cusp curve does move into A it means that for this P_β a subannulus or subrectangle of A lies in $\mathcal{A}(P_\beta)$, rather than A itself. This yields the additional fact that the closure of each $c(\mathcal{A})$ is a union of $c(\mathcal{A})$'s.

Thus the $c(\mathcal{A})$'s define a stratification of $c(B)$ into convex polyhedral substrata. The coordinate change transformations $\varphi_B^{-1} \circ \varphi_{B'}$ map the substrata in $c(B')$ linearly to the substrata in $c(B)$, as the proof of Proposition 4.1 shows, so there is induced on each piecewise linear stratum of $ML(M)$ a natural subdivision into substrata with natural linear structures.

Proposition 7.2. *If $\mathcal{S}(M)$ contains no surfaces of Euler characteristic zero, then the frontier of each linear substratum of $ML(M)$ is a union of linear substrata of lower dimension.*

References

[AB] R. Alperin and H. Bass, Length functions of group actions on Λ-trees. In: Combinatorial group theory and topology, 265-378, Ann. of Math. Stud. 111, Princeton Univ. Press, 1987.

[B] M. Brittenham, Essential laminations and Haken normal form: laminations with no holonomy. Comm. Anal. Geom. 3 (1995), 465-477.

[FH] W. Floyd and A. Hatcher, The space of incompressible surfaces in a 2-bridge link complement. Trans. Amer. Math. Soc. 305 (1988), 575-599.

[FO] W. Floyd and U. Oertel, Incompressible surfaces via branched surfaces. Topology 23 (1984), 117-125.

[GO] D. Gabai and U. Oertel, Essential laminations in 3-manifolds. Ann. of Math. 130 (1989), 41-73.

[H1] A. Hatcher, Measured lamination spaces for surfaces from the topological viewpoint. Topology Appl. 30 (1988), 63-88.

[H2] A. Hatcher, Some examples of essential laminations in 3-manifolds. Ann. Inst. Fourier (Grenoble) 42 (1992), 313-325.

[H3] A. Hatcher, On triangulations of surfaces. Topology Appl. 40 (1991) 189-194.

[HO] A. Hatcher and U. Oertel, Full laminations in 3-manifolds. Math. Proc. Cambridge Philos. Soc. 119 (1996), 73-82.

[O1] U. Oertel, Measured laminations in 3-manifolds. Trans. Amer. Math. Soc. 305 (1988), 531-573.

[O2] U. Oertel, Incompressible branched surfaces. Invent. Math. 76 (1984), 385-410.