Low dimensional topology

Name : Tanushree Shah
Student ID: 20131065
Supervised by Tejas Kalelkar
Indian Institute of Science Education and Research
Department of Mathematics

April 30, 2017

Contents

1 Curve system **3**
1.1 The complex of curve systems . . . 3
1.2 Train tracks . . . 4
1.3 Standard tracks . . . 5
1.4 The polyhedra ML(M) and PL(M) . . . 6
1.5 The global structure of ML(M) and PL(M) . . . 8

2 Measured Lamination **10**
2.1 Construction of measured laminations . . . 10
2.2 Length function . . . 11
2.3 Intrinsic structure via length function . . . 12
2.4 Normal form for measured laminations . . . 12

3 References **15**

The first two sections of my report are based on a paper by A.E.Hatcher titled ‘Measured lamination spaces for surfaces, from topological viewpoint’. The aim of this paper is to derive a significant part of Thurston’s fundamental theory of measured laminations topologically.

1 Curve system

1.1 The complex of curve systems

Let M be a compact surface, with or without boundary. A **curve system** in M means a finite collection of disjointly embedded curves which are either circles not bounding a disks in M and not isotopic to components of ∂M or arcs with endpoints in ∂M, not isotopic to arcs in ∂M relative to endpoints. Let $\mathcal{CS}(M)$ be the set of isotopy classes of curve systems in M. We can identify a nonempty curve system with anti positive number of parallel copies of itself which gives us the set $\mathcal{PS}(M)$, called projective isotopy classes of curve systems.
Every curve system can be expressed uniquely in the form $n_0C_0 + ... + n_kC_k$ where the $C_i's$ are connected, non-isotropic curve systems and n_iC_i denotes n_i parallel copies of $C_i, n_i > 0$. Let PS(M) be a simplicial complex with k-simplices corresponding to isotopy classes of such $(k + 1)$- tuples $[C_0, ..., C_k]$. The various faces of such a k-simplex are obtained by deleting one of the $C_i's$. $\mathcal{PS}(M)$ is the set of points of PS(M) having rational barycentric (projective) coordinates i.e. $n_0C_0+...+n_kC_k$ has coordinates $n_0[C_0]+...+n_k[C_k]$. For example, $3C_0+4C_1$ is the point on the edge $[C_0, C_1]$ four-sevenths of the way from C_0 to C_1.

Example1.1 Let M be pair of pants i.e. S^2 minus three disks. In this case curve systems can have only arcs. There are total six isotopy classes, namely, one joining each pair of boundary circles (see Fig 1.1(a)) and one joining each boundary circle to itself and separating the other two boundary circles (see Fig 1.1(b). Hence PS(M) has six vertices and it spans four triangles as shown in Fig 1.2.

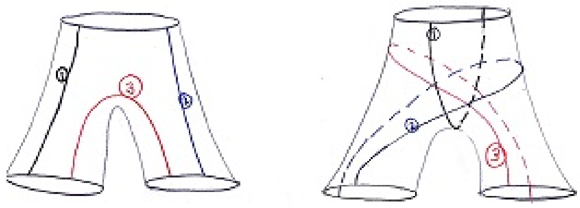

Figure 1.1 (a),(b)

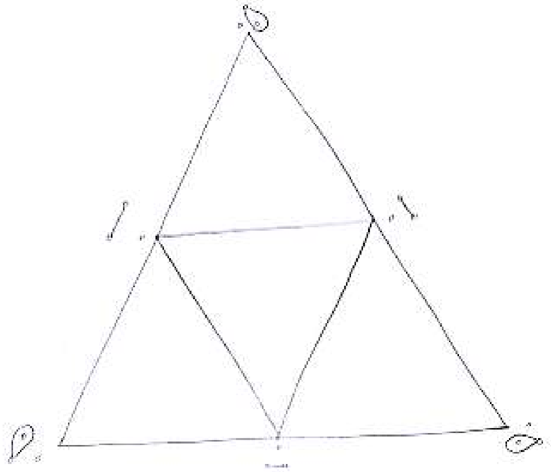

Figure 1.2

Example1.2 Let M be a twice punctured projective plane. Viewing this as a square with corners deleted and antipodal boundary points identified, as in Fig 1.3. There will be ten isotopy classes which will span a square as shown in Fig 1.3.

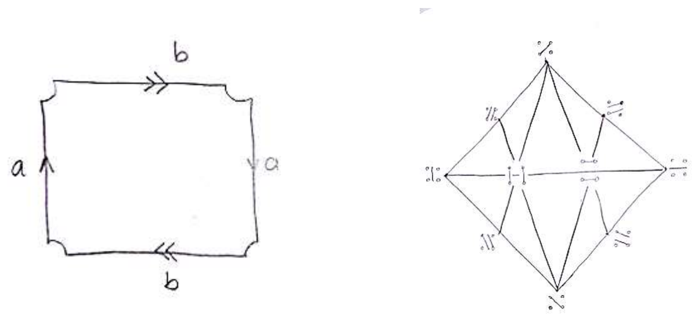

figure 1.3

1.2 Train tracks

A train track $\tau \subset M$ is a compact submanifold meeting ∂M transversely, except at finitely many branching points in the interior of M where two arcs merge into one, all three arcs having a common tangent direction. That is a train track is a closed subset locally diffeomorphic to Fig 1.4(a). A measure on a track τ is an assignment of weights $\alpha_i \geq 0$ to the components of the non-branchpoint locus of τ, satisfying the equation $\alpha_i = \alpha_j + \alpha_k$ at each branch point as in Fig 1.4(a). A measure $\alpha = (\alpha_1, ..., \alpha_n)$ with each α_i an integer, determines a curve system $S_\alpha \subset M$ by taking α_i parallel copies of the ith nonsingular arc of τ and matching these arcs at the branch points as shown in Fig 1.4 (b), using the branch equation. We say S_α is carried by τ.

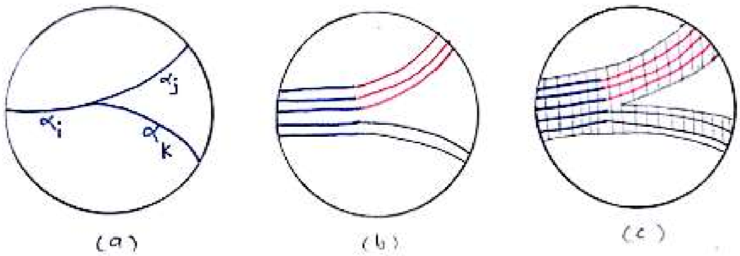

Figure 1.4

A good train track is the one which does not have the complementary regions as shown in Fig 1.5. Next we will see why these are called 'good'.

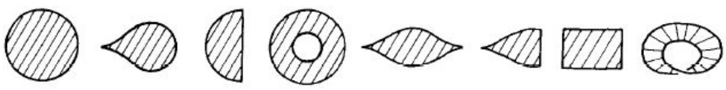

Figure 1.5

Lemma 1. *If τ is good and S_α is carried by τ, then $S_\alpha \in \mathcal{CS}(M)$.*

Proof. Let $N(\tau)$ be a fibered neighborhood of τ as in Fig 1.4(c), with fibers transverse to S_α. Extend the tangent linefield of S_α to a linefield on $N(\tau)$ transverse to fibers and tangent to $\partial N(\tau) - \partial M$. This linefield can be extended to a linefield on M having isolated singularities, and transverse to ∂M. For such linefields there is an index theory. Each isolated singularity is assigned an index, measuring how many times the linefield rotates as one goes around small circles enclosing the singularity; here, 180 degree counts as one rotation since the lines are not oriented. So for vector fields, the corresponding linefield index is twice the vector field index.

The key fact about linefield index is that the sum of the indices at all the singularities is a topological invariant, namely, twice the Euler characteristic. It is not hard to check that the excluded regions in Fig 1.5 are exactly those having nonnegative total index. In particular, we note that $\chi(M)$ must be negative if M contains a nonempty good track.

Returning now to the linefield we constructed on M, suppose S_α contains a circle bounding a disk in M. Then the total index in this disk is positive. The singularities of the linefield in this disk, however, have a negative sum since they determine the indices of the complementary regions of τ in this disk. So S_α can have no trivial circles components. Similarly, S_α contains no circle or arc parallel to ∂M. $\square$

Another useful consequence of linefield index theory is the fact that a subtrack of a good track is also good.

1.3 Standard tracks

We wish to construct a finite set of standard tracks which carry all curve systems on M. These standard tracks are not canonically defined, but depend on choosing disjoint circles $S_1, \ . \ . \ . \ , S_k$ which split M into pairs of pants $P_1, \ . \ . \ . \ , P_n (n = -\chi(M) > 0)$. Some S_j's will be one-sided if M is non-orientable; let $\tilde{S}_j$ denote the boundary of a Mobius band neighborhood of a one-sided S_j.

Consider an arbitrary curve systems in $\mathcal{CS}(M)$. Isotope S to minimize the number of intersection points with the splitting circles S_i. Then S is the disjoint union of systems S' and S'' where S' consists of circles parallel to the S_i's, and S'' meets each P_j in a curve system in $\mathcal{CS}(P_j)$. Each system $S'' \cap P_j$ is carried by one of the four basic tracks in Fig 1.6 (compare with Example 1). To reassemble S'' across a two-sided S_i, we must allow for possible twisting along S_j. A track which achieves this is obtained by inserting one of the two "connectors" in Fig 1.7(a) or (b), which contain S_j as a subset. For one-sided S_j no nontrivial twisting along S_i is possible, so the connector in Fig 1.7(c) suffices to carry S'' if S'' meets S_j in this case. In this way we form $2^t 4^n$ tracks in M, t being the number of two-sided S_i's. All of these tracks, together with their various subtracks, are called standard tracks. These standard tracks suffice to carry S'' and also the components of S' parallel to two-sided S_i's. To carry circles of S' parallel to one-sided S_i's, we enlarge our definition of *standard* to include tracks $\tau \cup S_j$ where τ is a standard track disjoint from the one-sided splitting circle S_j. Thus we obtain finitely many standard tracks carrying all curve systems on M.

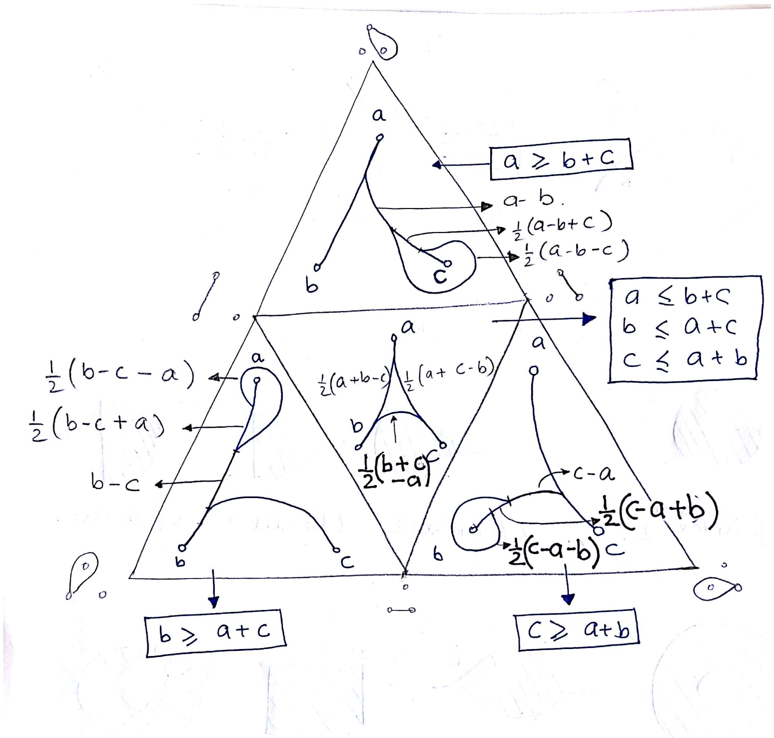

figure 1.6

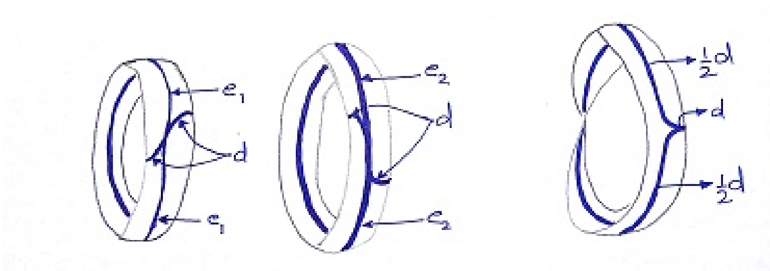

figure 1.7(a),(b),(c)

1.4 The polyhedra ML(M) and PL(M)

Each standard track τ has its cone $C(\tau)$ of measures, whose faces are identifiable with the cones $C(\tau')$ associated to the subtracks $\tau' \subset \tau$. Making all such identifications, there results a polyhedron ML(M). Projectivizing this construction by deleting $0 \in$ ML(M) (corresponding to the empty curve system) and factoring out by scalar multiplication, one obtains a finite polyhedron PL(M), with a decomposition into convex linear cells, the projectivized cones $C(\tau)$. Note that the simplices of PS(M) embed linearly into cells of PL(M), namely, $[C_0, \ . \ . \ . \ , \ C_k]$ embeds in the projectivization of $C(\tau)$ for T a standard track carrying $C_o \cup \cdots \cup C_k$.

If γ is any loop in M, not necessarily embedded, let $i_\gamma : \varphi\mathcal{S}(M) \to [\mathrm{O}, \ \infty)$ be the function which assigns to a curve system S the minimum number of intersection points of S with loops homotopic to γ.

Lemma 2. *For γ transverse to $S, i_\gamma(S) = |S \cap \gamma|$ iff no arc of $\gamma - S$ can be homotoped rel its endpoints into S.*

Proof. Suppose there is a homotopy $F : S^1 \times I \to M$ of $\gamma = F|S^1 \times \{0\}$ which decreases the number of intersection points with S. We may assume F is transverse to S, so $F^{-1}(S)$ is a one-dimensional submanifold of $S^1 \times I$ containing, by hypothesis, at least one arc with both endpoints on $S^1 \times \{0\}$. An innermost such arc cuts off from $S^1 \times I$ a half-disk, and restricting F to this half-disk gives a homotopy of an arc of $\gamma - S$ rel endpoints into S. $\square$

From this lemma it follows incidentally that if γ is an embedded loop, then $i_\gamma(S)$ is also the minimum number of intersection points of S with embedded loops isotopic to γ.

Lemma 3. *There exist finitely many embedded loops γ_m in M such that any two curve systems corresponding to distinct points of* $\mathrm{ML}(M)_{\mathbb{Z}}$ *are distinguished by their intersection numbers with one of the γ_m's.*

Proof. For a start, we take for γ_m's all the boundary curves of the pairs of pants P_j, i.e., the components of ∂M, the two-sided S_j's, and the $\tilde{S}_i$'s associated to one-sided S_i's. Intersection numbers with these γ_m's are given by the three boundary weights a, b, c of each of the basic tracks in Fig 1.6, by Lemma 2. Next consider a two-sided S_i. We shall choose two more curves γ_m such that intersection numbers with these two curves detect the "twist parameters" e_1 and e_2 in Fig 1.7(a). There are two subcases, according to whether the two "sides" of S_j belong to the same or different P_j's. One of the two choices for γ_m is shown in Fig 1.8(a). We are interested in how the twist parameter e_i affects the intersection number with this γ_m. We note that the intersection number of a curve system S carried by a standard track with γ_m can be computed using only the subsurface $M_0 \subset M$ shown in Fig 1.8(a), namely it is the intersection number of $S \cap M_0$ with γ_m in M_0 .

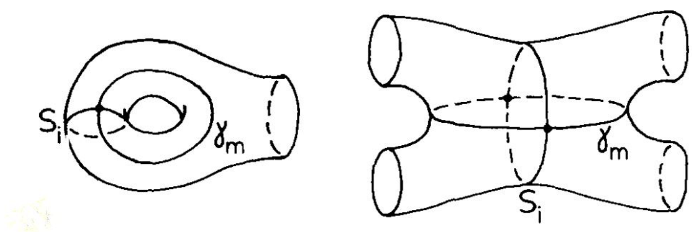

figure 1.8(a),(b)

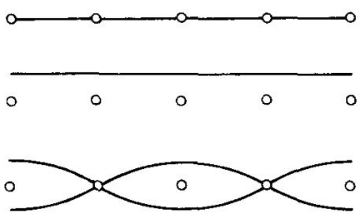

figure 1.9

To see the effect of varying e_j on this intersection number, it is convenient to lift to a covering space $\mathbb{R}^2 - \mathbb{Z}^2$ of M_o. In case M_0 is the punctured torus, this cover is the universal cover of the torus with the preimages of the puncture deleted. If M_0 is a 4-punctured sphere, there is the well-known 2-sheeted branched covering of a sphere by a torus, branched at the four punctures, and $\mathbb{R}^2 - \mathbb{Z}^2$ is again the universal cover of the torus with preimages of the four punctures deleted; the group of deck transformations in this case is generated by 180∘rotations about the points of $\mathbb{Z}^2$. We choose coordinates in $\mathbb{R}^2 - \mathbb{Z}^2$ so that S_j lifts to a slope ∞ line and γ_m lifts to a slope 0 line. For the punctured torus case, curve systems can be isotoped to consist of parallel copies of:

(1) arcs which lift to lines of rational slope passing through points of $\mathbb{Z}^2$, and

(2) circles which lift to lines of rational slope disjoint from $\mathbb{Z}^2$.

For the 4-punctured sphere case, there are also:

(3) arcs whose lifts are obtained from the lines in (1) by pushing of alternate points in $\mathbb{Z}^2$.

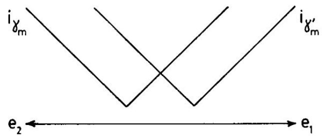

figure 1.10

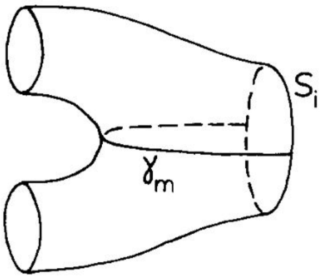

figure 1.11

Looking in this $\mathbb{R}^2 - \mathbb{Z}^2$ cover, we can see that increasing e_1 increases the slope of the lifted curves, while increasing e_2 decreases slopes. Intersection number with γ_m therefore increases with $e_{\rm i}$. So γ_m detects the amounts of twisting, with at most a two-fold ambiguity, the direction of twisting.

For this take in addition to γ_m, the loop γ'_m obtained from γ_m by Dehn twist along S_i. See Fig 1.10 for sketches of the graphs of these intersection numbers as functions of e_i. (One graph is a translate of the other.)

Consider a one-sided S_j. Here we are not concerned with twisting along S_j, but with detecting the presence of parallel copies of S_j in curve systems which are disjoint from $\tilde{S}_i$. This can be done by intersection number with one loop γ_m as shown in Fig 1.11. □

Let $\mathrm{ML}(M)_{\mathbb{Z}}$ denote the points of $\mathrm{ML}(M)$ having integer coordinates.

Proposition 1.1. *The map $ML(M)_{\mathbb{Z}} \to \mathcal{CS}(M)$ is a bijection.*

Proof. Surjectivity can be easily checked using previous construction. Injectivity follows from the above lemma. □

1.5 The global structure of ML(M) and PL(M)

Let b be the number of boundary components of M and let $\chi = \chi(M)$.

Proposition 1.2. $\mathrm{ML}(M)$ *is piecewise linearly homeomorphic to* $\mathbb{R}^{-3\chi-b} \times [0,\ \infty)^b$*, preserving scalar multiplication, with projection to the* $[0,\ \infty)$ *factors given by the weights at the b boundary circles of M. Consequently* $\mathrm{PL}(M)$ *is piecewise linearly homeomorphic to the join of a sphere* $S^{-3\chi-b-1}$ *with a simplex* Δ^{b-1}.

Proof. This will be by induction on k, the number of circles S_i in the splitting of M into pairs of pants. It is easy to check for pair of pants.

For the induction step, consider first the case of splitting M along a two-sided S_i to form the surface M'. Referring to Fig 1.8, we see that in passing from M' to M, two boundary weights d_1 and d_2 are set equal, and new weights e_1 and e_2 measuring the twisting around S_j are introduced. Thus the coordinates $(d_1,\ d_2) \in [0,\ \infty)^2$ are replaced by coordinates $(d,\ e_1)$ or $(d,\ e_2)$ .These two new quadrants $[0,\ \infty)^2$ have their boundaries identified as in Fig 1.11

to form, piecewise linearly, a full plane $\mathbb{R}^2$, since if either $d = 0$ or $e_1 = e_2 = 0$, the two alternative subtracks in Fig 1.7(a) coincide. Thus in $\mathbb{R}^{-3\chi-b} \times [0,\ \infty)^b$, a subproduct $[0,\ \infty)^2$ in the second factor shifts to an $\mathbb{R}^2$ in the first factor, completing the induction step in this case.

If S_j is one-sided, we form M from M' by adjoining the Mobius band in Fig 1.7(c). The boundary weight $d \in [0, \infty)$ for M' is carried along to the new track for M, becoming $\frac{1}{2}d$ around the new branch. In addition, if $d = 0$ we are allowed to enlarge our track by adding a copy of the loop S_i, with arbitrary weight $e \geq 0$. If both d and e are zero we have a common subtrack, so the new factor $[0,\ \infty) = \{e \geq 0\}$intersects the old $[0,\ \infty) = \{d \geq 0\}$ only in the origin, and the union of these two $[0,\ \infty)$'s is an $\mathbb{R}$. So in $\mathbb{R}^{-3\chi-b} \times [0,\ \infty)_b$, a factor $[0,\ \infty)$ shifts to an $\mathbb{R}$ factor when we pass from M' to M. This finishes the induction step.

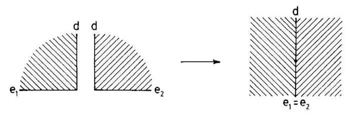

figure 1.12

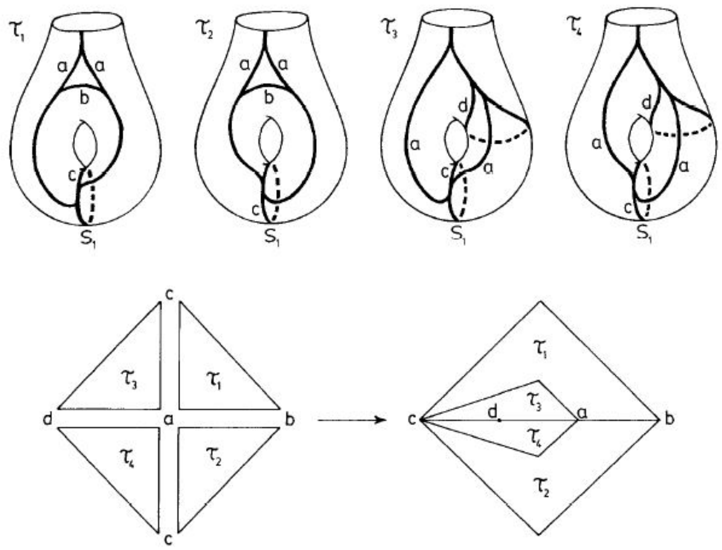

figure 1.13

□

2 Measured Lamination

The idea here is to provide an interpretation for the non-integral points of ML(M), and likewise for the irrational points of PL(M), as topological objects in M. At the same time we shall eliminate the dependence of these spaces on a choice of decomposition of M into pairs of pants, which was part of their original definition.

2.1 Construction of measured laminations

Given a track τ and a positive measure $\alpha \in C(\tau)$ i.e., with all weights $\alpha_i > 0$, we can construct a foliation N_α of the fibered neighborhood $N(\tau)$ as follows:
1.Decompose $N(\tau)$ as a union of rectangles lying over the nonsingular arcs of τ, the fibers of $N(\tau)$ giving the vertical direction in these rectangles.
2. Let the ith rectangle strip have width α_i. The branch equations $\alpha_i = \alpha_j + \alpha_k$ allows these paper rectangles to be matched together at their vertical ends to form a copy of $N(\tau)$, with the foliation N_α determined by the horizontal lines in the rectangles.
3.N_α has singularities at the cusp points of $\partial N(\tau)$. To eliminate these singularities, slit N_α open along its finitely many singular leaves, starting at the cusps. If all the weights α_j are rational, the singular leaves of N_α are compact, and the slitting open process is finite, yielding a thickening L_α of a curve system on M, with the product foliation (or twisted product, in the case of mobius band components of L_α). But if there are noncompact singular leaves of N_α, some care must be taken to damp down the magnitude of the slitting fast enough so that the process converges.

The result of this slitting is the **measured lamination** L_α, which lies in $N(\tau)$ transverse to the fibers. An alternative viewpoint which avoids the subtleties of slitting along noncompact leaves is to consider equivalence classes of N_α's under the equivalence relation generated by isotopy and by slitting only along compact arcs in leaves, starting at cusps as before.

Let $\mathcal{ML}(M)$ denote the set of equivalence classes of measured N_α's carried by good tracks in M, and let $\mathcal{PL}(M)$ be the projectivization of $M\mathcal{L}(M)$, taking non-empty N_α's and factoring out by scalar multiplication of α.

In each equivalence class in $M\mathcal{L}(M)$ there is a unique (up to isotopy) representative N'_α obtained by slitting completely all the compact singular leaves, since slitting a compact subarc of a noncompact singular leaf can be realized by isotopy. Certain components of N'_α are foliated by parallel compact leaves, forming a thickening of some curve system in M. The claim is that in the other components of N'_α there are no compact leaves, and the noncompact singular leaves (along which one would slit to form L_α) are dense. To see this, consider a short vertical arc in N'_α with one endpoint on a fixed nonsingular leaf. This arc can be translated horizontally along leaves to any prescribed distance without encountering cusps of N'_α if the vertical segment is short enough. If the given leaf is compact, a suitably short vertical arc therefore eventually returns to its initial position exactly, and the leaf has a neighborhood of compact leaves. So compact leaves in N'_α are open; they are also obviously closed, and so form certain components of N'_α. If translation of a fixed vertical arc could be continued indefinitely along a noncompact leaf without being obstructed by a cusp of N'_α, it would return infinitely often to the same fiber of N'_α with a vertical translation, and so give this fiber infinite measure, which is impossible. Thus translation of every vertical segment along a non-compact leaf must eventually meet a cusp, making the singular leaves dense in the noncompact-leaf components of N'_α. Two consequences of this are:

1. The fully slit open lamination L_α meets fibers of $N(\tau)$ only in intervals and Cantor sets.

2. Only the compact-leaf components of N'_α can meet ∂M, since obviously a noncompact singular leaf along which one would slit to form L_α cannot meet ∂M, the only endpoint of this leaf being a cusp point in the interior of M.

2.2 Length function

Given N_α as above and a loop γ in M, consider loops homotopic to γ which meet N_α in finitely many arcs which are either vertical (in fibers of $N(\tau)$) or horizontal (in leaves of N_α); such loops we call PVH-piecewise vertical or horizontal. PVH loops have a *length*, the total length of all their vertical segments. Define $l_\gamma(N_\alpha)$ to be the infimum of the lengths of all PVH loops homotopic to γ. (We will show that this infimum is always realized, assuming τ is good.) The notion of "vertical" depends on the vertical fiber structure of N_α, which is not really part of the data of N_α. To avoid this, we could instead use PTH paths: piecewise either transverse to N_α or horizontal. Length can be defined just as well for PTH paths, and leads to the same $l_\gamma(N_\alpha)$.

Note that $l_\gamma(N_\alpha)$ is linear with respect to scalar multiplication of α. Also, $l_\gamma(N_\alpha)$ is clearly constant on equivalence classes in $M\mathcal{L}(M)$, and so defines a map $l_\gamma : M\mathcal{L}(M) \to [0,\ \infty)$. If the measure α is integral, it is easy to see that $l_\gamma(N_\alpha) = i_\gamma(S_\alpha)$ where S_α is the curve system associated to $\alpha \in C(\tau)$. So l_γ extends i_γ from $\mathcal{BS}(M)$ to $M\mathcal{L}(M)$.

If α has some coordinates $\alpha_j = 0$, this corresponds to passing to a subtrack of τ on which α becomes positive, in other words to a face of $C(\tau)$. So we can regard l_γ as a function $C(\tau) \to [0,\ \infty)$.

Proposition 2.1. *For τ a good track and γ a loop in M, the function $l_\gamma : C(\tau) \to [0,\ \infty)$ is piecewise linear.*

Sketch of proof: We begin by putting γ into a minimal position with respect to τ. We consider γ's which are divided into finitely many smoothly immersed *segments* which lie either inside τ, or outside τ with endpoints on τ. For example, if γ is transverse to τ, the intersection points with τ determine a segmentation of γ. For a segmented γ, any backtracking along τ can easily be eliminated (by homotopy of γ) without adding new segments, so we may assume that at the endpoints of its segments, γ either leaves or enters τ, or γ stays in τ but reverses direction and switches branches at a cusp of τ. Possibly γ is a single segment consisting of a smoothly immersed circle in τ or $M - \tau$ a ("segment" without endpoints).

We choose γ within its homotopy class to have the minimum number of segments outside τ, and among γ's with this number of segments outside τ, the minimum number of segments inside τ.

figure 1.14

The track τ is obtained from N_α by collapsing fibers of N_α. Let γ_1 be a PVH loop which projects to γ under this collapse (and hence is homotopic to γ). Keeping γ_1PVH, we can by a finite number of vertical deformations pull it *taut* in N_α, so that it has no configuration as in Fig. 14 which can be shortened by pushing vertically a horizontal piece of γ_1.

A taut γ_1 can clearly be chosen to vary continuously with α, so $l_\gamma(N_\alpha)$ is at least a continuous function of α. To show piecewise linearity, look at a piece of γ_1 projecting to a segment of γ in τ. This piece of γ_1 moves monotonically through a string of rectangles of N_α formed by the fibers through cusp points, as shown in Fig. 18. The heights of these rectangles are given by weights α_j. The net vertical distance between any two horizontal edges in this chain of rectangles is a linear function of the α_j's with $\mathbb{Z}$ coefficients. This can be seen by induction on the number of rectangles between the two horizontal edges; in the induction step one either adds or subtracts an α_i, since two adjacent rectangles always have a horizontal edge on the same level.

Where two horizontal edges are at the same height therefore defines a hyperplane in $C(\tau)$ (if it is not all of $C(\tau)$). On the complementary components of the union of all such hyperplanes, the length of γ_1 is a $\mathbb{Z}$-linear function of α, since each stretch of γ_1 going monotonically (in the weak sense) up or down has length given by the vertical distance between two horizontal edges of rectangles. (This holds true even in the special case that γ is a smoothly immersed circle in τ and γ_1 is globally monotone up or down, since the length of γ_1 in this case is the vertical distance from a horizontal edge to itself as measured all the way around γ_1.)

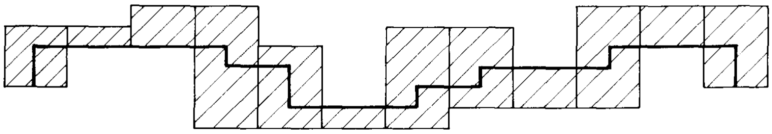

figure 1.15

We remark that we have proved slightly more than stated in proposition, since we have shown that where l_γ is linear, it has integer coefficients. This will be useful in proving Lemma below.

2.3 Intrinsic structure via length function

Call a collection of n loops $\gamma \subset M$ *injective* if the map $\mathcal{CS}(M) \to [0, \infty)^n$ having the associated intersection number functions i_γ as coordinates is injective. Previous lemma says that injective collections exist.

Lemma 4. *For an injective collection of loops γ, the map $l : \mathrm{ML}(M) \to [0, \infty)^n$ whose coordinates are the associated length functions l_γ is a piecewise linear homeomorphism onto its image.*

Proof. We can subdivide $\mathrm{ML}(M)$ into finitely many polyhedral cones on each of which l is $\mathbb{Z}$-linear, in particular $\mathbb{Q}$-linear. If l were non-injective on one cone, it would be non-injective on rational points, by linear algebra. (The kernel of a linear transformation with $\mathbb{Q}$ coefficients has the same dimension over $\mathbb{Q}$ as over $\mathbb{R}$.) Clearing denominators, l would then be non-injective on integer points, contrary to hypothesis. Similarly, if l took two points in different cones to the same point, this would already happen for a pair of integer points. So l is injective on all of $\mathrm{ML}(M)$. Since l is linear on the finitely many cones covering $\mathrm{ML}(M)$, it must be a piecewise linear homeomorphism onto its image. □

2.4 Normal form for measured laminations

Let a decomposition of M into pairs of pants be given, defining the polyhedron $\mathrm{ML}(M)$.

Lemma 5. *Given $N_\alpha \in M\mathcal{L}(M)$ and a curve system γ consisting of circles, then N_α can be slit and isotoped until each circle of γ either is a leaf of N_α or meets N_α in fibers of $N(\tau)$ whose total length realizes the minimum length in the homotopy class of that circle.*

Proof. Consider PVH embedded γ's which are divided into finitely many segments which lie either outside $N(\tau)$, with endpoints on $\partial N(\tau)$, or inside $N(\tau)$, proceeding monotonically (in the weak sense) through fibers, and transverse to fibers through cusp points of $\partial N(\tau)$. Define the complexity of such a γ to be the lexicographically ordered triple (# segments outside $N(\tau)$, # segments inside $N(\tau)$, # points of intersection with cusp fibers). Choose γ of minimal complexity within its isotopy class. This implies that γ passes monotonically through the rectangles of $N(\tau)$ bounded by the cusp fibers, except for arcs as in Fig 1.16(a).

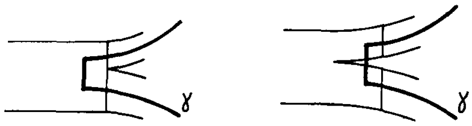

figure 1.16 (a),(b)

Now pull γ taut, by vertical *homotopy.* Again this taut γ achieves the minimum length within its homotopy class. The proof of this is virtually identical with the one given before, and will be left from the reader to verify in detail. The one minor difference (aside from the fact that we are now minimizing complexity over the isotopy class of γ, rather than the homotopy class) is that an edge most arc a with both endpoints on $\partial_+ D^2$, cutting off a disk from D^2 which maps to A, gives rise to the backtracking in Fig 1.16(a), which is not ruled out. The presence of such arcs a makes little difference to the rest of the argument however, which eventually eliminates them anyway.

Consider now the process of pulling γ taut. Slitting as in Fig 1.16(b) to eliminate non-monotonicity of γ with respect to fibers of $N(\tau)$ will not affect the process, so we may assume γ is monotonic. Eliminating configurations as in Fig. 14 may introduce self-intersections of γ, but these are of a special kind: they can be removed by small perturbations of γ.

Having γ in taut position, we first slit N_α along its compact singular leaves, which introduces some segments of γ outside $N(\tau)$. Achieving the desired conclusion for the compact-leaf components of the resulting N'_α is very easy, so we need only worry about the other components. In these, the singular leaves are dense, as observed earlier in this section, so we may slit N'_α until each vertical piece of γ has an endpoint on $\partial N'_\alpha$. Then with a little more slitting, the remaining horizontal pieces can be eliminated. Finally, γ can be perturbed back to an embedding. This is homotopic, hence isotopic, to the original γ. □

Proposition 2.2. *The map* $\mathrm{ML}(M) \to M\mathcal{L}(M)$ *is a bijection, hence* $l : \mathcal{ML}(M) \to [0, \infty)^\infty$ *is injective.*

Proof. The map $\mathrm{ML}(M) \to M\mathcal{L}(M)$ is injective since its composition with l is. Surjectivity will use the above lemma. Take N_α satisfying the conclusion of above lemma with γ the circles determining the given decomposition of M into pairs of pants P_j. After further slitting, N_α decomposes as the disjoint union of N'_α and N''_α where N'_α consists of all circle leaves of N_α parallel to circles of γ. Choose a track τ carrying N_α which is the disjoint union of tracks τ' and τ'' carrying N'_α and N_α where $\tau' \subset \gamma$ and τ'' is transverse to γ. We may take τ to be a splitting of a good track carrying the original N_α, so τ is good, hence also τ' and τ

The essential point is to see now that all leaves of $N''_\alpha \cap P$ are compact. Fixing j, we may assume the previous slitting of N_α eliminated any compact singular leaves of $N''_\alpha \cap P_j$. As observed earlier, all leaves of $N''_\alpha \cap P_j$ meeting ∂P_j must then be compact, forming certain foliated rectangle components of $N''_\alpha \cap P$ Let $\dot{N}_j$ be the union of the remaining components of $N''_\alpha \cap P_j \backslash$, carried by a subtrack τ_j of $\mathcal{T}'' \cap P$ contained in int (P_j). This τ must be good in P_j, since the only possible complementary region in Fig. 5 is an annulus with one boundary circle in τ_j and the other boundary circle in ∂P_j, but this would force N_j, and hence N_α to have circle leaves parallel to ∂P_j, contrary to the definition of N_α If $N \neq \emptyset$, then $C(\tau_i) \neq \emptyset$, so $C(\tau_j)$ would contain rational points corresponding to curve systems in P_j disjoint from ∂P_j. We conclude that $N = \emptyset$, verifying that all leaves of $N'_\alpha \cap P_j$ are compact arcs. Next we pinch N''_α (the reverse of the slitting operation), without disturbing N'_α , so that for all j, no complementary region of $N''_\alpha \cap P_j$ in P is one of the types in Fig. 5. The only bad regions which might occur are half-digons (sixth type in Fig. 5) and rectangles, since half-disks (third type) cannot occur, otherwise some component of γ would not achieve the minimum length in its homotopy class. Pinch according to the rule: half-digons first, then rectangles. Pinching rectangles might create half-digons, which should then be pinched before any more rectangles are pinched, to be sure that no digons are created. Since the total number of complementary regions in the P_j's decreases with each pinching, this is a finite process. After this pinching, each $N''_\alpha \cap P$ is carried by a standard track in P and so clearly N_α itself is carried by a standard track. □

Application to diffeomorphisms of compact surfaces without boundary.

Using the topological machinery developed so far, we can easily prove the following proposition.

Proposition 2.3. *Suppose M is compact and ∂M is nonempty. Then a diffeomorphism, f: $M \to M$ can be isotoped so that either*

(1) f has finite order, or

(2) f leaves invariant a curve system consisting of circles in M, or

(3) f takes a measured lamination in int (M) to a scalar multiple of itself.

Proof. A diffeomorphism f induces a homeomorphism of $M\mathcal{L}(M)$, given by a permutation of the coordinate length functions l_γ, and hence also a homeomorphism of $\mathcal{PL}(M)$. This is a ball if $\partial M \neq 0$, so by the Brouwer fixed point theorem, there must be a (nonempty) measured lamination $L_\alpha \subset M$ with $f(L_\alpha) = L_{\lambda\alpha}$ for some $\lambda > 0$. If L_α is disjoint from ∂M, we are in case (3), so we may assume L_α meets ∂M. Leaves of L_α meeting ∂M are compact. These yield a curve system consisting of arcs in M which is invariant under f. Let M' be M split open along these arcs, and let M'' be M' minus those components which are either disks, or annuli with one boundary circle in ∂M. If $M'' \neq \emptyset$ then the circles of $\partial M''$ which are not circles of ∂M give rise to a system of circles in M invariant under f, and we are in case (2). If $M'' = \emptyset, f$ can clearly be isotoped to be of finite order. □

In case (2), M and f can be split open along the invariant circle system to obtain a simpler situation which can be analyzed inductively. Only a small amount of information is lost in this splitting process: Dehn twists along the invariant circles. If we are in case (3) but not case (2), then the invariant lamination L_α has no compact leaves, and each complementary region of L_α must be a disk with some number of cusps on its boundary, and possibly one puncture (component of ∂M) in its interior.

The full strength of Thurston's theorem on surface diffeomorphisms which holds also when M is closed–is that in this last case, f also preserves another measured lamination L_{l3} transverse to L_α, with $f(L_\alpha) = L_{\lambda\alpha}$ and $f(L_\beta) = L_{\beta/\lambda}$ for some $\lambda \neq 1$. Since this approach of the proof is topological in nature it can be generalised for proving similar results for 3-manifolds.

3 References

[1] A.E.Hatcher, Measured lamination spaces for spaces from the topological viewpoint, Topology and it's Applications 30 (1988) 63-88.