

DEGENERATIONS OF HYPERBOLIC STRUCTURES ON SURFACES

Christopher J Leininger

*Department of Mathematics, University of Illinois, Urbana–Champaign, IL,
USA*

Email: clein@math.uiuc.edu

These are lecture notes from the summer school on *Geometry, Topology and Dynamics of Character Varieties* at the Institute for Mathematical Sciences, National University of Singapore in July of 2010. The objective is to give an introduction to some of the tools used in studying the degeneration of hyperbolic structures on surfaces as developed by Thurston, and in particular to describe his construction of a compactification of Teichmüller space.

1. Introduction

These notes are a continuation of lecture notes by J. Aramayona [2] which, beginning with an introduction to the geometry of the hyperbolic plane \mathbb{H}^2 , introduces the notion of a hyperbolic structure on S , a closed oriented surface of genus $g \geq 2$, and studies the space of equivalence classes of all hyperbolic structures, the Teichmüller space of S . In these notes we will provide a somewhat informal introduction to some of the work of W.P. Thurston on degenerations of hyperbolic structures and his compactification of Teichmüller space. After some preliminary discussion of length functions on Teichmüller space (Section 2), we introduce both measured foliations (Section 3) and measured geodesic laminations (Section 4), we explain how one can pass back and forth between these two types of objects (Section 5) as well as how to parametrize the space of all such (Section 6), and finally we explain how Thurston compactifies Teichmüller space via measured foliations (Section 7). The main references for the material discussed here are [16, 22, 13, 6, 24, 34, 12].

Acknowledgements. I would like to thank Javier Aramayona, David Du-

mas, Richard Kent, Dan Margalit and Caroline Series for helpful discussions throughout the writing of these notes, as well as Ser-Wei Fu, Pradthana Jaipong and the referee for many useful comments. I would especially like to thank Ser Peow Tan and the IMS at the National University of Singapore for their hospitality during the summer school on *Geometry, Topology and Dynamics of Character Varieties* in July 2010.

1.1. Notation, terminology and some conventions

We denote the Teichmüller space of S by

$$\mathcal{T}(S) = \{[f, X] \mid f : S \rightarrow X = \mathbb{H}^2/\Gamma\}$$

where $f : S \rightarrow X$ here is an orientation preserving homeomorphism, $\Gamma < \text{Isom}^+(\mathbb{H}^2) = \text{PSL}(2, \mathbb{R})$ is a discrete group, and $[f, X]$ represents the equivalence class via the equivalence relation $(f, X) \sim (h, Y)$ if and only if $h \circ f^{-1} \simeq \psi : X \rightarrow Y$, where ψ is an isometry. When there is no confusion, we write X instead of $[f, X]$. We also often think of a point in $\mathcal{T}(S)$ as a hyperbolic metric on S (up to isotopy). From $f : S \rightarrow X$ we obtain this hyperbolic metric $\rho = \rho_X$ by pulling back the metric on X by f .

The Teichmüller space naturally injects into the set of discrete faithful representations into $\text{PSL}(2, \mathbb{R})$, modulo conjugation within the full group of isometries

$$\mathcal{T}(S) \rightarrow \mathcal{DF}(\pi_1(S), \text{PSL}(2, \mathbb{R}))/\text{conj}$$

by the holonomy homomorphism $[f, X] \mapsto f_* : \pi_1(S) \rightarrow \Gamma < \text{PSL}(2, \mathbb{R})$. This is a bijection, and we topologize $\mathcal{T}(S)$ by declaring this to be a homeomorphism onto its image. See [22, 33] for more details.

The hyperbolic metric ρ is an example of a *path metric* which means that the distance between any two points is the infimum of lengths of paths between those points. In general, given a path metric on S , if we let $p : \tilde{S} \rightarrow S$ denote the universal covering, then there is a unique path metric on \tilde{S} for which the covering map is a local isometry. This is defined by declaring the lengths of paths in \tilde{S} to be the lengths of their images in S and we call this the *induced metric* on \tilde{S} .

Because S is closed, hence compact, it follows that any metric on S is complete. Therefore, any path metric on S is a *geodesic metric*: the distance between two points is the length of some path between the points. In fact, one can check that the induced metric on \tilde{S} is also a geodesic metric. We

refer the reader to [10] for a detailed discussion of path metrics and geodesic metrics.

We make a few remarks about terminology and set some definitions. First, we will use the same name for the metric on S and on \tilde{S} with the context determining which is intended. Second, we note that there are two uses of the word “geodesic” which we will need, and it is cumbersome to always adorn this word with all the appropriate adjectives. We therefore make the following convention. An isometric embedding of an interval in \mathbb{R} into \tilde{S} (or its image) is a *geodesic* in \tilde{S} . A *locally* isometric embedding of an interval or circle of some circumference (or its image) is a *geodesic* in S . In particular, geodesic has a different meaning depending on whether we are considering it in S or \tilde{S} . When we require another meaning, we will include the necessary adjectives (e.g. we may refer to a *minimizing geodesic segment* between two points of S to refer to an isometrically embedded arc between the two points).

The set of homotopy classes of (homotopically nontrivial) curves which have a simple representative will be denoted $\mathcal{S} = \mathcal{S}(S)$. The (*geometric*) *intersection number* $\iota(\alpha, \beta)$ is defined for any pair $\alpha, \beta \in \mathcal{S}(S)$ as the minimal number of intersection points among all simple representatives. If we fix a hyperbolic metric ρ on S , then the geodesic representatives of α and β are in fact simple and realize the geometric intersection number, meaning that they intersect in precisely $\iota(\alpha, \beta)$ points. See [13, 15], for example.

2. Length functions

Given $\alpha \in \mathcal{S}(S)$, we let

$$\ell(\alpha) : \mathcal{T}(S) \rightarrow \mathbb{R}_+$$

denote the *length function* of α :

$$\ell_{[f, X]}(\alpha) = \ell_X(\alpha) := \text{length}_X(f(\alpha))$$

where $\text{length}_X(f(\alpha))$ is the length of the geodesic representative of the homotopy class of α . Equivalently, this is the ρ_X -length of the geodesic representative of α in S .

Since the length of $f(\alpha)$ in X is equal to the translation length of $f_*(\alpha_0)$, where $\alpha_0 \in \pi_1(S)$ is any element whose conjugacy class represents α , an exercise in hyperbolic geometry shows that

$$\ell_{[f, X]}(\alpha) = 2 \cosh^{-1} \left(\frac{|\text{tr}(f_*(\alpha_0))|}{2} \right). \quad (2.1)$$

This last equation shows that $\ell(\alpha)$ is a continuous function on $\mathcal{T}(S)$ for all $\alpha \in \mathcal{S}(S)$.

Our goal in this section is to use the lengths of curves to completely describe the topology of $\mathcal{T}(S)$. More precisely, write

$$\ell_* : \mathcal{T}(S) \rightarrow \mathbb{R}^{\mathcal{S}}$$

for the map $X \mapsto \{\ell_X(\alpha)\}_{\alpha \in \mathcal{S}}$.

Theorem 2.1: *The map ℓ_* is a proper embedding. In fact, there exists $\alpha_1, \dots, \alpha_{9g-9} \in \mathcal{S}$ such that*

$$\mathcal{T}(S) \rightarrow \mathbb{R}^{9g-9}$$

given by $X \mapsto \{\ell_X(\alpha_i)\}_{i=1}^{9g-9}$ is a proper embedding.

Before we give the proof, first we recall the Fenchel-Nielsen coordinates as described in Aramayona's Lectures [2]; see also [22, 15, 33, 8] for other descriptions. Let $\mathcal{P} = \{\alpha_1, \dots, \alpha_{3g-3}\}$ be a pants decomposition with transverse curves $\beta_1, \dots, \beta_{3g-3}$ (so $\iota(\alpha_i, \beta_i)$ is 1 or 2, and $\iota(\alpha_i, \beta_j) = 0$ for $i \neq j$). The coordinates

$$F = F_{\mathcal{P}} : \mathcal{T}(S) \rightarrow \mathbb{R}_+^{3g-3} \times \mathbb{R}^{3g-3}$$

are given by

$$F(X) = (\ell_1(X), \dots, \ell_{3g-3}(X), t_1(X), \dots, t_{3g-3}(X))$$

where $\ell_i(X) = \ell_X(\alpha_i)$ and $t_i(X)$ is the twist coordinate of β_i in α_i for X .

To prove the theorem, we need to show that the twist coordinates can be determined from the length functions. Fix all length coordinates ℓ_j and all twist coordinates t_j but one, say t_i . Set $X_{t_i} = F^{-1}(\ell_1, \dots, \ell_{3g-3}, t_1, \dots, t_{3g-3})$. We are thinking of X_{t_i} as a function of the single variable t_i .

Lemma 2.2: *For any curve β with $\iota(\alpha_i, \beta) \neq 0$, the map*

$$t_i \mapsto \ell_{X_{t_i}}(\beta)$$

is a proper, strictly convex function.

We postpone the proof of this lemma and sketch the proof of Theorem 2.1.

Proof: [Proof of Theorem 2.1.] Let τ_i denote the Dehn twist in the curve α_i . Since $\iota(\beta_i, \alpha_j) = 0$ if $i \neq j$, we have $\tau_i(\beta_j) = \beta_j$ in this case.

A Dehn twist in α_i applied to a hyperbolic structure changes the i^{th} twist coordinate by subtracting the length ℓ_i of α_i , and therefore

$$\ell_{X_{t_i}}(\tau_i(\beta_i)) = \ell_{\tau_i^{-1}.X_{t_i}}(\beta_i) = \ell_{X_{t_i+\ell_i}}(\beta_i).$$

By Lemma 2.2 $\ell_{X_{t_i}}(\beta_i)$ is strictly convex as a function of t_i , and since $\ell_i > 0$, it follows that the map

$$t_i \mapsto (\ell_{X_{t_i}}(\beta_i), \ell_{X_{t_i+\ell_i}}(\beta_i)) = (\ell_{X_{t_i}}(\beta_i), \ell_{X_{t_i}}(\tau_i(\beta_i)))$$

is injective. Combining this with the fact that F is injective we see that the composition

$$\mathcal{T}(S) \xrightarrow{F} \mathbb{R}_+^{3g-3} \times \mathbb{R}^{3g-3} \rightarrow \mathbb{R}^{9g-9}$$

that sends $X \in \mathcal{T}(S)$ to

$$(\ell_X(\alpha_1), \dots, \ell_X(\alpha_{3g-3}), \ell_X(\beta_1), \dots, \ell_X(\beta_{3g-3}), \ell_X(\tau_1(\beta_1)), \dots, \ell_X(\tau_{3g-3}(\beta_{3g-3})))$$

is also injective, as required.

To complete the proof, we need only verify that these lengths determine a proper function. That is, if $X_n \rightarrow \infty$ in $\mathcal{T}(S)$, then we must check that the lengths of one of our curves tends to infinity. For this we use the fact that F is a homeomorphism onto $\mathbb{R}_+^{3g-3} \times \mathbb{R}^{3g-3}$. Therefore, up to subsequence, one of the following must happen:

- There exists some i so that $\ell_{X_n}(\alpha_i) = \ell_i(X_n) \rightarrow \infty$.
- There exists some i so that $\ell_{X_n}(\alpha_i) \rightarrow 0$. In this case we see that α_i has larger and larger collar neighborhoods in X_n (see [19, 23]), and hence $\ell_{X_n}(\beta_i) \rightarrow \infty$.
- There exists some i and $0 < c < C < \infty$ so that $c < \ell_{X_n}(\alpha_i) < C$ for all n and $t_i(X_n) \rightarrow \pm\infty$. In this case we have $\ell_{X_n}(\beta_i) \rightarrow \infty$.

In each case some curve has its length tending to infinity, and this completes the proof. \square

To prove the lemma, we take a short detour into a beautiful result of Kerckhoff [24] which gives an expression for the derivative of the length of β (which we will see is increasing so that the function is convex). Since we are considering just one index at a time, let $\alpha = \alpha_i$ and $t = t_i$. Then X_t is the 1-parameter family of hyperbolic structures obtained by *cutting open along* α (that is, deleting α and taking the completion of the associated path metric), twisting length t to the right, and regluing. We consider any curve β with $\iota(\alpha, \beta) = k \neq 0$, and we write $\theta_1(t), \dots, \theta_k(t)$ for the angles at

the k points of intersection, measured counterclockwise from the geodesic representative of α to that of β in X_t . Kerckhoff's result is the following.

Theorem 2.3: *With notation as above,*

$$\frac{d}{dt}(\ell_{X_t}(\beta)) = \sum_{i=1}^k \cos(\theta_i(t)).$$

Proof: Fix some $t_0 \in \mathbb{R}$, and we prove the theorem at this t_0 . It is useful to change perspective a bit. Imagine the curve β on S with α crossing it some number of times. We assume β is given by its geodesic representative on X_{t_0} . Now cut open along α , twist a distance $s = t - t_0$ and reglue to construct the surface $X_t = X_{t_0+s}$. Because we are not changing the hyperbolic structure in the complement of α , we can view β as a union of geodesic arcs with endpoints on α . We recover the homotopy class by connecting these arcs with arcs of length s running along α . See Figure 1.




Fig. 1. Upper left: original hyperbolic structure X_{t_0} with geodesics α and β . Upper right: hyperbolic structure X_{t_0+s} obtained from X_{t_0} by cutting, shearing to the right a distance s along α , then regluing. Bottom: the geodesic representative β_s of β in X_{t_0+s} .

Inside X_{t_0+s} straighten the result to obtain the geodesic representative for β which we denote β_s . Between consecutive intersection points of α

with β_s , we have arcs, $\beta_s^1, \dots, \beta_s^k$, and we can think of these as arcs whose interiors lie entirely inside the hyperbolic surface X_{t_0+s} with α removed. However, with α removed, all the surfaces X_{t_0+s} are isometric to each other, and hence we can think of β_s^j as a 1-parameter family of geodesic arcs in X_{t_0} with endpoints in α . See Figure 2. When $s = 0$, we are at our initial hyperbolic structure and so $\beta_0^1, \dots, \beta_0^k$ are the arcs of β cut open along α in X_{t_0} .




Fig. 2. The geodesic representatives β_s for X_{t_0+s} cut open and viewed back in X_{t_0} giving a 1-parameter family of geodesic arcs β_s^j . Endpoints of the consecutive arcs are distance exactly s apart.

Let $0 = a_0 < a_1 < \dots < a_k$ be such that $a_j - a_{j-1} = \ell_{X_{t_0}}(\beta_0^j)$, and assume we have parametrized the curves as $\beta_s(u)$ so that β_0 has unit speed and $\beta_s(a_j)$ is the j^{th} point of intersection (and $\beta_s(a_0) = \beta_s(a_k)$).

Claim. For each $j = 1, \dots, k$ we have

$$\frac{d}{ds} \Big|_{s=0} \ell(\beta_s^j) = \left\langle \beta_0'(a_j), \frac{d}{ds} \Big|_{s=0} \beta_s^j(a_j) \right\rangle - \left\langle \beta_0'(a_{j-1}), \frac{d}{ds} \Big|_{s=0} \beta_s^j(a_{j-1}) \right\rangle.$$

Proof: This follows from the first variation formula in Riemannian geometry (see e.g. [14]), but can also be seen using the hyperboloid model of the hyperbolic plane inside $(\mathbb{R}^{2,1}, \langle \cdot, \cdot \rangle)$ (see [3]) by appealing to the following two facts:

- $\cosh(d(x, y)) = -\langle x, y \rangle$ for any two points x and y in the hyperboloid.
- For any unit speed parametrization of a geodesic γ in the hyperboloid model and real numbers r, u we have

$$\gamma(r + u) = \gamma(r) \cosh(u) + \gamma'(r) \sinh(u).$$

Lifting to the hyperbolic plane in the hyperboloid model, the first equation implies

$$\cosh(\ell(\beta_s^j)) = -\langle \beta_s(a_{j-1}), \beta_s(a_j) \rangle.$$

Differentiating the left-hand side with respect to s at $s = 0$ yields

$$LHS' = \sinh(\ell(\beta_0^j)) \left. \frac{d}{ds} \right|_{s=0} (\ell(\beta_s^j)),$$

whereas differentiating with respect to s and applying the second fact to the geodesic β_0 twice, the right-hand side becomes

$$\begin{aligned} RHS' &= -\left\langle \left. \frac{d}{ds} \right|_{s=0} \beta_s(a_{j-1}), \beta_0(a_j) \right\rangle - \left\langle \beta_0(a_{j-1}), \left. \frac{d}{ds} \right|_{s=0} \beta_s(a_j) \right\rangle \\ &= -\left\langle \left. \frac{d}{ds} \right|_{s=0} \beta_s(a_{j-1}), \beta_0(a_{j-1}) \cosh(\ell(\beta_0^j)) + \beta_0'(a_{j-1}) \sinh(\ell(\beta_0^j)) \right\rangle \\ &\quad - \left\langle \beta_0(a_j) \cosh(-\ell(\beta_0^j)) + \beta_0'(a_j) \sinh(-\ell(\beta_0^j)), \left. \frac{d}{ds} \right|_{s=0} \beta_s(a_j) \right\rangle \\ &= -\left\langle \left. \frac{d}{ds} \right|_{s=0} \beta_s(a_{j-1}), \beta_0'(a_{j-1}) \sinh(\ell(\beta_0^j)) \right\rangle \\ &\quad + \left\langle \beta_0'(a_j) \sinh(\ell(\beta_0^j)), \left. \frac{d}{ds} \right|_{s=0} \beta_s(a_j) \right\rangle \end{aligned}$$

Here we've used the fact that for any point \mathbf{x} in the hyperboloid and any tangent vector \mathbf{v} to the hyperboloid at \mathbf{x} we have $\langle \mathbf{v}, \mathbf{x} \rangle = 0$. Then, factoring out $\sinh(\ell(\beta_0^j))$ from both sides we arrive at the desired equation. \square

Next observe that $\beta_s^j(a_j)$ and $\beta_s^{j+1}(a_j)$ are not the same points in X_{t_0} : the former point is exactly a distance s to the right of the latter point along the geodesic α ; see Figure 2. Therefore, if we let ξ_j denote the unit tangent vector to α at the j^{th} point of intersection, then it follows that

$$\left. \frac{d}{ds} \right|_{s=0} \beta_s^j(a_j) - \left. \frac{d}{ds} \right|_{s=0} \beta_s^{j+1}(a_j) = \xi_j$$

for $j = 1, \dots, k - 1$ and similarly

$$\left. \frac{d}{ds} \right|_{s=0} \beta_s^k(a_k) - \left. \frac{d}{ds} \right|_{s=0} \beta_s^1(a_0) = \xi_k.$$

Combining the claim with this and summing over j (noting that $\beta'(a_k) = \beta'(a_0)$) we have

$$\frac{d}{ds} \Big|_{s=0} \ell_{X_{t_0+s}}(\beta) = \frac{d}{ds} \Big|_{s=0} \sum_{j=1}^k \ell(\beta_s^j) = \sum_{j=1}^k \langle \beta'(a_j), \xi_j \rangle = \sum_{j=1}^k \cos(\theta_j)$$

as required. \square

Proof: [Proof of Lemma 2.2.] We now simply observe that as we increase t , the angles $\theta_i(t)$ are strictly decreasing; see Figure 3 and [25] for a proof. Since \cos is also a strictly decreasing function on the interval $[0, \pi]$, we see that $\frac{d}{dt} \ell_\beta(t)$ is strictly increasing, hence ℓ_β is strictly convex.




Fig. 3. For $s > 0$, $s \mapsto X_{t_0+s}$ shears everything to the right when viewed in the universal covering.

The fact that $\ell_\beta(t)$ is proper also follows from this picture. As $t \rightarrow \infty$ and $t \rightarrow -\infty$, we can see that the angle tends to 0 and π , respectively. Therefore, the sum of the cosines are bounded away from 0 as $t \rightarrow \pm\infty$, hence the derivatives are bounded away from 0, and $\ell_\beta(t)$ tends to ∞ as $t \rightarrow \pm\infty$. \square

According to Theorem 2.1, we can describe the topology on $\mathcal{T}(S)$ just in terms of lengths of simple closed curves. Moreover, because the embedding to \mathbb{R}^S is proper, we can use this theorem to describe how we go to infinity in $\mathcal{T}(S)$: if $X_n \rightarrow \infty$ in $\mathcal{T}(S)$, then the length of some curve tends to infinity. In particular, the image of the sequence in \mathbb{R}^S is tending toward infinity, and we can try to look for a subsequence that becomes asymptotic to some

ray—equivalently, we can projectivize and look for a convergent sequence. Thus, we would like to find a compactification of $\mathcal{T}(S)$ by a “projective light cone” just as in the hyperboloid model of hyperbolic space—or after projectivizing, as in the Klein model.

For the latter, we would like to know that the embedding $\mathcal{T}(S) \rightarrow \mathbb{R}^S$ remains injective after projectivizing. Let $\mathbb{P}\ell_* : \mathcal{T}(S) \rightarrow \mathbb{P}\mathbb{R}^S$ be the composition of ℓ_* with the projectivization $\mathbb{P} : \mathbb{R}^S - \{0\} \rightarrow \mathbb{P}\mathbb{R}^S$.

Proposition 2.4: *The map $\mathbb{P}\ell_* : \mathcal{T}(S) \rightarrow \mathbb{P}\mathbb{R}^S$ is also injective.*

Proof: This follows from Theorem 2.1, Equation (2.1), and the fact that for any two matrices $A, B \in \mathrm{SL}(2, \mathbb{R})$,

$$\mathrm{tr}(A) \mathrm{tr}(B) = \mathrm{tr}(AB) + \mathrm{tr}(AB^{-1}). \quad (2.2)$$

The point is that if we let α and β be two simple closed curves which intersect transversely in a single point and if we write $f_*(\alpha) = A$, $f_*(\beta) = B$ (up to sign), then letting τ_β denote the Dehn twist in β we have

$$f_*(\tau_\beta(\alpha)) = AB \text{ and } f_*(\tau_\beta^{-1}(\alpha)) = AB^{-1}.$$

So by (2.1) and (2.2) there is a nontrivial relation among the lengths of the four simple closed curves $\alpha, \beta, \tau_\beta(\alpha), \tau_\beta^{-1}(\alpha)$, and one can check that this is not respected by scaling for all choices of α and β . See [16] for details. \square

3. Measured foliations

Here we introduce measured foliations which provide one interpretation for degenerations of hyperbolic structures according to Thurston. An appropriate space of all such objects also embeds into \mathbb{R}^S (Theorem 3.5) disjoint from the image of $\mathcal{T}(S)$ (Proposition 3.6), and this will provide the “light cone” for the image of $\mathcal{T}(S)$ (Theorem 7.1). As our intention is to provide intuition rather than a complete treatment of the subject, we use a fairly restricted definition which allows us to use some geometric arguments. A theorem of Hubbard and Masur [21] proves that any measured foliation is equivalent (in the sense described below) to one by our definition. For a detailed discussion of measured foliations, see [16].

3.1. Definition by example

Before giving our definition, we start with a useful example that was described in Aramayona’s lectures [2] (see also Thurston’s original exposition

of this in [32]). Given $\alpha, \beta \in \mathcal{S}(S)$ we represent these homotopy classes by simple closed curves of the same names with exactly $\iota(\alpha, \beta)$ intersection points. Recall that α, β fill up S if

$$\iota(\alpha, \gamma) + \iota(\beta, \gamma) > 0$$

for any $\gamma \in \mathcal{S}(S)$. Equivalently, with our choice of representatives, $\alpha \cup \beta$ cut the surface up into topological disks, not one of which is a bigon.

We view $\alpha \cup \beta$ as the 1-skeleton of a cell structure on S with $\iota(\alpha, \beta)$ vertices, each of which is 4-valent. We take the dual cell decomposition and declare each cell to be a Euclidean square. We assume that the squares are glued together by isometries so that the metric on the squares induces a geodesic metric on S by measuring length of a curve as the sum of the lengths of the intersections with each of the Euclidean squares.

The resulting metric is a *Euclidean cone metric*: away from the vertices the metric is locally isometric to the Euclidean plane, while at some of the vertices, there are *cone singularities*, which we now explain. At each vertex v there is some even number, $2k = 2k(v)$, of corners of squares glued together, and because there are no bigons, $2k \geq 4$, so there are at least four corners of squares glued together. If there are exactly four squares, then there is a neighborhood of the vertex which is isometric to a disk in the Euclidean plane. If there are more than four squares, then the metric is singular at v , and v is called a cone point of cone angle $k\pi$; any sufficiently small circle of radius $r > 0$ centered at this point has circumference $k\pi r$. See Figure 4. We can represent α and β by geodesics in this metric which cut through as mid-lines in each square, intersecting each other exactly in the centers.

Because the metric is locally Euclidean, parallel transport (in the complement of the singularities) is locally independent of path. Because S is orientable and because any side of a square is either parallel to α or β (but not both), we see that parallel transport around a loop γ based at x defines a linear map on the tangent space $P_\gamma : T_x(S) \rightarrow T_x(S)$ which is either the identity I or $-I$. We describe this by saying that the holonomy of the metric lies in $\{\pm I\}$. Therefore, given a line L in the tangent space of a nonsingular point x , we can consider all maximal geodesics in the complement of the singularities *in the direction* L . These geodesics are the leaves of a foliation \mathcal{F}_0 on the surface minus the singularities. Adding the singular points back in we obtain a *singular foliation* \mathcal{F} on the entire surface, with a *k-prong singularity* at each cone point with cone angle $k\pi$. See Figure 5.

Away from the cone points we can find locally isometric coordinates




Fig. 4. α and β on S and the Euclidean cone metric obtained from dual cell structure.




Fig. 5. k -prong singularities: $k = 3$ on the left and $k = 4$ on the right.

(x, y) to \mathbb{R}^2 in which the leaves are sent to arcs of vertical lines. The horizontal variation defines a measure $\mu = \mu_\gamma$ on any path $\gamma : [a, b] \rightarrow S$:

$$\int_\gamma d\mu = \int_\gamma |dx| = \int_a^b |(x \circ \gamma)'(t)| dt.$$

The pair (\mathcal{F}, μ) is an example of a singular measured foliation.

We note that μ is really a family of measures, one for every path in S , which satisfies

- If γ_0 and γ_1 are two paths homotopic through a homotopy preserv-

- ing the intersection with leaves, then $\mu_{\gamma_0} = \mu_{\gamma_1}$, and
- if γ is the concatenation of two paths γ_0 and γ_1 , then the restriction of μ_γ to γ_i is μ_{γ_i} , $i = 0, 1$.

As is common, we simply write μ for μ_γ .

We now define a measured (singular) foliation to be a pair (\mathcal{F}, μ) obtained as above from *some* Euclidean cone metric q on S . More precisely, q is a Euclidean cone metric (not necessarily built from the construction described above) with cone angles of the form $k\pi$, $k \in \mathbb{Z}_{\geq 3}$ and holonomy in $\{\pm I\}$, the singular foliation \mathcal{F} is obtained from an honest foliation \mathcal{F}_0 on S minus the cone points by parallel geodesics in some direction, and μ is the variation orthogonal to the foliation, with respect to the metric q . We let $\widehat{\mathcal{MF}}(S)$ denote the set of all measured foliation on S .

Remark 3.1: The types of metrics q we are referring to are precisely the metrics induced by *holomorphic quadratic differentials* with respect to some complex structure, and the foliations are the *vertical foliations* for those quadratic differentials; see [17] for a detailed discussion of quadratic differentials.

Given a measured foliation (\mathcal{F}, μ) , we now explain what we mean by a *leaf* of \mathcal{F} . Let \mathcal{F}_0 be the honest foliation of S in some direction with the cone points removed. In \mathcal{F}_0 we know what leaves are: these are maximal, injective locally isometric immersion of (i) \mathbb{R} , (ii) a circle of some circumference, (iii) a bounded open interval, or (iv) an open ray in \mathbb{R} (in the given direction). We will call the types (i) and (ii) *nonsingular leaves* of \mathcal{F} . Suppose now that $\gamma : (a, b) \rightarrow S$ is one of types (iii) or (iv), with $b = \infty$ in the case of type (iv). Observe that because S is complete with respect to q , γ can be extended to $[a, b]$ in the first case and $[a, \infty)$ in the second. We call the images of these paths *separatrices*. A *singular leaf* of \mathcal{F} is an injective locally isometric immersion γ of \mathbb{R} or a circle of some circumference which is a concatenation of separatrices with the following property. After possibly reparametrizing to change the orientation, every time γ passes through a cone point (at the end of one separatrix and the beginning of the next), there is always a cone angle π on the right side of γ . Said differently, recall that all the cone points are k -prong singularities, and so we are requiring that when γ comes in to the singularity, it goes out along the right-most prong. Of course, reversing the orientation we have all left turns. A *leaf* of \mathcal{F} is either a singular leaf or a nonsingular leaf of \mathcal{F} .

3.2. Geometric intersection number I: measured foliations

Given a measured foliation (\mathcal{F}, μ) and $\alpha \in \mathcal{S}$, we define the geometric intersection number of α with μ , or with (\mathcal{F}, μ) , to be

$$\iota(\alpha, \mu) = \iota(\alpha, (\mathcal{F}, \mu)) := \inf_{\alpha_0 \in \alpha} \int_{\alpha_0} d\mu.$$

The next proposition provides the connection with the geometric intersection number between curves.

Proposition 3.2: *Given any $\alpha \in \mathcal{S}$, there exists a measured singular foliation (\mathcal{F}, μ) such that for all $\gamma \in \mathcal{S}$ we have*

$$\iota(\gamma, \alpha) = \iota(\gamma, \mu).$$

Proof: Let $\beta \in \mathcal{S}$ be any curve (or multicurve) which fills with α , and perform the construction on the pair α, β described at the beginning of Section 3.1 above to produce a Euclidean cone metric. The foliation by geodesics parallel to one of the sides of the squares is a foliation by curves each representing the homotopy class α . Indeed, if we consider the union of the singular leaves we get a graph Γ in S , and the complement of this is an open cylinder C foliated by curves in the homotopy class α ; see Figure 4.

Now given any $\gamma \in \mathcal{S}$, any representative $\gamma_0 \in \gamma$ must cross from one end of C to the other at least $\iota(\gamma, \alpha)$ times. Each time it crosses C , there is a contribution of at least 1 to the variation, and hence we have

$$\iota(\gamma, \mu) \geq \iota(\gamma, \alpha).$$

To prove the other inequality, let γ_0 be any representative which intersects one of the leaves α_0 of \mathcal{F} representing α exactly $\iota(\gamma, \alpha)$ times transversely. We can now apply a homotopy to γ_0 “pushing it into Γ ”, except for the $\iota(\gamma, \alpha)$ arcs which cross directly from one side of C to the other, each giving a contribution of exactly 1 to the total variation. Call this representative γ_1 . Since the part of γ_1 that runs around inside Γ does not contribute to the total variation we have

$$\iota(\gamma, \mu) \leq \int_{\gamma_1} d\mu = \iota(\gamma, \alpha)$$

which completes the proof. \square

Remark 3.3: We note that the construction depends very much on the choice of curve β , so that (\mathcal{F}, μ) is not unique.

Analogous to the length function used for Teichmüller space, geometric intersection number for curves and for foliations determines two maps

$$\iota_* : \mathcal{S} \times \mathbb{R}_+ \rightarrow \mathbb{R}^S \quad \text{and} \quad \iota_* : \widetilde{\mathcal{MF}}(S) \rightarrow \mathbb{R}^S$$

given by

$$\iota_*(t, \beta) = \{t\iota(\alpha, \beta)\}_{\alpha \in \mathcal{S}} \quad \text{and} \quad \iota_*(\mathcal{F}, \mu) = \{\iota(\alpha, \mu)\}_{\alpha \in \mathcal{S}},$$

respectively.

As a consequence of Proposition 3.2 we have the following.

Corollary 3.4:

$$\iota_*(\mathbb{R}_+ \times \mathcal{S}) \subset \iota_*(\widetilde{\mathcal{MF}}) \subset \mathbb{R}^S.$$

Proof: We note that by scaling the metric in the proof of the previous proposition by $t > 0$, we scale the geometric intersection numbers $\iota(\alpha, \mu)$ by t as well. \square

It is not too hard to see that the map ι_* on $\mathbb{R}_+ \times \mathcal{S}$ is injective. However, the map ι_* on $\widetilde{\mathcal{MF}}$ is *not* injective. First of all, we can apply an isotopy to any measured foliation. In addition, applying a *Whitehead move* (see Figure 6) will change the isotopy class of singular foliation, but will not change the geometric intersection numbers. We define an equivalence relation on $\widetilde{\mathcal{MF}}(S)$ by declaring two foliations (\mathcal{F}, μ) and (\mathcal{F}', μ') to be *measure equivalent* if they differ by isotopy and a sequence of Whitehead moves. We let $\mathcal{MF}(S)$ denote the set of measure equivalence classes of measured foliations.




Fig. 6. Whitehead move splitting up/combining singularities.

The following important theorem is beyond the scope of these notes, but it justifies the name “measure equivalent”; see [16].

Theorem 3.5: *Two foliations $(\mathcal{F}, \mu), (\mathcal{F}', \mu') \in \widetilde{\mathcal{MF}}(S)$ are equivalent if and only if $\iota_*(\mu) = \iota_*(\mu')$. Said differently, the equivalence relation is precisely the one for which the map ι_* descends to an injection*

$$\iota_* : \mathcal{MF}(S) \rightarrow \mathbb{R}^{\mathcal{S}}.$$

Using this theorem we give $\mathcal{MF}(S)$ the topology making this injection into an embedding. For our purposes we could just define the equivalence relation to be the one on $\widetilde{\mathcal{MF}}(S)$ determined by the fibers of the map ι_* .

Next, we observe that $\mathcal{MF}(S)$ and $\mathcal{T}(S)$ in $\mathbb{R}^{\mathcal{S}}$ are disjoint.

Proposition 3.6: *The image of $\mathcal{MF}(S)$ and $\mathcal{T}(S)$ are disjoint in $\mathbb{R}^{\mathcal{S}}$,*

$$\iota_*(\mathcal{MF}(S)) \cap \ell_*(\mathcal{T}(S)) = \emptyset.$$

Proof: We note that for any $X \in \mathcal{T}(S)$, we have a positive lower bound on $\ell_X(\alpha)$ over all $\alpha \in \mathcal{S}$ (in fact, $\{\ell_X(\alpha)\}_{\alpha \in \mathcal{S}}$ is a closed discrete subset of \mathbb{R} not containing 0). To prove the proposition, it therefore suffices to prove

Claim. Given (\mathcal{F}, μ) there exists a sequence of simple closed curves $\alpha_n \subset \mathcal{S}$ so that as $n \rightarrow \infty$ we have

$$\iota(\alpha_n, \mu) \rightarrow 0.$$

Proof: If \mathcal{F} has a closed nonsingular leaf, the isotopy class of this leaf, call it α , obviously has $\iota(\alpha, \mu)$, and we let $\alpha_n = \alpha$ for all n . If \mathcal{F} has no closed leaves, then there is a leaf locally isometric to \mathbb{R} in \mathcal{F} . This must accumulate at some point x on S . For any n we can construct a curve near x built from these arcs together with short arcs transverse to \mathcal{F} and having variation at most $1/n$ with μ . See [16], as well as [12] for a similar construction. \square

This completes the proof of the proposition. \square

3.3. Geodesics in the q -metric

We now explain one benefit of defining measured foliations with respect to a Euclidean cone metric q . First observe that the geodesics in the q -metric have a nice local description. Namely, in the complement of the singularities, these paths must be straight Euclidean segments (since a local isometry to

the Euclidean plane makes them length minimizing paths). Furthermore, whenever a geodesic enters and exits a singularity, it makes an angle $\geq \pi$ on both sides (if the angle were less than π on one side, you can easily find a nearby path which is shorter). An important point for us is that the geodesics in \tilde{S} with the induced metric are unique.

Proposition 3.7: *For any Euclidean cone metric q on S with cone angles $\geq 2\pi$, the induced metric q on \tilde{S} has the property that any two points are connected by a unique (distance minimizing) geodesic.*

This is a consequence of the fact that the q -metric on \tilde{S} is CAT(0), or *nonpositively curved* in the sense of Alexandrov together with the Cartan-Hadamard Theorem; see [1] or [10]. Instead of defining the CAT(0) condition, we sketch an elementary proof of this proposition using the Gauss-Bonnet Theorem.

Proof: Suppose we have two geodesics between a pair of points z and w . By passing to sub-arcs of these two geodesics we can assume that the interiors of the geodesics are disjoint. So, these two geodesics form a bigon as in Figure 7, and we can double this bigon over its boundary to produce a sphere with a Euclidean cone metric. While the cone angles may not all be greater than 2π , one can check that the only place where this can fail is at the two endpoints z and w . Now, taking a sufficiently high degree branched cover of the sphere over these two points, we again get a sphere with a Euclidean cone metric, and all cone angles are now greater than 2π . The Gauss-Bonnet Theorem shows that this is impossible.




Fig. 7. An (impossible) pair of distinct geodesics connecting two points in the Euclidean cone metric on \tilde{S} .

□

Remark 3.8: In fact, this proof (together with the existence of geodesics) also shows that the two properties of geodesics preceding the statement of the proposition are in fact both necessary and sufficient to guarantee that a

path is a geodesic. This also follows from the fact that the metric is CAT(0) and so local geodesics are global geodesics; see [10].

For a Riemannian metric a pair of compact geodesic segments which are not part of a larger single geodesic must intersect in only finitely many points. However, because the q -metric is singular, two geodesics γ_1 and γ_2 in \tilde{S} can run together for some part of their length without being contained in a larger geodesic. If γ_1 and γ_2 are disjoint for a while, then come together (possibly only at a point), then become disjoint again, then there are two possibilities: either γ_1 exits γ_2 on the opposite side it started on, and we say that γ_1 *crosses* γ_2 , or it exits on the same side.




Fig. 8. Two geodesics γ_1, γ_2 in \tilde{S} . The left and center figures show γ_1 crossing γ_2 , whereas γ_1 and γ_2 do not cross in the right figure.

Returning to the foliation \mathcal{F} , we also observe that this can be pulled back to the universal covering to obtain a foliation $\tilde{\mathcal{F}}$. The leaves of \mathcal{F} , hence also $\tilde{\mathcal{F}}$, satisfy the local conditions described above for being geodesic and so are indeed geodesics. It follows that any such leaf is proper in the universal covering. We will describe an even stronger property that q -geodesics in \tilde{S} enjoy in Section 5, but for now, we use properness to sketch a proof of the following; see [17].

Proposition 3.9: *Let (\mathcal{F}, μ) be a measured foliation on S defined in terms of some Euclidean cone metric q . Given a homotopy class of curve α on S , let α^* denote the q -geodesic representative. Then*

$$\iota(\alpha, \mu) = \int_{\alpha^*} d\mu.$$

That is, the geodesic representative realizes the geometric intersection number.

Proof: For any representative α_0 of α , $\int_{\alpha_0} d\mu$ measures “how many leaves” α_0 crosses. Of course, α_0 will generally cross infinitely many leaves, but the transverse measure provides a way to make this precise. Pulling this picture

back to the universal covering, we can choose lifts $\tilde{\alpha}^*$ a biinfinite geodesic and $\tilde{\alpha}_0$ a biinfinite path in \tilde{S} , each invariant under an element $\delta \in \pi_1(S)$, hence remaining a bounded distance from each other. Observe that since $\tilde{\alpha}^*$ is a geodesic, it will not cross any leaf of $\tilde{\mathcal{F}}$ more than once: if $\tilde{\alpha}^*$ crossed a leaf γ of $\tilde{\mathcal{F}}$, left it, then came back to cross it again, we would have two distinct geodesics between a pair of points (the arc of $\tilde{\alpha}^*$ and the arc of γ), which is impossible. Furthermore, if $\tilde{\alpha}^*$ crosses some leaf γ of $\tilde{\mathcal{F}}$, then its two ends must diverge from γ (this is another fact about CAT(0) spaces), and since $\tilde{\alpha}_0$ remains a bounded distance from $\tilde{\alpha}^*$, we see that $\tilde{\alpha}_0$ must also cross γ . So, α_0 crosses at least as many leaves as α^* and hence

$$\int_{\alpha^*} d\mu \leq \int_{\alpha_0} d\mu$$

as required. \square

4. Measured laminations

An alternative to measured singular foliations that serve equally well in many situations (and sometimes better) are measured geodesic laminations. To define these, we first fix a hyperbolic metric ρ on S . By a *complete simple geodesic* on S (with respect to ρ) we mean the image of an injective geodesic defined on either a circle or the entire real line. A *geodesic lamination* on S is a closed set which is a union of pairwise disjoint, complete simple geodesics

$$\mathcal{L} = \bigcup_{x \in \mathcal{L}} \ell_x.$$

Here each ℓ_x is a complete simple geodesic called a *leaf* of \mathcal{L} . Appealing to the Poincaré-Hopf Index Theorem, we see that the \mathcal{L} cannot be the entire surface, and in fact, it is nowhere dense and has zero Lebesgue measure. From this one sees that it can only be expressed as a disjoint union of complete simple geodesics in one way. In particular, the closed set \mathcal{L} determines the decomposition into a disjoint union of simple geodesics. See [13, 34, 12, 7] for more on geodesic laminations.

A *transverse measure* on \mathcal{L} can be defined as in the case of foliations as a measure on paths transverse to \mathcal{L} , so that the support of the measure is precisely the intersection of the path with \mathcal{L} . Furthermore we require that these are invariant under homotopy preserving transverse intersection with the leaves of \mathcal{L} and behave appropriately under concatenation (compare this with the properties of the transverse measure on a foliation described in

Section 3.1). However, there is a more concrete way to describe a transverse measure which we now explain; see [34, 6, 30] for a more on what follows.

As usual write ρ for the pull back of ρ to \tilde{S} , which makes \tilde{S} isometric to the hyperbolic plane \mathbb{H}^2 . If we use the Poincaré disk model of \mathbb{H}^2 , then the isometry $\tilde{S} \rightarrow \mathbb{H}^2$ provides a compactification of \tilde{S} by a *circle at infinity* which we denote S_∞^1 so that the isometry extends to a homeomorphism $\tilde{S} \cup S_\infty^1 \rightarrow \overline{\mathbb{H}^2}$. Geodesics in \tilde{S} determine a pair of points on S_∞^1 , the limit points of the two ends of the geodesic, and any pair of distinct points determines a geodesic in this way. This allows us to identify the set of all (unoriented biinfinite) ρ -geodesics in \tilde{S} with the set of unordered pairs of distinct points in S_∞^1 . We denote this $G(\tilde{S})$ so that

$$G(\tilde{S}) = (S_\infty^1 \times S_\infty^1 - \Delta) / \sim \quad (4.1)$$

where Δ is the diagonal and $(x, y) \sim (y, x)$. This also topologizes the set of geodesics in a natural way. For example, a sequence of geodesics γ_n converges to γ if and only if they can be parametrized to converge uniformly on compact sets; or alternatively, if and only if γ_n converges to γ in the *Chabauty topology* (see [12, 3]) on the set of closed subsets of \tilde{S} . We leave the verification of these claims to the reader.

Remark 4.1: At first glance, the notation $G(\tilde{S})$ and S_∞^1 seem to be imprecise since the metric ρ was instrumental in defining these, yet it does not appear in the notation. In fact, if ρ and $\hat{\rho}$ are any two hyperbolic metrics on S and S_∞^1 and \hat{S}_∞^1 are the two circle used to compactify \tilde{S} , then in fact the identity map $\tilde{S} \rightarrow \tilde{S}$ extends to a homeomorphism between the two compactifications; see e.g. [13]. So, abusing notation we do not mention the metric ρ in the notation.

The action of $\pi_1(S)$ on \tilde{S} by covering transformations extends to $\tilde{S} \cup S_\infty^1$. This action determines an action of $\pi_1(S)$ on $G(\tilde{S})$ (either from the fact that the action is by isometries, or by appealing to Equation (4.1)). A geodesic lamination \mathcal{L} on S gives rise to a geodesic lamination $\tilde{\mathcal{L}} = p^{-1}(\mathcal{L})$ on \tilde{S} which is invariant under the action of $\pi_1(S)$. The lamination $\tilde{\mathcal{L}}$ determines a closed $\pi_1(S)$ -invariant subset of $G(\tilde{S})$, which we also call $\tilde{\mathcal{L}}$ and we define a *transverse measure* on \mathcal{L} to be a $\pi_1(S)$ -invariant Radon measure on $G(\tilde{S})$ whose support is precisely this set $\tilde{\mathcal{L}}$ (recall that a Radon measure is a Borel measure which measures compact sets finitely). Since the support of μ is required to be $\tilde{\mathcal{L}}$, we note that μ actually determines $\tilde{\mathcal{L}}$ and hence also \mathcal{L} . As such, we sometimes refer to μ as a measured lamination, without

specifically mentioning the underlying lamination \mathcal{L} . We refer to \mathcal{L} as the *supporting lamination of μ* .

Suppose μ is a transverse measure in this sense. To see how this is related to the more geometric notion of a transverse measure, observe that given any path transverse to \mathcal{L} we can lift this to a path in \tilde{S} transverse to $\tilde{\mathcal{L}}$, and any two such lifts differ by an element of $\pi_1(S)$. For any sufficiently small arc of the path (so that the arc intersects any geodesic of $\tilde{\mathcal{L}}$ at most once), we look at the set of geodesics in $\tilde{\mathcal{L}}$ which intersect this arc, and define the measure of this arc to the measure of this set of geodesics. Invariance under $\pi_1(S)$ implies that we obtain a well-defined measure on the original path in S supported on the intersection of the path with \mathcal{L} .

We now give the basic example of a measured geodesic lamination. If α is a simple closed geodesic on S , then this is a geodesic lamination with one leaf, and the preimage $p^{-1}(\alpha) = \tilde{\alpha} \subset \tilde{S}$ defines a closed *discrete* set of geodesics on which $\pi_1(S)$ acts transitively. Up to scaling there is exactly one $\pi_1(S)$ -invariant Radon measure on $G(\tilde{S})$ supported on $\tilde{\alpha}$, namely the counting measure, and we denote this μ_α . This is the measure which assigns to a subset of geodesics $E \subset G(\tilde{S})$ the number of geodesics in $\tilde{\alpha}$ which lie in E . The associated measure on a path in S counts the number of times α crosses the path.

Not every geodesic lamination admits a transverse measure since we have required the measure to have support exactly equal to the lamination. For example, if (\mathcal{L}, μ) is a measured geodesic lamination, then any isolated leaf of \mathcal{L} —a leaf with a neighborhood in S intersecting \mathcal{L} only in that leaf—must be a simple closed geodesic. In particular, a lamination with an infinite isolated leaf, for example a “spiraling leaf” (see e.g. [7]), does not admit a transverse measure. Given a measured geodesic lamination (\mathcal{L}, μ) , let $\mathcal{L}_0 \subset \mathcal{L}$ be the union of the simple closed geodesics. We observe that \mathcal{L}_0 is necessarily a finite set of geodesics (indeed, it contains at most $3g - 3$ simple closed geodesics), and the measure μ decomposes as $\mu_0 + \mu_1$ with μ_0 supported on \mathcal{L}_0 and μ_1 supported on the lamination $\mathcal{L}_1 = \mathcal{L} - \mathcal{L}_0$. Since \mathcal{L}_1 has no closed leaves, the transverse measure μ_1 has no atoms in $G(\tilde{S})$ —that is, no one-point subset of $G(\tilde{S})$ has positive measure. Equivalently, the measure on any arc transverse to \mathcal{L}_1 has no atoms.

We let $\mathcal{ML}(S)$ denote the set of all measured geodesic laminations on S , which we can think of as a subset of the space of all measures on $G(\tilde{S})$. This latter space is given the weak* topology, and $\mathcal{ML}(S)$ is given the subspace topology. From the construction of a measured lamination μ_α from a simple closed curve α , we obtain an injection $\mathbb{R}_+ \times \mathcal{S}(S) \rightarrow \mathcal{ML}(S)$, by $(t, \alpha) \mapsto$

$t\mu_\alpha$. We note that $\mathcal{ML}(S)$ depends on the initial hyperbolic structure on S . However, as in Remark 4.1 we see that for any two hyperbolic metrics, the spaces so constructed are homeomorphic by a homeomorphism extending the inclusion of $\mathbb{R}_+ \times \mathcal{S}(S)$.

4.1. Geometric intersection number II: measured laminations

Given $\alpha \in \mathcal{S}(S)$ and $(\mathcal{L}, \mu) \in \mathcal{ML}(S)$, there is a geometric intersection number

$$\iota(\alpha, \mu) = \iota(\alpha, (\mathcal{L}, \mu)) = \inf_{\alpha_0 \in \alpha} \int_{\alpha_0 \in \alpha} d\mu_{\alpha_0}$$

The infimum is realized by the geodesic representative of α (which we denote by α). This is similar to the situation for measured foliations described in Section 3.3, and is perhaps easier to visualize since we have concrete models for $(\tilde{S}, \rho) \cong \mathbb{H}^2$. As before, the point is that any other representative of α must cross the leaves of \mathcal{L} at least as much as the geodesic representative, and hence the total variation is at least as large.

As with measured foliations, the geometric intersection number determines a map

$$\iota_* : \mathcal{ML}(S) \rightarrow \mathbb{R}^{\mathcal{S}}$$

by $\iota_*(\mathcal{L}, \mu) = \{\iota(\alpha, \mu)\}_{\alpha \in \mathcal{S}}$.

Though we will not use this below, we note that the geometric intersection number has the following nice extension due to Thurston. Recall that μ_α is the measured lamination determined by $\alpha \in \mathcal{S}(S)$.

Theorem 4.2: *There is a continuous, symmetric, bilinear function*

$$\iota : \mathcal{ML}(S) \times \mathcal{ML}(S) \rightarrow \mathbb{R}$$

such that for all $\alpha, \beta \in \mathcal{S}(S)$ and $t, s \in \mathbb{R}_+$

$$\iota(t\mu_\alpha, s\mu_\beta) = tst\iota(\alpha, \beta).$$

Bilinearity means that for all t, s, λ, μ, ν we have $\iota(t\mu, \lambda) = t\iota(\mu, \lambda) = \iota(\mu, t\lambda)$ and $\iota(\mu + \nu, \lambda) = \iota(\mu, \lambda) + \iota(\nu, \lambda)$, when $\mu + \nu$ is again a measured lamination. In fact, $\mathcal{ML}(S)$ is a subspace of infinite codimension in the space of all $\pi_1(S)$ -invariant Radon measures on $G(\tilde{S})$, called the space of *geodesic currents*, and Bonahon proved that ι has a continuous, symmetric, bilinear extension to this space as well; see [5, 6] for this and more on

currents. One can describe $\iota(\mu, \lambda)$ as the total mass of $\mu \times \lambda$ in an appropriate sense. For example, we can think of $\mu \times \lambda$ as defining a measure on S supported on the transverse intersections of the supporting laminations of μ and λ , and then $\iota(\mu, \lambda)$ is the total mass of this measure.

5. Measured foliations and measured laminations

We have seen that $\mathbb{R}_+ \times \mathcal{S}(S)$ can be naturally injected into both $\mathcal{MF}(S)$ as well as $\mathcal{ML}(S)$. The next theorem shows that these spaces are in some sense the same.

Theorem 5.1: *There is an \mathbb{R}_+ -invariant homeomorphism $\text{Str} : \mathcal{MF}(S) \rightarrow \mathcal{ML}(S)$ which is the identity on the image of $\mathbb{R}_+ \times \mathcal{S}(S)$ in each.*

One proof of this goes through an intermediate set of combinatorial objects called *train tracks*. See [30] and [11] for more on this.

We explain the idea for Theorem 5.1 here without mention of train tracks, leaving most of the details to the reader; see also [13, 26] for more on this approach.

To begin, recall that if (X, q) and (Y, ρ) are metric spaces, then given non-negative numbers K and C , a map $f : X \rightarrow Y$ is called a (K, C) -quasi-isometry if

$$\frac{1}{K}(\rho(x, y) - C) \leq q(x, y) \leq K\rho(x, y) + C$$

for all $x \in X$ and $y \in Y$, and if there exists $A > 0$ so that the A -neighborhood of $f(X)$ is the entire range, $Y = N_A(f(X))$.

We will need the following.

Proposition 5.2: *Given any q and ρ as above, there exists $K \geq 1$ so that the identity $(\tilde{S}, q) \rightarrow (\tilde{S}, \rho)$ is a $(K, 1)$ -quasi-isometry.*

Proof: This is really a special case of a much more general result, often called the Švarc-Milnor Lemma; see e.g. [10]. The basic idea is as follows. Since S is compact and since ρ and q determine the same topology, there exists $0 < \epsilon \leq \epsilon' \leq 1$ so that

$$B_\epsilon^\rho(x) \subset B_{\epsilon'}^q(x) \quad \text{and} \quad B_\epsilon^q(x) \subset B_{\epsilon'}^\rho(x) \quad (5.1)$$

for all $x \in \tilde{S}$.

Now connect any two points $z, w \in \tilde{S}$ by a ρ -geodesic and subdivide it into at most $1 + (\rho(z, w)/\epsilon)$ intervals of ρ -length at most ϵ . It follows that

the q -distance between consecutive endpoints of the subdivided intervals is at most ϵ' , and hence

$$q(z, w) \leq \epsilon' \left(1 + \frac{\rho(z, w)}{\epsilon} \right) \leq 1 + \frac{\epsilon' \rho(z, w)}{\epsilon}.$$

Setting $K = \epsilon'/\epsilon$ this proves one inequality and a symmetric argument proves the other. \square

From this it follows that any geodesic $\gamma : \mathbb{R} \rightarrow \tilde{S}$ in the q -metric is also a *quasi-geodesic* in the ρ -metric. That is, there exists constants $K, C > 0$ so that

$$\frac{1}{K}(|t - s| - C) \leq \rho(\gamma(t), \gamma(s)) \leq K|t - s| + C.$$

Since γ is a q -geodesic, we have $|t - s| = q(\gamma(s), \gamma(t))$, and so this is immediate from Proposition 5.2 for some K and $C = 1$. The importance of this lies in its application to the next fact.

Proposition 5.3: *There exists $A > 0$ so that for any q -geodesic γ in \tilde{S} , there is ρ -geodesic γ' so that*

$$\gamma \subset N_A^\rho(\gamma') \text{ and } \gamma' \subset N_A^\rho(\gamma).$$

If γ is compact (respectively, half-infinite), then γ' can be chosen to have the same endpoints (respectively, endpoint). With these choices γ' is unique. If γ is biinfinite, then γ' is unique.

The proposition is true for any quasi-geodesic where A depends only on the quasi-geodesic constants; see [10] for a proof of existence. The key fact that is used is that in the hyperbolic plane, closest point projection to a geodesic decreases arc-length by a factor which is exponential in the distance to the geodesic. The uniqueness statements follow from uniqueness of geodesics between any pair of distinct points in $\overline{\mathbb{H}^2}$. Given a q -geodesic γ , we call the corresponding ρ -geodesic the (ρ) -*straightening* of γ .

We now construct a map $\text{Str} : \mathcal{MF}(S) \rightarrow \mathcal{ML}(S)$, obtained by straightening all the leaves of a foliation to leaves of a lamination. Let (\mathcal{F}, μ) be a measured foliation defined via a Euclidean cone metric q , and let $\tilde{\mathcal{F}}$ be the foliation pulled back to \tilde{S} . For each leaf ℓ_q of $\tilde{\mathcal{F}}$, let ℓ_ρ denote the ρ -straightening. From this we see that ℓ_q of $\tilde{\mathcal{F}}$ has two endpoints on the circle at infinity S_∞^1 (this strengthens the properness of ℓ_q mentioned in Section 3.3).

Next, observe that no two leaves of $\tilde{\mathcal{F}}$ cross each other—they may however run together for a finite or infinite amount of time, but if and when they depart, they do so on the same side they started on: this is a consequence of our definition of leaves, specifically, the fact that they always “turn right”; see the end of Section 3.1. Therefore the endpoints on S_∞^1 of any two leaves do not *link* each other, meaning that given two pairs of points so that all four points are distinct, then the first pair lies in a single component of the complement of the other pair. From this it follows that the ρ -straightenings of distinct leaves are also either disjoint, or equal. Furthermore, if a sequence of these ρ -straightenings ℓ_{ρ_n} converge to a geodesic ℓ_ρ , one can check that (up to passing to a subsequence) the corresponding leaves ℓ_{q_n} also converge to a leaf ℓ_q which has ℓ_ρ as its ρ -straightening. Therefore $\tilde{\mathcal{L}}$, the union of all ρ -straightenings of leaves of $\tilde{\mathcal{F}}$, is a closed set of disjoint biinfinite geodesics, hence a geodesic lamination; see Figure 9.




Fig. 9. The left shows part of $\tilde{\mathcal{F}}$ and the right shows part of the straightening $\tilde{\mathcal{L}}$. The arc γ crossing the leaves of $\tilde{\mathcal{F}}$ is shown thickened.

Next observe that because $\tilde{\mathcal{F}}$ is invariant by $\pi_1(S)$, so is $\tilde{\mathcal{L}}$, and hence $\tilde{\mathcal{L}}$ descends to a geodesic lamination \mathcal{L} on S . We wish to impose a transverse measure λ on \mathcal{L} . For this, we let $[a, b] \times [c, d]$ denote any rectangle in $G(\tilde{S})$. Let $\Omega([a, b] \times [c, d])$ denote the (possibly empty) set of leaves of $\tilde{\mathcal{F}}$ which have one endpoint in $[a, b]$ and the other in $[c, d]$. Because no two leaves of $\tilde{\mathcal{F}}$ cross, $\Omega([a, b] \times [c, d])$ has a linear order, and we let ℓ_{q_0} and ℓ_{q_1} denote the two extreme leaves. Let γ be a q -geodesic arc between ℓ_{q_0} and ℓ_{q_1} , and observe that γ crosses each leaf in $\Omega([a, b] \times [c, d])$ exactly once (though it may run along one or more of these leaves); see Figure 9. Define

$$\lambda([a, b] \times [c, d]) = \mu(\gamma).$$

It is not hard to check that λ is a measure on $G(\tilde{S})$ with support $\tilde{\mathcal{L}}$, and hence (\mathcal{L}, λ) is a measured lamination. We define $\text{Str}(\mathcal{F}, \mu) = (\mathcal{L}, \lambda)$.

Generically each leaf of \mathcal{L} is the straightening of exactly one leaf of \mathcal{F} . The situation where this is not the case is when S contains an annulus foliated by parallel simple closed curves (for example, in the construction of a foliation from a simple closed curve). In this case, all the parallel leaves are straightened to a single simple closed curve.

We can also construct an inverse to Str . For this, we first observe that the straightening procedure has the effect of “inserting space” between all the leaves of a foliation to produce the lamination (this is a little bit like the operation of *Denjoying* a foliation, see [Example 4.14, [11]], in which more leaves are added to a foliation, except here we don’t add the leaves). To reverse this, we want to take a lamination and “collapse” all the space between the leaves. This is problematic when the lamination \mathcal{L} has simple closed curve components, as might be expected given the discussion of the previous paragraph. For this reason, we first modify our measured geodesic laminations as follows.

Given (\mathcal{L}, λ) , let $\mathcal{L}_0 = \beta_1 \cup \dots \cup \beta_k$ be the union of the simple closed geodesics in \mathcal{L} , so that β_i is assigned some t_i -times the transverse counting measure by λ for each $i = 1, \dots, k$. Replace each β_i with the ϵ -neighborhood of β_i for some very small ϵ , and foliate each neighborhood by curves parallel to the core curve, β_i . Finally, give each foliated annulus a transverse measure (with no atoms) so that the total measure of a transverse arc from one boundary component to the other is precisely t_i . Let \mathcal{L}' denote the resulting “partial foliation/lamination” and continue to denote the transverse measure by λ .

Remark 5.4: The point here is that the transverse measure assigned to any arc transverse to \mathcal{L} will have atoms at each point of intersection with one of the β_i . The procedure just described removes all the atoms.

We are now ready to begin describing the collapsing procedure. We make some choices to carry this out, which is not at all surprising: many different foliations will straighten to the same measured geodesic lamination (although all such measured foliations are equivalent to each other). Let $\mathcal{P} = \alpha_1, \dots, \alpha_{3g-3}$ denote any pants decomposition of S by geodesics for which \mathcal{L} transversely intersects each α_i . Let P_1, \dots, P_{2g-2} be the pants in S with geodesic boundary determined by \mathcal{P} . An exercise shows that any essential arc in a pair of pants is isotopic to one of the six shown in Figure 10; see for example [30]. By choosing the $\epsilon > 0$ used in the construction of

\mathcal{L}' from \mathcal{L} sufficiently small, we can assume that each leaf of \mathcal{L}' is transverse to \mathcal{P} , and that each arc of intersection of a leaf of \mathcal{L}' with each P_j is isotopic to one of these six.




Fig. 10. The six isotopy classes of arcs in a pair of pants.

For each $j = 1, \dots, 2g - 2$ consider the intersection of \mathcal{L}' with P_j . In P_j there are two or three rectangles which contain all arcs of intersection of leaves of \mathcal{L}' with P_j (or briefly, arcs of $\mathcal{L}' \cap P_j$), and up to homeomorphism, these rectangles must be as shown in one of the two pictures of Figure 11. We can choose these rectangles to be bounded on two sides by arcs of $\mathcal{L}' \cap P_j$ and on the other two sides by arcs in the boundary of P_j . Furthermore, all arcs of $\mathcal{L}' \cap P_j$ run between the latter two sides.




Fig. 11. All arcs of intersection of $\mathcal{L}' \cap P_j$ must be contained in the shaded rectangles.

Now we describe the collapsing procedure in each pair of pants. Fix one such $P = P_j$ and assume that we are in situation on the left in Figure 11, say, and that each rectangle contains a nonempty set of arcs of $\mathcal{L}' \cap R$.

We will explain how to modify this construction in the general case. Let R denote any one of the three rectangles in P .

We arbitrarily pick Euclidean coordinates on R providing a homeomorphism $G_R : R \rightarrow [0, 1] \times [0, 1]$ so that each arc of $\mathcal{L}' \cap R$ is of the form $[0, 1] \times \{y\}$, for some $y \in [0, 1]$. The transverse measure λ determines a measure (which we also denote λ) on the second factor $[0, 1]$. If the total variation of $[0, 1]$ with respect to this measure is r , we define a map

$$F_\lambda : [0, 1] \times [0, 1] \rightarrow [0, 1] \times [0, r]$$

by

$$F_\lambda(x, y) = \left(x, \int_0^y d\lambda \right).$$

Because λ has no atoms, it follows that this defines a continuous, surjective mapping. Observe that every arc $[0, 1] \times \{y\} \subset [0, 1] \times [0, r]$ is the F_λ -image of an arc of $\mathcal{L}' \cap R$. Composing we obtain a map $H = F_\lambda \circ G_R : R \rightarrow [0, 1] \times [0, r]$.

We do this for each of the three rectangles producing maps

$$H_i : R_i \rightarrow [0, 1] \times [0, r_i]$$

for $i = 1, 2, 3$. The three Euclidean rectangles $\{[0, 1] \times [0, r_i]\}_{i=1}^3$ can be glued together along two tripods *by isometries* as in Figure 12 and the result is a pair of pants \hat{P} with a Euclidean cone metric and a measured singular foliation (\mathcal{F}_P, μ_P) by geodesics. It is straightforward to construct a quotient map $H : P \rightarrow \hat{P}$ which is precisely H_i on each R_i . This map H sends each arc of $\mathcal{L}' \cap P$ to a leaf of the \mathcal{F}_P .




Fig. 12. The collapsing map H on a pair of pants. The image of the pants is a pair of pants expressed as the union of foliated Euclidean rectangles.

There are two other possibilities: there may be no arcs of $\mathcal{L}' \cap P$ in one of the rectangles or the rectangles may be as in Figure 11 on the right. For each of these situations we can construct a quotient map $H : P \rightarrow \hat{P}$ in a similar fashion with minor modifications. For example, in the case that all arcs are contained in just two rectangles, we glue these rectangles together along a single 4-prong graph (which can be seen as the degenerate case of the one we considered in more detail).

Because the measure λ was used to define the Euclidean metric on the arcs which met the boundary of the pants, it is straightforward to check that the resulting pants $\hat{P}_1, \dots, \hat{P}_{2g-1}$ can be glued together by isometries to produce a surface \hat{S} with a Euclidean cone metric and singular measured foliation (\mathcal{F}, μ) . Moreover, this gluing can be done so that the maps $H : P_j \rightarrow \hat{P}_j$ are the restrictions of a single map

$$H : S \rightarrow \hat{S}$$

which sends leaves of \mathcal{L}' to leaves of \mathcal{F} .

The map H is a homotopy equivalence, and as such is homotopic to a homeomorphism $\hat{H} : S \rightarrow \hat{S}$: this is straightforward from the construction here, but also follows from the Dehn–Nielsen–Baer Theorem; see [15]. Now we pull back the Euclidean cone metric and foliation to one on S by \hat{H} and we continue to denote it (\mathcal{F}, μ) . Then one can check that $\text{Str}(\mathcal{F}, \mu) = (\mathcal{L}, \lambda)$ so that at least this is a one-sided inverse. To prove that straightening, then collapsing gives an equivalent measured foliation is beyond the scope of these notes since we have not developed any of the tools for studying Whitehead equivalence; see [16] for a detailed treatment of Whitehead equivalence.

6. Dehn–Thurston coordinates

From now on we allow ourselves the freedom of passing back and forth between elements of $\mathcal{ML}(S)$ and $\mathcal{MF}(S)$ via Theorem 5.1. We now refer to an element of either space simply as μ when this is convenient. We will either explicitly state which we are referring to, or else let the context determine the meaning.

Our next goal is to describe the topology of the space $\mathcal{ML}(S)$ by constructing coordinates. This is a version of the Dehn–Thurston coordinates, using as much hyperbolic geometry as possible to simplify the discussion. See [30] for a more thorough discussion of Dehn–Thurston coordinates.

As with Fenchel–Nielsen coordinates, we begin with a pants decomposition $\mathcal{P} = \{\alpha_1, \dots, \alpha_{3g-g}\}$ for S . We choose a hyperbolic metric ρ on S so

that the lengths of the α_i are all small and equal. The Collar Lemma [23, 19] implies that by taking the α_i sufficiently small, we can find very wide collar neighborhoods of each of the α_i , so that the boundaries of these collars have length $1/2$, say. Note that this is not the geodesic length, but the actual length of the boundary curve. Now let C_1, \dots, C_{3g-3} be these collars and let P_1, \dots, P_{2g-2} be the pants which are the complements of these collars. We emphasize that the P_i are not taken to have geodesic boundary as is usually the case, but instead the pants share their boundaries with the collars.

We know that any complete simple geodesic ℓ on S will intersect any P_i in an arc which is isotopic to one of the 6 pictured in Figure 10. By choosing the length of the α_i sufficiently small, we can assume that if two such geodesics intersect the pants in isotopic arcs, then these arcs are very close to each other. More precisely, we observe the following.

Lemma 6.1: *Given $\epsilon > 0$ there is a $\delta > 0$ with the following property. Suppose we choose each α_i to have length $\delta_0 < \delta$ with the boundary of each P_j of length $1/2$. Then if ℓ_1 and ℓ_2 are two complete simple geodesics in S which intersect some P_j in isotopic arcs γ_1, γ_2 , then γ_1 and γ_2 are within ϵ of each other.*

Proof: For each small $\delta > 0$ look at a hyperbolic metric ρ_δ where the α_i all have length δ (the twisting parameter for the metrics is irrelevant when δ is small). Now suppose ℓ is a simple geodesic intersecting one of the pants P in an arc γ . Suppose the two boundary components of P are $\partial_0 P$ and $\partial_1 P$ and the corresponding geodesics in the pants decomposition are α_0 and α_1 , respectively. All of this can be pulled back to the universal covering to obtain geodesics $\tilde{\alpha}_0, \tilde{\alpha}_1, \tilde{\ell}$ and arcs $\tilde{\partial}_0 P, \tilde{\partial}_1 P$ as shown in Figure 13 in the disk model. These arcs $\tilde{\partial}_i P$ are equidistant lines to $\tilde{\alpha}_i$, for $i = 0, 1$.

Because we have chosen ∂P to have all components of length $1/2$, the distance between $\tilde{\partial}_0 P$ and $\tilde{\partial}_1 P$ is bounded above, independent of δ . It is now clear that if δ is chosen very small, the Euclidean diameter of the $\tilde{\alpha}_i$ will be very small, and any two geodesics $\tilde{\ell}_0, \tilde{\ell}_1$ which intersect both $\tilde{\alpha}_0, \tilde{\alpha}_1$ will be very close to each other between $\tilde{\partial}_0 P$ and $\tilde{\partial}_1 P$. \square

From this we can assume that the α_i have been chosen so small that every time a complete simple geodesic in S enters (or exits) any pants P_j through a component of ∂P_j it does so through one of four very small “windows” in that boundary component—which window it goes through depends on which boundary component it exits through. See Figure 14.




Fig. 13. Left: the lifted picture. Right: any two geodesics $\tilde{\ell}_0, \tilde{\ell}_1$ which intersect $\tilde{\alpha}_0$ and $\tilde{\alpha}_1$ must be close between $\partial_0 \widetilde{P}$ and $\partial_1 \widetilde{P}$.

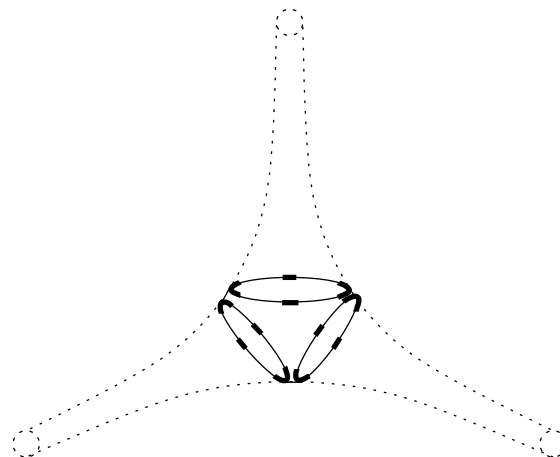


Fig. 14. Pants P_j and part of the adjacent collars (dotted). The dark lines are the windows through which all simple geodesics intersecting P_j must cross ∂P_j .

For each collar C_j , we choose a geodesic arc c_j from one boundary component to the other, with endpoints *outside* the windows. If (\mathcal{L}, μ) is any measured geodesic lamination, observe that \mathcal{L} will intersect each of these arcs c_j transversely (possibly not at all), and so determines a number $t_j(\mu)$ for each j :

$$t_j(\mu) = \int_{c_j} d\mu.$$

This measures how much (\mathcal{L}, μ) crosses c_j , and relative to $\iota(\alpha_i, \mu)$, this measures “how much (\mathcal{L}, μ) twists” inside of C_j . Moreover, if no α_i is a component of \mathcal{L} , then the orientation on the surface allows us to count this with a sign: as the arcs run from one side of C_j to the other, they either all cross left to right across c_j , or right to left. Call the former positive and the latter negative, and hence $t_j(\mu)$ takes values in \mathbb{R} , instead of just $\mathbb{R}_{\geq 0}$.

We also have the intersection number of any μ with α_i . Write this as

$$s_i(\mu) = i(\mu, \alpha_i)$$

for each $i = 1, \dots, 3g - 3$.

Let $\mathcal{ML}_{\mathcal{P}}(S)$ denote the set of measured geodesic laminations (\mathcal{L}, μ) for which no α_i is contained in \mathcal{L} . We have a map

$$\Phi_{\mathcal{P}} : \mathcal{ML}_{\mathcal{P}}(S) \rightarrow (\mathbb{R}_{\geq 0} \times \mathbb{R})^{3g-3}$$

given by

$$\Phi_{\mathcal{P}}(\mu) = \{(s_i(\mu), t_i(\mu))\}_{i=1}^{3g-3}.$$

We can extend $\Phi_{\mathcal{P}}$ to all of $\mathcal{ML}(S)$ as follows. First observe that if α_i is a component of \mathcal{L} , then (\mathcal{L}, μ) has

$$s_i(\mu) = i(\mu, \alpha_i) = 0.$$

We can still define $t_i(\mu)$ so that

$$|t_i(\mu)| = \int_{c_i} d\mu$$

but now $\mathcal{L} \cap C_i = \alpha_i$ so there is no obvious choice of sign for $t_i(\mu)$. Therefore, we simply give it *both signs*, and so to obtain a well-defined map, we must then modify the range. Namely, in each factor

$$\mathbb{R}_{\geq 0} \times \mathbb{R}$$

we identify $(0, t) \sim (0, -t)$. Note that the result is the same as the cone on \mathbb{RP}^1 which we denote $C\mathbb{RP}^1$, and this is homeomorphic to the plane: $C\mathbb{RP}^1 \cong \mathbb{R}^2$. We have therefore defined a map

$$\Phi_{\mathcal{P}} : \mathcal{ML}(S) \rightarrow (C\mathbb{RP}^1)^{3g-3} \cong \mathbb{R}^{6g-6}.$$

Theorem 6.2: *The map $\Phi_{\mathcal{P}}$ is a homeomorphism onto the complement of 0 (= the cone point), and furthermore $\Phi_{\mathcal{P}}(t\mu) = t\Phi_{\mathcal{P}}(\mu)$ for all $t > 0$, $\mu \in \mathcal{ML}(S)$. Consequently, $\mathbb{P}\mathcal{ML}(S) \cong S^{6g-7}$.*

We will just explain the idea, rather than carrying out the proof. It is not too hard to see that $\Phi_{\mathcal{P}}$ is continuous: the coordinates can all be described in terms of integrals of compactly supported continuous functions on $G(\tilde{S})$. To prove that $\Phi_{\mathcal{P}}$ is a bijection, we construct an inverse by building a “partial measured foliation” from foliated rectangles glued together according to the data specified by the coordinates $\{(s_i, t_i)\}_{i=1}^{3g-3}$, and then “straightening” this to a measured geodesic lamination (compare with Section 5). Finally, the map is proper by homogeneity, and hence a homeomorphism.

From this it is not too difficult to see that the map

$$\iota_* : \mathcal{ML}(S) \rightarrow \mathbb{R}^S$$

defined in Section 4.1 is injective. This is similar to the proof of Theorem 2.1. Indeed, the same $9g - 9$ curves

$$\alpha_1, \dots, \alpha_{3g-3}, \beta_1, \dots, \beta_{3g-3}, \tau_1(\beta_1), \dots, \tau_{3g-3}(\beta_{3g-3})$$

suffice for essentially the same reason.

As a consequence we have the following.

Corollary 6.3: *The subset of $\mathcal{ML}(S)$ consisting of (\mathcal{L}, μ) , where \mathcal{L} is a finite union of simple closed geodesics is dense in $\mathcal{ML}(S)$.*

Proof: If $\Phi_{\mathcal{P}}(\mathcal{L}, \mu)$ has all integer coordinates, then \mathcal{L} is a finite union of simple closed geodesics. The set of positive real multiples of integral points in $(C\mathbb{RP}^1)^{3g-3}$ is dense, hence the preimage of this set by $\Phi_{\mathcal{P}}$ is dense in $\mathcal{ML}(S)$. \square

With more work, one can see that $\mathbb{R}_+ \times \mathcal{S}(S)$ is dense in $\mathcal{ML}(S)$; see [16, 30].

7. Degenerations of hyperbolic structures

In this final section we explain Thurston’s Theorem which provides a compactification of Teichmüller space. We primarily follow the discussion in [16] in what follows.

Recall that by Proposition 2.4, the map

$$\mathbb{P}\ell_* : \mathcal{T}(S) \rightarrow \mathbb{P}\mathbb{R}^S$$

is injective. Similarly, Theorem 3.5 implies

$$\mathbb{P}\iota_* : \mathbb{P}\mathcal{MF}(S) \rightarrow \mathbb{P}\mathbb{R}^S$$

is injective. We identify $\mathcal{T}(S)$ and $\mathbb{P}\mathcal{MF}(S)$ with their images in $\mathbb{P}\mathbb{R}^S$. According to Proposition 3.6 these are disjoint in $\mathbb{P}\mathbb{R}^S$. We let $\overline{\mathcal{T}}(S)$ denote the closure of $\mathcal{T}(S)$ in $\mathbb{P}\mathbb{R}^S$.

Theorem 7.1: *The closure of $\mathcal{T}(S)$ in \mathbb{R}^S is given by*

$$\overline{\mathcal{T}}(S) = \mathcal{T}(S) \cup \mathbb{P}\mathcal{MF}(S)$$

and $(\overline{\mathcal{T}}(S), \mathbb{P}\mathcal{MF}(S)) \cong (\overline{B}^{6g-6}, S^{6g-7})$.

Here \overline{B}^{6g-6} is the closed ball in \mathbb{R}^{6g-6} with boundary the sphere S^{6g-7} . By Theorem 6.2 and Theorem 5.1 we already know that $\mathbb{P}\mathcal{MF}(S) \cong S^{6g-7}$.

According to this theorem, any sequence of hyperbolic structures X_n which exits every compact set has a subsequence (which we also denote X_n) with the following property. There exists $\mu \in \mathcal{ML}(S)$ and a sequence of numbers t_n so that for every $\alpha \in \mathcal{S}(S)$ we have

$$t_n \ell_{X_n}(\alpha) \rightarrow \iota(\mu, \alpha)$$

as $n \rightarrow \infty$. Alternatively we see that for any *two* curves α, β so that $\iota(\mu, \alpha) \neq 0$ we have

$$\frac{\ell_{X_n}(\beta)}{\ell_{X_n}(\alpha)} \rightarrow \frac{\iota(\mu, \beta)}{\iota(\mu, \alpha)}$$

as $n \rightarrow \infty$. Thus, the theorem says that, up to scaling, hyperbolic lengths are approximated by geometric intersection numbers with laminations/foliations.

7.1. The fundamental lemma

Thurston's proof of the theorem is based on making the last sentence of the previous section quantitative. Specifically, given a point of Teichmüller space X we wish to construct a measured foliation (\mathcal{F}_X, μ_X) so that lengths in X are well-approximated by intersection numbers with μ_X . To describe this construction, we let $\mathcal{P} = \{\alpha_1, \dots, \alpha_{3g-g}\}$ denote a pants decomposition of S . The similarity in the Dehn-Thurston coordinates on $\mathcal{ML}(S) = \mathcal{MF}(S)$ and Fenchel-Nielsen coordinates on $\mathcal{T}(S)$ provides the inspiration for the following construction.

Fix attention on one of the pairs of pants P and write $\alpha_1, \alpha_2, \alpha_3$ for the boundary curves (which may or may not all be distinct in S). For each $i = 1, 2, 3$ let $\ell_i = \ell_X(\alpha_i)$. Then we say that the triple (ℓ_1, ℓ_2, ℓ_3) *satisfies the triangle inequality* if for each $i = 1, 2, 3$ we have

$$\ell_i \leq \ell_{i+1} + \ell_{i+2}$$

with indices taken *mod* 3. If this is true with strict inequality for each i we say that the triple *satisfies the strict triangle inequality*. If this does not hold for all i then there is a unique i so that with indices take *mod* 3

$$\ell_i > \ell_{i+1} + \ell_{i+2}.$$

Call this α_i the *offending boundary component*.

In each of these three cases we first construct a partial measured foliation. Suppose for example that (ℓ_1, ℓ_2, ℓ_3) satisfies the strict triangle inequality. In this case, we let $\gamma_{12}, \gamma_{13}, \gamma_{23}$ denote the embedded hyperbolic geodesic arcs in P so that γ_{ij} is orthogonal to α_i and α_j , for each $i \neq j$. Let $r_{ij} = (\ell_i + \ell_j - \ell_k)/4$, where $k \notin \{i, j\}$, and consider the neighborhoods

$$N_{ij} = N_{r_{ij}}(\gamma_{ij}).$$

Observe that each of these neighborhoods is a “rectangle” in P with two boundary components in ∂P . Any two of the rectangles meet in exactly two points, both of which lie in the boundary. Specifically, $N_{ij} \cap N_{ik}$ is precisely two points in α_i . Foliate N_{ij} by arcs equidistant to γ_{ij} . See Figure 15.

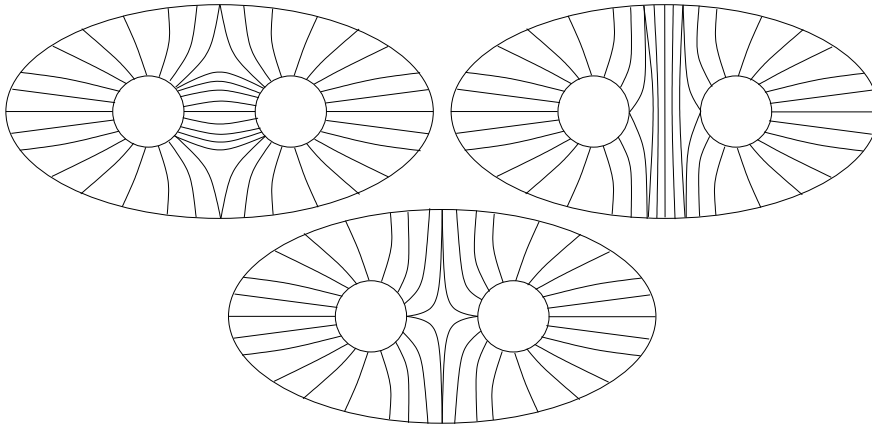


Fig. 15. The three types of partial measured foliations seen in a pair of pants. Upper left: strict triangle inequality. Upper right: no triangle inequality (with “outer boundary component” the offending component). Bottom: non-strict triangle inequality

When (ℓ_1, ℓ_2, ℓ_3) does not satisfy the triangle inequality, with offending boundary component α_1 , say, we can carry out a similar construction. In this case, we again consider the geodesic arcs γ_{12}, γ_{13} and neighborhoods

$$N_{1j} = N_{r_{1j}}(\gamma_{1j})$$

for $j = 2, 3$ where $r_{1j} = \ell_j/2$ here. Because $\ell_1 > \ell_2 + \ell_3$, we see that $\alpha_1 - (N_{12} \cup N_{13})$ is a pair of arcs, call them δ_1, δ_2 , each with length $(\ell_1 - \ell_2 - \ell_3)/2$. We let N_1 be the union of arcs running between δ_1 and δ_2 which are equidistant to γ_1 , the embedded geodesic arc which connects α_1 to itself orthogonally. We foliate N_1 by these arcs. See Figure 15. We caution the reader that γ_1 may not lie N_1 . The case of nonstrict triangle inequality is obviously the transition between the two, and if $\ell_1 = \ell_2 + \ell_3$, say, then we can carry out a similar construction using only the two geodesic arcs γ_{12} and γ_{13} and neighborhoods N_{12} and N_{13} ; see Figure 15.

We give this foliation a transverse measure which is just the distance between leaves. More precisely, given a tangent vector v let $\pi(v)$ denote the projection onto the orthogonal complement of the tangent space to the leaves (and $\pi(v) = 0$ if v is not based on a leaf), and define the transverse measure of a curve $\gamma : [0, 1] \rightarrow S$ by

$$\int_0^1 |\pi(\gamma'(t))| dt$$

where $|v|$ is the norm of the vector v using the hyperbolic metric ρ_X .

This produces a partial measured foliation in each pair of pants P_i for which the leaves meet the boundary orthogonally. From this and the definition of the transverse measure, it is straightforward to check that this is the restriction to each P_i of a partial measured foliation on the entire surface. Denote this $(\widehat{\mathcal{F}}_X, \widehat{\mu}_X)$.

This is a *partial* measured foliation because it is only defined on the complement of a collection of triangles and/or quadrilaterals (the latter occurring in the case of nonstrict triangle inequality). The triangles can be collapsed down to tripods (and 4-prong graphs for quadrilaterals) to produce a measured foliation (\mathcal{F}_X, μ_X) following the same idea as in the construction in the proof of Theorem 5.1. Here we build a Euclidean cone metric g_X from Euclidean rectangles, foliations on the rectangles as in the proof of Theorem 5.1, and a collapsing map, all done one pair of pants at a time. For example, in the case that (ℓ_1, ℓ_2, ℓ_3) satisfies the triangle inequality, we map N_{ij} to the Euclidean rectangle $[0, 1] \times [0, 2r_{ij}]$ with horizontal foliation (recall that $2r_{ij}$ is the hyperbolic width of N_{ij} described above) so that the transverse measures are preserved. By an abuse of notation, we think of the collapsing map as a map from S to itself, $f : S \rightarrow S$. This is a homeomorphism on the interiors of the rectangles, takes leaves of $\widehat{\mathcal{F}}_X$ to leaves of \mathcal{F}_X , and preserves transverse measures.

The key ingredient in the proof of Theorem 7.1 is the following “Funda-

mental Lemma” of [16]. Given $\epsilon > 0$, let $V(\mathcal{P}, \epsilon)$ denote the set of hyperbolic structures in which each curve α_i of \mathcal{P} has length at least ϵ

$$V(\mathcal{P}, \epsilon) = \{X \in \mathcal{T}(S) \mid \ell_X(\alpha_i) \geq \epsilon \text{ for all } \alpha_i \in \mathcal{P}\}.$$

Lemma 7.2: *Given ϵ and any curve $\gamma \in \mathcal{S}$ there exists $C \geq 0$ so that for all $X \in V(\mathcal{P}, \epsilon)$ we have*

$$\iota(\gamma, \mu_X) \leq \ell_X(\gamma) \leq \iota(\gamma, \mu_X) + C$$

where (\mathcal{F}_X, μ_X) is the measured foliation constructed above.

Observe that for the pants curves α_i , we can take $C = 0$ by construction of μ_X .

Proof: The first inequality is easy, and does not require $X \in V(\mathcal{P}, \epsilon)$. Let $f : S \rightarrow S$ be the collapsing map as described above. Let γ denote the geodesic representative with respect to X . If $\pi(\gamma'(t))$ is as defined above, then we have

$$\iota(\gamma, \mu_X) \leq \int_{f(\gamma)} d\mu_X = \int_0^1 |\pi(\gamma'(t))| dt \leq \int_0^1 |\gamma'(t)| dt = \ell_X(\gamma)$$

as required.

The second inequality is a bit more subtle, but we sketch the idea. The first fact that we need is the following.

Claim. For all $\epsilon > 0$, there exists $K > 0$ with the following property. If $X \in V(\mathcal{P}, \epsilon)$ and P is one of the pants determined by \mathcal{P} , then the length of the leaves of the partial measured foliation in P are bounded above by K .

Proof: Let $\alpha_1, \alpha_2, \alpha_3$ be the boundary geodesics of P so that $\ell_1, \ell_2, \ell_3 \geq \epsilon$. The proof is slightly different in each of the three cases, but the idea is the same. For concreteness, we consider the case that (ℓ_1, ℓ_2, ℓ_3) satisfies the triangle inequality and we use the notation as above. Without loss of generality, suppose $\ell_1 \geq \ell_2 \geq \ell_3 \geq \epsilon$. Observe first that the length of any leaf in a rectangle N_{ij} is bounded above by the lengths of the boundary leaves. Furthermore, if we bound the length of the boundary leaves for two of the rectangles, then it is not hard to see that the third is also bounded.

We will bound the lengths of the boundary leaves of N_{12} and N_{13} . First recall that the transverse boundary arc of N_{ij} has length

$$2r_{ij} = \frac{\ell_i + \ell_j - \ell_k}{2}$$

where i, j, k are all distinct. From the triangle inequality, and our assumptions on the lengths, it follows that this is at least $\epsilon/2$ for N_{12} and N_{13} . Fix one of these rectangles N and suppose that the boundary length is L . A calculation in hyperbolic geometry tells us that the length of an arc of the boundary of an R -neighborhood of a geodesic running between two geodesics orthogonal to the original geodesic is $\cosh(R)$ times the length of the geodesic segment between the orthogonals; see Figure 16. From this, a very crude lower bound for the area of N is $\epsilon L / (2 \cosh(\epsilon/2))$ (obtained by mapping via a 1-Lipschitz map to a Euclidean rectangle). However, N is contained in P which has area 2π , and hence

$$L \leq \frac{4\pi \cosh(\epsilon/2)}{\epsilon}.$$

This is the required bound in this case. The other two cases are similar. \square




Fig. 16. Lengths of equidistant curves increase exponentially.

Recall that the foliation (\mathcal{F}_X, μ_X) was described in terms of a Euclidean cone metric q_X , built by making each of the rectangles in the pants into Euclidean rectangles and collapsing out the complementary regions. A simple but useful observation about the metric q_X on each pair of pants P is that it can be obtained as the union of three Euclidean cylinders, each of height $1/2$ (and circumference equal to the X -length of the given curve). The three annuli are glued together along a spine for P as shown in Figure

17 in the case that (ℓ_1, ℓ_2, ℓ_3) satisfies the triangle inequality. To see this, consider the foliation orthogonal to \mathcal{F}_X (compare Figure 12).




Fig. 17. Pants built from Euclidean rectangles can also be obtained by gluing Euclidean cylinders. The union of the separatrices forms the spine.

Let γ be any simple closed curve on S , and let γ_0 to be a q_X -geodesic representative of γ . Every time γ_0 enters and exits a pair of pants P , the description of geodesics from Section 3.3 tells us that it does so in a relatively simple way: first it crosses one of the annuli, traverses at most one of the arcs of the spine, then it crosses another annulus and exits P . Observe that when γ_0 crosses an annulus, it generally does this “diagonally”. We replace each such diagonal subpath of γ_0 with a path having the same total μ_X -variation that first traverses around the boundary of the pants, then follows a leaf of \mathcal{F}_X to the spine. Similarly, if γ_0 traverses an edge of the spine, we can modify this, also keeping the μ_X -variation fixed, so that it runs along a leaf of \mathcal{F}_X to the boundary of the pants, runs along the boundary of the pants, then back along a leaf of \mathcal{F}_X . Let γ_1 denote the resulting representative of the homotopy class γ . We have

$$\int_{\gamma_1} d\mu_X = \int_{\gamma_0} d\mu_X = \iota(\gamma, \mu_X).$$

On the other hand, γ_1 determines a representative γ_2 which does essentially the same thing in the hyperbolic metric: it traverses the boundary of the pants in the same fashion as γ_1 and traverses leaves of the partial foliation $\hat{\mathcal{F}}_X$ whenever γ_1 does.

Now we observe that for a fixed curve γ , the number of times γ_1 , and hence γ_2 , traverses any arc of a leaf is uniformly bounded by some number $D > 0$ (depending only on the number of times γ crosses the pants curves).

The hyperbolic length of γ_2 is the sum of the lengths of the arcs contained in the pants curves plus the lengths of the arcs of leaves traversed. By the claim above the hyperbolic length of each arc of a leaf traversed is at most K , and hence

$$\ell_X(\gamma) \leq \ell_X(\gamma_2) \leq \int_{\gamma_1} d\mu_X + KD = \iota(\gamma, \mu_X) + KD.$$

Setting $C = KD$ completes the proof. \square

From this we have the following corollary. To avoid confusion, we do not identify $\mathcal{T}(S)$ and $\mathcal{PMF}(S)$ with their images in \mathbb{PR}^S , but refer to their respective embeddings $\mathbb{P}\ell_*$ and $\mathbb{P}\iota_*$ explicitly.

Corollary 7.3: *Let $X_n \in V(\mathcal{P}, \epsilon)$ with $X_n \rightarrow \infty$ in $\mathcal{T}(S)$. Then*

$$\{\mathbb{P}\ell_*(X_n)\}_{n=1}^\infty \text{ and } \{\mathbb{P}\iota_*(\mu_{X_n})\}_{n=1}^\infty$$

either both converge or both diverge. Furthermore, if they both converge, then they have the same limit.

Proof: From Theorem 2.1 we know that there is some curve γ so that $\ell_{X_n}(\gamma) \rightarrow \infty$ as $n \rightarrow \infty$. Suppose now that $\mathbb{P}\ell_*(X_n)$ converges to $\mathbb{P}(\{x_\alpha\}_{\alpha \in \mathcal{S}})$. Let $\{t_n\}_{n=1}^\infty$ be such that $t_n \ell_*(X_n)$ converges to $\{x_\alpha\}_{\alpha \in \mathcal{S}}$. By definition then we have

$$\lim_{n \rightarrow \infty} t_n \ell_{X_n}(\alpha) = x_\alpha$$

for all $\alpha \in \mathcal{S}(S)$. Applying this to $\alpha = \gamma$ we see that $t_n \rightarrow 0$.

For each $\alpha \in \mathcal{S}$, let $C(\alpha)$ be the constant from Lemma 7.2. Then $t_n C(\alpha) \rightarrow 0$ and so by the Lemma 7.2 and the squeeze theorem we have

$$\lim_{n \rightarrow \infty} t_n \iota(\mu_{X_n}, \alpha) = x_\alpha$$

as required.

The argument is easily reversed to prove that if $\{\mathbb{P}\iota_*(\mu_{X_n})\}$ converges, then $\{\mathbb{P}\ell_*(X_n)\}$ does as well and the limits are the same. \square

We are now ready for the proof of Theorem 7.1.

Proof: [Proof of Theorem 7.1.] Given any sequence $X_n \rightarrow \infty$ in $\mathcal{T}(S)$, we can find a pants decomposition $\mathcal{P} = \{\alpha_1, \dots, \alpha_{3g-3}\}$ an $\epsilon > 0$ and a subsequence (also called $\{X_n\}$) so that $\{X_n\} \subset V(\mathcal{P}, \epsilon)$. To find such a pants decomposition, start with any pants decomposition and consider one

of the curves. If there is a subsequence where the length of that curve is bounded below, pick that subsequence and move on to the next curve in the pants decomposition. If there is no such subsequence, then the length of that curve is tending to zero, and we can replace it with a curve that transversely intersects it. This has length tending to infinity, and we can then move on to the next curve. We continue this until we arrive at the desired \mathcal{P} and subsequence $\{X_n\}$.

Since $\mathbb{P}\mathcal{MF}(S) \cong S^{6g-7}$ is compact, we can pass to a subsequence so that μ_{X_n} converges to some $\mu \in \mathbb{P}\mathcal{MF}(S)$. From Corollary 7.3 we see that X_n also converges to μ . Therefore $\overline{\mathcal{T}}(S) \subset \mathcal{T}(S) \cup \mathbb{P}\mathcal{MF}(S)$.

To see that we get equality, we observe that by using Fenchel-Nielsen coordinates and Dehn-Thurston coordinates, the map $X \mapsto \mu_X$ maps $V(\mathcal{P}, \epsilon)$ onto the subset $\mathcal{MF}_{(\mathcal{P}, \epsilon)}(S) \subset \mathcal{MF}(S)$ consisting of those μ with $\iota(\mu, \alpha_i) \geq \epsilon$ for all $\alpha_i \in \mathcal{P}$. In particular, given any $\mu \in \mathcal{MF}(S)$, there is a pants decomposition \mathcal{P} and $\epsilon > 0$ so that $n\mu$ is in $\mathcal{MF}_{(\mathcal{P}, \epsilon)}(S)$ for all n . That is, $n\mu = \mu_{X_n}$ for some $X_n \in V(\mathcal{P}, \epsilon)$. The sequence $n\mu$ is constant projectively, but clearly $X_n \rightarrow \infty$ in $\mathcal{T}(S)$ —the lengths of all pants curves tend to ∞ . By Corollary 7.3 we have $X_n \rightarrow \mu$ and so

$$\mathcal{T}(S) \cup \mathbb{P}\mathcal{MF}(S) \subset \overline{\mathcal{T}}(S)$$

and hence these sets are equal.

Let $\mathbb{P}\mathcal{MF}_{\mathcal{P}}(S)$ be the set of measured foliations with $\iota(\alpha_i, \mu) \neq 0$ for all $\alpha_i \in \mathcal{P}$. Clearly, $\mathbb{P}\mathcal{MF}_{\mathcal{P}}(S)$ is the image of $\mathcal{MF}_{(\mathcal{P}, \epsilon)}(S)$ for any $\epsilon > 0$ under the projectivization \mathbb{P} . To see that $\overline{\mathcal{T}}(S)$ is a manifold with boundary equal to $\mathbb{P}\mathcal{MF}(S)$, we set

$$U_{\mathcal{P}} = \mathcal{T}(S) \cup \mathbb{P}\mathcal{MF}_{\mathcal{P}}(S) \subset \mathbb{P}\mathbb{R}^S.$$

Being an open subset of $\mathbb{P}\mathcal{MF}(S) \cong S^{6g-7}$, we see that $\mathbb{P}\mathcal{MF}_{\mathcal{P}}(S)$ is homeomorphic to an open subset $W \subset \mathbb{R}^{6g-7}$. We define

$$\Psi_{\mathcal{P}} : U_{\mathcal{P}} \rightarrow \mathbb{P}\mathcal{MF}_{\mathcal{P}}(S) \times (0, \infty]$$

by

$$\Psi_{\mathcal{P}}(X) = (\mathbb{P}(\mu_X), \min_{\alpha_i \in \mathcal{P}}(\ell_X(\alpha_i)))$$

for $X \in \mathcal{T}(S)$, and for $\mathbb{P}(\mu) \in \mathbb{P}\mathcal{MF}_{\mathcal{P}}(S)$ we set $\Psi_{\mathcal{P}}(\mathbb{P}(\mu)) = (\mathbb{P}(\mu), \infty)$. One can check that this is a homeomorphism. Then since

$$\mathbb{P}\mathcal{MF}_{\mathcal{P}}(S) \times (0, \infty] \cong W \times [0, \infty)$$

and since $\overline{\mathcal{T}}(S)$ is covered by such sets, we see that $\overline{\mathcal{T}}(S)$ is a manifold with boundary $\mathbb{P}\mathcal{MF}(S) \cong S^{6g-7}$ (the second countability and Hausdorff conditions follow from the embedding into $\mathbb{P}\mathbb{R}^S$)

That $\overline{\mathcal{T}}(S) \cong \overline{B}^{6g-6}$ now follows from this, the fact that $\mathcal{T}(S) \cong \mathbb{R}^{6g-6}$, and purely topological considerations; see [16] for more details. \square

7.2. Final comments on Theorem 7.1.

7.2.1. Application to the mapping class group

The utility of the Thurston compactification $\overline{\mathcal{T}}(S)$ comes from the fact that this is a mapping class group invariant compactification: the action of the mapping class group $\text{Mod}(S)$ on $\mathcal{T}(S)$ (see Aramayona's Lectures [2]) extends to an action on $\overline{\mathcal{T}}(S)$. This is immediate from the naturality of the construction. Indeed, $\text{Mod}(S)$ acts on \mathcal{S} , and hence on $\mathbb{P}\mathbb{R}^{\mathcal{S}}$ by acting on the coordinates

$$\phi \cdot \{x_\alpha\}_{\alpha \in \mathcal{S}} = \{x_{\phi \cdot \alpha}\}_{\alpha \in \mathcal{S}}.$$

It is straightforward to check that the inclusion $\mathbb{P}\ell_*$ is equivariant with respect to the action of $\text{Mod}(S)$ (as is $\mathbb{P}\iota_*$). Therefore, as we have identified $\mathcal{T}(S)$ with its image, we see that the induced action on $\mathcal{T}(S)$ as a subspace of $\mathbb{P}\mathbb{R}^{\mathcal{S}}$ agrees with the usual action, and hence the action extends to the closure $\overline{\mathcal{T}}(S)$ as the restriction of the action on all of $\mathbb{P}\mathbb{R}^{\mathcal{S}}$. This is a fundamental ingredient in Thurston's proof of his classification Theorem for mapping classes stated in Aramayona's lectures, [2]. Indeed, using the Brouwer Fixed Point Theorem, the action on $\overline{\mathcal{T}}(S)$ of any element of $\text{Mod}(S)$ has a fixed point, and a careful analysis of this fixed point provides all the necessary information needed to classify the elements of $\text{Mod}(S)$. See [16] for Thurston's original proof, or [13, 4, 20, 15, 9] for other proofs.

7.2.2. Other approaches

There are now several alternative constructions of the Thurston compactification of $\mathcal{T}(S)$ of Theorem 7.1. One goes by way of actions on \mathbb{R} -trees, and is due to Skora [31] (see also [18, 29, 28]), building on the work of Morgan and Shalen [29, 27]. Another proof is due to Bonahon [6], obtained by embedding $\mathcal{T}(S)$ into the space of projective geodesic currents mentioned in Section 4.1, and then showing that the closure is precisely the image of $\mathcal{T}(S)$ together with $\mathbb{P}\mathcal{ML}(S)$. It is worth mentioning however that these constructions do not give the full strength of Theorem 7.1. More precisely,

these provide a construction of the compactification, but not a description of $\overline{T}(S)$ as a closed ball.

References

1. S. B. Alexander and R. L. Bishop. The Hadamard-Cartan theorem in locally convex metric spaces. *Enseign. Math.* (2) 36 (1990), no. 3-4, 309–320.
2. J. Aramayona. *Hyperbolic structures on surfaces*. These proceedings.
3. R. Benedetti & C. Petronio. *Lectures on hyperbolic geometry*. Universitext, Springer-Verlag, Berlin, 1992. xiv+330 pp.
4. L. Bers. An extremal problem for quasiconformal mappings and a theorem by Thurston. *Acta Math.* 141 (1978), no. 1-2, 73–98.
5. F. Bonahon. Bouts des variétés hyperboliques de dimension 3. *Ann. of Math.* (2), 124(1):71–158, 1986.
6. F. Bonahon. The geometry of Teichmüller space via geodesic currents. *Invent. Math.*, 92(1):139–162, 1988.
7. F. Bonahon. Geodesic laminations on surfaces. In *Laminations and foliations in dynamics, geometry and topology (Stony Brook, NY, 1998)*, volume 269 of *Contemp. Math.*, pages 1–37. Amer. Math. Soc., Providence, RI, 2001.
8. M. Bestvina, K. Bromberg, K. Fujiwara, J. Souto. Shearing coordinates and convexity of length functions. Preprint, 2009.
9. M. Bestvina and M. Handel. Train-tracks for surface homeomorphisms. *Topology* 34 (1995), no. 1, 109–140.
10. M. R. Bridson & A. Haefliger. *Metric spaces of non-positive curvature*. Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences], 319. Springer-Verlag, Berlin,
11. D. Calegari. *Foliations and the geometry of 3-manifolds*. Oxford Mathematical Monographs. Oxford University Press, Oxford, 2007. xiv+363 pp.
12. R. D. Canary, D. B. A. Epstein & P. L. Green. *Notes on notes of Thurston With a new foreword by Canary*. London Math. Soc. Lecture Note Ser., 328, Fundamentals of hyperbolic geometry: selected expositions, 1–115, Cambridge Univ. Press, Cambridge, 2006,
13. A. J. Casson & S. A. Bleiler. *Automorphisms of surfaces after Nielsen and Thurston*. London Mathematical Society Student Texts, 9. Cambridge University Press, Cambridge, 1988. iv+105 pp
14. M. P. do Carmo. *Riemannian geometry*. Mathematics: Theory & Applications. Birkhäuser Boston, Inc., Boston, MA, 1992
15. B. Farb & D. Margalit. *A primer on mapping class groups*. Princeton University Press, to appear.
16. A. Fathi, F. Laudenbach, and V. Poenaru. *Travaux de Thurston sur les surfaces*, volume 66 of *Astérisque*. Société Mathématique de France, Paris, 1979. Séminaire Orsay, With an English summary.
17. F. P. Gardiner. *Teichmüller theory and quadratic differentials*. Pure and Applied Mathematics (New York). John Wiley & Sons Inc., New York, 1987. A Wiley-Interscience Publication.

18. H. Gillet & P. B. Shalen. Dendrology of groups in low \mathbf{Q} -ranks. *J. Differential Geom.* 32 (1990), no. 3, 605–712.
19. N. Halpern, A proof of the collar lemma. *Bull. London Math. Soc.* **13** (1981), no. 2, 141–144.
20. M. Handel & W. P. Thurston, New proofs of some results of Nielsen, *Adv. in Math.* 56 (1985), no. 2, 173–191.
21. J. Hubbard & H. Masur. Quadratic differentials and foliations. *Acta Math.*, 142(3-4):221–274, 1979.
22. Y. Iwayoshi & M. Taniguchi. *An introduction to Teichmüller spaces*. Springer-Verlag, Tokyo, 1992. xiv+279 pp.
23. L. Keen, Collars on Riemann surfaces. *Discontinuous groups and Riemann surfaces (Proc. Conf., Univ. Maryland, College Park, Md., 1973)*, pp. 263–268. Ann. of Math. Studies, No. 79, Princeton Univ. Press, Princeton, N.J., 1974.
24. S. P. Kerckhoff. The Nielsen realization problem. *Ann. of Math. (2)* **117** (1983), no. 2, 235–265.
25. S. P. Kerckhoff. Earthquakes are analytic. *Comment. Math. Helv.* **60** (1985), no. 1, 17–30.
26. G. Levitt. Foliations and laminations on hyperbolic surfaces. *Topology*, 22(2):119–135, 1983.
27. J. W. Morgan. Group Actions on Trees and the Compactification of the Spaces of Classes of $SO(n; 1)$ -representations. *Topology*, 25 (1986) 1–34.
28. J. W. Morgan and J.-P. Otal. Relative growth rates of closed geodesics on a surface under varying hyperbolic structures. *Comment. Math. Helv.* 68 (1993), no. 2, 171–208.
29. J. W. Morgan and P. B. Shalen. Degenerations of Hyperbolic Structures, I: Valuations, Trees and Surfaces. *Ann. of Math. (2)* 120 (1984), no. 3, 401–476.
30. R. C. Penner and J. L. Harer. *Combinatorics of train tracks*, volume 125 of *Annals of Mathematics Studies*. Princeton University Press, Princeton, NJ, 1992.
31. R. K. Skora. Splittings of surfaces. *J. Amer. Math. Soc.* 9 (1996), no. 2, 605–616.
32. W. P. Thurston. *On the geometry and dynamics of diffeomorphisms of surfaces*, Bull. Am. Math. Soc., New Ser. **19** (1988), 417–431.
33. W. P. Thurston. *Three-dimensional geometry and topology. Vol. 1*. Princeton Mathematical Series, 35. Princeton University Press, Princeton, NJ, 1997.
34. W. P. Thurston. *The Geometry and Topology of Three-Manifolds*, Princeton University course notes, available at www.msri.org/publications/books/gt3m/ (1980).