

Introduction to Teichmüller theory

Lecture notes

Bram Petri

Version: February 11, 2025

Contents

Preface	5
Lecture 1. Reminder on surfaces	7
1.1. Preliminaries on surface topology	7
1.2. Riemann surfaces	9
1.3. The uniformization theorem and automorphism groups	13
1.4. Quotients of the three simply connected Riemann surfaces	14
Lecture 2. Quotients, metrics, conformal structures	17
2.1. More on quotients	17
2.2. Riemannian metrics and Riemann surfaces	21
2.3. Conformal structures	23
Lecture 3. The Teichmüller space of the torus	27
3.1. Riemann surface structures on the torus	27
3.2. The Teichmüller and moduli spaces of tori	28
3.3. \mathcal{T}_1 as a space of marked structures	30
3.4. Markings by diffeomorphisms	32
3.5. The Teichmüller space of Riemann surfaces of a given type	34
Lecture 4. Markings, mapping class groups and moduli spaces	35
4.1. Teichmüller space in terms of markings	35
4.2. The mapping class group	38
4.3. Moduli space	38
4.4. Elements and examples of mapping class groups	39
Lecture 5. Beltrami differentials and a topology on Teichmüller space	45
5.1. Wirtinger derivatives	45
5.2. Beltrami coefficients	45
5.3. Quasiconformal mappings	49
5.4. The Teichmüller metric	51
5.5. Grötsch's theorem	52
Lecture 6. Measured foliations and quadratic differentials	55
6.1. Measured foliations	55
6.2. Holomorphic quadratic differentials	59
Lecture 7. Quadratic differentials, Teichmüller mappings, hyperbolic geometry	63
7.1. Quadratic differentials	63
7.2. Teichmüller mappings	64

7.3.	Hyperbolic surfaces	65
7.4.	The universal cover of a hyperbolic surface with boundary	70
7.5.	Pairs of pants and gluing	71
Lecture 8.	Pants decompositions and annuli	73
8.1.	Non-compact pairs of pants	73
8.2.	The geometry of isometries	74
8.3.	Geodesics and conjugacy classes	75
8.4.	Every hyperbolic surface admits a pants decomposition	77
8.5.	Annuli	78
Lecture 9.	Fenchel–Nielsen coordinates	81
9.1.	More on annuli	81
9.2.	Fenchel–Nielsen coordinates	82
Lecture 10.	Teichmüller’s theorem and the Nielsen–Thurston classification	87
10.1.	Teichmüller’s theorem	87
10.2.	The Nielsen–Thurston classification on the torus	90
10.3.	Definitions and statement of the classification theorem	90
10.4.	An overview of the proof	91
10.5.	The space of measured foliations	92
10.6.	Compactifying Teichmüller space	94
Lecture 11.	Landmark results, part I	95
11.1.	Compactifying Teichmüller space, part II	95
11.2.	Classifying mapping classes	97
Lecture 12.	Landmark results, part II	101
12.1.	The Nielsen realization problem	101
12.2.	Thurston’s hyperbolization of mapping tori	103
12.3.	The number of short simple closed geodesics	106
	Bibliography	113

Preface

These are the lecture notes for a course called *Introduction to Teichmüller theory*, taught in January and February 2024 and 2025 in the master's program *M2 de Mathématiques fondamentales* at Sorbonne University.

There are many references on various aspects of Teichmüller theory, like [IT92, Bus10, GL00, Zor06, Hub06, FM12, Baa21, Wri15]. All of these treat a lot more material than what we will have time for in the course, whence the present notes. Most of the material presented here is adapted from these references.

LECTURE 1

Reminder on surfaces

The Teichmüller space of a surface S is the deformation space of complex structures on S and can also be seen as a space of hyperbolic metrics on S . The aim of this course will be to study the geometry and topology of this space and its quotient: the moduli space of Riemann surfaces.

Before we get to any of this, we need to talk about surfaces themselves. So, today we will recall some of the basics on surfaces.

1.1. Preliminaries on surface topology

1.1.1. Examples and classification. A *surface* is a smooth two-dimensional manifold. We call a surface *closed* if it is compact and has no boundary. A surface is said to be of *finite type* if it can be obtained from a closed surface by removing a finite number of points and (smooth) open disks with disjoint closures. In what follows, we will always assume our surfaces to be orientable.

EXAMPLE 1.1.1. To properly define a manifold, one needs to not only describe the set but also give smooth charts. In what follows we will content ourselves with the sets.

- (a) The 2-*sphere* is the surface

$$\mathbb{S}^2 = \{ (x, y, z) \in \mathbb{R}^3 : x^2 + y^2 + z^2 = 1 \}.$$

- (b) Let \mathbb{S}^1 denote the circle. The 2-*torus* is the surface

$$\mathbb{T}^2 = \mathbb{S}^1 \times \mathbb{S}^1$$

- (c) Given two (oriented) surfaces S_1, S_2 , their *connected sum* $S_1 \# S_2$ is defined as follows. Take two closed sets $D_1 \subset S_1$ and $D_2 \subset S_2$ that are both diffeomorphic to closed disks, via diffeomorphisms

$$\varphi_i : \{ (x, y) \in \mathbb{R}^2 : x^2 + y^2 \leq 1 \} \rightarrow D_i, \quad i = 1, 2,$$

so that φ_1 is orientation preserving and φ_2 is orientation reversing.

Then

$$S_1 \# S_2 = \left(S_1 \setminus \mathring{D}_1 \sqcup S_2 \setminus \mathring{D}_2 \right) / \sim$$

where \mathring{D}_i denotes the interior of D_i for $i = 1, 2$ and the equivalence relation \sim is defined by

$$\varphi_1(x, y) \sim \varphi_2(x, y) \quad \text{for all } (x, y) \in \mathbb{R}^2 \text{ with } x^2 + y^2 = 1.$$

The figure below gives an example.




FIGURE 1. A connected sum of two tori.

Like our notation suggests, the manifold $S_1 \# S_2$ is independent (up to diffeomorphism) of the choices we make (the disks and diffeomorphisms φ_i). This is a non-trivial statement, the proof of which we will skip. Likewise, we will also not prove that the connected sum of surfaces is an associative operation and that $\mathbb{S}^2 \# S$ is diffeomorphic to S for all surfaces S .

A classical result from the 19th century tells us that the three simple examples above are enough to understand all finite type surfaces up to diffeomorphism.

THEOREM 1.1.2 (Classification of closed surfaces). *Every closed orientable surface is diffeomorphic to the connected sum of a 2-sphere with a finite number of tori.*

Indeed, because the diffeomorphism type of a finite type surface does not depend on where we remove the points and open disks (another claim we will not prove), the theorem above tells us that an orientable finite type surface is (up to diffeomorphism) determined by a triple of positive integers (g, b, n) , where

- g is the number of tori in the connected sum and is called the *genus* of the surface.
- b is the number of disks removed and is called the number of *boundary components* of the surface.
- n is the number of points removed and is called the number of *punctures* of the surface.

DEFINITION 1.1.3. The triple (g, b, n) defined above will be called the *signature* of the surface. We will denote the corresponding surface by $\Sigma_{g,b,n}$ and will write $\Sigma_g = \Sigma_{g,0,0}$.

1.1.2. Euler characteristic. The Euler characteristic is a useful topological invariant of a surface. There are multiple ways to define it. We will use triangulations. A *triangulation* $\mathcal{T} = (V, E, F)$ of a closed surface S will be the data of a finite set of points $V = \{v_1, \dots, v_k\} \subset S$ (called *vertices*), a finite set of arcs $E = \{e_1, \dots, e_l\}$ with endpoints in the vertices (called *edges*) so that the complement $S \setminus (\cup v_i \cup e_j)$ consists of a collection of disks $F = \{f_1, \dots, f_m\}$ (called *faces*) that all connect to exactly 3 edges.

Note that a triangulation \mathcal{T} here is a slightly more general notion than that of a simplicial complex (it's an example of what Hatcher calls a Δ -complex [Hat02, Page 102]). Figure 2 below gives an example of a triangulation of a torus that is not a simplicial complex.




FIGURE 2. A torus with a triangulation

DEFINITION 1.1.4. S be a closed surface with a triangulation $\mathcal{T} = (V, E, F)$. The *Euler characteristic* of S is given by

$$\chi(S) = |V| - |E| + |F|.$$

Because $\chi(S)$ can be defined entirely in terms of singular homology (see [Hat02, Theorem 2.4] for details), it is a homotopy invariant. In particular this implies it should only depend on the genus of our surface S . Indeed, we have

LEMMA 1.1.5. *Let S be a closed connected and oriented surface of genus g . We have*

$$\chi(S) = 2 - 2g.$$

PROOF. Exercise: prove this using your favorite triangulation. \square

For surfaces that are not closed, we can define

$$\chi(\Sigma_{g,b,n}) = 2 - 2g - b - n.$$

This can be computed with a triangulation as well. For surfaces with only boundary components, the usual definition still works. For surfaces with punctures there no longer is a finite triangulation, so the definition above no longer makes sense. There are multiple ways out. The most natural is to use the homological definition, which gives the formula above. Another option is to allow some vertices to be missing, that is, to allow edges to run between vertices and punctures. Both give the formula above.

1.2. Riemann surfaces

For the basics on Riemann surfaces, we refer to the lecture notes from the course by Elisha Falbel [Fal23] or any of the many books on them, like [Bea84, FK92]. For a text on complex functions of a single variable, we refer to [SS03].

1.2.1. Definition and first examples. A Riemann surface is a one-dimensional complex manifold. That is,

DEFINITION 1.2.1. A *Riemann surface* X is a connected Hausdorff topological space X , equipped with an open cover $\{U_\alpha\}_{\alpha \in A}$ of open sets and maps $\varphi_\alpha : U_\alpha \rightarrow \mathbb{C}$ so that

- (1) $\varphi_\alpha(U_\alpha)$ is open and φ_α is a homeomorphism onto its image.

(2) For all $\alpha, \beta \in A$ so that $U_\alpha \cap U_\beta \neq \emptyset$ the map

$$\varphi_\alpha \circ (\varphi_\beta)^{-1} : \varphi_\beta(U_\alpha \cap U_\beta) \rightarrow \varphi_\alpha(U_\alpha \cap U_\beta)$$

is holomorphic.

The pairs $(U_\alpha, \varphi_\alpha)$ are usually called *charts* and the collection $((U_\alpha, \varphi_\alpha))_{\alpha \in A}$ is usually called an *atlas*.

Note that we do not a priori assume a Riemann surface X to be a second countable space. It is however a theorem by Radó that every Riemann surface is second countable (for a proof, see [Hub06, Section 1.3]). Moreover every Riemann surface is automatically orientable (see for instance [GH94, Page 18]).

EXAMPLE 1.2.2. (a) The simplest example is of course $X = \mathbb{C}$ equipped with one chart: the identity map.

(b) We set $X = \mathbb{C} \cup \{\infty\} = \widehat{\mathbb{C}}$ and give it the topology of the one point compactification of \mathbb{C} , which is homeomorphic to the sphere \mathbb{S}^2 . The charts are

$$U_0 = \mathbb{C}, \quad \varphi_0(z) = z$$

and

$$U_\infty = X \setminus \{0\}, \quad \varphi_\infty(z) = 1/z.$$

So $U_0 \cap U_\infty = \mathbb{C} \setminus \{0\}$ and

$$\varphi_0 \circ (\varphi_\infty)^{-1}(z) = 1/z \quad \text{for all } z \in \mathbb{C} \setminus \{0\}$$

which is indeed holomorphic on $\mathbb{C} \setminus \{0\}$. $\widehat{\mathbb{C}}$ is usually called the *Riemann sphere*.

(c) $\widehat{\mathbb{C}}$ can also be identified with the projective line

$$\mathbb{P}^1(\mathbb{C}) = (\mathbb{C}^2 \setminus \{(0, 0)\}) / \mathbb{C}^*,$$

where $\mathbb{C}^* \curvearrowright \mathbb{C}^2 \setminus \{(0, 0)\}$ by $\lambda \cdot (z, w) = (\lambda \cdot z, \lambda \cdot w)$, for $\lambda \in \mathbb{C}^*$, $(z, w) \in \mathbb{C}^2 \setminus \{(0, 0)\}$. Indeed, we may equip $\mathbb{P}^1(\mathbb{C})$ with two charts

$$U_0 = \{[z : w] : w \neq 0\}, \quad \varphi_0([z : w]) = z/w$$

and

$$U_1 = \{[z : w] : z \neq 0\}, \quad \varphi_1([z : w]) = w/z.$$

The map

$$[z : w] \mapsto \begin{cases} z/w & \text{if } w \neq 0 \\ \infty & \text{if } w = 0 \end{cases}$$

then defines a biholomorphism $\mathbb{P}^1(\mathbb{C}) \rightarrow \widehat{\mathbb{C}}$.

(d) Recall that a *domain* $D \subset \widehat{\mathbb{C}}$ is any connected and open set in $\widehat{\mathbb{C}}$. Any such domain inherits the structure of a Riemann surface from $\widehat{\mathbb{C}}$.

1.2.2. Automorphisms. To get a larger set of examples, we will consider quotients. First of all, we need the notion of a holomorphic map:

DEFINITION 1.2.3. Let X and Y be Riemann surfaces, equipped with atlases $\{(U_\alpha, \varphi_\alpha)\}_{\alpha \in A}$ and $\{(V_\beta, \psi_\beta)\}_{\beta \in B}$ respectively. A function $f : X \rightarrow Y$ is called *holomorphic* if

$$\psi_\beta \circ f \circ \varphi_\alpha^{-1} : \varphi_\alpha(U_\alpha \cap f^{-1}(V_\beta)) \rightarrow \psi_\beta(f(U_\alpha) \cap V_\beta)$$

is holomorphic for all $\alpha \in A, \beta \in B$ so that $f(U_\alpha) \cap V_\beta \neq \emptyset$. A bijective holomorphism is called a *biholomorphism* or *conformal*. $\text{Aut}(X)$ will denote the *automorphism group* of X , the set of biholomorphisms $X \rightarrow X$.

The automorphism group of the Riemann sphere is

$$\text{Aut}(\mathbb{P}^1(\mathbb{C})) = \text{PGL}(2, \mathbb{C}) = \text{GL}(2, \mathbb{C}) / \left\{ \begin{pmatrix} \lambda & 0 \\ 0 & \lambda \end{pmatrix} : \lambda \neq 0 \right\}.$$

It acts on $\mathbb{P}^1(\mathbb{C})$ through the projectivization of the linear action of $\text{GL}(2, \mathbb{C})$ on $\mathbb{C}^2 \setminus \{(0, 0)\}$. We can also describe the action on $\widehat{\mathbb{C}}$. We have:

$$(1.2.1) \quad \begin{bmatrix} a & b \\ c & d \end{bmatrix} \cdot z = \begin{cases} \frac{az+b}{cz+d} & \text{if } z \neq -d/c \\ \infty & \text{if } z = -d/c \end{cases}$$

and

$$\begin{bmatrix} a & b \\ c & d \end{bmatrix} \cdot \infty = \begin{cases} \frac{a}{c} & \text{if } c \neq 0 \\ \infty & \text{if } c = 0. \end{cases}$$

These maps are called Möbius transformations.

Finally, we observe that

$$\text{PGL}(2, \mathbb{C}) \simeq \text{PSL}(2, \mathbb{C}) = \left\{ \begin{pmatrix} a & b \\ c & d \end{pmatrix} : a, b, c, d \in \mathbb{C}, \quad ad - bc = 1 \right\} / \left\{ \pm \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} \right\}$$

1.2.3. Quotients. Many subgroups of $\text{Aut}(\mathbb{P}^1(\mathbb{C}))$ give rise to Riemann surfaces:

THEOREM 1.2.4. Let $D \subset \widehat{\mathbb{C}}$ be a domain and let $G < \text{PSL}(2, \mathbb{C})$ such that

- (1) $g(D) = D$ for all $g \in G$
- (2) If $g \in G \setminus \{e\}$ then the fixed points of g lie outside of D .
- (3) For each compact subset $K \subset D$, the set

$$\{g \in G : g(K) \cap K \neq \emptyset\}$$

is finite.

Then the quotient space

$$D/G$$

has the structure of a Riemann surface.

A group that satisfies the second condition is said to act *freely* on D and a group that satisfies the third condition is said to act *properly discontinuously* on D . The proof of this theorem will be part of the exercises.

1.2.4. Tori. The theorem from the previous section gives us a lot of new examples. The first is that of tori. Consider the elements

$$g_1 := \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix}, g_\tau := \begin{bmatrix} 1 & \tau \\ 0 & 1 \end{bmatrix} \in \mathrm{PSL}(2, \mathbb{C}),$$

for some $\tau \in \mathbb{C}$ with $\mathrm{Im}(\tau) > 0$, acting on the domain $\mathbb{C} \subset \widehat{\mathbb{C}}$ by

$$g_1(z) = z + 1 \quad \text{and} \quad g_\tau(z) = z + \tau$$

for all $z \in \mathbb{C}$.

We define the group

$$\Lambda_\tau = \langle g_1, g_\tau \rangle < \mathrm{PSL}(2, \mathbb{C}).$$

A direct computation shows that

$$\begin{bmatrix} 1 & p + q\tau \\ 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & r + s\tau \\ 0 & 1 \end{bmatrix} = \begin{bmatrix} 1 & p + q + (r + s)\tau \\ 0 & 1 \end{bmatrix},$$

for all $p, q, r, s \in \mathbb{Z}$, from which it follows that

$$\Lambda_\tau = \left\{ \begin{bmatrix} 1 & n + m\tau \\ 0 & 1 \end{bmatrix} : m, n \in \mathbb{Z} \right\} \simeq \mathbb{Z}^2.$$

Let us consider the conditions from Theorem 1.2.4. (1) is trivially satisfied: Λ_τ preserves \mathbb{C} . Any non-trivial element in Λ_τ is of the form

$$\begin{bmatrix} 1 & n + m\tau \\ 0 & 1 \end{bmatrix}$$

and hence only has the point $\infty \in \widehat{\mathbb{C}}$ as a fixed point, which gives us condition (2). To check condition (3), suppose $K \subset \mathbb{C}$ is compact. Write $d_K = \sup \{ |z - w| : z, w \in K \} < \infty$. Given $g \in \Lambda_\tau$, write

$$T_g = \inf \{ |gz - z| : z \in \mathbb{C} \}$$

for the *translation length* of g . Note that $T_g = |gz - z|$ for all $z \in \mathbb{C}$ (this is quite special to quotients of \mathbb{C}). We have


$$\{ g \in \Lambda_\tau : g(K) \cap K \neq \emptyset \} \subset \{ g \in \Lambda_\tau : T_g \leq 2d_K \}$$

and the latter is finite. So \mathbb{C}/Λ_τ is indeed a Riemann surface.

We claim that this is a torus. One way to see this is to note that the quotient map $\pi : \mathbb{C} \rightarrow \mathbb{C}/\Lambda_\tau$ restricted to the convex hull

$$\begin{aligned} \mathcal{F} &= \mathrm{conv}(\{0, 1, \tau, 1 + \tau\}) \\ &:= \{ \lambda_1 + \lambda_2\tau + \lambda_3(1 + \tau) : \lambda_1, \lambda_2, \lambda_3 \in [0, 1], \lambda_1 + \lambda_2 + \lambda_3 \leq 1 \} \end{aligned}$$

is surjective. Figure 3 shows a picture of \mathcal{F} . On $\mathring{\mathcal{F}}$, π is also injective. So to understand what the quotient looks like, we only need to understand what happens to the sides of \mathcal{F} . Since the quotient map identifies the left hand side of \mathcal{F} with the right hand side and the top with the bottom, the quotient is a torus.

FIGURE 3. A fundamental domain for the action $\Lambda_\tau \curvearrowright \mathbb{C}$.

We can also prove that \mathbb{C}/Λ_τ is a torus by using the fact that for all $z \in \mathbb{C}$ there exist unique $x, y \in \mathbb{R}$ so that

$$z = x + y\tau.$$

The map $\mathbb{C}/\Lambda_\tau \rightarrow \mathbb{S}^1 \times \mathbb{S}^1$ given by

$$[x + y\tau] \mapsto (e^{2\pi i x}, e^{2\pi i y})$$

is a homeomorphism.

Note that we have not yet proven whether all these tori are distinct as Riemann surfaces. But it will turn out later that many of them are.

1.2.5. Hyperbolic surfaces. Set $\mathbb{H}^2 = \{z \in \mathbb{C} : \text{Im}(z) > 0\}$, the upper half plane. It turns out that the automorphism group of \mathbb{H}^2 is $\text{PSL}(2, \mathbb{R})$. We will see a lot more about this later during the course, but for now we will just note that there are many subgroups of $\text{PSL}(2, \mathbb{R})$ that satisfy the conditions of Theorem 1.2.4.

It also turns out that $\text{PSL}(2, \mathbb{R})$ is exactly the group of orientation preserving isometries of the metric

$$ds^2 = \frac{dx^2 + dy^2}{y^2}.$$

This is a complete metric of constant curvature -1 . So, this means that all these Riemann surfaces naturally come equipped with a complete metric of constant curvature -1 . We will prove some of these statements and treat a first example in the first problem sheet.

1.3. The uniformization theorem and automorphism groups

The Riemann mapping theorem tells us that any pair of simply connected domains in \mathbb{C} that are both not all of \mathbb{C} are biholomorphic. In the early 20th century this was generalized by Koebe and Poincaré to a classification of *all* simply connected Riemann surfaces:

THEOREM 1.3.1 (Uniformization theorem). *Let X be a simply connected Riemann surface. Then X is biholomorphic to exactly one of*

$$\widehat{\mathbb{C}}, \quad \mathbb{C} \quad \text{or} \quad \mathbb{H}^2.$$

PROOF. See for instance [FK92, Chapter IV]. □

This theorem implies that we can see obtain every Riemann surface as a quotient of one of three Riemann surfaces. Before we formally state this, we record the following fact:

PROPOSITION 1.3.2. $\bullet \operatorname{Aut}(\widehat{\mathbb{C}}) = \operatorname{PSL}(2, \mathbb{C})$ acting by Möbius transformations,

- $\bullet \operatorname{Aut}(\mathbb{C}) = \{ \varphi : z \mapsto az + b : a \in \mathbb{C} \setminus \{0\}, b \in \mathbb{C} \} \simeq \mathbb{C} \rtimes \mathbb{C}^*,$
- $\bullet \operatorname{Aut}(\mathbb{H}^2) = \operatorname{PSL}(2, \mathbb{R})$ acting by Möbius transformations.

PROOF. See for instance [Bea84, Chapter 5] or [IT92, Section 2.3]. \square

Note that in all three cases, we have

$$\operatorname{Aut}(X) = \left\{ g \in \operatorname{Aut}(\widehat{\mathbb{C}}) : g(X) = X \right\},$$

that is, all the automorphisms of \mathbb{C} and \mathbb{H}^2 extend to $\widehat{\mathbb{C}}$. However, not all automorphisms of \mathbb{H}^2 extend to \mathbb{C} .

COROLLARY 1.3.3. *Let X be a Riemann surface. Then there exists a group $G < \operatorname{Aut}(D)$, where D is exactly one of \mathbb{C} , $\widehat{\mathbb{C}}$ or \mathbb{H}^2 so that*

- $\bullet G$ acts freely and properly discontinuously on D and
- $\bullet X = D/G$ as a Riemann surface.

PROOF. Let \widetilde{X} denote the universal cover of X and $\pi_1(X)$ its fundamental group. The fact that X is a Riemann surface, implies that \widetilde{X} can be given the structure of a Riemann surface too, so that $\pi_1(X)$ acts freely and properly discontinuously on \widetilde{X} by biholomorphisms (see for instance [IT92, Lemma 2.6]) and such that

$$\widetilde{X}/\pi_1(X) = X.$$

Since \widetilde{X} is simply connected, it must be biholomorphic to exactly one of \mathbb{C} , $\widehat{\mathbb{C}}$ or \mathbb{H}^2 . \square

1.4. Quotients of the three simply connected Riemann surfaces

Now that we know that we can obtain all Riemann surfaces as quotients of one of three simply connected Riemann surfaces, we should start looking for interesting quotients.

1.4.1. Quotients of the Riemann sphere. It turns out that for the Riemann sphere there are none:

PROPOSITION 1.4.1. *Let X be a Riemann surface. The universal cover of X is biholomorphic to $\widehat{\mathbb{C}}$ if and only if X is biholomorphic to $\widehat{\mathbb{C}}$.*

PROOF. The “if” part is clear. For the “only if” part, note that every element in $\operatorname{PSL}(2, \mathbb{C})$ has at least one fixed point on $\widehat{\mathbb{C}}$ (this either follows by direct computation or from the fact that orientation-preserving self maps of the sphere have at least one fixed point, by the Brouwer fixed point theorem [Mil65, Problem 6]). Since, by assumption

$$X = \widehat{\mathbb{C}}/G,$$

where G acts properly discontinuously and freely, we must have $G = \{e\}$. \square

1.4.2. Quotients of the plane. In Section 1.2.4, we have already seen that in the case of the complex plane, the list of quotients is a lot more interesting: there are tori. This however turns out to be almost everything:

PROPOSITION 1.4.2. *Let X be a Riemann surface. The universal cover of X is biholomorphic to \mathbb{C} if and only if X is biholomorphic to either \mathbb{C} , $\mathbb{C} \setminus \{0\}$ or*

$$\mathbb{C} / \left\langle \begin{bmatrix} 1 & \lambda \\ 0 & 1 \end{bmatrix}, \begin{bmatrix} 1 & \mu \\ 0 & 1 \end{bmatrix} \right\rangle$$

for some $\lambda, \mu \in \mathbb{C} \setminus \{0\}$ that are linearly independent over \mathbb{R} .

PROOF. First suppose $X = \mathbb{C}/G$. We claim that, since G acts properly discontinuously, G is one of the following three forms:

- (1) $G = \{e\}$
- (2) $G = \langle \varphi_b \rangle$, where $\varphi_b(z) = z + b$ for some $b \in \mathbb{C} \setminus \{0\}$
- (3) $G = \langle \varphi_{b_1}, \varphi_{b_2} \rangle$ where $b_1, b_2 \in \mathbb{C}$ are independent over \mathbb{R} .

To see this, we first prove that G cannot contain any automorphism $z \mapsto az + b$ for $a \neq 1$. Indeed, if $a \neq 1$ then $b/(1-a)$ is a fixed point for this map, which would contradict freeness of the action. Moreover, since $z \mapsto z + b_1$ and $z \mapsto z + b_2$ commute for all $b_1, b_2 \in \mathbb{C}$, G is a free abelian group and

$$G \cdot z = \{z + b : \varphi_b \in G\}.$$

In particular, if G contains $\{z \mapsto z + b_1, z \mapsto z + b_2, z \mapsto z + b_3\}$ for $b_1, b_2, b_3 \in \mathbb{C}$ that are independent over \mathbb{Q} , then $\text{span}_{\mathbb{Z}}(b_1, b_2, b_3)$ is dense in \mathbb{C} . This means that we can find a sequence $((k_i, l_i, m_i))_i$ such that

$$\varphi_{b_1}^{k_i} \circ \varphi_{b_2}^{l_i} \circ \varphi_{b_3}^{m_i}(z) \rightarrow z \quad \text{as } i \rightarrow \infty,$$

thus contradicting proper discontinuity. On a side note, we could have also used the classification of surfaces (of potentially infinite type) in the last step: there is no surface that has \mathbb{Z}^k for $k \geq 3$ as a fundamental group.

We have already seen that the third case gives rise to tori. In the second case, the surface is biholomorphic to $\mathbb{C} \setminus \{0\}$. Indeed, the map

$$[z] \in \mathbb{C}/\langle \varphi_b \rangle \quad \mapsto \quad e^{2\pi iz/b} \in \mathbb{C} \setminus \{0\}$$

is a biholomorphism.

Now let us prove the converse. For $X = \mathbb{C}$ the statement is clear. Likewise, for $X = \mathbb{C} \setminus \{0\}$, we have just seen that the composition

$$\mathbb{C} \rightarrow \mathbb{C}/(z \sim z + 1) \simeq \mathbb{C} \setminus \{0\}$$

is the universal covering map. Finally, in the proposition, the tori are given as quotients of \mathbb{C} . \square

LECTURE 2

Quotients, metrics, conformal structures

2.1. More on quotients

2.1.1. Quotients of the complex plane, continued. We saw last time that any quotient Riemann surface of \mathbb{C} is either \mathbb{C} , $\mathbb{C} - \{0\}$ or a torus. It turns out that moreover every Riemann surface structure on the torus comes from the complex plane. We have seen above that the universal cover cannot be the Riemann sphere, which means that (using the uniformization theorem) all we need to prove is that it cannot be the upper half plane either.

The fundamental group of the torus is isomorphic to \mathbb{Z}^2 , so what we need to prove is that there is no subgroup of $\text{Aut}(\mathbb{H}^2) = \text{PSL}(2, \mathbb{R})$ that acts properly discontinuously and freely on \mathbb{H}^2 and is isomorphic to \mathbb{Z}^2 . We will state this as a lemma (in which we don't unnecessarily assume that the action is free, even if in our context that would suffice):

LEMMA 2.1.1. *Suppose $G < \text{PSL}(2, \mathbb{R})$ acts properly on \mathbb{H}^2 and suppose furthermore that G is abelian. Then either $G \simeq \mathbb{Z}$ or G is finite and cyclic.*

PROOF. We will use the classification of isometries of \mathbb{H}^2 that we shall prove in the exercises: an element $g \in \text{PSL}(2, \mathbb{R})$ has either

- a single fixed point in \mathbb{H}^2 , in which case it's called elliptic and can be conjugated into $\text{SO}(2)$
- a single fixed point on $\mathbb{R} \cup \{\infty\}$, in which case it's called parabolic and can be conjugated into $\left\{ \begin{pmatrix} 1 & t \\ 0 & 1 \end{pmatrix} : t \in \mathbb{R} \right\}$
- or two fixed points on $\mathbb{R} \cup \{\infty\}$, in which case it's called hyperbolic (or loxodromic) and can be conjugated into $\left\{ \begin{pmatrix} \lambda & t \\ 0 & \frac{1}{\lambda} \end{pmatrix} : \lambda > 0 \right\}$.

If $g_1, g_2 \in \text{PSL}(2, \mathbb{R})$ commute and $p \in \mathbb{H}^2 \cup \mathbb{R} \cup \{\infty\}$ is a fixed point of g_1 , then

$$g_1(g_2(p)) = g_2 \circ g_1(p) = g_2(p).$$

That is, $g_2(p)$ is also a fixed point of g_1 .

So if G contains an elliptic element g , then all other $g' \in G \setminus \{e\}$ are elliptic as well, with the same fixed point. Moreover, by proper discontinuity (and compactness of $\text{SO}(2)$), the angles of rotation of all elements in G must be rationally related rational multiples of π . This means that G is a finite cyclic group.

Now suppose G contains a parabolic element g . Then all other $g' \in G \setminus \{e\}$ are parabolic as well, with the same fixed point (which we may assume to be ∞). If

$$\begin{pmatrix} 1 & t_1 \\ 0 & 1 \end{pmatrix}, \begin{pmatrix} 1 & t_2 \\ 0 & 1 \end{pmatrix} \in G$$

for some $t_1, t_2 \in \mathbb{R}$ that are not rationally related, then G is not discrete, which contradicts proper discontinuity (see the exercises). So $G \simeq \mathbb{Z}$.

The argument in the hyperbolic case is essentially the same as in the parabolic case. \square

Combining this lemma with the uniformization theorem, we obtain:

COROLLARY 2.1.2. *Let X be a Riemann surface that is diffeomorphic to a torus. Then the universal cover of X is biholomorphic to \mathbb{C} .*

2.1.2. Quotients of the upper half plane. It will turn out that the richest family of Riemann surfaces is that of quotients of \mathbb{H}^2 . Indeed, looking at the classification of closed orientable surfaces, we note that we have so far only seen the sphere and the torus. It turns out that all the other closed orientable surfaces also admit the structure of a Riemann surface. In fact, they all admit lots of different such structures. The two propositions above imply that they must all arise as quotients of \mathbb{H}^2 .

We will not yet discuss how to construct all these surfaces but instead discuss an example (partially taken from [GGD12, Example 1.7]). Fix some distinct complex numbers a_1, \dots, a_{2g+1} and consider the following subset of \mathbb{C}^2 :

$$\mathring{X} = \{ (z, w) \in \mathbb{C}^2 : w^2 = (z - a_1)(z - a_2) \cdots (z - a_{2g+1}) \}.$$

Let X denote the one point compactification of \mathring{X} obtained by adjoining the point (∞, ∞) .

As opposed to charts, we will describe inverse charts, or *parametrizations* around every $p \in \mathring{X}$:

- Suppose $p = (z_0, w_0) \in \mathring{X}$ is so that $z_0 \neq a_i$ for all $i = 1, \dots, 2g+1$. Set

$$\varepsilon := \min_{i=1, \dots, 2g+1} \{ |z_0 - a_i| / 2 \}$$

Then define the map $\varphi^{-1} : \{ \zeta \in \mathbb{C} : |\zeta| < \varepsilon \} \rightarrow \mathring{X}$ by

$$\varphi^{-1}(\zeta) = \left(\zeta + z_0, \sqrt{(\zeta + z_0 - a_1) \cdots (\zeta + z_0 - a_{2g+1})} \right),$$

where the branch of the square root is chosen so that $\varphi^{-1}(0) = (z_0, w_0)$, gives a parametrization.

- For $p = (a_j, 0)$, we set

$$\varepsilon := \min_{i \neq j} \{ \sqrt{|a_j - a_i|} / 2 \}$$

Then define the map $\varphi^{-1} : \{\zeta \in \mathbb{C} : |\zeta| < \varepsilon\} \rightarrow \mathring{X}$ by

$$\varphi^{-1}(\zeta) = \left(\zeta^2 + a_j, \quad \zeta \sqrt{\prod_{i \neq j} (\zeta^2 + a_j - a_i)} \right).$$

The reason that we need to take different charts around these points is that

$$\sqrt{z - a_j}$$

is not a well defined holomorphic function near $z = a_j$.

Also note that the choice of the branch of the root does not matter. By changing the branch we would obtain a new parametrization $\tilde{\varphi}^{-1}$ that satisfies $\tilde{\varphi}^{-1}(\zeta) = \varphi^{-1}(-\zeta)$.

It's not hard to see that \mathring{X} is not bounded as a subset of \mathbb{C}^2 . This means in particular that it's not compact. We can however compactify it in a similar fashion to how we compactified \mathbb{C} in order to obtain the Riemann sphere. That is, we add a point (∞, ∞) and around this point define a parametrization:

$$\varphi_{\infty}^{-1}(\zeta) = \begin{cases} \left(\zeta^{-2}, \quad \zeta^{-(2g+1)} \sqrt{(1 - a_1 \zeta^2) \cdots (1 - a_{2g+1} \zeta^2)} \right) & \text{if } \zeta \neq 0 \\ (\infty, \infty) & \text{if } \zeta = 0, \end{cases}$$

for all $\zeta \in \{|\zeta| < \varepsilon\}$ and some appropriate $\varepsilon > 0$.

The reason that the resulting surface X is compact is that we can write it as the union of the sets

$$\left\{ (z, w) \in \mathring{X} : |z| \leq 1/\varepsilon^2 \right\} \cup \left(\left\{ (z, w) \in \mathring{X} : |z| \geq 1/\varepsilon^2 \right\} \cup \{(\infty, \infty)\} \right),$$

for some small $\varepsilon > 0$. The first set is compact because it's a bounded subset of \mathbb{C}^2 . The second set is compact because it's $\varphi_{\infty}^{-1}(\{|\zeta| \leq \varepsilon\})$.

To see that X is connected, we could proceed using charts as well. We would have to find a collection of charts that are all connected, overlap and cover X . However, it's easier to use complex analysis. Suppose $z_0 \neq a_i$ for all $i = 1, \dots, 2g+1$ and $z_0 \neq \infty$. In that case, we can define a path

$$z(t) \mapsto \left(z(t), \quad \sqrt{\prod_{i=1}^{2g+1} (z(t) - a_i)} \right)$$

where $z(t)$ is some continuous path in \mathbb{C} between z_0 and a_i and we pick a continuous branch of the square root, thus connecting any point $(z_0, w_0) \in X$ to $(0, a_i)$.

To figure out the genus of X , note that there is a map $\pi : X \rightarrow \widehat{\mathbb{C}}$ given by

$$\pi(z, w) = z \quad \text{for all } (z, w) \in X.$$

This map is two-to-one almost everywhere. Only the points $z = a_i$, $i = 1, \dots, 2g+1$ and the point $z = \infty$ have only one pre-image.

Now triangulate $\widehat{\mathbb{C}}$ so that the vertices of the triangulation coincide with the points $a_1, \dots, a_{2g+1}, \infty$. If we lift the triangulation to X using π , we can compute the Euler characteristic of X . Every face and every edge in the triangulation of $\widehat{\mathbb{C}}$ has two pre-images, whereas each vertex has only one. This means that:

$$\chi(X) = 2\chi(\widehat{\mathbb{C}}) - (2g + 2) = 2 - 2g.$$

Because X is an orientable closed surface, we see that it must have genus g (Lemma 1.1.5). In particular, if $g \geq 2$, these surfaces are quotients of \mathbb{H}^2 . Note that this also implies that for $g \geq 1$, the Riemann surface \hat{X} is also a quotient of \mathbb{H}^2 . Note that we could have also used the Riemann–Hurwitz formula for this calculation. Incidentally, this formula can be proved using a similar argument to what we just did above.

To get a picture of what X looks like, draw a closed arc α_1 between a_1 and a_2 on $\widehat{\mathbb{C}}$, an arc α_2 between a_3 and a_4 that does not intersect the first arc and so on, and so forth. The last arc α_{g+1} goes between a_{2g+1} and ∞ . Figure 1 shows a picture of what these arcs might look like.




FIGURE 1. $\widehat{\mathbb{C}}$ with some intervals removed.


Let

$$D = \widehat{\mathbb{C}} \setminus \left(\bigcup_{i=1}^{g+1} \alpha_i \right).$$

The map

$$\pi|_{\pi^{-1}(D)} : \pi^{-1}(D) \rightarrow D$$

is now a two-to-one map. Moreover on the arcs, it's two-to-one on the interior and one-to-one on the boundary. Because it's also smooth, this means that the pre-image of the arcs is a circle. So, X may be obtained (topologically) by cutting $\widehat{\mathbb{C}}$ open along the arcs, taking two copies of that, and gluing these along their boundary. Figure 2 depicts this process.

FIGURE 2. Gluing X out of two Riemann spheres.

Finally, we note that our Riemann surfaces come with an involution $\iota : X \rightarrow X$, given by

$$\iota(w) = \begin{cases} -w & \text{if } w \neq \infty \\ \infty & \text{if } w = \infty. \end{cases}$$

This map is called the *hyperelliptic involution* and the surfaces we described are hence called *hyperelliptic surfaces*. Note that $\pi : X \rightarrow \widehat{\mathbb{C}}$ is the quotient map $X \rightarrow X/\iota$.

2.2. Riemannian metrics and Riemann surfaces

We already noted that every Riemann surface comes with a natural Riemannian metric. Indeed the Riemann sphere has the usual round metric of constant curvature $+1$. Likewise, \mathbb{C} has a flat metric, its usual Euclidean metric $\text{Aut}(\mathbb{C})$ does not act by isometries. However, in the proof of Proposition 1.4.2, we saw that all the quotients are obtained by quotienting by a group that does act by Euclidean isometries. This means that the Euclidean metric descends. Finally, we proved in the exercises that $\text{Aut}(\mathbb{H}^2)$ also acts by isometries of the hyperbolic metric defined in Section 1.2.5. So every quotient of \mathbb{H}^2 comes with a natural metric of constant curvature -1 .

It turns out that we can also go the other way around. That is: Riemann surface structures on a given surface are in one-to-one correspondence with complete metrics of constant curvature.

One way to see this uses the Killing-Hopf theorem. In the special case of surfaces, this states that every oriented surface equipped with a Riemannian metric of constant curvature $+1$, 0 or -1 can be obtained as the quotient by a group of orientation preserving isometries acting properly discontinuously and freely on \mathbb{S}^2 equipped with the round metric, \mathbb{R}^2 equipped with the Euclidean metric or \mathbb{H}^2 equipped with the hyperbolic metric respectively (see [CE08, Theorem 1.37] for a proof). For a Riemannian manifold M , let us write

$$\text{Isom}^+(M) = \{ \varphi : M \rightarrow M : \varphi \text{ is an orientation preserving isometry} \}.$$

So, we need the fact that

- (1) $\text{Isom}^+(\mathbb{S}^2) = \text{SO}(2, \mathbb{R})$ and this has no non-trivial subgroups that act properly discontinuously on \mathbb{S}^2 .
- (2) $\text{Isom}^+(\mathbb{R}^2) = \text{SO}(2, \mathbb{R}) \ltimes \mathbb{R}^2$, where \mathbb{R}^2 acts by translations. The only subgroups of this group that act properly discontinuously and freely are the fundamental groups of tori and cylinders.

$$(3) \operatorname{Isom}^+(\mathbb{H}^2) = \operatorname{PSL}(2, \mathbb{R}).$$

Given the above, we get our one-to-one correspondence:

PROPOSITION 2.2.1. *Given an orientable surface Σ of finite type with $\partial\Sigma = \emptyset$, the identification described above gives a one-to-one correspondence of sets*

$$\left\{ \begin{array}{c} \text{Riemann surface} \\ \text{structures on } \Sigma \end{array} \right\} / \sim \leftrightarrow \left\{ \begin{array}{c} \text{Complete Riemannian} \\ \text{metrics of constant} \\ \text{curvature } \{-1, 0, +1\} \\ \text{on } \Sigma \end{array} \right\} / \sim,$$

where the equivalence on the left is biholomorphism and the equivalence on the right is isometry (and homothety in the Euclidean case).

PROOF SKETCH. From the above we see that a Riemann surface structure on Σ yields a metric of constant curvature and vice versa. We only need to check that biholomorphic Riemann surfaces yield isometric/homothetic metrics and vice versa.

Suppose $h : X \rightarrow Y$ is a biholomorphism. We may lift this to a biholomorphism $\tilde{h} : \tilde{X} \rightarrow \tilde{Y}$ of the universal covers \tilde{X} and \tilde{Y} of X and Y respectively. There are three cases to treat: $\tilde{X} \simeq \tilde{Y} \simeq \mathbb{C}, \hat{\mathbb{C}}, \mathbb{H}^2$. Because it's the most interesting case, we will treat the former, i.e. $\tilde{X} \simeq \tilde{Y} \simeq \mathbb{C}$. We will also assume X and Y are tori. If we write

$$X \simeq \mathbb{C}/\Lambda_1 \quad \text{and} \quad Y \simeq \mathbb{C}/\Lambda_2,$$

then we get that $\tilde{h} \in \operatorname{Aut}(\mathbb{C})$ is such that $\tilde{h}(\Lambda_1) = \Lambda_2$. Since all automorphisms of \mathbb{C} are of the form $z \mapsto az + b$ for $a \in \mathbb{C}^*$ and $b \in \mathbb{C}$, Λ_2 is obtained from Λ_1 by translating, scaling and rotating. This means that the quotient metrics are homothetic.

The proof of the reverse direction and both directions of all the remaining cases are similar. \square

Whether the curvature is 0, +1 or -1 is determined by the topology of Σ . This for instance follows from the discussion above. It can also be seen from the Gauss-Bonnet theorem. Recall that in the case of a closed Riemannian surface X , this states that

$$\int_X K \, dA = 2\pi \, \chi(\Sigma),$$

where K is the Gaussian curvature on X and dA the area measure. For constant curvature κ , this means that

$$\kappa \cdot \operatorname{area}(X) = 2\pi \, \chi(X)$$

So $\chi(X) = 0$ if and only if $\kappa = 0$ and otherwise $\chi(X)$ needs to have the same sign as κ . This last equality generalizes to finite type surfaces and we obtain:

LEMMA 2.2.2. *Let X be a hyperbolic surface homeomorphic to $\Sigma_{g,b,n}$ then*

$$\operatorname{area}(X) = 2\pi(2g + n + b - 2).$$

2.3. Conformal structures

There is another type of structures on a surface that is in one-to-one correspondence with Riemann surface structures, namely conformal structures.

We say that two Riemannian metrics ds_1^2 and ds_2^2 on a surface X are *conformally equivalent* if there exists a positive function $\rho : X \rightarrow \mathbb{R}_+$ so that

$$ds_1^2 = \rho \cdot ds_2^2.$$

So a conformal equivalence class of Riemannian metrics can be seen as a notion of angles on the surface.

We have already seen that a Riemann surface structure induces a Riemannian metric on the surface, so it certainly also induces a conformal class of metrics.

So, we need to explain how to go back. We will also only sketch this. First of all, suppose we are given a surface X with oriented charts $(U_j, (u_j, v_j))_j$ equipped with a Riemannian metric that in all local coordinates (u_j, v_j) is of the form

$$ds^2 = \rho(u_j, v_j) \cdot (du_j^2 + dv_j^2),$$

where $\rho : X \rightarrow \mathbb{R}_+$ is some smooth function. Consider the complex-valued coordinate

$$w_j = u_j + i v_j.$$

We claim that this is holomorphic. Indeed, applying a coordinate change on $U_j \cap U_k$, we have

$$ds^2 = \rho(u_k, v_k) \cdot \left[\left(\left(\frac{\partial u_j}{\partial u_k} \right)^2 + \left(\frac{\partial v_j}{\partial u_k} \right)^2 \right) du_k^2 + \left(\left(\frac{\partial u_j}{\partial v_k} \right)^2 + \left(\frac{\partial v_j}{\partial v_k} \right)^2 \right) dv_k^2 + 2 \left(\frac{\partial u_j}{\partial u_k} \frac{\partial u_j}{\partial v_k} + \frac{\partial v_j}{\partial u_k} \frac{\partial v_j}{\partial v_k} \right) du_k dv_k \right].$$

Our assumption implies that

$$(2.3.1) \quad \left(\frac{\partial u_j}{\partial u_k} \right)^2 + \left(\frac{\partial v_j}{\partial u_k} \right)^2 = \left(\frac{\partial u_j}{\partial v_k} \right)^2 + \left(\frac{\partial v_j}{\partial v_k} \right)^2$$

and

$$(2.3.2) \quad \frac{\partial u_j}{\partial u_k} \frac{\partial u_j}{\partial v_k} + \frac{\partial v_j}{\partial u_k} \frac{\partial v_j}{\partial v_k} = 0.$$

This last line can be written as

$$\det \begin{pmatrix} \partial u_j / \partial u_k & \partial v_j / \partial u_k \\ -\partial v_j / \partial u_k & \partial u_j / \partial v_k \end{pmatrix} = 0.$$

So this implies that

$$\begin{pmatrix} \partial u_j / \partial u_k \\ -\partial v_j / \partial u_k \end{pmatrix} = \lambda \cdot \begin{pmatrix} \partial v_j / \partial v_k \\ \partial u_j / \partial v_k \end{pmatrix}$$

for some $\lambda \in \mathbb{R}$. Filling this into (2.3.1), we obtain

$$\lambda^2 \cdot \left(\left(\frac{\partial u_j}{\partial v_k} \right)^2 + \left(\frac{\partial v_j}{\partial v_k} \right)^2 \right) = \left(\frac{\partial u_j}{\partial v_k} \right)^2 + \left(\frac{\partial v_j}{\partial v_k} \right)^2.$$

So $\lambda \in \{\pm 1\}$. Now using that our surface is oriented, i.e. that the determinant of the Jacobian of the chart transition is positive, we obtain that $\lambda = 1$. And hence

$$\frac{\partial u_j}{\partial u_k} = \frac{\partial v_j}{\partial v_k} \quad \text{and} \quad -\frac{\partial v_j}{\partial u_k} = \frac{\partial u_j}{\partial v_k},$$

the Cauchy-Riemann equations for the chart transition $w_k \circ w_j^{-1}$, which means that these coordinates are indeed holomorphic. The coordinates (U_j, w_j) are usually called *isothermal coordinates*.

Also note that we have not used the factor ρ , so any metric that is conformal to our metric will give us the same structure. Moreover, our usual coordinate ‘ z ’ on the three simply connected Riemann surfaces is an example of an isothermal coordinate, so if we apply the procedure above to the metric we obtain from our quotients, we find the same complex structure back.

This means that what we need to show is that for each Riemannian metric (that is not necessarily given to us in the form above), we can find a set of coordinates so that our metric takes this form. So, suppose our metric is given by

$$ds^2 = A dx^2 + 2B dx dy + C dy^2$$

in some local coordinates (x, y) .

Writing $z = x + iy$, we get that

$$ds^2 = \lambda |dz + \mu d\bar{z}|^2 := \lambda(dz + \mu d\bar{z})(d\bar{z} + \bar{\mu} dz),$$

where

$$\lambda = \frac{1}{4} \left(A + C + 2\sqrt{AC - B^2} \right) \quad \text{and} \quad \mu = \frac{A - C + 2iB}{A + C + 2\sqrt{AC - B^2}}.$$

We are looking for a coordinate $w = u + iv$ so that

$$ds^2 = \rho(du^2 + dv^2) = \rho |dw|^2 = \rho \cdot \left| \frac{\partial w}{\partial z} \right|^2 \cdot \left| dz + \frac{\partial w / \partial \bar{z}}{\partial w / \partial z} d\bar{z} \right|^2,$$

where

$$\frac{\partial w}{\partial z} = \frac{\partial w}{\partial x} - i \frac{\partial w}{\partial y} \quad \text{and} \quad \frac{\partial w}{\partial \bar{z}} = \frac{\partial w}{\partial x} + i \frac{\partial w}{\partial y}$$

are called *Wirtinger derivatives*, we will discuss these in slightly more detail in [Section 5.1](#).

This means that isothermal coordinates exist if there is a solution to the partial differential equation

$$\frac{\partial w}{\partial \bar{z}} = \mu \cdot \frac{\partial w}{\partial z}.$$

It turns out this solution does indeed exist on a surface, which means that we obtain a Riemann surface structure. Moreover, it turns out this map is one-to-one. In particular, holomorphic maps are conformal. So we obtain

PROPOSITION 2.3.1. *Given an orientable surface Σ of finite type with $\partial\Sigma = \emptyset$, the identification described above gives a one-to-one correspondence of sets*

$$\left\{ \begin{array}{l} \text{Riemann surface} \\ \text{structures on } \Sigma \end{array} \right\} \Big/ \text{biholom.} \leftrightarrow \left\{ \begin{array}{l} \text{Conformal classes} \\ \text{of Riemannian} \\ \text{metrics on } \Sigma \end{array} \right\} \Big/ \text{diffeomorphism.}$$

Combined with Proposition 2.2.1, the proposition above also implies that in every conformal class of metrics there is a metric of constant curvature that is unique (up to scaling if the metric is flat). This can also be proved without passing through the uniformization theorem, which comes down to solving a non-linear PDE on the surface. This was treated in Olivier Biquard's course *Introduction à l'analyse géométrique*.

LECTURE 3

The Teichmüller space of the torus

3.1. Riemann surface structures on the torus

The goal of the rest of this course is to understand the deformation spaces associated to Riemann surfaces: Teichmüller and moduli spaces.

In general, the Teichmüller space associated to a surface will be a space of *marked* Riemann surface structures on that surface and the corresponding moduli space will be a space of isomorphism classes of Riemann surface structures. As such, the moduli space associated to a surface will be a quotient of the corresponding Teichmüller space.

First of all, note that the uniformization theorem tells us that there is only one Riemann surface structure on the sphere. This means that the corresponding moduli space will be a point. It turns out that the same holds for its Teichmüller space. This means that the lowest genus closed surface for which we can expect an interesting deformation space is the torus.

So, let us parametrize Riemann surface structures on the torus. Recall from Proposition 1.4.2 that every Riemann surface structure on the torus is of the form

$$\mathbb{C} / \left\langle \begin{bmatrix} 1 & \lambda \\ 0 & 1 \end{bmatrix}, \begin{bmatrix} 1 & \mu \\ 0 & 1 \end{bmatrix} \right\rangle$$

for some $\lambda, \mu \in \mathbb{C} \setminus \{0\}$ that are linearly independent over \mathbb{R} .

First of all note that every such torus is biholomorphic to a torus of the form

$$R_\tau := \mathbb{C} / \Lambda_\tau,$$

for some $\tau \in \mathbb{H}^2$, where

$$\Lambda_\tau = \left\langle \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix}, \begin{bmatrix} 1 & \tau \\ 0 & 1 \end{bmatrix} \right\rangle.$$

Indeed, rotating and rescaling the lattice induce biholomorphisms on the level of Riemann surfaces (as we have already noted in the proof sketch of Proposition 2.2.1)

However, there are still distinct $\tau, \tau' \in \mathbb{H}^2$ that lead to holomorphic tori. We have:

PROPOSITION 3.1.1. *Let $\tau, \tau' \in \mathbb{H}^2$. The two tori R_τ and $R_{\tau'}$ are biholomorphic if and only if*

$$\tau' = \frac{a\tau + b}{c\tau + d}$$

for some $a, b, c, d \in \mathbb{Z}$ with $ad - bc = 1$.

PROOF. First assume R_τ and $R_{\tau'}$ are biholomorphic and let $f : R_{\tau'} \rightarrow R_\tau$ be a biholomorphism. Lift f to a biholomorphism $\tilde{f} : \mathbb{C} \rightarrow \mathbb{C}$. This means that

$$\tilde{f}(z) = \alpha z + \beta$$

for some $\alpha, \beta \in \mathbb{C}$. By postcomposing with a biholomorphism of \mathbb{C} , we may assume that $\tilde{f}(0) = 0$.

Because \tilde{f} is a lift, we know that both $\tilde{f}(1)$ and $\tilde{f}(\tau')$ are equivalent to 0 under Λ_τ . So

$$\tilde{f}(\tau') = \alpha\tau' = a\tau + b$$

$$\tilde{f}(1) = \alpha = c\tau + d$$

for some $a, b, c, d \in \mathbb{Z}$. So

$$\tau' = \frac{a\tau + b}{c\tau + d}.$$

So we only need to show that $ad - bc = 1$. Moreover, since $\tilde{f}(\Lambda_{\tau'}) = \Lambda_\tau$, $f(\tau') = a\tau + b$ and $f(1) = c\tau + d$ generate Λ_τ . This means that the map

$$m\tau + n \mapsto m \cdot (a\tau + b) + n \cdot (c\tau + d)$$

is an automorphism of Λ , and hence $\begin{pmatrix} a & b \\ c & d \end{pmatrix} \in \text{GL}(2, \mathbb{Z})$. So, we obtain $ad - bc = \pm 1$.

Since

$$\text{Im}(\tau') = \frac{ad - bc}{|c\tau + d|^2} > 0,$$

we get $ad - bc = 1$.

Conversely, if

$$\tau' = \frac{a\tau + b}{c\tau + d}$$

Then

$$f([z]) = [(c\tau + d)z]$$

gives a biholomorphic map $f : R_{\tau'} \rightarrow R_\tau$. □

3.2. The Teichmüller and moduli spaces of tori

Looking at Proposition 3.1.1, we see that we can parametrize all complex structures on the torus with the set


$$\mathcal{M}_1 = \text{SL}(2, \mathbb{Z}) \backslash \mathbb{H}^2 = \text{PSL}(2, \mathbb{Z}) \backslash \mathbb{H}^2.$$

Moreover this set is the quotient of the hyperbolic plane by a group ($\text{PSL}(2, \mathbb{Z})$) of isometries that acts properly discontinuously on it. However, the group doesn't quite act freely, so it's not directly a hyperbolic surface.

So, let us investigate the structure of this quotient. One way of doing this is to find a fundamental domain for the action of $\text{PSL}(2, \mathbb{Z})$ on \mathbb{H}^2 . Set

$$\mathcal{F} = \left\{ z \in \mathbb{H}^2 : |z| \geq 1 \text{ and } -\frac{1}{2} \leq \text{Re}(z) \leq \frac{1}{2} \right\}.$$

Figure 1 shows a picture of \mathcal{F} .

FIGURE 1. A fundamental domain for the action of $\mathrm{PSL}(2, \mathbb{Z})$ on \mathbb{H}^2 .

We claim

PROPOSITION 3.2.1. *For all $\tau \in \mathbb{H}^2$ there exists an element $g \in \mathrm{PSL}(2, \mathbb{Z})$ so that $g\tau \in \mathcal{F}$. Moreover,*

- if $\tau \in \mathring{\mathcal{F}}$ then

$$\left(\mathrm{PSL}(2, \mathbb{Z}) \cdot \tau \right) \cap \mathcal{F} = \{\tau\},$$

- if $\mathrm{Re}(\tau) = \frac{1}{2}$ then

$$\left(\mathrm{PSL}(2, \mathbb{Z}) \cdot \tau \right) \cap \mathcal{F} = \{g\tau, g\tau + 1\},$$

- if $\mathrm{Re}(\tau) = -\frac{1}{2}$ then

$$\left(\mathrm{PSL}(2, \mathbb{Z}) \cdot \tau \right) \cap \mathcal{F} = \{g\tau, g\tau - 1\}.$$


- and if $|\tau| = 1$ then

$$\left(\mathrm{PSL}(2, \mathbb{Z}) \cdot \tau \right) \cap \mathcal{F} = \{g\tau, -1/\tau\},$$

The proof of this proposition is part of this week's exercises.

Since $T : z \mapsto z + 1$ maps the line $\mathrm{Re}(z) = -1/2$ to the line $\mathrm{Re}(z) = 1/2$ and $S : z \mapsto -1/z$ fixes i and swaps $(-1 + \sqrt{3}i)/2$ and $(1 + \sqrt{3}i)/2$ (which are in turn the fixed points of ST). These turn out to be the only side identifications and thus the quotient looks like Figure

2:

FIGURE 2. A cartoon of \mathcal{M}_1 .

So \mathcal{M}_1 is a space that has the structure of a hyperbolic surface near almost every point. The only problematic points are the images of i and $(\pm 1 + \sqrt{3}i)/2$, where the \mathcal{M}_1 looks like a cone. The technical term for such a space is a hyperbolic *orbifold*.

\mathcal{M}_1 is called the *moduli space* of tori. $\mathcal{T}_1 = \mathbb{H}^2$ is called the *Teichmüller space* of tori.

Our next intermediate goal is to generalize this to all surfaces. To this end, we will introduce a different perspective on \mathcal{T}_1 , that generalizes naturally to higher genus surfaces.

3.3. \mathcal{T}_1 as a space of marked structures

Our objective in this section is to understand what the information is that is parametrized by \mathcal{T}_1 .

3.3.1. Markings as a choice of generators for $\pi_1(R)$. So, suppose $\tau \in \mathbb{H}^2$ and $\tau' = g\tau$ for some $g = \begin{bmatrix} a & b \\ c & d \end{bmatrix} \in \text{PSL}(2, \mathbb{Z})$. Let $f : R_{\tau'} \rightarrow R_{\tau}$ denote the biholomorphism from the proof of Proposition 3.1.1. We saw that we can find a lift $\tilde{f} : \mathbb{C} \rightarrow \mathbb{C}$ so that $\tilde{f}(z) = (c\tau + d)z$. In particular, using the relation between τ and τ' , we see that

$$\tilde{f}(\{1, \tau'\}) = \{c\tau + d, a\tau + b\}.$$

So, the biholomorphism corresponds to a base change (i.e. the change of a choice of generators) for Λ_{τ} .

Let us formalize this idea of a base change. First we take a base point $p_0 = [0] \in R_{\tau}$ for the fundamental group $\pi_1(R_{\tau}, p_0)$. The segments between 0 and 1 and between 0 and τ project to simple closed curves on R_{τ} and determine generators

$$[A_{\tau}], [B_{\tau}] \in \pi_1(R_{\tau}, p_0).$$

This now also gives us a natural choice of isomorphism $\Lambda_{\tau} \simeq \pi_1(R_{\tau}, p_0)$, mapping

$$1 \mapsto [A_{\tau}] \quad \text{and} \quad \tau \mapsto [B_{\tau}].$$

Likewise, for $R_{\tau'}$ we also obtain a natural system of generators $[A_{\tau'}], [B_{\tau'}] \in \pi_1(R_{\tau'}, p_0)$. Moreover, if $f_* : \pi_1(R_{\tau'}, p_0) \rightarrow \pi_1(R_{\tau}, p_0)$ denotes the map f induces on the fundamental group, then

$$f_*([A_{\tau'}]) \neq [A_{\tau}] \quad \text{and} \quad f_*([B_{\tau'}]) \neq [B_{\tau}].$$

Let us package these choices of generators:

DEFINITION 3.3.1. Let R be a Riemann surface homeomorphic to \mathbb{T}^2 .

- (1) A *marking* on R is a generating set $\Sigma_p \subset \pi_1(R, p)$ consisting of two elements.
- (2) Two markings Σ_p and $\Sigma_{p'}$ are called *equivalent* if there exists a continuous curve α from p to p' so that the corresponding isomorphism $T_\alpha : \pi_1(R, p) \rightarrow \pi_1(R, p')$ satisfies

$$T_\alpha(\Sigma_p) = \Sigma_{p'}.$$

Two pairs (R, Σ) and (R', Σ') of marked Riemann surfaces homeomorphic to \mathbb{T}^2 are called *equivalent* if there exists a biholomorphic mapping $h : R \rightarrow R'$ such that

$$h_*(\Sigma) \simeq \Sigma'.$$

Note that above we have *not* proved that $(R_{\tau}, \{[A_{\tau}], [B_{\tau}]\})$ and $(R_{\tau'}, \{[A_{\tau'}], [B_{\tau'}]\})$ are equivalent as marked Riemann surfaces, because our map f_* did not send the generators to each other, and in fact, they are not equivalent:

THEOREM 3.3.2. Let $\tau, \tau' \in \mathcal{T}_1$. Then the marked Riemann surfaces

$$(R_{\tau}, \{[A_{\tau}], [B_{\tau}]\}) \quad \text{and} \quad (R_{\tau'}, \{[A_{\tau'}], [B_{\tau'}]\})$$

are equivalent if and only if $\tau' = \tau$. Moreover, we have an identification

$$\mathcal{T}_1 = \left\{ (R, \Sigma_p) : \begin{array}{l} R \text{ a Riemann surface homemorphic to } \mathbb{T}^2 \\ p \in R, \Sigma_p \text{ a marking on } R \end{array} \right\} / \sim.$$

PROOF. We begin by proving part of the second claim: every marked complex torus is equivalent to a marked torus of the form $(R_{\tau}, \{[A_{\tau}], [B_{\tau}]\})$. So, suppose (R, Σ) is a marked torus. We know that R is biholomorphic to R_{τ} for some $\tau \in \mathcal{T}_1$. Moreover, since $\Sigma = \{[A], [B]\}$ is a minimal generating set for Λ_{τ} , we can find a lattice isomorphism $\varphi : \Lambda_{\tau} \rightarrow \Lambda_{\tau}$ so that

$$\varphi([A]) = 1.$$

Potentially switching the roles of $[A]$ and $[B]$, we can assume φ is an element of $\text{SL}(2, \mathbb{Z})$ and hence that $\varphi([B])$ lies in \mathbb{H}^2 . The torus $R_{\varphi([B])}$ is biholomorphic to R_{τ} . So (R, Σ) is equivalent to

$$(R_{\varphi([B])}, \{A_{\varphi([B])}, B_{\varphi([B])}\}).$$

So, to prove the theorem, we need to show that $(R_{\tau}, \{[A_{\tau}], [B_{\tau}]\})$ and $(R_{\tau'}, \{[A_{\tau'}], [B_{\tau'}]\})$ are equivalent if and only if $\tau = \tau'$. Of course, if $\tau = \tau'$ then the two corresponding marked surfaces are equivalent, so we need to show the other direction.

So let $h : R_{\tau'} \rightarrow R_{\tau}$ be a biholomorphism that induces the equivalence. We may assume that $h([0]) = [0]$ and take a lift $\tilde{h} : \mathbb{C} \rightarrow \mathbb{C}$ so that

$$\tilde{h}(0) = 0.$$

This means that $\tilde{h}(z) = \alpha z$ for some $\alpha \in \mathbb{C} \setminus \{0\}$. Hence $1 = \tilde{h}(1) = \alpha$, which implies that $\tau = \tilde{h}(\tau') = \tau'$. \square

Note that so far, our alternate description of Teichmüller space only recovers the set \mathcal{T}_1 and not yet its topology. Of course we can use the bijection to define a topology. However, there is also an intrinsic definition. We will discuss how to do this later.

3.4. Markings by diffeomorphisms

First, we give a third interpretation of \mathcal{T}_1 . This goes through another (equivalent) way of marking Riemann surfaces.

To this end, once and for all fix a surface S diffeomorphic to \mathbb{T}^2 . We define:

DEFINITION 3.4.1. Let R and R' be Riemann surfaces and let

$$f : S \rightarrow R \quad \text{and} \quad f' : S \rightarrow R'$$

be orientation preserving diffeomorphisms. We say that the pairs (R, f) and (R', f') are *equivalent* if there exists a biholomorphism $h : R \rightarrow R'$ so that

$$(f')^{-1} \circ h \circ f : S \rightarrow S$$

is homotopic to the identity.

Note that if we pick a generating set $\{[A], [B]\}$ for the fundamental group $\pi_1(S, p)$ then every pair (R, f) as above defines a point

$$(R, \{f_*([A]), f_*([B])\}) \in \mathcal{T}_1.$$

It turns out that this gives another description of the Teichmüller space of tori:

THEOREM 3.4.2. Fix S and $[A], [B] \in \pi_1(S, p)$ as above. Then the map

$$\left\{ (R, f) : \begin{array}{l} R \text{ a Riemann surface, } f : S \rightarrow R \\ \text{an orientation preserving diffeomorphism} \end{array} \right\} / \sim \rightarrow \mathcal{T}_1$$

given by

$$(R, f) \mapsto (R, \{f_*([A]), f_*([B])\}),$$

is a well-defined bijection.

PROOF. We start with well-definedness. Meaning, suppose (R, f) and (R', f') are equivalent. By definition, this means that there exists a biholomorphic map $h : R \rightarrow R'$ so that

$$h \circ f : S \rightarrow R' \quad \text{and} \quad f' : S \rightarrow R'$$

are homotopic. Now if α is a continuous arc between $f'(p)$ and $h(f(p))$, we see that T_α induces an equivalence between the markings

$$\{f'_*([A]), f'_*([B])\} \quad \text{and} \quad \{(h \circ f)_*([A]), (h \circ f)_*([B])\},$$

making $(R, \{f_*([A]), f_*([B])\})$ and $(R', \{f'_*([A]), f'_*([B])\})$ equivalent. This means that they correspond to the same point by the previous theorem. So, the map is well defined.

Moreover, the map is surjective. For any $\tau \in \mathcal{T}_1$ we can find an orientation preserving diffeomorphism $f : S \rightarrow R_\tau$. Indeed, we know that there exists some $\tau_0 \in \mathcal{T}_1$ such that $(S, \{[A], [B]\}) \sim (R_{\tau_0}, \{[A_{\tau_0}], [B_{\tau_0}]\})$ as marked surfaces. One checks that the map $f_\tau : \mathbb{C} \rightarrow \mathbb{C}$ given by

$$f_\tau(z) = \frac{(\tau - \bar{\tau}_0)z - (\tau - \tau_0)\bar{z}}{\tau_0 - \bar{\tau}_0}$$

descends to an orientation preserving diffeomorphism $R_{\tau_0} \rightarrow R_\tau$ that induces the marking $\{[A_\tau], [B_\tau]\}$ on R_τ .

For the injectivity, suppose that

$$\left[(R, \{f_*([A]), f_*([B])\}) \right] = \left[(R', \{f'_*([A]), f'_*([B])\}) \right].$$

Take $\tau_0 \in \mathcal{T}_1$ such that

$$\left[(S, \{[A], [B]\}) \right] = \left[(R_{\tau_0}, \{[A_{\tau_0}], [B_{\tau_0}]\}) \right].$$

Moreover, let $h : R \rightarrow R'$ be a holomorphism such that

$$h_*\{f_*([A]), f_*([B])\} = \{f'_*([A]), f'_*([B])\}.$$

We choose lattices $\Lambda, \Lambda' \subset \mathbb{C}$, generated by $(1, a)$ and $(1, a')$ respectively such that

$$R = \mathbb{C}/\Lambda \quad \text{and} \quad R' = \mathbb{C}/\Lambda',$$

and the generators induce the bases $\{f_*([A]), f_*([B])\}$ and $\{f'_*([A]), f'_*([B])\}$ respectively.

Now, let $\tilde{f}, \tilde{f}', \tilde{h} : \mathbb{C} \rightarrow \mathbb{C}$ be lifts. We may assume that

$$\tilde{f}(0) = \tilde{f}'(0) = \tilde{h}(0) = 0, \quad \tilde{f}(1) = \tilde{f}'(1) = \tilde{h}(1) = 1,$$

and

$$\tilde{f}(\tau_0) = a, \quad \tilde{f}'(\tau_0) = a' \quad \text{and} \quad \tilde{h}(a) = a'$$

So we obtain a homotopy $F_t : \mathbb{C} \rightarrow \mathbb{C}$ defined by

$$F_t(z) = (1 - t) \tilde{h} \circ \tilde{f}(z) + t \tilde{f}'(z)$$

between \tilde{f} and \tilde{f}' that descends to a homotopy between $f : S \rightarrow R'$ and $f' : S \rightarrow R'$. \square

3.5. The Teichmüller space of Riemann surfaces of a given type

The two description of the Teichmüller space of the torus above can be generalized to different Riemann surfaces. We will take the second one as a definition, as this is the most common definition in the literature. Moreover, it naturally leads to another key object in Teichmüller theory: the mapping class group.

DEFINITION 3.5.1. Let S be a surface of finite type. Then the *Teichmüller space* of S is defined as

$$\mathcal{T}(S) = \left\{ (X, f) : \begin{array}{l} X \text{ a Riemann surface, } f : S \rightarrow X \\ \text{an orientation preserving diffeomorphism} \end{array} \right\} / \sim,$$

where

$$(X, f) \sim (Y, g)$$

if and only if there exists a biholomorphism $h : X \rightarrow Y$ so that the map

$$g^{-1} \circ h \circ f : S \rightarrow S$$

is homotopic to the identity.

We will often write

$$\mathcal{T}(\Sigma_{g,n}) = \mathcal{T}_{g,n} \quad \text{and} \quad \mathcal{T}(\Sigma_g) = \mathcal{T}_g.$$

LECTURE 4

Markings, mapping class groups and moduli spaces

4.1. Teichmüller space in terms of markings

In order to get to the analogous definition to the space of marked tori, we need to single out particularly nice generating sets for the fundamental group, just like we did for tori. We will stick to closed surfaces. Recall that the fundamental group of a surface of genus g satisfies:

$$\pi_1(\Sigma_g, p) = \left\langle a_1, \dots, a_g, b_1, \dots, b_g \left| \prod_{i=1}^g [a_i, b_i] = e \right. \right\rangle.$$

In what follows, a generating set $A_1, \dots, A_g, B_1, \dots, B_g$ of $\pi_1(\Sigma_g, p)$ that satisfies

$$\prod_{i=1}^g [A_i, B_i] = e,$$

will be called a *canonical* generating set. Note that this includes the torus case.

DEFINITION 4.1.1. Let R be a closed Riemann surface.

- (1) A *marking* on R is a canonical generating set $\Sigma_p \subset \pi_1(R, p)$.
- (2) Two markings Σ_p and $\Sigma_{p'}$ are called *equivalent* if there exists a continuous curve α from p to p' so that the corresponding isomorphism $T_\alpha : \pi_1(R, p) \rightarrow \pi_1(R, p')$ satisfies

$$T_\alpha(\Sigma_p) = \Sigma_{p'}.$$

Two pairs (R, Σ) and (R', Σ') of marked closed Riemann surfaces are called *equivalent* if there exists a biholomorphic mapping $h : R \rightarrow R'$ so that

$$h_*(\Sigma) \simeq \Sigma'.$$

Just like in the case of the torus, the space of marked Riemann surfaces turns out to be the same as Teichmüller space:

THEOREM 4.1.2. Let S be a closed surface and Σ a marking on S . Then the map

$$\mathcal{T}(S) \rightarrow \left\{ (R, \Sigma_p) : \begin{array}{l} R \text{ a closed Riemann surface diffeomorphic to } S \\ p \in R, \Sigma_p \text{ a marking on } R \end{array} \right\} / \sim.$$

given by

$$[(R, f)] \mapsto [(R, f_*(\Sigma))]$$

is a bijection.

Before we sketch the proof of this theorem, we state and prove a lemma that will be of use in the study of mapping class groups as well:

LEMMA 4.1.3 (Alexander Lemma). *Let D be a 2-dimensional closed disk and $\phi : D \rightarrow D$ a homeomorphism that restricts to the identity on ∂D . Then ϕ is isotopic to the identity $D \rightarrow D$*

PROOF OF THE ALEXANDER LEMMA. Identify D with the closed unit disk in \mathbb{R}^2 and define the map $F : D \times [0, 1] \rightarrow D$ by

$$F_t(x) = \begin{cases} (1-t) \cdot \phi\left(\frac{x}{(1-t)}\right) & \text{if } \|x\| < 1-t \text{ and } t < 1 \\ x & \text{if } \|x\| > 1-t \text{ and } t < 1 \\ x & \text{if } t = 1. \end{cases}$$

This yields the isotopy we want. □

We can make this lemma work in the smooth category as well, but its proof is significantly less easy. It for instance follows from work by Smale [Sma59]. In this course we will generally gloss over the difference between homeomorphisms and diffeomorphisms.

PROOF SKETCH. Write $\Sigma = \{[A_1], \dots, [A_g], [B_1], \dots, [B_g]\}$, where A_i, B_i are simple closed curves based at a point $p_0 \in S$. Let us start with the injectivity. So, suppose

$$[(R, f_*(\Sigma))] = [(R', f'_*(\Sigma))].$$

This means that we can find a biholomorphic map $h : R \rightarrow R'$ and a self-diffeomorphism $g_0 : R' \rightarrow R'$ that is homotopic to the identity and such that

$$g_1 = g_0 \circ h \circ f$$

corresponds with f' on the curves $A_1, \dots, A_g, B_1, \dots, B_g$. The domain obtained by deleting these curves from S is a disk. This implies that f' and g_1 are homotopic (using the Alexander trick), which in turn means that

$$[(R, f)] = [(R', f')] \in \mathcal{T}(S).$$

For surjectivity, suppose we are given a marked Riemann surface (R, Σ_p) . So we need to find an orientation preserving homeomorphism $f : S \rightarrow R$ so that $f_*(\Sigma) = \Sigma_p$. So, let us take simple closed smooth curves $A'_1, \dots, A'_g, B'_1, \dots, B'_g$ such that

$$\Sigma_p = \{[A'_1], \dots, [A'_g], [B'_1], \dots, [B'_g]\}.$$

Moreover, we will set

$$C = \bigcup_{j=1}^g (A_j \cup B_j), \quad C' = \bigcup_{j=1}^g (A'_j \cup B'_j), \quad S_0 = S \setminus C, \quad \text{and} \quad R_0 = R \setminus C'.$$

R_0 and S_0 are diffeomorphic to polygons with $4g$ sides. So we can find a diffeomorphism by extending a diffeomorphism R_0, S_0 . For more details, see [IT92, Theorem 1.4]. □

4.1.1. Punctures and marked points. If $n \geq 1$, we can think of $\mathcal{T}(\Sigma_{g,n})$ as a space of surfaces with marked points (as opposed to punctures) as well:

PROPOSITION 4.1.4. *Let $n \geq 1$ and fix n distinct points $x_1, \dots, x_n \in \Sigma_g$. There is a bijection*

$$\mathcal{T}(\Sigma_{g,n}) \longrightarrow \{ (X, f) : f : \Sigma_g \rightarrow X \text{ an orientation preserving diffeomorphism} \} / \sim,$$

where $(X_1, f_1) \sim (X_2, f_2)$ if and only if there exists a biholomorphism $h : X_1 \rightarrow X_2$ such that

$$f_2^{-1} \circ h \circ f_1(x_i) = x_i \quad \text{for } i = 1, \dots, n$$

and $f_2^{-1} \circ h \circ f_1 : \Sigma_g \rightarrow \Sigma_g$ is homotopic to the identity through maps fixing x_1, \dots, x_n .

We leave the proof to the reader.

4.1.2. Basic examples. We have seen that the Teichmüller space of the torus can be identified with \mathbb{H}^2 (as a set for now). We will treat some further examples in this section.

PROPOSITION 4.1.5. *We have*

- (a) *Let S be diffeomorphic to $\Sigma_0, \Sigma_{0,1}, \Sigma_{0,2}$ or $\Sigma_{0,3}$, then $\mathcal{T}(S)$ is a point.*
- (b) *Let $\mathcal{T}(\Sigma_{1,1})$ can be identified with $\mathcal{T}(\Sigma_1)$.*

PROOF. For (a), suppose that $[X_1, f_1], [X_2, f_2] \in \mathcal{T}(\Sigma_{0,n})$ with $0 \leq n \leq 3$. We will think of these two as surfaces with marked points, coming from at most three marked points x_1, x_2, x_3 on \mathbb{S}^2 . By the uniformization theorem, we can identify X_1 and X_2 with the Riemann sphere $\hat{\mathbb{C}}$. Moreover (using that $n \leq 3$), there exists a Möbius transformation $\varphi : X_1 \rightarrow X_2$ such that

$$\varphi(f_1(x_i)) = f_2(x_i), \quad i = 1, \dots, n.$$

As such the diffeomorphism $f_2^{-1} \circ \varphi \circ f_1 : \mathbb{S}^2 \rightarrow \mathbb{S}^2$ fixes x_1, \dots, x_n . All we need to do, is show that this map is homotopic to the identity. If $n = 0$, we can perform a homotopy such that $f_2^{-1} \circ \varphi \circ f_1 : \mathbb{S}^2 \rightarrow \mathbb{S}^2$ fixes a point, which we shall call x_1 . This means that $f_2^{-1} \circ \varphi \circ f_1 : \mathbb{S}^2 \rightarrow \mathbb{S}^2$ can be restricted to a self homeomorphism of $\mathbb{S}^2 - \{x_1\} \simeq \mathbb{R}^2$, that we call $f : \mathbb{R}^2 \rightarrow \mathbb{R}^2$. The map $F : \mathbb{R}^2 \times [0, 1] \rightarrow \mathbb{R}^2$, defined by

$$F_t(x) = (1 - t) \cdot f(x) + t \cdot x$$

is a homotopy between f and the identity $\mathbb{R}^2 \rightarrow \mathbb{R}^2$. Because both $f_2^{-1} \circ \varphi \circ f_1$ and the identity fix $x_1 \in \mathbb{S}^2$, the homotopy above can be extended to \mathbb{S}^2 .

The proof for item (b) is similar. We again think in terms of surfaces with marked points. We have a surjective map

$$\mathcal{T}(\Sigma_{1,1}) \rightarrow \mathcal{T}(\Sigma_1),$$

mapping $[X, f] \in \mathcal{T}(\Sigma_{1,1})$ to $[X, f] \in \mathcal{T}(\Sigma_1)$. What we need to show is that this map is injective.

So, suppose $[X_1, f_1] = [X_2, f_2] \in \mathcal{T}(\Sigma_1)$. So there exists a biholomorphism $h : X_1 \rightarrow X_2$ such that $f_2^{-1} \circ h \circ f_1 : \Sigma_1 \rightarrow \Sigma_1$ is homotopic to the identity. We need to show that we can

modify h in such a way that $f_2^{-1} \circ h \circ f_1$ remains homotopic to the identity and also fixes our favorite point $x_1 \in \Sigma_1$. To this end, let's write $X_2 = \mathbb{C}/\Lambda$ for some lattice Λ . Suppose $[p], [q] \in X_2$. Observe that $h_0 : X_2 \rightarrow X_2$, defined by

$$h_0([z]) = h_0([z + q - p])$$

is a biholomorphic map $X_2 \rightarrow X_2$ that is homotopic to the identity and maps $[p]$ to $[q]$. So, if we set $[p] = h \circ f_1(x_1)$ and $[q] = f_2(x_1)$, then $h_0 \circ h : X_1 \rightarrow X_2$ is the biholomorphic map we're looking for. \square

4.2. The mapping class group

4.2.1. Defintion. Just like in the case of the torus, we have a natural group action on the Teichmüller space of a surface, by a group called the mapping class group:

DEFINITION 4.2.1. Let S_0 be a compact surface of finite type and $\Sigma \subset S_0$ a finite set. Set $S = S_0 \setminus \Sigma$. The *mapping class group* of S is given by

$$\text{MCG}(S) = \text{Diff}^+(S, \partial S, \Sigma) / \text{Diff}_0^+(S, \partial S, \Sigma)$$

where

$$\text{Diff}^+(S, \partial S, \Sigma) = \left\{ f : S_0 \rightarrow S_0 : \begin{array}{l} f \text{ an orientation preserving diffeomorphism that} \\ \text{acts as the identity on the boundary components} \\ \text{of } S_0 \text{ and preserves the elements of } \Sigma \text{ pointwise} \end{array} \right\}$$

and

$$\text{Diff}_0^+(S, \partial S, \Sigma) = \left\{ f \in \text{Diff}^+(S, \partial S, \Sigma) : \begin{array}{l} f \text{ homotopic to the identity} \\ \text{through a homotopy preserving} \\ \text{the elements of } \Sigma \text{ pointwise} \end{array} \right\}.$$

The group operation is induced by composition of functions.

Some authors let go of the condition that $\text{MCG}(S)$ fixes the punctures. The group we defined above is then often called the *pure mapping class group*.

4.3. Moduli space

Looking at Definition 3.5.1, we see there is a natural group action of the mapping class group of a surface on the corresponding Teichmüller space.

$$[g] \cdot [(R, f)] = [(R, f \circ g^{-1})].$$

The quotient is what will be called moduli space.

DEFINITION 4.3.1. Let S be a surface of finite type. The *moduli space* of S is the space

$$\mathcal{M}(S) = \mathcal{T}(S) / \text{MCG}(S).$$

We will write

$$\mathcal{M}(\Sigma_{g,n}) = \mathcal{M}_{g,n} \quad \text{and} \quad \mathcal{M}(\Sigma_g) = \mathcal{M}_g.$$

REMARK 4.3.2. Note that by using the convention that the mapping class group fixes boundary components and punctures, we leave these “marked”, i.e. if two surfaces are isometric, but any isometry between them permutes the punctures, these surfaces represent different points in moduli space.

4.4. Elements and examples of mapping class groups

4.4.1. **Basic examples.** We have:

PROPOSITION 4.4.1. *Let $n \leq 3$, then*

$$\mathrm{MCG}(\Sigma_{0,n}) = \{e\}.$$

PROOF. We start with the case $n \leq 1$. This is very similar to some of what we did in the proof of Proposition 4.1.5. Suppose $f : \Sigma_0 \rightarrow \Sigma_0$ is an orientation preserving diffeomorphism. We can (up to homotopy if $n = 0$) assume that f fixes a point $x \in \Sigma_0$. This allows us to restrict f to $\Sigma_0 - \{x\} \simeq \mathbb{R}^2$ and use a straight line homotopy to homotope $f|_{\mathbb{R}^2}$ to the identity. This extends to a homotopy between f and the identity on Σ_0 , because both fix x .

The proof of the cases $n \in \{2, 3\}$ is part of this week’s exercises. \square

4.4.2. **Dehn twists and the mapping class group of the annulus.** Before we move on, let us describe some non-trivial elements of the mapping class group. First, consider an annulus

$$A := [0, 1] \times \mathbb{R}/\mathbb{Z}.$$

Define a map $T : A \rightarrow A$ by

$$T(t, [\theta]) = (t, [\theta + t])$$

for all $t \in [0, 1]$, $\theta \in \mathbb{R}$. This map is called a Dehn twist. Note that this map fixes ∂A pointwise. Figure 1 shows that this map does to a segment connecting the two boundary components of the annulus.




FIGURE 1. A Dehn twist on an annulus.

Before we show how to turn T into a non-trivial element of a mapping class group of a different surface, we mention that T generates the mapping class group of the annulus:

PROPOSITION 4.4.2. *Let $A = [0, 1] \times \mathbb{R}/\mathbb{Z}$. Then*

$$\mathrm{MCG}(A) \simeq \mathbb{Z} = \langle [T] \rangle.$$

We will prove below..

Now let α be an essential (i.e. not homotopically trivial and not homotopic into a puncture or boundary component) simple closed curve on S . Let N be closed regular neighborhood of α . Identifying N with A , we can define a map $T_\alpha : S \rightarrow S$ by

$$T_\alpha(p) = \begin{cases} T(p) & \text{if } p \in N \\ p & \text{if } p \in S \setminus N \end{cases}$$

Because $T|_{\partial A}$ is the identity map, this is a continuous map. To obtain an element in $\mathrm{MCG}(S)$, we need to start with a smooth map. There are multiple ways out at the moment. We could smoothen T . Or we could use surface topology to argue that T_α is homotopic to a smooth map. Since for the mapping class group, we only care about diffeomorphisms up to homotopy, the element we get in $\mathrm{MCG}(S)$ will not depend on how we do this.

Figure 2 shows an example of a Dehn twist.




FIGURE 2. A Dehn twist on a surface of genus two.

We see that T_α maps a curve γ on the surface intersecting the defining curve α (of which we have only drawn the regular neighborhood) transversely to a curve that is not homotopic to γ . In particular, T_α is not homotopic to the identity and hence defines a non-trivial element in $\mathrm{MCG}(S)$.

PROOF OF PROPOSITION 4.4.2. We will first construct a homomorphism $\rho : \mathrm{MCG}(A) \rightarrow \mathbb{Z}$. Given an orientation preserving diffeomorphism $f : A \rightarrow A$ such that $f|_{\partial A} = \mathrm{Id}$, we can find a lift $\tilde{f} : [0, 1] \times \mathbb{R} \rightarrow [0, 1] \times \mathbb{R}$ such that $\tilde{f}(0, 0) = (0, 0)$. This means that

$$\tilde{f}|_{\{0\} \times \mathbb{R}} = \mathrm{Id}.$$

Because $f|_{\partial A} = \mathrm{Id}$, the restriction $\tilde{f}|_{\{1\} \times \mathbb{R}}$ is an integer translation. We let $\rho(f)$ be this integer.

ρ is surjective, because $\rho([T^n]) = n$. Now suppose $\rho([f]) = 0$. This means that \tilde{f} restricts to the identity on $\{0, 1\} \times \mathbb{R}$. We have that

$$\tilde{f}(n \cdot (t, x)) = f_*(n) \cdot \tilde{f}(t, x), \quad n \in \mathbb{Z}, (t, x) \in [0, 1] \times \mathbb{R},$$

where $f_* \in \mathrm{Aut}(\mathbb{Z}) = \{\pm \mathrm{Id}\}$. Because $\tilde{f}|_{\{0, 1\} \times \mathbb{R}} = \mathrm{Id}$, we need that $f_* = \mathrm{Id}$. Implying that

$$\tilde{f}(n \cdot (t, x)) = n \cdot \tilde{f}(t, x), \quad n \in \mathbb{Z}, (t, x) \in [0, 1] \times \mathbb{R}$$

and thus that the straight line homotopy

$$F_s((t, x)) = (1 - s) \cdot \tilde{f}(x, t) + s \cdot (x, t), \quad s \in [0, 1]$$

is a \mathbb{Z} -equivariant homotopy between \tilde{f} and the identity, that hence descends to A . This proves that ρ is injective and concludes the proof of the proposition. \square

4.4.3. The mapping class group of the torus. We briefly return to the torus. The question is of course whether the general definition on the mapping class group really corresponds to what happens in the case of the torus. We recall that

$$\mathcal{M}_1 = \mathrm{PSL}(2, \mathbb{Z}) \backslash \mathbb{H}^2.$$

This makes one wonder whether the mapping class group of the torus is maybe $\mathrm{PSL}(2, \mathbb{Z})$. This turns out to be almost correct. Indeed, we have the following theorem:

THEOREM 4.4.3. *We have*

$$\mathrm{MCG}(\mathbb{T}^2) \simeq \mathrm{SL}(2, \mathbb{Z}).$$

The action of $\mathrm{MCG}(\mathbb{T}^2)$ on \mathcal{T}_1 is that given by

$$\begin{pmatrix} a & b \\ c & d \end{pmatrix} \tau = \frac{a\tau - b}{-c\tau + d}$$

for all $\begin{pmatrix} a & b \\ c & d \end{pmatrix} \in \mathrm{SL}(2, \mathbb{Z})$ and $\tau \in \mathcal{T}_1$.

PROOF. We will identify

$$\mathbb{T}^2 = \mathbb{R}^2 / \mathbb{Z}^2$$

First observe that every element $A \in \mathrm{SL}(2, \mathbb{Z})$ induces a linear diffeomorphism $x \mapsto A \cdot x$ of \mathbb{R}^2 . Moreover, since $\mathrm{SL}(2, \mathbb{Z})$ preserves $\mathbb{Z}^2 \subset \mathbb{R}^2$, the action on \mathbb{R}^2 descends to an action by diffeomorphisms

$$\mathbb{T}^2 \rightarrow \mathbb{T}^2$$

that are orientation preserving because $\det(A) > 0$.

Our goal is to show that every orientation preserving diffeomorphism $\phi : \mathbb{T}^2 \rightarrow \mathbb{T}^2$ is homotopic to such a map. To this end, we may homotope ϕ so that it fixes $[0] \in \mathbb{T}^2$ and we can take a lift $\tilde{\phi} : \mathbb{R}^2 \rightarrow \mathbb{R}^2$ that fixes the origin of \mathbb{R}^2 . We have

$$\tilde{\phi}(x + (m, n)) = \tilde{\phi}(x) + \phi_*(m, n),$$

for all $(m, n) \in \mathbb{Z}^2$ where $\phi_* : \mathbb{Z}^2 \rightarrow \mathbb{Z}^2$ is an isomorphism, i.e. an element of $\mathrm{GL}(2, \mathbb{Z})$. For a general surface S , the map $[\phi] \in \mathrm{MCG}(S) \mapsto \phi_* \in \mathrm{Aut}(\pi_1(S))$ does not yield a homomorphism: we have chosen a homotopy to make ϕ fix a base point. Changing this choice a priori changes ϕ_* by an inner automorphism of $\pi_1(S)$. So we only obtain a map to $\mathrm{Out}(\pi_1(S))$. However, because \mathbb{Z}^2 is abelian, we have $\mathrm{Out}(\mathbb{Z}^2) \simeq \mathrm{Aut}(\mathbb{Z}^2)$. So in the case of the torus, we obtain a homomorphism $\mathrm{MCG}(\mathbb{T}^2) \rightarrow \mathrm{GL}(2, \mathbb{Z})$.

Write A_ϕ for the $\mathrm{GL}(2, \mathbb{Z})$ matrix corresponding to ϕ . Observe that

$$F_t(x) = tA_\phi \cdot x + (1 - t)\tilde{\phi}(x), \quad t \in [0, 1], x \in \mathbb{R}^2$$

gives a \mathbb{Z}^2 -equivariant homotopy between $\tilde{\phi}$ and the linear map $x \mapsto A_\phi \cdot x$. Since $\tilde{\phi}$ is orientation preserving, $\det(A_\phi) > 0$, and hence $A_\phi \in \mathrm{SL}(2, \mathbb{Z})$. So we obtain a map $\mathrm{MCG}(\mathbb{T}^2) \rightarrow \mathrm{SL}(2, \mathbb{Z})$. The map is surjective, because ϕ_A maps to A . Moreover, the map is injective, because if A_ϕ is the identity matrix, F_t gives a homotopy of $\tilde{\phi}$ to the identity. Since the action of $[\phi] \in \mathrm{MCG}(\mathbb{T}^2)$ on $\mathcal{T}(\mathbb{T}^2)$ is by precomposition with ϕ^{-1} , the action is as described. \square

REMARK 4.4.4. Note that the theorem above implies that the mapping class group action is not faithful. The kernel of the action is the center of $\mathrm{SL}(2, \mathbb{Z})$, i.e. the subgroup

$$\left\{ \pm \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} \right\} < \mathrm{SL}(2, \mathbb{Z}).$$

On the other hand, we do have

$$\mathbb{H}^2 / \mathrm{PSL}(2, \mathbb{Z}) = \mathbb{H}^2 / \mathrm{SL}(2, \mathbb{Z}).$$

This means that the mapping class group action is indeed a generalization of the situation for the torus case.

4.4.4. Mapping class groups in higher genus. We proved in the exercises that $\mathrm{SL}(2, \mathbb{Z})$ can be generated by the matrices

$$T = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix} \quad \text{and} \quad S = \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}.$$

We can also generate $\mathrm{SL}(2, \mathbb{Z})$ by the matrices

$$T = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix} \quad \text{and} \quad R = \begin{pmatrix} 1 & 0 \\ 1 & 1 \end{pmatrix}.$$

Indeed, a calculation shows that $S = T^{-1}RT^{-1}$.

Now identify $\mathbb{T}^2 = \mathbb{R}^2 / \mathbb{Z}^2$ again and write α and β for the closed curves in \mathbb{T}^2 that are the images of the straight line segments between the origin and $(0, 1)$ and $(1, 0)$ respectively. Tracing the proof of Theorem 4.4.3, we see that $T = [T_\alpha]$ and $R = [T_\beta]$. That is, $\mathrm{MCG}(\mathbb{T}^2)$ can be generated by two Dehn twists.

It actually turns out that an analogous statement holds for all mapping class groups. In the following theorem, a non-separating curve will be a curve α so that $S \setminus \alpha$ is connected. Figure 3 shows an example.




FIGURE 3. A separating curve (α) and a non-separating curve (β).

THEOREM 4.4.5 (Dehn - Lickorish theorem). *Let S be a surface of finite type, the mapping class group $\text{MCG}(S)$ is generated by finitely many Dehn twists about nonseparating simple closed curves.*

4.4.5. The action on homology. If S is a surface, then $\text{MCG}(S)$ acts on its homology $H_1(S)$. Indeed every diffeomorphism $f : S \rightarrow S$ induces an automorphism $f_* : H_1(S, \mathbb{Z}) \rightarrow H_1(S, \mathbb{Z})$. In this section, we briefly describe some aspects of this action. We will restrict to closed surfaces.

First of all, it turns out the action preserves some extra structure: the algebraic intersection number between oriented curves. In order to define it, let α and β be two oriented closed curves on an oriented surface S that intersect each other transversely at every intersection point. Then the *algebraic intersection number* between α and β is given by

$$i(\alpha, \beta) = \sum_{p \in \alpha \cap \beta} \text{sgn}(\omega(v_p(\alpha), v_p(\beta))),$$

where $\text{sgn} : \mathbb{R} \rightarrow \{\pm 1\}$ denotes the sign function, ω is any volume form that induces the orientation and $v_p(\alpha)$ and $v_p(\beta)$ denote the unit tangent vectors to α and β respectively at p . Note that

$$i(\beta, \alpha) = -i(\alpha, \beta).$$

Figure 4 shows an example of a positive contribution to the intersection number.




FIGURE 4. A positive contribution to $i(\alpha, \beta)$ if the orientation points out of the page.

We note that this form descends to homology. That is, it induces a form

$$i : H_1(S, \mathbb{Z}) \times H_1(S, \mathbb{Z}) \rightarrow \mathbb{Z}$$

called the *intersection form*, with the properties:

- (1) i is bilinear.
- (2) i is *alternating*, i.e.

$$i(a, b) = -i(b, a)$$

for all $a, b \in H_1(S, \mathbb{Z})$.

- (3) i is *non-degenerate*, i.e. if $a \in H_1(S, \mathbb{Z})$ is such that

$$i(a, b) = 0 \quad \text{for all } b \in H_1(S, \mathbb{Z})$$

then $a = 0$.

(see [FK92, Section III.1] for more details). Such a form is called a *symplectic form*.

First of all note that the image preserves the intersection form. Moreover, isotopic maps give rise to the same automorphism. So this gives us a representation

$$\text{MCG}(S) \rightarrow \text{Aut}(H_1(S, \mathbb{Z}), i)$$

called the *homology representation* of the mapping class group. Recall that if S is a closed orientable surface of genus g , then $H_1(S, \mathbb{Z}) \simeq \mathbb{Z}^{2g}$. Choosing an identification, the homology representation becomes a map

$$\mathrm{MCG}(S) \rightarrow \mathrm{Sp}(2g, \mathbb{Z}) = \{ A \in \mathrm{Mat}_{2g}(\mathbb{Z}) : i(Av, Aw) = i(v, w), \forall v, w \in \mathbb{Z}^{2g} \}.$$

It turns out that this representation is surjective (this can be proved using a finite generating set for $\mathrm{Sp}(2g, \mathbb{Z})$ consisting of transvections, which can be realized by Dehn twists), but generally highly non-injective. A notable exception is the case of the torus, there is an isomorphism

$$\mathrm{Sp}(2, \mathbb{Z}) \simeq \mathrm{SL}(2, \mathbb{Z})$$

and indeed the the homology representation $\mathrm{MCG}(\mathbb{T}^2) \rightarrow \mathrm{Sp}(2, \mathbb{Z})$ is an isomorphism.

LECTURE 5

Beltrami differentials and a topology on Teichmüller space

Our next goal is to put a topology on Teichmüller space (that coincides with the topology of the upper half plane in the case of the torus and the once punctured torus). Of course, the goal is that marked complex structures that are more similar should be closer to each other. So that raises the question how one measures the how similar two marked complex structures are.

5.1. Wirtinger derivatives

Let us first recall the notation for differentiation with respect to complex coordinates. If (U, z) is a complex coordinate chart on a surface S , we wrote $z = x + iy$ and

$$\frac{\partial}{\partial z} = \frac{1}{2} \left(\frac{\partial}{\partial x} - i \frac{\partial}{\partial y} \right) \quad \text{and} \quad \frac{\partial}{\partial \bar{z}} = \frac{1}{2} \left(\frac{\partial}{\partial x} + i \frac{\partial}{\partial y} \right).$$

Moreover, the equation $\frac{\partial f(z)}{\partial \bar{z}} = 0$ is equivalent to the Cauchy-Riemann equations for $f(z)$, i.e. this equation is equivalent to f being holomorphic. One readily verifies that in these coordinates, the product rule takes the form

$$\frac{\partial}{\partial z}(f \cdot g) = \frac{\partial f}{\partial z} \cdot g + f \cdot \frac{\partial g}{\partial z}.$$

Moreover, the chain rules read

$$\frac{\partial(f \circ g)}{\partial z}(z_0) = \frac{\partial f}{\partial z}(g(z_0)) \cdot \frac{\partial g}{\partial z}(z_0) + \frac{\partial f}{\partial \bar{z}}(g(z_0)) \cdot \frac{\partial \bar{g}}{\partial z}(z_0)$$

and

$$\frac{\partial(f \circ g)}{\partial \bar{z}}(z_0) = \frac{\partial f}{\partial z}(g(z_0)) \cdot \frac{\partial g}{\partial \bar{z}}(z_0) + \frac{\partial f}{\partial \bar{z}}(g(z_0)) \cdot \frac{\partial \bar{g}}{\partial \bar{z}}(z_0).$$

5.2. Beltrami coefficients

In Section 2.3 on conformal structures we saw that when trying to find an isothermal complex coordinate w on a Riemannian surface, equipped with a coordinate z , the fraction

$$\mu = \frac{\partial w / \partial \bar{z}}{\partial w / \partial z}$$

shows up as a measure of “how much” we need to “correct” the coordinate z in order to make the metric conformal to the standard Euclidean metric in w (and thus yielding holomorphic chart transitions).

In general, given a smooth map $f : S \rightarrow R$ between Riemann surfaces, we can use this expression to measure how far f is from holomorphic. Let us first prove that this yields a well defined object:

LEMMA 5.2.1. *Let S and R be Riemann surfaces and $f : S \rightarrow R$ a smooth map. Suppose that (U, z) is a holomorphic local coordinate on S and (V, w) one on R . Then the smooth function $\mu : z(U) \rightarrow \mathbb{C}$ defined by*

$$\mu = \left(\frac{\partial F}{\partial \bar{z}} \right) / \left(\frac{\partial F}{\partial z} \right),$$

where $F = w \circ f \circ z^{-1} : z(U) \rightarrow \mathbb{C}$ is independent of the choice of coordinate (V, w) .

PROOF. Suppose (V', w') is a different holomorphic local coordinate with $f(U) \subset V'$. Write $F' = w' \circ f \circ z^{-1} : z(U) \rightarrow \mathbb{C}$. By the chain rule,

$$\frac{\partial F'}{\partial \bar{z}} = \frac{\partial(w' \circ w^{-1} \circ F)}{\partial \bar{z}} = \left(\frac{\partial(w' \circ w^{-1})}{\partial z} \circ F \right) \cdot \frac{\partial F}{\partial \bar{z}},$$

where the second term disappears because $w' \circ w^{-1}$ is holomorphic. Likewise,

$$\frac{\partial F'}{\partial z} = \frac{\partial(w' \circ w^{-1} \circ F)}{\partial z} = \left(\frac{\partial(w' \circ w^{-1})}{\partial z} \circ F \right) \cdot \frac{\partial F}{\partial z}.$$

So when we divide the two, we obtain the same μ . □

Observe that μ *does* depend on the local coordinate (U, z) . Indeed, if (U', z') is a different holomorphic local coordinate and we write $F' = w \circ F \circ (z')^{-1}$, then

$$\begin{aligned} \frac{\partial F'}{\partial z'} &= \frac{\partial(F \circ z \circ (z')^{-1})}{\partial z'} = \left(\frac{\partial F}{\partial \bar{z}} \circ z \circ (z')^{-1} \right) \cdot \frac{\partial(\overline{z \circ (z')^{-1}})}{\partial z'} \\ &= \left(\frac{\partial F}{\partial \bar{z}} \circ z \circ (z')^{-1} \right) \cdot \overline{\left(\frac{\partial(z \circ (z')^{-1})}{\partial z'} \right)}, \end{aligned}$$

again using the fact that the coordinate change is holomorphic to conclude that the other term disappears, and

$$\frac{\partial F'}{\partial z'} = \frac{\partial(F \circ z \circ (z')^{-1})}{\partial z'} = \left(\frac{\partial F}{\partial z} \circ z \circ (z')^{-1} \right) \cdot \frac{\partial(z \circ (z')^{-1})}{\partial z'},$$

where we have used that

$$\frac{\partial(\overline{z \circ (z')^{-1}})}{\partial z'} = \overline{\left(\frac{\partial(z \circ (z')^{-1})}{\partial \bar{z}'} \right)} = 0.$$

In conclusion

$$\mu(z') = \mu(z) \cdot \overline{\left(\frac{\partial(z \circ (z')^{-1})}{\partial z'} \right)} / \frac{\partial(z \circ (z')^{-1})}{\partial z'}$$

and thus a smooth map $f : S \rightarrow R$ yields a differential μ_f of type $(-1, 1)$ associated to f . That is, it makes sense to write

$$\mu_f = \mu(z) \cdot \frac{d\bar{z}}{dz}.$$

So, we have a well defined differential μ_f associated to f that satisfies

$$\mu_f = 0 \quad \Leftrightarrow \quad f \text{ is biholomorphic.}$$

Indeed, if $\mu_f = 0$ then $\partial f / \partial \bar{z} = 0$ everywhere, so f is holomorphic. Since f is also invertible (and the inverse of a bijective holomorphic function is holomorphic), f is biholomorphic. We will call μ_f the *Beltrami coefficient* of f .

We also observe that the coordinate transition above implies that $z \mapsto |\mu_f(z)|$ is a well defined function on S . Moreover, it is uniformly bounded by 1 for all orientation preserving diffeomorphisms:

LEMMA 5.2.2. *Let S and R be Riemann surface and let $f : S \rightarrow R$ be an orientation preserving diffeomorphism, then*

$$|\mu_f(z)| < 1$$

for all $z \in S$.

PROOF. Since this is a local question, we can just think of f as map $U \rightarrow V$, where $U, V \subset \mathbb{C}$ are open subsets. Writing $f(x, y) = u(x, y) + iv(x, y)$, the fact that f is orientation preserving means that

$$\begin{aligned} 0 < \det(Df(x, y)) &= \det \begin{pmatrix} \partial u(x, y) / \partial x & \partial u(x, y) / \partial y \\ \partial v(x, y) / \partial x & \partial v(x, y) / \partial y \end{pmatrix} \\ &= \frac{\partial u(x, y)}{\partial x} \frac{\partial v(x, y)}{\partial y} - \frac{\partial u(x, y)}{\partial y} \frac{\partial v(x, y)}{\partial x} = \left| \frac{\partial f}{\partial z} \right|^2 - \left| \frac{\partial f}{\partial \bar{z}} \right|^2, \end{aligned}$$

hence proving the lemma. \square

A differential μ of type $(-1, 1)$ on a Riemann surface X whose L^∞ -norm satisfies

$$\|\mu\|_\infty = \sup_{z \in X} \{|\mu(z)|\} < 1$$

is called a *Beltrami differential*.

The following transformation formula for Beltrami coefficients will also be useful to us:

LEMMA 5.2.3. *Let S , X_1 and X_2 be Riemann surfaces and let*

$$S \xrightarrow{f} X_1 \xrightarrow{g} X_2$$

be orientation preserving diffeomorphisms. Then

$$\mu_g \circ f = \left(\frac{\partial f}{\partial z} / \overline{\left(\frac{\partial f}{\partial z} \right)} \right) \cdot \frac{\mu_{g \circ f} - \mu_f}{1 - \bar{\mu}_f \cdot \mu_{g \circ f}}.$$

PROOF. This is part of this week's exercises. \square

Indeed, we can derive from it that Beltrami coefficients can recognize biholomorphisms:

LEMMA 5.2.4. *Let X_1 and X_2 be Riemann surfaces and let*

$$f_1 : R \rightarrow X_1 \quad \text{and} \quad f_2 : R \rightarrow X_2$$

be orientation preserving diffeomorphisms. Then the map $f_2 \circ f_1^{-1} : X_1 \rightarrow X_2$ is biholomorphic if and only if

$$\mu_{f_1} = \mu_{f_2}.$$

PROOF. $f_2 \circ f_1^{-1}$ is biholomorphic if and only if $\mu_{f_2 \circ f_1^{-1}} = 0$ as a Beltrami differential on X_1 . This is true if and only if, as a Beltrami differential on S ,

$$0 = \mu_{f_2 \circ f_1^{-1}} \circ f_1 = \frac{\partial f_1 / \partial z}{\partial f_1 / \partial \bar{z}} \cdot \frac{\mu_{f_2} - \mu_{f_1}}{1 - \mu_{f_1} \cdot \mu_{f_2}},$$

where we have used the previous lemma. Because neither the first factor on the right, nor the denominator (because $|\mu_f| < 1$ for an orientation preserving diffeomorphism) can be 0, we obtain that the equation $\mu_{f_2} = \mu_{f_1}$ is equivalent to $\mu_{f_2 \circ f_1^{-1}} \circ f_1$ being 0. \square

5.2.1. Topologizing Teichmüller space. We are going to use Beltrami coefficients to topologize Teichmüller space. First of all, the following theorem implies that Riemann surface structures up to homotopy can be recognized using Beltrami differentials:

THEOREM 5.2.5. *Let S , X_1 and X_2 be Riemann surfaces and*

$$f_1 : S \rightarrow X_1 \quad \text{and} \quad f_2 : S \rightarrow X_2$$

be orientation preserving diffeomorphisms. Then there exists a biholomorphic mapping

$$h : X_1 \rightarrow X_2$$

if and only if

$$\mu_{f_1} = \mu_{f_2 \circ \varphi^{-1}}$$

for some $\varphi \in \text{Diff}^+(S)$. Moreover, the map

$$(f_2)^{-1} \circ h \circ f_1 : S \rightarrow S$$

is homotopic to the identity if and only if $\varphi \in \text{Diff}_0^+(S, \Sigma)$.

PROOF. First suppose that there exists a biholomorphic map $h : X_1 \rightarrow X_2$. Then we set

$$\varphi = (f_2)^{-1} \circ h \circ f_1 : S \rightarrow S.$$

Then

$$\mu_{f_2} = \mu_{h \circ f_1 \circ \varphi^{-1}} = \mu_{f_1 \circ \varphi^{-1}},$$

where we have used Lemma 5.2.3 for the second equality. Since $\varphi = (f_2)^{-1} \circ h \circ f_1$, the second claim is immediate.

Conversely, suppose there exists some $\varphi \in \text{Diff}^+(S)$ such that

$$\mu_{f_1} = \mu_{f_2 \circ \varphi^{-1}}$$

Lemma 5.2.3 then shows that $h = f_2 \circ \varphi \circ f_1^{-1} : X_1 \rightarrow X_2$ is biholomorphic. Again, the second claim follows from the form of φ . Indeed $f_2^{-1} \circ h \circ f_1 = \varphi$. \square

So, given a Riemann surface S , we can define the space

$$B(S)_1 := \left\{ \begin{array}{l} f : S \rightarrow R \text{ an orientation} \\ \mu_f : \text{preserving diffeomorphism,} \\ R \text{ a Riemann surface} \end{array} \right\}$$

That we equip with the L^∞ topology. This space admits an action of the group of orientation preserving diffeomorphisms $\text{Diff}^+(S)$ by

$$\varphi \cdot \mu_f = \mu_{f \circ \varphi^{-1}}, \quad \varphi \in \text{Diff}^+(S), \mu_f \in B(S)_1.$$

A direct consequence is the following:

COROLLARY 5.2.6. *The map from the set of marked Riemann surfaces defined by*

$$(R, f) \mapsto \mu_f$$

induces a bijections

$$\mathcal{T}(S) \rightarrow B(S)_1 / \text{Diff}_0^+(S, \Sigma)$$

and

$$\mathcal{M}(S) \rightarrow B(S)_1 / \text{Diff}^+(S, \Sigma).$$

In particular, since $B(S)_1$ is a topological space, these bijections equip $\mathcal{T}(S)$ and $\mathcal{M}(S)$ with a topology. It is not so hard to see that the choice of Riemann surface structure on S does not influence the topology on Teichmüller space. Indeed, if S and S' are two Riemann surfaces and $g : S' \rightarrow S$ is any orientation preserving diffeomorphism between them, then

$$[X, f] \in \mathcal{T}(S) \mapsto [X, f \circ g] \in \mathcal{T}(S')$$

is a homeomorphism.

5.3. Quasiconformal mappings

We will now first describe another viewpoint on the same topology: namely that of quasiconformal mappings and the Teichmüller metric.

Let us consider what it is that $|\mu_f(z)|$ measures. Let $f : D \rightarrow D$ denote an orientation preserving diffeomorphism of some domain $D \subset \mathbb{C}$ containing 0. And write

$$Df(0) \cdot z = \frac{\partial f}{\partial z}(0) \cdot z + \frac{\partial f}{\partial \bar{z}}(0) \cdot \bar{z} =: a \cdot z + b \cdot \bar{z}$$

denote the first order Taylor expansion of f at 0 (so $a = \frac{\partial f}{\partial z}(0)$ and $b = \frac{\partial f}{\partial \bar{z}}(0)$).

Let us consider the inverse image of a unit circle under this map. I.e. the solutions of the equation

$$|Df(0) \cdot z| = 1.$$

Let us write

$$z = r e^{i\theta}, \quad a = |a| e^{i\alpha}, \quad b = |b| e^{i\beta}.$$

The equation then becomes

$$r \cdot \left| (|a| + |b|) \cos \left(\theta + \frac{\alpha - \beta}{2} \right) + i(|a| - |b|) \sin \left(\theta + \frac{\alpha - \beta}{2} \right) \right| = 1.$$

This is the equation of an ellipse with

- minor axis at polar angle $\frac{\beta-\alpha}{2}$ of half length $\frac{1}{|a|+|b|}$
- major axis at polar angle $\frac{\beta-\alpha+\pi}{2}$ of half length $\frac{1}{|a|-|b|}$.

Note that the latter is positive since the Jacobian

$$\det(Df(0)) = |a|^2 - |b|^2 > 0,$$

as we have seen in the proof of Lemma 5.2.2. Figure 1 shows a picture of the situation.




FIGURE 1. An ellipse that gets mapped to a circle.

The ratio of the axes of the ellipse is

$$\frac{|a| + |b|}{|a| - |b|} = \frac{1 + |\mu_f(0)|}{1 - |\mu_f(0)|}.$$

In particular, if $\mu_f(0)$ is close to 1, then the preimage is very far away from a circle and if $\mu_f(0) = 0$ (i.e. if f is biholomorphic) then the preimage is a circle. So this means that

- Biholomorphic maps locally send circles to circles (i.e. they are conformal)
- The $\mu_f(0)$ measures how far away the preimage of a circle (an ellipse) is from a circle.

For this reason, the number

$$K_f(z) = \frac{1 + |\mu_f(z)|}{1 - |\mu_f(z)|}$$

is called the *quasiconformal dillatation* of f at z . Since the ratio between the major and minor axis of the ellipse is a measure of how far f is from being conformal, this leads to the following definition:

DEFINITION 5.3.1. Let X and Y be Riemann surfaces and let $f : X \rightarrow Y$ be an orientation preserving diffeomorphism. We say $f : X \rightarrow Y$ is a *K -quasiconformal mapping* if

$$K \geq K_f := \sup_{z \in X} \left\{ \frac{1 + |\mu_f(z)|}{1 - |\mu_f(z)|} \right\}.$$

K_f is called the quasiconformal dillatation of f .

REMARK 5.3.2. The notion of a quasiconformal map can be generalized to maps of much lower regularity (see eg. [IT92, Chapter 4]). We will consider slightly lower regularity (but not as low as can be found in the literature) soon as well.

We record two useful properties of quasiconformal maps in a lemma:

LEMMA 5.3.3. *Suppose X , Y and Z are Riemann surfaces and $f : X \rightarrow Y$ and $g : Y \rightarrow Z$ are orientation preserving diffeomorphisms. Then the following holds:*

(a) *We have that*

$$K_f \geq 1$$

with equality if and only if f is a biholomorphism.

(b) *We have that*

$$K_{g \circ f} \leq K_g \cdot K_f.$$

(c) *Finally,*

$$K_{f^{-1}} = K_f.$$

PROOF. This is part of this week's exercises. □

5.4. The Teichmüller metric

Lemma 5.3.3 implies that the following definition defines a metric:

DEFINITION 5.4.1. Let S be a Riemann surface. Then the *Teichmüller distance* between $[X_1, f_1], [X_2, f_2] \in \mathcal{T}(S)$ is

$$d_T([X_1, f_1], [X_2, f_2]) = \frac{1}{2} \log \left(\inf \left\{ K_g : \begin{array}{l} g : X_1 \rightarrow X_2 \text{ an orientation} \\ \text{preserving diffeomorphism} \\ \text{homotopic to } f_2 \circ f_1^{-1} \end{array} \right\} \right).$$

The factor $\frac{1}{2}$ is a convention. The first thing to observe is that the Teichmüller distance is a metric:

LEMMA 5.4.2. *Let S be a Riemann surface. Then the Teichmüller distance $d_T : \mathcal{T}(S) \times \mathcal{T}(S) \rightarrow [0, \infty)$ defines a metric.*

PROOF SKETCH. The fact that d_T is symmetric and satisfies the triangle inequality are direct from Lemma 5.3.3. In order to show non-degeneracy, suppose

$$d_T([X_1, f_1], [X_2, f_2]) = 0.$$

There are two ways to show that this implies that $[X_1, f_1] = [X_2, f_2]$:

- We can use Teichmüller's theorem (that we will state, but not prove later on in the course), which states that there is a map (of slightly lower regularity) that realizes d_T . This map must have $K_g = 1$ and hence is a biholomorphism.

- We can run an approximation argument. Suppose $g_n : X_1 \rightarrow X_2$ is a sequence of maps in the homotopy class of $f_2 \circ f_1^{-1}$ such that

$$\frac{1}{2} \log(K_{g_n}) \xrightarrow{n \rightarrow \infty} d_T([X_1, f_1], [X_2, f_2]) = 0.$$

This means that

$$K_{g_n} \xrightarrow{n \rightarrow \infty} 1,$$

which in turn implies that g_n locally uniformly converges to a holomorphic map $X_1 \rightarrow X_2$ (see [IT92, Proposition 4.36] for details). □

We could have used d_T to topologize Teichmüller space as well:

LEMMA 5.4.3. *The Teichmüller metric is compatible with the topology on $\mathcal{T}(S)$.*

PROOF. This is a matter of tracing the definitions. Two points in $[X_1, f_1], [X_2, f_2] \in \mathcal{T}(S)$ are close if we can make $\mu_{f_2 \circ f_1^{-1}}$ close to 0 in the L^∞ topology by precomposing f_1 with a homotopically trivial self-diffeomorphism $X_1 \rightarrow X_1$. This is the same as saying that K_g is small for g in the homotopy class of $f_2 \circ f_1^{-1}$. □

5.5. Grötsch's theorem

Teichmüller proved that the Teichmüller distance is realized by a unique map called the *Teichmüller map* (that has slightly lower regularity than the maps we're taking the infimum over, but we'll get to that).

As a warm up, we'll consider *Grötsch's problem*, that goes as follows. Suppose that $R_1 = [0, a] \times [0, 1]$ and $R_2 = [0, K \cdot a] \times [0, 1]$ are two rectangles in the plane, where $a > 0$ and $K \geq 1$. We can ask what the minimal quasi-conformal dilatation is of a map $R_1 \rightarrow R_2$ that preserves vertical sides and horizontal sides. Grötsch proved the following:

THEOREM 5.5.1. (*Grötsch's theorem*) *Suppose R_1 and R_2 are as above and $f : R_1 \rightarrow R_2$ is a homeomorphism that is smooth and orientation preserving away from a finite number of points. Then*

$$K_f \geq K$$

with equality if and only if f is the affine map

$$(x, y) \in R_1 \mapsto (K \cdot x, y) \in R_2.$$

Before we prove the theorem, we briefly comment on definitions. Formally, we defined K_f only for orientation preserving diffeomorphisms. In particular, in the points $z \in R_1$ where f is not smooth, $\mu_f(z)$ is not defined. We can remedy this by simply ignoring these points in the supremum.

PROOF SKETCH. Writing $K_f(x, y)$ for the quasiconformal dilatation of f at $(x, y) \in R_1$, we first claim that

$$(5.5.1) \quad \left| \frac{\partial f}{\partial x}(x, y) \right|^2 \leq K_f(x, y) \cdot \det(J_f(x, y)),$$

where $J_f(x, y)$ denotes the Jacobian matrix of f at $(x, y) \in R_1$. Indeed, writing

$$M = \sup \{ |df_{(x,y)}(v)| : v \in T_{(x,y)}^1 R_1 \}$$

and

$$m = \inf \{ |df_{(x,y)}(v)| : v \in T_{(x,y)}^1 R_1 \},$$

We have $K_f(x, y) = M/m$, $\det(J_f(x, y)) = M \cdot m$ and $\left| \frac{\partial f}{\partial x}(x, y) \right|^2 \leq M^2$, which proves (5.5.1).

The second inequality we will use is

$$(5.5.2) \quad \int_{R_1} \left| \frac{\partial f}{\partial x}(x, y) \right| dx dy \geq K \cdot \text{area}(R_1)$$

This is the observation that for almost all $y \in [0, 1]$, $\int_0^a \left| \frac{\partial f}{\partial x}(x, y) \right| dx \geq K \cdot a$, by the substitution rule. Integrating with respect to y gives the desired inequality.

We now have

$$\begin{aligned} (K \cdot \text{area}(R_1))^2 &\stackrel{(5.5.2)}{\leq} \left(\int_{R_1} \left| \frac{\partial f}{\partial x}(x, y) \right| dx dy \right)^2 \\ &\stackrel{(5.5.1)}{\leq} \left(\int_{R_1} \sqrt{K_f(x, y)} \cdot \sqrt{\det(J_f(x, y))} dx dy \right)^2 \\ &\stackrel{\text{Cauchy-Schwarz}}{\leq} \int_{R_1} K_f(x, y) dx dy \cdot \int_{R_1} \det(J_f(x, y)) dx dy \\ &\leq \text{area}(R_2) \cdot K_f \cdot \text{area}(R_1) \\ &= K \cdot \text{area}(R_1) \cdot K_f \cdot \text{area}(R_1). \end{aligned}$$

Thus yielding that $K_f \geq K$.

Equality for the affine map is a matter of computation, so now we need to show uniqueness. This can be proved using the fact that we need to have equality in all the inequalities above. \square

LECTURE 6

Measured foliations and quadratic differentials

6.1. Measured foliations

Teichmüller's characterization of maps realizing the Teichmüller distance is a generalization of Grötsch's theorem: optimizers are analogues of affine maps. In order to define what these are, we need to discuss measured foliations and quadratic differentials. We will start with the former.

6.1.1. The torus. We start with our favorite example: the torus $\mathbb{T}^2 = \mathbb{R}^2/\mathbb{Z}^2$. Given any straight line ℓ through $(0,0)$, we obtain a foliation $\tilde{\mathcal{F}}_\ell$ of \mathbb{R}^2 consisting of all lines (that we will call the *leaves* of $\tilde{\mathcal{F}}_\ell$) parallel to ℓ .

Because $\tilde{\mathcal{F}}_\ell$ is invariant under the \mathbb{Z}^2 -action, it defines a foliation \mathcal{F}_ℓ on \mathbb{T}^2 . We observe that if ℓ has a rational slope, then all the leaves of \mathcal{F}_ℓ are simple closed curves. If the slope is irrational, every leaf of \mathcal{F}_ℓ is dense in \mathbb{T}^2 .

The foliation \mathcal{F}_ℓ also naturally comes with a measure on trajectories that are transverse to it. Indeed, let

$$\nu_\ell : \mathbb{R}^2 \rightarrow \mathbb{R}$$

denote the (signed) distance to ℓ . The associated 1-form $d\nu_\ell$ is \mathbb{Z}^2 -invariant and hence descends to a 1-form w_ℓ on \mathbb{T}^2 . This induces a measure μ_ℓ on arcs α transverse to \mathcal{F}_ℓ , given by

$$\mu_\ell(\alpha) = \int_\alpha |w_\ell|.$$

By construction $\mu_\ell(\alpha)$ is invariant under isotopies of α that preserve the leaf that each point of α lies in. μ_ℓ is called a *transverse measure* to \mathcal{F}_ℓ . The pair $(\mathcal{F}_\ell, \mu_\ell)$ is called a *measured foliation*.

Observe that the 1-form w_ℓ completely determines the measured foliation: the leaves of \mathcal{F}_ℓ are the integral submanifolds corresponding to the subbundle $\ker(w_\ell)$ of the tangent bundle $T(\mathbb{R}^2/\mathbb{Z}^2)$.

We have seen that every mapping class $[\phi] \in \text{MCG}(\mathbb{T}^2)$ admits a representative of the form

$$\phi_A : [z] \mapsto [A \cdot z], \quad [z] \in \mathbb{T}^2,$$

for $A \in \text{SL}(2, \mathbb{Z})$. If A is hyperbolic, i.e. has two real eigenvalues $\lambda > 1$ and $1/\lambda$. The matrix

$$\begin{pmatrix} 2 & 1 \\ 1 & 1 \end{pmatrix}$$




FIGURE 1. A 3-pronged singularity.

is everyone's favorite example. Then A has two eigenspaces $E^+, E^- \subset \mathbb{R}^2$ respectively. These eigenspaces define two transverse measured foliations (\mathcal{F}^+, μ^+) and (\mathcal{F}^-, μ^-) on \mathbb{T}^2 . The map ϕ_A preserves these foliations and

$$(\phi_A)_*\mu^+ = \lambda \cdot \mu^+ \quad \text{and} \quad (\phi_A)_*\mu^- = \frac{1}{\lambda} \cdot \mu^-.$$

So we can think of A as stretching by a factor λ in the direction of \mathcal{F}^+ and contracting by a factor $1/\lambda$ in the direction of \mathcal{F}^- . As such, this is a natural analogue to affine maps on rectangles.

6.1.2. Higher complexity. For surfaces of higher complexity, the situation is more complicated. First of all, such surfaces do not admit smooth foliations, so we have to allow for singularities for the theory to work.

DEFINITION 6.1.1. Let S be a closed surface. A *singular foliation* \mathcal{F} on a S is a decomposition of S into a disjoint union of subsets of S , called the *leaves* of \mathcal{F} , and a finite set of points of S , called *singular points* of \mathcal{F} , such that the following conditions hold:

- (1) Around each nonsingular point $p \in S$, there is a smooth chart to \mathbb{R}^2 that takes leaves to horizontal line segments. The transition maps between any two of these charts are smooth maps that take horizontal lines to horizontal lines.
- (2) Around each singular point $p \in S$, there is a smooth chart to \mathbb{R}^2 that takes leaves to the level sets of a k -pronged singularity for some $k \geq 3$ (Figure 1 shows an example).

We have the following formula for the relation between the singularities and the Euler characteristic:

PROPOSITION 6.1.2 (Euler–Poincaré formula). *Given a singular foliation \mathcal{F} on a closed surface S , let P_s denote the number of prongs of the singular point $s \in S$. Then the Euler characteristic of S can be computed as*

$$2\chi(S) = \sum_{\substack{s \in S \\ \text{singular}}} 2 - P_s.$$

PROOF. The proof will be part of this week's exercises. □

In particular, if S has negative Euler characteristic, any foliation of it necessarily has singular points, thus justifying our claim above.




FIGURE 2. A foliation on a surface of genus two.

In order to define transverse measures, we need a notion of arcs *transverse* to a singular foliation. These are smooth arcs on the surface that do not go through the singular points of \mathcal{F} and are transverse to the leaves of \mathcal{F} .

DEFINITION 6.1.3. Let S be a closed surface and let \mathcal{F} be a singular foliation of S . A *transverse measure* μ on \mathcal{F} is an assignment of a positive real number $\mu(\alpha)$ to each arc α that is transverse to \mathcal{F} such that

- μ is invariant under leaf-preserving isotopy of smooth arcs
- μ is *regular*, that is, there exists a smooth chart to \mathbb{R}^2 around each non-singular point that maps the leaves of \mathcal{F} to horizontal lines and such that μ is induced by the line element $|dy|$.

The pair (\mathcal{F}, μ) is called a *measured foliation* on S .

6.1.3. Example 1: polygons. It's high time for an example. Let's start with a concrete one. Figure 2 shows a foliated octagon. We can turn this octagon into a surface X of genus 2 by using horizontal and vertical translations to identify its sides (an Euler characteristic computation shows that the resulting genus is indeed 2).

Moreover, because the foliation is translation invariant, it glues into a singular foliation \mathcal{F} of X . The singularity can be found at the image of the vertices of the polygon (that are all identified in a single vertex on X).

We claim that the singularity around this vertex is 6-pronged. One way to see this, is by using Euclidean geometry. Indeed, the surface X comes equipped with the structure of a Euclidean metric that is defined everywhere except at the image of the vertex. Indeed, adding up the angles in the corners of the octagon, we see that the total angle around the vertex is 6π instead of 2π . Every prong in the singularity adds an angle π (a 2-pronged singularity being a regular point). So, we indeed have a 6-pronged singularity.

A second way to see this, is just by using the Euler–Poincaré formula. If $s \in X$ denotes the image of the vertex, then we have

$$-4 = 2\chi(X) = 2 - P_s,$$

thus yielding $P_s = 6$.

Again the line-segment coming from the distance to any of the leaves yields a measure μ on \mathcal{F} .

Finally, none of this really used our explicit octagon (except of course the calculation of the number of prongs of the singularity). So in general, we can build a surface X , by taking a Euclidean polygon $P \subset \mathbb{R}^2$ and identifying pairs of parallel sides of the same length. After that, we take any foliation of \mathbb{R}^2 by straight lines, from which we obtain a singular foliation on X . Moreover, instead of only translations, we can also allow rotations of angle π . These also preserve the foliation. They however don't preserve any orientation on the leaves, so we obtain a non-orientable foliation.

6.1.4. Example 2: natural coordinates. Suppose our surface S and $P \subset S$ is a finite set of points. If the surface $S - P$ admits an atlas for which all the transition maps are of the form:

$$(x, y) \mapsto (f(x, y), c \pm y)$$

for some smooth $f : \mathbb{R}^2 \rightarrow \mathbb{R}$ and some constant $c \in \mathbb{R}$ (both depending on the pair of charts). That is, the transition maps send horizontal lines in \mathbb{R}^2 to horizontal lines. This means that the horizontal foliation of \mathbb{R}^2 induces a singular foliation on S . Since the line element $|dy|$ is preserved as well, this foliation comes with a measure too.

Let's again look at our surface X from the previous example. We now equip it with the horizontal foliation. Since the gluings are by translations, we can equip $X - s$ with an atlas for which the transition maps all take the form

$$(x, y) \mapsto (x + c_1, y + c_2)$$

for some $c_1, c_2 \in \mathbb{R}$. In particular, not only the horizontal, but also the vertical foliation of X falls under this example.

6.1.5. Example 3: simple closed curves. If we take a $(4g + 2)$ -gon and glue opposite sides together with orientation reversing homeomorphisms, we obtain a closed surface of genus g (this can for instance be shown using the Euler characteristic of the result). We can turn this $(4g + 2)$ -gon into a rectangle, by drawing two opposite sides vertically and all the other sides horizontally (see Figure 4 for an example of genus 2).

Again, since the identification is by translations, the foliation and the line element $|dy|$ descend to the quotient surface X . In this foliation, all the non singular leaves are simple closed curves, that are all homotopic to each other. Conversely, any simple closed curve on a surface defines such a foliation. We can change the measure by scaling the rectangle.




FIGURE 3. Another foliation on a surface of genus two.




FIGURE 4. Yet another foliation on a surface of genus two. We identify (combinatorially and not geometrically) opposite sides.

6.1.6. Boundary and punctures. We haven't discussed yet what we require of a singular measured foliation when the surface S has boundary of punctures.

In our definition:

- We will allow the puncture to be a singular point or a regular point. If it's singular it will also be allowed to be a 1-pronged singularity (see Figure 5)
- At the boundary, the leaves of the foliation are allowed to be transverse or parallel. Moreover, we allow for singularities on the boundary (see Figure 6 for what the local picture is allowed to look like).




FIGURE 5. A one-pronged singularity at a puncture.

6.2. Holomorphic quadratic differentials

Next up in our discussion are holomorphic quadratic differentials. Indeed, if our surface comes equipped with the structure of a Riemann surface, we can define measured foliations on it using quadratic differentials.




FIGURE 6. The behavior we allow on the boundary.

If X is a Riemann surface, then a *holomorphic quadratic differential* on X is a section of the symmetric square of the holomorphic cotangent bundle, i.e. a quadratic form on the tangent bundle that varies holomorphically, or more concretely, an object that, in a local coordinate (U, z) , can be expressed as

$$\varphi(z) \cdot dz^2.$$

So if (V, w) is another holomorphic chart in which the quadratic differential takes the form $\psi(w) \cdot dw^2$, then in the overlap of V and U , we have

$$\psi(w) \cdot \left(\frac{dw}{dz} \right)^2 = \phi(z).$$

It follows from the Riemann–Roch theorem (see [Fal23]) that, on a closed Riemann surface of genus g , the space of holomorphic quadratic differentials has complex dimension 1 if $g = 1$ and $3g - 3$ if $g \geq 2$.

6.2.1. Example 0: the complex plane and squares of 1-forms. The complex plane \mathbb{C} needs a single chart, so dz^2 is a well defined quadratic differential. Likewise $\alpha \otimes \alpha$ is quadratic differential for any holomorphic 1-form α on a Riemann surface X .

6.2.2. Example 1: polygons. Suppose $P \subset \mathbb{R}^2$ is a Euclidean polygon and we are given a set of side pairings that are all (compositions of) translations and rotations or angle π . Let X be the surface we obtain when we glue P together using these side pairings (just like we’ve seen in the examples of measured foliations above).

First of all, the surface X has a natural Riemann surface structure. Indeed, if we let $S \subset X$ be the set of images of vertices of P . The subsurface $X - S \subset X$ admits an atlas consisting of charts whose transition maps are translations and rotations (with angle π , but that’s irrelevant for this part). Since such maps are holomorphic, this equips $X - S$ with the structure of a Riemann surface. So, we only need to explain how to find charts around the points in S that are compatible with this atlas.

First we observe that X comes with a natural metric that is Euclidean outside of S . Because the side pairings are translations and rotations of angle π (now the angle does matter), the total angle around any point in S is $m\pi$ for some integer $m \geq 1$. Suppose s is such a singularity, then we build a chart as follows. We take some small disk around s in the piecewise Euclidean metric (small enough to avoid other singularities). We can think of this disk as a gluing of m half disks (see Figure 7).

Our goal is to map this disk into a small disk around the origin in \mathbb{C} . To this end, we label the half disks in order: H_0, \dots, H_{m-1} that we all parametrize

$$H_k = \{ r \cdot e^{i\theta} : 0 \leq r < r_0, 0 \leq \theta \leq \pi \}, \quad k = 1, \dots, m$$




FIGURE 7. A small disk around a singularity.

We map these half disks into sectors of total angle $2\pi/m$ in \mathbb{C} , by

$$re^{i\theta} \in H_k \mapsto r^{2/m} e^{i \cdot (2\theta + 2k\pi)/m} \in \mathbb{C}.$$

So locally, each of these maps is a branch of the map $z \mapsto z^{2/m}$. These maps glue together into a chart around the singularity. We claim that the atlas we obtain from such charts, combined with the charts of $X - S$ whose transition functions are translations combined with rotations of angle π . Indeed, the chart transitions are of the form $z \mapsto z^{2/m}$, that are holomorphic functions away from 0. Since 0 corresponds to the singular point, which does not lie in the overlap of two charts, the chart transitions are holomorphic where they're supposed to be. X , equipped with this Riemann surface structure, is called a *half-translation surface*. If the side pairings are translations (so no rotations of angle π are involved) we call X a *translation surface*.

The quadratic differential dz^2 on \mathbb{C} is preserved by translations and rotations of angle π , so it descends to a quadratic differential q on $X - S$. In order to obtain a quadratic differential on X , we need to say how we define q in the charts around the singularities. Let's write (U, w) for the singular chart around a singularity of total angle $m\pi$ and (V, z) for some non-singular chart that has a non-trivial overlap with it. So $q(z) = dz^2$. Let's write

$$q(w) = \phi(w) \cdot dw^2$$

for the form q takes in the w -coordinate. By the transformation rule for quadratic differentials,

$$\phi(w) = \left(\frac{dz}{dw} \right)^2 \cdot \psi(z) = \left(\frac{dz}{dw} \right)^2.$$

Because the transition function is of the form $z(w) = w^{m/2}$, we obtain that we need to set

$$q(w) = \phi(w)dw^2 = \left(\frac{d}{dw} w^{m/2} \right)^2 = \frac{m^2}{4} w^{m-2} dw^2.$$

If $m = 1$, this leads to a pole. We can either allow such poles, or leave these singularities out.

In conclusion, from a polygon and a set of side identifications given by translations and rotations of angle π , we obtain a half translation surface equipped with a quadratic differential and a singular flat metric. Moreover, the singularities of total angle $m \cdot \pi$ correspond to zeroes (or poles if $m = 1$) of order $m - 2$ of q .

Finally, we observe that if we use only translations for the gluings (and X is a translation surface), then all the angles are integer multiples of 2π (as opposed to π) and instead of

only dz^2 , the 1-form dz descends to $X - S$ and can be completed to a globally defined 1-form on X . The quadratic differential we just constructed is its square in this case.

LECTURE 7

Quadratic differentials, Teichmüller mappings, hyperbolic geometry

7.1. Quadratic differentials

7.1.1. Measured foliations coming from quadratic differentials. We've already seen that, besides differentials, polygon gluings yield foliations too. This is a general phenomenon: quadratic differentials induce measured foliations and the charts we've described above are natural coordinates.

Indeed, we can define two smooth vector fields V^h and V^v , where at $p \in X$, the vectors $V_p^h, V_p^v \in T_p X$ are such that

$$q(V_p^h) \in (0, \infty) \quad \text{and} \quad q(V_p^v) \in (-\infty, 0)$$

if p is not a zero or pole of q , and the zero vector if p is. The foliations \mathcal{F}^h and \mathcal{F}^v consisting of the integral lines of V^h and V^v are called the *horizontal* and *vertical* foliations associated to q .

These foliations come with transverse measures. Indeed, suppose $q(z) = \phi(z)dz^2$ in some local coordinate (U, z) on X , then we can define

$$\mu^h(\alpha) = \int_{\alpha} \left| \operatorname{Im} \left(\sqrt{\phi(z)} dz \right) \right| \quad \text{and} \quad \mu^v(\alpha) = \int_{\alpha} \left| \operatorname{Re} \left(\sqrt{\phi(z)} dz \right) \right|.$$

EXAMPLE 7.1.1. If q is a quadratic differential coming from a half translation structure as described above, then \mathcal{F}^h and \mathcal{F}^v are the foliations induced by the horizontal and vertical foliation of \mathbb{C} respectively. Moreover, the horizontal and vertical transverse measure are induced by $|dy|$ and $|dx|$ respectively.

A natural question is whether all measured foliations on X appear as the horizontal foliations of quadratic differentials. A theorem due to Hubbard–Masur [HM79] says that this is essentially (up to a certain equivalence of foliations) the case.

7.1.2. Natural coordinates. Just like measured foliations, quadratic differentials admit natural coordinates.

DEFINITION 7.1.2. Let X be a Riemann surface. A natural coordinate (U, z) for a quadratic differential q on X is a holomorphic chart such that $q(z) = z^k dz^2$ for some $k \geq -1$.

Such coordinates exist. For example, if $p \in X$ is a point at which q does not vanish and let (U, z) be a chart such that $z(p) = 0$ and U is small enough such that it doesn't contain

any zeroes of q . Moreover, suppose $q(z) = \phi(z) \cdot dz^2$. We can then compose $z : U \rightarrow \mathbb{C}$ with the function

$$\eta(z) = \int_0^z \sqrt{\phi(w)} dw,$$

where we take the branch of the square root that gives the correct value at 0. η is a well-defined biholomorphic function onto its image. Moreover

$$d\eta^2 = \left(\frac{d\eta}{dz}\right)^2 dz^2 = \phi(z) dz^2.$$

So it indeed does what we want. We refer to [Str84, Section 6] for details.

7.1.3. Euclidean metrics. The final fact we will record about quadratic differentials is that they *always* come with a singular Euclidean metric.

Indeed, if, locally, $q(z) = \phi(z) dz^2$ then

$$\frac{1}{2i} |\phi(z)| d\bar{z} \wedge dz = |\phi(z)| dx \wedge dy$$

and

$$|\phi(z)|^{1/2} |dz| = |\phi(z)|^{1/2} \sqrt{dx^2 + dy^2}$$

give a well-defined length and area element respectively, that vanish only at the zeroes of q .

7.2. Teichmüller mappings

We are now ready to define Teichmüller mappings:

DEFINITION 7.2.1. Let X and Y be two closed Riemann surfaces. We say that a homeomorphism $f : X \rightarrow Y$ is a *Teichmüller mapping* if there are holomorphic quadratic differentials q_X and q_Y on X and Y respectively, and $K > 0$ such that:

- (1) f maps the zeroes of q_X to the zeroes of q_Y .
- (2) If $p \in X$ is not a zero of q_X , then with respect to the natural coordinates for q_X and q_Y based at p and $f(p)$, f satisfies

$$f(x + iy) = \sqrt{K} x + i \frac{1}{\sqrt{K}} y.$$

In complex coordinates, the last equation reads:

$$f(z) = \frac{1}{2} \left(\frac{K+1}{\sqrt{K}} \cdot z + \frac{K-1}{\sqrt{K}} \cdot \bar{z} \right).$$

In particular, f is holomorphic as a function of z if and only if $K = 1$.

LEMMA 7.2.2. *If X and Y are closed Riemann surfaces and $f : X \rightarrow Y$ is a Teichmüller mapping, then the total area of (the singular Euclidean metrics corresponding to) q_X and q_Y are the same.*

PROOF. In the natural coordinates for q_X and q_Y , the area forms associated to these quadratic differentials are the usual Euclidean area forms. In these same coordinates, the Jacobian of f has determinant 1, which means that f is area preserving with respect to these two area forms. \square

We are now ready to state Teichmüller's theorem:

THEOREM 7.2.3 (Teichmüller's theorem). *Let X and Y be closed Riemann surfaces and let $f : X \rightarrow Y$ be a homeomorphism. Then the following holds.*

- *The homotopy class of f contains a Teichmüller mapping $h : X \rightarrow Y$.*
- *If f is quasiconformal then*

$$K_f \geq K_h$$

with equality if and only if $f \circ h^{-1}$ is a biholomorphism. In particular, if $g \geq 2$ this means that $f = h$.

If time permits, we will prove this in a future lecture. We will however first continue studying the topology of Teichmüller space.

7.3. Hyperbolic surfaces

Hyperbolic geometry is a powerful tool in the study of Teichmüller spaces of surfaces of higher genus. Indeed, we will use it to prove that Teichmüller space is a ball. Before we get to that, we start by recalling some of the basics of the hyperbolic geometry of surfaces. We have already seen some of what follows in the first problem set, so we refer to this problem set for some of the proofs. Our main goal will be to show how to use pants decompositions to build hyperbolic surfaces

7.3.1. The hyperbolic plane. Hyperbolic geometry originally developed in the early 19th century to prove that the parallel postulate in Euclidean geometry is independent of the other postulates. From this perspective, the hyperbolic plane can be seen as a geometric object satisfying a collection of axioms very similar to Euclid's axioms for Euclidean geometry, but with the parallel postulate replaced by something else. From a more modern perspective, hyperbolic geometry is the study of manifolds that admit a Riemannian metric of constant curvature -1 .

7.3.2. The upper half plane model. From the classical point of view, any concrete description of the hyperbolic plane is a *model* for two-dimensional hyperbolic geometry, in the same way that \mathbb{R}^2 is a model for Euclidean geometry.

As we've already mentioned in Lecture 1, the hyperbolic plane can be defined as follows.

DEFINITION 7.3.1. The hyperbolic plane \mathbb{H}^2 is the complex domain

$$\mathbb{H}^2 = \{ z \in \mathbb{C} : \text{Im}(z) > 0 \}$$

equipped with the Riemannian metric given by

$$ds^2 = \frac{1}{y^2}(dx^2 + dy^2)$$

at $x + iy \in \mathbb{H}^2$

The group of orientation preserving isometries of (\mathbb{H}^2, ds^2) coincides with the group of complex automorphisms of \mathbb{H}^2 . That is,

$$\text{Isom}^+(\mathbb{H}^2) = \text{PSL}(2, \mathbb{R}).$$

Moreover, we've already seen during the exercises that the associated distance function is given by

$$d(z, w) = \cosh^{-1} \left(1 + \frac{|z - w|^2}{2 \cdot \text{Im}(z) \cdot \text{Im}(w)} \right).$$

Finally, we have also seen in the exercises what geodesics look like:

PROPOSITION 7.3.2. *The image of a geodesic $\gamma : \mathbb{R} \rightarrow \mathbb{H}^2$ is a vertical line or a half circle orthogonal to \mathbb{R} . Moreover, every vertical line and half circle orthogonal to the real line can be parameterized as a geodesic.*

We will often forget about the parametrization and call the image of a geodesic a geodesic as well. Note that it follows from the proposition above that given any two distinct points $z, w \in \mathbb{H}^2$ there exists a unique geodesic $\gamma \subset \mathbb{H}^2$ so that both $z \in \gamma$ and $w \in \gamma$. Furthermore, it also follows given a point $z \in \mathbb{H}^2$ and a geodesic γ that does not contain it, there is a unique perpendicular from z to γ (a geodesic γ' that intersects γ once perpendicularly and contains z)

The final fact we will need about the hyperbolic plane is:

PROPOSITION 7.3.3. *Let $z \in \mathbb{H}$ and let $\gamma \subset \mathbb{H}^2$ be a geodesic so that $z \notin \gamma$. Then*

$$d(z, \gamma) := \inf \{ d(z, w) : w \in \gamma \}$$

is realized by the intersection point of the perpendicular from z to γ .

Likewise, any two geodesics that don't intersect and are not asymptotic to the same point in $\mathbb{R} \cup \{\infty\}$ have a unique common perpendicular. Moreover, this perpendicular minimizes the distance between them.

PROOF. The first claim follows from Pythagoras' theorem for hyperbolic triangles. Indeed, given three points in \mathbb{H}^2 so that the three geodesics through them form a right angled hyperbolic triangle with sides of length a , b and c (where c is the side opposite the right angle), we have

$$\cosh(a) \cosh(b) = \cosh(c).$$

Indeed, this can be computed directly by putting the triangle in some standard position and then using the distance formula, a computation that we leave to the reader. This means in particular that $c > b$.

So, any other point on γ is further away from z than the point w realizing the perpendicular. Because that other point forms a right angled triangle with w and z .

The second claim follows from the first. □




FIGURE 1. The Farey tessellation.

7.3.3. The disk model. Another useful model, especially if one likes compact pictures, is the disk model of the hyperbolic plane. Set

$$\Delta = \{z \in \mathbb{C} : |z| < 1\}.$$

The map $f : \mathbb{H}^2 \rightarrow \Delta$ given by

$$f(z) = \frac{z - i}{z + i}$$

is a biholomorphism. We can also use it to push forward the hyperbolic metric to Δ . A direct computation tells us that the metric we obtain is given by

$$ds^2 = 4 \frac{dx^2 + dy^2}{(1 - x^2 - y^2)^2}.$$

Since f is conformal, the angles in the disk model are still the same as Euclidean angles.

Using the fact that the map f above is a Möbius transformation and thus sends circles and lines to circles and lines, one can prove:

PROPOSITION 7.3.4. *The hyperbolic geodesics in Δ are*

- *straight diagonals through the origin $0 \in \Delta$*
- *$C \cap \Delta$ where $C \subset \mathbb{C}$ is a circle that intersects $\partial\Delta$ orthogonally.*

For example, Figure 1 shows a collection of geodesics in Δ , known as the Farey tessellation.

7.3.4. Hyperbolic surfaces. A hyperbolic surface will be a finite type surface equipped with a metric that locally makes it look like \mathbb{H}^2 .

Because we will want to deal with surfaces with boundary later on, we need half spaces. Let $\gamma \subset \mathbb{H}^2$ be a geodesic. $\mathbb{H}^2 \setminus \gamma$ consists of two connected components C_1 and C_2 . We will call $\mathcal{H}_i = C_i \cup \gamma$ a *closed half space* ($i = 1, 2$). So for example

$$\{z \in \mathbb{H}^2 : \operatorname{Re}(z) \leq 0\}$$

is a closed half space.

We formalize the notion of a hyperbolic surface as follows:

DEFINITION 7.3.5. A finite type surface S with atlas $(U_\alpha, \varphi_\alpha)_{\alpha \in A}$ is called a *hyperbolic surface* if $\varphi_\alpha(U_\alpha) \subset \mathbb{H}^2$ for all $\alpha \in A$ and

1. for each $p \in S$ there exists an $\alpha \in A$ so that $p \in U_\alpha$ and
 - If $p \in \partial S$ then

$$\varphi_\alpha(U_\alpha) = V \cap \mathcal{H}$$
 for some open set $V \subset \mathbb{H}^2$ and some closed half space $\mathcal{H} \subset \mathbb{H}^2$.
 - If $p \in \mathring{S}$ then $\varphi_\alpha(U_\alpha) \subset \mathbb{H}^2$ is open.
2. For every $\alpha, \beta \in A$ and for each connected component C of $U_\alpha \cap U_\beta$ we can find a Möbius transformation $A : \mathbb{H}^2 \rightarrow \mathbb{H}^2$ so that

$$\varphi_\alpha \circ \varphi_\beta^{-1}(z) = A(z)$$

for all $z \in \varphi_\beta(C) \subset \mathbb{H}^2$.

Note that every hyperbolic surface comes with a metric: every chart is identified with an open set of \mathbb{H}^2 which gives us a metric. Because the chart transitions are restrictions of isometries of \mathbb{H}^2 , this metric does not depend on the choice of chart and hence is well defined.

DEFINITION 7.3.6. A hyperbolic surface S is called *complete* if the induced metric is complete.

We have seen in Lecture 2 that complete hyperbolic surfaces without boundary (considered up to isometry) correspond one-to-one to Riemann surfaces (considered up to biholomorphism).


7.3.5. Right angled hexagons. Even though Definition 7.3.5 is a complete definition, it is not very descriptive. In what follows we will describe a concrete cutting and pasting construction for hyperbolic surfaces.

We start with right angled hexagons. A right angled hexagon $H \subset \mathbb{H}^2$ is a compact simply connected closed subset whose boundary consists of 6 geodesic segments, that meet each other orthogonally. The picture to have in mind is displayed in Figure 2.

It turns out that the lengths of three non-consecutive sides determine a right angled hexagon up to isometry.

PROPOSITION 7.3.7. *Let $a, b, c \in (0, \infty)$. Then there exists a right angled hexagon $H \subset \mathbb{H}^2$ with three non-consecutive sides of length a, b and c respectively. Moreover, if H' is another right angled hexagon with this property, then there exists a Möbius transformation $A : \mathbb{H}^2 \rightarrow \mathbb{H}^2$ so that*

$$A(H) = H'.$$


FIGURE 2. A right angled hexagon H .

PROOF. Let us start with the existence. Let γ_{im} denote the positive imaginary axis and set

$$B = \{ z \in \mathbb{H}^2 : d(z, \gamma_{im}) = c \}.$$

B is a one-dimensional submanifold of \mathbb{H}^2 . Because the map $z \mapsto \lambda z$ is an isometry that preserves γ_{im} for every $\lambda > 0$, it must also preserve B . This means that B is a (straight Euclidean) line.

Now construct the following picture:

FIGURE 3. Constructing a right angled hexagon $H(a, b, c)$.

That is, we take the geodesic through the point $i \in \mathbb{H}^2$ perpendicular to γ_{im} and at distance a draw a perpendicular geodesic γ . furthermore, for any $p \in B$, we draw the geodesic α that realizes a right angle with the perpendicular from p to γ_{im} . Now let

$$x = d(\alpha, \gamma) = \inf \{ d(z, w) : z \in \gamma, w \in \alpha \}.$$

Because of Proposition 7.3.3, x is realized by the common perpendicular to α and γ . By moving p over B , we can realize any positive value for x and hence obtain our hexagon $H(a, b, c)$.

We also obtain uniqueness from the picture above. Indeed, given any right angled hexagon H' with three non-consecutive sides of length a , b and c , apply a Möbius transformation $A : \mathbb{H}^2 \rightarrow \mathbb{H}^2$ so that the geodesic segment of length a starts at i and is orthogonal to the imaginary axis. This implies that the geodesic after a gets mapped to the geodesic γ . Furthermore, one of the endpoints of the geodesic segment of length c needs to lie on the line B . We now know that the geodesic α before that point needs to be tangent to B . Because α and β have a unique common perpendicular. The tangency point of α to B determines the picture entirely. Because the function that assigns the length x of the common perpendicular to the tangency point is injective, we obtain that there is a unique solution. \square

7.4. The universal cover of a hyperbolic surface with boundary

It will be useful to have a description of the Riemannian universal cover of a surface with boundary. To this end, we first prove:

PROPOSITION 7.4.1. *Let X be a hyperbolic surface with non-empty boundary that consists of closed geodesics. Then there exists a complete hyperbolic surface X^* without boundary in which X can be isometrically embedded so that X is a deformation retract of X^* .*

PROOF. For each $\ell \in (0, \infty)$, we define a hyperbolic surface

$$F_\ell = [0, \infty) \times \mathbb{R} / \{t \sim t + 1\},$$

equipped with the metric

$$ds^2 = d\rho^2 + \ell^2 \cosh^2(\rho) \cdot dt^2$$

for all $(\rho, t) \in F_\ell$. We will call such a surface a *funnel*. One can check that this is a metric of constant curvature -1 , in which the boundary is totally geodesic. Alternatively, we can identify

$$F_\ell = \left\{ z \in \mathbb{H}^2 : \operatorname{Re}(z) \geq 0 \right\} \Bigg/ \left\langle \begin{bmatrix} e^{\ell/2} & 0 \\ 0 & e^{\ell/2} \end{bmatrix} \right\rangle.$$

We can glue funnels of the right length along the boundary components, in a similar way to Example 8.1.1. Figure 4 shows an example.




FIGURE 4. Attaching funnels

Since both F_ℓ and X are complete, the resulting surface X^* is complete.

Moreover, since F_ℓ retracts onto its boundary component, X is a deformation retract of X^* . \square

See [Bus10, Theorem 1.4.1] for a version of the above to surfaces with more general types of boundary components.

Recall that a subset $C \subset M$ of a Riemannian manifold M is called *convex* if for all $p, q \in C$ there exists a length minimizing geodesic $\gamma : [0, d(p, q)] \rightarrow M$ such that

$$\gamma(0) = p, \quad \gamma(d(p, q)) = q \quad \text{and} \quad \gamma(t) \in C \quad \forall t \in [0, d(p, q)].$$

As a result of this construction we obtain:

PROPOSITION 7.4.2. *Let X be a complete hyperbolic surface with non-empty boundary that consists of closed geodesics. Then the universal Riemannian cover of \tilde{X} of X is isometric to a convex subset of \mathbb{H}^2 whose boundary consists of complete geodesics.*

PROOF. The Killing-Hopf theorem tells us that the universal cover of X^* is the hyperbolic plane \mathbb{H}^2 . Here X^* is the surface given by Proposition 7.4.1.

Let us denote the covering map by $\pi : \mathbb{H}^2 \rightarrow X^*$. Now let C be a connected component of $\pi^{-1}(X)$. The boundary of C consists of the lifts of ∂X and hence of a countable collection of disjoint complete geodesics in \mathbb{H}^2 . As such, it's a countable intersection of half spaces (which are convex) and hence convex. \square

7.5. Pairs of pants and gluing

One of our main building blocks for hyperbolic surfaces is the following:

DEFINITION 7.5.1. Let $a, b, c \in (0, \infty)$. A *pair of pants* is a hyperbolic surface that is diffeomorphic to $\Sigma_{0,3,0}$ such that the boundary components have length a , b and c respectively.

PROPOSITION 7.5.2. *Let $a, b, c \in (0, \infty)$ and let P and P' be pairs of pants with boundary curves of lengths a, b and c . Then there exists an isometry $\varphi : P \rightarrow P'$.*

PROOF SKETCH. There exists a unique orthogonal geodesic (this essentially follows from Proposition 7.3.3 below, the proof of Proposition 8.3.1 that we will do in full during the exercises, is similar) between every pair of boundary components of P .

These three orthogonals decompose P into right-angled hexagons out of which three non-consecutive sides are determined. Proposition 7.3.7 now tells us that this determines the hexagons up to isometry and this implies that P is also determined up to isometry. \square

Note that it also follows from the proof sketch above that the unique perpendiculars cut each boundary curve on P into two geodesic segments of equal length.

LECTURE 8

Pants decompositions and annuli

8.1. Non-compact pairs of pants

In order to deal with non-compact surfaces, we will need non-compact polygons. To this end, we note that, looking at Proposition 7.3.2, complete geodesics in \mathbb{H}^2 are parametrized by their *endpoints*: pairs of distinct point in

$$\partial\mathbb{H}^2 := \mathbb{R} \cup \{\infty\}$$

(or \mathbb{S}^1 if we use the disk model).

A (not necessarily compact) polygon now is a closed connected simply connected subset $P \subset \mathbb{H}^2$, whose boundary consists of geodesic segments.

If two consecutive segments “meet” at a point in $\partial\mathbb{H}^2$, this point will be called an *ideal* vertex of the boundary. Note that the angle at an ideal vertex is always 0. A polygon all of whose vertices are ideal is called an *ideal polygon*.

We can also make sense of a pair of pants where some of the boundary components have “length” 0. In this case, we obtain a complete hyperbolic structure on a surface with boundary and punctures so that

$$\#\text{punctures} + \#\text{boundary components} = 3.$$

Such pairs of pants can be obtained by gluing either

- two pentagons with one ideal vertex each and right angles at the other vertices,
- two quadrilaterals with two ideal vertices each right angles at the other vertices
or
- two ideal triangles.

Along the sides of infinite length there however is a gluing condition. We will come back to this later (see Proposition 8.3.3). Moreover, we obtain a similar uniqueness statement to the proposition above. As always in the non-compact case, the adjective “complete” does need to be added.

EXAMPLE 8.1.1. If P is a pair of pants and $\delta \subset \partial P$ is one of its boundary components, let us write $\ell(\delta)$ for the length of δ . Given two pairs of pants P_1 with boundary components δ_1, δ_2 and δ_3 and P_2 with boundary components γ_1, γ_2 and γ_3 so that

$$\ell(\delta_1) = \ell(\gamma_1),$$

we can choose an orientation reversing isometry $\varphi : \delta_1 \rightarrow \gamma_1$ and from that obtain a hyperbolic surface

$$S = P_1 \sqcup P_2 / \sim,$$

where $\varphi(x) \sim x$ for all $x \in \delta_1$. One way to see that this surface comes with a well defined hyperbolic structure, is that locally it's obtained by gluing two half spaces in \mathbb{H}^2 together along their defining geodesics. Note that S is diffeomorphic to $\Sigma_{0,4,0}$.

Repeating the construction above, we can build hyperbolic surfaces of any genus and any number of boundary components. In what follows we will prove that every hyperbolic surface can be obtained from this construction.

8.2. The geometry of isometries

Recall that we can classify isometries in $\mathrm{PSL}(2, \mathbb{R})$ into three different types:

DEFINITION 8.2.1. Let $g \in \mathrm{PSL}(2, \mathbb{R})$.

- (1) If $\mathrm{tr}(g)^2 < 4$ then g is called *elliptic*.
- (2) If $\mathrm{tr}(g)^2 = 4$ then g is called *parabolic*.
- (3) If $\mathrm{tr}(g)^2 > 4$ then g is called *hyperbolic*.

Note that, since trace is conjugacy invariant, conjugate elements in $\mathrm{PSL}(2, \mathbb{R})$ are of the same type. It turns out (as we will see below) that closed geodesics correspond exactly to conjugacy classes of hyperbolic elements.

We've seen during the exercises that the classification above can equivalently be described as:

LEMMA 8.2.2. Let $g \in \mathrm{PSL}(2, \mathbb{R})$. Then

- (1) g is elliptic if and only if g has a single fixed point inside \mathbb{H}^2 .
- (2) g is parabolic if and only if g has a single fixed point on $\mathbb{R} \cup \{\infty\}$.
- (3) g is hyperbolic if and only if g has two distinct fixed points on $\mathbb{R} \cup \{\infty\}$.

Given a hyperbolic isometry, we can define its translation distance as follows:

DEFINITION 8.2.3. Let $g \in \mathrm{PSL}(2, \mathbb{R})$ be hyperbolic. Then its *translation distance* is given by

$$T_g := \inf \{ z \in \mathbb{H}^2 : d(z, gz) \}.$$

Moreover, its *axis* is defined as

$$\alpha_g := \{ z \in \mathbb{H}^2 : d(z, gz) = T_g \}.$$

We have:

LEMMA 8.2.4. *Let $g \in \mathrm{PSL}(2, \mathbb{R})$ be hyperbolic with fixed points $x_1, x_2 \in \partial\mathbb{H}^2$. Then its axis α_g is the unique geodesic between x_1 and x_2 and its translation length is given by*

$$T_g = 2 \cosh^{-1} \left(\frac{|\mathrm{tr}(g)|}{2} \right).$$

PROOF. Since the claim is conjugacy invariant, we can conjugate g so that $x_1 = 0$ and $x_2 = \infty$. Which means that we can assume without loss of generality that

$$g = \begin{bmatrix} \lambda & 0 \\ 0 & \lambda^{-1} \end{bmatrix}$$

for some $\lambda \in (0, \infty)$. Using the fact that $2 \cosh(\frac{1}{2} \cosh^{-1}(x)) = \sqrt{2x+2}$, We get that

$$2 \cosh(d(z, gz)/2) = \sqrt{4 + \frac{(\lambda^2 - 1)^2 \cdot (\mathrm{Im}(z)^2 + \mathrm{Re}(z)^2)}{\lambda^2 \mathrm{Im}(z)^2}} \geq \sqrt{4 + \frac{(\lambda^2 - 1)^2}{\lambda^2}} = \lambda + \frac{1}{\lambda},$$

with equality if and only $\mathrm{Re}(z) = 0$, thus proving the lemma. \square

8.3. Geodesics and conjugacy classes

Recall that there is a one-to-one correspondence

$$\left\{ \begin{array}{c} \text{Conjugacy classes of} \\ \text{non-trivial elements in } \pi_1(X) \end{array} \right\} \leftrightarrow \left\{ \begin{array}{c} \text{free homotopy classes of} \\ \text{non-trivial closed curves on } X \end{array} \right\}.$$

We will call a curve puncture parallel if it can be homotoped into a puncture.

It turns out that on a hyperbolic surface (or more generally a negatively curved Riemannian manifold), every free homotopy class of essential closed curves contains a unique geodesic:

PROPOSITION 8.3.1. *Let X be a complete hyperbolic surface with totally geodesic boundary.*

- (1) *Then in every homotopy class of non-puncture parallel closed curves γ on X , there exists a unique geodesic that minimizes the length among all classes in the homotopy class.*
- (2) *Moreover, if the free homotopy class contains a simple closed curve, then the corresponding geodesic is also simple.*
- (3) *More generally, if γ and γ' are non-homotopic non-puncture parallel and non-trivial closed curves, then*
 - *The number of self-intersections of the unique geodesic $\bar{\gamma}$ homotopic to γ is minimal among all closed curves homotopic to γ and*
 - *$\#\bar{\gamma} \cap \bar{\gamma}'$ is minimal among all pairs of curves homotopic to γ and γ' respectively.*

PROOF. The proof will be part of this week's exercises. \square

We will use the proposition above to prove:

PROPOSITION 8.3.2. *Let X be a complete hyperbolic surface with totally geodesic boundary. Then there are one-to-one correspondences between the following three sets:*

$$\left\{ \begin{array}{l} \text{Non-trivial free homotopy classes of} \\ \text{non puncture-parallel closed curves on } X \end{array} \right\},$$

$$\left\{ \begin{array}{l} \text{Conjugacy classes of} \\ \text{hyperbolic elements in } \Gamma \end{array} \right\}$$

and

$$\left\{ \begin{array}{l} \text{Oriented, unparametrized} \\ \text{closed geodesics on } \mathbb{H}^2/\Gamma \end{array} \right\}.$$

PROOF. The correspondence between the last and the first set is given by the previous proposition, so we only need to show that conjugacy classes of hyperbolic elements correspond one-to-one to oriented, unparametrized geodesics.

In order to make our lives a little easier, we will assume X to be closed. The argument for the general case is similar. We will hence not worry about the assumption that the curve is non puncture parallel, nor about boundary components.

First of all consider a conjugacy class $K \subset \Gamma$ of hyperbolic elements. Let us pick an element $g \in K$, with axis $\alpha_g \subset \mathbb{H}^2$. The projection map $\pi : \mathbb{H}^2 \rightarrow X$ sends α_g to a closed geodesic of length T_g . Moreover, since

$$\pi(\alpha_{hgh^{-1}}) = \pi(h\alpha_g) = \pi(\alpha_g),$$

the resulting geodesic does not depend on the choice of g .

In the opposite direction, a closed geodesic on \mathbb{H}^2/Γ lifts to a countable union of geodesics in \mathbb{H}^2 (the orbit of a single such geodesic under Γ), each invariant under a cyclic group of deck transformations. These transformations need to fix the endpoints of the given geodesic, so they are hyperbolic. The action of Γ on the geodesics corresponds to conjugation of these hyperbolic elements. \square

Before we get to pants decompositions, we record what happens to curves that are parallel to a puncture.

PROPOSITION 8.3.3. *Let X be a complete hyperbolic surface. So that $X = C/\Gamma$ where C is a convex subset of \mathbb{H}^2 , bounded by complete geodesics and $\Gamma < \text{PSL}(2, \mathbb{R})$ acts properly discontinuously and freely on C . Then there are a one-to-one correspondences*

$$\left\{ \begin{array}{l} \text{Conjugacy classes of} \\ \text{parabolic elements in } \Gamma \end{array} \right\} \leftrightarrow \left\{ \begin{array}{l} \text{Oriented, unparametrized} \\ \text{puncture-parallel closed curves } \mathbb{H}^2/\Gamma \end{array} \right\}.$$

This proposition gives us the gluing condition we spoke about in Section 7.5: the gluing needs to be so that the resulting puncture parallel curves give rise to parabolic elements, this turns out to uniquely determine the gluing.

8.4. Every hyperbolic surface admits a pants decomposition

As an immediate consequence to Proposition 8.3.1 we get that hyperbolic surfaces admit pants decompositions.

DEFINITION 8.4.1. Let X be a complete, orientable hyperbolic surface of finite area. A pants decomposition of X is a collection of pairwise disjoint simple closed geodesics $\mathcal{P} = \{\alpha_1, \dots, \alpha_k\}$ in X so that each connected component of

$$X \setminus \left(\bigcup_{i=1}^k \alpha_i \right)$$

consists of hyperbolic pairs of pants whose boundary components have been removed.

We have the following:

LEMMA 8.4.2. *Let \mathcal{P} be a pants decomposition of a hyperbolic surface X that is homeomorphic to $\Sigma_{g,b,n}$ then*

- \mathcal{P} contains $3g + n + b - 3$ closed geodesics and
- $X \setminus \mathcal{P}$ consists of $2g + n + b - 2$ pairs of pants.

PROOF. This can be proved using the Euler characteristic. □

PROPOSITION 8.4.3. *Let X be a complete, orientable hyperbolic surface of finite area and totally geodesic boundary. Then X admits a pants decomposition.*

PROOF. Take any collection of simple closed curves on $\Sigma_{g,b,n}$ that decompose it into pairs of pants. Proposition 8.3.1 tells us that these curves can be realized by unique geodesics. □

Note that we actually get countably many such pants decompositions: given a pants decomposition we can apply a diffeomorphism not isotopic to the identity (of which we already know there are many) to obtain a new topological pants decomposition, that is realized by different geodesics.

Finally, we remark, that lengths alone are not enough to determine the hyperbolic metric:

EXAMPLE 8.4.4. φ in Example 8.1.1 is determined up to ‘twist’. That is, if we parameterize δ_1 by a simple closed geodesic $x : \mathbb{R}/(\ell(\delta_1)\mathbb{Z}) \rightarrow \delta_1$ and $\varphi' : \delta_1 \rightarrow \gamma_1$ is a different orientation reversing isometry, then there exists some $t_0 \in \mathbb{R}$ so that

$$\varphi'(x(t)) = \varphi(x(t_0 + t))$$

for all $t \in \mathbb{R}/(\ell(\delta_1)\mathbb{Z}) \rightarrow \delta_1$.

Summarizing the discussion above, we get the following parametrization of all hyperbolic surfaces:

THEOREM 8.4.5. *Let (g, b, n) be so that*

$$\chi(\Sigma_{g,b,n}) < 0.$$

If we fix a pants decomposition \mathcal{P} of $\Sigma_{g,b,n}$ and vary the lengths $\ell_i \in (0, \infty)$ and twist $\tau_i \in [0, \ell_i]$, we obtain all complete hyperbolic surfaces homeomorphic to $\Sigma_{g,b,n}$.

Note however that there is no guarantee that we don't obtain the same surface multiple times (and in fact we do).

8.5. Annuli

8.5.1. Hyperbolic annuli. Our goal is to use pants decompositions to define global coordinates on Teichmüller space. In order to prove continuity of the coordinates we obtain, we need to understand (to some degree) how the complex structure and the hyperbolic metric depend on each other. It turns out that understanding this for annuli will suffice. So, before we get to Fenchel–Nielsen coordinates, we will discuss annuli.

If $g \in \mathrm{PSL}(2, \mathbb{R})$ is a hyperbolic or parabolic isometry then the group $\langle g \rangle \simeq \mathbb{Z}$ acts on \mathbb{H}^2 properly discontinuously and freely. This means that

$$N_g = \mathbb{H}^2 / \langle \gamma \rangle$$

is an orientable hyperbolic surface with fundamental group \mathbb{Z} and hence an annulus. First we note that the geometry of the annulus only depends on the translation length of g . We record this as a lemma, the proof of which we leave to the reader.

LEMMA 8.5.1. *Let $g, h \in \mathrm{PSL}(2, \mathbb{R})$ be either both hyperbolic or both parabolic elements so that their translation lengths satisfy $T_g = T_h$. Then the annuli N_g and N_h are isometric. Moreover, every complete hyperbolic annulus is isometric to N_g for some parabolic or hyperbolic $g \in \mathrm{PSL}(2, \mathbb{R})$.*

Note that this includes the case where $T_g = T_h = 0$.

8.5.2. Complex annuli. The complex parametrization of annuli we will need is:

$$A_m := \{ z \in \mathbb{C} : 0 < \mathrm{Im}(z) < m \} / \mathbb{Z}$$

for all $m > 0$. Here the \mathbb{Z} -action is given by $k \cdot z = z + k$ for all $k \in \mathbb{Z}$, $z \in \mathbb{C}$.

We also record a version of Grötzsch's theorem for these annuli (the proof of which is a variation of the proof we saw in Section 5.5).

THEOREM 8.5.2 (Grötzsch's theorem). *Let $f : A_m \rightarrow A_{m'}$ be a K -quasiconformal map. Then*

$$\frac{1}{K} \leq \frac{m}{m'} \leq K.$$

Moreover, equality is realized if and only if f can be lifted to a map

$$\tilde{f} : \{ z \in \mathbb{C} : 0 < \mathrm{Im}(z) < m \} \rightarrow \{ z \in \mathbb{C} : 0 < \mathrm{Im}(z) < m' \}$$

given by

$$\tilde{f}(x + iy) = b + x + i \frac{m'}{m} y$$

for some $b \in \mathbb{R}$.

We observe that this theorem also implies that A_m and $A_{m'}$ are biholomorphic if and only if $m = m'$. The number m is called the *modulus* of the annulus.

LECTURE 9

Fenchel–Nielsen coordinates

9.1. More on annuli

The question now becomes whether N_g is biholomorphic to A_m for some m and if so, to which. In order to solve this question, we introduce a new (somewhat uncommon) model for the hyperbolic plane the *band model*. Set

$$\mathbb{B} = \left\{ z \in \mathbb{C} : |\operatorname{Im}(z)| < \frac{\pi}{2} \right\},$$

equipped with the metric

$$ds^2 = \frac{dx^2 + dy^2}{\cos^2(y)}.$$

This is another model for the hyperbolic plane, moreover the real line is a geodesic in \mathbb{B} . Maps of the form $\varphi_b : \mathbb{B} \rightarrow \mathbb{B}$ defined by

$$z \mapsto z + b$$

for some $b > 0$ are isometries for this metric. Moreover $\langle \varphi_b \rangle \simeq \mathbb{Z}$ acts on \mathbb{B} properly discontinuously, which means that

$$M_b = \mathbb{B} / \langle \varphi_b \rangle$$

is a hyperbolic annulus. Moreover, the translation length of φ_b is b , so using Lemma 8.5.1, we see that

$$M_b \simeq N_g$$

as hyperbolic surfaces, where $g \in \operatorname{PSL}(2, \mathbb{R})$ is any hyperbolic element with translation length b .

We now claim that:

LEMMA 9.1.1. *Let $m > 0$. The annuli A_m and $M_{\pi/m}$ are biholomorphic.*

PROOF. Since the map $z \mapsto z - i m/2$ is a biholomorphism of \mathbb{C} that commutes with the \mathbb{Z} -action. A_m is biholomorphic to

$$\left\{ z \in \mathbb{C} : |\operatorname{Im}(z)| < \frac{m}{2} \right\} / \mathbb{Z}.$$

The map $\left\{ z \in \mathbb{C} : |\operatorname{Im}(z)| < \frac{m}{2} \right\} \rightarrow \mathbb{B}$ given by $z \mapsto \frac{\pi}{m} z$ is a \mathbb{Z} -equivariant biholomorphism and hence descends to a biholomorphism

$$\left\{ z \in \mathbb{C} : |\operatorname{Im}(z)| < \frac{m}{2} \right\} / \mathbb{Z} \simeq M_{\pi/m}.$$

□

For the parabolic case we have:

LEMMA 9.1.2. *let $g \in \mathrm{PSL}(2, \mathbb{R})$ be parabolic. The annuli N_g and $\mathbb{D} \setminus \{0\}$ are biholomorphic.*

PROOF. Using Lemma 8.5.1, we may assume that

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

The map $\mathbb{H}^2 \rightarrow \mathbb{D}$ given by

$$z \mapsto e^{-2\pi iz}$$

induces the biholomorphism. □

9.2. Fenchel–Nielsen coordinates

Now we’re ready to introduce Fenchel–Nielsen coordinates on Teichmüller spaces of hyperbolic surfaces. In particular, in this section, we will assume that our base surface S admits a complete hyperbolic metric. Moreover, we will fix a (topological) pants decomposition \mathcal{P} on S .

9.2.1. Lengths. Given any essential closed curve γ on S , we obtain a function

$$\ell_\gamma : \mathcal{T}(S) \rightarrow \mathbb{R}_+$$

called a *length function*, defined as follows. Each $[R, f] \in \mathcal{T}(S)$ can be seen as a marked hyperbolic surface. So, Proposition 8.3.1 implies that the homotopy class of $f(\gamma)$ on R contains a unique geodesic. $\ell_\gamma([R, f])$ is the length of this geodesic.

Hence, given S and \mathcal{P} as above, we obtain a map

$$\ell_{\mathcal{P}} : \mathcal{T}(S) \rightarrow \mathbb{R}_+^{3g-3+n}$$

defined by

$$\ell_{\mathcal{P}}([R, f]) = \left(\ell_\gamma([R, f]) \right)_{\gamma \in \mathcal{P}}.$$

We have:

LEMMA 9.2.1 (Wolpert). *Let S and γ be as above. The function*

$$\log \circ \ell_\gamma : \mathcal{T}(S) \rightarrow \mathbb{R}$$

is 2-Lipschitz with respect to the Teichmüller metric, i.e.

$$|\log(\ell_\gamma([R, f])) - \log(\ell_\gamma([R', f']))| \leq 2 d_T([R, f], [R', f'])$$

for all $[R, f], [R', f'] \in \mathcal{T}(S)$.

PROOF. Fix a basepoint $p \in S$ so that we can identify γ with an element of $\pi_1(S, p)$, that we will also denote by γ . The infinite cyclic subgroup of $\pi_1(S, p)$ generated by γ induces a \mathbb{Z} -cover

$$S_\gamma \rightarrow S.$$

We will write A and A' for the corresponding covering spaces of R and R' . Just like in the proof of Proposition 8.3.1, these are annuli and by Lemma 9.1.1, they are biholomorphic

to $A_{\pi/\ell_\gamma}([R, f])$ and $A_{\pi/\ell_\gamma}([R', f'])$ respectively. K -quasiconformal maps between R and R' lift to K -quasiconformal maps of A and A' . So we have

$$\begin{aligned} 2 \, d_T([R, f], [R', f']) &= \inf_{\substack{g \text{ homotopic} \\ \text{to } f \circ f^{-1}}} \log(K_g) \\ &\geq \inf_{\substack{g \text{ homotopic} \\ \text{to } f' \circ f^{-1}}} \log(K_{\tilde{g}: A \rightarrow A'}) \\ &\geq \left| \log \left(\frac{\ell_\gamma([R, f])}{\ell_\gamma([R', f'])} \right) \right|, \end{aligned}$$

where the last line follows from Grötzsch's theorem (Theorem 8.5.2). \square

9.2.2. Twists. So, given S and \mathcal{P} as above, we have a continuous map

$$\ell_{\mathcal{P}} : \mathcal{T}(S) \rightarrow \mathbb{R}_+^{3g-3+n}.$$

It's however not quite injective. The problem is that we can still rotate the hyperbolic metric along the curves in the pants decomposition. Twist coordinates will remedy this.

First we pick a collection of disjoint simple closed curves Γ so that for each pair of pants P in $S \setminus \mathcal{P}$, $\Gamma \cap P$ consists of three arcs, each connecting a different pair of boundary components of P . Figure 1 shows an example.




FIGURE 1. A pants decomposition \mathcal{P} with a set of curves Γ .

Regardless of our choice of pants decomposition \mathcal{P} , such a system of curves Γ always exists.

Now let $\gamma \in \mathcal{P}$ be a pants curve. Then γ bounds either one P or two pairs of pants P_1 and P_2 in the decomposition. Let us assume the latter for simplicity, the other case is analogous. The left hand side of Figure 2 shows an example of such a curve γ .




FIGURE 2. The image of an arc under a diffeomorphism.

If $f : S \rightarrow R$ is an orientation preserving diffeomorphism, then it maps \mathcal{P} to some pants decomposition of R . Moreover, if η is one of the (two) components of $(P_1 \cup P_2) \cap \Gamma$ that intersects γ , then $f(\eta)$ is some arc between boundary components of $f(P_1)$ and $f(P_2)$ (like on the right hand side of Figure 2). Now

- δ will be the unique simple closed geodesic in the free homotopy class of $f(\gamma)$ on R .
- α_1 and α_2 the two unique perpendiculars between the boundary components between which $f(\eta)$ runs and δ (see Figure 2).

Then relative to the boundary of $f(P_1 \cup P_2)$, the arc $f(\eta)$ is freely homotopic to $\alpha_2 \cdot \delta^k \cdot \alpha_1$ for some $k \in \mathbb{Z}$.

The twist along γ is now

$$\tau_\gamma([R, f]) = k \cdot \ell_\gamma([R, f]) \pm d(p_1, p_2) \in \mathbb{R}$$

where

- p_1 and p_2 are the points where α_1 and α_2 hit δ .
- The signs are determined by the orientation of R in the following way. The orientation of R gives a notion of “twisting to the left” along δ . Left twists are counted positively and right twists negatively.

Let us prove that twists are continuous:

LEMMA 9.2.2. *Let S and γ be as above. The function*

$$\tau_\gamma : \mathcal{T}(S) \rightarrow \mathbb{R}$$

is continuous.

PROOF SKETCH. Suppose that

$$d_T([R, f], [R', f'])$$

is small. This means that the map $f' \circ f^{-1} : R \rightarrow R'$ is close to an isometry. Since it maps the curves and arcs used to define $\tau_\gamma([R, f])$ to those used to define $\tau_\gamma([R', f'])$. So, this

map lifts to a map $\widetilde{f' \circ f^{-1}} : \mathbb{H}^2 \rightarrow \mathbb{H}^2$ that is close to conformal and hence close to an isometry. This means that the numbers $\tau_\gamma([R, f])$ and $\tau_\gamma([R', f'])$ are close. \square

Putting the above together, we obtain a continuous map

$$\text{FN}_{\mathcal{P}} : \mathcal{T}(S) \rightarrow \mathbb{R}_+^{3g-3+n} \times \mathbb{R}^{3g-3+n}$$

defined by

$$\text{FN}_{\mathcal{P}}([R, f]) = \left(\ell_\gamma([R, f]), \tau_\gamma([R, f]) \right)_{\gamma \in \mathcal{P}}.$$

It turns out that the Fenchel Nielsen map is a homeomorphism:

THEOREM 9.2.3. *Let S be a surface of finite type such that $\chi(S) < 0$ and let \mathcal{P} be a pants decomposition of S . Then the map*

$$\text{FN}_{\mathcal{P}} : \mathcal{T}(S) \rightarrow \mathbb{R}_+^{3g-3+n} \times \mathbb{R}^{3g-3+n},$$

is a homeomorphism.

PROOF. Since we have already proved that lengths and twists are continuous, we only need to provide a continuous inverse to the map $\text{FN}_{\mathcal{P}}$.

Given a vector $(\ell_\gamma, \tau_\gamma)_{\gamma \in \mathcal{P}}$, we can use the gluing construction we discussed above in order to produce a hyperbolic surface R . The lengths give us the geometry of the pairs of pants and the gluing along a curve γ is determined by

$$\tau_\gamma^{(0)} = \tau_\gamma + k \cdot \ell_\gamma,$$

where k is such that $\tau_\gamma^{(0)} \in [0, \ell_\gamma)$. Call this surface $R((\ell_\gamma, [\tau_\gamma])_\gamma)$. In particular, by varying the twist τ_γ , we obtain the same surface countably many times.

The question however is what the marking, i.e. the map $f : S \rightarrow R((\ell_\gamma, [\tau_\gamma])_\gamma)$, should be. In order to do this, we fix open regular neighborhoods N_γ^S of the curves $\gamma \in \mathcal{P}$ on S so that

$$S \setminus \bigcup_{\gamma \in \mathcal{P}} N_\gamma$$

consists of disjoint pairs of pants P_1^S, \dots, P_k^S . We will once and for all parametrize the annuli

$$N_\gamma^S = (\mathbb{R}/\mathbb{Z}) \times (-1, 1).$$

On $R((\ell_\gamma, [\tau_\gamma])_\gamma)$ we pick such neighborhoods too and obtain neighborhoods N_γ^R and pairs of pants P_i^R . We will assume that

$$N_\gamma^R = \left\{ x \in R((\ell_\gamma, [\tau_\gamma])_\gamma) : d(x, \gamma) < \varepsilon \right\}$$

for some ε small enough. Moreover, we assume ε varies continuously as a function of $(\ell_\gamma, [\tau_\gamma])_\gamma$.

In order to build f , we now pick a parametrization

$$N_\gamma^R = \left(\mathbb{R}/\ell_\gamma \mathbb{Z} \right) \times (-1, 1)$$

where the subset

$$\left(\mathbb{R}/\ell_\gamma \mathbb{Z} \right) \times \{t\} \subset N_\gamma^R$$

is one of the (one or two) components of

$$\left\{ x \in R\left((\ell_\gamma, [\tau_\gamma])_\gamma\right) : d(x, \gamma) = |t| \cdot \varepsilon \right\},$$

parametrized by a constant multiple (depending on t) of arclength for all $t \in (-1, 1)$.

The map $f_\gamma : N_\gamma^S \rightarrow N_\gamma^R$ is now given by

$$f_\gamma(\theta, t) = \left(\ell_\gamma \cdot \theta + \tau_\gamma \cdot \frac{t+1}{2}, t \right).$$

The awkward $(t+1)/2$ is an artifact of choosing the interval $(-1, 1)$ instead of $(0, 1)$ (the latter would have made some of the previous equations more awkward).

For the complements of the annuli we choose arbitrary homeomorphisms and $f_i^P : P_i^S \rightarrow P_i^R$ that smoothly extend the f_γ .

This map is clearly an inverse and since we can make everything depend on the input continuously, it's continuous. \square

REMARK 9.2.4. Looking at the proof above, it's a natural question to ask whether we maybe get a fundamental domain for moduli space by only considering $\tau_\gamma \in [0, \ell_\gamma)$.

However, this not the case. To see this, take any $f \in \text{Diff}^+(S, \Sigma)$ (where $S = S_0 \setminus \Sigma$, S_0 is closed and Σ a finite set) that is not homotopic to the identity. Then we get a surface isometric to $R\left((\ell_\gamma, [\tau_\gamma])_{\gamma \in \mathcal{P}}\right)$ if we assign the lengths of the curves in $f(\mathcal{P})$ to the curves in \mathcal{P} instead (the isometry will be induced by f).

LECTURE 10

Teichmüller's theorem and the Nielsen–Thurston classification

10.1. Teichmüller's theorem

10.1.1. Building Teichmüller maps. The first question is how one builds Teichmüller maps. Given a Riemann surface X , a quadratic differential q_X and $K > 0$, let us describe how to build a Riemann surface Y , equipped with a quadratic differential q_Y and a Teichmüller map $X \rightarrow Y$ corresponding to this data.

Let X' denote the Riemann surface we obtain if we remove the zeroes of q_X from X . We may equip X' with an atlas consisting of natural coordinates for q_X . Now compose all of these coordinates with affine maps

$$f(x + iy) = \sqrt{K} \cdot x + \frac{1}{\sqrt{K}} \cdot iy$$

This yields a new Riemann surface Y' (homeomorphic to X' through a map we will also call f). Moreover, by Riemann's removable singularities theorem [SS03, Theorem 3.1], we may extend the Riemann surface structure to a closed Riemann surface Y homeomorphic to X . q_X induces a quadratic differential q_Y on Y and we obtain an Teichmüller map that we will also denote $f : X \rightarrow Y$.

Fixing X and q_X , but varying $K \in (0, \infty)$ yields a one parameter family of Riemann surfaces. If $[X, \phi] \in \mathcal{T}(S)$, the resulting family of points in $\mathcal{T}(S)$ is called a Teichmüller line.

10.1.2. An exponential map. Recall from Section 7.1.3 that, if in a local coordinate z , the quadratic differential q on X is given by $q(z) = \phi(z)dz^2$, then

$$|q(z)| = \frac{1}{2i} |\phi(z)| d\bar{z} \wedge dz = |\phi(z)| dx \wedge dy$$

is a well-defined area form on X . We set

$$||q|| = \int_X |q|.$$

Set

$$\text{QD}_1(X) = \{ q \text{ a quadratic differential on } X : ||q|| < 1 \}.$$

By the Riemann–Roch theorem, this has the topology of a ball in \mathbb{C}^{3g-3} .

Now given $q \in \text{QD}_1(X)$, we set

$$K_q = \frac{1 + ||q||}{1 - ||q||}.$$

So, using the procedure from the previous section, a quadratic differential $q \in \text{QD}_1(X)$ now yields a point $[Y, f] \in \mathcal{T}(X)$, where f is a K_q -quasiconformal Teichmüller map. We will denote this map by

$$\mathcal{E} : \text{QD}_1(X) \rightarrow \mathcal{T}(X).$$

This map will play an important role in our proof of Teichmüller's theorem.

10.1.3. A proof sketch. We now have all the set-up we need in order to prove Teichmüller's theorem, that we repeat here:

THEOREM 7.2.3 (Teichmüller's theorem). *Let X and Y be closed Riemann surfaces and let $f : X \rightarrow Y$ be a homeomorphism. Then the following holds.*

- *The homotopy class of f contains a Teichmüller mapping $h : X \rightarrow Y$.*
- *If f is quasiconformal then*

$$K_f \geq K_h$$

with equality if and only if $f \circ h^{-1}$ is a biholomorphism. In particular, if $g \geq 2$ this means that $f = h$.

PROOF SKETCH. The uniqueness part can be proved using a similar string of inequalities to the proof of Grötsch's theorem (Theorem 5.5.1) and we will skip it (see for instance [FM12, Section 11.6] for this).

We observe that existence is equivalent to the map $\mathcal{E} : \text{QD}_1(X) \rightarrow \mathcal{T}(X)$ being surjective. Indeed, if we want to find a Teichmüller map in the homotopy class of $f : X \rightarrow Y$ we need to show that the map \mathcal{E} hits the point $[Y, f] \in \mathcal{T}(X)$.

The idea is to use Brouwer's invariance of domain:

THEOREM 10.1.1. *Let $n \geq 1$, then any proper injective continuous map $\mathbb{R}^n \rightarrow \mathbb{R}^n$ is a homeomorphism.*

Indeed, \mathcal{E} is a map from one homeomorphic copy of \mathbb{R}^{6g-6} to another. So we need to show that it's proper, injective and continuous, which will then yield that it's surjective.

Uniqueness of the Teichmüller map in each homotopy class implies injectivity, so we need to show properness and continuity, the last of which is the hard part.

Indeed, there is no clear geometric reason why nearby quadratic differentials should yield geometrically similar Teichmüller maps: the behavior of the associated foliations can vary drastically: arbitrarily close to a quadratic differential with a horizontal foliation all of whose leaves are closed we can find quadratic differentials with horizontal foliations all of whose leaves are infinite (think of the rational versus irrational slope examples on the torus).

Continuity: In order to prove continuity of \mathcal{E} , we will decompose it into two maps \mathcal{E}_1 and \mathcal{E}_2 . The idea is to first transform our quadratic differential into a Beltrami differential. It makes sense to do this, because we've seen that Teichmüller space can be embedded in a space of Beltrami differentials. The second step is to go back from Beltrami differentials to points in Teichmüller space.

Let us detail both maps. So, the first map turns our quadratic differential $q \in \text{QD}_1(X)$ into an almost everywhere defined essentially bounded Beltrami differential on \mathbb{C} . Since we need only one chart to cover \mathbb{C} , we can think of the space of such differentials as $L^\infty(\mathbb{C})$. Let us assume that X has genus at least two and write it as $X = \Gamma \backslash \mathbb{H}^2$ (the case of tori is similar). So we can lift q to a Γ -invariant quadratic differential \tilde{q} on \mathbb{H}^2 . We now define $\mathcal{E}(q) \in L^\infty(\mathbb{C})^\Gamma$ by

$$\mathcal{E}_1(q)(z) = \begin{cases} \|q\| \cdot \tilde{q}(z) / |\tilde{q}(z)| & \text{if } \text{Im}(z) > 0 \\ \|q\| \cdot \overline{\tilde{q}(\bar{z})} / |\tilde{q}(\bar{z})| & \text{if } \text{Im}(z) < 0 \end{cases}.$$

In order to obtain a point in Teichmüller space from this Beltrami differential we now need to solve the Beltrami equation. Indeed, the idea is to find the quasiconformal map that induces $\mathcal{E}_1(q)$ as a Beltrami differential. Even if we identified Teichmüller space with a space of Beltrami differentials before, this step is necessary. Indeed, in our identification (Corollary 5.2.6), we identified Teichmüller space with a set of Beltrami differentials coming from quasiconformal maps. In order to go back, we need to know that we can go back: given a Beltrami differential, it comes from a quasiconformal map. To this end, we state the Riemann mapping theorem that we already mentioned once in the second lecture:

THEOREM 10.1.2. *Let $\mu \in L^\infty(\mathbb{C})$ with $\|\mu\|_\infty < 1$. There exists a unique quasiconformal homeomorphism $f^\mu : \hat{\mathbb{C}} \rightarrow \hat{\mathbb{C}}$ fixing 0, 1 and ∞ that almost everywhere satisfies the Beltrami equation*

$$\mu = \frac{\partial f^\mu / \partial \bar{z}}{\partial f^\mu / \partial z}.$$

Moreover, f^μ is smooth wherever μ is, and f^μ varies complex analytically with respect to μ .

Since $\|\mathcal{E}_1(q)\|_\infty = \|q\| < 1$, this theorem gives us a quasiconformal map $f^{\mathcal{E}_1(q)} : \hat{\mathbb{C}} \rightarrow \hat{\mathbb{C}}$. By uniqueness, $f^{\mathcal{E}_1(q)}$ has the same symmetries as $\mathcal{E}_1(q)$ and hence restricts to a Γ -invariant Beltrami differential on \mathbb{H}^2 and hence yields a Beltrami differential $\mu_{f^{\mathcal{E}_1(q)}}$ on X and hence a point in Teichmüller space $\mathcal{T}(X)$ (see Corollary 5.2.6), which we will call $\mathcal{E}_2(\mathcal{E}_1(q))$.

$\mathcal{E}_2 \circ \mathcal{E}_1$ is continuous because \mathcal{E}_1 is given by a continuous expression and continuity of \mathcal{E}_2 follows from the Riemann mapping theorem. What we still need to explain is that $\mathcal{E} = \mathcal{E}_2 \circ \mathcal{E}_1$.

To this end, suppose at a point q is given by $q = re^{i\theta} dz^2$ then $\mathcal{E}_1(q) = \|q\| e^{i\theta}$. So the Beltrami equation looks for a function f such that

$$\frac{\partial f / \partial \bar{z}}{\partial f / \partial z} = \|q\| e^{i\theta}$$

This is a stretch of $\frac{1+\|q\|}{1-\|q\|}$ in the direction $e^{i\theta}$ (dictated by q), which is exactly what the map $\mathcal{E}(q)$ does.

Properness: Because \mathcal{E} is continuous, we know that the pre-image of a closed set is closed. As such, we just need to check that the pre-image of a compact set in $\mathcal{T}(X)$ is bounded in $\text{QD}_1(X)$. This is a direct consequence of the fact that the topology of $\mathcal{T}(X)$ is defined in

terms of Beltrami differentials. Indeed, on a compact set in $\mathcal{T}(X)$, the minimal quasiconformal dilatation of maps homotopic to marking changes is uniformly bounded (because the ∞ -norm of their Beltrami differentials is uniformly bounded away from 1). So the pre-image yields a set of quadratic differentials whose norm is uniformly bounded away from 1. \square

10.2. The Nielsen–Thurston classification on the torus

The next goal of this class will be to describe Thurston's proof of the Nielsen–Thurston classification, based on a compactification of Teichmüller space using projective measured foliations. We note that there is also a proof using the Teichmüller metric, due to Bers, this can for instance be found in [FM12, Section 13.6] and [Hub16, Chapter 8].

First, we need to discuss the statement of the theorem. For some intuition, recall that if T denotes the torus, then $\mathrm{MCG}(T) \simeq \mathrm{SL}(2, \mathbb{Z})$ and elements in this group are either elliptic, parabolic or hyperbolic.

Thought of as self maps of the torus, both elliptic elements and $-\begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$ are homotopic to a map that can be realized as a homotopically non-trivial isometry/biholomorphism of the torus corresponding to their fixed point in $\mathcal{T}(T) \simeq \mathbb{H}^2$. We will call such mapping classes *periodic*.

Parabolic elements correspond to Dehn twists. Indeed, the matrices $\pm \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}$ represent the only conjugacy classes of parabolics in $\mathrm{SL}(2, \mathbb{Z})$ and they are both represented by Dehn twists. Moreover, if α is a simple closed curve on T and $[\varphi] \in \mathrm{MCG}(T)$, then the conjugate of the Dehn Twist T_α by $[\varphi]$ satisfies

$$[\varphi] \circ T_\alpha \circ [\varphi^{-1}] = T_{\varphi(\alpha)}.$$

As such all parabolics in $\mathrm{SL}(2, \mathbb{Z})$ correspond to Dehn twists. In particular, as mapping classes they are *reducible*: they fix a curve.

Finally, we have seen in Section 6.1 that if $A \in \mathrm{SL}(2, \mathbb{Z})$ is hyperbolic, like for instance $\begin{pmatrix} 2 & 1 \\ 1 & 1 \end{pmatrix}$, then it fixes two transverse smooth measured foliations, stretches along one and contracts by the same amount along the other. Such maps are called *Anosov*. They are also Teichmüller maps from a torus represented by a point on the axis of A in \mathbb{H}^2 (now thought of as a Riemann surface) to itself with equal initial and terminal quadratic differential.

10.3. Definitions and statement of the classification theorem

We now first need to generalize the definitions to higher genus surfaces. Singular foliations will be called transverse if their singularities coincide and their non-singular leaves are transverse to each other. Moreover, we generalize the terminology from the torus in the following way:

DEFINITION 10.3.1. Let $[\varphi] \in \mathrm{MCG}(\Sigma_g)$. Then

- $[\varphi]$ is called *periodic* if it has finite order in $\text{MCG}(\Sigma_g)$,
- $[\varphi]$ is called *reducible* if there exists some finite collection of disjoint pairwise non-homotopic homotopy classes of simple closed curves $\gamma_1, \dots, \gamma_k$ on Σ_g such that

$$[\varphi](\{\gamma_1, \dots, \gamma_k\}) = \{\gamma_1, \dots, \gamma_k\},$$

and

- $[\varphi]$ is called *pseudo-Anosov* if there exists a representative $\phi \in [\varphi]$, a real number $\lambda > 1$ and a pair of transverse measured foliations (\mathcal{F}^u, μ^u) and (\mathcal{F}^s, μ^s) called the *unstable* and the *stable foliation* respectively, such that

$$\phi \cdot (\mathcal{F}^u, \mu^u) = (\mathcal{F}^u, \lambda \mu^u) \quad \text{and} \quad \phi \cdot (\mathcal{F}^s, \mu^s) = (\mathcal{F}^s, \lambda^{-1} \mu^s).$$

We note that the pseudo-Anosov representative is not unique: we can isotope the foliations around and conjugate the map by that isotopy. However, it turns out that any two representatives are related by an isotopy [FLP79, Exposé 12, Thm III].

Again, the latter can be phrased in terms of Teichmüller mappings. $[\varphi]$ is pseudo-Anosov if and only if it contains a representative $\phi : \Sigma_g \rightarrow \Sigma_g$ for which there exists a point $[X, f] \in \mathcal{T}(\Sigma_g)$ such that

$$f \circ \phi \circ f^{-1} : X \rightarrow X$$

is a Teichmüller mapping.

The content of the Nielsen–Thurston classification is that the options above are the only options:

THEOREM 10.3.2 (Nielsen–Thurston classification of mapping classes). *Let $g \geq 1$ and $[\varphi] \in \text{MCG}(\Sigma_g)$. Then $[\varphi]$ is either periodic, reducible or pseudo-Anosov.*

,

10.4. An overview of the proof

Thurston’s full proof of Theorem 10.3.2 is worked out in [FLP79, FLP12] (see also [Baa21]). We will content ourselves with sketching the main ideas.

Let us start with an overview of the proof. The idea will be to try to mimic the classification on elements of $\text{MCG}(\Sigma_1) \simeq \text{SL}(2, \mathbb{Z})$ acting on $\mathcal{T}(\Sigma_1) \simeq \mathbb{H}^2$ into elliptic, parabolic and hyperbolic elements. This classification can be phrased in terms of the fixed points of elements of $\text{SL}(2, \mathbb{Z})$ on $\overline{\mathbb{H}^2} = \mathbb{H}^2 \cup \partial \mathbb{H}^2$, which is homeomorphic to a closed disk. In particular, Brouwer’s fixed point theorem guarantees fixed points for all elements of $\text{SL}(2, \mathbb{Z})$.

We have already seen that in the higher genus case, $\mathcal{T}(\Sigma_g)$ is homeomorphic to a closed ball. So, in order to run a similar argument, we need to find a way to attach a boundary to $\mathcal{T}(\Sigma_g)$ to which the mapping class group action continuously extends.

This boundary will be the space of *projective measured foliations* on Σ_g . We will first define the space of measured foliations $\mathcal{MF}(\Sigma_g)$, which as a set will be measured foliations up to isotopy and *Whitehead moves* (see Figure 1) that either collapse two singularities into a

single one or split a singularity into two singularities. Note that because of the invariance of the associated measure under leafwise isotopy, the measure is preserved by such a move. The projectivization $\mathbb{P}\mathcal{MF}(\Sigma_g) = \mathcal{MF}(\Sigma_g)/\mathbb{R}_{>0}$ of the space of measured foliations,




FIGURE 1. A Whitehead move, turning two 3-pronged singularities into a single 4-pronged singularity

where $\mathbb{R}_{>0}$ acts by multiplying the measure, turns out to be homeomorphic to a $6g - 7$ dimensional sphere.

After that we will consider the set

$$\mathcal{C} = \left\{ \text{essential simple closed curves on } \Sigma_g \right\} / \text{isotopy}$$

and the space $\mathbb{P}(\mathbb{R}_{>0}^{\mathcal{C}})$. It turns out that $\mathcal{T}(\Sigma_g)$ and $\mathbb{P}\mathcal{MF}(\Sigma_g)$ both inject into $\mathbb{P}(\mathbb{R}_{>0}^{\mathcal{C}})$. Moreover, in the image, $\mathbb{P}\mathcal{MF}(\Sigma_g)$ appears as the boundary of $\mathcal{T}(\Sigma_g)$ and the union is a closed $(6g - 6)$ -dimensional ball, on which $\text{MCG}(\Sigma_g)$ acts continuously. So, we get what we wanted: by Brouwer's fixed point theorem, every element will have a fixed point. Like in the case of $\text{SL}(2, \mathbb{Z})$, we obtain the classification by analyzing these fixed points.

10.5. The space of measured foliations

10.5.1. Measured foliations on a pair of pants. Let P be a pair of pants and let us define the set

$$\mathcal{MF}_0(P) = \left\{ (\mathcal{F}, \mu) \text{ a measured foliation on } P : \begin{array}{l} \text{no boundary component of} \\ P \text{ is a non-singular leaf} \end{array} \right\}.$$

We claim that $\mathcal{MF}_0(P)$ can be identified with $\mathbb{R}_{\geq 0}^3 - \{0\}$. Indeed, the claim is that for any $(a, b, c) \in \mathbb{R}_{\geq 0}^3$, there exists a unique measured foliation (\mathcal{F}, μ) , up to isotopy and Whitehead moves, such that the μ -lengths of the boundary components equals a , b and c respectively.

Indeed, up to Whitehead moves and isotopy and permuting the roles of the boundary components, the only two options for the foliation are shown in Figure 2.

Using the invariance of μ under leafwise isotopy, we see that $a + b \geq c$ for the foliation on the left and $a + b \leq c$ for the foliation on the right.

To see that all elements of $\mathcal{MF}_0(P)$ are indeed of this form, we use the Euler–Poincaré formula (Proposition 6.1.2), which implies that any foliation of P has at most two three-pronged singularity or one four-pronged singularity, which can be split into two three-pronged singularities using a Whitehead move. Afterwards it's a matter of where the singular leaves go.




FIGURE 2. The two options for the foliation underlying an element of $\mathcal{MF}_0(P)$.




FIGURE 3. A pants curve γ_i and two arcs δ_i and ε_i .

We obtain all elements of $\mathcal{MF}(P)$ by attaching annuli foliated by parallel circles to the elements of $\mathcal{MF}_0(P)$.

10.5.2. Coordinates on $\mathcal{MF}(\Sigma_g)$. In order to obtain global coordinates on $\mathcal{MF}(\Sigma_g)$, we will use a construction not dissimilar to that of Fenchel–Nielsen coordinates. Indeed, we will fix a pants decomposition $(\gamma_i)_{i=1}^{3g-3}$ of Σ_g .

Now fix an annulus A_i parallel to each pants curve γ_i such that γ_i is one of the boundary components of A_i . Moreover, fix two arcs δ_i and ε_i that triangulate A_i (see Figure 3).

Now given a measured foliation (\mathcal{F}, μ) , we can put it in *normal form* with respect to this choice of annuli (using isotopy and Whitehead moves). This means that:

- the foliation restricts to an element of $\mathcal{MF}_0(P)$ on the complements of the A_i , and
- in the annuli A_i , the foliation is either transverse to γ_i or it consists of leaves parallel to γ_i . In particular, it contains no singularities in A_i .

We obtain an injective map $\mathcal{MF}(\Sigma_g) \rightarrow \mathbb{R}_{\geq 0}^{9g-9}$ by recording the measures of the γ_i , δ_i and ε_i . Because any arc transverse to γ_i hits either ε_i or δ_i however, we have that

$$\mu(\gamma_i) = \mu(\delta_i) + \mu(\varepsilon_i),$$

so the image lands in a copy of \mathbb{R}^{6g-6} . This endows $\mathcal{MF}(\Sigma_g)$ with a topology and indeed gives $\mathbb{PMF}(\Sigma_g)$ the topology of \mathbb{S}^{6g-7} .

10.6. Compactifying Teichmüller space

10.6.1. Two maps. Like we said above, we will map both $\mathcal{T}(\Sigma_g)$ and $\mathbb{P}\mathcal{MF}(\Sigma_g)$ into $\mathbb{P}(\mathbb{R}_{\geq 0}^{\mathcal{C}})$ and will identify the latter as the image of the former.

The map from Teichmüller space is given by lengths of curves. That is, we will use the map $\ell : \mathcal{T}(\Sigma_g) \rightarrow \mathbb{P}(\mathbb{R}_{\geq 0}^{\mathcal{C}})$, given by

$$\ell([X, f]) = \left[(\ell_{\alpha}([X, f]))_{\alpha \in \mathcal{C}} \right].$$

To show that this is a homeomorphism onto its image, one can use Fenchel–Nielsen coordinates. Indeed, if we fix any pants decomposition, we immediately recover the lengths corresponding to it from the image, up to a constant factor. Using enough curves that intersect the curves in the pants decomposition, we can also recover the constant factor and the twists (it turns out that a finite number of curves suffices).

In order to define the map from $\mathbb{P}\mathcal{MF}(\Sigma_g)$ into $\mathbb{P}(\mathbb{R}_{\geq 0}^{\mathcal{C}})$, we need to define the length of a simple closed curve with respect to a measured foliation. We will define this by

$$I((\mathcal{F}, \mu), \alpha) = \inf \{ \mu(\tilde{\alpha}) : \tilde{\alpha} \text{ isotopic to } \alpha \}, \quad (\mathcal{F}, \mu) \in \mathcal{MF}(\Sigma_g), \alpha \in \mathcal{C}.$$

This defines a map

$$I : \mathbb{P}\mathcal{MF}(\Sigma_g) \rightarrow \mathbb{P}(\mathbb{R}_{\geq 0}^{\mathcal{C}}).$$

This map is a homeomorphism onto its image. This again comes down to proving that we can recover the coordinates we described above (up to a scaling factor) from the image of a point.

LECTURE 11

Landmark results, part I

11.1. Compactifying Teichmüller space, part II

We start this lecture by continuing our discussion of the Nielsen–Thurston classification theorem (Theorem 10.3.2). Just like last time, this is only meant to be a sketch of the ideas. Details can be found in [FLP79, FLP12, Baa21].

Now we need to show that the union of the images of the two maps $\ell : \mathcal{T}(\Sigma_g) \rightarrow \mathbb{P}(\mathbb{R}_{\geq 0}^c)$ and $I : \mathcal{MF}(\Sigma_g) \rightarrow \mathbb{P}(\mathbb{R}_{\geq 0}^c)$ described above is a closed $(6g - 6)$ -dimensional ball whose interior corresponds to $\mathcal{T}(\Sigma_g)$ and whose boundary is the image of $\mathbb{P}\mathcal{MF}(\Sigma_g)$.

We will first prove that the images are disjoint. To see this, we note that the *systole* – the length of the shortest closed geodesic – of a closed hyperbolic surface is a strictly positive number. As such, if a point in $\mathbb{P}(\mathbb{R}_{\geq 0}^c)$ comes from a point in $\mathcal{T}(\Sigma_g)$ then there is a uniform lower bound on the coordinates of any representative in $\mathbb{R}_{\geq 0}^c$ of that point.

On the other hand, we claim that this is not the case for points in the image of $\mathbb{P}\mathcal{MF}(\Sigma_g)$. Indeed, suppose $[(\mathcal{F}, \mu)] \in \mathbb{P}\mathcal{MF}(\Sigma_g)$ there are two cases to consider:

- (1) If \mathcal{F} has a closed leaf λ . This is a simple closed curve, the intersection of which with \mathcal{F} equals 0.
- (2) If \mathcal{F} doesn't have any closed leaves, then take any non-singular leaf of \mathcal{F} . Because Σ_g is compact, this leaf needs to come back arbitrarily close to itself. As such, we can create simple closed curves with arbitrarily small intersection with (\mathcal{F}, μ) by tracking such a non singular leaf and closing up the trajectory once we're very close to a point on the leaf we have already visited.

11.1.1. Projective measured laminations as the boundary of Teichmüller space.

In order to see that $I(\mathbb{P}\mathcal{MF}(\Sigma_g))$ is the boundary of $\ell(\mathcal{T}(\Sigma_g))$, we first construct a map from

$$p : \mathcal{T}(\Sigma_g) \longrightarrow \mathcal{MF}(\Sigma_g).$$

To do this, we pick a pants decomposition $(\alpha_i)_{i=1}^{3g-3}$ of Σ_g and measure Fenchel-Nielsen coordinates with respect to this pants decomposition.

The triple of lengths corresponding to a pair of pants give us an element of $\mathcal{MF}_0(P)$. We use the twists to decide how to glue these foliations together into a foliation of Σ_g .

The key observation is the following. Given $\varepsilon > 0$, define

$$U(\varepsilon) = \{ [X, f] \in \mathcal{T}(\Sigma_g) : \ell_{\alpha_i}([X, f]) > \varepsilon \text{ for } i = 1, \dots, 3g - 3 \}.$$

Now for every $\varepsilon > 0$ and every $\alpha \in \mathcal{C}$, there exists a $C(\alpha, \varepsilon) > 0$ such that

$$I(p([X, f]), \alpha) \leq \ell_\alpha([X, f]) \leq I(p([X, f]), \alpha) + C(\alpha, \varepsilon)$$

In particular, if a sequence $\left([X_n, f_n] \in U(\varepsilon)\right)_n$ is such that its image $\ell([X_n, f_n]) \in \mathbb{R}_{\geq 0}^{\mathcal{C}}$ tends to infinity, then the projectivized sequence

$$\left[\ell([X_n, f_n])\right] \in \mathbb{P}(\mathbb{R}_{\geq 0}^{\mathcal{C}})$$

converges if and only if the sequence

$$\left[I(p([X_n, f_n]))\right] \in \mathbb{P}(\mathbb{R}_{\geq 0}^{\mathcal{C}})$$

converges. Moreover, if they do, then the limits are the same. So this implies that we can identify limits of sequences of points in $U(\varepsilon)$ with points in $\mathcal{PMF}(\Sigma_g)$. By changing the pants decomposition, we can contain any sequence in $U(\varepsilon)$ for some $\varepsilon > 0$. This proves the identification.

We will write

$$\overline{\mathcal{T}}(\Sigma_g) = \ell(\mathcal{T}(\Sigma_g)) \cup I(\mathbb{PMF}(\Sigma_g)).$$

11.1.2. The topology of $\overline{\mathcal{T}}(\Sigma_g)$. It turns out that $\overline{\mathcal{T}}(\Sigma_g)$ is a manifold. Moreover, a theorem by Brown implies that $\mathbb{PMF}(\Sigma_g)$ admits a collar neighborhood in $\overline{\mathcal{T}}(\Sigma_g)$. The boundary of this collar is homeomorphic to \mathbb{S}^{6g-7} . Since it lies in $\overline{\mathcal{T}}(\Sigma_g)$, which is homeomorphic to \mathbb{R}^{6g-6} , the generalized Schönflies theorem, proved by Brown, Mazur and Morse, the interior of this copy of \mathbb{S}^{6g-7} is a ball (there is a regularity condition to check here, that we skip over). This implies that $\overline{\mathcal{T}}(\Sigma_g)$ is homeomorphic to a ball.

11.1.3. The case of the torus. In order to have a concrete example, we will treat the case of the torus. In the higher genus case, we used the fact that the lengths of finitely many curves suffice to determine a point in Teichmüller space.

In the case of the torus, three lengths suffice. That is, we identify $\Sigma_1 = \mathbb{R}^2/\mathbb{Z}^2$ and will let α_1, α_2 and α_3 be the curves corresponding to $(1, 0), (0, 1), (1, 1) \in \mathbb{Z}^2$ respectively. We obtain a map

$$(\ell_{\alpha_1}, \ell_{\alpha_2}, \ell_{\alpha_3}) : \mathcal{T}(\Sigma_1) \longrightarrow \mathbb{R}_{>0}^3.$$

Moreover, because, when we base them at the same point, α_1, α_2 and α_3 form a (non-degenerate) triangle, we have

$$\ell_{\alpha_i}([X, f]) < \ell_{\alpha_j}([X, f]) + \ell_{\alpha_k}([X, f]) \quad \text{for all distinct } i, j, k \in \{1, 2, 3\}$$

and for all $[X, f] \in \mathcal{T}(\Sigma_1)$. Moreover, we have also seen that we may rescale say the length of α_1 to be 1. This is equivalent to saying that the projectivization map $\mathbb{R}_{>0}^3 \rightarrow \mathbb{P}(\mathbb{R}_{>0}^3)$ is injective on the image of $\mathcal{T}(\Sigma_1)$. We can identify the image $\mathbb{P}(\mathbb{R}_{>0}^3)$ of the projectivization map with $\Delta = \{(x_1, x_2, x_3) \in \mathbb{R}_{>0}^3 : x_1 + x_2 + x_3 = 1\}$. Teichmüller space sits in here as the subset

$$\{(x_1, x_2, x_3) \in \Delta : x_i < x_j + x_k \text{ for all } i, j, k\}.$$

The map $\mathcal{MF}(\Sigma_1) \rightarrow \mathbb{R}_{\geq 0}^3$ defined by

$$[\mathcal{F}, \mu] \mapsto (\mu(\alpha_1), \mu(\alpha_2), \mu(\alpha_3))$$

is injective. Moreover, for the same reason as in Section 10.5.2,

$$\mu(\alpha_i) = \mu(\alpha_j) + \mu(\alpha_k) \text{ for one choice of distinct } i, j, k \in \{1, 2, 3\}.$$

So $\mathbb{P}\mathcal{MF}(\Sigma_1)$ can be identified with

$$\{(x_1, x_2, x_3) \in \Delta : x_i = x_j + x_k \text{ for some choice of } i, j, k\},$$

i.e. the boundary of the image of Teichmüller space. Moreover, their union is indeed a closed disk.

11.2. Classifying mapping classes

11.2.1. The mapping class group action. The mapping class group $\text{MCG}(\Sigma_g)$ acts both on Teichmüller space and on projective measured foliations. It turns out that these two actions combine into a continuous action on $\overline{\mathcal{T}}(\Sigma_g)$. Indeed, in the image, this is action is the action by permuting the elements of \mathcal{C} .

Since $\overline{\mathcal{T}}(\Sigma_g)$ Brouwer's fixed point theorem implies that every element of $\text{MCG}(\Sigma_g)$ has at least one fixed point in $\overline{\mathcal{T}}(\Sigma_g)$. Now we need to analyze the fixed points of elements of $\text{MCG}(\Sigma_g)$

11.2.2. Analyzing the fixed points. Let $[\varphi] \in \text{MCG}(\Sigma_g)$. We will distinguish the following cases:

- (1) $[\varphi] \cdot [X, f] = [X, f]$ for some $[X, f] \in \mathcal{T}(\Sigma_g)$.
- (2) The fixed point is some element $[\mathcal{F}, \mu] \in \mathbb{P}\mathcal{MF}(\Sigma_g)$. Since $\text{MCG}(\Sigma_g)$ acts on $\mathcal{MF}(\Sigma_g)$ as well, we may lift to a representative $[\mathcal{F}, \mu] \in \mathcal{MF}(\Sigma_g)$. This point is no longer necessarily a fixed point, because $[\varphi]$ might multiply the measure. As such, we make a further case distinction:
 - (a) $[\varphi] \cdot [\mathcal{F}, \mu] = [\mathcal{F}, \lambda \cdot \mu]$ for some $\lambda \in [1, \infty)$ and \mathcal{F} has at least one closed leaf.
 - (b) $[\varphi] \cdot [\mathcal{F}, \mu] = [\mathcal{F}, \mu]$ and \mathcal{F} has no closed leaves.
 - (c) $[\varphi] \cdot [\mathcal{F}, \mu] = [\mathcal{F}, \lambda \cdot \mu]$ for some $\lambda > 1$ and \mathcal{F} has no closed leaves.

In the above, we have assumed that $\lambda \geq 1$, which we may, by potentially replacing $[\varphi]$ by its inverse.

11.2.3. Case 1. In this case (X, f) and $(X, f \circ \varphi^{-1})$ are Teichmüller equivalent. This means that there exists an orientation preserving isometry (or equivalently a complex automorphism) $m : X \rightarrow X$ such that

$$(f \circ \varphi^{-1})^{-1} \circ m \circ f = \varphi \circ (f^{-1} \circ m \circ f)$$

is isotopic to the identity on Σ_g . Since $m : X \rightarrow X$ is a periodic map (because automorphism groups of higher genus surfaces are finite, by Hurwitz's theorem – see Theorem 12.1.1 below), this means that $[\varphi]$ is homotopic to a periodic map.




FIGURE 1. A full system of transversals near a three-pronged singularity.

11.2.4. Case 2a. The foliation \mathcal{F} will have at most finitely many isotopy classes of closed leaves. If φ preserves it, then it needs to permute these finitely many isotopy classes. In particular, this means that $[\varphi]$ is reducible.

11.2.5. Case 2b. If both the foliation \mathcal{F} and the measure μ are preserved by $[\varphi]$, then we can build a finite system of rectangles on Σ_g that are permuted by $[\varphi]$, which implies that $[\varphi]$ is periodic. This goes as follows.

First of all, a *horizontal rectangle* $R \subset \Sigma_g$ with respect to a singular foliation \mathcal{F} is an immersion $\phi : [0, 1]^2 \hookrightarrow \Sigma_g$ such that:

- The restriction of ϕ to $(0, 1)^2$ is an embedding.
- For all $t \in [0, 1]$, the set $\phi([0, 1] \times \{t\})$ is contained in a finite union of leaves of \mathcal{F} .
- For all $s \in [0, 1]$, the set $\phi(\{s\} \times [0, 1])$ is transverse to \mathcal{F} .

We will write $\partial^T R$ for the *transverse boundary* of R : the set $\phi(\{0, 1\} \times [0, 1])$.

We now choose transversal arcs τ_1, \dots, τ_n to \mathcal{F} such that:

- all these arcs have exactly one endpoint in the singularity,
- they have no other endpoints in common,
- there is exactly one arc between every pair of prongs of each singularity and
- the arcs based at a given singularity all have the same μ -length.

See Figure 1 for a local picture near a singularity. We will call this collection of arcs a *full system of transversals for \mathcal{F}* and we set $\tau = \tau_1 \cup \dots \cup \tau_n$.

We now claim:

LEMMA 11.2.1. *Given a singular foliation \mathcal{F} on Σ_g without any closed leaves and a full system of transversals τ , we can find a finite collection of horizontal rectangles R_1, \dots, R_N with respect to \mathcal{F} such that*

$$\bigcup_{i=1}^N R_i = \Sigma_g \quad \text{and} \quad \bigcup_{i=1}^N \partial^T R_i = \tau$$

PROOF. Using the full system of transversals $\tau = \cup_i \tau_i$, we define a finite set V of “vertices” as follows:

- Both endpoints of τ_i are elements V for all i ,

- the first intersection point with τ of any leaf ℓ of \mathcal{F} starting at any of the endpoints of any of the τ_i is an element of V .

The leaves running between the vertices (this includes all the finitely many singular leaves of \mathcal{F}) V delimit a finite system of rectangles, the union of whose transverse boundaries equals τ . Moreover, the union of these rectangles necessarily equals Σ_g . Indeed, if it didn't, the subsurface spanned by these rectangles would have a boundary component consisting of finitely many leaves. This boundary component would be homotopic to a closed curve contained in τ . This is impossible, because the components of τ need to be trees. The reason for this is that \mathcal{F} does not have any closed leaves. Indeed, any closed cycle in τ would need to be homotopically non-trivial, which would mean \mathcal{F} has a closed leaf homotopic to this cycle. \square

Now we go back to our mapping class $[\varphi]$ that preserves $[\mathcal{F}, \mu] \in \mathcal{MF}(\Sigma_g)$. We may take some power $[\varphi^n]$ that preserves the singularities and the singular leaves running between them. Since φ preserves the measure, we can assume that $\varphi^n(\tau) = \tau$. This means that φ^n maps each rectangle R_i to itself, fixing the boundary. By the Alexander lemma (Lemma 4.1.3), this means that φ^n is isotopic to the identity and hence that φ is periodic.

11.2.6. Case 2c. This, to some degree is the “generic” case and will yield pseudo-Anosov maps. In particular, we need to construct a singular measured foliation $(\widehat{\mathcal{F}}, \widehat{\mu})$ transverse to \mathcal{F} such that $[\varphi] \cdot (\widehat{\mathcal{F}}, \widehat{\mu}) = (\widehat{\mathcal{F}}, \frac{1}{\lambda} \widehat{\mu})$. The idea is to divide Σ_g into horizontal rectangles for \mathcal{F} which we will then foliate in the vertical direction to obtain $\widehat{\mathcal{F}}$. In order to determine $\widehat{\mu}$, we need to assign widths to these rectangles.

We can still construct a full system τ of transversals to the foliation \mathcal{F} . However, because $\lambda > 1$, we can only assume that $\varphi(\tau) \subset \tau$ as opposed to the equality we had before. Indeed, $\varphi_*\mu = \lambda \cdot \mu$, so φ needs to contract in the transverse direction to \mathcal{F} . We will not directly build our rectangles based on this, but transform the system of transversals first. This will guarantee that the rectangles we obtain have the right properties with respect to φ .

Again we let $n \geq 1$ be minimal such that φ^n preserves all the singularities of \mathcal{F} and maps the sectors of each singularity to themselves.

Because the leaves of \mathcal{F} are not closed, they are everywhere dense in Σ_g (compare to the case of foliations of the torus discussed in Section 6.1.1), once they do not run between two singularities of the foliation. In particular, every arc $\tau_i \subset \tau$ is intersected in its interior by some singular leaf ℓ_i – the existence of this leaf also uses that τ is a tree. We let α_i denote the segment of ℓ_i that runs between the singularity at which ℓ_i is based and the first point at which ℓ_i intersects τ_i .

This allows us to define

$$L = \cup_i \cup_{k=1}^n \varphi^k(\alpha_i).$$

Because φ stretches in the horizontal direction,

$$L \subset \varphi(L).$$

We will also shorten the arcs τ_i such that their endpoints lie on L .

Now we use Lemma 11.2.1 to fill Σ_g with a finite collection of horizontal rectangles R_1, \dots, R_N , using the arcs τ_i . We have that

$$\varphi(\partial^T R_i) \subset \bigcup_{i=1}^N \partial^T R_i.$$

Likewise, if we let $\partial^P R_i$ be the horizontal part of the boundary of R_i (parallel to the foliation), we have

$$\varphi^{-1}(\partial^P R_i) \subset \bigcup_{i=1}^N \partial^P R_i.$$

Now write

$$a_{ij} = \text{the number of connected components of } \varphi(\text{int}(R_i)) \cap \text{int}(R_j)$$

and we let $A = (a_{ij})_{ij}$ denote the corresponding \mathbb{N} -valued $N \times N$ matrix. Because the leaves of \mathcal{F} are dense in Σ_g , there is some $K \in \mathbb{N}$ such that A^K is a strictly positive matrix. By the Perron–Frobenius theorem, this means A has a unique eigenvalue ρ of maximal modulus. We let $v = (v_1, \dots, v_N)$ be any eigenvector corresponding to that eigenvalue, so by the same theorem, $v_1, \dots, v_N > 0$.

Now set the width of the rectangle R_i equal to v_i . This gives us a transverse measure $\hat{\mu}$ on the vertical foliation $\hat{\mathcal{F}}$ coming from the decomposition into rectangles. So,

$$\varphi \cdot [\hat{\mathcal{F}}, \hat{\mu}] = [\hat{\mathcal{F}}, \alpha \cdot \hat{\mu}].$$

To prove that $\alpha = \frac{1}{\lambda}$, we observe that μ and $\hat{\mu}$ together yield an area measure ν on Σ_g such that

$$\varphi_* \nu = \alpha \cdot \lambda \cdot \nu.$$

Since Σ is compact and φ is a diffeomorphism, this is only possible if $\alpha \cdot \lambda = 1$. So indeed φ is a pseudo-Anosov map.

LECTURE 12

Landmark results, part II

12.1. The Nielsen realization problem

We have seen above that if a non-trivial mapping class $[\varphi] \in \text{MCG}(\Sigma_g)$ has a fixed point $[X, f] \in \mathcal{T}(\Sigma_g)$, then $f \circ \varphi \circ f^{-1} : X \rightarrow X$ is isotopic to an orientation preserving isometry (or equivalently a conformal automorphism) of X . In particular, if X has a non-trivial isometry group G , then this yields a subgroup of the mapping class group. Such a subgroup is always finite because of the following classical result.

THEOREM 12.1.1 (Hurwitz). *Let X be a closed Riemann surface of genus $g \geq 2$ then*

$$\# \text{Aut}(X) = \# \text{Isom}^+(X) \leq 84(g - 1).$$

Here $\text{Isom}^+(X)$ denotes the group of orientation preserving isometries of the hyperbolic metric on X .

PROOF. We will use the hyperbolic metric for this proof. First we observe that if Y is an orientable hyperbolic 2-orbifold (like $\text{PSL}(2, \mathbb{Z}) \backslash \mathbb{H}^2$) then the area of Y equals

$$\text{area}(Y) = 2\pi(2g - 2 + n) + 2\pi \cdot \left(k - \sum_{i=1}^k \frac{1}{v_i} \right),$$

where g is the genus of Y , n its number of cusps, k its number of cone points and $2\pi/v_i$ the angle at the i^{th} cone point. We will prove this formula (in the closed case) in the exercises. It's also a consequence of the Gauss-Bonnet theorem. A case by case analysis of this formula shows that

$$\text{area}(Y) \geq \frac{\pi}{21}$$

and equality is uniquely realized by the sphere with three cone points such that $(v_1, v_2, v_3) = (2, 3, 7)$.

Our goal will be to show that $\text{Isom}^+(X) \backslash X$ is a closed orbifold. Indeed, then we can argue that

$$4\pi(g - 1) = \text{area}(X) = \# \text{Isom}^+(X) \cdot \text{area}(\text{Isom}^+(X) \backslash X) \geq \# \text{Isom}^+(X) \cdot \frac{\pi}{21},$$

which yields the claim.

In order to show that $\text{Isom}^+(X) \backslash X$ is indeed an orbifold, we need to show that $\text{Isom}^+(X)$ is finite. This we do by looking at the action of $\text{Isom}^+(X)$ on closed geodesics on X . Indeed, there can only be finitely many closed geodesics on X of bounded length – this is for instance a consequence of the proper discontinuity of the action of the fundamental

group of X on its universal cover, combined with the fact that every free homotopy class of non-trivial closed loops contains a unique geodesic.

So, because isometries preserve the lengths of geodesics, if γ is a closed geodesic on X , the orbit $\text{Isom}^+(X) \cdot \gamma$ consists of finitely many closed geodesics. Now we just need to prove that the stabilizer of a given geodesic is finite. To this end, we will use a geodesic γ with a single self intersection (like in Figure 1). Any isometry stabilizing γ needs to map




FIGURE 1. A closed curve with a single self intersection

its unique intersection point to itself. Moreover, it needs to map each of the four prongs emanating from the intersection point to another prong. In fact, the entire isometry is determined by where the prong goes. So, the size of the stabilizer is at most 4. \square

So, we conclude that isometry groups of surfaces yield finite subgroups of their mapping class groups. The *Nielsen realization problem* asks whether all finite subgroups of the mapping class group come from this construction. Indeed this is not obvious from the proof sketch above, because we had a second source of periodic maps (Case 2a).

Nielsen himself proved that his question has a positive answer if the finite subgroup is cyclic. The general case was settled by Kerckhoff [Ker83]:

THEOREM 12.1.2 (Kerckhoff). *Let $g \geq 2$ and let $G < \text{MCG}(\Sigma_g)$ be a finite subgroup. Then G can be lifted to a finite subgroup of $\text{Diff}^+(\Sigma_g)$ and moreover there exists a hyperbolic metric on Σ_g for which G acts by isometries.*

PROOF SKETCH. Let's start with proving the case in which $G < \text{MCG}(\Sigma_g)$ is a finite cyclic group. General theorems from topology (for instance using the fact that the group homology of G is non-trivial in infinitely many degrees) imply that G cannot act freely on a finite-dimensional contractible cell complex. In particular, if $h \in G$ is a generator, then there is some point $[X, f] \in \mathcal{T}(\Sigma_g)$ that is fixed by h (because Teichmüller space is a contractible cell complex). But since h is a generator, this means that all elements of G fix $[X, f]$ and hence that G can be lifted to a group of isometries of X .

For more general finite subgroups $G < \text{MCG}(\Sigma_g)$, the topological theorem is still true. So every element $h \in G$ has at least one fixed point. However, for the argument to work, some of these fixed points need to coincide. The mapping class group preserves much more structure on Teichmüller space than just its topology.




FIGURE 2. A filling collection of curves.

For instance it preserves the Teichmüller metric. If this were non-positively curved¹ then that would help us. For instance, let's consider the Euclidean metric. if some finite group G acts by isometries on \mathbb{R}^n then the circumcenter –the center of the unique ball of minimal radius that contains the orbit– of any orbit is a fixed point. This idea generalizes to more general non-positively curved spaces (including the hyperbolic plane and its higher-dimensional generalizations). Unfortunately, the Teichmüller metric is not non-positively curved.

So we need something else. The idea is to use sums of length functions over a filling collection of curves on Σ_g . Here, a set of non-trivial free homotopy classes of closed curves $\gamma_1, \dots, \gamma_k$ on Σ_g is said to be filling if the complement of the curves consists of disjoint disks (see Figure 2 for an example).

Given a finite subgroup $G < \text{MCG}(\Sigma_g)$, we take a G -invariant filling collection of curves $\Gamma = \{\gamma_1, \dots, \gamma_k\}$ on Σ_g . Such a collection can be obtained by taking the orbit under G of a collection of curves that is already filling. Now we define $\ell_\Gamma : \mathcal{T}(\Sigma_g) \rightarrow (0, \infty)$ by

$$\ell_\Gamma([X, f]) = \sum_{i=1}^k \ell_{\gamma_i}([X, f]).$$

The main theorem Kerckhoff proved is that such a function admits a unique minimum on Teichmüller space. Because ℓ_Γ itself is G -invariant, this minimum is G -invariant, i.e. a fixed point.

The way Kerckhoff proves that ℓ_Γ admits a unique minimum is by proving that it's a strictly convex function along deformations that are called *earthquakes*. These deformations are a generalization of twist deformations. Moreover, between any pair of points in Teichmüller space, there is a unique earthquake path. This implies that ℓ_Γ has a unique minimum. \square

12.2. Thurston's hyperbolization of mapping tori

The next big result we'll discuss is about 3-manifolds. Given an orientation preserving diffeomorphism $f : \Sigma_g \rightarrow \Sigma_g$, we can build a 3-manifold that we will call the *mapping torus* of f . This is the manifold

$$M_f = \left(\Sigma_g \times [0, 1] \right) / (x, 0) \sim (f(x), 0) \forall x \in \Sigma_g.$$

¹There is something to say about how to formalize this, because the Teichmüller metric is not a Riemannian.

Observe that such a manifold admits a map to the circle given by $[(x, t)] \mapsto [t] \in \mathbb{R}/\mathbb{Z}$. We also observe that if f and f' are isotopic, then M_f and M'_f are homeomorphic. Indeed, the isotopy is a map $F = \Sigma_g \times [0, 1] \rightarrow \Sigma$ such that $F(\cdot, 0) = f$ and $F(\cdot, 1) = f'$. So the homeomorphism can be built by attaching an extra “slab” homeomorphic to $\Sigma_g \times [0, 1]$ to the manifold before gluing.

Just like in dimension two, one can ask which type of geometric structures one can put on such a manifold. In dimension two, we had three types of geometries: spherical, euclidean and hyperbolic.

In dimension three, the situation is more complicated: there are eight types of geometries (three among them are still spherical, euclidean and hyperbolic) and it's no longer true that every manifold can be equipped with a geometric structure. Instead, every closed 3-manifold can be cut along spheres and tori in such a way that the pieces admit a geometric structure. This is the content of Perelman's geometrization theorem (which had been conjectured by Thurston).

We will discuss an important special case of this theorem, namely mapping tori. Let us first very briefly discuss 3-dimensional hyperbolic geometry. Like in dimension two, there are multiple models for hyperbolic 3-space, for instance

$$\mathbb{H}^3 = \{ (x, y, z) \in \mathbb{R}^3 : z > 0 \}, \quad ds_{\mathbb{H}^3}^2 = \frac{dx^2 + dy^2 + dz^2}{z^2}$$

and

$$\mathbb{B}^3 = \{ (x, y, z) \in \mathbb{R}^3 : x^2 + y^2 + z^2 < 1 \}, \quad ds_{\mathbb{B}^3}^2 = 4 \frac{dx^2 + dy^2 + dz^2}{(1 - x^2 - y^2 - z^2)^2}$$

are the 3-dimensional versions of the upper half plane model and the disk model respectively. More generally, the Killing–Hopf theorem also applies in three dimensions, so every complete simply connected Riemannian manifold is isometric to \mathbb{H}^3 . The boundary of \mathbb{H}^3 can be identified with the Riemann sphere $\widehat{\mathbb{C}}$. It turns out that the isometry group of \mathbb{H}^3 can be identified with the group of complex automorphisms of its boundary. That is $\text{Isom}^+(\mathbb{H}^3) \simeq \text{PSL}(2, \mathbb{C})$. Moreover, the non-trivial elements of $\text{PSL}(2, \mathbb{C})$ can still be classified into three categories:

- *elliptic elements* that have at least one fixed point in \mathbb{H}^3 ,
- *parabolic elements* that have infinite order, no fixed points in \mathbb{H}^3 and exactly one fixed point in $\widehat{\mathbb{C}}$
- *hyperbolic elements* that have positive translation length, no fixed points in \mathbb{H}^3 and exactly two fixed points in $\widehat{\mathbb{C}}$.

Moreover, the type of an element is again determined by its trace: if $\text{tr}(g) \in (-2, 2)$ then g is elliptic, if $\text{tr}(g) \in \{-2, 2\}$ then g is parabolic and if neither of the two is true then g is hyperbolic.

So if M is a closed hyperbolic 3-manifold, then there exists some $\Gamma < \text{PSL}(2, \mathbb{C})$ that is discrete and torsion free, such that M is isometric to $\Gamma \backslash \mathbb{H}^3$. In every non-trivial free homotopy class of closed curves in our closed manifold M , there exists a unique geodesic and the corresponding conjugacy class of Γ consists of hyperbolic elements.

In general, the study of hyperbolic 3-manifolds is very different from that of hyperbolic surfaces. Indeed, in dimension 3 (and higher) there is the *Mostow rigidity theorem* which states that if M and N are two closed hyperbolic 3-manifolds such that $\pi_1(M) \simeq \pi_1(N)$ as abstract groups, then M and N are isometric. In other words, the moduli space of hyperbolic metrics on any closed 3-manifold is a point.

Now we get back to our question of what kind of geometric structure our manifold M_f admits. We will ask when it admits a hyperbolic structure. First of all we observe:

LEMMA 12.2.1. *If $[f] \in \text{MCG}(\Sigma_g)$ is not pseudo-Anosov, then M_f does not admit a hyperbolic metric.*

PROOF. Our first claim is that, if M_f is not pseudo-Anosov, then M_f admits a finite cover $N \rightarrow M_f$ such that $\pi_1(N)$ has a subgroup that is isomorphic to \mathbb{Z}^2 .

Indeed, by the Nielsen–Thurston classification, if f is not pseudo-Anosov, then f is either periodic or reducible. We start with the periodic case. Suppose that $[f]^n$ is the identity in $\text{MCG}(\Sigma_g)$. Observe that in general M_{f^n} is an n -fold cover of M_f . If f^n is the identity map, that means that $M_{f^n} = \Sigma_g \times \mathbb{S}^1$. So its fundamental group is $\pi_1(\Sigma_g) \times \pi_1(\mathbb{S}^1)$, which contains plenty of copies of \mathbb{Z}^2 (generated by any pair of non trivial elements $g \in \pi_1(\Sigma_g)$ and $h \in \pi_1(\mathbb{S}^1) \simeq \mathbb{Z}$).

If $[f]$ is reducible, that implies it has some power $[f]^n$ that preserves a curve $\gamma \subset \Sigma_g$. Essentially for the same reason as before, any group element in the conjugacy class of $\pi(\Sigma_g)$ determined by γ and any non trivial element of $\pi_1(\mathbb{S}^1)$ together generate a copy of \mathbb{Z}^2 .

Now we need to show that the fundamental group Γ of a closed hyperbolic 3-manifold $\Gamma \backslash \mathbb{H}^3$ cannot contain a subgroup isomorphic to \mathbb{Z}^2 . There are many ways to do this. We will use the fact that all non-trivial elements of $\text{PSL}(2, \mathbb{C})$ are hyperbolic. Now if two hyperbolic elements $g, h \in \text{PSL}(2, \mathbb{C})$ commute, then they need to preserve each others fixed points. This means that they translate along the same axis. But then the group $\langle g, h \rangle$ generated by g and h can only be discrete if there exist some $m, n \in \mathbb{Z}$ such that $g^m = h^n$. That means that $\langle g, h \rangle \not\simeq \mathbb{Z}^2$. \square

The surprising fact is that this “obvious” obstruction is the only one:

THEOREM 12.2.2 (Thurston’s hyperbolization for mapping tori). *Let $g \geq 2$ and let $[f] \in \text{MCG}(\Sigma_g)$. Then M_f admits a hyperbolic metric if and only if $[f]$ is pseudo-Anosov.*

Otal [Ota96] was the first to provide a full proof of this theorem, based on the ideas by Thurston. We will not go into the proof, but will just mention that much more recently it was proved that mapping tori (i.e. 3-manifolds that fiber over the circle) are not as restrictive a class of manifolds as they may seem at first sight. Indeed, Agol [Ago13], using work by Wise and many others proved that every closed hyperbolic 3-manifold (by many measures the wildest class of 3-manifolds) admits a cover of finite degree by a mapping torus, thus confirming another conjecture by Thurston.

12.3. The number of short simple closed geodesics

The final big result we will discuss is Mirzakhani's result on the number of simple closed geodesics of bounded length on a hyperbolic surface.

In her thesis, Mirzakhani proved a sequence of groundbreaking results [Mir07a, Mir07b, Mir08]. She found a method to compute volumes of moduli spaces of hyperbolic surfaces (with respect to a volume form called the *Weil–Petersson volume form*), she found a new proof of Witten's conjecture on intersection numbers of tautological classes on $\mathcal{M}_{g,n}$ and she determined the asymptotics of the number of simple closed geodesics of bounded length on a hyperbolic surface. The first two results were discussed in Nalini Anantharaman's class on random hyperbolic surfaces. We will discuss Mirzakhani's result on counting simple closed geodesics.

For some context, let us first describe what happens when we count all short closed geodesics. For a hyperbolic surface X and $L > 0$, set

$$c_X(L) = \# \left\{ \begin{array}{l} \gamma \text{ a closed geodesic on } X \\ \text{of length } \leq L \end{array} \right\},$$

where we do not distinguish between geodesics γ and γ' if one can be obtained by reparametrizing the other, but do distinguish between two geodesics with opposite orientations.

A classical result, originally due to Huber [Hub56], states that on a closed hyperbolic surface²

$$c_X(L) \sim \frac{e^L}{L} \quad \text{as } L \rightarrow \infty.$$

This is a somewhat surprising result, because it is completely independent of the geometry and topology of X . Huber used the Selberg trace formula to derive his result.

Now let us discuss simple closed geodesics. Given a complete hyperbolic surface X of finite area, we will set:

$$s_X(L) = \# \left\{ \begin{array}{l} \gamma \text{ a simple closed geodesic} \\ \text{on } X \text{ of length } \leq L \end{array} \right\}.$$

In the function s_X , we will *not* distinguish between different orientations on our closed geodesics. This is just because this fits better with Mirzakhani's methods. In any event, the difference between the two functions is a factor 2.

It turns out that the number of simple geodesics grows much slower than all geodesics [Mir08, Theorem 1.1]:

THEOREM 12.3.1. *Let $X \in \mathcal{M}_{g,n}$ be a hyperbolic surface. Then there exists a constant $b_X > 0$ so that*

$$s_X(L) \sim b_X \cdot L^{6g-6+2n} \quad \text{as } L \rightarrow \infty.$$

The proof of this theorem is based on the convergence of a sequence of measures on the space of measured geodesic laminations on $\Sigma_{g,n}$. In particular, the methods behind which are wildly different from those behind Huber's result.

²Recall that the notation " $f(x) \sim g(x)$ as $x \rightarrow \infty$ " means that $\lim_{x \rightarrow \infty} f(x)/g(x) = 1$.

12.3.1. Measured laminations. We will very briefly sketch what measured laminations and their key properties are. For a more background, we refer to [CB88, PH92, AL17].

A *geodesic lamination* on a hyperbolic surface X is a closed subset that is a disjoint union of complete simple geodesics that do not intersect each other.

One example of a geodesic lamination is a collection of disjoint simple closed geodesics (a *multicurve*), like for instance a pants decomposition. Figure 3 shows another example.




FIGURE 3. A lamination consisting of two simple closed geodesics and one geodesic arc. The left hand side consists of three lifts of the components to \mathbb{H}^2

A *measured geodesic lamination* on X is a geodesic lamination on X , together with an assignment

$$\alpha \mapsto \lambda_\alpha$$

that assigns a Radon measure λ_α to each arc α that is transverse to the lamination. This assignment has to satisfy the following conditions:

- Subarcs: If $\alpha' \subset \alpha$ is a subarc, then $\lambda_{\alpha'}$ is the restriction of λ_α to α'
- Sliding: If α and α' are homotopic via a homotopy F_t so that $F_1 : \alpha \rightarrow \alpha'$ is a homeomorphism and $F_t(\alpha)$ is transverse to the leaves of the lamination for all t , then

$$\lambda_{\alpha'} = (F_1)_* \lambda_\alpha.$$

Note that it follows that the support of the measure λ_α is contained in the intersection of α with the lamination.

The set of measured laminations $\mathcal{ML}(X)$ can be topologized using weak star convergence. It turns out that this topology does not depend on the geometry of X . So we will often write $\mathcal{ML}_{g,n}$ for the space of measured laminations. There are coordinates, called Dehn-Thurston coordinates, that are somewhat similar to Fenchel-Nielsen coordinates and that provide a homeomorphism

$$\mathcal{ML}_{g,n} \rightarrow \mathbb{R}_+^{3g-3+n} \times \mathbb{R}^{3g-3+n}.$$

Unlike Fenchel-Nielsen coordinates, these coordinates do however not give $\mathcal{ML}_{g,n}$ the structure of a smooth manifold.

Note that \mathbb{R}_+ acts on $\mathcal{ML}_{g,n}$, by multiplying the measures with a scalar. $\text{MCG}(\Sigma_{g,n})$ acts on $\mathcal{ML}_{g,n}$ as well.

Every multicurve induces a measured lamination. The lamination is a multicurve and the measure assigned to an arc α is the sum Dirac masses on the intersections of α and the multicurve. We will denote this element of $\mathcal{ML}_{g,n}$ by

$$\sum_i \gamma_i \in \mathcal{ML}_{g,n},$$

where $\{\gamma_i\}_i$ are the geodesics on X that form the multicurve.

For us, the most important example of a measured lamination is a weighted multicurve, i.e. we scale the measures on arcs transverse to γ_i with a number $t_i \in \mathbb{R}_+$. These measured laminations will be denoted

$$\sum_i t_i \gamma_i \in \mathcal{ML}_{g,n}.$$

It turns out that the set of such measures is actually dense in $\mathcal{ML}_{g,n}$.

The length function on geodesics can naturally be extended to weighted multicurves, by setting

$$\ell\left(\sum_i t_i \gamma_i\right) = \sum_i t_i \ell(\gamma_i).$$

It turns out that ℓ extends to a continuous function

$$\ell : \mathcal{ML}_{g,n} \times \mathcal{T}_{g,n} \rightarrow \mathbb{R}_+$$

that satisfies:

- For any simple closed curve γ

$$\ell(\gamma, X) = \ell_\gamma(X).$$

- For all $t \in \mathbb{R}_+$, $\lambda \in \mathcal{ML}_{g,n}$, $X \in \mathcal{T}_{g,n}$

$$\ell(t\lambda, X) = t\ell(\lambda, X).$$

- For all $\varphi \in \text{MCG}(\Sigma_{g,n})$, $\lambda \in \mathcal{ML}_{g,n}$

$$\ell(\varphi\lambda, \varphi X) = \ell(\lambda, X)$$

Finally, $\mathcal{ML}_{g,n}$ can be equipped with a $\text{MCG}(\Sigma_{g,n})$ -invariant measure μ_{Th} called the *Thurston measure*. This measure satisfies

$$\mu_{\text{Th}}(tU) = t^{6g-6+2n} \mu_{\text{Th}}(U)$$

and [Mas85]:

THEOREM 12.3.2 (Masur). *Suppose ν is a measure on $\mathcal{ML}_{g,n}$ that is absolutely continuous³ with respect to μ_{Th} and $\text{MCG}(\Sigma_{g,n})$ -invariant. Then ν is a multiple of μ_{Th} .*

³recall that a measure λ is absolutely continuous with respect to another measure ν if $\nu(A) = 0$ implies that $\lambda(A) = 0$ for all measurable sets A .

12.3.2. The proof. The proof of Theorem 12.3.1 now goes as follows. Once and for all fix a simple closed curve γ on $\Sigma_{g,n}$ and for all $L > 0$, define a measure $\mu_{L,\gamma}$ on $\mathcal{ML}_{g,n}$ by:

$$\mu_{L,\gamma} = \frac{1}{L^{6g-6+2n}} \sum_{\alpha \in \text{MCG}(\Sigma_{g,n}) \cdot \gamma} \delta_{\frac{1}{L}\gamma},$$

where for $\lambda \in \mathcal{ML}_{g,n}$, δ_λ denotes the Dirac mass on λ .

For $X \in \mathcal{T}_{g,n}$, write

$$B_X = \{ \lambda \in \mathcal{ML}_{g,n} : \ell(\lambda, X) \leq 1 \}$$

We have

$$\begin{aligned} \mu_{L,\gamma}(B_X) &= \frac{1}{L^{6g-6+2n}} \sum_{\alpha \in \text{MCG}(\Sigma_{g,n}) \cdot \gamma} \mathbf{1}_{\frac{1}{L}\gamma \text{ has length } \leq 1} \\ &= \frac{1}{L^{6g-6+2n}} \sum_{\alpha \in \text{MCG}(\Sigma_{g,n}) \cdot \gamma} \mathbf{1}_{\gamma \text{ has length } \leq L} \\ &=: \frac{1}{L^{6g-6+2n}} s_X(L, \gamma). \end{aligned}$$

There are finitely many mapping class group orbits of simple closed curves on $\Sigma_{g,n}$, which means that $s_X(L)$ splits as a finite sum of measures of the form above.

So, if we can prove that, as $L \rightarrow \infty$ the measures $\mu_{L,\gamma}$ converge to fixed measures μ_γ , we would get

$$s_X(L, \gamma) \sim \mu_\gamma(B_X) L^{6g-6+2n}$$

as $L \rightarrow \infty$. Then, since $s_X(L)$ is a finite sum of the $s_X(L, \gamma)$'s, we would be done.

The proof of this goes as follows.

Step 1: One proves that there exists a constant $C(X, \gamma) > 0$ so that

$$\mu_{L,\gamma}(B_X) = \frac{s_X(L, \gamma)}{L^{6g-6+2n}} \leq C(X, \gamma).$$

Since any compact set K lies in $L \cdot B_X$ for L large enough, this proves that for any compact set K ,

$$\mu_{L,\gamma}(K)$$

is bounded as a function of L and as such (by Banach-Alaoglu) there is a subsequential limit

$$\mu_{L_n,\gamma} \rightarrow \nu.$$

Step 2: We need to prove that all such limits are multiples of the Thurston measure. It is not so hard to see that the limit is $\text{MCG}(\Sigma_{g,n})$ -invariant. So, we need to show that it is absolutely continuous with respect to μ_{Th} , so that we can invoke Masur's theorem (Theorem 12.3.2). For this, we refer to [Mir08].

Step 3: Now that we know all our subsequential limits are multiples of the Thurston measure, we need to prove that *which* multiple it is, does not depend on the subsequence. This is where Weil–Petersson volumes come in.

Let $(L_k)_k$ be so that

$$\mu_{L_k, \gamma} \rightarrow R_{(L_k)_k} \mu_{\text{Th}},$$

as $n \rightarrow \infty$.

Now let \mathcal{P} be a pants decomposition of $\Sigma_{g,n}$ and let $(\ell_{\alpha_i}, \tau_{\alpha_i})_i$ be the corresponding Fenchel–Nielsen coordinates on $\mathcal{T}(\Sigma_{g,n})$. Wolpert [Wol82] proved that the volume form

$$\bigwedge_{i=1}^{3g-3+n} d\ell_{\alpha_i} \wedge d\tau_{\alpha_i}$$

on Teichmüller space descends to a form of finite total volume on $\mathcal{M}_{g,n}$. Moreover, similar forms can be defined on Teichmüller and moduli spaces of surfaces with boundary whose boundary lengths are fixed.

In her thesis, Mirzakhani proved how to integrate a certain type of functions that she calls *geometric functions*. This construction is very similar to the Siegel transform in the context of the moduli space of lattices. So, it also makes sense to call this the *Mirzakhani transform*. Let $F : \mathbb{R}_{\geq 0}^k \rightarrow \mathbb{R}$ be a function and write $\mathcal{M}_{g,n}(L_1, \dots, L_n)$ for the moduli space of hyperbolic surfaces with n boundary components of length L_1, \dots, L_n respectively (we allow $L_i = 0$, in which case the corresponding boundary component is a cusp). Given an ordered k -tuple of simple closed curves $\Gamma = (\gamma_1, \dots, \gamma_k)$ on $\Sigma_{g,n}$, we can define the function $F^\Gamma : \mathcal{M}_{g,n}(L_1, \dots, L_n) \rightarrow \mathbb{R}$ by

$$F^\Gamma(X) = \sum_{(\alpha_1, \dots, \alpha_k) \in \text{MCG}(\Sigma_{g,n}) \cdot (\gamma_1, \dots, \gamma_k)} F(\ell_{\alpha_1}(X), \dots, \ell_{\alpha_k}(X)), \quad X \in \mathcal{M}_{g,n}(L_1, \dots, L_n).$$

Note that the length functions ℓ_{α_i} are only defined on Teichmüller space. However, because we sum over the whole mapping class group orbit, the function above is well-defined on moduli space.

Our function $X \mapsto s_X(L)$ is a finite sum such functions. Indeed, we again split our set of isotopy classes of simple closed curves on $\Sigma_{g,n}$ into finite number of mapping class group orbits. Given such an orbit, we pick a curve γ in it and set function $F = \chi_{[0, L]}$, the indicator function of the interval $[0, L]$, such that $s_X(L, \gamma) = F^\gamma(X)$.

Using a very nice covering argument, Mirzakhani proved:

THEOREM 12.3.3. *With the notation from above and for $L = (L_1, \dots, L_n)$:*

$$\int_{\mathcal{M}_{g,n}(L)} F^\Gamma(X) d\text{vol}_{\text{WP}}(X) = C_\Gamma \int_{\mathbb{R}_{\geq 0}^k} F(x) \cdot \text{vol}_{\text{WP}} \left(\mathcal{M}(\Sigma_{g,n} - \Gamma, x, x, L) \right) \cdot x_1 \cdots x_k dx_1 \cdots dx_k,$$

where $C_\Gamma > 0$ is a constant depending on Γ only, that can be computed effectively and $\mathcal{M}(\Sigma_{g,n} - \Gamma, x, x)$ denotes the moduli space of hyperbolic metrics on $\Sigma_{g,n} - \bar{\Gamma}$ whose boundary components corresponding to Γ have lengths given by (x, x) (they come in pairs) and those corresponding to $\Sigma_{g,n}$ have lengths given by L .

Using this, we can compute:

$$\begin{aligned}
R_{(L_k)_k} \int_{\mathcal{M}_{g,n}} \mu_{\text{Th}}(B_X) d \text{vol}_{\text{WP}}(X) &= \int_{\mathcal{M}_{g,n}} \lim_{k \rightarrow \infty} \mu_{L_k, \gamma}(B_X) d \text{vol}_{\text{WP}}(X) \\
&= \int_{\mathcal{M}_{g,n}} \lim_{k \rightarrow \infty} \frac{s_X(L_k, \gamma)}{L_n^{6g-6+2n}} d \text{vol}_{\text{WP}}(X) \\
&= \lim_{n \rightarrow \infty} \int_{\mathcal{M}_{g,n}} \frac{s_X(L_k, \gamma)}{L_n^{6g-6+2n}} d \text{vol}_{\text{WP}}(X) \\
&= \lim_{k \rightarrow \infty} \int_0^{L_k} \frac{1}{L_k^{6g-6+2n}} \text{vol}_{\text{WP}}(\mathcal{M}(\Sigma \setminus \gamma, x, x)) x d \text{vol}_{\text{WP}}(X).
\end{aligned}$$

Now we use that $\text{vol}_{\text{WP}}(\mathcal{M}(\Sigma \setminus \gamma, x, x))$ is a polynomial of the correct degree (another theorem from Mirzakhani's thesis) in x and see that the right hand side converges to a constant that is independent of the sequence L_k . The integral

$$\int_{\mathcal{M}_{g,n}} \mu_{\text{Th}}(B_X) d \text{vol}_{\text{WP}}(X)$$

also doesn't depend on the sequence, so neither can $R_{(L_k)_k}$, which finishes the proof.

Bibliography

- [Ago13] Ian Agol. The virtual Haken conjecture. *Doc. Math.*, 18:1045–1087, 2013. With an appendix by Agol, Daniel Groves, and Jason Manning.
- [AL17] Javier Aramayona and Christopher J. Leininger. Hyperbolic structures on surfaces and geodesic currents. In *Algorithmic and geometric topics around free groups and automorphisms*, Adv. Courses Math. CRM Barcelona, pages 111–149. Birkhäuser/Springer, Cham, 2017.
- [Baa21] Sebastian Baader. *Geometry and topology of surfaces*. Zurich Lectures in Advanced Mathematics. EMS Press, Berlin, [2021] ©2021.
- [Bea84] A. F. Beardon. *A primer on Riemann surfaces*, volume 78 of *London Mathematical Society Lecture Note Series*. Cambridge University Press, Cambridge, 1984.
- [Bus10] Peter Buser. *Geometry and spectra of compact Riemann surfaces*. Modern Birkhäuser Classics. Birkhäuser Boston, Inc., Boston, MA, 2010. Reprint of the 1992 edition.
- [CB88] Andrew J. Casson and Steven A. Bleiler. *Automorphisms of surfaces after Nielsen and Thurston*, volume 9 of *London Mathematical Society Student Texts*. Cambridge University Press, Cambridge, 1988.
- [CE08] Jeff Cheeger and David G. Ebin. *Comparison theorems in Riemannian geometry*. AMS Chelsea Publishing, Providence, RI, 2008. Revised reprint of the 1975 original.
- [Fal23] Elisha Falbel. Introduction to Riemann surfaces. Lecture notes, available at: <https://webusers.imj-prg.fr/~elisha.falbel/Surfaces/RS.html>, 2023.
- [FK92] H. M. Farkas and I. Kra. *Riemann surfaces*, volume 71 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, second edition, 1992.
- [FLP79] *Travaux de Thurston sur les surfaces*, volume 66-67 of *Astérisque*. Société Mathématique de France, Paris, 1979. Séminaire Orsay, With an English summary.
- [FLP12] Albert Fathi, François Laudenbach, and Valentin Poénaru. *Thurston’s work on surfaces*, volume 48 of *Mathematical Notes*. Princeton University Press, Princeton, NJ, 2012. Translated from the 1979 French original by Djun M. Kim and Dan Margalit.
- [FM12] Benson Farb and Dan Margalit. *A primer on mapping class groups*, volume 49 of *Princeton Mathematical Series*. Princeton University Press, Princeton, NJ, 2012.
- [GGD12] Ernesto Gironde and Gabino González-Diez. *Introduction to compact Riemann surfaces and dessins d’enfants*, volume 79 of *London Mathematical Society Student Texts*. Cambridge University Press, Cambridge, 2012.
- [GH94] Phillip Griffiths and Joseph Harris. *Principles of algebraic geometry*. Wiley Classics Library. John Wiley & Sons, Inc., New York, 1994. Reprint of the 1978 original.
- [GL00] Frederick P. Gardiner and Nikola Lakic. *Quasiconformal Teichmüller theory*, volume 76 of *Mathematical Surveys and Monographs*. American Mathematical Society, Providence, RI, 2000.
- [Hat02] Allen Hatcher. *Algebraic topology*. Cambridge University Press, Cambridge, 2002.
- [HM79] John Hubbard and Howard Masur. Quadratic differentials and foliations. *Acta Math.*, 142(3-4):221–274, 1979.
- [Hub56] Heinz Huber. Über eine neue Klasse automorpher Funktionen und ein Gitterpunktproblem in der hyperbolischen Ebene. I. *Comment. Math. Helv.*, 30:20–62 (1955), 1956.
- [Hub06] John Hamal Hubbard. *Teichmüller theory and applications to geometry, topology, and dynamics. Vol. 1*. Matrix Editions, Ithaca, NY, 2006. Teichmüller theory, With contributions by Adrien Douady, William Dunbar, Roland Roeder, Sylvain Bonnot, David Brown, Allen Hatcher, Chris Hruska and Sudeb Mitra, With forewords by William Thurston and Clifford Earle.

- [Hub16] John Hamal Hubbard. *Teichmüller theory and applications to geometry, topology, and dynamics. Vol. 2*. Matrix Editions, Ithaca, NY, 2016. Surface homeomorphisms and rational functions.
- [IT92] Y. Imayoshi and M. Taniguchi. *An introduction to Teichmüller spaces*. Springer-Verlag, Tokyo, 1992. Translated and revised from the Japanese by the authors.
- [Ker83] Steven P. Kerckhoff. The Nielsen realization problem. *Ann. of Math. (2)*, 117(2):235–265, 1983.
- [Mas85] Howard Masur. Ergodic actions of the mapping class group. *Proc. Amer. Math. Soc.*, 94(3):455–459, 1985.
- [Mil65] John W. Milnor. *Topology from the differentiable viewpoint*. University Press of Virginia, Charlottesville, VA, 1965. Based on notes by David W. Weaver.
- [Mir07a] Maryam Mirzakhani. Simple geodesics and Weil-Petersson volumes of moduli spaces of bordered Riemann surfaces. *Invent. Math.*, 167(1):179–222, 2007.
- [Mir07b] Maryam Mirzakhani. Weil-Petersson volumes and intersection theory on the moduli space of curves. *J. Amer. Math. Soc.*, 20(1):1–23, 2007.
- [Mir08] Maryam Mirzakhani. Growth of the number of simple closed geodesics on hyperbolic surfaces. *Ann. of Math. (2)*, 168(1):97–125, 2008.
- [Ota96] Jean-Pierre Otal. Le théorème d’hyperbolisation pour les variétés fibrées de dimension 3. *Astérisque*, (235):x+159, 1996.
- [PH92] R. C. Penner and J. L. Harer. *Combinatorics of train tracks*, volume 125 of *Annals of Mathematics Studies*. Princeton University Press, Princeton, NJ, 1992.
- [Sma59] Stephen Smale. Diffeomorphisms of the 2-sphere. *Proc. Amer. Math. Soc.*, 10:621–626, 1959.
- [SS03] Elias M. Stein and Rami Shakarchi. *Complex analysis*, volume 2 of *Princeton Lectures in Analysis*. Princeton University Press, Princeton, NJ, 2003.
- [Str84] Kurt Strebel. *Quadratic differentials*, volume 5 of *Ergebnisse der Mathematik und ihrer Grenzgebiete (3) [Results in Mathematics and Related Areas (3)]*. Springer-Verlag, Berlin, 1984.
- [Wol82] Scott Wolpert. The Fenchel-Nielsen deformation. *Ann. of Math. (2)*, 115(3):501–528, 1982.
- [Wri15] Alex Wright. Translation surfaces and their orbit closures: an introduction for a broad audience. *EMS Surv. Math. Sci.*, 2(1):63–108, 2015.
- [Zor06] Anton Zorich. Flat surfaces. In *Frontiers in number theory, physics, and geometry. I*, pages 437–583. Springer, Berlin, 2006.