MODULI SPACES OF RIEMANN SURFACES

ALEX WRIGHT

Contents

These are course notes for a one quarter long second year graduate class at Stanford University taught in Winter 2017.

All of the material presented is classical. For the softer aspects of Teichmüller theory, our reference was the book of Farb-Margalit [**FM12**], and for Teichmüller theory as a whole our comprehensive reference was the book of Hubbard [**Hub06**]. We also used Gardiner and Lakic's book [**GL00**] and McMullen's course notes [**McM**] as supplemental references. For period mappings our reference was the book [**CMSP03**]. For Geometric Shafarevich our reference was McMullen's survey [**McM00**].

The author thanks Steve Kerckhoff, Scott Wolpert, and especially Jeremy Kahn for many helpful conversations.

These are rough notes, hastily compiled, and may contain errors. Corrections are especially welcome, since these notes will be re-used in the future.

1. Different ways to build and deform Riemann surfaces

1.1. **Uniformization.** A Riemann surface is a topological space with an atlas of charts to the complex plane $\mathbb{C}$ whose translation functions are biholomorphisms. Any open subset of the Riemann sphere $\mathbb{P}^1$ is Riemann surface, as is the complex torus $\mathbb{C}/\mathbb{Z}[i]$. In this course we will be interested in moduli (parameter) spaces of Riemann surfaces, especially the space of all Riemann surfaces of fixed genus. Thus, typically we will restrict to closed connected surfaces, but for context it helps to start out more generally. Complex analysis starts to turn into Teichmüller theory with the following theorem and exercise.

Theorem 1.1 (Uniformization)**.** *Every simply connected Riemann surface is biholomorphic to $\mathbb{C}$, $\mathbb{P}^1$, or the upper half plane*

$$\mathbb{H} = \{z \in \mathbb{C} : \mathrm{Re}(z) > 0\}.$$

Thus in particular the moduli space of closed genus 0 Riemann surfaces is a point.

Exercise 1.2. Two annuli $\{z : r < |z| < R\}$ and $\{z : r' < |z| < R'\}$ are biholomorphic if and only if $R/r = R'/r'$. (The case when $r = 0$ and $R = \infty$ is an exception, since the punctured plane is not biholomorphic to the punctured disc.)

Thus, the moduli space of Riemann surfaces that are homeomorphic to annuli at least contains a copy of $(1, \infty]$; later you will show it is equal to $(1, \infty]$ (if we leave out the punctured plane).

The group of biholomorphic automorphisms of $\mathbb{P}^1$ is $PSL(2, \mathbb{C})$, acting via Möbius transformations. The group of biholomorphic automorphisms of $\mathbb{C}$ is the affine group, and the group of biholomorphic automorphisms of $\mathbb{H}$ is $PSL(2, \mathbb{R})$. The hyperbolic metric

$$\frac{dx^2 + dy^2}{y^2}$$

on $\mathbb{H}$ is preserved by $PSL(2, \mathbb{R})$, and in fact $PSL(2, \mathbb{R})$ is the group of orientation preserving isometries of $\mathbb{H}$.

Lemma 1.3. *Every Riemann surface of genus at least 2 has a metric of constant curvature -1 in its conformal class.*

Note, the conformal class of a Riemann surface is the set of all metrics on the surface that give the same angles between tangent vectors as the conformal structure. In a holomorphic coordinate $dz = dx + idy$, all metrics in the conformal class are of the form $e^h(dx^2 + dy^2)$, where h is a function.

Proof. The universal cover of the Riemann surface is $\mathbb{H}$, and the Deck group acts via orientation preserving biholomorphisms. Hence the Deck group acts via hyperbolic isometries. □

Exercise 1.4. Prove uniqueness of the hyperbolic metric.

Similarly one may show that every genus 1 Riemann surface has a (unique) flat metric.

We have shown that every Riemann surface of genus at least 2 is of the form $\mathbb{H}/\Gamma$, where $\Gamma \subset PSL(2,\mathbb{R})$ is the Deck group. Varying the matrix entries of generators of the Deck group, subject to the relation that holds on the generators, varies the Riemann surface.

A discrete subgroup of $PSL(2,\mathbb{R})$ is called a Fuchsian group.

1.2. **Algebraic curves.** Consider the zero set of a homogeneous polynomial p in $\mathbb{P}^2$. If the polynomial is non-singular, i.e. if the gradient and the polynomial do not vanish at the same time, the implicit function theorem gives that the zero set of the polynomial is a closed surface, and in fact the result is a Riemann surface. More generally, one can consider systems of equations on $\mathbb{P}^n$.

In the same way that an abstract surface is not the same thing as a surface embedded in $\mathbb{R}^n$, a Riemann surface is not the same thing as a Riemann surface cut out by specific polynomials in projective space; the later is called a (non-singular) algebraic curve.

Now suppose we change (perturb) the coefficients of the polynomial(s) defining the algebraic curve. In general, the Riemann surface may or may not change, although in some sense one expects it to. In the same way that an abstract surface can be embedded in $\mathbb{R}^n$ many different ways, a Riemann surface may have many distinct realizations as an algebraic curve. It is however a basic result in the theory of algebraic curves that every Riemann surface can be realized as an algebraic curve.

Philosophically, the situation can be compared to finite group theory. If one wishes to understand finite matrix groups, one may divide the study into two halves: classify all abstract groups, and then classify the linear representations of groups. For algebraic curves, the abstract portion is the study of moduli space, and the second portion is called Brill-Noether theory.

1.3. **Covers of $\mathbb{P}^1$.** Consider a finite degree branched cover of $\mathbb{P}^1$. Branched means that after removing a finite set of points from $\mathbb{P}^1$ and their pre-images, the map is an honest covering map. The points of the cover where the map is not a local diffeo are called the ramification points; their images are called the branch points. At each ramification point, the map looks locally like $z \mapsto z^k$.

Any finite degree branched cover of $\mathbb{P}^1$ may be assigned the structure of a Riemann surface in such a way that the covering map is holomorphic. Varying the branch points may or may not change the Riemann surface, although again one expects it to.

There is a Galois correspondence between finite branched covers of $\mathbb{P}^1$ and finite field extensions of $\mathbb{C}(z)$. The correspondence maps a branched cover to its field of meromorphic functions.

1.4. **Hyperbolic hexagons and pants.** See [**FM12**, Proposition 10.4] for an elegant and soft proof that, for any $\ell_1, \ell_2, \ell_3 > 0$, there is a unique hyperbolic hexagon with three non-adjacent sides of lengths ℓ_i. Gluing two such together along the remaining 3 sides gives a hyperbolic pair of pants with cuff lengths ℓ_1, ℓ_2, ℓ_3. Gluing together pants can give a hyperbolic surface.

Exercise 1.5. Prove that every closed hyperbolic surface of genus at least 2 can be obtaining in this way. Hint: You may use that every curve on a hyperbolic surface can be tightened to a unique geodesic. You may wish to consider minimal length curves between different cuffs.

The Riemann surface can be varied by changing the lengths of the cuffs, and changing the way they are glued together.

1.5. **Three manifolds.** Hyperbolic geometry plays a preeminent role in the topology of three manifolds; it turns out a great many three manifolds have hyperbolic structures. One may ask then if there is an analogue of Teichmüller theory for hyperbolic three manifolds, that studies the moduli space of hyperbolic structures on a given three manifold. Surprisingly, a closed three manifold may have at most one hyperbolic structure; this is the celebrated Mostow Rigidity Theorem.

Nonetheless, Teichmüller Theory plays a major role in three manifold theory. For example, the open manifolds given by a surface cross $\mathbb{R}$ admit a large space of hyperbolic structures parameterized by a product of Teichmüller spaces, and many closed three manifolds can be shown to have a hyperbolic metric using Teichmüller Theory.

2. Teichmüller space and moduli space

A good source for much of the material in this chapter is [**FM12**, Chapters 10, 12].

2.1. **Teichmüller space.** Fix an oriented closed surface S of genus g; this surface is only a topological surface (it isn't a Riemann surface), and it is sometimes called the reference surface. Teichmüller space $\mathcal{T}_g$ is defined to be the space pairs (X, ϕ), where X is a Riemann surface and $\phi : S \to X$ is an orientation preserving homeomorphism, up to to the following equivalence relation: $(X_1, \phi_1) \sim (X_2, \phi_2)$ if there is a biholomorphism $I : X_1 \to X_2$ such that ϕ_2 is homotopic to $I \circ \phi_1$. (At the moment $\mathcal{T}_g$ is a set, but shortly we will endow it with a topology.)

We may write

$$\pi_1(S) = \langle a_1, b_1, \dots a_g, b_g : \prod [a_i, b_i] = 1 \rangle.$$

A point in Teichmüller space can be viewed as a genus g Riemann surface on which one can refer to the homotopy class of any curve, for example a_1, or $a_1 b_2 a_3^{-1}$.

Exercise 2.1. Show that two homeomorphisms between closed surfaces are homotopic if and only if they induce the same map on π_1. (Because there is no fixed basepoint on X, the maps on π_1 should be considered up to conjugacy.)

Note: Soon we will see that Teichmüller space is connected. If in the definition ϕ were not required to be orientation preserving, then Teichmüller space would have two connected components.

2.2. **Teichmüller space of the torus.** Every genus 1 Riemann surface has universal cover $\mathbb{C}$. Indeed, the universal cover must be $\mathbb{C}$ or $\mathbb{H}$, and if it were $\mathbb{H}$ then the torus would have a metric of curvature -1, contradicting Gauss-Bonnet. (Alternative proof: $PSL(2, \mathbb{R})$ does not contain a discrete subgroup isomorphic to $\mathbb{Z}^2$.) The biholomorphisms of $\mathbb{C}$ are the affine maps, and the only affine maps without fixed points are translations. Therefore we get that the Deck group acts via translations, and the Riemann surface is of the form $\mathbb{C}/\Lambda$, where $\Lambda \subset \mathbb{C}$ is a discrete group acting via translations.

Typically for a point in Teichmüller space, we only have a conjugacy class of isomorphisms from fundamental group of the reference surface to the fundamental group of the Riemann surface. However, in genus 1 the fundamental group is abelian, so we get a well defined isomorphism.

Proposition 2.2. *The map*

$$\mathbb{H} \to \mathcal{T}_1, \qquad \tau \mapsto \mathbb{C}/\langle 1, \tau \rangle,$$

with the data of the marking given by the isomorphism

$$\pi_1(S) = \mathbb{Z}^2 \to \langle 1, \tau \rangle, \qquad (1,0) \mapsto 1, \quad (0,1) \mapsto \tau,$$

is a bijection.

Proof. We will construct the inverse. Any genus 1 Riemann surface may be written as $\mathbb{C}/\Lambda$. Hence any point in $\mathcal{T}_1$ may be written as $\mathbb{C}/\Lambda$ where we have a fixed identification of $\mathbb{Z}^2$ with π_1 induced by the marking. It will be important that the marking is orientation preserving. Assume the standard generators of $\mathbb{Z}^2$ are chosen so they have symplectic pairing (algebraic intersection number) 1. Let the images of the two generators be ν, τ respectively, so $\Lambda = \langle \nu, \tau \rangle$. By scaling and rotating, we may assume $\nu = 1$. Since ν and μ are linearly independent in $\mathbb{R}^2$, we must have that $\mathrm{Re}(\tau)$ is positive or negative. Using the fact that the intersection number of ν and μ is 1 and not -1, one can show that $\mathrm{Re}(\tau) > 0$. □

We can define the topology on $\mathcal{T}_1$ as the topology of $\mathbb{H}$.

In summary, we have described $\mathcal{T}_1$ by giving the representation from π_1 of the torus to the group of biholomorphisms of $\mathbb{C}$, modulo biholomorphisms of $\mathbb{C}$ (above this took the form of scaling or rotating both generators of the lattice). We will now do something similar for $\mathcal{T}_g, g > 1$.

2.3. **The character variety.** By uniformizing, a point of $\mathcal{T}_g, g > 1$ gives a map $\pi_1(S) \to PSL(2, \mathbb{R})$.

Exercise 2.3. Prove this map is well defined up to conjugation by an element of $PSL(2, \mathbb{R})$.

The map $\pi_1(S) \to PSL(2, \mathbb{R})$ is discrete (its image has no accumulation points) and faithful (injective). If we quotient $\mathbb{H}$ by the image of any such discrete and faithful map from $\pi_1(S)$ we get a Riemann surface with fundamental group group $\pi_1(S)$; by the classification of surfaces it must be a closed surface of genus g.

We begin by considering the set $DF(\pi_1(S), PSL(2, \mathbb{R}))$ of discrete faithful representations. We then take the quotient by conjugation to obtain

$$DF(\pi_1(S), PSL(2, \mathbb{R}))/PSL(2, \mathbb{R}),$$

the space of conjugacy classes of discrete faithful representations.

If $\pi_1(S) = \langle a_1, b_1, \ldots a_g, b_g : \prod [a_i, b_i] = 1 \rangle$, we may consider the representation variety of all homomorphisms from $\pi_1(S)$ to $SL(2, \mathbb{R})$. This can be identified with tuples of matrices $(A_1, B_1, \ldots, A_g, B_g)$ satisfying the polynomial equations that they are all determinant 1 and $\prod [A_i, B_i] = 1$. This has the induced topology as a subset of $SL(2, \mathbb{R})^{2g}$. Quotienting by ± 1 gives the $PSL(2, \mathbb{R})$ representation variety. The quotient of this representation variety by conjugation is called the character variety; it is endowed with the quotient topology, and contains $DF(\pi_1(S), PSL(2, \mathbb{R}))/PSL(2, \mathbb{R})$ as a subset.

In this way we obtain a topology on $\mathcal{T}_g$, which is a subset of

$$DF(\pi_1(S), PSL(2,\mathbb{R}))/PSL(2,\mathbb{R}).$$

In fact, $DF(\pi_1(S), PSL(2,\mathbb{R}))/PSL(2,\mathbb{R})$ has two connected components (corresponding to orientation preserving or reversing), and Teichmüller space is one of these two components.

Let $\gamma \in \pi_1(S)$, and $X \in \mathcal{T}_g$. Let $\rho_X \in DF(\pi_1(S), PSL(2,\mathbb{R}))$ denote a corresponding representation. The hyperbolic length of γ is given by

$$\cosh^{-1}(\mathrm{trace}(\rho_X(\gamma)/2)),$$

and so hence is a continuous function on Teichmüller space.

2.4. **Fenchel-Nielsen coordinates.** Any maximal collection of disjoint non-isotopic essential simple closed surfaces on a surface contains $3g-3$ curves and divides the surface up into pants. There is a length parameter for each curve, and a twist parameter. Hence Teichmüller space is homeomorphic to $\mathbb{R}^{3g-3}_{>0} \times \mathbb{R}^{3g-3} \simeq \mathbb{R}^{6g-6}$. In particular, it is a real manifold of dimension $6g-6$.

To see that the map $\mathbb{R}^{3g-3}_{>0} \times \mathbb{R}^{3g-3} \to \mathcal{T}_g$ is continuous, we may construct a continuous map from $\mathbb{R}^{3g-3}_{>0} \times \mathbb{R}^{3g-3}$ to $\mathrm{Hom}(\pi_1(S), PSL(2,\mathbb{R}))$ whose image gives the Fuchsian group defining (uniformizing) the hyperbolic surface.

Later we will see that $\mathcal{T}_g$ is a complex manifold; at the moment this may seem unexpected, since $PSL(2,\mathbb{R})$ is not a complex manifold (it has real dimension 3).

2.5. **The mapping class group.** The mapping class group $MCG(S)$ of a surface S is the group $\pi_0(\mathrm{Homeo}^+(S))$ of isotopy classes of homeomorphisms. An important example of a mapping class is a Dehn twist.

Proposition 2.4. *Let S be a torus. The action of a mapping class on $H_1(S,\mathbb{Z}) \simeq \mathbb{Z}^2$ gives an isomorphism*

$$\sigma : MCG(S) \to SL(2,\mathbb{Z}).$$

In higher genus it is unknown if the mapping class group is linear.

Proof. Any homeomorphism of S induces an invertible map

$$\phi_* : H_1(S,\mathbb{Z}) \to H_1(S,\mathbb{Z}).$$

If it is orientation preserving, its derivative can be shown to be 1 (instead of the only other possibility -1). Homotopic maps induce the same map on homology, so this gives the map $\sigma([\phi]) = \phi_*$.

The map σ is surjective by considering linear homeomorphisms. It is injective by $K(\pi,1)$ theory (Whitehead's Theorem). $\square$

In higher genus the mapping class group does not lift to $\mathrm{Homeo}^+(S)$.

There is a natural map $MCG(S) \to \mathrm{Out}\,\pi_1$, which is injective by $K(\pi, 1)$ theory. (In particular, this proves the mapping class group is countable.) The Dehn-Nielsen-Baer Theorem states that the map is surjective, if one uses the slightly larger extended mapping class group, which allows orientation reversing homeos.

An important result gives that the mapping class group is virtually torsion free (it has a finite index subgroup without finite order elements).

2.6. **Moduli space.** Moduli space $\mathcal{M}_g$ is the set of Riemann surfaces of genus g up to biholomorphism. There is a map $\mathcal{T}_g \to \mathcal{M}_g$ given by forgetting the marking: $(X, \phi) \mapsto X$. Moduli space caries the quotient topology.

$MCG(S)$ acts on $\mathcal{T}_g$ by precomposing the marking ϕ with the inverse of the mapping class. Two points $(X, \phi_1), (X, \phi_2)$ that map to the same point in $MCG(S)$ differ by the action of the mapping class $\phi_2^{-1} \circ \phi_1$. Thus, $\mathcal{M}_g$ is the quotient of $\mathcal{T}_g$ by $MCG(S)$.

Proposition 2.5. *If S has genus 1, the action of $MCG(S) = SL(2, \mathbb{Z})$ on $\mathcal{T}_1 = \mathbb{H}$ is by the usual action of fractional linear transformations conjugated by $z \mapsto -z$.*

In particular, $\mathcal{M}_1 = \mathbb{H}/SL(2, \mathbb{Z})$, which is called the modular curve.

Proof. We begin with the map

$$\mathbb{Z}^2 \to \langle 1, \tau \rangle, \qquad (1, 0) \mapsto 1, \quad (0, 1) \mapsto \tau,$$

and precompose with the linear map $\mathbb{Z}^2$ given by g^{-1}, for

$$g = \begin{pmatrix} a & b \\ c & d \end{pmatrix} \in SL(2, \mathbb{Z}).$$

The composite maps $(1, 0)$ to $d - c\tau$ and $(0, 1)$ to $-b + a\tau$. So the new point in Teichmüller space obtained is

$$-\frac{a(-\tau) + b}{c(-\tau) + d}.$$

□

In general, to see that $\mathcal{M}_g$ is a nice object, we need the following, which is [**FM12**, Theorem 12.2].

Proposition 2.6. *The action of the mapping class group on Teichmüller space is properly discontinuous.*

A group action on a locally compact space is called properly discontinuous if, for every compact set K, there are only finitely many group elements g such that $gK \cap K \neq \emptyset$.

Exercise 2.7. The quotient map by a properly discontinuous free action is a covering space map. (Free means that there are no fixed points.)

The quotient of a manifold by a properly discontinuous action that may have fixed points is an example of an orbifold; the orbifold we consider can be thought of as a manifold with "cone points", each of which is modeled on the quotient of $\mathbb{R}^n$ by the action of a finite group.

Since the mapping class group is virtually torsion free, we can find a subgroup $\Gamma \subset MCG(S)$ that has no torsion. Since the action on $\mathcal{T}_g$ is proper, the stabilizer of every point must be finite; hence Γ acts freely on $\mathcal{T}_g$ and the quotient $\mathcal{T}_g/\Gamma$ is a manifold. This manifold maps to moduli space: $\mathcal{T}_g/\Gamma \mapsto \mathcal{T}_g/MCG(S) = \mathcal{M}_g$. Thus $\mathcal{M}_g$ is the best sort of orbifold, the kind finitely covered by a manifold.

Proof of Proposition. Here we consider $\mathcal{T}_g$ as a subset of the character variety, and $MCG(S)$ as a subset of $Out(\pi_1(S))$, acting by precomposition on the character variety.

We will use the following fact, which is intuitive but will not be justified until later in the course: Let K be a compact subset of $\mathcal{T}_g$. Then for any $C > 0$ there are only finitely many isotopy classes of curves that have hyperbolic length at most C somewhere on K.

We will also use the following fact, which won't be justified at all in the course: There is a finite set $\mathcal{S}$ of isotopy classes of closed curves such that only the trivial mapping class fixes all of them.

The action of the mapping class group does not change the hyperbolic metric; it merely remarks the surface. Pick a point $(X, \phi) \in K$, and pick C so that every element of $\mathcal{S}$ has length at most C everywhere on K. Any $g \in MCG(S)$ with the property that gK intersects K must map each element of $\mathcal{S}$ to one of the other finitely many other isotopy classes that have length at most C somewhere on K. If mapping classes g_1, g_2 have the same action on $\mathcal{S}$, then $g_1 g_2^{-1}$ must fix each element of $\mathcal{S}$ and hence be the trivial mapping class, so $g_1 = g_2$. Since there are only finitely many maps from $\mathcal{S}$ to the finite set of isotopy classes that have length at most C somewhere on K, this proves that there can be only finitely many g such that gK intersects K, as desired. □

2.7. **Length functions on $\mathcal{T}_g$.** We have already commented that length functions are continuous on $\mathcal{T}_g$. Next, we remark that a point in Teichmüller space is determined by the length spectrum $\pi_1(S) \to \mathbb{R}_{>0}$.

This admits concrete proofs (and in fact it suffices to consider the lengths of $9g-9$ curves). Since lengths are determined by traces of elements of the Fuchsian group, it is also an instance of the general fact that a semisimple representation of *any* group over a field of characteristic zero is determined by its traces.

In fact, it is a theorem (called marked length spectrum rigidity) that any negatively curved surface is determined by the lengths of its closed geodesic: if there is a homomorphism between two such surfaces such that the geodesics representing corresponding simple closed curves have the same lengths on each surface, then the two surfaces are isometric.

2.8. **Marked points and the universal family over $\mathcal{T}_g$.** One may also consider the Teichmüller space $\mathcal{T}_{g,n}$ of marked Riemann surfaces with n distinguished points; in this case the reference surface S has n distinguished points, and the biholomorphisms are required to send distinguished points to distinguished points. One also calls the distinguished points marked points, but this is perhaps confusing since it is a distinct concept than the marking $\phi : S \to X$.

Puncturing the marked points and uniformizing, we find that the punctured surface has a complete hyperbolic metric; the punctures correspond to cusps of the metric. There are Fenchel-Nielsen coordinates, the only difference being that in addition to regular pants one must use degenerate pants, which are spheres with 2 boundary components and one puncture (cusp) or 1 boundary component and two punctures (cusps). One gets that $\mathcal{T}_{g,n}$ is homeomorphic to $\mathbb{R}^{6g-6+2n}$.

The forgetful map $\mathcal{T}_{g,1} \to \mathcal{T}_g$ is called the universal family (or curve) over $\mathcal{T}$. The fiber over a point (X, ϕ) is a copy of X.

$\mathcal{T}_g$ is what is called a *fine* moduli space. This means in particular that any family of marked Riemann surfaces over a base B can be obtained as the pull back of the universal family via a map $B \to \mathcal{T}_g$.

Exercise 2.8. Prove this when $B = S^1$.

The mapping class group $MCG(S)$ acts on the universal curve by remarking, and we get a map $\mathcal{T}_{g,1}/MCG(S) \to \mathcal{T}_g/MCG(S) = \mathcal{M}_g$. One wishes the fibers of this map were surfaces homeomorphic to S, however this isn't entirely true. Indeed, for a point of $\mathcal{T}_g$ that has a non-trivial stabilizer in $MCG(S)$, the stabilizer will act on the fiber of the universal curve, and the quotient of the fiber by this action will be an orbifold surface of lower genus.

Moreover, one can show that $\mathcal{M}_g$ is not a fine moduli space. For example, let X be a Riemann surface with an automorphism $f : X \to X$. Consider the bundle $X \times [0, 1]$, and then glue together the two ends $X \times \{0\}$ and $X \times \{1\}$ via f. This is a non-trivial bundle over S^1. Compare this to the trivial bundle $S^1 \times X$. Both give the same map $S^1 \to \mathcal{M}_g$, but they are different bundles!

In general, moduli spaces of objects that may have automorphisms are not fine moduli spaces. The reason $\mathcal{T}_g$ is a fine moduli space is that a marked Riemann surface may not have an automorphism *preserving the marking.*

3. The Deligne-Mumford compactification of moduli space

3.1. **The collar lemma.**

The following is [**FM12**, Lemma 13.5]; a proof can be found there.

Lemma 3.1. *Let γ be a simple closed geodesic on a hyperbolic surface. Then the set of points of distance at most*

$$\sinh^{-1}\left(\frac{1}{\sinh(\ell(\gamma)/2)}\right)$$

from γ is an embedded annulus, where $\ell(\gamma)$ is the length of γ.

When $\ell = \ell(\gamma)$ is small, the above expression is comparable to $-\log(\ell)$. In particular, the size of the collar goes to infinity as the length of the geodesic goes to zero.

Corollary 3.2. *There exists a constant $\delta > 0$ such that on any hyperbolic surface, any two distinct closed simple geodesics of length at most δ are disjoint.*

The corollary also applies for non-simple geodesics. The constant is called the Margulis constant.

The collar lemma is a special improvement over the following general theorems.

- Let X be a manifold of bounded non-positive curvature. Then there exists constants $C, \varepsilon > 0$ such that for any discrete subgroup Γ of the isometries of X, and any $x \in X$, the group generated by the set
$$\{\gamma \in \Gamma : d(x, \gamma x) < \varepsilon\}$$
contains a subgroup of index at most C that is nilpotent.
- If G is a semisimple Lie group then there is a neighbourhood U of the identity (called a Zassenhaus neighbourhood) and a $C > 0$ such that any discrete subgroup generated by elements of U contains a nilpotent subgroup of index at most C.

3.2. **Mumford's compactness criterion and the Bers constant.** Our goal is to give a generalization of Mahler's compactness criterion [**FM12**, Theorem 12.7], which says that the subset of $\mathcal{M}_1$ represented by unit area flat tori without a closed geodesic of length less than any $\varepsilon > 0$ is compact. In other words, the only way to degenerate a unit area flat torus is to create a short curve. The corresponding statement for $\mathcal{M}_g, g > 1$ is

Theorem 3.3. *The subset $\mathcal{M}_g^{\geq\varepsilon}$ of $\mathcal{M}_g$ consisting of surfaces without any closed geodesics of length less than ε is compact.*

The area of a ball of radius r in $\mathbb{H}$ is exponential in r. An area argument using Gauss-Bonnet gives that every genus g hyperbolic surface has a closed geodesic of length $O(\log(g))$. In fact, it is know that this $O(\log(g))$ is optimal. In particular, as the genus gets big, there are "fat" surfaces where all closed geodesics are long.

To establish the theorem we need the following more refined estimate, which is proved in [**FM12**, Prop 12.8].

Proposition 3.4. *For each $g > 1$, there is a constant L_g such that any $X \in \mathcal{M}_g$ has a pants decomposition where the length of each cuff is at most L_g.*

The optimal L_g is called the Bers constant. It is known that $L_g \leq 21(g-1)$, and conjectured that it is $O(\sqrt{g})$.

The proposition, together with the fact that there are only finitely many pants decompositions up to the action of the mapping class group, implies the theorem.

3.3. **Deligne-Mumford.** Isometries of $\mathbb{H}$ can be classified as follows. Consider $g \in PSL(2,\mathbb{R})$, not the identity.

- If $|\operatorname{trace}(g)| < 2$, then g has a unique fixed point in $\mathbb{H}$ and is called elliptic. Elliptic elements are conjugate to elements of $SO(2)$.
- If $|\operatorname{trace}(g)| = 2$, then g is called parabolic and is conjugate to $z \mapsto z + 1$.
- If $|\operatorname{trace}(g)| > 2$, then g is called hyperbolic and is conjugate to $z \mapsto \lambda z$, in which case the trace is $\lambda^{\frac{1}{2}} + \lambda^{-\frac{1}{2}}$.

If $\Gamma \subset PSL(2,\mathbb{R})$ uniformizes a compact surface $\mathbb{H}/\Gamma$, then it consists entirely of hyperbolics. The hyperbolic $z \mapsto \lambda z$ corresponds to a geodesic of length $\log \lambda$. (The axis, which in this case is the imaginary line, projects to the geodesic.) If the length of the geodesic goes to 0, then $\lambda \to 1$ and the trace goes to 2.

The loop around a cusp gives parabolic. Hence, we see intuitively that as a length goes to zero, the corresponding element of the Fuchsian group becomes parabolic, and the collar neighborhood converges to a cusp. For a more precise discussion, see [**Wol10**, Section 1.6].

Let $\overline{\mathcal{T}}_g$ be the augmented Teichüller space; it is obtained, for every possible pants decomposition, by formally adding in points where any subset of the cuffs lengths are zero and the corresponding twist parameter becomes undefined. Augmented Teichüller space is not compact; it is a bordification (partial compactification) of Teichmüller space. The points of $\overline{\mathcal{T}}_g$ may be considered to be nodal Riemann surfaces, marked by a map which is allowed to collapse curves to nodes.

We define the Deligne-Mumford compactification $\overline{\mathcal{M}}_g$ to be the quotient of augmented Teichmüller space by the mapping class group. Its points consist of unmarked nodal Riemann surfaces (the nodes are the cusps). It is compact, essentially by finiteness of the Bers constant and the fact that there are only finitely many pants decompositions up to the action of the mapping class group.

Combining pairs of length and twist coordinates in polar coordinates shows $\overline{\mathcal{M}}_g$ is a smooth orbifold. It is naturally stratified; the codimension one boundary strata parameterize surface with one pinched curve; there are $1 + \lfloor \frac{g}{2} \rfloor$ of them.

$\mathcal{M}_g$ has only one end, meaning that for any compact set, there is exactly one unbounded component of the complement. We haven't put any metric on $\mathcal{M}_g$ yet, but you should think of it is being finite volume, but cross sections of the cusps have diameter going to infinity.

4. Jacobians, Hodge structures, and the period mapping

The source for this is the excellent first chapter of [**CMSP03**].

4.1. The case of genus 1.

An Abelian differential on a Riemann surface is a differential one form with coefficients in the complex numbers that can locally be written as $f(z)dz$, where z is a holomorphic local coordinate and f is holomorphic.

Exercise 4.1. Show any Abelian differential is closed.

When we gave the isomorphism $\mathcal{T}_1 \simeq \mathbb{H}$, we implicitly used the Abelian differentials (holomorphic 1-forms) $\omega = dz$ on the torus $X = \mathbb{C}/\langle 1, \tau \rangle$. The parameter τ is the period (integral) of dz. More specifically, if α and β are a fixed symplectic basis of homology (coming from the marking surface), we have

$$\tau = \frac{\int_\beta \omega}{\int_\alpha \omega}.$$

The one forms ω and $\overline{\omega}$ are directly seen to be closed, linearly independent one forms, and hence span $H^1(X, \mathbb{C})$. It is also easy to see that any other holomorphic one form ω' on the torus is a multiple of ω; otherwise ω'/ω would be a non-constant holomorphic function.

Hence we that the space of holomorphic one forms $H^{1,0}(X)$ is a line in the two dimensional space $H^1(X, \mathbb{C})$. Using the marking, we may fix an isomorphism $H^1(X, \mathbb{C}) = \mathbb{C}^2 = \mathrm{span}(\alpha^*, \beta^*)$. Here (α^*, β^*) is the dual basis to (α, β). Normalizing ω as above we have $\omega = \alpha^* + \tau\beta^*$. Thus we see that τ measures the slope of the line $H^{1,0}(X)$ in $\mathbb{C}^2$.

The map from a marked torus to τ is called the period mapping. We have already proven it to be injective, but it is good to pause and appreciate that the Hodge theoretic information of the position of $H^{1,0}(X)$ inside $H^1(X, \mathbb{C})$ determines the Riemann surface.

4.2. **Higher genus and Siegel upper half space.** For any closed Riemann surface, either Hodge theory or Riemann-Roch gives that $H^1(X, \mathbb{C}) = H^{1,0}(X) \oplus H^{0,1}(X)$.

Exercise 4.2. Compute an explicit basis of Abelian differentials for the Riemann surface given by $y^2 = \prod_{i=1}^{2g-2}(z - z_i)$.

To understand how the Hodge structure varies we wish to understand how the g dimensional space $H^{1,0}(X)$ varies inside the $2g$ dimensional space $H^1(X, \mathbb{C})$. Before we do this, we should understand in more detail why the period mapping in genus 1 has image $\mathbb{H}$ instead of $\mathbb{P}^1$.

To start, recall that for any symplectic vector space with a symplectic basis $\alpha_1, \beta_1, \ldots, \alpha_g, \beta_g$ the symplectic pairing on the dual vector space is given by

$$\langle v, w\rangle = \sum_{i=1}^{g} v(\alpha_i)w(\beta_j) - v(\beta_i)w(\alpha_j).$$

You can check this on the standard basis of the dual, for example $\langle \alpha_i^*, \beta_j^*\rangle = \delta_{ij}$.

Lemma 4.3. *The dual of the standard symplectic form on $H_1(X, \mathbb{R})$ is given by*

$$\langle v, w\rangle = \int v \wedge w.$$

For any abelian differential ω, we have that $\frac{i}{2}\omega \wedge \overline{\omega}$ is a positive multiple of $z = dxdy$, in local coordinates $x + iy$. So, for the torus, we get

$$\frac{i}{2}\left(\int_{\alpha_1} \omega \int_{\beta_1} \overline{\omega} - \int_{\beta_1} \omega \int_{\alpha_1} \overline{\omega}\right) > 0.$$

If we normalize so $\int_{\alpha_1} \omega = 1$ (the inequality implies $\int_{\alpha_1} \omega \neq 0$) we get $\int_{\beta_1} \omega_1 \in \mathbb{H}$.

The above inequality in particular implies there is no Abelian differential all of whose α_i periods are zero. Hence it is possible to pick a basis $\omega_1, \ldots, \omega_g$ of $H^{1,0}(X)$ so that $\int_{\alpha_i} \omega_j = \delta_{ij}$. The matrix

$$\tau = \left(\int_{\beta_i} \omega_j\right)$$

is called the period matrix of X. It depends only on the choice of symplectic basis α_i, β_i for H_1.

Note that the space $H^{1,0}(X)$ is the graph of the period matrix mapping the span of the α_i^* to the span of the β_i^*; given a one form with specified α_i periods, the period matrix gives the β_j periods.

The inequality above gives that imaginary part of the period matrix is positive definite. Using $\int \omega_i \wedge \omega_j = 0$ gives that it is symmetric.

4.3. **Siegel upper half space.** If we pick a different symplectic basis for H_1, how does the period matrix change? The new symplectic basis $(\beta_1', \ldots, \beta_g', \alpha_1', \ldots, \alpha_g')$ may be given from the old basis $(\beta_1, \ldots, \beta_g, \alpha_1, \ldots, \alpha_g)$ via a symplectic matrix $g \in Sp(2g)$, and we may write

$$g = \begin{pmatrix} A & B \\ C & D \end{pmatrix}.$$

We wish to compute the matrix τ' that computes the β_i' periods from the α_i' periods, in terms of the matrix τ that computes the β_i periods from the α_i periods. We assume $\int_{\alpha_i} \omega_j = \delta_{ij}$. Then the α_i' periods are $C\tau + D$ and the β' periods are $A\tau + B$. Hence

$$\tau' = (A\tau + B)(C\tau + D)^{-1}.$$

The space $\mathcal{H}_g$ of g by g symmetric matrices with positive definite imaginary part is a homogeneous space for the action of $Sp(2g, \mathbb{R})$.

4.4. **Jacobians and the Torelli theorem.** We obtain a map $\mathcal{M}_g \to \mathcal{H}_g / Sp(2g, \mathbb{Z})$. To understand this map, we define the Jacobian of a Riemann surface X to be the complex torus

$$\mathrm{Jac}(X) = H^{1,0}(X)^* / H_1(X, \mathbb{Z})$$

plus the data of the symplectic form on $H_1(X, \mathbb{Z}) = H_1(\mathrm{Jac}(X), \mathbb{Z})$. Note that $H^{1,0}(X)^* \simeq H_1(X, \mathbb{R})$ as real vector spaces, but $H^{1,0}(X)^*$ carries a complex structure, whereas $H_1(X, \mathbb{R})$ does not. Without this complex structure, the Jacobian would depend only on the homeomorphism type of the Riemann surface.

If the period matrix of a Riemann surface is τ, then the Jacobian is $\mathbb{C}^g/(\mathbb{Z}^g + \tau\mathbb{Z}^g)$. The mapping

$$\tau \mapsto \mathbb{C}^g/(\mathbb{Z}^g + \tau\mathbb{Z}^g)$$

shows that the period mapping covers the map $X \mapsto \mathrm{Jac}(X)$.

The Torelli Theorem shows that $X \mapsto \mathrm{Jac}(X)$ is injective. The map is branched over the hyperelliptic locus. It extends to the set of full geometric genus Riemann surfaces, and it is proper on this locus.

The complex dimension of Siegel upper half space is $1+2+\cdots+g = (g+1)g/2$. Hence the real dimension is $6 = 6g - 6$ for $g = 2$ and $12 = 6g - 6$ for $g = 3$. However starting in genus 4, the dimension of Siegel upper half space is larger than that of $\mathcal{M}_g$. Determining the image is the notoriously difficult Schottky problem.

5. Quasiconformal maps

Given a map f from a subset of $\mathbb{C}$ to $\mathbb{C}$, how may we measure its failure to be holomorphic? The obvious answer is $f_{\overline{z}} = (f_x + if_y)/2$, because this quantity is zero if and only if f is holomorphic. However, given that multiplication by a scalar is holomorphic, we may not want to say that $2f$ is more or less holomorphic than f. Consider instead $\mu = f_{\overline{z}}/f_z$, which is called the complex dilatation of f.

To determine what geometric information μ encodes, consider the derivative of f at a point p, which is a linear map from $\mathbb{R}^2 \to \mathbb{R}^2$. (We write $\mathbb{R}^2$ instead of $\mathbb{C}$ to emphasize that, if f is not holomorphic, this map is only real linear.) It takes the circle $e^{i\theta}$ in the tangent space to p to the ellipse

$$f_z(p)e^{i\theta} + f_{\overline{z}}(p)e^{-i\theta} = f_z(p)(e^{i\theta} + \mu(p)e^{-i\theta}).$$

This achieves its maximum absolute value of $|f_z(p)|(1 + |\mu(p)|)$ at $\theta = \arg(\mu(p))/2$. The minimum value is $|f_z(p)|(1 - |\mu(p)|)$. Hence, the direction of greatest stretch is $\arg(\mu(p))/2$, and the ratio of major to minor axis of the image ellipse is

$$K_f(p) = \frac{1 + |\mu(p)|}{1 - |\mu(p)|}.$$

This is called the dilatation of f at p. The dilatation of f is defined to be

$$\sup_p K_f(p) = \frac{1 + \|\mu\|_\infty}{1 - \|\mu\|_\infty}.$$

A map f is called K-quasiconformal if $K_f < \infty$ and it is in the Sobolev space $W^{1,2}$. (In this generality, the supremum in the definition of $K_f(p)$ should be replaced with an essential supremum.)

It requires a fair bit of work, but it can be shown that the composition of a K_1-quasiconformal map and a K_2-quasiconformal is K_1K_2-quasiconformal, and that a 1 quasiconformal map is conformal. We will require only quasiconformal maps that are homeomorphisms and are differentiable outside of a discrete set of points. In this case the first fact is easy: an infinitesimal ellipse first gets mapped to an ellipse contained between two circles of radii K_1R and R. Each of these circles is mapped to an ellipse contained in between circles of radii K_1K_2RR' and K_1RR' (for the first) and K_2RR' and RR' (for the second). Hence the original circle gets mapped to an ellipse contained between circles of radii K_1K_2RR' and RR'.

One can also see that inverse of a K-qc map is K-qc.

The fact that a continuous 1-qc map that is differentiable outside a discrete set of points is holomorphic follows directly from the removable singularity theorem, since it must be holomorphic wherever it is differentiable.

One of the miracles of quasiconformal maps, which we will not use, is that the set of K-qc maps is closed under uniform limits. This is reminiscent of the fact that the set of holomorphic functions are closed under uniform limits.

The function $p \mapsto \mu(p)$ is called the Beltrami differential of f. We will now see how the Beltrami differential changes when we make a holomorphic change of coordinate in the domain. That is, we will suppose $z = z(w)$ is a holomorphic function of w, and compute the Beltrami differential in terms of the coordinate w. Write $dz/dw = e^{i\phi}r$

Recall the direction of greatest stretch is $\arg(\mu(p))/2$. Hence, in the w coordinate, the new direction of greatest stretch is

$$\arg(\mu(p))/2 - \phi = (\arg(\mu) - 2\phi)/2.$$

The new ratio of major to minor axes is the same, so, noting that

$$\overline{\frac{dz}{dw}} \Big/ \frac{dz}{dw} = e^{-2i\phi},$$

we see that the Beltrami differential in the new coordinate w is given by

$$\mu(z(w)) \, \overline{\frac{dz}{dw}} \Big/ \frac{dz}{dw} \, .$$

That is, Beltrami differentials are naturally $(-1,1)$ forms, i.e. tensors of type $\frac{d\bar{z}}{dz}$.

Postcomposing f with a holomorphic function does not change the Beltrami differential.

Theorem 5.1 (Measurable Riemann Mapping Theorem)**.** *If $\|\mu\|_\infty < 1$ there exists a unique quasiconformal homeo $f^\mu : \mathbb{P}^1 \to \mathbb{P}^1$ that fixes $0, 1, \infty$ whose complex dilatation is given by μ almost everywhere.*

Moreover, f^μ is smooth wherever μ is, and f^μ varies complex analytically with respect to μ.

For a proof, see [**Hub06**, Theorem 4.6.1]. The existence of f^μ when μ is analytic is not too hard and was shown by Gauss. The general theorem can be viewed as the 1D case of the Newlander-Nirenberg theorem.

A glimpse of the (co)tangent space to Teichmüller space. Kodaira-Spencer deformation theory predicts that the infinitesimal deformations of a complex manifold be given by the first Cech cohomology group with coefficients in the sheaf Θ of holomorphic vector fields. The idea is to specify on each overlap of charts a holomorphic vector field indicating how the two charts should be moved with respect to each other; infinitesimally, the new gluing map should be the old going map composed with an infinitesimal amount of the flow.

By Serre duality,

$$H^1(X, \Theta)^* = H^0(X, K - \Theta) = H^0(X, 2K) = Q(X),$$

where $Q(X)$ is the space of holomorphic quadratic differentials on X. This suggests that $Q(X)$ is the cotangent space to Teichmüller space. Later we will verify this, and that the tangent space is equal to the Beltrami differentials modulo the Beltrami differentials that pair trivially with all quadratic differentials.

6. The Teichmüller metric

The reference for this section is [**FM12**, Chapter 11].

The Teichmüller distance between two points of Teichmuller space $(X, \phi), (Y, \psi)$ is defined to be

$$d(X, Y) = \frac{1}{2} \log K,$$

where K is the infimum of the dilatations of all quasiconformal maps from X to Y which are differentiable outside a finite number of points and commute up to homotopy with the marking maps. Later we will show the infimum is uniquely realized and $d(X,Y)=0$ iff $X=Y$. We begin with a special case.

Whether or not to include the factor of $\frac{1}{2}$ is a matter of personal preference. Not including it makes the Teichmüller metric equal to the Kobayashi metric, and including it makes the formula for infinitesimal Teichmüller (co-)metric particularly nice.

Theorem 6.1 (Grötzsche's problem)**.** *If $f:[0,a]\times[0,1]\to[0,Ka]\times[0,1]$ is smooth outside a finite number of points and takes horizontal (resp. vertical) edges to horizontal (resp. vertical) edges, then*

$$K_f \geq K$$

with equality iff f is affine.

Proof. Integrating the inequality

$$Ka \leq \int_0^a |f_x(x,y)|dx$$

over $y \in [0,1]$ and squaring gives

$$K^2a^2 \leq \left(\int |f_x(x,y)|dA\right)^2.$$

Now, let m and M be the min and max of $|df(v)|$ at a fixed point (x,y), where v runs over unit tangent vectors. Then $K_f = M/m$ and $\mathrm{Jac}(f) = mM$ by definition, where Jac denotes the determinant of the derivative. Hence we get $|f_x(x,y)|^2 \leq K_f \,\mathrm{Jac}(f)$. Using Cauchy-Schwarz,

$$\begin{aligned} K^2a^2 &\leq \left(\int \sqrt{K_f(x,y)}\sqrt{\mathrm{Jac}(f)(x,y)}dA\right)^2 \\ &\leq \left(\int K_f(x,y)dA\right)\left(\int \mathrm{Jac}(f)(x,y)dA\right) \\ &\leq (K_f a)(Ka). \end{aligned}$$

To see uniqueness, note that for the first inequality to be equality, f must take horizontal lines to horizontal lines. For the second inequality to be equality, the horizontal direction must be the direction of maximal stretch everywhere. For the Cauchy-Schwarz to be an inequality, $\mathrm{Jac}(f)/K_f = m^2$ must be constant. From the final inequality, K_f must be constant, so M must be constant. Since the direction of max and min stretch are orthogonal, we get that f has constant derivative, and hence must be affine. $\square$

We now wish to generalize this, to claim that for two points in Teichmüller space there is a unique map that stretches and contracts a constant amount everywhere that achieves the infimum in the definition of the Teichmüller distance.

For the torus, it is clear what the candidate maps should be: since the mapping class group is $SL(2,\mathbb{Z})$, each mapping class is represented by a linear map. However, higher genus surfaces do not admit foliations, so it isn't even possible to continuously define the direction of maximal stretch.

For this purpose, consider a quadratic differential q, which is a section of the square of the holomorphic cotangent bundle of the Riemann surface. More concretely, in a local coordinate z it can be written as $q = f(z)(dz)^2$ for f holomorphic, and in a different local coordinate $w = w(z)$ it can be written as

$$q = f(z(w))z'(w)^2(dw)^2.$$

Every quadratic differential can be written in local coordinates as either $(dz)^2$ or $z^k dz$. For example, if $q = f(w)(dw)^2$, and $f(0) \neq 0$, the correct local coordinate near 0 is $z = \int_0^z \sqrt{f(w)}dw$. Then $dz = \sqrt{f(w)}dw$ by the fundamental theorem of calculus, and so $(dz)^2 = q$. Compare to [**Wri15**, Section 1] (which covers the closely related case of Abelian differentials) for more details.

A Teichmüller mapping is defined as a map that stretches the horizontal direction of a quadratic differential by a constant factor $\sqrt{K}$ and contracts the vertical by a constant factor $1/\sqrt{K}$.

Theorem 6.2 (Teichmüller's Uniqueness Theorem)**.** *Let $h : X \to Y$ be a Teichmüller map. If $f : X \to Y$ is quasiconformal and homotopic to h, then*

$$K_f \geq K_h$$

and equality holds if and only if $f = h$ (in genus 1, $f = g$ up to translations).

The proof will be identical to Grötzsche's problem, once we do some work to establish an analogue of the first inequality.

Lemma 6.3. *Let h be a homeomorphism of a compact geodesic metric space. Then there exists a constant $M \geq 0$ such that for any geodesic arc α, the length of $h(\alpha)$ is at least the length of α minus M.*

Proof. One may take M to be twice the maximum distance travelled by any point under an isotopy from α to the identity. $\square$

Lemma 6.4. *Let $h : X \to Y$ be a Teichmüller mapping with quadratic differential q and dilatation K. Let f be any homeo that is homotopic to h and is smooth outside a finite number of points. Then the average horizontal stretch of f is at least $\sqrt{K}$:*

$$\int |f_x| dA \geq \sqrt{K}\,\mathrm{Area}(q).$$

Note that if the dilatation is K, the horizontal stretch factor is $\sqrt{K}$.

Proof. Define $\delta : X \times \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ by

$$\delta(p, L) = \int_{-L}^{L} |f_x| dx,$$

which is the length $f(\alpha_{p,L})$, where $\alpha_{p,L}$ is the horizontal arc of length $2L$ centered at p. Here f_x is a vector (the directional derivative of f) and $|f_x|$ is its norm.

The Teichmüller map takes $\alpha_{p,L}$ to an arc of length $2L\sqrt{K}$. The previous lemma thus gives that $\delta_{p,L} \geq 2L\sqrt{K} - M$. Hence

$$\int \delta_{p,L} dA \geq (2L\sqrt{K} - M)\,\mathrm{Area}(q).$$

On the other hand,

$$\int \delta_{p,L} dA = 2L \int |f_x| dA.$$

Dividing by $2L$ and taking $L \to \infty$ gives the result. $\square$

The proof of Teichmüller's Uniqueness Theorem is now identical to Grötzsche's problem.

Theorem 6.5. *For every homeo $f : X \to Y$ between Riemann surfaces, there is exists a Teichmüller mapping $f : X \to Y$ homotopic to f.*

We need a the following corollary of the Measurable Riemann Mapping Theorem.

Corollary 6.6. *If μ is a Beltrami differential on a Riemann surface X, and $\|\mu\|_\infty < 1$, then there is a quasiconformal map from X to a Riemann surface Y with complex dilatation μ. Moreover, Y varies continuously with μ*

The techniques in the proof will be important later in the course.

Proof of Corollary. Lift μ to a Beltrami differential $\tilde{\mu} = b(z)d\overline{z}/dz$ on $\mathbb{H}$. Extend it to a measurable Beltrami differential $\tilde{\mu}_{ext}$ on $\mathbb{C}$ by setting $\tilde{\mu}_{ext} = \overline{b(\overline{z})}d\overline{z}/dz$ if z is on the lower half plane. ($\tilde{\mu}_{ext}$ is undefined on $\mathbb{R}$, which has measure zero.) This is the same thing as reflecting the ellipse field in the $\mathbb{R}$ axis.

Let $f : \mathbb{P}^1 \to \mathbb{P}^1$ be the function produced by the Measurable Riemann Mapping Theorem with complex dilation $\tilde{\mu}_{ext}$.

Claim 1: $f(\mathbb{R}) = \mathbb{R}$. Indeed, $z \mapsto \overline{f(\overline{z})}$ also has complex dilation $\tilde{\mu}_{ext}$ and fixes $0, 1, \infty$, so we must have $\overline{f(\overline{z})} = f(z)$ for all z.

Now, let Γ be the group of Deck transformations, so $X = \mathbb{H}/\Gamma$.

Claim 2: There is representation $\rho : \Gamma \to PSL(2, \mathbb{R})$ such that $f(\gamma z) = \rho(\gamma)f(z)$. Indeed, for $\gamma \in \Gamma$, the function $z \mapsto f(\gamma z)$ also has complex dilation $\tilde{\mu}_{ext}$, so it is equal to $\gamma' \circ f$ for a unique Möbius transformation $\gamma' \in PSL(2, \mathbb{C})$. It is easy to see this gives a representation $\rho : \Gamma \to PSL(2, \mathbb{C})$. The fact that $f(\mathbb{R}) = \mathbb{R}$ implies that the image of ρ is in $PSL(2, \mathbb{R})$.

Claim 3: ρ is discrete and faithful. Indeed, f induces a homeomorphisms between the topological space $\mathbb{H}/\Gamma$ and the topological space $\mathbb{H}/\rho(\Gamma)$.

Now f induces a map from $X = \mathbb{H}/\Gamma$ to $Y = \mathbb{H}/\rho(\Gamma)$. □

Proof of Theorem. Recall that the space $QD(X)$ of quadratic differentials on X has complex dimension $3g - 3$. There is a norm on $QD(X)$ given by

$$\|q\| = \int |q|,$$

which can be thought of as the area of the polygons defining q. Let $QD_1(X)$ denote the open unit ball.

For $q \in QD_1(X)$, set

$$K = \frac{1 + \|q\|}{1 - \|q\|}.$$

We can define a map $\Omega : QD_1(X) \to \mathcal{T}_g$ by constructing the Teichmuller mapping from X with quadratic differential Q and horizontal stretch factor K. We need to show Ω is surjective.

This follows by invariance of domain from two claims: (1) that it is continuous and (2) that it is proper.

The second claim follows from the continuity of the distance from X function on Teichmüller space. (To see this continuity, note that if Γ and Γ' are nearby discrete faithful representations, then they have fundamental domains that are almost equal, and it is possible to construct a quasiconformal map from one to the other with small dilatation.) To show properness, we need to show that the preimage of a compact set is compact. Since distance from X is continuous, the compact set is contained in a closed ball about X, and hence its preimage is contained in a closed ball in $QD_1(X)$, which is compact.

The first claim follows from the continuity statement in the corollary above by noting that the Beltrami differential $\|q\|\frac{\overline{q}}{q}$ of the Teichmüller map varies continuously with k.

To see that the Beltrami differential is $\|q\|\frac{\overline{q}}{q}$, it suffices to assume $q = (dz)^2$. ☐

Corollary 6.7. *The Teichmüller metric is a metric, rather than just a pseudo-metric.*

Proof. For any two distinct points of Teichmüller space, there is a Teichmüller mapping between them. If the horizontal stretch factor is $\sqrt{K}$, then the Teichmüller distance is $\log(K)$. If $K = 1$, the Teichmüller mapping is a biholomorphism. ☐

Corollary 6.8. *There is a unique geodesic through any two points of $\mathcal{T}_g$.*

Proof. The Teichmüller existence theorem gives one. To show it is unique, pick some point Z such that $d(X, Y) = d(X, Z) + d(Z, Y)$. Let f, g be the Teichmüller maps $X \to Z$ and $Z \to Y$ respectively. It must be that $g \circ f$ is a Teichmüller map with dilatation the product of the dilatations of f and g, so the terminal quadratic differential of f must be equal to the initial quadratic differential of g. ☐

Proposition 6.9. *The bijection $\mathbb{H} \to \mathcal{T}_1$ is an isometry from the hyperbolic metric to half the Teichmüller metric.*

Proof. The bundle of quadratic differentials above $\mathcal{T}_1$ is isomorphic to to the space of lattices in $\mathbb{C}$, *not* up to rescaling or rotation. $SL(2, \mathbb{R})$ acts transitively on the unit area locus via its usual linear action on $\mathbb{R}^2$, so we obtain a map from $SL(2, \mathbb{R})$ to the space of lattices given by $g \mapsto g\langle 1, i\rangle$. This map covers the bijection $\mathbb{H} = SL(2, \mathbb{R})/SO(2) \to \mathcal{T}_1$. (Actually this is bit tricky to get right, but it's not worth dwelling on. Technically the map which sends $g^{-1}(i, i)$ to the affine action of g times the square torus doesn't quite cover our old bijection $\mathbb{H}$ to $\mathcal{T}_1$, but rather that bijection composed with $\tau \mapsto -\tau$. This is assuming we want the action by "instructions" discussed below to be a right action, so isometries can be a left action.)

Note, the action of $SL(2, \mathbb{R})$ on marked lattices that we are considering does not correspond to Möbius transformations. One action is on the right and one is on the left. Some times people say one action is by isometries, and the other is via "instructions" (such as geodesic flow, horocycle flow, etc).

Every matrix $g \in SL(2,\mathbb{R})$ can be written as $g = k_1 g_t k_2$, where

$$g_t = \begin{pmatrix} e^{t/2} & 0 \\ 0 & e^{-t/2} \end{pmatrix},$$

and $k_1, k_2 \in SO(2)$.

$SL(2,\mathbb{R})$ is isomorphic to the unit tangent bundle of $\mathbb{H}$. The group $SO(2)$ acts via rotation unit tangent vectors, and g_t acts via geodesic flow by time t. Hence the hyperbolic distance between p and $k_1 g_t k_2$ is t. The linear action on lattices is a Teichmüller mapping with dilatation e^t, and so the Teichmüller distance is $t/2$. □

Remark 6.10. In general, the metric topology is the same as the previous topologies. We have already sketched why the Teichmüller distance is continuous for the previous topologies. Now it suffices to show that any neighbourhood of a point for the previous topologies contains a Teichmüller ball. This follows from the fact that Ω is a homeomorphism.

We conclude this chapter with a result that can be used to complete our proof that the action of MCG on $\mathcal{T}_g$ is properly discontinuous (Proposition 2.6).

Lemma 6.11 (Wolpert's Lemma). *If X_1, X_2 are hyperbolic surfaces and $\phi : X_1 \to X_2$ is K quasiconformal, then*

$$\frac{\ell_{X_1}(\gamma)}{K} \leq \ell_{X_2}(\gamma) \leq K\ell_{X_1}(\gamma).$$

Proof. By passing to covers, it suffices to assume X_1 and X_2 are annuli, and γ generates the fundamental group. We can write X_i as the strip $\mathbb{R} \times (0,\pi) \subset \mathbb{C}$ modulo $z \mapsto z + m_i$, and ϕ lifts to an equivariant map between these strips. The proof of Grötzsche gives $m_1/K \leq m_2 \leq Km_1$. The proof is completed by noting that m_i is the hyperbolic length. (The map $z \to \exp(z)$ maps the strip to the upper half plane, intertwining the action of $z \mapsto z + m$ with the action of $z \mapsto e^m z$.) □

7. Extremal length

The material in this section will not be used in other sections but is beautiful, important, and fairly easy to digest.

Suppose that Γ is a collection of rectifiable curves on a Riemann surface. The extremal length of Γ is defined as

$$EL(\Gamma) = \sup_\rho \frac{L_\rho(\Gamma)^2}{A_\rho},$$

where ρ ranges over all metrics in the given conformal class, A_ρ is the area of the metric, and $L_\rho(\Gamma)$ is the minimal ρ length of a curve in Γ.

Note that the fraction in the definition doesn't change if ρ is scaled, and that despite the name extremal length is better thought of as a squared length.

Lemma 7.1. *Consider the rectangle $R = (0, w) \times (0, h)$, and let Γ denote the set of all curves that go from the left edge to the right edge. Then*

$$EL(\Gamma) = w/h.$$

Proof. Taking ρ to be the standard Euclidean metric we have $L_\rho(\Gamma) = w$ and $A_\rho = wl$, so we conclude $EL(\Gamma) \geq w/h$. (It is a general feature of the definition that obtaining lower bounds for extremal length is often easy.)

Now consider the standard metric scaled by ρ, where $\rho : R \to [0, \infty]$ is an arbitrary Borel measurable function. Let $\ell = L_\rho(\Gamma)$. Note for any $0 \leq y \leq h$ we have

$$\ell \leq \int_0^w \rho(t + iy)dt,$$

hence

$$\ell h \leq \int_0^h \int_0^w \rho(t + iy)dtdy.$$

Now Cauchy-Schwarz with functions ρ and 1 gives

$$\left(\int_0^h \int_0^w \rho(t + iy)dtdy\right)^2 \leq wh \int_0^h \int_0^w \rho(t + iy)^2 dtdy = whA_\rho.$$

Hence $(\ell h)^2 \leq whA_\rho$, so $\ell^2/A_\rho \leq w/h$ as desired. $\square$

Similarly, one gets the following. See wikipedia for details.

Lemma 7.2. *Consider an annulus of modulus m. If Γ is the set of all curves from one side of the annulus to the other, we get that the extremal length across the annulus is m. If instead Γ is instead the set of all closed curves that go around the annulus once, we get that the extremal length around the annulus is $1/m$.*

Annuli are especially important since they suffice to give the hyperbolic length of any curve on any Riemann surface; associated to that curve there is a cover of the Riemann surface, and the hyperbolic length of the curve is one over the modulus of that annulus.

Here is an example of how one can use extremal length to give information on moduli. Consider a cylinder of circumference one and width m. Its modulus is then m. Suppose one makes a cut from each end, and now the largest sub cylinder of the cut cylinder has width w.

Exercise 7.3. The modulus m' of the cut cylinder satisfies $w \leq m' \leq w+1$. *Hint:* For one direction, use the fact that holomorphic maps contract. For the other direction, use the metric which is equal to the Euclidean metric within usual distance $\frac{1}{2}$ of the sub-cylinder, and 0 elsewhere.

Next consider a closed Riemann surface. We will refer to the extremal length of a simple closed curve, by which we mean the extremal length of the collection Γ of all such curves in the same isotopy class.

Theorem 7.4 (Existence and uniqueness of Strebel differentials.)**.** *For any closed curve on any Riemann surface there exists a unique holomorphic quadratic differential q that consists entirely of a horizontal cylinder, such that the core curve of this cylinder is isotopic to the given curve. The extremal length of the curve is equal to one over the modulus of the cylinder.*

In his thesis, Kerckhoff proved the following.

Theorem 7.5. *The Teichmüller distance between two Riemann surfaces X, Y is equal to*

$$\sup_{\alpha} \log\left(\frac{\sqrt{EL_X(\alpha)}}{\sqrt{EL_Y(\alpha)}}\right),$$

where α ranges over all simple closed curves, and $EL_X(\alpha)$ denotes the extremal length of α on X.

8. Nielsen-Thurston classification of mapping classes

The source for this section is [**FM12**, Chapter 13].

Genus 1. Every nonidentity element of $SL(2,\mathbb{R})$ is either elliptic, parabolic, or hyperelliptic, according to wether it fixes a point in $\mathbb{H}$, a single point in $\partial\mathbb{H}$, or a pair of points on $\partial\mathbb{H}$. These cases correspond to having absolute value of trace less than 2, equal to 2, or greater than 2, respectively.

The mapping class group in genus 1 is $SL(2,\mathbb{Z})$, and it acts on $\mathcal{T}_1 = \mathbb{H}$ via Möbius transformations. Suppose $A \in SL(2,\mathbb{Z})$.

- If A is elliptic, it fixes a point in $\mathbb{H}$. Since the action is properly discontinuous, it must be finite order.
- If A is parabolic, its eigenvalues are ± 1 and equal, and A fixes a line in $\mathbb{R}^2$ with rational slope. Hence A fixes the image of this line on $\mathbb{R}^2/\mathbb{Z}^2$. The image of the line is a simple closed curve, since the slope was rational.

- If A is hyperbolic, it acts on $\mathbb{R}^2/\mathbb{Z}^2$ via a Teichmüller mapping, which in this case is called an Anosov map.

Periodic mapping classes. We begin by studying finite order (periodic) mapping classes.

Lemma 8.1. *No nontrivial isometry of a closed hyperbolic surface is isotopic to the identity.*

Proof. An isometry that is isotopic to the identity lifts to an isometry of $\mathbb{H}$ that acts as the identity on the boundary. $\square$

Corollary 8.2. *The isometry group of a closed hyperbolic surface is finite.*

In fact, it has size at most $84(g-1)$, but we will not prove or need this.

Proof. It is easy to see that it is compact, and the previous lemma gives that it is discrete. See [**FM12**, Proposition 7.7] for details. $\square$

Lemma 8.3. *The stabilizer of a point $(X, \phi) \in \mathcal{T}_g, g > 1$ is isomorphic to $Isom(X)$.*

Proof. Let ψ be in the stabilizer. Then then there is a biholomorphism from (X, ϕ) to $(X, \phi \circ \psi^{-1})$, which can be viewed as an isometry of X isotopic to $\phi \circ \psi^{-1}\phi^{-1}$. This biholomorphism is unique by a previous lemma. The map sending ψ to this biholomorphism is an isomorphism from the stabilizer to $Isom(X)$. $\square$

Lemma 8.4. *Every finite order element of the mapping class group has a fixed point in $\mathcal{T}_g$.*

In fact, Kerckhoff proved that every finite order subgroup has a fixed point, solving Nielsen's realization problem.

Proof. Suppose f has order n. We induct on the number of prime factors of n, counting multiplicity.

Because finite groups cannot have finite dimensional $K(\pi, 1)$'s, it follows that $f^{n/p}$ has a fixed point for some prime factor p of n. If $n = p$, this establishes the base case. Otherwise, the fix point set of $f^{n/p}$ is isomorphic to the Teichmüller space of the quotient orbifold. Considering the action of f on this smaller Teichmüller space, which must have order at most n/p, produces a fixed point by induction. $\square$

The classification. A mapping class is called reducible if it fixes some simple multi-curve. It is called pseudo-Anosov if it is represented by a Teichmüller map with the same initial and terminal Riemann surface and quadratic differentials.

Theorem 8.5. *Every mapping class is periodic, reducible, or pseudo-Anosov.*

A pseudo-Anosov mapping class cannot be reducible. In some sense most mapping classes are pseudo-Anosov.

Pseudo-Anosovs are dynamically complicated: for any two curves α, β the intersection number $i(f^n(\alpha), \beta)$ grows exponentially. (Moreover, $f^n(\alpha)$ starts to look like the horizontal foliation of the quadratic differential.)

Every mapping class has a finite power which fixes some subsurfaces of the surface, and is either the identity or a pseudo-Anosov on each subsurface, and also may perform some Dehn twists on the boundaries of the subsurfaces.

Proof. Every mapping class f acts as an isometry on Teichmüller space. Let

$$\tau(f) = \inf_{X \in \mathcal{T}_g} d(X, f(X))$$

denote the translation distance of f.

If $\tau = 0$ and is realized, we have already proven f is periodic. We will show that if τ is not realized than f is reducible, and if $\tau > 0$ is realized then f is pseudo-Anosov.

First suppose τ is not realized, and suppose $(X_i, \phi_i) \in \mathcal{T}_g$ be such that $d((X_i, \phi_i), f(X_i, \phi_i)) \to \tau$.

Claim 1: $X_i \to \infty$ in $\mathcal{M}_g$.

Otherwise, we may find some h_i in the mapping class group so that $(X_i, \phi_i \circ h_i^{-1})$ converges to some point (X, ϕ). It is easy to see that $h_i f h_i^{-1}$ moves $(X_i, \phi_i \circ h_i^{-1})$ by only slightly more than τ, and hence also for (X, ϕ). (The distance moved minus τ goes to zero with i.) Since the action is properly discontinuous, this implies $h_i f h_i^{-1}$ is eventually constant (after passing to a subsequence) and hence that the distance from (X, ϕ) to $h_i f h_i^{-1}(X, \phi)$ is eventually equal to τ. Hence the distance between $h_i^{-1}(X, \phi)$ and $f h_i^{-1}(X, \phi)$ is eventually equal to τ, so τ is actually realized.

Claim 2: f is reducible.

By Mumford's compactness criterion, the length of the shortest hyperbolic geodesic on X_i goes to zero. We may fix $K > 0$ so that there is a K-quasi-conformal map from (X_i, ϕ_i) to $f(X_i, \phi_i)$ for all i. (Indeed, for i large enough the Teichmuller distance from (X_i, ϕ_i) to $f(X_i, \phi_i)$ is at most $\tau + 1$, so we can take $K = \exp(\tau + 1)$.) Let $\delta > 0$ be such that no two hyperbolic curves of length less than δ can cross. Eventually X_i has at least one curve α_i of length less than δ / K^{3g-3}. The orbit

$f^j(\alpha_i), k = 0, \ldots 3g-3$ consists of curves of length at most δ, hence two of them must be equal.

We now move on to the case where $\tau > 0$ is realized, say $d(X, f(X)) = \tau$. Let γ be the Teichmüller geodesic through X and $f(X)$.

Claim 3: $f(\gamma) = \gamma$.

First consider Y on the segment of γ from X to $f(X)$. Since $d(X, Y) + d(Y, f(X)) = d(X, f(X))$ and since $d(f(X), f(Y)) = d(X, Y)$, it is easy to see that $d(Y, f(Y)) = d(X, f(X)) = \tau$. Considering the Teichmüller mappings shows $f(Y) \in \gamma$. We conclude that $f(\gamma) = \gamma$.

Claim 4: The initial and terminal quadratic differentials for f on X are equal.

This follows from the fact that

$$d(f^2(X), X) = d(f(X), X) + d(f^2(X), f(X)),$$

which follows from the fact that γ is preserved. $\square$

It is true that the axis γ is unique, but this does not follow from the above proof. The above proof is due to Bers; one can see uniqueness from a prior proof due to Thurston, which is sketched in [**FM12**, Chapter 15].

There are many ways to construct examples of pseudo-Anosovs.

(1) Branched covers of a torus with an Anosov diffeo that fixes the branch points. In particular, square-tiled surfaces.
(2) The Thurston-Veech construction.
(3) Penner's construction: If $\{\alpha_i\}$ and $\{\beta_j\}$ are multi-curves that fill the surface, then any product of positive powers of the Dehn twists in the α_i and negative powers of the Dehn twists in the β_j will give a pseudo-Anosov.

9. Teichmüller space is a bounded domain

9.1. The Schwarzian derivative.

We follow the exposition in [**Hub06**, Chapter 6.3]. The following is [**Hub06**, Exercise 2.3.2].

Lemma 9.1. *Let X and Y be manifolds, and let $x \in X, y \in Y$ be points. Let $f, g : X \to Y$ be C^k maps with $f(x) = g(x) = y$. Then if in any local coordinates near x all partial derivatives of order at most $k-1$ vanish, then this is true in any local coordinates. Moreover, the partials of $f - g$ of order k provide a well-defined linear map $Sym^k(T_xX) \to T_yY$.*

Proof of first claim of lemma. It is equivalent to say that, putting Riemannian metrics on X and Y, the distance from $f(x + x')$ to $g(x + x')$

is $o(|x'|^{k-1})$. (Any two Riemannian metrics are locally bilipshitz.) The result follows because change of coordinate maps are bilipshitz.

If we wish only to check that the fact does not depend on coordinates in X, we can compose the $(k-1)$-st multivariable Taylor polynomials for the change of coordinate chart at $f-g$. (If we change coordinates in Y, then $f-g$ is not well defined.) $\square$

Note that T_{x_0} can be identified with first order differential operators. By the mixed partials theorem, these commute with each other. In local coordinates, the map is given by applying any homogeneous degree k polynomial in the partial derivatives to $f-g$.

As a first example, it is constructive to consider $k=1$. Then $Sym^k(T_{x_0}X) = T_{x_0}X$, and the well defined linear map is just the usual derivative.

As a second example, consider $Y=\mathbb{R}$, $k=2$. In this case a map from $Sym^2(T_xX)$ to $\mathbb{R}$ is an element of $Sym^2(T_x^*X)$, i.e. a symmetric bilinear form on T_xX. This is nothing other than the Hessian, and the above lemma recovers the fact, used frequently in Morse Theory, that the Hessian is well defined up to conjugacy (independent of coordinate chart) at a critical point.

Let $X=Y=\mathbb{R}$ and $k=2$, and h is a function with $h(0)=h'(0)=0$. We illustrate the computation showing that $h''(0)$ is an element of $Sym^2(T_0^*\mathbb{R})\otimes T_0\mathbb{R}$. First consider a change of coordinates on X; this corresponds to precomposing h with a function $\phi:\mathbb{R}\to\mathbb{R}$ with $\phi(0)=0$. Then $(h\circ\phi)' = (h'\circ\phi)\cdot\phi'$ and

$$(h\circ\phi)'' = ((h''\circ\phi)\cdot\phi')\cdot\phi' + (h'\circ\phi)\cdot\phi''.$$

So we get $(h\circ\phi)''(0) = h''(0)\phi'(0)^2$.

Next consider a change of coordinates on Y; this corresponds to post-composing h with ψ^{-1}, where $\psi:\mathbb{R}\to\mathbb{R}$ is the new local coordinate for Y. Then $(\psi^{-1}\circ h)' = ((\psi^{-1})'\circ h)\cdot h'$ and

$$(\psi^{-1}\circ h)'' = (((\psi^{-1})''\circ h)\cdot h')\cdot h' + ((\psi^{-1})'\circ h)\cdot h''.$$

So we get $(\psi^{-1}\circ h)''(0) = h''(0)/\psi'(0)$.

As a third example, consider $X=Y=\mathbb{R}$ and $k=3$. In this case $Sym^3(\mathbb{R})=\mathbb{R}$, and a linear map from $Sym^k(\mathbb{R})\to\mathbb{R}$ corresponds to a cubic map $\mathbb{R}$ to $\mathbb{R}$, i.e. a map T such that $T(\lambda v)=\lambda^3T(v)$.

We now wish the define the Schwarzian derivative $S\{f,g\}$ which for holomorphic functions f,g measures how far $g\circ f^{-1}$ is from being a Möbius transformation. For each z, there exists a unique Möbius transformation A such that f and $A\circ g$ have the same value and first and second derivatives. (Later we will compute A explicitely.)

Then the leading term $D^3(f - A \circ g)(z)$ is naturally a cubic map

$$T_zU \to T_{f(z)}\mathbb{P}^1.$$

We may compose this with the inverse of $Df(z)$ to get a cubic map $T_zU \to T_zU$.

We now use the fact that if V is a one dimensional vector space, quadratic maps $V \to \mathbb{C}$ are in bijection to cubic maps $V \to V$, via

$$\alpha \mapsto (w \mapsto \alpha(w)w).$$

(Both spaces of maps are 1 dimensional vector spaces.) Thus

$$Df(z)^{-1} \circ D^3(f - A \circ g)(z)$$

is naturally a quadratic form on T_zU, i.e. a quadratic differential. We define this to be $S\{f, g\}$.

Lemma 9.2. *Let $U \subset \mathbb{C}$. If $f : U \to \mathbb{C}$ is analytic with non-vanishing derivative, then*

$$S\{f, z\} = \left(\frac{f'''}{f'} - \frac{3}{2}\left(\frac{f''}{f'}\right)^2\right)(dz)^2.$$

Proof. Assume $z = 0$ and $f(z) = 0$, and write

$$f(z) = a_1 z + \frac{a_2}{2}z^2 + \frac{a_3}{6}z^3 + \cdots .$$

The Möbius transformation that best approximates f is

$$\frac{\alpha z}{1 + \beta z} = \alpha z - \alpha\beta z^2 + \alpha\beta^2 z^3 + \cdots ,$$

where $\alpha = a_1$ and $\beta = -\frac{a_2}{2a_1}$. The third derivative of $f(z) - A(z)$ is

$$6\left(\frac{a_3}{6} - \frac{a_2^2}{4a_1}\right).$$

Composing with $(Df)^{-1}$, i.e. dividing by a_1, gives

$$\frac{a_3}{a_1} - \frac{3a_2^2}{2a_1^2}.$$

□

Lemma 9.3. *Let $U \subset \mathbb{C}$ be connected and $f : U \to \mathbb{C}$. Denote $S(f) = S\{f, z\}$. Then*

1. *$S(f) = 0$ if and only if f is the restriction of a Möbius transformation to U.*
2. *$S(f \circ g) = g^*(S(f)) + S(g)$.*

The second condition is called the cocycle condition. There are simpler cocycles that you may know better. For example the log derivative $\log(f')$ satisfies

$$\log((f\circ g)') = \log(f'\circ g) + \log(g') = g^*(\log(f')) + \log(g')$$

and kills translations. The nonlinearity $N(f) = (\log(f'))'dz$ satisfies the same cocycle condition and kills complex linear maps.

Proof. That $S(f)=0$ when f is a Möbius transformation follows either from the definition or from direct computation.

We can write

$$S(f)/(dz)^2 = \left(\frac{f''}{f'}\right)' - \frac{1}{2}\left(\frac{f''}{f'}\right)^2.$$

Therefore if $S(f)=0$ then $y=f''/f'$ satisfies the separable ODE $y'=y^2/2$. This can be solved explicitely to show f is a Möbius transformation.

The cocycle condition follows from direct computation. □

Corollary 9.4. *If $S(f)=S(g)$, then $f=A\circ g$ for some Möbius transformation A.*

Proof. From the cocycle condition and the assumption we get

$$S(f\circ g^{-1}) = (g^{-1})^*(S(f)) + S(g^{-1}) = (g^{-1})^*(S(g)) + S(g^{-1}).$$

Applying the cocycle condition to $g\circ g^{-1}=1$ we get

$$0 = (g^{-1})^*(S(g)) + S(g^{-1}),$$

which gives $S(f\circ g^{-1})=0$. □

Lemma 9.5. *For any holomorphic quadratic differential q, there are solutions to $S(f)=q$. The solutions are meromorphic, and unique if the value and first two derivatives are specified at some point.*

Proof. Say $q=q(z)(dz)^2$. If w_1, w_2 are two solutions to the linear equation $w''+\frac{q}{2}w=0$, then if $f=w_1/w_2$ direct computation shows $S(f)=q$. □

Another nice formula, which we will not use, is

$$S(f) = 6\lim_{w\to z}\frac{\partial^2}{\partial w\partial z}\log\frac{f(w)-f(z)}{w-z}.$$

This can be used to obtain bounds for how much f distorts cross ratios, in terms of $S(f)$ [**GL00**, page 127].

9.2. **Complex projective structures.** A complex projective structure is an atlas of charts with transition maps in $PGL(2,\mathbb{C})$. There is a natural forgetful map from complex projective structures to Riemann surfaces, where you remember the holomorphic structure but forget the complex projective structure. Uniformization shows this map is surjective: every Riemann surface has a complex projective structure where the transition maps are actually in $PGL(2,\mathbb{R})$. But there are many other weird complex projective structures.

Given a complex projective structure on X, there is a natural holomorphic map from $\mathbb{H} = \tilde{X}$ to $\mathbb{P}^1$, which is well defined up postcomposition with Möbius transformations. This map, called the developing map, can be obtained by analytically continuing any chart. (Note that given two charts that overlap, the second can be adjusted by a Möbius transformation so they agree on the overlap. It is crucial that $\mathbb{H}$ is simply connected for the developing map to be well defined.)

For any $\gamma \in \pi_1(X)$, and any developing map f, then $f \circ \gamma$ is another developing map. Since the developing map is unique up to Möbius transformations, we have $f \circ \gamma = A_\gamma \circ f$ for some Möbius transformation A_γ. The map

$$\gamma \mapsto A_\gamma$$

is called the holonomy map. It is more concretely understood by analytically continuing a chart along the γ on X and comparing the new chart to the original.

For Fuchsian complex projective structures, i.e. those coming from uniformization, the developing map is a Möbius transformation, and X may be recovered from the holonomy by quotienting. However, for general complex projective structures, the holonomy map can have dense image, and the holonomy map fails to determine the complex projective structure.

Lemma 9.6. *There is a bijection from complex projective structures to representations $\pi_1(X) \to PSL(2,\mathbb{C})$ together with holomorphic immersions $\mathbb{H} \to \mathbb{P}^1$ that are equivariant for the given representation, modulo a simultaneous action by Möbius transformations. The bijection sends each complex projective structure to its holonomy representation and developing map.*

Proof. It suffices to build the inverse. For any small open subset of X, lift it to $\mathbb{H}$, and take the image under f; this will be a chart. We then need to see that different charts agree up to Möbius transformations, but this is basically by definition. $\square$

Proposition 9.7. *The set of marked complex projective structures is an affine bundle over $\mathcal{T}_g$ with fiber $QD(X)$ over X.*

Proof. Given two complex projective structures, one can take the Schwarzian derivative of a chart from one with respect to a chart from the other. ☐

9.3. **From $B(X)$ to $QD(X^*)$.** Let $B(X)$ denote Beltrami differentials on X of norm less than 1, and let $QD(X^*)$ denote quadratic differentials on the complex conjugate Riemann surface. Given $\mu \in B(X)$, lift it to $\tilde{\mu}$ on $\mathbb{H}$, and extend that to a Beltrami differential $\hat{\mu}$ on $\mathbb{P}^1$ by setting it to be zero on the lower half plane. Let f^μ denote the function given by the measurable Riemann mapping theorem with complex dilatation $\hat{\mu}$. As in the proof of Corollary 6.6, since μ is invariant under $\Gamma = \pi_1(X)$, we get that f^μ is equivariant with respect to some representation ρ of Γ. Note that $f(\mathbb{H})$ is invariant under $\rho(\Gamma)$ and has quotient Y_μ given by the Riemann surface in Corollary 6.6. Note also that $f(\mathbb{L})$ is invariant under $\rho(\Gamma)$ and has quotient X^*. (Here $\mathbb{L}$ denotes the lower half plane.) We say that $\rho(\Gamma)$ simultaneously uniformizes Y_μ and X (or X^*).

By taking the Schwarzian derivative of $f|_{\mathbb{L}}$ we get a quadratic differential on X^*.

The map from $B(X)$ to $QD(X^*)$ is a complex analytic map. (For basics on complex analytic maps between Banach spaces, see for example [**Hub06**, Appendix 5].)

9.4. **The Bers embedding.** Let f_μ (note the subscript instead of a superscript) be the map $f : X \to Y_\mu$ given by extending $\tilde{\mu}$ to the lower half plane via reflection.

Lemma 9.8. *$(f_\mu)|_{\mathbb{R}}$ is equal to $(f_\nu)|_{\mathbb{R}}$ up to Möbius transformations if and only if $Y_\nu = Y_\mu$.*

Proof. We have $Y_\mu = \mathbb{H}/\Gamma_\mu$, where $\Gamma_\mu = f_\mu \Gamma f_\mu^{-1}$. Any isometry of $\mathbb{H}$, and hence any Fuchsian group, is determined by its action on the boundary. The action of Γ_μ on the boundary is by definition the conjugate by $(f_\mu)|_{\mathbb{R}}$ of the action of Γ. ☐

The following is [**GL00**, page 133].

Lemma 9.9. *The following are equivalent.*

(1) $(f_\mu)|_{\mathbb{R}} = (f_\nu)|_{\mathbb{R}}$.
(2) $(f^\mu)|_{\mathbb{R}} = (f^\nu)|_{\mathbb{R}}$.
(3) $(f^\mu)|_{\mathbb{L}} = (f^\nu)|_{\mathbb{L}}$.

Proof. **2 implies 1:** Let g_μ be the conformal map from $f^\mu(\mathbb{H})$ to $\mathbb{H}$ normalized to fix $0, 1, \infty$. Then $g_\mu \circ f^\mu$ has the same Beltrami coefficient as f_μ, maps $\mathbb{H}$ to itself, and fixes $0, 1, \infty$, so we must have $f_\mu = g_\mu \circ f^\mu$. If $(f^\mu)|_{\mathbb{R}} = (f^\nu)|_{\mathbb{R}}$ then we get $f^\mu(\mathbb{H}) = f^\nu(\mathbb{H})$, and hence $g_\nu = g_\mu$.

1 implies 2: Define $h(z)$ by

$$h(z) = g_\nu^{-1} \circ g_\mu$$

in $f^\mu(\mathbb{H})$ and

$$h(z) = f^\nu \circ (f^\mu)^{-1}$$

elsewhere. By assumption we have $g_\mu \circ f^\mu = g_\nu \circ f^\nu$ on $\mathbb{R}$, so we get that $h(z)$ is continuous. Now, h is conformal on the complement of $f(\mathbb{R})$, and so hence must be conformal everywhere. (This uses that $f(\mathbb{R})$ is analytically removable for any quasi-conformal map f, a fact which we have not proved but is part of the basic theory of quasiconformal maps.) Since h is a holomorphic bijection of $\mathbb{C}$ fixing $0, 1$ we conclude that it is the identity.

3 implies 2: This is immediate, since the f^μ etc are homeos, and so in particular are continuous.

2 implies 3: This follows from the fact that holomorphic functions are determined by their boundary values (when the boundary values exist). (That is, apply the max modulus principle to $(f^\mu)|_{\mathbb{L}} - (f^\nu)|_{\mathbb{L}}$ and its negation.)

□

Corollary 9.10. *$Y^\nu = Y^\mu$ if and only if $S(f^\mu|_{\mathbb{L}}) = S(f^\nu|_{\mathbb{L}})$.*

Proof. First assume $Y^\nu = Y^\mu$. Then f_μ and f_ν agree on $\mathbb{R}$, so it follows from the previous lemma.

Next assume $S(f^\mu|_{\mathbb{L}}) = S(f^\nu|_{\mathbb{L}})$. Then $f^\mu|_{\mathbb{L}}$ and $f^\nu|_{\mathbb{L}}$ agree up to a Möbius transformation. Since they both fix $0, 1, \infty$ we must actually have $f^\mu|_{\mathbb{L}} = f^\nu|_{\mathbb{L}}$ □

The Bers embedding is the map $\mathcal{T}_g \to QD(X^*)$. There is one Bers embedding for each $X \in \mathcal{T}_g$. We now wish to understand what happens when we change X to X'. In doing so, we will show that $\mathcal{T}_g$ has a complex structure, and that the Bers embedding is a biholomorphism onto its image.

Lemma 9.11. *If $f_\mu = f_\rho \circ f_\lambda$, then*

$$\rho = \left(\frac{\mu - \lambda}{1 - \bar{\lambda}\mu} \frac{(f_\lambda)_z}{\overline{(f_\lambda)_z}} \right) \circ f_\lambda^{-1}.$$

We omit this computation, which can be viewed purely as a computation with linear maps $\mathbb{R}^2 \to \mathbb{R}^2$. We will use only that ρ is a holomorphic function of μ.

Theorem 9.12. *$\mathcal{T}_g$ has a well defined complex structure for which every Bers embedding is a biholomorphism to a domain in $\mathbb{C}^{3g-3}$.*

Proof. First we claim that any $X \in \mathcal{T}_g$ has a neighbourhood U_X with a section $U_X \to B(X)$ that is a one sided inverse to $\mu \mapsto Y_\mu$. This follows from the fact that the derivative of $B(X)$ to $QD(X^*)$ is surjective at $\mu = 0$, a fact which we will prove shortly in the proof of Theorem 9.14. Thus we can take a $3g-3$ dimensional complex subspace of $B(X)$ where the derivative is invertible at $\mu = 0$, and by the inverse function theorem this maps onto a neighbourhood of 0 in $QD(X^*)$. One can also explicitly construct the section (see Theorem 9.16).

Next we claim that the overlap maps U_X to $U_{X'}$ are holomorphic. Indeed, they are the composition of the holomorphic section , the change of basepoint map from $B(X)$ to $B(X')$, and the holomorphic map $B(X') \to QD((X')^*)$.

This shows that there is a well defined complex structure on $\mathcal{T}_g$. To conclude we comment that the Bers embedding is biholomorphic onto its image. Indeed, since it is injective and open, it suffices to show it is holomorphic. Fix the basepoint X. To show the Bers embedding is holomorphic at X', observe that it is the composition of a section $U_{X'}$ to $B(X')$, together with the change of basepoint map $B(X') \to B(X)$ and the map $B(X)$ to $QD(X^*)$. □

9.5. **The tangent space to Teichmüller space.** We now compute the kernel of the derivative of the map from Beltrami differentials on X to $QD(X^*)$. The quotient of $B(X)$ by this kernel is thus the tangent space to $\mathcal{T}_g$ at X.

Lemma 9.13. *Suppose $f(z,t)$ is holomorphic in z and differentiable in t. We think of this as a family of holomorphic functions indexed by time. Suppose $f(z,0) = z$. Then*

$$\frac{\partial}{\partial t}\Big|_{t=0} S(f) = f_t''',$$

where subscript denotes the t partial and prime denotes the z partial.

Proof. Compute

$$\frac{\partial}{\partial t}\left(\frac{f'''}{f'} - \frac{3}{2}\left(\frac{f''}{f'}\right)^2\right) = \frac{f_t''' f' - f''' f_t'}{(f')^2} - \frac{3}{2} \cdot 2 \cdot \frac{f''}{f'} \cdot \frac{\partial}{\partial t}\left(\frac{f''}{f'}\right)$$

and evaluate at $t = 0$ using $f'(0) = 1, f''(0) = 0$, $f'''(0) = 0$. □

Theorem 9.14. *The kernel of the derivative of the Bers embedding is the Beltrami differentials that pair trivially with all holomorphic quadratic differentials.*

Proof. We begin by showing that if a Beltrami differential μ is trivial, then it pairs trivially with $QD(X)$. Write

$$f^{t\mu}(z) = z + tv(z) + o(t),$$

so $\overline{\partial}$ of the vector field $v(z)/dz$ is equal to μ.

Note that since $f^{t\mu}$ fixes $0, 1, \infty$ we have that the vector field $v(z)/dz$ is zero at $0, 1, \infty$. Since $f^{t\mu}$ is holomorphic in the lower half plane, we have $v'''(z) = 0$ in the lower half plane. Hence v is quadratic on the lower half plane, and hence extends continuously to a quadratic function on $\mathbb{R}$.

Since the quadratic vector field $v|_{\mathbb{R}}\partial_x$ has three zeros (at $0, 1, \infty$), it must be zero. (Using a coordinate $u = 1/x$, one gets $du = -dx/x^2$, so $\partial_x = -u^2\partial_u$. The condition that $v|_{\mathbb{R}}\partial_x$ has a zero at infinity is equivalent to $v|_{\mathbb{R}}$ actually being linear rather than quadratic.) Now, consider $\gamma(v/dz)$, where $\gamma \in \Gamma$. This is again a solution to $\overline{\partial}\gamma(v/dz) = \mu$, and again we have $\gamma(v/dz)|_{\mathbb{R}} = 0$. Since any two solutions to the $\overline{\partial}$ equation differ by a holomorphic function, and since a holomorphic function cannot vanish on $\mathbb{R}$, we get that v/dz is Γ invariant, and hence descends to a vector field V on X.

We get that $\mu = \overline{\partial}(V)$ on X. The product rule (applied in local coordinates if desired) gives that

$$\overline{\partial}(qV) = \overline{\partial}(q)V + q\overline{\partial}(V).$$

Since qV is differential form of type dz, we have that $d(qV) = \overline{\partial}(qV)$, so

$$\int \overline{\partial}(qV) = 0.$$

We get

$$\int \mu q = \int (\overline{\partial}V)q = -\int V(\overline{\partial}q) = 0,$$

since q is holomorphic.

This shows that the space of trivial Beltrami differentials is contained in $QD(X)^{\perp}$. Now, $QD(X)^{\perp}$ has real codimension $6g-6$. (The pairing is non-degenerate, since $\overline{q}/|q|$ pairs non-trivially with q.) The space of trivial differentials cannot be smaller than $QD(X)^{\perp}$, since the quotient of all Beltrami differentials by the trivial ones has real dimension at most $6g-6$. □

We also now give a direct argument that if a Beltrami differential μ pairs trivially with $QD(X)$, then it is trivial.

The first observation is that $\tilde{\mu}$ pairs trivially with any finite area quadratic differential on $\mathbb{H}$. Indeed, averaging a finite area quadratic differential over the Deck group Γ produces a "locally finite area" quadratic differential on $\mathbb{H}$, which descends to a quadratic differential on $X = \mathbb{H}/\Gamma$.

Now, the first derivative of $S(f^{\mu t})$ is given by $(f^{\mu t})_t'''$. By mixed partials, we can take the t derivative first and then the three z derivatives. We may write

$$f^{t\mu}(z) = z + tv(z) + O(t^2),$$

where $v(z)\frac{\partial}{\partial z}$ is a vector field with $\overline{\partial} v = \mu$. This equation is often called the "infinitesimal Beltrami equation." Its solutions are unique only up to the addition of a holomorphic vector field. However, since all holomorphic vector fields on $\mathbb{P}^1$ are quadratic, this ambiguity will turn out not to matter.

The solution we pick is

$$v(z) = \frac{1}{\pi} \int \frac{\mu(w)}{(z-w)} |dw|^2.$$

(See [**Hub06**, Proposition A6.4.1] for a reminder that $1/(\pi z)$ is a fundamental solution to the $\overline{\partial}$ equation. We will only be interested in $z \in \mathbb{L}$, keeping in mind that μ is supported on $\mathbb{H}$.) Technically we should probably check that v is summable over Γ; this is because $\mu = \mu(z) d\overline{z}/dz$, where $\mu(z)y^2$ is bounded, and $\int \frac{(dw)^2}{y^2(z-w)}$ has finite area.

Taking the t derivative gives

$$\frac{1}{\pi} \int \frac{\mu(w)}{(z-w)} |dw|^2$$

and then taking the three z derivatives gives

$$\frac{6}{\pi} \int_{\mathbb{H}} \frac{\mu(w)}{(w-z)^4} |dw|^2.$$

This is the pairing of μ with a finite area quadratic differential. (Note $(dz)^2/(w-z)^4$ is not finite area on $\mathbb{C}$, but it is on $\mathbb{H}$, where μ is supported. This is because we only take the Schwarzian at $z \in \mathbb{L}$.) So we get that this derivative is 0. We have accomplished our claim.

9.6. **Supping up the Bers embedding.** The following shows the image of the Bers embedding is contained in the ball of radius $\frac{3}{2}$. For a proof see [**Hub06**, Theorem 6.3.9].

Theorem 9.15 (Nehari)**.** *If $f : \mathbb{H} \to \mathbb{C}$ is injective, then*

$$\|S(f)\|_\infty < \frac{3}{2}.$$

Like the Koebe quarter theorem, this follows easily from the area theorem (see the wikipedia entry for the Koebe quarter theorem). One can also use the area theorem to get a growth bound on injective functions from the disc to $\mathbb{C}$ with $f(0) = 0, f'(0) = 1$; Montel's Theorem then gives compactness of the space of such functions, which gives the weaker result that $\|S(f)\|_\infty$ is universally bounded by some constant.

The following shows the image contains a ball of radius $\frac{1}{2}$, which is [**Hub06**, Theorem 6.3.10].

Theorem 9.16 (Ahlfors-Weill)**.** *Let $q = q(z)(dz)^2$ be a quadratic differential on $\mathbb{H}$ with $\|q\|_\infty < \frac{1}{2}$. Set $\mu = 2y^2 q(\overline{z})\frac{d\overline{z}}{dz}$ if $z \in \mathbb{H}$ and $\mu = 0$ on $\mathbb{L}$. Then $S(f^\mu) = q$. Moreover if q is equivariant with respect to some Fuchsian group, then so is μ.*

9.7. **The Teichmüller metric is Finsler.** The following discussion comes from [**McM**, page 54]. For a more rigorous treatment, see Section 6.6 in [**Hub06**].

If $K(f)$ denotes the real dilatation, then

$$\frac{1}{2}\log(K(f_{t\mu})) = \frac{1}{2}\log\frac{1 + t\|\mu\|_\infty}{1 - t\|\mu\|_\infty} = t\|\mu\|_\infty + o(t^2).$$

Minimizing over all μ in the same equivalence class gives that the Teichmüller metric is the quotient metric. The Teichmüller cometric is thus L^1 norm on quadratic differentials, given by

$$\|q\| = \int |q| = Area(q).$$

Indeed, the dual of the space V of measurable quadratic differentials with the L^1 norm is the space of measurable Beltrami differentials with the L^∞ norm. If W is the space of holomorphic quadratic differentials, then the tangent space to Teichmuller space is $V^*/W^\perp$. A general fact in functional analysis says that $V^*/W^\perp$ is isometric to W^*, which gives the result.

9.8. **Isometries of $\mathcal{T}_g$.** Consider the map $X \to \mathbb{P}QD(X)$ given by mapping each point x of X to the line of quadratic differentials on X that vanish to the highest possible order at x. (This order depends on x.) One can show that the this map is analytic and injective, and that the metric on the unit ball of $QD(X)$ is less differential at points of the image than any other points.

This leads to Royden's Theorem, which says that the group biholomorphic isometries of $\mathcal{T}_g$ is the mapping class group. (Later we will see that a biholomorphism must automatically be an isometry.) See

[**Hub06**, Section 7.4] for details. The proof is slightly different when $g = 2$.

10. The Weil-Petersson Kähler structure

In this section we define the symplectic form ω on moduli space and calculate it in Fenchel-Nielsen coordinates, following [**Wol10**]. It is important for a number of reasons, including because ω/π^2 generates H^2 of the mapping class group (we will not prove this), and because it relates the hyperbolic geometry of individual Riemann surfaces to the geometry of moduli space.

10.1. **The symplectic form.** Recall that the space of quadratic differentials is the cotangent space to moduli space. The Weil-Petersson cometric on the space of quadratic differentials is the hermitian product

$$\langle q_1, q_2 \rangle = \int q_1 \overline{q}_2 (ds^2)^{-1},$$

where ds^2 is the hyperbolic area form. The dual metric on the tangent space to moduli space is a hermitian h metric called the WP hermitian product. The associated Riemannian metric $\operatorname{Re} h$ is called the WP metric, and the associated 2-form $\omega = -\operatorname{Im} h$ is the WP symplectic form. The 2-form associated to any hermitian metric is always non-degenerate, but we will need to prove ω is closed in order to know that it is symplectic.

If q is a quadratic differential, then $\overline{q}(ds^2)^{-1}$ is a Beltrami differential. So the map $q \mapsto \overline{q}(ds^2)^{-1}$ gives a map from the cotangent space to the tangent space of moduli space. As always, there is a pairing between tangent and cotangent vectors; recall that for moduli space the pairing of a Beltrami differential μ and a quadratic differential q is $\int \mu q$.

The WP hermitian product of two quadratic differentials q_1 and q_2 is equal to the pairing of q_1 with $\overline{q}_2(ds^2)^{-1}$. For any finite dimensional vector space with a hermitian product, there is an anti-linear map to the dual (or predual) given by $v \mapsto \langle \cdot, v \rangle$; we see that this map from quadratic differentials to the tangent space is given by $q \mapsto \overline{q}(ds^2)^{-1}$. By definition, for any finite dimensional vector space this map sends the hermitian product to the complex conjugate of the dual hermitian product. Hence we get that the WP hermitian product of two Beltrami differentials of the form $\mu_i = \overline{q}_i(ds^2)^{-1}$ is given by

$$\langle \mu_1, \mu_2 \rangle = \overline{\int \overline{\mu}_1 \mu_2 ds^2}.$$

Remark 10.1. Beltrami differentials of the form $\overline{q}(ds^2)^{-1}$ are called harmonic Beltrami differentials. The above shows that every tangent vector to moduli space can be represented by a harmonic beltrami differential. (The map $q \mapsto \overline{q}(ds^2)^{-1}$ must be injective because $\overline{q}(ds^2)^{-1}$ has nontrivial pairing with the quadratic differential q.)

Remark 10.2. It is a fact that for any two Beltrami differentials μ_1, μ_2, there are equivalent Beltrami differentials μ_1', μ_2' with disjoint support. This shows that the above formula is *not* true when the Beltrami differentials are not harmonic.

10.2. **Gardiner's formula.** Since we wish to arrive at a formula for the symplectic form in terms of differentials of geodesic lengths, we will need to understand the differential of the hyperbolic length of a curve as the Riemann surface changes in the direction of some Beltrami differential μ. We will use a dot over a variable to denote derivative with respect to the time variable t. In this subsection we prove the following result, which is [**Wol10**, Cor 2.6].

Theorem 10.3 (Gardiner's formula)**.** *Let X_t be a family of Riemann surfaces with quasi-conformal maps $f_t : X_0 \to X_t$ with Beltrami differentials μ_t, and assume $\dot{\mu}_0 = \mu$.*

Let ℓ_t denote the hyperbolic length of a fixed closed curve, and let $\mathcal{F}$ denote the cover of X_0 corresponding to the subgroup of the fundamental group generated by this curve. This cover is conformal to a unique cylinder of height π; let q denote the quadratic differential giving the Euclidean metric on this cylinder. Then

$$\dot{\ell}_0 = \frac{2}{\pi} \operatorname{Re} \int_{\mathcal{F}} \mu q.$$

Here we use μ to denote both the Beltrami differential on X_0 and its lift to $\mathcal{F}$.

Although stated in more complicated way, Gardiner's formula merely reflects the fact that the modulus of a cylinder under an infinitesimal quasi-conformal map changes according to how efficiently the beltrami differential pairs with the natural quadratic differential on the cylinder.

Gardiner's formula becomes more explicit if we assume that $X_0 = \mathbb{H}/\Gamma$ and that the closed geodesic in question is the image of the imaginary axis in $\mathbb{H}$. In this case $\mathcal{F}$ can be replaced with the fundamental domain

$$\mathcal{D} = \{z \in \mathbb{H} : 1 < |z| < e^{\ell}\}.$$

Noting that $\log(\mathcal{D}) = \{x + iy : 0 < x < \ell, 0 < y < \pi\}$, we get that q is the pull back of $(dz)^2$ via $z \mapsto \log z$, so $q = (dz/z)^2$. Gardiner's formula becomes

$$\dot{\ell}_0 = \frac{2}{\pi} \operatorname{Re} \int_{\mathcal{D}} \mu \left(\frac{dz}{z} \right)^2 .$$

Gardiner's formula is the statement that the differential $d\ell$ of geodesic length, which is a cotangent vector to moduli space, is equal to the quadratic differential $\frac{2}{\pi}\left(\frac{dz}{z}\right)^2$.

Proof of Gardiner's formula. The maps f_t induce maps h_t of the strip $\{z \in \mathbb{H} : \operatorname{Im}(z) < \pi\}$. These maps satisfy $h_t(z+\ell_0) = h(z) + \ell_t$. For any $0 < y < \pi$ and any t, we get

$$\ell_t = \operatorname{Re}\left(h_t(\ell_0 + iy) - h_t(iy)\right) = \operatorname{Re} \int_0^{\ell_0} \frac{d}{dx} h_t(x+iy) dx.$$

The maps f_t preserve the boundary of the strip, so it follows that

$$0 = \operatorname{Im} \frac{d}{dt}\left(h_t(x+i\pi) - h_t(x)\right) = \operatorname{Im} \int_0^{\pi} \frac{d^2}{dtdy} h_t(x+iy) dy.$$

Now, taking an equality that is true for any x and averaging it, we get

$$\begin{aligned} \frac{d}{dt}\ell_t &= \operatorname{Re} \int_0^{\ell_0} \frac{d^2}{dtdx} h_t(x+iy)dx \\ &= \frac{1}{\pi} \operatorname{Re} \int_0^{\pi} \int_0^{\ell_0} \frac{d^2}{dtdx} h_t(x+iy) dxdy. \end{aligned}$$

Adding zero (times i/π) to this we get

$$\begin{aligned} \frac{d}{dt}\ell_t &= \frac{1}{\pi} \operatorname{Re} \int_0^{\pi} \int_0^{\ell_0} \frac{d}{dt} \left(\frac{d}{dx} h_t(x+iy) + i\frac{d}{dy} h_t(x+iy) \right) dxdy. \\ &= \frac{2}{\pi} \operatorname{Re} \int_0^{\pi} \int_0^{\ell_0} \frac{d}{dt} \left(\frac{\partial}{\partial \bar{z}} h_t \right) dxdy. \end{aligned}$$

Now we use $\frac{\partial}{\partial \bar{z}} f_t(z) = \mu_t \frac{\partial}{\partial z} f_t(z)$, where μ_t is the Beltrami differential of f_t, to get

$$\frac{d}{dt}\left(\frac{\partial}{\partial \bar{z}} h_t \right)\Big|_{t=0} = \dot{\mu}_0.$$

(This can also be seen using a first order expansion $f_t(z) = z + tv + o(t)$, and computing that $\dot{\mu}_0 = \frac{\partial}{\partial \bar{z}} v$. First take the t derivative, to get v. Then take the $\frac{\partial}{\partial \bar{z}}$ derivative.) Since $\dot{\mu}_0 = \mu$ by definition, this gives

$$\frac{d}{dt}\ell_0 = \frac{2}{\pi} \operatorname{Re} \int_{\mathcal{D}} \mu.$$

$\square$

10.3. **Twist-length duality.** For a simple closed geodesic α on a Riemann surface of finite type, let ℓ_α denote its hyperbolic length, and let t_α be the vector field on Teichmüller space that generates the Fenchel-Nielsen twist about α; flowing for time t in this vector field increases the α twist parameter in Fenchel-Nielson coordinates by t. In this subsection we prove the following result, which is [**Wol10**, Theorem 3.3].

Theorem 10.4 (Twist-length duality)**.** *If J is the complex structure (multiplication by i) of Teichmüller space, and ω is the WP symplectic form, then*

$$2t_\alpha = J \operatorname{grad} \ell_\alpha.$$

Equivalently,

$$2\omega(\cdot, t_\alpha) = d\ell_\alpha.$$

Twist-length duality shows that the Hamiltonian flow of geodesic length is the twist.

The moral strategy of the proof is to first show these statements for the Teichmüller space of the annulus. However, the Teichmüller space of the annulus is one real dimensional, and hence its tangent space admits neither a complex structure or a symplectic form. Hence we take some care with the statements.

As above, instead of working with the annulus $\mathcal{F} = \mathbb{H}/\langle z \to e^\ell z\rangle$ we may apply log and work in the strip $\{z \in \mathbb{H} : \operatorname{Im}(z) < \pi\}$. In the annulus, there is a unique closed geodesic, which we call α. Using coordinates $z = x + iy$, the hyperbolic metric in the strip is

$$\frac{(dx)^2 + (dy)^2}{\sin(y)^2}.$$

Lemma 10.5. *Let $\phi : [0, \pi] \to [0, 1]$ be any C^1 monotone function that is 0 at 0 and 1 and π, with a decay condition on ϕ' at 0 and π. Consider the Beltrami differential*

$$\mu_\phi = \frac{i}{2}\phi'(y)$$

in the annulus. Let $\tilde{\mu}_\phi$ be the sum over cosets of $\langle\alpha\rangle$ in $\pi_1(X)$ of the lift of μ_ϕ to $\mathbb{H}$, so $\tilde{\mu}_\phi$ is $\pi_1(X)$ invariant. Then the corresponding Beltrami on X represents the infinitesimal twist in α.

Proof. We prove this in two steps.

Step 1: We first show this is true when $\phi'(y)$ is supported very close to $\pi/2$, so μ_ϕ is supported very close to the geodesic α.

Consider the family of maps $f_\varepsilon(z) = z + \varepsilon\phi(\operatorname{Im} z)$, which preserves the y coordinate. The Beltrami of this map is

$$\frac{\varepsilon\overline{\partial}\phi(\operatorname{Im} z)}{1+\varepsilon\partial\phi(\operatorname{Im} z)} = \frac{\varepsilon i/2\phi'(y)}{1-\varepsilon i/2\phi'(y)}.$$

So the corresponding infinitesimal Beltrami for this family of maps at $\varepsilon = 0$ is μ_ϕ.

Consider the effect of f_ε on a geodesic orthogonal to α, i.e. a vertical line. The image is homotopic rel endpoints to a vertical path to α, followed by a segment of α of Euclidian length ε, followed by a vertical path to the other side of the strip. Along α, the Euclidian metric is equal to the hyperbolic metric, so this segment of α also has hyperbolic length ε.

Suppose that $\phi'(y)$ is supported close enough to $\pi/2$ that μ_ϕ is supported on a collar neighbourhood of α. Using that disjoint collars are disjoint, it is easy to sum over cosets, and see that the corresponding Beltrami differential on X is supported on a collar, the Beltrami corresponding to ε changes the twist by exactly ε. Taking ε to 0 gives the result.

Step 2: Now, we show different choices of ϕ give equivalent Beltramis on X.

The decay condition is so that the preimage can be summed over cosets and the sum converges absolutely. (This is a technical point that can be avoided for our purposes, but we discuss it anyways.) In the strip, the line $y = \varepsilon$ is about $-\log(\varepsilon)$ hyperbolic distance away from α. At most e^r geodesics in the orbit of α intersect any ball of radius r in $\mathbb{H}$. So, we want $1/\varepsilon$ times $\phi'(\varepsilon)$ to go to 0; so the decay condition is just that $\phi' = 0$ at the endpoints.

To start, given ϕ_1, ϕ_2, we will construct a vector field $v(y)/dz$ whose $\overline{\partial}$ derivative is $\mu_{\phi_1} - \mu_{\phi_2}$. This is the solution to

$$\frac{i}{2}v' = \phi_1' - \phi_2',$$

so this is trivial to do, and v has the same decay condition.

Next, we can lift and sum over cosets to show that the Beltramis on X are $\overline{\partial}$ of a vector field on X. As in the proof of Theorem 9.14, this shows that these two Beltramis on X are equivalent. □

Lemma 10.6. *There is a choice of ϕ such that $2\tilde{\mu}_\phi$ represents $i \operatorname{grad} \ell_\alpha$.*

Proof. Suppose that μ_ℓ is a Beltrami differential realizing $i \operatorname{grad} \ell_\alpha$. The definition $(\operatorname{grad} f, \cdot) = df(\cdot)$ says that μ_ℓ is a Beltrami differential such that taking inner product with μ_ℓ has the same effect as pairing with $d\ell$.

Gardiner's formula says that $d\ell$ is integration against the pull back of the real part of $\frac{2}{\pi}$ times the flat metric $(dz)^2$ on the strip. The inner product is

$$(\mu_1, \mu_2) = \mathrm{Re} \int \mu_1 \overline{\mu}_2 ds^2 = \frac{1}{2} \mathrm{Re} \int \overline{\mu}_1 \mu_2 ds^2.$$

Hence we see that $\operatorname{grad} \ell$ is represented by the sum over cosets of the pull back of $\frac{2}{\pi}(d\overline{z})^2/ds^2$ to $\mathbb{H}$.

On the strip,

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

and hence

$$\frac{2}{\pi}(d\overline{z})^2/ds^2 = \frac{2}{\pi}\sin(y)^2 \overline{dz} dz.$$

So, comparing to $\mu_\phi = \frac{i}{2}\phi'(y)$ we wish to choose

$$\phi(y) = \frac{2}{\pi} \int_0^y \sin(y)^2.$$

This ϕ has that $\phi(\pi) = 1$, and gives the result. (It seems like a miracle that $\phi(\pi) = 1$, but if this isn't the case we'd still have nice formulas, just scaled by a constant.) □

The theorem follows immediately from the previous two lemmas.

10.4. **The symplectic form is closed.** This follows either by direct but unsatisfying calculation, or by constructing a one form on Teichmüller space whose exterior derivative is ω. The later uses "quasifuchsian reciprocity", a symmetry statement about the Bers embedding based at different points.

10.5. **The symplectic form in Fenchel-Nielsen coordinates.** In this subsection we prove the following result, which is [**Wol10**, Theorem 3.14].

Theorem 10.7. *In Fenchel-Nielsen coordinates the WP symplectic form is*

$$2\omega = \sum d\ell_i \wedge d\tau_i.$$

Proof. The first claim is that

$$2\omega \left(\frac{\partial}{\partial \ell_i}, \frac{\partial}{\partial \tau_j} \right) = \delta_{i,j}.$$

This follows from twist-length duality, which gives that the expression is $d\ell_j(\ell_i)$.

Each pair of pants has an orientation reversing isometry that sends each cuff to itself; it is given by exchanging the two hexagons. Taking a pants decomposition and adjusting the twist parameters gives a surface with an orientation reversing isometry ρ that maps each pants to itself. Note that since the symplectic form is invariant under twists, this does not change the coefficients of ω in Fenchel-Nielsen coordinates. (This uses that ω is closed, and that twists are Hamiltonian flows.)

The mapping class of ρ acts on Teichmüller space anti-holomorphically. Since it acts via remarking the surfaces, it preserves the WP metric g. Since it reverses the complex structure, this gives that $\rho^*(\omega) = -\omega$, since $\omega(u, v) = g(u, iv)$.

Now, $d\ell_i$ is invariant under ρ, and $d\tau$ is odd with respect to ρ. Hence the coefficients of $d\tau_i d\tau_j$ and $d\ell_i d\ell_j$ are zero. $\square$

11. Kobayashi hyperbolicity

Introduction. The Kobayashi (pseudo)metric on a complex manifold M is the largest (pseudo)-metric such that holomorphic maps from $\mathbb{H}$ to M are distance non-increasing. The key property of the Kobayashi metric is that holomorphic maps between complex manifolds M and N are always distance non-increasing with respect to the Kobayashi (pseudo)metrics on these M and N.

Example 11.1. The Kobayashi metric on $\mathbb{C}$ is 0. For any space that contains an holomorphically embedded copy of $\mathbb{C}$, the Kobayashi metric is not an honest metric (there are distinct points distance zero from each other).

Brody's Theorem asserts the converse. Spaces on which the Kobayashi metric is a metric are called Kobayashi hyperbolic.

Example 11.2. The Schwarz-Lemma gives that the Kobayashi metric on $\mathbb{H}$ is the hyperbolic metric.

Example 11.3. On a product, $d((x_1, y_1), (x_2, y_2)) \geq \max(d(x_1, x_2), d(y_1, y_2))$. For the product of discs, this is an equality.

Example 11.4. On the open unit ball in $\mathbb{C}^n$, the Kobayashi metric is the complex hyperbolic metric.

The Kobayashi metric is always Finsler (this is not obvious, but you can take it to be the definition for many purposes), with unit ball

$$\left\{ \frac{1}{2}\gamma'(0) \right\}$$

where γ runs over all holomorphic map from the disc into the manifold sending 0 to the given point.

Lemma 11.5. *Let B be the unit ball of some Banach space. Then the unit ball for the infinitesimal Kobayashi metric on T_0B is $\frac{1}{2}B$.*

Proof. It is clear that the unit ball contains $\frac{1}{2}B$, by looking at the maps $\lambda \mapsto \lambda e$ for any e of norm 1. Conversely, for any e of norm 1, then by Hahn-Banach there is a linear functional of norm one sending e to 1. Thus $\lambda e \mapsto \lambda$ under this linear functional. By the distance non-increasing property, the image of the unit ball under this map must be contained in the unit ball for the hyperbolic plane. □

Holomorphic motions. The reference for this subsection is [**Hub06**, Section 5.2], which contains complete proofs.

A holomorphic motion of a subset $X \subset \mathbb{P}^1$ indexed by the complex disc Δ is a function $\phi : \Delta \times X \to \mathbb{P}^1$ such that

(1) $\phi(0, x) = x$,
(2) $t \mapsto \phi(t, x)$ is holomorphic for each fixed x,
(3) $x \mapsto \phi(t, x)$ is injective for each fixed t.

The λ-Lemma of Mane-Sad-Sullivan shows that ϕ must be continuous, and moreover $x \mapsto \phi(t, x)$ must be "quasi-conformal". (One must use a different definition of quasi-conformal, since X is not assumed to be open.)

Theorem 11.6 (Slodkowski). *Any holomorphic motion of $X \subset \mathbb{P}^1$ extends to a holomorphic motion of $\mathbb{P}^1$.*

Surprisingly, it is not even clear that a holomorphic motion of n points can be extended to one of $n+1$ points – this key case is sometimes called the "holomorphic axiom of choice".

The Teichmüller metric. Because of the existence of Teichmüller discs, it is clear that the Teichmüller metric is greater than or equal to the Kobayashi metric. Because of the Bers embedding, it is clear that $\mathcal{T}_g$ is Kobayashi hyperbolic.

Theorem 11.7 (Royden). *The Teichmüller metric is the Kobayashi metric.*

Corollary 11.8. *Every biholomorphism of $\mathcal{T}_g$ is an isometry for the Teichmüller metric.*

Lemma 11.9. *Every holomorphic map from a disc to $\mathcal{T}_g$ can be lifted to a map to the Beltrami differentials on a fixed Riemann surface X.*

Proof of Theorem assuming Lemma. It suffices to show that any holomorphic $\Delta \to \mathcal{T}_g$ is distance non-increasing. We will do this infinitesimally: we wish to show that the Kobayashi unit ball is contained in the unit ball for the Teichmüller metric. This follows from the lifting: the derivative of Δ at 0 must be contained in the unit ball for the Teichmüller (i.e. L^∞) norm on Beltrami differentials. (This shows that the derivative at 0, which is an equivalence class of Beltrami differentials, has a representative in the unit ball for the L^∞ metric. The Teichmüller metric is the inf of the L^∞ norms of all Beltrami differentials in the equivalence class, so hence the Teichmuller norm must be at least as small as the L^∞ norm of the derivative of the lift.) □

Proof of Lemma. Let $\Delta \to \mathcal{T}_g$ be holomorphic. Fix $X \in \mathcal{T}_g$. Composing with the Bers embedding, we get $q_\lambda \in QD(X^*)$ for all $\lambda \in \Delta$. Let $F_\lambda : \mathbb{L} \to \mathbb{C}$ be the unique holomorphic map with $S(F_\lambda) = q_\lambda$ that fixes $0, 1, \infty$. We may view F as a holomorphic motion of the lower half plane union $\mathbb{R}$. By Slodkowski's Theorem, we can extend it to a holomorphic motion of $\mathbb{P}^1$.

Define μ_λ to be the complex dilatation of F_λ. Using an equivariant version of Slodkowski's theorem, we get that μ_λ is equivariant.

The map from Beltrami differentials to $\mathcal{T}_g$, if we view $\mathcal{T}_g$ as a subset of $QD(X^*)$ via the Bers embedding, is just taking Schwarzian derivative on the lower half plane (after extending by 0). Thus evidently μ_λ maps to q_λ and we have our desired section. □

12. Geometric Shafarevich and Mordell

The sources for this section are [**McM00**], [**McM**, Section 10] and notes from lectures of Benson Farb which are not publicly available. (These are not original sources: they are all expository.)

Statements. A family of Riemann surfaces over a base B is defined to be a holomorphic map $E \to B$ of complex manifolds whose fibers are all smooth Riemann surfaces. This induces a map $B \to \mathcal{M}_g$. In fact it is almost but not quite equivalent to giving such a map, and we can safely ignore the difference. (For the experts: one option is to declare that we are interested in families of Riemann surfaces with a level 3 structure.) We assume B is connected, and say that the family is truly varying if the map to $\mathcal{M}_g$ is nonconstant. Every non truly varying family is trivial (a product) after a pulling back the family along a finite cover $B' \to B$. (This is because the group of automorphisms of a higher genus Riemann surface is finite.)

Theorem 12.1 (Geometric Shafarevich)**.** *For each compact Riemann surface B, and each $g > 1$, there are only finitely many truly varying families of genus g Riemann surfaces over B.*

This is equivalent to saying there are only finitely many non-constant holomorphic maps $B \to \mathcal{M}_g$, and this is the perspective that we take.

Remark 12.2. The Lefschetz Hyperplane Theorem and Geometric Shafervich for one dimensional bases B imply Geometric Shafarevich for a higher dimensional base.

Remark 12.3. Since Teichmüller space is a bounded domain, there are *no* truly varying families when B is genus 0 or 1.

Remark 12.4. One can also allow B and the fibers to have finitely many punctures. Allowing the base to have punctures is more significant and more useful; it allows analysis of families over a compact base with some singular fibers, since these fibers can just be removed. Our proof will be only for the case when the base does not have punctures.

The case of B compact is due to Parshin in 1968, and the case with punctures is due to Arakelov in 1971.

Theorem 12.5 (Geometric Mordell)**.** *Each family $E \to B$ has only finitely many holomorphic sections $B \to E$.*

For example, perhaps $g = 2$, and the 6 Weierstrass points can be consistently labelled over the family. (This is always true up to a finite cover.) Then this gives six sections.

Remark 12.6. Using the "Parshin trick" (taking branched covers) one sees that Geometric Shafarevich implies Geometric Mordell (see [**McM00**]).

Analogy between function fields and number fields. Covers $B \to \mathbb{P}^1$ are in bijection with finite extensions of $\mathbb{C}(z)$; one maps B to its field of rational functions $\mathbb{C}(B)$, viewed as a finite extension of the rational functions $\mathbb{C}(z)$ of $\mathbb{P}^1$.

A curve defined over $K = \mathbb{C}(B)$ can be viewed as a family over B of curves defined over B. Indeed, we can view such a curve as the set in $\mathbb{P}^2(K)$ of solutions to a polynomial equation (planar case). One can then consider the subset of $E = \mathbb{P}^2(\mathbb{C}) \times B$ that satisfy the same equation (viewing elements of K as functions on B), and we have a natural projection $E \to B$. This gives a family of curves over B.

In particular, there is a bijection between holomorphic families over B and smooth projective curves over $\mathbb{C}(B)$. There is also a bijection between holomorphic sections of a family and $\mathbb{C}(B)$ rational points.

The non-geometric Mordell conjecture, now know as Faltings' Theorem, says that a higher genus curve defined over a number field K has only finitely many K rational points (and also L rational points, for any finite extension L of K).

The non-geometric Shafarevich conjecture, also proven by Faltings, says that there are only finitely many isomorphism classes of Abelian varieties of fixed dimension and fixed polarization degree over a fixed number field K with good reduction outside of a given finite set of points.

The proof, assuming a black box. We wish to show that there are only finitely many non-constant holomorphic maps $f : B \to \mathcal{M}_g$. We lift this to a map $\mathbb{D} \to \mathcal{T}_g$ which is equivariant for the monodromy representation. We view $\mathcal{T}_g$ as a bounded domain via the Bers embedding. Fatou's theorem (see wikipedia) says that boundary values exist for almost every direction in the disc, and that they can computed via any sequence contained in a wedge (i.e., any sequence approaching the boundary point non-tangentially).

The proof uses the following fact. For almost every boundary value q of f, and any sequence $X_n \in \mathcal{T}_g$ converging to q, we have

(1) $\ell_\alpha(X_n) \to \infty$ for all simple closed curves α and
(2) if $d_{Teich}(X_n, Y_n)$ is bounded, then $Y_n \to q$.

The second point is very intuitive (for example this is obvious if $\mathcal{T}_g$ is replaced with $\mathbb{H}^2$). The first point is a bit more mysterious, although it is perhaps reasonable that most ways of degenerating a marked Riemann surface should increase the lengths of all curves (keeping in mind that the boundary of the D-M bordification is smaller dimensional). (One example of such a sequence is given by precomposing the marking of a fixed point X_1 with powers of a pseudo-Anosov.) The proof of this fact uses that boundary points in the Bers embedding can be interpreted as Kleinian groups, and most of these Kleinian groups that occur on the boundary are "totally degenerate". We will revisit this after completing the proof of Geometric Shafarevich assuming the fact.

Lemma 12.7. *The monodromy $\pi_1(B) \to MCG$ does not fix any curve.*

Proof. Otherwise, the hyperbolic length of this curve would be a well defined continuous function on the compact space B. This contradicts the first item above. $\square$

Lemma 12.8. *For any B and any $b_0 \in B$, there is a compact subset K of $\mathcal{M}_g$ such that for any non-constant holomorphic map $f : B \to \mathcal{M}_g$ we have $f(b_0) \in K$.*

Proof. Fix $L > 0$ so that $\pi_1(B, b_0)$ can be generated by loops of length less than L. By the defining feature of the Kobayashi metric, these loops all map to loops of Kobayashi length at most L in $\mathcal{M}_g$ based at $f(b_0)$. Recall Wolpert's lemma, which says that if X_1, X_2 are hyperbolic surfaces and $\phi : X_1 \to X_2$ is e^L quasiconformal, then $\frac{\ell_{X_1}(\gamma)}{e^L} \leq \ell_{X_2}(\gamma) \leq e^L \ell_{X_1}(\gamma)$.

Recall also that Royden's theorem that the Kobayashi metric is the Teichmuller metric. So, the monodromy of the chosen generators of $\pi_1(B, b_0)$ can change hyperbolic lengths by at most a fixed factor e^L.

Now, let ε be the Margulis constant (so no two curves of length $\leq \varepsilon$ can cross), and suppose in order to find a contradiction that $f(b_0)$ has a curve of hyperbolic length less than $\varepsilon/(e^L)^{3g-2}$.

Let $\alpha_1, \alpha_2, \ldots, \alpha_{3g-3}, \alpha_{3g-2}$ be the shortest $3g-2$ simple closed curves on $f(b_0)$ in order of increasing length. Since the length of α_1 is less than $\varepsilon/(e^L)^{3g-2}$, we get that

$$\ell_{\alpha_{i+1}} > e^L \ell_{\alpha_i}$$

for some i. We then get that the set of curves $\{\alpha_1, \ldots, \alpha_i\}$ must be permuted by the generators of $\pi_1(B)$ and hence by $\pi_1(B)$. Passing to a finite cover of B we can assume that each curve is invariant under $\pi_1(B)$, which contradicts the previous lemma. □

Lemma 12.9. *The boundary values of f are determined by the monodromy representation. (i.e., for two different f with the same monodromy, the boundary values, which are only defined a.e. on S^1, agree a.e..)*

Proof. Every $z \in S^1 = \partial\mathbb{D}$ is the non-tangential limit of an orbit $\gamma_n z_0$, $\gamma_n \in \pi_1(B)$. So we get that the boundary value (defined for a.e. $z \in S^1$) is given by

$$\lim_{n\to\infty} f(\gamma_n z_0) = f_*(\gamma_n) f(z_0),$$

where $f_*(\gamma_n)$ is the monodromy of γ_n.

Suppose that h is another map $B \to \mathcal{M}_g$, and we continue to abuse notation by also using h to denote a lift of h to $\mathbb{D} \to \mathcal{T}_g$. Then we have

$$d(f(\gamma_n z_0), h(\gamma_n z_0)) = d(f_*(\gamma_n) f(z_0), h_*(\gamma_n) h(z_0)) = d(f(z_0), h(z_0))$$

since the mapping class group acts via isometries. It follows from the second point of the fact that the two sequences $f(\gamma_n z_0)$ and $h(\gamma_n z_0)$ converge to the same point in the boundary of the Bers embedding. □

Corollary 12.10. *f is determined by the monodromy representation.*

Proof. Two holomorphic functions from $\mathbb{D}$ to a bounded domain whose boundary values are equal a.e. must be equal. □

Proof of Geometric Shafarevich. First we claim that for any $L > 0$ and any compact set $K \subset \mathcal{M}_g$, only finitely many mapping classes can be realized by loops of Teichmüller length less than e^L based in K. Indeed, such loops must remain in a fixed compact part of $\mathcal{M}_g$. Any loop into a fixed compact part can be isotoped to be simplicial (for some subdivision of the compact part into tiny simplices) with at most a bounded multiplicative and additive length increase, which gives the result.

Now, given a map $B \to \mathcal{M}_g$, b_0 maps to a compact set. (If you are worried about basepoint issues, we can isotope B a bit so b_0 maps to a vertex is the simplicial triangulation, and there are only finitely many vertices.) Now, each generator of $\pi_1(B)$ maps to paths of bounded length, and there are only finitely many mapping classes realized by such paths. Hence there are only finitely many choices of where to send each generator of $\pi_1(B)$, and we get that there are only finitely many possible monodromy presentations. $\square$

The black box. We now revisit the two point fact that we used above: For almost every boundary value q of f, and any sequence $X_n \in \mathcal{T}_g$ converging to q, we have

(1) $\ell_\alpha(X_n) \to \infty$ for all simple closed curves α and
(2) if $d_{Teich}(X_n, Y_n)$ is bounded, then $Y_n \to q$.

To do this, we must study the boundary of the Bers embedding.

Lemma 12.11. *Let q_n be in the image of the Bers embedding, and suppose $q_n \to q$. Let f_n be injective holomorphic functions on the lower half plane with $S(f_n) = q_n$, normalized (using postcomposition with a Möbius transformation) so $f_n(-i) = -i$ and $f_n'(-i) = 1$. Then f_n converges to an injective holomorphic function f, well defined up to Möbius transformations, with $S(f) = q$.*

(Without adding additional normalization, different subsequences can converge to different functions, but they differ by a Möbius transformation.)

Proof. By the Koebe distortion theorem, the sequence f_n is pre-compact. Let f be any limit. Then $S(f) = q$, since the derivatives of f converge.

There is a unique solution to $S(f) = q$ up to Möbius transformations.

f is injective since the limit of injective holomorphic functions is either injective or constant. f cannot be a constant since $f'(-i) = 1$. $\square$

Lemma 12.12. *Let q be in the boundary of the Bers embedding based at $X = \mathbb{H}/\Gamma$, and let f be the function produced by the previous* *lemma. Then f is equivariant with respect to a homomorphism $\rho : \Gamma \to PSL(2, \mathbb{C})$.*

The map ρ is injective and has discrete image, and $f(\mathbb{L})/\rho(\Gamma) = X$.

Furthermore, the map $q \mapsto \rho(\Gamma)$ is a continuous and injective map to the conjugacy classes of discrete faithful representations of Γ into $PSL(2, \mathbb{R})$.

Proof. The equivariance comes from the fact that q is Γ invariant, so if $g \in \Gamma$, then f precomposed with g is again a solution to $S(f) = q$. There is a unique solution up to postcomposition by Möbius transformations, so it must be that $f \circ g = \rho(g) \circ f$.

Since f is injective, this forces ρ to be injective and to have discrete image. f induces a biholomorphism from $X = \mathbb{H}/\Gamma$ to $f(\mathbb{L})/\rho(\Gamma)$.

Continuity follows because solutions to $S(f) = q$ vary continuously. (The proof of injectivity is missing.) □

A limit point in the Bers embedding is called totally degenerate if the domain of discontinuity of $\rho(\Gamma)$ is $f(\mathbb{L})$ (rather than a larger set, containing $f(\mathbb{L})$ as another connected component) and $\rho(\Gamma)$ contains no parabolics. (Later we will see the first condition follows from the second.)

Lemma 12.13. *If q is totally degenerate, then $f(\mathbb{L})$ is dense. Hence, for every $\varepsilon > 0$ there is an $N > 0$ such that for $n > N$, the complement of $f_n(\mathbb{L})$ does not contain an ε ball (for some fixed nice metric on $\mathbb{P}^1$).*

Proof. The limit set of a discrete subgroup of $PSL(2, \mathbb{C})$ is the set of accumulation points of an orbit of the group. As long as the group is "non-elementary", it is the smallest closed invariant subset of the boundary, and the action on its complement is properly discontinuous.

The boundary of $f(\mathbb{L})$ is a closed set invariant under $\rho(\Gamma)$, hence it is equal to the limit set of $\rho(\Gamma)$. If $f(\mathbb{L})$ is not dense, then there must be another component of the domain of discontinuity. □

Lemma 12.14. *Suppose $X_n \to q$ and q is totally degenerate. Then $\ell_\alpha(X_n) \to \infty$ for every curve α.*

Proof. X_n may be presented as the complement of the closure of $f_n(\mathbb{L})$ modulo the action of Γ_n. Let $\gamma_n \in \Gamma_n$ be a representative of α. We assume $\gamma_n \to \gamma$. Assume that the fixed points of γ are z_1 and z_2. Suppose that z_1 and z_2.

Pick a circle around z_1 such that this circle together with its image under f together give an annulus between z_1 and z_2. Say the Euclidean distance between these circles is $d > 0$.

Then, since the ratio of the hyperbolic and euclidean metrics goes to infinity as distance to the boundary goes to zero (uniformly in the domain!), we get that the inf r_n of the ratio of the hyperbolic to Euclidean metrics on the complement of the closure of $f_n(\mathbb{L})$ goes to infinity. Since the translation length of γ_n is approximately at least dr_n, this gives the result. □

Theorem 12.15. *If* Γ *is a limit point in the Bers embedding, then if* Γ *has no parabolics it is totally degenerate.*

Comments on the proof. This follows directly from the next more general theorem.

However, it is good to consider the intuition. A point in the image of the Bers embedding gives a group $\Gamma_{X,Y} \subset PSL(2,\mathbb{C})$ that acts on $\mathbb{P}^1$. The limit set is a circle, which divides $\mathbb{P}^1$ into two topological discs Ω_1 and Ω_2. In our previous notation, the limit set is $f^\mu(\mathbb{R})$, and Ω_i are f^μ of the upper and lower half planes, and the group $\Gamma_{X,Y}$ is the conjugate of the Fuchsian group Γ by f^μ. Here μ arises as the dilatation of a map from X to Y, and we have $\Omega_1/\Gamma_{X,Y} = X^*$ and $\Omega_2/\Gamma_{X,Y} = Y$.

Moving $\Gamma_{X,Y}$ to the boundary of the Bers embedding should have the result of degenerating $\Omega_2/\Gamma_{X,Y} = Y$, whereas $\Omega_1/\Gamma_{X,Y} = X^*$ remains a constant. Hence, we expect to get a limit Kleinian group Γ' with a set Ω_1 on which it acts discontinuously with $\Omega_1/\Gamma' = X^*$. The question is, do we expect any other domain of discontinuity? If there was such another component Ω_2, we might expect Ω_2/Γ' to be some degeneration of Y. But every degeneration of Y has cusps, so then Γ' would have to have a parabolic. □

The following is [**McM**, Theorem 10.14].

Theorem 12.16. *Let* Γ *be a discrete subgroup of* $PSL(2,\mathbb{C})$ *isomorphic to the fundamental group of a surface of genus* $g > 1$*, and assume* Γ *has no parabolics. Then either*

(1) *the domain of discontinuity* Ω *consists of two connected connected components, and* Γ *is not in the boundary of the Bers embedding, or*
(2) Ω *is connected, or*
(3) *the limit set* Λ *is all of* $\mathbb{P}^1$*.*

Proof. We proceed in a number of steps, some of which use quite nontrivial theorems, which we will take as black boxes, but which are intuitive.

Step 1: Λ is connected. Otherwise, a general theorem (Stalling's Theorem) on boundaries of groups says that Γ would be a free product, and this is not the case. (Actually it would have to be either a free product amalgamated over a finite subgroup, or an HNN extension.) (I am not sure how easily this can be justified, if at all, but it may be helpful to think that the boundary of the abstract group Γ is a circle, and one expects the limit set of $\Gamma \subset PSL(2, \mathbb{C})$ to be the continuous image of that boundary circle.)

(Note that it may be tempting to think that in our setup above the limit set is obviously connected because it is the boundary of the topologically embedded disc $f(\mathbb{L})$. However, there are topologically embedded discs whose boundaries are not connected: for example, take a disc, and identify four disjoint intervals in pairs, and embed the result in the plane.)

Step 2: Ω has $0, 1$ or 2 connected components. The Alhfors' Finiteness Theorem states that, for any finitely generated Kleinian group Γ, we have that Ω/Γ has only finitely many connected components, each of which must be a surface of finite type (finite genus and finitely many punctures). (The idea of the proof is that the deformation space of Γ should be the Teichmüller space of Ω/Γ. (For example, in many situations where $\Omega = \mathbb{P}^1$, Mostow rigidity says the deformation space is trivial.) But since Γ is finitely generated, the deformation space, viewed as a subset of the character variety, is obviously finite dimensional. Hence Ω/Γ must have finitely dimensional Teichmüller space.)

By the first step, each component Ω' of Ω is a topological disc. Passing to a finite index subgroup of Γ, assume each is fixed by Γ. (Note Λ and hence Ω are unchanged by passing to finite index subgroups.) Note $\mathbb{H}^3/\Gamma$ is a three manifold with finitely many ends, each of which is a closed surface which maps π_1 isomorphically onto Γ.

If there are at least two ends, pick two. Since their fundamental groups both map isomorphically to the fundamental group of the ambient three manifold, there is an homotopy between two ends (because everything is a $K(\pi, 1)$ here, homotopy classes of maps of spaces correspond to maps of fundamental groups). Hence $\mathbb{H}^3/\Gamma$ has at most two ends. (The homotopy sweeps out a compact set whose complement is contained in the two given ends, so there are at most two ends.)

Step 3: The zero or one component cases. If Ω has zero components it is empty, so $\Lambda = \mathbb{P}^1$. If Ω has one component this means exactly that Ω is connected.

Step 4: The two component case. It is a general fact that any such group Γ is "quasifuchsian" [**McM**, Theorem 7.22] and hence in the image of the Bers embedding. □

To see that almost every limit point is totally degenerate, it now suffices to show the following.

Proposition 12.17. *The limit points of a holomorphic map $\mathbb{D} \to \mathcal{T}_g$ almost surely have no parabolics.*

Proof. Otherwise, there is a positive measure set where some specific $\gamma \in \Gamma$ is a parabolic, i.e. has trace 2. The trace of γ defines a holomorphic function on the set of quadratic differentials (which maps to subgroups of $PSL(2,\mathbb{C})$ via holonomy of the complex projective structure). It is a general fact (Privalov's uniqueness theorem for analytic functions) that if a holomorphic function on the disc is constant on a set of positive measure of boundary points, then it is the constant holomorphic function. This implies that the trace of γ must be 2 even on the interior of the Bers embedding, which is a contradiction. □

The final step to establish the black box is to show the following.

Proposition 12.18. *Suppose $X_n \to q$ and q is totally degenerate. If Y_n is bounded distance to X_n, then $Y_n \to q$.*

Proof. Otherwise, we can pass to a subsequence and assume that $Y_n \to q' \neq q$.

Since Y_n and X_n are K q.c., there is a K q.c. map from $\Gamma_{basepoint,X_n}$ and $\Gamma_{basepoint,Y_n}$. Indeed, consider a map from basepoint to X_n with Beltrami differential μ, and let ν be the Beltrami differential of the composition of that map with the K q.c. map form X_n to Y_n. Up to Möbius transformation, f^ν can constructed by composing with f^μ by the function with the dilatation of $X_n \to Y_n$ on $f^\mu(\mathbb{H})$ and conformal of $f^\mu(\mathbb{L})$. Then with these normalizations, f^μ composed $(f^\nu)^{-1}$ is K q.c. and we conjugate $\Gamma_{basepoint,X_n}$ by this map to get $\Gamma_{basepoint,Y_n}$. (Note the space of K-q.c. maps is compact, so there is a limit K q.c. map also.)

Then Sullivan rigidity states that the two limit groups are equal. Then we must also comment that the limiting groups being equal implies that the points in the Bers embedding are equal, which again follows since monodromy determines $f|_{\mathbb{R}}$. □

Remark 12.19. There is also an approach to Geometric Shafarevich using harmonic maps and the WP metric, see the nice survey "Harmonic mappings and moduli spaces of Riemann surfaces" the references therein.

References

[CMSP03] James Carlson, Stefan Müller-Stach, and Chris Peters, *Period mappings and period domains*, Cambridge Studies in Advanced Mathematics, vol. 85, Cambridge University Press, Cambridge, 2003.

[FM12] Benson Farb and Dan Margalit, *A primer on mapping class groups*, Princeton Mathematical Series, vol. 49, Princeton University Press, Princeton, NJ, 2012.

[GL00] Frederick P. Gardiner and Nikola Lakic, *Quasiconformal Teichmüller theory*, Mathematical Surveys and Monographs, vol. 76, American Mathematical Society, Providence, RI, 2000.

[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.

[McM] Curtis McMullen, *Riemann surfaces, dynamics and geometry*, `http://www.math.harvard.edu/~ctm/home/text/class/harvard/275/rs/rs.pdf`.

[McM00] Curtis T. McMullen, *From dynamics on surfaces to rational points on curves*, Bull. Amer. Math. Soc. (N.S.) **37** (2000), no. 2, 119–140.

[Wol10] Scott A. Wolpert, *Families of Riemann surfaces and Weil-Petersson geometry*, CBMS Regional Conference Series in Mathematics, vol. 113, Published for the Conference Board of the Mathematical Sciences, Washington, DC; by the American Mathematical Society, Providence, RI, 2010.

[Wri15] Alex Wright, *Translation surfaces and their orbit closures: an introduction for a broad audience*, EMS Surv. Math. Sci. **2** (2015), no. 1, 63–108.