Seminar
Harvard University — Spring 1999
Course Notes

1 Period matrices

1. The case of an annulus A. Using the Riemann mapping theorem, A is isomorphic to $A(R) = \{1 < |z| < R\}$. The natural invariant is the *modulus*, $\mathrm{mod}(A) = \log R$. This is both the ratio of height to radius, *and* a measurement of the height with respect to the unique invariant holomorphic 1-form with period 2π, namely dz/z. Thus we have one motivation for the use of 1-forms as moduli.

2. Holomorphic 1-forms. On a compact Riemann surface there are no holomorphic functions. However there may be functions well-defined up to translation, i.e. holomorphic *1-forms*.

3. Let $\Omega(X)$ denote the complex vector space of holomorphic 1-forms $\omega = \omega(z)\,dz$. Every such 1-form can be thought of as a local map $X \to \mathbb{C}$, well-defined up to an additive constant.

4. Periods. By integrating along loops any 1-form gives a map $\pi_1(X) \to \mathbb{C}$ called the *period homomorphism*. Since $\mathbb{C}$ is abelian, ω determines an element of $H^1(X, \mathbb{C})$. Thus we have a map $\Omega(X) \to H^1(X, \mathbb{C})$. This map is *injective* when X is compact; since a form with no periods can be integrated.

   We can also take periods of anti-holomorphic forms, $\overline{\Omega(X)} \to H^1(X, \mathbb{C})$. These $(1, 0)$ and $(0, 1)$ give *all* harmonic 1-forms (the exterior differentials of harmonic functions). Since the period map is injective on both, we find that for X compact:

   $$\dim \Omega(X) \leq g(X).$$

   In fact equality holds, but this is nontrivial.

5. The degree. Any type of holomorphic object (function, vector field, 1-form, section of a line bundle) on a compact Riemann surface has a well-defined *degree*: $\deg(f) = Z(f) - P(f)$, the number of zeros minus the number of poles.

For example: if f is a function, $\deg(f) = 0$. (One proof: apply the residue theorem to df/f.) If ω, ω' are two objects of the same type, then $f = \omega/\omega'$ is a meromorphic function, so $\deg(\omega) = \deg(\omega')$.

On the sphere, a vector field has degree two. Thus its dual, a 1-form, has degree -2. A quadratic differential has degree -4.

On a Riemann surface of genus g, the canonical bundle (whose sections are holomorphic 1-forms) has degree $2g - 2$.

6. Examples.

   (a) dz on $\mathbb{C}$; dz/z on $\mathbb{C}^*$. There are no holomorphic 1-forms on $\widehat{\mathbb{C}}$. (Otherwise you could integrate them to make a bounded holomorphic function).

   (b) $\omega = dz$ on $X = \mathbb{C}/\Lambda$. This is the unique 1-form up to scale. Why? If ω' is another, then ω'/ω would be a holomorphic function.

   Notice that the generators of Λ are *exactly* the periods of ω.

   (c) Let X be the curve $f(y) = g(x)$ in $\mathbb{P}^2$ (assumed to be smooth). Then $f'(y)dy = g'(x)dx$ on X, and we can consider

$$\omega = \frac{dx}{f'(y)} = \frac{dy}{g'(x)}.$$

   (One should think of x and y as two meromorphic functions on X, satisfying the relation $f(y) = g(x)$.

   Notice that at every point of X, at least one of f' and g' is non-vanishing (by smoothness), so this form is well-defined.

   (d) Hyperelliptic curves. As a particular case, take $y^2 = p(x)$, where $p(x)$ has degree $2d$. Then $\omega = dx/2y = dx/2\sqrt{p(x)}$. It was to consider integrals of such forms (which really live on a double branched cover of $\widehat{\mathbb{C}}$) that lead to the theory of Riemann surfaces and their moduli.

   (e) Elliptic curves again. If you look at the form $\omega = dx/p(x)$ for X the elliptic curve $y^2 = p(x) = (x-a)(x-b)(x-c)(x-d)$, the simple loops enclosing (a, b) and (b, c) lift to a basis for $H_1(X, \mathbb{Z})$, so the intervals of $\omega = dx/\sqrt{p(x)}$ over these curves allow one to find the lattice associated to X.

   Exercise: How to the periods behave as the cross-ratio $[a : b : c : d]$ tends to zero?

(f) Billiards. Consider an L-shaped pool table. Its double X is really the Riemann sphere with a quadratic differential ϕ (5 simple poles at the corners, one simple zero at the crook of the L). Letting P be the points where ϕ vanishes, there is a map of $\pi_1(X - P) \to \mathbb{Z}/2$ measuring the failure of ϕ to be of the form ω^2. Taking the 2-fold branched cover (which is a surface of genus 2), we get a holomorphic differential with an order 2 zero and no poles.

Exercise: compute its periods.

(g) Curves in the plane. Let $X \subset \mathbb{P}^2$ be a smooth curve of degree d. The adjunction formula relates the canonical bundles K_X and $K_{\mathbb{P}^2}$ – you get

$$K_X = \mathcal{O}(d-3)|X$$

(since $K_{\mathbb{P}^2} = \mathcal{O}(-3)$). Since the space of sections of $\mathcal{O}(n)$ is $(n+1)(n+2)/2$ dimensional, you find a curve of degree d has $(d-1)(d-2)/2$ holomorphic 1-forms in this way, which agrees with the genus.

Note that a canonical divisor on X is just the intersection of X with a curve of degree $d-3$. For quartics this means the intersection of X with a line, giving the canonical embedding.

Where do the g 1-forms come from? Theorem. If X is a compact Riemann surface of genus g, then $\dim \Omega(X) = g$.

Idea of the proof. Use the Hodge theorem: since the (1,0) and (0,1) parts of a harmonic form are holomorphic.

The proof must be somewhat hard since the ratio of 2 holomorphic 1-forms gives a nonconstant meromorphic function on X, and showing such a function exists is the main step in showing every compact Riemann surface is projective.

7. Quadratic differentials. For hyperelliptic surfaces you can't get holomorphic 1-forms on X by pulling back from $\widehat{\mathbb{C}}$ under a 2-1 map, since the push-forward does *not* create a pole; in fact the push-forward is always zero. Thus all holomorphic 1-forms on X are *antisymmetric.*

A better way to relate forms on X to those on $\widehat{\mathbb{C}}$ is by using their *squares*, i.e. quadratic differentials. These are *always* symmetric. For example, if $X \to \widehat{\mathbb{C}}$ is branched over 6 points, by taking the quadratic differential with simple poles at 4 of them, pulling back and taking the square-root, we obtain a holomorphic 1-form with simple zeros at the other two branch points.

8. The area. There is a natural size to a holomorphic 1-form, namely the total area it gives to X:

$$\|\omega\|^2 = \int_X |\omega(z)|^2 \, |dz|^2.$$

   Now each vector $v \in TX$ determines a linear functional on $\Omega(X)$, by $\omega \mapsto \omega(v)$, and we thus get a natural metric on X (the *Bergman metric*) defined by

$$g(v) = \sup_\omega |\omega(v)|^2 / \|\omega\|^2.$$

   It is a beautiful theorem of Kazhdan's that if one takes a suitable sequence of finite covers X_n of X, then the Bergman metrics (which are of course natural and descends back to X) converge to the hyperbolic metric of constant curvature -1.

9. A canonical metric on $\Omega(X)$. The size of a holomorphic 1-form ω can be measured by the area of $f(X)$, where $\omega = df$. This makes sense locally, and we just sum over local patches covering X. In this way we get a natural L^2-norm:

$$\|\omega\|^2 = \frac{i}{2} \int_X \omega \wedge \overline{\omega}.$$

   Notice that $(i/2) dz \wedge d\overline{z} = dx\, dy$ so the integral above does indeed measure the area of the image. Of course the Jacobian of f is always positive, so the above norm is non-degenerate.

10. Periods. Now it should not be unexpected that the norm of ω can also be calculated in terms of its periods. For example, if $X = \mathbb{C}/\Lambda$ and $\omega = dz$, then its periods $a, b \in \mathbb{C}$ on a basis for $\pi_1(X)$ determine the fundamental parallelogram and we have

$$\|\omega\|^2 = \operatorname{Re} a\overline{b}.$$

11. To treat a general Riemann surface, choose a homology basis

$$\langle \alpha_1, \ldots, \alpha_g, \beta_1, \ldots, \beta_g \rangle$$

with intersection numbers $\alpha_i \cdot \beta_j = \delta_{ij}$. For ease of visualization, we can take each pair (α_i, β_i) as simple loops meeting just once, in disjoint handles as i varies. The intersection number means we can think of α_i pointing along the positive real axis and β_i along the positive imaginary axis at the point of intersection.

Let a_i, b_i be the corresponding periods of ω. Cut X along the curves $(\alpha_i, \beta_i)_1^g$ to obtain a surface with boundary U. Then ω has no periods on Y so it integrates to a function $f : Y \to \mathbb{C}$. The key point is to notice, as in the case of a torus, that if $p, p' \in \partial U$ are two points that get glued together onto α_i, then $f(p') - f(p) = b_i$, the period along b_i. Similarly, if p, p' end up on b_i, then $f(p') - f(p) = -a_i$.

Now we use the fact that $d(z\, d\overline{z}) = dz\, d\overline{z}$. Thus by Stokes' theorem we have

$$\int_X \omega \wedge \overline{\omega} = \int_{\partial Y} f \wedge d\overline{\omega}.$$

The curve α_i, for example, occurs twice along ∂Y, giving a contribution of:

$$\int_{\alpha_i} (f(p) - f(p')) \wedge \overline{\omega} = -b_i \overline{a}_i.$$

Similarly from the curves β_i we get

$$\int_{\beta_i} (f(p) - f(p')) \wedge \overline{\omega} = a_i \overline{b}_i.$$

So altogether we have the famous formula:

$$\int_X \omega \wedge \overline{\omega} = \sum a_i \overline{b}_i - b_i \overline{a}_i. \tag{1.1}$$

12. The Riemann matrix. Notice from the beautiful formula above that if the a-periods of ω vanish, then so does ω. Since there are exactly g a-periods, the map $\omega \mapsto a_i(\omega)$ gives an isomorphism $\Omega(X) \cong \mathbb{C}^g$. Thus we can choose a canonical basis ω_i (adapted to our basis for the homology of X) such that $a_i(\omega_j) = \delta_{ij}$.

The *Riemann period matrix* of X is then given by

$$Z_{ij} = b_i(\omega_j).$$

Then if ω has a-period a_i, its b-periods are $b_i = \sum_j Z_{ij} a_j$. In other words, $b = Za$. Putting this into the formula for the norm we get:

$$\|\omega\|^2 = \frac{i}{2}\int \omega \wedge \overline{\omega} = \frac{i}{2}(a\overline{Z}\overline{a} - \overline{a}Za) = \operatorname{Im}(\overline{a}Za).$$

13. Symplectic forms. We can interpret the period formula (1.1) as follows. The periods give a map

$$H_1(X,\mathbb{Z}) \to \Omega(X)^*.$$

This map is nondegenerate, else there would be 1-forms with zero periods.

The symplectic structure coming from the intersection form on $H_1(X,\mathbb{Z})$, plus the complex structure on $\Omega(X)^*$, determines a metric on $\Omega(X)$. It works out like this: given a symplectic basis (a_i, b_i) for $H_1(X,\mathbb{Z})$, we think of a 1-form as a linear map $\alpha : H_1(X,\mathbb{Z}) \to \mathbb{C}$; then we set

$$\|\alpha\|_0^2 = \sum |\alpha(a_i) \wedge \alpha(b_i)|$$

where the wedge product is in $\mathbb{C}$. In other words we look at the sum of the (signed) areas of the images of the squares spanned by (a_i, b_i).

Since $|a \wedge b| = \operatorname{Im}(\overline{a}b) = (\overline{a}b - \overline{b}a)/2i$, we find

$$\|\alpha\|_0^2 = \frac{i}{2}\sum a_i\overline{b}_i - \overline{a}_i b_i = \frac{i}{2}\int \alpha \wedge \overline{\alpha} = \|\alpha\|^2,$$

the area norm we introduced already.

Summing up, the complex structure on $\Omega(X)^*$ together with the symplectic form on $H_1(X,\mathbb{Z})$ determines the norm on forms.

14. The Jacobian. The Jacobian variety of X, build from this data, is the complex torus

$$\operatorname{Jac}(X) = \Omega(X)^*/H_1(X,\mathbb{Z}).$$

The area metric, or equivalently the sympletic form on $H_1(\operatorname{Jac}(X),\mathbb{Z})$, is an additional piece of data called a *principal polarization* of $\operatorname{Jac}(X)$. It gives a shape to the Jacobian.

15. Examples: in genus 1, the Jacobian is X normalized to have area $= 1$.

16. There is a natural map $C \to \mathrm{Jac}(C)$ which is obtained by simultaneously integrating all the 1-forms.

    Exercise. The map $C \to \mathrm{Jac}(C)$ is the universal map to a complex torus; that is, any other map of C to a complex torus factors through this one. (Hint: pull back 1-forms from the target torus.)

17. Famous problem. Let C have genus $g \geq 2$; let C be embedded in a complex torus A; and let $G \subset A$ be a finitely generated subgroup. Prove that $G \cap C$ is finite.

18. Given the image of C in $\mathrm{Jac}(C)$, it is not too hard to see the polarization: namely given local projection $\pi : \mathrm{Jac}(C) \dashrightarrow \mathbb{C}$, we can calculate the area of $\pi(C)$ as a measurement of the size of π. This gives a metric on the cotangent space to the origin in $\mathrm{Jac}(C)$, and by translation and duality a metric on $\mathrm{Jac}(C)$.

19. The Torelli Theorem asserts that $(\mathrm{Jac}(X), \omega)$ determines X up to isomorphism.

20. Problem: show there are only finitely many possiblities for X, given $\mathrm{Jac}(X)$ without its polarization. That is, a given Abelian variety A admits only finitely many principal polarizations up to the action of $\mathrm{Aut}(A)$.

21. Analogue: a given real torus, $A = \mathbb{R}^n/\mathbb{Z}^n$, admits only finitely many unimodular metrics, up to isomorphism. By this we mean the square length of every closed loop should be an integer, and the total volume of the torus should be one.

    The first nontrivial example is $E_8 \subset \mathbb{R}^8$. This lattice is given by $\Lambda + p$, where $\Lambda \subset \mathbb{Z}^8$ consists of those points with $\sum x_i = 0 \bmod 2$, and $p = (1, 1, 1, 1, 1, 1, 1, 1)/2$. (Note that $p \cdot p = 1$).

22. Billiards on a rectangle. Consider a double cover $\pi : E \to \widehat{\mathbb{C}}$ branched over 4 points.

    The basic topological unit to know is that a 2-fold covering of D^2, branched over 2 points, is an annulus. From this it is easy to understand how the $(2, 2, 2, 2)$-orbifold is covered by a torus.

If we endow the sphere with a quadratic differential ϕ that makes it into the double of an $A \times B$ rectangle, then we see the periods of the 1-form with $\theta^2 = \pi^*\phi$ are just $2iA$ and $2B$. Thus $\tau = iA/B \in \mathbb{H}$. Also the area of the torus is clearly $4AB$ which is twice the area of the double of the rectangle.

23. Billiards on an L-shaped table. Now consider a double cover $\pi : X \to \widehat{\mathbb{C}}$ branched over 6 points on the real line.

    Choose a quadratic differential ϕ making the sphere into the double of an upside-down L, a rectangle of dimensions $A \times B$ (height by width) stacked on top of one of dimensions $D \times C$, with $C < B$. The boundary of the thickened L is combinatorially a hexagon. Disregarding two opposite edges (of lengths $A+D$ and $B-C$), the two pairs of remaining edges give a homology basis for the double cover. The periods of θ, with $\theta^2 = \pi^*(\phi)$ as before, are $(-iA, B, C, iD)$. Then we have:

$$\begin{aligned}\int_X |\omega|^2 &= \frac{i}{2}\int \omega \wedge \overline{\omega} = \frac{i}{2}\sum a_i\overline{b}_i - \overline{a}_i b_i \\ &= 2I[-iAB - iCD - iAB - iCD] = 4(AB + CD),\end{aligned}$$

    which is twice the area of the (doubled) L.

2 The Mandelbrot set is connected

1. Let $f_c(z) = z^2 + c$ be a monic quadratic polynomial with critical points at 0 and ∞. The *Mandelbrot set* is defined according to the behavior of the critical point $z = 0$, by

$$M = \{c \ : \ f_c^n(0) \text{ remains bounded as } n \to \infty\}.$$

   The *filled Julia set* $K(f_c)$ is the set of $z \in \widehat{\mathbb{C}}$ that do not tend to infinity under iteration; its boundary is the *Julia set* $J(f_c)$. Thus $c \in M$ iff $c \in K(f_c)$.

   It is easy to see that M is closed and bounded. Naive computer studies make it look like M has many tiny islands, but on closer inspection this islands are connected by filamentary bridges to the mainland.

   Our goal is to use the theory of Riemann surfaces to prove:

   Theorem. The Mandelbrot set is connected.

2. *Escape function.* Let $\Omega(f_c) = \widehat{\mathbb{C}} - K(f_c)$ denote the basin of attracting of infinity. Letting $\log^+|z| = \max(\log z, 0)$, we define the escape rate of z by

$$\phi_c(z) = \lim 2^{-n} \log |f_c^n(z)|.$$

For example, when $c = 0$ we have $\phi_c(z) = \log |z|$.

   (a) $\phi_c(z) \geq 0$ is harmonic on the basin of infinity $\Omega(f_c)$,

   (b) $\phi_c(z) = \log |z| + O(1)$ near ∞,

   (c) $\phi_c(z)$ extends continuously to be zero on $K(f_c)$, and

   (d) $\phi_c(f_c(z)) = 2\phi_c(z)$. That is, ϕ_c sends the dynamics of $f_c(z)$ to the dynamics of $x \mapsto 2x$ on $\mathbb{R}$.

3. *If $K(f_c)$ is connected.* Let

$$\Phi_c : (\widehat{\mathbb{C}} - K(f_c)) \to (\widehat{\mathbb{C}} - \overline{\Delta}) = \Sigma.$$

be the Riemann mapping normalized by $\Phi_c(\infty) = \infty$ and $\Phi_c'(\infty) = 1$. Then Φ_c conjugates f_c to a proper, degree two holomorphic map $S : \Sigma \to \Sigma$, fixing infinity; by the normalization of the derivative, we find $S(z) = z^2$. Thus:

$$\Phi_c(f_c(z)) = \Phi_c(z)^2,$$

and it is then easy to see:

$$\phi_c(z) = \log |\Phi_c(z)|.$$

Moreover we can construct $\Phi_c(z)$ by iteration:

$$\Phi_c(z) = \lim (f_c(z)^n)^{1/2^n}.$$

4. *A canonical metric and foliation.* Define

$$\omega_c = 2\overline{\partial}\phi_c.$$

Since ϕ_c is harmonic, ω_c is a holomorphic 1-form on $\Omega(f_c)$.

When $c = 0$, we have $\omega_c = dz/z$; When $K(f_c)$ is connected, we have $\omega_c = \Phi_c^*(dz/z)$.

In the $|\omega_c|$ metric, f_c acts by linear expansion by 2. The map f_c also preserves the oriented foliation attached to ω_c.

In this metric, $\Omega(f_c)$ is a half-infinite cylinder of radius 1, foliated by round circles; and $f_c : \Omega(f_c) \to \Omega(f_c)$ is a local similarity giving a double covering.

The leaves of this foliation are the closure of the small-orbit equivalence relation.

5. *If $K(f_c)$ is disconnected.* In this case, ϕ_c and ω_c are still globally defined, while Φ_c is only defined near infinity. Moreover the critical point must be in $\Omega(f_c)$, else successive preimages of a small disk around infinity would remain simply-connected and exhaust the basin.

   We can now draw a picture of the basin with the $|\omega_c|$ metric and foliation in place. Near infinity the leaves are closed loops. There is a dynamically determined annulus A_c bounded by the leaves through c and $f_c(c)$; let (M_c, θ_c) be its modulis and the angle between the marked boundary points c and $f_c(c)$.

   The preimage B_c^0 of A_c is bounded by a simple curve on the outside and a figure eight through the critical point $z = 0$ on the inside; it has modulus $M/2$. Its preimages B_c^1 consists of two annuli, one inside each lobe of the figure eight. Similarly, $B_c^n = f_c^{-n}(B_c^0)$ consists of 2^n annuli, all of the same modulus.

6. *Insulated Riemann surfaces.* Here are some basic results on annuli and Riemann surfaces.

   (a) If $K \subset U$ is a full continuum in a disk, then $\text{mod}(U - K) = \infty$ iff K is a point.

   (b) If disjoint annuli A_i are nested in B, then $\text{mod}(B) \geq \sum \text{mod}(A_i)$.

   (c) Any planar Riemann surface X embeds in $\widehat{\mathbb{C}}$.

   (d) A Riemann surface is *insulated* if every end is cut off by annuli A_i with $\sum A_i = \infty$.

   The embedding $X \subset \widehat{\mathbb{C}}$ for a planar, insulated Riemann surface is unique up to $\text{Aut}(\widehat{\mathbb{C}})$, and every end is a single point.

7. Theorem. For $c \notin M$, the set $K(f_c) = J(f_c)$ is a Cantor set.

   Proof. Every end is isolated by an infinite sequence of annuli $A_i \in B_c^i$ with $\text{mod}(A_i) = M_c/2$. ∎

8. Theorem. The map $c \mapsto (M_c, \theta_c)$ is a bijection from $\mathbb{C} - M$ to $\mathbb{R}_+ \times S^1$.

   Proof. We will construct an inverse. Given (M, θ), we can build the annulus A_0 of modulus M with 2 marked points x, x' that are θ apart. This A is the candidate for A_c. Attaching annuli A_i of modulus $2^i M$, twisted by $2^i\theta$, we obtain a Riemann surface $Y = \bigcup A_i$ with a map $F : Y \to Y$ sending A_i to A_{i+1} and x to x'.

   Next we extend backwards, by building the annulus $B_0 \cong A/2$ and pinching together the two preimages of x on the inner boundary to obtain a figure eight. We glue B_0 to A and B_{i+1} to B_i, where B_i consists of 2^i copies of B_0. In the end we obtain a *proper map* $F : X \to X$ defined on a *planar, insulated* Riemann surface $X = (\bigcup A_i) \cup (\bigcup B_i)$.

   Now embed X into $\widehat{\mathbb{C}}$, so that the A-end goes to infinity and so that x goes to zero. Let $\sqrt{X}$ denote the preimage of X. Since $F : X \to X$ is simply a degree two covering of $X - \{x\}$, there is an isomorphism $\iota : X \to \sqrt{X} \subset \widehat{\mathbb{C}}$ such that the diagram

$$\begin{array}{ccc} X & \xrightarrow{\ \iota\ } & \sqrt{X} \subset \widehat{\mathbb{C}} \\ {\scriptstyle F}\downarrow & & \downarrow{\scriptstyle z^2} \\ X & \longrightarrow & \widehat{\mathbb{C}} \end{array}$$

   commutes.

   Since the embedding of X is essentially unique, we have $\iota(z) = az + b$ for some a, b, and thus F is the restriction to X of a quadratic polynomial f_c. Up to conjugation, this polynomial can be uniquely written in the form $f_c(z) = z^2 + c$. ∎

9. *The Riemann mapping for M.* Alternatively, we find that

$$\exp(M_c + i\theta_c) = \Phi_c(c) = h(c).$$

   That is, if $z = \Phi_c(c)$ then A_c is isomorphic to the annulus bounded by the circles of radius $|z|$ and $|z|^2$, which has modulus $M = \log|z|$. Similarly, the marked points z and z^2 have angles differing by $\theta = \arg(z^2) - \arg(z) = \arg(z)$. Thus the (M_c, θ_c) parameterization of $\mathbb{C} - M$ agrees with the Riemann mapping.

10. Summary: $K(f_c)$ is connected if $c \in M$, else $K(f_c)$ is a Cantor set, determined uniquely by coordinates (M_c, θ_c).

    One can also show all f_c with $c \notin M$ are quasiconformally conjugate; the intrinsic construction of coordinates on $\mathbb{C} - M$ is the beginnings of a Teichmüller theory for rational maps.

    See [McS].

11. *Schottky groups.* A finitely-generated free group $\Gamma \subset \mathrm{Aut}(\widehat{\mathbb{C}})$ is called a *Schottky group.*

    To produce examples, take any compact Riemann surface X of genus $g \geq 1$, choose a homeomorphism $X \cong \partial H_g$ between X and the boundary of a handlebody of genus g, and let $Y \to X$ be the covering space determined by the kernel of the map

    $$\pi_1(X) \to \pi_1(H_g) \cong Z^{*g}.$$

    Then $\Gamma \cong \pi_1(H_g)$ acts on Y by deck transformations. In addition, Y is an insulated planar surface, so there is an essentially unique embedding $Y \subset \widehat{\mathbb{C}}$. By uniqueness, Γ extends to a discrete group of Möbius transformations (a *Kleinian group*) acting on $\widehat{\mathbb{C}}$.

    A deep theorem of Ahlfors implies that any Schottky group without parabolics and with limit set $\Lambda \neq \widehat{\mathbb{C}}$ arises in this way.

12. *Genus one Schottky groups.* In this case $\Gamma \cong \langle z \mapsto \lambda z \rangle$, $|\lambda| > 1$, and X is the complex torus with $\tau = 2\pi i / \log \lambda$.

13. *Reflection groups.* To obtain concrete examples of Schottky groups, take a collection $C_1, \ldots, C_{g+1}$ circles bounding disjoint closed disks in $\widehat{\mathbb{C}}$. Then the group Γ generated by reflections in the C_i has a nonorientable quotient orbifold consists of a planar surface with $g+1$ boundary components. Its double is a surface of genus g, uniformized by the orientation-preserving subgroup $\Gamma_0 \subset \Gamma$.

14. *Boundary of Schottky space.* As the circles come together and touch, the quotient surface develops cusps. For example, a chain of 4 tangent circles has a quotient surface consisting of 2 4-times punctured spheres (arising as a limit of a surface of genus 3). Each 4-times punctured sphere can further degenerate into a pair of 3-times punctured spheres.

This degeneration suggests compactifying the space of Riemann surfaces by allowing these kind of limits.

3 Teichmüller space

1. Let S be a compact, oriented surface, possibly with boundary. A Riemann surface X is of *finite type* if it is isomorphic to a compact surface with a finite number of points removed. (For example, every affine or projective curve is of finite type, and every finite type Riemann surface can be realized as such a curve.)

   A *marking* of X is a homeomorphism $f : \mathrm{int}(S) \to X$, preserving orientation. Two markings of X are *equivalent* if there induce the same isomorphism
   $$f_* : \pi_1(S) \to \pi_1(X).$$
   It is known that equivalent markings are isotopic.

   The space of *marked Riemann surfaces* $(f : S \to X)$ is the *Teichmüller space* of S. Two marked surfaces (f, X) and (g, Y) are equivalent if there is an isomorphism $\alpha : X \to Y$ such that the diagram

   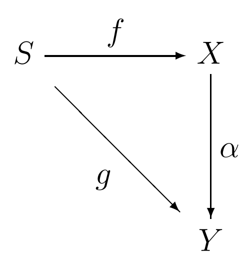

   commutes on the level of fundamental groups. That is, g and $\alpha \circ f$ should be equivalent markings of Y.

2. The *mapping-class group* or *modular group*, $\mathrm{Mod}(S) = \mathrm{Diff}^+(S)/\mathrm{Diff}_0(S)$, is the group of homotopy classes of homeomorphisms $h : S \to S$. It acts on $\mathrm{Teich}(S)$ by
   $$h \cdot (f : S \to X) = (f \circ h^{-1} : S \to X).$$

3. *Flat tori.* If X is a torus, then it admits a unique flat metric of area 1 in its conformal class. Up to rescaling, this metric is obtained from a

cylinder of circumference 1 and height y by gluing the top to bottom with a twist of length x.

If we fix S to be a standard torus, obtained from a cylinder of height 1 with no twist, then there is also a natural marking $f : S \to X$. Thus we have

$$\mathrm{Teich}(S) \cong \mathbb{H}$$

by $(f, X) \mapsto x + iy \in \mathbb{H}$.

The mapping-class group here is $\mathrm{Mod}(S) = SL_2(\mathbb{Z})$. Note that the center acts trivially.

4. *Fixed-points of mapping-classes.* In general, h fixes $X \in \mathrm{Teich}(S)$ iff h can be realized in $\mathrm{Aut}(X)$, the conformal automorphism group of X.

5. *Hyperbolic isometries.* We now consider a compact, hyperbolic Riemann surface X. Then $X = \mathbb{H}/\Gamma$ for some group $\Gamma \subset \mathrm{Isom}(\mathbb{H})$ with $\Gamma \cong \pi_1(X)$.

   Every isometry of hyperbolic space is fixes a point in $\overline{\mathbb{H}}$. For the universal cover of a compact hyperbolic surface, only hyperbolic elements arise; each stabilizes a unique geodesic. The stabilizer of a given geodesic is cyclic.

   In particular: *there exists a unique geodesic* in every nontrivial homotopy class.

   Also: *any simple loop is represented by a simple geodesic.*

   Exercise: *geodesics intersect as efficiently as possible.*

6. *Pairs of pants.* Theorem. For any lengths $A, B, C \geq 0$, there exists a unique hyperbolic pair of pants P with cuffs of length A, B, C. Here ∂P is totally geodesic, or a cusp when the length is zero.

7. *Spherical triangles.* As a warmup, we discuss:

   Theorem. Given angles $0 \leq \alpha, \beta, \gamma \leq \pi$ with $\pi < \alpha + \beta + \gamma$, there is a spherical triangle T (unique up to congruence) with these internal angles.

   Proof. Identify great circles on S^2 with unit vectors in $\mathbb{R}^3$. Then we are seeking unit vectors A, B, C with inner products $\langle A, B \rangle = \cos \gamma$, etc. Now *define* a quadratic form on $\mathbb{R}^3$ by

$$Q = \begin{pmatrix} 1 & \cos(\gamma) & \cos(\beta) \\ \cos(\gamma) & 1 & \cos(\alpha) \\ \cos(\beta) & \cos(\alpha) & 1 \end{pmatrix}.$$

The condition on the angles is equivalent to the condition that Q is positive definite. Thus we can make a change of coordinates sending Q to the standard Euclidean inner product; and then A, B, C map to the required vectors. ■

8. *Hyperbolic triangles.* Theorem. Given $0 < \alpha, \beta, \gamma < \pi$, with $\alpha+\beta+\gamma < \pi$, there is a unique hyperbolic triangle with these angles.

   Proof. Fixing α and β joined by an edge of length C, the remaining angle moves from $\pi - \alpha - \beta$ to 0 montonically as the length of C increases. ■

9. *Virtual triangles.* If we have 3 disjoint geodesics A, B and C in $\mathbb{H}$, the distances $\alpha = d(B, C)$, $\beta = d(C, A)$ and $\gamma = d(A, B)$ play the rule of angles. We can similarly construct A, B and C with any specified α, β, γ, unique up to congruence.

   The double of the right-angled hexagon so obtained gives the required pair of pants.

10. *Fenchel-Nielsen coordinates.* Pick $3g - 3$ *unoriented* simple closed curves $\gamma_1, \ldots, \gamma_{3g-3}$ on S, cutting S into $2g - 2$ pairs of pants. (Note that each pair of pants has $\chi(P) = -1$ and $\sum \chi(P_i) = \chi(S) = 2 - 2g$.)

    Given X in Teich(S) and a loop α on S, we can measure the hyperbolic length $\ell_\alpha(X)$ of the geodesic representative of α on X. Let $L_1, \ldots, L_{3g-3}$ be the functions

    $$L_i : \text{Teich}(S) \to \mathbb{R}_+$$

    defined by

    $$L_i(X) = \ell_{\gamma_i}(X).$$

    Now given the lengths $\langle L_i \rangle$, we can form a canonical surface Y by gluing together pairs of pants with cuffs of the same lengths, and with no twist. The twist condition means that on S we also fix loops on S meeting each pair of paints in 3 curves joining adjacent cuffs, and these curves become geodesics perpendicular to each cuff on Y.

To obtain X from Y, we must now twist along the cuffs. Let τ_i denote the total distance twisted along the cuff γ_i. This means X is obtained from Y by performing a *left earthquake* of length τ_i, meaning (for $\tau_i > 0$) that we make a shear along γ_i so that one side moves to the left relative to the other. Note that the direction of motion does *not* use an orientation of γ_i.

Then X is uniquely specified by the data (L_i, τ_i). Moreover every piece of data arises. We have thus shown:

Theorem (Fenchel-Nielsen) The coordinates $(L_i + i\tau_i)$ provide an isomorphism between $\mathrm{Teich}(S)$ and $\mathbb{H}^{3g-3}$. Thus Teichmüller space is a topological cell of real dimension $6g - 6$.

3.1 Automorphisms

11. *Finiteness of automorphisms.* Theorem. For any hyperbolic surface X of finite volume, $\mathrm{Aut}(X)$ is finite.

    Proof. By the Schwarz lemma, $\mathrm{Aut}(X)$ is a group of isometries, so it is compact. On the other hand, any isometry isotopic to the identity *is* the identity, so $\mathrm{Aut}(X)$ is discrete. ■

12. *Bounds for* $\mathrm{Aut}(X)$. The preceding proof is not very effective. For an effective bound, consider the orbifold $Y = X/\mathrm{Aut}(X)$. It has negative Euler characteristic; indeed, $\chi(Y) = \chi(X)/d < 0$ where $d = |\mathrm{Aut}(X)|$. Now the (closed orientable) orbifold Y of greatest negative Euler characteristic is the $(2,3,7)$-orbifold, with $\chi(Y) = -1 + 1/2 + 1/3 + 1/7 = -1/42$. Thus we find, for a closed surface of genus $g > 1$,

    $$|\mathrm{Aut}(X)| \le 42|\chi(X)| = 84(g-1).$$

    This orbifold fits into the sequence of nice manifolds $X(n) = \mathbb{H}/\Gamma(n)$, $n = 2,3,4,5,6,7$, with symmetry groups $PSL_2(\mathbb{Z}/n)$, giving as quotients the $(2,3,n)$ orbifolds. (Removing the n-point gives the $(2,3,\infty)$ orbifold $X(1)$.)

These groups are S_3, A_4, S_4, A_5, (acting the the sphere with 3, 4, 6, 12, punctures), then $PSL_2(\mathbb{Z}/6)$, $PSL_2(\mathbb{Z}/7)$, groups of order 72 and 168 acting on surfaces of genus 0 and genus 3.

Problem: what is the most symmetric surface of genus 2? (Hint: it has 48 automorphisms, not 84!) (Answer: it's the 2-fold cover of the octahedron, which is the most symmetric configuration of 6 points on the sphere.)

13. *Discreteness of length spectrum.* Theorem. Given X and $L > 0$, there are only a finite number of closed geodesic on X with $\ell_\gamma(X) < L$.

    Proof. Thinking of these geodesics as maps of S^1 into X, we again have compactness and discrteness (the latter because homotopic geodesics must be the same). ∎

14. *Moduli space exists.* Theorem. The space $\mathcal{M}(S) = \operatorname{Teich}(S)/\operatorname{Mod}(S)$ is a Hausdorff orbifold.

    Proof. We have finite stabilizers. The action is properly discontinuous by discreteness of the length spectrum. ∎

15. *Faithfulness of* Mod_g. Theorem. For genus $g \leq 2$, the center of the mapping class group is $Z = \mathbb{Z}/2$, and Mod_g/Z acts faithfully on $\mathcal{T}_g$.

    Proof. The most interesting point is that Mod_2 has nontrivial center. One proof by observing that S admits a system of simple closed curves, each invariant under the hyperelliptic involution, such that Dehn twists on these curves generate Mod_g. ∎

    Theorem. For $g \geq 3$ the mapping-class group has trivial center and acts faithfully on $\mathcal{T}_g$.

    Proof. Given S of genus $g \geq 3$, choose a pair of pants $P \subset S$ such that no component of $S - P$ is also a pair of pants. (This requires $g \geq 3$!)

    Now give S a hyperbolic structure such that the cuffs of P are short geodesics with distinct lengths, and no other geodesics on S are short. Then any isometry $f : S \to S$ must satisfy $f(P) = P$, and f must send each cuff of P to itself. Thus $f|P = \mathrm{id}$, and so $f = \mathrm{id}$ on S. Thus $\operatorname{Aut}(S)$ is trivial, and so Mod_g acts faithfully on $\mathcal{T}_g$. ∎

16. *Thick-thin.* Theorem. Let g be a simple closed geodesic on X. Then g admits an embedded collar of width $\asymp |\log L(g)|$ when $L(g)$ is small, and of width $\asymp e^{-L(g)}$ when $L(g)$ is large.

Proof. To every primitive hyperbolic element $g \in \pi_1(X)$, we can associate the tube $T(g) \subset \mathbb{H}$ (possibly empty) of points translated distance less than $\epsilon_0 > 0$.

We claim that for suitable ϵ_0, these tubes are disjoint or equal. Indeed, suppose $T(g)$ and $T(h)$ both contain p. Since $\Gamma \cdot p$ is discrete, we can assume $T(g)$ achieves the minimum translation of p over all elements of Γ. But then $[g,h]$ translates p much less than g or h, so $[g,h] = \text{id}$. But then g and h must stabilize each other's geodesic, so $T(g) = T(h)$ and $g = h^{\pm 1}$.

It follows that the quotient $T(g)/\langle g \rangle$ embeds in X; it is a tube of width about $\log(1/L(g))$, as can be seen by looking in the disk model. (If g acts on Δ, stabilizing the real axis, then g translates $\mathbb{R}$ about $L(g)$, and so its translation distance is $O(1)$ at distance $L(g)$ from S^1. Such points are hyperbolic distance $\log(1/L(g))$ from the center of the disk.)

Now consider a fairly long simple geodesic g, and the tube of width δ around it. If this tube fails to be embedded, then we obtain a geodesic of length $< \delta$ which does not admit a collar of width $L(g)$. It follows that $|\log(\delta)| < L(g)$, which says that $\delta > e^{-L(g)}$. Thus g admits a collar of width $e^{-L(g)}$. ■

3.2 The compactification of $\mathcal{M}_g$ by stable curves

17. *Bounded geodesics.* Theorem. Let X be a compact hyperbolic surface of genus $g \geq 2$. Then:

> (A) X carries a simple closed geodesic of length $O(\log g)$; and (B) X admits a pair-of-pants decomposition by simple geodesics each with length bounded by C_g.

Proof. (A) Otherwise a Dirichlet fundamental domain in the universal cover would contain a ball $B(p,r)$ of radius $\gg \log g$, and hence the area of $B(p,r)$, would be $\gg 2\pi(2g-2) = \text{area}(X)$.

(B) Apply induction, using the fact that on a surface with geodesic boundary, the simple curves cannot come too close to the boundary, and at the same time a long curve must nearly retrace its own path more than once if the injectivity radius of X is bounded below. ■

18. *Example: the 'hairy torus'.* The best known bounds for C_g have a large gap: $A\sqrt{g} \leq C_g \leq Bg$.

    To see the lower bound, puncture a square torus at the points of order n, replaces those punctures by very short geodesic boundary, and double. Then we have a surface of genus about $g = n^2$. After cutting the short curves, any homologically nontrivial curve on the torus has to have length $\geq An = A\sqrt{g}$.

    See [Bus].

19. *Trivalent graphs.* Up to the action of the mapping-class group, the choice of Fenchel-Nielsen coordinates corresponds to the choice of a pair of pants decomposition of S, which in turn corresponds to the choice of a connected trivalent graph G with $b_1(G) = g(S)$.

    Examples: For $g = 2$ there are 2 graphs: the theta and the dumbell. For $g = 3$ there are 4 graphs, with $0 - -3$ single-edge loops. (One is the Wheatstone bridge, i.e. the 1-skeleton of a tetrahedron.)

    Exercise: Show any two trivalent graphs of genus g are connected by a sequence of Whitehead moves.

20. *Mumford's compactness theorem.* Theorem. The set of Riemann surfaces with shortest geodesic of length $L(X) \geq r > 0$ is compact in $\mathcal{M}_g$.

    Proof. Given a sequence of Riemann surfaces $X_n \in \mathcal{M}_g$, choose a pair of pants decomposition of uniformly bounded length for each. Passing to a subsequence, we can assume the combinatorics of this decomposition is constant. Using the corresponding Fenchel-Nielsen coordinates, we can mark each surface is such a way that $L_i(X_n)$ is bounded above for all n. Applying Dehn twists, we can also arrange that $\tau_i(X_n)$ is bounded. Finally by assumption, $L_i(X_n) \geq r$ for all n. Thus the X_n, so marked, lie in a compact subset of Teichmüller space, so they have a convergent subsequence. ■

21. *Isospectral Riemann surfaces.* Theorem (Wolpert). The set of Riemann surfaces isospectral to a given one is finite.

Proof. Isospectral Riemann surfaces have the same length spectrum. A lower bound on the length confines us to a compact subset of moduli space, and then one must move a definite distance to reorganize the lengths of generators. ■

22. *Polar Fenchel-Nielsen coordinates.* Fix simple loops $\mathcal{P} = \langle \gamma_1, \ldots, \gamma_{3g-3} \rangle$ defining a pair-of-pants decomposition of S. We can then associate a Riemann surface with nodes to coordinates where some L_i may become zero.

Introduce coordinates $\theta_i = 2\pi\tau_i/L_i$ to represent the angle of twist; then $(L_i, \theta_i) \in \mathbb{R}^2 - \{0\}$ provide 'polar coordinates' for $\text{Teich}(S)/\mathbb{Z}^{3g-3}$, where we quotient by the subgroup of Dehn twists fixing $\mathcal{P}$. That is, we have a surjective map

$$\mathbb{H}^{3g-3}/\mathbb{Z}^{3g-3} \cong (\mathbb{R}^2 - \{0\})^{3g-3} \to \mathcal{M}_g.$$

Now letting $L_i(X)$ attain the value zero, we obtain a partial compactification $\overline{U}$ of U, corresponding to adding the origin to $\mathbb{R}^2$. The limiting points can be described as 'surfaces with nodes'. Indeed such a surface can almost be marked: that is, it can be marked up to Dehn twists about the nodes.

By Mumford's Theorem, any sequence of surfaces in $\mathcal{M}_g$ has a limit in $\overline{\mathcal{M}_g}$, and thus $\overline{\mathcal{M}_g}$ is compact. More precisely, we can always choose our pants decomposition so that the lengths L_i are bounded; since L_i plays the role of a radial coordinate on $\mathbb{R}^2$, we obtain compactness.

The space $\overline{\mathcal{M}_g}$ is the *Deligne-Mumford compactification* of moduli space.

23. *Divisors at infinity.* There are $[g/2] + 1$ divisors at infinity, which we can label as D_h, where $h = 0, 1, \ldots, [g/2]$. A generic stable curve in D_0 is connected; that is, a non-separating curve has been pinched. In D_h, $h > 0$, a generic curve has two components, of genus h and $g - h$ respectively.

3.3 Simple closed curves, measured foliations and laminations, and Thurston's boundary for $\text{Teich}(S)$

24. The number $N(L)$ of distinct *free* homotopy classes of closed loops on X of length L satisfies $N(L) \asymp e^L/L$.

    To see this, note that the pointed homotopy classes grow *exactly* like the area of a ball of radius R. Conjugacy reduces the number by a factor of $1/L$.

25. Now on a *punctured torus*, the number of simple closed curves of length L is comparable to L^2. This is because a simple curve is determined by its slope p/q, with $|p| + |q| \asymp L$.

26. On the other hand, the number of *simple* closed curves of length L satisfies $N_s(L) \asymp L^{6g-6}$.

    To see this, note that a simple closed curve is determined in each pair of pants by its intersection count with the cuffs.

    Give this data, one needs some additional (linear) twist parameters to describe the simple closed curve completely. The upshot is that the number of simple closed curves of length L is $O(L^{6g-6})$.

27. Corollary (Birman-Series): the closed set of all simple geodesics on X (or in T_1X) has Hausdorff dimension one.

28. *Train tracks.* The modern viewpoint realizes a simple closed curve by weights on a trivalent train track $\tau \subset S$ with $S - \tau$ consisting of cusped triangles.

    The number of complementary regions satisfies $F = 4g - 4$ (think of them as ideal triangles and apply Gauss-Bonnet). Thus

    $$2g - 2 = -\chi(S) = E - V - F = 4g - 4 - \chi(\tau).$$

    Therefore $\chi(\tau) = E - V = 6g - 6$. It turns out that the relations imposed by vertices are independent and thus the space of admissible weights has dimension $6g - 6$.

    A finite number of train tracks suffice to carry every simple closed curve, and thus $N_s(L) \asymp 6g - 6$.

3.4 The Nielsen realization problem

29. Let $G \subset \text{Mod}(S)$ be a finite group. The following statements are equivalent.

    (a) G lifts to $\text{Diff}(S)$.

    (b) G can be realized as a subgroup of the isometries of S for some Riemannian metric g on S.

    (c) G can be realized in $\text{Aut}(X)$ for some $X \in \text{Teich}(S)$.

    (d) G has a fixed-point on $\text{Teich}(S)$.

    The *Nielsen problem* is to prove that these assertions hold for *any* finite subgroup of $\text{Mod}(S)$, where S is a compact orientable surface.

30. *Groups of prime order.* If $G = \mathbb{Z}/p$ then G has no fixed-point free action on a simply-connected manifold M. Otherwise M/G would be a finite-dimensional $K(\mathbb{Z}/p, 1)$. But one knows $\mathbb{Z}/p$ has infinite cohomological dimension (as can be computed from the infinite-dimensional lens space $L_p^\infty = S^\infty/(\mathbb{Z}/p)$).

31. *Lengths and proper functions.* To begin the proof, we wish to find a collection of closed loops $\gamma_1, \ldots, \gamma_n$ on S such that

    $$F(X) = \sum \ell_{\gamma_i}(X)$$

    defines a *proper* map $F : \text{Teich}(S) \to \mathbb{R}$.

    This is easily done by choosing enough γ_i such that their lengths determine X and such that every γ_i has positive geometric intersection number with some γ_j. The intersecting condition implies $F(X)$ is bounded below, away from zero, and the determining condition implies $F(X_n) \to \infty$ if $X_n \to \infty$ in $\text{Teich}(S)$.

32. *Binding.* In fact this properness holds whenever the γ_i *bind* the surface S — that is, whenever $\sum i(\gamma_i, \delta) > 0$ for every simple closed curve δ.

    Exercise: show any surface S of genus g can be bound by just two simple closed curves.

    Solution (Grisha): Make the first curve non-separating; then $T = S - \gamma_1$ is an annulus with $g-1$ handles attached. Now in each handle H_i of T, draw 3 disjoint arcs between the boundary components of T, cutting H_i into disks. (This can be done with 2 of the arcs parallel). The we just need to glue the boundary components of T back togetheer so these $3g-3$ arcs join up to form a single simple closed curve γ_2. This can be achieved by regluing with a cyclic permutation in the correct direction.

33. *Invariance.* Since G is a finite group, we can make the set of loops γ_i G-invariant. Then F is G-invariant. We will show F has a *unique minimum.* Then this minimum must be fixed by G and the Nielsen problem is solved.

34. *Convexity and earthquakes.* The last step in the proof is to show:

    $\ell_\gamma(X)$ *is strictly convex along earthquake paths.*

    Thus $F(X)$ is also strictly continuous, and by joining two distinct minima of F with an earthquake path we would obtain a contradition.

    To prove convexity we first note that by an easy geometric argument, if $\tau_{t\alpha}$ is a left twist of length t along a simple closed geodesic α, then

$$\frac{d}{dt}\ell_\gamma(\tau_{t\alpha}(X)) = \sum_{p\in\alpha\cap\gamma} \cos(\theta_p),$$

    where $\theta_p \in [0,\pi]$ is the angle of intersection. On the other hand, it is easy to see $d\theta_p/dt < 0$, and therefore

$$\frac{d^2}{dt^2}\ell_\gamma(\tau_{t\alpha}(X)) = \sum_{p\in\alpha\cap\gamma} -\sin(\theta_p)\frac{d\theta_p}{dt} > 0,$$

    yielding convexity.

    Using measures, the same proof works when α is a *measured lamination,* thus verifying convexity. ■

References

[Bus] P. Buser. *Geometry and Spectra of Compact Riemann Surfaces.* Birkhauser Boston, 1992.

[McS] C. McMullen and D. Sullivan. Quasiconformal homeomorphisms and dynamics III: The Teichmüller space of a holomorphic dynamical system. *Adv. Math.* **135**(1998), 351–395.