Complex Analysis on Riemann Surfaces

Math 213b — Harvard University

Contents

1 Introduction

Scope of relations to other fields.

1. Topology: genus, manifolds. Algebraic topology, intersection for on $H^1(X,\mathbb{Z})$, $\int \alpha \wedge \beta$.

2. 3-manifolds. (a) Knot theory of singularities. (b) Isometries of $\mathbb{H}^3$ and $\operatorname{Aut}\widehat{\mathbb{C}}$. (c) Deformations of M^3 and ∂M^3.

3. 4-manifolds. (M,ω) studied by introducing J and then pseudo-holomorphic curves.

4. Differential geometry: every Riemann surface carries a conformal metric of constant curvature. Einstein metrics, uniformization in higher dimensions. String theory.

5. Algebraic geometry: compact Riemann surfaces are the same as algebraic curves. Intrinsic point of view: $x^2+y^2=1$, $x=1$, $y^2=x^2(x+1)$ are all 'the same' curve. Moduli of curves. $\pi_1(\mathcal{M}_g)$ is the mapping class group.

6. Arithmetic geometry: Genus $g \geq 2$ implies $X(\mathbb{Q})$ is finite. Other extreme: solutions of polynomials; $\mathbb{C}$ is an algebraically closed field.

7. Complex geometry: Sheaf theory; several complex variables. The Jacobian $\mathrm{Jac}(X) \cong \mathbb{C}^g/\Lambda$ determines an element of $\mathcal{H}_g/\mathrm{Sp}_{2g}(\mathbb{Z})$: arithmetic quotients of bounded domains.

8. Dynamics: unimodal maps exceedingly rich, can be studied by complexification: Mandelbrot set, Feigenbaum constant, etc.

Examples of Riemann surfaces.

1. $\mathbb{C}$, $\mathbb{C}^*$, Δ, Δ^*, $\mathbb{H}$.

2. The space $\widehat{\mathbb{C}} \cong \mathbb{P}^1$. The automorphism groups of these surfaces.

3. Quotients X/Γ, $\Gamma \subset \mathrm{Aut}(X)$. Covering spaces. $\mathbb{C}/(z \mapsto z+1) \cong \mathbb{C}^*$ by $\exp(2\pi i z)$; $\mathbb{H}/(z \mapsto z+1) \cong \Delta^*$.

   Theorem 1.1 *$\mathbb{H}/(z \mapsto \lambda z)$ is isomorphic to $A(r) = \{z : r < |z| < 1\}$ where $r = \exp(-2\pi^2/|\log \lambda|)$.*

   Warning. $A : \mathbb{C}^* \to \mathbb{C}^*$ by $A(x+iy) = (2x + iy/2)$ does not have a Hausdorff quotient! The action of $\langle A \rangle$ is free but not properly discontinuous (consider the images under A^n of the circle $|z|=1$). There are points on the real and imaginary axes whose neighborhoods in the quotient always intersect.

   Theorem 1.2 *If Γ is a discrete group of isometries acting freely, then the action is properly discontinuous and X covers the quotient manifold X/Γ.*

4. The group $\Gamma(2) \subset \mathrm{SL}_2(\mathbb{Z})$; isometric for hyperbolic metric; quotient is triply-punctured sphere.

5. Branched covers. $p_n : \Delta \to \Delta$ by $p_n(z) = z^n$. Covers of the 4-times punctured sphere.

6. Algebraic curves. A curve $C \subset \mathbb{C}^2$ defined by $f(x,y) = 0$ is *smooth* $df \neq 0$ along C. Example: $f(x,y) = y^2 - p(x)$ is smooth iff p has no multiple roots. Example: $y^2 = x^3$, $y^2 = x^2(x+1)$.

7. Symmetric examples. Square with opposite sides identified yields a curve $E \cong \mathbb{C}/\mathbb{Z}[i]$ of genus one. The quotient $E/(z \mapsto -z)$ is a sphere with four distinguished points B. The symmetry $z \mapsto iz$ gives a symmetry of the configuration B, fixing two points and interchanging two others. Thus we can take $B = \{0, \infty, \pm 1\}$.

8. A regular Euclidean octagon with opposite sides identified has an 8-fold automorphism $r : X \to X$. The quotient $X/(r^4) \cong \widehat{\mathbb{C}}$ is branched over six points B. The action of r gives an order 4 symmetry of B that fixes 2 points and cyclically permutes the remaining 4. Thus X is defined by $y^2 = x(x^4 - 1)$.

   Note: if one is worried about the vertices of the octagon, one can take a regular hyperbolic octagon with interior angles of 45°.

2 Maps between Riemann surfaces

Let $f : X \to Y$ be a nonconstant map. Then f is locally modeled on $z \mapsto z^m$. We write $m(f,p)$ for the multiplicity of f at p. It follows that:

1. The map f is *open* and *discrete*.

2. f satisfies the maximum principle.

3. If X is compact, then $f(X) = Y$.

4. An analytic function on a compact Riemann surface is constant.

Removable singularities. An isolated singularity of a bounded analytic function is removable. Consequently:

1. A bounded analytic function on a Riemann surface of finite type (compact with a finite number of points removed) is constant.

2. A bounded analytic function on $\mathbb{C}$ is constant. (Proof: it extends to $\widehat{\mathbb{C}}$.)

3. Every polynomial of degree $d \geq 1$ has a root in $\mathbb{C}$.

Covering maps. Let $\pi : X \to Y$ be a covering space of a Riemann surface Y. There is a unique complex structure on X such that π is holomorphic. The space X is determined up to isomorphism over Y by the subgroup $H \cong \pi_1(X, p) \subset \pi_1(Y, q)$.

The deck group is defined by $\Gamma = \mathrm{Deck}(X/Y) \subset \mathrm{Aut}(X)$ is the group of automorphisms α such that $\pi \circ \alpha = \pi$. We say X/Y is *normal* (Galois, regular) if the deck group acts transitively on the fibers of π.

In general $\mathrm{Deck}(X/Y) = N(H)/H$ where $N(H)$ is normalizer of H in G. To see this think of Y as $\widetilde{Y}/G$: then a deck transformation lifts to a g in G satisfying $gh_1x = h_2gx$, so it descends to $X = \widetilde{Y}/H$. This requires $gh_1g^{-1} = h_2$ and thus $g \in N(H)$.

The key property of covering space is an algebraic solution to the lifting problem for $f : Z \to Y$.

Branched coverings. Now let $\pi : X \to Y$ be a general nonconstant map between Riemann surfaces. Let $C = \{x : \mathrm{mult}(f, x) > 1\}$, let $B = f(C)$ and let $\widetilde{B} = \pi^{-1}(B)$. Let $X^* = X - \widetilde{B}$ and $Y^* = Y - B$.

We say π is a *branched covering* if B is discrete, $\pi : X^* \to Y^*$ is a covering map, and for any small loop γ around a single point of B, every component of $\pi^{-1}(\gamma)$ is compact.

Theorem 2.1 *A branched covering X is uniquely determined, up to isomorphism over Y, by a discrete set B and a subgroup $H \subset G = \pi_1(Y, p)$. Any subgroup that meets each peripheral subgroup of G with finite index determines such a covering.*

As before a covering is *normal* (Galois, regular) if $\mathrm{Deck}(X/Y)$ acts transitivity on fibers; this is equivalent to X^*/Y^* being normal.

Theorem 2.2 *Any properly discontinuous subgroup $\Gamma \subset \mathrm{Aut}(X)$ yields a quotient Riemann surface Y and a branched covering map $\pi : X \to Y$ with* $\mathrm{Deck}(X/Y) = \Gamma$.

Proper maps. Let $f : X \to Y$ be a proper, nonconstant map between Riemann surfaces. That is, assume K compact implies $f^{-1}(K)$ compact. Then:

1. f is closed: i.e. E closed implies $f(E)$ closed. (This requires only local connectivity of the base Y.)

2. f is surjective. (Since it is also open).

3. If $D \subset X$ is discrete, so is $f(D)$. (Since $f(D)$ meets any compact set K in a finite set, namely the image of $D \cap f^{-1}(K)$.)

4. In particular, the branch locus $B(f) \subset Y$ is discrete.

5. $f^{-1}(q)$ is finite for all $q \in Y$. (Since f is discrete.)

6. For any neighborhood U of $f^{-1}(q)$ there exists a neighborhood V of q whose preimage is contained in U. (Since $f(X - U)$ is closed and does not contain q.)

7. If f is a proper local homeomorphism, then it is a covering map. (This only requires that Y is locally compact.)

   Proof. A local homeomorphism is discrete, so given $q \in Y$ we can choose neighborhoods U_i of its preimages $p_1, \ldots, p_n$ such that $f : U_i \to V_i$ is a homeomorphism. Let V be a neighborhood of q such that $f^{-1}(V)$ is contained in $\bigcup U_i$. Then f evenly covers V. ■

8. For any $q \in Y$ with $f^{-1}(q) = \{p_1, \ldots, p_n\}$, there exists a disk V containing q such that $f^{-1}(V) = \bigcup U_i$ with U_i a disk, $p_i \in U_i$ and $f_i = f|U_i$ satisfies $f_i : (U_i, p_i) \to (V, q)$ is conjugate to $z \mapsto z^{d_i}$ on Δ.

   (Here one can use the Riemann mapping theorem to prove that if $g : \Delta \to \Delta$ is given by $g(z) = z^m$, and $V \subset \Delta$ is a simply-connected neighborhood of $z = 0$, then the pullback of g to V is also conjugate to $z \mapsto z^m$.)

9. The function $\sum_{f(p)=q} m(f, p)$ is independent of q. It is called the *degree* of f.

Properness and branched covers.

Theorem 2.3 *Any proper map is a branched covering. A branched covering is proper iff it has finite degree, i.e. the cardinality of one (and hence every) fiber is finite, i.e.* $X^* \to Y^*$ *has finite degree.*

Proof. For the first statement use the fact that a proper map has a degree set of branch values, and a proper local homeomorphism is a covering. For the second, use the fact that a finite covering map is proper. ■

Extension of proper maps. Let Y^* be the complement of a discrete set D in a Riemann surface Y, and let $f : X^* \to Y^*$ be a proper map. Then there is a unique way to complete X and f to obtain a proper map $f : X \to Y$.

Examples.

1. Any rational map $f : \widehat{\mathbb{C}} \to \widehat{\mathbb{C}}$ is a proper branched covering; it is a covering map iff it is a Möbius transformation. An entire function is proper iff it is a polynomial.

2. Any proper map of the disk to itself is given by a Blaschke product,
$$f(z) = \exp(i\theta) \prod_1^d \frac{z - a_i}{1 - \overline{a}_i z}.$$

3. The map $f : \widehat{\mathbb{C}} \to \widehat{\mathbb{C}}$ given by $f(z) = z^d$ is a degree d, proper, regular branched cover with $B = \{0, \infty\}$ and $\operatorname{Deck}(X/Y) = \mathbb{Z}/d$. Its restriction to a map $\mathbb{C}^* \to \mathbb{C}^*$ is a covering map. (As a map $\mathbb{C}^* \to \mathbb{C}$, it is not a covering map.)

4. The map $f(z) = z^d + 1/z^d$ is a regular branched cover, with deck group the dihedral group generated by $z \mapsto 1/z$ and $z \mapsto \exp(2\pi i/d)z$.

5. A basic example of an irregular cover is given by $f(z) = z^3 - 3z$. This map has branched points at ± 1 and branch values $B = \{\pm 2\}$; we have $\widetilde{B} = f^{-1}(B) = \{\pm 1, \pm 2\}$, and $f : \mathbb{C} - \widetilde{B} \to \mathbb{C} - B$ is an irregular degree three cover. It corresponds to the 'abbaab' triple cover of a bouquet of circles.

6. The map $f(z) = z^3 : \mathbb{H} \to \mathbb{C}$ is not proper, even though it is a local homeomorphism. Its extension to $\overline{\mathbb{H}}$ *is* proper, but not a local homeomorphism.

7. The map $f(z) = e^z : \mathbb{C} \to \mathbb{C}^*$ is a regular covering map, but not proper. Its deck group $\mathbb{Z}$ is generated by $z \mapsto z + 2\pi i$.

8. The map $f(z) = \tan(z) : \mathbb{C} \to f(\mathbb{C}) \subset \widehat{\mathbb{C}}$ is a covering map. In fact $\tan(z) = g(e^{2iz})$ where $g(z) = -i(z-1)/(z+1)$. Thus $f(\mathbb{C}) = \widehat{\mathbb{C}} - \{\pm i\}$.

9. The map $f(z) = \cos(z) : \mathbb{C} \to \mathbb{C}$ is a regular branched covering. Its branch values are $B = \{\pm 1\}$, and its branch points are $\widetilde{B} = \pi\mathbb{Z}$. The deck group is the *infinite* dihedral group generated by $z \mapsto -z$ and $z \mapsto z + 2\pi$. Note that $\cos(z) = g(e^{iz})$ where $g(z) = (z + 1/z)/2$.

The Riemann-Hurwitz formula. A Riemann surface X is of *finite type* if it is obtained from a compact Riemann surface by deleting a finite number of points. In this case $\chi(X) = 2 - 2g - n$ where g is the genus and n is the number of points removed.

Theorem 2.4 *If $f : X \to Y$ is a branched covering map between Riemann surfaces of finite type, then*

$$\chi(X) = d\chi(Y) - \sum_X (\mathrm{mult}(f, x) - 1).$$

Since $\chi(X) = 2 - 2g(X)$, this formula also relates the genera of X and Y.

Example: Let $f : \widehat{\mathbb{C}} \to \widehat{\mathbb{C}}$ be a map of degree d. The f has $2d-2$ critical points.

Example: The only compact Riemann surfaces admitting self maps of degree $d > 1$ are $\widehat{\mathbb{C}}$ and $\mathbb{C}/\Lambda$. The latter only admit self covering maps.

(Second proof: if $f : \mathbb{C} \to \mathbb{C}$ is a lift of a self-map of $\mathbb{C}/\Lambda$, then $f'(z)$ is a holomorphic, doubly-periodic function, hence constant.)

Hurwitz problem. Here is an unsolved problem. Let $f : \widehat{\mathbb{C}} \to \widehat{\mathbb{C}}$ be a rational map of degree d with critical values $B = (b_i)_1^n$. Let p_i be the partition of d corresponding to the fiber over b_i. What partitions $(p_1, \ldots, p_n)$ can be so realized?

This is really a problem in topology or group theory. We have to lift each p_i to an element $g_i \in S_d$ in the conjugacy class specified by p_i, in such a way that $g_1 \cdots g_n = e$.

Belyi's Theorem.

Theorem 2.5 *A compact Riemann surface X is defined over a number field iff it can be presented as a branched cover of $\widehat{\mathbb{C}}$, branched over just 3 points.*

We first need to explain what it means for X to be defined over a number field. For our purposes this means there exists a branched covering map $f : X \to \widehat{\mathbb{C}}$ with $B(f)$ consisting of algebraic numbers. Alternatively, X can be described as the completion of a curve in $\mathbb{C}^2$ defined by an equation $f(x, y) = 0$ with algebraic coefficients.

Grothendieck wrote that this was the most striking theorem he had heard since at age 10, in a concentration camp, he learned the definition of a circle as the locus of points equidistant from a given center.

Proof. I. Suppose X is defined over a number field, and $f : X \to \widehat{\mathbb{C}}$ is given with $B(f)$ algebraic. We will show there exists a polynomial $p : \widehat{\mathbb{C}} \to \widehat{\mathbb{C}}$ such that $B(p) \subset \{0, 1, \infty\}$ and $p(B(f)) \subset \{0, 1, \infty\}$. This suffices, since

$$B(p \circ f) = B(p) \cup p(B(f)).$$

We may assume $\infty \notin B(f)$. Let $\deg(B)$ denote the maximal degree over $\mathbb{Q}$ of the points of B, and let $z \in B(f)$ have degree $\deg(B)$. Then there exists a polynomial $p \in \mathbb{Q}[z]$ of degree d such that $p(z) = 0$. Moreover $\deg(B(p)) \leq d - 1$ since the critical values of B are the images of the zeros of $p'(z)$. Thus

$$B' = B(p) \cup p(B)$$

has fewer points of degree $\deg(B)$ over $\mathbb{Q}$. Iterating this process, we can reduce to the case where $B \subset \mathbb{Q}$.

Now comes a second beautiful trick. Consider the polynomial $p(z) = Cz^a(1-z)^b$. This polynomial has critical points at 0, 1 and $w = a/(a+b)$. By choosing the value of $C \in \mathbb{Q}$ correctly, we can arrange that $p(w) = 1$ and hence $B(p) \subset \{0, 1, \infty\}$. Thus if $w \in B$ is rational, we can choose p so that $B' = p(B) \cup B(p)$ has fewer points outside $\{0, 1, \infty\}$. Thus we can eventually eliminate all such points.

II. Let X be a Riemann surface presented as a branched covering $f : X \to \widehat{\mathbb{C}}$ with $B(f) = \{0, 1, \infty\}$. We need a nontrivial factor: there exists a second $g \in \mathcal{M}(X)$ and a cofinite set $X^* \subset X$ such that $(f, g) : X^* \to \mathbb{C}^2$ is an immersion with image the zero locus V_p of a polynomial $p(x, y) \in \mathbb{C}[x, y]$.

This polynomial has the property that the projection of (the normalization of) V_p to the first coordinate — which is just f — is branched over just 0, 1 and ∞. The set of all such polynomials of given degree is an algebraic subset W of the space of coefficients $\mathbb{C}^N$ for some large N. This locus W is defined by rational equations. Thus the component of W containing our given p also contains a polynomial q with coefficients in a number field. But V_q and V_p are isomorphic, since they have the same branch locus and the same covering data under projection to the first coordinate. Thus $X^* \cong V_p \cong V_q$ is defined over a number field. ∎

Corollary 2.6 *X is defined over a number field iff X can be built by gluing together finitely many unit equilateral triangles.*

Proof. If X can be built in this way, one can 2-color the barycentric subdivision and hence present X as a branched cover of the double of a

30-60-90 triangle. The converse is clear, since the sphere can be regarded as the double of an equilateral triangle.

Example: A square torus can be built by gluing together 8 isosceles right triangles. But these can also be taken to be equilateral triangles, with the same result! This is because the double of *any* two triangles is the same, as a Riemann surface.

3 Sheaves and analytic continuation

Presheaves and sheaves. A *presheaf of abelian groups* on X is a functor $\mathcal{F}(U)$ from the category of open sets in X, with inclusions, to the category of abelian groups, with homomorphisms.

It is a *sheaf* if (I) elements $f \in \mathcal{F}(U)$ are determined by their restrictions to an open cover U_i, and (II) any collection $f_i \in \mathcal{F}(U_i)$ with $f_i = f_j$ on U_{ij} for all i, j comes from an $f \in \mathcal{F}(U)$.

Example: applying (I) to the empty cover of the empty set, we see $\mathcal{F}(\emptyset) = (0)$. Thus the presheaf that assigns a fixed, nontrivial group G to every open set is not a sheaf.

Examples of presheaves and sheaves. The sheaves $\mathcal{C}$, $\mathcal{C}^\infty$, $\mathcal{O}$ and $\mathcal{M}$, of rings of continuous, smooth, holomorphic and meromorphic functions. The multiplicative group sheaves $\mathcal{O}^*$ and $\mathcal{M}^*$.

If G is a nontrivial abelian group, $\mathcal{F}(U) = G$ for U nonempty and $\mathcal{F}(\emptyset) = (0)$ is a presheaf. But it is not a sheaf: $\mathcal{F}(U_1 \sqcup U_2) \neq G \oplus G$.

To rectify this we can define $\mathcal{F}(U)$ to be the additive group of locally constant maps $f : U \to G$. This is now a sheaf; it is often denoted simply by G. (E.g. $\mathbb{R}$, $\mathbb{C}$, $\mathbb{Q}$, S^1, $\mathbb{Z}$.)

The presheaf $\mathcal{F}(U) = \mathcal{O}(U)/\mathbb{C}(U)$ is not a sheaf. For example, local logarithms do not assemble.

Stalks. Let $\mathcal{F}$ be a *presheaf*.

The *stalk* $\mathcal{F}_x$ is the direct limit of $\mathcal{F}(U)$ over the (directed) system of open sets containing $x \in X$. It can be described directly as the disjoint union of these groups modulo $f_1 \sim f_2$ if they have a common restriction near x. (Alternatively, $\mathcal{F}_x = (\oplus \mathcal{F}(U))/N$ where N is generated by elements of the form $(f|U_1) - (f|U_2)$, $f \in \mathcal{F}(U_1 \cup U_2)$.)

There is a natural map $\mathcal{F}(U) \to \mathcal{F}_x$ for any neighborhood U of x. We let f_x denote the image of $f \in \mathcal{F}(U)$ under this map. We have $f_x = 0$ iff there is a neighborhood V of x such that $f|V = 0$.

Example: $\mathcal{O}_a \cong \mathbb{C}\{\{z_a\}\}$ for any local chart (uniformizer) $z_a : U \to \mathbb{C}$, $z_a(a) = 0$.

Theorem 3.1 *Let $\mathcal{F}$ be a sheaf. Then $f \in \mathcal{F}(U)$ is zero iff $f_x = 0$ for all $x \in U$.*

Éspace étalé. The *éspace étalé* $|\mathcal{F}|$ of a presheaf is the disjoint union of the stalks $\mathcal{F}_x$, with a base for the topology given by sets of the form $[U, f] = \{f_x \ : \ x \in U\}$. It comes equipped with a natural projection $p : |\mathcal{F}| \to X$ which is a local homeomorphism.

We say $\mathcal{F}$ satisfies the *identity theorem* if whenever U is open and connected, $f, g \in \mathcal{F}(U)$ and $f_x = g_x$ for some $x \in U$, then $f = g$. Examples: $\mathcal{O}$ and $\mathcal{M}$.

Theorem 3.2 *$|\mathcal{F}|$ is Hausdorff if $\mathcal{F}$ satisfies the identity theorem.*

Structure of $|\mathcal{O}|$. There is a unique complex structure on $|\mathcal{O}|$ such that $p : |\mathcal{O}| \to X$ is an *analytic* local homeomorphism.

There is also a natural map $F : |\mathcal{O}| \to \mathbb{C}$ given by $F(f_x) = f(x)$. This map is analytic. However $|\mathcal{O}|$ is never connected, since the stalks $\mathcal{O}_x$ are uncountable.

For any holomorphic map $\pi : X \to Y$, there is a natural map $\pi^* : |\mathcal{O}_Y| \to |\mathcal{O}^X|$ compatible with projections to the base spaces X and Y. If $\pi(a) = b$ then $\pi^*(f_b) = f_b \circ \pi \in \mathcal{O}_a$.

There is also a pushforward map π_* when π is a local homeomorphism.

Path lifting. We now recall some results from the theory of covering spaces. Let $p : X \to Y$ be a local homeomorphism between *Hausdorff* spaces. Let $f : I = [0, 1] \to Y$ be a path, and $x \in X$ a point such that $p(x) = f(0)$. A *lifting* of f based at x is a path $F : I \to X$ such that $f = p \circ F$ and $F(0) = x$.

Theorem 3.3 *A lifting is unique if it exists.*

Proof. If we have two liftings, F_1 and F_2, the set of $t \in I$ such that $F_1(t) = F_2(t)$ is open since p is a local homeomorphism, closed since X is Hausdorff, and nonempty since $F_1(0) = F_2(0) = x$. By connectedness it is the whole interval. ■

Now let $f_s(t)$ be a homotopy of paths parameterized by $s \in [0, 1]$, such that $f_s(0) = y_0$ and $f_s(1) = y_1$ are constant. Suppose for every s there is a lift F_s of f_s based at x_0. We then have:

Theorem 3.4 (Monodromy Theorem) *The terminus $F_s(1) = x_1$ is independent of s. Moreover $F_s(t)$ is a lift of $f_s(t)$ as a function on $I \times I$. In particular, $F_0(t)$ and $F_1(t)$ are homotopic* rel *their endpoints.*

Non-unique path lifting. Let $\mathcal{F}$ be the space of continuous functions on $\mathbb{R}$. Consider the functions $f(x) = 0$ and $g(x) = \max(0, x)$. The germs of these functions determine a subspace Y of $|\mathcal{F}|$. The space $p : Y \to \mathbb{R}$ has 1 points over $(-\infty, 0)$ and two points over $[0, \infty)$. In particular the two solutions to $p(y) = 0$ cannot be separated by disjoint open sets, since these sets always meet over the negative real axis.

In this case $f_a = g_a$ for any $a < 0$, and the path $[a, 1] \to \mathbb{R}$ can be lifted to $|\mathcal{F}|$ in two different ways, both starting at this point.

Analytic continuation. Let $\mathcal{O}$ be the sheaf of analytic functions on $X = \mathbb{C}$. Let $\gamma : [0, 1] \to \mathbb{C}$ be path with $a = \gamma(0)$ and $b = \gamma(1)$. Let f_a be the germ of an analytic function at $z = a$, i.e. let $f \in \mathcal{O}_a$.

We say $f_b \in \mathcal{O}_b$ is obtained from f_a by *analytic continuation along* γ if there are analytic functions g_i on balls B_i containing $\gamma(t_i)$, $i = 0, \ldots, n$ such that:

1. $0 = t_0 < t_1 < \cdots < t_n = 1$;
2. B_i meets B_{i+1}, g_i and g_{i+1} agree there, and $\gamma[t_i, t_{i+1}] \subset B_i \cap B_{i+1}$; and
3. $g_0 = f_a$ and $g_1 = f_b$.

Example: For $a = 0$, $f_a(z) = \sum z^n/n = \log(1/(1-z))$ can be analytically continued to any point in $\mathbb{C}^*$, in many ways; the various possibilities for $f_b(z)$ differ by multiplies of $2\pi i$.

Observation: analytic continuation along γ is equivalent to path-lifting to the space $p : |\mathcal{O}| \to \mathbb{C}$.

Corollary 3.5 *An analytic continuation is unique if it exists.*

Corollary 3.6 (Monodromy theorem) *If analytic continuation from a to b is possible along a family of paths γ_s, $s \in I$, then the result f_b is always the same.*

Maximal analytic continuation (spreads). If $f : X \to Y$ is holomorphic and $f(a) = b$ then there is a natural pullback map $f^* : \mathcal{O}_b \to \mathcal{O}_a$ given by composition with f. If $\operatorname{mult}(f, a) = 1$ then f^* is an isomorphism and hence there is also a pushforward $f_* : \mathcal{O}_a \to \mathcal{O}_b$. (Remark: by summing over sheets, pushforward can also be defined at general points.)

To keep track of the potential multi-valuedness, we define an *analytic continuation* of $f_a \in \mathcal{O}_a$, $a \in \mathbb{C}$, to be a pointed Riemann surface (Y, b) endowed with a local analytic homeomorphism $p : (Y, b) \to (\mathbb{C}, a)$ and an $F \in \mathcal{O}(Y)$ such that $p_* F_b = F_a$.

Theorem 3.7 *Every germ of an analytic function f_a has a unique maximal analytic continuation, obtained by taking Y to be the connected component of $|\mathcal{O}|$ containing f_a, and restricting the natural maps $F : |\mathcal{O}| \to \mathbb{C}$ and $p : |\mathcal{O}| \to \mathbb{C}$ to Y.*

Proof. By map $y \mapsto p_*(F_y)$ gives an embedding of any analytic continuation into $|\mathcal{O}|$, which is then tautologically dominated by the maximal one described above. ∎

Examples.

1. The maximal domain of $f(z) = \sum z^{n!}$ is the unit disk.
2. For $\log(1/(1-z)) = \sum z^n/n$, the maximal analytic continuation is the Riemann surface $\mathbb{C}$ with $F : \mathbb{C} \to \mathbb{C}$ the identity map, and $p : \mathbb{C} \to \mathbb{C} - \{1\}$ the map $p(z) = 1 - e^{-z}$.
3. Let $f(z) = \sqrt{q(z)}$ where $q(z) = 4z^3 + az + b$. Then its maximal analytic continuation is given by $Y = E - E[2]$ with $E = \mathbb{C}/\Lambda$, with $p : Y \to \mathbb{C}$ given $p(z) = \wp(z)$, and with $F(z) = \wp'(z)$.

Remark. One can replace the base $\mathbb{C}$ of analytic continuation with any other Riemann surface X. Note however that the 'maximal' analytic continuation may become larger under an inclusion $X \hookrightarrow X'$.

4 Algebraic functions

We will develop two main results.

Theorem 4.1 *Let $\pi : X \to Y$ be a holomorphic branched covering of degree d. Then $\mathcal{M}(X)/\mathcal{M}(Y)$ is an algebraic field extension, of degree at most d.*

In fact the degree is exactly d, but to see this we need to know that $\mathcal{M}(X)$ separates the points of X.

Theorem 4.2 *Let $K/\mathcal{M}(Y)$ be a field extension of degree d. Then there is a unique degree d branched covering $\pi : X \to Y$ such that $K \cong \mathcal{M}(X)$ over $\mathcal{M}(Y)$.*

Moreover $\mathcal{M}(X)/\mathcal{M}(Y)$ is Galois iff X/Y is Galois, in which case there is a natural isomorphism

$$\mathrm{Gal}(\mathcal{M}(X)/\mathcal{M}(Y)) \cong \mathrm{Deck}(X/Y).$$

Putting these results together, we find for example that $X \mapsto \mathcal{M}(X)$ establishes an equivalence between (i) the category of finite-sheeted branched coverings of $\widehat{\mathbb{C}}$, and (ii) the category of finite field extensions of $\mathbb{C}(x)$.

Symmetric functions. The proof of the first result is based on the idea of *symmetric functions.* Recall that the 'elementary symmetric functions' s_i of $f_1, \ldots, f_d$ are related to the coefficients $c_i \in \mathbb{Z}[f_1, \ldots, f_d]$ of the polynomial

$$P(T) = \prod_1^d (T - f_i) = T^d + c_1 T^{d-1} + \cdots + c_d \tag{4.1}$$

by $s_i = (-1)^i c_i$. (Thus $s_1 = \sum f_i$, $s_2 = \sum_{i<j} f_i f_j$ and $s_d = \prod f_i$.) Clearly these polynomials s_i lie in $\mathbb{Z}[f_i]^{S_d}$ and in fact they generate the ring of invariant polynomials [La, §IV.6].

Now let $\pi : X^* \to Y^*$ be the unbranched part of the degree d branched covering $\pi : X \to Y$, and let $f \in \mathcal{M}(Y)$. By deleting more points if necessary, we can assume $f \in \mathcal{O}(Y^*)$.

Given $y \in Y^*$, let $\pi^{-1}(y) = \{x_1, \ldots, x_d\}$ and let $c_i(y)$ denote the coefficients of the polynomial with roots $f(x_i)_1^d$. Locally we can express $c_i(y)$ as the symmetric functions of the pullbacks f_i of f under the d branches of f^{-1}. Thus $c_i \in \mathcal{O}(Y^*)$, and these coefficients define a polynomial $P \in \mathcal{O}(Y^*)[T]$ such that $P(f) = 0$. (Here we have identified functions on Y^* with their pullbacks to X^*.)

We claim the coefficients c_i extends to $\mathcal{M}(Y)$. Suppose $b \in Y - Y^*$. If f is bounded on the fiber over b, then this follows by Riemann's removable singularities theorem. Otherwise, we can choose a local coordinate z_b vanishing at b and a $k > 0$ such that $z_b^k f$ is analytic over b, and thus the coefficients c_i' of $\prod(T - z_b^k f_i)$ are analytic at b. But c_i' and c_i differ only by a power of z_b, so c_i' is meromorphic at b.

By continuity, the function f is still a zero of the extended polynomial $P(T) \in \mathcal{M}(Y)[T]$. Thus f has degree at most n over $\mathcal{M}(Y)$. By the theorem of the primitive element, $\deg(\mathcal{M}(X)/\mathcal{M}(Y)) \leq n$. ■

Local algebraic functions. Now let us go backwards from a polynomial $P(T)$ of degree d to analytic functions $(f_i)_1^d$. We first work locally.

Let $P(T)$ be a monic polynomial of degree d with coefficients $c_i(z)$ in the local ring $\mathcal{O}_a$, $a \in Y$.

Theorem 4.3 *If $P(T, a) = T^d + c_1(a)T^{d-1} + \cdots + c_d(a)$ has simple zeros, then there exist $f_i \in \mathcal{O}_a$ such that $P(T) = \prod(T - f_i)$.*

Proof. This is simply the statement that the zeros of $P(T, z)$ vary holomorphically as its coefficients do. To make this precise, let $w_1, \ldots, w_d$ be the zeros of $P(T, a)$, and let $\gamma_1, \ldots, \gamma_d$ be the boundaries of disjoint disks in $\mathbb{C}$, centered at these zeros. We can then take:

$$f_i(z) = \frac{1}{2\pi i} \int_{\gamma_i} \frac{\zeta P'(\zeta, z)\, d\zeta}{P(\zeta, z)}.$$

Here $f_i(z)$ is analytic in z so long as $P(\zeta, z)$ never vanishes on γ_i. This is true for $P(\zeta, a)$, so by continuity it remains true in a small neighborhood of $z = a$. ∎

Resultant and discriminant. To go further, it is useful to recall the *resultant* $R(f, g)$.

Let K be a field, and let $f, g \in K[T]$ be nonzero polynomials with $\deg(f) = d$ and $\deg(g) = e$. Recall that $K[T]$ is a PID and hence a UFD. We wish to determine if f and g have a common factor, say h, of degree 1 or more. In this case $f = hf_1$, $g = hg_1$ and hence $g_1 f - f_1 g = 0$. Conversely, if we can find nonzero r and s with $\deg(r) < e$ and $\deg(s) < d$ such that $rf + sg = 0$, then f and g have a common factor.

The existence of such r and s is the same as a linear relation among the elements $(f, xf, \ldots, x^{e-1}f, g, xg, \ldots, x^{d-1}g)$, and hence it can be written as a determinant $R(f, g)$ which is simply a polynomial in the coefficients of f and g. We have $R(f, g) = 0$ iff f and g have a common factor.

The *discriminant* $D(f) = R(f, f')$ is nonzero iff f has simple zeros.

From a field extension to a branched cover. Now let K be a degree d field extension of $\mathcal{M}(Y)$. By the theorem of the primitive element, there exists a monic irreducible polynomial $P(T)$ with coefficients $c_i \in \mathcal{M}(Y)$ such that $K \cong \mathcal{M}(Y)[T]/(P(T))$.

By irreducibility, the discriminant $D(P) \in \mathcal{M}(Y)$ is not identically zero. Let $Y^* \subset Y$ be the complement of the zeros and poles of $D(P)$, and of the poles of the coefficients c_i.

Let $X^* \subset |\mathcal{O}_{Y^*}|$ be the set of germs f_a such that $P(f_a, a) = 0$. By the preceding local result, X^* is a degree d covering space of Y^*. The tautological map $f : X^* \to \mathbb{C}$ sending f_a to $f_a(a)$ is analytic. We can complete X^* to a branched covering $\pi : X \to Y$, and extend f to a meromorphic function on X using Riemann's removable singularities theorem.

By construction, $P(f) = 0$.

We must show that X is connected. If not, one of its connected components X_0/Y has degree $d_0 < d$. But then $f|X_0$ satisfies a polynomial

$Q \in \mathcal{M}(Y)[T]$ of degree $d_0 < d$, which is a factor of P. This contradicts irreducibility of P.

It remains to consider the Galois group. Suppose X/Y is Galois as a branched covering. Since $P_a(T)$ has distinct zeros for some $a \in Y$, the group $\mathrm{Deck}(X/Y)$ maps injectively into $\mathrm{Gal}(\mathcal{M}(X)/\mathcal{M}(Y))$; indeed, only the identity stabilizes the function f. By degree considerations, this map is surjective as well. Similarly, if $\mathcal{M}(X)/\mathcal{M}(Y)$ is Galois, then $G = \mathrm{Gal}(\mathcal{M}(X)/\mathcal{M}(Y))$ permutes the roots of the polynomial $P(T)$. These correspond to the sheets of X, so we get a map $G \to \mathrm{Deck}(X/Y)$ which is an isomorphism again by degree considerations. ■

The Riemann surface X can be regarded as a completion of the maximal analytic continuation of f_a, for any germ $f_a \in \mathcal{O}_a(Y)$ satisfying $P(f_a, a) = 0$ at a point where the discriminant of $P(T)$ is not zero.

Puiseux series.

> *Avec les series de Puiseux,*
> *Je marche comme sur des oeufs.*
> *Il s'ensuit que je les fuis*
> *Comme un poltron que je suis.*
> —A. Douady, 1996

(Variante: *Et me refugie dans la nuit.*)

Let $K = \mathcal{M}_p$ be the local field of a point p on a Riemann surface; it is isomorphic to the field of convergent Laurent series $\sum_{-n}^{\infty} a_i z^i$ in a local parameter $z \in \mathcal{O}_p$.

Theorem 4.4 *Every algebraic extension of K of degree d is of the form $L = K[\zeta]$, where $\zeta^d = z$.*

Proof. Let $P(T) \in K[T]$ be the degree d irreducible polynomial for a primitive element $f \in L$ (so $L = K(f)$).) Let $Y = \{|z| < r\}$ be a small neighborhood of P. Then we can find an $r > 0$ such that the coefficients c_i of $P(T)$ are well-defined on Y, only have poles at $z = 0$, and the discriminant $D(P)|Y$ vanishes only at $z = 0$ as well.

Note that $P(T)$ remains irreducible as a polynomial in $\mathcal{M}(Y)[T]$. Thus it defines a degree d branched covering $X \to Y$, branched only over the origin $z = 0$. But there is also an obvious branched covering $Y_d \to Y$ given by $\zeta \mapsto \zeta^d = z$. Since $\pi_1(Y^*) \cong \mathbb{Z}$ has a unique subgroup of index d, we have $Y_d \cong X$ over Y. ■

Corollary 4.5 *Any irreducible, degree d polynomial $P(T) = 0$ with coefficients in $\mathcal{M}_p$ has a solution of the form*

$$f(z) = \sum_{-n}^{\infty} a_i z^{i/d}.$$

In particular, any polynomial is locally 'solvable by radicals', i.e. its roots can be expressed in the form $f_i(z) = z^{1/d} g_i(z)$, where $g_i(z)$ is analytic. Equivalent, after the change of variables $z = \zeta^d$, the polynomial $P(T)$ factors into linear terms.

Example. For $P(T) = T^3 + T^2 - z = 0$, we have a solution

$$f(z) = z^{1/2} - \frac{z}{2} + \frac{5z^{3/2}}{8} - z^2 + \frac{231 z^{5/2}}{128} - \cdots$$

Why is this only of degree 1/2? Because the original polynomial has two distinct roots when $z = 0$: it is reducible over $\mathcal{M}_0$!

Construction as a curve. An alternative to the construction of $\pi : X \to Y$ is the following. Fix $P(T) \in \mathcal{M}(Y)[T]$, irreducible, and as before let Y^* be the locus where the discriminant is nonzero and where the coefficients of $P(T)$ are holomorphic. Then define

$$X^* = \{(x, y) \ : \ P_y(x) = 0\} \subset \mathbb{C} \times Y^*.$$

Let (F, π) be projections of X^* to its two coordinates. Then $\pi : X^* \to Y^*$ is a covering map, and $F : X^* \to \mathbb{C}$ is an analytic function satisfying $P(F) = 0$; and the remainder of the construction carries through as before.

Examples. For any polynomial $p(z) \in \mathbb{C}[z]$ which is not a square (i.e. which at least one root of odd order), we can form the Riemann surface $\pi : X \to \widehat{\mathbb{C}}$ corresponding to adjoining $f(z) = \sqrt{p(z)}$ to $K = \mathbb{C}(z) = \mathcal{M}(\widehat{\mathbb{C}})$. When p has $2n$ or $2n-1$ simple zeros, the map π is branched over $2n$ points (including infinity in the latter case), and hence X has genus $n-1$. Note that the polynomial $P(T) = T^2 - p(z)$ has discriminant $\mathrm{Res}(P, P') = -4p(z)$.

Valuations on $\widehat{\mathbb{C}}$. A *discrete valuation* on a field K is a surjective homomorphism $v : K^* \to \mathbb{Z}$ such that $v(f + g) \geq \min v(f), v(g)$.

Example: for $K = \mathbb{C}(z) = K(\widehat{\mathbb{C}})$, the order of zero (or pole) a rational function at a given point $p \in \widehat{\mathbb{C}}$ gives a *point valuation* $v_p(f)$.

Theorem 4.6 *Every valuation on $K(\widehat{\mathbb{C}})$ is a point valuation.*

Riemann surfaces	**Number fields**
$K = \mathcal{M}(\widehat{\mathbb{C}}) = \mathbb{C}(z)$	$K = \mathbb{Q}$
$A = \mathbb{C}[z]$	$A = \mathbb{Z}$
$p \in \mathbb{C}$	$p\mathbb{Z}$ prime ideal
uniformizer $z_p = (z - p) \in \mathbb{C}[z]$	uniformizer $p \in \mathbb{Z}$
order of vanishing of $f(z)$ at p	power of p dividing $n \in \mathbb{Z}$
$A_p = \mathcal{O}_p = \{\sum_0^\infty a_n z_p^n\}$	$A_p = \mathbb{Z}_p = \{\sum a_n p^n\}$
$m_p = z_p \mathcal{O}_p$	$m_p = p\mathbb{Z}_p$
residue field $k = A_p/m_p = \mathbb{C}$	residue field $k = A_p/m_p = \mathbb{F}_p$
value $f(p)$, $f \in \mathbb{C}[z]$	value $n \bmod p$, $n \in \mathbb{Z}$
$L = \mathcal{M}(X), \pi : X \to \widehat{\mathbb{C}}$	extension field $L/\mathbb{Q}$
$B = \{f$ holomorphic on $X_0 = \pi^{-1}(\mathbb{C})\}$	$B =$ integral closure of A in L
$P \in X_0 \, : \, \pi(P) = p$	prime P lying over p
$B_P \cdot m_p = m_P^e$; $e =$ ramification index	
$k' = B_P/m_P = \mathbb{C}$	$k' = B_P/m_P = \mathbb{F}_{p^f}$; $f =$ residue degree

Table 1. A brief dictionary.

Proof. Let $v : K^* \to \mathbb{Z}$ be a valuation. We claim v vanishes on the constant subfield $\mathbb{C} \subset K$. Indeed, if $f \in K^*$ has arbitrarily large roots $f^{1/n}$, then $v(f^{1/n}) = v(f)/n \to 0$ and thus $v(f) = 0$.

Next suppose $v(z-a) > 0$ for some a. We claim $v(z-b) = 0$ for all $b \neq a$. Indeed, $v(z-b) \geq \min v(z-a), v(b-a) = 0$. If in fact $v(z-b) > 0$ then we'd have $v(b-a) \geq \min v(z-a), v(z-b) > 0$, a contradiction. Since every rational function is a product of a constant and terms of the form $(z-c)$ and their reciprocals, the valuation is now determined by its value on $z-a$. Since the valuation is surjective, $v(z-a) = 1$ and thus $v = v_a$.

Finally suppose $v(z-a) \leq 0$ for all a. Using the fact that $(z-a)^{-1} - (z-b)^{-1}$ is a constant multiple of $((z-a)(z-b))^{-1}$, we find:

$$v((z-a)^{-1}) \geq \min v(((z-a)(z-b))^{-1}), v((z-b)^{-1}) = v((z-b)^{-1}).$$

Thus $v(z-c)$ is the same negative constant for all values of c. As before the constant must be -1, and thus $v = v_\infty$ measures the order of vanishing of f at infinity. ∎

Local rings and fields. For a summary of connection, see Table 1. Example: Let $L = \mathbb{Q}(\sqrt{D})$, where $D \in \mathbb{Z}$ is square-free. Then (forgetting the prime 2) we have ramification over the primes $p|D$. At these primes we have $B_P \cong \mathbb{Z}_p[p^{1/2}]$, just as for Puiseux series. For primes not dividing D, either:

(i) $T^2 - D$ is irreducible $\bmod p$, $f = 1$ and there is a unique P over p; or

(ii) D is a square mod p, and there are two primes P_1 and P_2 over p.

In case (i) the prime p behaves like a circle (closed string?) rather than a point, and P like a double cover of p.

In general when P/p has ramification index e and residue degree f, it can be thought of roughly as modeled on the map $(z, w) \mapsto (z^e, w^f)$ of $\Delta \times S^1$ to itself.

5 Holomorphic and harmonic forms

The cotangent space. Let $p \in X$ be a point on a Riemann surface, and let $m_p \subset \mathcal{C}^\infty_p$ be the ideal of smooth, complex-valued functions vanishing at p. Let m_p^2 be the ideal generated by products of element in m_p.

Theorem 5.1 *The ideal m_p^2 consists exactly of the smooth functions all of whose derivatives vanish at p.*

Proof. (Cf. [Hel, p.10]). Clearly the derivatives of fg, $f, g \in m_p$, all vanish.

For the converse, let us work locally at $p = 0 \in \mathbb{C}$, and suppose $f \in m_0$. Let f_1 and f_2 denote the partial derivatives of f with respect to x and y. Then we have:

$$\begin{aligned} f(x,y) &= \int_0^1 \frac{d}{dt} f(tx,ty)\,dt = x \int_0^1 f_1(tx,ty)\,dt + y \int_0^1 f_2(tx,ty)\,dt \\ &= x g_1(x,y) + y g_2(x,y) \end{aligned}$$

where $g_i(0,0) = f_i(0,0)$. Thus $g_1, g_2 \in m_p$ if both derivatives of f vanish, which implies $f \in m_p^2$. ∎

Corollary 5.2 *The vector space $T_p^{(1)} = m_p/m_p^2$ is isomorphic to $\mathbb{C}^2$.*

We refer to $T_p^{(1)}$ as the *complexified cotangent space* of the real surface X at p. The exterior differential of a function can then be regarded as the map defined by

$$(df)_p = [f(z) - f(p)] \in m_p/m_p^2.$$

The subspace $d\mathcal{O}_p$ is $T_p^{(1,0)} = T_p^* X \cong \mathbb{C}$, the *complex* (or holomorphic) tangent space. Its complex conjugate is $T^{(0,1)}$. In a local holomorphic coordinate $z = x + iy$, we have:

$$T_p^{(1)} = \mathbb{C} \cdot dx \oplus \mathbb{C} \cdot dy = \mathbb{C} \cdot dz \oplus \mathbb{C} \cdot d\overline{z} = T_p^{(1,0)} \oplus T_p^{(0,1)}.$$

The dual of T_p^* is T_p, the complex tangent space to X at p.

Exterior differentials. We have a natural splitting $d = \partial + \overline{\partial}$ where:

$$df = \frac{df}{dz}\,dz + \frac{df}{d\overline{z}}\,d\overline{z} = \partial f + \overline{\partial} f,$$

and

$$d(f\,dz + g\,d\overline{z}) = \left(-\frac{df}{d\overline{z}} + \frac{dg}{dz} \right) dz \wedge d\overline{z}.$$

(Recall that $dz \wedge d\overline{z} = 2i\,dx \wedge dy$.) Note that one 1-forms we have:

$$d(f\,dz) = (\overline{\partial} f)\,dz \quad \text{and} \quad d(g\,d\overline{z}) = (\partial g)\,d\overline{z}.$$

Holomorphic forms. A function $f : X \to \mathbb{C}$ is *holomorphic* if $\overline{\partial} f = 0$.

A $(1,0)$-form α on X is *holomorphic* if the following equivalent conditions hold:

1. $d\alpha = 0$; i.e. α is closed.

2. Locally $\alpha = df$ with f holomorphic.

3. Locally $\alpha = \alpha(z)\,dz$ with $\alpha(z)$ holomorphic.

4. $\overline{\partial}\alpha = 0$.

The space of all such forms is denoted $\Omega(X)$.

Examples: $\Omega(\widehat{\mathbb{C}}) = 0$; $\Omega(\mathbb{C}/\Lambda) = \mathbb{C} \cdot dz$; $dz/z \in \Omega(\mathbb{C}^*)$.

Meromorphic forms and the residue theorem. These are forms expressed locally as $\alpha = \alpha(z)\,dz$, with $\alpha(z)$ meromorphic.

The *residue* of α at $p \in X$ is defined by $\mathrm{Res}_p(\alpha) = a_1$ if locally $\alpha(z) = \sum a_n z^n\,dz$, where $z(p) = 0$; equivalently, by $\mathrm{Res}_p(\alpha) = (2\pi i)^{-1} \int_\gamma \alpha$ where γ is a small loop around p (inside of which α has only one pole).

Theorem 5.3 *If X is compact, then $\sum_p \mathrm{Res}_p(\alpha) = 0$.*

By considering $df/f = d\log f$, this gives another proof of:

Corollary 5.4 *If $f : X \to \widehat{\mathbb{C}}$ is a meromorphic function on a compact Riemann surface, then f has the same number of zeros as poles (counted with multiplicity).*

(The first proof was that $\sum_{f(x)=y} \mathrm{mult}(f,x) = \deg(f)$ by general considerations of proper mappings.)

Since the ratio α_1/α_2 of any two meromorphic 1-forms is a meromorphic function (so long as $\alpha_2 \neq 0$), we have:

Theorem 5.5 *The 'degree' of a meromorphic 1-form on a compact Riemann surface — that is, the difference between the number of zeros and the number of poles — is independent of the form.*

Example: meromorphic 1-forms on $\widehat{\mathbb{C}}$ have degree -2, i.e. they have 2 more poles than zeros.

Harmonic forms and periods. A function $u : X \to \mathbb{C}$ is *harmonic* if $\overline{\partial}\partial u = 0$; equivalently, if ∂u is a holomorphic 1-form.

A 1-form ω on X is *harmonic* if the following equivalent conditions hold:

1. $\partial\omega = \overline{\partial}\omega = 0$ (in particular, ω is closed);

2. Locally $\omega = du$ with u harmonic;

3. Globally $\omega = \alpha + \beta$ with $\alpha \in \Omega(X)$ and $\beta \in \overline{\Omega}(X)$.

The space of all such forms is denoted $\mathcal{H}^1(X)$. By the last condition we have $\mathcal{H}^1(X) = \Omega(X) \oplus \overline{\Omega}(X)$. In particular, every holomorphic 1-form is harmonic.

Example: dx and dy are harmonic 1-forms on $X = \mathbb{C}/\Lambda$.

The *period map* associates to ω the homomorphism $\gamma \mapsto \int_\gamma \omega$ on $\pi_1(X)$.

Theorem 5.6 *Let X be a compact Riemann surface of genus g. Then the period map*

$$\mathcal{H}^1(X) \to \mathrm{Hom}(\pi_1(X), \mathbb{C}) \cong \mathbb{C}^{2g}$$

is injective.

Proof. Any form in the kernel can be expressed globally as $\omega = du$ where u is harmonic — and hence constant — on the compact Riemann surface X. ■

Corollary 5.7 *The space $\Omega(X)$ has dimension $\leq g$, and the space $\mathcal{H}^1(X)$ has dimension $\leq 2g$.*

Example: the torus. For $X = \mathbb{C}/\Lambda$ we have $\Omega(X) = \mathbb{C}\omega$ with $\omega = dz$, and the period map sends $\pi_1(X) \cong \mathbb{Z}^2$ to Λ.

Example: the regular octagon. Here is concrete example of a compact Riemann surface X and a holomorphic 1-form $\omega \in \Omega(X)$. Namely let $X = Q/\equiv$ where Q is a regular octagon in the plane, with vertices at the 8th roots of unity, and $\equiv$ identifies opposite edges. Since translations preserve dz, the form $\omega = dz/\equiv$ is well-defined on X. It has a zero of order two at the single point $p \in X$ coming from the vertex. The edges of Q form a system of generators for $\pi_1(X)$, and the periods $\int_\alpha \omega$ are given by $\zeta^{i+1} - \zeta^i$ where ζ is a primitive 8th root of unity.

Example: hyperelliptic Riemann surfaces.

Theorem 5.8 *Let $B \subset \mathbb{C}$ be a finite set of cardinality $2n$, and let $p(T) = \prod_B (T - b)$. Let $\pi : X \to \widehat{\mathbb{C}}$ be the unique 2-fold covering of $\widehat{\mathbb{C}}$ branched over B. Then X has genus $g = n - 1$, and the forms*

$$\omega_i = \frac{z^i\, dz}{\sqrt{p(z)}}, \quad i = 0, 1, \ldots, n-2$$

form a basis for $\Omega(X)$. In particular $\dim \Omega(X) = g$.

Proof. To check this, we begin by investigating ω_0. Note that if $z = w^2$ then $dz = 2w\,dw$. Thus the pullback of dz to X has simple zeros at the branch points of π, which lie over the zeros of $p(z)$. Also $\sqrt{p(z)}$, as a meromorphic function on X, has simple zeros in the same locations. These zeros cancel when we form the quotient $dz/\sqrt{p(z)}$, and thus ω_0 is holomorphic except possibly over the two unbranched points $p_1, p_2 \in X$ lying over $z = \infty$. But now dz has a pole of order 2 at $z = \infty$, while $1/\sqrt{p(z)}$ has a zero of order n at $z = \infty$. We conclude:

> *The form ω_0 is holomorphic on X, with zeros of order $n-2$ at p_1 and p_2 and nowhere else.*

In particular, the degree of a meromorphic 1-form on X is $2n - 4 = 2g - 2$. Since z^i has a pole of order i at $z = \infty$, it follows that the g forms ω_i on X are holomorphic for $i \le n - 2$. ■

Note: periods. When $B = \{r_1, \ldots, r_{2n}\} \subset \mathbb{R}$, the periods of ω_0 can be expressed in terms of the integrals

$$\int_{r_i}^{r_{i+1}} \frac{dx}{\sqrt{(x-r_1)\cdots(x-r_{2n})}}.$$

We can also allow one point of B to become infinity. An example of a period that can be determined explicitly comes from the square torus:

$$\int_0^1 \frac{dx}{x(x^2-1)} = \frac{-2i\sqrt{\pi}\,\Gamma(5/4)}{\Gamma(3/4)}.$$

A more general type of period is:

$$\zeta(3) = \sum n^{-3} = \int_{0<x<y<z<1} \frac{dx\,dy\,dz}{(1-x)yz}.$$

See [KZ] for much more on periods.

General surfaces of genus g. The case of hyperelliptic curves suggests the following result, which we will later prove.

Theorem 5.9 *For any compact Riemann surface of genus g, we have* $\dim \Omega(X) = g$*; and any meromorphic 1-form on X has $2g-2$ more zeros than poles.*

The Hodge star operator. Let V be an n-dimensional real vector space with an inner product $\langle v_1, v_2\rangle$. Choose an orthonormal basis $e_1, \ldots, e_n$ for

V. Then the wedge products $e_I = e_{i_1} \wedge \cdots e_{i_k}$ provide an orthonormal basis for $\bigwedge^k V$. The *Hodge star* operator $* : \bigwedge^k V \to \bigwedge^{n-k} V$ is the unique linear map satisfying

$$*e_I = e_J \quad \text{where} \quad e_I \wedge *e_I = e_1 \wedge e_2 \wedge \cdots e_n.$$

Here J are the indices not occurring in I, ordered so the second equation holds. More generally we have:

$$v \wedge *w = \langle v, w \rangle e_1 \wedge e_2 \wedge \cdots e_n.$$

and thus $v \wedge *w = w \wedge *v$. Since $v \wedge *v = (-1)^{k(n-k)} *v \wedge v$, we have $*^2 = (-1)^{k(n-k)}$ on $\bigwedge^k V$. Equivalently, $*^2 = (-1)^k$ when n is even, and $*^2 = 1$ when n is odd.

Now let (M, g) be a compact Riemannian manifold. We can then try to represent each cohomology class by a closed form minimizing $\int_M \langle \alpha, \alpha \rangle$. Formally this minimization property implies:

$$d\alpha = d * \alpha = 0, \tag{5.1}$$

using the fact that a minimizer satisfies

$$\int_M \langle d\beta, \alpha \rangle = \int_M (d\beta) \wedge *\alpha = -\int_M \beta \wedge d * \alpha = 0$$

for all smooth $(k-1)$-forms β. Thus we call α *harmonic* if (5.1) holds, and let $\mathcal{H}^k(M)$ denote the space of all harmonic k-forms.

Theorem 5.10 (Hodge) *There is a natural isomorphism $\mathcal{H}^k(M) \cong H^k_{DR}(M)$.*

Adjoints. The *adjoint differential* $d^* : \mathcal{E}^k(M) \to \mathcal{E}^{k-1}(M)$ is defined so that:

$$\langle d\alpha, \beta \rangle = \langle \alpha, d^* \beta \rangle,$$

where $\langle \alpha, \beta \rangle = \int_M \alpha \wedge *\beta$. It is given by

$$d^*(\alpha) = \pm * d * \alpha$$

for a suitable choice of sign, since:

$$\langle d\alpha, \beta \rangle = \int d\alpha \wedge *\beta = -\int \alpha \wedge d * \beta = \pm \int \alpha \wedge *(*d * \beta).$$

Since $*^2 = 1$ on an odd-dimensional manifold, in that case we have $d^* = - * d*$. For a k-form β on an even dimensional manifold, we have instead:

$$d^*\beta = (-1)^k * d * \beta.$$

Here we have used the fact that $*^2 = (-1)^{k-1}$ on $n-(k-1)$ forms such as $d * \beta$.

Generalized Hodge theorem. Once the adjoint d^* is in play, the arguments of the Hodge theorem give a complete picture of *all* smooth k-forms on M.

Theorem 5.11 *The space of smooth k-forms has an orthogonal splitting:*

$$\mathcal{E}^k(M) = d\mathcal{E}^{k-1}(M) \oplus \mathcal{H}^k(M) \oplus d^*\mathcal{E}^{k+1}(M).$$

The Laplacian. Once we have a metric we can combine d and d^* to obtain the *Hodge Laplacian*

$$\Delta : \mathcal{E}^k(M) \to \mathcal{E}^k(M),$$

defined by

$$\Delta\alpha = (dd^* + d^*d)\alpha.$$

Theorem 5.12 *A form α on a compact manifold is harmonic iff $\Delta\alpha = 0$.*

Proof. Clearly $d\alpha = d * \alpha = 0$ implies $\Delta\alpha = 0$. Conversely if $\Delta\alpha = 0$ then:

$$\begin{aligned} 0 &= \int_M \langle \Delta\alpha, \alpha\rangle \\ &= \int_M \langle d\alpha, d\alpha\rangle + \langle d^*\alpha, d^*\alpha\rangle \end{aligned}$$

and so $d\alpha = d * \alpha = 0$ as well. ■

Note that on functions these definitions give

$$\Delta f = d^* df = - * d * df,$$

independent of the dimension of M. This satisfies $\int \langle f, \Delta f\rangle \geq 0$, but differs by a sign from the usual Euclidean Laplacian. (For example on S^1 we have $\int f f'' = -\int |f'|^2 \leq 0$ for the usual Laplacian.)

Riemann surfaces. Now suppose M has even dimension $n = 2k$. The Hodge star on the middle-dimensional k-forms is then *conformally invariant.*

Thus it makes sense to talk about harmonic k-forms when only a conformal structure is present.

In particular, the Hodge star is canonical for 1-forms on a Riemann surface, and can be expressed by $*dx = dy$ and $*dy = -dx$ for a local coordinate with $z = x + iy$. The pair of conditions $d\alpha = d * \alpha = 0$ are then the same as the pair of conditions $\partial\alpha = \overline{\partial}\alpha = 0$ for a 1-form.

Geometrically, for α to be closed means that the foliation defined by $\operatorname{Ker}\alpha$ admits a transverse invariant measure. The orthogonal foliation, defined by $*\alpha$, *also* admits such a measure iff α is harmonic.

Example: the level sets of $\operatorname{Re} f(z)$ and $\operatorname{Im} f(z)$ give the foliations associated to the form $\alpha = du$, $u = \operatorname{Re} f(z)$. The case $f(z) = z + 1/z$ in particular gives foliations by confocal ellipses and hyperboli, with foci ± 1 coming from the critical points of f.

Laplacian on Riemann surfaces. A harmonic function is one which satisfies $\Delta f = 0$; equivalently, $d * df = 0$. Since the Hodge star is natural for 1-forms on a Riemann surface, the harmonic functions are conformally invariant. For the same reason, harmonic 1-forms are conformally natural: the 2-form $d * \alpha = 0$ is independent of the conformal factor when α is a 1-form on a 2-manifold. In local coordinates $z = x + iy$ we have explicitly:

$$*dx = dy \quad \text{and} \quad *\, dy = -dx.$$

We also obtain a 'conformally natural' Laplacian sending functions to 2-forms (or measures); it is given by

$$\begin{aligned} \Delta f &= d * df = (f_{xx} + f_{yy})\, dx \wedge dy \\ &= i(\partial - \overline{\partial})(\partial + \overline{\partial}) f = 2i f_{z\overline{z}}\, dz \wedge d\overline{z}. \end{aligned}$$

(Note that $f_{z\overline{z}} = (1/4)(f_{xx} + f_{yy})$ and $dz \wedge d\overline{z} = -2i\, dx \wedge dy$, so indeed equality holds.)

On the other hand, the *spectrum* of the Laplacian on functions a Riemann surface with a metric depends very much on the metric. This is because we must divide by the volume form to get a map Δ sending functions to functions.

Since it is natural to take functions and forms on a Riemann surface to be complex valued, the *complexified* Hodge star includes composition with complex conjugation. This insures, for example, that

$$\int_M f \wedge *f = \int_M |f|^2\, dV.$$

With this convention, $*(i\alpha) = -i * \alpha$, and thus

$$*dz = *(dx + i\,dy) = dy + i\,dx = id\overline{z},$$

and $*d\overline{z} = -i\,dz$.

6 Cohomology of sheaves

Maps of sheaves; exact sequences. A map between sheaves is always specified at the level of open sets, by a family of compatible morphisms $\mathcal{F}(U) \to \mathcal{G}(U)$. A map of sheaves *induces* maps $\mathcal{F}_p \to \mathcal{G}_p$ between stalks. We say $\mathcal{F} \to \mathcal{G}$ is injective, surjective, an isomorphism, etc. iff $\mathcal{F}_p \to \mathcal{G}_p$ has the same property for each point p.

We say a sequence of sheaves $\mathcal{A} \to \mathcal{B} \to \mathcal{C}$ is *exact* at $\mathcal{B}$ if the sequence of groups

$$\mathcal{A}_p \to \mathcal{B}_p \to \mathcal{C}_p$$

is exact, for every p.

The exponential sequence. As a prime example: on any Riemann surface X, the sequence of sheaves

$$0 \to \mathbb{Z} \to \mathcal{O} \to \mathcal{O}^* \to 0$$

is exact. But it is only exact on the level of stalks! For every open set the sequence

$$0 \to \mathbb{Z}(U) \to \mathcal{O}(U) \to \mathcal{O}^*(U)$$

is exact, but the final arrow need not be surjective. (Consider $f(z) = z \in \mathcal{O}^*(\mathbb{C}^*)$; it cannot be written in the form $f(z) = \exp(g(z))$ with $g \in \mathcal{O}(\mathbb{C}^*)$.)

More generally, we have:

Theorem 6.1 *The global section functor is left exact. That is, for any short exact sequence of sheaves,* $0 \to \mathcal{A} \to \mathcal{B} \to \mathcal{C} \to 0$*, the sequence of global sections*

$$0 \to \mathcal{A}(U) \to \mathcal{B}(U) \to \mathcal{C}(U)$$

is also exact.

Sheaf cohomology is the derived functor which measures the failure of exactness to hold on the right.

Čech: the nerve of a covering. A precursor to sheaf cohomology is Čech cohomology. The idea here is that any open covering $\mathfrak{U} = (U_i)$ of X has an associated simplicial complex that is an approximation to the topology of X. The simplices in this complex are simply ordered finite sequence of indices I such that $\bigcap_I U_i \neq \emptyset$.

This works especially well if we require that all the multiple intersections are *connected*. Note that this is equivalent to requiring that $\mathbb{Z}(U_I) \cong \mathbb{Z}$ whenever $U_I \neq \emptyset$.

Cochains, cocycles and coboundaries. Now suppose we also have a sheaf $\mathcal{F}$ in play (for classical Čech cohomology, this sheaf is just $\mathbb{Z}$). We then put 'weights on our simplices' and define the space of q-cochains by:

$$C^q(\mathfrak{U}, \mathcal{F}) = \prod_{|I|=q+1} \mathcal{F}(U_I).$$

Here I ranges over *ordered* sets of indices $(i_0, \ldots, i_q)$, and

$$U_I = U_{i_0} \cap \cdots \cap U_{i_q}.$$

Examples: a 0-cochain is the data $f_i \in \mathcal{F}(U_i)$; a 0-cochain is the data $g_{ij} \in \mathcal{F}(U_i \cap U_j)$; etc.

Next we define a boundary operator

$$\delta : C^q(\mathfrak{U}, \mathcal{F}) \to C^{q+1}(\mathfrak{U}, \mathcal{F})$$

by setting $\delta f = g$ where, for $q = 0$:

$$g_{ij} = f_j - f_i;$$

for $q = 1$:

$$g_{ijk} = f_{jk} - f_{ik} + f_{ij},$$

and more generally

$$g_I = \sum_0^q (-1)^j f_{I_j}$$

where $I_j = (i_0, i_1, \ldots, \widehat{i_j}, \ldots, i_{q+1})$. When two indices are eliminated, they come with opposite sign, so $\delta^2 = 0$.

The kernel of δ is the group of *cocycles* $Z^q(\mathfrak{U}, \mathcal{F})$, its image is the group of *coboundaries* $B^q(\mathfrak{U}, \mathcal{F})$, and the qth cohomology group of $\mathcal{F}$ relative to the covering $\mathfrak{U}$ is defined by:

$$H^q(\mathfrak{U}, \mathcal{F}) = Z^q(\mathfrak{U}, \mathcal{F}) / B^q(\mathfrak{U}, \mathcal{F}).$$

Example: H^0. A 0-cocycle (f_i) is a coboundary iff $f_j - f_i = g_{ij} = 0$ for all i and j. By the sheaf axioms, this happens iff $f_i = f|U_i$, and thus:

$$H^0(\mathfrak{U}, \mathcal{F}) = \mathcal{F}(X).$$

Example: H^1. A 1-cocycle g_{ij} satisfies $g_{ii} = 0$, $g_{ij} = -g_{ji}$ and

$$g_{ij} + g_{jk} = g_{ik}$$

on U_{ijk}. It is a coboundary if it can be written in the form $g_{ij} = f_i - f_j$.

Refinement. Whenever $\mathfrak{V} = (V_i)$ is a finer covering than $\mathfrak{U} = (U_i)$, we can choose a refinement map on indices such that $V_i \subset U_{\rho i}$. Once ρ is specified, it determines maps $\mathcal{F}(U_{\rho I}) \to \mathcal{F}(V_I)$, and hence chain maps giving rise to a homomorphism

$$H^q(\mathfrak{U}, \mathcal{F}) \to H^q(\mathfrak{V}, \mathcal{F}).$$

Theorem 6.2 *The refinement map $H^q(\mathfrak{U}, \mathcal{F}) \to H^q(\mathfrak{V}, \mathcal{F})$ is independent of ρ.*

Definition. We define the cohomology of X with coefficients in $\mathcal{F}$ by:

$$H^q(X, \mathcal{F}) = \varinjlim H^q(\mathfrak{U}, \mathcal{F}),$$

where the limit is taken over the system of all open coverings, directed by refinement.

Theorem 6.3 *The refinement map $H^q(\mathfrak{U}, \mathcal{F}) \to H^q(\mathfrak{V}, \mathcal{F})$ is injective.*

Proof for $q = 1$. Suppose we are given coverings (U_i) and (V_i) with $V_i \subset U_{\rho i}$. Let g_{ij} be a 1-cocycle for the covering (U_i) that becomes trivial for (V_i). That means there exist $f_i \in \mathcal{F}(V_i)$ such that

$$g_{\rho i, \rho j} = f_i - f_j$$

on V_{ij}.

Our goal is to find $h_i \in \mathcal{F}(U_i)$ so $g_{ij} = h_i - h_j$. Note that on $V_{ij} \cap U_k$ we have:

$$g_{\rho i,k} + g_{k,\rho j} = f_i - f_j,$$

and thus we can define

$$h_k = f_i - g_{\rho i,k} = f_j - g_{\rho j,k}$$

consistently throughout U_k. We then have, on $U_{kl} \cap V_i$,

$$h_k - h_l = f_i - g_{\rho i,k} - f_i + g_{\rho i,l} = g_{kl}$$

as desired. ∎

Theorem 6.4 (Leray) *If $\mathfrak{U}$ is acyclic ($H^q(U_I, \mathcal{F}) = 0$ for all $q > 0$) then*

$$H^q(X, \mathcal{F}) \cong H^q(\mathfrak{U}, \mathcal{F}).$$

If just $H^1(U_i, \mathcal{F}) = 0$ for every i, $\mathfrak{U}$ can still be used to compute H^1.

Example: Let $S^1 = U_0 \cup U_1$ be a covering by a pair of intervals. Then $Z^1(\mathfrak{U}, \mathbb{Z}) = \mathbb{Z}(U_0 \cap U_1) = \mathbb{Z}^2$, since $U_0 \cap U_1$ has two components; while $B^1(\mathfrak{U}, \mathbb{Z}) = \mathbb{Z} = \{(a-b, a-b)\} \subset \mathbb{Z}^2$. Thus $H^1(S^1, \mathbb{Z}) = \mathbb{Z}$.

Vanishing theorem by dimension. Using the existence of fine coverings where the $(n+2)$-fold intersections are all empty, we have:

Theorem 6.5 *For any n-dimensional space X and any sheaf $\mathcal{F}$, $H^p(X, \mathcal{F}) = 0$ for all $p > n$.*

Vanishing theorems for smooth functions, forms, etc. Let $\mathcal{F}$ be the sheaf of C^∞ functions on a (paracompact) manifold X, or more generally a sheaf of modules over C^∞. We then have:

Theorem 6.6 *The cohomology groups $H^q(X, \mathcal{F}) = 0$ for all $q > 0$.*

Proof for $q = 1$. To indicate the argument, we will show $H^1(\mathfrak{U}, \mathcal{F}) = 0$ for any open covering $\mathfrak{U} = (U_i)$. We will use the fact that there exists a partition of unity $\rho_i \in C^\infty(X)$ subordinate to U_i: that is, a set of functions with $K_i = \operatorname{supp} \rho_i \subset U_i$, such that K_i forms a locally finite covering of X and $\sum \rho_i(x) = 1$ for all $x \in X$.

Let $g_{ij} \in Z^1(\mathfrak{U}, \mathcal{F})$ be a 1-cocycle. Then $g_{ii} = 0$, $g_{ij} = -g_{ji}$ and $g_{ij} + g_{jk} = g_{ik}$. Our goal is to write $g_{ij} = f_j - f_i$ (or $f_i - f_j$).

How will we ever get from g_{ij}, which is only define on U_{ij}, a function f_i define on all of U_i? The central observation is that:

$\rho_j g_{ij}$, extended by 0, is smooth on U_i.

This is because $\operatorname{supp} \rho_j g_{ij} \subset K_j \cap U_i$ is *closed* as a subset of U_i (even though it might not be closed as a subset of X.) Thus we can define:

$$f_i = \sum_k \rho_k g_{ik};$$

and then:

$$f_i - f_j = \sum_k \rho_k (g_{ik} - g_{jk}) = \sum_k \rho_k (g_{ik} + g_{kj}) = \sum_k \rho_k g_{ij} = g_{ij}.$$

■

The exact cohomology sequence; deRham cohomology. We can now explain how sheaf cohomology is used to capture global aspects of analytic problems that can be solved locally.

Let $\mathcal{E}^p$ denote the sheaf of smooth p-forms on a manifold X. Suppose $\alpha \in \mathcal{E}^1(X)$ is closed; then locally $\alpha = df$ for $f \in \mathcal{E}^0(X)$. When we can we find a global primitive for α?

To solve this problem, let $\mathfrak{U} = (U_i)$ be an open covering of X by disks. Then we can write $\alpha = df_i$ on U_i. On the overlaps, $g_{ij} = f_i - f_j$ satisfies $dg_{ij} = 0$, i.e. it is a constant function. Moreover we obviously have $g_{ij} + g_{jk} = g_{ik}$, i.e. g_{ij} is an element of $Z^1(\mathfrak{U}, \mathbb{C})$.

Now we may have chosen our f_i wrong to fit together, since f_i is not uniquely determined by the condition $df_i = \alpha_i$; we can always add a constant function c_i. But if we replace f_i by $f_i + c_i$, then g_{ij} will change by the coboundary $c_i - c_j$. Thus we can conclude:

$$\alpha = df \ \textit{iff} \ [g_{ij}] = 0 \ \textit{in} \ H^1(X, \mathbb{C}) \ .$$

The exact cohomology sequence. The conceptual theorem underlying the preceding discussion is the following:

Theorem 6.7 *Any short exact sequence of sheaves on X,*

$$0 \to \mathcal{A} \to \mathcal{B} \to \mathcal{C} \to 0,$$

gives rise to a long exact sequence

$$\begin{aligned} 0 \quad \to \quad & H^0(X, \mathcal{A}) \to H^0(X, \mathcal{B}) \to H^0(X, \mathcal{C}) \to \\ & H^1(X, \mathcal{A}) \to H^1(X, \mathcal{B}) \to H^1(X, \mathcal{C}) \to \\ & H^2(X, \mathcal{A}) \to H^2(X, \mathcal{B}) \to H^2(X, \mathcal{C}) \to \cdots \end{aligned}$$

on the level of cohomology.

Note: for *any* open set U, we get an exact sequence

$$0 \to \mathcal{A}(U) \to \mathcal{B}(U) \to \mathcal{C}(U), \qquad (6.1)$$

as can be checked using the sheaf axioms. Surjectivity of the maps $\mathcal{B}_x \to \mathcal{C}_x$ implies that for any $c \in \mathcal{C}(X)$, there is an open covering (U_i) and $b_i \in \mathcal{B}(U_i)$ mapping to c.

To obtain the *connecting homomorphism*

$$\delta^* : H^0(X, \mathcal{C}) \to H^1(X, \mathcal{A}),$$

we use the exactness of (6.1) to write $b_i - b_j$ as the image of a_{ij}. The resulting cocycle $[a_{ij}] \in H^1(X, \mathcal{A})$ is the image $\delta^*(c)$.

Example: deRham cohomology. Let $\mathcal{Z}^p$ denote the sheaf of closed p-forms, then (by the Poincaré lemma) we have an exact sequence of sheaves:

$$0 \to \mathcal{Z}^{p-1} \to \mathcal{E}^{p-1} \xrightarrow{d} \mathcal{Z}^p \to 0.$$

Recall that the deRham cohomology groups of X are given by:

$$H^p_{DR}(X) = \mathcal{Z}^p(X)/d\mathcal{E}^{p-1}(X).$$

Now let $p = 1$. Then $\mathcal{Z}^{p-1} = \mathbb{C}$. By examining the associate long exact sequence we find:

Theorem 6.8 *For any manifold X, we have $H^1_{DR}(X) \cong H^1(X, \mathbb{C})$.*

More generally, using all the terms in the exact sequence and all values of p, we find:

Theorem 6.9 *For any manifold X, we have*

$$H^p_{DR}(X) \cong H^1(X, \mathcal{Z}^{p-1}) \cong H^2(X, \mathcal{Z}^{p-2}) \cong \cdots H^p(X, \mathcal{Z}^0) = H^p(X, \mathbb{C}).$$

Corollary 6.10 *The deRham cohomology groups of homeomorphic smooth manifolds are isomorphic.*

(In fact one can replace 'homeomorphic' by 'homotopy equivalent' here.)

Finiteness. Now it is easy to prove that $H^p_{DR}(\mathbb{R}^n) = 0$ for all $p > 0$. We thus have, by taking a Leray covering:

Theorem 6.11 *For any compact manifold, the cohomology groups $H^p(X, \mathbb{C})$ are finite.*

Periods revisited. Finally we mention the fundamental group:

Theorem 6.12 *For any connected manifold X, we have*

$$H^1(X, \mathbb{C}) \cong H^1_{DR}(X) \cong \mathrm{Hom}(\pi_1(X), \mathbb{C}).$$

Proof. We have $\int_\gamma df = 0$ for all closed loops γ, so the period map is well-defined; if α has zero periods then $f(q) = \int_p^q \alpha$ is also well-defined and satisfies $df = \alpha$. To prove surjectivity, take a (Leray) covering of X by geodesically convex sets, and observe that (i) every element of $\pi_1(X)$ is represented by a 1-chain and (ii) every 1-boundary is a product of commutators. ∎

Remarks. One can define the period map $H^1(X,\mathbb{C}) \to \mathrm{Hom}(\pi_1(X),\mathbb{C})$ directly, using $\gamma : S^1 \to X$ to obtain from $\xi \in H^1(X,\mathbb{C})$ a class $\phi(\gamma) \in H^1(S^1,\mathbb{C}) \cong \mathbb{C}$. For more exotic topological spaces, however, this map need not be an isomorphism: e.g. the 'topologist's sine curve' is a compact, connected space X with $\pi_1(X) = 0$ but $H^1(X,\mathbb{Z}) = \mathbb{Z}$.

7 Cohomology on a Riemann surface

On a Riemann surface we have the notion of *holomorphic* functions and forms. Thus in addition to the sheaves $\mathcal{E}^p$ we have the important sheaves:

$\mathcal{O}$ — the sheaf of holomorphic functions; and
Ω — the sheaf of holomorphic 1-forms.

Let $h^i(\mathcal{F}) = \dim H^i(X,\mathcal{F})$. We will show that a compact Riemann surface satisfies:

$$\begin{aligned} h^0(\mathcal{O}) &= 1, \quad h^0(\Omega) = g, \\ h^1(\mathcal{O}) &= g, \quad h^1(\Omega) = 1. \end{aligned}$$

The symmetry of this table is not accidental: it is rather the first instance of Serre duality, which we will also prove.

The Dolbeault Lemma. Just as the closed forms can be regarded as the subsheaf $\mathrm{Ker}\, d \subset \mathcal{E}^p$, the holomorphic functions can be regarded as the subsheaf $\mathrm{Ker}\, \overline{\partial} \subset \mathcal{C}^\infty = \mathcal{E}^0$. So we must begin by studying the $\overline{\partial}$ operator.

We begin by studying the equation $df/d\overline{z} = g \in L^1(\mathbb{C})$. An example is given for each $r > 0$ by:

$$f_r(z) = \begin{cases} 1/z & \text{if } |z| > r, \\ \overline{z}/r^2 & \text{if } |z| \le r, \end{cases}$$

which satisfies $g_r = df/d\overline{z} = (1/r^2)\chi_{B(0,r)}(z)$. In particular $\int g_r = \pi$ is independent of r, which suggests the distributional equation:

$$\frac{d}{d\overline{z}} \frac{1}{z} = \pi \delta_0.$$

Using this fundamental solution (and the fact that $dx\, dy = 1/(2i) dz \wedge d\overline{z}$), we obtain:

Theorem 7.1 *For any $g \in C_c^\infty(\mathbb{C})$, a solution to the equation $df/d\overline{z} = g$ is given by:*

$$f(z) = g * \frac{1}{\pi z} = \frac{1}{2\pi i} \int_{\mathbb{C}} \frac{g(w)\, dw \wedge d\overline{w}}{z - w}.$$

Theorem 7.2 *For any $g \in \mathcal{C}^\infty(\Delta)$, there is an $f \in C^\infty(\Delta)$ with $df/d\overline{z} = g$.*

Proof. Write $g = \sum g_n$ where each g_n is smooth and compactly supported outside the disk D_n of radius $1 - 1/n$. Solve $df_n/d\overline{z} = g_n$. Then f_n is holomorphic on D_n. Expanding it in power series, we can find holomorphic functions h_n on the disk such that $|f_n - h_n| < 2^{-n}$ n D_n. Then $f = \sum(f_n - h_n) = \lim F_n$ converges uniformly, and $F - F_i$ is holomorphic on D_n for all $i > n$, so the convergence is also C^∞. ∎

Corollary 7.3 *For any $g \in \mathcal{C}^\infty(\Delta)$, there is an $f \in C^\infty(\Delta)$ with $\Delta f = g$.*

Proof. First solve $dh/dz = g$, then $df/d\overline{z} = h$. Then $\Delta(f/4) = d^2 f/dz\, d\overline{z} = g$. ∎

Note: the same results hold with Δ replaced by $\mathbb{C}$.

Dolbeault cohomology. Let us define, for any Riemann surface X,

$$H^{0,1}(X) = \frac{\mathcal{E}^{0,1}(X)}{\overline{\partial}\mathcal{E}^0(X)}.$$

The preceding results show the sequence of sheaves:

$$0 \to \mathcal{O} \to \mathcal{E}^0 \xrightarrow{\overline{\partial}} \mathcal{E}^{0,1} \to 0$$

is exact. Consequently we have:

Theorem 7.4 *For any Riemann surface X, $H^1(X, \mathcal{O}) \cong H^{0,1}(X)$.*

Thus the Dolbeault lemma can be reformulated as:

Theorem 7.5 *The unit disk satisfies $H^1(\Delta, \mathcal{O}) \cong H^{0,1}(\Delta) = 0$. The same is true for the complex plane.*

Corollary 7.6 *We have $H^1(\widehat{\mathbb{C}}, \mathcal{O}) = 0$.*

Proof. Let $U_1 \cup U_2$ be the usual covering by $U_1 = \mathbb{C}$ and $U_2 = \widehat{\mathbb{C}} - \{0\}$. By the preceding result, $H^1(U_i, \mathcal{O}) = 0$. Thus by Leray's theorem, this covering is sufficient for computing H^1: $H^1(\widehat{\mathbb{C}}, \mathcal{O}) = H^1(\mathfrak{U}, \mathcal{O})$. Suppose $g_{12} \in \mathcal{O}(U_{12}) = \mathcal{O}(\mathbb{C}^*)$ is given. Then $g_{12}(z) = \sum_{-\infty}^{\infty} a_n z^n$. Splitting this Laurent series into its positive and negative parts, we obtain $f_i \in \mathcal{O}(U_i)$ such that $g_{12} = f_2 - f_1$. ∎

Similarly, we define

$$H^{1,1}(X) = \frac{\mathcal{E}^{1,1}(X)}{\overline{\partial}\mathcal{E}^{1,0}(X)}.$$

Theorem 7.7 *We have $H^1(X,\Omega) = H^{1,1}(X)$.*

Proof. Consider the exact sequence of sheaves

$$0 \to \Omega \to \mathcal{E}^{1,0} \xrightarrow{\overline{\partial}} \mathcal{E}^{1,1} \to 0.$$

■

The Dolbeault isomorphism. Using the fact that $\overline{\partial}^2 = 0$ one can defined the Dolbeault cohomology groups for general complex manifolds, and prove using sheaf theory the following variant of the deRham theorem:

Theorem 7.8 *For any compact complex manifold X, we have $H^{p,q}_{\overline{\partial}}(X) \cong H^q(X,\Omega^p)$.*

Here Ω^q is the sheaf of holomorphic $(q,0)$-forms.

The residue map. Since $d = \overline{\partial}$ on $\mathcal{E}^{1,0}$, the *residue map*

$$\operatorname{Res}(\alpha) = \frac{1}{2\pi i}\int_X \alpha$$

gives a *canonical* map

$$\operatorname{Res} : H^1(X,\Omega) \cong H^{1,1}(X) \to \mathbb{C}.$$

(We will later see that this map is an isomorphism.)

Now suppose we have an element $\xi = [\alpha_{ij}] \in H^1(X,\Omega)$ that can be expressed in the form $\alpha_{ij} = \omega_i - \omega_j$, with $\omega_i \in \mathcal{M}^{(1)}(U_i)$. We then refer to (ω_i) as *Mittag-Leffler* data, and ξ as a *Mittag-Leffler coboundary.* We can then associate to ξ quantity

$$R(\xi) = \sum_p \operatorname{Res}_p(\omega_i).$$

Note that if $\alpha_{ij} = \beta_i - \beta_j$ then there exists a global meromorphic form given locally by $\eta = \omega_i - \beta_i$. Since $\sum \operatorname{Res}_p(\eta) = 0$, the value of $R(\xi)$ is well-defined.

Example. For any point $P \in X$ we can choose a local coordinate $z : U_1 \cong \Delta$, set $U_2 = X - \{P\}$, and define $\omega_1 = dz/z$ and $\omega_2 = 0$. Then $\alpha_{12} = dz/z$ on U_{12}, and $\xi = [\alpha_{ij}]$ satisfies $R(\xi) = 1$.

Theorem 7.9 *If $\xi = [\omega_i - \omega_j]$ is a Mittag-Leffler coboundary, then*

$$\mathrm{Res}(\xi) = \sum \mathrm{Res}_p(\omega_i).$$

Proof. Using the fact that $H^1(X, \mathcal{E}^{1,0}) = 0$, we can write the cocycle $\xi = (\alpha_{ij})$ as a coboundary in two ways: we have

$$\alpha_{ij} = \omega_i - \omega_j = g_i - g_j$$

with $g_i \in \mathcal{E}^{1,0}(U_i)$. Then $\eta = dg_i$ is globally well-defined and $\mathrm{Res}(\xi) = \int_X \eta$.

Now let $X^* \subset X$ be the complement of very small disks around the poles of the ω_i. Then on X^*, we have a global smooth 1-form $h = g_i - \omega_i$ on U_i (these definitions agree on the overlap by the above). Since ω_i is holomorphic, we have $dh = dg_i = \eta$, and thus:

$$2\pi i \,\mathrm{Res}(\xi) = \int_X \eta \approx \int_{X^*} dh = \int_{\partial X^*} h \approx 2\pi i \sum \mathrm{Res}_p(\omega_i).$$

In the last step we have taken into account the fact that ∂X^* is oriented to give loops that go negatively around the points p, and that $\int_{\partial X^*} g_i \approx 0$ since g_i is smooth at p. ∎

Finiteness. Recall that we already know $\dim \Omega(X) \geq g$, by period considerations; in particular, $\Omega(X)$ is finite-dimensional. We remark that a more robust proof of this finite-dimensionality can be given by endowing $\Omega(X)$ with a reasonable normal — e.g. $\|\alpha\|^2 = \int_X |\alpha|^2$ — and then observing that the unit ball is compact. (The same proof applies to show $\dim \mathcal{O}(X) < \infty$, without using the maximum principle. More generally the space of holomorphic sections of a complex line bundle over a compact space is finite-dimensional.)

Serre duality, special case. We can now use this finiteness to show finiteness of cohomology groups. (An alternative proof, again based on norms, is given in Forster.)

Theorem 7.10 *On any compact Riemann surface X, we have $H^1(X, \mathcal{O})^* \cong \Omega(X)$. In particular, $H^1(X, \mathcal{O})^*$ is finite-dimensional.*

We let $g_a = \dim \Omega(X)$, the *arithmetic genus* of X. We will eventually see that $g_a = g =$ the topological genus.

Proof. By Dolbeault we have

$$H^1(X, \mathcal{O}) \cong H^{0,1}(X) = \mathcal{E}^{(0,1)}(X) / \overline{\partial} \mathcal{E}^0(X).$$

We claim $\overline{\partial}\mathcal{E}^0(X)$ is closed in $\mathcal{E}^{(0,1)}(X)$ in the C^∞ topology. If not, there is a sequence $f_n \in \mathcal{E}^0(X)$ with $\overline{\partial} f_n \to \omega$ in the C^∞ topology, such that f_n has no convergent subsequence in $\mathcal{E}^0(X)$. Since bounded sets in $\mathcal{E}^0(X)$ are compact, the latter condition implies for some $k \geq 0$, $\|f_n\|_{C^k} \to \infty$. From this we will obtain a contradiction.

Dividing through by the C^k-norm, we can arrange that $\|f_n\|_{C^k} = 1$ in $\mathcal{E}^0/\mathbb{C}$ and $\overline{\partial} f_n \to 0$. Taking a bump function ρ in a chart, we have

$$\overline{\partial}(\rho f_n) = \rho(\overline{\partial} f_n) + (\overline{\partial}\rho) f_n.$$

Since f_n is bounded in C^k and $\overline{\partial} f_n$ tends to zero, the right-hand side is bounded in C^k. But since the solution to the $\overline{\partial}$ equation is given by convolution with $1/z$, a smoothing operator, we find that $\langle \rho f_n \rangle$ is precompact in C^k. Thus we can pass to a C^k convergent subsequence, $f_n \to g$. Then $\overline{\partial} g = 0$ so g is constant. But then $\|f_n\|_{C^k} \to 0$ in $\mathcal{E}^0(X)/\mathbb{C}$, a contradiction.

Since $\overline{\partial}\mathcal{E}^0(X)$ is closed, we have

$$H^1(X, \mathcal{O})^* = W \cong (\overline{\partial}\mathcal{E}^0(X))^\perp \subset (\mathcal{E}^{0,1}(X))^* = \mathcal{D}^{1,0}(X).$$

But any $(1,0)$-current $\omega \in W$ satisfies $\int \omega \wedge (\overline{\partial} f) = 0$ for all smooth f, and thus $\overline{\partial}\omega = 0$, which implies ω is holomorphic and thus $W = \Omega(X)$. ■

Corollary 7.11 *We have natural isomorphisms $H^{1,0}(X) \cong \Omega(X)$ and $H^{0,1}(X) \cong \overline{\Omega}(X)$.*

Proof. There is a natural isomorphism $\overline{\Omega}(X) \cong \Omega(X)^*$ coming from the pairing $\langle \alpha, \beta \rangle = \int_X \alpha \wedge \beta$, which is obviously nondegenerate since $\langle \alpha, \overline{\alpha} \rangle = \int |\alpha|^2$. ■

Taking stock. In preparation for Riemann-Roch, we now know $H^1(X, \mathcal{O})^* \cong \Omega(X)$, $g_a = \dim \Omega(X) \geq g$, and $\dim H^1(X, \Omega) \geq 1$ because of the residue map.

8 Riemann-Roch

One of the most basic questions about a compact Riemann surface X is: does there exist a nonconstant holomorphic map $f : X \to \widehat{\mathbb{C}}$? But as we have seen already in the discussion of the field $\mathcal{M}(X)$, it is desirable to ask more: for example, do the meromorphic functions on X separate points?

Even better, does there exist a holomorphic embedding $X \to \mathbb{P}^n$ for some n?

To answer these questions we aim to determine the *dimension* of the space of meromorphic functions with *controlled* zeros and poles. This is the Riemann-Roch problem.

Divisors. Let X be a compact Riemann surface. The group of *divisors* $\mathrm{Div}(X)$ is the free abelian group generated by the points of X. A divisor is sum $D = \sum a_P P$, where $a_P \in \mathbb{Z}$ and $a_P = 0$ for all but finitely many $P \in X$. We say $D \geq 0$ if $a_P \geq 0$ for all P. A divisor is *effective* if $D \geq 0$. Any divisor can be written unique as a difference of effective divisors, $D = D_+ - D_-$.

The *degree* of a divisor, $\deg(D) = \sum a_P$, defines a homomorphism $\deg : \mathrm{Div}(X) \to \mathbb{Z}$.

Note: one can also define a *sheaf* by $\mathrm{Div}(U) = \{\sum a_P P\}$ where the sum is *locally* finite. This is the *sheaf* canonically generated by the *presheaf* $\mathcal{M}^*/\mathcal{O}^*$.

Associated to any meromorphic function $f \neq 0$ is a divisor of degree zero recording its zeros and poles:

$$(f) = \sum_P \mathrm{ord}(f, P) \cdot P.$$

The divisors that arise in this way are said to be *principal.* Note that:

$$(fg) = (f) + (g),$$

so (f) defines a homomorphism from $\mathcal{M}^*(X)$ into $\mathrm{Div}_0(X)$. Note also that the *topological* degree of f is given by $\deg(f)_+$.

The sheaf $\mathcal{O}_D$ consists of meromorphic functions f such that $(f)+D \geq 0$. For example, $\mathcal{O}_{nP}(X)$ is the vector space of meromorphic functions on X with poles of order at most n at p.

Isomorphisms between sheaves. If $D - E = (f)$ is principal, we say D and E are *linearly equivalent.* Then the map $h \mapsto hf$ gives an isomorphism of sheaves:

$$\mathcal{O}_D = \mathcal{O}_{E+(f)} \cong \mathcal{O}_E,$$

because

$$(hf) + E = (h) + (f) + E = (h) + D.$$

The divisor $K = (\omega)$ of a nonzero meromorphic 1-form ω is defined similarly. Once we have such a *canonical* divisor, we get an isomorphism

$$\mathcal{O}_K \cong \Omega$$

by $h \mapsto h\omega$, because $h\omega$ is holomorphic iff

$$(h\omega) = (h) + K \geq 0$$

iff $h \in \mathcal{O}_K$.

Riemann-Roch Problem. The *Riemann-Roch problem* is to calculate or estimate

$$h^0(D) = \dim H^0(X, \mathcal{O}_D) = \dim \mathcal{O}_D(X).$$

Example: if X is a complex torus, we have $h^0(nP) = 1, 1, 2, 3, \ldots$. This can be explained by the fact that $\mathrm{Res}_P(f\,dz) = 0$.

It is a general principle that Euler characteristics are more stable than individual cohomology groups, and we define:

$$\chi(\mathcal{F}) = \sum (-1)^q h^q(\mathcal{F}).$$

We then have:

Theorem 8.1 (Riemann-Roch, Euler characteristic version) *For any divisor D, we have*

$$\chi(\mathcal{O}_D) = h^0(\mathcal{O}_D) - h^1(\mathcal{O}_D) = \deg D - g_a + 1.$$

Here $g_a = \dim \Omega(X)$ is the arithmetic genus.

For the proof we will use:

Theorem 8.2 *If $0 \to \mathcal{A} \to \mathcal{B} \to \mathcal{C} \to 0$ is an exact sequence of sheaves with finite-dimensional cohomology groups on a finite-dimensional space X, then we have:*

$$\chi(B) = \chi(A) + \chi(C).$$

Proof. Let the homomorphisms in dimension p for the resulting long exact sequence be denoted by α_p, β_p and δ_p. We then have:

$$\begin{aligned} \chi(\mathcal{A}) &= \sum (-1)^p (\dim \operatorname{Ker} \alpha_p + \dim \operatorname{Im} \alpha_p), \\ -\chi(\mathcal{B}) &= \sum (-1)^p (-\dim \operatorname{Ker} \beta_p - \dim \operatorname{Im} \beta_p), \quad \text{and} \\ \chi(\mathcal{C}) &= \sum (-1)^p (\dim \operatorname{Ker} \delta_p + \dim \operatorname{Im} \delta_p). \end{aligned}$$

By exactness we have:

$$\begin{aligned} 0 &= \sum(-1)^p(\dim \operatorname{Im} \alpha_p - \dim \operatorname{Ker} \beta_p), \\ 0 &= \sum(-1)^p(-\dim \operatorname{Im} \beta_p + \dim \operatorname{Ker} \delta_p), \quad \text{and} \\ 0 &= \sum(-1)^p(\dim \operatorname{Im} \delta_p - \dim \operatorname{Ker} \alpha_{p+1}), \end{aligned}$$

and thus $\chi(\mathcal{A}) - \chi(\mathcal{B}) + \chi(\mathcal{C}) = 0$. ∎

Skyscrapers. The *skyscraper sheaf* $\mathbb{C}_P$ is given by $\mathcal{O}(U) = \mathbb{C}$ if $P \in U$, and $\mathcal{O}(U) = 0$ otherwise. For any divisor D, we have the exact sequence:

$$0 \to \mathcal{O}_D \to \mathcal{O}_{D+P} \to \mathbb{C}_P \to 0, \tag{8.1}$$

where the final map records the leading coefficient of the polar part of f at P. It is easy to see (e.g. by taking fine enough coverings, without P in any multiple intersections):

Theorem 8.3 *We have $H^p(X, \mathbb{C}_P) = 0$ for all $p > 0$. In particular $\chi(\mathbb{C}_P) = h^0(\mathbb{C}_P) = 1$.*

Theorem 8.4 *The cohomology groups $H^p(X, \mathcal{O}_D)$ are finite-dimensional for $p = 0, 1$ and vanish for all $p \geq 2$.*

Proof. We have already seen the result is true for $D = 0$: the space $H^1(X, \mathcal{O}) \cong \Omega(X)^*$ is finite-dimensional, and using the Dolbeault sequence, one can show $H^p(X, \mathcal{O}) = 0$ for all $p \geq 2$. The result for general D then follows induction, using the skyscraper sheaf. ∎

Proof of Riemann-Roch. Since $h^1(\mathcal{O}) = \dim \Omega(X) = g_a$, the formula is correct for the trivial divisor. Using (8.1) and general properties of the Euler characteristic in long exact sequences, we find:

$$\chi(\mathcal{O}_{D+P}) = \chi(\mathcal{O}_D) + \chi(\mathbb{C}_P) = \chi(\mathcal{O}_D) + 1,$$

which implies Riemann-Roch for an arbitrary divisor D. ∎

Existence of meromorphic functions and forms. A more general form of Serre duality will lead to a more useful formulation of Riemann-Roch, but we can already deduce several useful consequences.

Theorem 8.5 *Any compact Riemann surface admits a nonconstant map $f : X \to \mathbb{P}^1$ with $\deg(f) \le g_a + 1$.*

Proof. We have $\dim H^0(X, \mathcal{O}_D) \ge \deg D - g_a + 1$. Once $\deg D > g_a$ this gives $\dim H^0(\mathcal{O}_D) > 1 = \dim \mathbb{C}$. ■

Corollary 8.6 *Any surface with $g_a = 0$ is isomorphic to $\mathbb{P}^1$.*

Proof. It admits a map to $\mathbb{P}^1$ of degree $g_a + 1 = 1$. ■

Corollary 8.7 *Any compact Riemann carries a nonzero meromorphic 1-form.*

Proof. Take $\omega = df$.

Corollary 8.8 *Canonical divisors $K = (\omega)$ exist, and satisfy $\mathcal{O}_K \cong \Omega$.*

The isomorphism is given by $f \mapsto f\omega$.

Arithmetic and topological genus; degree of canonical divisors; residues.

Corollary 8.9 *The degree of any canonical divisor is $2g - 2$, where g is the topological genus of X.*

Proof. Apply Riemann-Hurwitz to compute $\deg(df)$. ■

Corollary 8.10 *The topological and arithmetic genus of X agree: we have $g = g_a$; and $\dim H^1(X, \Omega) = 1$.*

Proof. Apply Riemann-Roch to a canonical divisor K: we get

$$h^0(K) - h^1(K) = g_a - h^1(K) = 1 - g_a + \deg(K) = 2g - g_a - 1,$$

or in other words:

$$2g_a - h^1(K) = 2g - 1.$$

Now we know $g_a \le g$ and $h^1(K) \ge 1$, so the only way equality can hold is if $g_a = g$ and $h^1(K) = 1$. ■

Corollary 8.11 *The residue map* $\mathrm{Res} : H^1(X, \Omega) \to \mathbb{C}$ *is an isomorphism.*

Theorem 8.12 (Hodge theorem) *On a compact Riemann surface every class in* $H^1_{DR}(X, \mathbb{C})$ *is represented by a harmonic 1-form. More precisely we have*

$$H^1_{DR}(X) = \Omega(X) \oplus \overline{\Omega(X)} = H^{1,0}(X) \oplus H^{0,1}(X) = \mathcal{H}^1(X).$$

Proof. We already know the harmonic forms inject into deRham cohomology, by considering their periods; since the topological and arithmetic genus agree, they also surject. ■

The space of smooth 1-forms. On a compact Riemann surface, the full Hodge theorem

$$\mathcal{E}^1(X) = d\mathcal{E}^0(X) \oplus \mathcal{H}^1(X) \oplus d^*\mathcal{E}^2(X)$$

becomes the statement:

$$\mathcal{E}^1(X) = (\partial + \overline{\partial})\mathcal{E}^0(X) \oplus (\Omega(X) \oplus \overline{\Omega}(X)) \oplus (\partial - \overline{\partial})\mathcal{E}^0(X).$$

Here $\partial - \overline{\partial}$ has the same image as $*d$, and hence the same image as d^*. This is a consequence of the isomorphism $H^{1,0}(X) \cong \Omega(X)$.

Remark: isothermal coordinates. The argument we have just given also proves the Hodge theorem for any compact, oriented Riemannian 2-manifold (X, g). To see this, however, we need to know that every Riemannian metric is locally conformally flat; i.e. that one can introduce 'isothermal coordinates' to make X into a Riemann surface with g a conformal metric.

Mittag-Leffler for 1-forms. The Mittag-Leffler problem is to construct meromorphic 1-form with prescribed principal parts.

The problem does not always have a solution. For example, if just a single simple pole is specified, then X would have to have genus 0 for f to exist.

The data is conveniently given by $\omega_i \in \mathcal{M}^{(1)}(U_i)$ such that $\delta\omega = (\delta\omega)_{ij}$ lies in $Z^1(\Omega)$, i.e. such that $\alpha_{ij} = \omega_j - \omega_i$ is holomorphic. Then the problem of constructing a global ω with the same principal parts reduces to showing that $(\delta\omega)$ is a coboundary in $H^1(X, \Omega)$.

We are thus lead to consider that cohomology group. Since we have shown that $h^1(K) = 1$ (above), we have:

Theorem 8.13 *There exists a meromorphic 1-form with prescribed principal parts iff the sum of its residues is equal to zero.*

Proof. The principal parts are specified by a Mittag-Leffler 0-cochain $\omega_i \in \mathcal{M}^{(1)}(U_i)$ with boundary

$$\omega_i - \omega_j = \alpha_{ij} \in \Omega(U_{ij});$$

and the class $\xi = [\alpha_{ij}] \in H^1(X, \Omega)$ is equal to zero iff $\mathrm{Res}(\xi) = 0$, in which case $\alpha_{ij} = \beta_i - \beta_j$ with $\beta_i \in \Omega(U_i)$, and $\omega_i + \beta_i = \eta$ defines a global meromorphic 1-form with the prescribed principal parts. ■

Corollary 8.14 *The Mittag-Leffler problem for 1-forms has a solution if and only if the sum of the residues of the principal parts is zero.*

Corollary 8.15 *Given any pair of distinct points $p_1, p_2 \in X$, there is a meromorphic 1-form ω with simple poles of residues $(-1)^i$ at p_i and no other singularities.*

This 'elementary differential of the third kind' is unique up to the addition of a global holomorphic differential. Example: the form dz/z works for $0, \infty \in \widehat{\mathbb{C}}$.

Corollary 8.16 *For any $p \in X$ and $n \geq 2$ there exists a meromorphic 1-form ω with a pole of order n at p (but vanishing residue) and no other singularities.*

This is an 'elementary differential of the second kind'.

Currents and residues. Here is another formulation of the Mittag-Leffler problem for 1-forms. Using sheaves of distributions and currents, we have an exact sequence

$$0 \to \Omega \to \mathcal{D}^{1,0} \xrightarrow{\bar{\partial}} \mathcal{D}^{1,1} \to 0,$$

which shows the isomorphism $H^1(X, \Omega) \cong \mathbb{C}$ can be computed using currents instead of smooth forms. Then a current in $\mathcal{D}^{1,1}$ representing the Mittag-Leffler cocycle $\omega_i - \omega_j$ is given by $\eta = d\omega_i$, which is supported just at the poles of ω_i and satisfies

$$\frac{1}{2\pi i} \int \eta = \mathrm{Res}(\omega_i).$$

Thus $\delta\omega_i = 0$ iff $\mathrm{Res}(\omega_i) = 0$.

9 Serre duality

Let $\mathcal{M}^{(1)}(X)$ denote the space of *meromorphic* 1-forms on X. It is a 1-dimensional vector space over the field $\mathcal{M}(X)$. We define:

$$\Omega_D(X) = \{\omega \in \mathcal{M}^{(1)}(X) \,:\, (\omega) + D \geq 0\}.$$

The sheaf $\Omega_D \cong \mathcal{O}_{D+K}$ is defined similarly. The goal of this section is to present:

Theorem 9.1 (Serre duality) *For any divisor D, we have a canonical isomorphism*

$$H^1(X, \mathcal{O}_{-D})^* \cong \Omega_D(X).$$

This result allows us to eliminate h^1 entirely from the statement of Riemann-Roch, and obtain:

Theorem 9.2 (Riemann-Roch, geometric version) *For any divisor D on a compact Riemann surface X, we have*

$$h^0(D) = \deg D - g + 1 + h^0(K - D).$$

Motivation: Mittag-Leffler for functions. Given a finite set of points $p_i \in X$, and the Laurent tails

$$f_i(z) = \frac{b_n}{z^n} + \cdots + \frac{b_1}{z}$$

of meromorphic functions f_i in local coordinates near P_i, when can we find a global meromorphic function f on X with the given principal parts?

Case $\boldsymbol{D \geq 0}$ and $\boldsymbol{K - D \geq 0}$. Let us consider first the simple case where $f_i(z) = a_i/z$ near P_i. Let $D = \sum P_i$. Clearly the data $(a_i)_1^d$, $d = \deg D$, determine f uniquely up to a constant; thus we have:

$$h^0(D) \leq \deg(D) + 1.$$

However, each holomorphic 1-form imposes a linear constraint: $\sum \operatorname{Res}(f_i\omega) = 0$. These conditions would reduce $h^0(D)$ by $g = \dim \Omega(X)$ if they were linearly independent. However, the forms in $\Omega_{-D}(X)$ vanish at all the points P_i and hence impose no conditions. Thus we get:

$$h^0(D) \leq \deg(D) + 1 - g + h^0(K - D).$$

Now if $K - D$ is also effective, we can interchange the roles of D and $K - D$ to obtain:

$$h^0(K-D) \le \deg(K-D) + 1 - g + h^0(D) = g - 1 - \deg(D) + h^0(D).$$

Summing these two equations, we get:

$$h^0(K-D) + h^0(D) \le h^0(K-D) + h^0(D).$$

Since in fact equality holds here, we conclude it held before, and thus we have the following special case of the geometric version of Riemann-Roch:

Theorem 9.3 *If both D and $K - D$ are effective divisors, we have:*

$$h^0(D) = \deg(D) + 1 - g + h^0(K-D).$$

Case $D \ge 0$. Returning to the original Mittag-Leffler problem, suppose we are given $f_i \in \mathcal{M}(U_i)$. Then the problem is to find $f \in \mathcal{M}(X)$ with the same principal parts.

Equivalently, we want to determine when $\delta f_i \in H^1(X, \mathcal{O})$ is a coboundary. By the case of Serre duality we have already proven, $H^1(X, \mathcal{O})$ is isomorphic to $\Omega(X)^*$, so there is a natural pairing between $(\delta f_i) \in H^1(X, \mathcal{O})$ and $\omega \in \Omega(X)$.

Recalling from Theorem 7.9 that (up to a constant) the pairing defining Serre duality is given by

$$\langle \delta f_i, \omega \rangle = \mathrm{Res}(\omega f_i),$$

we then have:

Corollary 9.4 *The Mittag-Leffler problem specified by (f_i) has a solution iff*

$$\sum \mathrm{Res}_p(f_i \omega) = 0$$

for every $\omega \in \Omega(X)$.

This proves the final Riemann-Roch theorem for *effective* divisors. For example, suppose $D = nP$. Then an element $f \in \mathcal{O}_D(X)$ is determined by its Laurent tail

$$f(z) = \frac{b_n}{z^n} + \cdots + \frac{b_1}{z} + b_0$$

in a coordinate system where $z(p) = 0$. The set of (b_i) that can arise is determined by the constraint $\mathrm{Res}(\omega f) = 0$ for all $\omega \in \Omega(X)$. But the residue vanishes trivially if ω belongs to $\Omega_{-nP}(X)$. Thus the number of constraints on $(b_0, \ldots, b_n)$ is $\dim \Omega(X) - \dim \Omega_{-nP}(X) = g - h^0(K - nP)$, and we find

$$h^0(nP) = n + 1 - g + h^0(K - nP)$$

in agreement with Riemann-Roch.

Proof of Serre duality: dimension counts. A proof of Serre duality can be given using the elliptic regularity, just as we did for the case $\Omega(X) \cong H^1(X, \mathcal{O})^*$. For more perspective we give a different argument, based loosely on Forster.

We begin with some useful qualitative dimension counts that follow immediately from the fact that $h^0(D) = 0$ if $\deg(D) < 0$, $h^1(D) \geq 0$ and the fact that $\Omega \cong \mathcal{O}_K$.

Theorem 9.5 *For* $\deg(D) > 0$*, we have*

$$\begin{aligned}
\dim \mathcal{O}_D(X) &\geq \deg(D) + O(1), \\
\dim \Omega_D(X) &\geq \deg(D) + O(1), \quad \textit{and} \\
\dim H^1(X, \mathcal{O}_{-D}) &= \deg(D) + O(1).
\end{aligned}$$

Pairings. Next we couple the product map

$$\mathcal{O}_{-D} \otimes \Omega_D \to \Omega$$

together with the residue map $\operatorname{Res} : H^1(X, \Omega) \to \mathbb{C}$ to obtain a map

$$H^1(X, \mathcal{O}_{-D}) \otimes \Omega_D(X) \to H^1(X, \Omega) \to \mathbb{C},$$

or equivalently a natural map

$$\Omega_D(X) \to H^1(X, \mathcal{O}_{-D})^*.$$

This map explicitly sends ω to the linear functional defined by

$$\phi(\xi) = \operatorname{Res}(\xi\omega).$$

Using the long exact sequence associated to equation (8.1) we obtain:

Theorem 9.6 *The inclusion* $\mathcal{O}_D \to \mathcal{O}_{D+P}$ *induces a surjection:*

$$H^1(X, \mathcal{O}_D) \to H^1(X, \mathcal{O}_{D+P}) \to 0.$$

Corollary 9.7 *We have a natural surjective map:*

$$H^1(X, \mathcal{O}_D) \to H^1(X, \mathcal{O}_E) \to 0$$

for any $E \geq D$.

Put differently, for $E \geq D$ we get a surjective map $H^1(X, \mathcal{O}_{-E}) \to H^1(X, \mathcal{O}_{-D})$ and thus an injective map on the level of duals. For organizational convenience we take the direct limit over increasing divisors and set

$$V(X) = \lim_{D \to +\infty} H^1(X, \mathcal{O}_{-D})^*.$$

Clearly $\Omega_D(X)$ maps into $V(X)$ for every D, so we get a map $\mathcal{M}^{(1)}(X) \to V(X)$.

Theorem 9.8 *The natural map $\mathcal{M}^{(1)}(X) \to V(X)$ is injective. Moreover a meromorphic form ω maps into $H^1(X, \mathcal{O}_{-D})^*$ iff $\omega \in \Omega_D(X)$.*

Proof. Both statements are easy to prove, because it is easy to produce elements of $H^1(X, \mathcal{O}_D)$. For the first, suppose $\omega \in \Omega_D(X)$ is a nonzero meromorphic form. Pick any point $P \in X$ with local coordinate $z : U_1 \to \Delta$, and let $U_2 = X - \{P\}$. Choose U_1 small enough so that the divisors D and (ω) have no points in $U_1 - \{P\}$. Then choose

$$k = -1 - \mathrm{ord}_P(\omega),$$

so that $\mathrm{Res}_P(z^k \omega) \neq 0$. Set $f_1 = z^k$ on U_1 and $f_2 = 0$ on U_2, and $\xi = (g_{ij}) = f_1 - f_2$. Then $\xi \in H^1(X, \mathcal{O}_{-D})$, and

$$\mathrm{Res}(\xi\omega) = \mathrm{Res}_P(z^k \omega) \neq 0.$$

This shows ω defines a nonzero element of $V(X)$.

The proof of the second statement is similar. If ω is in $H^1(X, \mathcal{O}_{-D})^*$, then it must vanish on all coboundaries for this group. Suppose however $-(k+1) = \mathrm{ord}_P(\omega) < -D(P)$ for some P. Then $k \geq D(P)$, so in the construction above we can arrange that $\xi = 0$ in $H^1(X, \mathcal{O}_{-D})$. This is a contradiction. Thus $\mathrm{ord}_P(\omega) \geq -D(P)$ for all P, i.e. $\omega \in \Omega_D(X)$. ∎

Completion of the proof of Serre duality. Note that both $\mathcal{M}^{(1)}(X)$ and $V(X)$ are vector spaces over the field of meromorphic functions $\mathcal{M}(X)$. (Indeed the former vector space is one-dimensional, generated by any meromorphic 1-form.)

Given $\phi \in H^1(X, \mathcal{O}_{-D})^* \subset V(X)$, we must show ϕ is represented by a meromorphic 1-form ω. (By the preceding result, ω will automatically lie in $\Omega_D(X)$.)

The proof will be by a dimension count. We note that for $n \gg 0$, we have

$$\dim H^1(X, \mathcal{O}_{-D-nP})^* = n + O(1).$$

On the other hand, this space contains $\Omega_{D+nP}(X)$ as well as $\mathcal{O}_{nP}(X) \cdot \phi$. Both of these spaces have dimension bounded below by $n + O(1)$. Thus for n large enough, we can write $f\phi = \omega$ where f is a meromorphic function and ω is a meromorphic form. But then $\phi = \omega/f$ is also a meromorphic form! ■

We can now round out the discussion by proving some results promised above.

Theorem 9.9 *For any divisor D we have*

$$H^0(X, \mathcal{O}_D) \cong H^1(X, \Omega_{-D})^*.$$

Proof. We have

$$H^0(X, \mathcal{O}_D) \cong H^0(X, \Omega_{D-K}) \cong H^1(X, \mathcal{O}_{K-D})^* \cong H^1(X, \Omega_{-D})^*$$

■

Corollary 9.10 *We have $H^1(X, \mathcal{O}_D) = 0$ as soon as $\deg(D) > \deg(K) = 2g - 2$.*

Proof. Because then $H^1(X, \mathcal{O}_D)^* \cong \Omega_{-D}(X) = 0$. ■

Corollary 9.11 *If $\deg(D) > 2g - 2$, then $h^0(D) = 1 - g + \deg(D)$.*

Corollary 9.12 *We have $H^1(X, \mathcal{M}) = H^1(X, \mathcal{M}^{(1)}) = 0$.*

Proof. Any representative cocycle (f_{ij}) for a class in $H^1(X, \mathcal{M})$ can be regarded as a class in $H^1(X, \mathcal{O}_D)$ for some D of large degree. But $H^1(X, \mathcal{O}_D) = 0$ once $\deg(D)$ is sufficiently large. Thus (f_{ij}) splits for the sheaf $\mathcal{O}_D$, and hence for $\mathcal{M}$. ■

Corollary 9.13 *Every element of $H^1(X, \mathcal{O}_D)$ can be represented as a Mittag-Leffler coboundary, $g_{ij} = f_i - f_j$ with meromorphic (f_i). Similarly for $H^1(X, \Omega_D)$.*

Mittag-Leffler for 1-forms, revisited. Here is another proof of the Mittag-Leffler theorem for 1-forms. Suppose we consider all possible principal parts with poles of order at most $n_i > 0$ at P_i, and let $D = \sum n_i P_i$. The dimension of the space of principal parts is then $n = \sum n_i = \deg D$. In addition, the principal part determines the solution to the Mittag-Leffler problem up to adding a holomorphic 1-form. That is, the solutions lie in the space $\Omega_D(X)$, and the map to the principal parts has $\Omega(X)$ as its kernel.

Thus the dimension of the space of principal parts that have solutions is:

$$k = \dim \Omega_D(X) - \dim \Omega(X).$$

But by Riemann-Roch we have

$$\begin{aligned} \dim \Omega_D(X) &= h^0(K + D) = h^0(-D) + \deg(K + D) - g + 1 \\ &= 2g - 2 + \deg D - g - 1 = g + \deg D - 1 = g + n - 1, \end{aligned}$$

so the solvable data has dimension $k = n - 1$. Thus there is one condition on the principal parts for solvability, and that condition is given by the residue theorem.

10 Maps to projective space

In this section we explain the connection between the sheaves $\mathcal{O}_D$, linear systems and maps to projective space.

Projective space. Let V be an $(n + 1)$-dimensional vector space over $\mathbb{C}$. The space of lines (one-dimensional subspaces) in V forms the *projective space*

$$\mathbb{P}V = (V - \{0\})/\mathbb{C}^*.$$

It has the structure of a complex n-manifold. The subspaces $S \subset V$ give rise to planes $\mathbb{P}S \subset \mathbb{P}V$; when S has codimension one, $\mathbb{P}S$ is a *hyperplane*.

The *dual* projective space $\mathbb{P}V^*$ parameterizes the hyperplanes in $\mathbb{P}V$, via the correspondence $\phi \in V^* \mapsto \mathrm{Ker}(\phi) = S \subset V$.

We can also form the quotient space $W = V/S$. Any line L in V that is not entirely contained in S projects to a line in W. Thus we obtain a natural map

$$\pi : (\mathbb{P}V - \mathbb{P}S) \to \mathbb{P}(V/S).$$

All the analytic automorphisms of projective space come from linear automorphisms of the underlying vector space: that is,

$$\mathrm{Aut}(\mathbb{P}V) = \mathrm{GL}(V)/\mathbb{C}^* = \mathrm{PGL}(V).$$

We let $\mathbb{P}^n = \mathbb{P}\mathbb{C}^{n+1}$ with homogeneous coordinates $[Z] = [Z_0 : \cdots : Z_n]$. It satisfies

$$\mathrm{Aut}(\mathbb{P}^n) = \mathrm{PGL}_{n+1}(\mathbb{C}).$$

The case $n = 1$ gives the usual identification of automorphisms of $\widehat{\mathbb{C}}$ with Möbius transformations.

Projective varieties. The zero set of a homogeneous polynomial $f(Z)$ defines an *algebraic hypersurface* $V(f) \subset \mathbb{P}^n$. A *projective algebraic variety* is the locus $V(f_1, \ldots, f_n)$ obtained as an intersection of hypersurfaces.

The ratio

$$F(Z) = f_1(Z)/f_0(Z) = [f_0(Z) : f_1(Z)]$$

of two homogeneous polynomials of the same degree defines a meromorphic 'function'

$$F : \mathbb{P}^n \dashrightarrow \mathbb{P}^1.$$

Its values are undetermined on the subvariety $V(f_0, f_1)$, which has codimension two if f_0 and f_1 are relatively prime.

The Hopf fibration. By considering the unit sphere in $\mathbb{C}^{n+1}$, we obtain the *Hopf fibration* $\pi : S^{2n+1} \to$
$cxproj^n$ with fibers S^1. This shows that projective space is compact. Moreover, for $n = 1$ the fibers of π are linked circles in S^3, and π generates $\pi_3(S^2) \cong \mathbb{Z}$.

Affine space. The locus $\mathbb{A}^n = \mathbb{P}^n - V(Z_0)$ is isomorphic to $\mathbb{C}^n$ with coordinates $(z_1, \ldots, z_n) = Z_i/Z_0$, while $V(Z_0)$ itself is a hyperplane; thus we have

$$\mathbb{P}^n \cong \mathbb{C}^n \cup \mathbb{P}^{n-1}.$$

By permuting the coordinates, we get a covering of $\mathbb{P}^n$ by $n+1$ affine charts.

Any ordinary polynomial $p(z_i)$ has a unique homogeneous version $P(Z_i)$ of the same maximal degree, such that $p(z_i) = P(1, z_1, \ldots, z_n)$. Thus any affine variety $V(p_1, \ldots, p_m)$ has a natural completion $V(P_1, \ldots, P_m) \subset \mathbb{P}^m$. This variety is *smooth* if it is a smooth submanifold in each affine chart.

Examples: Curves in $\mathbb{P}^2$.

1. The affine curve $x^2 - y^2 = 1$ meets the line at infinity in two points corresponding to its two asymptotes; its homogenization $X^2 - Y^2 = Z^2$

is smooth in every chart. In (y,z) coordinates it becomes $1 - y^2 = z^2$, which means the line at infinity $(z = 0)$ in two points.

2. The affine curve defined by $p(x,y) = y - x^3 = 0$ is smooth in $\mathbb{C}^2$, but its homogenization $P(X,Y,Z) = YZ^2 - X^3$ defines the curve $z^2 = x^3$ in the affine chart where $Y \neq 0$, which has a cusp.

3. The space of homogeneous polynomials of degree d in $\mathbb{P}^n$ has dimension given by
$$\dim P_d(\mathbb{C}^{n+1}) = \binom{d+n}{n}.$$
This can be seen by inserting n markers into a list of d symbols, and turning all the symbols up to the first marker into Z_0's, then the next stretch into Z_1's, etc. Thus the space of curves of degree d in $\mathbb{P}^2$ is itself a projective space $\mathbb{P}^N$, with $N = d(d+3)/2$. E.g. there is a 5-dimensional space of conics, a 9-dimensional space of cubics and a 14-dimensional space of quartics.

The degree of a plane curve. If f is irreducible, then $V(f)$ meets a typical line in exactly $d = \deg(f)$ points. Thus the degree of f is a visible property of $V(f)$.

If $f = f_1 \cdots f_n$ is a product of distinct irreducibles, then $V(f) = \bigcup V(f_i)$. These are also important examples of curves of degree d. For example, d distinct lines form a curve of degree d.

The normalization of a singular curve. Every *irreducible* homogeneous polynomial f on $\mathbb{P}^2$ determines a compact Riemann surface X together with a generically injective map $\nu : X \to V(f)$. The Riemann surface X is called the *normalization* of the (possibly singular) curve $V(f)$.

To construct X, use projection from a typical point $P \in \mathbb{P}^2 - V(f)$ to obtain a surjective map $\pi : V(f) \to \mathbb{P}^1$. After deleting a finite set from domain and range, including all the singular points of $V(f)$, we obtain an *open Riemann surface* $X^* = V(f) - C$ and a degree d *covering map*
$$\pi : X^* \to (\mathbb{P}^1 - B).$$
As we have seen, there is a canonical way to complete X^* and π to a compact Riemann surface X and a branched covering $\pi : X \to \mathbb{P}^1$. It is then easy to see that the isomorphism
$$\nu : X^* \to V(f) - C \subset \mathbb{P}^2$$
extends to a holomorphic map $\nu : X \to \mathbb{P}^2$ with image $V(f)$.

Example. The cuspidal cubic $y^2 = x^3$ is normalized by $\nu : \mathbb{P}^1 \to \mathbb{P}^2$ given by $\nu(t) = (t^2, t^3)$.

Maps to $\mathbb{P}^n$. Two divisors D, E are *linearly equivalent* if $D - E = (f)$ is principal; then $\mathcal{O}_D \cong \mathcal{O}_E$.

A divisor determines a natural map (if $n + 1 = h^0(D) > 0$):

$$\phi_D : X \to \mathbb{P}H^0(X, \mathcal{O}_D)^* \cong \mathbb{P}^n$$

by $\phi_D(x)(f) = f(x)$. Here the map should be renormalized near points x where elements of $H^0(X, \mathcal{O}_D)$ have zeros or poles, by dividing through by common zeros or poles.

Theorem 10.1 *The image $\phi_D(X)$ is not contained in any hyperplane.*

Linear systems and the base locus. The *linear system* $|D|$ determined by D is the collection of *effective* divisors E linearly equivalent to D. These are exactly the divisors of the form $E = (f) + D$, where $f \in H^0(X, \mathcal{O}_D)$.

The divisor E measures the extent to which $\mathcal{O}_x(f)$ *exceeds* $-D(x)$, i.e. it describes the points where f vanishes to higher order than necessary. We have a natural bijection

$$|D| = \mathbb{P}\mathcal{O}_D(X).$$

Equivalent divisors determine the *same* linear system.

Zeros and poles versus linear systems. Note: although the elements of $\mathcal{O}_D(X)$ are meromorphic functions, there is a profound shift in perspective when we pass to linear systems. Namely we no longer focus on the zeros and poles of f, but their *excess* $E = (f) + D$. This shift in perspective will eventually be made more systematic using the concept of a *line bundle*.

Base locus. The *base locus* B of a linear system $|D|$ is the largest divisor such that $0 \le B \le E$ for all $E \in |D|$. We say $|D|$ is *base-point free* if $B = 0$. This just means for all $P \in X$ there is an $E \in |D|$ which is supported in $X - \{P\}$.

We say $\mathcal{O}_D$ is *generated by global sections* if for each $x \in X$ we have a global section $f \in H^0(X, \mathcal{O}_D)$ such that the stalk $\mathcal{O}_{D,x}$, which is an $\mathcal{O}_x$-module, is generated by f; that is, $\mathcal{O}_{D,x} = \mathcal{O}_x \cdot f$.

Theorem 10.2 *The following are equivalent.*

1. *$|D|$ is base-point free.*

2. *$h^0(D - P) = h^0(D) - 1$ for all $P \in X$.*

3. $\mathcal{O}_D$ *is generated by global sections.*

4. *For any* $x \in X$ *there is an* $f \in \mathcal{O}_D(X)$ *such that* $\operatorname{ord}_x(f) = -D(x)$.

Example. Note that $h^0(nP) = 1$ for $n = 0$ and $= g$ for $n = 2g - 1$. Thus (for large genus) there must be values of n such that $h^0(nP) = h^0(nP - P)$. In this case, $D = nP$ is *not* globally generated.

Hyperplane sections.

Theorem 10.3 *If* D *is base-point free, then* $|D|$ *consists of the hyperplane sections* $\phi_D^{-1}(H)$.

Proof. Let $f_0, f_1, \ldots, f_n$ be a basis for $\mathcal{O}_D(X)$. Then the map to $\mathbb{P}^n$ is given locally near $p \in X$ by

$$\phi_D(z) = [z^d f_0(z) : z^d f_1(z) : \cdots : z^d f_n(z)]$$

where z is a local coordinate, $z(p) = 0$, and d is the maximum order of pole at p of the $f_i(z)$.

Since $\mathcal{O}_D$ is globally generated, we have $d = D(p)$. Now p belongs to the hyperplane at infinity $H = (Z_0 = 0)$, determined by coordinates $[Z_0 : \cdots : Z_n]$ on $\mathbb{P}^n$, if and only if $z^d f_0(p) = 0$. But this means exactly that $\operatorname{ord}_p(f_0) > -D(p)$. Therefore the divisor $E = (f_0) + D$ coincides with $\phi_D^{-1}(H)$ (where the preimage is counted with multiplicity).

Conversely, if $E = (f) + D \geq 0$, then f can be taken as a basis of element of $\mathcal{O}_D(X)$, determining in turn a hyperplane section giving E. ■

Corollary 10.4 *If* $|D|$ *is base-point free and* ϕ_D *is an embedding, then* $\phi_D(X)$ *is a curve of degree* $\deg D$.

Theorem 10.5 *Let* $\phi : X \to \mathbb{P}^n$ *be a map of* X *to projective space with the image not contained in a hyperplane. Then* $\phi = \pi \circ \phi_D$, *where* D *is the divisor of a hyperplane section,* $|D|$ *is a base-point free linear system, and where* $\pi : \mathbb{P}^N \dashrightarrow \mathbb{P}^n$ *is projection from a linear subspace* $\mathbb{P}S$ *disjoint from* $\phi_D(X)$.

Proof. Since hyperplanes can be moved, the linear system $|D|$ is base-point free, and all hyperplane sections are linearly equivalent to D. Thus ϕ can be regarded as the natural map from X to $\mathbb{P}S^*$, where S is a subspace of $H^0(X, \mathcal{O}_D)$, so ϕ can be factored through the map ϕ_D to $\mathbb{P}H^0(X, \mathcal{O}_D)^*$. ■

Embeddings into projective space.

Theorem 10.6 *Let $|D|$ be base-point free. Then for any effective divisor $P = P_1 + \ldots + P_n$ on X, the linear system*

$$P + |D - P| \subset |D|$$

consists exactly of the hyperplane sections $E = \phi_D(X) \cap H$ passing through $(P_1, \ldots, P_n)$. In particular, the dimension of the space of such hyperplanes is $\dim |D - P|$.

Proof. A hyperplane section $E \in |D|$ passes through P iff $E - P \geq 0$ iff $E = P + E'$ where $E' \in |D - P|$. ∎

Theorem 10.7 *If $h^0(D - P - Q) = h^0(D) - 2$ for any $P, Q \in X$, then $|D|$ provides a smooth embedding of X into projective space.*

Proof. This condition says exactly that the set of hyperplanes passing through $\phi_D(P)$ and $\phi_D(Q)$ has dimension 2 less than the set of all hyperplanes. Thus $\phi_D(P) \neq \phi_D(Q)$, so ϕ_D is $1 - 1$. The condition on $\phi(D - 2P)$ says that the set of hyperplanes containing $\phi_D(P)$ and $\phi_D'(P)$ is also 2 dimensions less, and thus ϕ_D is an immersion. Thus ϕ_D is a smooth embedding. ∎

Theorem 10.8 *The linear system $|D|$ is base-point free if* $\deg D \geq 2g$, *and gives an embedding into projective space if* $\deg D \geq 2g + 1$.

Proof. Use the fact that if $\deg D > \deg K = 2g - 2$, then we have $h^0(D) = \deg D - g + 1$, which is linear in the degree. ∎

Corollary 10.9 *Every compact Riemann surface embeds in projective space.*

Remark. By projecting we get an embedding of X into $\mathbb{P}^3$ and an immersion into $\mathbb{P}^2$.

Not every Riemann surface can be embedded into the plane! In fact a smooth curve of degree d has genus $g = (d - 1)(d - 2)/2$, so for example there are no curves of genus 2 embedded in $\mathbb{P}^2$.

Examples of linear systems.

1. *Genus 0.* On $\mathbb{P}^1$, we have $\mathcal{O}_D \cong \mathcal{O}_E$ if $d = \deg(D) = \deg(E)$. This sheaf is usually referred to as $\mathcal{O}(d)$; it is well-defined up to isomorphism, $|D|$ consists of all effective divisors of degree d. Note that $\mathcal{O}_{d\infty}(\mathbb{P}^1)$ is the $d+1$-dimensional space of polynomials of degree d. The corresponding map $\phi_D : \mathbb{P}^1 \to \mathbb{P}^d$ is given in affine coordinates by $\phi_D(t) = (t, t^2, \ldots, t^d)$. Its image is the *rational normal curve* of degree d. Particular cases are smooth conics and the twisted cubic.

2. *Genus 1.* On $X = \mathbb{C}/\Lambda$ with $P = 0$, the linear system $|P|$ is *not* base-point free, but $|2P|$ is, and $|3P|$ gives an embedding into the plane, via the map $z \mapsto (\wp(z), \wp'(z))$.

   Recall that $D = \sum m_i P_i$ is a principal divisor on X iff $\deg(D) = 0$ and $e(D) = \sum m_i P_i$ in $\mathbb{C}/\Lambda$ is zero. Thus $3Q \in |3P|$ iff $e(3(P-Q)) = 0$. This shows:

   > *A smooth cubic curve X in $\mathbb{P}^2$ has 9 flex points, corresponding to the points of order 3 in the group law on X.*

The canonical map. Note that the embedding of a curve X of genus 1 into $\mathbb{P}^2$ by the linear series $|3P|$ *breaks* the symmetry group of the curve: since $\mathrm{Aut}(X)$ acts transitively, the 9 flexes of $Y = \phi_{3P}(X)$ are not intrinsically special.

For genus $g \geq 2$ on the other hand there a more natural embedding — one which does not break the symmetries of X — given by the canonical linear system $|K|$.

The linear system $|K|$ corresponds to the map

$$\phi_K : X \to \mathbb{P}\Omega(X)^* \cong \mathbb{P}^{g-1}$$

given by $\phi(x) = [\omega_1(x), \ldots, \omega_g(x)]$, where ω_i is a basis for $\Omega(X)$.

Theorem 10.10 *The linear system $|K|$ is base-point free.*

Proof. Since X is not isomorphic to $\mathbb{P}^1$, we have $h^0(P) = h^0(0) = 1$. Thus

$$h^0(K-P) = h^0(P) + \deg(K-P) - g + 1 = h^0(K) - 1.$$

In other words, for any $P \in X$ we have $\omega(P) \neq 0$ for some $\omega \in \Omega(X)$. ■

Theorem 10.11 *Either $|K|$ gives an embedding of X into $\mathbb{P}^{g-1}$, or X is hyperelliptic.*

Proof. Suppose $|K|$ does not given an embedding. Then $h^0(K-P-Q) > h^0(K) - 2$ for some $P, Q \in X$, and thus $h^0(P+Q) > 1$. Thus there exists meromorphic function f with polar divisor $P+Q$, providing a degree two map $f : X \to \mathbb{P}^1$.

Now suppose X is hyperelliptic. Then there is a degree two holomorphic map $f : X \to \mathbb{P}^1$ branched over the zeros of a polynomial $p(z)$ of degree $2g+2$. A basis for the holomorphic 1-forms on X is given by

$$\omega_i = \frac{z^i\, dz}{\sqrt{p(z)}}$$

for $i = 0, \ldots, g-1$. That is, $\omega_i(x) = f(x)^i \omega_0$. It follows that the canonical map $\phi : X \to \mathbb{P}^{g-1}$ is given by

$$f(x) = [\omega_i(x)] = [\omega_0(x) f(x)^i] = [f(x)^i].$$

In other words, the canonical map factors as $\phi = \psi \circ f$, where $\psi : \mathbb{P}^1 \to \mathbb{P}^{g-1}$ is the rational normal curve of degree $g-1$. ∎

Canonical curves of genus two. We now describe more geometrically the canonical curves of genus two, three and four.

Theorem 10.12 *Any Riemann surface X of genus 2 is hyperelliptic, and any degree two map of X to $\mathbb{P}^1$ agrees with the canonical map (up to $\operatorname{Aut} \mathbb{P}^1$).*

Proof. In genus 2, we have $\deg K = 2g - 2 = 2$, so the canonical map $\phi : X \to \mathbb{P}^{g-1} = \mathbb{P}^1$ already presents X as a hyperelliptic curve. If $f : X \to \mathbb{P}^1$ is another such map, with polar divisor $P+Q$, then we have $h^0(P+Q) = 2 = h^0(K-P-Q) + 2 - 2 + 1$; thus there exists an ω with zeros just at P, Q and therefore $P+Q$ is a canonical divisor. ∎

Corollary 10.13 *The moduli space of curves of genus two is isomorphic to the 3-dimensional space $\mathcal{M}_{0,6}$ of isomorphism classes of unordered 6-tuples of points on $\mathbb{P}^1$. Thus $\mathcal{M}_2$ is finitely covered by $\mathbb{C}^3 - D$, where D consists of the hyperplanes $x_i = 0$, $x_i = 1$ and $x_i = x_j$.*

Remark. Here is a topological fact, related to the fact that every curve of genus two is hyperelliptic: if S has genus two, then the center of the mapping-class group $\operatorname{Mod}(S)$ is $\mathbb{Z}/2$, generated by any hyperelliptic involution. (In higher genus the center of $\operatorname{Mod}(S)$ is trivial.)

Canonical curves of genus 3. Let X be a curve of genus 3. Then X is either hyperelliptic, or its canonical map realizes it as a smooth plane quartic. We will later see that, conversely, any smooth quartic is a canonical curve (we already know it has genus 3). This shows:

Theorem 10.14 *The moduli space of curves of genus 3 is the union of the 5-dimensional space $\mathcal{M}_{0,8}$ and the 6-dimensional moduli space of smooth quartics, $(\mathbb{P}^{14} - D)/\operatorname{PGL}_3(\mathbb{C})$.*

Note: a smooth quartic curve that degenerates to a hyperelliptic one becomes a double conic. The eight hyperelliptic branch points can be thought of as the intersection of this conic with an infinitely near quartic curve.

Canonical curves of genus 4. Now let $X \subset \mathbb{P}^3$ be a canonical curve of genus 4 (in the non-hyperelliptic case). Then X has degree 6.

Theorem 10.15 *X is the intersection of an irreducible quadric and cubic hypersurface in $\mathbb{P}^3$.*

Proof. The proof is by dimension counting again. There is a natural linear from $\operatorname{Sym}^2(\Omega(X))$ into $H^0(X, \mathcal{O}_{2K})$. Since $\dim \Omega(X) = 4$, the first space has dimension $\binom{3+2}{2} = 10$, while the second has dimension $3g - 3 = 9$ by Riemann-Roch. Thus there is a nontrivial quadratic equation $Q(\omega_1, \ldots, \omega_4) = 0$ satisfied by the holomorphic 1-forms on X; equivalent X lies on a quadric. The quadric is irreducible because X does *not* lie on a hyperplane.

Carrying out a similar calculation for degree 3, we find $\dim \operatorname{Sym}^3(\Omega(X)) = \binom{3+3}{3} = 20$ while $\dim H^0(X, \mathcal{O}_{3K}) = 5g - 5 = 15$. Thus there is a 5-dimensional space of cubic relations satisfied by the (ω_i). In this space, a 4-dimensional subspace is accounted for by the product of Q with an arbitrary linear equation. Thus there must be, in addition, an irreducible cubic surface containing X. ∎

We will later see that the converse also holds.

Dimension counts for linear systems. Here is another perspective on the preceding proof. In intersection of surfaces of degree d in $\mathbb{P}^3$ with X gives a birational map between projective spaces,

$$|dH| \to |dK|.$$

Now note that in general a linear map $\phi : A \to B$ between vector spaces gives a birational map

$$\Phi : \mathbb{P}A \dashrightarrow \mathbb{P}B$$

which is projection from $\mathbb{P}C$ where $C = \operatorname{Ker}\phi$. Then $\dim A - \dim B \le \dim C$ and thus

$$\dim \mathbb{P}C \ge \dim \mathbb{P}A - \dim \mathbb{P}B - 1.$$

In the case at hand $\mathbb{P}C$ corresponds to the linear system $S_d(X)$ of surfaces of degree d containing X. Thus we get:

$$\dim S_d(X) \ge \dim |dH| - \dim |dK| - 1.$$

For $d = 2$ this gives

$$\dim S_2(X) \ge 9 - (3g - 4) - 1 = 9 - 8 - 1 = 0$$

which shows there is a unique quadric Q containing X. For $d = 3$ we get

$$\dim S_3(X) \ge 19 - (5g - 6) - 1 = 19 - 14 - 1 = 4.$$

Within $S_3(X)$ we have $Q + |H|$ which is 3-dimensional, and thus there must be a cubic surface C not containing Q in $S_3(X)$ as well.

This cubic is *not* unique, since we can move it in concert with $Q + H$.

Special divisors. An effective divisor D is *special* if $h^0(K - D) > 0$, i.e. if there is a holomorphic 1-form $\omega \neq 0$ vanishing at D.

In terms of the canonical map $\phi : X \to \mathbb{P}^{g-1}$, a divisor D is special iff $\phi(D)$ lies in a hyperplane H (determined by ω). (Moreover, the index of speciality, $i(D) = h^0(K - D)$, is one more than the dimension of the space of hyperplanes passing through D.)

Special divisors of degree g. The case of divisors of degree g is particularly interesting. Geometrically we see there exist plenty of such divisors – note that $|H \cap \phi(X)| = 2g - 2$, so a given hyperplane determines many such divisors. On the other hand, g typical points on $\phi(X)$ do *not* span a hyperplane, so these divisors really are special.

Proposition 10.16 *If $g \ge 2$ then their exist special divisors of degree g.*

Theorem 10.17 *An effective divisor D of degree g is special if and only if there is a nonconstant meromorphic function on X with $(f) + D \ge 0$.*

Proof. By Riemann-Roch, we have $h^0(D) = i(D) + \deg D - g + 1 = i(D) + 1 > 1$ iff $i(D) > 0$ iff D is special. ∎

Corollary 10.18 *If $g \geq 2$ then X admits a meromorphic function of degree $\leq g$.*

Example: genus 3. A curve of genus 3 either admits a map to $\mathbb{P}^1$ of degree two, or it embeds as a curve $X \subset \mathbb{P}^2$ of degree 4. In the latter case, projection to $\mathbb{P}^1$ from any $P \in X$ gives a map of degree $g = 3$.

The Wronskian and Weierstrass points. Now we focus on divisors of the form $D = gP$. We say P is a *Weierstrass point* if gP is special, i.e. if there is a meromorphic function $f : X \to \mathbb{P}^1$ with a pole of order at least one and at most g at P, and otherwise holomorphic.

Example: there are no Weierstrass points on a Riemann surface of genus 1. The branch points of every hyperelliptic surface of genus $g \geq 2$ are Weierstrass points.

To have $H \cap \phi(X) = D = gP$, the hyperplane H should contain not just P but the appropriate set of tangent directions at P, namely

$$(\phi(P), \phi'(P), \phi''(P), \ldots, \phi^{(g-1)}(P)).$$

For these tangent $(g-1)$ tangent directions to span a $(g-2)$-dimensional plane H through $\phi(P)$, there must be a linear relation among them; that is, the *Wronskian determinant* $W(P)$ must vanish.

In terms of a basis for Ω and a local coordinate z at P, the Wronskian is given by

$$W(z) = \det\left(\frac{d^j \omega_i}{dz^j}\right),$$

where $j = 0, \ldots, g-1$ and $i = 1, \ldots, g$.

We can see directly the vanishing of the Wronskian is equivalent to gP being special.

Theorem 10.19 *The Wronskian vanishes at P iff there is a holomorphic 1-form $\omega \neq 0$ with a zero at P of order at least g.*

Proof. The determinant vanishes iff there is a linear combination of the basis elements ω_i whose derivatives through order $(g-1)$ vanish at P. ■

The quantity $W = W(z)\, dz^N$ turns out to be independent of the choice of coordinate, where $N = 1 + 2 + \ldots + g = g(g+1)/2$. This value of N arises because the jth derivative of a 1-form behaves like dz^{j+1}.

Thus $W(z)$ is a section of $\mathcal{O}_{NK}$, so its number of zeros is $\deg NK = N(2g-2) = (g-1)g(g+1)$. This shows:

Theorem 10.20 *Any Riemann surface of genus g has $(g-1)g(g+1)$ Weierstrass points, counted with multiplicity.*

Weierstrass points of a hyperelliptic curve. These correspond to the branch points of the hyperelliptic map $\pi : X \to \mathbb{P}^1$, since the projective normal curve $\mathbb{P}^1 \to \mathbb{P}^{g-1}$ has no flexes.

Weierstrass points in genus 3. The Weierstrass points on a smooth quartic correspond to flexes; there are $2 \cdot 3 \cdot 4 = 24$ of them in general. For example, on the Fermat curve $x^4 + y^4 = 1$, there are 12 flexes altogether, each of multiplicity 2. Of these, 8 lie in the affine plane, and arise when one coordinate vanishes and the other is a 4th root of unity.

At the flexes we have $h^0(3P) = 2$. How can one go from a flex P to a degree 3 branched covering $f : X \to \mathbb{P}^1$? We can try projection f_P from P, but in general the line L tangent to X at P will meet X in a fourth point Q. Thus f_P will have a double pole at P and a simple pole at Q.

Instead, we project from Q! Then the line L through P and Q has multiplicity 3 at P, giving a triple order pole there.

Flexes of plane curves. In general, if C is defined by $F(X, Y, Z) = 0$, then the flexes of C are the locus where both F and the Hessian H of F vanish. For the Fermat curve, we have $F(X, Y, Z) = X^4 + Y^4 + Z^4$ and $H = 1728(XYZ)^2$. On a smooth curve of degree d the number of flexes is $3d(d-2)$.

The dimension of moduli space $\mathcal{M}_g$: Riemann's count. What is the dimension of $\mathcal{M}_g$? We know the dimension is 0, 1 and 3 for genus $g = 0$, 1 and 2 (using 6 points on $\mathbb{P}^1$ for the last computation).

Here is Riemann's heuristic. Take a large degree $d \gg g$, and consider the bundle $\mathcal{F}_d \to \mathcal{M}_g$ whose fibers are meromorphic functions $f : X \to \mathbb{P}^1$ of degree d. Now for a fixed X, we can describe $f \in \mathcal{F}_d(X)$ by first giving its polar divisor $D \geq 0$; then f is a typical element of $H^0(X, \mathcal{O}_D)$. (The parameters determining f are its principal parts on D.) Altogether with find

$$\dim \mathcal{F}_d(X) = d + h^0(D) = 2d - g + 1.$$

On the other hand, f has b critical points, where

$$\chi(X) = 2 - 2g = 2d - b,$$

so $b = 2d + 2g - 2$. Assuming the critical values are distinct, they can be continuously deformed to determine new branched covers (X', f'). Thus the dimension of the total space is given by

$$b = \dim \mathcal{F}_d = \dim \mathcal{F}_d(X) + \dim \mathcal{M}_g = 2d + 2g - 2 = 2d - g + 1 + \dim \mathcal{M}_g,$$

and thus $\dim \mathcal{M}_g = 3g - 3$. This dimension is in fact correct.

On the other hand, the space of hyperelliptic Riemann surfaces clearly satisfies

$$\dim \mathcal{H}_g = 2g + 2 - \dim \operatorname{Aut} \mathbb{P}^1 = 2g - 1,$$

since such a surface is branched over $2g+2$ points. Thus for $g > 2$ a typical Riemann surface is not hyperelliptic.

Tangent space to $\mathcal{M}_g$. As an alternative to Riemann's count, we note that the tangent space to the deformations of X is $H^1(X, \Theta)$, where $\Theta \cong \Omega^*$ is the sheaf of holomorphic vector fields on X. By Serre duality and Riemann-Roch, we have

$$\dim H^1(X, \Theta) = \dim H^1(X, \mathcal{O}_{-K}) = h^0(2K) = 4g - 4 - g + 1 = 3g - 3.$$

Serre duality also shows $H^1(X, \Theta)$ is naturally dual to the space of holomorphic quadratic differentials $Q(X)$.

Plane curves again. The space of homogeneous polynomials on $\mathbb{C}^{n+1}$ of degree d has dimension $N = \binom{n+d}{n}$. Thus the space of plane curves of degree d, up to automorphisms of $\mathbb{P}^2$, has dimension

$$N_d = \binom{2+d}{d} - 9.$$

We find

$$N_d = \begin{cases} -3 = \dim \operatorname{Aut} \mathbb{P}^1 & \text{for } d = 2, \\ 1 = \dim \mathcal{M}_1 & \text{for } d = 3, \\ 6 = \dim \mathcal{M}_3 & \text{for } d = 4, \\ 12 < \dim \mathcal{M}_6 = 15 & \text{for } d = 5. \end{cases}$$

Thus most curves of genus 6 can be realized as a plane curve. In a sense made precise by the Theorem below, there is no way to simply parameterize the moduli space of curves of high genus:

Theorem 10.21 (Harris-Mumford) *For g sufficiently large, $\mathcal{M}_g$ is of general type.*

11 Line bundles

Let X be a *complex manifold.* A *line bundle* $\pi : L \to X$ is a 1-dimensional holomorphic vector bundle.

This mean there exists a collection of trivializations of L over charts U_i on X, say $L_i \cong U_i \times \mathbb{C}$. Viewing L in two different charts, we obtain clutching data $g_{ij} : U_i \cap U_j \to \mathbb{C}^*$ such that $(x, y_j) \in L_j$ is equivalent to $(x, y_i) \in L_i$ iff $g_{ij}(x) y_j = y_i$. This data satisfies the cocycle condition $g_{ij} g_{jk} = g_{ik}$.

In terms of charts, a *holomorphic section* $s : X \to L$ is encoded by holomorphic functions $s_i = y_i \circ s(x)$ on U_i, such that

$$s_i(x) = g_{ij}(x) s_j(x).$$

Examples: the trivial bundle $X \times \mathbb{C}$; the canonical bundle $\wedge^n \mathrm{T}^* X$. Here the transition functions are $g_{ij} = 1$ for the trivial bundle and $g_{ij} = \det D\phi_i \circ \phi_j^{-1}$ for the canonical bundle, with charts $\phi_i : U_i \to \mathbb{C}^n$.

In detail, on a Riemann surface X, with local coordinates $z_i : U_i \to \mathbb{C}$, a section of the canonical bundle is locally given by $\omega_i = s_i(z_i)\, dz_i$; it must satisfy $s_i\, dz_i = s_j\, dz_j$, so $s_i = (dz_i / dz_j) s_j$.

Tensor powers. From L we can form the line bundle $L^* = L^{-1}$, and more generally L^d, with transition functions g_{ij}^d.

A line bundle is *trivial* if it admits a nowhere-vanishing holomorphic section (which then provides an isomorphism between L and $X \times \mathbb{C}$). Such a section exists iff there are $s_i \in \mathcal{O}^*(U_i)$ such that $s_i / s_j = g_{ij}$, i.e. iff g_{ij} is a coboundary.

Thus line bundles up to isomorphism over X are classified by the cohomology group $H^1(X, \mathcal{O}^*)$.

Sections and divisors. Now consider a divisor D on a Riemann surface X. Then we can locally find functions $s_i \in \mathcal{M}(U_i)$ with $(s_i) = D$. From this data we construction a line bundle L_D with transition functions $g_{ij} = s_i / s_j$. These transitions functions are chosen so that s_i is automatically a *meromorphic section* of L_D; indeed, a holomorphic section if D is effective.

Theorem 11.1 *The sheaf of holomorphic sections $\mathcal{L}$ of $L = L_D$ is isomorphic to $\mathcal{O}_D$.*

Proof. Choose a meromorphic section $s : X \to L$ with $(s) = D$. (On a compact Riemann surface, s is well-defined up to a constant multiple.) Then a local section $t : U \to L$ is holomorphic if and only if the meromorphic function $f = t/s$ satisfies

$$(t) = (fs) = (f) + D \geq 0$$

on U, which is exactly the condition that $f \in \mathcal{O}_D$. Thus the map $f \mapsto fs$ gives an isomorphism between $\mathcal{O}_D(U)$ and $\mathcal{L}(U)$. ∎

Line bundles on Riemann surfaces. Conversely, it can be shown that every line bundle L on a Riemann surface admits a non-constant meromorphic section, and hence $L = L_D$ for some D. More precisely, if $\mathcal{L}$ is the sheaf of sections of L one can show (see e.g. Forster, Ch. 29):

Theorem 11.2 *The group $H^1(X, \mathcal{L})$ is finite-dimensional.*

Corollary 11.3 *Given any $P \in X$, there exists a meromorphic section $s : X \to L$ with a pole of degree $\leq 1 + h^1(\mathcal{L})$ at P and otherwise holomorphic.*

Corollary 11.4 *Every line bundle has the form $L \cong L_D$ for some divisor D.*

From the point of view of sheaf theory, we have

$$0 \to \mathcal{O}^* \to \mathcal{M}^* \to \mathrm{Div} \to H^1(X, \mathcal{O}^*) \to H^1(X, \mathcal{M}^*) \to 0,$$

and since every line bundle is represented by a divisor, we find:

Corollary 11.5 *The group $H^1(X, \mathcal{M}^*) = 0$.*

Divisors and line bundles in higher dimensions. On complex manifolds of higher dimension, we can similarly construct line bundles from divisors. First, a divisor is simply an element of $H^0(\mathcal{M}^*/\mathcal{O}^*)$; this means it is locally a formal sum of analytic hypersurfaces, $D = \sum (f_i)$. Then the associated transition functions are $g_{ij} = f_i/f_j$ as before, and we find:

Theorem 11.6 *Any divisor D on a complex manifold X determines a line bundle $L_D \to X$ and a meromorphic section $s : X \to L$ with $(s) = D$.*

Failure of every line bundle to admit a nonzero section. However in general not every line bundle arises in this way. For example, there exist complex 2-tori $M = \mathbb{C}^2/\Lambda$ with no divisors but with plenty of line bundles (coming from characters $\chi : \pi_1(M) \to S^1$).

Degree. The degree of a line bundle, $\deg(L)$, is the degree of the divisor of any meromorphic section.

In terms of cohomology, the degree is associated to the exponential sequence: we have

$$\ldots H^1(X, \mathbb{Z}) \to H^1(X, \mathcal{O}) \to H^1(X, \mathcal{O}^*) \to H^2(X, \mathbb{Z}) \cong \mathbb{Z}.$$

This allows one to define the degree or *first Chern class*, $c_1(L) \in H^2(X,\mathbb{Z})$, for a line bundle on any complex manifold.

Projective space. For projective space, we have $H^1(\mathbb{P}^n,\mathbb{C}) = 0$ and so $H^1(X,\mathcal{O}) = 0$. It follows that line bundles on projective space are *classified by their degree*: we have

$$0 \to H^1(\mathbb{P}^n,\mathcal{O}^*) \to H^2(\mathbb{P}^n,\mathbb{Z}) \cong \mathbb{Z}.$$

We let $\mathcal{O}(d)$ denote the (sheaf of sections) of the unique line bundle of degree d. It has the property that for any meromorphic section $s \in H^0(\mathbb{P}^n,\mathcal{O}(d))$, the divisor $D = (s)$ represents $d[H] \in H^2(\mathbb{P}^n,\mathbb{Z})$, where $H \cong \mathbb{P}^{n-1}$ is a hyperplane.

Example: the tautological bundle. Let $\mathbb{P}^n$ be the projective space of $\mathbb{C}^{n+1}$ with coordinates $Z = (Z_0,\ldots,Z_n)$. The *tautological bundle* $\tau \to \mathbb{P}^n$ has, as its fiber over $p = [Z]$, the line $\tau_p = \mathbb{C}\cdot Z \subset \mathbb{C}^{n+1}$. Its total space is $\mathbb{C}^{n+1}$ with the origin blown up.

To describe τ in terms of transition functions, let $U_i = (Z_i \neq 0) \subset \mathbb{P}^n$. Then we can use the coordinate Z_i itself to trivialize $\tau|U_i$; in other words, we can map $\tau|U_i$ to $U_i \times \mathbb{C}$ the map

$$(Z_0,\ldots,Z_n) \mapsto ([Z_0 : \cdots : Z_n], Z_i).$$

(Here the origin must be blown up.) Then clearly the transition functions are given simply by $g_{ij} = Z_i/Z_j$, since they satisfy

$$Z_i = g_{ij} Z_j.$$

Theorem 11.7 *There is no nonzero holomorphic section of the tautological bundle.*

Proof. A section gives a map $s : \mathbb{P}^n \to \tau \to \mathbb{C}^{n+1}$ which would have to be constant because $\mathbb{P}^n$ is compact. But then the constant must be zero, since this is the only point in $\mathbb{C}^{n+1}$ that lies on every line through the origin. ■

As a typical *meromorphic* section, we can define $s(p)$ to be the intersection of τ_p with the hyperplane $Z_0 = 1$. In other words,

$$s([Z_0 : Z_1 : \cdots : Z_n]) = (1, Z_1/Z_0, \ldots, Z_n/Z_0).$$

Then $s_i = Z_i/Z_0$. Notice that this section is nowhere vanishing (since Z_i has no zero on U_i), but it has a pole along the divisor $H_0 = (Z_0)$.

Thus we have $\tau \cong \mathcal{O}(-1)$. Similarly, $\tau^* = \mathcal{O}(1)$.

Homogeneous polynomials. Note that the coordinates Z_i are sections of $\mathcal{O}(1)$. Indeed, any element in V^* naturally determines a function on the tautological bundle over $\mathbb{P}V$, linear on the fibers, and hence a section of the dual bundle. Similarly we have:

Theorem 11.8 *The space of global sections of $\mathcal{O}(d)$ over $\mathbb{P}^n$ can be naturally identified with the homogeneous polynomials of degree d on $\mathbb{C}^{n+1}$.*

Corollary 11.9 *The hypersurfaces of degree d in projective space are exactly the zeros of holomorphic sections of $\mathcal{O}(d)$.*

The canonical bundle. To compute the canonical bundle of project space, we use the coordinates $z_i = Z_i/Z_0$, $i = 1, \ldots, n$ to define a nonzero canonical form

$$\omega = dz_1 \cdots dz_n$$

on U_0. To examine this form in U_1, we use the coordinates $w_1 = Z_0/Z_1$, $w_i = Z_i/Z_1$, $i > 1$; then $z_1 = 1/w_1$ and $z_i = w_i/w_1$, $i > 1$, so we have

$$\omega = d(1/w_1)\, d(w_2/w_1) \cdots d(w_n/w_1) = -(dw_1 \cdots dw_n)/w_1^{n+1}.$$

Thus $(\omega) = (-n-1)H_0$ and thus the canonical bundle satisfies $K \cong \mathcal{O}(-n-1)$ on $\mathbb{P}^n$.

The adjunction formula.

Theorem 11.10 *Let $X \subset Y$ be a smooth hypersurface inside a complex manifold. Then the canonical bundles satisfy*

$$K_X \cong (K_Y \otimes L_X)|X.$$

Proof. We have an exact sequence of vector bundles on X:

$$0 \to \mathrm{T}X \to \mathrm{T}Y \to \mathrm{T}Y/\,\mathrm{T}X = NX \to 0,$$

where NX is the normal bundle. Now $(NX)^*$ is the sub-bundle of $\mathrm{T}^*Y|X$ spanned by 1-forms that annihilate TX. If X is defined in charts U_i by $f_i = 0$, then $g_{ij} = f_i/f_j$ defines L_X. On the other hand, df_i is a nonzero holomorphic section of $(NX)^*$. The 1-forms $df_i|X$, however, do not fit together on overlaps to form a global section of $(NX)^*$. Rather, on X we have $f_j = 0$ so

$$df_i = d(g_{ij} f_j) = g_{ij} df_j.$$

This shows (df_i) gives a global, nonzero section of $(NX)^* \otimes L_X$, and hence this bundle is trivial on X.

On the other hand $K_Y = K_X \otimes (NX)^*$, by taking duals and determinants. Thus $K_X = K_Y \otimes NX = K_Y \otimes L_X$. ∎

Smooth plane curves. Using the adjunction formula plus Riemann-Roch we can obtain some interesting properties of smooth plane curves $X \subset \mathbb{P}^2$ of degree d.

Theorem 11.11 *Let $f : X \to \mathbb{P}^n$ be a holomorphic embedding. Then f is given by a subspace of sections of the line bundle $L \to X$, where $L = f^*\mathcal{O}(1)$.*

Proof. The divisors of section of $\mathcal{O}(1)$ are hyperplanes. ∎

Theorem 11.12 *Every smooth plane curve of degree d has genus $g = (d-1)(d-2)/2$.*

Proof. We have $K_X \cong K_{\mathbb{P}^2} \otimes L_X = \mathcal{O}(d-3)$. Any curve Z of degree $d-3$ is the zero set of a section of $\mathcal{O}(d-3)$ and hence restricts to the zero set of a holomorphic 1-form on X. Thus we find $2g-2 = d(d-3)$. ∎

Corollary 11.13 *Every smooth quartic plane curve X is a canonical curve.*

Proof. We have $K_X \cong \mathcal{O}(d-3) = \mathcal{O}(1)$, which is the linear system that gives the original embedding of X into $\mathbb{P}^2$. ∎

Next note that the genus $g(X) = (d-1)(d-2)/2$ coincides with the dimension of the space of homogeneous polynomials on $\mathbb{C}^3$ of degree $d-3$. This shows:

Theorem 11.14 *Every effective canonical divisor on X has the form $K = X \cap Y$, where Y is a curve of degree $d-3$.*

Theorem 11.15 *Any $n+1$ distinct points in $\mathbb{P}^2$ impose independent condition on curves of degree n.*

Proof. Choose Y to be the union of n random lines through the first $k \le n$ points. Then Y is an example of a curve through the first k points that does not pass through the $k+1$st. This shows that adding the $k+1$st point imposes an additional condition on Y. ∎

Theorem 11.16 *A smooth curve X of degree $d > 1$ admits a nonconstant map to $\mathbb{P}^1$ of degree $d-1$, but none of degree $d-2$.*

Proof. For degree $d-1$, simply projection from a point on X. For the second assertion, suppose $f : X \to \mathbb{P}^1$ has degree $e \le d-2$. Let $E \subset X$ be a generic fiber of f. Then the $d-2$ points E impose independent conditions on the space of curves Y degree $d-3$. Consequently

$$h^0(K - E) = g - \deg E.$$

By Riemann-Roch we then have:

$$h^0(E) = 1 - g + \deg(E) - h^0(K - E) = 1,$$

so $|E|$ does not provide a map to $\mathbb{P}^1$. ∎

Hypersurfaces in products of projective spaces. Here are two further instances of the adjunction theorem.

Theorem 11.17 *Every smooth intersection X of a quadric Q and a cubic surface C is a canonical curve in $\mathbb{P}^3$.*

Proof. We have $K_Q \cong K_{\mathbb{P}^3} \otimes L_Q$ and thus

$$K_X \cong K_Q \otimes L_C \cong K_{\mathbb{P}^3} \otimes L_Q \otimes L_C \cong \mathcal{O}(-4 + 2 + 3) = \mathcal{O}(1).$$

∎

Theorem 11.18 *Every smooth (d, e) curve on $Q = \mathbb{P}^1 \times \mathbb{P}^1$ has genus $g = (d-1)(e-1)$.*

Proof. It is easy to see that $K_{X \times Y} = K_X \otimes K_Y$. Thus $K_Q = \mathcal{O}(-2, -2)$. Therefore $2g - 2 = C \cdot K_C$ and $K_C = \mathcal{O}(d-2, e-2)$, so $2g - 2 = d(e-2) + e(d-2)$, which implies the result. ∎

K3 surfaces. Manifolds with trivial canonical bundle are often interesting — in higher dimensions they are called *Calabi-Yau* manifolds.

For Riemann surfaces, K_X is trivial iff X is a complex torus. For 2-dimensional manifolds, complex tori also have trivial canonical, but they are not the only examples. Another class is provided by the *K3 surfaces*, which by definition are simply-connected complex surfaces with K_X trivial. Example:

Theorem 11.19 *Every surface of degree* $(2,2,2)$ *in* $X = \mathbb{P}^1 \times \mathbb{P}^1 \times \mathbb{P}^1$ *has trivial canonical bundle.*

Proof. Here $K_X = \mathcal{O}(-2,-2,-2)$. One can use the Lefschetz hyperplane theorem to show a hypersurface in a simply-connected complex 3-manifold is always itself simply-connected. ∎

12 Curves and their Jacobians

We now turn to the important problem of classifying line bundles on a Riemann surface X; equivalently, of classifying divisors modulo linear equivalence.

The Jacobian. Recall that a holomorphic 1-form on X is the same thing as a holomorphic map $f : X \to \mathbb{C}$ well-defined up to translation in $\mathbb{C}$. If the periods of f happen to generate a discrete subgroup Λ of $\mathbb{C}$, then we can regard f as a map to $\mathbb{C}/\Lambda$. However the periods are almost always indiscrete. Nevertheless, we can put all the 1-forms together and obtain a map to $\mathbb{C}^g/\Lambda$.

Theorem 12.1 *The natural map* $H_1(X,\mathbb{Z}) \to \Omega(X)^*$ *has as its image a lattice* $\Lambda \cong \mathbb{Z}^{2g}$.

Proof. If not, the image lies in a real hyperplane defined, for some nonzero $\omega \in \Omega(X)$, by the equation $\operatorname{Re}\alpha(\omega) = 0$. But then all the periods of $\operatorname{Re}\omega$ vanish, which implies the harmonic form $\operatorname{Re}\omega = 0$. ∎

The *Jacobian variety* is the quotient space $\operatorname{Jac}(X) = \Omega(X)^*/H_1(X,\mathbb{Z})$, the cycles embedded via periods.

Theorem 12.2 *Given any basepoint* $P \in X$*, there is a natural map* $\phi_P : X \to \operatorname{Jac}(X)$ *given by* $\phi_P(Q) = \int_P^Q \omega$.

We will later show this map is an embedding, and thus $\operatorname{Jac}(X)$, roughly speaking, makes X into a group.

Example: the pentagon. Let X be the hyperelliptic curve defined by $y^2 = x^5 - 1$. Geometrically, X can be obtained by gluing together two regular pentagons. Cleary X admits an automorphism $T : X \to X$ of order 5. Using the pentagon picture, one can easily show there is a cycle $C \in H_1(X,\mathbb{Z})$ such that its five images $T^i(C)$ span $H_1(X,\mathbb{Z}) \cong \mathbb{Z}^4$. This means $H_1(X,\mathbb{Z}) = A{\cdot}C$ is a free, rank one A-module, where $A = \mathbb{Z}[T]/(1+T+T^2+T^3+T^4)$.

Now T^* acts on $\Omega(X)$. Since $X/\langle T\rangle$ has genus zero, T has no invariant forms. Thus we can choose a basis (ω_1, ω_2) for $\Omega(X)$ such that

$$T^*\omega_i = \zeta_i \omega_i$$

where ζ_i is a primitive 5th root of unity.

Let us scale these ω_i so $\int_C \omega_i = 1$. Let

$$\pi : H_1(X,\mathbb{Z}) \to \Lambda \subset \mathbb{C}^2$$

be the period map, defined by

$$\pi(B) = \left(\int_B \omega_1, \int_B \omega_2 \right).$$

Since

$$\int_{T(B)} \omega = \int_B T^*\omega,$$

we have

$$\pi(TB) = \begin{pmatrix} \zeta_1 & 0 \\ 0 & \zeta_2 \end{pmatrix} \pi(B).$$

Since $H_1(X,\mathbb{Z}) = \mathbb{Z}[T] \cdot C$, we find that $\Lambda \subset \mathbb{C}^2$ is simply the image of the ring $\mathbb{Z}[T]$ under the ring homomorphism that sends T to (ζ_1, ζ_2).

Since Λ is a lattice, we cannot have $\zeta_2 = \overline{\zeta_1} = \zeta_1^4$, nor can we have $\zeta_2 = \zeta_1$. Thus $\zeta_2 = \zeta_1^2$ or ζ_1^3. In the latter case we can interchange the eigenforms to obtain the former case. This finally shows:

Theorem 12.3 *The Jacobian of the hyperelliptic curve $y^2 = x^5 - 1$ is isomorphic to $\mathbb{C}^2/\Lambda$, where Λ is the ring $\mathbb{Z}[T]/(1+T+T^2+T^3+T^4)$ embedded in $\mathbb{C}^2$ by*

$$T \mapsto (\zeta, \zeta^2),$$

and ζ is a primitive 5th root of unity.

This is an example of a Jacobian variety with *complex multiplication.*

The Picard group. Let $\mathrm{Pic}(X)$ denote the group of all line bundles on X. Since every line bundle admits a meromorphic section, there is a natural isomorphic between $\mathrm{Pic}(X)$ and $\mathrm{Div}(X)/(\mathcal{M}^*(X))$, where $\mathcal{M}^*(X)$ is the group of nonzero meromorphic functions, mapping to principal divisors in

$\mathrm{Pic}(X)$. Under this isomorphism, the degree and a divisor and of a line bundle agree.

The *Abel-Jacobi map* $\phi : \mathrm{Div}_0(X) \to \mathrm{Jac}(X)$ is defined by

$$\phi\left(\sum Q_i - P_i\right)(\omega) = \sum \int_{P_i}^{Q_i} \omega.$$

Because of the choice of path from P_i to Q_i, the resulting linear functional is well-defined only modulo cycles on X.

For example: given any basepoint $P \in X$, we obtain a natural map $f : X \to \mathrm{Jac}(X)$ by $f(Q) = \phi(Q - P) = \int_P^Q \omega$.

One of the most basic results regarding the Jacobian is:

Theorem 12.4 *The map ϕ establishes an isomorphism between* $\mathrm{Jac}(X)$ *and* $\mathrm{Pic}_0(X) = \mathrm{Div}_0(X)/$*(principal divisors).*

The proof has two parts: *Abel's theorem*, which asserts that ϕ is injective, and the *Jacobi inversion theorem*, which asserts that ϕ is surjective.

Abel's theorem. We turn to the proof of Abel's theorem, which states that a divisor D of degree zero is principal iff $\phi(D) = 0$. In more concrete terms this means:

Theorem 12.5 (Abel's theorem) *A divisor D is principal iff $D = \sum Q_i - P_i$ and*

$$\sum \int_{P_i}^{Q_i} \omega = 0$$

for all $\omega \in \Omega(X)$, for some choice of paths γ_i joining P_i to Q_i.

Proof in one direction. Suppose $D = (f)$. We can assume after multiplying f by a scalar, that none of its critical values are real. Let $\gamma = f^{-1}([0, \infty])$. Then we have

$$\sum \int_{P_i}^{Q_i} \omega = \int_\gamma \omega = \int_0^\infty f_*(\omega) = 0$$

since $f_*(\omega) = 0$, being a holomorphic 1-form on $\widehat{\mathbb{C}}$. (To see this, suppose locally $f(z) = z^d$. Then $f_*(z^i\,dz) = 0$ unless $z^i\,dz$ is invariant under rotation by the dth roots of unity. This first happens when $i = d - 1$, in which case $z^{d-1}\,dz = (1/d)d(z^d)$, so the pushforward is proportional to dz.) ∎

The curve in its Jacobian. Before proceeding to the proof of Abel's theorem, we derive some consequences.

Given $P \in X$, define

$$\phi_P : X \to \mathrm{Jac}(X)$$

by $\phi_P(Q) = (Q - P)$. Note that with respect to a basis ω_i for $\Omega(X)$, the derivative of $\phi_P(Q)$ in local coordinates is given by:

$$D\phi_P(Q) = (\omega_1(Q), \ldots, \omega_g(Q)).$$

This shows:

Theorem 12.6 *The canonical map $X \to \mathbb{P}\Omega(X)^*$ is the Gauss map of ϕ_P.*

Theorem 12.7 *For genus $g \geq 1$, the map $\phi_P : X \to \mathrm{Jac}(X)$ is a smooth embedding.*

Proof. If $Q - P = (f)$, then $f : X \to \mathbb{P}^1$ has degree 1 so $g = 0$. Since $|K|$ is basepoint-free, there is a nonzero-holomorphic 1-form at every point, and hence $D\phi_P \neq 0$. ∎

Theorem 12.8 (Jacobi) *The map $\mathrm{Div}_0(X) \to \mathrm{Jac}(X)$ is surjective. In fact, given $(P_1, \ldots, P_g) \in X^g$, the map*

$$\phi : X^g \to \mathrm{Jac}(X)$$

given by

$$\phi(Q_1, \ldots, Q_g) = \phi\left(\sum Q_i - P_i\right) = \left(\int_{P_i}^{Q_i} \omega_j\right)$$

is surjective.

Proof. It suffices to show that $\det D\phi \neq 0$ at some point, so the image is open. To this end, just note that $d\phi/dQ_i = (\omega_j(Q_i))$, and thus $\det D\phi(Q_1, \ldots, Q_g) = 0$ if and only if there is an ω vanishing simultaneously at all the Q_i, i.e. iff (Q_i) lies on a hyperplane under the canonical embedding. For generic Q_i's this will not be the case, and hence $\det D\phi \neq 0$ almost everywhere on X^g. ∎

Corollary 12.9 *We have a natural isomorphism:*

$$\mathrm{Pic}_0(X) = \mathrm{Div}_0(X)/(\mathcal{M}^*(X)) \cong \mathrm{Jac}(X).$$

Bergman metric. We remark that the space $\Omega(X)$, and hence its dual, carries a natural norm given by:

$$\|\omega\|_2^2 = \int_X |\omega(z)|^2\,|dz|^2.$$

This induces a canonical metric on $\mathrm{Jac}(X)$, and hence on X itself.

This metric ulimately comes from the intersection pairing or *symplectic form* on $H_1(X,\mathbb{Z})$, satisfying $a_i \cdot b_j = \delta_{ij}$.

Abel's theorem, proof I: The $\overline{\partial}$-equation. (Cf. Forster.) For the converse, we proceed in two steps. First we will construct a smooth solution to $(f) = D$; then we will correct it to become holomorphic.

Smooth solutions. Let us say a *smooth* map $f : X \to \widehat{\mathbb{C}}$ satisfies $(f) = D$ if near P_i (resp. Q_i) we have $f(z) = zh(z)$ (resp. $z^{-1}h(z)$) where h is a smooth function with values in $\mathbb{C}^*$, and if f has no other zeros or poles.

Note that for such an f, the *distributional* logarithmic derivative satisfies

$$\overline{\partial} \log f = \frac{\overline{\partial} f}{f} + \sum (p_i - q_i),$$

where q_i and p_i are δ functions (in fact currents), locally given by $\overline{\partial} \log z$, and $\overline{\partial} f/f$ is smooth.

Inspired by the proof in one direction already given, we first construct a smooth solution of $(f) = D$ which maps a disk neighborhood U_i of γ_i diffeomorphically to a neighborhood V of the interval $[0,\infty]$.

Lemma 12.10 *Given any arc γ joining P to Q on X, there exists a smooth solution to $(f) = Q - P$ satisfying*

$$\frac{1}{2\pi i} \int_X \omega \wedge \frac{\overline{\partial} f}{f} = \int_P^Q \omega$$

for all $\omega \in \Omega(X)$, where the integral is taken along γ.

Proof. First suppose P and Q are close enough that they belong to a single chart U, and γ is almost a straight line. Then we can choose the isomorphism $f : U \to V \subset \widehat{\mathbb{C}}$ so that $f(P) = 0$, $f(Q) = \infty$ and $f(\gamma) = [0,\infty]$ (altering γ

by a small homotopy rel endpoints). This f is already holomorphic on U, and it sends ∂U to a contractible loop in $\mathbb{C}^*$. Thus we can extend f to a smooth function sending $X - U$ into $\mathbb{C}^*$, which then satisfies $(f) = D$.

Now note that z admits a single-valued logarithm on the region $\widehat{\mathbb{C}} - [0, \infty]$, and thus $\log f(z)$ has a single-valued branch on $Y = X - \gamma$. Thus $\overline{\partial} f/f = \overline{\partial} \log f$ is an exact form on Y. However as one approaches γ from different sides, the two branches of $\log f$ differ by $2\pi i$. Applying Stokes' theorem, we find:

$$\int_X \omega \wedge \overline{\partial} \log f = 2\pi i \int_\gamma \omega.$$

To hand the case of well-separated P and Q, simply break γ up into many small segments and take the product of the resulting f's. ■

Taking the product of the solutions for several pairs of points, and using additivity of the logarithmic derivative, we obtain:

Corollary 12.11 *Given arc γ_i joining P_i to Q_i on X, there exists a smooth solution to $(f) = \sum Q_i - P_i$ satisfying*

$$\frac{1}{2\pi i} \int_X \omega \wedge \frac{\overline{\partial} f}{f} = \sum \int_{\gamma_i} \omega$$

for all $\omega \in \Omega(X)$.

From smooth to holomorphic. To complete the solution, it suffices to find a smooth function g such that fe^g is meromorphic. Equivalently, it suffices to solve the equation $\overline{\partial} g = -\overline{\partial} f/f$. (Note that $\overline{\partial} f/f$ is smooth even at the zeros and poles of f, since near there $f = z^n h$ where $h \neq 0$.) Since $H^{0,1}(X) \cong \Omega(X)^*$, such a g exists iff

$$\frac{1}{2\pi i} \int \omega \wedge \frac{\overline{\partial} f}{f} = \sum \int_{P_i}^{Q_i} \omega = 0$$

for all $\omega \in \Omega(X)$. This is exactly the hypothesis of Abel's theorem. ■

Abel's theorem, proof II: Symplectic forms. Our second proof makes the connection with the Jacobian and the symplectic form on $H^1(X, \mathbb{C})$ more transparent. (Cf. Lang, *Algebraic Functions*).

To try to construct f with $(f) = D$, we first construct a candidate for $\lambda = df/f$.

Theorem 12.12 *For any divisor with* $\deg D = 0$*, there exists a meromorphic differential* λ *with only simple poles such that* $\sum \mathrm{Res}_P(\lambda) \cdot P = D$*.*

Proof. By Riemann-Roch, for any $P, Q \in X$ we have $\dim H^0(K+P+Q) > \dim H^0(K)$. Thus there exists a meromorphic 1-form λ with a simple pole at one of P or Q. By the residue theorem, λ has poles at both points with opposite residues. Scaling λ proves the Theorem for $Q - P$, and a general divisor of degree zero is a sum of divisors of this form. ■

Alternate proof. By Mittag-Leffler for 1-forms, λ exists because the sum of its residues is zero. ■

Now if we can arrange that the periods of λ are all in the group $2\pi i\mathbb{Z}$, then $f(z) = \exp \int \lambda$ is a meromorphic function with $(f) = D$.

(Compare Mittag-Leffler's proof of Weierstrass's theorem on functions with prescribed zeros.)

The case of a torus. We now study the case where $X = \mathbb{C}/\mathbb{Z} \oplus \tau\mathbb{Z}$, $\tau \in \mathbb{H}$. Then $\Omega(X) = \mathbb{C} \cdot \omega$ where $\omega = dz$.

Let $F \subset \mathbb{C}$ be a fundamental polygon with sides a, b joining 0 to 1 and to τ, and with translates a', b' forming the rest of the boundary, so $\partial F = a + b' - a' - b$. The cycles a and b generate $H_1(X, \mathbb{Z})$, and the periods of ω are exactly

$$(\omega(a), \omega(b)) = (1, \tau).$$

Thus $\mathrm{Jac}(X)$ is isomorphic to X.

Reciprocity on a torus. Assume λ has poles only in the interior of F, and let $\lambda(a)$, $\lambda(b)$ denote the periods of λ around the loops a and b on X.

Let $\int \omega = z$ be a primitive for ω on F. Then by the residue theorem we have

$$\frac{1}{2\pi i} \int_{\partial F} z\lambda = \sum \mathrm{Res}_P(\lambda) z(P) = \sum \int_{P_i}^{Q_i} dz = \phi(D).$$

On the other hand, we find

$$\int_{a'-a} z\lambda = \lambda(a)\omega(b),$$

since z differs by $\omega(b)$ along corresponding points of a and a'. Similarly we have $\int_{b'-b} z\lambda = \lambda(b)\omega(a)$, and thus

$$\begin{aligned} 2\pi i\phi(D) &= \int_F z\lambda = \int_{a+b'-a'-b} z\lambda = \lambda(b)\omega(a) - \lambda(a)\omega(b) \\ &= \det \begin{pmatrix} \omega(a) & \omega(b) \\ \lambda(a) & \lambda(b) \end{pmatrix}. \end{aligned}$$

Conclusion of the proof on a torus. We now finish the proof with some linear algebra. First, suppose in the above formula the determinant vanishes. Then the periods of λ are proportional to those of ω. Thus for some $t \in \mathbb{C}$, the form $\lambda + t\omega$ has vanishing periods along a and b; replacing λ with this form, we obtain a meromorphic function with $(f) = D$.

Second, suppose we only know $\phi(D) = 0$ in $\mathrm{Jac}(X)$. That is, $\phi(D) = n_1\omega(a) + n_2\omega(b)$ for some integers (n_1, n_2). One way this can happen is if $(\lambda(a), \lambda(b)) = 2\pi i(-n_2, n_1)$. But after correcting λ by a multiple $t\omega$ as above, we can actually assume $(\lambda(a), \lambda(b)) = 2\pi i(-n_2, n_1)$. Then the periods of λ lie in $2\pi i\mathbb{Z}$, and once again $f = \exp \int \lambda$ satisfies $(f) = D$.

General surfaces: the symplectic form on $H^1(X, \mathbb{C})$. To treat the general case, we recall that a surface of genus g admits a basis for $H_1(X, \mathbb{Z})$ of the form (a_i, b_i), $i = 1, \ldots, g$, such that $a_i \cdot b_i = 1$ and all other products vanish.

This symplectic form on $H_1(X, \mathbb{Z})$ gives rise to one on the space of periods, $H^1(X, \mathbb{C})$, defined by:

$$[\alpha, \beta] = \sum \alpha(a_i)\beta(b_i) - \alpha(b_i)\beta(a_i).$$

Theorem 12.13 *Under the period isomorphism between $H^1_{DR}(X)$ and $H^1(X, \mathbb{C})$, we have*

$$\int_X \alpha \wedge \beta = [\alpha, \beta].$$

Proof. As before we cut X along the (a_i, b_i) curves to obtain a surface F with boundary, on which we can write $\alpha = dA$. Then we have $\int_F (dA) \wedge \beta = \int_{\partial F} A\,\beta$. As before, $\partial F = \sum a_i + b_i' - a_i' - b_i$, and thus

$$\int \alpha \wedge \beta = \sum \alpha(a_i)\beta(b_i) - \alpha(b_i)\beta(a_i).$$

General surfaces: Abel's theorem. With respect to this symplectic form, $\Omega(X) \subset H^1(X, \mathbb{C})$ is a Lagrangian subspace, and the bracket gives an isomorphism

$$\Omega(X)^* \cong H^1(X, \mathbb{C})/\Omega(X).$$

(To see that we have an isomorphism, just note that $\Omega(X)^\perp \supset \Omega(X)$ and $\dim \Omega(X)^\perp = 2g - \dim \Omega(X) = \dim \Omega(X)$.)

This isomorphism sends $H^1(X, \mathbb{Z})$ to the image of $H_1(X, \mathbb{Z})$ under the period mapping.

Then, in brief, we have

$$\langle 2\pi i\phi(D), \omega\rangle = 2\pi i \sum \int_{P_i}^{Q_i} \omega = [\lambda, \omega]$$

for every $\omega \in \Omega(X)$. Here the integral on the left is defined using paths between P_i and Q_i that lie in F, and the periods of λ on the right are along the chosen curves a_i, b_i forming the boundary of F.

The hypothesis of Abel's theorem is that $\phi(D) = 0$ in $\Omega(X)^*/H_1(X, \mathbb{Z})$. This mean there is a cycle $C = \sum n_i a_i + m_i b_i$ such that

$$\int_C \omega = [(2\pi i)^{-1}\lambda, \omega]$$

for all $\omega \in \Omega(C)$. In other words, if we let N be the integral vector $(-m_i, n_i)$ in $H^1(X, \mathbb{C})$, we have

$$[N, \omega] = [(2\pi i)^{-1}\lambda, \omega]$$

for all ω.

But this equality of brackets says exactly that $2\pi i N - \lambda$ is in $\Omega(X)^{\perp}$ with respect to the intersection form. Since $\Omega(X)$ is Lagrangian, we find there is an $\omega_0 \in \Omega(X)$ such that $2\pi i N$ represents the periods of $\lambda + \omega_0$.

In other words, after modifying λ by ω_0, the function $f = \exp \int \lambda$ becomes well-defined, and therefore we have solved for a meromorphic function such that $(f) = D$.

Letting $X^{(k)} = X^k/S_k \subset \mathrm{Div}_k(X)$, we have:

Theorem 12.14 *The fibers of the natural map $X^{(k)} \to \mathrm{Pic}_k(X)$ are projective spaces.*

Mordell's Conjecture.

Theorem 12.15 *Suppose X has genus $g \geq 2$. Then, given any finitely generated subgroup $A \subset \mathrm{Jac}(X)$, the set $X \cap A$ is finite.*

This theorem is in fact equivalent to Mordell's conjecture (Falting's theorem), which states that $X(K)$ is finite for any number field K.

The exponential sequence and the Jacobian. An alternative description of the Jacobian is via the exponential sequence which leads to the exact sequence

$$H^1(X, \mathbb{Z}) \to H^1(X, \mathcal{O}) \to H^1(X, \mathcal{O}^*) \to H^2(X, \mathbb{Z}) \to 0.$$

Under the isomorphisms $H^1(X,\mathbb{Z}) \cong H_1(X,\mathbb{Z})$ by cup product, $H^1(X,\mathcal{O}) \cong \Omega(X)^*$ by Serre duality, and $H^2(X,\mathbb{Z}) \cong \mathbb{Z}$ by degree, we obtain the isomorphism

$$\mathrm{Pic}_0(X) = \mathrm{Ker}(H^1(X,\mathcal{O}) \to H^2(X,\mathbb{Z})) \cong \Omega(X)^*/H_1(X,\mathbb{Z}) = \mathrm{Jac}(X).$$

The Siegel upper half-space. The Siegel upper half-space is the space $\mathcal{H}_g$ of symmetric complex $g \times g$ matrices P such that $\mathrm{Im}\, P$ is positive-definite.

The space $\mathcal{H}_g$ is the natural space for describing the g-dimensional complex torus $\mathrm{Jac}(X)$, just as $\mathbb{H}$ is the natural space for describing the 1-dimensional torus $\mathbb{C}/\Lambda$.

To describe the Jacobian via $\mathcal{H}_g$, we need to choose a symplectic basis (a_i, b_i) for $H_1(X,\mathbb{Z})$.

Theorem 12.16 *There exists a unique basis ω_i of $\Omega(X)$ such that $\omega_i(a_j) = \delta_{ij}$.*

Proof. To see this we just need to show the map from $\Omega(X)$ into the space of a-periods is injective (since both have dimension g. But suppose the a-periods of ω vanish. Then the same is true for the $(0,1)$-form $\overline{\omega}$, which implies

$$\int |\omega(z)|^2\, dz\, d\overline{z} = \int \omega \wedge \overline{\omega} = \sum \omega(a_i)\overline{\omega}(b_i) - \omega(b_i)\overline{\omega}(a_i) = 0,$$

and thus $\omega = 0$. ∎

Definition. The *period matrix* of X with respect to the symplectic basis (a_i, b_i) is given by

$$P_{ij} = \int_{b_j} \omega_i = \omega_i(b_j).$$

Theorem 12.17 *P is symmetric and $\mathrm{Im}\, P$ is positive-definite.*

Proof. To see symmetry we use the fact that for any i, j:

$$0 = \int \omega_i \wedge \omega_j = \sum \omega_i(a_k)\omega_j(b_k) - \omega_i(b_k)\omega_j(a_k) = \delta_{ij}P_{jk} - P_{ik}\delta_{jk} = P_{ji} - P_{ij}.$$

Similarly, we have

$$-\frac{i}{2}\int \omega \wedge \overline{\omega} = 2\int |\omega(z)|^2 |dz|^2 \geq 0,$$

since $dz\,d\overline{z} = 2i\,dx\,dy = 2i\,|dz|^2$. Thus the symmetric matrix $(i/2)Q$ is positive-definite, where

$$\begin{aligned} Q_{ij} &= \int \omega_i \wedge \overline{\omega}_j = \sum \omega_i(a_k)\overline{\omega}_j(b_k) - \omega_i(b_k)\overline{\omega}_j(a_k) = \delta_{ij}\overline{P}_{jk} - P_{ik}\delta_{jk} \\ &= \overline{P}_{ji} - P_{ij} = -2i\,\mathrm{Im}\,P_{ij}. \end{aligned}$$

Since $(i/2)Q = \mathrm{Im}\,P$, we find $\mathrm{Im}\,P$ is positive-definite. ∎

The *symplectic group* $\mathrm{Sp}_{2g}(\mathbb{Z})$ now takes over for $\mathrm{SL}_2(\mathbb{Z})$, and two complex tori are isomorphic as *principally polarized* abelian varieties if and only if there are in the same orbit under $\mathrm{Sp}_{2g}(\mathbb{Z})$.

With respect to the basis $(b_1, \ldots, b_g, a_1, \ldots, a_g)$ the symplectic group acts on the full matrix of a and b periods by

$$\begin{pmatrix} A & B \\ C & D \end{pmatrix} \begin{pmatrix} P \\ I \end{pmatrix} = \begin{pmatrix} AP + B \\ CP + D \end{pmatrix}.$$

We must then change the choice of basis for $\Omega(X)$ to make the a-periods, $CP + D$ into the identity matrix; and we find

$$g(P) = (AP + B)(CP + D)^{-1},$$

which shows $\mathrm{Sp}_{2g}(\mathbb{Z})$ acts on $\mathcal{H}_g$ by non-commutative fractional linear transformations.

Theta functions. The theory of θ-functions allows one to canonically attach a divisor

$$\Theta \subset A = \mathbb{C}^g/(\mathbb{Z}^g \oplus P(\mathbb{Z}^g))$$

to the principally polarized Abelian variety A determined by $P \in \mathcal{H}_g$. Using the fact that $\mathrm{Im}\,P \gg 0$, we define the entire θ-function $\theta : \mathbb{C}^g \to \mathbb{C}$ by

$$\theta(z) = \sum_{n \in \mathbb{Z}^g} \exp(2\pi i \langle n, z\rangle) \exp(\pi i \langle n, Pn\rangle).$$

The zero-set of θ is Λ-invariant, and descends to a divisor Θ on A.

For example, when $g = 1$ and $P = [\tau]$, $\tau \in \mathbb{H}$, we obtain:

$$\theta(z) = \sum_{\mathbb{Z}} \exp(2\pi i n z) \exp(\pi i n^2 \tau).$$

Clearly $\theta(z + 1) = \theta(z)$, and we have

$$\theta(z + \tau) = \theta(z) \exp(-2\pi i z - \pi i \tau).$$

Theorem 12.18 (Riemann) *We have* $\Theta = W_{g-1} + \kappa$ *for some* $\kappa \in \mathrm{Jac}(X)$, *where* $W_{g-1} \subset \mathrm{Jac}(X)$ *is the image of* X^{g-1} *under the Abel-Jacobi map.*

The Torelli theorem.

Theorem 12.19 *Suppose* $(\mathrm{Jac}(X), \omega)$ *and* $(\mathrm{Jac}(Y), \omega')$ *are isomorphic as (principally) polarized complex tori (Abelian varieties). Then* X *is isomorphic to* Y.

Sketch of the proof. Using the polarization, we can reconstruct the divisor $W_{g-1} \subset \mathrm{Jac}(X)$ up to translation, using θ-functions.

Now the tangent space at any point to $\mathrm{Jac}(X)$ is canonically identified with $\Omega(X)^*$, and hence tangent hyperplanes in $\mathrm{Jac}(X)$ give hyperplanes in $\mathbb{P}\Omega(X)$, the ambient space for the canonical curve. (For convenience we will assume X is not hyperelliptic.) In particular, the Gauss map on the smooth points of W_{g-1} defines a natural map

$$\gamma : W_{g-1} \to \mathbb{P}\Omega(X) \cong \mathbb{P}^{g-1}.$$

The composition of γ with the natural map $C^{g-1} \to W_{g-1}$ simply sends $g-1$ points P_i on the canonical curve to the hyperplane $H(P_i)$ they (generically) span.

Clearly this map is surjective. Since W_{g-1} and $\mathbb{P}^{g-1}$ have the same dimension, γ is a local diffeomorphism at most points. The *branch locus* corresponds to the hyperplanes that are *tangent* to the canonical curve $X \subset \mathbb{P}^{g-1}$. Thus from $(\mathrm{Jac}(X), \omega)$ we can recover the collection of hyperplanes X^* tangent to $X \subset \mathbb{P}^{g-1}$.

Finally one can show geometrically that $X^* = Y^*$ implies $X = Y$. The idea is that, to each point $P \in X$ we have a $(g-3)$ dimensional family of hyperplanes $H_P \subset X^*$ containing the tangent line $T_P(X)$. These hyperplanes must all be tangent to Y as well; but the only reasonable way this can happen is if $T_P(X) = T_Q(Y)$ for some Q on Y.

Now for genus $g > 3$ one can show a tangent line meets X in exactly one point; thus Q is unique and we can define an isomorphism $f : X \to Y$ by $f(P) = Q$. (For example, in genus 4 the canonical curve is a sextic in $\mathbb{P}^3$. If a tangent line $L \subset \mathbb{P}^3$ were to meet $2P$ and $2Q$, then the planes through L would give a complementary linear series of degree 2, so X would be hyperelliptic.) For $g = 3$ there are in general a finite number of 'bitangents' to the quartic curve $X \subset \mathbb{P}^2$, and away from these points we can define f, then extend. (For details see Griffiths and Harris). ∎

Question. What is the infinitesimal form of the Torelli theorem?

13 Hyperbolic geometry

(The B-side.)

Elements of hyperbolic geometry in the plane. The hyperbolic metric is given by $\rho = |dz|/y$ in $\mathbb{H}$ and $\rho = 2|dz|/(1-|z|^2)$ in Δ.

Thus $d(i, iy) = \log y$ in $\mathbb{H}$, and $d(0,x) = \log(x+1)/(x-1)$ in Δ. Note that $(x+1)/(x-1)$ maps $(-1,1)$ to $(0,\infty)$.

An important theorem for later use gives the hyperbolic distance from the origin to a hyperbolic geodesic γ which is an arc of a circle of radius r:

$$\sinh d(0,\gamma) = \frac{1}{r} \cdot$$

To see this, let x be the Euclidean distance from 0 to γ. Then we have, by algebra,

$$\sinh d(0,x) = \frac{2x}{1-x^2} \cdot \tag{13.1}$$

On the other hand, we have by a right triangle with sides 1, r and $x+r$. Thus $1+r^2 = (x+r)^2$ which implies $2x/(1-x^2) = 1/r$.

A more intrinsic statement of this theorem is that for any point $p \in \mathbb{H}$ and geodesic $\gamma \in \mathbb{H}$, we have

$$\sinh d(p,\gamma) = \cot(\theta/2),$$

where θ is the visual angle subtended by γ as seen from p.

Area of triangles and polygons. The area of an ideal triangle is π. The area of a triangle with interior angles $(0,0,\alpha)$ is $\pi-\alpha$. From these facts one can see the area of a general triangle is given by the angle defect:

$$T(\alpha,\beta,\gamma) = \pi - \alpha - \beta - \gamma.$$

To see this, one extends the edges of T to rays reaching the vertices of an ideal triangle I; then we have

$$T(\alpha,\beta,\gamma) = I - T(\pi-\alpha) - T(\pi-\beta) - T(\pi-\gamma)$$

which gives $\pi-\alpha-\beta-\gamma$ for the area.

Another formulation is that the area of a triangle is the sum of its exterior angles minus 2π. In this form the formula generalizes to polygons.

Right quadrilaterals with an ideal vertex.

Theorem 13.1 *Let Q be a quadrilateral with edges of lengths (a, b, ∞, ∞) and interior angles $(\pi/2, \pi/2, \pi/2, 0)$. Then we have*

$$\sinh(a)\sinh(b) = 1.$$

Proof. First we make a remark in Euclidean geometry: let Q' be an ideal hyperbolic quadrilateral, centered at the origin, with sides coming from circles of Euclidean radii (r, R, r, R). Then $rR = 1$.

Indeed, from this picture we can construct a right triangle with right-angle vertex 0, with hypotenuse of length $r + R$, and with altitude from the right-angle vertex of 1. By basic Euclidean geometry of similar triangles, we find $rR = 1^2 = 1$.

Now cut Q' into 4 triangles of the type Q. Then we have $\sinh(a) = 1/r$ and $\sinh(b) = 1/R$, by (13.1). Therefore $\sinh(a)\sinh(b) = 1$. ■

Right hexagons. For a right-angled hexagon H, the excess angle is $6(\pi/2) - 2\pi = \pi$, and thus $\mathrm{area}(H) = \pi$.

Theorem 13.2 *For any $a, b, c > 0$ there exists a unique right hexagon with alternating sides of lengths (a, b, c).*

Proof. Equivalently, we must show there exist disjoint geodesics α, β, γ in $\mathbb{H}$ with $a = d(\alpha, \beta)$, $b = d(\alpha, \gamma)$ and $c = d(\beta, \gamma)$. This can be proved by continuity.

Normalize so that α is the imaginary axis, and choose any geodesic β at distance a from α. Then draw the 'parallel' line L of constant distance b from α, on the same side as β. This line is just a Euclidean ray in the upper half-plane. For each point $p \in L_p$ there is a unique geodesic γ_p tangent to L_p at p, and consequently at distance b from α.

Now consider $f(p) = d(\gamma_p, \beta)$. Then as p moves away from the juncture of α and L, $f(p)$ decreases from ∞ to 0, with strict monotonicity since $\gamma_p \cup L$ separates β from γ_q. Thus there is a unique p such that $f(p) = c$. ■

Doubling H along alternating edges, we obtain a *pair of pants* P. Thus $\mathrm{area}(P) = 2\pi$.

Corollary 13.3 *Given any triple of lengths $a, b, c > 0$, there exists a pair of pants, unique up to isometry, with boundary components of lengths (a, b, c).*

Pairs of pants decomposition.

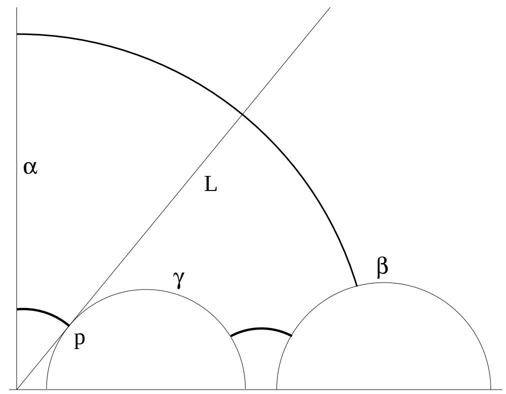

Figure 2. Right hexagons.

Theorem 13.4 *Any essential simple loop on a compact hyperbolic surface X is freely homotopic to a unique simple geodesic. Any two disjoint simple loops are homotopic to disjoint simple geodesics.*

Corollary 13.5 *Let X be a compact surface of genus g. Then X can be cut along $3g-3$ simple geodesics into $2g-2$ pairs of pants. In particular, we have*

$$\operatorname{area}(X) = 2\pi|\chi(X)|.$$

Parallels of a geodesic. There is a nice parameterization of the geodesic $|z| = 1$ in $\mathbb{H}$: namely

$$\delta(t) = \tanh t + i \operatorname{sech} t.$$

We have $\|\delta'(t)\| = 1$ in the hyperbolic metric.

Now given a closed simple geodesic γ on X, let $C(\gamma, r)$ be a parallel curve at distance r from γ. Then we have:

$$L(C(\gamma), r) = L(\gamma)\cosh(r).$$

Indeed, let γ can be covered by the imaginary axis $i\mathbb{R}_+$ in $\mathbb{H}$. Then $C(\gamma, r)$ is a ray from 0 to ∞ which passes through $\delta(r)$. Thus the Euclidean slope of $C(\gamma, r)$ is the same as that of the vector $(x, y) = (\tanh t, \operatorname{sech} t)$. Thus projection along Euclidean horizontal lines from $C(\gamma, r)$ to γ contracts by a factor of $y/\sqrt{x^2 + y^2} = \operatorname{sech}(t)$. Therefore $C(\gamma, r)$ is longer than γ by a factor of $\cosh(t)$.

The collar lemma.

Theorem 13.6 *Let α and β be disjoint simple geodesics on a compact Riemann surface, of lengths a and b respectively. Define A and B by* $\sinh(a/2)\sinh(A) = 1$ *and* $\sinh(b/2)\sinh(B) = 1$. *Then the collars of widths A and B about α and β are disjoint.*

Proof. We can assume that α and β are part of a pants decomposition of X, which reduces the result to the case where α and β are two cuffs of a pants P. By the Schwarz lemma, we can assume the lengths of the boundaries of P are $(a, b, 0)$.

Now cut P along a simple loop γ that begins and ends at its ideal boundary component, i.e. the cuff of length zero. Then the components of $P = \gamma$ are doubles of quadrilaterals with one ideal vertex and the remaining angles $\pi/2$. The quadrilateral meeting α has finite sides of lengths $a/2$ and A satisfying $\sinh(a/2)\sinh(A) = 1$, and similarly for the quadrilateral containing β. Thus these collars are disjoint. ■

Boundary of a collar.

Theorem 13.7 *The length of each component of the boundary of the standard collar around α with $L(\alpha) = a$ satisfies*

$$L(C(\alpha, r))^2 = a^2 \cosh^2(r) = \frac{a^2}{1 + \sinh^{-2}(a/2)} \to 4$$

as $a \to 0$. Thus the length of each component of the collar about α tends to 2 as $L(\alpha) \to 0$.

Proof. Apply the preceding formulas.

Check. The limiting case is the triply-punctured sphere, the double of an ideal triangle $T \subset \mathbb{H}$ with vertices $(-1, 1, \infty)$. The collars limit to the horocycles given by the circles of radius 1 resting on ± 1 together with the horizontal line segment H at height 2 running from $-1 + 2i$ to $+1 + 2i$. We have $L(H) = 1$, so upon doubling we obtain a collar boundary of length 2.

Corollary 13.8 *The collars about short geodesics on a compact hyperbolic surface cover the thin part of the surface.*

Proof. Suppose $x \in X$ lies in the thin part — that is, suppose there is a short essential loop δ through x. Then δ is homotopic to a closed geodesic γ, which is necessarily simple. But since δ is short, we see by the result above that δ must lie in the collar neighborhood of γ. ■

Thick-thin decomposition. There is a universal constant $r > 0$ such that any compact Riemann surface of genus g can be covered by a collection of $O(g)$ balls $B(x_i, r)$ and $O(g)$ standard collars about short geodesics.

Bers' constant.

Theorem 13.9 *There exists a constant L_g such that X admits a pants decomposition with no cuff longer than L_g.*

Theorem 13.10 *We can take $L_g = O(g)$, but there exist examples requiring at least one curve of length $> C\sqrt{g} > 0$.*

A finite-to-one map to $\mathcal{M}_g$.

Theorem 13.11 *For each trivalent graph of G with $b_1(G) = g$, there is a finite-to-one map*

$$\phi_G : ((0, L_g] \times S^1)^{3g-3} \to \mathcal{M}_g,$$

sending (r_i, θ_i) to the surface obtained by gluing together pants with cuffs of lengths r_i and twisting by θ_i, using pants and cuffs corresponding to vertices and edges of G.

The union of the images of the maps ϕ_G is all of $\mathcal{M}_g$.

Corollary 13.12 (Mumford) *The function $L : \mathcal{M}_g \to \mathbb{R}$ sending X to the length $L(X)$ of its shortest geodesic is proper.*

The Laplacian. Let M be a Riemannian manifold. The Laplace operator $\Delta : C_0^\infty(M) \to C_0^\infty(M)$ is defined so that

$$\int_M |\nabla f|^2 = \int_M f \Delta f,$$

both integrals taken with respect to the volume element on M.

For example, on $\mathbb{R}$ we find $\Delta f = -d^2 f/dx^2$ by integrating by parts. Similarly on $\mathbb{R}^n$ we obtain

$$\Delta f = -\sum \frac{d^2 f}{dx_i^2}.$$

Note this is the negative of the 'traditional' Laplacian.

In terms of the Hodge star we can write

$$\int \langle \nabla f, \nabla f \rangle = \int df \wedge *df = -\int f \wedge d * df = \int f(- * d * df)\, dV,$$

and therefore we have

$$\Delta f = - * d * df.$$

For example, on a Riemann surface with a conformal metric $\rho(z)|dz|$, we have $*dx = dy$, $*dy = -dx$, and

$$\Delta f = -\rho^{-2}(z)\left(\frac{d^2 f}{dx^2} + \frac{d^2 f}{dy^2}\right).$$

As a particular case, for $\rho = |dz|/y$ on $\mathbb{H}$ we see $\Delta y^\alpha = \alpha(1-\alpha)y^\alpha$, showing that y^α is an eigenfunction of the hyperbolic Laplacian.

The heat kernel. Let X be a compact hyperbolic surface. Enumerating the eigenvalues and eigenfunctions of the Laplacian, we obtain smooth functions satisfying

$$\Delta\phi_n = \lambda_n\phi_n,$$

$\lambda_n \geq 0$. The *heat kernel* $K_t(x, y)$ is defined by

$$K_t(x, y) = \sum e^{-\lambda_n t}\phi_n(x)\overline{\phi}_n(y),$$

The heat kernel is the fundamental solution to the *heat equation.* That is, for any smooth function f on X, the solution to the *heat equation*

$$\frac{df_t}{dt} = -\Delta f_t$$

with initial data $f_0 = f$ is given by $f_t = K_t * f$. Indeed, if $f(x) = \sum a_n\phi_n(x)$ then

$$f_t(x) = K_t * f = \sum a_n e^{-\lambda_n t}\phi_n(x)$$

clearly solves the heat equation and has $f_0(x) = f(x)$.

Note also that formally, convolution with K_t is the same as the operator $\exp(-\Delta)$, which acts by $\exp(-\lambda_n)$ on the λ_n-eigenspace.

Brownian motion. The heat kernel can also be interpreted using diffusion; namely, $K_t(x, \cdot)$ defines a probability measure on X that gives the distribution of a Brownian particle x_t satisfying $x_0 = x$.

For example, on the real line, the heat kernel is given by

$$K_t(x) = \frac{1}{\sqrt{4\pi t}}\exp(-x^2/(4t)).$$

Also we have $K_{s+t} = K_s * K_t$, as befits a Markov process.

To check this, note that K_t solves the heat equation, and that $\int K_t = 1$ for all t. Thus $K_t * f \to f$ as $t \to 0$, since K_t concentrates at the origin.

In terms of Brownian motion, the solution to the heat equation is given by $f_t(x) = E(f_0(x_t))$, where x_t is a random path with $x_0 = x$.

The trace. The *trace* of the heat kernel is the function

$$\operatorname{Tr} K_t = \int_X K_t(x,x) = \sum e^{-\lambda_n t}.$$

It is easy to see that the function $\operatorname{Tr} K_t$ determines the set of eigenvalues λ_n and their multiplicities.

Length spectrum and eigenvalue spectrum.

Theorem 13.13 *The length spectrum and genus of X determine the eigenvalues of the Laplacian on X.*

Proof. The proof is based on the trace of the heat kernel. Let $k_t(x,y)$ denote the heat kernel on the hyperbolic plane $\mathbb{H}$; it satisfies $k_t(x,y) = k_t(r)$ where $r = d(x,y)$. Then for $X = \mathbb{H}/\Gamma$ we have

$$K_t(x,y) = \sum_\Gamma k_t(x, \gamma y),$$

where we regard K_t as an equivariant kernel on $\mathbb{H}$.

Working more intrinsically on X, we can consider the set of pairs (x,δ) where δ is a loop in $\pi_1(X,x)$. Let $\ell_x(\delta)$ denote the length of the geodesic representative of δ based at x. Then we have:

$$K_t(x,x) = \sum_\delta k_t(\ell_x(\delta)).$$

Let $\mathcal{L}(X)$ denote the space of nontrivial free homotopy classes of maps $\gamma : S^1 \to X$. For each $\gamma \in \mathcal{L}(X)$ we can build a covering space $p : X_\gamma \to X$ corresponding to $\langle\gamma\rangle \subset \pi_1(X)$.

The points of X_γ correspond naturally to pairs (x,δ) on X with δ freely homotopic to γ. Indeed, given x' in X_γ, there is a unique homotopy class of loop δ' through x' that is freely homotopic to γ, and we can set $(x,\delta) = (p(x), p(\delta'))$. Conversely, given (x,δ), from the free homotopy of δ to γ we obtain a natural homotopy class of path joining x to γ, which uniquely determines the lift x' of x to X_γ.

For $x' \in X_\gamma$, let $r(x')$ denote the length of the unique geodesic through x' that is freely homotopic to γ. Then we have $\ell_x(\delta) = r(x')$. It follows that

$$\operatorname{Tr} K_t = \int_X K_t(x,x) = \int_X k_t(0) + \sum_{\mathcal{L}(X)} \int_{X_\gamma} r(x').$$

But $\int_X k_t(0) = \text{area}(X) k_t(0)$ depends only on the genus of X by Gauss-Bonnet, and the remaining terms depend only on the geometry of X_γ. Since the geometry of X_γ is determined by the length of γ, we see the length spectrum of X determines the trace of the heat kernel, and hence the spectrum of the Laplacian on X. ■

Remark. Almost nothing was used about the heat kernel in the proof. Indeed, the length spectrum of X determine the trace of any kernel $K(x, y)$ on X derived from a kernel $k(x, y)$ on $\mathbb{H}$ such that $k(x, y)$ depends only on $d(x, y)$.

Remark. In fact the genus is determined by the length spectrum.

Isospectral Riemann surfaces.

Theorem 13.14 *There exist a pair of compact hyperbolic Riemann surfaces X and Y, such that the length spectrum of X and Y agree (with multiplicities), but X is not isomorphic to Y.*

Isospectral subgroups. Here is a related problem in group theory. Let G be a finite group, and let H_1, H_2 be two subgroups of G. Suppose $|H_1 \cap C| = |H_2 \cap C$ for every conjugacy class C in G. Then are H_1 and H_2 conjugate in G?

The answer is *no* in general. A simple example can be given inside the group $G = S_6$. Consider the following two subgroups inside A_6, each isomorphic to $(\mathbb{Z}/2)^2$:

$$\begin{aligned} H_1 &= \langle e, (12)(34), (12)(56), (34)(56) \rangle, \\ H_2 &= \langle e, (12)(34), (13)(24), (14)(23) \rangle. \end{aligned}$$

Note that the second group actually sits inside A_4; it is related to the symmetries of a tetrahedron.

Now conjugacy classes in S_n correspond to permutations of n, i.e. cycle structures of permutations. Clearly $|H_i \cap C| = 3$ for the cycle structure $(ab)(cd)$, and $|H_i \cap C| = 0$ for other conjugacy classes (except that of the identity). Thus H_1 and H_2 are isospectral. But they are not conjugate ('internally isomorphic'), because H_1 has no fixed-points while H_2 has two.

Construction of isospectral manifolds.

Theorem 13.15 (Sunada) *Let $X \to Z$ be a finite regular covering of compact Riemannian manifolds with deck group G. Let $Y_i = X/H_i$, where H_1 and H_2 are isospectral subgroups of G. Then Y_1 and Y_2 are also isospectral.*

Proof. For simplicity of notation we consider a single manifold $Y = X/G$ and assume the geodesics on Z are discrete, as is case for a negatively curved manifold. Every closed geodesic on Y lies over a closed geodesic on Z.

Fixing a closed geodesic α on Y, we will show the set of lengths $\mathcal{L}$ of geodesics on Y lying over α depends only on the numbers $n_C = H \cap C$ for conjugacy classes C in G.

For simplicity, assume α has length 1. Let $\alpha_1, \ldots, \alpha_n$ denote the components of the preimage of α on X. Let $S_i \subset G$ be the stabilizer of α_i. The subgroups S_i fill out a single conjugacy class in G, and we have $S_i \cong \mathbb{Z}/m$ where $nm = |G|$. Each loop α_i has length m.

Let k be the index of $H \cap S_i$ in S_i. Then k is the length of α_i/H in Y. Moreover, the number of components α_j in the orbit $H \cdot \alpha_i$ is $|H|/|H \cap S_i| = |H||S_i|/k$, and of course all these components descend to a single loop on X/H. Thus the number of times the integer k occurs in $\mathcal{L}$ is exactly

$$|\mathcal{L}(k)| = \frac{kA_k}{m|H|},$$

where

$$A_k = |\{i \,:\, [S_i : S_i \cap H] = k\}|.$$

Thus to determine $\mathcal{L}$, it suffices to determine the integers A_k.

For example, let us compute A_1, the number of i such that we have $S_i \subset H$. Now H contains S_i if and only if H contains a generator g_i of S_i. We can choose the g_i's to fill out a single conjugacy class C, since the groups S_i are all conjugate. Then the proportion of i's satisfying $S_i \subset H$ is exactly $|H \cap C|/|C|$, and therefore

$$A_1 = \frac{n|H \cap C|}{|C|}.$$

An important point here: it can certainly happen that $S_i = S_j$ even when $i \neq j$. For example if G is abelian, then all the groups S_i are the same. But the number of i such that S_i is generated by a given element $g \in C$ is a constant, independent of g. Thus the *proportion* of S_i generated by an element of H is still $|H \cap C|/|C|$.

Now for $d|m$, let C_d be the dth powers of the elements in C. Then the subgroups of index d in the S_i's are exactly the cyclic subgroups generated by elements $g \in C_d$. Again, the correspondence is not exact, but constant-to-one; the number of i such that $\langle g \rangle \subset S_i$ is independent of $g \in C_d$. Thus the proportion of S_i's such that $H \cap S_i$ *contains* a subgroup of index d is exactly $|H \cap C_d|/|C_d|$, which implies:

$$\sum_{k|d} A_k = \frac{n|H \cap C_d|}{|C_d|}.$$

From these equations it is easy to compute A_k. ∎

Cayley graphs. The spaces in Sunada's construction do not have to be Riemannian manifolds. For example, we can take Z to be a bouquet of circles. Then X is the Cayley group of G, and Y_1 and Y_2 are coset graphs on which G acts. The coset graphs Y_1 and Y_2 also have the same length spectrum!

A small example. Let $G = (\mathbb{Z}/8)^* \ltimes \mathbb{Z}/8$ be the affine group of $A = \mathbb{Z}/8$, i.e. the group of invertible maps $f : A \to A$ of the form $f(x) = ax + b$. Let

$$\begin{aligned} H_1 &= \{x, 3x, 5x, 7x\}, \\ H_2 &= \{x, 3x+4, 5x+4, 7x\}. \end{aligned}$$

Then the subgroups H_1 and H_2 are isospectral.

In both cases, the coset space $Y_i = G/H_i$ can be identified with $\mathbb{Z}/8$; that is, $\mathbb{Z}/8 \times H_i = G$.

To make associated graphs, Y_1 and Y_2, we take $\langle x+1, 3x, 5x\rangle$ as generators for G. Note that $3x$ and $5x$ have order 2. Then the coset graph Y_1 is an octagon, coming from the generator $x+1$, with additional (unoriented, colored) edges joining x to $3x$ and $5x$. Similarly, Y_2 is also an octagon, but now the colored edges join x to the antipodes of $3x$ and $5x$, namely $3x+4$ and $5x+4$.

These graphs are isospectral. In counting the number of loops, it is important to regard the graphs as covering spaces. For this it is best to replace each colored edge which is *not* a loop by a pair of parallel edges with opposite arrows. Each colored loop should be replaced by a single oriented edge. Then the graphs become covering spaces of the bouquet of 3 circles, and the number of loops of length n is the same for both graphs.

Not isometric. Using short geodesics, we can arrange Z such that one can reconstruct the action of G on G/H_i from the intrinsic geometry of Y_i. Then Y_1 and Y_2 are isometric iff H_1 and H_2 are conjugate. So in this way we obtain isospectral, but non-isometric, Riemann surfaces.

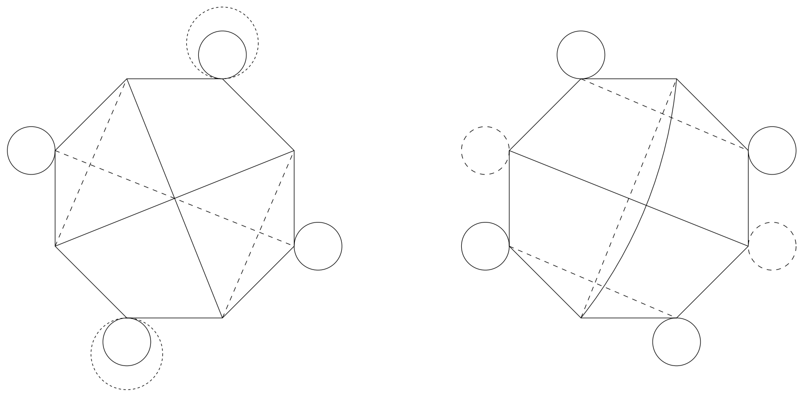

Figure 3. Isospectral graphs

14 Quasiconformal geometry

The measurable Riemann mapping theorem. For any μ on $\widehat{\mathbb{C}}$ with $\|\mu\|_\infty < 1$, there exists a quasiconformal map $f : \widehat{\mathbb{C}} \to \widehat{\mathbb{C}}$ with complex dilatation μ.

Corollary: Uniformization theorem. Evidentally we can uniformize at least one Riemann surface X_g of genus g, e.g. using a regular hyperbolic $4g$-gon. Now take any other surface Y of the same genus. By topology, there is a diffeomorphism $f : X_g \to Y$. Pulling back the complex structure to X_g and lifting to the universal cover, we obtain by qc conjugacy a Fuchsian group uniformizing Y.

References

[Hel] S. Helgason. *Differential Geometry, Lie Groups, and Symmetric Spaces.* Academic Press, 1978.

[KZ] M. Kontsevich and D. Zagier. Periods. In *Mathematics Unlimited — 2001 and Beyond*, pages 771–808. Springer-Verlag, 2001.

[La] S. Lang. *Algebra.* Addison-Wesley, 1984.