I · Background

1. Compact Riemann surfaces and algebraic curves

A Riemann surface is a topological space which in a neighbourhood of any point looks like the complex plane. (Some writers also require a Riemann surface to be connected.) More precisely, a Riemann surface is a Hausdorff space S which is endowed with a complex one-dimensional structure in the following way. For each point P of S we are given an open neighbourhood N of P and a homeomorphism ϕ from N to a disc $|z - \phi(P)| < c$ in the complex z-plane for some $c > 0$. These homeomorphisms satisfy the following consistency condition. Let P_1, P_2 be any two points of S such that their corresponding neighbourhoods N_1, N_2 overlap; then $\phi_1 \circ \phi_2^{-1}$ is holomorphic on $\phi_2(N_1 \cap N_2)$. If $f(z)$ is a function which is holomorphic and has non-zero derivative at $z = \phi(P)$, then we can replace ϕ by $f \circ \phi$ (with corresponding changes in N and c) without changing the complex structure of S. We call ϕ a local variable at P; clearly it is then a local variable at each point of N.

We can now transfer all the standard terminology of the theory of functions of a complex variable from the complex plane to Riemann surfaces. For example, a function ψ defined in a neighbourhood of P is holomorphic at P if $\psi \circ \phi^{-1}$ is holomorphic at $z = \phi(P)$, where ϕ is a local variable at P; the consistency condition is just what is needed to ensure that ψ is then holomorphic in a neighbourhood of P. Again, ψ has a zero of order n at P if $\psi \circ \phi^{-1}$ has a zero of order n at $z = \phi(P)$, and similarly for poles. A differential ω defined in a neighbourhood of P can be written in the form $\psi d\phi$; it is said to be holomorphic or meromorphic at P, or to have a zero or pole of order n at P, if ψ has the corresponding property. It is easy to see that none of these properties depends on the choice of the local variable ϕ. It must be remembered that a residue is associated with a differential rather than with a function. If $\omega = \psi d\phi$ has at worst an isolated singularity at P, then its residue at P is $(2\pi i)^{-1}\int\omega$ taken round a closed contour Γ which lies in N, goes once round P anticlockwise under the orientation induced in N by ϕ, and contains no singularity of ψ inside or on it except perhaps P itself. It follows that the residue is equal to the coefficient of $1/(\phi - \phi(P))$ in the Laurent series expansion of ψ in terms of $\phi - \phi(P)$, and also that it is independent of the choice of ϕ.

In general one is primarily interested in those functions and differentials on a Riemann surface which are everywhere meromorphic; and in what follows we shall confine ourselves to them. A differential is said to be of the first kind on a Riemann surface S if it is holomorphic everywhere on S, of the second kind if it is meromorphic on S and has residue zero at every pole, and of the third kind if it is meromorphic on S; differentials of any given kind form a **C**-vector space. A differential is called exact if it is of the form $d\psi$ where ψ is meromorphic on S; clearly an exact differential is of the second kind, but not every differential of the second kind is exact.

We can now translate Cauchy's theorem into the language of Riemann surfaces:

Theorem 1. Let Γ be a contour on S, not necessarily closed, and let ω be a meromorphic differential on S. Suppose that Γ varies continuously, with its endpoints if any remaining fixed. Then $\int_\Gamma\omega$ remains constant as long as Γ does not move across a singularity of ω, and it changes by $2\pi i$ times the residue at a singularity whenever Γ moves across that singularity.

In particular, if Γ is a closed curve and ω is of the first or second kind, the value of $\int_\Gamma\omega$ depends only on the homology class of Γ; this value is called the period of ω with respect to Γ. It is easy to see that ω is exact if and only if all its periods are equal to zero.

Henceforth we shall also assume that S is compact, and usually that it is connected. (A compact Riemann surface is the disjoint union of finitely many compact connected Riemann surfaces, so that results in the general case follow at once from those in the connected case; but they are often more complicated to state in the general case.) At each point of S there is a canonical local orientation, obtained by means of

ϕ^{-1} from the canonical orientation of the complex plane. These orientations are compatible, so S itself is oriented. Since it is compact it can be triangulated and the triangulation is finite; topologically S is just a sphere with g handles, for some $g \geq 0$, and its first homology group $H_1(S, \mathbf{Z})$ is a free abelian group on 2g generators.

Theorem 2. (i) If ω is a meromorphic differential on a compact Riemann surface S then ω has only finitely many poles on S and the sum of the residues of ω at these poles is 0.

(ii) If ψ is a non-constant meromorphic function on a compact Riemann surface S then ψ takes every value the same number of times (allowing for multiplicities). In particular ψ has as many poles as zeros.

Proof. If ω had infinitely many poles these would have a point of accumulation on S, by compactness, and ω would not be meromorphic at that point. Now triangulate S, integrate ω round each triangle and sum the results; (i) follows at once since in the sum ω has been integrated twice along each side of each triangle, once in each direction. Applying these results to the special case $\omega = d\psi/\psi$ we find that ψ has as many poles as zeros, and writing $\psi - c$ for ψ we find that ψ takes the value c as many times as it has poles. This completes the proof of the Theorem. The number of times ψ takes each value is called the valence of ψ. Constant functions are deemed to have valence 0.

Lemma 3. Let θ, ψ be non-constant meromorphic functions on a compact Riemann surface S, of valences m, n respectively. Then there is a polynomial F(X, Y) of degrees at most n in X and m in Y such that $F(\theta, \phi) = 0$. If moreover S is connected then F can be chosen to be irreducible.

Proof. Choose a complex number c which is not a value taken by ψ at any of the zeros or poles of θ, and let $P_1, \ldots, P_n$ be the points of S at which $\psi = c$. For any $r \geq 0$ the differential $\omega = \theta^r d\psi/(\psi - c)$ has a simple pole with residue $\theta^r(P_\nu)$ at each P_ν; its other poles are at the poles of θ and ψ, and at each of these the residue is a rational function of c. It now follows from Theorem 2(i) that $\Sigma\, \theta^r(P_\nu)$ is a rational function of c for each r, so that the same is true for the elementary symmetric functions of the $\theta(P_\nu)$. Thus θ satisfies an equation of degree n whose coefficients are rational functions of ψ; and hence the minimal equation connecting θ and ψ has degree at most n in θ and hence by symmetry has degree at most m in ψ.

Now suppose that S is connected and that $F(\theta, \psi) = 0$ is the minimal equation. If F is not irreducible then we can write $F = F_1 F_2$ for some polynomials F_1 and F_2. Choose a point P on S and let N be the canonical neighbourhood of P associated with some local variable at P; then

$$F_1(\theta, \psi) F_2(\theta, \psi) = 0 \text{ in } N,$$

and so (by the analogous result for the complex plane) either $F_1(\theta, \psi) = 0$ in N or $F_2(\theta, \psi) = 0$ in N. Assume the former; then by analytic continuation $F_1(\theta, \psi) = 0$ everywhere on S, which contradicts the minimality of F. This completes the proof of the Lemma.

Theorem 4. The meromorphic functions on a compact connected Riemann surface S form a finitely generated field of transcendence degree 1 over $\mathbf{C}$. Given any two points P_1 and P_2 on S, there is a meromorphic function ψ on S such that $\psi(P_1) \neq \psi(P_2)$. Moreover, given any point P on S there is a meromorphic function ϕ on S which is a local variable at P.

The difficult part of the proof is the second sentence, that there are enough meromorphic functions on S; in many practical cases this is irrelevant, since if one is presented with an explicit S it is likely to come equipped with plenty of functions. Given a non-constant meromorphic ψ on S, the proof of the first sentence is easy. For suppose that ψ has valence $n > 0$; then by Lemma 3 any meromorphic function on S is algebraic of degree at most n over $\mathbf{C}(\psi)$. Since the characteristic is zero, it follows from standard results in field theory that the field of all meromorphic functions on S is algebraic of degree at most n over $\mathbf{C}(\psi)$.

Corollary. Any compact connected Riemann surface can be regarded as an irreducible non-singular algebraic curve over **C**, and vice versa.

Because the functions on S separate points, S can be reconstructed from a knowledge of the field of meromorphic functions on it; and this field is of the sort that corresponds to an irreducible curve. The curve is non-singular because the field contains a local variable at each point. For the converse, we need only check that the Riemann surface corresponding to an irreducible curve is connected; if this were not so, the connected components of the Riemann surface would give rise to components of the curve.

The additive group of divisors on S is the free abelian group whose generators are the points of S. Let ψ be a meromorphic function on S which is not identically zero on any connected component of S, and let $P_1, \ldots, P_n$ be the poles and $Q_1, \ldots, Q_n$ be the zeros of ψ, repeated according to their multiplicities; then the divisor of ψ is defined to be

$$(\psi) = Q_1 + \ldots + Q_n - P_1 - \ldots - P_n.$$

The divisor of a differential ω is defined in a similar way and is written as (ω), though of course a differential need not have as many poles as zeros. Taking divisors is a homomorphism from the multiplicative group of non-zero meromorphic functions on S to the additive group of divisors. The kernel of this homomorphism is just the functions with no poles or zeros; these have valence 0 and are therefore constant on each connected component of S. In particular, if S is connected the divisor of a function on S determines the function up to multiplication by a non-zero constant. The image of the homomorphism, that is the group of divisors of functions, is called the group of principal divisors; and two divisors are said to be linearly equivalent if their difference is a principal divisor. The degree of a divisor $\Sigma n_\nu P_\nu$ is defined to be Σn_ν; thus linearly equivalent divisors have the same degree, but not necessarily vice versa. A divisor $\Sigma n_\nu P_\nu$ is said to be positive if every $n_\nu \geq 0$; this induces on the group of divisors a partial ordering which is compatible with addition. Since the quotient of two differentials is a function, and the product of a function and a differential is a differential, it is easy to see that the divisors of differentials precisely fill a linear equivalence class; this is called the canonical class and every divisor in it is called a canonical divisor.

There are two key theorems which tell one something about the structure of the group of divisors, the Riemann-Roch theorem which gives information about principal divisors and Abel's theorem which describes the group of divisor classes modulo linear equivalence. To state the Riemann-Roch theorem we need some notation. Let S be a connected compact Riemann surface, and $\mathfrak{a}$ a divisor on S. Denote by $L(\mathfrak{a})$ the $\mathbf{C}$-vector space consisting of those meromorphic functions ψ on S for which $(\psi) + \mathfrak{a} \geq 0$, together with the zero function, and write $l(\mathfrak{a}) = \dim L(\mathfrak{a})$; thus $l(\mathfrak{a})$ depends only on the linear equivalence class of $\mathfrak{a}$. Since $\deg((\psi)) = 0$ for any non-zero function ψ on S, $L(\mathfrak{a}) = \{0\}$ and $l(\mathfrak{a}) = 0$ whenever $\deg(\mathfrak{a}) < 0$.

Theorem 5 (Riemann-Roch). Let S be a compact connected Riemann surface. Then there is an integer $g \geq 0$, depending only on S, such that

$$l(\mathfrak{a}) = \deg(\mathfrak{a}) + 1 - g + l(\mathfrak{k} - \mathfrak{a}) \tag{1}$$

for any divisor $\mathfrak{a}$ and any canonical divisor $\mathfrak{k}$.

The condition that S should be connected may be dispensed with, if we suitably modify the definition of $L(\mathfrak{a})$ and allow g to be negative. However this is not a real generalization, for the equation (1) for general S can be obtained by addition of the corresponding equations for the connected components of S.

Corollary 1. $\mathrm{Deg}(\mathfrak{k}) = 2g - 2$ and $l(\mathfrak{k}) = g$, if S is connected.

Proof. Applying Theorem 5 to $\mathfrak{k} - \mathfrak{a}$ instead of $\mathfrak{a}$ gives

$$l(\mathfrak{k} - \mathfrak{a}) = \deg(\mathfrak{k}) - \deg(\mathfrak{a}) + 1 - g + l(\mathfrak{a}),$$

and comparing this with (1) gives $\deg(\mathfrak{k}) = 2g - 2$. Now take $\mathfrak{a} = \mathfrak{k}$ and note that $l(0) = 1$ because $L(0) = \mathbf{C}$; thus $l(\mathfrak{k}) = g$.

Cambridge University Press
978-0-521-20526-9 - Analytic Theory of Abelian Varieties
H. P. F. Swinnerton-Dyer
Excerpt
More information

The Riemann-Roch theorem is a duality theorem, relating $l(\mathfrak{a})$ and $l(\mathfrak{k} - \mathfrak{a})$; it can be written in self-dual form but nothing is gained by doing so. In the special case when $\deg(\mathfrak{a}) > 2g - 2$ we have $\deg(\mathfrak{k} - \mathfrak{a}) < 0$ and hence $l(\mathfrak{k} - \mathfrak{a}) = 0$; so (1) takes the form

$$l(\mathfrak{a}) = \deg(\mathfrak{a}) + 1 - g \quad \text{if} \quad \deg(\mathfrak{a}) > 2g - 2, \tag{2}$$

which is the form in which it is most frequently used.

Corollary 2. The differentials of the first kind form a **C**-vector space of dimension g, provided S is connected.

Proof. This is just the equation $l(\mathfrak{k}) = g$ in a new form. For let ω be a given non-zero differential; then the differentials are just the $\psi\omega$ where ψ runs through all meromorphic functions on S, and $\psi\omega$ is a differential of the first kind if and only if ψ is in $L((\omega))$, which has dimension g.

The statements of both these Corollaries need to be modified if S is not connected; in that case both $l(\mathfrak{k})$ and the dimension of the space of differentials of the first kind are equal to $(g - 1)$ plus the number of connected components of S.

The integer g in Theorem 5 is called the genus of S; if S is connected its genus is equal to the g which we have already defined topologically by the condition that $H_1(S, \mathbf{Z})$ is a free abelian group on $2g$ generators. To prove this we consider maps from one Riemann surface to another. Let S_1 and S_2 be compact connected Riemann surfaces; then a map $f : S_1 \to S_2$ is said to be holomorphic if for any point P_1 on S_1 and any local variable ϕ_2 at $f(P_1)$ on S_2 the function $\phi_2 \circ f$ is holomorphic at P_1. The point P_1 is said to be a ramification point of order r for f if $\{\phi_2 \circ f - \phi_2 \circ f(P_1)\}$ has a zero of order r at P_1 and $r > 1$; it follows from the compactness of S_1 that if f is non-constant it has only finitely many points of ramification. Moreover if P_2 is a point of S_2 then the degree of $f^{-1}(P_2)$ does not depend on the choice of P_2 provided that points of ramification in $f^{-1}(P_2)$ are taken with multiplicities equal to their orders of ramification; if this degree is n then we say that f is an n-to-1 map.

Cambridge University Press
978-0-521-20526-9 - Analytic Theory of Abelian Varieties
H. P. F. Swinnerton-Dyer
Excerpt
More information

Theorem 6. Let $f : S_1 \to S_2$ be a non-constant holomorphic n-to-1 map between two compact connected Riemann surfaces, and let $r_1, \ldots, r_m$ be the orders of the ramification points of f. Then

$$2g_1 - 2 = n(2g_2 - 2) + \Sigma(r_\mu - 1) \tag{3}$$

where g_1 is the genus of S_1 and g_2 the genus of S_2.

Proof. Let ω_2 be a differential on S_2 and assume for convenience of description that none of the images of points of ramification of f are zeros or poles of ω_2. By means of f we can pull ω_2 back to a differential ω_1 on S_1. With the notation above, if P_1 is a point of S_1 which is not a ramification point of f then $\phi_2 \circ f$ is a local variable at P_1; and it follows that ω_1 has a zero (pole) at P_1 if and only if ω_2 has a zero (pole) at $f(P_1)$, and both zeros (poles) have the same multiplicity. On the other hand, if P_1 is a ramification point of order r then $d(\phi_2 \circ f)$ has a zero of order $r - 1$ at P_1 and hence so has ω_1. Thus the divisor of ω_1 is equal to the sum of the pull-back of the divisor of ω_2 and all the ramification points of f, a point of order r being taken with multiplicity $r - 1$. To prove (3) we now evaluate the degree of (ω_1) both directly and through this decomposition.

In particular let S_2 be the complex plane including the point at infinity; then f can be any non-constant meromorphic function on S_1 and n is the valence of f. Moreover $g_2 = 0$ since the differential dz has a double pole at infinity; so (3) reduces to

$$2g_1 = 2 - 2n + \Sigma(r_\mu - 1). \tag{4}$$

If one now constructs a triangulation of S_2 which has all the images of ramification points among its vertices, and lifts it to a triangulation of S_1, a straightforward calculation shows that the rank of $H_1(S_1, \mathbf{Z})$ is $2g_1$.

In the classification of connected Riemann surfaces, the genus is the only discrete-valued invariant that occurs; all connected Riemann surfaces of the same genus form a single continuous system. How can one most easily calculate the genus of a given connected Riemann surface? If the surface is defined topologically then one should use the topological

Cambridge University Press
978-0-521-20526-9 - Analytic Theory of Abelian Varieties
H. P. F. Swinnerton-Dyer
Excerpt
More information

description of the genus; in particular, if the surface is given with a triangulation, as in modular function theory for example, one can use the formula

$$2g - 2 = e - v - f$$

where the triangulation has v vertices, e edges and f faces. If the surface is defined by means of its function field one should use either $\deg((\omega)) = 2g - 2$ for some suitably chosen differential ω or else one of the covering formulae (3) and (4); the various formulae that occur in algebraic geometry can be deduced without difficulty from these.

Now let Ω denote the **C**-vector space of differentials of the first kind on a compact connected Riemann surface S of genus $g > 0$. By Theorem 1 the pairing

$$(\omega, \Gamma) \to \int_\Gamma \omega$$

induces a map

$$\Omega \times H_1(S, \mathbf{Z}) \to \mathbf{C}$$

which is **C**-linear in the first argument and additive in the second; so this map induces a canonical homomorphism

$$H_1(S, \mathbf{Z}) \to \Omega^* = \mathrm{Hom}(\Omega, \mathbf{C}). \tag{5}$$

Call the image of this map Λ; it is obviously a free abelian group on at most 2g generators. Now let O be a fixed point and P an arbitrary point of S; to an arc OP corresponds an element of Ω^* given by $\omega \to \int_O^P \omega$. If we change the arc OP, leaving O and P the same, this element of Ω^* is changed by an element of Λ; thus to P corresponds an element of Ω^*/Λ, and by additivity this correspondence can be extended to a homomorphism

$$\{\text{Divisors on } S\} \to \Omega^*/\Lambda. \tag{6}$$

If we restrict this homomorphism to divisors of degree zero, it no longer depends on the choice of the point O; and it is just what we need to pick

CAMBRIDGE

Cambridge University Press
978-0-521-20526-9 - Analytic Theory of Abelian Varieties
H. P. F. Swinnerton-Dyer
Excerpt
More information

out the principal divisors.

Theorem 7 (Abel). With the notation above, Λ is a lattice in Ω^* - that is, a free abelian group on $2g$ generators which spans Ω^* considered as a real vector space. Moreover the homomorphism (6), restricted to divisors of degree zero, is onto and its kernel is just the group of principal divisors.

Another way of stating the first sentence is to say that (5) induces a homeomorphism between $H_1(S, Z)$ with the discrete topology and its image Λ with the topology induced on it as a subset of Ω^*.

Abel's theorem identifies the group of divisor classes of degree zero with the complex torus Ω^*/Λ of dimension g, which is also called the Jacobian of S. Not surprisingly, this complex torus is actually an abelian manifold - that is, it has enough meromorphic functions on it to separate points. This will be proved in §9; for the proof we shall need further information about Λ or, which comes to the same thing, about the periods associated with S. This information is contained in Riemann's relations; to state them it is convenient to choose a normalized base for $H_1(S, Z)$. So let $\Gamma_1, \ldots, \Gamma_g, \Gamma'_1, \ldots, \Gamma'_g$ be closed curves on S such that

(i) the corresponding classes generate $H_1(S, Z)$; and

(ii) no two of these curves have a point in common,

except that for each ν the curves Γ_ν and Γ'_ν have one common point at which they cross with the orientations shown in the diagram.

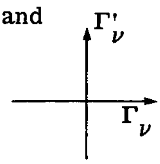

Such curves can be found, for if we consider S as a sphere with g handles we can choose Γ_ν to go once round the ν^{th} handle and Γ'_ν to run along the ν^{th} handle and back along the surface of the sphere. By considering intersection numbers we see first that the homology classes of these $2g$ curves are linearly independent and then that they form a base for $H_1(S, Z)$. Moreover any closed curve which does not meet any Γ_ν or Γ'_ν must be homologically trivial.

Theorem 8. (i) Let ω_1, ω_2 be differentials of the first kind on S and let $c_{\mu\nu}, c'_{\mu\nu}$ be the periods of ω_μ with respect to Γ_ν, Γ'_ν