RIEMANN–ROCH THEOREM

VALERIYA TALOVIKOVA

Abstract. This paper presents a brief introduction to algebraic geometry and provides several theorems and lemmas with proofs, along with many examples that allow one to gain a deeper intuitive understanding of the material. Finally, having worked through the prerequisites, this paper demonstrates an elementary proof of the Riemann–Roch Theorem, which is a vital tool to the fields of complex analysis and algebraic geometry. It is used for the computation of the dimension of the space of meromorphic functions with prescribed zeroes and allowed poles. It also relates the complex analysis of a compact, connected Riemann surface with the surface's purely topological genus, in a way that can be carried over into purely algebraic settings.

Contents

Some preliminaries for this article can be found at [1], Preliminaries. Necessary information on *coverings* can be read at [3], Ch.5.

1. Introduction to Riemann surfaces

In order to understand the statement and proof of the Riemann–Roch Theorem, one must first state several basic definitions about Riemann surfaces, hence:

Definition 1.1. A *Riemann surface* is a one-dimensional, connected, complex manifold. This means that every Riemann surface is a union of connected open subsets, which are homeomorphic to open sets in $\mathbb{C}$ (these homeomorphisms are called *[local] charts*, and the full collection of such charts is called an *[complex] atlas*). Moreover, for each two such neighborhoods having a non-empty intersection, there must exist a conformal isomorphism between the two images of the intersection created correspondingly by two individual neighborhood isomorphisms.

A more formal definition of a [complex] atlas and [local] charts can be found at [1], Definition 2.1.1.

Date: AUG 20, 2009.

Any conected open set in $\mathbb{C}$ is also is a non-compact Riemann surface. An atlas of such surfaces can contain just one chart and unlike more complicated surfaces there is no necessity in change of coordinates.

Remark 1.2. In the following, unless explicitly stated otherwise, we assume that our Riemann surfaces are compact topological spaces.

Definition 1.3. Any compact Riemann surface is homeomorphic to a sphere with $g > 0$ handles attached, and the number g is called the *genus* of the Riemann surface.

In particular, the simplest example of a compact Riemann surface is the Riemann sphere (i.e. a sphere with *zero* handles).

Fact 1.1 (Riemann Classification Theorem). Every compact Riemann surface is homeomorphic to a sphere with g handles for a certain g. A sphere with zero handles is simply a sphere. A sphere with one handle homeomorphic is to a torus. The proof of this Theorem can be found at [3], Ch 1. Another example of a compact Riemann surface is a torus. The proof that a torus is, in fact, a Riemann surface can be found at [1] (Example 2.2.3, Proposition 2.2.4).

Definition 1.4. The *Riemann sphere* is the set $\overline{\mathbb{C}} = \mathbb{C} \cup \infty$ with the neighborhoods $U_0 = \mathbb{C}$ and $U_1 = (\mathbb{C} \setminus \{0\}) \cup \{\infty\}$ and corresponding charts given by:

$$\phi_0(z) = z \quad \text{and} \quad \phi_1(z) = \begin{cases} \frac{1}{z}, & \text{if } z \neq \infty \\ 0, & \text{if } z = \infty \end{cases}$$

The proof that Riemann sphere is, indeed, a Riemann surface can be found at [1] (Proposition 2.2.2).

2. Holomorphic and meromorphic functions on Riemann surfaces

Definition 2.1. (Cf. [1], Definition 2.3.1) Let X and Y be Riemann surfaces that are not necessarily compact. A function $f : X \to Y$ is called *holomorphic* if for all charts $\phi : U_1 \to V_1$ on X, and $\psi : U_2 \to V_2$ on Y the function $\psi \circ f \circ \phi^{-1}$ is holomorphic (in usual sense) on $\phi(U_1 \cap f^{-1}(U_2))$.

Proposition 2.2. *Let X be a compact Riemann surface. If $f : X \to \mathbb{C}$ is holomorphic, then f is constant.*

Proof. Can be found at [1], Theorem 2.3.3. □

Definition 2.3. Let X be a Riemann surface that is not necessarily compact. A function f on X is called *meromorphic* if f is holomorphic on $X \backslash S$, where $S \subset X$ is a discrete (possibly empty) subset of X, and f has poles at every point $x \in S$.

Obviously, any holomorphic function is also meromorphic; the sums, differences, products, and quotients of two meromorphic functions are also meromorphic.

Proposition 2.4. *There exists a natural bijection between the set of meromorphic functions on a Riemann surface X, which is not necessarily compact, and the set of holomorphic functions (maps) from X to $\overline{\mathbb{C}}$.*

Proof. If f is our meromorphic function on X, and S is the set of its poles, we can define a map $\bar{f} : X \to \overline{\mathbb{C}}$ by the following rule: $\bar{f}(x) = \infty$, if $x \in S$, and $\bar{f}(x) = f(x)$, otherwise. It is easy to see that $\bar{f}$ is holomorphic: for $x \notin S$, there is no contradiction because f is already holomorphic at x. If $x \in S$, then the composition of f with the chart $\phi_1 : U_1 \to \mathbb{C}$ is $z \mapsto \frac{1}{f(z)}$, and if f has a pole at x, then $\frac{1}{f(z)}$ is bounded in the neighborhood of x and can be made holomorphic in accordance with the removable singularity theorem. Thus, $\bar{f}$ is also holomorphic (since its local representation in the neighborhood of x is.) The converse is also true: if we have a holomorphic map $\bar{f}$, we can easily build a meromorphic function f by restriction of $\bar{f}$ to pre-images of $\mathbb{C}$. Using same ideas we can see that f has poles in the points at the pre-images of ∞. □

Fact 2.1 (Riemann's Existence Theorem). Every compact Riemann surface admits a non-constant meromorphic function. The proof of the Theorem can be found at [10].

Definition 2.5. If $f : X \to Y$ is a holomorphic map of Riemann surfaces and a is a point of X, then we can find a local chart neighborhood U of a, defined by the coordinate z, such that a corresponds to $z = 0$. We can similarly find a chart for $f(U)$ under which $f(a)$ corresponds to zero. Under these coordinates, f in U will be represented by a Taylor series $c_n z^n + c_{n+1} z^{n+1} + \cdots$, where $n \geq 1$. We will call such coordinates *local coordinates*, and the corresponding Taylor series as the *local representation* for f.

3. Coverings

Definition 3.1. A local homeomorphism $\pi : M' \to M$, where M and M' are manifolds, is called a *covering* if each $x \in M$ has a connected neighborhood V such that every connected component of $\pi^{-1}(V)$ is mapped by π homeomorphically onto V. If π is clear from the context, we sometimes also call M' a *covering space* of M, and π is the *covering map.* The homeomorphic copies in M' of the connected neighborhood V are called *sheets.* For any point $x \in M$ the inverse image of x in M' is called the *fiber* over x.

For every $x \in M$, the fiber over x is a discrete subset of M'. If M is connected, the fibers are homeomorphic and so there is a discrete space F such that the fiber over each $x \in M$ is homeomorphic to F. Moreover, for every $x \in M$ there is a neighborhood V of x such that its full pre-image $\pi^{-1}(V)$ is homeomorphic to $V \times F$. In particular, the cardinality of the fiber over x is equal to the cardinality of F and it is called the *degree* of the cover π. Thus, if every fiber has n elements, we refer to it as an *n-fold covering.*

The following proposition gives a useful criterion to determine if a map is a covering; its proof can be found at page 56 of [2].

Proposition 3.2. *([2], Proposition 7.2) Let $f : M' \to M$ be a surjective continuous mapping of manifolds that satisfies the following conditions:*

(1) f is a local homeomorphism
(2) for each compact subset $K \subset M$, the set $f^{-1}(K) \subset M'$ is also compact.

Then, f is a covering map and for each $x \in M$ the fiber of x is finite.

Definition 3.3. Let f be a holomorphic map from a Riemann surface X to a Riemann surface Y and let a be a point of X. The winding number of $f(X)$ with respect to the point $f(a)$ is a positive integer called the *ramification index* of a. If the ramification index is greater than 1, then a is called a *ramification point* of f, and the corresponding value $f(a)$ is called a (algebraic) *branch point*. Equivalently, a is a ramification point with ramification index k if there exists a holomorphic function ϕ defined in a neighborhood of a such that $f(z) = \phi(z)(z-a)^k$ for some positive integer $k > 1$. If z is a local coordinate, f can be presented as $\phi(z)z^k$.

Proposition 3.4. *([2], pg. 55) Let f be a holomorphic map from a Riemann surface X to a Riemann surface Y; then:*

(1) $f(X) = Y$
(2) For each $y \in Y$, the fiber of y is a finite set.
(3) The number of ramification points on X is finite; If $S \subset Y$ is the set containing images of all ramification points of X, then restriction of f on $X \backslash f^{-1}(S) \to Y \backslash S$ is a covering map.
(4) If the degree of this covering map equals n and $f^{-1}(y) = \{x_1, \dots, x_k\}$ is the pre-image of a point $y \in Y$, then $\sum_{i=1}^{k} e_i = n$, where e_i is the ramification index of x_i.

Proof. (1) $f(X)$ is closed in Y, because X is compact and f is continuous, and $f(X)$ is open in Y, because f is holomorphic. Since Y is connected, it must be that $f(X) = Y$.

(2) $f^{-1}(y)$ cannot have an accumulation point x_0 in X (otherwise, f would be constant in a neighborhood of x_0 and, hence, everywhere). So, $f^{-1}(y)$ is discrete and (since X is a compact) finite.

(3) From the local representation viewpoint,a ramification point is a zero of a derivative of f (with f presented locally as $\phi(z)z^k$). By the same reasoning zeroes of the derivative cannot have an accumulation point (otherwise, the derivative will be zero everywhere, and f would be a constant). So, again this set is discrete and finite. To prove that restriction of f to $X \backslash f^{-1}(S) \to Y \backslash S$ is a covering map we can use 3.2, with $M' = X \backslash f^{-1}(S)$ and $M = Y \backslash S$. Second condition of 3.2 is satisfied because X is compact.

(4) We can find neighborhoods U_i for each x_i such that the local representation for f in U_i behaves like z^{e_i}, so restriction of f on $U_i \backslash \{x_i\}$ covers its image e_i times. Since, for all points y' close enough to y, we have $f^{-1}(y') \subset \cup_i U_i$, the equality $\sum_{i=1}^{k} e_i = n$ follows from the fact that all points $y' \neq y$ close enough to y have n pre-images.

□

Proposition 3.5. *The number of zeroes is equal to the number of poles for any meromorphic function on a Riemann surface.*

Proof. In virtue of 2.4 to the meromorphic function f corresponds a holomorphic map $\bar{f}$ from our Riemann surface X to $\overline{\mathbb{C}}$. By 3.4, to $\bar{f}$ corresponds a covering map from X to $\overline{\mathbb{C}}$ and the degree of the covering map is the number of pre-images for any point of $\overline{\mathbb{C}}$ counted with multiplicities (ramification indexes).

The degree is the same for every point, in particular, for 0 and ∞. Thus, the number of pre-images of 0 counted with multiplicities (zeroes of f) is the same as the number of pre-images of ∞ counted with multiplicities (poles of f). □

Theorem 3.1. *Fundamental Theorem of Algebra* Every non-constant, single-variable, degree-n polynomial with complex coefficients has n complex roots (counted with multiplicities).

Proof. Our polynomial f is a meromorphic function $f(z) = w$ on the Riemann surface $\overline{\mathbb{C}}$(the details can be found in [2], Lecture 6). Hence, f has to have as many zeroes as poles (counted with multiplicities). Clearly, f has just one pole of order n (at infinity), so the number of zeroes of f also is n. □

Definition 3.6. We can define the *Euler characteristic* for a Riemann surface X in the following way. Let us triangulate the Riemann surface (we assume that this is always possible). The Euler characteristic of the surface is then $\chi = V - E + F$, where V, E, and F are respectively the numbers of vertices, edges, and faces of our triangulation. It is well-known that Euler characteristic χ is a *topological invariant* and, in particular, does not depend on the triangulation itself.

It is also well-known that if the genus of our Riemann surface is equal to g, then $\chi = 2-2g$. (For the proof, we could describe our Riemann surface as a sphere with g handles, triangulate it, and find the Euler characteristic directly).

Theorem 3.2. *Riemann-Hurwitz formula* (Cf. [2], pp.57, 58) Let $f : X \to Y$ be a non-constant holomorphic map of one compact Riemann surface into another. If $\deg f = n$, the genus of X is equal to $g(X)$, the genus of Y is equal to $g(Y)$, and f is ramified at m points on X with indexes of ramification $e_1, \ldots, e_m$ correspondingly.

$$2 - 2g(X) = n(2 - 2g(Y)) - \sum_{i=1}^{m} (e_i - 1). \tag{3.7}$$

Proof. We will triangulate the surface Y in such way that all branch points are among the vertices of the triangulation. Then, f will provide for a branched covering map from X to Y. If V, E, and F are correspondingly the number of vertices, edges and faces of the triangulation, then the Euler characteristic $\chi = 2-2g(Y) = V - E + F$. If we consider the pre-image of our Y triangulation on X, we will get a triangulation on X, with the number of edges $E' = nE$, the number of faces

$F' = nF$ and the number of vertices $V' = nV - \sum_{i=1}^{m}(e_i - 1)$ (because a single branch point i will correspond to e_i regular points).

Now, we can calculate the Euler characteristic for X. It is:

$$2-2g(X) = V'-E'+F' = nV - \sum_{i=1}^{m}(e_i-1) - nE + nF = n(2-2g(Y)) - \sum_{i=1}^{m}(e_i-1).$$

□

4. Holomorphic and meromorphic 1-forms on Riemann surfaces

Definition 4.1. A *holomorphic 1-form* on a Riemann surface X is a complex differential form of degree 1 on X that can be written in local coordinates as $\omega = f dz$, where function f is a holomorphic function of the local coordinate z.

If we change our coordinate in a holomorphic way: $z = z(w)$, where $z(w)$ is a holomorphic function of argument w, then $f(z)dz = f(z(w))z'(w)dw$. So if a 1-form can be written as $f dz$ with holomorphic f for *some* coordinate z, it can be written this way for *any* coordinate.

Definition 4.2. A *meromorphic 1-form* on a Riemann surface X is a complex differential form of degree 1 on X that is *holomorphic* on $X \backslash S$, where $S \subset X$ is a discrete (possibly empty) subset of X, such that any point $a \in S$ has a neighborhood U on which the restriction of ω can be written as $f\omega'$, where f is a meromorphic function on U with a pole at a, and ω' is a holomorphic 1-form on U.

Informally, a meromorphic 1-form ω is something that can be written in local coordinates as $f(z)dz$, where f is a meromorphic function. *ω has a pole* at the point a if and only if $f(z)$ has a pole at a, provided that ω can be written as $f(z)dz$ in a neighborhood of the point a.

Proposition 4.3. *If ω_1 and ω_2 are meromorphic 1-forms on a compact Riemann surface X then there exists a meromorphic function f for which $\omega_1 = f\omega_2$.*

Proof. If in some neighborhood $U \subset X$ our 1-forms can be written as $\omega_1 = f_1 dz$, $\omega_2 = f_2 dz$, where f_1 and f_2 are meromorphic functions on U and z is a local coordinate, we will set f_U to be the meromorphic function f_1/f_2 on U. If in some other neighborhood $V \subset X$ having non-empty intersection with U, we have $\omega_1 = g_1 dw$, $\omega_2 = g_2 dw$, then on $U \cap V$ we have $g_1 = f_1 \cdot (dz/dw)$ and $g_2 = f_2 \cdot (dz/dw)$, so $f_1/f_2 = g_1/g_2$. Hence, the function $f = f_1/f_2$ is well-defined and $\omega_1 = f\omega_2$. □

Proposition 4.4. *The sum of the all residues for a meromorphic 1-form ω on any compact Riemann surface X is 0.*

Proof. Let $a_1, \ldots, a_n$ be set of all poles of the 1-form ω. Since this set is discrete we can surround every point a_i with a small disk D_i, which does not contain other poles. On the set $X' = X \backslash \cup \mathrm{int}(D_i)$, our 1-form is holomorphic and closed. Hence, by the Stokes Theorem we have:

$$\sum_i \int_{\partial D_i} \omega = \int_{\partial X'} \omega = \int_{X'} d\omega = 0.$$

□

Definition 4.5. The *principal part* at $z = a$ of a function: $f(z) = \sum_{k=-\infty}^{\infty} a_k(z-a)^k$ is the portion of the Laurent series consisting of terms with negative degree. That is, $\sum_{k=-\infty}^{-1} a_k(z-a)^k$ is the the principal part of f at a.

Proposition 4.6. Reverse Residue Theorem *If we have a set of points* $\{a_1, \ldots, a_n\}$ *on a compact Riemann surface and also a set of principal parts* $\{f_1, \ldots, f_n\}$*, then the following are equivalent:*

(1) There exists a meromorphic function f *that has principal part* f_i *at each point* a_i *and has no other poles.*
(2) $\sum_{i=1}^{n} Res_{a_i} f_i\omega = 0$ *for all holomorphic 1-forms* ω *on our Riemann surface.*

Proof. We will first prove the forward implication. If such a meromorphic function f exists, then for any holomorphic 1-form ω, we have $\sum_{i=1}^{n} \mathrm{Res}_{a_i} f\omega = 0$ by 4.4. Now, f only has poles in points $\{a_1, \ldots, a_n\}$, and the principal part of f at a_i is equal to f_i. Therefore the residue of f at the point a_i is the coefficient of the z_{-1} term of the Laurent series f_i. Taking into account that ω is a holomorphic 1-form, and so multiplying by ω can only increase degrees but never decrease them, we have $\mathrm{Res}_{a_i} f\omega = \mathrm{Res}_{a_i} f_i\omega$ and $\sum_{i=1}^{n} \mathrm{Res}_{a_i} f_i\omega = 0$.

As to the reverse implication – follows from the Serre Duality (cf. [4], pg. 188). □

We will need to make use of the following theorem, whose proof can be found at [10].

Theorem 4.1. If the genus of Riemann surface X is equal to g, there exist on X exactly g independent holomorphic 1-forms.

5. Divisors

Definition 5.1. A *divisor* is an element of a free abelian group generated by the points of Riemann surface. Or, more simply stated, a divisor is a linear combination of finite number of points with integer coefficients.

Since divisors are elements of an abelian group, they can be added, subtracted, and multipled by integers. If $D_1 = \sum_a n_a a$ and $D_2 = \sum_a m_a a$ are divisors, then $D_1 + D_2 = \sum_a (n_a + m_a)a$. We can also compare divisors: $D_1 \leq D_2$ if and only if $n_a \leq m_a$ for all a.

Definition 5.2. If $D = \sum_a n_a a$ is a divisor, then the *degree* of D is the sum $\deg(D) = \sum_a n_a$. (Do not confuse the degree of a divisor with the degree of a covering map!) Because $\deg(D_1 + D_2) = \deg(D_1) + \deg(D_2)$, we see that *deg* is a homomorphism from the abelian group of divisors to the abelian group $\mathbb{Z}$.

Definition 5.3. If f is a meromorphic function on a (not necessarily compact) Riemann surface X, and a is a point of X, we define the order of f at a as

$$ord_a(f) = \begin{cases} k, & f \text{ has a zero of multiplicity } k \text{ at } a \\ -k, & f \text{ has a pole of multiplicity } k \text{ at } a \end{cases}$$

For every meromorphic function f on a compact Riemann surface, there corresponds a divisor $\sum_a ord_a(f)a$ over the (discrete and hence finite) set of points where the function f has zeroes and poles.

Definition 5.4. The divisor $\sum_a ord_a(f)a$ is denoted as (f) and called the *principal divisor* for f.

It is obvious that $(f \cdot g) = (f) + (g)$ and $(f/g) = (f) - (g)$.

Lemma 5.1. If f is a meromorphic function on a compact Riemann surface, then $\deg((f)) = 0$.

Proof. As we proved earlier in the paper, any meromorphic function on a compact Riemann surface has as many poles (counted with multiplicities) as it has zeroes (counted with multiplicities). Thus, by definition of *ord* and *deg*, $\deg((f)) = 0$. □

Definition 5.5. If ω is a meromorphic 1-form on a (not necessarily compact) Riemann surface X, and a is a point of X, we define the order of ω at a to be $ord_a(\omega) = ord_a(f_a)$, where f_a is a local representation of ω at the point a, i.e. ω is represented locally as $f_a dz$ in a neighborhood of a.

Definition 5.6. The divisor $\sum_a ord_a(\omega)a$ is denoted as (ω) and called the *canonical divisor* of the meromorphic 1-form ω.

It is obvious from the given definitions that if f is a meromorphic function and ω is a meromorphic 1-form, then $(f \cdot \omega) = (f) + (\omega)$. Due to additivity of *deg* and the previous lemma, we see that $\deg((f\omega)) = \deg((f)) + \deg((\omega)) = \deg((\omega))$.

Since, any meromorphic 1-form can be presented as any other meromorphic 1-form multiplied by some meromorphic function, we see that all canonical divisors have the same degree.

Lemma 5.2. If ω is a meromorphic 1-form on a compact Riemann surface X, then $\deg((\omega)) = -\chi$, where χ is the Euler characteristic of the Riemann surface.

Proof. Due to the remark above, it is enough to prove the lemma for *any* 1-form. For our 1-form we choose df, where f is some non-constant meromorphic function on our Riemann surface, which exists due to Riemann's Existence Theorem (2.1). We will consider f as a holomorphic map from X to $\overline{\mathbb{C}}$.

We may assume that there are no ramification points among pre-images of $\infty \in \overline{\mathbb{C}}$ – otherwise, we can "rotate" the Riemann sphere (using some Moebius transformation) and consider the combination of f with this "rotation" and its inverse.

Let us count $\deg(df)$. We will denote the covering degree of f by n, and ramification indexes of f by $e_1, \dots, e_m$.

Our function f has exactly n simple poles (pre-images of ∞), and at each of them the 1-form df has a pole of order 2 (because, $d(1/z) = -1/z^2$). On the other hand, zeroes of df are ramification points of the map f, and if the ramification index at some point is equal to e, then df has a zero of order $e-1$ at this point. (Explanation: at this point the Taylor series for f starts with $c_e z^e$, so, instead of e pre-images, f has just one at this point; thus the local representation of df at the point will start with $ec_e z^{e-1}$, i.e., it has a zero of order $(e-1)$.)

So, df has n poles of order 2 and m zeroes of orders $e_1 - 1, \dots, e_m - 1$. Summing it all up, we get:

$$\deg(df) = \sum_{i=1}^{m} (e_i - 1) - 2n,$$

and this number is equal to $-\chi = 2g - 2$ by the Riemann-Hurwitz formula. □

Definition 5.7. If D is a divisor, we define $L(D)$ to be the set of all meromorphic functions f for which $(f) + D \geq 0$. Obviously, $L(D)$ is a vector space over $\mathbb{C}$, and we define $l(D)$ to be the dimension of this vector space. Similarly, $I(D)$ is defined as the set of all meromorphic 1-forms ω for which $\omega \geq D$. Obviously, $I(D)$ is a linear space over $\mathbb{C}$, and $i(D)$ is defined to be $dim_{\mathbb{C}}(I(D))$.

Corollary 5.3. If $\deg(D) < 0$, then $l(D) = 0$.

Proof. Since any meromorphic function satisfies $\deg((f)) = 0$, we see that $L(D) = \emptyset$ (there is no such function f such that $(f) + D \geq 0$, otherwise we would have $= \deg(\mathrm{D}) = \deg((\mathrm{f})) + \deg(\mathrm{D}) = \deg((f) + D) \geq 0$.) Thus, $l(D) = 0$. □

Corollary 5.4. If the divisor $D = 0$, then $l(D) = 1$.

Proof. Indeed, the space of meromorphic functions f such that $(f) \geq 0$ is just the space of all *holomorphic* functions (no poles), i.e. the space of constants, which is isomorphic to $\mathbb{C}$. □

Lemma 5.5. The linear space $I(D)$ is isomorphic to the space $L(K - D)$, where K is a canonical divisor. The linear space $I(K - D)$ is isomorphic to the space $L(D)$.

Proof. We shall prove that $I(K - D)$ is isomorphic to $L(D)$ first. We assume that $K = (\omega)$ for some meromorphic 1-form ω.

If f is a meromorphic function, then $(f) + D = (f\omega) - (\omega) + D = (f\omega) - (K - D)$. So, $f \in L(D) = \{g \mid (g) + D \geq 0\}$ if and only if $f\omega \in I(K - D) = \{\beta \mid (\beta) \geq K - D\}$. Thus, we have a map: $(L(D) \to I(K - D))$ given by $f \mapsto f\omega$. Since every meromorphic 1-form in $I(K - D)$ can be presented as $f\omega$ for some meromorphic function, we similarly have an inverse function $I(K - D) \to L(D)$. Since linearity of our map is obvious, we have an isomorphism between $I(K - D)$ and $L(D)$ for any divisor D.

Now, if we substitute $K - D$ instead of D, we will get isomorphism between $I(D)$ and $L(K - D)$. □

Lemma 5.6. The linear space $L(K)$, where K is a canonical divisor, is isomorphic to the linear space of all holomorphic 1-forms on the same Riemann surface.

From the previous lemma $L(K)$ is isomorphic to $I(K-K) = I(0) = \{\omega \mid (\omega) \geq 0\}$, which is obviously the space of all holomorphic 1-forms. From the previous lemma and the Theorem 4.1 it follows that $l(K) = dim(L(K)) = g$.

Definition 5.8. Two divisors D_1 and D_2 on a Riemann surface are called *linearly equivalent* if there exists a meromorphic function f defined on the same surface such that D_1 and D_2 differ by its principal divisor (f), i.e., if $D_2 = D_1 + (f)$.

Lemma 5.7. Linearly equivalent divisors have equal degrees.

Proof. If $D_2 = D_1 + (f)$, then $\deg(D_2) = \deg(D_1 + (f)) = \deg(D_1) + \deg((f)) = \deg(D_1)$, since $\deg((f)) = 0$. □

Lemma 5.8. If divisors D_1 and D_2 are linearly equivalent, the space $L(D_1)$ is isomorphic to $L(D_2)$, and the space $I(D_1)$ is isomorphic to $I(D_2)$.

Proof. Let us assume that $D_2 = D_1 + (f)$ for a certain meromorphic function f. If g is a meromorphic function and $g \in L(D_2)$, then $(fg) + D_1 = (f) + (g) + D_1 = (g) + D_1 + (f) = (g) + D_2$. It follows that $(g) + D_2 \geq 0$ if and only if $(fg) + D_1 \geq 0$. Therefore, we have a map $L(D_2) \to L(D_1)$. Since, f is a meromorphic function, its inverse f^{-1} also is a meromorphic function, which clearly provides us with the inverse map $L(D_2) \mapsto L(D_1)$. Since linearity is also obvious we have an isomorphism between $L(D_1)$ and $L(D_2)$. The fact that $I(D_1)$ is isomorphic to $I(D_2)$ is proved similarly.

□

6. Riemann–Roch Theorem

Theorem 6.1. *(Riemann–Roch Theorem)* If D is a divisor on a compact Riemann surface of genus g, then $l(D) - i(D) = \deg(D) + 1 - g$.

There are several obvious implications of Riemann–Roch Theorem. If we assume that it is valid (without proof), we actually can prove the theorem that there exist exactly g holomorphic 1-forms on the Riemann surface X with genus g.

Indeed, if we take $D = 0$ and substitute it in the Riemann–Roch Theorem, we will get:

$$l(0) - i(0) = \deg(0) + 1 - g.$$

Now, $l(0) = 1$ as we have established earlier, and $\deg(0)$ is simply 0, as it is the sum of coefficients of the zero divisor.

$i(0) = l(K)$, where K is a canonical divisor. As we know $L(K)$ is isomorphic to the space $I(0) = \{\omega \mid (\omega) \geq 0\}$, which is the space of all holomorphic 1-forms. Thus, the formula gives us:

$$1 - i(0) = 0 + 1 - g,$$

and we conclude that $i(0)$ is equal to the dimension of the space $I(0) = \{\omega \mid (\omega) \geq 0\}$ of all holomorphic 1-forms.

Furthermore, if we substitute $D = K$ into the Riemann–Roch formula, we will get:

$$l(K) - i(K) = \deg(K) + 1 - g.$$

Since $l(K) = i(0) = g$, and $i(K) = l(0) = 1$ as mentioned above, we see that

$$g - 1 = \deg(K) + 1 - g.$$

Therefore we obtain the same formula $\deg(K) = 2g - 2$ that was proven earlier.

Proof. The proof is done in three parts:

First we will prove the statement in the case where $\deg(D) \geq 0$. Consider a divisor

$$D = \sum_{i=1}^{n} m_i a_i$$

with $D \geq 0$.

Let V be the set of all tuples $\{f_1, \ldots, f_n\}$ of principal parts of the form

$$f_i = \frac{c_{m_i}}{z^{m_i}} + \cdots + \frac{c_{-1}}{z}.$$

Obviously, V is a linear space over $\mathbb{C}$ of dimension $\deg(D)$.

(Example: Say D $= 3a_1 + 2a_2$. Then $V = \left\{ \left(\frac{c_{-3}}{z^3} + \frac{c_{-2}}{z^2} + \frac{c_{-1}}{z}, \frac{d_{-2}}{z^2} + \frac{d_{-1}}{z} \right) \right\}$, and $dim_C(V) = 5 = \deg(D)$.)

Let us create a map $\Phi : L(D) \to V$ that sends $f \in L(D)$ to the tuple of principal parts of f at the points a_i.

Also, let us consider the kernel of Φ, which is the set of functions in $L(D)$ that are sent to 0. Since $D \geq 0$, a function f with $(f) \geq -D$ that is in the kernel has no principle parts at the points a_i and can have no other other poles. This implies that such an f is holomorphic and, therefore, constant. Thus $dim(Ker(\Phi))$ must equal 1, since only constant functions are in the kernel.

Now, if we let $Im(\Phi) = W$, then we have

$$l(D) = dim(Ker(\Phi)) + dim(Im(\Phi)) = 1 + dim(W)$$

.

Next, we need to find out what $dim(W)$ is: W is a set of $\{f_i\}$ principal parts, such that exist $f \in L(D)$ with set of tails equal to $\{f_i\}$. According to 4.6, such f exists if and only if for all holomorphic 1-forms ω on our Riemann surface $\sum_{i=1}^{n} \operatorname{Res}_{a_i} f_i \omega = 0$. For each holomorphic 1-form ω we will consider the linear map $\lambda_\omega : V \to \mathbb{C}$ defined by $\{f_1, \ldots, f_n\} \mapsto \sum_{i=1}^{n} Res_{a_i} f_i \omega$.

Now, $W = \bigcap Ker(\lambda_\omega)$ is the intersection of the kernels of λ_ω for all holomorphic 1-forms ω. It follows (due to the standard theorem from Linear Algebra), $dim(W) = dim(\bigcap Ker(\lambda_\omega)) = dim(V) - dim(\{\lambda_\omega\})$, where $\{\lambda_\omega\}$ is the linear space generated by all λ_ω.

We know that $dim(V) = \deg(D)$, so $dim(W) = \deg(D) - dim(\{\lambda_\omega\})$ and $l(D) = 1 + dim(W) = 1 + \deg(D) - dim(\{\lambda_\omega\})$.

Next, we need to find $dim(\{\lambda_\omega\})$. The dimension of $\{\lambda_\omega\}$ is less than or equal to the genus g, since we know that the number of independent holomorphic 1-forms ω is g. Thus, one needs to consider the 1-forms that turn all principal parts, the elements of our space V, to zero, because these will exactly correspond to maps λ_ω that do not contribute to the space generated by $\{\lambda_\omega\}$.

To make the principal part $\frac{c_{m_i}}{z^{m_i}} + \cdots$ at a_i to turn to zero we need to multiply it by an ω which has $ord_{a_i}(\omega) \geq m_i$. This is true if and only if $(\omega) \geq D$, i.e., if and only if ω is in $I(D)$. Thus, $\lambda_\omega = 0$ if and only if ω is in $I(D)$, and so $dim(\{\lambda_\omega\}) = g - i(D)$. From which follows, $l(D) = 1 + \deg(D) - g + i(D)$, as desired.

Next, we will show that for all D, $l(D) - i(D) \geq 1 + \deg(D) - g$. Let $a \in X$ be a point on our surface. Obviously, $\deg(D-a) = \deg(D)-1$, so, if the inequality above is true for some divisor D, it easy to see that $l(D-a)-i(D-a) \geq 1+\deg(D-a)-g = \deg(D) + 1 - g - 1 = l(D) - i(D) - 1$, and the inequality is true for the divisor $D - a$ too.

The idea is to take some divisor ≥ 0 and subtract from it point-by-point. In fact, taking into account the equality established in *Part I*, in order to establish our inequality, we just need to show that $l(D - a) - i(D - a) \geq (l(D) - i(D)) - 1$.

It is obvious that $l(D) \geq l(D - a) \geq l(D) - 1$ and $i(D) + 1 \geq i(D - a) \geq i(D)$ (since we only add/remove at most one coefficient/member at a single point - a), so the worst case scenario occurs when $l(D-a) = l(D)-1$ and $i(D)+1 = i(D-a)$, because only in this case $l(D-a)-i(D-a) = (l(D)-i(D))-2$ (in all remaing cases either $l(D-a)-i(D-a) = (l(D)-i(D))-1$ or $l(D-a)-i(D-a) = (l(D)-i(D))$).

We will show that this scenario is impossible: let us take f in $L(D) \setminus L(D-a)$ and ω in $I(D-a) \setminus I(D)$. Such f and ω exist by our assumption that $l(D-a) = l(D)-1$ and $i(D)+1 = i(D-a)$. Then, $(f) \geq -D$, $(f) < a-D$, $(\omega) \geq D-a$, and $(\omega) < D$. Say, $D = na + \ldots$. This implies that $-n+1 > ord_a(f) \geq -n$, which means that $ord_a(f) = -n$.

Similarly, $n > ord_a(\omega) \geq n-1$, which implies $ord_a(\omega) = n-1$. Thus, $ord_a(f\omega) = -n + n - 1 = -1$. In addition, for all $b \neq a, ord_b(f\omega) \geq 0$, because $(f) \geq -D$, and $(f\omega) \geq -a$, which implies that the poles may be in a only. Indeed, $ord_a(f\omega) = -1$.

But now we have both$\sum_{i=1}^{n} \text{Res}_{a_i} f\omega = 0$ and $\text{Res}_a(f\omega) = c_{-1} \neq 0$ (for some c_{-1}, which cannot be equal to 0 because we have a pole of exactly first order), with a being the only point where $(f\omega)$ may have (and has) non-zero residue; but this is a contradiction and implies that the worst case scenario never happens.

Thus, $l(D - a) - i(D - a) \geq (l(D) - i(D)) - 1$. This implies that for all divisors D we have

$$l(D) - i(D) \geq deg(D) + 1 - g.$$

Finally, for the last part of the proof we use a clever trick to obtain the desired equality for an arbitrary divisor D. We substitute $K - D$ for D in the inequality proved in Part II. Since $l(K - D) = i(D)$ (see 5.5).

$$\begin{aligned} l(K-D) - i(K-D) &\geq deg(K-D) + 1 - g \\ i(D) - l(D) &\geq deg(K-D) + 1 - g \\ i(D) - l(D) &\geq degK - deg(D) + 1 - g \end{aligned}$$

Since $degK = deg((\omega)) = -\chi = 2 - 2g$, we then have

$$\begin{aligned} i(D) - l(D) \geq & 2g - 2 - deg(D) + 1 - g = g - 1 - deg(D) \\ l(D) - i(D) \leq & deg(D) + 1 - g \\ l(D) - i(D) = & deg(D) + 1 - g. \end{aligned}$$

Combining this with the inequality from Part II proves the theorem for an arbitrary divisor D.

□

Acknowledgments. It is a pleasure to thank my mentors Spencer Dowdall and Ilya Gekhtman for all their assistance with this paper. I would also like to thank Andrew Ryan and Aleksandr Likhterman for their invaluable assistance in reviewing my work.

References

[1] Corrin Clarkson, *Riemann Surfaces, REU project*, available at `http://www.math.uchicago.edu/~may/VIGRE/VIGRE2007/REUPapers/FINALFULL/Clarkson.pdf`

[2] S. M. Lvovski, *Lectures on Complex Analysis (in Russian)*, available at `http://www.mccme.ru/free-books/lvo/can.html`

[3] William S. Massey, *Algebraic Topology: An Introduction*, New York, NY: Springer-Verlag, 1977.

[4] Rick Miranda, *Algebraic Curves and Riemann Surfaces (Graduate Studies in Mathematics, Vol 5)*, American Mathematical Society, Providence, RI, 1995.

[5] Francis Kirwan, *Complex Algebraic Curves*, LMS Student Texts 23, Cambridge 1992.

[6] V. V. Prasolov, and O. V. Schwarzman, *The Alphabet of Riemann Surfaces*, Fazis Publ. House, Moscow, 1999.

[7] M. Schlichenmaier, *An introduction to Riemann surfaces, algebraic curves and moduli spaces*, Lecture Notes in Phys., 1989.

[8] Jürgen Jost, *Compact Riemann Surfaces*, Berlin, New York: Springer-Verlag, 2006.

[9] William Fulton, *Algebraic Curves. Mathematics Lecture Note Series*, W.A. Benjamin, 1974, available at `http://www.math.lsa.umich.edu/~wfulton/CurveBook.pdf`

[10] George Springer, *Introduction To Riemann Surfaces*, Addison-Wesley Publishing, 1957.