Riemann – Roch and Projectivity of Compact Riemann Surfaces

Author. Jas Singh
Acknowledgements. I'd like to thank my friend Anish Chedalavada (at UIC) for his much appreciated aid in the editing process.

Abstract

In this report we will present the basic definitions and results in the study of compact Riemann surfaces and describe a proof of the Riemann – Roch theorem, as well as discuss its application to embedding compact Riemann surfaces into projective space. The proof we present will be a mostly immediate application of sheaf cohomology, and we give an abbreviated account of the definitions and theorems therein. Our primary reference for the topics covered is [Forster]. Many of the proofs presented here can be found in greater detail there. Another good reference on Riemann surfaces is [Miranda].

0 What is a Riemann surface?

0.1 Definition

Riemann surfaces are topological spaces locally modeled on open subsets of $\mathbb{C}$. As we will talk about embeddings into projective space, it will be convenient to define general complex manifolds before specializing to Riemann surfaces. For a good introduction to general complex manifolds, we point the reader to [Huybrechts, Ch. 2].

Definition 0.1. Let U be an open subset of $\mathbb{C}^n$. Give it coordinates $z_1, \ldots, z_n$ and write $z_j = x_j + iy_j$. Define the differential operator $\frac{\partial}{\partial \bar{z}_j} = \frac{1}{2}\left(\frac{\partial}{\partial x_j} + i\frac{\partial}{\partial y_j}\right)$. A smooth function $f : U \longrightarrow \mathbb{C}$ is called holomorphic if it satisfies the Cauchy – Riemann equation $\frac{\partial f}{\partial \bar{z}_j} = 0$ for all j. A map $f : U \longrightarrow \mathbb{C}^m$ with coordinates $f = (f_1, \ldots, f_m)$ is said to be holomorphic if each coordinate f_j is holomorphic.

Definition 0.2. Let X be a (nonempty) topological space. An n – dimensional holomorphic atlas $\mathcal{A}$ on X is a set of pairs (U, ϕ) with $U \subseteq X$ open and $\phi : U \longrightarrow \mathbb{C}^n$ a homeomorphism onto an open subset of $\mathbb{C}^n$. We call U a coordinate patch and ϕ a coordinate chart. The space X must be covered by coordinate patches, i.e. $\bigcup_{(U,\phi)\in\mathcal{A}} U = X$. Additionally, the charts must be holomorphically compatible in the following sense. For $(U, \phi), (V, \psi) \in \mathcal{A}$ such that $U \cap V$ is nonempty, we must have that the composition $\phi \circ \psi^{-1}$ is holomorphic in the sense defined above, where we suppress writing explicit restrictions.

Definition 0.3. A complex manifold of dimension n is a Hausdorff space equipped with an equivalence class of n – dimensional holomorphic atlases, where $\mathcal{A}$ and $\mathcal{B}$ are said to be equivalent if $\mathcal{A} \cup \mathcal{B}$ is also an atlas.

Definition 0.4. A map $f : X \longrightarrow Y$ between complex manifolds is holomorphic if $\psi \circ f \circ \phi^{-1}$ is holomorphic in the sense of complex variables for all charts ϕ on X and ψ on Y, where restrictions are again suppressed. A holomorphic map with holomorphic inverse is called a biholomorphism.

0.2 Examples

Example. Let $U \subseteq \mathbb{C}$ be open and nonempty with the holomorphic atlas $\{(U, \iota)\}$, where $\iota : U \longrightarrow \mathbb{C}$ is the inclusion map.

One can check that if U and V are open subsets of $\mathbb{C}$ with the above holomorphic atlases, then a map $f : U \longrightarrow V$ is holomorphic in the sense of complex analysis if and only if it is holomorphic in the sense of complex manifolds.

Our next example will be complex projective space. As we will be interested later in embedding compact Riemann surfaces into projective space, we'll need the definition in higher dimensions than just 1.

Example. Let $n \geq 1$ and consider the action of the unit group $\mathbb{C}^\times = \mathbb{C} - \{0\}$ on the complex vector space $\mathbb{C}^{n+1}$ via scalar multiplication. We define $\mathbb{P}^n$ as the quotient space $(\mathbb{C}^{n+1} - \{0\})/\mathbb{C}^\times$. Its points represent complex lines through the origin in $\mathbb{C}^{n+1}$. This is a topological space which we will give the structure of a complex manifold of dimension n.

We write the coordinates of $\mathbb{C}^{n+1}$ as $z_0, \dots, z_n$. Elements of $\mathbb{P}^n$ will be written as $[z_0 : \cdots : z_n]$, which we refer to as homogeneous coordiantes. Let $U_i = \{[z_0 : \cdots : z_n] \in \mathbb{P}^n : z_i \neq 0\}$. These form an open cover of $\mathbb{P}^n$. We can define charts on U_i by observing that any point in U_i can be written uniquely as $[z_0 : \cdots : 1 : \cdots : z_n]$, with 1 in the i^{th} coordinate. Then our chart $\phi_i : U_i \longrightarrow \mathbb{C}^n$ will take such a point to $(z_0, \dots, z_{i-1}, z_{i+1}, \dots, z_n)$.

It is important to note that $\mathbb{P}^n$ is compact. Indeed, we can view view $\mathbb{P}^n$ alternatively as S^{2n+1}/S^1, where we view these spheres as contained in $\mathbb{C}^{n+1}$ and $\mathbb{C}$ respectively.

Then $\mathbb{P}^1$ is a Riemann surface covered by two patches U_0 and U_1, both of which are biholomorphic to $\mathbb{C}$. It is worth pointing out that the transition map $\phi_1 \circ \phi_0^{-1} : \mathbb{C} - \{0\} \longrightarrow \mathbb{C} - \{0\}$ is given by $z \mapsto 1/z$. Let's note that U_0 consists of elements of the form $[z : 1]$ for $z \in \mathbb{C}$. We can see then that the complement of U_0 in $\mathbb{P}^1$ consists of the single point $[1 : 0]$. Hence, $\mathbb{P}^1$ is a one point compactification of the complex plane, and is therefore homeomorphic to the sphere S^2. It is therefore common to refer to $\mathbb{P}^1$ as the Riemann sphere. We will also often write $\mathbb{P}^1$ as $\mathbb{C} \cup \{\infty\}$, where $\mathbb{C}$ is identified with U_0 and $\infty = [1 : 0]$.

Example. Let v_1, v_2 form a basis of $\mathbb{C}$ as a real vector space. We let $\Lambda = \mathbb{Z}v_1 + \mathbb{Z}v_2$, which is a free abelian group generated by v_1, v_2. Λ is called a lattice in $\mathbb{C}$. The quotient group $\mathbb{C}/\Lambda$ can be given the structure of a Riemann surface with a holomorphic projection $\pi : \mathbb{C} \longrightarrow \mathbb{C}/\Lambda$. For some open subset $U \subseteq \mathbb{C}$ which is disjoint from all of its translates $\lambda + U$ for $\lambda \in \Lambda$, it follows that $\pi|_U : U \longrightarrow \pi[U]$ is a homeomorphism. We can define a chart on $\pi[U]$ via $\pi|_U^{-1}$. One can then check that this yields a holomorphic atlas on $\mathbb{C}/\Lambda$, endowing it with the structure of a Riemann surface. We call $\mathbb{C}/\Lambda$ a complex torus.

Complex torii are also compact. Indeed, consider $K = \{t_1 v_1 + t_2 v_2 : t_i \in [0, 1]\}$. Then the restriction of π to K maps onto $\mathbb{C}/\Lambda$ and K is compact.

Let's also note that all complex torii are diffeomorphic, but not necessarily biholomorphic. Indeed, one can show that $\mathbb{C}/\Lambda_1$ and $\mathbb{C}/\Lambda_2$ are biholomorphic if and only if $\Lambda_1 = a\Lambda_2$ for some $a \in \mathbb{C}$. This reflects the relative rigidity of complex geometry as opposed to real geometry.

0.3 First properties of Riemann surfaces

By passing to coordinate patches, it is roughly true that any fact about complex analysis will become a local fact about Riemann surfaces.

Proposition 0.1 (The identity theorem for Riemann surfaces). *Let X and Y be Riemann surfaces and let f and g be holomorphic functions $X \longrightarrow Y$. Suppose that X is connected. If the set $\{z \in X : f(z) = g(z)\}$ has a nonempty interior, then $f = g$.*

Proof. f and g will agree on a coordinate patch U by the identity theorem from complex analysis. Applying the identity theorem again will allow this equality to extend to any coordinate patch V which intersects U. A connectedness argument will extend this equality to all of X, so that $f = g$. □

Corollary 0.0.1 (The maximum principle for Riemann surfaces). *Let X be a compact, connected Riemann surface and let $f : X \longrightarrow \mathbb{C}$ be a holomorphic function. Then f is constant.*

Proof. By compactness of X, there is some $z_0 \in X$ such that $|f(z_0)|$ is maximized. Using the complex analytic maximum principle in some chart about z_0, we have that f is constant on a nonempty open set. By the identity theorem, it follows that f is constant. □

1 Meromorphic functions and divisors

It is often fruitful to understand the geometry of a space in terms of coordinate functions defined on it. In complex geometry, we expect our coordinate functions to be maps into $\mathbb{C}$. However, we saw in the maximum principle from the previous section that the only globally defined holomorphic functions to $\mathbb{C}$ on a compact Riemann surface are constant. Hence, we have no choice but to broaden our notion of a coordinate function. We will do this with meromorphic functions.

1.1 Meromorphic functions

Definition 1.1. Let X be a Riemann surface. A meromorphic function on X is a holomorphic map $X - S \longrightarrow \mathbb{C}$ where S is discrete and closed. Furthermore, we want the elements of S to represent poles of f, rather than potentially being essential singularities. Hence, we also require that $\lim_{z\to s} |f(z)| = \infty$ for $s \in S$. We denote the set of meromorphic functions by $\mathcal{M}(X)$.

As discussed in section 0, the projective line $\mathbb{P}^1$ can be thought of as $\mathbb{C} \cup \{\infty\}$. By including ∞ in our codomain, we can treat the singular points as the same as any other point.

Proposition 1.1. *Let X be a Riemann surface and $f : X - S \longrightarrow \mathbb{C}$ be a meromorphic function. Then f can be extended to a holomorphic function $F : X \longrightarrow \mathbb{P}^1$ via*

$$F(z) = \begin{cases} f(z) & z \notin S \\ \infty & z \in S \end{cases}$$

Conversely, any holomorphic function $F : X \longrightarrow \mathbb{P}^1$ which is not the constant function at ∞ yields a meromorphic function away from the singular set $S = F^{-1}[\infty]$.

Definition 1.2. Let $f \in \mathcal{M}(X)$ be a meromorphic function and let $P \in X$. Take a coordinate patch (U, ϕ) about P. Then $f \circ \phi^{-1}$ is a meromorphic function on $\phi[U]$, so we let $\mathrm{ord}_P(f)$ be the order of $f \circ \phi^{-1}$ at $\phi(P)$. This will be an element of $\mathbb{Z} \cup \{\infty\}$.

Of course, we must check that this definition is well defined, as a different coordinate chart could a priori yield a different order. However, this does not occur and is a routine check.

From now on we will mostly think of meromorphic functions as maps to $\mathbb{P}^1$ which are not identically ∞. Allowing our functions to attain ∞ is a significant boon to the flexibility of our coordinate functions. For instance, $\mathbb{P}^1$ itself has many nonconstant meromorphic functions on it, such as the identity map $\mathbb{P}^1 \longrightarrow \mathbb{P}^1$, and more generally, rational functions extended to $\mathbb{P}^1$. Complex torii also have interesting meromorphic functions on them through the theory of elliptic functions, which we will not explore here. The Riemann – Roch theorem will provide us with a remarkably fine control in prescribing meromorphic functions to Riemann surfaces with given data on the poles and roots. The nature of this "fine control" is formalized with divisors.

Convention. From here on out, X will always refer to a compact, connected Riemann surface.

1.2 Divisors

Definition 1.3. A divisor D on X is a finite formal sum of points in X. We may write divisors as $D = \sum_{P\in X} D_P P$, with $D_P \in \mathbb{Z}$. The D_P may be thought of as a weight attached to the point P. The group of divisors is written as $\mathrm{Div}(X)$.

Additionally, there is a group homomorphism $\deg : \mathrm{Div}(X) \longrightarrow \mathbb{Z}$ via $\deg(D) = \sum_{P\in X} D_P$

Example. Let $f \in \mathcal{M}(X)$ be a meromorphic function which is not identically 0. We associate to f a divisor (f), which we call a principal divisor, defined as follows. This sum is finite due to compactness.

$$(f) = \sum_{P \in X} \mathrm{ord}_P(f) P$$

Definition 1.4. Let D and D' be divisors on X. We write $D \leq D'$ if for all points $P \in X$, $D_P \leq D'_P$. At this point, it will also be convenient to notate $(0) = \infty$, where 0 denotes the constantly 0 function. This is not a divisor, but we will extend our partial order so that $D \leq \infty$ for all divisors D. We furthermore say that $\infty + D = \infty$ for all D.

Say D is a divisor on X and f is a meromorphic function. Let's ponder the situation $(f) + D \geq 0$. Consider $D = 2P - Q$, for $P \neq Q$. Then saying $(f) + D \geq 0$ is equivalent to $\mathrm{ord}_P(f) \geq -2$ and $\mathrm{ord}_Q(f) \geq 1$. That is, f has a pole of order at least 2 at P and a root of order at least 1 at Q. In this context, we can therefore view the statement $(f) + D \geq 0$ as saying that f has poles at least as prescribed by D. When we make statements like this, a pole of negative order means a root and vice versa.

For a general divisor D, it is not at all obvious that a meromorphic function with poles at least as prescribed by D exists at all. And in fact, for many divisors, such a thing will be impossible. The Riemann–Roch theorem will afford us a deep understanding of when such a meromorphic function exists.

2 The cohomology black box

This section will contain the technical heart of the coming proof of Riemann – Roch. As suggested by the title, we will have to treat all the results in this section as a black box. The material presented here can be found in the special case of Riemann surfaces in [Forster, §12]. For the more general theory, the classic paper by Serre [FAC] is a wonderful reference for Cech cohomology. One can also read [Hartshorne, Ch. 3] for the derived functor approach using injectives. The previous two references are especially focused on the algebraic case. For an in depth and general approach, [Bredon] is a more than comprehensive reference for what we need here. Furthermore, cohomology is used all over the place in [Huybrechts] and one can read Appendix B therein for a quick summary of the statements of sheaf cohomology.

2.1 Sheaves

Let's begin with the particular "sheaves" we care about before delving into the general theory. Let D be a divisor on a compact connected Riemann surface X. For an open subset U of X, we let $\mathcal{O}_D(U)$ be the set of meromorphic functions on U with poles at least as prescibed by D. Formally we define,

$$\mathcal{O}_D(U) = \{f \in \mathcal{M}(U) : (f) + D|_U \geq 0\}$$

Here, $D|_U = \sum_{P \in U} D_P P$. It's clear that $\mathcal{O}_D(U)$ is a $\mathbb{C}$ vector space for all U.

For example, $\mathcal{O}_0(U)$ is the space of holomorphic functions $U \longrightarrow \mathbb{C}$. We write $\mathcal{O} = \mathcal{O}_0$. We can now express the maximum principle as saying that $\mathcal{O}(X) = \mathbb{C}$.

The Riemann – Roch theorem will provide us with a formula for the dimension of $\mathcal{O}_D(X)$, which we shall see will grow roughly linearly in $\deg(D)$. In particular, we will often be able to ensure the existence of nonconstant meromorphic functions with certain data on the poles as prescribed by a divisor.

Definition 2.1. A sheaf (of $\mathbb{C}$ vector spaces) $\mathcal{F}$ on a topological space A is a choice of $\mathbb{C}$ vector space $\mathcal{F}(U)$ for all open subsets $U \subseteq A$ along with "restriction" homomorphisms $\rho_U^V : \mathcal{F}(V) \longrightarrow \mathcal{F}(U)$ for any inclusion of open sets $U \subseteq V$. These maps must be compatible in the sense that $\rho_U^V \circ \rho_V^W = \rho_U^W$ for $U \subseteq V \subseteq W$. We often write $f|_U$ for $\rho_U^V(f)$. The above data is also subject to the following two axioms regarding an open cover $U = \bigcup_{i \in I} U_i$ of an open subset U of A.

(i) If $f \in \mathcal{F}(U)$ satisfies $f|_{U_i} = 0$ for all $i \in I$, then $f = 0$.

(ii) Let's say we have $f_i \in \mathcal{F}(U_i)$ for all i such that $f_i|_{U_i \cap U_j} = f_j|_{U_i \cap U_j}$ for all $i, j \in I$. Then there is an element $f \in \mathcal{F}(U)$ such that $f|_{U_i} = f_i$.

Examples. 1. For a divisor D and a Riemann surface X, $\mathcal{O}_D$ as defined above is a sheaf with the restriction maps being actual restriction of functions.

2. Let $P \in X$. We define the skyscraper sheaf $\mathbb{C}_P$ as

$$\mathbb{C}_P(U) = \begin{cases} \mathbb{C} & P \in U \\ 0 & P \notin U \end{cases}$$

with restriction maps id or 0 when appropriate.

3. We have the trivial sheaf $\underline{0}(U) = 0$ for all U.

Definition 2.2. Let $\mathcal{F}$, $\mathcal{G}$ be sheaves on a space A. A morphism $\alpha : \mathcal{F} \longrightarrow \mathcal{G}$ is a collection of morphisms $\alpha_U : \mathcal{F}(U) \longrightarrow \mathcal{G}(U)$ which are compatible with respect to restriction in the sense that the following diagram commutes whenever $U \subseteq V$.

$$\begin{tikzcd}
\mathcal{F}(V) \arrow[r, "\alpha_V"] \arrow[d, "\rho_U^V"'] & \mathcal{G}(V) \arrow[d, "\rho_U^V"] \\
\mathcal{F}(U) \arrow[r, "\alpha_U"'] & \mathcal{G}(U)
\end{tikzcd}$$

2.2 Cohomology

Definition 2.3. Let V, W, U be $\mathbb{C}$ vector spaces with maps $f : V \longrightarrow W$ and $g : W \longrightarrow U$. We say that the sequence $V \xrightarrow{f} W \xrightarrow{g} U$ is exact if $\ker(g) = \operatorname{im}(f)$. Given a longer sequence like the following

$$\dots \longrightarrow V_{i-1} \longrightarrow V_i \longrightarrow V_{i+1} \longrightarrow \dots$$

we say this is exact if every triple $V_{i-1} \longrightarrow V_i \longrightarrow V_{i+1}$ is exact per the above definition.

For example, $V \xrightarrow{f} W \longrightarrow 0$ is exact if and only if f is surjective, and dually $0 \longrightarrow V \xrightarrow{f} W$ is exact if and only if f is injective.

We will now extend this notion to sheaves. One approach is to define an appropriate notion of images and kernels of sheaf maps. Images of maps of sheaves are subtle, so we instead define exactness using stalks.

Definition 2.4. Let $\mathcal{F}$ be a sheaf on a space A and let $a \in A$. We define the stalk of $\mathcal{F}$ at a concisely as the colimit $\mathcal{F}_a = \operatorname{colim}_{a \in U \subseteq A \text{ open}} \mathcal{F}(U)$. Explicitly, that means that $\mathcal{F}_a = \coprod_{a \in U \subseteq A \text{ open}} \mathcal{F}(U)/\sim$ where $\coprod$ is the disjoint union and the equivalence relation $\sim$ is defined as follows. Let $f \in \mathcal{F}(U)$, $g \in \mathcal{F}(V)$. We say $f \sim g$ if there is some nonempty open subset $W \subseteq U \cap V$ such that $f|_W = g|_W$. This can be given the structure of a $\mathbb{C}$ vector space.

There is a map $\mathcal{F}(U) \longrightarrow \mathcal{F}_a$ for any neighborhood $a \in U$. We write this map as $f \mapsto f_a$, and refer to f_a as the germ of f at a. Elements of the stalk $\mathcal{F}_a$ are therefore represented by elements of $\mathcal{F}(U)$ for some neighborhood U, and two such representatives are equivalent if they agree near a. The stalk therefore captures the local data about a.

Now, let $\alpha : \mathcal{F} \longrightarrow \mathcal{G}$ and let $a \in A$. We get an induced map $\alpha_a : \mathcal{F}_a \longrightarrow \mathcal{G}_a$, which arises concisely by functoriality of the colimit. Explicitly, this map takes a germ $f_a \in \mathcal{F}_a$ to the germ $(\alpha_U(f))_a$ for some representative $f \in \mathcal{F}(U)$ of the germ f_a. One must of course check that this is a well defined homomorphism, but this is routine and affords us the definition of exactness of a sequence of sheaf morphisms.

Definition 2.5. Let $\mathcal{F} \xrightarrow{\alpha} \mathcal{G} \xrightarrow{\beta} \mathcal{H}$ be a sequence of morphisms of sheaves over some space A. We say that this sequence is exact if it is stalkwise exact, that is for all $a \in A$ the associated sequence on stalks $\mathcal{F}_a \xrightarrow{\alpha_a} \mathcal{G}_a \xrightarrow{\beta_a} \mathcal{H}_a$ is exact in the usual sense for $\mathbb{C}$ vector spaces we defined above.

With the notion of exactness out of the way, we present without proof the fundamental property of the sheaf cohomology groups.

Proposition 2.1 (The long exact sequence of cohomology)**.** *Let $\underline{0} \longrightarrow \mathcal{F} \xrightarrow{\alpha} \mathcal{G} \xrightarrow{\beta} \mathcal{H} \longrightarrow \underline{0}$ be an exact sequence of sheaves on a space A. We call this a short exact sequence. There exists a "natural" exact sequence of $\mathbb{C}$ vector spaces*

$$0 \longrightarrow H^0(A,\mathcal{F}) \xrightarrow{\alpha_*} H^0(A,\mathcal{G}) \xrightarrow{\beta_*} H^0(A,\mathcal{H}) \longrightarrow H^1(A,\mathcal{F}) \xrightarrow{\alpha_*} H^1(A,\mathcal{G}) \xrightarrow{\beta_*} H^1(A,\mathcal{H}) \longrightarrow \dots$$

with this three-periodic pattern continuing indefinitely. We furthermore know that $H^0(X,\mathcal{F}) = \mathcal{F}(X)$.

We do not present the computational tools of injective resolutions or Cech cohomology, and instead point the reader to the above references. The next two results can be proven using Cech cohomology, as seen in [Forster, 16.6] and [Forster, 14.10] respectively, though the latter uses quite a bit of analysis whereas the former is much easier.

Lemma 2.1. *$H^0(X,\mathbb{C}_P) = \mathbb{C}$ and $H^1(X,\mathbb{C}_P) = 0$.*

Lemma 2.2. *$H^1(X,\mathcal{O})$ is finite dimensional.*

We will now present the key example of the long exact sequence of cohomology used for Riemann – Roch.

Lemma 2.3 (The key short exact sequence)**.** *Let D be a divisor on X. Let $P \in X$ and $D' = D + P$. Then we have the following short exact sequence.*

$$\underline{0} \longrightarrow \mathcal{O}_D \xrightarrow{\iota} \mathcal{O}_{D'} \xrightarrow{\pi} \mathbb{C}_P \longrightarrow \underline{0}$$

The map $\iota : \mathcal{O}_D \longrightarrow \mathcal{O}_{D'}$ is inclusion. Indeed, as $D' \geq D$, for f have to have poles at least as prescribed by D it must have poles at least D'. Formally, we have that $(f) + D'|_U \geq (f) + D|_U$, whence $f \in \mathcal{O}_D(U)$ implies $f \in \mathcal{O}_{D'}(U)$.

The map $\pi : \mathcal{O}_{D'} \longrightarrow \mathbb{C}_P$ is more subtle. Let $U \subseteq X$ and $f \in \mathcal{O}_{D'}(U)$. If $p \notin U$ then we are forced to take $\pi(f) = 0$. If $p \in U$, take coordinates ω about P. Then in these coordinates, $f(\omega)$ (meaning $f \circ \omega^{-1}$) has the Laurent series expansion $f(\omega) = \sum a_n \omega^n$. Now, let $k = D_P$, i.e. the multiplicity of P appearing in D. Then as $f \in \mathcal{O}_{D'}(U)$ and $D' = D + P$ it follows that the order of the pole of f at P is at most $k+1$. Hence, the Laurent series expansion is trivial below $-(k+1)$. That is, it's of the form $\sum_{n \geq -(k+1)} a_n \omega^n$. We then take $\pi(f) = a_{-(k+1)}$. One must of course check well definition and the like, but this is routine.

As for exactness, the only interesting stalk to check is at P. And indeed, at P we can see that $\pi(f) = 0$ holds if and only if f has a Laurent expansion of the form $\sum_{n \geq -k} a_n \omega^n$, so that the pole at P has order at most $-k$. That is, so that $f \in \mathcal{O}_D(U)$, which means that $f \in \mathrm{im}(\iota)$.

Corollary 2.3.1 (The key long exact sequence of cohomology)**.** *We have the exact sequence*

$$0 \longrightarrow H^0(X,\mathcal{O}_D) \xrightarrow{\iota_*} H^0(X,\mathcal{O}_{D'}) \xrightarrow{\pi_*} \mathbb{C} \longrightarrow H^1(X,\mathcal{O}_D) \xrightarrow{\iota_*} H^1(X,\mathcal{O}_{D'}) \xrightarrow{\pi_*} 0$$

3 Proving Riemann – Roch

We now finally have the requisite technology to state Riemann – Roch. First, some notation. If $\mathcal{F}$ is a sheaf of $\mathbb{C}$ vector spaces over a space A we let $h^i(A, \mathcal{F})$ denote the dimension of $H^i(A, \mathcal{F})$.

Theorem 3.1 (Riemann – Roch)**.** *Let X be a compact connected Riemann surface and D a divisor on X. We then have the formula*

$$h^0(X, \mathcal{O}_D) - h^1(X, \mathcal{O}_D) = 1 - h^1(X, \mathcal{O}) + \deg(D)$$

The number $h^1(X, \mathcal{O})$ is referred to as the genus of X, and we often write this as $g = h^1(X, \mathcal{O})$.

The strategy to prove Riemann – Roch will be to "induct upwards and downwards" on the divisor D. That is, we let $D' = D + P$ for some point P and show that Riemann – Roch holds for D if and only if it holds for D'. The "if" direction is the downwards induction going from D' to D and the "only if" direction is the upwards induction going from D to D'. All divisors arise by adding or subtracting points from one another in some finite fashion, so if we show the above double inductive step, it will suffice to prove Riemann – Roch for just one fixed divisor. And indeed, the result is obvious for the divisor $D = 0$, as $\mathcal{O}_0 = \mathcal{O}$ and $h^0(X, \mathcal{O}) = 1$ by the maximum principle.

Before we do this, we should take note that this formula does not even make sense unless we insist on finite dimensionality of the relevant cohomology groups, as otherwise we cannot make sense of subtraction. The proofs of finite dimensionality will follow a similar "double induction" pattern to the proof of Riemann – Roch itself.

Proposition 3.1. *The cohomology groups $H^0(X, \mathcal{O}_D)$ are finite dimensional for all divisors D.*

Proof. Let's start with $D = 0$. Then $H^0(X, \mathcal{O}_0) = H^0(X, \mathcal{O}) = \mathcal{O}(X)$, which is $\mathbb{C}$ by the maximum principle. To piece together our double induction, we will use the key long exact sequence of cohomology above:

$$0 \longrightarrow H^0(X, \mathcal{O}_D) \xrightarrow{\iota_*} H^0(X, \mathcal{O}_{D'}) \xrightarrow{\pi_*} \mathbb{C} \longrightarrow H^1(X, \mathcal{O}_D) \xrightarrow{\iota_*} H^1(X, \mathcal{O}_{D'}) \xrightarrow{\pi_*} 0$$

We can break off a piece of this to get a short exact sequence

$$0 \longrightarrow H^0(X, \mathcal{O}_D) \longrightarrow H^0(X, \mathcal{O}_{D'}) \longrightarrow \operatorname{im}(\pi_*) \longrightarrow 0$$

By rank nullity, $H^0(X, \mathcal{O}_D)$ is finite dimensional if and only if $H^0(X, \mathcal{O}_{D'})$ is finite dimensional, as $\operatorname{im}(\pi_*)$ is finite dimensional. Writing down the rigorous "double induction" from here will be tedious, so instead we illustrate the method with an example. The point is to traverse from D to 0, our base case, using the equivalence we just showed. Take $D = P - Q$. Then $h^0(X, \mathcal{O}_D)$ is finite if and only if $h^0(X, \mathcal{O}_{D-P})$ is finite. $D - P = -Q$, which is now closer to 0. Applying this equivalence once more yields $h^0(X, \mathcal{O}_{-Q})$ is finite if and only if $h^0(X, \mathcal{O}_0)$ is finite. We showed above that $h^0(X, \mathcal{O}_0) = 1$, so indeed $h^0(X, \mathcal{O}_D)$ is finite. □

As for finite dimensionality of $H^1(X, \mathcal{O}_D)$, the method is exactly the same as above. The "double inductive" step will again arise by splitting the long exact sequence of cohomology, this time using the following:

$$0 \longrightarrow \operatorname{cok}(\pi_*) \longrightarrow H^1(X, \mathcal{O}_D) \longrightarrow H^1(X, \mathcal{O}_{D'}) \longrightarrow 0$$

Here, $\operatorname{cok}(\pi_*)$ is the cokernel $\mathbb{C}/\operatorname{im}(\pi_*)$. This is finite dimensional, so as $h^1(X, \mathcal{O})$ is finite, we can proceed as above.

Proof of Riemann – Roch. As discussed above, Riemann – Roch clearly holds for the divisor $D = 0$, so we need to prove a double inductive step. Let D be a divisor and $D' = D + P$. We claim Riemann – Roch holds for D if and only if it holds for D'.

Repeatedly applying rank – nullity to the key long exact sequence yields the following identity

$$h^0(X, \mathcal{O}_D) - h^0(X, \mathcal{O}_{D'}) + 1 - h^1(X, \mathcal{O}_D) + h^1(X, \mathcal{O}_{D'}) = 0$$

Rearranging this yields

$$h^0(X, \mathcal{O}_D) - h^1(X, \mathcal{O}_D) + 1 = h^0(X, \mathcal{O}_{D'}) - h^1(X, \mathcal{O}_{D'})$$

Furthermore, $1 = \deg(D') - \deg(D)$, so we conclude

$$h^0(X, \mathcal{O}_D) - h^1(X, \mathcal{O}_D) - \deg(D) = h^0(\mathcal{O}_{D'}) - h^1(X, \mathcal{O}_{D'}) - \deg(D')$$

Riemann – Roch holding for D' means that the right hand side of this equation equals $1 - h^1(X, \mathcal{O})$, and Riemann – Roch holding for D means that the left hand side equals $1 - h^1(X, \mathcal{O})$. Thus, we have shown that Riemann – Roch holds for D if and only if it holds for D'. This with the result for $D = 0$ affords us a double induction to prove Riemann – Roch for all divisors. □

4 So what?

Our goal now is to explain in what sense Riemann – Roch is a meaningful theorem. One sense in which the theorem is currently unsatisfying is that the H^1 are seemingly defined by fiat to allow the formula to hold. We will later explain more in depth what these terms mean, but even without this we get a nontrivial result. Let's first state what Riemann – Roch tells us if we shut our eyes at the H^1 terms.

Corollary 4.0.1 (Minimalist form of Riemann – Roch)**.** *There is a constant C depending only on X so that $h^0(X, \mathcal{O}_D) \geq C + \deg(D)$.*

This itself is a significant result. It tells us that the dimension of $H^0(X, \mathcal{O}_D) = \mathcal{O}_D(X)$ grows at least linearly in $\deg(D)$. In particular, for sufficiently large $\deg(D)$ we actually know that $\mathcal{O}_D(X)$ has some nonconstant elements. That is,

Corollary 4.0.2. *Every compact connected Riemann surface X admits a nonconstant map $X \longrightarrow \mathbb{P}^1$.*

This stands in contrast to holomorphic coordinate functions on X, i.e. maps $X \longrightarrow \mathbb{C}$, which are all constant. That tells us that there is no direct holomorphic analog of the Whitney embedding theorem, which allows smooth manifolds to be embeded in $\mathbb{R}^N$, for compact Riemann surfaces. But by virtue of Riemann – Roch, meromorphic coordinate functions (maps $X \longrightarrow \mathbb{P}^1$) abound. Hence, if we replace our affine coordinate functions $X \longrightarrow \mathbb{C}$ with projective coordinate functions $X \longrightarrow \mathbb{P}^1$, we are led to the following embedding theorem.

Theorem 4.1 (The projective embedding theorem)**.** *Let X be a connected compact Riemann surface. Then there is some holomorphic embedding $F : X \longrightarrow \mathbb{P}^N$ for some N. By an embedding, we mean that F is injective and that its derivative is also injective.*

However, we will need a better understanding of the $h^1(X, \mathcal{O}_D)$ term in Riemann – Roch to show this. Specifically, we will need to know that this term vanishes for $\deg(D)$ sufficiently large to precisely control the growth of $h^0(X, \mathcal{O}_D)$.

Before this, let's discuss how to define maps into $\mathbb{P}^N$. Say we have meromorphic functions $f_0, \ldots, f_N$ on X. Since we have been calling these coordinate functions, we better be able to piece them together to get a map $F : X \longrightarrow \mathbb{P}^N$ whose projective coordinates are the f_i, that is $F = [f_0 : \cdots : f_N]$. We must however be a bit careful. The homogeneous coordinates on a point in $\mathbb{P}^N$ are not allowed to all be 0, so we must insist that not all f_i are identically 0. If $x \in X$ is a point such that not all $f_i(x) = 0$ and so that x is not a pole of any f_i, then $(f_0(x), \ldots, f_N(x))$ is in $\mathbb{C}^{N+1} - \{0\}$. Then this defines a point $F(x) = [f_0(x) : \cdots : f_N(x)] \in \mathbb{P}^N$.

We also want to extend F to be defined at points x so that all $f_i(x) = 0$ or that some $f_i(x) = \infty$. This can be done by factoring out sufficient powers of the coordinate z from the local Laurent expansions of the f_i about some point where F is not defined. Doing so will not change the value of F, as homogeneous coordinates are invariant under scaling, and will then allow for an extension to any such point. A more detailed form of this argument is in [Miranda, Ch. 5, Lemma 4.2].

4.1 On H^1

We'll first discuss $H^1(X, \mathcal{O})$. We previously referred to $h^1(X, \mathcal{O})$ as the genus of X. Referring to it as such requires justification, as the term genus is a pre – existing topological term. Indeed, as the taxonomical name suggests, the genus forms a classification of surfaces. Compact connected orientable surfaces are classified up to diffeomorphism by the number of holes they possess, and the number of holes is called the genus. For instance, spheres have genus 0 and torii have genus 1. This is a purely topological notion, and incredibly the topological notion of the genus via holes and the complex analytic notion of the genus via $g = h^1(X, \mathcal{O})$ agree. One can see [Miranda, Ch. 6, §3]. This enhances Riemann – Roch significantly, as one of the terms is now purely topological. Indeed, we can rewrite Riemann – Roch now as saying

$$h^0(X, \mathcal{O}_D) - h^1(X, \mathcal{O}_D) = 1 - g + \deg(D)$$

Both sides of course depend on D, but the left hand side includes terms on the complex analytic structure of X, whereas the right hand side includes terms on the topology of X. This provides a deep connection between the topology and analysis of compact connected Riemann surfaces.

Now we discuss $H^1(X, \mathcal{O}_D)$. For any compact connected Riemann surface X there is a special divisor K_X called the canonical divisor. The canonical divisor is defined as the principal divisor associated to a meromorphic 1 form on X, but we will not discuss the notion of differential calculus on Riemann surfaces so we will simply assert its existence and state some properties. This is discussed in more detail in [Forster, §17].

Proposition 4.1. *(a)* $\deg(K_X) = 2g - 2$.

(b) There is a canonical isomorphism $H^1(X, \mathcal{O}_D)^* \cong H^0(X, \mathcal{O}_{K-D})$.

It turns out that for a nonconstant meromorphic function f on a compact Riemann surface X that $\deg((f)) = 0$. See [Miranda, Ch. 2, Prop. 4.12]. So if $\deg(D) < 0$, we have that $H^0(X, \mathcal{O}_D) = 0$.

Corollary 4.1.1. $H^1(X, \mathcal{O}_D) = 0$ *if* $\deg(D) > 2g - 2$.

Corollary 4.1.2. *For* $\deg(D) > 2g - 2$*, Riemann – Roch says*

$$h^0(X, \mathcal{O}_D) = 1 - g + \deg(D)$$

This gives us a precise linear growth on $h^0(X, \mathcal{O}_D)$ for large enough $\deg(D)$, which we can use in the following refinement of the projective embedding theorem.

4.2 Projective embedding

Theorem 4.2. *Let D be a divisor with* $\deg(D) > 2g$ *and let $f_0, \ldots, f_N$ be a basis for $H^0(X, \mathcal{O}_D)$. Let $F : X \longrightarrow \mathbb{P}^N$ have homogeneous coordinates $F = [f_0 : \ldots f_N]$. Then F is an embedding of X into $\mathbb{P}^N$.*

Proof. We prove injectivity of F. Let $x_1 \neq x_2$ be distinct points. Now let $D' = D - x_2$ and $D'' = D' - x_1$. Then $\deg(D'') = \deg(D) - 2 > 2g - 2$. Hence, we have

$$\begin{aligned} h^0(X, \mathcal{O}_{D'}) &= 1 - g + \deg(D') \\ &= 1 - g + \deg(D'') + 1 \\ &= h^0(X, \mathcal{O}_{D''}) + 1 \end{aligned}$$

So there is some $f \in H^0(X, \mathcal{O}_{D'}) - H^0(X, \mathcal{O}_{D''})$. Then f satisfies $(f) + D' \geq 0$ and $(f) + D'' \ngeq 0$. In particular, $\mathrm{ord}_{x_1}(f) = -D_{x_1}$ and $\mathrm{ord}_{x_2}(f) \geq -D_{x_2} + 1$.

Let $k_j = \min_i \mathrm{ord}_{x_j}(f_i)$ Suppose that it were the case that $\mathrm{ord}_x(f_i) > -D_x$ for all i. Then $f_0, \dots, f_N$ would all be in $H^0(X, \mathcal{O}_{D-x})$, which is a codimension 1 subspace of $H^0(X, \mathcal{O}_D)$ by the same argument as for D' and D''. But they form a basis for $H^0(X, \mathcal{O}_D)$ so this is a contradiction. Hence, $k_j = -D_{x_j}$.

We take coordinates ω_j with $\omega_j(x_j) = 0$ defined on open neighborhoods U_j, which we can shrink to be disjoint as X is Hausdorff. $H^0(X, \mathcal{O}_{D'}) \subseteq H^0(X, \mathcal{O}_D)$, so we write $f = \sum \lambda_i f_i$ in our above basis. We consider then the meromorphic functions $g_{ij}(\omega) = f_i(\omega_j)/\omega_j^{k_j}$ and $g_j(\omega) = f(\omega)/\omega_j^{k_j}$. The purpose of this is that $\mathrm{ord}_{x_j}(g_j) = \mathrm{ord}_{x_j}(f) - k_j = \mathrm{ord}_{x_j}(f) + D_{x_j}$. Hence, $\mathrm{ord}_{x_1}(g_1) = 0$ so $g_1(x_1) \neq 0$ and $\mathrm{ord}_{x_2}(g_2) \geq 1$ so $g_2(x_2) = 0$.

As homogeneous coordinates are invariant under scaling,

$$F(x_j) = [f_0(x_j) : \cdots : f_N(x_j)] = [g_{0j}(x_j) : \cdots : g_{Nj}(x_j)]$$

Furthermore, $\sum_i \lambda_i g_{ij}(x_j) = g_j(x_j)$. Hence, $\sum_i \lambda_i g_{i1}(x_1) \neq 0$ and $\sum_i \lambda_i g_{i2}(x_2) = 0$. Thus, the points $[g_{01}(x_1) : \cdots : g_{N1}(x_1)]$ and $[g_{02}(x_2) : \cdots : g_{N2}(x_2)]$ cannot be equal. Hence, $F(x_1) \neq F(x_2)$.

The proof that F is an immersion is a similar pattern. One needs only show that $F'(x_0) \neq 0$ for any $x_0 \in X$. This can be done by considering the divisor $D' = D - x_0$, the details of which we omit. This proof can be found in [Forster, Thm. 17.22] □

Riemann – Roch has afforded us enough meromorphic coordinate functions to treat all compact Riemann surfaces as embedded submanifolds of $\mathbb{P}^N$. In fact, compact connected Riemann surfaces are embedded subvarieties of $\mathbb{P}^N$. That is, they are defined by the zero set of some finite collection of homogeneous polynomials. One can see [Griffiths and Harris, Ch. 1, §3] for the general result along these lines called Chow's theorem. It turns out then that the complex analytic theory of compact connected Riemann surfaces is precisely the same as the algebro – geometric theory of nonsingular projective curves over $\mathbb{C}$. This allows the simultaeneous use of both analytic and algebraic methods when studying these objects.

5 Conclusion

The Riemann – Roch theorem is a foundational result in the theory of compact connected Riemann surfaces. We have presented here just a fraction of its implications and meaning, and we invite the reader to consult [Miranda] and [Forster] for further study. In short, we have discussed the interplay between topology and analysis that Riemann – Roch implies. Furthermore, we have utilized Riemann – Roch to prove the projectivity of compact Riemann surfaces. As discussed, it is even true that compact Riemann surfaces are projective varieties. If we look back at the proof of Riemann – Roch, the method was wholly a computation with sheaf cohomology. This, together with algebraicity of compact Riemann surfaces, suggests that Riemann – Roch can be generalized to algebraic geometry far beyond $\mathbb{C}$.

References

[FAC] Jean-Pierre Serre. "Faisceaux Algebriques Coherents". In: *The Annals of Mathematics* 61.2 (Mar. 1955), p. 197. ISSN: 0003486X. DOI: 10.2307/1969915. URL: https://www.jstor.org/stable/1969915?origin=crossref.

[Hartshorne] Robin Hartshorne. *Algebraic Geometry*. Vol. 52. Graduate Texts in Mathematics. New York, NY: Springer New York, 1977. ISBN: 9781441928078 9781475738490. DOI: 10.1007/978-1-4757-3849-0. URL: http://link.springer.com/10.1007/978-1-4757-3849-0.

[Forster] Otto Forster. *Lectures on Riemann Surfaces*. Vol. 81. Graduate Texts in Mathematics. New York, NY: Springer New York, 1981. ISBN: 9781461259633 9781461259619. DOI: 10.1007/978-1-4612-5961-9. URL: http://link.springer.com/10.1007/978-1-4612-5961-9.

[Griffiths and Harris] Phillip Griffiths and Joe Harris. *Principles of algebraic geometry*. eng. Wiley Classics Library edition. Pure and applied mathematics. New York: Wiley, 1994. ISBN: 9780471050599.

[Miranda] Rick Miranda. *Algebraic curves and Riemann surfaces*. Graduate studies in mathematics v. 5. Providence, R.I: American Mathematical Society, 1995. ISBN: 9780821802687.

[Bredon] Glen E. Bredon. *Sheaf Theory*. Vol. 170. Graduate Texts in Mathematics. New York, NY: Springer New York, 1997. ISBN: 9781461268543 9781461206477. DOI: 10.1007/978-1-4612-0647-7. URL: http://link.springer.com/10.1007/978-1-4612-0647-7.

[Huybrechts] Daniel Huybrechts. *Complex geometry: an introduction*. eng. Universitext. Berlin New York: Springer, 2005. ISBN: 9783540266877.