THE RIEMANN-ROCH THEOREM AND SERRE DUALITY

JOHN HALLIDAY

Abstract. We introduce sheaves and sheaf cohomology and use them to prove the Riemann-Roch theorem and Serre duality. The main proofs follow the treatment in Forster [3].

Contents

1. Introduction

The theory of Riemann surfaces lies at the crossroads of algebra and geometry. A compact Riemann surface has two characterizations, both as a compact complex 1-manifold and as a nonsingular algebraic curve. The function field of a Riemann surface, the set of all meromorphic functions defined on it, has a geometric interpretation as the set of maps from the curve to the Riemann sphere. The Riemann-Roch theorem is a statement about the dimension of certain subsets of the function field, subsets where we have required that poles may only occur at certain locations, and that their poles must not be too high a negative degree. The Riemann-Roch theorem is a bridge from the genus, a characteristic of a surface as a topological space, to algebraic information about its function field. A more

Date: August 28, 2015.

elementary treatment would build the Riemann-Roch theorem directly from singular homology. Though this treatment is valuable, we use the language of sheaves and sheaf cohomology from the beginning. This method is more abstract, but the methods we use are very powerful and easily admit generalization. We attempt to give a comprehensive introduction to the machinery we utilize, and hopefully the geometric intuition of more elementary methods is preserved. Familiarity with basic Riemann surface theory and differential forms is assumed, but much of the paper may be readable without it. Whenever we do not explicitly state the underlying space of a sheaf, it can be assumed to be a compact connected Riemann surface.

2. Sheaves

We begin by introducing sheaves. Speaking informally, a sheaf is an object defined on a topological space which can capture local data on any open set, but can also give us information about the global structure of the object.

Definition 2.1. A presheaf [1] $\mathcal{F}$ of abelian groups on a topological space $(X, \mathcal{T})$ is a collection of abelian groups $(\mathcal{F}(U))_{U \in \mathcal{T}}$ along with a collection of group homomorphisms ρ_V^U whenever U, V are open and $V \subset U$, such that
i) $\rho_W^U = \rho_W^V \circ \rho_V^U$.
ii) $\rho_U^U = Id_U$.

In this paper, these groups will always be groups of functions, so a presheaf is a set of functions on each open set in X, with ρ_V^U the typical restriction map of functions on U to functions on V. We thus use the typical notation $f_{|V} = \rho_V^U \circ f$.

Definition 2.2. A sheaf of abelian groups is a presheaf of abelian groups such that
i) Given a collection $(U_i)_{i \in I}$ of open sets with $U = \bigcup_{i \in I} U_i$, and elements $f_i \in \mathcal{F}(U_i)$, if $f_{i|U_i \cap U_j} = f_{j\,|U_i \cap U_j}$ for all $i, j \in I$, there exists $f \in \mathcal{F}(U)$ such that $f_{|U_i} = f_i$ for all $i \in I$.
ii) if $f, g \in \mathcal{F}(U)$ and $f_{|U_i} = g_{|U_i}$ for all $i \in I$, then $f = g$.

A sheaf lets us patch functions that agree on the overlap of their domains into a unique larger function.

Examples 2.3. Examples of sheaves on a Riemann surface:
$\mathcal{O}$: The sheaf of holomorphic functions.
Ω : The sheaf of holomorphic 1-forms.
$\underline{\mathbb{C}}$: The sheaf of locally constant functions with values in $\mathbb{C}$.
$\mathcal{E}$: The sheaf of differentiable functions.
$\mathcal{E}^{(1)}$: The sheaf of differentiable 1-forms.
$\mathcal{E}^{1,0}$: The sheaf of differentiable 1-forms which locally look like $f dz$, aka have no $d\overline{z}$ term.
$\mathcal{Z}$: The sheaf of closed 1-forms.

We now want to define the stalk at a sheaf. speaking very nonrigorously, we want to take a single point $x \in X$ and see how much information in the sheaf can be seen from the vantage point of x. The stalk should tell us what the sheaf looks like

[1] In other words, a presheaf is a contravariant functor from the category with open sets in X as objects and inclusion maps as morphisms to the category of abelian groups.

locally around x. To translate this statement into mathematical rigor, we just take the direct limit over all neighborhoods of x.

Definition 2.4. The stalk of $\mathcal{F}$ at x, denoted $\mathcal{F}_x$ is defined as $\mathcal{F}_x = \varinjlim_{U \ni x} \mathcal{F}(U)$.

Some facts which follow simply from the definition of direct limits are relevant. For instance, if $[f, x] = [g, x]$, where $f \in \mathcal{F}(U_1)$, $g \in \mathcal{F}(U_2)$, and $[f, x]$ and $[g, x]$ are their corresponding elements in $\mathcal{F}_x$, then there exists a neighborhood of x $V \subset U_1 \cap U_2$ such that $f_{|V} = g_{|V}$. Looking at stalks allows us to discuss the local properties of sheaves in greater rigor. For instance, when we define exact sequences of sheaves, we will say that a sequence of sheaves is exact when it is "locally exact" which means the corresponding sequence of stalks at a point x is exact for every $x \in X$.

3. Sheaf Cohomology

Sheafs are fundamentally a local-to-global object, something that takes local information and builds it into a global structure. Along those same lines, we first define sheaf cohomology as relative to an open cover, then take the direct limit over all possible open covers.

Definition 3.1. given an open cover $\mathcal{U} = (U_i)_{i \in I}$ and a sheaf $\mathcal{F}$, We define the zeroth cochain group $C^0(\mathcal{U}, \mathcal{F}) = \prod_{i \in I} \mathcal{F}(U_i)$. The zeroth cochain group is a collection of local sections of the sheaf $\mathcal{F}$. An element of $C^0(\mathcal{U}, \mathcal{F})$ is denoted $(f_i)_{i \in I}$.

Definition 3.2. With $\mathcal{U}, \mathcal{F}$ as above, for any natural number q we define the q-th cochain group $C^q(\mathcal{U}, \mathcal{F}) = \prod_{i_0, \ldots i_q \in I} \mathcal{F}(U_{i_0} \cap \cdots \cap U_{i_q})$. An element of $C^q(\mathcal{U}, \mathcal{F})$ is denoted $(f_{i_0, \ldots i_q})_{i_0, \ldots i_q \in I}$.

In particular, the first cochain group $C^1(\mathcal{U}, \mathcal{F}) = \prod_{i,j \in I} \mathcal{F}(U_i \cap U_j)$.

From the point of view of cohomology theory, it is natural that we define a coboundary operator. It is easy to see how an element of C^0 maps to an element of C^1 by restricting the two functions f_i and f_j to $U_i \cap U_j$ and subtracting them. For our purposes we only need define its actions on C^0 and C^1, but the definition extends in the obvious way to $\delta : \bigcup_{i=0}^{\infty} C^i(\mathcal{U}, \mathcal{F}) \to \bigcup_{i=1}^{\infty} C^i(\mathcal{U}, \mathcal{F})$. The operator takes any element of $C^i(\mathcal{U}, \mathcal{F})$ to $C^{i+1}(\mathcal{U}, \mathcal{F})$.

Definition 3.3. We define $\delta : C^0(\mathcal{U}, \mathcal{F}) \to C^1(\mathcal{U}, \mathcal{F})$:
$\delta(f_i) = (g_{i,j})$, where $g_{i,j} = f_{j\,|U_i \cap U_j} - f_{i|U_i \cap U_j}$.

We define $\delta : C^1(\mathcal{U}, \mathcal{F}) \to C^2(\mathcal{U}, \mathcal{F})$:
$\delta(f_{i,j}) = (g_{i,j,k})$, where $g_{i,j,k} = f_{i,j\,|U_i \cap U_j \cap U_k} + f_{j,k\,|U_i \cap U_j \cap U_k} + f_{k,i\,|U_i \cap U_j \cap U_k}$.

Intuitively, the boundary operator on C^0 measures the failure of a 0-cochain to paste together into a global section. If $\delta(f) = 0$, then $f_{i|U_i \cap U_j} = f_{j\,|U_i \cap U_j}$ for all $i, j \in I$, so by the first sheaf axiom we can paste (f_i) together into a function $f \in \mathcal{F}(X)$. Elements in C^0 as above are called closed cochains, and the group of closed 0-cochains is denoted $Z^0(\mathcal{U}, \mathcal{F})$. Clearly $Z^0(\mathcal{U}, \mathcal{F}) = \ker(\delta) \cap C^0(\mathcal{U}, \mathcal{F})$. Similarly, we define $Z^1(\mathcal{U}, \mathcal{F}) = \ker(\delta) \cap C^1(\mathcal{U}, \mathcal{F})$. In other words, $Z^1(\mathcal{U}, \mathcal{F}) = \{(f_{i,j})_{i,j \in I} : \delta(f_{i,j}) = 0\}$.

We define the coboundary set $B^1(\mathcal{U},\mathcal{F}) = \text{Im}(\delta) \cap C^1(\mathcal{U},\mathcal{F}) = \{(f_{i,j}) \in C^1(\mathcal{U},\mathcal{F}) : \exists (g_i) \in C^0(\mathcal{U},\mathcal{F})\text{such that} f_{i,j} = g_{j|U_i \cap U_j} - g_{i|U_i \cap U_j}\}$ We define $B^0(\mathcal{U},\mathcal{F}) = \{0\}$.

It is clear that $\delta \circ \delta = 0$. Consider $\delta \circ \delta(f) = \delta(f_j - f_i)_{i,j \in I} = (f_j - f_i + f_k - f_j + f_i - f_k)_{i,j,k \in I} = 0$. This means that $B^1(\mathcal{U},\mathcal{F}) \subset Z^1(\mathcal{U},\mathcal{F})$. what we wish to do with these groups is measure the failure of a closed cochain to be a coboundary; in other words, the failure of δ to be exact as a cochain operator.

Definition 3.4. We define the ith cohomology group $H^i(\mathcal{U},\mathcal{F}) = Z^i(\mathcal{U},\mathcal{F})/B^i(\mathcal{U},\mathcal{F})$

Since $B^0(\mathcal{U},\mathcal{F}) = 0$, $H^0(\mathcal{U},\mathcal{F}) - Z^0(\mathcal{U},\mathcal{F})$. As we saw before H^0 is the set of global sections of the sheaf F. A lot of what we do next is devoted to developing an understanding of the first cohomology group $H^1(\mathcal{F})$. However, our definition of cohomology is still dependent on a specific open cover.

Definition 3.5. given two open covers $\mathcal{B} = (B_k)_{k \in K}$ and $\mathcal{U} = (U_i)_{i \in I}$, we say $\mathcal{B}$ is finer than $\mathcal{U}$, denoted $\mathcal{B} < \mathcal{U}$, if for all $B_k \in \mathcal{B}$, there exists $U_{\tau i} \in \mathcal{U}$ such that $B_k \subset U_{\tau k}$, where τ is a map from K to I.

In other words, every set in $\mathcal{B}$ must be contained in a set in $\mathcal{U}$.

Definition 3.6. $\tau^{\mathcal{U}}_{\mathcal{B}} : Z^1(\mathcal{U},\mathcal{F}) \to Z^1(\mathcal{B},\mathcal{F})$: $\tau^{\mathcal{U}}_{\mathcal{B}}((f_{ij})) = (g_{kl})$ where $g_{kl} = f_{\tau k, \tau l|V_k \cap V_l}$.

This is the natural restriction mapping of cochains induced by our map τ. Since $\tau^{\mathcal{U}}_{\mathcal{B}}$ clearly maps $B^1(\mathcal{U},\mathcal{F})$ to $B^1(\mathcal{B},\mathcal{F})$, it factors through to a map from $H^1(\mathcal{U},\mathcal{F})$ to $H^1(\mathcal{B},\mathcal{F})$.
We omit proofs of the following two lemmas, which are tedious.

Lemma 3.7. *The map $\tau^{\mathcal{U}}_{\mathcal{B}} : H^1(\mathcal{U},\mathcal{F}) \to H^1(\mathcal{B},\mathcal{F})$ doesn't depend on the choice of the map τ.*

Lemma 3.8. *The map $\tau^{\mathcal{U}}_{\mathcal{B}} : H^1(\mathcal{U},\mathcal{F}) \to H^1(\mathcal{B},\mathcal{F})$ is injective.*

With those two facts, we can define $H^1(X,\mathcal{F}) = \varinjlim H^1(\mathcal{U},\mathcal{F})$ where the direct limit is taken over all open covers of X. Because all the maps $\tau^{\mathcal{U}}_{\mathcal{B}}$ are injective, we know that every $H^1(\mathcal{U},\mathcal{F})$ injects into $H^1(X,\mathcal{F})$, so the first sheaf cohomology of a space is at least as large as the cohomology relative to any of its open covers. We will usually just write $H^1(\mathcal{F})$ instead of $H^1(X,\mathcal{F})$.

Theorem 3.9. *$H^1(\mathcal{E}) = 0$.*

Proof. This fact relies on the idea of partitions of unity. We use the fact that for any open cover $\mathcal{U} = (U_i)_{i \in I}$, there exists a set of functions $(\psi_i)_{i \in I}$ such that
i) $\text{supp}(\psi_i) \subset U_i$
ii) for any $x \in X$, there is a neighborhood V_x so V_x only interesects finitely many $\text{supp}(\psi_i)$
ii) $\sum_{i \in I} \psi_i(x) = 1$ for all $x \in X$.

Knowing this, we consider $(f_{ij}) \subset Z^1(\mathcal{U},\mathcal{E})$. Now we want to find g_i so $g_i - g_j = f_{ij}$ on $U_i \cap U_j$. We set $g_i = \sum_{k \in I} \psi_k f_{ik}$. This can be defined as a differentiable element of $\mathcal{E}(U_i)$ since each function $\psi_i f_{ij}$ can be seen as zero outside of $\text{supp}(\psi_i)$, and because the sum is differentiable since all of its terms are and it is locally finite. Now, we have

$$\begin{aligned} g_i - g_j &= \sum_{k \in I} \psi_k [f_{ik} - f_{jk}] \\ &= \sum_{k \in I} \psi_k f_{ij} \\ &= f_{ij} \end{aligned}$$

□

A near-identical version of this proof shows that $\mathcal{E}^{1,0}$, $\mathcal{E}^{0,1}$, $\mathcal{E}^{(1)}$, and $\mathcal{E}^{(2)}$ all have trivial first cohomology.

4. Exact Sequences of Sheaves

Much of the reason sheaf cohomology is interesting is because of how it relates to exact sequences of sheaves. To define an exact sequence of sheaves, we must first define a sheaf homomorphism.

Definition 4.1. a sheaf homomorphism $\alpha : \mathcal{F} \to \mathcal{G}$ is a collection of maps $\alpha_U : \mathcal{F}(U) \to \mathcal{G}(U)$ such that the maps commute with restriction maps, meaning that for all open sets U, V so $V \subset U$,

$$\begin{CD} \mathcal{F}(U) @>{\alpha_U}>> \mathcal{G}(U) \\ @V{\rho^U_V}VV @VV{\widetilde{\rho}^U_V}V \\ \mathcal{F}(V) @>>{\alpha_V}> \mathcal{G}(V) \end{CD}$$

where the restriction maps of $\mathcal{F}$ are represented by ρ and the restriction maps of $\mathcal{G}$ are represented by $\widetilde{\rho}$.

This map induces homomorphisms of stalks $\mathcal{F}_x \xrightarrow{\alpha_x} \mathcal{G}_x$. We define the map to be a monormprhism of sheaves if it is a monomorphism on each stalk, and an epimorphism of sheaves if it is an epimorphism on each stalk. We define an exact sequence of sheaves $\mathcal{F} \xrightarrow{\alpha} \mathcal{G} \xrightarrow{\beta} \mathcal{H}$ to be exact if for all $x \in X$, the sequence

$$\mathcal{F}_x \xrightarrow{\alpha_x} \mathcal{G}_x \xrightarrow{\beta_x} \mathcal{H}_x$$

is exact. A short exact sequence $0 \to \mathcal{F} \xrightarrow{\alpha} \mathcal{G} \xrightarrow{\beta} \mathcal{H} \to 0$ is then exact if it induces a short exact sequence on stalks. We are interested in how this definition of exact sequences interacts with the induced maps on cohomology. The simplest example of this is just the maps induced on the zeroth cohomology groups, aka the global sections of X. Before we study that, however, we prove more general facts about the induced maps $\mathcal{F}(U) \to \mathcal{G}(U) \to \mathcal{H}(U)$.

Lemma 4.2. *suppose $\alpha : \mathcal{F} \to \mathcal{G}$ is a sheaf monomorphism. Then every map $\alpha_U : \mathcal{F}(U) \to \mathcal{G}(U)$ is a monomorphism.*

Proof. since α is a sheaf monomorphism, α_x is injective for all $x \in X$ suppose $\alpha_U(f) = 0$. Then for all $x \in U$, $\alpha_x([f, x]) = 0$. Thus, $[f, x] = [0, x]$ for all $x \in U$. Thus, around each x there exists an open neighborhood $V_x \ni x$ so that $f_{|V_x} = 0$. By sheaf axiom II, $f = 0$ on U. □

This lemma tells us that going from a sheaf homomorphism to a homomorphism of local sections preserves injectivity. However, it does not preserve surjectivity, as we can easily see from looking at the exponential map on $\mathbb{C}$.

$$\mathcal{O} \xrightarrow{\exp} \mathcal{O}^*$$

where $\exp(f) = e^{2\pi i f}$. That this is locally but not globally surjective is fundamentally a restatement about properties of the logarithm from complex analysis. We know around any point there is a simply connected neighborhood of the point contained in $\mathbb{C}^*$ within which exp is invertible. But on some non-simply connected neighborhoods, in particular $\mathbb{C}^*$ itself, the logarithm is not gloablly defined, so the function $\mathrm{Id} \in \mathcal{O}^*(\mathbb{C}^*)$ has no preimage under exp.

This example tells us that we cannot always get a short exact sequence of sections from a short exact sequence of sheaves, but as our next theorem shows we can at least preserve part of the exactness.

Theorem 4.3. *If $0 \to \mathcal{F} \xrightarrow{\alpha} \mathcal{G} \xrightarrow{\beta} \mathcal{H}$ is an exact sequence of sheaves, then for any open set U,*

$$0 \to \mathcal{F}(U) \xrightarrow{\alpha} \mathcal{G}(U) \xrightarrow{\beta} \mathcal{H}(U)$$

is exact. [2]

Proof. (i) That α is injective we proved in theorem 4.3.

(ii) $\mathrm{Im}(\alpha) \subset \mathrm{Ker}(\beta)$: Consider $f \in \mathcal{F}(U)$, $g = \alpha(f) \in \mathcal{G}(U)$. On any stalk $\mathcal{G}_x$, $[g, x] = \alpha_x[f, x]$, so $\beta_x[g, x] = [0, x]$. by basic properties of the direct limit, this means that on some open set V_x about x, $\beta(g|V_i) = 0$. By Sheaf axiom I, these $\beta(g|V_i)$ paste together into a global function 0, so $\beta(g) = 0$.

(iii) $\mathrm{Ker}(\beta) \subset \mathrm{Im}(\alpha)$. This is pretty much the same idea as (ii). consider $g \in \mathcal{G}(U)$ so $\beta(g) = 0$. Then on any stalk $[g, x] = \alpha_x[f_x, x]$, and so $g|V_x = \alpha(f_x)$ for some neighorhood $V_x \ni x$ and $f_x \in \mathcal{F}(V_x)$. This is an open cover, and $\alpha(f_{x|V_x \cap V_y}) = \alpha(f_{y|V_x \cap V_y}) = g_{|V_x \cap V_y}$. Since α is injective, this means $f_{x|V_x \cap V_y}) = f_{y|V_x \cap V_y}$, so they can be pasted together into a global function $f \in \mathcal{F}(U)$ so that $\alpha(f) = g$ □

We are most interested in this lemma when the original sequence is actually a short-exact sequence, meaning β is surjective, and when $U = X$. We then have that $0 \to H^0(\mathcal{F}) \xrightarrow{\alpha_0} H^0(\mathcal{G}) \xrightarrow{\beta_0} H^0(\mathcal{H})$. The amazing fact is that this sequence extends to a long exact sequence on cohomology, namely that in a canonical way we can attach the sequence $H^1(\mathcal{F}) \xrightarrow{\alpha_1} H^1(\mathcal{G}) \xrightarrow{\beta_1} H^1(\mathcal{H})$ to the end of this one.

Definition 4.4. if we have $0 \to \mathcal{F} \xrightarrow{\alpha} \mathcal{G} \xrightarrow{\beta} \mathcal{H} \to 0$, then we define the connecting homomorphism $\delta^* : H^0(\mathcal{H}) \to H^1(\mathcal{F})$ as follows: We consider $h \in H^0(\mathcal{H})$. since β is surjective on stalks, we proceed as in Theorem 4.3 (part iii of the proof) and construct an open cover $\mathcal{V} = (V_x)$, and a corresponding cochain $g = (g_x) \in C^0(\mathcal{V}, \mathcal{G})$ so that $\beta(g) = h_{|V_x}$. But since $\beta(g_x - g_y) = 0$ on $V_x \cap V_y$, we apply theorem 4.3 with $U = V_x \cap V_y$ and get an element $f = (f_{xy}) \in \mathcal{F}(V_x \cap V_y)$ so that $\alpha(f_{xy}) = g_x - g_y$. Doing this for all pairs (x, y) gets us a cochain $(f_{ij}) \in C^1(\mathcal{V}, \mathcal{F})$. Clearly (f_{ij})

[2] In other words, the functor taking a sheaf on a space to its group on an open set is left-exact.

is closed. We denote $\delta^* h = [f_{xy}]$, where $[f_{xy}]$ is the cohomology class of (f_{xy}) in $H^1(\mathcal{F})$.

That δ^* is well-defined is tedious but simple to check. Finally, we have reached the most important theorem of this section:

Theorem 4.5. *With $0 \to \mathcal{F} \xrightarrow{\alpha} \mathcal{G} \xrightarrow{\beta} \mathcal{H} \to 0$, δ^* as above, The sequence*

$$0 \longrightarrow H^0(\mathcal{F}) \xrightarrow{\alpha_0} H^0(\mathcal{G}) \xrightarrow{\beta_0} H^0(\mathcal{H}) \xrightarrow{\delta^*}$$

$$\dots \longrightarrow H^1(\mathcal{F}) \xrightarrow{\alpha_1} H^1(\mathcal{G}) \xrightarrow{\beta_1} H^1(\mathcal{H})$$

is exact.

The proof of this theorem is long but doesn't involve any original ideas, so we omit it.[3]. This theorem is important because it allows us to get information about H^1 groups, and write them in terms of quotients of H^0 groups, giving a concrete interpretation to H^1 groups.

Corollary 4.6. *If $0 \to \mathcal{F} \xrightarrow{\alpha} \mathcal{G} \xrightarrow{\beta} \mathcal{H} \to 0$ is exact, and $H^1(\mathcal{G}) = 0$, we have $H^1(\mathcal{F}) = H^0(\mathcal{H})/\beta H^0(\mathcal{G})$.*

Proof. We have the long exact sequence

$$0 \to H^0(\mathcal{F}) \xrightarrow{\alpha} H^0(\mathcal{G}) \xrightarrow{\beta} H^0(\mathcal{H}) \to H^1(\mathcal{F}) \to 0.$$

□

We will use this technique several times in the next section.

5. De Rham Cohomology

5.1. **A Few Exact Sequences.** We consider a few exact sequences of sheaves. First, we define two operators:

Definition 5.1. We know that $d(f) = f_z dz + f_{\overline{z}} d\overline{z}$. We define $\partial f = f_z dz$, and $\overline{\partial} f = f_{\overline{z}} d\overline{z}$. This definition on functions extends to forms in the same way the normal definition of the differential does. Where d maps $\mathcal{E}^{p,q}$ to $\mathcal{E}^{p+1,q} \oplus \mathcal{E}^{p,q+1}$, ∂ maps $\mathcal{E}^{p,q}$ to $\mathcal{E}^{p+1,q}$ and $\overline{\partial}$ maps $\mathcal{E}^{p,q}$ to $\mathcal{E}^{p,q+1}$. Clearly we have $d = \partial + \overline{\partial}$.

A function is holomorphic if $\overline{\partial} f = 0$, for example.

Theorem 5.2. *The following sequences are exact.*

$$0 \to \underline{\mathbb{C}} \to \mathcal{E} \xrightarrow{d} \mathcal{Z} \to 0$$
$$0 \to \underline{\mathbb{C}} \to \mathcal{O} \xrightarrow{d} \Omega \to 0$$
$$0 \to \mathcal{O} \to \mathcal{E} \xrightarrow{\overline{\partial}} \mathcal{E}^{0,1} \to 0$$
$$0 \to \Omega \to \mathcal{E}^{1,0} \xrightarrow{\overline{\partial}} \mathcal{E}^{(2)} \to 0$$

Proof. In the first sequence, since closed forms are locally exact, on any stalk, a closed 1-form will be locally the differential of some function. The other sequences are easy to check. □

[3] The complete proof can be found in chapter 15 of [3].

5.2. The De Rham Groups and a Special Case of Serre Duality.

Definition 5.3. We define the de Rham cohomology group $H^1_{DR}(X) = H^0(\mathcal{Z})/dH^0(\mathcal{E})$.

The de Rham cohomology is the quotient of closed forms over exact forms. The higher de Rham groups are defined analogously, but we will have no need for them.

Theorem 5.4. *De Rham's Theorem for sheaves:* $H^1_{DR}(X) = H^1(\underline{\mathbb{C}})$.

Proof. Since we have the exact sequence $0 \to \underline{\mathbb{C}} \to \mathcal{E} \xrightarrow{d} \mathcal{Z} \to 0$, this extends to a long-exact sequence on cohomology. Since $H^1(\mathcal{E}) = 0$, we have the sequence $0 \to \mathbb{C} \to H^0(\mathcal{E}) \to H^0(\mathcal{Z}) \to H^1(\underline{\mathbb{C}}) \to 0$ is exact, and the proof follows. □

We now want to show $H^0(\Omega) = H^1(\mathcal{O})^*$. This is extremely important, not only because it is a special case of Serre duality, but because it allows us to discuss the idea of a genus in a surface in a coherent way. Our proof relies on a theorem from analysis which we don't prove [4], but whose statement should be relatively intuitive.

Lemma 5.5. *if X is a compact connected Riemann surface, ρ a 2-form on X, then there exists a function f such that $\partial\overline{\partial} f = \rho$ if and only if $\int_X \rho = 0$.*[5]

Since the sequence $0 \to \mathcal{O} \to \mathcal{E} \xrightarrow{\overline{\partial}} \mathcal{E}^{0,1} \to 0$ is exact, and $H^1(\mathcal{E}) = 0$, we have $H^1(\mathcal{O}) = \mathcal{E}^{0,1}/d\mathcal{E}$. In this manner we treat $H^1(\mathcal{O})$ as cosets of $(0,1)$ forms. We take the map $\sigma : \sigma(\alpha) = [\overline{\alpha}]$, where the brackets represents taking a form to its equivalence class.

Lemma 5.6. *σ is surjective.*

Proof. We consider some coset $[\beta]$. Now, we want to pick some representative $\beta' = \overline{\partial} f$ so that $\partial\beta' = 0$. Since β' locally looks like $fd\overline{z}$, and $\partial\beta' = f_z d\overline{z}dz$, if $\partial\beta' = 0$ then $\overline{\partial\beta'} = 0$, and so $\overline{\beta'}$ is a holomorphic 1-form as desired. But, $\partial\beta' = 0$ is equivalent to $\partial\overline{\partial} f = -\overline{\partial}\beta$, and since by Stokes' theorem the latter side has zero integral, we can apply our lemma and find an f such that this is possible. Thus, σ is surjective.

□

Theorem 5.7. $H^0(\Omega) = H^1(\mathcal{O})^*$.

Proof. Still viewing $H^1(\mathcal{O})$ as cosets of $(0,1)$ forms, we define the bilinear map $B : H^0(\Omega) \oplus H^1(\mathcal{O}) \to \mathbb{C}$: $B(\alpha, [\beta]) = \int_X \alpha \wedge \beta$. Since $\int_X \alpha \wedge \overline{\partial} f = \int_X \overline{\partial}(\alpha f) = 0$ by Stokes' theorem, this is a well-defined map. Now, we define $\langle\ ,\ \rangle : H^0(\Omega) \oplus H^0(\Omega) \to \mathbb{C}$: $\langle \alpha, \beta \rangle = B(\alpha, \sigma(\beta))$.

But if $\alpha = fdz$ locally, $\langle \alpha, \alpha \rangle = \int_X fdz \wedge \overline{f}d\overline{z} = \int_X |f|^2 dz \wedge d\overline{z}$, which is positive if α is not identically 0. So, given an element $\alpha \in H^0(\Omega)$, it acts nontrivially on its image $\sigma\alpha$, so it is a nonzero element of its dual. Thus B establishes a dual pairing. □

The above proof additionally proves that σ is in fact an isomorphism, since it takes nonzero elements to nonzero elements and is therefore injective. We now therefore have an isomrophism $H^1(\mathcal{O}) = \overline{H^0(\Omega)}$. We now want to prove that the de Rham cohomology is equal to the direct sum of these two groups.

[4] A complete proof can be found in Chapter 9 of [2].

[5] The operator $\partial\overline{\partial}$ is locally the laplacian $\partial\overline{\partial} f = 4(\Delta f)dzd\overline{z}$.

Theorem 5.8. $H^0(\Omega) \oplus H^1(\mathcal{O}) \cong H^1_{DR}$.

Proof. Since $H^1(\mathcal{O}) = \overline{H^0(\Omega)}$, we can view this as $H^0(\Omega) \oplus \overline{H^0(\Omega)} \cong H^1_{DR}$. Now, we have the map $\Psi(\alpha, \overline{\beta}) = [\alpha + \overline{\beta}]$. This is injective: suppose $\alpha + \overline{\beta} = df$. locally, $\alpha = adz$, a holomorphic, and $\overline{\beta} = bd\overline{z}$, b antiholomorphic. Thus we would need to have $a = f_z$, $b = f_{\overline{z}}$. But this would mean f was both holomorphic and antiholomorphic, so f is constant and $\alpha = \beta = 0$.

The inverse map can be defined like this: we consider any closed 1-form ω, which locally equals $adz + bd\overline{z}$. We can split $a = a' + a''$, $b = b' + b''$. where a', b' are holomorphic, a'', b'' antiholomorphic. Then $d\omega = a''_{\overline{z}} - b'_z dz \wedge d\overline{z}$. since this must be zero if ω is closed, necesarily $a''_{\overline{z}} = b'_z$, and a holomorphic and antiholomorphic function are only equal if they are the same constant, so $a'' = c\overline{z} + d$, $b' = cz + e$. So, $a''dz + b'd\overline{z}$ is clearly exact, and the values of a' and b'' are uniquely determined by the cohomology class of ω. Thus the map $\omega \to (a'dz, b''d\overline{z})$ (in local coordinates) is injective, and is also an inverse of the above map Ψ.

□

6. Genus

An important aspect of the theory of Riemann surfaces is the idea of genus. Heuristically, genus measures the number of holes in a surface. A sphere has genus 0, a torus has genus 1, etc. Since any Riemann surface, ignoring the complex structure, is equivalent to a torus with some number of handles, this is the definition we want to keep in our mind. Defining this rigorously, we get that on a Riemann surface, (or orientable surface in general), the genus is half the dimension of the first singular cohomology group with respect to $\mathbb{Z}$.[6] We don't define singular cohomology in this paper, but the fact that it corresponds to the "number of holes" in this way is very intuitive. It is a theorem known as de Rham's theorem which states that on a smooth manifold (such as a Riemann surface) the de Rham cohomology and singular cohomology are equal in dimension. A reader familiar with Cech cohomology will also know that the first Cech cohomology of the constant sheaf is equal to the singular cohomology. So can compute this topological idea of genus using the de Rham cohomology or cohomology of the constant sheaf.

The second definition is that of geometric genus, which is $\dim H^0(\Omega)$, the dimension of the space of holomorphic 1-forms on the surface. The third is $\dim H^1(\mathcal{O})$, the dimension of the first cohomology group of the sheaf of holomorphic functions. By our theorems in the previous section, we have seen that these dimensions are all equal: $\dim H^0(\Omega) = \dim H^1(\mathcal{O}) = \frac{1}{2}\dim H^1_{DR}$. One of the most important aspects of the Riemann-Roch theorem is its ability to translate topological information about a Riemann surface, that is, its genus, into algebraic information about subgroups of its function field.

7. The Riemann-Roch Theorem

The Riemann-Roch Theorem is stated in the language of divisors. A divisor on a compact Riemann surface is essentially a way of encoding information about zeros and poles of a collection of meromorphic functions or forms.

[6]For more on this, look at [1].

Definition 7.1. A divisor on a compact Riemann surface X is a finite formal sum $\sum_{k=1}^{n} c_k s_k$, where $c_k \in \mathbb{Z}$ and $s_k \in X$.

We denote the set of divisors on X by $\mathrm{Div}(X)$. The most important example of a divisor is the divisor of a function. We define the function $() : \mathcal{M}(X) \to \mathrm{Div}(X)$; $(f) = \sum_{k=1}^{n} c_k s_k$ where locally around s_k, $f(z) = \sum_{i=c_k}^{\infty} a_i z^i$. For example, if $f : \mathbb{P}^1 \to \mathbb{P}^1$, $f(z) = z^2, (f) = 2[1,0] - 2[0,1]$, because z^2 has a zero of order 2 at the bottom of the Riemann sphere and a pole of order 2 at the point at infinity. If a divisor is the divisor of a function, it is called a principal divisor.

Definition 7.2. We can take the set $\mathrm{Div}(X)$ and quotient out all principal divisors. The resulting group is called the divisor class group, and is denoted $\mathrm{Cl}(X) = \mathrm{Div}(X)/\mathrm{Prin}(X)$. [7]

Given a meromorphic 1-form ω, we cover X with coordinate charts, so that locally $\omega = f dz$. we define its divisor as the zeroes and poles of f. For example, consider the form dz on the Riemann sphere. It has no zeros, but in fact has a pole of order 2 at the point at infinity. We have the standard coordinate charts on the Riemann-sphere, with coordinates z on $\mathbb{P}^1\backslash[1,0]$ and ξ on $\mathbb{P}^1\backslash[0,1]$, $\xi = 1/z$. So, around the point at infinity $[0,1]$ we have $d\xi = \frac{\partial \xi}{\partial z} dz = \frac{-1}{z^2} dz = \xi^2 dz$, so $dz = \frac{1}{\xi^2} d\xi$, so in local coordinates we have a double pole at the point at infinity.

A divisor is called effective if $c_k \geq 0$ for all k. A divisor of a holomorphic function is always effective, for instance. We can define the addition and subtraction of divisors by adding the formal sums. We say $D \geq D'$ if $D - D'$ is effective.

Definition 7.3. We define a function $\mathrm{Deg} : \mathrm{Div}(X) \to \mathbb{Z} : \mathrm{Deg}(\sum c_k s_k) = \sum c_k$.

If D is principal, $\mathrm{Deg}(D) = 0$, since on a compact Riemann surface the zeroes and poles of a function must sum to zero.

Definition 7.4. given two meromorphic 1-forms ω_1 and ω_2, their quotient ω_1/ω_2 is a well-defined meromorphic function in a natural way. Select some coordinate chart (U_z) and locally we get $\omega_1/\omega_2 = f_1 dz/f_2 dz = f_1/f_2$. Since a different coordinate chart changes the values of f_1 and f_2 proportianally, the function doesn't depend on the choice of chart and so is well-defined. This means that any two meromorphic 1-forms lie in the same coset of the divisor class group. This is an extremely important divisor, and is called the canonical divisor, denoted K.

We now want to associate a sheaf to a divisor.

Definition 7.5. Given a divisor D, we define the sheaf $\mathcal{O}_D$ on X, so $\mathcal{O}_D(U) = f \in \mathcal{M}(U) : (f) \geq -D$. With the restriction maps just restriction of functions, this is a sheaf.

The sheaf of a divisor D is the collection of functions that have their order bounded below by $-D$. If $-D$ is negative at a point, the functions may have a pole there, but only of order greater than or equal to that of the divisor. if $-D$ is positive at a point, the functions must all have a zero of at least that order at that point. If D has no points in an open set U, then $\mathcal{O}_D(U) = \mathcal{O}_D(U)$. If $D = 0$ then $\mathcal{O}_D = \mathcal{O}$. If $D \leq D'$, then $\mathcal{O}_D(U) \subset \mathcal{O}'_{\mathcal{D}}(U)$ for all open sets $U \subset X$.

---

[7] This is also known as the Picard group and is denoted by Pic(X).

Lemma 7.6. *If D and D' lie in the same coset of the divisor class-group, then $\mathcal{O}_D \cong \mathcal{O}_{D'}$.*

Proof. If D and D' are equivalent modulo principal divisors, then $D' - D = (g)$ for some $(g) \in \mathcal{M}(X)$. Now, for any $U \subset X$, the map $\phi_U : \mathcal{O}_D(U) \to \mathcal{O}_{D'}(U) : \phi_U(f) = g_{|U} f$ is an isomorphism. Clearly this commutes with restriction so is an isomorphism of sheaves. □

Lemma 7.7. $\mathcal{O}_{D+K} \cong \Omega_D$.

Proof. Since $\mathcal{O}_{(\omega_1)} \cong \mathcal{O}_{(\omega_2)}$, where ω_1 and ω_2 are both meromorphic 1-forms, we can take an arbitrary meromorphic 1-form ω and say $K = (\omega)$. Now, the map $\psi_U : \mathcal{O}_{D+K}(U) \to \Omega_D(U) : \psi_U(f) = \omega_{|U} f$ is an isomorphism. □

The global sections of a divisor sheaf of are interest to us. These are the meromorphic functions that are "bounded" by the divisor. We control the degree and location of their poles, and insist on zeroes of at least some degree at certain locations. Calculating the dimension of $H^0(X, \mathcal{O}_D)$ is the most immediate use of the Riemann-Roch theorem. By being able to find the dimension of functions that are bounded in this way, we get very tight control over the meromorphic functions on X. [8]

Example 7.8. If $\mathrm{Deg}(D) < 0$, then $\dim H^0(X, \mathcal{O}_D)$. For suppose $(f) \in H^0(X, \mathcal{O}_D)$. then $(f) \geq -D$, so $\mathrm{Deg}((f)) \geq \mathrm{Deg}(-D) > 0$, which is impossible unless $f \equiv 0$. Thus $H^0(X, \mathcal{O}_D) = \{0\}$.

We now want to look at sheaf monomorphisms of the form $\mathcal{O}_D \to \mathcal{O}_{D'}$, where $D \leq D'$. Let us first consider the case $D' = D + p$, ie the case where the two divisors differ only by a single point. It is clear we then have the inclusion map $\mathcal{O}_D \to \mathcal{O}_{D+p}$. It is clear that locally around p the functions in $\mathcal{O}_{D+p}$ have one additional degree of freedom than those in $\mathcal{O}_D$, but on an open set not containing p the two are equal. Can we find another sheaf in order to make this an exact sequences of sheaves?

Definition 7.9. We define the skyscraper sheaf $\mathbb{C}_p$: $\mathbb{C}_p(U) = \begin{cases} \mathbb{C} & \text{if } p \in U \\ \{0\} & \text{if } p \notin U. \end{cases}$

Lemma 7.10. *For any Riemann surface X, $H^0(X, \mathbb{C}_p) = \mathbb{C}$, $H^1(X, \mathbb{C}_p) = \{0\}$.*

Proof. (i) Consider any open cover $\mathcal{U}$ of X. Since it can be refined so that only one U_i contains p, we assume without loss of generality that this is the case. Then $H^0(\mathcal{U}, \mathbb{C}_p) = \mathbb{C}$, and taking the direct limit, $H^0(\mathbb{C}_p) = \mathbb{C}$.

(ii) Consider any element $\xi \in H^1(X, \mathbb{C}_p)$. For any open cover $\mathcal{U}$, ξ has a representative $\eta \in Z^1(\mathcal{U}, \mathbb{C}_p)$. But we can choose a refinement of $\mathcal{U}$, $\mathcal{B}$, such that p is only contained in a single open set. But then no set $B_i \cap B_j$ contains p, so $Z^1(\mathcal{B}, \mathbb{C}_p) = \{0\}$. Since a refinement produces an injection, $Z^1(\mathcal{U}, \mathbb{C}_p) = \{0\}$, so $\eta = 0$ and $\xi = 0$. Thus $H^1(\mathbb{C}_p) = \{0\}$. □

[8] That these dimensions are finite is a nontrivial theorem from analysis, and we don't include it. All we require is that $H^0(\mathcal{O})$ is finite, and the rest follow from the induction we will do in the proof of Riemann-Roch. For a proof, see chapter 14 of [3].

We define a map $\beta_p : \mathcal{O}_{D+p} \to \mathbb{C}_p$: if $f \in \mathcal{O}_{D+p}(U)$ and $p \in U$, we choose a coordinate chart (V, z) such that $V \subset U$, $z(V) = \mathbb{D}$ and $z(p) = 0$. Then locally around p, $f = \sum_{n=-k-1}^{\infty} a_n z^n$, where D has value k at p and $D+p$ has value $k+1$ at p. Then, we set $\beta(f) = a_{-k-1}$ for $f \in \mathcal{O}_{D+p}(U)$, and for any U so that $p \notin U$ set β to be the zero homomorphism. Clearly the sequence

$$0 \to \mathcal{O}_D \xrightarrow{i} \mathcal{O}_{D+p} \xrightarrow{\beta} \mathbb{C}_p \to 0$$

is exact. This sequence is interesting to us because it induces a long exact sequence on cohomology

$$0 \longrightarrow H^0(X, \mathcal{O}_D) \xrightarrow{i_0} H^0(X, \mathcal{O}_{D+p}) \xrightarrow{\beta_0} \mathbb{C} \xrightarrow{\delta^*}$$

$$\ldots \longrightarrow H^1(X, \mathcal{O}_D) \xrightarrow{i_1} H^1(X, \mathcal{O}_{D+p}) \xrightarrow{\beta_1} 0.$$

Theorem 7.11. ***The Riemann-Roch Theorem:*** *Suppose D is a divisor on a compact Riemann surface X of genus g. Then*[9]

$$\dim H^0(X, \mathcal{O}_D) - \dim H^1(X, \mathcal{O}_D) = 1 - g + Deg(D)$$

Proof. We will do a proof by induction. We first prove the theorem is true for $D = 0$, then prove that if it is true for D, then it is true for $D + p$ and for $D - p$. Since every divisor equals $p_1 + \cdots + p_n - q_1 \cdots - q_m$, we can reach any divisor with a finite number of additions and subtractions of points.
Step 1: Suppose $D = 0$. If $D = 0$, $\mathcal{O}_D = \mathcal{O}$, and $H^0(X, \mathcal{O})$ consists of the constant functions, so has dimension 1, and $H^1(X, \mathcal{O}) = g$ by definition. Thus we have $1 - g = 1 - g$ and we are done.
Step 2: Suppose our theorem holds for a divisor D. We consider the divisor $D + p$. As above, we have the long exact sequence

$$0 \longrightarrow H^0(X, \mathcal{O}_D) \xrightarrow{i_0} H^0(X, \mathcal{O}_{D+p}) \xrightarrow{\beta_0} \mathbb{C} \xrightarrow{\delta^*}$$

$$\ldots \longrightarrow H^1(X, \mathcal{O}_D) \xrightarrow{i_1} H^1(X, \mathcal{O}_{D+p}) \xrightarrow{\beta_1} 0.$$

We want to split the sequence in two at $\mathbb{C}$. This means we want to replace $\mathbb{C}$ with $\text{Im}(\beta_0)$ in the top diagram, and with $\mathbb{C}/\text{Im}(\beta_0)$ in the bottom diagram.

$$0 \longrightarrow H^0(X, \mathcal{O}_D) \xrightarrow{i_0} H^0(X, \mathcal{O}_{D+p}) \xrightarrow{\beta_0} \text{Im}(\beta_0) \longrightarrow 0$$

$$0 \longrightarrow \mathbb{C}/\text{Im}(\beta_0) \xrightarrow{\widetilde{\delta^*}} H^1(X, \mathcal{O}_D) \xrightarrow{i_1} H^1(X, \mathcal{O}_{D+p}) \xrightarrow{\beta_1} 0$$

[9] A characterization of Riemann-Roch which we won't use but should at least mention is that of the Euler characteristic of a sheaf. We define the Euler characteristic in the natural way as an alternating sequence on cohomology $\chi(\mathcal{F}) = \sum_i (-1)^i \dim H^i(\mathcal{F})$. So, the Riemann-Roch theorem states

$$\chi(msO_D) = 1 - g + \text{Deg}(D)$$

If we wanted to generalize Riemann-Roch to higher degree algebraic varieties, this would be the form we would use.

Since $\mathbb{C}$ is 1-dimensional, we know that either $\text{Im}(\beta_0) = \mathbb{C}$ and $\mathbb{C}/\text{Im}(\beta_0) = 0$ or vice versa. Since $\text{Im}(\beta_0) = Ker\delta^*$, δ^* factors through $\mathbb{C}/\text{Im}(\beta_0)$ and induces the injective map $\widetilde{\delta^*}$. Now that we have two short exact sequences, we can do dimension counting: by basic linear algebra in a short exact sequence of vector spaces $0 \to A \to B \to C \to 0$, $\dim(B) = \dim(A) + \dim(C)$. So, we know

$$\dim H^0(X, \mathcal{O}_{D+p}) = \dim H^0(X, \mathcal{O}_D) + \dim \text{Im}((\beta_0)),$$

and

$$\dim H^1(X, \mathcal{O}_D) = \dim H^1(X, \mathcal{O}_{D+p}) + \dim \mathbb{C}/\text{Im}((\beta_0)).$$

Adding the two together, we get

$$\dim H^0(X, \mathcal{O}_{D+p}) + \dim H^1(X, \mathcal{O}_D) = \dim H^0(X, \mathcal{O}_D) + \dim H^1(X, \mathcal{O}_{D+p}) + 1.$$

Reorganizing and applying Riemann-Roch to the case D, we get

$$\begin{aligned}\dim H^0(X, \mathcal{O}_{D+p}) - \dim H^1(X, \mathcal{O}_{D+p}) &= \dim H^0(X, \mathcal{O}_D) - \dim H^1(X, \mathcal{O}_D) + 1\\ &= 1 - g + \text{Deg}(D) + 1\\ &= 1 - g + \text{Deg}(D + p).\end{aligned}$$

To prove that the case D implies the case $D - p$, we replace D with $D - p$ above and reverse the equation, and get

$$\begin{aligned}\dim H^0(X, \mathcal{O}_{D-p}) - \dim H^1(X, \mathcal{O}_{D-p}) &= \dim H^0(X, \mathcal{O}_D) - \dim H^1(X, \mathcal{O}_D) - 1\\ &= 1 - g + \text{Deg}(D) - 1\\ &= 1 - g + \text{Deg}(D - p).\end{aligned}$$

This completes the proof. □

8. Serre Duality

Serre duality states that $H^0(\Omega_D) = H^1(\mathcal{O}_{-D})^*$. Before we attempt a proof, we look at the use this statement has to us. First of all and most importantly, the two must have equal dimension. This lets us rephrase the Riemann-Roch theorem without the use of cohomology, making it a statement entirely about global sections of functions bounded by two divisors. In this form, We get the statement

$$\dim H^0(\mathcal{O}_D) - H^0(\mathcal{O}_{K-D}) = 1 - g + \text{Deg}(D).$$

This gives us a much more concrete statement of Riemann-Roch. Even if we have no concrete interpretation of cohomology we get a powerful statement about dimensions of spaces of either functions or forms.

8.1. **Residue as a Linear Functional on $H^1(\mathcal{O})$.**

We want to take the linear functional Res on meromorphic functions and use it in quite a different manner. Attempting to define Res on a Riemann surface, we quickly see that there is no coordinate-invariant way to define it on a function. However, it is well behaved as a functional on meromorphic 1-forms.

Definition 8.1. Residue of a meromorphic 1-form: given ω, to compute $\mathrm{Res}_a(\omega)$, we select a local coordinate chart (U, z) about a so $z(U) = \mathbb{D}$ and $z(a) = 0$. Then in this chart, $\omega = \left(\sum_{n=-k}^{\infty} a_n z^n\right) dz$. We define

$$\mathrm{Res}_a(\omega) = \int_{|z|=\varepsilon} \omega = a_{-1}.$$

Because under a coordinate change, z and dz transform inversely to each other, a simple computation shows that this is invariant under choice of coordinates.

Definition 8.2. Given some open cover, $\mathcal{U} = (U_i)$ a cochain $(\omega_i) \in C^0(\mathcal{U}, \mathcal{M}^{(1)})$ is a Mittag-Leffler distribution if $\omega_i - \omega_j$ is a holomorphic 1-form on $U_i \cap U_j$. This means $\delta(\omega_i) \in Z^1(\mathcal{U}, \Omega)$.

The poles of a Mittag-Leffler distribution are either contained in exactly one open set, or that the elements of (ω_i) have the same principal parts whenever they overlap at a pole. Given a Mittag-Leffler distribution $\omega = (\omega_i)_{i \in I}$, and given $a \in X, a \in U_j$, we define $\mathrm{Res}_a(\omega) = \mathrm{Res}_a(\omega_j)$. Since $\omega_i - \omega_j$ is holomorphic, $\mathrm{Res}_a(\omega_j) = \mathrm{Res}_a(\omega_i)$ for all $U_i \cap U_j \ni a$, so this is well-defined.

What we want to do is extend residue to an operator on $H^1(\Omega)$. This is an intuitive thing to do if we view a Mittag-Leffler distribution as the natural place in which to talk about residue. We want some definition of Res such that $\mathrm{Res}(\delta\nu) = \mathrm{Res}(\nu)$ for a Mittag-Leffler distribution ν.

We have the short exact sequence $0 \to \Omega \to \mathcal{E}^{1,0} \xrightarrow{d} \mathcal{E}^{(2)} \to 0$. Since by Theorem 3.9, $H^1(\mathcal{E}^{1,0}) = H^1(\mathcal{E}^{(2)}) = \{0\}$ $H^1(X, \Omega) = \mathcal{E}^{(2)}/d(\mathcal{E}^{1,0})$. In order to continue, we must create an explicit isomorphism between $H^1(X, \Omega)$ and $\mathcal{E}^{(2)}/d(\mathcal{E}^{1,0})$. $\delta\nu \in Z^1(\mathcal{U}, \Omega) \subset B^1(\mathcal{U}, \mathcal{E}^{1,0})$ (since $H^1(\mathcal{E}^{1,0}) = 0$), we know $\omega_i - \omega_j = \sigma_i - \sigma_j$, where $(\sigma_i) \subset C^0(\mathcal{U}, \mathcal{E}^{1,0})$ since $\omega_i - \omega_j$ is holomorphic, $d(\sigma_i - \sigma_j) = 0$, so $d\sigma_i = d\sigma_j$ on $U_i \cap U_j$. We can therefore paste together $d\sigma_i$ into a global function $\tau \subset \mathcal{E}^{(2)}(X)$. Looking back at the construction of δ^*, one can check that taking $\delta\nu$ to τ is the inverse of the map δ^*, which means that τ is a representative of $\delta^*\nu$'s class in $H^1(\mathcal{U}, \Omega)$.

Theorem 8.3. *With τ and ν defined as above, $\frac{1}{2\pi i}\int_X \tau = Res(\nu)$.*

Proof. We set the poles of ν as $\{a_1 \dots a_n\}$ and we set $X' = X \backslash \{a_1, \dots, a_n\}$. We want to find a function σ so that $d\sigma = \tau$ everywhere except the poles of ν, but we want to preserve the residues of ν around these poles. To do this, we observe that $\sigma_i - \omega_i = \sigma_j - \omega_j$ on $U_i \cap U_j$, and we paste them together into a global function $\sigma \in \mathcal{E}^{1,0}(X)$. Since $d\omega_i = 0$ wherever ω_i is holomorphic, $d\sigma = \tau$ on X'. Now we can select elements of $\mathcal{U} : U_{i(k)} \ni a_k$, and local coordinate charts (V_k, z_k) around each pole. We then construct functions f_k so that $f_k = 1$ in a neighborhood around a_k and $\mathrm{supp}(f_k) \subset V_k$. setting $g = 1 - \sum f_k$, we see that $g = 0$ in a neighborhood of each pole, so $g(\sigma)$ can be continued across all poles, making it into a globally defined 1-form. Now we can integrate:

$$\begin{aligned}\int_X \tau &= \int_X d(\sigma) \\ &= \int_X d(g\sigma) + \sum_{k=1}^{n} \int_X d(f_k\sigma).\end{aligned}$$

By Stokes' theorem, $\int_X d(g\sigma) = 0$. Since $d(f_k\sigma) = 0$ wherever f_k is constant, it is zero in some ϵ neighborhood N_k of a_k on which $f_k = 1$ and some larger neighorhood M_k on which $f_k = 0$. So, we have

$$\begin{aligned}\sum_{k=1}^{n} \int_X d(f_k\sigma) &= \sum \int_{M_k \setminus N_k} d(f_k\sigma) \\ &= \sum \int_{\partial M_k} f_k\sigma - \int_{\partial N_k} f_k\sigma \\ &= -\sum \int \partial N_k \left(\sigma_{i(k)} - \omega_{i(k)}\right) \\ &= \sum 2\pi i \mathrm{Res}_{a_k} \omega_{i(k)} \\ &= 2\pi i \mathrm{Res}(\nu).\end{aligned}$$

□

Now we can define $\mathrm{Res}([\delta\nu]) = \mathrm{Res}(\nu)$, and we have successfully extended Res to a linear functional on $H^1(\mathcal{O})$.

8.2. **The Proof of Serre Duality.** Given a divisor D and an open $U \subset X$, if $\xi \in \Omega_{-D}(X)$, $\eta \in \mathcal{O}_D(U)$, $\xi_{|U}\eta \in \mathcal{O}(U)$. Thus, There is a map induced by ξ on cochains, and since this map takes coboundaries to coboundaries it induces a map $H^0(\Omega_D) \times H^1(\mathcal{O}_{-D}) \to H^1(\mathcal{O})$. composing this with Res gets us a bilinear map $\langle\ ,\ \rangle : \langle \eta, \xi \rangle = \mathrm{Res}(\eta\xi)$. This induces a map $i_D : H^0(\Omega_D) \to H^1(\mathcal{O}_{-D})^*$. We want to show that this map is an isomorphism. Doing this in full is a long and technical proof. We follow the proof in Chapter 17 of [3], and any details we leave out can be found there.

Theorem 8.4. *The map i_D is injective.*

Proof. We pick $\omega \in H^0(\Omega_D)$, and consider any point $p \notin D \cup (\omega)$. We want to show that ω is nonzero as an element of $H^1(\mathcal{O}_{-D})^*$, so we find some open cover and a cochain ξ in that cover such that $\mathrm{Res}(\omega\xi) \neq 0$. Since $D \cup (\omega)$ is an isolated, set, we can pick an open set and coordinate (U_0, z), such that U_0 doesn't intersect $D \cup (\omega)$ at all. So, locally $\omega = f(z)dz$, where $f \in \mathcal{O}^*(U_0)$, a holomorphic function which is never zero in U_0. Now we set $U_1 = X \cap p$, and $\mathcal{U} = \{U_0, U_1\}$. Then we define $(f_i) \in C^0(\mathcal{U}, \mathcal{M})$, where $f_0 = 1/(zf)$ and $f_1 = 0$. Then $\delta(f_i) \subset Z^1(\mathcal{U}, \mathcal{O}_{-D})$. But $\omega\delta(f_i) = dz/z$ on $U_0 \cap U_1$, so $\mathrm{Res}(\omega\delta(f_i)) = 1$. □

We know from the long exact sequence of cohomology groups that if we have an inclusion $0 \to \mathcal{O}_D \to \mathcal{O}_{D'}$ then we have a surjection $H^1(\mathcal{O}_{D'}) \to H^1(\mathcal{O}_D)$.[10] Taking the dual, we have a monomorphism of the dual groups

$$0 \to H^1(\mathcal{O}_D)^* \xrightarrow{i_{D'}^D} H^1(\mathcal{O}_{D'})^*.$$

We assert the following detail-checking lemmas without proof:

[10]This is the inverse of the monomorphism induced by the long exact sequence.

Lemma 8.5. *The following diagram commutes.*

$$\begin{CD} 0 @>>> H^1(\mathcal{O}_D)^* @>{i^D_{D'}}>> H^1(\mathcal{O}_{D'})^* \\ @. @A{i_D}AA @A{i_{D'}}AA \\ 0 @>>> H^0(\Omega_{-D}) @>>> H^0(\Omega_{-D'}) \end{CD}$$

Lemma 8.6. *If $\alpha \in H^1(\mathcal{O}_D)^*$ and $\omega \in H^0(\Omega_{-D'})$ are such that*

$$i^D_{D'}(\alpha) = i_{D'}(\omega)$$

then ω is also an element of $H^0(\Omega_{-D})$ and $\alpha = i_D(\omega)$.

Given a function $\psi \in \mathcal{O}_B$, there is an induced isomorphism $\psi : \mathcal{O}_{D-B} \to \mathcal{O}_D$. This induces a linear mapping $\psi : H^1(\mathcal{O}_{D-B}) \to H^1(\mathcal{O}_D)$, and a dual mapping

$$\psi^* : H^1(\mathcal{O}_D)^* \to H^1(\mathcal{O}_{D-B})^*.$$

Since ψ^* is a dual mapping induced by the linear form Res, We know that $\operatorname{Res}((\psi^*\omega)\xi) = \operatorname{Res}(\omega(\psi\xi))$. So, we know that the following diagram commutes:

$$\begin{CD} H^1(\mathcal{O}_D)^* @>{\psi^*}>> H^1(\mathcal{O}_{D-B})^* \\ @A{i_D}AA @A{i_{D-B}}AA \\ H^0(\Omega_{-D}) @>>> H^0(\Omega_{-D+B}) \end{CD}$$

It is clear that unless $\psi = 0$, ψ^* is injective. This is because we can factor the map $\mathcal{O}_{D-B} \xrightarrow{\psi} \mathcal{O}_D$ into

$$\mathcal{O}_{D-B} \to \mathcal{O}_{D+(\psi)} \xrightarrow{\psi} \mathcal{O}_D$$

Since the map $\mathcal{O}_{D-B} \to \mathcal{O}_{D+(\psi)}$ is an inclusion, we have an epimorphism on the H^1 groups and so a monomorphism on the duals.

With these lemmas out of the way, which all amounted to verifying things we would expect to be true, we can now complete the proof. We proved injectivity of i_D, so we just need to prove surjectivity.

Theorem 8.7. *Serre Duality: the map $i_D : H^0(\Omega_{-D}) \to H^1(\mathcal{O}_D)^*$ is an isomorphism.*

Proof. We choose an arbitrary nonzero element $\alpha \in H^1(\mathcal{O}_D)^*$, and we want to show that there exists some element of $H^0(\Omega_{-D})$ such that $i_D(\omega) = \lambda$. We choose a point $p \in X$. Set $D_n = D - n(p)$. We want this divisor because we can pick an n large enough that $H^0(\mathcal{O}_{D_n})$ is trivial. Our strategy is to use dimension counts to prove that the image of two maps must nontrivially intersect.

We choose an arbitrary $\psi \in H^0(\mathcal{O}_{np})$. Then this induces an injection $\psi^* : H^1(\mathcal{O}_D)^* \to H^1(\mathcal{O}_{D_n})^*$. Since this map is injective, the map $\widetilde{\lambda} : \psi \to \psi^*\lambda$ which maps $H^0(\mathcal{O}_{np}) \to H^1(\mathcal{O}_{D_n})^*$ is injective (checking this is just working with the previous commutative diagram). Set $V = \operatorname{Im}(\lambda^*)$. This is isomorphic to $H^0(\mathcal{O}_{np})$. By Riemann-Roch

$$\dim V \geq 1 - g + n.$$

What we want to do is intersect this vector space with the image of i_{D_n}. Then we can use lemma 8.7 to show that this is actually an element of $\mathrm{Im}(i_D)$.

We set $W = \mathrm{Im}(i_{D_n})$. We have

$$\dim W = \dim H^0(\Omega_{-D_n}) \geq 1 - g + \mathrm{Deg}(-D_n) = 1 - g + n - \mathrm{Deg}(D).$$

Now, if we have n so $\mathrm{Deg}(D_n) < 0$, then $H^0(\mathcal{O}_{D_n})$ is trivial, so

$$\dim H^1(\mathcal{O}_{D_n})^* = \dim H^1(\mathcal{O}_{D_n}) = g - 1 + n - \deg D$$

But now we are nearly done, because we have

$$\begin{aligned}\dim V + \dim W &\geq 2 - 2g - \mathrm{Deg}(D) + 2n \\ &> n + g - 1 - \deg D \\ &= \dim H^1(\mathcal{O}_{D_n})^*\end{aligned}$$

Thus, there is a nonzero element in $V \cap W$. That is, there exist elements $\psi \in H^0(\mathcal{O}_{np})$, $\omega \in H^0(\Omega_{-D_n})$, such that $\lambda^*\psi = i_{D_n}(\omega)$.
It is a fact that $\frac{1}{\psi} \in H^0(\mathcal{O}_{(\psi)})$ We set $D' = D - (\psi)$. Then

$$\begin{aligned}i^D_{D'}(\lambda) &= \frac{1}{\psi}(\widetilde{\lambda\psi}) \\ &= \frac{1}{\psi}(\psi^*\lambda) \\ &= \frac{1}{\psi} i_{D_n}(\omega) \\ &= i_{D'}\left(\frac{1}{\psi}\omega\right).\end{aligned}$$

Then by lemma 8.7, $\frac{1}{\psi}\omega$ is our desired element of $H^0(\Omega_{-D})$.

□

We will now outline a very beautiful alternate proof of Serre duality, one which is no less technical but may provide more intuition for the result. More details can be found in [2]. We will proceed by induction. We must show that if Serre duality holds for D and $-D$. Then $H^0(\Omega_{-D'}) = H^1(\mathcal{O}_{D'})^*$ for D' and $-D'$, where $D' = D + p$ or $D - p$.
We consider the exact sequence of sheaves $0 \to \mathcal{O}_D \to \mathcal{O}_{D+p} \to \mathbb{C}_p \to 0$. This induces a long exact sequence

$$0 \to H^0(\mathcal{O}_D) \to H^0(\mathcal{O}_{D+p}) \to \mathbb{C} \to H^1(\mathcal{O}_D) \to H^1(\mathcal{O}_{D+p}) \to 0.$$

We rewrite this as

$$0 \to H^0(\Omega_{D-K}) \to H^0(\Omega_{D+p-K}) \to \mathbb{C} \to H^1(\Omega_{D-K}) \to H^1(\Omega_{D+p-K}) \to 0.$$

Now, replacing D with $K - D - p$, we have the long exact sequence

$$0 \to H^0(\mathcal{O}_{K-D-p}) \to H^0(\mathcal{O}_{K-D}) \to \mathbb{C} \to H^1(\mathcal{O}_{K-D-p}) \to H^1(\mathcal{O}_{K-D}) \to 0$$

and the corresponding long exact sequence on the duals

$$0 \to H^1(\mathcal{O}_{K-D})^* \to H^1(\mathcal{O}_{K-D-p})^* \to \mathbb{C} \to H^0(\mathcal{O}_{K-D})^* \to H^0(\mathcal{O}_{K-D-p})^* \to 0.$$

Using our maps from before, we have the chain complex

$$\begin{array}{ccccccccc}
0 & \longrightarrow & H^0(\Omega_{D-K}) & \longrightarrow & H^0(\Omega_{D+p-K}) & \longrightarrow & \mathbb{C} & \longrightarrow & H^1(\Omega_{D-K}) & \longrightarrow & \dots \\
\downarrow & & \downarrow{\scriptstyle i_{K-D}} & & \downarrow{\scriptstyle i_{K-D-p}} & & \downarrow & & \downarrow{\scriptstyle i^*_{D-K}} & & \\
0 & \longrightarrow & H^1(\mathcal{O}_{K-D})^* & \longrightarrow & H^1(\mathcal{O}_{K-D-p})^* & \longrightarrow & \mathbb{C} & \longrightarrow & H^0(\mathcal{O}_{K-D})^* & \longrightarrow & \dots
\end{array}$$

That this diagram commutes must be verified. Then, since by our induction hypothesis i_{K-D} is an isomorphism, and the other i maps are injective and their duals surjective by theorem 8.5, by the five lemma from homological algebra [11] we can conclude that i_{K-D+p} is an isomorphism. The same argument on the complex formed from the last five elements of each exact sequence

$$\begin{array}{ccccccccccc}
\dots & \longrightarrow & H^0(\Omega_{D+p-K}) & \longrightarrow & \mathbb{C} & \longrightarrow & H^1(\Omega_{D-K}) & \longrightarrow & H^1(\Omega_{D+p-K}) & \longrightarrow & 0 \\
 & & \downarrow{\scriptstyle i_{K-D-p}} & & \downarrow & & \downarrow{\scriptstyle i^*_{D-K}} & & \downarrow{\scriptstyle i^*_{D+p-K}} & & \downarrow \\
\dots & \longrightarrow & H^1(\mathcal{O}_{K-D-p})^* & \longrightarrow & \mathbb{C} & \longrightarrow & H^0(\mathcal{O}_{K-D})^* & \longrightarrow & H^0(\mathcal{O}_{K-D-p})^* & \longrightarrow & 0
\end{array}$$

tells us that i_{D-K-p} is also an isomorphism. Setting $B = K - D$, we have that Serre duality holding for B and $-B$ implies Serre duality holding for $B + p$ and $-(B + p)$. We can do a completely analogous argument for $B - p$, and then we have our theorem.

9. Applications

9.1. the Degree of K and the Riemann-Hurwitz Formula.

The first thing we want to do is obtain information about the degree of the canonical divisor K. As it happens, there is a very beautiful correspondence between the canonical divisor of a surface and the Euler characteristic of the surface.

Theorem 9.1. *Deg*$(K) = 2g - 2$.

Proof. By Riemann-Roch, we have

$$\begin{aligned}
\mathrm{Deg}(K) &= \dim H^0(\mathcal{O}_K) - \dim H^1(\mathcal{O}_K) + g - 1 \\
&= \dim H^0(\Omega) - \dim H^0(\Omega_{-K}) + g - 1 \\
&= \dim H^0(\Omega) - \dim H^0(\mathcal{O}) + g - 1 \\
&= g - 1 + g - 1 \\
&= 2g - 2.
\end{aligned}$$

Where we apply Serre duality to $\dim H^0(\mathcal{O}_K)$ and then the the mappings $\mathcal{O}_{D+K} = \Omega_D$.

□

Using this fact, we get this aesthetically pleasing version of the Riemann-Roch theorem:

$$\dim H^0(\mathcal{O}_D) - H^0(\mathcal{O}_{K-D}) = \frac{\mathrm{Deg}(D)}{2} - \frac{\mathrm{Deg}(K-D)}{2}.$$

[11] See [5].

The degree of K is extremely important. In fact, if we had chosen to prove it in a manner not requiring Serre duality, could immediately derive a weak version of Serre duality from it. We do not get an explicit isomorphism, but we can prove that the groups in question have the same dimension. We only require theorem 8.5: that i_D is injective.

Theorem 9.2. *Weak version of Serre duality:* $\dim H^0(\Omega_D) = \dim H^1(\mathcal{O}_{-D})$.

Proof. Since i_D is injective, we have $\dim H^0(\Omega_D) \leq \dim H^1(\mathcal{O}_{-D})$, so $\dim H^0(\mathcal{O}_{D+K}) \leq \dim H^1(\mathcal{O}_{-D})$. By Riemann-Roch, we know that

$$\dim H^0(\mathcal{O}_{D+K}) - \dim H^1(\mathcal{O}_{D+K}) = 1 - g + \mathrm{Deg}(D+K)$$

and

$$\dim H^0(\mathcal{O}_{-D}) - \dim H^1(\mathcal{O}_{-D}) = 1 - g + \mathrm{Deg}(-D)$$

Adding these together, we get

$$\dim H^0(\mathcal{O}_{D+K}) - \dim H^1(\mathcal{O}_{-D}) + \dim H^0(\mathcal{O}_{-D}) - \dim H^1(\mathcal{O}_{D+K}) = \mathrm{Deg}(K) + 2 - 2g = 0.$$

But since by above both $\dim H^0(\mathcal{O}_{D+K}) - \dim H^1(\mathcal{O}_{-D})$ and $\dim H^0(\mathcal{O}_{-D}) - \dim H^1(\mathcal{O}_{D+K})$ must be nonpositive, they both must be zero, so we get $\dim H^0(\mathcal{O}_{D+K}) - \dim H^1(\mathcal{O}_{-D}) = 0$ as desired. □

We can use this degree to investigate the degree and branching of a function between any two Riemann surfaces. We define the valency $v(f,x)$ of f at a point x to be the multiplicity with which f takes the value $f(x)$ at the point x. In other words, in local coordinates f looks like z^k where $k = v(f,x)$. We then define $B(f) = \sum_{x \in X} v(f,x) - 1$. at a typical point, there will be zero branching. If we are on a compact Riemann surface, it is elementary Riemann surface theory that there can only be a finite number of branch points, and so $B(f)$ is finite.

This next formula relates the genus two Riemann surfaces with the degree and branching of a mapping between them.

Theorem 9.3. *The Riemann-Hurwitz Formula: If* $f : X \to Y$ *is a analytic map between Riemann surfaces, then* $g_X - 1 = \deg(f)(g_Y - 1) + \frac{1}{2}B(f)$.

Proof. We can pull back a differential form ω from Y to X. from above, we know $\mathrm{Deg}((\omega)) = 2g_Y - 2$ and $\mathrm{Deg}((f^*\omega)) = 2g_X - 2$.

We select $x \in X$, designate $f(x) = y$ and pick a coordinate chart (U, z) around x and (V, w) around y around which $f = w = z^k$ where $k = v(f,x)$. Then on V, we set $\omega = \phi(w)dw$. $f^*\omega = \phi(z^k)kz^{k-1}dz$. Transforming w to z^k multiplies the order of ω by k and the z^{k-1} term adds $k - 1 = b(f,x)$, so we get

$$\mathrm{ord}_x(f^*\omega) = b(f,x) + v(f,x)\mathrm{ord}_y(\omega).$$

$v(f,x)$, when taken over all $x = f^{-1}(y)$, must total to the degree of f, so we get

$$\sum_{x = f^{-1}(y)} \mathrm{ord}_x(f^*\omega) = \sum_{x = f^{-1}(y)} b(f,x) + \deg f \mathrm{ord}_y(\omega).$$

Summing this over all $y \in Y$, we get

$$\sum_{x \in X} \mathrm{ord}_x(f^*\omega) = \sum_{x \in X} b(f,x) + n \sum_{y \in Y} \mathrm{ord}_y(\omega),$$

which simplifies to $2g_X - 2 = \deg f(2g_Y - 2) + B(f, x)$, or

$$g_X - 1 = \deg(f)(g_Y - 1) + \frac{1}{2}B(f).$$

□

Example 9.4. Let's look at the Riemann surface generated by the equation $\sqrt[n]{(1 - z^n)}$, where $z : \mathbb{P}^1 \to \mathbb{P}^1$ is the identity on the Riemann sphere. We want to determine the genus of this surface. Rewriting the equation as $\prod_{i=1}^{n}(z = \zeta_i)^{1/n}$, it is clear there are n branch points, one at each of the nth roots of unity, and each one has valency n, and so branching number $n - 1$. since $\deg(f) = n$, we get

$$\begin{aligned} g_X - 1 &= -n + \frac{n(n-1)}{2} \\ g_X &= -(n-1) + \frac{n(n-1)}{2} \\ &= \frac{(n-2)(n-1)}{2}. \end{aligned}$$

This result is a special case of the degree-genus formula.

9.2. **Applications to Riemann Surfaces.** Serre duality and Riemann-Roch are beautiful results in and of themselves, but where they really shine is their usefulness in computations. We give a few examples.

Theorem 9.5. *Any Riemann surface X of genus g is at most a $g + 1$-branched cover of the Riemann-sphere.*

Proof. This statement is equivalent to there being a non-constant f of degree at most $g + 1$ on X (remember that a meromorphic function is just an analytic map to the Riemann sphere.) If we consider the divisor $D = (g + 1)p$ for some $p \in X$, we have

$$\dim H^0(\mathcal{O}_D) \geq 1 - g + \mathrm{Deg}(D) = 2.$$

So there must be a non-constant element f which is a meromorphic function with only one pole of degree at most $g + 1$. □

In particular, since a degree one covering map is an isomorphism, we have proven that every genus zero algebraic curve is isomorphic to the Riemann sphere.

Theorem 9.6. *Every genus one Riemann surface is isomorphic to a complex torus.*

Proof. Call our surface X. A complex torus is a quotient $\mathbb{C}$ by the group actions $\langle z \to z + \omega_1 \rangle, \langle z \to z + \omega_2 \rangle$ where $\omega_1 \omega_2$ are $\mathbb{R}$-linear independent. Since our surface when seen as a real 2-manifold has genus one and is therefore topologically a torus, we know it has a universal cover $\hat{X} = \mathbb{R}^2$, with $\pi : \hat{X} \to X$ the covering map. All we have to show is that $\hat{X} = \mathbb{C}$ as a Riemann surface and we have our proof, since the Deck transformations must take the form of the group actions above.

We consider a meromorphic 1-form ω on X. since $\mathrm{Deg}((\omega)) = 0$ since it is a canonical divisor, we have by Riemann-Roch

$$\dim H^0(\mathcal{O}_{(\omega)}) - \dim H^0(\mathcal{O}) = 1 - g + \deg D \dim H^0(\mathcal{O}_{(\omega)}) = 1.$$

If we select $f \in H^0(\mathcal{O}_{(\omega)})$ we know $(f) \geq -(\omega)$, so $(f\omega) \geq 0$, which means $f\omega$ is a holomorphic 1-form. since $\mathrm{Deg}(f\omega) = 0$ since it is a meromorphic function multiplied by a meromorphic 1-form, we know that $(f\omega) = 0$, so it is a 1-form with no zeroes or poles. Now, consider the pullback $\pi^*(f\omega)$. fix a point $p_0 \in \hat{X}$ and define $\phi(p) = \int_{p_0}^{p} \pi^*(f\omega)$. Since $\pi^*(f\omega)$ is still a holomorphic 1-form, the given integral only depends on its endpoints, so is well-defined. The integral of a holomorphic 1-form is holomorphic, so ϕ is an isomorphism between $\hat{X}$ and $\mathbb{C}$ as desired. □

A curve is called hyperelliptic if it can be seen as a double-cover of the Riemann sphere.

Theorem 9.7. *Any genus two Riemann surface is hyperellipic.*

Proof. Clearly a curve is hyperelliptic if it admits a degree two meromorphic function. For a genus two curve, we have

$$\begin{aligned} \dim H^0(\Omega) &= \dim H^0(\mathcal{O}_K) \\ &= 1 - g + \mathrm{Deg}(K) + \dim H^0(\mathcal{O})) \\ &= 1 - 2 + 2 + 1 = 2 \end{aligned}$$

so there is a non-constant holomorphic 1-form ω. Now say $K = \omega$. This means $K \geq 0$, and since as above we know $\dim H^0(\mathcal{O}_K) = 2$, there is a nonconstant meromorphic function $f \in H^0(\mathcal{O}_K)$. f can have a pole of order at most 2, and cannot have solely poles of degree 1 or it would be isomorphic to the Riemann sphere, so we have found our double-cover. □

10. Conclusion

The Riemann-Roch theorem and Serre duality are very powerful tools in complex and algebraic geometry. Additional structures such as the Jacobi variety and the Picard group are natural things to investigate after learning this material, and provide further geometric tools to work with. The theory of vector bundles is a very useful persepctive that we regrettably could not discuss in this paper. Learning the vector bundle perspective on this material lets one view any of the sheave $\mathcal{O}_D$ as the sheaf of sections of some vector bundle. The theory of sheaves and sheaf cohomology is very important in algebraic geometry, but its most general form is more complicated than the version we have presented here. For more on this, see [4]. The Riemann-Roch theorem and Serre duality also have powerful generalizations. For higher dimensional complex manifolds, the Riemann-Roch theorem still appears as a statement about the Euler characteristic of a sheaf, but more powerful machinery than we develop here is required to state any of its generalizations. We see the true "duality" provided by Serre duality when viewing its statement on a complex manifold of dimension n: $H^i(\mathcal{O}_D) \cong H^{n-i}(\mathcal{O}_{K-D})$. Studying this material from either a complex geometric or algebraic perspective reveals powerful generalizations of all of the material we have discussed, but the extremely well-behaved theory of Riemann surfaces is an excellent place to first see these ideas in their more elementary form.

11. Acknowledgments

It is a pleasure to thank my mentor, Sergei Sagatov for his excellent guidance and assistance over the summer. I would like to thank Peter May for hosting the REU and making this paper possible.

References

[1] Raoul Bott and Loring W. Tu. Differential Forms in Algebraic Topology. Springer, 1982.
[2] Simon Donaldson. Riemann Surfaces. Oxford University Press. 2011
[3] Otto Forster. Lectures in Riemann Surfaces Springer, 1981.
[4] Robin Hartshorne. Algebraic Geometry. Springer 1977
[5] Williams S. Massey.f A Basic Course in Algebraic Topology. Springer, 1991.
[6] Rick Miranda. Algebraic Curves and Riemann Surfaces American Mathematical Society, 1995.
[7] Wilhelm Schlag. A Course in Complex Analysis and Riemann Surfaces. American Mathematical Society, 2014.