Riemann-Roch theorem: Foundations and Applications

Mentees: Souleymane Diallo, Zakaria Chakir
Mentor : Rudy Ariaz

June 2025

1 Introduction

Algebraic geometry is one of the oldest subjects in mathematics, dating back to the Greeks' attempts to solve algebraic equations using geometry. But it was only in the 19th and 20th centuries that the field really began to adopt its more modern approach. From the Italian school's study of the geometry of solution sets to polynomial equations, and continuing with the French school's reformulation of the subject in the 1950s and 60s, the field has both sparked and benefited from developments in a wide variety of mathematical disciplines. In particular, going against its name, the subject of complex algebraic geometry, which for a long time was the only "satisfactory" theory, and whose results geometers attempted (and still attempt) to generalize to other base fields, greatly benefited from the use of analytic methods. Specifically, in the context of curves, much of what we know is due to the 19th-century contributions of Riemann and his theory of Riemann surfaces, which are defined analytically as one-dimensional complex manifolds. Among the central tools and results Riemann introduced is the celebrated Riemann–Roch theorem, relating the dimension of the vector space of meromorphic functions with prescribed poles and zeroes on a Riemann surface to topological invariants of the manifold. The importance of this theorem cannot be overstated. Since its initial formulation in the 1850s, it has inspired generalizations such as the Grothendieck-Hirzebruch–Riemann–Roch theorem and the Atiyah–Singer index theorem. These, in turn, have yielded deep results across algebraic geometry, number theory (e.g., in the arithmetic of elliptic curves) and even theoretical physics (e.g., in the study of Calabi–Yau manifolds in string theory and quantum field theory).

In this report, after introducing some topological and complex-analytic preliminaries, we define Riemann surfaces and present several common examples. We then introduce holomorphic/meromorphic functions and differential forms on Riemann surfaces, as well as holomorphic maps between Riemann surfaces. Finally, we present the notion of a divisor on a compact Riemann surface and the associated spaces of meromorphic functions and forms, which leads us to the statement and proof of the Riemann-Roch theorem. We conclude by discussing two applications of the Riemann-Roch theorem, showcasing its power and importance, namely, the algebraicity of compact Riemann surfaces and the existence of a Weierstrass equation for elliptic curves over the complex numbers.

2 Acknowldgements

We would like to give a special thanks to our mentor, Rudy Ariaz, for taking the time to guide us through this rich and intricate theory and to the organizers of the DRP program for giving us the opportunity to learn more about this fundamental and beautiful result of algebraic geometry.

Contents

3 Preliminaries

3.1 Basic topology

Before beginning our study of Riemann surfaces, we give a brief overview of some of the recurring notions from topology and complex analysis, upon which rests the theory of Riemann surfaces.

Definition 3.1. A *topological space* is a set X together with a collection of so-called *open subsets* $\mathfrak{T} \subseteq \mathcal{P}(X)$ called a *topology*, satisfying the following:

(i) $\emptyset, X \in \mathfrak{T}$

(ii) For any family of open sets $\{U_i\}_{i\in I}$, $\bigcup_{i\in I} U_i \in \mathfrak{T}$.

(iii) For any *finite* family $\{U_i\}_{i=1}^n$ of open sets, $\bigcap_{i\in I} U_i \in \mathfrak{T}$.

A subset $A \subset X$ of a topological space is said to be *closed* if its complement is open, i.e, if $X \backslash A \in \mathfrak{T}$. Closed sets satisfy the following properties:

(iv) For any family $\{V_i\}_{i\in I}$ of closed sets, $\bigcap_{i\in I} V_i$ is closed.

(v) For any *finite* family $\{V_i\}_{i\in I}$ of closed sets, $\bigcup_{i\in I} V_i$ is closed.

We note that, unlike their usual meanings in everyday language, the topological notions of "open" and "closed" are not opposites. In fact, an arbitrary subset of a topological space can be both open and closed as well as neither. Since the open sets of a topology completely determine the closed sets, and vice versa, a topology on a set X can also be defined by specifying a collection $\mathfrak{C}$ of closed sets which satisfies conditions (i), (iv) and (v).

We list some common properties a topological space may satisfy.

Definition 3.2. Let $(X, \mathfrak{T})$ be a topological space

1. A *base* for the topology $\mathfrak{T}$ is a collection of open sets $\mathfrak{B} \subset \mathfrak{T}$ such that every open set may be expressed as a union of elements of $\mathfrak{B}$. A topological space is *second-countable* if its topology admits a countable base.

2. X is *Hausdorff* or *separated* if for any pair of points $p, q \in X$, there exists nonempty open sets U_p, U_q containing p and q respectively such that $U_p \cap U_q = \emptyset$.

3. X is *connected* if it cannot be written as a union of open sets $U_1 \cup U_2$ such that $U_1 \cap U_2 = \emptyset$.

Definition 3.3. An *open cover* of a subset $A \subset X$ is a collection of open sets $\{U_i\}_{i\in I}$ such that $A \subset \bigcup_{i\in I} U_i$

Definition 3.4. A subset $A \subset X$ is *compact* if for any open cover $\{U_i\}_{i\in I}$ of A, there exists a finite index set $I_n \subset I$ such that $\{U_i\}_{i\in I_n}$ is a cover of A.

Given a subset A of a topological space X, it is natural to consider A as a topological space of its own by giving it a topology induced from the one on X. Similar considerations can be made for $X/\sim$, the quotient of X under an equivalence relation $\sim$.

Definition 3.5. Let A be a subset of a topological space $(X, \mathfrak{T})$. The *subspace topology* on A is the collection of sets $\{U \cap A | U \in \mathfrak{T}\}$.

Definition 3.6. Let $\sim$ be an equivalence relation on X and consider the set of equivalence classes $X/\sim$. If we denote by $\pi : X \to X/\sim$ the projection map, then the *quotient topology* is the collection of sets $\{U \subset X/\sim | \pi^{-1}(U) \in \mathfrak{T}\}$.

The relevant maps between topological spaces are the *continuous maps*. In essence, these are the maps that preserve the topological structure.

Definition 3.7. A map $f : X \to Y$ between two topological spaces is *continuous at a point* $p \in X$ if for every open neighbourhood V of $f(p)$, there exists an open neighbourhood U of p such that $f(U) \subset V$. If f is continuous at every point of X, we simply say f is globally continuous or continuous.

Lemma 3.8. *A map $f : X \to Y$ is continuous if and only if, for every open set $V \subset Y$, the preimage $f^{-1}(V)$ is open in X. Equivalently, f is continuous if and only if the preimage of every closed set is closed.*

As mentioned, continuous maps behave nicely with regard to topological properties of spaces, namely:

Proposition 3.9. *Let $A \subset X$ and $f : A \to Y$ a continuous map, then*

- *If A is connected then $f(A)$ is connected.*
- *If A is compact then $f(A)$ is compact.*

The definition of a topological isomorphism is the natural one

Definition 3.10. Let $f : X \to Y$ be a continuous map between topological spaces. We say f is a *homeomorphism* or a *topological isomorphism* if it is bijective and its inverse $f^{-1} : Y \to X$ is continuous.

By definition, the continuity of a map is completely determined by the topologies on its domain and codomain. Conversely, given a function $f : X \to Y$ from a set X into a topological space Y, one can equip X with a topology $\mathfrak{T}_f$, defined as the coarsest[1] topology that makes f continuous.

Definition 3.11. Let X be a set, $(Y_i)_{i\in I}$ a family of topological spaces and $\mathcal{F} = (f_i)_{i\in I}$ be a family of maps $f_i : X \to Y_i$. We define the *weak topology on X induced by $\mathcal{F}$* as the topology induced by the set

$$\{f_i^{-1}(U) | U \subset Y_i \text{ is open}, i \in I\},$$

that is, the set of all finite intersections of arbitrary unions of sets of the form given above.

There is also the notion of gluing topological spaces along homeomorphic open subspaces. We give the definition for the simple case of the glueing of two topological spaces along a common subspace; the general definition for an arbitrary family of topological spaces $(X_i)_{i\in I}$ is a natural generalization which only differs by requiring additional compatibility conditions between the pairwise homeomorphic open subsets.

Definition 3.12. Let X_1, X_2 be topological spaces, $U_1 \subset X_1$ and $U_2 \subset X_2$ be open subsets and $\phi : U_1 \to U_2$ be a homeomorphism. We define the *glueing of X_1 and X_2 along U_1 and U_2* as $(X_1 \coprod X_2)\backslash \sim$, the quotient of the disjoint union $X_1 \coprod X_2$ under the equivalence induced by $x_1 \sim x_2$ if $\phi(x_1) = x_2$.

3.2 Functions of a Complex Variable

We give a brief overview of some standard definitions and results of complex analysis which will also be crucial to the rest of this paper. For reasons of time and scope, the results will be stated without proof; readers wishing for a more detailed treatment of functions of a complex variable may refer to the first few chapters of [SS03] or [NN01].

We denote by $\mathbb{C}$ the complex plane, i.e., the set $\{a + bi \mid a, b \in \mathbb{R}\}$ where $i = \sqrt{-1}$. We view $\mathbb{C}$ as a normed vector space, the norm of a complex number z being its modulus $|z| = \sqrt{z\bar{z}}$. We equip $\mathbb{C}$ with the metric topology induced from $|\cdot|$: a base for the topology of $\mathbb{C}$ is then given by sets of the form $B_p(r) = \{z \in \mathbb{C} \mid |z - p| < r\}$ for all $p \in \mathbb{C}$.

In the following, we let U denote a connected open subset of $\mathbb{C}$.

[1] A topology $\mathfrak{T}$ is said to be *weaker* or *coarser* than another $\mathfrak{T}'$ if $\mathfrak{T} \subset \mathfrak{T}'$; conversely we say that $\mathfrak{T}'$ is *finer* or *stronger* than $\mathfrak{T}$.

Definition 3.13. A function $f : U \to \mathbb{C}$ is *holomorphic* or *complex differentiable* at a point $z_0 \in U$ if the following limit exists:

$$\lim_{z\to z_0} \frac{f(z) - f(z_0)}{z - z_0}.$$

If the condition holds for all $z \in U$ we say that f is *holomorphic on* U, or simply *holomorphic* if $U = \mathbb{C}$.

We note that, as we may separate a complex number $z = x + iy$ into its real and imaginary parts x and y, we may express a complex valued function $f : U \to \mathbb{C}$ in terms of 2 real valued functions $u(x, y)$ and $v(x, y)$ such that

$$f(x, y) = u(x, y) + iv(x, y).$$

If f is holomorphic, then u and v will be real continuously differentiable (C^1); however, the converse is not always true[2]. A necessary and sufficient condition for a pair of continuously differentiable (C^1) functions (u, v) to correspond to a holomorphic function $f(x, y) = u(x, y) + iv(x, y)$ is the satisfaction of the so-called *Cauchy-Riemann equations*:

$$\frac{\partial u}{\partial x} = \frac{\partial v}{\partial y}, \quad \frac{\partial u}{\partial y} = -\frac{\partial v}{\partial x}.$$

Holomorphicity is a much stronger property than real-differentiability, and it implies stronger consequences, a few of which we list here:

1. If f is holomorphic at a point $z_0 \in \mathbb{C}$, then it is *infinitely differentiable* at z_0.
2. If f is holomorphic at z_0, then it is *analytic* at z_0; that is, there exists a power series $\sum_{m=0}^{\infty} a_m(z - z_0)^m$ and some $r > 0$ such that $f_{|B_{z_0}(r)}(z) = \sum_{m=0}^{\infty} a_m(z - z_0)^m$.
3. (*Cauchy's theorem*) If f is holomorphic on U and γ is a closed path whose interior is contained U, then
$$\int_\gamma f(z)dz = 0.$$[3]

While holomorphic functions are the primary object of study of complex analysis, it is useful and often necessary to enlarge the class of functions considered by allowing functions which, at certain points, blow up to infinity. However, for those functions to retain certain analytic properties, we must remain selective about the types of singularities we allow.

Definition 3.14. Let f be a holomorphic function defined in a punctured neighbourhood U of a point $z_0 \in \mathbb{C}$ (f is defined on $U\backslash\{z_0\}$):

1. If $|f|$ is bounded on $U\backslash\{z_0\}$, we say f has a *removable singularity* at z_0.
2. If $\lim_{z\to z_0} |f| = \infty$, then f has a *pole* at z_0.
3. In any other case, we say f has an *essential singularity* at z_0.

Definition 3.15. A function $f : U \setminus \{z_0\} \to \mathbb{C}$ is said to be *meromorphic at* $z_0 \in U$ if either f has a removable singularity at z_0, in which case it can be extended to a holomorphic function on U, or if f has a pole at p.

The set of zeroes and poles of a nonzero meromorphic function is discrete. Therefore, while a meromorphic function f may not be well-defined at a point z_0, it is perfectly well-defined and holomorphic in a punctured neighbourhood of z_0. Additionally, the pole at z_0 is, in a way, "well-behaved", so that we can expand the function at z_0 into a power series which contains negative exponents.

[2]Consider the complex conjugation function taking $z = x + iy$ to $\bar{z} = x - iy$.

[3]If $\gamma(t) : [0, 1] \to U$ is a continuous path, we define the integral $\int_\gamma f(z)\,dz$ to be $\int_0^1 f(\gamma(t))\gamma'(t)dt$, the latter being the "usual" Riemann integral on the unit interval.

Definition 3.16. A *Laurent series* centred at a point $z_0 \in \mathbb{C}$ is a formal series of the form

$$\sum_{m \in \mathbb{Z}}^{\infty} a_m (z - z_0)^m$$

where the coefficients a_m are complex numbers.

Proposition 3.17. *If $f : U \to \mathbb{C}$ is meromorphic at a point $z_0 \in U$, then in some punctured neighbourhood of z_0, f may expressed as a Laurent series*

$$f(z) = \sum_{m=m_0}^{\infty} a_m (z - z_0)^m.$$

where $m_0 \in \mathbf{Z}$.

Definition 3.18. Let $f : U \to \mathbb{C}$ be meromorphic at a point $z_0 \in U$. The *order* of f at z_0 denoted $\operatorname{ord}_{z_0}(f)$ is the smallest unique integer n for which $(z - z_0)^n f(z)$ is holomorphic and nonzero at z_0. In other words, it is the index of the first nonzero term in the Laurent series expansion of f at z_0.

The following properties of the order function follow from basic computations of the Laurent series of the associated operations on the meromorphic functions.

Proposition 3.19. *Let f, g be meromorphic functions and let $f(z_0) \neq 0$.*

1. $\operatorname{ord}_{z_0}(fg) = \operatorname{ord}_{z_0}(f) + \operatorname{ord}_{z_0}(g)$
2. $\operatorname{ord}_{z_0}(1/f) = -\operatorname{ord}_{z_0}(f)$
3. $\operatorname{ord}_{z_0}(f \pm g) \geq \min\{\operatorname{ord}_{z_0}(f), \operatorname{ord}_{z_0}(g)\}$

One key difference between holomorphic and meromorphic functions lies in their behaviour under integration along closed loops in the complex plane. For holomorphic functions, the integral around any closed contour within a simply connected domain[4] always vanishes, by Cauchy's theorem. In contrast, meromorphic functions may have isolated singularities, and their integrals around closed curves enclosing such points generally yield nonzero values.

Definition 3.20. Let $f : U \to \mathbb{C}$ be a meromorphic function with Laurent series expansion $f(z) = \sum_{m=m_0}^{\infty} a_m (z - z_0)^m$ at a point $z_0 \in U$. The *residue* of the function at z_0 is the quantity

$$\operatorname{Res}(f, z_0) = a_{-1}.$$

This quantity is quite important due to the following theorem.

Theorem 3.21 (*Cauchy Residue formula*)**.** *Let $U \subset \mathbb{C}$ be a simply connected open subset, and let f be a function holomorphic on U except for a finite number of isolated singularities $w_1, w_2, \ldots, w_n \in U$. Let γ be a well-defined, positively oriented, simple closed contour in U that encloses isolated singularities $w_1, w_2, \ldots, w_k \in U$. Then:*

$$\int_{\gamma} f(z)\, dz = 2\pi i \sum_{j=1}^{k} \operatorname{Res}(f, a_j),$$

where $\operatorname{Res}(f, a_k)$ is as defined above.

[4] As, sadly, we will not discuss of homotopies and fundamental groups, one may take as a definition of a simply connected domain, that of an open subset of $\mathbb{C}$ which is homeomorphic to the open unit ball.

4 Definition and important examples of Riemann surfaces

In essence, the theory of Riemann surfaces aims to utilize the powerful tools of complex analysis to study properties of spaces which locally "look" like open subsets of $\mathbb{C}$. To do so, we must first define precisely what it means for a topological space to locally "look" like $\mathbb{C}$ and to that effect, we use complex charts.

Definition 4.1. A *complex chart* or simply a *chart* on a topological space X is a pair (U, ϕ) where $U \subset X$ is open and $\phi : U \to V$ a homeomorphism onto an open subset $V \subset C$; we say ϕ is centred at $p \in X$, if $\phi(p) = 0$.

We can think of a chart ϕ as a continuous invertible deformation of U into an open subset of the complex plane, attaching to each point $p \in U$ a unique local complex coordinate $z \in \mathbb{C}$.

Example 4.2. Let $X = \mathbb{R}^2$ and for every open set $U \subset \mathbb{R}^2$, let $\phi_U : U \to V, (x, y) \mapsto x + iy$. These charts map each point in $\mathbb{R}^2$ to its corresponding vector in $\mathbb{C}$. $\mathbb{R}^2$ equipped with these charts is called the *complex plane.*

Proposition 4.3. *Given a topological space X, a complex chart $\phi : U \to V$, an open set $U_1 \subset U$ and a holomorphic bijection $\psi : V \to W$ where $V, W \subset \mathbb{C}$ is open.*

1. *The restriction $\phi|_{U_1} : U_1 \to \phi(U_1)$ is a chart and called a sub-chart of ϕ*
2. *The composition $\psi \circ \phi : U \to W$ is a chart and called a change of coordinates*

Definition 4.4. Two complex charts on a topological space X, $\phi_1 : U_1 \to V_1$ and $\phi_2 : U_2 \to V_2$, where $U_1, U_2 \subset X$ and $V_1, V_2 \subset \mathbb{C}$, are *compatible* if either

1. $U_1 \cap U_2 = \emptyset$ or
2. $\phi_2 \circ \phi_1^{-1} : \phi_1(U_1 \cap U_2) \to \phi_2(U_1 \cap U_2)$ is *holomorphic* on $\mathbb{C}$.

The function $T = \phi_2 \circ \phi_1^{-1}$ displayed in the second condition is called a *transition function* between both charts; it is always a bijection. The compatibility condition for charts is essential to the development of the entire theory. As one can guess, we will later declare that a function $f : U \to \mathbb{C}$ is holomorphic if, given a chart (U, ϕ), the composition $f \circ \phi^{-1} : \phi(U) \to \mathbb{C}$ is holomorphic. Forgoing the compatibility condition, it may happen that under another chart (U, ψ), $f \circ \psi^{-1} : \psi(U) \to \mathbb{C}$ is not. If we wish for holomorphicity to be an intrinsic property of f, independent of the choice of chart on U, we must impose the compatibility condition.

Definition 4.5. A *complex atlas* $\mathcal{A}$ on X is a collection of pairwise compatible complex charts (ϕ_i, U_i) such that $X = \bigcup_{i \in I} U_i$.

The set of all atlases on a X is partially ordered under inclusion. Thus, given an atlas $\mathcal{A}$, we may consider the maximal atlas $\mathcal{A}'$ containing $\mathcal{A}$. A maximal atlas $\mathcal{A}'$ on X is called a complex structure.

We are now ready to give the definition of a Riemann surface.

Definition 4.6. A *Riemann surface* is a second countable, connected, Hausdorff topological space X with a complex structure.

We now present a few examples of important Riemann surfaces which we will frequently revisit throughout this paper.

4.1 The Riemann Sphere

The most well-known Riemann surface, after the complex plane itself, is the Riemann Sphere, denoted by $\mathbb{C}_\infty$, which consists of the usual complex plane with an added point at infinity. Concretely $\mathbb{C}_\infty$ consists of the standard 2-sphere, $S^2 \subset \mathbb{R}$ with charts given by stereographic projection from either the south and north pole [5], onto the plane $\{(x, y, z) \mid z = 0\}$ which is

[5] This is one of the two complex charts of from the north pole: $Proj(p) = u + iv = \frac{x}{1-z} + i\frac{y}{1-z}$ $\forall p \in \mathbb{S}^2 - \{(1, 0, 0)\}$

naturally identified with $\mathbb{C}$. The change of coordinate function between the two charts is given by $T(z) = (\phi_2 \circ \phi_1^{-1})(z) = \frac{1}{z}$.

4.2 Smooth Projective Curves

Let $\mathbb{P}^n$ denote n-dimensional complex projective space, i.e., the set of all one-dimensional linear subspaces of $\mathbb{C}^{n+1}$. Given $x = (x_0, x_1, \ldots, x_n) \in \mathbb{C}^{n+1} \setminus \{0\}$, we denote by $[x_0 : \cdots : x_n]$, the projective point representing the span of x. This representation is unique up to multiplication by a nonzero complex number, as a non-zero scalar multiple of p generates the same span so that $[x_0 : \cdots : x_n] = \{(\lambda x_0, \ldots, \lambda x_n) \mid \lambda \in \mathbb{C}^*\}$. Thus, we can represent projective space as

$$\mathbb{P}^n = (\ \mathbb{C}^{n+1} - \{0\}) \setminus \mathbb{C}^*.$$

We can cover $\mathbb{P}^n$ with $n+1$ open sets $U_0, \ldots, U_n$ where $U_i = \{[x_0 : \cdots : x_n] \mid x_i \neq 0\}$. Then, for a given U_i, the map $\phi_i : U_i \to \mathbb{C}^n$ mapping $[x_i, \ldots, x_k] \mapsto (\frac{x_j}{x_i}, \ldots, \frac{x_k}{x_i})$ gives a bijection between U_i and $\mathbb{C}^n$. If we equip $\mathbb{P}^n$ with the weak topology induced from the collection of maps (ϕ_i), these maps become homeomorphisms. Moreover, if we consider the closed unit polydisc in $\mathbb{C}^n$ defined as

$$D = \{(z_1, \ldots, z_n) \in \mathbb{C}^n \mid |z_i| \leq 1 \forall i \in \{1, \ldots, n\}\}$$

we find that $\mathbb{P}^n = \bigcup_i \phi_i^{-1}(D)$. As D is compact and ϕ_i^{-1} is continuous, we get that $\mathbb{P}^n$ is compact.

Now, we present the actual definition of a smooth projective curve.

Definition 4.7. Let X be a Riemann surface, which is a subset of a projective space $\mathbb{P}^n$. We say that X is holomorphically embedded[6] in $\mathbb{P}^n$ if for every point p on X there is a homogeneous coordinate z_j; such that:

- $z_j \neq 0$ at p;
- for every k, the ratio z_k/z_j is a holomorphic function on X near p;
- there is a homogeneous coordinate z_i; such that the ratio z_i/z_j is a local coordinate on X near p.

A Riemann surface which is holomorphically embedded in projective space is called a *smooth projective curve*. In the special case where $X \subset \mathbb{P}^2$, we say that X is *smooth projective plane curve*.

Definition 4.8. A polynomial $F(x_0, \ldots, x_n) \in \mathbb{C}[x_0, \ldots, x_n]$ is *homogeneous of degree d* if, for all $\lambda \in \mathbb{C}^*$, $F(\lambda x_0, \ldots, \lambda x_n) = \lambda^d F(x_0, \ldots, x_n)$.

While the evaluation of a homogeneous polynomial F at an arbitrary projective point $p \in \mathbb{P}^n$ is not well-defined, the homogeneity condition implies that the locus of F is a well-defined subset of $\mathbb{P}^n$.

Definition 4.9. A homogeneous polynomial $F \in \mathbb{C}[x_0, \ldots, x_n]$ is *singular* at a point p if

$$F(p) = \frac{\partial F}{\partial x_0}(p) = \frac{\partial F}{\partial x_1}(p) = \cdots = \frac{\partial F}{\partial x_n}(p) = 0;$$

we also say that p is a singular point of F. F is *non-singular* if there exists no such point $p \in \mathbb{P}^n$.

In general, under the topology induced from $\mathbb{P}^n$, the vanishing set $X_F \subset \mathbb{P}^n$ of a non-singular polynomial is *complex manifold* of dimension $n-1$, which is a higher-dimensional analogue of a Riemann surface[7]. In particular, if $F \in \mathbb{C}[x_0, x_1, x_2]$, so that $X_F \subset \mathbb{P}^2$, we will see that X_F is a Riemann surface.

[6] Holomorphic/Smooth embeddings are defined in a much more general context, and later on we will provide a more general definition in Definition 7.28, but this definition suffices to define a projective curve.

[7] The charts on a complex manifold of dimension n are local homeomorphims to $\mathbb{C}^n$, rather than to $\mathbb{C}$.

Example 4.10. Consider the cubic curve defined by the equation $F(x_0, x_1, x_2) = x_2x_1^2 - x_0^3 - x_0^2x_2$ and let X be its set of solutions in $\mathbb{P}^2$. Notice that at the point $p = [0 : 0 : 1]$ we have that

$$\frac{\partial F}{\partial x_0}(p) = 3x_0^2 + 2x_0x_2 = 0, \quad \frac{\partial F}{\partial x_1}(p) = 2x_2x_1 = 0, \quad \frac{\partial F}{\partial x_2}(p) = x_1^2 + x_0 = 0$$

p is therefore a singular point. By intersecting $X \cap U_2$, we recover the graph of the affine plane curve $x_1^2 = x_0^3 + x_0^2$ in $\mathbb{C}^2$. While an incomplete representation, the restriction of the curve to $\mathbb{R}^2$ allows us to visualize the singularity at the origin. Looking at Figure 1, we see that around any neighbourhood of $(0, 0)$, the projection of the curve onto either axis is not one-to-one.

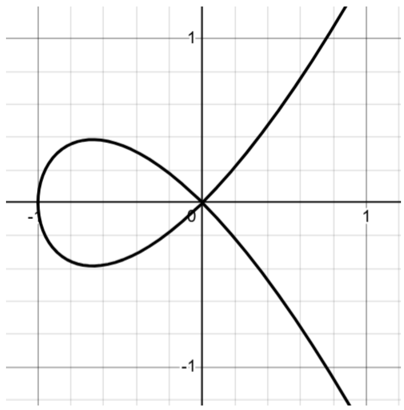

Figure 1: Graph of the elliptic curve $x_1^2 - x_0^3 - x_0^2$ in $\mathbb{R}^2$

As was mentioned in the above example, given a smooth projective curve $X \subset \mathbb{P}^2$, for every i, the intersection $X_i = X \cap U_i$ is a smooth affine curve on $\mathbb{C}^2$. Moreover, we have that X is smooth if and only if X_i is smooth for every i. Using the *implicit function theorem* in $\mathbb{C}^2$, locally around each point of X_i equipped with the subspace topology of $\mathbb{C}^2$, we obtain a homeomorphism into $\mathbb{C}$ by projecting along one of the two (complex) axes. We can therefore define the charts on X by first projecting X onto some X_i and then taking local projections from the affine curve to $\mathbb{C}$. We therefore have the following.

Proposition 4.11. *Let $F \in \mathbb{C}[x_0, x_1, x_2]$ be a smooth homogeneous polynomial. Then, its locus of zeros $X_F \subset \mathbb{P}^2$ is a compact Riemann surface.*

The fact that X is connected is a consequence of a theorem of algebraic geometry, which states that every smooth homogeneous polynomial is irreducible. X being Hausdorff is a result of the affine curves $X_i := X_F \cap U_i$ being Hausdorff with respect to the subspace topology on $\mathbb{C}^2$, inherited by ϕ_i. The charts on X are constructed by applying the implicit function theorem, locally at each point; this is possible as X is smooth. The curve X can then be represented as the glueing [3.12] of the affine curves X_i along their pairwise intersections. This allows us to transport the charts on the $X_i's$ to X according to the following diagram

$$\begin{tikzcd} & X_F & \\ X_0 \arrow[ur, "\phi_0"] \arrow[dr, Rightarrow] & X_1 \arrow[u, "\phi_1"] \arrow[d, shift left] \arrow[d, shift right] & X_2 \arrow[ul, "\phi_2"'] \arrow[dl, Rightarrow] \\ & \mathbb{C} & \end{tikzcd}$$

where the maps ϕ_i are the natural inclusions, and double arrows represent the 2 possible projections onto either axis of $\mathbb{C}^2$. We therefore have a copious amount of compatibility conditions to verify. We will only show one case, as the others are proven by an analogous argument. We therefore let $p \in X_0 \cap X_1$ such that $\phi_0(p) = [1 : \frac{x_1}{x_0} : \frac{x_2}{x_0}]$ and $\phi_1(p) = [\frac{x_0}{x_1}, 1, \frac{x_2}{x_1}]$ where $x_0, x_1 \neq 0$.

Suppose that in the first chart, the local coordinate is given by $z = \frac{x_1}{x_0}$ and in the second given by $w = \frac{x_2}{x_1}$ such that locally $\phi_0(p) = [1 : z : h(z)]$ and $\phi_1(p) = [g(w) : 1 : w]$ for holomorphic functions h and g. We then find that $\phi_1 \circ \phi_0^{-1}$ is given by $z \mapsto h(z)/z$, which is holomorphic as z is nonzero in a neighbourhood of p. Our charts are therefore compatible. The compactness of X_F comes from the fact that for each i, X_i is a closed subset in $\mathbb{C}^2$ which is homeomorphic to U_i and thus closed in $\mathbb{P}^2$. Since $X_F = \bigcup_i X_i$, it is a closed subset of $\mathbb{P}^2$ and thus compact.

4.3 Complex Tori

Let $\mathbb{H} = \{z \in \mathbb{C} \mid \mathrm{Im}(z) > 0\}$ denote the complex upper half plane and fix $\tau \in \mathbb{H}$. Define the lattice $\mathcal{L} = \mathbb{Z} + \mathbb{Z}\tau = \{a + b\tau \mid a, b \in \mathbb{Z}\}$ which is a subgroup of the additive group $(\mathbb{C}, +)$ and consider the quotient $X = \mathbb{C}/\mathcal{L}$. Using the canonical projection map $\pi : \mathbb{C} \to X$ we can endow X with the quotient topology such that $U \subset X$ is open $\iff$ $\pi^{-1}(U)$ is open in $\mathbb{C}$. By construction, π is continuous and since $\mathbb{C}$ is connected, so is X. Moreover, π is an open mapping as given an open set $V \subset \mathbb{C}$, $\pi^{-1}(\pi(V)) = \bigcup_{\omega \in \mathcal{L}}(\omega + V)$ which is a union of translates of V in the complex plane.

For a fixed $z \in \mathbb{C}$, we construct the following closed and "half-closed" parallelograms

$$P_z = \{z + \lambda_1 + \lambda_2\tau \mid \lambda_i \in [0, 1]\}$$

$$\tilde{P}_z = \{z + \lambda_1 + \lambda_2\tau \mid [0, 1[\}.$$

We can think of $\tilde{P}_z$ as being obtained from P_z by identifying the two pairs of opposite sides of the closed parallelogram together. Under this identification, any point $z \in \mathbb{C}$ is congruent to a unique point of $\tilde{P}_z$ under the action of $\mathcal{L}$; $\tilde{P}_z$ is thus, by definition, a fundamental domain of the action of $\mathcal{L}$ on $\mathbb{C}$. Also we note that $\pi : P_z \to X$ is a surjective open mapping and since P_z is compact, so is X.

We now define a set of complex charts on X. Since $\mathcal{L} \subset \mathbb{C}$ is discrete, there exists some $\varepsilon > 0$ such that for all $w \in \mathcal{L}$, $|w| > 2\varepsilon$. This means for a given point $p_0 \in \mathbb{C}$, the open ball $B(p_0, \varepsilon)$ does not contain a pair of $\mathcal{L}$-congruent points and thus the mapping $\pi_{|B(p_0,\varepsilon)} : B(p_0, \varepsilon) \to \pi(B(p_0, \varepsilon))$ is a homeomorphism. We can then construct a complex atlas on X by considering the collection of inverse charts

$$\phi_{p_0} : \pi(B(p_0, \varepsilon)) \to B(p_0, \varepsilon)$$

where ϕ_{p_0} is defined to be the inverse of $\pi_{|B(p_0,\varepsilon)}$ such that

$$\mathcal{A} = \{(\phi_p, \pi(B(p, \varepsilon)) \mid p \in \mathbb{C}\}.$$

To check that the charts of $\mathcal{A}$ are compatible, we take two points $p_1, p_2 \in \mathbb{C}$ and let $U = \pi(B(p_1, \varepsilon)) \cap \pi(B(p_2, \varepsilon))$. If the intersection is empty, then the charts ϕ_{p_1} and ϕ_{p_2} are trivially compatible. We then consider the transition function $T(z) : \phi_{p_1}(U) \to \mathbb{C}$, where $T(z) = \phi_{p_2}(\phi_{p_1}^{-1}(z)) = \phi_{p_2}(\pi(z))$. By construction, on its domain $\pi(T(z)) = \pi(z)$ which means that $T(z) - z = \omega(z)$ where $\omega(z) : \phi_{p_1}(U) \to \mathcal{L}$ is an element of the lattice dependent on z since both z and $T(z)$ are mapped to the same point on the torus. As $T(z)$ is continuous, $\omega(z)$ is also continuous. However, by continuity of ω, for every z, since singletons in the discrete space $\mathcal{L}$ are open, the preimage of $\{\omega(z)\}$ is open in $\phi_{p_1}(U)$, which means that ω is locally constant. Therefore, around every point of its domain $T(z) = z + \omega(z)$ and since $w(z)$ is holomorphic, we have that T is holomorphic on $\phi_{p_1}(U)$.

Altogether, $(X, \mathcal{A})$ is a compact Riemann surface of genus one. Heuristically, one can visualize this isomorphism by considering the fundamental parallelogram P_z and "gluing" the pairs of opposite sides together. Complex tori are a very important object of study, mainly because any given complex torus is isomorphic as a complex Lie group[8] to the solution set $E(\mathbb{C})$ of some elliptic curve defined over $\mathbb{C}$. The group of points of E is isomorphic to the additive group $\mathbb{C}/\mathcal{L}$.

[8] A complex lie group G is a complex manifold equipped with a group law $G \times G \to G$ which is holomorphic and an inversion $G \to G, g \to g^{-1}$ that is also holomorphic.

5 Functions and Differential forms on Riemann surfaces

5.1 Functions on Riemann surfaces

As previously mentioned, part of the motivation for studying Riemann surfaces comes from the desire to generalize results from complex analysis to more general complex manifolds. In this chapter, we will see that many of these results indeed extend to this broader setting. However, particularly when the manifolds in question are compact, their topological differences from $\mathbb{C}$ will impose significant restrictions on the kinds of global holomorphic and meromorphic functions they can admit. We begin by defining the holomorphicity and meromorphicity of functions on a Riemann surface.

Definition 5.1. Let X be a Riemann surface and $U \subset X$ an open subset. A function $f : U \to \mathbb{C}$ is *holomorphic (resp. meromorphic) at a point* $p \in U$ if there exist a chart (ϕ, U_p), with $p \in U_p$ such that the function

$$f \circ \phi^{-1} : \phi(U_p) \to \mathbb{C}$$

is holomorphic (resp. meromorphic) at $\phi(p)$. We say f is *holomorphic (resp. meromorphic) on* U if the condition holds for every point of U; When $U = X$, we simply say f is *holomorphic (resp. meromorphic)*.

While the above criterion requires the existence of a single chart at p under which f is holomorphic (resp. holomorphic), the compatibility of charts on X then implies that this condition is satisfied for every other chart containing p in its domain. Furthermore, using a local coordinate z centred at p, we may express a meromorphic function, f, as a Laurent series

$$f(z) = \sum_{m=m_0}^{\infty} a_m z^m$$

allowing us to speak of the order of a zero or a pole at p as shown in Definition 3.18.

We have the following notation for an open subset $U \subset X$ of a Riemann surface X:

$$\mathcal{O}(U) = \{f : U \to \mathbb{C} | f \text{ is holomorphic}\}$$

$$\mathcal{M}(U) = \{f : U \to \mathbb{C} | f \text{ is meromorphic}\}$$

We note that both spaces are $\mathbb{C}$-algebras and if U is connected, then $\mathcal{M}(U)$ is a field.

We now give a non-exhaustive list of classical theorems from complex analysis which hold true for Riemann surfaces. In the following, we suppose that f and g are meromorphic functions on a *connected* open subset W of a Riemann surface X.

Theorem 5.2 (Identity theorem)**.** *If $f, g : W \to \mathbb{C}$ agree on a subset $S \subset W$ which has a limit point in W, then f and g are identical on W.*

Theorem 5.3 (Discreteness of zeroes and poles)**.** *The set of poles and zeroes of f forms a discrete closed set. In particular, if X is compact, this set is finite.*

Theorem 5.4 (Maximum modulus principle)**.** *If there exists a point $p \in W$ such that $|f(x)| \leq |f(p)|$ for all $x \in W$, then f is constant on W.*

This final theorem implies the following striking result.

Corollary 5.5. *Let X be a compact Riemann surface, then every global holomorphic function $f : X \to \mathbb{C}$ is constant.*

Proof. Let $f : X \to \mathbb{C}$ be holomorphic. As X is compact and $|f|$ is continuous, it attains a global maximum at some point $p \in X$. X is also open and connected, and thus by the maximum modulus principle, f is constant on X. ☐

It is important to emphasize that this result applies only to global holomorphic functions. Locally, using coordinate charts, one can easily construct non-constant holomorphic functions on neighbourhoods of points. However, this result marks a key distinction between compact and non-compact Riemann surfaces. This alone motivates the study of meromorphic functions on compact Riemann surfaces, as they are the only "interesting" global functions available. Even then, establishing the existence of non-constant meromorphic functions on an arbitrary compact Riemann surface requires far more advanced tools than those covered in this paper.

We now provide a few examples of meromorphic functions on the most standard Riemann surfaces.

Example 5.6. Let f be a meromorphic function on the Riemann sphere $\mathbb{C}_\infty$,then f is given by a rational function $f(z) = p(z)/q(z)$ for some polynomials $p, q \in \mathbb{C}[z]$.

Proof. We first prove the converse: that every rational function $g(z)$ on $\mathbb{C}$ admits a canonical extension to a meromorphic function on the Riemann sphere $\mathbb{C}_\infty$. Since g is meromorphic on the complex plane, it admits a Laurent series expansion around 0 of the form:

$$g(z) = \sum_{m=m_0}^{\infty} a_m z^m.$$

Let $w = T(z) = 1/z$ denote the transition function near the point at infinity. Then g is meromorphic at ∞ if and only if $g \circ T^{-1}(w)$ is meromorphic at $w = 0$. However, since g and $1/z$ are rational functions, so is $g(1/z) = g \circ T^{-1}$ making it meromorphic at $w = 0$. This shows that g extends meromorphically to $\mathbb{C}_\infty$. Now let $f : \mathbb{C}_\infty \to \mathbb{C}$ be an arbitrary (non-constant) meromorphic function. Since $\mathbb{C}_\infty$ is compact[9], the function f has a finite set of zeroes and poles, denoted $\{p_i\}_{i\in I}$, in the complex plane $\mathbb{C}_\infty \backslash \{\infty\}$. We then construct the rational function

$$r(z) = \prod_i (z - p_i)^{e_i},$$

where $e_i = \mathrm{ord}_{p_i}(f)$ denotes the order of f at p_i. Since both f and r are meromorphic on $\mathbb{C}_\infty$, their ratio f/r is also meromorphic. By construction, f/r has neither zeroes nor poles on the finite complex plane, so it is holomorphic there and admits a Taylor expansion:

$$\frac{f(z)}{r(z)} = \sum_{m=0}^{\infty} a_m z^m.$$

On the other hand, because f/r is meromorphic on the whole sphere, we can also express it near ∞ using the coordinate $w = 1/z$ and the Taylor expansion above:

$$(f/r) \circ T^{-1}(w) = \sum_{m=0}^{\infty} a_m w^{-m}.$$

For this expression to define a meromorphic function at $w = 0$, the Laurent series must be finite, that is, only finitely many a_m are nonzero. This implies that f/r is a polynomial in z, given by the finitely many a_m in its Taylor series. But since f/r has no zeroes on $\mathbb{C}$, by the Fundamental Theorem of Algebra, it must in fact be constant. Hence, there exists $\alpha \in \mathbb{C}$ such that $f(z) = \alpha \cdot r(z)$. Therefore, f is a rational function, as claimed. □

Finally, we present an interesting and crucial application of the Residue Theorem and the Maximum Modulus Principle to meromorphic (and holomorphic) functions on compact Riemann surfaces.

[9]Compactness implies that infinitely many poles cluster to some point, meaning that f has a non-isolated singularity and hence is not meromorphic. A similar argument can be made for zeros using the coordinate $w = 1/z$.

Example 5.7. Let X be a compact Riemann surface and let f be some non-identically zero meromorphic function on X. Using a collection of complex charts for X, we have that the Residue theorem tells us, since X is compact, that :

$$\sum_{p \in X} \mathrm{Res}(\frac{df}{f}; p) = 0$$

By direct computation of the Laurent series at any point $p \in X$, we see that the residue of $\frac{df}{f}$ is $\mathrm{ord}_p(f)$. So, we obtain that on a compact Riemann surface

$$\sum_{p \in X} \mathrm{ord}_p(f) = 0$$

5.2 Differential forms on Riemann surfaces

After studying the basic properties of holomorphic and meromorphic functions on a Riemann surface X , it is natural to turn to differential forms on X, with the aim of developing an integration theory analogous to that on $\mathbb{C}$. Differential forms also play a central role in the formulation and proof of the Riemann-Roch theorem and Serre duality, which relate the dimensions of spaces of functions on X to those of certain spaces of differential forms. We begin by recalling the definition of holomorphic and meromorphic forms on $\mathbb{C}$ before transporting it to a general Riemann surface.

Definition 5.8. A holomorphic (resp. meromorphic) differential form of degree one or simply a 1-form ω on $\mathbb{C}$ is an expression of the form

$$\omega = f(z)dz$$

where f is a holomorphic (resp. meromorphic) function on $\mathbb{C}$. For the purpose of this paper, dz may be viewed as a formal symbol[10].

Definition 5.9. Let $\omega = f(z)dz$ and $\eta = g(w)dw$ be two 1-forms on respective open sets $U, V \subset \mathbb{C}$ and let $T : V \to U$ be a holomorphic mapping. We say ω *transforms* to η under T if

$$g(w) = f(T(w))T'(w)$$

One typically views transformations of 1-forms as a change of coordinates by setting $z = T(w)$, which implies that $dz = T'(w)dw$.

The purpose of a 1-form on $\mathbb{C}$ is to be integrated along a path γ in the complex plane. Using local charts on a Riemann surface, we may give an analogous definition.

Definition 5.10. Let X be a Riemann surface, *a holomorphic (resp. meromorphic) 1-form* is a collection of holomorphic (resp. meromorphic) 1-forms $\{\omega_\phi\}$, one for each chart ϕ in an atlas $\mathcal{A}$ of X. We also require that if two charts ϕ_1 and ϕ_2 have overlapping domains, their associated forms ω_1 and ω_2 transform under the transition function $T = \phi_1 \circ \phi_2^{-1}$ between the 2 charts.

Note that once a 1-form has been defined for some atlas $\mathcal{A}$, the compatibility of the charts on X, uniquely determines it on the rest of the charts on X. Moreover, this compatibility also allows us to define the order of a meromorphic 1-form at a point.

Definition 5.11. Let $\omega = f(z)dz$ be a meromorphic 1-form on X we define the *order of* ω at a point $p \in X$ denoted $\mathrm{ord}_p(\omega)$ to be the order of the function f at p. Note that this quantity is independent of the chart used to define ω near p.

As every Riemann surface is also a 2-dimensional C^∞ real manifold, we may broaden the class of 1-forms to consider C^∞ 1-forms.

[10]To be precise, a holomorphic (resp. meromorphic) 1-form ω on a complex analytic manifold X is a holomorphic (resp. meromorphic) section of the cotangent bundle $T^*(X)$. Analogously, one can define a k-form on X as a section of the k'th exterior power of the cotangent bundle $\bigwedge^k T^*(X)$ of X.

Definition 5.12. A C^∞ 1-form ω on a 2-dimensional real analytic manifold X is an expression locally of the form

$$\omega = f(x, y)dx + g(x, y)dy$$

where f and g are C^∞ functions on X.

We have the following notation for an open set $U \subset X$, where X is a complex manifold:

$$\Omega^{(1)}(U) = \{\text{holomorphic 1-forms on } U\}$$
$$\mathcal{M}^{(1)}(U) = \{\text{meromorphic 1-forms on } U\}$$
$$\mathcal{E}^{(1)}(U) = \{C^\infty \text{ 1-forms on } U\}$$

We note that the spaces $\mathcal{M}^{(1)}(U)$ and $\Omega^{(1)}(U)$ are $\mathcal{O}(U)$-modules and if U is connected, $\mathcal{M}^{(1)}(U)$ is a vector space over the field $\mathcal{M}(U)$.

6 Maps between Riemann surfaces

Now that we have defined holomorphic functions on Riemann surfaces, it is natural to consider maps between Riemann surfaces X and Y that preserve their complex structures. Such maps would allow us to pullback holomorphic and meromorphic functions from Y to X. Moreover, if bijective, the existence of such maps would provide a criterion for isomorphisms between Riemann surfaces and, therefore, open the door to their classification. For this purpose, we give the following definition.

Definition 6.1. Let X and Y be Riemann surfaces. A map $F : X \to Y$ is *holomorphic at* $p \in X$ if there exist a pair of complex charts ϕ_X and ϕ_Y on each Riemann surface such that $\phi_Y \circ F \circ \phi_X^{-1}$ is holomorphic at $z = \phi_X(p)$. F is a *holomorphic map* if it is holomorphic at all $p \in X$.

While according to the above definition, the holomorphicity of F at a p is contingent on the existence of a single pair of charts (ϕ_X, ϕ_Y) satisfying the condition, the compatibility of charts on X and Y implies that F is holomorphic if and only if the condition is satisfied for all pairs of charts containing p and $F(p)$ respectively.

Now, we present a lemma concerning some properties of maps between Riemann surfaces.

Lemma 6.2. *Let $F : X \to Y$ be a holomorphic map. Then;*

- *If $f : Y \to \mathbb{C}$ is a holomorphic (resp. meromorphic) function, then the composition $f \circ F : X \to \mathbb{C}$ is a holomorphic (resp. meromorphic) function on X.*
- *If Z is another Riemann surface and $G : Y \to Z$ is a holomorphic map, then the composition $G \circ F : X \to Z$ is a holomorphic map.*
- *If F is a not identically constant, F is a continuous open mapping.*

We may now define an isomorphism of Riemann surfaces :

Definition 6.3. A holomorphic map $F : X \to Y$ is an *isomorphism* if it is bijective and its inverse $F^{-1} : Y \to X$ is holomorphic. If such a map exists, we say that X and Y are isomorphic.

Example 6.4. The Riemann sphere $\mathbb{C}_\infty$ and the projective line $\mathbb{P}^1$ are isomorphic as Riemann surfaces.

Proof. It suffices to check that the map $F : \mathbb{P} \to \mathbb{C}_\infty$ defined by

$$F([z : w]) \mapsto (2\Re(z\bar{w}), 2\Im(z\bar{w}), |z|^2 - |w|^2)/(|z|^2 + |w|^2)$$

where $\Re(z)$ and $\Im(z)$ denote the real and imaginary part of any $z \in \mathbb{C}$, is indeed a holomorphic bijection. □

We now state two important structural properties of holomorphic maps, which are essential for the development of the rest of this section.

Proposition 6.5. *Let X be a compact Riemann surface and $F : X \to Y$ a non-constant holomorphic map. Then, F is surjective and Y is a compact Riemann surface.*

Proof. X is open and from Lemma 6.2, F is an open map, which means that $F(X)$ is open in Y. Also, since X is compact and F is continuous, $F(X)$ is compact and thus closed since Y is Hausdorff. Then $F(X)$ is both closed and open, which implies that $F(X) = Y$ by connectedness of Y *(and the fact that $F(X) \neq \emptyset$).* This also implies that Y is compact. □

Proposition 6.6. *Let $F : X \to Y$ be a holomorphic map. Then, for every point $p \in Y$, the preimage $F^{-1}(\{p\})$ is a discrete subset of X.*

Proof. Let $p \in Y$ and let $x \in F^{-1}(\{p\})$. By the definition of holomorphic maps, there exists a local coordinate z centred at x and a local coordinate w centred at p such that F is given by a holomorphic function $f(z)$. Moreover, f has a zero at x and since zeros of holomorphic functions are discrete, there exists a neighbourhood U of x in which x is the only zero and thus $U \cap F^{-1}(\{p\}) = \{x\}$. Therefore, $F^{-1}(\{p\})$ is discrete. □

Corollary 6.7. *If $F : X \to Y$ is a non-constant holomorphic map between compact Riemann surfaces, then for all $p \in Y$, $F^{-1}(\{p\})$ is a finite set.*

Proof. This follows from the previous proposition and the fact that closed discrete subsets of compact sets are finite, by the finite sub-covering definition of compactness □

Example 6.8. [Maps to the Riemann Sphere] One important class of holomorphic maps are those arising from meromorphic functions on a Riemann surface. Indeed, a meromorphic function $f : X \to \mathbb{C}$ induces a holomorphic map $F : X \to \mathbb{C}_\infty$ given by

$$F(p) = \begin{cases} f(p), & \text{if } f \text{ is holomorphic at } p \\ \infty, & \text{if } f \text{ has a pole at } p \end{cases}$$

It is clear that F is holomorphic at points where f is holomorphic. Now, suppose that f has a pole at a point p. Since poles are isolated, we can choose a neighbourhood U of p such that f has no other zeroes or poles in U. Let z be a local coordinate centred at p, so that $z(p) = 0$, and let w denote the standard coordinate on $\mathbb{C}_\infty \setminus \{\infty\}$ (i.e. on the complex plane $\mathbb{C}_\infty \setminus \{\infty\} \cong \mathbb{C}$).

On the punctured neighbourhood $U \setminus \{p\}$, we have $F(z) = f(z)$, i.e., F is locally given by the function $w = f(z)$. To analyze the behaviour at p, we switch to the chart at infinity on $\mathbb{C}_\infty$, where the transition function is $w \mapsto 1/w$. In this local chart, F is given by

$$\tilde{F}(z) = \frac{1}{f(z)}.$$

Since f has no zeroes in our chosen neighbourhood U, the function $\tilde{F}(z)$ is holomorphic on $U \setminus \{p\}$, and it extends holomorphically at $z = p$ by setting $\tilde{F}(p) = 0$, as the singularity is removable. Therefore, in local coordinates, F is given by

$$F(z) = \begin{cases} \dfrac{1}{f(z)}, & \text{if } z \in U \setminus \{p\}, \\ 0, & \text{if } z = p, \end{cases}$$

which is holomorphic in a neighbourhood of p when viewed in the coordinate chart at infinity.

6.1 Global Properties of Holomorphic maps

We now come to discuss the global properties of holomorphic maps between Riemann surfaces. Historically, the original development of the theory by Riemann was in large part motivated by the study of these properties. Riemann was interested in the graphs of "multi-valued" functions on $\mathbb{C}$ such as $z^{\frac{1}{n}}$ and $\log(z)$. In the first case, the graph is given by the vanishing set $X_F \subset \mathbb{C}^2$ of the polynomial $F(x, y) = y - x^n$. As seen in Section 4.2, this set is a Riemann surface equipped with a projection map $X_F \xrightarrow{\pi} \mathbb{C}$. A branch of $z^{1/n}$ on an open set $U \subset \mathbb{C}$ is given by a local section of the map π, as the following diagram illustrates.

$$\begin{tikzcd} \mathbb{C}^2 & X_F \arrow[l, hook'] \arrow[r, "\pi"] & U \arrow[l, bend right, dashed, "z^{1/n}"'] \end{tikzcd}$$

If we define X_F for a more general multi-valued function f, i.e. $F(x, y) = y - f(x)$. Then, under the map π, the graph of f can be viewed as a covering space projecting onto an open subset of $U \subset \mathbb{C}$ whose fiber above each point $z \in U$ consists of all possible values of f.[11]

In this section, we develop this notion of a local n-fold covering of a Riemann surface by another.

Theorem 6.9 (Local Normal Form). *Let $F : X \to Y$ be a non-constant holomorphic map between Riemann surfaces. For every $p \in X$, there exists a unique integer $n_p \geq 1$ such that for every chart (ϕ_Y, U_Y) centred at $F(p)$, there exist a chart (ϕ_X, U_X) centred at p such that $(\phi_Y \circ F \circ \phi_X^{-1})(z) = z^{n_p}$.*

Proof. Let z and w be local coordinates centred at p and $F(p)$ respectively under the charts ϕ_X and ϕ_Y. F is holomorphic at p and thus $w = (\phi_Y \circ F \circ \phi_X^{-1})(z) = f(z)$ for some holomorphic function f whose Taylor expansion, with $m_0 \geq 0$ is of the form

$$f(z) = \sum_{m=m_0}^{n} a_m z^m.$$

As f vanishes at 0, we must have $m_0 \geq 1$ and thus

$$f(z) = z^{m_0} h(z)$$

where h is a holomorphic function such that $h(0) \neq 0$. Thus, there exists locally a holomorphic function $g(z)$ in a neighbourhood of $z = 0$ such that $g(z)^{m_0} = h(z)$ and $g(0) \neq 0$. In the above (simply-connected) local coordinate z, since h is non-vanishing on a some neighbourhood $U \in \mathbb{C}$ of p, by continuity, we can use $g(z) := \exp(m_0 L(z))$, where $L(z)$ is a branch of the complex logarithm, as the root to our desired function $h(z)$. Thus we may then rewrite

$$f(z) = z^{m_0} g(z)^{m_0} = (zg(z))^{m_0}.$$

Letting $\alpha(z) = zg(z)$, we notice that : $\quad \alpha'(z) = g(z) + zg'(z) \implies \alpha'(0) = g(0) \neq 0$

which implies that in a neighbourhood of zero, α is bijective (*Implicit function theorem*). By Proposition 4.3, the composition $\psi = \alpha \circ \phi_X$ is a chart centred at p giving a new local coordinate $\zeta = \alpha(z)$. In a neighbourhood of p, under ψ, F is given by

$$\begin{aligned} F(z) &= ((\phi_Y \circ F \circ \psi^{-1})(z) \\ &= (\phi_Y \circ F \circ \phi_X^{-1} \circ \alpha^{-1})(\zeta) \\ &= (zg(z))^{m_0} \\ &= \zeta^{m_0} \end{aligned}$$

Letting $n_p := m_0$, we have hence proven the theorem. $\square$

[11] In complex analysis, where the multi-valued functions $log(z)$ and $z^{\frac{1}{n}}$ are first seen, one would usually define a branch cut to render them single valued (restricting their domain of definition), allowing us to work them as any other globally holomorphic function.

Note furthermore that this integer n_p is unique as it records the number of pre-images in a neighbourhood of p of points in a neighbourhood of $F(p)$, which is independent of the choice of chart. We may therefore give the following definitions.

Definition 6.10. Let $p \in X$ and $F : X \to Y$ be a non-constant holomorphic map:

- We define the *multiplicity of* F *at* p denoted $\mathrm{mult}_p(F)$ to be the positive integer exponent n_p in the local normal form of F at p.
- We say p is a *ramification point* if $\mathrm{mult}_p(F) \geq 2$; the image of a ramification point is called a *branch point.*

The ramification points of a map are those at which it fails to be a local homeomorphism. For example, if we refer back to the diagram in Section 6.1 and consider the map $F : \mathbb{C} \to \mathbb{C}$ given by $F(z) = z^2$, viewed as a map of Riemann surfaces, we find that F has a unique ramification point of multiplicity 2 at $z = 0$. Aside from the image of this point, namely $F(0) = 0$, every point $p \in \mathbb{C}$ can be enclosed in a neighbourhood V_p whose preimage under F consists of two disjoint open sets, each of which is homeomorphic to V_p.

More generally, one can also view ramification as the failure of a map to be a local n-fold covering. This information is crucial for the study of Riemann surfaces and can be expressed as follows.

Definition 6.11. Let $F : X \to Y$ be a holomorphic map between Riemann surfaces X, Y and let $y \in F(X)$. We define the degree of F at y, denoted $d_y(F)$, as:

$$d_y(F) = \sum_{p \in F^{-1}(y)} \mathrm{mult}_p(F)$$

This definition is particularly useful when X and Y are compact because the degree becomes an invariant of F, independent of the point y.

Proposition 6.12. *Let $F : X \to Y$ be a holomorphic map between compact Riemann surfaces. Then $d_y(F)$ is independent of the point y.*

Proof. We will show that the function $y \mapsto d_y(F)$ is locally constant. Then, since Y is connected, it will be globally constant.

We let $y \in Y$ be arbitrary and consider its preimage under F which is a finite set $\{x_1, \ldots, x_n\}$ by Corollary 6.7. As Y is Hausdorff, for each x_i, we may select a centred local coordinate z_i in a pairwise disjoint neighbourhood U_i. Using the local normal form theorem, in each U_i, F may be expressed as the map $z \mapsto z^{m_i}$. Thus, any given y' in a small neighbourhood of y has exactly $\sum_i m_i = d_y(F)$ preimages of multiplicity one in $\bigcup_i U_i$.To prove the result, it is left to show that all preimages of y' are contained in the union $\bigcup_i U_i$.

We proceed by contradiction and suppose that in each neighbourhood of y, we may find a point y_i for which there exists $p_i \in F^{-1}(y_i) \setminus \bigcup_j U_j$. We may hence construct a sequence $(y_i)_{i \in \mathbb{N}}$ by choosing each y_i from a nested, *local* countable base $\{V_i\}_{i \in \mathbb{N}}$ at y.[12] By construction, the sequence $(y_i)_i$ converges to y. Also, considering the sequence $(p_i)_{i \in I}$ of points associated to y_i, we may find a subsequence (p_{i_k}) of (p_i) which converges to some point $p \in X$ by compactness of X. By continuity of F, we have that $F(p) = y$, which implies that $p = x_j$ j. However, this contradicts our assumption that none of the p_i's were contained in any of the neighbourhoods U_i. Therefore, the function $y \mapsto d_y(F)$ is locally constant and thus constant on Y.

□

[12] A Riemann surface is second countable, so it admits a (global) base of the topology that is countable. Hence, there is a local base at y, which consists only of elements of the base containing y, that is countable, making our choice of $\{V_i\}_{i \in \mathbb{N}}$ possible.

As it is independent of the point in Y at which we calculate it, we refer to the quantity $d_y(F)$ as the *degree of map* F and denote it by $\deg(F)$. The degree of a map F is one of, if not the most, important invariants! It allows us to deduce global properties of F, X and Y, by simply calculating what happens at a single point.

Corollary 6.13. *Let $F : X \to Y$ be a holomorphic map between compact Riemann surfaces. Then F is an isomorphism if and only if* $\deg(F) = 1$.

Proof. We have already proven that F is surjective, and a map of degree one is precisely one where the preimage of every point is a singleton. □

Corollary 6.14. *Let X be a compact Riemann surface and f a meromorphic function on X with a single pole of order one. Then X is isomorphic to the Riemann sphere.*

Proof. From Example 6.8, we may extend f to a holomorphic map $F : X \to \mathbb{C}_\infty$. As f has a single pole of order one, the preimage $F^{-1}(\infty)$ is a single point. Thus, $d_\infty(F) = 1$. By the previous results, this implies that $\deg(F) = 1$ and thus an isomorphism. □

We now present an important consequence of Proposition 6.12, which imposes a restriction on the types of meromorphic functions on compact Riemann surfaces.

Before doing so, we introduce an important correspondence between the order of a function on a Riemann surface and the multiplicity of a map between Riemann surfaces.

Lemma 6.15. *Let $f : X \to \mathbb{C}$ be a meromorphic function on a Riemann surface X, with the associated holomorphic map $F : X \to \mathbb{C}_\infty$, using the construction presented in Example 6.8. Then,*

- *If $p \in X$ is a zero of f, then* $\mathrm{mult}_p(F) = \mathrm{ord}_p(f)$
- *If $p \in X$ is a pole of f, then* $\mathrm{mult}_p(F) = -\mathrm{ord}_p(f)$
- *Otherwise,* $\mathrm{mult}_p(F) = \mathrm{ord}_p(f - f(p))$

Proof. For any $p \in X$, we may choose a local coordinate z centred at p, by Theorem 6.9, so that F is simply given by $z \to z^{m_p}$ for $m_p \in \mathbb{N}$. Then, by the definition of multiplicity of the local normal form (Definition 6.10), $\mathrm{mult}_p(F) = m_p$.

- If $F(p) = 0$, i.e. f has a zero at p, then using the standard stereographic chart of $\mathbb{C}_\infty$ near 0, $\zeta(w) = 1/w$, we have that locally $f(z) = z^{m_p}$. Thus, by Definition 3.18 of the order, $\mathrm{ord}_p(f) = m_p$
- If $F(p) = \infty$, i.e. f has a pole at p then using instead the standard stereographic chart of $\mathbb{C}_\infty$ near ∞, $\zeta(w) = 1/w$, we have that in local coordinates, $1/f(z) = z^{m_p}$. So, $f(z) = z^{-m_p}$ and hence by Definition 3.18, $\mathrm{ord}_p(f) = -m_p$
- Otherwise, we may apply the exact same argument as for p being a zero of f, but to the function $f - f(p)$ instead.

□

Now, we prove that the sum of the orders of meromorphic functions at all points $p \in X$ must vanish, where X is a compact Riemann surface. This was done in Example 5.7 using an analytical approach, but will be done here using the Residue theorem and a topological approach, with the tools we defined above.

Proposition 6.16. *Let $f : X \to \mathbb{C}$ be a meromorphic function on a compact Riemann surface X. Then,* $\sum_{p \in X} \mathrm{ord}_p(f) = 0$.

Proof. Given a meromorphic function $f : X \to \mathbb{C}$, we may view it as a holomorphic map $F : X \to \mathbb{C}_\infty$ (recall Example 6.8). The quantity $\sum_{p \in X} \mathrm{ord}_p(f)$ corresponds to the sums of the zeroes and poles of f counted with multiplicity. Under F, the zeros and poles of f correspond to, respectively, the preimages of the points 0 and ∞ of $\mathbb{C}_\infty$. Then, the absolute value of the order of those points is given by the local multiplicity of F. Therefore,

$$\begin{aligned}\sum_{p\in X} \mathrm{ord}_p(f) &= \sum_{p\in F^{-1}(0)} \mathrm{mult}_p(F) \; - \sum_{q\in F^{-1}(\infty)} \mathrm{mult}_q(F) \\ &= d_0(F) - d_\infty(F) = \deg(F) - \deg(F) = 0\end{aligned}$$

where $\deg(F)$ is constant since X and $\mathbb{C}_\infty$ are compact Riemann surfaces (See Proposition 6.12). $\square$

6.2 Hurwitz formula

We conclude this section with a brief mention of the Hurwitz formula, which relates the degree of a map $F : X \to Y$ to the genera g_X and g_Y of the respective Riemann surfaces. While we will not elaborate on the proof of the formula, one can arrive at it by triangulating X and Y [Mir95, p. 50-52]. Informally speaking, a *triangulation* T on a surface M is a finite partition $M = \bigcup_i M_i$ into closed subsets which are individually homeomorphic to a triangle $\Delta \subset \mathbb{R}^2$ and whose pairwise intersections $M_i \cap M_j$ are either empty, a single point or the image of an edge of Δ. If we let E, V, F respectively denote the set of edges, vertices and faces of T we define the *Euler characteristic* $e_T(M)$ as

$$e_T(M) = |V| - |E| + |F|.$$

In fact, one can show that $e_T(M)$ is an invariant of M independent of the triangulation T [Mir95, p. 51]. In the case where M is a compact orientable surface without boundary (which is the case for compact Riemann surfaces), the genus of M is given by $g_M = 1 - \frac{1}{2}e_T(M)$. We may then state the Hurwitz formula

Theorem 6.17 (Hurwitz formula)**.** *Let X and Y be compact Riemann surfaces and $F : X \to Y$ be a non-constant holomorphic map. Let g_X, g_Y be the respective genera of X and Y. Then, we have the following identity;*

$$2g_X - 2 = \deg(F)\,(2g_Y - 2) + \sum_{p\in X} [\mathrm{mult}_p(F) - 1]$$

Proof. *See [Mir95, p. 52]* $\square$

7 Divisors

Finally, we present the concept of divisors on a Riemann surface, which are essential for the formulation and understanding of the Riemann-Roch theorem. We begin by giving the definition of a divisor, its degree, and some important examples. We then study linear systems of divisors and their associated spaces of meromorphic functions. In all that follows, we let X refer to any *compact* Riemann surface.

Definition 7.1. A divisor on X is a function $D : X \to \mathbb{Z}$ whose support, i.e. the set of points $p \in X$ such that $D(p) \neq 0$, is a discrete closed subset of X. We denote by $\mathrm{Div}(X)$ the set of all divisors on X, noting that it is a group under point-wise addition. It is convenient to represent a divisor as a formal sum of integral multiples of points $p \in X$ with coefficients where $D(p)$ at each point:

$$D = \sum_{p\in X} D(p) \cdot p$$

As we assumed X was compact, any divisor $D : X \to \mathbb{Z}$ has finite support. The group $\mathrm{Div}(X)$ of divisors on X is then exactly the free abelian group generated by the points of X.

Definition 7.2. The degree of a divisor D on X, is defined as

$$\deg(D) = \sum_{p\in X} D(p).$$

The degree map is a group homomorphism, $\deg : \mathrm{Div}(X) \to \mathbb{Z}$. We denote its kernel by $\mathrm{Div}_0(X)$.

There is a natural partial ordering on the group of divisors $\mathrm{Div}(X)$ associated to a compact Riemann surface X. We declare that for $D_1, D_2 \in \mathrm{Div}(X)$, $D_1 \geq D_2$ if for all $p \in X, D_1(p) \geq D_2(p)$; an analogous definition is given for the strict inequality sign. If we let 0 be the zero divisor, it can easily be seen that any divisor D may be expressed as a sum $D = E + F$ where $E \geq 0$ and $F \leq 0$.

7.1 Types of divisors

At first glance, the general definition of a divisor may appear abstract and lacking in concrete meaning. However, its generality is justified once one realizes that divisors naturally arise from many concrete objects defined on a Riemann surface X. Divisors then provide an elegant tool to both formulate and prove theorems about said objects. We give some examples of important types of divisors.

Definition 7.3. Let f be a meromorphic function on X. Then, the divisor of f is defined as:

$$\mathrm{div}(f) = \sum_{p \in X} \mathrm{ord}_p(f) \cdot p$$

Divisors of this form are called principal divisors on X, and they form the subgroup $\mathrm{PDiv}(X) \subset \mathrm{Div}(X)$. As X is compact, it is a subgroup of $\mathrm{Div}_0(X)$ by Proposition 6.16.

Definition 7.4. Let ω be a meromorphic 1-form on X. Then, the divisor of ω is defined as:

$$\mathrm{div}(\omega) = \sum_{p \in X} \mathrm{ord}_p(\omega) \cdot p$$

The divisors of this form are called canonical divisors on X, and we denote their collection by $\mathrm{KDiv}(X)$. Importantly, the set $\mathrm{KDiv}(X)$ does not form a group.

We now present a few useful properties of the divisor of meromorphic functions and 1-forms.

Proposition 7.5. *Let f and g be non-zero meromorphic functions on X and let ω be a non-zero meromorphic 1-form on X. Then :*

- $\mathrm{div}(fg) = \mathrm{div}(f) + \mathrm{div}(g)$
- $\mathrm{div}(1/f) = -\mathrm{div}(f)$
- $\deg(\mathrm{div}(f)) = 0$, if X is a *compact* Riemann surface.
- $\mathrm{KDiv}(X) = \mathrm{div}(\omega) \; + \; \mathrm{PDiv}(X)$

Proof. The first two identities are restatements of the properties of the order function presented in Proposition 3.19, while the third is a restatement of Proposition 6.16 on the vanishing sum of orders of meromorphic functions on compact Riemann surfaces.

For the final identity, we consider another non-zero 1-form on X, $\eta \in \Omega^1(X)$. Then, use a local coordinate z centred at a point $p \in X$ to locally write $\eta = g(z)dz$ and $\omega = f(z)dz$ for some locally defined meromorphic functions f and g. We consider the function $h(z) := f(z)/g(z)$, which is also meromorphic in a neighbourhood of p. If w is another local coordinate centred at p, and T is the transition function such that $T(w) = z$, we have:

$$\omega = f(T(w))T'(w)dw \quad \text{and} \quad \eta = g(T(w))T'(w)dw$$

By taking the quotient of the 2 meromorphic function representing both 1-forms, we obtain:

$$f(T(w))T'(w)/g(T(w))T'(w) = f(z)/g(z)$$

This shows that the function h is independent of the chart chosen to represent ω and η, and thus, it can be extended to a global meromorphic function on X such that $\omega = h\eta$. The last identity is then obtained from the following property of the order function: $\mathrm{ord}_p(f\omega) = \mathrm{ord}_p(f) + \mathrm{ord}_p(\omega)$ □

Definition 7.6. Let $F : X \to Y$ be a non-constant holomorphic mapping of compact Riemann surfaces and $q \in Y$, a point. We define the *inverse image divisor*, $F^*(q)$, of F at the point q as the divisor:

$$F^*(q) = \sum_{p \in F^{-1}(q)} \text{mult}_p(F) \cdot p.$$

Since the multiplicity of a map is a well-defined integer associated with each point of its domain, $F^*(q)$ is indeed a divisor on X. Moreover, we have that from the definition above, it follows that $\deg(F^*(q)) = \deg(F)$, where $\deg(F)$ is the degree of the holomorphic map $F : X \to Y$ as defined in Definition 6.12.

As with functions and differential forms, via a non-constant holomorphic map $F : X \to Y$, one can pull-back a divisor $D \in \text{Div}(Y)$ to a divisor $F^*(D) \in \text{Div}(X)$ in the following way:

$$F^*(D) = \sum_{q \in Y} D(q) F^*(q),$$

Note that the pull-back map $F^* : \text{Div}(Y) \to \text{Div}(X)$ is a group homomorphism and additionally obeys the identity:

$$\deg(F^*(D)) = \deg(F)\deg(D).$$

7.2 Space of functions and forms attached to a divisor

Definition 7.7. Let D be a divisor on a *compact* Riemann surface X. We define the complex vector space $L(D)$ of *meromorphic functions with poles bounded by* D:

$$L(D) = \{f \in \mathcal{M}(X) \mid \text{div}(f) \geq -D\}$$

Concretely, the space $L(D)$ consists of all the meromorphic functions on X, which for a given point p, have a pole at p of at most order $D(p)$ or a zero of any order, if $D(p) > 0$, and a zero of order at least $-D(p)$, if $D(p) < 0$.

Similarly, the vector space of meromorphic 1-forms with poles bounded by D, $L^{(1)}(D)$, is:

$$L^{(1)}(D) = \{\omega \in \mathcal{M}^{(1)}(X) \mid \text{div}(\omega) \geq -D\}$$

We present some lemmas that will prove to be very handy in later chapters.

Lemma 7.8. *Let X be a compact Riemann surface and D_1, D_2 be divisors on X. If $D_1 \leq D_2$, then $L(D_1) \subseteq L(D_2)$.*

Proof. We may define the following divisor:

$$E = D_2 - D_1 = \sum_{p \in X} (D_2(p) - D_1(p)) \cdot p$$

Thus, we can write $D_2 = D_1 + E$ where $E \geq 0$. If $f \in L(D_1)$, then $\text{div}(f) + D_1 \geq 0$, so $\text{div}(f) + D_2 = \big(\text{div}(f) + D_1\big) + E \geq 0$. Hence $f \in L(D_2)$, and $L(D_1) \subseteq L(D_2)$. □

Lemma 7.9. *Let X be a compact Riemann surface. If D is a divisor on X and* $\deg(D) < 0$, *then* $L(D) = \{0\}$.

Proof. We proceed by contradiction. Suppose we take $f \in L(D)$ and f is not identically zero. Consider the following divisor $E = \text{div}(f) + D$. Since $f \in L(D)$, we know that $\text{div}(f) \geq -D$, so $\text{div}(f) + D \geq 0$, so $E \geq 0$.

On the other hand, by the third statement of Proposition 7.5, we know that $\deg(\text{div}(f)) = 0$ since X is compact. Hence, we have that by linearity of the degree of a divisor:

$$\deg(E) = \deg(D) + \deg(\text{div}(f)) = \deg(D)$$

However, $\deg(D) < 0$, so we have that $\deg(E) < 0$, which is a contradiction and thus, $L(D)$ is the trivial vector space $L(D) = \{0\}$. □

Lemma 7.10. *Let $D \in \mathrm{Div}(X)$ be a divisor on a compact Riemann surface X and let $p \in X$. Then, either $L(D-p) = L(D)$ or $L(D-p)$ is of codimension one in $L(D)$.*

Proof. Let $n := -D(p)$ and let z be a local coordinate centred at p. With respect to this coordinate, a function $f \in L(D)$ has a local Laurent series expression of the form

$$f(z) = a_n z^n + \sum_{k=n+1}^{\infty} a_k z^k$$

Consider the linear map $\alpha : L(D) \to \mathbb{C}$ given by $f \mapsto a_n$. Its kernel consists of the functions $f \in L(D)$ such that $\mathrm{ord}_p(f) \geq n+1$. So, $\mathrm{div} f - p \geq -D$, which implies that $f \in L(D-p)$. By the rank-nullity theorem,

$$\dim(L(D)/L(D-p)) = \dim \alpha(L(D)) = 0 \text{ or } 1.$$

So, if α is onto, then $\dim \alpha(L(D)) = \dim \mathbb{C} = 1$. In which case, $L(D-p)$ has codimension one in $L(D)$. Otherwise, $\dim \alpha(L(D)) = \dim\{0\} = 0$. In which case, $L(D) = L(D-p)$ since the latter is a subspace of the former by Lemma 7.8.

□

This lemma has the following important consequence.

Corollary 7.11. *For any divisor $D \in \mathrm{Div}(X)$, let $P, N \in \mathrm{Div}(X)$ be the unique non-negative divisors with disjoint support such that $D = P - N$. Then $\dim L(D) \leq 1 + \deg(P)$. In particular, $L(D)$ is **finite**-dimensional.*

Proof. First, note that $L(0)$ is the space $\mathcal{O}(X)$ of holomorphic functions on X which, from [5.7], is of dimension one, consisting only of the constant functions. If D is an arbitrary divisor and P, N are as given above, by recursively applying Lemma 7.10 to P, we have that $\dim L(P) \leq 1 + \deg P$. As $D \leq P$, implies $L(D) \subset L(P)$ by Lemma 7.8, we obtain that $\dim L(D) \leq \dim L(P) \leq 1 + \deg(P)$. □

Lemma 7.12. *Let $K = \mathrm{div}(\omega)$ be a canonical divisor on X where $\omega \neq 0$, let D be a divisor on X and f be a meromorphic function such that $f \in L(D+K)$. The multiplication map $\mu_\omega : L(D+K) \to L^{(1)}(D)$, $f \mapsto f\omega$ is an isomorphism of vector spaces. In particular, we have that :*

$$\dim L(D+K) = \dim L^{(1)}(D)$$

Proof. The map is clearly linear by the distributivity of the multiplication of a meromorphic function with a meromorphic 1-form, and is clearly injective since $f\omega = 0$ if and only if $f = 0$, as we assumed ω was nonzero. For surjectivity, if we choose a 1-form $v \in L^{(1)}(D)$, then by the last statement of Proposition 7.5, we know that we can find $g \in \mathrm{PDiv}(X)$ such that $g\omega = v$. Moreover we see that $g \in L(D+K)$ since :

$$\mathrm{div}(g) + K = \mathrm{div}(g) + \mathrm{div}(\omega) = \mathrm{div}(g\omega) = \mathrm{div}(v) \geq -D \text{ since } v \in L^{(1)}(D)$$

Hence, we have that there exists $g \in L(D+K)$ such that for any $v \in L^{(1)}(D)$, $\mu_\omega(g) = v$. By the rank-nullity theorem, the map μ_ω is hence an isomorphism. □

Lemma 7.13. *Let X be a compact Riemann surface of genus g. Additionally, suppose the existence of a non-constant global meromorphic function $f \in \mathcal{M}(X)$. Then, any canonical divisor D on X has degree $2g-2$.*

Proof. We give a sketch of the proof found in [Mir95, p. 132-133].
Using the function f, we construct a holomorphic map $F : X \to \mathbb{C}_\infty$, in the sense of [6.8]. We then consider the 1-form $\omega = dz$ on $\mathbb{C}_\infty$ which only has a pole of order 2 at infinity. We then consider the pullback $F^*\omega \in \Omega^1(X)$ and use the Hurwitz formula [6.17] to show that $\deg(F^*\omega) = 2g-2$. Since any two canonical divisors on X differ by a principal divisor and $\deg(\mathrm{div}(f)) = 0$ by Proposition 7.5, it follows that all canonical divisors on X have the same degree, namely $2g-2$. □

7.3 Complete linear systems of divisors and holomorphic maps to $\mathbb{P}^n$

As suggested by the last statement of Proposition 7.5, there is some importance to analyzing divisors (not necessarily arising from 1-forms) of a Riemann surface X that differ by the divisor of some meromorphic function on X. We call this relation between divisors linear equivalence.

Definition 7.14. Two divisors D_1, D_2 on a Riemann surface X are said to be *linearly equivalent*, denoted $D_1 \sim D_2$, if their difference is a principal divisor. Linear equivalence is an equivalence relation.

Definition 7.15. The complete linear system of a divisor D, denoted by $|D|$ is the set:

$$|D| = \{E \in Div(X) \ : \ E \sim D \text{ and } E \geq 0\}.$$

Recall that if V is a vector space over $\mathbb{C}$, we define the projectivization $\mathbb{P}(V)$ of V as the set of all one-dimensional subspaces of V. We have a map $\text{span}(\cdot) : V \setminus \{0\} \to \mathbb{P}(V)$ and denote by $[v]$ the image of an element $v \in V \setminus \{0\}$.

As multiplying a meromorphic function on X by a non-zero complex number does not affect its order at points, for a divisor $D \in \text{Div}(X)$, the map $S : L(D) \to \text{Div}(X)$ mapping f to $\text{div}(f) + D$ descends to a map $S : \mathbb{P}(L(D)) \to \text{Div}(X)$. In particular, since $\text{div}(f) \geq -D$, the image of the map is contained in the complete linear system $|D|$.

Proposition 7.16. *Let X be a compact Riemann surface and D a divisor on X. Then, the map $S : \mathbb{P}\big(L(D)\big) \to |D|$ is a bijection.*

Proof. The map is clearly injective since if $S([f]) = S([g])$ for $f, g \in L(D)$, then $\text{div(f/g)} = 0$ which means that $f/g = \alpha$ for some $\alpha \in \mathbb{C}^*$,i.e., $[f] = [g]$. The map is surjective as since $E \in |D|$, $E - D = \text{div}(h)$ for $h \in \mathcal{M}(X)$. As $E \geq 0$, $\text{div}(h) \geq D$, that is, $h \in L(D)$ and thus $S([h]) = \text{div}(h) + D = (E - D) + D = E$. Therefore, the map S is bijective. $\square$

We thus have a 1-1 correspondence between one-dimensional subspaces of $L(D)$ and elements of the complete linear system $|D|$. We use this correspondence to define *linear (not necessarily complete) systems.*

Definition 7.17. A *linear system* is a subset $E \subset |D|$ of a complete linear system whose pre-image under the map S is a linear subspace of $L(D)$.

We note that, as a complex manifold, $\mathbb{P}(L(D))$ is abstractly isomorphic to $\mathbb{P}^n$. Under such an isomorphism, a linear system E is a subset of $|D|$ whose preimage in $\mathbb{P}^n$ is a linear subspace of some dimension $k < n$. This property does not depend on the choice of isomorphism to $\mathbb{P}^n$.

Other than simply grouping divisors together, the above construction has an important role in defining holomorphic maps to projective space, which will be useful in Chapter 10.1, where we will show that, in fact, any compact Riemann surface can be embedded into projective space via a holomorphic map. Let us now define what is meant by a holomorphic map to projective space.

Definition 7.18. Let X be a compact Riemann surface. A map $\phi : X \to \mathbb{P}^n$ is a *holomorphic at* $p \in X$, if, near p, there are locally-defined holomorphic functions $g_0, g_1, \ldots, g_n$, which do not all vanish at p and such that $\phi(x) = [g_0(x), g_1(x), \ldots, g_n(x)]$ for x near p. The map ϕ is a *holomorphic map* if this condition holds for all $p \in X$.

Since this definition requires the functions $g_0, \ldots, g_n$ to be holomorphic, and there exist no non-constant global holomorphic functions on a compact Riemann surface, the task of constructing a non-constant holomorphic map ϕ to $\mathbb{P}^n$ can seem difficult. A priori, one needs to take an open cover of X, define ϕ on members of the cover using locally-defined holomorphic functions and then verify that this local data glues together to form a globally defined map. The task would be much easier if one could construct ϕ from functions defined on all of X. The following proposition shows that this is possible: one can define a holomorphic map from globally defined meromorphic functions.

Proposition 7.19. *This proposition is separated into two parts:*

[Existence] *Let $\mathcal{F} := (f_0, \dots, f_n)$ be a sequence of non-zero global meromorphic functions which are not all identically zero and, on the subset of X where $\phi_{\mathcal{F}}(x)$ is defined, let $\phi_{\mathcal{F}}(x) = [f_0(x) : \dots : f_n(x)]$. Then, the map $\phi_{\mathcal{F}}$ extends to a holomorphic map to $\mathbb{P}^n$ and, conversely, any holomorphic map $\phi : X \to \mathbb{P}^n$ is of the form $\phi_{\mathcal{F}}$, for some sequence of meromorphic functions $\mathcal{F}$.*

[Uniqueness] *If two sequences $\mathcal{F}$ and $\mathcal{G}$ produce the same holomorphic maps, i.e., if $\phi_{\mathcal{F}} = \phi_{\mathcal{G}}$, then there exists a meromorphic function $\lambda \in \mathcal{M}(X)$ such that $f_i = \lambda g_i$ for all $0 \leq i \leq n$.*

Proof. We begin by proving the existence part of the proposition. Let $\mathcal{F} = (f_0, \dots, f_n)$ and $\phi_{\mathcal{F}}$ be as above; let $p \in X$ be a point where $\phi_{\mathcal{F}}$ is not defined, i.e., at least one function f_i has a pole at p or all the functions f_i have a zero at p. Let $n := \min_i(\operatorname{div}(f_i)(p))$ and let z be a local coordinate centred at p. As the combined set of poles and zeroes of elements of $\mathcal{F}$ is discrete, we may assume that in a sufficiently small neighbourhood of p, the functions f_i have no other poles or zeroes outside of p. Thus, in that neighbourhood, the quantity $[f_0(z)/z^n : \dots : f_n(z)/z^n]$ is well-defined as a projective point and extends the map $\phi_{\mathcal{F}}$ to p. Repeating this procedure at all points where $\phi_{\mathcal{F}}$ is not defined gives an extension of the map to the whole of X.

Conversely, if $\phi : X \to \mathbb{P}^n$ is a holomorphic map, for all $0 \leq i \leq n$, let x_i be the coordinate function on $\mathbb{P}^n$. Since $\phi(X)$ is not identically the zero element of $\mathbb{P}^n$, $[0, \dots, 0]$, there is some x_j that is not identically zero. By relabelling our coordinate functions, we may, without loss of generality, suppose that x_0 is not identically zero on $\phi(X)$. Thus, for any $i \in \{0, \dots, n\}$, we may define the global functions $f_i := \phi^*(x_i/x_0)$ as the pullback of x_i/x_0. Let p be any point in X. By definition, since ϕ is a holomorphic map, for any r in a neighbourhood of p, we have that $\phi(r) = [g_0(r), \dots, g_n(r)]$ for some locally holomorphic functions g_i. Thus, near p, $f_i = \phi^*(x_i/x_0) = g_i/g_0$, which is meromorphic since f_i is the ratio of holomorphic functions g_i, g_0. Since this holds for any $p \in X$, and for any $i \in \{0, \dots, n\}$, we have that $\phi = \phi_{\mathcal{F}}$, where $\mathcal{F} = \{1, f_1, \dots, f_n\}$, i.e. the meromorphic functions $f_i = \phi^*(x_i/x_0)$, for $0 \leq i \leq n$, generate the map ϕ.

To show the uniqueness part of the proposition, let $\mathcal{F}$ and $\mathcal{G}$ be sequences of global meromorphic maps as above. Since $\phi_{\mathcal{G}} = \phi_{\mathcal{F}}$ are holomorphic maps, without loss of generality, we may assume that g_0 is not identically zero and that f_0 is not identically zero. So, we can define $\lambda : X \to \mathbb{C}_\infty/\{0\}$, $\lambda(p) = f_0(p)/g_0(p)$ at every $p \in X$, where g_0 and f_0 have no pole nor zero. Then, since λ is the quotient of meromorphic functions, it is meromorphic on $X/\{p \in X : \operatorname{ord}_p(g_0) \neq 0, \operatorname{ord}_p(f_0) \neq 0\}$. Now, let $U = \{p \in X : \forall h \in \mathcal{F} \cup \mathcal{G}, \operatorname{ord}_p(h) \neq 0\}$. Then, since $\phi_{\mathcal{G}} = \phi_{\mathcal{F}}$, we have that $[f_0, \dots f_n] = [g_0, \dots, g_n]$ and since $f_0(q) = \lambda g_0(q)$ for any $q \in U$, we have that $f_i = \lambda g_i$ on U for all $0 \leq i \leq n$. Thus, since the set of zeroes and poles of a meromorphic function is discrete and U is our Riemann surface except the zeroes and poles of $2n + 2$ meromorphic functions in $\mathcal{F} \cup \mathcal{G}$, U is a dense set. So, since $f_i = \lambda g_i$ on U for all $0 \leq i \leq n$ on a dense set, by the Identity theorem [5.2], we must have that λ can be extended to a meromorphic function on X and $f_i = \lambda g_i$ for any $0 \leq i \leq n$.

□

Using this proposition, we can associate with each holomorphic map ϕ a linear system $|\phi|$ which will be well-defined, independent of the choice of meromorphic functions used to represent ϕ. We proceed as follows.

Let $\phi : X \to \mathbb{P}^n$ be a holomorphic map and let $\mathcal{F} = (f_0, \dots, f_n)$ be a sequence of meromorphic functions defining the map. Let $D_{\mathcal{F}} \in \operatorname{Div}(X)$ be the divisor given by $D_{\mathcal{F}}(p) = -\min_i(\operatorname{div}(f_i)(p))$ so that $V_{\mathcal{F}} := \operatorname{span}_{\mathbb{C}}(f_0, \dots, f_n)$ is a linear subspace of $L(D_{\mathcal{F}})$. By Proposition 7.16, the vector subspace $V_{\mathcal{F}}$ is associated to a linear system which we use to define $|\phi|$.

Definition 7.20. Given a holomorphic map $\phi : X \to \mathbb{P}^n$ with non-degenerate[13] image, i.e., ϕ does not factor through a map $\phi' : X \to \mathbb{P}^m$ for $m < n$ (or in simpler terms, the image of ϕ is not contained in any proper hyperplane of $\mathbb{P}^n$), we define the linear system $|\phi|$ as the linear system associated to $V_{\mathcal{F}}$ defined above.

Concretely, $|\phi| = \{\mathrm{div}(g) + D_{\mathcal{F}} : g \in V_{\mathcal{F}}\}$ which is well-defined, independently of $\mathcal{F}$, as if λ is a meromorphic function, $\{\mathrm{div}(h) + D_{\mathcal{F}} : h \in V_{\mathcal{F}}\} = \{\mathrm{div}(\lambda g) + D_{\lambda\mathcal{F}} : g \in V_{\mathcal{F}}\}$.

7.3.1 Base points

At this point, we've managed to attach a linear system $|\phi|$ to a holomorphic map $\phi : X \to \mathbb{P}^n$. Dually, one can also ask whether a given linear system Q on X is the linear system attached to a holomorphic map to $\mathbb{P}^n$. In fact, there is a unique obstruction to a linear system being associated with a holomorphic map: base points.

Definition 7.21. Let Q be a linear system on a compact Riemann surface X. A point $p \in X$ is said to be a *base point* of Q if for every divisor $E \in Q$, $E \geq p$.

The presence of base points for a linear system, Q, of divisors provides an obstruction to Q being the linear system attached to a holomorphic map to $\mathbb{P}^n$ due to the following.

Lemma 7.22. *Let X be a compact Riemann surface. Let $\phi : X \to \mathbb{P}^n$ be a holomorphic map. Then $|\phi|$ is a base-point-free linear system.*

Proof. Let's prove the lemma by contradiction. Suppose that $p \in X$ is a base point of the holomorphic map ϕ. Then, choosing an appropriate set of $n+1$ meromorphic functions, denoted $\mathcal{F}$, we can represent ϕ as $\phi = [f_0 : \cdots : f_n]$. Recall that when constructing the linear system associated to a holomorphic map, we used the divisor $D_{\mathcal{F}} = -\min_i\{\mathrm{div}(f_i)\}$. Suppose that the minimum order of the f_i's at the point p is obtained by the function f_j for $j \in \{0, \dots, n\}$. Then, we can create the divisor $E := \mathrm{div}(f_j) + D_{\mathcal{F}}$. This divisor will belong to the linear system $|\phi|$, but we have $E(p) = \mathrm{ord}_p(f_j) + D(p) = \mathrm{ord}_p(f_j) - \mathrm{ord}_p(f_j) = 0$. So E doesn't have p in its support, which means that p cannot be a base point.

□

7.3.2 Defining holomorphic maps with a divisor

Finally, we come to the culmination of our considerations of holomorphic maps and divisors: defining a holomorphic map via a linear system. We've just shown that a necessary condition for the existence of a holomorphic map attached to a linear system is the absence of base points. This condition is, in fact, sufficient. Moreover, the pairing between base point free linear systems and holomorphic maps to $\mathbb{P}^n$ is unique.

Proposition 7.23. *Let $Q \subset |D|$ be a base-point-free linear system on a compact Riemann surface X. Then, there is a unique (up to a choice of coordinates in $\mathbb{P}^n$) holomorphic map $\phi : X \to \mathbb{P}^n$, for some n such that $Q = |\phi|$.*

Proof. Suppose that the linear system Q corresponds to a vector subspace $V \subset L(D)$. By the one-to-one correspondence shown in Proposition 7.16, the elements of Q are of the form $\mathrm{div}(f) + D$ for $f \in V$. Then, since V is finite-dimensional, we can choose a basis $(f_0, \dots, f_n)$ for V. Since Q has no base points, $|\phi| = Q$. To show uniqueness, suppose that two holomorphic maps ϕ, φ have the same associated linear system, i.e. $|\phi| = |\varphi|$. Then, there exists sequences $\mathcal{F} = (f_0, \dots, f_n) \subset L(D_{\mathcal{F}})$ and $\mathcal{G} = (g_0, \dots, g_n) \subset L(D_{\mathcal{G}})$ such that, by changing coordinates for $\mathcal{G}$ if necessary, for all $0 \leq i \leq n$, $\mathrm{div}(f_i) + D_{\mathcal{F}} = \mathrm{div}(g_i) + D_{\mathcal{G}}$. Thus, the functions $h_i := g_i/f_i$ all have the same divisor;

$$\mathrm{div}(h_i) = \mathrm{div}(g_i) - \mathrm{div}(f_i) = D_{\mathcal{F}} - D_{\mathcal{G}}$$

So there exists $\alpha_i \in \mathbb{C}^*$ such that $g_i = \alpha_i f_i$. Consequently, $[g_0 : \cdots : g_n] = [\alpha_0 f_0 : \cdots : \alpha_n f_n]$ so if we perform the change of coordinates $x_i \mapsto \alpha_i^{-1} x_i$ on $\mathbb{P}^n$ for the map ϕ, the maps ϕ and φ are equal. □

[13] This condition is equivalent to the family $\mathcal{F}$ of functions used to define ϕ being linearly independent

The procedure outlined above to construct a holomorphic map ϕ from a divisor D required that it was base point free. However, it is still possible to construct a holomorphic map from a divisor with base points. The procedure goes as follows.

Definition 7.24. The fixed part F of a complete linear system $|D|$ is the divisor defined by $F(p) = \min\{E(p)|E \in |D|$. In other words, the fixed part of $|D|$ is the largest divisor "contained" in all divisors of the system $|D|$.

Lemma 7.25. *If F is the fixed part of the complete linear system $|D|$, then $L(D - F) = L(D)$. In particular, $L(D) = L(D - p)$ if and only if p is a base point of $|D|$.*

Proof. First, as $F \geq 0$, $D - F \leq D$, and so by Lemma 7.8, $L(D - F) \subset L(D)$. Conversely, suppose $f \in L(D)$, i.e. $\text{div}(f) + D \geq 0$. By minimality of F, this implies that $0 \leq F \leq \text{div}(f) + D$, so $\text{div}(f) + D - F \geq 0$, i.e., $f \in L(D - F)$. Altogether, $L(D - F) = L(D)$. The second statement follows from Proposition 7.10. □

Therefore, removing the fixed part of a divisor does not affect its corresponding space of functions.

Lemma 7.26. *Let D be a divisor and let F be its fixed part. Then, $|D - F|$ is base point free.*

Proof. We proceed by contradiction. Suppose $p \in X$ is a base point of $|D - F|$; that is, for all $f \in L(D - F)$, $(\text{div}(f) + D - F)(p) > 0$. From the previous lemma, this condition holds for all $f \in L(D)$, which means that $\min\{E(p)|E \in |D|\} > F(p)$, contradicting the defining property of F. Thus $|D - F|$ is base point free. □

This fact allows to attach a holomorphic map to any linear system as follows.

Definition 7.27. Let $|D|$ be a linear system of divisors on X and let F be the fixed part of $|D|$. We define the map $\phi_D : X \to \mathbb{P}^n$ as the map $\phi_D(x) = \phi_{D-F}(x)$, where $|D - F|$ is the associated base-point free divisor. Note that if D is base point free, $D - F = D$, so this definition is consistent with the one in Proposition 7.23.

7.3.3 Holomorphic embeddings into $\mathbb{P}^n$

Among the class of holomorphic maps to $\mathbb{P}^n$, it is natural to restrict our attention to embeddings, $\phi : X \hookrightarrow \mathbb{P}^n$, in which X is isomorphic to its image $\phi(X) \subset \mathbb{P}^n$. We make this definition more precise.

To do so, we first define, more precisely than in Definition 4.7, what a *holomorphic immersion* and a *holomorphic embedding* are.

Definition 7.28. Let $\phi : X \to \mathbb{P}^n$ be a holomorphic map. For any $p \in X$, let $\zeta : U \to \mathbb{C}$ and $\eta : V \to \mathbb{C}^n$ be local coordinate charts near $p \in U \subset X$ and $\phi(p) \in V \subset \mathbb{P}^n$, respectively. Assume furthermore that we order the homogeneous coordinates $f_0, \dots f_n$ so that f_0 does not vanish on some neighbourhood of p.

Let $z_0 := \zeta(p)$. The local representation of the map ϕ is hence given by $\tilde{\phi}(z) := \left(\eta \circ \phi \circ \zeta^{-1}\right)(z) = \left(\frac{f_1(z)}{f_0(z)}, \dots, \frac{f_n(z)}{f_0(z)}\right)$. The differential $d\phi_p : T_pX \to T_{\phi(p)}\mathbb{P}^n$ is locally identified by the map:

$$d\tilde{\phi}_{z_0} : T_{z_0}\mathbb{C} \to T_{\tilde{\phi}(z_0)}\mathbb{C}^n,\ d\tilde{\phi}_{z_0}(r) = r\left(\frac{d}{dz}\left(\frac{f_1}{f_0}\right), \frac{d}{dz}\left(\frac{f_2}{f_0}\right), \dots, \frac{d}{dz}\left(\frac{f_n}{f_0}\right)\right)\Bigg|_{z=z_0}$$

We then say that ϕ is a *holomorphic immersion* at p if the map $d\tilde{\phi}_{z_0}$ is injective, i.e. the derivatives of f_i/f_0 are not all zero at $z_0 = \zeta(p)$.

Also, ϕ is a *holomorphic embedding* if it is an immersion at each $p \in X$ and a homeomorphism onto its image. Since X is compact, and ϕ is a holomorphic map, it suffices to have that ϕ is an injection and an immersion at each point p.

Notice that if $\phi : X \to \mathbb{P}^n$ is a holomorphic embedding, then $\phi(X) \subset \mathbb{P}^n$ is a complex manifold and, X and $\phi(X)$ are isomorphic as complex manifolds.[14]

One can ask whether certain conditions can be imposed on a divisor D for the associated holomorphic map ϕ_D to be a holomorphic embedding. To alleviate this last part of this chapter, we will only treat the case where the complete linear system $|D|$ has no base points; the other case can be treated similarly using the fixed part of $|D|$ discussed above.

Proposition 7.29. *Let X be a compact Riemann surface and let D be a divisor such that $|D|$ has no base points. Then ϕ_D is an injective holomorphic map and an isomorphism onto its image, i.e. a holomorphically embedded Riemann surface in $\mathbb{P}^n$, if and only if for every p and q in X (p and q could be the same point!), we have* $\dim L(D-p-q) = \dim L(D) - 2$.

We proceed in two steps: first, showing the criterion for the map ϕ_D to be injective; second, showing the criteria for ϕ_D to be a holomorphic immersion at p.

Lemma 7.30. *Let X be a compact Riemann surface, p be a point of X, and $|D|$ be a base point free linear system. A basis $\{f_0, \dots, f_n\}$ of $L(D)$ satisfies* $\mathrm{ord}_p(f_i) > -D(p)$ *for all $1 \le i \le n$ and* $\mathrm{ord}_p(f_0) = -D(p)$ *if and only if $\{f_1, \dots, f_n\}$ is a basis of $L(D-p)$. Moreover, such a basis exists.*

Proof. First, as $|D|$ is base point free, from Lemma 7.25 and Proposition 7.10, $\dim L(D-p) = \dim L(D) - 1 = n$. Now, suppose, we have a basis $\mathcal{F} = \{f_1, \dots, f_n\}$ of $L(D-p)$. Consider any extension of $\mathcal{F}$ to a basis $\{f_0, \dots, f_n\}$ of $L(D)$. This extended basis satisfies $\mathrm{ord}_p(f_i) > -D(p)$ for all $1 \le i \le n$ since $f_i \in L(D-p)$. Moreover, since $f_0 \in L(D)$, $\mathrm{ord}_p(f_0) \ge -D(p)$. But, since $f_0 \notin \mathrm{span}\,(f_1, \dots, f_n) := L(D-p)$, we have that $\mathrm{ord}_p(f_0) \le -D(p)$. Thus, $\mathrm{ord}_p(f_0) = -D(p)$.

Conversely, suppose that a basis $\{f_0, \dots, f_n\}$ of $L(D)$ satisfies $\mathrm{ord}_p(f_0) = -D(p)$ and $\mathrm{ord}_p(f_i) > -D(p)$ for all $1 \le i \le n$. Then, $f_1, \dots, f_n \in L(D-p)$ and since they are linearly independent and $\dim L(D-p) = n$; they form a basis of $L(D-p)$. $\square$

Proposition 7.31. *Let X be a compact Riemann surface, $|D|$ a base point free linear system on X and $\phi_D : X \to \mathbb{P}^n$ its associated holomorphic map. Let $p, q \in X$ be two distinct points. Then, $\phi_D(p) = \phi_D(q)$ if and only if $L(D-p-q) = L(D-p) = L(D-q)$. Equivalently, by Lemma 7.10, ϕ_D is injective if and only if* $\dim L(D-p-q) = \dim L(D) - 2$.

Proof. By Lemma 7.30 applied at p, choose a basis $f_0, \dots, f_n$ of $L(D)$ with $\{f_1, \dots, f_n\}$ a basis of $L(D-p)$. In this basis, $\mathrm{ord}_p(f_0) = -D(p)$ and $\mathrm{ord}_p(f_i) > -D(p)$ for $i \ge 1$, so, by locally scaling the coordinates near q by $z^{D(q)}$,

$$\phi_D(p) = [f_0(p) : f_1(p) : \cdots : f_n(p)] = [1 : 0 : \cdots : 0].$$

Suppose that $\phi_D(p) = \phi_D(q)$. Then, $\phi_D(q) = \phi_D(p) = [f_0(p) : f_1(p) : \cdots : f_n(p)] = [1 : 0 : \cdots : 0]$. This means that for all $1 \le i \le n$, $\mathrm{ord}_q(f_i) > \mathrm{ord}_q(f_0)$. As $|D|$ is base point free and $f_0, \dots, f_n$ form a basis of $L(D)$, we must have $\mathrm{ord}_q(f_0) = -D(q)$ and thus the functions $f_1, \dots, f_n$ form a basis of $L(D-q)$. These functions, however, also form a basis of $L(D-p)$ by Lemma 7.30. Thus, since $L(D-p)$ and $L(D-q)$ share the same basis, they are equal. Moreover, since L(D-p-q) is a subspace of L(D-p) and L(D-q) by Lemma 7.8 and the functions $f_1, \dots, f_n$ also form a basis of $L(D-p-q)$, we have; $L(D-p) = L(D-q) = L(D-p-q)$.

Conversely, suppose that $L(D-p-q) = L(D-p) = L(D-q)$. Then, choosing any basis $(f_1, \dots, f_n)$ of $L(D-p)$, we have that it is also a basis of $L(D-q)$. By Lemma 7.30 applied at q, $\mathrm{ord}_q(f_0) = -D(q)$ and $\mathrm{ord}_q(f_i) > -D(q)$ for all $1 \le i \le n$. Hence, since f_0 has strictly smaller order at q than any other basis element, by locally scaling the coordinates near q by $z^{D(q)}$, $f_0(q) \ne 0$ and $f_i(q) = 0$ for all $1 \le i \le n$. By the same reasoning for p, we arrive at:

$$\phi_D(q) = [1 : 0 : \cdots : 0] = \phi_D(p).$$

$\square$

[14] This follows from the Inverse function theorem, giving a local holomorphic inverse to ϕ

Proposition 7.32. *Let $\phi_D : X \to \mathbb{P}^n$ be an injective holomorphic map where $|D|$ is a base point free complete linear system and let p be a point of X. Then, ϕ_D is a holomorphic immersion at p if and only if* $\dim L(D-2p) = \dim L(D) - 2$.

Proof. First, we show that $\dim L(D-2p) = \dim L(D) - 2$ is equivalent to $\dim L(D-2p) = \dim L(D-p) - 1$. By Lemma 7.25, since $|D|$ is base point free, $L(D) \neq L(D-p)$. But, by Lemma 7.10, this must mean that $\dim L(D-p) = \dim L(D) - 1$. Thus, we have that $\dim L(D-p) - 1 = \dim L(D) - 2 = \dim L(D-2p)$.

By Lemma 7.30, let $f_0, \dots, f_n$ be a sequence of representatives for ϕ_D such that $\mathrm{ord}_p(f_0) = -D(p)$ and, for all $1 \leq i \leq n$, $\mathrm{ord}_p(f_i) > -D(p)$.

Now, suppose that $\dim L(D-2p) = \dim L(D-p) - 1$. Since $L(D-2p) \subseteq L(D-p)$, by Lemma 7.8, there exists some $j \in \{1, \dots, n\}$ such that $f_j \in L(D-p) \setminus L(D-2p)$. Note that the converse is also true: by Lemma 7.10, if there is some $f_j \in L(D-p) \setminus L(D-2p)$, then $\dim L(D-2p) = \dim L(D-p) - 1$.

So, we have that $\dim L(D-2p) = \dim L(D-p) - 1$ if and only if here exists some $j \in \{1, \dots, n\}$ such that $f_j \in L(D-p) \setminus L(D-2p)$. But, the latter is equivalent to $\mathrm{ord}_p(f_j) = 1 - D(p)$, by definition of the space $L(D)$. Let z be a local coordinate near p given by the local chart $\zeta : X \to \mathbb{C}$. We then observe that;

$$\mathrm{ord}_p\left(\left(f_j(z) \cdot z^{D(p)}\right)'\right) = \mathrm{ord}_p(f_j) + D(p) - 1$$

Then, after locally scaling by $z^{D(p)}$, we have $\phi_D(p) = [1 : \cdots : 0]$, where the j^{th} coordinate has $\mathrm{ord}_p((f_j(z) \cdot z^{D(p)})' = 0$, i.e. the j^{th} coordinate has a simple zero at p. Thus, $(f_j(z)/z^{D(p)})'(\zeta(p)) \neq 0$. Since X has one complex dimension, this is equivalent to saying that the map $d\phi_p$ is injective at p, which means by definition that the map ϕ_D is a holomorphic immersion at p.

□

With this, we may name the type of divisor giving rise to a holomorphic embedding ϕ_D.

Definition 7.33. A divisor D such that $|D|$ has no base points and ϕ_D is an embedding is called a *very ample* divisor

We have finally shown the only condition for ϕ_D to be an embedding. We refer to Section 10.1 for some further direct applications related to embeddings.

8 Algebraic curves

The above discussion mainly viewed Riemann surfaces as 1-dimensional complex manifolds, as we defined them. However, to fully appreciate the Riemann-Roch theorem and have the tools to prove it, we need to view Riemann surfaces as algebraic curves (1-dimensional varieties over a field), a basic object of algebraic geometry. This is exactly the aim of this short section.

We first start by making an essential assumption: any compact Riemann surface has a sufficient number of meromorphic functions so that the following arguments can make sense. This is actually guaranteed for all compact Riemann surfaces, but is well beyond the scope of this chapter.[15]

Definition 8.1. Let S be a set of meromorphic functions on a compact Riemann surface X. We say that S separates points of X if for every pair of distinct points p and q in X there is a meromorphic function $f \in S$ such that $f(p) \neq f(q)$.

Definition 8.2. Let S be a set of meromorphic functions on a compact Riemann surface X. We say that S *separates tangents* of X if for every point $p \in X$ there is a meromorphic function $f \in S$ which has multiplicity one at p, i.e., by Lemma 6.15, f is holomorphic at p and has a simple zero, has a simple pole at p or $\mathrm{ord}_p(f - f(p)) = 1$.

[15] We refer any curious readers on the proof of this statement to the following source: Existence of meromorphic functions on Riemann surface.

We adopt a non-standard definition of an algebraic curve that differs from the standard one in algebraic geometry, where it refers to a 1-dimensional algebraic variety[16] over $\mathbb{C}$. In Section 10.1, we show that, for Riemann surfaces, our definition coincides with the standard one.

Definition 8.3. A compact Riemann surface X is an *algebraic curve* if the field $\mathcal{M}(X)$ of global meromorphic functions on X separates points and tangents.

Since all Riemann surfaces have enough meromorphic functions (so that the set of global meromorphic functions separates points and tangents), we have that every compact Riemann surface is an algebraic curve, in the sense of the previous definition.

Theorem 8.4. *Every compact Riemann surface is an algebraic curve.*

8.1 Laurent tail divisors and the space $L(D)$

Utilizing the Laurent expansions of functions in the space $L(D)$, we can construct a special map from meromorphic functions to a group of special divisors called *Laurent tail divisors*. It just so happens that on top of their essential role in describing the structure of the field of meromorphic functions $\mathcal{M}(X)$, these divisors allow us to discover a relation between the space $L(D)$ and the kernel of the map α_D presented below that will prove to be crucial for the statement and the proof of the Riemann-Roch theorem.

We first define more precisely the notion of a Laurent tail divisor and the group of said divisors, denoted by $\mathcal{T}(X)$.

Definition 8.5. A Laurent tail divisor on X is a finite formal sum defined as follows;

$$\sum_{p\in X} r_p(z_p)\cdot p$$

where $r_p(z_p)$ is a Laurent polynomial, i.e. it has finitely many terms, and z_p is a fixed local coordinate centred at p. Furthermore, the set of these divisors forms a group under pointwise addition and is denoted as $\mathcal{T}(X)$.

We can then further filter the group of divisors $\mathcal{T}(X)$ into the following subgroups: [17]

$$\mathcal{T}[D](X) = \{\sum_{p\in X} r_p(z_p)\cdot p \mid \deg(r_p) < -D(p)\ \forall p \in X \text{ with } r_p \neq 0\}$$

With this subgroup of divisors, we can define a truncation and a multiplication map that are presented below:

Definition 8.6. Let D_1, D_2, D be divisors with $D_1 \leq D_2$. Then, we can define the following two natural truncation maps:

$$t_D : \mathcal{T}(X) \to \mathcal{T}[D](X) \text{ and } t_{D_2}^{D_1} : \mathcal{T}[D_1](X) \to \mathcal{T}[D_2](X)$$

These maps act pointwise. The first map, t_D, acts on a Laurent tail divisor $\sum r_p(z_p)\cdot p$ by removing all terms in the Laurent series $r_p(z_p)$ of order greater or equal to $-D(p)$. Similarly, the second map, $t_{D_2}^{D_1}$, removes all terms of order greater or equal to $-D_2(p)$, which makes sense since we assume that $D_2 \geq D_1$.

Definition 8.7. Let f be a meromorphic function and D be any divisor, then we define the multiplication map :

$$\mu_f^D : \mathcal{T}[D](X) \to \mathcal{T}[D - \mathrm{div}(f)](X)$$

by multiplying each finite Laurent polynomial, r_p, by f so that $\sum_p r_p \cdot p$ is sent to the truncation of $\sum (fr_p)\cdot p$ with respect to $D - \mathrm{div}(f)$. Note that this map is an isomorphism since it has as inverse the map $\mu_{1/f}^{D-\mathrm{div}(f)}$.

[16] A 1-dimensional projective algebraic variety is a 1-dimensional space cut out by polynomial equations in projective space.

[17] The degree of the Laurent polynomial r_p refers to its highest power in this case.

With the above definitions, we can finally present the map, α_D, mentioned at the beginning of this section.

Definition 8.8. Fixing any divisor D on X, we define the map:

$$\alpha_D : \mathcal{M}(X) \to \mathcal{T}[D](X)$$

where we send a meromorphic function f to the sum $\sum_p r_p \cdot p$, where r_p is the truncated Laurent expansion of f at z_p, with all the terms of order equal to or higher than $-D(p)$ removed.

It is worth noting that this map commutes with the truncation map above, in the sense that $\alpha_{D_2} = t_{D_2}^{D_1} \circ \alpha_{D_1}$.

With this map now defined, the kernel of the map $\ker(\alpha_D)$, consists of all meromorphic functions whose Laurent expansion at each point $p \in X$ doesn't contain any term of order lower than $-D(p)$. This is exactly the definition of the space $L(D)$ given in Definition 7.7. So, we obtain the following correspondence;

$$L(D) = \ker(\alpha_D).$$

8.2 The space $H^1(D)$

One fundamental problem in the function theory of any algebraic curve X is the *Mittag-Leffler* problem, which relates to knowing whether a Laurent tail divisor $Z \in \mathcal{T}[D](X)$ is in the image of α_D. In other words, it tries to find a globally meromorphic function with the exact Laurent tails given by Z. Algebraically, we can actually create a space, the $H^1(D)$ space, that will serve as a measure of the failure of being able to solve the *Mittag-Leffler* problem for a given divisor D and a given compact Riemann surface X.

Definition 8.9. Let D be any divisor. Then, the space $H^1(D)$ is defined as follows:

$$H^1(D) = \operatorname{coker}(\alpha_D) := \mathcal{T}[D](X) \big/ \operatorname{image}(\alpha_\mathrm{D})$$

where the cokernel of any map is defined as the quotient of the target space by the image.

We can then see that a Laurent tail divisor Z is in the image of α_D if and only if it is in the same equivalence class as 0 in the quotient space $H^1(D)$.

To further analyze the space $H^1(D)$ and be able to derive a Lemma that is crucial to the first form of the Riemann-Roch theorem, we need to present the concept of a sequence and, more precisely, of a short exact sequence.

Definition 8.10. A sequence of $\mathbb{C}$-linear maps is written as follows :

$$U_1 \xrightarrow{a_1} U_2 \xrightarrow{a_2} \ldots \xrightarrow{a_n} U_{n+1}$$

where U_i are vector space over $\mathbb{C}$ and a_i are $\mathbb{C}$-linear maps.

Such a sequence is said to exact at U_i (for $1 < i < n$) if $\ker(a_i) = \operatorname{image}(a_{i-1})$. Moreover, such a sequence is said to be a short exact sequence if it is of the form;

$$0 \to U_1 \to U_2 \to U_3 \to 0$$

Example 8.11. For any divisor D on a compact Riemann surface X, we have the following exact sequence:

$$0 \to L(D) \xrightarrow{\iota} \mathcal{M}(X) \xrightarrow{\alpha_D} \mathcal{T}[D](X) \xrightarrow{\pi} H^1(D) \to 0$$

The first map in the sequence is the unique zero map from the zero vector space to $L(D)$, $0 : 0 \to 0_{L(D)}$, and the final map is the unique projection from $H^1(D)$ to the zero vector space. The ι map is an inclusion map, $\iota : L(D) \hookrightarrow \mathcal{M}(X), f \to f$. The map α_D, defined previously, sends a meromorphic function to its truncated Laurent tail divisor. Finally, the map π is the canonical projection of $\mathcal{T}[D](X)$ onto to the quotient space $\mathcal{T}[D](X) \big/ \operatorname{image}(\alpha_D)$.

We claim that this sequence is exact. First, $\operatorname{image}(0) = 0 = \ker(\iota)$, since any meromorphic function bounded by D is sent to itself, but in the vector space of all meromorphic functions, so the map must be injective, i.e. $\ker(\iota) = 0$. Then, we know that the image of $\operatorname{image}(\iota) = L(D)$ and we saw previously that $\ker(\alpha_D) = L(D)$, so we have that $\ker(\alpha_D) = \operatorname{image}(\iota)$. Then, we have that the kernel of the projection map π is $\operatorname{image}(\alpha_D)$, since the quotient space to which it maps is $H^1(D) := \mathcal{T}[D](X) \big/ \operatorname{image}(\alpha_D)$. Finally, the image of π is simply the space $H^1(D)$, while the kernel of the map onto the zero space is the domain of the map, which is $H^1(D)$.

Moreover, we can write the above exact sequence as a short exact sequence by using the first isomorphism theorem and the quotient space $\mathcal{M}(X) \big/ L(D)$ to obtain a sequence starting and ending by the vector space $\{0\}$ and only consisting of five vector spaces and four maps:

$$0 \to \mathcal{M}(X) \big/ L(D) \to \mathcal{T}[D](X) \to H^1(D) \to 0$$

We will now present an essential lemma, without proof, the *Snake Lemma*, coming from the theory of *homological algebra*. It will be our main tool to analyze and compare H^1 spaces defined by different divisors.

Lemma 8.12 (Snake Lemma)**.** *Let the following be a commutative diagram of abelian groups, with rows that form exact sequences.*

$$\begin{tikzcd}
0 \arrow[r] & A \arrow[r, "f"] \arrow[d, "a"] & A' \arrow[r, "f'"] \arrow[d, "b"] & A'' \arrow[r] \arrow[d, "c"] & 0 \\
0 \arrow[r] & B \arrow[r, "g"] & B' \arrow[r, "g'"] & B'' \arrow[r] & 0
\end{tikzcd}$$

Then there exists an exact sequence

$$0 \longrightarrow \ker a \longrightarrow \ker b \longrightarrow \ker c \xrightarrow{\delta} \operatorname{coker} a \longrightarrow \operatorname{coker} b \longrightarrow \operatorname{coker} c \longrightarrow 0$$

where δ is a homomorphism, called the connecting homomorphism.

Now, using the sequence presented in [8.11], we can construct the following commutative diagram where D_1 and D_2 are divisors on X and $D_2 \geq D_1$, so that the truncation map $t := t_{D_2}^{D_1}$ is defined. Recall that $L(D_1) \subseteq L(D_2)$ by Lemma 7.8.

$$\begin{tikzcd}
0 \arrow[r] & \mathcal{M}(X)/L(D_1) \arrow[r, "\alpha_{D_1}"] \arrow[d, "\zeta"'] & \mathcal{T}[D_1](X) \arrow[r, "c_1"] \arrow[d, "\iota"] & H^1(D_1) \arrow[r] \arrow[d, "\eta"] & 0 \\
0 \arrow[r] & \mathcal{M}(X)/L(D_2) \arrow[r, "\alpha_{D_2}"] & \mathcal{T}[D_2](X) \arrow[r, "c_2"] & H^1(D_2) \arrow[r] & 0
\end{tikzcd}$$

By the snake lemma applied to this commutative diagram, we obtain the following exact sequence:

$$0 \longrightarrow \ker \zeta \longrightarrow \ker t \longrightarrow \ker \eta \xrightarrow{\delta} \operatorname{coker} \zeta \longrightarrow \operatorname{coker} t \longrightarrow \operatorname{coker} \eta \longrightarrow 0$$

However, since $L(D_1) \subseteq L(D_2)$ and $D_2 \geq D_1$, the vertical maps are surjective. Hence, the cokernels of the maps ζ, η, and t are zero vector spaces, since their images are equal to their target spaces.

Hence, we can simplify the above long exact sequence, since $\operatorname{coker}(\eta) = 0$, to the following short exact sequence

$$0 \longrightarrow \ker \zeta \longrightarrow \ker t \longrightarrow \ker \eta \to 0$$

The advantage is that these kernels are equal to spaces we have already presented, and we know, by the property of exact sequences, what the kernel and image of the maps between them are. This will hence allow us to find a relation between the dimensions of these kernels.

First, we see that obviously $\ker(\zeta) = L(D_2)/L(D_1)$, so: $\dim\ker(\zeta) = \dim L(D_2) - \dim L(D_1)$, since the space $L(D)$ for any divisor D is finite dimensional, by Corollary 7.11.

We now compute the kernel of the truncation map t. Recalling the definition of $t := t_{D_2}^{D_1}$, a Laurent polynomial r_p lies in the kernel of t at p if and only if all its nonzero terms have order$k$ satisfying $-D_2(p) \leq k < -D_1(p)$. The monomials $(z_p)^k$ form a basis of the kernel at p. Thus summing this contribution over all $p \in X$;

$$\dim\ker(t) = \sum_{p\in X}(D_2(p) - D_1(p)) = \sum_{p\in X} D_2(p) - \sum_{p\in X} D_1(p) := deg(D_2) - deg(D_1)$$

Finally, let us denote the kernel of the map η between the space $H^1(D_1)$ and $H^1(D_2)$ as: $\ker(\eta) = H^1(D_1/D_2)$. In this case, we see that the short exact sequence can be written as :

$$0 \to L(D_2)/L(D_1) \xrightarrow{f} \ker(t) \xrightarrow{g} H^1(D_1/D_2) \to 0$$

Using the fact that this sequence is exact, we must have that $\ker(f) = 0$. Moreover, since f is a linear map between vector spaces, we have by the rank nullity and the property of exact sequences that:

$$\begin{aligned}\dim(\ker(g)) &= \dim(\mathrm{Im}(f)) \\ &= \dim(L(D_2)/L(D_1)) - \dim(\ker(f)) \\ &= \dim(L(D_2)) - \dim(L(D_1))\end{aligned}$$

Moreover, since $H^1(D_1/D_2)$ is mapped to the zero vector space, the kernel of that trivial map must be its whole domain. So, by exactness of the sequence, $\mathrm{Image}(g) = H^1(D_1/D_2)$. With this, we can finally, using once again rank nullity and the property of exact sequences, obtain that:

$$\begin{aligned}&\dim\ker(t) = \dim\ker(g) + \dim\mathrm{Im}(g) \\ \implies &\deg(D_2) - \deg(D_1) = \dim(L(D_2)) - \dim(L(D_1)) + \dim H^1(D_1/D_2)\end{aligned} \tag{1}$$

It is worth noting that this also shows that the space $H^1(D_1/D_2)$ is finite-dimensional, since its dimension is the sum of the dimensions of finite-dimensional vector spaces.

8.2.1 Finite dimensionality of the space $H^1(D)$

To conclude this section, we show that $H^1(D)$ is finite-dimensional.

Proposition 8.13. *Let X be an **algebraic curve**. There is some divisor A_0 such that the vector space $H^1(A_0)$ is finite dimensional*

Proof. [Sketch] We will only provide a comprehensive sketch of the full proof shown in [Mir95, p.182-184].

First, using the fact that X is an algebraic curve, we fix some non-constant global meromorphic function f on X. Let $D = \mathrm{div}_\infty(f)$. Then, for large m, the dimension of $H^1(0/mD)$ is independent of f and of m and is finite. This statement is proven using some uniform bound arguments and the Laurent Series Approximation Lemma, which is as follows;

Lemma 8.14 (Laurent Series Approximation)**.** *Suppose that X is an algebraic curve. Fix a finite number of points $p_i, \ldots, p_n$ in X, choose a local coordinate z_i at each p_i, and finally choose Laurent polynomials $r_i(z_i)$ for each i. Then there is a global meromorphic function f on X such that for every i, f has r_i as a Laurent tail at p_i.*

The above lemma relies heavily on the fact that algebraic curves separate tangents, justifying why X needs to be an algebraic curve. Then, using the dimension relation (1) obtained from Snake lemma, for the divisors 0 and mD, we have :

$$\deg(mD) - \dim L(mD) = \dim H^1(0/mD) - 1$$

So, since $H^1(0/mD)$ eventually has constant dimension as we increase m, and it can also be proven that its dimension increases when m increases, we have that $H^1(0/mD)$ is finite dimensional for any $m \geq 0$. This implies that there is some M such that $\deg(mD) - \dim L(mD) \leq M$ for the divisor of the poles of any meromorphic function.

Now, we fix some non-constant global meromorphic function f on X and let M be the upper bound of $\deg(mD) - \dim L(mD)$ for any $m \geq 0$. Take any divisor A on X and let S_A denote the support of A. By letting,

$$g = \Pi_{i \in S_A}(f - f(p_i))^{A(p_i)}$$

one can prove that for some $m_A \geq 0$, we have that $B = A - \mathrm{div}(g) \leq m_A D$. Then, from the dimension equality proven from the Snake Lemma, (1), we can show that $\deg(A) - \dim L(A) \leq M$. Since $\deg(A) - \dim L(A)$ can only take on integer values, we must have that there is some divisor A_0 with maximal $\deg(A_0) - \dim L(A))$.

Then, one can prove that $H^1(A_0) = 0$ by contradiction. The argument goes like this:

If, $H^1(A_0) \neq 0$, then there is a Laurent tail divisor $Z \in \mathcal{T}[A_0](X)$ that isn't of the form $\alpha_{A_0}(f)$. We can then increase A_0 to a divisor B so that $t(Z) = 0$, where t is the truncation map from A_0 to B. Thus, the equivalence class of Z, under the quotient $H^1(A_0)$, is in the kernel $H^1(A_0/B)$. So, since Z doesn't correspond to $\alpha_{A_0}(f)$, which is equivalent to zero in$H^1(A_0/B)$, we have that $\dim H^1(A_0/B) \geq 1$ *Thus, using 1, we have that*

$$1 \leq \dim H^1(A_0/B) = (\deg(B) - \dim L(B)) - (\deg(A_0) - \dim L(A_0))$$

But, $\deg(A_0) - \dim L(A_0)$ *is maximal so* $\dim H^1(A_0/B) < 0$. Thus, we would arrive at a contradiction and hence prove $H^1(A_0) = 0$.

□

Then, we show that $H^1(D)$ is finite dimensional for any divisor D. Notice that we may write the following: $D - A_0 = P - N$, where P, N are non-negative divisors with disjoint support. Since $A_0 \leq A_0 + P$, $\mathcal{T}[A_0](X) \to \mathcal{T}[A_0 + P](X)$ is a surjection, and hence, the induced map between the equivalence classes $H^1(A_0)$ and $H^1(A_0 + P)$ is also surjective. So, since $H^1(A_0) = 0$, by the rank nullity theorem, $H^1(A_0 + P) = 0$. So, the map $H^1(A_0 + P - N) \to H^1(A_0 + P)$ has kernel $H^1(A_0 + P - N/A_0 + P) = H^1(A_0 + P - N)$ since all elements are mapped to the null vector space $H^1(A_0 + P) = 0$.Since $D = A_0 + P - N$, we have proven that $H^1(D)$ is finite-dimensional.

Since the space $H^1(D)$ is finite-dimensional, we can write Equation 1 in a slightly clearer way. Since the map $d : H^1(D_1) \to H^1(D_2)$ is surjective and it's a linear map between vector spaces, we have, by rank nullity, that :

$$\dim H^1(D_1/D_2) = \dim \ker(d) = \dim H^1(D_1) - \dim H^1(D_2)$$

So we conclude this section by summarizing the computations above:

Lemma 8.15. *Suppose that D_1 and D_2 are some divisors on a compact Riemann surface X, with $D_1 \leq D_2$. Then,*

$$\dim L(D_2) - \dim H^1(D_2) - \deg(D_2) = \dim L(D_1) - \dim H^1(D_1) - \deg(D_1)$$

9 Riemann Roch theorem

We now finally reach the central result of this paper: the Riemann–Roch theorem. It provides a precise and powerful tool for computing the dimensions of spaces of meromorphic functions associated with divisors, $L(D)$.

9.1 First form of Riemann-Roch

The first form of the Riemann-Roch theorem is simply a generalization of Lemma 8.15 for any two divisors on an algebraic curve X, without the need to have that $D_1 \leq D_2$. To do so, given two arbitrary divisors D_1 and D_2 on an algebraic curve X, we can construct a divisor $D_{max} = \sum_p max\{D_1(p), D_2(p)\} \cdot p$. It is easy to see that $D_{max} \geq D_1, D_2$, so we can apply Lemma 8.15 to D_1 and D_{max}, and to D_2 and D_{max}. We then obtain that:

$$\begin{aligned} \dim L(D_i) - \dim H^1(D_i) - \deg(D_i) &= \dim L(D_{\max}) - \dim H^1(D_{\max}) - \deg(D_{\max}), \text{ for } i = 1, 2 \\ \implies \dim L(D_2) - \dim H^1(D_2) - \deg(D_2) &= \dim L(D_1) - \dim H^1(D_1) - \deg(D_1) \end{aligned}$$

So, we see that the quantity $\dim L(D) - \dim H^1(D) - \deg(D)$ is indeed constant for all divisors. The first form of the Riemann-Roch theorem consists, then, in comparing the value of the above quantity for any arbitrary divisor with that of the 0 divisor. We can actually find a simplified expression in the case of the zero divisor since $\deg(0) = \sum_p 0(p) = 0$ and $\dim L(0) = \dim \mathbb{C} = 1$, where in the latter we use the fact that the only global holomorphic functions on X are the constant functions, so $L(0) \cong \mathbb{C}$. Hence, we have that

$$\dim L(0) - \dim H^1(0) - \deg(0) = 1 - \dim H^1(0)$$

which allows us to then state the first form of the Riemann-Roch theorem for any algebraic curve, recalling that a compact Riemann surface is an algebraic curve, as stated in Theorem 8.4.

Theorem 9.1. ***(The Riemann-Roch Theorem: First form).***
Let D be a divisor on an algebraic curve X. Then,

$$\dim L(D) - \dim H^1(D) - \deg(D) = 1 - \dim H^1(0)$$

The problem with this first form of the Riemann-Roch theorem is that we simply shifted our need to compute the dimension of $L(D)$ to the computation of $\dim H^1(0)$ and $\dim H^1(D)$. This theorem only becomes actually useful when we introduce its second form. However, for that we need to identify the H^1's and to do so we will need Serre duality.

9.2 Serre Duality

First, before discussing Serre duality, we need to define residue maps. Recall why we defined the $H^1(D)$ space: to measure the failure of a Laurent tail divisor, Z, to be realized by a global meromorphic function on an algebraic curve X. So, we want to know if there is a meromorphic function f such that $\alpha_D(f) = Z$. We first consider $D = 0$. In this case, we look for a meromorphic function with principal parts[18] prescribed at each point by Z and that is holomorphic outside the support of Z. If we suppose such an f exists, writing $\alpha_0(f) = \sum r_p \cdot p$, for any given holomorphic 1-form ω, the forms $f\omega$ and $r_p\omega$ have the same principal part at p. Using the Residue theorem, we can then see, since $f\omega$ is a meromorphic 1-form on a compact Riemann surface, that we obtain:

$$\sum_p \mathrm{Res}(f\omega) = 0 \implies \sum_p \mathrm{Res}(r_p\omega) = 0$$

Therefore, the above is a necessary condition on $Z = \sum_p r_p \cdot p$ to find a function f such that $Z = \alpha_0(f)$. The Serre Duality theorem will show that this condition is, in fact, also sufficient.

Suppose that D is some divisor on X and that $\omega \in L^{(1)}(-D)$. Then, by the definition of $L^{(1)}(-D)$:

$$\mathrm{div}(\omega) \geq -(-D) \implies \mathrm{ord}_p(\omega) \geq D(p) \ \forall p \in X$$

Thus, we see that the lowest order term of the Laurent series expansion of ω at p is at least $D(p)$. So, for any $p \in X$, we may write ω in terms of local coordinate z at p and coefficients c_n, as:

[18] The principal part of a Laurent series is the finite sum of its negative order terms.

$$\omega = \left(\sum_{n=D(p)}^{\infty} c_n z^n \right) dz.$$

For any (global) meromorphic function, f, on X, its Laurent series expansion at $p \in X$, in terms of the same local coordinate as above, is given by: $f = \sum_{j=N_p}^{\infty} a_j z^j$, where $N_p \in \mathbb{Z}$. By direct computation of the residue of $f\omega$, we obtain:

$$\mathrm{Res}_p(f\omega) = \sum_{n=D(p)}^{\infty} c_n a_{-(n+1)}$$

Since the meromorphic 1-form ω is fixed and only the function f is variable, the residue at each point only depends on the coefficients a_i where $N_p \leq i < -D(p)$, meaning that these coefficients can be encoded in the Laurent tail divisor given by $\alpha_D(f)$. We may then construct a residue map with respect to a given meromorphic 1-form to formalize the above discussion.

Definition 9.2. Let D be some divisor on X and fix some meromorphic 1-form $\omega \in L^{(1)}(-D)$. The residue map of ω is then defined as:

$$\mathrm{Res}_\omega : \mathcal{T}[D](X) \to \mathbb{C}, \;\; \mathrm{Res}_\omega \left(\sum_{p \in X} r_p \cdot p \right) = \sum_{p \in X} \mathrm{Res}_p(r_p \omega)$$

The residue of $f\omega$ being encoded by the Laurent tail divisor $\alpha_D(f)$ can hence be expressed as :

$$\mathrm{Res}_\omega(\alpha_D(f)) = \sum_p \mathrm{Res}_p(f\omega) = 0$$

where the last equality is given by the Residue theorem applied to the meromorphic 1-form $f\omega$. Notice also that the map Res_ω is linear since Res_p is linear. Hence, since the image of α_D vanishes under Res_ω, we can descend the map Res_ω to the following map :

$$\mathrm{Res}_\omega : \mathcal{T}[D](X) \big/ \operatorname{image}(\alpha_D) \to \mathbb{C}$$

Recalling that $H^1(D) := \mathcal{T}[D](X) \big/ \operatorname{image}(\alpha_D)$, we remark that Res_ω is a linear functional in the dual space of $H^1(D)$, denoted by $H^1(D)^*$. We may hence define **the** Residue map:

$$\mathrm{Res} : L^{(1)}(-D) \to H^1(D)^*, \;\; \omega \to \mathrm{Res}_\omega$$

It is with this map that we will be able to state Serre duality and extract the relation between the dimensions of the H^1 and $L^{(1)}$ spaces.

Theorem 9.3. ***(Serre Duality).***
For any divisor D on an algebraic curve X, the map

$$\mathrm{Res} : L^{(1)}(-D) \to H^1(D)^*$$

is an isomorphism of complex vector spaces. In particular, for any canonical divisor K on X,

$$\dim H^1(D) = \dim L^{(1)}(-D) = \dim L(K - D)$$

Proof. This proof is divided into two parts: proving the injectivity of the residue map and proving its surjectivity. We omit the proof of surjectivity [Mir95, p.188-191], and will only present the proof of injectivity to provide an example of how the residue map can be applied, and then show how the equality of the dimension arises.

To prove the map is injective, let D be a divisor on X and take any $\omega \in L^{(1)}(-D) \setminus \{0\}$, such that the linear functional $\mathrm{Res}(\omega) := \mathrm{Res}_\omega$ is the zero map on $H^1(D)$.

We proceed by contradiction. Recall that $H^1(D) = \mathcal{T}[D](X)/\operatorname{image}(\alpha_D)$. Since Res_ω is identically zero on this quotient space, it can be naturally lifted to the entire space of tail divisors, where it must be identically zero. That is:

$$\mathrm{Res}_\omega\left(\sum_{p\in X} r_p \cdot p\right) = \sum_{p\in X} \mathrm{Res}_p(r_p\omega) = 0 \ \text{ for all } \sum_{p\in X} r_p \cdot p \in \mathcal{T}[D](X)$$

Choosing a local coordinate z centred at a point $p \in X$, we have $\mathrm{ord}_p(\omega) \geq D(p)$ since $\omega \in L^{(1)}(-D)$. Define $j := \mathrm{ord}_p(\omega)$ and notice that $j \geq D(p)$. Writing the meromorphic 1-form in its Laurent series expansion centred at p, we have:

$$\omega = \left(\sum_{n=j}^{\infty} c_n z^n\right) dz$$

where $c_j \neq 0$, since the order of ω at p is j. Now, consider the tail divisor $J := z^{-1-j} \cdot p \in \mathcal{T}(X)$.Since, $j \geq D(p)$ implies that $j > D(p) - 1$, which in turn means that $-1 - j < -D(p)$, we have $J \in \mathcal{T}[D](X)$. Using the natural lift of Res_ω from $H^1(D)$ to $\mathcal{T}[D](X)$, we have that $\mathrm{Res}_\omega(J) = 0$.

However, $\mathrm{Res}_\omega(J) = \mathrm{Res}_p\left(z^{-1-j}\sum_{n=j}^{\infty} c_n z^n\right) = \mathrm{Res}_p\left(\sum_{n=-1}^{\infty} c_{j+1+n} z^n\right) = c_j \neq 0$

We therefore arrive at a contradiction. Since assuming there is a non-zero element in the kernel of the residue map leads to a contradiction, its kernel must be trivial, i.e. $\ker(\mathrm{Res}) = \{0\}$. Thus, the residue map is injective.

Using the injectivity of the residue map proven above, its surjectivity (see [Mir95, p.188-191]) and its linearity, we obtain that the residue map forms an isomorphism and hence $\dim H^1(D)^* = \dim L^{(1)}(-D)$. Since $H^1(D)^*$ is finite-dimensional, by the isomorphism between a finite-dimensional vector space and its dual, we have: $\dim H^1(D) = \dim H^1(D)^* = \dim L^{(1)}(-D)$. The equality $\dim L^{(1)}(-D) = \dim L(K - D)$ follows from Lemma 7.12. □

9.3 Second form of Riemann-Roch

Before proceeding to the second form of the *Riemann–Roch* theorem, we can appreciate the importance of the Serre Duality theorem by using it conjointly with the first form of the *Riemann–Roch* theorem to show an interesting relation between the genus and the dimension of the H^1 space.

First, we know from Lemma 7.13 that if K is a canonical divisor, $\deg(K) = 2g - 2$. Then, applying the Serre duality to the divisor K itself, we have that $\dim H^1(K) = \dim L^1(-K) = \dim L(0) = 1$, where the last equality follows from Example 5.7. Moreover, applying the Serre Duality to the zero divisor instead, we have that $\dim H^1(0) = \dim L^1(0) = \dim L(K)$. Finally, if we apply the first form of the Riemann-Roch theorem to K, we have that :

$$\begin{aligned} \dim L(K) - \dim H^1(K) &= \deg(K) + 1 - \dim H^1(0) \\ \implies \dim H^1(0) - 1 &= 2g - 2 + 1 - \dim H^1(0) \\ \implies \dim H^1(0) &= g. \end{aligned}$$

This identity is of great importance because it unifies three different notions of genus:

- The topological genus, g, which is defined as the number of "holes" in the surface; it's a purely topological concept, explaining why it is given the name of topological genus.
- The arithmetic genus, $\dim H^1(0)$, which measures the obstruction to the existence of a global meromorphic function with prescribed Laurent tails.

- The analytic genus, $\dim L^{(1)}(0)$, which is the dimension of the space of holomorphic 1-forms on X.

This shows that we have analytical and algebraic invariants for compact Riemann surfaces with the same topological genus. In higher dimensions, these perspectives generalize in more sophisticated ways, forming the basis for powerful theorems such as the *Hirzebruch–Riemann–Roch* theorem, which we will not present in this paper.

We may now present *la pièce de résistance* of this paper;

Theorem 9.4. ***(The Riemann-Roch Theorem: Second form).***
Let X be an algebraic curve (compact Riemann surface) of genus g. Then, for any divisor D and any canonical divisor K, we have :

$$\dim L(D) - \dim L(K - D) = \deg(D) + 1 - g$$

Proof. After all the work done above, the proof is trivial. Starting with the statement of the first form of the Riemann-Roch theorem, we have :

$$\dim L(D) - \dim H^1(D) - \deg(D) = 1 - \dim H^1(0)$$

We saw above that $\dim H^1(0) = g$ and that $\dim H^1(D) = \dim L^{(1)}(-D) = \dim L(K - D)$ for any canonical divisor K. Plugging those expressions into the first form of the *Riemann-Roch* theorem, we obtain the second form of said theorem. □

Now that we have proven the *Riemann–Roch* theorem, we must mention a corollary which greatly simplifies the computation of the dimension of $L(D)$, when the degree of D is sufficiently large.

Corollary 9.5. *Let D be a divisor on an algebraic curve X with* $\deg(D) \geq 2g - 1$*. Then,*

$$\dim L(D) = \deg(D) + 1 - g$$

Proof. First, notice that if $\deg(D) \geq 2g - 1$, then, $\deg(K - D) \leq -1$. But, as seen in Lemma 7.9, since $\deg(K - D) < 0$, we have that $\dim L(D - K) = 0$, which by direct computation using the second form of the *Riemann–Roch* theorem shows the above. □

10 Applications of the Riemann-Roch theorem

We conclude this paper by showcasing a few applications of the Riemann–Roch theorem. These examples will hopefully highlight just how far-reaching this fundamental result is in the study of **compact** Riemann surfaces.

10.1 Algebraic curves and projective space

The first application we present of the Riemann-Roch theorem is that any algebraic curve, and hence any compact Riemann surface, can be holomorphically embedded into projective space.

Indeed, using the Riemann-Roch theorem, we can give a criterion in terms of the degree of a divisor D, for the map ϕ_D to projective space, presented in Section 7.3.2, to be a holomorphic embedding.

Proposition 10.1. *Let X be an algebraic curve of genus g. Then, any divisor D on X with degree* $\deg(D) \geq 2g + 1$ *is very ample, meaning that $|D|$ has no base points and ϕ_D is a holomorphic embedding onto a smooth projective curve of degree equal to* $\deg(D)$*.*

Proof. Recalling Proposition 7.29, we know that it suffices to show that for every $p, q \in X$, we have that $\dim L(D - p - q) = \dim L(D) - 2$.

Let $p, q \in X$ be any points on an algebraic curve X. First, we know that, by hypothesis, $\deg(D-p-q) = \deg(D) - 2 \geq 2g - 1$ and $\deg(D) \geq 2g - 1$. So, applying Corollary 9.5 to both D and $D - p - q$ we have that :

$$\dim L(D-p-q) = \deg(D-p-q) + 1 - g \quad \text{and} \quad \dim L(D) = \deg(D) + 1 - g$$

Since $\deg(D-p-q) = \deg(D) - 2$, combining the above expressions, we obtain:

$$\dim L(D-p-q) = \dim L(D) - 2$$

□

With the above criterion for the map ϕ_D to be a holomorphic embedding into projective space, we can see that all algebraic curves are, in reality, smooth projective curves.

Proposition 10.2. *Every algebraic curve X can be holomorphically embedded into projective space.*

Proof. The proof is quite trivial. We simply need to construct a divisor D of degree at least $2g+1$. To do so, we can simply take $D = (2g+1) \cdot p$ for any point $p \in X$ and by the above proposition, we have that ϕ_D is a holomorphic embedding. □

10.1.1 The image of a holomorphic embedding into $\mathbb{P}^n$ can be given by a polynomial equation.

Now that we have that, all compact Riemann surfaces are holomorphically embedded in a projective space, we can go further and show their algebraicity, giving rise to algebraic tools that aren't usually available to general manifolds.

Let $\varphi : X \to \mathbb{P}^n$ be a holomorphic embedding of a compact Riemann surface X, which exists by Proposition 10.2. Let f be a non-constant meromorphic function on X. Then we have the following theorem, which is necessary to what we want to show;

Theorem 10.3. *[**NN01**] Given a non-constant meromorphic function $f : X \to \mathbb{C}_\infty$, there exists an integer $q > 0$ such that for any meromorphic function F on X, there exist single complex variable rational functions $a_1, ..., a_q : \mathbb{C}_\infty \to \mathbb{C}_\infty$ such that :*

$$F(z)^q + a_1(f(z)) \cdot F(z)^{q-1} + a_q(f(z)) = 0 \; \forall z \in X \tag{2}$$

Proof. To prove the above theorem, we will start by showing that $f : X \to \mathbb{C}_\infty$ is actually a covering map onto $\mathbb{P}^1$ and has fibers of constant size, which we will denote by q, except at finitely many ramification points of f. First, we can easily see that since f is meromorphic, it can be viewed as a holomorphic function from X to $\mathbb{C}_\infty$. However, as seen in Example 6.4, $\mathbb{C}_\infty$ and $\mathbb{P}^1$ are isomorphic, so we can view f as a map from X to $\mathbb{P}^1$. Furthermore, using the inverse function theorem and the complex charts of X, we can show that f is a local homeomorphism and hence is a covering map of X except at the finitely many ramification points of f.

Now, since f is a covering map except at ramification points, it has the same number of preimages away from ramification points. So, we first need to show that f has finitely many critical points. Notice that branch points of a function occur when $f'(x) = 0$. That locus must be finite since X is compact and ramification points of f are isolated for meromorphic functions, i.e. they can't accumulate.

Now, to finalize this proof, let's define the functions $a_k : \mathbb{C}_\infty \to \mathbb{C}_\infty$ for $1 \leq k \leq q$. Recall that we defined q as the constant number of sheets in the branched cover $f : X \to \mathbb{C}_\infty \cong \mathbb{P}^1$ away from finitely many ramification points. Note that this definition of q is equivalent to the degree of the holomorphic map f. Then, let $f^{-1}(w) = \{x_1(w), x_2(w), \ldots, x_q(w)\}$ be the preimages of $w \in \mathbb{P}^1 \setminus B$[19], where B is the finite set of branch points of f and the images of poles of F, under the map f. With this, we define the functions a_k as the following symmetric polynomials :

[19]The image of f is both open, by the open mapping theorem and closed, since X is compact, so since $\mathbb{P}^1$ is connected, f is surjective

$$a_k(w) = \sum_{1 \leq i_1 < i_2, \ldots, < i_k \leq q} F(x_{i_1}(w)) \cdot F(x_{i_2}(w)) \cdot \ldots \cdot F(x_{i_k}(w))$$

Since a_k is just a composition of a polynomial with locally holomorphic functions (the x_i), it will be holomorphic near any $w \in \mathbb{P}^1 \setminus B$. Thus, for any k, the composition $a_k \circ f$ is holomorphic on $\mathbb{P}^1 \setminus B$.

A defining property of symmetric polynomials is that if σ_k denote the symmetric polynomial of some set of variables $\{x_1, \ldots x_q\}$, then the following polynomial in T has the property:

$$\sum_{i=1}^{q} T^{(q-i)} \, (-1)^i \, \sigma_i(x_1, x_2, ..., x_q) = \prod_{i=1}^{q} (T - x_i)$$

For example, $\sigma_1(x_1, x_2, ..., x_q) = \sum_{i=1}^{q} x_i$ and $\sigma_q(x_1, x_2, ..., x_q) = \prod_{i=1}^{q} x_i$.

So, by applying this to $T = F(z)$ and to the preimages $f^{-1}(f(z)) = \{x_1(w), x_2(w), \ldots, x_q(w)\}$ for any $z \in X \setminus \{\text{ramification points of } f \text{ and poles of } F\}$, we obtain :

$$F(z)^q - a_1(f(z)) \cdot F(z)^{q-1} + \cdots + (-1)^q \, a_q(f(z)) = 0 \tag{3}$$

At this point in the proof, we just need to show that Equation 2 holds not only for the points where we know that $a_k \circ f$ is holomorphic, but for all points of X. Using Riemann's removable singularity theorem, one can prove that the a_k are meromorphic on $\mathbb{P}^1$. First, by Example 5.6, meromorphic functions on $\mathbb{C}_\infty \cong \mathbb{P}^1$ are rational functions. Since the left-hand side of (2) is a polynomial, it can be seen as a holomorphic map $X \to \mathbb{C}_\infty$ which is hence continuous. Thus, by continuity of (2), since (3) holds everywhere on X except at some isolated points, we must have that (3) holds for all $z \in X$.

□

Now, we remind ourselves of the goal of this section: to show that the image of a holomorphic embedding is given by polynomial equations.

To prove said result, since φ is a holomorphic embedding, locally, we may choose a suitable sequence $\{\bar{f}_0, \ldots, \bar{f}_n\}$ of global meromorphic functions f_i on X such that: $\varphi(z) = [\bar{f}_0(z), \ldots, \bar{f}_n(z)]$. To simplify further computations, since projective coordinates are scale invariant, we may divide by $\bar{f}_0$ all the $\bar{f}_i$ to obtain $\phi_D(z) = [1, f_1(z), \ldots, f_n(z)]$, where $f_i := \bar{f}_i/\bar{f}_0$ for any $i > 1$.

Pick any $\nu, \mu \in \{1, \ldots, n\}$, $\nu < \mu$ and apply Theorem 10.3 with $F = f_\nu$ and $f = f_\mu$. We will then obtain rational equations of the following form:

$$f_\nu(z)^q + a_1(f_\mu(z)) f_\nu(z)^{q-1} + \cdots + a_q(f_\mu(z)) = 0 \quad \text{for any } z \in X \tag{4}$$

Then, since the functions a_k are meromorphic on $\mathbb{P}^1 \cong \mathbb{C}^\infty$, and hence are rational functions, we may clear the denominators by multiplying each equation by the least common multiple of all the denominators appearing in said rational equation. Thus, since $\varphi(X) \subseteq \mathbb{P}^n$, we obtain a, not necessarily homogeneous, polynomial equation in terms of both $f_\nu(z)$ and $f_\mu(z)$;

$$P(f_\nu(z), f_\mu(z)) = 0 \;\; \forall z \in X \tag{5}$$

Now, using projective coordinates $[Z_0, Z_1, \ldots, Z_n]$ on $\mathbb{P}^n$, we have that f_i corresponds to Z_i/Z_0 for any $1 \leq i \leq n$. Thus, in projective coordinates, our system of polynomial equations can be written as a system of **rational** equations in terms of Z_0, Z_ν, Z_μ :

$$P\left(\frac{Z_\nu}{Z_0}, \frac{Z_\mu}{Z_0}\right) = 0 \tag{6}$$

Then, to obtain a polynomial equation of the form $P(Z_0, Z_\nu, Z_\mu) = 0$, we can multiply both sides of Equation (6) by the highest power of Z_0 in the denominator appearing in it. Note that the resulting polynomial equation will be **homogeneous**, with degree equal to that of the highest power of Z_0 in the denominator, since the terms of (6) have degree zero in terms of the local coordinates Z_i.

Thus, since we built these polynomials from equations that were satisfied by $f_i(z)$ $\forall z \in X$, we know that $\varphi(X)$ is in the locus of said polynomials.

Moreover, using the fact that the embedding φ strictly separates points, its coordinate functions uniquely identify every location on X, one can also show that the locus of the system of homogeneous polynomial equations is exactly $\varphi(X)$. Hence, the image $\varphi(X) \subset \mathbb{P}^n$ is the locus of the zeroes of a polynomial expression, showing the algebraicity of Riemann surfaces.[20]

10.2 Existence of Weierstrass equation for an elliptic curve

As was briefly mentioned in Section 4.3, compact Riemann surfaces of genus one are of central interest in fields such as algebraic geometry and number theory since any complex torus X is isomorphic (as a Lie group) to the set of complex points $E(\mathbb{C})$ of an elliptic curve defined over $\mathbb{C}$. Typically, such a curve is given as the set of solutions in $\mathbb{P}^2(\mathbb{C})$ of the homogenization of an equation in so-called *Weierstrass* form:

$$Y^2 = aX^3 + bX + c$$

where a, b, c are complex numbers and X, Y are variables. In the broader context of algebraic geometry, one would be correct to wonder why we restrict our study to such a specific type of curve among the seemingly larger class of curves of genus one. However, using the Riemann-Roch theorem, we can prove that in fact, every smooth projective curve of genus one may be given as the solution set in $\mathbb{P}^2(\mathbb{C})$ of an equation in *Weierstrass* form.

Theorem 10.4. *[Sil09] Let E be a smooth algebraic curve of genus one with a distinguished point $p \in E$. Then there exists an embedding $\phi : E \hookrightarrow \mathbb{P}^2(\mathbb{C})$ such that $\operatorname{Im}(X) \subset \mathbb{P}^2(\mathbb{C})$ is the solution set to the homogenization of an equation in Weierstrass form. In particular, $\phi(p) = [0 : 1 : 0]$ and the images of the rest of the points of E correspond to the solutions to the affine equation.*

Proof. We will separate this proof into 2 parts by first proving that the embedding of E into $\mathbb{P}^2$ is the set of solutions of an equation in *generalized Weierstrass form*:

$$Y^2 + aXY + bY = X^3 + dX^2 + eX + f$$

and then, using a change of variables, we will transform this equation into the desired *Weierstrass form*.
Given the distinguished point $p \in E$, we consider the divisor $D = p$. As $\deg(D) = 1 \geq 2g - 1$, for all $n \geq 1$, we can use corollary.9.5 to get

$$\begin{aligned} \dim L(n \cdot D) &= \deg(n \cdot D) + g - 1 \\ &= \deg(n \cdot D) + 1 - 1 \\ &= n \end{aligned}$$

Therefore, $L(2 \cdot D) = \operatorname{span}\{1, X\}$ where X is a meromorphic function with a pole of order 2 at p: it cannot have a pole of order 1 as $\dim L(D) = 1$. Similarly, $\dim L(3 \cdot D) = 3$ implies that $L(3 \cdot D) = \operatorname{span}\{1, X, Y\}$ where Y is a meromorphic function with a pole of order 3 at p. From those 3 functions, we may construct the following set of 7 meromorphic functions.

$$H = \{1, X, Y, Y^2, XY, X^2, X^3\}$$

Note that all functions in H have a single pole of order at most 6 at p, so $\operatorname{span}(H)$ is a subspace of $L(6 \cdot D)$. However, we've demonstrated that $\dim L(6 \cdot D) = 6$, which means that H is a linearly dependent set. Therefore, there exist complex numbers a, b, c, d, e, f, g, not all zero, such that

$$aY^2 + bXY + cY = dX^3 + eX^2 + fX + g.$$

[20] A comprehensive summary of this proof can be found in the following MathOverflow post: https://mathoverflow.net/a/410040

We note that neither a nor d can equal zero as if that were the case, every function with a non-zero term would have a different order at p, making the inequality impossible[21]. We therefore construct the map to projective space given by $\phi = [X : Y : 1]$. By construction, ϕ is a holomorphic map from E to C and is thus surjective by Proposition 6.5. To show that $\phi : E \to C$ is a holomorphic embedding, we note that ϕ is the holomorphic map associated to the divisor $D_3 := 3 \cdot p$. Since $\deg(D_3) = 3 \geq 2g+1$, Proposition 10.1 implies that D_3 is very ample and thus that, by Definition 7.33 of a very ample divisor, ϕ is a holomorphic embedding.

It remains only to show that the generalized Weierstrass equation for C can be transformed into "proper" Weierstrass form via a clever change of variables. However, we also wish to emphasize the strength of Proposition 10.1, which allows us to deduce that the map ϕ is a holomorphic embedding using an easy-to-check numerical criterion. Therefore, after transforming the equation for C, we provide a second, longer proof of this statement that requires less machinery.

Starting from the generalized form

$$aY^2 + bXY = cY = dX^3 + eX^2 + fX + g$$

We begin by completing the square on the right-hand side to get

$$(\sqrt{a}Y + \frac{bX+c}{2})^2 = dX^3 + (e + b^2/4)X^2 + (f + bc)X + \frac{c^2}{4} + g.$$

We now perform the transformation $X' := X + \dfrac{(e + b^2/4)}{3d}$ and note the following:

$$d(X')^3 d = d\left(X + \frac{e + b^2/4}{3d}\right)^3 = dX^3 + (e + b^2/4)X^2 + 3d\left(\frac{e + b^2/4}{d}\right)^2 X + d\left(\frac{e + b^2/4}{d}\right)^3$$

so that the cubic and quadratic terms in the previous equations match. Therefore, letting $A := f + bc - 3d\left(\frac{e+b^2/4}{d}\right)^2$ and $B := c^2/4 + g - d\left(\frac{e+b^2/4}{d}\right)^3$, we may express the equation in the form

$$(\sqrt{a}Y + \frac{bX+c}{2})^2 = dX'^3 + AX + B \tag{7}$$

We perform the final affine transformation $Y' := (\sqrt{a}Y + \dfrac{bX+c}{2})$ which finally yields the equation in Weierstrass form:

$$Y'^2 = dX'^3 + AX + B. \tag{8}$$

Note that, under the change of variables we made, X' and Y' respectively have the same orders at p as X and Y, and we may assume that $d = 1$ in (8). The curve C is non-singular if and only if the polynomial $F(x, y) = y^2 - x^3 - Ax - B$ is non-singular. Letting $P(x) := x^3 + Ax + B$ and computing derivatives, we find that singular points of F are of the form $(\alpha, 0)$, where α is a root of both $P(x)$ and $P'(x)$. A polynomial in a single variable shares a root with its derivative if and only if it is inseparable, i.e., if it has a root of multiplicity greater or equal to 2; we show that P has no multiple root [Mir95][p. 153, Exercise V.3.F.11]. Suppose $\alpha \in \mathbb{C}$ is a multiple root of $P(x)$ and consider the meromorphic function $f := \dfrac{Y'}{(X' - \alpha)}$ which has a pole of order one at p. Taking its square we find that

$$\frac{Y'^2}{(X' - \alpha)^2} = \frac{P(X')}{(X' - \alpha)^2} = (X' - \beta)$$

where β is the third root of $P(x)$. Thus, $f^2 \in L(2 \cdot p)$, which means that $f \in L(p)$. However, $L(p)$ is of dimension one, so f must be a constant function, contradicting the assertion that it has a pole of order one at p. The polynomial $P(x)$ is thus separable, which means that C is a non-singular curve, and thus that $\phi : E \to \mathbb{P}^2$ is a holomorphic embedding into projective space. This completes the proof.

□

[21] We would be able to prove that a function with a pole of order $n+1$ at p is contained in $L(n \cdot D)$ which is a contradiction.

References

[Mir95] Rick Miranda. *Algebraic Curves and Riemann Surfaces*, volume 5 of *Graduate Studies in Mathematics*. American Mathematical Society, Providence, RI, 1995.

[NN01] Raghavan Narasimhan and Yves Nievergelt. *Complex Analysis in One Variable*. Springer Book Archive. Birkhäuser, Boston, MA, 2 edition, 2001.

[Sil09] Joseph H. Silverman. *The Arithmetic of Elliptic Curves*, volume 106 of *Graduate Texts in Mathematics*. Springer, New York, NY, 2 edition, 2009.

[SS03] Elias M. Stein and Rami Shakarchi. *Complex Analysis*, volume 2 of *Princeton Lectures in Analysis*. Princeton University Press, Princeton, NJ, 2003.