RIEMANN SURFACES, BRANCHED COVERINGS, AND FIELD EXTENSIONS

ZACH MARKS

Abstract. This paper discusses and proves a correspondence between branched coverings of compact Riemann surfaces and field extensions of their function fields. In addition to proving this relationship, we establish a few consequences of this result that help characterize those surfaces.

Contents

1. Introduction

The central theorem of Galois Theory is Galois' famous correspondence, proving a direct analogy between Galois groups and field extensions. In doing so, he showed that understanding one problem was equivalent to understanding a (seemingly) very different one - to enumerate intermediate field extensions is the same as to enumerate subgroups of a Galois group. Such correspondences between seemingly disparate areas are found in many other places, often providing new insight into a problem by correlating it to another. In this paper, we seek to work similarly, proving a correspondence between field extensions and covering spaces, as well as some direct consequences.

2. What is a Branched Cover? What is a Riemann Surface? AKA Introductory Results

We begin by discussing some preliminary results of topology.

Definition 2.1. A connected one-dimensional complex manifold is called a Riemann Surface.

Certain authors say that a Riemann surface does not need to be connected, but for our purposes we shall consider one to be.

Date: August 29, 2022.

Example 2.2. The following are all examples of Riemann Surfaces:
The complex plane, $\mathbb{C}$
The extended complex plane, AKA the Riemann Sphere, $\hat{\mathbb{C}} := \mathbb{C} \cup \{\infty\}$
A hyperelliptic surface, $\mathcal{C} := \{(x, y) | y^2 = x^5 - x\}$

Definition 2.3. A continuous map $p : Y \to X$ is said to be a covering of X if for every $x \in X$, there exists an open neighborhood U of x such that $p^{-1}(U)$ is the disjoint union of open sets $\{V_i\}_{i \in \mathbb{R}}$, where p restricted to V_i is a homeomorphism from V_i to U.

In other words, a covering p of X maps some neighborhood around each point from disjoint copies of that same neighborhood in the preimage.

Proposition 2.4. *If X is connected, and $p : Y \to X$ is a covering of X, then the cardinality of $p^{-1}(x)$ is the same over all $x \in X$.*

Proof. Pick $x \in X$. Then there exists an open set $U \ni x$ such that $p^{-1}(U)$ is the disjoint union of (potentially infinitely many) open sets $\{V_i\}$ in Y. We can thus define functions f_i such that $f_i : V_i \to U$ is a homeomorphism. In particular, $f_i = p|_{V_i}$. We then note that if $x' \in U$, there is an open set $U' \ni x'$ such that $p^{-1}(U')$ is the disjoint union of (potentially infinitely many) open sets $\{V_i'\}$ in Y, with similarly defined homeomorphisms $f_i = p|_{V_i'}$. Let $U'' = U \cap U'$ be an open set in X. Now, because U'' is a subset of U, the homeomorphisms f_i can be inverted over U'', letting us define $V_i'' = f_i^{-1}(U'')$. This implies that $p^{-1}(U'')$ is expressible as the disjoint union of $\{V_i''\}$, and in particular

$$|\{V_i''\}| = |\{V_i\}| = |p^{-1}(x)|.$$

However, we note that we can just as easily say the same for U', since U'' being a subset of U' means we can define $V_i'' = (f_i')^{-1}(U'')$. This definition of V_i'' will be the same, because $f_i^{-1}, (f_i')^{-1}$ are both restrictions of the same continuous homeomorphic map to open sets U and U', so they must agree on their intersection, U''. As a result, we may likewise conclude that

$$|\{V_i''\}| = |\{V_i'\}| = |p^{-1}(x')|.$$

From here, we can define S_c to be the set of points $x \in X$ such that $|p^{-1}(x)| = c$. We have just shown that S_c is open. In particular, we can see that $X = \bigsqcup_{i=1}^{\infty} S_i$. Since X is connected, this implies that all but one of the sets S_i must be empty, as X cannot be expressed as the disjoint union of open sets. Thus, $X = S_c$ for some c, which proves the proposition. $\square$

If X is connected, and $p : Y \to X$ is a covering with $|p^{-1}(x)| = n < \infty$, we say that p is an n-sheeted covering of X.

Definition 2.5. A map $p : Y \to X$ is said to be a branched covering of X if there exist finite subsets $S \subset Y, T \subset X$ such that $p|_{Y \setminus S} : Y \setminus S \to X \setminus T$ is a covering of $X \setminus T$.

The finite points $T \subset X$ over which p is not a covering of the same sheet as it is over $X \setminus T$ are called the branch points of p.

Remark 2.6. Branched coverings from one Riemann surface to another take the form of holomorphic maps, maps which locally look like x^k.

In essence, a branched covering is a covering outside of a finite set of points. In fact, over compact Riemann surfaces, this comparison becomes more than just analogy, as we see in the following proposition.

Proposition 2.7. *Let Y be a compact Riemann surface, $S \subset Y$ a finite subset, and X° a connected space such that $p' : X^\circ \to Y \setminus S$ is a finite-sheeted covering of $Y \setminus S$. Then X° can be embedded into a Riemann surface X such that p' extends into a branched covering $p : X \to Y$.*

The proof of this proposition can be found in [1], from pages 268 to 271.
Next, we introduce the concept of a sheaf.

Definition 2.8. A sheaf of sets over a topological space X is a rule $\mathcal{F}$ such that for each open set $U \subset X$, $\mathcal{F}(U)$ is a set, and for $V \subset U$ open sets, there exists a map $f_{UV} : \mathcal{F}(U) \to \mathcal{F}(V)$ such that f_{UU} is the identity on U and $f_{UW} = f_{VW} \circ f_{UV}$ for all $W \subset V \subset U$ open sets in X. Furthermore, if U is a nonempty open set in X with a (potentially infinite) open cover $\{U_i\}$, then the following hold:
i) If $s, t \in \mathcal{F}(U)$ satisfy $s|_{U_i} = t|_{U_i}$ for all i, then $s = t$.
ii) If we have a set of representatives $\{s_i : s_i \in \mathcal{F}(U_i)\}$ satisfying $s_i|_{U_i \cap U_j} = s_j|_{U_i \cap U_j}$ for all i, j, then there exists a unique $s \in \mathcal{F}(U)$ satisfying $s|_{U_i} = s_i$.

We can similarly define a sheaf of groups, abelian groups, rings, etc. The two conditions defined above are commonly known as the axioms of 'locality' and 'gluing', respectively. Intuitively, locality tells us that understanding an element $s \in \mathcal{F}(U)$ on a cover of U is enough to understand s completely, and gluing tells us that elements which agree on intersections of covers can be 'glued' together naturally into an element $s \in \mathcal{F}(U)$.

There exists a correspondence relating certain sheaves to coverings, as we see in these next few results, that makes sheaves a very useful tool for the discussion of coverings.

Definition 2.9. A sheaf over a topological space X is called locally constant if for all $x \in X$, there exists an open set $U \ni x$ such that for V an open set contained in U, $\mathcal{F}(U) = \mathcal{F}(V)$ and f_{UV} is the identity map.

Proposition 2.10. *There is a correspondence between coverings of X and locally constant sheaves of X.*

This result is proven by introducing the concept of a stalk, what is in essence the direct limit of a sheaf at a point. Formally, the stalk of a sheaf $\mathcal{F}$ at a point x is defined by a set of equivalence classes $\mathcal{F}_x = \{[s, U] : x \in U, s \in \mathcal{F}(U)\}$ under the equivalence relation $[s, U] \sim [t, V]$ if and only if there is a $W \subset U \cap V$ such that $s|_W = t|_w$. It can then be shown that if $\mathcal{F}$ is a locally constant sheaf, the map $p : \sqcup_{x \in X} \mathcal{F}_x \to X$ defined by $p(\mathcal{F}_x) = x$ is a well-defined branch cover, and further that this construction can be inverted for an arbitrary branched covering to generate a locally constant sheaf.

We end this section with a brief discussion of meromorphic functions, in particular over Riemann surfaces.

Definition 2.11. A meromorphic function over a Riemann surface is a function $f : X \to \mathbb{C} \cup \{\infty\}$ which is complex differentiable and is not identically ∞.

Proposition 2.12. *The set of meromorphic functions over a Riemann Surface X forms a field.*

This proposition is a known result; proof can be found at page 72 in [3]. It suffices to show the field axioms hold for this set; to this end, it is critical both that the surface is connected, and (as we shall soon see), poles and zeroes of each function are relatively well behaved.

We denote this field as $\mathcal{M}(X)$.

Remark 2.13. $\mathcal{M}(\hat{\mathbb{C}}), \cong \mathbb{C}(x)$
$\mathcal{M}(\mathcal{C}) \cong \mathbb{C}(x, \sqrt{x^5 - x})$

The former portion of the remark is derived by showing the meromorphic functions over $\hat{\mathbb{C}}$ are exactly the rational functions over $\mathbb{C}$. A complete and general proof can be found in [2] at pages 13 and 14. The latter result is derived via some calculations using later results in this paper. We elaborate later on this derivation in section 4.

Proposition 2.14. *If X is a compact Riemann Surface, and f is a meromorphic function over X that is not identically zero, the poles and zeroes of f each form a discrete set.*

Proof. Let f be an element of $\mathcal{M}(X)$ with non-discrete zeroes, and let $v(f)$ be the set of zeroes of f. As $v(f)$ is not discrete, it has a limit point. Then we note that

$$f(x) = 0(x), \ \forall x \in v(f),$$

where $0(x)$ is the meromorphic function on X which is zero everywhere. Since $v(f)$ is a subset of X with a limit point, the Identity Principle says that since $f = 0$ over $v(f)$, $f = 0$ over some open neighborhood of X.

Let S be the set of all such x such that f is zero on a neighborhood of x. Then S is open. But as every limit point of $v(f)$ is contained in S, and S is a subset of $v(f)$, S is closed. Thus by the connectivity of Riemann surfaces, $S = X$, which means $f(x) = 0$ for all $x \in X$. Hence, any non-zero f in $\mathcal{M}(X)$ has discrete zeroes.

We then observe that $\frac{1}{f} \in \mathcal{M}(X)$ for f a non-zero meromorphic function on X. We just showed that $\frac{1}{f}$ has non-discrete zeroes, and the zeroes of $\frac{1}{f}$ are exactly the poles of f. Therefore we conclude that for every nonzero meromorphic function f on X, f has discrete zeroes and poles. ◻

3. Branched Coverings and Field Extensions

This section shall largely follow the proof from section 3.3 of [3].

Lemma 3.1. *Let X, Y be Riemann surfaces, $p : Y \to X$ an n-sheeted branched covering, and $f \in \mathcal{M}(Y)$, then f satisfies a polynomial in $\mathcal{M}(X)[t]$ of degree n.*

Proof. Let S be the set of branch points of p. Pick $x \in X \setminus S$. Then there exists an open set $U \ni x$ such that $p^{-1}(U) = \sqcup_{i=1}^{n} V_i$ for V_i an open subset of Y which is homeomorphic to U for $i = 1, ..., n$. We can then define n homeomorphisms $p_i : U \to V_i$. We now consider the polynomial

$$q(x,t) = \prod_{i=1}^{n} (t - f \circ p_i(x)).$$

We can see that f satisfies this polynomial over each of the sets V_i by construction. As the functions p_i are homeomorphisms, $f \circ p_i$ will necessarily be meromorphic on U. As the coefficients of $q(x,t)$ will be symmetric polynomials in $f \circ p_i$, we can conclude that the coefficients of $q(x,t)$ will also be meromorphic on U.

We can then choose $x' \in X \setminus S$, with associated open set U' such that $p^{-1}(U') = \sqcup_{i=1}^{n} V_i'$. We note that the construction of $q(x,t)$ over U' must agree with the construction of $q(x,t)$ over U on their intersection $U \cap U'$. Since the two agree on the open set, they agree everywhere. Thus following from the connectivity of X, we conclude that q is well-defined over $Y \setminus p^{-1}(S)$ with coefficients in $\mathcal{M}(X \setminus S)$.

We then only need to show that the coefficients remain meromorphic on S. Let $q(x,t) = a_n(x)t^n + ... + a_1(x)t + a_0(x)$. For any point x in S, since X is a Riemann surface, there is an open neighborhood U of x such that $\phi : U \to \mathbb{C}$ is a coordinate chart to $\mathbb{C}$ and $\phi(x) = 0$. Then $\phi \circ p$ defines a holomorphic function over some neighborhood of each $y \in p^{-1}(x)$ such that $(\phi \circ p)(y) = 0$. We thus can find $k > 0$ such that $(\phi \circ p)^k f$ is bounded in a punctured neighborhood of each $y \in p^{-1}(x)$. Then $\phi^k(f \circ p_i)$ is bounded on $U \setminus \{x\}$. In essence, ϕ^k is serving here as a coordinate chart. As a_i are symmetrical polynomials in $f \circ p_i$, this implies the existence of $m_i > 0$ for $i = 1, ..., n$ such that $\phi^{m_i} a_i$ is bounded on $U \setminus \{x\}$. By Riemann's Removable Singularity Theorem, $\phi^{m_i} a_i$ thus extends to a holomorphic function over U, which means $\{a_i\}$ are meromorphic on U. We conclude that $\{a_i\}$ are in $\mathcal{M}(X)$, and thus f satisfies $q(x,t)$, a degree n polynomial with coefficients in $\mathcal{M}(X)$. □

Theorem 3.2. *(Riemann Existence Theorem) If X is a compact Riemann surface, and $x, y \in X$, then there exists a meromorphic function $f \in \mathcal{M}(X)$ such that $f(x) \neq f(y)$.*

The proof of this theorem is long, difficult, and outside the scope of this paper. For further reading, curious readers are encouraged to read pages 275-286 of [4]. However, as an illustrative example, we consider our earlier example $\mathcal{C}$ of a compact Riemann surface. Over $\mathcal{C}$, the function $f(x,y) = x$ is such a non-constant meromorphic function. In fact, for any hyperelliptic compact Riemann surface (surface of form $y^2 = h(x)$ where h is an irreducible polynomial in x), the function $f(x,y) = x$ likewise suffices as such a non-constant meromorphic function. As we shall see moving forward, finding a non-constant meromorphic function for a general compact Riemann surface provides characterization for that surface.

Lemma 3.3. *If X is a compact Riemann surface, $x_1, ..., x_n \in X$, and $c_1, ..., c_n \in \mathbb{C}$, then there exists $f \in \mathcal{M}(X)$ such that $f(x_i) = c_i$ for $i = 1, ..., n$.*

Proof. This proof proceeds by induction on n. The base case of $n = 1$ is guaranteed trivially, as the meromorphic function $f(x) = c_1 \in \mathcal{M}(X)$ satisfies it.

For the further inductive step, we are guaranteed a function f such that $f(x_i) = c_i$ for $i = 1, ..., n-1$. We also know by Lemma 3.2 that for each pair (i, n), there exists a function $f_{in} \in \mathcal{M}(X)$, such that $f_{in}(x_i) \neq f_{in}(x_n)$. From this, we can generate a function $g_{in} \in \mathcal{M}(X)$ such that $g_{in}(x_i) = 0, g_{in}(x_n) \neq 0$. If x_i is a pole of f_{in}, we can shift f_{in} such that $f_{in}(x_n) \neq 0$ and then take the reciprocal; if not, we can just take $f_{in}(x) - f_{in}(x_i)$. We can then let $g = \prod_{i=1}^{n} g_{in}$, and g is a function such that $g(x_i) = 0$ for $i = 1, ..., n-1$, and $g(x_n) \neq 0$. We can then take linear combinations of f, g to get the desired function. □

We are now ready to prove the main theorem of this section.

Theorem 3.4. *If X, Y are compact Riemann surfaces, and $p : Y \to X$ is an n-sheeted branched covering, then $\mathcal{M}(Y)/\mathcal{M}(X)$ is a finite field extenstion of degree n.*

Proof. Let $x' \in X$ be a non-branch point of p - in particular, $p^{-1}(x')$ is a set of n distinct points $x_1, ..., x_n \in Y$. Then Lemma 3.3 tells us that there exists a function $f \in \mathcal{M}(Y)$ satisfying $f(x_i) = i$ for each $i = 1, ..., n$.

We then consider the polynomial $q(x,t) = \prod_{i=1}^{n}(t - f \circ p_i)$, as constructed in the proof of Lemma 3.1. As Lemma 3.1 tells us, f satisfies q as a polynomial in $\mathcal{M}(X)[t]$. We then note that

$$q(x', t) = \prod_{i=1}^{n}(t - f \circ p_i(x')) = \prod_{i=1}^{n}(t - f(x_i)) = \prod_{i=1}^{n}(t - i).$$

This means (x', i) is a root of q for $i = 1, ..., n$. We can use this to prove q is the minimal polynomial of f over $\mathcal{M}(X)$. After all, we know by the construction of f, there must necessarily be n distinct roots of the minimal polynomial when the coefficients are evaluated at x', as we have chosen f to have n distinct values on the n distinct values of $p^{-1}(x')$. However, if q were not the minimal polynomial of f, the true minimal polynomial r would have to be of degree strictly smaller than n in $\mathcal{M}(X)[t]$, which means by the fundamental theorem of algebra it *cannot* take on n distinct roots. This means that q itself is the minimal polynomial of f over $\mathcal{M}(X)$.

We then consider $\mathcal{M}(X)(f)$. This is, by construction, a finite field extension of $\mathcal{M}(X)$ of degree n. We also note that since $f \in \mathcal{M}(Y)$, $\mathcal{M}(X)(f) \subset \mathcal{M}(Y)$. This means that $[\mathcal{M}(Y) : \mathcal{M}(X)] \geq n$. However, Lemma 3.1 tells us that every element of $\mathcal{M}(Y)$ satisfies a polynomial of degree at most n, which means that $[\mathcal{M}(Y) : \mathcal{M}(X)] \leq n$. From this, we can conclude that $[\mathcal{M}(Y) : \mathcal{M}(X)] = n$. $\square$

With this, we have proven one direction of what we shall soon see is a two-way correspondence between field extensions and branched coverings.

4. Finite Field Extensions and Branched Coverings

We begin this section by proving some lemmas.

Lemma 4.1. *Let X be a a compact Riemann surface. If $q(x,t) \in \mathcal{M}(X)[t]$ is an irreducible polynomial, then the set of points $x \in X$ such that there exists $t \in \mathbb{C}$ satisfying $q(x,t) = 0$ and $\frac{\partial q}{\partial t}(x,t) = 0$ is finite.*

Proof. Let I be the ideal generated by q and $\frac{\partial q}{\partial t}$. As $\mathcal{M}(X)$ is a field, $\mathcal{M}(X)[t]$ is a Principal Ideal Domain. Consequently, I is generated by some single element $d(x,t)$. Since $q \in I$, this implies the existence of some $c(x,t) \in \mathcal{M}(X)[t]$ such that $q(x,t) = c(x,t)d(x,t)$. However, as q is irreducible, this implies that d is either associate to q, or is a unit.

However, $\frac{\partial q}{\partial t} \in I = (d)$ implies that the degree of d is at most the degree of $\frac{\partial q}{\partial t}$. As all associates to q must be of the same degree as q, and the degree of $\frac{\partial q}{\partial t}$ is strictly less than that of q, we conclude that d cannot be associate to q. Thus, d is a unit of the ring $\mathcal{M}(X)[t]$, which means the ideal generated by d is exactly the ideal generated by $1(x,t)$ (the function which is identically 1 everywhere).

From here, we may apply Bézout's Lemma for Principal Ideal Domains, which says that since $(q, \frac{\partial q}{\partial t}) = (1)$, there exist $a, b \in \mathcal{M}(X)[t]$ such that

$$a(x,t)q(x,t) + b(x,t)\frac{\partial q}{\partial t}(x,t) = 1(x,t) = 1.$$

We then note that if (x,t) is a point such that $q(x,t)=0$ and $\frac{\partial q}{\partial t}=0$, then

$$a(x,t)*0+b(x,t)*0=1.$$

Thus, the only way for this to be maintained is for x to be a pole of at least one of the coefficients of a,b (recall that a and b are polynomials in t with coefficients in $\mathcal{M}(X)$), since otherwise the left hand side of the equation would have to be 0. However, we know by Proposition 2.14 that there are necessarily only finitely many poles. We thus conclude that the number of points x such that there exists $t \in \mathbb{C}$ satisfying $q(x,t)=0$ and $\frac{\partial q}{\partial t}=0$ is finite. □

Lemma 4.2. *Let X be a a compact Riemann surface, and let $q(x,t)=\sum_{i=0}^{n} a_i(x)t^i \in \mathcal{M}(X)[t]$ be irreducible. Then for all but finitely many values of $y \in X$, $\sum_{i=0}^{n} a_i(y)t^i$ is separable.*

Proof. First, by Proposition 2.14 there are only finitely many poles of each of the $n+1$ coefficients of the polynomial $q(x,t)$, as they are meromorphic functions in a compact Riemann Surface. Call the set of such points P. Since P is finite, it is sufficient to prove the claim on $X \setminus P$. Over $X \setminus P$, $\sum_{i=0}^{n} a_i(y)t^i$ is a polynomial in $\mathbb{C}[t]$. Such a polynomial is inseparable if and only if it shares a root with its derivative. But

$$\frac{d}{dt}\sum_{i=0}^{n} a_i(y)t^i = \sum_{i=1}^{n} i a_i(y)t^{i-1} = \frac{\partial q}{\partial t}(y,t),$$

meaning $\sum_{i=0}^{n} a_i(y)t^i$ is inseparable if and only if $q(y,t)=0$ and $\frac{\partial q}{\partial t}(y,t)=0$ for some value of t in $\mathbb{C}$. By Lemma 4.1, there are only finitely many such values of y.

Consequently, we conclude that there are only finitely many values of y in X such that $\sum_{i=0}^{n} a_i(y)t^i$ is inseparable. □

Lemma 4.3. *Let X be a compact Riemann surface. If $q(x,t) \in \mathcal{M}(X)[t]$ is a polynomial of degree n, then for all but finitely many points $x \in X$, there is an open set $U \ni x$ such that there are exactly n distinct holomorphic functions $l_1,...,l_n$ on U satisfying $q(y,l_i(y))=0$ for $i=1,...,n, y \in U$.*

Proof. Let $q(x,t)=\sum_{i=0}^{n} a_i(x)t^i$. By Lemma 4.1 we know that there are only finitely many values of $y \in X$ such that there exists an s in $\mathbb{C}$ simultaneously satisfying $q(y,s)=0$ and $\frac{\partial q}{\partial t}(y,s)=0$. We call the set of such values of y $P_1 \subset X$. Similarly, Lemma 4.2 tells us that there are only finitely many values of $y \in X$ such that $\sum_{i=0}^{n} a_i(y)t^i$ is inseparable. We call the set of such values $P_2 \subset X$. Lastly, we denote the set of zeroes of a_n and poles of $\{a_0,...,a_n\}$ as P_3. As these are all meromorphic functions over the compact Riemann surface X, and a_n is not identically zero, Proposition 2.14 tells us that P_3 is finite.

Let x be an element of $X \setminus (P_1 \cup P_2 \cup P_3)$. As x is not contained in P_2, $\sum_{i=0}^{n} a_i(x)t^i$ is a separable polynomial. In particular, since x is not contained in P_3, this is a separable polynomial of degree exactly n in $\mathbb{C}[t]$. By the Fundamental Theorem of Algebra, such a polynomial has n distinct complex roots, which we denote by $x_1,...,x_n$.

We note first that by definition this means $q(x,x_i)=0$ for $i=1,...,n$. Furthermore, we note that as x is not contained within P_1, $q(x,x_i)=0$ means that $\frac{\partial q}{\partial t}(x,x_i) \neq 0$ for $i=1,...,n$. Consequently, for each such value of x_i, the holomorphic implicit function theorem says that there is a holomorphic function l_i on some open subset $U_i \ni x$ of X such that $l_i(x)=x_i$ and $q(y,l_i(y))=0$ for all $y \in U_i$. We then let $U = \cap_{i=1}^{n} U_i$, and we can see that all of the l_i functions are holomorphic over $U \subset X$, and that U is indeed an open subset of X containing x.

As for the condition of there being 'exactly' n such functions, this follows from the uniqueness criterion of the Holomorphic Implicit Function Theorem, i.e. that for each value x_i, l_i is the *unique* holomorphic function satisfying $l_i(x) = x_i$ and $q(x, l_i(x)) = 0$ over U.

Consequently, we conclude that for all $x \in X$ outside the finite set $P_1 \cup P_2 \cup P_3$, there is an open subset U in X containing x such that there are exactly n distinct holomorphic functions $l_1, ..., l_n$ satisfying $q(y, l_i(y)) = 0$ for all $y \in U$, for $i = 1, ..., n$, as we desired to show. $\square$

We are now prepared to prove the second half of our correspondence.

Theorem 4.4. *Let X be a compact Riemann surface. If L is a field extension of $\mathcal{M}(X)$ of degree n, then*
(i) $L \cong \mathcal{M}(Y)$ for some compact Riemann surface Y, and
(ii) There exists an n-sheeted branched cover $p : Y \to X$.

Proof. As L is a finite degree field extension of $\mathcal{M}(X)$, a field of characteristic 0, L is a separable field extension. By the Primitive Element Theorem, $L \cong \mathcal{M}(X)(l)$ for some $l \in L$. We denote the minimal polynomial of l over $\mathcal{M}(X)$ as $q(x, t)$. The polynomial $q(x, t)$ is expressible as $t^n + a_{n-1}(x)t^{n-1} + ... + a_1(x)t + a_0(x)$, where $a_i \in \mathcal{M}(X), t \in \mathbb{C}$.

We then present two different manners of constructing the branched cover.

For the first method, let S be the finite subset of X under which Lemma 4.3 does not hold. Fix $x \in X \setminus S$. This means there is an open set U containing x such that there are n distinct holomorphic functions $l_1, ..., l_n$ satisfying $q(x, l_i(x)) = 0$ over U, for $i = 1, ..., n$. We then let $Y' = (x, l_i(x))$, for each $i = 1, ..., n$. As this is a surface created by adding the solutions to n holomorphic functions, Y' can be understood to be a one-dimensional complex manifold. But we note the map $p' : Y' \to X \setminus S$ defined by

$$p'(x, y) = x, \forall (x, y) \in Y'$$

is an n-sheeted covering map, as each point in $X \setminus S$ will have precisely n preimages, by Lemma 4.3. Further, the preimage of U as defined above will be n open sets containing the n preimages of x, as over U the values y in (x, y) vary holomorphically.

We then show that Y' must be connected. If C' is a connected component of Y', then $p'|_{C'}$ must be a covering of $X \setminus S$, and thus it must be a covering of some finite sheet m; in particular, we can see that $m \leq n$, as $p'|_{C'}^{-1}(x) \subset p'^{-1}(x)$ for all points $x \in X \setminus S$. By Proposition 2.7, $p'|_{C'} : C' \to X \setminus S$ can be extended to an m-sheeted branched covering $p|_C : C \to X$.

We now define the function $f(x, y) = y$ over C'. f is holomorphic over C', as we can see by how Y' is constructed from adding holomorphic functions. In fact, we can extend f meromorphically over C adding points at infinity for values $(x, y) \in C$ where $a_i(x)$ is a pole for some $i = 1, ..., n$. This means that $f \in \mathcal{M}(C)$, which by Lemma 3.1 means it satisfies a polynomial of degree m in $\mathcal{M}(X)[t]$. However, we note that f also satisfies q, since every element of C is (x, y) such that $q(x, y) = 0$. Since q is irreducible, we are forced to conclude that $m = n$, and consequently that $C = Y$, meaning that Y is connected, and thus a Riemann surface. Now, by

Proposition 2.7, since this is a covering map of X minus some finite set, we can thus extend this into an n-sheeted branched cover $p : Y \to X$.

For the second method, we define the sheaf $\mathcal{F}(U)$ thusly:

$$\mathcal{F}(U) = \{f : U \to \mathbb{C} : f \text{ holomorphic, } q(x, f(x)) = 0 \ \forall x \in U.\}$$

By examination, this can be seen to meet the criteria of being a sheaf. In particular, the inclusion map f_{UV} is the restriction map, so that for $V \subset U$, the map $f_{UV} : \mathcal{F}(U) \to \mathcal{F}(V)$ is defined as

$$f_{UV}(g) = g|_V \ \forall g \in \mathcal{F}(U),$$

which can then be confirmed to satisfy both f_{UU} being the identity morphism and $f_{UW} = f_{UV} \circ f_{VW}$. Thus, $\mathcal{F}$ is a sheaf on X. But in particular, Lemma 4.3 says that outside of a finite set of points, it is a locally constant sheaf. Formally, if P is the finite set of points under which the conclusion of Lemma 4.3 does not apply, then for $x \in X \setminus P$, there is an open set $U \ni x$ such that there exist exactly n distinct holomorphic functions $l_1, ..., l_n$ such that $q(x, l_i(x)) = 0$ for all $x \in U$, for $i = 1, ..., n$. Then $\mathcal{F}(U) = \{l_1, ..., l_n\}$. Furthermore, for an open set $V \subset U$, $\mathcal{F}(V) = \{l_1|_V, ..., l_n|_V\}$ - as a consequence of Lemma 4.3, these are the only possible elements of $\mathcal{F}(V)$. In particular, for all $y \in U$, there is some open set V satisfying $y \in V \subset U$, where $\mathcal{F}(V)$ must be $\{l_1|_V, ..., l_n|_V\}$. We conclude that

$$\mathcal{F}(y) = \{l_1|_y, ..., l_n|_y\}, \ \forall y \in U,$$

and thus that $\mathcal{F}$ is a locally constant sheaf on $X \setminus P$.

Thus by Proposition 2.10, we can construct a cover $p' : Y' \to X \setminus P$, in particular an n-sheeted cover. From there, we can extend this into an n-sheeted branched cover $p : Y \to X$ for some Riemann surface Y.

Now all that remains is to show $\mathcal{M}(Y) \cong L$. However, it is first clear that L is isomorphic to some subset of $\mathcal{M}(Y)$. This is because $L = \mathcal{M}(X)(l)$, where l is a term satisfying $q(x, l(p^{-1}(x))) = 0$. In particular, $l(x, y) = y$ for all $(x, y) \in Y$ is such an element satisfying q and thus will generate L. Since we constructed Y by adjoining the n holomorphic root functions of $q(x, t)$ at each non-branched point of X, this function l is holomorphic outside a finite set of points on Y, and can thus be seen to be meromorphic on Y. As l is a primitive element of L over $\mathcal{M}(X)$, $l \in \mathcal{M}(Y)$ means $L \subseteq \mathcal{M}(Y)$. Formally, this implies

$$\begin{array}{c} \mathcal{M}(Y) \\ \Big| \\ L \\ \Big|{\scriptstyle n} \\ \mathcal{M}(X). \end{array}$$

However, Theorem 3.4 says that the existence of the branched cover $p : Y \to X$ implies $[\mathcal{M}(Y) : \mathcal{M}(X)] = n$. Since $[L : \mathcal{M}(X)] = n$, and $\mathcal{M}(Y)$ is a field extension of L, we conclude that

$$[\mathcal{M}(Y) : L] = \frac{[\mathcal{M}(Y) : \mathcal{M}(X)]}{[L : \mathcal{M}(X)]} = \frac{n}{n} = 1,$$

and consequently that $L \cong \mathcal{M}(Y)$. $\square$

From this Theorem and Theorem 3.4 we can conclude the following:

Theorem 4.5. *If X is a compact Riemann surface, then there is an equivalence of categories between the categories of field extensions of $\mathcal{M}(X)$ and branched coverings $p : Y \to X$ for Y a compact Riemann surface.*

This is the correspondence we sought to prove. We end with some immediate consequences of this theorem.

Corollary 4.6. *If X is a compact Riemann surface, then $\mathcal{M}(X)$ is a field extension of $\mathbb{C}$ of transcendence degree one.*

Proof. We recall that by Lemma 3.2 there exists a non-constant meromorphic function $f \in \mathcal{M}(X)$. We claim that $f : X \to \hat{\mathbb{C}}$ is a branched covering. This is because f being meromorphic over X means it is holomorphic over X on all but finitely many points. Then we can see by Proposition 2.7 that it can be extended to a branched covering $X \to \mathbb{C} \cup \infty = \hat{\mathbb{C}}$ by considering the poles to be points at infinity. We can then see that X has a branched covering of $\hat{\mathbb{C}}$, meaning that by Theorem 4.5 $\mathcal{M}(X)$ is an algebraic field extension of $\mathcal{M}(\hat{\mathbb{C}})$. As we know that $\hat{\mathbb{C}}$ is isomorphic to $\mathbb{C}(x)$, a field extension of $\mathbb{C}$ of transcendence degree one, we conclude likewise that $\mathcal{M}(X)$ must be an algebraic field extension of $\mathbb{C}(x)$, and thus also be a field extension of $\mathbb{C}$ of transcendence degree one. $\square$

As $\mathcal{M}(X)$ is a field of transcendence degree one over $\mathbb{C}$, we can see that finding a single non-constant element is sufficient to allow for the computation of $\mathcal{M}(X)$. This aligns with our proof of Theorem 4.4, as we saw constructing the primitive element l of $\mathcal{M}(Y)$ was enough to completely reverse engineer the structure of the compact Riemann Surface Y.

Indeed, this allows us to justify our calculation of $\mathcal{M}(\mathcal{C}) \cong \mathbb{C}(x, \sqrt{x^5 - x}) \cong \mathbb{C}(x, y)/(y^2 = x^5 - x)$ from Remark 2.13. We recall that $f(x, y) = x$ is a non-constant meromorphic function on $\mathcal{C}$, and in particular this is a 2-sheeted branched covering of $\hat{\mathbb{C}}$, meaning $\mathcal{M}(\mathcal{C})$ is a quadratic extension of $\mathbb{C}(x)$. From there, we can see that f will be contained in $\mathbb{C}(x, \sqrt{x^5 - x})$, which is a quadratic extension of $\mathbb{C}(x)$, allowing us to see this equivalence. In fact, this leads us to the following statement.

Corollary 4.7. *There is a correspondence between hyperelliptic compact Riemann surfaces and quadratic field extensions of $\mathbb{C}(x)$.*

Acknowledgments

I would like to thank my mentor, Carlos Serván, for suggesting this topic and for constant guidance throughout the REU. I would like to thank Professor May for the organization of the REU, and all of the lecturers for donating their time to the many topics and discussions therein.

5. bibliography

References

[1] William Fulton, *Algebraic Topology: A First Course*, Graduate Texts in Mathematics, vol. 153, Springer-Verlag, New York, 1995.

[2] Ernesto Girondo, Gabino González-Diez, *Introduction to Compact Riemann Surfaces and Dessins d'Enfants*, London Mathematical Society Student Texts, vol. 79, Cambridge University Press, New York, 2012.

[3] Tamás Szamuely, *Galois Groups and Fundamental Groups*, Cambridge studies in advanced mathematics, vol. 117, Cambridge University Press, Cambridge, 2009.
[4] David Harbater, *The Legacy of Bernhard Riemann After 150 Years Volume II*, Advanced Lectures in Mathematics, vol. 35.2, Higher Education Press and International Press, Beijing-Boston, 2015, pp. 275-286.
[5] Rick Miranda, *Algebraic Curves and Riemann Surfaces*, Graduate Studies in Mathematics, vol. 5, American Mathematical Society, Rhode Island, 1995.