Riemann Existence Theorem

Ryan Patrick Catullo

July 2024

Abstract

In this paper, we study a specific version of Riemann's Existence Theorem regarding branched covers of $\mathbb{P}^1_{\mathbb{C}}$ realizing a given monodromy representation. We use this to expand on a known proof of Belyi's theorem, which describes covers over $\mathbb{P}^1_{\mathbb{C}}$ branched over three points as precisely those admitting a model over $\overline{\mathbb{Q}}$. In the latter sections, we describe a possible inductive strategy to algebraically prove RET by gluing together simple branched covers of the sphere at a node, and then resolving the singularity.

Contents

1 Introduction

The ***Riemann Existence Theorem*** (RET) is a statement or collection of related statements that lies between analysis, topology, and the geometry of algebraic curves. One statement of the theorem is that on any compact Riemann surface X, there exists a nonconstant meromorphic function $f \in \mathcal{M}(X)$ separating any two points. By definition, we have a branched cover $f\colon X \to \mathbb{P}^1_{\mathbb{C}}$ whose branch locus $\Delta \subset \mathbb{P}^1_{\mathbb{C}}$ is finite. Recall that around any $P \in X$, there are holomorphic local charts under which f is given by $f(z) = z^{e_P}$ near P for some integer e_P. This integer is called the ramification index of P, and we say P ramifies if $e_P > 1$, in which case we also say $f(P) = Q$ is a branch point of f. Equivalently $f'(P)$ vanishes since $f'(z) = e_P z^{e_P - 1}$ near P, which vanishes at $z = 0$ if and only if $e_P > 1$. Note that, in this chart, no other z has $f'(z) = 0$ and so there is an open neighborhood around P with no other ramification points. Therefore, the set of ramification points is discrete, and since X is compact the set of ramification points is finite. We can show that this implies $\mathcal{M}(X)$ is a finite extension of $\mathbb{C}(t)$, and therefore X can be recognized as a complex algebraic curve. In essence, we have an equivalence of categories

$$\left\{\begin{array}{c}\text{Compact Riemann Surfaces } X\\ \text{holomorphic maps } f: X \to Y\end{array}\right\} \leftrightarrow \left\{\begin{array}{c}\text{Complex Algebraic Curves } C\\ \mathbb{C}\text{-morphisms } f: C \to C'\end{array}\right\}$$

Therefore, showing this equivalence is reduced to demonstrating the existence of the nonconstant meromorphic function on X. There are several proofs, notably using local-to-global methods in the complex metric topology or Serre's GAGA to pass from the algebraic to analytic setting.

There is another similar statement, which says given a finite subset Δ of $\mathbb{P}^1_{\mathbb{C}}$ (the branch points) and a permutation representation of $\pi_1(\mathbb{P}^1_{\mathbb{C}} \setminus \Delta) = \langle \gamma_1, \ldots, \gamma_n \mid \gamma_1 \cdots \gamma_n = 1 \rangle$, there is a compact Riemann surface X and a morphism $X \to \mathbb{P}^1_{\mathbb{C}}$ which is a covering space away from Δ having this representation as its monodromy. This is the version we will discuss in this paper, and monodromy representations and the proof of the theorem are expanded in more detail in Section 2.1.

2 Riemann's Existence Theorem

Recall some basic properties of $\mathbb{P}^1_{\mathbb{C}}$. Namely, let $\Delta = \{x_1, \ldots, x_n\}$ be a finite subset of $\mathbb{P}^1_{\mathbb{C}}$, and consider the topological space given by removing Δ from $\mathbb{P}^1_{\mathbb{C}}$. One can show that the fundamental group is generated by loops γ_i around each x_i subject to the relation $\gamma_1 \cdots \gamma_n = 1$, as the image below demonstrates.

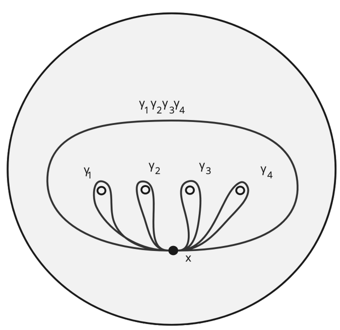

Figure 1: Visual demonstration that $\pi_1(\mathbb{P}^1_{\mathbb{C}} \setminus \Delta, x) = \langle \gamma_1, \ldots, \gamma_n \mid \gamma_1 \cdots \gamma_n = 1 \rangle$

Bringing the product loop around the top of the sphere, down the back, and back up to the front from the bottom allows us to contract it to the basepoint, hence it is the trivial loop. The idea of a monodromy representation is that knowing how these loops act on the covering space $X \setminus \varphi^{-1}(\Delta) \to \mathbb{P}^1_{\mathbb{C}} \setminus \Delta$ by permuting the fiber over the basepoint (recall the action of the fundamental group on the covering space by path lifting) should tell us the ramification data needed to "fill in" the points of X that we removed to get an unbranched covering.

2.1 Monodromy Representations

Let $f\colon Y \to X$ be a morphism of degree d ramified over some finite branch locus $\Delta = \{x_1, \dots, x_n\} \subset X$, and let $x \in X$ be a chosen regular value. Let $\Omega = f^{-1}\Delta$ be the ramification locus. Then we can associate to f a group homomorphism

$$M_f\colon \pi_1(X \setminus \Delta, x) \to \mathrm{Bij}(f^{-1}(x))$$

by sending $\gamma \mapsto \sigma_\gamma^{-1}$, where σ_γ is defined as follows. Note that $f\colon Y \setminus \Omega \to X \setminus \Delta$ is an unbranched covering, and so $\pi_1(X \setminus \Delta, x)$ acts on $y \in f^{-1}(x)$ by lifting a loop γ at x to a path $\tilde{\gamma}$ starting at y and ending at some other $y' \in f^{-1}(x)$, and letting $y \cdot \gamma = y'$. Note that this is a right action since path concatenation is performed in opposite order of function composition. We let σ_γ just be the bijection on $f^{-1}(x)$ which sends $\sigma_\gamma(y) = y \cdot \gamma$.

The reason why we map $M_f(\gamma) = \sigma_\gamma^{-1}$ is so that this is a group homomorphism since path concatenation and function composition are done in reverse orders. Indeed, if $y \cdot \alpha = y'$ and $y' \cdot \beta = y''$, then $\sigma_{\alpha\beta}^{-1}(y'') = y$ and $\sigma_\alpha^{-1} \circ \sigma_\beta^{-1}(y'') = \sigma_\alpha^{-1}(y') = y$.

If we number the fiber over x, i.e. choose a bijection $\varphi\colon \{1, \dots, d\} \to f^{-1}(x)$, then we can construct a homomorphism to Σ_d, the symmetric group on d letters.

$$\rho\colon \pi_1(X \setminus \Delta, x) \to \Sigma_d$$

where $\gamma \mapsto \varphi^{-1} \circ M_f(\gamma) \circ \varphi$. Note that if $\psi\colon \{1, \dots, d\} \to f^{-1}(x)$ is another labelling of the fiber over x giving ρ', then $\rho' = \tau^{-1} \circ \rho \circ \tau$ where $\tau = \varphi^{-1} \circ \psi \in \Sigma_d$, and therefore they are conjugate.

Additionally, if x' is another regular value basepoint, then the homomorphism it produces ρ' will also be conjugate to ρ. To see this, let δ be a path from x to x' and note that

$$\pi_1(X \setminus \Delta, x) = \delta \cdot \pi_1(X \setminus \Delta, x') \cdot \delta^{-1}$$

where the operation is concatenation, not composition. Then for $y \in f^{-1}(x)$, we have $y \cdot \delta \in f^{-1}(x')$, and so the labeling $\varphi\colon \{1, \dots, d\} \to f^{-1}(x)$ gives $\varphi_\delta\colon \{1, \dots, d\} \to f^{-1}(x')$ by $i \mapsto \varphi(i) \cdot \delta$. Then it's a short computation to check $\rho(\gamma) = \rho'(\delta^{-1}\gamma\delta)$ with these labellings, and therefore the homomorphisms are conjugate.

Definition 2.1 Let $f\colon Y \to X$ be a degree d morphism ramified over some finite $\Delta \in X$, with $x \in X$ a fixed regular value. Fix a labelling $\varphi\colon \{1, \dots, d\} \to f^{-1}(x)$ of the fiber over x. Then the group homomorphism

$$\rho\colon \pi_1(X \setminus \Delta, x) \to \Sigma_d$$

defined above is called the *monodromy* of f, and is well-defined up to conjugation by elements of Σ_d. Let $\mathrm{Mon}(f)$ be the image of ρ in Σ_d, which we will refer to as the *monodromy group* of f. When Y is connected, the monodromy group is transitive.

We can now state the content of the main theorem.

Theorem 2.1 (Riemann's Existence Theorem) Let $\Delta \subset \mathbb{P}^1_{\mathbb{C}}$ be a finite subset of size d, and fix a point $x \in \mathbb{P}^1_{\mathbb{C}} \setminus \Delta$. Let ρ be a morphism

$$\rho\colon \pi_1(\mathbb{P}^1_{\mathbb{C}} \setminus \Delta, x) \to \Sigma_d$$

with transitive image. Then there is a compact Riemann surface X and morphism $\varphi\colon X \to \mathbb{P}^1_{\mathbb{C}}$ of degree d such that ρ is the monodromy of φ.

We will recall the Galois correspondence between connected covering spaces of a pointed space (X, x) and subgroups of the fundamental group $\pi_1(X, x)$, which will make the method of proving this theorem more apparent. Usually, for a semi-locally simply connected and path-connected space X (i.e. such that X has a universal cover), the equivalence gives actually an equivalence of categories between path-connected covering spaces $p\colon (Y, y) \to (X, x)$ and subgroups H of $\pi_1(X, x)$. Here, $p\colon (Y, y) \to (X, x)$ maps to $p_*\pi_1(Y, y)$. In the other direction, the group H acts on the universal cover $\tilde{X}$ by unique path lifting since its elements are loops in $\pi_1(X)$. Then we can take the quotient space $\tilde{X}/H =: X_H$ such that $p\colon (X_H, y) \to (X, x)$ is a covering space where y is any point in the fiber such that $p_*\pi_1(X_H, y) = H$.

We will use a different form of this equivalence that is useful for monodromy representations. Let X be a path-connected semi-locally sipmly connected space with a point $x \in X$, and let $\mathrm{Cov}(X)$ denote the category whose objects are covering spaces $p\colon Y \to X$ and whose morphisms are maps of covering spaces $\sigma\colon Y' \to Y$ that commute with p' and p. Let Set^{π_1} denote the category of sets equipped with a right $\pi_1(X,x)$-action. We will describe a functor $\mathrm{Fib}_x\colon \mathrm{Cov}(X) \to \mathrm{Set}^{\pi_1}$ as follows. Given a covering space $p\colon Y \to X$, we have a π_1-set given by the fiber $p^{-1}(x)$ and action given by path lifting loops at x to an element of the fiber.

Lemma 2.2 The functor $\mathrm{Fib}_x\colon \mathrm{Cov}(X) \to \mathrm{Set}^{\pi_1}$ is an equivalence of categories.

Proof. [Sza09, Theorem 2.3.4] □

The functor in the opposite direction is given as follows. Given a π_1-set F, we can construct Y as follows. Recall that $\tilde{X} = \{[\delta]\colon \text{paths } \delta \text{ on X with initial point } x\}$ with the correct topology is the universal cover of X. Then we want somehow F copies of $\tilde{X}$ such that our covering has degree equal to the size of F, and such that the fiber over x is F. Therefore, we let $Y = F \times \tilde{X}/\sim$ where if $\gamma \in \pi_1(X,x)$, we let $(y\gamma, \delta) \sim (y, \gamma \cdot \delta)$. The obvious projection $p\colon Y \to X$ is given by $p(\delta, y) = \delta(1)$, the other endpoint of δ, since then the fiber over x is points (γ, y) where $\gamma \in \pi_1(X,x)$, and hence $(\gamma, y) \sim (1, \gamma y) \in F$.

Proof of Theorem 2.1. Suppose we are given some representation as above

$$\rho\colon \pi_1(\mathbb{P}^1_{\mathbb{C}} \setminus \Delta, x) \to \Sigma_d$$

Let $F = \{1, \ldots, d\}$ and note that ρ induces a right action of $\pi_1(\mathbb{P}^1_{\mathbb{C}} \setminus \Delta, x)$ on F, so by Lemma 2.2 there is a covering space $p\colon X_0 \to \mathbb{P}^1_{\mathbb{C}} \setminus \Delta$ such that if F is a labeling of the fiber over x, the action of $\pi_1(\mathbb{P}^1_{\mathbb{C}} \setminus \Delta, x)$ on F by path lifting is the same as the action given by ρ. For any point $z \in \Delta$, we can find a small neighborhood near z such that in holomorphic charts it looks like $U = \mathbb{D}^*$, the punctured open disk, and the fiber over it under p are $k \leq d$ disjoint copies $\sqcup_{i=1}^k V_i = \sqcup_{i=1}^k \mathbb{D}^*$ (this is by definition of a covering space). Since changing the basepoint is just conjugation by a path, we can assume the basepoint lies in U. Let $x_i \in V_i$ be an (not necessarily unique) element of the fiber over $x \in U$. Let n_i be the size of the orbit of x_i under the action of $\pi_1(U,x)$ on the fiber. Since, locally, the map $V_i = \mathbb{D}^* \to \mathbb{D}^* = U$ is given by $z \mapsto z^k$ for some integer k, it must be that $k = n_i$ since k is the number of preimages of $x \neq 0$ in V_i, which is also equal to the size of its π_1-orbit since any two preimages of x are path-connected, hence this path descends to a loop at x whose action maps the points to each other. Therefore, we can extend this map by adding 0 to both disks and letting the map be locally given by $z \mapsto z^{n_i}$.

Therefore, we have described a topological branched cover $p\colon X \to \mathbb{P}^1_{\mathbb{C}}$ which is unramified away from Δ, and realizing ρ as its monodromy. However, we require X to be a Riemann surface. Note that p is a local homeomorphism, and we have holomorphic charts on $\mathbb{P}^1_{\mathbb{C}}$, and so by giving X these same holomorphic charts we get a unique Riemann surface structure on X making p a holomorphic branched covering of Riemann surfaces. This concludes the proof of the theorem. □

There are some interesting consequences of RET, besides its implications in the GAGA equivalence between complex algebraic curves and Riemann surfaces. One such surprising example is Belyi's theorem, a classification statement about covers of $\mathbb{P}^1_{\mathbb{C}}$ branched only over $\overline{\mathbb{Q}}$ points.

2.2 Belyi's Theorem

A smaller problem one can consider is to consider branched covers $p\colon X \to \mathbb{P}^1_{\mathbb{C}}$ such that the branch locus is just $\{0, 1, \infty\}$, i.e. just three points up to some Mobius transformation of the line. It turns out that, if X is a complex algebraic curve, there exists a function $f\colon X \to \mathbb{P}^1_{\mathbb{C}}$ branched over $\{0, 1, \infty\}$ if and only if X is defined over $\overline{\mathbb{Q}}$. When we say X is defined over $\overline{\mathbb{Q}}$, this means there is some curve C over $\overline{\mathbb{Q}}$ and morphism $f'\colon X' \to \mathbb{P}^1_{\overline{\mathbb{Q}}}$ ramified over the same points as f and such that (X, f) is the pullback of (X', f') along $\mathbb{P}^1_{\mathbb{C}} \to \mathbb{P}^1_{\overline{\mathbb{Q}}}$. Both directions are proved in separate ways, so we state them as separate theorems. The following result was known before Belyi, and so is often referred to in the literature as the "obvious" direction of the theorem, although obvious is a misnomer. The idea is to take a cover $\varphi\colon X \to \mathbb{P}^1_{\mathbb{C}}$ and widen it into a family of smooth covers branched over the same points. Then we use Riemann's existence theorem to show every cover in the family is isomorphic, and one such element of the family is a curve defined over $\overline{\mathbb{Q}}$.

Theorem 2.3 (Belyi's Theorem 1) Let X be a compact Riemann surface, i.e. a complex algebraic curve, and let $\varphi\colon X \to \mathbb{P}^1_{\mathbb{C}}$ be a branched covering ramified over some finite subset $\Delta \subset \mathbb{P}^1_{\overline{\mathbb{Q}}}$. Then there is a curve X' and morphism $\varphi'\colon X' \to \mathbb{P}^1_{\overline{\mathbb{Q}}}$ both over $\overline{\mathbb{Q}}$ as above, such that the below diagram is a fiber diagram.

$$\begin{array}{ccc} X & \longrightarrow & X' \\ \Big\downarrow{\scriptstyle \varphi} & & \Big\downarrow{\scriptstyle \varphi'} \\ \mathbb{P}^1_{\mathbb{C}} & \longrightarrow & \mathbb{P}^1_{\overline{\mathbb{Q}}} \end{array}$$

Proof. Let $X \subset \mathbb{P}^n_{\mathbb{C}}$ be cut out by $Z(F_1, \ldots, F_r)$ where the F_i are homogeneous polynomials in n variables with coefficients in $\mathbb{C}$. Let L be the minimal field over $\overline{\mathbb{Q}}$ containing the coefficients of the Z_i and of φ, and note $\mathbb{C}/L/\overline{\mathbb{Q}}$. Then it is a basic fact in field theory that there are $z_1, \ldots, z_m \in L$ such that $\overline{\mathbb{Q}}(z_1, \ldots, z_m)/\overline{\mathbb{Q}}$ is purely transcendental, and $L/\overline{\mathbb{Q}}(z_1, \ldots, z_m)$ is finite, hence simple by the primitive element theorem. That is, there is $w \in L$ such that $L = \overline{\mathbb{Q}}(z_1, \ldots, z_m, w)$. Suppose the minimal polynomial of w over $\overline{\mathbb{Q}}(z_1, \ldots, z_m)$ is given by

$$w^d + \frac{p_{d-1}(z_1, \ldots, z_m)}{q(z_1, \ldots, z_m)} w^{d-1} + \ldots + \frac{p_0(z_1, \ldots, z_m)}{q(z_1, \ldots, z_m)}$$

after multiplying by an appropriate polynomial to get a common denominator of $q(z_1, \ldots, z_m)$. Then by making the substitution $w = q(z_1, \ldots, z_m)w$ which still generates L, we can assume the minimal polynomial of w is $w^d + p_{d-1}w^{d-1} + \ldots + p_0$.

Let $A = \overline{\mathbb{Q}}[z_1, \ldots, z_m, w]/(w^d + p_{d-1}w^{d-1} + \ldots + p_0)$ and let $V = \operatorname{Spec} A$ be a variety such that its function field $K(V) = L$. Then by localizing at the product of the coefficients, we can assume the coefficients of $F_1, \ldots, F_r, \varphi$ lie in A. Thus, we can define $\mathsf{X} = Z(F_1, \ldots, F_r) \subset \mathbb{P}^n_V = \mathbb{P}^n_{\overline{\mathbb{Q}}} \times_{\operatorname{Spec}\overline{\mathbb{Q}}} V$ which is equipped with a morphism $\Phi\colon \mathsf{X} \to \mathbb{P}^1_V$ given by the equation for φ. Note that, by construction, the below diagram is a fiber product.

$$\begin{array}{ccc} X & \longrightarrow & \mathsf{X} \\ \Big\downarrow{\scriptstyle \varphi} & & \Big\downarrow \\ \mathbb{P}^1_{\mathbb{C}} & \longrightarrow & \mathbb{P}^1_V \\ \Big\downarrow & & \Big\downarrow \\ \operatorname{Spec}(\mathbb{C}) & \longrightarrow & V \end{array}$$

The bottom map $\bar{\eta}\colon \operatorname{Spec}\mathbb{C} \to V$ is the inclusion of the generic point, given by the inclusion $A \to L \to \mathbb{C}$, and the fiber over $\operatorname{Spec}\mathbb{C}$ under Π is X (note this is the inclusion of the generic point of V). By using the fact that this fiber is a smooth curve whose map φ is ramified over Δ, we will show we can find an open neighborhood of $\operatorname{Spec}\mathbb{C}$ in V such that this still holds, i.e. the fiber of Φ over a point of $\mathbb{P}^1_V$ will be a smooth curve ramified over the same points Δ as $\varphi\colon X \to \mathbb{P}^1_{\mathbb{C}}$. This statement makes sense since, by assumption, $\Delta \subset \mathbb{P}^1_{\overline{\mathbb{Q}}}$ and $\mathbb{P}^1_V = \mathbb{P}^1_{\overline{\mathbb{Q}}} \times_{\operatorname{Spec}\overline{\mathbb{Q}}} V$.

Let $\operatorname{Sing}(\mathsf{X}/V)$ be the singular locus of X, i.e. the locus $Z(F_1, \ldots, F_r, K_i) \subset \mathsf{X}$ where K_i are the $(n-1) \times (n-1)$ minors of the Jacobian $J = \left(\frac{\partial F_i}{\partial x_j}\right)$. Since $\mathbb{P}^1_V \to V$ is proper, $\Pi(\operatorname{Sing}(\mathsf{X}/V)) \subset V$ is closed. Further, $\operatorname{Spec}\mathbb{C} \in V \setminus \Pi(\operatorname{Sing}(\mathsf{X}/V))$ since the fiber under Π is the smooth curve X, hence contains no singular points. Thus, the open set $U = V \setminus \Pi(\operatorname{Sing}(\mathsf{X}/V)) \subset V$ is nonempty, and so after localizing we can assume without loss of generality that the fibers over points of V are smooth curves.

To show that Φ is ramified only over $\Delta \times_{\operatorname{Spec}\overline{\mathbb{Q}}} V$, we use a similar argument. We have a map $d\Phi\colon \Phi^*\Omega_{\mathbb{P}^1_V/V} \to \Omega_{\mathsf{X}/V}$ of cotangent bundles over X whose cokernel is $\Omega_{\mathsf{X}/\mathbb{P}^1_V}$,. Then $\operatorname{Supp}(\Omega_{\mathsf{X}/\mathbb{P}^1_V}) \subset \mathsf{X}$

is the ramification locus, i.e. the points where $d\Phi$ is not an isomorphism, which is a closed subscheme since it can be described locally as the vanishing locus of the determinant of the Jacobian of Φ. Then let $R = \mathrm{Supp}(\Omega_{X/\mathbb{P}^1_V}) \cap \Phi^{-1}(\mathbb{P}^1_V \setminus (\Delta \times_{\mathrm{Spec}\,\overline{\mathbb{Q}}} V))$ be the set of ramification points lying over $\Delta \times_{\mathrm{Spec}\,\overline{\mathbb{Q}}} V$, which is a closed subscheme of X. Since Π is proper, $\Pi(R) \subseteq V$ is closed and so $U = V \setminus \Pi(R)$ is an open subset. Further, it is nonempty since it contains $\mathrm{Spec}\,\mathbb{C}$, as its fiber under Π is X which ramifies over $\Delta \times_{\mathrm{Spec}\,\overline{\mathbb{Q}}} \mathrm{Spec}\,\mathbb{C} \subseteq \Delta \times_{\mathrm{Spec}\,\overline{\mathbb{Q}}} V$. Therefore, after localizing appropriately we can assume Φ ramifies only over $\Delta \times_{\mathrm{Spec}\,\overline{\mathbb{Q}}} V$.

Now we are done localizing and we use the main trick of this theorem, which is to invoke RET. Since this is a statement about complex curves, we have to base change to $\mathbb{C}$ first. We have maps $V \to \mathrm{Spec}\,\overline{\mathbb{Q}}$ and $\mathrm{Spec}\,\mathbb{C} \to \mathrm{Spec}\,\overline{\mathbb{Q}}$ given by inclusions, so we can base change to get $V_{\mathbb{C}} = V \times_{\mathrm{Spec}\,\overline{\mathbb{Q}}} \mathrm{Spec}\,\mathbb{C}$. Similarly, we have $\Phi_{\mathbb{C}} \colon X_{\mathbb{C}} \to \mathbb{P}^1_{V_{\mathbb{C}}}$ and since everything behaves well with respect to base change, we have that the fibers are complex smooth curves and $\Phi_{\mathbb{C}}$ ramifies only over $\Delta_{\mathbb{C}} \times_{\mathrm{Spec}\,\mathbb{C}} V_{\mathbb{C}}$ where $\Delta_{\mathbb{C}} = \Delta \times_{\mathrm{Spec}\,\overline{\mathbb{Q}}} \mathrm{Spec}\,\mathbb{C}$. Consider the family of smooth curves $\{\Phi_q \colon X_q \to \mathbb{P}^1_{\mathbb{C}}\}_{q \in V_{\mathbb{C}}}$ where $X_q \subseteq X_{\mathbb{C}}$ is the fiber under $\Pi_{\mathbb{C}}$ of $q \in V_{\mathbb{C}}$. Each is a branched covering ramified only over the same set $\Delta_{\mathbb{C}}$, and so is determined up to isomorphism by its monodromy representation $\rho_q \colon \pi_1(\mathbb{P}^1_{\mathbb{C}} \setminus \Delta_{\mathbb{C}}, x) \to \Sigma_d$ up to conjugacy. We show that for any $q, q' \in V_{\mathbb{C}}$, these monodromies are conjugate and hence the covers are isomorphic. Let $x_1, \ldots, x_d$ be a labeling of the fiber over q, and let γ be a path from q to q' in $(\mathbb{P}^1_{\mathbb{C}} \setminus \Delta_{\mathbb{C}}) \times V_{\mathbb{C}}$. Then γ acts on $x_1, \ldots, x_d$ in $X \setminus \Phi^{-1}(\Delta_{\mathbb{C}} \times V_{\mathbb{C}})$ by path lifting, giving $x_i \cdot \gamma = x'_{\sigma(i)}$ where $x'_1, \ldots, x'_d$ are the fiber over q' and $\sigma \in \Sigma_d$. Hence, $\rho_q = \sigma^{-1}\rho_{q'}\sigma$ and so every branched cover in the family $\{\Phi_q \colon X_q \to \mathbb{P}^1_{\mathbb{C}}\}_{q \in V_{\mathbb{C}}}$ is isomorphic by RET. If $q \in \overline{\mathbb{Q}}$ our cover is defined over $\overline{\mathbb{Q}}$ and isomorphic to $\varphi \colon X \to \mathbb{P}^1_{\mathbb{C}}$ by the above argument, so we are done. □

Example 2.1 Polynomials of degree d are branched coverings $\mathbb{P}^1_{\mathbb{C}} \to \mathbb{P}^1_{\mathbb{C}}$ which have a single preimage of ∞, namely ∞ which ramifies to degree d. Thus, a natural first place to find examples is polynomials whose critical values are only $0, 1$. Consider the polynomial

$$f_{a,b}(x) = 1/\beta(a,b) \int_0^x t^{a-1}(1-t)^{b-1}dt$$

where $\beta(a,b) = (a-1)!(b-1)!/(a+b-1)!$ (this is the cdf of the *beta distribution*). We can see that $f'_{a,b}(x) = 0$ if and only if $x = 0$ or 1, so these are the critical points. The critical values are also $0, 1$ since $f_{a,b}(0) = 0$, $f_{a,b}(1) = 1$.

2.3 Dessins D'Enfants

In this section, we continue the discussion of Belyi morphisms by looking at a bijective correspondence between these morphisms and bipartite graphs on X. The bijection is established by the following idea. Take a degree d Belyi morphism $f \colon X \to \mathbb{P}^1_{\mathbb{C}}$ branched over $0, 1, \infty$. Since away from the branch points, this map is a covering map, we can lift the segment $(0,1)$ to d segments in X. These are the edges of our graph. We can add back in the preimages of 0 and 1 to be the vertices, and note that if two vertices are adjacent, one is a preimage of 0 and the other is a preimage of 1. Thus, the graph admits a 2-coloring and thus is bipartite. Also implicit in this description is an orientation of the edges induced by the orientation on X. Amazingly, given a bipartite graph on X, we can recover a Belyi morphism $f \colon X \to \mathbb{P}^1_{\mathbb{C}}$. These graphs are called *Dessin's D'Enfants*, which is french for child's drawings.

Secretly, this is just RET. Suppose we have a bipartite graph on X with vertices labeled 0 and 1. Let our basepoint $x = 1/2$ and label the fiber $x_1, \ldots, x_d$. Let γ_0 be the loop at x that goes around 0, and let γ_1 be the loop at x that goes around 1. Note that $\pi_1(\mathbb{P}^1_{\mathbb{C}} \setminus \{0, 1, \infty\}, x)$ is freely generated by γ_0, γ_1. Let $\sigma_0 \in \Sigma_d$ be the permutation given by the following. Around each vertex labeled 0, permute the edges incident to it according to the orientation induced on them by X.

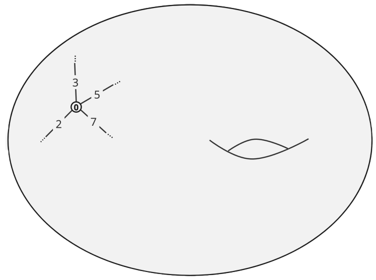

Figure 2: Diagram shows σ_0 has cycle (3 2 7 5) induced by the orientation on X.

The same definition is given to σ_1. Then let $\rho\colon \pi_1(\mathbb{P}^1_{\mathbb{C}} \setminus \{0, 1, \infty\}, x) \to \Sigma_d$ be the homomorphism given by $\gamma_0 \mapsto \sigma_0$ and $\gamma_1 \mapsto \sigma_1$. By Theorem 2.1, there is a branched covering $\varphi\colon X \to \mathbb{P}^1_{\mathbb{C}}$ realizing this monodromy. That is, the action of γ_0 on the fiber over x is σ_0^{-1}, the same for γ_1. One can check that these constructions are indeed inverses of each other. That is, the action of γ_0 on an element of the fiber is permutation of the edge label around the vertex, and thus these two covers have conjugate monodromies (given by a relabeling of edges) and hence are isomorphic. Further, the lift of $(0, 1)$ in the cover corresponding the monodromy constructed from the graph returns the same graph since the path from 0 to 1 is contained in the loop $\gamma_0\gamma_1$, and lifting this loop gives a path on the graph which contains a unique edge.

We can also note the following. Consider $\mathbb{P}^1_{\mathbb{C}} \setminus [0, 1]$, that is the Riemann sphere minus a line segment. This naturally deformation retracts to an open disk containing the point ∞, and its preimage in X is the faces of the dessin. Therefore, the faces correspond to the fiber over ∞.

Note A graph with a single face is a graph with no loops, i.e. a tree. Therefore, trees correspond to Belyi morphisms with a single point of maximal ramification index over ∞. In the case $X = \mathbb{P}^1_{\mathbb{C}}$, these are precisely polynomials branched over only 0 and 1, as we discussed in Example 2.1. These are called Shabat polynomials.

Remark Let g be the genus of X, and $\varphi\colon X \to \mathbb{P}^1_{\mathbb{C}}$ be our degree d Belyi morphism. By Riemann-Hurwitz, we have

$$2g - 2 = -2d + \sum_{x \in X} (e_x - 1)$$

where e_x is the ramification index of x. By the above description, we can see that if x is a point lying over 0 or 1, i.e. a vertex of the dessin, then e_x is the degree of the vertex (since near the vertex, a point has $\deg(x)$ preimages given by one for each edge incident to x). If x is a point lying over ∞, it corresponds to a face of the dessin and we can see that the ramification index e_x is half the number of edges incident to the face. The degree of the map d is the number of edges of the graph, and so if our abstract graph is given by V, E, F with v, e, f vertices, edges, faces resp., we are left with

$$\begin{aligned} 2g - 2 &= -2e + \sum_{x \in V} (\deg(x) - 1) + \sum_{x \in F} (\deg(x)/2 - 1) \\ &= -2e + 2e - v + e - f \\ &= -v + e - f \end{aligned}$$

This is the usual Euler characteristic of the graph, so everything seems to check out numerically.

In the next section, we discuss the more abstract notion of the moduli space of simple branched covers of the Riemann sphere in an effort to find an inductive structure on the number of branch points.

3 Inductive Structure on Simple Branched Covers

In this section, we describe an inductive structure on the space of simple branched coverings of $\mathbb{P}^1_{\mathbb{C}}$ starting with three branch points. Given two coverings $\varphi\colon X \to \mathbb{P}^1_{\mathbb{C}}$ and $\varphi'\colon X' \to \mathbb{P}^1_{\mathbb{C}}$ branched over three points, we construct a nodal curve joining X, X' which projects to $\mathbb{P}^1_{\mathbb{C}}$ with four special points, namely three marked points and a single node given by merging one branch point from each of φ, φ'. Before describing this, we start with preliminary notions of moduli spaces.

3.1 Moduli of Genus g Curves

We give a quick overview of definitions, but a more comprehensive overview is given in [HM98]. Recall that a moduli problem is a functor $F\colon \mathrm{Sch}_{\mathbb{C}} \to \mathrm{Set}$, and we say F is *representable* if there is a scheme M such that F is isomorphic (as a functor) to the functor of points of M, i.e. the functor $h_M : \mathrm{Sch}_{\mathbb{C}} \to \mathrm{Set}$ where $h_M(T) = \mathrm{Hom}(T, M)$. For example, if we want to describe the functor of points of the moduli space of curves of genus g, we could let

$$F(T) = \left\{ \begin{array}{c} C \to T \text{ proper, flat family with} \\ \text{geometric fibers smooth, connected,} \\ \text{curves of genus } g \end{array} \right\} \Big/ \sim$$

where $\sim$ is isomorphism of families, i.e. isomorphisms

$$\begin{array}{ccc} C & \longrightarrow & C' \\ & \searrow \quad \swarrow & \\ & T & \end{array}$$

in the category of such families of curves over T. Note that the image of a moduli functor is usually a set of objects in a category modulo equivalence under isomorphism. This leads us to generalize to the notion of a *stack*, which is roughly a category $\mathcal{X}$ fibered in groupoids along with some extra assumptions, i.e. a category $\mathcal{X}$ and functor $p\colon \mathcal{X} \to \mathrm{Sch}_{\mathbb{C}}$ such that for any $T \in \mathrm{Sch}_{\mathbb{C}}$, the fiber over T (the set of objects $t \in \mathcal{X}$ with $p(t) = T$ and morphisms given by arrows $f\colon t \to t'$ such that $p(f) = \mathrm{id}_T$) is a groupoid. An *algebraic stack* has a smooth, surjective morphism $U \to \mathcal{X}$. If this morphism is étale, we say $\mathcal{X}$ is a *Deligne-Mumford stack* or DM stack, for short. Roughly speaking, a stack is a generalization of a scheme where we allow points to have nontrivial automorphism groups, and DM stacks can be thought of as algebraic stacks with no infinitesimal automorphisms.

Example 3.1 Let X be a scheme over $\mathbb{C}$. Then there is a stack associated to X, denoted $\underline{X}$, whose objects are schemes over X, i.e. morphisms $Y \to X$, and whose morphisms are morphisms of schemes that commute with the maps to X. Then $p\colon \underline{X} \to \mathrm{Sch}_{\mathbb{C}}$ is given by $p(Y \to X) = Y$ since $Y \to X \to \mathrm{Spec}\,\mathbb{C}$ gives Y the structure of a $\mathbb{C}$-scheme, and $p(Y \to Y') = Y \to Y'$. This is why is makes sense to consider maps $U \to \mathcal{X}$, since we really mean U to be the stack associated to the scheme U.

Example 3.2 Let $\mathcal{M}_g$ be the category whose objects are proper, flat morphisms $C \to T$ whose geometric fibers are smooth, connected curves of genus g. The morphisms from $C \to T$ to $C' \to T'$ are pullback diagrams

$$\begin{array}{ccc} C & \longrightarrow & C' \\ \downarrow & & \downarrow \\ T & \longrightarrow & T' \end{array}$$

The functor $p\colon \mathcal{M}_g \to \mathrm{Sch}_{\mathbb{C}}$ takes $C \to T$ to T and a morphism as above to the map $T \to T'$. This gives $\mathcal{M}_g$ the structure of a stack. In fact, it is a Deligne-Mumford stack, so named after the original paper establishing the irreducibility of this space in [DM69].

The notion of an algebraic stack can be seen to be a generalization of a moduli functor in the following sense.

$$\begin{array}{ccc} & & \text{Stacks} \\ & \overset{X \mapsto \underline{X}}{\nearrow} & \Big\downarrow {\scriptstyle X \mapsto (T \mapsto p^{-1}(T)/\text{iso.})} \\ \mathrm{Sch}_{\mathbb{C}} & \xrightarrow[X \mapsto h_X]{} & \text{Moduli Functors} \end{array}$$

This equates our two notions of $\mathcal{M}_g$, since the image of the stack is the functor sending a scheme T to the set of proper, flat families $C \to T$ with geometric fibers that are smooth, connected curves of genus g modulo isomorphism.

Question What is the dimension of $\mathcal{M}_g$?

We can do a quick count of how many parameters determine a genus g curve to try and figure out what it should be, which is what Riemann initially did to try and compute this number. The idea is to count branched covers of the projective line by smooth curves in two ways.

Fix a curve C of genus g, and recall that a degree d map $C \to \mathbb{P}^1_{\mathbb{C}}$ is the same as giving two linearly independent sections of a degree d line bundle $\mathcal{L} \in \mathrm{Pic}^d(C)$ on C. Then $h^1(C, \mathcal{L}) = h^0(C, K_C \otimes \mathcal{L}^\vee)$ by Serre duality, and the latter bundle has degree $2g - 2 - d$. Thus, if $d > 2g - 2$ this vanishes, and so by Riemann-Roch we have $h^0(C, \mathcal{L}) = d + 1 - g$. Thus, there are $d + 1 - g$ first sections we can choose, and $d - g$ second sections after we remove the sections not linearly independent to the first. Hence, there are $2d + 1 - 2g$ degree d covers $C \to \mathbb{P}^1_{\mathbb{C}}$ for fixed $C, \mathcal{L}$. Then there are $\dim \mathrm{Pic}^d(C) = \dim \mathrm{Pic}^0(C) = g$ choices for $\mathcal{L}$, and there are $\dim \mathcal{M}_g$ choices for C, hence the dimension of such branched covers of degree d is

$$2d + 1 - 2g + g + \dim \mathcal{M}_g = 2d + 1 - g + \dim \mathcal{M}_g$$

Now we count a second way. By Riemann-Hurwitz, a cover $C \to \mathbb{P}^1_{\mathbb{C}}$ of degree d has $b = 2d + 2g - 2$ branch points (with multiplicity). By Riemann's existence theorem, this classifies covers and so by equating the two counts, we get $2d + 1 - g + \dim \mathcal{M}_g = 2d + 2g - 2$. Therefore,

$$\dim \mathcal{M}_g = 3g - 3 \tag{1}$$

Now that we know what it should be, we can compute it rigorously using an idea called ***deformation theory***. The idea is to study infinitesimal variations of a curve C by studying the functor of points of $\mathcal{M}_g$ over the dual numbers, $D = \mathrm{Spec}\, \mathbb{C}[\varepsilon]/(\varepsilon^2)$. The following is a summary of the exposition found in Section 7.2 of [Sch23], but a good discussion of the ideas can also be found in [HM98]. We want to show that the tangent space $T_C \mathcal{M}_g$ at a curve C has dimension $3g - 3$. The key insight is the following exercise.

Exercise 13.1.G in [Vak24]. Let $\mathcal{M}$ be a complex scheme. Then there is a natural bijection between $\mathrm{Mor}_{\mathbb{C}}(\mathrm{Spec}\, \mathbb{C}[\varepsilon]/(\varepsilon^2), \mathcal{M})$ to the data of a $\mathbb{C}$-point p and tangent vector $v \in T_p \mathcal{M}$.

The idea is that a map $\mathrm{Spec}\, \mathbb{C}[\varepsilon]/(\varepsilon^2) \to \mathcal{M}$ gives us a point p which is the image of the unique prime ideal (ε), and induces a map on local rings $\mathcal{O}_{\mathcal{M},p} \to \mathbb{C}[\varepsilon]/(\varepsilon^2)$ sending $\mathfrak{m}_{\mathcal{M},p} \mapsto (\varepsilon)$ and $\mathfrak{m}^2_{\mathcal{M},p} \mapsto (\varepsilon)^2 = 0$. Thus, we get a map $\mathfrak{m}_{\mathcal{M},p}/\mathfrak{m}^2_{\mathcal{M},p} \to \mathbb{C}[\varepsilon]/(\varepsilon^2) \to \mathbb{C}$ where the latter map is given by $a + b\varepsilon \mapsto b$. This map is an element of $(\mathfrak{m}_{\mathcal{M},p}/\mathfrak{m}^2_{\mathcal{M},p})^\vee = T_p \mathcal{M}$.

Now elements of $T_C \mathcal{M}_g$ correspond to morphisms $\mathrm{Spec}\, \mathbb{C}[\varepsilon]/(\varepsilon^2) \to \mathcal{M}_g$ with image C. By the description of $\mathcal{M}_g$ as a moduli stack, this is equivalent to flat, proper families $\mathcal{C} \to \mathrm{Spec}\, \mathbb{C}[\varepsilon]/(\varepsilon^2)$ of smooth, connected, genus g curves such that C is the pullback of $\mathcal{C}$ along the closed embedding $\mathrm{Spec}\, \mathbb{C} \to \mathrm{Spec}\, \mathbb{C}[\varepsilon]/(\varepsilon^2)$.

$$\begin{array}{ccc} C & \hookrightarrow & \mathcal{C} \\ \downarrow & & \downarrow \\ \mathrm{Spec}\, \mathbb{C} & \hookrightarrow & \mathrm{Spec}\, \mathbb{C}[\varepsilon]/(\varepsilon^2) \end{array}$$

Figure 3: A ***first-order deformation*** of C.

There are some observations we can make about this diagram [Sch23].

(a) Since $\operatorname{Spec}\mathbb{C}[\varepsilon]/(\varepsilon^2)$ and $\operatorname{Spec}\mathbb{C}$ are points, this is an isomorphism of topological spaces. Since isomorphisms behave well under fiber product, C is isomorphic to $\mathcal{C}$ which means the Zariski open sets of C correspond to the open subsets of $\mathcal{C}$. In fact, affines correspond to affines.

(b) A first-order deformation of an affine scheme U is trivial, namely if $\mathcal{U} \to \operatorname{Spec}\mathbb{C}[\varepsilon]/(\varepsilon^2)$ is a deformation of $U \to \operatorname{Spec}\mathbb{C}$, then $\mathcal{U} = U \times \operatorname{Spec}\mathbb{C}[\varepsilon]/(\varepsilon^2)$. Thus, if we choose a cover U_i of C, then each corresponding open subset in the cover $\mathcal{U}_i$ of $\mathcal{C}$ will be a trivial deformation. Therefore, the data of a deformation $\mathcal{C} \to \operatorname{Spec}\mathbb{C}[\varepsilon]/(\varepsilon^2)$ is the data of transition maps

$$\varphi_{ij} : U_{ij} \times \operatorname{Spec}\mathbb{C}[\varepsilon]/(\varepsilon^2) \to U_{ij} \times \operatorname{Spec}\mathbb{C}[\varepsilon]/(\varepsilon^2) \tag{2}$$

describing how $\mathcal{U}_i$ and $\mathcal{U}_j$ glue on the overlaps. They must satisfy the cocycle condition $\varphi_{jk} \circ \varphi_{ij} = \varphi_{ik}$ and

$$\varphi_{ij}|_{U_{ij} \times \operatorname{Spec}\mathbb{C}} = \mathrm{id}_{U_{ij}} \tag{3}$$

since this is how they are glued in the fiber over $\operatorname{Spec}\mathbb{C}$, which is C.

(c) Since C is separated, $U_{ij} = \operatorname{Spec} B$ is affine and so $U_{ij} \times \operatorname{Spec}\mathbb{C}[\varepsilon]/(\varepsilon^2) = \operatorname{Spec} B[\varepsilon]/(\varepsilon^2)$. One can check then that the isomorphisms φ_{ij} are of the form $B[\varepsilon]/(\varepsilon^2) \to B[\varepsilon]/(\varepsilon^2)$ sending $a+\varepsilon b \mapsto a+\varepsilon(b+\eta_{ij}(a))$ where $\eta_{ij} \in \mathrm{Der}_{\mathbb{C}}(B,B)$ is a $\mathbb{C}$-linear derivation, or equivalently a section $v_{ij} \in H^0(U_{ij}, T_{U_{ij}})$.

(d) Therefore, for a fixed open cover U_i of C we can associate a system

$$(v_{ij} \in H^0(U_{ij}, T_{U_{ij}}))_{ij}$$

of sections on the overlaps $U_i \cap U_j$ to any first order deformation $\mathcal{C}$. The fact that φ_{ij} satisfy a cocycle condition is equivalent to v_{ij} satisfying a cocycle condition. Therefore, first-order deformations $\mathcal{C}$ of C are equivalent to elements

$$(v_{ij})_{ij} \in H^1(C, T_C)$$

in the first Čech cohomology group of the tangent bundle on C.

Therefore,

$$T_C \mathcal{M}_g = \left\{ \begin{array}{l} \mathcal{C} \to \operatorname{Spec}\mathbb{C}[\varepsilon]/(\varepsilon^2) \text{ first-order} \\ \text{deformation of } C \end{array} \right\} = H^1(C, T_C) \tag{4}$$

Then by Serre duality, $h^1(C,T_C) = h^0(C, \Omega_C^{\otimes 2})$, and by Riemann-Roch $h^0(C,\Omega_C^{\otimes 2}) - h^1(C,\Omega_C^{\otimes 2}) = (4g-4)+1-g = 3g-3$. Note that, again by Serre duality, $h^1(C,\Omega_C^{\otimes 2}) = h^0(C,T_C)$ and $\deg T_C = 2-2g < 0$ for $g > 1$, so

$$\dim T_C \mathcal{M}_g = h^0(C, T_C) = 3g-3 \tag{5}$$

This concludes our discussion of $\mathcal{M}_g$, but there is a rich theory of higher order deformations and obstructions, as well as the discussion of the moduli space of *stable nodal curves* that compatifies $\mathcal{M}_g$. Now we discuss the moduli space of simple branched covers, which can be used to show $\mathcal{M}_g$ is irreducible.

3.2 Hurwitz Schemes

We follow the exposition given in [Cas19]. The Hurwitz scheme is a quasiprojective variety parameterizing simple branched coverings of $\mathbb{P}^1$. By a simple branched cover, we mean a cover $X \to \mathbb{P}^1$ such that, over any branch point, there is exactly one point that ramifies to degree 2 and the rest of the points in the preimage are unramified. The benefit of this description, then, is partly that by Riemann-Hurwitz,

$$2g-2 = -2d+b \tag{6}$$

where b is the number of branch points.

Let X be a connected compact Riemann surface, and let $X^{(b)}$ denote the b-fold symmetric product of X with itself, i.e. X^b quotiented by the action of Σ_b, viewed as a complex manifold of dimension b. Let Δ_b be the discriminant locus, i.e. the closed subspace of unordered b-tuples where not all the points are distinct. Then an element $B \in X^{(b)} \setminus \Delta_b$ is a subset of b distinct points of X. Let $\mathcal{H}_d(B)$ be the set of connected

d-sheeted branched covers of X with branch locus B. Let $H_{d,b}$ be the set of d-sheeted branched covers of X with b branch points. Then there is a natural map $\delta \colon H_{d,b} \to X^{(b)} \setminus \Delta_b$ given by mapping a cover to its branch locus. By definition, $\delta^{-1}(B) = H_d(B)$. From now on, we let $X = \mathbb{P}^1$ and denote $H^{d,b}$ by the set of connected d-sheeted branched covers of $\mathbb{P}^1$ with b branch points.

By Proposition 2.2. in [Cas19], there is a topology on $H^{d,b}$ such that δ is a covering map. We will show that $H^{d,b}$ is a connected quasiprojective variety (hence irreducible) and the natural map $H^{d,b} \to M_g$ given by taking a branched cover $C \to \mathbb{P}^1$ to C is surjective for $d \geq g+1$, where $b = 2d + 2g - 2$ as in (6). It follows that M_g is irreducible.

Theorem 3.1 The space of simple degree d branched covers of $\mathbb{P}^1$ with exactly b branch points $H^{d,b}$ is connected.

We first reduce this to a combinatorial problem. Note that $H^{d,b}$ is connected if and only if any two points in $\delta^{-1}(B) = H_d(B)$ are connected by a path. Pick some $x \in \mathbb{P}^1$ away from B, and pick a loop C that starts at x and passes through the points of B. This gives us an ordering $w_1, \ldots, w_b$ of the points in B.

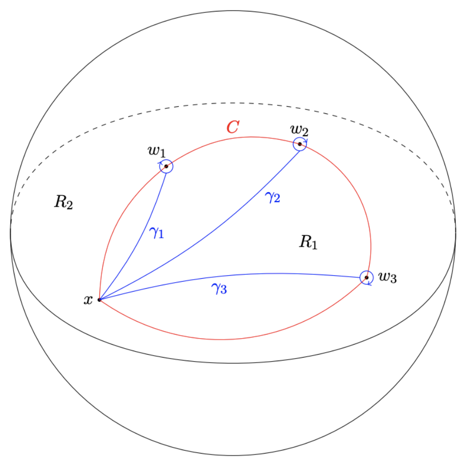

Figure 4: The loop C at a basepoint x through points of B [Cas19].

C bounds two open sets in $\mathbb{P}^1$, labeled R_1 and R_2. We can choose small non-intersecting loops γ_i at x that travel in a clockwise loop around w_i through R_1. Then $\pi_1(\mathbb{P}^1 \setminus B, x) = \langle \gamma_i \mid \prod \gamma_i = 1 \rangle$. Recall that Riemann's existence theorem gives a bijection between monodromies $\rho \colon \pi_1(\mathbb{P}^1 \setminus B, x) \to \Sigma_d$ with transitive image and connected coverings of $\mathbb{P}^1$ branched over B. Then, under this equivalence, the requirement that the covering be simple equates to required that the image of a generator γ_i in Σ_d is a transposition, since the map is given by the action of path lifting on the labeled vertices and, since the one ramification point has degree 2, the size of the orbit (or the length of the cycle) is therefore also 2. Let $A_{d,b}$ be the set of ordered b-tuples $[t_1, \ldots, t_b]$ of transpositions that generate a transitive subgroup of Σ_d and such that $t_1 \cdots t_b = 1$. It is clear that $H^{d,b}$ is in bijection with $A_{d,b}$ modulo the equivalence relation given by simultaneous conjugation by an element of Σ_d, since this is exactly the equivalence given by RET. Fix a tuple $[t_1, \ldots, t_b] \in A_{d,b}$.

Let $1 \leq i \leq b$. Choose a path ψ_i going from w_i to w_{i+1} in R_2, and a path ψ_{i+1} going from w_{i+1} to w_i in R_1. Then construct a loop $\psi \in \pi_1(\mathbb{P}^b \setminus \Delta_b, B)$ by the formaula

$$\psi(s) = \{w_1, \ldots, w_{i-1}, \psi_i(s), \psi_{i+1}(s), w_{i+2}, \ldots, w_b\}$$

Then ψ lifts to a path $\tilde{\psi}$ in $H^{d,b}$ such that

$$\tilde{\psi}(0) = [t_1, \ldots, t_b]$$

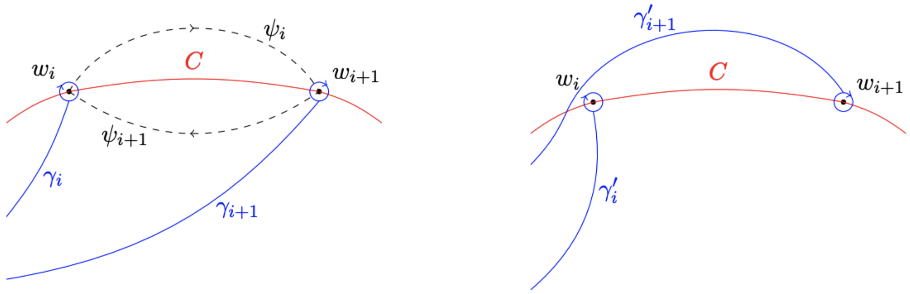

Figure 5: Computing $\phi(1)$ by deforming γ_i along ψ_i [Cas19].

To compute $\phi(1)$, deform γ_i along ψ_i through a continuous family of loops ending at $\gamma'_{i+1} \in \pi_1(\mathbb{P}^1 \setminus B, x)$. Construct γ'_i similarly, and note that $\gamma'_{i+1} = \gamma_i \gamma_{i+1} \gamma_i^{-1}$ and $\gamma'_i = \gamma_i$. It follows that

$$\phi(1) = [t_1, \ldots, t_{i-1}, t_{i+1}, t_{i+1}^{-1} t_i t_{i+1}, t_{i+2}, \ldots, t_b]$$

All that remains is to show that any two elements of $A_{d,b}$ are related by a sequence of transformations of the above form. This is Theorem 3.3 in [Cas19], but we omit it as it is not relevant to the discussion of the geometry of $H^{d,b}$. This concludes the proof of Theorem 3.1.

4 Conclusion

The continuation of this discussion lies both in formulating the compactification $\overline{H}^{d,b}$ of the Hurwitz scheme to the space of *admissible covers*, in such a way that we have a fiber diagram.

$$\begin{array}{ccc} \overline{H}^{d,b} & \longrightarrow & \overline{M}_{g,n} \\ \downarrow & & \downarrow \\ H^{d,b} & \longrightarrow & M_g \end{array}$$

Above, $\overline{M}_{g,n}$ is the moduli space of n-pointed stable nodal curves of genus g. This notion compactifies the space M_g, since it resolves the pathological example of a sequence of curves of the same genus whose limit is not a compact curve. Instead, it requires that the curves have enough fixed points such that the automorphism group is finite, and therefore excludes such pathological examples and is compact. A similar construction exists for $H^{d,b}$ but is more complicated to construt and requires further discussion.

References

[Cas19] Robert Cass. "Hurwitz Schemes". MA thesis. Harvard University, 2019.

[DM69] Pierre Deligne and David Mumford. "The irreducibility of the space of curves of given genus". In: *Publications Mathematiques de l'IHES* 36 (1969).

[HM98] Joe Harris and Ian Morrison. *Moduli of Curves*. Vol. 187. New York: Springer Verlag, 1998.

[Sch23] Johannes Schmitt. *The Moduli Space of Curves*. 2023. url: https://johannesschmitt.gitlab.io/ModCurves/Script.pdf.

[Sza09] Tamás Szamuely. *Galois Groups and Fundamental Groups*. Cambridge University Press, 2009.

[Vak24] Ravi Vakil. *Foundations of Algebraic Geometry*. 2024. url: https://math.stanford.edu/~vakil/216blog/FOAGfeb2124public.pdf.