HURWITZ SCHEMES

ROBERT CASS

ABSTRACT. We give an introduction to the theory of Hurwitz schemes. Our main focus is on proving that Hurwitz schemes are irreducible. We then use this fact to show that the moduli space of curves of genus g is irreducible. We conclude with a glimpse of the compactifications of Hurwitz schemes.

1. INTRODUCTION

A Hurwitz scheme is a variety which parameterizes simple branched coverings of $\mathbb{P}^1$. Recall that a connected branched covering $f \colon C \to \mathbb{P}^1$ is *simple* if for every branch point $w \in \mathbb{P}^1$, there is exactly one point in $f^{-1}(w)$ with ramification index 2, and all other fibers are unramified. Given a simple branched covering, one can associate three discrete invariants: the genus g of C, the degree d of the covering, and the number b of branch points in $\mathbb{P}^1$. By the Riemann-Hurwitz formula, these are related by the equation $b = 2d + 2g - 2$. Consequently, we will construct a scheme $\mathcal{H}^{d,b}$ that parameterizes simple d-sheeted branched coverings of $\mathbb{P}^1$ with b branch points.

In Section 2 we first construct $\mathcal{H}^{d,b}$ as a complex manifold by realizing $\mathcal{H}^{d,b}$ as a finite cover $\delta \colon \mathcal{H}^{d,b} \to \mathbb{P}^b - \Delta_b$ of the space of possible branch loci. This involves giving a complete description of a branched cover from the associated monodromy representation.

By forgetting the map to $\mathbb{P}^1$, the Hurwitz scheme $\mathcal{H}^{d,b}$ admits a map to the coarse moduli space M_g of curves of genus g. For $d \geq g + 1$ this map is dominant, and so we can hope to learn about M_g by studying $\mathcal{H}^{d,b}$, which is significantly easier to construct. We will show in Theorem 3.1 that $\mathcal{H}^{d,b}$ is connected, which we will later use to prove that M_g is irreducible. Our approach dates back to the 19th century and involves showing $\pi_1(\mathbb{P}^b - \Delta_b)$ acts transitively on the fibers of δ. Ultimately this amounts to a purely combinatorial statement.

In Section 4 we give a moduli interpretation of the functor of points of $\mathcal{H}^{d,b}$. After Theorem 3.1, the key input to showing that M_g is irreducible is the fact that if C is a curve of genus g, then for $d \geq g + 1$ there is a simple d-sheeted branched covering $C \to \mathbb{P}^1$. We give a careful proof of this result, following the argument given by Fulton in [Ful69]. The proof makes use of the Picard scheme of C.

In order to apply methods of projective geometry, it is often fruitful to compactify moduli spaces. One way to compactify M_g is by allowing stable curves instead of only smooth curves. This idea suggests that one should compactify $\mathcal{H}^{d,b}$ by considering nodal covers of stable curves of genus 0. In Section 5 we introduce stable curves and use them to define admissible covers. This allows us to give a moduli description of a compactification of $\mathcal{H}^{d,b}$.

2. Construction of $\mathcal{H}^{d,b}$

In this section we treat the theory of Hurwitz schemes from the complex analytic perspective, following the modern treatment in [Ful69]. Let X be a compact Riemann surface. If $Y \to X$ is a topological covering then Y naturally has a complex analytic structure coming from X. Recall that a morphism between covers is a map $Y \to Y'$ such that the following diagram commutes.

A morphism between covers is automatically analytic for the complex structures on Y and Y' coming from X. Two covers Y and Y' are *equivalent* if there is an isomorphism $Y \to Y'$ that is also a map of covers. Given two groups G and H, we define two homomorphisms α, $\beta \colon G \to H$ to be *equivalent* if there exists $h \in H$ such that $\beta(g) = h\alpha(g)h^{-1}$ for all $g \in G$. Let S_d be the symmetric group acting on $\{1, \ldots, d\}$.

Proposition 2.1. Let $B \subset X$ be a finite subset, $x \in X - B$, and let d be a positive integer. Then there is a natural bijection between the following sets:

(1) Equivalence classes of connected d-sheeted branched coverings of X with branch locus contained in B.
(2) Equivalence classes of homomorphisms $\pi_1(X - B, x) \to S_d$ whose images act transitively on $\{1, \ldots, d\}$.

Proof. Given a connected d-sheeted branched cover $f \colon Y \to X$ with branch locus contained in B, let $Y^\circ = f^{-1}(X - B)$. Then Y° is a covering space of $X - B$. Pick a labelling of the preimages $y_1, \ldots, y_d \in f^{-1}(x)$. For a loop $\gamma \in \pi_1(X - B, x)$ and a given y_i, the homotopy lifting property allows us to lift γ to a path from y_i to some y_j. This procedure induces a permutation of the y_i, and gives rise to a homomorphism $\pi_1(X - B, x) \to S_d$. The equivalence class of this homomorphism is independent of the choice of the ordering on $f^{-1}(y)$. As Y° is connected the image of this homomorphism acts transitively on $\{1, \ldots, d\}$.

Conversely, given a homomorphism $\pi_1(X - B, x) \to S_d$ whose images act transitively on $\{1, \ldots, d\}$, let $K \subset \pi_1(X - B, x)$ be the subgroup whose image fixes 1. By the theory of covering spaces there is a d-sheeted connected cover $Y^\circ \to X - B$ associated to K. Moreover, the procedure described above will produce from $Y^\circ \to X - B$ the original homomorphism $\pi_1(X - B, x) \to S_d$ up to equivalence. To complete Y° to a branched cover of X, choose disjoint open disks D_i centered at each of the points $w_i \in B$. Let $D_i^\circ = D_i - \{w_i\}$. As the fundamental group of D_i° is infinite cyclic, then, $f^{-1}(D_i^\circ)$ consists of a disjoint union of copies $D_{i,1}^\circ, \ldots, D_{i,k}^\circ$ of D_i°. Up to conformal equivalence, over each $D_{i,k}^\circ$ the map f is of the form $z \mapsto z^{e_{i,k}}$ for a unique positive integer $e_{i,k}$. If we add a center $p_{i,k}$ to each $D_{i,k}^\circ$, the map $D_{i,k}^\circ \to D_i^\circ$ extends uniquely to an analytic map $D_{i,k} \to D_i$ via the formula $z \mapsto z^{e_{i,k}}$.

This shows how to extend f uniquely to a branched cover $f\colon Y \to X$. The integer $e_{i,k}$ is the ramification index of f at $p_{i,k}$. $\qquad\square$

Fix a positive integer b, and let $X^{(b)}$ be the b-fold symmetric product of X with itself, viewed as a complex manifold of complex dimension b. Let $\Delta_b \subset X^{(b)}$ be the discriminant locus, that is, the closed subspace where the points are not all distinct. Then we may naturally view a point $B \in X^{(b)} - \Delta_b$ as a subset of X consisting of b distinct points. Let $\mathcal{H}_d(B)$ be the set of connected d-sheeted branched covers of X with branch locus equal to B. Finally, let $\mathcal{H}_{d,b}$ be the set of connected d-sheeted branched covers of X with exactly b branch points. There is a natural map $\delta\colon \mathcal{H}_{d,b} \to X^{(b)} - \Delta_b$ which sends a branched cover to its branch locus. By definition, $\delta^{-1}(B) = \mathcal{H}_d(B)$.

Proposition 2.2. There is a natural topology on $\mathcal{H}_{d,b}$ such that $\delta\colon \mathcal{H}_{d,b} \to X^{(b)} - \Delta_b$ is a covering map.

Proof. Let $B \in X^{(b)} - \Delta_b$ and pick disjoint open disks $U_i \subset X$ around the points $w_i \in B$. Then the products of the form $N(U_1, \ldots, U_b) := \prod_{i=1}^{b} U_i$ form a basis for the topology on $X^{(b)} - \Delta_b$. Now pick another point $B' \in N(U_1, \ldots, U_b)$. We will produce a natural bijection between $\mathcal{H}_d(B)$ and $\mathcal{H}_d(B')$. Let U be the union of the U_i. As $X - U$ is a deformation retract of $X - B$, then the inclusion $X - U \subset X - B$ induces an isomorphism on fundamental groups. The same is true when we replace B by B'. Thus by Proposition 2.1 we get a natural bijection $\mathcal{H}_d(B) \cong \mathcal{H}_d(B')$. Now fix a point $y \in \mathcal{H}_d(B)$. As we let $B' \in N(U_1, \ldots, U_b)$ vary and look at the images of y under the bijections $\mathcal{H}_d(B) \cong \mathcal{H}_d(B')$, we get a subset of $\mathcal{H}_{d,b}$ that is naturally in bijection with $N(U_1, \ldots, U_b)$ via the map δ. We give this subset the topology coming from $N(U_1, \ldots, U_b)$. As our bijections $\mathcal{H}_d(B) \cong \mathcal{H}_d(B')$ do not depend on the choice of the U_i, this procedure globalizes and allows us to equip $\mathcal{H}_{d,b}$ with the desired topology. $\qquad\square$

From now on we will only consider the case $X = \mathbb{P}^1$. Then $X^{(b)} = \mathbb{P}^b$ (see [Ful69, §5]). For $B \in \mathbb{P}^b - \Delta_b$, let $\mathcal{H}^d(B)$ be the set of equivalence classes of simple d-sheeted branched coverings of $\mathbb{P}^1$ with branch locus equal to B. Let $\mathcal{H}^{d,b}$ be the set of all simple d-sheeted branched covers of $\mathbb{P}^1$ with b branching points. This is a closed and open submanifold of $\mathcal{H}_{d,b}$.

3. $\mathcal{H}^{d,b}$ IS CONNECTED

3.1. **Geometry.** Our goal in this section is to prove the following theorem.

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

We will again follow the treatment in [Ful69], though the results in this section are due to Clebsch [Cle73] and Hurwitz [Hur91]. We first reduce the problem to one that is purely combinatorial. Recall that we have a covering map $\delta\colon \mathcal{H}^{d,b} \to \mathbb{P}^b - \Delta_b$. Fix $B \in \mathbb{P}^b - \Delta_b$. Then $\mathcal{H}^{d,b}$ is connected if and only if any two points in $\delta^{-1}(B) = \mathcal{H}^d(B)$ are connected by a path.

Pick $x \in \mathbb{P}^1$ distinct from the points in B and pick a loop C that starts at x and passes successively through the points of B before returning back to x. This induces an ordering
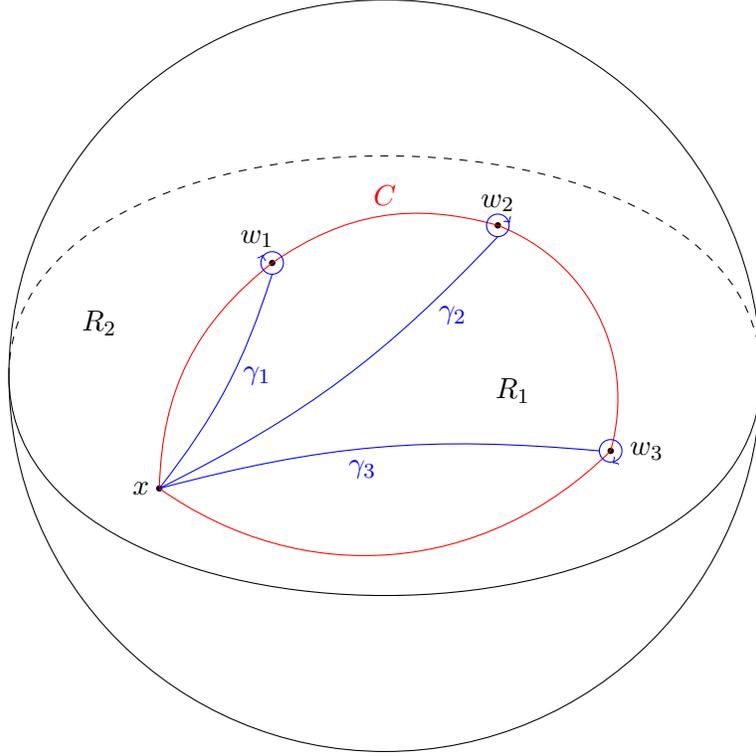
FIGURE 1

$w_1, \ldots, w_b$ of the points in B. The loop C bounds two open sets in $\mathbb{P}^1$; label them R_1 and R_2. Choose non-intersecting loops γ_i starting at x which travel toward w_i in R_1 and then each make a small clockwise loop around w_i before traveling back along the same path back to x. The fundamental group $\pi_1(\mathbb{P}^1 - B, x)$ is isomorphic to the free group generated by the γ_i modulo the relation $\prod_{i=1}^{b} \gamma_i = 1$.

Under the bijection of Proposition 2.1, the simple branched covers correspond to homomorphisms $\varphi \colon \pi_1(\mathbb{P}^1 - B, x) \to S_d$ such that $\varphi(\gamma_i)$ is a transposition for all i. Let $A_{d,b}$ be the set of ordered b-tuples $[t_1, \ldots, t_b]$ of transpositions in the symmetric group S_d such that $\prod_{i=1}^{b} t_i = \mathrm{id}$ and the subgroup of S_d generated by the t_i acts transitively on $\{1, \ldots, d\}$. Then $\mathcal{H}^{d,b}$ is naturally in bijection with the set $A_{d,b}$ modulo the equivalence relation which comes from simultaneous conjugation by elements of S_d. Fix such a tuple $[t_1, \ldots, t_b] \in A_{d,b}$.

Let $1 \le i < b$. 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 $\psi \in \pi_1(\mathbb{P}^b - \Delta_b, B)$ by the formula

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

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

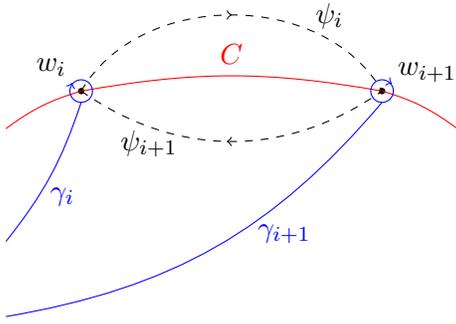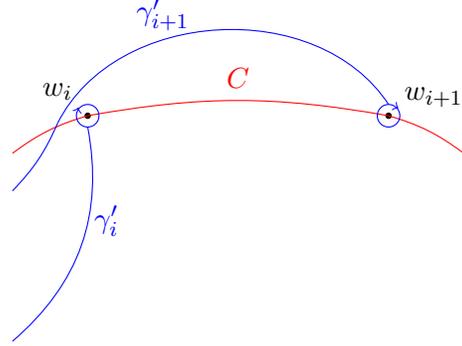$$\tilde{\psi}(0) = [t_1, \ldots, t_b].$$

FIGURE 2

FIGURE 3

To determine $\tilde{\psi}(1)$, deform γ_i along ψ_i through a continuous family of loops ending at $\gamma'_{i+1} \in \pi_1(\mathbb{P}^1 - B, x)$ as shown below. Construct γ'_i similarly. It is not hard to check that $\gamma'_{i+1} = \gamma_i \gamma_{i+1} \gamma_i^{-1}$ and $\gamma'_i = \gamma_i$ in $\pi_1(\mathbb{P}^1 - B, x)$. It follows that

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

Theorem 3.1 now follows from Theorem 3.3 in the next section, which says that any two elements in $A_{d,b}$ are related by a sequence of transformations of the above form.

3.2. **Combinatorics.** Let b and d be positive integers. Consider the collection of ordered b-tuples $[t_1, \ldots, t_b]$ of transpositions in the symmetric group S_d.

Definition 3.2. An *elementary move of type I* sends the tuple $[t_1, \ldots, t_b]$ to

$$[t_1, \ldots, t_{i-1}, t_{i+1}, t_{i+1}^{-1} t_i t_{i+1}, t_{i+2}, \ldots, t_b].$$

We define two tuples t, t' to be *equivalent*, and we write $t \sim t'$, if they are related by a sequence of elementary moves of type I.

As in Section 3.1, we let $A_{d,b}$ be the set of ordered b-tuples $[t_1, \ldots, t_b]$ of transpositions in the symmetric group S_d such that $\prod_{i=1}^b t_i = \mathrm{id}$ and the subgroup of S_d generated by the t_i acts transitively on $\{1, \ldots, d\}$. Elementary moves of type I induce well-defined automorphisms of $A_{d,b}$.

Theorem 3.3. Suppose $t \in A_{d,b}$. Then

$$t \sim [\underbrace{(12), (12), \ldots, (12)}_{b-2(d-2) \text{ times}}, (23), (23), (34), (34), \ldots, ((d-1)d), ((d-1)d)].$$

Though we will follow the the arguments in [Moc95], the results in this section were originally published in 1873 [Cle73]. To prove Theorem 3.3 it will be convenient to introduce a second type of elementary move.

Definition 3.4. An *elementary move of type II* sends the tuple $[t_1, \ldots, t_b]$ to

$$[t_1, \ldots, t_{i-1}, t_i t_{i+1} t_i^{-1}, t_i, t_{i+2}, \ldots, t_b].$$

The original tuple $[t_1, \ldots, t_b]$ can be obtained by applying an elementary move of type I to the above expression, so elementary moves of type II preserve equivalence classes.

Lemma 3.5. Let $\alpha \in \{1, \ldots, d\}$, and let $t = [t_1, \ldots, t_b]$ be a sequence of transpositions in S_d such that the subgroup generated by $t_1, \ldots, t_b$ acts transitively on $\{1, \ldots, d\}$. Then t is equivalent to a tuple whose first $d-1$ transpositions are of the form

$$(\alpha x_1), (x_1 x_2), \ldots, (x_{d-2} x_{d-1}),$$

where $\{\alpha, x_1, \ldots, x_{d-1}\} = \{1, \ldots, d\}$.

Proof. Let $x_0 = \alpha$. By the transitivity assumption there is some transposition $t_i = (\alpha x_1)$ in t containing α. By applying elementary moves of type I to move t_i to the front we can assume $t_1 = (\alpha x_1)$. Now suppose we have found $t' \sim t$ such that the first r transpositions of t' are of the form

$$(\alpha x_1), (x_1 x_2), \ldots, (x_{r-1} x_r),$$

where the x_i are distinct from each other and α. If $r < d-1$, we will show how to find an equivalent tuple whose first $r+1$ transpositions are of the desired form. By transitivity there is some $t_i' = (yz)$ with $y \in \{\alpha, x_1, \ldots, x_r\}$ and $z \notin \{\alpha, x_1, \ldots, x_r\}$. If $y = x_r$ then we use elementary moves of type I to move $(x_r z)$ to the right of $(x_{r-1} x_r)$. Otherwise, if $y = x_i$ for $0 \le i < r$ then we use elementary moves of type I to move $(x_i z)$ to the right of $(x_i x_{i+1})$. This does not disturb the $(x_j x_{j+1})$ for $j \ge i+1$. Now we are done in this case because

$$[(x_i x_{i+1}), (x_i z)] \sim [(x_i z), (z x_{i+1})].$$

$\square$

Lemma 3.6. Let $\alpha, \beta \in \{1, \ldots, d\}$ be distinct and let t be as in Lemma 3.5. Then there exists a tuple t' equivalent to t such that $t_1' = (\alpha \beta)$.

Proof. By Lemma 3.5 we can assume the first r transpositions of t are of the form

$$(\alpha x_1), (x_1 x_2), \ldots, (x_{r-1} \beta)$$

where the x_i are distinct from each other and α, β. If $r = 1$ we are done. Otherwise, note that

$$[(\alpha x_1), (x_1 x_2)] \sim [(\alpha x_2), (\alpha x_1)].$$

This takes care of the case $r = 2$, and in general we can use elementary moves of type II to move $(x_1 x_2)$ to the right of $(x_{r-1} \beta)$, thus reducing r by one. Now we conclude by induction on r. $\square$

Proof of Theorem 3.3. We will first show that we can write

$$t \sim [\underbrace{(12), (12), \ldots, (12), (12)}_{r}, t_{r+1}', \ldots, t_b'] \tag{3.1}$$

where r is even and t_i' does not contain 1 for $i > r$. By Lemma 3.6 we may assume $t_1 = (12)$. As $\prod_{i=1}^b t_i = \mathrm{id}$, then after removing t_1 the remaining transpositions still act transitively on $\{1, \ldots, d\}$, so by applying Lemma 3.6 again we may assume $t_2 = (12)$ as well.

Now by applying elementary moves of type I to move things to the left we can arrange so that there is some integer r such that t_i contains 1 if and only if $i \leq r$. We claim that we can arrange so that $t_3 = t_4 = \cdots = t_r$ (for possibly different r). Indeed, if $a \neq b$ then

$$[(1a), (1b)] \sim [(1b), (ab)].$$

Now moving (ab) to the right with elementary moves of type II, and noting that

$$[(ab), (1a)] \sim [(1b), (ab)],$$

we can reduce r by 1. Proceeding in this manner we eventually get $t_3 = t_4 = \cdots = t_r$. As $\prod_{i=1}^{b} t_i = \mathrm{id}$, then r must be even. If $r > 0$, suppose $t_3 = (1a)$ with $a \neq 2$. Then

$$[(12), (1a), (1a)] \sim [(1a), (2a), (1a)] \sim [(1a), (12), (2a)] \sim [(12), (2a), (2a)].$$

Thus if $r = 4$ we have reached the form (3.1). Otherwise, by moving $(2a)$ twice to the right we get that

$$[(2a), (2a), (1a), (1a)] \sim [(2a), (12), (12), (2a)].$$

Thus, after moving the two new copies of (12) to left we can assume $t_i = (12)$ for $1 \leq i \leq 4$. Now by repeatedly applying the same algorithm as in the beginning of this paragraph, we eventually reach a tuple of the form (3.1).

Once we have an equivalence as in (3.1), we observe that the the tuple $[t'_{r+1}, t'_{r+2}, \ldots, t'_b]$ acts transitively on $\{2, \ldots, d\}$ and satisfies $\prod_{i=r+1}^{b} t'_i = \mathrm{id}$. Thus, by applying the same algorithm as above we get

$$t \sim [\underbrace{(12), \ldots, (12)}_{r_1}, \underbrace{(23), \ldots, (23)}_{r_2}, \ldots, \underbrace{((d-1)d), \ldots, ((d-1)d)}_{r_{d-1}}].$$

At the moment we have no control over the size of the blocks, except that $r_i > 0$ is even for all i. By repeated application of the following lemma, we can ensure $r_i = 2$ for $i \geq 1$. □

Lemma 3.7. Let a, b, $c \in \{1, \ldots, d\}$ be distinct. Then

$$[(ab), (ab), (bc), (bc), (bc), (bc)] \sim [(ab), (ab), (ab), (ab), (bc), (bc)].$$

Proof. We sketch some of the steps and leave the details to the reader. To ease the notation we will write 1, 2, 3 instead of a, b, c. First, by moving t_5 to the front we have

$$[(12), (12), (12), (12), (23), (23)] \sim [(23), (13), (13), (13), (13), (23)].$$

Moving t_2 to the left and t_5 to the right gives

$$[(13), (12), (13), (13), (12), (13)].$$

Applying this operation one more time produces

$$[(12), (23), (13), (13), (23), (12)].$$

After moving each (13) to the left once we get

$$[(12), (13), (13), (23), (23), (12)].$$

We leave it to the reader to show that

$$[(23), (23), (12)] \sim [(13), (13), (12)].$$

Now the lemma follows by moving (12) to the left four times, which gives

$$[(12), (13), (13), (13), (13), (12)] \sim [(12), (12), (23), (23), (23), (23)].$$

$\square$

4. APPLICATION TO MODULI OF CURVES

We would like to upgrade $\mathcal{H}^{d,b}$ to a scheme by describing its functor of points. To do this we must first define a family of branched covers of $\mathbb{P}^1$, following [Ful69]. All of our schemes in this section will be locally noetherian $\mathbb{C}$-schemes. Let us fix a finite flat morphism $f \colon Y \to X$. Every point $x \in X$ has an affine neighborhood $\mathrm{Spec}(A)$ such that $f^{-1}(\mathrm{Spec}(A))$ corresponds a finite free A-module B of rank $n(x)$. The function $x \mapsto n(x)$ is locally constant on X, and if it is constant we call the integer n the *degree* of f.

Now suppose B is a finite free A-algebra. The multiplication by b map $B \xrightarrow{b} B$ is an A-linear map of finite free A-modules, so it has a trace, which we denote $\mathrm{tr}(b)$. More generally, we can construct a trace morphism $\mathrm{tr} \colon f_*(\mathcal{O}_Y) \to \mathcal{O}_X$. This is a map of $\mathcal{O}_X$-modules.

The *discriminant* of $f \colon Y \to X$ is the locally principal ideal sheaf on X defined as follows. Over an open affine $\mathrm{Spec}(A) \subset X$ with preimage corresponding to a free A-module B of rank n, let b_i be a basis for B over A. Then the determinant of the matrix $\mathrm{tr}(b_i b_j)$ is independent of the choice of basis up to multiplication by a unit in A. We define the discriminant of f by requiring that its restriction to $\mathrm{Spec}(A)$ is the ideal sheaf generated by $\det(\mathrm{tr}(b_i b_j))$. Let $\delta(f)$ be the closed subscheme of X defined by the the discriminant ideal. We say that f is a *covering* if $\delta(f)$ is a Cartier divisor on X. The support of $\delta(f)$ is the *branch locus* of f. Note that f is étale if and only if $\delta(f)$ is empty.

Example 4.1. Let $X = \mathrm{Spec}(A)$ be affine, and let $Y = \mathrm{Spec}(B)$ where $B = A[T]/(T^2 - aT - b)$ for some $a, b \in A$. Let $f \colon Y \to X$ correspond to the inclusion $A \subset B$. Note that B is a free A-module of rank two with basis $1, T$. One can compute $\mathrm{tr}(1) = 2$, $\mathrm{tr}(T) = a$, and $\mathrm{tr}(T^2) = a^2 + 2b$. Then $\delta(f)$ is the closed subscheme of X defined by the ideal $(a^2 + 4b)$. Thus f is a covering of degree 2 if and only if $a^2 + 4b$ is not a zero-divisor.

Now suppose S is a fixed scheme. Then we have the projective space $\mathbb{P}^1_S$ over S. Recall that a relative Cartier divisor on $\mathbb{P}^1_S/S$ is an effective Cartier divisor $D \subset \mathbb{P}^1_S$ such that D is finite and flat over S. We call D a divisor of degree b on $\mathbb{P}^1$ over S if it has degree b over S. Let Div^b be the functor on schemes which sends S to the set $\mathrm{Div}^b(S)$ of relative Cartier divisors of degree b on $\mathbb{P}^1_S/S$. This functor is represented by $\mathbb{P}^b$ ([Ful69, 5.2]).

A divisor $D \in \mathrm{Div}^b(S)$ is *simple* if its geometric fibers over points in S consist of b distinct points. Let SDiv^b be the functor which sends S to the set of $\mathrm{SDiv}^b(S)$ of simple relative Cartier divisors of degree b on $\mathbb{P}^1_S/S$. There is a closed subscheme $\Delta_b \subset \mathbb{P}^b$ such that SDiv^b is represented by $\mathbb{P}^b - \Delta_b$. To describe Δ_b, note that to a point $[x_0 : x_1 : \ldots : x_b]$ in $\mathbb{P}^b(\mathbb{C})$ we can associate (up to scaling) a degree b polynomial $x_0 X^b + x_1 X^{b-1} + \cdots + x_b$. The closed subscheme Δ_b is defined by the vanishing locus of the discriminant of this polynomial.

A simple d-sheeted branched covering of $\mathbb{P}^1_S$ with b branch points is a covering $f \colon Y \to \mathbb{P}^1_S$ of degree d with geometrically connected fibers over S such that the discriminant $\delta(f)$ is simple relative Cartier divisor of degree b on $\mathbb{P}^1_S/S$. The notion of equivalence of such coverings is the same as in the topological setting. Let $\mathcal{H}^{d,b}$ be the functor which sends S

to the set $\mathcal{H}^{d,b}(S)$ of equivalence classes of d-sheeted branched covering of $\mathbb{P}^1_S$ with b branch points. There is a morphism $\delta\colon \mathcal{H}^{d,b} \to \mathbb{P}^b - \Delta_b$ which sends a covering to its discriminant. The following proposition justifies our use of the same notation as in the analytic setting.

Proposition 4.2 ([Ful69, §6, §7]). If $d > 2$ then the functor $\mathcal{H}^{d,b}$ is represented by a scheme, and the morphism $\delta\colon \mathcal{H}^{d,b} \to \mathbb{P}^b - \Delta_b$ is a finite étale cover. The analytic space associated to $\mathcal{H}^{d,b}$ is canonically isomorphic to the analytic space of Section 2 with the same name.

As a corollary of Theorem 3.1, we see that the scheme $\mathcal{H}^{d,b}$ is irreducible. The next proposition is a classical result of Severi [Sev]. We will give a modern proof due to Fulton [Ful69, 8.1], adding details in certain places.

Proposition 4.3. Let C be a smooth projective curve of genus g and let $d \geq g+1$ be an integer. Then there exists a simple d-sheeted branched covering $C \to \mathbb{P}^1$.

Proof. Let $C^{(d)}$ be the dth symmetric power of C. There are morphisms $\alpha_1\colon C^{d-1} \to C^{(d)}$, $\alpha_2\colon C^{d-2} \to C^{(d)}$, and $\alpha_3\colon C^{d-2} \to C^{(d)}$ defined by the formulas

$$\alpha_1(P_1, \ldots, P_{d-1}) = 2P_1 + P_2 + P_3 + \cdots + P_{d-1},$$
$$\alpha_2(P_1, \ldots, P_{d-2}) = 2P_1 + 2P_2 + P_3 + \cdots + P_{d-2},$$
$$\alpha_3(P_1, \ldots, P_{d-2}) = 3P_1 + P_2 + P_3 + \cdots + P_{d-2}.$$

Define

$$D := \operatorname{im}(\alpha_1) \subset C^{(d)}, \qquad T := \operatorname{im}(\alpha_2) \cup \operatorname{im}(\alpha_3) \subset C^{(d)}.$$

Let Pic^d be the Picard scheme which parameterizes degree d line bundles on C. This is a g-dimensional smooth proper scheme over $\mathbb{C}$. There is a surjective morphism $\varphi\colon C^{(d)} \to \operatorname{Pic}^d$ which sends $\sum P_i$ to $\mathcal{O}(\sum P_i)$. For a line bundle $L \in \operatorname{Pic}^d(\mathbb{C})$, let

$$S_L := \varphi^{-1}(L), \qquad D_L := S_L \cap D, \qquad T_L := S_L \cap T.$$

Then S_L consists of effective divisors whose associated line bundles are isomorphic to L. Recall that there is a canonical identification between S_L and the projective space associated to $H^0(C, L)$. Under this bijection a section $s \in H^0(C, L)$ is sent to its divisor of zeros $(s)_0 \in S_L$.

Let U be as in Lemma 4.4 below and fix $L \in U$. For a vector space V and an integer n let $\operatorname{Gr}(n, V)$ be the Grassmannian parameterizing n-dimensional subspaces of V. This is a variety of dimension $n(\dim V - n)$. Note that $\operatorname{Gr}(2, H^0(C, L))$ parameterizes lines in S_L. For every point $P \in C$, let $|L|_P \subset S_L$ be the hyperplane corresponding to the inclusion $H^0(C, L(-P)) \subset H^0(C, L)$.

We claim there is a non-empty open subset $V \subset \operatorname{Gr}(2, H^0(C, L))$ such that every line $\ell \in V$ satisfies the following:

(1) ℓ is not contained in any $|L|_P$.
(2) ℓ meets D_L in finitely many points.
(3) ℓ does not meet T_L.

To prove the claim, let $\operatorname{Gr}(d - g, H^0(C, L))$ be the variety of hyperplanes in S_L, and let

$$I := \{(\ell, H) \ : \ \ell \subset H\} \subset \operatorname{Gr}(2, H^0(C, L)) \times \operatorname{Gr}(d - g, H^0(C, L)).$$

Then there are projections $\pi_\ell \colon I \to \mathrm{Gr}(2, H^0(C, L))$ and $\pi_H \colon I \to \mathrm{Gr}(d - g, H^0(C, L))$. There is also an algebraic map $C \to \mathrm{Gr}(d - g, H^0(C, L))$ sending P to $|L|_P$; let Z be the image of this map. The fibers of π_H are isomorphic to $\mathrm{Gr}(2, \mathbb{C}^{d-g})$, and so they have dimension $2(d - g - 2)$. Thus $\dim \pi_H^{-1}(Z) \leq 2(d - g - 2) + 1$. Observe that

$$\dim \pi_\ell(\pi_H^{-1}(Z)) < \dim \mathrm{Gr}(2, H^0(C, L)) = 2(d - g - 1).$$

Hence to satisfy (1) it suffices to exclude the closure of $\pi_\ell(\pi_H^{-1}(Z))$.

For (2), we note that as D_L is a hypersurface then any line ℓ in S_L meets D_L. Thus, as long as ℓ is not contained in D_L, then $\ell \cap D_L$ is finite. There is a proper closed subvariety of $\mathrm{Gr}(2, H^0(C, L))$ which parameterizes lines contained in D_L. To satisfy (2) we may simply exclude this subvariety of $\mathrm{Gr}(2, H^0(C, L))$.

For (3), let

$$J := \{(\ell, x) \ : \ x \in \ell\} \subset \mathrm{Gr}(2, H^0(C, L)) \times S_L.$$

There are projections $\pi_\ell \colon J \to \mathrm{Gr}(2, H^0(C, L))$ and $\pi_{S_L} \colon J \to S_L$. Then $\pi_\ell(\pi_{S_L}^{-1}(T_L)) \subset \mathrm{Gr}(2, H^0(C, L))$ is the collection of lines meeting T_L. The fibers of π_{S_L} have dimension $\dim S_L - 1 = d - g - 1$. This is because a line through a given point is determined by a second point on the line, and two points determine the same line if they are collinear with the original point. As $\dim T_L \leq d - g - 2$ then

$$\dim \pi_\ell(\pi_{S_L}^{-1}(T_L)) \leq 2d - 2g - 3 < \dim \mathrm{Gr}(2, H^0(C, L)).$$

Thus for (3) it suffices to exclude the closure of $\pi_\ell(\pi_{S_L}^{-1}(T_L))$. This proves the claim.

Now fix a line ℓ satisfying (1), (2), and (3), and let f_0, $f_1 \in H^0(C, L)$ be sections such that ℓ is spanned by $\lambda_1 f_0 - \lambda_0 f_1$ for λ_0, $\lambda_1 \in \mathbb{C}$ not both zero. Let $V \subset H^0(C, L)$ be the subspace spanned by f_0 and f_1. The sections f_0, f_1 define a rational map $f \colon C \to \mathbb{P}(V^*)$. We now recall the construction of f. Let $P \in C$ and let L_P be the restriction of L at P. Then L_P is a one-dimensional vector space. Define the linear functional $\mathrm{ev}_P \colon V \to L_P$ by sending a section to its restriction at P. The map f sends P to the kernel of ev_P, which is a hyperplane in V and hence corresponds to a line in V^* (since $\dim V = 2$ we did not need to take the dual). The map f is defined on the locus of points where at least one of f_0, f_1 is nonzero.

We claim that f is a simple d-sheeted branched covering. First, note that $s \in H^0(C, L)$ vanishes at $P \in C$ if and only if $s \in |L|_P$. Thus by (1) we see that f is defined everywhere on C. By construction, $f^{-1}([\lambda_0 : \lambda_1]) = (\lambda_1 f_0 - \lambda_0 f_1)_0$. Hence by (2) this fiber consists of exactly d points over all but finitely many $[\lambda_0 : \lambda_1]$. By (3) this fiber always contains at least $d - 1$ points, so f is a simple d-sheeted branched covering.

$\square$

Lemma 4.4. There is a non-empty open subset $U \subset \mathrm{Pic}^d$ such that for every line bundle $L \in U$,

(1) $\dim S_L = d - g$.
(2) $\dim D_L \leq d - g - 1$.
(3) $\dim T_L \leq d - g - 2$.
(4) $H^0(C, L(-P)) \subset \dim H^0(C, L)$ has codimension one for all $P \in C$.

Proof. The existence of U satisfying (1) follows from Lemma 4.5. For (2), apply Lemma 4.5 to the composite morphism $\varphi \circ \alpha_1$ to get another open subset $V \subset \mathrm{Pic}^d$, and take the intersection $U \cap V$. Here we are using that $d \geq g+1$ so that $\varphi \circ \alpha_1$ is dominant. If $d = g+1$ then neither $\varphi \circ \alpha_2$ nor $\varphi \circ \alpha_3$ is dominant, so we may simply remove the closure of their images. Otherwise, if $d > g+1$ then we proceed as in (2). At this point we have an open subset $U \subset \mathrm{Pic}^d$ satisfying (1), (2), and (3).

Now consider the map $\psi \colon \mathrm{Pic}^d \times C \to \mathrm{Pic}^{2g-d-1}$ defined by $\psi(L,P) = \Omega \otimes L^{-1}(P)$, where Ω is the sheaf of differentials on C. Let Z be the image of the natural map $C^{(2g-d-1)} \to \mathrm{Pic}^{2g-d-1}$ (we set $Z = \emptyset$ if $2g - d - 1 < 0$). As $d \geq g+1$ then $\dim Z \leq g-2$. As the fibers of ψ are curves, then $\dim \psi^{-1}(Z) \leq g-1$. Let $\pi \colon \mathrm{Pic}^d \times C \to \mathrm{Pic}^d$ be the projection, and choose $U \subset \mathrm{Pic}^d$ satisfying (1), (2), (3) and such that $U \cap \pi(\psi^{-1}(Z)) = \emptyset$. By this last condition, $H^0(C, \Omega \otimes L^{-1}(P)) = 0$ for every $L \in U$ and $P \in C$. By Serre duality, $H^1(C, L(-P)) = 0$.

Recall that if D is a divisor on C then $\mathcal{O}(D)$ is the sheaf whose sections over an open set $U \subset C$ are those functions $U \to \mathbb{P}^1$ with $\mathrm{Div}(f) + D \geq 0$. Hence if $P \in C$ there is a natural map $\mathcal{O}(-P) \to \mathcal{O}$. By tensoring with L we get a map $L(-P) \to L$. This map is injective with cokernel a one-dimensional skyscraper sheaf δ_P at P. By using that $H^1(C, L(-P)) = 0$ and the long exact sequence associated to

$$0 \to L(-P) \to L \to \delta_P \to 0,$$

we see that

$$\dim \mathrm{cok}(H^0(C, L(-P)) \to H^0(C, L)) = 1.$$

$\square$

Lemma 4.5. Let $f \colon X \to Y$ be a dominant map between integral $\mathbb{C}$-schemes of finite type. Then there is a nonempty open subset $U \subset Y$ such that $\dim f^{-1}(y) = \dim X - \dim Y$ for every $y \in Y(\mathbb{C})$.

Proof. One way to prove this is to use the upper semi-continuity theorem ([Gro66, 13.1.3]). Another way is to use that f is generically flat ([Gro66, 8.9.4]) and that fibers of flat morphisms have the expected dimension ([Gro65, 6.1.4]).

$\square$

Corollary 4.6. The coarse moduli space M_g of curves of genus g over $\mathbb{C}$ is irreducible.

Proof. By the Riemann-Hurwitz formula, a simple d-sheeted branched covering $f \colon C \to \mathbb{P}^1$ from a curve C of genus g has $b = 2d + 2g - 2$ branch points. There is a morphism $\mathcal{H}^{d,b} \to M_g$ which forgets the morphism to $\mathbb{P}^1$. As $\mathcal{H}^{d,b}$ is smooth, then since it is connected (Theorem 3.1) it is irreducible. As long as $d \geq g+1$ then Proposition 4.3 guarantees the map $\mathcal{H}^{d,b} \to M_g$ is surjective. Hence M_g is irreducible as it is the image of an irreducible variety.

$\square$

5. ADMISSIBLE COVERS

We conclude by indicating how to compactify the Hurwitz scheme $\mathcal{H}^{d,b}$. By a compactification we mean a projective scheme $\overline{\mathcal{H}}^{d,b}$ that contains $\mathcal{H}^{d,b}$ as a dense open subset.

Moreover, we would like $\overline{\mathcal{H}}^{d,b}$ to a have a modular description, i.e. its points should parameterize some type of coverings. We will primarily follow Harris and Morrison [HM98] in this section.

Let $D^\circ := \mathrm{Spec}(\mathbb{C}[t, t^{-1}])$ be the punctured disk, and let $C \to D^\circ$ be a family of simple branched covers of $\mathbb{P}^1$ parameterized by D°. As a first attempt, one may try to understand what happens to the family as $t \to 0$. In the case where two of the branch points collide, there are three possibilities depending on the monodromy of the two points:

(1) If the transpositions corresponding to the two colliding branch points are equal then the curves C_t degenerate to a limit curve C_0 with a simple node.
(2) If the transpositions are non-commuting then the limit curve C_0 is smooth and has a triple ramification point.
(3) If the transpositions are disjoint then the limit curve C_0 is smooth and has two points of double ramification lying over a single point of $\mathbb{P}^1$.

If we allow more branch points to collide, then the possibilities quickly become more complicated. A more elegant way to attack the problem is to not allow the branch points to collide by instead considering covers of stable curves of genus zero.

A *stable curve* is a proper connected curve C that has only nodes as singularities and has only finitely many automorphisms. Since we will be interested in branched covers, we also need to keep track of an ordered set of points on a stable curve. A *stable n-pointed curve* is a proper connected curve C that has only nodes as singularities together with n distinct smooth points $p_1, \ldots, p_n \in C$ such that the tuple $(C, p_1, \ldots, p_n)$ has only finitely many automorphisms.

As with smooth curves, the genus $g = \dim H^1(C, \mathcal{O}_C)$ is an important invariant of a stable curve C. If C has δ nodes and ν irreducible components of geometric genera $g_1, \ldots, g_\nu$ then

$$g = \sum_{i=1}^{\nu} g_i + \delta - \nu + 1.$$

This can be proved by using the following facts. If $\tilde{C}$ is the normalization of C, then the normalization map $\varphi \colon \tilde{C} \to C$ is finite. In particular, $\varphi_* \mathcal{O}_{\tilde{C}}$ has the same cohomology as $\mathcal{O}_{\tilde{C}}$. Moreover, if the nodes of C are $q_1, \ldots, q_\delta$ we have an exact sequence

$$0 \to \mathcal{O}_C \to \varphi_* \mathcal{O}_{\tilde{C}} \to \oplus_{i=1}^{\delta} \mathbb{C}_{q_i} \to 0.$$

The claim can be proved by considering the associated long exact sequence.

There exist coarse moduli spaces $\overline{M}_g$ and $\overline{M}_{g,n}$ of stable curves and n-pointed stable curves. These are projective varieties. In particular, given a family $C \to X - P$ of stable curves, where X is a smooth curve and $P \in X$, there exists a unique way to extend this to a family over X.

Example 5.1. Let X be a smooth curve of genus $g > 0$ and pick a point $P \in X$. Consider the family $C \to X - P$ where the curve C_Q for $Q \in X - P$ is the nodal curve of genus $g + 1$ obtained by gluing P and Q together. Then the limit curve C_P appears to be a cuspidal curve, which is not stable. To determine the stable limit, consider the projection $p_1 \colon (X - P) \times X \to X - P$. The original family is obtained from this family by identifying the fibers over $X - P$ of the images L_1 of the diagonal map $\Delta \colon X - P \to (X - P) \times X$

and the horizontal slice $L_2 := (X - P) \times \{P\}$. Now let Y be the blow-up of $X \times X$ at the point (P, P). The fiber of the map $Y \to X \times X \xrightarrow{p_1} X$ over P is isomorphic to C glued at the point P to $\mathbb{P}^1$. Note that L_1 and L_2 extend uniquely to complete curves $\overline{L}_1, \overline{L}_2$ in Y meeting the fiber over P in different points of $\mathbb{P}^1 - P$. We can complete the original family by identifying $\overline{L}_1$ and $\overline{L}_2$. The result is that the stable limit is a copy of X glued at the point P to an irreducible nodal curve of geometric genus 0 having a unique node distinct from P.

Example 5.2. Let $(B, w_1(t), \ldots, w_b(t))$ be a family of stable b-pointed curves over D° where the only thing that changes is the locations of the marked points. Suppose that in the limit $t \to 0$, the first k points $w_1(t), \ldots w_k(t)$ remain disjoint and the other points $w_{k+1}(t), \ldots, w_b(t)$ all converge at comparable speeds to a single point P. To determine the stable limit, we blow up the point in the family $B \times D$ where the marked points meet. The limit curve then consists of the curve $(B, w_1(0), \ldots, w_k(0))$ glued at the point P to $(\mathbb{P}^1, w_{k+1}(0), \ldots, w_b(0))$ where all of the marked points are disjoint from P.

We are now ready to extend the notion of a branched covering of $\mathbb{P}^1$ to coverings of stable curves of genus zero. For a variety X, let X_{ns} be the locus of non-singular points and let X_{sing} be the locus of singular points.

Definition 5.3. Let $(B, w_1, \ldots, w_b)$ be a b-pointed stable curve of genus zero, and let $q_1, \ldots, q_k$ be the nodes of B. An admissible cover of B is a nodal curve C (possibly disconnected) with a morphism $f \colon C \to B$ such that

(1) $\pi^{-1}(B_{\mathrm{ns}}) = C_{\mathrm{ns}}$ and the morphism $C_{\mathrm{ns}} \to B_{\mathrm{ns}}$ is a simple branched cover with branch locus $\{w_1, \ldots, w_b\}$.
(2) $\pi^{-1}(B_{\mathrm{sing}}) = C_{\mathrm{sing}}$ and for every node q of B with a node r lying over it, there exists an integer $m > 0$ (depending on q, r) such that there exist local coordinates x, y on B and u, v on C satisfying the following: the nodes q, r are defined by $xy = 0$, $uv = 0$, and

$$\pi^* x = u^m, \qquad \pi^* y = v^m.$$

Let $b = 2d + 2g - 2$. Then there exists a coarse moduli space $\overline{\mathcal{H}}^{d,b}$ of admissible covers of genus g. Moreover, all of the moduli spaces we have described so far fit into a natural diagram

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

Example 5.4. We return to Example 5.2 in the case where $B = \mathbb{P}^1$. Suppose we have a family C_t of simple d-sheeted branched coverings of $\mathbb{P}^1$ with branch locus $w_1(t), \ldots, w_b(t)$. The monodromy of each $w_i(t)$ is an element $t_i \in S_d$ that is independent of $t \neq 0$. Topologically, the stable limit of $(B, w_1(t), \ldots, w_b(t))$ can be viewed as the result of contracting a great circle γ along $\mathbb{P}^1$ to a point, where $w_1(t), \ldots w_k(t)$ lie on one hemisphere and $w_{k+1}(t), \ldots, w_b(t)$ lie on the other hemisphere. Let us choose our basepoint x to lie on this great circle.

In the notation of Section 3, recall that we have loops γ_i around each of the w_i. Then, with the appropriate choice of direction, $\gamma \in \pi_1(B, x)$ is

$$\gamma = \gamma_1 \gamma_2 \cdots \gamma_k.$$

Write γ as a product of disjoint cycles in S_d of length $a_1, \ldots, a_\delta$. Then the nodal curve C_0 has δ nodes lying over P with ramification degrees $a_1, \ldots, a_\delta$. Here the ramificiation degree of a node is the integer m appearing in Definition 5.3 (2). Over $B - P$ the curve C_0 is possibly disconnected, and it is a simple branched cover with components corresponding to the orbits of $t_1, \ldots, t_k$ on $\{1, \ldots, d\}$. Similarly over the other copy of $\mathbb{P}^1 - P$ the curve C_0 is a simple branched cover with components corresponding to the orbits of $t_{k+1}, \ldots, t_b$ on $\{1, \ldots, d\}$. In general the curve C_0 may not be stable. We refer the reader to Section 3G of [HM98] for several example computations, and also for what happens when the branch points come together in more complicated ways.

References

[Cle73] A. Clebsch. Zur Theorie der Riemann'schen Fläche. *Math. Ann.*, 6(2):216–230, 1873.

[Ful69] William Fulton. Hurwitz schemes and irreducibility of moduli of algebraic curves. *Ann. of Math. (2)*, 90:542–575, 1969.

[Gro65] A. Grothendieck. Éléments de géométrie algébrique. IV. Étude locale des schémas et des morphismes de schémas. II. *Inst. Hautes Études Sci. Publ. Math.*, (24):231, 1965.

[Gro66] A. Grothendieck. Éléments de géométrie algébrique. IV. Étude locale des schémas et des morphismes de schémas. III. *Inst. Hautes Études Sci. Publ. Math.*, (28):255, 1966.

[HM98] Joe Harris and Ian Morrison. *Moduli of curves*, volume 187 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1998.

[Hur91] A. Hurwitz. Ueber Riemann'sche Flächen mit gegebenen Verzweigungspunkten. *Math. Ann.*, 39(1):1–60, 1891.

[Moc95] Shinichi Mochizuki. The geometry of the compactification of the Hurwitz scheme. *Publ. Res. Inst. Math. Sci.*, 31(3):355–441, 1995.

[Sev] F. Severi. Vorlesungen über algebraische Geometrie (tr. by E. Loffler). Teubuer, Leibzig, 1921.