

BRANCHED COVERS OF THE RIEMANN SPHERE

BRIAN OSSERMAN

We introduce Riemann surfaces and branched covers of the Riemann sphere, and describe the relationship to topology and group theory.

1. COMPLEX MANIFOLDS

A (real) manifold of dimension n is a (Hausdorff, second countable) space which is locally homeomorphic to an open subset of \mathbb{R}^n . If we wish to make a definition of a complex manifold, we could replace \mathbb{R}^n by \mathbb{C}^n , but then we see that we have simply given the definition of an even-dimensional real manifold. As with real differentiable manifolds, the key is to add some additional structure via a well-behaved atlas.

Definition 1.1. Given a topological space X , a **chart** for X is a homeomorphism $\varphi : V \rightarrow U$, where V is an open subset of \mathbb{C}^n for some n , and U is an open subset of X . A **holomorphic atlas** for X is a collection of charts $\varphi_i : V_i \rightarrow U_i$, with the $V_i \subseteq \mathbb{C}^n$ for some fixed n , such that:

- (1) the U_i cover X ;
- (2) for each $i \neq j$, the **transition map**

$$\varphi_{i,j} : \varphi_i^{-1}(U_i \cap U_j) \xrightarrow{\varphi_i} U_i \cap U_j \xrightarrow{\varphi_j^{-1}} \varphi_j^{-1}(U_i \cap U_j)$$

is holomorphic.

A **complex manifold** of dimension n is a (Hausdorff, second countable) topological space X together with a holomorphic atlas.

Proposition 1.2. *A complex manifold X is naturally oriented when considered as a topological manifold.*

Sketch of proof. The complex numbers have a natural orientation, which induces one on \mathbb{C}^n . Because holomorphic maps preserve this orientation, we get an induced orientation on any complex manifold. \square

Using the atlas and the fact that we have a definition of holomorphic functions on open subsets of \mathbb{C}^n , we can define holomorphic functions on complex manifolds.

Definition 1.3. Let X be a complex manifold with atlas $\{\varphi_i : V_i \rightarrow U_i\}$, and $U \subseteq X$ open. A function $f : U \rightarrow \mathbb{C}$ is **holomorphic** if for each i , the composed function

$$f \circ \varphi_i : \varphi_i^{-1}(U \cap U_i) \rightarrow \mathbb{C}$$

is holomorphic.

This definition makes sense whether or not we require the transition maps $\varphi_{i,j}$ to be holomorphic. However, the additional hypothesis implies that we don't need to look at more than one U_i if $U \subseteq U_i$.

Exercise 1.4. If $U \subseteq U_i$ for some i , then $f : U \rightarrow \mathbb{C}$ is holomorphic if and only if $f \circ \varphi_i : \varphi_i^{-1}(U) \rightarrow \mathbb{C}$ is holomorphic.

Using the notion of holomorphic functions, we can then define a holomorphic map between two complex manifolds.

Definition 1.5. Given two holomorphic manifolds X, X' , a continuous map $\varphi : X \rightarrow X'$ is **holomorphic** if for all $U \subseteq X'$ open, and all $f : U \rightarrow \mathbb{C}$ holomorphic, the composed function

$$f \circ \varphi : \varphi^{-1}(U) \rightarrow \mathbb{C}$$

is also holomorphic.

This is equivalent to a definition which uses the atlas explicitly:

Exercise 1.6. Given holomorphic manifolds X, X' with atlas $\{\varphi_i : V_i \rightarrow U_i\}$ and $\{\varphi'_i : V'_i \rightarrow U'_i\}$, a map $\varphi : X \rightarrow X'$ is holomorphic if and only if for all i, j the composed map

$$(\varphi'_j)^{-1} \circ \varphi \circ \varphi_i : \varphi_i^{-1}(\varphi^{-1}(U'_j) \cap U_i) \rightarrow V'_j$$

is holomorphic.

We also have:

Exercise 1.7. Given a complex manifold X , and $U \subseteq X$ open, a function $f : U \rightarrow \mathbb{C}$ is holomorphic if and only if it gives a holomorphic mapping when \mathbb{C} is considered as a complex manifold via the trivial atlas.

Definition 1.8. A **Riemann surface** is a complex manifold of dimension 1.

Example 1.9. The Riemann sphere \mathbb{CP}^1 is a compact Riemann surface with underlying topological space S^2 . If we realize as S^2 as the unit sphere in \mathbb{R}^3 , let U_1 be the complement of the “north pole” $(0, 0, 1)$ and U_2 the complement of the “south pole” $(0, 0, -1)$. Both V_1 and V_2 are the complex plane. We then define the φ_i as follows:

$$\varphi_1(x + iy) = \frac{1}{x^2 + y^2 + 1}(2x, 2y, x^2 + y^2 - 1)$$

and

$$\varphi_2(x + iy) = \frac{1}{x^2 + y^2 + 1}(2x, -2y, 1 - x^2 - y^2).$$

Then we compute

$$(\varphi_2^{-1} \circ \varphi_1)(x + iy) = \frac{x - iy}{x^2 + y^2} = \frac{1}{x + iy},$$

which is holomorphic, and similarly for $\varphi_1^{-1} \circ \varphi_2$. (Note that φ_1^{-1} is stereographic projection from the north pole, and φ_2^{-1} is stereographic projection from the south pole composed with complex conjugation)

Example 1.10. We often picture complex manifolds not in terms of a topological space with an atlas, but in terms of gluing together the pieces V_i along biholomorphisms of open subsets. For instance, we picture \mathbb{CP}^1 as obtained from two copies of \mathbb{C} by gluing along the open subsets $\mathbb{C} \setminus \{0\}$ via the map $1 \mapsto 1/z$. However, when working this way, one has to be careful about the Hausdorff condition.

For instance, if we instead glue two copies of \mathbb{C} along the open subsets $\mathbb{C} \setminus \{0\}$ via the identity map, we obtain the “plane with a doubled origin”, which is locally biholomorphic to open subsets of \mathbb{C} , but is not Hausdorff, as the two origins do not have disjoint open neighborhoods.

2. BRANCHED COVERS

Although branched covers are interesting more generally, we will restrict our attention to branched covers of the Riemann sphere, which already have a rich and elegant structure.

Definition 2.1. A **branched cover** (of the Riemann sphere) is a pair (C, f) where C is a compact, connected Riemann surface, and $f : C \rightarrow \mathbb{CP}^1$ is a non-constant holomorphic map.

Definition 2.2. We say two branched covers (C, f) and (C', f') are **equivalent** if there is a biholomorphism $g : C \rightarrow C'$ such that $f = f' \circ g$.

The following is a basic fact from complex analysis:

Proposition 2.3. *Suppose (C, f) is a branched cover, and $P \in C$. Then there are open neighborhoods U, U' of P and $f(P)$ respectively, and open neighborhoods V, V' of 0 in \mathbb{C} , and biholomorphisms $g : U \rightarrow V, g' : U' \rightarrow V'$ sending P and $f(P)$ respectively to 0 , such that the map $g' \circ f \circ g^{-1} : V \rightarrow V'$ is equal to $z \mapsto z^{e_P}$ for some $e_P \geq 1$.*

Definition 2.4. If the e_P from Proposition 2.3 is strictly greater than 1, we say that P is a **ramification point** of f , with **ramification index** e_P . In this case, $f(P)$ is a **branch point** of f . A neighborhood U of P such that there exist U', V, V', g, g' as in Proposition 2.3 is a **standard neighborhood** of P .

It is clear that any neighborhood contained in a standard neighborhood is again a standard neighborhood.

Corollary 2.5. *Let (C, f) be a branched cover. Then there are only finitely many ramification and branch points of f .*

Proof. Given a neighborhood U of a point $P \in C$ as in Proposition 2.3, we see that on U , the map f is locally one-to-one everywhere except possibly at P , where it is locally e_P -to-one. By compactness of C , there is a finite open cover of C by such neighborhoods, so we conclude that f is locally one-to-one at all but finitely many points, and hence can be ramified at only finitely many points. The statement on branch points follows. \square

A stronger version of Proposition 2.3, which makes use of the compactness of the cover, is:

Lemma 2.6. *Given a branched cover (C, f) , and $Q \in \mathbb{C}\mathbb{P}^1$, the fiber $f^{-1}(Q)$ is finite, and there exists a connected neighborhood U' of Q such that every connected component of $f^{-1}(U')$ contains a unique point $P \in f^{-1}(Q)$, surjects onto U' , and is a standard neighborhood of P .*

Proof. C is covered by finitely many standard neighborhoods by compactness, and on every standard neighborhood, f is finite-to-one, so we conclude the finiteness of $f^{-1}(Q)$.

We next observe that f is a closed map, because C is compact and $\mathbb{C}\mathbb{P}^1$ is Hausdorff, and also an open map, by the open mapping theorem in complex analysis (or more concretely, because the map $z \mapsto z^e$ is visibly open for any e). Then for any open subset $U' \subseteq \mathbb{C}\mathbb{P}^1$, the induced map $f^{-1}(U') \rightarrow U'$ is also open and closed. It follows that if $W \subseteq f^{-1}(U')$ is a connected component, its image is open and closed in U' , and hence is all of U' if U' is connected.

Let P_1, \dots, P_m be the preimages of Q , and U_1, \dots, U_m be standard neighborhoods of P_1, \dots, P_m . We may further suppose that they are chosen to be disjoint from one another. Each boundary ∂U_i is then a closed set with $f(\partial U_i)$ closed and not containing Q . We can then choose U' a connected neighborhood of Q disjoint from all the $f(\partial U_i)$. Let U be a connected component of $f^{-1}(U')$; by the above, U surjects onto U' , so contains (at least) one of the P_i . To see that U is a standard neighborhood and does not contain P_j for $j \neq i$, it is enough to prove $U \subseteq U_i$. We have $U_i \cap f^{-1}(U')$ and $\overline{U}_i \cap f^{-1}(U')$ open and closed respectively in $f^{-1}(U')$ by definition. But $\partial U_i \cap f^{-1}(U') = \emptyset$ by construction, so $U_i \cap f^{-1}(U') = \overline{U}_i \cap f^{-1}(U')$ is also closed. Since $U \cap U_i \neq \emptyset$ and U is connected, we conclude $U \subseteq U_i$, as desired. \square

We now conclude:

Corollary 2.7. *Let (C, f) be a branched cover. Then f induces a (finite-degree) topological covering map after removing the branch points and their preimages.*

Proof. After removing the ramification points, it is clear that f is a local homeomorphism, and more specifically, a homeomorphism onto its image for every standard neighborhood. The corollary thus follows immediately from Lemma 2.6. \square

Definition 2.8. The **degree** of a branched cover is the degree of the induced covering map after removing the branch points and their preimages.

We also conclude the following result, which says that we can think of a branched cover as having degree d even over the branch points, if we count with appropriate multiplicities.

Corollary 2.9. *If (C, f) is a branched cover of degree d , then for any $Q \in \mathbb{CP}^1$, we have*

$$\sum_{P \in f^{-1}(Q)} e_P = d.$$

Proof. Take $Q' \neq Q$ a point which lies in the neighborhood U' provided by Lemma 2.6 (this is then necessarily not a branch point). Given $P \in f^{-1}(Q)$, if U is the connected component of $f^{-1}(U')$ containing P , then Proposition 2.3 implies that Q' has e_P preimages in U . We conclude that the number of points in $f^{-1}(Q')$ is equal to $\sum_{P \in f^{-1}(Q)} e_P$. On the other hand, by Proposition 2.7, the number of points in $f^{-1}(Q')$ is equal to d , giving the desired identity. \square

This motivates the following definition:

Definition 2.10. Given a branched cover (C, f) of degree d , and Q a branch point, the **branch type** of f at Q is the partition of d given by $\{e_P : P \in f^{-1}(Q)\}$.

Thus, a branched cover comes with the combinatorial data of its degree d and a tuple of partitions of d determined by its branching. We often assume that the branch points are marked in an ordered fashion Q_1, \dots, Q_r ; in this case, we can place the combinatorial data into a tuple $(d; T_1, \dots, T_r)$ where each T_i is the branch type at Q_i ; this is called the **type** of the cover (C, f) .

We now return to the discussion of the topological covering map induced by (C, f) , according to Corollary 2.7. Note that removing a finite set of points from a connected surface leaves it connected, so the covering space coming from Corollary 2.7 is also connected. Remarkably, we have the following converse result, known as the Riemann existence theorem:

Theorem 2.11 (Riemann existence). *Given distinct points $Q_1, \dots, Q_r \in \mathbb{CP}^1$, write $\widehat{\mathbb{CP}}^1 := \mathbb{CP}^1 \setminus \{Q_1, \dots, Q_r\}$, and suppose \hat{C} is a connected topological space and $\hat{f} : \hat{C} \rightarrow \widehat{\mathbb{CP}}^1$ is a topological covering map of finite degree. Then there exists a unique branched cover (C, f) (up to equivalence) such that \hat{C} and \hat{f} are obtained by removing the Q_i and their preimages.*

This gives a basic equivalent definition of branched covers in terms of finite topological covering maps of the punctured sphere. The proof is not too difficult, following from the study of connected topological coverings of the punctured disk, but we omit it.

Remark 2.12. Riemann proved rather more than this, but this is what will be necessary for our immediate purposes.

3. THE MONODROMY OF A BRANCHED COVER

Because we have related branched covers to topological covering spaces, standard results from topology allow us to use the fundamental groups to describe branched covers in terms of group theory.

If we fix distinct points $Q_1, \dots, Q_r \in \mathbb{CP}^1$, write $\widehat{\mathbb{CP}}^1 := \mathbb{CP}^1 \setminus \{Q_1, \dots, Q_r\}$, and choose also a point $Q \in \widehat{\mathbb{CP}}^1$. Suppose we have a branched cover (C, f) branched only at the Q_i , with branch

type T_i at each Q_i . Let $\hat{f} : \hat{C} \rightarrow \widehat{\mathbb{C}\mathbb{P}^1}$ be the topological covering space of degree d we obtain from Corollary 2.7. Then if γ is a loop in $\widehat{\mathbb{C}\mathbb{P}^1}$ based at Q , and $P \in f^{-1}(Q)$, it is a standard result of topology that there is a unique lift $\tilde{\gamma}_P$ of γ to a path in \hat{C} which starts at P . Then $\tilde{\gamma}_P$ necessarily ends in $f^{-1}(Q)$, but not necessarily at P . Denote the endpoint by $\mu(\gamma)(P)$. Thus, we have a function

$$\mu(\gamma) : f^{-1}(P) \rightarrow f^{-1}(P).$$

We see immediately that this is invertible, since if γ^{-1} is the same loop as γ with its direction reversed, we clearly have $\mu(\gamma^{-1}) = \mu(\gamma)^{-1}$. Thus, $\mu(\gamma)$ is a permutation of $f^{-1}(Q)$. This is invariant under homotopy, and also compatible with composition. If we choose a labeling of $f^{-1}(Q)$, we get an isomorphism $\text{Sym}(f^{-1}(Q)) \xrightarrow{\sim} S_d$, so we conclude:

Proposition 3.1. *Given a branched cover (C, f) with branch points Q_i , and a labeling of $f^{-1}(Q)$, then lifting of loops based at Q induces a homomorphism*

$$\mu : \pi_1(\widehat{\mathbb{C}\mathbb{P}^1}, Q) \rightarrow S_d.$$

Definition 3.2. The homomorphism μ is the **monodromy map**, and $\mu(\gamma)$ is the **monodromy** of (C, f) around γ .

The next step involves a good understanding of the fundamental group of $\widehat{\mathbb{C}\mathbb{P}^1}$. A standard calculation from topology gives the following:

Proposition 3.3. *There exist loops $\gamma_1, \dots, \gamma_r \subseteq \widehat{\mathbb{C}\mathbb{P}^1}$ based at Q , satisfying:*

- (i) *The γ_i generate $\pi_1(\widehat{\mathbb{C}\mathbb{P}^1}, Q)$.*
- (ii) *The only relation among the γ_i in $\pi_1(\widehat{\mathbb{C}\mathbb{P}^1}, Q)$ is that $\gamma_1 \cdots \gamma_r = 1$.*
- (iii) *Each γ_i is homotopic to a small loop around Q_i .*

Note that (ii) implies that $\pi_1(\widehat{\mathbb{C}\mathbb{P}^1}, Q)$ is the free group generated by any $r - 1$ of the γ_i .

Remark 3.4. Choice of such γ_i is not unique. Indeed, this non-uniqueness will be very important to us when we discuss braid actions and connected components of Hurwitz spaces.

We now suppose we have fixed Q_1, \dots, Q_r and Q , but not necessarily any branched cover.

Proposition 3.5. *Suppose (C, f) is a branched cover with branch points Q_1, \dots, Q_r , and branch type T_i at Q_i for $i = 1, \dots, r$, and with a labeling of $f^{-1}(Q)$. Set $\sigma_i = \mu(\gamma_i) \in S_d$ for $i = 1, \dots, r$. Then we have:*

- (i) *The subgroup of S_d generated by the σ_i is transitive.*
- (ii) $\sigma_1 \cdots \sigma_r = 1$.
- (iii) *The cycle decomposition of σ_i agrees with T_i , for $i = 1, \dots, r$.*

Recall that a subgroup $G \subseteq S_d$ is transitive if for all $i, j \in \{1, \dots, d\}$, there is some $\sigma \in G$ with $\sigma(i) = j$. For $\sigma \in S_d$, the decomposition of σ into disjoint cycles yields a partition of d by considering the lengths of the cycles (if we consider all fixed points of σ to lie in cycles of length 1), and thus it makes sense to compare to the T_i .

Proof. (i) By Proposition 3.5 (i), this is equivalent to showing that $\mu(\pi_1(\widehat{\mathbb{C}\mathbb{P}^1}, Q))$ is transitive. This follows from the connectedness of C : first observe that $\hat{C} := C \setminus (\bigcup_i f^{-1}(Q_i))$ is still connected, and indeed path connected, so given $P, P' \in f^{-1}(Q)$, there is a path $\tilde{\gamma}$ in \hat{C} connecting them. Then $f(\tilde{\gamma})$ is a loop in $\widehat{\mathbb{C}\mathbb{P}^1}$ based at Q , and by definition $\tilde{\gamma}$ is its unique lift starting at P . Thus $\mu(f(\tilde{\gamma}))$ sends P to P' , and since P, P' were arbitrary in $f^{-1}(Q)$, we conclude the desired transitivity.

(ii) This follows immediately from Proposition 3.5 (ii), and the fact that μ is a homomorphism.

(iii) By Proposition 3.5 (iii), it is enough to consider the case that Q lies on such a small loop, say γ . By Lemma 2.6, we see that $f^{-1}(\gamma)$ is a disjoint union of paths, each contained in a standard neighborhood of some $P \in f^{-1}(Q_i)$. Fix such a standard neighborhood, so that we are simply looking at lifts of a loop around 0 under the map $z \mapsto z^{e_P}$. Now, Q has e_P preimages under this map, differing by powers of the e_P th root of unity $\zeta_{e_P} := e^{2\pi i/e_P}$, and we see that a lift of γ starting at one point P' will end at $\zeta_{e_P} P'$. Thus $\mu(\gamma)$ cyclically permutes the e_P points lying over Q in the standard neighborhood of P , and the desired statement follows. \square

This motivates:

Definition 3.6. A tuple $(\sigma_1, \dots, \sigma_r) \in (S_d)^r$ is a **Hurwitz factorization** of type $(d; T_1, \dots, T_r)$ if it satisfies (i)-(iii) of Proposition 3.5.

Thus, the proposition says that a branched cover of type $(d; T_1, \dots, T_r)$, together with a labeling of $f^{-1}(Q)$, yields a Hurwitz factorization of the same type. What happens if we choose a different labeling? We change all the σ_i by a fixed relabeling of $\{1, \dots, d\}$, or equivalently, conjugate them all by a fixed element of S_d .

Definition 3.7. Two Hurwitz factorizations $(\sigma_1, \dots, \sigma_r)$ and $(\sigma'_1, \dots, \sigma'_r)$ are **equivalent** if there exists a relabeling of $\{1, \dots, d\}$ sending σ_i to σ'_i for all i , or equivalently, if there exists $\tau \in S_d$ such that $\sigma_i = \tau \sigma'_i \tau^{-1}$ for all i .

We thus have that a branched cover yields a well-defined equivalence class of Hurwitz factorizations. On the other hand, given a Hurwitz factorization, by Proposition 3.5 (i) and (ii) we obtain a homomorphism $\pi_1(\widehat{\mathbb{C}\mathbb{P}^1}, Q) \rightarrow S_d$ with transitive image, and a basic theorem of topology relating the fundamental group to covering maps then implies that we obtain a corresponding topological covering map $\hat{C} \rightarrow \widehat{\mathbb{C}\mathbb{P}^1}$ of degree d , with \hat{C} connected. Applying the Riemann existence theorem, we conclude:

Theorem 3.8. *The monodromy map induces a bijection between equivalence classes of branched covers of type $(d; T_1, \dots, T_r)$ with branch points Q_1, \dots, Q_r , and equivalence classes of Hurwitz factorizations of the same type.*

4. HURWITZ THEORY

We immediately conclude the following from Theorem 3.8:

Corollary 4.1. *If we fix $Q_1, \dots, Q_r \in \mathbb{C}\mathbb{P}^1$, and a type $(d; T_1, \dots, T_r)$, then there are only finitely many equivalence classes of branched covers of type $(d; T_1, \dots, T_r)$ with branch points Q_1, \dots, Q_r . Moreover, this number does not depend on the Q_i .*

This leads us to the following:

Definition 4.2. The **Hurwitz number** $h(d; T_1, \dots, T_r)$ is the number of equivalence classes of branched covers of type $(d; T_1, \dots, T_r)$ with branched points Q_1, \dots, Q_r , for any choice of distinct $Q_i \in \mathbb{C}\mathbb{P}^1$. Equivalently, $h(d; T_1, \dots, T_r)$ is the number of equivalence classes of Hurwitz factorizations of type $(d; T_1, \dots, T_r)$.

Remark 4.3. We have given equivalent characterizations of branched covers in terms of complex geometry, topology, and group theory. One can also give an equivalent characterization in terms of algebraic geometry, as Chow's theorem states that any compact Riemann surface has a (unique) structure of a projective nonsingular curve, and holomorphic maps correspond to algebraic morphisms. Using this definition, we can generalize branched covers from the complex numbers to any field.

From this point of view, independence of the choice of Q_i is not at all a trivial fact: in fact, there are simple examples demonstrating that the same statement fails if we work instead over fields of positive characteristic.

Remark 4.4. In fact, the Hurwitz number is usually defined slightly differently: we use a weighted count of branched covers, dividing by the size of the automorphism group of each cover. The automorphisms of a cover correspond to the relabelings that leave the corresponding Hurwitz factorization unchanged, so group-theoretically, this corresponds to the number of Hurwitz factorizations divided by $d!$. We use the above definition because its motivation is clearer, but it turns out the standard definition is in many ways more natural. In any case, in many interesting cases there will not be any automorphisms of the covers in question, so the two definitions will agree.

Now we consider what happens if we allow the Q_i to vary. Given a choice of $\gamma = (\gamma_1, \dots, \gamma_r)$ satisfying the conditions of Proposition 3.5, let $\mathcal{U} = (U_1, \dots, U_r)$ be a tuple of small open disks centered at Q_1, \dots, Q_r respectively, and with U_i contained in the interior of γ_i for each i . Then for any choices of $Q'_i \in U_i$, we still have γ_i homotopic to a small loop around Q'_i . Denote by $\Delta_{\gamma, \mathcal{U}}(C, f)$ the set of all branched covers (C', f') with labeled branch points $Q'_i \in U_i$, and such that the equivalence classes of Hurwitz factorizations associated to each by γ are the same.

We can then define:

Definition 4.5. Given a type $(d; T_1, \dots, T_r)$, let the **Hurwitz space** $\mathcal{H}(d; T_1, \dots, T_r)$ of type $(d; T_1, \dots, T_r)$ be the set of branched covers (C, f) with labeled branch points and type $(d; T_1, \dots, T_r)$, equipped with the topology with base consisting of the sets $\Delta_{\gamma, \mathcal{U}}(C, f)$ as γ, \mathcal{U} and (C, f) vary over all possibilities.

We have:

Proposition 4.6. *The topology on $\mathcal{H}(d; T_1, \dots, T_r)$ is well defined; that is, the $\Delta_{\gamma, \mathcal{U}}(C, f)$ satisfy the conditions for a base of a topology.*

In order to examine this topology further, let $\mathcal{U}_r \subseteq (\mathbb{C}\mathbb{P}^1)^r$ denote the open subset consisting of r -tuples of distinct points in $\mathbb{C}\mathbb{P}^1$; thus, a choice of (Q_1, \dots, Q_r) corresponds to a point of \mathcal{U}_r . If we have a branched cover together with a labeling of its branch points, we obtain a point of \mathcal{U}_r , so we have a map $\mathcal{H}(d; T_1, \dots, T_r) \rightarrow \mathcal{U}_r$.

We then have the following consequence of the Riemann existence theorem:

Proposition 4.7. *The natural map*

$$\mathcal{H}(d; T_1, \dots, T_r) \rightarrow \mathcal{U}_r$$

is a topological covering map, of degree equal to $h(d; T_1, \dots, T_r)$.

Proof. It is clear that for any Q_1, \dots, Q_r , a tuple of open neighborhoods $\mathcal{U} = (U_1, \dots, U_r)$ gives an open neighborhood of (Q_1, \dots, Q_r) inside \mathcal{U}_r . Moreover, if we fix any choice of $\gamma = (\gamma_1, \dots, \gamma_r)$ as above, we see immediately that the preimage of \mathcal{U} in $\mathcal{H}(d; T_1, \dots, T_r)$ is, as a set, the disjoint union over all (C, f) with branch points Q_i of the sets $\Delta_{\gamma, \mathcal{U}}(C, f)$. It is moreover clear that each $\Delta_{\gamma, \mathcal{U}}(C, f)$ maps bijectively to $\mathcal{U} \subseteq \mathcal{U}_r$, so it is enough to see that $\Delta_{\gamma, \mathcal{U}}(C, f)$ maps homeomorphically to \mathcal{U} , and each $\Delta_{\gamma, \mathcal{U}}(C, f)$ is in fact a connected component of the preimage of \mathcal{U} . Since \mathcal{U} is connected, the latter follows from the former together with the observation that each $\Delta_{\gamma, \mathcal{U}}(C, f)$, being an element of a base for the topology, is open in $\mathcal{H}(d; T_1, \dots, T_r)$.

We thus prove that $\Delta_{\gamma, \mathcal{U}}(C, f)$ maps homeomorphically to \mathcal{U} . If we consider all $\mathcal{U}' = (U'_1, \dots, U'_r)$ with each U'_i an open disk contained in U_i , then the \mathcal{U}' give a base for the topology of \mathcal{U} , and their preimages in $\Delta_{\gamma, \mathcal{U}}(C, f)$ are precisely $\Delta_{\gamma, \mathcal{U}'}(C, f)$, so we conclude that the map $\Delta_{\gamma, \mathcal{U}'}(C, f)$ is continuous. The argument that it is open involves comparing different choices of γ , and is similar to the proof of Proposition 4.6.

Finally, the degree is by definition given by $h(d; T_1, \dots, T_r)$. \square

Since \mathcal{U}_r is naturally a complex manifold of (complex) dimension r , we immediately conclude:

Corollary 4.8. *The Hurwitz space $\mathcal{H}(d; T_1, \dots, T_r)$ is a (complex) manifold of dimension r .*

A basic question one can ask is the following:

Question 4.9. How can one characterize the connected components of $\mathcal{H}(d; T_1, \dots, T_r)$? In particular, for which $(d; T_1, \dots, T_r)$ is $\mathcal{H}(d; T_1, \dots, T_r)$ connected?

Just as with the Hurwitz number, the connected components of Hurwitz spaces can be interpreted purely in terms of group theory. In order to give this interpretation, we observe that there is a simple way to move from one Hurwitz factorization to others.

Definition 4.10. Given a Hurwitz factorization $(\sigma_1, \dots, \sigma_r)$ of type $(d; T_1, \dots, T_r)$, and i with $1 \leq i \leq r-1$, let $B_i(\sigma_1, \dots, \sigma_r)$ be the Hurwitz factorization of type $(d; T_1, \dots, T_{i-1}, T_{i+1}, T_i, T_{i+2}, \dots, T_r)$ given by

$$(\sigma_1, \dots, \sigma_{i-1}, \sigma_{i+1}, \sigma_{i+1}^{-1} \sigma_i \sigma_{i+1}, \sigma_{i+2}, \dots, \sigma_r).$$

Let the free group F_{r-1} act on the set of Hurwitz factorizations by the i th generator g_i acting as B_i ; in general, this permutes the indices in the type. If we define a homomorphism $F_{r-1} \rightarrow S_r$ by sending g_i to the transposition $(i, i+1)$, then the kernel P_r of this homomorphism does not permute the indices of the type when it acts on Hurwitz factorizations, so we obtain an action of P_r on Hurwitz factorizations of type $(d; T_1, \dots, T_r)$. This is called the **pure braid action**.

Example 4.11. If we simply repeat B_i twice for any given i , we obtain an action on Hurwitz factorizations of fixed type $(d; T_1, \dots, T_r)$ which is in general non-trivial.

We observe that the action of P_r is well-defined on equivalence classes. The main fact is then the following:

Proposition 4.12. *Two branched covers of type $(d; T_1, \dots, T_r)$, with fixed branch points Q_1, \dots, Q_r , lie in the same connected component of $\mathcal{H}(d; T_1, \dots, T_r)$ if and only if for some (equivalently all) choice of γ_i , their associated equivalence classes of Hurwitz factorizations lie in the same pure braid orbit.*

Sketch of proof. We first give a geometric interpretation of B_i . If we fix a choice of the γ_i , then the action of B_i is precisely what is obtained by switching Q_i with Q_{i+1} , while having the γ_i follows the points as they move, without letting them pass through any of the Q_j . It follows that for any combination of the B_i giving a pure braid action, we obtain a path in $\mathcal{H}(d; T_1, \dots, T_r)$ connecting each cover to its image under B_i , and thus the pure braid orbits are contained in the same connected component. Arguing as in Proposition 3.5 (i), proving the converse amounts to showing that the loops in \mathcal{U}_r obtained in this way from the pure braid group in fact generate the fundamental group of \mathcal{U}_r . \square