Columbia Journal of Mathematics **1** (2024),
8–27

A topological proof of the Riemann–Hurwitz formula

By Mrinalini Sisodia Wadhwa

Abstract

The Riemann–Hurwitz formula is generally given as a result from algebraic geometry that provides a means of constraining branched covers of surfaces via their Euler characteristic. By restricting to the special case of compact Riemann surfaces, we develop an alternative proof of the formula that draws on topology and manifold theory as opposed to more advanced algebraic machinery. We first discuss the foundation in manifold theory, defining Riemann surfaces and providing an example of the complex projective line. We then discuss the local topological structure of holomorphic maps between Riemann surfaces, introducing the notion of a *branched cover* and of *branch points*. Next, we discuss triangulations of a topological space and use this to introduce the Euler characteristic of Riemann surfaces. Using these definitions, we explicate and prove the Riemann–Hurwitz formula on compact Riemann surfaces. To conclude, we discuss consequences of this formula for adjacent fields such as algebraic topology. We provide visual intuition and examples throughout, drawing primarily on Szameuly's *Galois Groups and Fundamental Groups* (2009), as well as Forster's *Lectures on Riemann Surfaces* (1981), Guillemin and Pollack's *Differential Topology* (1974), and a few other supplementary sources. The main prerequisite for this paper is a background in topology and covering spaces.

1. **Introduction**

We begin with a topological problem. Suppose we have two surfaces, each with certain properties—such as holes, punctures, boundaries, and so forth—that are invariant under homeomorphism. We call such properties *topological invariants*. Can we always obtain a surjective map between these two surfaces that preserves their local structure? In fact, we cannot: as we will see at the end of this paper, there is no such surjective map from the complex projective plane (a surface with no holes) to the torus (a surface with one hole).

This prompts a second question: do the topological invariants of these surfaces tell us something about whether or not we can obtain such a map between the two surfaces? There is, in fact, an intricate relationship between the surfaces' topological invariants and the existence of a surjective map between them that preserves their structure. This relationship is given by the Riemann–Hurwitz formula, first proposed by Bernard Riemann (1826–66) in his 1857 *Theorie der Abel'schen Functionen* [*Theory of Abelian Functions*] [Rie57, §7]. We give a preliminary statement of the formula below and will define the terms used in the formula statement carefully in subsequent sections of the paper.

THEOREM 1.1 (Riemann–Hurwitz Formula). *Let $\varphi\colon Y \to X$ be a holomorphic map of compact Riemann surfaces with degree d as a branched cover. The Euler characteristics $\chi(X)$ and $\chi(Y)$ of X and Y are related by the formula*

$$\chi(Y) = d \cdot \chi(X) - \sum_{y} (e_y - 1),$$

where the sum is over the branch points of φ and e_y is the ramification index corresponding to each branch point $y \in Y$.

A *branched cover* is a particularly well-structured surjective map between Y and X, and the ramification index corresponds to "sheets" of the cover intersecting with one another. As seen in the above statement, Riemann was looking at a specific class of surfaces and maps between them—namely, Riemann surfaces and holomorphic maps, concepts he had introduced in his 1851 doctoral dissertation that now serve as the foundation for the field of complex analysis. He appears to have died without offering a proof of this formula [Oor16, p.568–69]. The first proof was likely Adolf Hurwitz's (1858–1919) argument in his 1891 paper, *Über Riemann'sche Fläche mit gegebenen Verzweigungspunkten* [*On Riemann surfaces with Given Branch Points*] [Hur91, p.375–76].

This formula has subsequently been generalized to an algebraic-geometric version that takes X and Y to be smooth curves (rather than Riemann surfaces) and φ to be a morphism between them (rather than a holomorphic map) [Oor16, p.573–74]. It is in this abstract form—within the context of algebraic geometry—that most students now encounter the Riemann–Hurwitz formula. The usual proof of this version of the formula, given in [Sta18, Tag 0C1B], relies upon spectral sequences and other abstract-algebraic machinery. To avoid getting lost in the thickets of algebraic geometry, we will restrict to original case of Riemann surfaces and holomorphic maps between them. From this, we can develop a proof of the Riemann–Hurwitz formula that uses topology and manifold theory, providing the elusive topological and geometric intuition for the formula that draws us back to our motivating problem—how to understand the relationship between a surjective map between two surfaces and their topological invariants. The goal of this paper is to offer such a proof, following [Sza09, §3.6].

We proceed in four sections. Section 2 provides a foundation in manifold theory, defining Riemann surfaces and discussing the complex projective line as an example. Section 3 discusses holomorphic maps between Riemann surfaces and their local topological structure, introducing the notion of branch points and a branched cover. Section 4 discusses triangulation and the Euler characteristic of Riemann surfaces. Finally, Section 5 completes the proof of the Riemann–Hurwitz formula and discusses some interesting corollaries for algebraic topology.

This paper assumes a background in topology—specifically point-set topology and covering spaces. The reader does not need an extensive background in manifold theory or complex analysis. Rather, the relevant concepts from these fields—specifically holomorphisms and complex manifolds—are explained in the Section 2 with reference to the real case, identifying $\mathbb{R}^2$ with $\mathbb{C}$. It should also be noted that because complex analysis is not the main focus of this paper, we either assume or sketch complex-analytic results as needed to complete the major proofs in this paper, particularly in Section 4. Wherever possible, we provide pictures and visual intuition for definitions and proofs.

2. **Riemann surfaces**

This section grounds this paper in the relevant manifold theory, drawing on [For07, p.1–12] and [Sza09, §3.1–3.2]. We build up to a definition of Riemann surfaces and discuss some examples.

We begin by defining a holomorphic map between subsets of $\mathbb{C}^n$ and a complex atlas on a manifold.

Definition 2.1. Let $U \subset \mathbb{C}^n$ and $V \subset \mathbb{C}^m$. A map $f: U \to V$ is **holomorphic** if, for every $x \in U$, there exists a neighborhood $U_x \subset U$ of x such that f is complex-differentiable everywhere in U_x.

This is the complex analogue of a smooth map in real analysis, although being holomorphic is a vastly stronger condition: a complex-differentiable map is both infinitely differentiable and analytic (unlike in the real case, where C^1 maps are not necessarily C^∞, and where C^∞ maps are not necessarily analytic).

Definition 2.2. Let X be a topological 2-manifold. A **complex chart** on X is a pair $(U_i \subset X, f_i: U_i \to f_i(U_i) \subset \mathbb{C})$ such that U_i is an open subset of X and f_i is a homeomorphic mapping from U_i onto its image $f(U_i) \subset \mathbb{C}$.

We say a chart (U_i, f_i) is **centered** at $x \in U_i$ if $f_i(x) = 0$. Two charts (U, f), (V, g) are **holomorphically compatible** if their transition maps $f \circ g^{-1}$ and $g \circ f^{-1}$ are holomorphic where defined. This is illustrated in Figure 1.

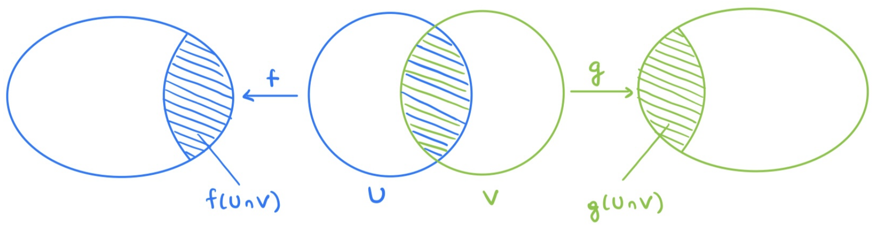

Figure 1. Two charts (U, f) and (V, g), with transition map $f \circ g^{-1}$ defined on $g(U \cap V)$ and $g \circ f^{-1}$ defined on $f(U \cap V)$, based on the illustration of the real case in [Tu07, §5.2, Fig 5.2].

Definition 2.3. A **complex atlas** $\mathcal{U}$ on X is a collection of holomorphically compatible charts (U_i, f_i) such that the $\{U_i\}$ form an open cover of X.

We say two atlases (U_i, f_i), (V_j, g_j) on X are **equivalent** if their union, defined by taking all U_i and V_j as a covering of X and all complex charts, is also a complex atlas on X. In particular, this implies that $f_i \circ g_j^{-1}$ and $g_j \circ f_i^{-1}$ are holomorphic on their respective domains for all i, j. We now proceed to define a Riemann surface by placing a complex structure on X, in a manner analogous to how [Tu07, §2.5] discusses placing a smooth structure in the real case.

Definition 2.4. A **Riemann surface** is a topological 2-manifold X with an equivalence class of complex atlases (which we call a **complex structure** on X).

As a trivial example, consider any open subset $U \subset \mathbb{C}$. Then U is a Riemann surface with the complex atlas $(U, i\colon U \hookrightarrow \mathbb{C})$, where i is the inclusion map. We consider one nontrivial example, the complex projective line, which we return to in subsequent sections of this paper.

Example 2.5 (Complex projective line $\mathbb{CP}^1$). Let $\mathbb{CP}^1 = \mathbb{C} \cup \{\infty\}$, where ∞ is an extra point not included in $\mathbb{C}$. We topologize $\mathbb{CP}^1$ as follows: the open sets are the usual open sets $U \subset \mathbb{C}$ and sets of the form $V \cup \{\infty\}$, where $V \subset \mathbb{C}$ is the complement of a compact set $K \subset \mathbb{C}$. We call $\mathbb{CP}^1$ with this topology the **complex projective line**, and we see that it is homeomorphic to the 2-sphere $S^2 \subset \mathbb{R}^3$ with antipodal points identified with 0 and ∞, as shown in Figure 2.

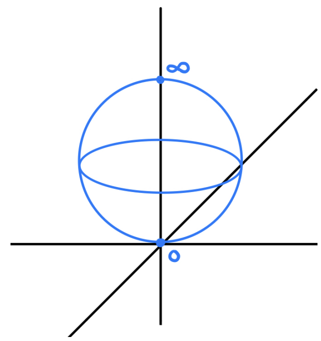

Figure 2. $\mathbb{CP}^1$ homeomorphic to S^2.

Now we define a complex atlas on $\mathbb{CP}^1$. Let $U_1 := \mathbb{CP}^1 \setminus \{\infty\} = \mathbb{C}$, and let f_1 be the identity map z on U_1. Then, let $U_2 := \mathbb{CP}^1 \setminus \{0\}$, and define the map f_2 as follows:

$$f_2(z) = \begin{cases} \frac{1}{z} & z \in U_2 \setminus \{\infty\} \\ 0 & z = \infty. \end{cases}$$

Then both f_1 and f_2 are well-defined homeomorphisms onto their images. The charts $(U_1, f_1), (U_2, f_2)$ cover $\mathbb{CP}^1$ and are holomorphically compatible, as their transition maps

$$f_1 \circ f_2^{-1} = f_2 \circ f_1^{-1} \colon U_1 \cap U_2 = \mathbb{C} \setminus \{0\} \to U_1 \cap U_2 = \mathbb{C} \setminus \{0\}$$

are given by $z \mapsto \frac{1}{z}$. Thus, they form a complex atlas on $\mathbb{CP}^1$.

If we consider $\mathbb{CP}^1$ under the homeomorphism to S^2, then the maps f_1 and f_2 correspond to the stereographic projection from $\mathbb{R}^2 \times \{0\}$ and $\mathbb{R}^2 \times \{1\}$, respectively, as shown in Figure 3.

$\Diamond$

3. Holomorphic maps and branched covers

We now define holomorphic maps between Riemann surfaces, discuss their local topological structure, and introduce the notion of a branched cover.

Definition 3.1. Let X and Y be Riemann surfaces. A continuous map $\varphi\colon Y \to X$ is **holomorphic** if for every pair of charts (U, f) on Y and (V, g) on X such that $\varphi(U) \subset V$, the map $g \circ \varphi \circ f^{-1} \colon f(U) \subset \mathbb{C} \to g(V) \subset \mathbb{C}$ is holomorphic in the usual sense (i.e., the sense of Definition 2.1).

This definition is visualized in Figure 4. Note that this corresponds to our definition of a smooth map between smooth manifolds in the real case in [Tu07, §2.6].

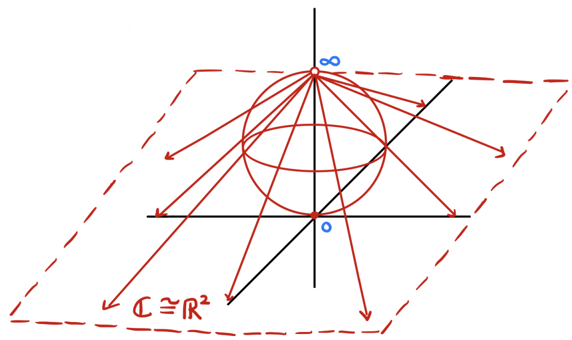

(a) The chart (U_1, f_1) on $\mathbb{CP}^1$.

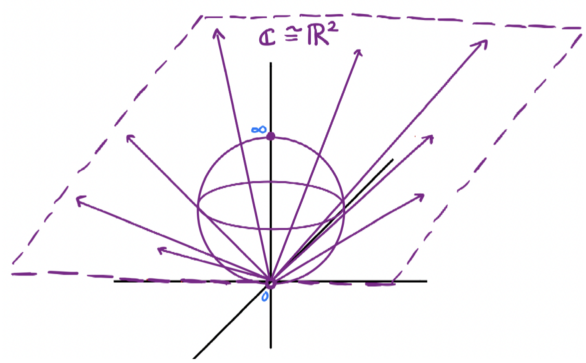

(b) The chart (U_2, f_2) on $\mathbb{CP}^1$.

Figure 3. A complex atlas on $\mathbb{CP}^1$, visualized under homeomorphism to S^2.

Henceforth, to avoid the trivial case, we assume that all holomorphic maps between Riemann surfaces in this paper are nonconstant on all connected components—i.e., that they do not map an entire connected component to a single point. We remarked after Definition 2.1 that being holomorphic is a stronger condition than being smooth, as holomorphic maps are both infinitely differentiable and analytic on subsets of $\mathbb{C}^n$. As we shall see, this means that we actually know a great deal more about the local structure of holomorphic maps than we do about smooth maps in the real case. This is summarized in the below proposition, which tells us that locally, every holomorphic map is just exponentiation.

PROPOSITION 3.2. *Let $\varphi\colon Y \to X$ be a holomorphic map of Riemann surfaces and $y \in Y$ with image $\varphi(y) = x$ in X. Then there exist open neighborhoods $V_y \subset Y$ and $U_x \subset X$ of y and x respectively satisfying $\varphi(V_y) \subset U_x$,*

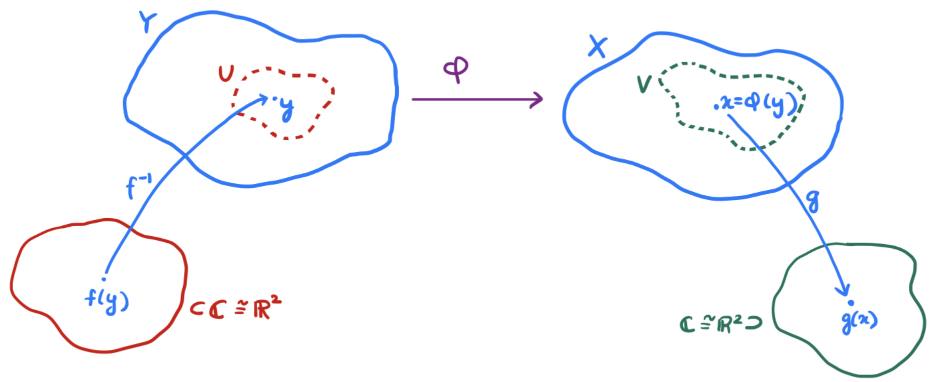

Figure 4. A holomorphic map $\varphi\colon Y \to X$ with charts (U, f) on Y and (V, g) on X, based on the illustration of the real case in [Tu07, §2.6, Figure 6.3].

as well as homeomorphisms $g_y\colon V_y \to g_y(V_y) \subset \mathbb{C}$ *and* $f_x\colon U_x \to f_x(U_x) \subset \mathbb{C}$ *satisfying* $f_x(x) = g_y(y) = 0$ *such that the diagram*

$$\begin{tikzcd} V_y \arrow[r, "\varphi"] \arrow[d, "g_y"'] & U_x \arrow[d, "f_x"] \\ \mathbb{C} \arrow[r, "z\mapsto z^{e_y}"'] & \mathbb{C} \end{tikzcd}$$

commutes for an appropriate positive integer e_y *chosen with respect to* y *that does not depend upon the choice of* g_y *or* f_x.

Figure 5 provides a geometric visualization of the commutative diagram. We proceed to sketch its proof, drawing on some results from complex analysis.

Proof sketch of Proposition 3.2. First, by selecting and shrinking neighborhoods U_x and V_y as necessary and performing linear transformations in $\mathbb{C}$, we can find charts (V_y, g'_y) and (U_x, f_x) centered at y and x, respectively. We will now modify these in order for the diagram to commute. As φ is a holomorphism from Y to X, we know by Definition 3.1 that $f_x \circ \varphi \circ g'^{-1}_y$ is holomorphic in a neighborhood of 0 and vanishes at 0. As holomorphic maps are necessarily analytic, complex analysis tells us that $f_x \circ \varphi \circ g'^{-1}_y$ must be of the form $z \mapsto z^{e_y} H(z)$, where H is a holomorphic function such that $H(0) \neq 0$.

We denote by log a fixed branch of the logarithm function in a neighborhood of $H(0)$. Now we apply complex analysis results to conclude: we shrink the neighborhood V_y as necessary so that $h := \exp((1/e_y) \log H)$ defines a holomorphic function h on $g'_y(V_y)$ such that $h^{e_y} = H$, and then we define g_y to be the composition of g'_y with the map $z \mapsto zh(z)$. This yields charts (V_y, g_y) and (U_x, f_x) centered at y and x respectively such that the diagram commutes.

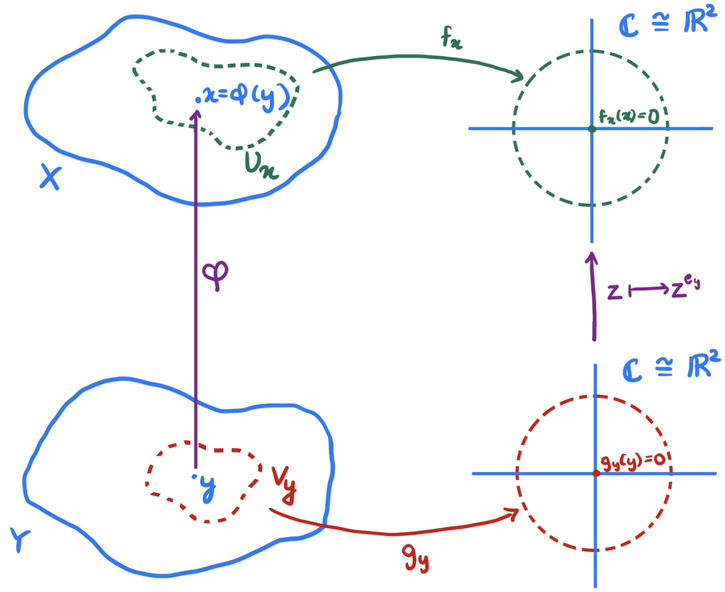

Figure 5. A geometric visualization of the local structure on holomorphic maps.

We observe moreover that e_y, defined in relation to an invertible holomorphic map, is necessarily a positive integer independent of the choice of g_y, f_x. ☐

Having established this local structure, we introduce the notions of ramification index, branch points, and branched cover, following [Sza09, §3.2].

Definition 3.3. The positive integer e_y in Proposition 3.2 is called the **ramification index** of φ at y. The points $y \in Y$ such that $e_y > 1$ are called the **branch points** of φ. We denote the set of branch points of φ by S_φ.

Remark 3.4. Note that S_φ is a discrete closed subset of Y. This follows from Proposition 3.2: given any $y \in Y$, there exists a punctured open neighborhood V_y of y that contains no branch points where φ has finitely many points in its preimage (due to the local structure of the map $z \mapsto z^{e_y}$).

From this observation, we proceed to introduce the notion of a branched cover and relate it to this local structure on holomorphic maps. First we must define a proper map.

Definition 3.5. A continuous map of locally compact topological spaces $\varphi\colon N \to M$ is **proper** if the preimage of each compact subset of M under φ is compact in N.

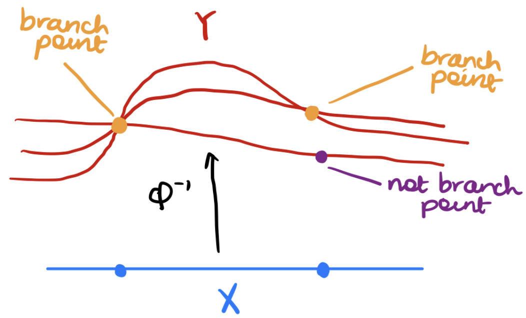

(a) A finite branched cover φ.

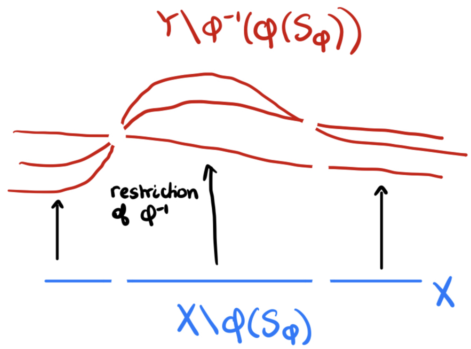

(b) Restricting φ to obtain a covering space.

Figure 6. Visualization of a branched cover.

Definition 3.6. Given locally compact Hausdorff spaces N and M, a proper surjective map $\varphi \colon N \to M$ is a **finite branched cover** if it restricts to a finite cover (of M) outside a discrete closed subset (of N).

Its **degree** is defined to be the degree of the finite cover obtained by its restriction.

We can think of a finite branched cover as essentially a covering space at all but a small number of points (namely, the branch points, which lie within the discrete closed subset). At the branch points, we can visualize the sheets of the cover merging together, so that the sheets of the cover "branch out" from them. Thus, when we remove these points, we obtain a covering space in its regular topological sense, as shown in Figure 6.

Finally, we relate this notion of a branched cover to holomorphic maps between Riemann surfaces with the following rather wonderful result.

THEOREM 3.7. *Let X be a connected Riemann surface, and let $\varphi\colon Y \to X$ be a proper holomorphic map. Then φ is a finite branched cover.*

Proof of Theorem 3.7. This result follows from Proposition 3.2 and Remark 3.4.

First, by definition, X and Y are locally compact Hausdorff spaces because they are Riemann surfaces.

Second, we claim φ is surjective because it is holomorphic and proper. Proposition 3.2 implies that as a holomorphic map between Riemann surfaces, φ is in fact an open map, since the map $z \mapsto z^{e_y}$ is open and f_x, g_y are homeomorphisms and therefore open maps. Thus $\varphi(Y)$ is open in X. Moreover, because φ is proper and X, Y are Hausdorff and locally compact, φ is a closed map, because in a locally compact Hausdorff space a subset is closed if and only if its intersection with every compact subset is closed. Thus $\varphi(Y)$ is closed in X. Then $\varphi(Y)$ is a nonempty clopen subset of X, so we must have $\varphi(Y) = X$ as X is connected, proving that φ is surjective.

Third, we have from Remark 3.4 that S_φ is a discrete closed subset of Y.

To conclude, we claim that the restriction of the map φ to $Y \setminus \varphi^{-1}(\varphi(S_\varphi))$ is a finite topological cover of $X \setminus \varphi(S_\varphi)$. This follows again from Proposition 3.2: given $x \in X \setminus \varphi(S_\varphi)$, each of the finitely many points in the preimage $\varphi^{-1}(x)$ has an open neighborhood that maps homeomorphically onto an open neighborhood of x. The intersection of these open neighborhoods is an open neighborhood of x that satisfies the definition of a finite topological cover (demonstrated in Figure 6). $\square$

In light of this result, we will now take as a given that a holomorphic map φ as above yields a finite branched cover in subsequent sections of this paper.

4. **Triangulation of Riemann surfaces**

This section completes the setup for our topological proof of the Riemann–Hurwitz formula, carefully defining and providing geometric intuition for the various terms used in the formula statement. We define triangulation on a compact topological 2-manifold (and thus on any Riemann surface), prove that every compact Riemann surface has a triangulation, and introduce the concept of the Euler characteristic of a compact Riemann surface.

Intuitively, a triangulation divides up a space into smaller "triangles"—closed subsets of the space that map homeomorphically onto unit triangles in $\mathbb{R}^2$—that are glued together at edges or vertices. We formalize this notion below.

Definition 4.1. Let X be a compact topological 2-manifold. A **triangulation** of X consists of a finite system $\mathcal{T} = \{T_1, \ldots, T_n\}$ of closed subsets of

X whose union is the whole of X, and homeomorphisms $\varphi_i \colon \Delta \to T_i$, where Δ is the unit triangle in $\mathbb{R}^2$.

We say that the T_i are the **faces** of the triangulation, and that the images of the vertices (respectively edges) under φ_i of Δ are the **vertices** (respectively **edges**) of the triangulation. These must satisfy the following conditions:

(1) Each vertex (respectively edge) of $\mathcal{T}$ contained in a face T_i should be the image of a vertex (respectively edge) of Δ under φ_i;
(2) Any two different faces must either be disjoint, or intersect at a single vertex, or intersect at a single edge.

As an example, we consider a triangulation on the 2-sphere S^2, the underlying topological structure for the complex projective line $\mathbb{CP}^1$ discussed in Example 2.5.

Example 4.2 (Triangulation on the 2-sphere). By cutting S^2 along the equator and two meridians, we obtain a triangulation $\mathcal{T}$ with 6 vertices, 8 faces, and 12 edges. Figure 7 provides a visualization of $\mathcal{T}$ and the homeomorphic map from the unit triangle to one of its closed subsets T_1. Since S^2 is homeomorphic to $\mathbb{CP}^1$, this implies that there is a corresponding triangulation of $\mathbb{CP}^1$.

$\Diamond$

Proposition 4.3 (Refinement of a triangulation). *Given a particular triangulation $\mathcal{T}$ of a compact topological space X and a point $x \in X$ that is not a vertex of $\mathcal{T}$, we can refine $\mathcal{T}$ to include x as a vertex.*

Proof of Proposition 4.3. There are two cases: either x lies in the interior of a face of $\mathcal{T}$ or it lies on an edge of $\mathcal{T}$.

Case 1: Take the face $\varphi_i(\Delta)$ that contains x, and consider the natural subdivision of Δ that arises from joining $\varphi_i^{-1}(x)$ to the vertices and replace φ_i with its restrictions to the smaller triangles Δ_1, Δ_2 and Δ_3 that arise from the subdivision (where each Δ_i is homeomorphic to the unit triangle Δ in $\mathbb{R}^2$).

Case 2: Take the two faces $\varphi_i(\Delta)$ and $\varphi_j(\Delta)$ that meet at the edge on which x lies, and repeat the same process, considering the natural subdivision of Δ that arises from joining $\varphi_i^{-1}(x) = \varphi_j^{-1}(x)$ to the vertices and replace φ_i and φ_j with their restrictions to the smaller triangles $\Delta_1, \Delta_2, \Delta_3$ and Δ_4 (where, likewise, each Δ_i is homeomorphic to the unit triangle Δ in $\mathbb{R}^2$).

This process is illustrated in Figure 8.

$\square$

We will now prove an important result, following [Sza09, §3.6], which will set up our definition of the Euler characteristic.

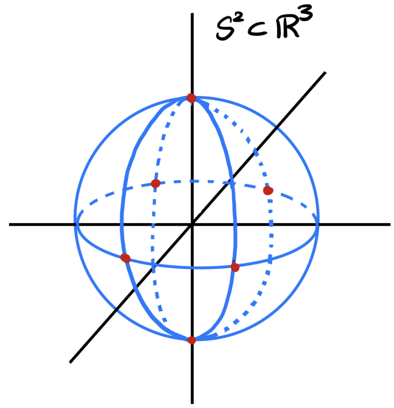

(a) A triangulation
$\mathcal{T}$of S^2.

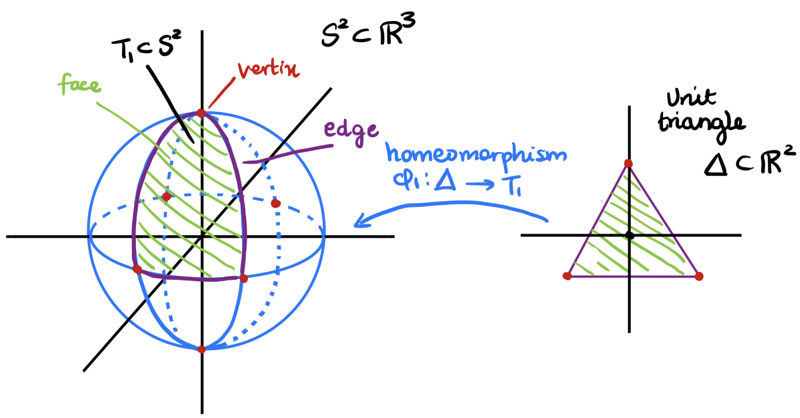

(b) Homeomorphic mapping of the unit triangle Δ to $T_1 \in \mathcal{T}$.

Figure 7. Illustration of Example 4.2.

THEOREM 4.4. *Every compact Riemann surface has a triangulation.*

To prove this theorem, we begin with an arbitrary compact Riemann surface, and use results from complex analysis and topology to reduce this to the case of $\mathbb{CP}^1$, for which we know there is a triangulation by Example 4.2. We proceed in three steps. First, we show that a triangulation can be canonically lifted via a finite branched cover. Second, we sketch a proof using complex analysis that any compact Riemann surface yields a finite branched cover of $\mathbb{CP}^1$. Finally, we relate these findings and Example 4.2 to conclude.

LEMMA 4.5. *Let $\varphi\colon Y \to X$ be a finite branched cover of compact Riemann surfaces Y and X (in particular, following Theorem 3.7, consider the*

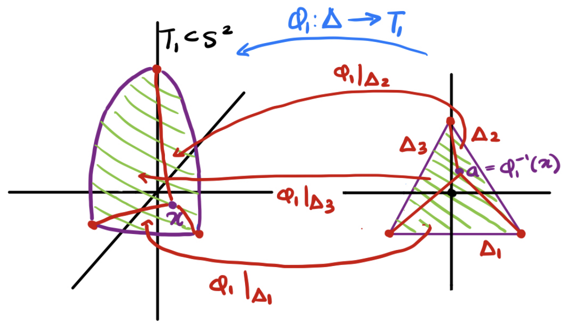

(a) Case 1, where x does not lie on an edge of $\mathcal{T}$.

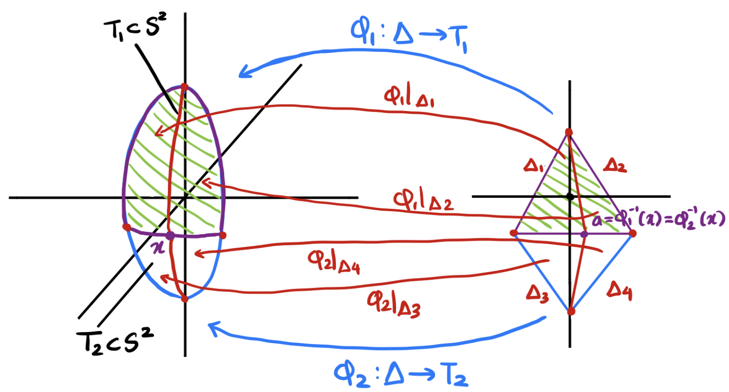

(b) Case 2, where x lies on an edge of $\mathcal{T}$.

Figure 8. Two possible cases for refining a triangulation.

case where X is connected and φ is a proper surjective holomorphic map $Y \to X$). Then every triangulation of X can be lifted canonically to a triangulation of Y.

Proof of Lemma 4.5. Take some triangulation $\mathcal{T}$ on X with faces $\{T_i\}$ and homeomorphisms $\psi_i : \Delta \to T_i$, and let S_0 be the set of all vertices of $\mathcal{T}$. Following the process described in Proposition 4.3, we can refine $\mathcal{T}$ as necessary so that S_0 contains all images of branch points, i.e., for every $x \in X$ such that $x = \varphi(y)$ for some $y \in S_\varphi$, we have $x \in S_0$. Then the definition of a finite branched cover implies that the restriction of φ to $X \setminus \varphi^{-1}(S_0)$ is a cover.

Let Δ' be the subset of Δ obtained by omitting all vertices. We observe that Δ' is simply connected because it is contractible: as the triangle is filled-in and convex, we can take the straight-line homotopy to contract it to its center point. Then the fundamental group of Δ' is trivial and thus the restriction of the branched cover $\varphi\colon Y \to X$ above each of $\psi_i(\Delta')$ is trivial. This implies that we can canonically lift the restriction of each ψ_i to Δ' to each sheet of the cover $X \setminus \varphi^{-1}(S_0)$. We can also canonically lift all vertices of $\psi_i(\Delta)$ that are not the images of branch points. This demonstrates that the triangulation of X gives rise to a triangulation of Y away from the branch points.

It remains to show that we also have a triangulation of Y at the branch points. We revisit Proposition 3.2 regarding the local structure of holomorphic maps to consider the behaviour of φ near branch points. For any branch point $y \in Y$, Proposition 3.2 implies that there is a neighborhood of y on which φ locally looks like the continuous open map $z \mapsto z^{e_y}$. Then we can apply the process outlined in Proposition 4.3 to refine the triangulation again, adding each branch point y as a vertex. This yields the desired triangulation of Y.

□

This process of lifting a triangulation is illustrated in Figure 9, where we lift a triangle from X to Y via a branched cover of degree 3. We shall revisit this process and diagram in Section 5 while proving the Riemann–Hurwitz theorem.

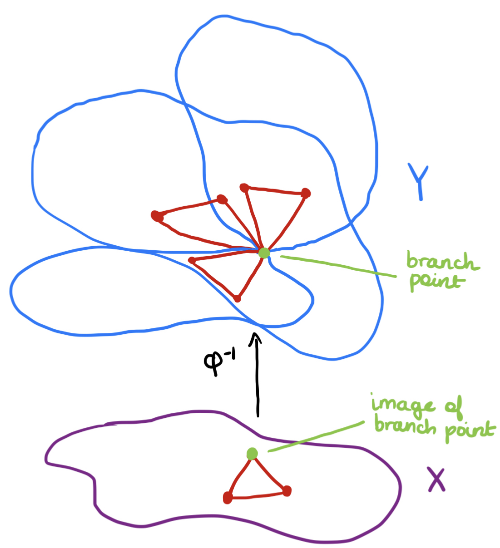

Figure 9. A triangulation of X lifts canonically to a triangulation of Y when $\varphi\colon Y \to X$ is a finite branched cover.

We move on to the second step of the proof of Theorem 4.4.

Lemma 4.6. *Given a connected compact Riemann surface Y, there exists a nonconstant holomorphic map $Y \to \mathbb{CP}^1$.*

The proof of this lemma utilizes two results from complex analysis that we will state below but not prove. Full proofs are given in [For07, Corollary 14.13, Theorem 1.8].

Lemma 4.7 (Riemann's existence theorem). *Let X be a compact Riemann surface, $x_1, \dots, x_n \in X$ a finite set of points, and $a_1, \dots, a_n$ a sequence of complex numbers. Then there exists a function f on X that satisfies the following conditions*:

(1) *f is holomorphic everywhere on $X \setminus S$, where $S \subset X$ is a discrete closed subset, and for all complex charts $(U, \varphi\colon U \to \mathbb{C})$, the complex function $f \circ \varphi^{-1}$ is holomorphic everywhere except on a discrete closed subset of the domain*[1];
(2) *f is holomorphic at all the x_i, with $f(x_i) = a_i$ for all i from 1 to n.*

□

Lemma 4.8 (Riemann's removable singularities theorem). *Let U be an open subset of a Riemann surface X, let $a \in U$, and let f be some function that is holomorphic on $U \setminus \{a\}$. Suppose f is bounded in some neighborhood of a. Then f can be extended uniquely to a function f' that is holomorphic on U.* □

Proof sketch of Lemma 4.6. Since Y is a compact Riemann surface, Lemma 4.7 gives a nonconstant function $f\colon Y \to \mathbb{C}$ that satisfies conditions (1) and (2) in its statement. We now define a map $\varphi_f\colon Y \to \mathbb{CP}^1 = \mathbb{C} \cup \{\infty\}$ as follows:

$$\varphi_f(y) = \begin{cases} f(y) & y \neq 0, \infty \\ \infty & y = 0, \infty. \end{cases}$$

We know f is holomorphic on all but a discrete set of points by condition (1), so given some $y \in Y$, we can choose a chart $(U, g\colon U \to \mathbb{C})$ centered around y such that f is holomorphic on $U \setminus \{y\}$ (shrinking U as necessary). Recall from Example 2.5 that the two standard complex charts on $\mathbb{CP}^1$ are given by z and $\frac{1}{z}$ on $\mathbb{CP}^1 \setminus \{\infty\} = \mathbb{C}$ and $\mathbb{CP}^1 \setminus \{0\}$, respectively. Now there are two cases. If f is holomorphic at y, then $z \circ \varphi_f \circ g^{-1}$ is holomorphic on $g(U)$. If f is not holomorphic at y, we apply Lemma 4.8: $(\frac{1}{z}) \circ \varphi_f \circ g^{-1}$ maps $g(U \setminus \{y\})$ to a bounded open subset of $\mathbb{C}$ and thus extends to a holomorphic function on $g(U)$. We conclude that φ_f is holomorphic. Moreover, since f is nonconstant, φ is also nonconstant by construction. □

[1] This is equivalent to f being a **meromorphic** function on X, which [Sza09, §3.3] discusses formally. We eschew further discussion of meromorphic functions here to avoid losing sight of the goal of this section: proving Theorem 4.4.

We finally piece together these results to complete the proof of Theorem 4.4.

Proof of Theorem 4.4. Take any compact Riemann surface Y, and consider its components (necessarily finitely many as Y is compact), which are connected compact Riemann surfaces $Y_1, \dots, Y_n$. By Lemma 4.6, there exist nonconstant holomorphic maps $\varphi_1, \dots, \varphi_n$ such that each φ_i maps Y_i into $\mathbb{CP}^1$. By Example 4.2, there is a triangulation on $\mathbb{CP}^1$. By Lemma 4.5, since each φ_i is a holomorphic map from a compact connected Riemann surface Y_i to $\mathbb{CP}^1$, another compact connected Riemann surface, we have that φ_i is a branched cover and that our triangulation of $\mathbb{CP}^1$ from Example 4.2 lifts to a triangulation of Y_i, say $\mathcal{T}_i$. We can then piece together these triangulations by taking their union to obtain a triangulation $\mathcal{T}$ on all of Y, completing the proof. □

Finally, we introduce the notion of an Euler characteristic, following the definition given in [Sza09, §3.6].

Definition 4.9. Given a triangulation $\mathcal{T}$ of a compact Riemann surface X, denote by S_0, S_1, and S_2 the set of vertices, edges, and faces of $\mathcal{T}$, respectively. Let s_0, s_1, and s_2 be their respective cardinalities. Then we define the **Euler characteristic** of X to be $\chi(X) := s_0 - s_1 + s_2$.

This is the classical definition of the Euler characteristic. It is in fact equivalent to the definition based on intersection theory in [GP78, p.116], though proving this equivalence casts beyond this paper's scope. Intuitively, this definition offers us a means of classifying compact Riemann surfaces based on their triangulations.

Note that we need Theorem 4.4 to ensure that the Euler characteristic is defined for all compact Riemann surfaces, as any compact Riemann surface must have a triangulation by Theorem 4.4 and therefore its Euler characteristic can be computed using the given formula. Moreover, the Euler characteristic is well-defined independent of the choice of triangulation on a given compact Riemann surface. To see this, notice that the Euler characteristic remains unchanged under the process of refining a triangulation described in Proposition 4.3 and illustrated in Figure 8. In both Case 1 (where we added x as a vertex when x was not on an edge) and Case 2 (where we added x as a vertex when x was on an edge), the Euler characteristic with the refined triangulation is $(s_0+1)-(s_1+3)+(s_2+2) = s_0-s_1+s_2$, the same as the original. Then, given any two triangulations of a compact Riemann surface, we can take their common refinement and thereby obtain the same value for its Euler characteristic throughout.

As an example, we return to the case of S^2, homeomorphic to $\mathbb{CP}^1$ as discussed in Examples 2.5 and 4.2.

Example 4.10. The triangulation $\mathcal{T}$ given in Example 4.2 and illustrated in Figure 7 has 8 faces, 6 vertices, and 12 edges, which implies that $\chi(S^2) = 6 - 12 + 8 = 2$. Indeed, any other triangulation of S^2 yields the same calculation for the Euler characteristic. For example, Figure 10 shows another triangulation $\mathcal{T}'$ of S^2 obtained by cutting along the equator and twice in the upper hemisphere. This triangulation has 4 faces, 4 vertices, and 6 edges, so again we compute $\chi(S^2) = 4 - 6 + 4 = 2$.

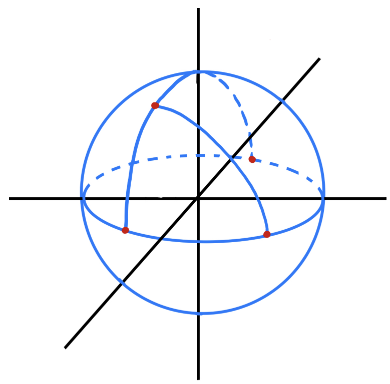

Figure 10. Another triangulation $\mathcal{T}'$ of S^2.

$\diamondsuit$

5. **Proof of Riemann–Hurwitz formula**

Finally, we move to prove the Riemann–Hurwitz formula, given as Theorem 1.1 in the introduction, which we restate for convenience. Let $\varphi \colon Y \to X$ be a holomorphic map of compact Riemann surfaces with degree d as a branched cover. The Euler characteristics $\chi(X)$ and $\chi(Y)$ of X and Y are related by the formula

$$\chi(Y) = d \cdot \chi(X) - \sum_{y} (e_y - 1),$$

where the sum is over the branch points of φ and e_y is the ramification index corresponding to each branch point $y \in Y$.

The proof follows from closely revisiting the process of lifting a triangulation via a branched cover discussed in Lemma 4.5 and illustrated in Figure 9.

Proof of Theorem 1.1 (*The Riemann–Hurwitz formula*). Take any triangulation on X, and let s_0, s_1 and s_2 be the number of vertices, faces, and edges, respectively. Consider its canonical lifting to a triangulation of Y via the finite branched cover φ, given by the process outlined in the proof of Lemma 4.5. Notice that, by construction, all branch points $y \in Y$ of φ correspond to vertices of the lifted triangulation and therefore do not lie on edges or faces.

Thus, edges and faces are lifted canonically on the cover of degree d, so the number of edges and the number of faces of the lifted triangulation are equal to ds_1 and ds_2, respectively. For vertices on the lifted triangulation, there are two cases. Vertices that do not correspond to the images of branch points have d preimages as well, as the covering space is of degree d. However, at any branch point y, we have to account for the e_y sheets of the branched cover merging together, and thus the number of preimages is instead $d - (e_y - 1)$. Thus the number of vertices of the lifted triangulation can be written as $ds_0 - \sum_{y \in S_\varphi}(e_y - 1)$, and we can compute the Euler characteristic as follows:

$$\begin{aligned}\chi(Y) &= \left(ds_0 - \sum_{y \in S_\varphi}(e_y - 1)\right) - ds_1 + ds_2 \\ &= d(s_0 - s_1 + s_2) - \sum_{y \in S_\varphi}(e_y - 1) \\ &= d \cdot \chi(X) - \sum_{y \in S_\varphi}(e_y - 1).\end{aligned}$$

□

This is visually illustrated in the lifting of a triangle in Figure 9, where φ is a branched cover of degree 3. We notice that the 3 edges of the triangle in X are each lifted to 3 edges (for a total of 9 edges) in Y, and likewise that the 1 face of the triangle in X is lifted to 3 faces in Y. The 2 vertices in X that do not correspond to branch points are each lifted to 3 vertices in Y, but the vertex that corresponds to a branch point (with ramification index 3, as 3 sheets merge) is only lifted to $1 = 3 \cdot 1 - (3 - 1)$ vertex in Y. This gives a total of $7 = 3 \cdot 3 - (3 - 1)$ vertices in the lifted triangle, providing visual intuition for the proof.

This gives us the major result of this paper, and we conclude with a brief discussion of its implications. Because we are working with compact Riemann surfaces (instead of smooth curves in the algebraic geometry setting), we can apply some results from algebraic topology to restate the statement of Theorem 1.1 in more specific terms. In particular, any compact Riemann surface X is homeomorphic to a torus with g holes. The proof of this result, given in [Ful95, Theorem 17.4], utilizes the fact that compact Riemann surfaces are orientable topological 2-manifolds and a method of "cutting and pasting." We call g the **genus** of X, and can thus classify compact Riemann surfaces in terms of their genera, as depicted in Figure 11: compact Riemann surfaces of genus 0 are homeomorphic to S^2 and $\mathbb{CP}^1$, those of genus 1 are homeomorphic to the torus, those of genus 2 are homeomorphic to the 2-torus, and so forth.

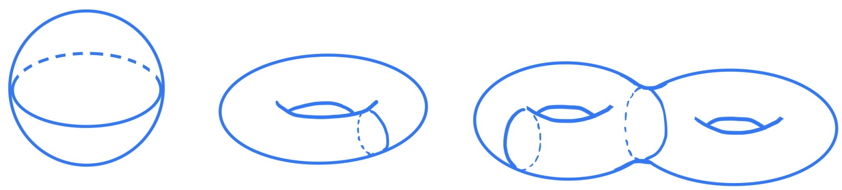

Figure 11. Tori with genera 0, 1, and 2, respectively.

Moreover, the genus of a compact Riemann surface gives us information about its Euler characteristic: a compact Riemann surface of genus g has Euler characteristic $2-2g$. This algebraic topology result, proven in [Ful95, p.244], follows by taking $g=0$ and $g=1$ as base cases and inducting on the genus. Note that we have already shown the $g=0$ case in Example 4.10, since we computed the Euler characteristic of $\mathbb{CP}^1$ to be $2=2-2\cdot 0$, where $g_{\mathbb{CP}^1}=0$. By restating Theorem 1.1 in these terms, we obtain the following corollary.

Corollary 5.1. *Let $\varphi\colon Y\to X$ be a holomorphic map of compact Riemann surfaces with degree d as a branched cover. Then*

$$2g_Y-2=d(2g_X-2)+\sum_y(e_y-1),$$

where the sum is over the branch points of φ, e_y is the ramification index corresponding to each branch point $y\in Y$, and g_X and g_Y are the genera of X and Y, respectively. $\square$

This restatement of the Riemann–Hurwitz formula has a number of implications, one of which is discussed below, relating to our previous discussion of the case of $\mathbb{CP}^1$ in Examples 2.5 and 4.2.

Corollary 5.2. *If X is a compact Riemann surface of genus $g>0$, then there are no nonconstant holomorphic maps $\mathbb{CP}^1\to X$.*

Proof of Corollary 5.2. Suppose to the contrary that φ is a nonconstant holomorphic map $\mathbb{CP}^1\to X$. By Theorem 3.7, φ induces a branched cover, so by Corollary 5.1,

$$2g_{\mathbb{CP}^1}-2=d(2g-2)+\sum_y(e_y-1).$$

As $g_{\mathbb{CP}^1} = 0$, the left-hand side equals $2 \cdot 0 - 2 = -2$. But the right-hand side must be a positive value, as $g > 0$ by assumption, so

$$d(2g-2) + \sum_y (e_y - 1) > 0.$$

This gives us a contradiction, so no such φ can exist. □

This result is particularly interesting, as it reveals that the reverse of Lemma 4.6 does not hold: while, for any connected compact Riemann surface Y, we can have a nonconstant holomorphic map from Y into $\mathbb{CP}^1$, we cannot necessarily have a nonconstant holomorphic map out of $\mathbb{CP}^1$ into Y.

References

[For07] Otto Forster. *Lectures on Riemann Surfaces*. Springer, 2007.

[Ful95] William Fulton. *Algebraic Topology: A First Course*. Springer, 1995.

[GP78] Victor Guillemin and Alan Pollack. *Differential Topology*. AMS Chelsea Publishing, Providence, RI, 1978.

[Hur91] Adolf Hurwitz. Uber Riemann'sche Fläche mit gegebenen Verzweigungspunkten. *Mathematische Annalen*, 1891.

[Oor16] Frans Oort. The Riemann-Hurwitz Formula. *The Legacy of Bernhard Riemann After One Hundred and Fifty Years*, 2016.

[Rie57] Bernard Riemann. Theorie der Abel'schen Functionen. *Journal für die reine und angewandte Mathematik*, 1857.

[Sta18] The Stacks Project Authors. *Stacks Project*. https://stacks.math.columbia.edu, 2018.

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

[Tu07] Loring W. Tu. *An Introduction to Manifolds*. Springer, 2007.

Department of Mathematics, Columbia University
E-mail: mw2706@columbia.edu
https://journals.library.columbia.edu/index.php/cjum