An introduction to Riemann surfaces

Chris Evans

1 Riemann Surfaces as complex 1-manifolds

A surface may be broadly defined as a topological space which locally resembles a disc in the complex plane. The following definition formalises this notion:

Definition 1.1. *A* ***surface*** *S is a connected, second countable, Hausdorff topological space with a family of homeomorphisms $\phi_\alpha : U_\alpha \to D_\alpha$ from domains U_α that form an open cover of S to open subsets D_α in the complex plane.*

The homeomorphisms are called ***charts****, and the family they belong to is called an* ***atlas*** *A on S. We write $A = \{(\phi_\alpha, U_\alpha)\}$.*

If the chart maps onto the open disc $\Delta = \{z \in \mathbb{C} : |z| < 1\}$, we call it a **coordinate disc**, since we can associate all points in the domain with points in Δ, hence providing a coordinate system to use within the domain on our surface.

Such a space is often referred to as a **real 2-manifold**; that is to say it is locally homeomorphic to the Euclidean space of 2 dimensions, $\mathbb{R}^2$, which is of course trivially homeomorphic to $\mathbb{C}$. However, a further addition to the definition is required to transform our surface into a **complex 1-manifold**. Though our charts maps domains of S into subsets of the complex plane, this alone does not guarantee that our surface has a complex structure.

The issue lies in the intersection of the domains of S: If $U_{\alpha_1} \cap U_{\alpha_2} \neq \emptyset$, then $\phi_{\alpha_2} \circ \phi_{\alpha_1}^{-1} : \phi(U_{\alpha_1}) \to \phi(U_{\alpha_2})$ is, by composition, a homeomorphism. Such a map is called a **transition map**. The nature of the transition maps determines the structure of S.

Definition 1.2. *If R is a surface with an atlas A, and if all transition maps determined by the atlas are holomorphic, then R is called a* ***Riemann surface****. A is then known as an* ***analytic atlas****.*

This definition motivates the obvious question of whether all surfaces admit a complex structure. It is in fact impossible to construct a complex structure on a non-orientable surface, such as the Möbius strip, the Klein bottle, or the real projective plane $\mathbb{P}^2$. This is essentially a consequence of holomorphic maps preserving orientation (recall that maps such as $z \mapsto \bar{z}$ are not holomorphic on $\mathbb{C}$ as a simple application of the Cauchy-Riemann equations.)

2 Functions on Riemann Surfaces

With our definition for an atlas on a Riemann Surface in mind, we notice that we potentially could give a surface two different complex structures. For example, by shifting the domains D_α in the plane, we would have two different atlases on R; yet for such a trivial change, we would hope that they were still compatible somehow. This motivates the following definition:

Definition 2.1. *Let R_1 and R_2 be Riemann surfaces with atlases $A_1 = \{(\phi_\alpha, U_\alpha)\}$ and $A_2 = \{(\psi_\beta, V_\beta)\}$ respectively. A function $f : R_1 \to R_2$ is called* ***holomorphic*** *if the composition $\psi_\beta \circ f \circ \phi_\alpha^{-1} : \phi(U_\alpha \cap f^{-1}(V_\beta)) \to \psi(V_\beta)$ is holomorphic for each α and β.*

This is, in some ways, a generalisation of the definition of transition maps. We would hope that we may use such a definition to form an equivalence class on Riemann surfaces.

Indeed this is the case: if R_1 and R_2 are Riemann surfaces with a **biholomorphic** map f between them (that is to say, a holomorphic map with a holomorphic inverse), we say that R_1 and R_2 are **conformally equivalent**.

Proposition 2.2. *Conformal equivalence is an equivalence relation.*

Note that if we have two equivalent atlases on a surface S, then their union is also an atlas on S and belongs to the same equivalence class. By combining all atlases on a surface, we may obtain a **maximal atlas** on a surface.

Conformal equivalence is for Riemann surfaces the means of classification, similar to how homeomorphisms are used to classify topological spaces and diffeomorphisms are used to classify smooth manifolds. Of course, conformal equivalence requires the spaces to be homeomorphic, so it is at least as strong a condition. Indeed, we shall see in a later example that there are important spaces that are homeomorphic but not conformally equivalent.

3 Examples of Riemann Surfaces

At this point we introduce the three most important Riemann surfaces, and illustrate the primary method by which one can construct other more complicated Riemann surfaces.

It is immediately obvious from the definition that the **complex plane** $\mathbb{C}$ is a Riemann surface; the atlas containing the single chart, the identity map, is analytic.

Similarly, the **unit disc**, $\Delta = \{z \in \mathbb{C} : |z| < 1\}$ is also a Riemann surface under a similar atlas. In fact, any open connected subset of the complex plane is a Riemann surface.

Figure 1: Stereographic projection - Identifying A on the sphere and Z in the complex plane

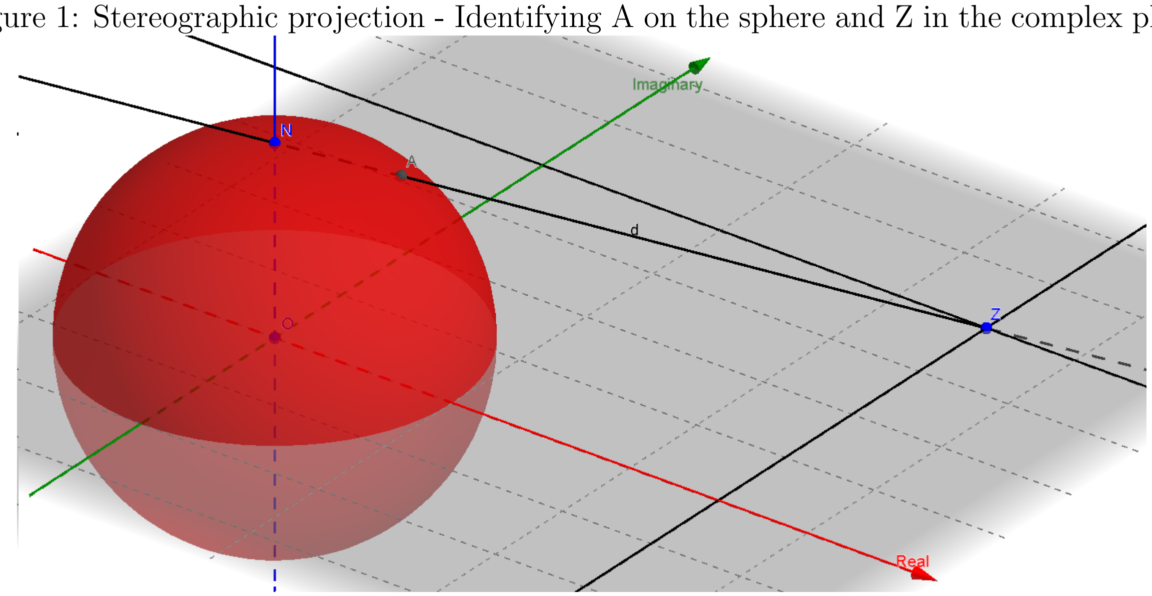

Our final example comes from the one point compactification of $\mathbb{C}$. By adding a point at infinity, we construct a compact topological space $\mathbb{C}_\infty$, known as the **Riemann sphere**. Under an appropriately constructed atlas, this too may be shown to be a Riemann surface.

It may not be immediately obvious why we use the term sphere to describe this abstract space. Figure 1 shows how the standard 2-sphere

$$S^2 = \{(x, y, z) \in \mathbb{R}^3 : x^2 + y^2 + z^2 = 1\}$$

may be identified with the complex plane through **stereographic projection**; the south pole of the sphere takes the value 0 and the north pole takes the value ∞.

General Riemann surfaces may be constructed using the quotient topology, as the following theorem demonstrates:

Theorem 3.1. *Let R be a Riemann surface, and let G be a group of holomorphic automorphisms acting on R. Then R/G under the quotient topology inherits a natural analytic atlas.*

Figure 2: (Torus Construction)

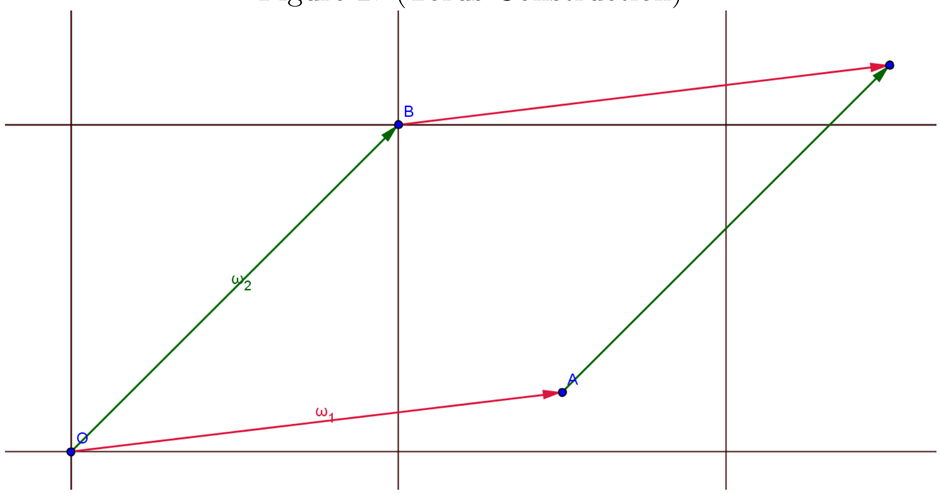

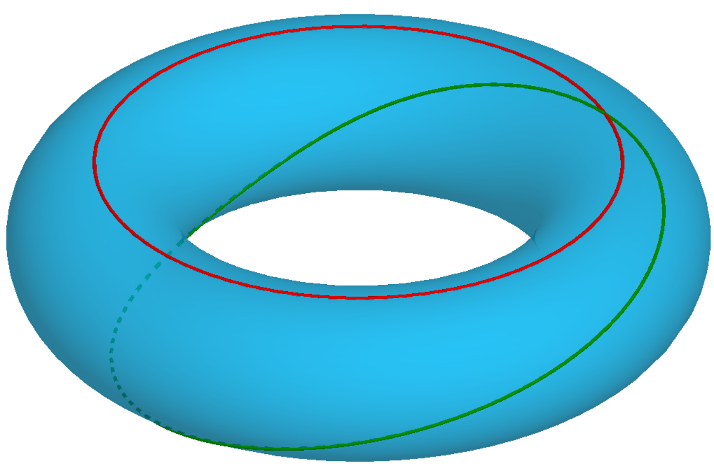

The construction of the **torus** T^2 as a quotient illustrates this idea: Let ω_1 and ω_2 be complex numbers linearly independent over $\mathbb{R}$, $G = \{\gamma(z) = z + m\omega_1 + n\omega_2 : m, n \in \mathbb{Z}\}$. G is then a **lattice group** in $\mathbb{C}$, and the quotient $\mathbb{C}/G$ gives a torus, which by the above theorem inherits an analytic atlas through the quotient topology. This process can be thought of as an identification of sides of a parallelogram, as shown in Figure 2.

One may ask the question of whether any of the surfaces we have introduced so far are homeomorphic or conformally equivalent. Understanding when two surfaces are conformally equivalent is the major focus of this document. Understanding equivalence in a topological sense is somewhat easier. The primary tools used to tackle such questions topologically are **topological invariants**. These are properties invariant under homeomorphisms; therefore, if S_1 and S_2 do not share such a property, they are not homeomorphic. Later we shall construct a more complicated topological invariant, the fundamental group. For now, we show compactness is a topological invariant:

Proposition 3.2. *Compactness is a topological invariant.*

Proof. Suppose A and B are homeomorphic topological spaces, and suppose A is compact. Let $\{M_\beta\}$ be an open cover for B. If $f : B \to A$ is a homeomorphism, $\{f(M_\beta)\}$ is an open cover for A, and therefore has a finite subcover. Mapping this subcover under f^{-1} gives a finite subcover for B. $\square$

Therefore, since $\mathbb{C}_\infty$ is compact, the Riemann sphere is not homeomorphic to either the plane or to any open subset of the complex plane. Of course, they are therefore not conformally equivalent either. However such a line of reasoning will not be able to distinguish between the plane $\mathbb{C}$ and an open subset of the complex plane such as the unit disc Δ. In fact, these spaces are topologically equivalent: this follows from the well known existence of a bijection between the unit interval $[0, 1)$ and the positive reals $[0, \infty)$. However, the two are not conformally equivalent, as the argument below shows:

Recall Liouville's Theorem from complex analysis: Any entire function must be unbounded or constant. Suppose $f : \mathbb{C} \to \Delta$ is a conformal equivalence. Then f is by definition holomorphic and thus is an entire function. Therefore since $|f(z)| < 1$, by Liouville's theorem f must be constant, a contradiction since f is by definition a surjection.

4 Homotopy and the Fundamental Group

Definition 4.1. *A **curve** in a surface S is a continuous map $\gamma : [0, 1] \to S$. A curve is **closed** if $\gamma(0) = \gamma(1)$, and **simple** if $\gamma(x) = \gamma(y) \implies x = y$ for $x \in (0, 1)$ (that is to say, there are no other points of self-intersection).*

From now on we shall assume if unstated that γ denotes a curve from the unit interval.

Suppose we have two curves γ_1 and γ_2 that satisfy $\gamma_1(1) = \gamma_2(0)$. By **concatenation**, we can join these two curves together:

$$\gamma_2 * \gamma_1(t) = \begin{cases} \gamma_1(2t), & 0 \leq t \leq 1/2 \\ \gamma_2(2t-1), & 1/2 \leq t \leq 1 \end{cases}$$

We want to try and characterise the topology of a surface by the nature of the curves that lie within it. As a motivating example, consider two curves in $\mathbb{C}$:

$$\gamma_1 : [0, 1] \to \mathbb{C}, \gamma_1(t) = \exp(i\pi t)$$

$$\gamma_2 : [0,1] \to \mathbb{C}, \gamma_2(t) = \exp(-i\pi t)$$

Both are circular arcs traversed from 1 to -1, and we see that by a process of continuous deformation, we can move one to the other. In that sense, the curves are equivalent. However, if we were to replace $\mathbb{C}$ with $\mathbb{C} - \{0\}$, the same deformation is not possible, since the path of deformation runs through the origin. The following definition makes these notions precise:

Definition 4.2. *(For maps between general spaces) Let X, Y be topological spaces. Let $f, g : X \to Y$ be continuous. A* ***homotopy*** *is a continuous map $H : [0,1] \times X \to Y$ such that for all $x \in X$, $H(0,x) = f(x)$ and $H(1,x) = g(x)$. If such a map exists, f and g are said to be* ***homotopic****.*

(For maps from the unit interval, i.e. curves) Let $\gamma_0 : [0,1] \to S$ and $\gamma_1 : [0,1] \to S$ be curves in a surface S. A ***homotopy*** *(of curves) is a continuous map $H : [0,1] \times [0,1] \to S$ such that $H(0,t) = \gamma_0(t)$ and $H(1,t) = \gamma_1(t)$. If such a map exists, γ_0 and γ_1 are called* ***homotopic****.*

We only need homotopy of curves in order to develop the fundamental group.

Homotopy forms an equivalence relation $\sim$ between curves on a space. We will denote the equivalence class by $[\gamma] = \{\mu : [0,1] \to S : \mu \sim \gamma\}$. We now construct the fundamental group:

Definition 4.3. *Let S be a surface, and let $z \in S$ be the* ***base point****. Let $\Gamma = \{\gamma : [0,1] \to S : \gamma(0) = \gamma(1) = z\}$ (i.e. closed curves with z as the base point). Under concatenation $*$, Γ forms a group. We define the* ***fundamental group of S based at z*** *to be $\pi_1(S,z) = \Gamma / \sim$.*

Simply put, the fundamental group is the group of closed curves based at a point in a space with respect to the homotopy class. In order for the quotient to be well-defined, we should have that $[\gamma] * [\mu] = [\gamma * \mu]$, and indeed this is so. We have that the null curve is the identity element, and the inverse element is simply the same curve traced in the opposite direction.

We would like the fundamental group to be a topological invariant, yet this seems unlikely from the definition since it seems it depends on the base point chosen. The following result shows that this is not a concern for surfaces and other path-connected spaces.

Theorem 4.4. *If S is a path-connected space, then for any $z_0, z_1 \in S$, we have that $\pi_1(S, z_0) = \pi_1(S, z_1)$. Hence, we shall write simply $\pi_1(S)$.*

Proof. Let $\gamma : [0,1] \to S$ satisfy $\gamma(0) = z_0$ and $\gamma(1) = z_1$, and let $\tilde{\gamma}(t) = \gamma(1-t)$. If $[\mu]$ is a homotopy class at z_0, $[\gamma * \mu * \tilde{\gamma}]$ is a homotopy class at z_1, and vice-versa. □

If the fundamental group is trivial (i.e. all curves are homotopic to the base point) then we say that the space is **simply connected**. Of the surfaces we have met so far, $\mathbb{C}_\infty$, $\mathbb{C}$ and Δ are simply connected. In fact, these three surfaces are, up to conformal equivalence, the only simply connected Riemann surfaces. This is the content of the **uniformisation theorem**, an incredibly important result which shall be the main focus of the upcoming sections.

5 Covering Spaces, Lifts, the Universal Cover

Though we shall only consider covering spaces in the context of Riemann surfaces, we introduce them more generally as a topological concept:

Definition 5.1. *Let S be a space. A* ***covering space*** *is a connected surface $\tilde{S}$ with a* ***projection map*** *$p : \tilde{S} \to S$ such that p is continuous and surjective, and for every point $x \in S$, there exists an open neighbourhood U of x with $p^{-1}(U)$ being a union of* ***disjoint*** *open sets in $\tilde{S}$. Such a neighbourhood is called* ***evenly covered****. On each of these disjoint open sets, p is a homeomorphism (so we say that the spaces are* ***locally homeomorphic****).*

Covering surfaces may be thought of as layered, with each layer being projected onto the base surface. This is best illustrated with the example of the torus. From the construction of a torus as a quotient of the plane with a lattice group, we deduce that the torus is covered by the plane. The layers are the quadrilaterals of the lattice.

It is obvious from the definition that the covering space for a surface is also a surface. The following argument shows that complex structure also transfers, and in fact does so uniquely:

Let S be a Riemann surface with a given atlas A, and let it be covered by $p : \tilde{S} \to S$. Let $\phi : U \to D$ be an arbitrary chart in A, and $\tilde{U} \subset \tilde{S}$ be an open set such that p is a homeomorphism when restricted to $\tilde{U}$. Then define an atlas on $\tilde{S}$ with members the charts given by $\phi \circ p : \tilde{U} \to D$. This is obviously an analytic atlas on $\tilde{S}$. Then $\tilde{S}$ is a Riemann surface, and in fact p is holomorphic. This complex structure is also unique: p is biholomorphic locally (since coverings are local homeomorphisms), so the identity map between any two complex structures is biholomorphic also.

Definition 5.2. *Let $p : \tilde{S} \to S$ be a cover, and let $\gamma : [0, 1] \to S$ be a curve in S. Then a* ***lift*** *of γ is a function $\tilde{\gamma} : [0, 1] \to \tilde{S}$ such that $\gamma = p \circ \tilde{\gamma}$*

If, for any curve $\gamma : [0, 1] \to S$ and any $x \in \tilde{S}$ with $p(x) = \gamma(0)$, there exists a lift of γ to $\tilde{S}$ with initial point x, then we say the covering surface is ***regular****.*

We shall assume from now on that all coverings are regular.

Lifts provide us with a way to analyse the path-related properties of a surface (such as the fundamental group or analytic continuation) through consideration of the covering surface. The following lemma characterises the unique placement of lifts in the layers of the cover, and is essential for further discussion of lifts.

Lemma 5.3. *In the set-up of the previous definition, suppose $\tilde{\gamma}_1$ and $\tilde{\gamma}_2$ are both lifts of γ, and suppose $\tilde{\gamma}_1(0) = \tilde{\gamma}_2(0)$. Then $\tilde{\gamma}_1 = \tilde{\gamma}_2$.*

Proof. Suppose $\tilde{\gamma}_1(a) = \tilde{\gamma}_2(a)$ for some $a \in [0, 1]$. Then for some $\epsilon > 0$, $\gamma([a, a+\epsilon))$ lies in an evenly covered neighbourhood, and so $\tilde{\gamma}_1(t) = \tilde{\gamma}_2(t)$ for $t \in [a, a + \epsilon)$. Therefore, if $\tilde{\gamma}_1 \neq \tilde{\gamma}_2$, there is a first point where they are not equal, say $b \in [0, 1]$. However, at such a point, we can find an evenly covered neighbourhood, then we have that $\tilde{\gamma}_1(b)$ and $\tilde{\gamma}_2(b)$ are in disjoint open sets, yet by assumption for $\epsilon > 0$ we have that $\tilde{\gamma}_1(b - \epsilon) = \tilde{\gamma}_2(b - \epsilon)$. So we find that either $\tilde{\gamma}_1$ or $\tilde{\gamma}_2$ is discontinuous, a contradiction. $\square$

Ideally, in order to discuss concepts like the fundamental group in the context of covers, we would have that homotopy be preserved by lifts, and indeed we find this is so:

Lemma 5.4. *Let γ_0 and γ_1 be curves on a surface S that are homotopic by $H : [0,1]\times[0,1] \to S$, and let $p : \tilde{S} \to S$ be a cover. Then if $\tilde{\gamma_0}$ is a lift of γ_0, there is a unique lift $\tilde{H}$ of H with $\tilde{H}(0,t) = \tilde{\gamma_0}$.*

The proof is somewhat more involved than that of the previous lemma though it follows a similar route, and so is omitted. A proof may be found in [9].

With these two lemmas in mind, we are now able to prove the monodromy theorem, an important result that demonstrates the equivalence of homotopy between curves and their lifts.

Theorem 5.5 (Monodromy)**.** *Let γ_0 and γ_1 be curves on a surface S with $\gamma_0(0) = \gamma_1(0)$ and $\gamma_0(1) = \gamma_1(1)$, and let $p : \tilde{S} \to S$ be a cover. Let $\tilde{\gamma_0}$ and $\tilde{\gamma_1}$ be lifts with the same starting point. Then $\gamma_0 \sim \gamma_1$ iff $\tilde{\gamma_0} \sim \tilde{\gamma_1}$; if $\gamma_0 \sim \gamma_1$, then the lifts have the same end point.*

Proof. Projection trivially preserves homotopy, so one way is obvious. If $\gamma_0 \sim \gamma_1$, then we have a homotopy H that may be lifted by the homotopy lemma to a homotopy $\tilde{H}$ between $\tilde{\gamma_0}$ and a $\tilde{\gamma_1}$ by the lifting lemma. The second statement follows trivially. □

Now that we have the concept of lifts, we can consider how the fundamental group interacts with covering spaces. If $p : \tilde{S} \to S$ is a cover, it induces a map p_* from the fundamental group of $\tilde{S}$, $\pi_1(\tilde{S}, x)$, to the fundamental group of S, $\pi_1(S, p(x))$. By the monodromy theorem, this map is well-defined (indeed, a homomorphism) and injective. Therefore, the fundamental group of a covering surface is isomorphic to a subgroup of the fundamental group of the base surface.

We now define an important interaction of the fundamental group and a covering map. The **monodromy action** is given by letting $\pi_1(S, x)$ act on the set $p^{-1}(x)$: If $[\gamma] \in \pi_1(S, x)$, we define the monodromy action as $[\gamma] * y = \tilde{\gamma}_y(1)$, where $\tilde{\gamma}_y(1)$ is the end point of the lift of γ starting at $y \in p^{-1}(x)$. The monodromy action is well-defined as a consequence of the monodromy theorem. It is also transitive (that is to say, for any pair $a, b \in p^{-1}(x)$ there is a unique element in $\pi_1(S, x)$ that takes a to b) as a consequence of $\tilde{S}$ being path connected.

The action has an important consequence: The stabiliser subgroup of $y \in p^{-1}(x)$ under the monodromy action is given by $p_*\pi_1(\tilde{S}, y)$. As a corollary, we obtain that if the base surface S is simply connected, p is a homeomorphism.

If a covering surface is simply connected, we call it the **universal cover**. This alludes to the uniqueness of the simply connected cover and suggests that it is somehow the "largest" cover for our original surface. The following three theorems demonstrate maximality, uniqueness and existence of simply connected surfaces, hence justifying the use of the word "universal". The proofs are omitted; they may be found in a general topological sense in [9] or in the case of Riemann surfaces in [4]. The proofs of maximality and uniqueness are rather straightforward given the material covered above. The proof of existence however is fairly long and technical, and only an outline is presented below. In particular, the construction of the space will prove useful later.

Theorem 5.6 (Maximality). *Let $p : \tilde{S} \to S$ be a cover with $\tilde{S}$ simply connected, and let $p' : \tilde{S}' \to S$ also be a cover. Then there exists a unique covering map $p'' : \tilde{S} \to \tilde{S}'$ such that $p = p' \circ p''$.*

Theorem 5.7 (Uniqueness). *If $p : \tilde{S} \to S$ is a cover with $\tilde{S}$ simply connected, then $\tilde{S}$ is unique up to conformal equivalence.*

Theorem 5.8 (Existence). *Every Riemann surface S has a simply connected cover.*

Proof. This proof involves constructing a space consisting of equivalence classes of paths in the base space with a projection dependent on the end point of the paths, then demonstrating that this space satisfies the requirements: it is topological, Hausdorff, path-connected, simply connected, and justifying that the projection map is indeed a cover.

The space is constructed as follows:

Fix a base point in R, say x_0. We form equivalence classes $[\gamma, x_1]$ of curves in R starting at x_0 and ending at x_1. Two such pairs are equivalent if they end at the same point, and the paths are homotopic. We claim this space of equivalence classes is a simply connected cover for R with the projection $p([\gamma, x_1]) = x_1 \in R$. □

6 The Riemann Mapping Theorem and The Uniformisation Theorem

The uniformisation theorem is a broad classification of all Riemann surfaces. It gives us the three unique simply connected surfaces and, in addition, gives us the means to classify all other Riemann surfaces as quotients of their universal cover. This section shall largely deal with the uniformisation theorem and its corollaries and consequences, but first we tackle an important precursor: The Riemann mapping theorem.

Theorem 6.1 (Riemann Mapping Theorem). *Let S be a simply connected domain in $\mathbb{C}$ such that $S \neq \mathbb{C}$. Then S is conformally equivalent to Δ. In fact, for any $z_0 \in S$, one can find an analytic map $f : S \to \Delta$ such that $f(z_0) = 0$*

Proof. (Sketch[(1)]) We start with $\mathcal{F} = \{f : S \to \Delta, f \text{ analytic, injective, } f(z_0) = 0\}$. It suffices to show there exists a surjective member of $\mathcal{F}$.

First, we show $\mathcal{F}$ is non-empty. We use the fact that a point (arbitrarily 0 by translation) is missing, and that the complement $\mathbb{C}_\infty \backslash S$ is simply connected to construct a branch cut from 0 to ∞. This allows us to define a square root function f analytic on S. A square root function necessarily only hits half of the complex plane; that is to say, if z is in the image, $-z$ is not. By the open mapping theorem, we may take find an open disc about $-z \notin f(S)$ disjoint from $f(S)$, and then apply Möbius transformations composed with f to form g such that the disc under g is Δ. Then $1/g$ maps S into Δ. By a Möbius transformation, we may be assured that $\mathcal{F}$ is non-empty.

---

(1) A particularly clear and straightforward version of this proof may be found in full at `http://people.reed.edu/~jerry/311/rmt.pdf`

The next step involves showing that if a function in $\mathcal{F}$ has the maximal derivative at z_0, then it is a surjection. Supposing $f \in \mathcal{F}$ fails to surject at $w \in \Delta$, we may pursue an argument by contradiction, again exploiting properties of the square and square root functions.

The final step is the application of the Arzelà-Ascoli theorem to $\mathcal{F}$ to find a subsequence converging to the maximal function. This requires first showing that $\mathcal{F}$ is equicontinuous, which may be demonstrated by use of the Cauchy integral formula. Then we may use the Hurwitz theorem to guarantee the injectivity of the limit function and the Weierstrass theorem to guarantee that the limit function is analytic. □

The Riemann mapping theorem is a very powerful result, and immediately hints at the more general uniformisation theorem.

Theorem 6.2 (Uniformisation Theorem)**.** *Let S be a simply connected Riemann surface. Then S is conformally equivalent to either $\mathbb{C}_\infty$, $\mathbb{C}$ or Δ.*

Furthermore, if R is any Riemann surface, then R is conformally equivalent to S/G where S is one of the three simply connected surfaces and G is a discrete subgroup of $\mathrm{Aut}(S)$*. In addition, $G \cong \pi_1(R)$.*

The first half of the theorem is significantly more difficult to prove than the second half, though the second half is very important in its own right since it goes a long way towards classifying all Riemann surfaces, provided we have sufficient understanding of the automorphism group and its subgroups. This is the content of a later chapter. The theorem distinctly classifies all surfaces by their universal cover.

We present now a (very much abridged) version of the standard proof of the first half of the uniformisation theorem. Since the proof isn't particularly illuminating, and the tools developed are not of interest for the rest of the paper, many of the details are skated over or omitted completely. Full proofs using this method may be found in [2] and [6], a briefer proof along the same lines in [4]. The original proof using this approach was due to Koebe.

Proof. Note we have already shown via Liouville's theorem and simple topological arguments that no two of the three spaces are equivalent. We first introduce the central idea in the proof: **Green's functions**. The motivation comes from a somewhat physical viewpoint: it is possible to solve the partial differential equation (Poisson's equation, a generalisation of Laplace's equation)

$$\nabla^2 \phi = f$$

for complex functions ϕ and f under suitable boundary conditions using Green's functions. This is also known as the **Dirichlet problem**, and can be summarised in complex analytic terms as follows: for a domain D, f continuous on ∂D, can we find F harmonic on D such that $F = f$ on ∂D? The boundary conditions are of interest here: We find ourselves able to solve the equation on subsets of the complex place such as Δ or $H = \{z \in \mathbb{C} : \Im(z) > 0\}$ but not on the whole plane $\mathbb{C}$ (or for that matter $\mathbb{C}_\infty$). This suggests that maybe only proper subsets of the complex plane support Green's functions.

We define **subharmonic function**: a continuous (real-valued) function g is subharmonic on a surface S if for every harmonic function u on a domain $U \subset S$, $g = u$ or $g < u$ on U. We construct a **Perron family** of $\mathcal{F}$ subharmonic functions defined by two properties:

(i) if two functions g_1 and g_2 are in $\mathcal{F}$ so is $\max(g_1, g_2)$, and (ii) if h is harmonic on a domain and agrees with $g \in \mathcal{F}$ on the boundary of a domain and on the complement of the domain, then $h \in \mathcal{F}$. Perron families have the key property that the function u defined by $u(p) = \sup\{g(p) : g \in \mathcal{F}\}$ is harmonic or positive infinite everywhere.

We need a special Perron family to define Green's functions on S: take a point $p \in S$ contained in a coordinate disc (ϕ, U) with $\phi(p) = 0$. Define $\mathcal{F}_p$ as the family of subharmonic functions on $S - \{p\}$ such that each has compact support and if $u \in \mathcal{F}_p$ then $v(z) = u(z) + log|\phi(z)|$ is subharmonic on a neighbourhood of p. A Green's function with singularity at p is defined by $g_p(z) = \sup\{u(z) : u \in \mathcal{F}_p\}$, provided the supremum is strictly less than infinity on S.

This lengthy set-up leads to the key proposition[(2)]: If S is a simply connected surface with a Green's function g_p for any $p \in S$[(3)], then there exists a conformal map from S to Δ.

What then remains to classify is the surfaces which do not admit a Green's function. This may be done by defining an alternative to a Green's function with singularities at two points. This may then be used to construct a meromorphic function on S, the image of which may be shown to be either empty or a single point. If empty, the surface is $\mathbb{C}_\infty$. If a single point, we may move the point to infinity by a Möbius transform, and so the surface is $\mathbb{C}$. □

We will refer to each of the three cases $\{\mathbb{C}_\infty, \mathbb{C}, \Delta\}$ as elliptic, parabolic and hyperbolic respectively. Of the surfaces we have met so far, only $\mathbb{C}_\infty$ is covered by $\mathbb{C}_\infty$. $\mathbb{C}$, $\mathbb{C} - 0$ and T^2, the torus, are covered by $\mathbb{C}$. All other surfaces are covered by Δ, so the vast majority of surfaces are of hyperbolic type.

We postpone the second half of the proof pending a discussion of some group theory in the next section.

7 Kleinian Groups And Fuchsian Groups

From the statement of second half of the uniformisation theorem, it is necessary to understand the automorphism structure of a given Riemann surface. We begin with the automorphism groups for the three simply connected surfaces, which are all groups of Möbius transformations. Möbius transforms may be associated with complex matrices by the following correspondence:

$$\left\{ f(z) = \frac{az+b}{cz+d} : a, b, c, d \in \mathbb{C}, ad - bc \neq 0 \right\}$$
$$\updownarrow$$
$$\left\{ \begin{pmatrix} a & b \\ c & d \end{pmatrix} : a, b, c, d \in \mathbb{C}, ad - bc \neq 0 \right\} = GL(2, \mathbb{C})$$

The restriction to $ad - bc = 1$ gives the **special linear group** $SL(2, \mathbb{C})$. Since we obtain the same Möbius transform with $\{-a, -b, -c, -d\}$ as with $\{a, b, c, d\}$, one can quotient to

[(2)]See `http://www.math.mcgill.ca/gantumur/math580f11/downloads/uniformisation.pdf`, p15 for a detailed proof using this method.

[(3)]In fact, if a Green's function exists for one point p in S, it exists for all points in S.

obtain the **projective special linear group** $PSL(2,\mathbb{C}) = SL(2,\mathbb{C})/\{I,-I\}$[(4)]. This is the group of automorphisms of the Riemann sphere C_∞. The automorphism groups for the three simply connected Riemann surfaces are summarised below:

$$\mathrm{Aut}(\mathbb{C}_\infty) \cong PSL(2,\mathbb{C})$$
$$\mathrm{Aut}(\mathbb{C}) = \{f(z) = az + b : a, b \in \mathbb{C}, a \neq 0\} \cong P\Delta(2,\mathbb{C})$$
$$\mathrm{Aut}(\Delta) = \left\{ f(z) = \frac{az+b}{\bar{b}z+\bar{a}} : a, b \in \mathbb{C}, |a|^2 + |b|^2 = 1 \right\} \cong SU(1,1)$$
$$\cong Aut(H) = \left\{ f(z) = \frac{az+b}{cz+d} : a, b, c, d \in \mathbb{R}, ad - bc = 1 \right\} / \{I, -I\} \cong PSL(2,\mathbb{R})$$

Note the last two are automatically isomorphic since due to the uniformisation theorem, Δ is conformally equivalent to $H = \{z \in \mathbb{C} : \Im(z) > 0\}$. We shall in general refer to the automorphism group of Δ as $PSL(2,\mathbb{R})$ rather than $SU(1,1)$.

Definition 7.1. *A **topological group** is a group G with a Hausdorff*[(5)] *topology on G that respects the group structure: that is to say, the group operation and inverses are continuous functions on G.*

Of course, all the groups encountered in this section are topological groups. For example, $GL(2,\mathbb{C})$ may be given a topology by identifying it with $\mathbb{C}^2$. $SL(2,\mathbb{C})$ inherits the subspace topology from $GL(2,\mathbb{C})$, and $PSL(2,\mathbb{C})$ is a topological group under the quotient topology. In fact, the structure of these topological groups can be shown to be that of a manifold. Most important are the discrete subgroups of the automorphism groups (recall from the uniformisation theorem that all Riemann surfaces are given as quotients of a simply connected surface S with a discrete subgroup of $\mathrm{Aut}(S)$).

Definition 7.2. *Let H be a subgroup of G, a topological group. H is said to be a **discrete subgroup** of G if it inherits a discrete topology from G (that is to say, all subsets of H are open under the topology of G).*

In particular, special names are given to the discrete subgroups of $\mathrm{Aut}(\mathbb{C}_\infty)$ and $\mathrm{Aut}(\Delta)$:

Definition 7.3. *A **Kleinian group** is a discrete subgroup of $\mathrm{Aut}(\mathbb{C}_\infty) \cong PSL(2,\mathbb{C})$. A **Fuchsian group** is a discrete subgroup of $\mathrm{Aut}(\Delta) \cong PSL(2,\mathbb{R})$.*

The discrete topology encodes notions of the separation of elements (in this case, functions on the Riemann sphere). A definition that categorises the action of a group on a set in terms of separation of elements is the following:

Definition 7.4. *Let G be a subgroup of $\mathrm{Aut}(S)$. G acts **properly discontinuously** on S if for any compact subset K of S there are finitely many $g \in G$ such that $g(K) \cap K \neq \varnothing$.*

[(4)]In general, the projective subgroup is defined by quotienting out scalar transformations. In fact, $PGL(2,\mathbb{C}) = PSL(2,\mathbb{C})$. This follows from the fact that $\mathbb{C}$ is an algebraically closed field.

[(5)]The requirement that the topology be Hausdorff is included for ease of use.

It is a consequence of the definition that for a subgroup to act properly discontinuously, it must be discrete. For the specific case of $\mathrm{Aut}(H) = PSL(2,\mathbb{R})$, the inverse holds:

Proposition 7.5. *Let G be a subgroup of $\mathrm{Aut}(H) = PSL(2,\mathbb{R})$. G is Fuchsian (discrete) if and only if it acts properly discontinuously on H.*

Proof. Suppose G does not act properly discontinuously on H. Then there is a K, a compact set, and a sequence of distinct maps $g_n \in G$ such that for each n, $g_n(K) \cap K$ is non-empty. Take a pair of points $(x_n, y_n) \in K \times K$ for each n with $g_n(x_n) = y_n$. Since K is compact, so is $K \times K$ and so (potentially after reducing to subsequences) we have that $(x_n, y_n) \to (x, y) \in K \times K$, and since g_n are holomorphic, $g_n(x) \to y$. By considering G acting on Δ now, we see that $\{g_n\}$ is locally bounded and hence converges uniformly on compact sets by Morera's theorem to a holomorphic function g on Δ (and hence to a holomorphic function g on H[(6)]).

To see that $g \in \mathrm{Aut}(H)$, again consider the functions as acting on Δ. The maximum principle implies that either $|g| < 1$ or g is constant with $|g| = 1$. If the former, $g \in \mathrm{Aut}(H)$ since g_n^{-1} also converges uniformly (after subsequences) to a holomorphic function g^{-1} with $|g^{-1}| < 1$ and $g \circ g^{-1} = g^{-1} \circ g = id$. If the latter, the functions converge to a boundary point of H, but by assumption, $g(x) = y \in K \subset H$. So $g \in \mathrm{Aut}(H)$.

Finally, we show this implies G is not Fuchsian. We have that $g_n \to g \in \mathrm{Aut}(H)$, and so $g_n^{-1} \to g^{-1} \in \mathrm{Aut}(H)$. Since g_n are pairwise distinct, $id \neq g_n^{-1} \circ g_{n+1} \in G$ for any n, yet $g_n^{-1} \circ g_{n+1} \to id$, and therefore id is not isolated, which implies G is not Fuchsian. $\square$

This is not however a general property. The Kleinian group below does not act properly discontinuously on $\mathbb{C}_\infty$, since (for example) 0 is fixed by infinitely many transformations:

$$\left\{ f(z) = \frac{az+b}{cz+d} : a,b,c,d \in \mathbb{Z} + i\mathbb{Z}, ad - bc \neq 0 \right\} \subset PGL(2,\mathbb{C}) = \mathrm{Aut}(\mathbb{C}_\infty)$$

We now return to the proof of the second half of the uniformisation theorem. First, note that each Riemann surface R is covered by a unique simply connected surface S by Theorems 5.7 and 5.8. Let $p : S \to R$ be such a covering. We define the **universal covering transformation group** G to be the automorphisms of S that commute with p: that is to say, $g \in G$ if and only if $p \circ g = p$ (such transformations are known as **deck transformations**). Clearly G is a subgroup of $\mathrm{Aut}(S)$. It is an important fact that if $p(x_1) = p(x_2)$ then there exists a unique transformation $g \in G$ such that $g(x_1) = x_2$.

Proposition 7.6 (Uniformisation)**.** *Let R be a Riemann surface, and let $p : S \to R$ be a universal cover. Then R is conformally equivalent to S/G, and G is a subgroup of $\mathrm{Aut}(S)$ acting properly discontinuously on S (and hence G is discrete).*

Proof. For each $z \in R$, the set $p^{-1}(z)$ is an equivalence class in S/G. Therefore the projection map $z \mapsto p^{-1}(z)$ is a well-defined map from R to S/G. It is surjective automatically. It is injective since if $p^{-1}(z_1) = p^{-1}(z_2)$ then $p \circ p^{-1}(z_1) = p \circ p^{-1}(z_2)$ and therefore $z_1 = z_2$. Finally, it is conformal since it is the composition of two conformal maps: p^{-1} and canonical projection (which is conformal under the quotient topology).

[(6)] This is an abuse of notation, the functions g on Δ and g on H are of course conjugate rather than equal.

G acts properly discontinuously on S by an argument similar in nature to Proposition 7.5. $\square$

Two subgroups of $\mathrm{Aut}(S)$ acting discontinuously on S generate conformally equivalent surfaces after quotienting if and only if they are conjugate subgroups.

Theorem 7.7. *Let $p : S \to R$, with S simply connected. Then the universal covering transformation group G of p is isomorphic to $\pi_1(R)$.*

Proof. [7] Let x_0 be any point in R, . Recall the construction of S as a space with points of the form $[\gamma, x_1]$, with γ a curve in R with initial point x_0 and end point x_1. We define an action of G on S as follows: For $[\gamma, x_1] \in S$, and any $[\mu] \in \pi_1(R, x_0)$, define an action $*$ by concatenation of curves:

$$[\mu] * ([\gamma, x_1]) = [\mu * \gamma, x_1]$$

This is a deck transformation of S since it does not alter the projection. Hence, this is indeed an action of G for each $[\mu] \in \pi_1(R, x_0)$. This generates a correspondence: $[\mu] \mapsto [\mu]*$. We show it is an isomorphism of $\pi_1(R, x_0)$ and G.

First, it is a homomorphism by construction. It is injective since the kernel is trivial: if $[\mu]*$ is the identity in G, it is homotopic to the zero path, and hence $[\mu]$ is identity in $\pi_1(R, x_0)$.

Finally, it is surjective by a lifting argument: Let $g \in G$ and let I_0 be the zero curve in R at x_0 (i.e. $I_0 = x_0$). Let $\tilde{\gamma}$ be a curve in S with initial point $[I_0, x_0]$ and final point $g([I_0, x_0])$. Project $\tilde{\gamma}$ to R to get a curve γ, a closed path on R with base point x_0, and hence is in $\pi_1(R, x_0)$. By lifting γ to S, we get that $[\gamma]* = g$. $\square$

8 Classifying Riemann Surfaces of Exceptional Type

We have seen in the previous section how Riemann surfaces are generated as quotients of their fundamental group with one of the three simply connected surfaces. Our principle aims are to classify the type and number (up to conformal equivalence) of Riemann surfaces by analysing group structure. We have already seen the only Riemann surfaces with a trivial fundamental group are $\mathbb{C}_\infty$, $\mathbb{C}$ and H. Since they cover all other Riemann surfaces, this a good starting point.

Of particular importance to the classification are the fixed points of Möbius transformations.

We begin with two important but elementary lemmata that will prove useful for further discussion. The first demonstrates the lack of fixed points in the cover group, the second the number of fixed points of a general member of the cover group.

Lemma 8.1. *Let $p : S \to R$ be a universal cover, with universal cover transformation group G. Then for each $x \in S$, one can find a neighbourhood $U \subset S$ such that $g(U) \cap U \neq \varnothing \implies g = id$.*

[7] Here we follow the line of proof in [7].

Lemma 8.2. *Let $g : \mathbb{C}_\infty \to \mathbb{C}_\infty$ be a Möbius transformation. Then g has either one or two fixed points. For $h \in \mathrm{Aut}(\mathbb{C}_\infty)$, $g = h \circ f \circ h^{-1}$ (i.e. is conjugate to) one of two Möbius transforms: If one fixed point, $f(z) = z + b, b \neq 0$. If two fixed points, $f(z) = az, a \neq 0, 1$.*

Instantly, we can use these two lemmata to classify all the Riemann surfaces of elliptic type:

Theorem 8.3. *The only Riemann surface covered by $\mathbb{C}_\infty$ is $\mathbb{C}_\infty$.*

Proof. Suppose R is covered by $\mathbb{C}_\infty$, with universal cover group G. By Lemma 8.2, each element in G has fixed points. By Lemma 8.1, each element in G is therefore the identity, which implies G is trivial and hence $R = \mathbb{C}_\infty/G = \mathbb{C}_\infty$. □

What about the surfaces covered by $\mathbb{C}$? The situation is more complicated, but still relatively simple:

Theorem 8.4. *The only Riemann surfaces covered by $\mathbb{C}$ are the whole plane itself $\mathbb{C}$, the punctured plane $\mathbb{C} - \{0\}$ and tori given by $T^2 = \mathbb{C}/G$, where G is a lattice group.*

Proof. Suppose R is covered by $\mathbb{C}$, with universal cover group G. Recall elements in G are of the form $g(z) = az + b$, with $a, b \in \mathbb{C}, a \neq 0$. By Lemma 8.1, g can have no fixed points in $\mathbb{C}$ so g of the form (technically, conjugate to) $g(z) = z + b$ by Lemma 8.2. So all elements in G are translations (i.e. they fix infinity when considered over the whole complex plane).

1. If G is trivial, we obviously recover $R = \mathbb{C}$.

2. If all the translations in G are parallel, then since G is discrete, it is generated by a translation $g(z) = z + b$. Then clearly $G = \mathbb{Z}$, and we get $R = \mathbb{C} - \{0\}$. (Think of the quotienting as creating an infinite cylinder, which is then homeomorphic to a sphere with two points removed, say the north and south pole.)

3. If we have two translations in G not parallel, we can find two generating translations, as in the construction of the lattice group in section 3. Then $G = \mathbb{Z} \oplus \mathbb{Z}$ and R is a torus.

It is a simple exercise to show that the first two of these surfaces have $\mathbb{C}$ as their universal cover.

For the torus[(8)], if the universal cover is not $\mathbb{C}$ then, by Theorem 8.3, it must be H. Then $G = \mathbb{Z} \oplus \mathbb{Z}$ is a discrete subgroup of $PSL(2, \mathbb{R})$. Suppose f and g generate G. Since G is abelian, they commute with one and other. Then $f(z) = z + b_1 \implies g(z) = z + b_2$, and $f(z) = a_1 z \implies g(z) = a_2 z$. If the former, we can assume $b_1 = 1$. Then b_2 must be irrational, else $pf = qg$, for $b_2 = p/q$. But if b_2 is irrational, G is certainly not discrete!

If the latter, we must have $a_1, a_2 > 0$ and $a_1^n \neq a_2^m$ for $n, m \neq 0$. Taking logarithms, we get $\log f(z) = \log a_1 + \log z$ and $\log g(z) = \log a_2 + \log z$, with the condition that $n \log a_1 \neq m \log a_2$. This is the same problem as above. We can always find non-discrete pairs generated by $\log f$ and $\log g$, and hence G cannot be discrete either (if we can make points of $\log f$ and $\log g$ arbitrarily close, we can do so for f and g also). □

[(8)] We follow the proof given on pg.210 of [6]

What of the surfaces covered by H? All Riemann surfaces with the exception of the sphere, the plane, the punctured plane and tori are covered by H - the vast majority of Riemann surfaces are of hyperbolic type. They are far more diverse and difficult to classify. The classification theorem for compact surfaces[(9)] gives every compact, connected, (real)-2-manifold as either the sphere, a connected sum of tori T^2 or a connected sum of projective planes $\mathbb{P}^2$. Since projective planes (and sums of) are non-orientable, we cannot give them a complex structure. So the only compact Riemann surfaces are the sphere, the torus, and the n-holed torus, for $n \geq 2$. So we have classified the compact Riemann surfaces to some extent (we don't know yet whether we can impose two conformally distinct structures on topologically equivalent surfaces, so our understanding is far from complete.) For the non-compact hyperbolic surfaces, the situation is significantly more difficult. Though topological classification theorems do exist, they are significantly more complex.

There is one final point to make on this subject. All of the surfaces of elliptic and parabolic type have Abelian fundamental groups. The 2-holed torus, like the majority of hyperbolic surfaces, does not[(10)]. However, there are two remaining surfaces which do have Abelian fundamental groups.

Theorem 8.5. *The punctured disc $\Delta_0 = \Delta - \{0\}$ and the annulus $\Delta_r = \{z \in \mathbb{C} : r < z < 1\}$ are hyperbolic Riemann surfaces, and have fundamental group $\mathbb{Z}$.*

Theorem 8.6. *All other hyperbolic Riemann surfaces have non-Abelian fundamental groups.*

9 The Moduli Problem

In the following sections[(11)], all surfaces shall be assumed to be compact (i.e. n-holed tori). We call n the **genus** of the surface.

A (simplified) statement of the moduli problem is this: when are two conformal structures on topologically equivalent surfaces conformally equivalent? We have already seen in detail how topological equivalence is necessary but not sufficient to guarantee conformal equivalence. To be more precise, suppose R_1 and R_2 are homeomorphic, and have universal cover S, with covering maps p_1 and p_2. Can we find a conformal map $f : R_1 \to R_2$ with a lift $\tilde{f} : S \to S$ such that the diagram below commutes?

$$\begin{tikzcd} S \arrow[r, "\tilde{f}"] \arrow[d, "p_1"'] & S \arrow[d, "p_2"] \\ R_1 \arrow[r, "f"] & R_2 \end{tikzcd}$$

(9) See for example [9], [10] or any other introductory topology book

(10) The fundamental group of the 2-holed torus may be calculated using the Seifert-van Kampen theorem; it is $(\mathbb{Z}*\mathbb{Z}*\mathbb{Z}*\mathbb{Z})/N$, where $*$ is the free product, and N is the normal subgroup generated by $aba^{-1}b^{-1}dcd^{-1}c^{-1}$. Clearly this is pretty far from Abelian!

(11) We present a short introduction to the basic concepts of moduli and Teichmüller theory, as discussed in the opening chapters of [7].

We denote by $\mathcal{M}_g$ the set of all equivalence classes of a (compact) Riemann surface of genus g. We call $\mathcal{M}_g$ the **moduli space** of genus g. It's not obvious at this point that the use of the word space is justified.

Firstly, by the uniqueness of universal covers, all the simply connected surfaces are unique up to conformal equivalence. In particular, the moduli space of the Riemann sphere is trivial.

We consider the remaining surfaces of parabolic type. It is obvious that any two punctured planes are conformally equivalent (consider a Möbius transformation mapping the two punctures to each other whilst preserving infinity.) It is less obvious whether the same is true for tori. The following theorem shows there is not one, but an uncountable number of conformally distinct tori. First however, it helps to normalise the lattice group given in section 3. If $\mathbb{C}/G$ is a torus with lattice group $G = \{m\omega_1 + n\omega_2 : m, n \in \mathbb{Z}\}$, it is conformally equivalent to a torus given by lattice group $\Gamma = \{m + n\tau : m, n \in \mathbb{Z}\}$, with $\tau \in H$.

Theorem 9.1. *Let $\mathbb{C}/G_1$ and $\mathbb{C}/G_2$ be two tori given respectively by the lattice groups $G_1 = \{m + n\tau_1 : m, n \in \mathbb{Z}\}$ and $G_2 = \{m + n\tau_2 : m, n \in \mathbb{Z}\}$. Then they are conformally equivalent if and only if*

$$\tau_2 = \frac{a\tau_1 + b}{c\tau_1 + d}, \quad a, b, c, d \in \mathbb{Z}, ad - bc = 1$$

.

Before the proof, note that the condition above is equivalent to saying that $\tau_2 = g(\tau_1)$, where $g \in PSL(2, \mathbb{Z})$. $PSL(2, \mathbb{Z})$ is a discrete subgroup of $PSL(2, \mathbb{R})$, and hence a Fuchsian group. Therefore, as a corollary to the above theorem we obtain that the moduli space of the torus is given by the quotient:

$$\mathcal{M}_1 \cong H/PSL(2, \mathbb{Z})$$

Proof. The proof of the reverse direction is simple: we can construct a map $\tilde{f} : \mathbb{C} \to \mathbb{C}$ with $\tilde{f}(z) = (c\tau_1 + d)z$. This projects to a biholomorphic map $f : \mathbb{C}/G_2 \to \mathbb{C}/G_1$.

For the forward direction, assume we have such a map $f : \mathbb{C}/G_2 \to \mathbb{C}/G_1$. Using the monodromy theorem, we obtain a lift $\tilde{f}$ in $\mathbb{C}$ that commutes with the projections and f as per the diagram above. $\tilde{f} \in \mathrm{Aut}(\mathbb{C})$, so we may assume $\tilde{f}(z) = \alpha z, \alpha \in \mathbb{C}$ (we may safely drop the constant term). $f(\tau_2) = f(1) = f(0)$, so $\tilde{f}(\tau_2), \tilde{f}(1) \in G_1$. Then:

$$\tilde{f}(\tau_2) = \alpha\tau_2 = a\tau_1 + b,$$
$$\tilde{f}(1) = \alpha = c\tau_1 + d$$

for $a, b, c, d \in \mathbb{Z}$. The result follows after a standard argument to show $ad - bc = 1$. □

It can be shown through a cutting and pasting argument that $\mathcal{M}_1 \cong H/PSL(2, \mathbb{Z})$ is in fact the complex plane. See [8] for details. Though the moduli space has been proven to be relatively simple to obtain in the case of a torus, the situation is bleak for surfaces of higher genus. For this purpose, we introduce the related Teichmüller space.

10 Teichmüller Space

We begin with a discussion of the Teichmümller space for the torus[12].

In the above argument when calculating the moduli space, two conformally equivalent tori (with a conformal map f between them) were generated by the action of a lattice group on $\mathbb{C}$. The map f lifted to a map $\tilde{f}$ on $\mathbb{C}$, which transformed the lattice group of one torus to the lattice group of the other.

The lattice group is generated by two translations: by assumption, these were $z \mapsto z+1$ and $z \mapsto z+\tau$. When projected down to the torus, they become generators for the fundamental group of the torus. Fix a base point x for the fundamental group. The map f induces an isomorphism f_* on the fundamental group of the torus (with base point x), mapping pairs of generators to each other.

For any torus R, take a pair of generators $\Sigma_x = \{[A_1], [B_1]\}$ for the fundamental group $\pi_1(R_\tau, x)$, where $\tau \in H$ gives the lattice generating the fundamental group. Σ_x is called a **marking** on R. The construction of the **Teichmüller space** $\mathcal{T}_1$ for genus 1 goes as follows: consider the space with points $[(R, \Sigma_x)]$ (i.e. equivalence classes of marked tori.) Two markings Σ_x and Σ_y are equivalent if and only if there is a continuous curve C_0 joining x and y on R which induces an isomorphism T_{C_0} between the respective fundamental groups with $[C] \mapsto [C_0^{-1} * C * C_0]$ under T_{C_0}, sending the generators to each other.

Figure 3: A (potential) equivalence construction for markings on a torus

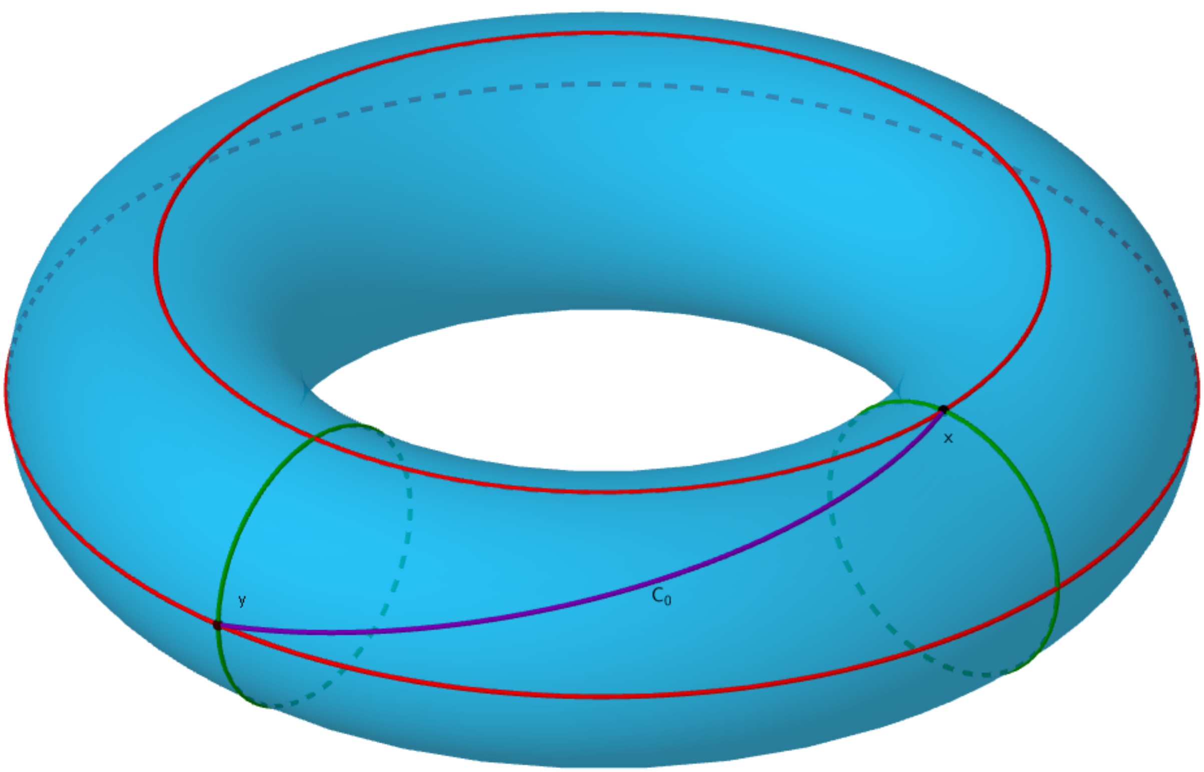

[12] Again, we follow [7].

Two marked tori, (R_1, Σ_{x_1}) and (R_2, Σ_{x_2}) are then equivalent if and only if there is a conformal equivalence $f : R_1 \to R_2$, with the induced map f_* on the fundamental group producing a marking on R_2 that is equivalent to Σ_{x_2}.

Now, suppose we have a compact Riemann surface of arbitrary genus g. The fundamental group for a g-holed torus is given by a system of $2g$ generators: $\{[A_1], [B_1], ..., [A_g], [B_g]\}$. The Teichmüller space $\mathcal{T}_g$ of genus g is constructed in an analogous way to that of the torus; its points are equivalence classes of marked g-holed tori, with the equivalence set-up the same as with the torus.

We can see that the definition Teichmüller space is more restrictive than that of the moduli space: for two surfaces to be equivalent, we require that not only are they conformally equivalent, but also that the conformal equivalence preserves the marking on the torus. Therefore there are points equivalent in the moduli space that are inequivalent in Teichmüller space. The map $\phi : \mathcal{T}_g \to \mathcal{M}_g$, $\phi([R, \Sigma_x]) = [R]$ simply forgets the marking and is an injection. In fact, the moduli space may be given as a quotient of the Teichmüller space by the mapping class group[13], Mod_g :

$$\mathcal{M}_g = \mathcal{T}_g \,/\, \mathrm{Mod}_g$$

To conclude, we sketch how the Teichmüller space $\mathcal{T}_g$ $(g > 1)$ can be associated with a subspace of $\mathbb{R}^{6g-6}$ via **Fricke coordinates**[14].

Let R be a compect surface of genus $g > 1$. Then R is covered by H by a covering $p : H \to R$. Recall from Theorem 7.7 that the universal covering transformation group G of p is isomorphic to $\pi_1(R)$. Since the universal covering transformation group is necessarily a discrete subgroup of $PSL(2, \mathbb{R})$, we shall call it a **Fuchsian model** of R. The Fuchsian model is therefore generated by $2g$ transformations $\{\alpha_1, \beta_1, ..., \alpha_g, \beta_g\}$ corresponding to the generators of the fundamental group of R.

Under a certain normalisation condition on the Fuchsian model (we insist that the transformations have certain fixed points to remove ambiguity caused by equivalence of conjugation in $\mathrm{Aut}(H)$, see [7]), we find that a system of generators is uniquely determined by a point $[R, \Sigma_x]$ in $\mathcal{T}_g$.

We write each transformation for $j = 1, ..., g-1$ in the Fuchsian model as a Möbius transformation:

$$\alpha_j = \frac{a_j z + b_j}{c_j z + d_j}, \;\; a_j, b_j, c_j, d_j \in \mathbb{R}, \; c_j > 0, \; a_j d_j - b_j c_j = 1$$

$$\beta_j = \frac{a'_j z + b'_j}{c'_j z + d'_j}, \;\; a'_j, b'_j, c'_j, d'_j \in \mathbb{R}, \; c'_j > 0, \; a'_j d'_j - b'_j c'_j = 1$$

We then define the Fricke coordinations $\mathcal{F}_g : \mathcal{T}_g \to \mathbb{R}^{6g-6}$ by:

$$\mathcal{F}_g([R, \Sigma_x]) = (a_1, c_1, d_1, a'_1, c'_1, d'_1, ..., a_{g-1}, c_{g-1}, d_{g-1}, a'_{g-1}, c'_{g-1}, d'_{g-1})$$

(13) For a detailed exposition of the relationship between Teichmüller space, moduli space and the mapping class group, see [5].

(14) In fact, Teichmüller space can be given a complex manifold structure of (complex) dimension $3g-3$. For a detailed version of this proof, see [7].

We claim the Fricke coordinates are injective. First, note b_j is determined by the relation $a_jd_j - b_jc_j = 1$, and thus α_j (and similarly β_j) is uniquely determined by a point in $\mathbb{R}^{6g-6}$, for $j = 1, ..., g-1$.

We now have to show that α_g and β_g are determined also. This follows (with some work) from the normalisation conditions and the fundamental group giving the following relation on G:

$$\alpha_1\beta_1\alpha_1^{-1}\beta_1^{-1}...\alpha_g\beta_g\alpha_g^{-1}\beta_g^{-1} = id$$

Thus we may relate the Teichmüller space to points in the much simpler real space $\mathbb{R}^{6g-6}$, the **Fricke space** $\mathcal{F}_g(\mathcal{T}_g)$. Remarkably, the Fricke space is simply connected, and we can give $\mathcal{T}_g$ a topology simply by associating it with $\mathcal{F}_g(\mathcal{T}_g) \subset \mathbb{R}^{6g-6}$.

11 Acknowledgements

Many thanks to Jason Lotay for his suggestions for material to read and directions to take on this project.

References

[1] W. Abikoff, The Real Analytic Theory of Teichmüller Space, 1980: Springer-Verlag.

[2] L. V. Ahlfors, Conformal Invariants: Topics in Geometric Function Theory, 2010: Ams Chelsea Publishing.

[3] L. V. Ahlfors, Complex Analysis: An Introduction to the Theory of Analytic Functions of One Complex Variable, Second Edition, 1966: McGraw-Hill.

[4] A. F. Beardon, A Primer On Riemann Surfaces, 1984: Cambridge University Press.

[5] B. Farb, D. Margalit, A Primer on Mapping Class Groups, 2011: Princeton University Press.

[6] H. M. Farkas, I. Kra, Riemann Surfaces, Second Edition, 1991: Springer-Verlag.

[7] Y. Imayoshi, M. Taniguchi, Introduction to Teichmüller Spaces, 1992: Springer-Verlag.

[8] G. A. Jones, D. Singerman, Complex Functions: An Algebraic and Geometric Viewpoint, 2010: Cambridge University Press.

[9] J. M. Lee, Introduction to Topological Manifolds, Second Edition, 2011: Springer-Verlag.

[10] J. Munkres, Topology, Second Edition, 2000: Pearson.

[11] M. Seppälä, T. Sorvali , Geometry of Riemann Surfaces and Teichmüller Spaces, 1992: Elsevier.

[12] G. Springer, Introduction to Riemann Surfaces, Second Edition, 1999: AMS Chelsea Publishing.