Riemann Surfaces and Teichmüller Theory

Master's thesis, 12th of September 2017

Author:

Toni Ikonen

Adviser:

Kai Rajala

JYVÄSKYLÄN YLIOPISTO
UNIVERSITY OF JYVÄSKYLÄ

Tiivistelmä

Ikonen, Toni
Riemannin pinnat ja Teichmüller-teoriaa
Pro gradu -tutkielma
Matematiikan laitos, Jyväskylän yliopisto, 2017, 78 pages.

Tämän työn päämääränä on määritellä Riemannin pintojen Teichmüller-avaruudet sekä tutkia niiden geometrisia ominaisuuksia. Ensin työssä kehitetään peiteavaruuksien ja toimintojen teoriaa, jota sovelletaan Möbius-kuvauksista koostuviin ryhmiin. Tämän jälkeen kvasikonformaalikuvaukset määritellään Riemannin pinnoille ja niiden yhteyttä yhdesti yhtenäisten Riemannin avaruuksien kvasikonformikuvauksiin tutkitaan. Näitä tietoja sekä yhdesti yhtenäisten Riemannin pintojen uniformisaatiolausetta hyödyntämällä todistetaan yleisten Riemannin pintojen uniformisaatiolause. Tämä tulos liittää pinnat Möbius-kuvauksien toimintoihin yhdesti yhtenäisillä Riemannin pinnoilla.

Yleisten Riemannin pintojen uniformaatioteoreema mahdollistaa työssä käytetyt Teichmüllerin avaruuksien määritelmät. Näille avaruuksille annetaan useampi ekvivalentti määritelmä. Tämän jälkeen Teichmüllerin avaruuksiin määritellään teorian kannalta luonnollinen etäisyysfunktio, joka tekee avaruuksista geodeettisen ja täydellisen. Lisäksi osoitetaan että Riemannin pintojen väliset kvasikonformaalikuvaukset indusoivat surjektiivisen isometrian pintojen Teichmüllerin avaruuksien välille. Lopuksi yhdesti yhtenäisten Riemannin pintojen, punkteerattujen kompaktien Riemannin pintojen sekä topologisten sylintereiden Teichmüller -avaruudet karakterisoidaan. Yhdesti yhtenäisistä pinnoista vain hyperbolisella tasolla osoittautuu olevan epätriviaali Teichmüllerin avaruus. Topologisten sylintereiden tapauksessa havaitaan kolme erilaista Teichmüllerin avaruutta, jotka vastaavat punkteerattua tasoa, punkteerattua kiekkoa ja rengasta.

Avainsanat: toiminto, peiteavaruus, peitekuvaus, nosto, Möbius-kuvaus, Riemannin pinta, kvasikonformikuvaus, Teichmüllerin avaruus, Teichmüllerin metriikka

Abstract

Ikonen, Toni
Riemann surfaces and Teichmüller theory
Master's thesis
Department of Mathematics, University of Jyväskylä, 2017, 78 pages.

The main objective of this work is to develop the necessary tools to define the Teichmüller spaces of Riemann surfaces and study their geometric properties. Firstly, some theory of covering spaces and topological actions will be studied and the results applied to Möbius transformations. Secondly, quasiconformal maps between Riemann surfaces will be defined and they will be characterized using quasiconformal maps between simply-connected Riemann surfaces. These results and the Uniformization Theorem of simply-connected Riemann surfaces will be used to prove a Uniformization Theorem for general Riemann surfaces. Such surfaces will be linked to actions of Möbius transformations on simply-connected Riemann surfaces.

The Uniformization Theorem of Riemann surfaces will be used to define Teichmüller spaces. A couple of equivalent definitions will be introduced. After that a natural distance function is defined on Teichmüller spaces which makes them geodesic and complete. It will be shown that quasiconformal maps between Riemann surfaces induce isometries between their Teichmüller spaces. Finally, the Teichmüller spaces of Riemann surfaces that are either simply-connected, punctured compact Riemann surfaces, or topological cylinders will be characterized. In the simply-connected case, only the hyperbolic plane has a non-trivial Teichmüller space. The topological cylinders have three distinct Teichmüller spaces each of which correspond to exactly one of the following: the once-punctured plane, the once-punctured disk, or annuli.

Keywords: Action, covering space, covering map, lift, Möbius transformation, Riemann surface, quasiconformal map, Teichmüller space, Teichmüller metric

Contents

Introduction

Riemann surface is a topological surface together with a maximal atlas where the coordinate transformations consist of conformal maps between open subsets of the complex plane – such an atlas is called a conformal structure. This is more restrictive compared to smooth manifolds, where the coordinate transformations are only required to be diffeomorphisms. This can be seen from the fact that every Riemann surface is orientable since the Jacobians of conformal maps are strictly positive. A Möbius band is an example of a smooth surface that is non-orientable [15, Chapter 15].

Consider the Uniformization Theorem: Every simply-connected Riemann surface is conformally equivalent to exactly one of the following: the Riemann sphere $\hat{\mathbb{C}}$, the Euclidean plane $\mathbb{C}$, or the hyperbolic plane $\mathbb{H}$. This means that there are three types of simply-connected Riemann surfaces. The following special case of the theorem is known as the Riemann Mapping Theorem: every simply-connected open proper subset of the Euclidean plane is conformally equivalent to $\mathbb{H}$. In particular, the unit disk $\mathbb{D}$ and $\mathbb{H}$ are conformally equivalent. Consider the map $(r, \exp(it)) \mapsto \tan\left(\frac{\pi}{2}r\right)\exp(it)$. It provides an orientation-preserving diffeomorphism from the unit disk $\mathbb{D}$ onto the Euclidean plane $\mathbb{C}$. This means that every simply-connected Riemann surface is diffeomorphic to the Riemann sphere $\hat{\mathbb{C}}$ or the Euclidean plane $\mathbb{C}$. This is an example why being conformally equivalent is not the same as being diffeomorphic.

For this introduction X refers to the simply-connected Riemann surfaces $\hat{\mathbb{C}}$, $\mathbb{C}$, or $\mathbb{H}$, and G refers to a subgroup of conformal automorphisms of X. As a reminder, the conformal automorphisms of X are Möbius transformations mapping X onto itself.

Riemann surfaces that are not simply-connected can be studied by developing a theory of topological actions, covering spaces, and actions of certain types of groups G acting on X. This is the topic of the first two chapters. On the first half of the third chapter, the following characterization is shown: given a Riemann surface M, there exists a unique simply-connected Riemann surface X and a subgroup G of conformal automorphisms of X acting on X such that M is conformally equivalent to the Riemann surface X/G. This follows from basic results of covering spaces of surfaces and the Uniformization Theorem.

The latter half of the third chapter is spent on studying quasiconformal maps and the basic definition is as follows: Given open sets Ω and Ω' in $\mathbb{C}$, a quasiconformal map from Ω onto Ω' is an orientation-preserving homeomorphism satisfying the equation

$$\partial_{\bar{z}}\phi = \mu \partial_z \phi$$

almost everywhere for some measurable function μ with essential supremum strictly less than one. The partial differential equation is called the Beltrami PDE and the coefficient μ is called the Beltrami differential. Quasiconformal maps with $\mu = 0$ are conformal maps and conformal maps are quasiconformal maps with $\mu = 0$. The definition above generalizes naturally to Riemann surfaces using coordinate charts.

It will be shown in the third chapter that if ϕ is a quasiconformal map between Riemann surfaces M and N and if M is conformally equivalent to X/G, then N is conformally equivalent to X/H for some subgroup H of conformal automorphisms of X and ϕ induces a group isomorphism between G and H. This shows that Riemann surfaces come in three distinct families: the ones conformally equivalent to $\hat{\mathbb{C}}/G$, to $\mathbb{C}/G$, or to $\mathbb{H}/G$. These families can be characterized and it is one of the main results of the first three chapters. The first family consists of Riemann surfaces that are conformally equivalent to the Riemann sphere $\hat{\mathbb{C}}$ and the second family consists of those Riemann surfaces conformally equivalent to the Euclidean plane, the (once-)punctured Euclidean plane, or a torus of genus one. The third family, called the hyperbolic Riemann surfaces, has by far the richest structure and most Riemann surfaces are of the form $\mathbb{H}/G$. As an example, the compact hyperbolic Riemann surfaces are orientable compact surfaces, hence they are homeomorphic to a finite connected sum of tori of genus one (Proposition 6.20 of [14] – the classification of orientable compact surfaces).

Consider the relationship between homeomorphic, quasiconformally equivalent, and conformally equivalent Riemann surfaces. Even though the punctured disk, an annulus, and the punctured plane are homeomorphic, the first pair cannot be quasiconformally equivalent to the punctured plane as the first two are hyperbolic and the punctured plane is conformally equivalent to $\mathbb{C}/G$ for some G. This means that two homeomorphic Riemann surfaces are not necessarily quasiconformally equivalent. Furthermore, it turns out that any two annuli are quasiconformally equivalent but they are conformally equivalent if and only if the quotient of their inner and outer radii coincide.

A question arises whether an annulus and a punctured disk can be quasiconformally equivalent. This can be answered using the notion of ideal boundaries of Riemann surfaces. The ideal boundary ∂M of a Riemann surface M is related to the fact that M is conformally equivalent to a quotient X/G, and the boundary provides an extension of the Riemann surface in some sense. This is made rigorous in the latter half of the third chapter. The boundary is compatible with quasiconformal maps in the following sense: a quasiconformal map from a Riemann surface M onto N admits a continuous extension to $M \cup \partial M$, where ∂M is the ideal boundary of M, such that ∂M is mapped homeomorphically onto ∂N. For an annulus the ideal boundary is the topological boundary, i.e. the disjoint union of two circles whereas the ideal boundary of the punctured disk is a circle. This implies that annuli and punctured disk cannot be quasiconformally equivalent even though they are homeomorphic hyperbolic Riemann surfaces.

As a conclusion, if M and N are homeomorphic they are not necessarily quasiconformally equivalent even if they are of the form X/G and X/H. Furthermore, two quasiconformally equivalent Riemann surfaces are not necessarily conformally equivalent. The other direction of these implications does hold in general.

The notions of homeomorphic Riemann surfaces and quasiconformally equivalent Riemann surfaces agree, in a sense, for the following type of surfaces: A Riemann surface M is of type (g, n) if there exist a compact Riemann surface M' of genus g such that M can be conformally embedded into M' such that $M' \setminus M$ has cardinality n – basically M is a compact Riemann surface of genus g with n points removed. The cardinality is allowed to be zero but always finite. Given two Riemann surfaces M and N of types (g, n) and (g', n'), respectively, there exists an orientation-preserving homeomorphism ϕ from M to N if and only if $(g, n) = (g', n')$ and there exists a quasiconformal map ϕ' from M' to N' that restricts to a quasiconformal map between M and N that is homotopic to ϕ. This is discussed in some detail in the last chapter.

The first half of the fourth chapter is spent on studying the deformation space of Riemann surfaces: Given a Riemann surface M, the elements of the deformation space of M are pairs (N, ϕ), where N is a Riemann surface and $\phi\colon M \to N$ is a quasiconformal map. The latter half is spent on studying the Teichmüller space of Riemann surfaces. It is the deformation space modulo a certain type of equivalence relation related to the notion of ideal boundaries. There is a natural distance on Teichmüller spaces, called the Teichmüller distance, that makes it a complete and geodesic metric space. Furthermore, quasiconformally equivalent Riemann surfaces have isometric Teichmüller spaces. This means that the Teichmüller space captures something quasiconformally invariant about a Riemann surface.

If a Riemann surface M is conformally equivalent to $\hat{\mathbb{C}}/G$ or $\mathbb{C}/G$, its Teichmüller space can be interpreted as the space of all possible conformal structures on M. Given a Riemann surface M of type (g, n), every orientation-preserving homeomorphism to another Riemann surface of type (g, n) can be interpreted to be an element of the deformation space of M. Furthermore, it induces an equivalence class to the Teichmüller space. For such Riemann surfaces, the Teichmüller space can be used to characterize all possible homeomorphisms of this type.

The last section of the last chapter is spent on a characterization of the Teichmüller space of Riemann surfaces for a few special cases. A characterization is given for simply-connected Riemann surfaces, the Riemann surfaces homeomorphic to a topological cylinder, and Riemann surfaces of type (g, n).

Chapter 1

Topology

1.1 Group action

Remark 1.1.1*:*
The goal of this section is to give the algebraic definition of group actions, introduce some related terminology and to give the definition of a topological action. The essential result of this section is Proposition 1.1.7.

Definition 1.1.2 (Action)**:**
Let X be a non-empty set and G a group. A (left) action of G on X is a map $X \times G \to X$, denoted by $(x, g) \mapsto g \cdot x$, satisfying $e \cdot x = x$ for the identity element e and $h \cdot (g \cdot x) = (hg) \cdot x$ for every $g, h \in G$ and every $x \in X$. If such an action exists, and the action is not the trivial action $(x, g) \mapsto x$, it is said that the group G acts on X.

The subset $G_x = \{g \in G \mid g \cdot x = x\}$ of G is the stabilizer of $x \in X$. The orbit of $x \in X$ is the subset $G \cdot x = \{g \cdot x \mid g \in G\}$ of X. The set X/G denotes the union $\bigcup_{x \in X} G \cdot x$ and it is called the orbit space of G.

Remark 1.1.3 *(Basic properties of actions):*
Suppose that G acts on X. The following properties are readily verified:

(a) The stabilizer G_x of $x \in X$ is a subgroup of G.

(b) Let $x, y \in X$. Then the orbits $G \cdot x$ and $G \cdot y$ are equal if and only if there exists $h \in G$ such that $h \cdot x = y$.

(c) Furthermore, the orbits $G \cdot x$ and $G \cdot y$ are either equal or disjoint.

(d) If H is a subgroup of G, then H acts on X.

Definition 1.1.4 (Free action)**:**
Let G act on X. Then G acts on X freely if for any $x \in X$, the stabilizer of x is the trivial subgroup.

A topological group $G = (G, \tau)$ is a group G with a topology τ such that the group product of G is continuous. Formally, the map $G \times G \to G$, where $(g, h) \mapsto g \circ h$ is a continuous map in the product topology of $G \times G$.

Definition 1.1.5 (Action of a topological group)**:**
Let X be a topological space and G a topological group. An action of G on X is a continuous action, if the action is continuous in the product topology of $X \times G$.

An action of G on X is a covering space action (or covering action) if it is a continuous action, it acts freely on X, and given $x \in X$, there exists a neighbourhood U_x of X such that $(g \cdot U_x) \cap U_x \neq \emptyset$ only for finitely many $g \in G$.

Remark 1.1.6*:*
It is a straightforward consequence of Remark 1.1.3 that given $x \in X$ and its stabilizer group H, that for any $x \in X$, the map $gH \mapsto g \cdot x$ is a bijective map from the cosets G/H to the orbit $G \cdot x$ of x. In particular, if G acts freely, then G and the orbit of x are bijective.

It is readily seen from this that an action G on X is free if and only if for every $x \in X$, the map $g \mapsto g \cdot x$ is bijective. Note that given $g \in G$, the map $x \mapsto g \cdot x$ is always bijective as its inverse is given by $x \mapsto g^{-1} \cdot x$.

An action is continuous if and only if the maps $x \mapsto g \cdot x$ and $g \mapsto g \cdot x$ are continuous. This means that given a covering space action of G on X, the map $x \mapsto g \cdot x$ is a homeomorphism from X to X for any $g \in G$. In this case, the group G can be identified with a subgroup of homeomorphisms of X onto itself.

Proposition 1.1.7 (Characterization of covering actions)**:**
Let X be a Hausdorff space and suppose that G acts on X continuously.

(a) The action of G on X is a covering space action if and only if for every $x \in X$ there exists a neighbourhood U_x such that $(g \cdot U_x) \cap U_x \neq \emptyset$ is equivalent to $g = e$.

(b) Let $F\colon X/G \to X$ be a right inverse of $\pi(x) := G \cdot x$ and for every $z \in X/G$ let $V_{F(z)}$ be a neighbourhood of $F(z)$ contained in $U_{F(z)}$. Then the sets $V_{g \cdot F(z)} = g \cdot V_{F(z)}$, for the index set $G \times (X/G)$, form an open cover of X with the property
$$V_{g \cdot F(z)} = V_{h \cdot F(z)}$$
if and only if $g = h$. Furthermore, for every $x \in X$ there exists $g \in G$ such that $x = g \cdot F(G \cdot x)$.

Proof:
Part (a): The "if" direction is just a stronger version of the definition. The "only if" direction is a simple construction, which uses the facts that X is Hausdorff, every map of the form $x \mapsto h \cdot x$ is a homeomorphism, and intersection of finite number of open neighbourhoods is still an open neighbourhood.

Part (b): Let $x \in X$ and V_x be an open neighbourhood of x contained in the neighbourhood U_x given by the assumptions. Let $F\colon X/G \to X$ be a right-inverse of $x \mapsto G \cdot x$. For every $z \in X/G$, define $V_{g \cdot F(z)}$ as $g \cdot V_{F(z)}$; since g is a homeomorphism, the set $V_{g \cdot F(z)}$ is an open neighbourhood of $g \cdot F(z)$. Observe that for every $e \neq g \in G$ it follows that

$$V_{g \cdot F(z)} \cap V_{F(z)} = (g \cdot V_{F(z)}) \cap V_{F(z)} \subset (g \cdot U_{F(z)}) \cap U_{F(z)} = \emptyset.$$

Suppose that $h \in G$ such that $(h \cdot (g \cdot V_{F(z)})) \cap (g \cdot V_{F(z)}) \neq \emptyset$. This can be restated equivalently as

$$((g^{-1}hg)V_{F(z)}) \cap V_{F(z)} \neq \emptyset.$$

Since $V_{F(z)} \subset U_{F(z)}$, the definition of $U_{F(z)}$ implies that $g^{-1}hg = e$, therefore $h = e$.

Remark 1.1.3 Part (c) shows that every element $x \in X$ is contained in some neighbourhood $V_{g \cdot F(z)}$, where $(g, z) \in G \times (X/G)$. It follows that they form the desired cover. □

1.2 Covering spaces

Remark 1.2.1*:*
The goal of this section is to construct necessary topological tools to find a link between certain types of conformal automorphism subgroups of $\hat{\mathbb{C}}$, $\mathbb{C}$ and $\mathbb{H}$ with Riemann surfaces – more on this in Chapters 2 and 3. This link can be established by studying the connections of covering spaces and covering space actions.

As a reminder, given a topological space X and a point $x \in X$, the fundamental group $\pi(X, x)$ is the group of homotopy classes of paths $[\phi]$ defined on the interval $[0, 1]$ that start and end at x, where the homotopy classes are defined rel $\{0, 1\}$. The group structure of the fundamental group is introduced in [11] and [14].

This means that two closed paths starting and ending at x are in the same homotopy class, if there exists a homotopy $h_t\colon [0, 1] \to X$ such that $x = h_0(0) = h_t(0)$ and $x = h_0(1) = h_t(1)$ for any $t \in [0, 1]$. A topological space is said to be simply-connected if it is path-connected and its fundamental group is trivial at some point (equivalently, at any point [11, Proposition 1.5]). For any $x \in X$, any continuous map $\phi\colon X \to Y$ induces a homomorphism ϕ_*, where $[\gamma] \mapsto [\phi \circ \gamma]$, between the fundamental groups $\pi(X, x)$ and $\pi(Y, \phi(x))$.

Definition 1.2.2 (Covering space)**:**
Let X and X' be topological spaces, where X' is a path-connected and locally path-connected Hausdorff space. Let $\pi\colon X' \to X$ be a surjective map such that for every $x \in X$ there exists a neighbourhood U_x of x for which

$$\pi^{-1}(U_x) = \coprod_{y \in \pi^{-1}(x)} W_y$$

is a disjoint union of open neighbourhoods W_y of y each of which is mapped homeomorphically to U_x by π.

The map π is called a covering map and the pair (X', π) is called a covering space of X. A covering space (X', π) is a universal cover if X' is simply-connected.

Remark 1.2.3*:*
It is not always required that X' is connected nor Hausdorff; the Hausdorff assumption is not required in [14] nor in [11], and the connectivity of X' is not required in [11]. For the purposes of this work, the added connectivity and Hausdorff assumptions make many of the statements more clear. It should be noted that the path-connectedness assumption on X' is equivalent to assuming that it is connected.

The covering map π is continuous and a local homeomorphism, in particular, it is an open and closed map. Furthermore, a homeomorphism is a special case of a covering map. It is also clear that a composition of a covering map and a homeomorphism is a covering map, but this may not be true for a composition of two covering maps [11, Section 1.3, Exercise 6].

Since X' is connected and locally path-connected, it follows that X is always Hausdorff and path-connected. In fact, if every point of X (or X') has a neighbourhood basis with a topological property that is preserved by homeomorphisms, then every point of X' (or X, respectively) has a neighbourhood basis of the same type. In particular, X is always locally path-connected.

As a reminder, it is said that a topological space X is locally a Banach space E, if for every $x \in X$ there exists a chart (U, ϕ), i.e. an open neighbourhood U of x, an open set V of E, and a homeomorphism onto its image $\phi\colon U \to V$. A topological Banach manifold refers to a topological space X that is locally a Banach space E and Hausdorff. The dimension of a manifold refers to the cardinality of the vector basis of E. If $E = \mathbb{R}^n$ or $E = \mathbb{C}^n$, it is required that M is second countable.

Now it is also clear that a simply-connected covering space X' has a topological Banach manifold structure if and only if X has a topological manifold structure of the same dimension. Notice that if X' is second countable, then X is second countable. However, it is not clear that X' is second countable if X is second countable. If X' is locally $\mathbb{R}^n$ or $\mathbb{C}^n$ and Hausdorff, this follows from a general result known as Poincaré-Volterra theorem [6, p.186].

Every connected and locally simply-connected topological space has a universal cover [14, Theorem 11.43]. This includes connected Banach manifolds, since coordinate balls are contractible. The discussion above shows that any universal cover of a Banach manifold has the structure of a Banach manifold of the same dimension. This relationship will be studied more deeply in the category of Riemann surfaces in Chapter 3.

Lemma 1.2.4 (Covering space from an action)**:**
Suppose that X is a connected and locally connected Hausdorff space with a covering space action by a group G. Then the quotient map $\mu_G\colon X \to X/G$, $x \mapsto G \cdot x$ is a covering map making X a covering space of X/G.

Proof:
Note that by defining $x \sim y$ if $y \in G \cdot x$ gives an equivalence relation; see Remark 1.1.3. Endow X/G with the quotient topology, i.e. the finest topology that makes μ_G continuous. The claim follows readily from Proposition 1.1.7. □

Remark 1.2.5*:*
From now on, the notation μ_G will refer to the canonical covering map $x \mapsto G \cdot x$ introduced in Lemma 1.2.4.

Definition 1.2.6:
Let (X', π) be a covering space of X and $f\colon Y \to X$ a continuous map. If there exists a continuous map $g\colon Y \to X'$ such that $\pi \circ g = f$, the map g is said to be a lift of f along π. The map f is said to be a descension of g along π.

If (Y', μ) is a covering space of Y and $g\colon Y' \to X'$ is a continuous map such that $\pi \circ g = f \circ \mu$, the map g is said to be a lift of f along π and μ. Conversely, the map f is said to be a descension of g along π and μ.

Theorem 1.2.7 (Unique lifting of homotopies)**:**
Let (X', π) be a covering space of X and let Y be a locally connected space. Suppose that $f_t\colon Y \to X$ is a homotopy and $g\colon Y \to X'$ is a lift of f_0. Then there exists a unique lift $F_t\colon Y \to X'$ of f_t such that $g = F_0$.

Proof:
The proof can be found in [14], more specifically Theorem 11.13. It is also proven in [11]. □

Theorem 1.2.8 (Unique lifting theorem)**:**
Let (X', π) be a covering space of X and let Y be a connected and locally path-connected space, and $\phi\colon Y \to X$ a continuous map. Given any $x' \in X'$ and $y \in Y$ such that $\phi(y) = \pi(x') =: x$, the map ϕ has a lift $\phi'\colon Y \to X'$ satisfying $\phi'(y) = x'$ if and only if $\phi_* \pi_1(Y, y) \subset \pi_* \pi_1(X, x')$. Furthermore, if two lifts of ϕ along π agree at a single point of Y, they agree on all of Y.

Proof:
The statements are combined from Propositions 1.33 and 1.34 of [11]. The proofs can also be found there. They are also proven in [14]; see Theorems 11.18 and 11.12. □

Definition 1.2.9 (Covering group)**:**
Let $X' = (X', \pi)$ be a covering space of X. A homeomorphism $f\colon X' \to X'$ is a covering transformation of π if $\pi = \pi \circ f$. The set G of all covering transformations – also called deck transformations or covering automorphisms [14, p.308] – of π is a subgroup of the automorphism group of X'. It is called the covering group of π.

Remark 1.2.10*:*
Let G' be a covering group of the covering map $\mu_G\colon X \to X/G$. Then the map $f\colon G \to G'$, where $f(g)(x) = g \cdot x$, is well-defined and an isomorphism: It is readily seen that f is a homomorphism as the action is well-defined. The map f is injective since the action is free. Surjectivity of f follows from the fact that $\mu_G^{-1}(x) = G \cdot x$ and the uniqueness part of Theorem 1.2.8.

Proposition 1.2.11:
Let (X', π) be a covering space of X. Then every point $x \in X$ has a path-connected neighbourhood U_x and a collection of path-connected neighbourhoods W_y of $y \in \pi^{-1}(x)$ satisfying

$$\pi^{-1}(U_x) = \coprod_{y \in \pi^{-1}(x)} W_y.$$

Let g be an element of the covering group of π and $y \in \pi^{-1}(x)$. Then $g(y) \in \pi^{-1}(x)$ and g is an homeomorphism from W_y onto $W_{g(y)}$ satisfying

$$g|_{W_y} = \left(\pi|_{W_{g(y)}}\right)^{-1} \circ \pi|_{W_y}.$$

If h is an element of the covering group of π, then $W_{g(y)} \cap W_{h(y)} \neq \emptyset$ if and only if $g = h$.

Proof:
Let $x \in X$ and U'_x be a neighbourhood of x given by the definition of a covering map. Since X is locally path-connected, there exists a path-connected neighbourhood $x \in U_x \subset U'_x$. The preimage of U_x consists of disjoint path-connected sets W_y, $y \in \pi^{-1}(x)$ for which

$$\pi^{-1}(U_x) = \coprod_{y \in \pi^{-1}(x)} W_y$$

and every W_y is mapped homeomorphically to U_x by π. Let $y \in \pi^{-1}(x)$ and $g \in G$. Since $\pi \circ g = \pi$, it follows that $g(W_y) \subset \pi^{-1}(U_x)$ and $g(y) \in \pi^{-1}(x)$. Consequently, the set $g(W_y)$ is contained in $W_{g(y)}$ and it is open, because g is a homeomorphism.

The same composition identity shows that $\pi|_{g \cdot W_y} : g(W_y) \to U_x$ is a surjective map. Since $\pi|_{W_{g(y)}} : W_{g(y)} \to U_x$ is a bijective map, it follows that $W_{g(y)} = g(W_y)$. If $W_{g(y)} \cap W_{h(y)} \neq \emptyset$ for some $h \in G$, the uniqueness part of Theorem 1.2.8 implies that $g = h$. □

Theorem 1.2.12 (Covering action)**:**
Let (X', π) be a covering space of X. If the covering group G of π is given the discrete topology, it acts as a covering space action on X.

Proof:
Let G be the covering group of π. Consider the map $X' \times G \to X'$, where $(x, g) \mapsto g(x) =: g \cdot x$. If G is given the discrete topology, the action is continuous. Proposition 1.2.11 and Proposition 1.1.7 show that the action is a covering space action. □

Definition 1.2.13 (Normal covering spaces and transitivity)**:**
Let (X', π) be a covering space of X. The corresponding covering action is said to be transitive, if for any $x' \in X'$ and every $y \in \pi^{-1}(\pi(x'))$, there exists g in the covering group of π such that $g(x') = y$.

Equivalently, for any $x \in X$ and every $y \in \pi^{-1}(x)$, the orbit of y under the covering action is equal to $\pi^{-1}(x)$. If the covering action of π is transitive, the covering space (X', π) of X is said to be normal covering (space) of X.

Remark 1.2.14*:*
Proposition 1.39 of [11] states that given a universal cover (X', π) of X, the covering group of π and the fundamental group of X are isomorphic. This isomorphism will later be used to characterize all the possible Riemann surface structures of a surface that is topologically a torus of genus one.

Remark 1.2.10 implies that every covering space action is transitive. The next result shows the link between normal covering spaces and covering space actions.

Theorem 1.2.15 (Normal covering action)**:**
Let (X', π) be a normal covering space of X with a covering group G. Then the map $f\colon X'/G \to X$ satisfying $\pi = f \circ \mu_G$ is well-defined and a homeomorphism.

Proof:
Transitivity of the covering action combined with Remark 1.2.10 shows that the covering group of μ_G is exactly G. Then it is readily checked that $\mu_G(x') = \mu_G(y')$ if and only if $\pi(x') = \pi(y')$; this requires the transitivity assumption. This shows that f is well-defined and injective.

Surjectivity of f follows as $\pi = f \circ \mu_G$ and π is surjective. Since π and μ_G are local homeomorphisms, it is clear that f is a local homeomorphism. A bijective local homeomorphism is a homeomorphism, hence the claim follows. □

Corollary 1.2.16:
Let (X', π) be a normal covering space of X, G the covering group of π and $F\colon X \to X'$ a right-inverse of π.

Suppose that for every $w \in X$, the point $F(w)$ has a local basis with some property $\mathcal{X}$ preserved by the homeomorphisms $g \in G$. Then there exists an open cover

$$\left\{g \cdot V_{F(w)}\right\}_{(g,w)\in G\times X}$$

of X' with property $\mathcal{X}$ satisfying $V_{g\cdot F(w)} = g \cdot V_{F(w)}$. Moreover, if $g, h \in G$, then

$$(g \cdot V_{F(w)}) \cap (h \cdot V_{F(w)}) \neq \emptyset$$

if and only if $h = g$. Furthermore, $\pi(h \cdot V_{F(w)}) = \pi(V_{F(w)})$ is a neighbourhood of w and

$$\pi^{-1}\left(\pi(V_{F(w)})\right) = \coprod_{g\in G} g \cdot V_{F(w)}, \quad h|_{V_{g\cdot F(w)}} = \left(\pi|_{V_{(hg)\cdot F(w)}}\right)^{-1} \circ \pi|_{V_{g\cdot F(w)}}.$$

Proof:
This is a corollary of Theorem 1.2.15, Proposition 1.1.7 Part (b) and Proposition 1.2.11. □

Proposition 1.2.17:
Let (X', π) and (Y, μ) be covering spaces of X.

(a) Suppose that X' is simply-connected and $x \in X'$. Then given any $y \in \mu^{-1}(\pi(x))$, there exists a unique lift of π along μ satisfying $\phi(x) = y$.

(b) If ϕ is a lift of π along μ, then ϕ is a covering map.

(c) Let $\psi\colon Z' \to Z$ be a covering map and suppose that Z is simply-connected. Then Z' is simply-connected and ψ is a homeomorphism.

In particular, if X is simply-connected, and ϕ is as above, the covering maps π, μ and ϕ are homeomorphisms, and Y is simply-connected. If Y is simply-connected, then ϕ is a homeomorphism and X' is simply-connected.

Proof:
Part (a): Theorem 1.2.8 is always applicable as X' is simply-connected.

Part (b): Consider surjectivity of ϕ. Let $x \in X'$ and $y = \phi(x)$. Fix $z \in Y$ and consider a path θ starting at y and ending at y. The path $\mu \circ \theta$ has a lift γ along π such that $\gamma(0) = x$ (Theorem 1.2.8). It is readily checked that $\phi \circ \gamma$ and θ are lifts of $\pi \circ \gamma$ along μ that agree at $t = 0$, hence they are equal everywhere by uniqueness of lifts. Computing θ at 1 implies the surjectivity of ϕ. The rest of the claim is a straight-forward corollary of Proposition 1.2.11, and the continuity and surjectivity of ϕ.

Part (c): If ψ is injective, it follows that it is a homeomorphism. Furthermore, if Z is simply-connected and ψ is a homeomorphism, then Z' is simply-connected. The rest of the claim is a corollary of the first part. Thus it is sufficient to show that ψ is injective.

Let $z, w \in Z'$ such that $\psi(z) = \psi(w)$. Let γ be a path starting from z and ending at w. Then $\theta = \psi \circ \gamma$ is a closed path and it is homotopic to the constant path $t \mapsto \theta(0)$, where $h_1 = \theta$ and the homotopy h_t fixes the basepoint $\theta(0)$. This means that $t \mapsto h_t(0)$, $t \mapsto h_t(1)$ and $s \mapsto h_1(s)$ are constant paths. As constant paths lift to constant paths and the homotopy h_t lifts to a homotopy H_t between γ and H_1 (Theorem 1.2.7), the injectivity of ψ follows. □

Remark 1.2.18*:*
Suppose that (X', π) is a universal covering of X. Then (X', π) is a normal covering space of X as a corollary of Proposition 1.2.17: let $(Y, \mu) = (X', \pi)$ and consider lifts of π along π.

Definition 1.2.19 (Conjugation)**:**
Let $\psi\colon X \to Y$ be a homeomorphism and let $\phi\colon X \to X$ be a map. Then $F_\psi(\phi) = \psi \circ \phi \circ \psi^{-1}$ is the conjugation of ϕ by ψ.

Remark 1.2.20*:*
The composition of conjugations satisfy $F_{\psi\circ\psi'} = F_\psi \circ F_{\psi'}$ whenever $\psi \circ \psi'$ is well-defined. Furthermore, $F_{\mathrm{id}} = \mathrm{id}$, $F_{\psi^{-1}} = \left(F_\psi\right)^{-1}$ and $F_\psi(\phi^n) = \left(F_\psi(\phi)\right)^n$.

Theorem 1.2.21 (Homeomorphisms between covering spaces)**:**
Let (X', π_X) and (Y', π_Y) be universal covers of X and Y, respectively. Let G_X and G_Y denote the covering groups of π_X and π_Y, respectively.

(a) Let $\phi\colon X \to Y$ be a homeomorphism. Then for every $x \in X$ and $y \in \pi_Y^{-1}(\phi(x))$, there exists a unique map $\psi\colon X' \to Y'$ that is a lift of ϕ along π_X and π_Y, $\psi(x) = y$ and ψ is a homeomorphism. Additionally, the conjugation map

$$F_\psi\colon G_X \to G_Y$$

is a well-defined isomorphism between the covering groups G_X and G_Y.

(b) Suppose that $\psi\colon X' \to Y'$ is a continuous map and $F\colon G_X \to G_Y$ is a group homomorphism for which $F(g) \circ \psi = \psi \circ g$ for every $g \in G_X$. Then there exists a unique continuous map $\phi\colon X \to Y$ that is a descension of ψ along π_X and π_Y.

In particular, if ψ is a homeomorphism and F is an isomorphism, the continuous map ϕ is a homeomorphism and $F = F_\psi$.

(c) Given homotopic homeomorphisms $\phi\colon X \to Y$ and $\phi'\colon X \to Y$ and ψ, a lift of ϕ satisfying (a), there exists a unique ψ' satisfying (a) for ϕ' such that $F_\psi(g) = F_{\psi'}(g)$ for every $g \in G_X$.

Furthermore, the homotopy between ϕ and ϕ' lifts to a homotopy between ψ and ψ'. If the homotopy between ϕ and ϕ' is rel $A \subset X$, the homotopy between ψ and ψ' is rel $\pi_X^{-1}(A)$.

Proof:
Part (a): The maps π_Y and $\phi \circ \pi_X$ are covering maps of Y as ϕ is a homeomorphism. Let $x \in X'$ and $y \in \pi_Y^{-1}(\phi \circ \pi_X(x))$. Then Proposition 1.2.17 Part (a) shows that there exists a unique lift ψ of $\phi \circ \pi_X$ along π_Y mapping x to y. Parts (b) and (c) show that ψ is a homeomorphism; the maps π_Y and $\phi \circ \pi_X$ are covering maps and Y' is simply-connected.

Consider the conjugation map $F_\psi \colon G_X \to G_Y$, $g \mapsto \psi \circ g \circ \psi^{-1}$. It is clear that if it is well-defined, it is a group homomorphism. If $g \in G_X$, then

$$\phi^{-1} \circ \pi_Y \circ (\psi \circ g) = (\phi^{-1} \circ \pi_Y \circ \psi) \circ g = \pi_X \circ g = \pi_X = \phi^{-1} \circ \pi_Y \circ \psi.$$

This implies that $\pi_Y \circ \psi = \pi_Y \circ (\psi \circ g)$, therefore $F_\psi(g) \in G_Y$. This means that F_ψ is well-defined and a homomorphism from G_X to G_Y. Symmetry in the argument shows that $F_{\psi^{-1}} \colon G_Y \to G_X$ is well-defined, a homomorphism, and the inverse of F_ψ.

Part (b): The goal is to construct a continuous map $\phi \colon X \to Y$ just by using ψ and π_Y. Let $y \in \pi^{-1}(x)$ for some $x \in X'$. Since $g \mapsto F(g)$ is a homomorphism between G_X and G_Y, and G_X acts transitively, it is clear that $\pi_Y \circ \psi(y) = \pi_Y \circ \psi(x)$. As $\pi_X(y) = \pi_X(x)$, the map ϕ can be defined as $\pi_Y \circ \psi = \phi \circ \pi_X$. Since π_X is surjective, it follows that ϕ is well-defined and unique. Since ψ is continuous and the maps π_X and π_Y are surjective local homeomorphisms, the map ϕ is continuous.

If ψ is a homeomorphism and F is an isomorphism, it is clear that $F = F_\psi$. The assumptions of the first portion apply for ψ and ψ^{-1}. Let ϕ and ϕ' denote the descensions of ψ and ψ^{-1}, respectively. By applying the uniqueness of the first part to the pairs $(\phi \circ \phi', \mathrm{id}_{Y'})$ and $(\phi' \circ \phi, \mathrm{id}_{X'})$, it follows that $\phi' = \phi^{-1}$.

Part (c): Let $h_t \colon X \to Y$ be a homotopy between two homeomorphisms $\phi = h_0$ and $\phi' = h_1$. Part (a) implies that ϕ lifts to a homeomorphism along π_X and π_Y. Then Theorem 1.2.7 shows that h_t lifts to a unique homotopy $H_t \colon X' \to Y'$ along π_X and π_Y satisfying $H_0 = \psi$ (apply the theorem for $h_t \circ \pi_X$). The map $H_1 := \psi'$ is a lift of ϕ and Part (a) implies that it is a homeomorphism. If the homotopy h_t fixes some set A, the map $h_t \circ \pi_X$ fixes $\pi_X^{-1}(A)$. As it was deduced in the proof of Proposition 1.2.17 Part (c), the homotopy H_t is fixed in the set $\pi_X^{-1}(A)$.

Fix $g \in G_X$ and define $G_t = F_\psi(g) \circ H_t \circ g^{-1}$. Since $F_\psi(g) \in G_Y$, it is clear that $\pi_Y \circ G_t = h_t \circ \pi_X$. Since $G_0 = \psi$, the uniqueness stated in Part (a) shows that $F_\psi(g) \circ \psi' \circ g^{-1} = \psi'$. This can be restated as $F_\psi(g) = F_{\psi'}(g)$. Since this holds for any $g \in G_X$, the claim follows. □

Chapter 2

Algebra

2.1 Automorphism groups of planar domains

Remark 2.1.1*:*
It is known from complex analysis that Möbius transformations are the conformal automorphisms from the Riemann sphere onto itself. In this text, this group is the automorphism group of the Riemann sphere $\hat{\mathbb{C}}$ and it is denoted by $\operatorname{Aut}\left(\hat{\mathbb{C}}\right)$. They are of the form

$$\operatorname{Aut}(\hat{\mathbb{C}}) = \left\{ z \mapsto \frac{az+b}{cz+d} \mid a,\, b,\, c,\, d \in \mathbb{C},\, ad-bc=1 \right\}.$$

The topology of $\operatorname{Aut}\left(\hat{\mathbb{C}}\right)$ is given by the identification of $\operatorname{Aut}\left(\hat{\mathbb{C}}\right)$ with $PSL(2,\mathbb{C}) = SL(2,\mathbb{C})/\{I,-I\}$. The identification is made using the action of $SL(2,\mathbb{C})$ on $\operatorname{Aut}\left(\hat{\mathbb{C}}\right)$ defined by

$$\begin{bmatrix} a & b \\ c & d \end{bmatrix} \mapsto \left(z \mapsto \frac{az+b}{cz+d} \right).$$

Every another subgroup of $\operatorname{Aut}\left(\hat{\mathbb{C}}\right)$ is given the subspace topology as subsets of $\operatorname{Aut}\left(\hat{\mathbb{C}}\right)$; see [13, p. 35] and [2, Section 3.7] for more details.

The conformal automorphisms of the Riemann sphere fixing ∞ are identified with the conformal automorphisms of $\mathbb{C}$, denoted by $\operatorname{Aut}(\mathbb{C})$, and they are characterized as

$$\operatorname{Aut}(\mathbb{C}) = \{ z \mapsto az+b \mid a \in \mathbb{C} \setminus \{0\}\,,\, b \in \mathbb{C} \}\,.$$

Furthermore, in the case of the upper half-plane, it is clear that

$$\operatorname{Aut}(\mathbb{H}) = \left\{ z \mapsto \frac{az+b}{cz+d} \mid a,\, b,\, c,\, d \in \mathbb{R},\, ad-bc=1 \right\}.$$

This can be deduced from the fact that every element of this group must map $\mathbb{R} \cup \{\infty\}$ onto itself and the upper half-plane onto itself, thus the coefficients must be real. Sometimes it is more convenient to study $\mathbb{D}$ instead of $\mathbb{H}$. These are conformally equivalent, where a conformal map from the disk to $\mathbb{H}$ is given by $z \mapsto i\frac{1-z}{1+z}$. The conformal automorphisms of the disk are

$$\mathrm{Aut}(\mathbb{D}) = \left\{ z \mapsto \frac{az+b}{\bar{b}z+\bar{a}} \mid a, b \in \mathbb{C}, |a|^2 - |b|^2 = 1 \right\}.$$

For future reference, given $\alpha \in \mathbb{C} \setminus \{0\}$, let T_α denote the translation $z \mapsto z + \alpha$ and g_α denote the dilation $z \mapsto \alpha z$.

Any Möbius transformation can be identified with its representation in $SL(2, \mathbb{C})$. The trace of a Möbius transformation is the trace of its representative in $SL(2, \mathbb{C})$. The trace is denoted by Tr. The trace squared is independent of the chosen representative.

Definition 2.1.2 (Classification of Möbius transformations)**:**
Let $\mathrm{id} \neq \phi \in \mathrm{Aut}(\hat{\mathbb{C}})$. The Möbius transformation ϕ is

(a) parabolic if there exists $\psi \in \mathrm{Aut}(\hat{\mathbb{C}})$ such that $F_\psi(\phi) = T_\alpha$ for some $\alpha \neq 0$;

(b) elliptic if there exists $\psi \in \mathrm{Aut}(\hat{\mathbb{C}})$ such that $F_\psi(\phi) = z \mapsto \exp(i\theta) z$ for some $\theta \in (0, 2\pi) + 2\pi\mathbb{Z}$;

(c) hyperbolic if there exists $\psi \in \mathrm{Aut}(\hat{\mathbb{C}})$ such that $F_\psi(\phi) = g_\lambda$ for some $1 \neq \lambda > 0$.

If ϕ is hyperbolic or none of the above type, it is loxodromic.

A fixed point $x \in \hat{\mathbb{C}}$ of a Möbius transformation ψ is attracting if given $z \in \hat{\mathbb{C}}$, the limit of $\psi^n(z)$ as $n \to \infty$ is x and repelling if $\psi^n(z) \to x$ when $n \to -\infty$.

Lemma 2.1.3 (Conjugation and fixed points)**:**
Let $\phi, \phi' \colon \hat{\mathbb{C}} \to \hat{\mathbb{C}}$ be Möbius transformations and $L \colon \hat{\mathbb{C}} \to \hat{\mathbb{C}}$ a homeomorphism. If $F_L(\phi) = \phi'$ and $x \in \mathrm{Fix}(\phi)$ is a fixed point of ϕ, then $L(x)$ is a fixed point of ϕ' and

$$L(\mathrm{Fix}(\phi)) = \mathrm{Fix}(\phi').$$

If $x \in \hat{\mathbb{C}}$ is an attracting or repelling fixed point of ϕ, then $L(x)$ is a fixed point of ψ' of the same type.

Furthermore, if $\mathrm{id}_{\hat{\mathbb{C}}} \neq \phi, \phi' \in \mathrm{Aut}(\hat{\mathbb{C}})$ are two Möbius transformations, then ϕ and ϕ' commute if and only if $\phi(\mathrm{Fix}(\phi')) = \mathrm{Fix}(\phi')$ and $\phi'(\mathrm{Fix}(\phi)) = \mathrm{Fix}(\phi)$.

Proof:
The identity $F_L(\phi) = \phi'$ can be restated as $L \circ \phi = \phi' \circ L$ and $L^{-1} \circ \phi' = \phi \circ L^{-1}$. The first equality implies that every fixed point of ϕ is mapped to a fixed point of ϕ' by L. The second equality shows the same for the fixed points of ϕ' and L^{-1}. Since L is bijective, the equality follows. The basic properties of conjugation show that $L \circ \phi^n = (\phi')^n \circ L$, therefore the claim about conjugation preserving the types of fixed points is clear. The characterization of commutative Möbius transformations can be found in [2], see Theorem 4.3.6. □

Lemma 2.1.4 (Classification by Trace)**:**
An element $\mathrm{id} \neq \phi \in \mathrm{Aut}(\hat{\mathbb{C}})$ is parabolic if and only if $\mathrm{Tr}^2(\phi) = 4$, elliptic if and only if $\mathrm{Tr}^2(\phi) \in [0, 4)$, hyperbolic if and only if $\mathrm{Tr}^2(\phi) \in (4, \infty)$ and loxodromic if and only if $\mathrm{Tr}^2(\phi) \in \mathbb{C} \setminus [0, 4]$.

Proof:
This follows from the fact that Tr^2 and the type is invariant under conjugation by Möbius transformation, see [2, Theorem 4.3.1]. The square of the trace for translations, rotations, and dilations is readily computed. This implies the classification result and the details are shown in Theorem 4.3.4 of [2] and also in Section 2.3.3 of Imayoshi and Taniguchi [13]. □

Remark 2.1.5*:*
Let $\mathrm{id} \neq \phi \in \mathrm{Aut}(\hat{\mathbb{C}})$. The fixed point equation $\phi(z) = z$ is equivalent to studying the zeroes of a polynomial of degree one or two, depending on whether ϕ fixes ∞ or not. Then it is clear from the fundamental theorem of algebra that every non-trivial Möbius transformation has at least a single fixed point in $\hat{\mathbb{C}}$, and at most two fixed points in $\hat{\mathbb{C}}$. In particular, a Möbius transformation has three or more fixed points if and only if it is the identity.

By considering the Möbius transformations of $\mathrm{Aut}(\mathbb{C})$, i.e. maps of the form $z \mapsto az + b$, it is clear that such a map does not have a fixed point in $\mathbb{C}$ if and only if $\phi = T_b$ for some $b \in \mathbb{C} \setminus \{0\}$. Furthermore, if $\phi \in \mathrm{Aut}(\mathbb{H})$, then ϕ is either parabolic, elliptic or hyperbolic; the square of the trace of ϕ is real and positive.

If $\mathrm{id} \neq \phi \in \mathrm{Aut}(\mathbb{H})$, then ϕ is parabolic if and only if it has a single fixed point in $\mathbb{R} \cup \{\infty\}$, elliptic if and only if it has two fixed points z_1 and z_2 such that $z_1 \in \mathbb{H}$ and $z_2 = \overline{z_1}$, or hyperbolic if and only if it has two fixed points in $\mathbb{R} \cup \{\infty\}$ [13, Lemma 2.10]. In the parabolic and hyperbolic case the conjugation F_ψ in the Definition 2.1.2 can be done by an element of $\mathrm{Aut}(\mathbb{H})$; see the discussion after Lemma 2.9 of [13]. Beardon has a more extensive discussion on the topic [2].

It is worth noting that a parabolic element of $\mathrm{Aut}(\mathbb{H})$ can be conjugated by an element of $\mathrm{Aut}(\mathbb{H})$ to exactly one of T_1 or T_{-1} and a hyperbolic element of $\mathrm{Aut}(\mathbb{H})$ to exactly one of g_λ or $g_{\lambda^{-1}}$, where $\lambda > 1$. This follows from the fact that there does not exists $h \in \mathrm{Aut}(\mathbb{H})$ that conjugates T_1 to T_{-1} nor g_λ to $g_{\lambda^{-1}}$.

Definition 2.1.6 (Fuchsian group)**:**
A group G is a Fuchsian group if it is a discrete subgroup of $\mathrm{Aut}(\mathbb{H})$.

Remark 2.1.7*:*
An equivalent definition for a Fuchsian group is that every point of $\mathbb{H}$ satisfies the neighbourhood property discussed in Definition 1.1.5 without requiring that the action is free. This is shown in Lemma 2.16 of [13]. Some texts define a Fuchsian group as a discrete subgroup of $\mathrm{Aut}(\mathbb{D})$ [17] [12]. An even more general definition is given in [2], where $\mathbb{D}$ or $\mathbb{H}$ can be replaced by any disk that is mapped onto $\mathbb{H}$ by a Möbius transformation.

The first paragraph of this remark implies that a Fuchsian group defines a covering space action if and only if its elements do not have fixed points in $\mathbb{H}$, i.e. it does not contain any elliptic elements (Remark 2.1.5). Furthermore, given $g \in G$ from a Fuchsian group that defines a covering space action on $\mathbb{H}$, either g or its inverse can be conjugated by an element ϕ of $\mathrm{Aut}(\mathbb{H})$ to T_1 in the parabolic case and g_λ for $\lambda > 1$ in the hyperbolic case.

A non-trivial Fuchsian group acting freely on $\mathbb{H}$ is Abelian if and only if it is cyclic and generated either by parabolic or hyperbolic elements [13, Lemma 2.14].

Proposition 2.1.8 (Non-Abelian Fuchsian groups)**:**
Let G be a Fuchsian group containing no elliptic elements. If $g_1, g_2 \in G$ such that $g_1 \circ g_2 \neq g_2 \circ g_1$, then the following holds:

(a) The element $g_3 = g_1 \circ g_2 \in G$ does not commute with g_1 nor g_2.

(b) The elements g_1, g_2 and g_3 have distinct fixed points and all of them are contained in $\mathbb{R} \cup \{\infty\}$.

(c) There exists a conformal map $\psi \in \mathrm{Aut}(\mathbb{H})$ such that the conformal map $F_\psi(g_1)$ fixes 0, $F_\psi(g_2)$ fixes 1, and $F_\psi(g_3)$ fixes ∞.

Proof:
Part (a) is clear because of Lemma 2.1.3. Consider Part (b): Let $\{g_i\}_{i=1}^3$ be as in Part (a). Since G is a Fuchsian group, the elements g_1, g_2, g_3 are either hyperbolic or parabolic and their fixed points are contained in $\mathbb{R} \cup \{\infty\}$ (Remark 2.1.5). Let $i, j = 1, 2, 3$ with $i \neq j$. If g_i is hyperbolic, then either g_i and g_j share all of their fixed points or all of their fixed points are distinct [13, Lemma 2.20] – this is the hard part of the proof. This also holds if g_i and g_j are parabolic. Lemma 2.1.3 shows that the fixed points of g_i and g_j must be distinct as they do not commute.

Part (c): For $i = 1, 2, 3$, let p_i be a fixed point of g_i. The unique Möbius transformation obtained by solving the cross-ratio $[\psi(z), 0, 1, \infty] = [z, p_1, p_2, p_3]$ is in $\mathrm{Aut}(\mathbb{H})$ as $p_i \in \mathbb{R} \cup \{\infty\}$. Then Lemma 2.1.3 implies the claim. □

2.2 Covering groups

Remark 2.2.1*:*
The point of this section is to characterize the covering groups of $\hat{\mathbb{C}}$, $\mathbb{C}$, and $\mathbb{H}$. The results of the previous section and Chapter 1 give an intermediate result Theorem 2.2.2, but a more useful characterization in the context of this work is Corollary 2.2.4. The latter result will play a major role in Chapters 3 and 4.

Theorem 2.2.2 (Covering groups)**:**
Let $X = \hat{\mathbb{C}}$, $\mathbb{C}$ or $\mathbb{H}$ and let G be a subgroup of $\mathrm{Aut}(X)$.

(a) If $X = \hat{\mathbb{C}}$, the group G is a covering group if and only if it is the trivial group.

(b) If $X = \mathbb{C}$, the group G is a covering group if and only if it is the trivial group, it is generated by a single non-trivial translation T_α, or it is generated by two non-trivial translations T_α, T_β, where α and β are $\mathbb{R}$-linearly independent.

(c) If $X = \mathbb{H}$, the group G is a covering group if and only if it is a Fuchsian group containing no elliptic elements.

If G is Abelian, the group is the trivial group, or a cyclic group generated by a parabolic or a hyperbolic element of $\mathrm{Aut}(\mathbb{H})$.

If G is non-Abelian, there are two possibilities. Either G contains a subgroup H generated by a hyperbolic element with index $[G\colon H]$ equal to two, or every subgroup generated by a hyperbolic element has infinite index and G contains a subgroup H that consists entirely of hyperbolic elements such that H is isomorphic to the free group of two generators.

Proof:
Consider the case $X = \hat{\mathbb{C}}$. The trivial group is the only subgroup of $\mathrm{Aut}(X)$ that acts freely on X; every Möbius transformation has a fixed point in X as discovered in Remark 2.1.5, therefore Remark 1.1.6 gives the result.

If $X = \mathbb{C}$, then Remark 2.1.5 shows that G can be the trivial group or generated by translations. Since $\mathbb{C}$ is an $\mathbb{R}$-linear vector space of dimension two, it follows that G must be generated by one element or at most two $\mathbb{R}$-linearly independent elements: Consider this claim. The group G can be identified with the G-orbit of $\{0\}$, i.e. $G = G \cdot 0$, and since Proposition 1.1.7 Part (a) holds, the set $G \subset \mathbb{C}$ is a discrete subset of $\mathbb{C}$. Then $\mathbb{Z}$ acts naturally on G as $(T_\alpha)^n = T_{n\alpha}$ holds for every translation.

The set $G \cap \overline{B}(0, R)$ is finite for every $R > 0$ by compactness of Euclidean balls and discreteness of G. If this intersection is equal to G for every $R > 0$, it follows that G is the identity group. If this is the case, the claim is done. Otherwise, let $0 \neq \alpha$ be the element closest to the origin for some $R > 0$. The minimality of α implies that if $\gamma \in \mathbb{R} \cdot \alpha$, where $\mathbb{R}\cdot$ refers to the vector product in $\mathbb{C}$, then $\gamma \in \mathbb{Z} \cdot \alpha$. Otherwise it would be possible to find an element $0 \neq \alpha' \in G$ closer to the origin than α.

Now if $G \cap \overline{B}(0, R) = \mathbb{Z} \cdot \alpha$ for every $R > 0$, the claim is done as α is a generator of G. Thus let $R' \geq R$ such that the intersection contains an element not in the subgroup generated by α and let $\beta \in G \setminus \mathbb{Z} \cdot \alpha$ be closest to the origin. Note that $\beta \notin \mathbb{R} \cdot \alpha$, and if $\gamma \in \mathbb{R} \cdot \beta$, then $\gamma \in G$ if and only if $\gamma \in \mathbb{Z} \cdot \beta$.

The minimality of β, the facts $\beta \notin \mathbb{R} \cdot \alpha$ and $|\alpha| \leq |\beta|$, and triangle inequality imply that $\gamma = \lambda_1 \alpha + \lambda_2 \beta \in G$ if and only if $(\lambda_1, \lambda_2) \in \mathbb{Z} \times \mathbb{Z}$. Since α and β form a $\mathbb{R}$-linear vector basis over $\mathbb{C}$, every element $\gamma \in G$ can be represented in the form $\lambda_1 \alpha + \lambda_2 \beta$. The claim follows.

Consider the claim about $X = \mathbb{H}$. The first part follows from Remark 2.1.7 and the fact that elliptic elements have fixed points in $\mathbb{H}$. The non-Abelian characterization result is shown in Parts 3 and 4 of Proposition 3.1.2 in [12]. It is the hardest part of the proof. □

Definition 2.2.3 (Standard covering groups)**:**
Let $G \subset \mathrm{Aut}(X)$ be the covering group of X, where $X = \hat{\mathbb{C}}$, $\mathbb{C}$ or $\mathbb{H}$. Then G is said to be standardized if one of the following holds:

(a) The group G is the trivial group.

(b) If $X = \mathbb{C}$, then G is generated either by a translation T_1 or by two translations T_1, T_t with $t \in \mathbb{H}$.

(c) If $X = \mathbb{H}$ and G is Abelian, it is generated either by T_1 or g_λ for some $\lambda > 1$. If G is non-Abelian, then for every $x = 0, 1, \infty$, there exists an element h_x of G not equal to the identity which fixes x.

Corollary 2.2.4 (Existence of standard covering groups)**:**
Suppose that $G \subset \mathrm{Aut}(X)$ is a covering group of $X = \hat{\mathbb{C}}$, $\mathbb{C}$ or $\mathbb{H}$. Then there exists a Möbius transformation $\psi \in \mathrm{Aut}(X)$ such that $F_\psi(G)$ is a standard covering group.

Proof:
If G is trivial, there is nothing to show. Thus suppose that G is non-trivial. If $X = \mathbb{C}$, then either G is generated by T_α or T_α and T_β, where α and β are $\mathbb{R}$-linearly independent. In the latter case, it can be assumed that $\mathrm{Im}\left(\frac{\beta}{\alpha}\right) > 0$. Then the dilation $\psi := g_{\frac{1}{\alpha}} : \mathbb{C} \to \mathbb{C}$ shows the claim.

If $X = \mathbb{H}$ and G is Abelian, then G is cyclic and generated either by a parabolic or hyperbolic element g of $\mathrm{Aut}(\mathbb{H})$ (Remark 2.1.7). By replacing g by g^{-1}, if need be, there exists $\psi \in \mathrm{Aut}(\mathbb{H})$ such that $F_\psi(g) = T_1$ in the parabolic case and $F_\psi(g) = g_\lambda$ for $\lambda > 1$ in the hyperbolic case (Remark 2.1.5). Then ψ is the desired isomorpism. If G is not Abelian, and if G does not contain elements of the wanted form, the desired isomorphism can be constructed using Proposition 2.1.8. □

Remark 2.2.5:
Let $X = \hat{\mathbb{C}}$, $\mathbb{C}$ or $\mathbb{H}$, and let $G \subset \mathrm{Aut}(X)$ be a covering group acting on X. Consider the action of G on $\hat{\mathbb{C}}$. A point $z \in \hat{\mathbb{C}}$ is a limit point of G, whenever there exists $w \in \hat{\mathbb{C}}$ and a sequence $g_n \in G$ of distinct elements such that $g_n(w) \to z$ in $\hat{\mathbb{C}}$. The limit set of G is the set of all limits points of G, and it is denoted by $L(G)$.

It is readily seen that $g \cdot L(G) = L(G)$ for every $g \in G$. Similarly, whenever $C\colon \hat{\mathbb{C}} \to \hat{\mathbb{C}}$ is a homeomorphism and F_C is an isomorphism between G and $F_C(G) \subset \mathrm{Aut}(X)$, then $z \in L(G)$ implies $C(z) \in L(F_C(G))$ and $L(F_C(G)) = C(L(G))$.

If $X = \hat{\mathbb{C}}$, the limit set is empty since G is the trivial group. If $X = \mathbb{C}$, the limit set is empty if G is the trivial group, otherwise it is ∞, and in fact, given any $z \in \mathbb{C}$, there exists a sequence of distinct elements of G such that $g_n(z) \to \infty$, where $g_n = T_{n\alpha}$ for some translation T_α for $\alpha \neq 0$. If $X = \mathbb{H}$, it is sufficient to consider elements of $w \in \mathbb{H} \cup \partial\mathbb{H}$ as every element of G is reflection symmetric along the extended real line.

Theorem 4A of [16, Chapter 3] shows that $L(G) = \overline{G \cdot z} \cap \partial\mathbb{H}$, whenever $z \in \mathbb{H}$ (Proposition 1.1.7 implies that every $z \in \mathbb{H}$ satisfies property (1) of the quoted theorem). The previous paragraphs show that $L(G) = \overline{G \cdot z} \cap \partial X$ for any $z \in X$, whenever $X = \hat{\mathbb{C}}$, $\mathbb{C}$ or $\mathbb{H}$.

Since G does not contain elliptic elements if it acts in X as a covering space action, Remark 1L and Theorem 3L [16, Chapter 3] imply that the set $\overline{X} \setminus L(G)$ is the largest subset of $\overline{X}$, where G acts as a covering space action. Then the implication of Theorem 4A can actually be strengthened: if $z \in \overline{X} \setminus L(G)$, then $L(G) = \overline{G \cdot z}$. The set $I(G) = \partial X \setminus L(G)$ is called the ideal boundary of G. It is non-empty only if $X = \mathbb{H}$.

Consider $X = \mathbb{H}$. If $\{\mathrm{id}_X\} \neq G$ is Abelian, the limit set consists of the fixed points of the generators, which is readily seen from Corollary 2.2.4. Thus the limit set is a singleton, if the group is generated by a parabolic element, and contains two points, if the group is generated by a hyperbolic element.

If G is non-Abelian, then $X = \mathbb{H}$ and the group G always contains hyperbolic elements by Theorem 2.2.2. The limit set always contains the fixed points of hyperbolic elements. In fact, it is shown in Theorem 4H of [16, Chapter 3] that the limit set is equal to the closure of the fixed points of the hyperbolic elements of G. Theorem 4I [16, Chapter 3] shows that if G contains a parabolic element, the limit set is also equal to the closure of the set of parabolic fixed points of G.

Chapter 3

Quasiconformal maps and Riemann surfaces

3.1 Quasiconformal maps between planar domains

Definition 3.1.1 (Beltrami coefficient)**:**
Let Ω and Ω' be open sets in $\mathbb{C}$. A measurable function $\mu\colon \Omega \to \mathbb{C} \in L^\infty(\Omega)$ is a Beltrami coefficient, if $\|\mu\|_{\infty,\Omega} < 1$. A function $\psi\colon \Omega \to \Omega' \in W^{1,2}_{\text{loc.}}(\Omega)$ is a K-quasiregular map if there exists a Beltrami coefficient μ with $\|\mu\|_{\infty,\Omega} \leq k = \frac{K-1}{K+1}$ such that ψ is a solution to

$$\psi_{\bar{z}} = \mu\psi_z$$

almost everywhere. This partial differential equation is called the Beltrami PDE. The Beltrami coefficient μ is said to be Beltrami differential of ψ and it is denoted by μ_ψ. If ψ is a homeomorphism and the map ψ is K-quasiregular, the map ψ is said to be a K-quasiconformal map.

Remark 3.1.2*:*
The continuity of quasiregular maps follow from Section 5.4.1 of [1]. Theorem 2.5.4 combined with Theorem 3.7.7 of [1] imply that given a homeomorphic K-quasiconformal map, its inverse is a K-quasiconformal map. The same result shows that a locally injective K-quasiregular map is locally K-quasiconformal.

The coefficient $k_\psi := \|\mu\|_{\infty,\Omega}$ defines the smallest possible constant $K_\psi = \frac{1+k_\psi}{1-k_\psi}$ for which ψ is a K-quasiregular. The factor K_ψ is said to be the maximal dilatation of ψ. Furthermore, a map $\psi\colon \Omega \to \Omega'$ is 1-quasiregular if and only if $\mu_\psi = 0$ if and only if it is analytic [1, p.27 and Weyl's Lemma A.6.10].

Lemma 3.1.3 (Chain rule of Beltrami differentials)**:**
Let $f\colon \Omega \to \Omega'$ and $g\colon \Omega' \to \Omega''$ be two quasiregular maps between open subsets of $\mathbb{C}$ with Beltrami differentials μ_f and μ_g, respectively. Then

$$\mu_{g\circ f} = \frac{\mu_f + r_f(\mu_g \circ f)}{1 + r_f\overline{\mu_f}(\mu_g \circ f)}, \text{ where } r_f = \frac{\overline{f_z}}{f_z}.$$

In particular, if g and $h\colon \Omega''' \to \Omega$ are analytic, then $\mu_{g\circ f\circ h} = (\mu_f \circ h) r_h$ implying that $k_{g\circ f\circ h} = k_f$ and $K_{g\circ f\circ h} = K_f$. Furthermore, if f is a quasiregular map and $h\colon \Omega \to \Omega'$ is a quasiconformal map, then

$$\mu_{f\circ h^{-1}} \circ h = \frac{\mu_f - \mu_h}{1 - \overline{\mu_h}\mu_f}\left(\frac{h_z}{|h_z|}\right)^2.$$

Proof:
The first claim is shown in [5, p.6]. The proof does not need the fact that the maps are invertible. The form of $\mu_{g\circ f\circ h}$ whenever g and h are analytic follows from the first result as analytic maps have vanishing Beltrami differentials, see Remark 3.1.2. The last claim follows by applying the composition rule of Beltrami differentials for $h^{-1}\circ h$ and $h \circ h^{-1}$. □

Lemma 3.1.4:
Let f be a quasiconformal, h analytic, and $g = F_f(h)$. Then g is analytic if and only if $\mu_f = (\mu_f \circ h)\, r_h$.

If $f\colon \Omega \to \Omega'$ and $g\colon \Omega' \to \Omega''$ are K- and K'-quasiregular maps, then $g \circ f$ is a KK'-quasiregular map. Additionally, if $f\colon \Omega \to \Omega'$ is a quasiconformal map, then $K_f = K_{f^{-1}}$.

Proof:
The definition of g can be restated as $g \circ f = f \circ h$. Since $g = (f \circ h) \circ f^{-1}$, Lemma 3.1.3 implies that $\mu_g = 0$ if and only if $\mu_{f\circ h} = \mu_f$. Now $\mu_{f\circ h} = (\mu_f \circ h)\, r_h$, therefore the claim about g follows.

The latter part is shown in Theorem 1.4.1 [5] – the proof holds for quasiregular maps as well. Lemma 3.1.3 implies that $k_{f^{-1}} = k_f$, hence $K_{f^{-1}} = K_f$. □

3.2 Quasiconformal maps between Riemann surfaces

Definition 3.2.1 (Riemann surface)**:**
A Riemann surface is a connected topological surface of complex dimension one with a conformal structure. A conformal structure refers to a maximal atlas, where the coordinate transformations (also called transition maps) are conformal maps between open subsets of the complex plane $\mathbb{C}$.

Remark 3.2.2*:*
As a consequence of Zorn's lemma, every atlas with conformal coordinate transformations induces a conformal structure. Every Riemann surface is necessarily orientable since conformal maps of the plane are sense-preserving. Furthermore, every orientable C^2-surface smoothly embeddable into $\mathbb{R}^3$ can be made into a Riemann surface [17, p.134].

It is also clear that every open subset of $\hat{\mathbb{C}}$ has a natural Riemann surface structure. The natural structure makes the inclusion to $\hat{\mathbb{C}}$ an analytic embedding. Every open subset Ω of $\mathbb{C}$ will be represented using the chart $(\Omega, \mathrm{id}_\Omega)$. The Riemann sphere $\hat{\mathbb{C}}$ has the standard charts $(\mathbb{C}, \mathrm{id}_\mathbb{C})$ and $\left(\hat{\mathbb{C}} \setminus \{0\}, \frac{1}{z}\right)$.

Remark 3.2.3*:*
Let M be a Riemann surface and μ a measurable $\mathbb{C}$-antilinear endomorphism of the tangent bundle TM. The set of all such maps is denoted by $L^\infty_*(TM)$. The definition means that for almost every $z \in M$, there exists an $\mathbb{C}$-antilinear map $\mu_z \colon T_zM \to T_zM$. By identifying T_zM with $\mathbb{C}$, the map μ_z is of the form $z \mapsto \lambda\bar{z}$ for some $\lambda \in \mathbb{C}$.

Given an atlas $\{(U_i, \phi_i)\}_{i\in I}$ of M, μ can be represented by measurable functions $\mu_i \in L^\infty(\phi_i(U_i))$, where the change of coordinate from μ_i to μ_j satisfies

$$\mu_i = \mu_j \circ \left(\phi_j \circ \phi_i^{-1}\right) r_{\phi_j \circ \phi_i^{-1}}.$$

The notation corresponds to Lemma 3.1.3, where $r_g = \frac{\overline{g_z}}{g_z}$. Then the norm

$$\|\mu\|_\infty = \sup_{i\in I} \|\mu_i\|_{\Omega_i,\infty}$$

is well-defined and Lemma 3.1.3 implies that it is independent of the used atlas. If there exists $0 \le k < 1$ such that $\|\mu\|_\infty \le k$, then μ is said to be a k-Beltrami differential of M. The collection of all such elements is denoted by $B(M)$. Further reading on the topic can be found in [12], [5], [13], and [17].

Definition 3.2.4:
Let M and N be two Riemann surfaces and $f\colon M \to N$ a map. Then f is K-quasiregular, if for every $p \in M$, there exist charts (U_p, ϕ) and $\left(W_{f(p)}, \psi\right)$ for which $f(U_p) \subset W_{f(p)}$ and $\psi \circ f \circ \phi^{-1}\colon \phi(U_p) \to \psi(W_{f(p)})$ is a K-quasiregular map. The map $\psi \circ f \circ \phi^{-1}$ is said to be a coordinate representation of f.

If $f\colon M \to N$ is bijective and K-quasiregular, it is K-quasiconformal. The map f is analytic if its coordinate representations are analytic, and conformal if it is bijective and its coordinate representations are analytic.

Remark 3.2.5*:*
Let $\psi\colon M \to N$ be a K-quasiregular map between Riemann surfaces and fix atlases $\{(U_i, \phi_i)\}_{i\in I}$ for M and $\{(V_i, \theta_i)\}_{i\in I}$ for N with $\psi(U_i) \subset V_i$. The coordinate representations $\theta_i \circ \psi \circ \phi_i^{-1}$ define Beltrami differentials $\mu_i \in L^\infty(\phi_i(U_i))$.

Since composing from the right by a conformal map does not change a Beltrami differential (Lemma 3.1.3), it follows that μ_i are independent of the atlas chosen for N. By composing from the right by coordinate transformations of M shows that μ_i transforms in a correct way to define a Beltrami differential on M as a consequence of Lemma 3.1.3. This means that every K-quasiregular map defines a k-Beltrami differential of M. It will later be shown that every k-Beltrami differential on M is a Beltrami differential of some K-quasiconformal map, see Theorem 4.2.5.

The definitions here agree with the definitions given in the planar case (Remark 3.1.2). In the planar case a map is 1-quasiregular if and only if it is analytic and this holds here as well. It immediately follows from the planar case that a locally injective K-quasiregular map between Riemann surfaces is locally K-quasiconformal, and the inverse of a homeomorphism that is K-quasiregular is K-quasiregular. The discussion above motivates the next result.

Lemma 3.2.6:
Let $f\colon M \to N$ be a K-quasiregular map. Then there exists a k-Beltrami differential such that $\partial_{\bar{z}} f = \mu\,(\partial_z f)$ almost everywhere. This Beltrami differential is denoted by μ_f.

Proposition 3.2.7:
The space $(L^\infty_*(TM), \|\ldots\|_\infty)$ is a Banach space. In particular, given a Cauchy sequence of k-Beltrami differentials, the limit is a k-Beltrami differential.

Proof:
Given an element $\mu \in L^\infty_*(TM)$, it can represented locally as $\mu_i \in L^\infty(\Omega_i)$. Thus, it is clear that $\|\ldots\|_\infty$ is a norm. Completeness follows from the property that convergence in $L^\infty(\Omega_i)$-norm implies convergence pointwise almost everywhere. Then the limits of $L^\infty_*(TM)$ transform in the desired way to define an element of $L^\infty_*(TM)$. □

Definition 3.2.8 (Pullback)**:**
If $\pi\colon M \to N$ is analytic map, and μ is a Beltrami differential on N, the pullback Beltrami differential $\pi^*\mu$ on M is defined as follows: Let $\{(U_i, \phi_i)\}_{i\in I}$ be an atlas of M. Then there exists an atlas $\{(W_i, \psi_i)\}_{i\in I}$ for which $f(U_i) \subset W_i$.

Let $\tilde{f}_i = \psi_i \circ f \circ \phi_i^{-1}$ and μ_i be the coordinate representative of f and μ in W_i. Then the coordinate representation $(f^*\mu)_i$ is defined by the formula $(f^*\mu)_i = (\mu_i \circ \tilde{f}_i) r_{\tilde{f}_i}$.

Remark 3.2.9*:*
It is readily checked using Lemma 3.1.3 that the definition above defines a Beltrami differential on M. In particular, the composition rule $r_{g\circ h} = (r_g \circ h) r_h$ holds if g and h are analytic, where $r_g = \frac{\overline{g_z}}{g_z}$. Furthermore, if f and g are analytic maps and $g \circ f$ is well-defined, then $(g \circ f)^* = f^* \circ g^*$, which is readily verified. A pullback formula for quasiconformal maps is discussed in [12].

3.3 Standard covering spaces

Definition 3.3.1 (Analytic covering space)**:**
If X and M are Riemann surfaces and (X, π) is a covering space of M such that π is analytic, the covering space (X, π) is said to be an analytic covering space of M.

Remark 3.3.2*:*
In the category of Riemann surfaces it is sufficient to require that π is a covering map and analytic to get that every point $x \in M$ has a neighbourhood U_x such that its preimage is a disjoint union of open sets each of which is mapped to U_x conformally by π since a locally injective analytic map is locally conformal (Remark 3.2.5).

However, in the category of smooth manifolds (and Riemannian manifolds), it is usual to require directly that the preimages of U_x are mapped diffeomorphically to U_x by π [7, Definition 1.83]. This comes down to the fact that not every smooth locally injective map is a local diffeomorphism.

The proof of Theorem 3.3.3 goes through even if the category of Riemann surfaces is switched to category of smooth manifolds. And actually, analogous result holds in the category of Riemannian manifolds; see [7, Theorem 1.95] for the smooth and [7, Proposition 2.28] for the Riemannian case.

Theorem 3.3.3 (Conformal structure of covering spaces)**:**
Suppose that (X, π) is a covering space of M. Then the following holds:

(a) If (X, π) is an analytic covering space of M, the covering transformations of π are conformal.

(b) If M is a Riemann surface, the covering space (X, π) has a unique conformal structure making it an analytic covering space of M.

(c) Suppose that X is a Riemann surface, (X, π) is a normal covering space of M, and its covering transformations are conformal mappings. Then M has a unique conformal structure making (X, π) an analytic covering space of M.

(d) If (X, π) is a normal analytic covering space of M, the covering map $\mu_G\colon X \to X/G$ defined by $x \mapsto G \cdot x$ induces a unique conformal structure making X/G and M conformally equivalent. The map $f\colon X/G \to M$ satisfying the commutative property $f \circ \mu_G = \pi$ is a conformal map.

Proof:
It was already noted in Remark 1.2.3 that X has a topological surface structure if and only if M has one. Thus it is sufficient to focus on constructing atlases in (b) and (c) inducing the desired conformal structures. Furthermore, the uniqueness in Parts (c) and (d) follows from the uniqueness shown in Part (b).

Part (a): This is a consequence of Proposition 1.2.11, since every element $g \in G$ can be locally represented as a composition of two analytic maps. This shows that g is analytic. This argument holds for g^{-1} as well, hence g is actually conformal.

Part (b): Let $\{(U_x, \phi_x)\}_{x \in M}$ be an atlas of M such that every U_x is connected and $\pi^{-1}(U_x) = \coprod_{w \in \pi^{-1}(x)} W_{x,w}$ and π maps every $W_{x,w}$ homeomorphically onto U_x. Then

$$\left\{ \left(W_{x,w}, \phi_x \circ \pi|_{W_{x,w}} \right) \right\}$$

forms an atlas of X inducing a conformal structure for X. The only unclear part is whether the coordinate transformations are conformal maps.

Suppose that $W_{x,w} \cap W_{x',w'} \neq \emptyset$. The restriction of π to $W_{x,w} \cap W_{x',w'}$ is a homeomorphism onto its image. In this case the image is $U_x \cap U_{x'}$. In fact, if $w \in \phi_x \circ \pi|_{W_{x',w'}} (W_{x,w} \cap W_{x',w'})$, then

$$\left(\phi_{x'} \circ \pi|_{W_{x,w}} \right) \circ \left(\phi_x \circ \pi|_{W_{x,w}} \right)^{-1} (w) = \phi_{x'} \circ \phi_x^{-1}(w).$$

It follows that the coordinate transformations of X are conformal. Note that every point of X has charts associated with it giving a coordinate representation of π that is the identity implying that π is an analytic map.

Suppose that X is a Riemann surface such that π is analytic. Let (U_x, ϕ_x) and $W_{x,w}$ be as in the beginning of the claim. Then $W_{x,w}$ maps π is conformally to U_x. Let (W, ψ) be a chart of X with $W_{x,w} \cap W \neq \emptyset$. Then given $z \in \psi(W_{x,w} \cap W)$, it follows that

$$\left(\phi_x \circ \pi|_{W_{x,w}}\right) \circ \psi^{-1}(z) = \phi_x \circ \pi \circ \psi^{-1}(z).$$

The right-hand side is a conformal map, since π is locally conformal. This shows that $\phi_x \circ \pi|_{W_{x,w}}$ is compatible with the given structure. The uniqueness part of the claim follows.

Part (c): This is an application of Corollary 1.2.16. Let $\Lambda \colon M \to X$ be a right-inverse of π and $x \in M$. Using Corollary 1.2.16, there exists a path-connected chart (V_x, ψ_x) of $\Lambda(x)$ such that

$$\pi^{-1}\left(\pi(V_x)\right) = \coprod_{g \in G} g \cdot V_x.$$

Let $\psi_{g \cdot x} := \psi_x \circ g^{-1}$. Since the covering transformations of π are conformal maps, the charts

$$\left\{\left(g \cdot V_x, \psi_{g \cdot x}\right)\right\}_{(g,x) \in G \times M}$$

form an atlas of X compatible with the given structure. Define $W_x := \pi(g \cdot V_x)$ (independent of g) and $\phi_x := \psi_x \circ \left(\pi|_{V_x}\right)^{-1}$.

The claim is that $\{(W_x, \phi_x)\}_{x \in M}$ forms an atlas of M with conformal coordinate transformations. It is sufficient to check that the coordinate transformations are conformal maps.

Suppose that $x, y \in M$ such that $W_x \cap W_y \neq \emptyset$. Let $z \in \phi_x(W_x \cap W_y)$ and define $w = \phi_x^{-1}(z)$. The definitions of V_x and V_y imply that there exists unique $g, f \in G$ such that $x' := g \cdot \Lambda(w) \in V_x$ and $y' := f \cdot x' \in V_y$.

Let $U_{x'}$ be an open neighbourhood of x' contained in

$$\left(g \cdot V_{\Lambda(w)}\right) \cap \left(\pi|_{V_x}\right)^{-1} (W_x \cap W_y).$$

Corollary 1.2.16 implies that there exists a covering transformation $h \in G$ such that

$$\phi_y \circ \phi_x^{-1}\Big|_{K_z} = \psi_y \circ h \circ \psi_x^{-1}\Big|_{K_z}.$$

The right-hand side is a conformal map because h is a conformal map. This means that $\phi_y \circ \phi_x^{-1}$ is locally conformal. This implies that its inverse is also analytic, hence $\{(W_x, \phi_x)\}_{x \in M}$ is an atlas for M inducing a conformal structure for M. The map π is an analytic map as every point $z \in X$ and $\pi(z) \in M$ has charts that make the coordinare representation of π an identity map.

Part (d): Theorem 1.2.15 implies the existence of the maps μ_G and f as introduced in the claim. Part (c) can be applied to give X/G a conformal structure such that μ_G is an analytic map. Since the covering maps are locally conformal and locally $f = \pi \circ (\mu_G)^{-1}$, and likewise $f^{-1} = \mu_G \circ \pi^{-1}$, it follows that f and f^{-1} are analytic. This means that f is a conformal map. □

Theorem 3.3.4 (Uniformization theorem)**:**
Every simply-connected Riemann surface is conformally equivalent to exactly one of $\hat{\mathbb{C}}$, $\mathbb{C}$ and $\mathbb{H}$.

Proof:
This is proven in [12]. The proof does not fit the scope of this work. □

Remark 3.3.5*:*
The three model spaces $\hat{\mathbb{C}}$, $\mathbb{C}$ and $\mathbb{H}$ have a lot more structure. They have their canonical conformal structure as open subsets of $\hat{\mathbb{C}}$. Furthermore, they can be induced with a Riemannian surface structure: the Riemann sphere $\hat{\mathbb{C}}$ can be induced with the metric pulled back from the sphere via stereographic projection, the plane $\mathbb{C}$ with the Euclidean metric and the hyperbolic plane $\mathbb{H}$ with the metric defined as in [13], [12], and [1]: $\mathrm{d}s^2 = \frac{|dz|^2}{\mathrm{Im}(z)^2}$. The disk $\mathbb{D}$ can be induced with the pull-back metric from $\mathbb{H}$ via the conformal map $z \mapsto i\frac{1-z}{1+z}$. This makes $\mathbb{H}$ and $\mathbb{D}$ isomorphic as Riemann and Riemannian surfaces.

The conformal automorphisms of these spaces are precisely the ones stated in Remark 2.1.1. In the hyperbolic case every single Möbius transformation is an orientation-preserving isometry of $\mathbb{H}$ (or $\mathbb{D}$) [12, Proposition 2.1.2]. In the case of $\mathbb{C}$, the only elements of $\mathrm{Aut}(\mathbb{C})$ that are isometries in the Euclidean metric are the translations and rotations.

The hyperbolic plane has Gaussian curvature -1 with the metric defined as above [12, Proposition 2.1.12]. The definition of the metric differs by a factor of two from the metric used in [17], which is a standard reference on the topic of Teichmüller spaces. The same difference of a factor of two will be in the definition of the Teichmüller metric, which will be defined in Chapter 4.

Proposition 2.1.7 of [12] shows that the hyperbolic plane $\mathbb{H}$ and the hyperbolic disk $\mathbb{D}$ are uniquely geodesic. Also $\hat{\mathbb{C}}$ is a geodesic space. Then a corollary of the classical Hopf-Rinow theorem [7, Corollary 2.105] shows that $\hat{\mathbb{C}}$, $\mathbb{C}$ and $\mathbb{H}$ are complete and proper metric spaces. Since $\mathbb{C}$, $\mathbb{H}$ and $\mathbb{D}$ are proper and uniquely geodesic, a direct application of Arzela-Ascolí's theorem shows that the map $(x, y, t) \mapsto \gamma_{x,y}(t)$ is continuous, where $\gamma_{x,y}\colon [0, 1] \to X$ is the unique constant speed geodesic between x and y and X is $\mathbb{C}$, $\mathbb{H}$ or $\mathbb{D}$.

If G is a standard covering group of $\mathbb{C}$ or $\mathbb{H}$ – see Definition 2.2.3 – it follows that an element $g \in G$ satisfies $g \circ \gamma_{x,y} = \gamma_{g(x),g(y)}$ since g is an isometry and the spaces of interest are uniquely geodesic. This will be useful later on.

The distance between two points μ and μ' of $\mathbb{D}$ in the hyperbolic distance is given by

$$d_{\mathbb{D}}(\mu, \mu') = \log \frac{|1 - \overline{\mu}\mu'| + |\mu - \mu'|}{|1 - \overline{\mu}\mu'| - |\mu - \mu'|},$$

which is derived in [1]. Since $z \mapsto i\frac{1-z}{1+z}$ is an isometry between $\mathbb{D}$ and $\mathbb{H}$, it follows that given t and t' from $\mathbb{H}$, then

$$d_{\mathbb{H}}(t, t') = \log \left(\frac{|t' - \bar{t}| + |t - t'|}{|t' - \bar{t}| - |t - t'|} \right).$$

These distances will play an important role later in this text.

Definition 3.3.6:
Let $X = \hat{\mathbb{C}}, \mathbb{C}$ or $\mathbb{H}$ and M be a Riemann surface with a Riemannian metric. Let (X, π) be a covering space of M, and let G denote the covering group of π.

If π is an analytic covering map and a local isometry in the Riemannian sense, and G is a standard covering group of $\mathrm{Aut}(X)$, then (X, π) is said to be a standard covering of M.

Corollary 3.3.7:
Let G be a standard covering group of $X = \hat{\mathbb{C}}$, $\mathbb{C}$ or $\mathbb{H}$. Then X/G has a unique Riemann and Riemannian surface structures making the covering space (X, μ_G) a standard covering of X/G.

Proof:
The Riemann structure part follows immediately from the definition of standard covering groups Definition 2.2.3, the characterization of covering space actions Lemma 1.2.4, and Theorem 3.3.3. Since standard covers consist of isometries of X – see Definition 2.2.3 and the discussion in Remark 3.3.5 – the claim follows from Proposition 2.28 [7]. □

Remark 3.3.8:
From now on if G is a standard covering group of X, the associated surface X/G will have the unique structure making μ_G an analytic map and a local isometry.

Theorem 3.3.9 (Standard covers of Riemann surfaces)**:**
Let M be a Riemann surface. Then M has a unique Riemannian surface structure such that it admits a standard covering (X, π), where X is exactly one of $\hat{\mathbb{C}}, \mathbb{C}$, or $\mathbb{H}$. Let G be the covering group of π. Then the map $f\colon X/G \to M$ satisfying $f \circ \mu_G = \pi$ is conformal and an isometry.

Proof:
Every topological manifold has a universal cover as discussed in Remark 1.2.3. Theorem 3.3.3 gives the universal cover a structure of a Riemann surface such that the covering map is analytic. Theorem 3.3.4 and the fact that a composition of an analytic covering map and a conformal map is a covering map shows that it can be assumed that X is exactly one of $\mathbb{C}$, $\hat{\mathbb{C}}$ or $\mathbb{H}$.

Let (X, π') be the constructed analytic covering map of M and let H denote its covering group. Theorem 3.3.3 shows that the elements of H are conformal maps and H is a covering group as shown in Theorem 1.2.12. Corollary 2.2.4 implies that there exists a conformal map $\psi \in \mathrm{Aut}(X)$ for which $G := F_\psi(H)$ is a standard covering group. Define $\pi = \pi' \circ \psi^{-1}$. This is still an analytic covering map of M and it is readily seen that G is its covering group. Theorem 3.3.3 shows that M is conformally equivalent to X/G.

Then M can be induced with the Riemannian metric pulled backed from X/G using the conformal map f^{-1}; see Corollary 3.3.7. Since μ_G is a local isometry and f is an isometry, it follows that π is a local isometry. The uniqueness part of Proposition 2.28 [7] shows that this is the unique Riemannian surface structure for M such that π is a local isometry. □

Corollary 3.3.10 (Uniformization of Standard Covers)**:**
Let $\phi\colon M \to N$ be a K-quasiconformal map between Riemann surfaces, and let (X, π_M) and (Y, π_N) be standard coverings of X and Y, respectively. Then $X = Y$. In particular, every simply-connected Riemann surface is quasiconformally equivalent to exactly one of $\hat{\mathbb{C}}$, $\mathbb{C}$ or $\mathbb{H}$.

Proof:
Theorem 1.2.21 shows that there exists a homeomorphism $\psi\colon X \to Y$ such that $\pi_N \circ \psi = \phi \circ \pi_N$. Since π_N and π_M are locally conformal, it follows that ψ is a K-quasiconformal map. Note that X is compact if and only if Y is compact as ψ is a homeomorphism. This shows that $X = Y$ if either one them is $\hat{\mathbb{C}}$.

Suppose that $X = \mathbb{C}$. Since ψ is a quasiconformal map, Theorem 3.6.3 of [1] shows that $Y = \mathbb{C}$. This means that $X = \mathbb{C}$ if and only if $Y = \mathbb{C}$. This concludes the proof. □

Remark 3.3.11 *(Ideal boundaries):*
Let M be a Riemann surface with a standard covering (X, π) and a covering group G. Suppose that M has a boundary ∂M and that π extends to a continuous map from $X \cup I(G)$, where $I(G)$ is the ideal boundary of G (Remark 2.2.5), to $M \cup \partial M$ such that $\pi(I(G)) = \partial M$. Then it is readily seen that $\pi\colon X \cup I(G) \to M \cup \partial M$ is a topological covering map of $M \cup \partial M$ with G as its covering group. Furthermore, as (X, π) is a normal covering space of M, the same holds for its extension.

Theorem 1.2.15 applies in this case, and the homeomorphism f given by that result restricted to X/G is a conformal map and an isometry between X/G and M (Corollary 3.3.7), and it maps $I(G)/G$ homeomorphically onto ∂M. The boundary ∂M is said to be the ideal boundary of M.

If M is a Riemann surface such that π does not extend continuously to $I(G)$, the ideal boundary of M can be constructed such that π extends to $I(G)$: Let $M' = M$ and $\partial M' = I(G)/G$. Then define

$$\pi(z) = \begin{cases} \pi(z), & z \in X \\ \mu(z), & z \in I(G). \end{cases}$$

Extending $f\colon X/G \to M$ as identity to $I(G)$ implies that $\mu_G = f \circ \pi$. By inducing $M' \cup \partial M'$ with the quotient topology (use f as the quotient map) immediately makes f a homeomorphism. Then $\pi = f^{-1} \circ \mu_G$ is a covering map as a composition of a homeomorphism and a covering map, and it is the continuous extension of π.

This means that M can always be assumed to be a Riemann surface with a(n ideal) boundary ∂M such that its standard covering (X, π) extends continuously to $X \cup I(G)$ as discussed in the second paragraph of this remark.

Remark 3.3.12*:*
The culmination point of the theory developed so far is Theorem 3.3.9 and Remark 3.3.11. They show that every Riemann surface has a standard covering (X, π), where X is exactly one of $\hat{\mathbb{C}}$, $\mathbb{C}$ or $\mathbb{H}$. Furthermore, the Riemann surface has a boundary that is related to the covering group of π. The significance of the boundary will be seen in the next chapter. Since π is a local isometry and the curvature of the Riemann sphere is 1, the plane has curvature 0 and the hyperbolic plane has curvature -1, Riemann surfaces can be classified using their curvature.

Suppose that (X, π) is a standard covering of M. The fact that π is a local isometry implies that the lifts of geodesics of M are geodesics in X and the descensions of geodesics of X are geodesics of M ([7, Proposition 2.81]). This implies that the length metric d_M of M and d_X of X induced by the Riemannian metrics are characterized by

$$d_M(x, y) = \inf \left\{ d_X(x', y') \mid x' \in \pi^{-1}(x), y' \in \pi^{-1}(y) \right\}.$$

In particular, the map π is a local isometry in the metric sense and also 1-Lipschitz. It is worth noting that, in general, Riemannian length metrics do not change the original topology of a manifold, see [15, Theorem 13.29]. It can be checked that M is geodesically complete, hence M is complete and proper as a metric space [7, Corollary 2.105 (classical Hopf-Rinow theorem)].

A lot more details about the geometric structure of hyperbolic Riemann surfaces can be found in [13] and [12]. Also a more general look can be found in [7].

3.4 Quasiconformal maps of the Riemann sphere

Remark 3.4.1*:*
The next topic is to introduce some basic existence and compactness results for quasiconformal maps. Some of their corollaries will be of interest in the study of Teichmüller theory.

Remark 3.4.2*:*
Every element $B(\hat{\mathbb{C}})$, a Beltrami differential of the Riemann sphere, can be identified with its representation in the chart $(\mathbb{C}, \mathrm{id}_{\mathbb{C}})$. In fact, given $\mu \in L^\infty(\mathbb{C})$, by defining

$$\mu'(z) = \mu\left(\frac{1}{z}\right)\left(\frac{z}{|z|}\right)^4 \in L^\infty(\Omega),$$

the pair μ and μ' defines an element of $B(\hat{\mathbb{C}})$. From now on, every element of $B(\hat{\mathbb{C}})$ will be represented in $(\mathbb{C}, \mathrm{id}_{\mathbb{C}})$.

The Measurable Mapping Theorem – see [1, Theorem 5.3.4] – shows that given $\mu \in L^\infty(\mathbb{C})$ with $\|\mu\|_{\infty,\mathbb{C}} = k < 1$, there exists a quasiconformal map ψ of the Riemann sphere that fixes 0, 1, and ∞ such that $\mu = \mu_\psi$. The quasiconformal map can be identified with its coordinate representation in $(\mathbb{C}, \mathrm{id}_{\mathbb{C}})$. Furthermore, given a quasiconformal map from $\mathbb{C}$ onto itself fixing 0 and 1 defines a normalized quasiconformal map of the Riemann sphere. These observations are combined in the first part of the next result.

Theorem 3.4.3:
Let μ be a k-Beltrami differential of $\mathbb{C}$. Then there exists a K-quasiconformal map $\psi\colon \hat{\mathbb{C}} \to \hat{\mathbb{C}}$ normalized by the conditions $\psi(0) = 0$, $\psi(1) = 1$, and $\psi(\infty) = \infty$ solving the Beltrami PDE defined by μ. Every K-quasiconformal map of $\mathbb{C}$ onto itself fixing 0 and 1 extends to a normalized K-quasiconformal map of the Riemann sphere.

If $\Omega, \Omega' \subset \hat{\mathbb{C}}$ are open sets and $g\colon \Omega \to \Omega'$ is a K-quasiconformal map satisfying $\mu_g = \mu$ almost everywhere, there exists a conformal map $h\colon \psi(\Omega) \to \Omega'$ such that $g = h \circ \psi|_{\Omega}$. In particular, if $\Omega = \hat{\mathbb{C}}$ and g is normalized, then $g = \psi$.

If the boundaries of Ω and Ω' are Jordan curves in $\hat{\mathbb{C}}$, the quasiconformal map g has a continuous extension to the boundary of Ω that is a homeomorphism between the closures of Ω and Ω' in $\hat{\mathbb{C}}$.

Proof:
The first part follows from the discussion in Remark 3.4.2. The second part is a straight-forward consequence of Theorem 5.5.1 and Corollary 5.5.4 of [1].

Consider the extension result. The fact that g extends to the boundary follows from the observation that ψ maps Jordan domains to Jordan domains of $\hat{\mathbb{C}}$. Then the classical Carathéodory's theorem [9, Theorem 3.1], combined with the Uniformization Theorem 3.3.4, implies that h extends continuously to the boundary of $\partial\psi(\Omega) = \psi(\partial\Omega)$. The extension of h maps $\psi(\partial\Omega)$ homeomorphically onto the boundary of Ω'. Then the extension of g can be defined using the formula $h \circ \psi\colon \overline{\Omega} \to \overline{\Omega'}$. □

Remark 3.4.4*:*
The last part of Theorem 3.4.3 is a generalization of the classical Carathéodory's theorem. An alternative proof for the generalization can be found in [1, Section 5.9].

Lemma 3.4.5:
Let $\{\mu_n\}_{n=1}^{\infty}$ be a Cauchy sequence of k-Beltrami coefficients of $\mathbb{C}$. Then there exists a k-Beltrami coefficient μ of $\mathbb{C}$ such that

$$\mu(z) = \lim_{n\to\infty} \mu_n(z)$$

almost every $z \in \mathbb{C}$. Let f_n be the normalized solutions to the Beltrami PDE defined by μ_n. Then the limit

$$f(z) = \lim_{n\to\infty} f_n(z)$$

exists everywhere in $\hat{\mathbb{C}}$, the convergence of f_n to f is locally uniform in the Riemann sphere, and f is the normalized solution of the Beltrami PDE defined by μ. Additionally, $k_{f\circ f_n^{-1}}$ converges to zero when $n \to \infty$.

Proof:
This is a consequence of the completeness of $L^\infty(\mathbb{C})$, Theorem 5.3.5 [1] and Theorem 3.4.3. The chain rule of Beltrami differentials 3.1.3 implies that $k_{f\circ f_n^{-1}}$ converges to zero. □

Definition 3.4.6 (Extensions and normalizations)**:**
Let $X = \hat{\mathbb{C}}$, $\mathbb{C}$ or $\mathbb{H}$ and $f\colon X \to X$ be a quasiconformal map. A quasiconformal map $\tilde{f}$ of the Riemann sphere is an extension of f if the restriction of $\tilde{f}$ to X is equal to f and in the case $X = \mathbb{H}$, $\tilde{f}$ is reflection symmetric with respect to the extended real line.

The extension of $\mu \in L^\infty(X)$ is defined to be μ if $X = \mathbb{C}$. If $X = \mathbb{H}$, the extension of μ is defined as $\mu(z) = \overline{\mu}(\overline{z})$ for every $\overline{z} \in \mathbb{H}$.

The family of quasiconformal maps from X onto X, whose extensions fix 0, 1 and ∞ are denoted by $\mathcal{F}_X$. The elements of $\mathcal{F}_X$ are said to be normalized.

Theorem 3.4.7 (Extensions of Beltrami differentials)**:**
Let $X = \hat{\mathbb{C}}$, $\mathbb{C}$ or $\mathbb{H}$ and let μ be a Beltrami differential of X. Then there exists a unique normalized quasiconformal map $\psi \in \mathcal{F}_X$ such that $\mu_\psi = \mu$.

The map ψ is the restriction of the normalized solution of the Beltrami PDE defined by the extension of μ. If $\psi'\colon X \to X$ is a quasiconformal map with $\mu_{\psi'} = \mu$ almost everywhere, there exists a Möbius transformation $h \in \mathrm{Aut}(X)$ such that $\psi' = h \circ \psi$.

Proof:
Theorem 3.4.3 handles the first two cases. Consider the case $X = \mathbb{H}$. Let ψ be the normalized solution of the extension PDE defined by the extension of μ. The goal is to show that $\psi \in \mathcal{F}_{\mathbb{H}}$. The boundary of $\psi(\mathbb{H})$ is a Jordan curve, therefore the claim follows straight-forwardly from Theorem 3.3.4, Theorem 3.4.3, and Lemma 3.1.3 as seen from below.

First extend μ as $\mu(z) = \overline{\mu}(\overline{z})$ for every $\overline{z} \in \mathbb{H}$, and let $\psi\colon \hat{\mathbb{C}} \to \hat{\mathbb{C}}$ be the normalized solution solving that Beltrami PDE in the Riemann sphere. Let $\Omega = \psi(\mathbb{H})$. The boundary of Ω is a Jordan curve, therefore Theorem 3.3.4 implies the existence of a conformal map $h\colon \Omega \to \mathbb{H}$. Now $\phi := h \circ \psi\colon \mathbb{H} \to \mathbb{H}$ is a quasiconformal map as a composition of two such maps.

Theorem 3.4.3 shows that ϕ can be extended to the boundary of $\mathbb{H}$ and the extension maps it homeomorphically onto itself. Then there exists a conformal map $g \in \mathrm{Aut}(\mathbb{H})$ such that $g \circ \phi\colon \mathbb{H} \cup \partial\mathbb{H} \to \mathbb{H} \cup \partial\mathbb{H}$ is normalized. The map $g \circ \phi$ can be extended to the Riemann sphere by reflecting it along the extended real line.

Lemma 3.1.3 shows that the Beltrami differential of $g \circ \phi$ and ψ agree on $\mathbb{H}$, and since they are both reflection symmetric, the Beltrami differentials agree in the

Riemann sphere. Uniqueness of normalized solutions implies that $g \circ \phi = \psi$. This means that $\psi \in \mathcal{F}_{\mathbb{H}}$. The uniqueness follows from Theorem 3.4.3. □

Corollary 3.4.8 (Extensions of quasiconformal maps)**:**
Let $f\colon X \to X$ be as in Definition 3.4.6. Then the extension $\tilde{f}$ exists and is unique. If $g, h \in \mathrm{Aut}(X)$ such that $f \circ g = h \circ f$, the extension of f satisfies $\tilde{f} \circ g = h \circ \tilde{f}$ in $\hat{\mathbb{C}}$.

Proof:
After the discussion in Remark 3.4.2, the only unclear case is $X = \mathbb{H}$. This is straight-forward to verify after Theorem 3.4.7. □

Remark 3.4.9*:*
The notation $\tilde{f}$ will be dropped as f's extension is unique, thus f can be identified with $\tilde{f}$. Whenever an extension of f is mentioned, it refers to this specific one.

3.5 Lifting characterization of quasiconformal maps

Remark 3.5.1*:*
From now on X will denote $\hat{\mathbb{C}}$, $\mathbb{C}$ or $\mathbb{H}$. Let $\mathcal{G}_X$ denote the standard covering groups of X and $\mathcal{R}_X$ denote the set of Riemann surfaces M admitting a standard covering (X, π_M) with an ideal boundary ∂M.

For every $M \in \mathcal{R}_X$ there exists a standard covering (X, π_M) such that π_M extends continuously to $I(G_M)$ and $\partial M = \pi_M(I(G_M))$ as discussed in Remark 3.3.11. From now on (X, π_M) will refer to such a standard covering and G_M to its covering group. If $M = X/G$, it will be assumed that $\pi_M := \mu_G$. Theorem 3.3.9 (and Remark 3.3.11) show that given $M \in \mathcal{R}_X$, then X/G_M and M are related by

$$\begin{array}{ccc} X \cup I(G_M) & \xrightarrow{\mu_{G_M}} & (X \cup I(G_M))/G_M \\ & \searrow^{\pi_M} & \downarrow f_M \\ & & M \cup \partial M \end{array} \tag{3.5.1}$$

The map f_M is a homeomorphism mapping $I(G_M)/G_M$ homeomorphically onto ∂M, and mapping X conformally and isometrically onto M. The ideal boundary will become relevant in the next chapter.

Remark 3.5.2:
The next theorem plays a major role in the Teichmüller theory developed in the next chapter. It will also be utilized in establishing the existence result for the solutions of Beltrami PDE's in Riemann surfaces.

Theorem 3.5.3 (Lifts of quasiconformal maps)**:**
Let $\phi\colon M \to N$ be a homeomorphism between Riemann surfaces and let (X, π_M) and (X, π_N) be standard covers of M and N, respectively.

(a) For every $x \in X$ and $y \in \pi_N^{-1}(\phi(\pi_M(x)))$, there exists a unique homeomorphism $\psi\colon X \to X$ such that $\pi_N \circ \psi(y) = \phi \circ \pi_M(x)$ and ψ is a lift of ϕ along π_M and π_N.

(b) The map ψ is K-quasiconformal if and only if ϕ is.

(c) If ϕ and ψ are as above, and G_M and G_N are the covering groups of π_M and π_N, respectively, the map $F_\psi\colon G_M \to G_N$ is well-defined and a group isomorphism.

(d) Suppose that $\psi\colon X \to X$ is a quasiconformal map and the conjugation map $F_\psi\colon G_M \to G_N$ is well-defined and an isomorphism. Then there exists a unique quasiconformal map $\phi\colon M \to N$ such that ψ is a lift of ϕ along π_M and π_N.

(e) Suppose that ϕ and ϕ' are homotopic quasiconformal maps and ψ is a lift of ϕ. Then there exists a unique lift ψ' of ϕ' for which $F_\psi = F_{\psi'}$.

   Furthermore, if $X = \hat{\mathbb{C}}$, any two quasiconformal maps ϕ and ϕ' are homotopic and $F_\psi = F_{\psi'}$ are trivial maps. If $X = \mathbb{C}$ or $\mathbb{H}$ and there exists a lift ψ' of ϕ' such that $F_\psi = F_{\psi'}$, then ϕ and ϕ' are homotopic.

Proof:
Part (a), Part (c), Part (d) and first part of (e) follow from Theorem 1.2.21. Part (b) is trivial, because analytic covering maps are locally conformal. Thus the second part of (e) is the only part that requires more justification.

Suppose that $X = \hat{\mathbb{C}}$. Theorem 2.2.2 shows that π_M and π_N are conformal maps and the isomorphisms $F_\psi = F_{\psi'}$ are trivial maps. Since quasiconformal maps are orientation-preserving homeomorphisms, it follows that ψ and ψ' have the same degree. Hopf Degree Theorem [10, Chapter 3, p. 146] shows that ψ and ψ' are homotopic. Since $\phi = \pi_N^{-1} \circ \psi \circ \pi_M$ and $\phi' = \pi_N^{-1} \circ \psi' \circ \pi_M$, it follows that ϕ and ϕ' are also homotopic.

Suppose that $X = \mathbb{H}$ or $\mathbb{C}$. Let $\gamma_{xy}\colon [0, 1] \to X$ be the family of geodesics as discussed in Remark 3.3.5. The map $(x, t) \mapsto \theta_t(x) := \gamma_{\psi(x)\psi'(x)}(t)$ is a continuous function as a composition of two continuous functions.

Notice that $\theta_0(x) = \psi(x)$ and $\theta_1(x) = \psi'(x)$, i.e. θ is a homotopy between ψ and

ψ'. Let $g \in G_M$. Then $F_{\psi'}(g)$ is an element of G_N, therefore it is an isometry – see Remark 3.3.5 – hence

$$F_{\psi'}(g)(\theta_t(x)) = \gamma_{F_{\psi'}(g)\circ\psi(x), F_{\psi'}(g)\circ\psi(x)}.$$

Since $F_{\psi'} = F_\psi$, the definitions of the conjugation map F and the homotopy θ_t show that this is equal to $\theta_t(g(x))$. It follows that

$$F_{\psi'}(g) \circ \theta_t = \theta_t \circ g.$$

By considering Theorem 1.2.21 Part (b) for the covering maps $\pi_M \times \mathrm{id}_{[0,1]}$ and $\pi_N \times \mathrm{id}_{[0,1]}$, there exists a unique continuous map h_t for which $h_t \circ \pi_M = \pi_N \circ \theta_t$. The uniqueness of descensions implies that $h_0 = \phi$ and $h_1 = \phi'$, therefore h_t is the desired homotopy. □

Definition 3.5.4:
Let $X = \hat{\mathbb{C}}$, $\mathbb{C}$ or $\mathbb{H}$ and G a subgroup of $\mathrm{Aut}(X)$. A Beltrami differential μ of X is G-invariant, if $\mu = g^*\mu$ for every $g \in G$. The collection of G-invariant Beltrami differentials of X is denoted by $B_X(G)$.

Remark 3.5.5:
Let μ be a Beltrami differential on M. Then the pullback $\eta := \pi_M^*\mu$ defines a Beltrami differential on X as discussed in Definition 3.2.8. Since $\pi \circ g = \pi$ for every $g \in G_M$, the differential η satisfies $g^*\eta = \eta$. Conversely, if there is a Beltrami differential η on X that satisfies $g^*\eta = \eta$ for every $g \in G_M$, then there exists a unique Beltrami differential μ on M such that $\pi_M^*\mu = \eta$.

If $X = \hat{\mathbb{C}}$ or $\mathbb{C}$, the Beltrami differential η can be identified with its coordinate representation in $\mathbb{C}$. Furthermore, in the case $X = \mathbb{H}$, η can be represented using its coordinate representation in $\mathbb{H}$. In the Riemann sphere case, the Beltrami differential η has no useful symmetries as G is trivial – this follows from Theorem 2.2.2 – however, if X is the Euclidean plane or the hyperbolic plane, the fact that $g^*\eta = \eta$ implies that $\eta = (\eta \circ g) r_g = (\eta \circ g)\frac{\overline{g_z}}{g_z}$ almost everywhere for every $g \in G$.

Theorem 3.5.3 and Lemma 3.1.4 show that a quasiconformal map between two Riemann surfaces M and N in $\mathcal{R}_X$ lifts to a quasiconformal map. Its Beltrami differential is G_M-invariant and its inverse's Beltrami differential is G_N-invariant.

3.6 Examples of Riemann surfaces

Definition 3.6.1:
Let $\arg\colon \mathbb{C} \to (-\pi, \pi]$ denote the branch of arg such that $\arg \mathbb{R}_+ = 0$ and $\arg \mathbb{R}_- = \pi$. For $r > 1$ define $\lambda_r = \exp\left(\frac{2\pi^2}{\log r}\right)$ and let $A(1, r)$ denote the open annulus centered at the origin with inner radius 1 and outer radius r. Let $\pi_r\colon \mathbb{H} \to A(1, r)$ denote the map defined by $\pi_r(z) = \exp\left(-2\pi i \frac{\log z}{\log \lambda_r}\right)$. For $t \in \mathbb{H}$, let $G_{1,t}$ denote the subgroup of $\mathrm{Aut}(\mathbb{C})$ generated by T_1 and T_t.

Remark 3.6.2:
It is readily seen that $\pi_r\colon \mathbb{H} \to A(1, r)$ is an analytic covering map and its covering group is generated by g_{λ_r}. Furthermore, Theorem 3.3.9 shows that $A(1, r)$ is conformally equivalent to $\mathbb{H}/\langle g_{\lambda_r}\rangle$. It follows that $(\mathbb{H}, \pi_r)$ is a standard covering of $A(1, r)$, when the annulus is endowed with the pushforward metric from $\mathbb{H}$. The map $\lambda_r\colon (1, \infty) \to (1, \infty)$, where $r \mapsto \lambda_r$ is surjective, therefore a Riemann surface with a covering group generated by a hyperbolic element is conformally equivalent to some annulus $A(1, r)$, where $r > 1$. This follows from Corollary 2.2.4.

As π_r extends continuously to $\mathbb{R} \setminus \{0\}$, where the positive real line is mapped onto $S(0, 1)$, and the negative real line onto $S(0, r)$, the discussion in Corollary 3.3.7 and Remark 3.3.11 show that the extension of π_r is a covering map of the closed annulus. In this case the ideal boundary $I(\langle g_{\lambda_r}\rangle)$ is mapped to a disjoint union of 1-manifolds.

Every Riemann surface $\mathbb{C}/G_{1,t}$ for $t \in \mathbb{H}$ is topologically a torus of genus one. They can be endowed with Riemannian metrics that make $(\mathbb{C}, \mu_{G_{1,t}})$ a standard covering of $\mathbb{C}/G_{1,t}$. Remark 1.2.14 and Theorem 2.2.2 imply that every Riemann surface that is topologically a torus of genus one is conformally equivalent to some such surface.

If the standard covering group of M is $\langle T_1\rangle$ in $\mathbb{C}$ or in $\mathbb{H}$, then M is conformally equivalent to $\mathbb{C} \setminus \{0\}$ or $\mathbb{D} \setminus \{0\}$, respectively. A standard covering is given by $z \mapsto \exp(2\pi i z)$, when $\mathbb{C} \setminus \{0\}$ or $\mathbb{D} \setminus \{0\}$ are endowed with the pushforward metric from $\mathbb{C}$ and $\mathbb{H}$, respectively. The covering map of the punctured disk extends to the real line continuously and it is the ideal boundary of the covering group $\langle T_1\rangle$. The ideal boundary is mapped to the unit circle by the covering map. The ideal boundary of the punctured plane is empty.

If the standard covering group of an analytic covering of a Riemann surface is trivial, the Riemann surface is conformally equivalent to its universal cover.

The Riemann surfaces discussed so far have had Abelian covering groups. Actually, every another family of Riemann surfaces have non-Abelian covering groups. The standard covering space is $\mathbb{H}$, i.e. all of these Riemann surfaces are hyperbolic with Gaussian curvature -1.

The rest of compact Riemann surfaces are connected sums of 1-tori, thus such surfaces have genus $g \geq 2$ (Proposition 6.20 of [14] – the classification of orientable compact surfaces). It is shown in Theorem 2.22 of [13] that $G \subset \text{Aut}(\mathbb{H})$ consists entirely of hyperbolic elements and in Proposition 2.19 that the limit set $L(G)$ is the extended real line, i.e. $I(G)$ is empty. All compact Riemann surfaces are now classified.

Another example of interest are Riemann surfaces of finite analytic type (g, n): A Riemann surface M is of finite analytic type (g, n) if there exists a compact Riemann surface M' of genus g and an analytic embedding $\phi_M \colon M \to M'$ such that the cardinality of $M' \setminus \phi_M(M)$ is n. It is later shown that if $g = 0$ and $n \geq 3$, $g = 1$ and $n \geq 1$, or $g \geq 2$ and $n \geq 0$, the Riemann surface is hyperbolic. Otherwise the surface is non-hyperbolic. Furthermore, the type (g, n) turns out to be a topological invariant in the following sense: if Riemann surfaces of type (g, n) and (g', n') are homeomorphic, then $(g, n) = (g', n')$. Additionally, any pair of homeomorphic surfaces of type (g, n) are actually quasiconformally equivalent. This is discussed in the last chapter.

Chapter 4

Teichmüller theory of Riemann surfaces

4.1 Deformation space of standard covering groups

Definition 4.1.1:
Let $X = \hat{\mathbb{C}}$, $\mathbb{C}$ or $\mathbb{H}$. Given a standard covering group $G \in \mathcal{G}_X$ of X, let $\mathrm{Def}_X(G)$ denote the collection of pairs (H, ψ) such that $H \in \mathcal{G}_X$ is a standard covering group of X, the conjugation map $F_\psi \colon G \to H$ is an isomorphism, and $\psi \colon X \to X$ is a quasiconformal map. The collection $\mathrm{Def}_X(G)$ is the deformation space of G.

Given $(G', \psi') \in \mathrm{Def}_X(G)$ and $(H, \psi) \in \mathrm{Def}_X(G)$, the function $(G', \psi')^\#$ is defined as

$$(G', \psi')^\# \colon \mathrm{Def}_X(G) \to \mathrm{Def}_X(G'),\ (H, \psi) \mapsto (H, \psi \circ (\psi')^{-1}).$$

The composition of two such functions, whenever $\psi'' \circ \psi'$ is well-defined, is defined as

$$(G'', \psi'')^\# \circ (G', \psi')^\# = (G'', \psi'' \circ \psi')^\#.$$

Remark 4.1.2*:*
It is clear that the functions $(\dots)^\#$ are well-defined, the identity element (G, id_X) induces the identity map $(G, \mathrm{id}_X)^\#$, and if $(H, \psi) \in \mathrm{Def}_X(G)$, then $(G, \psi^{-1})^\#$ is the inverse function of $(H, \psi)^\#$.

The motivation for this section is given by Theorem 3.5.3, Corollary 3.3.10, and Section 3.4. Given two Riemann surfaces M and N with standard coverings (X, π_M) and (Y, π_N) (Definition 3.3.6), and a quasiconformal map $\phi \colon M \to N$, then $X = Y$, and Theorem 3.5.3 shows that ϕ admits a lift $\psi \colon X \to X$ along π_M and π_N. It also shows that ψ defines an isomorphism $F_\psi \colon G_M \to G_N$, where $g \mapsto \psi \circ g \circ \psi^{-1}$ is the conjugate isomorphism discussed in Definition 1.2.19.

This means that the properties of $\mathrm{Def}_X(G)$ tell a lot about the properties of quasiconformal maps between Riemann surfaces. Even more can be said and that will be the topic of the next section. Before that some basic properties of $\mathrm{Def}_X(G)$ will be studied.

Proposition 4.1.3:
Suppose that X is $\hat{\mathbb{C}}$, $\mathbb{C}$, or $\mathbb{H}$ and $G \in \mathcal{G}_X$ is a standard covering group of X.

(a) If $\psi \in \mathcal{F}_X$ is a normalized quasiconformal map and $F_\psi(G)$ is a subgroup of $\mathrm{Aut}(\hat{\mathbb{C}})$, then $F_\psi(G)$ is a standard covering group of X and of the same type as G. In particular, $\big(F_\psi(G), \psi\big) \in \mathrm{Def}_X(G)$.

If G contains T_1, then $F_\psi(T_1) = T_1$ and if G contains T_t for $t \in \mathbb{H}$, then $F_\psi(T_t) = T_{\psi(t)}$, where $\psi(t) \in \mathbb{H}$. If G is generated by g_λ, then $F_\psi(g_\lambda) = g_{\psi(\lambda)}$ generates $F_\psi(G)$ and $\psi(\lambda) > 1$. If G is non-Abelian and for every $x = 0, 1, \infty$, the element $h_x \in G$ fixes x, then $F_\psi(h_x)$ fixes x and x is a same type of fixed point for h_x and $F_\psi(h_x)$.

(b) Let $\psi \in \mathcal{F}_X$ such that $\big(F_\psi(G), \psi\big) \in \mathrm{Def}_X(G)$. If $\psi' \in \mathcal{F}_X$ is a normalized quasiconformal map and there exists $C \in \mathrm{Aut}(X)$ such that $F_C = F_{\psi'\circ\psi^{-1}}$, then $F_{\psi'} = F_\psi \colon G \to F_\psi(G)$. In particular, $\big(F_\psi(G), \psi'\big) \in \mathrm{Def}_X(G)$.

Proof:
Part (a): If G is the trivial group, there is nothing to do, thus the non-triviality can be assumed. Suppose that G contains an element $\mathrm{id}_X \neq h_x$ fixing x, where $x = 0$, 1, or ∞. Define $f_x = F_\psi(h_x)$. Since ψ is normalized, Lemma 2.1.3 shows that x is fixed by f_x. It also shows that if h_x has a single fixed point, then f_x has only one fixed point, and in this case it is x. If h_x has an attractive fixed point at x and a repulsive fixed point and x', then f_x has an attractive fixed point at x and a repulsive one at $\psi(x')$. This actually implies the non-Abelian cases of the claim.

Consider the Abelian cases. Suppose that T_α is contained in G. Then $T_{n\alpha} \in G$ for every $n \in \mathbb{Z}$ and $f^n := F_\psi(T_{n\alpha})$ for some $f \in \mathrm{Aut}(X)$. The discussion above shows that f^n fixes ∞ and no other points. Since ψ is normalized, it follows that

$$f^n(0) = \psi(n\alpha), \quad f^n(1) = \psi((n+1)\alpha).$$

It is readily verified that this can happen if and only if $f = T_{\psi(\alpha)}$. This implies that if G is generated by T_1, the group $F_\psi(G)$ is generated by $T_1 = F_\psi(T_1)$. If G is generated by T_1 and T_t, where $t \in \mathbb{H}$, then $F_\psi(G)$ is generated by $T_1 = F_\psi(T_1)$ and $T_{\psi(t)} = F_\psi(T_t)$. It is shown in Lemma 4.5.6 that $\psi(t) \in \mathbb{H}$ (independent result).

If G is generated by g_λ for $\lambda > 1$, then $g_\mu = F_\psi(g_\lambda)$ for some $\mu > 1$; the Möbius transformation $F_\psi(g_\lambda)$ has an attractive fixed point at ∞ and a repulsive one at 0, therefore such a μ exists. Computing $g_\mu = F_\psi(g_\lambda)$ at $z = 1$ shows that $\mu = \psi(\lambda)$.

Part (b): First observe that $F_{\psi'} = F_C \circ F_\psi$. Since $F_\psi(G)$ is a subgroup of $\mathrm{Aut}(X)$ and C is a conformal map, it follows that $F_{\psi'}(G)$ is a subgroup of $\mathrm{Aut}(X)$ as well.

If f_x and h_x are as in the beginning of the proof of the first part and $\theta_x = F_{\psi'}(f_x)$, the assumption on C shows that $F_C(h_x) = \theta_x$. The claim follows from Part (a) and Lemma 2.1.3. □

Proposition 4.1.4:
Let $X = \hat{\mathbb{C}}$, $\mathbb{C}$ or $\mathbb{H}$, and let H and G be subgroups of $\mathrm{Aut}(X)$. If $F\colon H \to G$ is a group isomorphism, the subfamily of normalized K-quasiconformal maps from

$$\{\psi \in \mathcal{F}_X \mid F_\psi = F\}$$

is a normal family. In particular, the K-quasiconformal maps from $\mathcal{F}_X$ form a normal family.

Proof:
By considering the extensions of K-quasiconformal maps of $\mathcal{F}_X$, the normalicity follows from Theorem 3.9.4 [1]. □

Proposition 4.1.5:
Suppose that X is $\hat{\mathbb{C}}$, $\mathbb{C}$ or $\mathbb{H}$ and G is a subgroup of $\mathrm{Aut}(X)$. Let $(H_n, \psi_n) \in \mathrm{Def}_X(G)$ be a sequence of normalized K-quasiconformal maps such that $\left(\mu_{\psi_n}\right)_{n=1}^{\infty}$ is a Cauchy sequence in $L^\infty(\mathbb{C})$.

Then the normalized solution ψ of the limit Beltrami differential μ of μ_{ψ_n} is a K-quasiconformal map, it is the locally uniform limit of ψ_n, $k_{\psi\circ\psi_n^{-1}}$ converges to 0 when $n \to \infty$, and $F_\psi(G)$ is a standard covering group of X, i.e. $\left(F_\psi(G), \psi\right) \in \mathrm{Def}_X(G)$.

Proof:
By extending ψ_n to the Riemann sphere, Lemma 3.4.5 shows the existence of ψ and μ with the wanted properties. Proposition 4.1.4 shows that the limit is an element of $\mathcal{F}_X$ and $F_\psi(G)$ is a subgroup of $\mathrm{Aut}(X)$ as a corollary of Lemma 3.1.4 and Lemma 3.1.3. Since G is a standard covering group, Proposition 4.1.3 shows that so is $F_\psi(G)$. □

Proposition 4.1.6:
Let (H, ψ) and (H', ψ') be two normalized quasiconformal maps of $\mathrm{Def}_X(G)$. Then the maps ψ and ψ' are homotopic via $(H_t, \psi_t) \in \mathrm{Def}_X(G)$ such that $\psi_0 = \psi$, $\psi_1 = \psi'$, and ψ_t are normalized quasiconformal maps.

The homotopy can be chosen in such a way that $\log K_{\psi_t\circ\psi_s^{-1}} = |t-s| \log K_{\psi'\circ\psi^{-1}}$, in addition to satisfying $\mu_{\psi_s\circ\psi_t^{-1}} \to 0$ and $\mu_{\psi_t} \to \mu_{\psi_s}$ in $L^\infty(X)$ as $t \to s$. Furthermore, the map $(t, z) \mapsto \psi_t^{-1}(z)$ is a homotopy between ψ^{-1} and $(\psi')^{-1}$.

Proof:
First consider the case $\psi' = \mathrm{id}_X$ and $F_\psi\colon G \to H$ is an isomorphism. Consider the Beltrami differential

$$\mu_t = \frac{(1+|\mu_\psi|)^t - (1-|\mu_\psi|)^t}{(1+|\mu_\psi|)^t + (1-|\mu_\psi|)^t}\frac{\mu_\psi}{|\mu_\psi|}.$$

The expression is defined for almost every $z \in X$. The following expressions in the proof are evaluated at some such point. The path $t \mapsto \mu_t$ defines the constant speed radial geodesic between 0 and μ_ψ in the hyperbolic metric $d_{\mathbb{D}}$ introduced in Remark 3.3.5.

Let $\psi_t \in \mathcal{F}_X$ denote the normalized solution to the Beltrami PDE μ_t given by Theorem 3.4.7. The Beltrami differential of ψ is G-invariant since F_ψ is an isomorphism, hence so is μ_t by construction. Lemma 3.1.4 shows that $H_t := F_{\psi_t}(G)$ is a group of Möbius transformations and the normalicity of ψ_t and standardness of G combined with Proposition 4.1.3 shows that H_t is a standard covering group of X. Uniqueness of normalized solutions implies that $\mathrm{id}_X = \psi_0$ and $\psi = \psi_1$. The chain rule of Beltrami differentials 3.1.3 implies that

$$d_{\mathbb{D}}\left(0, \mu_{\psi_t \circ \psi_s^{-1}}\right) = d_{\mathbb{D}}\left(\mu_t, \mu_s\right) = |t-s|\, d_{\mathbb{D}}\left(0, \mu_\psi\right).$$

If $s \to t$, Lemma 3.4.5 shows that ψ_s converges to ψ_t locally uniformly. The left-hand side implies that $\psi_t \circ \psi_s^{-1}$ converges to the identity locally uniformly as $|t-s| \to 0$. Also, by taking the essential supremum from both sides, it is readily checked that

$$\log K_{\psi_t \circ \psi_s^{-1}} = |t-s| \log K_\psi.$$

This shows that $\psi_t \circ \psi_s^{-1}$ converges to the identity and ψ_s to ψ_t locally uniformly as $s \to t$. The map $z \mapsto \psi_t^{-1}(z)$ is locally Hölder continuous, where the Hölder exponent depends only on the supremum of $K_{\psi_t^{-1}} \leq K_\psi$ (Corollary 3.10.3 of [1]). As a consequence, the maps ψ_s^{-1} converge to ψ_t^{-1} locally uniformly as $s \to t$. The locally uniform convergences imply that the maps $(t, z) \mapsto \psi_t^{-1}(z)$ and $(t, z) \mapsto \psi_t(z)$ are continuous.

For the general case, let $f = \psi' \circ \psi^{-1}$ and construct f_t as above. Then $\psi_t := f_t \circ \psi \in \mathcal{F}_X$ is a homotopy between $\psi_0 = \psi$ and $\psi_1 = \psi'$. The map ψ_t satisfies the desired properties. □

4.2 Deformation space of Riemann surfaces

Remark 4.2.1*:*
Let $X = \hat{\mathbb{C}}$, $\mathbb{C}$ or $\mathbb{H}$. Then $M \in \mathcal{R}_X$ refers to a Riemann surface admitting a standard covering (X, π_M), where π_M extends to the ideal boundary $I(G_M)$ and it is mapped to the ideal boundary ∂M of M – the group G_M is the covering group of π_M. Furthermore, the covering map π_M extends to a normal topological covering map $(X \cup I(G_M), \pi_M)$ of $M \cup \partial M$ with the covering group G_M. The notations are introduced in Remark 3.5.1. Corollary 3.3.10 shows that if $\phi\colon M \to N$ is a quasiconformal map, then $N \in \mathcal{R}_X$.

Definition 4.2.2 (Deformation space)**:**
Let $M \in \mathcal{R}_X$. The collection $\mathrm{Def}(M)$, called the deformation space of M, denotes the pairs (N, ϕ) where $\phi\colon M \to N$ is a quasiconformal map and $N \in \mathcal{R}_X$. Given $(M', \phi') \in \mathrm{Def}(M)$ and $(N, \phi) \in \mathrm{Def}(M)$, then

$$(M', \phi')_\#\colon \mathrm{Def}(M) \to \mathrm{Def}(M'),\ (N, \phi) \mapsto (N, \phi \circ (\psi')^{-1}).$$

The composition of $(\dots)_\#$ is defined as in Definition 4.1.1.

Remark 4.2.3*:*
It is clear that the functions $(M', \phi')_\#$ are well-defined, the identity element is the function defined by (M, id_M), and every function $(\dots)_\#$ is invertible.

Given a standard covering (X, π_M) of M, the connection between $\mathrm{Def}_X(G_M)$ and $\mathrm{Def}(M)$ is discussed in Remark 4.1.2. Even more can be said. If $(H, \psi) \in \mathrm{Def}_X(G_M)$, Corollary 3.3.7 shows that $X/H \in \mathcal{R}_X$, and Theorem 3.5.3 shows that there exists a unique quasiconformal map $(X/H, \phi) \in \mathrm{Def}(M)$ that is the descension of ψ along π_M and $\pi_{X/H} = \mu_H$. Theorem 3.5.3 shows that lifting and descending quasiconformal maps does not change their maximal dilatations.

Remark 4.2.4*:*
Given a standard covering (X, π_M) of M with a covering group G_M, the pullback π_M^*, defined in Definition 3.2.8, induces a bijection from $B(M)$ – the Beltrami differentials of M – onto $B_X(G_M)$, the G_M-invariant Beltrami differentials of X (Remark 3.5.5).

Let $\mu \in B(M)$ be a Beltrami differential on M and $\pi_M^*(\mu) \in B_X(G_M)$. The notation $(H^\mu, \psi^\mu) \in \mathrm{Def}_X(G_M)$ refers to the normalized quasiconformal map $\psi^\mu \in \mathcal{F}_X$ such that $\mu_{\psi^\mu} = \pi_M^*(\mu)$, where the standard covering group H^μ is equal to $F_{\psi^\mu}(G_M)$.

Then N_μ refers to the Riemann surface X/H^μ with standard covering $(X, \pi_{N_\mu} = \mu_{H^\mu})$, and $\phi_\mu\colon M \to N_\mu$ to the quasiconformal map solving the Beltrami PDE defined by μ such that ψ^μ is its lift along π_M and $\pi_{N_\mu} = \mu_{H^\mu}$. The existence of these is established in Theorem 4.2.5.

Theorem 4.2.5 shows that the maps $\mu \mapsto \pi_M^*(\mu)$, $\mu \mapsto (H^\mu, \psi^\mu)$, and $\mu \mapsto (N_\mu, \phi_\mu)$ are injective and well-defined. The idea of the next section is to define appropriate equivalence relations on $B(M)$, $B_X(G_M)$, $\mathrm{Def}(M)$, and $\mathrm{Def}_X(G_M)$ to make the above maps bijective modulo the equivalence relations, but still have some meaningful structure on the spaces. Each of these spaces will highlight different kind of properties of the same structure of interest.

Theorem 4.2.5 (Measurable Riemann mapping theorem)**:**
Let μ be a Beltrami differential on M. Given a standard covering (X, π_M) of M, there exists a unique normalized quasiconformal map $\psi^\mu \in \mathcal{F}_X$ such that $\pi_M^*(\mu) = \mu_{\psi^\mu}$. Let $H^\mu := F_{\psi^\mu}(G_M)$ and $N_\mu = X/H^\mu$.

Then N_μ is a Riemann surface, (X, μ_{H^μ}) its standard cover, and there exists a unique quasiconformal map $(N_\mu, \phi_\mu) \in \mathrm{Def}(M)$ such that $\mu = \mu_{\phi_\mu}$, and ϕ_μ is the descension of ψ^μ along π_M and μ_{H^μ}.

Let N' be a Riemann surface and $(X, \pi_{N'})$ its standard cover. Suppose that $(N', \phi') \in \mathrm{Def}(M)$, and $(G_{N'}, \psi') \in \mathrm{Def}_X(G_M)$ is its lift along π_M and $\pi_{N'}$. Then $\mu_{\phi'} = \mu$ if and only if $(N_\mu, \phi_\mu)_\#(N', \phi') \in \mathrm{Def}(N_\mu)$ is a conformal map if and only if $(H^\mu, \psi^\mu)^\#(G_{N'}, \psi') \in \mathrm{Def}_X(H^\mu)$ is an element of $\mathrm{Aut}(X)$.

Proof:
Theorem 3.4.7 shows that there exists a unique $\psi^\mu \in \mathcal{F}_X$ such that $\mu_{\psi^\mu} = \pi_M^*(\mu)$. Since $\pi^*(\mu)$ is G_M-invariant, Corollary 3.4.8 shows that $H^\mu := F_{\psi^\mu}(G_M)$ defines a subgroup of $\mathrm{Aut}(X)$. Proposition 4.1.3 says that it is a standard covering group of M. Then $N_\mu := X/H^\mu$ is a Riemann surface with a standard covering $(X, \pi_{N_\mu} := \mu_{H^\mu})$ as shown in Corollary 3.3.7. Theorem 3.5.3 implies the existence of a unique quasiconformal $(N_\mu, \phi_\mu) \in \mathrm{Def}(M)$ that is the descension of ψ^μ along π_M and π_{N_μ}. It is clear that $\mu_{\phi_\mu} = \mu$.

Suppose that $(N', \phi') \in \mathrm{Def}(M)$ is a quasiconformal map. Let $(X, \pi_{N'})$ be a standard covering of N'. Then Theorem 3.5.3 shows that ϕ' lifts to a quasiconformal map $(G_{N'}, \psi') \in \mathrm{Def}_X(G_M)$, i.e. $\pi_{N'} \circ \psi' = \phi' \circ \pi_M$. Now $\mu = \mu_{\phi'}$ if and only if $\mu_{\psi^\mu} = \mu_{\psi'}$ (bijectivity of π_M^*) if and only if $\psi' \circ (\psi^\mu)^{-1} \in \mathrm{Aut}(X)$ (Theorem 3.4.7) if and only if $\phi' \circ \phi_\mu^{-1} \colon N_\mu \to N'$ is a conformal map (Theorem 3.5.3). The claim follows. □

Remark 4.2.6*:*
Let M and N be two Riemann surfaces of $\mathcal{R}_X$. If $X = \hat{\mathbb{C}}$ or $X = \mathbb{C}$, the ideal boundaries of M and N are always empty. If $X = \mathbb{H}$, the ideal boundary may or may not be empty, see Remark 3.6.2. If the ideal boundary happens to be empty, the following remarks have obvious interpretations.

Let $\phi \colon M \to N$ be a quasiconformal map between Riemann surfaces, and let (X, π_M) and (X, π_N) be standard covers of M and N, respectively. Suppose that M and N have ideal boundaries ∂M and ∂N as discussed in Remark 3.3.11, and let ψ be a lift of ϕ along π_M and π_N. Then ψ extends to the boundary of ∂X such that $F_\psi\colon G_M \to G_N$ is an isomorphism, when G_M and G_N are considered to be acting on $X \cup I(G_M)$ and $X \cup I(G_N)$, respectively. This is a consequence of Corollary 3.4.8.

It is readily seen from the definition of the limit set (Remark 2.2.5) that if $w \in L(G_M)$, then $\psi(w) \in L(G_N)$ and $\psi(L(G_M)) = L(G_N)$. Theorem 1.2.21 shows that there exists a unique homeomorphism $\tilde{\phi}\colon M \cup \partial M \to N \cup \partial N$ that is a descension of ψ along (the extensions of) π_M and π_N. As $\tilde{\phi}$ agrees with ϕ in M, the extension $\tilde{\phi}$ maps ∂M homeomorphically onto ∂N. This yields the following result.

Proposition 4.2.7:
Let $(N, \phi) \in \mathrm{Def}(M)$ with ideal boundaries ∂M and ∂N, respectively. Then ϕ has a unique continuous extension $\phi\colon M \cup \partial M \to N \cup \partial N$ that is a homeomorphism.

Definition 4.2.8:
Let $G \in \mathcal{G}_X$ be a standard covering group of X, and let (H, ψ) and (H, ψ') be two quasiconformal elements of $\mathrm{Def}_X(G)$. Then (H, ψ) and (H, ψ') are said to be homotopic rel ∂X, if $F_\psi = F_{\psi'}\colon G \to H$ in $X \cup I(G)$ and $\psi = \psi'$ in ∂X.

Let $(N, \phi) \in \mathrm{Def}(M)$ be a quasiconformal map between Riemann surfaces. The extension of ϕ refers to the unique continuous extension of ϕ to the ideal boundaries of ∂M and ∂N, respectively. As the extension is unique, ϕ will be identified with its extension.

If (N, ϕ) and (N, ϕ') are two quasiconformal maps in $\mathrm{Def}(N)$, they are said to be homotopic rel ∂M, if there exists a homotopy $h_t\colon M \cup \partial M \to N \cup \partial N$ between ϕ and ϕ' rel ∂M.

Remark 4.2.9:
The homotopy rel ∂X and homotopy ∂M defined above define equivalence relations and it is straight-forwardly checked. The next result is the motivation behind the definition of homotopy in Def_X.

Proposition 4.2.10:
Suppose that (N, ϕ) and (N, ϕ') are two quasiconformal maps of $\mathrm{Def}(M)$, and let (X, π_M) and (X, π_N) be standard coverings of M and N, respectively. Let $(G_N, \psi) \in \mathrm{Def}_X(G_M)$ be a lift of ϕ along π_M and π_N. The following are equivalent:

(a) The maps (N, ϕ) and (N, ϕ') are homotopic rel ∂M.

(b) There exists a lift $(G_N, \psi') \in \mathrm{Def}_X(G_M)$ of (N, ϕ') such that (G_N, ψ) and (G_N, ψ') are homotopic rel ∂X.

If the lift exists, it is unique. If $X = \mathbb{H}$, these are equivalent to

(c) There exists a lift $(G_N, \psi') \in \mathrm{Def}_X(G_M)$ of (N, ϕ') such that $\psi' = \psi$ in ∂X.

(d) There exists a lift $(G_N, \psi') \in \mathrm{Def}_X(G_M)$ of (N, ϕ') and a homotopy $(G_N, \psi_t) \in \mathrm{Def}_X(G_M)$ between (G_N, ψ) and (G_N, ψ') rel ∂X.

(e) There exists a homotopy $(N, \phi_t) \in \mathrm{Def}(M)$ rel ∂M between the maps (N, ϕ) and (N, ϕ').

If a lift $(G_N, \psi') \in \mathrm{Def}_X(G_M)$ satisfies any one of Part (b)-(d), it satisfies all of them. Especially, the lift (G_N, ψ') satisfying any one of these properties is unique.

Proof:
If $X = \hat{\mathbb{C}}$ or $\mathbb{C}$, the equivalence of Parts (a) and (b) follows from Theorem 3.5.3. It also shows uniqueness. This means that it is sufficient to consider the case $X = \mathbb{H}$. It is clear that Part (d) implies Part (e) (Theorem 1.2.21 and Theorem 3.5.3). The implications (e) $\Rightarrow$ (a) and (b) $\Rightarrow$ (c) are readily verified. This means that it is sufficient to show that Part (a) implies Part (b) and Part (c) implies Part (d).

If ϕ and ϕ' are homotopic rel ∂M, Theorem 1.2.21 shows that there exists a unique lift ψ' for which $F_\psi = F_{\psi'}\colon G_M \to G_N$ in $X \cup I(G_M)$, and the maps ψ and ψ' agree in $I(G_M)$. If $L(G_M)$ contains two or less points, the continuity of ψ and ψ' in ∂X implies that $\psi = \psi'$ in ∂X. If $L(G_M)$ contains more than two points, the discussion in Remark 2.2.5 implies that it is infinite and it is equal to the closure of the hyperbolic fixed points of ψ and ψ'. As $F_\psi = F_{\psi'}$ in X, Lemma 2.1.3 implies that ψ and ψ' agree on the hyperbolic fixed points of the elements of G_M, hence they agree in $L(G_M)$. Thus Part (a) implies Part (b).

Consider the implication Part (c) $\Rightarrow$ Part (d). The idea of this proof comes originally from [4]. There exists a conformal map $(F_C(G_N), C) \in \mathrm{Def}_X(G_N)$ so that $(F_C(G_N), C \circ \psi)$ and $(F_C(G_N), C \circ \psi')$ are normalized quasiconformal maps of $\mathrm{Def}_X(G_M)$; this follows from Corollary 3.4.8 and the fact that ψ and ψ' agree in ∂X. Let $(H_t, \tilde{\psi}_t) \in \mathrm{Def}_X(G_M)$ be the homotopy between $C \circ \psi$ and $C \circ \psi'$ given by Proposition 4.1.6.

The homotopy $\tilde{\psi}_t^{-1}$ depends continuously on t, hence its restriction $\tilde{\psi}_t^{-1}$ to ∂X is a continuous function of t. The restriction defines a continuous family of quasisymmetric maps [17, Sections I.5.1-I.5.2]. The restriction of $\tilde{\psi}_t^{-1}$ to ∂X can be extended to a normalized quasiconformal map f_t of $\mathcal{F}_X$ via the so-called Douady-Earle extension [3]. The Douady-Earle extension f_t has the following properties: it is equal to $(\tilde{\psi}_t)^{-1}$ in ∂X; the extension depends continuously on t; $(G_M, f_t) \in \mathrm{Def}_X(H_t)$; the isomorphism $F_{f_t \circ \tilde{\psi}_t}$ is the identity isomorphism in $X \cup I(G_M)$; $f_t \circ \tilde{\psi}_t$ is the identity when restricted to ∂X; and $f_0 = f_1$ as $C \circ \psi = \tilde{\psi}_0$ and $C \circ \psi' = \tilde{\psi}_1$ agree in ∂X.

The homotopy ψ_t defined by the formula $C \circ \psi_t = f_0^{-1} \circ f_t \circ \tilde{\psi}_t$ shows the last implication. It is clear from the proof that if $(G_N, \psi') \in \mathrm{Def}_X(G_M)$ satisfies any of the properties discussed here, it satisfies all of them. The uniqueness of the lift is a consequence of Theorem 1.2.21. □

Corollary 4.2.11:
Let M and N be Riemann surfaces from $\mathcal{R}_X$, (X, π_M) and (X, π_N) their standard coverings, $(N, \phi) \in \mathrm{Def}(M)$ quasiconformal, and a lift $(G_N, \psi) \in \mathrm{Def}_X(G_M)$ of (N, ϕ) along π_M and π_N.

Then there exists a quasiconformal map $(N, \phi') \in \mathrm{Def}(M)$ homotopic to (N, ϕ) rel ∂N minimizing the maximal dilatation in the homotopy class of (N, ϕ). The minimizer can be assumed to admit a lift $(G_N, \psi') \in \mathrm{Def}_X(G_M)$ such that ψ and ψ' are normalized by the same Möbius transformation $h \in \mathrm{Aut}(X)$ and (G_N, ψ') minimizes the maximal dilatation in the rel-homotopy class of (N, ϕ).

Proof:
Let (G_N, ψ) be a lift of (N, ϕ). Proposition 4.2.10 and Theorem 3.5.3 show that there is a one-to-one correspondence between the quasiconformal maps $(G_N, \psi') \in \mathrm{Def}_X(G_M)$ homotopic rel ∂X to (G_N, ψ) and the quasiconformal maps $(N, \phi) \in \mathrm{Def}(M)$ homotopic to (N, ϕ) rel ∂M. The corresponce is given by lifts and descensions, which leave the maximal dilatation invariant (Theorem 3.5.3). This means that it is sufficient to find the minimizer in the rel-homotopy class of (G_N, ψ).

The claim can be reduced to the compactness result Proposition 4.1.4: Suppose that (G_N, ψ') and (G_N, ψ) are homotopic rel ∂X. This means that $F_{\psi'} = F_\psi$ in $X \cup \partial X$ and $\psi' = \psi$ in ∂X. Corollary 3.4.8 shows the existence of C and h in $\mathrm{Aut}(X)$ for which $C \circ \psi'$ and $h \circ \psi$ are normalized quasiconformal maps. It is readily checked using the properties of conjugations that

$$F_{C \circ h^{-1}} = F_{(C \circ \psi') \circ (h \circ \psi)^{-1}}$$

in X. Proposition 4.1.3 Part (a) shows that $F_{C \circ \psi'}(G_M)$ and $F_{h \circ \psi}(G_M)$ are standard covering groups of X and by Part (b), the conjugations $F_{C \circ \psi'}$ and $F_{h \circ \psi}$ agree in X. Since $F_{\psi'} = F_\psi$ in X, it follows that $F_C = F_h$ in X.

If $X = \hat{\mathbb{C}}$ or $\mathbb{C}$, this is sufficient to conclude that $(G_N, h^{-1} \circ (C \circ \psi'))$ is homotopic to (G_N, ψ) rel ∂X (Proposition 4.2.10 Part (a)-Part (b)). If $X = \mathbb{H}$, then ψ' and ψ are equal to one another in ∂X, hence C must be equal to h. This means that $(G_N, h^{-1} \circ (C \circ \psi'))$ is homotopic to (G_N, ψ) rel ∂X in this case as well.

Since composing by a conformal map does not change maximal dilatations (Lemma 3.1.3), this implies that to find the minimizer, it is sufficient to consider quasiconformal maps $(G_N, \psi') \in \mathrm{Def}_X(G_M)$ homotopic rel ∂X to (G_N, ψ) that are normalized by the same conformal map as (G_N, ψ).

A minimizer of the maximal dilatation of such a family exists due to Proposition 4.1.4, Theorem 3.5.3, and Proposition 4.2.10: given a sequence of quasiconformal maps approaching the infimum from above, the first claim shows that some subsequence of it converges to a quasiconformal map that induces the same isomorphism between G_M and G_N in X. Theorem 3.5.3 shows that such a map descends to a quasiconformal map between M and N. This suffices for the cases $X = \hat{\mathbb{C}}$ and $\mathbb{C}$ due to Proposition 4.2.10 Part (a) and Part (b) as ∂X is empty. If $X = \mathbb{H}$, then the restrictions of the minimizing sequence to ∂X are the same maps ∂X, hence the limit is also the same map. Then Proposition 4.2.10 Part (c) shows the desired result. The claim follows. □

4.3 Teichmüller spaces

Remark 4.3.1*:*
Now there is enough theory developed to meaningfully study the structure on Teichmüller spaces. The next definition gives four notions of equivalence relations, one for each of the following spaces: the deformation space $\mathrm{Def}_X(G)$ of a standard covering group, the deformation space $\mathrm{Def}(M)$ of a Riemann surface, the Beltrami differentials $B(M)$ on a Riemann surface, and the space of G-invariant Beltrami differentials $B_X(G)$.

The notion of equivalence in $\mathrm{Def}(M)$ is the standard definition found in [17] and [12], and in special cases in [13] and [5] (when the ideal boundary is empty). The notion of equivalence in $\mathrm{Def}_X(G)$ is motivated by Proposition 4.2.10. The idea for the equivalence relation in $B(M)$ arises from Theorem 4.2.5, and the notion of equivalence in $B_X(G)$ is motivated by the bijectivity of the pullback map defined by the standard covering maps as discussed in Remark 3.5.5.

Definition 4.3.2:
Equivalence relations:

(a) Let $G \in \mathcal{G}_X$ and $(H_1, \psi_1), (H_2, \psi_2) \in \mathrm{Def}_X(G)$. Then $(H_1, \psi_1) \simeq (H_2, \psi_2)$ if and only if there exists a conformal map $(H_1, C) \in \mathrm{Def}_X(H_2)$ such that (H_1, C) is homotopic to $(H_1, \psi_1)^{\#} (H_2, \psi_2)$ rel ∂X. The equivalence classes are denoted by $[[(H_1, \psi_1)]]$. The short-hand $[[\psi_1]]$ is also utilized at times.

(b) Let $M \in \mathcal{R}_X$ and $(N_1, \phi_1), (N_2, \phi_2) \in \mathrm{Def}(M)$. Then $(N_1, \phi_1) \sim (N_2, \phi_2)$ if and only if $(N_2, \phi_2)_{\#}(N_1, \phi_1)$ is homotopic to a conformal map $(N_1, c) \in \mathrm{Def}(N_2)$ rel ∂N_1. The equivalence classes are denoted by $[(N_1, \phi_1)]$, and sometimes by $[\phi_1]$.

(c) Two Beltrami differentials μ_1 and μ_2 of $B(M)$ are equivalent under $\sim_B$ if $(H^{\mu_1}, \psi^{\mu_1})$ is homotopic to $(H^{\mu_2}, \psi^{\mu_2})$ rel ∂X and $(N_{\mu_1}, \phi_{\mu_1})$ is homotopic to $(N_{\mu_2}, \phi_{\mu_2})$ rel ∂M. The equivalence class of μ is denoted by $[\mu]_B$.

(d) Two G-invariant Beltrami differentials η_1 and η_2 of $B_X(G)$ are equivalent (denoted by $\sim_{B_X}$) if their inverse images under $\pi^*_{X/G}$ are equivalent elements of $B(X/G)$. The equivalence class of η is $[\eta]_{B_X}$.

Remark 4.3.3:
The goal of this remark is to deduce that the relations are in fact equivalence relations, and see how they are related to one another.

Let $\mu_1, \mu_2 \in B(M)$ and assume that $(H^{\mu_1}, \psi^{\mu_1}) \simeq (H^{\mu_2}, \psi^{\mu_2})$. The idea is to show that this is equivalent to $(N_{\mu_1}, \phi_{\mu_2}) \sim (N_{\mu_2}, \phi_{\mu_2})$ and these are equivalent to $\mu_1 \sim_B \mu_2$. This will take a couple of steps and it will be useful in establishing that $\sim$ and $\simeq$ are equivalence relations.

The definition of $\simeq$ implies that there exists a conformal map (H^{μ_2}, C) of the deformation space $\mathrm{Def}_X(H^{\mu_1})$ that is homotopic to $(H^{\mu_1}, \psi^{\mu_1})^{\#}(H^{\mu_2}, \psi^{\mu_2})$ rel ∂X. As $(H^{\mu_1}, \psi^{\mu_1})^{\#}(H^{\mu_2}, \psi^{\mu_2})$ is a lift of $(N_{\mu_1}, \phi_{\mu_1})_{\#}(N_{\mu_2}, \phi_{\mu_2})$, Proposition 4.2.10 shows that the descension (N_{μ_2}, c) of (H^{μ_2}, C) – given by Theorem 1.2.21 – is homotopic to $(N_{\mu_1}, \phi_{\mu_1})_{\#}(N_{\mu_2}, \phi_{\mu_2})$ rel ∂N_{μ_1}. Theorem 3.5.3 shows that the restriction of (N_{μ_2}, c) to M is a conformal map onto N.

Conversely, if $(N_{\mu_1}, \phi_{\mu_1})_{\#}(N_{\mu_2}, \phi_{\mu_2})$ is homotopic to a conformal map (N_{μ_2}, c) rel ∂N_{μ_1}, Proposition 4.2.10 shows that (N_{μ_2}, c) admits a lift $(H^{\mu_2}, C) \in \mathrm{Def}_X(H^{\mu_1})$ that is homotopic to $(H^{\mu_1}, \psi^{\mu_1})^{\#}(H^{\mu_2}, \psi^{\mu_2})$ rel ∂X.

Proposition 4.1.3 implies that such a C can exists only if $H^{\mu_2} = H^{\mu_1}$, $F_{\psi^{\mu_1}} = F_{\psi^{\mu_2}}$ in X, and $N_{\mu_2} = N_{\mu_1}$. If $X = \hat{\mathbb{C}}$ or $\mathbb{C}$, Proposition 4.1.3 shows that actually $(H^{\mu_1}, \psi^{\mu_1})$ is homotopic to $(H^{\mu_2}, \psi^{\mu_2})$ rel ∂X. If $X = \mathbb{H}$, the fact that ψ^{μ_1} and ψ^{μ_2} are normalized shows that C must fix 0, 1 and ∞ (Proposition 4.2.10 Part (c)). It follows that $C = \mathrm{id}_X$, and $(H^{\mu_1}, \psi^{\mu_1})$ is homotopic to $(H^{\mu_2}, \psi^{\mu_2})$ rel ∂X in this case as well. Proposition 4.2.10 shows that this happens if and only if $(N_{\mu_1}, \phi_{\mu_1})$ is homotopic to $(N_{\mu_2}, \phi_{\mu_2})$ rel ∂M. It follows that $\mu_1 \sim_B \mu_2$. The goal of the second paragraph is achieved.

If $X = \hat{\mathbb{C}}$ or $\mathbb{C}$, it is clear from Theorem 4.2.5 and Proposition 4.2.10 that $(N_1, \phi_1) \sim (N_2, \phi_2)$ if and only if

$$(N_{\mu_{\phi_1}}, \phi_{\mu_{\phi_1}}) \sim (N_{\mu_{\phi_2}}, \phi_{\mu_{\phi_2}})$$
$$(H^{\mu_{\phi_1}}, \psi^{\mu_{\phi_1}}) \simeq (H^{\mu_{\phi_2}}, \psi^{\mu_{\phi_2}}).$$

This happens if and only if $\mu_{\phi_1} \sim_B \mu_{\phi_2}$ as discussed before.

If $X = \mathbb{H}$, then $(N_1, \phi_1) \sim (N_2, \phi_2)$ if only if given lifts (G_{N_1}, ψ_1) and (G_{N_2}, ψ_2) of (N_1, ϕ_1) and (N_2, ϕ_2), respectively, there exists a conformal map C – a lift of the conformal map given by the definition of $\sim$ – such that $\psi_2 = C \circ \psi_1$ in ∂X (Theorem 4.2.5 and Parts (a) and (c) of Proposition 4.2.10). Now

$$\psi^{\mu_{\phi_2}} = \left(\psi^{\mu_{\phi_2}} \circ \psi_2^{-1} \circ C\right) \circ \psi_1$$

in ∂X. As the left-hand side fixes 0, 1, ∞ and $\psi^{\mu_{\phi_2}} \circ \psi_2^{-1} \circ C$ is a Möbius transformation of $\mathrm{Aut}(X)$, the uniqueness part of Corollary 3.4.8 shows that the right-hand side must be equal to $\psi^{\mu_{\phi_1}}$. Proposition 4.2.10 and the beginning part of this remark imply that this happens if and only if $\mu_{\phi_1} \sim_B \mu_{\phi_2}$. It follows that $(N_1, \phi_1) \sim (N_2, \phi_2)$ if and only if $\mu_{\phi_1} \sim_B \mu_{\phi_2}$.

To show that $\simeq$ is an equivalence relation for an arbitrary standard covering group $G \in \mathcal{G}_X$, just consider the Riemann surface $M = X/G$ and the standard covering (X, μ_G). Then it can be checked that $(H_1, \psi_1) \simeq (H_2, \psi_2)$ implies that their descensions $(X/H_1, \phi_1)$ and $(X/H_2, \phi_2)$ are equivalent elements of $\mathrm{Def}(X/G)$. The converse also holds as a consequence of Proposition 4.2.10. This means that $(H_1, \psi_1) \simeq (H_2, \psi_2)$ if and only if $\mu_{\phi_1} \sim \mu_{\phi_2}$. Since being homotopic is an equivalence relation, it is readily checked that $\sim_B$ defines an equivalence relation. It follows that $\sim_B$, $\sim_{B_X}$, $\sim$ and $\simeq$ are equivalence relations. These observations imply the next result.

Proposition 4.3.4:
Characterization of equivalence relations:

(a) Given a quasiconformal map $(N, \phi) \in \mathrm{Def}(M)$, the map $(N_{\mu_\phi}, \phi_{\mu_\phi})_\#(N, \phi)$ is conformal in N_{μ_ϕ}. In particular, $(N, \phi) \sim (N_{\mu_\phi}, \phi_{\mu_\phi})$, and there exists a lift $(G_N, \psi) \in \mathrm{Def}_X(G_M)$ of (N, ϕ) along π_M and π_N such that $(H^{\mu_\phi}, \phi^{\mu_\phi})^\#(G_N, \psi)$ is a Möbius transformation of $\mathrm{Aut}(X)$, and $(G_N, \psi) \simeq (H^{\mu_\phi}, \phi^{\mu_\phi})$.

(b) Given two quasiconformal maps (N, ϕ) and (N', ϕ') of the deformation space $\mathrm{Def}(M)$, then $(N, \phi) \sim (N', \phi')$ if and only if $\mu_{\phi_1} \sim_B \mu_{\phi_2}$.

If $X = \mathbb{H}$, then $(N, \phi) \sim (N', \phi')$ if and only if given a lift $(G_N, \psi) \in \mathrm{Def}_X(G_M)$ of (N, ϕ), there exists a lift $(G_{N'}, \psi') \in \mathrm{Def}_X(G_M)$ of (N', ϕ') and a conformal map $(G_{N'}, C) \in \mathrm{Def}_X(G_N)$ such that $C \circ \psi = \psi'$ in ∂X.

(c) The relations $\sim_B$, $\sim_{B_X}$, $\sim$ and $\simeq$ are equivalence relations.

Proposition 4.3.5 (Quasiconformal invariance)**:**
Let $(H, f), (H', f') \in \mathrm{Def}_X(G)$ and suppose that (F, ψ) and (F, ψ') are homotopic quasiconformal maps of $\mathrm{Def}_X(G)$ rel ∂X. Then $[[H, f)]] = [[(H', f')]]$ if and only if

$$\left[\left[(F, \psi)^\#(H, f)\right]\right] = \left[\left[(F, \psi')^\#(H', f')\right]\right].$$

Similarly, let (N, f) and (N', f') be two quasiconformal maps of $\mathrm{Def}(M)$. Suppose that (M', ϕ) and (M', ϕ') are two quasiconformal elements of $\mathrm{Def}(M)$ homotopic rel ∂M. Then $[(N, f)] = [(N', f')]$ if and only if $[(M', \phi)_\#(N, f)] = [(M', \phi')_\#(N', f')]$.

Proof:
If $X = \hat{\mathbb{C}}$ or $\mathbb{C}$, this is just a direct computation from the definitions using the given assumptions. If $X = \mathbb{H}$, Proposition 4.3.4 and Proposition 4.2.10 imply the claim. □

Definition 4.3.6 (Teichmüller spaces)**:**
Let G be a standard covering group of $\mathcal{G}_X$. The quotient $B_X(G)/{\sim_X}$ is called the Teichmüller space of G-invariant Beltrami differentials and $\mathrm{Def}_X(G)/{\simeq}$ is called the Teichmüller space of G. They are denoted by $\mathcal{B}_X(G)$ and $T_X(G)$.

If M is a Riemann surface in $\mathcal{R}_X$, the quotients $B(M)/{\sim}$ and $\mathrm{Def}(M)/{\sim}$ are called the Teichmüller space of Beltrami differentials of M and the Teichmüller space of M, and they are denoted by $\mathcal{B}(M)$ and $T(M)$, respectively.

Theorem 4.3.7:
The maps $\pi_M^*\colon \mathcal{B}(M) \to \mathcal{B}_X(G_M)$, where $[\mu]_B \mapsto [\pi_M^*(\mu)]_{B_X}$, $S^M\colon \mathcal{B}(M) \to T_X(G_M)$, where $[\mu]_B \mapsto [[(H^\mu, \psi^\mu)]]$, and $P_M\colon \mathcal{B}(M) \to T(M)$, where $[\mu]_B \mapsto \left[(N_\mu, \phi_\mu)\right]$ are well-defined and bijective.

Proof:
The claim is clear from Remark 4.3.3, Proposition 4.3.4, and the bijectivity of π_M^* on the level Beltrami differentials $B(M)$ of M and the G_M-invariant Beltrami differentials $B_X(G_M)$ of X. □

Definition 4.3.8:
The lift operator $\mathcal{L}_M\colon T(M) \to T_X(G_M)$ is defined as $\mathcal{L}_M = S^M \circ P_M^{-1}$. The descension operator $\mathcal{D}_M\colon T_X(G_M) \to T(M)$ is defined by $\mathcal{D}_M = (S^M \circ P_M)^{-1}$.

Let $(N, \phi) \in \mathrm{Def}(M)$. The map $(N, \phi)_*\colon T(M) \to T(N)$ is defined by $[(N', \phi')] \mapsto [(N, \phi)_\#(N', \phi')]$. Furthermore, given $(H, \psi) \in \mathrm{Def}_X(G)$, the map $(H, \psi)^*$ from $T_X(G)$ to $T_X(H)$ is defined as

$$[[(H', \ \psi')]] \mapsto \left[\left[(H, \psi)^\#(H', \psi')\right]\right].$$

Let $[\mu] \in \mathcal{B}(M)$ and $\eta \in [\mu]$. Then $[\mu]^* : T_X(G_M) \to T_X(H^\mu)$ is defined as

$$[[(H, \psi)]] \mapsto (H^\eta, \psi^\eta)^* \, [[(H, \psi)]],$$

and $[\mu]_* : T(M) \to T(N_\mu)$ is defined as $[(N, \phi)] \mapsto \left(N_\eta, \phi_\eta\right)_* [(N, \phi)]$.

Corollary 4.3.9:
The lift operator $\mathcal{L}_M$, descend operator $\mathcal{D}_M$, $(N,\phi)_*$, $(H,\psi)^*$, and the maps $[\mu]^* : T_X(G_M) \to T_X(H^\mu)$ and $[\mu]_* : T(M) \to T(X/H^\mu)$ are well-defined and bijective.

Proof:
The lift and descend operators are well-defined and bijective as compositions of bijective maps. Furthermore, they are inverses of one another by construction. Proposition 4.3.5 shows that the maps $(N,\phi)_*$ and $(H,\psi)^*$ are well-defined. If $\eta \in [\mu]$, then $\eta \sim_B \mu$ by definition. Proposition 4.3.5 implies that the maps $[\mu]^*$ and $[\mu]_*$ are well-defined. The maps $(H,\psi)^\#$ and $(N,\phi)_\#$ are invertible, therefore the maps $(\dots)^*$, $(\dots)_*$, $[\mu]^*$, and $[\mu]_*$ are invertible as well. □

4.4 Teichmüller spaces as metric spaces

Remark 4.4.1*:*
If (L,l) is an element of $\mathrm{Def}_X(G)$ or $\mathrm{Def}(M)$, the maximal dilatation of (L,l), denoted by $K_{(L,l)}$ refers to the maximal dilatation K_l of l. At first, a distance function will be defined for $\mathcal{B}(M)$ and some of its properties will be studied. After that, a notion of distance for the Teichmüller spaces $T_X(G)$ and $T(M)$ will be introduced. After that the maps introduced in Corollary 4.3.9 will turn out to be isometries.

Definition 4.4.2:
Given two Beltrami differentials $[\mu]$ and $[\mu']$ of $\mathcal{B}(M)$, their distance is defined as

$$\begin{aligned} d_{\mathcal{B}(M)}\left([\mu],[\mu']\right) &= \inf\left\{\log K_{(N_\eta,\phi_\eta)} \mid \eta \in B(M),\ [(N_\eta,\phi_\eta)] = [\mu]_* P_M([\mu'])\right\} \\ &= \inf\left\{\log K_{(H^\eta,\psi^\eta)} \mid \eta \in B(M),\ [[(H^\eta,\psi^\eta)]] = [\mu]^* S^M([\mu'])\right\}. \end{aligned}$$

Remark 4.4.3*:*
The two distinct definitions of $d_{\mathcal{B}(M)}$ require some unpacking. Let $\eta \in B(M)$ such that $[(N_\eta,\phi_\eta)] = [\mu]_* P_M([\mu'])$. This equality shows that the equivalence classes of (N_η,ϕ_η) and $(N_{\mu'},\phi_{\mu'}\circ(\phi_\mu)^{-1})$ are the same. This is equivalent to stating that (H^η,ψ^η) and $(H^{\mu'},\psi^{\mu'}\circ(\psi^\mu)^{-1})$ are in the same equivalence class, i.e. $[[(H^\eta,\psi^\eta)]] = [\mu]^* S^M([\mu'])$.

The latter one happens if and only if $H^\eta = H^{\mu'}$ and (H^η,ψ^η) is homotopic to $(H^\mu,\psi^\mu)^\#(H^{\mu'},\psi^{\mu'})$ rel ∂X. This happens if and only if $N_\eta = N_{\mu'}$ and the maps (N_η,ϕ_η) and $(N_\mu,\phi_\mu)_\#(N_{\mu'},\phi_{\mu'})$ are homotopic rel ∂N_μ. It follows that the infimum is actually taken in the rel-homotopy classes of $(H^\mu,\psi^\mu)^\#(H^{\mu'},\psi^{\mu'})$ and

$(N_\mu, \phi_\mu)_\#(N_{\mu'}, \phi_{\mu'})$, respectively. Theorem 3.5.3 shows that ψ^η and ϕ_η have the same maximal dilatation, hence these two notions of distances in $\mathcal{B}(M)$ agree.

Proposition 4.4.4 (Extremal representatives)**:**
Let M be a Riemann surface, and let μ and μ' be two Beltrami differentials of $B(M)$. Then there exists a Beltrami differential $\eta \in B(N_\mu)$ such that

$$d_{\mathcal{B}(M)}\left([\mu], [\mu']\right) = \log K_{(N_\eta, \phi_\eta)} = \log K_{(H^\eta, \psi^\eta)},$$

where $N_\eta = N_{\mu'}$, $H^\eta = H^{\mu'}$, and the maps (N_η, ϕ_η) and $(N_\mu, \phi_\mu)_\#(N_{\mu'}, \phi_{\mu'})$ are homotopic rel ∂N_μ, and (H^η, ψ^η) is homotopic to $(H^\mu, \psi^\mu)^\#(H^{\mu'}, \psi^{\mu'})$ rel ∂X.

Proof:
After the discussion in Remark 4.4.3, the claim follows from Corollary 4.2.11 and the bijectivity of $\pi^*_{N_\mu} \colon B(N_\mu) \to B_X(H^\mu)$. □

Corollary 4.4.5:
The Teichmüller space of Beltrami differentials $\left(\mathcal{B}(M), d_{\mathcal{B}(M)}\right)$ is a metric space.

Proof:
The latter part of Lemma 3.1.4, i.e. the fact that $K_{a \circ b} \leq K_a K_b$ for quasiconformal maps, implies that $d_{\mathcal{B}(M)}$ satisfies the triangle inequality. It also shows that that the metric is symmetric. Proposition 4.4.4 shows that $d_{\mathcal{B}(M)}$ is a distance: The distance between two equivalence classes of Beltrami differentials $[\mu]$ and $[\mu']$ is zero if and only if the extremal representative $(H^{\mu'}, \psi^\eta) \in \mathrm{Def}_X(H^\mu)$ is conformal. Since $(H^{\mu'}, \psi^\eta)$ is normalized, this can be happen if and only if ψ^η is the identity. It follows that $[\mu]$ and $[\mu']$ are distance zero apart if and only if (H^μ, ψ^μ) and $(H^{\mu'}, \psi^{\mu'})$ are homotopic rel ∂X, i.e. $[\mu] = [\mu']$. □

Definition 4.4.6 (Teichmuller metrics)**:**
Let G be a standard covering group in $\mathcal{G}_X$. The map $d_{T_X(G)} \colon T_X(G) \times T_X(G) \to \mathbb{R} \cup \{\infty\}$ defined as

$$\left([[\psi]], [[\psi']]\right) \mapsto \inf\left\{\log K_{(H',f')^\#(H,f)} \mid ((H, f), (H', f')) \in [[\psi]] \times [[\psi']]\right\}$$

is the Teichmüller distance between $[[\psi]]$ and $[[\psi']]$.

Let M be a Riemann surface in $\mathcal{R}_X$. The map $d_{T(M)} \colon T(M) \times T(M) \to \mathbb{R} \cup \{\infty\}$ defined as

$$\left([\phi], [\phi']\right) \mapsto \inf\left\{\log K_{(N',f')_\#(N,f)} \mid ((N, f), (N', f')) \in [\phi] \times [\phi']\right\}$$

is the Teichmüller distance between the equivalence classes $[\phi]$ and $[\phi']$.

Remark 4.4.7:
The discussion in Remark 4.4.3 implies that

$$d_{T(M)}\left(P_M([\mu]), P_M([\mu'])\right) \leq d_{\mathcal{B}(M)}\left([\mu], [\mu']\right)$$
$$d_{T(G_M)}\left(S^M([\mu]), S^M([\mu'])\right) \leq d_{\mathcal{B}(M)}\left([\mu], [\mu']\right).$$

The inequalities "$\geq$" are corollaries of Remark 4.3.3 and Lemma 3.1.3: every quasiconformal map of the deformation space $\mathrm{Def}_X(G_M)$ is of the form (H^μ, ψ^μ) modulo a conformal map (similar result holds for (N_μ, ϕ_μ)). Then the inequality "$\geq$" follows from Lemma 3.1.3 – composing by conformal maps from left and right leave maximal dilatations unchanged.

As P_M and S^M are bijective due to Theorem 4.3.7 and $\mathcal{B}(M)$ is a metric space as shown in Corollary 4.4.5, it follows that $d_{T(M)}$ and d_G define distance functions, and that P_M and S^M are isometries.

Proposition 4.4.4 implies that the distance between two equivalence classes $[[\psi]]$ and $[[\psi']]$ of $T_X(G)$ is realized by a normalized quasiconformal map $(F_{\psi'}(G), f)$ contained in $\mathrm{Def}_X\left(F_\psi(G)\right)$. These observations, together with some corollaries, are gathered to the next result.

Corollary 4.4.8:
The Teichmüller distances are distances and the maps

$$P_M \colon \mathcal{B}(M) \to T(M),\ S^M \colon \mathcal{B}(M) \to T_X(M),$$
$$\mathcal{L}_M \colon T(M) \to T_X(G_M),\ \mathcal{D}_M \colon T_X(G_M) \to T(M),$$

and the maps $(N, \phi)_*$ and $(H, \psi)^*$, together with $[\mu]^*$ and $[\mu]_*$ are isometries. Furthermore, the Teichmüller spaces are complete and geodesic.

Proof:
The Teichmüller distances are distance functions due to Remark 4.4.7, therefore it is sufficient to consider the rest of the claim. Consider the claim about the isometries. The first two isometries follows from the discussion in Remark 4.4.7, and the lift operator $\mathcal{L}_M$ and the descension operator $\mathcal{D}_M$ are isometries as compositions of isometries (Definition 4.3.8). The last four cases follow from Corollary 4.3.9 and the observation that $(f \circ g^{-1}) \circ (h \circ g^{-1})^{-1} = f \circ h^{-1}$.

As isometries preserve completeness and geodesics, it is sufficient to show that $T_X(G)$ is complete and geodesic. Let $[[\psi_n]]$ be a Cauchy sequence in $T_X(G)$. It can be assumed without loss of generality that $\psi_n \in \mathcal{F}_X$ is a normalized quasiconformal map – every equivalence class has such a representative. For every n, let $f_n \in \mathcal{F}_X$ such that $(F_{f_n}(G), f_n)$ is homotopic to the quasiconformal map $(F_{\psi_n}(G), \psi_n)^\#(F_{\psi_{n+1}}(G), \psi_{n+1})$ rel ∂X and

$$d_{T_X(G)}\left([[\psi_n]], [[\psi_{n+1}]]\right) = \log K_{f_n}.$$

By passing to a subsequence and reindexing, it can be assumed that $k_{f_n} \leq 2^{-n}$. This implies that $K_{f_n} \leq 1 + 2^{2-n}$. Let $g_1 = \psi_1$ and $g_n = f_{n-1} \circ g_{n-1}$. Induction shows that $(F_{g_n}(G), g_n)$ is homotopic to $(F_{\psi_n}(G), \psi_n)$ rel ∂X. The sum over 2^{2-k} is finite, therefore by induction and Lemma 3.1.4

$$K_{g_n} \leq K_{\psi_1} \prod_{k=1}^{\infty} K_{f_k} \leq K_{\psi_1} \prod_{k=1}^{\infty} \left(1 + 2^{2-k}\right) < \infty.$$

This means that $K = \sup K_{g_n}$ is finite. The chain rule of Beltrami differentials Lemma 3.1.3 implies that

$$\begin{aligned} \left\|\mu_{g_{n+1}} - \mu_{g_n}\right\|_\infty &= \left\|1 - \overline{\mu_{g_n}} \mu_{g_{n+1}}\right\|_\infty \left\|\mu_{g_{n+1} \circ g_n^{-1}}\right\|_\infty \\ &\leq 2 \cdot \left\|\mu_{f_n}\right\|_\infty \leq 2^{1-n}. \end{aligned}$$

As the sum $\sum_{n=1}^{\infty} \left\|\mu_{g_{n+1}} - \mu_{g_n}\right\|_\infty$ is finite, the sequence $\{\mu_{g_n}\}_{n=1}^{\infty}$ is a Cauchy sequence. Proposition 4.1.5 shows that there exists $(F_g(G), g) \in \mathrm{Def}_X(G)$ with

$$\begin{aligned} d_{T_X(G)}\left([[\psi_n]], [[g]]\right) = d_{T_X(G)}\left([[g_n]], [[g]]\right) &\leq \log K_{g \circ g_n^{-1}} \\ &= \log \frac{1 + k_{g \circ g_n^{-1}}}{1 - k_{g \circ g_n^{-1}}} \to \log 1 = 0. \end{aligned}$$

This means that $[[g]]$ is the limit of the Cauchy sequence. It follows that the Teichmüller space $T_X(G)$ is complete.

Consider the claim about $T_X(G)$ being geodesic. Since $[[\ ]]^*$ are isometries, it is sufficient to show that every point $[[\psi']] = [[(H', \psi')]]$ can be connected to $[[\mathrm{id}_X]] = [[(G, \mathrm{id}_X)]]$ via a geodesic. Let $(H, \psi) \in \mathrm{Def}_X(G)$ be a normalized quasiconformal map realizing the distance between $[[\psi']]$ and $[[\mathrm{id}_X]]$. Proposition 4.1.6 and reverse triangle inequality imply the claim: the homotopy constructed in Proposition 4.1.6 is a unit speed geodesic between $[[\mathrm{id}_X]]$ and $[[\psi']]$. □

Remark 4.4.9:
It is not clear from Proposition 4.4.4 when an extremal representative is unique. Some clarification should be made. Uniqueness refers to the following property: If the infimum in Proposition 4.4.4 is reached by η and η' with $\eta \sim_B \eta'$, then $\eta = \eta'$. It is clear from the proof of Corollary 4.4.8 that uniqueness holds if and only if the Teichmüller spaces are uniquely geodesic. Moreover, this is equivalent to the statement that there exists a unique geodesic between the origin and any other equivalence class. This is not true in general, extremal representatives need not be unique, see [5, Example 8.1.3].

The uniqueness of extremal representatives can be established in several cases. The proofs do not fit the scope of this work. However, here are some known results: Uniqueness of extremal representatives can be established when the

Teichmüller space is finite dimensional – see Theorem 8 of Section 6.7 and Sections 6.1-6.3 of [8]. This includes hyperbolic Riemann surfaces of finite analytic type (g, n) for $3g - 3 + n > 0$ and $(g, n) = (0, 3)$ [8, Section 1.11]. It will be shown in the next section that every such surface is quasiconformally equivalent to every other one of the same type. This means that their Teichmüller spaces are isometric.

The non-hyperbolic Riemann surfaces of finite analytic type will also be characterized – the uniqueness of extremal representatives, in the sense above, also holds. The uniqueness in the case $g = 1$ and $n = 0$ follows from the results of [8, Section 2.2], specifically Theorem 2. It is also shown in some detail in [17, Chapter V Section 6.4]. More on that in the next section. Uniqueness of an extremal representative of an equivalence class can be established in a more general setting as well, even if the Teichmüller space is not finite dimensional, see Theorem 9 of Section 6.8 of [8].

4.5 Examples of Teichmüller spaces

Remark 4.5.1*:*
Corollary 4.4.8 shows that quasiconformally equivalent Riemann surfaces have isometric Teichmüller spaces. This means that it is sufficient to characterize Teichmüller spaces for a smaller class of Riemann surfaces. Additionally, Corollary 4.4.8 gives a way to conclude when two Riemann surfaces cannot be quasiconformally equivalent if their Teichmüller spaces is known. As discussed in Remark 4.2.1, another way to see whether two Riemann surfaces can be quasiconformally equivalent is to study their Gaussian curvature induced by the natural structure as in Theorem 3.3.9; Corollary 3.3.10 shows that quasiconformal maps do not change the curvature of these structures.

By concentrating on the level of Def_X, a necessary condition that two Riemann surfaces can be quasiconformally equivalent is that their standard covering groups are isomorphic, where the isomorphism maps parabolic elements to parabolic elements and hyperbolic elements to hyperbolic elements (Theorem 3.5.3 and Lemma 2.1.3). In particular, the ideal boundaries of the Riemann surfaces must be homeomorphic on the level of $I(G_M)$ and ∂M; this follows from Proposition 4.2.7 and Proposition 1.2.17.

As an example, consider Riemann surfaces with cyclic non-trivial standard covering groups. There are essentially three different cases as shown in Remark 3.6.2. First of all, if a Riemann surface is quasiconformally equivalent to the punctured plane, its Gaussian curvature is zero, whereas the punctured disk and annuli have negative curvature. Furthermore, the Riemann surfaces quasiconformally equivalent to annuli have ideal boundary homeomorphic to a disjoint union of circles, whereas the ideal boundary of surfaces quasiconformally equivalent to the punctured disk is homeomorphic to a circle. Thus there are (at least) three families of conformal structures on a topological cylinder of dimension two.

A question arises whether two Riemann surfaces are quasiconformally equivalent if and only if they are conformally equivalent. If a Riemann surface has a singleton Teichmüller space, the answer turns out to be yes as seen later in this section. The "only if" direction does not hold – consider for example annuli $A(1, r)$ and $A(1, r')$ with $r \neq r'$. It is observed in the next subsection that any two such annuli are quasiconformally equivalent, but two such annuli are conformally equivalent if and only if $r = r'$. Furthermore, even though the Teichmüller space of a punctured disk is non-trivial, any Riemann surface quasiconformally equivalent to a punctured disk is conformally equivalent to it.

Another question arises whether two Riemann surfaces of finite analytic type (g, n) and (g', n') are quasiconformally equivalent if and only if $g = g'$ and $n = n'$. Neither direction is clear initially, but it turns out that both directions hold. Some work needs to be done to show this. First some easier examples.

Simple examples

Lemma 4.5.2:
Let M be conformally equivalent to either $\hat{\mathbb{C}}$, $\mathbb{C}$, or the punctured plane. Then the Teichmüller space $T(M)$ is a singleton.

Proof:
Let G_M be the standard covering group of M. In the first two cases it is trivial and in the last case it is generated by T_1. Every equivalence class of $T_X(G_M)$ has a normalized representative, hence Proposition 4.1.3 shows that the normalized representative induces an automorphism of G_M, more specifically the automorphism is equal to F_{id}. Thus the equivalence class contains id by Proposition 4.3.4 and Proposition 4.1.3.

This means that for $\hat{\mathbb{C}}$, $\mathbb{C}$, and the punctured plane, the Teichmüller space $T_X(G_M)$ is a singleton. Corollary 4.3.9 shows that $T_X(G_M)$ and $T(M)$ are bijective, hence the claim follows. □

Remark 4.5.3:
It is clear from the proof for the punctured plane and Proposition 4.1.3 that if M is a Riemann surface that is quasiconformally equivalent to the punctured disk, then M is actually conformally equivalent to it. This does not mean that the Teichmüller space of a punctured disk is a singleton – the ideal boundary of the punctured disk is not empty.

Remark 4.5.4*:*
Suppose that M is conformally equivalent to $\mathbb{H}$. Then $T(M)$ is isometric to $T_{\mathbb{H}}(\{\mathrm{id}_{\mathbb{H}}\})$, which is bijective with the space of normalized quasisymmetric maps of the extended real line: This is clear due to Proposition 4.2.10, the fact that the restriction of a quasiconformal map $\psi\colon \mathbb{H} \to \mathbb{H}$ to the extended real line defines a quasisymmetric map, every equivalence class of $T_{\mathbb{H}}(\{\mathrm{id}_{\mathbb{H}}\})$ is represented by a normalized quasiconformal map (Proposition 4.3.4), and every normalized quasisymmetric map extends to a quasiconformal map $\psi\colon \mathbb{H} \to \mathbb{H}$ via the Douady-Earle extension discussed in Proposition 4.2.10 and in [3].

The space $T_{\mathbb{H}}(\{\mathrm{id}_{\mathbb{H}}\})$ is called the universal Teichmüller space. It is clear that whenever $G \in \mathcal{G}_{\mathbb{H}}$ is a standard covering group, the space $T_{\mathbb{H}}(G)$ is a subset of the universal Teichmüller space. In fact, it turns out that the inclusion map is a topological embedding, but not an isometry in general as shown in Chapter III Section 3 of [17].

Consider the following examples: If $M = \mathbb{D} \setminus \{0\}$, every normalized quasisymmetric map of the extended real line satisfying $g(z+n) = g(z) + n$ for every $z \in \mathbb{R}$ and $n \in \mathbb{Z}$ defines an equivalence class of $T(G_M)$ as a consequence of Proposition 4.1.3. If g and g' are two different normalized quasisymmetric maps as above, they define distinct equivalence classes in the Teichmüller space of $T(G_M)$. Furthermore, every equivalence class of $T(G_M)$ can be represented by such a normalized quasisymmetric map.

For $r, r' \in (1, \infty)$, let $\psi_{r,r'}(z) = |z|^{K-1} z$ with $K = \frac{\log r'}{\log r}$. It is a normalized quasiconformal map of $\mathcal{F}_{\mathbb{H}}$ and its maximal dilatation is $\max\{K, K^{-1}\}$ [1, p. 29]. The discussion in Remark 3.6.2 implies that

$$\left(\langle g_{\lambda_{r'}} \rangle, \psi_{r,r'}\right) \in \mathrm{Def}_{\mathbb{H}}(\langle g_{\lambda_r} \rangle),$$

therefore the annuli $A(1, r)$ and $A(1, r')$ are quasiconformally equivalent. The quasiconformal map $\psi_{r,r'}$ is an extremal representative of its equivalence class since every other quasiconformal map between $A(1, r)$ and $A(1, r')$ has a bigger maximal dilatation [5, Theorem 2.2.1].

Riemann surfaces of finite analytic type

Remark 4.5.5:
The next topic is the Teichmüller space of tori of genus one. For $t, t' \in \mathbb{C} \setminus \mathbb{R}$, let $u_{t,t'}(z) = \lambda(z + \mu\bar{z})$ be the unique $\mathbb{R}$-linear map fixing 0, 1 and mapping t to t'. The determinant of $u_{t,t'}$ is $|\lambda|^2 (1 - |\mu|^2)$. As $\{1, t\}$ and $\{1, t'\}$ are $\mathbb{R}$-linear basis of $\mathbb{C}$, it follows that $|\lambda| \neq 0$ and $|\mu| \neq 1$ as $u_{t,t'}$ is invertible.

The map $u_{t,t'}$ is sense-preserving if and only if $|\mu| < 1$, which can be seen by considering the sign of the determinant of $u_{t,t'}$. This is equivalent to the inequality $0 < \operatorname{Im}(t) \operatorname{Im}(t')$. A direct computation shows that the maximal dilatation of $u_{t,t'}$ satisfies $\log K_{u_{t,t'}} = d_{\mathbb{D}}(0, \mu) = d_{\mathbb{H}}(t, t')$, where the distances are as in Remark 3.3.5.

Lemma 4.5.6 (Existence of extremal maps between tori)**:**
Let $t \in \mathbb{H}$ and $\psi\colon \mathbb{C} \to \mathbb{C}$ be a normalized quasiconformal map satisfying $\psi \circ T_{n_1+n_2 t} = T_{n_1+n_2 t'} \circ \psi$. Then $t' = \psi(t) \in \mathbb{H}$, $F_{u_{t,t'}} = F_\psi$, and $K_{u_{t,t'}} \leq K_\psi$.

Proof:
Let ψ and $u := u_{t,t'}$ be as in the claim. Define the set $Q_t := [0, 1] + t\,[0, 1]$. As u is $\mathbb{R}$-linear, it is clear that $F_\psi = F_u$. The periodicity of ψ and u combined with the compactness of Q_t implies that

$$L := \|\psi - u\|_{\infty,\mathbb{C}} = \|\psi - u\|_{\infty, Q_t} < \infty.$$

Observe that $\psi(t) = t'$ follows from $\psi(t) = \psi \circ T_t(0) = T_{t'} \circ \psi(0) = t'$. For each $k \in \mathbb{N}$ define $\psi_k(z) = \frac{\psi(kz)}{k}$, and note that $\psi_k \circ T_{n_1+n_2 t} = T_{n_1+n_2 t'} \circ \psi_k$ and $\mu_{\psi_k}(z)$ is equal to $\mu_\psi(kz)$. The first equality is a direct computation using the corresponding assumption of ψ. The second one follows from Lemma 3.1.3 and it implies $K_{\psi_k} = K_\psi$. The periodicity of ψ and u, in addition to u being $\mathbb{R}$-linear, implies that

$$\|\psi_k - u\|_{\mathbb{C},\infty} = \frac{1}{k}\,\|\psi - u\|_{Q_t,\infty} = \frac{L}{k}.$$

The limit $k \to \infty$ shows that ψ_k converge to u uniformly in $\mathbb{C}$. This implies that u is a quasiconformal map and $K_u \leq K_\psi$ (Proposition 4.1.4). In particular, u is orientation-preserving. Remark 4.5.5 shows that this is equivalent to $t' \in \mathbb{H}$. □

Remark 4.5.7:
As a reminder, the group $G_{1,s}$ refers to the group generated by the translations T_1 and T_s. Since Remark 4.5.5 and Lemma 4.5.6 hold, it is clear that the Teichmüller space $T_X(G_{1,t})$ (and $T(X/G_{1,t})$) is isometric to $\mathbb{H}$ for any $t \in \mathbb{H}$ as the maps $(G_{1,s'}, u_{s,s'}) \in \operatorname{Def}_X(G_{1,s})$ are normalized extremal representatives of their equivalence classes with maximal dilatations $\exp(d_{\mathbb{H}}(s, s'))$, see Proposition 4.4.4. The discussion in Remark 4.4.9 implies that it is the unique normalized extremal representative of its equivalence class.

Every torus of genus one is quasiconformally equivalent to every other torus of genus one. This follows from Remark 3.6.2 and the fact that $u_{t,t'}$ is a quasiconformal map whenever $t, t' \in \mathbb{H}$. Furthermore, this holds true for every family of Riemann surfaces studied so far. The examples include the non-hyperbolic Riemann surfaces of finite analytic type (g, n).

Consider a Riemann surface of finite analytic type (g, n) (hyperbolic or non-hyperbolic). It is readily seen that an orientation-preserving homeomorphism between Riemann surfaces of type (g, n) and (g', n') extends to an orientation-preserving homeomorphism between the compact Riemann surfaces of genus g and g' in such a way that $g = g'$ and $n = n'$. Such a homeomorphism is homotopic to a diffeomorphism, where the homotopy can be taken to fix the punctures n. The proof of the existence of such a diffeomorphism consists of the following facts. The following gives a rough idea behind the proof.

Firstly, the Douabdy-Earle extension [3] shows that a homeomorphism mapping the unit circle onto itself extends to a diffeomorphism of the disk onto itself. This implies that a homeomorphism of the closed unit disk onto itself mapping the unit circle onto itself is homotopic to a diffeomorphism where the homotopy fixes the unit circle.

Secondly, given an orientation-preserving homeomorphism of the closed unit disk onto itself which fixes the origin and is a diffeomorphism of the punctured closed disk onto itself, there exists a diffeomorphism of the closed unit disk onto itself agreeing with the original map at the origin and at the unit circle. The maps are homotopic where the homotopy can be chosen in such a way that it fixes the origin and the unit circle. Such a diffeomorphism can be constructed using a slight modification of Alexander's trick, where the modification uses the fact that two orientation-preserving diffeomorphisms of the unit circle onto itself are homotopic via h_t, where the map $t \mapsto h_t(z)$ is smooth and the maps $z \mapsto h_t(z)$ are orientation-preserving diffeomorphisms.

Given these two facts, the claim itself relies on covering the compact Riemann surface of genus g with appropriate coordinate balls in two steps. Firstly, a homotopy between the given homeomorphism and another homeomorphism is constructed. The latter homeomorphism is a diffeomorphism outside the n punctures and homotopic to the original homeomorphism in such a way that the homotopy fixes the n punctures. This can be achieved using the first fact. Secondly, given the newly-obtained homeomorphism, there exists a diffeomorphism that agrees with the new homeomorphism outside small neighbourhoods of the punctures and a homotopy between these maps that fixes the punctures. The diffeomorphism and the homotopy can be constructed using the second fact. Then the original homeomorphism is homotopic to a diffeomorphism, where the homotopy fixes the punctures. Since an orientation-preserving diffeomorphism between compact Riemann surfaces is a quasiconformal map, the claim is done.

The claim of the third paragraph of this remark and the classification of compact surfaces discussed in Remark 3.6.2 implies that two Riemann surfaces of $(g, n = 0)$ and $(g', n' = 0)$ are quasiconformally equivalent if only if the corresponding compact Riemann surfaces of genus g and g' are quasiconformally equivalent. The same holds for $n = n' > 0$: the case $n = 0 = n'$ implies that it is sufficient to construct a quasiconformal map from a Riemann surface of genus g onto itself mapping any given set of cardinality n to any other other set with the same cardinality.

Fix a Riemann surface M of genus g and two sets P and Q with cardinality $n > 0$. There are two cases. If $g = 0$, by taking advantage of Möbius transformations, it is sufficient to show that there exists a normalized quasiconformal map of the Riemann sphere mapping a set $P \subset \mathbb{C}$ of cardinality n to a set $Q \subset \mathbb{C}$ with the same cardinality. If $g > 0$, by the classification of compact surfaces, there exists a $4g$-gon Ω that represents the Riemann surface of genus g in the sense of classification of compact surfaces, see [14, Example 6.13]. It can be assumed without loss of generality that the punctures P and Q are contained in the interior of Ω. Then it is sufficient to show that there exists an orientation-preserving homeomorphism mapping P onto Q and fixing the boundary of Ω.

These statements can be deduced from the following fact: given an open convex subset Ω of $\mathbb{C}$ and two sets P and Q of cardinalities n contained in Ω, there exists an orientation-preserving diffeomorphism mapping P to Q, the set Ω onto itself, and fixing the complement of Ω. This can be shown using the theory of smooth vector fields – the basic results of [15, Chapter 9] suffices. The idea is illustrated in Figure 4.1: given two points z_1 and z_2, a set P contained in the line between z_1 and z_2 positive distance apart from the endpoints, and a small open neighbourhood U of the line $[z_1, z_2]$, there exists an orientation-preserving diffeomorphism mapping z_1 to z_2 which fixes the set P and the complement of U.

Remark 4.5.8:
Remark 4.5.7 shows that two surfaces of a fixed type (g, n) are quasiconformally equivalent if and only if the corresponding compact Riemann surfaces of genus g are quasiconformally equivalent. In this case, their Teichmüller spaces are isometric. The extremal representatives are of a very special form, which allows the computation of the dimension of the Teichmüller space for these spaces. As discussed in Remark 4.4.9, the Teichmüller spaces of such surfaces are uniquely geodesic. It is shown in Section 6.7 of [8] that if $3g - 3 + n > 0$ or $(g, n) = (0, 3)$, the Teichmüller space of a Riemann surface of type (g, n) is homeomorphic to $\mathbb{R}^{6g-6+2n}$. If $g = 0$ and $n < 3$, then it was also already shown that the Teichmüller space is a singleton. If $g = 1$ and $n = 0$, the Teichmüller space is isometric to the hyperbolic plane. This means that the Teichmüller spaces of Riemann surfaces of type (g, n) are characterized.

Theorem 3.3.3 shows that if $(N, \phi) \in \mathrm{Def}(M)$, there exists a conformal structure

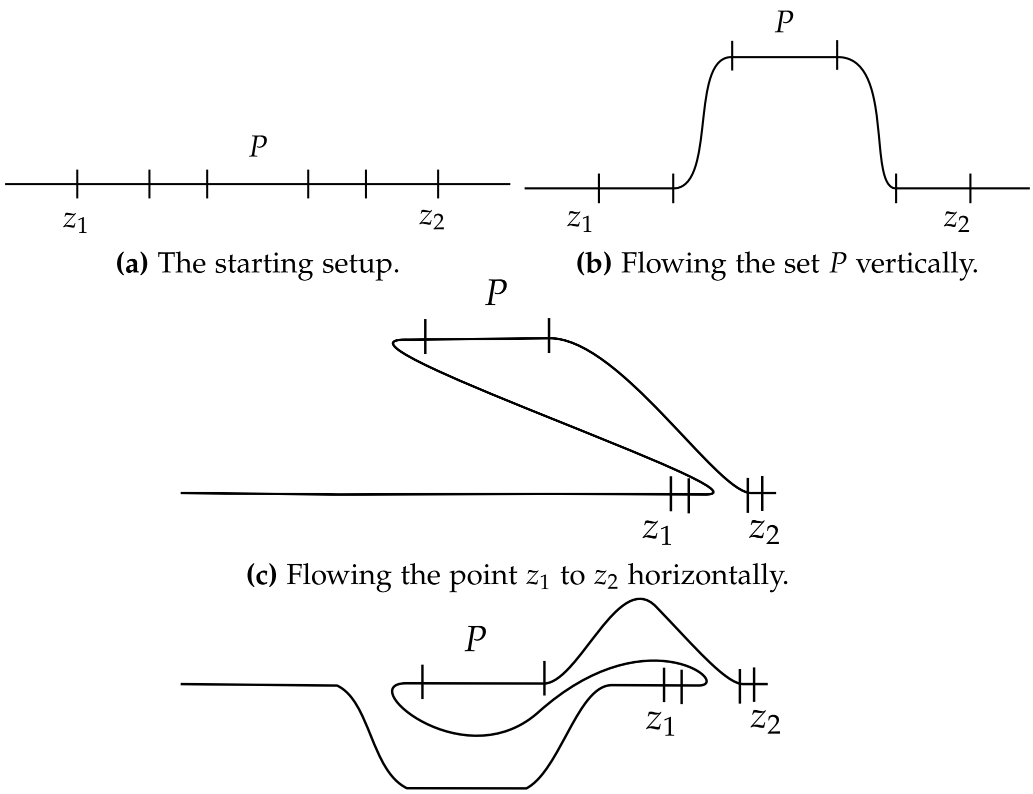

(a) The starting setup.

(b) Flowing the set P vertically.

(c) Flowing the point z_1 to z_2 horizontally.

(d) Returning the set P to its original position.

Figure 4.1. Illustration of the discussion in Remark 4.5.7. Given two points z_1 and z_2 from $\mathbb{C}$, a set P contained in the open interval (z_1, z_2), and a small neighbourhood U of the line segment $[z_1, z_2]$, there exists an orientation-preserving diffeomorphism mapping z_1 to z_2, fixing the set P, and equal to the identity outside the complement of U.

on M that makes ϕ a conformal map and this structure is unique. The results of this section show that for Riemann surfaces of finite analytic type (g, n), every conformal structure on them is obtained this way. This is not true in general. Consider for example topological cylinders (the punctured plane, annuli, and the punctured disk), the topological disks (the Euclidean plane and the hyperbolic plane), or a more exotic example 6.1.2 of [12], where a topological surface is constructed which has uncountably many distinct families of quasiconformal structures. This means that a topological surface can have different conformal structures with wildly different Teichmüller spaces as seen readily from the case of topological cylinders.

Bibliography

[1] Kari Astala, Tadeusz Iwaniec and Gaven Martin. *Elliptic Partial Differential Equations and Quasiconformal Mappings in the Plane*. Princeton University Press, 2009.

[2] Alan F. Beardon. *The Geometry of Discrete Groups*. Springer-Verlag New York Inc., 1983.

[3] Adrien Douady and Clifford J. Earle. 'Conformally natural extension of homeomorphisms of the circle'. In: *Acta Math.* 157 (1986), pp. 23–48. DOI: 10.1007/BF02392590. Available at: http://dx.doi.org/10.1007/BF02392590.

[4] Clifford J. Earle and Curt McMullen. 'Quasiconformal isotopies'. In: *Holomorphic functions and moduli, Vol. I (Berkeley, CA, 1986)*. Volume 10. Math. Sci. Res. Inst. Publ. Springer, New York, 1988, pp. 143–154. DOI: 10.1007/978-1-4613-9602-4_12. Available at: http://dx.doi.org/10.1007/978-1-4613-9602-4_12.

[5] Alastair Fletcher and Vladimir Markovic. *Quasiconformal Maps and Teichmüller Theory*. Oxford University Press, 2007.

[6] Otto Forster. *Lectures on Riemann Surfaces*. Springer-Verlag New York Inc., 1981.

[7] Sylvestre Gallot, Dominique Hulin and Jacques Lafontaine. *Riemannian Geometry*. Edition 2. Springer-Verlag Berlin Heidelberg, 1990.

[8] Frederick P. Gardiner. *Teichmüller Theory and Quadratic Differential*. John Wiley & Sons, Inc., 1987.

[9] John B. Garnett and Donald E. Marshall. *Harmonic Measure*. Cambridge University Press, 2005.

[10] Victor Guillemin and Alan Pollack. *Differential Topology*. Prentice-Hall, Inc., Englewood Cliffs, New Jersey, 1974.

[11] Allen Hatcher. *Algebraic Topology*. Available at: https://www.math.cornell.edu/~hatcher/AT/AT.pdf (visited on 16/07/2017).

[12] John H. Hubbard. *Teichmüller Theory and Applications to Geometry, Topology, and Dynamics*. Volume 1. Matrix Editions, 2006.

[13] Y. Imayoshi and M. Taniguchi. *An Introduction to Teichmüller Spaces*. Springer-Verlag Tokyo, 1992.

[14] John M. Lee. *Introduction to Topological Manifolds*. Edition 2. Springer Science+Business Media LLC, 2011.

[15] John M. Lee. *Introduction to Smooth Manifolds*. Edition 2. Springer Science+Business Media New York, 2013.
[16] Joseph Lehner. *Discontinuous Groups and Automorphic Functions*. American Mathematical Society, 1964.
[17] Olli Lehto. *Univalent Functions and Teichmüller Spaces*. Springer-Verlag New York Inc., 1987.