SJÄLVSTÄNDIGA ARBETEN I MATEMATIK

MATEMATISKA INSTITUTIONEN, STOCKHOLMS UNIVERSITET

Introduction to Hyperbolic Geometry and Fuchsian Groups

av

Lydia Wegerman Gemzöe

2021 - No K35

MATEMATISKA INSTITUTIONEN, STOCKHOLMS UNIVERSITET, 106 91 STOCKHOLM

Introduction to Hyperbolic Geometry and Fuchsian Groups

Lydia Wegerman Gemzöe

Självständigt arbete i matematik 15 högskolepoäng, grundnivå

Handledare: Gregory Arone

2021

Abstract

This thesis is an introduction to hyperbolic geometry and Fuchsian groups. We will introduce the Poincaré models of the hyperbolic plane and give a matrix representation of the group of hyperbolic isometries. A Fuchsian group is a discrete group of orientation-preserving hyperbolic isometries. We will give a definition of a fundamental domain for a Fuchsian group and describe the relation between Fuchsian groups and hyperbolic tessellations. One of the main results of this work is the Poincaré Polygon Theorem, which states that given a hyperbolic polygon we can find, provided that certain conditions are met, a Fuchsian group which has this polygon as a fundamental domain.

Acknowledgements

My sincere thanks to my supervisor Gregory Arone for inspiring me to study this topic, for illuminating to me unknown paths, and for allowing this thesis to grow organically; to Sven Raum for inducing its final fairy form; and to my dear friend and colleague Claudia Kalhori for her keen eye for detail, and for sharing my unhealthy obsession with the aesthetics of spacing.

1 Introduction

There are many reasons to study hyperbolic geometry. The charm of straight lines that do not appear to be straight; a small tweak of a metric making trivial patterns beautifully complex. The sheer visual beauty of hyperbolic tessellations may be enough a reason, but beyond this, one finds a particularly lavish and spectacular example of non-euclidean geometry.

The hyperbolic plane is the long sought after anti-sphere: a complete surface of constant negative curvature. The famous Uniformization Theorem states that the universal cover of a Riemann surface is either the euclidean plane, the sphere, or the hyperbolic plane. In the light of this result hyperbolic geometry arises as an infinite source of surfaces of genus two or higher. Hyperbolic geometry does exist outside the mathematical universe too, from the intricacy of crinkled lettuce leaves and jellyfish tentacles, to the vast landscape of special relativity and, possibly, the geometry of *our* universe.[1]

The richness of hyperbolic geometry seeps into several areas of mathematics, with the notion of polygonal tessellations being at the intersection of these areas, and at the heart of this thesis. A tessellation is, naively, a covering of a space with a repeating pattern of shapes, with no overlaps or gaps. Tessellations by regular polygons of the euclidean plane are limited to triangles, squares and hexagons, and of the sphere to the five regular polyhedra. If we want to explore regular tessellations further, we have to look towards the hyperbolic plane.

The first part of this thesis is a brief introduction to hyperbolic geometry. Hyperbolic isometries will be studied from two perspectives: as elements of matrix groups, and as reflections in hyperbolic lines. Towards the end of the first part we seek to harmonise these two perspectives into a complete classification of hyperbolic isometries. Our ultimate goal is to, given a hyperbolic polygon, construct a tessellation of the hyperbolic plane. We will discover that Fuchsian groups, that is, discrete groups of orientation-preserving hyperbolic isometries, play a key role to our endeavour. A fundamental domain is a subset of the hyperbolic plane that tessellates the plane under the action of a Fuchsian group. Part two is dedicated to this special kind of isometry groups and their fundamental domains. The Poincaré Polygon Theorem, which will be our main concern in the third part, provides conditions for a polygon to generate a tessellation. We will not treat this theorem in its most general form, but give a proof in the case of compact hyperbolic polygons.

[1] A random forum post. Yes, I've done my research.

2 The Playground

Let us delve daringly into the magical world of hyperbolic geometry! We will spend most of our time with the hyperbolic plane; this is where the magic happens. To bring this mystical creature down to earth, or rather, to the euclidean plane where things are more familiar, we will present two models in which the hyperbolic plane sits as a subspace of the complex plane $\mathbb{C}$. Many of the characteristics of the hyperbolic plane are uncovered by studying its isometries. We will arrive at the isometry group through two different routes, picking up anything that piques our curiosity along the road (or in the gutter[2]).

2.1 The Poincaré Models

`Excuse Me, Where Am I?`

There are several models of the hyperbolic plane; we will work with two closely related models – the Poincaré half-plane and the Poincaré disc – as they give us exactly what we want (and more, if we're feeling greedy). The underlying space for both models is the complex plane $\mathbb{C}$. The half-plane model is the upper half-plane

$$\mathbb{H} = \{z \in \mathbb{C} \mid \mathrm{Im}(z) > 0\}$$

equipped with the element of arc-length

$$ds = \frac{|dz|}{\mathrm{Im}(z)}. \tag{2.1}$$

We find the hyperbolic length of paths by integrating the hyperbolic element of arc-length (2.1) as follows.

Definition 2.1. Let $\gamma : [0, 1] \to \mathbb{H}$ be a piecewise differentiable path given by $\gamma(t) = z(t) = x(t) + iy(t)$. The **hyperbolic length** of γ is given by

$$\mathrm{length}_{\mathbb{H}}(\gamma) = \int_\gamma ds = \int_\gamma \frac{|dz|}{\mathrm{Im}(z)} = \int_0^1 \frac{\left|\frac{dz}{dt}\right| dt}{y(t)}.$$

From this we derive a metric on $\mathbb{H}$.

[2]Maths is hard sometimes.

Definition 2.2. The **hyperbolic distance** between two points z and w in $\mathbb{H}$, denoted by $d_{\mathbb{H}}(z, w)$, is the infimum of the lengths of all piecewise differentiable paths connecting the two points.

Verifying that Definition 2.2 defines a metric is fairly straightforward; thus $(\mathbb{H}, d_{\mathbb{H}})$ is a metric space.[3]

We now introduce a type of complex functions that will play a key role throughout this work.

Definition 2.3. A **Möbius transformation** is a transformation of the complex plane of the form

$$z \mapsto \frac{az+b}{cz+d} \quad \text{with } a, b, c, d \in \mathbb{C} \text{ and } ad - bc \neq 0. \tag{2.2}$$

The domain of Möbius transformations extends to the **extended complex plane** $\mathbb{C} \cup \{\infty\}$. For reasons that will be clear later, it is helpful to treat $\mathbb{H}$ as a subspace of $\mathbb{C} \cup \{\infty\}$. Moreover, introducing the point at infinity allows us to generalize the notion of a circle in $\mathbb{C}$ as either a circle or the union of a line with $\{\infty\}$, thus treating lines as 'circles with infinite radius'. A quick check confirms that Möbius transformations are bijections of $\mathbb{C} \cup \{\infty\}$ and that the set of Möbius transformations forms a group under function composition.

Recall that a map is **conformal** if it preserves angles. An important feature of Möbius transformations is that they are conformal, and furthermore, that they preserve circles in $\mathbb{C} \cup \{\infty\}$.[4]

Proposition 2.4. *Any Möbius transformation*

(i) preserves angles, and

(ii) takes circles in $\mathbb{C} \cup \{\infty\}$ to circles in $\mathbb{C} \cup \{\infty\}$.

The disc model is constructed by mapping $\mathbb{H}$ onto the unit disc

$$\mathbb{D} = \{z \in \mathbb{C} \mid |z| < 1\}$$

using the Möbius transformation

$$J(z) = \frac{iz+1}{z+i}. \tag{2.3}$$

J maps the extended real line $\mathbb{R} \cup \{\infty\}$ onto the unit circle, and takes points above the real axis to interior points of the unit circle. In particular, i is

[3] [1], Theorem 3.16.
[4] [1], Theorem 2.5 and Theorem 2.23.

mapped to 0, ∞ is mapped to i, and ± 1 are fixed. Each model has its advantages, we will thus work with the two interchangeably, swapping back and forth via the map J.

We define the distance between points w_1 and w_2 in $\mathbb{D}$ to be the distance between their preimages $J^{-1}(w_1)$ and $J^{-1}(w_2)$ in $\mathbb{H}$. It follows that the isometries of $\mathbb{D}$ are conjugates of isometries of $\mathbb{H}$, that is, any $\mathbb{D}$-isometry is of the form JhJ^{-1}, where h is an $\mathbb{H}$-isometry.

The preimage of a point $w \in \mathbb{D}$ is given by

$$J^{-1}(w) = \frac{-iw+1}{w-i}.$$

Plugging this into (2.1) we find the element of arc-length on $\mathbb{D}$ to be

$$\begin{aligned} ds &= \left| d\frac{-iw+1}{w-i} \right| \Bigg/ \operatorname{Im}\left(\frac{-iw+1}{w-i}\right) \\ &= \left| \frac{-2dw}{(w-i)^2} \right| \Bigg/ \operatorname{Im}\left(\frac{(1-iw)(\bar{w}+i)}{|w-i|^2}\right) \\ &= \frac{|2dw|}{1-|w|^2}. \end{aligned}$$

We use the usual notion of angle that $\mathbb{H}$ inherits from $\mathbb{C}$. As J is conformal, this definition is carried over to $\mathbb{D}$ without modification.

We will later see that the geodesics in $\mathbb{H}$ (i.e. the shortest paths between points) are euclidean lines and semicircles orthogonal to the real axis (Theorem 2.9). These subsets of $\mathbb{H}$, which we will call hyperbolic lines, are preserved by a certain subgroup of Möbius transformations (Theorem 2.7). The notion of a circle in the extended complex plane gracefully unifies the two seemingly different kinds of hyperbolic lines.

Definition 2.5. A **hyperbolic line** is the intersection of $\mathbb{H}$ with a circle in $\mathbb{C} \cup \{\infty\}$ orthogonal to the real axis.

As the map J takes the real axis onto the unit circle, Proposition 2.4 reveals the lines in $\mathbb{D}$ – the images of lines in $\mathbb{H}$ under J – as circles orthogonal to the unit circle.

As in the euclidean plane, hyperbolic lines extend indefinitely. This is because the length of paths tends to infinity as points approach the boundary. Accordingly, we will refer to the boundary of $\mathbb{H}$ – the extended real line $\mathbb{R} \cup \{\infty\}$ – as the **circle at infinity**. In $\mathbb{D}$ this is of course the unit circle.

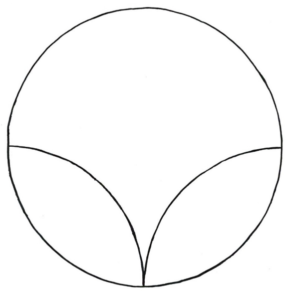

Figure 2.1. Asymptotic lines.

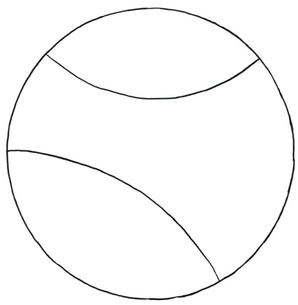

Figure 2.2. Ultraparallel lines.

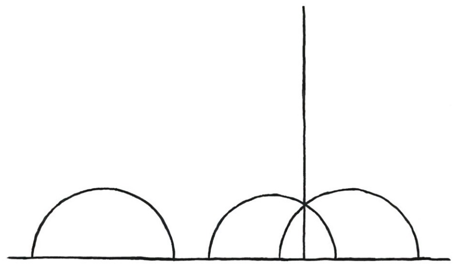

Figure 2.3. Several parallel lines.

Another way in which hyperbolic lines behave as euclidean is that for any two points in the hyperbolic plane there is a unique hyperbolic line connecting the two points. Through a euclidean lens, it is easy to see that there is a unique circle (in the general sense) orthogonal to the real axis passing through the two points.

We say that two hyperbolic lines are **parallel** if they are disjoint. The notion of parallel lines showcases the particularities of hyperbolic geometry. Given any line in the hyperbolic plane and any point not on the line, there are infinitely many lines through the given point and parallel to the given line. This is an example of how the hyperbolic plane diverts from the euclidean in the opposite way to the sphere, in which there are no parallel lines.

In addition to this excess of parallel lines, some lines enjoy a special status: they are lines that meet on the circle at infinity. In this case we say that the lines are **asymptotic**. Asymptotic lines are strictly speaking disjoint since the common point does not lie in the hyperbolic plane. To distinguish this special kind of parallel lines from the properly disjoint kind, we refer to the latter as **ultraparallel**.

In the remaining sections we will work our way towards a suitable presentation of the isometry group of $\mathbb{H}$. Let us denote it by $\mathrm{Iso}(\mathbb{H})$. It turns out that Möbius transformations serve as an excellent starting point.

2.2 The Group $\mathrm{PSL}(2, \mathbb{R})$

The Group with Capital G

Building on the preceding discussion, we impose conditions on Möbius transformations to map $\mathbb{H}$ onto itself. First we note that a map preserving $\mathbb{H}$ must preserve its boundary $\mathbb{R} \cup \{\infty\}$. For a Möbius transformation (2.2) this is the case when the coefficients are real. Let $T(z) = \frac{az+b}{cz+d}$ with $a, b, c, d \in \mathbb{R}$. Since

$$T(z) = \frac{(az+b)(c\bar{z}+d)}{|cz+d|^2} = \frac{acz\bar{z} + adz + bc\bar{z} + bd}{|cz+d|^2},$$

we have that

$$\mathrm{Im}\,(T(z)) = \frac{1}{|cz+d|^2}\mathrm{Im}(z)(ad-bc). \tag{2.4}$$

It follows that $\mathrm{Im}(z) > 0$ implies $\mathrm{Im}(T(z)) > 0$ if and only if $ad - bc > 0$. We may multiply numerator and denominator by a constant to obtain $ad - bc = 1$, as this represents the same transformation. Hence the Möbius transformations preserving $\mathbb{H}$ are of the form

$$z \mapsto \frac{az+b}{cz+d} \quad \text{with } a, b, c, d \in \mathbb{R} \text{ and } ad - bc = 1. \tag{2.5}$$

The set of functions (2.5) form a group under function composition. Taking a closer look, one realises that the group structure is very much similar to that of the special linear group $\mathrm{SL}(2, \mathbb{R})$. More precisely, we may define a surjective homomorphism ψ from $\mathrm{SL}(2, \mathbb{R})$ onto the group of transformations (2.5) by

$$\psi\left(\begin{bmatrix} a & b \\ c & d \end{bmatrix}\right) = \left(z \mapsto \frac{az+b}{cz+d}\right). \tag{2.6}$$

Since the kernel of this map is $\pm\mathrm{Id}$, where Id is the identity in $\mathrm{SL}(2, \mathbb{R})$, the group of Möbius transformations preserving $\mathbb{H}$ is isomorphic to $\mathrm{SL}(2, \mathbb{R})/\{\pm\mathrm{Id}\}$. We call this group the **projective special linear group**, denoted by $\mathrm{PSL}(2, \mathbb{R})$.

Theorem 2.6. $\mathrm{PSL}(2, \mathbb{R})$ *acts on* $\mathbb{H}$ *by isometries.*

Proof. We show that $\mathrm{PSL}(2, \mathbb{R})$ preserves the length of paths, it then follows from the definition that the hyperbolic distance is preserved. Let $\gamma : [0, 1] \to \mathbb{H}$ be a piecewise differentiable path given by $z(t) = x(t) + iy(t)$, let T be an element of $\mathrm{PSL}(2, \mathbb{R})$, and $w = T(z) = \frac{az+b}{cz+d}$. Put $T(z(t)) = w(t) = u(t)+iv(t)$. We have that

$$\frac{dw}{dz} = \frac{1}{(cz + d)^2}.$$

By (2.4) $v = \frac{y}{|cz+d|^2}$, so $\left|\frac{dw}{dz}\right| = \frac{1}{|cz+d|^2} = \frac{v}{y}$. Hence

$$\mathrm{length}_{\mathbb{H}}(T(\gamma)) = \int_0^1 \frac{\left|\frac{dw}{dt}\right| dt}{v(t)} = \int_0^1 \frac{\left|\frac{dw}{dz}\frac{dz}{dt}\right| dt}{v(t)} = \int_0^1 \frac{\left|\frac{dz}{dt}\right| dt}{y(t)} = \mathrm{length}_{\mathbb{H}}(\gamma).$$

■

As elements of $\mathrm{PSL}(2, \mathbb{R})$ preserve angles by Proposition 2.4, $\mathrm{PSL}(2, \mathbb{R})$ is a group of **orientation-preserving** isometries of $\mathbb{H}$.

Theorem 2.7. *The set of lines in* $\mathbb{H}$ *is invariant under* $\mathrm{PSL}(2, \mathbb{R})$.

Proof. As $\mathrm{PSL}(2, \mathbb{R})$ maps $\mathbb{R} \cup \{\infty\}$ onto itself, this is an immediate consequence of Proposition 2.4. ■

The following lemma will greatly facilitate our work as it allows us to swiftly transfer a situation to a simpler one by change of coordinates.

Lemma 2.8. $\mathrm{PSL}(2, \mathbb{R})$ *acts transitively on the set of hyperbolic lines.*

Proof. We prove the equivalent statement that for any hyperbolic line L there exists an element of $\mathrm{PSL}(2, \mathbb{R})$ mapping L to the positive imaginary axis I. Let α, α' be the endpoints at infinity of L. We show that $T(z) = -(z - \alpha)^{-1} + \beta$ takes L to I for a suitable choice of β. First note that a matrix representation of T is

$$\begin{bmatrix} \beta & -\alpha\beta - 1 \\ 1 & -\alpha \end{bmatrix}$$

which is clearly in $\mathrm{PSL}(2, \mathbb{R})$. By Theorem 2.7, it suffices to check that the endpoints of L are mapped to the endpoints 0 and ∞ of I. If L is a euclidean line with $\alpha \in \mathbb{R}$ and $\alpha' = \infty$ we set $\beta = 0$. If L is a euclidean semicircle with $\alpha, \alpha' \in \mathbb{R}$, $\alpha < \alpha'$ we set $\beta = \frac{1}{\alpha' - \alpha}$. ■

We now turn to proving that hyperbolic lines behave as one would expect lines to behave, that is, that they are the shortest paths between points in the plane.

Theorem 2.9. *The distance between two distinct points in $\mathbb{H}$ is the length of the unique hyperbolic line segment joining the two points.*

Proof. First, we consider points ia and ib on the positive imaginary axis with $a < b$. Let $\gamma : [0,1] \to \mathbb{H}$ be any piecewise differentiable path connecting ia and ib, defined by $\gamma(t) = (x(t), y(t))$. We have

$$\begin{aligned} \text{length}_{\mathbb{H}}(\gamma) &= \int_0^1 \frac{\sqrt{\left(\frac{dx}{dt}\right)^2 + \left(\frac{dy}{dt}\right)^2}dt}{y(t)} \\ &\geq \int_0^1 \frac{\left|\frac{dy}{dt}\right| dt}{y(t)} \geq \int_0^1 \frac{\frac{dy}{dt}dt}{y(t)} \\ &= \ln \frac{b}{a}. \end{aligned} \tag{2.7}$$

$\ln \frac{b}{a}$ is precisely the length of the segment of the imaginary axis joining ia and ib, and so the result holds in this case.

For arbitrary $z_1, z_2 \in \mathbb{H}$, let L be the unique hyperbolic line joining the two points. By Lemma 2.8 there is a transformation mapping L to the imaginary axis. By the invariance of the hyperbolic metric under $\mathrm{PSL}(2,\mathbb{R})$, the problem is reduced to the first case, and we conclude that the segment of L joining z_1 and z_2 is the shortest path between the two points. ∎

Letting a approach zero and letting b approach ∞ in the proof of Theorem 2.9 shows that hyperbolic lines extend indefinitely. We say that the hyperbolic lines are **geodesics**; we call the unique line segment joining two distinct points z and w the **geodesic segment** joining z and w, denoted by $[z,w]$.

Corollary 2.10. *Let z, w be distinct points in $\mathbb{H}$. Then*

$$d_{\mathbb{H}}(z,w) = d_{\mathbb{H}}(z,\xi) + d_{\mathbb{H}}(\xi,w)$$

if and only if $\xi \in [z,w]$.

We now derive a formula for the hyperbolic distance.

Proposition 2.11.

$$d_{\mathbb{H}}(z,w) = \ln \frac{|z-\bar{w}| + |z-w|}{|z-\bar{w}| - |z-w|} \tag{2.8}$$

for any $z, w \in \mathbb{H}$.

Proof. The left-hand side of (2.8) is invariant under $\mathrm{PSL}(2, \mathbb{R})$ by Theorem 2.6. We prove that the right-hand side is invariant. Lemma 2.8 then reduces the problem to showing that (2.8) holds for points on the imaginary axis, and a short calculation confirms that this is the case. Let $T \in \mathrm{PSL}(2, \mathbb{R})$. Using $ad - bc = 1$ we get

$$\begin{aligned} T(z) - T(w) &= \frac{az+b}{cz+d} - \frac{aw+b}{cw+d} \\ &= \frac{(z-w)(ad-bc)}{(cz+d)(cw+d)} \\ &= \frac{(z-w)}{(cz+d)(cw+d)}. \end{aligned}$$

Substituting into the right-hand side of (2.8) we see after some not so elegant algebraic manipulation that

$$\ln \frac{|T(z) - \overline{T(w)}| + |T(z) - T(w)|}{|T(z) - \overline{T(w)}| - |T(z) - T(w)|} = \ln \frac{|z - \bar{w}| + |z - w|}{|z - \bar{w}| - |z - w|}.$$

■

We will now look at some specific isometries that will come in handy later.

Example 2.12.

(i) Euclidean rotation

$$r_\theta(w) = e^{i\theta} w \quad \text{with } \theta \in \mathbb{R}$$

around 0 in $\mathbb{D}$ is a $\mathbb{D}$-isometry. The absolute value $|w|$ of points $w \in \mathbb{D}$ is preserved by rotation around 0. Since the element of arc-length on $\mathbb{D}$ depends only on $|w|$, distance is preserved.

(ii) Euclidean translation

$$t_\alpha(z) = z + \alpha \quad \text{with } \alpha \in \mathbb{R}$$

is given by the matrix $\begin{bmatrix} 1 & \alpha \\ 0 & 1 \end{bmatrix}$ in $\mathrm{PSL}(2, \mathbb{R})$, hence it is an $\mathbb{H}$-isometry.

(iii) Euclidean dilatation

$$d_\rho(z) = \rho z \text{ with positive } \rho \in \mathbb{R}$$

is given by the matrix $\begin{bmatrix} \sqrt{\rho} & 0 \\ 0 & \frac{1}{\sqrt{\rho}} \end{bmatrix}$ in $\mathrm{PSL}(2, \mathbb{R})$.

Although (i) can be seen as a hyperbolic rotation, (ii) is not a hyperbolic translation, but an isometry unique to hyperbolic geometry called limit rotation, and (iii) is in fact a hyperbolic translation. This terminology will make sense once we understand isometries as reflections in hyperbolic lines. A definition will be given in Section 2.4 (Definition 2.28). With a hint of foreshadowing, the isometries in Example 2.12 give us a very good idea of how to visualise the action of $\mathrm{PSL}(2, \mathbb{R})$ on the hyperbolic plane.

Each isometry permutes a certain set of lines, and leaves invariant curves orthogonal to the permuted lines. r_θ permutes the diameters of $\mathbb{D}$, and leaves invariant circles centred at 0. In $\mathbb{H}$, the permuted lines are euclidean circles through i, and the invariant curves are hyperbolic circles with centre i. t_α permutes the lines $x =$ constant and leaves invariant the curves $y =$ constant. These curves are called **horocycles**. In $\mathbb{D}$ the permuted lines are hyperbolic lines through i, and the invariant horocycles are euclidean circles tangential to the circle at infinity. d_ρ permutes euclidean semicircles with centre 0, and leaves invariant euclidean lines $y = x \times$ constant. In $\mathbb{D}$ the permuted lines are euclidean circles orthogonal to the imaginary axis, and the invariant curves are euclidean circles passing through $\pm i$. Figure 2.4, Figure 2.5, and Figure 2.6 picture some of the permuted lines and invariant curves for r_θ, t_α, and d_ρ respectively, in the disc model and in the half-plane model.

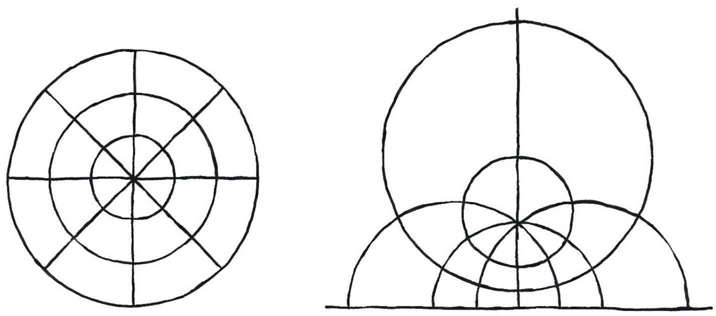

Figure 2.4.

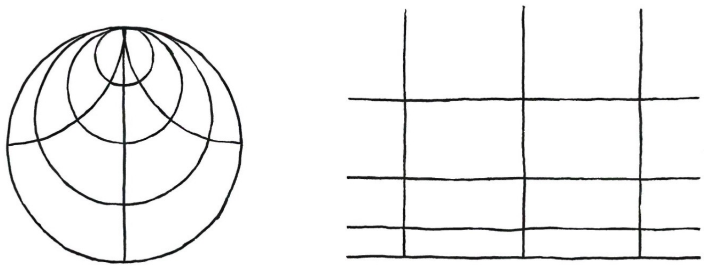

Figure 2.5.

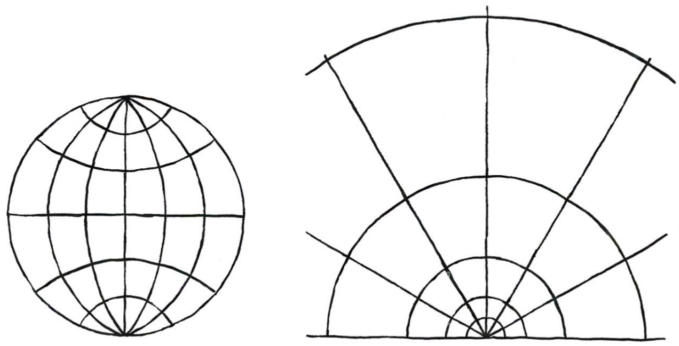

Figure 2.6.

2.3 Reflections

Upside Down and Inside Out (We Are All Mad Here)

In this section we will explore hyperbolic reflections. The definition is analogous to the usual notion of reflection.

Definition 2.13. Let L be a hyperbolic line and let P be a point not on L. Let M be the unique line through P intersecting L orthogonally. Let Q be the intersection point. The **mirror image** of P with respect to L is the point P' such that

(i) P' lies on M, and

(ii) $d_{\mathbb{H}}(P, Q) = d_{\mathbb{H}}(P', Q)$.

We define **hyperbolic reflection** in L to be the map fixing points on L and sending points not on L to their mirror images with respect to L.

Remark 2.14. To see that M is uniquely determined by P, we take L to be the imaginary axis. Then M is a circle with centre 0, and P is a point on M, so the radius of M is the euclidean distance from 0 to P.

When the line of reflection is a euclidean line, hyperbolic reflection coincides with euclidean reflection.

Example 2.15. Euclidean reflection in the imaginary axis $\bar{r}_I(z) = -\bar{z}$ is a $\mathbb{H}$-reflection. Figure 2.7 shows a point P in $\mathbb{H}$ and its mirror image under $\bar{r}_I$.

Proposition 2.16. *Reflection in the line $x = \epsilon$ is given by $t_\epsilon \bar{r}_I t_\epsilon^{-1}$.*

To establish what it means to 'reflect' in a circle, we need the following definition from euclidean geometry.

Definition 2.17. Let C be a euclidean circle with centre O and radius ρ, and let P be a point different from O. The **inverse** of P with respect to C is the point P' such that

(i) P' lies on the line through O and P, and

(ii) $|OP'||OP| = \rho^2$.

We define **inversion** in C to be the map sending points to their inverses with respect to C.

Remark 2.18. As $P \neq O$, $|OP|$ is nonzero and so P' is the point on the line through O and P with distance $\rho^2/|OP|$ from O. Note that if P lies on the circle C, then $P' = P$. Furthermore, we can extend the definition to $\mathbb{C} \cup \{\infty\}$ by letting inversion in C interchange O and ∞.

Example 2.19. Inversion in the unit circle is given by $I(z) = 1/\bar{z}$. To see this, rewrite $I(z)$ as $z/|z|^2$. As each point z is scaled by a factor $|z|^2$, condition (i) of Definition 2.17 is satisfied. For the unit circle condition (ii) reads $\frac{|z|}{|z|^2}|z| = 1$, which obviously holds for all z different from 0.

Lemma 2.20. *Inversion in a circle centred on the real axis is conjugate to inversion in the unit circle.*

Proof. Consider the circle $C_{\epsilon,\rho}$ with centre $\epsilon \in \mathbb{R}$ and radius ρ. Let P be some point in $\mathbb{H}$ with P' its image under inversion in $C_{\epsilon,\rho}$. We apply the isometry $t_\epsilon^{-1} : z \mapsto z - \epsilon$ to move the centre of the circle to the origin. Let $P_1 = t_\epsilon^{-1}(P), P_1' = t_\epsilon^{-1}(P')$. It is clear that P_1, P_1' satisfy conditions (i) and (ii) of Definition 2.17 for inversion in the circle with radius ρ centred at 0.

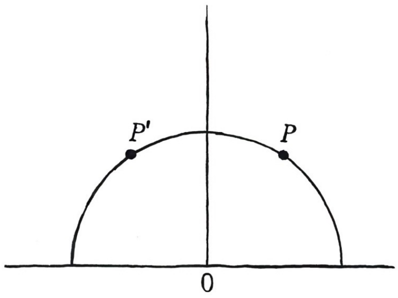

Figure 2.7.

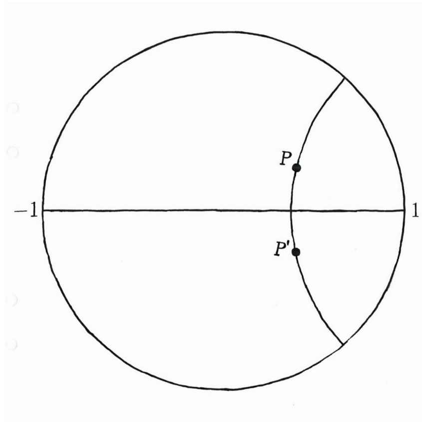

Figure 2.8.

Now we apply the isometry $d_\rho^{-1} = z \mapsto z/\rho$. This moves points along lines radiating from the origin, so condition (i) still holds. As P_1, P_1', and ρ are scaled by the same factor, condition (ii) holds. Hence inversion in the circle $C_{\epsilon,\rho}$ is the map $t_\epsilon d_\rho I d_\rho^{-1} t_\epsilon^{-1}$. ■

Proposition 2.21. *Inversion in a euclidean circle centred on the real axis is a $\mathbb{H}$-reflection.*

Proof. By Lemma 2.20, it suffices to prove that inversion in the unit circle is a hyperbolic reflection. Passing to $\mathbb{D}$, we see that JIJ^{-1} is the map $w \mapsto \bar{w}$, which is reflection in the real axis. This is clearly a $\mathbb{D}$-reflection. Figure 2.8 shows a point P in $\mathbb{D}$ and its mirror image under reflection in the real axis. ■

Proposition 2.21 and Proposition 2.16 cover all hyperbolic reflections. A calculation similar to the proof of Theorem 2.6 proves that inversion in the unit circle is an isometry of $\mathbb{H}$. It is clear that $\bar{r}_I$ is an isometry. As any hyperbolic reflection is a conjugate of either of these two reflections, we see that hyperbolic reflections are isometries.

It is a standard result that inversions are are anti-conformal and map circles to circles.[5] However, with the tools at hand, proving the following is an easy task.

[5] [6], p.55.

Proposition 2.22. *Reflection in a hyperbolic line*

(i) maps hyperbolic lines to hyperbolic lines, and

(ii) preserves angles, but reverses their signs.

Proof. Let L be the line of reflection. As elements of $\mathrm{PSL}(2, \mathbb{R})$ take hyperbolic lines to hyperbolic lines and preserve angles, we can apply Lemma 2.8 to map L to the imaginary axis, for which the result is trivial. ■

Proposition 2.22 tells us that reflections are **orientation-reversing**. Consequently, a product of two reflections is orientation-preserving.

We have seen that hyperbolic lines are length-minimising paths; the following lemma provides another characterization of lines.

Lemma 2.23. *The set of points equidistant from two points $P, P' \in \mathbb{H}$ is a line L, and reflection in L exchanges P and P'.*

Proof. By rotation and euclidean translation we can take the points to be mirror images in the imaginary axis. Reflection $\bar{r}_I$ in the imaginary axis exchanges P and P' and fixes the imaginary axis. Hence any point on the imaginary axis is equidistant from P, P', so that the imaginary axis is contained in the set of equidistant points. Suppose there is a point R equidistant from P, P' not on the imaginary axis. Then the mirror image R' of R is equidistant from P, P'. Without loss of generality, we may assume a situation like that in Figure 2.9. Let Q be the point of intersection of the lines $P'R$ and PR' (which obviously lies on the imaginary axis). Then we have

$$\begin{aligned} d_{\mathbb{H}}(P', R') &= d_{\mathbb{H}}(P, R) \quad \text{by reflection} \\ &= d_{\mathbb{H}}(P', R) \quad \text{by assumption} \\ &= d_{\mathbb{H}}(P', Q) + d_{\mathbb{H}}(Q, R) \quad \text{by Corollary 2.10} \\ &= d_{\mathbb{H}}(P', Q) + d_{\mathbb{H}}(Q, R') \quad \text{by reflection.} \end{aligned}$$

Since Q is not on the line segment joining P' and R', this contradicts Corollary 2.10. We conclude that the imaginary axis is the equidistant set of P, P'. ■

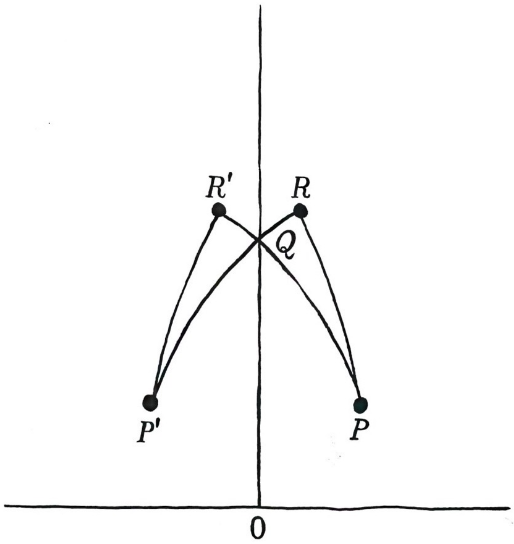

Figure 2.9.

The following theorem is of importance. It shows that Iso($\mathbb{H}$) is generated by reflections in hyperbolic lines.

Theorem 2.24. *Any hyperbolic isometry can be expressed as a product of one, two, or three reflections.*

Proof. We begin with noting that any $\mathbb{H}$-isometry is uniquely determined by the images of three points not in a line. This is because each point of $\mathbb{H}$ is determined by its distance to three points A, B, C not in a line: if there were two points P, P' with the same distance to A, B, C then these points lie in the equidistant set of P, P', contradicting Lemma 2.23. To prove the theorem let $f \in \text{Iso}(\mathbb{H})$ and pick A, B, C as above.

Case 1: f fixes two of the points, say A and B. Then $A = f(A)$ and $B = f(B)$ are contained in the equidistant set of C and $f(C)$. Hence reflecting in the line joining A and B sends C to $f(C)$, and we are done.

Case 2: f fixes one of the points, say A. Then A lies on the line of points equidistant from B and $f(B)$. Hence reflection $\bar{g}$ in this line sends B to $f(B)$ and A to $f(A)$. If this leaves C at $f(C)$ we are done. If not, reflection $\bar{h}$ in the line through $f(A)$ and $f(B)$ sends $\bar{g}(C)$ to $f(C)$.

Case 3: None of the points are fixed by f. Then we construct f as above with one, two, or three reflections $\bar{g}, \bar{h}, \bar{i}$ in lines equidistant from A and $f(A)$, $\bar{g}(B)$ and $f(B)$, and $\bar{h}\bar{g}(C)$ and $f(C)$ respectively. ■

We are now ready to complete the task set out earlier – to identify the isometry group of $\mathbb{H}$. Let $\mathrm{PS}^*\mathrm{L}(2,\mathbb{R}) = \mathrm{S}^*\mathrm{L}(2,\mathbb{R})/\{\pm\mathrm{Id}\}$, where $\mathrm{S}^*\mathrm{L}(2,\mathbb{R})$ is the group of real matrices with determinant ± 1. $\mathrm{PSL}(2,\mathbb{R})$ is a subgroup of $\mathrm{PS}^*\mathrm{L}(2,\mathbb{R})$ of index 2.

Theorem 2.25. $\mathrm{Iso}(\mathbb{H})$ *is generated by* $\mathrm{PSL}(2,\mathbb{R})$ *together with* $\bar{r}_I(z) = -\bar{z}$ *and is isomorphic to* $\mathrm{PS}^*\mathrm{L}(2,\mathbb{R})$. $\mathrm{PSL}(2,\mathbb{R})$ *is the subgroup of orientation-preserving isometries.*

Proof. Let
$$f(z) = \frac{az+b}{cz+d} \quad \text{with } ad - bc = 1, \tag{2.9}$$
and
$$\bar{f}(z) = \frac{a\bar{z}+b}{c\bar{z}+d} \quad \text{with } ad - bc = -1. \tag{2.10}$$
Note that the maps $\bar{f} \in \mathrm{Iso}(\mathbb{H})$ are orientation-reversing isometries as they are compositions of $\bar{r}_I$ with orientation-preserving isometries (2.9). Composing two isometries (2.10) will again result in an isometry of the form (2.9).

We claim that the map $\psi : \mathrm{PS}^*\mathrm{L}(2,\mathbb{R}) \to \mathrm{Iso}(\mathbb{H})$ defined by
$$\psi\left(\begin{bmatrix} a & b \\ c & d \end{bmatrix}\right) = \begin{cases} z \mapsto f(z) \text{ if } ad - bc = 1 \\ z \mapsto \bar{f}(z) \text{ if } ad - bc = -1 \end{cases}$$
is an isomorphism.

A routine calculation shows that ψ is a homomorphism. We see that $\bar{r}_I$ is in the image of ψ. As $\mathrm{PSL}(2,\mathbb{R})$ acts transitively on hyperbolic lines, any reflection is a conjugate of $\bar{r}_I$ and so is of the form (2.10). By Theorem 2.24, this covers all isometries. Hence ψ is surjective. It is clear that the kernel is trivial, so that ψ is in fact an isomorphism. ■

2.4 Classification of Hyperbolic Isometries

Mad About Fixed Points

We have seen how to describe isometries as complex functions represented by matrices, and as reflections in hyperbolic lines. We now set out to give a complete classification of hyperbolic isometries and to merge the two different perspectives.

We borrow the following classification of elements of $\mathrm{PSL}(2,\mathbb{R})$ from linear algebra, depending on the absolute value of the trace $tr(T) = a + d$.

Definition 2.26. Let $T \in \mathrm{PSL}(2, \mathbb{R})$.

(i) If $|tr(T)| < 2$ then T is called **elliptic**.

(ii) If $|tr(T)| = 2$ then T is called **parabolic**.

(iii) If $|tr(T)| > 2$ then T is called **hyperbolic**.

We solve the equation

$$z = \frac{az + b}{cz + d} \tag{2.11}$$

to find the fixed points of T. If $c = 0$ then there is either one solution $z = \infty$ when $a = d$, or two solutions $z = \infty, b/(d - a)$ when $a \neq d$. If $c \neq 0$ then using $ad - bc = 1$ we have the solutions

$$z = \frac{a - d \pm \sqrt{(a + d)^2 + 4(bc - ad)}}{2c} = \frac{a - d \pm \sqrt{(a + d)^2 - 4}}{2c}.$$

This concludes that the trace squared determines whether T fixes one or two points on the circle at infinity, or two complex conjugate points, corresponding to one in $\mathbb{H}$. In particular, this tells us that a nontrivial orientation-preserving isometry can have no more than two fixed points. We summarise this in the following proposition.

Proposition 2.27.

(i) An elliptic transformation has one fixed point in $\mathbb{H}$.

(ii) A parabolic transformation has one fixed point in $\mathbb{R} \cup \{\infty\}$.

(iii) A hyperbolic transformation has two fixed points in $\mathbb{R} \cup \{\infty\}$.

As any orientation-preserving isometry $f(z) \in \mathrm{Iso}(\mathbb{H})$ is the product of an even number of reflections, Theorem 2.24 infers that $f(z)$ is the product of *two* reflections. The lines of reflection are either intersecting, asymptotic, or ultraparallel. This gives rise to the following definition.

Definition 2.28.

(i) A hyperbolic **rotation** is a composition of reflections in two intersecting lines.

(ii) A **limit rotation** is a composition of reflections in two asymptotic lines

(iii) A a hyperbolic **translation** is a composition of reflections in two ultraparallel lines.

We wish to identify these different types of isometries with the classification of elements of $\mathrm{PSL}(2,\mathbb{R})$ given in Definition 2.26. As you may have sensed by now, the link between $\mathbb{H}$-isometries expressed as reflections in hyperbolic lines and their matrix representations is provided by their fixed points.

Before moving on, we need to remark that we are using the term 'fixed point' a little irresponsibly. In order to distinguish the different types of isometries, we allow fixed points on the circle at infinity. However, points on the circle at infinity are not actual fixed points of their action on $\mathbb{H}$. Any isometry with fixed points on the circle at infinity is in reality fixed point free. As elliptic elements are the only isometries with a fixed point in $\mathbb{H}$, groups of isometries containing elliptic elements are the only subgroups of $\mathrm{PSL}(2,\mathbb{R})$ not acting freely on $\mathbb{H}$.

Proposition 2.29. *Let L, M be two distinct hyperbolic lines. Let $\bar{r}_L$ denote reflection in L and $\bar{r}_M$ reflection in M.*

- *(i) If L, M intersect in a point $P \in \mathbb{H}$ then P is the unique fixed point of the rotation $\bar{r}_M \bar{r}_L$.*
- *(ii) If L, M are asymptotic then the unique fixed point of the limit rotation $\bar{r}_M \bar{r}_L$ is the common endpoint of L and M.*

Proof. The proofs of the two cases are identical, so we take L, M to be either intersecting or asymptotic at P. It is clear that P is fixed. Suppose $\bar{r}_M \bar{r}_L$ fixes a point Q different from P. Any such point cannot be fixed by both reflections and hence we must have that both $\bar{r}_L$ and $\bar{r}_M$ interchanges Q and some point Q'. This implies that M and L belong to the equidistant set of Q, Q' and so by Lemma 2.23 we have $M = L$, a contradiction. ■

To determine the fixed points of a translation we need two lemmas.

Lemma 2.30. *Suppose L, M are two distinct lines. Then reflection in L preserves M if and only if M intersects L orthogonally. Furthermore, if M is mapped onto itself by reflection in L, then the endpoints of M are interchanged.*

Proof. As $\mathrm{PSL}(2,\mathbb{R})$ preserves angles, we may apply Lemma 2.8 to map L to the positive imaginary axis. Then M is preserved by $\bar{r}_I$ if and only if it is a semicircle centred at 0. The second claim of the lemma is evident. ■

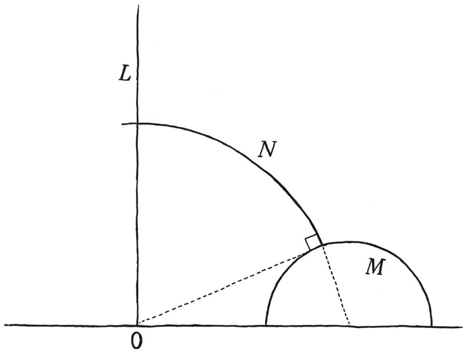

Figure 2.10.

Lemma 2.31. *Suppose L, M are ultraparallel lines. Then there is a unique line intersecting both L and M orthogonally.*

Proof. As in the proof of the previous lemma we take L to be the positive imaginary axis. Then M is a euclidean circle which does not intersect the imaginary axis (see Figure 2.10). Let ϵ denote the centre of M and ρ its radius. Suppose N is a common perpendicular to L and M. Any hyperbolic line perpendicular to L is a euclidean circle with centre 0. Let r denote the radius of N. For N to intersect M orthogonally the equation $\epsilon^2 = r^2 + \rho^2$ must be satisfied, which has a unique positive solution for r. ■

Proposition 2.32. *Let L, M be two ultraparallel lines, and let N be the unique common perpendicular of L and M. Then the endpoints of N are the unique fixed points of the translation $\bar{r}_M \bar{r}_L$.*

Proof. It follows from the above lemmas that $\bar{r}_M \bar{r}_L$ leaves N invariant. As each reflection $\bar{r}_L, \bar{r}_M$ interchanges the endpoints of N, the endpoints are fixed by $\bar{r}_M \bar{r}_L$. Since an orientation-preserving isometry has at most two fixed points, we are done. ■

We call the common perpendicular the **axis** of the translation. We will now see how Example 2.12 ties into the discussion.

Proposition 2.33.

(i) Each rotation in $\mathbb{H}$ is conjugate to $\frac{\cos(\theta)z+\sin(\theta)}{-\sin(\theta)z+\cos(\theta)}$ for some $\theta \in \mathbb{R}$.

(ii) Each limit rotation in $\mathbb{H}$ is conjugate to $t_\alpha(z) = z + \alpha$ for some $\alpha \in \mathbb{R}$.

(iii) *Each translation in $\mathbb{H}$ is conjugate to $d_\rho(z) = \rho z$ for some positive $\rho \in \mathbb{R}$.*

Proof. The property of pairs of lines being intersecting, asymptotic, or ultraparallel is preserved by isometries. Hence by transitivity of $\mathrm{PSL}(2, \mathbb{R})$ we may apply isometries to obtain one of the standard isometries introduced in Example 2.12.

Let $f(z) = \bar{r}_M \bar{r}_L$ for two distinct lines L, M.

(i) In this case L and M are intersecting, and we can apply suitable transformations t_α and d_ρ to map the intersection point to i. Passing to $\mathbb{D}$, reflecting in lines intersecting at $J(i) = 0$ gives a rotation r_θ in $\mathbb{D}$, as these lines are diameters of $\mathbb{D}$. Conjugating by J gives the formula for rotation around i in $\mathbb{H}$:

$$f(z) = J^{-1} r_\theta J(z) = \frac{\cos(\theta)z + \sin(\theta)}{-\sin(\theta)z + \cos(\theta)}.$$

(ii) In this case L and M are asymptotic at, say, $a \in \mathbb{R}$. We may apply $z \mapsto \frac{1}{a-z}$ to map the common endpoint to ∞. Then L and M are euclidean lines and $f(z) = t_\alpha(z)$ for some $\alpha \in \mathbb{R}$.

(iii) In this case L and M are ultraparallel. We may apply an isometry that takes the axis of $f(z)$ to the imaginary axis. Then L and M are semicircles with centre 0, and $f(z) = d_\rho(z)$ for some $\rho > 0$.

■

At this point, we have everything we need to give a complete classification of orientation-preserving $\mathbb{H}$-isometries.

Theorem 2.34 (Classification)**.** *Let $T \in \mathrm{PSL}(2, \mathbb{R})$.*

A. *The following are equivalent.*

(i) *T is a rotation.*

(ii) *T is elliptic.*

(iii) *T fixes one point in $\mathbb{H}$.*

(iv) *T is conjugate to $\begin{bmatrix} cos\theta & \sin\theta \\ -\sin\theta & \cos\theta \end{bmatrix}$ for some $\theta \in \mathbb{R}$.*

B. *The following are equivalent.*

(i) T is a limit rotation.

(ii) T is parabolic.

(iii) T fixes one point in $\mathbb{R} \cup \{\infty\}$.

(iv) T is conjugate to $\begin{bmatrix} 1 & \alpha \\ 0 & 1 \end{bmatrix}$ for some $\alpha \in \mathbb{R}$.

C. *The following are equivalent.*

(i) T is a translation.

(ii) T is hyperbolic.

(iii) T fixes two points in $\mathbb{R} \cup \{\infty\}$.

(iv) T is conjugate to $\begin{bmatrix} \rho & 0 \\ 0 & 1/\rho \end{bmatrix}$ for some nonzero $\rho \in \mathbb{R}$.

We now introduce a new type of hyperbolic isometry.

Definition 2.35. A **glide reflection** is a translation (possibly trivial) composed with a reflection in the axis of the translation.

We will prove that all orientation-reversing isometries are glide reflections. First we need a little piece of information about the fixed points of an orientation-reversing isometry.

Lemma 2.36. *An orientation-reversing $\mathbb{H}$-isometry has two fixed points on the circle at infinity.*

Proof. Using the fact that $x = \bar{x}$ for $x \in \mathbb{R} \cup \{\infty\}$, we find the fixed points on the circle at infinity of an orientation-reversing isometry (2.10) by again solving (2.11), this time for real values. If $c = 0$ the solutions are $x = b/(d-a), \infty$. Otherwise, we have the solutions

$$\begin{aligned} x &= \frac{-(d-a) \pm \sqrt{(d-a)^2 + 4bc}}{2c} \\ &= \frac{a-d \pm \sqrt{d^2 - 4ad + 2ad + a^2 + 4bc}}{2c} \\ &= \frac{a-d \pm \sqrt{(a+d)^2 + 4(bc-ad)}}{2c} \\ &= \frac{a-d \pm \sqrt{(a+d)^2 + 4}}{2c}, \end{aligned}$$

using $ad - bc = -1$ in the last step. This gives two distinct real solutions since $(a + d)^2 + 4 > 0$. ■

Theorem 2.37. *Any orientation-reversing isometry is a glide reflection.*

Proof. Let $\bar{f} \in \mathrm{Iso}(\mathbb{H})$ be orientation-reversing. By Lemma 2.36 $\bar{f}$ has two fixed points on the circle at infinity. Let L be the unique geodesic connecting the two fixed points. Now the isometry $f = \bar{f}\bar{r}_L$ is orientation-preserving and fixes the endpoints of L, so f must be a translation with axis L (possibly trivial). It follows that $\bar{f} = f\bar{r}_L$ is a glide reflection with axis L. ■

3 The Tools

Now that we have located the playground and gotten somewhat confident with navigating around it, how do we play? There are naturally many ways to go about this. In order to not stray too far away from the path and get completely lost, we will keep our gaze fixed at tessellations, whilst keeping an open mind.

3.1 Basics

`Topology Bonanza or the Not so Basic Basics`

Before we dig through the art supplies, we need to brush up our topology skills. Here we will introduce concepts that will be frequently used, and state a few important results that we will refer to later on. We do assume the reader to be somewhat familiar with these concepts and keep the discussion brief. Everything presented in this section can be found in [4].

3.1.1 Topological Groups

Definition 3.1. A **topological group** G is a group and a topological space such that the group operation map $G \times G \to G$ defined by $(x, y) \mapsto xy$, and the inversion map $G \to G$ defined by $x \mapsto x^{-1}$ are continuous.

The following is a direct consequence of the continuity of these maps.

Proposition 3.2. *Let G be a topological group and let $a \in G$ be some fixed element of G. Then the map $G \to G$ defined by $x \mapsto ax$ is a homeomorphism.*

It follows that a topological group is **homogeneous**, that is, given any two points $x, y \in G$, there is a homeomorphism of G mapping x to y (namely $t \mapsto yx^{-1}t$). What this means is that local properties of G can be verified by considering only one element of G. This gives us a nice condition for discreteness of topological groups.

Proposition 3.3. *A topological group is discrete if and only if the identity element is isolated.*

3.1.2 Local Properties

Definition 3.4. Let X be a topological space. X is called **locally compact** if for every point x of X, there exists a compact set $K \subset X$ containing a neighbourhood of x.

Definition 3.5. A family $\mathcal{A}$ of subsets of a topological space X is said to be **locally finite** if each point of X has a neighbourhood V that intersects at most finitely many of the sets in $\mathcal{A}$.

For a locally compact space, we may rephrase Definition 3.5 as follows.

Proposition 3.6. *Let X be a locally compact topological space. Then a family $\mathcal{A}$ of subsets of X is locally finite if and only if for any compact subset $K \subset X$,*

$$K \cap A \neq \varnothing \text{ for at most finitely many } A \in \mathcal{A}. \tag{3.1}$$

Proof. Suppose (3.1) holds. Let $x \in X$. By local compactness of X, there is a compact set K and a neighbourhood V with $x \in V \subset K$. Then the set $\{A \in \mathcal{A} \mid V \cap A \neq \varnothing\}$ is a subset of the finite set $\{A \in \mathcal{A} \mid K \cap A \neq \varnothing\}$, so is finite. So $\mathcal{A}$ is locally finite.

Conversely, suppose $\mathcal{A}$ is locally finite and let K be a compact subset of X. For all $x \in K$, there is a neighbourhood V_x intersecting at most finitely many elements of $\mathcal{A}$. We have that $K \subset \bigcup_{x \in K} V_x$. Since K is compact, finitely many of the neighbourhoods V_x cover K. Thus K is contained in a finite union of sets, each of which intersects finitely many of the sets in $\mathcal{A}$. So (3.1) holds. ■

3.1.3 Covering Spaces

Here we present two results that will be needed for the proof of the Poincaré Polygon Theorem (Theorem 4.6). First, we need a few definitions.

Definition 3.7. Let E and X be topological spaces and $q : E \to X$ a continuous map. An open subset U of X is said to be **evenly covered** by q if the preimage of U is a disjoint union of connected open subsets of E, called the **sheets** of the covering, and if q restricts to a homeomorphism from each sheet to U.

Definition 3.8. A **covering map** is a continuous map $q : E \to X$ such that every point of X has an evenly covered neighbourhood. We call E a **covering space** of X, and X the **base** of the covering.

Lemma 3.9. *Suppose M and N are metric spaces, $f : M \to N$, and suppose there exists an $\epsilon > 0$ such that for every $x \in M$, f restricts to an isometry $B_\epsilon(x) \to B_\epsilon(f(x))$. Then f is a covering map.*

Proof. Let $y \in N$ and let $f^{-1}(y)$ be the preimage of y. We want to prove that $B_\epsilon(y)$ is evenly covered. More precisely, we prove that the preimage of $B_\epsilon(y)$

is the disjoint union of the balls $B_\epsilon(x)$ for each $x \in f^{-1}(y)$. By assumption $B_\epsilon(x) \subseteq f^{-1}(B_\epsilon(y))$ for all $x \in f^{-1}(y)$.

Suppose z is a point in the preimage of $B_\epsilon(y)$. Then $B_\epsilon(f(z))$ contains y. By assumption $B_\epsilon(f(z)) = f(B_\epsilon(z))$. So $f(B_\epsilon(z))$ contains y. So $B_\epsilon(z)$ contains a point in the preimage of y. That is, $x \in B_\epsilon(z)$ for some $x \in f^{-1}(y)$, which is equivalent to $z \in B_\epsilon(x)$. We have proved that $f^{-1}(B_\epsilon(y)) = \bigcup_{x \in f^{-1}(y)} B_\epsilon(x)$. It remains to prove that the balls $B_\epsilon(x)$ are disjoint. Suppose $B_\epsilon(x_i) \cap B_\epsilon(x_j) \neq \varnothing$ for some $x_i, x_j \in f^{-1}(y)$. Pick $z \in B_\epsilon(x_i) \cap B_\epsilon(x_j)$. Then $x_i, x_j \in B_\epsilon(z)$. Then $d_M(x_i, x_j) = d_N(f(x_i), f(x_j)) = d_N(y, y) = 0$, so $x_i = x_j$.

We conclude that the preimage of $B_\epsilon(y)$ is the disjoint union $\bigcup_{x \in f^{-1}(y)} B_\epsilon(x)$. Each sheet $B_\epsilon(x)$ is mapped homeomorphically onto $B_\epsilon(y)$, as any isometry is also a homeomorphism. ■

The following theorem is a standard consequence of the classification of covering spaces.

Theorem 3.10. *Suppose X is simply connected, locally simply connected, and locally path-connected, and suppose E is path-connected and nonempty. Then every covering map $q : E \to X$ is a homeomorphism.*

3.2 Fuchsian Groups

`For the Love of Abstract Algebra`

Landing softly once more on hyperbolic ground, we endow $\mathbb{H}$ with the metric topology. The next result shows that the hyperbolic plane and the euclidean plane are homeomorphic.

Proposition 3.11. *The hyperbolic metric and the euclidean metric induce the same topology on $\mathbb{H}$.*

Proof. We show that a hyperbolic disc in the upper half-plane is a euclidean disc and vice versa. The metric topology is generated by the set of all open discs, and the result follows. Euclidean circles in $\mathbb{D}$ with centre 0 are clearly hyperbolic circles. Mapping to $\mathbb{H}$, this is enough to identify all hyperbolic circles. However, we may also find the circles in $\mathbb{H}$ directly. Consider the hyperbolic circle with radius r and centre $w \in \mathbb{H}$. We use the formula for the hyperbolic distance given in Proposition 2.11. For any point z on the circle we have

$$\begin{aligned}
r = \ln \frac{|z-\bar{w}|+|z-w|}{|z-\bar{w}|-|z-w|} &\iff \\
e^r = \frac{|z-\bar{w}|+|z-w|}{|z-\bar{w}|-|z-w|} &\iff \\
e^r(|z-\bar{w}|-|z-w|) = |z-\bar{w}|+|z-w| &\iff \\
|z-\bar{w}|(e^r-1) = |z-w|(e^r+1) &\iff \\
\frac{|z-w|}{|z-\bar{w}|} = \frac{e^r-1}{e^r+1}&.
\end{aligned}$$

So the ratio of the distance between z and w to z and $\bar{w}$ is a constant between 0 and 1. Hence z lies on the Apollonian circle with focus points w and $\bar{w}$; furthermore, the circle lies entirely in $\mathbb{H}$. ∎

Hence to our convenience $\mathbb{H}$ inherits many useful properties of $\mathbb{C}$. For instance, we deduce that $\mathbb{H}$ is second countable, locally compact and simply connected.

$\mathrm{SL}(2,\mathbb{R})$ can be identified with the subset $\{(a,b,c,d) \in \mathbb{R}^4 \mid ad-bc=1\}$ of $\mathbb{R}^4$. In other words we identify the matrix

$$A = \begin{bmatrix} a & b \\ c & d \end{bmatrix} \in \mathrm{SL}(2,\mathbb{R})$$

with the point $(a,b,c,d) \in \mathbb{R}^4$. We define a norm on $\mathrm{SL}(2,\mathbb{R})$ by $||A|| = \sqrt{a^2+b^2+c^2+d^2}$. Then $\mathrm{SL}(2,\mathbb{R})$ is a topological space with respect to the metric $d(A,B) = ||A-B||$. The map (2.6) from $\mathrm{SL}(2,\mathbb{R})$ to $\mathrm{PSL}(2,\mathbb{R})$ defines an equivalence relation $A \sim -A$ on $\mathrm{SL}(2,\mathbb{R})$ and $\mathrm{PSL}(2,\mathbb{R})$ is topologized as the quotient space under this equivalence relation. As orientation–reversing isometries are given by matrices in $\mathrm{GL}(2,\mathbb{R})$ with determinant -1, $\mathrm{Iso}(\mathbb{H})$ can be topologized using the same metric. We define convergence in $\mathrm{Iso}(\mathbb{H})$ with respect to the metric as usual.

Now that we have entered the realm of topology, we can make sense of the title of this section.

Definition 3.12. A **Fuchsian group** is a discrete subgroup of $\mathrm{PSL}(2,\mathbb{R})$. In other words, Fuchsian groups are discrete groups consisting of orientation-preserving $\mathbb{H}$-isometries.

We give an example of a Fuchsian group that we will return to later, as well as a non-example.

Example 3.13. The subgroup of $\mathrm{PSL}(2,\mathbb{R})$ consisting of matrices with integer coefficients is clearly discrete. We call this group the **modular group**, denoted by $\mathrm{PSL}(2,\mathbb{Z})$.

Example 3.14. The extension $\mathrm{PS^*L}(2,\mathbb{Z})$ of the modular group to include orientation-reversing isometries with integer coefficients is an example of a discrete group of $\mathbb{H}$-isometries which is not Fuchsian.

We will now zoom in on the action on $\mathbb{H}$ by Fuchsian groups. Let Γ be a subgroup of $\mathrm{PSL}(2,\mathbb{R})$ and let Id denote the identity in Γ. Of special interest are the orbits under Γ. We first remind ourselves of the definition.

Definition 3.15. The **Γ-orbit** of a point z in $\mathbb{H}$ is the family $\{T(z) \mid T \in \Gamma\}$, denoted by Γz. We count each point of Γz with multiplicity equal to the order of the stabilizer $\mathrm{Stab}_\Gamma(z)$ of z in Γ. We say that two points $z_1, z_2 \in \mathbb{H}$ are Γ-**equivalent** if they belong to the same orbit.

Recall the definition of a locally finite collection of subsets (Definition 3.5). As $\mathbb{H}$ is locally compact, we can use Proposition 3.6 to define what it means for an orbit to be locally finite.

Definition 3.16. Γz is called **locally finite** if for any compact subset $K \subset \mathbb{H}$, the set $\{T \in \Gamma \mid T(z) \cap K \neq \varnothing\}$ is finite.

Definition 3.17. Γ acts **properly discontinuously** on $\mathbb{H}$ if the Γ-orbit of any point of $\mathbb{H}$ is locally finite.

The aim of this section is to prove that Fuchsian groups act properly discontinuously on $\mathbb{H}$. Definition 3.17 may seem a bit technical, the following proposition gives equivalent characterizations of properly discontinuous actions.

Proposition 3.18. *The following are equivalent.*

(i) Γ *acts properly discontinuously on* $\mathbb{H}$.

(ii) For any $z \in \mathbb{H}$ *the orbit* Γz *is discrete and the stabilizer* $\mathrm{Stab}_\Gamma(z)$ *is finite.*

(iii) Each point $z \in \mathbb{H}$ *has a neighbourhood* V *such that*

$$T(V) \cap V \neq \varnothing \text{ for only finitely many } T \in \Gamma. \tag{3.2}$$

Proof. Let $z \in \mathbb{H}$.

(i) $\Longrightarrow$ (ii): Suppose the stabilizer of z is infinite. Let K be a compact subset of $\mathbb{H}$ containing z (such as $\{z\}$). Then $T(z) = z \in K$ for infinitely many $T \in \Gamma$, and so the orbit of z is not locally finite. Now, suppose the orbit of z is not discrete. Then there is $x \in \Gamma z$ such that any neighbourhood of x contains infinitely many distinct elements of Γz. Let K be a compact subset of $\mathbb{H}$ containing a neighbourhood of x. Then K contains infinitely many distinct elements of Γz, and so Γz is not locally finite.

(ii) $\Longrightarrow$ (i): Suppose (ii) holds. Let K be a compact subset of $\mathbb{H}$. Since Γz is discrete, the set $\Gamma z \cap K$ is finite. For each point x in Γz, the stabilizer is finite and so the set $\{T \in \Gamma \,|\, T(z) = x\}$ is finite. It follows that the set $\{T \in \Gamma \,|\, T(z) \in K \cap \Gamma z\} = \{T \in \Gamma \,|\, T(z) \in K\}$ is a finite union of finite sets, so is finite.

(ii) $\Longrightarrow$ (iii): Suppose (ii) holds. Since Γz is discrete, there exists a neighbourhood V of z containing no points of Γz other than z. Then $T(V) \cap V \neq \varnothing$ implies that $T \in \mathrm{Stab}_\Gamma(z)$. This is possible only for finitely many T, or the stabilizer would be infinite.

(iii) $\Longrightarrow$ (ii): Suppose Γz is not discrete. Then there is $x \in \Gamma z$ such that any neighbourhood of x contains infinitely many elements of Γz, contradicting (3.2). Now suppose the stabilizer of z is infinite. Then any neighbourhood of z contains infinitely many of its images, again contradicting (3.2).

■

For the proof of the main result of this section we need two lemmas.

Lemma 3.19. *Let $z_0 \in \mathbb{H}$ and K a compact subset of $\mathbb{H}$. Then the set*

$$E = \{T \in \mathrm{PSL}(2,\mathbb{R}) \,|\, T(z_0) \in K\}$$

is compact.

Proof. Since $\mathrm{PSL}(2,\mathbb{R})$ is topologized as a quotient space of $\mathrm{SL}(2,\mathbb{R})$ the quotient map $\psi : \mathrm{SL}(2,\mathbb{R}) \to \mathrm{PSL}(2,\mathbb{R})$ defined by (2.6) is continuous. We prove that the set

$$E_1 = \{T \in \mathrm{SL}(2,\mathbb{R}) \,|\, T(z_0) \in K\}$$

is compact, it then follows that $E = \psi(E_1)$ is compact. By continuity of Möbius transformations, the map $\beta : \mathrm{SL}(2,\mathbb{R}) \to \mathbb{H}$ defined by

$$\beta(A) = \psi(A)(z_0)$$

is continuous. Since $E_1 = \beta^{-1}(K)$ and K is closed, so is E_1. It remains to prove that E_1 is bounded. Since K is bounded there exists $M_1 > 0$ such that for all $\begin{bmatrix} a & b \\ c & d \end{bmatrix} \in E_1$

$$\left| \frac{az_0 + b}{cz_0 + d} \right| < M_1. \tag{3.3}$$

Since K is compact in $\mathbb{H}$ there exists $M_2 > 0$ such that for all $\begin{bmatrix} a & b \\ c & d \end{bmatrix} \in E_1$

$$\operatorname{Im}\left(\frac{az_0 + b}{cz_0 + d} \right) \geq M_2.$$

Using (2.4) we rewrite this as

$$|cz_0 + d| \leq \sqrt{\operatorname{Im}(z_0)/M_2}.$$

Combining with (3.3) gives

$$|az_0 + b| \leq M_1 \sqrt{\operatorname{Im}(z_0)/M_2}.$$

We have proved that $|az_0+b|$ and $|cz_0+d|$ are bounded. It follows that a, b, c, d are bounded. To see this, suppose that a is unbounded. Then $\operatorname{Im}(az_0 + b)$ is unbounded and hence $|az_0 + b|$ is. So a must be bounded. Then suppose b is unbounded. It follows that $\operatorname{Re}(az_0 + b)$ is unbounded and hence $|az_0 + b|$ is. We conclude that E_1 is bounded. ■

Lemma 3.20. *Let Γ be a subgroup of $\mathrm{PSL}(2, \mathbb{R})$ acting properly discontinuously on $\mathbb{H}$ and suppose $p \in \mathbb{H}$ is fixed by some element of Γ. Then there is a neighbourhood V of p such that no other point of V is fixed by a nontrivial element of Γ.*

Proof. Suppose T fixes p for some nonidentity element T of Γ, and assume for contradiction that any neighbourhood of p contains fixed points of transformations in Γ. This means that there is a sequence $\{p_n\}$ in $\mathbb{H}$ converging to p with $T_n(p_n) = p_n$ for $T_n \in \Gamma$. Let $\epsilon > 0$. Since Γ acts properly discontinuously on $\mathbb{H}$ the set $\{T \in \Gamma \,|\, T(p) \in \overline{B_\epsilon(p)}\}$ is finite. These two claims imply that there exists an integer N such that for all $n > N$, $d_{\mathbb{H}}(p_n, p) < \epsilon/2$ while $d_{\mathbb{H}}(T_n(p), p) > \epsilon$. Using the triangle inequality and the invariance of the hyperbolic metric we reach a contradiction: for all $n > N$ we have

$$\begin{aligned} d_{\mathbb{H}}(T_n(p), p) &\leq d_{\mathbb{H}}(T_n(p), T_n(p_n)) + d_{\mathbb{H}}(T_n(p_n), p) \\ &= d_{\mathbb{H}}(p, p_n) + d_{\mathbb{H}}(p_n, p) < \epsilon. \end{aligned}$$

■

Lemma 3.20 enables us to pick a point of $\mathbb{H}$ that is not fixed by any element of Γ other than the identity. This is crucial for the construction of the Dirichlet polygon which we will discuss in the next section.

Theorem 3.21. *A subgroup* Γ *of* $\mathrm{PSL}(2,\mathbb{R})$ *is Fuchsian if and only if* Γ *acts properly discontinuously on* $\mathbb{H}$.

Proof. Suppose Γ is Fuchsian. Let $z \in \mathbb{H}$ and K a compact subset of $\mathbb{H}$. Consider the set

$$\{T \in \Gamma \,|\, T(z) \in K\} = \{T \in \mathrm{PSL}(2,\mathbb{R}) \,|\, T(z) \in K\} \cap \Gamma.$$

The first set in the intersection is compact by Lemma 3.19, the second is discrete by assumption. It follows that the intersection is finite. So Γ acts properly discontinuously on $\mathbb{H}$. For the converse, suppose Γ is not discrete. Then there is a sequence $\{T_k\}$ of distinct elements in Γ converging to the identity. By Lemma 3.20 there is a point $s \in \mathbb{H}$ not fixed by any element of the sequence. Then $\{T_k(s)\}$ is a sequence of points distinct from s converging to s. Hence every closed hyperbolic disc centred at s contains infinitely many points of the orbit, i.e. Γ does not act properly discontinuously on $\mathbb{H}$. ∎

Corollary 3.22. *A subgroup* Γ *of* $\mathrm{PSL}(2,\mathbb{R})$ *acts properly discontinuously on* $\mathbb{H}$ *if and only if for all* $z \in \mathbb{H}$, Γz *is a discrete subset of* $\mathbb{H}$.

Proof. Let $z \in \mathbb{H}$. If Γ acts properly discontinuously on $\mathbb{H}$ then by Proposition 3.18 it follows that Γz is discrete. Suppose the action of Γ is not properly discontinuous. Then by Theorem 3.21, Γ is not discrete. We pick $s \in \Gamma z$ and as in the proof above construct a sequence $\{T_k(s)\}$ of points distinct from s converging to s. Thus any neighbourhood of s contains infinitely many elements of Γz and so Γz is not discrete. ∎

We close this section with a nice property of Fuchsian groups.

Proposition 3.23. *Any Fuchsian group is countable.*

Proof. Let Γ be a Fuchsian group and assume for contradiction that Γ is uncountable. Pick $z \in \mathbb{H}$. By Proposition 3.18 and Theorem 3.21 the Γ-orbit of z is discrete and the stabilizer finite. The order of the orbit is equal to the index of the stabilizer, why the orbit must be uncountable. Let $\mathcal{B}$ be a countable basis for the topology on $\mathbb{H}$. We reach a contradiction by constructing an injection from the uncountable set Γz to the countable set $\mathcal{B}$. Since Γz is discrete, each point $x \in \Gamma z$ is contained in a basis element $B_x \in \mathcal{B}$

containing no other points of Γz. Thus the map $\Gamma z \to \mathcal{B}$ which sends $x \in \Gamma z$ to the basis element B_x is injective. ■

3.3 Fundamental Polygons

`Playing Tetris with Group Theory`

Fundamental polygons are the building blocks of tessellations by Fuchsian groups. In this section we will see that any Fuchsian group possesses a fundamental polygon.

Definition 3.24. Let $\mathcal{H} = \{H_\alpha\}_{\alpha \in A}$ be a collection of half-planes in $\mathbb{H}$. For each $\alpha \in A$, let L_α be its bounding line. We say that $\mathcal{H}$ is **locally finite** if each point z in $\mathbb{H}$ has a neighbourhood V intersecting only finitely many of the bounding lines L_α

Definition 3.25. A **hyperbolic polygon** D is a closed subset of $\mathbb{H}$ with nonempty interior that can be expressed as the intersection of a locally finite collection of half-planes.

Let D be a hyperbolic polygon and let L be a hyperbolic line. If the intersection $\partial D \cap L$ of L with the boundary ∂D of D is nonempty, we call it a **vertex** if it is a single point, otherwise we call it an **edge** or a **side**. Each side of D is contained in one of the bounding lines of the half-planes that define D, and each vertex is the intersection of two sides.[6] If two sides are asymptotic, the corresponding vertex on the circle at infinity is called an **ideal vertex**. Note that it follows from Proposition 3.6 that a compact polygon has finitely many sides.

Definition 3.26. We say that a polygon is **regular** if its sides are of equal length, and all of its interior angles are equal.

Definition 3.27. A **fundamental domain** for a Fuchsian group Γ is an open set F in $\mathbb{H}$ with the following properties.

(i) No two points of F are Γ-equivalent.

(ii) Every point in $\mathbb{H}$ is Γ-equivalent to a point in the closure $\overline{F}$ of F in $\mathbb{H}$.

A **fundamental polygon** for Γ is a polygon whose interior is a fundamental domain for Γ.

[6] [1], p.158 and Lemma 5.9.

We will now present a method of constructing a fundamental polygon for a Fuchsian group. Let Γ be a Fuchsian group, let Γ' be the set of nontrivial elements of Γ, and let p be a point not fixed by any element of Γ'. For each $T \in \Gamma'$, let L_T be the equidistant set of p and $T(p)$. Let H_T be the closed half-plane determined by L_T containing p. We define the **Dirichlet polygon** centred at p to be the intersection

$$D_\Gamma(p) = \bigcap_{T \in \Gamma'} H_T.$$

Note that

$$H_T = \{w \in \mathbb{H} \mid d_{\mathbb{H}}(p, w) \leq d_{\mathbb{H}}(T(p), w)\},$$

so

$$D_\Gamma(p) = \{w \in \mathbb{H} \mid d_{\mathbb{H}}(p, w) \leq d_{\mathbb{H}}(T(p), w) \text{ for all } T \in \Gamma'\}$$

is an equivalent definition of the Dirichlet polygon.

Proposition 3.28. *$D_\Gamma(p)$ is a fundamental polygon for Γ.*

Proof. As Γp is discrete, $D_\Gamma(p)$ contains a neighbourhood of p and so has nonempty interior. $D_\Gamma(p)$ is defined as the intersection of a collection of closed sets, so it is closed. To prove that $D_\Gamma(p)$ is a polygon, it remains to prove that the collection of half-planes $\{H_T \mid T \in \Gamma'\}$ is locally finite. For any $\epsilon > 0$, we have by construction of L_T that $L_T \cap \overline{B_\epsilon(p)} \neq \varnothing$ if and only if $T(p) \in \overline{B_{2\epsilon(p)}}$. As the orbit of p is discrete, $\Gamma p \cap \overline{B_{2\epsilon(p)}}$ contains only finitely many elements of Γp. So $\overline{B_\epsilon(p)} \cap L_T \neq \varnothing$ for finitely many $T \in \Gamma'$. Let $z \in \mathbb{H}$, $\delta > 0$, and $\epsilon = d_{\mathbb{H}}(p, z) + \delta$. Then $B_\delta(z) \subseteq \overline{B_\epsilon(p)}$. So $B_\delta(z)$ intersects finitely many of the bounding lines L_T. This completes the proof that $D_\Gamma(p)$ is a polygon.

Next, we prove that $D_\Gamma(p)$ contains representatives of every orbit of Γ. Let $w \in \mathbb{H}$. As Γp is discrete, there exists a point $T_w(p) \in \Gamma p$ with smallest distance to w. That is,

$$d_{\mathbb{H}}(w, T_w(p)) \leq d_{\mathbb{H}}(w, T(p)) \text{ for all } T \in \Gamma.$$

By the invariance of the hyperbolic metric under Γ we may rewrite this as

$$d_{\mathbb{H}}(T_w^{-1}(w), p) \leq d_{\mathbb{H}}(T_w^{-1}(w), T_w^{-1}T(p)) \text{ for all } T \in \Gamma.$$

Any element of Γ can be written as $T_w^{-1}T$ for some $T \in \Gamma$, so this implies that $T_w^{-1}(w)$ is contained in $D_\Gamma(p)$.

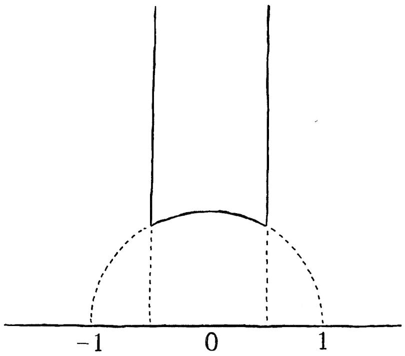

Figure 3.1. Fundamental polygon for $\mathrm{PSL}(2,\mathbb{Z})$.

Lastly, we prove that $\mathrm{Int}D_\Gamma(p)$ contains no two points from the same orbit. Assume the contrary. Then for some $T \in \Gamma'$, $\mathrm{Int}D_\Gamma(p) \cap T(\mathrm{Int}D_\Gamma(p))$ contains a point x. Since $x \in \mathrm{Int}D_\Gamma(p)$ we have

$$d_{\mathbb{H}}(x,p) < d_{\mathbb{H}}(x,T(p)). \tag{3.4}$$

Since $x \in T(\mathrm{Int}D_\Gamma(p))$ we have $x = T(y)$ for some $y \in \mathrm{Int}D_\Gamma(p)$. So

$$d_{\mathbb{H}}(y,p) < d_{\mathbb{H}}(y,T^{-1}(p)).$$

T is an isometry so this is equivalent to

$$d_{\mathbb{H}}(x,T(p)) < d_{\mathbb{H}}(x,p),$$

contradicting (3.4). ■

Example 3.29 (Modular Group)**.** In this example we construct the Dirichlet polygon for $\Gamma = \mathrm{PSL}(2,\mathbb{Z})$ centred at $2i$. It is easily verified that $2i$ is not fixed by any non-identity element of Γ. We claim that the closure $\overline{F}$ of the set

$$F = \{z \in \mathbb{H} \,|\, |z| > 1, |\mathrm{Re}(z)| < \frac{1}{2}\}$$

is the Dirichlet polygon for Γ centred at $2i$. Note that $\overline{F}$ is the (asymptotic) triangle with angles $0, \frac{\pi}{3}, \frac{\pi}{3}$ bounded by the lines $\mathrm{Re}(z) = \pm\frac{1}{2}$ and $|z| = 1$. The isometries $s(z) = -\frac{1}{z}$ and $t(z) = z + 1$ are clearly in Γ. As above, let L_s denote the equidistant set of $2i$ and $s(2i)$, and so on. We have that $|z| = 1$ is the line L_s, $\mathrm{Re}(z) = \frac{1}{2}$ is the line L_t, and $\mathrm{Re}(z) = -\frac{1}{2}$ is the line $L_{t^{-1}}$. The intersection of the half-planes associated to these lines is precisely $\overline{F}$, and so $D_\Gamma(2i) \subset \overline{F}$. This shows that any point of $\mathbb{H}$ is Γ-equivalent to a point in $\overline{F}$. To complete the proof that $\overline{F}$ is the Dirichlet polygon for Γ, we prove that $F \cap T(F) \neq \varnothing$ implies $T = \mathrm{Id}$. Suppose there is $T \in \Gamma'$ with $z, T(z) \in F$. Put $T(z) = \frac{az+b}{cz+d}$ and recall that $\mathrm{Im}(T(z)) = \mathrm{Im}(z)/|cz+d|^2$ by (2.4). Then since $|z| > 1$ and $\mathrm{Re}(z) > -\frac{1}{2}$ we have

$$|cz+d|^2 = c^2|z|^2 + 2\mathrm{Re}(z)cd + d^2 > c^2 - |cd| + d^2 = (|c| - |d|)^2 + |cd| \geq 1.$$

The last inequality holds since $|cd|$ is nonzero and $c, d \in \mathbb{Z}$. Replacing z with $T(z)$ and T with T^{-1} in the argument above yields $\mathrm{Im}(z) = \mathrm{Im}(T(z))/|cz-a|^2$, and using $|z| > 1$ and $\mathrm{Re}(z) < \frac{1}{2}$ we obtain

$$|cz-a|^2 = c^2|z|^2 - 2\mathrm{Re}(z)ca + a^2 > c^2 - |ca| + d^2 = (|c| - |a|)^2 + |ca| \geq 1.$$

Putting it all together gives $\mathrm{Im}(z) < \mathrm{Im}(T(z)) < \mathrm{Im}(z)$, a contradiction.

4 The Art

Let us take a look at all the tools scattered on the hyperbolic floor. We have seen how to, given a Fuchsian group, construct a fundamental polygon for its action on $\mathbb{H}$. Ideally, we would like to reverse this process: start with a polygon D, and find a discrete group which has D as a fundamental domain. This is the content of Poincaré's Polygon Theorem, which we will now give our undivided attention to.

4.1 Tessellations

`Enough with the Secrecy!`

We have avoided it long enough. Let us come clean and state what we mean by a tessellation of the hyperbolic plane. For a clear view, it is helpful to broaden the perspective and allow the tessellated space to be either the euclidean plane, the sphere, or the hyperbolic plane. Let X be any of the aforementioned.

Definition 4.1. A **tessellation** $\mathcal{T}$ of X is a family of congruent polygons in X such that

(i) $\mathrm{Int}D \cap \mathrm{Int}D' = \varnothing$ for distinct $D, D' \in T$, and

(ii) $\bigcup_{D \in T} D = X$.

Each individual polygon in $\mathcal{T}$ is called a **tile**. We say that a tessellation is **regular** if the tiles are regular polygons.

For $\mathcal{T}$ to be a tessellation of X it is necessary that at any point $p \in X$ where polygons of $\mathcal{T}$ meet at their vertices, the angles of the polygons at p sum to 2π, so that the union of the polygons meeting at p contains a neighbourhood of p. In the case of a regular tessellation, this imposes the condition that the angle of the regular polygon generating the tessellation is a submultiple of 2π. This is why the regular tessellations of the euclidean plane and the sphere are limited to a finite number. In the hyperbolic case the angle condition imposes no threat and the possibilities are endless.

It is becoming very obvious that fundamental polygons are so intimately related to tessellations that we may equally well define the former in terms of the latter.

Proposition 4.2. *Let Γ be a Fuchsian group and let F be a hyperbolic polygon. Then F is a fundamental polygon for Γ if and only if the family $\{T(F) \mid T \in \Gamma\}$*

is a tessellation of $\mathbb{H}$.

Note that the polygons $T(F)$ are congruent since length, area[7] and angles are preserved by $\mathrm{PSL}(2,\mathbb{R})$. We have already found an example of a hyperbolic tessellation by the Dirichlet polygon of the modular group given in Example 3.29. In the next section we will present the Poincaré Polygon Theorem, which will enable us to construct a multitude of tessellations with ease.

4.2 The Poincaré Polygon Theorem

`Endgame`

Before stating the theorem, there are a few things that need to be defined.

Definition 4.3. Let D be a hyperbolic polygon and s a side of D. An element $g_s \in \mathrm{Iso}(\mathbb{H})$ is called a **side pairing transformation** associated to s if there is a side s' of D (not necessarily distinct from s) so that

(i) $g_s(s') = s$, and

(ii) $g_s(D) \cap D = s$.

Definition 4.4. A **side pairing** Φ of a polygon D is a choice of side pairing transformations g_s for each side s such that $g_{s'} = g_s^{-1}$. If $g_s \in \Phi$ with $g_s(s) = s'$ for sides s, s' in D, then points $w \in s$ and $w' = g_s(w) \in s'$ are said to be **identified** by g_s. A **vertex cycle** is a chain of identified vertices $\{v_1, \ldots, v_k\}$ where v_i is identified with v_{i+1} by some $g_i \in \Phi$.

Now, let D be a compact hyperbolic polygon equipped with a side pairing Φ. Let v_1 be a vertex of D. Pick one of the sides intersecting at v_1 and call it s_1. Let s_1' be the side of D paired with s_1 by $g_1 \in \Phi$, so that $s_1' = g_1(s_1)$. Set $v_2 = g_1(v_1)$. s_1' shares the endpoint v_2 with another side of D which we label as s_2. Let s_2' be the side of D paired with s_2 by $g_2 \in \Phi$ so that $s_2' = g_2(s_2)$. Set $v_3 = s_2(v_2)$. v_3 is an endpoint of s_2', and of another side of D which we label as s_3. Continuing in this manner, we obtain a sequence $\{g_i\}$ of side pairing transformations, a sequence $\{v_i\}$ of identified vertices, and a sequence $\{(s_i, s_i')\}$ of pairs of sides, where s_i' and s_{i+1} are adjacent. Since D has finitely many sides, we will eventually reach the side we started off with, and $g_n(v_n) = v_1$ for some n. Set

$$\phi = g_n \circ \cdots \circ g_1.$$

[7][3], Theorem 1.4.1.

Definition 4.5. We call ϕ a **cycle transformation** associated to the vertex v_1. We let τ denote the order of ϕ. We define the **angle sum** of v_1 to be

$$\text{sum}(v_1) = \sum_{k=1}^{n} \alpha(v_k),$$

where v_k is a vertex in the vertex cycle and $\alpha(v_k)$ is the interior angle of D at v_k.

We can view the side pairing transformations as a recipe for gluing together translates of D. The translates of D that share the vertex v_1 are

$$g_1^{-1}(D), g_1^{-1}g_2^{-1}(D), \dots, \phi^{-1}(D), g_1^{-1}\phi^{-1}(D), g_1^{-1}g_2^{-1}\phi^{-1}(D), \dots, \phi^{-\tau}(D). \tag{4.1}$$

For these translates to contain a neighbourhood of v_1, the interior angles at v_1 must sum to 2π. We say that D satisfies the **cycle condition** if for any vertex v of D, $\tau\text{sum}(v) = 2\pi$.

Theorem 4.6 (Poincaré). *Let D be a compact hyperbolic polygon equipped with a side pairing Φ satisfying the cycle condition. Then the subgroup of* $\text{Iso}(\mathbb{H})$ *generated by Φ is discrete and has D as a fundamental polygon.*

Let us take a moment to contemplate. Proposition 4.2 tells us that what Poincaré's Theorem amounts to is claiming that when the conditions are met, $\mathbb{H}$ is tessellated by images of D under Γ. The purpose of the cycle condition is to guarantee that there is a local tessellation at any vertex of D; the theorem claims that this is enough to guarantee a tessellation of all of $\mathbb{H}$. However, let us not get too lost in intuition before we lay out the proof.

4.3 Proof of Poincaré's Theorem

Let's Get Our Hands Dirty

We go about this in a somewhat roundabout way, by constructing a space $\mathbb{H}^*$ that is tessellated by D under the action of the group Γ generated by Φ, and then showing that this space is homeomorphic to $\mathbb{H}$. With a little help from Section 3.1, the last step boils down to proving that the natural map from $\mathbb{H}^*$ to $\mathbb{H}$ satisfies the conditions of Lemma 3.9, and so is a covering map. As the spaces in question are sufficiently nice, Theorem 3.10 applies, from which it follows that they are in fact homeomorphic.

Consider the space $\Gamma \times D$. We may view this space as the disjoint union of copies of D, indexed by elements of Γ. We will define an equivalence relation on this space generated by Φ, which has the effect of gluing copies of D at their adjacent edges. Let s, s' be sides of D with $g_s \in \Phi$ taking s to s'. For any $h \in \Gamma$ and $x \in s$, define $\sim$ by $(hg_s, x) \sim (h, g_s x)$. For interior points of D we define $(h, y) \sim (g, x)$ if and only if $h = g$ and $x = y$. This relation is symmetric and reflexive, but fails to be transitive. Let $*$ be the transitive closure of $\sim$. That is, $(g, x) * (h, y)$ if and only if there is a chain of points (g_i, x_i) of $\Gamma \times D$ with

$$(g, x) = (g_1, x_1) \sim (g_2, x_2) \sim \cdots \sim (g_n, x_n) = (h, y).$$

We denote the equivalence class of (g, x) by $\langle g, x \rangle$. Note that

$$\langle g, x \rangle = \langle h, y \rangle \implies gx = hy, \tag{4.2}$$

and for $x, y \in \mathrm{Int} D$

$$\langle g, x \rangle = \langle h, y \rangle \iff g = h \text{ and } x = y. \tag{4.3}$$

We give Γ the discrete topology, D the subspace topology and $\Gamma \times D$ the product topology. Once we prove that Γ has a fundamental domain in $\mathbb{H}$, it follows that Γ is discrete in $\mathrm{Iso}(\mathbb{H})$, which is what we ultimately want. We now define $\mathbb{H}^*$ to be the quotient of $\Gamma \times D$ by $*$, endowed with the quotient topology. The natural map

$$\begin{aligned} \alpha : \mathbb{H}^* &\to \mathbb{H} \\ \langle g, x \rangle &\mapsto gx \end{aligned}$$

is well-defined by (4.2).

The action map $\Gamma \times \mathbb{H} \to \mathbb{H}$ restricts to a map

$$\begin{aligned} \gamma : \Gamma \times D &\to \mathbb{H} \\ (g, x) &\mapsto gx. \end{aligned}$$

γ factors through the quotient map

$$\begin{aligned} \beta : \Gamma \times D &\to \mathbb{H}^* \\ (g, x) &\mapsto \langle g, x \rangle \end{aligned}$$

and the following diagram commutes.

$$\begin{tikzcd} \Gamma \times D \arrow[r, "\beta"] \arrow[dr, "\gamma"'] & \mathbb{H}^* \arrow[d, "\alpha"] \\ & \mathbb{H} \end{tikzcd}$$

The quotient map β is automatically continuous. We prove that γ is continuous. Let A be an open subset of $\mathbb{H}$. The preimage of A under γ is the set $\bigcup_{g\in\Gamma}\{g\} \times (g^{-1}(A) \cap D) = \{(g,x) \in \Gamma \times D \mid gx \in A\}$. By continuity of g, $g^{-1}(A)$ is open in $\mathbb{H}$, so $g^{-1}(A) \cap D$ is open in D. Finally, $\{g\}$ is open in Γ. It now follows from the definition of the product topology that $\gamma^{-1}(A)$ is open in $\Gamma \times D$. $\alpha^{-1}(A)$ is open in $\mathbb{H}^*$ if and only if $\beta^{-1}\alpha^{-1}(A)$ is open in $\Gamma \times D$. Hence continuity of α follows from continuity of γ.

We define a metric on $\mathbb{H}^*$ as follows. Let $a, b \in \mathbb{H}^*$, and consider sequences of points in $\Gamma \times D$ of the form

$$p_1, q_1, p_2, q_2, \ldots, p_n, q_n$$

where $a = \beta(p_1)$, $\beta(q_i) = \beta(p_{i+1})$, $b = \beta(q_n)$, and p_i, q_i belong to the same copy of D in $\Gamma \times D$. We define the distance between points $(g,x),(g,y)$ in $\Gamma \times D$ to be $d((g,x),(g,y)) = d_{\mathbb{H}}(x,y)$. Then the distance between a and b in $\mathbb{H}^*$ is

$$d_*(a,b) = \inf \Sigma_{i=1}^{n} d(p_i, q_i),$$

where the infimum is taken over all such possible sequences of points.

Note that

$$d_*(a,b) \geq d_{\mathbb{H}}(\alpha(a), \alpha(b)), \tag{4.4}$$

since the image of such a path in $\mathbb{H}$ will be a polygonal path, for which the length is greater than or equal to the geodesic segment connecting $\alpha(a)$ and $\alpha(b)$.

Each $f \in \Gamma$ induces a map

$$\begin{aligned} f^* : \mathbb{H}^* &\to \mathbb{H}^* \\ \langle g, x\rangle &\mapsto \langle fg, x\rangle. \end{aligned}$$

If $\langle g,x\rangle = \langle h,y\rangle$, it follows from 4.2 that $\langle fg,x\rangle = \langle fh,y\rangle$, so f^* is well-defined. Furthermore, the maps f^* are bijections of $\mathbb{H}^*$, and as d_* is defined in terms of the hyperbolic metric, which is preserved by Γ, we see that the maps f^* are in fact isometries of $\mathbb{H}^*$.

Set

$$\langle D\rangle = \{\langle \mathrm{Id}, x\rangle \mid x \in D\}$$

and similarly for the interior of D. As $g^*\langle D\rangle = \{\langle g,x\rangle \mid x \in D\}$, we see that

$$\bigcup_{g\in\Gamma} g^*\langle D\rangle = \mathbb{H}^*. \tag{4.5}$$

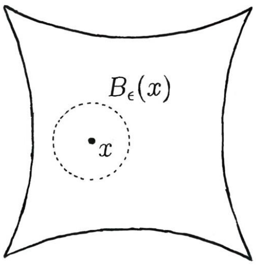

Figure 4.1.

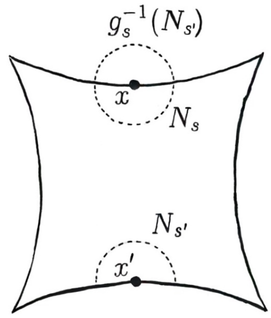

Figure 4.2.

Furthermore, for distinct $g, h \in \Gamma$ we have by (4.3) that

$$g^*\langle \mathrm{Int}D\rangle \cap h^*\langle \mathrm{Int}D\rangle = \varnothing.$$

In other words, $\langle D\rangle$ is a fundamental polygon for Γ. From this we see that if α is surjective then $\bigcup_{g\in\Gamma} g(D) = X$, and if α is injective then for distinct g, h in Γ, $g(\mathrm{Int}D) \cap h(\mathrm{Int}D) = \varnothing$. Thus if we can prove that α is a homeomorphism, we find that the Γ-images of D tessellate $\mathbb{H}$.

This is where Lemma 3.9 comes into play. We will prove that α is a covering map by proving that it is a local isometry, and that there is a lower bound on the radius of neighbourhoods on which α restricts to an isometry. Once we know that α is a covering map it follows that it is a homeomorphism, since $\mathbb{H}$ is simply connected.

To this end, we construct a system of neighbourhoods of points in D, and simultaneously a system of neighbourhoods of points in $\langle D\rangle$. More precisely, we find for each point x in D a neighbourhood in $\mathbb{H}^*$ such that α restricts to a bijection from this neighbourhood to $B_\epsilon(x)$. There are three cases to consider: (a) interior points; (b) points on single edges of D; and (c) vertices of D. Cases (a) and (b) are illustrated in Figure 4.1 and Figure 4.2 respectively.

(a) $x \in \mathrm{Int}D$. Let $\epsilon > 0$ be less than the distance from x to ∂D, so that $B_\epsilon(x) \subset \mathrm{Int}D$. Then $\alpha^{-1}(B_\epsilon(x)) = \{\langle \mathrm{Id}, y\rangle \,|\, y \in B_\epsilon(x)\}$, and by (4.3), $\langle \mathrm{Id}, y\rangle)$ contains only (Id, y). It is clear that α restricts to a bijection from $\langle \mathrm{Id}, B_\epsilon(x)\rangle$ in $\mathbb{H}^*$ to $B_\epsilon(x)$ in $\mathbb{H}$.

(b) x is contained in the interior of a side s of D. Then there is $g_s \in \Phi$ with $g_s(s) = s'$ for some side s' of D. Set $x' = g_s(x)$. x is contained in exactly one other copy of D, namely $g_s^{-1}(D)$. Let $\epsilon_1 > 0$ be less than the distance from x to $\partial D - s$, $\epsilon_2 > 0$ be less than the distance from x' to $\partial D - s'$, and set $\epsilon = \min(\epsilon_1, \epsilon_2)$. Let $N_s = B_\epsilon(x) \cap D$, and $N_{s'} = B_\epsilon(x') \cap D$. Then $B_\epsilon(x) = N_s \cup g_s^{-1}(N_{s'})$. Note that Definition 4.3 (ii) ensures that N_s and $g_s^{-1}N_{s'}$ overlap only on the boundary. We have that $B_\epsilon(x) = \alpha(\langle \mathrm{Id}, N_s \rangle \cup \langle \mathrm{Id}, g_s^{-1}(N_{s'}) \rangle)$. This set consists of single point equivalence classes of interior points e.g. (Id, x) for $x \in N_s$ and (g_s^{-1}, x) for $x \in N_{s'}$, and two-point equivalence classes of points in s, for example $\langle \mathrm{Id}, x \rangle$ contains (Id, x) and (g_s^{-1}, x'). Again we see that α restricts to a bijection from $\langle \mathrm{Id}, N_s \rangle \cup \langle \mathrm{Id}, g_s^{-1}(N_{s'}) \rangle$ to $B_\epsilon(x)$.

(c) x is a vertex of D. This is where we need to make use of the cycle condition. Let ϕ be a cycle transformation associated to x. Consider the translates of D under ϕ sharing the vertex x as in (4.1). Relabel them as $D_1, D_2, \ldots, D_k = \phi^{-\tau}(D) = D$. Let $h_i \in \Gamma$ be the element taking D to D_i (for example, as $D_1 = g_1^{-1}(D)$, we have $h_1 = g_1^{-1}$). Set $x_i = h_i^{-1}(x)$. Let s_i, s_{i-1} be the sides adjacent to x_i. Let $\epsilon_i > 0$ be less than the distance between x_i and $\partial D - \{s_i \cup s_{i-1}\}$, and let $\epsilon = \min(\epsilon_1, \ldots, \epsilon_n)$. Set $N_i = B_\epsilon(x_i) \cap D$. Again, the sets $h_i(N_i)$ overlap only on their boundaries. Furthermore, the cycle condition guarantees that $B_\epsilon(x) = \bigcup_{i=1}^n h_i(N_i)$. As in case (b), the preimage of $B_\epsilon(x)$ contains one-point equivalence classes of interior points of N_i, and two-point equivalence classes of pairs of points on sides of D. It also contains the equivalence class $\langle \mathrm{Id}, x \rangle$, which consists of the points (h_i, x_i). As in the previous cases we find that α restricts to a bijection from $\langle \mathrm{Id}, B_\epsilon(x) \rangle$ to $B_\epsilon(x)$.

We want to prove that the restriction of α to each of these neighbourhoods is an isometry. The next step is thus to examine how it treats distance between points in these neighbourhoods. This is accomplished by employing (4.4), which gives us a lower bound on the distance between points in $\mathbb{H}^*$. We construct a path connecting two points in $\mathbb{H}^*$ whose length is precisely the hyperbolic distance between their images under α. We can then infer that this is the distance between the points in $\mathbb{H}^*$.

(a) In this case $B_\epsilon(x)$ is contained in D, so all points of $\langle \mathrm{Id}, B_\epsilon(x) \rangle$ belong to the same copy of D. So for any $a, b \in B_\epsilon(x)$, we choose the geodesic segment in $\Gamma \times D$ from (Id, a) to (Id, b) whose length is $d_{\mathbb{H}}(a, b)$. Then by (4.4) we have $d_*(\alpha^{-1}(a), \alpha^{-1}(b)) = d_{\mathbb{H}}(a, b)$.

(b) Here we have to consider points that belong to different copies of D, say, $a \in N_s, b \in g_s^{-1}(N_{s'})$. Let c be the point where the geodesic segment $[a, b]$ intersects ∂D. Set $c' = g_s(c)$. Consider the polygonal path with breakpoints

$$(\mathrm{Id}, a), (\mathrm{Id}, c), (g_s^{-1}, c'), (g_s^{-1}, b).$$

The length of this path is $d_{\mathbb{H}}(a, b)$, so $d_*(\alpha^{-1}(a), \alpha^{-1}(b)) = d_{\mathbb{H}}(a, b)$.

(c) For $a, b \in B_\epsilon(x)$ we have two new cases to consider: 1) the geodesic segment $[a, b]$ intersects more than two copies of D, and 2) $[a, b]$ intersects the common vertex x. As the number of copies of D intersecting at x is finite, we can proceed as in (b) to construct a finite sequence of points in $\Gamma \times D$, giving us a polygonal path whose length is $d_{\mathbb{H}}(a, b)$. Whenever the geodesic segment $[a, b]$ passes through the vertex x, it will do so from one copy of D to another, so this case reduces to the case in (b).

We have found that α is a local isometry when restricted to a neighbourhood of $\langle D \rangle$. We know that Γ acts on $\langle D \rangle$ by isometries. We also know by (4.5) that the translates of $\langle D \rangle$ by elements of Γ cover $\mathbb{H}^*$. This concludes that α is a local isometry.

To apply Lemma 3.9, we need to show that the radius of neighbourhoods for which α restricts to an isometry can be defined globally on $\mathbb{H}^*$, independent of x. Here we use the assumption that D is compact. Similarly to the argument above, we first find a $\delta > 0$ such that the restriction of α to $B_\delta(x)$ is an isometry for points in $\langle D \rangle$; it then follows from (4.5) that this holds for all points of $\mathbb{H}^*$.

The collection of $\epsilon-$neighbourhoods of points of D is an open cover of D. As D is compact, by Lebesgue's Number Lemma there exists $\delta > 0$ such that every subset of D with a diameter less than δ is contained in some set in the cover. Passing to the neighbourhoods in $\mathbb{H}^*$, this gives us an open cover $\mathcal{U}$ of $\langle D \rangle$ such that each $U \in \mathcal{U}$ has a diameter at least δ and is mapped isometrically onto $\alpha(U)$. Now it follows from Lemma 3.9 that α is a covering map.

To conclude that α is a homeomorphism, we need to know that $\mathbb{H}^*$ is path-connected. We have already seen that the neighbourhoods constructed above are path-connected, so that $\mathbb{H}^*$ is locally path-connected. As a locally path-connected space is path-connected if and only if it is connected, it suffices to prove that $\mathbb{H}^*$ is connected.

As D is connected, so are (g, D) and $\langle (g, D) \rangle$. For any $h \in \Gamma$ and any $g_s \in \Phi$ we have

$$\langle hg_s, D \rangle \cap \langle h, D \rangle \neq \varnothing,$$

since for $x \in s$, both equivalence classes contain the point $(hg_s, x) * (h, g_s x)$. It follows that for any sequence $g_1, g_2, \ldots, g_n$ of elements of ϕ with $h_1 = hg_1 g_2 \cdots g_n$, we have that $\langle h_1, D\rangle$ and $\langle h, D\rangle$ belong to the same connected component of $\mathbb{H}^*$. As any $h \in \Gamma$ is a product of elements of ϕ, it follows that $\mathbb{H}^*$ is connected. This proves that $\mathbb{H}^*$ is path-connected. $\mathbb{H}$ is simply connected, and so the conditions of Theorem 3.10 are satisfied, and we conclude that α is a homeomorphism.

4.4 Examples

The Plane Is Our Canvas

It is time to put the theorem to work. We will look at two examples in the disc model.

Example 4.7. Let D be a compact regular octagon in $\mathbb{D}$ whose interior angles are $\frac{1}{4}\pi$. Label the sides counterclockwise as $s_1, s_2, \ldots$, let g_k be the translation whose axis is the perpendicular bisector of s_k and s_{4+k} and that satisfies $g_k(s_{4+k}) = s_k$. Each g_k is a side pairing transformation for D, since $g_k(s_{4+k}) = s_k$, $g_{s_{4+k}} = g_{s_k}^{-1}$, and $g_k(D) \cap D = s_k$. Let v_k be the common vertex of s_k and s_{k+1}. We determine the cycle transformation associated to v_1, starting with s_1. s_1 is paired with s_5 by g_1^{-1}, which takes v_1 to v_4. The other side adjacent to v_4 is s_4. s_4 is paired with s_8 by g_4^{-1}, which takes v_4 to v_7, and so on. We obtain the vertex cycle

$$v_1 \xrightarrow{g_1^{-1}} v_4 \xrightarrow{g_4^{-1}} v_7 \xrightarrow{g_3} v_2 \xrightarrow{g_2^{-1}} v_5 \xrightarrow{g_1} v_8 \xrightarrow{g_4} v_3 \xrightarrow{g_3^{-1}} v_6 \xrightarrow{g_2} v_1,$$

and the cycle transformation associated to v_1 is

$$\phi = g_2 \circ g_3^{-1} \circ g_4 \circ g_1 \circ g_2^{-1} \circ g_3 \circ g_4^{-1} \circ g_1^{-1}.$$

This is the identity, so $\tau = 1$. The angle sum at v_1 is $\operatorname{sum}(v_1) = \frac{8}{4}\pi = 2\pi$. If we start with any other vertex, we will get a permutation of the transformations in the composition above, which will again be the identity. Hence the conditions of Theorem 4.6 are satisfied, and so the group generated by the side pairings is discrete and has D as a fundamental polygon.

We have centred the discussion around Fuchsian groups, but we have not assumed the side pairing transformations to be orientation-preserving, and the statements made about Fuchsian groups apply to all discrete subgroups of $\operatorname{Iso}(\mathbb{H})$. In the following example the side pairing transformations are reflections.

Example 4.8. Consider the hyperbolic triangle with interior angles $\pi/2$,$\pi/3$, $\pi/7$. Let v_1, v_2, v_3 be the vertices with interior angles $\pi/2$,$\pi/3$, and $\pi/7$, respectively. Let s_i be the side opposite v_i. Let g_i be reflection in the line containing the side s_i. Note that g_i is a side pairing transformation as $g_i(s_i) = s_i$, $g_i^{-1} = g_i$, and $g_i(D) \cap D = s_i$. We determine the cycle transformation associated to v_1, starting with the side s_2. The corresponding side pairing transformation is g_2, fixing both s_2 and v_1. The other side adjacent to v_1 is s_3. The corresponding side pairing transformation is g_3, again fixing v_1 and s_3. The other side adjacent to v_1 is s_2, and we are done. Thus the cycle transformation associated to v_1 is $\phi = g_3 \circ g_2$. The angle sum at v_1 is the sum of the interior angles at $g_2(v_1) = v_1$ and $g_3 \circ g_2(v_1) = v_2$. So $\mathrm{sum}(v_1) = \pi$. As ϕ is composed by reflections in two intersecting lines, it is a rotation. If we take the vertex to be at 0 in $\mathbb{D}$, we see that reflection in lines which intersect at angle θ is rotation through 2θ. So ϕ is rotation through π and thus $\tau = 2$. So $\tau\mathrm{sum}(v_1) = 2\pi$, as required. The construction of a cycle transformation associated to the remaining vertices is analogous, and so the conditions of Theorem 4.6 are satisfied.

The group generated in Example 4.7 is Fuchsian, as it is generated by translations. The group generated in Example 4.8 is not Fuchsian, as it contains reflections. However, the index 2 subgroup of orientation-preserving isometries is Fuchsian, and a fundamental polygon for this group is given by the union of one triangle with its reflection with respect to one of its sides. In Figure 4.4 this is the union of a black and a white triangle. Furthermore, we see that the union of the fourteen triangles meeting at the origin is a regular heptagon, so this triangular tessellation also gives a regular tessellation by heptagons.

There is a version of the Poincaré Polygon Theorem for noncompact polygons, but this imposes further conditions to ensure that the translates cover all of $\mathbb{H}$. For example, the fundamental polygon for the modular group found in example 3.29 is noncompact, but the transformations $s(z)$ and $t(z)$ defined in the example provide a side pairing for this polygon, and they are in fact generators of $\mathrm{PSL}(2, \mathbb{Z})$.

Polygons as the octagon and triangle in the examples above do not exist in euclidean or spherical geometry. In hyperbolic geometry, their existence is guaranteed by the Gauss-Bonnet Formula, which relates the sum of the interior angles of a polygon to its hyperbolic area. Example 4.8 is an example of a more general construction of tessellations by triangles. For a triangle with angles $\pi/p, \pi/q, \pi/r$, the cycle condition is satisfied as long as p, q, r are integers. When

$$\frac{1}{p} + \frac{1}{q} + \frac{1}{r} < 1, \tag{4.6}$$

there are infinitely many possibilities, and the triangle with interior angles $\pi/p, \pi/q, \pi/r$ is hyperbolic, as the angle sum is less than π (the angle sum of a euclidean triangle is equal to π, and the angle sum of a spherical triangle is greater than π). This explains why there are infinitely many hyperbolic tessellations.

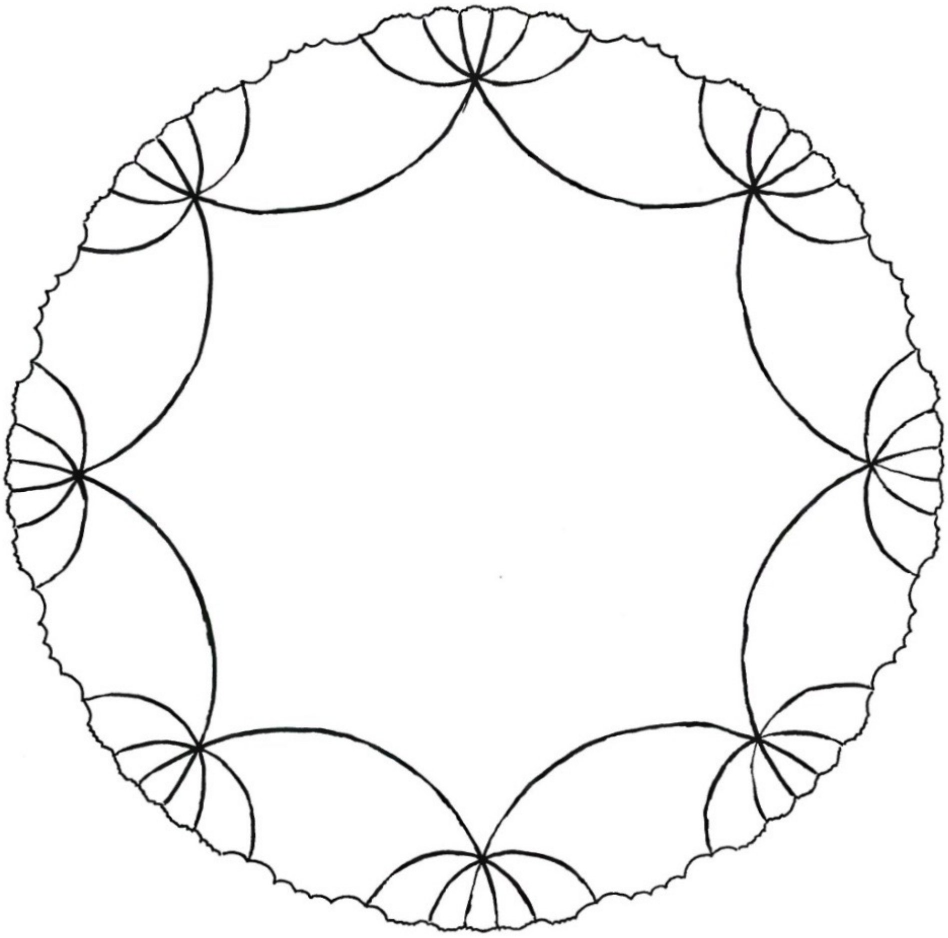

Figure 4.3. Hyperbolic tessellation by regular octagons.

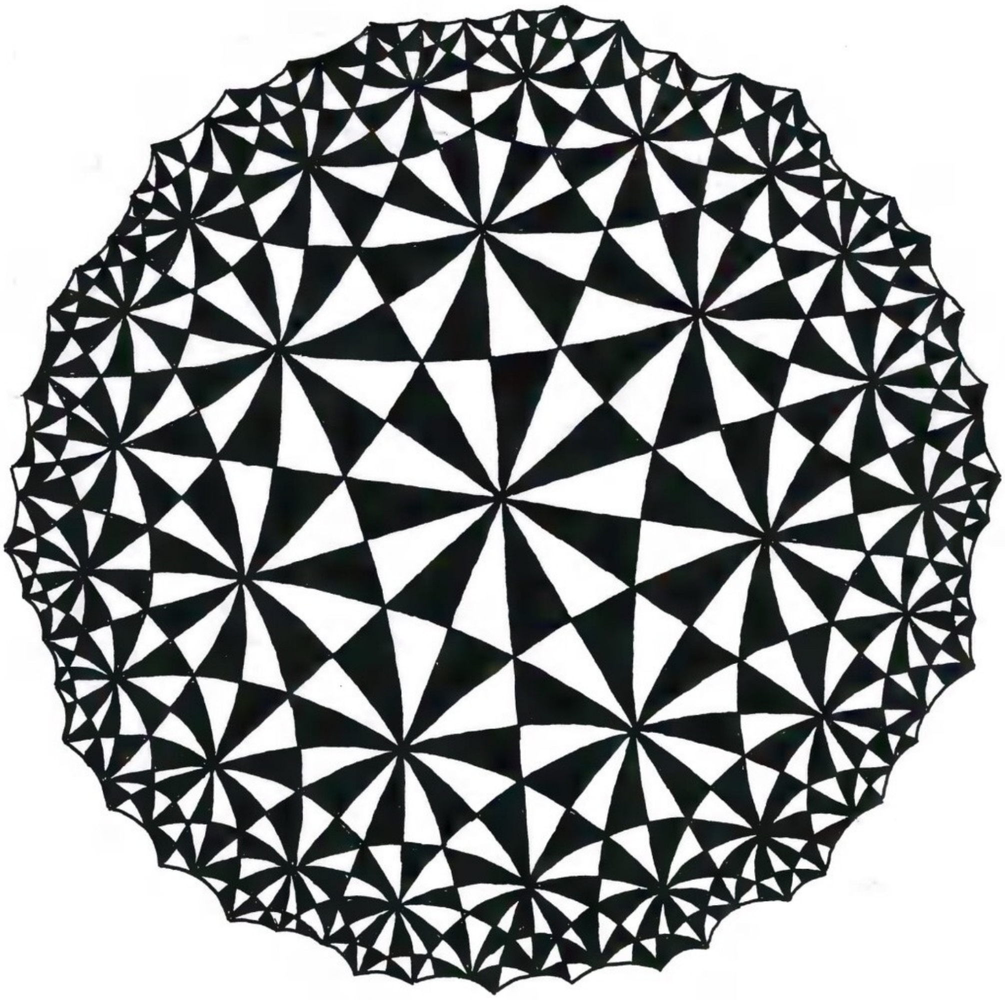

Figure 4.4. Hyperbolic tessellation by triangles.

5 The Aftermath

5.1 Conclusion

`After Maths Comes More Maths`

In our search for pretty patterns, the Poincaré Polygon Theorem emerges as a recipe for creating a tessellation of the hyperbolic plane with a given polygon. In the process, we have gained access to an endless supply of Fuchsian groups.

The construction of Fuchsian groups from polygons is tightly intertwined with the construction of hyperbolic surfaces. There is a version of the Uniformization Theorem mentioned in the introduction stating that any Riemann surface is the quotient of the euclidean plane, the sphere, or the hyperbolic plane by a discrete, fixed point free group of isometries, which in turn can be viewed as a quotient space of a fundamental polygon. As an example of this relationship, the quotient of the hyperbolic plane by the Fuchsian group generated by the octagonal tessellation in Example 4.7 is a compact hyperbolic surface of genus 2: the 'double torus'. Since the modular group and the Fuchsian group generated by the triangle in Example 4.8 contain rotations, these groups do not act freely on the hyperbolic plane, and so the quotient spaces are not hyperbolic surfaces. Instead, they are examples of a more general structure called an *orbifold*. But that is a topic for another story time.

References

[1] James W. Anderson. *Hyperbolic Geometry*. Springer-Verlag, London, 2005.

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

[3] Svetlana Katok. *Fuchsian Groups*. The University of Chicago Press, Chicago, 1992.

[4] John M. Lee. *Introduction to Topological Manifolds*. Springer Science+Business Media, New York, 2011.

[5] Bernard Maskit. "On Poincaré's Theorem for Fundamental Polygons". *Advances in Mathematics* 7 (3): 219-230. 1971.

[6] John Stillwell. *Geometry of Surfaces*. Springer Science+Business Media, New York, 1992.