Ruprecht-Karls-Universität Heidelberg
Fakultät für Mathematik und Informatik
Institut für Mathematik

On The Multiplicity of the Smallest Positive Eigenvalue of the Laplacian on the Klein Quartic

This bachelor thesis was written by Benno Wendland
under the supervision of Jr. Prof. Dr. Maria Beatrice Pozzetti.
Date of submission: August 14, 2023.

Abstract. This thesis discusses various ideas related to the proof of the statement *the multiplicity* $m_1(K)$ *of the smallest positive eigenvalue* $\lambda_1(K)$ *of the Laplacian on the hyperbolic surface called the Klein Quartic is equal to* 8. The statement was proven a few years ago in the paper [1]. First, we take a look at how the Selberg Trace Formula can be used to obtain bounds for $m_1(K)$. Second, we show the existence of a representation whose decomposition into irreducible summands implies an integer equation for the multiplicity $m_1(K)$ of $\lambda_1(K)$. Third, we investigate the conditions under which a polyhedron and its reflection group provides a tessellation of hyperbolic space. Finally, we study the macro structure of the proof of the statement $m_1(K) = 8$ and how the different strategies are combined to prove the desired result. This work contains no original results.

Zusammenfassung. In dieser Bachelorarbeit werden einige Aussagen behandelt die für den Beweis der Aussage *Die Multiplizität* $m_1(K)$ *des kleinsten positiven Eigenwerts* $\lambda_1(K)$ *des Laplace-Operators auf der hyperbolischen Fläche die den Namen Kleinsche Quartic trägt ist gleich* 8 von Relevanz sind. Dieser Satz wurde in dem Paper [1] bewiesen. Das soeben erwähnte Paper ist die primäre Quelle der Arbeit.
Wir beschäftigen uns zuerst kurz mit der Selbergschen Spurformel und beschreiben eine Strategie wie aus der Formel Abschätzungen für $m_1(K)$ gewonnen werden können. Im darauffolgenden Abschnitt beschreiben wir eine Darstellung der Isometriengruppe der Kleinschen Quartik und erläutern wie aus dieser eine Gleichung aus natürlichen Zahlen für die Zahl $m_1(K)$ zustande kommt. Daraufhin betrachten wir Kachelungen des hyperbolischen Raums welche aus einem Polytop und dessen Translaten gegeben sind. Der letzte Abschnitt enthält eine Liste der Ergebnisse welche für die Bestimmung von $m_1(K)$ benötigt werden und eine Erklärung dessen wie die partiellen Aussagen zusammengesetzt werden können um $m_1(K)$ zu berechnen.

Contents

1 Notation

- $\lambda_1(M)$ and $\mathrm{m}_1(M)$ refer to the smallest positive eigenvalue and its multiplicity of the Laplacian on a Riemannian manifold M, respectively.
- Let X be a metric space. For any $x \in X$, $r > 0$, the expression $B_r(x)$ denotes the open ball of radius r around x. Analogously, $S_r(x)$ is the distance sphere of radius r.
- For any $r > 0$ and $n \in \mathbb{N}$, $\mathbb{S}^n_r \subset \mathbb{R}^{n+1}$ denotes the Riemannian manifold that is the sphere of radius r centered around 0 endowed with the induced metric.
- The interior of a subset S of a topological space X is denoted by either $\mathrm{int}(S)$ or $\mathring{S}$.
- The symbol $\mathbb{H}^n$ denotes n dimensional hyperbolic space. The Poincare ball model of $\mathbb{H}^n$ is denoted by B^n.

2 Introduction

This thesis concerns itself primarily with the multiplicity of the smallest positive eigenvalue of the Laplacian $\triangle$ on the Klein Quartic, but some of the techniques are applicable to other closed connected Riemannian surfaces S of genus $g \geq 2$ endowed with a hyperbolic metric. The operator we want to study is defined by

$$\begin{aligned} \triangle : \mathcal{C}^2(S, \mathbb{R}) &\to \mathcal{C}^0(S, \mathbb{R}) \\ f &\mapsto -\mathrm{div}(\mathrm{grad}(f)). \end{aligned}$$

It is a fact that the set of eigenvalues of the Laplacian is countable, contained in $[0, \infty) \subset \mathbb{R}$ and does not have any accumulation points, so that the eigenvalues can be sorted in increasing order

$$0 \leq \lambda_0(S) \leq \lambda_1(S) \leq \dots .$$ [1]

In the paper [1], the authors show that the multiplicity $m_1(K)$ of the smallest positive eigenvalue on the Klein Quartic K is 8. Furthermore, they show that this is the maximal value for $m_1(S)$ amongst all closed hyperbolic surfaces S of genus $g = 3$ [ibid Theorem 1.2]. This thesis describes some ideas which are more or less directly relevant to the determination of $m_1(K)$.

One way of describing the Klein Quartic is as the quotient of the hyperbolic plane by a subgroup of the reflection group of the hyperbolic triangle with interior angles $\frac{\pi}{2}, \frac{\pi}{3}, \frac{\pi}{7}$. This description leads to a fundamental domain of the group's action on the hyperbolic plane which e.g. can be useful for the construction of some closed geodesics. One section of this thesis deals with the proof of the statement that certain polyhedra in hyperbolic space $\mathbb{H}^n$ are fundamental domains for their induced reflection groups. In other words, in certain circumstances, the translates of a polyhedron by the isometries in the induced reflection group form a tiling of the space $\mathbb{H}^n$.

The other sections of this text deal with two different approaches to the goal of obtaining upper and lower bounds on the multiplicity $m(\lambda)$ of eigenvalues λ of the Laplacian on a compact connected hyperbolic surface S. First, we discuss the Selberg trace formula. The formula comes up in this context because it relates a series of values which are summed over the multiset $\sigma(S)$ of eigenvalues of the Laplacian to a different quantity. This other quantity involves summation over the closed geodesics in the surface S. Therefore, we spend some time proving the fact that there are only countably many closed geodesics in such a surface - this statement obviously requires the exclusion of constant paths as well as an identification of geodesics which differ only by a reparameterization. At this point, it is worth emphasizing the distinction between geodesics $\gamma : [a, b] \to S$ which form a loop in the sense that $\gamma(a) = \gamma(b)$ and geodesics which are closed in the sense that $\gamma(a) = \gamma(b)$ and $\dot{\gamma}(a) = \dot{\gamma}(b)$. The argument for the countability of the set of closed geodesics proceeds by establishing a correspondence between the conjugacy classes of the group of deck transformations and certain equivalence classes of closed geodesics. It works in the general case of a compact, connected Riemannian manifold with strictly negative sectional curvature. The compactness and connectedness assumption implies that any deck transformation leaves the image of some geodesic invariant. The curvature assumption guarantees the uniqueness (up to reparameterisation) of such a geodesic for any fixed deck transformation.

The second approach involved in determining the multiplicity $m_1(S)$ of the smallest positive eigenvalue has its starting point in the observation that the isometry group $\mathrm{Iso}(S)$ acts on the eigenspace E_λ: any isometry h induces a linear map

$$\begin{aligned} L_h : E_\lambda &\to E_\lambda \\ f &\mapsto f \circ h^{-1}. \end{aligned}$$

[1] See, for example, [3, chapter 7].

Using the basic result from the representation theory of finite groups which states that any representation decomposes into a sum of irreducible representations, one quickly arrives at the equation

$$m_1(S) = n_1 d_1 + n_2 d_2 + \cdots + n_k d_k,$$

where k is the cardinality of the collection of irreducible representations of the finite group, $d_1, \ldots, d_k$ are the dimensions of these irreducible representations and $n_1, \ldots, n_k \in \mathbb{N}_0$ are some natural numbers. It is at first glance an entirely useless equation, since the trivial representation is an irreducible representation of dimension 1. However, in the case of the specific representation we are concerned with, the one-dimensional irreducible representations do not occur in the representation's direct sum decomposition [1, Corollary 5.2.].

3 Bounds on the Multiplicities of Eigenvalues of the Laplacian Derived From the Selberg Trace Formula

We want to count - with multiplicities - the eigenvalues of the Laplacian which lie in a certain interval $[a,b] \subset \mathbb{R}$. Let $\sigma(M)$ denote the multiset whose elements are the eigenvalues of the Laplacian on M, each occurring according to its multiplicity, i.e. the dimension of the associated eigenspace. The approach for finding an upper and lower bound is based on expressions of the form

$$\sum_{\lambda \in \sigma(M)} f(\lambda)$$

for certain special functions f. For example, suppose that f is bounded from below on the interval $[a,b]$ by a constant c and non-negative on $\mathbb{R} \setminus [a,b]$. Then we have the inequality

$$c \cdot \#\sigma(M) \cap [a,b] = \sum_{\lambda \in \sigma(M) \cap [a,b]} c \leq \sum_{\lambda \in \sigma(M) \cap [a,b]} c + \sum_{\lambda \in \sigma(M) \cap (\mathbb{R} \setminus [a,b])} f(\lambda) \leq \sum_{\lambda \in \sigma(M)} f(\lambda).$$

Analogously, if we assume f to be bounded from above by a constant c on the interval $[a,b]$ and be non-positive on $\mathbb{R} \setminus [a,b]$, we obtain a lower bound for the number of eigenvalues counting multiplicity in the interval $[a,b]$:

$$c \cdot \#\sigma(M) \cap [a,b] = \sum_{\lambda \in \sigma(M) \cap [a,b]} c \geq \sum_{\lambda \in \sigma(M) \cap [a,b]} c + \sum_{\lambda \in \sigma(M) \cap (\mathbb{R} \setminus [a,b])} f(\lambda) \geq \sum_{\lambda \in \sigma(M)} f(\lambda).$$

The difficulty lies in finding a useful way to bound a sum of the form

$$\sum_{\lambda \in \sigma(M)} f(\lambda),$$

as well as functions which satisfy the conditions under which such a bound holds. A useful expression for the sum $\sum_{\lambda \in \sigma(M)} f(\lambda)$ is given by the Selberg trace formula. One method for proving the existence of appropriate functions uses interval arithmetic, which we will not discuss here.

The Selberg trace formula involves a series where the summation index is the set of closed geodesics. Since attempting to add uncountably many (non-zero) real numbers is not a sensible idea, the trace formula can only be plausible if there are at most countably many closed geodesics.

3.1 Deck Transformations and Closed Geodesics

In this section, (M,g) is a compact connected Riemannian manifold with negative sectional curvature and $(\tilde{M}, \tilde{g}, \pi)$ is its universal Riemannian cover. Note that $\tilde{M}$ is a simply connected, complete Riemannian manifold with non-positive sectional curvature. Such manifolds are sometimes referred to as Cartan-Hadamard manifolds.

We show that the set of closed geodesics in M is countable by relating the closed geodesics to deck transformations of the universal cover $\tilde{M}$. As mentioned in the introduction, this requires the exclusion of the constant geodesics and an identification of closed geodesics which are reparameterizations of each other. The main ingredient is the statement, which holds given the above assumptions, that any non-trivial deck transformation of the Riemannian universal cover leaves exactly one geodesic invariant [6, Lemmas 12.21 and 12.22]. We use the result without proof.

Definition 3.1 (closed geodesics)**.** A *closed geodesic* is a geodesic $\gamma : [a,b] \to M$ with the property $\gamma(a) = \gamma(b)$ and $\dot{\gamma}(a) = \dot{\gamma}(b)$. It is sometimes convenient to interpret closed geodesics as smooth maps $\mathbb{S}^1 \to M$. This, for one thing, has the consequence that reparamerisations to different constant speeds are not possible, which would be the case if we allowed intervals in $\mathbb{R}$. Furthermore, the description of a change of the point at which the closed geodesic is based is simplified.

Define $\mathcal{C}(M)$ as the set of equivalence classes of oriented non-trivial closed geodesics of M up to changes of the base point. More precisely:

$$\mathcal{C}(M) := \{\, \gamma : S^1 \to M \mid \gamma \text{ is not a constant map and } \gamma \text{ is a closed geodesic} \,\} / \sim$$

where the equivalence relation $\sim$ is given by

$$\gamma \sim \tau \iff \exists c \in S^1 \forall t \in S^1 : \gamma(t) = \tau(t + c).$$

Definition 3.2. A deck transformation $\Phi : \tilde{M} \to \tilde{M}$ is a continuous map which satisfies $\pi = \pi \circ \Phi$. The group of deck transformations is denoted by $A(\tilde{M})$. With $\sim$ denoting the equivalence relation where two maps are equivalent if they are conjugate, the set of conjugacy classes is denoted by $A(\tilde{M})/\sim$.

An axis for a deck transformation Φ is a geodesic $\gamma : \mathbb{R} \to \tilde{M}$ for which there exists a constant $c \in \mathbb{R} \setminus \{\, 0 \,\}$ such that $\Phi \circ \gamma(t) = \gamma(t + c)$ for all $t \in \mathbb{R}$.

Lemma 3.3. *Let $\Phi \in A(\tilde{M})$ be a non-trivial deck transformation. Then Φ has an axis γ which is unique up to reparameterisation. These reparameterisations are of the form $t \mapsto \pm ct + a$ for some $a \in \mathbb{R}$ and $c \in \mathbb{R} \setminus \{\, 0 \,\}$. If γ is an axis of Φ, then the constant*

$$c \in \mathbb{R} \setminus \{\, 0 \,\} \text{ with } \gamma(t + c) = \Phi(\gamma(t)) \text{ for all } t \in \mathbb{R} \tag{1}$$

depends on the parameterisation of γ, but is unique once a parameterisation is fixed. The axes of the trivial deck transformation, that is the identity on $\tilde{M}$, are exactly the constant paths.

Proof. The first statement is the content of [6, Lemmas 12.21 and 12.22].

The latter two statements are immediate consequences of the fact that the non-constant geodesics of Cartan-Hadamard manifolds do not self-intersect. This holds, because the injectivity radius is infinite [6, Prop. 12.9]:

Let γ be a non-constant geodesic which is an axis for Φ and suppose $c_1, c_2 \in \mathbb{R} \setminus \{\, 0 \,\}$ satisfy $\gamma(t + c_1) = \Phi(\gamma(t)) = \gamma(t + c_2)$. Then the fact that γ does not have any self-intersections implies $c_1 = c_2$.

Any axis γ of the identity map must satisfy $\gamma(t + c) = \gamma(t)$ for a non-zero constant $c \in \mathbb{R}$, whence γ has a self-intersection and must be a constant geodesic. □

Let γ be an axis for Φ. The previous theorem guarantees the existence and uniqueness (up to reparameterisation) of such geodesics. We will subsequently denote the axis of Φ by γ_Φ. It is worth noting that the constant c in the above lemma distinguishes between different deck transformations which have the same axis. More precisely, if γ is an axis for two deck transformations Φ and Ψ, then the constants c_Φ, c_Ψ for which Equation (1) holds are the same if and only if Φ is equal Ψ. This follows from Part (1) of the next lemma, which consists of a few properties concerning deck transformations and closed geodesics which we will need in order to show that there is a correspondence between the two.

Lemma 3.4.

1. *A deck transformation Φ of $\tilde{M}$ is uniquely determined by its value at any one point, i.e.*

$$\Big(\Phi, \Psi \in A(\tilde{M}) \text{ and } \exists p \in \tilde{M} : \Phi(p) = \Psi(p)\Big) \implies \Phi = \Psi.$$

2. *Let $\gamma_\Phi : \mathbb{R} \to \tilde{M}$ be an (up to reparameterisation unique) axis of $\Phi \in A(\tilde{M}) \setminus \{\, \mathrm{Id} \,\}$. Let $c \in \mathbb{R} \setminus \{\, 0 \,\}$ be the unique constant by which Φ translates γ_Φ. If $c < 0$, then we switch to the reparameterization of γ_Φ given by $t \mapsto \gamma_\Phi(-t)$ to ensure that $c > 0$. In particular, we have $\Phi(\gamma_\Phi(0)) = \gamma_\Phi(c)$. Define*

$$\tau_\Phi : [0, c] \to M, t \mapsto \pi \circ \gamma_\Phi(t).$$

Then τ_Φ is a closed geodesic which is uniquely determined up to constant speed orientation preserving reparameterisations and a different choice of starting point.

3. *The associated closed geodesics τ_Φ, τ_Ψ are the same if and only if the deck transformations are conjugate, i.e.*
$$\tau_\Phi = \tau_\Psi \iff \exists \eta \in A(\tilde{M}) : \Phi = \eta \circ \Psi \circ \eta^{-1}.$$
4. *For any closed geodesic $\gamma : [0, b] \to M \in \mathcal{C}(M)$ there exists at least one deck transformation $\Phi \in A(\tilde{M})$ such that $\gamma = \tau_\Phi$. One such deck transformation is obtained by choosing a lift $\tilde{\gamma}$ of γ and defining Φ via $\Phi(\tilde{\gamma}(0)) = \tilde{\gamma}(b)$.*

Proof. Part 1: Let $p \in \tilde{M}$ and $\Phi \in A(\tilde{M})$. Let U be an evenly covered neighborhood of $\pi(p)$ in M, and denote the components of $\pi^{-1}(\{\pi(p)\})$ which contain p and $\Phi(p)$ by V_p and $V_{\Phi(p)}$, respectively. By continuity of Φ, there exists a neighborhood $W \subset V_p$ of p which is mapped into $V_{\Phi(p)}$. Since Φ is a deck transformation, it follows that
$$\Phi\big|_W = \left(\pi\big|_{V_{\Phi(p)}}\right)^{-1} \circ \pi\big|_W. \tag{2}$$

Suppose Ψ is a deck transformation which satisfies $\Psi(p) = \Phi(p)$. Then the above equation shows that the two maps agree on a neighborhood of p. Since p was arbitrary, the set of points in $\tilde{M}$ on which Ψ agrees with Φ is open. It is also closed, since both maps are continuous. Due to the connectedness of $\tilde{M}$, this suffices to conclude that $\Phi = \Psi$. Note that the above formula also implies that deck transformations are isometries, because π is a local isometry.

Part 2: Let γ_Φ be a parameterisation of the axis of Φ such that the translation constant c is positive. Then any version of τ_Φ is given by
$$t \mapsto \pi \circ \gamma_\Phi(at + b)$$
for some constants $a > 0$ and $b \in \mathbb{R}$, because those are the only admissible reparameterisations of γ_Φ. The translation constant for $t \mapsto \gamma_\Phi(at + b)$ is given by $\frac{c}{a} > 0$, so that the domain of this version of τ_Φ is given by $[0, \frac{c}{a}]$.

The fact that τ_Φ is a closed geodesic follows directly from the local representation of Φ stated in Equation (2).

Part 3: First, we show that conjugate deck transformations induce the same closed geodesic. For this purpose, let $\Phi, \eta \in A(\tilde{M})$ and $c > 0$ with $\gamma_\Phi(t + c) = \Phi \circ \gamma_\Phi(t)$ for all $t \in \mathbb{R}$. Then $\eta \circ \gamma_\Phi$ is an axis of $\Psi := \eta \circ \Phi \circ \eta^{-1}$ whose translation constant is also given by c. To see this, note that, for all $t \in \mathbb{R}$,
$$\Psi(\eta \circ \gamma_\Phi(t)) = \eta \circ \Phi \circ \left(\eta^{-1} \circ \eta\right) \circ \gamma_\Phi(t) = \eta \circ \gamma_\Phi(t + c).$$
Consequently, the curves $\pi \circ \gamma_\Phi$ and $\pi \circ \gamma_\Psi$ restricted to the domain $[0, c]$ are specific representatives of the equivalence classes τ_Φ and τ_Ψ, respectively. For these parameterisations, we therefore have
$$\tau_\Phi = \pi \circ \gamma_\Phi\big|_{[0,c]} = (\pi \circ \eta) \circ \gamma_\Phi\big|_{[0,c]} = \pi \circ (\eta \circ \gamma_\Phi)\big|_{[0,c]} = \pi \circ \gamma_\Psi\big|_{[0,c]} = \tau_\Psi.$$

Suppose τ_Φ is equal to τ_Ψ. We want to show that Ψ agrees with a conjugate of Φ at the point $\gamma_\Psi(0)$ and apply Part (1) to deduce that Ψ is a conjugate of Φ.

Our assumption means, in particular, that domains of τ_Φ and τ_Ψ are the the same. The domain is the interval $[0, c]$ for some $c > 0$. By construction, we can choose parameterisations of the axes γ_Φ and γ_Ψ such that γ_Φ restricted to $[0, c]$ is a lift of τ_Φ and γ_Ψ restricted to $[0, c]$ is a lift of τ_Ψ. Furthermore, it holds that $\Phi(\gamma_\Phi(0)) = \gamma_\Phi(c)$ and $\Psi(\gamma_\Psi(0)) = \gamma_\Psi(c)$. Let η be the deck transformation which maps $\gamma_\Phi(0)$ onto $\gamma_\Psi(0)$. We know that such a deck transformation exists, because $\tilde{M}$ is simply connected. See for example [2, p. 148]. We conclude that γ_Ψ and $\eta \circ \gamma_\Phi$ are two lifts of the same path which start at the same point, whence they are equal. In particular, $\eta^{-1}(\gamma_\Psi(0)) = \gamma_\Phi(0)$ and $\gamma_\Psi(c) = \eta(\gamma_\Phi(c))$. Combining these observations, we obtain
$$\eta \circ \Phi \circ \eta^{-1}(\gamma_\Psi(0)) = \eta \circ \Phi(\gamma_\Phi(0)) = \eta(\gamma_\Phi(c)) = \gamma_\Psi(c).$$
We conclude that $\eta \circ \Phi \circ \eta^{-1}$ and Ψ are the same deck transformation.

Part 4: Let $\gamma : [0, b] \to M$ be a closed geodesic with a lift $\tilde{\gamma}$. Let Φ be the unique deck transformation with $\Phi(\tilde{\gamma}(0)) = \tilde{\gamma}(b)$. We claim that the extension of $\tilde{\gamma}$ to $\mathbb{R}$ is an axis for Φ. More specifically, it holds that $\tilde{\gamma}(kb + t) = \Phi \circ \tilde{\gamma}((k-1)b + t)$ for all $k \in \mathbb{Z}$ and $t \in [0, b]$. This can be shown by induction. We have, for example,

$$\Phi^{-1} \circ \tilde{\gamma}(b + t) = \tilde{\gamma}(t)$$

for $t \in [0, b]$, which holds as the curves on the left and right side of the equality are geodesics with the same starting point and the same initial velocity. The latter statement about the velocities is a consequence of the relationship between the derivative of $\tilde{\gamma}$ when approaching 0 from above and the derivative of $\tilde{\gamma}$ when approaching b from below, which stems from the fact that $\gamma'(0) = \gamma'(b)$. More precisely, one can compute

$$\begin{aligned}\left(\Phi^{-1} \circ \tilde{\gamma}\right)'(b) &= d_{\tilde{\gamma}(b)}\Phi^{-1}(\tilde{\gamma}'(b)) = d\left(\pi_{|V_0}\right)^{-1}\left(d\pi_{|V_b}(\tilde{\gamma}'(b))\right) = d\left(\pi_{|V_0}\right)^{-1}\left((\pi_{|V_b} \circ \tilde{\gamma})'(b)\right) \\ &= d\left(\pi_{|V_0}\right)^{-1}(\gamma'(b)) = d\left(\pi_{|V_0}\right)^{-1}(\gamma'(0)) = d\left(\pi_{|V_0}\right)^{-1}\left((\pi_{|V_0} \circ \tilde{\gamma})'(0)\right) = \tilde{\gamma}'(0)\end{aligned}$$

for appropriate neighborhoods V_b of $\tilde{\gamma}(b)$ and V_0 of $\tilde{\gamma}(0)$. $\square$

Lemma 3.5. *There exists a bijection between the non-trivial conjugacy classes of the group of deck transformations of $\tilde{M}$ and the set of equivalence classes of non-constant closed geodesics in M, i.e.*

$$\left(A(\tilde{M})/\sim\right) \setminus \{\,[\mathrm{Id}_{\tilde{M}}]\,\} \overset{set}{\cong} \mathcal{C}(M).$$

Since the group of deck transformations is countable, this implies

$$\#\mathcal{C}(M) = \#A(\tilde{M})/\sim = \#\mathbb{N}.$$

Proof. We claim that the two maps

$$\begin{aligned}\mathcal{C}(M) &\to \left(A(\tilde{M})/\sim\right) \setminus \{\,[\mathrm{Id}_{\tilde{M}}]\,\} \\ \gamma &\mapsto [\Phi] \text{ where } \Phi \text{ is any deck transformation with } \tau_\Phi = \gamma\end{aligned}$$

and

$$\begin{aligned}\left(A(\tilde{M})/\sim\right) \setminus \{\,[\mathrm{Id}_{\tilde{M}}]\,\} &\to \mathcal{C}(M) \\ [\Phi] &\mapsto \tau_\Phi.\end{aligned}$$

are well defined and each others inverses.

To show that the first map is well defined, we need to prove that there exists a non-trivial deck transformation associated to any closed geodesic in $\mathcal{C}(M)$ and that any two such deck transformations are conjugate. The existence is part (4) of lemma (3.4), the uniqueness up to conjugacy is part (3) of the same lemma.

It remains to be seen that the deck transformation specified in Part (4) is not the trivial one. Given the formula for Φ specified in the lemma, it is clear that the associated deck transformation is trivial if and only if the closed geodesic is null-homotopic. Therefore, we must show that $\mathcal{C}(M)$ does not contain null-homotopic closed geodesics. This is a consequence of $\tilde{M}$ being a Cartan-Hadamard manifold. We claim that the only null-homotopic closed geodesics in M are the constant paths themselves; these are not contained in the set $\mathcal{C}(M)$. To see this, note that a null-homotopic closed geodesic γ in M lifts to a geodesic $\tilde{\gamma}$ in $\tilde{M}$ whose start and end points are the same. The only self-intersecting geodesics in a Cartan-Hadamard manifold are the constant paths, whence $\tilde{\gamma}$ must be constant. Consequently, γ is constant, and we conclude that all null-homotopic closed geodesics in M are constant paths.

To show that the second map is well defined, we need to note three things. First, the closed geodesic τ_Φ induced by a non-trivial deck transformation Φ is uniquely defined up to orientation preserving reparameterisation. This is the reason why we exclude the conjugacy class of the identity: it does not have a unique axis, since all constant paths qualify for the position. Second, τ_Φ is not constant because it lifts to a path between the distinct points p and $\Phi(p)$, where p is any element of the fiber $\pi^{-1}(\{ \tau_\Phi(0) \})$. Here we use the fact that a deck transformation which is not the identity map on $\tilde{M}$ has no fixed points, a consequence of the previous lemma's part (1). Third, the conjugate deck transformations induce the same closed geodesic by part (3) of the aforementioned lemma.

The fact that the maps are inverse to each other is obvious.

The cardinality of the group of deck transformations can be bounded using part (1). Let $p \in \tilde{M}$. Since any deck transformations is determined by the image of p, the evaluation map

$$\begin{aligned} A(\tilde{M}) &\to \pi^{-1}(\{ \pi(p) \}) \\ \Phi &\mapsto \Phi(p) \end{aligned}$$

is an injection. The codomain is a countable set, since the preimage of an evenly covered neighborhood U of $\pi(p)$ under π is the union of a collection of disjoint open sets indexed by $\pi^{-1}(\{ \pi(p) \})$. Since the universal cover is a manifold, and as such is a second countable topological space, it follows that the set $\pi^{-1}(\{ \pi(p) \})$ is countable. It is worth mentioning that there exists an isomorphism between the group of deck transformations and the fundamental group, so the above cardinality statement simply amounts to the statement that the fundamental group is countable, which is true for topological manifolds in general. We implicitly used this fact when we said that the universal cover is second countable. □

3.2 The Selberg Trace Formula and Bounds on The Multiplicity of Eigenvalues of the Laplacian

In this subsection, M is a compact Riemannian manifold of dimension 2, genus $g \geq 2$ and constant negative sectional curvature equal -1.

Definition 3.6 (Two notions of length for closed geodesics)**.** For any geodesic γ, the expression γ^n, is inductively defined to be the concatenation of the curve γ^{n-1} with $\gamma^1 := \gamma$. For any non-trivial closed geodesic γ, the set

$$\{ n \in \mathbb{N} \mid \exists : \eta \text{ closed geodesic } : \gamma = \eta^n \}$$

is nonempty, as it contains 1. We define $l(\gamma)$ as the length of γ. Let m be the maximum of

$$\{ n \in \mathbb{N} \mid \exists : \eta \text{ closed geodesic} : \gamma = \eta^m \}.$$[2]

Define $\Lambda(\gamma)$ as the length of η, where η is uniquely determined by the condition $\eta^m = \gamma$. If $m = 1$, then γ is called a primitive closed geodesic.

[2]From the discussion in the subsection on closed geodesics, we know that there are no null-homotopic geodesics in M other than the constant paths. Therefore, given a non-constant closed geodesic $\gamma : [0, b] \to M$, there exists a lower bound $\epsilon > 0$ for the set $\{ t \in [0, b] \mid \gamma(0) = \gamma(t) \}$ of times t at which γ returns to its starting point. It follows that any geodesic η with $\gamma = \eta^n$ satisfies $l(\eta^m) = ml(\eta) > m\epsilon > l(\gamma)$ for sufficiently large $m \in \mathbb{N}$.

We obtain a lower bound as follows: Let $p \in \tilde{M}$ such that $\pi(p) = \gamma(0)$ and let $\tilde{\gamma}$ be the lift of γ starting at p. We know that $\tilde{\gamma}$ is not a constant geodesic, since γ is not null-homotopic, which implies that the endpoints af $\tilde{\gamma}$ are distinct. Because $\tilde{M}$ is a Cartan-Hadamard manifold, $\tilde{\gamma}$ has no self-intersections. Let U be an evenly covered neighborhood of p and V_p the component of $\pi^{-1}(U)$ which contains p. Then there exists an $\epsilon > 0$ such that $B_\epsilon(p) \subset V_p$. Since π is injective on V_p, it follows that γ does not self intersect on the interval $[0, \frac{\epsilon}{c}]$, where c is the constant speed of γ.

As a sidenote, a version of the argument above should also be usable to prove that the Lebesgue number l of a cover of M by evenly covered balls $\left((B_{\epsilon(p)}(p)\right)_{p \in M}$ should be a lower bound for the lengths of non-constant closed geodesics in M.

We need to introduce some additional definitions in order to formulate the Selberg trace formula. Define a map

$$\tau : [0,\infty) \to \mathbb{R} \cup i[0,\frac{1}{2}]$$
$$\lambda \mapsto \begin{cases} i\sqrt{\frac{1}{4}-\lambda} & \text{if } 0 \leq \lambda < \frac{1}{4} \\ \sqrt{\lambda - \frac{1}{4}} & \text{else} \end{cases}.$$

Define the Fourier transform $\hat{f} : D \to \mathbb{C}$ of an integrable map $f : \mathbb{R} \to \mathbb{C}$ by the expression

$$\hat{f}(\xi) := \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} f(x) e^{-i\xi x} dx,$$

where D is the subset of $\mathbb{C}$ for which the integral is defined.

Theorem 3.7 (Selberg Trace Formula)**.** *Let $f : \mathbb{R} \to \mathbb{C}$ be a function with the following properties:*

1. *f is integrable*
2. *f is even*
3. *there exists an $\epsilon > 0$ such that Fourier transform $\hat{f}$ of f defines a holomorphic map on the set $\{\, \xi \in \mathbb{C} \mid |\mathrm{Im}\xi| < \frac{1}{2} + \epsilon \,\}$ and $\hat{f} = \mathcal{O}((1+|\xi|^2)^{-(1+\epsilon)})$.*

Then the equation

$$\sum_{\lambda \in \sigma(M)} \hat{f}(\tau(\lambda)) = 2(g-1) \int_0^{\infty} r \hat{f}(r) \tanh(\pi r) dr + \frac{1}{\sqrt{2\pi}} \sum_{\gamma \in \mathcal{C}(M)} \frac{\Lambda(\gamma)}{2 \sinh\left(\frac{l(\gamma)}{2}\right)} f(l(\gamma))$$

holds. The series on both sides of the equality converge absolutely.

The formulation of the theorem was taken from [[1]]. For the proof, they refer to [[3]].

What is the idea behind the Selberg trace formula? One uses f to define an integral kernel on M and then evaluates the kernel's trace in two different ways. One computation of the integral is based on an expansion of the kernel by eigenfunctions of the Laplacian. This is how the spectrum of the Laplacian gets involved.

From the Selberg trace formula we can obtain a formula for both a lower and an upper bound on the number of eigenvalues counting multiplicities which lie in certain intervals.

Lemma 3.8. *[Upper bound formula [1, Lemma 3.2]] Assume f satisfies the conditions required for the Selberg trace formula to hold. Assume $\hat{f} \circ \tau$ is bounded from below on the interval $[a,b]$ by a constant c and non-negative on (b,∞). Here, we additionally assume that $0 < a \leq \lambda_1(M)$ and $a < b$. Finally, assume $f(l(\gamma))$ is non-positive for all $\gamma \in \mathcal{C}(M)$. Let N be an arbitrary subset of $\mathcal{C}(M)$.*

Then we have the upper bound

$$\#\sigma(M) \cap [a,b] \leq \frac{1}{c} \left(2(g-1) \int_0^{\infty} r \hat{f}(r) \tanh(\pi r) dr + \frac{1}{\sqrt{2\pi}} \sum_{\gamma \in N} \frac{\Lambda(\gamma)}{2 \sinh\left(\frac{l(\gamma)}{2}\right)} f(l(\gamma)) - \hat{f}(\tau(0)) \right).$$

Proof. Given our assumptions, the argument from the beginning of this chapter is almost applicable. We just need to deal with the eigenvalue 0, since this eigenvalue is (the only one which is) not in the set $\mathbb{R} \setminus [a,b]$. The modified inequalities are

$$\begin{aligned} c \cdot \#\sigma(M) \cap [a,b] + \hat{f}(\tau(0)) &= \sum_{\lambda \in \sigma(M) \cap [a,b]} c + \hat{f}(\tau(0)) \\ \leq \sum_{\lambda \in \sigma(M) \cap [a,b]} c + \sum_{\lambda \in \sigma(M) \cap (\mathbb{R} \setminus (\{\,0\,\} \cup [a,b]))} \hat{f}(\tau(\lambda)) + \hat{f}(\tau(0)) &\leq \sum_{\lambda \in \sigma(M)} \hat{f}(\lambda), \end{aligned}$$

which gives us the inequality

$$\#\sigma(M) \cap [a,b] \leq \frac{1}{c}\left(\sum_{\lambda\in\sigma(M)} \hat{f}(\lambda) - \hat{f}(\tau(0))\right).$$

The Selberg trace formula immediately delivers the inequality

$$\sum_{\lambda\in\sigma(M)} \hat{f}(\tau(\lambda)) \leq 2(g-1)\int_0^\infty r\hat{f}(r)\tanh(\pi r)dr + \frac{1}{\sqrt{2\pi}}\sum_{\gamma\in N}\frac{\Lambda(\gamma)}{2\sinh\left(\frac{l(\gamma)}{2}\right)}f(l(\gamma)),$$

since the fact that $\hat{f}(l(\gamma))$ is non-positive implies that we obtain an upper bound by leaving out some terms from the sum

$$\frac{1}{\sqrt{2\pi}}\sum_{\gamma\in\mathcal{C}(M)}\frac{\Lambda(\gamma)}{2\sinh\left(\frac{l(\gamma)}{2}\right)}f(l(\gamma)).$$

By combining the two inequalities, we arrive at the desired upper bound. $\square$

There is an analogous formula which provides a lower bound for the number of eigenvalues in a compact interval counting mulitplicities, see [1, Lemma 5.9].

4 An Approach Based on a Representation of the Surface's Isometry Group

All representations in this section are assumed to be finite dimensional.

Associated to each eigenvalue λ of the Laplacian on a Riemannian manifold (M, g), there is a representation of the isometry group $\mathrm{Iso}(M)$ on the eigenspace E_λ. This has to do with the fact that the composition of an eigenvector of the Laplacian with an isometry of M is again an eigenvector of the Laplacian. Our first goal is to prove said statement.

Lemma 4.1. *Let A be a covariant k-tensor field, $h : M \to M$ be a diffeomorphism, and $V \in \mathcal{X}(M)$ be a vector field. Then*

$$\mathcal{L}_{h_*^{-1}V} h^* A = h^* \mathcal{L}_V A.$$

Remark. This is in a sense analogous to the formula

$$h_*[X, Y] = [h_* X, h_* Y]$$

for the Lie derivative of vector fields, where $h : M \to M$ is a diffeomorphism and $X, Y \in \mathcal{X}(M)$ are vector fields. We prove the lemma by reducing it to this statement.

Proof. We verify the equality using the formula

$$\big(\mathcal{L}_V A\big)(X_1, \ldots, X_k) = V\big(A(X_1, \ldots, X_k)\big) + \sum_i A(X_1, \ldots, X_{i-1}, [V, X_i], X_{i+1}, \ldots, X_k),$$

which holds for any covariant k-tensor field A and vector fields $X_1, \ldots, X_k$ on M.[3] Showing that two covariant k-tensor fields are the same is equivalent to checking that the pairing of the tensor fields with any ordered tuple of vector fields $X_1, \ldots, X_k$ produces the same real valued function.

Note that we have

$$\begin{aligned}\big(h^* \mathcal{L}_V A\big)(X_1, \ldots, X_k)(p) &= \big(\mathcal{L}_V A\big)_{h(p)}(d_p h(X_1(p)), \ldots, d_p h(X_k(p))) \\ &= \big(\mathcal{L}_V A\big)_{h(p)}(h_* X_1(h(p)), \ldots, h_* X_k(h(p))),\end{aligned}$$

for all points $p \in M$, which globally can be written as an equality of real valued functions on M:

$$\big(h^* \mathcal{L}_V A\big)(X_1, \ldots, X_k) = \big(\mathcal{L}_V A\big)(h_* X_1, \ldots, h_* X_k) \circ h.$$

This is a useful description, because we now have a pairing of the Lie derivative of a covariant tensor field with vector fields on the right hand side, so that we can use the formula for the Lie derivative stated above. Simply inserting into said formula gives us the two expressions

$$\begin{aligned}\big(h^* \mathcal{L}_V A\big)(X_1, \ldots, X_k) =& \big(\mathcal{L}_V A\big)(h_* X_1, \ldots, h_* X_k) \circ h \\ =& V\big(A(h_* X_1, \ldots, h_* X_k)\big) \circ h \\ &+ \sum_i A(h_* X_1, \ldots, h_* X_{i-1}, [V, h_* X_i], h_* X_{i+1}, \ldots, h_* X_k) \circ h,\end{aligned}$$

and

$$\begin{aligned}\big(\mathcal{L}_{h_*^{-1}V} h^* A\big)(X_1, \ldots, X_k) =& h_*^{-1} V\big((h^* A)(X_1, \ldots, X_k)\big) \\ &+ \sum_i (h^* A)(X_1, \ldots, X_{i-1}, [h_*^{-1} V, X_i], X_{i+1}, \ldots, X_k).\end{aligned}$$

[3] See, for example, [7, Corollary 12.33].

We claim that the i-th summand in the formula for $\big(h^*\mathcal{L}_V A\big)(X_1,\dots,X_k)$ agrees with the i-th summand in the formula for $\big(\mathcal{L}_{h_*^{-1}V} h^*A\big)(X_1,\dots,X_k)$. First, we consider the first summand evaluated at a point $p \in M$. The computation is mostly an exercise in inserting the definition of the push forward of a vector field:

$$
\begin{aligned}
&V\big(A(h_*X_1,\dots,h_*X_k)\big)\circ h(p) && = V_{h(p)}\big(A(h_*X_1,\dots,h_*X_k)\big)\\
=&V_{h(p)}\big(A(dh(X_1)\circ h^{-1},\dots,dh(X_k)\circ h^{-1})\big) && = V_{h(p)}\big((h^*A)(X_1,\dots,X_k)\circ h^{-1}\big)\\
\overset{\text{chain rule}}{=}&d_{h(p)}h^{-1}(V_{h(p)})\big((h^*A)(X_1,\dots,X_k)\big) && = h_*^{-1}V\big((h^*A)(X_1,\dots,X_k)\big)(p).
\end{aligned}
$$

Second, we consider the $(i+1)$-th summand

$$
\begin{aligned}
&A(h_*X_1,\dots,h_*X_{i-1},[V,h_*X_i],h_*X_{i+1},\dots,h_*X_k)\circ h\\
=&A(h_*X_1,\dots,h_*X_{i-1},h_*[h_*^{-1}V,X_i],h_*X_{i+1},\dots,h_*X_k)\circ h\\
=&(h^*A)(X_1,\dots,X_{i-1},[h_*^{-1}V,X_i],X_{i+1},\dots,X_k),
\end{aligned}
$$

where we used

$$h_*[h_*^{-1}V,X_i] = h_*[h_*^{-1}V,h_*^{-1}h_*X_i] = h_*h_*^{-1}[V,h_*X_i] = [V,h_*X_i].$$

This completes the proof. □

Lemma 4.2. *[Representation associated to an eigenvalue] Let $\lambda \in \sigma(M)$. Define, for any isometry $h : M \to M$, the map*

$$
\begin{aligned}
L_h : E_\lambda &\to E_\lambda\\
f &\mapsto f\circ h^{-1}.
\end{aligned}
$$

The map L_h is well defined and linear. The map

$$
\begin{aligned}
\tau_\lambda : \mathrm{Iso}(M) &\to \mathrm{GL}(E_\lambda)\\
h &\mapsto L_h
\end{aligned}
$$

is a group morphism, or, in other words, a representation of $\mathrm{Iso}(M)$.

Proof. We have to show that $L_h \in \mathrm{GL}(E_\lambda)$. The linearity of L_h is clear. What remains to be seen is the fact that if f is an eigenvector of the Laplacian with eigenvalue λ then so is $f\circ h$. First, we note that the Laplacian can be described in terms of the Lie derivative. One characterisation of the divergence $\mathrm{div}(X)$ of a vector field X on M is the equation

$$d(X\lrcorner\mathrm{vol}) = \mathrm{div}(X)\mathrm{vol},$$

where d denotes the exterior derivative, $\lrcorner$ denotes the interior product, and vol denotes the volume form with respect to the Riemannian metric g. In this situation, Cartan's magic formula states that

$$\mathcal{L}_X\mathrm{vol} = d(X\lrcorner\mathrm{vol}) + X\lrcorner d(\mathrm{vol}) = d(X\lrcorner\mathrm{vol}),$$

because the exterior derivative of forms of degree equal to the dimension of the manifold is zero. We have thus obtained the equation

$$-\triangle(f)\mathrm{vol} = \mathrm{div}(\boldsymbol{\nabla} f)\mathrm{vol} = \mathcal{L}_{\boldsymbol{\nabla} f}\mathrm{vol},$$

for any smooth $f : M \to \mathbb{R}$. Now, let us consider the composition of a smooth function $f : M \to \mathbb{R}$ and an isometry h of M. First, we note that the gradient of $f\circ h$ is the pushforward of the gradient of f by h^{-1}. To see this, let $p \in M$ and $v \in T_pM$ and compute

$$g(\boldsymbol{\nabla}(f\circ h)(p),v) = d_{h(p)}f\circ d_ph(v) = g(\boldsymbol{\nabla} f(h(p)),d_ph(v)) = g\big(d_{h(p)}h^{-1}\big(\boldsymbol{\nabla} f(h(p))\big),v\big),$$

where the last step makes use of the fact that h^{-1} is an isometry. We conclude that

$$\boldsymbol{\nabla}(f \circ h)(p) = d_{h(p)}h^{-1}\big(\boldsymbol{\nabla} f(h(p))\big)$$

holds for all $p \in M$, which is to say

$$\boldsymbol{\nabla}(f \circ h) = h_*^{-1}\boldsymbol{\nabla} f.$$

If f is an eigenfunction of the Laplacian with eigenvalue λ, it follows that

$$\mathcal{L}_{\boldsymbol{\nabla}(f\circ h^{-1})}\mathrm{vol} = \mathcal{L}_{h_*\boldsymbol{\nabla} f}(h^{-1})^*\mathrm{vol} \overset{\text{Lemma}(4.1)}{=} (h^{-1})^*\mathcal{L}_{\boldsymbol{\nabla} f}\mathrm{vol} = (h^{-1})^*(-\lambda f\mathrm{vol}) = (-\lambda f \circ h^{-1})\mathrm{vol},$$

where we used the fact $(h^{-1})^*\mathrm{vol} = \mathrm{vol}$, which holds for any isometry h^{-1} of M, twice. This allows us to conclude that

$$\triangle(f \circ h^{-1}) = \lambda f \circ h^{-1}.$$

Finally, we note that τ_λ satisfies

$$\tau_\lambda(gh)f = f \circ (gh)^{-1} = (f \circ h^{-1}) \circ g^{-1} = \tau_\lambda(g) \circ \tau_\lambda(h)f$$

for any $g, h \in \mathrm{Iso}(M)$ and $f \in E_\lambda$. Hence, τ_λ is a representation of the isometry group of M. The equation also proves that L_h is invertible with inverse $L_{h^{-1}}$. □

Suppose we are in a situation where it is known that any representation of $\mathrm{Iso}(M)$ splits into a sum of irreducible representations, and furthermore have a list of all the irreducible representations. In this case, we immediately obtain an integer equation for the dimension of E_λ, i.e. the multiplicity of the eigenvalue λ. In principle this approach is applicable when the isometry group is compact, but we will only need the statement for finite groups later. Of course, the equation will be of no use unless one can exclude the trivial representation from the decomposition of the representation.

Definition 4.3. A complex representation $\rho : G \to \mathrm{GL}(V)$ is called *realizable over* $\mathbb{R}$ if there exists a real representation $\tilde{\rho} : G \to \mathrm{GL}(W)$ such that ρ and $\tilde{\rho} \otimes \mathbb{C}$ are isomorphic as complex representations.

Definition 4.4. Let $\rho : G \to \mathrm{GL}(V)$ be a representation and W an invariant subspace. Define the restriction

$$\rho_{\big|W} : G \to \mathrm{GL}(W)$$

of ρ to W by setting

$$\rho_{\big|W} g \coloneqq (\rho g)_{\big|W}$$

for all $g \in G$.

Let $\rho : G \to \mathrm{GL}(V)$ be a complex representation. Then we can interpret V as a real vector space by restricting the scalar multiplication to $\mathbb{R}$. It has real dimension $2\dim_{\mathbb{C}} V$. This is what we mean by the phrase *ρ interpreted as a real representation.*

Theorem 4.5 ([10, Theorem 1, page 6])**.** *Let $\rho : G \to \mathrm{GL}(V)$ be a representation, G a finite group, and V a vector space over a field $\mathbb{F}$ such that $\#G = \sum_{i=1}^{\#G} 1 \in \mathbb{F}$ is not equal to 0 in $\mathbb{F}$. Then any ρ-invariant subspace W of V has a ρ-invariant complement U.*

Proof. Let U be a vector space complement of W in V, that is to say $U \oplus W = V$ and define the projection $\pi : V \to V, u + w \mapsto w$. By averaging π over the elements of G, we obtain a projection whose image is W and whose kernel is ρ-invariant. As is true for any projection, the map's kernel is complementary to its image, which completes the proof.

Define the map $\tilde{\pi}$ by

$$\tilde{\pi} \coloneqq \frac{1}{\#G}\sum_{g\in G} \rho(g) \circ \pi \circ \rho(g^{-1}).$$

First, we verify that $\tilde{\pi}$ is a projection, which is to say that $\tilde{\pi} \circ \tilde{\pi} = \tilde{\pi}$, or, equivalently, $\tilde{\pi}(w) = w$ for all $w \in \mathrm{Im}(\tilde{\pi})$. This property is immediately inherited from π. To see this, let $v \in V$ and compute

$$\begin{aligned}
\tilde{\pi}(\tilde{\pi}(v)) &= \frac{1}{\#G} \sum_{g \in G} \rho(g) \circ \pi \circ \rho(g^{-1}) \left(\frac{1}{\#G} \sum_{h \in G} \rho(h) \circ \pi \circ \rho(h^{-1})(v) \right) \\
&= \frac{1}{\#G} \sum_{g \in G} \sum_{h \in G} \frac{1}{\#G} \rho(g) \circ \pi \circ \rho(g^{-1}h) \circ \pi \circ \rho(h^{-1})(v) \\
&= \frac{1}{\#G} \sum_{g \in G} \sum_{h \in G} \frac{1}{\#G} \rho(gg^{-1}h) \circ \pi \circ \rho(h^{-1})(v) \\
&= \frac{1}{\#G} \sum_{g \in G} \tilde{\pi}(v) = \tilde{\pi}(v).
\end{aligned}$$

Second, we have to check that the kernel $\ker(\tilde{\pi})$ is a ρ-invariant, i.e. we have to show that $\tilde{\pi}(\rho(h)(v)) = 0$ for all $v \in \ker(\tilde{\pi})$ and $h \in G$. It suffices to show that $\rho(h)$ and $\tilde{\pi}$ commute. The computation is based on the observation that left multiplication is a bijection of the group. Note that

$$\begin{aligned}
\rho(h^{-1}) \circ \tilde{\pi} \circ \rho(h) &= \frac{1}{\#G} \sum_{g \in G} \rho(h^{-1}g) \circ \pi \circ \rho(g^{-1}h) \\
&= \frac{1}{\#G} \sum_{g \in G} \rho(g) \circ \pi \circ \rho(g^{-1}) \\
&= \tilde{\pi}.
\end{aligned}$$

Finally, it remains to be seen that $\mathrm{Im}(\tilde{\pi}) = W$. The inclusion '$\subset$' follows from the fact that $\rho(g) \circ \pi \circ \rho(g^{-1})(v)$ is an element of W for any $v \in V$ and $g \in G$. This holds because π is a projection onto W, and W is a ρ-invariant subspace of V. The inclusion '$\supset$' follows from the fact that $\tilde{\pi}(w) = w$, which is a direct consequence of $\rho(g) \circ \pi \circ \rho(g^{-1})(w) = \rho(gg^{-1})(w) = w$ for all $g \in G$ and $w \in W$. $\square$

Theorem 4.6. *Let $\rho : G \to \mathrm{GL}(V)$ be a representation. If any ρ-invariant subspace has a ρ-invariant complement, then ρ can be decomposed into a sum of irreducible representations.*

Proof. The proof is by induction on the dimension of the vector space on which G acts. There is nothing to show in the case $n = 1$, since $\{0\}$ and V are the only subspaces of any 1-dimensional vector space V.

Suppose the theorem holds for representations of G in vector spaces of dimension less or equal n. Let ρ be of dimension $n+1$. If ρ is irreducible, we need not show anything. If ρ has an invariant subspace W, then there exists a ρ-invariant complement U. Hence, $\rho = \rho\big|_W \oplus \rho\big|_U$ and the induction hypothesis applies to both summands, since the dimension of both U and W is at most n. Therefore, the sum of the decomposition of $\rho\big|_W$ and the decomposition of $\rho\big|_U$ is a decomposition of ρ given by irreducible representations. $\square$

The next two lemmata describe the relationship between real and complex irreducible representations.

Lemma 4.7. *[Real invariant subspaces of irreducible complex representation] Let G be a finite group. Let $\rho : G \mapsto \mathrm{GL}(V)$ be a complex irreducible representation of $dim_{\mathbb{C}} V = n$. Suppose that ρ interpreted as a real representation is not an irreducible real representation. Then there exists an n dimensional real subspace W of V that is invariant under ρ and such that the restriction $\rho|_W$ is an irreducible real representation.*

Furthermore, for any two invariant non-zero proper real subspaces $W_1, W_2 \subset V$, the representations $\rho|_{W_1}$ and $\rho|_{W_2}$ are isomorphic.

Proof. We consider the situation of a complex irreducible representation ρ with a non-trivial real invariant subspace W of V. Since ρ is irreducible, it must hold that

$$\{ \sum_i \lambda_i w_i \mid k \in \mathbb{N}, w_1, \dots, w_k \in W, \lambda_1, \dots, \lambda_k \in \mathbb{C} \},$$

which is a non-zero invariant complex subspace of V, is the entire space V. Therefore, the real dimension of W is at least n. As G is a finite group, an invariant subspace W always has a complementary invariant subspace U, i.e. $V = W \oplus U$. This subspace U must also have real dimension of at least n. This observation, combined with the integer equation $2 \dim_{\mathbb{C}} V = 2n = \dim_{\mathbb{R}} W + \dim_{\mathbb{R}} U$, implies that $\dim_{\mathbb{R}} W = n$. We have just shown that any non-trivial real invariant subspace of V has real dimension equal to n. Since any real invariant subspace of W is certainly a real invariant subspace of V, and therefore is either equal $\{0\}$ or of dimension n, it follows that W has no non-trivial real invariant subspaces. In other words, the restriction of ρ to W is an irreducible real representation.

In order to prove that any two invariant subspaces are isomorphic, we want to use character theory for representations in vector spaces over the field $\mathbb{R}$. Specifically, we want to apply the statement *if characters of two real representations are equal as maps from G to $\mathbb{R}$ then the representations are isomorphic.* Let W_1, W_2 be invariant real subspaces of V. We have already noted that $\rho_{|W_1} \otimes \mathbb{C} = \rho = \rho_{|W_2} \otimes \mathbb{C}$. This is useful, because the character of the complexification of a real representation is given by the inclusion of $\mathbb{R}$ into $\mathbb{C}$ composed with the character of the real representation. One way to see this, is the fact that any basis $B \coloneqq w_1, \dots, w_k$ of the real vector space W induces a basis $\tilde{B} \coloneqq w_1 \otimes 1, \dots, w_k \otimes 1$ of the complex vector space $W \otimes \mathbb{C}$ with the property that the matrix representation of $\rho(g)$ w.r.t. B is the same as the matrix representation of $\rho(g) \otimes \mathbb{C}$ w.r.t. $\tilde{B}$ for any element $g \in G$. In summary, we have seen that

$$\iota \circ \chi_{\rho_{|W_2}} = \chi_{\rho_{|W_2} \otimes \mathbb{C}} = \chi_\rho = \chi_{\rho_{|W_1} \otimes \mathbb{C}} = \iota \circ \chi_{\rho_{|W_1}},$$

where $\iota : \mathbb{R} \to \mathbb{C}$ is the inclusion map. Because the inclusion map is injective, we can conclude that the characters of $\rho_{|W_1}$ and $\rho_{|W_2}$ are the same, and the two representations are isomorphic. ☐

Lemma 4.8. *[Complexification of irreducible real representation] Let G be a finite group. Let $\rho : G \mapsto \mathrm{GL}(V)$ be an irreducible real representation. Then the complex representation $\rho \otimes \mathbb{C}$ is irreducible or has an invariant complex subspace W of complex dimension $\frac{\dim_{\mathbb{R}} V}{2}$. In the former case, there exists no real $\rho \otimes \mathbb{C}$-invariant subspace of $V \otimes \mathbb{C}$ whose real dimension is less than* $\dim_{\mathbb{R}} V$. *In the latter case, all other non-trivial invariant complex subspaces are of the same dimension, which implies that W is irreducible. Furthermore, $(\rho \otimes \mathbb{C})_{|W}$ and ρ are isomorphic when interpreted as real representations. Therefore, the fact that ρ is an irreducible real representation implies that $(\rho \otimes \mathbb{C})_{|W}$ interpreted as a real representation is irreducible.*

Proof. Suppose $V = \mathbb{R}^n$ and ρ is an irreducible real representation whose complexification $\rho \otimes \mathbb{C}$ is not irreducible. Let W be a (non-zero) invariant complex subspace of $\mathbb{R}^n \otimes \mathbb{C} = \mathbb{C}^n$. In the following, by abuse of notation, we refer to the image of $\mathbb{R}^n$ under the conventional inclusion of $\mathbb{R}^n$ in $\mathbb{C}^n$ by $\mathbb{R}^n$. The space $W \cap \mathbb{R}^n$ is a real invariant subspace and must therefore equal either $\{0\}$ or $\mathbb{R}^n$. In the latter case, the space W is equal to $\mathbb{C}^n$. Suppose now that $W \cap \mathbb{R}^n = \{0\}$. Then $U \coloneqq \big(W \oplus \kappa(W)\big) \cap \mathbb{R}^n$ is a non-zero real invariant subspace, wherefore its dimension must be n. Thus, it holds that $W + \kappa(W)$ is a complex subspace which contains $\mathbb{R}^n$. Hence, it must be equal $\mathbb{C}^n$. It follows that $n = \dim_{\mathbb{C}} W + \kappa(W) = 2 \dim_{\mathbb{C}} W$, which is to say that the complex dimension of W is $\frac{n}{2}$. The last equality in the previous sentence follows from the observation that $W \cap \kappa(W)$ is a subset of $W \cap \mathbb{R}^n = \{0\}$.

The claim about the two representations $(\rho \otimes \mathbb{C})_{|W}$ and ρ being isomorphic real representations is again shown by computing the characters. First, note that

$$\chi^{\mathbb{C}}_{(\rho\otimes\mathbb{C})_{|W}} + \overline{\chi^{\mathbb{C}}_{(\rho\otimes\mathbb{C})_{|W}}} = \chi^{\mathbb{C}}_{(\rho\otimes\mathbb{C})_{|W}} + \chi^{\mathbb{C}}_{(\rho\otimes\mathbb{C})_{|\kappa(W)}} = \chi^{\mathbb{C}}_{\rho\otimes\mathbb{C}} = \iota \circ \chi^{\mathbb{R}}_{\rho},$$

so that

$$2\,\mathrm{Re}\left(\chi^{\mathbb{C}}_{(\rho\otimes\mathbb{C})_{|W}}\right) = \chi^{\mathbb{R}}_{\rho}.$$

Second, we need to compare the trace of a complex endomorphism and the trace of the same endomorphism considered as a map between real vector spaces. For this purpose, consider a matrix representation $A+iB$ of $(\rho\otimes\mathbb{C})_{|W}(g)$ with respect to some complex basis of W for some $g \in G$ and $A, B \in \mathrm{Mat}(\dim_{\mathbb{C}} W \times \dim_{\mathbb{C}} W, \mathbb{R})$. Now, there is an obvious real basis induced by the previously chosen complex basis of W. With respect to this basis, the matrix representation of $(\rho \otimes \mathbb{C})_{|W}(g)$ is an element of $\mathrm{Mat}(2\dim_{\mathbb{C}} W \times 2\dim_{\mathbb{C}} W, \mathbb{R})$ and it is described in block form by $(A, -B, B, A)$. Here, the first entry is the block in the upper left corner, the second entry is the block in the upper right corner, and the third entry is the block in the lower left corner. The trace of this matrix is two times the trace of A. This happens to be equal to two times the real part of the trace of $A + iB$. We can now conclude

$$\chi^{\mathbb{R}}_{(\rho\otimes\mathbb{C})_{|W}} = 2\,\mathrm{Re}\left(\chi^{\mathbb{C}}_{(\rho\otimes\mathbb{C})_{|W}}\right) = \chi^{\mathbb{R}}_{\rho},$$

whence the two representations are isomorphic. □

Lemma 4.9 (Integer equation for the dimension of a representation)**.** *Let G be a finite group and $\rho : G \to \mathrm{GL}(V)$ a representation of G on the finite dimensional vector space V. Let S denote the set of irreducible representations of G. Let $d : S \to \mathbb{N}$ be the map that maps any representation to its dimension. Then there exists a finite subset $\{\rho_1, \dots, \rho_k\}$ of S of cardinality $k \in \mathbb{N}$ and $n_1, \dots, n_k \in \mathbb{N}$ such that*

$$\dim(V) = \sum_{i=1}^{k} n_i d(\rho_i).$$

Proof. The statement follows immediately from the existence of a decomposition of ρ into a sum of irreducible representations. □

5 Tessellations of the Hyperbolic Plane

Let $n \in \mathbb{N}$. In this section, X is either the n-dimensional hyperbolic space $\mathbb{H}^n$ or the n-dimensional sphere $\mathbb{S}^n_\epsilon$ of radius ϵ for some $\epsilon > 0$. In the proofs, we mostly consider only the case of hyperbolic space and the sphere of radius 1. Since most definitions, statements and proofs are taken from [9], we frequently mention the corresponding theorem or on which page one can find the corresponding statement in that book. For the sake of brevity, all references without an explicitly stated source refer to this book.

5.1 An Assortment of Definitions and Some Lemmata

Definition 5.1 (Tessellation)**.** Let I be a set and $M_i \subset X$ for each $i \in I$. Then the collection of sets $\{ M_i \mid i \in I \}$ is a tessellation of X if its elements cover X and have pair wise disjoint interiors, and if the collection is locally finite. In other words, a tessellation satisifies

i $\bigcup_{i \in I} M_i = X$

ii $\forall i, j \in I : \mathring{M}_i \cap \mathring{M}_j = \emptyset$

iii $\forall x \in X \exists \epsilon > 0 : \# \{ i \in I \mid B_\epsilon(x) \cap M_i \neq \emptyset \} < \infty$.

Definition 5.2 (Fundamental domain (p. 236))**.** Let Γ be a group acting on X. A fundamental domain for Γ is a connected, open subset $F \subset X$ for which there exists a set $F \subset M \subset \overline{F}$ such that M contains exactly one representative of any orbit $z \in X/\Gamma$.

Another way to phrase the above conditions is to say: a fundamental domain F for Γ is an open connected set such that $(g\overline{F})_{g\in\Gamma}$ cover X and the sets $(gF)_{g\in\Gamma}$ are pairwise disjoint. The condition that F contains at most one representative of any orbit $z \in X/\Gamma$ is equivalent to the pair-wise disjointness of the sets $(gF)_{g\in\Gamma}$. The condition that the sets $(g\overline{F})_{g\in\Gamma}$ cover X is equivalent to the existence of a representative of any orbit $z \in X/\Gamma$ in $\overline{F}$. So a set M can be obtained by choosing one representative in $\overline{F} \setminus F$ for each orbit of Γ which is not represented in F.

Definition 5.3 (Convex set (p.195))**.** A subset C of H^n is convex, if the unique geodesic between any two points $x, y \in C$ is contained in C. A subset C of S^n is convex, if the unique shortest geodesic between x and y, where $x, y \in C$ are two non-antipodal points, is contained in C.

Note that, according to this definition, all sets of the form $\{ x, -x \}$ for any point $x \in \mathbb{S}^n$ are convex subsets of $\mathbb{S}^n$.

Definition 5.4 (Side of convex set (p. 198))**.** A side of a convex set C is a non-empty maximal convex subset of its topological boundary ∂C.

Definition 5.5 (Polyhedron)**.** A subset P of X is *a polyhedron in* X, if P is non-empty, closed, convex (with respect to X), and the set of sides of P is locally finite.

Definition 5.6 (Planes and Dimension [p.117, p.123, p.195])**.** If $X = \mathbb{S}^n \subset \mathbb{R}^{n+1}$, then an m-plane is the intersection of X with a $(m+1)$-dimensional vector subspace of $\mathbb{R}^{n+1}$. Equivalently, we can define an m-plane as the image of $\operatorname{span}\{ e_1, \dots, e_{m+1} \} \cap \mathbb{S}^n$ under an isometry of $\mathbb{S}^n$, where e_i refers to the i-th standard basis vector of $\mathbb{R}^{n+1}$.

If $X = \mathbb{H}^n$, we define m-planes for the Poincare ball model B^n as follows: An m-plane is the image of the set $\operatorname{span}\{ e_1, \dots, e_m \} \cap B^n$ under an isometry of B^n, where e_i refers to the i-th standard basis vector of $\mathbb{R}^n$.[4]

[4]There are more explicit descriptions of the m-planes of hyperbolic space for any of its models, but since the isometry group acts transitively on the set of m-planes, see [9, Proof of Theorem 4.5.3.], our definition is equivalent to the one used in the textbook.

Let $C \subset X$ be a non-empty convex set. The dimension of C is defined as the minimal integer m such that C is contained in an m plane P of X, i.e.

$$\dim_X C := \min \{ m \in \mathbb{N}_0 \mid \exists P \ m\text{-plane of } X : C \subset P \}.$$

It is a fact that there exists exactly one ($\dim_X C$)-plane, subsequently referred to as $< C >$, which contains C.

We will state - without proof - a handful of properties of polyhedrons which we will need for the proof of a later theorem. The first of these is the following

Theorem 5.7 (Theorem 6.2.3, p. 197)**.** *Let $C \subset X$ be a non-empty convex set. The interior $\mathring{C}$ of C viewed as a subset of $< C >$ is nonempty.*

Proof. Omitted. □

This is incredibly useful, because it implies

Lemma 5.8. *If C is a non-empty convex set of the same dimension as X, then C is contained in $\overline{\mathring{C}}$. An equivalent way of phrasing this is*

$$\forall x \in C \forall k \in \mathbb{N} : B_{\frac{1}{k}}(x) \cap \mathring{C} \neq \emptyset.$$

In fact, the somewhat stronger statement

$$\forall x \in C \exists K > 0 \forall 0 < \epsilon < K : S_\epsilon(x) \cap \mathring{C} \neq \emptyset$$

holds.

Proof. Let $x \in C$ and $\epsilon > 0$. Our task is to produce an element $y \in \mathring{C}$ whose distance to x is ϵ. Since the dimensions of C and X are the same, the plane $< C >$ generated by C is the entire space X. By the previous lemma, the interior of C as a subset of X is non-empty, i.e. there exists a $z \in C$ and $\delta > 0$ such that $B_\delta(z) \subset C$. We may assume that $x \notin B_\delta(z)$, for if this is the case, the lemma is obviously true.

The idea is to pull this open ball towards x along the geodesic segments between x and the elements of the ball. If X is the unit sphere, we can assume without loss of generality that z and x are not antipodal and that $-x \notin B_\delta(z)$. This can be achieved by choosing a different point in the ball and a smaller radius, if necessary. Given these preliminaries, we can solve the problem by rescaling the preimage under the exponential map at x. More precisely, the exponential map $\exp_x$ restricted to the set $B_\pi(0) \subset T_x\mathbb{S}^n$ is a diffeomorphism onto $\mathbb{S}^n \setminus \{ -x \}$. In the case where X is a hyperbolic space, the exponential map $\exp_x : T_xX \to X$ is a diffeomorphism. Now let $r := d(x, z)$ and move the preimage of $B_\delta(z)$ under the exponential map towards the origin, i.e. consider

$$U := \frac{\epsilon}{r} \exp_x^{-1}(B_\delta(z)).$$

We claim that

$$y := \exp_x(\frac{\epsilon}{r} \exp_x^{-1}(z)) \text{ satisfies } y \in \mathring{C} \text{ and } d(x, y) = \epsilon.$$

Obviously, the distance between y and x is equal ϵ, as long as $\epsilon < \pi$. If $\epsilon \leq r$, then y lies in C, because C contains the unique minimizing geodesic between the points x and z and this geodesic is given by $[0, 1] \to X, t \mapsto \exp_x(t \exp_x^{-1}(z))$. The same argument shows that the entire open subset $\exp_x(U)$ is contained in C. Hence, y lies in the interior of C.

By inspecting the proof, one can conclude that the constant K from the statement of the lemma can be chosen in dependence on x as the minimum of r and π. □

Another helpful statement is the next

Theorem 5.9 (Theorems 6.3.1 and 6.3.4, p. 201)**.** *The side of a m-dimensional polyhedron is a $(m-1)$-dimensional polyhedron.*

Proof. Omitted. $\square$

Definition 5.10 (Dihedral angle)**.** Suppose P is a one-dimensional polyhedron in $\mathbb{S}^1$. Then P is either $\mathbb{S}^1$ or a geodesic segment whose length is less than or equal π. Thus, P has less than two sides unless its length is strictly less than π, in which case it has exactly two sides. The two sides S, T of P are the endpoints $x, y \in S^1$ of the geodesic segment P, i.e. $S = \{x\}, T = \{y\}$. The dihedral angle $\theta(S,T)$ between the two sides is defined by

$$\theta(S,T) \in (0,\pi) \text{ and } \cos(\theta(S,T)) = (x,y),$$

where $(\cdot,\cdot)$ denotes the euclidean inner product on $\mathbb{R}^2$.

Suppose P is a compact two-dimensional polyhedron in $\mathbb{H}^2$. The purpose of the compactness assumption is to exclude the possibility of so called ideal vertices, that is vertices which lie on the boundary of the disk in the Poincare disk model. Then the dihedral angle $\theta(S,T)$ is defined for any two distinct sides S, T of P which intersect. Note that the sides are one-dimensional convex sets, whence they must be geodesic segments. Suppose that S intersects T, and let $x \in S \cap T$. By translating x to the origin 0, we can assume without loss of generality that S and T are subsets of straight lines through the origin. Therefore, there exist $u, v \in \mathbb{S}^1$ such that $S \subset \{tu \mid t \in [0,1]\}$ and $T \subset \{tv \mid t \in [0,1]\}$. We define the dihedral angle $\theta(S,T)$ of S and T by

$$\theta(S,T) \in (0,\pi) \text{ and } \cos(\theta(S,T)) = (u,v),$$

where $(\cdot,\cdot)$ denotes the euclidean inner product on $\mathbb{R}^2$.

For the definition for general polyhedra, we refer to [9, p. 213].

Definition 5.11 (Reflections)**.** We define reflections for $(n-1)$-planes P in X. Note that P is a co-dimension one submanifold of X, wherefore the tangent space of X at a point $x \in P$ can be decomposed into $T_xX \cong T_xP \oplus \mathbb{R}v$ where $v \in T_xX$ is any vector orthogonal to T_xP. The reflection in P is an isometry Φ_P of X which restricts to the identity on P and whose derivative at any point $x \in P$ is given by

$$\begin{aligned} T_xP \oplus \mathbb{R}v &\to T_xP \oplus \mathbb{R}v \\ w + tv &\mapsto w - tv. \end{aligned} \tag{3}$$

There exists a unique map with the above properties, see the following lemma.

Let Q be an n-dimensional polyhedron in X and S a side of Q and P the unique $(n-1)$-plane in X which contains S. The reflection $g_S : X \to X$ in the side S is defined as the reflection Φ_P in the plane P.

Lemma 5.12. *For any $(n-1)$-plane P in X, there exists a unique reflection in P.*

Proof. Let $p \in P$. The uniqueness of a reflection in the plane P is a direct consequence of the surjectivity of the exponential map $\exp_p$ and the equation

$$\phi \circ \exp_p = \exp_p \circ d_p\phi$$

which holds for any isometry ϕ that fixes p.

The above equation is of little use when it comes to defining an isometry. While we know that an isometry is completely determined if its derivative at a point is known, the equation says little about the existence of isometries. Given a complete connected Riemannian manifold M, a point $p \in M$, and an orthogonal linear endomorphism L of T_pM, one cannot expect an isometry whose derivative a p is L to exist. Of course, one can compute the derivative $d_v \exp_p$ of $\exp_p$ at the point $v \in T_pM$, use the equation to obtain an explicit description of the derivative of Φ at $\exp_p(v)$, and check if it is an isometry. We choose a different approach to avoid these calculations.

We claim that it suffices to show the existence of reflections in the plane P_X, which is defined as

$$P_X := \begin{cases} P_{\mathbb{H}^n} := \operatorname{span}\{\, e_1, \dots, e_{n-1} \,\} \cap B^n & \text{if } X = \mathbb{H}^n \\ P_{\mathbb{S}^n} := \operatorname{span}\{\, e_1, \dots, e_n \,\} \cap \mathbb{S}^n & \text{if } X = \mathbb{S}^n, \end{cases}$$

where e_i is the i-th standard basis vector in $\mathbb{R}^n$ and $\mathbb{R}^{n+1}$, respectively.

Suppose that P is an $(n-1)$-plane in X. By our definition of planes in X, there exists an isometry Ψ which maps P_X onto P. Suppose that a reflection Φ_{P_X} in the plane P_X exists. Then the conjugation of Φ_{P_X} by Ψ is a reflection in the plane P: The map $\Psi \circ \Phi_{P_X} \circ \Psi^{-1}$ restricts to the identity on P and its derivative at the point $x \in P$ satisfies Equation (3). The statement about the derivative is a consequence of the fact that the derivative of the isometry Ψ maps the orthogonal decomposition $T_x P_X \oplus (T_x P_X)^\perp$ onto the orthogonal decomposition $T_x P \oplus (T_x P)^\perp$.

We complete the proof by providing an explicit formula for the reflection in P_X. The maps

$$\begin{aligned} \Phi_{P_{\mathbb{S}^n}} : \mathbb{S}^n &\to \mathbb{S}^n \\ \sum_{i=1}^{n} \lambda_i e_i + \lambda_{n+1} e_{n+1} &\mapsto \sum_{i=1}^{n} \lambda_i e_i - \lambda_{n+1} e_{n+1} \end{aligned} \tag{4}$$

and

$$\begin{aligned} \Phi_{P_{\mathbb{H}^n}} : \mathbb{B}^n &\to \mathbb{B}^n \\ \sum_{i=1}^{n-1} \lambda_i e_i + \lambda_n e_n &\mapsto \sum_{i=1}^{n-1} \lambda_i e_i - \lambda_n e_n \end{aligned} \tag{5}$$

are isometries and it is not hard to check that they have the desired properties. □

5.2 Links of a Polyhedron - Some Local Considerations

Lemma 5.13. *The distance sphere around any $x \in X$ of a sufficiently small radius $r > 0$ in X is a Euclidean sphere, i.e. $S_r(x) := \{\, y \in X \mid d(x, y) = r \,\} \subset X$ is isometric to the $(n-1)$-dimensional sphere $\mathbb{S}^{n-1}_{\tilde{R}}$ of some radius $\tilde{R} > 0$.*

Proof. Case $X = \mathbb{H}^n$: We use the Poincaré ball model B of hyperbolic space. It suffices to consider the distance spheres $S_r(0)$ around the origin. For any $p \in B$, there exists an isometry which maps p to 0. This isometry restricts to an isometry between $S_r(p)$ and $S_r(0)$. Thus, if $S_r(0)$ is a Euclidean sphere, the same is true for $S_r(p)$.

In order to show that $S_r(0)$ is a Euclidean sphere, it suffices to show the equality of sets $S_r(0) = \{\, z \in B \mid \|z\|_{\text{euclid}} = R \,\}$. This is a consequence of the Riemannian metric of X at the point $x \in X$ being given by

$$(v, w) \mapsto \frac{4}{\sum_i x_i^2} \sum_i v_i w_i,$$

where we use the standard identification of the tangent space at x with $\mathbb{R}^n$. The description clearly shows that the metric only depends on the Euclidean norm of the point. If the above equality of sets holds, it follows that the metric on $S_r(0)$ is a constant multiple of the metric on $\mathbb{S}^{n-1}_r$.

Let $x \in S_r(0)$ for some radius $r > 0$ and define $R := \|x\|_{\text{Euclid}}$. We claim that any element $y \in B$ with $\|y\|_{\text{Euclid}} = R$ also lies in the distance sphere $S_r(0)$. It is clear that the action of the orthogonal group $O(n)$ on $\mathbb{R}^n$ leaves the unit ball invariant and preserves the Riemannian metric. Thus, they restrict to isometries of the Poincaré Ball. Furthermore, the group $O(n)$ acts transitively on the Euclidean distance sphere $\{\, z \in \mathbb{R}^n \mid \|z\|_{\text{Euclid}} = R \,\}$ and the action of any element of the group fixes the origin $0 \in B$. Finally, we observe that if $\Phi \in O(n)$ maps x onto y, then $d(y, 0) = d(\Phi(x), \Phi(0)) = d(x, 0)$. Thereby, we have proven the inclusion

$$\{ y \in B \mid \|y\|_{\text{Euclid}} = R \} \subset S_r(0). \tag{6}$$

We claim that the converse inclusion also holds, i.e.

$$S_r(0) \subset \{ y \in B \mid \|y\|_{\text{Euclid}} = R \} .$$

It suffices to show that $y, z \in B$ with $\|y\|_{\text{Euclid}} \neq \|z\|_{\text{Euclid}}$ also satisfy $d(y, 0) \neq d(z, 0)$. Without loss of generality, we may assume that $\|y\|_{\text{Euclid}} > \|z\|_{\text{Euclid}}$.

Let $\eta : [0, c] \to X$ be a minimizing geodesic segment from 0 to y parameterized at unit speed. Since the complement of $\{ v \in \mathbb{R}^n \mid \|v\|_{\text{Euclid}} = \|z\|_{\text{Euclid}} \}$ in B is disconnected and the path η starts in one component and ends in the other, the path must intersect the aforementioned set at some time $t_0 < c$. That is to say, $\|\eta(t_0)\|_{\text{Euclid}} = \|z\|_{\text{Euclid}}$. Since η is distance preserving, we conclude that

$$d(z, 0) \overset{\text{Equation (6)}}{=} d(\eta(t_0), 0) \overset{t_0 < c}{<} d(\eta(c), 0) = d(y, 0).$$

Case $X = \mathbb{S}^n$: An argument similar to the above should work, which is why we worded the above proof without using uniqueness of geodesics. We omit the proof and refer to the general literature. □

Theorem 5.14 (Theorem 6.4.1, p. 214)**.** *Suppose that P is a polyhedron of dimension $m > 1$. Let $x \in P$, and $\epsilon > 0$ sufficiently small. Then the set $P \cap S_\epsilon(x)$ is an $(m-1)$-dimensional polyhedron in $S_\epsilon(x)$. Let $\mathcal{S}(P \cap S_\epsilon(x))$ denote the set of sides of the polyhedron $P \cap S_\epsilon(x)$. The set of sides of said polyhedron is given by*

$$\mathcal{S}(P \cap S_\epsilon(x)) = \{ S \cap S_\epsilon(x) \mid S \text{ is a side of } P \text{ and } x \in S \} .$$

The dihedral angle between any two adjacent sides S, T of P is equal to the dihedral angle between the adjacent sides $S \cap S_\epsilon(x)$ and $T \cap S_\epsilon(x)$ of $P \cap S_\epsilon(x)$.[5]

Proof. We consider only the case where P is a compact 2-dimensional polyhedron in the hyperbolic plane $\mathbb{H}^2$ and refer to the textbook for the general case. There are three distinct possibilities for a point $x \in P$. It either lies in the interior of P, in exactly one side of P, or exactly two sides of P. To see this, note that P is the union of its boundary and its interior, and the boundary is covered by the sides of P. Therefore, if x is not in $\mathring{P}$, it must lie in at least one side of P. All that remains to be shown to prove the claim is the fact that there are at most two sides S_1 and S_2 of P which contain x. This is illustrated in the Figure 1.

Suppose that $x = 0 \in B^2$ lies in the sides S_1 and S_2. Let S_3 be a third side with $x \in S_3$. We will show that S_3 is not a subset of the boundary ∂P. Any side of P is a one-dimensional polyhedron and must therefore be a segment of a geodesic line. The planes $< S_1 >$ and $< S_2 >$ are straight lines through the origin. They divide the disk B^2 into four quadrants, one of which contains P, see [Theorem 6.3.2., p.202]. Let us label the quadrant which contains the polyhedron by the name Q_1. Choose two points B and C which lie in the sides S_1 and S_2, respectively. Consider the triangle with vertices $A \coloneqq 0$, B, and C. Since the polyhedron P is convex, the triangle is contained in P. This can easily be seen to be true in the projective disk model of the hyperbolic plane, where the geodesics are all contained in straight lines. The side S_3 is a subset of some ray $\mathcal{L}$ which starts in the origin and lies in the quadrant Q_1. In particular, the side S_3 intersects the interior of the triangle with vertices A, B, and C. Therefore, the intersection of S_3 and $\mathring{P}$ is non-empty, which is to say that S_3 is not a subset of the boundary ∂P, and, hence, is not a side of P.

The case $x \in \mathring{P}$ is trivial, because $P \cap S_\epsilon(x) = S_\epsilon(x)$ for sufficiently small ϵ.

[5] We have not yet defined the term *adjacent* for two sides of a polyhedron. For a compact 2-dimensional polyhedron P in $\mathbb{H}^2$, two distinct sides S, T are adjacent if they intersect. If this is the case, S and T must be geodesic segments which have exactly one endpoint in common. For a compact 1-dimensional polyhedron P in $\mathbb{S}^1$, two sides S and T of P are adjacent if $S \neq P$. For the general definition, see [p.213].

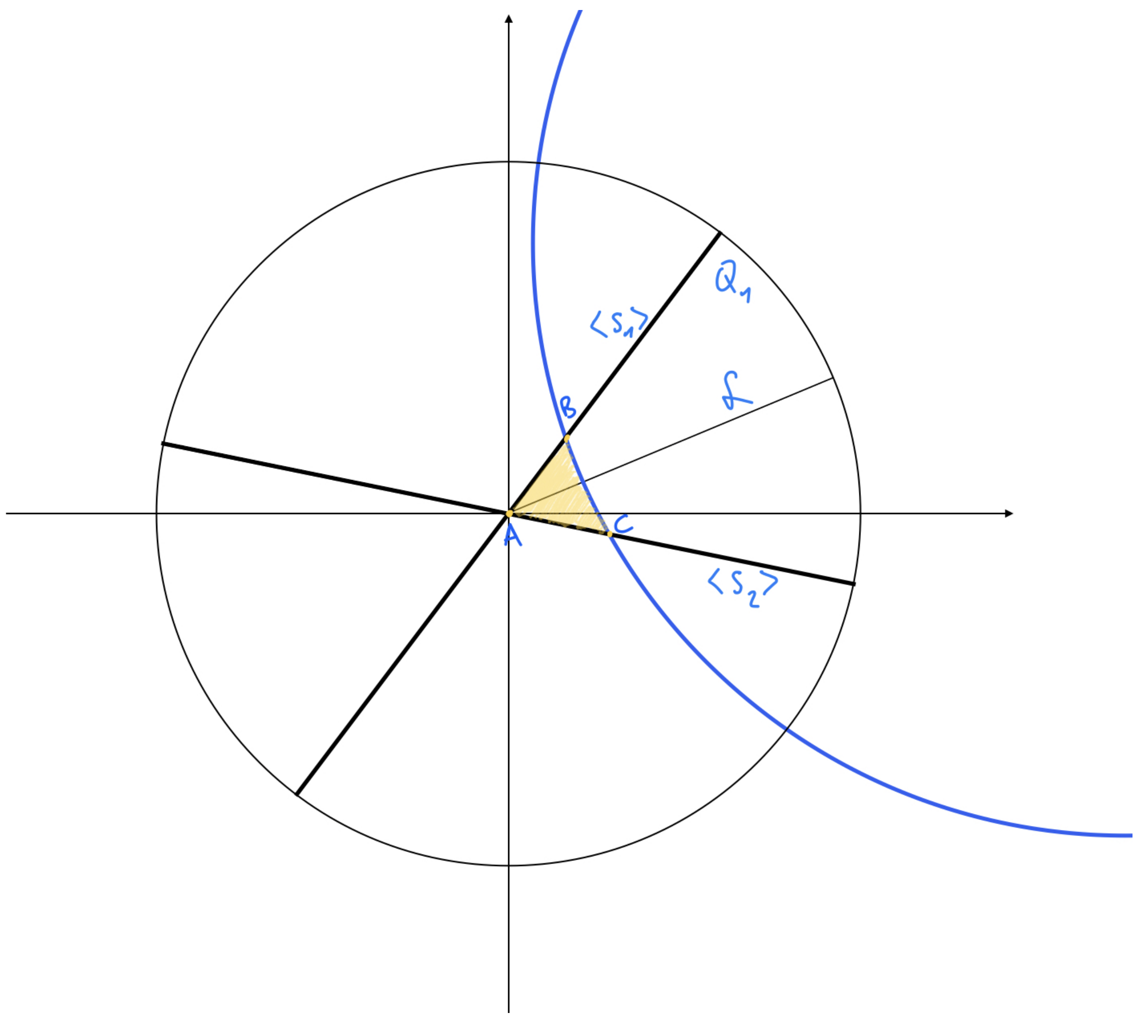

FIGURE 1: The 2-dimensional polyhedron P is contained in the quadrant Q_1 created by the straight lines $< S_1 >$ and $< S_1 >$, which are the 1-planes generated by the sides S_1 and S_2 of P. The point $A \in P$ can lie in at most two sides S_1 and S_2 of P.

If x lies on an edge S of P and is not a vertex, then there exists a neighborhood of x which intersects only one edge of P. One can show that $S_\epsilon(x)$ intersects the edge that contains x at antipodal points of the circle $S_\epsilon(x)$. One way to see this is to translate x to the origin $0 \in B$ so that the edge S is mapped to a straight line through the origin and $S_\epsilon(x)$ becomes a Euclidean circle centered at the origin. By the definition of convexity, the polyhedron $P \cap S_\epsilon(x)$ has exactly one side, namely the set $S_\epsilon(x) \cap S$.

The final case is the situation that x is a vertex, i.e. x lies on two edges. By applying a translation, if necessary, we can assume that x is the origin. Then the distance sphere $S_\epsilon(0)$ is a Euclidean circle centered at 0, and the edges are straight lines through the origin. This should imply that $P \cap S_\epsilon(x)$ is the segment between the two points on the circle which are given by the intersection of either edge with the circle. $\square$

Lemma 5.15. *Let S be a side of an n-dimensional polyhedron Q. Suppose $x \in S$ and $\epsilon := \epsilon(x) > 0$ is sufficiently small. Then*

$$g_S\big|_{S_\epsilon(x)} = g_{S \cap S_\epsilon(x)},$$

i.e. the restriction of the reflection g_S in the side S of Q to the sphere $S_\epsilon(x)$ is the reflection in the side $S \cap S_\epsilon(x)$ of the polyhedron $Q \cap S_\epsilon(x)$. The latter set is interpreted as a polyhedron in the geometric sphere $\mathbb{S}_R^{n-1}$, with $R > 0$ being chosen such that $S_\epsilon(x)$ is isometric to $\mathbb{S}_R^{n-1}$.

Proof. We only prove the lemma for $X = \mathbb{H}^n$. We begin by considering the special case $x = 0 \in B^n$ and $< S >= P_{\mathbb{H}^n}$, whereafter we reduce the general case to this situation. Let $\epsilon > 0$ be sufficiently small for Lemma 5.13 and Theorem 5.14 to hold. This guarantees the existence of an $R > 0$ such that $S_\epsilon(0)$ is isometric to $\mathbb{S}_R^{n-1}$. Furthermore, it implies that $Q \cap S_\epsilon(0)$ is an $(n-1)$-dimensional polyhedron in $\mathbb{S}_R^{n-1}$.

The explicit descriptions in Equations (4) and (5) show that the reflection $\Phi_{P_{\mathbb{H}^n}}$ in the plane $P_{\mathbb{H}^n}$ restricted to the sphere $S_\epsilon(0)$ is the reflection in the plane $P_{\mathbb{S}_R^{n-1}} = S_\epsilon(0) \cap \operatorname{span}(e_1, \ldots, e_{n-1})$, i.e.

$$g_S\big|_{S_\epsilon(0)} = \Phi_{<S>}\big|_{S_\epsilon(0)} = \Phi_{P_{\mathbb{H}^n}}\big|_{\mathbb{S}_R^{n-1}} = \Phi_{P_{\mathbb{S}_R^{n-1}}}.$$

To prove the equality $g_S\big|_{S_\epsilon(0)} = g_{S \cap S_\epsilon(0)}$, we need only show that the plane $< S \cap S_\epsilon(0) >$ generated by the side $S \cap S_\epsilon(0)$ is given by $P_{\mathbb{S}_R^{n-1}}$. This would imply

$$\Phi_{P_{\mathbb{S}_R^{n-1}}} = \Phi_{< S \cap S_\epsilon(0) >} = g_{S \cap S_\epsilon(0)},$$

which combined with the previous equation results in the desired conclusion. We proceed to proving the equality of the planes $< S \cap S_\epsilon(0) >$ and $P_{\mathbb{S}_R^{n-1}}$. First, note that S is contained in the vector subspace $\operatorname{span}(e_1, \ldots, e_{n-1})$ of $\mathbb{R}^n$, whence $S \cap S_\epsilon(0)$ is contained in $\operatorname{span}(e_1, \ldots, e_{n-1}) \cap S_\epsilon(0) = P_{\mathbb{S}_R^{n-1}}$. The latter object is an $(n-2)$-plane. Since the side $S \cap S_\epsilon(0)$ of the polyhedron $Q \cap S_\epsilon(0)$ is of dimension $n-2$ and there exists only one $(n-2)$-plane which contains a specific $(n-2)$-dimensional convex subset, it follows that $< S \cap S_\epsilon(0) >= P_{\mathbb{S}_R^{n-1}}$.

Now, consider the general case of a side S of Q which is not necessarily contained in $P_{\mathbb{H}^n}$, a point $x \in S$, and an $\epsilon > 0$ which is sufficiently small. We claim that there exists an isometry Ψ which maps $< S >$ onto $P_{\mathbb{H}^n}$ and x onto 0. According to our definition of hyperbolic planes, there exists an isometry $\tilde{\Psi}$ which maps the plane $P_{\mathbb{H}^n}$ onto the $(n-1)$-plane $< S >$ generated by the side S of Q. By composing $\tilde{\Psi}$ with a translation τ of $\tilde{\Psi}(x)$ to 0 which maps $P_{\mathbb{H}^n}$ onto $P_{\mathbb{H}^n}$, we obtain an isometry which maps $< S >$ onto $P_{\mathbb{H}^n}$ and x onto 0. A map with the properties required of τ exists: any translation between the two points maps $P_{\mathbb{H}^n}$ onto the intersection of B^n and an $(n-1)$-dimensional vector subspace V of $\mathbb{R}^n$. We can subsequently move $V \cap B^n$ into the position of $P_{\mathbb{H}^n}$ by applying an orthogonal transformation. Therefore, the composition of the translation and the orhogonal transformation is an isometry of hyperbolic space which maps $\tilde{\Psi(x)}$ to 0 and $P_{\mathbb{H}^n}$ onto itself. In summary, the map $\Psi := \tau \circ \tilde{\Psi}$ has the desired properties.

The reflection in $< S >$ is obtained from the reflection in $\Psi^{-1}(< S >) = P_{\mathbb{H}^n}$ by conjugation with Ψ, as was discussed in detail in the proof of Lemma 5.12. We conclude that

$$\begin{aligned} g_S\big|_{S_\epsilon(x)} &= \Psi \circ g_{\Psi^{-1}(S)} \circ \Psi^{-1}\big|_{S_\epsilon(x)} \\ &\overset{\text{Previous discussion}}{=} \Psi \circ g_{\Psi^{-1}(S)\cap S_\epsilon(0)} \circ \Psi^{-1}\big|_{S_\epsilon(x)} = g_{S\cap S_\epsilon(x)}, \end{aligned}$$

where the final equality stems from the following consideration: First, the conjugation by $\Psi\big|_{S_\epsilon(0)}$ of the reflection in the plane $< \Psi^{-1}(S) \cap S_\epsilon(0) >$ is the reflection in the plane $\Psi(< \Psi^{-1}(S) \cap S_\epsilon(0) >)$. Second, the latter plane is no other than $< \Psi(\Psi^{-1}(S) \cap S_\epsilon(0)) > = < S \cap S_\epsilon(x) >$, which is a consequence of a uniqueness argument similar to one we employed in the first part of the proof. $\square$

Definition 5.16. For any polyhedron P in X, define Γ_P as the subgroup of the isometry group of X generated by the reflections in the sides of P, i.e.

$$\Gamma_P := < g_S \mid \text{S is a side of P} > .$$

Let $x \in P$. Define $\Gamma_P(x)$ as the subgroup of Γ_P generated by the reflections in those sides of P which contain x, i.e.

$$\Gamma_P(x) := < g_S \mid \text{S is a side of P and } x \in S > .$$

Lemma 5.17. *Let $x \in P$ and $\epsilon > 0$ sufficiently small. Recall that $P \cap S_\epsilon(x)$ is a polyhedron in the sphere $S_\epsilon(x)$. In this situation, the homomorphism*

$$\begin{aligned} \rho : \Gamma_P(x) &\longrightarrow \Gamma_{P\cap S_\epsilon(x)} \\ g &\mapsto g\big|_{S_\epsilon(x)} \end{aligned}$$

is well defined and bijective.

Proof. The proof is based on the following equality of sets: If a subgroup H of a group G is generated by a subset $F \subset H \subset G$, then

$$H = \{ \prod_{i=1}^{k} h_i \mid k \in \mathbb{N} \text{ and } h_1, \ldots, h_k \in F \cup F^{-1} \} .$$

The inclusion '$\supset$' is obvious: finite products of elements of a subgroup are themselves elements of the subgroup. The inclusion '$\subset$' follows from the fact that

$$H = \bigcap_{J \in \{ I \subset G \mid I \text{ is subgroup of } G \text{ and } F \subset I \}} J$$

and the fact that

$$\{ \prod_{i=1}^{k} h_i \mid k \in \mathbb{N} \text{ and } h_1, \ldots, h_k \in F \cup F^{-1} \}$$

is a subgroup of G that contains F.

ρ is well defined:

Let $g \in \Gamma_P(x)$, then there exist sides $S_1, \ldots, S_k$ of P, all of which contain x, such that $g = \prod_{i=1}^k g_{S_i}$. It follows that

$$g\big|_{S_\epsilon(x)} = \left(\prod_{i=1}^{k} g_{S_i} \right)\Bigg|_{S_\epsilon(x)} = \prod_{i=1}^{k} g_{S_i}\big|_{S_\epsilon(x)} \overset{\text{lemma5.15}}{=} \prod_{i=1}^{k} g_{S_i \cap S_\epsilon(x)}. \tag{7}$$

Since the sets $S_i \cap S_\epsilon(x)$ are sides of the polyhedron $P \cap S_\epsilon(x)$, the maps $g_{S_i \cap S_\epsilon(x)}$ are elements of $\Gamma_{P\cap S_\epsilon(x)}$. We have shown that $g\big|_{S_\epsilon(x)}$ is the product of elements of $\Gamma_{P\cap S_\epsilon(x)}$, which implies that the map itself is an element of the group.

ρ is surjective:
Let $g\prime \in \Gamma_{P\cap S_\epsilon(x)}$. There exist sides $S_1', \ldots, S_k'$ of $P \cap S_\epsilon(x)$ such that $g\prime$ is equal $\prod_{i=1}^k g_{S_i'}$. By Theorem 5.14, the sides S_i' are given by $S_i \cap S_\epsilon(x)$ for some sides S_i of P which contain x. This tells us that $g\prime$ is equal to $\prod_{i=1}^k g_{S_i \cap S_\epsilon(x)}$, and we obtain

$$g\prime = \left(\prod_{i=1}^k g_{S_i} \right)\Bigg|_{S_\epsilon(x)}$$

from the latter two equalities of equation (7). Since the maps $g_{S_i} \in \Gamma_P(x)$, this shows that $g\prime \in \mathrm{im}(\rho)$. We conclude that ρ is surjective.

ρ is injective:
It suffices to show that an isometry g of X which fixes the point x is uniquely determined by its values on the sphere $S_\epsilon(x)$. This, in turn, would follow if we could show that g is determined on $B_{\frac{\epsilon}{2}}(x)$: In this case, the derivative of g at x is determined, and this determines the isometry via the exponential map, which is surjective, since X is complete.

Let $y \in B_{\frac{\epsilon}{2}}(x)$. Assume that there exists a geodesic γ and an interval $0 \in I \subset \mathrm{dom}(\gamma)$ such that $\gamma(I) \subset B_\epsilon(x)$, $\gamma(0) = y$ and there exist $t_1, t_2 \in I$ such that $\gamma(t_1)$ and $\gamma(t_2)$ are distinct points in $S_\epsilon(x)$. The isometry g maps γ onto the geodesic $g\gamma$ with the property $g\gamma(I) \subset B_\epsilon(x)$. A geodesic inside a sufficiently small ball connecting two points in said ball is uniquely determined (up to reparameterisation to different constant speeds) by those two points. It follows that

$$g\gamma\big|_{[t_1,t_2]}$$

is equal to the unique unit speed geodesic $\tilde{\gamma}$ which connects the two points $g\gamma(t_1)$ and $g\gamma(t_2)$ and lies in $B_\epsilon(x)$. This proves that the values of g are determined on $B_{\frac{1}{2}\epsilon}(x)$, since

$$gy = g\gamma(0) = \tilde{\gamma}(0)$$

and $\tilde{\gamma}$ depends only on $g\big|_{S_\epsilon(x)}$.

About the existence of such geodesics: Let τ be a unit speed geodesic with $\tau(0) = y$. Note that we have

$$d(\tau(t), y) = |t|$$

for sufficiently small values t and

$$B_\epsilon(x) \subset B_{\frac{3\epsilon}{2}}(y)$$

by the triangle inequality. Therefore, $\tau(\frac{3\epsilon}{2}), \tau(-\frac{3\epsilon}{2}) \in X \setminus B_{\frac{3\epsilon}{2}}(y) \subset X \setminus B_\epsilon(x)$. The two maps

$$\begin{aligned} f_\pm : \mathbb{R} &\longrightarrow \mathbb{R} \\ t &\longmapsto d(\tau(\pm t), x) \end{aligned}$$

are continuous and satisfy

$$f_\pm(0) = d(y, x) < \frac{\epsilon}{2} \text{ and } f_\pm(\frac{3\epsilon}{2}) > \epsilon.$$

By the intermediate value theorem, there exist $t_1, t_2 > 0$ with

$$f_+(t_1) = f_-(t_2) = \epsilon,$$

which is just another way of saying $\tau(t_1), \tau(-t_2) \in S_\epsilon(x)$. Note also, that $B_{\frac{\epsilon}{2}}(y) \subset B_\epsilon(x)$, which implies $\tau\left((-\frac{\epsilon}{2}, -\frac{\epsilon}{2})\right) \subset B_\epsilon(x)$. If we let

$$t_1 := \inf\{\, t \in (0,\infty) \mid \tau(t) \in S_\epsilon(x)) \,\} \text{ and } t_2 := \inf\{\, t \in (0,\infty) \mid \tau(-t) \in S_\epsilon(x)) \,\},$$

then it holds that $\frac{\epsilon}{2} < t_1, t_2 < \frac{3\epsilon}{2}$, $\tau\left((-t_2, t_1)\right) \subset B_\epsilon(x)$, and $\tau(t_1), \tau(-t_2) \in S_\epsilon(x)$. Upon closer inspection, it is clear that our argument does not really require $\tau(t_1) \neq \tau(-t_2)$. However, in the above, we did claim the existence of a geodesic through y which intersects the distance sphere $S_\epsilon(x)$ in two distinct points. Indeed, this is the case, as there are no self-intersecting geodesics of length strictly less than 2π in $\mathbb{S}^n$ or $\mathbb{H}^n$. $\square$

5.3 An Example: Tessellations of the Circle

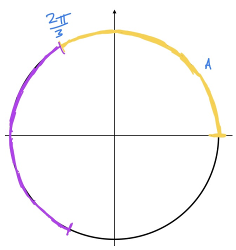

FIGURE 2: An example of the case where the length of the segment A is an odd submultiple of 2π. The purple segment is the image of A under both a reflection and a rotation contained in the reflection group Γ_A generated by A.

Let us return to the situation of a polyhedron P in the hyperbolic plane $\mathbb{H}^2$. Let $\epsilon > 0$ and $x \in P$. The previous subsection explained that the sets $P \cap S_\epsilon(x)$ are 1-dimensional polyhedrons in a geometric sphere. The idea behind the proof showing that

$$\{\, gP \mid g \in \Gamma_P \,\}$$

is a tessellation of the hyperbolic plane, is to show that

$$\{\, gP \cap S_\epsilon(x) \mid g \in \Gamma_P(x) \,\}$$

is a tessellation of S^1_ϵ and use this to understand the 'local' properties of the collection of translates of P in $\mathbb{H}^2$.

First, we note an obvious necessary condition for a family of translates of a polyhedron to tessellate the circle. We interpret the circle S^1 as a subset of the complex plane. By a segment of the circle, we mean any set of the form $A(c_1, c_2) := \{\, \mathrm{e}^{i\theta} \mid c_1 \leq \theta \leq c_2 \,\}$, where c_1, c_2 are positive real numbers with $|c_2 - c_1| < 2\pi$ and $c_1 < c_2$. The length l of any such segment is equal to $c_2 - c_1$. If a collection $(A_i)_{i \in I}$ of segments of the same length tessellate the circle, then

$$2\pi = \sum_{i \in I} l(g_i(A)) = l(A_1)\#I,$$

because the overlap between two different segments consists of one point and has length zero. Therefore, the cardinality of I is finite and the length of A is equal to $\frac{2\pi}{\#I}$.

Lemma 5.18. *Let $c_1, c_2 \geq 0$ with $l := c_2 - c_1 = \frac{2\pi}{m}$ for some integer $m > 1$ and consider the segment $A := A(c_1, c_2)$. Then the the collection*

$$\{\, gA \mid g \in \Gamma_A \,\}$$

is a tessellation of S^1 if and only if m is even. In that case,

$$\Gamma_A = \{\, \mathrm{e}^{2ikl} \mid k \in \{\, 1, \ldots, \frac{m}{2} \,\}\,\} \cup \{\, g_{\mathrm{e}^{i(c_1 - kl)}} \mid k \in \{\, 1, \ldots, \frac{m}{2} \,\}\,\}.$$

Proof. For the sake of readibility, we introduce some labels: $a := e^{ic_1}$, $b := e^{ic_2}$. We mention that the groups $\Gamma_A(\{\, a \,\})$ and $\Gamma_A(\{\, b \,\})$ generated by the reflection in the point a and b, respectively, contain only two elements. These are the reflection of S^1 in the respective point and the identity. This follows from the fact that reflections are self inverse.

Next, we compute the group Γ_A. As a first step, we note that the underlying set can be described as

$$\{\, 1 \,\} \cup \{\, (g_a g_b)^k \mid k \in \mathbb{N} \,\} \cup \{\, (g_b g_a)^k \mid k \in \mathbb{N} \,\} \cup \{\, g_b (g_a g_b)^k \mid k \in \mathbb{N} \,\} \cup \{\, g_a (g_b g_a)^k \mid k \in \mathbb{N} \,\}. \quad (8)$$

We claim that the reflections in the boundary points of the segment are given by

$$g_a \mathrm{e}^{i\theta} = \mathrm{e}^{i(c_1-(\theta-c_1))} \quad \text{and} \quad g_b \mathrm{e}^{i\theta} = \mathrm{e}^{i(c_2-(\theta-c_2))} = \mathrm{e}^{i(2c_2-\theta)},$$

for any $\theta \in \mathbb{R}$. One way to see that g_a is the reflection around a, is to note that it is an isometry since it preserves the distance between any two points. Futhermore, the path $t \mapsto \mathrm{e}^{i(c_1+t)}$ is mapped onto the path $t \mapsto \mathrm{e}^{i(c_1-t)}$. The velocities at $t = 0$ are given by $i\mathrm{e}^{ic_1}$ and $-i\mathrm{e}^{ic_1}$, respectively. This shows that the derivative of g_a at a is $-\mathrm{Id}_{T_a S^1}$. The same argument applies to g_b.

Visually, it is clear that the product of an even number of reflections is a rotation and the product of an odd number of reflections is a reflection. We proceed with an explicit descriptions of the products. First, observe that

$$g_b g_a \mathrm{e}^{i\theta} = g_b \mathrm{e}^{i(2c_1-\theta)} = \mathrm{e}^{i(2c_2-2c_1+\theta)} = \mathrm{e}^{2i(c_2-c_1)}\mathrm{e}^{i\theta} = \mathrm{e}^{2il}\mathrm{e}^{i\theta}$$

for any $\theta \in \mathbb{R}$. That is to say, $g_b g_a$ is the rotation by the angle $2l$. We will notate this by

$$g_b g_a = \mathrm{e}^{2il},$$

where we interpret the complex number e^{2ikl} as an isometry of S^1. Since $g_a g_b$ is the inverse of $g_b g_a$, the equality $g_a g_b = \mathrm{e}^{-2il}$ follows immediately. Another immediate consequence is the equation

$$(g_b g_a)^k = \mathrm{e}^{2ikl}.$$

Assume that m is odd, i.e. $m = 2p + 1$ for some $p \in \mathbb{N}$. It is not hard to see that the powers of the composition $g_b g_a$ of the reflections in the boundary points of the segment A are exactly the rotations by the angles $2\pi \frac{k}{2p+1}$, where $k = 0, \ldots, 2p$.[6] Therefore, the cardinality of the set $\{ (g_a g_b)^k \mid k \in \mathbb{N} \}$ is equal to m. The group Γ_A also contains at least one reflection. Thus, its cardinality is at least $m + 1$. This shows that

$$l \# \Gamma_A \geq \frac{2\pi}{m}(m+1) > 2\pi,$$

which allows us to conclude that the collection in (5.18) is not a tessellation. See Figure 2.

[6] Since this argument is not terribly interesting, it is relegated to a footnote. First, we note that $2ikl = 2\pi i \frac{2k}{2p+1}$, wherefore

$$\{ (g_b g_a)^k \mid k \in \mathbb{N} \} = \{ \mathrm{e}^{2\pi i \frac{2k}{2p+1}} \mid k \in \mathbb{N} \}.$$

Next, we take a closer look at $\frac{2k}{2p+1}$ mod 1. We note the following equalities of sets:

$$\begin{aligned}
&\{ \frac{2k}{2p+1} \bmod 1 \mid k =\in \mathbb{N} \} = \{ \frac{2k}{2p+1} \mid k = 0, \ldots, 2p \} \\
=& \{ \frac{2k}{2p+1} \bmod 1 \mid k = 0, \ldots, p \} \cup \{ \frac{2p}{2p+1} + \frac{2k}{2p+1} \bmod 1 \mid k = 1, \ldots, p \} \\
=& \{ \frac{2k}{2p+1} \bmod 1 \mid k = 0, \ldots, p \} \cup \{ \frac{2k-1}{2p+1} \bmod 1 \mid k = 1, \ldots, p \} \\
=& \{ \frac{k}{2p+1} \bmod 1 \mid k = 0, \ldots, 2p \}.
\end{aligned}$$

From the above, the equality

$$\{ (g_b g_a)^k \mid k \in \mathbb{N} \} = \{ \mathrm{e}^{2\pi i \frac{k}{2p+1}} \mid k = 0, \ldots, 2p \}$$

follows immediately.

Suppose now that m is even, i.e. $m = 2p$ for some $p \in \mathbb{N}$. We want to describe the group Γ_A explicitly, for which purpose we note that

$$\begin{aligned}
(g_a g_b)^k &= \mathrm{e}^{-2ikl} = \mathrm{e}^{-2\pi i \frac{k}{p}} \\
g_a (g_b g_a)^k \mathrm{e}^{i\theta} &= \mathrm{e}^{i(2c_1 - 2kl - \theta)} = \mathrm{e}^{i(2(c_1 - \pi\frac{k}{p}) - \theta)} = g_{\mathrm{e}^{i(c_1 - \pi\frac{k}{p})}} \mathrm{e}^{\theta} \\
g_b (g_a g_b)^k \mathrm{e}^{i\theta} &= \mathrm{e}^{i(2(c_2 - \pi\frac{k}{p}) - \theta)} = g_{\mathrm{e}^{i(c_2 - \pi\frac{k}{p})}} \mathrm{e}^{\theta}
\end{aligned}$$

for any $k \in \mathbb{N}$ and $\theta \in \mathbb{R}$. These descriptions imply

$$g_b (g_a g_b)^k = g_a (g_b g_a)^{k-1}$$

for any $k \in \mathbb{N}$, because of the fact that

$$c_2 - \pi\frac{k}{p} = c_1 + l - \pi\frac{k}{p} = c_1 - \pi\frac{k-1}{p}.$$

The reflection around a point $x = \mathrm{e}^{i\theta_0}$ in S^1 is the same map as the reflection around its antipodal point $-x$, because $\mathrm{e}^{i(2\theta_0 - \theta)} = \mathrm{e}^{i(2(\theta_0 + \pi) - \theta)}$ for any $\theta \in \mathbb{R}$. In other words,

$$\{\, g_{\mathrm{e}^{i(c_1 - \pi\frac{k}{p})}} \mid k \in \mathbb{N}_0 \,\} = \{\, g_{\mathrm{e}^{i(c_1 - \pi\frac{k}{p})}} \mid k = 1, \dots p \,\}.$$

To summarize the above discussion, we have seen that

$$\{\, (g_a g_b)^k \mid k \in \mathbb{N} \,\} = \{\, \mathrm{e}^{-2\pi i \frac{k}{p}} \mid k = 1, \dots, p \,\} = \{\, \mathrm{e}^{2\pi i \frac{k}{p}} \mid k = 1, \dots, p \,\} = \{\, (g_b g_a)^k \mid k \in \mathbb{N} \,\}$$

and

$$\{\, g_b (g_a g_b)^k \mid k \in \mathbb{N} \,\} = \{\, g_a (g_b g_a)^k \mid k \in \mathbb{N}_0 \,\} = \{\, g_{\mathrm{e}^{i(c_1 - \pi\frac{k}{p})}} \mid k = 1, \dots, p \,\}.$$

By inserting these results into the the description of Γ_A given in (8), we arrive at the conclusion

$$\Gamma_A = \{\, \mathrm{e}^{2\pi i \frac{k}{p}} \mid k = 1, \dots, p \,\} \cup \{\, g_{\mathrm{e}^{i(c_1 - \pi\frac{k}{p})}} \mid k = 1, \dots, p \,\}.$$

The circle is tessellated by the $2p$ segments of length $\frac{\pi}{p}$. For every segment given as the image of A under a rotation, the adjacent segments are given as the image of A under a reflection. See Figure 3. □

5.4 Sufficient Conditions for a Polyhedron to Induce a Tessellation

The next lemma, affirming the existence of a so called Lebesgue number for open covers of a compact metric space, is a basic statement from topology. We will use it in the main proof of this section. The reason the theorem is included here is that it is usually stated for compact metric spaces. We, however, will apply it to a cover of a compact subset of a metric space *by open subsets of the ambient metric space.*

Lemma 5.19 (Lebesgue number, see [2]). *Let (M, d) be a metric space and $K \subset M$ a compact subset. Let I be a set and let $(U_i)_{i \in I}$ be a collection of open subsets of M which cover K. Then it holds that*

$$\exists l > 0 \forall x \in K \exists i \in I : B_l(x) \subset U_i.$$

Proof. Omitted. The proof in the referenced textbook requires only minor modifications.

□

Before we prove the main theorem of this section, we briefly discuss the local situation of tessellations induced by a compact hyperbolic triangle P whose angles are sub-multiples of π. Consider the situation depicted in Figure 4. If the point $p := p_2$ lies in the interior of the triangle, there is no need to involve any reflections to cover a neighborhood of p_2. If a point $p := p_3$ lies on exactly one side of the triangle, a neighborhood of p in $\mathbb{H}^2$ is given by the union of P and its image under the reflection in the side that contains p. The interesting case are the vertices of the triangle, i.e. points $p := p_1$ at which two sides intersect. In this situation, we look at sufficiently small distance spheres $S_{\epsilon_1}(p_1)$ and apply the discussion from Lemma 5.18 to the segment of the circle colored in purple in the figure, namely the intersection of P and $S_{\epsilon_1}(p_1)$. This tells us that $S_{\epsilon_1}(p_1)$ is tessellated by the images of the segment $S_{\epsilon_1}(p_1) \cap P$ under a finite subgroup of the isometry group of $\mathbb{H}^2$. Since the same holds for circles of all radii $0 < \epsilon < \epsilon_1$, we can conclude that the ball $B_{\epsilon_1}(p_1)$ is tessellated by the images of $B_{\epsilon_1}(p_1) \cap P$ under a finite subgroup of the isometry group. The idea of reducing the local situation to a tessellation of a lower dimensional sphere is explained in detail in the proof of the next theorem and works for hyperbolic space of all dimensions $n \geq 2$. The difficulty lies in making sure that the reflection group of the polyhedron does not contain any isometries which map the polyhedron onto a set which intersects the ball $B_{\epsilon_1}(p_1)$ except for the isometries in the subgroup determined in the local argument.

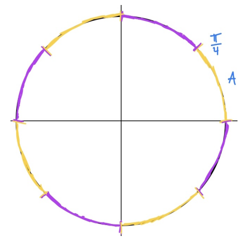

FIGURE 3: An example of the case where the length of the segment A is an even sub-multiple of 2π. The purple segments are images of A under a unique reflection of the reflection group Γ_A. The yellow segments are images of A under a unique rotation of the reflection group Γ_A.

Theorem 5.20 (Theorem 7.1.3, p.265)**.** *Let P be an n-dimensional compact polyhedron in either one of the two spaces $X := \mathbb{S}^n$ or $X := \mathbb{H}^n$. If the dihedral angle between any two adjacent sides S and T of P is an element of $\{ \frac{\pi}{n} \mid n \in \mathbb{N} \}$, then*

$$\{ gP \mid g \in \Gamma_P \}$$

is a tessellation of the space.

Proof. We begin by introducing some notation used subsequently: Let $\mathcal{S}$ denote the set whose elements are the sides of P. Let $\mathcal{S}(x)$ denote the set whose elements are the sides of P which contain x.

The proof proceeds by induction on the dimension of the space X. The case of $\mathbb{S}^1$ was described in detail in an earlier example (see Lemma 5.18). The case of $\mathbb{H}^1$ is omitted. Suppose now that $n \in \mathbb{N}$ with $n > 1$.

Part One: *Let $x \in P$. Then there exists $\epsilon := \epsilon(x) > 0$ such that $\{ gP \cap B_\epsilon(x) \mid g \in \Gamma_P(x) \}$ is a tessellation of $B_\epsilon(x)$. Furthermore, $\Gamma_P(x)$ is finite.*

The statement is trivial if $x \in \mathring{P}$: Since some neighborhood of x is contained in $\mathring{P} = P \setminus \partial P = P \setminus (\cup_{S \in \mathcal{S}} S)$, x is not contained in any side of P. Therefore, $\Gamma_P(x)$ is the trivial subgroup. For an $\epsilon > 0$ small enough to guarantee $B_\epsilon(x) \subset P$, the claim reduces to

$$\{ B_\epsilon(x) \} \text{ is a tessellation of } B_\epsilon(x),$$

which is obviously true.

Now, let $x \in \partial P$. Let $\epsilon > 0$ be sufficiently small so that $B_\epsilon(x)$ only intersects the sides of P that contain x and $\epsilon < \frac{\pi}{2}$. For any $0 < \delta < \epsilon$, Theorem 5.14 implies that the set $P \cap S_\delta(x)$ is a polyhedron in an $(n-1)$-dimensional sphere whose dihedral angles are still submultiples of pi.

Until now, we have used the notation gU for the image of $U \subset X$ under the map $g : X \to X$. For the next couple of sentences the notation $g(U)$ is used instead for the sake of clarity. The induction hypothesis applied to this polyhedron implies that

$$\{ g\prime(P \cap S_\delta(x)) \mid g\prime \in \Gamma_{P \cap S_\delta(x)} \}$$

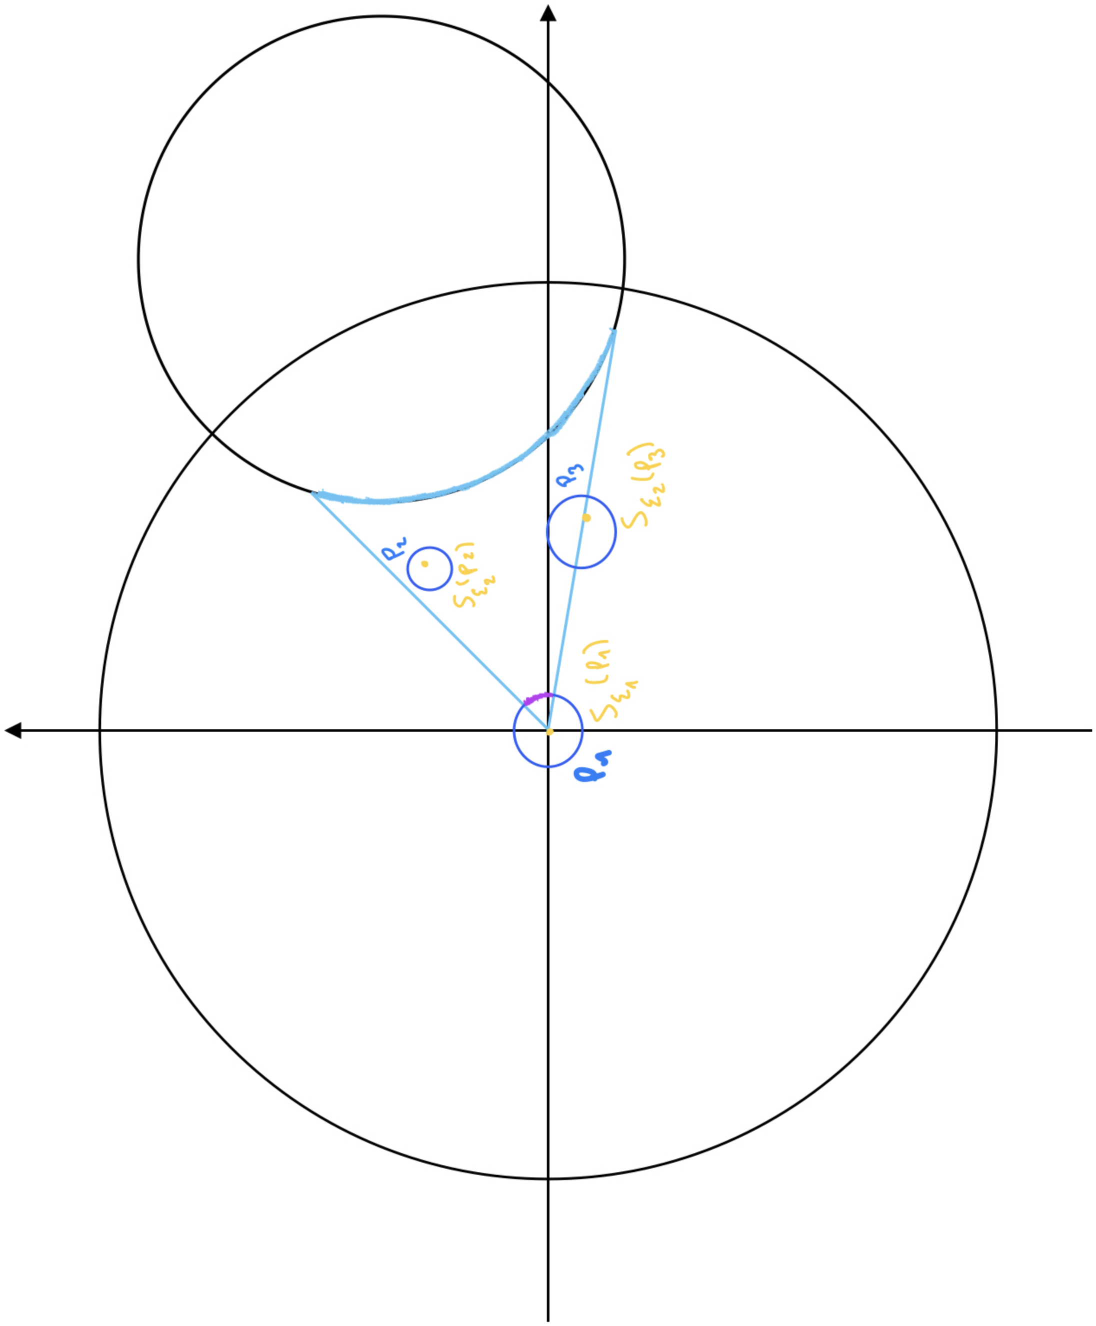

Figure 4: The three distinct possibilities for the local situation of tessellation induced by a compact 2-dimensional polyhedron (in the figure, a triangle) in the hyperbolic plane depicted in the Poincare disk model.

is a tessellation of $S_\delta(x)$. To show the equality of sets

$$\{ g\prime(P \cap S_\delta(x)) \mid g\prime \in \Gamma_{P\cap S_\delta(x)} \} = \{ g(P) \cap S_\delta(x) \mid g \in \Gamma_P(x) \},$$

we use the map ρ from lemma (5.17). The inclusion $' \subset'$ follows from the equality

$$g\prime(P \cap S_\delta(x)) = \rho^{-1}(g\prime)(P) \cap S_\delta(x),$$

for any $g\prime \in \Gamma_{P\cap S_\delta(x)}$, and the inclusion $' \supset'$ is a consequence of

$$g(P) \cap S_\delta(x) = \rho(g)(P \cap S_\delta(x)),$$

for any $g \in \Gamma_P(x)$. Since the two collections of sets are the same and one is a tessellation, the same holds for the other, i.e.

$$\{ g(P) \cap S_\delta(x) \mid g \in \Gamma_P(x) \} \text{ is a tessellation of } S_\delta(x). \tag{9}$$

This information allows us to deduce that $\{ gP \cap B_\epsilon(x) \mid g \in \Gamma(x) \}$ is a tessellation of $B_\epsilon(x)$. First, note that

$$\begin{aligned}
\bigcup_{g\in\Gamma(x)} gP \cap B_\epsilon(x) \quad &= \bigcup_{g\in\Gamma(x)} gP \cap \left(\{ x \} \cup \bigcup_{0<\delta<\epsilon} S_\delta(x) \right) \\
= \{ x \} \cup \bigcup_{0<\delta<\epsilon} \bigcup_{g\in\Gamma(x)} gP \cap S_\delta(x) \quad &\overset{\text{Equation9}}{=} \{ x \} \cup \bigcup_{0<\delta<\epsilon} S_\epsilon(x) \\
= B_\epsilon(x).
\end{aligned}$$

Second, let $g, h \in \Gamma_P(x)$ and suppose the set

$$\left(g\mathring{P} \cap B_\epsilon(x) \right) \cap \left(h\mathring{P} \cap B_\epsilon(x) \right)$$

is not empty. Then the image of this set under g^{-1}, given by $\mathring{P} \cap B_\epsilon(x) \cap g^{-1}h\mathring{P}$, is non-empty. Let $y \in \mathring{P} \cap B_\epsilon(x) \cap g^{-1}h\mathring{P}$. Since $x \in \partial P$ and $y \in \mathring{P}$, we have $y \neq x$. Therefore, $0 < d_X(x, y) < \epsilon$ and

$$y \in \left(\mathring{P} \cap S_{d_X(x,y)}(x) \right) \cap \left(g^{-1}h\mathring{P} \cap S_{d_X(x,y)}(x) \right) = \emptyset,$$

where the last equality follows from the statement in Equation (9). This contradicts the fact that the empty set does not contain any elements. Hence the interior of the elements of our collection are pairwise disjoint.

Third, again from the statement in Equation (9), it follows that $\{ g\mathring{P} \cap S_\delta(x) \mid g \in \Gamma_P(x) \}$ is a collection of pairwise disjoint open subsets of a compact space. Hence, it is finite. From Lemma 5.8, we know that $\mathring{P} \cap S_\delta(x) \neq \emptyset$. It follows that all sets of the form $g\mathring{P} \cap S_\delta(x)$ for maps g which fix x are non-empty. Since the sets of that form are pairwise disjoint and non-empty, they are certainly pairwise distinct. Hence, the finite collection is in bijection with $\Gamma_P(x)$, whence the latter is finite, also.

Part Two: *The collection of sets* $\{ gP \mid g \in \Gamma_P \}$ *covers* X.
The idea is as follows: We show that there exists a positive real number l such that $B_l(x)$ is contained in $\bigcup_{g\in\Gamma_P} gP$ for all $x \in \bigcup_{g\in\Gamma_P} gP$. Then we conclude by showing that this property implies $\bigcup_{g\in\Gamma_P} gP = X$.

By the result of the previous part, we know that there exists an $\epsilon(x) > 0$ for any $x \in P$ such that

$$B_{\epsilon(x)}(x) = \bigcup_{g\in\Gamma_P(x)} gP \cap B_{\epsilon(x)}(x) = \left(\bigcup_{g\in\Gamma_P(x)} gP \right) \cap B_{\epsilon(x)}(x) \subset \bigcup_{g\in\Gamma_P} gP. \tag{10}$$

Applying the Lebesgue number lemma to the open cover $\{ B_{\epsilon(x)}(x) \mid x \in P \}$ of the compact set P gives us an $l > 0$ with the following property:

$$\forall x \in P \exists y \in P : B_l(x) \subset B_{\epsilon(y)}(y). \tag{11}$$

The previous two numbered equations combined state that the ball of radius l around any element $x \in P$ is contained in the union of the translates of P, i.e.

$$\forall x \in P : B_l(x) \subset \bigcup_{g \in \Gamma_P} gP.$$

Note that the set $\bigcup_{g\in\Gamma_P} gP$ is invariant under any $h \in \Gamma_P$, i.e. $h \bigcup_{g\in\Gamma_P} gP \subset \bigcup_{g\in\Gamma_P} gP$. This allows us to deduce that the Property (5.4) extends to all elements of $\bigcup_{g\in\Gamma_P} gP$: If $x \in \bigcup_{g\in\Gamma_P} gP$, then $x = gp$ for some $g \in \Gamma$ and $p \in P$. Since g is an isometry, it maps a ball around p onto the ball of the same radius around gp. In particular, $B_l(x) = B_l(gp) = gB_l(p)$. Since $B_l(p)$ is contained in $\bigcup_{g\in\Gamma_P} gP$, so is $gB_l(p)$. Another way to phrase the above is to say

$$\left(y \in X \text{ and } \exists x \in \bigcup_{g \in \Gamma_P} gP : d(x, y) < l \right) \implies y \in \bigcup_{g \in \Gamma_P} gP.$$

The only subset of X with this property is X itself. To be more specific, if a subset $S \subset X$ satisfies

$$(y \in X \text{ and } \exists x \in S : d(x, y) < \epsilon) \implies y \in S \tag{12}$$

for some $\epsilon > 0$, then S is equal X.

To prove this, it suffices to show that the ball $B_{n\epsilon}(x)$ is contained in S for some fixed $x \in S$ and all $n \in \mathbb{N}$. We show this by induction on $n \in \mathbb{N}$. Let $x \in S$. It is clear that $B_\epsilon(x)$ is contained in S. Now let $n > 1$ and assume that $B_{(n-1)\epsilon}(x) \subset S$. We want to show that any element $y \in B_{n\epsilon}(x)$ is contained in S. To do so, we produce an element $z \in X$ which satisfies $d(z, y) < \epsilon$ and $d(x, z) < (n-1)\epsilon$. Any such z has the properties $z \in S$ and $y \in B_\epsilon(z)$, whence property (12) guarantees that $y \in S$. We return to the existence of a point with the above properties. Let $\gamma : [0, d(x, y)] \to X$ be a geodesic from x to y whose length is equal to the distance between x and y. This geodesic must also be the shortest path from $\gamma(t_1)$ to $\gamma(t_2)$ for any elements t_1, t_2 of its domain. Heuristically, the idea is to use the point $\gamma(d(x, y) - \epsilon)$, but its distance to the point y is ϵ, which means it won't work for our purposes. [7] We need to get a little closer to y without leaving the ball $B_{(n-1)\epsilon}(x)$. The point $z := \gamma(d(x, y) - \epsilon + \frac{n\epsilon - d(x,y)}{2})$ does the job: From the inequality $d(x, y) < n\epsilon$, it follows that

$$d(z, y) = \epsilon - \frac{n\epsilon - d(x, y)}{2} < \epsilon \quad \text{and} \quad d(x, z) = \frac{d(x, y) + n\epsilon}{2} - \epsilon < (n-1)\epsilon.$$

Part Three: *Assume that for any $x \in P$ and $g \in \Gamma_P$ it holds that*

$$gP \cap B_{\epsilon(x)}(x) \neq \emptyset \implies g \in \Gamma_P(x). \tag{13}$$

Then the collection $\{ gP \mid g \in \Gamma_P \}$ is locally finite, and the interiors of its elements are pairwise disjoint.

We leave the general case $x \in X$ to the reader and deal only with a point $x \in P$. By assumption,

$$\{ g \in \Gamma_P \mid gP \text{ intersects } B_{\epsilon(x)}(x) \}$$

is a subset of the finite set $\Gamma_P(x)$. Hence, only finitely many members of the collection $\{ gP \mid g \in \Gamma_P \}$ intersect $B_{\epsilon(x)}(x)$. Since x was an arbitrary element of X, the collection is locally finite.

Next, suppose that $g, h \in \Gamma$ are two elements of Γ_P such that $g\mathring{P}$ intersects $h\mathring{P}$. We want to show that $g = h$. Note that $\mathring{gP} = g\mathring{P}$, because g is a homeomorphism. The previous assumption implies that $\mathring{P}$ intersects $h^{-1}g\mathring{P}$. Let $x \in \mathring{P} \cap h^{-1}g\mathring{P}$. Then $h^{-1}gP$ intersects $B_{\epsilon(x)}(x)$ and Equation (13) implies that $h^{-1}g \in \Gamma_P(x)$. The fact that $x \in \mathring{P}$ tells us that x does not lie in any side of P, whence $\Gamma_P(x)$ is the

[7] Formally, this does not make sense if $y \in B_\epsilon(x)$, as $d(x, y) - \epsilon$ is less than zero and does not lie in the domain of γ. In this case, however, there was nothing to show in the first place.

subgroup of Γ generated by the empty set, and as such only contains the identity element. It follows that $g = h$, and we conclude that the interiors of the elements of the collection are pairwise disjoint.

Part Four: *The Condition (13) holds.*
The goal is to find a description of the elements of $(g,p),(h,q) \in \Gamma_P \times P$ for which $gp = hq$ holds. Our approach consists of defining an equivalence relation $\sim$, in the hope that

$$\kappa : \tilde{X} \longrightarrow X$$
$$[g,p] \longmapsto gp$$

defined on the set of equivalence classes $\tilde{X} := (\Gamma_P \times P)/\sim$ will turn out to be injective. We will endow $\tilde{X}$ with the quotient topology coming from the projection and the product topology on $\Gamma_P \times P$, while Γ_P and $P \subset X$ are endowed with the discrete and the subspace topology, respectively. Note that for κ to be well defined, the equivalence relation $\sim$ must satisfy

$$(g,p) \sim (h,q) \implies gp = hq$$

for any two points $(g,p),(h,q) \in \Gamma \times P$.

We define the equivalence relation by explicitly stating a partition of $\Gamma_P \times P$. For $(g,p) \in \Gamma_P \times P$, define the set

$$[g,p] := \{ (gh,p) \mid h \in \Gamma_P(p) \} = g\Gamma_P(p) \times \{ p \}$$

It is not hard to verify that

$$\tilde{X} := \{ [g,p] \mid (g,p) \in \Gamma_P \times P \}$$

is a partition of $\Gamma_P \times P$. Clearly, for any element $(h,q) \in [g,p]$, it holds that $gp = hq$: we have $p = q$ and $hq = gjp = gp$ for some $j \in \Gamma_P(p)$. Therefore, κ is well defined.

If κ were injective, we could proceed as follows to conclude that Condition (13) holds: Suppose $x \in X$, $g \in \Gamma_P$, $y \in gP \cap B_{\epsilon(x)}(x)$ and $p \in P$ such that $y = gp$. We have to show that $g \in \Gamma_P(x)$. By part 1, we know that $gp = hq$ for some $h \in \Gamma_P(x)$ and $q \in P \cap B_\epsilon(x)$. Since κ is assumed to be injective, it holds that $[g,p] = [h,q]$, which is equivalent to

$$(g,p) \in [h,q] = h\Gamma_P(q) \times \{ q \}\,.$$

We claim that $\Gamma_P(q) \subset \Gamma_P(x)$. To see this, note that any side S of P which contains q intersects $B_\epsilon(x)$, and this ball has the property that it only intersects sides of P which contain x. An immediate consequence of this is that $\mathcal{S}(q)$ is a subset of $\mathcal{S}(x)$ and $\{ g_S \mid S \in \mathcal{S}(q) \}$ is a subset of $\{ g_S \mid S \in \mathcal{S}(x) \}$. The claim follows, since the inclusion of a set of generators of a subgroup in the set of generators of another subgroup implies the corresponding inclusion of the generated subgroups. The previous statements imply

$$g \in h\Gamma_P(q) \subset h\Gamma_P(x) \subset \Gamma_P(x),$$

which is what we wanted to show.

To finish the inductive step and thereby the entire proof, it remains to show that κ is injective. We will show that κ is a covering map. More specifically, this entails checking that the domain and codomain of κ are Hausdorff and path-connected, and that every element of the codomain has an evenly covered neighborhood. This suffices, since any covering map of a simply connected space is injective.[8] To see this, let $z \in X$ and $x, y \in \kappa^{-1}(\{ z \})$, and consider a path $\gamma : [0,1] \to \tilde{X}$ from x to y. The path γ is mapped onto a loop $\kappa \circ \gamma$ in X. Since X is simply connected, the loop is homotopic to the constant loop via a homotopy with fixed endpoints. Since κ is a covering map, this homotopy lifts to a homotopy with fixed endpoints between γ and the constant path $\tau_x : [0,1] \to \tilde{X}, t \mapsto x$. The fact that the homotopy fixes the endpoints of all curves implies

$$y = \gamma(1) = \tau_x(1) = x.$$

[8]Since a covering map is both surjective and a local homeomorphism, the phrase 'injective covering map' is synonymous to the word 'homeomorphism'.

In other words, the cardinality of $\kappa^{-1}(\{ z \}) \leq 1$ for any $z \in X$, i.e. κ is injective.

We claim that $B_l(z)$ (with l as in the earlier part of the proof) is an evenly covered neighborhood of z for any $z \in X$ and that $\tilde{X}$ is a connected Hausdorff space. As a consequence of this, it follows that κ is a covering map. For any $[g,p] \in \tilde{X}$ and any $r > 0$, define the set

$$\tilde{B}([g,p],r) \coloneqq \{ [gh,q] \mid h \in \Gamma_P(p) \text{ and } q \in P \cap B_r(p) \}.$$

To prove the claim, it suffices to show

(i) $\tilde{B}([g,p],r)$ is an open neighborhood of $[g,p]$ for any $[g,p] \in \tilde{X}$ and $r > 0$,

(ii) $\kappa\big|_{\tilde{B}([g,p],r)} : \tilde{B}([g,p],r) \to B_r(gp)$ is a homeomorphism for any $[g,p] \in \tilde{X}$ and $0 < r \leq \epsilon(p)$,

(iii) $\tilde{X}$ is Hausdorff,

(iv) $\tilde{X}$ is connected,

(v) $\bigcup_{[g,p]\in\kappa^{-1}(\{ z \})} \tilde{B}([g,p],l) = \kappa^{-1}(B_l(z))$ for all $z \in X$, and

(vi) if $z \in P$ and $x, y \in \kappa^{-1}(\{ z \})$ with $x \neq y$, then the sets $\tilde{B}(x,l)$ and $\tilde{B}(y,l)$ are disjoint.

Point (i). *The set $\tilde{B}([g,p],r)$ is an open neighborhood of $[g,p]$ for any $[g,p] \in \tilde{X}$ and $r > 0$.*
Let $[g,p] \in \tilde{X}$ and $r > 0$. In order to show that $\tilde{B}([g,p],r)$ is open in $\tilde{X}$, we need to show that the preimage of $\tilde{B}([g,p],r)$ under the quotient map $[\cdot] : \Gamma_P \times P \to \tilde{X}$ is open. Note that, for any $q \in B_{\epsilon(p)}(p)$ it holds that $\Gamma_P(q) \subset \Gamma_P(p)$, and, consequently, $\Gamma_P(q)\Gamma_P(p) = \Gamma_P(p)$ and

$$\{ ghj \mid h \in \Gamma_P(p) \text{ and } j \in \Gamma_P(q) \} \times \{ q \} = g\Gamma_P(p)\Gamma_P(q) \times \{ q \} = g\Gamma_P(p) \times \{ q \}.$$

Therefore, the preimage under the projection is given by

$$\begin{aligned}[\cdot]^{-1}\big(\tilde{B}([g,p],r)\big) &= \{ (ghj,q) \mid h \in \Gamma_P(p), q \in P \cap B_r(p), \text{ and } j \in \Gamma_P(q) \} \\ &= g\Gamma_P(p) \times (P \cap B_r(p)).\end{aligned}$$

This description of $[\cdot]^{-1}\big(\tilde{B}([g,p])\big)$ allows us to conclude that it is an open subset of $\Gamma_P \times P$, since it is given as a product of open subsets of Γ_P and P, respectively.

Point (ii). *The map $\kappa\big|_{\tilde{B}([g,p],r)} : \tilde{B}([g,p],r) \to B_r(gp)$ is a homeomorphism for any $[g,p] \in \tilde{X}$ and $0 < r \leq \epsilon(p)$.*
We begin by verifying the continuity of the map $m : \Gamma_P \times P \to X, (g,p) \mapsto gp$. For any $z \in X$ and $r > 0$

$$m^{-1}(B_r(z)) = \bigcup_{g\in\Gamma} \{ g \} \times \big(P \cap g^{-1}(B_r(z))\big)$$

is the union of open subsets and therefore itself open. By the universal property of the quotient topology, the continuity of m implies the continuity of κ, since $m = \kappa \circ [\cdot]$.

Next, we check that the restrictions are well defined and surjective. For any $(g,p) \in \Gamma_P \times P$ and $0 < r < \epsilon(p)$, part one of this proof describes a certain tessellation of $B_r(p)$. In particular, we have seen that the sets $hP \cap B_r(p)$, with $h \in \Gamma_P(p)$, cover $B_r(p)$. Therefore, it holds that

$$\begin{aligned}B_r(gp) = gB_r(p) \quad &= g \bigcup_{h\in\Gamma_P(p)} hP \cap B_r(p) &&= \bigcup_{j\in g\Gamma_P(p)} j\,(P \cap B_r(p)) \\ = m\big(g\Gamma_P(p) \times (P \cap B_r(p))\big) \quad &= \kappa\big(\,[g\Gamma_P(p) \times (P \cap B_r(p))]\,\big) &&= \kappa(\tilde{B}([g,p],r)).\end{aligned}$$

This tells us that κ maps $\tilde{B}([g,p],r)$ onto $B_r(gp)$, so that the restrictions defined in Point (ii) are well defined and surjective.

It still needs to be proven that $\kappa\big|_{\tilde{B}([g,p],r)}$ is injective. The idea is to reduce the injectivity of κ to a weakened 'local' version of Property (13). This version states that

$$(q_1, q_2 \in P \cap B_r(p),\ h \in \Gamma_P(p), \text{ and } q_1 = hq_2) \implies h \in \Gamma_P(q_1). \tag{14}$$

First, let us take a closer look at any two elements of $\tilde{X}$ which are mapped onto the same point by κ. Let $(h_1, q_1), (h_2, q_2) \in \Gamma_P(p) \times (P \cap B_r(p))$ and assume

$$\kappa([gh_1, q_1]) = gh_1q_1 = gh_2q_2 = \kappa([gh_2, q_2]).$$

We want to show $[h_1, q_1] = [h_2, q_2]$, which is equivalent to $q_1 = q_2$ and $h_1^{-1}h_2 \in \Gamma_P(q_1)$. Obviously, it holds that $q_1 = h_1^{-1}h_2q_2$ and $h_1^{-1}h_2 \in \Gamma_P(p)$. So we have two points in P which are close to eachother and a map in $\Gamma_P(p)$ which maps one point onto the other. We claim that Equation (14) is a sufficient condition for the injectivity of $\kappa\big|_{\tilde{B}([g,p],r)}$. To see this, note that the equation implies $h_1^{-1}h_2 \in \Gamma_P(q_1)$. Therefore, both the map itself and its inverse fix the point q_1. Consequently, it holds that

$$q_1 = (h_1^{-1}h_2)^{-1}q_1 = (h_1^{-1}h_2)^{-1}h_1^{-1}h_2q_2 = q_2.$$

This proves the injectivity of $\kappa\big|_{\tilde{B}([g,p],r)}$.

We shall presently verify Equation (14). The main ingredient are yet again the tessellations of small balls constructed in part one of the proof. Let q_1, q_2 and h be as in equation (14). Let $s > 0$ such that $B_s(q_1) \subset B_r(p)$. Then the collection

$$\{\, jP \cap B_s(q_1) \mid j \in \Gamma_P(q_1) \,\}$$

covers the ball $B_s(q_1)$. First, note that $h\mathring{P} \cap B_s(q_1)$ is non-empty. This follows from Lemma 5.8, as hP is a non-empty convex set which contains q_1. Second, observe that the sets $j\partial P \cap B_s(q_1)$, with $j \in \Gamma_P(q_1)$, do not cover any open subset of X. The latter statement holds, because it is a finite collection of $(n-1)$-dimensional sets, and such a collection can not cover an n-dimensional set. These observations allow us deduce the existence of a $k \in \Gamma_P(q_1)$ such that the intersection

$$\big(h\mathring{P} \cap B_s(q_1)\big) \cap \big(k\mathring{P} \cap B_s(q_1)\big)$$

is non-empty. It follows that the superset

$$\big(h\mathring{P} \cap B_r(p)\big) \cap \big(k\mathring{P} \cap B_r(p)\big)$$

is non-empty as well. This fact allows us to conclude that $h = k$ by noting that $h, k \in \Gamma_P(p)$ and

$$\{\, jP \cap B_r(p) \mid j \in \Gamma_P(p) \,\}$$

is a tessellation of $B_r(p)$. Thus, h is an element of $\Gamma_P(q_1)$.

Finally, we turn our attention to the continuity of the $\kappa\big|_{\tilde{B}([g,p],r)}{}^{-1}$. Let $U \subset \tilde{B}([g,p],r)$ be an open subset and $[j,x] \in U$. The preimage of U under the map in question is given by the set $\kappa(U)$. To show that $\kappa(U)$ is open in X, it suffices to prove the existence of a $s > 0$ with $B_s(jx) \subset \kappa(U)$. By definition of the quotient topology, the set $[\cdot]^{-1}(U)$ is open in $\Gamma_P \times P$ and contains $[\cdot]^{-1}([j,x]) = j\Gamma_P(x) \times \{\, x \,\}$. By definition of the product topology, there exists a family of open neighborhoods $(V_h)_{h \in \Gamma_P(x)}$ of x in P with the property

$$\bigcup_{h \in \Gamma_P(x)} \{\, jh \,\} \times V_h \subset [\cdot]^{-1}(U).$$

The crucial property we exploit in the next step is the finite cardinality of the group $\Gamma_P(x)$. It implies that the set $\bigcap_{h\in\Gamma_P(x)} V_h$ is open in P, and, consequently, contains the set $P \cap B_s(x)$ for some $0 < s < \epsilon(jx)$. This, in turn, allows us to conclude

$$\begin{aligned} U = [\cdot]\big([\cdot]^{-1}(U)\big) \qquad\qquad & \supset [\cdot]\Big(\bigcup_{h\in\Gamma_P(x)} \{\, jh \,\} \times V_h\Big) \\ \supset \{\, [jh, q] \mid h \in \Gamma_P(x) \text{ and } q \in P \cap B_s(x) \,\} \qquad & = \tilde{B}([j, x], s) \end{aligned}$$

and, finally,

$$\kappa(U) \supset \kappa(\tilde{B}([j, x], s)) = B_s(jx).$$

Point (iii). *The space $\tilde{X}$ is Hausdorff.*
Let $[g, p], [h, q] \in \tilde{X}$ with $[g, p] \neq [h, q]$. Since $[g, p] = [h, q]$ if and only if $p = q$ and $g^{-1}h \in \Gamma_P(q)$, our assumption is equivalent to the statement '$q \neq p$ or $g^{-1}h \notin \Gamma_P(q)$'. We claim that there exist $r_1, r_2 > 0$ such that $\tilde{B}([g, p], r_1)$ does not intersect $\tilde{B}([h, q], r_2)$. Since the quotient map $[\cdot]$ is surjective, it suffices to show that the preimage of the intersection is empty. Using the description of the preimages from earlier, we see that

$$\begin{aligned} [\cdot]^{-1}\left(\tilde{B}([g, p], r_1) \cap \tilde{B}([h, q], r_2)\right) &= \big(g\Gamma_P(p) \times (P \cap B_{r_1}(p))\big) \cap \big(h\Gamma_P(q) \times (P \cap B_{r_2}(q))\big) \\ &= \big(h\Gamma_P(q) \cap g\Gamma_P(p)\big) \times \big(P \cap B_{r_1}(p) \cap B_{r_2}(q)\big). \end{aligned}$$

We immediately notice that, in the case $q \neq p$, the fact that we can separate points in X by choosing sufficiently small radii r_1, r_2 implies that the set above is empty. In the case $q = p$ and $g^{-1}h \notin \Gamma_P(q)$, the two left cosets $g\Gamma_P(p)$, $h\Gamma_P(p)$ are distinct, and we conclude that the set above is empty for arbitrary $r_1, r_2 > 0$.

Point (iv). *The space $\tilde{X}$ is connected.*
Let $U, V \subset \tilde{X}$ be open, disjoint subsets which cover $\tilde{X}$. The goal is to prove that one of the two subsets is the whole space and the other is empty. We are working explicitly with the definition of connectedness. First, we show that

$$\{\, [g, x] \mid x \in P \,\}$$

is contained either in U or V for any $g \in \Gamma_P$. The sets

$$(\{\, g \,\} \times P) \cap [\cdot]^{-1}(U) \text{ and } (\{\, g \,\} \times P) \cap [\cdot]^{-1}(V),$$

are disjoint open subsets of $\{\, g \,\} \times P$ which cover $\{\, g \,\} \times P$. Since the set P is connected[9] and homeomorphic to $\{\, g \,\} \times P$, one of the sets in the above equation must be empty. Therefore, $\{\, [g, x] \mid x \in P \,\}$ is contained in either U or V.

Without loss of generality, we can assume that $\{\, [1, x] \mid x \in P \,\}$ is contained in U. We now show that $U = \tilde{X}$. For this purpose, let $[g, x] \in \tilde{X}$, let $S_1, \ldots, S_k$ be a finite sequence of sides of P with $g = \prod_{i=1}^{k} g_{S_i}$, and denote the identity map on X by g_0. We show that $W_j := \{\, [\prod_{i=0}^{j} g_i, x] \mid x \in P \,\}$ is contained in U for $j \in \{\, 0, \ldots, k \,\}$ by induction. We already discussed the case $j = 0$. Let $j \in \{\, 1, \ldots, k \,\}$ and assume that W_{j-1} is contained in U. For any $y \in S_j$ we know that $g_{S_j} \in \Gamma_P(y)$. Therefore, $(\prod_{i=1}^{j} g_{S_i})^{-1} \prod_{i=1}^{j-1} g_{S_i} = g_{S_j} \in \Gamma_P(y)$. In other words, $[\prod_{i=1}^{j} g_{S_i}, y] = [\prod_{i=1}^{j-1} g_{S_i}, y]$. Recall that this point lies in U by the induction hypothesis. Since one element of the set W_j is contained in U, and the set is contained in one of the disjoint sets U or V, it follows that W_j is a subset of U. This concludes the induction, and thereby proves that $[g, x] \in W_k \subset U$. Since $[g, x]$ was an arbitrary element of $\tilde{X}$, we have shown the equality of $\tilde{X}$ and U.

For the proof of Points (v) and (vi), we need the uniqueness of lifts of any short geodesic given a choice of lift of its starting point. First, we discuss the existence. If $\gamma : [a, b] \to X$ is a geodesic of length less

[9] All polyhedra P are connected with the sole exception being the polyhedra consisting of two antipodal points in $\mathbb{S}^n$. This exceptional case cannot occur in our situation, because the dimension of P is not zero.

than l, then γ is contained in $B_l(\gamma(a))$. Given a choice of $x \in \kappa^{-1}(\gamma(a))$, one lift of the geodesic is given by $\tilde{\gamma}_1 : {\kappa_{|\tilde{B}(x,l)}}^{-1} \circ \gamma$. This proves the existence of lifts. Suppose now that $\tilde{\gamma}_1 : [a, b] \to \tilde{X}$ is another lift of γ which also starts at the x. Our approach consists of showing that

$$A \coloneqq \{\, t \in [a, b] \mid \tilde{\gamma}_1(t) = \tilde{\gamma}_2(t) \,\}$$

is both open and closed. Since it is non-empty and $[a, b]$ is connected, this will imply that it is equal to $[a, b]$, which allows us to conclude $\tilde{\gamma}_1 = \tilde{\gamma}_2$. To see that A is closed, consider a convergent sequence $(s_i)_{i\in\mathbb{N}}$ in A. Note that the expression $\lim_{i\to\infty} \tilde{\gamma}_1(s_i)$ refers to a single point in $\tilde{X}$, as the Hausdorff property implies that a sequence can have at most one limit. Therefore we can meaningfully say

$$\tilde{\gamma}_1(\lim_{i\to\infty} s_i) = \lim_{i\to\infty} \tilde{\gamma}_1(s_i) = \lim_{i\to\infty} \tilde{\gamma}_2(s_i) = \tilde{\gamma}_2(\lim_{i\to\infty} s_i),$$

which is to say $\lim_{i\to\infty} s_i \in A$. To see that A is open, consider a $t_0 \in A$. Since $\tilde{\gamma}_1(t_0)$ lies in $\tilde{B}(x, l)$ by construction and the two lifts agree at t_0, the same holds for $\tilde{\gamma}_2$. By continuity of the lifts, there exists an open neighborhood I of t_0 in $[a, b]$ which is mapped into $\tilde{B}(x, l)$ by both lifts. The injectivity of $\kappa_{|\tilde{B}(x,l)}$ together with the fact that γ is a path in $B_l(\gamma(a))$ implies that any two lifts agree while they lie in the ball $\tilde{B}(x, l)$. Therefore, I is contained in A, which shows that A is open.

Point (v). *We have the equality of sets $\bigcup_{[g,p]\in\kappa^{-1}(\{\, z \,\})} \tilde{B}([g, p], l) = \kappa^{-1}(B_l(z))$ for all $z \in X$.* We have already seen the inclusion '$\subset$'. For the converse, let $x \in \kappa^{-1}(B_l(z))$. We want to find $y \in \kappa^{-1}(\{\, z \,\})$ such that x lies in $\tilde{B}(y, l)$. Let $b \coloneqq d(\kappa(x), z)$ and consider the shortest path $\gamma : [0, b] \to X$ from $\kappa(x)$ to z whose image lies in $\tilde{B}(y, l)$. Its length is certainly less than l. Therefore, it has a unique lift $\tilde{\gamma}$ starting at x. Note that $y \coloneqq \tilde{\gamma}(b) \in \kappa^{-1}(\{\, z \,\})$ and $\kappa_{|\tilde{B}(y,l)}$ is a homeomorphism onto $B_l(z)$. The lift of γ given by

$${\kappa_{|\tilde{B}(y,l)}}^{-1} \circ \gamma$$

is equal to $\tilde{\gamma}$. To see this, note that $t \mapsto \gamma(b - t)$ is a path from z to $\kappa(x)$, and

$$t \mapsto {\kappa_{|\tilde{B}(y,l)}}^{-1} \circ \gamma(b - t)$$

and

$$t \mapsto \tilde{\gamma}(b - t)$$

are both lifts of $t \mapsto \gamma(b - t)$ which start at y. The uniqueness of such lifts implies that the two lifts are the same. Since reparameterisation is a reversible operation, we conclude that

$$x = \tilde{\gamma}(0) = {\kappa_{|\tilde{B}(y,l)}}^{-1} \circ \gamma(0),$$

which certainly means x lies in the ball $\tilde{B}(y, l)$. This proves the inclusion '$\supset$'.

Point (vi). *If $z \in P$ and $x, y \in \kappa^{-1}(\{\, z \,\})$ with $x \neq y$, then the sets $\tilde{B}(x, l)$ and $\tilde{B}(y, l)$ are disjoint.* It is worth nothing that we already studied the intersection of two sets of the form $\tilde{B}(x, r)$ for $x \in \tilde{X}$ and $r > 0$ to prove that $\tilde{X}$ is Hausdorff. There, our upper bound on the radii depended on the two points in $\tilde{X}$. Here, $l > 0$ is fixed for all elements in the preimage of a point $z \in X$. Suppose the intersection is non-empty, and let $w \in \tilde{B}(x, l) \cap \tilde{B}(y, l)$. Let γ be the shortest path from $\kappa(w)$ to z which is contained in $B_l(z)$. Then the fact that the restrictions of κ to both $\tilde{B}(x, l)$ and $\tilde{B}(y, l)$ are homeomorphisms allows us to produce two lifts of γ which both start at w. These lifts are not the same, because one ends at x, the other ends at y, and we assumed x and y to be distinct. But the length of γ is less than l, so there exists only one lift starting at w. This is a contradiction. We conclude that the intersection is empty. $\square$

6 The Multiplicity $m_1(K)$ of the Smallest Positive Eigenvalue of the Laplacian on the Klein Quartic K

In this section, K denotes the Klein Quartic. We refer to the literature for a definition of K.

We will give an overview of how the various partial results concerning the multiplicity $m_1(K)$ of the smallest positive eigenvalue $\lambda_1(K)$ of the Laplacian on K are combined to arrive at the result $m_1(K) = 8$. We have discussed two ingredients in this document: the integer equation for $m_1(K)$ which stems from the representation theory argument and the lower and upper bound formulae for the number of eigenvalues that lie in an interval which stem from the Selberg trace formula. There are lots of details which we will not delve into here. We also need some results whose proofs are based on different strategies. Specifically, we need the following results, which we quote without proof:

1. If $\lambda \in (0, 7.85]$, then any decomposition of τ_λ into a sum of irreducible representations does not include a 1-dimensional representation. [10]
2. The first positive eigenvalue satisfies the inequality
$$\lambda_1(K) \leq 2(4 - \sqrt{7}) \approx 2.708497.^{11}$$
3. There is a theorem which roughly says *If S is a closed hyperbolic surface of genus $g \geq 2$, then all sufficiently small eigenvalues of the Laplacian have multiplicity less or equal $2g-1$.* In the case of the Klein quartic, the precise statement for an eigenvalue λ of the Laplacian is
$$\lambda \leq \frac{1}{4} + \left(\frac{\pi}{4\text{arcsin}(\sqrt{2}}\right) \approx 0.719512 \implies m(\lambda) \leq 5.^{12}$$

For the applications of Lemma 3.8, the existence of functions with the desired properties has to be shown. Second, some information about the closed geodesics of K and their lengths has to be obtained. We refer to the literature. The results given by applications of the lemma are statements of the kind "If λ lies in some interval then the number of eigenvalues counting multiplicity is greater equal or lesser equal some $n \in \mathbb{N}$." The authors of the paper [1] show that

4. $\lambda \in [0.71, 2.575] \implies \text{m}(\lambda) < 6$
5. $\lambda \in [2.575, 5.5] \implies \text{m}(\lambda) < 12$
6. $\lambda \in [2.575, 5.5] \implies \text{m}(\lambda) > 7$

in Propositions 5.5, 5.6, and 5.7, respectively

Before we can put together the pieces of the puzzle, we need one final statement, which is the integer equation for the multiplicity of an eigenvalue of the Laplacian derived from a decomposition of τ_λ. For this purpose, we need some information about the irreducible real representations of the isometry group of the Klein Quartic. It turns out that the character table of the isometry group, which is isomorphic to the projective general linear group $\text{PGL}(2,7)$, provides sufficient data. It is specified on the page

`https://people.maths.bris.ac.uk/~matyd/GroupNames/321/PGL(2,7).html`.

Lemma 6.1 ([1, Corollary 5.3, p.16]). *The character table of* Iso(K) *provides enough information to assess if an irreducible complex representation is realizable over* $\mathbb{R}$. *More specifically, there are 9 irreducible complex*

[10] [1, Proposition 4.4 and Corollary 5.2]
[11] [1, p. 3]
[12] [1, Theorem 1.1 and Remark on p.7]

representations $\rho_1, \dots, \rho_9$. Their (complex) dimensions are $1, 6, 7,$ and 8. They are all realizable over $\mathbb{R}$. Therefore, the (real) dimensions of the irreducible real representations of $\mathrm{Iso}(K)$ *are also $1, 6, 7,$ and 8.*

Because the direct sum decomposition of the representation $\tau_{\lambda_1(K)}$ does not include any one-dimensional representation, we obtain an integer equation

$$m_1(K) = n_1 6 + n_2 7 + n_3 8 \tag{15}$$

for some $n_1, n_2, n_3 \in \mathbb{N}_0$.

Proof. According to [10, Proposition 39, p. 109], the question of the realizability of a representation χ over $\mathbb{R}$ can be answered by looking at the sum

$$\sum_{y \in \mathrm{Iso}(K)} \chi(y^2) = \sum_{[y] \text{ conjugacy class of } \mathrm{Iso}(K)} \#[y]\chi(y^2).$$

Since the characters all have real values, the value of this sum is equal to $\#\mathrm{Iso}(K)$ if and only if χ is realizable over $\mathbb{R}$ and equal to $-\#\mathrm{Iso}(K)$ if and only if χ is not realizable over $\mathbb{R}$.

The character of a representation is a class function, i.e. the image of y depends only on its conjugacy class. (This just amounts to the statement that the trace of an endomorphism is invariant under conjugation.) If x, y are conjugate group elements, then the squares x^2, y^2 are conjugate as well. Therefore, the map $y \mapsto \chi(y^2)$ is a class function. Next, note that we can determine the order of y^2 from the order of y. Namely, if the order of y is odd, then the order of y and y^2 are equal. If the order of y is even, then the order of y^2 is half the order of y. In the case of the group $\mathrm{Iso}(K)$, this almost suffices to determine the class function $y \mapsto y^2$. Namely, we have the following

$$\begin{array}{lll} 1 \mapsto 1 & 2A \mapsto 1 & 2B \mapsto 1 \\ 3 \mapsto 3 & 4 \mapsto 2A \text{ or } 2B & 6 \mapsto 3 \\ 7 \mapsto 7 & 8A \mapsto 4 & 8B \mapsto 4, \end{array}$$

with the image of the conjugacy class 4 being the only entry which is not fully determined by the order.

We consider ρ_3 and attempt to compute the sum

$$\sum_{[y] \text{ conjugacy class of } \mathrm{Iso}(K)} \#[y]\chi_{\rho_3}(y^2).$$

Given the ambiguous information about the map $y \mapsto y^2$, we get two possible results for the sum:

$$\begin{aligned} &\sum_{[y] \in \{\text{ conjugacy classes of } \mathrm{Iso}(K)\ \}} \#[y]\chi_{\rho_3}(y^2) \\ =&1 \cdot 6 + 21 \cdot 6 + 28 \cdot 6 + 56 \cdot 0 + 42 \cdot (-2) + 56 \cdot 0 + 48 \cdot (-1) + 42 \cdot 2 + 42 \cdot 2 \\ \text{or }& 1 \cdot 6 + 21 \cdot 6 + 28 \cdot 6 + 56 \cdot 0 + 42 \cdot 0 + 56 \cdot 0 + 48 \cdot (-1) + 42 \cdot 2 + 42 \cdot 2 \\ =&336 \text{ or } 420. \end{aligned}$$

This allows us to resolve the ambiguity. Since the sum must be equal to ± 336, it must hold that $4 \mapsto 2A$. Furthermore, we can conclude that ρ_3 is realizable over $\mathbb{R}$. Next, we compute the sum for ρ_4

$$\begin{aligned} &\sum_{[y] \in \{\text{ conjugacy classes of } \mathrm{Iso}(K)\ \}} \#[y]\chi_{\rho_4}(y^2) \\ =&1 \cdot 6 + 21 \cdot 6 + 28 \cdot 6 + 56 \cdot 0 + 42 \cdot 2 + 56 \cdot 0 + 48 \cdot (-1) + 42 \cdot 0 + 42 \cdot 0 \\ =&336 \end{aligned}$$

and conclude that ρ_4 is realizable over $\mathbb{R}$. In the expression for ρ_5, the summands are the same as for ρ_4, whence it too is realisable over $\mathbb{R}$. The corresponding computations for the other six irreducible representations are omitted.

We claim that the complexification of any irreducible real representation of the isometry group of K is an irreducible complex representations. This follows from Lemma 4.8 and the fact that all the irreducible complex representations are realizable over $\mathbb{R}$. For if the complexification of some irreducible real representation were not complex irreducible, then there would exist an irreducible complex representation which is not realizable over $\mathbb{R}$. In particular, this shows that for any irreducible real representation of (real) dimension d, there exists an irreducible complex representation of (complex) dimension d - namely its complexification. Thus, the dimensions of the real irreducible representations are $1, 6, 7$, and 8. Since the decomposition of $\tau_{\lambda_1(K)}$ does not involve one-dimensional representations, any decomposition of $\tau_{\lambda_1(K)}$ into a sum of irreducible real representations implies the integer equation stated in Equation (15). $\square$

We are finally in a position where we can explain how the pieces of the puzzle collected on the last couple of pages fit together.

The results stated in Points 1 and 2, allowed us to obtain the integer equation Equation (15). In particular, this equation shows that $m_1(K)$ greater equal 6. Therefore, the upper bound on the multiplicity of small eigenvalues from Point (3) tells us that

$$\lambda_1(K) \geq \frac{1}{4} + (\frac{\pi}{4 \arcsin \sqrt{2}})^2 \approx 0.719512\ldots.$$ [13]

A first application of Lemma 3.8, Point (4) implies that the multiplicity $m(\lambda)$ is strictly less than 6 for eigenvalues $\lambda \in [0, 2.575]$. Consequently, it holds that

$$\lambda_1(K) > 2.575.$$

Taken together with the upper bound for from Point (2), we know that $\lambda_1(K) \in [2.575, 5.5]$. Now two further applications of Lemma 3.8 and its lower bound variant show that the multiplicity $m_1(K)$ of $\lambda_1(K)$ satisfies

$$7 < m_1(K) < 12.$$

The integer Equation (15) implies that

$$m_1(K) = 8.$$

This is the approach used to prove the

Theorem 6.2 ([1, Theorem 1.2.]). *The multiplicity $m_1(K)$ of the smallest positive eigenvalue $\lambda_1(K)$ of the Laplacian on the Klein Quartic is equal 8.*

[13] [1, Theorem 1.1 and Corollary 5.4]

7 Conclusion

In this thesis, we have had a look at how the multiplicity of the smallest positive Laplacian eigenvalue on the Klein Quartic was determined.

Along the way, we skipped lots of details some of which might warrant a closer look in the future. In particular, the linear programming needed for the application of the lemma derived from the Selberg trace formula could be an interesting topic.

A natural question is to what extent the methods employed in the main paper can be extended to determine the maximal multiplicity of the smallest positive eigenvalues of the Laplacian for surfaces of any genus $g \geq 2$. We also did not delve into the history of either questions or results concerning the Laplacian and its spectrum. Getting an overview of the main theorems and the open questions would be a worthwhile task.

Another obvious question is why anyone would care to know about these properties of the Laplacian's spectrum on closed hyperbolic surfaces. Specifically, one might want to understand how these properties relate to other questions of more explicitly geometric nature.

References

[1] M.F. Bourque and B. Petri. *The Klein Quartic Maximizes the Multiplicity of the First Positive Eigenvalue of the Laplacian*. 2023. arXiv: `2111.14699 [math.GT]`.

[2] G.E. Bredon. *Topology and Geometry*. Springer, 1993.

[3] P. Buser. *Geometry and Spectra of Compact Riemann Surfaces*. Birkhäuser, 2010.

[4] M.F. Do Carmo. *Riemannian Geometry*. Birkhäuser, 1992.

[5] G. Faraco. *Tessellations of surfaces*. 2023. arXiv: `2303.17263 [math.HO]`.

[6] J. M. Lee. *Introduction to Riemannian Manifolds*. Springer, 2018.

[7] J. M. Lee. *Introduction to Smooth Manifolds*. Springer, 2012.

[8] P. Petersen. *Riemannian Geometry*. Springer, 2016.

[9] J.G Ratcliffe. *Foundations of Hyperbolic manifolds*. Springer, 2006.

[10] J.-P. Serre. *Linear Representations of Finite Groups*. Springer, 1977.

[11] L.W. Tu. *Differential Geometry*. Springer, 2017.

Erklärung

Ich versichere hiermit, dass ich die vorliegende Arbeit selbständig verfasst und keine anderen als die angegebenen Hilfsmittel und Quellen benutzt habe. Sowohl inhaltlich als auch wörtlich entnommene Inhalte wurden als solche kenntlich gemacht. Die Arbeit ist in gleicher oder vergleichbarer Form noch bei keiner anderen Prüfungsbehörde eingereicht worden. Die an das Prüfungssekretariat gesandte elektronische Version der Arbeit stimmt mit der gedruckten Abgabe überein. Die Arbeit darf mit entsprechender Software auf Plagiate überprüft werden.

Heidelberg, August 14, 2023,

B. Wendland

Acknowledgements

I am grateful for Professor Pozzetti's help throughout the entire process, in particular their willingness to answer frequent questions and recommend suitable literature.