University of Zurich[UZH]

Institute of Mathematics
University of Zurich

MASTER THESIS

Platonic Surfaces and their Automorphism Groups

Author:
Colin Meckes

Supervisors:
Prof. Corinna Ulcigrai
Yuriy Tumarkin

Submitted on
May 4, 2026

Imprint

Project:	Master Thesis
Title:	Platonic Surfaces and their Automorphism Groups
Author:	Colin Meckes
Date:	May 4, 2026
Keywords:	Topology, Algebraic Topology, Platonic surfaces, Regular maps, Monodromy group, Translation surfaces
Copyright:	University of Zurich

Supervisor 1:
Prof. Corinna Ulcigrai
University of Zurich
Email: corinna.ulcigrai@math.uzh.ch
Web: Link

Supervisor 2:
Yuriy Tumarkin
University of Zurich
Email: yuriy.tumarkin@math.uzh.ch
Web: Link

Abstract

We study the automorphism group of a Platonic surface by analyzing the monodromy (or deck) group associated to the cover of a regular polygon or double polygon by the unfolded Platonic surface. To each Platonic surface, we associate its unfolding, defined as the minimal branched covering on which all holonomy becomes trivial, so that the unfolding is a translation surface. We show that a normal subgroup of the rotation group of the Platonic surface, and in some cases even the whole rotation group, is isomorphic to the monodromy group. Explicit bounds on group orders are derived in terms of the Schläfli symbol $\{p, q\}$, and several examples, including the cube, octahedron, hemicube, and hemi–octahedron, are examined in detail to illustrate the theory, where the latter ones are an extension to non-orientable Platonic surfaces.

Contents

Abstract **iii**

1 Introduction **1**

2 Preliminaries **3**
2.1 Geometry & Surfaces . . . 3
2.1.1 Surfaces and Orientability . . . 4
2.1.2 Translation Surface . . . 4
2.1.3 Platonic Surface . . . 6
2.2 Coverings . . . 9
2.2.1 Lifting Property . . . 9
2.2.2 Branched Covering Maps . . . 10
2.3 Singular Homology . . . 10
2.3.1 Exact Sequences . . . 11
2.4 Unfolding . . . 13
2.5 Groups . . . 15
2.5.1 Deck Group . . . 15
2.5.2 Rotation Subgroup . . . 16
2.5.3 Monodromy Representation . . . 17

3 On Orientable Surfaces **19**
3.1 Main Theorems . . . 19
3.2 Some Bounds . . . 24
3.3 Examples . . . 27
3.3.1 Cube . . . 27
3.3.2 Octahedron . . . 29

4 Extension to Non-Orientable Surfaces **32**
4.1 Main Idea . . . 32
4.2 Non-orientable Theorems . . . 33
4.3 Examples . . . 34
4.3.1 Hemicube . . . 34
4.3.2 Hemi-Octahedron . . . 36
4.3.3 Non-orientable Surface N4:$\{4,6\}_6$ . . . 39
4.4 Alternative Approaches . . . 41

5 Conclusion **45**

Bibliography **46**

Chapter 1

Introduction

Platonic solids are classical geometric structures that have been studied since ancient Greece. Due to their high degree of symmetry, they are important objects of study at the intersection of topology, algebra, and geometry. In mathematics, it is natural to ask whether this notion can be extended beyond regular polyhedra in three dimensional space. Thus, giving rise to the study of Platonic surfaces. Each Platonic surface is characterized by the property that its symmetry group (the automorphism group) acts transitively on flags of faces, edges, and vertices. This serves as a way to preserve its maximal symmetry. Hence, Platonic surfaces can be viewed as higher genus analogs of Platonic solids.

Platonic surfaces can be constructed via topological identifications on collections of regular Euclidean polygons. In our analysis, we will utilize results from the theory of translation surfaces. We will examine two primary group actions: the deck transformation group induced by the covering map from the unfolded Platonic surface, and the rotation group of the surface itself. This approach is based on David Aulicino's paper "A New Approach to the Automorphism Group of a Platonic Surface" [Aul19], which we will follow closely.

The concept of an unfolding originates from the field of dynamical systems, particularly in the study of rational billiards. In that context, straight-line trajectories are maintained by reflecting the polygon (the billiard table) rather than the ball itself. In our geometric setting, the goal is different: we aim to "flatten" the surface so that every cone angle becomes an integer multiple of 2π.

When considering the unfolding of a Platonic surface, a crucial distinction must be drawn between its *net* and its *unfolding*. The net is an arrangement of polygons in the plane, with prescribed identifications, that assembles into the original surface. Since the local geometry of a Platonic surface is flat almost everywhere, except in neighborhoods of its vertices, we can effectively hide these points of concentrated curvature within the corners of the net polygons. Observe that the interiors of the edges of the polygons do not see a curvature issue. Thus, they are locally isomorphic to a disk in the plane.

Our analysis will also draw on results from related fields, notably the study of regular maps, covering theory and singular homology. We will utilize two key covering maps: $\pi_D \colon \widetilde{D} \to D$ and $\pi_\Pi \colon \widetilde{D} \to \Pi_p$. Here, $\widetilde{D}$ denotes the unfolding, D the original Platonic surface, and Π_p a certain p-gon (or double p-gon) which, as we will later establish, is a translation surface. Precise definitions will be given in the next chapter.

Using these covering maps, we will prove the main theorems of this paper. These results establish an isomorphism between the monodromy group $\mathrm{Mon}(\pi_\Pi)$ and a normal subgroup of the rotation group $\mathrm{Rot}(D)$. Furthermore, under specific geometric conditions on D, we obtain the stronger isomorphism $\mathrm{Mon}(\pi_\Pi) \cong \mathrm{Rot}(D)$.

Later, we will extend the discussion to non-orientable surfaces. After establishing the theory for the orientable case, we demonstrate how it could adapt to non-orientable analogs such as the hemicube and the hemi-octahedron. We will also briefly discuss several approaches we took to formalize our theory. The main idea is to take the net of a non-orientable Platonic surface and reflect it across one of its edges, resulting in a degree-2 cover that is an orientable Platonic surface. If two edges were glued via translation in the original non-orientable surface, we preserve that identification. However, if they were glued via a reflection, we instead find an edge in the reflected copy that can be identified via translation. In this way, we obtain an orientable Platonic surface, allowing us to apply our previously developed theory.

Chapter 2

Preliminaries

This chapter provides the background and definitions required for the thesis. Section 2.1 presents basic definitions, including that of a *Platonic surface*. This is followed by a brief discussion of *coverings* in Section 2.2, *singular homology* in Section 2.3, and of *unfoldings* in Section 2.4. Finally, Section 2.5 introduces the *deck group* and the *monodromy representation*, both of which are central to this thesis.

2.1 Geometry & Surfaces

Our framework is built upon the theory of regular polyhedra. We begin with their formal definition, taken from *Hartshorne's Geometry: Euclid and Beyond*[Har00].

Definition 2.1.1. A *regular polyhedron* is a convex polyhedron whose faces are all the same regular p-gon, and at each vertex, the same number of faces meet.

The only regular polyhedra in $\mathbb{R}^3$ are the five Platonic solids, namely the *tetrahedron, cube, octahedron, dodecahedron and icosahedron*, see Figure 2.1.

We will be interested in *Platonic surfaces*, which are natural higher-genus generalizations of the Platonic solids. A formal definition will be given later.

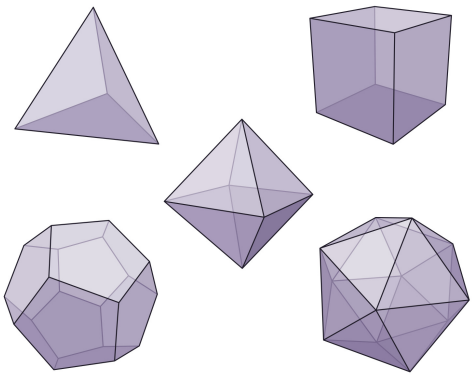

FIGURE 2.1: The five Platonic solids [Dru19]

2.1.1 Surfaces and Orientability

A topological surface is a 2-dimensional topological manifold. Such a surface is called *orientable* if it admits a consistent choice of local orientation at every point. Intuitively, this means that the surface has two distinct sides. A non-orientable surface, on the other hand, has only one side. It is an asymmetric object where a closed path can return to its starting point in a mirrored state. The most common example is the Möbius strip, where a path along its centerline returns to the origin, but with the local orientation reversed. It is a standard result that a compact manifold M is non-orientable if and only if it contains a Möbius strip.

Many surfaces can be constructed from a simple polygon by identifying its edges in pairs. A classic example is the torus, which is represented by a square whose opposite sides are identified, see Figure 2.2. Furthermore, observe that extracting a strip from the Klein bottle's square representation (again, see Figure 2.2) yields a Möbius strip, thus reinforcing the aforementioned result.

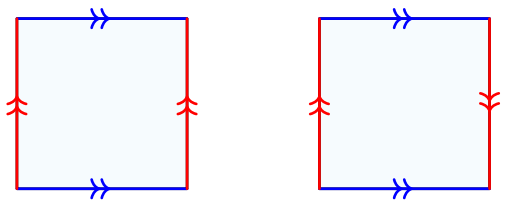

FIGURE 2.2: The tours and the Klein bottle

2.1.2 Translation Surface

For this section, we will refer to [Via06] and [Jay24].

Our first object of interest that we will introduce is the *translation surface*. A *translation surface* S is a topological surface, that can be obtained by the following construction:

Consider a polygon in $\mathbb{R}^2$ with an even number $2d \geq 4$ of sides

$$s_1, ..., s_d, s'_1, ...s'_d$$

such that for every $i = 1, ..., d$, the sides s_i and s'_i are parallel, non-adjacent, of the same length, and pointing in the same direction. See Figure 2.3 for an example with $d = 4$.

We construct a surface by orienting each side so that the Euclidean translation identifying a side with its partner is orientation-preserving. For each pairing, the polygon on one side of an oriented edge must lie to the left, and the polygon on the other side must lie to the right. The resulting structure is a closed orientable surface, called *the translation surface* S.

For a given vertex, consider a point on an adjacent side and trace a small closed loop around the vertex by moving along the side identifications. The total angle traversed by this loop is the *cone angle* at the corresponding vertex (or *cone singularity*). All cone angles occurring on a translation surface are integer multiples of 2π. This will be demonstrated in the following examples.

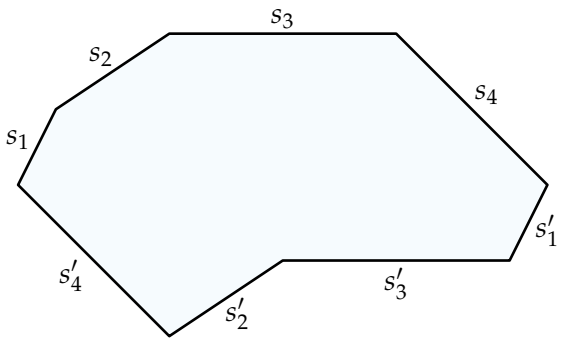

FIGURE 2.3: Translation surface with $d = 4$

Consider Π_4, the torus (also called the flat torus when represented as a translation surface). Its polygonal representation has two sides, with the interior lying both to the left and right of the identified sides. In contrast, the Klein bottle in Figure 2.2 fails to meet this criterion and is therefore not a translation surface. Returning to the torus, all four vertices collapse to a single point. The angles sum to 2π, which equals the cone angle around the resulting singularity; see Figure 2.4.

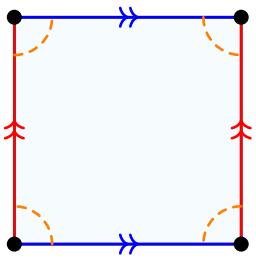

FIGURE 2.4: The cone angle of Π_4

Another example is an L-shaped translation surface, shown in Figure 2.5. Gluing the parallel, similarly colored sides by translation gives a genus-2 surface. As before, the vertices are identified to a single point, but the cone angle is 6π here.

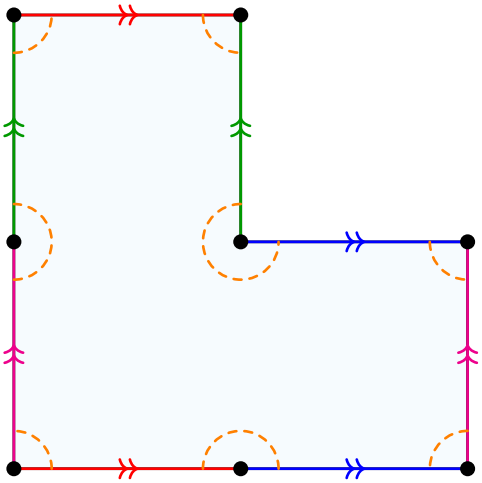

FIGURE 2.5: The cone angle of the L-shaped surface

2.1.3 Platonic Surface

We begin by defining the automorphism group and the notion of an isometry. Next we introduce the concept of a flag. All three definitions are required for the subsequent definition of a Platonic surface.

Definition 2.1.2. Let (S, g) be a surface equipped with a Riemannian metric. An *isometry* of S is a bijection $f : S \to S$ such that for every piecewise smooth curve γ on S,

$$\text{Length}(\gamma) = \text{Length}(f \circ \gamma),$$

where the length is measured with respect to the metric g.

Definition 2.1.3. Let D be a closed, orientable surface. The *automorphism group* of D, denoted $\text{Aut}(D)$, is the group of all isometries $f : D \to D$, with group operation given by composition.

Remark 2.1.4. We can also denote the group of all isometries by or $\text{Isom}(D)$, which we will use interchangeably with the notation $\text{Aut}(D)$.

Definition 2.1.5. A *flag* is a triple (f, e, v) of a p-gon, consisting of:

- A vertex v
- An edge e incident to v
- A face f incident to e

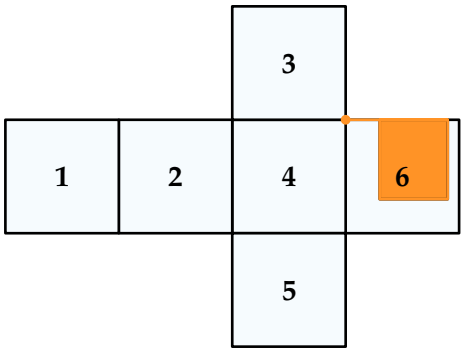

FIGURE 2.6: A flag on the cube net

Note that Figure 2.6 is somewhat misleading, as the whole face and the whole edge should be colored orange, rather than only portions of them. This is done solely to illustrate where the name *flag* comes from.

Definition 2.1.6. Let D be a closed, orientable surface formed by identifying the edges of a collection of congruent, regular p-gons. If the automorphism group of D acts transitively on its flags, then D is called a *Platonic surface*.

Let us examine this statement more closely. Flag-transitivity means that for any two flags (f, e, v) and (f', e', v'), there is an automorphism (a symmetry) that maps one flag to the other. This is the strongest form of symmetry for a tiling.

This makes it easy to depict surfaces of higher genus with "strange" curvatures, by

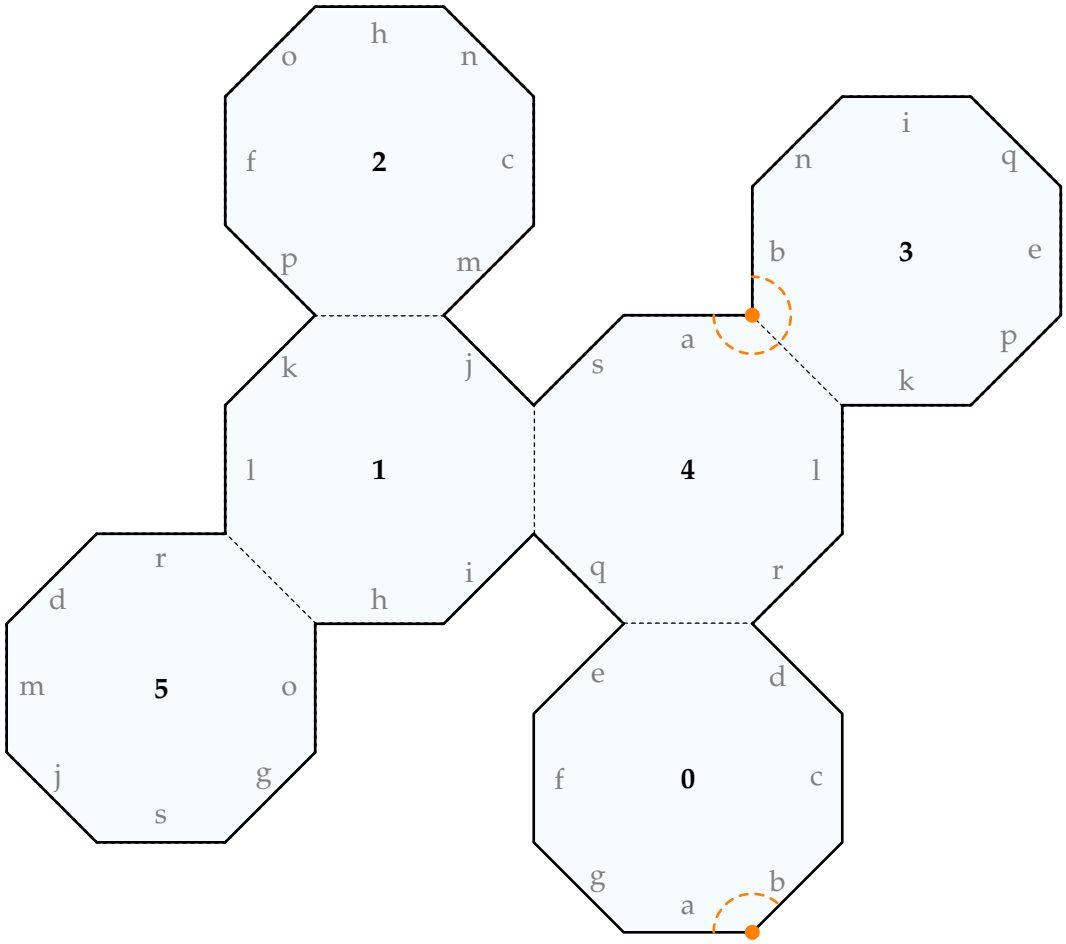

FIGURE 2.7: The Bolza surface

hiding the points with angle greater than 2π in the corners of the presentation. Consider Figure 2.7, the Bolza surface. Notice that if we look at the vertex at the intersection of the edges a and b, highlighted by the orange dot, we can figure out the angle around said point by taking any small closed curve around the point and then computing the cone angle. In this example, it should yield $\frac{3}{2}\pi + \frac{3}{4}\pi = \frac{9}{4}\pi$. It should also be noted that the Bolza surface cannot be smoothly embedded in 3-dimensions, so this presentation is very useful.

Remark 2.1.7. Later in the paper, we will frequently work with pairs (f, e) consisting of a face and an incident edge, rather than full flags. It suffices to consider only such a pair, since, by symmetry, the vertex does not provide additional information.

Definition 2.1.8. The *Schläfli symbol* of a Platonic surface is a notation $\{p, q\}$, where each face has p edges and each vertex has degree q.

For example, each face of the cube has $p = 4$ edges and every vertex is incident with $q = 3$ faces. So, the cube has Schläfli symbol $\{4, 3\}$. See table 2.1 for completeness, with regard to the Platonic solids.

Definition 2.1.9. A *regular p-gon-tiled surface* is a closed, orientable surface constructed by identifying the edges of a collection of regular Euclidean p-gons. The identification must satisfy two conditions:

Object	Schläfli symbol
Tetrahedron	$\{3,3\}$
Cube	$\{4,3\}$
Octahedron	$\{3,4\}$
Dodecahedron	$\{5,3\}$
Icosahedron	$\{3,5\}$

TABLE 2.1: The Schläfli symbols of the Platonic solids

- the resulting surface must have no boundary, and
- the total angle around every point must be an integral multiple of 2π.

The resulting surface then is a translation surface that we can endow with a choice of horizontal direction, since we have flat metric on translation surfaces.

The regular p-gon-tiled surfaces considered in this thesis are, up to a choice of horizontal direction, examples of Platonic surfaces. The only deviation from the definition of a Platonic surface is the specification of the horizontal direction, thus forgetting this additional structure yields a genuine Platonic surface.

Remark 2.1.10. We also have to make distinctions depending on the parity of p.

- if p is even, a single regular p-gon possesses opposite parallel sides, which makes the identification of the sides clear and natural. This leads to an example of a translation surface denoted by Π_p.
- if p is odd, a single polygon lacks this property, therefore we will construct Π_p from two regular p-gons, identifying each side with its unique opposite parallel side on the other polygon.

In both constructions, the result is a translation surface and the notation Π_p captures both cases, see Figure 2.8.

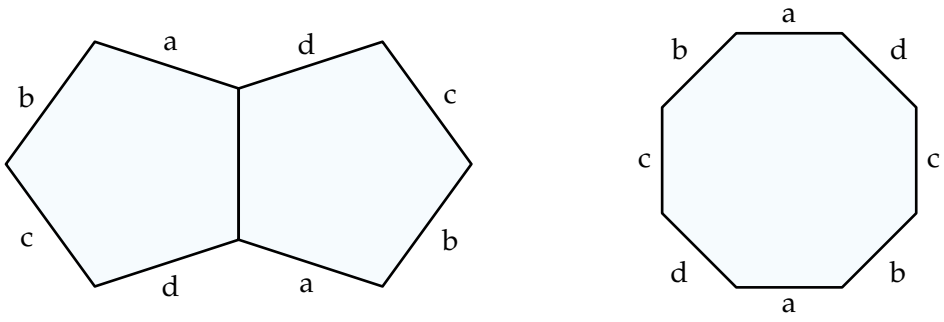

FIGURE 2.8: Examples for Π_5 and Π_8

2.2 Coverings

We will refer to [Hat01] for most of this section.

Definition 2.2.1. A space E is a *covering space* of a space X if there exists a *covering map* $p : E \to X$ with the following property: every point $x \in X$ has an open neighborhood U that is *evenly covered*. This means $p^{-1}(U)$ is a union of disjoint open sets in E, called *sheets*, each of which is mapped homeomorphically onto U by p.

Definition 2.2.2. The *skeleton* of a polyhedral tiling is the graph formed by its vertices and edges, ignoring the faces.

For example, the skeleton of a cube is the graph formed by its 8 vertices and 12 edges.

Definition 2.2.3. The skeleton of a Platonic surface is called a *regular map*.

An equivalent definition for the notion of a *regular map* stems from topological graph theory. There, a *regular map* is a cellular embedding of a connected graph on a closed, compact surface that is maximally symmetric in the following three ways:

1. *Vertex-transitive:* For any two vertices v_1, v_2, there exists an automorphism of the embedding sending v_1 to v_2.

2. *Edge-transitive:* For any two edges e_1, e_2, there exists an automorphism sending e_1 to e_2.

3. *Face-transitive:* For any two faces f_1, f_2, there exists an automorphism sending f_1 to f_2.

Moreover, the automorphism group acts freely on flags. This additional condition makes the map regular in the strongest sense.

2.2.1 Lifting Property

Definition 2.2.4. Let $p : \widetilde{X} \to X$ be a covering map. Given a continuous map $f : Y \to X$, a *lift* of f is a continuous map $\tilde{f} : Y \to \widetilde{X}$ such that:

$$p \circ \tilde{f} = f$$

That is, the following diagram commutes:

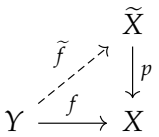

The following Propositions are well-known and their proofs can be found in *Allen Hatcher's "Algebraic Topology"* [Hat01].

Proposition 2.2.5 (Homotopy Lifting Property)**.** *Let $p : \widetilde{X} \to X$ be a covering map. Given:*

- *a homotopy* $f_t : Y \to X$,
- *a lift* $\tilde{f}_0 : Y \to \widetilde{X}$ *of the initial map* f_0,

then there exists a unique homotopy $\tilde{f}_0 : Y \to \widetilde{X}$ *of* $\tilde{f}_0$ *that lifts* f_t.

Remark 2.2.6. Taking Y to be a point yields the *path lifting property* for a covering space $p : \widetilde{X} \to X$. This states that for any path $f : I \to X$ and any lift $\tilde{x}_0$ of the starting point x_0, there exists a unique path $\tilde{f} : I \to \widetilde{X}$ lifting f with $\tilde{x}_0$.

A natural question that arises is: when can we lift a map, and is such a lift always unique? The following proposition addresses the existence, while Proposition 2.2.8 addresses its uniqueness.

Proposition 2.2.7 (Lifting Criterion)**.** *Given a covering space* $p : (\widetilde{X}, \tilde{x}_0) \to (X, x_0)$, *a map* $f : (Y, y_0) \to (X, x_0)$, *where* Y *is path-connected and locally path-connected, we denote by* $f_* : \pi_1(Y, y_0) \to \pi_1(X, x_0)$ *and* $p_* : \pi_1(\widetilde{X}, \tilde{x}_0) \to \pi_1(X, x_0)$ *the homomorphisms induced on the fundamental groups. Then, the lift* $\tilde{f} : (Y, y_0) \to (\widetilde{X}, \tilde{x}_0)$ *exists if and only if*

$$f_*(\pi_1(Y, y_0)) \subset p_*(\pi_1(\widetilde{X}, \tilde{x}_0)).$$

This can be illustrated with the following diagram:

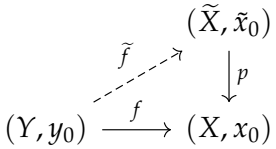

Proposition 2.2.8 (Unique Lifting Property)**.** *Let* $p : \widetilde{X} \to X$ *be a covering map. If* Y *is connected and* $\tilde{f}_1, \tilde{f}_2 : Y \to \widetilde{X}$ *are two lifts of the same map* $f : Y \to X$ *(i.e.,* $p \circ \tilde{f}_1 = p \circ \tilde{f}_2 = f$*) that agree at one point of* Y, *then* $\tilde{f}_1 = \tilde{f}_2$ *on all of* Y.

2.2.2 Branched Covering Maps

Definition 2.2.9. A map $p : Y \to X$ is said to be a *branched covering* of X if there exist finite subsets $S \subset Y, \Sigma \subset X$ such that $p|_{Y \setminus S} : Y \setminus S \to X \setminus \Sigma$ is a covering of $X \setminus \Sigma$.

In essence, a branched covering is a covering outside of a finite set of points [Mar22]. For our purposes, the set Σ is typically the set of vertices with an angle unequal to 2π, which we call *cone singularities*.

2.3 Singular Homology

This paper uses only basic notions from algebraic topology that can be found in any algebraic topology textbook, e.g. *Allen Hatcher's Algebraic Topology* [Hat01]. For a topological space X, let $H_1(X; \mathbb{Z})$ denote its *first singular homology group* with integer coefficients. Since we are dealing with surfaces, we will only define the first singular homology. Informally, $H_1(X)$ is an abelian group that "counts" one-dimensional

holes in X. For instance, $H_1(S^1) \cong \mathbb{Z}$ and $H_1(\text{point}) = 0$. To give a proper definition, we must first introduce *cycles* and *boundaries*.

Definition 2.3.1. A *singular* 1*-cycle* in a topological space X is a finite formal sum of continuous maps

$$\sigma_i \colon S^1 \to X,$$

called *closed curves*. The group of all singular 1-cycles in X is denoted by $Z_1(X)$.

Definition 2.3.2. A singular 1-cycle is called a *boundary* if it arises as the boundary of a singular 2-chain, that is, if it bounds a continuous map of a disk

$$D^2 \to X.$$

The subgroup of $Z_1(X)$ consisting of all boundaries is denoted by $B_1(X)$.

Definition 2.3.3. The *first singular homology group* of X is then defined as the quotient

$$H_1(X;\mathbb{Z}) := Z_1(X)/B_1(X).$$

Thus, two closed curves define the same homology class if their difference bounds a surface in X. Intuitively, $H_1(X)$ measures closed curves in X up to the relation of bounding a disk.

If X is path-connected, the first singular homology group admits an alternative description in terms of the fundamental group. There is a natural isomorphism

$$H_1(X;\mathbb{Z}) \cong \pi_1(X)^{\text{ab}},$$

where $\pi_1(X)^{\text{ab}} = \pi_1(X)/[\pi_1(X), \pi_1(X)]$ denotes the abelianization of the fundamental group. Under this identification, passing from $\pi_1(X)$ to $H_1(X)$ corresponds to imposing commutativity of loops. Thus, we are forgetting non-abelian homotopical information while retaining all one-dimensional homological information.

Every continuous map $f \colon X \to Y$ induces a homomorphism $f_* \colon H_1(X) \to H_1(Y)$, defined by sending a (homology class of a) cycle σ to the (class of the) cycle $f \circ \sigma$. We will call f_* the *induced map of f on homological level*. The construction is functorial: $(\text{id}_X)_* = \text{id}_{H_1(X)}$ and $(g \circ f)_* = g_* \circ f_*$ whenever the compositions make sense.

For the purposes of this paper, the following facts suffice:

- If $f, g \colon X \to Y$ are homotopic, then $f_* = g_*$ on H_1.
- If X is contractible, then $H_1(X) = 0$.
- If $f \colon X \to Y$ is a homotopy equivalence, then $f_* \colon H_1(X) \to H_1(Y)$ is an isomorphism.

2.3.1 Exact Sequences

From this section we will only need the notion of a *short exact sequence*. To reach that point, we must first define exactness, introduce the concept of a *long exact sequence*, and introduce the boundary operator.

Definition 2.3.4. A sequence of groups and group homomorphisms

$$... \longrightarrow A_{n+1} \xrightarrow{f_{n+1}} A_n \xrightarrow{f_n} A_{n-1} \xrightarrow{f_{n-1}} ...$$

is said to be *exact* if $\text{Ker}(f_n) = \text{Im}(f_{n+1})$ for each n.

To understand the definition, it is helpful to consider simple finite sequences that begin or end with the trivial group. Following conventional additive notation, the trivial group is denoted by 0.

Consider the following sequences:

- $0 \xrightarrow{0} A \xrightarrow{f} B$. The image of the zero map is zero. Therefore, the sequence is exact if and only if $\text{Ker}(f) = 0$ and that is if and only if f is injective.
- $A \xrightarrow{f} B \xrightarrow{g} 0$. The kernel of g is B. Therefore, the sequence is exact if and only if $\text{Im}(f) = B$ and that is if and only if f is surjective.
- $0 \to A \xrightarrow{f} B \xrightarrow{g} 0$. This sequence is exact if and only if f is an isomorphism, by the previous two statements.

Definition 2.3.5. *Short exact sequences* are exact sequences of the form

$$0 \to A \xrightarrow{f} B \xrightarrow{g} C \to 0.$$

As shown previously, in any short exact sequence, f is injective, g is surjective, and $\text{Im}(f) = \text{Ker}(g)$.

One can think of A as a subobject of B (using the inclusion f) and of C as the quotient object B/A. The map g then yields the isomorphism

$$C \cong B/\,\text{Im}(f) = B/\,\text{Ker}(g).$$

Definition 2.3.6. A *long exact sequence* is a sequence of Abelian groups and homomorphisms

$$\cdots \longrightarrow A_{n+1} \xrightarrow{d_{n+1}} A_n \xrightarrow{d_n} A_{n-1} \longrightarrow \cdots$$

such that the sequence is *exact* at every term, i.e.

$$\text{im}(d_{n+1}) = \text{ker}(d_n) \quad \text{for all } n.$$

Definition 2.3.7. Given a pair (X, A) with $A \subset X$, there is a homomorphism

$$\partial\colon H_1(X, A) \to H_0(A),$$

called the *boundary operator*, which measures how a relative 1-cycle in X fails to be a cycle in A.

In general, the boundary operator is defined for all homology groups H_n, not just H_1. However, since we have only defined, and only require H_1, this definition is sufficient for our purposes.

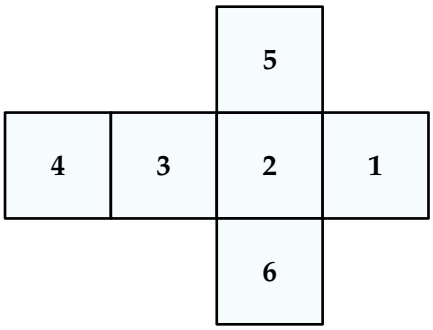

FIGURE 2.9: A cube net with labeled squares

Definition 2.3.8. For a pair (X, A) of topological spaces with $A \subset X$, there is a *long exact sequence* in singular homology. In particular, the segment

$$H_1(A) \xrightarrow{i_*} H_1(X) \xrightarrow{j_*} H_1(X, A) \xrightarrow{\partial} H_0(A) \xrightarrow{i_*} H_0(X)$$

is exact. If A is path-connected and the inclusion $A \hookrightarrow X$ induces an injective map on H_0, this yields a short exact sequence

$$0 \longrightarrow H_1(A) \xrightarrow{i_*} H_1(X) \xrightarrow{j_*} H_1(X, A) \longrightarrow 0.$$

2.4 Unfolding

The general notion of an *unfolding* comes from the study of billiards. There, when the ball hits a side, the entire space is mirrored across that side. From the ball's perspective, it continues moving in a straight line, but it is now in a reflected copy of the original space.

We also employ the term *unfolding*, though with a slightly different meaning. Here we "flatten" the surface by cutting along the edges of the regular p-gons, making the entire structure planar. In our context, *unfoldings* are simply branched coverings in which the cone angles at the singularities are forced to be integer multiples of 2π. We will give a detailed definition later on. Before we get there, we need to make an important distinction from an *unfolding* to a *net*.

Definition 2.4.1. A *net* of a Platonic surface is a collection of edge-connected polygons arranged in the plane without overlaps, which can be folded along the edges to form the Platonic surface.

One can think of this as "cutting up" a Platonic surface along some of the edges. Notice that the sequence in which we cut up our Platonic surface does not matter, since the surface is symmetric and the p-gons are arranged identically. Starting with any one will break the local symmetry in the same way. The resulting planar arrangement of the p-gons is therefore independent of these choices, up to isometry. For an example, see Figure 2.9 and try to convince yourself of the equivalences of different *cube nets*.

In order to give a rigorous definition of the *unfolding of a Platonic surface*, we first need to define what a *holonomy* is.

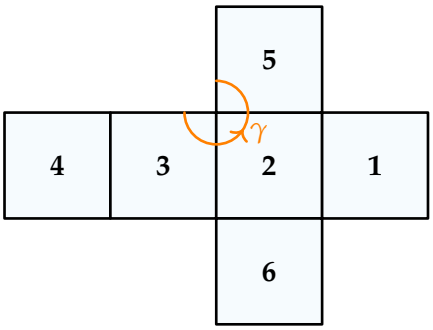

FIGURE 2.10: Holonomy of a closed curve γ on the cube net

Definition 2.4.2. The *holonomy* of a closed curve is defined as the net rotation of a tangent vector under parallel transport around that curve.

This concept is closely related to the *cone angle* introduced earlier. In fact, the relationship can be made explicit: if γ is a small loop around a cone vertex and θ is the cone angle at that vertex, then

$$\text{holonomy along } \gamma \;=\; \text{rotation by } 2\pi - \theta.$$

The holonomy of a closed curve in the plane is zero. On a Platonic surface, the holonomy rarely lies in $2\pi\mathbb{Z}$, but always in $\frac{2\pi}{k}\mathbb{Z}$, for some $k \in \mathbb{N}$. For example, consider the holonomy of a closed curve γ on the cube, which is $\frac{1}{2}\pi$, see Figure 2.10. Nevertheless, one can always lift to an appropriate degree-k branched cover where it does. And this is exactly how we want to define the *unfolding of the Platonic surface*.

Definition 2.4.3. We define the *unfolding* $\widetilde{D}$ of a Platonic surface D, with *branched covering map* $\pi_D : \widetilde{D} \to D$, to be the smallest cover such that the *holonomy* of the cover lies in $2\pi\mathbb{Z}$ and it is a translation surface. In other words, every curve on D lifts to a curve on $\widetilde{D}$ with holonomy in $2\pi\mathbb{Z}$. Let k denote the degree of π_D.

We will always assume the unfolding is a regular p-gon tiled surface, which is achieved by choosing a horizontal direction. This arises because, by an appropriate rotation of its presentation, every edge is replicated in k unique orientations. Thus, each edge has a unique parallel copy that can be identified.

This yields a *natural branched covering map* $\pi_\Pi : \widetilde{D} \to \Pi_p$ with extra structure. Away from the vertices, it induces a bijection on unit tangent directions. Therefore, directions on $\widetilde{D}$ and Π_p correspond precisely modulo 2π. In contrast, this directional correspondence fails for π_D whenever $k > 1$.
Consider the cube unfolding shown in Figure 2.11. This is a degree-4 cover where the sheets are labeled with numbers, indicated by gray circles. We adopt the convention of specifying a location by a pair (x, y), where x denotes the sheet number and y denotes the specific p-gon. For example, $(3, 6)$ refers to square 6 in the third sheet. Here, the directional correspondence fails because $k = 4$. This can be seen by considering, a horizontal vector on $(3, 2)$ and a horizontal vector on $(4, 2)$. The covering map π_Π sends both vectors to horizontal vectors on the translation surface Π_4, whereas π_D maps them to square 2 on the cube. There, these vectors differ by a rotation of $\frac{\pi}{2}$, which breaks the correspondence.

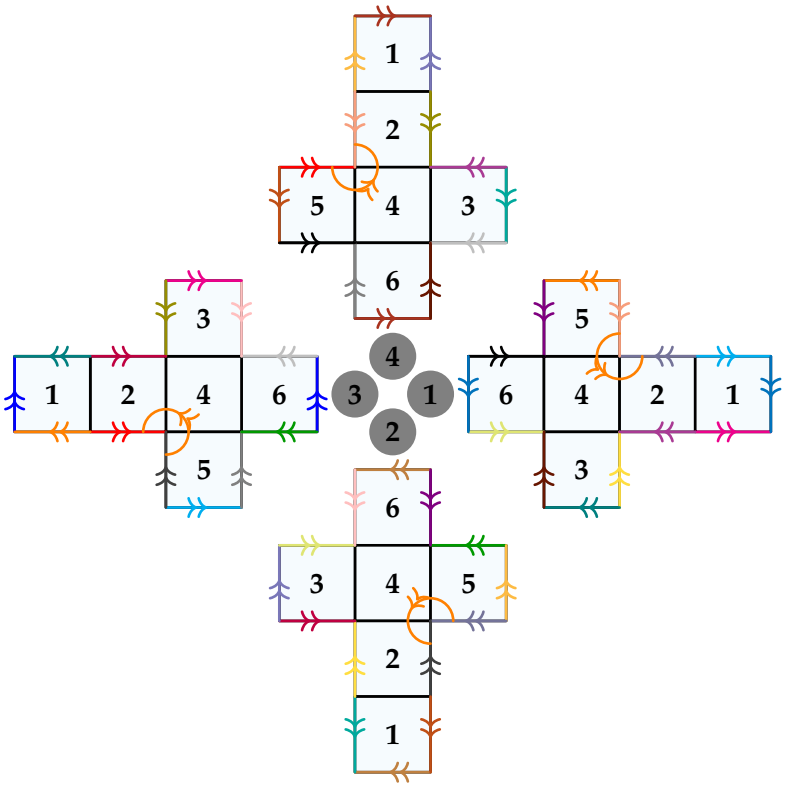

FIGURE 2.11: The cube unfolding with a closed curve γ

Let us return to our example of the closed curve γ on the cube (Figure 2.10). Now, consider the same curve γ on $\widetilde{D}$ (Figure 2.11). We observe that on $\widetilde{D}$, the curve γ begins on sheet 4, moves successively to sheets 3, 2, and 1, and returns to sheet 4. Hence, its cone angle is equal to $4 \cdot \frac{3}{2}\pi = 6\pi$. Thus, its holonomy is -4π, which is an integer multiple of 2π.

2.5 Groups

2.5.1 Deck Group

Definition 2.5.1. Let $p : E \to X$ be a covering. A *deck transformation* is a homeomorphism $d : E \to E$ such that the diagram of continuous maps commutes:

$$\begin{array}{ccc} E & \xrightarrow{\quad d \quad} & E \\ & {\scriptstyle p}\searrow \quad \swarrow{\scriptstyle p} & \\ & X & \end{array}$$

Definition 2.5.2. Together with the composition of maps, the set of *deck transformations* forms a group $\mathrm{Deck}(p)$, the *Deck group*.

Intuitively, the *Deck group* is the group of symmetries of a covering space. The *Deck group* $\mathrm{Deck}(p)$ is also denoted by $\mathrm{Aut}(p)$.

Definition 2.5.3. Let $p : E \to X$ be a covering map. If $Deck(p)$ acts transitively on every fiber, we call p *normal* or *regular*.

In other words, if p is normal, then for any point $x \in X$ and for any two points e_1, e_2 in the fiber $p^{-1}(x)$, there exists a deck transformation d such that $d(e_1) = e_2$.

2.5.2 Rotation Subgroup

Definition 2.5.4. For a Platonic surface D with automorphism group $\mathrm{Aut}(D)$, the *rotation subgroup* $\mathrm{Rot}(D)$ is defined as the subgroup of orientation-preserving automorphisms.

Proposition 2.5.5. *$\mathrm{Rot}(D)$ is a subgroup of $\mathrm{Aut}(D)$ of index two. Thus, it is a normal subgroup of $\mathrm{Aut}(D)$.*

Proof. Every automorphism of a Platonic surface either preserves or reverses orientation. Let

$$\alpha : \mathrm{Aut}(D) \longrightarrow \mathbb{Z}/2\mathbb{Z},$$

with

$$\alpha(f) = \begin{cases} 0, & \text{if } f \text{ preserves orientation,} \\ 1, & \text{if } f \text{ reverses orientation.} \end{cases}$$

for any $f \in \mathrm{Aut}(D)$. We claim that α is a group homomorphism. Let $f, g \in \mathrm{Aut}(D)$. The composition $f \circ g$ preserves orientation if and only if either both f and g preserve orientation or both reverse orientation. Therefore,

$$\alpha(f \circ g) = \alpha(f) + \alpha(g) \mod 2.$$

Thus, α is indeed a group homomorphism.

Next, we will show that α is surjective. The identity automorphism is orientation preserving, so $0 \in \mathrm{Im}(\alpha)$. Fix a global orientation on D and let $F = (f, e, v)$ be a flag on D. Locally, the flag F is either compatible with the global orientation, in which case it is called a *positive flag*, or incompatible, in which case it is called a *negative flag*. The set of all flags splits into two non-empty classes of positive and negative flags. Since D is a Platonic surface, $\mathrm{Aut}(D)$ acts transitively on flags. Thus, for any two flags F and F', $\exists f \in \mathrm{Aut}(D)$ such that $f(F) = F'$. In particular, choose $F = (f, e, v_1)$ to be a positive flag and let $F' = (f, e, v_2)$ be the same flag, but with the other vertex incident to the edge e. Then, F' is a negative flag. Thus, f is orientation reversing, which implies that $\alpha(f) = 1 \in \mathrm{Im}(\alpha)$. Therefore α is surjective.

By definition of the kernel, we have $\mathrm{Ker}(\alpha) = \{f \in \mathrm{Aut}(D) | \alpha(f) = 0\}$, i.e. the orientation preserving maps. Thus, $\mathrm{Ker}(\alpha) = \mathrm{Rot}(D)$. By the first isomorphism theorem, we know that

$$\mathrm{Aut}(D)/\mathrm{Ker}(\alpha) = \mathrm{Aut}(D)/\mathrm{Rot}(D) \cong \mathrm{Im}(\alpha) = \mathbb{Z}/2\mathbb{Z}.$$

Since $\mathbb{Z}/2\mathbb{Z}$ has two elements and α is surjective, we have

$$|\mathrm{Aut}(D) : \mathrm{Rot}(D)| = 2.$$

□

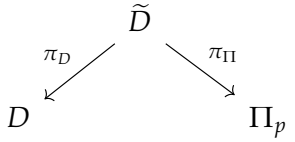

FIGURE 2.12: Covering maps and spaces of interest

Space or Map	Symbol
Unfolding of D	$\widetilde{D}$
Platonic Surface	D
Translation Surface	Π_p
Branched Covering Map	π_D
Natural Branched Covering Map	π_Π

TABLE 2.2: Notations in the diagram

2.5.3 Monodromy Representation

Definition 2.5.6. Let $p : Y \to X$ be a branched covering map of degree n, where x_0 is the fixed base point of X. Let $p^{-1}(x_0) = \{y_1, ..., y_n\}$. For $[\gamma] \in \pi_1(X, x_0)$, let $\tilde{\gamma}_i$ be the lift of γ with $\tilde{\gamma}_i(0) = y_i$. Then $\tilde{\gamma}_i(1) = y_{\sigma(i)}$, for some $\sigma(i) \in \{1, ..., n\}$. The *monodromy representation* is the natural homomorphism

$$\psi_p : \pi_1(X, x_0) \to S_n$$

$$[\gamma] \mapsto \sigma$$

Intuitively, this map labels the sheets of the covering and records the permutation of the loops affected by the lifting of these loops to Y. The homotopy lifting property shows that this map is well-defined. Moreover, if we equip $\pi_1(X, x_0)$ with the group operation $\alpha \circ \beta := \alpha * \beta$, where $\alpha * \beta$ is the concatenation of paths, then ψ_p becomes a group homomorphism.

Definition 2.5.7. The *monodromy group* $\mathrm{Mon}(\pi_\Pi)$ is the image of the *monodromy representation* ψ_{π_Π}.

Remark 2.5.8. There are different notions of *monodromy* and especially of the *monodromy group* in the study of regular maps, but this is not what we are studying here. Our notions are the typical ones found in topology.

Let us summarize the covering maps and spaces of interest in the diagram 2.12 and Table 2.2.

The proof for the following theorem follows [AAH19, Thm. 4.5].

Theorem 2.5.9. *The covering $\pi_\Pi : \widetilde{D} \to \Pi_p$ is regular. Furthermore, $\widetilde{D}$ is a Platonic surface. In other words, $\mathrm{Rot}(\widetilde{D})$, the group of orientation preserving isometries of $\widetilde{D}$, is transitive on pairs.*

Proof. First we show that $\widetilde{D}$ is a Platonic surface. Let $(\tilde{f}_1, \tilde{e}_1, \tilde{v}_1)$ and $(\tilde{f}_2, \tilde{e}_2, \tilde{v}_2)$ be flags on $\widetilde{D}$. We have a covering map $\pi_D : \widetilde{D} \to D$ that sends

$$(\tilde{f}_i, \tilde{e}_i, \tilde{v}_i) \mapsto (f_i, e_i, v_i), \forall i \in \{1, 2\}.$$

We know that D is a Platonic surface, so $\mathrm{Aut}(D)$ is transitive on flags. Therefore, there exists an isometry g that sends

$$(f_1, e_1, v_1) \mapsto (f_2, e_2, v_2).$$

Next, we need to find a lift that takes g to an isometry $\tilde{g}$ satisfying

$$\pi_D \circ \tilde{g}(\tilde{f}_1, \tilde{e}_1, \tilde{v}_1) = (f_2, e_2, v_2).$$

To do this, we recall the notion of parallel transport on a flat surface. Let $x_0 \in D^* := D \setminus \Sigma$, where Σ denotes the set of cone singularities. Since D^* is flat, vectors can be parallel transported along any smooth path in D^* by identifying tangent spaces via local Euclidean coordinates. Parallel transport around loops based at x_0 determines a homomorphism

$$\rho : \pi_1(D^*, x_0) \longrightarrow SO(2, \mathbb{R}),$$

where $SO(2, \mathbb{R})$ acts on the tangent space $T_{x_0}D^*$. Because $SO(2, \mathbb{R})$ is abelian, the homomorphism ρ factors through the abelianization of $\pi_1(D^*, x_0)$, and hence through the first integral homology group. We therefore obtain a well-defined homomorphism

$$PT\colon H_1(D^*; \mathbb{Z}) \longrightarrow SO(2, \mathbb{R}),$$

which we call the parallel transport map.

The isometry g satisfies $PT \circ g_* = PT$, where g_* is the induced map on $H_1(D^*)$. Next, we will apply the lifting criterion (Proposition 2.2.7), which says that the lift $\tilde{g} : (D^*, x_0) \to (\widetilde{D}, \tilde{x}_0)$ exists if and only if $g_*(\pi_1(D^*, x_0)) \subset (\pi_D)_*(\pi_1(\widetilde{D}, \tilde{x}_0))$, which we know is true. Hence, g lifts to an isometry $\tilde{g} \in \mathrm{Isom}(\widetilde{D})$ fulfilling

$$\pi_D \circ \tilde{g}(\tilde{f}_1, \tilde{e}_1, \tilde{v}_1) = (f_2, e_2, v_2).$$

Now, since π_D is a regular covering, we have $\delta \in \mathrm{Deck}(\pi_D)$, where $\delta : \widetilde{D} \to \widetilde{D}$ with $\delta \circ \tilde{g}(\tilde{f}_1, \tilde{e}_1, \tilde{v}_1) = (\tilde{f}_2, \tilde{e}_2, \tilde{v}_2)$. The composition $\delta \circ \tilde{g}$ is an isometry and proves that $\widetilde{D}$ is Platonic.

To see that the covering $\pi_\Pi : \widetilde{D} \to \Pi_p$ is regular, fix a flag (f, e, v) in Π_p. Consider the set of flags in $\widetilde{D}$ that are preimages of (f, e, v). As noted previously, $\mathrm{Isom}(\widetilde{D})$ acts transitively on these flags, which is exactly the deck transformation group of the covering π_Π. □

Chapter 3

On Orientable Surfaces

3.1 Main Theorems

This chapter is based on David Aulicino's work *A New Approach to the Automorphism Group of a Platonic Surface* [Aul19].

Lemma 3.1.1. *Consider two Platonic surfaces Y and X. Let $\pi : Y \to X$ be a regular covering. An orientation-preserving isometry $F : X \to X$ induces an automorphism $\widetilde{F} : Y \to Y$ such that the following diagram commutes.*

$$\begin{array}{ccc} Y & \overset{\widetilde{F}}{\dashrightarrow} & Y \\ {\scriptstyle\pi}\downarrow & & \downarrow{\scriptstyle\pi} \\ X & \xrightarrow{F} & X \end{array}$$

Proof. We want to construct a lift for $F : X \to X$. First, we select a pair (f_0, e_0) on X and its image under F, (f_1, e_1). Then, we lift both pairs to Y, obtaining $(\tilde{f}_0, \tilde{e}_0)$ and $(\tilde{f}_1, \tilde{e}_1)$. Now, since Y is a Platonic surface, there exists a unique isometry $\widetilde{F} : Y \to Y$ that takes $(\tilde{f}_0, \tilde{e}_0)$ to $(\tilde{f}_1, \tilde{e}_1)$. Then, $\pi \circ \widetilde{F}(\tilde{f}_0, \tilde{e}_0) = (f_1, e_1)$. Since the choice of (f_0, e_0) was arbitrary, we have $\widetilde{F} \circ \pi = F$. Now, $\widetilde{F}$ is a lift, by definition, so the existence of the automorphism $\widetilde{F}$ is guaranteed. □

Proposition 3.1.2. *Let $\pi : Y \to X$ be a regular covering of Platonic surfaces Y and X. This induces a surjective homomorphism*

$$\phi_\pi : \mathrm{Rot}(Y) \to \mathrm{Rot}(X),$$

with the property that

$$\ker(\phi_\pi) = \mathrm{Deck}(\pi) \cong \mathrm{Mon}(\pi).$$

Proof. We will explicitly construct the homomorphism $\phi_\pi : \mathrm{Rot}(Y) \to \mathrm{Rot}(X)$. First, we will show that ϕ_π is a well-defined map. Let $\widetilde{F} \in \mathrm{Rot}(Y)$ and choose a pair $(\tilde{f}_1, \tilde{e}_1)$ on Y such that $\widetilde{F}(\tilde{f}_1, \tilde{e}_1) = (\tilde{f}_2, \tilde{e}_2)$. Now consider $(f_i, e_i) = (\pi(\tilde{f}_i), \pi(\tilde{e}_i))$, for $i \in \{1, 2\}$. Then, we can define F to be the unique rotation in X that takes (f_1, e_1) to

(f_2, e_2). So, we have the following diagram.

$$\begin{array}{ccc} (\tilde{f}_1, \tilde{e}_1) & \xrightarrow{\widetilde{F}} & (\tilde{f}_2, \tilde{e}_2) \\ {\scriptstyle\pi}\downarrow & & \downarrow{\scriptstyle\pi} \\ (f_1, e_1) & \xrightarrow{F} & (f_2, e_2) \end{array}$$

By Lemma 3.1.1, we can lift F to $\widetilde{F}' \in \mathrm{Rot}(Y)$. Then $\widetilde{F}'$ rotates the same elements as before differently.

$$\begin{array}{ccc} (\tilde{f}_1', \tilde{e}_1') & \xrightarrow{\widetilde{F}'} & (\tilde{f}_2', \tilde{e}_2') \\ {\scriptstyle\pi}\downarrow & & \downarrow{\scriptstyle\pi} \\ (f_1, e_1) & \xrightarrow{F} & (f_2, e_2) \end{array}$$

Our goal now is to show that $\widetilde{F} = \widetilde{F}'$. We know that π is a regular covering, so the deck group is transitive. Also, since the lift is unique up to actions of the deck group, there exists an isometry taking $(\tilde{f}_i', \tilde{e}_i')$ to $(\tilde{f}_i, \tilde{e}_i)$. We will call these elements $\delta_i \in \mathrm{Deck}(\pi)$. From this we get a new diagram:

$$\begin{array}{ccc} (\tilde{f}_1', \tilde{e}_1') & \xrightarrow{\widetilde{F}'} & (\tilde{f}_2', \tilde{e}_2') \\ {\scriptstyle\delta_1}\downarrow & & \downarrow{\scriptstyle\delta_2} \\ (\tilde{f}_1, \tilde{e}_1) & \xrightarrow{\widetilde{F}} & (\tilde{f}_2, \tilde{e}_2) \end{array}$$

Therefore, the composition $\delta_2 \circ \widetilde{F}' \circ \delta_1^{-1}$ acts as a rotation on Y, and it sends $\widetilde{F}(\tilde{f}_1, \tilde{e}_1)$ to $(\tilde{f}_2, \tilde{e}_2)$. Because a rotation is uniquely determined by its action on a single pair, we obtain $\delta_2 \circ \widetilde{F}' \circ \delta_1^{-1} = \widetilde{F}$. Hence, the construction of F from $\widetilde{F}$ is independent of the initial choice of pair, so ϕ_π is well-defined.

The fact that ϕ_π is a homomorphism follows from the commutativity of the following diagram.

$$\begin{array}{ccccc} Y & \xrightarrow{\widetilde{F}_1} & Y & \xrightarrow{\widetilde{F}_2} & Y \\ {\scriptstyle\pi}\downarrow & & \downarrow{\scriptstyle\pi} & & \downarrow{\scriptstyle\pi} \\ X & \xrightarrow{F_1} & X & \xrightarrow{F_2} & X \end{array}$$

Finally, the kernel of ϕ_π consists precisely of those maps for which $\widetilde{F}(\tilde{f}_1, \tilde{e}_1) = (\tilde{f}_2, \tilde{e}_2)$. This condition is exactly that two maps differ by a deck transformation. Finally, surjectivity follows by considering lifts of all possible pairs (f, e) on X. □

Remark 3.1.3. Notice that we can now apply Proposition 3.1.2 to ϕ_Π, since Π_p is a Platonic surface. Further, by Lemma 3.1.1, π_Π is a regular cover and $\widetilde{D}$ is a Platonic surface. Therefore Proposition 3.1.2 also applies to ϕ_D.

Summarized, $\pi_\Pi : \widetilde{D} \to \Pi_p$ induces $\phi_\Pi : \mathrm{Rot}(\widetilde{D}) \to \mathrm{Rot}(\Pi_p)$, and $\pi_D : \widetilde{D} \to D$ induces $\phi_D : \mathrm{Rot}(\widetilde{D}) \to \mathrm{Rot}(D)$.

The following Proposition comes from [Dyd11, Prop. 2.10].

Proposition 3.1.4. *Let $\pi_\Pi : \widetilde{D} \to \Pi_p$ be a regular covering map. Then $\forall \tilde{d}_0 \in \widetilde{D}$ there exists a natural monomorphism*

$$\chi : \mathrm{Deck}(\pi_\Pi) \to \mathrm{Mon}(\pi_\Pi)$$

$$\tilde{d}_0 \mapsto x_0,$$

where $x_0 \in \mathrm{Mon}(\pi_\Pi)$.

The monomorphism χ is an isomorphism if $\mathrm{Deck}(\pi_\Pi)$ acts transitively on the fibers of π_Π.

Proof. We need to show that χ is an injective homomophism. $\forall \delta_1 \in \mathrm{Deck}(\pi_\Pi)$, we choose a path γ_{δ_1} in $\widetilde{D}$ from $\tilde{d}_0$ to $\delta_1(\tilde{d}_0)$. Since π_Π is a regular covering, the equivalence class $[\pi_\Pi \circ \gamma_{\delta_1}]$ does not depend on the choice of γ_{δ_1}. $[\pi_\Pi \circ \gamma_{\delta_1}]$ is a class of loops in Π_p based at x_0.
If $\delta_2 \in \mathrm{Deck}(\pi_\Pi)$, then the concatenation $\gamma_{\delta_2} * \delta_2(\gamma_{\delta_1})$ is a path from $\tilde{d}_0$ to $\delta_2(\delta_1(\tilde{d}_0))$ and

$$[\pi_\Pi(\gamma_{\delta_2} * \delta_2(\gamma_{\delta_1}))] = [\pi_\Pi(\gamma_{\delta_2}) * \pi_\Pi(\gamma_{\delta_1})].$$

Since the composition is preserved on the level of equivalence classes, χ is a homomorphism.
Next, we will show that χ is injective. Suppose $\chi(\delta_1) = \chi(\delta_2)$ in $\mathrm{Mon}(\pi_\Pi)$. Then, $[\pi_\Pi \circ \gamma_{\delta_1}] = [\pi_\Pi \circ \gamma_{\delta_2}]$. So, they are homotopic in Π_p. By the homotpy lifting property, the unique lifts starting at $\tilde{d}_0$ will have the same endpoints. But, $\pi_\Pi \circ \gamma_{\delta_1}$ lifts to γ_{δ_1}, ending at $\delta_1(\tilde{d}_0)$ and $\pi_\Pi \circ \gamma_{\delta_2}$ lifts to γ_{δ_2}, ending at $\delta_2(\tilde{d}_0)$. Thus, $\delta_1(\tilde{d}_0) = \delta_2(\tilde{d}_0)$. Now, δ_1, δ_2 are deck transformations agreeing at one point $\tilde{d}_0$. For covering spaces, deck transformations are uniquely determined by their value at a single point if $\widetilde{D}$ is connected, which it is. Therefore, $\delta_1 = \delta_2$, so χ is injective.

Assume, $\mathrm{Deck}(\pi_\Pi)$ acts transtively on the fibers of π_Π and $[\gamma] \in \mathrm{Mon}(\pi_\Pi)$, then we can lift γ to $\tilde{\gamma}$ and pick a deck transformation δ s.t. $\delta(\tilde{\gamma}(0)) = \delta(\tilde{\gamma}(1))$. Now, $\chi(\delta) = [\pi_\Pi \circ \tilde{\gamma}] = [\gamma]$. Therefore, χ is surjective and an isomorphism. □

Remark 3.1.5. The Proposition 3.1.4 was originally given in a more general case. For our purposes, and the way we have stated it, χ is always an isomorphism. For any connected covering $p : Y \to X$, the deck group $\mathrm{Deck}(p)$ acts transitively on every fiber if and only if X is path-connected and the covering p is normal. Since $\widetilde{D}$ is path-connected and π_Π is normal, by Theorem 2.5.9, we have

$$\mathrm{Mon}(\pi_\Pi) \cong \mathrm{Deck}(\pi_\Pi).$$

Lemma 3.1.6. *There exists a monomorphism*

$$\widetilde{\Psi} : \mathrm{Mon}(\pi_\Pi) \hookrightarrow \mathrm{Rot}(\widetilde{D}).$$

Proof. Remark 3.1.5 tells us that $\mathrm{Deck}(\pi_\Pi) \cong \mathrm{Mon}(\pi_\Pi)$. Furthermore, since deck transformations are, by definition, homeomorphisms of $\widetilde{D}$ that permute the sheets of the covering π_Π, they act as isometries of $\widetilde{D}$. Permuting the sheets of the covering can be interpreted as a rotation of the covering space. Therefore, $\mathrm{Deck}(\pi_\Pi)$ forms a

subgroup of $\mathrm{Rot}(\widetilde{D})$. Thus, the homomorphism $\mathrm{Deck}(\pi_\Pi) \cong \mathrm{Mon}(\pi_\Pi) \to \mathrm{Rot}(\widetilde{D})$ is injective. □

Lemma 3.1.7. *Let $\delta \in \ker(\phi_D)$, and suppose δ maps a pair (f, e) to a pair (f', e'). Let θ be the angle between the edges e and e' on the surface $\widetilde{D}$. Then, $\theta \equiv 0 \mod (2\pi)$ if and only if δ is the identity.*

Proof. If δ is the identity, the statement is trivial. Recall that $\ker(\phi_D)$ is exactly the deck group of the covering $\pi_D : \widetilde{D} \to D$, which has degree k. By definition, this group consists of rotations of $\widetilde{D}$ by integer multiples of $2\pi/k$. Thus, any non-identity element δ rotates edges by a non-zero multiple of $2\pi/k$. □

The following theorem is the main theorem of this section. To prove it, we apply results from finite group theory to the following diagram. The reader is advised to refer to the following diagram as a visual guide throughout the argument.

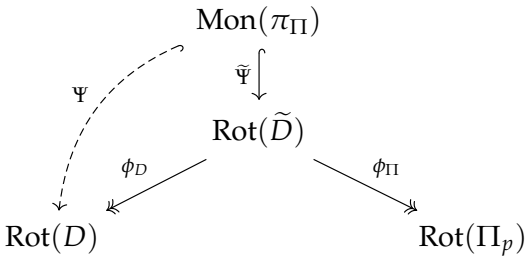

Theorem 3.1.8. *There is an isomorphism between the monodromy group $\mathrm{Mon}(\pi_\Pi)$ and a normal subgroup of $\mathrm{Rot}(D)$. Moreover, the corresponding quotient group*

$$\mathrm{Rot}(D)/\phi_D(\mathrm{Mon}(\pi_\Pi))$$

is cyclic.

Proof. The Lemma 3.1.6 tells us that there exists a monomorphism $\widetilde{\Psi} : \mathrm{Mon}(\pi_\Pi) \hookrightarrow \mathrm{Rot}(\widetilde{D})$. Since π_Π is a normal covering, we can apply Proposition 3.1.2 and Remark 3.1.3. Thus, we know that π_Π induces a surjective homomorphism $\phi_\Pi : \mathrm{Rot}(\widetilde{D}) \twoheadrightarrow \mathrm{Rot}(\Pi_p)$. Combining these statements, we can notice that we have a short exact sequence:

$$0 \longrightarrow \mathrm{Mon}(\pi_\Pi) \xhookrightarrow{\widetilde{\Psi}} \mathrm{Rot}(\widetilde{D}) \xrightarrow{\phi_\Pi} \mathrm{Rot}(\Pi_p) \longrightarrow 0$$

Now, since we also know that in any group homomorphism the kernel is always a normal subgroup, we have that $\ker(\phi_\Pi) = \widetilde{\Psi}(\mathrm{Mon}(\pi_\Pi))$ is a normal subgroup of $\mathrm{Rot}(\widetilde{D})$.
Furthermore, since π_D is also a normal covering, we can once again apply Proposition 3.1.2 and Remark 3.1.3 to induce a surjective homomorphism $\phi_D : \mathrm{Rot}(\widetilde{D}) \twoheadrightarrow \mathrm{Rot}(D)$. Finally, we claim that since ϕ_D is surjective and $\widetilde{\Psi}(\mathrm{Mon}(\pi_\Pi))$ is normal in $\mathrm{Rot}(\widetilde{D})$, this implies that $\phi_D(\widetilde{\Psi}(\mathrm{Mon}(\pi_\Pi)))$ is normal in $\mathrm{Rot}(D)$.

Since ϕ_D is surjective, $\forall \tau \in \mathrm{Rot}(D)$, $\exists \sigma \in \mathrm{Rot}(\widetilde{D})$, such that $\phi_D(\sigma) = \tau$. Let $\delta \in \widetilde{\Psi}(\mathrm{Mon}(\pi_\Pi))$, then

$$\begin{aligned}\tau^{-1}\phi_D(\delta)\tau &= \phi_D(\sigma)^{-1}\phi_D(\delta)\phi_D(\sigma)\\ &= \phi_D(\sigma^{-1})\phi_D(\delta)\phi_D(\sigma)\\ &= \phi_D(\sigma^{-1}\delta\sigma).\end{aligned}$$

Now, $\sigma^{-1}\delta\sigma \in \widetilde{\Psi}(\mathrm{Mon}(\pi_\Pi))$, since $\widetilde{\Psi}(\mathrm{Mon}(\pi_\Pi))$ is normal in $\mathrm{Rot}(\widetilde{D})$. Thus, $\phi_D(\sigma^{-1}\delta\sigma) \in \phi_D(\widetilde{\Psi}(\mathrm{Mon}(\pi_\Pi)))$. Therefore, $\tau^{-1}\phi_D(\delta)\tau \in \phi_D(\widetilde{\Psi}(\mathrm{Mon}(\pi_\Pi)))$. Since τ and $\phi_D(\delta)$ were arbitrary, we have shown that $\phi_D(\widetilde{\Psi}(\mathrm{Mon}(\pi_\Pi)))$ is normal in $\mathrm{Rot}(D)$, which proves the claim.

Next, we claim that $\phi_D \circ \widetilde{\Psi} : \mathrm{Mon}(\pi_\Pi) \to \mathrm{Rot}(D)$ is injective. If this claim holds, then $\mathrm{Mon}(\pi_\Pi) \cong \phi_D(\widetilde{\Psi}(\mathrm{Mon}(\pi_\Pi)))$, the first part of the proof would be done. Since $\widetilde{\Psi}$ is already injective, the claim can only fail, if

$$\widetilde{\Psi}(\mathrm{Mon}(\pi_\Pi)) \cap \ker(\phi_D) \neq \{\mathrm{id}\}.$$

We know that $\ker(\phi_D)$ is $\mathrm{Deck}(\pi_D)$, by Proposition 3.1.2. Also, by Lemma 3.1.7, every non-trivial element of $\ker(\phi_D)$ rotates edges by a non-zero angle. An element of $\mathrm{Mon}(\pi_\Pi)$ permutes sheets of the covering, it does not rotate faces, it only relabels them, i.e. it preserves angles. Thus, every non-trivial deck transformation rotates edges and every monodromy element never rotates edges. Therefore, their intersection is the identity. Hence,

$$\ker(\phi_D \circ \widetilde{\Psi}) = \{\mathrm{id}\} \Rightarrow \phi_D \circ \widetilde{\Psi} \text{ is injective.}$$

Finally, let $\nabla : \mathrm{Rot}(D) \to \mathrm{SO}(2)$ be the derivative map, where $SO(2)$ is the special orthogonal group of degree 2. In a local chart $\varphi : U \subset D \to \mathbb{R}^2$, the map f can be written as $\varphi \circ f \circ \varphi^{-1}(x) = Ax + b$, where $A \in \mathrm{SO}(2)$ and $b \in \mathbb{R}^2$. The linear part A is independent of the choice of a chart, while b may vary. We define $\nabla(f) := A$.

Again, elements of $\mathrm{Mon}(\pi_\Pi)$ do not rotate faces. Thus, their derivatives are trivial. Therefore,

$$\mathrm{Mon}(\pi_\Pi) = \ker(\nabla).$$

The image of ∇ is a finite subgroup of $\mathrm{SO}(2)$, and since every finite subgroup of $\mathrm{SO}(2)$ is cyclic, $\mathrm{Im}(\nabla)$ is cyclic. Therefore,

$$\mathrm{Rot}(D)/\phi_D(\mathrm{Mon}(\pi_\Pi)) \cong \mathrm{Im}(\nabla)$$

is cyclic.

□

This theorem is significant because $\mathrm{Mon}(\pi_\Pi)$ is generally easy to compute for a Platonic surface D. We can represent it as a subgroup of a large symmetric group. Although this does not always yield the full automorphism group of the regular map, it is sometimes sufficient to completely determine $\mathrm{Rot}(D)$.

3.2 Some Bounds

Recall that k is the degree of the covering $\pi_D : \widetilde{D} \to D$. Even though explicit values for k are difficult to determine due to their dependence on the global geometry of the surface, the bounds can be derived in terms of the Schläfli symbol $\{p, q\}$. This then also bounds the algorithms that are needed for the group computations.

Remark 3.2.1. For the remainder of this paper, we adopt the following conventions:

- Let $d = \gcd(p, q)$,
- Let m denote the number of faces of D.

Proposition 3.2.2. *Let D be a Platonic surface with Schläfli symbol $\{p, q\}$. Then, the cone angle at each vertex equals*

$$\frac{q(p-2)}{p}\pi.$$

Proof. Any regular p-gon can be decomposed into p congruent isosceles triangles by connecting each vertex with the barycenter. The vertex angle at the barycenter is $\frac{2\pi}{p}$. The remaining two angles of the isosceles triangles sum to $\pi - \frac{2\pi}{p} = \frac{p-2}{p}\pi$. At any vertex of the surface, q such interior angles meet. Therefore, the cone angle equals $q \cdot \frac{(p-2)}{p}\pi$. □

Proposition 3.2.3. *Consider a Platonic surface D with Schläfli symbol $\{p, q\}$, and define k' to be the smallest positive integer that satisfies the following equation:*

$$k' \frac{q(p-2)}{p}\pi \equiv 0 \mod 2\pi$$

Let k be the degree of the covering of the unfolded Platonic surface. Then $k' \mid k$. Moreover,

- *if p is even, then $k' \leq k \leq p$,*
- *if p is odd, then $k' \leq k \leq 2p$.*

Proof. By Proposition 3.2.2, k' is the minimal integer required to ensure that each cone point in the unfolding has a cone angle that is an integer multiple of 2π. Any positive integer failing this condition does not produce such cone points. Thus, $k' \mid k$. Next, we analyze the holonomy of a closed curve on D. Let γ be a closed curve on D. The surface is constructed from regular p-gons, each having interior angle $(p-2)\pi/p$. Therefore, as γ enters and exits a face, its tangent direction changes by $m \cdot \pi/p$ for some integer m. Summing over all faces gives total holonomy in $(\pi/p)\mathbb{Z}$. Lifting to an appropriate cyclic $2p$-fold cover then ensures that every curve's holonomy lies in $2\pi\mathbb{Z}$.
Now, assume that p is even. Then $2 - p$ is even, so one factor of 2 can be taken out from the numerator $p - 2$. Thus, each curve crossing a polygon changes the holonomy by an angle that is a multiple of $2\pi/p$. Therefore, a p-fold cover suffices for the maximal degree of the cover. □

p \ q	odd	0 mod 4	2 mod 4
odd	$\frac{2p}{d}$	$\frac{p}{d}$	$\frac{p}{d}$
0 mod 4	$\frac{p}{d}$	$\frac{p}{d}$	$\frac{p}{d}$
2 mod 4	$\frac{p}{2d}$	$\frac{p}{d}$	$\frac{p}{d}$

TABLE 3.1: Values for k'

Lemma 3.2.4. *For a surface D with Schläfli symbol $\{p, q\}$, the value of k' from Proposition 3.2.3 is provided in Table 3.1.*

Proof. Consider

$$\frac{p}{d}\left(\frac{q(p-2)}{p}\right) = \frac{p}{d}(p-2) =: C.$$

Since $d = \gcd(p, q)$, we know that $d \mid q$. Thus, $\frac{q}{d}$ is an integer. The parity of q is also important.

Case 1: q is even $\implies$ C is even.

Case 2: q is odd, then either

- p is also odd $\implies$ C is odd and $k' = \frac{2p}{d}$.
- $4 \mid p \implies 2 \mid (p-2)$ and $4 \nmid (p-2)$. Thus, $k' = \frac{p}{d}$.
- $4 \mid (p-2) \implies C/2$ is even and $2 \nmid d$ since q is odd. Thus, $k' = \frac{p}{2d}$.

□

The following Proposition stems from [Bre93, Chapter III, Cor. 5.2].

Proposition 3.2.5. *Let j be the degree of the covering $\pi_\Pi : \widetilde{D} \to \Pi_p$. The order of* $\mathrm{Mon}(\pi_\Pi)$ *equals j.*

Proof. Choose a base point x_0 in $\Pi_p^* = \Pi_p \setminus \Sigma$, where Σ denotes the set of branch points. Then $\pi_\Pi^{-1}(x_0)$ has j points. Since π_Π is regular, we know that $\forall x \in \pi_\Pi^{-1}(x_0)$ there exists a unique $\gamma \in \mathrm{Mon}(\pi_\Pi)$ such that $\gamma(\tilde{x}_0) = x$, where $\tilde{x}_0 \in \pi_\Pi^{-1}(x_0)$ is a fixed lift of x_0. No non-trivial element of $\mathrm{Mon}(\pi_\Pi)$ fixes a point in the fiber. Therefore, the action is free. Thus, each element of the fiber is reached by a unique element of the monodromy group. From this, we can conclude that

$$|\mathrm{Mon}(\pi_\Pi)| = |\pi_\Pi^{-1}(x_0)| = j.$$

□

Proposition 3.2.6. *Let D be a Platonic surface and let D have Schläfli symbol $\{p, q\}$. The order of the rotation group* $\mathrm{Rot}(D)$ *equals mp.*

p \ q	odd	0 mod 4	2 mod 4
odd	$\frac{mp}{d}$	$\frac{mp}{2d}$	$\frac{mp}{2d}$
0 mod 4	$\frac{mp}{d}$	$\frac{mp}{d}$	$\frac{mp}{d}$
2 mod 4	$\frac{mp}{2d}$	$\frac{mp}{d}$	$\frac{mp}{d}$

TABLE 3.2: Lower bounds for $|\mathrm{Mon}(\pi_\Pi)|$

p \ q	odd	0 mod 4	2 mod 4
odd	d	$2d$	$2d$
0 mod 4	d	d	d
2 mod 4	$2d$	d	d

TABLE 3.3: Upper bounds for $|\mathrm{Rot}(D)/\phi_D(\mathrm{Mon}(\pi_\Pi))|$

Proof. Every Platonic surface has m faces and p edges. Thus, there are mp pairs. By the definition of a Platonic surface, the rotation group acts transitively on these pairs. Each transformation is uniquely determined by the image of a single pair. Thus, $|\mathrm{Rot}(D)| = mp$. □

Proposition 3.2.7. *Let D be a Platonic surface with Schläfli symbol $\{p,q\}$. The lower bounds for the order of the monodromy group $\mathrm{Mon}(\pi_\Pi)$ are given in Table 3.2.*

Proof. The Proposition 3.2.5 gives $|\mathrm{Mon}(\pi_\Pi)| = km$ if p is even, and if p is odd, $|\mathrm{Mon}(\pi_\Pi)| = km/2$. By Proposition 3.2.3, $k' \mid k$ and by Lemma 3.2.4, k' is given in Table 3.1.
Furthermore, if p is odd, Π_p consists of two polygons and mk is an even number. Thus, we can note that all the lower bounds given in Table 3.2 are integers. □

Corollary 3.2.8. *For a Platonic surface D with Schläfli symbol $\{p,q\}$, Table 3.3 gives an upper bound for the order of the cyclic group $\mathrm{Rot}(D)/\phi_D(\mathrm{Mon}(\pi_\Pi))$.*

Theorem 3.2.9. *Let D be a Platonic surface with Schläfli symbol $\{p,q\}$. If $\gcd(p,q) = 1$ and either*

- *4 divides p, or*
- *p and q are odd,*

then

$$\mathrm{Mon}(\pi_\Pi) \cong \mathrm{Rot}(D).$$

Proof. We now consult Table 3.3. For our purposes, we require the upper bound of $|\mathrm{Rot}(D)/\phi_D(\mathrm{Mon}(\pi_\Pi))|$ to be 1. This occurs if and only if the relevant entry in

Table 3.3 is 1, which is equivalent to $d = \gcd(p,q) = 1$. Thus, the only entries of Table 3.3 that concern us are those equal to d.

Observe that p and q cannot both be even, as this would imply $\gcd(p,q) \geq 2$. It follows that at least one of p or q must be odd. The surviving cases are therefore p and q both odd, or $p \equiv 0 \mod 4$ together with q odd. These are exactly the two cases considered in Theorem 3.2.9. □

3.3 Examples

3.3.1 Cube

The first polyhedron we examine in detail is the cube. Our aim is to obtain a concrete example that illustrates the isomorphism stated in Theorem 3.2.9. The cube has Schläfli symbol $\{4,3\}$ and since $\gcd(4,3) = 1$ and p is divisible by 4, the theorem applies directly.

First, we identify the required translation surface Π_4, which corresponds to a torus. The fundamental group of the surface Π_4 is generated by two elements, which we denote by h and v with a basepoint x_0, see Figure 3.1.

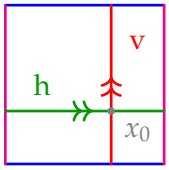

FIGURE 3.1: Translation surface Π_4 with two generators

As seen previously, when defining the holonomy of a closed curve (Figure 2.10), we know that the cone angle of the cube is $\frac{3}{4}\pi$. Thus, the holonomy of a closed curve on the cube is $\frac{1}{2}\pi$. Following the definition of an unfolding of a Platonic surface, we know that the unfolding is the smallest cover such that the holonomy lies in $2\pi\mathbb{Z}$. Therefore, we have 4 rotated copies of the cube net in our unfolding, see Figure 3.2.

Next, we want to trace the action of the generators in the unfolding. To do this, we need to consider the lifts of h and v starting at the various pre-images of the basepoint x_0. The cube unfolding is a 24-sheet cover of Π_4. Thus, there are 24 pre-images of x_0 in the unfolding. The lifts of the generators h or v then start at one of the various pre-images of x_0 and there exists a permutation that maps the sheet of a starting point of a lift to the sheet of the endpoint of the same lift.

For example, consider $\tilde{h}_1$, the lift of h with $\tilde{h}_1(0) = y_{(1,5)}$, where $y_{(1,5)}$ is a pre-image of x_0 in the sheet $(1,5)$. Recall that $(1,5)$ denotes the sheet "5" of the first copy of the unfolding. Thus, the first coordinate indicates which of the four rotated copies of the net we are in, and the second indicates the corresponding face. Now, the starting point of $\tilde{h}_1$ is $\tilde{h}_1(0) = y_{(1,5)}$, and the endpoint is $\tilde{h}_1(1)$. Next, we need to figure out where $\tilde{h}_1(1)$ is located. For that, we can follow $\tilde{h}_1$, which leads to the boundary of the sheet $(1,5)$. If we were to map this trajectory back onto the cube, the corresponding adjacent face is labeled "2". Taking orientation into account, the path naturally continues in a straight line to the face $(4,2)$, because this face is positioned so that a face labeled "5" lies to its left, maintaining the straight path of $\tilde{h}_1$. Thus, $\tilde{h}_1(1) = y_{(4,2)}$. Continuing this process, the path eventually closes, returning to its

starting sheet $(1,5)$:

$$(1,5) \mapsto (4,2) \mapsto (3,3) \mapsto (2,6) \mapsto (1,5).$$

By iterating over all possible pre-images of x_0 with all possible lifts of h and v, we obtain all possible permutations for the generators h and v in the unfolding. This is needed to calculate the order of the monodromy group, which can be done with Sage for example, see Listing 3.1. Applying the Sage functions will yield a group of order 24. Thus, we have that the order of $\mathrm{Mon}(\pi_\Pi)$ is 24.

```
phi_h_cube = [((1,5), (4,2), (3,3), (2,6)),
((1,6), (1,4), (1,2), (1,1)),
((1,3), (2,2), (3,5), (4,6)),
((2,4), (2,5), (4,1), (2,3)),
((2,1), (4,5), (4,4), (4,3)),
((3,1), (3,2), (3,4), (3,6))]

phi_v_cube = [((1,1), (3,5), (3,4), (3,3)),
((1,2), (2,5), (3,6), (4,3)),
((1,3), (1,4), (1,5), (3,1)),
((1,6), (4,5), (3,2), (2,3)),
((2,1), (2,2), (2,4), (2,6)),
((4,6), (4,4), (4,2), (4,1))]

GMon_cube = PermutationGroup([phi_h_cube, phi_v_cube])
GMon_cube.order()
```

LISTING 3.1: Calculations for the cube monodromy group

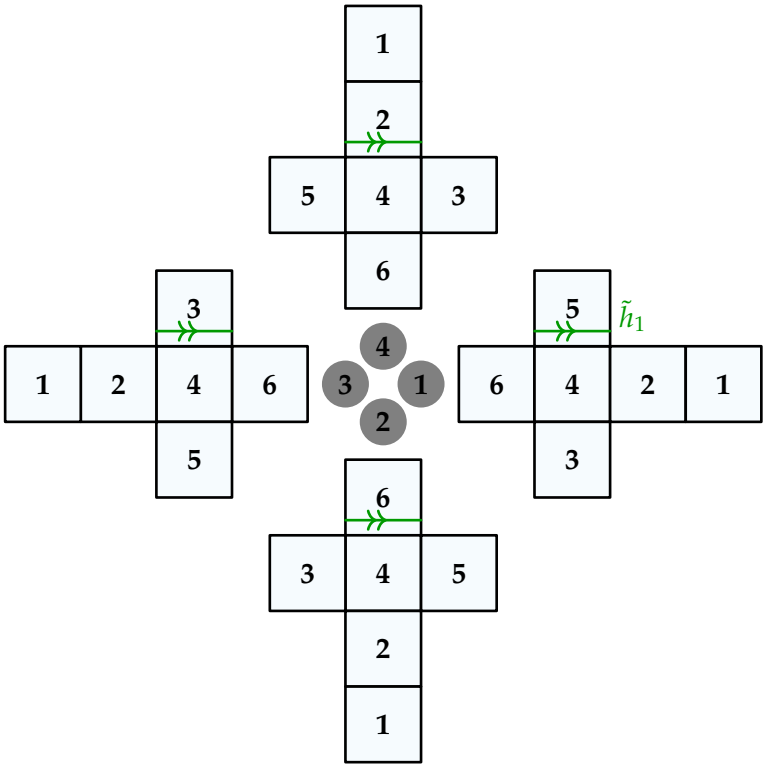

FIGURE 3.2: The cube unfolding with the action of $\tilde{h}_1$

Applying the isomorphism stated in Theorem 3.2.9, the order of $\mathrm{Rot}(D)$ should be the same as the order of $\mathrm{Mon}(\pi_\Pi)$. Since the cube is a rather easy example, we could have also noticed that each face of the cube can be rotated in 4 different ways, giving $|\mathrm{Rot}(D)| = 24$. This result is consistent with the isomorphism stated in Theorem 3.2.9 and the calculations produce $\mathrm{Rot}(D)$ as a faithful representation into S_{24}.

Finally, let us examine the action of the monodromy group in detail. Starting from the square $(1,1)$, we act by the identity element e, which leaves us at $(1,1)$. Thus, we have the correspondence $e \leftrightarrow (1,1)$. Acting by $\psi(h)$ once sends $(1,1)$ to $(1,6)$, so

$$\psi(h) \leftrightarrow (1,6).$$

Similarly, $\psi(h)^2 \leftrightarrow (1,4)$ and $\psi(h)^{-1} \leftrightarrow (1,2)$. Naturally, we may also consider compositions of these actions. For instance, if v corresponds to moving upwards, then $\psi(v) \leftrightarrow (3,5)$. Consequently, the composition $\psi(v)\psi(h)\psi(v)^{-1}$ corresponds to the square $(4,1)$.

3.3.2 Octahedron

Another important example to consider is the octahedron. This is the only Platonic solid that is excluded from Theorem 3.2.9, since $q = 0 \mod 4$, and p is odd. In the discussion of this example, we will see that the isomorphism does in fact not hold. First, we will consider the octahedron net, shown in Figure 3.3.

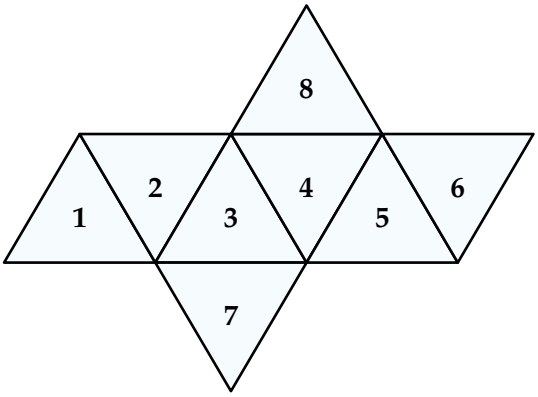

FIGURE 3.3: The octahedron net

The Schläfli number of the octahedron is $\{3,4\}$. Since $p = 3$, the resulting translation surface is Π_3. By Remark 2.1.10, Π_3 is constructed from two regular p-gons. Notice, the fundamental group of Π_3, once again, has two generators, which we will denote by h and v, see Figure 3.4.

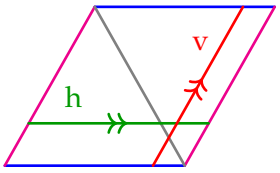

FIGURE 3.4: Generators of Π_3

Next, we will consider the unfolding of the octahedron with π_D, which is shown in Figure 3.5. The cone angle of the octahedron is $\frac{\pi}{3}$. Thus, the holonomy of a closed curve is equal to $2\pi - \frac{\pi}{3} = \frac{5}{3}\pi$. For a closed curve to have holonomy in $2\pi\mathbb{Z}$, we

would need $6 \cdot \frac{5}{3}\pi = 10\pi$. Consequently, the branched covering map $\pi_D : \widetilde{D} \to D$ should be a degree-6 covering. However, as we will see by the computations done with π_Π, it is enough to consider the degree-3 unfolding of the octagon, as the two sets of cycles we can generate are completely disjoint. Moreover, each set individually generates the same group size as they do together. Thus, it suffices to consider only the 3-sheeted cover.

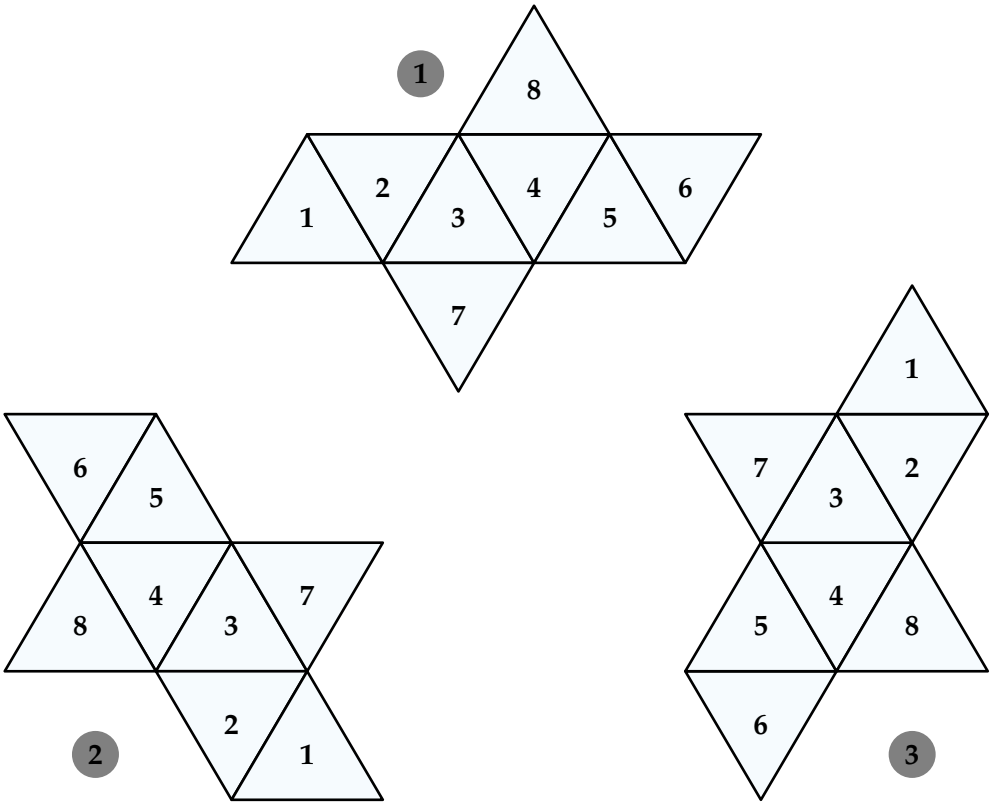

FIGURE 3.5: The octahedron unfolding with π_D

Furthermore, we need to consider the natural branched covering map $\pi_\Pi : \widetilde{D} \to \Pi_3$. Here, Π_3 is being covered. Thus, we must not trace the lifts of the triangular faces of the octahedron, but Π_3 itself (or rather the lifts of its generators). Therefore, our initial way of numbering the faces makes the labeling less efficient. Thus, we will simply label every two faces, such that it stays consistent with the unfolding and the identification of its edges. This can be seen in Figure 3.6. It follows that π_Π is a degree-12 covering map.

Now, we can trace the lifts of its generators h and v in the unfolded octahedron with π_Π, figure out the cycles, and run a similar code in Sage, as described in Listing 3.2.

```
phi_h_cube = [(1,2,3), (4,11,6), (5,7,12), (8,9,10)]
phi_v_cube = [(1,8,7), (2,5,4), (3,6,9), (10,11,12)]
GMon_cube = PermutationGroup([phi_h_cube, phi_v_cube])
GMon_cube.order()
```

LISTING 3.2: Calculations for the octahedron monodromy group

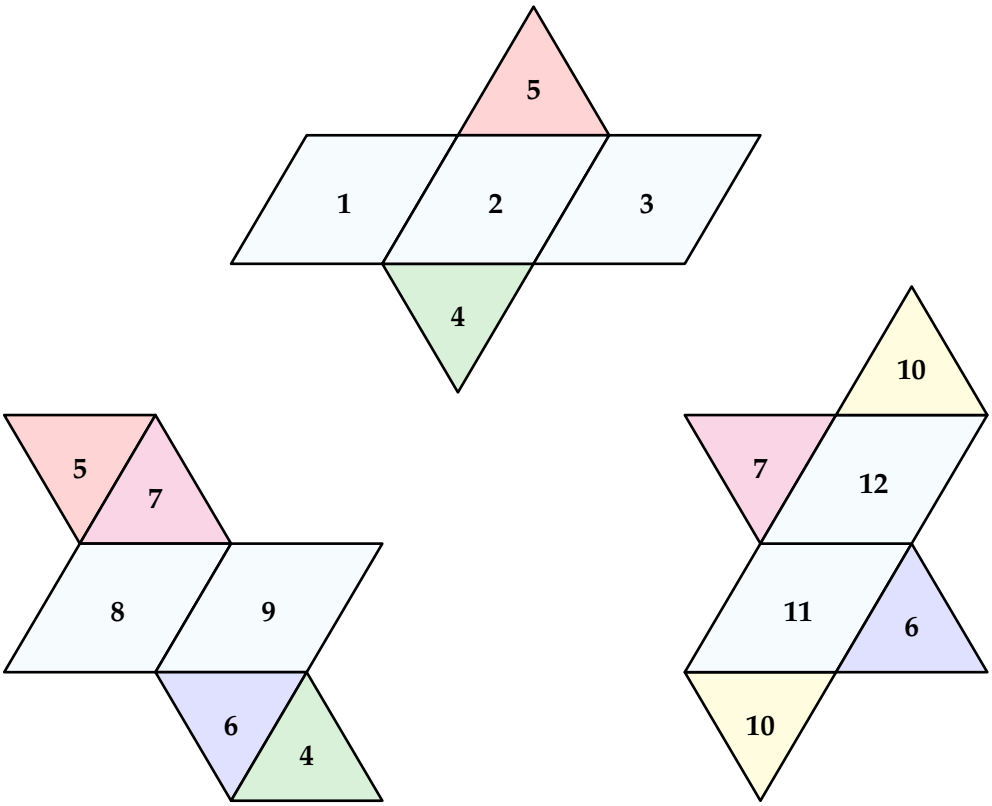

FIGURE 3.6: The octahedron unfolding with π_{Π}

This will yield $|\mathrm{Mon}(\pi_{\Pi})| = 12$. Now, since any face of the octahedron can be rotated in three different ways, and there are 8 total faces, we know that

$$|\mathrm{Rot}(D)| = 8 \cdot 3 = 24 \neq 12 = |\mathrm{Mon}(\pi_{\Pi})|.$$

This result is exactly what we should expect, given that Theorem 3.2.9 does not apply to the octahedron. Indeed, there is no isomorphism between $\mathrm{Rot}(D)$ and $\mathrm{Mon}(\pi_{\Pi})$.

Chapter 4

Extension to Non-Orientable Surfaces

4.1 Main Idea

In any polygonal presentation of a non-orientable Platonic surface, there exists at least one edge for which the two corresponding faces lie on the same side. The Klein bottle, illustrated in Figure 2.2, provides an example, as the blue-tinted surface appears to the right of both red arrows. This contrasts with the torus, also shown in Figure 2.2, where for each arrow the surface lies once to the left and once to the right, thus establishing a well-defined directional flow.

To resolve this issue, we first consider the net of the non-orientable Platonic surface N and reflect the polygon across one of its edges. This reflection creates two new edges, each lying on the opposite side of the surface from its original counterpart. Identifying these edges with their reflections, we obtain a clear directional flow. For edges that were already identified in a clear manner, we preserve the original identifications, while the two new edges in the reflected copy are identified with each other, see Figure 4.1 for an example.

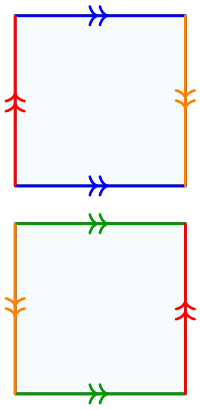

FIGURE 4.1: The double cover of the Klein bottle

This process yields a two-fold cover of the non-orientable Platonic surface, which is now orientable (as will be proven in Section 4.2). We denote this surface by S. Then, we can apply all the previously developed theory and treat the new orientable Platonic surface exactly as before. It should be noted, however, that the Klein bottle

is not a Platonic surface. This example is included only to illustrate the idea of reflecting along an edge to resolve non-orientability. Furthermore, we will not always indicate exactly which edges are identified with each other using distinct colors, as this is directly implied by the side on which the face lies relative to the edge.

4.2 Non-orientable Theorems

In this section, we will formally define the required objects and prove some of the most essential statements needed for the examples discussed in section 4.3.

Definition 4.2.1. Let S be the orientable double cover of a non-orientable Platonic surface N, with the covering map $\eta : S \to N$. The double cover is defined by reflecting the polygons and identifying their edges accordingly. Let r denote the reflection that sends each polygon to its mirror image. If, in N, the edges e_1 and e_2 are glued, then we glue e_1 to $r(e_2)$ and e_2 to $r(e_1)$ when the polygons lie on the same sides of the edges. Otherwise, if the polygons already lie on the opposing sides, we glue e_1 to e_2 and $r(e_1)$ to $r(e_2)$.

Lemma 4.2.2. *The double cover S of a non-orientable Platonic surface N is itself a Platonic surface.*

Proof. We need to show that $\mathrm{Aut}(S)$ acts transitively on its flags. Consider the flag (v, e, f). Let us define three isometries:

$$\mu_v : (v, e, f) \mapsto (v', e, f)$$

$$\mu_e : (v, e, f) \mapsto (v, e', f)$$

$$\mu_f : (v, e, f) \mapsto (v, e, f')$$

where μ_v changes the vertex but fixes the edge and face, μ_e changes the edge but fixes the vertex and face, and μ_f changes the face but fixes the vertex and edge. Thus, any one of these operations moves to an adjacent flag. We can do this by the regularity of the faces. Now, by a sequence of moves, we can reach any flag on S. Since our faces are regular p-gons, we can rotate to any face by compositions of μ_e and μ_f. Thus, there exists an automorphism that is formed by composition of μ_v, μ_e, μ_f that maps (v, e, f) to (v', e', f'), which proves the lemma. □

Lemma 4.2.3. *The automorphism group of a non-orientable Platonic surface* $\mathrm{Aut}(N)$ *is isomorphic to the rotation subgroup of the double cover* $\mathrm{Rot}(S)$.

Proof. There exists a deck transformation $d : S \to S$, such that $d^2 = \mathrm{id}$, $\eta \circ d = \eta$, and d reverses orientation. We define ψ as follows. Let $\psi : \mathrm{Rot}(S) \to \mathrm{Aut}(N)$. Take $g \in \mathrm{Rot}(S)$ and define $\psi(g)(x) := \eta(g(\tilde{x}))$, where $\tilde{x} \in S$ is any lift of $x \in N$.
If $\tilde{x}$ and $d(\tilde{x})$ are two lifts, then

$$\eta(g(d(\tilde{x}))) = \psi(g \circ d)(\tilde{x}) = \psi(d \circ g)(\tilde{x}) = \eta(g(\tilde{x})).$$

Since g commutes with d, the second equality holds. Thus, ψ is well-defined. Furthermore, ψ is a homomorphism, since $\forall g_1, g_2 \in \mathrm{Rot}(S)$, $\psi(g_1 g_2) = \psi(g_1) \circ \psi(g_2)$.
Now, assume $\psi(g) = \mathrm{id}$. Then, $\forall \tilde{x} \in S$, $\eta(g(\tilde{x})) = \eta(\tilde{x})$. This implies that either $g(\tilde{x}) = \tilde{x}$, or $g(\tilde{x}) = d(\tilde{x})$. But, since g preserves orientation and d reverses it, we

know that $g = \mathrm{id}$. Therefore, ψ is injective.
Next, take $h \in \mathrm{Aut}(N)$. There exists a lift $\tilde{h} : S \to S$, such that $\eta \circ \tilde{h} = \tilde{h} \circ \eta$. Also, there exist two types of lifts; $\tilde{h}$ and $d \circ \tilde{h}$, i.e. the orientation preserving and reversing ones. Choose $\tilde{h} \in \mathrm{Rot}(S)$, then $\psi(\tilde{h}) = h$. Thus, ψ is surjective, which proves that $\mathrm{Aut}(N) \cong \mathrm{Rot}(S)$. □

Definition 4.2.4. We define the unfolding $\tilde{N}$ of a non-orientable Platonic surface N to be the same as the unfolding $\tilde{S}$ of its orientable double cover S, with the branched covering map $\pi_S : \tilde{S} \to S$. We then define $\pi_N := \eta \circ \pi_S$, which produces a branched covering map $\pi_N \colon \tilde{N} \to N$.

Analogously to the orientable section, this yields a natural branched covering map $\pi_\Pi : \tilde{S} \to \Pi_p$. Thus, we have the following diagram:

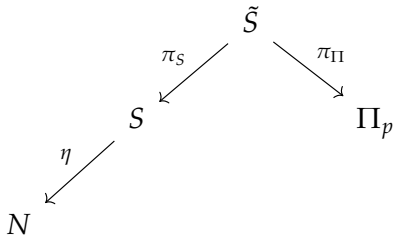

Theorem 4.2.5. *The monodromy group* $\mathrm{Mon}(\pi_\Pi)$ *is isomorphic to a normal subgroup of* $\mathrm{Aut}(N)$.

Proof. By Lemma 4.2.2, S is a Platonic surface. Hence, we can apply Theorem 3.1.8, which gives an isomorphism between $\mathrm{Mon}(\pi_\Pi)$ and a normal subgroup of $\mathrm{Rot}(S)$. Moreover, by Lemma 4.2.3, we have $\mathrm{Rot}(S) \cong \mathrm{Aut}(N)$, so the result follows. □

4.3 Examples

4.3.1 Hemicube

The first non-orientable Platonic surface we consider is the *hemicube*. A *hemicube* is a regular polyhedron obtained by slicing a cube with a plane through two opposite corners and the midpoints of two edges. It can also be visualized by removing three of the six faces of a cube and identifying opposite edges with reversed orientation. This construction yields a surface with 4 distinct vertices, 3 faces, and 6 edges, 3 on the inside, and 3 on the outside, as shown in Figure 4.2.

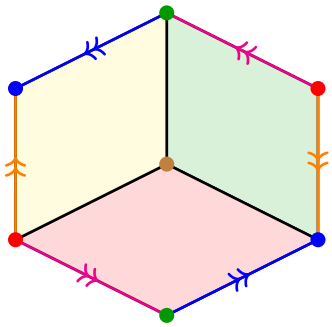

FIGURE 4.2: The hemicube

Next, we consider the net of the hemicube. Focusing on the edge between the green and red faces, we can cut the hemicube open along this edge to flatten it. We then identify the separated edge, now denoted by the green edge. We also label the faces by numbers, see Figure 4.3.

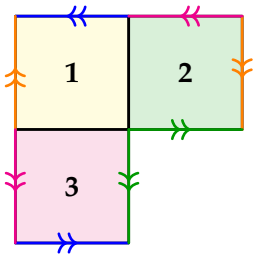

FIGURE 4.3: The net of the hemicube

Following our theory, we need to reflect the net along one of its edges. We chose the orange edge of the second face. We will also continue to label the faces, see Figure 4.4. After this reflection, we need to differentiate the edges so that the identifications are well-defined. Thus, we modify some of the edge colorings to ensure that for each edge, the face appears once on its left and once on its right.

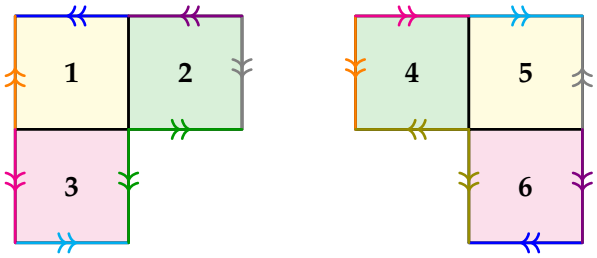

FIGURE 4.4: Orientable degree-2 cover of the hemicube

Now, we have an orientable Platonic surface. Next, we need to consider the unfolding of the orientable Platonic surface, which we will simply call the unfolding of the hemicube, see Figure 4.5. This will act as a cover of Π_4, as we have previously done in Figure 3.1. Note that the disconnected Platonic surface shown in Figure 4.4 can be glued together along the identified edge to produce a valid net of the cube itself. For instance, we can glue the third face to the fifth face along the light blue edge, which

produces one of the 11 valid cube nets. The calculations for the monodromy group are then analogous to those for the cube, but we include them for completeness.

```
gen1_hemicube = [(1,2,14,15), (3,22,13,21), (4,5,16,17),
   (6,10,18,9), (7,8,11,12), (19,20,24,23)]
gen2_hemicube = [(1,6,5,3), (2,19,4,12), (7,22,23,9), (8,16,24,15),
    (10,20,21,11), (13,17,18,14)]
GMon_hemicube = PermutationGroup([gen1_hemicube, gen2_hemicube])
GMon_hemicube.order()
```

LISTING 4.1: Calculations for the hemicube monodromy group

This then yields $|\mathrm{Mon}(\pi_{\Pi})| = 24$. By Theorem 3.2.9, we can conclude that the order of the rotation group is also 24. Intuitively, this also makes sense, as the hemicube has three faces that can each be rotated in 4 different ways, and since there is no sense of orientation, we may also apply reflections, giving $3 \cdot 4 \cdot 2 = 24$.

4.3.2 Hemi-Octahedron

Similarly to the hemicube, the *hemi-octahedron* is a non-orientable Platonic surface that contains half the faces of the octahedron. It consists of 3 triangular faces, 3 distinct vertices, and 6 edges, 3 outside, and 3 inside, see Figure 4.6.

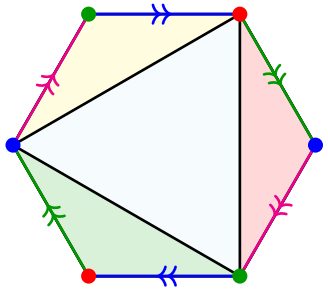

FIGURE 4.6: The hemi-octrahedron

Once again, we first consider the net of the hemi-octahedron.

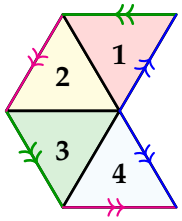

FIGURE 4.7: The hemi-octahedron net

Next, we reflect the net along a chosen edge. Specifically, we reflect along the edge at the bottom of the fourth face.

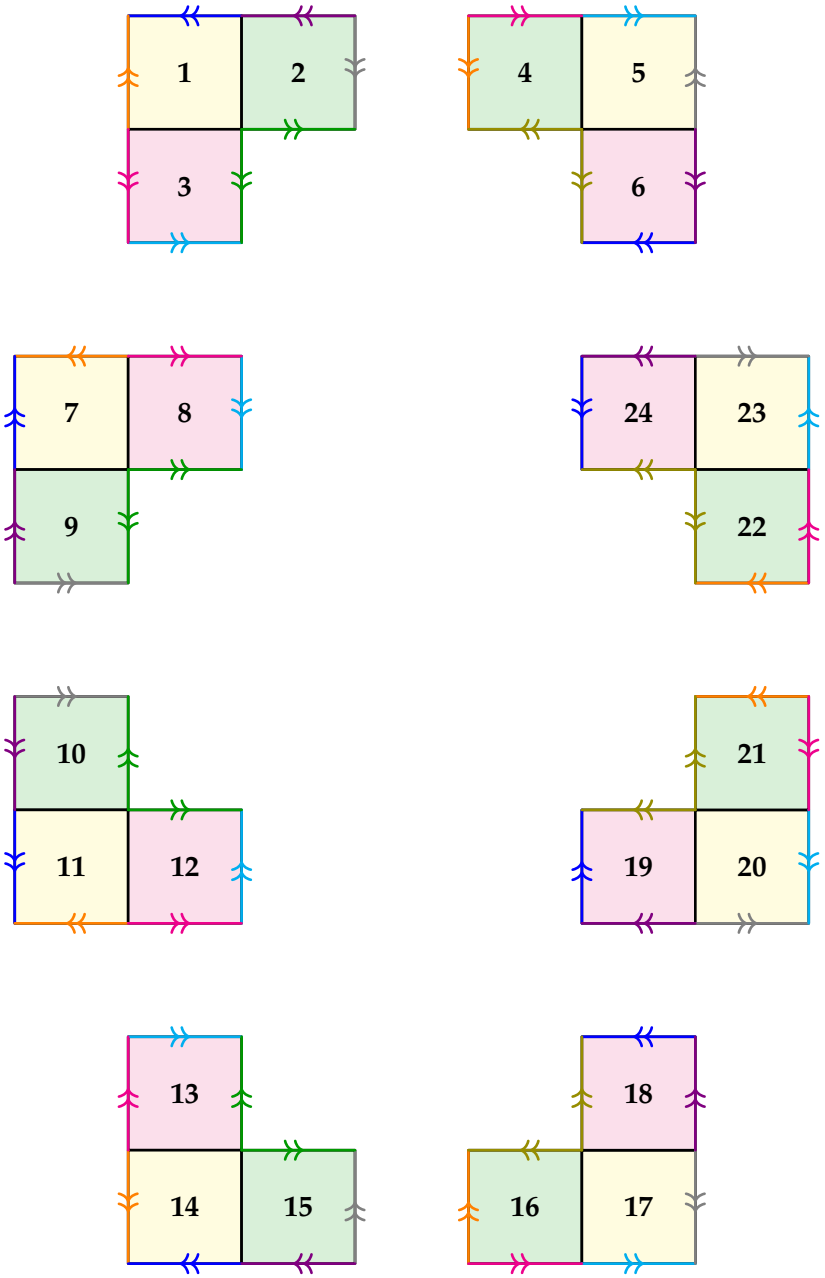

FIGURE 4.5: The hemicube unfolded

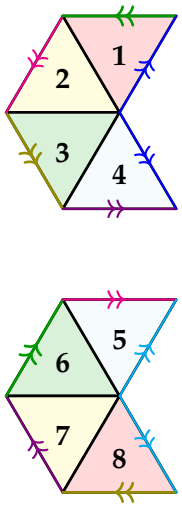

FIGURE 4.8: Orientable 2-degree cover of the hemi-octahedron

Analogous to the octahedron, the holonomy of a closed curve on the 2-fold cover of the hemi-octahedron is $\frac{5}{3}\pi$. Here again, considering 3 rotations of the Platonic surface suffices in place of 6, as the two sets of cycles we would get are disjoint.

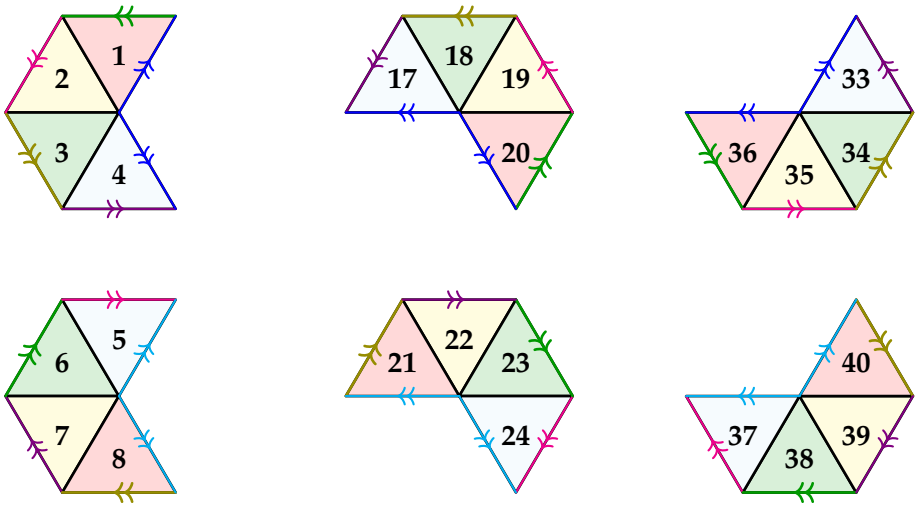

FIGURE 4.9: The hemi-octahedron unfolding with π_D

Next, we need to take the hemi-octahedron unfolding as a cover of π_Π in order to compute the order of the monodromy group.

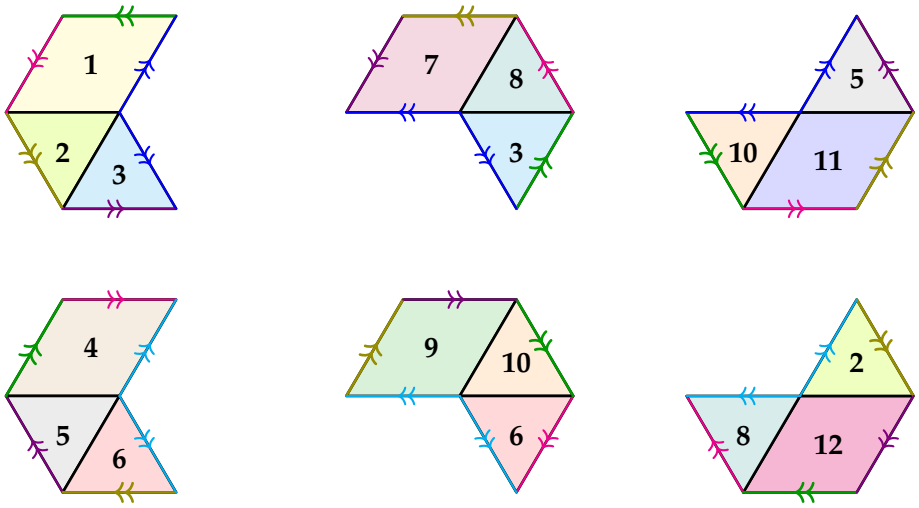

FIGURE 4.10: The hemi-octahedron unfolding as a cover of Π_3

For completeness, we include the Sage code for the computations. The order of the monodromy group is 12.

```
gen1_hemiocta = [(1,5,6), (2,3,4), (7,8,12), (11,9,10)]
gen2_hemiocta = [(1,12,2), (3,8,9), (7,6,10), (11,5,4)]
GMon_hemiocta = PermutationGroup([gen1_hemiocta, gen2_hemiocta])
GMon_hemiocta.order()
```

LISTING 4.2: Calculations for the hemi-octahedron monodromy group

4.3.3 Non-orientable Surface N4:$\{4,6\}_6$

Both the hemicube and the hemi-octahedron are non-orientable surfaces of genus 1. Our next and final example is the non-orientable surface N4:$\{4,6\}_6$, which is of genus 4. It has 4 distinct vertices, 6 faces, and 12 edges, with Schläfli symbol $\{4,6\}$. Thus, we have 6 distinct square faces, and at every vertex all 6 faces meet.

Figure 4.11 presents the surface in a two-dimensional polygonal representation. Because 6 squares meet at a single vertex, we employ *crosscaps* at the unavoidable intersections of the faces, denoted by purple and magenta circles. Under this construction, any two points directly across a crosscap from each other are identified. To avoid using crosscaps, we can consider the net of N4:$\{4,6\}_6$, see Figure 4.12

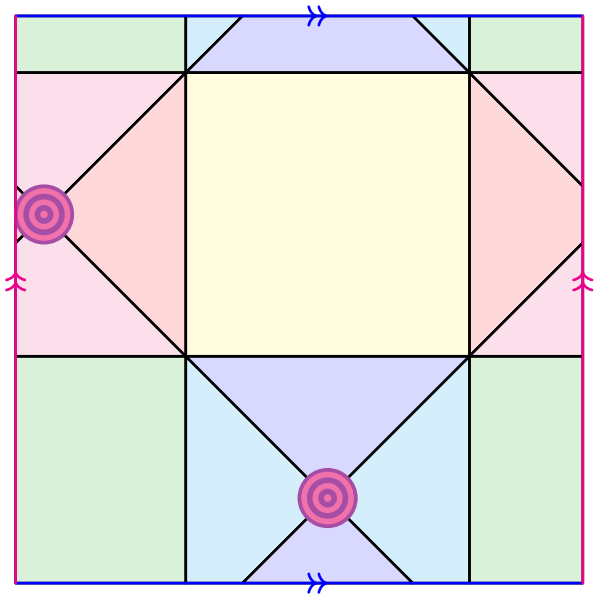

FIGURE 4.11: Non-orientable surface N4:$\{4,6\}_6$

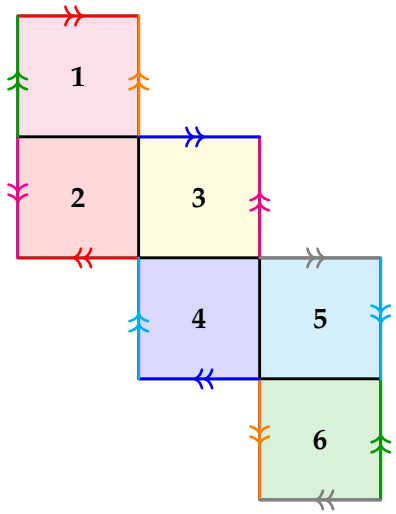

FIGURE 4.12: The net of N4:$\{4,6\}_6$

Next, we will reflect the net along a chosen edge to get an orientable Platonic surface, see Figure 4.13. Note that earlier, we recolored the arrows for which a face layed twice on the same side. Here, the identifications are more straightforward, so we will not draw this distinction.

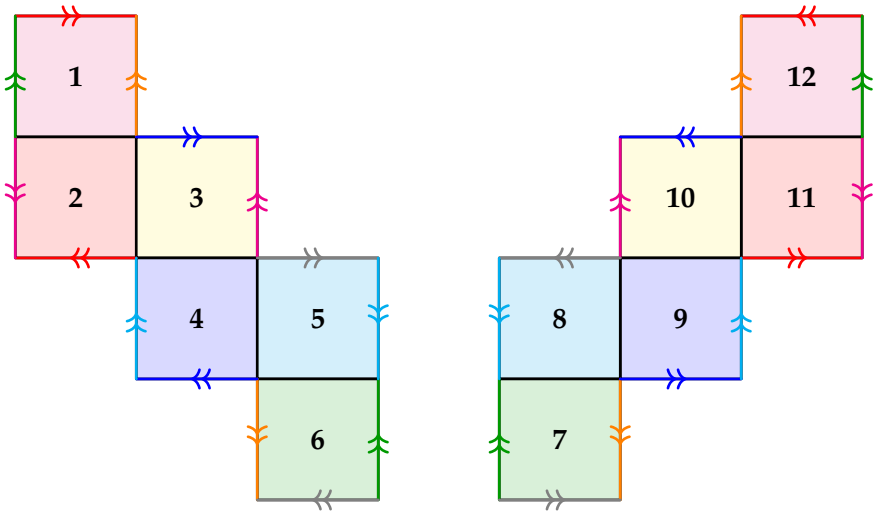

FIGURE 4.13: N4:$\{4,6\}_6$ as a Platonic Surface

As a Platonic surface, N4:$\{4,6\}_6$ has 8 distinct vertices. The cone angle for any closed loop around one of the vertices is 3π, which makes the holonomy equal to -1π. Therefore, we need a degree-2 cover to get the unfolding, see Figure 4.14. The Schläfli formula for this Platonic surface is also $\{4,6\}$, so we cannot utilize Theorem 3.2.9, as $\gcd(4,6) \neq 1$. This is a degree-24 cover of Π_4, the torus.

Finally, we once again trace the action of the generators of Π_4 in the unfolding, by taking the lifts of h and v, respectively. Denoting all the cycles and computing the monodromy group will give an order of 24, see Listing 4.3.

```
gen1_N4 = [(1,13,18,6), (2,3,17,16), (4,5,15,14), (7,24,19,12),
    (8,9,20,21), (10,11,22,23)]
gen2_N4 = [(1,11,12,2), (3,9,10,4), (5,7,8,6), (13,20,19,14),
    (15,22,21,16), (17,24,23,18)]
GMon_N4 = PermutationGroup([gen1_N4,gen2_N4])
GMon_N4.order()
```

LISTING 4.3: Calculations for the N4:$\{4,6\}_6$ monodromy group

4.4 Alternative Approaches

Before ultimately arriving at an extension of the theory to non-orientable surfaces, we explored a variety of alternative approaches, most of which proved to be unsuccessful but nevertheless guided our understanding of the underlying constraints.

A large portion of these attempts was centered around *half-translation surfaces*. There, instead of restricting ourselves to identifying polygonal edges purely via translations, we also allowed identifications involving half-turn rotations. This additional flexibility initially appeared promising, as it naturally includes certain symmetries that arise in non-orientable contexts.

Moreover, we investigate constructions based on *non-orientable translation surfaces*, which are defined in a manner similar to classical translation surfaces, but with certain restrictions relaxed. Specifically, we define them via a polygon with an even number of sides, where each pair of corresponding sides is parallel, non-adjacent,

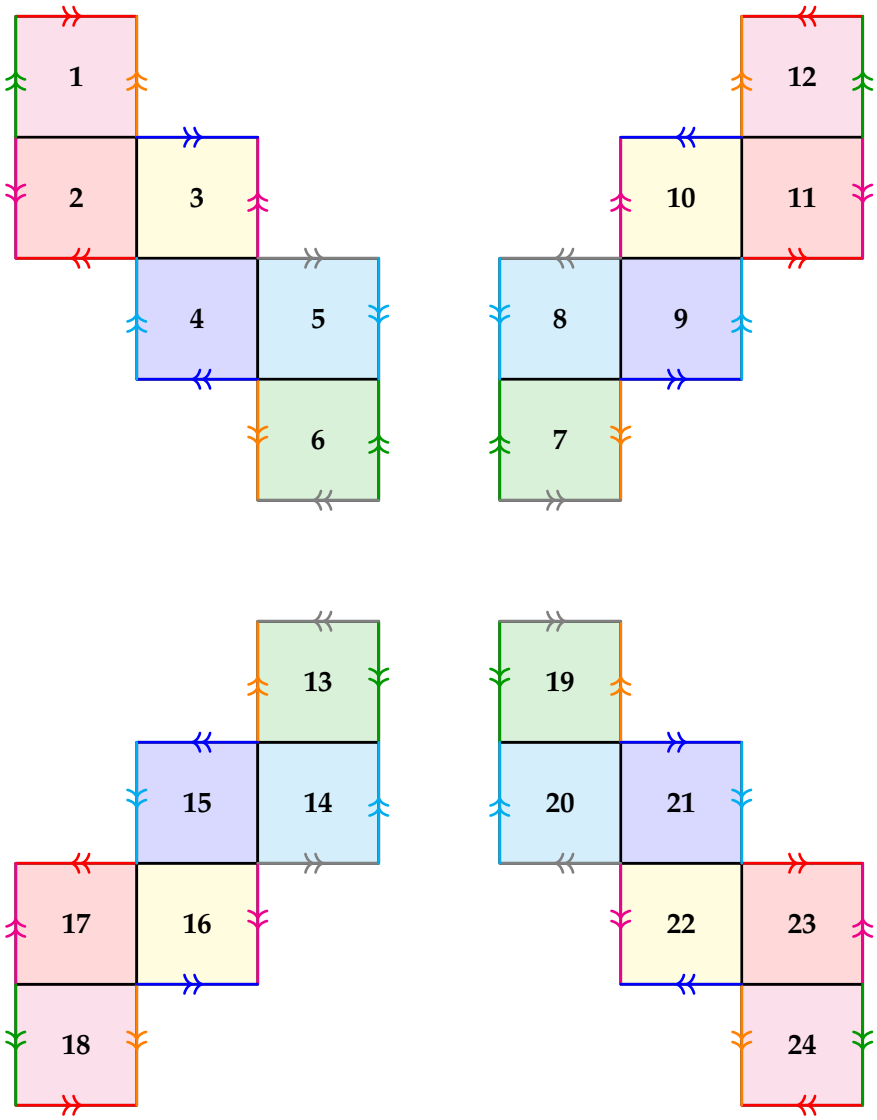

FIGURE 4.14: The unfolding of N4:$\{4,6\}_6$

and of equal length. Notably, we do not require the paired sides to have the same orientation, nor do we impose the condition that the polygon lies once to the left and once to the right of each pair of sides.

We used this to examine whether the fundamental polygon of the Klein bottle could serve as a suitable non-orientable translation surface for a covering of the hemicube. Unfortunately, this proved to be unsuccessful. Furthermore, we examined an alternative model in which the Klein bottle is realized as a quotient of two squares, see Figure 4.15, with the hope that this models a more traceable covering space of the hemicube.

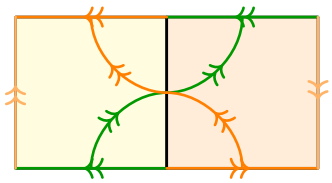

FIGURE 4.15: The Klein bottle made from two squares with its two generators

In fact, this construction worked. The unfolding can be viewed as a degree 6 cover of the Klein bottle constructed from two squares. The covering is branched over the red singularity (see Figure 4.16), while the pre-image of the gray singularity consists of the three remaining singularities, namely the green, blue, and yellow ones (see Figure 4.17). Note that, in this realization of the unfolding of the hemicube, it is necessary to consider four different symmetries of the Klein bottle formed from two squares in order for the unfolding to be a branched covering.

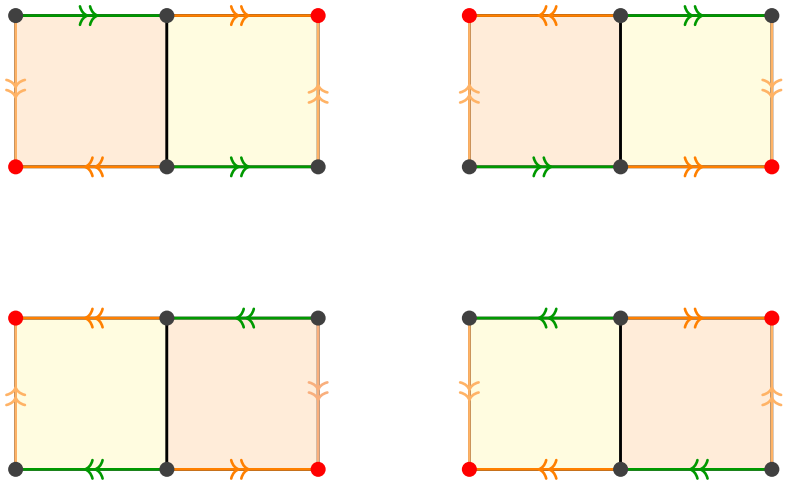

FIGURE 4.16: Four symmetries of the Klein bottle made from two squares

The fundamental group of the Klein bottle has two generators, so the same is true for the Klein bottle made from two squares, see Figure 4.15. If we trace the lift of one of the generators in the unfolding, we would get the green lift of figure 4.17, for example. The Deck transformations of π_Π fix the red vertex, and they descend to automorphisms of the hemicube that fix the red vertex. Thus, if we let D be the hemicube and R be the red vertex, then $\mathrm{Stab}_{\mathrm{Aut}(D)}(R) \cong \mathrm{Mon}(\pi_\Pi)$. This result is, of

course, less informative than the isomorphisms obtained previously. Nevertheless, it suggests that non-orientable translation surfaces may still serve as viable base spaces for the coverings, which could prove to be useful for future research.

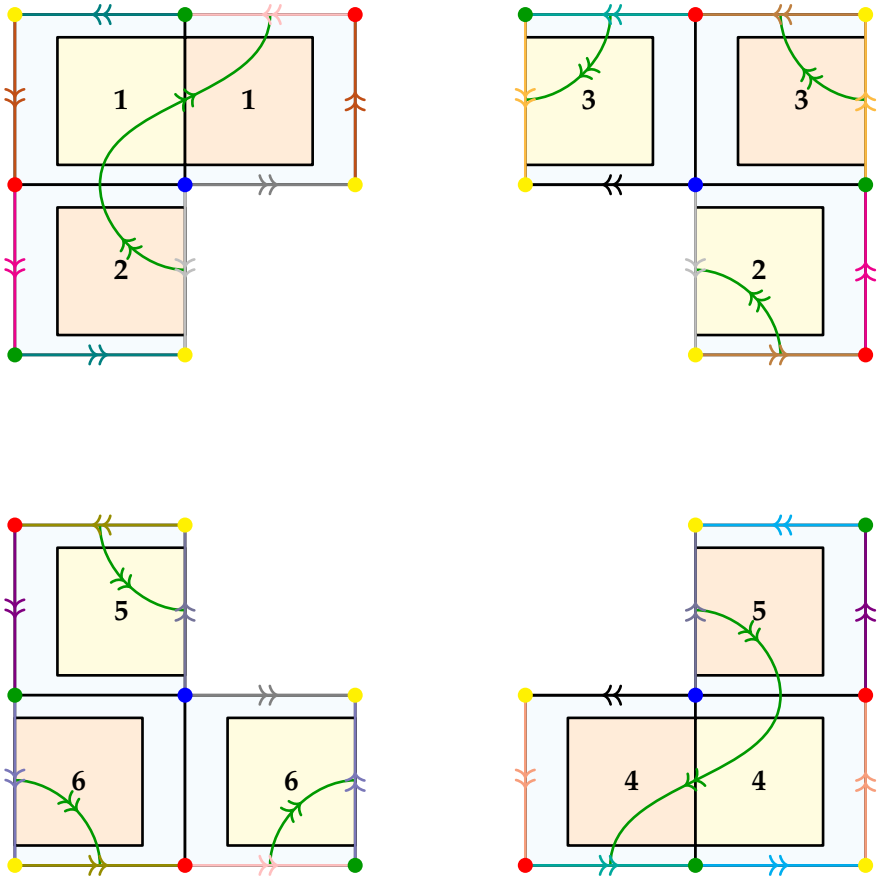

FIGURE 4.17: The hemicube unfolded as a cover of the Klein bottle made from two squares

Chapter 5

Conclusion

We developed a framework for studying Platonic surfaces through the interplay of covering theory and group actions. A central role in our approach is played by the construction of the unfolding, which allows us to associate to each Platonic surface a translation surface that serves as a natural covering space. More precisely, given a Platonic surface with Schläfli symbol $\{p, q\}$, we considered its unfolding as a branched cover of the associated regular p-gon tiled surface. This construction enables us to pass from a setting with non-trivial holonomy to one in which the holonomy becomes trivial.

This allows us to define and analyze the monodromy group $\mathrm{Mon}(\pi_\Pi)$, which encodes how loops on the base surface permute the sheets of the covering. One of the main results of this work is that this monodromy group is always isomorphic to a normal subgroup of the rotation group of the original Platonic surface. This establishes a direct and explicit link between the combinatorial data of the covering and the geometric symmetries of the surface. Moreover, under suitable conditions on the Schläfli symbol $\{p, q\}$, this relationship becomes even stronger. In these cases, the monodromy group is in fact isomorphic to the entire rotation group. This can be a powerful tool, as the monodromy group is typically easier to compute, whereas determining the full rotation group can be considerably more challenging.

We also extended the theory to include non-orientable Platonic surfaces. By constructing appropriate orientable double covers via reflections of their polygonal nets, we were able to reduce many questions in the non-orientable case to the orientable framework developed earlier. This extension demonstrates that the methods introduced in this thesis are sufficiently robust to handle a broader class of surfaces, and it suggests further avenues for exploration in the study of symmetry and coverings of more general geometric structures.

Bibliography

[Aul19] David Aulicino. "A New Approach to the Automorphism Group of a Platonic Surface". In: *Preprint* (2019). The author was partially supported by NSF DMS-1738381, DMS-1600360, and PSC-CUNY Grants 60571-00 48 and 61639-00 49. URL: https://arxiv.org/abs/1811.07052.

[Har00] Robin Hartshorne. *Geometry: Euclid and Beyond*. Springer-Verlag, 2000. ISBN: 0-387-98650-2. URL: https://doi.org/10.1007/978-0-387-22676-7.

[Dru19] Drummyfish. *Platonic Solids Transparent*. Wikimedia Commons. Online image; CC0 1.0 Universal Public Domain Dedication. 2019. URL: https://commons.wikimedia.org/w/index.php?curid=77742585.

[Via06] Marcelo Viana. "Ergodic Theory of Interval Exchange Maps". In: *Journal of Modern Dynamics* 1.2 (2006), pp. 1–54. URL: https://arxiv.org/abs/math/0508114.

[Jay24] Howard Masur Jayadev S. Athreya. *Translation Surfaces*. American Mathematical Society, 2024. URL: https://www.math.tau.ac.il/~barakw/translation_surfaces/Athreya_Masur.pdf.

[Hat01] Allen Hatcher. *Algebraic Topology*. Cambridge University Press, 2001. ISBN: 0-521-79160-X. URL: https://pi.math.cornell.edu/~hatcher/AT/AT.pdf.

[Mar22] Zach Marks. "Branched Coverings and Field Extensions of Compact Riemann Surfaces". In: *Preprint* (2022). URL: https://math.uchicago.edu/~may/REU2022/REUPapers/Marks.pdf.

[AAH19] Jayadev S. Athreya, David Aulicino, and W. Patrick Hooper. "Platonic solids and high genus covers of lattice surfaces". In: *Preprint* (2019). With an appendix by Anja Randecker. URL: https://arxiv.org/abs/1811.04131.

[Dyd11] Jerzy Dydak. *Coverings and fundamental groups: a new approach*. 2011. arXiv: 1108.3253 [math.GN]. URL: https://arxiv.org/abs/1108.3253.

[Bre93] Glen E. Bredon. *Topology and Geometry*. Springer Verlag, 1993. URL: https://virtualmath1.stanford.edu/~ralph/math215b/Bredon.pdf.