Kai-Uwe Bux

Groups and Spaces

Volume II: Spaces

June 2, 2007: Preliminary version, do not cite, do not distribute.

Preface

This volume contains those parts of topology that are essential to the understanding of geometric group theory.

Contents

Part I

Graphs, Trees, and $\mathbb{R}$-Trees

Chapter 1

Groups Acting on Trees

1.1 Graphs of Groups and Spaces

Definition 1.1.1. A graph of spaces $\mathcal{X}$ is a graph Γ together with the following data:

1. For every vertex v, we are given a group $\mathcal{X}_v$, called a vertex group.
2. For each unoriented edge e, we are given a group $\mathcal{X}_e$ called an edge group.
3. For every oriented edge $\vec{e}$, we have a π_1-injective continuous map

$$f_{\vec{e}} : \mathcal{G}_e \to \mathcal{G}_v$$

where e is the unoriented edge underlying $\vec{e}$ and v is the terminal vertex of $\vec{e}$.

Remark 1.1.2. From a category theoretical point of fiew, it might be more natural to think in terms of the first barycentric subdivision of the graph Γ. This is a graph too. All vertices in this graph are decorated by the vertex and edge spaces. In this subdivision, we orient all the edges such that they point away from the subdivision vertices. Now these edges can be decorated with the inclusions

$$f : \mathcal{X}_e \to \mathcal{X}_v.$$

Thus, we have constructed a diagram of groups all of whose morphisms are injections.

Note that diagrams over a given oriented graph form a category. Therefore, graphs of groups over Γ form a category. Of course, the same holds for graphs of spaces over Γ:

Definition 1.1.3. Let $\mathcal{X}$ be a graph of spaces over the graph Γ. As we have remarked in (1.1.2), a graph of spaces gives rise to a diagram of topological spaces. The homotopy colimit

$$|\mathcal{X}| := \text{ho-colim}(\mathcal{X})$$

of this diagram is called the geometric realization $\mathcal{X}$.

Now, this nonsence is way too abstract. So let us describe the construction in detail: We start with the disjoint union of all vertex and edge spaces. For each oriented edge $\vec{e}$, we construct the mapping cyclinder $M_{\vec{e}}$ of the map

$$f_{\vec{e}} : \mathcal{X}_e \to \mathcal{X}_v$$

and attach it to the spaces $\mathcal{X}_e$ and $\mathcal{X}_v$.

Note that there is a natural projection

$$\pi : |\mathcal{X}| \to \Gamma$$

sending the mapping cone parts to the corresponding half edges of Γ.

Lemma 1.1.4. *The realization $|\mathcal{X}|$ of a graph of spaces over a tree T all of whose edge and vertex spaces are m-connected is m-connected itself.*

Proof. Consider the projection

$$\pi : |\mathcal{X}| \to T$$

and note that every vertex and every edge in T has an m-connected preimage. Thus the claim follows from (**??**). **q.e.d.**

Exercise 1.1.5. Give less high-powered proof of (1.1.4) along the following lines: argue connectivity directly, apply van Kampen's theorem for simple connectivity, and use Mayer-Vietoris and Hurewicz for the higher homotopy groups.

Observation 1.1.6. *Let $\mathcal{X}$ be a graph of spaces over the graph Γ. Any covering space of $|\mathcal{X}|$ is the realization of a graph of spaces $\mathcal{Y}$ satisfying:*

1. *The underlying graph Δ of $\mathcal{Y}$ is a covering $\pi : \Delta \to \Gamma$.*

2. *Every vertex space Y_v is a covering space of the vertex space $\mathcal{X}_{\pi_v}$.*

3. *Every edge space Y_e is a covering space of the edge space $\mathcal{X}_{\pi_e}$.*

Theorem 1.1.7. *Let $\mathcal{X}$ be a graph of spaces. Suppose all edge and vertex spaces in $\mathcal{X}$ are Eilenberg-Maclane spaces. Then the realization $|\mathcal{X}|$ is an Eilenberg-Maclane space, too.*

Proof. We will construct the universal cover of $|\mathcal{X}|$ hands on and see that it is contractible. Let $v \in \Gamma$ be a vertex, and let $I(v)$ be the set of oriented edges in Γ pointing towards v. Put

$$Y_v := \biguplus_{\vec{e} \in I(v)} \mathcal{X}_e$$

and observe that the maps $\mathcal{X}_e \to \mathcal{X}_v$ induce a map $Y_v \to \mathcal{X}_v$. Now, we define the space

$$X_v := \mathrm{MC}(Y_v \to \mathcal{X}_v)$$

whose universal cover will be denoted by

$$\tilde{X}_v.$$

Note that $\tilde{X}_v$ is contractible since X_v deformation retracts onto the Eilenberg-Maclane space $\mathcal{X}_v$. Moreover, all the maps $f_{\vec{e}} : \mathcal{X}_e \to \mathcal{X}_v$ are π_1-injective and so are the inclusions $\mathcal{X}_e \hookrightarrow X_v$.

q.e.d.

Definition 1.1.8. A graph of groups $\mathcal{G}$ is a graph Γ together with the following data:

1. For every vertex v, we are given a group $\mathcal{G}_v$, called a vertex group.
2. For each unoriented edge e, we are given a group $\mathcal{G}_e$ called an edge group.
3. For every oriented edge $\vec{e}$, we have a monomorphism

   $$\iota_{\vec{e}} : \mathcal{G}_e \hookrightarrow \mathcal{G}_v$$

   where e is the unoriented edge underlying $\vec{e}$ and v is the terminal vertex of $\vec{e}$.

Observation 1.1.9. *The Eilenberg-MacLane space construction can be made functorial. The functor* $\mathrm{K}(-,1)$ *extends to a functor taking graphs of groups to graphs of spaces.* **q.e.d.**

Definition 1.1.10. The fundamental group of a graph of groups $\mathcal{G}$ is the fundamental group

$$\pi_1(|\mathrm{K}(\mathcal{G},1)|)$$

of the geometric realization of its associated graph of spaces.

Proposition 1.1.11. *Let $\mathcal{G}$ be a graph of groups. Then all its vertex and edge groups inject into $\pi_1(\mathcal{G})$.*

Proof. Let $\mathcal{X} := \mathrm{K}(\mathcal{G},1)$ be the graph of Eilenberg-Maclane spaces $\mathcal{X}_v := \mathrm{K}(\mathcal{G}_v,1)$ associated to $\mathcal{G}$. We will construct the universal cover of $|\mathcal{X}|$.

q.e.d.

1.2 Splitting Groups into Free Products

Free products arise from wedge sums of topological spaces: Let Y_1 and Y_2 be connected well-behaved topological spaces (e.g., CW-complexes) with base points. Then

$$\pi_1(Y_1 \wedge Y_2) = \pi_1(Y_1) * \pi_1(Y_2)\,.$$

Here we glue the spaces at their basepoints. Any closed loop in $Y_1 \wedge Y_2$ centered at the basepoint decomposes into an alternating product of loops in Y_1 and Y_2. Since any bounding disc of a null-homotopic loop also gets squeezed at the basepoint, one can construct bounding discs for all factors. Therefore, an alternating word with non-trivial factors cannot be trivial unless it is empty.

We stretch the basepoint of the wedge to an edge and construct the space

$$X := !!!PICTURE!!!$$

which is obviously a homotopy equivalent to $Y_1 \wedge Y_2$. Let us consider the universal cover of the space X. This cover $\tilde{X}$ looks decidedly treelike: We will find in $\tilde{X}$ many copies of $\tilde{Y_1}$ and $\tilde{Y_2}$ connected by edges:

$$\tilde{X} = !!!PICTURE!!!$$

Collapsing the chunks $\tilde{Y_1}$ and $\tilde{Y_2}$ to vertices, we obtain a tree together with an action of $\pi_1(X)$ which is inherited from the action of this group on $\tilde{X}$ by deck transformations. Let us collect some fact about this tree-action:

Proposition 1.2.1. *The free product $A * B$ acts on a simplicial tree such that the following hold:*

1. *The action is free and transitive on the set of edges.*
2. *The action is rigid, i.e., the stabilizer of an edge does not flip it. (This is clear, since stabilizers are trivial. This condition, however, will become more important later.)*
3. *There are precisely two orbits of vertices.*
4. *There is one edge with endpoints v and w such that*

$$\begin{aligned} A &= \mathrm{Stab}(v) \\ B &= \mathrm{Stab}(w) \end{aligned}$$

Moreover, A acts transitively on the edges emanating from v, and B acts transitively on the edges emanating from w.

Thus, we have found a way to construct a nice tree-action from a free product. Let us see, whether we can reverse this.

Proposition 1.2.2. *Let G act on the tree T such that the following hold:*

1. *The action on the set of edges is transitive.*
2. *The action on the set of edges is free.*
3. *There are precisely two orbits of vertices.*

Then group is a free product $G = A * B$ where A and B are stabilzers of two adjacent vertices.

Proof. Let e be a fixed edge with midpoint M and vertices v and w. Put $A := \mathrm{Stab}(v)$ and $B := \mathrm{Stab}(w)$.

Note that A acts transitively on the edges emanating from v: The action is transitive on edges and does not take w to v.

We define a length function on G as follows:

$$\ell(g) := d(M, gM)$$

We claim that for every non-trivial element $g \neq 1$, there is an element $h \in A \cup B$ such that

$$\ell(hg) < \ell(g)\,.$$

!!! PICTURE !!!

It follows that A and B generate G.

Finally, we show that any product of non-trivial elements alternatingly taken from A and B is non-trivial unless it is empty.

!!! PICTURE !!!

... **q.e.d.**

Free splittings of finitely generated groups are completely understood: The rank $\mathrm{rk}(G)$ of a finitely generated group G is the smalles number of elements that generate G.

Theorem 1.2.3 (Grushko). *If A and B are finitely generated groups, then*

$$\mathrm{rk}(A * B) = \mathrm{rk}(A) + \mathrm{rk}(B)\,.$$

Corollary 1.2.4. *The process of splitting a finitely generated group as a free product terminates.*

Theorem 1.2.5 (Kurosh). *Let*

$$\begin{aligned} G &= F * A_1 * \cdots * A_r \\ &= F' * B_1 * \cdots * B_s \end{aligned}$$

be two free decompositions of G *into a free group and freely indecomposible factors. Then* $r = s$ *and the factors can be ordered such that* A_i *is conjugate to* B_i*. Moreover, the free groups have equal ranks.*

Chapter 2

Groups Acting on $\mathbb{R}$-Trees

Part II

Surfaces

Chapter 3

Preview

This lecture will deal predominantly with closed orientable surfaces, i.e., one of the following:

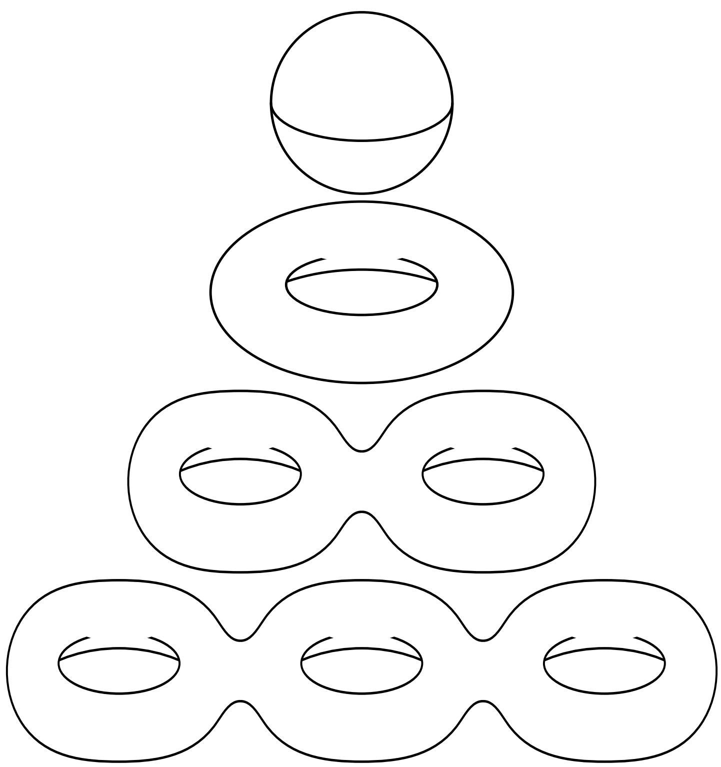

etc.

So let Σ be one of these. We will:

1. Classify all curves in Σ

- up to homotopy.
- up to isotopy (for simple curves).[1]

We shall see that, for surfaces, there is no difference in these two notions. We are lead to the fundamental group $\pi_*(\Sigma)$.

2. Classify all homeomorphisms $\zeta : \Sigma \to \Sigma$

   - up to homotopy.
   - up to isotopy.[2]

   This leads us to the mapping class group $\mathcal{M}(\Sigma)$.

3. We will show that $\mathcal{M}(\Sigma) = \mathrm{Out}(\pi_*(\Sigma))$.

To state our other goals, let us briefly discuss the torus.

Fact. *Let $T = T_2$ be the two-dimensional torus.*

1. $\pi_*(T) = C_\infty \times C_\infty$.
2. $\mathcal{M}(T) = \mathrm{GL}_2(\mathbb{Z})$.

Elements $M \in \mathrm{GL}_2(\mathbb{Z})$ can be sorted according to trace. There are three cases

$|\mathrm{tr}(M)| < 2$ *(elliptic): In this case, M has finite order. The corresponding homeomorphisms are called finite.*

$|\mathrm{tr}(M)| = 2$ *(parabolic): In this case, there is a homeomorphism $\zeta : T \to T$ represented by M that leaves a simple closed curve on T fixed. The homeomorphisms in this isotopy class are called reducible*

To see that there is a fixed curve, we consider the universal cover $\mathbb{R} \times \mathbb{R}$. Here the matrix defines a linear map $M : \mathbb{R}^2 \to \mathbb{R}^2$ which preserves the lattice $\mathbb{Z}^2$. Therefore the map directly descends to the map on T. Since the trace has absolute value 2, the matrix M has a double eigenvalue at 1 or -1. The corresponding eigenspace has rational slope and therefore descends to a closed curve on T, which is fixed.

[1] An isotopy of an embedding is a homotopy that stays an embedding during the deformation.
[2] An isotopy of a homeomorphism is a homotopy that stays a homeomorphism.

$|\mathrm{tr}(M)| > 2$ *(hyperbolic): In this case, the matrix M represents a so called Anosov-homeomorphism. The matrix M has two different eigenvalues. The corresponding eigenspaces have irrational slope. Thus they descend to a pair of space filling curves on T that intersect transversally. This is the prime example of a pair of transverse foliations. The homeomorphism is stretching in the direction of one foliation and shrinking with respect to the other one.*

Our first main goal will be to extend this to other surfaces:

The Nielsen-Thurston Classification. *Let Σ be a higher genus surface. Then every homeomorphism is either finite, i.e., isotopic to a homeomorphism of finite order, or reducible, i.e., isotopic to a homeomorphism that leaves a multi-curve fixed, or Pseudo-Anosov to be defined later. Roughly speaking, a Pseudo-Anosov homeomorphism locally looks like a map*

$$(x, y) \mapsto (Cx, \frac{y}{C}).$$

To achieve this goal, we will need to introduce the Teichmüller space of Σ. The mapping class group acts on Teichmüller space, and it is a good understanding of this action that will allow us to prove the Nielsen-Thurston classification.

Another space upon which the mapping class group acts is the Curve Complex. This is a simplicial complex whose vertices are simple closed curves on Σ. A set of those forms a simplex, if they do not intersect pairwise.

We will study these and other spaces to prove various result about the mapping class group. In particular, we will see that it is finitely generated.

Along the way, we will have to prove many classical results of planar geometric topology:

1. The Jordan Curve Theorem: A simple closed curve separates the plane into two regions one of which is bounded.

2. Schönflies' Theorem: The bounded region in the Jordan Curve Theorem actually is a disc.

3. The Hauptvermutung for surfaces: Given a finite system of arcs in a surface, there always is a homeomorphism that takes them simultaneously to polygonal arcs – here polygonal is meant to be detected in the universal cover: either $\mathbb{R}^2$ or the hyperbolic plane.

Chapter 4

The Jordan-Schönflies Theorem

The goal of this chapter is to prove the following two fundamental theorems about the topology of the plane.

Theorem 4.2.3 (Jordan-Veblen). *A simple closed curve separates the sphere into exactly two regions.*

Theorem 4.3.7 (Schönflies-Brouwer). *These two regions are discs.*

Marie Ennemond Camille Jordan[1] was the first to realize that the separation theorem actually requires proof. The one he gave in the third edition of his calculus textbook was, however, false. The first correct proof is due to Oswald Veblen.

4.1 Geodesic Homology

Heresy 4.1.1. *Some maps are not continuous. We call them functions.*

Let U be an open subset of $\mathbb{E}^2$ or S_2.

4.1.1 Dimension 0

Definition 4.1.2. A 0-cycle in U is a $\mathbb{R}$-valued function $\zeta^0 : U \to \mathbb{R}$ that vanishes at all but finitely many points. A 0-cycle is reduced, if

$$\sum_{x \in U} \zeta^0(x) = 0.$$

[1] This Jordan was French. He is the Jordan in Jordan-Hölder and the Jordan Normal Form. Wilhelm Jordan is the one from the Gauss-Jordan Elimination Method. Jordan Algebras finally are named after Pascual Jordan, a German Physicist.

It is obvious that the set of all 0-cycles is an $\mathbb{R}$-vector space.

There is an obvious way to generate reduced 0-cycles: Take an oriented geodesic segment and assign -1 to its starting point and 1 to its terminal point. We call a 0-cycle a 0-boundary if it is a finite linear combination of such geodesic cycles. Obviously, the boundaries form a linear subspace.

We call to 0-cycles equivalent if they differ by a boundary. Since the boundaries form a linear subspace, the set

$$\mathrm{H}_0(U)$$

of equivalence classes of 0-cycles is a $\mathbb{R}$-vector space. Since boundaries are always reduced cycles, we can also form the vector space

$$\tilde{\mathrm{H}}_0(U)$$

of equivalence classes of reduced cycles.

The vector spaces $\mathrm{H}_0(U)$ and $\tilde{\mathrm{H}}_0(U)$ are called the homology and reduced homology of U in dimension 0.

The homology of a space is intimately related to its connected components:

Proposition 4.1.3. *Let R be a system of points representing the components of U, i.e., there R contains precisely one point from each component. Then the set*

$$\left\{ \left[\zeta_x^0 : y \mapsto \begin{cases} 1 & \text{if } y = x \\ 0 & \text{if } y \neq x \end{cases} \right] \;\middle|\; x \in R \right\}$$

is a basis of $\mathrm{H}_0(U)$.

Proof. We have to prove that each 0-cycle is equivalent to a function with support in R.

We consider a 0-cycle as an assignment of charges to finitely many points. For each of these points, it follows from (A.2.2) that we can use a broken geodesic to transfer the charge to a unique point in R. **q.e.d.**

Corollary 4.1.4. *If $\mathrm{H}_0(U)$ has finite dimension, then this dimension equals the number of components of U. Thus, we have $|\{\text{components of } U\}| = \dim_{\mathbb{R}}(\mathrm{H}_0(U)) = \dim_{\mathbb{R}}\left(\tilde{\mathrm{H}}_0(U)\right) + 1$.*

Exercise 4.1.5. Let $\mathrm{H}^0(U)$ be the $\mathbb{R}$-vector space of all locally constant $\mathbb{R}$-valued functions on U. Prove that $\mathrm{H}^0(U)$ is the dual vector space of $\mathrm{H}_0(U)$.

4.1.2 Dimension 1

Definition 4.1.6. A 1-cycle on U is a "balanced flow with support in a finite geodesic graph". More precisely:

A flow on a graph is a $\mathbb{R}$-valued function on the set of oriented edges such that the value of an oriented edge and its reverse always add up to 0. A flow is balanced if it adds up to 0 on the outgoing edges around any vertex; it is non-trivial if no oriented edge is assigned 0.

A geodesic graph is a graph that is embedded in U such that all edges are geodesic segments. Such a graph is minimal if it does not contain isolated vertices or vertices of degree 2 that span an angle of π.

Now, we can say, a 1-cycle is a (possibly empty) minimal geodesic graph that carries a non-trivial balanced flow.

Remark 4.1.7. There is an obvious way of adding 1-cycles by super-position of the corresponding graphs and flows. Of course, one might have to introduce new vertices to keep the graphs embedded, where edges overlap, the flows are added, indeed cancellations might result in the deletion edges.

Moreover, we can multiply 1-cycles by scalars. Thus, the set of all 1-cycles carries a the structure of a real vector space.

Definition 4.1.8. There is an obvious way of creating 1-cycles. Take a geodesic triangle or any other convex geodesic polygon and flow with unit speed counter clockwise around its boundary. Finite linear combinations of these flows are called 1-boundaries.

Two 1-cycles are equivalent if they differ by a boundary. The set of equivalence classes is a real vector space, denoted by

$$\mathrm{H}_1(U)$$

and called the homology of U in dimension 1.

4.1.3 The Mayer-Vietoris Sequence

Let U and V be two open subsets of S_2 or $\mathbb{E}^2$. Put

$$\begin{aligned} X &:= U \cup V \\ Y &:= U \cap V \end{aligned}$$

Given elements $\zeta \in \mathrm{H}_i(U)$ and $c' \in \mathrm{H}_i(V)$, we can take the difference $\zeta - c'$ in $\mathrm{H}_i(X)$. The reason is that super-positioning pictures in U and V will not exceed the union X. Thus, we defined two maps:

$$\begin{aligned} \mathrm{H}_0(U) \oplus \mathrm{H}_0(V) &\to \mathrm{H}_0(X) \\ \mathrm{H}_1(U) \oplus \mathrm{H}_1(V) &\to \mathrm{H}_1(X) \end{aligned}$$

We can view a cycle in Y as a cycle in any open set containing Y: Just draw the picture in Y and realize that the same picture defines a cycle in bigger open sets too. Thus there are four obvious linear maps:

$$\begin{aligned}
\mathrm{H}_0(Y) &\rightarrow \mathrm{H}_0(U) \\
\mathrm{H}_0(Y) &\rightarrow \mathrm{H}_0(V) \\
\mathrm{H}_1(Y) &\rightarrow \mathrm{H}_1(U) \\
\mathrm{H}_1(Y) &\rightarrow \mathrm{H}_1(V)\,.
\end{aligned}$$

We combine:

$$\begin{aligned}
\mathrm{H}_0(Y) &\rightarrow \mathrm{H}_0(U) \oplus \mathrm{H}_0(V) \\
\mathrm{H}_1(Y) &\rightarrow \mathrm{H}_1(U) \oplus \mathrm{H}_1(V)\,.
\end{aligned}$$

So far, we have two sequences

$$\mathrm{H}_1(Y) \rightarrow \mathrm{H}_1(U) \oplus \mathrm{H}_1(V) \rightarrow \mathrm{H}_1(X)$$

and

$$\mathrm{H}_0(Y) \rightarrow \mathrm{H}_0(U) \oplus \mathrm{H}_1(V) \rightarrow \mathrm{H}_0(X)\,.$$

It is one of the most useful observations in topology that we can paste these sequences by means of a slick linear map

$$\mathrm{H}_1(X) \rightarrow \mathrm{H}_0(Y)\,.$$

Geometric Intuition 4.1.9. A 1-cycle in X is a balanced flow with support in a finite embedded graph. This is a network of wires with batteries moving charges around. A 0-cycle in Y is determined by one number per component of Y. So from the given flow, we have to cook up one number for each component in Y.

There is no source nor sink in the network. However, when we restrict our attention to a component of Y, we might find that the part of the network in this area has a certain net throughput from U toward V. It is precisely this throughput from U to V that we assign to the given component of Y. So in figure 4.1, the components in the intersection, from top to bottom, are assigned the numbers: -5, 3, 2, -5, and 5.

Formally, we subdivide our graph so that each edge is contained completely within one of the open sets U or V. Then the throughput of a component Y^0 in Y is the sum of the flows of those edges that lie in U and have precisely one of their endpoints in Y^0. Of course, signs have to be taken care of.

Table 4.1: The left two claws are U

Theorem 4.1.10. *With the linear maps as defined above, the* Mayer-Vietoris *sequences*

$$\mathrm{H}_1(Y) \to \mathrm{H}_1(U) \oplus \mathrm{H}_1(V) \to \mathrm{H}_1(X) \to \mathrm{H}_0(Y) \to \mathrm{H}_0(U) \oplus \mathrm{H}_0(V) \to \mathrm{H}_0(X)$$

and

$$\mathrm{H}_1(Y) \to \mathrm{H}_1(U) \oplus \mathrm{H}_1(V) \to \mathrm{H}_1(X) \to \tilde{\mathrm{H}}_0(Y) \to \tilde{\mathrm{H}}_0(U) \oplus \mathrm{H}_0(V) \to \tilde{\mathrm{H}}_0(X)$$

are exact.

Proof. We will deal with unreduced homology only. The arguments apply to reduced homology with no or little change. Also, we will only discuss exactness at the slots where the connecting map is involved. The two other positions are left as an exercise.

Exactness at $\mathrm{H}_0(Y)$**:** First we observe that any 1-cycle in $\mathrm{H}_1(X)$ dies in $\mathrm{H}_0(U) \oplus \mathrm{H}_0(V)$. The reason is that 1-cycles are balanced flows. Hence there is no source or sink in any connected component of U. Thus for each of these components, the through-puts of the bordering pieces in $U \cap V$ have to add up to 0. By the same argument, all components of V will be assigned 0.

So let us fix a 0-cycle $\zeta^0 \in \mathrm{H}_0(Y)$ concentrated in a set of representatives R for the components of Y. We assume that ζ^0 dies in $\mathrm{H}_0(U) \oplus \mathrm{H}_0(V)$. Let U^0 be a component of U. Since ζ^0 dies in $\mathrm{H}_0(U)$, the designated through-puts of all components in Y bordering U^0 will cancel out. So we fix a point x in U^0 and connect it to the bordering points of R. Each of these paths is assigned the flow required to generate the throughput. Since the values cancel out, the point x will not be a source. So we created part of a balanced flow. We do this for all components of U and all components of V. Finally we constructed a 1-cycle that matches the specifications.

Exactness at $\mathrm{H}_1(X)$**:** Let ζ^1 be the difference of two balanced flows with supports in U and V respectively. Since a balanced flow in U cannot possibly generate a non-trivial throughput to V in any of the components of Y, it follows that a linear combination of those flows cannot either. Thus, ζ^1 dies in $\mathrm{H}_0(Y)$.

On the other hand, suppose $\zeta^1 \in \mathrm{H}_1(X)$ dies in $\mathrm{H}_0(Y)$. If necessary, we subdivide the edges of ζ^1 to ensure that each edge is completely contained in U or V (or both). Pick your favorite component Y^0 of Y. Let us call those edges in U *transient* that have precisely one terminal vertex in Y^0. The flows of the transient edges cancel out since ζ^1 dies in $\mathrm{H}_0(Y)$. We choose a point $y \in Y^0$ and redirect the flow on the transient edges to pass through y. This is possible since Y^0 is path-connected.

We do this in all components of Y and use the points we constructed as "cut vertices" to separated the flow ζ^1 into two flows with support in U and V respectively. **q.e.d.**

Exercise 4.1.11. Show that the Mayer-Vietoris sequence is exact at the second and fifth term.

Definition 4.1.12. U is called <u>1-acyclic</u> if it is non-empty and $\mathrm{H}_1(U) = \tilde{\mathrm{H}}_0(U) = \{0\}$.

Example 4.1.13. Star shaped open regions are 1-acyclic.

Proof. Star-shaped means, there is a cone point that can be geodesically connected to any other point in the region. Thus any edge in a given flow spans a triangle with the cone point. Adding flows around those triangles will cancel everything. **q.e.d.**

Definition 4.1.14. A space X is called simply connected if every loop is homotopic to a constant loop. Equivalently, X is simply connected if every map $f : \mathbb{S}^1 \to X$ can be extended to a map $\tilde{f} : \mathbb{B}^2 \to X$.

A space is called 1-connected if it is non-empty, path-connected and simply connected.

Exercise 4.1.15. Suppose $X = U \cup V$ where U and V are open, and suppose U, V, and $U \cap V$ are all 1-connected. Prove that X is 1-connected.

Exercise 4.1.16. Call a planar rectangle *admissible* if at least one of its sides has an integer length, i.e., the height or width or both are integers. Show that only admissible rectangles can be tiled by admissible rectangles.

Corollary 4.1.17. *Suppose $X = U \cup V$ where U and V are open, and suppose U, V, and $U \cap V$ are all* 1*-acyclic. Then X is* 1*-acyclic.*

Proof. It is clear that X is non-empty and path-connected. It remains to prove that $\mathrm{H}_1(X) = 0$. We write down the Mayer-Vietoris sequence and observe that almost all terms vanish:

$$0 \to 0 \oplus 0 \to \mathrm{H}_1(X) \to 0 \to 0 \oplus 0 \to \mathrm{H}_0(X)$$

The claim is immediate. **q.e.d.**

Example 4.1.18. Let X be the complement of two points $x, y \in S_2$, then we have

$$\begin{aligned} \mathrm{H}_0(X) &= \mathbb{R} \\ \mathrm{H}_1(X) &= \mathbb{R}. \end{aligned}$$

Proof. Let U be the complement of the geodesic arc joining x and y. Let V be the open disc that has this arc as a diameter. Both of these regions are star-shaped and, therefore, 1-acyclic. Moreover, $Y := U \cap V$ has two components: The intermediate value theorem implies that on each path from the norther half disc to the southern half disc crosses height 0. Thus $\tilde{\mathrm{H}}_0(Y) = \mathbb{R}$. We are now ready to compute $\mathrm{H}_1(X)$ using the Mayer-Vietoris sequence. The relevant part is

$$0 \oplus 0 \to \mathrm{H}_1(X) \to \mathbb{R} \to 0,$$

and it follows that $\mathrm{H}_1(X) = \mathbb{R}$.

It is obvious that X is connected whence $\mathrm{H}_0(X) = \mathbb{R}$. **q.e.d.**

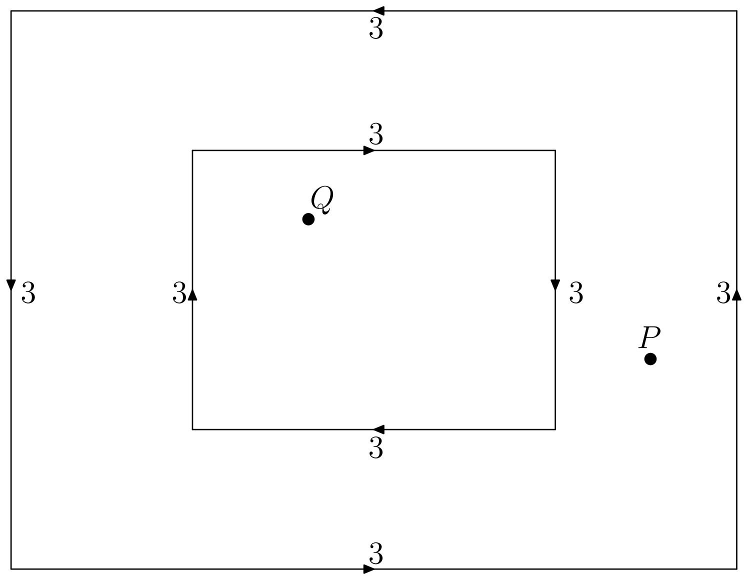

Table 4.2: A planar 1-cycle

4.1.4 The Carrier of a 1-cycle

We have seen that the plane is 1-acyclic. In particular, every 1-cycle is a 1-boundary. The proof was dumb: we did not have to be very smart to construct the geodesic triangles needed for the 1-boundary carefully. Let us try to do better. This will be necessary to answer the following natural question:

Problem 4.1.19. Given a 1-cycle ζ^1 in an open subset U of the plane, how can we decide whether ζ^1 represents the trivial element in $\mathrm{H}_1(U)$?

In talking about this problem, some additional terminology proves convenient:

Definition 4.1.20. A finite formal $\mathbb{R}$-linear combination of solid geodesic triangles is called a <u>2-chain</u>. If $\sum_i a_i \Delta_i$ is a two chain, we call $C := \bigcup_{a_i \neq 0} \Delta_i$ its <u>carrier</u>. There is an obvious way of associating a 1-boundary to any 2-chain. If a 1-cycle is the boundary of a certain 2-chain, we say, the 2-chain <u>bounds</u> the 1-cycle.

Example 4.1.21. Consider the planar 1-cycle in figure 4.2. It is obvious that we can write this flow as a boundary of a 2-chain whose carrier does avoid the point Q. It is, however, hard to imagine such a 2-chain avoiding P. Our problem is, how can we prove that every 2-chain that bounds this cycle has to have P in its carrier?

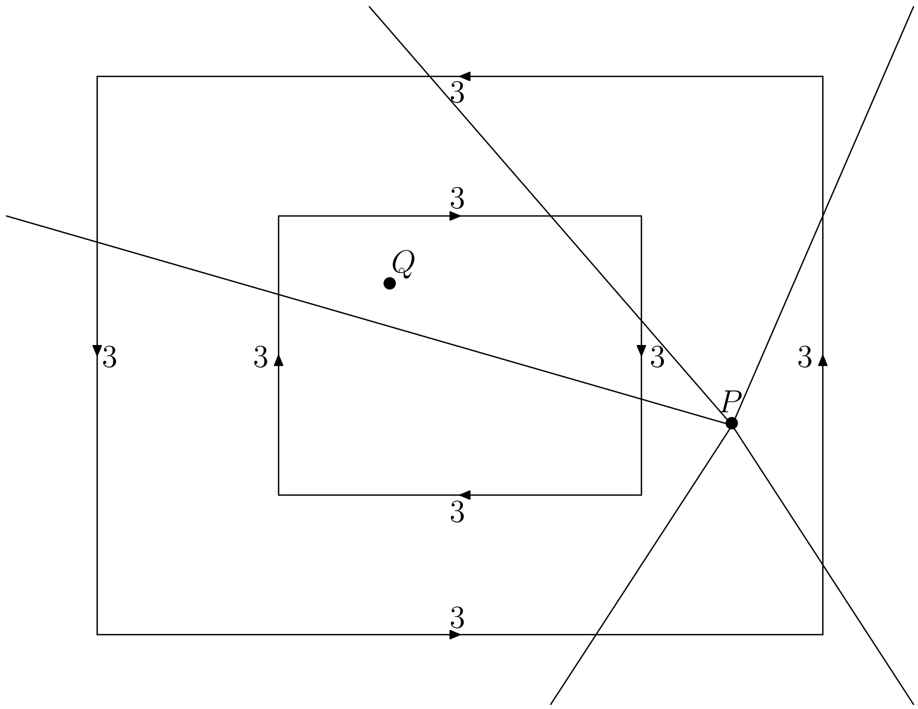

Table 4.3: The rotation number

The tool that we need is the rotation of a 1-cycle around a point not contained in its support. So let ζ^1 be a planar 1-cycle and P be a point outside the support of ζ^1. By

$$\mathrm{rot}_{\zeta^1}(P)$$

denote the flow around P induced by ζ^1. It is the throughput through any infinite ray emanating from P. Note that the throughput does not depend on the ray since any pair of rays creates two regions none of which contains a source nor a sink – recall that ζ^1 is a balanced flow. Compare figure 4.3 Thus, what moves in through one ray has to leave through the other.

Observation 4.1.22. *The following elementary properties of rotation numbers are immediate:*

1. *If P is outside the supports of ζ^1 and ξ^1 then*

$$\mathrm{rot}_{\zeta^1+\xi^1}(P) = \mathrm{rot}_{\zeta^1}(P) + \mathrm{rot}_{\xi^1}(P)\,.$$

2. *The map* rot_{ζ^1} *is locally constant.*

3. *If P is outside the triangle Δ then*

$$\mathrm{rot}_{\partial(\Delta)}(P) = 0.$$

Thus, if the carrier of a 2-chain ζ^2 does not contain P then

$$\mathrm{rot}_{\partial(\zeta^2)}(P) = 0.$$

Corollary 4.1.23. *If* $\mathrm{rot}_{\zeta^1}(P) \neq 0$ *then P is in the carrier of every 2-chain that bounds ζ^1.* **q.e.d.**

The converse also holds so that we have:

Theorem 4.1.24. *Every planar 1-cycle ζ^1 is the boundary of a 2-chain whose carrier is the closure of the set* $\{P \in \mathbb{E}^2 \mid \mathrm{rot}_{\zeta^1}(P) \neq 0\}$.

Corollary 4.1.25. *Let ζ^1 be a 1-cycle in $U \subseteq \mathbb{E}^2$. Then ζ^1 is trivial in* $\mathrm{H}_1(U)$ *if and only if* $\{P \in \mathbb{E}^2 \mid \mathrm{rot}_{\zeta^1}(P) \neq 0\} \subset U$. *In particular, if ζ^1 is a 1-cycle in $U \cap V$ and is trivial in* $\mathrm{H}_1(U)$ *and* $\mathrm{H}_1(V)$ *then it is also trivial in* $\mathrm{H}_1(U \cap V)$. *This is to say that for $U, V \subseteq \mathbb{E}^2$,*

$$\mathrm{H}_1(U \cap V) \to \mathrm{H}_1(U) \oplus \mathrm{H}_1(V)$$

is injective. **q.e.d.**

Proof of (4.1.24). We start with a triangulation of $\mathbb{E}^2$ into equilateral triangles. Consider an edge in the support of ζ^1. It extends to a straight line that cuts some triangles in the triangulation. We subdivide the triangles that are cut to obtain e refined triangulation of $\mathbb{E}^2$ that contains the edge of ζ^1 in its 1-skeleton. We proceed the same way for all edges in ζ^1 – recall that there are only finitely many – and obtain a triangulation $\mathcal{T}$ of $\mathbb{E}^2$ that has the support of ζ^1 in its 1-skeleton. Note that this triangulation has only finitely many triangles in the convex hull of ζ^1.

The function rot_{ζ^1} is defined in the interior of every triangle in our final triangulation. Since it is locally constant and triangles are connected, there is a well defined number $\mathrm{rot}_{\zeta^1}(\Delta)$ for each triangle in $\mathcal{T}$. Since points outside the convex hull of ζ^1 have 0 rotation, these numbers vanish for all but fintely many triangles in the triangulation $\mathcal{T}$.

The proof is completed by the following:

Exercise 4.1.26. The 2-chain

$$\zeta^2 := \sum_{\Delta \in \mathcal{T}} \mathrm{rot}_{\zeta^1}(\Delta)\,\Delta$$

bounds ζ^1. **q.e.d.**

Exercise 4.1.27. Since $\mathbb{E}^2$ is homeomorphic to a punctured sphere $S_2 - \{P\}$, one might expect that the results of this section also apply to 1-cycles and 2-chains in $S_2 - \{P\}$ where a fixed point $P \in S_2$ is removed.

1. Devise a notion of rotation such that the arguments in this section extend to the spherical setting.
2. Show that the results do not hold for the unpunctured sphere. Find a 1-cycle in the sphere that is bounded by two disjoint 2-chains.

4.1.5 Concluding Remarks on Geodesic Homology

Geodesic homology of open sets in $\mathbb{E}^2$ or S_2 is defined in terms of geodesically embedded finite graphs. The reason to require these graphs to be geodesically embedded comes from our need to add them. Since two finite graphs that are just topologically embedded may intersect in infinitely many points, we might be in trouble when adding those. Singular homology is especially designed to handle the technicalities that arise from considering topologically embedded graphs.

Avoiding technical issues comes at a price. We do not know whether our homology groups $\mathrm{H}_1(U)$ are invariants of U considered as a topological space, or whether these groups depend on the particular embedding $U \hookrightarrow \mathbb{E}^2$ or $U \hookrightarrow S_2$. Even more down to earth problems are still open: The stereographic projection induces a map

$$\iota : \mathbb{E}^2 \to S_2$$

which we can use to identify open subsets of $\mathbb{E}^2$ with open subsets of the (punctured) sphere. That gives us two ways of looking at a planar open set U. We can either regard it as an open subset of $\mathbb{E}^2$ and compute its homology this way, or we can identify it with an open subset of S_2 via the stereographic projection and compute its homology groups that way. A priori, there can be a difference since the stereographic projection does not identify planar geodesics with spherical geodesics. In fact, however, there is no real difference:

Exercise 4.1.28. Let $\iota : \mathbb{E}^2 \to S_2$ denote the stereographic projection. Show that for each open subset $U \subseteq \mathbb{E}^2$,

$$\mathrm{H}_1^{\mathbb{E}^2}(U) = \mathrm{H}_1^{S_2}(\iota(U)) .$$

One way of attacking the problem is to introduce a second homology theory. Geodesic homology is based on geodesically embedded graphs. We could, however, allow for more general embeddings as follows: A graph embedded in $\mathbb{E}^2$ or S_2 is weakly circle-like, if all of its edges edges are either geodesic segments or circular arcs. It is clear, that two weakly circle-like graphs intersect in only finitely many points. Thus, we can construct homology groups $\overline{\mathrm{H}}_0(U)$ and $\overline{\mathrm{H}}_1(U)$ for planar and spherical open sets. As the stereographic projection takes planar geodesics/circles to spherical geodesics/circles, it is easy to solve exercise 4.1.28 for these homology theories. That, however leaves the problem to prove

Exercise 4.1.29. Show that for a planar open set $U \subseteq \mathbb{E}^2$ and a spherical open set $V \subseteq S_2$,

$$\begin{aligned} \mathrm{H}_1^{\mathbb{E}^2}(U) &= \overline{\mathrm{H}}_1(U) \\ \mathrm{H}_1^{S_2}(V) &= \overline{\mathrm{H}}_1(V) . \end{aligned}$$

Exercise 4.1.30. Consider a rectangle that is tiled by subrectangles. Each of the subrectangles has is corners colored by red and blue according to one of the following six patterns:

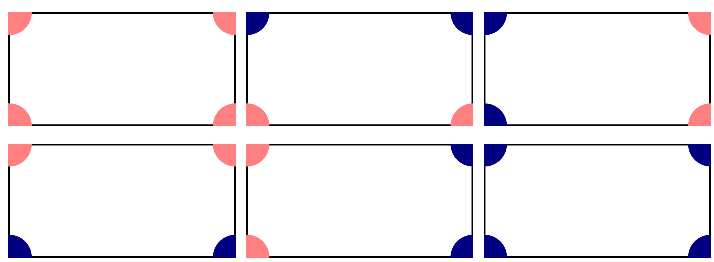

Prove that the induced coloring of the four corners of the big rectangle also fits one of these patterns provided that in the interior the colors around vertices that belong to different rectangles match. Here is an example:

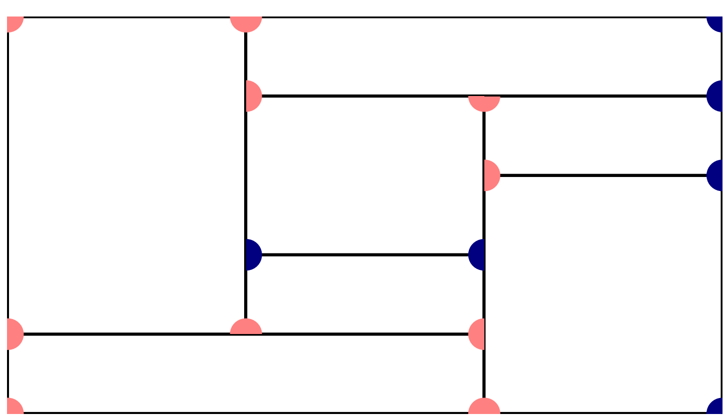

Exercise 4.1.31. This is a similar coloring problem concerning triangles. Assume, you have a triangle which is triangulated, i.e., it is tiled by triangles that intersect

only in vertices and edges such that a vertex of one triangle does not lie in the interior of an edge of another triangle. Assume that the vertices are colored in red, blue, and green such that the following conditions are satisfied:

1. The three vertices of the big triangle are given pairwise different colors.

2. Along each of the three edges of the big triangle, only two colors are used for the vertices in this edge.

Prove that you will always find a tricolored triangle in the subdivision.

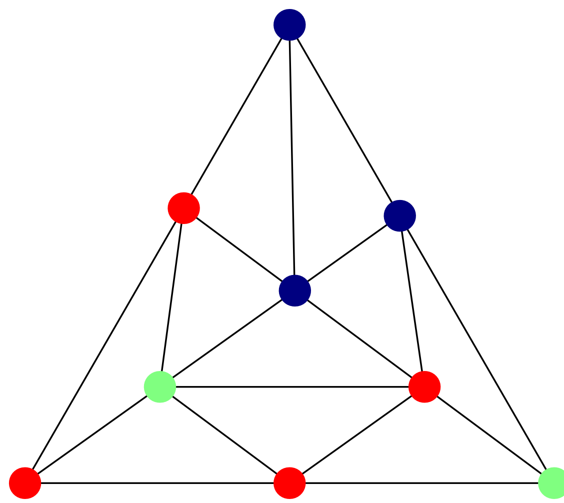

Exercise 4.1.32. Use (4.1.31) to prove Brouwer's fix point theorem for the 2-dimensional disc: Every continous map $f : \mathbb{B}^2 \to \mathbb{B}^2$ has a fixed point.

4.2 Proof of the Curve Theorem

Theorem 4.2.1. *The complement of an embedded broken geodesic arc A in S_2 is acyclic.*

Proof. Let P be a vertex in the interior of the arc. Since the arc is embedded, P splits the arc into two halves $A_{\mathbf{L}}$ and $A_{\mathbf{R}}$ that intersect in $\{P\}$. We put

$$\begin{aligned} X &:= S - \{P\} \\ U &:= S - A_{\mathbf{L}} \\ V &:= S - A_{\mathbf{R}} \\ Y &:= S - A \end{aligned}$$

and it is apparent that $X = U \cup V$ and $Y = U \cap V$. Thus, from the Mayer-Vietoris sequence, we have:

$$0 = \mathrm{H}_1(X) \to \mathrm{H}_0(Y) \to \mathrm{H}_0(U) \oplus \mathrm{H}_0(V)\,.$$

Thus the map

$$\mathrm{H}_0(Y) \to \mathrm{H}_0(U) \oplus \mathrm{H}_0(V)$$

is injective. By the spherical analogue of (4.1.25) the map

$$\mathrm{H}_1(Y) \to \mathrm{H}_1(U) \oplus \mathrm{H}_0(V)$$

is also injective.

Thus equipped, we can argue by contradiction. Suppose ζ was a non-trivial cycle in $\mathrm{H}_i(Y)$ then it would remain non-trivial in at least one of the vector spaces $\mathrm{H}_i(U)$ or $\mathrm{H}_i(V)$. Thus, if we have a witness for the complement of A not being acyclic, the same witness would prove one of the pieces $A_{\mathbf{L}}$ or $A_{\mathbf{R}}$ to have a non-acyclic complement.

Now we finish the proof by induction: The cut vertex P can be chosen so that $A_{\mathbf{L}}$ and $A_{\mathbf{R}}$ both have fewer segments than A. Since we already know that one-segment arcs have acyclic complements, the claim follows. **q.e.d.**

Essentially the same proof yields the stronger statement:

Theorem 4.2.2. *The complement of an embedded topological arc A in S_2 is acyclic.*

Proof. The proof above applies almost literally, the only difficulty is the induction part at the end: we might keep subdividing our arc, getting smaller and smaller pieces, but the process never stops. The point is that, morally, the process converges to a point which gives the start of the induction as a single point has an acyclic complement.

To make this precise let us assume that the cycle ζ is a witness for $S_2 - A$ not being acyclic. We argued above that, given any subdivision of A into two pieces $A_{\mathbf{L}}$ and $A_{\mathbf{R}}$ that intersect in a single point, the witness ζ will either prove $S_2 - A_{\mathbf{L}}$ or $S_2 - A_{\mathbf{R}}$ to be non-acyclic. We choose the evil part an keep subdividing. This way, we obtain a sequence of subarcs $A = A_0 \supset A_1 \supset A_2 \supset \cdots$ that converges to a single point $P = \bigcap A_i$.

We know, that the complement of P is acyclic. So ζ is a boundary in $S_2 - \{P\}$. The chain that allows us to write ζ as a boundary, however, has a compact carrier that avoids P and thus an open neighborhood of P. This open neighborhood already contains almost all A_i . Thus, ζ is a boundary in $S_2 - A_i$ for sufficiently large i; but that contradicts the way we constructed the sequence A_i keeping ζ a witness for all these arcs to be evil. **q.e.d.**

Corollary 4.2.3 (Jordan's Curve Theorem). *The complement of a embedded topological loop γ in S_2 has two components.*

Proof. We write the loop as the union $\gamma = A_{\mathbf{L}} \cup A_{\mathbf{R}}$ of two arcs that intersect in the two point set $\{P, Q\}$. Put

$$\begin{aligned} X &:= S - \{P, Q\} \\ U &:= S - A_{\mathbf{L}} \\ V &:= S - A_{\mathbf{R}} \\ Y &:= S - \gamma \end{aligned}$$

and write down the Mayer-Vietoris sequence:

$$\mathrm{H}_1(U) \oplus \mathrm{H}_1(V) \to \mathrm{H}_1(X) \to \tilde{\mathrm{H}}_0(Y) \to \tilde{\mathrm{H}}_0(U) \oplus \tilde{\mathrm{H}}_0(Y)\,.$$

We know:

$$\begin{aligned} \mathrm{H}_1(U) \oplus \mathrm{H}_1(V) &= 0 \\ \mathrm{H}_1(X) &= \mathbb{R} \\ \mathrm{H}_0(U) \oplus \tilde{\mathrm{H}}_0(Y) &= 0. \end{aligned}$$

Thus, $\mathrm{H}_0(Y) = \mathbb{R}$ which implies that the loop complement Y has precisely two components. **q.e.d.**

Exercise 4.2.4. Let C be a compact set in the plane $\mathbb{E}^2$. Suppose the complement has precisely two components and each point in C is *arcwise accessible* from both components. (A point in C is arcwise accessible from a given complementary component if it is the endpoint of an arc that is, away from this endpoint, completely contained in a complementary component.) Show that any two point subset of C disconnects C.

4.3 Schönflies' Theorem

4.3.1 The Schönflies Theorem for Polygons

In this section, we follow [Mois77].

Definition 4.3.1. Let $\mathcal{T}$ be a triangulation of a planar compact region whose boundary curve is a closed polygon P. A triangle $\Delta \in \mathcal{T}$ is called free if $\partial(\Delta) \cap P$ consists of one or two edges of Δ.

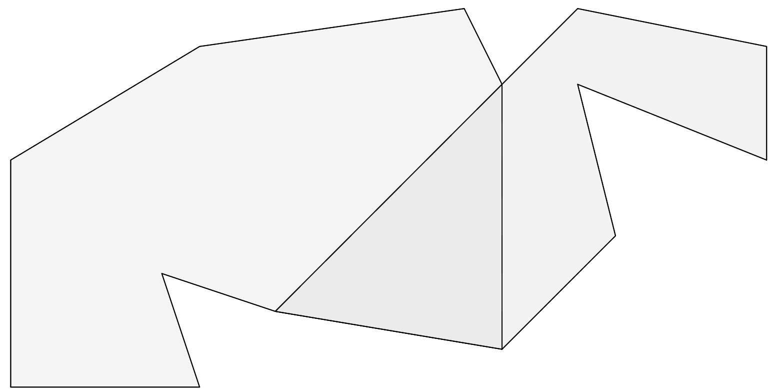

Table 4.4: Splitting a polygon

Lemma 4.3.2. *Let $\mathcal{T}$ be a triangulation of a planar compact region whose boundary curve is a closed polygon P. Suppose that $\mathcal{T}$ consists of two or more triangles. Then at least two of them are free.*

Proof. We proceed by induction. The case $|\mathcal{T}| = 2$ is clear: Both triangles are free.

Since P has more than three edges, there are at least two triangles in $\mathcal{T}$ that have an edge in P. If both of them are free, we are done. Otherwise, there is a triangle Δ that has an edge in P but is not free. In this case, the picture 4.4 captures the essence of the situation: The triangulation can be split into two parts both of which have fewer triangles and contain the triangle Δ as a free triangle. By induction, both parts contain free triangles that will be free in $\mathcal{T}$ too. **q.e.d.**

Exercise 4.3.3. Let P be a simple closed polygon, which cuts the plane into two regions. The closure of the bounded region can be triangulated and has P as its boundary.

Theorem 4.3.4. *Let P be a simple closed polygon, which cuts the plane into two regions. There is a homeomorphism of the plane that takes the closure of the bounded region B to a triangle.*

Proof. First, we find a triangulation of the plane that has the polygon P in its 1-skeleton. This is possible since P has only finitely many edges and the lines containing these segments chop up the plane into finitely many convex regions which in turn can be subdivided into triangles. Obviously, we can arrange things so that the interior region cut out by P covers only finitely many triangles. Thus, let $\mathcal{T}$ be a finite triangulation of this interior region.

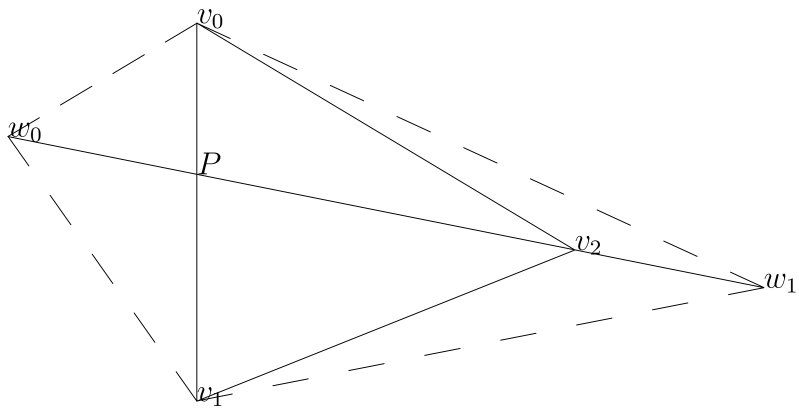

Table 4.5: Neighborhood of a free triangle

If $|\mathcal{T}| = 1$, we are done. Otherwise, we will find a homeomorphism of the plane that removes a free triangle from $\mathcal{T}$. Then we find the next homeomorphism, and we continue in this fashion until we transformed P into a triangle.

It remains to present the homeomorphism that gets rid of a free triangle $\Delta \in \mathcal{T}$. Free triangles come in two flavors:

Δ has two edges in P: Let v_0, v_1, and v_2 be the vertices of Δ. Suppose that the edge v_0v_1 is inner. Then you can find two points w_0 and w_1 collinear with v_2 such that:

1. The point w_0 is inner.
2. The line segments w_0v_2 and v_0v_1 intersect. Let P be the point of intersection.
3. The point w_1 is out.

Thus, we have the picture shown in figure 4.5 By affine extension, we find a homeomorphism that leaves the outer quadrilateral of this figure fixed pointwise and takes v_2 to P. We extend this to the plane. This homeomorphism visibly kills the triangle.

Δ has one edge in P: In this case, we assume that v_2 is the inner vertex of Δ. Here everything works as above, except that w_0 is out and w_1 is inner. The final homeomorphism also works the other way: it takes P to v_2. **q.e.d.**

Remark 4.3.5. From the proof, it is clear that we can construct the homeomorphism so that it is the identity outside a previously chosen small open neighborhood of B.

Corollary 4.3.6. *Every simple closed polygon in the plane is the boundary of an embedded disc.* **q.e.d.**

4.3.2 The General Case

Theorem 4.3.7 (Schönflies). *Let γ be a simple closed curve in the plane. Then the closure of the bounded complementary region is a disc.*

Lemma and Definition 4.3.8. *Let γ be a simple closed curve in the plane with interior (bounded complementary region) I, and let A be an open arc in γ. Then there is a line segment vw with $v \in A$ and $vw - \{v\} \subset I$.*

In this case, we say that v is linearly accessible from I and vw is a driveway for v.

Proof. Let v_0 be any point in A and let w be a point in I that close to v_0 such that the closed disc centered at v_0 through w does not contain any points of γ outside A. Draw the line segment v_0w and let v be the last point on that segment that belongs to γ. Observe that $v \in A$. **q.e.d.**

Lemma 4.3.9. *Let γ be a simple closed curve in the plane with interior region I and let A be a closed arc in γ with linearly accessible endpoints $v_{\mathbf{L}}$ and $v_{\mathbf{R}}$. Let $v_{\mathbf{L}}w_{\mathbf{L}}$ and $v_{\mathbf{R}}w_{\mathbf{R}}$ be driveways for $v_{\mathbf{L}}$ and $v_{\mathbf{R}}$. For any positive number ε, there is a broken line joining the segment $v_{\mathbf{L}}w_{\mathbf{L}}$ to $v_{\mathbf{R}}w_{\mathbf{R}}$ that is contained in I, stays within ε-distance to A, and intersects each of the two driveways in precisely one point.*

Proof. We can always find a broken line β connecting the two driveways inside I satisfying all conditions but ε-closeness. Let $w'_{\mathbf{L}}$ and $w'_{\mathbf{R}}$ be the points where β meets the driveways. Now, if β is not ε-close to A, then we can shrink ε and assume that β actually misses the ε-neighborhood of A altogether and moreover lies in its unbounded complementary region.

We draw a square grid whose square are so small that none of them contains two points of distance ε. Moreover, we do not want the grid to have a vertex on one of the driveways. Let Q be the union of those squares that meet A. Note that Q is connected and ε-close to A. Thus β lies in the unbounded complementary region of Q. See figure 4.6.

The part of the boundary of Q that borders the unbounded complementary region is a simple closed polygon because Q is *gallery connected*: any two squares in Q can be joined by a chain of squares in which neighboring squares share an edge.

Exercise 4.3.10. Justify this way of reasoning.

Let P be this simple closed polygon bordering the unbounded complementary region. We consider the simple closed curve

$$\gamma' := v_{\mathbf{L}} w'_{\mathbf{L}} \cup \beta \cup v_{\mathbf{R}} w'_{\mathbf{R}} \cup A$$

and observe that its interior I intersects P in a finite number of broken lines. We will show that one of them connects the driveways.

So suppose none of the parts of $\overline{I} \cap P$ connects the two driveways. Consider the polygon P. Its interior contains A whereas β lies in the exterior. Thus P separates A and β. Thus $\overline{I} \cap P$ separates A and β inside $\overline{I}$. Hence any broken line π from β to A will intersect at least one of the pieces in $\overline{I} \cap P$. So pick an π with the minimum number of intersection points. We claim that this particular π does not intersect $\overline{I} \cap P$ at all, which is a contradiction.

To see that, we consider the first point on π that belongs to P. The particular arc in $\overline{I} \cap P$ does not connect the two driveways. Thus both endpoints of this arc are on the same driveway. It follows that there is another intersection of this arc with π. We now modify π as indicated in figure 4.7 to get rid of two intersections. This proves the claim. **q.e.d.**

Exercise 4.3.11. Let $f : S_1 \to \mathbb{E}^2$ be continuous and injective. Show that f is a homeomorphism from S_1 onto its image.

Proof of Schönflies' Theorem.

Claim A. There is a sequence $\mathcal{P}_0, \mathcal{P}_1, \mathcal{P}_2, \ldots$ such that:

- Each $\mathcal{P}_i$ is a finite decomposition of γ into arcs intersecting in their endpoints.
- These endpoints are all linearly accessible.
- Each arc has diameter $\leq \frac{1}{i}$.
- $\mathcal{P}_{i+1}$ is a refinement of $\mathcal{P}_i$, i.e., The arcs in $\mathcal{P}_i$ are unions of consecutive arcs in $\mathcal{P}_{i+1}$. For technical reasons, we require that each arc in $\mathcal{P}_i$ really is subdivided in the next step. Thus, the number of arcs in the decompositions at least doubles in each step.

Proof. By exercise 4.3.11, we can pretend that γ is a circle in the topology induced by the metric inherited from $\mathbb{E}^2$. Since each open arc contains linearly accessible points, these points are dense in γ. Now it is easy to find the first partition and then refine it successively. □

Claim B. The is a sequence $\mathcal{D}_0, \mathcal{D}_1, \mathcal{D}_2, \ldots$ such that

- Each $\mathcal{D}_i$ is a collection of driveways, containing precisely one driveway leading to each of the partition points in $\mathcal{P}_i$.
- No two driveways in $\mathcal{D}_i$ intersect.
- $\mathcal{D}_i \subset \mathcal{D}_{i+1}$.

PROOF. Since $\mathcal{P}_1$ is finite, we can construct $\mathcal{D}_1$. Since the $\mathcal{P}_i$ stay finite, we can continue. □

Now, we use Lemma 4.3.9 to construct a polygonal disc C_0 and a sequence of annuli $R_1, R_2, \ldots$ where R_i is decomposed into polygonal discs $C_i^{(j)}$. We proceed as follows:

1. For every arc $A \in \mathcal{P}_0$, we create broken line β_A connecting the two driveways that stays within distance 1 of A. With only little extra effort, we arrange that two broken lines of our collection do not intersect away from the driveways and that two adjacent broken lines hit their intermediate driveway in the same point. Thus the broken lines form a simple closed polygon. By the polygonal Schönflies Theorem, its interior region is a disc C_0.

2. Similarly, for every $A' \in \mathcal{P}_1$, we create a broken line $\beta_{A'}$ that connects the driveways for A' and stays within distance $\frac{1}{2}$ of A'. Again, we avoid that these broken lines have unwanted intersections among each other or with the driveways. We also avoid intersections with C_0 which we can easily achieve by creating these lines even closer to their arcs. These broken lines form a new simple closed polygon. The disc cut out by this polygon contains the middle disc C_0. Thus we have constructed an annulus R_1. The driveways for $\mathcal{P}_0$ chop this annulus into discs $C_1^{(j)}$. Here again, we use the Schönflies Theorem for polygonal curves to see that the regions $C_1^{(j)}$ are, in fact, discs.

3. Now we turn to $\mathcal{P}_2$. Here we will have the broken lines within distance $\frac{1}{3}$ of their corresponding arcs. We continue the process for all partitions $\mathcal{P}_i$ using a maximum distance of $\frac{1}{i}$. Thus, we create annuli R_i chopped up into discs $C_i^{(j)}$ by the driveways in $\mathcal{D}_{i-1}$. Observe that the disc $C_i^{(j)}$ is contained in an $\frac{1}{i}$-neighborhood of its arc, which in turn has diameter $\leq \frac{1}{i}$.

We copy this picture into the unit disc where we subdivide each arc of one generation evenly within the next generation. Thus we have a middle disc $\tilde{C}_0$ surrounded by annuli $\tilde{R}_1, \tilde{R}_2, \ldots$ chopped into discs $\tilde{C}_i^{(j)}$. See figure 4.8

We can easily define a homeomorphism on the 1-skeleton of these pictures that matches corresponding arcs. This homeomorphism extends to the 2-skeleta since any homeomorphism between the boundaries of two discs extends to their interiors. Thus, we have a homeomorphism

$$\tilde{C}_0 \cup \bigcup_i^j \tilde{C}_i^{(j)} \to C_0 \cup \bigcup_i^j C_i^{(j)}.$$

It remains to see that this homeomorphism extend continuously to the boundary. This, however, follows from the fact that an "ith-generation" disc is small and close to the boundary: The discs $\tilde{C}_i^{(j)}$ are small since their number per annulus at least doubles each step. The disc $C_i^{(j)}$ in the image is small by construction: It is within $\frac{1}{i}$-distance from its corresponding arc. This arc, in turn has diameter $\leq \frac{1}{i}$. **q.e.d.**

4.4 Consequences

Corollary 4.4.1. *Let γ be a simple closed curve in S_2. Then γ cuts S_2 into two discs intersecting in γ.*

Corollary 4.4.2. *Let γ be a simple closed curve in S_2. Then any homeomorphism $\zeta : \gamma \to S_1 \subset S_2$ extends to a homeomorphism $S_2 \to S_2$.*

Proof. This is clear since any homeomorphism of the boundaries of two discs extends to the interiors. The result follows from applying this to the two discs in $S_2 - \gamma$. **q.e.d.**

Corollary 4.4.3 (Schönflies Theorem, second form). *Let γ be a simple closed curve in $\mathbb{E}^2$. Then every homeomorphism $\gamma \to S_1 \subset \mathbb{E}^2$ extends to a homeomorphism $\mathbb{E}^2 \to \mathbb{E}^2$.*

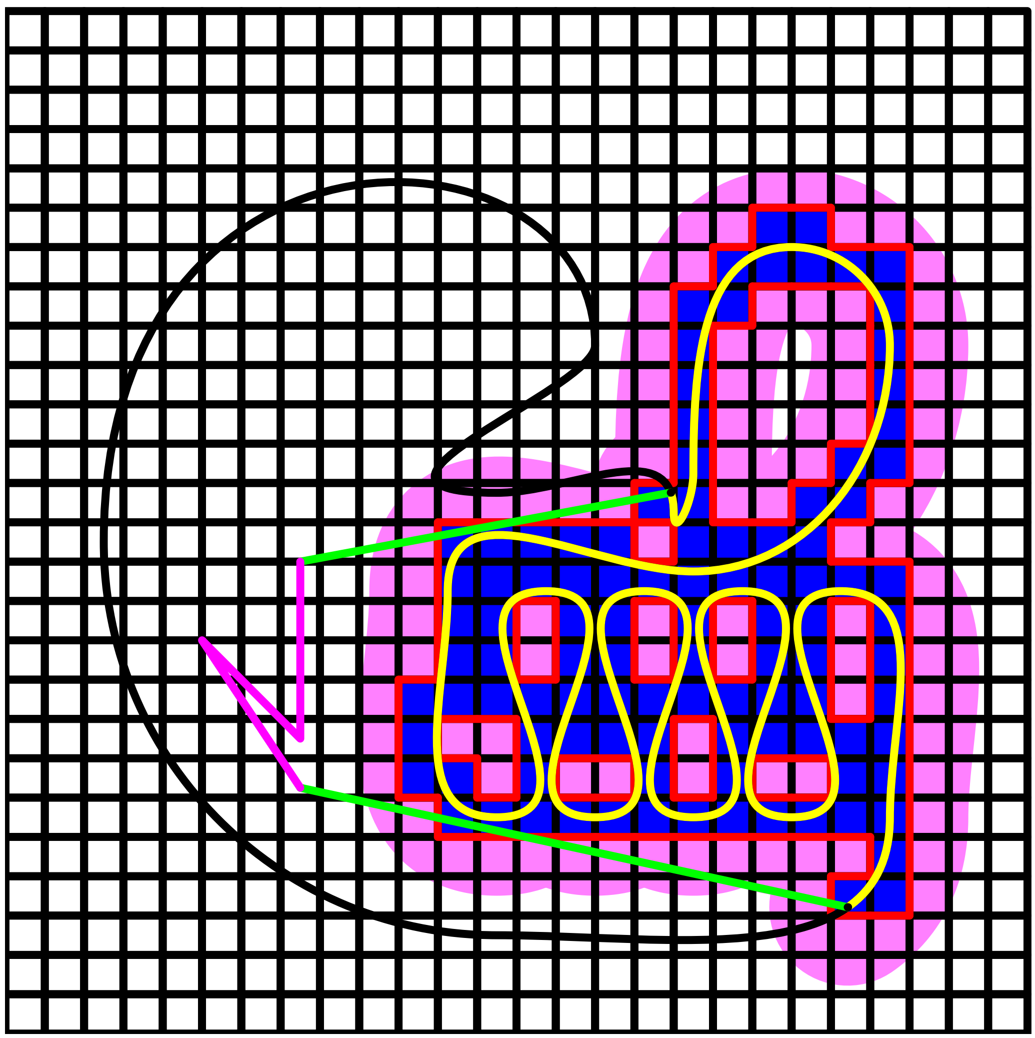

Table 4.6: The outer red boundary is P

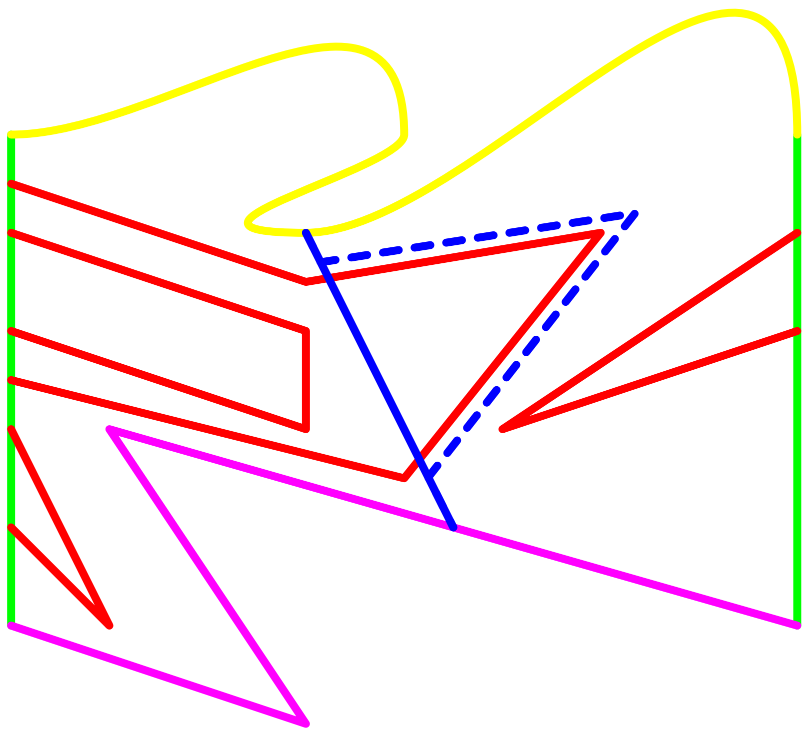

Table 4.7: Eliminating two intersections

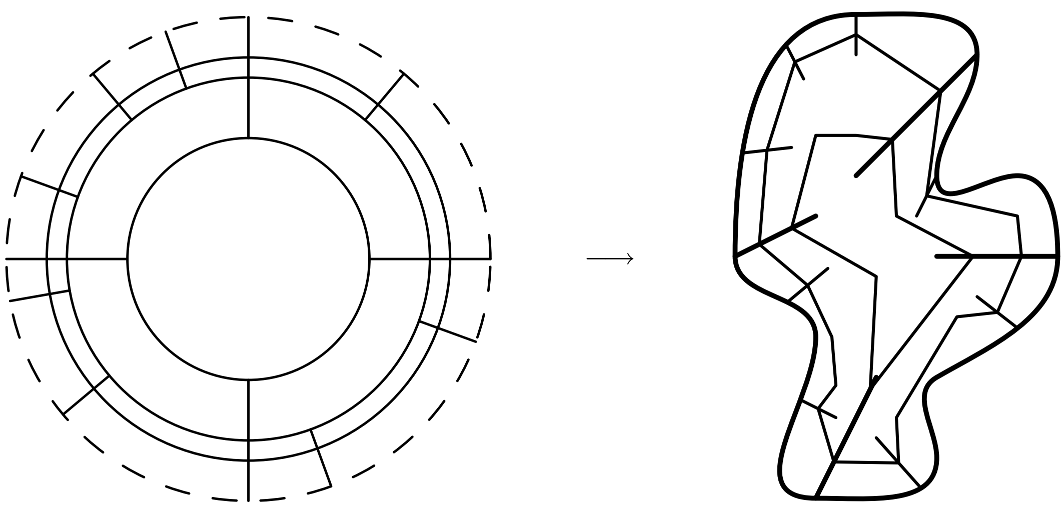

Table 4.8: Construction of $\tilde{C}_i^{(j)}$ and $C_i^{(j)}$

Chapter 5

Classification of Closed Surfaces

5.1 Manifolds

Definition 5.1.1. A (topological) m-manifold is a second countable Hausdorff space M wherein each point has a neighborhood that is homeomorphic to an open subset of $\mathbb{R}^m$.

A chart is a pair $(U, \varphi : U \to \tilde{U})$ where U and $\tilde{U}$ are connected open subset of M and $\mathbb{R}^m$ respectively and where φ is a homeomorphism. A collection of charts whose domains cover M is called an atlas.

Let $(U_0, \varphi_0 : U_0 \to \mathbb{R}^m)$ and $(U_1, \varphi_1 : U_0 \to \mathbb{R}^m)$ be two charts. Put $V := U_0 \cap U_1$. Then the map

$$\begin{array}{rcl} \xi : \varphi_0(V) & \to & \varphi_1(V) \\ x & \mapsto & \varphi_1\big(\varphi_0^{-1}(x)\big) \end{array}$$

is a homeomorphism called change of coordinates.

Example 5.1.2. Euclidean spaces of finite dimension are manifolds. So are spheres and real projective spaces. Direct products of manifolds are manifolds. In particular, the torus is a manifold.

5.2 Euler Characteristic

Definition 5.2.1. An abstract simplicial complex is a set $\mathcal{V}$ of vertices together with a collection $\mathcal{S}$ of non-empty finite subsets called simplices containing all singleton subsets of $\mathcal{V}$ and satisfying the condition that any non-empty subset of a simplex is also a simplex. A simplicial map between abstract simplicial complexes is a map

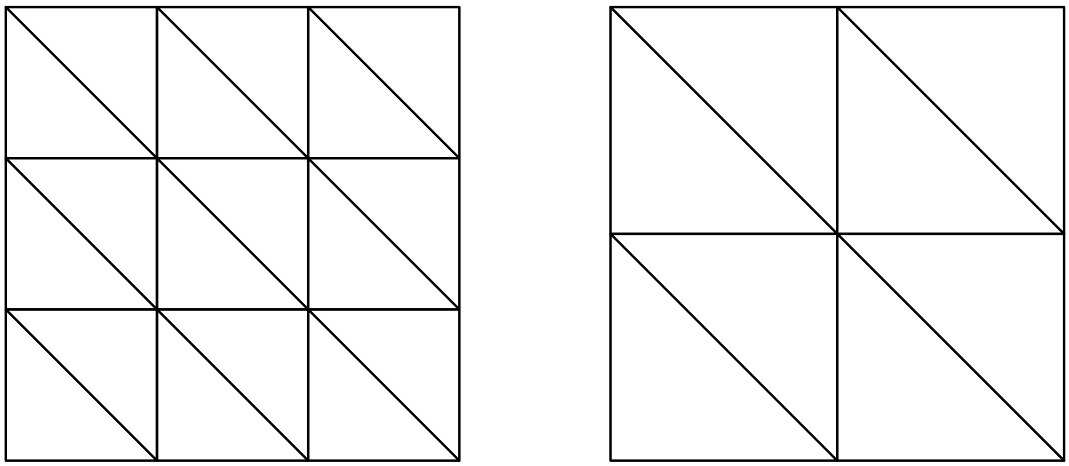

Table 5.1: Left: a triangulation of the torus; Right: a non-triangulation of the torus.

between their vertex sets such that the image of any simplex is a simplex in the target complex.

The realization of a simplicial complex $K = (\mathcal{V}, \mathcal{S})$ is defined as

$$|K| := \bigcup_{\sigma \in \mathcal{S}} |\sigma|$$

where $|\sigma|$ is the convex hull of the vertices in σ in $\mathbb{R}^{\mathcal{V}}$. Since $|K|$ is defined as a union, we will endow it with the weak topology. If you do not know what that is, never mind: we will consider finite simplicial complexes only, and for these, the weak topology coincided with the subspace topology inherited from $\mathbb{R}^{\mathcal{V}}$. Note that every simplicial map induces a continuous map between realizations.

A simplicial complex is the realization of an abstract simplicial complex sometimes with and sometimes without added information on what the vertices and simplices are.

Definition 5.2.2. Let X be a topological space. A triangulation of X is a simplicial complex that is homeomorphic to X. Sometimes, we will call the homeomorphism the triangulation.

Example 5.2.3. Figure (5.1) shows a triangulation and a non-triangulation of the torus.

Exercise 5.2.4. Find a triangulation of the torus that uses as few triangles as possible.

Definition 5.2.5. The closed star of a simplex τ in a simplicial complex K is the subcomplex of all simplices containing τ. The link of τ is the boundary of the star. Equivalently, it is the subcomplex of all those simplices σ such that $\sigma \cap \tau = \emptyset$ although $\sigma \cup \tau$ is a simplex.

There is an obvious way for a simplicial complex to be a manifold: All simplex links are spheres of the appropriate dimension. Those complexes are called combinatorial manifolds. In dimension 2, this is the only possibility.

Theorem 5.2.6. *Let K be a triangulated 2-manifold. Then K is a combinatorial 2-manifold, i.e., the link of each vertex is a subdivided circle.*

Proof. Every point in K has a neighborhood homeomorphic to an open disc. Hence, there are no isolated vertices. Moreover, every edge borders at least one triangle: Otherwise an interior point of that edge would separate every sufficiently small neighborhood, which is impossible in a 2-manifold as it does not happen in the plane. Similarly, the fact that no semicircle can separate the plane implies that every edge is, in fact, contained in at least two triangles.

Now we show that each edge is contained in at most two triangles. So suppose the edge e was in the intersection of at least three triangles. Then a point in the interior of e has a circle around it passing through two of these triangles. But that circle does not separate since you can bypass it along the third triangle. This contradicts the Jordan curve theorem which should hold near every point.

It follows that the link of every vertex is a disjoint union of circles. Since no point in a manifold can separate its neighborhoods, the link consists of one circle only. **q.e.d.**

Remark 5.2.7. In higher dimension, all sorts of bad things happen. There are manifolds that do not admit a combinatorial triangulation although they have a triangulation. In particular, that some links in a simplicial complex are not spheres does not imply that the complex is not a manifold. It is really hard to think how a vertex with a non-sphere link can have a neighborhood that is an open disc.

Definition 5.2.8. Let K be an abstract simplicial complex. A subdivision of K is an abstract simplicial complex L such that

1. The vertices of L are points in $|K|$.

2. Every simplex of L is contained in the realization of a simplex of K.

3. The induced linear map $|L| \to |K|$ is a homeomorphism.

Two simplicial complexes are called combinatorially equivalent if they have isomorphic subdivisions.

Example 5.2.9. Any two subdivisions of a given simplicial complex K are combinatorially equivalent. In fact, they have a common refinement.

Proof. Since all simplices of the subdivisions are contained in simplices of K and look straight therein, their intersections, if non-empty, are convex. From here, a common finer subdivision is easily found. **q.e.d.**

Fact 5.2.10. *Every compact* 2*-manifold* Σ *has a triangulation and any two triangulations of* Σ *are combinatorially equivalent.*

Remark 5.2.11. In higher dimensions, it is not true that all manifolds have triangulations, and there are manifolds that admit combinatorially inequivalent triangulations. Even when we restrict ourselves to combinatorial triangulations to begin with, there are inequivalent ones. !!! give a reference !!!

Corollary 5.2.12. *Any invariant of* 2*-manifolds defined in terms of combinatorial equivalence classes of triangulations is, in fact, a topological invariant of the manifold.*

Example 5.2.13. The Euler characteristic of a simplicial complex is the alternating sum of the numbers of simplices in different dimensions, i.e.,

$$\chi(K) := \sum_{m \geq 0} (-1)^m \left|\{\sigma \in K \mid \dim(\sigma) = m\}\right|.$$

If L is a subdivision of K, then $\chi(L) = \chi(K)$. Hence, the Euler characteristic of a surface is a topological invariant.

Proof. In dimension 2 the "deletion proof" works: Inside the triangles, delete edges one by one decreasing the number of regions and edges by one. If there is only one region left, delete interior vertices along with their edges (push in free faces!). Finally, delete vertices in the 1-skeleton. See figure (5.2)

Warning: This proof does not work in higher dimensions. Removing the 1-dimensional material is possible only because we can find terminal vertices. In dimension 3, we would be left with the task of removing 2-complexes. However, we might run into something like Bing's house (see figure 5.3) where we do not find any "free faces" to push in. **q.e.d.**

Exercise 5.2.14. Give a proof for the invariance of the Euler-characteristic with respect to subdivisions that works in all dimensions.

Remark 5.2.15. Since the Euler characteristic can also be computed from the rank of singular homology groups, it turns out, that the Euler characteristic is a topological invariant for all triangulable spaces, i.e., any two triangulations of the same space have the same Euler characteristic, even if they are not combinatorially equivalent.

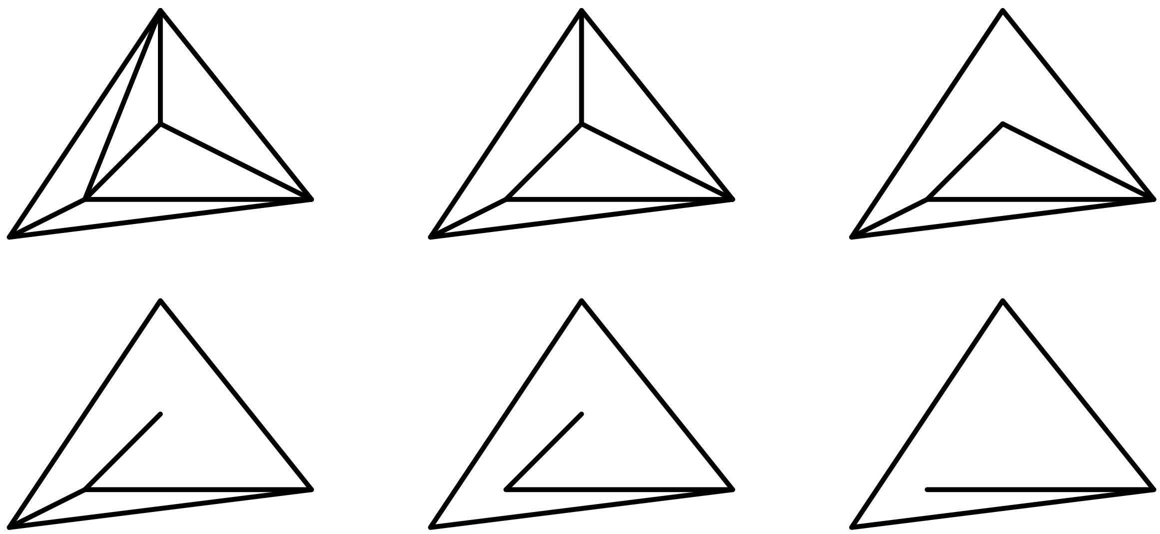

Table 5.2: The deletion proof

5.3 Triangulability of Surfaces

Theorem 5.3.1 (Rado 1925). *Every 2-manifold Σ has a triangulation.*

We present Rado's original proof since it is efficient and not technical.

Proof of theorem 5.3.1. We need a little bit of terminology. Call an embedded closed disc J in Σ a notionJordan domain if it is contained in a chart. Since the topology of Σ has a countable basis, Σ is covered by countably many charts. Each chart, in turn, allows for a countable set of Jordan domains whose interiors cover the chart. Thus, we find a sequence $J_1, J_2, \ldots$ of Jordan whose interiors cover Σ.

Claim A. There is a sequence $J_1^*, J_2^*, \ldots$ of Jordan domains whose interiors cover Σ such that $\partial(J_i^*) \cap \partial\big(J_j^*\big)$ is finite for $i \neq j$.

Proof. Put $J_1^* := J_1$. Now suppose that $J_1^*, J_2^*, \ldots, J_r^*$ have already been constructed such that:

1. The regions J_i^* "thicken" the original domains J_i, i.e., we have $J_i \subseteq J_i^*$ for $i = 1, 2, \ldots, r$.
2. The set of intersections

$$M_r := \bigcup_{i<j\leq r} \partial(J_i^*) \cap \partial\big(J_j^*\big)$$

is finite.

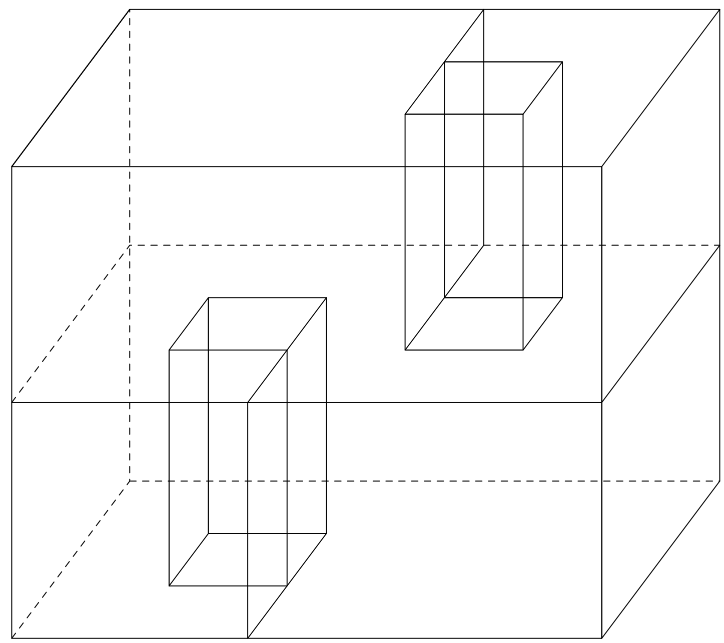

Table 5.3: Bing's house

Our task will be to find the next term J_{r+1}^* such that the above two conditions are preserved.

Some more local definitions will ease the argument. We call the points in M_r *crossings*, and a path in Σ is *admissible* if it intersects the set

$$B := \bigcup_{i \leq r} \partial(J_i^*)$$

in only finitely many points. Given an open set $U \subseteq \Sigma$, call two points U-*equivalent* if they can be connected by an admissible path in U. Note that U-equivalence is an equivalence relation.

Let P be a point outside M_r. Any neighborhood U of P contains a subneighborhood V such that any two point in V are U-equivalent. Indeed, if P does not lie in B, we can choose V so that it is connected and does not intersect B. In this case, any two points in V are even V-equivalent. If $P \in B - M_r$ then we chose V to intersect only one of the boundary curves $\partial(J_i^*)$. The Schönflies Theorem implies that we can pretend this curve is the unit circle in the plain. In this picture, however, the claim is obvious.

It follows that for any open set U that does not contain any crossings, the U-equivalence classes are open. Thus, they coincide with the components of U.

In particular, if U is connected, any two points in U can be connected by an admissible path.

Now, we can proceed to construct J^*_{r+1}. Consider a chart that contains J_{r+1}. By the Schönflies Theorem we can assume that J_{r+1} is represented as the closed unit disc in this chart. Let R be an open annulus around J_{r+1} that does not contain any crossings. Thus, R intersects B in a union of disjoint arcs. The set $R - B$ is non-empty and open. Thus there are two points $P, Q \in R$ such that:

1. The two radii from the center of J^*_{r+1} through P and Q do not intersect and hence separate the annulus into two topological rectangles R_0 and R_1.
2. Both points have open neighborhoods in $R - B$. Thus, we can find points P_0, P_1 close to P and points Q_0, Q_1 close to Q such that the segments PP_i and QQ_i are contained in the rectangle R_i and do not intersect B.

Since the rectangles are connected and do not contain any crossings, we know that P_i and Q_i can be connected inside R_i by a path p_i that intersects B only finitely many times.

Now we concatenate: the path

$$(Q_1 Q Q_0) \rightarrowtail p_0 \rightarrowtail (P_0 P P_1) \rightarrowtail p_1$$

is a closed loop inside R surrounding J_{r+1}. Deleting pieces if necessary to avoid self-intersection, we find a simple closed curve inside R that intersects B only finitely many times and whose interior contains J_{r+1}. This is our choice for J^*_{r+1}. □

Now, we define a sequence

$$P_1^{(1)}, P_2^{(1)}, \ldots, P_2^{(s_2)}, P_3^{(1)}, \ldots, P_3^{(s_3)}, \ldots$$

of closed discs that cover Σ such that the following hold:

1. For any r, we have $\bigcup_{i \leq r} J_i^* = \bigcup_{i \leq r} \bigcup_{j=1}^{s_i} P_i^{(j)}$.
2. The interiors of the $P_i^{(j)}$ are pairwise disjoint.
3. Each point is contained in only finitely many $P_i^{(j)}$. Note that by compactness of closed discs, this condition implies that any of these discs meets only finitely many other discs of the sequence.

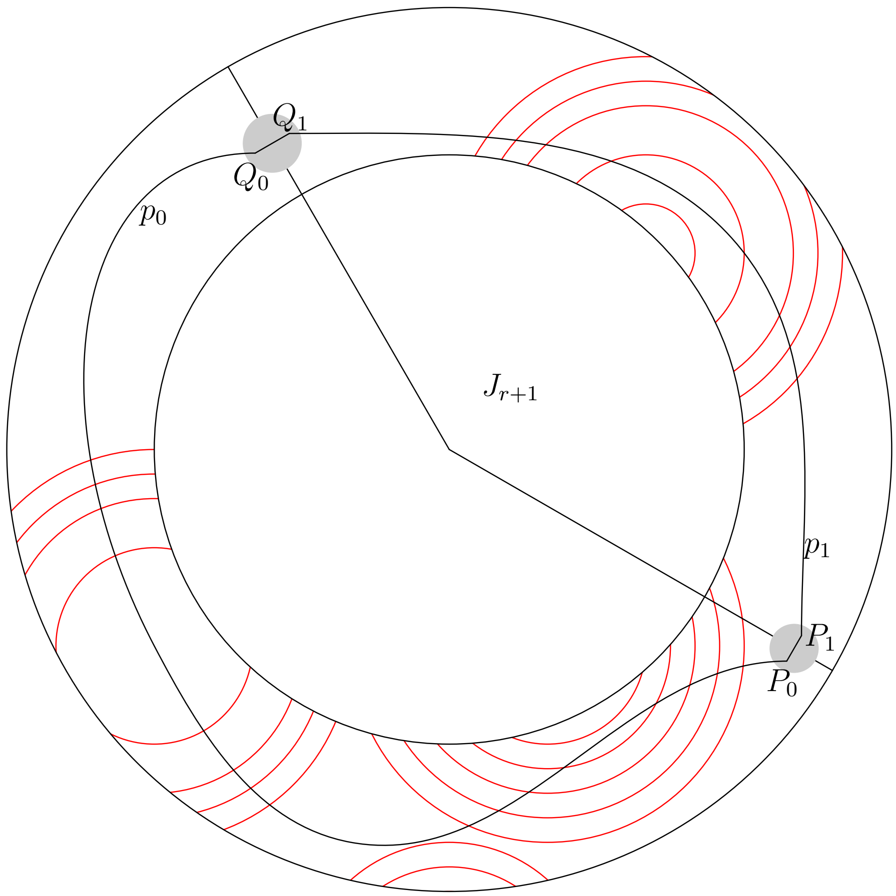

Table 5.4: The construction of J^*_{r+1}

Thus, we can think of these $P_i^{(j)}$ as a polygonal decomposition of Σ which is easily turned into an honest triangulation.

Hence, we are reduced to proving the existence of the sequence $P_1^{(1)}, P_2^{(1)}, \dots, P_2^{(s_2)}, P_3^{(1)}, \dots, P_3^{(s_3)}, \dots$. Put $P_1^{(1)} := J_1^*$. Suppose already have constructed

$$P_1^{(1)}, P_2^{(1)}, \dots, P_2^{(s_2)}, P_3^{(1)}, \dots, P_3^{(s_3)}, \dots P_r^{(1)}, \dots, P_r^{(s_r)}.$$

The Jordan domain J^*_{r+1} is chopped up into regions by the boundary curves $\partial(J_i^*)$ for $i \leq r$. Some of these regions might not be discs but contain finitely many holes. We further subdivide and arrive at a decomposition of J^*_{r+1} into finitely many discs. Among these we chose as $P_{r+1}^{(1)}, \dots, P_{r+1}^{(s_{r+1})}$ precisely those that do not contain any interior point of $\bigcup_{i \leq r} J_i^*$.

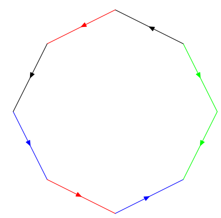

Table 5.5: A polygon diagram

Of the three requirements our sequence is supposed to meet, only (3) requires proof. So let P be a point in Σ. There is a Jordan domain J_k^* containing P as an interior point. Let U be a neighborhood of P in J_k^*. A disc P_i^j can intersect U only if $i \leq k$. This establishes (3) and completes the proof. **q.e.d.**

5.4 The Classification Theorem

We saw that each surface has a triangulation. Compact surfaces have finite triangulations. In this section, we shall see that one can put these combinatorial data into a standard form.

The torus is obtained from the square by identifying opposite edges. In general, a polygon diagram is a polygon whose edges are marked with orientation arrows and colors such that each color occurs exactly twice, see figure (5.5). From a polygon diagram, we obtain a topological space by gluing edges of the same color together so that their arrows match up.

Exercise 5.4.1. Show that the space defined by a polygon diagram is a closed surface.

Proposition 5.4.2. *Every closed surface Σ can be described by a polygon diagram.*

The proof is an interpolation between two-dimensional simplicial complexes and polygon diagrams. Thus, we need a notion that generalizes both.

Definition 5.4.3. A polygon complex is a collection of polygons whose edges are colored and marked with orientation arrows.

Observation 5.4.4. *The following are obvious:*

1. *A two-dimensional simplicial complex is a polygon complex if and only if every vertex and every edge are contained in a two-simplex. In general, simplicial complexes whose maximal simplices all have the same dimension are called chamber complexes.*

2. *The polygon diagrams are precisely those polygonal complexes that consist of just one polygon.*

3. *Any polygon complex gives rise to a topological space by identifying edges of the same color respecting the orientation of the edges.*

4. *Any polygonal complex can be subdivided to yield a two-dimensional chamber complex.*

Proof of (5.4.2). Since Σ can be triangulated, there is a polygon complex realizing Σ. Now suppose, we had a polygon complex realizing Σ with more than one polygon. Since Σ is connected, there is a pair of equi-colored edges in two different tiles. We reduce the polygon complex by gluing these two tiles along their this pair of edges. Thereby, we form a bigger polygon. Since this process decreases the number of polygons in the complex, it will stop and we arrive at a polygon diagram for Σ. **q.e.d.**

We can improve upon this quite a bit. Recall that a polygon diagram represents a surface by identification of its edges. Thus certain points on the boundary of the polygon represent identical points in the surface. We call any two such boundary points in a polygon diagram equivalent. We call to edges equivalent if their mid-points are equivalent.

Proposition 5.4.5. *Any surface that is not homeomorphic to the sphere has a polygon diagram all of whose corners are equivalent.*

Definition 5.4.6. Let us call a polygon diagram a one-vertex-diagram if all corners are equivalent.

Proof. Color the corners of the polygon diagram according to their equivalence class. Suppose you need more than one color. In this case, a bigon represents the sphere. Thus, we assume that the polygon has at least four edges.

We will give a procedure for getting rid of any specified color. Suppose, we want to eliminate green. As green is not the only color, there will be an edge connecting a green corner to a corner of a different color, say blue. This edge has a color, say red, which specifies a partner edge. There are two cases. Either the two red edges have a corner in common or not.

Suppose the two edges have a corner in common. Then their arrow either point both toward that corner or away from that corner – this follows from the coloring of the vertices. We can than "swallow" that common corner into the interior of the polygon diagram. The case, where the green vertex is swallowed is shown. In that

Figure 5.1: First case – the two edges have the green corner in common.

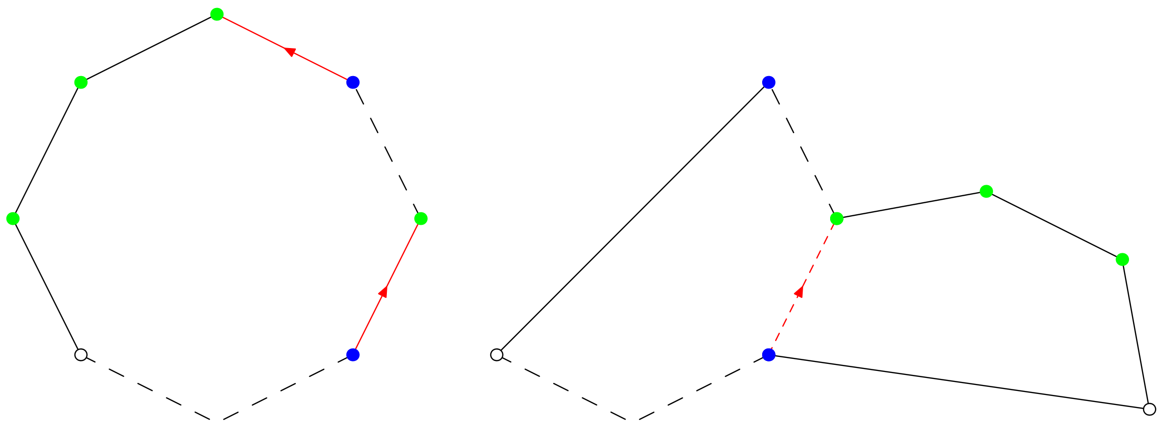

Figure 5.2: Second case – the two edges do not overlap.

case, we reduce the number of green vertices by two. If the blue vertex is swallowed, the number of green vertices decreases by one.

Suppose the two red edges have no common corner. Then pick one of the red edges and move along this edge starting in its blue corner. The next corner you meet is the green corner of this edge. Continue your path along the polygon until you reach the next non-green corner. Cut of this region and glue it to the other red edge. This reduces the number of green corners by one.

Continue this process until all green corners are gone. If you need still more than one color, rid the picture of the next. **q.e.d.**

Definition 5.4.7. A simplicial complex is orientable if all its simplices can be given orientations compatible with the inclusion of faces as subsimplices. Note that subdivisions of simplicial complexes inherit orientations. Thus, orientability of a triangulated surface does not depend on the triangulation. We call a surface orientable if it has an orientable triangulation.

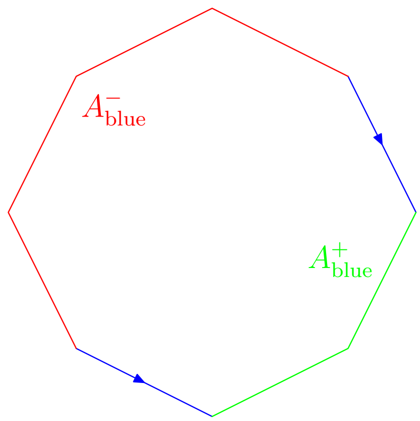

Figure 5.3: The two arcs defined by a pair of equivalent edges.

Remark 5.4.8. Let us discuss orientability of surfaces. Think of a realization of the surface as a polygon complex. Take some big sheet of paper whose two sides are colored red and yellow. Cut out the polygons of the complex. If the edge identifications allow you to glue the pieces so that crossing an edge will never get you from a red side to a green side, then you obtain an oriented surface.

For a polygon diagram, the criterion for orientability given in (5.4.8) is also necessary:

Exercise 5.4.9. Prove: A polygon diagram describes an orientable surface if and only if, for each edge-color a, the two edges of color a are oriented oppositely in the boundary circle of the polygon diagram.

Corollary 5.4.10. *In a one-vertex-diagram for an orientable surface, adjacent edges are inequivalent.*

Proof. Suppose we had a pair of equivalent adjacent edges. Since the surface is orientable, these edges are oppositely oriented. In this case, however, the corner spanned by these two edges is inequivalent to any other corner. Thus, we are not dealing with a one-vertex-diagram. **q.e.d.**

Thus, in a one-vertex-diagram for an orientable surface, any color a defines two edges with opposite orientations that cut the boundary into two non-empty arcs: The arc A^+_a, toward which the edges of color a point, and the arc A^-_a, away from which the edges point. See figure 5.3.

Observation 5.4.11. *In any one-vertex-diagram for an orientable surface and any edge-color a, there is a pair of equivalent edges such that one of them lies on A_a^+ and the other one lies on A_a^-.*

Proof. Suppose not, then the corners at the edges of color a would fall into two distinct equivalence classes. **q.e.d.**

Definition 5.4.12. The genus g standard polygon diagram is the regular $4g$ polygon whose edges are colored with $2g$ colors $a_1, \ldots, a_g$ and $b_1, \ldots, b_g$ and marked so that the boundary reads the word

$$\bullet \xrightarrow{a_1} \bullet \xrightarrow{b_1} \bullet \xleftarrow{a_1} \bullet \xleftarrow{b_1} \bullet \xrightarrow{a_2} \bullet \xrightarrow{b_2} \bullet \xleftarrow{a_2} \bullet \xleftarrow{b_2} \bullet \cdots \bullet \xrightarrow{a_g} \bullet \xrightarrow{b_g} \bullet \xleftarrow{a_g} \bullet \xleftarrow{b_g}$$

The g-torus is the surface obtained from the genus g standard polygon diagram.

Theorem 5.4.13. *Every closed oriented surface is either a sphere or a g-torus for some $g \geq 1$.*

Proof. Let us start with a one-vertex-diagram for the surface. We will use cut and paste to transform the diagram until we obtain a genus g standard polygon diagram.

Let us call a sequence of four edges a *run* if it has the form

$$\bullet \xrightarrow{a} \bullet \xrightarrow{b} \bullet \xleftarrow{a} \bullet \xleftarrow{b} \bullet.$$

If every edge occurs in a run then we have the a standard polygon for some genus. Thus, we want to eliminate edge colors that do not occur in runs. Let a be a color whose corresponding edges do not form a run. By (5.4.11), we know that there is another color b such that the polygon diagram looks essentially like this:

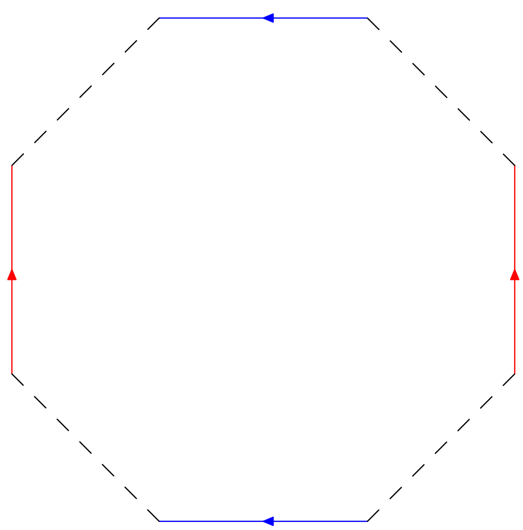

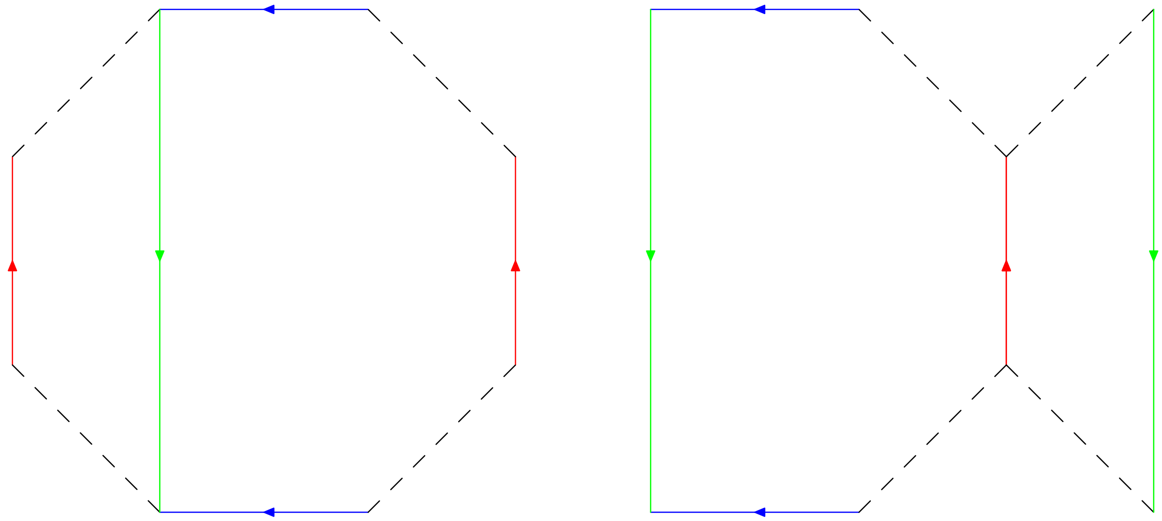

Figure 5.4: The first cut.

We cut and past as illustrated in figure 5.4. Note that the runs in the dashed arcs are not destroyed.

The first cut put us in a situation where we have three edges of a run but the forth partner of the middle edge might be somewhere:

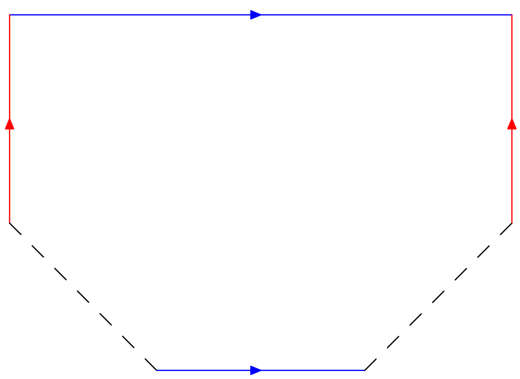

We can create a run by cut and past ash shown in figure 5.5. Again, we do not destroy any runs previously created. **q.e.d.**

Exercise 5.4.14. Show that any non-orientable surface has a one-vertex-diagram whose boundary reads the colors

$$\bullet \xrightarrow{a_1} \bullet \xrightarrow{a_1} \bullet \xrightarrow{a_2} \bullet \xrightarrow{a_2} \bullet \xrightarrow{a_3} \bullet \xrightarrow{a_3} \bullet \cdots \bullet \xrightarrow{a_g} \bullet \xrightarrow{a_g}$$

for some $g \geq 0$.

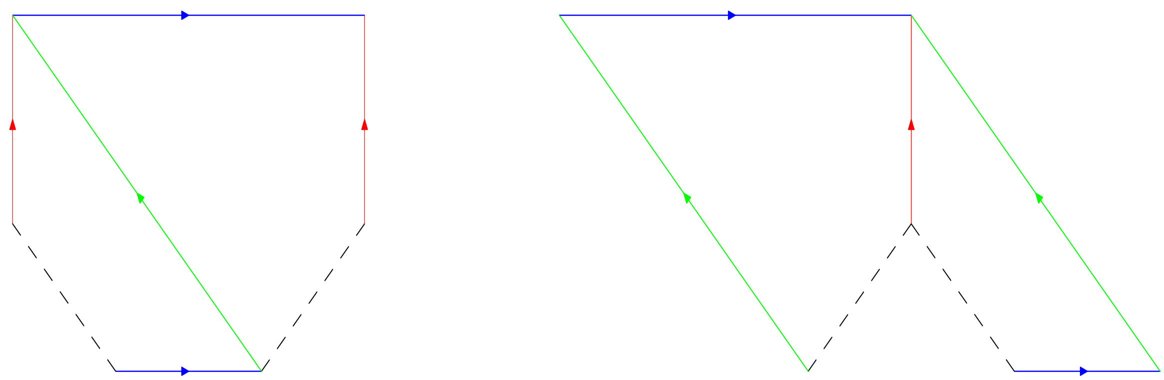

Figure 5.5: The second cut.

Chapter 6

The Torus

6.1 Geometric Structures

Definition 6.1.1. A differentiable structure on a manifold is an atlas maximal with respect to the restriction that all coordinate changes are differentiable maps.

A complex structure on a manifold is an atlas maximal with respect to the restriction that all coordinate changes are holomorphic maps. A map $\xi : \mathbb{R}^2 \to \mathbb{R}^2$ is holomorphic if it is differentiable and at every point, its derivative is a matrix of the form

$$\begin{pmatrix} a & b \\ -b & a \end{pmatrix}.$$

A Euclidean structure on a manifold is an atlas maximal with respect to the restriction that all coordinate changes are locally Euclidean isometries: A map $\xi : U \to V$ between open sets in a Euclidean space $\mathbb{E}$ is locally isometric if each point $x \in U$ has an open neighborhood U_x such that

$$\xi \mid_{U_x} = \lambda \mid_{U_x}$$

for some isometry $\lambda : \mathbb{E} \to \mathbb{E}$.

Given a fixed homeomorphic identification of $\mathbb{R}^m$ with hyperbolic m-space, we can define a hyperbolic structure on a manifold as an atlas maximal with respect to the restriction that all coordinate changes are locally hyperbolic isometries. Here locally isometric maps are defined analogously.

Example 6.1.2. Construct the torus by identifying opposite edges of the unit square. This construction imposes a Euclidean structure on the torus. You can realize it with four charts as show in figure 6.1. The blue chart is drawn completely. For each chart, the dashed area in the plane is the range of the chart map. Observe that intersections

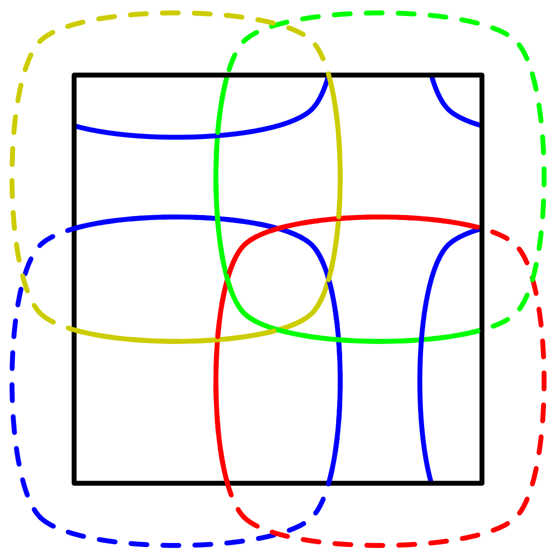

Table 6.1: The "unit square torus"

of chart domains are in general not connected and that the coordinate changes are translations on the components of the intersections.

Example 6.1.3. There is no Euclidean structure on the sphere S_2.

Proof. Suppose there was a Euclidean structure. Since the sphere is compact, we could find a cover of the sphere by finitely many flat triangles (i.e., triangles that are completely contained in a Euclidean chart and look straight in this chart). We find a common subdivision of all these triangles so that we end up with a flat triangulation of the sphere. Let V, E, and F be the number of vertices, edges, and triangles, respectively.

The angle sum around each vertex is 2π, since this is true in $\mathbb{E}^2$. Thus

$$2\pi V = \text{sum of all angles in the triangulation.}$$

On the other hand, we can sum the angles sorted according to the triangles in which they occur. Since the angle sum in a Euclidean triangle is π, we have

$$\pi F = \text{sum of all angles in the triangulation} = 2\pi V.$$

Finally, since each edge has two neighboring triangles and each triangle contains three edges, we have

$$3E = 2F.$$

From this is follows that

$$\begin{aligned} 2\pi\chi(S_2) &= 2\pi V - 2\pi E + 2\pi F \\ &= \pi F - 2\pi\frac{3}{2}F + 2\pi F \\ &= 0 \end{aligned}$$

However, $\chi(S_2) = 2$. **q.e.d.**

Remark 6.1.4. This proof contains some insights that are useful:

1. Any closed surface with a geometric structure has a finite triangulation by geodesic triangles.
2. Only a surface with Euler characteristic 0 can support a Euclidean structure.
3. Only a surface with negative Euler characteristic can support a hyperbolic structure.

It follows from the classification of surfaces, that the only closed surface that admits a Euclidean structure is the torus.

6.1.1 $(\mathcal{I}, \mathcal{X})$-manifolds

Euclidean and hyperbolic structures are just examples of a more general notion.

Definition 6.1.5. Let $\mathcal{X}$ be a fixed m-manifold and let $\mathcal{I}$ be a group of homeomorphisms of $\mathcal{X}$. A $(\mathcal{I}, \mathcal{X})$-chart on an m-manifold M is a pair $(U, \zeta : U \to V \subseteq \mathcal{X})$ where U and V are open sets in M and $\mathcal{X}$ respectively, and $\zeta : U \to V$ is a homeomorphism. A collection of $(\mathcal{I}, \mathcal{X})$-charts forms a $(\mathcal{I}, \mathcal{X})$-atlas if the charts cover M and all coordinate changes

$$\xi : V_0 \to V_1$$

are locally $\mathcal{I}$-maps, i.e., for each point $x \in V_0$ there is a homeomorphism $\xi : \mathcal{X} \to \mathcal{X}$ in $\mathcal{I}$ that equals ξ in an open neighborhood of x.

Exercise 6.1.6. Show that every $(\mathcal{I}, \mathcal{X})$-atlas for M is contained in a unique maximal $(\mathcal{I}, \mathcal{X})$-atlas.

Definition 6.1.5 (continued). A $(\mathcal{I}, \mathcal{X})$-structure for M is a maximal $(\mathcal{I}, \mathcal{X})$-atlas. A $(\mathcal{I}, \mathcal{X})$-manifold is a manifold together with a $(\mathcal{I}, \mathcal{X})$-structure.

Definition 6.1.7. A group $\mathcal{I}$ of homeomorphisms of a topological space $\mathcal{X}$ is rigid for $\mathcal{X}$ if any two homeomorphism ξ_0 and ξ_1 coincide if the coincide on an open subset of $\mathcal{X}$.

Exercise 6.1.8. Show that the full isometry group of Euclidean m-space is rigid for $\mathbb{E}^m$.

Definition 6.1.9. A continuous map $f : M_0 \to M_1$ between Euclidean (hyperbolic) manifolds is geometric if it looks locally like an isometry in local coordinates. That is, for every point $P \in M_0$ there are charts $\varphi_0 : U_0 \to \mathbb{E}^m$ and $\varphi_1 : U_1 \to \mathbb{E}^m$ with $P \in U_0$ and $f(P) \in U_1$ such that

$$\varphi_1 \circ f \circ \varphi_0^{-1}$$

is a local isometry in Euclidean space (hyperbolic space).

Two Euclidean (hyperbolic) structures $\mathcal{G}_0$ and $\mathcal{G}_1$ on M are equivalent if there is a geometric homeomorphism

$$(M, \mathcal{G}_0) \to (M, \mathcal{G}_1).$$

Observe that

$$\varphi_1 \circ f \circ \varphi_0^{-1} = \varphi_1 \circ \left(\varphi_0 \circ f^{-1}\right)^{-1}$$

describes the coordinate change between the given chart in M_1 and a hypothetical chart whose coordinates are given by $\varphi_0 \circ f^{-1}$. Thus, we have a slick way of phrasing this:

Definition 6.1.10. A map $f : M_0 \to M_1$ is a $(\mathcal{X}, \mathcal{I})$-map if around each point $P_0 \in M_0$ there exists a chart $(U_0, \varphi_0 : U_0 \to \mathcal{X})$ such that f maps U_0 homeomorphically to an open set $U_1 \subseteq M_1$ that forms a chart together with the coordinate map

$$\varphi_0 \circ f^{-1} : U_1 \to \mathcal{X}.$$

Remark 6.1.11. It is easy to construct inequivalent Euclidean structures on the torus by rescaling.

Definition 6.1.12. A similarity of a metric space (X, d) is a map

$$\sigma : X \to X$$

for which there is a constant $L > 0$ such that

$$d(\sigma(x)\,, \sigma(y)) = Ld(x, y)\,.$$

Definition 6.1.13. A continuous map $f : M_0 \to M_1$ between Euclidean manifolds is a similarity if it looks locally like an similarity in local coordinates. That is, for every point $P \in M_0$ there are charts $\varphi_0 : U_0 \to \mathbb{E}^m$ and $\varphi_1 : U_1 \to \mathbb{E}^m$ with $P \in U_0$ and $f(P) \in U_1$ such that

$$\varphi_1 \circ f \circ \varphi_0^{-1}$$

extends to a similarity of Euclidean space.

Two Euclidean structures $\mathcal{G}_0$ and $\mathcal{G}_1$ on M are similar if there is a homeomorphism

$$(M, \mathcal{G}_0) \to (M, \mathcal{G}_1)$$

that is a similarity.

Exercise 6.1.14. Find two non-similar Euclidean structures on the 2-dimensional torus.

6.1.2 Developing and Holonomy

From now on, we assume that $\mathcal{I}$ is rigid for $\mathcal{X}$. Let

- P be a point in M and let
- $\varphi : U \to \mathcal{X}$ be a chart around P. Moreover, let
- $p : \mathbb{I} \to M$ be a path in M starting at $p(0) = P$.

We will demonstrate how these data give rise to a unique path in $\mathcal{X}$.

Since $\mathbb{I}$ is compact, there is a partition

$$0 = t_0 < t_1 < t_2 < \cdots < t_r = 1$$

and charts

$$(\varphi_i : U_i \to \mathcal{X})_{i=1}^{r}$$

such that $\varphi = \varphi_1$ and $p([t_{i-1}, t_i]) \subseteq U_i$. Put

$$x_i := p(t_i)$$

and

$$p_i := p\mid_{[t_{i-1}, t_i]} .$$

Since $\mathcal{I}$ is rigid for $\mathcal{X}$, there is a unique

$$\xi_i \in \mathcal{I}$$

that extends the coordinate change $\varphi_i \circ \varphi_{i+1}^{-1}$ around a neighborhood of x_i. Thus, we have

$$\varphi_i = \xi_i \varphi_{i+1}$$

around x_i. Then

$$\tilde{p}_i := \xi_1 \circ \cdots \circ \xi_i \circ \varphi_{i+1} \circ p_i : [t_i, t_{i+1}] \to \mathcal{X}$$

is a path in $\mathcal{X}$. Note that

$$\begin{aligned} \tilde{p}_{i+1}(t_i) &= \xi_1 \circ \cdots \circ \xi_i \circ \xi_{i+1} \circ \varphi_{i+2}(x_i) \\ &= \xi_1 \circ \cdots \circ \xi_i \circ \varphi_{i+1}(x_i) \\ &= \tilde{p}_i(t_i)\,. \end{aligned}$$

Thus, the path $\tilde{p}_{i+1}$ continues where $\tilde{p}_i$ ends. Therefore, we can connect these pieces and obtain a path

$$\tilde{p} : \mathbb{I} \to \mathcal{X}.$$

Definition 6.1.15. This path $\tilde{p}$ is called the continuation of p along the chart $\varphi : U \to \mathcal{X}$.

Lemma 6.1.16. *The continuation $\tilde{p}$ only depends on p and $\varphi : U \to \mathcal{X}$.*

It does not depend on the partition or the chain of charts covering the partition.

Proof. First, we prove that the continuation does not depend on the chain of charts once a partition is fixed. So suppose $\psi : V \to \mathcal{X}$ is another sequence of chart with $\varphi = \psi_1$ and $p([t_{i-1}, t_i]) \subseteq V_i$. Let ζ_i be the induced sequence in $\mathcal{I}$. We claim that

$$\xi_1 \circ \cdots \circ \xi_i \circ \varphi_{i+1} = \zeta_1 \circ \cdots \circ \zeta_i \circ \psi_{i+1}. \tag{6.1}$$

The proof is by induction. The statement is true for $i = 0$ since $\varphi = \varphi_1 = \psi_1$.

Now suppose (6.1) holds for i. Then the diagram (6.2) commutes in the component of $p(t_i)$ in $U_i \cap V_i \cap U_{i+1} \cap V_{i+1}$. Now, the path p_{i+1} is connected whence the diagram commutes all along the segment up to $p(t_{i+1})$.

Now suppose, we also change the partition. Since any two partitions have a common refinement and we can use the charts for a given partition for any refinement, as well, we can find a common third to see that the continuations arising this way are actually equal. **q.e.d.**

Proposition 6.1.17. *Let $\varphi : U \to \mathcal{X}$ be a chart and let $p : \mathbb{I} \to M$ and $q : \mathbb{I} \to M$ be two paths whose endpoints coincide. If p and q are homotopic relative to their endpoints then the endpoints of their continuations coincide and their continuations are homotopic relative to their endpoints.*

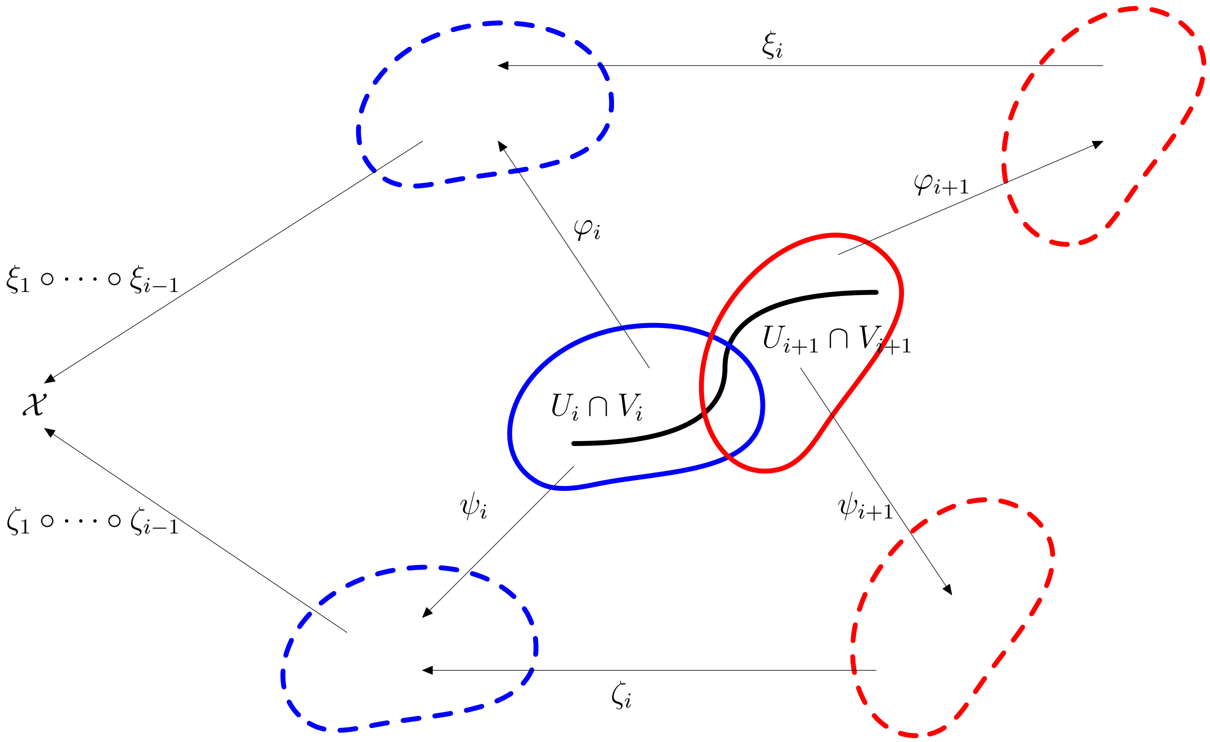

Table 6.2: The diagram commutes near the path because the path is connected.

Proof. This is obvious if the homotopy stays within one coordinate chart. If this is not the case, subdivide the homotopy into small pieces that do. **q.e.d.**

Corollary 6.1.18. *Let M be a 1-connected $(\mathcal{I}, \mathcal{X})$-manifold and $\varphi : U \to \mathcal{X}$ be a chart. Then there is a unique $(\mathcal{I}, \mathcal{X})$-map*

$$\tilde{\varphi} : M \to \mathcal{X}$$

that extends the chart φ.

Proof. Let P_0 be a fixed point in U. For any point $P \in M$, there is a path p connecting P_0 to P. Continuing φ along p, we obtain a value $\tilde{\varphi}(P)$ that is actually independent of p since any two path from P_0 to P are homotopic relative to their endpoints. The map

$$\tilde{\varphi} : M \to \mathcal{X}$$

is clearly a $(\mathcal{I}, \mathcal{X})$-map. This proves existence.

To see that $\tilde{\varphi}$ is unique, suppose that $\nu : M \to \mathcal{X}$ is a $(\mathcal{I}, \mathcal{X})$-map that extends φ. Note that for each open set V in M, the restriction

$$\nu \mid_V : V \to \mathcal{X}$$

is a chart since ν is a $(\mathcal{I}, \mathcal{X})$-map. Thus, we may use these charts to compute the continuation of φ along p. Since we just restrict from a globally defined map, there is no need ever to move path segments to make them fit. Thus all the patching homeomorphisms taken from $\mathcal{I}$ will be trivial and the continuation we obtain from these charts will evaluate to $\nu(P)$ at that end of p. However, the continuation is independent of the charts used to compute it. Thus $\nu(P) = \tilde{\varphi}(P)$. **q.e.d.**

Observation 6.1.19. *Let M be a connected $(\mathcal{I}, \mathcal{X})$-manifold, and let*

$$\nu_0, \nu_1 : M \to \mathcal{X}$$

be two $(\mathcal{I}, \mathcal{X})$-maps. Then there is a unique element $\xi \in \mathcal{I}$ such that

$$\nu_1 = \xi \circ \nu_0.$$

Proof. Let us prove uniqueness first: Suppose we had two elements ξ and ζ such that

$$\nu_1 = \xi \circ \nu_0 \qquad \text{and} \qquad \nu_1 = \zeta \circ \nu_0.$$

Then ξ and ζ agree at least on a small open set in the image of ν_0. Thus, by rigidity of $\mathcal{I}$, we have $\xi = \zeta$.

For existence, observe that ν_0 and ν_1 are both charts (with fairly big domains). Therefore, $\nu_1 \nu_0^{-1}$ is a coordinate change. Hence it is locally represented by elements of $\mathcal{I}$. Since M is connected, there is an element ξ that represents the coordinate change globally. **q.e.d.**

Observation 6.1.20. *Let $f : N \to M$ be a local homeomorphism and M be a $(\mathcal{I}, \mathcal{X})$-manifold. Then there is a unique $(\mathcal{I}, \mathcal{X})$-structure on N that renders f to be a $(\mathcal{I}, \mathcal{X})$-map. This structure is induced by charts of the form*

$$V \xrightarrow{f} U \xrightarrow{\varphi} \mathcal{X}$$

where V is open and homeomorphic via f to the domain of the chart $\varphi : U \to \mathcal{X}$.

In particular, if $\pi : \bar{M} \to M$ is a covering space of a $(\mathcal{I}, \mathcal{X})$-manifold M, then $\bar{M}$ is a $(\mathcal{I}, \mathcal{X})$-manifold in a canonical way.

Let $\bar{M}$ be a 1-connected $(\mathcal{I}, \mathcal{X})$-manifold. Let us suppose that a group G acts on $\bar{M}$ by $(\mathcal{I}, \mathcal{X})$-maps. Moreover, we assume that the action is topologically free, i.e., every point $\bar{P} \in \bar{M}$ has an open neighborhood U such that $gU \cap U = \emptyset$ for all non-trivial $g \in G$. By (6.1.20), the quotient $G \backslash \bar{M}$ is a $(\mathcal{I}, \mathcal{X})$-manifold.

Let $\bar{\varphi} : \bar{U} \to \mathcal{X}$ be a chart in M. By (6.1.18), the chart extends to a $(\mathcal{I}, \mathcal{X})$-map

$$\tilde{\varphi} : \bar{M} \to \mathcal{X}.$$

For any $g \in G$, the map $\tilde{\varphi} \circ g : \bar{M} \to \mathcal{X}$ is also a $(\mathcal{I}, \mathcal{X})$-map. Since $\bar{M}$ is connected, there is a unique $\xi_g \in \mathcal{I}$ such that

$$\tilde{\varphi} \circ g = \xi_g \circ \tilde{\varphi}$$

(see 6.1.19). Note that, by uniqueness,

$$\xi_g \circ \xi_{h^{-1}} \circ \tilde{\varphi} = \tilde{\varphi} \circ g \circ h^{-1} = \tilde{\varphi} \circ \left(gh^{-1}\right) = \xi_{gh^{-1}} \circ \tilde{\varphi}$$

implies

$$\xi_{gh^{-1}} = \xi_g \circ \xi_{h^{-1}}.$$

Clearly, $\xi_1 = \mathrm{id}_{\mathcal{X}}$. Thus, we have defined a group homomorphism

$$\eta_{\tilde{\varphi}} : G \to \mathcal{I}.$$

Definition 6.1.21. The homomorphism η is called the holonomy of M determined by $\tilde{\varphi} : \bar{M} \to \mathcal{X}$

Note that the construction does not require the map $\tilde{\varphi}$ to be a lift of a chart. For any $(\mathcal{I}, \mathcal{X})$-map $\nu : \bar{M} \to \mathcal{X}$, we obtain a holonomy η_ν.

Exercise 6.1.22. Show that for two $(\mathcal{I}, \mathcal{X})$-maps $\nu_0, \nu_1 : \bar{M} \to \mathcal{X}$, there is a unique $g \in \mathcal{I}$ such that

$$\eta_{\nu_0}(h) = g \eta_{\nu_1}(h)\, g^{-1}.$$

Thus, two holonomies differ by an inner automorphism of $\mathcal{I}$.

Example 6.1.23. If M is path connected, then the universal cover $\tilde{M}$ is a 1-connected manifold. Pull back an atlas for M to put a $(\mathcal{I}, \mathcal{X})$-structure on $\tilde{M}$. Then any chart induces a holonomy

$$\eta : \mathrm{Cov}\Big(\tilde{M}/M\Big) \to \mathcal{I}.$$

Thus, we can construct geometric representations of the fundamental group of M.

Definition 6.1.24. Let M be a path connected $(\mathcal{I}, \mathcal{X})$-manifold, and let $\delta : \tilde{M} \to \mathcal{X}$ be a $(\mathcal{I}, \mathcal{X})$-map. Then the induced map

$$\eta_\delta : \mathrm{Cov}\Big(\tilde{M}/M\Big) \to \mathcal{I}$$

characterized by the equation

$$\eta_{\,\delta}(\tau) \circ \delta = \delta \circ \tau \qquad \text{for all } \tau \in \mathrm{Cov}\Big(\tilde{M}/M\Big)$$

is called the holonomy associated to δ. That this equation determines a map follows from (6.1.22). It is easy to check that it is a homomorphism.

If we are given a base point $\underline{\mathrm{P}}$ in M, we have an isomorphism

$$\pi_1(M, \underline{\mathrm{P}}) = \mathrm{Cov}\Big(\tilde{M}/M\Big).$$

Thus, we obtain a homomorphism

$$\eta_\delta : \pi_1(M, \underline{\mathrm{P}}) \to \mathcal{I},$$

which we also call the holonomy.

6.2 Defining Teichmüller Space

Our goal is to classify Euclidean structures on the torus up to equivalence. It turns out that it is easier to classify them up to similarity. So let us build the set up. Let

- $\mathrm{Homeo}(T)$ be the group of self-homeomorphisms on the torus T, and let
- $\mathrm{Homeo}_1(T)$ be the normal subgroup of those homeomorphisms that are homotopic to the identity. The factor group
- $M(T) := \mathrm{Homeo}(T)\,/\mathrm{Homeo}_1(T)$ is called the mapping class group of T.

  Let
- $\mathrm{Isom}(\mathbb{E}^2)$ be the isometry group of the plane. This is a normal subgroup in the group
- $\mathrm{Sim}(\mathbb{E}^2)$ of similarities. Note that $\mathrm{Sim}(\mathbb{E}^2)$ acts from the left on
- $\mathcal{E}(T)$, the set of Euclidean structures on T. The action is given by modifying all the charts, appending the similarity $\sigma \in \mathrm{Sim}(\mathbb{E}^2)$. Put

- $\mathcal{S}(T) := \mathrm{Sim}(\mathbb{E}^2) \setminus \mathcal{E}(T)$.

  Note that Homeo(T) acts on $\mathcal{E}(T)$ from the right as follows: For a homeomorphism $\zeta : T \to T$, a given Euclidean structure $\mathcal{E}$ on T and a chart $\varphi : U \to \mathbb{E}^2$ for this structure, define a corresponding chart

$$\varphi \circ \zeta : \zeta^{-1}(U) \to \mathbb{E}^2.$$

  All these charts form a new atlas for T and define a different Euclidean structure $\mathcal{E}\zeta$. Note that

$$\zeta : (T, \mathcal{E}) \to (T, \mathcal{E}\zeta)$$

  is an equivalence of Euclidean structures. This action induces an action of Homeo(T) on Sim(T).

  The quotient

- $\mathcal{M}_T := \mathcal{S}(T)\,/\mathrm{Homeo}(T)$ is called the moduli space of T and the quotient

- $\mathcal{T}_T := \mathcal{S}(T)\,/\mathrm{Homeo}_1(T)$ is called the Teichmüller space of T. Note that there is a natural action of $M(T)$ on $\mathcal{T}_T$ such that

$$\mathcal{M}_T = \mathcal{T}_T / M(T)\,.$$

Lemma 6.2.1. *A similarity σ of a complete metric X space with scale factor $\neq 1$ has a unique fixed point.*

Proof. Passing to σ^{-1} if necessary, we can assume that the scale factor is < 1. Then the sequence

$$\sigma^i(x)$$

is Cauchy. Its limit is a fixed point. Moreover, the fixed point is unique since the distance between two fixed points cannot shrink under σ. **q.e.d.**

Lemma 6.2.2. *Let M be a connected $(\mathcal{I}, \mathcal{X})$-manifold. Let $\mathcal{H}$ be a subgroup of $\mathcal{I}$. Then the $(\mathcal{I}, \mathcal{X})$-structure on M contains an $(\mathcal{H}, \mathcal{X})$-structure for M if and only if the image of a holonomy $\eta : \pi_1(M) \to \mathcal{I}$ is contained in $\mathcal{H}$.*

Proof. The condition is clearly necessary.

So let us suppose that we have a holonomy $\eta : \pi_1(M) \to \mathcal{H}$ defined by some developing map $\nu : \tilde{M} \to \mathcal{X}$.

Let $\{\varphi_i : U_i \to \mathcal{X}\}$ be an $(\mathcal{I}, \mathcal{X})$-atlas wherein each U_i is evenly covered. The sheets above the chart define a $(\mathcal{I}, \mathcal{X})$-atlas for $\tilde{M}$.

Now, define new charts

$$\psi_i : U_i \to \mathcal{X}$$

by

$$\psi_i := \nu \circ \zeta_i$$

where $\zeta : U_i \to \tilde{M}$ is a homeomorphism identifying U_i with a sheet above it.

Let us consider a coordinate change:

$$\psi_i \circ \psi_j^{-1} = \nu \circ \zeta_i \circ \zeta_j^{-1} \circ \nu^{-1}.$$

Since $\zeta_i \circ \zeta_j^{-1}$ is a covering transformation, this coordinate change is locally $\nu \circ \tau \circ \nu^{-1}$ which is an element of H. Thus, $\{\psi_i : U_i \to \mathcal{X}\}$ is an $(\mathcal{H}, \mathcal{X})$-atlas for M. It is easy to check that it is compatible with the given $(\mathcal{I}, \mathcal{X})$-structure. **q.e.d.**

Proposition 6.2.3. *The set $\mathcal{S}(T)$ can be identified with the set of all* $(\mathrm{Sim}(\mathbb{E}^2), \mathbb{E}^2)$*-structures on T.*

Proof. Every Euclidean structure on T induces a $(\mathrm{Sim}(\mathbb{E}^2), \mathbb{E}^2)$-structure in an obvious way. Moreover, two Euclidean structures that represent the same class in $\mathcal{S}(T)$ clearly define identical $(\mathrm{Sim}(\mathbb{E}^2), \mathbb{E}^2)$-structures. Thus, it remains to prove that every $(\mathrm{Sim}(\mathbb{E}^2), \mathbb{E}^2)$-structure is induced by a Euclidean structure on T.

We claim that the maximal atlas for any $(\mathrm{Sim}(\mathbb{E}^2), \mathbb{E}^2)$-structure contains a subatlas that defines a Euclidean structure. By (6.2.2), we have to show that a holonomy takes values in $\mathrm{Isom}(\mathbb{E}^2)$.

Let $\nu : \tilde{T} \to \mathbb{E}^2$ be any $(\mathrm{Sim}(\mathbb{E}^2), \mathbb{E}^2)$-map. We turn it into an isometry by pulling back the metric from $\mathbb{E}^2$. So let

$$\tau : \tilde{T} \to \tilde{T}$$

be a deck transformation. Thus, τ is a similarity in local coordinates. Since $\tilde{T}$ is connected, all the scale factors agree. Using the triangle inequality for τ and τ^{-1}, we see that τ is a similarity of $\tilde{M}$. If τ is non-trivial, it has no fixed points and must be an isometry, being a similarity already.

The holonomy defined by ν sends a deck transformation τ to the similarity $\sigma : \mathbb{E}^2 \to \mathbb{E}^2$ that satisfies

$$\nu \circ \tau = \sigma\nu.$$

However, ν and τ are local isometries. Hence, so is σ. **q.e.d.**

6.3 The Dehn-Nielsen Theorem for the Torus

We fix a base point $\underline{t}$ on the torus. Let $\zeta : T \to T$ be a homeomorphism, and let $p : \mathbb{I}$ be a path from $\underline{t}$ to $\zeta(\underline{t})$. Then, we have two group homomorphisms:

$$\begin{aligned} \zeta_* : \pi_1(T, \underline{t}) &\to \pi_1(T, \zeta(\underline{t})) \\ [\gamma] &\mapsto [\zeta \circ \gamma] \end{aligned}$$

and

$$\begin{aligned} p_* : \pi_1(T, \zeta(\underline{t})) &\to \pi_1(T, \underline{t}) \\ [\gamma] &\mapsto [p \rightarrowtail \zeta \circ \gamma \rightarrowtail p^{\mathrm{rev}}] \,. \end{aligned}$$

Observation 6.3.1. *If q is another path from $\underline{t}$ to $\zeta(\underline{t})$, the two homomorphisms p_* and q_* differ by an inner automorphism of $\pi_1(T, \underline{t})$ given by the loop $p \rightarrowtail q^{\mathrm{rev}}$.* **q.e.d.**

Observation 6.3.2. *If $\xi : T \to T$ is a homeomorphism homotopic to ζ via a homotopy $\Phi : T \times \mathbb{I} \to T$, then*

$$p_* \circ \zeta_* = (p \rightarrowtail q)_* \circ \xi_*$$

where q is the path from $\zeta \underline{t}$ to $\xi(\underline{t})$ given by $\Phi(\underline{t}, -)$. **q.e.d.**

Thus, we obtain a well defined map

- $\nu : \mathcal{M}(T) \to \mathrm{Out}(\pi_1(T, \underline{t}))$.

Theorem 6.3.3 (Dehn-Nielsen). *The map ν is an isomorphism of groups.*

Proof. First, let us check that ν is a homomorphism of groups. So let ζ and ξ be two homeomorphism of the torus T. We choose paths p and q from $\underline{t}$ to $\zeta(\underline{t})$ and $\xi(\underline{t})$, respectively. Then

$$\begin{aligned} \nu(\zeta)\,\nu(\xi) &= [p_* \circ \zeta_*]\,[q_* \circ \xi_*] \\ &= \big[(p \rightarrowtail \zeta \circ q)_* (\zeta \circ \xi)_*\big] \\ &= \nu(\zeta \circ \xi)\,. \end{aligned}$$

Now, we show that ν is injective. So let $\zeta : T \to T$ be a homeomorphism with $\nu([\zeta]) = 1$. So, for any path p from $\underline{t}$ to $\zeta(\underline{t})$, the homomorphism $p_* \circ \zeta_*$ is an inner automorphism of $\pi_1(T, \underline{t})$. We have to show that ζ is homotopic to the identity.

Let γ be a loop based at $\underline{t}$ such that $p_* \circ \zeta_*$ is conjugation by γ. This is to say that

$$\gamma_* = p_* \circ \zeta_*.$$

Put $q := \gamma^{\text{rev}} \rightarrowtail p$. Then

$$q_* \circ \zeta_* = {\gamma^{\text{rev}}}_* \circ p_* \circ \zeta_* = 1.$$

Thus, for any loop γ', the curve $q \rightarrowtail \zeta \circ \gamma' \rightarrowtail q^{\text{rev}}$ is homotopic to γ'. We apply this result to the two fundamental curves γ_1 and γ_2 on T. We obtain the following map on the surface of a cube:

On the front, we have the standard identification map $\mathbb{I}^2 \to T$. In the back, we have the composition $\mathbb{I}^2 \to T \xrightarrow{\zeta} T$. The four faces in the *boundary annulus* are filled by the homotopies

$$\gamma_i \sim q \rightarrowtail \zeta \circ \gamma_i \rightarrowtail q^{\text{rev}}.$$

This is a map defined on the two-dimensional sphere $S_2 \to T$. Since T is aspherical, it extends to a map on the ball. Moreover, note that opposite faces along the boudary annulus are mapped identically, we actually can make face identifications and obtain a map

$$T \times \mathbb{I} \to T$$

that visibly gives a homotopy from the identity (front) to ζ (back).

Finally, we observe that ν is onto. We know that $\mathrm{Out}(\pi_1(T, \underline{t})) = \mathrm{GL}_2(\mathbb{Z})$. The action of $\mathrm{GL}_2(\mathbb{Z})$ on the plane $\mathbb{R}^2$ immediately descends to an action on the torus by homeomorphisms. This gives an inverse to ν. **q.e.d.**

Lemma 6.3.4. *The torus is aspherical.*

Proof. The sphere is 1-connected. Hence any map to the torus lifts to the universal cover, which is the plane. The lift extends to a map on the ball, and so does the original map. **q.e.d.**

Corollary 6.3.5. *Let $\tilde{\zeta} : \tilde{T} \to \tilde{T}$ be a homeomorphism that commutes with all deck transformations, i.e., the following diagram commutes for all deck transformations $\tau : \tilde{T} \to \tilde{T}$:*

$$\begin{array}{ccc} \tilde{T} & \xrightarrow{\tilde{\zeta}} & \tilde{T} \\ \uparrow{\scriptstyle\tau} & & \uparrow{\scriptstyle\tau} \\ \tilde{T} & \xrightarrow{\tilde{\zeta}} & \tilde{T} \end{array}$$

Then $\tilde{\zeta}$ induces a homeomorphism $\zeta : T \to T$, which is homotopic to the identity.

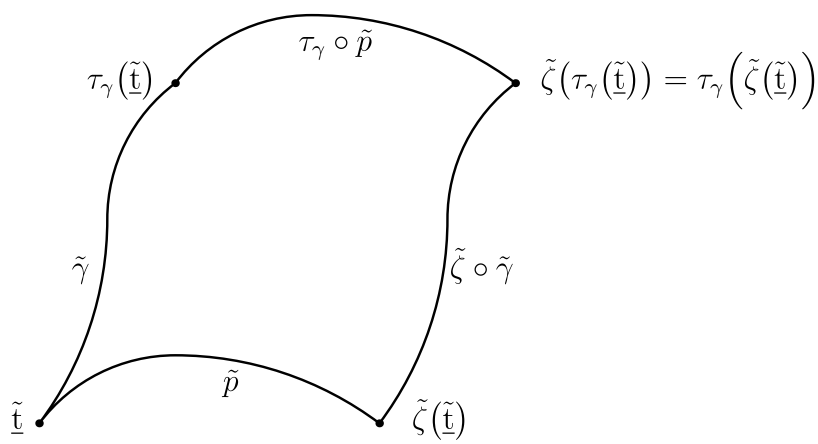

Figure 6.1: A closed path in $\tilde{T}$

Proof. It is easy to see that $\tilde{\zeta}$ induces a homeomorphism ζ of T. We will only show that ζ is homotopic to the identity. By the Dehn-Nielsen Theorem (6.3.3), it suffices to prove that $\nu(\zeta)$ is the class of inner automorphisms of $\pi_1(T)$.

Fix a path $\tilde{p}$ in $\tilde{T}$ from the base point $\underline{\tilde{t}}$ to $\tilde{\zeta}(\underline{\tilde{t}})$. For any loop γ in T based at $\underline{t} = \pi(\underline{\tilde{t}})$, let $\tilde{\gamma}$ be the lift of γ based at $\underline{\tilde{t}}$. This lift is a path from $\underline{\tilde{t}}$ to $\tau_\gamma(\underline{\tilde{t}})$ where τ_γ is the deck transformation corresponding to γ.

From
$$\left(\tilde{\zeta} \circ \tau_\gamma\right)(\underline{\tilde{t}}) = \left(\tau_\gamma \circ \tilde{\zeta}\right)(\underline{\tilde{t}}),$$
it follows that
$$\tilde{\gamma} \rightarrowtail \tau_\gamma \circ \tilde{p} \rightarrowtail \left(\tilde{\zeta} \circ \tilde{\gamma}\right)^{\mathrm{rev}} \rightarrowtail \tilde{p}^{\mathrm{rev}}$$
is a closed path in $\tilde{T}$. See figure 6.1. Thus,
$$\gamma \sim p \rightarrowtail \zeta \circ \gamma \rightarrowtail \gamma^{\mathrm{rev}}.$$
Thus, $p_* \circ \zeta_*$ is the identity automorphism of $\pi_1(T, \underline{t})$. **q.e.d.**

Definition 6.3.6. Let X and Y be topological spaces. A map $f : X \to Y$ is a homotopy equivalence if there is a map $h : Y \to X$ such that $h \circ f : X \to X$ is homotopic to the identity on X and $f \circ h : Y \to Y$ is homotopic to the identity on Y. (I.e., a homotopy equivalence is a map that induces an isomorphism in the homotopy category :-)

Two spaces X and Y are called homotopy equivalent if there is a homotopy equivalence $f : X \to Y$.

Note that any homotopy equivalences $f : X \to Y$ induces an isomorphism

$$f_* : \pi_1(X, \underline{x}) \to \pi_1(Y, f(\underline{x})) .$$

In particular, any self homotopy equivalence $f : X \to X$ induces an outer automorphism $\nu(f)$ of $\pi_1(X)$ in the same way that a self homeomorphism does.

Exercise 6.3.7. Let $f : T \to T$ be a self homotopy equivalence of the torus. Assume $\nu(f)$ is the class of inner automorphisms of $\pi_1(T)$. (This is, given a path p from $\underline{t}$ to $f(\underline{t})$, the homomorphism $p_* \circ f_*$ is an inner automorphism of $\pi_1(T, \underline{t})$.) Prove that f is homotopic to the identity.

Exercise 6.3.8. Let $f : T \to T$ be a self homotopy equivalence of the torus. Show that f is homotopic to a homeomorphism.

Exercise 6.3.9. Construct an example of a self homotopy equivalence of a finite graph Γ that is not homotopic to a homeomorphism.

Exercise 6.3.10. Let Γ be a finite graph. Let $f_0, f_1 : \Gamma \to \Gamma$ be two self homotopy equivalences of Γ. Prove that f_0 is homotopic to f_1 if $\nu(f_0) = \nu(f_1)$.

Exercise 6.3.11. Let R_n be the graph with one vertex v and n loops attached to the vertex. Show that every automorphism of $F_n := \pi_1(R_n, v)$ arises as an f_* for some homotopy equivalence

$$f : R_n \to R_n.$$

Exercise 6.3.12. Show that two finite graphs are homotopy equivalent if they have isomorphic fundamental groups.

6.4 Calculating Teichmüller Space

Theorem 6.4.1. *The map*

$$\begin{aligned} \Psi : \mathcal{T}_T &\to \mathcal{D}_T \\ [\mathcal{E}] &\mapsto \left[\eta_{\mathcal{E}}^{\delta}\right] \end{aligned}$$

is a bijection.

Proof of Injectitivity. Suppose we have two Euclidean structures $\mathcal{E}_1$ and $\mathcal{E}_2$ on T such that

$$\left[\eta_{\mathcal{E}_1}^{\delta_1}\right] = \left[\eta_{\mathcal{E}_2}^{\delta_2}\right] .$$

Then there is a similarity $\sigma : \mathbb{E}^2 \to \mathbb{E}^2$ such that, for each loop γ, the following diagram commutes:

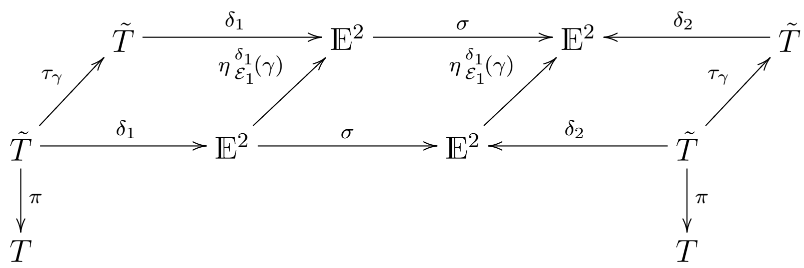

By (6.3.5), it follows that $\delta_2^{-1} \circ \sigma \circ \delta_1$ induces a homeomorphism $\zeta : T \to T$ that is homotopic to the identity. It is easy to check that all these diagrams add up to:

$$\sigma \mathcal{E}_1 \zeta = \mathcal{E}_2$$

Thus, $[\mathcal{E}_1] = [\mathcal{E}_2]$. **q.e.d.**

Exercise 6.4.2. Let T_1 and T_2 be two tori, and let $\phi : \operatorname{Cov}\left(\tilde{T}_1/T_1\right) \to \operatorname{Cov}\left(\tilde{T}_1/T_2\right)$ be an isomorphism. Show that there exists a homeomorphism $\tilde{\zeta} : \tilde{T}_1 \to \tilde{T}_2$ of the universal covers such that $\tilde{\zeta} \circ \tau = \phi(\tau) \circ \tilde{\zeta}$ holds for each deck transformation $\tau \in \operatorname{Cov}(T_1)$,

Proof of Surjectivity. Let

$$\eta : \pi_1(T) \to \operatorname{Isom}\left(\mathbb{E}^2\right)$$

be a discrete, injective homomorphism. We have seen already that η factors through the group of translations. Thus, the image $G := \operatorname{im}(\eta)$ is a free abelian group generated by two linearly independent translations that acts on $\mathbb{E}^2$ topologically freely. Hence the quotient $G \backslash \mathbb{E}^2$ is a torus. This torus comes with a Euclidean structure. The idea is, of course, to transfer this structure to T.

By (6.4.2), there is a homeomorphism

$$\tilde{\zeta} : \tilde{T} \to \mathbb{E}^2$$

such that

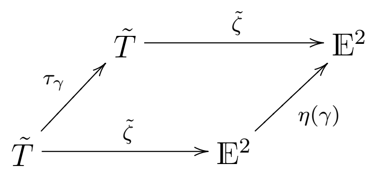

commutes for any loop γ. Thus, we can use $\tilde{\zeta}$ to define a Euclidean structure on $\tilde{T}$ which actually descends to a Euclidean Structure on T. Using $\tilde{\zeta}$ as our developing map, we see that this structure induces the holonomy η. **q.e.d.**

6.5 Classification of Homeomorphisms

Let us consider orientation preserving homeomorphisms of the torus T up to homotopy. They form the group

$$\mathrm{SL}_2(\mathbb{Z}).$$

A matrix $M \in \mathrm{SL}_2(\mathbb{Z})$ can be studied by looking at its trace.

Definition 6.5.1. M is elliptic if $|\mathrm{tr}(M)| < 2$.
M is parabolic if $|\mathrm{tr}(M)| = 2$.
M is hyperbolic if $|\mathrm{tr}(M)| > 2$.

The significance lies in the fact that the characteristic polynomial of M is given by:

$$\mathrm{Det}(M) - \lambda\,\mathrm{tr}(M) + \lambda^2.$$

Thus we have:

M is elliptic: In this case, we have two complex conjugate eigenvalues λ_1, λ_2. Thus there is a fixed point in Teichmüller space. There are only three possibilities:

$\mathrm{tr}(M) = 0$: We find $\lambda_1 = \mathrm{i}$ and $\lambda_2 = \mathrm{i}$. Thus, the matrix has order four. This homeomorphism is the rotation by $\frac{\pi}{2}$. The standard Euclidean structure (identify opposite edges of a unit square) is the fixed point in Teichmüller space.

$\mathrm{tr}(M) = -1$: Here λ_i is a sixth root of unity and M has order six. Here, we expect the homeomorphism to be a rotation by $\frac{\pi}{3}$. Moreover, the Euclidean structure should correspond to a shape of a fundamental domain. Thus, we represent the torus as a regular hexagon with opposite edges identifed. A rotation around the center is our homeomorphism. It might take you some time to convice yourself that a hexagon really gives a torus when you identify opposite edges. You can see this, however, from the induced tessalation of the Euclidean plane by hexagons.

$\mathrm{tr}(M) = 1$: Finally, we find λ_i is a third root of unity, and M has order three. This homeomorphism is the square of the previous one.

Thus, elliptic elements are periodic. They have finite order.

Moreover, since one the eigenvalues lies in the upper half plane, there is a fixed point in Teichmüller space.

M is parabolic: Here $\lambda_1 = \lambda_2 = 1$. Consider the action on $\mathbb{R}^2$. There is an eigenspace. Since the eigenvalue is 1, this line is fixed point wise. Will it descend to a closed curve on the torus $\mathbb{R}^2/\mathbb{Z}^2$?

We have to check if the slope of this line is rational. This follows from the fact that all entries in M are rational whence we can interpret the singularity of $M - \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$ over $\mathbb{Q}$. Thus, this homeomorphism fixes a closed curve on the torus.

Note that $M(T) = \mathrm{SL}_2(\mathbb{Z})$ acts transitively on $\mathbb{P}^1(\mathbb{Q})$. Thus, all parabolic elements are conjugate to elements of the form

$$\pm \begin{pmatrix} 1 & s \\ 0 & 1 \end{pmatrix}.$$

Thus, parabolic elements are powers of Dehn twists.

M is hyperbolic: Now we have two real eigenvalues. Their product is $\mathrm{Det}(M) = 1$. The eigenvalues are not rational. For suppose the eigenspaces had rational slope. Then they would descend to closed curves on the torus and in the universal cover these would be represented by a family of parallel lines. However, the contracting eigenvalue would have to shrink the lattice of intersections with its eigenspace. Hence the pattern upstairs can only be invariant if it is dense.

Thus, the invariant lines descend to bi-infinite geodesic curves on the torus. This is the most elementary example of a geodesic lamination. Thus, a hyperbolic homeomorphism leaves invariant a pair of geodesic laminations.

Chapter 7

Higher Genus Surfaces

7.1 The Main Result

We will outline two proofs of the main theorem:

Theorem 7.1.1. *Let Σ be a closed oriented surface of genus $g > 1$. Then every homotopy class of homeomorphisms has a representative $\zeta : \Sigma \to \Sigma$ satisfying one of the following conditions:*

elliptic case: *The homeomorphism has finite order, i.e., $\zeta^k = \mathrm{id}_\Sigma$.*

hyperbolic case: *The homeomorphism leaves a pair of geodesic laminations on Σ invaraint.*

parabolic case: *There is a non-empty collection of simple closed cuves on Σ that is left invariant as a subset of Σ. In this case, a power of ζ fixes the curves point wise.*

Definition 7.1.2. For a closed oriented surface of genus $g > 1$, Teichmüller space is defined as

$$\mathcal{T}_\Sigma = {\{\text{hyperbolic structures on } \Sigma\}} \Big/ {\mathrm{Homeo}_1(\Sigma)} \; .$$

The main problem to overcome in both proofs is that the action of $M(\Sigma)$ on $\mathcal{T}_\Sigma$ is not cocompact. There are two main strategies to overcome this obstacle:

- Restrict your attention to a cocompact subspace of $\mathcal{T}_\Sigma$.
- Compactify $\mathcal{T}_\Sigma$ so that the action of $M(\Sigma)$ extends to the compactification.

7.1.1 First Proof: Cutting off Infinity

Promise 1. There is a metric on Teichmüller space $\mathcal{T}_\Sigma$ such that:

1. $\mathcal{T}_\Sigma$ is a geodesic metric space.
2. Geodesics are unique.
3. Local geodesics are global.
4. The action of $M(\Sigma)$ on $\mathcal{T}_\Sigma$ is by isometries.

Thus, $\mathcal{T}_\Sigma$ is a proper metric space and uniquely geodesic.

Definition 7.1.3. Let X be a metric space and $\lambda : X \to X$ be an isometry. The displacement function of λ is

$$\begin{aligned} D_\lambda : X &\to \mathbb{R} \\ x &\mapsto d_X(x, \lambda(x)) \,. \end{aligned}$$

The displacement of λ is

$$D(\lambda) := \inf_{x \in X} D_\lambda(x) \,.$$

The displacement is realized if there is a point $x \in X$ such that

$$D(\lambda) = D_\lambda(x) \,.$$

Fix a homeomorphism

$$\zeta : \Sigma \to \Sigma,$$

which induces an isometry λ_ζ on Teichmüller space by

$$\lambda_\zeta : [\mathcal{H}] \mapsto [\mathcal{H}\zeta] \,.$$

There are three cases:

- The displacement is realized and equals 0.
- The displacement is realized and strictly positive.
- The displacement is not realized.

The Displacement is Realized and Vanishes

Let $\mathcal{H}$ be a hyperbolic structure on Σ such that $[\mathcal{H}] \in \mathcal{T}_\Sigma$ realizes the displacement 0. Note that this point is a fixed point of ζ:

$$[\mathcal{H}] = [\mathcal{H}\zeta] \,.$$

Thus there is a homeomorphism $\xi : \Sigma \to \Sigma$ homotopic to the identity such that

$$\mathcal{H}\xi = \mathcal{H}\zeta.$$

Therefore, $\zeta \circ \xi^{-1}$ is an isometry of $(\Sigma, \mathcal{H})$. Since ξ is homotopic to the identity, we conclude that ζ is homotopic to an isometry of $(\Sigma, \mathcal{H})$. This isometry has finite order:

Promise 2. Any isometry of an oriented closed hyperbolic surface has finite order.

The Displacement is Realized and Strictly Positive

Our first goal is to construct a geodesic that is fixed by λ_ζ:

Lemma 7.1.4. *Let X be a geodesic metric space and $\lambda : X \to X$ be an isometry whose displacement is strictly positive and realized at a point $x \in X$. Then*

$$l := \bigcup_{k\in\mathbb{Z}} \left[X, \lambda^k(x)\right] \lambda^{k+1}(x) = \bigcup_{k\in\mathbb{Z}} \lambda^k \left[X, x\right] \lambda(x)$$

is locally a geodesic.

Proof. We know that l is geodesic at all points in the interior of $[x, \lambda(x)]$. Since λ preserves being locally geodesic, it suffices to show that l is geodesic at $\lambda(x)$.

Consider the midpoint y of $[x, \lambda(x)]$. Observe that

$$D(\lambda) \leq d(y, \lambda(y)) \leq d(y, \lambda(x)) + d(\lambda(x)\,, \lambda(y)) \leq d(x, \lambda(x)) = D(\lambda)\,.$$

Thus l is geodesic at $\lambda(x)$. **q.e.d.**

This construction applies to Teichmüller space and yields are global bi-infinite geodesic C by (1(3)). Note that this geodesic is invariant with respect to λ_ζ.

This is the hyperbolic case:

Promise 3. Every geodesic in Teichmüller space $\mathcal{T}_\Sigma$ gives rise to a pair of transverse geodesic laminations.

The Displacement is Not Realized

Definition 7.1.5. A metric space is proper if closed balls are compact.

Exercise 7.1.6. Show that a metric space is proper if an only if:

$$\text{compact} \Longleftrightarrow \text{closed and bounded}$$

Exercise 7.1.7. Show that a geodesic metric space is proper if it is complete and locally compact.

Definition 7.1.8. A group G acts properly discontinuously on a topological space X if for every compact subset $C \subseteq X$, the set

$$\{g \in G \mid gC \cap C \neq \emptyset\}$$

is finite.

Remark 7.1.9. A properly discontinuous action is a topological analogue of an action with finite stabilizers.

We already know that the mapping class group does not act freely on Teichmüller space.

Promise 4. Teichmüller space is a complete, locally compact, proper metric space, and the action of the mapping class group acts properly discontinuously on Teichmüller space.

We need a big theorem. For any $\varepsilon > 0$ let $\mathcal{T}_\varepsilon$ be the subset of $\mathcal{T}_\Sigma$ of those hyperbolic structures for which the length of all closed geodesics in Σ are bounded from below by ε. Note that $\mathcal{T}_\varepsilon$ is $M(\Sigma)$-invariant.

Promise 5 (Mumford's Compactness Theorem). For each $\varepsilon > 0$, there is a compact subset $C_\varepsilon \subset \mathcal{T}_\Sigma$ such that

$$\mathcal{T}_\varepsilon = C_\varepsilon M(\Sigma)\,.$$

In fact, C_ε can be taken to be a fundamental domain for the action.

Let us choose a sequence of hyperbolic structures $(\mathcal{H}_i)$ such that

$$d([\mathcal{H}_i]\,,[\mathcal{H}_i\zeta]) \to D\lambda_\zeta \qquad \text{as } i \to \infty.$$

Lemma 7.1.10. *There is no $\varepsilon > 0$ such that $[\mathcal{H}_i] \in \mathcal{T}_\varepsilon$ for all i.*

Proof. We argue by contradiction. So suppose $[\mathcal{H}_i] \in \mathcal{T}_\varepsilon$ for all i. Then we can find a sequence $\xi_i \in M(\Sigma)$ such that

$$[\mathcal{H}_i \xi_i] \in C_\varepsilon.$$

Note that the sequence

$$d([\mathcal{H}_i], [\mathcal{H}_i \zeta]) = d([\mathcal{H}_i \xi_i], [\mathcal{H}_i \zeta \xi_i])$$

is bounded. Thus the points

$$[\mathcal{H}_i \zeta \xi_i] = \left[\mathcal{H}_i \xi_i \circ \xi_i^{-1} \circ \zeta \circ \xi_i\right]$$

stays within bounded distance from the compact set C_ε. Thus we can pass to a subsequence such that simultaneously

$$[\mathcal{H}_i \xi_i] \to \mathcal{H}_+$$

and

$$\left[\mathcal{H}_i \xi_i \circ \xi_i^{-1} \circ \zeta \circ \xi_i\right] \to \mathcal{H}_*.$$

Observe that the isometries $\xi_i^{-1} \circ \zeta \circ \xi_i$ take points close to $\mathcal{H}_+$ to points close to $\mathcal{H}_*$. Since the mapping class group acts properly discontinuously on Teichmüller space, it follows that there are only finitely many elements in $M(\Sigma)$ that do this. By the box principle, one of these occurs infinitely many times in the sequence $\xi_i^{-1} \circ \zeta \circ \xi_i$. Let this isometry be $\xi^{-1} \circ \zeta \circ \xi$. Since

$$d([\mathcal{H}_*], [\mathcal{H}_*]) = D(\zeta)$$

it follows that the displacement of ζ is realized at

$$\left[\mathcal{H}_+ \xi^{-1}\right].$$

q.e.d.

Definition 7.1.11. The spectrum of a hyperbolic structure $\mathcal{H}$ on Σ is the set

$$\Sigma(\mathcal{H}) := \{\ln(\gamma) \mid \gamma \text{ is a simple closed geodesic in } \Sigma\}.$$

Promise 6. For any hyperbolic surface, closed geodesics of length less than $3 + \sqrt{2}$ do not intersect.

Promise 7. Any collection of pairwise non intersecting non-homotopic loops on a surface of genus g has at most $3g - 3$ elements.

Corollary 7.1.12. *For any hyperbolic structure $\mathcal{H}$,*

$$\left|\Sigma(\mathcal{H}) \cap \left(-\infty, \ln\left(3+\sqrt{2}\right)\right]\right| \leq 3g-3. \qquad \textbf{q.e.d.}$$

Promise 8. Let γ be a simple closed curve on Σ that is not homotopically trivial. For each hyperbolic structure $\mathcal{H}$, there is a unique geodesic $\gamma_{\mathcal{H}}$ homotopic to γ. Moreover, the map

$$\ell_\gamma : [\mathcal{H}] \mapsto \ln(\text{lenght of } \gamma_{\mathcal{H}})$$

is well defined and satisfies the inequality

$$|\ell_\gamma([\mathcal{H}_1]) - \ell_\gamma([\mathcal{H}_2])| \leq d_{\mathcal{T}_\Sigma}([\mathcal{H}_1], [\mathcal{H}_2]).$$

Choose L greater than all $D_{\lambda_\zeta}([\mathcal{H}_i])$. Since no $\mathcal{T}_\varepsilon$ contains all $[\mathcal{H}_i]$, it follows that there is an index i for which

$$\Sigma(\mathcal{H}_i) = M \uplus N$$

with

- $M \neq \emptyset$.
- $\sup(M) < \ln\left(3+\sqrt{2}\right)$.
- $\sup(M) + L < \inf N$.

We claim that the curves from which the lengths in M arise form an invariant system. Let Δ denote the set of homotopy classes of those closed geodesics.

Observe that

$$\Sigma(\mathcal{H}) = \Sigma(\mathcal{H}\zeta) = M \uplus N.$$

Thus, we may ask whether ζ respects the decomposition into M and N. The answer is "yes" because of (8): The curves γ in Δ are those with logarithmic length relative to $\mathcal{H}_i$ in M:

$$\ell_\gamma \mathcal{H}_i \in M.$$

Since

$$|\ell_\gamma \mathcal{H}_i - \ell_\gamma \mathcal{H}_i \zeta| \leq d(\mathcal{H}_i, \mathcal{H}_i \zeta) \leq L,$$

it follows from $\sup(M) + L < \inf N$ that

$$\ell_\gamma \mathcal{H}_i \zeta = \ell_{\zeta \circ \gamma} \mathcal{H} \in M.$$

Thus, ζ permutes the homotopy classes in Δ. A final fact proves the ζ is reducible:

Promise 9. If a homeomorphism ζ permutes a finte set Δ of non-parallel, pairwise disjoint simple closed curves then these homotopy classes can simultaneously realized by simple closed curves which are permuted by a homeomorphism homotopic to ζ.

7.1.2 Second Proof: Compactifying Teichmüller Space

...

Exercise 7.1.13. Show that for any three angles α_1, α_2, and α_3 with $0 < \alpha_i < \pi$ and $\alpha_1 + \alpha_2 + \alpha_3 < \pi$ there exists a triangle in the hyperbolic plane with these interior angles. Moreover, this triangle is unique up to isometry.

Exercise 7.1.14. Let Σ be a surface, let P be a polygon, and let $f : P \to \Sigma$ be a map that realizes Σ by identifying the edges of P in pairs. Prove that the universal cover $\tilde{\Sigma}$ is naturally tiled with copies of P that intersect only along their boundaries.

Exercise 7.1.15. Let P be an n-polygon that has positive interior angles assigned to its corners. Prove that this polygon can be drawn in the hyperbolic plane provided the angles add up to strictly less then $(n-2)\pi$.

Exercise 7.1.16. Let the Poincaré disc

$$\mathbb{D}^2 := \left\{ (x,y) \mid x^2 + y^2 < 1 \right\}$$

be the unit disc endowed with the Riemannian metric

$$\mathrm{d}\, s^2 := 4 \frac{\mathrm{d}\, x^2 + \mathrm{d}\, y^2}{\left(1 - x^2 + y^2\right)^2}.$$

Show that $\mathbb{D}^2$ and $\mathbb{H}^2$ are isometric.

Exercise 7.1.17. Show that geodesics in the Poincaré disc look like circles perpendicular to the unit disc.

Exercise 7.1.18. Let P and Q be two points in the hyperbolic plane. Give ruler and compass constructions for the geodesic through P and Q in the upper half plane model and in the Poincare disc model.

7.2 Poincaré's Theorem

Theorem 7.2.1 (Poincaré). *Let D be a polygon diagram drawn in the hyperbolic plane such that the lengths of its edges and the interior angles at its corners satisfy the following two conditions:*

1. *Equivalent edges have the same length.*

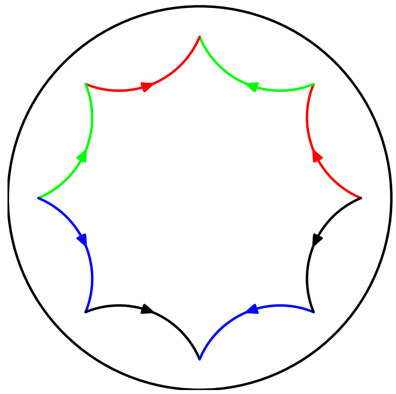

Figure 7.1: The genus 2 standard polygon diagram can be drawn in the hyperbolic plane so that all edges have equal length and all interior angles are $\frac{\pi}{4}$. This gives rise to a tiling of $\mathbb{H}^2$ by regular 8-gons. The group of coloring preserving symmetries of this tiling is the fundamental group of the 2-torus.

2. *The angles of all corners in an equivalence class sum up to* 2π.

Then there is a tiling of the hyperbolic plane by isometric copies of D *such that each at edge of two copies of* D *meet along a pair of equivalent edges. Moreover, the coloring preserving symmetries of this tiling are a group of hyperbolic isometries of* $\mathbb{H}^2$ *isomorphic to the fundamental group of the surface defined by* D.

Remark 7.2.2. The conditions say that the polygon diagram D can tile the hyperbolic plane locally around edges and vertices. Thus, they are clearly necessary conditions for the existence of a global tiling. The theorem says, if a tile tiles locally it tiles globally.

Remark 7.2.3. Although the theorem is stated for polygon diagrams in the hyperbolic plane, it also holds for polygons in the Euclidean plane and even in the sphere. The proof carries over to these cases unchanged.

Proof of (7.2.1). Let Σ be the surface defined by D and let $\tilde{\Sigma}$ its universal cover. By (7.1.14), $\tilde{\Sigma}$ is tiled by topological copies of D in the way the theorem requires. Our strategy will be to put a hyperbolic structure on $\tilde{\Sigma}$ and prove that it is isometric to $\mathbb{H}^2$.

Recall that the tiles are defined as lifts $D \to \tilde{\Sigma}$ that take D homeomorphically to its image. Moreover, these lifts make the following diagram commute:

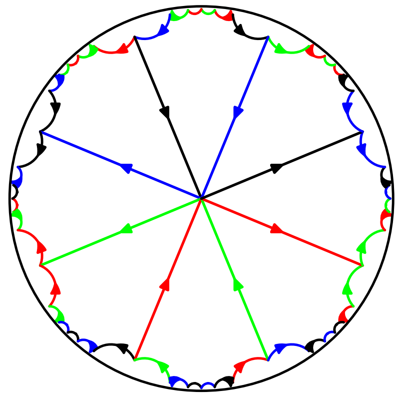

Figure 7.2: The neighborhood of a vertex in a tiling by regular genus 2 standard polygons.

$$\begin{array}{ccc} & & \tilde{\Sigma} \\ & \nearrow & \downarrow \\ \mathbb{H}^2 \supset D & \longrightarrow & \Sigma \end{array}$$

Now, we define a hyperbolic structure by three types of charts; see figure 7.4.

type I: In the interior $D^{(2)}$ of a tile $\tilde{D}$, we use the fact that we have a continuous inverse $\varphi_{\tilde{D}} : D^{(2)} \to D \subset \mathbb{H}^2$, which we declare to be a chart map. We say that the interior of D is the *defining piece* for the type I charts.

type II: The second type of charts will give us neighborhoods of edges. Let a be an edge-color and let e and e_- be the two open edges of color a. Fix two disjoint neighborhood U_{e_+} and U_{e_-} of e_+ and e_- in D that do not contain the end points of these edges. We form a subset $U_a \subset \mathbb{H}^2$ by gluing together two hyperbolic translates of U_{e_+} and U_{e_-} – recall that D is drawn in $\mathbb{H}^2$ and that e_+ and e_- have the same length. Note that each edge $e_\bullet$ in $\tilde{\Sigma}$ of color a has a neighborhood $V_{e_\bullet}$ that is homeomorphically identified with U via the map $U_{e_+} \cup U_{e_-} \to U_a$. The induced homeomorphisms

$$\varphi_{e_\bullet} : V_{e_\bullet} \to U_a$$

are our second collection of coordinate charts. We say that the open sets U_{e_+} and U_{e_-} are the defining pieces for the type II charts.

type III: To define the third type of charts, fix a positive real number R such that the hyperbolic discs of radius R around all the corners of $D \subset \mathbb{H}^2$ are disjoint. Now

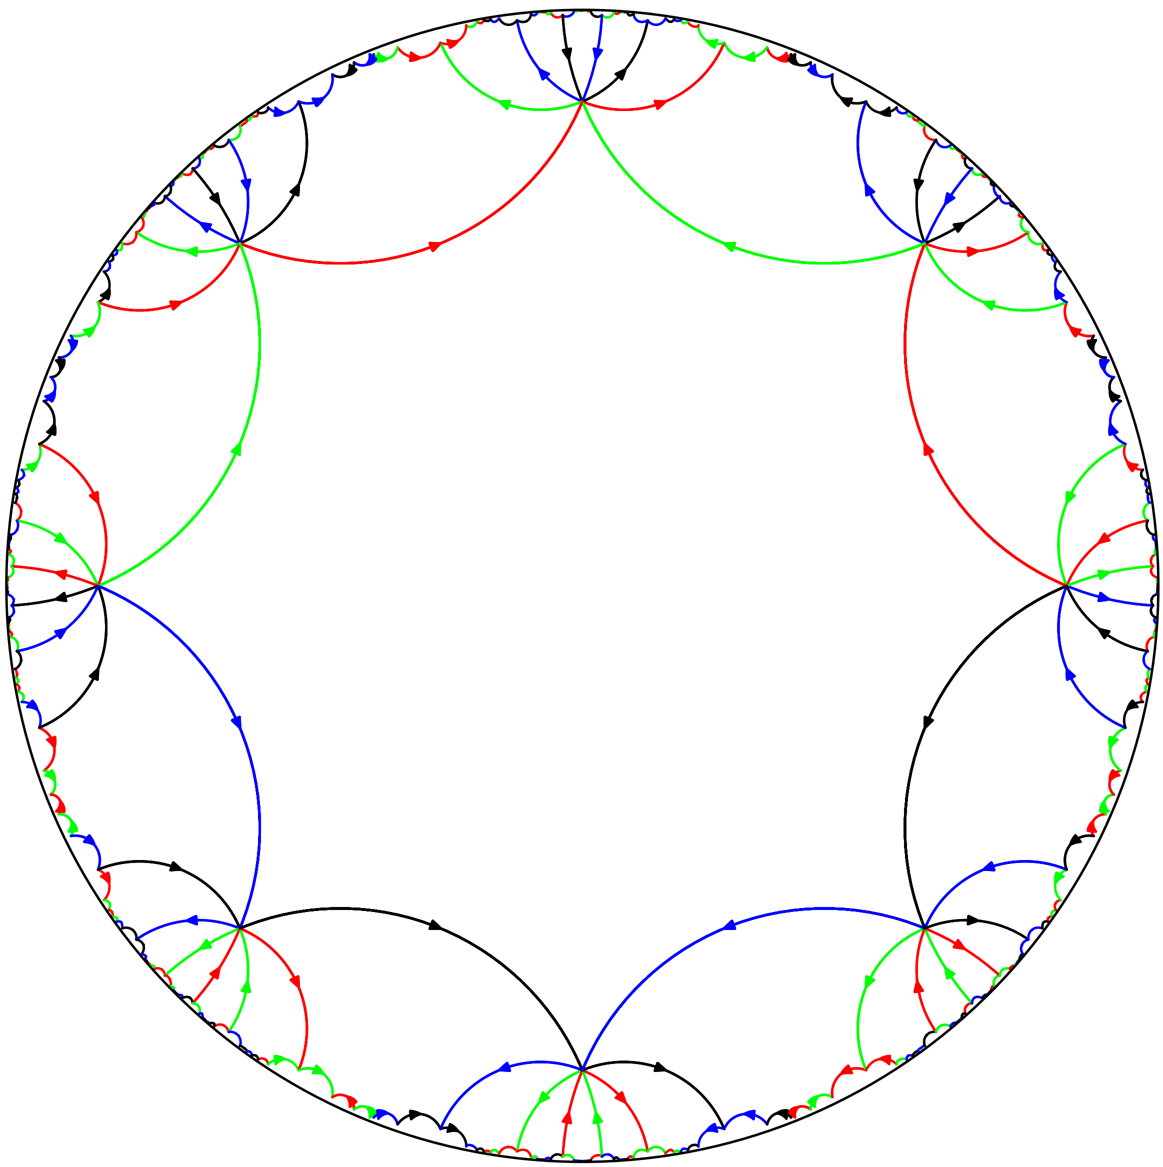

Figure 7.3: An impression of the tiling.

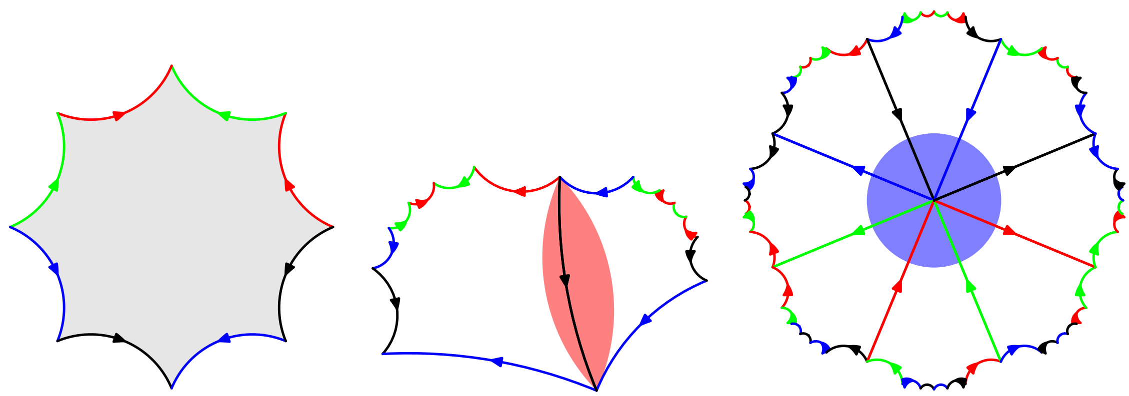

Figure 7.4: Chart types.

fix an equivalence class (vertex-colors) a of corners in D. Translate the open R-neighborhood of these corners in the hyperbolic plane so that they form a local picture $U \subset \mathbb{H}^2$ for neighborhoods V_w of vertices w of color a in $\tilde{\Sigma}$. The canonical homeomophisms

$$\varphi_w : V_w \to U_a$$

will be our chart. The R-neighborhood of corners in D are the defining pieces for the type III charts.

The domains of these charts form an open cover of $\tilde{\Sigma}$. This follows since the defining pieces form an open cover of D Note that charts of type II and III are assebled by moving pieces of D via hyperbolic isometries. It follows that coordinate changes are hyperbolic isometries. Thus, we have defined a hyperbolic structure on $\tilde{\Sigma}$. Deck transformations of $\tilde{\Sigma}$ move the tiles and respect the gluing pattern. Thus, by construction of the hyperbolic structure, deck transformations become isometries with respect to this structure. Equivalently, we could say that we have, in fact, constructed a hyperbolic structure on Σ.

The hyperbolic structure on $\tilde{\Sigma}$ is complete. This follows since the cover of D by the defining pieces (figure 7.5) has a positive Lebesgue number. As a consequence, we infer that the simply connected cover $\tilde{\Sigma}$ is isometric to $\mathbb{H}^2$. Thus, the tiling of $\tilde{\Sigma}$ *is* the tiling of $\mathbb{H}^2$ that we were looking for. **q.e.d.**

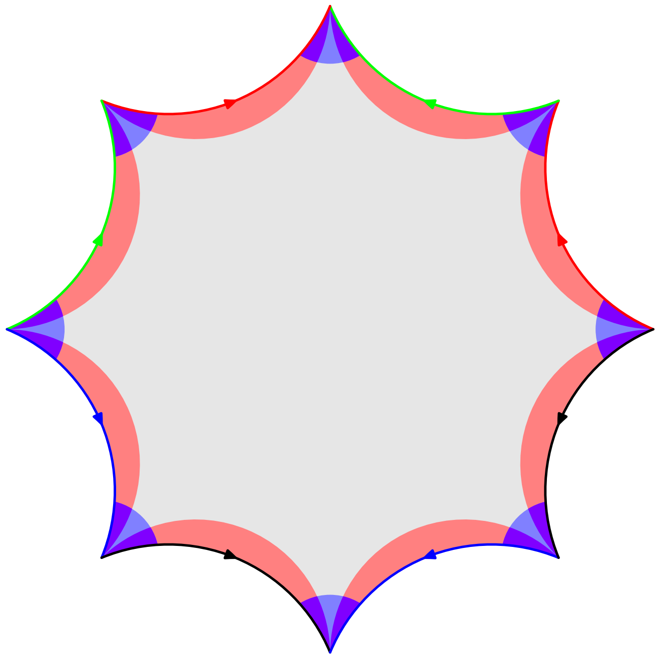

Figure 7.5: The cover by defining pieces.

7.3 The Dehn-Nielsen Theorem

Recall that any map

$$f : X \to Y$$

induces a homomorphism

$$f_* : \pi_1(X, x) \to \pi_1(Y, f(x)) .$$

This homomorphism depends only on the homotopy class of f. If f is a homotopy equivalence, this map is an isomorphism.

A path p in X induces a "change of basepoint" isomorphism:

$$\begin{array}{rcl} p_* : \pi_1(X, p(1)) & \to & \pi_1(X, p(0)) \\ [\gamma] & \mapsto & [p \twoheadrightarrow \gamma \twoheadrightarrow p^{\mathrm{rev}}] . \end{array}$$

Fix a basepoint $\underline{\mathrm{P}}$ in the closed oriented surface Σ. For any self-homotopy equivalence $f : \Sigma \to \Sigma$, let p be any path from $\underline{\mathrm{P}}$ to $f(\underline{\mathrm{P}})$. The isomorphism

$$p_* \circ f_* : \pi_1(\Sigma, \underline{\mathrm{P}}) \to \pi_1(\Sigma, \underline{\mathrm{P}})$$

depends on p, but the induced outer automorphism

$$\nu(f) := [p_* \circ f_*] \in \mathrm{Out}(\pi_1(\Sigma, \underline{\mathrm{P}}))$$

does not. Thus, the map $\nu : f \mapsto \nu(f)$ induces a well defined map

$$\nu : M(\Sigma) \to \mathrm{Out}(\pi_1(\Sigma, \underline{\mathrm{P}}))\,.$$

Theorem 7.3.1 (Dehn-Nielsen). *Let Σ be a closed oriented surface with negative Euler characteristic. Then, the map $\nu : M(\Sigma) \to \mathrm{Out}(\pi_1(\Sigma, \underline{P}))$ is an isomorphism of groups. Moreover, every mapping class is realized by a self-homeomorphism of Σ.*

There is a slightly different phrasing of this result in terms of the group of deck transformations. For any homotopy equivalence $f : \Sigma \to \Sigma$ we can choose a lift $\tilde{f} : \tilde{\Sigma} \to \tilde{\Sigma}$. This lift induces an isomorphism $\tilde{f}_* : \mathrm{Cov}\left(\tilde{\Sigma}/\Sigma\right) \to \mathrm{Cov}\left(\tilde{\Sigma}/\Sigma\right)$ defined by the requirement that

$$\begin{array}{ccc} \tilde{\Sigma} & \xrightarrow{\tilde{f}} & \tilde{\Sigma} \\ {\scriptstyle\tau}\uparrow & & \uparrow{\scriptstyle\tilde{f}_*(\tau)} \\ \tilde{\Sigma} & \xrightarrow{\tilde{f}} & \tilde{\Sigma}_2 \end{array}$$

commutes for every deck transformation $\tau \in \mathrm{Cov}\left(\tilde{\Sigma}/\Sigma\right)$.

Different choices of the lift $\tilde{f}$ yield different isomorphisms, however, these isomorphisms differ only by an inner automorphism. Thus, $\left[\tilde{f}_*\right] \in \mathrm{Out}\left(\mathrm{Cov}\left(\tilde{\Sigma}/\Sigma\right)\right)$ only depends on f and we have a well defined map

$$\begin{array}{rcl} \tilde{\nu} : M(\Sigma) & \to & \mathrm{Out}\left(\mathrm{Cov}\left(\tilde{\Sigma}/\Sigma\right)\right) \\ f & \mapsto & \left[\tilde{f}_*\right]. \end{array}$$

Exercise 7.3.2. Recall that every choice of a point $\tilde{\underline{\mathrm{P}}}$ in the fiber above $\underline{\mathrm{P}}$ defines an isomorphism

$$\pi_1(\Sigma, \underline{\mathrm{P}}) \to \mathrm{Cov}\left(\tilde{\Sigma}/\Sigma\right).$$

Show that, independent of the choice of $\tilde{\underline{\mathrm{P}}}$, we obtain a well defined isomorphism

$$\Phi : \mathrm{Out}(\pi_1(\Sigma, \underline{\mathrm{P}})) \to \mathrm{Out}\left(\mathrm{Cov}\left(\tilde{\Sigma}/\Sigma\right)\right)$$

and that this isomorphism makes the diagram

$$\begin{array}{ccc} & & \mathrm{Out}(\pi_1(\Sigma, \underline{P})) \\ & \overset{\nu}{\nearrow} & \\ M(\Sigma) & & \downarrow \Phi \\ & \underset{\tilde{\nu}}{\searrow} & \\ & & \mathrm{Out}\Big(\mathrm{Cov}\Big(\tilde{\Sigma}/\Sigma\Big)\Big) \end{array}$$

commute.

Now (7.3.1) implies:

Corollary 7.3.3. *The map* $\tilde{\nu} : M(\Sigma) \to \mathrm{Out}\Big(\mathrm{Cov}\Big(\tilde{\Sigma}/\Sigma\Big)\Big)$ *is an isomorphism of groups. Moreover, every mapping class is induced by a self-homeomorphism of* $\tilde{\Sigma}$. **q.e.d.**

Corollary 7.3.4. *Let* Σ *be a closed oriented surface with negative Euler characteristic. Let* $\tilde{\Sigma}$ *be its universal cover and let* $\tilde{\zeta} : \tilde{\Sigma} \to \tilde{\Sigma}$ *be a homeomorphism that commutes with all deck transformations, i.e., the following diagram commutes for all deck transformations* $\tau : \tilde{\Sigma} \to \tilde{\Sigma}$*:*

$$\begin{array}{ccc} \tilde{\Sigma} & \xrightarrow{\tilde{\zeta}} & \tilde{\Sigma} \\ \tau \uparrow & & \uparrow \tau \\ \tilde{\Sigma} & \xrightarrow{\tilde{\zeta}} & \tilde{\Sigma} \end{array}$$

Then $\tilde{\zeta}$ *induces a homeomorphism* $\zeta : \Sigma \to \Sigma$*, which is homotopic to the identity.*

Proof. Observe that $\tilde{\zeta}_*$ is the identity. Thus, ζ is in the kernel of $\tilde{\nu}$ and therefore homotopic to the identity. **q.e.d.**

Exercise 7.3.5. Let Σ be a closed oriented surface of Euler characteristic χ. Show that $\pi_1(\Sigma)^{\mathrm{ab}} = \mathbb{Z}^{2-\chi}$.

Corollary 7.3.6. *Let* Σ_1 *and* Σ_2 *be two closed oriented surfaces of negative Euler characteristic. Let* $\tilde{\Sigma}_i$ *be the universal cover of* Σ_i*, and let* $\phi : \mathrm{Cov}\Big(\tilde{\Sigma}_1/\Sigma_1\Big) \to \mathrm{Cov}\Big(\tilde{\Sigma}_1/\Sigma_2\Big)$ *be an isomorphism. Then there exists a homeomorphism* $\tilde{\zeta} : \tilde{\Sigma}_1 \to \tilde{\Sigma}_2$ *of the universal covers that makes the diagram*

$$
\begin{array}{ccc}
\tilde{\Sigma}_1 & \xrightarrow{\tilde{\zeta}} & \tilde{\Sigma}_2 \\
{\scriptstyle \tau}\uparrow & & \uparrow{\scriptstyle \phi(\tau)} \\
\tilde{\Sigma}_1 & \xrightarrow{\tilde{\zeta}} & \tilde{\Sigma}_2
\end{array}
$$

commute for each deck transformation $\tau \in \mathrm{Cov}\left(\tilde{\Sigma}_1/\Sigma_1\right)$.

Proof. Note that Σ_1 and Σ_2 are two closed oriented surfaces with isomorphic fundamental groups. By (7.3.5), the two surfaces have the same Euler characteristic and are therefore homeomorphic by the classification of closed oriented surfaces. Thus, there is a homeomorphism

$$\xi : \Sigma_2 \longrightarrow \Sigma_1,$$

which has a lift

$$\tilde{\xi} : \tilde{\Sigma}_2 \longrightarrow \tilde{\Sigma}_1.$$

By (7.3.3), the isomorphism $\tilde{\xi}_* \circ \phi : \mathrm{Cov}[\Sigma_1]\tilde{\Sigma}_1 \to \mathrm{Cov}[\Sigma_1]\tilde{\Sigma}_1$ is induced by a homeomorphism $\tilde{\zeta}_1 : \tilde{\Sigma}_1 \to \tilde{\Sigma}_1$. We have the commutative diagram:

$$
\begin{array}{ccccc}
\tilde{\Sigma}_1 & \xrightarrow{\tilde{\zeta}_1} & \tilde{\Sigma}_1 & \xleftarrow{\tilde{\xi}} & \tilde{\Sigma}_2 \\
{\scriptstyle \tau}\uparrow & & \uparrow{\scriptstyle \tilde{\xi}_*(\phi(\tau))} & & \uparrow{\scriptstyle \phi(\tau)} \\
\tilde{\Sigma}_1 & \xrightarrow[\tilde{\zeta}_1]{} & \tilde{\Sigma}_1 & \xleftarrow[\tilde{\xi}]{} & \tilde{\Sigma}_2
\end{array}
$$

It follows that $\zeta := \xi^{-1} \circ \zeta_1$ satisfies our needs. **q.e.d.**

The remainder of this section is devoted to the proof of the Dehn-Nielsen theorem (7.3.1). We have to check three statements:

1. The map ν is a group homomorphism.
2. This homomorphism is injective.
3. Every outer automorphism of $\pi_1(\Sigma, \underline{P})$ is induced by a homeomorphism of Σ. This statement implies that ν is onto and that every mapping class is realized by a homeomorphism.

The first two statements are proved in the same way as for the torus. However, since we phrased the statement here for self-homotopy equivalences rather than self-homeomorphisms, let us go through the argument again.

First, let us verify that ν is a homomorphism of groups. So let f and h be self-homotopy equivalences of the surface Σ, and let p and q be paths from the basepoint $\underline{P}$ to $f(\underline{P})$ and $h(\underline{P})$, respectively. Then

$$\begin{aligned} \nu(f)\,\nu(h) &= [p_* \circ f_*]\,[q_* \circ h_*] \\ &= \big[(p \rightarrowtail f \circ q)_* (f \circ h)_*\big] \\ &= \nu(f \circ h)\,. \end{aligned}$$

7.3.1 Injectivity

Lemma 7.3.7. *Every closed oriented surface with negative Euler characteristic is aspherical.*

Proof. The universal cover of such a surface is the hyperbolic plane and therefore contractible. **q.e.d.**

Proposition 7.3.8. *Let Σ be a closed oriented surface with negative Euler characteristic, let $f : \Sigma \to \Sigma$ be a homotopy equivalence, and let p be any path from $\underline{P}$ to $f(\underline{P})$. If $p_* \circ f_*$ is an inner automorphism of $\pi_1(\Sigma, \underline{P})$, then f is homotopic to the identity.*

Proof. Let γ be a loop such that the inner automorphism induced by γ equals $p_* \circ f_*$. That is, for any loop γ', we have

$$[\gamma \rightarrowtail \gamma' \rightarrowtail \gamma^{\mathrm{rev}}] = [p \rightarrowtail f \circ \gamma' \rightarrowtail p^{\mathrm{rev}}]\,.$$

Conjugating by γ^{rev}, we obtain the equation

$$[\gamma'] = [\gamma^{\mathrm{rev}} \rightarrowtail p \rightarrowtail f \circ \gamma' \rightarrowtail p^{\mathrm{rev}} \rightarrowtail \gamma] = [(\gamma^{\mathrm{rev}} \rightarrowtail p) \rightarrowtail f \circ \gamma' \rightarrowtail (\gamma^{\mathrm{rev}} \rightarrowtail p)^{\mathrm{rev}}]\,. \tag{7.1}$$

Replacing p by $\gamma^{\mathrm{rev}} \rightarrowtail p$, we assume w.l.o.g. that $p_* \circ f_*$ is actually the identity.

We have to construct a homotopy

$$\Phi : \Sigma \times \mathbb{I} \to \Sigma$$

from $f = \Phi(-,0)$ to the identity $\mathrm{id}_\Sigma = \Phi(-,1)$. To this end, let D be a genus g standard polygon diagram for Σ, and let $\pi : D \to \Sigma$ be the projection that realizes the identifications by which D describes Σ. We will construct a map

$$\Psi : D \times \mathbb{I} \to \Sigma$$

that will induce the desired homotopy $\Phi : \Sigma \times \mathbb{I} \to \Sigma$.

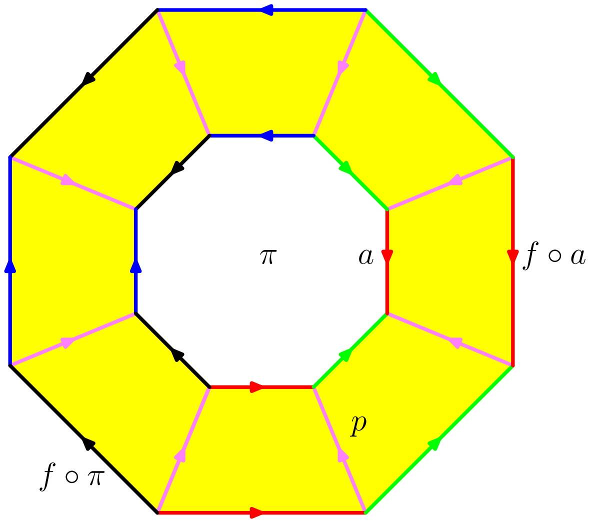

Figure 7.6: The drum (front view).

Note that $\Sigma \times \mathbb{I}$ is obtained from the drum by making identifications along the yellow boundary annulus. These identifications of faces are induced by the identifications of edges in the polygon diagram D. We will define Ψ in such a way that it is compatible with those identifications. Thus, Ψ will descend to $\Sigma \times \mathbb{I}$.

We define $\Psi : D \times \mathbb{I} \to \Sigma$ as indicated in figure 7.6:

- In the back face $D \times \{1\}$ of the *drum* $D \times \mathbb{I}$, we define Ψ to be the composition

$$D \times \{0\} \to D \xrightarrow{\pi} \Sigma.$$

- In the front face $D \times \{0\}$, we define Ψ to be the composition

$$D \times \{0\} \to D \xrightarrow{\pi} \Sigma \xrightarrow{f} \Sigma.$$

- The pink edges are of the form $x \times \mathbb{I}$. We map them all to the path p. Formally, on these edges, we define Ψ to be the composition

$$x \times \mathbb{I} \to \mathbb{I} \xrightarrow{p} \Sigma.$$

- Now, we fill in the yellow squares along the boundary of the drum. Note that we already have defined Ψ on the boundary circles of theses square 2-cells. The boundary of a square whose edge-color is a is mapped to the loop

$$f \circ a \twoheadrightarrow p \twoheadrightarrow a^{\text{rev}} \twoheadrightarrow p^{\text{rev}}.$$

By equation 7.1, this loop is null-homotopic in Σ. Thus, we can extend Ψ to the two yellow squares with this edge-color. Moreover, we can choose the extension so that Ψ is compatible with the face-identifications on the drum.

- As of now, we have defined Ψ on the whole boundary sphere of the drum. By asphericity of Σ, we can extend Ψ to all of $D \times \mathbb{I}$.

Since Ψ is compatible with the face-identifications on $D \times \mathbb{I}$, we can define Φ by the diagram:

$$\begin{array}{ccc} D \times \mathbb{I} & \xrightarrow{\Psi} & \Sigma \\ \downarrow & \nearrow_{\Phi} & \\ \Sigma \times \mathbb{I} & & \end{array}$$

This is the desired homotopy of f and id_Σ. **q.e.d.**

Corollary 7.3.9. *The homomorphism ν is injective.* **q.e.d.**

7.3.2 Surjectivity and the "Moreover,..." Clause

...

Exercise 7.3.10. Prove: In a closed surface with a fixed hyperbolic structure, every closed curve is freely homotopic to a unique closed geodesic – here, a closed geodesic need not be simple.

Definition. Let G be a group with a fixed generating system Σ. The Cayley graph $\Gamma_\Sigma(G)$ is a directed graph whose vertices are the elements of G. For each vertex g and each generator $x \in \Sigma$, there is an edge from g to gx. We ignore the orientation of these edges and define a metric on the vertex set by declaring all edges to have length 1: The metric

$$d_\Sigma : G \times G \to \mathbb{R}$$

is then given by shortest paths – note that $\Gamma(G)$ is connected since Σ generated G.

Exercise 7.3.11. Let G and H be groups generated by the *finite* generating sets Σ and Ξ, respectively. Let $\varphi : G \to H$ be a group homomorphism. Show that there is a constant C such that for all $g, h \in G$,

$$d_\Xi(\varphi(g), \varphi(h)) \leq C d_\Sigma(g, h).$$

Definition. Two metric space X and Y are called quasi-isometric if there exist two non-negative constants K and C and a function

$$\varphi : X \to Y$$

such that:

1. For all $x, y \in X$,
$$\frac{1}{C} d_X(x,y) - K \leq d_Y(\varphi(x), \varphi(y)) \leq C d_X(x,y) + K.$$

2. Every point in Y is within distance K of the image of φ.

Exercise 7.3.12. Show that quasi-isometry is an equivalence relation on the class of metric spaces.

Exercise 7.3.13. Let Σ be a closed oriented surface with negative Euler characteristic. Show that the Cayley graph of $\pi_1(\Sigma)$ with respect to any finite generating set is quasi-isometric to $\mathbb{H}^2$.

7.4 Calculation of Teichmüller Space

Let Σ be a closed oriented surface of negative Euler characteristic. Teichmüller space $\mathcal{T}_\Sigma$ is the space of all hyperbolic structures on Σ up to equivalence: Let

- $\mathrm{Homeo}(\Sigma)$ be the group of self-homeomorphisms on the torus Σ, and let

- $\mathrm{Homeo}_1(\Sigma)$ be the normal subgroup of those homeomorphisms that are homotopic to the identity. The factor group

- $M(\Sigma) := \mathrm{Homeo}(\Sigma) / \mathrm{Homeo}_1(\Sigma)$ is called the mapping class group of Σ.

  Let

- $\mathrm{Isom}(\mathbb{H}^2)$ be the isometry group of the hyperbolic plane. Note that $\mathrm{Isom}(\mathbb{H}^2)$ acts from the left on

- $\mathcal{H}(\Sigma)$, the set of hyperbolic structures on Σ. The action is given by modifying all the charts, appending the isometry $\lambda \in \mathrm{Isom}(\mathbb{H}^2)$.

  Note that $\mathrm{Homeo}(\Sigma)$ acts on $\mathcal{H}(\Sigma)$ from the right as follows: For a homeomorphism $\zeta : \Sigma \to \Sigma$, a given hyperbolic structure $\mathcal{H}$ on Σ and a chart $\varphi : U \to \mathbb{H}^2$ for this structure, define a corresponding chart

$$\varphi \circ \zeta : \zeta^{-1}(U) \to \mathbb{H}^2.$$

All these charts form a new atlas for Σ and define a different hyperbolic structure $\mathcal{H}\zeta$. Note that

$$\zeta : (\Sigma, \mathcal{H}) \to (\Sigma, \mathcal{H}\zeta)$$

is an equivalence of hyperbolic structures. This action induces an action of $\mathrm{Homeo}(\Sigma)$ on $\mathcal{H}(\Sigma)$.

The double quotient

- $\mathcal{M}_\Sigma := \mathrm{Isom}(\mathbb{H}^2) \backslash \mathcal{H}(\Sigma) / \mathrm{Homeo}(\Sigma)$ is called the moduli space of Σ and the quotient

- $\mathcal{T}_\Sigma := \mathrm{Isom}(\mathbb{H}^2) \backslash \mathcal{H}(\Sigma) / \mathrm{Homeo}_1(\Sigma)$ is called the Teichmüller space of Σ. Note that there is a natural action of $M(\Sigma)$ on $\mathcal{T}_\Sigma$ such that

$$\mathcal{M}_\Sigma = \mathcal{T}_\Sigma / M(\Sigma) .$$

Furthermore, the quotient

- $\mathcal{D}_\Sigma := \mathrm{Isom}(\mathbb{H}^2) \backslash \mathrm{Hom}^{\mathrm{i,d}}(\pi_1(\Sigma, \underline{\mathrm{P}}), \mathrm{Isom}(\mathbb{H}^2))$ is called the deformation space of Σ, where $\mathrm{Hom}^{\mathrm{i,d}}(\pi_1(\Sigma, \underline{\mathrm{P}}), \mathrm{Isom}(\mathbb{H}^2)) := \{\varphi : \pi_1(\Sigma, \underline{\mathrm{P}}) \to \mathrm{Isom}(\mathbb{H}^2) \mid \varphi \text{ is injective and has discrete image.}\}$.

Theorem 7.4.1. *The map*

$$\begin{array}{rcl} \Psi : \mathcal{T}_\Sigma & \to & \mathcal{D}_\Sigma \\ [\mathcal{E}] & \mapsto & [\eta^{\delta}_{\mathcal{E}}] \end{array}$$

is a bijection.

Proof of Injectitivity. Suppose we have two hyperbolic structures $\mathcal{H}_1$ and $\mathcal{H}_2$ on Σ such that

$$\left[\eta^{\delta_1}_{\mathcal{H}_1}\right] = \left[\eta^{\delta_2}_{\mathcal{H}_2}\right] .$$

Then there is an isometry $\lambda : \mathbb{H}^2 \to \mathbb{H}^2$ such that, for each deck transformation τ, the following diagram commutes:

$$\begin{array}{ccccccc} \tilde{\Sigma} & \xrightarrow{\delta_1} & \mathbb{H}^2 & \xrightarrow{\lambda} & \mathbb{H}^2 & \xleftarrow{\delta_2} & \tilde{\Sigma} \\ {\scriptstyle\tau}\uparrow & & \uparrow{\scriptstyle \eta^{\delta_1}_{\mathcal{H}_1}(\gamma)} & & \uparrow{\scriptstyle \eta^{\delta_1}_{\mathcal{H}_1}(\gamma)} & & \uparrow{\scriptstyle\tau} \\ \tilde{\Sigma} & \xrightarrow[\delta_1]{} & \mathbb{H}^2 & \xrightarrow[\lambda]{} & \mathbb{H}^2 & \xleftarrow[\delta_2]{} & \tilde{\Sigma} \end{array}$$

By (7.3.4), it follows that $\delta_2^{-1} \circ \lambda \circ \delta_1$ induces a homeomorphism $\zeta : \Sigma \to \Sigma$ that is homotopic to the identity. It is easy to check that all these diagrams add up to:

$$\lambda \mathcal{H}_1 \zeta = \mathcal{H}_2$$

Thus, $[\mathcal{H}_1] = [\mathcal{H}_2]$. **q.e.d.**

Exercise 7.4.2. Show that the fundamental group of any non-compact surface is free.

Hint 7.4.3. First, consider the case of a punctured surface Σ, i.e., a closed surface with some discrete set of points removed. Show that there is a graph inside the surface onto which Σ deformation retracts. Then, the fundamental group of the surface is the fundamental group of the graph and hence free.

For the general case, consider a triangulation of the surface Σ. Show that there is a graph Γ inside the 1-skeleton of the triangulation whose complementary components are all infinite, simply-connected, and one-ended: A space X is called one-ended if every compact subset is contained in another compact subset that has a connected complement. Show that Σ deformation retracts onto Γ.

Exercise 7.4.4. Let Σ be a closed, non-orientable surface. Prove that $\pi_1(\Sigma)^{\mathrm{ab}}$ contains an element of order 2.

Corollary 7.4.5. *Two non-homeomorphic closed surfaces have non-isomorphic fundamental groups.*

Proof. This follows from the classification of closed surfaces, (7.3.5), and (7.4.4). Indeed, already the abelianizations of their fundamental groups differ. **q.e.d.**

Proof of surjectivity. Let

$$\eta : \pi_1(\Sigma) = \mathrm{Cov}\Big(\tilde{\Sigma}/\Sigma\Big) \to \mathrm{Isom}\big(\mathbb{H}^2\big)$$

be an injective homomorphism with discrete image $G := \mathrm{im}(\eta) \leq \mathrm{Isom}(\mathbb{H}^2)$. Note that $G\backslash\mathbb{H}^2$ is a surface with fundamental group G which is isomorphic to $\pi_1(\Sigma)$. Since G cannot be free, $G\backslash\mathbb{H}^2$ is a closed surface. By (7.4.5), the surface $G\backslash\mathbb{H}^2$ is homeomorphic to Σ. Note that $G\backslash\mathbb{H}^2$ comes with a canonical hyperbolic structure. The idea is to pull this one over to Σ.

By (7.3.6), there is a homeomorphism

$$\tilde{\zeta} : \tilde{\Sigma} \to \mathbb{H}^2$$

such that

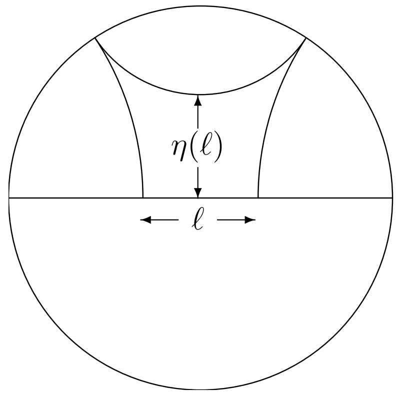

Figure 7.7: The width function.

$$\begin{array}{ccc} \tilde{\Sigma} & \xrightarrow{\tilde{\zeta}} & \mathbb{H}^2 \\ \uparrow \tau & & \uparrow \eta(\tau) \\ \tilde{\Sigma} & \xrightarrow[\tilde{\zeta}]{} & \mathbb{H}^2 \end{array}$$

commutes for each deck transformation τ. Thus, we use $\tilde{\zeta}$ to define a hyperbolic structure on $\tilde{\Sigma}$, which visibly descends to a hyperbolic structure $\mathcal{H}$ on Σ. The homeomorphism $\tilde{\zeta}$ is a developing map for $\mathcal{H}$, and using this developing map, we see that η is the holonomy representation induced by the hyperbolic structure $\mathcal{H}$. **q.e.d.**

7.5 Short Geodesics

We will prove that short simple closed geodesics are disjoint unless they coincide. We follow a proof given by John H. Hubbard.

Definition 7.5.1. The width function $\eta : \mathbb{R}^+ \to \mathbb{R}^+$ is defined as follows. For any positive real number ℓ, draw a line segment of length ℓ on you favorite geodesic in $\mathbb{H}^2$. At its endpoints, draw the perpendiculars and extend them into one side of the hyperbolic plane until they hit the boundary. This way, you obtain two points on the boundary, one for each perpendicular. Join these boundary points by a geodesic. The distance from this geodesic to your favorite one is $\eta(\ell)$. (See figure 7.7.)

Observation 7.5.2. *The width function is monotonically decreasing.* **q.e.d.**

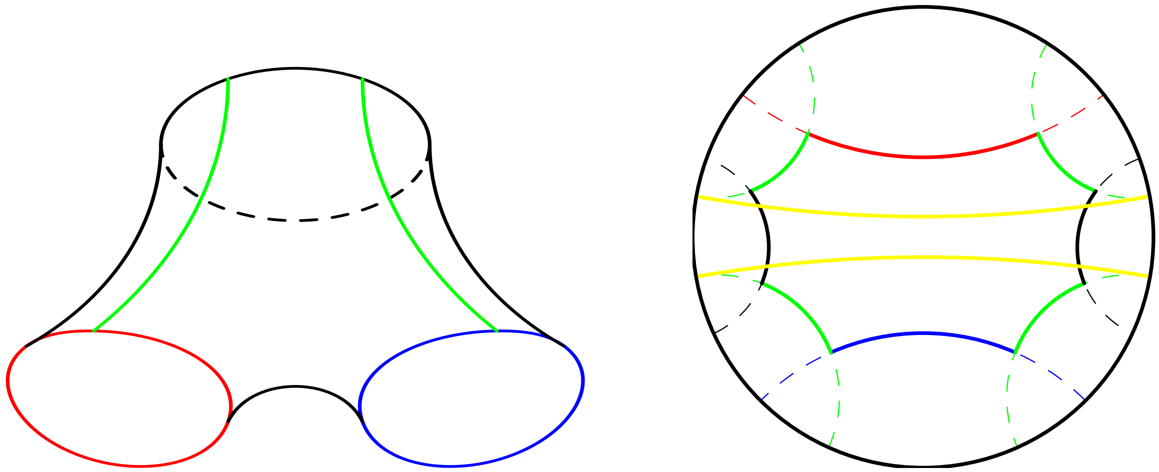

Figure 7.8: A hyperbolic pair of pants with cut lines.

Lemma 7.5.3. *In a hyperbolic pair of pants with totally geodesic boundary circles γ_1, γ_2, and γ_3 of lengths ℓ_1, ℓ_2, and ℓ_3, respectively, you can draw an annulus of width $\eta(\ell_i)$ around γ_i and all three annuli will be pairwise disjoint.*

Proof. Consider the pair of pants in figure 7.8 and let the red boundary circle be γ_1 and the blue circle be γ_2. Fix shortest lines from the third circle to γ_1 and γ_2. Note that these two green arcs will be geodesics, they will be perpendicular to the boundary, and they will be disjoint. Cut along the green geodesic arcs. We obtain a right-angled octagon as shown in the right figure. Note that the green geodesics do not intersect since they have common perpendiculars. Thus, the two yellow geodesics that determine $\eta(\ell_1)$ and $\eta(\ell_2)$ do not intersect. The claim now follows. **q.e.d.**

Theorem 7.5.4. *Let $\{\gamma_1, \gamma_2, \ldots\}$ be a set of pairwise non-homotopic simple closed geodesics in a closed hyperbolic surface with lengths $\ell_1, \ell_2, \ldots$. Then the open $\eta(\ell_i)$-neighborhoods of the loops γ_i are pairwise disjoint.*

Proof. Extend the set of curves to a complete pair of pants decomposition of the surface and apply (7.5.3). **q.e.d.**

Corollary 7.5.5. *Let γ_1 and γ_2 be two non-homotopic, simple, closed geodesics of lengths ℓ_1 and ℓ_2 on a closed hyperbolic surface. Suppose these two loops intersect n of times. Then If $\ell_2 \geq 2nn\eta(\ell_1)$.*

Proof. The two loops intersect transversally. For each intersection point, we find a segment of length $2\eta(\gamma_1)$ on γ_2 centered at this intersection. These segments do not overlap since γ_2 is simple. The lower bound for ℓ_2 now follows. **q.e.d.**

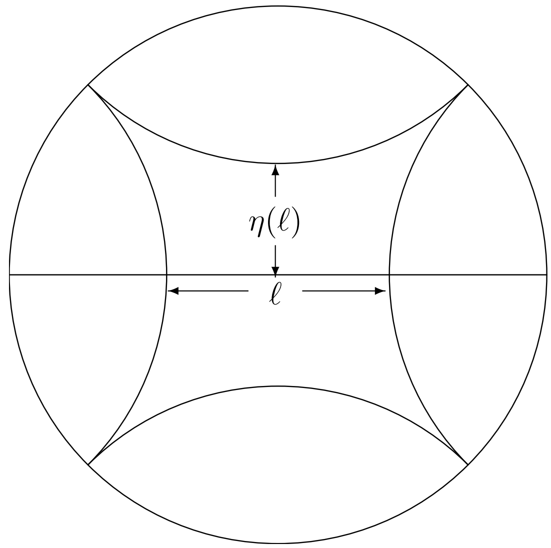

Figure 7.9: Solving $\ell = 2\eta(\ell)$.

Corollary 7.5.6. *Let γ_1 and γ_2 be two simple closed geodesics of lengths ℓ_1 and ℓ_2 on a closed hyperbolic surface. If $\ell_2 < 2\eta(\ell_1)$ then the loops are either disjoint or coincide.* **q.e.d.**

Corollary 7.5.7. *Two simple closed geodesics intersect only finitely many times.* **q.e.d.**

Exercise 7.5.8. Show that the number $\ln(3 + 2\sqrt{2})$ is the unique solution to the equation $\ell = 2\eta(\ell)$. (Hint: look at figure 7.9.)

Corollary 7.5.9. *If two non-homotopic simple closed geodesics in a closed hyperbolic surface have both length $< \ln(3 + 2\sqrt{2})$, then these loops are disjoint.*

Proof. Let γ_1 and γ_2 be two simple closed geodesic curves of length ℓ_1 and ℓ_2. We suppose $\ell_1, \ell_2 < \ln(3 + 2\sqrt{2})$. Thus, we have

$$\ell_2 < \ln(3 + 2\sqrt{2}) < 2\eta(\ell_1)$$

whence the claim follows from (7.5.6). **q.e.d.**

Lemma 7.5.10. *The area of an oriented closed hyperbolic surface of Euler characteristik χ is ...*

Lemma 7.5.11. *Let Σ be an oriented closed hyperbolic surface of diameter D and area A; let ℓ be the length of the shortest non-trivial closed geodesic on Σ. Then*

$$D\ell \leq A.$$

Proof. ... **q.e.d.**

7.6 Isometries of Closed Hyperbolic Surfaces

Lemma 7.6.1. *Let Σ by an oriented, closed, hyperbolic surface. Then, every isometry $\lambda : \Sigma \to \Sigma$ has has a power that fixes a point.*

Proof. By (7.3.13), for any given bound $C > 0$, there are only finitely many simple, closed geodesics in Σ whose length is bounded from above by C. Thus, the set

$$\Gamma_C := \{\gamma \text{ simple, closed geodesic in } \Sigma \mid |\gamma| \leq C\}$$

is finite, whence the set of their intersections

$$\mathcal{P}_C := \bigcup_{\substack{\gamma, \gamma' \in \Gamma_C \\ \gamma \neq \gamma'}} \gamma \cap \gamma'$$

is finite, too. Note that $\mathcal{P}_C$ is non-empty for sufficiently large C.

Since $\mathcal{P}_C$ is defined entirely in metric terms, λ stabilizes this set, i.e., the isometry permutes the points in $\mathcal{P}_C$. Thus the claim follows from $\mathcal{P}_C$ being a finite set: It is a finite union of intersections, each of which is finite by (7.5.7). **q.e.d.**

Theorem 7.6.2. *Every isometry of any oriented, closed, hyperbolic surface has finite order.*

Proof. In view of (7.6.1), we only need to prove the claim for isometries that have a fixed point.

We lift λ to the universal cover and onbtain an isometry $\tilde{\lambda} : \mathbb{H}^2 \to \mathbb{H}^2$. Replacing $\tilde{\lambda}$ with $\tau \circ \tilde{\lambda}$ where τ is a suitably chosen deck transformation, we may assume that $\tilde{\lambda}$ fixes a point in $\mathbb{H}^2$. Since $\tilde{\lambda}$ also preserves a tessalation of $\mathbb{H}^2$, some power of $\tilde{\lambda}$ is the identity. Thus, λ has finite order. **q.e.d.**

7.7 Mumford's Compactness Theorem

Let us actually define a topology on Teichmüller space. We will actually work with deformation space

$$\mathcal{D}_\Sigma = \mathrm{Isom}\big(\mathbb{H}^2\big) \setminus \mathrm{Hom}^{\mathrm{i,d}}\big(\pi_1(\Sigma, \underline{\mathrm{P}}) \,, \mathrm{Isom}\big(\mathbb{H}^2\big)\big) \,.$$

Note that $\mathrm{Isom}(\mathbb{H}^2)$ is a Lie group, in fact, it is a matrix group. Elements are described by a bunch of real numbers and this bunch is finite. There is no doubt about the topology on $\mathrm{Isom}(\mathbb{H}^2)$. Moreover, note that $\pi_1(\Sigma, \underline{\mathrm{P}})$ is finitely generated: The boundary edges of the standard polygon diagram represent loops that generate the fundamental group. Thus, there is a natural topology on

$$\mathrm{Hom}\big(\pi_1(\Sigma, \underline{\mathrm{P}}) \,, \mathrm{Isom}\big(\mathbb{H}^2\big)\big)$$

because we have:

Observation 7.7.1. *Let H be a topological group (e.g., a group of real or complex matrices) and let $G = \langle x_1, \dots, x_r \rangle$ be a finitely generated group. Every homomorphism*

$$\varphi : G \to H$$

is uniquely determined by the tupel

$$(\varphi(x_1) \,, \dots, \varphi(x_r)) \,.$$

This gives a system of H-valued coordinates (in case of a matrix group, actually reals or complex coordinates) for

$$\mathrm{Hom}(G, H) \,.$$

Thus, this space carries a natural topology.

In fact, this topology does not depend on the chosen generating set: Let $\{y_1, \dots, y_s\}$ be another finite generating set for G. Since we can express every y_j as a product of the x_i and their inverses, we obtain a continuous map translating the old coordinates into the coordinates relative to the alternative generating set. Expressing the old generator in terms of the new ones, we see that this map is bi-continuous. Thus, the topologies defined by these coordinate systems coincide.

We endow

$$\mathrm{Hom}^{\mathrm{i,d}}\big(\pi_1(\Sigma, \underline{\mathrm{P}}) \,, \mathrm{Isom}\big(\mathbb{H}^2\big)\big)$$

with the subspace topology. Finally, deformation space

$$\mathcal{D}_\Sigma = \mathrm{Isom}\big(\mathbb{H}^2\big) \setminus \mathrm{Hom}^{\mathrm{i,d}}\big(\pi_1(\Sigma, \underline{\mathrm{P}}) \,, \mathrm{Isom}\big(\mathbb{H}^2\big)\big)$$

carries the quotient topology.

Observation 7.7.2. *For any element* $\gamma \in \pi_1(\Sigma, \underline{P})$*, the map*

$$\ell_\gamma : \mathcal{D}_\Sigma \quad \rightarrow \quad \mathbb{R}^+$$

is continuous. **q.e.d.**

Part III

Three Manifolds

Part IV

CW-Complexes

Part V

Appendices

Appendix A

Notions from Topology

Credo. *All maps are continuous.*

A.1 Paths, Curves, and Homotopies

We denote by

- $\mathbb{I} := [0, 1]$ the unit interval. By
- $\mathbb{B}^m$ we denote the unit ball in $\mathbb{E}^m$, and
- $\mathbb{S}^m$ represents the unit sphere in $\mathbb{E}^{m+1}$.

Definition A.1.1. Let X be a topological space and $S \subset X$ be a subset. To maps $f_0, f_1 : X \to Y$ are homotopic relative to S if there is a homotopy relative to S, i.e., a map

$$\Phi : X \times \mathbb{I} \to Y$$

that satisfies:

1. $f_0^- = \Phi^{-,0}$.
2. $f_1^- = \Phi^{-,1}$.
3. For any $x \in S$, the function $\Phi^{x,-}$ is constant.

For $t \in \mathbb{I}$,

$$f_t := \Phi^{-,t}$$

denotes the homotopy at time t.

Two embeddings (homeomorphisms) $f_0, f_1 : X \hookrightarrow Y$ are isotopic relative to S if there is a homotopy relative to S such that f_t is an embedding (homeomorphism).

To avoid wordiness, we drop the “relative to” phrase if S is empty.

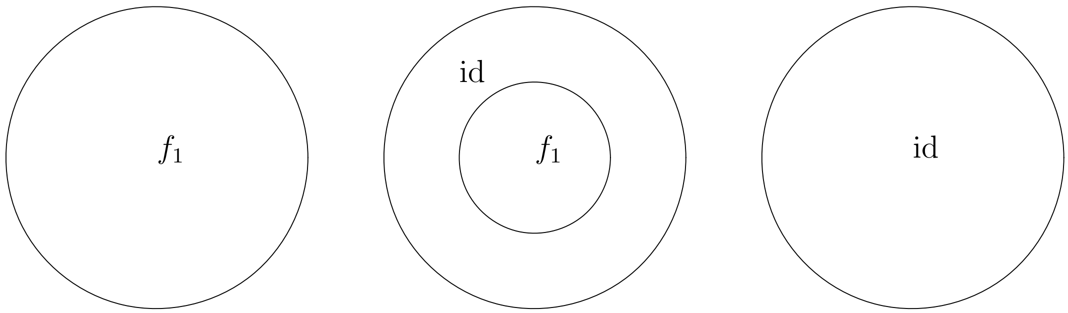

Table A.1: Alexander's trick

Remark A.1.2. If X is locally compact and Hausdorff, then two maps are homotopic if and only if they belong to the same connected component of the function space $\mathrm{Map}(X, Y)$ with the compact open topology.

Definition A.1.3. A path or a curve in X is a map $p : \mathbb{I} \to X$. A path is an arc if it is an embedding. A loop or closed curve is a map $p : \mathbb{S}^1 \to X$. A curve is simple if it is an embedding.

Example A.1.4. A self-homeomorphism $f_1 : \mathbb{B}^{m+1} \to \mathbb{B}^{m+1}$ that restricts to the identity on the boundary sphere $\mathbb{S}^m$ is isotopic to the identity relative to $\mathbb{S}^m$.

Proof. The problem is that the straight line homotopy might not yield an isotopy. Alexander overcame this problem. He put

$$f_t^x := \begin{cases} x & \text{for } |x| \geq t \\ tf\left(\frac{x}{t}\right) & \text{for } |x| < t \end{cases}.$$

Figure A.1 shows how this works: There is an annulus growing from the boundary where we have the identity. The central disc that remains to be filled gets a rescaled version of f. **q.e.d.**

Let

$$\zeta : \mathbb{I} \to \mathbb{I}$$

be a homeomorphism that is the identity on the boundary $\{0, 1\}$ – this is to say ζ preserves the orientation of $\mathbb{I}$. We can now reparameterize any path p by passing to $p \circ \zeta$.

Corollary A.1.5. *Any orientation preserving homeomorphism $\zeta : \mathbb{I} \to \mathbb{I}$ is isotopic to the identity. Thus, reparameterization of paths does not change their homotopy or isotopy class.*

Exercise A.1.6. Prove that there are precisely two homeomorphism of the circle $\mathbb{S}^1$ up to isotopy. Prove that these two are not homotopic.

A.2 Connectivity

Definition A.2.1. A space X is called connected if every cover by disjoint open sets contains at most one non-empty member. The maximal connected subspaces of X are called the (connected) components of X.

The space X is path-connected if any two points in X can be joined by a path. The maximal path-connected subspaces of X are called path-components of X.

Exercise A.2.2. Let X be a locally path-connected space, i.e., the path-components of open subsets in X are open sets. Show that the following are equivalent:

1. X is connected.
2. X is path connected.
3. Every locally constant function on X is globally constant.

Moreover, show that any two points in a connected open subset of $\mathbb{E}^m$ or S_m can be connected by a broken geodesic, i.e., a path that consists of finitely many geodesic segments. (In $\mathbb{E}^m$, this is called a polygonal arc.)

A.3 Covering Spaces

Definition A.3.1. Let $\pi : \bar{X} \to X$ be a continuous map. An open set $U \subseteq X$ is evenly covered by π if the preimage $\pi^{-1}(U)$ is a disjoint union of open sets $U_i' \subseteq \bar{X}$, called sheets, such that

$$\pi|_{U_i'} : U_i' \to U$$

is a homeomorphism for each i.

The map $\pi : \bar{X} \to X$ is a covering projection if every point in X has a neighborhood that is evenly covered. The space $\bar{X}$ together with the map π is a covering space for X if $\bar{X}$ is path connected and π is a covering projection.

Observation A.3.2. *Every covering projection is an open surjection. In particular, any that space has a covering space is path connected.* **q.e.d.**

Remark A.3.3. Every covering projection is a local homeomorphism. The converse is not true.

Exercise A.3.4. Give an example of a local homeomorphism that is not a covering projection.

Example A.3.5. Let $\bar{X}$ be a path connected topological space and G be a group that acts topologically free on $\bar{X}$, i.e., every point has an open neighborhood U such that $gU \cap U = \emptyset$ for all $g \in G - \{1\}$. Then

$$\pi : \bar{X} \to X := {}_{G}\backslash^{\bar{X}}$$

turns $\bar{X}$ into a covering space for the quotient space X.

Exercise A.3.6. Prove the statement of example (A.3.5)

A.3.1 Lifting Maps

Lemma and Definition A.3.7. *Let $\pi : \bar{X} \to X$ be a covering space. For every path $p : \mathbb{I} \to X$ in X issuing from $x = p(0)$ and every point $\bar{x} \in \pi^{-1}(x)$ in the fiber over x there is a unique lift, i.e., a map $\bar{p} : \mathbb{I} \to \bar{X}$ that makes the diagram*

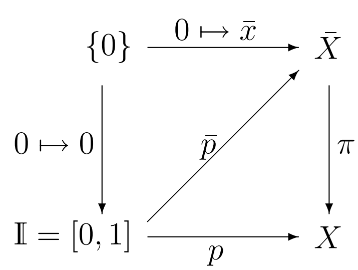

commute.

Proof. We show uniqueness first. Let $\bar{p}_0$ and $\bar{p}_1$ be two lifts of p both starting at $\bar{x}$. Then the set

$$M := \{t \in \mathbb{I} \mid \bar{p}_0(t) = \bar{p}_1(t)\}$$

is visibly closed and non-empty as it contains 0. However, it is also open: Suppose $\bar{p}_0(t) = \bar{p}_1(t)$. Let $U \subseteq X$ be an open neighborhood of $p(t)$ that is evenly covered. Let U' be the sheet above U that contains $\bar{p}_0(t) = \bar{p}_1(t)$. Since $\pi : U' \to U$ is a homeomorphism, we see that the lifts $\bar{p}_0$ and $\bar{p}_1$ agree in an open interval around t. Being an open, closed, and non-empty subset of $\mathbb{I}$, the set M equals $\mathbb{I}$, whence $\bar{p}_0 = \bar{p}_1$.

Now, we turn to existence. Subdivide $\mathbb{I}$ by points

$$0 = t_0 < t_1 < t_2 < \cdots < t_{r-1} < t_r = 1$$

such that for each $i \in \{1, \ldots, r\}$, there is an evenly covered open set U_i in X containing the path segment $p\,|_{[t_{i-1},t_i]}$. Since $\bar{x}$ determines a sheet U_0' above U_0, it is easy to lift the path up to t_1. However, now U_2 takes over, and since t_1 has already been lifted, we are given a sheet above U_2 to continue the construction. Iterating the procedure, we clearly arrive at a lift of p. **q.e.d.**

Corollary A.3.8. *Let Y be a topological space. Suppose the diagram*

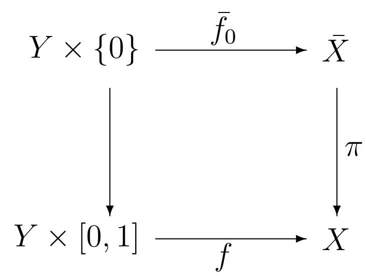

commutes. Then there exists a unique continuous map

$$\bar{f} : Y \times [0,1] \to \bar{X}$$

such that

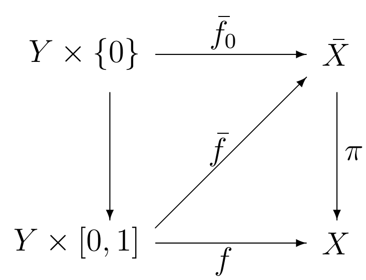

commutes.

Proof. For any fixed $y \in Y$, we have the path $p_y := f(y, -)$. Since these paths have unique lifts, we see that there is a unique function $\bar{f} : Y \times [0,1] \to \bar{X}$ making the diagram commute. The only task is to show that this function is continuous. This is left as an exercise. **q.e.d.**

Exercise A.3.9. Prove that the map $\bar{f}$ is continuous.

Exercise A.3.10 (Covering Homotopy Lemma). Let $\bar{X}$ be a covering space for X with covering projection $\pi : \bar{X} \to X$. Let x_0 and x_1 be two points in X connected by two paths $p, q : \mathbb{I} \to X$ starting at x_0 and ending at x_1 that are homotopic relative to their endpoints. Furthermore, let $\bar{x}_0$ be a point in the fiber over x_0. Consider the two lifts $\bar{p}$ and $\bar{q}$ starting at $\bar{x}_0$. Prove that $\bar{p}(1) = \bar{q}(1)$ and that the paths $\bar{p}$ and $\bar{q}$ are homotopic relative to their endpoints.

Corollary A.3.11. *Every lift of a homotopically trivial loop in any covering space is a homotopically trivial loop.* **q.e.d.**

A.3.2 The Fundamental Group

Notation A.3.12. If $p : \mathbb{I} \to X$ and $q : \mathbb{I} \to X$ are paths such that $p(1) = q(0)$, we can concatenate these paths:

$$\begin{aligned} p \rightarrowtail q : \mathbb{I} &\to X \\ t &\mapsto \begin{cases} p(2t) & \text{for } 0 \leq t \leq \frac{1}{2} \\ q(2t-1) & \text{for } \frac{1}{2} \leq t \leq 1. \end{cases} \end{aligned}$$

Definition A.3.13. Let X be a topological space with base point $\underline{x}$. The <u>fundamental group</u> $\pi_1(X, \underline{x})$ of X based at $\underline{x}$ is the set of homotopy classes of paths starting and ending at $\underline{x}$ relative to their endpoints, i.e:

$$\pi_1(X, \underline{x}) := \{p : \mathbb{I} \to X \mid p(0) = p(1) = \underline{x}\} \Big/ \begin{array}{l} p = q \text{ if } p \text{ is homotopic to } q \text{ relative to} \\ \text{both endpoints.} \end{array}$$

Since all these paths are closed loops, we can concatenate these loops and it is clear that concatenation of loops descends to a well defined multiplication on homotopy classes. This way, we define a multiplication in $\pi_1(X, \underline{x})$. It is easy to check that this multiplication turns the set $\pi_1(X, \underline{x})$ into a group.

Definition A.3.14. If $f : (X, \underline{x}) \to (Y, \underline{y})$ is a map of base pointed spaces, then there is an induced map

$$\pi_1(f) : \pi_1(X, \underline{x}) \to \pi_1(Y, \underline{y})$$

defined by pushing loops in X to Y via f.

Observation A.3.15. *If $h : (Y, \underline{y}) \to (Z, \underline{z})$ is another map of base pointed spaces, then*

$$\pi_1(h \circ f) = \pi_1(h) \circ \pi_1(f)\,.$$

In other words, the fundamental group is a covariant functor.

Example A.3.16. Let $(\bar{X}, \bar{\underline{\mathrm{x}}})$ be a covering space of the base pointed space $(X, \underline{\mathrm{x}})$ with base point preserving covering projection $\pi : \bar{X} \to X$. Then the induced homomorphism

$$\pi_1(\pi) : \pi_1\big(\bar{X}, \bar{\underline{\mathrm{x}}}\big) \to \pi_1(X, \underline{\mathrm{x}})$$

is injective because of (A.3.11).

Observation A.3.17. *Let $\underline{x}_0$ and $\underline{x}_1$ be two points in X and fix a path p from $\underline{x}_0$ to $\underline{x}_1$. Then p induces an isomorphism*

$$\Phi_p : \pi_1(X, \underline{x}_1) \to \pi_1(X, \underline{x}_0)$$

by means of the lollipop construction:

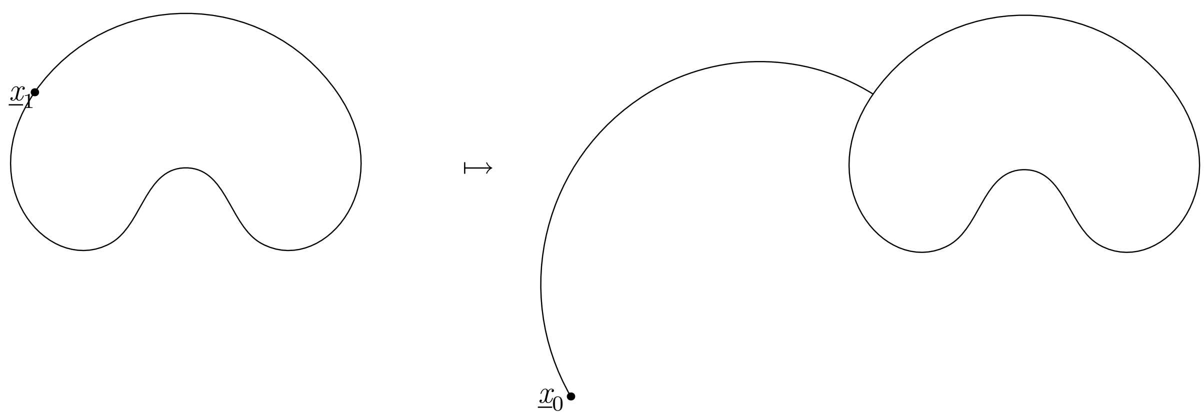

A.3.3 Morphisms of Covering Spaces

Proposition A.3.18. *Let $(Y, \underline{y})$ be 1-connected and locally path connected and let $\bar{X}$ be a covering space with covering projection $\pi : \bar{X} \to X$. For each map $f : (Y, \underline{y}) \to (X, \underline{x})$ of base pointed spaces and each choice of a lift $\bar{\underline{x}} \in \pi^{-1}(\underline{x})$, there is a unique continuous lift*

$$\bar{f} : (Y, \underline{y}) \to (\bar{X}, \bar{\underline{x}}).$$

Proof. Since Y is path-connected, we know precisely how the lift $\bar{f}$ has to be defined: any point $y \in Y$ is the endpoint of some path p_y in Y starting at $\underline{\mathrm{y}}$. The path $f \circ p_y$ is a path in X starting at $\underline{\mathrm{x}}$ and has a unique lift in $\bar{X}$ starting at $\bar{\underline{\mathrm{x}}}$. The endpoint of that lift has to be $\bar{f}(y)$.

Since Y is simply connected, this value $\bar{f}(y)$ is well defined: Any two paths p_y and q_y connecting $\underline{\mathrm{y}}$ to y are homotopic relative to their endpoints in Y. The same holds true for $f \circ p_y$ and $f \circ q_y$. Thus the corresponding lifts have identical endpoints by (A.3.10).

It remains to show that $\bar{f}$ is continuous. We will prove that $\bar{f}$ is continuous everywhere. So let y be any point in Y. We find an open neighborhood U in X around $f(y)$ that is evenly covered. Let U' be the sheet above U that contains the lift $\bar{f}(y)$. Let $\zeta : U \to U'$ be the local homeomorphism inverse to π. Since f is open, the preimage $f^{-1}(U)$ is an open neighborhood V around y. It would be nice, if we could argue that

$$\bar{f}\,|_V = \zeta \circ f\,|_V$$

since this would visibly imply that $\bar{f}$ is continuous near y. However, life is not that easy. The problem is that in evaluating $\bar{f}$ at, say, $y' \in V$, we will want to use a path from y to y' that we can push to U via f and lift via ζ. Unfortunately there might be no path from y to y' that stays inside V.

To fix this problem, we use the hypotheses that Y is locally path connected. Thus, we can shrink V to $V' \subseteq V$ such that any point $y' \in V'$ can be connected to y by a path in V. Now, these paths will be pushed into U via f and therefore lift nicely by means of ζ. Thus

$$\bar{f}\,|_{V'} = \zeta \circ f\,|_{V'}$$

which also shows $\bar{f}$ to be continuous near y. **q.e.d.**

Definition A.3.19. Let $\bar{X}_0$ and $\bar{X}_1$ be two covering spaces over X with projections $\pi_i : \bar{X}_i \to X$. A map

$$f : \bar{X}_0 \to \bar{X}_1$$

is a X-map if the diagram

$$\begin{array}{ccc} \bar{X}_0 & \xrightarrow{f} & \bar{X}_1 \\ \downarrow{\scriptstyle \pi_0} & & \downarrow{\scriptstyle \pi_1} \\ X & = & X \end{array}$$

commutes. Note that compositions of X-maps are X-maps.

The two covering spaces $\bar{X}_0$ and $\bar{X}_1$ are equivalent if there are X-maps $f : \bar{X}_0 \to \bar{X}_1$ and $h : \bar{X}_1 \to \bar{X}_0$ such that $f \circ h = \mathrm{id}_{\bar{X}_1}$ and $h \circ f = \mathrm{id}_{\bar{X}_0}$.

For any covering space $\bar{X}$, let

$$\mathrm{Cov}\big(\bar{X}/X\big)$$

denote the group of self-X-maps of $\bar{X}$.

Observation A.3.20. *Every X-map is a covering projection. To see this, consider an open subset U' of $\bar{X}_1$ that small enough so that its image in X is evenly covered. Then $f^{-1}(U')$ is contained in ${\pi_0}^{-1}(\pi_1(U'))$. It follows that U' is evenly covered.*

Definition A.3.21. A simply connected base pointed covering space $(\tilde{X}, \underline{\tilde{x}})$ with covering projection $\tilde{\pi} : (\tilde{X}, \underline{\tilde{x}}) \to (X, \underline{x})$ is called a universal cover of X.

Observation A.3.22. *Let $(\bar{X}, \underline{\bar{x}})$ be a covering space for the base pointed space $(X, \underline{x})$ with covering projection $\pi : (\bar{X}, \bar{X}) \to (X, \underline{x})$ be a covering space. Then there is a unique X-map $\tilde{X} \to \bar{X}$ that takes $\underline{\tilde{x}}$ to $\underline{\bar{x}}$. This follows immediately from (A.3.18).* **q.e.d.**

Corollary A.3.23. *Up to equivalence, there is at most one universal cover.* **q.e.d.**

Theorem A.3.24. *If $(\tilde{X}, \underline{\tilde{x}})$ is a universal cover for a locally path connected space $(X, \underline{x})$, then*

$$\mathrm{Cov}\left(\tilde{X}/X\right) = \pi_1(X, \underline{x}).$$

Proof. We give a pair of mutually inverse homomorphisms. The homomorphism

$$\mathrm{Cov}\left(\tilde{X}/X\right) \to \pi_1(X, \underline{x})$$

is given as follows: Let $f : \tilde{X} \to \tilde{X}$ be a deck transformation. There is path from $\underline{\tilde{x}}$ to $f(\underline{\tilde{x}})$. This path projects down to a loop in X centered at $\underline{x}$. This construction yields a well defined map since any two paths in $\tilde{X}$ connecting $\underline{\tilde{x}}$ to $f(\underline{\tilde{x}})$ are homotopic relative to their endpoints. The map is visibly a homomorphism: Composition of deck transformation turns into concatenation of paths and then concatenation of loops.

For the other direction

$$\pi_1(X, \underline{x}) \to \mathrm{Cov}\left(\tilde{X}/X\right)$$

we start with a loop γ in X based at $\underline{x}$. It lifts to a unique path p in $\tilde{X}$ starting at $\underline{\tilde{x}}$. By (A.3.18), there is a unique lift f of $\tilde{\pi}$ that takes $\underline{\tilde{x}}$ to the endpoint of p. This construction is easily seen to yield a homomorphism.

It is routine to check that these two homomorphisms are mutually inverse. **q.e.d.**

Corollary A.3.25. $\pi_1(\mathbb{S}^1) = C_\infty$. **q.e.d.**

Definition A.3.26. A space X is semilocally simply connected if every point has a neighborhood such that every closed path in that neighborhood is homotopically trivial in X.

Theorem A.3.27 (Existence of the Universal Cover). *Every connected, locally path connected, semilocally simply connected space* $(X, \underline{x})$ *has a universal cover. Indeed, the quotient*

$$\tilde{X} := \{p : \mathbb{I} \to X \mid p \text{ is continuous and } p(0) = \underline{x}\}/\sim$$

is a universal cover, where the equivalence relation $p \sim q$ *holds if and only if* p *and* q *are homotopic relative endpoints. The topology on* $\mathcal{P}_{\underline{x}}(X) := \{p : \mathbb{I} \to X \mid p \text{ is continuous and } p(0) = \underline{x}\}$ *is defined by the basic open sets* $U_{p,V}$ *where* V *is an open neighborhood of the endpoint* $p(1)$*. The basic open set* $U_{p,V}$ *is defined as*

$$U_{p,V} := \{p \rightarrowtail q \mid q \text{ is a path in } V \text{ starting at } p(1)\}.$$

The base point of $\tilde{X}$ *is given as the homotopy class of the constant path that stays at* $\underline{x}$.

The covering projection is induced by the endpoint map $p \mapsto p(1)$.

Proof. Let us see that $\tilde{X}$ is simply connected. So let

$$\begin{array}{rcl} p : \mathbb{I} & \to & \mathcal{P}_{\underline{\mathrm{x}}}(X) / \sim \\ t & \mapsto & p_t \end{array}$$

be a closed path in $\tilde{X}$. By (A.3.28), any representative of the class p_t is homotopic relative endpoints to the path

$$\begin{array}{rcl} q_t : \mathbb{I} & \to & X \\ s & \mapsto & p_t(ts). \end{array}$$

In particular, q_t is a representative of the class p_t. Moreover, q_t is just an initial segment of q_1. Since $p_1 = p_0$, the path q_1 is homotopic to the constant path. A contracting homotopy for q_1 restricts to all the initial segments q_t and, therefore, induces a contracting homotopy of p in $\tilde{X}$. **q.e.d.**

Exercise A.3.28. Let

$$\begin{array}{rcl} p : \mathbb{I} & \to & \mathcal{P}_{\underline{\mathrm{x}}}(X) / \sim \\ t & \mapsto & p_t \end{array}$$

be a path. Prove that any representative of the class p_1 is homotopic relative endpoints to the path

$$\begin{array}{rcl} q : \mathbb{I} & \to & X \\ t & \mapsto & p_t(1). \end{array}$$

Exercise A.3.29. Prove that the endpoint map $p \mapsto p(1)$ induces a covering projection $(\mathcal{P}_{\underline{\mathrm{x}}}(X) / \sim) \to X$.

Appendix B

Hyperbolic Geometry

B.1 The Upper Half Plane Model

Definition B.1.1. The hyperbolic plane is the set

$$\mathbb{H}^2 := \{(x, y) \mid y > 0\} = \{z \in \mathbb{C} \mid \Im(z) > 0\}$$

with the metric

$$\mathrm{d}\, s^2 = \frac{\mathrm{d}\, x^2 + \mathrm{d}\, y^2}{y^2}.$$

Remark B.1.2. This means that the length of a path $p(t) = (x(t)\,, y(t))$ traversed from time t_0 to t_1 is given by the integral

$$\int_{t_0}^{t_1} \sqrt{\frac{x'(t)^2 + y't^2}{y(t)^2}}\,\mathrm{d}\, t = \int_{t_0}^{t_1} \frac{|p'(t)|}{y(t)}\,\mathrm{d}\, t.$$

In other words, the lengths of tangent vectors at height y are rescaled by $\frac{1}{y}$, regardless of their direction. At each point, we see just a rescaled version of the standard metric.

Remark B.1.3. The area element is given by $\mathrm{d}\, A = \frac{\mathrm{d}\, x\, \mathrm{d}\, y}{y^2}$. **q.e.d.**

B.1.1 Some Orientation Preserving Isometries

Note that the upper half plane is also the northern hemisphere in the Reimann-sphere, $\mathbb{P}(\mathbb{C})$. Note that the equator, consisting of real numbers and the point at infinity, is an embedded copy of $\mathbb{P}(\mathbb{R})$. Thus $\mathrm{GL}_2(\mathbb{R}) \leq \mathrm{GL}_2(\mathbb{C})$ acts on $\mathbb{P}(\mathbb{C})$ stabilizing the equator. Hence every element of $\mathrm{GL}_2(\mathbb{R})$ either interchanges the northern and southern hemispheres or leaves them stabile. The index-2-subgroup $\mathrm{GL}_2^+(\mathbb{R})$ of matrices with positive determinant therefore acts on the upper half plane. Since the complex number z represents the projective point $(z : 0)$, it follows that the action is by Möbius transformations, i.e., the matrix $\begin{pmatrix} a & b \\ c & d \end{pmatrix}$ sends z to $\frac{az+b}{cz+d}$. Our goal in this section is to verify that this is an action by isometries relative to the given metric.

The metric above is defined by specifying an inner product in the tangent space of each point of the hyperbolic plane – in fact, we have specified a positive definite quadratic form which induces an inner product by means of the parallelogram identity. Note that isometries of $\mathbb{H}^2$ are precisely those diffeomorphims whose derivatives are isometries of tangent spaces. Thus, it is in principle an easy matter of Calculus to check whether a map is an isometry. However, for some maps, one does not even need to compute the derivatives since everything is plain obvious.

Observation B.1.4. *The horizontal translations*

$$\begin{aligned} \begin{pmatrix} 1 & b \\ 0 & 1 \end{pmatrix} : z &\mapsto z + b \\ (x, y) &\mapsto (x + b, y) \end{aligned}$$

are isometries of $\mathbb{H}^2$: They do not affect the y-coordinate, and all tangentspaces on the same height had their metrics rescaled by the same factor. **q.e.d.**

Observation B.1.5. *The central dilations*

$$\begin{pmatrix} a & 0 \\ 0 & d \end{pmatrix} : z \mapsto \frac{a}{d} z$$

are isometries of $\mathbb{H}^2$: they do not affect the y-coordinate, and all tangentspaces on the same height had their metrics rescaled by the same factor. **q.e.d.**

Observation B.1.6. *The map*

$$\begin{aligned} z &\mapsto -\frac{1}{\overline{z}} \\ (x, y) &\mapsto \left(\frac{x}{x^2 + y^2}, \frac{y}{x^2 + y^2}\right) \end{aligned}$$

is an orientation reversing isometry of $\mathbb{H}^2$. Therefore, it cannot be realised as a Möbius transformation. To see that this map is an isometry, we will understand this map geometrically: First, note that it stabilizes the lines through the origin $(0, 0)$. Moreover, it sends the half-circle centered at $(0, 0)$ of radius r to the half-circle of radius $\frac{1}{r}$. With this information, it is easy to see what the derivative does to a suitably *chosen orthonormal basis of the tangent space at (x, y). We choose one vector pointing away from the origin and the other one tangent to the half-circle through (x, y). The scaling is exactly as desired.*

Exercise B.1.7. Infer that the map

$$\begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix} : z \mapsto \frac{-1}{z}$$

is an orientation preserving isometry of $\mathbb{H}^2$.

Corollary B.1.8. *The group* $\mathrm{SL}_2(\mathbb{R})$ *acts on* $\mathbb{H}^2$ *as a group of orientation preserving isometries.*

Proof. By the preceding observations and the above exercise, the claim follows from the fact that

$$\mathrm{SL}_2(\mathbb{R}) = \left\langle \begin{pmatrix} 1 & \star \\ 0 & 1 \end{pmatrix}, \begin{pmatrix} \star & 0 \\ 0 & \star \end{pmatrix}, \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix} \right\rangle \tag{B.1}$$

where all the matrices on the right hand are supposed to have determinant 1.

Indeed, the strict upper triangular matrices and the diagonal matrixes in $\mathrm{SL}_2(\mathbb{R})$ generate the group of upper triangular matrices in $\mathrm{SL}_2(\mathbb{R})$. This group is the stabilizer of $\infty \in \mathbb{P}^1(\mathbb{R})$. The matrix $\begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}$ swaps ∞ and 0 in $\mathbb{P}^1(\mathbb{R})$. Thus, we find the stabilizer of 0 in the group generated on the right hand of (B.1). This stablizer is the group of lower triangular matrices. Since every matrix can be written as a product of upper and lower triangular matrices, the claim follows. **q.e.d.**

B.1.2 Geodesics

Geodesics in a metric space X are isometrically embedded curves, i.e., maps $\lambda : [0, \ell] \to X$ such that $d(\lambda(t_0), \lambda(t_1)) = |t_0 - t_1|$. Such curves connect their endpoints via a shortes possible path. In general, we cannot expect points to be connected via a geodesic, nor will geodesics be automatically unique. Note however that we did not yet define a metric on $\mathbb{H}^2$, we just defined a Riemannian metric. The associated notion of distance is given as follows:

Definition B.1.9. The distance of any two points x and y in $\mathbb{H}^2$ is the infimum of lengths of curves connecting x and y.

Remark B.1.10. The operational word in this definition is "infimum". We do not actually know a priory that a length-minimizing curve exists. Thus, we may not say minimum. The notion of distance above clearly satisfies the triangle inequality. However, it might not be a metric for a silly reason: two distinct points could have distance 0. Of course, this does not happen in $\mathbb{H}^2$.

In this section we shall describe all geodesics in the hyperbolic plane. We shall see that any two points in $\mathbb{H}^2$ are joined by a unique geodesic segment.

Observation B.1.11. *Any two points on the imaginary axis are joined by a unique geodesic segment, namely the interval that they span on the imaginary axis. Moreover,*

$$\begin{array}{rcl} \mathbb{R} & \longrightarrow & \mathbb{H}^2 \\ t & \mapsto & \exp(t) \end{array}$$

is an isometric embedding. Thus, the imaginary axis is a geodesic line in $\mathbb{H}^2$*.*

Proof. The curcial observation is that the map

$$(x, y) \mapsto (0, y)$$

is length-nonincreasing and strictly length-decreasing for any curve that is not a straight vertical line. **q.e.d.**

Exercise B.1.12. Show that any two points can be simultaneously taken to the imaginary axis by a Möbius transformation.

Corollary B.1.13. *Any two points in* $\mathbb{H}^2$ *are connected by a unique geodesic segment, i.e.,* $\mathbb{H}^2$ *is a* <u>*geodesic space*</u>*. Moreover, any geodesic segment can be extended to a bi-infinite geodesic line, i.e.,* $\mathbb{H}^2$ *is* <u>*geodesically complete*</u>*.* **q.e.d.**

We obtain other geodesics in $\mathbb{H}^2$ by applying isometries.

Observation B.1.14. *Horizontal shifts show that every straight vertical line is a geodesic line.*

We apply the orientation reversing isometry

$$\begin{array}{rcl} z & \mapsto & -\frac{1}{\bar{z}} \\ (x, y) & \mapsto & (\frac{x}{x^2 + y^2}, \frac{y}{x^2 + y^2}) \end{array}$$

to the vertical line $x = 1$ *and see that the half-circle centered at* $x = \frac{1}{2}$ *is a geodesic. By horzontal shifts and central dilatations, we find that all half-circles with centers on the real line are geodesics.*

Finally, these are all geodesics, because geodesic segments are unique. **q.e.d.**

Exercise B.1.15. Show that, in the upper half plane model, the map

$$z \mapsto \frac{1}{\bar{z}}$$

is a reflection along a geodesic.

Exercise B.1.16. Show that hyperbolic circles in the upper half plane model are Euclidean circles. (The Euclidean and hyperbolic center, however, do not coincide.)

Exercise B.1.17 (Gauss-Bonet). In a hyperbolic triangle, we define angles by drawing Euclidean tangent lines and measuring their Euclidean angles. Show that the area of a triangle with angles α, β, and γ is $\pi - \alpha - \beta - \gamma$.

Appendix C

References

[Mois77] *E.E. Moise:* Geometric Topology in Dimensions 2 and 3. Springer GTM 47