/ university of groningen

faculty of science and engineering

Delaunay triangulations on hyperbolic surfaces



Master Project Mathematics

August 21, 2017

Student: Y.M. Ebbens

First supervisor: Prof.dr. G. Vegter

Second supervisor: Dr. A.E. Sterk

Abstract

There exist several algorithms to compute the Delaunay triangulation of a point set in a Euclidean space. In the literature the incremental algorithm has been extended to Euclidean orbifolds and hyperbolic surfaces, but it is only guaranteed to work in a finitely sheeted covering space. Bounds for the minimum number of sheets that are needed are only known for the Bolza surface. In this thesis we will first prove that the lower bound for the number of sheets is $\Omega(g)$ in general and $\Omega(g^3)$ if the surface can be represented by a fundamental region with 4g concircular vertices, where g is the genus of the surface. Then we will prove that there does not exist an upper bound for the number of sheets necessary for surfaces of genus 2.

The systole of a surface plays an important role in determining the complexity of triangulating the corresponding surface. We will state a conjecture for the systole of hyperbolic surfaces of genus g represented by regular 4g-gons.

Finally, to avoid many sheeted covering spaces a different method uses a well chosen set of points, which guarantees that the output is simplicial. We will show a lower bound for the number of points of such a point set, which is of order $\Omega(\sqrt{g})$ in general and of order $\Omega(g)$ if the systole of a family of surfaces is bounded.

Keywords: hyperbolic geometry, Delaunay triangulation, hyperbolic surface, Fuchsian group, Teichmüller space, systole

Contents

1	Introduction	5
	1.1 Context and state of the art	5
	1.2 Applications	6
	1.3 Structure	6
2	Statement of the problem	7
3	Summary of results	9
4	Hyperbolic geometry	11
	4.1 History	11
	4.2 Models of hyperbolic geometry	12
	4.2.1 Upper half-plane model	12
	4.2.2 Poincaré disk model	13
	4.2.3 Relation	14 14
	4.5 Classification of Mobius transformations	14 15
	4.5 Trigonometric optimization problems	17
	4.6 Circle packings	20
	4.7 Fuchsian groups	$\frac{-\circ}{22}$
	4.8 Hyperbolic surfaces	23
	4.9 Closed geodesics on a hyperbolic surface	24
	4.10 Teichmüller space	26
	4.10.1 Pairs of pants \ldots	27
	4.10.2 Cubic graphs	28
	4.10.3 Twist parameters	29
	4.10.4 Fenchel-Nielsen coordinates	30
	4.10.5 Teichmüller space \ldots	31
	4.10.6 Mapping class group	32
	4.10.7 Zieschang-Vogt-Coldewey coordinates	32
5	Triangulations	35
6	Lower bound for the number of sheets	38
7	Upper bound for the number of sheets	41
8	Systole of surfaces corresponding to regular polygons	45
	8.1 From systole to optimization problem	45
	8.2 Upper bound for the systole	46
	8.3 Towards a complete proof	47
	8.4 Summary and speculations	52
	8.5 Examples	53
9	Lower bound for the number of points in a Delaunay triangulation	55

10 Future work	58	
11 Acknowledgements	58	
A Algebra	59	

1 Introduction

1.1 Context and state of the art

In computational geometry, geometric objects are often described or approximated by using discrete methods. For example, representing surfaces in a computer can be done by computing a triangulation of a point set on the surface. To avoid long and skinny triangles, a Delaunay triangulation is used, since among all triangulations it maximizes the minimum angle between adjacent edges. In a Delaunay triangulation the interior of the circumscribed circle of each triangle does not contain any points from the point set. It is closely related to the Voronoi diagram of the point set. The Voronoi diagram partitions the ambient space into Voronoi regions, each of which contains all points closer to some given point from the point set than the other points. As cell complexes, the Delaunay triangulation and the Voronoi diagram are dual. This means that faces in one graph correspond to vertices in the other and two faces share an edge if and only if the corresponding vertices in the dual graph are joined by an edge.

There are several algorithms to compute Delaunay triangulations in *n*-dimensional Euclidean space (see, e.g., [11, 20]). The incremental algorithm, first described in [13, 49], inserts the points one at a time, deletes the simplices containing the added point and retriangulates this region. In the PhD thesis of Manuel Caroli [16] and the following paper [17], this algorithm was extended to closed Euclidean orbifolds, the sphere and spherical orbit spaces. In the PhD thesis of Mikhail Bogdanov [9] and following articles [10, 8], the algorithm was extended to *n*-dimensional hyperbolic space and closed hyperbolic surfaces. For Euclidean orbifolds and hyperbolic surfaces the algorithm is only guaranteed to work in a finite-sheeted covering space of the original space, but the minimum number of sheets that are needed is not known. Namely, the output of the algorithm could fail to be a triangulation due to loops or double edges in the graph. In [8] it is shown that for the Bolza surface, arguably the simplest hyperbolic surface of genus 2 to consider, the number of sheets is between 33 and 128.

In this thesis we will look at Delaunay triangulations of closed hyperbolic surfaces. Because hyperbolic surfaces are locally isometric to open subsets of the hyperbolic plane, they come equipped with a Riemannian metric of constant Gaussian curvature -1. The hyperbolic structure of a hyperbolic surface induces a conformal structure, i.e., the structure of a Riemann surface. We will turn to this again in Section 2 and explain the terminology in detail in Section 4.8.

As mentioned before, the minimum number of sheets of a suitable covering space of a given hyperbolic surface is not known, even in simple cases such as the Bolza surface. This thesis focuses on finding bounds for the number of sheets for other hyperbolic surfaces, for example with higher genus and different conformal structures. Because the number of sheets is probably too large to be of practical use, a second method uses dummy point sets to avoid many sheeted covering spaces. Here the algorithm is started with a well chosen set of points, which guarantees that the output is simplicial. Then the points from the point set are added. In most cases, the points from the dummy point set can be removed afterwards. Both problems use the notion of systole, i.e., the length of the shortest homotopically non-trivial closed geodesic, so a third problem is working on

bounds for systoles of classes of surfaces.

1.2 Applications

Even though our primary focus is mathematical, there are several applications for Delaunay triangulations on hyperbolic surfaces, usually for fields which use periodic objects or objects with periodic boundary conditions. For example, some cosmological models use that the structure of the universe is considered to be periodic on a sufficiently large scale to be able to replicate a sample of the structure periodically [31]; here the Bolza surface is studied as a cosmological model with non-trivial topology [4]. Furthermore, the periodic orbits of the Bolza surface are studied as a model for quantum chaos [3, 5, 41]. In molecular dynamics, the Bolza surface is used to model periodic boundary conditions to study the fragility of a glassforming liquid [39, 40]. In mathematical neuroscience, patterns on hyperbolic surfaces are used as models for the neural organization of the brain, in particular visual texture perception [18, 21]. Of course, Delaunay triangulations in Euclidean space are used in geometric modelling as well, for example to approximate surfaces [38].

1.3 Structure

This thesis is structured in the following way. In Section 2 we will give a more precise statement of the problem, followed by a summary of our results in Section 3. In Sections 4 and 5 we will introduce the necessary background on hyperbolic geometry and triangulations, respectively. In Sections 6 to 9 we will look at the results in turn: Section 6 will give a lower bound for the number of sheets of order $\Omega(g)$ in general and of order $\Omega(g^3)$ if the surface can be represented by a fundamental region with 4g concircular vertices, where g is the genus of the surface; Section 7 will show that there does not exist an upper bound for the number of sheets necessary for surfaces of genus 2, i.e. even when we restrict ourselves to surfaces of genus 2, the minimum number of sheets can be made arbitrarily large; Section 8 will give a conjecture on the systole of surfaces obtained from regular polygons, in particular that this is bounded with respect to g; finally, Section 9 will give a lower bound for the size of a dummy point set, which is of order $\Omega(\sqrt{g})$ in general and of order $\Omega(g)$ for families of surfaces with bounded systole. Section 10 will conclude with suggestions for future work.

2 Statement of the problem

A triangulation is a subdivision of a topological space into subsets, each of which is homeomorphic to a Euclidean triangle, where the intersection of two of these subsets is homeomorphic to the empty set or a vertex or edge or face of a triangle. In Euclidean, spherical or hyperbolic space we can speak of a Delaunay triangulation: a triangulation is Delaunay if the circumscribed disk of every triangle does not contain any vertex of the triangulation in its interior. For every point set in such spaces there exists a Delaunay triangulation with this point set as vertices and if there are no subsets of at least four concircular vertices, then this Delaunay triangulation is unique. Among all triangulations, the Delaunay triangulation stands out, since it maximizes the minimum angle between adjacent edges among all triangulations. Hence, long and skinny triangles, which might cause problems in applications, are avoided as far as possible.

In this thesis we consider only triangulations of closed hyperbolic surfaces. Let \mathbb{D}^2 be the Poincaré disk model of the hyperbolic plane [6]. A hyperbolic surface is a connected 2dimensional manifold, which is locally isometric to an open subset of \mathbb{D}^2 . Every hyperbolic surface M can be written as a quotient space $M = \mathbb{D}^2/\Gamma$, where Γ is a Fuchsian group, i.e., a discrete group consisting of isometries of the hyperbolic plane. It has been shown that the set of hyperbolic surfaces can be parametrized in a so called Teichmüller space [14]. To find the Delaunay triangulation of a finite set of points S on a hyperbolic surface M, one can instead compute the periodic Delaunay triangulation of ΓS , i.e., the images of S under Γ , in \mathbb{D}^2 . If we project this triangulation using the projection $\pi : \mathbb{D}^2 \to M$, then we obtain the "Delaunay triangulation" of S in M. However, the resulting object is not always a triangulation as defined above: it can happen that different vertices of one triangle project to the same point, leading to loops or double edges in the graph.

To avoid this problem, we can instead project to a covering space M' of M with universal covering projection $\pi' : \mathbb{D}^2 \to M'$. It has been shown that there exists a finite-sheeted covering space M' of M such that the projection of the Delaunay triangulation of ΓS in \mathbb{D}^2 under π' is indeed a triangulation. Such a covering space must satisfy the inequality syst $(M') > 2\delta_S$: here syst(M') denotes the systole of M', which is the length of the shortest, homotopically non-trivial, closed geodesic of the surface; δ_S denotes the diameter of the largest disk in \mathbb{D}^2 not containing any points of ΓS in its interior. It follows that we can speak of the minimum number of sheets necessary for such a covering space. In [9] it was shown that the minimum number of sheets for the Bolza surface is between 33 and 128. As far as we know, there are no results for other surfaces, either with higher genus or with different conformal structures.

As we will see in Section 6, the minimum number of sheets of a suitable covering space is of order $\Omega(g)$ in general and of order $\Omega(g^3)$ in a special case (see the next section for the notation). For applications this would not be efficient. Therefore, we also investigate the second approach mentioned in [8]. Here we initialize the triangulation with a well chosen fixed point set P on the surface M for which $syst(M) > 2\delta_P$. This inequality will continue to hold for larger point sets, so we can then use the incremental algorithm to add the given point set. In most cases, the dummy points can be removed afterwards. Of course, we want our dummy point set to have the smallest possible cardinality, as each point adds to the complexity of the algorithm. For the Bolza surface a point set consisting of 14 points is given in [8].

3 Summary of results

As our results make use of notation regarding computational complexity, we will briefly explain this notation. Let $f : \mathbb{N} \to \mathbb{R}$ be a function. We say that $h : \mathbb{N} \to \mathbb{R}$ is of order $\Omega(f(g))$ if there exists a constant $c \in \mathbb{R}$ and $g_0 \in \mathbb{N}$ such that $h(g) \ge cf(g)$ for all $g \ge g_0$. Intuitively, this means that asymptotically and up to constants, h grows at least as fast as f. Secondly, we say that $h \sim f$ if $h(g)/f(g) \to 1$ as $g \to \infty$. Usually we will use g as variable, since this variable denotes the genus of a given surface.

Now we will state our results. As mentioned in the previous section, the incremental algorithm can be used to compute Delaunay triangulations of point sets on hyperbolic surfaces, but it is only guaranteed to work in a finitely sheeted covering space. We will show in Section 6 that the number of sheets necessary for such a covering space is of order $\Omega(g)$ in general and of order $\Omega(g^3)$ if the surface can be represented by a fundamental region with 4g concircular vertices. This result improves and generalizes the currently known lower bound for the number of sheets.

Theorem 3.1. Let M be any hyperbolic surface of genus $g \ge 2$. Let M' be a k-sheeted covering space of M such that $syst(M') > 2\delta_M$. Then

$$k > \frac{\pi}{9} \cdot \frac{\cot^2(\frac{\pi}{12g-6}) - 3}{g-1} \sim \frac{16}{\pi} \cdot g$$

Furthermore, if M can be represented by a fundamental region with 4g concircular vertices, then we have the higher upper bound

$$k > \frac{\pi}{3} \cdot \frac{\cot^4 \frac{\pi}{4g} - 1}{g - 1} \sim \frac{256}{3\pi^3} \cdot g^3.$$

In Section 7 we will prove that there does not exist an upper bound for the number of sheets necessary for surfaces of genus 2.

Theorem 3.2. For all $B \in \mathbb{R}$ there exists a hyperbolic surface M of genus 2 such that if M' is a k-sheeted covering space of M such that $syst(M') > \delta_M$, then k > B.

Here, 'all hyperbolic surfaces' refers to all elements of the Teichmüller space of hyperbolic surfaces of genus 2.

The systole of a surface is an important concept to determine the complexity of computing a Delaunay triangulation. In Section 8 we state the following conjecture about the systole of hyperbolic surfaces M_g of genus $g \ge 2$ corresponding to regular 4g-gons. In particular, the systole of this family of surfaces is bounded, by which we mean that the set {syst(M_g) | $g \in \mathbb{N}, g \ge 2$ } is bounded as a subset of \mathbb{R} .

Conjecture 3.3. The systole of the surface M_g corresponding to the regular 4g-gon satisfies

$$\cosh\left(\frac{\operatorname{syst}(M_g)}{2}\right) = 1 + 2\cos(\frac{\pi}{2g}).$$

As a partial result we prove the following theorem.

Theorem 3.4. The systole of the surface M_g corresponding to the regular 4g-gon satisfies

$$\cosh\left(\frac{\operatorname{syst}(M_g)}{2}\right) \le 1 + 2\cos(\frac{\pi}{2g}),$$

with equality for g = 2, 3.

The inequality for arbitrary genus is new. The systole for M_2 was already known in the literature, but the method developed in Section 8 leads to a new proof, as well as to a proof for g = 3.

For applications the number of sheets necessary is usually too large. A second method initializes the triangulation with a well chosen point set, so that the output is guaranteed to be simplicial. In Section 9 we will prove the following proposition. We will also show that this implies that the number of points of such a dummy point set is of order $\Omega(\sqrt{g})$ in general and of order $\Omega(q)$ if the systole of a family of surfaces is bounded.

Proposition 3.5. Let M be a hyperbolic surface of genus $g \ge 2$. Let P be a set of points in M such that $syst(M) > 2\delta_P$. Then

$$|P| > \left(\frac{\pi}{\pi - 6\operatorname{arccot}(\sqrt{3}\operatorname{cosh}(\frac{1}{4}\operatorname{syst}(M)))} - 1\right) \cdot 2(g-1).$$

4 Hyperbolic geometry

4.1 History

Around 300 B.C. Euclid wrote in the first book of his $Elements^1$:

- 1. Let it be postulated that from every point to every point we can draw a straight line,
- 2. and that from a bounded straight line we can produce an unbounded straight line,
- 3. and that for every center and distance we can draw a circle,
- 4. and that all right angles are identical to each other,
- 5. and that, if a straight line intersecting two other straight lines makes the interior angles on one side less than two right angles, then the two straight lines, extended to infinity, intersect on the side where the angles are less than two right angles.

Given the first four postulates, the fifth postulate can be shown to be equivalent with the parallel postulate: given a line and a point not on this line, there exists precisely one line through the point parallel to the given line. The first four of Euclid's postulates are intuitively clear, but the fifth has been cause for much debate. It has been regarded as not self-evident enough to be assumed without proof, but for over two thousand years it could not be proved from the other postulates.

In the first half of the nineteenth century², the construction of so called non-Euclidean geometry by Lobachevsky and Bolyai (independently) proved that the attempts would be fruitless from the start. In both Euclidean and non-Euclidean geometry the first four of Euclid's postulates hold, but in the latter the fifth postulate does not hold. This early non-Euclidean geometry is usually called Bolyai-Lobachevsky geometry and formed the basis of hyperbolic geometry. It should be noted that several years before Lobachevsky and Bolyai published their findings Gauss described similar ideas in a letter, but he never published his construction.

Initially, the study of non-Euclidean geometry existed separately from the rest of mathematics. However, in 1868 Beltrami showed that two-dimensional non-Euclidean geometry coincides with the study of suitable surfaces of constant negative curvature, in this way connecting non-Euclidean and Riemannian geometry. His idea can be illustrated as follows. Consider all points inside the unit disk in \mathbb{R}^2 . Identify each (x, y) in the unit disk with the point $(x, y, \sqrt{1 - x^2 - y^2})$ on the unit hemisphere in \mathbb{R}^3 equipped with the Riemannian metric

$$ds^{2} = \frac{dx^{2} + dy^{2} + dz^{2}}{z^{2}}$$

If we project orthogonally onto the xy-plane, then geodesics in the hemisphere project onto straight line segments in the unit disk in the xy-plane. It can be shown that the

¹The translation is mine. For a translation of the complete work, see [25].

²The following discussion is primarily based on [34], but we refer to [12] for a more detailed treatise. For an extensive bibliography on the history of non-Euclidean geometry, see [36, 33ff.].

fifth postulate does not hold in this situation. This construction is called the Beltrami-Klein model of the hyperbolic plane. We will not discuss this model further. However, in Section 4.2 we will discuss two other models of hyperbolic geometry: the upper half-plane model and the Poincaré disk model. The former can be obtained from the hemisphere in the above construction by stereographic projection from (0, 0, -1) onto the plane z = 0; the latter by stereographic projection from (0, 0, -1) onto the plane z = 1.

The study of the geometry arising from the models mentioned above is usually called hyperbolic geometry to distinguish it from spherical geometry, another form of non-Euclidean geometry. Where hyperbolic geometry violates the parallel postulate by having multiple lines through a point parallel to a given line, spherical geometry violates the parallel postulate by having no parallel lines at all.

The embedding of hyperbolic geometry in Riemannian geometry by using these models enabled the development of a theory of hyperbolic geometry. In 1882, Poincaré described the isometries of the hyperbolic plane by using the upper half-plane model. Furthermore, he stressed the importance of discrete subgroups of isometries, leading to the theory of Fuchsian groups. In the beginning of the twentieth century the notion of a smooth manifold was rigorously defined and this led to the definition of hyperbolic manifolds. In the following subsections we will treat each of these topics in more detail.

4.2 Models of hyperbolic geometry

To prove that the parallel postulate is independent of the other postulates several models of hyperbolic geometry have been constructed. In this subsection we will discuss the upper half-plane model \mathbb{H}^2 and the Poincaré disk model \mathbb{D}^2 , primarily using [6, 14]. For other models, such as the Beltrami-Klein disk model or the hyperboloid model, see [19, 37, 45]. For more details on the classification of Möbius transformations, see [6, 48].

4.2.1 Upper half-plane model

The upper half-plane is given by $\mathbb{H}^2 = \{z \in \mathbb{C} \mid \text{Im}(z) > 0\}$. The points on the Euclidean boundary $\{z \in \mathbb{C} \mid \text{Im}(z) = 0\}$ together with a point ' ∞ ' are usually called points at infinity and denoted by $\partial \mathbb{H}^2$. Topologically these points at infinity form a circle, as we will see later as well in the Poincaré disk model. Equipped with the Riemannian metric

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

the upper-half plane becomes a model for hyperbolic geometry. The lines in this model are the geodesics for this metric, namely open rays and open semicircles emanating from and orthogonal to the real axis (see Figure 1). The hyperbolic distance d(z, w) between points $z, w \in \mathbb{H}^2$ is given by

$$d(z, w) = \log \frac{|z - \bar{w}| + |z - w|}{|z - \bar{w}| - |z - w|}.$$



Figure 1: Geodesics in \mathbb{H}^2

The group of orientation preserving isometries of \mathbb{H}^2 is denoted by $\text{Isom}^+(\mathbb{H}^2)$ and each $\varphi \in \text{Isom}^+(\mathbb{H}^2)$ is of the form

$$\varphi(z) = \frac{az+b}{cz+d}$$

for $a, b, c, d \in \mathbb{R}$ with ad - bc > 0. Reversely, every map of this form is an orientation preserving isometry of \mathbb{H}^2 .

4.2.2 Poincaré disk model

If we equip the unit disk $\mathbb{D}^2 = \{z \in \mathbb{C} \mid |z| < 1\}$ with the Riemannian metric

$$ds^* = \frac{2|dz|}{1 - |z|^2}$$

we obtain the Poincaré disk model of hyperbolic geometry. In this case the points at infinity are given by the Euclidean boundary $\partial \mathbb{D}^2 = \{z \in \mathbb{C} \mid |z| = 1\}$. The lines in this model are given by the geodesics for this metric, namely the diameters of \mathbb{D}^2 and the arcs of (Euclidean) circles orthogonal to $\partial \mathbb{D}^2$ (see Figure 2). The hyperbolic distance $d^*(z, w)$ between points $z, w \in \mathbb{D}^2$ is given by

$$d^*(z, w) = \log \frac{|1 - z\bar{w}| + |z - w|}{|1 - z\bar{w}| - |z - w|}$$

The group of orientation preserving isometries of \mathbb{D}^2 is denoted by $\text{Isom}^+(\mathbb{D}^2)$ and each $\varphi \in \text{Isom}^+(\mathbb{D}^2)$ is of the form

$$\varphi(z) = \frac{az+b}{\bar{b}z+\bar{a}}$$

for $a, b \in \mathbb{C}$ with $|a|^2 - |b|^2 = 1$. Conversely, every map of this form is an orientation preserving isometry of \mathbb{D}^2 .



Figure 2: Geodesics in \mathbb{D}^2

4.2.3 Relation

The upper half-plane model and the Poincaré disk model are both conformal models of hyperbolic geometry, meaning that they preserve angles. They are even more intimately related: the map $f : \mathbb{H}^2 \to \mathbb{D}^2$ given by

$$f(z) = \frac{z-i}{iz-1}$$

maps bijectively \mathbb{H}^2 to \mathbb{D}^2 and $\partial \mathbb{H}^2$ to $\partial \mathbb{D}^2$. Furthermore, it is an isometry between (\mathbb{H}^2, d) and (\mathbb{D}^2, d^*) . From now on we will use ds for both ds, ds^* and d for both d, d^* , since usually the context will make clear which is meant.

4.3 Classification of Möbius transformations

We start with a note regarding terminology. Isometries of the hyperbolic plane, as introduced in the previous subsection, are examples of Möbius transformations. Möbius transformations are defined in [27] as linear fractional transformations on the extended complex plane, i.e., maps of the form

$$z \mapsto \frac{az+b}{cz+d},$$

where $a, b, c, d \in \mathbb{C}$ with $ad - bc \neq 0$. Observe that in this case the coefficients a, b, c, d can take any complex value instead of only real values as for isometries of the upper halfplane. In [6], Möbius transformations are defined as compositions of inversions in spheres or half-planes, which is equivalent with the definition in [27]. On the other hand, the definition of Möbius transformations in [48] coincides with our definition of orientation preserving isometries of \mathbb{H}^2 . We will follow the definitions of [6, 27]. In that case, Möbius transformations are usually classified into four types: elliptic, parabolic, hyperbolic and loxodromic transformations, where hyperbolic transformations are a special case of loxodromic transformations. We are only interested in the first three types, so we will explain these in further detail below.

Every orientation preserving isometry of the hyperbolic plane is either elliptic, parabolic or hyperbolic. We can distinguish between these three by the number and the location of their fixed points. Since \mathbb{H}^2 and \mathbb{D}^2 are isometric, it suffices to only look at transformations in Isom⁺(\mathbb{D}^2). If $\gamma \in$ Isom⁺(\mathbb{D}^2) has

- one fixed point in \mathbb{D}^2 , then it is called elliptic (or a rotation);
- one fixed point in $\partial \mathbb{D}^2$, then it is called parabolic (or a dilation);
- two fixed points in $\partial \mathbb{D}^2$, then it is called hyperbolic (or a translation).

In the last case, the geodesic connecting the two fixed points is called the axis X_{γ} of γ . We have that $d(x, \gamma(x))$ is constant for all $x \in X_{\gamma}$ and this constant is called the translation length of γ , denoted by $l(\gamma)$.

To a transformation

$$\gamma(z) = \frac{az+b}{\bar{b}z+\bar{a}}$$

we can associate an equivalence class of matrices

$$\begin{bmatrix} a & b \\ \bar{b} & \bar{a} \end{bmatrix},$$

where A is equivalent to B if and only if $A = \pm B$. Working with equivalence classes is necessary here, since we could write $\gamma(z) = (-az - b)/(-\bar{b}z - \bar{a})$ as well. We will use the same notation for the transformation itself and the corresponding matrix. Composition in the group Isom⁺(\mathbb{D}^2) is then given by matrix multiplication. Furthermore, it can be seen that the classification of elements of Isom⁺(\mathbb{D}^2) into elliptic, parabolic and hyperbolic transformations corresponds to $|\operatorname{tr}(\gamma)| < 2$, $|\operatorname{tr}(\gamma)| = 2$ and $|\operatorname{tr}(\gamma)| > 2$ respectively. Here $|\operatorname{tr}(\gamma)|$ denotes the absolute value of the trace of (a matrix corresponding to) γ , which is well defined, because γ determines its corresponding matrix up to sign. An explicit formula for the translation length is given by

$$\cosh\left(\frac{l(\gamma)}{2}\right) = \frac{1}{2}|\operatorname{tr}(\gamma)|.$$

4.4 Trigonometry

In this subsection we will briefly state some results about hyperbolic trigonometry, using primarily [6, 29]. These trigonometric formulas hold in both \mathbb{H}^2 and \mathbb{D}^2 , so we will refer to either model by the hyperbolic plane.

Since in both models the first four postulates of Euclid hold, in particular there exists for distinct z, w in the hyperbolic plane a unique geodesic segment [z, w] joining z to w. For distinct, non-collinear z_1, z_2, z_3 in the hyperbolic plane, the hyperbolic triangle with vertices z_1, z_2, z_3 is $[z_1, z_2, z_3] = [z_1, z_2] \cup [z_2, z_3] \cup [z_3, z_1]$. More generally, $[z_1, z_2, \ldots, z_n] = [z_1, z_2] \cup \ldots \cup [z_{n-1}, z_n] \cup [z_n, z_1]$ denotes the hyperbolic *n*-gon with vertices z_1, \ldots, z_n . Let [A, B, C] be a hyperbolic triangle with angles α, β, γ at A, B, C respectively and let a, b, c be the length of the opposite edges (Figure 3). First assume that $\gamma = \frac{\pi}{2}$. In this case the hypotenuse c is given by

$\cosh c = \cosh a \cosh b,$

which is called the hyperbolic Pythagorean Theorem. Equations for the angles in terms



Figure 3: Hyperbolic triangle

of two of the sides are given by

 $\sin \alpha = \frac{\sinh a}{\sinh c},$ $\cos \alpha = \frac{\tanh b}{\tanh c},$ $\tan \alpha = \frac{\tanh a}{\sinh b}.$

Now, let $\gamma \in [0, \pi)$ be arbitrary. The hyperbolic sine rule is given by

$$\frac{\sinh a}{\sin \alpha} = \frac{\sinh b}{\sin \beta} = \frac{\sinh c}{\sin \gamma}$$

and it is the analogue of the Euclidean sine rule. The first hyperbolic cosine rule is given by

 $\cosh c = \cosh a \cosh b - \sinh a \sinh b \cos \gamma$

and given the form of the hyperbolic Pythagorean Theorem, it is the analogue of the Euclidean cosine rule. The second hyperbolic cosine rule is given by

$$\cosh c = \frac{\cos \alpha \cos \beta - \cos \gamma}{\sin \alpha \sin \beta}$$

and it has no analogue in Euclidean geometry. It implies that hyperbolic triangles with identical angles are isometric, which is not true in Euclidean geometry due to scaling. The hyperbolic area of T = T(A, B, C) is given by

$$\operatorname{area}(T) = \pi - \alpha - \beta - \gamma.$$

In particular, $\operatorname{area}(T) \leq \pi$ with identity if and only if T is an ideal triangle, i.e. A, B, C are points at infinity. Similarly, if an *n*-gon P has interior angles $\alpha_1, \ldots, \alpha_n$, then its area is given by

$$\operatorname{area}(P) = (n-2)\pi - \sum_{i=1}^{n} \alpha_i.$$

4.5 Trigonometric optimization problems

In our results we will use several isoperimetric-like inequalities. The following inequality gives a lower bound for the radius of the circumscribed circle of a polygon with a given area.

Lemma 4.1. Let P be a convex hyperbolic n-gon with area A. If P has a circumscribed circle C with radius R and center $O \in P \setminus \partial P$, then $R \ge R(A)$, where

$$\cosh R(A) = \cot\left(\frac{\pi}{n}\right)\cot\left(\frac{(n-2)\pi - A}{2n}\right),$$

with equality if and only if P is regular.

A proof is given in [35]. As a direct consequence we find a lower bound for the furthest vertex in terms of the area of the polygon.

Corollary 4.2. Let P be a convex hyperbolic n-gon with area A containing O. The distance R between O and the vertex of P furthest from O is at least R(A).

Proof. Construct $P' \supseteq P$ such that the vertices of P' lie on the circle C with center O and radius R, for example by letting the vertices of P' be the intersection points of the rays from O to the vertices of P with C (see Figure 4). Since $\operatorname{area}(P') \ge A$, we have by Lemma 4.1 that $R \ge R(\operatorname{area}(P')) \ge R(A)$.

The next inequality gives an upper bound for the area of a polygon given the radius of the circumscribed circle, so in a sense it is the dual of the first statement.

Lemma 4.3. Let P be a convex n-gon for $n \ge 3$ with all vertices on a circle with radius R. Then the area of P attains its maximal value A(R) if and only if P is regular and in this case

$$\cosh R = \cot\left(\frac{\pi}{n}\right)\cot\left(\frac{(n-2)\pi - A(R)}{2n}\right).$$

Proof. For the proof we use the same approach as [35] does for Lemma 4.1.

Consider n = 3. Divide P into three pairs of right-angled triangles with angles θ_i at the center of the circumscribed circle, angles α_i at the vertices and right angles at the edges of P (see Figure 5). By the second hyperbolic cosine rule

$$\cosh R = \cot \theta_i \cot \alpha_i$$

for i = 1, 2, 3. Furthermore, $\sum_{i=1}^{3} \theta_i = \pi$ and $A = \pi - 2 \sum_{i=1}^{3} \alpha_i$. Therefore, maximizing A reduces to minimizing

$$f(\theta_1, \theta_2, \theta_3) = \sum_{i=1}^3 \operatorname{arccot}(\cosh R \tan \theta_i)$$
(1)



Figure 4: Construction of P' from P

subject to the constraint $\sum_{i=1}^{3} \theta_i = \pi, 0 \leq \theta_i < \pi$, i.e., minimizing (1) over the triangle in \mathbb{R}^3 with vertices $(\pi, 0, 0), (0, \pi, 0), (0, 0, \pi)$. Parametrize this triangle as follows

$$\theta_1 = s + t, \theta_2 = s - t, \theta_3 = \pi - 2s$$

for $0 < s < \pi/2$ and $|t| \le s$. By (1), we can view f as a function of s and t. First we fix s and minimize over t. We have

$$\begin{split} \frac{\partial}{\partial t} f(\theta_1(s,t),\theta_2(s,t),\theta_3(s,t)) &= \sum_{i=1}^3 \frac{-\sec^2 \theta_i}{1+\cosh^2 R \tan^2 \theta_i} \frac{\partial \theta_i}{\partial t}, \\ &= \frac{\sec^2 \theta_2}{1+\cosh^2 R \tan^2 \theta_2} - \frac{\sec^2 \theta_1}{1+\cosh^2 R \tan \theta_1}, \\ &= \frac{1}{1+(\cosh^2 R - 1)\sin^2 \theta_2} - \frac{1}{1+(\cosh^2 R - 1)\sin^2 \theta_1}. \end{split}$$

Therefore, a minimum is obtained if and only if $\theta_1 = \theta_2$, i.e., if and only if t = 0. In a similar way we minimize over s.

$$\begin{aligned} \frac{\partial}{\partial s} f(\theta_1(s,t), \theta_2(s,t), \theta_3(s,t)) &= \sum_{i=1}^3 \frac{-\sec^2 \theta_i}{1 + \cosh^2 R \tan^2 \theta_i} \frac{\partial \theta_i}{\partial s}, \\ &= \frac{2}{1 + (\cosh^2 R - 1) \sin^2 \theta_3} - \frac{2}{1 + (\cosh^2 R - 1) \sin^2 \theta_1}, \end{aligned}$$



Figure 5: Division of P into three pairs of right-angled triangles

and it follows that a minimum is obtained for $\theta_1 = \theta_3$. Therefore, the area of P obtains its maximal value A(R) if and only if $\theta_1 = \theta_2 = \theta_3 = \pi/3$, i.e., if and only if P is a regular triangle. In this case,

$$\mathbf{SO}$$

$$\alpha_1 = \alpha_2 = \alpha_3 = \frac{\pi - A(R)}{6},$$

$$\cosh(R) = \cot \theta_i \cot \alpha_i = \cot \left(\frac{\pi}{3}\right) \cot \left(\frac{\pi - A(R)}{6}\right).$$

For arbitrary $n \ge 3$, the proof that maximal area is obtained for a regular polygon is the same but with more parameters. In this case $\theta_i = \pi/n$ and

$$A(R) = (n-2)\pi - 2n\alpha_i,$$

so the area A(R) of the regular polygon is given by

$$\cosh(R) = \cot \theta_i \cot \alpha_i = \cot \left(\frac{\pi}{n}\right) \cot \left(\frac{(n-2)\pi - A(R)}{2n}\right).$$

We can use this lemma to prove an upper bound for the area of a triangle in terms of the radius of its circumscribed circle.

Corollary 4.4. Let T be a hyperbolic triangle with a circumscribed disk of radius r. Then

$$\operatorname{area}(T) \le \pi - 6 \operatorname{arccot}(\sqrt{3} \cosh(r)).$$

Proof. By Lemma 4.3, we have that $\operatorname{area}(T) \leq A(r)$ for A(r) satisfying

$$\cosh r = \cot\left(\frac{\pi}{3}\right)\cot\left(\frac{\pi - A(r)}{6}\right).$$

Then

$$A(r) = \pi - 6 \operatorname{arccot}(\sqrt{3} \cosh(r)),$$

which finishes the proof.

4.6 Circle packings

In this section we will discuss the packing density of circle packings in the hyperbolic plane. This section will only be used to improve the lower bound for the number of sheets in Section 6. To illustrate packing density we will first consider circle packings in the Euclidean plane. For an introduction to circle packings, see [43]. We base the discussion of circle packings in the hyperbolic plane on [46].

A circle packing is a set of disks with mutually disjoint interiors. We will only consider circle packings where the disk are all congruent. The packing density of a circle packing in the Euclidean plane can be defined in the following way. Let P be a circle packing in the Euclidean plane. Fix a point O in the plane and consider a circle C(R) with center Oand radius R. Let $\operatorname{area}_P(C(R))$ be the sum of the areas of all circles of the circle packing which lie entirely in C(R). Let $\operatorname{area}(C(R))$ be the area of C(R). Then the packing density D(P) of the circle packing P is defined as

$$D(P) = \lim_{R \to \infty} \frac{\operatorname{area}_P(C(R))}{\operatorname{area}(C(R))}.$$

It makes no difference if we take $\operatorname{area}(P \cap C(R))$ instead of $\operatorname{area}_P(C(R))$, i.e., if we also consider circles of P that lie partially in C(R). Namely, if the circles of P have radius r, then all circles partially in C(R) lie in the strip $C(R+2r) \setminus C(R-2r)$. Therefore, their area is at most the area of this strip, which is

area
$$(C(R+2r))$$
 - area $(C(R-2r)) = \pi (R+2r)^2 - \pi (R-2r)^2 = 8\pi r R.$

It follows that

$$0 \le \lim_{R \to \infty} \frac{\operatorname{area}(P \cap C(R)) - \operatorname{area}_P(C(R))}{\operatorname{area}(C(R))} \le \lim_{R \to \infty} \frac{8\pi rR}{\pi R^2} = 0,$$

so the contribution of these circles to the density is negligible in the limit. It is well known that the hexagonal packing H (see Figure 6) is the circle packing with congruent circles in the Euclidean plane with the maximum packing density, namely

$$D(H) = \frac{\pi}{2\sqrt{3}}.$$



Figure 6: Hexagonal packing in the Euclidean plane

In the hyperbolic plane we cannot use this definition of packing density, as the contribution from circles partially in C(R) is not negligible anymore. Instead, packing density of a circle packing in the hyperbolic plane is defined in [46] by triangulating the hyperbolic plane and taking the mean density over all triangles. Then the density within each triangle is computed using the area of all (partial) circles within the triangle. In this way, the maximum packing density for a circle packing with circles of radius r is given by

$$D(a) = \frac{3\csc\frac{\pi}{a} - 6}{a - 6},$$

where a is defined by

$$\csc\frac{\pi}{a} = 2\cosh r.$$

Note that the maximum packing density increases monotonically as a function of r. For $r \to \infty$ we have $a \to \infty$ and $D(a) \to \frac{3}{\pi}$, which is the maximum packing density for a circle packing in the hyperbolic plane independent of the radius of the circles.

4.7 Fuchsian groups

In this section we will first define Fuchsian groups and then look at fundamental domains. Proofs of propositions will be omitted. For more details we refer to [6, 29, 48].

Recall that a subset of a topological space is called discrete if the subspace topology on this set is the discrete topology, i.e., the topology where every subset is open and closed. The identification of elements of $\text{Isom}^+(\mathbb{D}^2)$ with matrices induces the structure of a topological space on $\text{Isom}^+(\mathbb{D}^2)$. We can now give the definition of a Fuchsian group.

Definition 4.5. A discrete subgroup of $\text{Isom}^+(\mathbb{D}^2)$ is called a Fuchsian group.

Again, since \mathbb{H}^2 and \mathbb{D}^2 are isometric, we will only consider subgroups of Isom⁺(\mathbb{D}^2). For a Fuchsian group Γ and a point $x \in \mathbb{D}^2$, the orbit of x under Γ is defined as $\Gamma(x) = \{\gamma(x) \mid \gamma \in \Gamma\} \subset \mathbb{D}^2$.

Proposition 4.6. Let Γ be a subgroup of $\text{Isom}^+(\mathbb{D}^2)$. The following statements are equivalent:

- 1. Γ is a Fuchsian group,
- 2. For all $x \in \mathbb{D}^2$, $\Gamma(x)$ is a discrete subset of \mathbb{D}^2 .
- 3. For all $x \in \mathbb{D}^2$, there exists a neighbourhood N, such that $\gamma(N) \cap N \neq \emptyset$ for only finitely many $\gamma \in \Gamma$.

If Γ satisfies the third statement, we usually say that Γ acts properly discontinuously on \mathbb{D}^2 , even though definitions of properly discontinuous may vary in the literature. Denote the interior of a subset F of a topological space by \mathring{F} . We can then define a fundamental region for the action of a Fuchsian group Γ on \mathbb{D}^2 .

Definition 4.7. Given a Fuchsian group Γ , a fundamental domain F for Γ is a closed subset of \mathbb{D}^2 such that

1.
$$\bigcup_{\gamma \in \Gamma} \gamma(F) = \mathbb{D}^2,$$

2. For all $\gamma_1, \gamma_2 \in \Gamma$ we have: if $\gamma_1 \neq \gamma_2$, then $\gamma_1(\mathring{F}) \cap \gamma_2(\mathring{F}) = \emptyset$.

A priori, we do not know that there actually exists a fundamental region for a given Fuchsian group Γ , but later on we will explicitly construct a fundamental region called the Dirichlet region. Different fundamental domains can look very different, but the following proposition states that their area is always the same.

Proposition 4.8. Let Γ be a Fuchsian group with fundamental domains F, F' such that $\operatorname{area}(\partial F) = \operatorname{area}(\partial F') = 0$ and $\operatorname{area}(F) < \infty$. Then

$$\operatorname{area}(F) = \operatorname{area}(F').$$

Naturally, a subgroup of a Fuchsian group is a Fuchsian group as well.

Proposition 4.9. Let Γ be a Fuchsian group with fundamental domains F such that $\operatorname{area}(\partial F) = 0$ and let $\Gamma' < \Gamma$ be a subgroup of Γ of index k with fundamental domain F'. Then

$$\operatorname{area}(F') = k \operatorname{area}(F).$$

It can be shown that for every Fuchsian group Γ there exists a point $p \in \mathbb{D}^2$ such that $\gamma(p) \neq p$ for all $\gamma \in \Gamma \setminus \{\text{Id}\}$, i.e. there exists a point that is not fixed by any non-trivial element of Γ . We can now define the Dirichlet region.

Definition 4.10. Let Γ be a Fuchsian group and let p be a point that is not fixed by any non-trivial element of Γ . Define the Dirichlet region $D_p(\Gamma)$ of p with respect to Γ as

$$D_p(\Gamma) = \{ x \in \mathbb{D}^2 \mid d(x, p) \le d(x, \gamma(p)) \text{ for all } \gamma \in \Gamma \}.$$

Intuitively, $D_p(\Gamma)$ can be seen as the collection of points that are closer to p than to the other elements of $\Gamma(p)$. Indeed, $D_p(\Gamma)$ is a fundamental domain.

Proposition 4.11. Let Γ be a Fuchsian group and let p be a point that is not fixed by any non-trivial element of Γ . The Dirichlet region $D_p(\Gamma)$ is a fundamental domain for Γ . If $\operatorname{area}(D_p(\Gamma)) < \infty$, then $D_p(\Gamma)$ is a convex hyperbolic polygon with finitely many sides.

Suppose that in the situation above there exists a side s of $D_p(\Gamma)$ and $\gamma \in \Gamma$, such that $\gamma(s)$ is also a side of $D_p(\Gamma)$. Then we call such a γ a side pairing transformation. Indeed, sides are paired, since γ^{-1} maps the side $\gamma(s)$ back to the side s. In fact, for a Dirichlet region $D_p(\Gamma)$ we can find such a side pairing transformation for every side s of $D_p(\Gamma)$. Namely, every side is a piece of the perpendicular bisector of the segment $[p, \gamma(p)]$ for some $\gamma \in \Gamma \setminus \{\mathrm{Id}\}$ and it can be shown that γ^{-1} maps s to another side of $D_p(\Gamma)$. Hence, to any Dirichlet region we can associate a set of side pairing transformations.

4.8 Hyperbolic surfaces

A Fuchsian group Γ naturally acts on \mathbb{D}^2 , so we can form the quotient space \mathbb{D}^2/Γ . We saw in the previous section that with the Dirichlet region of Γ we can associate a set of side pairing transformations. These side pairing transformations can be seen as 'glueing' the Dirichlet region along paired sides. In this way we can see \mathbb{D}^2/Γ as a surface which locally looks like a part of the hyperbolic plane. Such a surface will be called a hyperbolic surface. In this section we will see that the converse holds as well: a given hyperbolic surface is isometric to a quotient \mathbb{D}^2/Γ for some Fuchsian group Γ . Another reverse construction is provided by Poincaré's Theorem: in this case from a polygon and a set of side pairing transformations the corresponding Fuchsian group is constructed, provided some conditions are satisfied. We will not discuss this in detail, see instead [29, 48].

First we give the definition of hyperbolic surface.

Definition 4.12. A hyperbolic surface is a connected 2-dimensional manifold that is locally isometric to an open subset of \mathbb{D}^2 .

Again, we will only consider \mathbb{D}^2 , since \mathbb{H}^2 and \mathbb{D}^2 are isometric. Because hyperbolic surfaces are defined to be locally isometric to open subsets of the hyperbolic plane, they

have an induced Riemannian metric of constant Gaussian curvature -1. As such, they cannot be embedded in \mathbb{R}^3 ; for the first proof of this, see [26]. However, for visualization we will still draw hyperbolic surfaces as if they were surfaces in \mathbb{R}^3 . Furthermore, we will always assume that the surface is orientable. Hyperbolic surfaces can be obtained as a quotient space under the action of a Fuchsian group.

Proposition 4.13. For every hyperbolic surface M there exists a Fuchsian group Γ acting on \mathbb{D}^2 without fixed points, such that M is isometric to \mathbb{D}^2/Γ .

Since elliptic Möbius transformations have a fixed point in \mathbb{D}^2 , a Fuchsian group as in the proposition does not contain any elliptic elements.

A compact hyperbolic surface is called closed. A Fuchsian group is called cocompact, if \mathbb{D}^2/Γ is compact. It can be shown that a cocompact Fuchsian group cannot contain any parabolic elements; see, e.g., [29].

Corollary 4.14. For every closed hyperbolic surface M there exists a Fuchsian group Γ of which all non-trivial elements are hyperbolic, such that M is isometric to \mathbb{D}^2/Γ .

A note on the literature: this section is mostly based on [44], since the main statement of this section is stated there explicitly. Works on Teichmüller spaces, such as [27, 42], often focus more on the classification of Riemann surfaces. For Riemann surfaces a conformal structure is a maximal atlas such that all transition maps are holomorphic. Buser [14] gives a proof that the classifications of conformal structures on Riemann surfaces of genus $g \ge 2$ and atlases for hyperbolic surfaces coincide. Namely, since the transition maps of a hyperbolic atlas are restrictions of Möbius transformations, a hyperbolic atlas naturally induces a conformal structure. Reversely, given a Riemann surface M of genus $g \ge 2$ there exists by the Uniformization Theorem a universal covering map $\pi : \mathbb{D}^2 \to M$. The covering transformations are conformal self-mappings of \mathbb{D}^2 , so the local inverses of π can be used as parametrizations for a hyperbolic surface. Beardon [6] evades this distinction: initially he considers Riemann surfaces, but then he introduces 'Riemann surfaces of hyperbolic type', which are *defined* to be of the form \mathbb{D}^2/Γ .

4.9 Closed geodesics on a hyperbolic surface

In the criterium for well-behaved triangulations of hyperbolic surfaces that we will discuss in Section 5, a central role is played by the systole of a surface. Therefore we will look in this section at closed geodesics, the systole and the length spectrum of hyperbolic surfaces, following primarily [36, 44]. The relevant homotopy theory can be found in an introduction on algebraic topology; see, e.g., [23].

Let $M = \mathbb{D}^2/\Gamma$ be a closed hyperbolic surface. By using the metric on \mathbb{D}^2 and the fact that Γ consists of isometries, we obtain a metric on M, so we can speak of geodesics on M. Define the circle $\mathbb{S}^1 = \mathbb{R}/\sim$, where $x \sim x + 1$ for all $x \in \mathbb{R}$. A closed curve on M is a continuous map $c : \mathbb{S}^1 \to M$. We will always assume differentiability, except maybe at a point: a curve $c : [0, 1] \to M$ such that c(0) = c(1) which is differentiable at (0, 1) is called a loop. Closed curves $c_0, c_1 : \mathbb{S}^1 \to M$ are called freely homotopic if there exists a continuous map $H : \mathbb{S}^1 \times [0, 1] \to M$ such that

$$H(x,0) = c_0(x), \quad H(x,1) = c_1(x)$$

for all $x \in \mathbb{S}^1$. Curves which are homotopic to a point are called homotopically trivial. We will always consider a closed geodesic together with a parametrization; in this way we can distinguish between a closed geodesic c and the closed geodesic c^2 obtained by traversing c twice. We now define the systole.

Definition 4.15. The length of the shortest homotopically non-trivial closed curve on a closed hyperbolic surface M is called the systole of M and denoted by syst(M).

The systole of a surface is attained as the length of some curve. Clearly, the shortest closed curves on M are simple, i.e. they have no self-intersections except at the endpoints. By the following proposition it is sufficient to consider only (simple) closed geodesics.

Proposition 4.16. Every homotopically non-trivial (simple) closed curve on a closed hyperbolic surface is freely homotopic to a unique (simple) closed geodesic, which is the shortest curve in the corresponding free homotopy class.

A proof can be found in [36, Thm. 9.6.4 & 9.6.5]. Closed geodesics are closely related to the structure of the corresponding Fuchsian group.

Proposition 4.17. Closed geodesics of a closed hyperbolic surface $M = \mathbb{D}^2/\Gamma$ are in one-to-one correspondence with conjugacy classes of elements of Γ .

For a proof, see [36, Thm. 9.6.2]. In the above correspondence, the axis X_{γ} of $\gamma \in \Gamma$ in \mathbb{D}^2 projects onto its corresponding closed geodesic c under the projection map $\mathbb{D}^2 \to \mathbb{D}^2/\Gamma$. It follows that the length of c is given by the distance $d(x, \gamma(x))$ for $x \in X_{\gamma}$, as this is the distance that γ moves $x \in X_{\gamma}$ along X_{γ} until it reaches the next point in the orbit. Therefore, the length of c is equal to the translation length $l(\gamma)$. Then finding the systole reduces to the following optimization problem:

$$\cosh\left(\frac{\operatorname{syst}(\mathbb{D}^2/\Gamma)}{2}\right) = \min_{\substack{\gamma \in \Gamma \\ \gamma \neq \operatorname{Id}}} \frac{1}{2} |\operatorname{tr}(\gamma)|.$$

The following proposition gives an upper bound for the systole; the proof is from [14].

Proposition 4.18. Let M be a closed hyperbolic surface of genus $g \ge 2$. Then

$$\operatorname{syst}(M) \le 2\log(4g-2).$$

Proof. Let c be the shortest homotopically non-trivial closed curve on M and fix $p \in c$. We have that $D_r = \{q \in M \mid d((p,q) < r)\}$ is a hyperbolic disk of radius r as long as $r < \operatorname{syst}(M)/2$. Then

$$4\pi(g-1) = \operatorname{area}(M) > \operatorname{area}(D_r) = 2\pi(\cosh(r) - 1).$$

Taking the limit $r \to \operatorname{syst}(M)/2$ we obtain

 $\cosh(\operatorname{syst}(M)/2) \le 2g - 1.$

Since $\frac{1}{2} \exp(\operatorname{syst}(M)/2) < \cosh(\operatorname{syst}(M)/2)$, we have

$$\exp(\operatorname{syst}(M)/2) < 4g - 2,$$

which proves the result.

The above upper bound is up to a constant the best possible. Namely, in [15] a family of hyperbolic surfaces $\{M(g)\}$ is constructed with genus $g \to \infty$ and

$$syst(M(g)) \ge \frac{4}{3}\log g - c$$

for some constant c. The same bound is found in [30] for a different family of surfaces. In both cases the families of surfaces are constructed by considering principal congruence subgroups of arithmetic Fuchsian groups. We will not discuss arithmeticity here, but refer to [29]. In [33] it is shown that the coefficient $\frac{4}{3}$ is the best possible for surfaces obtained in this way. Whether there are families of surfaces that yield larger coefficients, is currently not known. A universal nonzero lower bound for the systole of hyperbolic surfaces does not exist: in [7] it is shown³ that for any $\epsilon > 0$ there exists a hyperbolic surface M with syst $(M) < \epsilon$.

As a final remark, the systole is the first element of a certain sequence called the length spectrum.

Definition 4.19. The length spectrum $\mathcal{L}(M)$ of a closed hyperbolic surface M is the ascendingly ordered sequence of lengths of closed geodesics on M.

In light of the discussion before, the length spectrum of a hyperbolic surface $M = \mathbb{D}^2/\Gamma$ is equal to the ordered sequence of translation lengths of conjugacy classes of Γ . Since isometries preserve the length of every geodesic on the surface, isometric surfaces have the same length spectrum. However, the converse is not true: there exist non-isometric surfaces which are isospectral, i.e., they have the same length spectrum. In [47], upper bounds are given for the number of non-isometric, isospectral surfaces in terms of genus and systole.

4.10 Teichmüller space

Classifying the isomorphism classes of Riemann surfaces is known as the moduli problem. By the remark in Section 4.8, this is equivalent to the description of all isometry classes of hyperbolic surfaces. This problem is currently unsolved, even though there are solutions for low genera. In Section 4.10.6, we will see an example of such a solution for surfaces of signature (0, 3), i.e., of genus 0 with 3 punctures.

First we will discuss pairs of pants and cubic graphs, the building blocks and skeletons, respectively, of hyperbolic surfaces. Then we will discuss the twist parameters, extra degrees of freedom that arise when we glue pairs of pants together. These ingredients will be combined to define the Fenchel-Nielsen coordinates, a natural model for the Te-ichmüller space, which consists of equivalence classes of marked hyperbolic surfaces. We will see that isometry classes of hyperbolic surfaces have multiple representatives in the Teichmüller space and this multiplicity is described in the mapping class group. Finally we will introduce the Zieschang-Vogt-Coldewey coordinates, another useful parametrization of the Teichmüller space. We will follow [14] closely, but see also [27, 36, 42].

³In fact, they show the corresponding statement for hyperbolic *n*-dimensional manifolds. We will see in the definition of the Fenchel-Nielsen coordinates that the statement for hyperbolic surfaces is trivial.

4.10.1 Pairs of pants

Let *H* be a right-angled hyperbolic hexagon with consecutive sides $b_1, s_1, b_2, s_2, b_3, s_3$ (see Figure 7).



Figure 7: Right-angled hyperbolic hexagon

Let H' be a copy of H with sides b'_i, s'_i . We will glue H and H' together along the seams s_1, s_2, s_3 and s'_1, s'_2, s'_3 (see Figure 8). Parametrize the sides with constant speed to obtain



Figure 8: Construction of pair of pants

 $t \mapsto s_i(t), t \mapsto s'_i(t), t \in [0, 1]$. Let $Y = H \sqcup H' / \sim$ be the disjoint union of H and H'modulo the glueing condition \sim , where $p \sim q$ for $p \in H$ and $q \in H'$ if and only if there exists $i \in \{1, 2, 3\}$ and $t \in [0, 1]$ such that $p = s_i(t)$ and $q = s'_i(t)$. The resulting Y will be called a pair of pants. It is a hyperbolic surface with boundary⁴ homeomorphic to a thrice-punctured sphere. There are three boundary curves, namely $b_i \cup b'_i$, i = 1, 2, 3. These boundary curves are closed geodesics, since all angles in the hexagons are right angles.

Now, let Y be a pair of pants. For every pair of boundary geodesics of Y there exists a unique simple common perpendicular. These perpendiculars, i.e., the seams, are mutually disjoint and divide the boundary geodesics in two arcs of the same length. Therefore, by cutting Y open along the seams we obtain two isometric right-angled hexagons.

By the construction above, we see that for any triple l_1, l_2, l_3 of positive real numbers, there exists a pair of pants with boundary geodesics of lengths l_1, l_2, l_3 . By the decomposition of a pair of pants into hexagons and the fact that the lengths of b_1, b_2, b_3 determine the hexagon up to isometry, we see that such a pair of pants is unique up to isometry.

4.10.2 Cubic graphs

Recall that a graph G = (V, E) consists of a set of vertices V = V(G) and edges E = E(G). Denote the number of vertices and edges of G by v(G), e(G) respectively. We will assume that the graph is undirected and loops and double edges are allowed. A graph is connected if for all $v, w \in G$, there exists a sequence $v = v_1, v_2, \ldots, v_k = w$ such that $(v_i, v_{i+1}) \in E$ for all $i = 1, \ldots, k - 1$.

For our purpose it is useful to interpret each edge as two half-edges. A cubic graph is a graph where every vertex has three emanating half-edges. Therefore, a cubic graph Gcontains 3v(G) half-edges, so 3v(G) = 2e(G). This means that the number of vertices of a cubic graph is always even, say v(G) = 2g - 2. Denote the vertices of G by v_1, \ldots, v_{2g-2} and its edges by e_1, \ldots, e_{3g-3} . Denote the half-edges emanating from v_i by $e_{i\alpha}$, $\alpha = 1, 2, 3$. Each edge $e_k = (v_i, v_j)$ is interpreted as the union of two half-edges $e_k = e_{i\alpha} \cup e_{j\beta}$ for some $\alpha, \beta \in \{1, 2, 3\}$. Then the list

$$e_k = e_{i\alpha} \cup e_{j\beta}, \quad k = 1, \dots, 3g - 3$$

completely describes the graph G.

Example 4.20. In Figure 9 we see an example of a cubic graph. It is completely described by the following list:

 $e_{1} = e_{11} \cup e_{21},$ $e_{2} = e_{12} \cup e_{22},$ $e_{3} = e_{13} \cup e_{31},$ $e_{4} = e_{23} \cup e_{32},$ $e_{5} = e_{33} \cup e_{41},$ $e_{6} = e_{42} \cup e_{43}.$

 $^{^4{\}rm Technically},$ we did not define hyperbolic surfaces with boundary, but the definition is similar to manifolds with boundary.



Figure 9: Example of a cubic graph

4.10.3 Twist parameters

Let Y, Y' be pairs of pants with boundary geodesics $b_i : \mathbb{S}^1 \to Y, b'_i : \mathbb{S}^1 \to Y'$ parametrized with constant speed. Assume that $\ell(b_1) = \ell(b'_1)$, where $\ell(b_1)$ denotes the length of b_1 . We will glue Y and Y' together along the boundaries b_1 and b'_1 (see Figure 10). For any $a \in \mathbb{R}$, let $X_a = Y \sqcup Y' / \stackrel{a}{\sim}$ be the disjoint union of Y and Y' modulo the glueing condition $\stackrel{a}{\sim}$, where $p \stackrel{a}{\sim} q$ for $p \in Y$ and $q \in Y'$ if and only if there exists $t \in \mathbb{S}^1$ such that $p = b_1(t)$ and $q = b'_1(a - t)$. Observe that we took \mathbb{S}^1 to be a quotient space with base space \mathbb{R} , so the expression a - t makes sense. Furthermore, we use $b'_1(a - t)$ instead of $b'_1(a + t)$ to preserve the orientation. The resulting X_a is a hyperbolic surface with boundary homeomorphic to a sphere with four punctures. Of course we can continue this glueing procedure with X_a and another pair of pants and we will do so in the next part. The parameter a is called a twist parameter and can be seen as the amount of twisting used in the glueing of Y and Y'.



Figure 10: Glueing of pairs of pants with twist parameter $\frac{1}{4}$

4.10.4 Fenchel-Nielsen coordinates

Now we have all the ingredients to construct hyperbolic surfaces. Intuitively, the construction is as follows: suppose we are given a connected cubic graph. Each vertex corresponds to a pair of pants. Two pairs of pants are glued together along a pair of their boundary geodesics if and only if there is an edge between the corresponding vertices. A loop in the graph means that two boundary geodesics of one and the same pair of pants are glued together. The isometry class of the resulting surface will depend on the lengths of the boundary geodesics and the twist parameters.

Example 4.21. In Figure 11 we see an example of the construction of a hyperbolic surface, where the underlying structure is given by the cubic graph of Figure 9. In this case vertex v_i corresponds to pair of pants Y_i . Indeed, Y_1 and Y_2 are glued together along two of their boundary geodesics as there is a double edge between them in the cubic graph. In the same way the other edges are represented by glueing.



Figure 11: Construction of hyperbolic surfaces

Now we will formally describe this procedure. Let $g \ge 2$ be an integer. Let G be a connected cubic graph with v(G) = 2g - 2, which can be completely described by

$$e_k = e_{i\alpha} \cup e_{j\beta}, \quad k = 1, \dots, 3g - 3.$$

Choose $l_1, \ldots, l_{3g-3} \in \mathbb{R}_+$ and $a_1, \ldots, a_{3g-3} \in \mathbb{R}$. Associate to each vertex v_i with halfedges $e_{i\alpha}, \alpha = 1, 2, 3$ a pair of pants Y_i with boundary geodesics $b_{i\alpha}, \alpha = 1, 2, 3$ such that for the pairs in the list above

$$l_k = \ell(b_{i\alpha}) = \ell(b_{j\beta}), \quad k = 1, \dots, 3g - 3$$

Let

$$M = \bigsqcup_{k=1}^{3g-3} Y_k \Big/ \bigsqcup_{k=1}^{3g-3} \overset{a_k}{\sim}$$

be the disjoint union of the Y_k modulo all the glueing conditions $\stackrel{a_k}{\sim}$, where each $\stackrel{a_k}{\sim}$ is understood to apply to the corresponding Y_i and Y_j from the list above. In this way we obtain a hyperbolic surface M of genus g. We call the sequence $(l_1, \ldots, l_{3g-3}, a_1, \ldots, a_{3g-3})$ the Fenchel-Nielsen coordinates of the closed hyperbolic surface M. The following theorem shows the usefulness of the above construction. **Theorem 4.22.** Let G be a fixed connected cubic graph with v(G) = 2g - 2. Then every closed hyperbolic surface of genus g can be obtained by the construction above with underlying graph G.

We immediately see that closed hyperbolic surfaces can be constructed in multiple ways using the procedure above, for example by constructing one hyperbolic surface from different graphs. In Section 4.10.6 we will discuss this further.

4.10.5 Teichmüller space

The Fenchel-Nielsen coordinates provide a natural model for the Teichmüller space $\mathcal{T}_{g,n}$. To formally define the Teichmüller space we need to introduce marked hyperbolic surfaces. For each signature (g, n), where g is the genus and n the number of punctures, define a fixed closed hyperbolic surface $B_{g,n}$ of genus g with n punctures such that its boundary components are smooth closed curves.

Definition 4.23. A marked hyperbolic surface (M, φ) of signature (g, n) consists of a closed hyperbolic surface M of signature (g, n) and a homeomorphism $\varphi : B_{g,n} \to M$, which is called the marking homeomorphism.

Marked hyperbolic surfaces are considered to be 'the same' if they are marking equivalent.

Definition 4.24. Two marked hyperbolic surfaces $(M, \varphi), (M', \varphi')$ are called marking equivalent if there exists an isometry $f: M \to M'$ such that φ' and $f \circ \varphi$ are isotopic.

Recall that homeomorphisms $f_0, f_1 : X \to Y$ of topological spaces are isotopic if there exists a continuous map $J : [0, 1] \times X \to Y$ such that

$$J(0, x) = f_0(x), \quad J(1, x) = f_1(x)$$

for all $x \in X$ and $J(t, \cdot) : X \to Y$ is a homeomorphism for all $t \in [0, 1]$.

Definition 4.25. The Teichmüller space $\mathcal{T}_{g,n}$ of signature (g,n) is the set of all marking equivalence classes of marked hyperbolic surfaces. We write \mathcal{T}_q instead of $\mathcal{T}_{q,0}$.

To see that the Fenchel-Nielsen coordinates are a model of the Teichmüller space, for a given connected cubic graph G, set B_G equal to the hyperbolic surface with underlying graph G and Fenchel-Nielsen coordinates $l_k = 1, a_k = 0$ for $k = 1, \ldots, 3g - 3$. Then we can construct the marking homeomorphism from B_G to the hyperbolic surface with arbitrary Fenchel-Nielsen coordinates, which consists of stretching (to make the l_k larger) and twisting (to make the a_k larger). We will not elaborate on this; see [14] for more details. The set of marked hyperbolic surface with such a marking homeomorphism, with base surface B_G and with underlying graph G, is called \mathcal{T}_G .

Theorem 4.26. Let G be a fixed connected cubic graph with v(G) = 2g - 2. Then for every marked hyperbolic surface (M, φ) there exists a unique $(M', \varphi') \in \mathcal{T}_G$, which is marking equivalent to (M, φ) .

It follows that we indeed have a bijection between \mathcal{T}_g and the Fenchel-Nielsen coordinates for a fixed connected cubic graph G.

4.10.6 Mapping class group

The Teichmüller space does not solve the moduli problem. Indeed, an isometry class of hyperbolic surfaces of signature (g, n) has multiple representatives in $\mathcal{T}_{g,n}$.

For example, consider the Teichmüller space of signature (0,3), i.e. pairs of pants. Two marked pairs of pants $(Y, \varphi), (Y', \varphi')$ are marking equivalent if and only if $\varphi \circ \varphi'$ fixes each of the boundary components. Therefore, a pair of pants with boundary lengths (1, 2, 3)for the labeled boundary geodesics b_1, b_2, b_3 is not marking equivalent to a pair of pants with boundary lengths (2, 1, 3), even though they are isometric. It is clear that in this case the moduli space of isometry classes of hyperbolic surfaces is given by $\mathcal{T}_{0,3}/S_3$, where S_3 is the permutation group permuting the labels of the boundary geodesics.

More generally, consider the mapping class group.

Definition 4.27. For a fixed signature (g, n) and base surface $B_{g,n}$ the mapping class group $\mathfrak{M}_{g,n}$ is the group of all isotopy classes of homeomorphisms $B_{g,n} \to B_{g,n}$.

Each such homeomorphism f induces an action $m(f) : \mathcal{T}_{g,n} \to \mathcal{T}_{g,n}$ on marked hyperbolic surfaces by the following rule:

$$m(f)(M,\varphi) = (M,\varphi \circ f).$$

Definition 4.28. The Teichmüller modular group $\mathcal{M}_{q,n}$ is the group of transformations

$$\mathcal{M}_{g,n} = \{ m(f) \mid f \in \mathfrak{M}_{g,n} \}.$$

Then we can see that the Teichmüller modular group plays the role that S_3 plays in the case of pairs of pants.

Proposition 4.29. Marked hyperbolic surfaces $(M, \varphi), (M', \varphi') \in \mathcal{T}_{g,n}$ are isometric if and only if there exists $\mu \in \mathcal{M}_{g,n}$ such that $\mu(M, \varphi) = (M', \varphi')$.

It follows that the moduli space $\mathcal{R}_{g,n}$ of isometry classes of hyperbolic surfaces of signature (g, n) is

$$\mathcal{R}_{g,n} = \mathcal{T}_{g,n} / \mathcal{M}_{g,n},$$

so describing $\mathcal{R}_{g,n}$ is equivalent to finding a fundamental domain for the action of $\mathcal{M}_{g,n}$ on $\mathcal{T}_{g,n}$. For some signatures this has succeeded, but in general this problem is unsolved.

4.10.7 Zieschang-Vogt-Coldewey coordinates

There are several other systems of coordinates for \mathcal{T}_g besides Fenchel-Nielsen coordinates, which use the relation between closed hyperbolic surfaces and Fuchsian groups (Bers' coordinates) or hyperbolic polygons (Zieschang-Vogt-Coldewey coordinates). We will discuss the latter in more detail.

Fix $g \geq 2$. Consider a convex geodesic hyperbolic 4g-gon with consecutive side lengths $b_1, b_2, b'_1, b'_2, \ldots, b_{2g}, b'_{2g}$ and consecutive interior angles $\zeta_1, \zeta_2, \zeta'_1, \zeta'_2, \ldots, \zeta_{2g}, \zeta'_{2g}$, where ζ_i is the angle following b_i and ζ'_i the angle following b'_i . See Figure 12 for an example with g = 3.



Figure 12: Geodesic polygon for g = 3

Definition 4.30. A polygon P as above is called a canonical polygon if

• $b_i = b'_i$ for all i = 1, ..., 2g,

•
$$\zeta_1 + \ldots + \zeta_{2g} + \zeta'_1 + \ldots + \zeta'_{2g} = 2\pi$$

If in addition

• $\zeta_1 + \zeta_2 = \zeta'_1 + \zeta_2 = \pi$,

then P is called a normal canonical polygon.

We can form the set of equivalence classes of normal canonical polygons.

Definition 4.31. Two normal canonical polygons P, \tilde{P} with sides b_1, \ldots, b'_{2g} and $\tilde{b}_1, \ldots, \tilde{b}'_{2g}$ are called equivalent if there exists an isometry $f: P \to \tilde{P}$ sending b_1 to \tilde{b}_1 and b_2 to \tilde{b}_2 . The set of all equivalence classes of normal canonical 4g-gons is denoted by \mathcal{P}_q .

A canonical polygon P as above can be made into a hyperbolic surface by glueing b_i to b'_i for $i = 1, \ldots, 2g$. The vertices of P are all glued together to form a base point p. The first condition makes sure that all sides can be glued and the second condition that all angles at the base point add up to 2π . Following this procedure for an arbitrary normal canonical 4g-gon we obtain a base surface B_q for the Teichmüller space \mathcal{T}_q .

Now, let (M, φ) be a marked closed hyperbolic surface of genus g. Since each pair b_i, b'_i used in the construction of B_g corresponds to a loop β_i on B_g with base point p, we have that each $\varphi \circ \beta_i$ is a loop on M with base point $\varphi(p)$.

Theorem 4.32. For a marked closed hyperbolic surface (M, φ) of genus g with loops $\varphi \circ \beta_i, i = 1, \ldots, 2g$ there exist a homeomorphism $h : M \to M$ and loops $\alpha_1, \ldots, \alpha_{2g}$ such that

- h is isotopic to the identity map,
- $h \circ \beta_i = \alpha_i$ for $i = 1, \ldots, 2g$,
- α_1 and α_2 are closed geodesics with intersection point q,
- $\alpha_3, \ldots, \alpha_{2g}$ are geodesic loops with base point q.

Such loops α_i are called a normal canonical dissection of M.

By cutting M open along a normal canonical dissection we obtain a normal canonical 4g-gon.

Theorem 4.33. The map $\mathcal{T}_g \to \mathcal{P}_g$ which cuts open a marked closed hyperbolic surface of genus g along a normal canonical dissection to obtain a normal canonical polygon, is a bijection.

For $P \in \mathcal{P}_g$ we call the sequence $(b_3, \ldots, b_{2g}, \zeta_3, \ldots, \zeta_{2g}, \zeta'_3, \ldots, \zeta'_{2g})$ the Zieschang-Vogt-Coldewey coordinates of P. Note that $b_1, b_2, \zeta_1, \zeta_2, \zeta'_1, \zeta'_2$ are not included. By looking at equivalence classes, there is no dependence on these parameters any more. Moreover, we see that we have 6g - 6 parameters left, which is no surprise, since there are also 6g - 6Fenchel-Nielsen parameters.

5 Triangulations

In this section we will discuss triangulations, the Delaunay property and triangulations in quotient spaces and covering spaces. For more details on triangulations in general we refer to [24, 32]. The remainder of this section is based on [9, 10, 17], where proofs of the propositions can be found as well.

Since we only consider two-dimensional surfaces, we do not have to introduce simplicial complexes to define triangulations. Let $\Delta = \{(x, y) \in \mathbb{R}^2 \mid x \ge 0, y \ge 0, x + y \le 1\}$ be a triangle in \mathbb{R}^2 . In a natural way we can speak of the vertices and edges of Δ .

Definition 5.1. Let M be a two-dimensional topological manifold. A triangulation $T = \{(M_i, \varphi_i)\}_{i \in \mathcal{I}}$ of M consists of subsets $M_i \subset M$ and homeomorphisms $\varphi_i : M_i \to \Delta$ such that

- $\bigcup_{i \in \mathcal{T}} M_i = M$,
- for all $i, j \in \mathcal{I}$ with $i \neq j$, $\varphi_i(M_i) \cap \varphi_j(M_j)$ is either the empty set or a vertex or an edge of Δ .

The vertices, edges and faces of T are the images of the vertices of Δ , edges of Δ and entire domain under $\varphi_i^{-1}, i \in \mathcal{I}$. The sets of vertices, edges and faces of T are denoted by V(T), E(T) and F(T) respectively.

Given a point set S on a two-dimensional manifold M, a triangulation of S in M is a triangulation T of M such that V(T) = S. If M is also a metric space and S is a bounded set, then for every triangulation $\{(M_i, \varphi_i)\}_{i \in \mathcal{I}}$ of S in M we have that $\bigcup_{i \in \mathcal{I}} M_i$ is compact. Therefore, strictly speaking there do not exist triangulations of bounded point sets in \mathbb{D}^2 (or other non-compact metric spaces). By abuse of terminology, we define a triangulation of a bounded point set S in \mathbb{D}^2 to be a triangulation of S in the hyperbolic convex hull of S in \mathbb{D}^2 . Here we recall that the hyperbolic convex hull conv(S) of a set S is the (inclusion-wise) smallest convex set containing S, i.e., the smallest set containing S such that $[x, y] \in \text{conv}(S)$ for all $x, y \in \text{conv}(S)$. Furthermore, we will always assume that the edges of a triangulation in \mathbb{D}^2 are geodesic segments.

Definition 5.2. A triangulation T of a point set S in \mathbb{D}^2 is called a Delaunay triangulation if each face of T has a hyperbolic circumscribed circle in \mathbb{D}^2 such that the corresponding disk does not contain any points of S in its interior.

For every point set S in \mathbb{D}^2 there exists a Delaunay triangulation, but in general it is not unique. Namely, if S contains a subset S' of at least 4 concircular points such that the corresponding disk does not contain any points of S, then each triangulation of S' in \mathbb{D}^2 yields a Delaunay triangulation. We will denote any Delaunay triangulation of S in \mathbb{D}^2 by $\mathrm{DT}_{\mathbb{D}}(S)$.

Now, let $M = \mathbb{D}^2/\Gamma$ be a closed hyperbolic surface and let S be a finite point set on M. To be able to compute a Delaunay triangulation of S in M, we can instead compute the Delaunay triangulation $\mathrm{DT}_{\mathbb{D}}(\Gamma S)$ of ΓS in the universal covering space \mathbb{D}^2 . Intuitively, we then want to use the universal covering map $\pi : \mathbb{D}^2 \to M$ to define the Delaunay triangulation $\mathrm{DT}_M(S)$ of S in M, i.e. " $\mathrm{DT}_M(S) := \pi(\mathrm{DT}_{\mathbb{D}}(\Gamma S))$ ". However, this is not always a triangulation in the sense of Definition 5.1. For example, it could happen that two vertices of one face of $DT_{\mathbb{D}}(\Gamma S)$ project to the same point under π , with the result that the face after projection is not homeomorphic to Δ any more. This leads to the following definition.

Definition 5.3. Let $M = \mathbb{D}^2/\Gamma$ be a closed hyperbolic surface with universal covering map $\pi : \mathbb{D}^2 \to M$ and let S be a finite point set on M. If $\pi(\mathrm{DT}_{\mathbb{D}}(\Gamma S))$ is a triangulation of M, then it is called the Delaunay triangulation of S in M and denoted by $\mathrm{DT}_M(S)$.

As mentioned above, $\pi(\mathrm{DT}_{\mathbb{D}}(\Gamma S))$ can fail to be a triangulation if vertices of one face project to the same point. More generally, if we look at the graph consisting of the vertices and edges of $\pi(\mathrm{DT}_{\mathbb{D}}(\Gamma S))$, then we have the following characterization of $\pi(\mathrm{DT}_{\mathbb{D}}(\Gamma S))$ being a triangulation.

Proposition 5.4. If the graph of $\pi(DT_{\mathbb{D}}(\Gamma S))$ does not contain cycles of length at most 2, then $\pi(DT_{\mathbb{D}}(\Gamma S))$ is a triangulation of M.

Even though the proof given in [17] is for closed Euclidean surfaces, it is purely combinatorial, so it generalizes immediately to closed hyperbolic surfaces. Cycles of length 1 are loops in the graph and cycles of length 2 are occurrences of multiple edges in the graph. Both correspond to closed curves in M which are freely homotopic to a closed geodesic.

We want to find a geometric criterion equivalent to Proposition 5.4 involving only the surface M (and possibly its corresponding Fuchsian group Γ) and the point set S. Given a point set $S \subset M = \mathbb{D}^2/\Gamma$ we can define δ_S to be the diameter of the largest disk in \mathbb{D}^2 that does not contain any point of ΓS in its interior. Instead of $\delta_{\{p\}}$ for $p \in M$ we write δ_p . For $p \in M$ largest empty disks are centered at the furthest vertex of $D_p(\Gamma)$. Define $\delta_M = \sup\{\delta_p \mid p \in D_O(\Gamma)\}$, where O denotes the origin.

Proposition 5.5. Let $M = \mathbb{D}^2/\Gamma$ be a closed hyperbolic surface and S a finite point set on M. If $syst(M) > \delta_S$, then the graph of $\pi(DT_{\mathbb{D}}(\Gamma S))$ does not contain a cycle of length 1. If $syst(M) > 2\delta_S$, then $\pi(DT_{\mathbb{D}}(\Gamma S))$ does not contain a cycle of length 2.

Note that if $\operatorname{syst}(M) > 2\delta_M$, then the graph $\pi(\operatorname{DT}_{\mathbb{D}}(\Gamma S))$ does not contain cycles of length 1 or 2 for any finite point set S on M, since $\delta_S \leq \delta_p \leq \delta_M$ for every $p \in S$. Unfortunately, this is not always true. Therefore, we will look at covering spaces of M.

Let $\Gamma' < \Gamma$ be a subgroup of Γ . We have seen before that Γ' is a Fuchsian group as well, so $M' = \mathbb{D}^2/\Gamma'$ is a closed hyperbolic surface with universal covering map $\pi' : \mathbb{D}^2 \to M'$. It can be seen that M' is a covering space of M. Instead of projecting $\mathrm{DT}_{\mathbb{D}}(S)$ onto M, we can project $\mathrm{DT}_{\mathbb{D}}(S)$ onto M'. In this case we have a similar proposition.

Proposition 5.6. Let the notation be as above. If $syst(M') > \delta_M$, then the graph of $\pi'(DT_{\mathbb{D}}(\Gamma S))$ does not contain a cycle of length 1 for any finite point set S on M. If $syst(M') > 2\delta_M$, then $\pi'(DT_{\mathbb{D}}(\Gamma S))$ does not contain a cycle of length 2 for any finite point set S on M.

In principle, we can work with any covering space of M. For example, if we take $\Gamma' = {\text{Id}}$, then $M' = \mathbb{D}^2$, so $\pi'(\text{DT}_{\mathbb{D}}(\Gamma S)) = \text{DT}_{\mathbb{D}}(\Gamma S)$ certainly is a triangulation. However, this is not useful for practical applications. In this case we want to restrict to covering spaces M' of M with finitely many sheets. **Proposition 5.7.** There exists a finite-sheeted covering space M' of M such that $syst(M') > 2\delta_M$. There exists a finite-sheeted covering space M'' of M, such that $syst(M'') > 2\delta_M$.

The conclusion is then the following.

Corollary 5.8. Let $M = \mathbb{D}^2/\Gamma$ be a closed hyperbolic surface with universal covering map $\pi : \mathbb{D}^2 \to M$ and let S be a finite point set on M. There exists a finite-sheeted covering space M' of M with universal covering map $\pi' : \mathbb{D}^2 \to M'$ such that $\pi'(DT_{\mathbb{D}}(\Gamma S))$ is a triangulation of M'.

The proof uses the fact that for any B > 0 there are only finitely many closed geodesics with length bounded by B and the existence of subgroups Γ' of finite index which exclude the group elements corresponding to these short geodesics. In this way, the systole of the covering space M' can be made arbitrarily large.

6 Lower bound for the number of sheets

The main result of this section is the following theorem. We emphasize that this generalizes the result for the Bolza surface given in [8] in two ways: it gives a lower bound for all genera; it gives a lower bound for all conformal structures, by which we mean all elements of the corresponding Teichmüller space. Furthermore, it increases the lower bound for the Bolza surface from 32 to 34 sheets.

Theorem 6.1. Let M be any hyperbolic surface of genus $g \ge 2$. Let M' be a k-sheeted covering space of M such that $syst(M') > 2\delta_M$. Then

$$k > \frac{\pi}{9} \cdot \frac{\cot^2(\frac{\pi}{12g-6}) - 3}{g-1} \sim \frac{16}{\pi} \cdot g$$

Furthermore, if M can be represented by a fundamental region with 4g concircular vertices, then we have the higher upper bound

$$k > \frac{\pi}{3} \cdot \frac{\cot^4 \frac{\pi}{4g} - 1}{g - 1} \sim \frac{256}{3\pi^3} \cdot g^3.$$

Proof. Let $M = \mathbb{D}^2/\Gamma$. We will first give a lower bound for δ_O , since this will give a lower bound for δ_M . As already mentioned in [8, p. 20:7], largest empty disks are centered at the furthest vertices of the Dirichlet region $D_O(\Gamma)$. Let R be the distance between O and the furthest vertex of $D_O(\Gamma)$. Let P be a normal canonical polygon representing M (see Section 4.10.7). By Proposition 4.8, we have $\operatorname{area}(D_O(\Gamma)) = \operatorname{area}(P) = 4\pi(g-1)$. Then by Corollary 4.2

$$\cosh R \ge \cosh R(4\pi(g-1)) = \cot\left(\frac{\pi}{n}\right)\cot\left(\frac{(n-2)\pi - (4g-4)\pi}{2n}\right),\tag{2}$$

where n is the number of sides of $D_O(\Gamma)$. It can be shown that $4g \le n \le 12g - 6$; see, e.g., [6, Thm. 10.5.1]. Then

$$\cot\left(\frac{\pi}{n}\right)\cot\left(\frac{(n-2)\pi-(4g-4)\pi}{2n}\right) \ge \cot\left(\frac{\pi}{12g-6}\right)\cot\left(\frac{(12g-6-2)\pi-(4g-4)\pi}{2(12g-6)}\right),$$
$$=\cot\left(\frac{\pi}{12g-6}\right)\cot(\frac{\pi}{3}),$$
$$=\cot\left(\frac{\pi}{12g-6}\right)\cdot\frac{1}{3}\sqrt{3},$$

because the left-hand side is a decreasing function in n. Because $\delta_O = 2R$, we have the following bound:

$$\delta_O = 2R \ge 2 \operatorname{arccosh}\left(\operatorname{cot}\left(\frac{\pi}{12g-6}\right) \cdot \frac{1}{3}\sqrt{3}\right)$$

Therefore,

$$\cosh \delta_M \ge \cosh \delta_O = \cosh(2R) = 2\cosh^2 R - 1 \ge \frac{2}{3}\cot^2(\frac{\pi}{12g-6}) - 1.$$

Let D be a maximal disk in M', i.e. its diameter is equal to syst(M'). Then

$$\operatorname{area}(D) = 2\pi \left(\cosh\left(\frac{1}{2}\operatorname{syst}(M')\right) - 1\right) > 2\pi \left(\cosh\delta_M - 1\right).$$

Assume for a contradiction that $\operatorname{area}(M') < \frac{\pi}{3}\operatorname{area}(D)$. Observe that D and γD have disjoint interiors for $\gamma \in \Gamma \setminus \{\operatorname{Id}\}$, since D is an embedded disk. Furthermore, D and γD are isometric, since γ is an isometry. Therefore, ΓD is a circle packing with congruent circles in the hyperbolic plane. Because we can tile the hyperbolic plane by translates of a fundamental region under Γ and each tile contains a single circle (possibly in pieces), the packing density of this circle packing is equal to

$$\frac{\operatorname{area}(D)}{\operatorname{area}(M')} > \frac{3}{\pi}.$$

This contradicts the fact that the maximum packing density in the hyperbolic plane is $\frac{3}{\pi}$ (see Section 4.6). Hence, $\operatorname{area}(M') \geq \frac{\pi}{3} \operatorname{area}(D)$. It follows that

$$k = \frac{\operatorname{area}(M')}{\operatorname{area}(M)},$$

$$\geq \frac{\pi}{3} \cdot \frac{\operatorname{area}(D)}{\operatorname{area}(M)},$$

$$\geq \frac{\pi}{3} \cdot \frac{2\pi(\cosh \delta_M - 1)}{4\pi(g - 1)},$$

$$\geq \frac{\pi}{9} \cdot \frac{\cot^2(\frac{\pi}{12g - 6}) - 3}{g - 1}.$$

For sufficiently large g, we can use the approximations

$$\cos\left(\frac{\pi}{12g-6}\right) = 1 - \frac{1}{2} \left(\frac{\pi}{12g-6}\right)^2 + \text{h.o.t.},$$
$$\sin\left(\frac{\pi}{12g-6}\right) = \frac{\pi}{12g-6} + \text{h.o.t.},$$

to obtain

$$\cot(\frac{\pi}{12g-6}) = \frac{1 - \frac{1}{2} \left(\frac{\pi}{12g-6}\right)^2}{\frac{\pi}{12g-6}} + \text{h.o.t.} = \frac{12g}{\pi} + \text{h.o.t.},$$

 \mathbf{SO}

$$\frac{\pi}{9} \cdot \frac{\cot^2(\frac{\pi}{12g-6}) - 3}{g-1} = \frac{16}{\pi} \cdot g + \text{h.o.t.},$$

where h.o.t. denotes 'higher order terms'. Now, suppose that M can be represented by a fundamental region F with 4g concircular vertices. Let O be the center of the circle passing through the vertices. We have that $F = D_O(\Gamma)$, since each vertex of F is equidistant from at least three surrounding points in the orbit of O under Γ . Namely, if vertex v has neighbours x, w, then v is equidistant from $O, g_x(O), g_w(O)$, where g_x, g_w denote the side pairing transformations mapping a side of the fundamental region to [v, x], [v, w]respectively. It follows that M has a Dirichlet region with 4g vertices. Therefore, we can now take n = 4g in equation (2), from which we get

$$\cosh R \ge \cot\left(\frac{\pi}{4g}\right)\cot\left(\frac{(4g-2)\pi - (4g-4)\pi}{2\cdot 4g}\right) = \cot^2\frac{\pi}{4g}.$$

By using the same reasoning as before, we obtain $\cosh \delta_M > 2 \cot^4 \frac{\pi}{4a} - 1$ and

$$k > \frac{\pi}{3} \cdot \frac{\cot^4 \frac{\pi}{4g} - 1}{g - 1}.$$

For sufficiently large g we get in this case that

$$\cot \frac{\pi}{4g} = \frac{4g}{\pi} + \text{h.o.t.},$$

 \mathbf{SO}

$$\frac{\pi}{3} \cdot \frac{\cot^4 \frac{\pi}{4g} - 1}{g - 1} = \frac{256}{3\pi^3} \cdot g^3 + \text{h.o.t.}.$$

Remark 6.2. The lower bound for δ_O is sharp: namely if P is a regular 4g-gon, then by the reasoning above $P = D_O(\Gamma)$. Therefore, the distance between O and the furthest vertex of $D_O(\Gamma)$ is the distance between O and any vertex of P. The origin together with two vertices of P forms an isosceles triangle with angles $\frac{\pi}{2g}, \frac{\pi}{4g}, \frac{\pi}{4g}$ (see Figure 13). Adding a perpendicular from O onto the opposite side yields a right-angled triangle with angles $\frac{\pi}{4g}, \frac{\pi}{4g}$. Its hypotenuse c is the distance between O and a vertex of P and we have $\cosh c = \cot^2(\frac{\pi}{4g})$.



Figure 13: Distance to a vertex in a regular 4g-gon

Remark 6.3. Theorem 6.1 gives the following lower bounds for hyperbolic surfaces represented by fundamental regions with 4g concircular vertices for $2 \le g \le 10$:

7 Upper bound for the number of sheets

In this section we will show that there exists no upper bound for the minimum number of sheets needed for a suitable covering space for hyperbolic surfaces of genus 2.

Theorem 7.1. For all $B \in \mathbb{R}$ there exists a hyperbolic surface M of genus 2 such that if M' is a k-sheeted covering space of M such that $syst(M') > \delta_M$, then k > B.

We will use [1] to explicitly construct surfaces with vertices far from the center. In this article, closed hyperbolic surfaces of genus 2 are identified with octagons of a certain form. Namely, first choose $z_1, z_2, z_3 \in \mathbb{D}^2$ with

$$0 < \arg(z_1) < \arg(z_2) < \arg(z_3) < \pi.$$

It is shown that if

$$\operatorname{Im}((1-\bar{z}_1)(1-z_1\bar{z}_2)(1-z_2\bar{z}_3)(1+z_3)) < 0, \tag{3}$$

then there exists a unique $z_0 \in (0,1) \subset \mathbb{D}^2$ such that $[z_0, z_1, z_2, z_3, z_4, z_5, z_6, z_7]$ (with $z_i = -z_{i-4}, i = 4, 5, 6, 7$) forms a hyperbolic octagon with area 4π . If an octagon of the form above satisfies (3), it is called an admissible octagon. By symmetry in the origin we naturally obtain a set of side pairing transformations: the transformation mapping the side between z_{i+3} and z_{i+4} to the side between z_i and z_{i-1} will be denoted by g_i for i = 1, 2, 3, 4 (note: $z_8 = z_0$). See Figure 14. As we have seen before, from this we can obtain the corresponding Fuchsian group and Dirichlet region. This yields a bijection between triples z_1, z_2, z_3 satisfying the condition and the Teichmüller space for hyperbolic surfaces of genus 2.

In general, the octagons constructed in this way are not equal to the corresponding Dirichlet region $D_O(\Gamma)$. In the proof we will solve this issue by showing that the distance between O and the furthest vertex of $D_O(\Gamma)$ is at least half the distance between O and the furthest vertex of the octagon we started with.

Proof. (of Theorem 7.1) We will first prove that there exist admissible octagons where the furthest vertex is arbitrarily far away. Let $z_j = (1 - \epsilon)e^{\pi i j/4}$ for $0 < \epsilon < 1$ and j = 1, 2, 3. We will show that if ϵ is sufficiently small, then equation (3) holds. We compute

$$1 - \bar{z_1} = 1 - (1 - \epsilon)e^{-\pi i/4},$$

$$= 1 - e^{-\pi i/4} + o(\epsilon),$$

$$1 - z_1 \bar{z_2} = 1 - (1 - \epsilon)e^{\pi i/4} \cdot (1 - \epsilon)e^{-2\pi i/4},$$

$$= 1 - e^{-\pi i/4} + o(\epsilon),$$

$$1 - z_2 \bar{z_3} = 1 - (1 - \epsilon)e^{2\pi i/4} \cdot (1 - \epsilon)e^{-3\pi i/4},$$

$$= 1 - e^{-\pi i/4} + o(\epsilon),$$

$$1 + z_3 = 1 + (1 - \epsilon)e^{3\pi i/4},$$

$$= 1 + e^{3\pi i/4} + o(\epsilon),$$



Figure 14: Admissible octagon

 \mathbf{SO}

$$(1 - \bar{z_1})(1 - z_1 \bar{z_2})(1 - z_2 \bar{z_3})(1 + z_3) = (1 - e^{-\pi i/4})^4 + o(\epsilon),$$

= $(4\sqrt{2} - 6)i + o(\epsilon).$

Therefore $\operatorname{Im}((1 - \overline{z_1})(1 - z_1\overline{z_2})(1 - z_2\overline{z_3})(1 + z_3)) < 0$ for sufficiently small ϵ . Observe that $d(O, z_2) \to \infty$ when $\epsilon \to 0$, so this distance can be made arbitrarily large.

We will show that it follows that the distance δ_O between O and the furthest vertex of $D_O(\Gamma)$ can be made arbitrarily large as well. To do this, set $r = \frac{1}{2}(1-\epsilon)$ and let M be the midpoint of O and z_2 . We will show that $M \in D_O(\Gamma)$, which implies that $\delta_O \geq r$. We have to show that $d(M, O) < d(M, \gamma O)$ for all $\gamma \in \Gamma \setminus \{\text{Id}\}$, where Γ is the Fuchsian group generated by the before mentioned side pairing transformations g_1, \ldots, g_4 . It is sufficient to consider only the translates $g_1M, \ldots, g_4M, g_1^{-1}M, \ldots, g_4^{-1}M$, since these are the closest to O. By symmetry in the origin it is sufficient to look only at the upper half and by symmetry across the imaginary axis it is sufficient to look only at the upper left quadrant (see Figure 15). We will first show that $d(M, O) < d(M, g_3O)$. Let $A = g_3O$. Since

$$[A, z_2] = [g_3O, g_3z_7] = g_3([O, z_7]),$$

we know $d(A, z_2) = d(O, z_7) = 1 - \epsilon = 2r$. Therefore, $[O, z_2, A]$ is an isosceles triangle. Use the following notation:

$$\angle AOM = \tau, \quad d(A, M) = t, \quad \angle OAM = \rho$$



Figure 15: Upper left quadrant

Since $[O, z_2, A]$ is isosceles, we have $\rho < \tau$. Then by the sine rule we have

$$\sinh r = \frac{\sin \rho}{\sin \tau} \cdot \sinh t < \sinh t,$$

which means $d(M, O) < d(M, g_3 O)$.

We will now show $d(M, O) < d(M, g_4O)$. Let $B = g_4O$. Let C be the midpoint of Oand B. Since g_4 maps the midpoint of $[z_7, z_0]$ to the midpoint of $[z_3, z_4]$, we find that Cis the midpoint of $[z_3, z_4]$ with $d(C, z_3) = d(C, z_4) =: b$. As $\epsilon \to \infty$, the interior angles of the octagon at all vertices except z_0, z_4 approach zero, so the interior angles at z_0, z_4 approach π . Since $\angle z_4Oz_3 = \frac{\pi}{4}$ is fixed, this means that $\angle z_3z_4O > \angle z_4Oz_3$ for sufficiently small ϵ . By the same reasoning as above, but in this case with the sine rule in triangle $[O, z_3, z_4]$, this implies b < r for sufficiently small ϵ . Let d(O, C) = x and $\angle Cz_3O = \xi$. By applying the first hyperbolic cosine rule to triangle $[O, C, z_3]$ we see

$$\cosh x = \cosh 2r \cosh b - \sinh 2r \sinh b \cos \xi,$$

$$\geq \cosh 2r \cosh b - \sinh 2r \sinh b,$$

$$= \cosh(2r - b).$$

We have shown above that b < r for sufficiently small ϵ , so in that case

$$\cosh x \ge \cosh(2r - b) > \cosh r.$$

Hence, x > r for sufficiently small ϵ . By the triangle inequality,

$$d(M,B) \ge d(O,B) - d(O,M) = 2x - r > r = d(M,O),$$

from which we conclude $d(M, g_4 O) > d(M, O)$. It follows that $M \in D_O(\Gamma)$, so $\delta_O \geq \frac{1}{2}(1-\epsilon)$, which can be made arbitrarily large.

Let $B \in \mathbb{R}$ be arbitrary. By the argument above, for all $B' \in \mathbb{R}$ there exists a hyperbolic surface M of genus 2 such that $\delta_M \geq \delta_O \geq B'$. Choose $B' = \operatorname{arccosh}(2B+1)$. Now, we can use a similar argument as for the lower bound. Let M' be a k-sheeted covering space of M such that $\operatorname{syst}(M') > 2\delta_M$. Let D be a maximal disk in M', i.e. its diameter is equal to $\operatorname{syst}(M')$. Then

$$\operatorname{area}(M') > \operatorname{area}(D) = 2\pi(\cosh(\frac{1}{2}\operatorname{syst}(M') - 1)) > 2\pi(\cosh\delta_M - 1).$$

We conclude that

$$k = \frac{\operatorname{area}(M')}{\operatorname{area}(M)} > \frac{2\pi(\cosh \delta_M - 1)}{4\pi} > \frac{\cosh B' - 1}{2} = B.$$

Remark 7.2. In the proof of Theorem 7.1 we used angles $\pi j/4$ for convenience. However, the proof can be directly generalized to $z_j = (1 - \epsilon)e^{\theta_j i}$, j = 1, 2, 3 if $\theta_3 > \frac{\pi}{2} > \theta_1$. Again, first we show that for sufficiently small ϵ , this triple is admissible. We have

$$1 - \bar{z_1} = 1 - e^{-\theta_1 i} + o(\epsilon),$$

$$1 - z_1 \bar{z_2} = 1 - e^{(\theta_1 - \theta_2)i} + o(\epsilon),$$

$$1 - z_2 \bar{z_3} = 1 - e^{(\theta_2 - \theta_3)i} + o(\epsilon),$$

$$1 + z_3 = 1 + e^{\theta_3 i} + o(\epsilon).$$

After a straightforward computation we obtain

$$\frac{\operatorname{Im}((1-\bar{z_1})(1-\bar{z_1}\bar{z_2})(1-\bar{z_2}\bar{z_3})(1+\bar{z_3}))}{2} = \sin\theta_1 - \sin\theta_2 + \sin\theta_3 + \sin(\theta_2 - \theta_1) + \sin(\theta_3 - \theta_2) + \sin(\theta_1 - \theta_3) + \sin(\theta_2 - \theta_1 - \theta_3) + o(\epsilon) = -8\sin\left(\frac{\theta_1}{2}\right)\cos\left(\frac{\theta_3}{2}\right)\sin\left(\frac{\theta_1 - \theta_2}{2}\right)\sin\left(\frac{\theta_2 - \theta_3}{2}\right) + o(\epsilon).$$

The right-hand side can be seen to be negative for sufficiently small ϵ , which yields an admissible octagon.

To show that the furthest vertex of $D_O(\Gamma)$ has arbitrarily large distance to O, we still consider only the upper left quadrant: in this case the octagon is not necessarily symmetric in the imaginary axis, but the argument is exactly the same for both sides. The proof that $d(M, O) < d(M, g_3 O)$ is exactly the same as above, since we only used that $[O, z_3, g_3 O]$ is an isosceles triangle. For the proof that $d(M, O) < d(M, g_4 O)$, recall that the only thing we used here is that b < r for sufficiently small ϵ . We assumed that $\theta_3 > \frac{\pi}{2}$ so $\angle z_4 O z_3 < \frac{\pi}{2}$ and as $\epsilon \to 0$ we have $\angle z_3 z_4 O \to \frac{\pi}{2}$, so indeed $\angle z_3 z_4 O > \angle z_4 O z_3$ for sufficiently small ϵ . Hence, b < r for sufficiently small ϵ . This finishes the proof.

8 Systole of surfaces corresponding to regular polygons

In this section we will discuss the following conjecture.

Conjecture 8.1. The systole of the surface M_g corresponding to the regular 4g-gon satisfies

$$\cosh\left(\frac{\operatorname{syst}(M_g)}{2}\right) = 1 + 2\cos(\frac{\pi}{2g}).$$

As a partial result we have

Theorem 8.2. The systole of the surface M_g corresponding to the regular 4g-gon satisfies

$$\cosh\left(\frac{\operatorname{syst}(M_g)}{2}\right) \le 1 + 2\cos(\frac{\pi}{2g}),$$

with equality for g = 2, 3.

The inequality for arbitrary genus is new and we will prove it in Section 8.2. The remainder of the section will describe our progress in obtaining a complete proof of the conjecture. At the end we will look at three examples, namely g = 2, 3, 6. The systole of M_2 was already known in the literature, but the method developed in this section leads to a new proof, as well as to a proof of the case g = 3. The case g = 6 is included to show the difficulties in this approach.

8.1 From systole to optimization problem

According to Section 4.9 the systole of M_g satisfies

$$\cosh\left(\frac{\operatorname{syst}(M_g)}{2}\right) = \min_{\substack{\gamma \in \Gamma_g \\ \gamma \neq \operatorname{Id}}} \frac{1}{2} |\operatorname{tr}(\gamma)|,$$

where Γ_g denotes the Fuchsian group corresponding to the surface M_g . The group Γ_g is generated by the side pairing transformations which pair opposite sides, which can be represented by the matrices

$$A_k = \begin{bmatrix} \cosh(\frac{l_0}{2}) & \sinh(\frac{l_0}{2})\exp(\frac{ik\pi}{2g}) \\ \sinh(\frac{l_0}{2})\exp(-\frac{ik\pi}{2g}) & \cosh(\frac{l_0}{2}) \end{bmatrix}$$

for $k = 0, 1, \ldots, 2g - 1$. Here l_0 is a constant satisfying

$$\cosh(\frac{l_0}{2}) = \cot(\frac{\pi}{4q}).$$

The inverses of $A_k, k = 0, 1, ..., 2g - 1$ are $A_k, k = 2g, ..., 4g - 1$, respectively. Every $M \in \Gamma_g$ can be written as a composition

$$M = B_1 B_2 \cdots B_n,$$

where $B_i \in \{A_0, \ldots, A_{4g-1}\}$. Therefore the minimization problem

$$\label{eq:min_states} \begin{split} \min \frac{1}{2} |\operatorname{tr}(\gamma)|, \\ \text{subject to } \gamma \in \Gamma_g \setminus \{\operatorname{Id}\} \end{split}$$

is equivalent to the minimization problem

$$\min \frac{1}{2} |\operatorname{tr}(M)|,$$

subject to $M \in M(2, \mathbb{C}),$
 $M = B_1 B_2 \cdots B_n \text{ for } B_i \in \{A_0, \dots, A_{4g-1}\}.$

8.2 Upper bound for the systole

To prove an upper bound for the optimal value of a minimization problem, it is sufficient to find a feasible point which attains this upper bound. By the discussion above, it is sufficient to find $M = B_1 B_2 \cdots B_n$ with $B_i \in \{A_0, \ldots, A_{4g-1}\}$ such that $\frac{1}{2} |\operatorname{tr}(M)| = 1 + 2\cos(\frac{\pi}{2g})$ to prove that

$$\cosh\left(\frac{\operatorname{syst}(M_g)}{2}\right) \le 1 + 2\cos(\frac{\pi}{2g}).$$

Consider the upper left entry of $A_k A_{k+2g-1}$, which is given by

$$\begin{bmatrix} \cot(\frac{\pi}{4g}) & \exp(\frac{ik\pi}{2g})\sqrt{\cot^2(\frac{\pi}{4g}) - 1} \end{bmatrix} \begin{bmatrix} \cot(\frac{\pi}{4g}) \\ \exp(-\frac{i(k+2g-1)\pi}{2g})\sqrt{\cot^2(\frac{\pi}{4g}) - 1} \end{bmatrix} = \\ \cot^2(\frac{\pi}{4g}) + (\cot^2(\frac{\pi}{4g}) - 1)\exp(-\frac{(2g-1)\pi}{2g}).$$

Recall from Section 4.2.2 that every orientation preserving isometry φ of \mathbb{D}^2 is of the form

$$\varphi(z) = \frac{az+b}{\bar{b}z+\bar{a}}$$

for $a, b \in \mathbb{C}$ with $|a|^2 - |b|^2 = 1$. Therefore,

$$|\operatorname{tr}(\varphi)| = 2\operatorname{Re}(a).$$

In our case

$$\frac{1}{2}|\operatorname{tr}(A_k A_{k+2g-1})| = \operatorname{cot}^2(\frac{\pi}{4g}) + (\operatorname{cot}^2(\frac{\pi}{4g}) - 1)\cos(-\frac{(2g-1)\pi}{2g}).$$

By using half-angle formulas we can rewrite this in the following way:

$$\cot^{2}(\frac{\pi}{4g}) + (\cot^{2}(\frac{\pi}{4g}) - 1)\cos(-\frac{(2g-1)\pi}{2g}) = \frac{1 + \cos(\frac{\pi}{2g})}{1 - \cos(\frac{\pi}{2g})} - \left(\frac{1 + \cos(\frac{\pi}{2g})}{1 - \cos(\frac{\pi}{2g})} - 1\right)\cos(\frac{\pi}{2g}) = \frac{1 + \cos(\frac{\pi}{2g}) - 1 + \cos(\frac{\pi}{2g})}{1 - \cos(\frac{\pi}{2g})}\cos(\frac{\pi}{2g}) = \frac{-2\cos^{2}(\frac{\pi}{2g}) + \cos(\frac{\pi}{2g}) + 1}{1 - \cos(\frac{\pi}{2g})} = \frac{-2(\cos(\frac{\pi}{2g}) - 1)(\cos(\frac{\pi}{2g}) + \frac{1}{2})}{1 - \cos(\frac{\pi}{2g})} = 1 + 2\cos(\frac{\pi}{2g}).$$

This proves the desired inequality for the systole. Intuitively, that $A_k A_{k+2g-1}$ give the shortest curves would not be surprising: these transformations consist of a translation in a certain fixed direction (dependent on k), followed by the translation that is almost the inverse of A_k . Namely, A_k has inverse A_{k+2g} and A_{k+2g-1} is the translation of which the axis intersects the axis of the inverse with the smallest angle.

8.3 Towards a complete proof

We will show our progress in obtaining a complete proof of the conjecture in a series of lemmas. For a short introduction to the algebra we use, we refer to the Appendix. Recall that $\mathbb{Q}(\zeta_{4g})$ denotes the 4g-th cyclotomic field and denote $K = \mathbb{Q}(\zeta_{4g})$. We will always take $\zeta_{4g} = \exp(\frac{\pi i}{2g})$ as primitive 4g-th root of unity. First we will show that $\cot(\frac{\pi}{4g}) \in K$. By a half-angle formula, we can write

$$\cot(\frac{\pi}{4g}) = \frac{1 + \cos(\frac{\pi}{2g})}{\sin(\frac{\pi}{2g})}$$

We know

$$\cos\left(\frac{\pi}{2g}\right) = \frac{\exp\left(\frac{\pi i}{2g}\right) + \exp\left(-\frac{\pi i}{2g}\right)}{2} \in K,$$
$$\sin\left(\frac{\pi}{2g}\right) = \frac{\exp\left(\frac{\pi i}{2g}\right) - \exp\left(-\frac{\pi i}{2g}\right)}{2\exp\left(\frac{g\pi i}{2g}\right)} \in K,$$

so $\cot(\frac{\pi}{4g}) \in K$. To prove that $\cot(\frac{\pi}{4g})$ is an algebraic integer, we will in fact prove the stronger claim that $\cot(\frac{\pi}{2g}) \in \mathcal{O}_K$. The proof is based on a comment in [50, p. 107].

Lemma 8.3. For all $g \in \mathbb{N}$, $\cot(\frac{\pi}{2q}) \in \mathcal{O}_K$.

Proof. We immediately see that $\cot(\frac{\pi}{2g}) \in K$, since $\cos(\frac{\pi}{2g})$, $\sin(\frac{\pi}{2g}) \in K$, as we have seen before. By definition,

$$\cot\left(\frac{\pi}{2g}\right) = \frac{\cos\left(\frac{\pi}{2g}\right)}{\sin\left(\frac{\pi}{2g}\right)},$$
$$= i \cdot \frac{\exp\left(\frac{\pi i}{2g}\right) + \exp\left(-\frac{\pi i}{2g}\right)}{\exp\left(\frac{\pi i}{2g}\right) - \exp\left(-\frac{\pi i}{2g}\right)},$$
$$= i \cdot \frac{\exp\left(\frac{\pi i}{g}\right) + 1}{\exp\left(\frac{\pi i}{g}\right) - 1}.$$

Denote $\lambda = \exp(\frac{\pi i}{g})$. Since $i \in \mathcal{O}_K$, we only have to show that $\frac{\lambda+1}{\lambda-1} \in \mathcal{O}_K$. Observe that

$$\left(\frac{\lambda+1}{\lambda-1}+1\right)^g + \left(\frac{\lambda+1}{\lambda-1}-1\right)^g = \left(\frac{\lambda+1+\lambda-1}{\lambda-1}\right)^g + \left(\frac{\lambda+1-\lambda+1}{\lambda-1}\right)^g,$$
$$= \frac{2^g(\lambda^g+1)}{(\lambda-1)^g},$$
$$= 0,$$

so $\frac{\lambda+1}{\lambda-1}$ is a root of $f := (X+1)^g + (X-1)^g$. This polynomial is not monic, but has leading coefficient 2. However, it is easily seen that every other coefficient is divisible by 2, which means that $\frac{1}{2}f$ is a monic polynomial with integer coefficients, of which $\frac{\lambda+1}{\lambda-1}$ is a root. Therefore, $\frac{\lambda+1}{\lambda-1}$ is an algebraic integer and we conclude that $\cot(\frac{\pi}{2g}) \in \mathcal{O}_K$. \Box

To be able to describe the elements of the Fuchsian group more precisely, we also need the following result.

Lemma 8.4. For all $g \in \mathbb{N}$, $\cot^2(\frac{\pi}{4g}) - 1 \in 2\mathcal{O}_K$.

Proof. By a double-angle formula

$$\cot\left(\frac{\pi}{2g}\right) = \frac{\cot^2\left(\frac{\pi}{4g}\right) - 1}{2\cot\left(\frac{\pi}{4g}\right)},$$

 \mathbf{SO}

$$\cot^2\left(\frac{\pi}{4g}\right) - 1 = 2\cot\left(\frac{\pi}{2g}\right)\cot\left(\frac{\pi}{4g}\right).$$

By Lemma 8.3, we know that $\cot(\frac{\pi}{2g}), \cot(\frac{\pi}{4g}) \in \mathcal{O}_K$. It follows that $\cot^2(\frac{\pi}{4g}) - 1 \in \mathcal{O}_K$. \Box

We can now give a description of the elements of the Fuchsian group.

Proposition 8.5. Every element of the Fuchsian group Γ_q can be written as

$$\begin{bmatrix} \alpha & \beta \sqrt{\cot^2(\frac{\pi}{4g}) - 1} \\ \bar{\beta} \sqrt{\cot^2(\frac{\pi}{4g}) - 1} & \bar{\alpha} \end{bmatrix},$$

where $\alpha, \beta \in \mathcal{O}_K$ for $K = \mathbb{Q}(\zeta_{4g})$ such that $\alpha - 1 \in 2\mathcal{O}_K$ for all products of an even number of generators of Γ_g and $\alpha - \cot(\frac{\pi}{4g}) \in 2\mathcal{O}_K$ for all products of an odd number of generators of Γ_g .

Remark 8.6. We emphasize that we do not claim that all matrices of the given form correspond to an element of the Fuchsian group. On the contrary, we will show in Example 8.7 that this is not the case.

Proof. We will first prove that all element of Γ_g are of the given form with $\alpha, \beta \in \mathcal{O}_K$. Recall that the generators of Γ_g are given by

$$\begin{bmatrix} \cosh(\frac{l_0}{2}) & \sinh(\frac{l_0}{2})\exp(\frac{ik\pi}{2g}) \\ \sinh(\frac{l_0}{2})\exp(-\frac{ik\pi}{2g}) & \cosh(\frac{l_0}{2}) \end{bmatrix} = \begin{bmatrix} \cot(\frac{\pi}{4g}) & \exp(\frac{ji\pi}{2g})\sqrt{\cot^2(\frac{\pi}{4g}) - 1} \\ \exp(-\frac{ji\pi}{2g})\sqrt{\cot^2(\frac{\pi}{4g}) - 1} & \cot(\frac{\pi}{4g}) \end{bmatrix}$$

In this case, we have $\alpha = \cot(\frac{\pi}{4g}) \in \mathcal{O}_K$ and $\beta = \exp(\frac{ji\pi}{2g}) \in \mathcal{O}_K$. Let M, M' be matrices of the given form with $\alpha, \beta \in \mathcal{O}_K$ and $\alpha', \beta' \in \mathcal{O}_K$ respectively. Then

$$MM' = \begin{bmatrix} \alpha & \beta \sqrt{\cot^2(\frac{\pi}{4g}) - 1} \\ \bar{\beta} \sqrt{\cot^2(\frac{\pi}{4g}) - 1} & \bar{\alpha} \end{bmatrix} \begin{bmatrix} \alpha' & \beta' \sqrt{\cot^2(\frac{\pi}{4g}) - 1} \\ \bar{\beta}' \sqrt{\cot^2(\frac{\pi}{4g}) - 1} & \bar{\alpha}' \end{bmatrix}$$
$$= \begin{bmatrix} \alpha \alpha' + \beta \bar{\beta}' (\cot^2(\frac{\pi}{4g}) - 1) & (\alpha \beta' + \bar{\alpha}' \beta) \sqrt{\cot^2(\frac{\pi}{4g}) - 1} \\ (\alpha' \bar{\beta} + \bar{\alpha} \bar{\beta}') \sqrt{\cot^2(\frac{\pi}{4g}) - 1} & \bar{\alpha} \bar{\alpha}' + \bar{\beta} \beta' (\cot^2(\frac{\pi}{4g}) - 1) \end{bmatrix}.$$

Since $\alpha, \beta, \alpha', \beta', \cot(\frac{\pi}{4g}) \in \mathcal{O}_K$, we also have $\alpha \alpha' + \beta \bar{\beta}' (\cot^2(\frac{\pi}{4g}) - 1), \alpha \beta' + \bar{\alpha}' \beta \in \mathcal{O}_K$. Therefore, we have shown that set of matrices of the given form with $\alpha, \beta \in \mathcal{O}_K$ is closed under multiplication. Since all generators of Γ_g are of this form, we deduce that all elements of Γ_g have the given form with $\alpha, \beta \in \mathcal{O}_K$.

Note that the upper left entry of the product of two generators can be given by

$$\alpha = \cot^2\left(\frac{\pi}{4g}\right) + \exp\left(\frac{(j-k)i\pi}{2g}\right)\left(\cot^2\left(\frac{\pi}{4g}\right) - 1\right)$$

for some $1 \leq j, k \leq 4g - 1$. Therefore, in this case we have

$$\alpha - 1 = (\cot^2(\frac{\pi}{4g}) - 1)(1 + \exp(\frac{(j-k)i\pi}{2g})).$$

Since $\cot^2(\frac{\pi}{4g}) - 1 \in 2\mathcal{O}_K$ by Lemma 8.4 and since $1 + \exp(\frac{(j-k)i\pi}{2g}) \in \mathcal{O}_K$, we have $\alpha - 1 \in 2\mathcal{O}_K$. Let M, M' be matrices of the given form with $\alpha, \beta \in \mathcal{O}_K$ such that $\alpha - 1, \alpha' - 1 \in 2\mathcal{O}_K$. As we have seen before, the upper left entry of MM' is given by

$$\gamma := \alpha \alpha' + \beta \bar{\beta}' (\cot^2(\frac{\pi}{4g}) - 1).$$

Then

$$\gamma - 1 = \alpha(\alpha' - 1) + (\alpha - 1) + \beta \overline{\beta}'(\cot^2(\frac{\pi}{4g}) - 1).$$

Because $\alpha - 1$, $\alpha' - 1$, $\cot^2(\frac{\pi}{4g}) - 1 \in 2\mathcal{O}_K$ and the remaining terms are in \mathcal{O}_K , we see that $\gamma - 1 \in 2\mathcal{O}_K$. Therefore, multiplication preserves the property that $\alpha - 1 \in 2\mathcal{O}_K$. Since all products of two generators satisfy this property, we see that all products of an even number of generators satisfy $\alpha - 1 \in 2\mathcal{O}_K$.

Finally, to show the statement for products of an odd number of generators, we know that every product of an odd number of generators is either a generator itself or the product of a generator with a product of an even number of generators. For all generators, we have $\alpha = \cot(\frac{\pi}{4g})$, so clearly $\alpha - \cot(\frac{\pi}{4g}) \in 2\mathcal{O}_K$ in this case. If M is a generator and M' a product of an even number of generators with $\alpha', \beta' \in \mathcal{O}_K$ such that $\alpha' - 1 \in 2\mathcal{O}_K$, then the upper left entry of MM' is given by

$$\gamma := \cot(\frac{\pi}{4q})\alpha' + \bar{\beta}' \exp(\frac{ji\pi}{2q})(\cot^2(\frac{\pi}{4q}) - 1)$$

for some $1 \leq j \leq 4g - 1$. Then

$$\gamma - \cot(\frac{\pi}{4g}) = \cot(\frac{\pi}{4g})(\alpha' - 1) + \bar{\beta}' \exp(\frac{ji\pi}{2g})(\cot^2(\frac{\pi}{4g}) - 1).$$

Because $\alpha' - 1, \cot^2(\frac{\pi}{4g}) - 1 \in 2\mathcal{O}_K$ and the remaining terms are in \mathcal{O}_K , we see that $\gamma - \cot(\frac{\pi}{4g}) \in 2\mathcal{O}_K$. This finishes the proof.

Proposition 8.5 shows that all elements of the Fuchsian group Γ_g have a certain structure. However, not all matrices having this structure are elements of the Fuchsian group. We will illustrate this with an example.

Example 8.7. Assume that $g \geq 3$. Consider $\alpha_n = 1 + (4 - 2\zeta_{4g} - 2\zeta_{4g}^{-1})^n$. First of all, we have $4 - 2\zeta_{4g} - 2\zeta_{4g}^{-1} = 2(2 - \zeta_{4g} - \zeta_{4g}^{-1}) \in 2\mathcal{O}_K$. This means that $\alpha_n - 1 \in 2\mathcal{O}_K$. Furthermore, $\zeta_{4g} + \zeta_{4g}^{-1} = 2\cos(\frac{\pi}{2g})$, so $\alpha_n \in \mathbb{R}$ and $\alpha_n > 1$ for all $n \in \mathbb{N}$. However, $\cos(\frac{\pi}{2g}) \geq \cos(\frac{\pi}{6}) = \frac{1}{2}\sqrt{3} > \frac{3}{4}$, so $4 - 2\zeta_{4g} - 2\zeta_{4g}^{-1} = 4(1 - \cos(\frac{\pi}{2g})) < 1$, which means that $\alpha_n \to 1$ as $n \to \infty$. If all α_n correspond to an element of the Fuchsian group, then Γ_g would contain hyperbolic elements with trace arbitrarily close to 2, which is not possible by discreteness of Γ_g .

From the previous example it is clear that we need an extra restriction on α, β . Such a restriction is given by the following lemma. For an illustration of the idea behind this lemma and its usefulness we refer to the examples in Section 8.5.

Lemma 8.8. Let

$$M = \begin{bmatrix} \alpha & \beta \sqrt{\cot^2(\frac{\pi}{4g}) - 1} \\ \bar{\beta} \sqrt{\cot^2(\frac{\pi}{4g}) - 1} & \bar{\alpha} \end{bmatrix}$$

be an element of the Fuchsian group Γ_g . Let $\psi_k \in \text{Gal}(K : \mathbb{Q})$ be the automorphism sending $\exp(\frac{\pi}{2g})$ to $\exp(\frac{k\pi}{2g})$, in particular we have $\gcd(k, 4g) = 1$. If g < k < 3g, then

$$|\psi_k(\alpha)| < 1, \quad |\psi_k(\beta)| < (\cot^2(\frac{k\pi}{4g}) - 1)^{-1}.$$

Proof. Since det M = 1, we have

$$|\alpha|^2 + |\beta|^2 (1 - \cot^2(\frac{\pi}{4g})) = 1$$

This is a formula for elements of K, so it still holds if we apply an automorphism to it:

$$|\psi_k(\alpha)|^2 + |\psi_k(\beta)|^2 (1 - \cot^2(\frac{k\pi}{4g})) = 1.$$
(4)

Note that indeed $\psi_k(\cot(\frac{\pi}{4g})) = \cot(\frac{k\pi}{4g})$. For g < k < 3g we have $\frac{\pi}{4} < \frac{k\pi}{4g} < \frac{3\pi}{4}$, so $-1 < \cot(\frac{k\pi}{4g}) < 1$. Then all terms in (4) are positive, so each of the summands must be smaller than 1, from which we conclude that

$$|\psi_k(\alpha)| < 1, \quad |\psi_k(\beta)| < (1 - \cot^2(\frac{k\pi}{4g}))^{-1}.$$

To find the systole we are primarily interested in the values $\operatorname{Re}(\alpha)$ can take for elements of the Fuchsian group. In Proposition 8.5 we have seen that either $\alpha - 1 \in 2\mathcal{O}_K$ or $\alpha - \cot(\frac{\pi}{4g}) \in 2\mathcal{O}_K$. To reduce these two cases to one, we first give an expression for a sum of cosines.

Lemma 8.9. For all $x \in \mathbb{R}$ and $n \in \mathbb{N}$,

$$\sum_{j=0}^{n} \cos(jx) = \frac{1}{2} + \frac{1}{2} \cdot \frac{\sin((n+\frac{1}{2})x)}{\sin(\frac{1}{2}x)}.$$

Proof. For fixed $0 \le j \le n$ we have

$$\sin((j+\frac{1}{2})x) - \sin((j-\frac{1}{2})x) = 2\sin(\frac{1}{2}x)\cos(jx).$$

Consider the telescoping sum

$$\sum_{j=0}^{n} 2\sin(\frac{1}{2}x)\cos(jx) = \sum_{j=0}^{n} \left(\sin((j+\frac{1}{2})x) - \sin((j-\frac{1}{2})x)\right),$$
$$= -\sin((0-\frac{1}{2})x) + \sin((n+\frac{1}{2})x),$$
$$= \sin(\frac{1}{2}x) + \sin((n+\frac{1}{2})x).$$

The result is obtained by dividing both sides by $2\sin(\frac{1}{2}x)$.

_	_	
	1	
	1	
	- 1	
_	_	

By using this expression we can rewrite $\cot(\frac{\pi}{4g})$ in the following way. Lemma 8.10. For all $g \in \mathbb{N}$,

$$\cot(\frac{\pi}{4g}) = 1 + 2\sum_{j=1}^{g-1} \cos(\frac{j\pi}{2g}).$$

Proof. By Lemma 8.9 with $x = \frac{\pi}{2g}$ and n = g - 1 we have

$$\sum_{j=0}^{g-1} \cos(\frac{j\pi}{2g}) = \frac{1}{2} + \frac{1}{2} \cdot \frac{\sin((g-1+\frac{1}{2})\frac{\pi}{2g})}{\sin(\frac{1}{2} \cdot \frac{\pi}{2g})}.$$

Observe that

$$\sin((g-1+\frac{1}{2})\frac{\pi}{2g}) = \sin(\frac{\pi}{2} - \frac{\pi}{4g}) = \cos\frac{\pi}{4g},$$

from which we see that

$$\sum_{j=0}^{g-1} \cos(\frac{j\pi}{2g}) = \frac{1}{2} + \frac{1}{2} \cdot \frac{\cos(\frac{\pi}{4g})}{\sin(\frac{\pi}{4g})} = \frac{1}{2} + \frac{1}{2} \cot(\frac{\pi}{4g}).$$

Then

$$\cot(\frac{\pi}{4g}) = -1 + 2\sum_{j=0}^{g-1} \cos(\frac{j\pi}{2g}) = 1 + 2\sum_{j=1}^{g-1} \cos(\frac{j\pi}{2g}).$$

We find the following characterization of α .

Lemma 8.11. Let $\alpha \in \mathcal{O}_K$ be such that $\alpha - 1 \in 2\mathcal{O}_K$ or $\alpha - \cot(\frac{\pi}{4g}) \in 2\mathcal{O}_K$. Then we can write

$$\operatorname{Re}(\alpha) = 1 + 2\sum_{j=0}^{g-1} a_j \cos(\frac{j\pi}{2g})$$

for $a_j \in \mathbb{Z}$.

Proof. Recall that $\mathcal{O}_K = \mathbb{Z}[\zeta_{4g}]$. If $\alpha - 1 \in 2\mathcal{O}_K$, we can write

$$\alpha = 1 + 2\sum_{j=0}^{4g-1} b_j \zeta_{4g}^j$$

for $b_j \in \mathbb{Z}$. Then

$$\operatorname{Re}(\alpha) = 1 + 2 \sum_{j=0}^{4g-1} b_j \cos(\frac{j\pi}{2g}).$$

Because

$$\cos\left(\frac{(2g-j)\pi}{2g}\right) = -\cos\left(\frac{j\pi}{2g}\right),\\ \cos\left(\frac{(2g+j)\pi}{2g}\right) = -\cos\left(\frac{j\pi}{2g}\right),\\ \cos\left(\frac{(4g-j)\pi}{2g}\right) = \cos\left(\frac{j\pi}{2g}\right),$$

we can rewrite this as

$$\operatorname{Re}(\alpha) = 1 + 2b_0 - 2b_{2g} + 2\sum_{j=1}^{g} (b_j - b_{2g-j} - b_{2g+j} + b_{4g-j}) \cos(\frac{j\pi}{2g}).$$

The indices require a word of caution: by choosing $1 \leq j \leq g$, we can combine coefficients in groups of four. However, this leads to missing out on the terms with coefficients b_0, b_{2g} , which we added again, and the terms with coefficients b_g, b_{3g} , which do not matter, since $\cos(\frac{g\pi}{2g}) = \cos(\frac{3g\pi}{2g}) = 0$. It is clear that we obtain an expression of the form

$$\operatorname{Re}(\alpha) = 1 + 2\sum_{j=0}^{g-1} a_j \cos(\frac{j\pi}{2g}).$$

If $\alpha - \cot(\frac{\pi}{4g}) \in 2\mathcal{O}_K$, then we can write

$$\alpha = \cot(\frac{\pi}{4g}) + 2\sum_{j=0}^{4g-1} b_j \zeta_{4g}^j$$

for $b_j \in \mathbb{Z}$. By Lemma 8.10 this is equal to

$$\alpha = 1 + 2\sum_{j=1}^{g-1} \cos(\frac{j\pi}{2g}) + 2\sum_{j=0}^{4g-1} b_j \zeta_{4g}^j.$$

Then

$$\alpha = 1 + 2\sum_{j=1}^{g-1} \cos(\frac{j\pi}{2g}) + 2\sum_{j=0}^{4g-1} b_j \cos(\frac{j\pi}{2g}).$$

The result is obtained by rewriting the second sum in the same way as before. \Box

8.4 Summary and speculations

In Proposition 8.5 we have given a characterization of the elements of Γ_g . In Lemma 8.8 we have given a number of restrictions on these elements. We did not show⁵ that all matrices of the form of Proposition 8.5 satisfying the restrictions of Lemma 8.8 correspond to an element of Γ_g . However, we have shown that the set consisting of these matrices includes the set of feasible matrices for the optimization problem for finding the systole (see Section 8.1). We conjecture that the minimum of $\frac{1}{2}|\operatorname{tr}(M)|$ over this larger set of matrices is equal to the proposed value of the systole. If that is the case, it follows that the solution to the optimization problem in Section 8.1 is also equal to this value, because it is attained in the feasible set. Hence, this would complete the proof of our conjecture.

So far, we could not find the minimum of $\frac{1}{2}|\operatorname{tr}(M)|$ over this larger set of matrices. The main difficulty is working with the restrictions of Lemma 8.8, because it is not

⁵In fact, in its current form this statement is already not true for g = 2. Namely, in [2] it is shown that extra conditions are necessary to ensure that every matrix of this form belongs to Γ_2 . It is not known if the statement is true for other genera, but it seems probable that this is not the case.

immediately apparent what the set of matrices satisfying these restrictions looks like.

In the last part of the previous section we have looked at $\operatorname{Re}(\alpha)$ instead. We can apply Lemma 8.8 to the case where we only look at $\operatorname{Re}(\alpha)$. Namely, because $|\psi_k(\alpha)| < 1$ for all $\psi_k \in \operatorname{Gal}(K : \mathbb{Q})$ with g < k < 3g, we also have

$$|\psi_k(\operatorname{Re}(\alpha))| = |\operatorname{Re}(\psi_k(\alpha))| < 1.$$

If we would have that the optimal value of the optimization problem

$$\min 1 + 2\sum_{j=0}^{g-1} a_j \cos(\frac{j\pi}{2g}),$$

subject to $a_j \in \mathbb{Z}$,

$$1 + 2\sum_{j=0}^{g-1} a_j \cos(\frac{j\pi}{2g}) > 1,$$

$$\left| 1 + 2\sum_{j=0}^{g-1} a_j \cos(\frac{kj\pi}{2g}) \right| < 1 \text{ for } g < k < 3g, \gcd(k, 4g) = 1,$$

is equal to our proposed value of the systole, then we could use the same reasoning as before to conclude that the conjecture is true. In the next section we will see examples where this is indeed the case (g = 2, 3) and an example where this is not the case (g = 6). Therefore, we cannot use this simplification in general. However, in numerical simulations it seems to be the case that all a_j have the same sign, which suggests that we maybe need to introduce some extra restrictions.

8.5 Examples

In [2] the systole of M_2 was computed using modular arithmetic. By using the method from this section we obtain a new proof.

Example 8.12. Consider g = 2. In [2] and [3] it is shown using modular arithmetic that if M is a product of elements of Γ_2 , then $\frac{1}{2}|\operatorname{tr}(M)|$ has the form $|m(n) + n\sqrt{2}|$ where given $n \in \mathbb{Z}$, m(n) is the unique odd integer minimizing $|m - n\sqrt{2}|$. From this it follows that $\cosh(\operatorname{syst}(M_2)/2) = 1 + \sqrt{2}$. We will show that our method yields the same solution.

By Lemma 8.11 we can write

$$\frac{1}{2}|\operatorname{tr}(M)| = |\operatorname{Re}(\alpha)| = |1 + 2a_0 + 2a_1 \cos(\frac{\pi}{4})| = |1 + 2a_0 + a_1 \sqrt{2}|$$

for $a_0, a_1 \in \mathbb{Z}$, where α is the upper left entry of an arbitrary element of Γ_2 . By Lemma 8.8 we have $|\psi_k(\alpha)| < 1$ for k = 3 and k = 5. For k = 3 we have $\zeta_8 \mapsto \zeta_8^3$, so $\exp(\frac{\pi i}{4}) \mapsto \exp(\frac{3\pi i}{4})$, so $\cos(\frac{\pi}{4}) \mapsto \cos(\frac{3\pi}{4})$, so $\sqrt{2} \mapsto -\sqrt{2}$. Then the condition $|\psi_3(\operatorname{Re}(\alpha))| < 1$ can be rewritten as

$$|1 + 2a_0 - a_1\sqrt{2}| < 1.$$

The case k = 5 yields the same constraint if we only look at the real part of α . If we identify $m(n) = 1 + 2a_0$ and $n = a_1$ we immediately see that this yields the same solution as was given in the literature.

Now, we show that we can use the same method to obtain the systole of M_3 . As far as we know, this result was not yet known.

Example 8.13. Consider g = 3. By Lemma 8.11 we can write

$$\frac{1}{2}|\operatorname{tr}(M)| = |\operatorname{Re}(\alpha)| = |1 + 2a_0 + 2a_1 \cos(\frac{\pi}{6}) + 2a_2 \cos(\frac{\pi}{3})|,$$
$$= |1 + 2a_0 + a_1\sqrt{3} + a_2|.$$

By Lemma 8.8 we have $|\psi_k(\alpha)| < 1$ for k = 5 and k = 7. Again, when we only look at the real part of α , these constraints are identical. We have $\psi_5 : \zeta_{12} \mapsto \zeta_{12}^5$, so $\cos(\frac{\pi}{6}) \mapsto \cos(\frac{5\pi}{6})$, so $\sqrt{3} \mapsto -\sqrt{3}$. Then the condition $|\psi_5(\operatorname{Re}(\alpha))| < 1$ can be rewritten as

$$|1 + 2a_0 - a_1\sqrt{3} + a_2| < 1.$$

To satisfy this constraint, $1 + 2a_0 + a_2$ and $a_1\sqrt{3}$ should have the same sign. The case $|\operatorname{Re}(\alpha)| = 0$ is not interesting, so without loss of generality $1 + 2a_0 + a_2 > 0$ and $a_1 > 0$. The minimum is obtained for $1+2a_0+a_2 = a_1 = 1$. Therefore, $\cosh(\operatorname{syst}(M_3)/2) = 1+\sqrt{3}$.

The last example is meant to show that in general it is not sufficient to look only at the real part of α .

Example 8.14. Consider g = 6. By Lemma 8.11 we can write

$$\begin{aligned} \frac{1}{2} |\operatorname{tr}(M)| &= |\operatorname{Re}(\alpha)|, \\ &= |1 + 2a_0 + 2a_1 \cos(\frac{\pi}{12}) + 2a_2 \cos(\frac{\pi}{6}) + 2a_3 \cos(\frac{\pi}{4}) + 2a_4 \cos(\frac{\pi}{3}) + 2a_5 \cos(\frac{5\pi}{12})|, \\ &= |1 + 2a_0 + \frac{1}{2}a_1(\sqrt{6} + \sqrt{2}) + a_2\sqrt{3} + a_3\sqrt{2} + a_4 + \frac{1}{2}a_5(\sqrt{6} - \sqrt{2})|. \end{aligned}$$

By Lemma 8.8 we have $|\psi_k(\alpha)| < 1$ for k = 7, 11, 13, 17. When we only look at the real part of α , we only have to consider k = 7, 11. We have $\psi_7 : \zeta_{24} \mapsto \zeta_{24}^7$, so $\cos(\frac{\pi}{12}) \mapsto \cos(\frac{7\pi}{12})$, so $\frac{1}{4}(\sqrt{6} + \sqrt{2}) \mapsto \frac{1}{4}(-\sqrt{6} + \sqrt{2})$, which corresponds to the automorphism $\sqrt{3} \mapsto -\sqrt{3}$. In a similar way, ψ_{11} corresponds to $\sqrt{2} \mapsto -\sqrt{2}$. Consider $\operatorname{Re}(\alpha) = 1 + \frac{1}{2}(\sqrt{6} + \sqrt{2}) - \sqrt{2}$. Clearly, $1 < \operatorname{Re}(\alpha) < 1 + 2\cos(\frac{\pi}{12})$ and

$$|\psi_7(\operatorname{Re}(\alpha))| = |1 + \frac{1}{2}(-\sqrt{6} + \sqrt{2}) - \sqrt{2}| < 1,$$

$$|\psi_{11}(\operatorname{Re}(\alpha))| = |1 + \frac{1}{2}(\sqrt{6} - \sqrt{2}) + \sqrt{2}| < 1.$$

Hence, we have a solution with $|\operatorname{Re}(\alpha)|$ smaller than our proposed value of the systole. This does not show that our conjecture is false; it shows that it is not sufficient to only look at the real value of α .

9 Lower bound for the number of points in a Delaunay triangulation

The main result of this section is the following proposition.

Proposition 9.1. Let M be a hyperbolic surface of genus $g \ge 2$. Let P be a set of points in M such that $syst(M) > 2\delta_P$. Then

$$|P| > \left(\frac{\pi}{\pi - 6\operatorname{arccot}(\sqrt{3}\operatorname{cosh}(\frac{1}{4}\operatorname{syst}(M)))} - 1\right) \cdot 2(g-1).$$

Proof. Let |P| = v. Compute a Delaunay triangulation of P in M with e edges and f faces. We know that 3f = 2e, since every triangle consists of three edges and every edge belongs to two triangles. By Euler's formula,

$$v - e + f = 2 - 2g,$$

 \mathbf{SO}

$$f = 4g - 4 + 2v.$$

Denote the triangles in the Delaunay triangulation by F_i , their circumscribed disks by C_i and the diameters of these disks by d_i . By Corollary 4.4 we have

$$\operatorname{area}(F_i) \le \pi - 6 \operatorname{arccot}(\sqrt{3} \cosh(\frac{1}{2}d_i)).$$

Since all C_i do not have any points of P in their interior due to the triangulation being Delaunay, we have $d_i < \delta_P < \frac{1}{2} \operatorname{syst}(M)$. It follows that

$$\operatorname{area}(M) = \sum_{i=1}^{f} \operatorname{area}(F_i),$$

$$= \sum_{i=1}^{f} \pi - 6 \operatorname{arccot}(\sqrt{3} \operatorname{cosh}(\frac{1}{2}d_i)),$$

$$< \sum_{i=1}^{f} \pi - 6 \operatorname{arccot}(\sqrt{3} \operatorname{cosh}(\frac{1}{4} \operatorname{syst}(M))),$$

$$= f(\pi - 6 \operatorname{arccot}(\sqrt{3} \operatorname{cosh}(\frac{1}{4} \operatorname{syst}(M)))),$$

$$= (4g - 4 + 2v)(\pi - 6 \operatorname{arccot}(\sqrt{3} \operatorname{cosh}(\frac{1}{4} \operatorname{syst}(M))))$$

Observe that $\operatorname{area}(M) = 4\pi(g-1)$, so

$$v > \frac{2\pi(g-1)}{\pi - 6\operatorname{arccot}(\sqrt{3}\operatorname{cosh}(\frac{1}{4}\operatorname{syst}(M)))} - 2g + 2,$$

= $\left(\frac{\pi}{\pi - 6\operatorname{arccot}(\sqrt{3}\operatorname{cosh}(\frac{1}{4}\operatorname{syst}(M)))} - 1\right) \cdot 2(g-1).$

Г		
L		
L		

Remark 9.2. Proposition 9.1 gives the following lower bounds for the surfaces $M_g, 2 \le g \le 10$ corresponding to a regular 4g-gon:

Here we used the upper bound for the systole of these surfaces shown in the previous section.

Remark 9.3. In the previous remark, the lower bound grows approximately linearly as a function of g. More generally, since $\operatorname{arccot}(\sqrt{3} \cosh(\frac{1}{4} \operatorname{syst}(M)))$ is a decreasing function of $\operatorname{syst}(M)$ which tends to zero as $\operatorname{syst}(M) \to \infty$, we see that the coefficient

$$\frac{\pi}{\pi - 6\operatorname{arccot}(\sqrt{3}\operatorname{cosh}(\frac{1}{4}\operatorname{syst}(M)))} - 1$$

is a decreasing function of $\operatorname{syst}(M)$ as well, which also tends to zero as $\operatorname{syst}(M) \to \infty$. For sequences of surfaces with bounded systole, this coefficient is bounded away from zero, so in this case $|P| = \Omega(g)$, where we recall that the systole of a family of surfaces $\{M(g) : g \in \mathbb{N}, g \geq 2\}$ is bounded if the set $\{\operatorname{syst}(M(g)) : g \in \mathbb{N}, g \geq 2\}$ is bounded. For families of surfaces where the systole is not necessarily bounded, we only have the bound $\operatorname{syst}(M(g)) \leq 2\log(4g-2)$ (see Section 4.9). To see the behaviour of the above coefficient, define

$$f(x) = \operatorname{arccot}(\sqrt{3}\cosh(\frac{1}{4} \cdot 2\log(4x - 2))).$$

Then

$$\cot f(x) = \sqrt{3}\cosh(\frac{1}{2}\log(4x-2)).$$

For x sufficiently large, $\cosh(x) \approx \frac{1}{2}e^x$, so $\cot f(x) \approx \frac{1}{2}\sqrt{12x-6}$. On the other hand, for x sufficiently large $f(x) \approx 0$, so

$$\cot f(x) = \frac{\cos f(x)}{\sin f(x)} \approx \frac{1}{f(x)}.$$

It follows that for sufficiently large x,

$$f(x) \approx \frac{2}{\sqrt{12x - 6}}.$$

This means that

$$\begin{aligned} \frac{\pi}{\pi - 6 \operatorname{arccot}(\sqrt{3} \operatorname{cosh}(\frac{1}{4} \operatorname{syst}(M(g))))} &-1 \ge \frac{\pi}{\pi - 6 \operatorname{arccot}(\sqrt{3} \operatorname{cosh}(\frac{1}{2} \log(4g - 2)))} - 1, \\ &= \frac{\pi}{\pi - 6f(g)} - 1, \\ &\approx \frac{\pi\sqrt{12g - 6}}{\pi\sqrt{12g - 6} - 12} - \frac{\pi\sqrt{12g - 6} - 12}{\pi\sqrt{12g - 6} - 12}, \\ &= \frac{12}{\pi\sqrt{12g - 6} - 12}, \end{aligned}$$

with equality in the limit $g \to \infty$. We conclude that in general the number of points is of order $\Omega(\sqrt{g})$.

In [28] the bound

$$|P| \ge \frac{7 + \sqrt{1 + 48g}}{2}$$

is given, based on a purely combinatorial argument. This bound is in practice better, but asymptotically of the same order. However, our new bound is not redundant. Indeed, we have already argued that it is of order $\Omega(g)$ for families of surfaces with bounded systole. Secondly, to arrive at a complexity of $\Omega(\sqrt{g})$ we used the upper bound syst $(M(g)) \leq 2\log(4g-2)$. Currently there are no families of surfaces known which attain this bound, even asymptotically. As stated in Section 4.9, the current maximum of the systole of a family of surfaces is (up to constants) $\frac{4}{3}\log g$. Therefore, it makes sense to compute the order of the number of points with $\frac{4}{3}\log x$ instead of $2\log(4x-2)$ in the definition of fabove. Asymptotically we get in this case that $\cot f(x) \approx \frac{1}{2}\sqrt{3}x^{1/3}$, so $f(x) \approx \frac{2}{3}\sqrt{3}x^{-1/3}$. Then the number of points is of order $\Omega(g^{2/3})$, which is sharper than the bound given in [28].

10 Future work

Future work could follow several paths. First of all, a sharper lower bound for the minimum number of sheets for the hyperbolic surface $M = \mathbb{D}^2/\Gamma$ can be obtained by obtaining a sharper lower bound for δ_O , the diameter of the largest empty disk in \mathbb{D}^2 that does not contain any point of ΓO in its interior. In this way the lower bound for the asymptotic complexity of the general case may be increased.

Secondly, there are currently no upper bounds known for the number of sheets, with the exception of 128 for the Bolza surface, where the proof makes use of a computer. It would be interesting to find a simpler proof of an upper bound for the Bolza surface in particular and an upper bound for hyperbolic surfaces with different genera and conformal structures in general.

Thirdly, our conjecture on the systole of hyperbolic surfaces represented by regular 4g-gons remains to be proven.

Finally, to obtain an upper bound for the number of points of a dummy point set, one could explicitly construct a suitable point set. So far, this has only been done for the Bolza surface with a point set of cardinality 14.

11 Acknowledgements

I would like to thank my supervisor Gert Vegter for introducing me to this subject and for his guidance during the process. He gave me a sufficient amount of freedom to find out solutions on my own, while simultaneously keeping an eye on the general direction the research was going. Furthermore, I would like to thank my second supervisor Alef Sterk for evaluating my work. I would also like to thank Monique Teillaud for her feedback on my work during her visit to Groningen and Iordan Iordanov with his suggestions for making the figures in this thesis. Lastly, I would like to thank family and friends for their support.

A Algebra

In our proofs regarding the systole of surfaces corresponding to regular 4g-gons, we make use of some concepts from algebraic number theory. In this section we will give a brief introduction. For more details we refer to [22].

We start with the definition of a number field.

Definition A.1. A number field is a finite degree field extension of \mathbb{Q} .

We are primarily interested in cyclotomic fields.

Definition A.2. The n-th cyclotomic field is the number field obtained by adjoining a primitive n-th root of unity to \mathbb{Q} .

We denote the *n*-th cyclotomic field by $\mathbb{Q}(\zeta_n)$, where $\zeta_n = \exp(\frac{2\pi i}{n})$ is the primitive *n*-th root of unity with the smallest positive argument.

Algebraic integers generalize the idea of integers.

Definition A.3. Let $\alpha \in \mathbb{C}$. Then α is called an algebraic integer if there exists a monic polynomial $f \in \mathbb{Z}[X]$ such that $f(\alpha) = 0$.

Note that $\mathbb{Z}[X]$ is the ring of polynomials in the indeterminate X with coefficients in \mathbb{Z} and $f \in \mathbb{Z}[X]$ is called monic if its leading coefficient is 1. We immediately see that all integers are algebraic integers, since for $n \in \mathbb{Z}$ we have f(n) = 0 for f = X - n. The set of all algebraic integers is denoted by \mathcal{O} . Given a number field K, the set of algebraic integers contained in K is denoted by \mathcal{O}_K , i.e. $\mathcal{O}_K = \mathcal{O} \cap K$. This set is called the ring of integers of K. That this name makes sense follows from the following proposition.

Proposition A.4. The set $\mathcal{O}_K = \mathcal{O} \cap K$ of algebraic integers contained in a number field K forms a ring. In particular, if $\alpha, \beta \in \mathcal{O}_K$, then $\alpha + \beta, \alpha\beta \in \mathcal{O}_K$.

For a proof, see [22, I.2.4d].

Consider the cyclotomic field $K = \mathbb{Q}(\zeta_n)$. Since $f(\zeta_n) = 0$ for $f = X^n - 1$, we have $\zeta_n \in \mathcal{O}_K$. By the proposition above, this means that $\mathbb{Z}[\zeta_n] \subseteq \mathcal{O}_K$. In fact, we have equality.

Proposition A.5. The ring of integers \mathcal{O}_K of $K = \mathbb{Q}(\zeta_n)$ is given by $\mathbb{Z}[\zeta_n]$.

For a proof, we refer to [22, VI.1.Thm.46].

References

- Aline Aigon-Dupuy, Peter Buser, Michel Cibils, Alfred F Künzle, and Frank Steiner. Hyperbolic octagons and Teichmüller space in genus 2. *Journal of mathematical physics*, 46(3):033513, 2005.
- [2] Ralf Aurich, Eugene B. Bogomolny, and Franz Steiner. Periodic orbits on the regular hyperbolic octagon. *Physica D: Nonlinear Phenomena*, 48(1):91–101, 1991.
- [3] Ralf Aurich and Frank Steiner. On the periodic orbits of a strongly chaotic system. *Physica D: Nonlinear Phenomena*, 32(3):451–460, 1988.
- [4] Agnès Bachelot-Motet. Wave computation on the hyperbolic double doughnut. *Journal of Computational Mathematics*, pages 790–806, 2010.
- [5] N.L. Balazs and A. Voros. Chaos on the pseudosphere. *Physics reports*, 143(3):109–240, 1986.
- [6] Alan F. Beardon. The geometry of discrete groups, volume 91 of Graduate Texts in Mathematics. Springer-Verlag, 2012.
- [7] Mikhail V. Belolipetsky and Scott A. Thomson. Systoles of hyperbolic manifolds. *Algebraic & Geometric Topology*, 11(3):1455–1469, 2011.
- [8] M. Bogdanov, M. Teillaud, and G. Vegter. Delaunay triangulations on orientable surfaces of low genus. In Leibniz International Proceedings in Informatics, editor, 32nd International Symposium on Computational Geometry (SoCG 2016), pages 20:1–20:17, 2016.
- [9] Mikhail Bogdanov. Delaunay triangulations of spaces of constant negative curvature. PhD thesis, Université Nice Sophia Antipolis, 2013.
- [10] Mikhail Bogdanov and Monique Teillaud. Delaunay triangulations and cycles on closed hyperbolic surfaces. Technical Report RR-8434, INRIA, 2013.
- [11] Jean-Daniel Boissonnat and Mariette Yvinec. Algorithmic geometry. Cambridge University Press, 1998.
- [12] Roberto Bonola. Non-Euclidean geometry: A critical and historical study of its development. Courier Corporation, 1955.
- [13] Adrian Bowyer. Computing Dirichlet tessellations. *The Computer Journal*, 24(2):162–166, 1981.
- [14] Peter Buser. Geometry and spectra of compact Riemann surfaces. Springer-Verlag, 2010.
- [15] Peter Buser and Peter Sarnak. On the period matrix of a Riemann surface of large genus (with an Appendix by J.H. Conway and N.J.A. Sloane). *Inventiones mathematicae*, 117(1):27–56, 1994.

- [16] Manuel Caroli. Triangulating Point Sets in Orbit Spaces. PhD thesis, Université de Nice Sophia Antipolis, 2010.
- [17] Manuel Caroli and Monique Teillaud. Delaunay triangulations of point sets in closed Euclidean d-manifolds. In *Proceedings of the 27th Annual Symposium on Computational Geometry*, pages 274–282, 2011.
- [18] Pascal Chossat and Olivier Faugeras. Hyperbolic planforms in relation to visual edges and textures perception. *Plos Computational Biology*, 5(12):e1000625, 2009.
- [19] Harold Scott Macdonald Coxeter. Non-Euclidean geometry. Cambridge University Press, 1998.
- [20] Mark De Berg, Otfried Cheong, Marc Van Kreveld, and Mark Overmars. Computational Geometry: Introduction. Springer, 2008.
- [21] Grégory Faye, Pascal Chossat, and Olivier Faugeras. Analysis of a hyperbolic geometric model for visual texture perception. The Journal of Mathematical Neuroscience, 1(1):1–51, 2011.
- [22] Albrecht Fröhlich, Martin J Taylor, and Martin J Taylor. Algebraic number theory, volume 27 of Cambridge studies in advanced mathematics. Cambridge University Press, 1993.
- [23] William Fulton. Algebraic topology: a first course, volume 153 of Graduate Texts in Mathematics. Springer-Verlag, 1995.
- [24] Ronald L Graham, Martin Grotschel, and László Lovász, editors. Handbook of combinatorics. Elsevier, 1995.
- [25] Thomas Little Heath et al. The thirteen books of Euclid's Elements. Courier Corporation, 1956.
- [26] D. Hilbert. Über Flächen von konstanter Gaußscher Krümmung. Transactions of the American Mathematical Society, 2:87–99, 1901.
- [27] Yoichi Imayoshi and Masahiko Taniguchi. An introduction to Teichmüller spaces. Springer-Verlag, 2012.
- [28] Mark Jungerman and Gerhard Ringel. Minimal triangulations on orientable surfaces. Acta Mathematica, 145(1):121–154, 1980.
- [29] Svetlana Katok. *Fuchsian groups*. University of Chicago Press, 1992.
- [30] Mikhail G Katz, Mary Schaps, Uzi Vishne, et al. Logarithmic growth of systole of arithmetic Riemann surfaces along congruence subgroups. *Journal of Differential Geometry*, 76(3):399–422, 2007.
- [31] William C. Keel. The road to galaxy formation. Springer-Verlag, 2007.
- [32] John Lee. Introduction to topological manifolds, volume 940 of Graduate Texts in Mathematics. Springer-Verlag, 2010.

- [33] Shotaro Makisumi. A note on Riemann surfaces of large systole, 2012. arXiv:1206.2965 [math.DG].
- [34] John W Milnor. Hyperbolic geometry: the first 150 years. Bulletin of the American Mathematical Society, 6(1):9–24, 1982.
- [35] M. Näätänen. Regular n-gons and Fuchsian groups. Ann. of the Finn. Acad. of Sci., Ser. A, I, 7:291–300, 1982.
- [36] John Ratcliffe. Foundations of hyperbolic manifolds, volume 149 of Graduate Texts in Mathematics. Springer-Verlag, 2006.
- [37] William F. Reynolds. Hyperbolic geometry on a hyperboloid. The American mathematical monthly, 100(5):442–455, 1993.
- [38] Guodong Rong, Miao Jin, and Xiaohu Guo. Hyperbolic centroidal Voronoi tessellation. In Proceedings of the 14th ACM Symposium on Solid and Physical Modeling, pages 117–126, 2010.
- [39] François Sausset, Gilles Tarjus, and David R Nelson. Structure and dynamics of topological defects in a glassy liquid on a negatively curved manifold. *Physical Review E*, 81(3):031504, 2010.
- [40] François Sausset, Gilles Tarjus, and Pascal Viot. Tuning the fragility of a glassforming liquid by curving space. *Physical review letters*, 101(15):155701, 2008.
- [41] Charles Schmit. Quantum and classical properties of some billiards on the hyperbolic plane. *Chaos and Quantum Physics*, pages 335–369, 1991.
- [42] Mika Seppälä and Tuomas Sorvali. Geometry of Riemann surfaces and Teichmüller spaces, volume 169 of North Holland Mathematics Studies. Elsevier, 2011.
- [43] Kenneth Stephenson. Introduction to circle packing: The theory of discrete analytic functions. Cambridge University Press, 2005.
- [44] John Stillwell. *Geometry of surfaces*. Springer-Verlag, 1992.
- [45] John Stillwell. Sources of hyperbolic geometry, volume 10 of History of Mathematics. American Mathematical Society, 1996.
- [46] L. Fejes Tóth. Kreisausfüllungen der hyperbolischen ebene. Acta Mathematica Hungarica, 4(1-2):103–110, 1953.
- [47] Weston Ungemach. A bound on length isospectral families of hyperbolic surfaces, 2013. arXiv:1310.8276 [math.MG].
- [48] Charles Walkden. Hyperbolic Geometry. Lecture notes, 2012.
- [49] David F. Watson. Computing the n-dimensional Delaunay tessellation with application to Voronoi polytopes. *The Computer Journal*, 24(2):167–172, 1981.

[50] Don Bernard Zagier. Equivariant Pontrjagin classes and applications to orbit spaces: applications of the G-signature theorem to transformation groups, symmetric products and number theory, volume 290 of Lecture notes in Mathematics. Springer-Verlag, 2006.