University of Groningen

Delaunay triangulations of hyperbolic surfaces

Ebbens, Yde Matthijs

DOI:
[10.33612/diss.181607575](https://doi.org/10.33612/diss.181607575)

IMPORTANT NOTE: You are advised to consult the publisher's version (publisher's PDF) if you wish to cite from it. Please check the document version below.

Document Version
Publisher's PDF, also known as Version of record

Publication date:
2021

[Link to publication in University of Groningen/UMCG research database](#)

Citation for published version (APA):
Ebbens, Y. M. (2021). *Delaunay triangulations of hyperbolic surfaces*. [Thesis fully internal (DIV), University of Groningen]. University of Groningen. <https://doi.org/10.33612/diss.181607575>

Copyright

Other than for strictly personal use, it is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license (like Creative Commons).

The publication may also be distributed here under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license. More information can be found on the University of Groningen website: <https://www.rug.nl/library/open-access/self-archiving-pure/taverne-amendment>.

Take-down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Downloaded from the University of Groningen/UMCG research database (Pure): <http://www.rug.nl/research/portal>. For technical reasons the number of authors shown on this cover page is limited to 10 maximum.

Delaunay triangulations of hyperbolic surfaces

Yde Matthijs Ebbens

PhD thesis, University of Groningen, the Netherlands

Copyright 2021 Y.M. Ebbens

Printed by: Gildeprint

rijksuniversiteit
 groningen

Delaunay triangulations of hyperbolic surfaces

Proefschrift

ter verkrijging van de graad van doctor aan de
Rijksuniversiteit Groningen
op gezag van de
rector magnificus prof. dr. C. Wijmenga
en volgens besluit van het College voor Promoties.

De openbare verdediging zal plaatsvinden op

maandag 18 oktober 2021 om 14:30 uur

door

Yde Matthijs Ebbens

geboren op 5 januari 1994
te Harlingen

Promotor

Prof. dr. G. Vegter

Copromotor

Dr. A.E. Sterk

Beoordelingscommissie

Prof. dr. X. Goaoc

Prof. dr. T. Müller

Prof. dr. J.-M. Schlenker

Contents

Contents	v
Introduction	ix
Background and motivation	ix
Previous work	xi
Statement of the problem	xiii
Contribution	xiv
Publications	xvii
1 Mathematical preliminaries	1
1.1 Hyperbolic geometry	1
1.1.1 History	1
1.1.2 The Poincaré disk	2
1.1.3 Trigonometry	3
1.1.4 Classification of Möbius transformations	5
1.1.5 Fuchsian groups	6
1.1.6 Hyperbolic surfaces	7
1.1.7 Generalized Bolza surfaces	8
1.1.8 Geodesics on a hyperbolic surface	11
1.1.9 Teichmüller space	14
1.2 Triangulations	20
1.2.1 Simplicial complexes	20
1.2.2 Delaunay triangulations	21
2 Systoles of hyperbolic surfaces	23
2.1 Introduction	23
2.2 Systole of generalized Bolza surfaces	24
2.2.1 Representation of a simple closed geodesic by a sequence of segments	24
2.2.2 Upper bound for the systole	26
2.2.3 Lower bound for the systole	26
2.3 Systole of hyperelliptic surfaces	32
2.3.1 Neighborhood of hyperelliptic surfaces	32

2.3.2	Systole of hyperelliptic surfaces in a neighborhood of bounded size	33
2.3.3	Trigonometry of admissible polygons	35
2.3.4	Line segments between sides of admissible polygons	39
2.3.5	Proof of Theorem 2.8	40
2.4	Word length of systoles	41
3	Computing Delaunay triangulations of generalized Bolza surfaces	45
3.1	Introduction	45
3.2	Computing Delaunay triangulations	47
3.2.1	Bowyer's algorithm in the Euclidean plane	47
3.2.2	Delaunay triangulations of points on hyperbolic surfaces . .	48
3.3	Bounds on the number of dummy points	50
3.4	Local validity condition	54
3.5	Computation of dummy point sets	57
3.5.1	Refinement algorithm	61
3.5.2	Symmetric algorithm	63
3.5.3	Structured algorithm	68
3.5.4	Experimental results for small genus	72
3.6	Data structure, predicates, and implementation	73
3.6.1	Canonical representatives	73
3.6.2	Data structure	76
3.6.3	Degree of predicates	79
3.6.4	Implementation and experimental results	81
3.7	Conclusion and open problems	84
4	Minimal Delaunay triangulations of hyperbolic surfaces	85
4.1	Introduction	85
4.2	Linear upper bound for the number of vertices of a minimal distance Delaunay triangulation	86
4.2.1	Distance Delaunay triangulations of hyperbolic cylinders . .	87
4.2.2	Constructing a distance Delaunay triangulation of a hyperbolic surface with few vertices	91
4.2.3	Proof of Theorem 4.2	98
4.3	Classes of hyperbolic surfaces attaining the order of the upper bound	99
4.4	Lower bound	105
A	Omitted proofs	109
A.1	Proofs omitted in Chapter 2	109
A.2	Proofs omitted in Chapter 3	116
A.3	Proofs omitted in Chapter 4	124
	Conclusion and open problems	129

<i>CONTENTS</i>	vii
Samenvatting	133
Acknowledgements	137
Bibliography	139

Introduction

Background and motivation

Triangulations are among the most important and well-studied objects in computational geometry. One of their applications is *meshing*, in which a continuous shape is approximated by subdividing it into cells with a discrete representation. This allows the use of computer algorithms to analyze the geometry of the shape or perform simulations.

A particular kind of triangulation called the *Delaunay triangulation* is characterized by the property that the circumscribed disk of every triangle does not contain any vertices in its interior. An example is shown in Figure 1. Due to their favorable properties for meshing, Delaunay triangulations are used for example in height interpolation in the modelling of terrain [27, Chapter 9].

Figure 1: Left: Delaunay triangulation of a point set in the Euclidean plane, together with three circumscribed circles. Right: Triangulation of the same point set that is not a Delaunay triangulation.

Delaunay triangulations were initially studied for point sets in the Euclidean plane and higher-dimensional Euclidean spaces, but a similar notion exists for Riemannian manifolds [53]. More specifically, in the last decade algorithms for computing Delaunay triangulations have been extended to point sets in *hyperbolic* spaces [22, 13, 44].

Hyperbolic geometry is an example of *non-Euclidean geometry*. It was introduced in the first half of the nineteenth century as a negative answer to the long-standing open problem whether Euclid's parallel postulate was a consequence of his other four postulates [56]. One of the models that is used to represent the hyperbolic plane is the Poincaré disk model (see Figure 2). In this model, the hyperbolic plane is represented by a disk of radius 1 and hyperbolic lines are given by diameters of the disk or circle segments that intersect the boundary circle orthogonally. The properties of the hyperbolic plane, and in particular its constant Gaussian curvature $K = -1$, allow its application in describing shapes or structures that, intuitively speaking, cannot be “flattened” in the Euclidean plane.

Figure 2: Hyperbolic lines in the Poincaré disk model: lines L_1 and L_3 are disjoint, lines L_2 and L_3 intersect within the Poincaré disk and lines L_1 and L_2 intersect at a point of the boundary circle, i.e., they are parallel.

In this thesis, we consider Delaunay triangulations of *hyperbolic surfaces*. Roughly speaking, a hyperbolic surface is a surface that locally “looks like” the hyperbolic plane. Every hyperbolic surface can be represented as a hyperbolic polygon together with a set of side-pairing transformations, i.e. isometries of the hyperbolic plane that pair sides of the polygon. For example, the hyperbolic octagon in Figure 3 together with the set of side-pairing transformations that pair opposite sides corresponds to a hyperbolic surface of genus 2 called the Bolza surface. We will discuss this and other representations of hyperbolic surfaces in more detail in Chapter 1.

The representation of hyperbolic surfaces mentioned above allows us to interpret triangulations of hyperbolic surfaces as periodic triangulations of infinite periodic point sets in the hyperbolic plane. In our case, the infinite periodic point set is the image of a set of points in the hyperbolic plane under the action of the group generated by the side-pairing transformations (see Figure 4). Under certain conditions on the hyperbolic polygon, this group is a so-called Fuchsian group and the infinite periodic point set is discrete. See Chapter 1 for more details. Furthermore, the triangulation of a point set on a hyperbolic surface corresponds

Figure 3: A representation of the Bolza surface, a hyperbolic surface of genus 2.

to a periodic triangulation of the corresponding infinite periodic point set in the hyperbolic plane.

Figure 4: Left: point set on (a representation of) the Bolza surface. Middle: the infinite periodic point set that is the image of the point set shown in the left figure under the action of the group generated by the side-pairing transformations. Right: a Delaunay triangulation of the infinite periodic point set shown in the middle figure.

Even though our point of view is primarily mathematical, let us mention that periodic point sets and hyperbolic surfaces find many applications in different fields of science, like materials science [58], cosmology [51, 6], neuromathematics [26, 36] and quantum chaos theory [8, 68].

Previous work

In previous work, Delaunay triangulations of hyperbolic surfaces were mostly studied from an algorithmic point of view. For instance, Bowyer’s incremental algorithm [19] for computing Delaunay triangulations of point sets in the Euclidean plane was generalized to hyperbolic surfaces and implemented for the Bolza surface [15, 16, 45]. Lawson’s flip algorithm [52] has also been shown to generalize

on hyperbolic surfaces [30]. Here, we restrict our discussion to the former, since experience gained in the CGAL project has shown that it leads to cleaner code that is easier to maintain (see also Chapter 3).

Let us illustrate the idea of Bowyer's algorithm by considering a set of points in the Euclidean plane (see Figure 5). Bowyer's algorithm inserts the points one by one and after each insertion the triangulation is updated. To update the triangulation, we first delete the triangles that contain the inserted point in their circumscribed disk. The triangulation is then repaired by forming new triangles with the inserted point as a vertex.

Figure 5: Insertion of a point in a Delaunay triangulation with Bowyer's incremental algorithm.

A necessary condition for Bowyer's algorithm to work is that the triangulation is a *simplicial complex*. In our context, that means that it should not contain any loops or double edges. In a triangulation in the Euclidean plane where edges are given by straight lines this condition is automatically satisfied, but this is not necessarily the case for hyperbolic surfaces, as Figure 6 illustrates.

Figure 6: Surface of genus 2 with two triangles containing two double edges (top left) and a triangle containing a loop and a double edge (bottom right).

Observe that the loop in Figure 6 is a non-contractible closed curve on the surface. The length of the shortest non-contractible closed curve on a hyperbolic surface (which is a closed geodesic) is called the *systole* of that surface. If the length of every edge in a triangulation is shorter than the systole, then the triangulation does not contain any loops. Adding a factor $\frac{1}{2}$, i.e., requiring that the length of every edge is shorter than half the systole, prevents double edges as well. This leads to the so-called *validity condition*. We note that, given a set of points, the

validity condition is usually stated as a condition on the diameter of the largest empty disk instead of the length of edges. We refer to Chapter 3 for more details.

The validity condition is satisfied when the input point set is sufficiently large and well-distributed. However, to make sure that the validity condition is satisfied during all insertions as well, the algorithm is initialized by computing the Delaunay triangulation of an appropriate (but small) set of points that satisfies the validity condition by construction. These points are called *dummy points*. After sufficiently many and well-distributed input points have been inserted, the validity condition is satisfied without the dummy points, which can then be removed.

Statement of the problem

In this thesis we will study Delaunay triangulations of hyperbolic surfaces from two points of view. First, we discuss how to extend the results for the Bolza surface [45] to a larger class of hyperbolic surfaces so that we can compute Delaunay triangulations of point sets on these surfaces using Bowyer's algorithm. More specifically, we will look at the so-called *generalized Bolza surfaces*. As mentioned before, the Bolza surface can be represented by a regular hyperbolic octagon where opposite sides are paired. Similarly, the generalized Bolza surface of genus $g \geq 2$ can be represented by a regular hyperbolic $4g$ -gon where opposite sides are paired (see Figure 7).

Figure 7: A representation of the Bolza surface (left) and the generalized Bolza surface of genus 3 (right).

Even though there is a generalization of Bowyer's algorithm to arbitrary hyperbolic surfaces in the literature [15], the validity condition in this generalization depends on the value of the systole, which is not known in general. There exists an algorithm to compute the systole of a given hyperbolic surface [2], but since the complexity of this algorithm has not been analyzed, it is not clear whether it can be used in practice. Therefore, one part of our extension of Bowyer's algorithm to

the generalized Bolza surfaces will be to find a method for computing the systole of these surfaces.

Our second point of view is more theoretical and concerns the minimal number of vertices of Delaunay triangulations of hyperbolic surfaces in terms of their genus. To the extent of our knowledge, this problem has not been studied before in this context, even though the more general combinatorial problem of finding the minimal number of vertices of a simplicial triangulation of a *topological* surface has been studied extensively [48].

Contribution

After providing an introduction to hyperbolic geometry and triangulations in Chapter 1, we present the contribution of this thesis in three chapters.

In Chapter 2 we discuss a number of results on the systole of hyperbolic surfaces. First, we show that the value of the systole of the generalized Bolza surface of genus g is given by $2 \operatorname{arccosh}(1 + 2 \cos(\frac{\pi}{2g}))$. To illustrate the idea of the proof, consider the two oriented line segments in Figure 8 indicated by the arrows. Be-

Figure 8: The sequence of two oriented hyperbolic line segments indicated by the black arrows is a representation of a closed geodesic on the Bolza surface. The starting point of each arrow is identified with the endpoint of the other arrow through the transformations that pair opposite sides.

cause opposite sides of the octagon are paired, the head of each of the arrows is identified with the tail of the other arrow. Therefore, the union of these two oriented line segments represents a closed curve on the Bolza surface. More generally, we show that every closed geodesic on the generalized Bolza surface of genus g can be represented in this way as a sequence of oriented hyperbolic line segments

between the sides of the corresponding hyperbolic $4g$ -gon. Then, the result follows by obtaining bounds on the lengths of these line segments.

Second, we discuss the question whether our method for computing the systole of generalized Bolza surfaces can also be applied to different hyperbolic surfaces. To do this, we slightly perturb the vertices of the regular hyperbolic $4g$ -gon representing the generalized Bolza surface of genus g (see Figure 9) and look at the corresponding hyperbolic surface obtained by pairing opposite sides. The surfaces that we construct in this way are so-called *hyperelliptic surfaces* in a neighborhood of the generalized Bolza surfaces in the corresponding Teichmüller space. For sufficiently small perturbations, the geometry of the polygons is quite similar, so it seems reasonable that the geometry of the corresponding hyperbolic surfaces is similar as well. Indeed, we show that for sufficiently small perturbations our method for computing the systole of the generalized Bolza surfaces also works for the hyperbolic surfaces resulting from this construction.

Figure 9: The Bolza surface (black) with a hyperbolic surface (red) obtained from slightly perturbing the vertices of the regular octagon.

Third, motivated by our computations of the systole, we look at the so-called word length of systoles. It is known that every closed geodesic on a closed hyperbolic surface corresponds to an element of the group generated by the side-pairing transformations, up to conjugacy. The minimal number of generators of a group that need to be multiplied to obtain a specific element of that group is called the word length of that element. The elements corresponding to a systole of a generalized Bolza surface are products of precisely two of the side-pairing transformations. This raised the question whether there exists a general upper bound for the word length of the elements corresponding to systoles of a given hyperbolic surface, but we proved that there is no such bound.

In Chapter 3 we describe the properties of the generalized Bolza surfaces that allow Bowyer's algorithm for computing Delaunay triangulations to be applied to these surfaces. First, to be able to check whether a given point set on a generalized Bolza surface satisfies the validity condition, it is necessary to know the value of the systole, which was already computed in Chapter 2.

Second, recall that triangulations of a hyperbolic surface can be seen as periodic triangulations of infinite periodic point sets in the hyperbolic plane, as was illustrated in Figure 4. In practice, the Delaunay triangulation of a point set on the surface is usually obtained by computing a Delaunay triangulation of a corresponding point set in the hyperbolic plane and projecting the result to the hyperbolic surface. However, it is not possible to compute a triangulation of an infinite set of points. Therefore, we show that during the execution of the algorithm it is sufficient to consider the Delaunay triangulation of only a small set of copies of the input points (see Figure 10).

Figure 10: To compute a Delaunay triangulation of a point set on the Bolza surface, it is sufficient to consider the copies of the input points that lie in the shaded area. This area consists of the regular hyperbolic octagon used to represent the Bolza surface, together with all images of this hyperbolic octagon that are adjacent to it under the action of a side-pairing transformation.

Third, we mentioned that a set of dummy points is used to make sure that the validity condition is satisfied throughout the execution of Bowyer's algorithm. Therefore, we create several algorithms to construct dummy point sets for the generalized Bolza surfaces and analyze the resulting number of points of each. More generally, we prove upper and lower bounds for the cardinality of a dummy point set for *arbitrary* hyperbolic surfaces.

One disadvantage of the validity condition from the literature [15] is that it uses the value of the systole as criterion to determine whether a point set is suffi-

ciently dense and well-distributed. If a hyperbolic surface has only one short closed geodesic and the remainder of the surface is relatively “thick”, then a dummy point set satisfying the validity condition might be denser than necessary in the parts of the surface far away from the short closed geodesic. Therefore, we will also briefly discuss a *local* validity condition. The local validity condition does not depend on the value of the globally shortest non-contractible curve, i.e., the systole, but instead considers for each point on the surface the shortest non-contractible curve (which is a geodesic loop) passing through that point.

Finally, in a preliminary implementation of the described algorithm [44] it turns out that our extension of Bowyer’s algorithm to the generalized Bolza surfaces is computationally quite expensive for genus $g \geq 3$. We discuss the occurring numerical issues and derive a bound for the degree of the predicates used in the computations.

In Chapter 4 we discuss the minimal number of vertices of Delaunay triangulations of hyperbolic surfaces. It is known that the minimal number of vertices of a simplicial triangulation of a topological surface is $\Theta(\sqrt{g})$ [48]. We show that this lower bound is tight for Delaunay triangulations of hyperbolic surfaces as well. Moreover, we show that every hyperbolic surface of genus $g \geq 2$ has a Delaunay triangulation with at most $151g$ vertices. Finally, we construct a class of hyperbolic surfaces for which the number of vertices of any Delaunay triangulation is $\Omega(g)$, which implies that the $O(g)$ upper bound is optimal.

Publications

Several parts of this thesis have previously appeared as conference papers or preprints:

M. Ebbens, I. Iordanov, M. Teillaud, G. Vegter. Delaunay triangulations of generalized Bolza surfaces. Extended abstract in *35th European Workshop on Computational Geometry*, 15:1-15:8, 2019. <http://www.eurocg2019.uu.nl/papers/15.pdf>. Full version accepted up to revision at *Journal of Computational Geometry*. Preprint: arXiv:2103.05960.
(Section 2.2 of Chapter 2 and Chapter 3)

M. Ebbens, H. Parlier, G. Vegter. Minimal Delaunay triangulations of hyperbolic surfaces. Extended abstract in *Proceedings of the 37th International Symposium on Computational Geometry*, 31:1-31-16, 2021. doi:10.4230/LIPIcs.SoCG.2021.31. Full version accepted up to revision at *Discrete and Computational Geometry*. Preprint: arXiv:2011.09847.
(Chapter 4)

Chapter 1

Mathematical preliminaries

1.1 Hyperbolic geometry

1.1.1 History

Around 300 B.C. Euclid wrote in the first book of his Elements¹:

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

¹The translation is mine. For a translation of the complete work, see [40].

²The following discussion is primarily based on [56], but we refer to [18] for a more detailed treatise. For an extensive bibliography on the history of non-Euclidean geometry, see [64, 33ff.].

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 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 1.1.2 we will discuss a different model of hyperbolic geometry: the Poincaré disk model. This model can be obtained from the hemisphere in the above construction 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.

1.1.2 The Poincaré disk

The model of the hyperbolic plane that we use in this thesis is the Poincaré disk [10, 20]. This model is given by the open unit disk $\mathbb{D}$ in the complex plane

equipped with the Riemannian metric

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

of constant Gaussian curvature $K = -1$. The Euclidean boundary $\mathbb{D}_\infty$ of the unit disk consists of the points at infinity or *ideal points* of the hyperbolic plane (which do not belong to $\mathbb{D}$). The geodesic segment $[z, w]$ between points $z, w \in \mathbb{D}$ is the (unique) shortest curve connecting z and w . A hyperbolic line (i.e., a geodesic for the given metric) in this model is a curve which contains the geodesic segment between any two of its points. These geodesics are diameters of $\mathbb{D}$ or circle arcs whose supporting lines or circles intersect $\mathbb{D}_\infty$ orthogonally (see the leftmost frame of Figure 1.1). A circle in the hyperbolic plane is a Euclidean circle in the Poincaré disk, in general with a hyperbolic center and radius that are different from their Euclidean counterparts.

Figure 1.1: Left: the Poincaré disk model $\mathbb{D}$ of the hyperbolic plane, with some geodesics. The boundary $\mathbb{D}_\infty$ does not belong to $\mathbb{D}$, but consists of ideal points of $\mathbb{D}$. Geodesics L_0 and L_1 are parallel (have an ideal point in common), L_2 and L_3 are intersecting and L_4 is disjoint from the other geodesics. The points z and w are connected by a hyperbolic segment.

Right: A hyperbolic translation f has two fixed points on the boundary $\mathbb{D}_\infty$ of the Poincaré disk $\mathbb{D}$. The axis of f is the (unique) geodesic connecting the fixed points of f . The orbit of point w is contained in the axis of f . The orbit of point z , which does not lie on the axis of f , is contained in an equidistant of the axis (an arc of a Euclidean circle through the fixed points). The red region containing the point z is mapped by f to the red region containing $f(z)$.

1.1.3 Trigonometry

In this subsection we will briefly state some results about hyperbolic trigonometry [10, 49]. Since in both models the first four postulates of Euclid hold, in particular there exists for distinct z and w in the hyperbolic plane a unique geodesic

segment $[z, w]$ joining z to w . For distinct, non-collinear points 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, $[z_1, z_2, \dots, z_n] = [z_1, z_2] \cup \dots \cup [z_{n-1}, z_n] \cup [z_n, z_1]$ denotes the hyperbolic n -gon with vertices $z_1, \dots, z_n$. 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 1.2). 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

Figure 1.2: Hyperbolic triangle

terms of two of the sides are given by

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

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

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

In particular, $\text{area}(T) \leq \pi$ 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 $\alpha_1, \dots, \alpha_n$, then its area is given by

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

1.1.4 Classification of Möbius transformations

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

$$f(z) = \frac{az + b}{\bar{b}z + \bar{a}} \quad (1.1)$$

for $a, b \in \mathbb{C}$ with $|a|^2 - |b|^2 = 1$. Conversely, every map of this form is an orientation preserving isometry of $\mathbb{D}$. From now on, we will abbreviate orientation preserving isometry as *isometry*.

Every isometry of the hyperbolic plane is either elliptic, parabolic or hyperbolic. We distinguish between these three by the number and the location of their fixed points. If an isometry f of $\mathbb{D}$ has

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

In the last case, the geodesic connecting the two fixed points is called the axis X_f of f . The distance $d(x, f(x))$ is constant for all $x \in X_f$ and this constant is called the translation length of f , denoted by $\ell(f)$.

To a transformation of the form given in Equation (1.1) 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 $f(z) = (-az - b)/(-\bar{b}z - \bar{a})$ as well. We will use the same notation for the transformation itself and the corresponding matrix. Composition in the group of isometries is then given by matrix multiplication. Furthermore, it can be seen that the classification of isometries into elliptic,

parabolic and hyperbolic transformations corresponds to $|\operatorname{Tr}(f)| < 2$, $|\operatorname{Tr}(f)| = 2$ and $|\operatorname{Tr}(f)| > 2$ respectively. Here $|\operatorname{Tr}(f)|$ denotes the absolute value of the trace of (a matrix corresponding to) f , which is well defined, because f determines its corresponding matrix up to sign. An explicit formula for the translation length $\ell(f)$ is given by

$$\cosh\left(\frac{\ell(f)}{2}\right) = \frac{1}{2}|\operatorname{Tr}(f)|.$$

1.1.5 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 [10, 49, 74].

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 $\operatorname{Isom}^+(\mathbb{D})$ with matrices induces the structure of a topological space on $\operatorname{Isom}^+(\mathbb{D})$. We now give the definition of a Fuchsian group.

Definition 1.1. A discrete subgroup of $\operatorname{Isom}^+(\mathbb{D})$ is called a Fuchsian group.

For a Fuchsian group Γ and a point $x \in \mathbb{D}$, the orbit of x under Γ is defined as $\Gamma(x) = \{f(x) \mid f \in \Gamma\} \subset \mathbb{D}$.

Proposition 1.2. Let Γ be a subgroup of $\operatorname{Isom}^+(\mathbb{D})$. The following statements are equivalent:

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

If Γ satisfies the third statement, we usually say that Γ acts *properly discontinuously* on $\mathbb{D}$, even though definitions of properly discontinuous may vary in the literature. Denote the interior of a subset F of a topological space by $\mathring{F}$.

Definition 1.3. Given a Fuchsian group Γ , a fundamental domain F for Γ is a closed subset of $\mathbb{D}$ such that

1. $\bigcup_{f \in \Gamma} f(F) = \mathbb{D}$,
2. For all $f_1, f_2 \in \Gamma$ we have: if $f_1 \neq f_2$, then $f_1(\mathring{F}) \cap f_2(\mathring{F}) = \emptyset$.

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

Proposition 1.4. *Let Γ be a Fuchsian group with fundamental domains F, F' such that $\text{area}(\partial F) = \text{area}(\partial F') = 0$ and $\text{area}(F) < \infty$. Then*

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

It can be shown that for every Fuchsian group Γ there exists a point $\mathbf{p} \in \mathbb{D}$ such that $f(\mathbf{p}) \neq \mathbf{p}$ for all $f \in \Gamma \setminus \{1\}$, i.e. there exists a point that is not fixed by any non-trivial element of Γ . We now define the Dirichlet region.

Definition 1.5. Let Γ be a Fuchsian group and let $\mathbf{p}$ be a point that is not fixed by any non-trivial element of Γ . Define the Dirichlet region $D_{\mathbf{p}}(\Gamma)$ of $\mathbf{p}$ with respect to Γ as

$$D_{\mathbf{p}}(\Gamma) = \{\mathbf{x} \in \mathbb{D} \mid d(\mathbf{x}, \mathbf{p}) \leq d(\mathbf{x}, f(\mathbf{p})) \text{ for all } f \in \Gamma\}.$$

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

Proposition 1.6. *Let Γ be a Fuchsian group and let $\mathbf{p}$ be a point that is not fixed by any non-trivial element of Γ . The Dirichlet region $D_{\mathbf{p}}(\Gamma)$ is a fundamental domain for Γ . If $\text{area}(D_{\mathbf{p}}(\Gamma)) < \infty$, then $D_{\mathbf{p}}(\Gamma)$ is a convex hyperbolic polygon with finitely many sides.*

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

1.1.6 Hyperbolic surfaces

A Fuchsian group Γ naturally acts on $\mathbb{D}$, so we can form the quotient space $\mathbb{D}/\Gamma$. We saw in the previous section that we can associate a set of side pairing transformations to a Dirichlet region of Γ . These side pairing transformations can be seen as ‘glueing’ the Dirichlet region along paired sides. In this way we can see $\mathbb{D}/\Gamma$ 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 $\mathbb{D}/\Gamma$ for some Fuchsian group Γ . Another 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 [49, 74].

First we give the definition of hyperbolic surface.

Definition 1.7. A hyperbolic surface is a connected 2-dimensional manifold that is locally isometric to an open subset of $\mathbb{D}$.

We emphasize that points in the hyperbolic plane $\mathbb{D}$ will be denoted by $\mathbf{z}, \mathbf{w}, \mathbf{p}, \mathbf{q}$ and so on, whereas the corresponding points on the surface $\mathbb{D}/\Gamma$ are denoted by z, w, p, q and so on. 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 $\mathbb{R}^3$ [41]. 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 complete (as a metric space) and orientable. Hyperbolic surfaces can be obtained as a quotient space under the action of a Fuchsian group.

Proposition 1.8. *For every hyperbolic surface $\mathbb{M}$ there exists a Fuchsian group Γ acting on $\mathbb{D}$ without fixed points, such that $\mathbb{M}$ is isometric to $\mathbb{D}/\Gamma$.*

Since elliptic Möbius transformations have a fixed point in $\mathbb{D}$, a Fuchsian group as in the previous proposition does not contain any elliptic elements. A compact hyperbolic surface is called closed. A Fuchsian group is called cocompact, if $\mathbb{D}/\Gamma$ is compact. It can be shown that a cocompact Fuchsian group does not contain any parabolic elements [49].

Corollary 1.9. *For every closed hyperbolic surface $\mathbb{M}$ there exists a Fuchsian group Γ of which all non-trivial elements are hyperbolic, such that $\mathbb{M}$ is isometric to $\mathbb{D}/\Gamma$.*

A note on the literature: this section is mostly based on [70], since the main statement of this section is stated there explicitly. Works on Teichmüller spaces, such as [43, 69], 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 [20] gives a proof that the classifications of conformal structures on Riemann surfaces of genus $g \geq 2$ 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 $\mathbb{M}$ of genus $g \geq 2$ there exists by the Uniformization Theorem a universal covering map $\pi : \mathbb{D} \rightarrow \mathbb{M}$. The covering transformations are conformal self-mappings of $\mathbb{D}$, so the local inverses of π can be used as parametrizations for a hyperbolic surface. Beardon [10] 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 $\mathbb{D}/\Gamma$.

1.1.7 Generalized Bolza surfaces

The Fuchsian group Γ_g . The generalized Bolza group of genus g , $g \geq 2$, is the Fuchsian group Γ_g defined in the following way. Consider the regular hyperbolic $4g$ -gon D_g with angle-sum 2π . The counterclockwise sequence of vertices is

Figure 1.3: The side-pairings $f_0, \dots, f_3$ of the Bolza surface (of genus 2) pair opposite edges of the fundamental octagon (a regular octagon in $\mathbb{D}$ with angles $\frac{1}{4}\pi$). Their inverses $f_4, \dots, f_7$ satisfy $f_{k+4} = f_k^{-1}$. The side-pairings generate the Fuchsian group Γ_2 . All vertices are in the same Γ_2 -orbit. The composition $f_0 f_5 f_2 f_7 f_4 f_1 f_6 f_3$ is the identity $\mathbb{1} \in \Gamma_2$.

$\mathbf{v}_0, \dots, \mathbf{v}_{4g-1}$, where the midpoint of edge $[\mathbf{v}_0, \mathbf{v}_1]$ lies on the positive real axis. See Figure 1.3 for $g = 2$. The sides of D_g are $\mathbf{s}_j, j = 0, \dots, 4g - 1$, where $\mathbf{s}_j$ is the side with vertices $\mathbf{v}_j$ and $\mathbf{v}_{j+1}$ (counting indices modulo $4g$). The orientation preserving isometries $f_0, \dots, f_{4g-1}$ pair opposite sides of D_g . More precisely, f_j maps $\mathbf{s}_{j+2g}$ to $\mathbf{s}_j$, and $\mathbf{s}_j = f_j(D_g) \cap D_g$. According to (1.1) the side-pairing $f_j, j = 0, 1, \dots, 4g - 1$, is represented by any of two matrices $\pm A_j$ with determinant 1. Using some elementary hyperbolic geometry it can be seen that A_j is given by [5]

$$A_j = \begin{pmatrix} \cot(\frac{\pi}{4g}) & \exp(\frac{ij\pi}{2g})\sqrt{\cot^2(\frac{\pi}{4g}) - 1} \\ \exp(-\frac{ij\pi}{2g})\sqrt{\cot^2(\frac{\pi}{4g}) - 1} & \cot(\frac{\pi}{4g}) \end{pmatrix}. \quad (1.2)$$

By Poincaré's Theorem ([10, Chapter 9.8] and [64, Chapter 11.2]) these side-pairings generate a Fuchsian group, the generalized Bolza group Γ_g , all non-identity elements of which are hyperbolic translations. The polygon D_g is a fundamental domain for the action of this group, and it is even the Dirichlet region of the origin.

Since $\mathbf{v}_j = f_j f_{j+1}^{-1}(\mathbf{v}_{j+2})$, we see that the element $f_0 f_1^{-1} f_2 f_3^{-1} \dots f_{4g-2} f_{4g-1}^{-1}$ of Γ_g maps $\mathbf{v}_{4g}$ to $\mathbf{v}_0$. In other words, $\mathbf{v}_0$ is a fixed point of this element. Since all non-identity elements of Γ_g are hyperbolic translations, and, hence, without fixed points in $\mathbb{D}$, this element is the identity $\mathbb{1}$ of Γ_g :

$$f_0 f_1^{-1} f_2 f_3^{-1} \dots f_{4g-2} f_{4g-1}^{-1} = \mathbb{1}. \quad (1.3)$$

For even j we have $f_j = f_{j(2g+1)}$, since we are counting indices modulo $4g$. Simi-

larly, $f_j^{-1} = f_{j+2g} = f_{j(2g+1)}$, for odd j . Therefore, we can rewrite (1.3) as

$$\prod_{j=0}^{4g-1} f_{j(2g+1)} = f_0 f_{2g+1} f_{2(2g+1)} \cdots f_{(4g-1)(2g+1)} = \mathbb{1}. \quad (1.4)$$

The order of the factors in this product does matter since the group Γ is not abelian. Equation (1.4) is usually called the *relation* of Γ_g . In addition to (1.3) and (1.4), there are many other ways to write the relation. By rotational symmetry of D_g , conjugating $\prod_{j=0}^{4g-1} f_{j(2g+1)}$ with the rotation by angle $k\pi/2g$ around the origin yields the relation $\prod_{j=0}^{4g-1} f_{k+j(2g+1)} = \mathbb{1}$. The latter expression can be rewritten as

$$f_k f_{k+1}^{-1} f_{k+2} f_{k+3}^{-1} \cdots f_{k+4g-2} f_{k+4g-1}^{-1} = \mathbb{1}. \quad (1.5)$$

Neighbors of vertices of the fundamental polygon. In the clockwise sequence of Dirichlet regions $h_1(D_g), h_2(D_g), \dots, h_{4g}(D_g)$ around vertex $\mathbf{v}_k$ the element $h_j \in \Gamma_g$ is the prefix of length j in the left-hand side of (1.5):

$$h_j = \begin{cases} f_k f_{k+1}^{-1} \cdots f_{k+j-2} f_{k+j-1}^{-1}, & \text{if } j \text{ is even,} \\ f_k f_{k+1}^{-1} \cdots f_{k+j-2}^{-1} f_{k+j-1}, & \text{if } j \text{ is odd.} \end{cases} \quad (1.6)$$

Prefixes h_j of length $j \geq 2g$ can be reduced to a word of length $4g - j$ in $f_k, f_{k+1}^{-1}, \dots, f_{k+2g-1}^{-1}$ using relation (1.5) (where the empty word $-$ of length zero $-$ corresponds to $\mathbb{1}$) and the fact that $f_j = f_{j-2g}^{-1}$ for $j \geq 2g$. More precisely, h_{4g-j} is the prefix of length j in $f_{k+2g-1}^{-1} f_{k+2g-2} \cdots f_{k+1}^{-1} f_k$, for $j = 0, \dots, 2g$. Figure 1.4 depicts the neighbors of $\mathbf{v}_k$ for the case $g = 2$.

Figure 1.4: Enumeration of the regions around a vertex $\mathbf{v}_k$, $k = 0, 1, \dots, 7$, for the case $g = 2$. Note that $f_{k+3}^{-1} f_{k+2} f_{k+1}^{-1} f_k = f_k f_{k+1}^{-1} f_{k+2} f_{k+3}^{-1}$ by the group relation.

The hyperbolic surface $\mathbb{M}_g$. The *generalized Bolza surface* of genus g is the hyperbolic surface $\mathbb{D}/\Gamma_g$, denoted by $\mathbb{M}_g$. The projection map is $\pi_g : \mathbb{D} \rightarrow \mathbb{D}/\Gamma_g$. The surface $\mathbb{M}_2$ is the classical Bolza surface [17, 4].

1.1.8 Geodesics on a hyperbolic surface

In this section we will discuss geodesics, the systole and the length spectrum of hyperbolic surfaces, following primarily [64, 70]. The relevant homotopy theory can be found in an introduction on algebraic topology; see, e.g., [37].

Let $\mathbb{M} = \mathbb{D}/\Gamma$ be a closed hyperbolic surface with projection $\pi : \mathbb{D} \rightarrow \mathbb{M}$. By using the metric on $\mathbb{D}$ and the fact that Γ consists of isometries, we obtain a metric on $\mathbb{M}$, so we can speak of geodesics on $\mathbb{M}$. We define triangles and circles on a hyperbolic surface in a similar way. A *triangle* t on $\mathbb{M}$ is the π -image of a triangle $\mathbf{t}$ in $\mathbb{D}$ such that π is injective on $\mathbf{t}$. Clearly, the vertices of t are the projections of the vertices of $\mathbf{t}$ and the edges of t are geodesic segments. A *circle* on $\mathbb{M}$ is the π -image of a circle in the hyperbolic plane. In this case, we do not require π to be injective on the circle, so the image may have self-intersections.

Next, we consider closed geodesics. Define the circle $\mathbb{S}^1 = \mathbb{R}/\sim$, where $x \sim x+1$ for all $x \in \mathbb{R}$. A closed curve on $\mathbb{M}$ is a continuous map $\gamma : \mathbb{S}^1 \rightarrow \mathbb{M}$. We will always assume differentiability, except maybe at a point: a curve $\gamma : [0, 1] \rightarrow \mathbb{M}$ such that $\gamma(0) = \gamma(1)$ which is differentiable at $(0, 1)$ is called a loop. Closed curves $\gamma_0, \gamma_1 : \mathbb{S}^1 \rightarrow \mathbb{M}$ are called freely homotopic if there exists a continuous map $H : \mathbb{S}^1 \times [0, 1] \rightarrow \mathbb{M}$ such that

$$H(x, 0) = \gamma_0(x), \quad H(x, 1) = \gamma_1(x)$$

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 γ and the closed geodesic γ^2 obtained by traversing γ twice.

Closed geodesics are closely related to the structure of the corresponding Fuchsian group.

Proposition 1.10 ([64, Thm. 9.6.2]). *Closed geodesics of a closed hyperbolic surface $\mathbb{M} = \mathbb{D}/\Gamma$ are in one-to-one correspondence with conjugacy classes of elements of Γ .*

In the above correspondence, the projection $\pi : \mathbb{D} \rightarrow \mathbb{D}/\Gamma$ maps (oriented) geodesics of $\mathbb{D}$ to (oriented) geodesics of $\mathbb{M} = \mathbb{D}/\Gamma$, and it maps the axis of a hyperbolic translation $f \in \Gamma$ to a *closed* geodesic of $\mathbb{M}$. Every (oriented) closed geodesic γ on $\mathbb{M}$ arises in this way, i.e., there is a hyperbolic translation $f \in \Gamma$ such that γ *lifts* to the axis of f . The axes of two hyperbolic translations $f, f' \in \Gamma$ project to the same closed geodesic of $\mathbb{M}$ if and only if f' is conjugate to f in Γ (i.e., iff there is an $h \in \Gamma$ such that $f' = h^{-1}fh$). It follows that the length of γ is given by the distance $d(x, f(x))$ for $x \in X_f$, as this is the distance that f moves

x along the axis X_f until it reaches the next point in the orbit. Therefore, the length of γ is equal to the translation length $\ell(f)$.

Around a simple closed geodesic γ , the local geometry of a surface is given by its so-called collar. Roughly speaking, for small enough r , the set

$$C_\gamma(r) = \{x \in \mathbb{M} \mid d(x, \gamma) \leq r\}$$

is an embedded cylinder. A bound on how large one can take the r to be while retaining the cylinder topology is given by the Collar Lemma:

Lemma 1.11 ([20, Theorem 4.1.1]). *Let γ be a simple closed geodesic on a closed hyperbolic surface $\mathbb{M}$. The collar $C_\gamma(w(\gamma))$ of width $w(\gamma)$ given by*

$$w(\gamma) = \operatorname{arcsinh} \left(\frac{1}{\sinh(\frac{1}{2}\ell(\gamma))} \right) \quad (1.7)$$

is an embedded hyperbolic cylinder isometric to $[-w(\gamma), w(\gamma)] \times \Gamma^1$ with the Riemannian metric $ds^2 = d\rho^2 + \ell(\gamma)^2 \cosh^2(\rho) dt^2$ at (ρ, t) . Furthermore, if two simple closed geodesics γ and γ' are disjoint, then the collars $C_\gamma(w(\gamma))$ and $C_{\gamma'}(w(\gamma'))$ are disjoint as well.

In the algorithm for computing Delaunay triangulations of hyperbolic surfaces that we will discuss in Chapter 3, a central role is played by the systole of a surface.

Definition 1.12. The length of the shortest homotopically non-trivial closed curve on a closed hyperbolic surface $\mathbb{M}$ is called the systole of $\mathbb{M}$ and denoted by $\operatorname{sys}(\mathbb{M})$.

In Chapter 2 we will state a number of results on systoles, among which the exact value of the systole of generalized Bolza surfaces. The systole of a surface is attained as the length of some closed curve. Clearly, the shortest closed curves on $\mathbb{M}$ 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 1.13 ([64, Thm. 9.6.4 & 9.6.5]). *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.*

The following proposition gives an upper bound for the systole.

Proposition 1.14 ([20, Theorem 5.2.1]). *Let $\mathbb{M}$ be a closed hyperbolic surface of genus $g \geq 2$. Then*

$$\operatorname{sys}(\mathbb{M}) \leq 2 \log(4g - 2).$$

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

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

Taking the limit $r \rightarrow \text{sys}(\mathbb{M})/2$ we obtain

$$\cosh(\text{sys}(\mathbb{M})/2) \leq 2g - 1.$$

Since $\frac{1}{2} \exp(\text{sys}(\mathbb{M})/2) < \cosh(\text{sys}(\mathbb{M})/2)$, it follows that

$$\exp(\text{sys}(\mathbb{M})/2) < 4g - 2,$$

which proves the result. $\square$

The above upper bound is up to a constant the best possible. Namely, in [21] a family of hyperbolic surfaces $\mathbb{M}(g)$ is constructed with genus $g \rightarrow \infty$ and

$$\text{sys}(\mathbb{M}(g)) \geq \frac{4}{3} \log g - c$$

for some constant c . The same bound is found in [50] 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 [49]. In [55] 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 [11] it is shown³ that for any $\varepsilon > 0$ there exists a hyperbolic surface $\mathbb{M}$ with $\text{sys}(\mathbb{M}) < \varepsilon$.

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

Definition 1.15. The length spectrum $\mathcal{L}(\mathbb{M})$ of a closed hyperbolic surface $\mathbb{M}$ is the ascendingly ordered sequence of lengths of closed geodesics on $\mathbb{M}$.

In light of the discussion before, the length spectrum of a hyperbolic surface $\mathbb{M} = \mathbb{D}/\Gamma$ 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 [73], upper bounds are given for the number of non-isometric, isospectral surfaces in terms of genus and systole.

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.

1.1.9 Teichmüller space

Classifying the isomorphism classes of Riemann surfaces is known as the moduli problem. By the remark in Section 1.1.6, 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 1.1.9, 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 Teichmüller 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. We will follow [20] closely, but see also [43, 64, 69].

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 1.5).

Figure 1.5: 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 1.6).

Parametrize the sides with constant speed to obtain $t \mapsto s_i(t), t \mapsto s'_i(t), t \in [0, 1]$. 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

Figure 1.6: Construction of pair of pants

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.

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 $v = v_1, v_2, \dots, v_k = w$ such that $(v_i, v_{i+1}) \in E$ for all $i = 1, \dots, k-1$.

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 G contains $3v(G)$ 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 $v_1, \dots, v_{2g-2}$ and its edges by $e_1, \dots, e_{3g-3}$. 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

⁴Technically, we did not define hyperbolic surfaces with boundary, but the definition is similar to manifolds with boundary.

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 1.16. In Figure 1.7 we see an example of a cubic graph. It is completely described by the following list:

$$e_1 = e_{11} \cup e_{21},$$

$$e_2 = e_{12} \cup e_{22},$$

$$e_3 = e_{13} \cup e_{31},$$

$$e_4 = e_{23} \cup e_{32},$$

$$e_5 = e_{33} \cup e_{41},$$

$$e_6 = e_{42} \cup e_{43}.$$

Figure 1.7: Example of a cubic graph

Twist parameters

Let Y and Y' be pairs of pants with boundary geodesics $b_i : \mathbb{S}^1 \rightarrow Y, b'_i : \mathbb{S}^1 \rightarrow Y'$ 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 1.8). For any $a \in \mathbb{R}$, let $X_a = Y \sqcup Y' / \overset{a}{\sim}$ be the disjoint union of Y and Y' modulo the glueing condition $\overset{a}{\sim}$, where $p \overset{a}{\sim} q$ 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 1.8: Gluing of pairs of pants with twist parameter $\frac{1}{4}$

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 1.17. In Figure 1.9 we see an example of the construction of a hyperbolic surface, where the underlying structure is given by the cubic graph of Figure 1.7. 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 1.9: Construction of hyperbolic surfaces

Now we will formally describe this procedure. Let $g \geq 2$ 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 $l_1, \dots, l_{3g-3} \in \mathbb{R}_+$ and $a_1, \dots, a_{3g-3} \in \mathbb{R}$. Associate to each vertex v_i with half-edges $e_{i\alpha}, \alpha = 1, 2, 3$ 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

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

be the disjoint union of the Y_k modulo all the glueing conditions $\sim^{a_k}$, where each $\sim^{a_k}$ is understood to apply to the corresponding Y_i and Y_j from the list above. In this way we obtain a hyperbolic surface $\mathbb{M}$ of genus g . We call the sequence $(l_1, \dots, l_{3g-3}, a_1, \dots, a_{3g-3})$ the Fenchel-Nielsen coordinates of the closed hyperbolic surface $\mathbb{M}$. The following theorem shows the usefulness of the above construction.

Theorem 1.18. *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 1.1.9 we will discuss this further.

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

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

Definition 1.20. Two marked hyperbolic surfaces $(\mathbb{M}, \varphi), (\mathbb{M}', \varphi')$ are called marking equivalent if there exists an isometry $f : \mathbb{M} \rightarrow \mathbb{M}'$ such that φ' and $f \circ \varphi$ are isotopic.

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

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

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

Definition 1.21. 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 $\mathcal{T}_g$ instead of $\mathcal{T}_{g,0}$.

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 $k = 1, \dots, 3g-3$. 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 [20] 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 1.22. *Let G be a fixed connected cubic graph with $v(G) = 2g-2$. Then for every marked hyperbolic surface $(\mathbb{M}, \varphi)$ there exists a unique $(\mathbb{M}', \varphi') \in \mathcal{T}_G$, which is marking equivalent to $(\mathbb{M}, \varphi)$.*

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

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 1.23. 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 $B_{g,n} \rightarrow B_{g,n}$.

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

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

Definition 1.24. The Teichmüller modular group $\mathcal{M}_{g,n}$ is the group of transformations

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

We see that the Teichmüller modular group plays the role that S_3 plays in the case of pairs of pants.

Proposition 1.25. *Marked hyperbolic surfaces $(\mathbb{M}, \varphi), (\mathbb{M}', \varphi') \in \mathcal{T}_{g,n}$ are isometric if and only if there exists $\mu \in \mathcal{M}_{g,n}$ such that $\mu(\mathbb{M}, \varphi) = (\mathbb{M}', \varphi')$.*

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.

1.2 Triangulations

In this section we will discuss triangulations in general, the Delaunay property and how to extend the notion of Delaunay triangulation to hyperbolic surfaces. An algorithm to compute Delaunay triangulations of some specific hyperbolic surfaces will be discussed in Chapter 3 and the minimal number of vertices of a Delaunay triangulation of an arbitrary hyperbolic surface will be treated in Chapter 4.

1.2.1 Simplicial complexes

In short, a triangulation is a subdivision of a 2-dimensional topological manifold into vertices, edges and triangles such that distinct triangles either have no intersection or intersect only in a vertex or an edge. More specifically, in this thesis we assume that every triangulation is a *simplicial complex*. Here, a collection $\mathcal{K}$ of vertices, edges, and triangles (together called *simplices*) is called a simplicial complex if it satisfies the following two conditions (cf [3, Chapter 6] and [60, Chapter 1]):

- each face of a simplex of $\mathcal{K}$ is also an element of $\mathcal{K}$;
- the intersection of two simplices of $\mathcal{K}$ is either empty or is a simplex of $\mathcal{K}$.

Note that, as we will always assume that the set of vertices is finite, all triangulations considered in this thesis are locally finite, so, we can skip the local finiteness in the above conditions (see also the discussion in [24, Section 2.1]).

In our case, an embedding of a graph into a surface is a simplicial complex if and only if it does not contain any 1- or 2-cycles. In particular, a geodesic triangulation of a point set in the Euclidean or hyperbolic planes is always a simplicial complex. This is because there are no geodesic monogons or bigons.

1.2.2 Delaunay triangulations

We recall that given a set of vertices in the Euclidean plane a triangle is called a *Delaunay triangle* if its circumscribed disk does not contain any vertex in its interior. A triangulation of a set of vertices in the Euclidean plane is a *Delaunay triangulation* if all triangles are Delaunay triangles (Figure 1.10).

Figure 1.10: Delaunay triangulation of a point set in the Euclidean plane, together with three circumscribed circles.

We define Delaunay triangulations in the hyperbolic plane in the same way as in the Euclidean plane, where we use the fact that hyperbolic circles are Euclidean circles. Delaunay triangulations on hyperbolic surfaces are defined by lifting vertices on a hyperbolic surface $\mathbb{M}$ to the universal cover $\mathbb{D}$ [15, 30]. More specifically, let $\mathcal{P}$ be a set of vertices on $\mathbb{M}$ and let $\pi : \mathbb{D} \rightarrow \mathbb{D}/\Gamma$ be the projection of the hyperbolic plane $\mathbb{D}$ to the hyperbolic surface $\mathbb{M} = \mathbb{D}/\Gamma$. A triangle (v_1, v_2, v_3) with $v_i \in \mathcal{P}$ is called a *Delaunay triangle* if there exist pre-images $v'_i \in \pi^{-1}(\{v_i\})$ such that the circumscribed disk of the triangle (v'_1, v'_2, v'_3) in the hyperbolic plane does not contain any point of $\pi^{-1}(\mathcal{P})$ in its interior. A triangulation of $\mathcal{P}$ on $\mathbb{M}$ is a *Delaunay triangulation* if all triangles are Delaunay triangles.

A Delaunay triangulation of a point set on a hyperbolic surface $\mathbb{M}$ is related to a Delaunay triangulation in $\mathbb{D}$ as follows [15]. Given a point set $\mathcal{P}$ on $\mathbb{M}$, we consider a Delaunay triangulation T' of the infinite point set $\pi^{-1}(\mathcal{P})$. Then, we let $T = \pi(T')$. By definition, T is a Delaunay triangulation. Moreover, because every triangulation in $\mathbb{D}$ is a simplicial complex, T' is a simplicial complex. However, T is not necessarily a simplicial complex, because projecting T' to $\mathbb{M}$ might introduce 1- or 2-cycles. In Chapter 3 we will discuss the so-called validity condition

that ensures that the result after projection to $\mathbb{M}$ is simplicial. In Chapter 4 we will show explicitly that the projection to $\mathbb{M}$ is simplicial for a specific Delaunay triangulation (that does not necessarily satisfy the mentioned validity condition).

consider specific triangulations and show that will show explicitly that the result after projection of specific Delaunay triangulations is simplicial. We will use the correspondence between Delaunay triangulations in $\mathbb{D}$ and in $\mathbb{M}$ in Definition 4.9 and the proof of Theorem 4.2 and show explicitly that in these cases the result after projecting to $\mathbb{M}$ is simplicial.

To make sure that $T = \pi(T')$ is a well-defined triangulation, we will assume without loss of generality that T' is Γ -invariant, i.e., the image of any Delaunay triangle in T' under an element of Γ is a Delaunay triangle. Otherwise, it is possible that in so-called *degenerate cases* T contains edges that intersect in a point that is not a vertex [16]. Namely, suppose that T' contains a polygon $P = \{p_1, p_2, \dots, p_k\}$ consisting of $k \geq 4$ concircular vertices and let T_P be the Delaunay triangulation of P in T' . Because the Delaunay triangulation of a set of at least four concircular vertices is not uniquely defined, assume that there exists $A \in \Gamma$ such that the Delaunay triangulation $T_{A(P)}$ of $A(P)$ in T' is not equal to $A(T_P)$. Because $\pi(P) = \pi(A(P))$, there exists an edge of $\pi(T_{A(P)})$ and an edge of $\pi(A(T_P))$ that intersect in a point that is not a vertex.

Chapter 2

Systoles of hyperbolic surfaces

2.1 Introduction

In Section 1.1.8, we have introduced the notion of the systole of a hyperbolic surface and stated a well-known upper bound in terms of the genus of the surface. In general, an expression for the systole of a hyperbolic surface in terms of, for example, its Fenchel-Nielsen coordinates is unknown. There exists an algorithm for computing the systole of a given hyperbolic surface [2], but it is not clear how efficient this algorithm is.

Apart from being studied from a purely theoretical viewpoint, the systole also plays a central role in a certain *validity condition* that is used in the algorithm for computing Delaunay triangulations that we present in Chapter 3. This chapter provides three results concerning the systoles of hyperbolic surfaces in general and generalized Bolza surfaces in particular.

First, we give an explicit value for the systole of the generalized Bolza surface $\mathbb{M}_g$ in Section 2.2:

Theorem 2.1. *The systole of the surface $\mathbb{M}_g$ is given by ς_g , where ς_g is defined as*

$$\varsigma_g := 2 \operatorname{arccosh} \left(1 + 2 \cos\left(\frac{\pi}{2g}\right) \right).$$

The same result was also recently obtained by Bai *et al.* [7]. Our proof gives more insight into the representation in D_g of closed geodesics on $\mathbb{M}_g$. Note also that this result generalizes an earlier work of Aurich and Steiner [5], which was restricted to the Bolza surface $\mathbb{M}_2$.

Second, we show in Section 2.3 that the method used to compute the value of the systole of generalized Bolza surfaces can also be applied to hyperbolic surfaces of genus 2 in some neighborhood of the Bolza surface $\mathbb{M}_2$.

From the proofs of the results in Sections 2.2 and 2.3 we will see that systoles of these particular surfaces correspond to products of precisely two generators of the corresponding Fuchsian group. However, this is not the case in general. In

Section 2.4 we will show that for any hyperbolic surface $\mathbb{M} = \mathbb{D}/\Gamma$ there exists a set of generators of Γ such that elements of Γ that correspond to systoles of $\mathbb{M}$ are the product of an arbitrarily large number of generators.

2.2 Systole of generalized Bolza surfaces

This section is devoted to proving Theorem 2.1 stated in the introduction, which gives the value of the systole for the generalized Bolza surfaces $\mathbb{M}_g$. As a preparation for the proof we show in Section 2.2.1 how to represent a simple closed geodesic γ on $\mathbb{M}_g$ by a sequence γ of pairwise disjoint hyperbolic line segments between sides of the fundamental domain D_g . The length of γ is equal to the sum of the lengths of the line segments in γ .

The proof consists of two parts. In Section 2.2.2 we show that $\text{sys}(\mathbb{M}_g) \leq \varsigma_g$ by constructing a simple (non-contractible) closed geodesic of length ς_g . In Section 2.2.3 we show that $\text{length}(\gamma) \geq \varsigma_g$ for all closed geodesics γ by a case analysis based on the line segments contained in the sequence γ representing γ . This shows that $\text{sys}(\mathbb{M}_g) \geq \varsigma_g$.

2.2.1 Representation of a simple closed geodesic by a sequence of segments

Consider a simple closed geodesic γ on the generalized Bolza surface $\mathbb{M}_g$. Because D_g is compact, there is a finite number, say m , of pairwise disjoint hyperbolic lines intersecting D_g in the preimage $\pi_g^{-1}(\gamma)$ of γ . See the leftmost panel in Figure 2.1. These hyperbolic lines are the axes of conjugated elements of Γ_g . Therefore, the intersection of $\pi_g^{-1}(\gamma)$ with D_g consists of m pairwise disjoint hyperbolic line segments between the sides of D_g , the union of which we denote by γ . See the rightmost panel in Figure 2.1. These line segments are oriented and their orientations are compatible with the orientation of γ . In particular, every line segment has a starting point and an endpoint. Since D_g is a fundamental domain for Γ_g , the π_g -images of these line segments form a covering of the closed geodesic γ by m closed subsegments with pairwise disjoint interiors. In other words, these projected segments lie side-by-side on γ , so they form a (cyclically) ordered sequence. This cyclic order lifts to an order $\gamma_1, \dots, \gamma_m$ of the m segments in D_g , which together represent the simple closed geodesic γ . More precisely:

Definition 2.2. An oriented simple closed geodesic γ on $\mathbb{M}_g$ is *represented by a sequence of oriented geodesic segments* $\gamma_1, \dots, \gamma_m$ in D_g if (i) the starting point and endpoint of each segment lie on different sides of ∂D_g , and (ii) the projections $\pi_g(\gamma_1), \dots, \pi_g(\gamma_m)$ are oriented closed subsegments of γ that cover γ , have pairwise disjoint interiors, and lie side-by-side on γ in the indicated order.

We now discuss in more detail how such a sequence is obtained from a hyperbolic isometry the axis of which intersects D_g and projects onto the simple closed

Figure 2.1: Left: Connected components of the preimage of an oriented simple closed geodesic γ on the Bolza surface intersecting the fundamental octagon D_2 . The geodesic is covered by $f = f_0 f_5 f_0$ (and all its conjugates in the Bolza group), which has axis $\mathbf{l}_1$. Right: The cyclic sequence of geodesic segments $\gamma_1, \gamma_2, \gamma_3$ in D_g represents γ . The segments γ_1^*, γ_2^* and γ_3^* are the successive intersections of the axis of f with $F_0(D_g)$, $F_1(D_g)$ and $F_2(D_g)$, where $F_0 = \mathbb{1}$, $F_1 = f_0$ and $F_2 = f_0 f_5$. The endpoint of γ_3^* is $F_3(\mathbf{p}_0^*)$, where $F_3 = f$ and $\mathbf{p}_0^*$ is the starting point of γ_1^* . The endpoint of γ_k is paired with the starting point of γ_{k+1} by the side-pairing $F_k^{-1} F_{k-1}$. In this example $F_1^{-1} F_0 = f_0^{-1} = f_4$, $F_2^{-1} F_1 = f_5^{-1} = f_1$, and $F_3^{-1} F_2 = f_0^{-1} = f_4$.

geodesic. Let $\mathbf{l}_1$ be an arbitrary oriented geodesic in the set of m connected components of $\pi_g^{-1}(\gamma)$ that intersect D_g . The oriented segment γ_1 is the intersection $\mathbf{l}_1 \cap D_g$. Let $f \in \Gamma_g$ be the hyperbolic isometry that covers γ and has axis $\mathbf{l}_1$. More precisely, if $\mathbf{p}_0^*$ is the starting point of γ_1 , then the segment $[\mathbf{p}_0^*, f(\mathbf{p}_0^*)]$ projects onto γ , and π_g is injective on the interior of this segment.

Let $F_0(D_g), F_1(D_g), \dots, F_{m-1}(D_g)$, be the *sequence* of successive Dirichlet domains intersected by the segment $[\mathbf{p}_0^*, f(\mathbf{p}_0^*)]$. Here $F_0 = \mathbb{1}$ and $F_0, F_1, \dots, F_{m-1}$ are distinct elements of Γ_g . This sequence consists of m regions, since the segments $F_0^{-1}(\mathbf{l}_1), F_1^{-1}(\mathbf{l}_1), \dots, F_{m-1}^{-1}(\mathbf{l}_1)$ are the (pairwise disjoint) geodesics in $\pi_g^{-1}(\gamma)$ that intersect D_g . This implies that the segment $[\mathbf{p}_0^*, f(\mathbf{p}_0^*)]$ is covered by the *sequence* of closed segments $\gamma_1^*, \dots, \gamma_m^*$ in which $\mathbf{l}_1$ intersects these m regions, i.e., $\gamma_k^* = F_{k-1}(D_g) \cap \mathbf{l}_1$, for $k = 1, \dots, m$. (Note that $\gamma_1^* = \gamma_1$.) The segments $\gamma_k = F_{k-1}^{-1}(\gamma_k^*)$, $k = 1, \dots, m$, lie in D_g and project onto the same subsegment of γ as γ_k^* . In other words, $\pi_g(\gamma_1), \dots, \pi_g(\gamma_m)$ lie side-by-side on the closed geodesic γ and cover γ . Therefore, the simple closed geodesic γ is represented by the sequence $\gamma_1, \dots, \gamma_m$.

It is convenient to consider $f(\mathbf{p}_0^*)$ as the starting point of the segment $\mathbf{l}_1 \cap f(D_g)$, which we denote by γ_{m+1}^* . Taking $F_m = f$, we see that $\gamma_{m+1}^* = \mathbf{l}_1 \cap F_m(D_g)$. Extending our earlier definition $\gamma_k = F_{k-1}^{-1}(\gamma_k^*)$ to $k = m+1$, we see that γ_{m+1} is the subsegment of $\pi_g^{-1}(\gamma) \cap D_g$ with starting point $F_m^{-1}(f(\mathbf{p}_0^*)) = \mathbf{p}_0^*$, so $\gamma_{m+1} =$

γ_1 .

Finally, we show that the endpoint of γ_k is mapped to the starting point of γ_{k+1} by a side-pairing transformation of D_g , for $k = 1, \dots, m$. Since $F_{k-1}(D_g) \cap F_k(D_g)$ is a side of $F_k(D_g)$, for $k = 1, \dots, m$, the intersection $F_k^{-1}F_{k-1}(D_g) \cap D_g$ is a side of D_g , say s_{j_k} . Then $F_k^{-1}F_{k-1} = f_{j_k}$, since $s_{j_k} = f_{j_k}(D_g) \cap D_g$. Let $\gamma_k^* = [p_{k-1}^*, p_k^*]$, then $\gamma_k = [F_{k-1}^{-1}(p_{k-1}^*), F_{k-1}^{-1}(p_k^*)]$. Therefore, f_{j_k} maps the endpoint $F_{k-1}^{-1}(p_k^*)$ of γ_k to the starting point $F_k^{-1}(p_k^*)$ of γ_{k+1} , since $f_{j_k}F_{k-1}^{-1} = F_k^{-1}$. See the rightmost panel in Figure 2.1.

2.2.2 Upper bound for the systole

To show that $\text{sys}(\mathbb{M}_g) \leq \varsigma_g$ it is sufficient to prove the following lemma.

Lemma 2.3. *There is a simple closed geodesic on $\mathbb{M}_g$ of length ς_g .*

Proof. The axis of the hyperbolic translation $f_{2g+1}f_0$ projects onto a simple closed geodesic γ on $\mathbb{M}_g$ with length equal to the translation length $\ell(f_{2g+1}f_0)$. See Section 1.1.8. Since $f_{2g+1}f_0$ is represented by the matrix $A_{2g+1}A_0$, with A_j given by (1.2), we see that

$$\cosh \frac{1}{2}\ell(f_{2g+1}f_0) = \frac{1}{2} |\text{Tr}(A_{2g+1}A_0)| = 1 + \cos\left(\frac{\pi}{2g}\right).$$

Since $1 + \cos\left(\frac{\pi}{2g}\right) = \cosh \frac{1}{2}\varsigma_g$ we conclude that γ has length ς_g . $\square$

Remark 2.4. Two connected components of the pre-image $\pi_g^{-1}(\gamma)$ of the simple closed geodesic γ , appearing in the proof of Lemma 2.3, intersect the fundamental polygon D_g : the axis l_1 of $f = f_{2g+1}f_0$, and the geodesic $l_2 = f_{2g+1}^{-1}(l_1)$, which is the axis of f_0f_{2g+1} . The geodesic γ is represented by the segments $\gamma_1 = l_1 \cap D_g$ and $\gamma_2 = l_2 \cap D_g$. The first segment connects the midpoint m_{2g} of s_{2g} and the midpoint m_{2g+1} of s_{2g+1} , whereas the second segment connects the midpoints of s_0 and s_1 . See Figure 2.2.

This can be seen as follows. Since $f = f_{2g+1}f_{2g}^{-1}$ and the axes of f_{2g} and f_{2g+1} intersect at the origin O , (the proof of) Theorem 7.38.6 of [10] implies that the axis of f passes through the midpoint of the segment $[O, f_{2g}(O)]$ and the midpoint of $[O, f_{2g+1}(O)]$. But these midpoints coincide with m_{2g} and m_{2g+1} , respectively, so $[m_{2g}, m_{2g+1}] = \gamma_1$. This theorem also implies that the length of the latter segment is half the translation length of f , i.e., $\frac{1}{2}\varsigma_g$. A similar argument shows that the length of γ_2 is $\frac{1}{2}\varsigma_g$.

2.2.3 Lower bound for the systole

We now prove that the length of every simple closed geodesic of $\mathbb{M}_g$ is at least ς_g , or, equivalently, that the total length of the segments representing such a geodesic is at least ς_g . To this end we consider different types of closed geodesics based on

Figure 2.2: Two geodesics in the pre-image of γ intersect the fundamental polygon D_g (left) The intersections are the segments γ_1 and γ_2 , which represent γ (right). The figure illustrates the situation for the Bolza surface ($g = 2$).

which “kind” of segments are contained in the sequence. We say that an oriented hyperbolic line segment between two sides of D_g is a k -segment, $1 \leq k \leq 4g - 1$, if its starting point and endpoint are contained in s_j and s_{j+k} , respectively, for some j with $0 \leq j \leq 4g - 1$, where indices are counted modulo $4g$. Furthermore, we say that the segment is k -separated or has separation k , $1 \leq k \leq 2g$, if either the segment itself or the segment with the opposite orientation is a k -segment. Equivalently, a k -separated segment is either a k -segment or a $(4g - k)$ -segment. For example, both segments in Figure 2.2 - Right are 1-separated, but γ_1 is a 1-segment while γ_2 is a 7-segment.

In the derivation of the lower bound for the systole we will use the following lemma. This lemma will be used in the proof of Proposition 3.18 in Section 3.6 as well.

Lemma 2.5. *For geodesic segments between the sides of D_g the following properties hold:*

1. *The length of a segment that has separation at least 4 is at least ς_g .*
2. *The length of a segment that has separation at least 2 is at least $\frac{1}{2}\varsigma_g$.*
3. *Every pair of consecutive 1-separated segments consists of exactly one 1-segment and one $(4g - 1)$ -segment.*
4. *The length of two consecutive 1-separated segments is at least $\frac{1}{2}\varsigma_g$.*
5. *A sequence of segments consisting of precisely two 1-separated segments has length ς_g .*

Proof. We prove the different properties in the same order as the statement of the lemma.

1. Consider a segment γ_j of separation $k \geq 4$. By symmetry of D_g , we can assume that γ_j is a segment between s_0 and s_k . The length of γ_j is greater than or equal to the distance between s_0 and s_k , which is given as the length of the common orthogonal line segment $\gamma_j^\perp$ between s_0 and s_k (see Figure 2.3).

Figure 2.3: Computing the length of $\gamma_j^\perp$.

To find an expression for $\text{length}(\gamma_j^\perp)$, we draw line segments between the origin O and s_0 and between O and s_k . In this way, we obtain a hyperbolic pentagon with four right-angles and remaining angle $\frac{k\pi}{2g}$. The line segment from O to s_0 is a non-hypotenuse side of an isosceles triangle with angles $\frac{\pi}{4g}, \frac{\pi}{4g}, \frac{\pi}{2}$, as shown in Figure 2.3. Therefore, [10, Theorem 7.11.3(i)]

$$\cosh(d(O, s_0)) = \frac{\cos(\frac{\pi}{4g})}{\sin(\frac{\pi}{4g})} = \cot(\frac{\pi}{4g}).$$

Drawing a line segment from O orthogonal to $\gamma_j^\perp$, we obtain two quadrilaterals, each of which has three right angles and remaining angle $\frac{k\pi}{4g}$. It follows that [10, Theorem 7.17.1(ii)]

$$\cosh(\frac{1}{2} \text{length}(\gamma_j^\perp)) = \cosh(d(O, s_0)) \sin(\frac{k\pi}{4g}) = \cot(\frac{\pi}{4g}) \sin(\frac{k\pi}{4g}) \quad (2.1)$$

The lower bound for the length of a segment of separation at least 4 follows from $\sin(\frac{k\pi}{4g}) \geq \sin(\frac{\pi}{g})$ and a direct computation using properties of trigonometric functions.

2. Now, consider a segment γ_j of separation $k \geq 2$. Using the same argument as in Part 1, we see that formula (2.1) still holds. The lower bound for the length of γ_j follows from $\sin(\frac{k\pi}{4g}) \geq \sin(\frac{\pi}{2g})$ and a direct computation using properties of trigonometric functions.
3. Consider a pair γ_1, γ_2 of consecutive 1-separated segments. By symmetry, we can assume without loss of generality that γ_1 is a 1-segment between s_0 and s_1 (Figure 2.4).

Figure 2.4: Construction in the proof of Part 3 of Lemma 2.5.

The side-pairing transformation f_{2g+1} maps the endpoint p of γ_1 to the starting point p' of γ_2 . Both p and p' lie on the same equidistant curve L of the axis of f_{2g+1} . The curve L intersects all geodesics that are perpendicular to $\text{axis}(f_{2g+1})$ orthogonally. It can be seen that the angle α between γ_1 and s_1 is acute and that γ_1 and $f_{2g+1}^{-1}(\gamma_2)$ lie on opposite sides of L . Moreover, the parts of $\mathbb{D}$ separated by L are f_{2g+1} -invariant (see also Figure 1.1 - Right). Hence, γ_1 and γ_2 lie on opposite sides of L as well. We conclude that the endpoint of γ_2 lies on s_{2g} so γ_2 is a $(4g-1)$ -segment.

4. As in Part 3, denote the two segments by γ_1 and γ_2 . By Part 3, we know that one of γ_1 and γ_2 is a 1-segment and the other is a $(4g-1)$ -segment. Hence, we can assume without loss of generality that γ_1 is a 1-segment between s_0 and s_1 and γ_2 is a $(4g-1)$ -segment between s_{2g+1} and s_{2g} (see Figure 2.5).

Let x be the distance between $\mathbf{p}_1$ and $\mathbf{v}_1$ and let α_1 be the angle between γ_1 and $\mathbf{s}_0$. The distance between $\mathbf{p}'_1$ and $\mathbf{v}_{2g+1}$ is $\ell - x$, where the length ℓ of the sides satisfies $\cosh(\frac{1}{2}\ell) = \cot(\frac{\pi}{4g})$. Let α_2 be the angle between γ_2 and $\mathbf{s}_{2g}$.

Figure 2.5: Construction in the proof of Part 4 of Lemma 2.5.

By the hyperbolic sine rule,

$$\frac{\sinh(\text{length}(\gamma_1))}{\sin(\frac{\pi}{2g})} = \frac{\sinh(x)}{\sin(\alpha_1)},$$

so

$$\sinh(\text{length}(\gamma_1)) \geq \sinh(x) \sin(\frac{\pi}{2g}).$$

In a similar way, we obtain

$$\sinh(\text{length}(\gamma_2)) \geq \sinh(\ell - x) \sin(\frac{\pi}{2g}).$$

We minimize

$$f(x) := \text{arcsinh}(\sinh(x) \sin(\frac{\pi}{2g})) + \text{arcsinh}(\sinh(\ell - x) \sin(\frac{\pi}{2g}))$$

subject to $0 < x < \ell$, as this provides a lower bound for $\text{length}(\gamma_1 \cup \gamma_2)$. Because

$$\frac{d^2}{dx^2} \text{arcsinh}(\sinh(x) \sin(\frac{\pi}{2g})) = \frac{\sin(\frac{\pi}{2g}) \cos^2(\frac{\pi}{2g}) \sinh(x)}{(\sin^2(\frac{\pi}{2g}) \sinh^2(x) + 1)^{3/2}} > 0$$

for all $x > 0$, the function $x \mapsto \operatorname{arcsinh}(\sinh(x) \sin(\frac{\pi}{2g}))$ is strictly convex. It follows that f is also strictly convex, so it has a unique global minimum. The derivative of f is given by

$$f'(x) = \frac{\sin(\frac{\pi}{2g}) \cosh(x)}{(\sin^2(\frac{\pi}{2g}) \sinh^2(x) + 1)^{1/2}} - \frac{\sin(\frac{\pi}{2g}) \cosh(\ell - x)}{(\sin^2(\frac{\pi}{2g}) \sinh^2(\ell - x) + 1)^{1/2}}.$$

It is clear that $f'(\frac{1}{2}\ell) = 0$, so $x = \frac{1}{2}\ell$ is the unique minimizer with minimum value $f(\frac{1}{2}\ell) = 2 \operatorname{arcsinh}(\sinh(\frac{1}{2}\ell) \sin(\frac{\pi}{2g}))$. By the discussion above, this implies that

$$\sinh(\frac{1}{2} \operatorname{length}(\gamma_1 \cup \gamma_2)) \geq \sinh(\frac{1}{2}\ell) \sin(\frac{\pi}{2g}).$$

Then

$$\begin{aligned} \cosh(\operatorname{length}(\gamma_1 \cup \gamma_2)) &= 2 \sinh^2(\frac{1}{2} \operatorname{length}(\gamma_1 \cup \gamma_2)) + 1, \\ &\geq 2 \sinh^2(\frac{1}{2}\ell) \sin^2(\frac{\pi}{2g}) + 1, \\ &= 2(\cot^2(\frac{\pi}{4g}) - 1) \sin^2(\frac{\pi}{2g}) + 1, \\ &= (1 + 2 \cos(\frac{\pi}{2g}))^2, \end{aligned}$$

from which we conclude that $\operatorname{length}(\gamma_1 \cup \gamma_2) > \frac{1}{2}\varsigma_g$.

5. Now, using the notation from Part 4, assume that $\gamma = \gamma_1 \cup \gamma_2$. By the argument in Part 4, the length of γ is minimal when $\mathbf{p}_1$ is the midpoint of $\mathbf{s}_1$ and $\mathbf{p}'_1$ is the midpoint of $\mathbf{s}_{2g+1}$, given any location of the starting point of γ_1 and the endpoint of γ_2 . By symmetry of the argument, it follows that $\operatorname{length}(\gamma)$ is minimal when the starting point of γ_1 is the midpoint of $\mathbf{s}_0$ and the endpoint of γ_2 is the midpoint of $\mathbf{s}_{2g}$. It can be seen that the resulting curve is the curve constructed in the proof of Lemma 2.3, the length of which is ς_g . This finishes the proof. □

The lower bound for the systole follows from the following result.

Lemma 2.6. *Every closed geodesic on $\mathbb{M}_g$ has length at least ς_g .*

Proof. It is sufficient to show that every sequence of segments representing a closed geodesic on $\mathbb{M}_g$ has length at least ς_g . Let γ be a sequence of segments. We distinguish between the following four types:

1. γ contains at least one segment that has separation at least 4,
2. γ contains at least two segments that have separation 2 or 3 and all other segments are 1-separated,
3. γ contains exactly one segment that has separation 2 or 3 and all other segments are 1-separated,

4. all segments of γ are 1-separated.

It is straightforward to check that every sequence of segments is of precisely one type.

First, suppose that γ is of Type 1 or 2. Then, it follows directly from Part 1 and 2 of Lemma 2.5 that $\text{length}(\gamma) \geq \varsigma_g$.

Second, suppose that γ is of Type 3. It is not possible to form a closed geodesic with a segment of separation 2 or 3 and just one segment of separation 1, so we can assume that there are at least two 1-separated segments. In the cyclic ordering of the segments, these 1-separated segments are consecutive, so it follows from Part 4 of Lemma 2.5 that their combined length is at least $\frac{1}{2}\varsigma_g$. By Part 2 of Lemma 2.5 the length of the segment of separation 2 or 3 is at least $\frac{1}{2}\varsigma_g$ as well, so we conclude that $\text{length}(\gamma) \geq \varsigma_g$.

Finally, suppose that γ is of Type 4. By Part 3 of Lemma 2.5 every 1-segment is followed by a $(4g-1)$ -segment and reversely, so in particular the number of 1-segments and $(4g-1)$ -segments is identical. Therefore, the number of 1-separated segments in γ is even (and at least two). If the number of 1-separated segments is exactly two, then $\text{length}(\gamma) = \varsigma_g$ by Part 5 of Lemma 2.5. If the number of 1-separated segments is at least four, then $\text{length}(\gamma) \geq \varsigma_g$, since every pair of consecutive 1-separated segments has combined length at least $\frac{1}{2}\varsigma_g$ by Part 4 of Lemma 2.5.

This finishes the proof. $\square$

2.3 Systole of hyperelliptic surfaces

In the previous section, we computed the value of the systole of generalized Bolza surfaces. These surfaces are examples of *hyperelliptic surfaces*, i.e., hyperbolic surfaces of genus g that have an isometry ψ with $\psi^2 = \text{id}$ and with $2g+2$ fixed points. In particular, all hyperbolic surfaces of genus 2 are hyperelliptic [67]. In this section, we will show that we can use the same reasoning as in Section 2.2 to compute the systole of hyperelliptic surfaces in some neighborhood of the generalized Bolza surfaces (Theorem 2.8).

2.3.1 Neighborhood of hyperelliptic surfaces

Before we state the main result of this section in Section 2.3.2, we first describe more precisely what we mean by a neighborhood of hyperelliptic surfaces.

Every hyperelliptic surface of genus g can be represented by a $4g$ -gon in the Poincaré disk in the following way, based on a similar representation of hyperbolic surfaces of genus 2 [1]. Given a set of $2g$ points $z_0, z_1, \dots, z_{2g-1}$ in $\mathbb{D}$ with arguments

$$0 = \arg(z_0) < \arg(z_1) < \dots < \arg(z_{2g-1}) < \pi,$$

set $z_j = -z_{j-2g}$ for $j = 2g, \dots, 4g-1$ and let $P = [z_0, z_1, \dots, z_{4g-1}]$ be the resulting geodesic $4g$ -gon. We call P admissible if the interior angle sum is 2π . By

pairing opposite sides of an admissible polygon P we obtain a smooth hyperelliptic surface $\mathbb{M}_P$. Moreover, it is known that every hyperelliptic surface $\mathbb{M}$ of genus g corresponds to an admissible $4g$ -gon P , i.e., $\mathbb{M} = \mathbb{M}_P$ for some admissible $4g$ -gon P [67, Theorem 14].

More precisely, there is a one-to-one correspondence between the set of all admissible $4g$ -gons and the subset of the Teichmüller space $\mathcal{T}_g$ consisting of marked hyperelliptic surfaces. In this correspondence, the marking of a hyperelliptic surface is given by the closed geodesics represented by the sides of the admissible polygon. Henceforth, we will denote the marked hyperelliptic surface corresponding to the admissible polygon P by $\mathbb{M}_P$. Note that different admissible polygons may correspond to the same isometry class of hyperbolic surfaces, just like multiple marked hyperbolic surfaces may correspond to the same isometry class. For example, rotations around the origin of an admissible polygon that map z_j for $j \in \{1, \dots, 2g-1\}$ to the real axis do not change the isometry class of the corresponding hyperelliptic surface.

We define an ε -neighborhood of the generalized Bolza surfaces in the set of (marked) hyperelliptic surfaces in the following way. In the definition, $\text{dist}_{\mathbb{M}_g} = \text{arccosh}(\cot^2(\frac{\pi}{4g}))$ denotes the distance between O and a vertex of D_g .

Definition 2.7. Let $0 < \varepsilon < 1$. The ε -neighborhood of $\mathbb{M}_g$ in the set of hyperelliptic surfaces is the set of (marked) hyperelliptic surfaces $\mathbb{M}_P$ represented by an admissible $4g$ -gon P satisfying

$$\begin{aligned} |z_j| &\in ((1 - \varepsilon) \text{dist}_{\mathbb{M}_g}, (1 + \varepsilon) \text{dist}_{\mathbb{M}_g}), \\ \arg(z_{j+1}) - \arg(z_j) &\in ((1 - \varepsilon) \frac{\pi}{2g}, (1 + \varepsilon) \frac{\pi}{2g}). \end{aligned}$$

We denote the ε -neighborhood of $\mathbb{M}_g$ by $\mathbb{M}_g(\varepsilon)$.

An illustration of a hyperelliptic surface in a neighborhood of $\mathbb{M}_g$ for $g = 2$ is given in Figure 2.6.

2.3.2 Systole of hyperelliptic surfaces in a neighborhood of bounded size

The next theorem is the main result of this section.

Theorem 2.8. *There exists a positive constant $c \in \mathbb{R}$ such that for any $g \geq 2$ and $\mathbb{M}_P \in \mathbb{M}_g(c(\log(g))^{-1})$ a systole of $\mathbb{M}_P$ is the projection onto $\mathbb{M}_P$ of a single segment or a sequence of two 1-separated segments between the sides of P .*

We checked numerically that $c = 0.04$ is sufficient for $2 \leq g \leq 10^4$. For more details on obtaining an estimate for c , we refer to the discussion at the end of Section 2.3.4. Moreover, we do not claim that the $O((\log(g))^{-1})$ order of growth is necessary; it could be the case that the statement is also true for $\mathbb{M}_g(\varepsilon)$ where ε has a larger order of growth.

Figure 2.6: The Bolza surface $\mathbb{M}_2$ (black) with a hyperelliptic surface (red) in a neighborhood of $\mathbb{M}_2$.

To illustrate the idea of the proof, let us recall the proof of Theorem 2.1 stating that $\text{sys}(\mathbb{M}_g) = \varsigma_g$. We first found a closed geodesic with length ς_g (Lemma 2.3) and then showed that $\text{sys}(\mathbb{M}_g) \geq \varsigma_g$ (Lemma 2.6). The main ingredients of the proof of Lemma 2.6 were lower bounds for the length of a k -segment for $k \geq 2$ and the length of a pair of consecutive 1-segments presented in Lemma 2.5. The trigonometric computations from which these lower bounds followed, were relatively simple due to the rotational symmetry of D_g .

However, an admissible polygon in some neighborhood of $\mathbb{M}_g$ does not necessarily have the same rotational symmetry. Therefore, we will derive lower bounds on several trigonometric quantities of an arbitrary admissible polygon in a neighborhood $\mathbb{M}_g(\varepsilon)$ of $\mathbb{M}_g$ (as function of ε) in Section 2.3.3. Using these lower bounds, we will present an extension of Lemma 2.5 to hyperelliptic surfaces in a neighborhood of $\mathbb{M}_g$ in Section 2.3.4. Finally, the proof of Theorem 2.8 is given in Section 2.3.5.

Remark 2.9. In Section 3.6.2 we will introduce the *inclusion property*, an important ingredient for defining the data structure of the algorithm to compute Delaunay triangulations of the generalized Bolza surfaces. The proof of the inclusion property follows more or less directly from the statements in Lemma 2.5 that provide lower bounds for the lengths of k -segments for $k \geq 2$ and pairs of 1-separated segments. We will see in the proof of Theorem 2.8 that these lower bounds are also satisfied for all admissible polygons corresponding to hyperelliptic surfaces in $\mathbb{M}_g(c(\log(g))^{-1})$. Hence, the inclusion property is satisfied for these

hyperelliptic surfaces as well.

2.3.3 Trigonometry of admissible polygons

As mentioned, the main ingredients of the proof of $\text{sys}(\mathbb{M}_g) \geq \varsigma_g$ in Lemma 2.6 were lower bounds for the length of a k -segment for $k \geq 2$ and the length of a pair of consecutive 1-segments in Lemma 2.5. The lower bound for the length of a k -segment for $g \geq 2$ was proved using the angle between the line segments connecting O with consecutive midpoints and the distance between O and the sides of D_g . By the rotational symmetry of D_g , these lengths and angles were all identical, but for an admissible polygon in $\mathbb{M}_g(\varepsilon)$ this is not necessarily the case. Therefore, for a given admissible polygon P let $\mathbf{m}_j$ be the orthogonal projection of O onto side $[z_j, z_{j+1}]$ and let $m_j = d(O, \mathbf{m}_j)$ for $j = 0, \dots, 4g - 1$. Note that if $\mathbb{M}_P = \mathbb{M}_g$ then the $\mathbf{m}_j$ are the midpoints of the sides of $P = D_g$. Furthermore, we define ξ_j to be the angle between $[O, \mathbf{m}_j]$ and $[O, \mathbf{m}_{j+1}]$ (see Figure 2.7 (left)).

Figure 2.7: Definitions of $\mathbf{m}_j$ and ξ_{jk} (left) and d_j and α_j (right).

Similarly, the lower bound of the length of a pair of consecutive 1-segments was proved using the length of the sides and the angle between adjacent sides of D_g . Hence, let $d_j = d(z_j, z_{j+1})$ and let α_j be the angle between sides $[z_{j-1}, z_j]$ and $[z_j, z_{j+1}]$ for $j = 0, \dots, 4g - 1$ (see Figure 2.7 (right)).

The following lemma states lower bounds $m_{\min}, \xi_{\min}, d_{\min}$ and $\alpha_{\min}$ for m_j, ξ_j, d_j and α_j , respectively, as function of ε .

Lemma 2.10. *Let $0 < \varepsilon < 1$ and let P be an arbitrary admissible polygon in $\mathbb{M}_g(\varepsilon)$. Then the following four properties hold:*

1. The distances from O to the sides satisfy $m_j > m_{\min}$ for all $j = 0, \dots, 4g-1$, with

$$m_{\min} = \operatorname{arctanh}(\tanh((1-\varepsilon) \operatorname{dist}_{\mathbb{M}_g}) \cos((1+\varepsilon) \frac{\pi}{4g})). \quad (2.2)$$

2. The angles between consecutive midpoints satisfy $\xi_j > \xi_{\min}$ for all $j = 0, \dots, 4g-1$, with

$$\xi_{\min} = 2\eta, \quad (2.3)$$

where η is the solution of

$$\cos(\eta) = \frac{\tanh((1+\varepsilon) \operatorname{dist}_{\mathbb{M}_g})}{\tanh((1-\varepsilon) \operatorname{dist}_{\mathbb{M}_g})} \cdot \cos\left((1-\varepsilon) \frac{\pi}{2g} - \eta\right).$$

3. The lengths of the sides satisfy $d_j > d_{\min}$ for all $j = 0, \dots, 4g-1$, with

$$d_{\min} = \operatorname{arccosh}(\cosh^2((1-\varepsilon) \operatorname{dist}_{\mathbb{M}_g}) - \sinh^2((1-\varepsilon) \operatorname{dist}_{\mathbb{M}_g}) \cos((1-\varepsilon) \frac{\pi}{2g})). \quad (2.4)$$

4. The angles between the sides satisfy $\alpha_j > \alpha_{\min}$ for all $j = 0, \dots, 4g-1$, with

$$\alpha_{\min} = 2 \arcsin \left(\frac{\sinh(\operatorname{arctanh}(\tanh((1+\varepsilon) \operatorname{dist}_{\mathbb{M}_g}) \cos(\eta'))))}{\sinh((1+\varepsilon) \operatorname{dist}_{\mathbb{M}_g})} \right), \quad (2.5)$$

where η' is the solution of

$$\cos(\eta') = \frac{\tanh((1-\varepsilon) \operatorname{dist}_{\mathbb{M}_g})}{\tanh((1+\varepsilon) \operatorname{dist}_{\mathbb{M}_g})} \cdot \cos\left((1+\varepsilon) \frac{\pi}{2g} - \eta'\right).$$

Proof.

1. Let $j = 0, \dots, 4g-1$ and consider the triangle $[O, \mathbf{z}_j, \mathbf{z}_{j+1}]$ (see Figure 2.8). By symmetry of the argument, we may assume without loss of generality that $|\mathbf{z}_j| \geq |\mathbf{z}_{j+1}|$. For fixed $|\mathbf{z}_{j+1}|$ and any angle $\arg(\mathbf{z}_{j+1}) - \arg(\mathbf{z}_j)$ the distance m_j to side $[\mathbf{z}_j, \mathbf{z}_{j+1}]$ is monotonically increasing as function of $|\mathbf{z}_{j+1}|$. Therefore, m_j is minimal when $|\mathbf{z}_{j+1}|$ is minimal, i.e., when $|\mathbf{z}_j| = |\mathbf{z}_{j+1}|$. Then, we know that in the right-angled triangle $[O, \mathbf{m}_j, \mathbf{z}_j]$

$$\tanh(m_j) = \tanh(|\mathbf{z}_j|) \cos(\tfrac{1}{2}(\arg(\mathbf{z}_{j+1}) - \arg(\mathbf{z}_j))),$$

so m_j is monotonically increasing as function of $|\mathbf{z}_j|$ and monotonically decreasing as function of $\arg(\mathbf{z}_{j+1}) - \arg(\mathbf{z}_j)$. Therefore, a lower bound $m_{\min}$ is obtained by substituting $|\mathbf{z}_j| = (1-\varepsilon) \operatorname{dist}_{\mathbb{M}_g}$ and $\arg(\mathbf{z}_{j+1}) - \arg(\mathbf{z}_j) = (1+\varepsilon) \frac{\pi}{2g}$. The result follows.

Figure 2.8: Trigonometry in a triangle contained in P .

2. Again, consider the triangle $[O, z_j, z_{j+1}]$ for $j = 0, \dots, 4g - 1$. For $k = 0, 1$, let η_{jk} be the angle between $[O, m_j]$ and $[O, z_{j+k}]$. To show that $\xi_j > \xi_{\min} = 2\eta$, it is sufficient to show that $\eta_{j1} > \eta$ by symmetry of the argument. In the right-angled triangles $[O, m_j, z_j]$ and $[O, m_j, z_{j+1}]$ we know that

$$\begin{aligned} \cos(\eta_{j0}) &= \frac{\tanh(d(O, m_j))}{\tanh(|z_j|)}, \\ \cos(\eta_{j1}) &= \frac{\tanh(d(O, m_j))}{\tanh(|z_{j+1}|)}, \end{aligned}$$

so that

$$\cos(\eta_{j1}) = \frac{\tanh(|z_j|)}{\tanh(|z_{j+1}|)} \cos(\eta_{j0}).$$

Because $\arg(z_{j+1}) - \arg(z_j) = \eta_{j0} + \eta_{j1}$, we see that in this case η_{j1} is the solution η of

$$\cos(\eta) = \frac{\tanh(|z_j|)}{\tanh(|z_{j+1}|)} \cos(\arg(z_{j+1}) - \arg(z_j) - \eta).$$

Therefore, η_{j1} is minimal when $|z_j|$ is maximal, $|z_{j+1}|$ is minimal and $\arg(z_{j+1}) - \arg(z_j)$ is minimal, i.e., when $|z_j| = (1 + \varepsilon) \text{dist}_{\mathbb{M}_g}$, $|z_{j+1}| = (1 - \varepsilon) \text{dist}_{\mathbb{M}_g}$ and $\arg(z_{j+1}) - \arg(z_j) = (1 - \varepsilon) \frac{\pi}{2g}$. This proves the desired formula.

3. Let $j = 0, \dots, 4g - 1$ and consider the triangle $[O, z_j, z_{j+1}]$. As in the proof of the formula for $m_{\min}$, we may assume that $|z_j| \geq |z_{j+1}|$ by symmetry of

the argument. For fixed $|\mathbf{z}_{j+1}|$ and any angle $\arg(\mathbf{z}_{j+1}) - \arg(\mathbf{z}_j)$ the side $[\mathbf{z}_j, \mathbf{z}_{j+1}]$ has minimal length when $|\mathbf{z}_j|$ is minimal, i.e., when $|\mathbf{z}_j| = |\mathbf{z}_{j+1}|$. Applying the hyperbolic cosine rule yields

$$\cosh(d_j) = \cosh^2(|\mathbf{z}_j|) - \sinh^2(|\mathbf{z}_j|) \cos(\arg(\mathbf{z}_{j+1}) - \arg(\mathbf{z}_j)).$$

It can be seen that d_j is a monotonically increasing function of $|\mathbf{z}_j|$ and $\arg(\mathbf{z}_{j+1}) - \arg(\mathbf{z}_j)$. Therefore, a lower bound $d_{\min}$ is obtained by substituting $|\mathbf{z}_j| = (1 - \varepsilon) \text{dist}_{\mathbb{M}_g}$ and $\arg(\mathbf{z}_{j+1}) - \arg(\mathbf{z}_j) = (1 - \varepsilon) \frac{\pi}{2g}$. The result follows.

4. Again, consider the triangle $[O, \mathbf{z}_j, \mathbf{z}_{j+1}]$ for $j = 0, \dots, 4g - 1$. For $k = 0, 1$, let α_{jk} be the angle between $[O, \mathbf{z}_{j+k}]$. To prove that $\alpha_j > \alpha_{\min}$ for $\alpha_{\min}$ as in Equation (2.5) it is sufficient to show that

$$\alpha_{j0} > \arcsin \left(\frac{\sinh(\arctanh(\tanh((1 + \varepsilon) \text{dist}_{\mathbb{M}_g}) \cos(\eta'))))}{\sinh((1 + \varepsilon) \text{dist}_{\mathbb{M}_g})} \right),$$

where η' is as stated in the result. By the hyperbolic sine rule we know that

$$\frac{\sin(\alpha_{j0})}{\sinh(|\mathbf{z}_{j+1}|)} = \frac{\sin(\arg(\mathbf{z}_{j+1}) - \arg(\mathbf{z}_j))}{\sinh(d(\mathbf{z}_j, \mathbf{z}_{j+1}))}.$$

For any fixed $|\mathbf{z}_{j+1}|$ and $\arg(\mathbf{z}_{j+1}) - \arg(\mathbf{z}_j)$, $d(\mathbf{z}_j, \mathbf{z}_{j+1})$ is an increasing function of $|\mathbf{z}_j|$, so α_{j0} is a decreasing function of $|\mathbf{z}_j|$. Therefore, α_{j0} is minimal for $|\mathbf{z}_j| = (1 + \varepsilon) \text{dist}_{\mathbb{M}_g}$. By a similar argument, it can be seen that α_{j0} is an increasing function of $|\mathbf{z}_{j+1}|$ for any fixed $|\mathbf{z}_j|$ and $\arg(\mathbf{z}_{j+1}) - \arg(\mathbf{z}_j)$, so α_{j0} is minimal for $|\mathbf{z}_{j+1}| = (1 - \varepsilon) \text{dist}_{\mathbb{M}_g}$. Hence, to find a lower bound for α_{j0} we can assume without loss of generality that $|\mathbf{z}_j| = (1 + \varepsilon) \text{dist}_{\mathbb{M}_g}$ and $|\mathbf{z}_{j+1}| = (1 - \varepsilon) \text{dist}_{\mathbb{M}_g}$. Now, because

$$\cosh(|\mathbf{z}_j|) = \cot(\alpha_{j0}) \cot(\eta_{j0}),$$

and because we have fixed $|\mathbf{z}_j|$, α_{j0} is minimal for maximal η_{j0} . Since we have fixed $|\mathbf{z}_j|$ and $|\mathbf{z}_{j+1}|$, it can be seen by a similar reasoning as in the proof of $\xi_{\min}$ that η_{j0} is maximal when $\arg(\mathbf{z}_{j+1}) - \arg(\mathbf{z}_j)$ is maximal, i.e., when $\arg(\mathbf{z}_{j+1}) - \arg(\mathbf{z}_j) = (1 + \varepsilon) \frac{\pi}{2g}$. Then η_{j0} is the solution η' of

$$\cos(\eta') = \frac{\tanh((1 - \varepsilon) \text{dist}_{\mathbb{M}_g})}{\tanh((1 + \varepsilon) \text{dist}_{\mathbb{M}_g})} \cdot \cos \left((1 + \varepsilon) \frac{\pi}{2g} - \eta' \right).$$

Finally, substituting the expression for m_j

$$\tanh(m_j) = \tanh(|\mathbf{z}_j|) \cos(\eta_{j0}),$$

in

$$\sin(\alpha_{j0}) = \frac{\sinh(m_j)}{\sinh(|\mathbf{z}_j|)}$$

and plugging in $|\mathbf{z}_j| = (1 + \varepsilon) \text{dist}_{\mathbb{M}_g}$ and $\eta_{j0} = \eta'$ yields the result.

□

2.3.4 Line segments between sides of admissible polygons

The lower bounds for m_j, ξ_j, d_j and α_j given in Lemma 2.10 allow us to state the following extension of Lemma 2.5 to hyperelliptic surfaces in the neighborhood $\mathbb{M}_g(\varepsilon)$. Here, segments of separation k between the sides of an admissible polygon P are defined in a similar way as for segments between the sides of D_g .

Lemma 2.11. *Let $0 < \varepsilon < 1$ and let $m_{\min}, \xi_{\min}, d_{\min}$ and $\alpha_{\min}$ be as in Lemma 2.10. If*

$$2 \sin \left(\frac{(g+1)\pi}{12g} \right) \cosh(m_{\min}) \sin(\xi_{\min}) > 1, \quad (2.6)$$

$$4 \sin^2 \left(\frac{(g+1)\pi}{12g} \right) (1 + \sinh^2(\frac{1}{2}d_{\min}) \sin^2(\alpha_{\min})) > 1, \quad (2.7)$$

then for every $\mathbb{M}_P \in \mathbb{M}_g(\varepsilon)$, the admissible polygon P satisfies the following two properties:

1. The length of a segment of separation at least 2 is larger than $\frac{1}{2} \text{sys}(\mathbb{M}_P)$.
2. The length of two consecutive 1-separated segment is larger than $\frac{1}{2} \text{sys}(\mathbb{M}_P)$.

Proof. First, we show that Inequality (2.6) implies that the length of a segment of separation at least 2 is at least $\frac{1}{2} \text{sys}(\mathbb{M}_P)$. Let γ be a segment of separation $k \geq 2$ between (without loss of generality) sides $[z_0, z_1]$ and $[z_k, z_{k+1}]$. The length of γ is at least the length of the common orthogonal $\gamma^\perp$ of sides $[z_0, z_1]$ and $[z_k, z_{k+1}]$. By symmetry of the argument, we can assume without loss of generality that $m_0 \leq m_k$. Note that for fixed $\xi_0 + \xi_1 + \dots + \xi_{k-1}$, the length of $\gamma^\perp$ is monotonically increasing as function of m_k . Therefore, we may assume without loss of generality that $m_k = m_0$. By a similar computation as in the proof of Part 1 of Lemma 2.5 we see that

$$\cosh(\frac{1}{2} \text{length}(\gamma^\perp)) = \cosh(m_0) \sin \left(\frac{1}{2} \sum_{j=0}^{k-1} \xi_j \right).$$

Because $m_0 > m_{\min}$ and $\xi_j > \xi_{\min}$, it follows that

$$\cosh(\frac{1}{2} \text{length}(\gamma^\perp)) > \cosh(m_{\min}) \sin(\frac{1}{2} k \xi_{\min}).$$

Since $\text{length}(\gamma) \geq \text{length}(\gamma^\perp)$ and $k \geq 2$, we conclude that

$$\cosh(\frac{1}{2} \text{length}(\gamma)) > \cosh(m_{\min}) \sin(\xi_{\min}). \quad (2.8)$$

It is known that [9]

$$\cosh(\frac{1}{4} \text{sys}(\mathbb{M}_P)) \leq \frac{1}{2 \sin \left(\frac{(g+1)\pi}{12g} \right)},$$

which, combined with Inequalities (2.6) and (2.8), yields $\text{length}(\gamma) > \frac{1}{2} \text{sys}(\mathbb{M}_P)$, as we wanted to show.

Second, we show that Inequality (2.7) implies that the length of two consecutive 1-separated segments is at least $\frac{1}{2} \text{sys}(\mathbb{M}_P)$. Let $\gamma_1 \cup \gamma_2$ be a pair of consecutive 1-separated segments. The proof of Part 3 of Lemma 2.5 does not use the metric of the generalized Bolza surfaces, so this property holds for all hyperelliptic surfaces as well, i.e., every pair of consecutive 1-separated segments consists of exactly one 1-segment and one $(4g-1)$ -segment. Therefore, we can assume without loss of generality that γ_1 is a 1-segment between $[z_0, z_1]$ and $[z_1, z_2]$ and γ_2 is a $(4g-1)$ -segment between $[z_{2g+1}, z_{2g+2}]$ and $[z_{2g}, z_{2g+1}]$. By the same reasoning as in Part 4 of Lemma 2.5, it can be shown that

$$\sinh(\tfrac{1}{2} \text{length}(\gamma_1 \cup \gamma_2)) \geq \sinh(\tfrac{1}{2} d_1) \sin(\alpha_1).$$

Since $d_1 > d_{\min}$ and $\alpha_1 > \alpha_{\min}$, it follows that

$$\cosh(\tfrac{1}{2} \text{length}(\gamma_1 \cup \gamma_2)) \geq \sqrt{1 + \sinh^2(\tfrac{1}{2} d_{\min}) \sin^2(\alpha_{\min})} \quad (2.9)$$

Combining

$$\cosh(\tfrac{1}{4} \text{sys}(\mathbb{M}_P)) \leq \frac{1}{2 \sin\left(\frac{(g+1)\pi}{12g}\right)},$$

with Inequalities (2.7) and (2.9) yields $\text{length}(\gamma_1 \cup \gamma_2) > \frac{1}{2} \text{sys}(\mathbb{M}_P)$, as we wanted to show. $\square$

Given g and ε it is straightforward to check numerically whether Inequalities (2.6) and (2.7) are satisfied. Using a computer algebra program we found that they are satisfied for $2 \leq g \leq 10^4$ and $\varepsilon = 0.04(\log(g))^{-1}$. The constant 0.04 is optimal up to two decimal digits: if $\varepsilon = 0.05(\log(g))^{-1}$, then Inequality (2.7) fails for $g \geq 14$. In the next section we will see that Inequalities (2.6) and (2.7) being satisfied is sufficient for proving the result of Theorem 2.8. Therefore, taking $c = 0.04$ in the statement of Theorem 2.8 is sufficient for $2 \leq g \leq 10^4$.

2.3.5 Proof of Theorem 2.8

Proof. (Theorem 2.8)

If for some $0 < \varepsilon < 1$ Inequalities (2.6) and (2.7) are satisfied, then segments of separation at least 2 and pairs of consecutive 1-separated segments have length at least $\frac{1}{2} \text{sys}(\mathbb{M}_P)$ by Lemma 2.11. Therefore, we can reuse the proof of Lemma 2.6 to show that a systole of $\mathbb{M}_P$ corresponds to either a single segment (of separation $2g$) or a pair of consecutive 1-separated segments between the sides of P .

Hence, it is sufficient to prove that Inequalities (2.6) and (2.7) are satisfied for $\varepsilon = c(\log(g))^{-1}$ for some positive $c \in \mathbb{R}$. We do this by approximating each of the terms in these inequalities depending on ε by their 0-th order Taylor expansion and giving an upper bound for the estimation error. The remainder of the proof is given in Appendix A.1. $\square$

2.4 Word length of systoles

Recall from the proof of Theorem 2.1 (in particular, from the construction in Lemma 2.3) that the closed geodesics with length $\text{sys}(\mathbb{M}_g)$ on the generalized Bolza surface $\mathbb{M}_g$ of genus $g \geq 2$ correspond to elements of Γ_g consisting of the product of precisely *two* generators. More generally, given an arbitrary hyperbolic surface $\mathbb{M} = \mathbb{D}/\Gamma$, one could ask whether there exists an upper bound for the number of generators that need to be multiplied to obtain an element of Γ corresponding to a systole of $\mathbb{M}$. The number of generators that need to be multiplied is called the *word length* of the element of Γ (see Definition 2.14). For convenience, we abbreviate ‘element of Γ corresponding to a systole of $\mathbb{M}$ ’ to *systole element of Γ* . The following theorem states that there is no bound on the word length of a systole element if there are no restrictions on the generating set of Γ .

Theorem 2.12. *Let $\mathbb{M} = \mathbb{D}/\Gamma$ be a closed hyperbolic surface of genus $g \geq 2$. For every $N_g \in \mathbb{N}$ there exists a generating set $\mathcal{G}$ of Γ consisting of hyperbolic translations, such that the word length of every systole element of Γ with respect to $\mathcal{G}$ is at least N .*

The idea of the proof is the following. First, we consider a ‘standard’ generating set $\mathcal{A}$ of Γ satisfying the commutator relation. We will modify this generating set to obtain another generating set $\mathcal{G}_m$ with $m \in \mathbb{N}$. Second, it is known that there are finitely many conjugacy classes of systole elements of Γ . We define the notion of a *G-chain* for a generator $G \in \mathcal{A}$ and use the length of *G-chains* for different values of G to describe systole elements. This allows us to find an m such that $\mathcal{G} = \mathcal{G}_m$ has the desired properties.

We will first introduce some notions from geometric group theory. First, note that the Fuchsian group Γ is the fundamental group of the closed hyperbolic surface $\mathbb{M} = \mathbb{D}/\Gamma$ of genus $g \geq 2$. Let $\mathcal{A} = \{A_1, B_1, \dots, A_g, B_g\}$ be the usual generating set for the fundamental group that satisfies the relation $[A_1, B_1] \dots [A_g, B_g] = \mathbb{1}$, where $[A_i, B_i]$ denotes the commutator of A_i and B_i . Every element $X \in \Gamma$ can be represented as a word $X = X_1^{\epsilon_1} \dots X_\ell^{\epsilon_\ell}$, where $X_i \in \mathcal{A}$ and $\epsilon_i = \pm 1$.

Definition 2.13. We call $X_1^{\epsilon_1} \dots X_\ell^{\epsilon_\ell}$ a *reduced word* if for all $i = 1, \dots, \ell - 1$ either $X_i \neq X_{i+1}$ or $\epsilon_i = \epsilon_{i+1}$. We call $X_1^{\epsilon_1} \dots X_\ell^{\epsilon_\ell}$ a *cyclically reduced word* if every cyclic permutation of the word is reduced.

Any word can be written as a reduced word after reduction, i.e., deleting all instances of GG^{-1} or $G^{-1}G$ for $G \in \mathcal{A}$. Similarly, any word can be written as a cyclically reduced word by cyclic reduction. These representations are not unique as we can add or delete instances of the group relation. Furthermore, from the process of reduction it is clear that a word is reduced if and only if all subwords consisting of two consecutive generators are reduced.

We now define the word length of an element of $\mathcal{A}$.

Definition 2.14. If $X = X_1^{\epsilon_1} \cdots X_\ell^{\epsilon_\ell}$ is a cyclically reduced word, then we call $\ell_{\mathcal{A}}(X) := \ell$ the word length of X with respect to the generating set $\mathcal{A}$.

In some texts, the word length of X is defined using the reduction of X instead of the cyclic reduction of X . Sometimes, word length is defined without any reduction whatsoever. For example, the word length of $A_1 B_1 B_1^{-1} A_2 A_1^{-1}$ is 5, 3 or 1 when the word length is defined without reduction, with reduction and with cyclic reduction respectively. Since cyclic permutation of a word can be obtained as a conjugation in the group Γ , there is a natural correspondence between cyclically reduced words and conjugacy classes of words in the group Γ . Note, however, that this is not a one-to-one correspondence, as instances of the group relation can be added or deleted without changing the conjugacy class. Nevertheless, this is the reason that we use cyclic reduction to define word length.

Instead of $\mathcal{A}$, consider a different generating set $\mathcal{G} = \{S_1, T_1, \dots, S_g, T_g\}$. Since by definition S_i and T_i are elements of Γ , we can write S_i and T_i as words in $\mathcal{A}$. Now, given a representation of a word X in $\mathcal{G}$, we can write X as word in $\mathcal{A}$ by substituting the representations of S_i and T_i in $\mathcal{A}$. Note that X being reduced as word in $\mathcal{G}$ does not mean that X is also reduced as word in $\mathcal{A}$.

Example 2.15. Let $g = 1$. Let $S_1 = A_1$ and $T_1 = A_1 B_1$. Clearly, $\mathcal{G} = \{S_1, T_1\}$ is a generating set for the group generated by $\mathcal{A} = \{A_1, B_1\}$. Consider the word $X = S_1^{-1} T_1$. Observe that X is reduced as word in $\mathcal{G}$. Using the representations of S_1 and T_1 as words in $\mathcal{A}$, we see that $X = A_1^{-1} A_1 B_1$. In particular, X is not reduced as word in $\mathcal{A}$.

To prove Theorem 2.12 we use the generating set defined below.

Definition 2.16. Let $\mathcal{A}$ be the generating set defined above. Let $m \in \mathbb{N}$. Define $\mathcal{G}_m = \{S_1, T_1, \dots, S_g, T_g\}$, where

$$\begin{aligned} S_i &= B_i^m A_i, \\ T_i &= (B_i^m A_i)^m B_i. \end{aligned}$$

It is straightforward to check that $\mathcal{G}_m$ is indeed a generating set. The generating set $\mathcal{G}_m$ was defined in such a way that S_i and T_i contain sequences of m consecutive letters B_i . More generally, we define a G -chain in a word X as a sequence of consecutive letters G in a cyclically reduced representation of X .

Definition 2.17. Let $X = X_1^{\epsilon_1} \cdots X_\ell^{\epsilon_\ell}$ be cyclically reduced. We say that X contains a G -chain of length L if there exists $i = 1, \dots, \ell$ such that $X_j = G$ for $j = i, i+1, \dots, i+L-1$ and $\epsilon_i = \epsilon_{i+1} = \dots = \epsilon_{i+L-1} = 1$, where the indices are counted modulo ℓ .

A G^{-1} -chain is defined similarly, but with $\epsilon_i = \dots = \epsilon_{i+L-1} = -1$. The existence of G -chains of a certain length may depend on the representation of X .

Example 2.18. Consider $g = 1$. Consider the word $X = A_1^2$. Clearly, X contains a A_1 -chain of length 2. However, since $A_1 B_1 A_1^{-1} B_1^{-1} = 1$, X can also be represented by $A_1^2 = A_1 \cdot A_1 = A_1 B_1 A_1 B_1^{-1}$. With this representation, X only contains a A_1 -chain of length 1 (in fact, it contains two such chains). Note that it is important that both representations are cyclically reduced. For example, $A_1^3 A_1^{-1}$ would also be a representation of X , but it is not reduced. We want to avoid this, because in this way we would be able to construct A_1 -chains of arbitrarily large lengths.

The generators $S_i^{\pm 1}$ and $T_i^{\pm 1}$ of $\mathcal{G}_m$ contain $B_i^{\pm 1}$ -chains of length m . The following lemma states that the same holds for products of at most m generators in $\mathcal{G}_m$.

Lemma 2.19. *Let $X = X_1^{\epsilon_1} \cdots X_\ell^{\epsilon_\ell}$ with $X_i \in \mathcal{G}_m$ be a cyclically reduced word in $\mathcal{G}_m$. Let $\tilde{X}$ be the representation of X as word in $\mathcal{A}$ obtained from substituting the representations of the generators S_i and T_i as words in $\mathcal{A}$. If $\ell \leq m$, then there exists a cyclical reduction of $\tilde{X}$ that contains a $B_j^{\pm 1}$ -chain of length at least m for some $j = 1, \dots, g$.*

Proof. Because X is a non-trivial word, there exists j such that X contains $S_j^{\pm 1}$ or $T_j^{\pm 1}$. Because each of $S_j^{\pm 1}$ and $T_j^{\pm 1}$ contains at least one $B_j^{\pm 1}$ -chain of length m , it is sufficient to show that at least one of these $B_j^{\pm 1}$ -chains remains after cyclically reducing $\tilde{X}$.

First, suppose that X contains at least one T_j and no T_j^{-1} (the proof for X containing at least one T_j^{-1} and no T_j is similar, so we omit it). Every T_j contains m B_j -chains of length m . Because X does not contain any T_j^{-1} , every B_j^{-1} -chain of length m that may possibly reduce one of the B_j -chains of length m of T_j is contained in some S_j^{-1} . Because X is the product of at most m generators, it contains at most $m - 1$ generators S_j^{-1} in addition to T_j , so at most $m - 1$ B_j -chains of length m are possible reduced in this way. It follows that at least one B_j -chain of length m remains in a cyclical reduction of $\tilde{X}$.

Second, suppose that X contains at least one T_j and at least one T_j^{-1} . The B_j^{-1} -chains of length m in T_j^{-1} can only be involved in a reduction with the B_j -chains of length m in T_j if T_j and T_j^{-1} are separated by a subword W of X that reduces to 1 when cyclically reducing $\tilde{X}$. Because X is cyclically reduced, this can only occur if W is the group relation (or a product of group relations), but a straightforward computation shows that the group relation in terms of $\mathcal{G}$ consists of (far) more than m generators. Therefore, the B_j^{-1} -chains of length m in T_j^{-1} are not involved in a reduction with the B_j -chains of length m in T_j . Again, because every S_j^{-1} in X reduces at most one B_j -chain of length m in T_j , we conclude that at least one B_j -chain of length m remains in $\tilde{X}$.

Third, suppose that X does not contain any T_j nor T_j^{-1} . Then X contains S_j or S_j^{-1} . We will only treat the case in which X contains at least one S_j , since

the proof for X containing S_j^{-1} is similar. The S_j in X contains one B_j -chain of length m . Because X does not contain any T_j , the B_j -chain of length m in S_j can only be reduced by a B_j^{-1} -chain in S_j^{-1} . Similar to the argument above for T_j and T_j^{-1} , a B_j -chain of length m in S_j can only be involved in a reduction with a B_j^{-1} -chain of length m in S_j^{-1} if they are separated by a subword of X that reduces to $\mathbb{1}$ when cyclically reducing $\tilde{X}$. Because X is a product of at most m generators, this is not possible for X , so we conclude that at least one B_j -chain of length m remains in $\tilde{X}$ after cyclical reduction. $\square$

Now, we give a proof of Theorem 2.12.

Proof. Let $\mathcal{A} = \{A_1, B_1, \dots, A_g, B_g\}$ be the standard generating set as before. Let $N \in \mathbb{N}$. Up to conjugation there are only finitely many elements of Γ , say $W_1, \dots, W_n$, corresponding to systoles of $\mathbb{M}$. For each W_j , let $L_G(W_j)$ be the length of the maximal G -chain contained in any cyclical reduction of W_j . This number is well-defined, since it is not possible to construct arbitrary long G -chains in a cyclical reduction of W_j by adding or subtracting instances of the group relation. Then, choose $m \in \mathbb{N}$ such that

$$m > \max\{N, L_G(W_j) : G \in \mathcal{A}, j = 1, \dots, n\}.$$

We claim that $\mathcal{G} := \mathcal{G}_m$ is the generating set that we are looking for. Consider an arbitrary W_j . We will show that every representation of W_j in $\mathcal{G}$ is the product of at least N generators. Represent $W_j = X_1^{\epsilon_1} \cdots X_\ell^{\epsilon_\ell}$ as a cyclically reduced (non-trivial) word in $\mathcal{G}$, i.e., $X_i \in \mathcal{G}$. Consider any $B_i \in \mathcal{A}$. By definition of m , we know that the longest B_i -chain in W_j has length smaller than m . By Lemma 2.19, we see that $\ell > m$. Since $m > N$, we conclude that $\ell > N$, which finishes the proof. $\square$

Chapter 3

Computing Delaunay triangulations of generalized Bolza surfaces

3.1 Introduction

Lawson's well-known incremental algorithm that computes Delaunay triangulations using edge flips in the Euclidean plane [52] has recently been proved to generalize on hyperbolic surfaces [30]. However, the experience gained in the CGAL project for many years has shown that Bowyer's algorithm [19] leads to a cleaner code, much easier to maintain; there is actually work in progress in CGAL to replace Lawson's flip algorithm, in triangulation packages that are still using it, by Bowyer's algorithm. In the context of quotient spaces Bowyer's algorithm was used already in the CGAL packages for 3D flat quotient spaces [23] and for the Bolza surface [46]. To the best of our knowledge, the latter package is the only available software for a hyperbolic surface. The advantages of Bowyer's algorithm largely compensate the constraint that it intrinsically requires that the Delaunay triangulation be a simplicial complex.

In this chapter, we study the extension of this approach to the generalized Bolza surfaces. For a set of points $\mathcal{Q} \subset \mathbb{M}$, let $\delta(\mathcal{Q})$ be the diameter of the largest disks in $\mathbb{D}$ that do not contain any point in $\pi^{-1}(\mathcal{Q})$. In Section 3.2 we recall the following *validity condition* [24, 16]: If a finite set $\mathcal{Q}$ of points on the surface $\mathbb{M}$ satisfies the inequality

$$\delta(\mathcal{Q}) < \frac{1}{2} \text{sys}(\mathbb{M}) \quad (\text{condition (3.1) in Proposition 3.2})$$

then Bowyer's algorithm can be extended to the computation of the Delaunay triangulation of any finite set of points $\mathcal{S}$ on $\mathbb{M}$ containing $\mathcal{Q}$. Before actually inserting the input points, the algorithm performs a preprocessing step consisting of computing the Delaunay triangulation of an appropriate (but small) set $\mathcal{Q}$

satisfying this validity condition; following the terminology of previous papers [16, 24], we refer to the points of $\mathcal{Q}$ as *dummy points*. When sufficiently many and well-distributed input points have been inserted, the validity condition is satisfied without the dummy points, which can then be removed. This approach was used in the CGAL package for the Bolza surface $\mathbb{M}_2$ [45, 46].

Other practical approaches for (flat) quotient spaces start by computing in a finite-sheeted covering space [24] or in the universal covering space [62], thus requiring the duplication of some input points. In contrast to this approach, our algorithm proceeds directly on the surface, thus circumventing the need for duplicating any input points. While the number of copies of duplicated points in approaches using covering spaces is small, the number of duplicated input points is always much larger than the number of dummy points that could instead be added to the set of input points in our approach. Moreover, to the best of our knowledge, the number of required copies in the case of hyperbolic surfaces is largely unknown; first bounds have been obtained recently [29].

Results.

We describe the extension of Bowyer’s algorithm to the case of the generalized Bolza surface $\mathbb{M}_g$ in Section 3.2. Then, we derive bounds on the number of dummy points necessary to satisfy the validity condition (Propositions 3.4 and 3.5 in Section 3.3), yielding the following result:

Theorem 3.1. *The number of dummy points that must be added on $\mathbb{M}_g$ to satisfy the validity condition (3.1) grows as $\Theta(g)$.*

Because the validity condition depends on the systole of a hyperbolic surface, that is, on a global property of the surface, the resulting dummy point set might contain more dummy points than necessary in parts of the surface that are far from small non-contractible closed curves. Therefore, we will consider a *local validity condition* that does not depend on the systole of the hyperbolic surface, but on the *local systole*, i.e., the length of the smallest non-contractible closed curve passing through a given point on the hyperbolic surface. In Section 3.4 we briefly discuss a local validity condition and give a lower bound for the number of resulting dummy points.

In Section 3.5 we propose two algorithms to compute dummy points. For ease of computation, these algorithms depend on the validity condition (not on the local validity condition). The first algorithm is based on the well-known Delaunay refinement algorithm for mesh generation [66]. Using a packing argument, we prove that it provides an asymptotically optimal number of dummy points (Theorem 3.12). The second algorithm modifies the refinement algorithm so as to yield a symmetric dummy point set, at the expense of a slightly larger output size $\Theta(g \log g)$ (Theorem 3.13); this symmetry may be interesting for some applications [25]. The two algorithms have been implemented and we quickly present results for small genera $g = 2$ and $g = 3$.

Finally, in Section 3.6, we describe the data structure that we are using to support the extension of Bowyer’s algorithm to generalized Bolza surfaces. We also discuss the algebraic degree of the predicates used in the computations and present experimental results.

3.2 Computing Delaunay triangulations

3.2.1 Bowyer’s algorithm in the Euclidean plane

There exist various algorithms to compute Delaunay triangulations in Euclidean spaces. Bowyer’s algorithm [19, 75] has proved its efficiency in CGAL [47].

Let us focus here on the two-dimensional case. Let $\mathcal{P}$ be a set of points in the Euclidean plane $\mathbb{E}$ and $\text{DT}_{\mathbb{E}}(\mathcal{P})$ its Delaunay triangulation. Let $\mathbf{p} \notin \mathcal{P}$ be a point in the plane to be inserted in the triangulation. Bowyer’s algorithm performs the insertion as follows.

1. Find the set of triangles of $\text{DT}_{\mathbb{E}}(\mathcal{P})$ that are *in conflict* with $\mathbf{p}$, i.e., whose open circumscribing disks contain $\mathbf{p}$;
2. Delete each triangle in conflict with $\mathbf{p}$; this creates a “hole” in the triangulation;
3. Repair the triangulation by forming new triangles with $\mathbf{p}$ and each edge of the hole boundary to obtain $\text{DT}_{\mathbb{E}}(\mathcal{P} \cup \{\mathbf{p}\})$.

Degeneracies can be resolved using a symbolic perturbation technique, which actually works in any dimension [32].

An illustration is given in Figure 3.1. The first step of the insertion of $\mathbf{p}$ uses

Figure 3.1: Insertion of a point in a Delaunay triangulation with Bowyer’s incremental algorithm.

geometric computations, whereas the next two are purely combinatorial. This is another reason why this algorithm is favored in CGAL: it allows for a clean separation between combinatorics and geometry, as opposed to an insertion by flips, in which geometric computations and combinatorial updates would alternate.

Note that the combinatorial part heavily relies on the fact that *the union of the triangles of $\text{DT}_{\mathbb{E}}(\mathcal{P})$ in conflict with $\mathbf{p}$ is a topological disk*. We will discuss this essential property in the next section.

3.2.2 Delaunay triangulations of points on hyperbolic surfaces

Let $\mathbb{M} = \mathbb{D}/\Gamma$ be a hyperbolic surface, as introduced in Section 1.1.6, with the associated projection map $\pi : \mathbb{D} \rightarrow \mathbb{M}$, and $F \subset \mathbb{D}$ a fundamental domain.

Let us consider a triangle t and a point p on $\mathbb{M}$. The triangle t is said to be *in conflict* with p if the circumscribing disk of one of the triangles in $\pi^{-1}(t)$ is in conflict with a point of $\pi^{-1}(p)$ in the fundamental domain. As noticed in the literature [14], the notion of conflict in $\mathbb{D}$ is the same as in $\mathbb{E}$, since for the Poincaré disk model, hyperbolic circles are Euclidean circles (see Section 1.1.2).

Let us now consider a finite set $\mathcal{P}$ of points on the surface $\mathbb{M}$ and a partition of $\mathbb{M}$ into triangles with vertex set $\mathcal{P}$. Assume that the triangles of the partition have no conflict with any of the vertices. Let $p \notin \mathcal{P}$ be a point on $\mathbb{M}$. The region C_p formed by the union of the triangles of the partition that are in conflict with p might not be a topological disk (see Figure 3.2). In such a case, the last step of Bowyer's algorithm could not directly apply, as there are multiple geodesics between p and any given point on the boundary of C_p .

Figure 3.2: Bowyer's insertion is not well defined when the conflict region is not a topological disk.

In order to be able to use Bowyer's algorithm on $\mathbb{M}$, the triangles on $\mathbb{M}$ without conflict with any vertex, together with their edges and their vertices, should form a *simplicial complex*. In other words, the graph of edges of the triangles should have no loops (1-cycles) or multiple edges (2-cycles).

For a set of points $\mathcal{Q} \subset \mathbb{M}$ we denote by $\delta(\mathcal{Q})$ the diameter of the largest disks in $\mathbb{D}$ that do not contain any point in $\pi^{-1}(\mathcal{Q})$. We will reuse the following result.

Proposition 3.2 (Validity condition [16]). *Let $\mathcal{Q} \subset \mathbb{M}$ be a set of points such that*

$$\delta(\mathcal{Q}) < \frac{1}{2} \text{sys}(\mathbb{M}). \quad (3.1)$$

Then, for any set of points $\mathcal{S} \subset \mathbb{M}$ such that $\mathcal{Q} \subseteq \mathcal{S}$, the graph of edges of the projection $\pi(\text{DT}_{\mathbb{D}}(\pi^{-1}(\mathcal{S})))$ has no 1- or 2-cycles.

This condition is stronger than just requiring that the Delaunay triangulation of $\mathcal{Q}$ be a simplicial complex: if only the latter condition holds, inserting more points could create cycles in the triangulation [24, Figure 3]; see also Remark 3.8 below.

The proof is easy, we include it for completeness.

Proof. Assume that condition (3.1) holds. For each edge e of the (infinite) Delaunay triangulation $\text{DT}_{\mathbb{D}}(\pi^{-1}(\mathcal{Q}))$ in $\mathbb{D}$, there exists an empty ball having the end-points of e on its boundary, so, the length of e is not larger than $\delta(\mathcal{Q})$. Assume now that there is a 2-cycle formed by two edges $\pi(e_1)$ and $\pi(e_2)$ in $\pi(\text{DT}_{\mathbb{D}}(\pi^{-1}(\mathcal{Q})))$, then the length of the non-contractible loop that they form is the sum of the lengths of e_1 and e_2 , which is at most $2\delta(\mathcal{Q})$ and smaller than $\text{sys}(\mathbb{M})$. This is impossible by definition of $\text{sys}(\mathbb{M})$, so, there is no 2-cycle in $\pi(\text{DT}_{\mathbb{D}}(\pi^{-1}(\mathcal{Q})))$.

As the diameter of the largest empty disks does not increase with the addition of new points, the same holds for any set $\mathcal{S} \supseteq \mathcal{Q}$. $\square$

The most obvious example of a set that does not satisfy the validity condition is a single point: each edge of the projection is a 1-cycle. The condition is satisfied when the set contains sufficiently many and well-distributed points.

Definition 3.3. Let $\mathcal{S} \subset \mathbb{M}$ be a set of points satisfying the validity condition (3.1). The *Delaunay triangulation of $\mathbb{M}$ defined by $\mathcal{S}$* is then defined as $\pi(\text{DT}_{\mathbb{D}}(\pi^{-1}(\mathcal{S})))$ and denoted by $\text{DT}_{\mathbb{M}}(\mathcal{S})$.

As for the Bolza surface [16], Proposition 3.2 naturally suggests a way to adapt Bowyer's algorithm to compute $\text{DT}_{\mathbb{M}}(\mathcal{P})$ for a given set $\mathcal{P}$ of n points on $\mathbb{M}$:

- Initialize the triangulation as the Delaunay triangulation $\text{DT}_{\mathbb{M}}(\mathcal{Q})$ of $\mathbb{M}$ defined by an artificial set of *dummy points* $\mathcal{Q}$ that satisfies condition (3.1);
- Compute incrementally the Delaunay triangulation $\text{DT}_{\mathbb{M}}(\mathcal{Q} \cup \mathcal{P})$ by inserting the points $p_1, p_2, \dots, p_k, \dots, p_n$ of $\mathcal{P}$ one by one, i.e., for each new point p_k :
 - find all triangles of the Delaunay triangulation $\text{DT}_{\mathbb{M}}(\mathcal{Q} \cup \{p_1, \dots, p_{k-1}\})$ that are in conflict with p_k ; let C_{p_k} denote their union; since $\mathcal{Q}$ satisfies the validity condition, C_{p_k} is a topological disk;
 - delete the triangles in C_{p_k} ;
 - repair the triangulation by forming new triangles with p_k and each edge of the boundary of C_{p_k} ;
- Remove from the triangulation all points of $\mathcal{Q}$ whose removal does not violate the validity condition.

We ignore degeneracies here; they can be resolved as in the case of flat orbit spaces [24]. Depending on the size and distribution of the input set $\mathcal{P}$, the final Delaunay triangulation of $\mathbb{M}$ might have some or all of the dummy points as vertices. If $\mathcal{P}$ already satisfies the validity condition then no dummy point is left.

3.3 Bounds on the number of dummy points

In the following proposition we show that a dummy point set exists and give an upper bound for its cardinality. The proof is non-constructive, but we will construct dummy point sets for generalized Bolza surfaces in Section 3.5.

Proposition 3.4. *Let $\mathbb{M}$ be a hyperbolic surface of genus g with systole $\text{sys}(\mathbb{M})$. Then there exists a point set $\mathcal{Q} \subset \mathbb{M}$ satisfying the validity condition (3.1) with cardinality*

$$|\mathcal{Q}| \leq \frac{2(g-1)}{\cosh(\frac{1}{8} \text{sys}(\mathbb{M})) - 1}.$$

Proof. Let $\mathcal{Q}$ be a maximal set of points such that for all distinct $p, q \in \mathcal{Q}$ we have $d(p, q) \geq \frac{1}{4} \text{sys}(\mathbb{M})$. By maximality, we know that for all $x \in \mathbb{M}$ there exists $p \in \mathcal{Q}$ such that $d(x, p) < \frac{1}{4} \text{sys}(\mathbb{M})$: if this is not the case, i.e., if there exists $x \in \mathbb{M}$ such that $d(x, p) \geq \frac{1}{4} \text{sys}(\mathbb{M})$ for all $p \in \mathcal{Q}$, then we can add x to $\mathcal{Q}$, which contradicts maximality of $\mathcal{Q}$. Hence, for any $x \in \mathbb{M}$ the largest disk centered at x and not containing any points of $\mathcal{Q}$ has diameter less than $\frac{1}{2} \text{sys}(\mathbb{M})$, which implies $\delta(\mathcal{Q}) < \frac{1}{2} \text{sys}(\mathbb{M})$.

To prove the statement on the cardinality of $\mathcal{Q}$, denote the open disk centered at $p \in \mathcal{Q}$ with radius R by $B_p(R)$. The disk $B_p(\frac{1}{8} \text{sys}(\mathbb{M}))$ for $p \in \mathcal{Q}$ is embedded in $\mathbb{M}$, because its radius is smaller than $\frac{1}{2} \text{sys}(\mathbb{M})$. Furthermore, by construction of $\mathcal{Q}$,

$$B_p(\frac{1}{8} \text{sys}(\mathbb{M})) \cap B_q(\frac{1}{8} \text{sys}(\mathbb{M})) = \emptyset$$

for all distinct $p, q \in \mathcal{Q}$. Hence, the cardinality of $\mathcal{Q}$ is bounded from above by the number of disjoint embedded disks of radius $\frac{1}{8} \text{sys}(\mathbb{M})$ that we can fit in $\mathbb{M}$. We obtain

$$|\mathcal{Q}| \leq \frac{\text{area}(\mathbb{M})}{\text{area}(B_p(\frac{1}{8} \text{sys}(\mathbb{M})))} = \frac{4\pi(g-1)}{2\pi(\cosh(\frac{1}{8} \text{sys}(\mathbb{M})) - 1)} = \frac{2(g-1)}{\cosh(\frac{1}{8} \text{sys}(\mathbb{M})) - 1}.$$

□

Similarly, in the next proposition we state a lower bound for the cardinality of a dummy point set.

Proposition 3.5. *Let $\mathbb{M}$ be a hyperbolic surface of genus $g \geq 2$. Let $\mathcal{Q}$ be a set of points in $\mathbb{M}$ such that the validity condition (3.1) holds. Then*

$$|\mathcal{Q}| > \left(\frac{\pi}{\pi - 6 \operatorname{arccot}(\sqrt{3} \cosh(\frac{1}{4} \text{sys}(\mathbb{M})))} - 1 \right) \cdot 2(g-1).$$

The proof uses the following lemma.

Lemma 3.6. *Let T be a hyperbolic triangle with a circumscribed disk of radius R . Then*

$$\text{area}(T) \leq \pi - 6 \operatorname{arccot}(\sqrt{3} \cosh(R)).$$

Lemma 3.6 is the special case $m = 3$ of the following lemma. A proof was given in my master's thesis [33], but for completeness it is included here as well.

Lemma 3.7. *Let P be a convex hyperbolic m -gon for $m \geq 3$ 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*

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

Proof. A lower bound for the circumradius of a polygon given the area of the polygon is given in the literature [61]. We use the same approach to prove Lemma 3.7.

Consider $m = 3$. Divide P into three pairs of right-angled triangles with angles θ_j at the center of the circumscribed circle, angles α_j at the vertices and right angles at the edges of P (see Figure 3.3).

Figure 3.3: Division of P into three pairs of right-angled triangles

By the second hyperbolic cosine rule

$$\cosh R = \cot \theta_j \cot \alpha_j$$

for $j = 1, 2, 3$. Furthermore, $\sum_{j=1}^3 \theta_j = \pi$ and $A = \pi - 2 \sum_{j=1}^3 \alpha_j$. Therefore, maximizing A reduces to minimizing

$$f(\theta_1, \theta_2, \theta_3) = \sum_{j=1}^3 \operatorname{arccot}(\cosh R \tan \theta_j) \quad (3.2)$$

subject to the constraints $\sum_{j=1}^3 \theta_j = \pi$ and $0 \leq \theta_j < \pi$, i.e., minimizing (3.2) over the triangle in $\mathbb{R}^3$ with vertices $(\pi, 0, 0), (0, \pi, 0), (0, 0, \pi)$. We parametrize

this triangle as follows

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

for $0 < s < \pi/2$ and $|t| \leq s$. By (3.2), we can view f as a function of s and t . First, we fix s and minimize over t . Then

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

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 .

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

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,

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

so

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

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

$$A(R) = (m - 2)\pi - 2m\alpha_j,$$

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

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

□

We continue with the proof of Proposition 3.5.

Proof. Denote the number of vertices, edges and triangles in the (simplicial) Delaunay triangulation $\text{DT}_{\mathbb{M}}(\mathcal{Q})$ of $\mathbb{M}$ by k_0, k_1 and k_2 , respectively. We know that $3k_2 = 2k_1$, since every triangle consists of three edges and every edge belongs to two triangles. By Euler's formula,

$$k_0 - k_1 + k_2 = 2 - 2g,$$

so

$$k_2 = 4g - 4 + 2k_0. \quad (3.3)$$

Consider an arbitrary triangle t in $\text{DT}_{\mathbb{M}}(\mathcal{Q})$. Because the validity condition holds, the circumradius of t is smaller than $\frac{1}{2}$ sys. By Lemma 3.6, $\text{area}(t) < \pi - 6 \operatorname{arccot}(\sqrt{3} \cosh(\frac{1}{4} \text{sys}))$. Because $\mathbb{M}$ has area $4\pi(g-1)$, it follows that

$$k_2 > \frac{4\pi(g-1)}{\pi - 6 \operatorname{arccot}(\sqrt{3} \cosh(\frac{1}{2} \text{sys}))}. \quad (3.4)$$

Combining (3.3) and (3.4) yields the result. $\square$

To show that our lower and upper bounds are meaningful, we consider the asymptotics of these bounds for a family of surfaces of which the systoles are 1. contained in a compact subset of $\mathbb{R}_{>0}$, 2. arbitrarily close to zero, or 3. arbitrarily large.

1. If the systoles of the family of surfaces are contained in a compact subset of $\mathbb{R}_{>0}$, which is the case for the generalized Bolza surfaces, then the upper bound is of order $O(g)$ and the lower bound of order $\Omega(g)$. Hence, a minimum dummy point set has cardinality $\Theta(g)$.
2. If $\text{sys}(\mathbb{M}) \rightarrow 0$, then $\cosh(\frac{1}{8} \text{sys}(\mathbb{M})) - 1 \sim \frac{1}{2}(\frac{1}{8} \text{sys}(\mathbb{M}))^2$, so our upper bound is of order $g \cdot O(\text{sys}(\mathbb{M})^{-2})$. In a similar way, it can be shown that

$$\left(\frac{\pi}{\pi - 6 \operatorname{arccot}(\sqrt{3} \cosh(\frac{1}{4} \text{sys}(\mathbb{M})))} - 1 \right) \sim \frac{64\pi}{3\sqrt{3} \text{sys}(\mathbb{M})^2},$$

which means that our lower bound is of order $g \cdot \Omega(\text{sys}(\mathbb{M})^{-2})$. It follows that in this case a minimum dummy point set has cardinality $g \cdot \Theta(\text{sys}(\mathbb{M})^{-2})$.

3. Finally, consider the case when $\text{sys}(\mathbb{M}) \rightarrow \infty$ when $g \rightarrow \infty$. Since $\text{sys}(\mathbb{M}) \leq 2 \log(4g-2)$ for all hyperbolic surfaces of genus g (see Lemma 1.14 in Section 1.1.8), we only consider the case where $\text{sys}(\mathbb{M}) \sim C \log g$ for some C with $0 < C \leq 2$. In this case, we can use $\cosh x \sim \frac{1}{2}e^x$ to deduce that our upper bound reduces to an expression of order $O(g^{1-\frac{1}{8}C})$. Similarly, by considering the Taylor expansion of the coefficient in the lower bound we see that the lower bound has cardinality $\Omega(g^{1-\frac{1}{4}C})$. Of our three cases, this is the only case in which there is a gap between the stated upper and lower bound.

Remark 3.8. Note that the validity condition (3.1) is stronger than just requiring that the Delaunay triangulation of $\mathcal{Q}$ be a simplicial complex. This can also be seen in the following way. In Chapter 4 we will show that every hyperbolic surface of genus g has a simplicial Delaunay triangulation with at most $151g$ vertices. In particular, this upper bound does not depend on $\text{sys}(\mathbb{M})$. Since the coefficient of $g - 1$ in the lower bound given in Proposition 3.5 goes to infinity as $\text{sys}(\mathbb{M})$ goes to zero, the minimal number of vertices of a set $\mathcal{Q}$ satisfying the validity condition is strictly larger than the number of vertices needed for a simplicial Delaunay triangulation of a hyperbolic surface with sufficiently small systole. Moreover, in the same chapter it is shown that for infinitely many genera g there exists a hyperbolic surface $\mathbb{M}$ of genus g which has a simplicial Delaunay triangulation with $\Theta(\sqrt{g})$ vertices. Hence, the number of vertices needed for a simplicial Delaunay triangulation and a point set satisfying the validity condition differs asymptotically as well.

3.4 Local validity condition

We have seen in the previous section that the number of dummy points needed to satisfy the validity condition (3.1) on a hyperbolic surface $\mathbb{M}$ with $\text{sys}(\mathbb{M}) \rightarrow 0$ is $g \cdot \theta(\text{sys}(\mathbb{M})^{-2})$. Recall that a dummy point set satisfies the validity condition if it is sufficiently dense and well-distributed, where dense and well-distributed are interpreted with respect to $\frac{1}{2} \text{sys}(\mathbb{M})$. In particular, if a hyperbolic surface has only one closed geodesic γ of length $\text{sys}(\mathbb{M})$ and the remainder of the surface does not contain ‘short’ closed geodesics, then a dummy point set that is constructed to satisfy the validity condition might be denser than necessary in the parts of the surface far from γ .

In this section, we will introduce a *local validity condition* and give a lower bound on the number of dummy points needed to satisfy this local validity condition. We emphasize that at this stage the local validity condition that we present here is not practically usable in an implementation. For further details, we refer to the discussion at the end of the section.

Let us now state the local validity condition. Given a set of points $\mathcal{Q}$ on a hyperbolic surface $\mathbb{M}$ and a point p on $\mathbb{M}$ (not necessarily in $\mathcal{Q}$), let $\delta_p(\mathcal{Q})$ be the diameter of the largest disk centered at a point in $\pi^{-1}(p)$ that does not contain any point of $\pi^{-1}(\mathcal{Q})$ in its interior. Furthermore, let the local systole $\text{sys}_p(\mathbb{M})$ of $\mathbb{M}$ at p be the length of the shortest non-contractible closed curve on $\mathbb{M}$ passing through p .

Proposition 3.9 (Local validity condition). *Let $\mathcal{Q}$ be a finite point set on a hyperbolic surface $\mathbb{M} = \mathbb{D}/\Gamma$. Suppose that for each $p \in \mathbb{M}$ at least one of the following conditions holds:*

- $\delta_p(\mathcal{Q}) < \frac{1}{2} \text{sys}(\mathbb{M})$,
- $\delta_p(\mathcal{Q}) < C \text{sys}_p(\mathbb{M})$, for $C = 0.447$.

Then $\delta_p(\mathcal{Q}') \leq \delta_p(\mathcal{Q})$ for all $p \in \mathbb{M}$ and $\mathcal{Q}' \supseteq \mathcal{Q}$ and $\pi(\text{DT}_{\mathbb{D}}(\pi^{-1}(\mathcal{Q})))$ does not contain 1- or 2-cycles.

Proof. By definition of $\delta_p(\mathcal{Q})$, the statement that $\delta_p(\mathcal{Q}') \leq \delta_p(\mathcal{Q})$ for all $p \in \mathbb{M}$ and $\mathcal{Q}' \supseteq \mathcal{Q}$ is immediate. To prove the second statement, we will first prove that $\ell(e) < \frac{1}{2} \text{sys}_q(\mathbb{M})$ for all edges $e = (u, v)$ in $\text{DT}_{\mathbb{D}}(\pi^{-1}(\mathcal{Q}))$, where q is the midpoint of e and $q = \pi(q)$.

Figure 3.4: Sketch of trigonometry used in the proof of Proposition 3.9.

Consider the circumcenter r of one of the triangles in $\text{DT}_{\mathbb{D}}(\Gamma\mathcal{Q})$ containing e . The triangles $[u, q, r]$ and $[v, q, r]$ are congruent and have a right angle at q (see Figure 3.4). The disk centered at r of radius $d(r, u)$ is the circumdisk of a triangle in $\text{DT}_{\mathbb{D}}(\pi^{-1}(\mathcal{Q}))$, so it does not contain any points of $\pi^{-1}(\mathcal{Q})$. Therefore, $d(r, u) = \frac{1}{2}\delta_r(\mathcal{Q})$, where $r = \pi(r)$. By the hyperbolic Pythagorean law

$$\cosh(\tfrac{1}{2}\ell(e)) = \cosh(d(u, q)) = \frac{\cosh(d(r, u))}{\cosh(d(q, r))} = \frac{\cosh(\tfrac{1}{2}\delta_r(\mathcal{Q}))}{\cosh(d(q, r))}.$$

By assumption one of the following holds:

- $\delta_r(\mathcal{Q}) < \frac{1}{2} \text{sys}(\mathbb{M})$,
- $\delta_r(\mathcal{Q}) < C \text{sys}_r(\mathbb{M})$.

In the first case,

$$\cosh(\tfrac{1}{2}\ell(e)) = \frac{\cosh(\tfrac{1}{2}\delta_r(\mathcal{Q}))}{\cosh(d(q, r))} \leq \cosh(\tfrac{1}{2}\delta_r(\mathcal{Q})) < \cosh(\tfrac{1}{4} \text{sys}(\mathbb{M})),$$

so $\ell(e) < \frac{1}{2} \text{sys}(\mathbb{M}) \leq \frac{1}{2} \text{sys}_q(\mathbb{M})$, which we wanted to prove. In the second case,

$$\text{sys}_r(\mathbb{M}) \leq \text{sys}_q(\mathbb{M}) + 2d(q, r),$$

since the curve consisting of following $[r, q]$, a local systole at q and $[q, r]$ is a closed curve passing through r . Therefore,

$$\cosh(\tfrac{1}{2}\ell(e)) = \frac{\cosh(\tfrac{1}{2}\delta_r(\mathcal{Q}))}{\cosh(d(q, r))} < \frac{\cosh(\tfrac{1}{2}C \text{sys}_r(\mathbb{M}))}{\cosh(d(q, r))} \leq \frac{\cosh(C(\tfrac{1}{2} \text{sys}_q(\mathbb{M}) + d(q, r)))}{\cosh(d(q, r))}.$$

To prove $\ell(e) < \frac{1}{2} \text{sys}_q(\mathbb{M})$ it is sufficient to show that

$$\frac{\cosh(C(\frac{1}{2} \text{sys}_q(\mathbb{M}) + d(\mathbf{q}, \mathbf{r})))}{\cosh(d(\mathbf{q}, \mathbf{r}))} \leq \cosh(\frac{1}{4} \text{sys}_q(\mathbb{M})). \quad (3.5)$$

Equation (3.5) is equivalent to

$$C \leq \frac{\text{arccosh}(\cosh(\frac{1}{4} \text{sys}_q(\mathbb{M})) \cosh(d(\mathbf{q}, \mathbf{r})))}{\frac{1}{2} \text{sys}_q(\mathbb{M}) + d(\mathbf{q}, \mathbf{r})}.$$

It can be verified with a computer algebra program that

$$\min_{\text{sys}_q(\mathbb{M}), d(\mathbf{q}, \mathbf{r}) > 0} \frac{\text{arccosh}(\cosh(\frac{1}{4} \text{sys}_q(\mathbb{M})) \cosh(d(\mathbf{q}, \mathbf{r})))}{\frac{1}{2} \text{sys}_q(\mathbb{M}) + d(\mathbf{q}, \mathbf{r})} \approx 0.447...$$

Because $C = 0.447$, Equation (3.5) is satisfied and we conclude that $\ell(e) < \frac{1}{2} \text{sys}_q(\mathbb{M})$ in the second case as well. Hence, $\ell(e) < \frac{1}{2} \text{sys}_q(\mathbb{M})$ for all edges e in $\text{DT}_{\mathbb{D}}(\pi^{-1}(\mathcal{Q}))$.

Now, suppose for a contradiction that there exists a 2-cycle in $\pi(\text{DT}_{\mathbb{D}}(\pi^{-1}(\mathcal{Q})))$ consisting of the edges e_1, e_2 . By the above reasoning, we know that $\ell(e_i) < \frac{1}{2} \text{sys}_{\mathbf{q}_i}(\mathbb{M})$, where $\mathbf{q}_i$ is the midpoint of e_i . Assume without loss of generality that $\text{sys}_{\mathbf{q}_1}(\mathbb{M}) \geq \text{sys}_{\mathbf{q}_2}(\mathbb{M})$. Then $\ell(e_1) + \ell(e_2) < \text{sys}_{\mathbf{q}_1}(\mathbb{M})$. Since 2-cycles correspond to homotopically non-trivial closed curves on $\mathbb{M}$, and since this is a curve through $\mathbf{q}_1$ with length less than $\text{sys}_{\mathbf{q}_1}(\mathbb{M})$, we obtain a contradiction. Therefore, there are no 2-cycles in $\pi(\text{DT}_{\mathbb{D}}(\pi^{-1}(\mathcal{Q})))$. By a similar reasoning, there are no 1-cycles as well. This finishes the proof. $\square$

Because the local validity condition is weaker than the (global) validity condition (3.1), the number of dummy points necessary to satisfy the local validity condition is in general smaller. Since the density of a point set satisfying the global validity condition depends only on the (global) systole, the local validity condition seems to be particularly useful for surfaces with a small systole.

This can also be seen from the following example. Consider a continuous deformation of a hyperbolic surface that is defined by monotonically decreasing the length of a fixed closed geodesic on the surface. As the length of this fixed closed geodesic is decreasing, a point set satisfying the global validity condition will need an increasing number of points and these extra points will need to be inserted across the entire surface. On the other hand, a point set satisfying the local validity condition only needs more points near the fixed closed geodesic with decreasing length, since the remainder of the surface is unchanged.

Unfortunately, so far we were not able to prove the existence of a point set satisfying the local validity condition with a cardinality that is (asymptotically) strictly smaller than the bound given in Proposition 3.4. Simulations suggest that it may be sufficient to place $O(\text{sys}(\mathbb{M})^{-1})$ points near any short closed geodesic

(of which there are $O(g)$) and $O(g)$ points in the remainder of the surface. This would lead to a point set of cardinality $g \cdot O(\text{sys}(\mathbb{M})^{-1})$ instead of $g \cdot O(\text{sys}(\mathbb{M})^{-2})$ for point sets satisfying the global validity condition, but we emphasize that at this stage it is no more than a conjecture.

Moreover, an important practical drawback of the local validity condition is that in general we do not know the local systole at every point on a given hyperbolic surface (or, at a sufficiently dense set of points on the surface). Therefore, constructing a dummy point set that satisfies the local validity condition is generally harder than constructing a dummy point set that satisfies the (global) validity condition. For this reason, and because we cannot precisely quantify the benefit of the local validity condition over the global validity condition, we will not consider the local validity condition when computing dummy point sets for the generalized Bolza surfaces in Section 3.5.

3.5 Computation of dummy point sets

In this section we present three algorithms for constructing a dummy point set $\mathcal{Q}_g$ satisfying the validity condition (3.1) for $\mathbb{M}_g$ and give the growth rate of the cardinality of $\mathcal{Q}_g$ as a function of g .

All algorithms use the set $\mathcal{W}_g$ of the so-called Weierstrass points of $\mathbb{M}_g$. In the fundamental domain D_g , the Weierstrass points are represented by the origin, the vertices and the midpoints of the sides. In the original domain $\tilde{D}_g$, where there is only one point of each orbit under the action of Γ_g , this reduces to $2g + 2$ points: the origin, the midpoint of each of the $2g$ closed sides, and the vertex v_0 . Some special properties of Weierstrass points are known in Riemann surface theory [34], however we will not use them in this paper.

Each of the algorithms has its own advantages and drawbacks. The *refinement algorithm* (Section 3.5.1) yields a point set with optimal asymptotic cardinality $\Theta(g)$ (Proposition 3.5). The idea is borrowed from the well-known Delaunay refinement algorithm for mesh generation [66]. The *symmetric algorithm* (Section 3.5.2) uses the Delaunay refinement algorithm as well. However, instead of inserting one point in each iteration, we insert its images by all rotations around the origin by angle $k\pi/2g$ for $k = 1, \dots, 4g$. In this way, we obtain a dummy point set that preserves the symmetries of D_g , at the cost of increasing the asymptotic cardinality to $\Theta(g \log g)$. The approach of the *structured algorithm* (Section 3.5.3) is fundamentally different from the refinement and symmetric algorithms: the dummy point set and the corresponding Delaunay triangulation are exactly described. As in the symmetric algorithm, the resulting dummy point set preserves the symmetries of D_g and is of order $\Theta(g \log g)$.

Let us now elaborate on the refinement algorithm. The set $\mathcal{Q}_g$ is initialized as $\mathcal{W}_g$ and the triangulation as $\text{DT}_{\mathbb{M}_g}(\mathcal{W}_g)$. Then, all non-admissible triangles in $\text{DT}_{\mathbb{D}}(\pi_g^{-1}(\mathcal{Q}_g))$ are removed by inserting the projection onto $\mathbb{M}_g$ of their circumcenter, while updating the set $\mathcal{Q}_g$ of vertices of the triangulation. The following

proposition shows that $\text{DT}_{\mathbb{D}}(\pi_g^{-1}(\mathcal{Q}_g) \cap D_{\mathcal{N}_g})$ contains at least one representative of each face of $\text{DT}_{\mathbb{D}}(\pi_g^{-1}(\mathcal{Q}_g))$, thus providing the refinement algorithm with a finite input.

Proposition 3.10. *For any finite set of points $\mathcal{Q}_g$ on $\mathbb{M}_g$ containing $\mathcal{W}_g$, each face in $\text{DT}_{\mathbb{D}}(\pi_g^{-1}(\mathcal{Q}_g))$ with at least one vertex in $\tilde{D}_g$ is contained in $D_{\mathcal{N}_g}$.*

The proof will use the following lemma.

Lemma 3.11. *Let C_q be a Euclidean disk centered at O and passing through a point q (Figure 3.5). Let H_1 and H_2 denote the two open half-planes bounded by the Euclidean line through O and q . Let H_0 be a half-plane that contains q , bounded by another Euclidean line passing through O but not through q . Let $R_j = (H_0 \cap H_j) \setminus C_q, j = 1, 2$, and let $p_j \in R_j, i = 1, 2$. The disk $C(p_1, p_2)$ through O, p_1 , and p_2 contains q .*

Figure 3.5: Illustrations for the proof of Lemma 3.11.

Proof. It is easy to verify that there exist pairs of points $(p_1, p_2) \in R_1 \times R_2$ for which the point q lies inside the disk $C(p_1, p_2)$. For instance, consider a line perpendicular to the line through O and q so that q is closer to O than to their intersection point, as shown in Figure 3.5 - Left. If p_1 lies on this perpendicular line and p_2 is the reflection of p_1 in the line through O and q , then the disk $C(p_1, p_2)$ contains q . Since this disk varies continuously when (p_1, p_2) ranges over $R_1 \times R_2$, it is sufficient to prove that there are no pairs $(p_1^*, p_2^*) \in R_1 \times R_2$ for which q lies on the boundary of $C(p_1^*, p_2^*)$.

Suppose, for a contradiction, that there exists a pair $(p_1^*, p_2^*) \in R_1 \times R_2$ for which $C(p_1^*, p_2^*)$ is a disk with q and O on its boundary. Consider the disk C_q centered at O and passing through q . Let F be the intersection of the disk with

diameter $[O, q]$ with the half-plane H_1 , as shown in Figure 3.5 - Right. For any point $p \in H_2 \setminus C_q$, the circle C_p through O, q , and p has a non-empty intersection with H_1 , which is completely included in F , so in particular $C(p_1^*, p_2^*)$ intersects H_1 inside the disk with diameter $[O, q]$. By a symmetric observation, $C(p_1^*, p_2^*)$ also intersects H_2 inside the same disk. Therefore, $C(p_1^*, p_2^*)$ is the disk with diameter $[O, q]$. This implies that both p_1 and p_2 lie in the disk C_q , which is a contradiction. Therefore, there exists no pair $(p_1^*, p_2^*) \in R_1 \times R_2$ for which $C(p_1^*, p_2^*)$ has q on its boundary. This finishes the proof. $\square$

Note that Lemma 3.11 can be directly used in the Poincaré disk because hyperbolic circles are represented as Euclidean circles, and hyperbolic geodesics through the origin O are supported by Euclidean lines.

We now proceed with the proof of Proposition 3.10.

Figure 3.6: Illustration for the proof of Proposition 3.10 for $g = 2$.

Proof. We show that each edge in $\text{DT}_{\mathbb{D}}(\pi_g^{-1}(\mathcal{Q}_g))$ with one endpoint in $\tilde{D}_g$ has its other endpoint inside $D_{\mathcal{N}_g}$.

Let e be a segment with an endpoint in $\tilde{D}_g$ and an endpoint outside $D_{\mathcal{N}_g}$. We will prove that every disk passing through the endpoints of e contains a point in $\mathcal{W}_g$. There are two cases to consider: e either crosses only one image of $\tilde{D}_g$ under an element of $\mathcal{N}_g \setminus \{1\}$, or it crosses several of its images. We examine each case separately.

Case A: The edge e crosses only one image of $\tilde{D}_g$ before leaving $D_{\mathcal{N}_g}$. This case is illustrated in Figure 3.6 - Left. Let $f(\tilde{D}_g), f \in \mathcal{N}_g \setminus \{1\}$ be the Dirichlet region that e crosses. The image $e' = f^{-1}(e)$ of e then crosses $\tilde{D}_g$, intersecting two of its non-adjacent sides s_j and s_k in the points p_j and p_k , respectively.

We can assume without loss of generality that the hyperbolic segment $[\mathbf{p}_j, \mathbf{p}_k]$ does not contain the origin, since in that case any disk through $\mathbf{p}_j$ and $\mathbf{p}_k$ clearly contains the origin. Then, there exists a line l_O through O such that $\mathbf{p}_j$ and $\mathbf{p}_k$ are contained in the same half-space. Let $\mathbf{m}_\ell$ be the midpoint of a side between $\mathbf{s}_j$ and $\mathbf{s}_k$ in the same half-space as $\mathbf{p}_j$ and $\mathbf{p}_k$. Consider the disk centered at O that passes through $\mathbf{m}_\ell$ (and, of course, through all the other midpoints $\mathbf{m}_k$ as well), and consider also the line through O and $\mathbf{m}_\ell$. By Lemma 3.11, the disk $C(\mathbf{p}_j, \mathbf{p}_k)$ passing through $O, \mathbf{p}_j$, and $\mathbf{p}_k$ contains $\mathbf{m}_\ell$. Since O and $\mathbf{m}_\ell$ are on both sides of the segment $[\mathbf{p}_j, \mathbf{p}_k]$, any disk through $\mathbf{p}_j$ and $\mathbf{p}_k$ contains either $\mathbf{m}_\ell$ or O , therefore there is no empty disk that passes through $\mathbf{p}_j$ and $\mathbf{p}_k$. Assume now that there is an empty disk that passes through the endpoints of $\mathbf{e}'$. This empty disk can then be shrunk continuously so that it passes through $\mathbf{p}_j$ and $\mathbf{p}_k$. The shrunk version of the disk must be also empty, which is a contradiction. Therefore, there is no empty disk passing through the endpoints of $\mathbf{e}'$, which implies that $\mathbf{e}'$ (and, by consequence, $\mathbf{e}$) cannot be an edge in $\text{DT}_{\mathbb{D}}(\pi_g^{-1}(\mathcal{Q}_g))$.

Case B: The edge $\mathbf{e}$ crosses several images of $\tilde{D}_g$ before leaving $D_{\mathcal{N}_g}$. This case is illustrated in Figure 3.6 - Right. There exist multiple images of $\mathbf{e}$ in $\text{DT}_{\mathbb{D}}(\pi_g^{-1}(\mathcal{Q}))$ that intersect $\tilde{D}_g$, in fact as many as the number of Dirichlet regions it intersects. Each one of these images intersects two adjacent sides of $\tilde{D}_g$. Let $\mathbf{e}'$ be an image of $\mathbf{e}$ that intersects two adjacent sides $\mathbf{s}_j$ and $\mathbf{s}_{j+1}$ of $\tilde{D}_g$ so that the hyperbolic line supporting $\mathbf{e}'$ separates O and the midpoint $\mathbf{m}_{j+1}$. Note that such an image of $\mathbf{e}$ exists always: $\mathbf{e}$ either separates O and the midpoint $\mathbf{m}_{j+1}$, or it separates an image of O under some translation f of Γ_g and $\mathbf{m}_{j+1}$; in the second case, $f^{-1}(\mathbf{e})$ separates O and the midpoint $f^{-1}(\mathbf{m}_{j+1})$. The edge $\mathbf{e}'$ intersects also the side $\mathbf{s}_k$ adjacent to $\mathbf{s}_{j+1}$ in the Dirichlet region that shares the side $\mathbf{s}_{j+1}$ with $\tilde{D}_g$ (see Figure 3.6 - Right). Let $\mathbf{p}_j$ and $\mathbf{p}_k$ be the intersection points of $\mathbf{e}'$ with $\mathbf{s}_j$ and $\mathbf{s}_k$, respectively. Consider the circle centered at the origin that passes through $\mathbf{m}_{j+1}$. Consider also the line through O and $\mathbf{m}_{j+1}$ and the line l_0 through O perpendicular to it. By Lemma 3.11, the disk $C(\mathbf{p}_j, \mathbf{p}_k)$ passing through $O, \mathbf{p}_j$, and $\mathbf{p}_k$ contains $\mathbf{m}_{j+1}$. Since O and $\mathbf{m}_{j+1}$ are on both sides of the segment $[\mathbf{p}_j, \mathbf{p}_k]$, any disk through $\mathbf{p}_j$ and $\mathbf{p}_k$ contains either $\mathbf{m}_{j+1}$ or O , therefore there is no empty disk that passes through $\mathbf{p}_j$ and $\mathbf{p}_k$. By the same reasoning as in Case A, there is no empty disk passing through the endpoints of $\mathbf{e}'$ either, which implies that $\mathbf{e}'$ (and, by consequence, $\mathbf{e}$) cannot be an edge in $\text{DT}_{\mathbb{D}}(\pi_g^{-1}(\mathcal{Q}_g))$.

In conclusion, no edge of $\text{DT}_{\mathbb{D}}(\pi_g^{-1}(\mathcal{Q}_g))$ can have an endpoint in $\tilde{D}_g$ and an endpoint outside $D_{\mathcal{N}_g}$, therefore all faces with at least one vertex in $\tilde{D}_g$ are included in $D_{\mathcal{N}_g}$. $\square$

The set $\pi_g^{-1}(\mathcal{Q}_g) \cap D_{\mathcal{N}_g}$ is obtained as follows: we first consider the set of canonical representatives (as defined in Section 1.1.7) of the points of $\mathcal{Q}_g$, which is $\pi_g^{-1}(\mathcal{Q}_g) \cap \tilde{D}_g$. Then, we obtain $\pi_g^{-1}(\mathcal{Q}_g) \cap D_{\mathcal{N}_g}$ by computing the images of

$\pi_g^{-1}(\mathcal{Q}_g) \cap \tilde{D}_g$ under the elements in $\mathcal{N}_g$. In other words, $\pi_g^{-1}(\mathcal{Q}_g) \cap D_{\mathcal{N}_g}$ can be computed as $\mathcal{Q}_{\mathcal{N}_g} = \{f(\pi_g^{-1}(\mathcal{Q}_g) \cap \tilde{D}_g), f \in \mathcal{N}_g\}$.

3.5.1 Refinement algorithm

Following the refinement strategy introduced above and using Proposition 3.10, we insert the circumcenter of each triangle in $\text{DT}_{\mathbb{D}}(\mathcal{Q}_{\mathcal{N}_g})$ having a non-empty intersection with the domain $\tilde{D}_g$ and whose circumradius is at least $\frac{1}{2} \text{sys}(\mathbb{M}_g)$ (see Algorithm 1). Figure 3.7 illustrates the computation of $\text{DT}_{\mathbb{D}}(\mathcal{Q}_{\mathcal{N}_3})$.

Input : hyperbolic surface $\mathbb{M}_g$
Output: finite point set $\mathcal{Q}_g \subset \mathbb{M}_g$ such that $\delta(\mathcal{Q}_g) < \frac{1}{2} \text{sys}(\mathbb{M}_g)$

- 1 Initialize: let $\mathcal{Q}_g$ be the set $\mathcal{W}_g$ of Weierstrass points of $\mathbb{M}_g$.
- 2 Compute $\text{DT}_{\mathbb{D}}(\mathcal{Q}_{\mathcal{N}_g})$.
- 3 **while** there exists a triangle Δ in $\text{DT}_{\mathbb{D}}(\mathcal{Q}_{\mathcal{N}_g})$ with circumdiameter at least $\frac{1}{2} \text{sys}(\mathbb{M}_g)$ and $\Delta \cap D_g \neq \emptyset$ **do**
- 4 | Add to $\mathcal{Q}_g$ the projection onto $\mathbb{M}_g$ of the circumcenter of Δ
- 5 | Update $\text{DT}_{\mathbb{D}}(\mathcal{Q}_{\mathcal{N}_g})$
- 6 **end**

Algorithm 1: Refinement algorithm

We can now show that the cardinality of the resulting dummy point set is linear in the genus g .

Theorem 3.12. *The refinement algorithm terminates and the resulting dummy point set $\mathcal{Q}_g$ satisfies the validity condition (3.1). The cardinality $|\mathcal{Q}_g|$ is bounded as follows*

$$5.699(g-1) < |\mathcal{Q}_g| < 27.061(g-1).$$

Proof. We first prove that the hyperbolic distance between two distinct points of $\mathcal{Q}_g$ is greater than $\frac{1}{4} \text{sys}(\mathbb{M}_g)$. It is known that the distance between any pair of Weierstrass points is larger than $\frac{1}{2} \text{sys}(\mathbb{M}_g)$ [9]. Furthermore, every point added after the initialization is the projection of the circumcenter of an empty disk in $\mathbb{D}$ of radius at least $\frac{1}{4} \text{sys}(\mathbb{M}_g)$, so the distance from the added point to any other point in $\mathcal{Q}_g$ is at least $\frac{1}{4} \text{sys}(\mathbb{M}_g)$. For arbitrary $p \in \mathcal{Q}_g$, consider the disk D_p in $\mathbb{M}_g$ of radius $\frac{1}{8} \text{sys}(\mathbb{M}_g)$ centered at p , i.e., the set of points in $\mathbb{M}_g$ at distance at most $\frac{1}{8} \text{sys}(\mathbb{M}_g)$ from p . Every disk of radius at most $\frac{1}{2} \text{sys}(\mathbb{M}_g)$ is embedded in $\mathbb{M}_g$, so in particular D_p is an embedded disk. Because the distance between any pair of points of $\mathcal{Q}_g$ is at least $\frac{1}{4} \text{sys}(\mathbb{M}_g)$, the disks D_p and D_q of radius $\frac{1}{8} \text{sys}(\mathbb{M}_g)$ centered at p and q , respectively, are disjoint for every distinct $p, q \in \mathcal{Q}_g$. For fixed g , the area of such disks is fixed, as is the area of $\mathbb{M}_g$, so only a finite number of points can be added. Hence, the algorithm terminates.

Observe that the algorithm terminates if and only if the while loop ends, i.e. $\mathcal{Q}_g$ satisfies the validity condition.

After initialization

First insertion

After first insertion

After last insertion

Figure 3.7: Several steps in the refinement algorithm (genus 3)

Finally, we bound for the cardinality of $\mathcal{Q}_g$. From the above argument we know that the cardinality of $\mathcal{Q}_g$ is bounded above by the number of disjoint disks D of radius $\frac{1}{8} \text{sys}(\mathbb{M}_g)$ that fit inside $\mathbb{M}_g$. Hence,

$$|\mathcal{Q}_g| \leq \frac{\text{area}(\mathbb{M}_g)}{\text{area}(D)} = \frac{4\pi(g-1)}{2\pi \left(\cosh\left(\frac{1}{8} \text{sys}(\mathbb{M}_g)\right) - 1 \right)} = \frac{2(g-1)}{\cosh\left(\frac{1}{8} \text{sys}(\mathbb{M}_g)\right) - 1}.$$

Proposition 3.5 gives a lower bound. The coefficients of $g-1$ in these upper and lower bounds decrease as a function of g , so the announced bounds can be obtained by plugging in the value of $\text{sys}(\mathbb{M}_g)$ (see Theorem 2.1) for $g \rightarrow \infty$ and $g = 2$ respectively. This finishes the proof. $\square$

3.5.2 Symmetric algorithm

This algorithm is similar to the refinement algorithm. However, instead of adding one point at every step in the while loop, it uses the $4g$ -fold symmetry of the fundamental polygon D_g to add $4g$ points at every step (see Algorithm 2). Figure 3.8 illustrates the computation of $\text{DT}_{\mathbb{D}}(\mathcal{Q}_{\mathcal{N}_3})$.

<p>Input : hyperbolic surface $\mathbb{M}_g$</p> <p>Output: finite point set $\mathcal{Q}_g \subset \mathbb{M}_g$ such that $\delta(\mathcal{Q}_g) < \frac{1}{2} \text{sys}(\mathbb{M}_g)$</p> <ol style="list-style-type: none"> 1 Initialize: let $\mathcal{Q}_g$ be the set $\mathcal{W}_g$ of Weierstrass points of $\mathbb{M}_g$. 2 Compute $\text{DT}_{\mathbb{D}}(\mathcal{Q}_{\mathcal{N}_g})$. 3 while there exists a triangle Δ in $\text{DT}_{\mathbb{D}}(\mathcal{Q}_{\mathcal{N}_g})$ with circumdiameter at least $\frac{1}{2} \text{sys}(\mathbb{M}_g)$ do <li style="padding-left: 20px;">4 for $k = 0, \dots, 4g - 1$ do <li style="padding-left: 40px;">5 Let $\mathbf{p}_k$ be the circumcenter of Δ rotated around the origin by angle $\frac{k\pi}{2g}$. <li style="padding-left: 40px;">6 Add $\pi_g(\mathbf{p}_k)$ to $\mathcal{Q}_g$. <li style="padding-left: 20px;">7 end <li style="padding-left: 20px;">8 Update $\text{DT}_{\mathbb{D}}(\mathcal{Q}_{\mathcal{N}_g})$. 9 end

Algorithm 2: Symmetric algorithm

By using the symmetry of the regular $4g$ -gon we obtain a more symmetric dummy point set, which may be interesting for some applications [25]. However, asymptotically the resulting point set is larger than the point set obtained from the refinement algorithm.

Theorem 3.13. *The symmetric algorithm terminates and the resulting dummy point set satisfies the validity condition (3.1). Its cardinality is of order $\Theta(g \log g)$.*

Proof. The first two statements follow directly from the proof of Theorem 3.12, so we only have to prove the claim on the cardinality of $\mathcal{Q}_g$.

First, we prove that $|\mathcal{Q}_g|$ is of order $O(g \log g)$. Again, the distance between the Weierstrass points is more than $\frac{1}{4} \text{sys}(\mathbb{M}_g)$. We claim that the distance between points that are added in different iterations of the while loop is at least $\frac{1}{4} \text{sys}(\mathbb{M}_g)$. Namely, by the same reasoning as in the proof of Theorem 3.12, the distance between the circumcenter of an empty disk of radius at least $\frac{1}{4} \text{sys}(\mathbb{M}_g)$ and any other point in $\mathcal{Q}_g$ is at least $\frac{1}{4} \text{sys}(\mathbb{M}_g)$. Because $\mathcal{Q}_g$ is invariant under symmetries of D_g , it follows that the distance between an image of the circumcenter under a rotation around the origin and any other point in $\mathcal{Q}_g$ is at least $\frac{1}{4} \text{sys}(\mathbb{M}_g)$ as well.

However, the distance between points in $\mathcal{Q}_g$ can be smaller than $\frac{1}{4} \text{sys}(\mathbb{M}_g)$ if they are added simultaneously in some iteration of the while loop. Denote the points added to $\mathcal{Q}_g$ in iteration j by $\pi_g(\mathbf{p}_k^j)$ where $k = 0, \dots, 4g - 1$. In particular, $\mathbf{p}_k^j$ is the circumcenter of a triangle in $\text{DT}_{\mathbb{D}}(\mathcal{Q}_{\mathcal{N}_g})$, i.e., in the hyperbolic plane.

Figure 3.8: Several steps in the symmetric algorithm (genus 3)

Let $D(p, r)$ be the hyperbolic disk with center p and radius r , where p is either a point in $\mathbb{H}^2$ or in $\mathbb{M}_g$. For each iteration j , define

$$U_j = \bigcup_{k=0}^{4g-1} D\left(p_k^j, \frac{1}{8} \text{sys}(\mathbb{M}_g)\right)$$

and let $U_j = \pi_g(U_j)$. Let $a_j = \text{area}(U_j)$. Denote the area of a hyperbolic circle of radius $\frac{1}{8} \text{sys}(\mathbb{M}_g)$ by a , i.e.

$$a := 2\pi \left(\cosh\left(\frac{1}{8} \text{sys}(\mathbb{M}_g)\right) - 1 \right).$$

Observe that $a \leq a_j \leq 4ga$, where the lower bound is in the limiting case where all disks are equal and the upper bound in the case where all disks are disjoint.

Define

$$\mathcal{I} = \{j \mid a_j < 2ga\}$$

and denote its complement by $\mathcal{I}^c$. We give upper bounds for $|\mathcal{I}|$ and $|\mathcal{I}^c|$. To see for which j the inequality $a_j < 2ga$ holds, we first look at the area of U_j (see Figure 3.9a). The amount of overlap between $D(\mathbf{p}_k^j)$ and $D(\mathbf{p}_{k+1}^j)$ can be written as a strictly decreasing function of $d(\mathbf{p}_k^j, \mathbf{p}_{k+1}^j)$, which can be written as a strictly increasing function of $d(O, \mathbf{p}_k^j)$. Therefore, there exists a constant $d_g > 0$ such that $\text{area}(U_j) < 2ga$ if and only if $d(O, \mathbf{p}_k^j) < d_g$ for all $k = 0, \dots, 4g - 1$.

We claim that $j \in \mathcal{I}$ if and only if there exists $k \in \{0, \dots, 4g - 1\}$ such that either $d(O, \mathbf{p}_k^j) < d_g$ or $d(\mathbf{v}_0, \mathbf{p}_k^j) < d_g$. First, assume that such a k exists. If $d(O, \mathbf{p}_k^j) < d_g$ (Figure 3.9a), then $\text{area}(U_j) < 2ga$ by definition of d_g , so $j \in \mathcal{I}$. Now, assume that $d(\mathbf{v}_0, \mathbf{p}_k^j) < d_g$ (Figure 3.9b). By symmetry $d(\mathbf{v}_\ell, \mathbf{p}_{k+\ell}^j) = d(\mathbf{v}_0, \mathbf{p}_k^j)$ for all $\ell = 0, \dots, 4g - 1$ (counting modulo $4g$). Recall that f_0 is the side-pairing transformation that maps s_{2g} to s_0 . Then

$$\begin{aligned} d(f_0(\mathbf{p}_{k+2g+1}^j), \mathbf{v}_0) &= d(f_0^{-1}(f_0(\mathbf{p}_{k+2g+1}^j)), f_0^{-1}(\mathbf{v}_0)), \\ &= d(\mathbf{p}_{k+2g+1}^j, \mathbf{v}_{2g+1}), \\ &= d(\mathbf{p}_k^j, \mathbf{v}_0). \end{aligned}$$

Therefore, the circle C_j centered at $\mathbf{v}_0$ and passing through $\mathbf{p}_k^j$ passes through $f_0(\mathbf{p}_{k+2g+1}^j)$ as well. By induction, for every pair of adjacent fundamental domains $f(D_g)$ and $f'(D_g)$ that contain $\mathbf{v}_0$ there exists an $\ell \in \{0, \dots, 4g - 1\}$ such that $f(\mathbf{p}_\ell^j)$ and $f'(\mathbf{p}_{\ell+2g+1}^j)$ are equidistant from $\mathbf{v}_0$. There are $4g$ fundamental domains that have $\mathbf{v}_0$ as one of their vertices. Because $2g + 1$ and $4g$ are co-prime, it follows that C_j contains exactly one translate of $\mathbf{p}_\ell^j$ for every $\ell = 0, \dots, 4g - 1$. Hence, if we translate the union of disks of radius $\frac{1}{8} \text{sys}(\mathbb{M}_g)$ centered at the translates of $\mathbf{p}_\ell^j, \ell = 0, \dots, 4g - 1$ on C_j by the hyperbolic translation that maps $\mathbf{v}_0$ to the origin, we obtain a union of disks of radius $\frac{1}{8} \text{sys}(\mathbb{M}_g)$ at distance $d(\mathbf{v}_0, \mathbf{p}_k^j) < d_g$ from the origin. By definition of d_g , it follows that $a_j < 2ga$.

Second, assume that $d(O, \mathbf{p}_k^j) \geq d_g$ and $d(\mathbf{v}_0, \mathbf{p}_k^j) \geq d_g$ for all $k \in \{0, \dots, 4g - 1\}$. If $d(\mathbf{p}_0^j, \partial D_g) \geq \frac{1}{8} \text{sys}(\mathbb{M}_g)$, then U_j is completely contained in D_g . Because $d(O, \mathbf{p}_k^j) \geq d_g$, it follows that $a_j \geq 2ga$ by definition of d_g , so $j \in \mathcal{I}^c$. Now, assume that $d(\mathbf{p}_0^j, \partial D_g) < \frac{1}{8} \text{sys}(\mathbb{M}_g)$. If $\mathbf{p}_0^j$ is close to the midpoint of a side of D_g , then $D(\mathbf{p}_0^j, \frac{1}{8} \text{sys}(\mathbb{M}_g))$ can only overlap with a translate of $D(\mathbf{p}_{2g}^j, \frac{1}{8} \text{sys}(\mathbb{M}_g))$ (Figure 3.9c). Then, U_j contains at least $2g$ pairwise disjoint disks, so $a_j \geq 2ga$. Therefore, $j \in \mathcal{I}^c$. Hence, the only way that $D(\mathbf{p}_0^j, \frac{1}{8} \text{sys}(\mathbb{M}_g))$ can overlap with multiple other disks is when $\mathbf{p}_0^j$ is sufficiently close to a vertex of D_g . Consider again the circle C_j centered at $\mathbf{v}_0$ and passing through a translate of $\mathbf{v}_\ell^j$ for all $\ell \in \{0, \dots, 4g - 1\}$. Because now $d(\mathbf{v}_0, \mathbf{p}_k^j) \geq d_g$, it follows that $a_j \geq 2ga$ by definition of d_g .

(a) $d(O, p_k^j) < d_g$ for all $k = 0, \dots, 4g - 1$. (b) $d(v_0, p_k^j) < d_g$ for some $k \in \{0, \dots, 4g - 1\}$.

(c) $d(O, p_k^j) \geq d_g$ and $d(v_0, p_k^j) \geq d_g$ for all $k \in \{0, \dots, 4g - 1\}$ and $d(p_0^j, \partial D_g) < \frac{1}{8} \text{sys}(\mathbb{M}_g)$.

Figure 3.9: Schematic drawings of different cases. In the first two drawings, the minimum width of the corresponding annulus is marked in red. In the second drawing, only the disks with center in D_g or sufficiently close to v_0 are drawn. In the third drawing, only the disks with center in D_g together with the unique disk that overlaps $D(p_0^j, \frac{1}{8} \text{sys}(\mathbb{M}_g))$ are drawn.

We conclude that $j \in \mathcal{I}$ if and only if there exists $k \in \{0, \dots, 4g-1\}$ such that either $d(O, \mathbf{p}_k^j) < d_g$ or $d(\mathbf{v}_0, \mathbf{p}_k^j) < d_g$. We have also shown that if $d(O, \mathbf{p}_k^j) < d_g$, then U_j is a topological annulus around the origin. If $d(\mathbf{v}_0, \mathbf{p}_k^j) < d_g$, then $\pi_g^{-1}(U_j)$ contains a topological annulus around $\mathbf{v}_0$. In either case, the boundary of such an annulus consists of two connected components. Let the minimum width of an annulus be given by the distance between these connected components. Suppose, for a contradiction, that the minimum width of an annulus corresponding to $j \in \mathcal{I}$ can be arbitrarily close to 0. Then the disks in U_j have arbitrarily small overlap, so a_j is arbitrarily close to $4ga$. However, this is not possible, since $a_j < 2ga$ for all $j \in \mathcal{I}$. Therefore, there exists $\varepsilon > 0$ (independent of the output of the algorithm) such that the minimum width of an annulus corresponding to $j \in \mathcal{I}$ is at least ε .

To find an upper bound for $|\mathcal{I}|$, consider the line segment $[O, \mathbf{v}_0]$ between the origin and $\mathbf{v}_0$. By the above discussion, $[O, \mathbf{v}_0]$ crosses the annulus corresponding to any $j \in \mathcal{I}$ exactly once. Because the annuli are pairwise disjoint and each annulus has minimum width ε , there are at most $\text{length}([O, \mathbf{v}_0])/\varepsilon$ annuli, where

$$\text{length}([O, \mathbf{v}]) = \text{arccosh} \left(\cot^2 \left(\frac{\pi}{4g} \right) \right).$$

Therefore,

$$|\mathcal{I}| \leq \frac{\text{arccosh} \left(\cot^2 \left(\frac{\pi}{4g} \right) \right)}{\varepsilon}.$$

Because $\cot^2(\frac{\pi}{4g}) \sim \frac{16}{\pi^2}g^2$ for $g \rightarrow \infty$, it follows that $|\mathcal{I}|$ is of order $O(\log g)$.

Now, consider $\mathcal{I}^c$. Because the disks of radius $\frac{1}{8} \text{sys}(\mathbb{M}_g)$ centered at points of $\mathcal{Q}_g$ that correspond to different iterations of the while loop are disjoint, we see that

$$\text{area}(\mathbb{M}_g) \geq \text{area}(\cup_{j \in \mathcal{I}^c} U_j) = \sum_{j \in \mathcal{I}^c} \text{area}(U_j) = \sum_{j \in \mathcal{I}^c} a_j \geq |\mathcal{I}^c| \cdot 2ga.$$

Since $\text{area}(\mathbb{M}_g) = 4\pi(g-1)$ and a is constant, $|\mathcal{I}^c|$ is of order $O(1)$.

Because the number of iterations is given by $|\mathcal{I}| + |\mathcal{I}^c|$, the number of iterations is of order $O(\log g)$. Each iteration adds $4g$ points, so the resulting dummy point set has cardinality of order $O(g \log g)$.

Secondly, we show that $|\mathcal{Q}_g|$ is of order $\Omega(g \log g)$. As before, the points added to $\mathcal{Q}_g$ in iteration j of the while loop are denoted by $\pi_g(\mathbf{p}_k^j)$ where $k = 0, \dots, 4g-1$. Fix an arbitrary vertex $\mathbf{v}$ of D_g . Let $P = \langle O, \mathbf{p}_{k_1}^{j_1}, \mathbf{p}_{k_2}^{j_2}, \dots, \mathbf{p}_{k_n}^{j_n}, \mathbf{v} \rangle$ be a shortest path from the origin to $\mathbf{v}$ in the Delaunay graph of $\pi_g^{-1}(\mathcal{Q}_g)$. We claim that all indices j_h are distinct, i.e. P contains at most one element of each of the sets $\{\mathbf{p}_k^j \mid k = 0, \dots, 4g-1\}$ (see Figure 3.10).

Suppose, for a contradiction, that there exist l and m with $l < m$, such that $j_l = j_m$. We will construct a path P' from O to $\mathbf{v}$ that is shorter than P . We know that $\mathbf{p}_{k_l}^{j_l} \neq \mathbf{p}_{k_m}^{j_m}$, because otherwise the shortest path would contain a cycle,

Figure 3.10: The left figure shows a path P from the origin to v that visits two vertices from the same iteration, namely p_3^2 and p_4^2 . The right figure shows a shorter path from the origin to v . In this case, $j_1 = j_2 = 2$ and $k_1 = 4$ and $k_2 = 3$. The subdivision of P into three parts is given by $P_1 = \langle O, p_4^2 \rangle$, $P_2 = \langle p_4^2, p_3^2 \rangle$ and $P_3 = \langle p_3^2, p_1^2, p_3^2, v \rangle$. The path P' is defined as $P'_1 \cup P_3$, where P'_1 is obtained by rotating P_1 around the origin by angle $-\frac{\pi}{2g}$, i.e., $P'_1 = \langle O, p_3^2 \rangle$.

so in particular $k_l \neq k_m$. Subdivide P into three paths: the path P_1 from O to $p_{k_l}^{j_l}$, the path P_2 from $p_{k_l}^{j_l}$ to $p_{k_m}^{j_m}$, and the path P_3 from $p_{k_m}^{j_m}$ to v . Now, let P'_1 be the image of P_1 after rotation around O by angle $(k_m - k_l) \cdot \frac{\pi}{2g}$. It is clear that P'_1 is a path from O to $p_{k_m}^{j_m}$ of the same length of P_1 . It follows that $P' := P'_1 \cup P_3$ is a path from O to v that is shorter than P . This is a contradiction, so all indices j_h are distinct. Therefore, the number of vertices of the graph that P visits is smaller than the number of iterations of the while loop. Each edge in the path P is the side of a triangle with circumdiameter smaller than $\frac{1}{2} \text{sys}(\mathbb{M}_g)$, so in particular the length of each edge is smaller than $\frac{1}{2} \text{sys}(\mathbb{M}_g)$. The length of P is at least

$$\text{length}([O, v]) = \text{arccosh} \left(\cot^2 \left(\frac{\pi}{4g} \right) \right) \sim 2 \log g.$$

As $\frac{1}{2} \text{sys}(\mathbb{M}_g)$ is bounded as a function of g (Theorem 2.1), the number of edges in P is of order $\Omega(\log g)$. Then, the number of iterations of the while loop is of order $\Omega(\log g)$, so $|\mathcal{Q}_g|$ has cardinality of order $\Omega(g \log g)$. The result follows by combining the lower and upper bounds. $\square$

3.5.3 Structured algorithm

Like the symmetric algorithm, this algorithm respects the $4g$ -fold symmetry of the Dirichlet region of $\mathbb{M}_g$. Before we give the algorithm in pseudocode, we first explain the idea and the notation. See Figure 3.11 for an illustration of the dummy points within one slice of the $4g$ -gon.

Figure 3.11: Dummy points within one slice of the $4g$ -gon.

1. As in the other algorithms, initially $\mathcal{Q}_g$ consists of the Weierstrass points of $\mathbb{M}_g$: the origin, the vertex and the midpoints of the sides. As usual, the origin and the vertices of D_g are denoted by O and v_k respectively. The midpoint of side s_k is denoted by p_k^0 . Because the sides are paired to obtain $\mathbb{M}_g$, we only consider $k = 0, \dots, 2g - 1$ to avoid that several points are actually the same on the surface. Just as the side s_k is obtained from s_0 by rotating it around the origin by angle $\frac{k\pi}{2g}$, so p_k^0 is obtained from p_0^0 by rotating it in the same way. Hence, we use the lower index as ‘rotation’ index in the definition of the other points as well.
2. Secondly, the projections $\pi_g(p_k^*)$ of the midpoints p_k^* of the geodesic segment in $\mathbb{H}^2$ connecting consecutive pairs (p_k^0, p_{k+1}^0) of midpoints are added to $\mathcal{Q}_g$. These points are unique in the sense that they will be the only points in $\mathcal{Q}_g$ that have their pre-image in $\tilde{D}_g$ not on some line segment $[O, p_k^0]$. This is the reason why they have a star as superscript.
3. Thirdly, points p_0^j are consecutively added on $[O, p_0^{j-1}]$ in such a way that the distance between consecutive points p_0^{j-1} and p_0^j is given by $d(p_0^j, p_0^{j-1}) = \frac{1}{4} \text{sys}(\mathbb{M}_g)$, until $d(p_0^j, O) \leq \frac{1}{4} \text{sys}(\mathbb{M}_g)$. By rotating the points in the same way as before we obtain the points p_k^j . The projections $\pi_g(p_k^j)$ are added to $\mathcal{Q}_g$. Here, the upper index denotes the ‘iteration’ index. Notice that the midpoints of the sides were denoted p_k^0 to initialize this process. Since $d(O, p_0^0) = \text{arccosh}(\cot(\frac{\pi}{4g}))$, this step consists of

$$m = \left\lceil \frac{\text{arccosh}(\cot(\frac{\pi}{4g}))}{\frac{1}{4} \text{sys}(\mathbb{M}_g)} \right\rceil - 1$$

iterations.

4. Finally, observe that the triangles $[O, p_k^0, p_{k+1}^0]$ and $[v, p_k^0, p_{k+1}^0]$ are congruent under reflection in $[p_k^0, p_{k+1}^0]$. To establish the same congruence in

$\mathcal{Q}_g$, we want to apply this reflection to all points in $[O, \mathbf{p}_k^0, \mathbf{p}_{k+1}^0]$ that are currently in $\pi_g^{-1}(\mathcal{Q}_g) \cap D_g$. However, if we do so directly, we obtain several pairs of points projecting to the same point on $\mathbb{M}_g$. To avoid this, we reflect the points $\mathbf{p}_k^j$ in one half of the fundamental polygon across the line through $\mathbf{p}_k^0$ and $\mathbf{p}_{k+1}^0$ and the points $\mathbf{p}_k^j$ in the other half of the fundamental polygon across the line through $\mathbf{p}_{k-1}^0$ and $\mathbf{p}_k^0$. In each case, the image of the $\mathbf{p}_k^j$ after reflection is denoted by $\mathbf{q}_k^j$ and its projection $\pi_g(\mathbf{q}_k^j)$ is added to $\mathcal{Q}_g$.

One of the major advantages of this algorithm is that the combinatorics of the resulting Delaunay triangulation can be explicitly described. Below we describe these combinatorics. The proof that this is the Delaunay triangulation of the dummy point set will be given in Lemma 3.16. Again, see Figure 3.11 for an illustration of the dummy points and the triangulation $\mathcal{T}$ within one slice of the $4g$ -gon.

Definition 3.14. Define the infinite triangulation $\mathcal{T}$ of $\pi_g^{-1}(\mathcal{Q}_g)$ as follows:

- As vertices take $\pi_g^{-1}(\mathcal{Q}_g)$.
- The edges completely contained in the hyperbolic triangle $[O, \mathbf{p}_0^0, \mathbf{p}_1^0]$ are given by the following list.

$$\begin{aligned} &(\mathbf{p}_k^j, \mathbf{p}_k^{j+1}) & j = 0, \dots, m-1, & k = 0, 1, \\ &(\mathbf{p}_0^j, \mathbf{p}_1^{j+1}) & j = 1, \dots, m-1, \\ &(\mathbf{p}_k^m, O) & k = 0, 1, \\ &(\mathbf{p}_0^j, \mathbf{p}_1^j) & j = 1, \dots, m, \\ &(\mathbf{p}_k^j, \mathbf{p}_0^*) & j = 0, 1, & k = 0, 1. \end{aligned}$$

The other edges can be obtained as the images of the edges in the list above under the following maps:

- rotation around the origin by angle $\frac{k\pi}{2g}$,
- reflection in the line through $\mathbf{p}_0^0, \mathbf{p}_1^0$,
- reflection in the line through $\mathbf{p}_0^0, \mathbf{p}_1^0$, followed by rotation around the origin by angle $\frac{k\pi}{2g}$,
- any one of the above maps, followed by an element of Γ_g .

Algorithm 3 shows the algorithm in pseudocode. We refer to this algorithm as the structured algorithm.

The main difference between the structured algorithm and the other two algorithms is that there is no **while** loop in the structured algorithm, only **for** loops. As a result, the cardinality of the resulting dummy point set is known precisely (see Theorem 3.17). On the other hand, for the cardinality of the dummy point sets for the refinement or symmetric algorithm we can only give an estimate and the exact number of points depends on the implementation.

	<p>Input : hyperbolic surface $\mathbb{M}_g$</p> <p>Output: finite point set $\mathcal{Q}_g \subset \mathbb{M}_g$ such that $\text{diam } \mathcal{Q}_g < \frac{1}{2} \text{sys}(\mathbb{M}_g)$</p> <ol style="list-style-type: none"> 1 Initialize: let $\mathcal{Q}_g$ be the set $\mathcal{W}_g$ of Weierstrass points of $\mathbb{M}_g$. 2 Label the vertex by v and the origin by O. 3 For all $k = 0, \dots, 4g - 1$, label the midpoint of side s_k by p_k^0. 4 For all $k = 0, \dots, 4g - 1$, label the midpoint of p_k^0 and p_{k+1}^0 by p_k^* and add $\pi_g(p_k^*)$ to $\mathcal{Q}_g$. 5 $m \leftarrow \lceil 4 \operatorname{arccosh}(\cot(\frac{\pi}{4g})) / \text{sys}(\mathbb{M}_g) \rceil - 1$. 6 for $i = 1, \dots, m$ do 7 Let p_0^j be the point on $[p_0^{j-1}, O]$ with $d(p_0^j, p_0^{j-1}) = \frac{1}{4} \text{sys}(\mathbb{M}_g)$ and add $\pi_g(p_0^j)$ to $\mathcal{Q}_g$. 8 for $k = 1, \dots, 2g - 1$ do 9 Let p_k^j be p_0^j rotated clockwise around the origin by angle $\frac{k\pi}{2g}$. 10 Let q_k^j be the reflection of p_k^j in the line through p_k^0, p_{k+1}^0. 11 Add $\pi_g(p_k^j)$ and $\pi_g(q_k^j)$ to $\mathcal{Q}_g$. 12 end 13 for $k = 2g, \dots, 4g - 1$ do 14 Let p_k^j be p_0^j rotated clockwise around the origin by angle $\frac{k\pi}{2g}$. 15 Let q_k^j be the reflection of p_k^j in the line through p_{k-1}^0, p_k^0. 16 Add $\pi_g(p_k^j)$ and $\pi_g(q_k^j)$ to $\mathcal{Q}_g$. 17 end 18 end
--	---

Algorithm 3: Structured algorithm

The following two lemmas show that the circumdiameters of triangles $\mathcal{T}$ are smaller than $\frac{1}{2} \text{sys}(\mathbb{M}_g)$ and that $\mathcal{T}$ is a Delaunay triangulation. The proofs are given in Appendix A.1.

Lemma 3.15. *The circumdiameters of triangles in $\mathcal{T}$ are smaller than $\frac{1}{2} \text{sys}(\mathbb{M}_g)$.*

Lemma 3.16. *The triangulation $\mathcal{T}$ is a Delaunay triangulation of $\pi_g^{-1}(\mathcal{Q}_g)$.*

From these two lemmas the main statement of this subsection follows directly.

Theorem 3.17. *The structured algorithm terminates. The resulting dummy point $\mathcal{Q}_g$ set satisfies $\text{diam } \mathcal{Q}_g < \frac{1}{2} \text{sys}(\mathbb{M}_g)$ and its cardinality $|\mathcal{Q}_g|$ is equal to*

$$6g + 2 + 8g(\lceil 4 \operatorname{arccosh}(\cot(\frac{\pi}{4g})) / \text{sys}(\mathbb{M}_g) \rceil - 1).$$

A Delaunay triangulation $\text{DT}_{\mathbb{D}}(\pi_g^{-1}(\mathcal{Q}_g))$ is given by $\mathcal{T}$.

Proof. Termination of the algorithm is trivial. By Lemma 3.15, the resulting dummy point set satisfies $\text{diam } \mathcal{Q}_g < \frac{1}{2} \text{sys}(\mathbb{M}_g)$. The cardinality of $\mathcal{Q}_g$ can be

computed as follows. In line 1, $\mathcal{Q}_g$ contains the $2g + 2$ Weierstrass points of $\mathbb{M}_g$. In line 4 we add $4g$ points to $\mathcal{Q}_g$. There are

$$m = \lceil 4 \operatorname{arccosh}(\cot(\frac{\pi}{4g})) / \operatorname{sys}(\mathbb{M}_g) \rceil - 1$$

iterations of the **for** loop in line 6, each adding $8g$ points $\mathcal{Q}_g$. The cardinality of $\mathcal{Q}_g$ is obtained by adding these expressions. By Lemma 3.16, $\mathcal{T}$ is a Delaunay triangulation. This finishes the proof. $\square$

3.5.4 Experimental results for small genus

The refinement algorithm and the symmetric algorithm have been implemented. The implementation uses the `CORE::Expr` number type [77] to represent coordinates of points, which are algebraic numbers.

For the Bolza surface (genus 2), both algorithms compute a set of 22 dummy points. In Figure 3.12 we have shown the dummy point set computed by the symmetric algorithm. However, a smaller set, consisting of 14 dummy points, was proposed earlier [16]: in addition to the six Weierstrass points, it contains the eight midpoints of the segments $[O, v_k]$, $k = 0, 1, \dots, 7$ (see Figure 3.12).

Figure 3.12: Set of 22 dummy points for the Bolza surface computed by the symmetric algorithm (left) and set of 14 dummy points constructed by hand [16] (right).

The computation does not terminate for higher genus after seven hours of computations when performing the computations exactly. To be able to obtain a result, we impose a finite precision to `CORE::Expr`.

For genus 3, we obtain sets of dummy points with both strategies with precision $512 \times g$ bits (chosen empirically). The refinement algorithm yields a set of 28 dummy points (Figure 3.7), while the symmetric algorithm leads to 32 dummy points (Figure 3.8). Computing dummy point sets for Bolza surfaces of higher

genus poses a challenge regarding the evaluation of algebraic expressions. Our experiments show that we have to design a new strategy for arithmetic computations, which goes beyond the scope of this paper.

3.6 Data structure, predicates, and implementation

In this section, we detail two major aspects of Bowyer's algorithm for generalized Bolza surfaces. On the one hand, the combinatorial aspect, i.e., the data structure and the way it supports the algorithm, is studied in Section 3.6.2. On the other hand, the algebraic degree of the predicates based on which the decisions are made by the algorithm is analyzed in Section 3.6.3. Finally, we report on our implementation and experimental results in Section 3.6.4.

Let us first define a unique canonical representative for each vertex and triangle of a triangulation, which is a major ingredient for the data structure.

3.6.1 Canonical representatives

Let the *original domain* $\tilde{D}_g$ be a subset of D_g containing exactly one representative of each point on the surface $\mathbb{M}_g$, i.e., of each orbit under Γ_g . The original domain $\tilde{D}_g$ is constructed from the fundamental domain D_g as follows (see Figure 3.13): $\tilde{D}_g$ and D_g have the same interior; the only vertex of D_g belonging to $\tilde{D}_g$ is the vertex v_0 ; the $2g$ sides $s_{2g}, \dots, s_{4g-1}$ of D_g preceding v_0 (in counter-clockwise order) are in $\tilde{D}_g$, while the subsequent $2g$ sides are not. For a point p on $\mathbb{M}_g$, the

Figure 3.13: Original domain $\tilde{D}_g$ for $g = 2$. Only vertex v_0 and the solid sides are included in $\tilde{D}_g$.

canonical representative of p is the unique point of the orbit $\pi_g^{-1}(p)$ that lies in $\tilde{D}_g$.

To determine a unique canonical representative for each orbit of a triangle in $\mathbb{D}$ under the action of Γ_g , we first define the set $\mathcal{N}_g$ of *neighboring translations* as:

$$\mathcal{N}_g = \{f \in \Gamma_g \mid f(D_g) \cap D_g \neq \emptyset\}.$$

Each Dirichlet region sharing an edge or a vertex with the (closed) domain D_g is the image of D_g under the action of a translation in $\mathcal{N}_g$, which is used to label the region. Also see Figure 1.4. We denote the union of these *neighboring regions* of D_g by $D_{\mathcal{N}_g}$, so

$$D_{\mathcal{N}_g} = \bigcup_{f \in \mathcal{N}_g} f(D_g).$$

Note that we slightly abuse terminology in the sense that the identity is an element of $\mathcal{N}_g$, and, therefore, a neighboring translation, even though it is not a hyperbolic translation. Also note that D_g is a neighboring region of itself.

We consider all the neighboring regions to be ordered counterclockwise around 0, starting with the Dirichlet region

$$\prod_{j=0}^{2g-1} f_{j(2g+1)}(D_g) = f_0 f_{2g+1} f_{2(2g+1)} \cdots f_{(2g+1)^2}(D_g)$$

(where indices are taken modulo $4g$) incident to v_0 , which gives an ordering of $\mathcal{N}_g \setminus \{\mathbb{1}\}$. An illustration for genus 2 is shown in Figure 3.14.

Figure 3.14: The ordering of $\mathcal{N}_2$ starts with $f_0 f_1^{-1} f_2 f_3^{-1} = f_0 f_5 f_2 f_7$.

We say that a triangle in $\mathbb{D}$ is *admissible* if its circumdiameter is less than half the systole of $\mathbb{M}_g$. We can prove the following property:

Proposition 3.18 (Inclusion property). *If at least one vertex of an admissible triangle is contained in $\tilde{D}_g$, then the whole triangle is contained in $D_{\mathcal{N}_g}$.*

Proof. It is sufficient to show that the distance between the boundary ∂D_g of D_g and the boundary $\partial D_{\mathcal{N}_g}$ of $D_{\mathcal{N}_g}$ is at least $\frac{1}{2} \text{sys}(\mathbb{M}_g)$. Consider points $\mathbf{p} \in \partial D_g$ and $\mathbf{q} \in \partial D_{\mathcal{N}_g}$. We will show that $d(\mathbf{p}, \mathbf{q}) \geq \frac{1}{2} \text{sys}(\mathbb{M}_g)$. By symmetry of D_g , we can assume without loss of generality that $\mathbf{p} \in s_0$. In Section 2.2.3, we gave a definition for a k -segment and a k -separated segment, where the segment is a hyperbolic line segment between sides of D_g . This definition extends naturally to line segments between sides of a translate of D_g .

Figure 3.15: Cases in the proof of Proposition 3.18.

Recall that f_0 is the side-pairing transformation that maps s_{2g} to s_0 . First, assume that $\mathbf{q} \in f_0(D_g)$. Because $\mathbf{q} \in \partial D_{\mathcal{N}_g}$, $[\mathbf{p}, \mathbf{q}]$ is a segment of separation at least 2 (see Figure 3.15a). By Part 2 of Lemma 2.5, $d(\mathbf{p}, \mathbf{q}) \geq \frac{1}{2} \text{sys}(\mathbb{M}_g)$.

Second, assume that $\mathbf{q} \notin f_0(D_g)$. Without loss of generality, we may assume that $\mathbf{q}$ is contained in a translate of D_g that contains either $\mathbf{v}_0$ or $\mathbf{v}_1$ as a vertex. If $\mathbf{p}$ is either $\mathbf{v}_0$ or $\mathbf{v}_1$, then again $[\mathbf{p}, \mathbf{q}]$ is a segment of separation at least 2 (see Figure 3.15b), so $d(\mathbf{p}, \mathbf{q}) \geq \frac{1}{2} \text{sys}(\mathbb{M}_g)$ by Part 2 of Lemma 2.5. If $\mathbf{p}$ is not a vertex of D_g , then $[\mathbf{p}, \mathbf{q}]$ intersects one of the two sides in $f_0(D_g) \setminus \partial D_{\mathcal{N}_g}$, say in a point $\mathbf{r}$ (see Figure 3.15c). In particular, $[\mathbf{p}, \mathbf{r}]$ is a 1-separated segment. If $[\mathbf{r}, \mathbf{q}]$ is a segment of separation at least 2, then $d(\mathbf{r}, \mathbf{q}) \geq \frac{1}{2} \text{sys}(\mathbb{M}_g)$ by Part 2 of Lemma 2.5, so $d(\mathbf{p}, \mathbf{q}) \geq \frac{1}{2} \text{sys}(\mathbb{M}_g)$ as well. If $[\mathbf{r}, \mathbf{q}]$ is a 1-separated segment, then $d(\mathbf{p}, \mathbf{r}) + d(\mathbf{r}, \mathbf{q}) \geq \frac{1}{2} \text{sys}(\mathbb{M}_g)$ by Part 4 of Lemma 2.5.

We have shown that in all cases $d(\mathbf{p}, \mathbf{q}) \geq \frac{1}{2} \text{sys}(\mathbb{M}_g)$, which finishes the proof. $\square$

Remark 3.19. As mentioned in Remark 2.9 in Section 2.3.2, the inclusion property also holds for hyperelliptic surfaces in a neighborhood $\mathbb{M}_g(c(\log(g))^{-1})$ of the generalized Bolza surfaces.

Let now $\mathcal{S} \subset \mathbb{M}_g$ be a set of points satisfying the validity condition (3.1). By definition, all triangles in the Delaunay triangulation $\text{DT}_{\mathbb{D}}(\pi_g^{-1}(\mathcal{S}))$ are admissible and thus satisfy the inclusion property. Let t be a face in the Delaunay triangulation $\text{DT}_{\mathbb{M}_g}(\mathcal{S})$.

By definition of $\tilde{D}_g$, each vertex of t has a unique preimage by π_g in $\tilde{D}_g$, so, the set

$$\Sigma = \left\{ t \in \pi_g^{-1}(t) \mid t \text{ has at least one vertex in } \tilde{D}_g \right\} \quad (3.6)$$

contains at most three faces. See Figure 3.16. When Σ contains only one face, then this face is completely included in $\tilde{D}_g$, and we naturally choose it to be the canonical representative t^c of t . Let us now assume that Σ contains two or three faces. From Proposition 3.18, each face $t \in \Sigma$ is contained in $D_{\mathcal{N}_g}$. So, for each vertex u of t , there is a unique translation $T(u, t)$ in $\mathcal{N}_g$ such that u lies in $T(u, t)(\tilde{D}_g)$. This translation is such that

$$T(u, t)(u^c) = u.$$

Considering the triangles in $\mathbb{D}$ to be oriented counterclockwise, for $t \in \Sigma$, we denote as u_t^* the first vertex of t that is not lying in $\tilde{D}_g$. Using the ordering on $\mathcal{N}_g$ defined above, we can now choose t^c as the face of Σ for which $T(u_t^*, t^c)$ is closest to $f_0 f_{2g+1} f_{2(2g+1)} \dots f_{(2g+1)^2}$ for the counterclockwise order on $\mathcal{N}_g$.

Figure 3.16: Examples (for $g = 2$) of faces of $\text{DT}_{\mathbb{D}}(\pi_g^{-1}(S))$ with one (left), two (middle) and three (right) vertices in $\tilde{D}_g$ that project to the same face on $\mathbb{M}_g$. Their respective vertices drawn as a dot project to the same vertex on $\mathbb{M}_g$ (same for cross and square). The canonical representative is the shaded face.

To summarize, we have shown that:

Proposition 3.20. *Let $S \subset \mathbb{M}_g$ be a set of points satisfying the validity condition (3.1). For any face t in $\text{DT}_{\mathbb{M}_g}(S)$, there exists a unique canonical representative $t^c \subset D_{\mathcal{N}_g}$ in $\text{DT}_{\mathbb{D}}(\pi_g^{-1}(S))$.*

Using a slight abuse of vocabulary, for a triangle t in $\mathbb{D}$, we will sometimes refer to the canonical representative t^c of its projection $t = \pi_g(t)$ as the canonical representative of t .

3.6.2 Data structure

Proposition 3.20 allows us to propose a data structure to represent Delaunay triangulations of generalized Bolza surfaces.

A triangulation of a point set $\mathcal{S} \subset \mathbb{M}_g$ is represented via its vertices and triangular faces. Each vertex u stores its canonical representative $\mathbf{u}^c$ in $\tilde{D}_g$ and gives access to one of its incident triangles. Each triangle t is actually storing information to construct its canonical representative $\mathbf{t}^c$: it gives access to its three incident vertices u_0, u_1 , and u_2 and its three adjacent faces; it also stores the three translations $T(u_j, t) := T(\mathbf{u}_j, \mathbf{t}^c)$, $j = 0, 1, 2$ in $\mathcal{N}_g$ as defined in Section 3.6.1, so that applying each translation to the corresponding canonical point yields the canonical representative $\mathbf{t}^c$ of t , i.e.,

$$\mathbf{t}^c = (T(u_0, t)(\mathbf{u}_0^c), T(u_1, t)(\mathbf{u}_1^c), T(u_2, t)(\mathbf{u}_2^c)).$$

In the rest of this section, we show how this data structure supports the algorithm that was briefly sketched in Section 3.2.2.

Finding conflicts. The notion of conflict defined in section 3.2.2 can now be made more explicit: a triangle $t \in \text{DT}_{\mathbb{M}_g}(\mathcal{S})$ is in conflict with a point $p \in \mathbb{M}_g$ if the circumscribing disk of one of the (at most three) triangles in Σ is in conflict with $\mathbf{p}^c$, where Σ is the set defined by relation (3.6).

By the correspondence between Euclidean circles and hyperbolic circles in the Poincaré disk model, the triangle in the Delaunay triangulation in $\mathbb{D}$ whose associated Euclidean triangle contains the point $\mathbf{p}^c$ is in conflict with this point; these Euclidean and hyperbolic triangles will both be denoted as $\mathbf{t}_p$, which should not introduce any confusion. To find this triangle, we adapt the so-called *visibility walk* [31]: the walk starts from an arbitrary face, then, for each visited face, it visits one of its neighbors, until the face whose associated Euclidean triangle contains $\mathbf{p}^c$ is found. This walk will be detailed below.

We first need some notation. Let t, t' be two adjacent faces in $\text{DT}_{\mathbb{M}_g}(\mathcal{S})$. We define the *neighbor translation* $T^{\text{nbr}}(\mathbf{t}'^c, \mathbf{t}^c)$ from $\mathbf{t}'^c$ to $\mathbf{t}^c$ as the translation of Γ_g such that $T^{\text{nbr}}(\mathbf{t}'^c, \mathbf{t}^c)(\mathbf{t}'^c)$ is adjacent to $\mathbf{t}^c$ in $\text{DT}_{\mathbb{D}}(\pi^{-1}(\mathcal{S}))$. See Figure 3.17. Let u be a vertex common to t and t' , and let $\mathbf{u}_j$ and $\mathbf{u}_{j'}$ be the vertices of $\mathbf{t}^c$ and $\mathbf{t}'^c$ that project on u by π_g . We can compute the neighbor translation from $\mathbf{t}'^c$ to $\mathbf{t}^c$ as $T^{\text{nbr}}(\mathbf{t}'^c, \mathbf{t}^c) = T(u_j, t)(T(u_{j'}, t'))^{-1}$. It can be easily seen that $T^{\text{nbr}}(\mathbf{t}'^c, \mathbf{t}^c) = T(u_j, t)(T(u_{j'}, t'))^{-1} = (T(u_{j'}, t')(T(u_j, t))^{-1})^{-1} = (T^{\text{nbr}}(\mathbf{t}^c, \mathbf{t}'^c))^{-1}$.

Figure 3.17: Translating $\mathbf{t}'^c$ by $T^{\text{nbr}}(\mathbf{t}'^c, \mathbf{t}^c)$ gives a face adjacent to $\mathbf{t}^c$.

Finally, we define the *location translation* T_p^{loc} as the translation that moves the canonical face $\mathbf{t}_p^c$ to $\mathbf{t}_p$. This translation is computed during the walk. The walk starts from a face containing the origin. As this face is necessarily canonical, T_p^{loc} is initialized to $\mathbb{1}$. Then, for each visited face $\mathbf{t}$ of $\text{DT}_{\mathbb{D}}(\pi^{-1}(\mathcal{S}))$, we consider the Euclidean edge $\mathbf{e}$ defined by two of the vertices of $\mathbf{t}$. We check whether the Euclidean line supporting $\mathbf{e}$ separates $\mathbf{p}^c$ from the vertex of $\mathbf{t}$ opposite to $\mathbf{e}$. If this is the case, the next visited face is the neighbor $\mathbf{t}'$ of $\mathbf{t}$ through $\mathbf{e}$; the location translation is updated: $T_p^{\text{loc}} := T_p^{\text{loc}} T_p^{\text{nbr}}(\mathbf{t}'^c, \mathbf{t}^c)$. The walk stops when it finds the Euclidean triangle $\mathbf{t}_p$ containing $\mathbf{p}^c$. Then the canonical face $\mathbf{t}_p^c$ in conflict with $\mathbf{p}^c$ is $(T_p^{\text{loc}})^{-1}(\mathbf{t}_p)$. See Figure 3.18 for an example. Here the walk first visits canonical faces and reaches the face $\mathbf{t}_D \subset D_g$; up to that stage, T_p^{loc} is unchanged. Then the walk visits the non-canonical neighbor $\mathbf{t}'$ of $\mathbf{t}_D$, and T_p^{loc} is updated to $T_p^{\text{nbr}}(\mathbf{t}'^c, \mathbf{t}_D)$. The next face visited by the walk is $\mathbf{t}_p$, which contains $\mathbf{p}^c$; as $\mathbf{t}_p^c$ and $\mathbf{t}'^c$ are adjacent, T_p^{loc} is left unchanged.

Let us now present the computation of the set $\mathbf{C}_p$ of faces of $\text{DT}_{\mathbb{D}}(\pi^{-1}(\mathcal{S}))$ in conflict with $\mathbf{p}^c$. Starting from $\mathbf{t}_p$, for each face of $\text{DT}_{\mathbb{D}}(\pi^{-1}(\mathcal{S}))$ in conflict with $\mathbf{p}^c$ we recursively examine each neighbor (obtained with a neighbor translation) that has not yet been visited, checking it for conflict with $\mathbf{p}^c$. When a face is found to be in conflict, we temporarily store directly in each of its vertices the translation that moves its corresponding canonical point to it (we cannot store such translations in the face itself, since this face will be deleted by the insertion). Since the union of the faces of $\mathbf{C}_p$ is a topological disk by definition, the resulting translation for a given vertex is the same for all faces of $\mathbf{C}_p$ incident to it, so this translation is well defined for each vertex. The temporary translations will be used during the insertion stage described below. We store the set $\mathbf{C}_p^c$ of canonical faces corresponding to faces of $\mathbf{C}_p$. Note that $\mathbf{C}_p^c$ is not necessarily a connected region in $\mathbb{D}$, as illustrated in Figure 3.18(Right).

Inserting a point. To actually insert the new point p on $\mathbb{M}_g$, we first create a new vertex storing $\mathbf{p}^c$. We store $\mathbb{1}$ as the temporary translation in the new vertex.

For each edge $\mathbf{e}$ on the boundary of $\mathbf{C}_p$, we create a new face $\mathbf{t}_e$ on $\mathbb{M}_g$ corresponding to the triangle $\mathbf{t}_e$ in $\mathbb{D}$ formed by the new vertex and the edge $\mathbf{e}$. The neighbor of $\mathbf{t}_e$ through $\mathbf{e}$ is the neighbor through $\mathbf{e}$ of the face in $\mathbf{C}_p^c$ that is incident to $\mathbf{e}$. Two new faces consecutive along the boundary of $\mathbf{C}_p$ are adjacent. We now delete all faces in $\mathbf{C}_p$. The triangle $\mathbf{t}_e$ is not necessarily the canonical representative of $\mathbf{t}_e$; we must now compute the three translations to be stored in $\mathbf{t}_e$ to get $\mathbf{t}_e^c$. To this aim, we first retrieve the translations temporarily stored in its vertices $u_j, j = 0, 1, 2$ and we respectively initialize the translations $T(u_j, \mathbf{t}_e)$ in $\mathbf{t}_e$ to them. If all translations are equal to $\mathbb{1}$, then the face is already canonical and there is nothing more to do. Otherwise, the translations stored in the face are updated following Section 3.6.1: $T(u_j, \mathbf{t}_e) := (T(u_k, \mathbf{t}_e))^{-1} T(u_j, \mathbf{t}_e), j = 0, 1, 2$, where k is the index in $\{0, 1, 2\}$ for which $\mathbf{u}_k = \mathbf{u}_{\mathbf{t}_e}^*$.

Figure 3.18: Left: The shaded faces are the (not necessarily canonical) faces in C_p , i.e., faces in conflict with the red point p . Their union is a topological disk. Right: The region C_p^c of the (shaded) corresponding canonical faces is not connected in $\mathbb{D}$.

Once this is done for all new faces, temporary translations stored in vertices can be removed.

3.6.3 Degree of predicates

Following the celebrated exact geometric computation paradigm [76], the correctness of the combinatorial structure of the Delaunay triangulation relies on the exact evaluation of predicates. The main two predicates are

- **ORIENTATION**, which checks whether an input point p in $\tilde{D}_g$ lies on the right side, the left side, or on an oriented Euclidean segment.
- **INCIRCLE**, which checks whether an input point p in $\tilde{D}_g$ lies inside, outside, or on the boundary of the disk circumscribing an oriented triangle.

Input points, which lie in $\tilde{D}_g$, are tested against canonical triangles of the triangulation, whose vertices are images of input points by translations in $\mathcal{N}_g$. If points are assumed to have rational coordinates, then evaluating the predicates boils down to determining the sign of polynomial expressions whose coefficients are lying in some extension field of the rationals. Proposition 3.21 gives an upper bound on the degree of these polynomial expressions. For the special case of the Bolza surface ($g = 2$), it improves the previously known upper bound from 72 [45, Proposition 1], which was proved using symbolic computations in MAPLE, to 48.

Proposition 3.21. *For the generalized Bolza surface of genus $g \geq 2$, the predicates can be evaluated by determining the sign of rational polynomial expressions of total degree at most $12\varphi(4g) \leq 24g$ in the coordinates of the input points, where φ is the Euler totient function.*

Recall that the Euler totient function $\varphi(n)$ counts the number of integers up to a given integer n that are relatively prime to n [54].

Proof. We will only consider the INCIRCLE predicate; the strategy for determining the maximum degree for the ORIENTATION predicate is similar and the resulting maximum degree is lower. The INCIRCLE predicate is given by the sign of

$$\begin{aligned} \text{INCIRCLE}(\mathbf{p}_1, \mathbf{p}_2, \mathbf{p}_3, \mathbf{p}_4) &= \begin{vmatrix} x_1 & y_1 & x_1^2 + y_1^2 & 1 \\ x_2 & y_2 & x_2^2 + y_2^2 & 1 \\ x_3 & y_3 & x_3^2 + y_3^2 & 1 \\ x_4 & y_4 & x_4^2 + y_4^2 & 1 \end{vmatrix} \\ &= x_3 y_4 (x_1^2 + y_1^2) - x_3 y_4 (x_2^2 + y_2^2) + \dots, \end{aligned}$$

for points $\mathbf{p}_j = x_j + y_j i, j = 1, \dots, 4$, in $\mathbb{D}$. In the second equality, we have just written 2 of the 24 terms to illustrate what the terms look like. This will be used later in the proof to determine their maximum degree. Since at least one of the four points is contained in $\bar{D}_g$, we will assume without loss of generality that $x_1, y_1 \in \mathbb{Q}$. The other points are images of points with rational coordinates under some elements of Γ_g . We know that Γ_g is generated by f_k for $k = 0, \dots, 4g - 1$. The translation f_k can be represented by the matrix A_k (see Equation (1.2) in Section 1.1.7). The entries of A_k are contained in the extension field

$$L = \mathbb{Q} \left[\zeta_{4g}, \sqrt{\cot^2\left(\frac{\pi}{4g}\right) - 1} \right],$$

where $\zeta_{4g} = \exp(\frac{\pi i}{2g})$ is a primitive $4g$ -th root of unity. The field L is an extension field of degree 2 of the cyclotomic field $\mathbb{Q}[\zeta_{4g}]$, which is an extension field of degree $\varphi(4g)$ of $\mathbb{Q}$ so the total degree of L as an extension field of $\mathbb{Q}$ is $2\varphi(4g)$. Later in the proof, we will actually look at the degree of the field $L \cap \mathbb{R}$ over $\mathbb{Q}$. Because $L \cap \mathbb{R}$ is the fixed field of L under complex conjugation, L is a quadratic extension of $L \cap \mathbb{R}$. Therefore, the degree of $L \cap \mathbb{R}$ as an extension field of $\mathbb{Q}$ is $\varphi(4g)$. Since each translation in Γ can be represented by a product of matrices A_k , it follows that for $j = 2, 3, 4$ we can write

$$x_j + y_j i = \frac{\alpha_j(x_j^c + y_j^c i) + \beta}{\bar{\beta}(x_j^c + y_j^c) + \bar{\alpha}},$$

where x_j^c and y_j^c are the (rational) coordinates of the canonical representative of $x_j + y_j i$ and where α_j and β_j are elements of L . As usual, we can get rid of the i in the denominator by multiplying numerator and denominator by the complex conjugate of the denominator. Because both the numerator and denominator are linear as function of x_j^c and y_j^c , we obtain

$$x_j + y_j i = \frac{P_j(x_j^c, y_j^c) + Q_j(x_j^c, y_j^c)i}{R_j(x_j^c, y_j^c)},$$

where P_j, Q_j and R_j are polynomials in x_j^c and y_j^c of total degree at most 2 with coefficients in $L \cap \mathbb{R}$. Note that we indeed know that the coefficients are real numbers, since by construction we have already split the real and imaginary parts. Hence, suppressing the dependencies on x_j^c and y_j^c , we see that

$$\begin{aligned} & \text{INCIRCLE}(\mathbf{p}_1, \mathbf{p}_2, \mathbf{p}_3, \mathbf{p}_4) \\ &= \frac{P_3 Q_4 (x_1^2 + y_1^2)}{R_3 R_4} - \frac{P_3 Q_4 (P_2^2 + Q_2^2)}{R_2^2 R_3 R_4} + \dots, \\ &= \frac{R_2^2 P_3 R_3 Q_4 R_4 (x_1^2 + y_1^2) - P_3 R_3 Q_4 R_4 (P_2^2 + Q_2^2) + \dots}{R_2^2 R_3^2 R_4^2}. \end{aligned}$$

Now, testing whether $\text{INCIRCLE}(\mathbf{p}_1, \mathbf{p}_2, \mathbf{p}_3, \mathbf{p}_4) > 0$ amounts to testing whether

$$R_2^2 P_3 R_3 Q_4 R_4 (x_1^2 + y_1^2) - P_3 R_3 Q_4 R_4 (P_2^2 + Q_2^2) + \dots > R_2^2 R_3^2 R_4^2.$$

Since all P_j, Q_j and R_j are polynomials in x_j^c and y_j^c of degree at most 2, this reduces to evaluating a polynomial of total degree at most 12 in the coordinates of the input points, with coefficients in $L \cap \mathbb{R}$. Because $L \cap \mathbb{R}$ is an extension field of $\mathbb{Q}$ of degree $\varphi(4g)$, we conclude that evaluating $\text{INCIRCLE}(\mathbf{p}_1, \mathbf{p}_2, \mathbf{p}_3, \mathbf{p}_4)$ amounts to determining the sign of a polynomial of total degree at most $12\varphi(4g)$ with rational coefficients. To prove $12\varphi(4g) \leq 24g$, we write $g = 2^k g'$ where g' is odd. Then

$$\varphi(4g) = \varphi(2^{k+2} g') = \varphi(2^{k+2}) \varphi(g') = (2^{k+2} - 2^{k+1}) \varphi(g') = 2^{k+1} \varphi(g').$$

If $g' = 1$, then $\varphi(g') = 1$, so $\varphi(4g) = 2^{k+1} = 2g$. If $g' > 1$, then $\varphi(g') \leq g' - 1$, so $\varphi(4g) \leq 2^{k+1}(g' - 1) \leq 2(g - 1)$. Hence, in both cases $\varphi(4g) \leq 2g$. This finishes the proof. $\square$

3.6.4 Implementation and experimental results

The algorithm presented in Section 3.2 was implemented in C++, with the data structure described in Section 3.6.2. The preprocessing step consists in computing dummy points that serve for the initialization of the data structure, following the two options presented in Section 3.5. The implementation also uses the value of the systole given by Theorem 2.1.

Let us continue the discussion on predicates. In practice, the implementation relies on the `CORE::Expr` number type [77], which provides us with exact and filtered computations. As for the computation of dummy points (Section 3.5.4), the evaluation exceeds the capabilities of `CORE` for genus bigger than 2, due to the barriers raised by their very high algebraic degree, so, only a non-robust implementation of the algorithm can be obtained.

The rest of this section is devoted to the implementation for the Bolza surface, for which a fully robust implementation has been integrated in `CGAL` [46]. All

details can be found in Iordanov's PhD thesis [44]. We only mention a few key points here.

To avoid increasing further the algebraic degree of predicates, the coordinates of dummy points are rounded to rationals (see Table 3.1). We have checked that the validity condition (3.1) still holds for the rounded points, and that the combinatorics of the Delaunay triangulations of exact and rounded points are identical.

Table 3.1: Exact and rational expressions for the dummy points for the Bolza surface. The midpoint of side s_j of the fundamental domain is denoted as $\mathbf{m}_j$. The midpoint of segment $[0, \mathbf{v}_j]$ is denoted as $\mathbf{p}_j$.

Point	Expression	Rational approximation
$\mathbf{v}_0$	$\left(\frac{2^{3/4}\sqrt{2+\sqrt{2}}}{4}, -\frac{2^{3/4}\sqrt{2-\sqrt{2}}}{4} \right)$	$(97/125, -26/81)$
$\mathbf{m}_4$	$\left(-\sqrt{\sqrt{2}-1}, 0 \right)$	$(-9/14, 0)$
$\mathbf{m}_5$	$\left(-\frac{\sqrt{2}\sqrt{\sqrt{2}-1}}{2}, -\frac{\sqrt{2}\sqrt{\sqrt{2}-1}}{2} \right)$	$(-5/11, -5/11)$
$\mathbf{m}_6$	$\left(0, -\sqrt{\sqrt{2}-1} \right)$	$(0, -9/14)$
$\mathbf{m}_7$	$\left(\frac{\sqrt{2}\sqrt{\sqrt{2}-1}}{2}, -\frac{\sqrt{2}\sqrt{\sqrt{2}-1}}{2} \right)$	$(5/11, -5/11)$
$\mathbf{p}_0$	$\left(\frac{2^{1/4}\sqrt{2+\sqrt{2}}}{2\sqrt{2}+2\sqrt{2-\sqrt{2}}}, -\frac{2^{1/4}\sqrt{2-\sqrt{2}}}{2\sqrt{2}+2\sqrt{2-\sqrt{2}}} \right)$	$(1/2, -4/19)$
$\mathbf{p}_1$	$\left(\frac{2^{3/4}(\sqrt{2+\sqrt{2}}+\sqrt{2-\sqrt{2}})}{4\sqrt{2}+4\sqrt{2-\sqrt{2}}}, \frac{2^{3/4}(\sqrt{2+\sqrt{2}}-\sqrt{2-\sqrt{2}})}{4\sqrt{2}+4\sqrt{2-\sqrt{2}}} \right)$	$(1/2, 4/19)$
$\mathbf{p}_2$	$\left(\frac{2^{1/4}\sqrt{2-\sqrt{2}}}{2\sqrt{2}+2\sqrt{2-\sqrt{2}}}, \frac{2^{1/4}\sqrt{2+\sqrt{2}}}{2\sqrt{2}+2\sqrt{2-\sqrt{2}}} \right)$	$(4/19, 1/2)$
$\mathbf{p}_3$	$\left(\frac{2^{3/4}(\sqrt{2-\sqrt{2}}-\sqrt{2+\sqrt{2}})}{4\sqrt{2}+4\sqrt{2-\sqrt{2}}}, \frac{2^{3/4}(\sqrt{2+\sqrt{2}}+\sqrt{2-\sqrt{2}})}{4\sqrt{2}+4\sqrt{2-\sqrt{2}}} \right)$	$(-4/19, 1/2)$
$\mathbf{p}_4$	$\left(-\frac{2^{1/4}\sqrt{2+\sqrt{2}}}{2\sqrt{2}+2\sqrt{2-\sqrt{2}}}, \frac{2^{1/4}\sqrt{2-\sqrt{2}}}{2\sqrt{2}+2\sqrt{2-\sqrt{2}}} \right)$	$(-1/2, 4/19)$
$\mathbf{p}_5$	$\left(-\frac{2^{3/4}(\sqrt{2+\sqrt{2}}+\sqrt{2-\sqrt{2}})}{4\sqrt{2}+4\sqrt{2-\sqrt{2}}}, \frac{2^{3/4}(\sqrt{2-\sqrt{2}}-\sqrt{2+\sqrt{2}})}{4\sqrt{2}+4\sqrt{2-\sqrt{2}}} \right)$	$(-1/2, -4/19)$
$\mathbf{p}_6$	$\left(-\frac{2^{1/4}\sqrt{2-\sqrt{2}}}{2\sqrt{2}+2\sqrt{2-\sqrt{2}}}, -\frac{2^{1/4}\sqrt{2+\sqrt{2}}}{2\sqrt{2}+2\sqrt{2-\sqrt{2}}} \right)$	$(-4/19, -1/2)$
$\mathbf{p}_7$	$\left(\frac{2^{3/4}(\sqrt{2+\sqrt{2}}-\sqrt{2-\sqrt{2}})}{4\sqrt{2}+4\sqrt{2-\sqrt{2}}}, -\frac{2^{3/4}(\sqrt{2-\sqrt{2}}+\sqrt{2+\sqrt{2}})}{4\sqrt{2}+4\sqrt{2-\sqrt{2}}} \right)$	$(4/19, -1/2)$

Attention has also been paid to the manipulation of translations. As seen in Section 3.6.2, translations are composed during the execution of the algorithm. To avoid performing the same multiplications of matrices several times, we actually

represent a translation as a word on the elements of $\mathbb{Z}_8$, where $\mathbb{Z}_8$ is considered as an alphabet and each element corresponds to a generator of Γ_2 . The composition of two translations corresponds to the concatenation of their two corresponding words. Section 3.6.2 showed that only the finitely many translations in $\mathcal{N}_2$ must be stored in the data structure. Moreover, words that appear during the various steps of the algorithm can be reduced by Dehn’s algorithm [28, 38], yielding a finite number of words to be stored, so, a map can be used to associate a matrix to each word. Dehn’s algorithm terminates in a finite number of steps and its time complexity is polynomial in the length of the input word. From Sections 3.6.1 and 3.6.2, words to be reduced are formed by the concatenation of two or three words corresponding to elements of $\mathcal{N}_2$, whose length is not more than four, so, the longest words to be reduced have length 12.

Running times have been measured on a MacBook Pro (2015) with processor Intel Core i5, 2.9 GHz, 16 GB and 1867 MHz RAM, running MacOS X (10.10.5). The code was compiled with clang-700.1.81. We generate 1 million points in the half-open octagon $\tilde{D}_2$ and construct four triangulations:

- a CGAL Euclidean Delaunay triangulation with `double` as number type,
- a CGAL Euclidean Delaunay triangulation with `CORE::Expr` as number type,
- our Delaunay triangulation of the Bolza with `double` as number type,
- our Delaunay triangulation of the Bolza surface with `CORE::Expr` as number type,

Note that the implementations using `double` are not robust and are only considered for the purpose of this experimentation. The insertion times are averaged over 10 executions. The results are reported in Table 3.2.

	Runtime (in seconds)
Euclidean DT (<code>double</code>)	1
Euclidean DT (<code>CORE::Expr</code>)	24
Bolza DT (<code>double</code>)	16
Bolza DT (<code>CORE::Expr</code>)	55

Table 3.2: Runtimes for the computation of Delaunay triangulations of 1 million random points in the half-open octagon $\tilde{D}_2$.

The experiments confirm the influence of the algebraic demand for the Bolza surface: almost two thirds of the running time is spent in predicate evaluations. Also, it was observed that only 0.76% calls to predicates involve translations in $\mathcal{N}_2$, but these calls account for 36% of the total time spent in predicates.

Note also that the triangulation can quickly be cleared of dummy points: in most runs, all dummy points are removed from the triangulation after the insertion of 30 to 70 points.

3.7 Conclusion and open problems

We have extended Bowyer’s algorithm to the computation of Delaunay triangulations of point sets on generalized Bolza surfaces, a particular type of hyperbolic surfaces. A challenging open problem is the generalization of our algorithm to arbitrary hyperbolic surfaces.

One of the main ingredients of our extension of Bowyer’s algorithm is the validity condition (3.1), and to be able to say whether it holds or not we need to know the value of the systole of the hyperbolic surface. For general hyperbolic surfaces an explicit value, or a ‘reasonable’ lower bound of the systole, is not known, and there are no efficient algorithms to compute or approximate it. The effective procedure presented in [2] is based on the construction of a pants decomposition of a hyperbolic surface, and computes the systole from the Fenchel-Nielsen coordinates associated with this decomposition. However, the complexity of this algorithm does not seem to be known, and it is not clear how to turn this method into an efficient and robust algorithm.

If the systole is known, then it seems that we can use the refinement algorithm presented in Section 3.5.1 to compute a dummy point set satisfying the validity condition. However, in the case of generalized Bolza surfaces it is sufficient to consider only the translates of vertices in $D_{\mathcal{N}_g}$ by Proposition 3.10, whereas it is not clear how many translates are needed for an arbitrary hyperbolic surface.

A more modest attempt towards generalization could focus on hyperbolic surfaces represented by a ‘nice’ fundamental polygon. Hyperelliptic surfaces have a point-symmetric fundamental polygon (See [67]), so these surfaces are obvious candidates for future work.

Chapter 4

Minimal Delaunay triangulations of hyperbolic surfaces

4.1 Introduction

In this chapter we consider the minimal number of vertices of a simplicial Delaunay triangulation of a closed hyperbolic surface of genus g . Motivated by the interest in embeddings where edges are shortest paths between their endpoints [35, 42], which have applications in for example the field of graph drawing [71], we restrict ourselves to *distance* Delaunay triangulations, where edges are distance paths, i.e., shortest paths between their endpoints.

Our main result is the upper bound on the number of vertices with sharp order of growth:

Theorem 4.1. *An orientable closed hyperbolic surface of genus $g \geq 2$ has a distance Delaunay triangulation with at most $O(g)$ vertices. Furthermore, there exists a family of surfaces, X_g , $g \geq 2$, such that the number of vertices of any distance Delaunay triangulation of them grows like $\Omega(g)$.*

The above result is a compilation of Theorems 4.2 and 4.17 where explicit upper and lower bounds are given.

Another reason to study triangulations whose edges are distance paths, comes from the study of moduli spaces $\mathcal{M}_g$. As mentioned in Section 1.1.9, these spaces admit natural coordinates associated to pants decompositions, the so-called Fenchel-Nielsen coordinates. It is a classical theorem of Bers [12] that any surface admits a short pants decomposition, meaning that the length of each of its simple closed geodesics is bounded by a function that only depends on the topology of the surface (but not its geometry). As these curves provide a local description of the surface, one might hope that they are also geodesically convex, meaning that the shortest distance path between any two points of a given curve is contained in the curve. It is perhaps surprising that *most* surfaces admit no short

pants decompositions with geodesically convex curves. Indeed it is known that *any* pants decomposition of a random surface (chosen with respect to a natural probability measure on $\mathcal{M}_g$) has at least one curve of length on the order of $g^{\frac{1}{6}-\varepsilon}$ as g grows (for any fixed $\varepsilon > 0$) [39]. And it is a theorem of Mirzakhani that these same random surfaces are also of *diameter* on the order of $\log(g)$ [57]. Hence the longest curve of any pants decomposition of a random surface is not convex.

The lengths of edges in a given triangulation are another parameter set for $\mathcal{M}_g$. By the theorem above, such a parameter set can be chosen with a reasonable number of vertices such that the edges are all convex. Using the moduli space point of view, one has a function $\omega : \mathcal{M}_g \rightarrow \mathbb{N}$ which associates to a surface the minimal number of vertices of any of its distance Delaunay triangulations. The above result implies that

$$\limsup_{g \rightarrow \infty} \max_{X \in \mathcal{M}_g} \frac{\omega(X)}{g}$$

is finite and strictly positive, but for instance we do not know whether the actual limit exists.

The examples we exhibit are geometrically quite simple, as they are made by gluing hyperbolic pants, with bounded cuff lengths, in something that resembles a line as the genus grows. One might wonder whether all surfaces have this property, but we show this is not the case by exploring the quantity $\min_{X \in \mathcal{M}_g} \omega(X)$. This quantity has a precise lower bound on the order of $\Theta(\sqrt{g})$ because we ask that our triangulations be simplicial [48]. We show how to use the celebrated Ringel-Youngs construction [65] to construct a family of hyperbolic surfaces that attain this bound for infinitely many genera (Theorem 4.25), showing that one cannot hope for better than the simplicial lower bound in general.

Although our results provide a good understanding on the extremal values of ω , there are still plenty of unexplored questions. For example, what is the behavior of ω for a random surface (using Mirzakhani's notion of randomness [57] alluded to above)?

This chapter is structured as follows. In Section 4.2, we prove our linear upper bound for the number of vertices of a minimal distance Delaunay triangulation. In Section 4.3, we construct classes of hyperbolic surfaces attaining the order of this linear upper bound. Finally, in Section 4.4, we construct a family of hyperbolic surfaces attaining the general $\Theta(\sqrt{g})$ lower bound.

4.2 Linear upper bound for the number of vertices of a minimal distance Delaunay triangulation

As our first result, we prove that for every hyperbolic surface there exists a distance Delaunay triangulation whose cardinality grows linearly as a function of the genus. For convenience, the set of all distance Delaunay triangulations of a closed hyperbolic surface $\mathbb{M}$ is denoted by $\mathcal{D}(\mathbb{M})$.

Theorem 4.2. *For every closed hyperbolic surface $\mathbb{M}$ of genus g there exists a distance Delaunay triangulation $T \in \mathcal{D}(\mathbb{M})$ with at most $151g$ vertices.*

Note that the constant 151 is certainly not optimal. The idea of the proof is the following. Given a hyperbolic surface $\mathbb{M}$, we construct a vertex set $\mathcal{P}$ on $\mathbb{M}$ consisting of at most $151g$ vertices such that the projection T of a Delaunay triangulation of $\pi^{-1}(\mathcal{P})$ in $\mathbb{D}$ to $\mathbb{M}$ is a distance Delaunay triangulation of $\mathbb{M}$.

It is known that T is a simplicial complex if $\mathcal{P}$ is sufficiently dense and well-distributed [15]. More precisely, there are no 1- or 2-cycles in T if the diameter of the largest disk in $\mathbb{D}$ not containing any points of $\pi^{-1}(\mathcal{P})$ is less than $\frac{1}{2} \text{sys}(\mathbb{M})$, where $\text{sys}(\mathbb{M})$ is the systole of $\mathbb{M}$, i.e. the length of the shortest homotopically non-trivial closed curve. However, the systole of a hyperbolic surface can be arbitrarily close to zero, which means that we would need an arbitrarily dense set $\mathcal{P}$ to satisfy this condition.

Instead, for a constant $\varepsilon > 0$ we subdivide $\mathbb{M}$ into its ε -thick part

$$\mathbb{M}_{\text{thick}}^\varepsilon = \{x \in \mathbb{M} \mid \text{injrads}(x) > \varepsilon\}$$

and its ε -thin part $\mathbb{M}_{\text{thin}}^\varepsilon = \mathbb{M} \setminus \mathbb{M}_{\text{thick}}^\varepsilon$, where $\text{injrads}(x)$ is the injectivity radius at x , i.e., the radius of the largest embedded open disk centered at x . Note that the minimum of $\text{injrads}(x)$ over all $x \in \mathbb{M}$ is given by $\frac{1}{2} \text{sys}(\mathbb{M})$. We will see in Section 4.2.1 that, for sufficiently small ε , $\mathbb{M}_{\text{thin}}^\varepsilon$ is a collection of hyperbolic cylinders (see Figure 4.1). In these hyperbolic cylinders we want to construct a set of vertices the cardinality of which does not depend on $\text{sys}(\mathbb{M})$. To do this, we put three vertices on the “waist” and each of the two boundary components of the cylinders that are “long and narrow”. In the cylinders that are not “long and narrow” it suffices to place three vertices on its waist only. The notions of “waist” and “long and narrow” will be specified in Section 4.2.1. Because $\text{injrads}(x) > \varepsilon$ for all $x \in \mathbb{M}_{\text{thick}}^\varepsilon$, we can construct a sufficiently dense and well-distributed point set in $\mathbb{M}_{\text{thick}}^\varepsilon$ whose cardinality does not depend on $\text{sys}(\mathbb{M})$ but only on ε . In Section 4.2.2 we will describe how we combine the vertices placed in the hyperbolic cylinders with the dense and well-distributed set of vertices in $\mathbb{M}_{\text{thick}}^\varepsilon$. Finally, the proof of Theorem 4.2 is given in Section 4.2.3.

4.2.1 Distance Delaunay triangulations of hyperbolic cylinders

We now describe our construction of a set of vertices for the ε -thin part $\mathbb{M}_{\text{thin}}^\varepsilon$ of the hyperbolic surface $\mathbb{M}$. The following lemma describes $\mathbb{M}_{\text{thin}}^\varepsilon$ in more detail.

Lemma 4.3 ([20, Theorem 4.1.6]). *If $\varepsilon < \text{arcsinh}(1)$ then $\mathbb{M}_{\text{thin}}^\varepsilon$ is a collection of at most $3g - 3$ pairwise disjoint hyperbolic cylinders.*

The following description of the geometry of the hyperbolic cylinders in $\mathbb{M}_{\text{thin}}^\varepsilon$ is based primarily on a similar description in the context of colourings of hyperbolic surfaces [63]. Each hyperbolic cylinder C in $\mathbb{M}_{\text{thin}}^\varepsilon$ consists of points with injectivity

Figure 4.1: Decomposition of a hyperbolic surface into a thick part consisting of two connected components and two narrow hyperbolic cylinders (in red).

radius at most ε and the boundary curves γ^+ and γ^- consist of all points with injectivity radius equal to ε . Every point on the boundary curves is the base point of an embedded geodesic loop of length 2ε (Figure 4.2), which is completely contained in the hyperbolic cylinder. All points on the boundary curves have the same distance K_C to a closed geodesic γ (called the *waist* of C), where K_C only depends on ε and the length $\ell(\gamma)$ of γ . To see this, fix a point p on γ^+ and consider a distance path ξ from p to γ (Figure 4.2). Cutting along γ, ξ and the loop of length 2ε with base point p yields a hyperbolic quadrilateral. The common orthogonal of γ and the geodesic loop subdivides this quadrilateral into two congruent quadrilaterals, each with three right angles. Applying a standard result from hyperbolic trigonometry yields [20, Formula Glossary 2.3.1(v)]

$$\sinh(\varepsilon) = \sinh\left(\frac{1}{2}\ell(\gamma)\right) \cosh(\ell(\xi)).$$

Because $K_C = \ell(\xi)$, it follows that

$$K_C = \operatorname{arccosh}\left(\frac{\sinh(\varepsilon)}{\sinh(\frac{1}{2}\ell(\gamma))}\right). \quad (4.1)$$

We see that γ^+ consists of points that are equidistant to γ . By symmetry, the distance between a point on γ^- and γ is equal to K_C as well. Moreover, γ^+ and γ^- are smooth.

Recall the notion of a collar from Section 1.1.8. In particular, each hyperbolic cylinder C in $\mathbb{M}_{\text{thin}}^\varepsilon$ is a collar of width K_C , i.e., $C = C_\gamma(K_C)$. Comparing equation (4.1) for K_C with equation (1.7) in the statement of the Collar Lemma, we see that $w(\gamma) > K_C$, because $\sinh \varepsilon < 1$. This inequality will be used in the proof of Lemma 4.5 to give a lower bound for the distance between distinct hyperbolic cylinders in $\mathbb{M}_{\text{thin}}^\varepsilon$.

We distinguish between two kinds of hyperbolic cylinders in $\mathbb{M}_{\text{thin}}^\varepsilon$, namely ε' -thin cylinders and ε' -thick cylinders, where $\varepsilon' = 0.99\varepsilon$. An ε' -thick cylinder with

Figure 4.2: Computing K_C .

waist γ satisfies $2\varepsilon' \leq \ell(\gamma) \leq 2\varepsilon$, since γ is contained in the ε -thin part. An ε' -thin cylinder satisfies $\ell(\gamma) < 2\varepsilon'$.

Lemma 4.12 in Section 4.2.2 states that the triangulation depicted in Figure 4.3 is a Delaunay triangulation for ε' -thin cylinders. We call this triangulation a *standard triangulation* and describe it in more detail in the following definition. For ε' -thick cylinders we use a different construction defined in Definition 4.8.

Definition 4.4. Let $\mathbb{M}$ be a closed hyperbolic surface. Let C be an ε' -thin hyperbolic cylinder in $\mathbb{M}_{\text{thin}}^\varepsilon$ with waist γ and boundary curves γ^+, γ^- . Place three equally-spaced points $x_i, i = 1, 2, 3$ on γ (see Figure 4.3). Then, place three points $x_i^+, i = 1, 2, 3$ on γ^+ and three points $x_i^-, i = 1, 2, 3$ on γ^- such that the projection of $x_i^\pm$ on γ is equal to x_i for $i = 1, 2, 3$. Let V be the set consisting of x_i, x_i^- and x_i^+ for $i = 1, 2, 3$. Let E be the set of edges of one of the forms

$$(x_i^-, x_{i+1}^-), (x_i^-, x_i), (x_i^-, x_{i+1}^+), (x_i, x_{i+1}^+), (x_i, x_i^+), (x_i, x_{i+1}^+), (x_i^+, x_{i+1}^+)$$

for $i = 1, 2, 3$ (counting modulo 3), where the embedding of an edge in C is as shown in Figure 4.3. We call (V, E) a *standard triangulation* of C .

We not only have to prove that a standard triangulation of an ε' -thin cylinder is a Delaunay triangulation, we also have to show that its edges are distance paths. Corollary 4.7 states that all edges in a standard triangulation are distance paths if $\varepsilon \leq 0.72$. Before we can prove Corollary 4.7, we first need the following lemma.

Lemma 4.5. *Let $\mathbb{M}$ be a closed hyperbolic surface and let $\varepsilon \leq 0.72$. For each pair of distinct closed geodesics γ_1 and γ_2 in $\mathbb{M}_{\text{thin}}^\varepsilon$ the collars $C_{\gamma_1}(K_{C_1} + \frac{1}{3}\varepsilon)$ and $C_{\gamma_2}(K_{C_2} + \frac{1}{3}\varepsilon)$ are embedded and disjoint.*

Figure 4.3: Standard triangulation of an ε' -thin cylinder.

Remark 4.6. The value 0.72 was found experimentally and is optimal up to two decimal digits, i.e., the statement is not true for $\varepsilon = 0.73$. More specifically, if $\varepsilon \geq 0.73$ then there exists a closed hyperbolic surface $\mathbb{M}$ with disjoint closed geodesics γ_1 and γ_2 in $\mathbb{M}_{\text{thin}}^\varepsilon$ such that $C_{\gamma_1}(K_{C_1} + \frac{1}{3}\varepsilon)$ and $C_{\gamma_2}(K_{C_2} + \frac{1}{3}\varepsilon)$ are not disjoint.

Proof. See Figure 4.4. We will show that $w(\gamma_i) - K_{C_i} \geq \frac{1}{3}\varepsilon$ for $i = 1, 2$. Namely, this implies that $C_{\gamma_i}(K_{C_i} + \frac{1}{3}\varepsilon) \subseteq C_{\gamma_i}(w(\gamma_i))$. Because $C_{\gamma_1}(w(\gamma_1))$ and $C_{\gamma_2}(w(\gamma_2))$ are embedded and disjoint by the Collar Lemma, it follows that $C_{\gamma_1}(K_{C_1} + \frac{1}{3}\varepsilon)$ and $C_{\gamma_2}(K_{C_2} + \frac{1}{3}\varepsilon)$ are embedded and disjoint as well.

Comparing expression (4.1) for K_{C_i} and expression (1.7) for $w(\gamma_i)$, we see that $w(\gamma_i) - K_{C_i}$ is a positive number, with infimum when $\ell(\gamma_i) \rightarrow 0$ [63]. A straightforward computation shows that for $\varepsilon = 0.72$ this infimum is equal to $0.24\dots > \frac{1}{3}\varepsilon$. Since $w(\gamma_i) - K_{C_i}$ is decreasing as a function of ε , it follows that $w(\gamma_i) - K_{C_i} \geq \frac{1}{3}\varepsilon$ for all $\varepsilon \leq 0.72$. $\square$

Corollary 4.7. Let $\mathbb{M}$ be a closed hyperbolic surface and let $\varepsilon \leq 0.72$. All edges in a standard triangulation of an ε' -thin cylinder in $\mathbb{M}_{\text{thin}}^\varepsilon$ are distance paths.

Proof. It is clear that edges of the form $(x_i^-, x_{i+1}^-), (x_i, x_{i+1}), (x_i^+, x_{i+1}^+)$ for $i = 1, 2, 3$ are distance paths. Now, consider the edge of length K_C between x_i and x_i^+ . Because we know the metric of the cylinder, it can be shown explicitly that there are no shorter paths completely contained in the cylinder. Furthermore,

Figure 4.4: Illustration of the collars $C_{\gamma_i}(K_{C_i}) \subset C_{\gamma_i}(K_{C_i} + \frac{1}{3}\varepsilon) \subseteq C_{\gamma_i}(w(\gamma_i))$.

because the collar $C_{\gamma}(K_C + \frac{1}{3}\varepsilon)$ is embedded by Lemma 4.5, any path that leaves the top half of the cylinder and returns through the bottom half has length at least $K_C + \frac{2}{3}\varepsilon$. It follows that the edges of the form (x_i, x_{i+1}^+) are distance paths. By symmetry, the edges of the form (x_i^-, x_i) are distance paths as well.

Finally, consider the edge between x_i and x_{i+1}^+ . Because $d(x_i, x_{i+1}) = \frac{1}{3}\ell(\gamma) < \frac{2}{3}\varepsilon$ and $d(x_{i+1}, x_{i+1}^+) = K_C$, we see from the triangle inequality that $d(x_i, x_{i+1}^+) < K_C + \frac{2}{3}\varepsilon$. Because any path that leaves the top half of the cylinder and returns through the bottom part of the cylinder has length at least $K_C + \frac{2}{3}\varepsilon$ by the same reasoning as above, it follows that edges of the form (x_i, x_{i+1}^+) are distance paths. By symmetry, edges of the form (x_i^-, x_{i+1}) are distance paths as well. $\square$

For ε' -thick cylinders, we see from Equation (4.1) for K_C that the width K_C is close to zero. It turns out that we do not need to place three points on its waist and on each of its two boundary curves. Instead, three vertices on its waist suffice.

Definition 4.8. Let $\mathbb{M}$ be a closed hyperbolic surface. Let C be a ε' -thick hyperbolic cylinder in $\mathbb{M}_{\text{thin}}^\varepsilon$ with waist γ . Place three equally-spaced points $x_i, i = 1, 2, 3$ on γ . Let $V = \{x_i \mid i = 1, 2, 3\}$ and $E = \{(x_1, x_2), (x_2, x_3), (x_3, x_1)\}$. We call (V, E) a *standard cycle* of C .

4.2.2 Constructing a distance Delaunay triangulation of a hyperbolic surface with few vertices

After constructing sets of vertices in the cylinders in the ε -thin part $\mathbb{M}_{\text{thin}}^\varepsilon$, we construct a sufficiently dense and well-distributed set of vertices in the remainder

of the surface. The following definition shows more precisely how we construct a set of vertices in $\mathbb{M}_{\text{thick}}^\varepsilon$ and a corresponding Delaunay triangulation.

Definition 4.9. Set $\varepsilon = 0.72$ and $\varepsilon' = 0.99\varepsilon$. Let $\mathbb{M}$ be a closed hyperbolic surface. Let $\mathcal{P}_1$ be the set consisting of the vertices of a standard triangulation of every ε' -thin cylinder in $\mathbb{M}_{\text{thin}}^\varepsilon$ together with the vertices of a standard cycle for every ε' -thick cylinder in $\mathbb{M}_{\text{thin}}^\varepsilon$. Let T_j be the union of triangles in a standard triangulation (V_j, E_j) of an ε' -thin cylinder C_j . For every ε' -thick cylinder C_j , set $T_j = \emptyset$. Define $\mathcal{P}_2$ to be a maximal set in $\mathbb{M} \setminus \cup_j T_j$ such that $d(p, q) \geq \frac{1}{2}\varepsilon$ for all distinct $p \in \mathcal{P}_1 \cup \mathcal{P}_2, q \in \mathcal{P}_2$. Denote the union $\mathcal{P}_1 \cup \mathcal{P}_2$ by $\mathcal{P}$ and let T be the Delaunay triangulation of $\mathcal{P}$ on $\mathbb{M}$ obtained after projecting a Delaunay triangulation of $\pi^{-1}(\mathcal{P})$ in $\mathbb{D}$ to $\mathbb{M}$. We call T a *thick-thin Delaunay triangulation* of $\mathbb{M}$. The vertices in $\mathcal{P}_1$ and $\mathcal{P}_2$ are called the *cylinder vertices* and *non-cylinder vertices* of T , respectively.

Remark 4.10. Because by Corollary 4.7 all edges in a standard triangulation of any ε' -thin cylinder are distance paths if we choose $\varepsilon \leq 0.72$, we have chosen $\varepsilon = 0.72$ in Definition 4.9. Namely, we will see in the proof of Theorem 4.2 that the larger we choose ε , the smaller the constant (in our case 151) in the upper bound for the number of vertices. As in Section 4.2.1 we will fix $\varepsilon = 0.72$ and $\varepsilon' = 0.99\varepsilon$ throughout this subsection.

The edges between vertices on the same boundary curve of C_j are not equal to the boundary curves of C_j (because the latter are not geodesics), so T_j is strictly contained in C_j . We define $\mathcal{P}_2$ as a point set in $\mathbb{M} \setminus \cup_j T_j$ instead of in $\mathbb{M} \setminus \cup_j C_j$ to simplify our proof of Lemma 4.15, where we show that a thick-thin Delaunay triangulation of $\mathbb{M}$ is a simplicial complex.

The definition of $\mathcal{P}$ does not explicitly forbid placing vertices of $\mathcal{P}_2$ in ε' -thick cylinders. However, we will see in the next lemma that there are no vertices of $\mathcal{P}_2$ in ε' -thick cylinders, because then they would be too close to the vertices of a standard cycle.

Lemma 4.11. *Let $\mathbb{M}$ be a closed hyperbolic surface and let T be a thick-thin Delaunay triangulation of $\mathbb{M}$. Every vertex of T contained in an ε' -thick cylinder in $\mathbb{M}_{\text{thin}}^\varepsilon$ is a cylinder vertex.*

Proof. Let $\mathcal{P}_1$ be the set of cylinder vertices and $\mathcal{P}_2$ the set of non-cylinder vertices. Let C be an arbitrary ε' -thick cylinder with waist γ and standard cycle (V, E) . We will show that the union U of the disks of radius $\frac{1}{2}\varepsilon$ centered at the vertices of V covers C completely. Namely, this implies that every point of C has distance at most $\frac{1}{2}\varepsilon$ to a vertex of V . Because $d(p, q) \geq \frac{1}{2}\varepsilon$ for all $p \in \mathcal{P}_1$ and $q \in \mathcal{P}_2$, it follows that there are no vertices of $\mathcal{P}_2$ contained in C .

To prove that U covers C completely, first observe that $d(x_i, x_{i+1}) = \frac{1}{3}\ell(\gamma) < \frac{2}{3}\varepsilon$ for all $i = 1, 2, 3$ (counting modulo 3). Therefore, the circles of radius $\frac{1}{2}\varepsilon$ centered at x_i and x_{i+1} intersect in two points, of which we call one p . Since the

collar $C_\gamma(d(\gamma, p))$ is contained in U , it suffices to show that $K_C < d(\gamma, p)$, because then $C = C_\gamma(K_C) \subset C_\gamma(d(\gamma, p)) \subset U$. From equation (4.1) for K_C we know that

$$\cosh(K_C) = \frac{\sinh(\varepsilon)}{\sinh(\frac{1}{2}\ell(\gamma))} \leq \frac{\sinh(\varepsilon)}{\sinh(\varepsilon')} \leq 1.02,$$

where we substituted $\varepsilon' = 0.99\varepsilon$ and $\varepsilon = 0.72$ in the last step. On the other hand, the hyperbolic Pythagorean theorem yields

$$\cosh(d(\gamma, p)) = \frac{\cosh(\frac{1}{2}\varepsilon)}{\cosh(\frac{1}{6}\ell(\gamma))} \geq \frac{\cosh(\frac{1}{2}\varepsilon)}{\cosh(\frac{1}{3}\varepsilon)} \geq 1.03,$$

(see Figure 4.5) where again we substituted $\varepsilon = 0.72$ in the last step. We conclude that $K_C < d(\gamma, p)$, which finishes the proof.

Figure 4.5: Computing $d(\gamma, p)$.

□

Even though the set of vertices of a thick-thin Delaunay triangulation of $\mathbb{M}$ contains the vertices of a standard triangulation (V_j, E_j) for every ε' -thin cylinder C_j , a priori it is not clear that the edges in E_j are edges in T as well. In the next lemma, we will show that for every ε' -thin cylinder the triangles in a standard triangulation are Delaunay triangles with respect to the set of vertices of any thick-thin Delaunay triangulation of $\mathbb{M}$. Namely, if this holds, then there exists a Delaunay triangulation of $\mathcal{P}$ on $\mathbb{M}$ containing a standard triangulation of every ε' -thin cylinder in $\mathbb{M}_{\text{thin}}^\varepsilon$.

Lemma 4.12. *Let $\mathbb{M}$ be a closed hyperbolic surface. Let T be a thick-thin Delaunay triangulation of $\mathbb{M}$ with vertex set $\mathcal{P}$ and let C be an ε' -thin cylinder in $\mathbb{M}_{\text{thin}}^\varepsilon$ with waist γ . Let (V, E) be a standard triangulation of C such that $V \subset \mathcal{P}$. Then all triangles of (V, E) are Delaunay triangles with respect to the point set $\mathcal{P}$.*

Proof. To prove that the triangles of (V, E) are Delaunay triangles, we will show that every circumscribed disk does not contain any point of $\mathcal{P}$ in its interior. By symmetry, it is sufficient to consider the top half of the cylinder. Let $i = 1, 2, 3$ be arbitrary and denote the disk passing through $x_i^+, x_{i+1}^+, x_i, x_{i+1}$ by D_i . That D_i does not contain any $p \in V$ in its interior is clear. The remainder of the proof

consists of showing that p is not contained in the interior of D_i for all $p \in \mathcal{P} \setminus V$. Take $p \in \mathcal{P} \setminus V$ arbitrarily. Let c_i be the center of D_i . If $d(c_i, p) > d(c_i, x_i)$, then p is not contained in the interior of D_i .

Observe that $d(p, x_i^\pm) \geq \frac{1}{2}\varepsilon$ for $i = 1, 2, 3$. Namely, if $p \in \mathcal{P}_2$, where $\mathcal{P}_2$ is the subset of $\mathcal{P}$ constructed in $\mathbb{M}_{\text{thick}}^\varepsilon$, then by definition $d(p, x_i^\pm) \geq \frac{1}{2}\varepsilon$ for $i = 1, 2, 3$. On the other hand, if $p \in \mathcal{P}_1$, then p is a vertex in some hyperbolic cylinder $C' \neq C$ with waist γ' in $\mathbb{M}_{\text{thin}}^\varepsilon$. By Lemma 4.5, the collars $C_\gamma(K_C + \frac{1}{3}\varepsilon)$ and $C_{\gamma'}(K_{C'} + \frac{1}{3}\varepsilon)$ are disjoint, so the distance between C and C' is at least $\frac{2}{3}\varepsilon$. Hence, $d(p, x_i) \geq \frac{1}{2}\varepsilon$ for $i = 1, 2, 3$.

Figure 4.6: Construction to show that we can assume without loss of generality that $d(x_i^+, p) = d(x_{i+1}^+, p) = \frac{1}{2}\varepsilon$ in the proof of Lemma 4.12.

Now, we claim that we can assume without loss of generality that $d(x_i^+, p) = d(x_{i+1}^+, p) = \frac{1}{2}\varepsilon$. Let m_i be the midpoint of x_i and x_{i+1} and let m_i^+ be the midpoint of x_i^+ and x_{i+1}^+ , as in Figure 4.6. Consider the curve γ_c consisting of points of distance $d(c_i, \gamma)$ from γ and let p_c be the point of γ_c closest to p . Because $c_i \in \gamma_c$, we know that $d(p, c_i) \geq d(p, p_c)$. Now, note that a standard result from hyperbolic trigonometry in the quadrilateral (x_i, x_i^+, m_i^+, m_i) with three right angles [20, Formula Glossary 2.3.1(v)] states that

$$\sinh\left(\frac{1}{2}d(x_i^+, x_{i+1}^+)\right) = \sinh\left(\frac{1}{6}\ell(\gamma)\right) \cosh(K_C) = \frac{\sinh\left(\frac{1}{6}\ell(\gamma)\right) \sinh(\varepsilon)}{\sinh\left(\frac{1}{2}\ell(\gamma)\right)},$$

where the last equality follows from expression (4.1) for K_C . It can be deduced that $d(x_i^+, x_{i+1}^+) < \varepsilon$. Because $d(x_i^+, x_{i+1}^+) < \varepsilon$, the circles of radius $\frac{1}{2}\varepsilon$ centered

Equation (4.4) follows from applying the hyperbolic Pythagorean theorem in triangle (x_i^+, p, m_i^+) . Moreover, applying the hyperbolic Pythagorean theorem in triangle (x_i, c_i, m_i) yields

$$\begin{aligned} \cosh(d(c_i, x_i)) &= \cosh\left(\frac{1}{6}\ell(\gamma)\right) \cosh(d(c, m_i)), \\ &= \cosh\left(\frac{1}{6}\ell(\gamma)\right) \cosh\left(\operatorname{arctanh}\left(\frac{\cosh(\frac{1}{6}\ell(\gamma))}{\coth(\frac{1}{2}K_C)}\right)\right), \end{aligned} \quad (4.5)$$

where we used equation (4.3) in the second line.

When we substitute the expressions for K_C and $d(x_i^+, x_{i+1}^+)$ into equations (4.2), (4.3), (4.4) and (4.5), we find expressions for $d(m_i, m_i^+)$, $d(m_i, c_i)$, $d(m_i^+, p)$ and $d(c_i, x_i)$ in terms of ε and $\ell(\gamma)$. As $\varepsilon = 0.72$ is fixed, we can treat these as functions of $\ell(\gamma)$. By a straightforward (but tedious) computation, it can be shown that $d(m_i, m_i^+) - d(m_i, c_i) - d(c_i, x_i)$ is strictly decreasing as a function of $\ell(\gamma)$ with minimum $-0.180\dots$ for $\ell(\gamma) = 2\varepsilon'$. By a similar computation, $d(m_i^+, p)$ is strictly increasing as a function of $\ell(\gamma)$ with minimum $0.247\dots$ for $\ell(\gamma) \rightarrow 0$. We conclude that

$$d(m_i, m_i^+) - d(m_i, c_i) - d(c_i, x_i) + d(m_i^+, p) \geq -0.180\dots + 0.247\dots > 0,$$

from which it follows that $d(c_i, p) = d(m_i, m_i^+) - d(m_i, c_i) + d(m_i^+, p) > d(c_i, x_i)$. Hence, p is not contained in D_i . This finishes the proof. $\square$

Remark 4.13. We note that in the proof it is shown as an intermediate step that $d(x_i^\pm, x_{i+1}^\pm) < \varepsilon$ for all $i = 1, 2, 3$. This inequality is used once more in the proof of Lemma 4.14.

Henceforth, we will assume that for each ε' -thin cylinder the vertices and edges of a standard triangulation are contained in a thick-thin Delaunay triangulation of $\mathbb{M}$. To show that $T \in \mathcal{D}(\mathbb{M})$, we must show that T is a simplicial complex, i.e. it does not contain any 1- or 2-cycles, and that its edges are distance paths.

In the next lemma, we show that any edge that intersects $\mathbb{M}_{\text{thick}}^\varepsilon$ has length smaller than ε . Moreover, we show that it follows that all edges that intersect $\mathbb{M}_{\text{thick}}^\varepsilon$ are distance paths and that there are no 1- and 2-cycles consisting of edges intersecting $\mathbb{M}_{\text{thick}}^\varepsilon$.

Lemma 4.14. *Let $\mathbb{M}$ be a closed hyperbolic surface and let T be a thick-thin Delaunay triangulation of $\mathbb{M}$. Any edge of T that intersects $\mathbb{M}_{\text{thick}}^\varepsilon$ has length smaller than ε and is a distance path. Moreover, there are no 1- or 2-cycles that intersect $\mathbb{M}_{\text{thick}}^\varepsilon$ and consist of edges of length smaller than ε .*

Proof. Let (u, v) be an edge of T with non-empty intersection with $\mathbb{M}_{\text{thick}}^\varepsilon$. Assume that (u, v) is contained in a triangle (u, v, w) in T with circumradius r and circumcenter c . We will first show that $\ell(u, v) < \varepsilon$. We consider two cases, depending on which set c is contained in. First, if $c \in T_j$ for some ε' -thin cylinder C_j , then at

least one of u and v is contained in $\mathcal{P}_1$. If both are contained in $\mathcal{P}_1$, then (u, v) is contained in T_j , because the edges of a standard triangulation of an ε' -thin cylinder are distance paths by Corollary 4.7. This contradicts $(u, v) \cap \mathbb{M}_{\text{thick}}^\varepsilon \neq \emptyset$, so we can assume that only one of u and v , say v , is contained in $\mathcal{P}_1$. Then without loss of generality the situation is as depicted in Figure 4.8, where $\{v, w\} = \{x_i^+, x_{i+1}^+\}$.

Figure 4.8: Circumscribed disk of a triangle (u, x_i^+, x_{i+1}^+) .

The distance between either x_i^+ or x_{i+1}^+ and a point on the shortest arc of the circumscribed circle between x_i^+ and x_{i+1}^+ is less than the distance between x_i^+ and x_{i+1}^+ . In particular, $\ell(u, v) < d(x_i^+, x_{i+1}^+)$. Because $d(x_i^+, x_{i+1}^+) < \varepsilon$ by Remark 4.13, it follows that $\ell(u, v) < \varepsilon$. Second, if $c \in \mathbb{M} \setminus \cup_{j \in I} T_j$, then we can deduce that $r < \frac{1}{2}\varepsilon$. Namely, if we suppose for a contradiction that $r \geq \frac{1}{2}\varepsilon$, then $d(c, p) \geq \frac{1}{2}\varepsilon$ for all $p \in \mathcal{P}$, because the circumcircle of (u, v, w) is empty. Then we could add c to $\mathcal{P}_2$, which contradicts its maximality. We conclude that $r < \frac{1}{2}\varepsilon$. Because (u, v) is contained in a circle of radius $r < \frac{1}{2}\varepsilon$, it follows that $\ell(u, v) < \varepsilon$. Because $\ell(u, v) < \varepsilon$ in both cases, the first claim of the lemma follows.

To show that (u, v) is a shortest distant path between its endpoints, suppose for a contradiction that it is not. Then there exists a geodesic γ from u to v , such that $\ell(\gamma) < \ell(u, v)$. This means that $(u, v) \cup \gamma$ is a homotopically non-trivial closed curve of length smaller than $2\ell(u, v) < 2\varepsilon$. However, because $\text{inrad}(x) > \varepsilon$ for all $x \in \mathbb{M}_{\text{thick}}^\varepsilon$, every homotopically non-trivial closed curve γ intersecting $\mathbb{M}_{\text{thick}}^\varepsilon$ has length at least 2ε , which contradicts $\ell((u, v) \cup \gamma) < 2\varepsilon$. We conclude that (u, v) is a distance path between its endpoints.

A 1- or 2-cycle in T corresponds to a homotopically non-trivial closed curve on $\mathbb{M}$ [16]. By the same argument as before, the length of a 1- or 2-cycle σ intersecting $\mathbb{M}_{\text{thick}}^\varepsilon$ is at least 2ε . Therefore, there are no 1- or 2-cycles that intersect $\mathbb{M}_{\text{thick}}^\varepsilon$ and consist of edges of length smaller than ε . $\square$

Using the previous lemma, we show that a thick-thin Delaunay triangulation of $\mathbb{M}$ is a distance Delaunay triangulation.

Lemma 4.15. *Every thick-thin Delaunay triangulation of a closed hyperbolic surface is a distance Delaunay triangulation.*

Proof. Let $\mathbb{M}$ be a closed hyperbolic surface and let T be a thick-thin Delaunay triangulation of $\mathbb{M}$. By definition, T is a Delaunay triangulation. We will show that T does not contain any 1- or 2-cycles to prove that it is a simplicial complex. We know from Lemma 4.14 that any edge (u, v) such that $(u, v) \cap \mathbb{M}_{\text{thick}}^\varepsilon \neq \emptyset$ is not a 1-cycle. Because by construction there are no 1-cycles in a standard triangulation or standard cycle in $\mathbb{M}_{\text{thin}}^\varepsilon$ as well, we conclude that T contains no 1-cycles.

To prove that T does not contain any 2-cycles, consider two distinct edges (u, v) and (v, w) of T with at least one shared endpoint. There are three cases, depending on whether two, one or zero of the edges (u, v) and (v, w) intersect $\mathbb{M}_{\text{thick}}^\varepsilon$.

First, if (u, v) and (v, w) both intersect $\mathbb{M}_{\text{thick}}^\varepsilon$, then they do not form a 2-cycle by Lemma 4.14.

Second, if precisely one of (u, v) and (v, w) , say (u, v) , intersects $\mathbb{M}_{\text{thick}}^\varepsilon$, then $\ell(u, v) < \varepsilon$ and (v, w) is an edge contained in a hyperbolic cylinder of $\mathbb{M}_{\text{thin}}^\varepsilon$. If (v, w) is an edge contained in an ε' -thick cylinder C with waist γ , then (v, w) is one of the edges of the standard cycle of C , because there are no other vertices in C by Lemma 4.11. Then $\ell(v, w) = \frac{1}{3}\ell(\gamma) < \frac{2}{3}\varepsilon$, so (u, v) and (v, w) do not form a 2-cycle by Lemma 4.14. Next, assume that (v, w) is an edge in an ε' -thin cylinder with waist γ . Then either w lies on γ and v lies on one of the boundary curves of C or v and w both lie on the same boundary curve of C . If w lies on γ and v on a boundary curve of C , then (u, v) and (v, w) do not form a 2-cycle, because u does not lie on γ . If v and w both lie on the same boundary geodesic, then $\ell(v, w) < \varepsilon$ by Remark 4.13, so (u, v) and (v, w) do not form a 2-cycle by Lemma 4.14.

Third, if neither (u, v) nor (v, w) intersects $\mathbb{M}_{\text{thick}}^\varepsilon$, then (u, v) and (v, w) are both contained in a hyperbolic cylinder in $\mathbb{M}_{\text{thin}}^\varepsilon$. They are contained in the same cylinder, because different cylinders are separated by $\mathbb{M}_{\text{thick}}^\varepsilon$. Because by construction standard triangulations and standard cycles do not contain any 2-cycle, (u, v) and (v, w) do not form a 2-cycle.

This finishes the case analysis and we conclude that T is a simplicial complex.

To prove that all edges of T are distance paths, we know from Lemma 4.14 that any edge that intersects $\mathbb{M}_{\text{thick}}^\varepsilon$ is a distance path. Because all edges in a standard triangulation are distance paths by Corollary 4.7 and because all edges in a standard cycle are distance paths by construction, we conclude that all edges in T are distance paths. $\square$

4.2.3 Proof of Theorem 4.2

Proof. (Theorem 4.2)

Let $\mathbb{M}$ be an arbitrary hyperbolic surface of genus g and let T be a thick-thin Delaunay triangulation of $\mathbb{M}$. By definition, T is a Delaunay triangulation. By Lemma 4.15, T is a simplicial complex and all edges of T are distance paths. Hence, $T \in \mathcal{D}(\mathbb{M})$.

We will show here that the number of vertices of T is smaller than $151g$. By Lemma 4.3, $\mathbb{M}_{\text{thin}}^\varepsilon$ consists of at most $3g - 3$ cylinders and each of these cylinders contains either 9 vertices (if it is ε' -thin) or 3 vertices (if it is ε' -thick). Therefore, $|\mathcal{P}_1| \leq 27g - 27$.

To find an upper bound for the cardinality of $\mathcal{P}_2$, observe that for distinct $p, q \in \mathcal{P}_2$ the disks $B_p(\frac{1}{4}\varepsilon)$ and $B_q(\frac{1}{4}\varepsilon)$ of radius $\frac{1}{4}\varepsilon$ centered at p and q , respectively, are embedded and disjoint. Therefore, the cardinality of $\mathcal{P}_2$ is bounded above by the number of disjoint, embedded disks of radius $\frac{1}{4}\varepsilon$ that we can fit in $\mathbb{M}$. Because the area of a hyperbolic disk of radius $\frac{1}{4}\varepsilon$ is $2\pi(\cosh(\frac{1}{4}\varepsilon) - 1)$ [10] and because the area of $\mathbb{M}$ is $4\pi(g - 1)$ [70], we obtain

$$|\mathcal{P}_2| \leq \frac{2(g - 1)}{\cosh(\frac{1}{4}\varepsilon) - 1}.$$

Therefore,

$$|\mathcal{P}| \leq 27g - 27 + \frac{2(g - 1)}{\cosh(\frac{1}{4}\varepsilon) - 1} \leq 151g.$$

This finishes the proof. $\square$

Remark 4.16. The constant 151 is not optimal. We can obtain the stronger upper bound $|\mathcal{P}| \leq 124g$ by looking more precisely at the upper bounds of $|\mathcal{P}_1|$ and $|\mathcal{P}_2|$ but because we are mainly interested in the the order of growth, we will not provide any details.

4.3 Classes of hyperbolic surfaces attaining the order of the upper bound

As our second result, we show that there exists a class of hyperbolic surfaces which attains the order of the upper bound presented in Theorem 4.2. We will first introduce this class of hyperbolic surfaces and then state the precise result in Theorem 4.17.

Recall from the Preliminaries that cutting a hyperbolic surface along $3g - 3$ disjoint simple closed geodesics decomposes the surface into $2g - 2$ pairs of pants and that each pair of pants decomposition has an associated 3-regular graph. Conversely, define L_g as the trivalent graph depicted in Figure 4.9. Here, every vertex v_i corresponds to a pair of pants Y_i . There is one edge from v_1 to itself and similarly from v_{2g-2} to itself. Moreover, for $1 \leq i \leq 2g - 3$ there is one edge between v_i and v_{i+1} if i is odd and there are two edges if i is even.

Now, fix some interval $[a, b] \subset \mathbb{R}$ with $0 < a < b$. Let $S_g(a, b)$ be the subset of $\mathcal{T}_g$ with underlying graph L_g such that all length parameters are contained in $[a, b]$. In particular, $S_g(a, b)$ contains an open subset of $\mathcal{T}_g$, showing that having a linear number of vertices in terms of genus is relatively stable in this part of Teichmüller space. We will now state the result of this section.

Figure 4.9: Trivalent graph L_g with corresponding pair of pants decomposition.

Theorem 4.17. *There exists a constant $B > 0$ depending only on a, b such that a minimal distance Delaunay triangulation of any hyperbolic surface in $S_g(a, b)$ has at least Bg vertices.*

The idea of the proof is the following. Let a hyperbolic surface $\mathbb{M} \in S_g(a, b)$ and a triangulation $T \in \mathcal{D}(\mathbb{M})$ be given. Euler's formula implies $v - \frac{1}{3}e = 2 - 2g$ for triangulations of a surface of genus g , where v and e are the number of vertices and edges of the triangulation. We prove that $e \leq B'v$ for some constant $B' > 3$ only depending on a, b , which implies that

$$v \geq \frac{6g - 6}{B' - 3}.$$

This implies the result of Theorem 4.17. Hence, the argument consists mostly in finding an upper bound for the number of edges in terms of the number of vertices.

Before we continue with the proof of Theorem 4.17, we will look at our construction of $S_g(a, b)$ in more detail. By definition, every boundary geodesic of a pair of pants in the pair of pants decomposition of $\mathbb{M} \in S_g(a, b)$ with respect to L_g has length in $[a, b]$. As explained in Section 1.1.9, the geometry of a pair of pants depends continuously on the lengths of its three boundary geodesics. In particular, the diameter $\text{diam}(Y)$ of a pair of pants Y as well as the minimal distance $\text{mindist}(Y)$ between any two of its boundary geodesics depend continuously on the lengths of its boundary geodesics. Because $[a, b]$ is a compact set, we obtain as an immediate consequence the following lemma.

Lemma 4.18. *There exist positive numbers $m(a, b)$ and $M(a, b)$ depending on a and b such that $m(a, b) \leq \text{mindist}(Y) < \text{diam}(Y) \leq M(a, b)$ for every pair of pants Y whose boundary geodesics have length in $[a, b]$.*

Remark 4.19. It is not too difficult to compute bounds for $\text{mindist}(Y)$ and $\text{diam}(Y)$ in terms of the lengths of the boundary geodesics of Y . This would give explicit expressions for $m(a, b)$ and $M(a, b)$ in terms of a and b . As we are only

interested in the order of growth, to avoid further technical details, we do not provide details.

From now on, the numbers $m = m(a, b)$ and $M = M(a, b)$ will be fixed. Furthermore, a *cluster* in a hyperbolic surface $\mathbb{M}$ is a subset of $\mathbb{M}$ consisting of a number of consecutive pairs of pants, where *consecutive* is with respect to the ordering of L_g . Consider $T \in \mathcal{D}(\mathbb{M})$. A *k-gap* is a cluster consisting of k consecutive empty pairs of pants, where *empty* means that the pairs of pants do not contain any vertices of T . If a vertex of T is contained in two pairs of pants, i.e., if the vertex lies on a boundary geodesic, then we only count it as a vertex of the pair of pants with the lowest index in L_g . We say that an edge of T *crosses* a cluster if the pairs of pants containing its endpoints are separated by the cluster. Note that the cluster need not contain all pairs of pants which separate the two endpoints.

The next lemma states that if an edge of a distance Delaunay triangulation crosses many pairs of pants, then “many” of these pairs of pants are empty.

Lemma 4.20. *Let $\mathbb{M} \in S_g(a, b)$ and define $N = N(a, b)$ as*

$$N(a, b) := \left\lceil \frac{M(a, b)}{m(a, b)} \right\rceil + 1.$$

Then, for every $T \in \mathcal{D}(\mathbb{M})$ the following statements hold:

1. *If an edge of T crosses a cluster consisting of at least $3N$ pairs of pants, this cluster contains an N -gap.*
2. *If an edge of T crosses a cluster in which the first N and the last N pairs of pants are empty, then all pairs of pants in the cluster are empty.*

Proof. Let $T \in \mathcal{D}(\mathbb{M})$ and let (u, v) be an edge of T with $u \in Y_i$ and $v \in Y_j$. We will show that the cluster consisting of the union of all Y_k with $i + N + 1 \leq k \leq j - N - 1$ is empty. In other words, only the first N and last N pairs of pants are possibly non-empty. In particular, this implies that the two properties of the lemma hold.

Now, because (u, v) is a Delaunay edge, it is contained in some empty disk D passing through u and v . Consider Y_k with $i + N + 1 \leq k \leq j - N - 1$ and take $p \in Y_k$ arbitrarily (see Figure 4.10). We will show that the distance between p and the center c of D satisfies $\text{dist}(p, c) < \text{dist}(u, c)$. This implies that p is contained in the interior of D , so it cannot be a vertex of T . Therefore, Y_k is empty.

Let γ_u be a distance path from c to u . First assume that $\gamma_u \cap Y_k \neq \emptyset$. Let x_k be the intersection of γ_u with one of the boundary geodesics of Y_k . By the triangle inequality and the definition of M , we know that

$$\text{dist}(c, p) \leq \text{dist}(c, x_k) + \text{dist}(x_k, p) \leq \text{dist}(c, x_k) + M.$$

To give an upper bound for $\text{dist}(c, x_k)$, observe that the part of γ_u from x_k to u passes through $Y_{k-1}, \dots, Y_{i+1}$ before reaching $u \in Y_i$. By definition of m , the

Figure 4.10: Depiction of the construction in the proof of Lemma 4.20.

length of the part of γ_u within each of these $k - 1 - i$ pair of pants is at least m , so $\text{dist}(x_k, u) \geq (k - 1 - i)m$. This means that

$$\text{dist}(c, x_k) = \text{dist}(c, u) - \text{dist}(x_k, u) \leq \text{dist}(c, u) - (k - 1 - i)m.$$

It follows that

$$\text{dist}(c, p) \leq \text{dist}(c, u) + M - (k - 1 - i)m.$$

Because $k - 1 - i \geq N$ and $Nm > M$, we see that $M - (k - 1 - i)m < 0$, so

$$\text{dist}(c, p) < \text{dist}(c, u),$$

so $Y_k \subseteq D$, hence Y_k does not contain any vertices of T . Note that we assumed that $\gamma_u \cap Y_k \neq \emptyset$. If $\gamma_u \cap Y_k = \emptyset$, then we consider a distance path γ_v from c to v instead of γ_u . The rest of the proof is analogous. We conclude that Y_k is empty for $i + N + 1 \leq k \leq j - N - 1$. This finishes the proof. $\square$

The following lemma states that we can construct a set of clusters which has as one of its properties that every edge of the distance Delaunay triangulation has its endpoints in the same cluster, or in two consecutive clusters.

Lemma 4.21. *Let $\mathbb{M} \in S_g(a, b)$ be a hyperbolic surface and let $N = N(a, b)$ be as defined in Lemma 4.20. Let $T \in \mathcal{D}(\mathbb{M})$. There are interior-disjoint clusters with the following properties:*

1. *Each cluster consists of at most $6N$ consecutive pairs of pants;*
2. *Every cluster contains at least one vertex of $\mathcal{T}$, and every vertex of $\mathcal{T}$ belongs to exactly one cluster;*
3. *Every edge of $\mathcal{T}$ has its endpoints in the same cluster, or in two consecutive clusters.*

Proof. A wide gap in the sequence of $2g - 2$ pairs of pants is a maximal subsequence consisting of at least N consecutive empty pairs of pants, where maximality is defined with respect to inclusion. The complement of the collection of wide gaps of $\mathbb{M}$ consists of a number of sequences of consecutive pairs of pants, that we call

superclusters (see Figure 4.11). To obtain clusters that satisfy the properties of the lemma, each supercluster will be chopped up into one or more subsequences of length at most $6N$ in the following way:

- Each supercluster consisting of at most $3N$ pairs of pants is a cluster. Such a cluster is said to be a *short cluster*.
- Each supercluster consisting of more than $3N$ pairs of pants is chopped up into non-overlapping subsequences of length $3N$, followed by a subsequence of length between $3N$ and $6N - 1$.

More precisely, let l be the length of such a supercluster. Since $l > 3N$ there are integers k and r , with $0 \leq r < 3N$ and $k \geq 1$, such that $l = 3kN + r$. Define the integers $l_1, \dots, l_k$ by setting $l_1 = \dots = l_{k-1} = 3N$ and $l_k = 3N + r$. Then $l_1 + \dots + l_k = l$ and $3N \leq l_i < 6N$. Therefore, the supercluster is the concatenation of k subsequences of length $l_1, \dots, l_k$.

This construction enforces Property 1.

Figure 4.11: Illustration of the construction in the proof of Lemma 4.21. Each square represents a pair of pants in the pair of pants decomposition of a hyperbolic surface of genus 17. A pair of pants corresponding to a black square contains vertices of T , whereas a pair of pants corresponding to a white square does not contain vertices of T . We assume that $N = 2$. The superclusters and clusters that are defined during the construction are indicated.

Proof of Property 2. Each supercluster contains at least one vertex. Therefore, short clusters contain at least one vertex. The clusters obtained by chopping up a supercluster have length at least $3N$. Since a supercluster contains no N -gap, each of these clusters contains at least one vertex.

Since all vertices belong to some supercluster, they belong to at least one cluster. Since the clusters are interior disjoint, every vertex belongs to exactly one cluster. We note again that vertices on a boundary geodesic of some pair of pants are considered to belong to the cluster containing the pair of pants with the lowest index in L_g . This completes the proof of Property 2.

Proof of Property 3. Suppose the property does not hold. Then there is an edge of T with vertices in non-adjacent clusters. In other words, this edge crosses another cluster, say C . The construction of clusters implies that C does not contain an N -gap. By Part 1 of Lemma 4.20 the cluster C consists of less than $3N$ pairs of pants.

Therefore, C is a short cluster. Since this short cluster is neither the first nor the last in the sequence of superclusters, it is preceded by a wide gap and succeeded by a wide gap. Since a wide gap contains an N -gap, Part 2 of Lemma 4.20 implies that cluster C is empty. This contradicts Property 2. Therefore, Property 3 holds. $\square$

In the following corollary, we denote the number of vertices of $T \in \mathcal{D}(\mathbb{M})$ contained in a subset U of $\mathbb{M}$ by $v(U)$. Likewise, let $e(U, W)$ be the number of edges of T with one endpoint in $U \subset \mathbb{M}$ and one endpoint in $W \subset \mathbb{M}$.

Corollary 4.22. *Let $\mathbb{M} \in S_g(a, b)$ be a hyperbolic surface and let $T \in \mathcal{D}(\mathbb{M})$. Let $\{\Gamma_i \mid i = 1, \dots, n\}$ be a collection of clusters satisfying the properties of Lemma 4.21 for some $n \in \mathbb{N}$. If v and e are the number of vertices and edges of T , respectively, then*

$$\begin{aligned} n &\leq v, \\ v &= \sum_{i=1}^n v(\Gamma_i), \\ e &= \sum_{i=1}^n e(\Gamma_i, \Gamma_i) + \sum_{i=1}^{n-1} e(\Gamma_i, \Gamma_{i+1}). \end{aligned}$$

Proof. Because every cluster contains at least one vertex, the number of clusters is at most the number of vertices, which proves the first equation. The second equation follows from the property that every vertex is contained in a cluster. Because every edge has its endpoints in the same cluster, or in two consecutive clusters, the third equation holds. $\square$

Recall that we want to find a linear upper bound for the number of edges of a distance Delaunay triangulation in terms of the number of vertices. By Corollary 4.22, it suffices to find upper bounds for $e(\Gamma_i, \Gamma_i)$ and $e(\Gamma_i, \Gamma_{i+1})$ for clusters Γ_i satisfying the properties of Lemma 4.21. We will do this in the next lemma. The proof is given in Appendix A.3.

Lemma 4.23. *Let the notation be as in Corollary 4.22. Then, the following upper bounds hold:*

1. $e(\Gamma_i, \Gamma_i) \leq 3v(\Gamma_i) + 18N(N+1)$ for all $i = 1, \dots, n$,
2. $e(\Gamma_i, \Gamma_{i+1}) \leq 18v(\Gamma_i \cup \Gamma_{i+1}) + 216N(N+1)$ for all $i = 1, \dots, n-1$.

We now continue with the proof of Theorem 4.17.

Proof. (Theorem 4.17)

Take $\mathbb{M} \in S_g(a, b)$ arbitrary and let $T \in \mathcal{D}(\mathbb{M})$ be arbitrary. Let $\{\Gamma_i \mid i =$

$1, \dots, n\}$ be a collection of clusters satisfying the properties of Lemma 4.21. By Corollary 4.22,

$$e = \sum_{i=1}^n e(\Gamma_i, \Gamma_i) + \sum_{i=1}^{n-1} e(\Gamma_i, \Gamma_{i+1}).$$

Substituting the upper bounds for $e(\Gamma_i, \Gamma_i)$ and $e(\Gamma_i, \Gamma_{i+1})$ from Lemma 4.23, we obtain

$$\begin{aligned} e &\leq \sum_{i=1}^n (3v(\Gamma_i) + 18N(N+1)) + \sum_{i=1}^{n-1} (18v(\Gamma_i \cup \Gamma_{i+1}) + 216N(N+1)), \\ &\leq 39 \sum_{i=1}^n (v(\Gamma_i) + 6N(N+1)). \end{aligned}$$

From Corollary 4.22, we know that $\sum_{i=1}^n v(\Gamma_i) = v$ and $n \leq v$. Hence,

$$e \leq 39(1 + 6N(N+1))v.$$

Euler's formula for triangulations $v - \frac{1}{3}e = 2 - 2g$ implies that

$$\begin{aligned} 39(1 + 6N(N+1))v &\geq e = 3v + 6g - 6, \\ v &\geq \frac{g-1}{6 + 39N(N+1)}, \end{aligned}$$

which finishes the proof. $\square$

4.4 Lower bound

In this section, we will look at a general lower bound for the minimal number of vertices of a distance Delaunay triangulation of a hyperbolic surface of genus g .

In the more general situation of a simplicial triangulation of a topological surface of genus g , one has an immediate lower bound on the minimal number of vertices. The fact that this lower bound is sharp is the following classical theorem of Jungerman and Ringel:

Theorem 4.24. [48, Theorem 1.1] *The minimal number of vertices of a simplicial triangulation of a topological surface of genus g is*

$$\left\lceil \frac{7 + \sqrt{1 + 48g}}{2} \right\rceil.$$

We show that the same result holds for the minimal number of vertices of a distance Delaunay triangulation of a hyperbolic surface of genus g for infinitely many values of g .

Theorem 4.25. *For any $g \geq 2$ of the form*

$$g = \frac{(n-3)(n-4)}{12}$$

for some $n \equiv 0 \pmod{12}$, the minimal number of vertices of a distance Delaunay triangulation of a hyperbolic surface of genus g is

$$n = \frac{7 + \sqrt{1 + 48g}}{2}.$$

Proof. Because every distance Delaunay triangulation of a hyperbolic surface is a simplicial triangulation of the corresponding topological surface, it follows from Theorem 4.24 that the minimal number of vertices is at least

$$\left\lceil \frac{7 + \sqrt{1 + 48g}}{2} \right\rceil.$$

In the remainder of the proof, we will construct for a given hyperbolic surface a distance Delaunay triangulation with the required number of vertices, inspired by a similar construction in the context of the chromatic number of hyperbolic surfaces [63].

Let $n \equiv 0 \pmod{12}$ and assume that $n \neq 0$. The complete graph K_n on n vertices can be embedded in a topological surface S_g of genus

$$g = \frac{(n-3)(n-4)}{12},$$

which is the smallest possible genus [65]. Because we have assumed that $n \equiv 0 \pmod{12}$, we know that the embedding of K_n into S_g is a triangulation T [72]. To turn T into a distance Delaunay triangulation, we will add a hyperbolic metric to the topological surface as follows. Every triangle in T is replaced by the unique equilateral hyperbolic triangle with all three angles equal to $\frac{2\pi}{n-1}$. In the complete graph K_n every vertex has $n-1$ neighbouring vertices. This means that in every vertex $n-1$ equilateral triangles meet, so the total angle at each vertex is 2π . Therefore, the result after replacing all triangles in T by hyperbolic triangles is a smooth hyperbolic surface Z_g .

It remains to be shown that $T \in \mathcal{D}(Z)$. By construction, T is a simplicial complex. It has also been shown that all edges are distance paths [63]. We will show here that T is a Delaunay triangulation of Z_g . Consider an arbitrary triangle (u, v, w) in T with circumcenter c and let $p \notin \{u, v, w\}$ be an arbitrary vertex of T (Figure 4.12). Consider a distance path γ from c to p . We can regard γ as the concatenation of simple segments that each pass through an individual triangle.

The first of these simple segments starts from c and leaves the triangle (u, v, w) , so its length is at least the distance between c and a side of (u, v, w) . Therefore, denoting by x the projection of c on one of the edges as shown in Figure 4.12, the length of the first segment is at least $d(c, x)$. The last of the simple segments passes

Figure 4.12: Schematic overview of the proof of T being a Delaunay triangulation.

through a triangle, say Δ , before arriving at p , so it has to pass through the side of Δ opposite to p . Therefore, its length is at least the distance between p and the opposite side of Δ . It is known that the distance between a vertex and the opposite side of an equilateral triangle is at least $\frac{1}{2}\ell$, where ℓ denotes the length of the sides of the equilateral triangle [63]. Hence, $d(c, p) = \ell(\gamma) \geq d(c, x) + \frac{1}{2}\ell$. By the triangle inequality in triangle (c, w, x) we see that $d(c, w) \leq d(c, x) + d(x, w) = d(c, x) + \frac{1}{2}\ell$, so we conclude that $d(c, p) \geq d(c, w)$. This means that p is not contained in the interior of the circumcircle of (u, v, w) , which shows that (u, v, w) is a Delaunay triangle. By symmetry, it follows that all triangles are Delaunay triangles, which finishes the proof. $\square$

Appendix A

Omitted proofs

A.1 Proofs omitted in Chapter 2

Proof. (Theorem 2.8 (continued))

First, consider Inequality (2.6). The terms depending on ε are $\cosh(m_{\min})$ and $\sin(\xi_{\min})$. The 0-th order Taylor expansion of $\cosh(m_{\min})$ around $\varepsilon = 0$ is

$$\cosh(m_{\min}) = \cosh(m_{\min})|_{\varepsilon=0} + E_m(\varepsilon, g), \quad (\text{A.1})$$

where $E_m(\varepsilon, g)$ is the error term. Technically speaking, we only defined $m_{\min}, \xi_{\min}, d_{\min}$ and $\alpha_{\min}$ for $\varepsilon > 0$, but it is clear that the formulas given in Lemma 2.10 are well-defined for $\varepsilon = 0$ as well. The constant term in the Taylor expansion is equal to $\cosh(d(O, \mathbf{s}_j))$, where $\mathbf{s}_j$ is a side of D_g , i.e., it is equal to $\cot(\frac{\pi}{4g})$, as was shown in the proof of Lemma 2.5. The Lagrange error bound for Taylor expansions implies that if

$$\left| \frac{d \cosh(m_{\min})}{d\varepsilon} \right| \leq M_m(g)$$

for all ε and g , then

$$|E_m(\varepsilon, g)| \leq M_m(g)\varepsilon. \quad (\text{A.2})$$

Therefore, we will compute an upper bound for $|\frac{d}{d\varepsilon} \cosh(m_{\min})|$. Differentiating

$$\begin{aligned} \cosh(m_{\min}) &= \cosh(\operatorname{arctanh}(\tanh((1 - \varepsilon) \operatorname{dist}_{\mathbb{M}_g}) \cos((1 + \varepsilon) \frac{\pi}{4g}))), \\ &= \left(1 - \tanh^2((1 - \varepsilon) \operatorname{dist}_{\mathbb{M}_g}) \cos^2((1 + \varepsilon) \frac{\pi}{4g}) \right)^{-1/2} \end{aligned}$$

yields

$$\begin{aligned} \frac{d \cosh(m_{\min})}{d\varepsilon} &= - \frac{\operatorname{dist}_{\mathbb{M}_g} \tanh((1 - \varepsilon) \operatorname{dist}_{\mathbb{M}_g}) \cos^2((1 + \varepsilon) \frac{\pi}{4g})}{\cosh^2((1 - \varepsilon) \operatorname{dist}_{\mathbb{M}_g}) (1 - \tanh^2((1 - \varepsilon) \operatorname{dist}_{\mathbb{M}_g}) \cos^2((1 + \varepsilon) \frac{\pi}{4g}))^{3/2}} \\ &\quad - \frac{\tanh^2((1 - \varepsilon) \operatorname{dist}_{\mathbb{M}_g}) \sin((1 + \varepsilon) \frac{\pi}{4g}) \cos((1 + \varepsilon) \frac{\pi}{4g})}{(1 - \tanh^2((1 - \varepsilon) \operatorname{dist}_{\mathbb{M}_g}) \cos^2((1 + \varepsilon) \frac{\pi}{4g}))^{3/2}} \cdot \frac{\pi}{4g}. \end{aligned}$$

Because

$$0 \leq \tanh((1 - \varepsilon) \operatorname{dist}_{\mathbb{M}_g}) \leq \tanh(\operatorname{dist}_{\mathbb{M}_g}) \leq 1$$

and

$$0 \leq \cos((1 + \varepsilon) \frac{\pi}{4g}) \leq \cos(\frac{\pi}{4g}) \leq 1,$$

we see that

$$1 - \tanh^2((1 - \varepsilon) \operatorname{dist}_{\mathbb{M}_g}) \cos^2((1 + \varepsilon) \frac{\pi}{4g}) \geq 1 - \tanh^2(\operatorname{dist}_{\mathbb{M}_g}) \cos^2(\frac{\pi}{4g}) = \tan(\frac{\pi}{4g}).$$

Therefore,

$$\left| \frac{d \cosh(m_{\min})}{d\varepsilon} \right| \leq \frac{\operatorname{dist}_{\mathbb{M}_g} \cot^3(\frac{\pi}{4g})}{\cosh^2((1 - \varepsilon) \operatorname{dist}_{\mathbb{M}_g})} + \sin((1 + \varepsilon) \frac{\pi}{4g}) \cot^3(\frac{\pi}{4g}) \cdot \frac{\pi}{4g}. \quad (\text{A.3})$$

Note that $\cosh(x) \sim \frac{1}{2}e^x$ and $\operatorname{arccosh}(x) \sim \log(2x)$, so

$$\begin{aligned} \cosh((1 - \varepsilon) \operatorname{dist}_{\mathbb{M}_g}) &= \cosh((1 - \varepsilon) \operatorname{arccosh}(\cot^2(\frac{\pi}{4g}))), \\ &\sim \frac{1}{2} \exp((1 - \varepsilon) \log(2 \cot^2(\frac{\pi}{4g}))), \\ &= \frac{1}{2} (\sqrt{2} \cot(\frac{\pi}{4g}))^{2(1-\varepsilon)}. \end{aligned}$$

We assume without loss of generality that $\varepsilon \leq \frac{1}{4}$, so that $\cosh((1 - \varepsilon) \operatorname{dist}_{\mathbb{M}_g}) = \Omega(g^{3/2})$, since $\cot(\frac{\pi}{4g}) = \Theta(g)$. Then the leading term in the right-hand side of Inequality (A.3) is the second term, which is $\Theta(g)$, since $\sin(\frac{\pi}{4g}) = \Theta(g^{-1})$. Hence,

$$\left| \frac{d \cosh(m_{\min})}{d\varepsilon} \right| \leq C_m g$$

for some constant $C_m \in \mathbb{R}$, which means that we can take $M_m(g) := C_m g$ in Inequality (A.2).

Now, consider $\sin(\xi_{\min})$. We write

$$\sin(\xi_{\min}) = \sin(\xi_{\min})|_{\varepsilon=0} + E_\xi(\varepsilon, g) \quad (\text{A.4})$$

and we want to find $M_\xi(g)$ such that

$$\left| \frac{d \sin(\xi_{\min})}{d\varepsilon} \right| \leq M_\xi(g)$$

for all ε and g , so that

$$|E_\xi(\varepsilon, g)| \leq M_\xi(g) \varepsilon. \quad (\text{A.5})$$

Recall that $\xi_{\min} = 2\eta$, where η is the solution of

$$\cos(\eta) = \frac{\tanh((1 + \varepsilon) \operatorname{dist}_{\mathbb{M}_g})}{\tanh((1 - \varepsilon) \operatorname{dist}_{\mathbb{M}_g})} \cdot \cos\left((1 - \varepsilon) \frac{\pi}{2g} - \eta\right). \quad (\text{A.6})$$

For $\varepsilon = 0$, we see that $\eta = \frac{\pi}{4g}$, so $\sin(\xi_{\min})|_{\varepsilon=0} = \sin(\frac{\pi}{2g})$. Furthermore,

$$\frac{d \sin(\xi_{\min})}{d\varepsilon} = \cos(\xi_{\min}) \cdot \frac{d\xi_{\min}}{d\varepsilon} = 2 \cos(\xi_{\min}) \frac{d\eta}{d\varepsilon},$$

so

$$\left| \frac{d \sin(\xi_{\min})}{d\varepsilon} \right| \leq 2 \left| \frac{d\eta}{d\varepsilon} \right|. \quad (\text{A.7})$$

Differentiating both sides of Equation (A.6) with respect to ε we obtain after some simplifications

$$\begin{aligned} \frac{d\eta}{d\varepsilon} = & -\beta^{-1} \sin((1-\varepsilon)\frac{\pi}{2g} - \eta) \cdot \frac{\pi}{2g} - \frac{\beta^{-1} \operatorname{dist}_{\mathbb{M}_g} \cos((1-\varepsilon)\frac{\pi}{2g} - \eta)}{\tanh((1-\varepsilon) \operatorname{dist}_{\mathbb{M}_g}) \cosh^2((1+\varepsilon) \operatorname{dist}_{\mathbb{M}_g})} \\ & - \frac{\beta^{-1} \operatorname{dist}_{\mathbb{M}_g} \cos((1-\varepsilon)\frac{\pi}{2g} - \eta) \tanh((1+\varepsilon) \operatorname{dist}_{\mathbb{M}_g})}{\sinh^2((1-\varepsilon) \operatorname{dist}_{\mathbb{M}_g})}, \end{aligned}$$

where we abbreviated

$$\beta := \sin(\eta) + \frac{\tanh((1+\varepsilon) \operatorname{dist}_{\mathbb{M}_g})}{\tanh((1-\varepsilon) \operatorname{dist}_{\mathbb{M}_g})} \sin((1-\varepsilon)\frac{\pi}{2g} - \eta).$$

Since

$$\beta^{-1} \sin((1-\varepsilon)\frac{\pi}{2g} - \eta) \leq (\beta - \sin(\eta))^{-1} \sin((1-\varepsilon)\frac{\pi}{2g} - \eta) = \frac{\tanh((1-\varepsilon) \operatorname{dist}_{\mathbb{M}_g})}{\tanh((1+\varepsilon) \operatorname{dist}_{\mathbb{M}_g})} \leq 1$$

and in a similar way

$$\beta^{-1} \cos((1-\varepsilon)\frac{\pi}{2g} - \eta) \leq \cot((1-\varepsilon)\frac{\pi}{2g} - \eta) \leq \cot((1-\varepsilon)\frac{\pi}{4g}),$$

where the last step follows from $(1-\varepsilon)\frac{\pi}{2g} - \eta \geq (1-\varepsilon)\frac{\pi}{4g}$, we see that

$$\left| \frac{d\eta}{d\varepsilon} \right| \leq \frac{\pi}{2g} + \frac{\operatorname{dist}_{\mathbb{M}_g}}{\tanh((1-\varepsilon) \operatorname{dist}_{\mathbb{M}_g}) \cosh^2((1+\varepsilon) \operatorname{dist}_{\mathbb{M}_g})} + \frac{\operatorname{dist}_{\mathbb{M}_g}}{\sinh^2((1-\varepsilon) \operatorname{dist}_{\mathbb{M}_g})}.$$

Looking at the asymptotics in the same way as for $\cosh(m_{\min})$, we see that the leading term is $\frac{\pi}{2g}$, which is $\Theta(g^{-1})$. Using Inequality (A.7), it follows that

$$\left| \frac{d \sin(\xi_{\min})}{d\varepsilon} \right| \leq C_\xi g^{-1}$$

for some constant $C_\xi \in \mathbb{R}$, which means that we can take $M_\xi(g) := C_\xi g^{-1}$ in Inequality (A.5).

Plugging Equations (A.1) and (A.4) into Inequality (2.6) we obtain

$$2 \sin\left(\frac{(g+1)\pi}{12g}\right) \left(\cot\left(\frac{\pi}{4g}\right) + E_m(\varepsilon, g)\right) \left(\sin\left(\frac{\pi}{2g}\right) + E_\xi(\varepsilon, g)\right) > 1.$$

Because $\cot(\frac{\pi}{4g})\sin(\frac{\pi}{2g}) = 1 + \cos(\frac{\pi}{2g})$, we can rewrite this as

$$\begin{aligned} & 2 \sin\left(\frac{(g+1)\pi}{12g}\right) (1 + \cos(\frac{\pi}{2g})) - 1 > \\ & - 2 \sin\left(\frac{(g+1)\pi}{12g}\right) \left(\sin(\frac{\pi}{2g})E_m(\varepsilon, g) + \cot(\frac{\pi}{4g})E_\xi(\varepsilon, g) \right). \end{aligned} \quad (\text{A.8})$$

Note that the left-hand side is monotonically decreasing for $g \geq 2$ with infimum

$$\lim_{g \rightarrow \infty} 2 \sin\left(\frac{(g+1)\pi}{12g}\right) (1 + \cos(\frac{\pi}{2g})) - 1 = \sqrt{6} - \sqrt{2} - 1 \approx 0.035 \dots$$

In particular, the left-hand side is positive for all $g \geq 2$. Because $\sin(\frac{\pi}{2g}) = \Theta(g^{-1})$, $|E_m(\varepsilon, g)| \leq C_m g \varepsilon$, $\cot(\frac{\pi}{4g}) = \Theta(g)$ and $|E_\xi(\varepsilon, g)| \leq C_\xi g^{-1} \varepsilon$ we know that

$$2 \sin\left(\frac{(g+1)\pi}{12g}\right) \left(\sin(\frac{\pi}{2g})|E_m(\varepsilon, g)| + \cot(\frac{\pi}{4g})|E_\xi(\varepsilon, g)| \right) \leq C_{\varepsilon,1} \varepsilon$$

for some constant $C_{\varepsilon,1} \in \mathbb{R}$. Using $-E_m(\varepsilon, g) \leq |E_m(\varepsilon, g)|$ and $-E_\xi(\varepsilon, g) \leq |E_\xi(\varepsilon, g)|$, we conclude that Inequality (A.8), and hence Inequality (2.6), is satisfied for any $g \geq 2$ if

$$\varepsilon < c_{\varepsilon,1} := \frac{\sqrt{6} - \sqrt{2} - 1}{C_{\varepsilon,1}}.$$

Second, consider Inequality (2.7). The terms depending on ε are $\sinh^2(\frac{1}{2}d_{\min})$ and $\sin^2(\beta_{\min})$. As before, we write

$$\sinh^2(\tfrac{1}{2}d_{\min}) = \sinh^2(\tfrac{1}{2}d_{\min})|_{\varepsilon=0} + E_d(\varepsilon, g) \quad (\text{A.9})$$

and we want to find $M_d(g)$ such that

$$\left| \frac{d \sinh^2(\frac{1}{2}d_{\min})}{d\varepsilon} \right| \leq M_d(g)$$

for all ε and g , so that

$$|E_d(\varepsilon, g)| \leq M_d(g)\varepsilon. \quad (\text{A.10})$$

The constant term $\sinh^2(\frac{1}{2}d_{\min})|_{\varepsilon=0}$ is equal to $\sinh^2(\frac{1}{2}\ell)$, where ℓ is the length of the sides of D_g . It was shown in the proof of Lemma 2.5 that $\ell = 2 \operatorname{arccosh}(\cot(\frac{\pi}{4g}))$, so

$$\sinh^2(\tfrac{1}{2}d_{\min})|_{\varepsilon=0} = \cosh^2(\tfrac{1}{2}\ell) - 1 = \cot^2(\tfrac{\pi}{4g}) - 1.$$

Furthermore,

$$\sinh^2(\tfrac{1}{2}d_{\min}) = \frac{\cosh(d_{\min}) - 1}{2},$$

so

$$\left| \frac{d \sinh^2(\frac{1}{2}d_{\min})}{d\varepsilon} \right| = \frac{1}{2} \left| \frac{d \cosh(d_{\min})}{d\varepsilon} \right|.$$

Since

$$d_{\min} = \operatorname{arccosh}(\cosh^2((1 - \varepsilon) \operatorname{dist}_{\mathbb{M}_g}) - \sinh^2((1 - \varepsilon) \operatorname{dist}_{\mathbb{M}_g}) \cos((1 - \varepsilon) \frac{\pi}{2g})),$$

a straightforward computation shows that

$$\begin{aligned} \frac{d \cosh(d_{\min})}{d\varepsilon} &= -4 \operatorname{dist}_{\mathbb{M}_g} \sin^2((1 - \varepsilon) \frac{\pi}{4g}) \sinh((1 - \varepsilon) \operatorname{dist}_{\mathbb{M}_g}) \cosh((1 - \varepsilon) \operatorname{dist}_{\mathbb{M}_g}) \\ &\quad - \sinh^2((1 - \varepsilon) \operatorname{dist}_{\mathbb{M}_g}) \sin((1 - \varepsilon) \frac{\pi}{2g}) \cdot \frac{\pi}{2g} \end{aligned}$$

Using $\sin((1 - \varepsilon) \frac{\pi}{4g}) \leq \sin(\frac{\pi}{4g})$ and similar inequalities for $\cosh$ and $\sinh$ and substituting $\operatorname{dist}_{\mathbb{M}_g} = \operatorname{arccosh}(\cot^2(\frac{\pi}{4g}))$, we obtain

$$\left| \frac{d \sinh^2(\frac{1}{2} d_{\min})}{d\varepsilon} \right| \leq 2 \operatorname{dist}_{\mathbb{M}_g} \cos^2(\frac{\pi}{4g}) \sqrt{\cot^4(\frac{\pi}{4g}) - 1} - \cot(\frac{\pi}{4g}) (\cot^2(\frac{\pi}{4g}) - 1) \cdot \frac{\pi}{2g}$$

Here, the first term is the leading term, which is $\Theta(g^2 \log(g))$. Therefore,

$$\left| \frac{d \sinh^2(\frac{1}{2} d_{\min})}{d\varepsilon} \right| \leq C_d g^2 \log(g)$$

for some constant $C_d \in \mathbb{R}$, which means that we can take $M_d(g) := C_d g^2 \log(g)$ in Inequality (A.10).

Now consider $\sin^2(\alpha_{\min})$. We write

$$\sin^2(\alpha_{\min}) = \sin^2(\alpha_{\min})|_{\varepsilon=0} + E_\alpha(\varepsilon, g) \quad (\text{A.11})$$

and we want to find $M_\alpha(g)$ such that

$$\left| \frac{d \sin^2(\alpha_{\min})}{d\varepsilon} \right| \leq M_\alpha(g)$$

for all ε and g , so that

$$|E_\alpha(\varepsilon, g)| \leq M_\alpha(g) \varepsilon. \quad (\text{A.12})$$

Because the angle between any two adjacent sides of D_g is $\frac{\pi}{2g}$, the constant term in Equation (A.11) is $\sin^2(\frac{\pi}{2g})$. Because Lemma 2.10 gives an expression for $\sin(\frac{1}{2}\alpha_{\min})$ instead of $\sin(\alpha_{\min})$ we write

$$\sin^2(\alpha_{\min}) = 4 \sin^2(\frac{1}{2}\alpha_{\min}) \cos^2(\frac{1}{2}\alpha_{\min}) = 4 \sin^2(\frac{1}{2}\alpha_{\min}) - 4 \sin^4(\frac{1}{2}\alpha_{\min}),$$

so that

$$\begin{aligned} \frac{d \sin^2(\alpha_{\min})}{d\varepsilon} &= 8 \sin(\frac{1}{2}\alpha_{\min}) \cdot \frac{d \sin(\frac{1}{2}\alpha_{\min})}{d\varepsilon} - 16 \sin^3(\frac{1}{2}\alpha_{\min}) \cdot \frac{d \sin(\frac{1}{2}\alpha_{\min})}{d\varepsilon}, \\ &= 8 \sin(\frac{1}{2}\alpha_{\min}) \cos(\alpha_{\min}) \cdot \frac{d \sin(\frac{1}{2}\alpha_{\min})}{d\varepsilon}. \end{aligned}$$

Because $\sin(\alpha_{\min}) \leq \sin(\frac{\pi}{4g})$ and $\cos(\alpha_{\min}) \leq 1$, we see that

$$\left| \frac{d \sin^2(\alpha_{\min})}{d\varepsilon} \right| \leq 8 \sin(\frac{\pi}{4g}) \cdot \left| \frac{d \sin(\frac{1}{2}\alpha_{\min})}{d\varepsilon} \right|. \quad (\text{A.13})$$

Differentiating

$$\begin{aligned} \sinh(\tfrac{1}{2}\alpha_{\min}) &= \frac{\sinh(\operatorname{arctanh}(\tanh((1+\varepsilon)\operatorname{dist}_{\mathbb{M}_g})\cos(\eta')))}{\sinh((1+\varepsilon)\operatorname{dist}_{\mathbb{M}_g})}, \\ &= \frac{\tanh((1+\varepsilon)\operatorname{dist}_{\mathbb{M}_g})\cos(\eta')}{\sinh((1+\varepsilon)\operatorname{dist}_{\mathbb{M}_g})\sqrt{1-\tanh^2((1+\varepsilon)\operatorname{dist}_{\mathbb{M}_g})\cos^2(\eta')}}}, \\ &= \frac{\cos(\eta')}{\sqrt{\cosh^2((1+\varepsilon)\operatorname{dist}_{\mathbb{M}_g})-\sinh^2((1+\varepsilon)\operatorname{dist}_{\mathbb{M}_g})\cos^2(\eta')}}}, \end{aligned}$$

we obtain

$$\begin{aligned} \frac{d \sin(\frac{1}{2}\alpha_{\min})}{d\varepsilon} &= -\frac{\sin(\eta')}{\delta^{1/2}} \cdot \frac{d\eta'}{d\varepsilon} - \frac{\sin(\eta')\cos^2(\eta')\sinh^2((1+\varepsilon)\operatorname{dist}_{\mathbb{M}_g})}{\delta^{3/2}} \cdot \frac{d\eta'}{d\varepsilon} \\ &\quad - \frac{\operatorname{dist}_{\mathbb{M}_g}\sin^2(\eta')\cos(\eta')\cosh((1+\varepsilon)\operatorname{dist}_{\mathbb{M}_g})\sinh((1+\varepsilon)\operatorname{dist}_{\mathbb{M}_g})}{\delta^{3/2}}, \end{aligned}$$

where we abbreviated

$$\delta := \cosh^2((1+\varepsilon)\operatorname{dist}_{\mathbb{M}_g}) - \sinh^2((1+\varepsilon)\operatorname{dist}_{\mathbb{M}_g})\cos^2(\eta').$$

Using

$$\begin{aligned} \delta &\geq \cosh^2(\operatorname{dist}_{\mathbb{M}_g}) - \sinh^2(\operatorname{dist}_{\mathbb{M}_g})\cos^2(\frac{\pi}{4g}), \\ &= \cot^2(\frac{\pi}{4g}) \end{aligned}$$

in the first term and

$$\begin{aligned} \delta &= \sin^2(\eta')\cosh^2((1+\varepsilon)\operatorname{dist}_{\mathbb{M}_g}) + \cos^2(\eta'), \\ &\geq \sin^2(\eta')\cosh^2((1+\varepsilon)\operatorname{dist}_{\mathbb{M}_g}) \end{aligned}$$

in the second and third term, we see that

$$\begin{aligned} \left| \frac{d \sin(\frac{1}{2}\alpha_{\min})}{d\varepsilon} \right| &\leq \sin(\eta')\tan(\frac{\pi}{4g}) \cdot \left| \frac{d\eta'}{d\varepsilon} \right| + \frac{\operatorname{dist}_{\mathbb{M}_g}}{\sin(\eta')\cosh((1+\varepsilon)\operatorname{dist}_{\mathbb{M}_g})} \\ &\quad + \frac{1}{\sin^2(\eta')\cosh((1+\varepsilon)\operatorname{dist}_{\mathbb{M}_g})} \cdot \left| \frac{d\eta'}{d\varepsilon} \right| \end{aligned}$$

Since

$$\sin(\frac{\pi}{4g}) \leq \sin(\eta') \leq \sin(\frac{\pi}{2g}) \leq 2 \sin(\frac{\pi}{4g})$$

and

$$\cosh((1 + \varepsilon) \operatorname{dist}_{\mathbb{M}_g}) \geq \cosh(\operatorname{dist}_{\mathbb{M}_g}) = \cot^2\left(\frac{\pi}{4g}\right),$$

it follows that

$$\left| \frac{d \sin(\frac{1}{2} \alpha_{\min})}{d\varepsilon} \right| \leq 2 \sin\left(\frac{\pi}{4g}\right) \tan\left(\frac{\pi}{4g}\right) \cdot \left| \frac{d\eta'}{d\varepsilon} \right| + \frac{\operatorname{dist}_{\mathbb{M}_g} \tan\left(\frac{\pi}{4g}\right)}{\cos\left(\frac{\pi}{4g}\right)} + \frac{1}{\cos^2\left(\frac{\pi}{4g}\right)} \cdot \left| \frac{d\eta'}{d\varepsilon} \right|.$$

By similar computations as for η , it can be shown that

$$\left| \frac{d\eta'}{d\varepsilon} \right| \leq \frac{\pi}{2g} + \frac{\operatorname{dist}_{\mathbb{M}_g}}{\tanh((1 - \varepsilon) \operatorname{dist}_{\mathbb{M}_g}) \cosh^2((1 + \varepsilon) \operatorname{dist}_{\mathbb{M}_g})} + \frac{\operatorname{dist}_{\mathbb{M}_g}}{\sinh^2((1 - \varepsilon) \operatorname{dist}_{\mathbb{M}_g})}.$$

In particular, the leading term of the right-hand side is $\Theta(g^{-1})$. It follows that the leading term of the upper bound for $\left| \frac{d}{d\varepsilon} \sin\left(\frac{1}{2} \alpha_{\min}\right) \right|$ is the second term, which is $\Theta(g^{-1} \log(g))$, since $\tan\left(\frac{\pi}{4g}\right) = \Theta(g^{-1})$. Using Inequality (A.13) and $\sin\left(\frac{\pi}{4g}\right) = \Theta(g^{-1})$, it follows that

$$\left| \frac{d \sin^2(\alpha_{\min})}{d\varepsilon} \right| \leq C_\alpha g^{-2} \log(g)$$

for some constant $C_\alpha \in \mathbb{R}$, which means that we can take $M_\alpha(g) = C_\alpha g^{-2} \log(g)$ in Inequality (A.12).

Plugging Equations (A.9) and (A.11) into Inequality (2.7) and rearranging the terms as we did for Inequality (2.6), we obtain

$$\begin{aligned} & 4 \sin^2\left(\frac{(g+1)\pi}{12g}\right) (1 + 4 \cos^2\left(\frac{\pi}{4g}\right) \cos\left(\frac{\pi}{2g}\right)) - 1 > \\ & - 4 \sin^2\left(\frac{(g+1)\pi}{12g}\right) \left(\sin^2\left(\frac{\pi}{2g}\right) E_d(\varepsilon, g) + (\cot^2\left(\frac{\pi}{4g}\right) - 1) E_\alpha(\varepsilon, g) \right). \end{aligned} \quad (\text{A.14})$$

Note that the left-hand side is monotonically decreasing for $g \geq 2$ with infimum

$$\lim_{g \rightarrow \infty} 4 \sin^2\left(\frac{(g+1)\pi}{12g}\right) (1 + 4 \cos^2\left(\frac{\pi}{4g}\right)) - 1 = 5(2 - \sqrt{3}) - 1 \approx 0.340 \dots$$

In particular, the left-hand side is positive for all $g \geq 2$. Because $\sin\left(\frac{\pi}{2g}\right) = \Theta(g^{-1})$, $|E_d(\varepsilon, g)| \leq C_d g^2 \log(g) \varepsilon$, $\cot\left(\frac{\pi}{4g}\right) = \Theta(g)$ and $|E_\alpha(\varepsilon, g)| \leq C_\alpha g^{-2} \log(g) \varepsilon$, we know that

$$4 \sin^2\left(\frac{(g+1)\pi}{12g}\right) \left(\sin^2\left(\frac{\pi}{2g}\right) |E_d(\varepsilon, g)| + (\cot^2\left(\frac{\pi}{4g}\right) - 1) |E_\alpha(\varepsilon, g)| \right) \leq C_{\varepsilon, 2} \log(g) \varepsilon$$

for some constant $C_{\varepsilon, 2} \in \mathbb{R}$. Using $-E_d(\varepsilon, g) \leq |E_d(\varepsilon, g)|$ and $-E_\alpha(\varepsilon, g) \leq |E_\alpha(\varepsilon, g)|$, we conclude that Inequality (A.14), and hence Inequality (2.7), is satisfied for any $g \geq 2$ if $\varepsilon < c_{\varepsilon, 2} (\log(g))^{-1}$, where

$$c_{\varepsilon, 2} := \frac{9 - 5\sqrt{3}}{C_{\varepsilon, 2}}.$$

Combining the results for Inequalities (2.6) and (2.7), we see that both inequalities are satisfied for $\varepsilon = c(\log(g))^{-1}$ with $c = \min\{c_{\varepsilon,1}, c_{\varepsilon,2}\}$. This finishes the proof. $\square$

A.2 Proofs omitted in Chapter 3

Proof. (Lemma 3.15)

For convenience, we show the figure illustrating the dummy points within one slice of the $4g$ -gon once again in Figure A.1.

Figure A.1: Dummy points within one slice of the $4g$ -gon.

By symmetry it is sufficient to consider the circumscribed disks of the triangles

$$[p_0^0, p_0^1, p_0^*], [p_0^1, p_1^1, p_0^*], [p_0^m, p_1^m, O], [p_0^j, p_1^j, p_1^{j+1}],$$

for $1 \leq j \leq m-1$. For easy reference, the used lengths and angles satisfy the following relations. Here we denote $\text{length}([p, q])$ by abuse of notation by $[p, q]$.

$$[p_0^0, p_0^*] = [p_0^0, p_0^1] = \frac{1}{4} \text{sys}(\mathbb{M}_g), \quad (\text{A.15})$$

$$\cosh(\frac{1}{2} \text{sys}(\mathbb{M}_g)) = 1 + 2 \cos(\frac{\pi}{2g}), \quad (\text{A.16})$$

$$\angle(OP_0^0 p_0^*) = \frac{\pi}{4}, \quad (\text{A.17})$$

$$\sinh(\frac{1}{2} [p_0^1, p_0^*]) = \sinh(\frac{1}{4} \text{sys}(\mathbb{M}_g)) \sin(\frac{1}{2} \angle(OP_0^0 p_0^*)), \quad (\text{A.18})$$

$$\sin \angle(p_0^1 p_0^* p_0^0) = \frac{\sinh(\frac{1}{4} \text{sys}(\mathbb{M}_g)) \sin \angle(OP_0^0 p_0^*)}{\sinh([p_0^1, p_0^*])}, \quad (\text{A.19})$$

$$\angle(p_0^1 p_0^* p_1^1) = \pi - 2\angle(p_0^1 p_0^* p_0^0). \quad (\text{A.20})$$

Equations (A.15) and (A.16) follow from the construction in the proof of Lemma 2.3 and from Theorem 2.1, respectively. Equation (A.17) holds because $\angle(OP_0^0 v_1) = \frac{\pi}{2}$ and the triangles $[O, p_0^0, p_0^1]$ and $[v_1, p_0^0, p_0^1]$ are congruent. Equation (A.18) follows from the application of [10, Theorem 7.11.2(ii)] to the triangle with vertices p_0^1, p_0^0 and the midpoint of p_0^0 and p_0^* : this is a right-angled triangle, because

Figure A.2: Deriving a formula for the circumradius of an isosceles triangle.

$[\mathbf{p}_0^1, \mathbf{p}_0^0, \mathbf{p}_0^*]$ is isosceles. Equation (A.19) is the sine rule [10, Chapter 7.12] in triangle $[\mathbf{p}_0^1, \mathbf{p}_0^0, \mathbf{p}_0^*]$. Finally, Equation (A.20) holds by symmetry.

The circumradius R of an isosceles triangle with legs of length b and vertex angle α (see Figure A.2) satisfies

$$\tanh(R) = \frac{\tanh(\frac{1}{2}c)}{\cos(\frac{1}{2}\alpha)}, \quad (\text{A.21})$$

which can be derived by applying [10, Theorem 7.1..2(iii)] to the interior triangle with edges of length R and $\frac{1}{2}b$.

First, consider the circumradius $R(\mathbf{p}_0^0, \mathbf{p}_0^1, \mathbf{p}_0^*)$ of $[\mathbf{p}_0^0, \mathbf{p}_0^1, \mathbf{p}_0^*]$. By the above observation,

$$\begin{aligned} \tanh(R(\mathbf{p}_0^0, \mathbf{p}_0^1, \mathbf{p}_0^*)) &= \frac{\tanh(\frac{1}{8} \text{sys}(\mathbb{M}_g))}{\cos(\frac{1}{2} \angle(O\mathbf{p}_0^0\mathbf{p}_0^*))}, \\ &= \sqrt{4 - 2\sqrt{2} \tanh(\frac{1}{8} \text{sys}(\mathbb{M}_g))}. \end{aligned}$$

We want to prove that $R(\mathbf{p}_0^0, \mathbf{p}_0^1, \mathbf{p}_0^*) < \frac{1}{4} \text{sys}(\mathbb{M}_g)$, or, equivalently, that

$$\tanh(R(\mathbf{p}_0^0, \mathbf{p}_0^1, \mathbf{p}_0^*)) < \tanh(\frac{1}{4} \text{sys}(\mathbb{M}_g)).$$

Because $\tanh(\frac{1}{8} \text{sys}(\mathbb{M}_g))$ is strictly increasing as function of g , we know that

$$\sqrt{4 - 2\sqrt{2} \tanh(\frac{1}{8} \text{sys}(\mathbb{M}_g))} < \sqrt{4 - 2\sqrt{2}} \lim_{g \rightarrow \infty} \tanh(\frac{1}{8} \text{sys}(\mathbb{M}_g)) \approx 0.448.$$

For the same reason,

$$\tanh(\tfrac{1}{4} \text{sys}(\mathbb{M}_g)) \geq \tanh(\tfrac{1}{4} \text{sys}(\mathbb{M}_g))|_{g=2} \approx 0.643.$$

It follows that $\tanh(R(\mathbf{p}_0^0, \mathbf{p}_0^1, \mathbf{p}_0^*)) < \tanh(\tfrac{1}{4} \text{sys}(\mathbb{M}_g))$ holds. This proves that the circumdiameter of $[\mathbf{p}_0^0, \mathbf{p}_0^1, \mathbf{p}_0^*]$ is smaller than $\tfrac{1}{2} \text{sys}(\mathbb{M}_g)$. Second, consider the circumradius $R(\mathbf{p}_0^1, \mathbf{p}_1^1, \mathbf{p}_0^*)$ of $[\mathbf{p}_0^1, \mathbf{p}_1^1, \mathbf{p}_0^*]$. By a similar computation as above, we see that

$$\begin{aligned} \tanh(R(\mathbf{p}_0^1, \mathbf{p}_1^1, \mathbf{p}_0^*)) &= \frac{\tanh(\tfrac{1}{2}[\mathbf{p}_0^1, \mathbf{p}_0^*])}{\cos(\tfrac{1}{2}\angle(\mathbf{p}_0^1 \mathbf{p}_0^* \mathbf{p}_1^1))}, \\ &\stackrel{(A.20)}{=} \frac{\tanh(\tfrac{1}{2}[\mathbf{p}_0^1, \mathbf{p}_0^*])}{\sin \angle(\mathbf{p}_0^1 \mathbf{p}_0^* \mathbf{p}_0^0)}. \end{aligned}$$

By using

$$\begin{aligned} \sin \angle(\mathbf{p}_0^1 \mathbf{p}_0^* \mathbf{p}_0^0) &\stackrel{(A.19)}{=} \frac{\sinh(\tfrac{1}{4} \text{sys}(\mathbb{M}_g)) \sin \angle(\mathbf{Op}_0^0 \mathbf{p}_0^*)}{\sinh([\mathbf{p}_0^1, \mathbf{p}_0^*])}, \\ &= \frac{\sinh(\tfrac{1}{4} \text{sys}(\mathbb{M}_g)) \cdot 2 \sin(\tfrac{1}{2}\angle(\mathbf{Op}_0^0 \mathbf{p}_0^*)) \cos(\tfrac{1}{2}\angle(\mathbf{Op}_0^0 \mathbf{p}_0^*))}{\sinh([\mathbf{p}_0^1, \mathbf{p}_0^*])}, \\ &\stackrel{(A.18)}{=} \frac{\sinh(\tfrac{1}{2}[\mathbf{p}_0^1, \mathbf{p}_0^*]) \cdot 2 \cos(\tfrac{1}{2}\angle(\mathbf{Op}_0^0 \mathbf{p}_0^*))}{2 \sinh(\tfrac{1}{2}[\mathbf{p}_0^1, \mathbf{p}_0^*]) \cosh(\tfrac{1}{2}[\mathbf{p}_0^1, \mathbf{p}_0^*])}, \\ &= \frac{\cos(\tfrac{1}{2}\angle(\mathbf{Op}_0^0 \mathbf{p}_0^*))}{\cosh(\tfrac{1}{2}[\mathbf{p}_0^1, \mathbf{p}_0^*])}, \end{aligned}$$

we can rewrite this as

$$\begin{aligned} \tanh(R(\mathbf{p}_0^1, \mathbf{p}_1^1, \mathbf{p}_0^*)) &= \frac{\tanh(\tfrac{1}{2}[\mathbf{p}_0^1, \mathbf{p}_0^*]) \cosh(\tfrac{1}{2}[\mathbf{p}_0^1, \mathbf{p}_0^*])}{\cos(\tfrac{1}{2}\angle(\mathbf{Op}_0^0 \mathbf{p}_0^*))}, \\ &= \frac{\sinh(\tfrac{1}{2}[\mathbf{p}_0^1, \mathbf{p}_0^*])}{\cos(\tfrac{1}{2}\angle(\mathbf{Op}_0^0 \mathbf{p}_0^*))}, \\ &\stackrel{(A.18)}{=} \sinh(\tfrac{1}{4} \text{sys}(\mathbb{M}_g)) \tan(\tfrac{1}{2}\angle(\mathbf{Op}_0^0 \mathbf{p}_0^*)), \\ &= (\sqrt{2} - 1) \sinh(\tfrac{1}{4} \text{sys}(\mathbb{M}_g)). \end{aligned}$$

By the same reasoning as before,

$$\tanh(R(\mathbf{p}_0^1, \mathbf{p}_1^1, \mathbf{p}_0^*)) < (\sqrt{2} - 1) \lim_{g \rightarrow \infty} \sinh(\tfrac{1}{4} \text{sys}(\mathbb{M}_g)) \approx 0.414,$$

which shows that $\tanh(R(\mathbf{p}_0^1, \mathbf{p}_1^1, \mathbf{p}_0^*)) < \tanh(\tfrac{1}{4} \text{sys}(\mathbb{M}_g))$. This proves that the circumdiameter of $[\mathbf{p}_0^1, \mathbf{p}_1^1, \mathbf{p}_0^*]$ is smaller than $\tfrac{1}{2} \text{sys}(\mathbb{M}_g)$.

Third, consider the circumradius $R(\mathbf{p}_0^m, \mathbf{p}_1^m, O)$ of $[\mathbf{p}_0^m, \mathbf{p}_1^m, O]$. We know that $[O, \mathbf{p}_0^m] = x \leq \frac{1}{4} \text{sys}(\mathbb{M}_g)$. By a similar computation as above, we see that

$$\begin{aligned} \tanh(R(\mathbf{p}_0^m, \mathbf{p}_1^m, O)) &= \frac{\tanh(\frac{1}{2}x)}{\cos(\frac{\pi}{4g})}, \\ &\leq \frac{\tanh(\frac{1}{8} \text{sys}(\mathbb{M}_g))}{\cos(\frac{\pi}{8})}, \\ &= \tanh(R(\mathbf{p}_0^0, \mathbf{p}_0^1, \mathbf{p}_0^*)), \\ &< \tanh(\frac{1}{4} \text{sys}(\mathbb{M}_g)), \end{aligned}$$

from which we conclude that $R(\mathbf{p}_0^m, \mathbf{p}_1^m, O) < \frac{1}{4} \text{sys}(\mathbb{M}_g)$. This proves that the circumdiameter of $[\mathbf{p}_0^m, \mathbf{p}_1^m, O]$ is smaller than $\frac{1}{2} \text{sys}(\mathbb{M}_g)$.

Finally, let $1 \leq j \leq m-1$ and consider the circumradius $R(\mathbf{p}_0^j, \mathbf{p}_1^j, \mathbf{p}_1^{j+1})$ of $[\mathbf{p}_0^j, \mathbf{p}_1^j, \mathbf{p}_1^{j+1}]$. In fact, the points $\mathbf{p}_0^j, \mathbf{p}_1^j, \mathbf{p}_0^{j+1}, \mathbf{p}_1^{j+1}$ are concircular due to symmetry, so the center of the circumscribed circle of $[\mathbf{p}_0^j, \mathbf{p}_1^j, \mathbf{p}_1^{j+1}]$ lies on the line segment $[O, \mathbf{p}_0^*]$. It follows that the circumradius of this disk decreases when the distance of $\mathbf{p}_0^j$ to $[O, \mathbf{p}_0^*]$ decreases, or equivalently, when the distance $[O, \mathbf{p}_0^j]$ decreases. Hence, for all $1 \leq j \leq m-1$, $R(\mathbf{p}_0^j, \mathbf{p}_1^j, \mathbf{p}_1^{j+1}) < R(\mathbf{p}_0^1, \mathbf{p}_1^1, \mathbf{p}_1^2)$. Therefore, it is sufficient to show that $R(\mathbf{p}_0^1, \mathbf{p}_1^1, \mathbf{p}_1^2) < \frac{1}{4} \text{sys}(\mathbb{M}_g)$.

In this case, we cannot use equation (A.21), since $[\mathbf{p}_0^1, \mathbf{p}_1^1, \mathbf{p}_1^2]$ is not an isosceles triangle. There exists a more general expression for the circumradius of an arbitrary (not necessarily isosceles) triangle, but this will lead to unnecessarily long expressions. Instead, consider the circumcenter $\mathbf{c}$ of $[\mathbf{p}_0^1, \mathbf{p}_1^1, \mathbf{p}_1^2]$. See also Figure A.3 for a more detailed view of the relevant triangles. By the discussion above, we

Figure A.3: Close-up of situation at $[\mathbf{p}_0^1, \mathbf{p}_1^1, \mathbf{p}_1^2]$

know that $\mathbf{c} \in [O, \mathbf{p}_0^*]$. Let $\mathbf{c}'$ be the orthogonal projection of $\mathbf{c}$ onto $[O, \mathbf{p}_0^*]$. Since

$[\mathbf{c}, \mathbf{p}_1^1, \mathbf{p}_1^2]$ is isosceles, we know that

$$[O, \mathbf{c}'] = [O, \mathbf{p}_1^2] + \frac{1}{8} \text{sys}(\mathbb{M}_g) = [O, \mathbf{p}_1^0] - \frac{3}{8} \text{sys}(\mathbb{M}_g) = \text{arccosh}(\cot(\frac{\pi}{4g})) - \frac{3}{8} \text{sys}(\mathbb{M}_g),$$

from which it can be seen that $[O, \mathbf{c}']$ is strictly increasing as function of g . Furthermore, [10, Theorem 7.11.2(i)]

$$\tanh([\mathbf{c}, \mathbf{c}']) = \sinh([O, \mathbf{c}']) \tanh(\frac{\pi}{4g}),$$

which after substitution of our expression for $[O, \mathbf{c}']$ can be rewritten as

$$\tanh([\mathbf{c}, \mathbf{c}']) = \sqrt{1 - \tanh^2(\frac{\pi}{4g})} \cosh(\frac{3}{8} \text{sys}(\mathbb{M}_g)) - \sinh(\frac{3}{8} \text{sys}(\mathbb{M}_g)).$$

By the Pythagorean law in $[\mathbf{c}, \mathbf{c}', \mathbf{p}_1^2]$ we know

$$\cosh(R(\mathbf{p}_0^1, \mathbf{p}_1^1, \mathbf{p}_1^2)) = \cosh(\frac{1}{8} \text{sys}(\mathbb{M}_g)) \cosh([\mathbf{c}, \mathbf{c}']).$$

Using this expression of $R(\mathbf{p}_0^1, \mathbf{p}_1^1, \mathbf{p}_1^2)$, it can be seen that $R(\mathbf{p}_0^1, \mathbf{p}_1^1, \mathbf{p}_1^2)$ is strictly increasing as function of g . Therefore,

$$\cosh(R(\mathbf{p}_0^1, \mathbf{p}_1^1, \mathbf{p}_1^2)) < \lim_{g \rightarrow \infty} \cosh(R(\mathbf{p}_0^1, \mathbf{p}_1^1, \mathbf{p}_1^2)) \approx 1.140,$$

which can be obtained by computing the corresponding limits of $\cosh(\frac{1}{8} \text{sys}(\mathbb{M}_g))$ and $\cosh([\mathbf{c}, \mathbf{c}'])$. Hence, $\tanh(R(\mathbf{p}_0^1, \mathbf{p}_1^1, \mathbf{p}_1^2)) < 0.480$, so $\tanh(R(\mathbf{p}_0^1, \mathbf{p}_1^1, \mathbf{p}_1^2)) < \tanh(\frac{1}{4} \text{sys}(\mathbb{M}_g))$. This proves that the circumdiameter of $[\mathbf{p}_0^1, \mathbf{p}_1^1, \mathbf{p}_1^2]$ is smaller than $\frac{1}{2} \text{sys}(\mathbb{M}_g)$. Since the circumdiameters of all four triangles are smaller than $\frac{1}{2} \text{sys}(\mathbb{M}_g)$ and since by symmetry every triangle is congruent to one of these four, this concludes the proof. $\square$

Proof. (Lemma 3.16)

We show that the circumdisks of triangles in $\mathcal{T}$ do not have vertices of $\mathcal{T}$ in their interior. Denote the circumscribed circle and *open* circumscribed disk of a triangle $[\mathbf{p}, \mathbf{q}, \mathbf{r}]$ by $C(\mathbf{p}, \mathbf{q}, \mathbf{r})$ and $D(\mathbf{p}, \mathbf{q}, \mathbf{r})$ respectively, and similarly when more than three points are concircular. Observe in particular that $C(\mathbf{p}, \mathbf{q}, \mathbf{r}) \cap D(\mathbf{p}, \mathbf{q}, \mathbf{r}) = \emptyset$. Denote the hyperbolic line through points $\mathbf{p}, \mathbf{q}$ by $L(\mathbf{p}, \mathbf{q})$ and the open and closed line segments connecting $\mathbf{p}, \mathbf{q}$ by $(\mathbf{p}, \mathbf{q})$, $[\mathbf{p}, \mathbf{q}]$ respectively. By symmetry it is sufficient to consider only $D(\mathbf{p}_0^0, \mathbf{p}_1^0, \mathbf{p}_0^*)$, $D(\mathbf{p}_0^1, \mathbf{p}_1^1, \mathbf{p}_0^*)$, $D(\mathbf{p}_0^m, \mathbf{p}_1^m, O)$ and $D(\mathbf{p}_0^j, \mathbf{p}_1^j, \mathbf{p}_1^{j+1})$ for $1 \leq j \leq m-1$. For convenience, we treat the cases for each circumdisk in a fixed order, namely

1. O, v
2. $\mathbf{p}_k^*, k = 0, \dots, 4g-1$,
3. $\mathbf{p}_k^j, k = 0, \dots, 4g-1, j = 0, \dots, m$,

4. $\mathbf{q}_k^j, k = 0, \dots, 4g-1, j = 1, \dots, m$.

First, consider $D(\mathbf{p}_0^0, \mathbf{p}_0^1, \mathbf{p}_0^*)$.

1. Clearly, O and $\mathbf{v}$ are too far away from $\mathbf{p}_0^*$ to be inside $D(\mathbf{p}_0^0, \mathbf{p}_0^1, \mathbf{p}_0^*)$.
2. Since $\mathbf{p}_0^* \in C(\mathbf{p}_0^0, \mathbf{p}_0^1, \mathbf{p}_0^*)$, we know that $\mathbf{p}_0^* \notin D(\mathbf{p}_0^0, \mathbf{p}_0^1, \mathbf{p}_0^*)$. Since the center of $D(\mathbf{p}_0^0, \mathbf{p}_0^1, \mathbf{p}_0^*)$ lies inside $D(\mathbf{p}_0^0, \mathbf{p}_0^1, \mathbf{p}_0^*)$ on the bisector of angle $\angle(\mathbf{p}_0^1 \mathbf{p}_0^0 \mathbf{p}_0^*)$ we see that $\mathbf{p}_{4g-1}^*$ is farther away from this center than $\mathbf{p}_0^*$. Since $\mathbf{p}_0^* \notin D(\mathbf{p}_0^0, \mathbf{p}_0^1, \mathbf{p}_0^*)$, it follows that $\mathbf{p}_{4g-1}^* \notin D(\mathbf{p}_0^0, \mathbf{p}_0^1, \mathbf{p}_0^*)$ as well. Since $D(\mathbf{p}_0^0, \mathbf{p}_0^1, \mathbf{p}_0^*) \cap C(\mathbf{p}_0^*, \dots, \mathbf{p}_{4g-1}^*)$ is contained in the shortest open chord of $C(\mathbf{p}_0^*, \dots, \mathbf{p}_{4g-1}^*)$ between $\mathbf{p}_0^*$ and $\mathbf{p}_{4g-1}^*$, we see that $\mathbf{p}_k^* \notin D(\mathbf{p}_0^0, \mathbf{p}_0^1, \mathbf{p}_0^*)$ for $k = 0, \dots, 4g-1$.
3. Since the center of $D(\mathbf{p}_0^0, \mathbf{p}_0^1, \mathbf{p}_0^*)$ is in the interior of $[\mathbf{p}_0^0, \mathbf{p}_0^1, \mathbf{p}_0^*] \subset [O, \mathbf{p}_0^0, \mathbf{p}_0^*]$, we know that $\mathbf{p}_0^j$ is closer to the center of $D(\mathbf{p}_0^0, \mathbf{p}_0^1, \mathbf{p}_0^*)$ than $\mathbf{p}_k^j$ for $j = 0, \dots, m$ and $k \neq 0$. Because $L(O, \mathbf{p}_0^0) \cap D(\mathbf{p}_0^0, \mathbf{p}_0^1, \mathbf{p}_0^*) = (\mathbf{p}_0^0, \mathbf{p}_0^1)$, we know that $\mathbf{p}_0^j \notin D(\mathbf{p}_0^0, \mathbf{p}_0^1, \mathbf{p}_0^*)$ for all $j = 0, \dots, m$. Therefore, $\mathbf{p}_k^j \notin D(\mathbf{p}_0^0, \mathbf{p}_0^1, \mathbf{p}_0^*)$ for $j = 0, \dots, m$ and $k = 0, \dots, 4g-1$.
4. By a reasoning similar to above, $\mathbf{p}_k^j$ is closer to the center of $D(\mathbf{p}_0^0, \mathbf{p}_0^1, \mathbf{p}_0^*)$ than $\mathbf{q}_k^j$. Since by the previous step $\mathbf{p}_k^j \notin D(\mathbf{p}_0^0, \mathbf{p}_0^1, \mathbf{p}_0^*)$ for $j = 1, \dots, m$ and $k = 0, \dots, 4g-1$, it follows that $\mathbf{q}_k^j \notin D(\mathbf{p}_0^0, \mathbf{p}_0^1, \mathbf{p}_0^*)$ for $j = 1, \dots, m$ and $k = 0, \dots, 4g-1$ as well.

Second, consider $D(\mathbf{p}_0^1, \mathbf{p}_1^1, \mathbf{p}_0^*)$.

1. Clearly, O and $\mathbf{v}$ are too far away from $\mathbf{p}_0^*$ to be inside $D(\mathbf{p}_0^1, \mathbf{p}_1^1, \mathbf{p}_0^*)$.
2. The circle $C(\mathbf{p}_0^1, \mathbf{p}_1^1, \mathbf{p}_0^*)$ is tangent to $C(\mathbf{p}_0^*, \dots, \mathbf{p}_{4g-1}^*)$, because both circles have their center on $[O, \mathbf{p}_0^*]$ and pass through $\mathbf{p}_0^*$. Since the radius of $C(\mathbf{p}_0^1, \mathbf{p}_1^1, \mathbf{p}_0^*)$ is smaller than the radius of $C(\mathbf{p}_0^*, \dots, \mathbf{p}_{4g-1}^*)$, this means that $D(\mathbf{p}_0^1, \mathbf{p}_1^1, \mathbf{p}_0^*) \subseteq D(\mathbf{p}_0^*, \dots, \mathbf{p}_{4g-1}^*)$. Therefore, $\mathbf{p}_k^* \notin D(\mathbf{p}_0^1, \mathbf{p}_1^1, \mathbf{p}_0^*)$ for all $k = 0, \dots, 4g-1$.
3. First, to prove that $\mathbf{p}_0^j \notin D(\mathbf{p}_0^1, \mathbf{p}_1^1, \mathbf{p}_0^*)$ for all $j = 0, \dots, m$, we now show that $D(\mathbf{p}_0^1, \mathbf{p}_1^1, \mathbf{p}_0^*) \cap (O, \mathbf{p}_0^1) = \emptyset$. First observe that the line $L(O, \mathbf{p}_0^0)$ intersects $C(\mathbf{p}_0^1, \mathbf{p}_1^1, \mathbf{p}_0^*)$ in one or two points. If the intersection consists of one point, then it has to be $\mathbf{p}_0^1$ and we are done. If the intersection consists of two points, then it is sufficient to show that $\mathbf{p}_0^1$ is the closest of these two to O . Let $\mathbf{p}_0^M$ denote the midpoint of $\mathbf{p}_0^1$ and $\mathbf{p}_1^1$. It is sufficient to show that $[\mathbf{p}_0^*, \mathbf{p}_0^M] \geq R(\mathbf{p}_0^1, \mathbf{p}_1^1, \mathbf{p}_0^*)$, since then the center of $D(\mathbf{p}_0^1, \mathbf{p}_1^1, \mathbf{p}_0^*)$ is contained in $[\mathbf{p}_0^1, \mathbf{p}_1^1, \mathbf{p}_0^*]$. It is known that [10, Theorem 7.11.2(iii)]

$$\begin{aligned} \tanh([\mathbf{p}_0^*, \mathbf{p}_0^M]) &= \tanh([\mathbf{p}_0^1, \mathbf{p}_0^*]) \cos \frac{1}{2} \angle(\mathbf{p}_0^1 \mathbf{p}_0^* \mathbf{p}_1^1), \\ &= \tanh([\mathbf{p}_0^1, \mathbf{p}_0^*]) \sin \angle(\mathbf{p}_0^1 \mathbf{p}_0^* \mathbf{p}_0^0), \end{aligned}$$

where the second equality follows from Equation (A.20) in the proof of Lemma 3.15. Therefore, $[\mathbf{p}_0^*, \mathbf{p}_0^M] \geq R(\mathbf{p}_0^1, \mathbf{p}_1^1, \mathbf{p}_0^*)$ is equivalent with the following sequence of inequalities:

$$\begin{aligned} \tanh([\mathbf{p}_0^*, \mathbf{p}_0^M]) &\geq \tanh(R(\mathbf{p}_0^1, \mathbf{p}_1^1, \mathbf{p}_0^*)), \\ \tanh([\mathbf{p}_0^1, \mathbf{p}_0^*]) \sin \angle(\mathbf{p}_0^1 \mathbf{p}_0^* \mathbf{p}_0^0) &\geq \frac{\tanh(\frac{1}{2}[\mathbf{p}_0^1, \mathbf{p}_0^*])}{\sin \angle(\mathbf{p}_0^1 \mathbf{p}_0^* \mathbf{p}_0^0)}, \\ \sin^2 \angle(\mathbf{p}_0^1 \mathbf{p}_0^* \mathbf{p}_0^0) &\geq \frac{\tanh(\frac{1}{2}[\mathbf{p}_0^1, \mathbf{p}_0^*])}{\tanh([\mathbf{p}_0^1, \mathbf{p}_0^*])}. \end{aligned}$$

Since

$$\frac{\tanh(\frac{1}{2}[\mathbf{p}_0^1, \mathbf{p}_0^*])}{\tanh([\mathbf{p}_0^1, \mathbf{p}_0^*])} = \frac{1}{2} \tanh^2(\frac{1}{2}[\mathbf{p}_0^1, \mathbf{p}_0^*]) + \frac{1}{2},$$

and since $[\mathbf{p}_0^1, \mathbf{p}_0^*]$ is strictly increasing by (A.18) in the proof of Lemma 3.15, we find that

$$\frac{\tanh(\frac{1}{2}[\mathbf{p}_0^1, \mathbf{p}_0^*])}{\tanh([\mathbf{p}_0^1, \mathbf{p}_0^*])} \leq \lim_{g \rightarrow \infty} \frac{\tanh(\frac{1}{2}[\mathbf{p}_0^1, \mathbf{p}_0^*])}{\tanh([\mathbf{p}_0^1, \mathbf{p}_0^*])} \approx 0.542.$$

Furthermore, since

$$\begin{aligned} \sin^2 \angle(\mathbf{p}_0^1 \mathbf{p}_0^* \mathbf{p}_0^0) &\stackrel{(A.19)}{=} \frac{\sinh^2(\frac{1}{4} \text{sys}(\mathbb{M}_g)) \sin^2 \angle(O \mathbf{p}_0^0 \mathbf{p}_0^*)}{\sinh^2([\mathbf{p}_0^1, \mathbf{p}_0^*])}, \\ &\stackrel{(A.18)}{=} \frac{\sinh^2(\frac{1}{2}[\mathbf{p}_0^1, \mathbf{p}_0^*]) \sin^2 \angle(O \mathbf{p}_0^0 \mathbf{p}_0^*)}{\sin^2(\frac{1}{2} \angle(O \mathbf{p}_0^0 \mathbf{p}_0^*)) \sinh^2([\mathbf{p}_0^1, \mathbf{p}_0^*])}, \\ &= \frac{\cos^2(\frac{1}{2} \angle(O \mathbf{p}_0^0 \mathbf{p}_0^*))}{\cosh^2(\frac{1}{2}[\mathbf{p}_0^1, \mathbf{p}_0^*])}, \end{aligned}$$

and since $\angle(O \mathbf{p}_0^0 \mathbf{p}_0^*)$ is constant and $[\mathbf{p}_0^1, \mathbf{p}_0^*]$ strictly increasing, we see that $\sin^2 \angle(\mathbf{p}_0^1 \mathbf{p}_0^* \mathbf{p}_0^0)$ is strictly decreasing, so

$$\sin^2 \angle(\mathbf{p}_0^1 \mathbf{p}_0^* \mathbf{p}_0^0) \geq \lim_{g \rightarrow \infty} \sin^2 \angle(\mathbf{p}_0^1 \mathbf{p}_0^* \mathbf{p}_0^0) \approx 0.744.$$

From this we can conclude that

$$\sin^2 \angle(\mathbf{p}_0^1 \mathbf{p}_0^* \mathbf{p}_0^0) \geq \frac{\tanh(\frac{1}{2}[\mathbf{p}_0^1, \mathbf{p}_0^*])}{\tanh([\mathbf{p}_0^1, \mathbf{p}_0^*])}$$

holds, which by the chain of equivalent inequalities means that $[\mathbf{p}_0^*, \mathbf{p}_0^M] \geq R(\mathbf{p}_0^1, \mathbf{p}_1^1, \mathbf{p}_0^*)$. It follows that if $L(O, \mathbf{p}_0^0) \cap C(\mathbf{p}_0^1, \mathbf{p}_1^1, \mathbf{p}_0^*)$ consists of two points, then $\mathbf{p}_0^1$ is the closest of these two. This implies that $D(\mathbf{p}_0^1, \mathbf{p}_1^1, \mathbf{p}_0^*) \cap (O, \mathbf{p}_0^1) = \emptyset$. We conclude that $\mathbf{p}_0^j \notin D(\mathbf{p}_0^1, \mathbf{p}_1^1, \mathbf{p}_0^*)$ for all $j = 0, \dots, m$. By symmetry, we see that $\mathbf{p}_1^j \notin D(\mathbf{p}_0^1, \mathbf{p}_1^1, \mathbf{p}_0^*)$ for all $j = 0, \dots, m$.

Second, by the reasoning above we see that $D(\mathbf{p}_0^1, \mathbf{p}_1^1, \mathbf{p}_0^*)$ is contained in the union of the triangle $[O, \mathbf{p}_0^0, \mathbf{p}_1^0]$ and the (open) annulus

$$D(\mathbf{p}_0^*, \dots, \mathbf{p}_{4g-1}^*) \setminus (D(\mathbf{p}_0^1, \dots, \mathbf{p}_{4g-1}^1) \cup C(\mathbf{p}_0^1, \dots, \mathbf{p}_{4g-1}^1))$$

centered at O , with boundary passing through $\mathbf{p}_0^*$ on one side and through $\mathbf{p}_0^1$ on the other side. Combining this with $\mathbf{p}_k^j \notin D(\mathbf{p}_0^1, \mathbf{p}_1^1, \mathbf{p}_0^*)$ for $j = 0, \dots, m$ and $k = 0, 1$, we can immediately conclude that $\mathbf{p}_k^j \notin D(\mathbf{p}_0^1, \mathbf{p}_1^1, \mathbf{p}_0^*)$ for $j = 0, \dots, m$ and $k = 0, \dots, 4g - 1$.

4. As we have seen before, $C(\mathbf{p}_0^1, \mathbf{p}_1^1, \mathbf{p}_0^*)$ is tangent to $C(\mathbf{p}_0^*, \dots, \mathbf{p}_{4g-1}^*)$, so $D(\mathbf{p}_0^1, \mathbf{p}_1^1, \mathbf{p}_0^*)$ is contained in the interior of the $4g$ -gon $[\mathbf{p}_0^0, \dots, \mathbf{p}_{4g-1}^0]$. Therefore, $\mathbf{q}_k^j \notin D(\mathbf{p}_0^1, \mathbf{p}_1^1, \mathbf{p}_0^*)$ for $j = 1, \dots, m$ and $k = 0, \dots, 4g - 1$.

Third, consider $D(\mathbf{p}_0^m, \mathbf{p}_1^m, O)$.

1. Since $O \in C(\mathbf{p}_0^m, \mathbf{p}_1^m, O)$, we know that $O \notin D(\mathbf{p}_0^m, \mathbf{p}_1^m, O)$. Clearly, $\mathbf{v}$ is too far away from O to be inside $D(\mathbf{p}_0^m, \mathbf{p}_1^m, O)$.
2. Clearly, the points $\mathbf{p}_k^*$ for $k = 0, \dots, 4g - 1$ are too far away from O to be inside $D(\mathbf{p}_0^m, \mathbf{p}_1^m, O)$.
3. Since $D(\mathbf{p}_0^m, \mathbf{p}_1^m, O)$ is contained in the union of the triangle $[O, \mathbf{p}_0^0, \mathbf{p}_1^0]$ and the disk $D(\mathbf{p}_0^m, \mathbf{p}_1^m, \dots, \mathbf{p}_{4g-1}^m)$, we see that $\mathbf{p}_k^j \notin D(\mathbf{p}_0^m, \mathbf{p}_1^m, O)$ for all $j = 0, \dots, m$ and $k \neq 0, 1$. Since $L(O, \mathbf{p}_0^m) \cap D(\mathbf{p}_0^m, \mathbf{p}_1^m, O) = (O, \mathbf{p}_0^m)$, it follows that $\mathbf{p}_0^j \notin D(\mathbf{p}_0^m, \mathbf{p}_1^m, O)$ for $j = 0, \dots, m$. Similarly, $\mathbf{p}_1^j \notin D(\mathbf{p}_0^m, \mathbf{p}_1^m, O)$ for $j = 0, \dots, m$. Therefore, $\mathbf{p}_k^j \notin D(\mathbf{p}_0^m, \mathbf{p}_1^m, O)$ for all $j = 0, \dots, m$ and $k = 0, \dots, 4g - 1$.
4. Clearly, the points $\mathbf{q}_k^j$ for $j = 1, \dots, m$ and $k = 0, \dots, 4g - 1$ are too far away from O to be inside $D(\mathbf{p}_0^m, \mathbf{p}_1^m, O)$.

Finally, let $1 \leq j \leq n - 1$ and consider $D(\mathbf{p}_0^j, \mathbf{p}_1^j, \mathbf{p}_1^{j+1})$.

1. Clearly, $\mathbf{v}$ is too far away from the center of $D(\mathbf{p}_0^j, \mathbf{p}_1^j, \mathbf{p}_1^{j+1})$ to be inside $D(\mathbf{p}_0^j, \mathbf{p}_1^j, \mathbf{p}_1^{j+1})$. Moreover, $D(\mathbf{p}_0^j, \mathbf{p}_1^j, \mathbf{p}_1^{j+1})$ does not contain O because $L(O, \mathbf{p}_1^0) \cap D(\mathbf{p}_0^j, \mathbf{p}_1^j, \mathbf{p}_1^{j+1}) = (\mathbf{p}_1^j, \mathbf{p}_1^{j+1})$.
2. Of the set of disks $\{D(\mathbf{p}_0^j, \mathbf{p}_1^j, \mathbf{p}_1^{j+1}), j = 1, \dots, m - 1\}$, the one that is closest to $\mathbf{p}_0^*$ is $D(\mathbf{p}_0^1, \mathbf{p}_1^1, \mathbf{p}_1^2)$, i.e., if $\mathbf{p}_0^* \notin D(\mathbf{p}_0^1, \mathbf{p}_1^1, \mathbf{p}_1^2)$, then $\mathbf{p}_0^* \notin D(\mathbf{p}_0^j, \mathbf{p}_1^j, \mathbf{p}_1^{j+1})$ for $j = 1, \dots, m - 1$. Observe that $C(\mathbf{p}_0^1, \mathbf{p}_1^1, \mathbf{p}_1^2)$ and $C(\mathbf{p}_0^*, \mathbf{p}_0^1, \mathbf{p}_1^1)$ intersect in the points $\mathbf{p}_0^1, \mathbf{p}_1^1$. Since the center of $C(\mathbf{p}_0^1, \mathbf{p}_1^1, \mathbf{p}_1^2)$ is closer to O than the center of $C(\mathbf{p}_0^*, \mathbf{p}_0^1, \mathbf{p}_1^1)$, we can conclude that $\mathbf{p}_0^* \notin D(\mathbf{p}_0^1, \mathbf{p}_1^1, \mathbf{p}_1^2)$. Therefore, $\mathbf{p}_0^* \notin D(\mathbf{p}_0^j, \mathbf{p}_1^j, \mathbf{p}_1^{j+1})$ for $j = 1, \dots, m - 1$. It follows that $D(\mathbf{p}_0^j, \mathbf{p}_1^j, \mathbf{p}_1^{j+1}) \subseteq D(\mathbf{p}_0^*, \dots, \mathbf{p}_{4g-1}^*)$ for $j = 1, \dots, m - 1$, which implies that $\mathbf{p}_k^* \notin D(\mathbf{p}_0^j, \mathbf{p}_1^j, \mathbf{p}_1^{j+1})$ for $j = 1, \dots, m - 1$ and $k = 0, \dots, 4g - 1$.

3. Since $\mathbf{p}_0^j, \mathbf{p}_1^j, \mathbf{p}_0^{j+1}, \mathbf{p}_1^{j+1}$ are concircular, we see that $D(\mathbf{p}_0^j, \mathbf{p}_1^j, \mathbf{p}_1^{j+1})$ is contained in the union of the (closed) triangle $[O, \mathbf{p}_0^j, \mathbf{p}_1^j]$ and the (open) annulus

$$D(\mathbf{p}_0^j, \dots, \mathbf{p}_{4g-1}^j) \setminus (D(\mathbf{p}_0^{j+1}, \dots, \mathbf{p}_{4g-1}^{j+1}) \cup C(\mathbf{p}_0^{j+1}, \dots, \mathbf{p}_{4g-1}^{j+1}))$$

centered at O , with boundary passing through $\mathbf{p}_0^j$ on one side and through $\mathbf{p}_0^{j+1}$ on the other side. Therefore, $\mathbf{p}_k^j \notin D(\mathbf{p}_0^j, \mathbf{p}_1^j, \mathbf{p}_1^{j+1})$ for $j = 0, \dots, m$ and $k \neq 0, 1$. Furthermore, since $L(O, \mathbf{p}_0^j) \cap D(\mathbf{p}_0^j, \mathbf{p}_1^j, \mathbf{p}_1^{j+1}) = (\mathbf{p}_0^j, \mathbf{p}_0^{j+1})$, we see that $\mathbf{p}_0^j \notin D(\mathbf{p}_0^j, \mathbf{p}_1^j, \mathbf{p}_1^{j+1})$ for $j = 0, \dots, m$. Similarly, $\mathbf{p}_1^j \notin D(\mathbf{p}_0^j, \mathbf{p}_1^j, \mathbf{p}_1^{j+1})$ for $j = 0, \dots, m$. We conclude that $\mathbf{p}_k^j \notin D(\mathbf{p}_0^j, \mathbf{p}_1^j, \mathbf{p}_1^{j+1})$ for $j = 0, \dots, m$ and $k = 0, \dots, 4g - 1$.

4. Clearly, $D(\mathbf{p}_0^j, \mathbf{p}_1^j, \mathbf{p}_1^{j+1})$ is contained in the $4g$ -gon $[\mathbf{p}_0^j, \dots, \mathbf{p}_k^j, \dots, \mathbf{p}_{4g-1}^j]$, which means that $\mathbf{q}_k^j \notin D(\mathbf{p}_0^j, \mathbf{p}_1^j, \mathbf{p}_1^{j+1})$ for $j = 1, \dots, m$ and $k = 0, \dots, 4g - 1$.

Since each triangle of the infinite triangulation $\mathcal{T}$ is congruent to one of the triangles above and since the circumdisk of each of the above triangles is empty, it follows that $\mathcal{T}$ is a Delaunay triangulation. $\square$

A.3 Proofs omitted in Chapter 4

Proof. (Lemma 4.23)

Throughout the proof, we denote the set of vertices of T contained in a subset U of $\mathbb{M}$ by $V(U)$. Likewise, let $E(U, W)$ be the set of edges with one endpoint in $U \subset \mathbb{M}$ and one endpoint in $W \subset \mathbb{M}$.

Part 1. Consider the graph $G_i = (V(\Gamma_i), E(\Gamma_i, \Gamma_i))$. Let g_i be the genus of G_i , i.e., the minimal genus of a surface onto which G_i can be embedded. It is known that [59, Proposition 4.4.4]

$$g_i \geq \left\lceil \frac{e(\Gamma_i, \Gamma_i)}{6} - \frac{v(\Gamma_i)}{2} + 1 \right\rceil,$$

or, equivalently, that

$$e(\Gamma_i, \Gamma_i) \leq 6g_i + 3v(\Gamma_i) - 6. \quad (\text{A.22})$$

We will show that the embedding of G_i into $\mathbb{M}$ intersects at most $6N(N+1) + 2$ pairs of pants, which implies that $g_i \leq 3N(N+1) + 1$. Certainly, $V(\Gamma_i)$ is contained in Γ_i , which consists of at most $6N$ consecutive pairs of pants. Now, let $e \in E(\Gamma_i, \Gamma_i)$. Because the diameter of Γ_i is at most $6NM$, we know that there exists a path of length at most $6NM$ between the endpoints of e . Because e is a distance path, it follows that

$$\ell(e) \leq 6NM < 6N^2m.$$

Suppose that e intersects exactly k pairs of pants that are not contained in Γ_i and denote the farthest pair of pants that it intersects by Y^* . Here, farthest is defined

with respect to the distance along the trivalent graph L_g . Because e is an edge between vertices in Γ_i , e has to traverse at least $k - 1$ pairs of pants to reach Y^* and similarly at least $k - 1$ pairs of pants to return to Γ_i . As the length of e within each of these pairs of pants is at least m , we know

$$\ell(e) \geq 2(k - 1)m.$$

It follows that

$$2(k - 1)m < 6N^2m,$$

which implies that $k < 3N^2 + 1$. We conclude that G_i is embedded in a surface consisting of at most $(3N^2 + 1) + 6N + (3N^2 + 1) = 6N(N + 1) + 2$ pairs of pants. It follows that $g_i \leq 3N(N + 1) + 1$. Hence

$$e(\Gamma_i, \Gamma_i) \leq 6(3N(N + 1) + 1) + 3v(\Gamma_i) - 6 = 3v(\Gamma_i) + 18N(N + 1),$$

which finishes the proof.

Part 2. We consider two cases.

Case 1: there are at most $6N^2 + 2$ pairs of pants between Γ_i and Γ_{i+1} .

Consider the graph $(V(\Gamma_i \cup \Gamma_{i+1}), E(\Gamma_i \cup \Gamma_{i+1}, \Gamma_i \cup \Gamma_{i+1}))$. We have shown in Part 1 that edges in $E(\Gamma_i, \Gamma_i)$ and $E(\Gamma_{i+1}, \Gamma_{i+1})$ can traverse at most $3N^2 - 1$ pairs of pants that are not contained in Γ_i and Γ_{i+1} respectively. Therefore, $E(\Gamma_i, \Gamma_i) \cup E(\Gamma_{i+1}, \Gamma_{i+1})$ is contained in a surface consisting of at most

$$(3N^2 + 1) + 6N + (3N^2 + 1) + (3N^2 + 1) + 6N + (3N^2 + 1) = 12N^2 + 12N + 4$$

pairs of pants. With a similar argument it can be shown that $E(\Gamma_i, \Gamma_{i+1})$ is contained in this surface as well. Therefore, $(V(\Gamma_i \cup \Gamma_{i+1}), E(\Gamma_i \cup \Gamma_{i+1}, \Gamma_i \cup \Gamma_{i+1}))$ is embedded in a surface consisting of at most $12N^2 + 12N + 4$ pairs of pants. Replacing Γ_i by $\Gamma_i \cup \Gamma_{i+1}$ in Inequality (A.22) yields

$$\begin{aligned} e(\Gamma_i \cup \Gamma_{i+1}, \Gamma_i \cup \Gamma_{i+1}) &\leq 6(6N^2 + 6N + 4) + 3v(\Gamma_i \cup \Gamma_{i+1}) - 6, \\ &= 3v(\Gamma_i \cup \Gamma_{i+1}) + 36N(N + 1) + 18. \end{aligned}$$

Because $e(\Gamma_i, \Gamma_{i+1}) \leq e(\Gamma_i \cup \Gamma_{i+1}, \Gamma_i \cup \Gamma_{i+1})$, the desired inequality follows.

Case 2: there are more than $6N^2 + 2$ pairs of pants between Γ_i and Γ_{i+1} .

We show that there are integers $g_{i,1}$ and $g_{i,2}$ with

$$g_{i,1} \leq g_{i,2}, \tag{A.23}$$

such that

$$e(\Gamma_i, \Gamma_{i+1}) \leq 2v(\Gamma_i) + 2v(\Gamma_{i+1}) + 4g_{i,1} - 4 \tag{A.24}$$

and

$$e(\Gamma_i, \Gamma_{i+1}) \geq 3v(\Gamma_i) + 3v(\Gamma_{i+1}) - 3e(\Gamma_i, \Gamma_i) - 3e(\Gamma_{i+1}, \Gamma_{i+1}) + 6g_{i,2} - 6. \tag{A.25}$$

Combining (A.24) and (A.25) yields

$$e(\Gamma_i, \Gamma_{i+1}) \leq 6e(\Gamma_i, \Gamma_i) + 6e(\Gamma_{i+1}, \Gamma_{i+1}) + 12g_{i,1} - 12g_{i,2}.$$

Using the upper bound $e(\Gamma_j, \Gamma_j) \leq 3v(\Gamma_j) + 18N(N+1)$ from Part 1 for $j = i$ and $j = i+1$, together with (A.23) yields

$$e(\Gamma_i, \Gamma_{i+1}) \leq 18v(\Gamma_i \cup \Gamma_{i+1}) + 216N(N+1),$$

which is the desired inequality.

The number $g_{i,1}$ is the genus of the graph $G_{i,1}$ with edge set $E(\Gamma_i, \Gamma_{i+1})$ and vertex set $V_{i,1}$ consisting of all vertices incident to some edge in $E(\Gamma_i, \Gamma_{i+1})$. The number $g_{i,2}$ is the genus of the graph $G_{i,2}$ with edge set $E(\Gamma_i \cup \Gamma_{i+1}, \Gamma_i \cup \Gamma_{i+1})$ and vertex set $V(\Gamma_{i+1} \cup \Gamma_{i+1})$. Since $G_{i,1}$ is a subgraph of $G_{i,2}$, inequality (A.23) holds. Therefore, it remains to be proved that (A.24) holds for this value of $g_{i,1}$, and that (A.25) holds for this value of $g_{i,2}$.

To prove that (A.24) holds, we apply a result in graph theory about bipartite graphs to $G_{i,1}$. By construction, $G_{i,1}$ contains no cycle of length 3. We claim that $G_{i,1}$ is connected. Then, [59, Prop. 4.4.4, eq. 4.13]

$$e(\Gamma_i, \Gamma_{i+1}) \leq 2|V_{i,1}| + 4g_{i,1} - 4.$$

Observing that $|V_{i,1}| \leq v(\Gamma_i) + v(\Gamma_{i+1})$ yields the desired inequality.

To prove that $G_{i,1}$ is connected, consider a pair of pants Y between Γ_i and Γ_{i+1} such that there are at least $3N^2 + 1$ pairs of pants between Y and Γ_i and between Y and Γ_{i+1} . Such a pair of pants Y exists because there are more than $6N^2 + 2$ pairs of pants between Γ_i and Γ_{i+1} . Let γ be the boundary geodesic of Y such that $M \setminus \gamma$ consists of exactly two connected components. Every edge in $E(\Gamma_i, \Gamma_{i+1})$ intersects γ , because γ separates Γ_i and Γ_{i+1} . Furthermore, every edge in $E(\Gamma_i, \Gamma_{i+1})$ intersects γ *exactly once*, because the edges in $E(\Gamma_i, \Gamma_{i+1})$ and γ are geodesics and there are no hyperbolic bigons. No edge in $E(\Gamma_i, \Gamma_i) \cup E(\Gamma_{i+1}, \Gamma_{i+1})$ intersects γ , because by the reasoning in Part 1 edges in $E(\Gamma_j, \Gamma_j)$ for $j = i$ and $j = i+1$ intersect fewer than $3N^2 + 1$ pairs of pants that are not contained in Γ_j . Because T has no edges between non-adjacent clusters, it follows that the only edges of T that intersect γ are edges in $E(\Gamma_i, \Gamma_{i+1})$ and that each edge intersects γ in exactly one point. Therefore, we can write $E(\Gamma_i, \Gamma_{i+1}) = \{e_1, \dots, e_k\}$ for $k = |E(\Gamma_i, \Gamma_{i+1})|$, where the indices of the edges correspond to the order in which they intersect γ . For every $i = 1, \dots, k$, the edges e_i and e_{i+1} are contained in a triangle of T consisting of e_i, e_{i+1} and an edge in either $E(\Gamma_i, \Gamma_i)$ or $E(\Gamma_{i+1}, \Gamma_{i+1})$. Hence, e_i and e_{i+1} share an endpoint. Therefore, there is a path in $G_{i,1}$ between any endpoint of e_i and any endpoint of e_{i+1} . This implies that there is a path in $G_{i,1}$ between any endpoint of e_i and e_j for all $i, j = 1, \dots, k$. Because $V_{i,1}$ consists of the vertices incident to some edge in $E(\Gamma_i, \Gamma_{i+1})$, it follows that $G_{i,1}$ is connected. This concludes the proof of Inequality (A.24).

We continue with the proof of Inequality (A.25). Recall that the graph $G_{i,2}$ has edge set $E(\Gamma_i \cup \Gamma_{i+1}, \Gamma_i \cup \Gamma_{i+1})$ and vertex set $V(\Gamma_{i+1} \cup \Gamma_{i+1})$. The union of

the triangles of T with edges and vertices in $G_{i,2}$ define a topological surface with boundary components¹. Each boundary component consists of a finite number of edges of $G_{i,2}$. By adding a face to each boundary component, we obtain an embedding of $G_{i,2}$ in a closed surface S_i . When we speak of faces of $G_{i,2}$, we will always refer to faces with respect to the embedding of $G_{i,2}$ in S_i .

Denote the number of edges that are contained in exactly zero, one or two triangles in $G_{i,2}$ by δ_0, δ_1 and δ_2 , respectively. The total number of edges is given by $e_{i,2} = \delta_0 + \delta_1 + \delta_2$. Now, let f_Δ be the number of triangular faces of $G_{i,2}$. As the total number of faces $f_{i,2}$ is at least the number of triangular faces, we know that $f_{i,2} \geq f_\Delta$. Since $3f_\Delta = 2\delta_2 + \delta_1$, it follows that

$$3f_{i,2} \geq 2\delta_2 + \delta_1.$$

As $e_{i,2} = \delta_0 + \delta_1 + \delta_2$, we obtain

$$3f_{i,2} \geq 2(e_{i,2} - \delta_1 - \delta_0) + \delta_1 = 2e_{i,2} - \delta_1 - 2\delta_0.$$

Because

$$e_{i,2} = e(\Gamma_i, \Gamma_i) + e(\Gamma_i, \Gamma_{i+1}) + e(\Gamma_{i+1}, \Gamma_{i+1}), \quad (\text{A.26})$$

we see

$$3f_{i,2} \geq 2e(\Gamma_i, \Gamma_i) + 2e(\Gamma_i, \Gamma_{i+1}) + 2e(\Gamma_{i+1}, \Gamma_{i+1}) - \delta_1 - 2\delta_0. \quad (\text{A.27})$$

We will now bound the right-hand side from below by proving that

$$\delta_1 + 2\delta_0 \leq 2e(\Gamma_i, \Gamma_i) + 2e(\Gamma_{i+1}, \Gamma_{i+1}).$$

We claim that every edge in $E(\Gamma_i, \Gamma_{i+1})$ is part of two triangles. To prove this, let $e = (u, v) \in E(\Gamma_i, \Gamma_{i+1})$ and consider a triangle (u, v, w) in T containing e . Because every edge has its endpoints in either the same cluster or consecutive clusters, there is only an edge between u and w if $w \in \Gamma_j$ for $j = i - 1, i, i + 1$ and there is only an edge between v and w if $w \in \Gamma_j$ for $j = i, i + 1, i + 2$. Since (u, v, w) is a triangle in T , it follows that $w \in \Gamma_i \cup \Gamma_{i+1}$, so $w \in G_{i,2}$. This means that (u, v, w) is a triangle in $G_{i,2}$ and since we have chosen it arbitrarily, both triangles in T containing e are triangles in $G_{i,2}$. It follows that the edges that are contained in no triangles or exactly one triangle in $G_{i,2}$ are contained in $E(\Gamma_i, \Gamma_i) \cup E(\Gamma_{i+1}, \Gamma_{i+1})$. This means that

$$\delta_0 + \delta_1 \leq e(\Gamma_i, \Gamma_i) + e(\Gamma_{i+1}, \Gamma_{i+1}),$$

so

$$2\delta_0 + \delta_1 \leq 2e(\Gamma_i, \Gamma_i) + 2e(\Gamma_{i+1}, \Gamma_{i+1}).$$

¹In fact, these triangles define a hyperbolic surface with boundary components consisting of a finite number of geodesic segments. However, for our argument we do not use any metric properties.

Combining this upper bound with Equation (A.27) we obtain

$$3f_{i,2} \geq 2e(\Gamma_i, \Gamma_{i+1}). \quad (\text{A.28})$$

To conclude, we will look at Euler's formula for the graph $G_{i,2}$, which is given by

$$v_{i,2} - e_{i,2} + f_{i,2} = 2 - 2g'_{i,2}, \quad (\text{A.29})$$

where $g'_{i,2}$ is the genus of the embedding of $G_{i,2}$ in S_i . Because $g_{i,2}$ is the minimal genus of a surface onto which $G_{i,2}$ can be embedded, in particular $g_{i,2} \leq g'_{i,2}$. Substituting Equation (A.26) and Inequality (A.28) into Euler's formula (A.29), we obtain after some simplifications

$$e(\Gamma_i, \Gamma_{i+1}) \geq 3v(\Gamma_i \cup \Gamma_{i+1}) - 3e(\Gamma_i, \Gamma_i) - 3e(\Gamma_{i+1}, \Gamma_{i+1}) + 6g_{i,2} - 6.$$

This finishes the proof. □

Conclusion and open problems

In this thesis we studied Delaunay triangulations of hyperbolic surfaces. We considered the following three topics:

1. computing the systole of a specific class of hyperbolic surfaces,
2. describing the properties of a given class of hyperbolic surfaces that are needed to compute Delaunay triangulations of point sets on these surfaces,
3. providing upper and lower bounds for the minimal number of vertices of a Delaunay triangulation of a hyperbolic surface.

Here, we give a summary of the results and state some open problems.

In Chapter 2 we studied the systole of hyperbolic surfaces. We showed that the systole of the so-called generalized Bolza surface of genus g is given by $2 \operatorname{arccosh}(1 + 2 \cos(\frac{\pi}{2g}))$. In the proof, closed geodesics on the surface are represented by hyperbolic line segments between the sides of a fundamental polygon representing the generalized Bolza surface. The result follows by obtaining bounds on the lengths of these line segments by using straightforward hyperbolic trigonometry.

The fundamental polygon used to represent the generalized Bolza surface of genus g is a regular hyperbolic $4g$ -gon. By perturbing the vertices of this regular $4g$ -gon while keeping point-symmetric symmetry and considering the corresponding hyperbolic surfaces, we obtain an open neighborhood of the generalized Bolza surfaces in the set of hyperelliptic surfaces. For sufficiently small perturbations, the geometry of the resulting polygons is similar to the geometry of the regular $4g$ -gon. In particular, we proved that for perturbations of size $O((\log(g))^{-1})$ the method for computing the systole of generalized Bolza surfaces works for the hyperelliptic surfaces in the corresponding neighborhood as well.

Motivated by the computation of the systole of generalized Bolza surfaces and hyperelliptic surfaces, we looked at the so-called word length of systoles. It is known that every closed geodesic on a closed hyperbolic surface corresponds to an element of the corresponding Fuchsian group, up to conjugacy. The minimal number of generators of a group that need to be multiplied to obtain a specific element of that group is called the word length of that element. The elements of the Fuchsian group corresponding to a systole of a generalized Bolza surface

(or hyperelliptic surfaces in some neighborhood) are products of precisely two of the side-pairing transformations. This raised the question whether there exists a general upper bound for the word length of the elements corresponding to systoles of a given hyperbolic surface, but we proved that there is no such bound.

Our study of systoles has multiple directions for future research. First, we have shown that the method for computing the systole of the generalized Bolza surfaces can be applied to hyperelliptic surfaces in some neighborhood as well, but our derivation of an admissible size of the perturbations is rather indirect. One can ask whether there is a more explicit upper bound on the admissible size of perturbations. Moreover, a next question could be whether the same sort of reasoning can be used for hyperbolic surfaces outside this neighborhood as well, for example by giving a more detailed description of the geometry of the fundamental polygon. In this way, we might obtain a better understanding of the behavior of the systole as function on the Teichmüller space of hyperbolic surfaces of a given genus.

Second, even though we have shown that in general there is no upper bound for the word length of a systole of a hyperbolic surface, the method used in our proof leads to a set of generators that belong to a very elongated fundamental domain. A follow-up question is whether there is still no upper bound when we pose additional conditions on the fundamental domain, for example being a Dirichlet region.

In Chapter 3 we described the properties of the generalized Bolza surfaces that allow Bowyer's algorithm to be applied to these surfaces. First, to be able to check whether a given point set on a generalized Bolza surface satisfies the validity condition, it is necessary to know the value of the systole, which was already computed in Chapter 2. Using a similar argument as in the computation of the value of the systole, we proved that during the execution of the algorithm it is sufficient to consider a small set of copies of the input points, instead of all (infinitely many) copies in the orbit space, which cannot be done in practice. Moreover, we created several algorithms to construct a dummy point set with which Bowyer's algorithm can be initialized and analyzed the resulting number of points of each. In a more general setting, we proved upper and lower bounds for the cardinality of a dummy point set for arbitrary hyperbolic surfaces. Apart from describing point sets that satisfy the validity condition known from the literature, we also briefly discussed a local validity condition, that is based on the injectivity radius of points of the hyperbolic surface instead of the systole. Finally, we looked at numerical issues for the implementation of the described algorithm and stated a bound for the degree of the predicates used in the computations.

A natural (but challenging) open problem is to generalize Bowyer's algorithm to arbitrary hyperbolic surfaces. Since our extension of Bowyer's algorithm relies on the validity condition, for which we need to know the value of the systole, and since the value of the systole of a hyperbolic surface is not known in general, a first step could be to consider a subclass of hyperbolic surfaces. We have already looked at hyperelliptic surfaces in some neighborhood of the generalized Bolza

surfaces in Chapter 2, so we could start by generalizing the algorithm to these surfaces.

In Chapter 4 we looked at the minimal number of vertices of a Delaunay triangulation of a hyperbolic surface. By subdividing an arbitrary hyperbolic surface into its ‘thick’ and ‘thin’ parts, we were able to show that any hyperbolic surface of genus g admits a Delaunay triangulation with at most $151g$ vertices such that its edges are distance paths. Similarly, there exists a family of hyperbolic surfaces such that the number of vertices of any Delaunay triangulation where edges are distance paths grows like $\Omega(g)$. The examples that we used to prove the latter statement are geometrically quite simple, as they are made by gluing pairs of pants with boundary lengths chosen in some closed interval in a pattern that somewhat resembles a ‘line’. Finally, one can wonder whether all hyperbolic surfaces satisfy this $\Omega(g)$ bound, but we have shown that this is not the case. Namely, using the Ringel-Youngs construction we constructed a family of surfaces for which there exist Delaunay triangulations with $\Theta(\sqrt{g})$ vertices.

From the moduli space point of view, we have a function from moduli space to the set of natural numbers assigning to each hyperbolic surface the minimal number of vertices of any of its Delaunay triangulations. Even though the results that we obtained so far give a good understanding of the extremal values of this function, a further question is how this function behaves in different ‘regions’ of moduli space. For example, one can wonder how the number of vertices behaves for a random hyperbolic surface, chosen with respect to a natural probability measure on moduli space.

Moreover, even though we found an example of a family of hyperbolic surfaces attaining the general lower bound $\Omega(\sqrt{g})$, the geometry of the resulting surfaces is quite special. It is not immediately clear whether in a small neighborhood of these surfaces in moduli space the minimal number of vertices is still $\Theta(\sqrt{g})$.

Samenvatting

Triangulaties zijn een veelvuldig bestudeerd onderwerp binnen de computationele meetkunde. Een *triangulatie* van een oppervlak is een opdeling van dat oppervlak in driehoeken, een van de meest eenvoudige wiskundige vormen. Het opdelen van een oppervlak in driehoeken maakt het mogelijk om computeralgoritmes toe te passen en op deze manier de eigenschappen van het oppervlak te analyseren.

Delaunay triangulaties zijn een specifiek soort triangulaties die vanwege hun speciale eigenschappen geschikt zijn voor vele toepassingen, bijvoorbeeld het modelleren van hoogteverschillen van een bepaald terrein. Oorspronkelijk werden Delaunay triangulaties bestudeerd voor verzamelingen van punten in het tweedimensionale (Euclidische) vlak, maar ze kunnen ook in andere soorten ruimtes gedefinieerd worden. Zo zijn er in de afgelopen jaren verschillende algoritmes om Delaunay triangulaties te construeren generaliseerd naar *hyperbolische ruimtes*.

Hyperbolische meetkunde kwam op in de eerste helft van de negentiende eeuw. Een model dat vaak gebruikt wordt om het hyperbolische vlak weer te geven is het *Poincaré schijfmodel*. Hierin wordt het hele hyperbolische vlak gerepresenteerd door een schijf, zoals in Figuur 1. We hebben in Figuur 1 ook drie hyperbolische lijnen getekend. Hoewel twee van deze hyperbolische lijnen vanuit het Euclidische perspectief waaraan we gewend zijn gekromd lijken, zijn ze het equivalent van de rechte lijnen in het Euclidische vlak. Dit illustreert dat het hyperbolische vlak overal negatief gekromd is, in tegenstelling tot het Euclidische vlak dat niet gekromd is. De eigenschappen van het hyperbolische vlak zorgen ervoor dat het gebruikt kan worden om vormen of structuren te beschrijven die, intuïtief gesproken, niet in het Euclidische vlak “plat neergelegd” kunnen worden.

In dit proefschrift bestuderen we Delaunay triangulaties van *hyperbolische oppervlakken*. Hyperbolische oppervlakken zijn oppervlakken die lokaal op het hyperbolische vlak “lijken”. Ze hebben toepassingen in diverse wetenschapsgebieden, zoals materiaalwetenschappen, cosmologie, wiskundige neurologie en quantumchaostheorie.

In eerder werk is *Bowyers incrementele algoritme*, dat oorspronkelijk ontwikkeld is om Delaunay triangulaties in het Euclidische vlak te construeren, generaliseerd naar hyperbolische oppervlakken. Er is een implementatie van dit algoritme voor het Bolzaoppervlak, een specifiek hyperbolisch oppervlak.

Een voorwaarde voor de juiste werking van het algoritme is dat de triangulatie

Figuur 1: Het Poincaré schijfmodel van het hyperbolische vlak met drie hyperbolische lijnen L_1 , L_2 en L_3 .

te allen tijde *simpliciaal* is tijdens het uitvoeren. Dit betekent dat geen van de zijden van de driehoeken in de triangulatie hetzelfde begin- als eindpunt mag hebben, en dat tussen elk paar hoekpunten van de triangulatie maar één zijde van een driehoek mag liggen. Aan deze voorwaarde is voldaan wanneer de verzameling van hoekpunten van de triangulatie voldoende groot en gelijkmatig over het oppervlak verdeeld is.

Om te controleren of de verzameling van hoekpunten inderdaad voldoende groot en gelijkmatig is, gebruiken we een parameter genaamd de *systole*. Dit is de lengte van de kortste gesloten kromme op het oppervlak die niet samengetrokken kan worden. De waarde van de systole is in het algemeen, dat wil zeggen voor een willekeurig hyperbolisch oppervlak, niet bekend. Er bestaat een algoritme om de systole van een gegeven hyperbolisch oppervlak uit te rekenen, maar het is niet duidelijk of dit efficiënt genoeg is om in de praktijk te gebruiken.

We hebben de resultaten van dit proefschrift onderverdeeld in drie hoofdstukken.

In hoofdstuk 2 berekenen we de waarde van de systole van de *gegeneraliseerde Bolzaoppervlakken*, een specifieke klasse van hyperbolische oppervlakken. Verder bewijzen we dat de methode die we gebruiken om de systole van deze klasse van oppervlakken uit te rekenen ook gebruikt kan worden voor andere oppervlakken, zolang ze niet te veel verschillen van de gegeneraliseerde Bolzaoppervlakken.

In hoofdstuk 3 beschrijven we de eigenschappen van de gegeneraliseerde Bolzaoppervlakken die nodig zijn om Bowyers incrementele algoritme te kunnen toepassen op deze oppervlakken. Een van de karakteristieke kenmerken van Bowyers algoritme is dat het een *incrementeel algoritme* is: het begint met een Delaunay triangulatie van een klein aantal hoekpunten, voegt vervolgens de overige hoekpunten een voor een toe en past na elke toevoeging van een hoekpunt de triangulatie aan. Om ervoor te zorgen dat de triangulatie ook in het begin simpliciaal is, introduceren we meerdere algoritmes die een verzameling van hoekpunten construeren

die als startpunt kunnen dienen voor Bowyers incrementele algoritme.

In hoofdstuk 4 behandelen we Delaunay triangulaties van hyperbolische oppervlakken met een minimaal aantal hoekpunten. We leiden een bovengrens af voor dit minimale aantal hoekpunten, dat wil zeggen dat elk hyperbolisch oppervlak een Delaunay triangulatie heeft waarbij het aantal hoekpunten niet groter is dan deze bovengrens. Verder geven we voorbeelden van oppervlakken waarvoor het minimale aantal hoekpunten van een Delaunay triangulatie dezelfde orde van grootte heeft als de bovengrens. De orde van grootte van de bovengrens wordt gegeven in termen van het *geslacht* van het oppervlak, dat ruwweg het aantal “gaten” in het oppervlak weergeeft. Zo heeft het oppervlak in Figuur 2 geslacht 3.

Figuur 2: Oppervlak van geslacht 3 onderverdeeld in 4 broeken.

Een hulpmiddel bij de analyse van het minimale aantal hoekpunten van Delaunay triangulaties van hyperbolische oppervlakken is de onderverdeling van hyperbolische oppervlakken in zogenaamde *broeken*. Een broek is een oppervlak met drie openingen: twee voor de benen en een voor het middel. Het oppervlak in Figuur 2 hebben we onderverdeeld in 4 broeken. Het gaat hierbij alleen om het aantal openingen; het is voor ons niet van belang of de broeken ook een vorm hebben zodat ze daadwerkelijk aangetrokken kunnen worden.

Naast de bovengrens leiden we ook een ondergrens af. Tevens geven we voorbeelden van oppervlakken waarvoor het minimale aantal hoekpunten van een Delaunay triangulatie dezelfde orde van grootte heeft als deze ondergrens.

Acknowledgements

First and foremost I would like to thank my supervisors.

Gert Vegter, thank you for being my daily supervisor. I think that the fact that I continued with a PhD after I did my bachelor's and master's theses under your supervision speaks to the quality of your guidance. I fondly remember the discussions about research we had during which we filled many a whiteboard. Over the years I have learned a lot from you, not only as a researcher, but also as a person.

Alef Sterk, thank you for being my second supervisor. In particular, I would like to thank you for your feedback on the introduction, which, as you probably know, is often one of the hardest parts to write.

I would also like to express my gratitude to the members of my assessment committee Xavier Goaoc, Tobias Müller and Jean-Marc Schlenker for the time they invested in my thesis and defence.

I am grateful to Monique Teillaud and Iordan Iordanov for the collaboration leading to the results in Chapters 2.2 and 3, and to Hugo Parlier for the collaboration leading to the results in Chapter 4. I also wish to thank the other researchers that I met at INRIA Nancy and the University of Luxembourg with whom I had many interesting discussions on hyperbolic surface theory and discrete and computational geometry.

I would like to thank my office mates Nikolay, Réka and Federico for the enjoyable times we had in our office. Special thanks to Nikolay for being one of my paranymphs and for helping with the design of this thesis by providing various templates.

Furthermore, I would like to thank my friends and colleagues from the University of Groningen. In particular, I wish to thank Anna, Ben, Bohuan, Chengtao, Deepan, Dijs, Dirk, Eduardo, Eric, Gabriel, Georg, Giovanni, Hildeberto, Jaap, Jian, Jorge, Kevin, Lorenzo, Mahdi, Markus, Martine, Paul, Pieter, Riccardo, Teun and Yongjiao.

Ten slotte wil ik mijn familie bedanken. Zonder de stabiele basis die zij voor mij vormen denk ik niet dat ik zo ver had kunnen komen.

Bibliography

- [1] 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):33513, 2005.
- [2] H. Akrouit. Un processus effectif de détermination des systoles pour les surfaces hyperboliques. *Geometriae Dedicata*, 121(1):1–8, 2006. doi:10.1007/s10711-006-9076-x.
- [3] M. A. Armstrong. *Basic Topology*. Undergraduate Texts in Mathematics. Springer New York, 2013.
- [4] R. Aurich, E. B. Bogomolny, and F. Steiner. Periodic orbits on the regular hyperbolic octagon. *Physica D: Nonlinear Phenomena*, 48(1):91–101, 1991. doi:10.1016/0167-2789(91)90053-C.
- [5] R. Aurich and F. Steiner. On the periodic orbits of a strongly chaotic system. *Physica D: Nonlinear Phenomena*, 32(3):451–460, 1988. doi:10.1016/0167-2789(88)90068-1.
- [6] Agnès Bachelot-Motet. Wave computation on the hyperbolic double doughnut. *Journal of Computational Mathematics*, pages 790–806, 2010.
- [7] Sheng Bai, Yue Gao, and Shicheng Wang. Systoles of hyperbolic surfaces with big cyclic symmetry. *Science China Mathematics*, 64(2):421–442, 2021. doi:10.1007/s11425-019-1655-8.
- [8] N.L. Balazs and A. Voros. Chaos on the pseudosphere. *Physics reports*, 143(3):109–240, 1986.
- [9] Christophe Bavard. La systole des surfaces hyperelliptiques. Technical report, École Normale Supérieure Lyon, 1992.
- [10] Alan F. Beardon. *The geometry of discrete groups*. Springer-Verlag New York, 1 edition, 1983. doi:10.1007/978-1-4612-1146-4.
- [11] Mikhail V. Belolipetsky and Scott A. Thomson. Systoles of hyperbolic manifolds. *Algebraic & Geometric Topology*, 11(3):1455–1469, 2011.

- [12] Lipman Bers. An inequality for Riemann surfaces. In *Differential geometry and complex analysis*, pages 87–93. Springer, Berlin, 1985.
- [13] Mikhail Bogdanov. *Delaunay triangulations of spaces of constant negative curvature*. PhD thesis, Université Nice Sophia Antipolis, 2013.
- [14] Mikhail Bogdanov, Olivier Devillers, and Monique Teillaud. Hyperbolic Delaunay complexes and Voronoi diagrams made practical. *Journal of Computational Geometry*, 5:56–85, 2014. doi:10.20382/jocg.v5i1a4.
- [15] Mikhail Bogdanov and Monique Teillaud. Delaunay triangulations and cycles on closed hyperbolic surfaces. Technical Report RR-8434, INRIA, 2013.
- [16] Mikhail Bogdanov, Monique Teillaud, and Gert Vegter. Delaunay triangulations on orientable surfaces of low genus. In *Proceedings of the Thirty-second International Symposium on Computational Geometry*, pages 20:1–20:15, 2016. doi:10.4230/LIPIcs.SoCG.2016.20.
- [17] Oskar Bolza. On Binary Sextics with Linear Transformations into Themselves. *American Journal of Mathematics*, 10:47–70, 1887. doi:10.2307/2369402.
- [18] Roberto Bonola. *Non-Euclidean geometry: A critical and historical study of its development*. Courier Corporation, 1955.
- [19] A. Bowyer. Computing Dirichlet tessellations. *The Computer Journal*, 24(2):162–166, 1981. doi:10.1093/comjnl/24.2.162.
- [20] P. Buser. *Geometry and Spectra of Compact Riemann Surfaces*. Progress in Mathematics Series. Birkhäuser, 1992.
- [21] P Buser and P 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. doi:10.1007/BF01232233.
- [22] Manuel Caroli. *Triangulating Point Sets in Orbit Spaces*. PhD thesis, Université de Nice Sophia Antipolis, 2010.
- [23] Manuel Caroli and Monique Teillaud. 3D periodic triangulations. In *CGAL User and Reference Manual*. CGAL Editorial Board, 3.5 edition, 2009. URL: <http://doc.cgal.org/latest/Manual/packages.html#PkgPeriodic3Triangulation3Summary>.
- [24] Manuel Caroli and Monique Teillaud. Delaunay triangulations of closed Euclidean d-orbifolds. *Discrete & Computational Geometry*, 55(4):827–853, 2016. URL: <https://hal.inria.fr/hal-01294409>, doi:10.1007/s00454-016-9782-6.

- [25] P. Chossat, G. Faye, and O. Faugeras. Bifurcation of hyperbolic planforms. *Journal of Nonlinear Science*, 21:465–498, 2011. doi:10.1007/s00332-010-9089-3.
- [26] Pascal Chossat and Olivier Faugeras. Hyperbolic planforms in relation to visual edges and textures perception. *Plos Computational Biology*, 5(12):e1000625, 2009.
- [27] Mark De Berg, Marc Van Kreveld, Mark Overmars, and Otfried Schwarzkopf. *Computational Geometry: Algorithms and Applications*. Springer, 1997.
- [28] M. Dehn. Transformation der Kurven auf zweiseitigen Flächen. *Mathematische Annalen*, 72(3):413–421, 1912. doi:10.1007/BF01456725.
- [29] Vincent Despré, Benedikt Kolbe, and Monique Teillaud. Half-minimizers and Delaunay triangulations on closed hyperbolic surfaces. Technical report, INRIA, 2020. URL: <https://hal.inria.fr/hal-03045921>.
- [30] Vincent Despré, Jean-Marc Schlenker, and Monique Teillaud. Flipping geometric triangulations on hyperbolic surfaces. In *Proceedings 36th International Symposium on Computational Geometry*, pages 35:1–35:16, 2020. doi:10.4230/LIPIcs.SocG.2020.35.
- [31] Olivier Devillers, Sylvain Pion, and Monique Teillaud. Walking in a triangulation. *International Journal of Foundations of Computer Science*, 13:181–199, 2002. URL: <https://hal.inria.fr/inria-00102194>, doi:10.1142/S0129054102001047.
- [32] Olivier Devillers and Monique Teillaud. Perturbations for Delaunay and weighted Delaunay 3D Triangulations. *Computational Geometry: Theory and Applications*, 44:160–168, 2011. URL: <http://hal.inria.fr/inria-00560388/>, doi:10.1016/j.comgeo.2010.09.010.
- [33] Matthijs Ebbens. Delaunay triangulations on hyperbolic surfaces. Master’s thesis, University of Groningen, 2017. URL: <http://fse.studenttheses.ub.rug.nl/id/eprint/15727>.
- [34] Hershel M. Farkas and Irwin Kra. Riemann surfaces. In *Riemann surfaces*, pages 9–31. Springer, 1992.
- [35] István Fáry. On straight-line representation of planar graphs. *Acta scientiarum mathematicarum*, 11(229-233):2, 1948.
- [36] 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.

- [37] William Fulton. *Algebraic topology: a first course*, volume 153 of *Graduate Texts in Mathematics*. Springer-Verlag, 1995.
- [38] Martin Greendlinger. Dehn’s algorithm for the word problem. *Communications on Pure and Applied Mathematics*, 13(1):67–83, 1960. doi:10.1002/cpa.3160130108.
- [39] Larry Guth, Hugo Parlier, and Robert Young. Pants decompositions of random surfaces. *Geom. Funct. Anal.*, 21(5):1069–1090, 2011.
- [40] Thomas Little Heath et al. *The thirteen books of Euclid’s Elements*. Courier Corporation, 1956.
- [41] D. Hilbert. Über Flächen von konstanter Gaußscher Krümmung. *Transactions of the American Mathematical Society*, 2:87–99, 1901.
- [42] Alfredo Hubard, Vojtěch Kaluža, Arnaud De Mesmay, and Martin Tancer. Shortest path embeddings of graphs on surfaces. *Discrete & Computational Geometry*, 58(4):921–945, 2017.
- [43] Yoichi Imayoshi and Masahiko Taniguchi. *An introduction to Teichmüller spaces*. Springer-Verlag, 2012. doi:10.1007/978-4-431-68174-8.
- [44] Iordan Iordanov. *Delaunay triangulations of a family of symmetric hyperbolic surfaces in practice*. PhD thesis, Université de Lorraine, 2019. URL: <https://tel.archives-ouvertes.fr/tel-02072155>.
- [45] Iordan Iordanov and Monique Teillaud. Implementing Delaunay triangulations of the Bolza surface. In *Proceedings of the Thirty-third International Symposium on Computational Geometry*, pages 44:1–44:15, 2017. doi:10.4230/LIPIcs.SoCG.2017.44.
- [46] Iordan Iordanov and Monique Teillaud. 2D periodic hyperbolic triangulations. In *CGAL User and Reference Manual*. CGAL Editorial Board, 4.14 edition, 2019. URL: <https://doc.cgal.org/latest/Manual/packages.html#PkgPeriodic4HyperbolicTriangulation2>.
- [47] Clément Jamin, Sylvain Pion, and Monique Teillaud. 3D triangulations. In *CGAL User and Reference Manual*. CGAL Editorial Board. URL: <https://doc.cgal.org/latest/Manual/packages.html#PkgTriangulation3>.
- [48] Mark Jungerman and Gerhard Ringel. Minimal triangulations on orientable surfaces. *Acta Mathematica*, 145(1):121–154, 1980.
- [49] Svetlana Katok. *Fuchsian groups*. University of Chicago Press, 1992.

- [50] Mikhail G. Katz, Mary Schaps, and Uzi Vishne. Logarithmic growth of systole of arithmetic riemann surfaces along congruence subgroups. *J. Differential Geom.*, 76(3):399–422, 07 2007. URL: <http://projecteuclid.org/euclid.jdg/1180135693>.
- [51] William C. Keel. *The road to galaxy formation*. Springer-Verlag, 2007.
- [52] Charles L. Lawson. Software for C^1 surface interpolation. In *Symposium on Mathematical Software*, 1977. NASA Technical Report JPL-PUBL-77-30. URL: <https://ntrs.nasa.gov/citations/19770025881>.
- [53] Greg Leibon and David Letscher. Delaunay triangulations and Voronoi diagrams for Riemannian manifolds. In *Proceedings of the Sixteenth International Symposium on Computational Geometry*, 2000.
- [54] T Long Calvin. Elementary introduction to number theory. *Lexington: DC Heath and Company, LCCN*, 77171950, 1972.
- [55] Shotaro Makisumi. A note on Riemann surfaces of large systole, 2012. arXiv:1206.2965 [math.DG].
- [56] John W Milnor. Hyperbolic geometry: the first 150 years. *Bulletin of the American Mathematical Society*, 6(1):9–24, 1982.
- [57] Maryam Mirzakhani. Growth of Weil-Petersson volumes and random hyperbolic surfaces of large genus. *J. Differential Geom.*, 94(2):267–300, 2013.
- [58] Maarten Moesen. *Modeling of the geometry and mechanical behavior of bone scaffolds*. PhD thesis, Katholieke Universiteit Leuven, 2009.
- [59] Bojan Mohar and Carsten Thomassen. *Graphs on surfaces*, volume 10. JHU Press, 2001.
- [60] J. R. Munkres. *Elements of Algebraic Topology*. Advanced book classics. Perseus Books, 1984.
- [61] Marjatta Näätänen. Regular n-gons and Fuchsian groups. *Annales Academiae Scientiarum Fennicae, Series A I Mathematica*, 7:291–300, 1982. doi:10.5186/aasfm.1982.0724.
- [62] Georg Osang, Mael Rouxel-Labbé, and Monique Teillaud. Generalizing CGAL periodic Delaunay triangulations. In *Proceedings 28th European Symposium on Algorithms*, pages 75:1–75:17, 2020. (Best Paper Award - Track B: Engineering and Applications). doi:10.4230/LIPIcs.ESA.2020.75.
- [63] Hugo Parlier and Camille Petit. Chromatic numbers of hyperbolic surfaces. *Indiana University Mathematics Journal*, pages 1401–1423, 2016.

- [64] John Ratcliffe. *Foundations of hyperbolic manifolds*, volume 149 of *Graduate Texts in Mathematics*. Springer-Verlag, 2006.
- [65] Gerhard Ringel and John W.T. Youngs. Solution of the Heawood map-coloring problem. *Proceedings of the National Academy of Sciences*, 60(2):438–445, 1968.
- [66] Jim Ruppert. A Delaunay refinement algorithm for quality 2-dimensional mesh generation. *Journal of algorithms*, 18(3):548–585, 1995. doi:10.1006/jagm.1995.1021.
- [67] Paul Schmutz Schaller. Teichmüller space and fundamental domains of Fuchsian groups. *L’Enseignement Mathématique*, 45:169–187, 1999.
- [68] Charles Schmit. Quantum and classical properties of some billiards on the hyperbolic plane. *Chaos and Quantum Physics*, pages 335–369, 1991.
- [69] Mika Seppälä and Tuomas Sorvali. *Geometry of Riemann Surfaces and Teichmüller Spaces*. North-Holland Mathematics Studies. Elsevier Science, 2011.
- [70] John Stillwell. *Geometry of surfaces*. Springer-Verlag, 1992.
- [71] Roberto Tamassia. *Handbook of graph drawing and visualization*. Chapman and Hall/CRC, 2013.
- [72] Charles M. Terry, Lloyd R. Welch, and John W.T. Youngs. The genus of K_{12s} . *Journal of Combinatorial Theory*, 2(1):43–60, 1967.
- [73] Weston Ungemach. A bound on length isospectral families of hyperbolic surfaces, 2013. arXiv:1310.8276 [math.MG].
- [74] Charles Walkden. Hyperbolic Geometry. *Lecture notes*, 2012.
- [75] D. F. Watson. Computing the n-dimensional Delaunay tessellation with application to Voronoi polytopes. *The Computer Journal*, 24(2):167–172, 1981. doi:10.1093/comjnl/24.2.167.
- [76] C. K. Yap and T. Dubé. The exact computation paradigm. In D.-Z. Du and F. K. Hwang, editors, *Computing in Euclidean Geometry*, volume 4 of *Lecture Notes Series on Computing*, pages 452–492. World Scientific, Singapore, 2nd edition, 1995. doi:10.1142/9789812831699_0011.
- [77] Chee Yap *et al.* The CORE library project. URL: http://cs.nyu.edu/exact/core_pages/intro.html.