



Delaunay triangulations of generalized Bolza surfaces

Matthijs Ebbens, Iordan Iordanov, Monique Teillaud, Gert Vegter

► To cite this version:

Matthijs Ebbens, Iordan Iordanov, Monique Teillaud, Gert Vegter. Delaunay triangulations of generalized Bolza surfaces. 2020. hal-03080125v1

HAL Id: hal-03080125

<https://inria.hal.science/hal-03080125v1>

Preprint submitted on 17 Dec 2020 (v1), last revised 11 Mar 2021 (v2)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Delaunay triangulations of generalized Bolza surfaces ^{*}

Matthijs Ebbens[†] Iordan Iordanov[‡] Monique Teillaud[§] Gert Vegter[¶]

December 17, 2020

Abstract

The Bolza surface can be seen as the quotient of the hyperbolic plane under the action of the group generated by the hyperbolic isometries identifying opposite edges of the regular octagon in the Poincaré disk that is centered at the origin. We consider *generalized* Bolza surfaces \mathbb{M}_g , where the octagon is replaced by the regular $4g$ -gon, leading to a genus g surface. We propose an extension of Bowyer’s algorithm to these surfaces. In particular, we compute the value of the systole of \mathbb{M}_g ; we also propose algorithms computing sets of points on \mathbb{M}_g whose Delaunay triangulation is a simplicial complex.

1 Introduction

The well-known incremental algorithm that computes Delaunay triangulations using edge flips in the Euclidean plane [25] has recently been proved to generalize on hyperbolic surfaces [13]. However, the experience gained in the CGAL project for many years has shown that Bowyer’s algorithm [7] 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 [9] and for the Bolza surface [23]. To the best of our knowledge, the latter package is the only available software for a hyperbolic surface.

In this paper, we study the extension of this approach to what we call *generalized Bolza surfaces*. A closed orientable hyperbolic surface \mathbb{M} is isometric to a quotient \mathbb{D}/Γ , where Γ is a discrete group of orientation preserving isometries acting on the hyperbolic plane, represented here as the Poincaré disk \mathbb{D} . See Section 2 for some mathematical background on the hyperbolic plane and hyperbolic surfaces. The universal cover of such a surface is the hyperbolic plane, with associated projection map $\pi : \mathbb{D} \rightarrow \mathbb{D}/\Gamma$. The *generalized Bolza group* Γ_g , $g \geq 2$, is the (discrete) group generated by the orientation preserving isometries that pair opposite sides of the regular $4g$ -gon D_g , centered at the origin and with angle sum 2π (unique up to rotations). The *generalized Bolza surface* \mathbb{M}_g , of genus g , is defined as the hyperbolic surface \mathbb{D}/Γ_g , with projection map $\pi_g : \mathbb{D} \rightarrow \mathbb{D}/\Gamma_g$. In particular, \mathbb{M}_2 is the usual Bolza surface.

We denote by $\text{sys}(\mathbb{M})$ the systole of a closed hyperbolic surface \mathbb{M} , i.e., the length of a shortest non-contractible curve on \mathbb{M} , and, for a set of points $\mathcal{Q} \subset \mathbb{M}$, by $\delta(\mathcal{Q})$ the diameter of the largest disks in \mathbb{D} that do not contain any point in $\pi^{-1}(\mathcal{Q})$. In Section 3 (Proposition 3), we will recall the following *validity condition* [6]: If \mathcal{Q} denotes a set of points on the surface \mathbb{M} such that

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

^{*}This work was partially supported by the grant(s) ANR-17-CE40-0033 of the French National Research Agency ANR (project SoS) and INTER/ANR/16/11554412/SoS of the Luxembourg National Research fund FNR.

[†]Bernoulli Institute for Mathematics, Computer Science and Artificial Intelligence, University of Groningen, Netherlands. y.m.ebbens@rug.nl

[‡]Université de Lorraine, CNRS, Inria, LORIA, F-54000 Nancy, France. i.m.iordanov@gmail.com

[§]Université de Lorraine, CNRS, Inria, LORIA, F-54000 Nancy, France. Monique.Teillaud@inria.fr

[¶]Bernoulli Institute for Mathematics, Computer Science and Artificial Intelligence, University of Groningen, Netherlands. g.vegter@rug.nl

then Bowyer’s algorithm can be extended to the computation of the Delaunay triangulation of any set of points \mathcal{S} on \mathbb{M} containing \mathcal{Q} . Before actually inserting the input points, a preprocessing step consists in computing the Delaunay triangulation of an appropriate (but small) set \mathcal{Q} satisfying this validity condition; following the terminology of previous papers [10, 6], we will 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 [22, 23].

Other practical approaches for (flat) quotient spaces start by computing in a finite-sheeted covering space [10] or in the universal covering space [27], 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 [15].

Results.

We describe the extension of Bowyer’s algorithm to the case of a generalized Bolza surface \mathbb{M}_g in Section 3, and we prove bounds on the number of necessary dummy points to satisfy the validity condition (Propositions 5 and 6 in Section 3.3), yielding the following result:

Theorem 1. *The number of dummy points that must be added on \mathbb{M}_g to satisfy the validity condition (5) grows as $\Theta(g)$.*

In Section 4, we give an explicit value for the systole of \mathbb{M}_g :

Theorem 2. *The systole of the surface \mathbb{M}_g is given by ς_g , where ς_g is defined as*

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

This is a generalization of a result of Aurich and Steiner [2] who derived the identity $\cosh \frac{1}{2} \operatorname{sys}(\mathbb{M}_2) = 1 + \sqrt{2}$ for the systole of the Bolza surface ($g = 2$), with a method that is quite different from our proof.

Then, in Section 5, we propose two algorithms to compute dummy points. The first algorithm uses the well-known Delaunay refinement algorithm for mesh generation [28]. Using a packing argument, we prove that it provides an asymptotically optimal number of dummy points (Theorem 18). The second algorithm yields a symmetric dummy point set, at the expense of a slightly larger output size $\Theta(g \log g)$ (Theorem 19); this symmetry may be interesting for some applications [11]. The two algorithms have been implemented and we quickly present results for small genus $g = 2$ and 3.

Finally, in Section 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 computations and present experimental results.

2 Mathematical preliminaries

2.1 The Poincaré disk

The model of the hyperbolic plane we use is the Poincaré disk, the open unit disk \mathbb{D} in the complex plane equipped with a Riemannian metric of constant Gaussian curvature $K = -1$ [4]. The Euclidean boundary \mathbb{D}_∞ of the unit disk consists of the points at infinity or *ideal points* of the hyperbolic plane (which do not belong to \mathbb{D}). 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}_∞

orthogonally (see the leftmost frame of Figure 1). The geodesic segment $[z, w]$ between points $z, w \in \mathbb{D}$ is the (unique) shortest curve connecting z and w . 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.

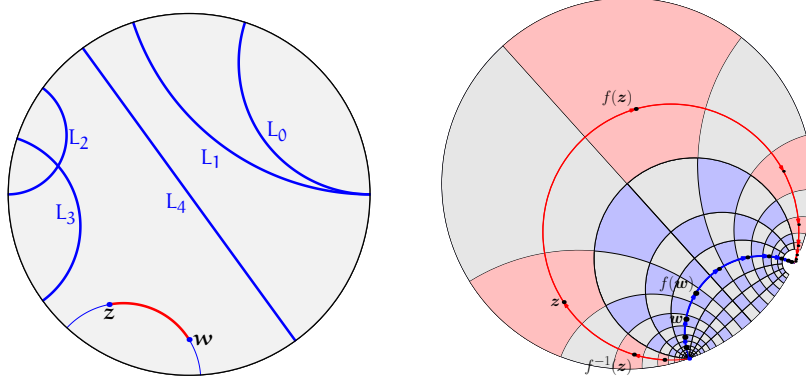


Figure 1: Left: the Poincaré disk model \mathbb{D} of the hyperbolic plane, with some geodesics. The boundary \mathbb{D}_∞ does not belong to \mathbb{D} , but consists of ideal points of \mathbb{D} . Geodesics L_0 and L_1 are parallel (have an ideal point in common), L_2 and L_3 are intersecting and 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}_∞ 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 $f(z)$.

We only consider orientation-preserving isometries of \mathbb{D} , called *isometries* from now on, which are linear fractional transformation of the form

$$f(z) = \frac{az + b}{\bar{b}z + \bar{a}}, \quad (1)$$

with $a, b \in \mathbb{C}$ such that $|a|^2 - |b|^2 = 1$. By representing isometries of the form (1) by either of the two matrices

$$\pm \begin{pmatrix} a & b \\ \bar{b} & \bar{a} \end{pmatrix}, \quad (2)$$

with $|a|^2 - |b|^2 = 1$, composition of isometries corresponds to multiplication of either of their representing matrices. The only non-identity isometries we consider are *hyperbolic translations*, which are characterized by having two distinct fixed points on \mathbb{D}_∞ . An isometry of the form (1) is a hyperbolic translation if and only if $\text{Tr}^2(f) > 4$, where the trace-squared $\text{Tr}^2(f)$ of f is the square of the trace $\pm 2 \text{Re } a$ of the matrices representing f .

The f -orbit $\{f^n(z) \mid n \in \mathbb{Z}\}$ of a point $z \in \mathbb{D}$ is contained in a Euclidean circle through the fixed points of the hyperbolic translation f . See Figure 1. Let d denote the distance in the hyperbolic plane. The *translation length* $\ell(f)$ of a hyperbolic translation f is the minimal value of the displacement function $z \mapsto d(z, f(z))$, which is attained at all points z on the geodesic connecting the two fixed points of f . This geodesic is the *axis of f* . The translation length is given by $\cosh^2(\frac{1}{2}\ell(f)) = \frac{1}{4} \text{Tr}^2(f)$, or, in terms of the matrices (2) representing f :

$$\cosh(\frac{1}{2}\ell(f)) = |\text{Re } a|. \quad (3)$$

2.2 Hyperbolic surfaces, closed geodesics and systoles

In our setting a *hyperbolic surface* is a two-dimensional orientable manifold without boundary which is locally isometric to the hyperbolic plane. In particular, it has constant Gaussian curvature

-1. By the uniformization theorem [1] a hyperbolic surface \mathbb{M} has \mathbb{D} as its universal covering space. The surface \mathbb{M} is isometric to the quotient surface \mathbb{D}/Γ of the hyperbolic plane \mathbb{D} under the action of a Fuchsian group, i.e., a discrete group Γ of orientation preserving isometries of \mathbb{D} . The orbit Γz of a point $z \in \mathbb{D}$ is a discrete subset of \mathbb{D} . Note that $\Gamma z = \pi^{-1}(z)$, with $z = \pi(\mathbf{z}) \in \mathbb{M}$. We emphasize that points in the hyperbolic plane \mathbb{D} are denoted by $\mathbf{z}, \mathbf{w}, \mathbf{p}, \mathbf{q}$ and so on, whereas the corresponding points on the surface \mathbb{D}/Γ are denoted by z, w, p, q and so on. Since \mathbb{D}/Γ is a smooth closed surface all non-identity elements of Γ are hyperbolic translations.

The covering projection $\pi : \mathbb{D} \rightarrow \mathbb{D}/\Gamma$ is a local isometry. The distance between points p and q on a hyperbolic surface is given by $\min\{d(\mathbf{p}, \mathbf{q}) \mid \mathbf{p} \in \pi^{-1}(p), \mathbf{q} \in \pi^{-1}(q)\}$ and, abusing notation, is denoted by $d(p, q)$. The projection π 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 $f' = h^{-1}fh$).

A *simple closed geodesic* of \mathbb{M} is the π -image of a hyperbolic segment $[z, f(z)]$ on the axis of a hyperbolic translation f such that π is injective on the open segment $(z, f(z))$. The length of this geodesic is equal to the translation length $\ell(f)$ of f . For every $L > 0$ the number of simple closed geodesics of \mathbb{M} with length less than L is finite, so there is at least one with minimal length. This minimal length is the *systole* of \mathbb{M} , denoted by $\text{sys}(\mathbb{M})$. It is known that

$$\text{sys}(\mathbb{M}) \leq 2 \log(4g - 2) \quad (4)$$

for every hyperbolic surface \mathbb{M} of genus g [8, Theorem 5.2.1].

A *triangle* t on a hyperbolic surface 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 a hyperbolic surface 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.

2.3 Fundamental domain for the action of a Fuchsian group

The *Dirichlet region* $D_{\mathbf{p}}(\Gamma)$ of a point $\mathbf{p} \in \mathbb{D}$ with respect to the Fuchsian group Γ is the the closure of the open cell of p in the Voronoi diagram $\text{VD}_{\mathbf{p}}(\Gamma\mathbf{p})$ of the infinite discrete set of points $\Gamma\mathbf{p}$ in \mathbb{D} . In other words, $D_{\mathbf{p}}(\Gamma) = \{\mathbf{x} \in \mathbb{D} \mid d(\mathbf{x}, \mathbf{p}) \leq d(\mathbf{x}, f(\mathbf{p})) \text{ for all } f \in \Gamma\}$. Since \mathbb{D}/Γ is compact, every Dirichlet region is a compact convex hyperbolic polygon with finitely many sides. Each Dirichlet region $D_{\mathbf{p}}(\Gamma)$ is a *fundamental domain* for the action of Γ on \mathbb{D} , i.e., (i) $D_{\mathbf{p}}(\Gamma)$ contains at least one point of the orbit $\Gamma\mathbf{p}$, and (ii) if $D_{\mathbf{p}}(\Gamma)$ contains more than one point of $\Gamma\mathbf{p}$ then all these points of $\Gamma\mathbf{p}$ lie on its boundary.

2.4 Generalized Bolza surfaces

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 orientation preserving isometries f_0, \dots, f_{4g-1} , with $f_{2g+k} = f_k^{-1}$ for $k = 0, \dots, 2g-1$, pair opposite sides of D_g . An explicit expression for f_k as a matrix A_k is known, but we will only use this in the proof of Lemma 12 in Appendix B. By Poincaré's Theorem these side-pairings generate a Fuchsian group, the generalized Bolza group Γ_g , all non-identity elements of which are hyperbolic translations.

It can be seen that the element

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

of the group Γ_g , where indices are counted modulo $4g$, is the identity map $\mathbb{1}$. The equation $\prod_{j=0}^{4g-1} f_{j(2g+1)} = \mathbb{1}$ is usually called the *relation* of Γ_g . However, there are many ways to write the




Figure 2: 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 $\mathbb{1} \in \Gamma_2$.

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 $\prod_{j=0}^{4g-1} f_{k+j(2g+1)} = \mathbb{1}$. Similarly, taking inverses we see that $\prod_{j=0}^{4g-1} f_{k+j(2g-1)}$ is the identity map as well. The regular $4g$ -gon D_g is the Dirichlet region of the origin.

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

Let us now define the *original domain* \tilde{D}_g , which contains exactly one representative of each point on the surface \mathbb{M}_g , i.e., of each orbit under Γ_g . We denote as $s_j, j = 0, \dots, 2g-1$ the sides of D_g , where s_j lies between v_j and v_{j+1} ; with this notation, the side-pairing f_j maps s_{j+2g} to s_j (counting modulo $4g$). The original domain \tilde{D}_g is constructed from the fundamental D_g as follows (see Figure 3): \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 following $2g$ sides are not. For a point p on \mathbb{M}_g , the *canonical* representative of p




Figure 3: Original domain \tilde{D}_g for $g = 2$. Only vertex v_0 and the solid sides are included in \tilde{D}_g .

is the unique point of the orbit $\pi_g^{-1}(p)$ that lies in \tilde{D}_g . We then 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 can be used to label the region. Lastly, we denote the union of these *neighboring regions* of D_g as

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

3 Computing Delaunay triangulations

3.1 Bowyer's algorithm in the Euclidean plane

There exist various algorithms to compute Delaunay triangulations in Euclidean spaces. Bowyer's algorithm [7, 29] has proved its efficiency in CGAL [24].

Let us focus here on the two-dimensional case. Let \mathcal{P} be a set of points in the Euclidean plane \mathbb{C} and $\text{DT}_{\mathbb{C}}(\mathcal{P})$ their Delaunay triangulation. Let $\mathbf{p} \notin \mathcal{P}$ 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 $\text{DT}_{\mathbb{C}}(\mathcal{P})$ 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 $\text{DT}_{\mathbb{C}}(\mathcal{P} \cup \{\mathbf{p}\})$.

We ignore degeneracies here; they can be solved as in the case of flat orbit spaces [10].

An illustration is given in Figure 4. The first step of the insertion of \mathbf{p} uses geometric compu-




Figure 4: Insertion of a point in a Delaunay triangulation with Bowyer's incremental algorithm.

tations, 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 $\text{DT}_{\mathbb{C}}(\mathcal{P})$ in conflict with \mathbf{p} is a topological disk*. We will refer to this essential property in the next section.

3.2 Delaunay triangulations of points on closed surfaces

Let $\mathbb{M} = \mathbb{D}/\Gamma$ be a closed hyperbolic surface, as introduced above, 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 [5], the notion of conflict in \mathbb{D} is the same as in \mathbb{C} , since for the Poincaré disk model, hyperbolic circles are Euclidean circles (see Section 2.1).

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 5). 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 5: 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, their edges, and their vertices should form a *simplicial complex*. Here, a collection \mathcal{K} of vertices, edges, and triangles is called a simplicial complex if it satisfies the following two conditions:

- 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} .

In other words, the graph of edges of the triangles should have no loops (1-cycles) or multiple edges (2-cycles). Note that, as the set \mathcal{P} is finite, all triangulations considered in this paper are locally finite, so, we can skip the local finiteness in the above conditions (see also [10]).

For a set of points $\mathcal{Q} \subset \mathbb{M}$, we denote as $\delta(\mathcal{Q})$ 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 (Validity condition [6]). *Let $\mathcal{Q} \subset \mathbb{M}$ be a set of points such that*

$$\delta(\mathcal{Q}) < \frac{1}{2} \text{sys}(\mathbb{M}). \quad (5)$$

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 [10, Figure 3]; see also Remark 7 below.

The proof is easy, we include it for completeness.

Proof. Assume that condition (5) 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 endpoints 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 4. Let $\mathcal{S} \subset \mathbb{M}$ be a set of points satisfying the validity condition (5). 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 as $\text{DT}_{\mathbb{M}}(\mathcal{S})$.

As for the Bolza surface [6], Proposition 3 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 (5);

- 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.

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 will 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 5.

Proposition 5. *Let \mathbb{M} be a closed 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 (5) 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 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 will give a lower bound for the cardinality of a dummy point set.

Proposition 6. *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 (5) holds. Then*

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

Proof. Denote the number of vertices, edges and triangles in the (simplicial) Delaunay triangulation $\pi(\text{DT}_{\mathbb{D}}(\pi^{-1}(\mathcal{Q})))$ 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

$$k_2 = 4g - 4 + 2k_0. \quad (6)$$

Consider an arbitrary triangle t in $\pi(\text{DT}_{\mathbb{D}}(\pi^{-1}(\mathcal{Q})))$. Because the validity condition holds, the circumradius of t is smaller than $\frac{1}{2}\text{sys}$. It can be shown that it follows that $\text{area}(t) < \pi - 6 \arctan(\sqrt{3} \cosh(\frac{1}{4}\text{sys}))$; this is Lemma 23 in Appendix A. Because \mathbb{M} has area $4\pi(g-1)$, it follows that

$$k_2 > \frac{4\pi(g-1)}{\pi - 6 \arctan(\sqrt{3} \cosh(\frac{1}{2}\text{sys}))}. \quad (7)$$

Combining (6) and (7) yields the result. \square

To show that our lower and upper bound are meaningful, we will consider the asymptotics of these bounds for a family of surfaces with genus going to infinity. We consider several cases.

1. If the systoles of the family of surfaces are contained in a compact subset of $\mathbb{R}_{>0}$, then the upper bound is of order $O(g)$ and the lower bound of order $\Omega(g)$. Hence, a minimum dummy point set will have cardinality $\Theta(g)$.
2. If $\text{sys}(\mathbb{M}) \rightarrow 0$ when $g \rightarrow \infty$, 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 $O(g \text{sys}(\mathbb{M})^{-2})$. In a similar way, it can be shown that

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

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

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 Equation (4) in Section 2.2), we will only consider the case where $\text{sys}(\mathbb{M}) \sim C \log g$ for some $0 < C \leq 2$. 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 it 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 7. Note that the validity condition (5) 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. It has been shown that every hyperbolic surface of genus g has a simplicial Delaunay triangulation with at most $151g$ vertices [17]. 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 6 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 5 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 work it was 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 5 differs asymptotically as well.

4 Proof of Theorem 2: Systole of generalized Bolza surfaces

In the previous section we have seen that the validity condition (5) $\delta(\mathcal{Q}) < \text{sys}(\mathbb{M})$ on a point set $\mathcal{Q} \subset \mathbb{M}$ guarantees that the projection $\pi(\text{DT}_{\mathbb{D}}(\pi^{-1}(\mathcal{Q})))$ is a simplicial complex (Proposition 3). To be able to verify that this condition holds, we must know the value of the systole for the given hyperbolic surface.

This section is devoted to proving Theorem 2 stated in Introduction, which gives the value of the systole for the generalized Bolza surfaces \mathbb{M}_g defined in Section 2.

As preparation for the proof, we represent in Section 4.1 a given simple closed geodesic γ on \mathbb{M}_g as a sequence γ of pairwise disjoint hyperbolic line segments between the sides of D_g . The length of γ is equal to the sum of the lengths of the line segments in γ . We will also show how we can determine an element of Γ_g corresponding to the homotopy class of γ using the sequence γ .

Then, the proof consists of two parts. In Section 4.2 we will show that $\text{sys}(\mathbb{M}_g) \leq \varsigma_g$ by constructing a sequence σ of hyperbolic line segments such that $\text{length}(\sigma) = \varsigma_g$. Moreover, we show that the homotopy class of the corresponding closed geodesic σ is non-trivial.


In Section 4.3, 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 . This shows that $\text{sys}(\mathbb{M}_g) \geq \varsigma_g$. Theorem 2 follows from the arguments in Section 4.2 and 4.3.

4.1 Representation of a closed geodesic as a sequence of line segments


Consider a simple closed geodesic γ on the hyperbolic surface $\mathbb{M}_g = \mathbb{D}/\Gamma_g$, with projection map $\pi_g : \mathbb{D} \rightarrow \mathbb{D}/\Gamma_g$. 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 Figure 6a). These hyperbolic lines are the axes of conjugated elements of Γ_g . The intersection of $\pi_g^{-1}(\gamma)$ with D_g consists of m pairwise disjoint hyperbolic line segments between the sides of D_g (see Figure 6b), of which we denote the union by γ . We assume that these line segments are oriented and that their orientations are compatible with the orientation of γ . In particular, every line segment has a starting point and an endpoint. Consider a point $p \in \partial D_g$ and let p' be the point on ∂D_g identified with p under the action of Γ_g . We claim that if p is contained in one of the oriented line segments, then there exists a unique line segment in γ containing p' . Namely, because the union of line segments projects to the closed geodesic γ , there exists at least one line segment containing p' . Because the line segments are pairwise disjoint, it follows that there is exactly one line segment containing p' .

We order the oriented line segments up to cyclic permutation in the following way (see again Figure 6b). First, we pick γ_1 arbitrarily. To inductively define γ_{j+1} given γ_j , let p_j be the endpoint of γ_j . Let p'_j be the point on ∂D_g identified with p_j under the action of Γ_g . We define γ_{j+1} as the unique oriented line segment with starting point p'_j . For every j , the collection $\{\gamma_1, \gamma_2, \dots, \gamma_j\}$ of segments ordered up to step j has no duplicates. Therefore, we will return to γ_1 after exactly m steps. As γ_1 was chosen arbitrarily, the line segments are ordered up to cyclic permutation. In particular, we let $\gamma_m = \gamma_0$ and $\gamma_{m+1} = \gamma_1$.

Now we show how to determine an element of Γ_g corresponding to the homotopy class of γ using the sequence γ . As every hyperbolic line in $\pi_g^{-1}(\gamma)$ is the axis of a representative of the conjugacy class in Γ_g corresponding to γ , it is sufficient to find the element of Γ_g that has the line l_1 containing γ_1 as axis. To find the element of Γ_g with axis l_1 , we look for the first point on l_1 after p_1 that is an image of p_1 under some $f \in \Gamma_g$. To do this, we simply trace the fundamental regions that l_1 passes through (see Figure 7). Denote the edge containing the endpoint p_ℓ of γ_ℓ by s_{j_ℓ} . The line l_1 leaves D_g through p_1 , so through side s_{j_1} . Therefore, the first translate of D_g that l_1 passes through after D_g is $f_{j_1}(D_g)$. The segment of l_1 contained in $f_{j_1}(D_g)$ is the translate $f_{j_1}(\gamma_2)$ of γ_2 . Therefore, l_1 leaves $f_{j_1}(D_g)$ through the point $f_{j_1}(p_2)$ contained in side $f_{j_1}(s_{j_2})$. Then, the second translate of D_g that l_1 passes through is $f_{j_2}(f_{j_1}(D_g))$. Continuing this process, we end up at the translate $f_{j_m}f_{j_{m-1}} \cdots f_{j_1}(p_1)$ of p_1 after passing through m fundamental regions. We conclude that $f_{j_m}f_{j_{m-1}} \cdots f_{j_1}$ is a representative of the conjugacy class in Γ_g corresponding to γ .



(a) Preimage of a closed geodesic γ on the Bolza surface.



(b) Sequence of oriented hyperbolic line segments.

Figure 6: Sequence of hyperbolic line segments corresponding to a closed geodesic.

4.2 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 8. *There exists a sequence of segments σ in D_g corresponding to a homotopically non-trivial closed curve σ on \mathbb{M}_g such that $\text{length}(\sigma) = \varsigma_g$.*

Proof. Let σ be the sequence of segments depicted in Figure 6b, consisting of two segments γ_1 and γ_2 such that p_1, p_2, p'_1 and p'_2 are the midpoints of s_{2g+1}, s_0, s_1 and s_{2g} , respectively. The conjugacy class of the element of Γ_g corresponding to σ is $f_0 f_{2g+1}$, so σ is homotopically non-trivial. A straightforward computation in hyperbolic trigonometry shows that $\text{length}(\sigma) = \varsigma_g$. \square

4.3 Lower bound for the systole

Recall that to find the systole of \mathbb{M}_g we have to find the minimum length of a simple closed geodesic on \mathbb{M}_g . Because we have represented every simple closed geodesic as a sequence of segments between the sides of D_g , it suffices to find the minimum length of these sequences of segments instead. We consider different cases based on which “kind” of segments are contained in the sequence. For an integer $1 \leq k \leq 4g - 1$, we say that an oriented hyperbolic line segment between the sides of D_g is a k -segment if its starting point and endpoint are contained in s_j and s_{j+k} , respectively, for some $0 \leq j \leq 4g - 1$, where indices are counted modulo $4g$. Furthermore, we say that the segment is k -separated or has separation k for an integer $1 \leq k \leq 2g$ 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 6b are 1-separated, but γ_1 is a 1-segment while γ_2 is a 7-segment.

To show that $\text{sys}(\mathbb{M}_g) \geq \varsigma_g$, we distinguish between the following four cases:

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.




Figure 7: Finding the element of Γ_g corresponding to γ .

It is straightforward to check that every sequence of segments belongs to precisely one of these cases. Before examining them in detail we state two lemmas.

For Cases 1 and 2 we need the following lemma. This lemma will be used in the proof of Proposition 20 in Section 6 as well.

Lemma 9.

1. The length of a segment that has separation at least 2 is at least $\frac{1}{2}\varsigma_g$.
2. The length of a segment that has separation at least 4 is at least ς_g .

The proof is a straightforward use of hyperbolic trigonometry and can be found in Appendix B.

In the proofs for Case 3 and 4, we need the following lemma. Like Lemma 9, this lemma will be used again in the proof of Proposition 20 in Section 6.

Lemma 10. *The length of two consecutive 1-separated segments of which one is a 1-segment and the other is a $(4g - 1)$ -segment is at least $\frac{1}{2}\varsigma_g$. Furthermore, if a sequence of segments γ consists of only one 1-segment and one $(4g - 1)$ -segment, then $\text{length}(\gamma) \geq \varsigma_g$.*

The proof can be found in Appendix B.

Cases 1 and 2.

As a direct consequence of Lemma 9, we find the following corollary.

Corollary 11. *If a sequence of segments in D_g is in Case 1 or 2, then its length is at least ς_g .*

In the remainder of this section, we will show that the minimum length of a sequence of segments belonging to Case 3 or 4 is at least ς_g .

Case 3: γ contains exactly one segment that has separation 2 or 3 and all other segments are 1-separated.

For Case 3, we show that it is sufficient to consider only four possibilities for the homotopy class of γ . The length of γ can be computed from its homotopy class as the trace of the corresponding element of Γ_g (represented as a matrix). Therefore, we will compute the value of the trace of these specific elements in the following lemma.

Lemma 12. Let $m \in \{3, 2g - 2, 2g + 2, 4g - 3\}$ and define

$$Y_m = \prod_{j=1}^m f_{j(2g+1)}.$$

Then

$$\frac{1}{2} |\text{Tr}(Y_m)| \geq 1 + 2 \cos\left(\frac{\pi}{2g}\right).$$

The proof can be found in Appendix B.

Recall that we denoted the edge containing the endpoint \mathbf{p}_ℓ of a segment γ_ℓ by \mathbf{s}_{j_ℓ} . Suppose that γ_ℓ is a k_ℓ -segment. Then, the starting point of γ_ℓ is contained in $\mathbf{s}_{j_\ell - k_\ell}$. On the other hand, by definition the starting point of γ_ℓ is contained in the edge identified with $\mathbf{s}_{j_{\ell-1}}$ under the action of Γ_g , i.e., $\mathbf{s}_{j_{\ell-1} - 2g}$. It follows that $j_\ell - k_\ell = j_{\ell-1} - 2g$. Hence, $j_\ell - j_{\ell-1} = k_\ell - 2g$. We use this observation in the proof of Lemma 13.

Lemma 13. If a sequence of segments in D_g is in Case 3, then its length is at least ς_g .

Proof. Let γ be a sequence of segments belonging to Case 3. Changing the orientation of γ turns every 1-segment into a $(4g - 1)$ -segment and reversely, while keeping the length of the curve unchanged. Therefore, we can assume without loss of generality that the number m_{4g-1} of $(4g - 1)$ -segments in γ is at least the number m_1 of 1-segments. Furthermore, we order γ in such a way that γ_1 is the unique segment of separation 2 or 3. If γ_1 is a k -segment for $k = 2, 3, 4g - 3, 4g - 2$, then $j_1 - j_m = k - 2g$. Similarly, $j_\ell - j_{\ell-1} = 1 - 2g$ if γ_ℓ is a 1-segment, and $j_\ell - j_{\ell-1} = 2g - 1$ if γ_ℓ is a $(4g - 1)$ -segment for all $\ell \in \{2, 3, \dots, m\}$. Hence,

$$\sum_{\ell=1}^m (j_\ell - j_{\ell-1}) \equiv (m_{4g-1} - m_1)(2g - 1) + k - 2g \pmod{4g}.$$

Because the left-hand side is a telescoping sum, we see

$$\sum_{\ell=1}^m (j_\ell - j_{\ell-1}) = j_m - j_0 \equiv 0 \pmod{4g}.$$

Therefore,

$$\begin{aligned} (m_{4g-1} - m_1)(2g - 1) &\equiv 2g - k \pmod{4g}, \\ m_{4g-1} - m_1 &\equiv (2g - k)(2g - 1) \pmod{4g}. \end{aligned}$$

Recall that $m_{4g-1} \geq m_1$. If $m_1 \geq 1$, then there exists at least one pair $\gamma_\ell, \gamma_{\ell+1}$ for $\ell \in \{2, \dots, m - 1\}$ of consecutive segments consisting of one 1-segment and one $(4g - 1)$ -segment. By Lemma 10, $\text{length}(\gamma_\ell \cup \gamma_{\ell+1}) \geq \frac{1}{2}\varsigma_g$. Because $\text{length}(\gamma_1) \geq \frac{1}{2}\varsigma_g$ by Part 1 of Lemma 9, it follows that $\text{length}(\gamma) \geq \varsigma_g$.

If $m_1 = 0$, then

$$m_{4g-1} \equiv (2g - k)(2g - 1) \pmod{4g} \equiv \begin{cases} 2g + 2 & \pmod{4g} \text{ if } k = 2, \\ 3 & \pmod{4g} \text{ if } k = 3, \\ 4g - 3 & \pmod{4g} \text{ if } k = 4g - 3, \\ 2g - 2 & \pmod{4g} \text{ if } k = 4g - 2 \end{cases}.$$

Recall that Γ_g satisfies the relation $\prod_{j=1}^{4g} f_{j(2g+1)} = \mathbb{1}$, so every sequence of $4g$ consecutive $(4g - 1)$ -segments corresponds to a homotopically trivial curve. Hence, every sequence of segments in Case 3 with $m_1 = 0$ is homotopically equivalent to the sequence of segments with $m_{4g-1} \in \{3, 2g - 2, 2g + 2, 4g - 3\}$. By rotational symmetry of D_g , we can assume without loss of generality that $j_m = 2g + 1$. Then, the conjugacy class of Γ_g corresponding to γ can be represented by

$$f_{j_m} f_{j_{m-1}} \cdots f_{j_1} = \prod_{j=1}^{m_{4g-1}} f_{j(2g+1)},$$

since $j_{\ell-1} = j_\ell - 2g + 1 \equiv j_\ell + 2g + 1$ for all $\ell \in \{2, \dots, m\}$. By Lemma 12,

$$\frac{1}{2} \left| \text{Tr} \left(\prod_{j=1}^{m_{4g-1}} f_{j(2g+1)} \right) \right| \geq 1 + 2 \cos\left(\frac{\pi}{2g}\right).$$

It follows that $\text{length}(\gamma) \geq \varsigma_g$, which concludes the proof. \square

Case 4: all segments of γ are 1-separated.

We immediately state the result for Case 4.

Lemma 14. *If a sequence of segments in D_g is in Case 4 and represents a homotopically non-trivial curve on \mathbb{M}_g , then its length is at least ς_g .*

Remark 15. This is the only case in which we explicitly specify that γ represents a *homotopically non-trivial* curve on \mathbb{M}_g . We will give an example of a homotopically trivial curve for which the result of the lemma does not hold in Example 16 below.

Proof. Let γ be sequence of segments belonging to Case 4. As in the proof of Lemma 13, we assume without loss of generality that the number m_{4g-1} of $(4g-1)$ -segments in γ is at least the number m_1 of 1-segments. Then,

$$\sum_{\ell=1}^m (j_\ell - j_{\ell-1}) \equiv (m_{4g-1} - m_1)(2g-1) \pmod{4g}.$$

Because the left-hand side is a telescoping sum, we see

$$\sum_{\ell=1}^m (j_\ell - j_{\ell-1}) = j_m - j_0 \equiv 0 \pmod{4g}.$$

Therefore,

$$(m_{4g-1} - m_1)(2g-1) \equiv 0 \pmod{4g},$$

which implies that

$$m_{4g-1} \equiv m_1 \pmod{4g}.$$

If $m_1 \geq 2$, then there exist at least two disjoint pairs of consecutive segments, each consisting of one 1-segment and one $(4g-1)$ -segment. By Lemma 10, each of these pairs of segments has length at least $\frac{1}{2}\varsigma_g$, so in this case $\text{length}(\gamma) \geq \varsigma_g$.

If $m_1 = 1$, then $m_{4g-1} \equiv 1 \pmod{4g}$. Because the segments are ordered cyclically, all $(4g-1)$ -segments are consecutive (in this cyclic ordering). As in the proof of Lemma 13, the group relation implies that every sequence of $4g$ consecutive $(4g-1)$ -segments corresponds to a homotopically trivial curve. Hence, every sequence of segments in Case 4 with $m_1 = 1$ is homotopically equivalent to the sequence of segments with $m_{4g-1} = m_1 = 1$. By Lemma 10, it follows that $\text{length}(\gamma) \geq \varsigma_g$.

If $m_1 = 0$, then $m_{4g-1} \equiv 0 \pmod{4g}$. Because $\prod_{j=1}^{4g} f_{j(2g+1)} = \mathbb{1}$, a sequence of segments in Case 4 with $m_1 = 0$ is homotopically trivial. We conclude that $\text{length}(\gamma) \geq \varsigma_g$ in all cases where γ is homotopically non-trivial, which finishes the proof. \square

Example 16. We will show that there exists a sequence of segments corresponding to a homotopically trivial curve of length smaller than ς_g . Consider the sequence of segments γ in Figure 8, consisting of eight 1-segments (or $(4g-1)$ -segments, depending on the orientation of the curve). The distance between any starting point or endpoint of a segment and the closest vertex of D_g is fixed, say $\varepsilon > 0$.

This sequence of segments represents a circle on \mathbb{M}_g , which is a homotopically trivial curve. The length of γ depends continuously on ε , so for sufficiently small ε , the length of γ is smaller than ς_g .




Figure 8: Sequence of segments representing a homotopically trivial curve.

5 Computation of dummy points

Let \mathbb{M}_g be the generalized Bolza surface of genus $g \geq 2$ as defined earlier. In this section we will present two algorithms for constructing a dummy point set \mathcal{Q}_g satisfying the validity condition (5) for \mathbb{M}_g and give the growth rate of the cardinality of \mathcal{Q}_g as a function of g .

Both 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 [18], however we will not use them in this paper.

Each of the algorithms has its own advantages and drawbacks. The *refinement algorithm* (Section 5.1) yields a point set with optimal asymptotic cardinality $\Theta(g)$ (Proposition 6). The idea is borrowed from the well-known Delaunay refinement algorithm for mesh generation [28]. The *symmetric algorithm* (Section 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 $\Theta(g \log g)$.

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 17. *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 is given in Appendix C.

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 2.4) 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 $\mathcal{Q}_{\mathcal{N}_g} = \{f(\pi_g^{-1}(\mathcal{Q}_g) \cap \tilde{D}_g), f \in \mathcal{N}_g\}$.

Apart from the two algorithms, detailed below, we have also looked at the *structured algorithm* [21], which can be found in Appendix D. Its approach is fundamentally different from the simple 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 $\Theta(g \log g)$.

5.1 Refinement algorithm

Following the refinement strategy introduced above and using Proposition 17, 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 9 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 18. *The refinement algorithm terminates and the resulting dummy point set \mathcal{Q}_g satisfies the validity condition (5). The cardinality $|\mathcal{Q}_g|$ is bounded as follows*

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

Proof. We will first prove that the hyperbolic distance between two distinct points of \mathcal{Q}_g is greater than $\frac{1}{4} \text{sys}(\mathbb{M}_g)$. The distance between any pair of Weierstrass points is larger than $\frac{1}{4} \text{sys}(\mathbb{M}_g)$ (see Lemma 26 in Appendix C), of which a proof can be found in Appendix C.


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.


Finally, we will find bounds 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 6 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) for $g \rightarrow \infty$ and $g = 2$ respectively. This finishes the proof. \square




After initialization



First insertion



After first insertion



After last insertion

Figure 9: Several steps in the refinement algorithm (genus 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 10 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 $\frac{1}{2} \text{sys}(\mathbb{M}_g)$
 do
- 4 **for** $k = 0, \dots, 4g - 1$ **do**
- 5 Let \mathbf{p}_k be the circumcenter of Δ rotated around the origin by angle $\frac{k\pi}{2g}$.
- 6 Add $\pi_g(\mathbf{p}_k)$ to \mathcal{Q}_g .
- 7 **end**
- 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 [11]. However, asymptotically the resulting point set is larger than the point set obtained from the refinement algorithm.

Theorem 19. *The symmetric algorithm terminates and the resulting dummy point set satisfies the validity condition (5). Its cardinality is of order $\Theta(g \log g)$.*

Proof. The first two statements follow directly from the proof of Theorem 18, so we only have to prove the claim on the cardinality of \mathcal{Q}_g .

First, we will 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 18, 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 \mathbf{p}_k^j where $k = 0, \dots, 4g - 1$.

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

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

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 will give upper bounds for $|\mathcal{I}|$ and $|\mathcal{I}^c|$. To see for which j the inequality $a_j < 2ga$ holds, we will first look at the area of U_j (see Figure 11a). The amount of




After initialization



First insertion



After first insertion



After second (also last) insertion

Figure 10: Several steps in the symmetric algorithm (genus 3)

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 there exists $k \in \{0, \dots, 4g-1\}$ such that either $d(O, \mathbf{p}_k^j) < d_g$ or $d(\mathbf{v}_0, \mathbf{p}_k^j) < d_g$. If $d(O, \mathbf{p}_k^j) < d_g$ (Figure 11a), 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 11b). 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 regions $f(D_g)$ 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 regions that have \mathbf{v}_0 as one of their vertices. Because $2g+1$ and $4g$ are coprime, it follows that C_j contains exactly one translate of \mathbf{p}_ℓ^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 11c). 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}_ℓ^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 .


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 $d(\mathbf{v}_0, \mathbf{p}_k^j) \geq d_g$. 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 $d(\mathbf{v}_0, \mathbf{p}_k^j) \geq d_g$, 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


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

Therefore,

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



(a) $d(O, \mathbf{p}_k^j) < d_g$ for all $k = 0, \dots, 4g - 1$. (b) $d(\mathbf{v}_0, \mathbf{p}_k^j) < d_g$ for some $k \in \{0, \dots, 4g - 1\}$.



(c) $d(O, \mathbf{p}_k^j) \geq d_g$ and $d(\mathbf{v}_0, \mathbf{p}_k^j) \geq d_g$ for all $k \in \{0, \dots, 4g - 1\}$ and $d(\mathbf{p}_0^j, \partial D_g) < \frac{1}{8} \text{sys}(\mathbb{M}_g)$.

Figure 11: 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 \mathbf{v}_0 are drawn. In the third drawing, only the disks with center in D_g together with the unique disk that overlaps $D(\mathbf{p}_0^j, \frac{1}{8} \text{sys}(\mathbb{M}_g))$ are drawn.

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

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

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 will show that $|\mathcal{Q}_g|$ is of order $\Omega(g \log g)$. As before, denote the points added to \mathcal{Q}_g in iteration j of the while loop by \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 12).




Figure 12: The left figure shows a path P from the origin to \mathbf{v} that visits two vertices from the same iteration, namely \mathbf{p}_3^2 and \mathbf{p}_4^2 . The right figure shows a shorter path from the origin to \mathbf{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 $P_1 = \langle O, \mathbf{p}_4^2 \rangle$, $P_2 = \langle \mathbf{p}_4^2, \mathbf{p}_3^2 \rangle$ and $P_3 = \langle \mathbf{p}_3^2, \mathbf{p}_1^1, \mathbf{p}_2^2, \mathbf{v} \rangle$. 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 $-\frac{\pi}{2g}$.

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, so in particular $k_l \neq k_m$. Subdivide P into three paths: the path P_1 from O to $\mathbf{p}_{k_l}^{j_l}$, the path P_2 from $\mathbf{p}_{k_l}^{j_l}$ to $\mathbf{p}_{k_m}^{j_m}$, and the path P_3 from $\mathbf{p}_{k_m}^{j_m}$ to \mathbf{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 $\mathbf{p}_{k_m}^{j_m}$ of the same length of P_1 . It follows that $P' := P'_1 \cup P_3$ is a path from O to \mathbf{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

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

As $\frac{1}{2} \text{sys}(\mathbb{M}_g)$ is bounded as a function of g , 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

5.3 Experimental results for small genus

The refinement algorithm and the symmetric algorithm have been implemented. The implementation uses the `CORE::Expr` number type [31] 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 13 we have shown the dummy point set computed by the symmetric algorithm. However, a smaller set, consisting of 14 dummy points, was proposed earlier [6]: in addition to the six Weierstrass points, it contains the eight midpoints of the segments $[O, \mathbf{v}_k]$, $k = 0, 1, \dots, 7$ (see Figure 13).




Figure 13: Set of 22 dummy points for the Bolza surface computed by the symmetric algorithm (left) and set of 14 dummy points constructed by hand [6] (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 9), while the symmetric algorithm leads to 32 dummy points (Figure 10). 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.

6 Data structure and implementation

6.1 Canonical representatives

We have defined in Section 2.4 the canonical representative of a point on the surface \mathbb{M}_g . Let us now determine a unique canonical representative for each orbit of a triangle in \mathbb{D} under the action of Γ_g .

We consider all the neighboring regions, i.e., the images of D_g by a translation in \mathcal{N}_g (see Section 2.4, 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 \mathbf{v}_0 , which gives an ordering of \mathcal{N}_g . An illustration for genus 2 is shown in Figure 14.

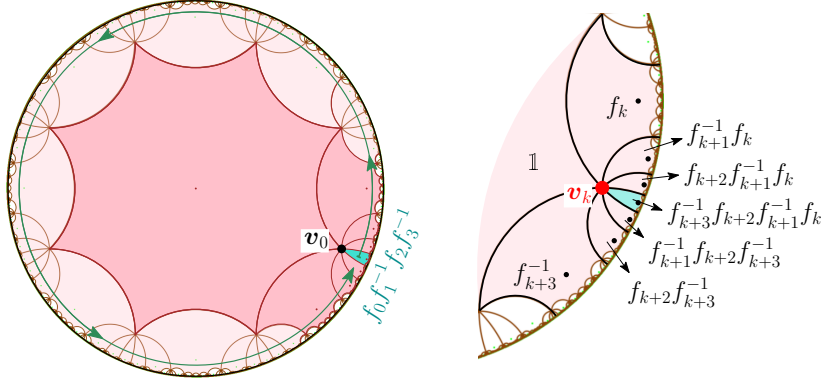


Figure 14: The neighboring translations for genus $g = 2$. The ordering starts with $f_0 f_1^{-1} f_2 f_3^{-1} = f_0 f_5 f_2 f_7$. Right: enumeration of the regions around a vertex $\mathbf{v}_k, k \in \{0, 1, 2, 3\}$. Note that $f_{k+3}^{-1} f_{k+2} f_{k+1}^{-1} f_k = f_k f_{k+1}^{-1} f_{k+2} f_{k+3}^{-1}$ by the group relation.

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 20 (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 $\mathbf{p} \in \mathbf{e}_0$. In Section 4.3, 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 .

Recall that f_0 is the side-pairing transformation that maps \mathbf{e}_{2g} to \mathbf{e}_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 15a). By Lemma 9, $d(\mathbf{p}, \mathbf{q}) \geq \frac{1}{2} \text{sys}(\mathbb{M}_g)$.

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 15b), so $d(\mathbf{p}, \mathbf{q}) \geq \frac{1}{2} \text{sys}(\mathbb{M}_g)$ by Lemma 9. 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 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 Lemma 9, so $d(\mathbf{p}, \mathbf{q}) \geq \frac{1}{2} \text{sys}(\mathbb{M}_g)$ as well. If




Figure 15: Cases in the proof of Proposition 20.

$[r, q]$ is a 1-separated segment, then it is clear that one of $[p, r]$ and $[r, q]$ is a 1-segment and the other a $(4g - 1)$ -segment. By Lemma 10, $d(p, r) + d(r, q) \geq \frac{1}{2} \text{sys}(\mathbb{M}_g)$.

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

Let now $\mathcal{S} \subset \mathbb{M}_g$ be a set of points satisfying the validity condition (5). 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

$$\Sigma = \left\{ t \in \pi_g^{-1}(t) \mid t \text{ has at least one vertex in } \tilde{D}_g \right\} \quad (8)$$

contains at most three faces. See Figure 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 20, each face $t \in \Sigma$ is contained in D_{N_g} . So, for each vertex u of t , there is a unique translation $T(u, t)$ in $N_g \cup \{1\}$ 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 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 $f_0 f_{2g-1} f_{2(2g-1)} \cdots f_{(2g-1)^2}$ for the counterclockwise order on N_g .

To summarize, we have shown that:

Proposition 21. *Let $\mathcal{S} \subset \mathbb{M}_g$ be a set of points satisfying the validity condition (5). For any face t in $\text{DT}_{\mathbb{M}_g}(\mathcal{S})$, there exists a unique canonical representative $t^c \in D_{N_g}$ in $\text{DT}_{\mathbb{D}}(\pi_g^{-1}(\mathcal{S}))$.*

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 .

6.2 Data structure

Proposition 21 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 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




Figure 16: Examples (for $g = 2$) of faces of $\text{DT}_{\mathbb{D}}(\pi_g^{-1}(\mathcal{S}))$ 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.

\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 $T(\mathbf{u}_j, \mathbf{t}^c), j = 0, 1, 2$ in \mathcal{N}_g as defined above, so that applying each translation to the corresponding canonical point yields the canonical representative \mathbf{t}^c of \mathbf{t} , i.e.,

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

In the rest of this section, we show how this data structure supports the algorithm that was briefly sketched in Section 3.2.

Finding conflicts. The notion of conflict 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 (8).

Since hyperbolic circles are Euclidean circles, the Euclidean Delaunay triangle containing the point \mathbf{p}^c gives us a triangle in the Delaunay triangulation in \mathbb{D} that 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* [14]: 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 \mathbf{t}, \mathbf{t}' be two adjacent faces in $\text{DT}_{\mathbb{M}_g}(\mathcal{S})$ and \mathbf{t}, \mathbf{t}' two of their preimages by π_g in $\text{DT}_{\mathbb{D}}(\pi^{-1}(\mathcal{S}))$. We define the *neighbor translation* $T^{\text{nbr}}(\mathbf{t}', \mathbf{t})$ from \mathbf{t}' to \mathbf{t} as the translation of Γ_g such that $T^{\text{nbr}}(\mathbf{t}', \mathbf{t})(\mathbf{t}')$ is adjacent to \mathbf{t} in $\text{DT}_{\mathbb{D}}(\pi^{-1}(\mathcal{S}))$.

Let u be a vertex common to \mathbf{t} and \mathbf{t}' , and let \mathbf{u}_t and $\mathbf{u}_{t'}$ the vertices of \mathbf{t} and \mathbf{t}' that project on u by π_g . We can compute the neighbor translation from \mathbf{t} to \mathbf{t}' as

$$T^{\text{nbr}}(\mathbf{t}', \mathbf{t}) = T(\mathbf{u}_t, \mathbf{t})(T(\mathbf{u}_{t'}, \mathbf{t}'))^{-1}.$$

Figure 17 illustrates the neighbor translation of the canonical representatives of \mathbf{t} and \mathbf{t}' . It can be easily seen that $T^{\text{nbr}}(\mathbf{t}', \mathbf{t}) = T(\mathbf{u}_t, \mathbf{t})(T(\mathbf{u}_{t'}, \mathbf{t}'))^{-1} = (T(\mathbf{u}_{t'}, \mathbf{t}')(T(\mathbf{u}_t, \mathbf{t})))^{-1} = (T^{\text{nbr}}(\mathbf{t}, \mathbf{t}'))^{-1}$.

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: $T_p^{\text{loc}} = T_p^{\text{loc}} T^{\text{nbr}}(\mathbf{t}', \mathbf{t})$. 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 $(T_p^{\text{loc}})^{-1}(\mathbf{t}_p)$.




Figure 17: Translating t'^c by $T^{\text{nbr}}(t'^c, t^c)$ gives a face adjacent to t^c .

See Figure 18 for an example. Here the walk first visits canonical faces and reaches the face $t_D \subset D_g$; up to that stage, T_p^{loc} is unchanged. Then the walk visits the non-canonical neighbor t' of t_D , and T_p^{loc} is updated to $T^{\text{nbr}}(t', t_D)$. The next face visited by the walk is t_p , which contains p^c ; as t_p^c and t'^c are adjacent, T_p^{loc} is left unchanged.

The location translation T_p^{loc} is then used when looking for all triangles in conflict with p . Starting from t_p^c , for each face in conflict with p^c , we recursively examine the translated image under T_p^{loc} of each neighbor (obtained with a neighbor translation) that has not yet been visited. We store the set C_p^c of canonical faces in conflict with p^c . Note that C_p^c is not necessarily a connected region in \mathbb{D} , as illustrated in Figure 18(Right). However, the projection of C_p^c onto \mathbb{M}_g is a topological disk, and by construction, the union of the translated faces $T_p^{\text{loc}} T^{\text{nbr}}(t, t_p^c)(t)$, for all $t \in C_p^c$ is a topological disk C_p in \mathbb{D} .




Figure 18: Left: The shaded faces are the (not necessarily canonical) faces in conflict with the red point. Their union C_p is a topological disk. Right: The union C_p^c of the (shaded) canonical faces in conflict with the red point is not a connected region in \mathbb{D} .

Inserting a point. To actually insert the new point p on \mathbb{M}_g , we can first create a new vertex storing p^c . For each face t in C_p^c , the translations $T_p^{\text{loc}} T^{\text{nbr}}(t, t_p^c) T(u_j, t)$, $j = 0, 1, 2$ can temporarily be stored directly in its three vertices (not in the face itself, since it will be deleted). Note that since C_p is a topological disk, the resulting translation for a given vertex is the same for all faces of C_p incident to it. We store $\mathbb{1}$ in the new vertex.

For each edge e on the boundary of C_p , we create a new face t_e formed by the new vertex and the edge e . Note that t_e is not necessarily canonical. The neighbor of t_e outside C_p^c is the neighbor through e of the face in C_p^c incident to e . Two new faces consecutive along the boundary of C_p are adjacent. We can now delete all faces in C_p .

All that is left to do now is to compute the translations to be stored in each new face t_{new} to get its canonical representative. We retrieve the translations temporarily stored in its vertices. If

all translations are equal to $\mathbb{1}$, then the face is already canonical and there is nothing more to do. Otherwise, one of the vertices of the newly created representative \mathbf{t}_{new} is the new vertex; without loss of generality it is its vertex \mathbf{u}_0 , and $T(\mathbf{u}_0, \mathbf{t}_{new}) = \mathbb{1}$. For $j = 0, 1, 2$ we can easily compute the index d_j of $T(\mathbf{u}_{t_j}^*, \mathbf{t}_j)$ in the ordering on \mathcal{N}_g , where \mathbf{t}_j is the image of \mathbf{t}_{new} under $(T(\mathbf{u}_j, \mathbf{t}_{new}))^{-1}$: $\mathbf{u}_{t_j}^*$ is the first vertex of \mathbf{t}_j such that $T(\mathbf{u}_{t_j}^*, \mathbf{t}_j) \neq \mathbb{1}$. This can be done by computations on the translations. Finally, the translations $(T(\mathbf{u}_k, \mathbf{t}_{new}))^{-1}T(\mathbf{u}_j, \mathbf{t}_{new}), j = 0, 1, 2$, are stored in the face \mathbf{t}_{new} , where k is the index for which d_k is minimal. The data structure now gives access to the canonical representative of \mathbf{t}_{new} .

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

6.3 Experiments

The algorithm presented in Section 3 was implemented in C++, with the data structure described in the previous section. The preprocessing step consists in computing dummy points that serve for the initialization of the data structure, following the two options presented in Section 5. The implementation also uses the value of the systole given by Theorem 2.

Following the celebrated exact geometric computation paradigm [30], 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 \mathbf{p} in \tilde{D}_2 lies on the right side, the left side, or on an oriented Euclidean segment.
- **INCIRCLE**, which checks whether an input point \mathbf{p} in \tilde{D}_2 lies inside, outside, or on the boundary of the disk circumscribing an oriented triangle.

Input points, which lie in \tilde{D}_2 , are tested against canonical triangles of the triangulation, whose vertices are images of input points by translations in \mathcal{N}_2 . 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 22 ([22, 21]). *In the case of the Bolza surface, the predicates can be evaluated by determining the sign of rational polynomial expressions of total degree at most 72 in the coordinates of input points.*

The proof consists of symbolic computations with MAPLE, organized along a case analysis on the different possible positions of the canonical triangles, i.e., on the possible translations of \mathcal{N}_2 that can be involved in each predicate. Its heaviness does not allow for an extension to higher genus. It is expected that the algebraic degree of predicates is much higher than 72 already for genus 3. In practice, the implementation relies on the `CORE::Expr` number type [31], which provides with exact and filtered computations. As for the computation of dummy points (Section 5.3), the evaluation exceeds the capabilities of CORE for genus bigger than 2, due the barriers raised by their very high algebraic degree, so, only a non-robust implementation of the algorithm can be obtained.

A fully robust implementation for the Bolza surface has been integrated in CGAL [23]. All details on the implementation can be found in Iordanov's PhD thesis [21]. We only mention a few key points here.

Let us first continue the discussion on predicates. To avoid increasing further the algebraic degree of predicates, the coordinates of dummy points are rounded to rationals (see Table 1). We have checked that the validity condition (5) still holds for the rounded points, and that the combinatorics of the Delaunay triangulations of exact and rounded points are identical.

Attention has also been paid to the manipulation of translations. As seen in Section 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

Table 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	$\left(\frac{2^{3/4}(\sqrt{2+\sqrt{2}}-\sqrt{2-\sqrt{2}})}{4\sqrt{2}+4\sqrt{2-\sqrt{2}}}, -\frac{2^{3/4}(\sqrt{2-\sqrt{2}}+\sqrt{2+\sqrt{2}})}{4\sqrt{2}+4\sqrt{2-\sqrt{2}}} \right)$	$(4/19, -1/2)$

of Γ_2 . The composition of two translations corresponds to the concatenation of their two corresponding words. Section 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 [12, 19], 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 6.1 and 6.2, the longest 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 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 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.

A Proof of Lemma 23

Lemma 23. *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 23 is the special case $m = 3$ of the following lemma. A proof was given in Ebbens's master's thesis [16], but for completeness we have included it here as well.

Lemma 24. *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 [26]. We will use the same approach to prove Lemma 24.

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 19).

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

$$f(\theta_1, \theta_2, \theta_3) = \sum_{j=1}^3 \operatorname{arccot}(\cosh R \tan \theta_j) \tag{9}$$

subject to the constraints $\sum_{j=1}^3 \theta_j = \pi$ and $0 \leq \theta_j < \pi$, i.e., minimizing (9) over the triangle in \mathbb{R}^3 with vertices $(\pi, 0, 0)$, $(0, \pi, 0)$, $(0, 0, \pi)$. We parametrize this triangle as follows

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




Figure 19: Division of P into three pairs of right-angled triangles

for $0 < s < \pi/2$ and $|t| \leq s$. By (9), 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

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

□

B Hyperbolic trigonometry of generalized Bolza surfaces

In this section of the appendix, we give the proofs of results stated in Section 4. First, we give the proof of Lemma 9.

Proof. (of Lemma 9)

Consider a segment γ_j of separation $k \geq 2$. 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 γ_j^\perp between s_0 and s_k (see Figure 20).




Figure 20: Computing the length of γ_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 20. Therefore, [4, 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 γ_j^\perp , we obtain two quadrilaterals, each of which has three right angles and remaining angle $\frac{k\pi}{4g}$. It follows that [4, Theorem 7.17.1(ii)]

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

The lower bounds for the length of γ_j follow from $\sin(\frac{k\pi}{4g}) \geq \sin(\frac{\pi}{2g})$ for $k \geq 2$ and $\sin(\frac{k\pi}{4g}) \geq \sin(\frac{\pi}{g})$ for $k \geq 4$, and a direct computation using properties of trigonometric functions. □

Second, we give the proof of Lemma 10.

Proof. (of Lemma 10)

Denote the two segments by γ_1 and γ_2 . Without loss of generality, we will assume 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 21). Let x be the distance between p_1 and v_1 and let α_1 be the angle between γ_1 and s_1 . The distance between p'_1 and 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 v_{2g+1} .




Figure 21: Construction in the proof of Lemma 10. We have denoted the starting point of γ_1 by v'_m , since γ may consist of an arbitrary number, say m , of segments.

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 will minimize

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

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

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

for all $x > 0$, the function $x \mapsto \text{arsinh}(\sinh(x) \sin(\frac{\pi}{2g}))$ 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 $f(\frac{1}{2}\ell) = 2 \operatorname{arsinh}(\sinh(\frac{1}{2}\ell) \sin(\frac{\pi}{2g}))$. 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 $\operatorname{length}(\gamma_1 \cup \gamma_2) > \frac{1}{2}\varsigma_g$.

Now, assume that $\gamma = \gamma_1 \cup \gamma_2$. By the argument above, 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 \mathbf{p}_2 and \mathbf{p}'_2 . By symmetry of the argument, it follows that $\operatorname{length}(\gamma)$ is minimal when \mathbf{p}_2 is the midpoint of \mathbf{s}_{2g} and \mathbf{p}'_2 is the midpoint of \mathbf{s}_1 . By hyperbolic trigonometry, it can be seen that the length of $\gamma_1 \cup \gamma_2$ where \mathbf{p}_1 and \mathbf{p}_2 are the midpoints of their sides is ς_g . This finishes the proof. \square

Third, we give the proof of Lemma 12.

Proof. (of Lemma 12)

The side pairing transformation f_k is represented by the matrix $\pm A_k$ with

$$A_k = \begin{pmatrix} \cot(\frac{\pi}{4g}) & \exp(\frac{ik\pi}{2g})\sqrt{\cot^2(\frac{\pi}{4g}) - 1} \\ \exp(-\frac{ik\pi}{2g})\sqrt{\cot^2(\frac{\pi}{4g}) - 1} & \cot(\frac{\pi}{4g}) \end{pmatrix}$$

for $k = 0, 1, \dots, 2g - 1$ with $f_k^{-1} = f_{k+2g}$ [3]. Hence,

$$\frac{1}{2} |\operatorname{Tr}(Y_m)| = \frac{1}{2} \operatorname{Tr} \left(\prod_{j=1}^m A_{j(2g+1)} \right),$$

where the indices $j(2g - 1)$ are counted modulo $4g$.

Because Y_m for $m = 3$ is a product of just three matrices, it is straightforward to show that $\frac{1}{2} |Y_3| > 1 + 2 \cos(\frac{\pi}{2g})$. Furthermore,

$$\begin{aligned} \prod_{j=1}^{4g-3} A_{j(2g+1)} &= \prod_{j=1}^{4g} A_{j(2g+1)} \cdot A_{4g(2g+1)}^{-1} A_{(4g-1)(2g+1)}^{-1} A_{(4g-2)(2g+1)}^{-1}, \\ &= A_{4g(2g+1)}^{-1} A_{(4g-1)(2g+1)}^{-1} A_{(4g-2)(2g+1)}^{-1} \end{aligned}$$

by the group relation. Hence, we can find $\frac{1}{2} |\operatorname{Tr}(Y_m)|$ for $m = 4g - 3$ in a similar way by computing the trace of a product of three matrices.

In the remainder of the proof we treat the cases $m = 2g \pm 2$. For convenience, let $\alpha = \frac{\pi}{4g}$. The eigenvalues of A_k are roots of the characteristic polynomial $\chi(t) = t^2 - t \cdot \operatorname{Tr}(A_k) + 1$. Because $\operatorname{Tr}(A_k) = 2 \cot \alpha > 2$, the eigenvalues are distinct, real and positive, and do not depend on k . Therefore, we can diagonalize $A_k = S_k \Lambda S_k^{-1}$, where Λ is the diagonal matrix containing the eigenvalues of A_k and

$$S_k = \frac{1}{\sqrt{2}} \begin{pmatrix} e^{ik\alpha} & e^{ik\alpha} \\ -e^{-ik\alpha} & e^{-ik\alpha} \end{pmatrix}.$$

Then

$$Y_m = S_{2g+1} \Lambda S_{2g+1}^{-1} \cdot S_{2(2g+1)} \Lambda S_{2(2g+1)}^{-1} \cdots S_{(m-1)(2g+1)} \Lambda S_{(m-1)(2g+1)}^{-1} \cdot S_{m(2g+1)} \Lambda S_{m(2g+1)}^{-1}.$$

A straightforward computation shows that for all $j \in \{0, \dots, m\}$

$$S_{j(2g+1)}^{-1} S_{(j+1)(2g+1)} = T_{2g+1},$$

where

$$T_k := \begin{pmatrix} \cos k\alpha & i \sin k\alpha \\ i \sin k\alpha & \cos k\alpha \end{pmatrix}.$$

Therefore,

$$\begin{aligned} Y_m &= S_{2g+1}(\Lambda T_{2g+1})^m T_{2g+1}^{-1} S_{m(2g+1)}^{-1}, \\ &= S_{2g+1}(\Lambda T_{2g+1})^m T_{2g+1}^{-1} S_{m(2g+1)}^{-1} S_{2g+1} S_{2g+1}^{-1}, \\ &= S_{2g+1}(\Lambda T_{2g+1})^m T_{-m(2g+1)} S_{2g+1}^{-1}, \end{aligned}$$

where $T_{2g+1}^{-1} S_{m(2g+1)}^{-1} S_{2g+1} = T_{-m(2g+1)}$ in the last step can be verified with a direct computation. Writing $X = \Lambda T_{2g+1}$ and $Z_m = T_{-m(2g+1)}$, it follows that

$$\text{Tr}(Y_m) = \text{Tr}(X^m Z_m),$$

because Y_m and $X^m Z_m$ are conjugate matrices. It is known that [20, Equation (6)]

$$\text{Tr}(X^m Z_m) = \text{Tr}(X Z_m) \cdot \frac{\sin m\theta}{\sin \theta} - \text{Tr}(Z_m) \cdot \frac{\sin(m-1)\theta}{\sin \theta}, \quad (10)$$

where θ is the angle such that $\text{Tr}(X) = 2 \cos \theta$. To simplify Equation (10), we compute $\text{Tr}(X Z_m)$, $\text{Tr}(Z_m)$ and θ . Note that $T_k T_\ell = T_{k+\ell}$ for all k and ℓ , so

$$X Z_m = \Lambda T_{2g+1} T_{-m(2g+1)} = \Lambda T_{-(m-1)(2g+1)}.$$

Since Λ is a diagonal matrix, $\text{Tr}(\Lambda T_k) = \text{Tr} \Lambda \cdot \cos k\alpha$ for all k . Furthermore, $\text{Tr} \Lambda = \text{Tr} A_k = 2 \cot \alpha$. Therefore,

$$\text{Tr}(X Z_m) = \text{Tr}(\Lambda T_{-(m-1)(2g+1)}) = 2 \cot \alpha \cos((m-1)(2g+1)\alpha).$$

Using $2g\alpha = \frac{1}{2}\pi$ and the fact that $m = 2g \pm 2$ is even, we get

$$\cos((m-1)(2g+1)\alpha) = \cos(\frac{1}{2}(m-1)\pi + (m-1)\alpha) = (-1)^{m/2} \sin((m-1)\alpha).$$

Similarly,

$$\text{Tr}(Z_m) = \text{Tr}(T_{-m(2g+1)}) = 2 \cos(m(2g+1)\alpha).$$

Using again $2g\alpha = \frac{1}{2}\pi$ and the fact that $m = 2g \pm 2$ is even, we get

$$\cos(m(2g+1)\alpha) = \cos(\frac{1}{2}m\pi + m\alpha) = (-1)^{m/2} \cos(m\alpha).$$

Finally,

$$\text{Tr}(X) = \text{Tr}(\Lambda T_{2g+1}) = 2 \cot \alpha \cos((2g+1)\alpha) = 2 \cot \alpha \cdot -\sin \alpha = 2 \cos(\pi - \alpha),$$

so $\theta = \pi - \alpha$. Substituting the expressions for $\text{Tr}(X Z_m)$, $\text{Tr}(Z_m)$ and θ into Equation 10, we obtain

$$\begin{aligned} \frac{1}{2} \text{Tr}(Y_m) &= (-1)^{m/2} \left(\cot \alpha \sin((m-1)\alpha) \frac{\sin(m(\pi - \alpha))}{\sin(\pi - \alpha)} - \cos(m\alpha) \frac{\sin((m-1)(\pi - \alpha))}{\sin(\pi - \alpha)} \right) \\ &= (-1)^{m/2} \left(\cot \alpha \sin((m-1)\alpha) \frac{-\sin(m\alpha)}{\sin \alpha} - \cos(m\alpha) \frac{\sin((m-1)\alpha)}{\sin \alpha} \right) \\ &= (-1)^{m/2} \frac{-\sin(m-1)\alpha}{\sin^2 \alpha} (\cos \alpha \sin(m\alpha) + \cos(m\alpha) \sin \alpha) \\ &= (-1)^{(m+2)/2} \frac{\sin(m-1)\alpha \sin(m+1)\alpha}{\sin^2 \alpha}. \end{aligned}$$

Now, a straightforward trigonometric computation shows that $\frac{1}{2} |\text{Tr}(Y_m)| \geq 1 + 2 \cos(\frac{\pi}{2g})$ for $m = 2g \pm 2$, which finishes the proof. \square

C Proofs omitted in Section 5

In this section of the appendix, we give the proofs of results stated in Section 5.

First, we give the proof of Proposition 17, which states that for any finite set of points \mathcal{Q}_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 a geometric lemma.

Lemma 25. *Let C_q be a Euclidean disk centered at O and passing through a point q . 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 22: Illustrations for the proof of Lemma 25.

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 22 - 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 22 - 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. \square

Note that Lemma 25 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 can now proceed with the proof of Proposition 17.




Figure 23: Illustration for the proof of Proposition 17 for $g = 2$.

Proof. (Proposition 17)

We will 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 Γ_g in $D_{\mathcal{N}_g} \setminus \tilde{D}_g$, 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 23 - Left. Let $f(\tilde{D}_g), f \in \mathcal{N}_g$ 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 $[p_j, p_k]$ does not contain the origin, since in that case any disk through p_j and p_k clearly contains the origin. Then, there exists a line l_O through O such that p_j and p_k are contained in the same half-space. Let m_ℓ be the midpoint of a side between s_j and s_k in the same half-space as p_j and p_k . Consider the disk centered at O that passes through m_ℓ (and, of course, through all the other midpoints m_k as well), and consider also the line through O and m_ℓ . By Lemma 25, the disk $C(p_j, p_k)$ passing through O, p_j , and p_k contains m_ℓ . Since O and m_ℓ are on both sides of the segment $[p_j, p_k]$, any disk through p_j and p_k contains either m_ℓ or O , therefore there is no empty disk that passes through p_j and p_k . Assume now that there is an empty disk that passes through the endpoints of e' . This empty disk can then be shrunk continuously so that it passes through p_j and 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 e' , which implies that e' (and, by consequence, e) cannot be an edge in $\text{DT}_{\mathbb{D}}(\pi_g^{-1}(\mathcal{Q}_g))$.

Case B: The edge e crosses several images of \tilde{D}_g before leaving $D_{\mathcal{N}_g}$. This case is illustrated in Figure 23 - Right. There exist multiple images of 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 e' be an image of e that intersects two adjacent sides s_j and s_{j+1} of \tilde{D}_g so that the hyperbolic line supporting e' separates O and the midpoint m_{j+1} . Note that such an image of e exists always: e either separates O and the midpoint 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, $f^{-1}(e)$ separates O and the midpoint $f^{-1}(\mathbf{m}_{j+1})$. The edge e' intersects also the side s_k adjacent to s_{j+1} in the Dirichlet region that shares the side s_{j+1} with \tilde{D}_g (see Figure 23 - Right). Let \mathbf{p}_j and \mathbf{p}_k be the intersection points of e' with s_j and 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 25, 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 e' either, which implies that e' (and, by consequence, 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

Second, we compute the distance between any pair of Weierstrass points, as mentioned in the proof of Theorem 18.

Lemma 26. *The distance between any pair of distinct Weierstrass points of \mathbb{M}_g is at least $\frac{1}{4} \text{sys}(\mathbb{M}_g)$.*

Proof. Recall that the Weierstrass points of \mathbb{M}_g are represented by the origin, the vertices and the midpoints of the sides of D_g . The distance between midpoints of non-consecutive sides is clearly larger than the distance between midpoints of consecutive sides. Hence, by symmetry, it is sufficient to consider the region in Figure 24.




Figure 24: Computing the distances between Weierstrass points of \mathbb{M}_g .

Let \mathbf{m}_k be the midpoint of side s_k and let \mathbf{p}_k^* be the midpoint of \mathbf{m}_{k-1} and \mathbf{m}_k . First, in the right-angled triangle $[O, \mathbf{m}_0, \mathbf{v}_0]$ we know that [4, Theorem 7.11.3]

$$\cosh(d(O, \mathbf{m}_0)) = \cosh(d(\mathbf{m}_0, \mathbf{v}_0)) = \frac{\cos(\frac{\pi}{4g})}{\sin(\frac{\pi}{4g})} = \cot(\frac{\pi}{4g}),$$

$$\cosh(d(O, \mathbf{v}_0)) = \cot^2(\frac{\pi}{4g}).$$

Second, because $d(O, \mathbf{m}_0) = d(\mathbf{m}_0, \mathbf{v}_1)$, $[O, \mathbf{m}_0, \mathbf{p}_1^*]$ and $[\mathbf{v}_1, \mathbf{m}_0, \mathbf{p}_1^*]$ are congruent triangles. In particular, $\angle(O\mathbf{m}_0\mathbf{p}_1^*) = \angle(\mathbf{v}_1\mathbf{m}_0\mathbf{p}_1^*) = \frac{\pi}{4}$. It follows that

$$\cosh(\frac{1}{2}d(\mathbf{m}_0, \mathbf{m}_1)) = \cosh(d(\mathbf{m}_0, \mathbf{p}_1^*)) = \frac{\cos(\frac{\pi}{4g})}{\sin(\frac{\pi}{2})} = \sqrt{2} \cos(\frac{\pi}{4g}). \quad (11)$$

The above formulas yield expressions for $d(O, \mathbf{m}_0)$, $d(O, \mathbf{v}_0)$ and $d(\mathbf{m}_0, \mathbf{m}_1)$ and comparing these to the expression for $\text{sys}(\mathbb{M}_g)$ (see Theorem 2) yields the result. \square

D 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 will first explain the idea and the notation. See Figure 25 for an illustration of the dummy points within one slice of the $4g$ -gon.




Figure 25: 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 \mathbf{v}_j respectively. The midpoint of side \mathbf{s}_k is denoted by \mathbf{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 \mathbf{s}_k is obtained from \mathbf{s}_0 by rotating it around the origin by angle $\frac{k\pi}{2g}$, so \mathbf{p}_k^0 is obtained from \mathbf{p}_0^0 by rotating it in the same way. Hence, we will use the lower index as ‘rotation’ index in the definition of the other points as well.
2. Secondly, the projections $\pi_g(\mathbf{p}_k^*)$ of the midpoints \mathbf{p}_k^* of the geodesic segment in \mathbb{H}^2 connecting consecutive pairs $(\mathbf{p}_k^0, \mathbf{p}_{k+1}^0)$ 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 $[O, \mathbf{p}_k^0]$. This is the reason why they have a star as superscript.
3. Thirdly, points \mathbf{p}_0^j are consecutively added on $[O, \mathbf{p}_0^{j-1}]$ in such a way that the distance between consecutive points $\mathbf{p}_0^{j-1}, \mathbf{p}_0^j$ is given by $d(\mathbf{p}_0^{j-1}, \mathbf{p}_0^j) = \frac{1}{4} \text{sys}(\mathbb{M}_g)$, until $d(\mathbf{p}_0^j, O) \leq \frac{1}{4} \text{sys}(\mathbb{M}_g)$. By rotating the points in the same way as before we obtain the points \mathbf{p}_k^j . Here, the upper index denotes the ‘iteration’ index. Notice that the midpoints of the sides were denoted \mathbf{p}_k^0 to initialize this process. Since $d(O, \mathbf{p}_0^0) = \text{arcosh}(\cot(\frac{\pi}{4g}))$, this step consists of

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

iterations.

4. Finally, observe that the triangles $[O, \mathbf{p}_k^0, \mathbf{p}_{k+1}^0]$ and $[\mathbf{v}, \mathbf{p}_k^0, \mathbf{p}_{k+1}^0]$ are congruent under reflection in $[\mathbf{p}_k^0, \mathbf{p}_{k+1}^0]$. 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 will describe these combinatorics. The proof that this is the Delaunay triangulation of the dummy point set will be given in Lemma 29. Again, see Figure 25 for an illustration of the dummy points and the triangulation \mathcal{T} within one slice of the $4g$ -gon.

Definition 27. 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.

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

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 will 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 30). 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.

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. As the proofs are rather technical, the reader may want to skip them at a first reading.

Lemma 28. *The circumdiameters of triangles in \mathcal{T} are smaller than $\frac{1}{2} \text{sys}(\mathbb{M}_g)$.*

Proof. By symmetry it is sufficient to consider the circumscribed disks of the triangles

$$[\mathbf{p}_0^0, \mathbf{p}_0^1, \mathbf{p}_0^*], [\mathbf{p}_0^1, \mathbf{p}_1^1, \mathbf{p}_0^*], [\mathbf{p}_0^m, \mathbf{p}_1^m, O], [\mathbf{p}_0^j, \mathbf{p}_1^j, \mathbf{p}_1^{j+1}],$$

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 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 \mathbf{p}_k^0 .
- 4 For all $k = 0, \dots, 4g - 1$, label the midpoint of \mathbf{p}_k^0 and \mathbf{p}_{k+1}^0 by \mathbf{p}_k^* and add $\pi_g(\mathbf{p}_k^*)$ to \mathcal{Q}_g .
- 5 $m \leftarrow \lceil 4 \text{arcosh}(\cot(\frac{\pi}{4g})) / \text{sys}(\mathbb{M}_g) \rceil - 1$.
- 6 **for** $i = 1, \dots, m$ **do**
 - 7 Let \mathbf{p}_0^j be the point on $[\mathbf{p}_0^{j-1}, O]$ with $d(\mathbf{p}_0^j, \mathbf{p}_0^{j-1}) = \frac{1}{4} \text{sys}(\mathbb{M}_g)$ and add $\pi_g(\mathbf{p}_0^j)$ to \mathcal{Q}_g .
 - 8 **for** $k = 1, \dots, 2g - 1$ **do**
 - 9 Let \mathbf{p}_k^j be \mathbf{p}_0^j rotated clockwise around the origin by angle $\frac{k\pi}{2g}$.
 - 10 Let \mathbf{q}_k^j be the reflection of \mathbf{p}_k^j in the line through $\mathbf{p}_k^0, \mathbf{p}_{k+1}^0$.
 - 11 Add $\pi_g(\mathbf{p}_k^j), \pi_g(\mathbf{q}_k^j)$ to \mathcal{Q}_g .
 - 12 **end**
 - 13 **for** $k = 2g, \dots, 4g - 1$ **do**
 - 14 Let \mathbf{p}_k^j be \mathbf{p}_0^j rotated clockwise around the origin by angle $\frac{k\pi}{2g}$.
 - 15 Let \mathbf{q}_k^j be the reflection of \mathbf{p}_k^j in the line through $\mathbf{p}_{k-1}^0, \mathbf{p}_k^0$.
 - 16 Add $\pi_g(\mathbf{p}_k^j), \pi_g(\mathbf{q}_k^j)$ to \mathcal{Q}_g .
 - 17 **end**
- 18 **end**

Algorithm 3: Structured algorithm

for $1 \leq j \leq m - 1$. For easy reference, the used lengths and angles satisfy the following relations. Here we will denote $\text{length}([\mathbf{p}, \mathbf{q}])$ by abuse of notation by $[\mathbf{p}, \mathbf{q}]$.

$$[\mathbf{p}_0^0, \mathbf{p}_0^*] = [\mathbf{p}_0^0, \mathbf{p}_0^1] = \frac{1}{4} \text{sys}(\mathbb{M}_g), \quad (12)$$

$$\cosh(\frac{1}{2} \text{sys}(\mathbb{M}_g)) = 1 + 2 \cos(\frac{\pi}{2g}), \quad (13)$$

$$\angle(O\mathbf{p}_0^0\mathbf{p}_0^*) = \frac{\pi}{4}, \quad (14)$$

$$\sinh(\frac{1}{2}[\mathbf{p}_0^1, \mathbf{p}_0^*]) = \sinh(\frac{1}{4} \text{sys}(\mathbb{M}_g)) \sin(\frac{1}{2}\angle(O\mathbf{p}_0^0\mathbf{p}_0^*)), \quad (15)$$

$$\sin \angle(\mathbf{p}_0^1\mathbf{p}_0^*\mathbf{p}_0^0) = \frac{\sinh(\frac{1}{4} \text{sys}(\mathbb{M}_g)) \sin \angle(O\mathbf{p}_0^0\mathbf{p}_0^*)}{\sinh([\mathbf{p}_0^1, \mathbf{p}_0^*])}, \quad (16)$$

$$\angle(\mathbf{p}_0^1\mathbf{p}_0^*\mathbf{p}_0^1) = \pi - 2\angle(\mathbf{p}_0^1\mathbf{p}_0^*\mathbf{p}_0^0). \quad (17)$$

In general, the circumradius R of a hyperbolic triangle with sides a, b, c and angles α, β, γ satisfies

$$\tanh(R) = \frac{4 \sinh(\frac{1}{2}a) \sinh(\frac{1}{2}b) \sinh(\frac{1}{2}c)}{\sinh(b) \sinh(c) \sin \alpha} = \frac{\sinh(\frac{1}{2}a)}{\cosh(\frac{1}{2}b) \cosh(\frac{1}{2}c) \sin \alpha}. \quad (18)$$

In the specific case of an isosceles hyperbolic triangle with legs of length c and vertex angle α we know

$$\tanh(R) = \frac{\tanh(\frac{1}{2}c)}{\cos(\frac{1}{2}\alpha)}.$$

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 $\tanh(R(\mathbf{p}_0^0, \mathbf{p}_0^1, \mathbf{p}_0^*)) < \tanh(\frac{1}{4} \text{sys}(\mathbb{M}_g))$. 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,

$$\tanh(\frac{1}{4} \text{sys}(\mathbb{M}_g)) \geq \tanh(\frac{1}{4} \text{sys}(\mathbb{M}_g))|_{g=2} \approx 0.643.$$

It follows that $\tanh(R(\mathbf{p}_0^0, \mathbf{p}_0^1, \mathbf{p}_0^*)) < \tanh(\frac{1}{4} \text{sys}(\mathbb{M}_g))$ holds. This proves that the circumdiameter of $[\mathbf{p}_0^0, \mathbf{p}_0^1, \mathbf{p}_0^*]$ is smaller than $\frac{1}{2} \text{sys}(\mathbb{M}_g)$.

Secondly, 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

$$\begin{aligned} \tanh(R(\mathbf{p}_0^1, \mathbf{p}_1^1, \mathbf{p}_0^*)) &= \frac{\tanh(\frac{1}{2}[\mathbf{p}_0^1, \mathbf{p}_0^*])}{\cos(\frac{1}{2}\angle(\mathbf{p}_0^1 \mathbf{p}_0^* \mathbf{p}_1^1))}, \\ &\stackrel{(17)}{=} \frac{\tanh(\frac{1}{2}[\mathbf{p}_0^1, \mathbf{p}_0^*])}{\sin \angle(\mathbf{p}_0^1 \mathbf{p}_0^* \mathbf{p}_0^*)}. \end{aligned}$$

By using

$$\begin{aligned} \sin \angle(\mathbf{p}_0^1 \mathbf{p}_0^* \mathbf{p}_0^0) &\stackrel{(16)}{=} \frac{\sinh(\frac{1}{4} \text{sys}(\mathbb{M}_g)) \sin \angle(O \mathbf{p}_0^0 \mathbf{p}_0^*)}{\sinh([\mathbf{p}_0^1, \mathbf{p}_0^*])}, \\ &= \frac{\sinh(\frac{1}{4} \text{sys}(\mathbb{M}_g)) \cdot 2 \sin(\frac{1}{2} \angle(O \mathbf{p}_0^0 \mathbf{p}_0^*)) \cos(\frac{1}{2} \angle(O \mathbf{p}_0^0 \mathbf{p}_0^*))}{\sinh([\mathbf{p}_0^1, \mathbf{p}_0^*])}, \\ &\stackrel{(15)}{=} \frac{\sinh(\frac{1}{2}[\mathbf{p}_0^1, \mathbf{p}_0^*]) \cdot 2 \cos(\frac{1}{2} \angle(O \mathbf{p}_0^0 \mathbf{p}_0^*))}{2 \sinh(\frac{1}{2}[\mathbf{p}_0^1, \mathbf{p}_0^*]) \cosh(\frac{1}{2}[\mathbf{p}_0^1, \mathbf{p}_0^*])}, \\ &= \frac{\cos(\frac{1}{2} \angle(O \mathbf{p}_0^0 \mathbf{p}_0^*))}{\cosh(\frac{1}{2}[\mathbf{p}_0^1, \mathbf{p}_0^*])}, \end{aligned}$$

we can rewrite this as

$$\begin{aligned} \tanh(R(\mathbf{p}_0^1, \mathbf{p}_1^1, \mathbf{p}_0^*)) &= \frac{\tanh(\frac{1}{2}[\mathbf{p}_0^1, \mathbf{p}_0^*]) \cosh(\frac{1}{2}[\mathbf{p}_0^1, \mathbf{p}_0^*])}{\cos(\frac{1}{2} \angle(O \mathbf{p}_0^0 \mathbf{p}_0^*))}, \\ &= \frac{\sinh(\frac{1}{2}[\mathbf{p}_0^1, \mathbf{p}_0^*])}{\cos(\frac{1}{2} \angle(O \mathbf{p}_0^0 \mathbf{p}_0^*))}, \\ &\stackrel{(15)}{=} \sinh(\frac{1}{4} \text{sys}(\mathbb{M}_g)) \tan(\frac{1}{2} \angle(O \mathbf{p}_0^0 \mathbf{p}_0^*)), \\ &= (\sqrt{2} - 1) \sinh(\frac{1}{4} \text{sys}(\mathbb{M}_g)). \end{aligned}$$

By the same reasoning as before,

$$\tanh(R(\mathbf{p}_0^1, \mathbf{p}_1^1, \mathbf{p}_0^*)) < (\sqrt{2} - 1) \lim_{g \rightarrow \infty} \sinh(\frac{1}{4} \text{sys}(\mathbb{M}_g)) \approx 0.414,$$

which shows that $\tanh(R(\mathbf{p}_0^1, \mathbf{p}_1^1, \mathbf{p}_0^*)) < \tanh(\frac{1}{4} \text{sys}(\mathbb{M}_g))$. This proves that the circumdiameter of $[\mathbf{p}_0^1, \mathbf{p}_1^1, \mathbf{p}_0^*]$ is smaller than $\frac{1}{2} \text{sys}(\mathbb{M}_g)$.

Thirdly, 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 will not use equation (18), as 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 26 for a more detailed view of the relevant triangles. By the discussion above, we know that $\mathbf{c} \in [O, \mathbf{p}_0^*]$. Let \mathbf{c}' be the orthogonal




Figure 26: Close-up of situation at $[\mathbf{p}_0^1, \mathbf{p}_1^1, \mathbf{p}_1^2]$

projection of \mathbf{c} onto $[O, \mathbf{p}_1^0]$. 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,

$$\tanh([c, c']) = \sinh([O, c']) \tanh(\frac{\pi}{4g}),$$

which after substitution of our expression for $[O, \mathbf{c}']$ can be rewritten as

$$\tanh([c, c']) = \sqrt{1 - \tan^2(\frac{\pi}{4g})} \cosh(\frac{3}{8} \text{sys}(\mathbb{M}_g)) - \sinh(\frac{3}{8} \text{sys}(\mathbb{M}_g)).$$

By the Pythagorean law in $[\mathbf{c}, \mathbf{c}', \mathbf{p}_1^2]$ we know

$$\cosh(R(\mathbf{p}_0^1, \mathbf{p}_1^1, \mathbf{p}_1^2)) = \cosh(\frac{1}{8} \text{sys}(\mathbb{M}_g)) \cosh([c, c']).$$

Since $[O, \mathbf{c}']$ monotonically increases, $[c, \mathbf{c}']$ increases as well, so $R(\mathbf{p}_0^1, \mathbf{p}_1^1, \mathbf{p}_1^2)$ is also strictly increasing. 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 $\cosh([c, c'])$. It follows that $\tanh(R(\mathbf{p}_0^1, \mathbf{p}_1^1, \mathbf{p}_1^2)) < 0.480$, 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 $[p_0^1, p_1^1, p_1^2]$ 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

Lemma 29. *The triangulation \mathcal{T} is a Delaunay triangulation of $\pi_g^{-1}(\mathcal{Q}_g)$.*

Proof. We will 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 $[p, q, r]$ by $C(p, q, r)$ and $D(p, q, r)$ respectively, and similarly when more than three points are concircular. Observe in particular that $C(p, q, r) \cap D(p, q, r) = \emptyset$. Denote the hyperbolic line through points p, q by $L(p, q)$ and the open and closed line segments connecting p, q by $(p, q), [p, q]$ respectively. By symmetry it is sufficient to consider only $D(p_0^0, p_0^1, p_0^*)$, $D(p_0^1, p_1^1, p_0^*)$, $D(p_0^m, p_1^m, O)$, $D(p_0^j, p_1^j, p_1^{j+1})$, for $1 \leq j \leq m-1$. For convenience, we will treat the cases for each circumdisk in a fixed order, namely

1. O, v
2. $p_k^*, k = 0, \dots, 4g-1$,
3. $p_k^j, k = 0, \dots, 4g-1, j = 0, \dots, m$,
4. $q_k^j, k = 0, \dots, 4g-1, j = 1, \dots, m$.

First consider $D(p_0^0, p_0^1, p_0^*)$.

1. Clearly, O and v are too far away from p_0^* to be inside $D(p_0^0, p_0^1, p_0^*)$.
2. Since $p_0^* \in C(p_0^0, p_0^1, p_0^*)$, we know that $p_0^* \notin D(p_0^0, p_0^1, p_0^*)$. Since the center of $D(p_0^0, p_0^1, p_0^*)$ lies inside $D(p_0^0, p_0^1, p_0^*)$ on the bisector of angle $\angle(p_0^0 p_0^1 p_0^*)$ we see that p_1^* is farther away from this center than p_0^* . Since $p_0^* \notin D(p_0^0, p_0^1, p_0^*)$, it follows that $p_1^* \notin D(p_0^0, p_0^1, p_0^*)$ as well. Since $D(p_0^0, p_0^1, p_0^*) \cap C(p_0^*, \dots, p_{4g-1}^*)$ is contained in the shortest open chord of $C(p_0^*, \dots, p_{4g-1}^*)$ between p_0^* and p_1^* , we see that $p_k^* \notin D(p_0^0, p_0^1, p_0^*)$ for $k = 0, \dots, 4g-1$.
3. Since the center of $D(p_0^0, p_0^1, p_0^*)$ is in the interior of $[p_0^0, p_0^1, p_0^*] \subset [O, p_0^0, p_0^*]$, we know that p_0^j is closer to the center of $D(p_0^0, p_0^1, p_0^*)$ than p_k^j for $j = 0, \dots, m$ and $k \neq 0$. Because $L(O, p_0^0) \cap D(p_0^0, p_0^1, p_0^*) = (p_0^0, p_0^1)$, we know that $p_0^j \notin D(p_0^0, p_0^1, p_0^*)$ for all $j = 0, \dots, m$. Therefore, $p_k^j \notin D(p_0^0, p_0^1, p_0^*)$ for $j = 0, \dots, m$ and $k = 0, \dots, 4g-1$.
4. By a reasoning similar to above, p_k^j is closer to the center of $D(p_0^0, p_0^1, p_0^*)$ than q_k^j . Since by the previous step $p_k^j \notin D(p_0^0, p_0^1, p_0^*)$ for $j = 1, \dots, m$ and $k = 0, \dots, 4g-1$, it follows that $q_k^j \notin D(p_0^0, p_0^1, p_0^*)$ for $j = 1, \dots, m$ and $k = 0, \dots, 4g-1$ as well.

Secondly, consider $D(p_0^1, p_1^1, p_0^*)$.

1. Clearly, O and v are too far away from p_0^* to be inside $D(p_0^1, p_1^1, p_0^*)$.
2. By mirror symmetry in the line $L(O, p_0^*)$, $C(p_0^1, p_1^1, p_0^*)$ is tangent to $C(p_0^*, \dots, p_{4g-1}^*)$. Since the radius of $C(p_0^1, p_1^1, p_0^*)$ is smaller than the radius of $C(p_0^*, \dots, p_{4g-1}^*)$, this means that $D(p_0^1, p_1^1, p_0^*) \subseteq D(p_0^*, \dots, p_{4g-1}^*)$. Therefore, $p_k^* \notin D(p_0^1, p_1^1, p_0^*)$ for all $k = 0, \dots, 4g-1$.
3. First, to prove that $p_0^j \notin D(p_0^1, p_1^1, p_0^*)$ for all $j = 0, \dots, m$, we will now show that $D(p_0^1, p_1^1, p_0^*) \cap (O, p_0^1) = \emptyset$. First observe that the line $L(O, p_0^0)$ will intersect $C(p_0^1, p_1^1, p_0^*)$ in one or two points. If the intersection consists of one point, then it has to be p_0^1 and we are done. If the intersection consists of two points, then it is sufficient to show that p_0^1 is the closest of these two to O . Let p_0^M denote the midpoint of p_0^1 and p_1^1 . It is sufficient to show that $[p_0^*, p_0^M] \geq R(p_0^1, p_1^1, p_0^*)$. We have

$$\tanh([p_0^*, p_0^M]) = \tanh([p_0^1, p_0^*]) \cos \angle(p_0^1 p_0^* p_1^1) = \tanh([p_0^1, p_0^*]) \sin \angle(p_0^1 p_0^* p_0^0),$$

where the second equality is by equation (17) in the proof of Lemma 28. 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 (15) in the proof of Lemma 28, 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

$$\begin{aligned} \sin^2 \angle(\mathbf{p}_0^1 \mathbf{p}_0^* \mathbf{p}_0^0) &\stackrel{(16)}{=} \frac{\sinh^2(\frac{1}{4} \text{sys}(\mathbb{M}_g)) \sin^2 \angle(O \mathbf{p}_0^0 \mathbf{p}_0^*)}{\sinh^2([\mathbf{p}_0^1, \mathbf{p}_0^*])}, \\ &\stackrel{(15)}{=} \frac{\sinh^2(\frac{1}{2}[\mathbf{p}_0^1, \mathbf{p}_0^*]) \sin^2 \angle(O \mathbf{p}_0^0 \mathbf{p}_0^*)}{\sin^2(\frac{1}{2} \angle(O \mathbf{p}_0^0 \mathbf{p}_0^*)) \sinh^2([\mathbf{p}_0^1, \mathbf{p}_0^*])}, \\ &= \frac{\cos^2(\frac{1}{2} \angle(O \mathbf{p}_0^0 \mathbf{p}_0^*))}{\cosh^2(\frac{1}{2}[\mathbf{p}_0^1, \mathbf{p}_0^*])}, \end{aligned}$$

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 $j = 0, \dots, m$. Secondly, 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 $C(\mathbf{p}_0^*, \dots, \mathbf{p}_{4g-1}^*)$, which means that $D(\mathbf{p}_0^1, \mathbf{p}_1^1, \mathbf{p}_0^*)$ is contained in the interior of the $4g$ -gon $[\mathbf{p}_0^0, \dots, \mathbf{p}_k^0, \dots, \mathbf{p}_{4g-1}^0]$. Therefore, $\mathbf{q}_k^j \notin D(\mathbf{p}_0^1, \mathbf{p}_1^1, \mathbf{p}_0^*)$ for $j = 1, \dots, m$ and $k = 0, \dots, 4g-1$.

Thirdly, 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 disk $D(\mathbf{p}_0^m, \mathbf{p}_1^m, \dots, \mathbf{p}_{4g-1}^m)$ and the (closed) triangle $[O, \mathbf{p}_0^0, \mathbf{p}_1^0]$, we immediately 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 $L(O, \mathbf{p}_0^0) \cap D(\mathbf{p}_0^m, \mathbf{p}_1^m, O) = (O, \mathbf{p}_0^0)$, 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, $O \notin D(\mathbf{p}_0^j, \mathbf{p}_1^j, \mathbf{p}_1^{j+1})$, since $L(O, \mathbf{p}_1^0) \cap D(\mathbf{p}_0^j, \mathbf{p}_1^j, \mathbf{p}_1^{j+1}) = (\mathbf{p}_1^j, \mathbf{p}_1^{j+1})$.
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 $C(\mathbf{p}_0^*, \mathbf{p}_1^0, \mathbf{p}_1^1)$ 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 $C(\mathbf{p}_0^*, \mathbf{p}_1^0, \mathbf{p}_1^1)$, 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 $[O, \mathbf{p}_0^0, \mathbf{p}_1^0]$ 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 $L(O, \mathbf{p}_0^0) \cap D(\mathbf{p}_0^j, \mathbf{p}_1^j, \mathbf{p}_1^{j+1}) = (\mathbf{p}_0^j, \mathbf{p}_0^{j+1})$, 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 $[\mathbf{p}_0^0, \dots, \mathbf{p}_k^0, \dots, \mathbf{p}_{4g-1}^0]$, 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

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

Theorem 30. *The structured algorithm terminates. The resulting dummy point \mathcal{Q}_g set satisfies $\delta(\mathcal{Q}_g) < \frac{1}{2} \text{sys}(\mathbb{M}_g)$ 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 28 the resulting dummy point set satisfies $\delta(\mathcal{Q}_g) < \frac{1}{2} \text{sys}(\mathbb{M}_g)$. 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

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

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 29, \mathcal{T} is a Delaunay triangulation. This finishes the proof. \square

References

- [1] William Abikoff. The Uniformization Theorem. *The American Mathematical Monthly*, 88(8):574–592, 1981. doi:10.2307/2320507.
- [2] 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.
- [3] Ralf Aurich and Frank Steiner. On the periodic orbits of a strongly chaotic system. *Physica D: Nonlinear Phenomena*, 32(3):451–460, 1988.
- [4] Alan F. Beardon. *The geometry of discrete groups*. Springer-Verlag New York, 1 edition, 1983. doi:10.1007/978-1-4612-1146-4.
- [5] 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.
- [6] 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.
- [7] A. Bowyer. Computing Dirichlet tessellations. *The Computer Journal*, 24(2):162–166, 1981. doi:10.1093/comjnl/24.2.162.
- [8] P. Buser. *Geometry and Spectra of Compact Riemann Surfaces*. Progress in Mathematics Series. Birkhäuser, 1992.
- [9] 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>.
- [10] 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.
- [11] 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.
- [12] M. Dehn. Transformation der Kurven auf zweiseitigen Flächen. *Mathematische Annalen*, 72(3):413–421, 1912. doi:10.1007/BF01456725.
- [13] 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.

- [14] 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>.
- [15] Vincent Despré, Benedikt Kolbe, and Monique Teillaud. Half-minimizers and Delaunay triangulations on closed hyperbolic surfaces. Technical report. URL: <https://hal.archives-ouvertes.fr/hal-03045921>.
- [16] Matthijs Ebbens. Delaunay triangulations on hyperbolic surfaces. Master’s thesis, University of Groningen, 2017. URL: <http://fse.studenttheses.ub.rug.nl/id/eprint/15727>.
- [17] Matthijs Ebbens, Hugo Parlier, and Gert Vegter. Minimal Delaunay triangulations of hyperbolic surfaces, 2020. [arXiv:2011.09847](https://arxiv.org/abs/2011.09847).
- [18] Hershel M Farkas and Irwin Kra. Riemann surfaces. In *Riemann surfaces*, pages 9–31. Springer, 1992.
- [19] 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.
- [20] H. Helling. Diskrete Untergruppen von $SL(2, \mathbb{R})$. *Inventiones mathematicae*, 17(3):217–229, 1972. doi:10.1007/BF01425449.
- [21] 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>.
- [22] 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.
- [23] 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>.
- [24] 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>.
- [25] 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>.
- [26] 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.
- [27] 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.
- [28] Jim Ruppert. A Delaunay refinement algorithm for quality 2-dimensional mesh generation. *Journal of algorithms*, 18(3):548–585, 1995.
- [29] 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.

- [30] 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.
- [31] Chee Yap *et al.* The CORE library project. URL: http://cs.nyu.edu/exact/core_pages/intro.html.