



HAL
open science

Delaunay triangulations of generalized Bolza surfaces

Matthijs Ebbens, Jordan Iordanov, Monique Teillaud, Gert Vegter

► **To cite this version:**

Matthijs Ebbens, Jordan Iordanov, Monique Teillaud, Gert Vegter. Delaunay triangulations of generalized Bolza surfaces. *Journal of Computational Geometry*, 2022, 13 (1), pp.125-177. 10.20382/jocg.v13i1a5 . hal-03664678

HAL Id: hal-03664678

<https://inria.hal.science/hal-03664678v1>

Submitted on 11 May 2022

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[†] Jordan Iordanov[‡] Monique Teillaud[§] Gert Vegter[¶]

ABSTRACT. The Bolza surface can be seen as the quotient of the hyperbolic plane, represented by the Poincaré disk model, under the action of the group generated by the hyperbolic isometries identifying opposite sides of a regular octagon centered at the origin. We consider *generalized* Bolza surfaces \mathbb{M}_g , where the octagon is replaced by a 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 small sets of points on \mathbb{M}_g that are used to initialize Bowyer’s algorithm.

1 Introduction

Delaunay triangulations in d -dimensional Euclidean space are among the most important structures in computational geometry; they find applications in a vast array of domains of science and technology. They are well-studied and their mathematical properties are well understood. Many algorithms to compute them have been proposed and CGAL, the Computational Geometry Algorithms Library [18], includes packages providing efficient implementations; the impact of these packages extends far beyond computational geometry.¹ Work on the computation of Delaunay triangulations of hyperbolic surfaces is motivated by applications in various fields such as physics [49, 7], neuro-mathematics [19], and materials science [30, 31, 32].

Lawson’s well-known incremental algorithm that computes Delaunay triangulations using edge flips in the Euclidean plane [41] has recently been proved to generalize on hyperbolic surfaces [22]. However, the experience gained in the CGAL project for many years has shown that Bowyer’s algorithm [12] leads to a cleaner code, much easier to maintain: in particular, it allows for a strict separation between combinatorial operations and geometric computations, which are handled by separate classes. 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. Delaunay triangulations on the sphere are also using Bowyer’s algorithm [15, 47].

*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. ymebbens@gmail.com

[‡]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

¹See <http://www.cgal.org/projects.html> for some application fields.

In the context of quotient spaces Bowyer’s algorithm is used already in the CGAL packages for 3D flat quotient spaces [16] (it will also be used in the forthcoming package handling general cases [45]) and for the Bolza surface [38]. To the best of our knowledge, the latter package is the only available software for a hyperbolic surface. The advantages of Bowyer’s algorithm largely compensate for the constraint that it intrinsically requires that the Delaunay triangulation be a simplicial complex.

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 we recall the following *validity condition* [17, 10]: If a finite set \mathcal{Q} of points on the surface \mathbb{M} satisfies the inequality

$$\delta(\mathcal{Q}) < \frac{1}{2} \text{sys}(\mathbb{M}) \quad (\text{condition (10) in Proposition 3})$$

then Bowyer’s algorithm can be extended to the computation of the Delaunay triangulation of any finite set of points \mathcal{S} on \mathbb{M} containing \mathcal{Q} . Before actually inserting the input points, the algorithm performs a preprocessing step consisting of computing the Delaunay triangulation of an appropriate (but small) set \mathcal{Q} satisfying this validity condition; following the terminology of previous papers [10, 17], we refer to the points of \mathcal{Q} as *dummy points*. When sufficiently many and well-distributed input points have been inserted, the validity condition is satisfied without the dummy points, which can then be removed. This approach was used in the CGAL package for the Bolza surface \mathbb{M}_2 [37, 38].

Other practical approaches for (flat) quotient spaces start by computing in a finite-sheeted covering space [17] or in the universal covering space [45], 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 [21].

Results.

We describe the extension of Bowyer’s algorithm to the case of the generalized Bolza surface \mathbb{M}_g in Section 3, and we derive bounds on the number of dummy points necessary 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 (10) grows as $\Theta(g)$.*

In Section 4, we give a proof for the value of the systole of \mathbb{M}_g , announced in [26] and [27, Theorem 3.1]:

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

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

The same result was also recently obtained by Bai *et al.* [6]. Our proof gives more insight into the representation in D_g of closed geodesics on \mathbb{M}_g . Note also that this result generalizes an earlier work of Aurich and Steiner [5], which was restricted to the Bolza surface \mathbb{M}_2 .

Then, in Section 5, we propose two algorithms to compute dummy points. The first algorithm is based on the well-known Delaunay refinement algorithm for mesh generation [48]. Using a packing argument, we prove that it provides an asymptotically optimal number of dummy points (Theorem 14). The second algorithm modifies the refinement algorithm so as to yield a symmetric dummy point set, at the expense of a slightly larger output size $\Theta(g \log g)$ (Theorem 15); this symmetry may be interesting for some applications [19]. The two algorithms have been implemented and we quickly present results for small genera $g = 2$ and $g = 3$.

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 predicates used in the computations and present experimental results.

2 Mathematical preliminaries

In this section we define our notation and present a short introduction on hyperbolic geometry and hyperbolic surfaces [8, 46].

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 $K = -1$ [8]. 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}). The geodesic segment $[\mathbf{z}, \mathbf{w}]$ between points $\mathbf{z}, \mathbf{w} \in \mathbb{D}$ is the (unique) shortest curve connecting \mathbf{z} and \mathbf{w} . A

hyperbolic line (i.e., a geodesic for the given metric) in this model is a curve which contains the geodesic segment between any two of its points. These geodesics are diameters of \mathbb{D} or circle arcs whose supporting lines or circles intersect \mathbb{D}_∞ orthogonally (see the leftmost frame of Figure 1). A circle in the hyperbolic plane is a Euclidean circle in the Poincaré disk, in general with a hyperbolic center and radius that are different from their Euclidean counterparts.

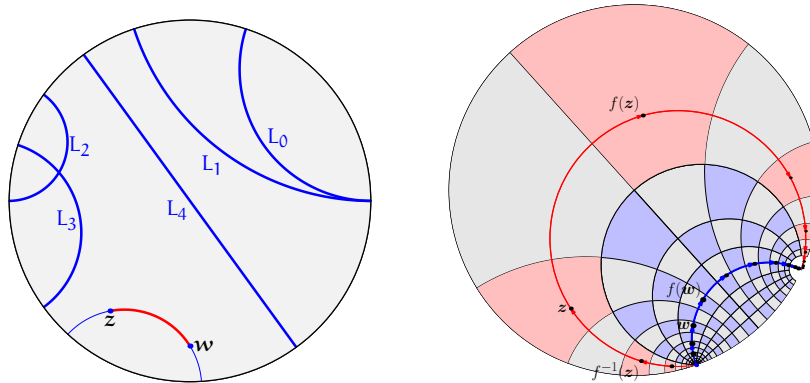


Figure 1: Left: the Poincaré disk model \mathbb{D} of the hyperbolic plane, with some geodesics. The boundary \mathbb{D}_∞ does not belong to \mathbb{D} , but consists of ideal points of \mathbb{D} . Geodesics L_0 and L_1 are parallel (have an ideal point in common), L_2 and L_3 are intersecting and 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 transformations 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 $\mathbf{z} \mapsto d(\mathbf{z}, f(\mathbf{z}))$, which is attained at all points \mathbf{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 . We will always assume our hyperbolic surfaces to be compact. 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 covering projection $\pi : \mathbb{D} \rightarrow \mathbb{D}/\Gamma$ is a local isometry. The orbit $\Gamma\mathbf{z}$ of a point $\mathbf{z} \in \mathbb{D}$ is a discrete subset of \mathbb{D} . Note that $\Gamma\mathbf{z} = \pi^{-1}(z)$, 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 hyperbolic surface all non-identity elements of Γ are hyperbolic translations.

The distance between points p and q on a hyperbolic surface is given by the distance $\min\{d(\mathbf{p}, \mathbf{q}) \mid \mathbf{p} \in \pi^{-1}(p), \mathbf{q} \in \pi^{-1}(q)\}$ between the orbits of p and q in \mathbb{D} 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 $[\mathbf{z}, f(\mathbf{z})]$ on the axis of a hyperbolic translation f such that π is injective on the open segment $(\mathbf{z}, f(\mathbf{z}))$. 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 [13, 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 closure of the open cell of p in the Voronoi diagram 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.

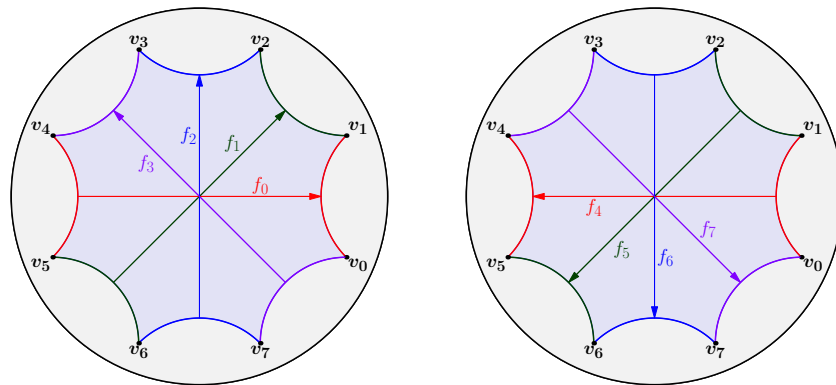


Figure 2: The side-pairings f_0, \dots, f_3 of the Bolza surface (of genus 2) pair opposite edges of the fundamental octagon (a regular octagon in \mathbb{D} with angles $\frac{1}{4}\pi$). Their inverses f_4, \dots, f_7 satisfy $f_{k+4} = f_k^{-1}$. The side-pairings generate the Fuchsian group Γ_2 .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

2.4 Generalized Bolza surfaces

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

$$A_j = \begin{pmatrix} \cot(\frac{\pi}{4g}) & \exp(\frac{ij\pi}{2g})\sqrt{\cot^2(\frac{\pi}{4g}) - 1} \\ \exp(-\frac{ij\pi}{2g})\sqrt{\cot^2(\frac{\pi}{4g}) - 1} & \cot(\frac{\pi}{4g}) \end{pmatrix}. \tag{5}$$

By Poincaré’s Theorem ([8, Chapter 9.8] and [46, Chapter 11.2]) these side-pairings generate a Fuchsian group, the generalized Bolza group Γ_g , all non-identity elements of which are

hyperbolic translations. The polygon D_g is a fundamental domain for the action of this group, and it is even the Dirichlet region of the origin.

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

$$f_0 f_1^{-1} f_2 f_3^{-1} \cdots f_{4g-2} f_{4g-1}^{-1} = \mathbb{1}. \quad (6)$$

For even j we have $f_j = f_{j(2g+1)}$, since we are counting indices modulo $4g$. Similarly, $f_j^{-1} = f_{j+2g} = f_{j(2g+1)}$, for odd j . Therefore, we can rewrite (6) as

$$\prod_{j=0}^{4g-1} f_{j(2g+1)} = f_0 f_{2g+1} f_{2(2g+1)} \cdots f_{(4g-1)(2g+1)} = \mathbb{1}. \quad (7)$$

The order of the factors in this product does matter since the group Γ is not abelian. Equation (7) is usually called the *relation* of Γ_g . In addition to (6) and (7), there are many other ways to write the relation. By rotational symmetry of D_g , conjugating $\prod_{j=0}^{4g-1} f_{j(2g+1)}$ with the rotation by angle $k\pi/2g$ around the origin yields the relation $\prod_{j=0}^{4g-1} f_{k+j(2g+1)} = \mathbb{1}$. The latter expression can be rewritten as

$$f_k f_{k+1}^{-1} f_{k+2} f_{k+3}^{-1} \cdots f_{k+4g-2} f_{k+4g-1}^{-1} = \mathbb{1}. \quad (8)$$

Neighbors of vertices of the fundamental polygon. In the clockwise sequence of Dirichlet regions $h_1(D_g), h_2(D_g), \dots, h_{4g}(D_g)$ around vertex \mathbf{v}_k the element $h_j \in \Gamma_g$ is the prefix of length j in the left-hand side of (8):

$$h_j = \begin{cases} f_k f_{k+1}^{-1} \cdots f_{k+j-2} f_{k+j-1}^{-1}, & \text{if } j \text{ is even,} \\ f_k f_{k+1}^{-1} \cdots f_{k+j-2}^{-1} f_{k+j-1}, & \text{if } j \text{ is odd.} \end{cases} \quad (9)$$

Prefixes h_j of length $j \geq 2g$ can be reduced to a word of length $4g-j$ in $f_k, f_{k+1}^{-1}, \dots, f_{k+2g-1}^{-1}$ using relation (8) (where the empty word – of length zero – corresponds to $\mathbb{1}$) and the fact that $f_j = f_{j-2g}^{-1}$ for $j \geq 2g$. More precisely, h_{4g-j} is the prefix of length j in the product $f_{k+2g-1}^{-1} f_{k+2g-2} \cdots f_{k+1}^{-1} f_k$, for $j = 0, \dots, 2g$. Figure 3 depicts the neighbors of \mathbf{v}_k for the case $g = 2$.

The ordering of neighbors of the vertices of D_g yields an ordering of all regions around D_g , which will play an important role in the data structure for the representation of Delaunay triangulations in Section 6. More precisely, we define the set \mathcal{N}_g of *neighboring translations* as:

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

Each Dirichlet region sharing an edge or a vertex with the (closed) domain D_g is the image of D_g under the action of a translation in \mathcal{N}_g , which is used to label the region. Also see Figure 3. We denote the union of these *neighboring regions* of D_g by $D_{\mathcal{N}_g}$, so

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

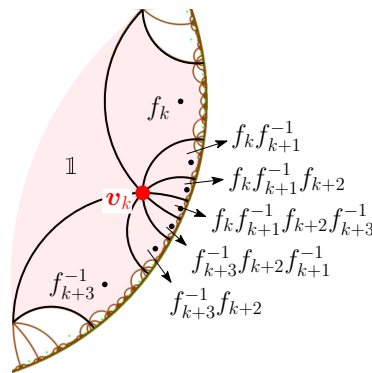


Figure 3: Enumeration of the regions around a vertex v_k , $k = 0, 1, \dots, 7$, for the case $g = 2$. Note that $f_{k+3}^{-1} f_{k+2} f_{k+1}^{-1} f_k = f_k f_{k+1}^{-1} f_{k+2} f_{k+3}^{-1}$ by the group relation.

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

The hyperbolic surface \mathbb{M}_g . The *generalized Bolza surface* of genus g is the hyperbolic surface \mathbb{D}/Γ_g , denoted by \mathbb{M}_g . The projection map is $\pi_g : \mathbb{D} \rightarrow \mathbb{D}/\Gamma_g$. The surface \mathbb{M}_2 is the classical Bolza surface [11, 4].

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

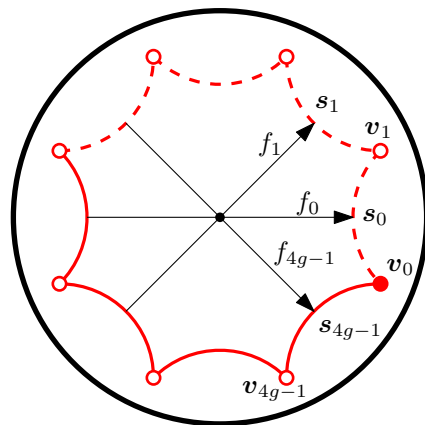


Figure 4: Original domain \tilde{D}_g for $g = 2$. Only vertex v_0 and the solid sides are included in \tilde{D}_g .

For a point p on \mathbb{M}_g , the *canonical* representative \mathbf{p}^c of p is the unique point of the orbit $\pi_g^{-1}(p)$ that lies in \tilde{D}_g . Similarly, the canonical representative of a point \mathbf{p} in \mathbb{D} is the unique point of the orbit $\Gamma_g \mathbf{p}$ that lies in \tilde{D}_g , which, by a slight abuse of notation, is also denoted by \mathbf{p}^c .

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 [12, 51] has proved its efficiency in CGAL [39].

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

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

Degeneracies can be resolved using a symbolic perturbation technique, which actually works in any dimension [24].

An illustration is given in Figure 5.

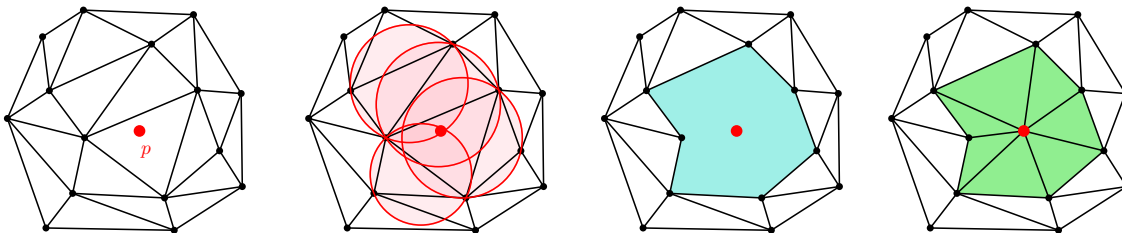


Figure 5: Insertion of a point in a Delaunay triangulation with Bowyer's incremental algorithm.

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

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

3.2 Delaunay triangulations of points on hyperbolic surfaces

Let $\mathbb{M} = \mathbb{D}/\Gamma$ be a hyperbolic surface, as introduced in Section 2.2, 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 [9], the notion of conflict in \mathbb{D} is the same as in \mathbb{E} , since for the Poincaré disk model, hyperbolic circles are Euclidean circles (see Section 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 6). 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 .

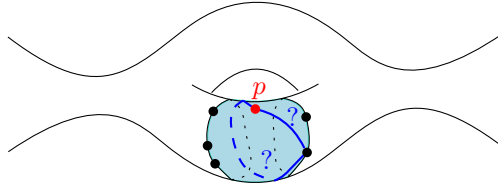


Figure 6: Bowyer's insertion is not well defined when the conflict region is not a topological disk.

In order to be able to use Bowyer's algorithm on \mathbb{M} , the triangles on \mathbb{M} without conflict with any vertex, together with their edges and their vertices, should form a *simplicial complex*. Here, a collection \mathcal{K} of vertices, edges, and triangles (together called *simplices*) is called a simplicial complex if it satisfies the following two conditions (cf [3, Chapter 6] and [43, Chapter 1]):

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

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 the discussion in [17, Section 2.1]).

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

Proposition 3 (Validity condition [10]). *Let $\mathcal{Q} \subset \mathbb{M}$ be a set of points such that*

$$\delta(\mathcal{Q}) < \frac{1}{2} \text{sys}(\mathbb{M}). \quad (10)$$

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

The proof is easy, we include it for completeness.

Proof. Assume that condition (10) 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 (10). The *Delaunay triangulation of \mathbb{M} defined by \mathcal{S}* is then defined as $\pi(\text{DT}_{\mathbb{D}}(\pi^{-1}(\mathcal{S})))$ and denoted by $\text{DT}_{\mathbb{M}}(\mathcal{S})$.

As for the Bolza surface [10], 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 (10);
- Compute incrementally the Delaunay triangulation $\text{DT}_{\mathbb{M}}(\mathcal{Q} \cup \mathcal{P})$ by inserting the points $p_1, p_2, \dots, p_k, \dots, p_n$ of \mathcal{P} one by one, i.e., for each new point p_k :
 - find all triangles of the Delaunay triangulation $\text{DT}_{\mathbb{M}}(\mathcal{Q} \cup \{p_1, \dots, p_{k-1}\})$ that are in conflict with p_k ; let C_{p_k} denote their union; since \mathcal{Q} satisfies the validity condition, C_{p_k} is a topological disk;
 - delete the triangles in C_{p_k} ;
 - repair the triangulation by forming new triangles with p_k and each edge of the boundary of C_{p_k} ;
- Remove from the triangulation all points of \mathcal{Q} whose removal does not violate the validity condition.

We ignore degeneracies here; they can be resolved as in the case of flat orbit spaces [17]. Depending on the size and distribution of the input set \mathcal{P} , the final Delaunay triangulation of \mathbb{M} might have some or all of the dummy points as vertices. If \mathcal{P} already satisfies the validity condition then no dummy point is left.

3.3 Bounds on the number of dummy points

In the following proposition we show that a dummy point set exists and give an upper bound for its cardinality. The proof is non-constructive, but we will construct dummy point sets for generalized Bolza surfaces in Section 5.

Proposition 5. *Let \mathbb{M} be a hyperbolic surface of genus g with systole $\text{sys}(\mathbb{M})$. Then there exists a point set $\mathcal{Q} \subset \mathbb{M}$ satisfying the validity condition (10) with cardinality*

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

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

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

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

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

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

□

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

Proposition 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 (10) holds. Then*

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

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

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

so

$$k_2 = 4g - 4 + 2k_0. \quad (11)$$

Consider an arbitrary triangle t in $\text{DT}_{\mathbb{M}}(\mathcal{Q})$. Because the validity condition holds, the circumradius of t is smaller than $\frac{1}{2} \text{sys}$. It can be shown that $\text{area}(t) < \pi - 6 \operatorname{arccot}(\sqrt{3} \cosh(\frac{1}{4} \text{sys}))$; this is Lemma 19 in Appendix A. Because \mathbb{M} has area $4\pi(g-1)$, it follows that

$$k_2 > \frac{4\pi(g-1)}{\pi - 6 \operatorname{arccot}(\sqrt{3} \cosh(\frac{1}{2} \text{sys}))}. \quad (12)$$

Combining (11) and (12) yields the result. \square

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

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

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

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

3. Finally, consider the case when $\text{sys}(\mathbb{M}) \rightarrow \infty$ when $g \rightarrow \infty$. Since $\text{sys}(\mathbb{M}) \leq 2 \log(4g-2)$ for all hyperbolic surfaces of genus g (see Equation (4) in Section 2.2), we only consider the case where $\text{sys}(\mathbb{M}) \sim C \log g$ for some C with $0 < C \leq 2$. In fact, there exist families of surfaces for which $\text{sys}(\mathbb{M}) > \frac{4}{3} \log g - c$ for some constant $c \in \mathbb{R}$ for infinitely many genera g . See [14, page 45] and [40]. In this case, we can use $\cosh x \sim \frac{1}{2}e^x$ to deduce that our upper bound reduces to an expression of order $O(g^{1-\frac{1}{8}C})$. Similarly, by considering the Taylor expansion of the coefficient in the lower bound we see that the lower bound has cardinality $\Omega(g^{1-\frac{1}{4}C})$. Of our three cases, this is the only case in which there is a gap between the stated upper and lower bound.

Remark 7. Note that the validity condition (10) 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 [29]. 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 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 differs asymptotically as well.

4 Proof of Theorem 2: Systole of generalized Bolza surfaces

In the previous section we have recalled the validity condition (10), allowing us to define the Delaunay triangulation $DT_{\mathbb{M}}(\mathcal{S})$. 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 the introduction, which gives the value of the systole for the generalized Bolza surfaces \mathbb{M}_g defined in Section 2.4.

As a preparation for the proof we show in Section 4.1 how to represent a simple closed geodesic γ on \mathbb{M}_g by a sequence γ of pairwise disjoint hyperbolic line segments between sides of the fundamental domain D_g . The length of γ is equal to the sum of the lengths of the line segments in γ .

The proof consists of two parts. In Section 4.2 we show that $\text{sys}(\mathbb{M}_g) \leq \varsigma_g$ by constructing a simple (non-contractible) closed geodesic of length ς_g . In Section 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 representing γ . This shows that $\text{sys}(\mathbb{M}_g) \geq \varsigma_g$.

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

Consider a simple closed geodesic γ on the generalized Bolza surface \mathbb{M}_g . Because D_g is compact, there is a finite number, say m , of pairwise disjoint hyperbolic lines intersecting D_g in the preimage $\pi_g^{-1}(\gamma)$ of γ . See the leftmost panel in Figure 7. These hyperbolic lines are the axes of conjugated elements of Γ_g . Therefore, the intersection of $\pi_g^{-1}(\gamma)$ with D_g consists of m pairwise disjoint hyperbolic line segments between the sides of D_g , the union of which we denote by γ . See the rightmost panel in Figure 7. These line segments are oriented and their orientations are compatible with the orientation of γ . In particular, every line segment has a starting point and an endpoint. Since D_g is a fundamental domain for Γ_g , the π_g -images of these line segments form a covering of the closed geodesic γ by m closed subsegments with pairwise disjoint interiors. In other words, these projected segments lie side-by-side on γ , so they form a (cyclically) ordered sequence. This cyclic order lifts to an order $\gamma_1, \dots, \gamma_m$ of the m segments in D_g , which together represent the simple closed geodesic γ . More precisely:

Definition 8. An oriented simple closed geodesic γ on \mathbb{M}_g is *represented by a sequence of oriented geodesic segments* $\gamma_1, \dots, \gamma_m$ in D_g if (i) the starting point and endpoint of each segment lie on different sides of ∂D_g , and (ii) the projections $\pi_g(\gamma_1), \dots, \pi_g(\gamma_m)$ are oriented closed subsegments of γ that cover γ , have pairwise disjoint interiors, and lie side-by-side on γ in the indicated order.

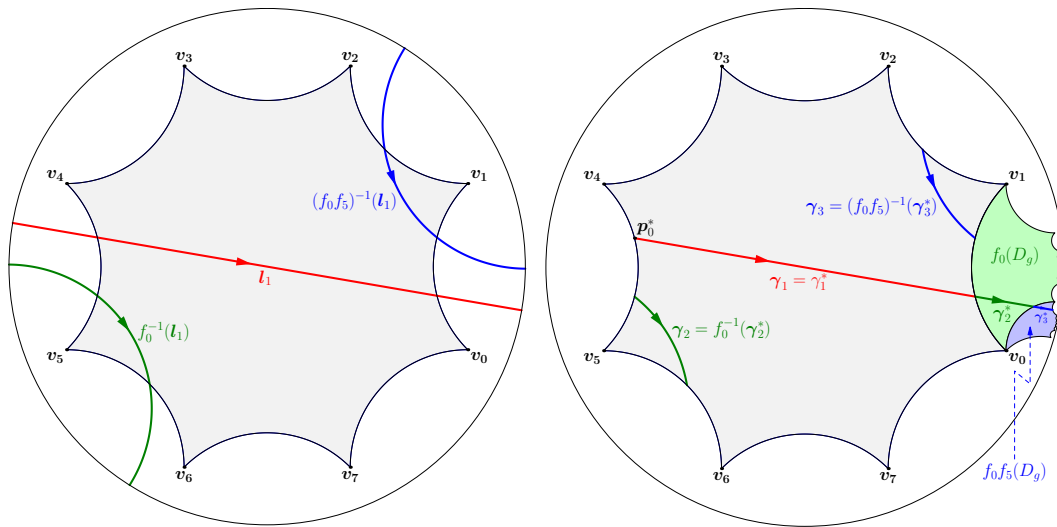


Figure 7: Left: Connected components of the preimage of an oriented simple closed geodesic γ on the Bolza surface intersecting the fundamental octagon D_2 . The geodesic is the projection onto the Bolza surface of the axis \mathbf{l}_1 of $f = f_0 f_5 f_0$ (and of the axes of all the conjugates of f in the Bolza group).

Right: The cyclic sequence of geodesic segments $\gamma_1, \gamma_2, \gamma_3$ in D_g represents γ . The segments γ_1^*, γ_2^* and γ_3^* are the successive intersections of the axis of f with $F_0(D_g), F_1(D_g)$ and $F_2(D_g)$, where $F_0 = \mathbb{1}, F_1 = f_0$ and $F_2 = f_0 f_5$. The endpoint of γ_3^* is $F_3(\mathbf{p}_0^*)$, where $F_3 = f$ and \mathbf{p}_0^* is the starting point of γ_1^* . The endpoint of γ_k is paired with the starting point of γ_{k+1} by the side-pairing $F_k^{-1} F_{k-1}$. In this example $F_1^{-1} F_0 = f_0^{-1} = f_4, F_2^{-1} F_1 = f_5^{-1} = f_1$, and $F_3^{-1} F_2 = f_0^{-1} = f_4$.

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

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

It is convenient to consider $f(\mathbf{p}_0^*)$ as the starting point of the segment $\mathbf{l}_1 \cap f(D_g)$, which we denote by γ_{m+1}^* . Taking $F_m = f$, we see that $\gamma_{m+1}^* = \mathbf{l}_1 \cap F_m(D_g)$. Extending

our earlier definition $\gamma_k = F_{k-1}^{-1}(\gamma_k^*)$ to $k = m + 1$, we see that γ_{m+1} is the subsegment of $\pi_g^{-1}(\gamma) \cap D_g$ with starting point $F_m^{-1}(f(\mathbf{p}_0^*)) = \mathbf{p}_0^*$, so $\gamma_{m+1} = \gamma_1$.

Finally, we show that the endpoint of γ_k is mapped to the starting point of γ_{k+1} by a side-pairing transformation of D_g , for $k = 1, \dots, m$. Since $F_{k-1}(D_g) \cap F_k(D_g)$ is a side of $F_k(D_g)$, for $k = 1, \dots, m$, the intersection $F_k^{-1}F_{k-1}(D_g) \cap D_g$ is a side of D_g , say \mathbf{s}_{j_k} . Then $F_k^{-1}F_{k-1} = f_{j_k}$, since $\mathbf{s}_{j_k} = f_{j_k}(D_g) \cap D_g$. Let $\gamma_k^* = [\mathbf{p}_{k-1}^*, \mathbf{p}_k^*]$, then $\gamma_k = [F_{k-1}^{-1}(\mathbf{p}_{k-1}^*), F_{k-1}^{-1}(\mathbf{p}_k^*)]$. Therefore, f_{j_k} maps the endpoint $F_{k-1}^{-1}(\mathbf{p}_k^*)$ of γ_k to the starting point $F_k^{-1}(\mathbf{p}_k^*)$ of γ_{k+1} , since $f_{j_k}F_{k-1}^{-1} = F_k^{-1}$. See the rightmost panel in Figure 7.

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

Proof. The axis of the hyperbolic translation $f_{2g+1}f_0$ projects onto a simple closed geodesic γ on \mathbb{M}_g with length equal to the translation length $\ell(f_{2g+1}f_0)$. See Section 2.2. Since $f_{2g+1}f_0$ is represented by the matrix $A_{2g+1}A_0$, with A_j given by (5), we see that

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

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

Remark 10. Two connected components of the pre-image $\pi_g^{-1}(\gamma)$ of the simple closed geodesic γ , appearing in the proof of Lemma 9, intersect the fundamental polygon D_g : the axis \mathbf{l}_1 of $f = f_{2g+1}f_0$, and the geodesic $\mathbf{l}_2 = f_{2g+1}^{-1}(\mathbf{l}_1)$, which is the axis of f_0f_{2g+1} . The geodesic γ is represented by the segments $\gamma_1 = \mathbf{l}_1 \cap D_g$ and $\gamma_2 = \mathbf{l}_2 \cap D_g$. The first segment connects the midpoint \mathbf{m}_{2g} of \mathbf{s}_{2g} and the midpoint \mathbf{m}_{2g+1} of \mathbf{s}_{2g+1} , whereas the second segment connects the midpoints of \mathbf{s}_0 and \mathbf{s}_1 . See Figure 8.

This can be seen as follows. Since $f = f_{2g+1}f_{2g}^{-1}$ and the axes of f_{2g} and f_{2g+1} intersect at the origin O , (the proof of) Theorem 7.38.6 of [8] implies that the axis of f passes through the midpoint of the segment $[O, f_{2g}(O)]$ and the midpoint of $[O, f_{2g+1}(O)]$. But these midpoints coincide with \mathbf{m}_{2g} and \mathbf{m}_{2g+1} , respectively, so $[\mathbf{m}_{2g}, \mathbf{m}_{2g+1}] = \gamma_1$. This theorem also implies that the length of the latter segment is half the translation length of f , i.e., $\frac{1}{2}\varsigma_g$. A similar argument shows that the length of γ_2 is $\frac{1}{2}\varsigma_g$.

4.3 Lower bound for the systole

We now prove that the length of every simple closed geodesic of \mathbb{M}_g is at least ς_g , or, equivalently, that the total length of the segments representing such a geodesic is at least ς_g . To this end we consider different types of closed geodesics based on which “kind” of segments are contained in the sequence. We say that an oriented hyperbolic line segment between two sides of D_g is a *k-segment*, $1 \leq k \leq 4g - 1$, if its starting point and endpoint are contained in \mathbf{s}_j and \mathbf{s}_{j+k} , respectively, for some j with $0 \leq j \leq 4g - 1$, where indices are

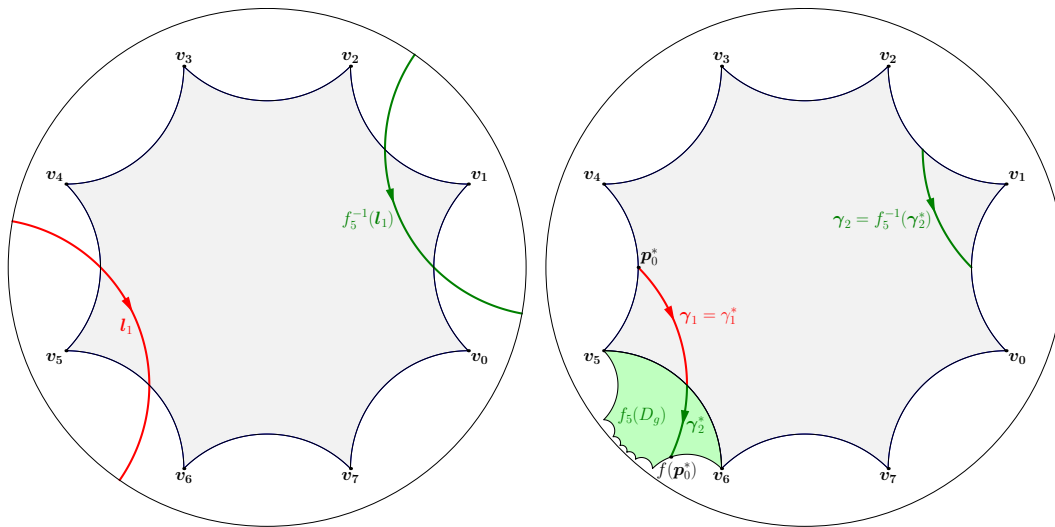


Figure 8: Two geodesics in the pre-image of γ intersect the fundamental polygon D_g (left) The intersections are the segments γ_1 and γ_2 , which represent γ (right). The figure illustrates the situation for the Bolza surface ($g = 2$).

counted modulo $4g$. Furthermore, we say that the segment is k -separated or has separation k , $1 \leq k \leq 2g$, if either the segment itself or the segment with the opposite orientation is a k -segment. Equivalently, a k -separated segment is either a k -segment or a $(4g - k)$ -segment. For example, both segments in Figure 8 - Right are 1-separated, but γ_1 is a 1-segment while γ_2 is a 7-segment.

In the derivation of the lower bound for the systole we will use the following lemma. This lemma will be used in the proof of Proposition 16 in Section 6 as well.

Lemma 11. *For geodesic segments between the sides of D_g the following properties hold:*

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

The proof can be found in Appendix B. The lower bound for the systole follows from the following result.

Lemma 12. *Every closed geodesic on \mathbb{M}_g has length at least ς_g .*

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

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

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

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

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

First, suppose that γ is of Type 1 or 2. Then, it follows directly from Part 1 and 2 of Lemma 11 that $\text{length}(\gamma) \geq \varsigma_g$.

Second, suppose that γ is of Type 3. It is not possible to form a closed geodesic with a segment of separation 2 or 3 and just one segment of separation 1, so we can assume that there are at least two 1-separated segments. First, assume that γ contains at least one 1-segment and at least one $(4g - 1)$ -segment. Then by the cyclic ordering of the segments it also contains at least one pair of consecutive segments consisting of one 1-segment and one $(4g - 1)$ -segment, which, by Part 3 of Lemma 11, have combined length at least $\frac{1}{2}\varsigma_g$. By Part 2 of Lemma 11 the length of the segment of separation 2 or 3 is at least $\frac{1}{2}\varsigma_g$ as well, so we conclude that $\text{length}(\gamma) \geq \varsigma_g$. Second, assume that the 1-separated segments in γ are either all 1-segments or all $(4g - 1)$ -segments. By possibly changing the orientation of γ we may assume without loss of generality that all 1-separated segments in γ are 1-segments. Furthermore, for convenience let us label the segments of γ by $\gamma_0, \gamma_1, \dots, \gamma_m$, where γ_0 is the unique segment of separation 2 or 3 and the segments $\gamma_1, \dots, \gamma_m$ are all 1-segments. By rotational symmetry of D_g we can assume without loss of generality that the endpoint of γ_0 is contained in \mathbf{s}_0 . Because the segments $\gamma_1, \dots, \gamma_m$ are all 1-segments, the successive Dirichlet regions that are intersected by the axis of the element of Γ_g corresponding to γ are given by $h_1(D_g), h_2(D_g), \dots, h_m(D_g)$, using the notation introduced in Section 2.4 for the clockwise sequence of Dirichlet regions around a vertex. In this case, the sequence of Dirichlet regions is around the vertex \mathbf{v}_0 , so

$$h_m = f_0 f_{2g+1} \cdots f_{m(2g+1)}$$

and we know by the argument in Section 4.1 that h_m is the element of Γ_g corresponding to γ . We will now derive the possible values of m . It can be shown that γ_0 is a k -segment if and only if the difference between the index of the first and last term in h_m is $2g + k$ (counting mod $4g$), using a similar argument as in Section 4.1 for 1-separated segments. Because γ_0 is a segment of separation 2 or 3, we find

$$-m(2g + 1) \equiv 2g + k \pmod{4g}$$

for $k \in \{2, 3, 4g - 2, 4g - 3\}$. Solving for m we obtain

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

Since by the group relation every sequence of $4g$ consecutive 1-segments corresponds to a homotopically trivial curve, it is sufficient to consider m such that $1 \leq m < 4g$. Therefore,

we can assume without loss of generality that $m \in \{3, 2g - 2, 2g + 2, 4g - 3\}$. Then, a direct computation (see Lemma 21 in Appendix B) shows that the length of γ is at least ς_g .

Finally, suppose that γ is of Type 4. By possibly changing the orientation of γ we may assume without loss of generality that the number of 1-segments contained in γ is larger than the number of $(4g - 1)$ -segments. If γ contains at least two $(4g - 1)$ -segments (and hence, at least two 1-segments), then by the cyclic ordering of the segments there are at least two pairs of consecutive segments each consisting of precisely one 1-segment and one $(4g - 1)$ -segment. As each of these pairs of segments has combined length at least $\frac{1}{2}\varsigma_g$ by Part 3 of Lemma 11, the length of γ is at least ς_g . Now assume that γ contains precisely one $(4g - 1)$ -segment. Since all but one of the segments of γ are 1-segments, we can use a similar reasoning as in the proof for sequences of segments of Type 3 to show that the number of 1-segments is $1 \pmod{4g}$. Again, since by the group relation sequences of $4g$ consecutive 1-segments correspond to a homotopically trivial curve, we can assume without loss of generality that γ contains precisely one 1-segment. Then the length of γ is ς_g by Part 4 of Lemma 11. Finally, if γ does not contain any $(4g - 1)$ -segment, i.e., if all segments in γ are 1-segments, then it is not difficult to see that γ is homotopically trivial, which is not possible since γ is a closed geodesic.

This finishes the proof. \square

5 Computation of dummy points

In this section we present two algorithms for constructing a dummy point set \mathcal{Q}_g satisfying the validity condition (10) 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 \mathbf{v}_0 . Some special properties of Weierstrass points are known in Riemann surface theory [33], 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 [48]. 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 13. *For any finite set of points \mathcal{Q}_g on \mathbb{M}_g containing \mathcal{W}_g , each face in the Delaunay triangulation $\text{DT}_{\mathbb{D}}(\pi_g^{-1}(\mathcal{Q}_g))$ 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*, a description of which can be found in the arXiv version [28, Appendix D].² Its approach is fundamentally different from the refinement and symmetric algorithms: the dummy point set and the corresponding Delaunay triangulation are exactly described. As in the symmetric algorithm, the resulting dummy point set preserves the symmetries of D_g and is of order $\Theta(g \log g)$.

5.1 Refinement algorithm

Following the refinement strategy introduced above and using Proposition 13, 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**

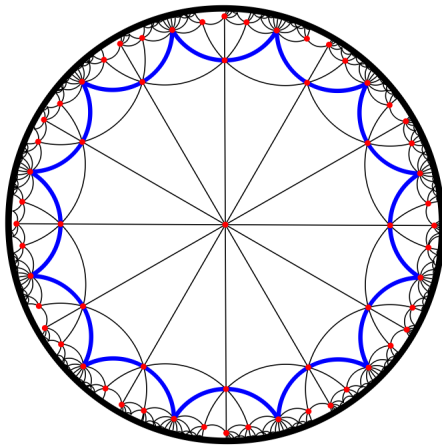
Algorithm 1: Refinement algorithm

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

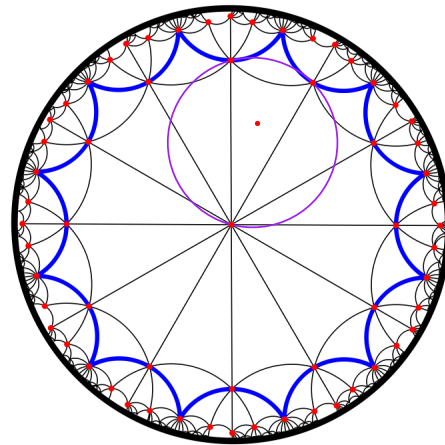
Theorem 14. *The refinement algorithm terminates and the resulting dummy point set \mathcal{Q}_g satisfies the validity condition (10). The cardinality $|\mathcal{Q}_g|$ is bounded as follows*

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

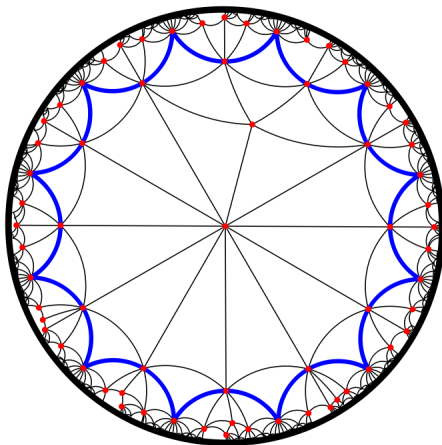
²Note that Appendix D is the only difference between the arXiv version and this version.



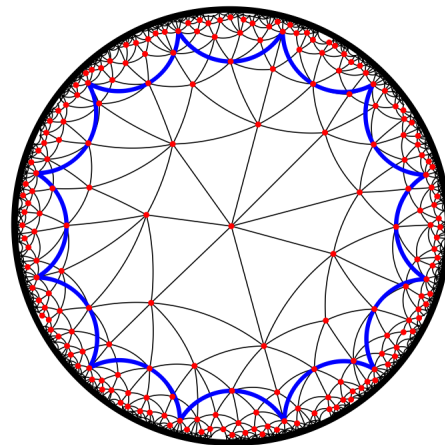
After initialization



First insertion



After first insertion



After last insertion

Figure 9: Several steps in the refinement algorithm (genus 3)

Proof. We first prove that the hyperbolic distance between two distinct points of \mathcal{Q}_g is greater than $\frac{1}{4} \text{sys}(\mathbb{M}_g)$. Let us first show that the distance between any pair of Weierstrass points is larger than $\frac{1}{4} \text{sys}(\mathbb{M}_g)$. 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 10.

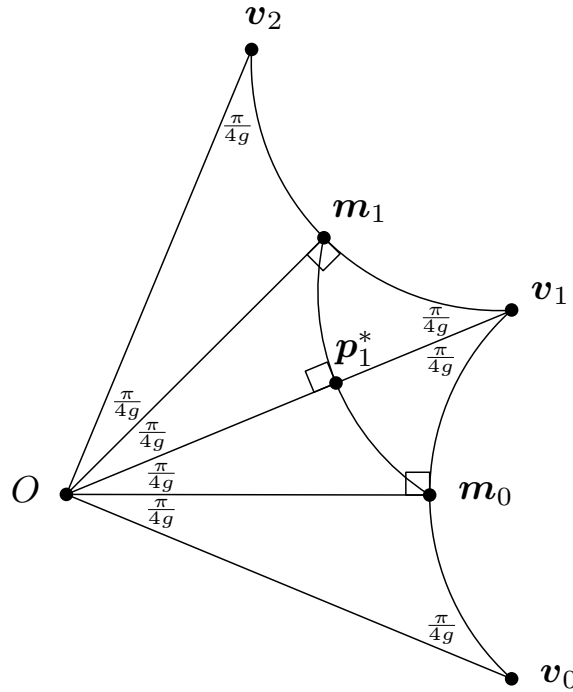


Figure 10: Computing the distances between Weierstrass points of \mathbb{M}_g .

Let m_k be the midpoint of side s_k and let p_k^* be the midpoint of m_{k-1} and m_k . First, in the right-angled triangle $[O, m_0, v_0]$ we know that [8, Theorem 7.11.3]

$$\begin{aligned} \cosh(d(O, m_0)) &= \cosh(d(m_0, v_0)) = \frac{\cos(\frac{\pi}{4g})}{\sin(\frac{\pi}{4g})} = \cot(\frac{\pi}{4g}), \\ \cosh(d(O, v_0)) &= \cot^2(\frac{\pi}{4g}). \end{aligned}$$

Second, because $d(O, m_0) = d(m_0, v_1)$, $[O, m_0, p_1^*]$ and $[v_1, m_0, p_1^*]$ are congruent triangles. In particular, $\angle(Om_0p_1^*) = \angle(v_1m_0p_1^*) = \frac{\pi}{4}$. It follows that

$$\cosh(\frac{1}{2}d(m_0, m_1)) = \cosh(d(m_0, p_1^*)) = \frac{\cos(\frac{\pi}{4g})}{\sin(\frac{\pi}{4})} = \sqrt{2} \cos(\frac{\pi}{4g}). \tag{13}$$

The above formulas yield expressions for $d(O, m_0)$, $d(O, v_0)$ and $d(m_0, m_1)$ and comparing these to the expression for $\text{sys}(\mathbb{M}_g)$ (see Theorem 2) yields the result.

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

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

Proposition 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

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 11 illustrates the computation of $\text{DT}_{\mathbb{D}}(\mathcal{Q}_{\mathcal{N}_3})$.

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

Algorithm 2: Symmetric algorithm

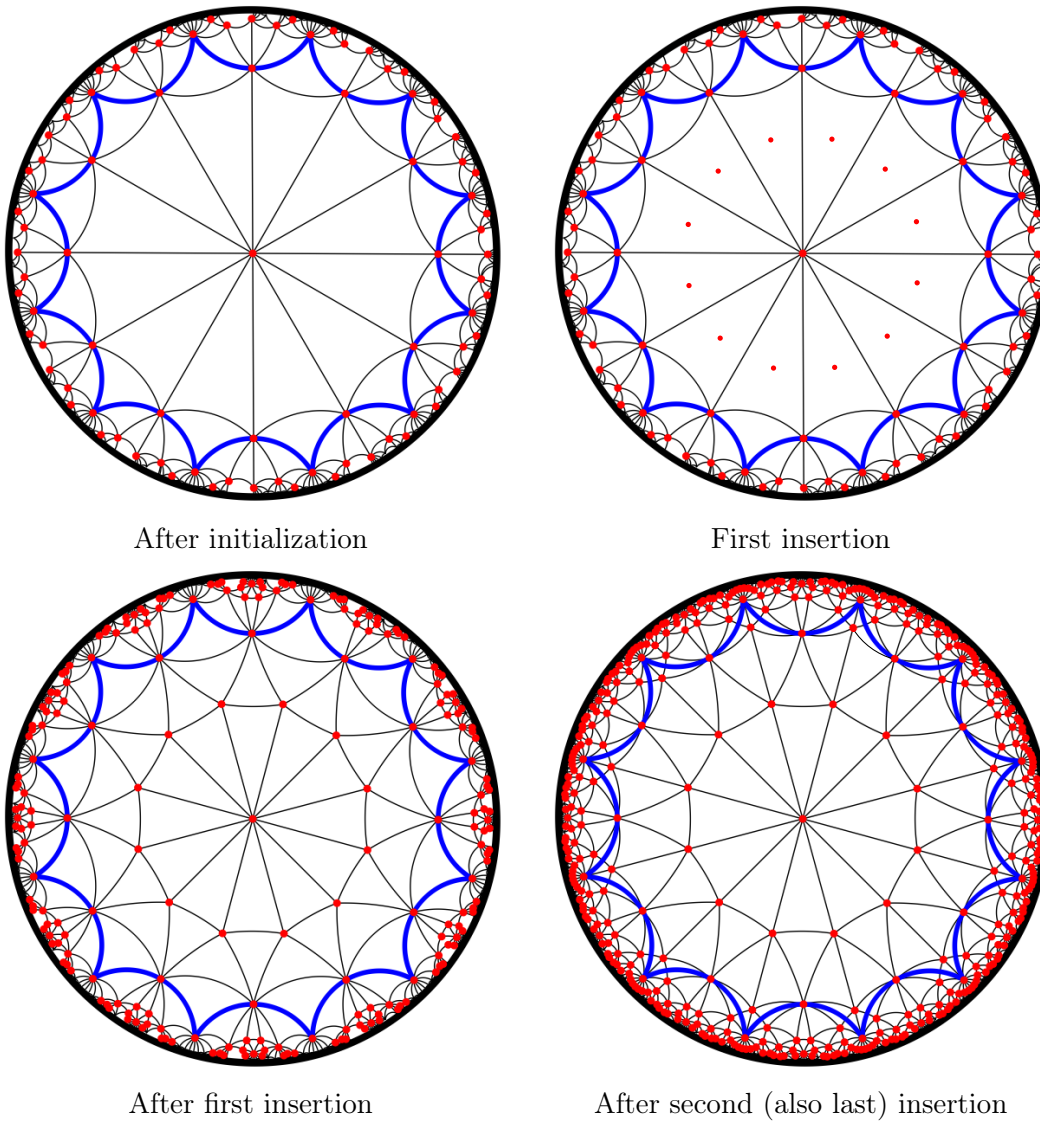


Figure 11: Several steps in the symmetric algorithm (genus 3)

By using the symmetry of the regular $4g$ -gon we obtain a more symmetric dummy point set, which may be interesting for some applications [19]. However, asymptotically the resulting point set is larger than the point set obtained from the refinement algorithm.

Theorem 15. *The symmetric algorithm terminates and the resulting dummy point set satisfies the validity condition (10). Its cardinality is of order $\Theta(g \log g)$.*

Proof. The first two statements follow directly from the proof of Theorem 14, so we only have to prove the claim on the cardinality of \mathcal{Q}_g .

First, we prove that $|\mathcal{Q}_g|$ is of order $O(g \log g)$. Again, the distance between the Weierstrass points is more than $\frac{1}{4} \text{sys}(\mathbb{M}_g)$. We claim that the distance between points that are added in different iterations of the while loop is at least $\frac{1}{4} \text{sys}(\mathbb{M}_g)$. Namely, by the same reasoning as in the proof of Theorem 14, the distance between the circumcenter of an empty disk of radius at least $\frac{1}{4} \text{sys}(\mathbb{M}_g)$ and any other point in \mathcal{Q}_g is at least $\frac{1}{4} \text{sys}(\mathbb{M}_g)$. Because \mathcal{Q}_g is invariant under symmetries of D_g , it follows that the distance between an image of the circumcenter under a rotation around the origin and any other point in \mathcal{Q}_g is at least $\frac{1}{4} \text{sys}(\mathbb{M}_g)$ as well.

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

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

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

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

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

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

Define

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

and denote its complement by \mathcal{I}^c . We give upper bounds for $|\mathcal{I}|$ and $|\mathcal{I}^c|$. We will show that $j \in \mathcal{I}$, i.e., that the inequality $a_j < 2ga$ holds, if \mathbf{p}_k^j is sufficiently close to O or \mathbf{v}_0 for some $k \in \{0, \dots, 4g - 1\}$. To quantify this “sufficiently close”, let us consider the area of U_j (see Figure 12a). The amount of overlap between $D(\mathbf{p}_k^j)$ and $D(\mathbf{p}_{k+1}^j)$ can be written as a strictly decreasing function of $d(\mathbf{p}_k^j, \mathbf{p}_{k+1}^j)$, which can be written as a strictly increasing function of $d(O, \mathbf{p}_k^j)$. Therefore, there exists a constant $d_g > 0$ such that $\text{area}(U_j) < 2ga$ if and only if $d(O, \mathbf{p}_k^j) < d_g$ for all $k = 0, \dots, 4g - 1$.

We claim that $j \in \mathcal{I}$ if and only if there exists $k \in \{0, \dots, 4g - 1\}$ such that either $d(O, \mathbf{p}_k^j) < d_g$ or $d(\mathbf{v}_0, \mathbf{p}_k^j) < d_g$. First, assume that such a k exists. If $d(O, \mathbf{p}_k^j) < d_g$ (Figure 12a), 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 12b). 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 \mathbf{s}_{2g} to \mathbf{s}_0 . Then

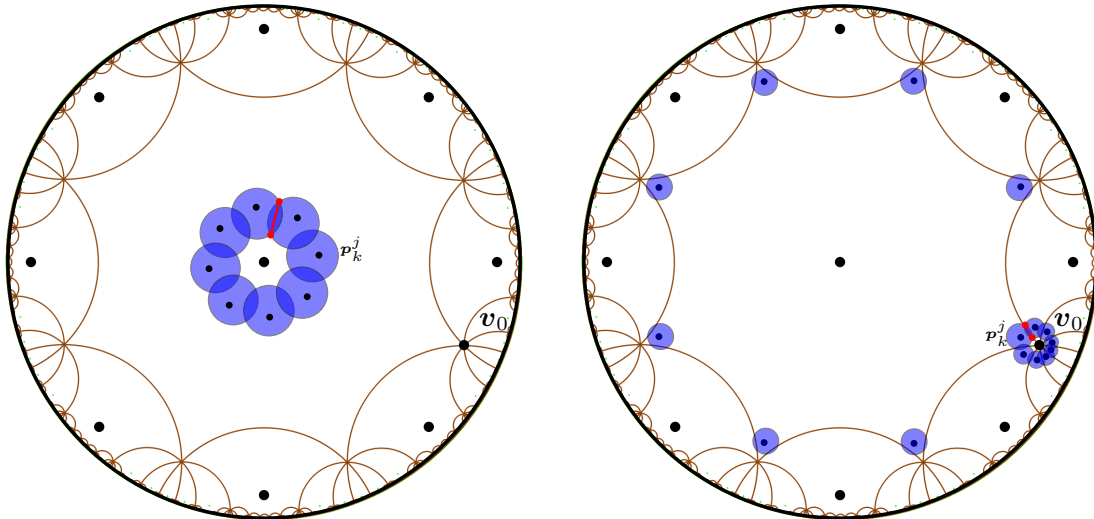
$$\begin{aligned} d(f_0(\mathbf{p}_{k+2g+1}^j), \mathbf{v}_0) &= d(f_0^{-1}(f_0(\mathbf{p}_{k+2g+1}^j)), f_0^{-1}(\mathbf{v}_0)), \\ &= d(\mathbf{p}_{k+2g+1}^j, \mathbf{v}_{2g+1}), \\ &= d(\mathbf{p}_k^j, \mathbf{v}_0). \end{aligned}$$

Therefore, the circle \mathcal{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 co-prime, it follows that \mathcal{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 \mathcal{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 $a_j < 2ga$.

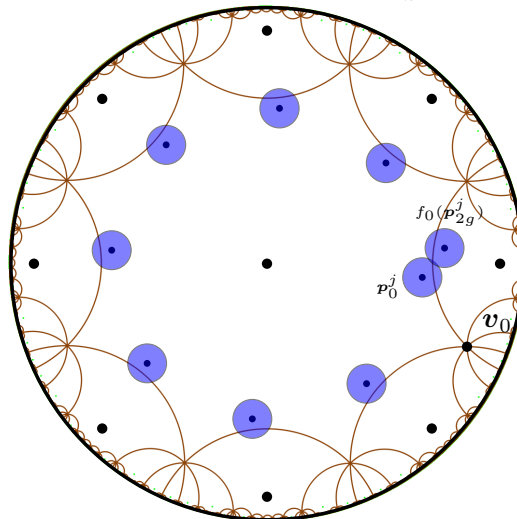
Second, assume that $d(O, \mathbf{p}_k^j) \geq d_g$ and $d(\mathbf{v}_0, \mathbf{p}_k^j) \geq d_g$ for all $k \in \{0, \dots, 4g-1\}$. If $d(\mathbf{p}_0^j, \partial D_g) \geq \frac{1}{8} \text{sys}(\mathbb{M}_g)$, then U_j is completely contained in D_g . Because $d(O, \mathbf{p}_k^j) \geq d_g$, it 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 12c). 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 \mathcal{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 $d(\mathbf{v}_0, \mathbf{p}_k^j) < d_g$. We have also shown that if $d(O, \mathbf{p}_k^j) < d_g$, then U_j is a topological annulus around the origin. If $d(\mathbf{v}_0, \mathbf{p}_k^j) < d_g$, then $\pi_g^{-1}(U_j)$ contains a topological annulus around \mathbf{v}_0 . In either case, the boundary of such an annulus consists of two connected components. Let the minimum width of an annulus be given by the distance between these connected components. Suppose, for a contradiction, that the minimum width of an annulus corresponding to $j \in \mathcal{I}$ can be arbitrarily close to 0. Then the disks in U_j have arbitrarily small overlap, so a_j is arbitrarily close to $4ga$. However, this is not possible, since $a_j < 2ga$ for all $j \in \mathcal{I}$. Therefore, there exists $\varepsilon > 0$ (independent of the output of the algorithm) such that the minimum width of an annulus corresponding to $j \in \mathcal{I}$ is at least ε .

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



(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 12: 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.

ε , there are at most $\text{length}([O, \mathbf{v}_0])/\varepsilon$ annuli, where

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

Therefore,

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

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

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

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

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

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

Secondly, we show that $|\mathcal{Q}_g|$ is of order $\Omega(g \log g)$. As before, the points added to \mathcal{Q}_g in iteration j of the while loop are denoted by $\pi_g(\mathbf{p}_k^j)$ where $k = 0, \dots, 4g-1$. Fix an arbitrary vertex \mathbf{v} of D_g . Let $P = \langle O, \mathbf{p}_{k_1}^{j_1}, \mathbf{p}_{k_2}^{j_2}, \dots, \mathbf{p}_{k_n}^{j_n}, \mathbf{v} \rangle$ be a shortest path from the origin to \mathbf{v} in the Delaunay graph of $\pi_g^{-1}(\mathcal{Q}_g)$. We claim that all indices j_h are distinct, i.e. P contains at most one element of each of the sets $\{\mathbf{p}_k^j \mid k = 0, \dots, 4g-1\}$ (see Figure 13).

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 (Theorem 2), 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

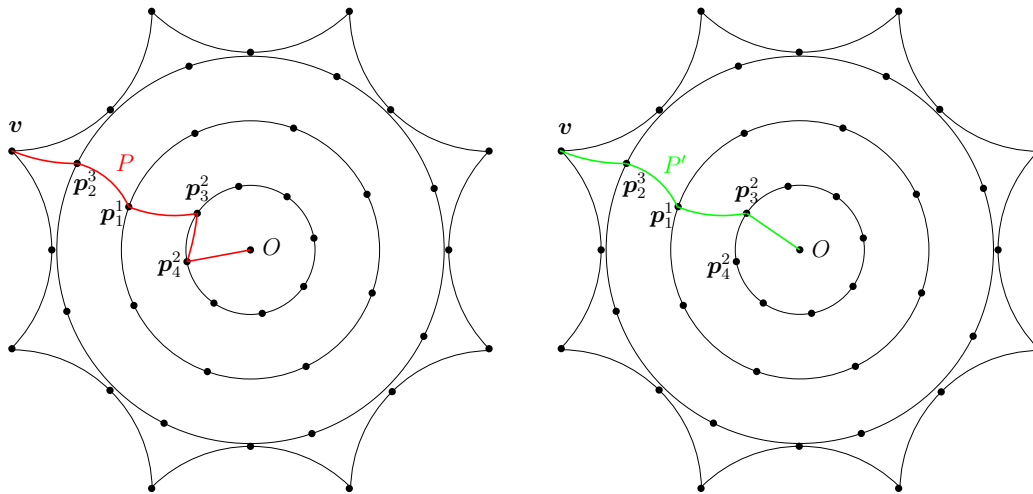


Figure 13: The left figure shows a path P from the origin to v that visits two vertices from the same iteration, namely p_3^2 and p_4^2 . The right figure shows a shorter path from the origin to v . In this case, $j_1 = j_2 = 2$ and $k_1 = 4$ and $k_2 = 3$. The subdivision of P into three parts is given by $P_1 = \langle O, p_4^2 \rangle$, $P_2 = \langle p_4^2, p_3^2 \rangle$ and $P_3 = \langle p_3^2, p_1^1, p_2^3, v \rangle$. The path P' is defined as $P'_1 \cup P_3$, where P'_1 is obtained by rotating P_1 around the origin by angle $-\frac{\pi}{2g}$, i.e., $P'_1 = \langle O, p_3^2 \rangle$.

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 [53] 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 14 we have shown the dummy point set computed by the symmetric algorithm. However, a smaller set, consisting of 14 dummy points, was proposed earlier [10]: in addition to the six Weierstrass points, it contains the eight midpoints of the segments $[O, v_k]$, $k = 0, 1, \dots, 7$ (see Figure 14).

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 11). 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.

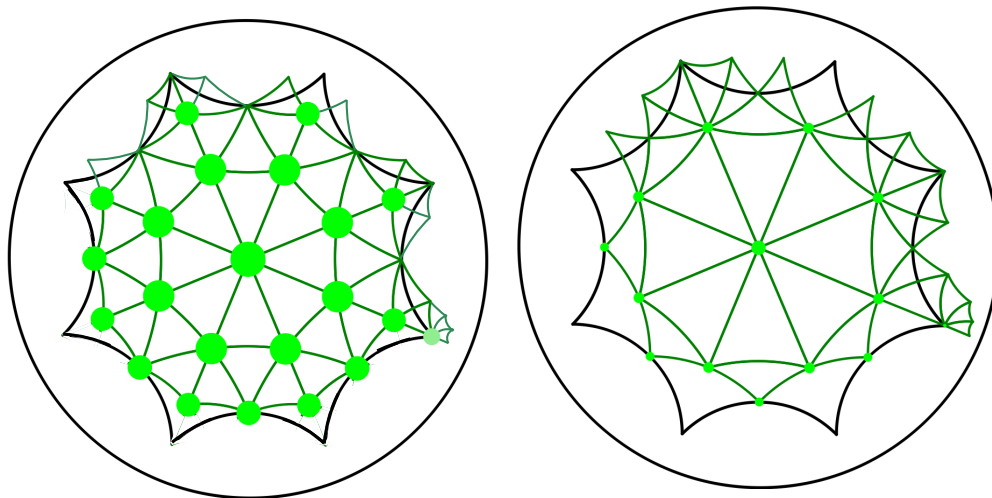


Figure 14: Set of 22 dummy points for the Bolza surface computed by the symmetric algorithm (left) and set of 14 dummy points constructed by hand [10] (right).

6 Data structure, predicates, and implementation

In this section, we detail two major aspects of Bowyer’s algorithm for generalized Bolza surfaces. On the one hand, the combinatorial aspect, i.e., the data structure and the way it supports the algorithm, is studied in Section 6.2. On the other hand, the algebraic degree of the predicates based on which the decisions are made by the algorithm is analyzed in Section 6.3. Finally, we report on our implementation and experimental results in Section 6.4.

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

6.1 Canonical representatives

We have defined in Section 2.4 the canonical representative $\mathbf{p}^c \in \tilde{D}_g$ of a point p on the surface \mathbb{M}_g or a point \mathbf{p} in \mathbb{D} . 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 $\mathcal{N}_g \setminus \{\mathbb{1}\}$ (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 $\mathcal{N}_g \setminus \{\mathbb{1}\}$. An illustration for genus 2 is shown in Figure 15.

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:

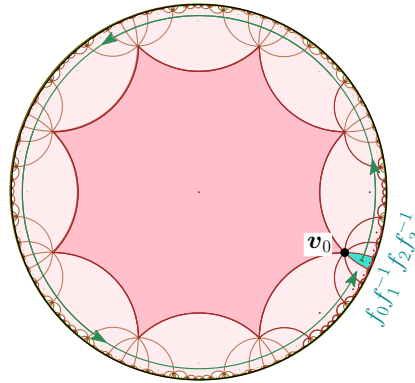


Figure 15: The ordering of \mathcal{N}_2 starts with $f_0 f_1^{-1} f_2 f_3^{-1} = f_0 f_5 f_2 f_7$.

Proposition 16 (Inclusion property). *If at least one vertex of an admissible triangle is contained in \tilde{D}_g , then the whole triangle is contained in $D_{\mathcal{N}_g}$.*

Proof. It is sufficient to show that the distance between the boundary ∂D_g of D_g and the boundary $\partial D_{\mathcal{N}_g}$ of $D_{\mathcal{N}_g}$ is at least $\frac{1}{2} \text{sys}(\mathbb{M}_g)$. Consider points $\mathbf{p} \in \partial D_g$ and $\mathbf{q} \in \partial D_{\mathcal{N}_g}$. We will show that $d(\mathbf{p}, \mathbf{q}) \geq \frac{1}{2} \text{sys}(\mathbb{M}_g)$. By symmetry of D_g , we can assume without loss of generality that $\mathbf{p} \in s_0$. In Section 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 .

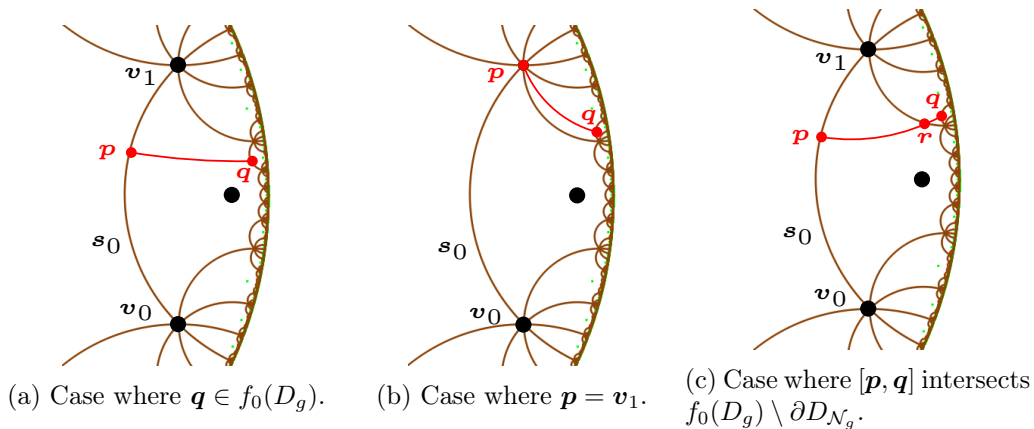


Figure 16: Cases in the proof of Proposition 16.

Recall that f_0 is the side-pairing transformation that maps s_{2g} to s_0 . First, assume that $\mathbf{q} \in f_0(D_g)$. Because $\mathbf{q} \in \partial D_{\mathcal{N}_g}$, $[\mathbf{p}, \mathbf{q}]$ is a segment of separation at least 2 (see Figure 16a). By Part 2 of Lemma 11, $d(\mathbf{p}, \mathbf{q}) \geq \frac{1}{2} \text{sys}(\mathbb{M}_g)$.

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 16b), so $d(\mathbf{p}, \mathbf{q}) \geq \frac{1}{2} \text{sys}(\mathbb{M}_g)$ by Part 2 of Lemma 11. 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 16c). In particular, $[\mathbf{p}, \mathbf{r}]$ is a 1-separated segment. If $[\mathbf{r}, \mathbf{q}]$ is a segment of separation at least 2, then $d(\mathbf{r}, \mathbf{q}) \geq \frac{1}{2} \text{sys}(\mathbb{M}_g)$ by Part 2 of Lemma 11, so $d(\mathbf{p}, \mathbf{q}) \geq \frac{1}{2} \text{sys}(\mathbb{M}_g)$ as well. If $[\mathbf{r}, \mathbf{q}]$ is a 1-separated segment, then $d(\mathbf{p}, \mathbf{r}) + d(\mathbf{r}, \mathbf{q}) \geq \frac{1}{2} \text{sys}(\mathbb{M}_g)$ by Part 3 of Lemma 11.

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

Let now $\mathcal{S} \subset \mathbb{M}_g$ be a set of points satisfying the validity condition (10). 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

$$\Sigma = \left\{ t \in \pi_g^{-1}(t) \mid t \text{ has at least one vertex in } \tilde{D}_g \right\} \tag{14}$$

contains at most three faces. See Figure 17. 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 16, each face $t \in \Sigma$ is contained in $D_{\mathcal{N}_g}$. So, for each vertex \mathbf{u} of t , there is a unique translation $T(\mathbf{u}, t)$ in \mathcal{N}_g such that \mathbf{u} lies in $T(\mathbf{u}, t)(\tilde{D}_g)$. This translation is such that

$$T(\mathbf{u}, t)(\mathbf{u}^c) = \mathbf{u},$$

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

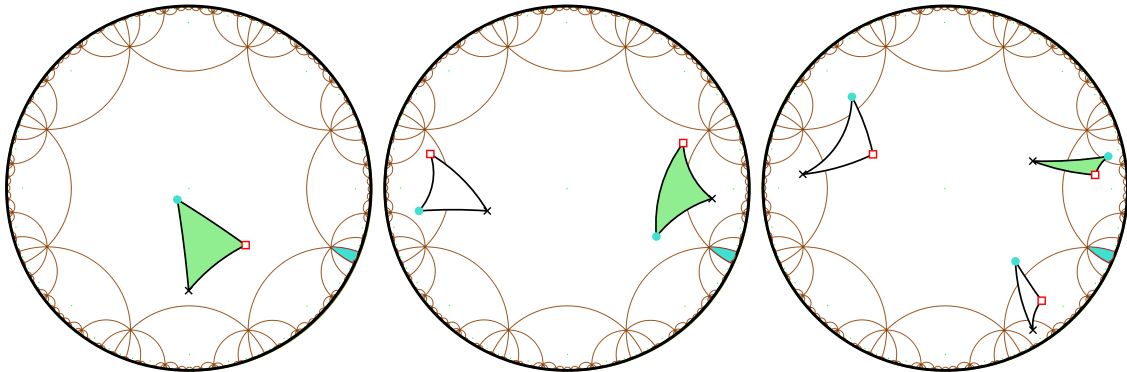


Figure 17: Examples (for $g = 2$) of faces of $\text{DT}_{\mathbb{D}}(\pi_g^{-1}(\mathcal{S}))$ with one (left), two (middle) and three (right) vertices in \tilde{D}_g that project to the same face on \mathbb{M}_g . Their respective vertices drawn as a dot project to the same vertex on \mathbb{M}_g (same for cross and square). The canonical representative is the shaded face.

To summarize, we have shown that:

Proposition 17. *Let $\mathcal{S} \subset \mathbb{M}_g$ be a set of points satisfying the validity condition (10). For any face t in $\text{DT}_{\mathbb{M}_g}(\mathcal{S})$, there exists a unique canonical representative $t^c \subset D_{\mathcal{N}_g}$ in $\text{DT}_{\mathbb{D}}(\pi_g^{-1}(\mathcal{S}))$.*

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

6.2 Data structure

Proposition 17 allows us to propose a data structure to represent Delaunay triangulations of generalized Bolza surfaces.

A triangulation of a point set $\mathcal{S} \subset \mathbb{M}_g$ is represented via its vertices and triangular faces. Each vertex u stores its canonical representative \mathbf{u}^c in \tilde{D}_g and gives access to one of its incident triangles. Each triangle t is actually storing information to construct its canonical representative \mathbf{t}^c : it gives access to its three incident vertices u_0, u_1 , and u_2 and its three adjacent faces; it also stores the three translations $T(u_j, t) := T(\mathbf{u}_j, \mathbf{t}^c)$, $j = 0, 1, 2$ in \mathcal{N}_g as defined in Section 6.1, so that applying each translation to the corresponding canonical point yields the canonical representative \mathbf{t}^c of t , i.e.,

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

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

Finding conflicts. The notion of conflict defined in section 3.2 can now be made more explicit: a triangle $t \in \text{DT}_{\mathbb{M}_g}(\mathcal{S})$ is in conflict with a point $p \in \mathbb{M}_g$ if the circumscribing disk of one of the (at most three) triangles in Σ is in conflict with \mathbf{p}^c , where Σ is the set defined by relation (14).

By the correspondence between Euclidean circles and hyperbolic circles in the Poincaré disk model, the triangle in the Delaunay triangulation in \mathbb{D} whose associated Euclidean triangle contains the point \mathbf{p}^c is in conflict with this point; these Euclidean and hyperbolic triangles will both be denoted as \mathbf{t}_p , which should not introduce any confusion. To find this triangle, we adapt the so-called *visibility walk* [23]: the walk starts from an arbitrary face, then, for each visited face, it visits one of its neighbors, until the face whose associated Euclidean triangle contains \mathbf{p}^c is found. This walk will be detailed below.

We first need some notation. Let t, t' be two adjacent faces in $\text{DT}_{\mathbb{M}_g}(\mathcal{S})$. We define the *neighbor translation* $T^{\text{nbr}}(\mathbf{t}'^c, \mathbf{t}^c)$ from \mathbf{t}'^c to \mathbf{t}^c as the translation of Γ_g such that $T^{\text{nbr}}(\mathbf{t}'^c, \mathbf{t}^c)(\mathbf{t}'^c)$ is adjacent to \mathbf{t}^c in $\text{DT}_{\mathbb{D}}(\pi^{-1}(\mathcal{S}))$. See Figure 18. Let u be a vertex common to t and t' , and let \mathbf{u}_j and $\mathbf{u}_{j'}$ be the vertices of \mathbf{t}^c and \mathbf{t}'^c that project on u by π_g . We can compute the neighbor translation from \mathbf{t}'^c to \mathbf{t}^c as $T^{\text{nbr}}(\mathbf{t}'^c, \mathbf{t}^c) = T(u_j, t)(T(u_{j'}, t'))^{-1}$. It can be easily seen that $T^{\text{nbr}}(\mathbf{t}'^c, \mathbf{t}^c) = T(u_j, t)(T(u_{j'}, t'))^{-1} = (T(u_{j'}, t')(T(u_j, t))^{-1})^{-1} = (T^{\text{nbr}}(\mathbf{t}^c, \mathbf{t}'^c))^{-1}$.

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

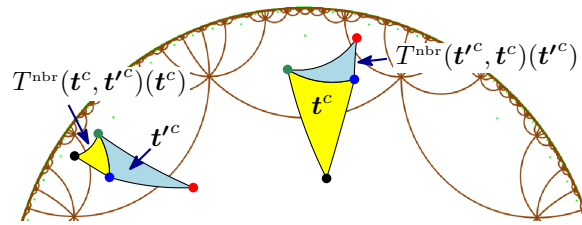


Figure 18: Translating t'^c by $T^{\text{nbr}}(t'^c, t^c)$ gives a face adjacent to t^c .

two of the vertices of t . We check whether the Euclidean line supporting e separates p^c from the vertex of t opposite to e . If this is the case, the next visited face is the neighbor t' of t through e ; the location translation is updated: $T_p^{\text{loc}} := T_p^{\text{loc}} T^{\text{nbr}}(t'^c, t^c)$. The walk stops when it finds the Euclidean triangle t_p containing p^c . Then the canonical face t_p^c in conflict with p^c is $(T_p^{\text{loc}})^{-1}(t_p)$. See Figure 19 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 $T^{\text{nbr}}(t'^c, t_D)$. 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.

Let us now present the computation of the set C_p of faces of $\text{DT}_{\mathbb{D}}(\pi^{-1}(\mathcal{S}))$ in conflict with p^c . Starting from t_p , for each face of $\text{DT}_{\mathbb{D}}(\pi^{-1}(\mathcal{S}))$ in conflict with p^c we recursively examine each neighbor (obtained with a neighbor translation) that has not yet been visited, checking it for conflict with p^c . When a face is found to be in conflict, we temporarily store directly in each of its vertices the translation that moves its corresponding canonical point to it (we cannot store such translations in the face itself, since this face will be deleted by the insertion). Since the union of the faces of C_p is a topological disk by definition, the resulting translation for a given vertex is the same for all faces of C_p incident to it, so this translation is well defined for each vertex. The temporary translations will be used during the insertion stage described below. We store the set C_p^c of canonical faces corresponding to faces of C_p . Note that C_p^c is not necessarily a connected region in \mathbb{D} , as illustrated in Figure 19(Right).

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

For each edge e on the boundary of C_p , we create a new face t_e on \mathbb{M}_g corresponding to the triangle t_e in \mathbb{D} formed by the new vertex and the edge e . The neighbor of t_e through e is the neighbor through e of the face in C_p^c that is incident to e . Two new faces consecutive along the boundary of C_p are adjacent. We now delete all faces in C_p . The triangle t_e is not necessarily the canonical representative of t_e ; we must now compute the three translations to be stored in t_e to get t_e^c . To this aim, we first retrieve the translations temporarily stored in its vertices $u_j, j = 0, 1, 2$ and we respectively initialize the translations $T(u_j, t_e)$ in t_e to them. If all translations are equal to $\mathbb{1}$, then the face is already canonical and there is nothing more to do. Otherwise, the translations stored in the face are updated following Section 6.1: $T(u_j, t_e) := (T(u_k, t_e))^{-1} T(u_j, t_e), j = 0, 1, 2$, where k is the index in $\{0, 1, 2\}$ for which $u_k = u_{t_e^c}^*$.

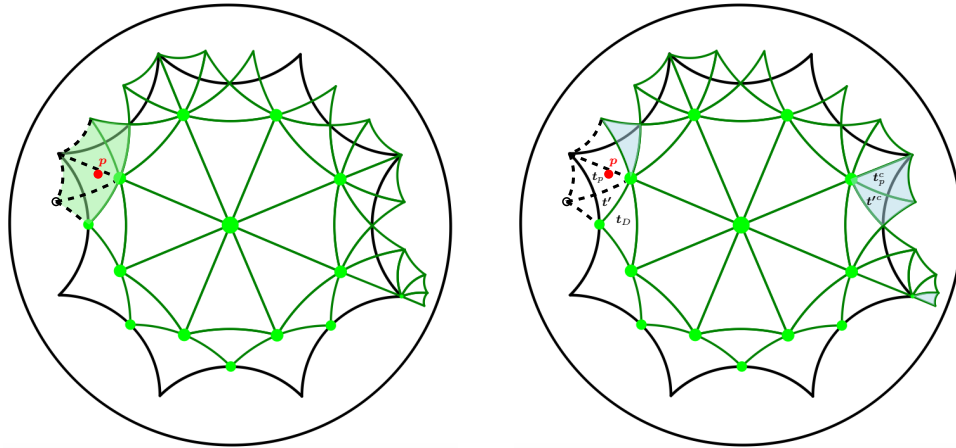


Figure 19: Left: The shaded faces are the (not necessarily canonical) faces in \mathcal{C}_p , i.e., faces in conflict with the red point p^c . Their union is a topological disk. Right: The region \mathcal{C}_p^c of the (shaded) corresponding canonical faces is not connected in \mathbb{D} .

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

Removing a vertex. We mimick the practical trick already used in CGAL in the case of the 3D flat torus [16]. Let us just give a short sketch. The idea is to compute an external (Euclidean) Delaunay triangulation of the set of points given by the vertices that are adjacent to the vertex to be removed. Then the relevant part of this external triangulation is “sewed” into the hole created by the removal of v , and translations to be stored can be retrieved. Technicalities can be found in Jordanov’s PhD thesis [36, Sections III.3.2 and III.3.3].

6.3 Degree of predicates

Following the celebrated exact geometric computation paradigm [52], the correctness of the combinatorial structure of the Delaunay triangulation relies on the exact evaluation of predicates. The main two predicates are

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

Input points, which lie in \tilde{D}_g , are tested against canonical triangles of the triangulation, whose vertices are images of input points by translations in \mathcal{N}_g . If points are assumed to have rational coordinates, then evaluating the predicates boils down to determining the sign of polynomial expressions whose coefficients are lying in some extension field of the rationals. Proposition 18 gives an upper bound on the degree of these polynomial expressions. For the

special case of the Bolza surface ($g = 2$), it improves the previously known upper bound from 72 [37, Proposition 1], which was proved using symbolic computations in MAPLE, to 48.

Proposition 18. *For the generalized Bolza surface of genus $g \geq 2$, the predicates can be evaluated by determining the sign of rational polynomial expressions of total degree at most $12\varphi(4g) \leq 24g$ in the coordinates of the input points, where φ is the Euler totient function.*

Recall that the Euler totient function $\varphi(n)$ counts the number of integers up to a given integer n that are relatively prime to n [42].

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

$$\begin{aligned} \text{INCIRCLE}(\mathbf{p}_1, \mathbf{p}_2, \mathbf{p}_3, \mathbf{p}_4) &= \begin{vmatrix} x_1 & y_1 & x_1^2 + y_1^2 & 1 \\ x_2 & y_2 & x_2^2 + y_2^2 & 1 \\ x_3 & y_3 & x_3^2 + y_3^2 & 1 \\ x_4 & y_4 & x_4^2 + y_4^2 & 1 \end{vmatrix} \\ &= x_3y_4(x_1^2 + y_1^2) - x_3y_4(x_2^2 + y_2^2) + x_2y_4(x_3^2 + y_3^2) - x_2y_3(x_4^2 + y_4^2) \dots, \end{aligned}$$

for points $\mathbf{p}_j = x_j + y_ji, j = 1, \dots, 4$, in \mathbb{D} . In the second equality, we have just written 4 of the 24 terms to illustrate what the terms look like. This will be used later in the proof to determine their maximum degree. Since at least one of the four points is contained in \tilde{D}_g , we will assume without loss of generality that $x_1, y_1 \in \mathbb{Q}$. The other points are images of points with rational coordinates under some elements of Γ_g . We know that Γ_g is generated by f_k for $k = 0, \dots, 4g - 1$. The translation f_k can be represented by the matrix A_k (see Equation (5) in Section 2.4). The entries of A_k are contained in the extension field

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

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

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

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

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

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

$$\begin{aligned} & \text{INCIRCLE}(\mathbf{p}_1, \mathbf{p}_2, \mathbf{p}_3, \mathbf{p}_4) \\ &= \frac{P_3 Q_4 (x_1^2 + y_1^2)}{R_3 R_4} - \frac{P_3 Q_4 (P_2^2 + Q_2^2)}{R_2^2 R_3 R_4} + \frac{P_2 Q_4 (P_3^2 + Q_3^2)}{R_2 R_3^2 R_4} - \frac{P_2 Q_3 (P_4^2 + Q_4^2)}{R_2 R_3 R_4^2} + \dots, \\ &= \frac{R_2^2 P_3 R_3 Q_4 R_4 (x_1^2 + y_1^2) - P_3 R_3 Q_4 R_4 (P_2^2 + Q_2^2) + \dots}{R_2^2 R_3^2 R_4^2} \end{aligned}$$

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

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

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

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

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

6.4 Implementation and experimental results

The algorithm presented in Section 3 was implemented in C++, with the data structure described in Section 6.2. The preprocessing step consists in computing dummy points that serve for the initialization of the data structure, following the two options presented in Section 5. The implementation also uses the value of the systole given by Theorem 2.

Let us continue the discussion on predicates. In practice, the implementation relies on the `CORE::Expr` number type [53], which provides us 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 to the barriers raised by their very high algebraic degree, so, only a non-robust implementation of the algorithm can be obtained.

The rest of this section is devoted to the implementation for the Bolza surface, for which a fully robust implementation has been integrated in `CGAL` [38]. All details can be found in Iordanov's PhD thesis [36]. We only mention a few key points here.

To avoid increasing further the algebraic degree of predicates, the coordinates of dummy points are rounded to rationals (see Table 1). We have checked that the validity

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

condition (10) 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 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 [20, 34], 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, 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.

7 Conclusion and open problems

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

One of the main ingredients of our extension of Bowyer’s algorithm is the validity condition (10), and to be able to say whether it holds or not we need to know the value of the systole of the hyperbolic surface. For general hyperbolic surfaces an explicit value, or a ‘reasonable’ lower bound of the systole, is not known, and there are no efficient algorithms to compute or approximate it. The effective procedure presented in [2] is based on the

construction of a pants decomposition of a hyperbolic surface, and computes the systole from the Fenchel-Nielsen coordinates associated with this decomposition. However, the complexity of this algorithm does not seem to be known, and it is not clear how to turn this method into an efficient and robust algorithm.

If the systole is known, then it seems that we can use the refinement algorithm presented in Section 5.1 to compute a dummy point set satisfying the validity condition. However, in the case of generalized Bolza surfaces it is sufficient to consider only the translates of vertices in $D_{\mathcal{N}_g}$ by Proposition 13, whereas it is not clear how many translates are needed for an arbitrary hyperbolic surface.

A more modest attempt towards generalization could focus on hyperbolic surfaces represented by a ‘nice’ fundamental polygon. Hyperelliptic surfaces have a point-symmetric fundamental polygon (See [50]), so these surfaces are obvious candidates for future work.

References

- [1] William Abikoff. The Uniformization Theorem. *The American Mathematical Monthly*, 88(8):574–592, 1981. doi:10.2307/2320507.
- [2] H. Akrouf. Un processus effectif de détermination des systoles pour les surfaces hyperboliques. *Geometriae Dedicata*, 121(1):1–8, 2006. doi:10.1007/s10711-006-9076-x.
- [3] M. A. Armstrong. *Basic Topology*. Undergraduate Texts in Mathematics. Springer New York, 2013.
- [4] R. Aurich, E. B. Bogomolny, and F. Steiner. Periodic orbits on the regular hyperbolic octagon. *Physica D: Nonlinear Phenomena*, 48(1):91–101, 1991. doi:10.1016/0167-2789(91)90053-C.
- [5] R. Aurich and F. Steiner. On the periodic orbits of a strongly chaotic system. *Physica D: Nonlinear Phenomena*, 32(3):451–460, 1988. doi:10.1016/0167-2789(88)90068-1.
- [6] Sheng Bai, Yue Gao, and Shicheng Wang. Systoles of hyperbolic surfaces with big cyclic symmetry. *Science China Mathematics*, 64(2):421–442, 2021. doi:10.1007/s11425-019-1655-8.
- [7] N.L. Balazs and A. Voros. Chaos on the pseudosphere. *Physics Reports*, 143(3):109 – 240, 1986. doi:10.1016/0370-1573(86)90159-6.
- [8] Alan F. Beardon. *The geometry of discrete groups*. Springer-Verlag New York, 1 edition, 1983. doi:10.1007/978-1-4612-1146-4.
- [9] 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.

- [10] 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.
- [11] Oskar Bolza. On Binary Sextics with Linear Transformations into Themselves. *American Journal of Mathematics*, 10:47–70, 1887. doi:10.2307/2369402.
- [12] A. Bowyer. Computing Dirichlet tessellations. *The Computer Journal*, 24(2):162–166, 1981. doi:10.1093/comjnl/24.2.162.
- [13] P. Buser. *Geometry and Spectra of Compact Riemann Surfaces*. Progress in Mathematics Series. Birkhäuser, 1992.
- [14] P. Buser and P. Sarnak. On the period matrix of a Riemann surface of large genus (with an Appendix by J.H. Conway and N.J.A. Sloane). *Inventiones mathematicae*, 117:27–56, 1994. doi:10.1007/BF01232233.
- [15] Manuel Caroli, Pedro M. M. de Castro, Sébastien Lorient, Olivier Rouiller, Monique Teillaud, and Camille Wormser. Robust and efficient Delaunay triangulations of points on or close to a sphere. In *9th International Symposium on Experimental Algorithms*, volume 6049 of *Lecture Notes in Computer Science*, pages 462–473, 2010. URL: <http://hal.inria.fr/inria-00405478/>, doi:10.1007/978-3-642-13193-6_39.
- [16] 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>.
- [17] 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.
- [18] CGAL, The Computational Geometry Algorithms Library. URL: <http://www.cgal.org>.
- [19] 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.
- [20] M. Dehn. Transformation der Kurven auf zweiseitigen Flächen. *Mathematische Annalen*, 72(3):413–421, 1912. doi:10.1007/BF01456725.
- [21] Vincent Despré, Benedikt Kolbe, and Monique Teillaud. Half-minimizers and Delaunay triangulations on closed hyperbolic surfaces. Technical report, INRIA, 2020. URL: <https://hal.inria.fr/hal-03045921>.
- [22] 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.

- [23] Olivier Devillers, Sylvain Pion, and Monique Teillaud. Walking in a triangulation. *International Journal of Foundations of Computer Science*, 13:181–199, 2002. URL: <https://hal.inria.fr/inria-00102194>, doi:10.1142/S0129054102001047.
- [24] Olivier Devillers and Monique Teillaud. Perturbations for Delaunay and weighted Delaunay 3D Triangulations. *Computational Geometry: Theory and Applications*, 44:160–168, 2011. URL: <http://hal.inria.fr/inria-00560388/>, doi:10.1016/j.comgeo.2010.09.010.
- [25] Matthijs Ebbens. Delaunay triangulations on hyperbolic surfaces. Master’s thesis, University of Groningen, 2017. URL: <http://fse.studenttheses.ub.rug.nl/id/eprint/15727>.
- [26] Matthijs Ebbens, Iordan Iordanov, Monique Teillaud, and Gert Vegter. Systole of regular hyperbolic surfaces with an application to Delaunay triangulations. In *9th International Conference on Curves and Surfaces*, Arcachon, France, June 2018. URL: <https://hal.inria.fr/hal-01803443>.
- [27] Matthijs Ebbens, Iordan Iordanov, Monique Teillaud, and Gert Vegter. Delaunay triangulations of symmetric hyperbolic surfaces. In *Abstracts 35th European Workshop on Computational Geometry*, pages 15:1–15:8, 2019. <http://www.eurocg2019.uu.nl/papers/15.pdf>. URL: <https://hal.inria.fr/hal-02940717>.
- [28] Matthijs Ebbens, Iordan Iordanov, Monique Teillaud, and Gert Vegter. Delaunay triangulations of generalized Bolza surfaces, 2021. URL: <https://hal.inria.fr/hal-03080125>, arXiv:2103.05960.
- [29] Matthijs Ebbens, Hugo Parlier, and Gert Vegter. Minimal Delaunay triangulations of hyperbolic surfaces, 2020. arXiv:2011.09847.
- [30] Myfanwy E. Evans, Vanessa Robins, and Stephen T. Hyde. Periodic entanglement I: networks from hyperbolic reticulations. *Acta Crystallographica Section A*, 69(3):241–261, May 2013. doi:10.1107/S0108767313001670.
- [31] Myfanwy E. Evans, Vanessa Robins, and Stephen T. Hyde. Periodic entanglement II: weavings from hyperbolic line patterns. *Acta Crystallographica Section A*, 69(3):262–275, May 2013. doi:10.1107/S0108767313001682.
- [32] Myfanwy E. Evans and Gerd E. Schröder-Turk. In a material world: Hyperbolic geometry in biological materials. *Asia Pacific Mathematics Newsletter*, 5(2):21–30, 2015. URL: http://www.asiapacific-mathnews.com/05/0502/0021_0030.pdf.
- [33] Hershel M Farkas and Irwin Kra. Riemann surfaces. In *Riemann surfaces*, pages 9–31. Springer, 1992.
- [34] 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.
- [35] H. Helling. Diskrete Untergruppen von $SL(2, \mathbb{R})$. *Inventiones mathematicae*, 17(3):217–229, 1972. doi:10.1007/BF01425449.

- [36] Jordan 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>.
- [37] Jordan 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.
- [38] Jordan 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>.
- [39] 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>.
- [40] Mikhail G. Katz, Mary Schaps, and Uzi Vishne. Logarithmic growth of systole of arithmetic riemann surfaces along congruence subgroups. *J. Differential Geom.*, 76(3):399–422, 07 2007. URL: <http://projecteuclid.org/euclid.jdg/1180135693>.
- [41] 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>.
- [42] T Long Calvin. Elementary introduction to number theory. *Lexington: DC Heath and Company, LCCN*, 77171950, 1972.
- [43] J. R. Munkres. *Elements of Algebraic Topology*. Advanced book classics. Perseus Books, 1984.
- [44] 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.
- [45] 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.
- [46] J. Ratcliffe. *Foundations of Hyperbolic Manifolds*. Number vol. 10 in Graduate Texts in Mathematics. Springer, 2006.
- [47] Mael Rouxel-Labbé, Monique Teillaud, and Claudia Werner. 2D triangulations on the sphere. In *CGAL User and Reference Manual*. CGAL Editorial Board, 5.3 edition, 2021. URL: <https://doc.cgal.org/latest/Manual/packages.html#PkgTriangulationOnSphere2>.
- [48] Jim Ruppert. A Delaunay refinement algorithm for quality 2-dimensional mesh generation. *Journal of algorithms*, 18(3):548–585, 1995. doi:10.1006/jagm.1995.1021.

- [49] François Sausset, Gilles Tarjus, and Pascal Viot. Tuning the fragility of a glassforming liquid by curving space. *Physical Review Letters*, 101:155701(1)–155701(4), 2008. doi: [10.1103/PhysRevLett.101.155701](https://doi.org/10.1103/PhysRevLett.101.155701).
- [50] Paul Schmutz Schaller. Teichmüller space and fundamental domains of Fuchsian groups. *L'Enseignement Mathématique*, 45:169–187, 1999.
- [51] 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](https://doi.org/10.1093/comjnl/24.2.167).
- [52] 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](https://doi.org/10.1142/9789812831699_0011).
- [53] Chee Yap *et al.* The CORE library project. URL: http://cs.nyu.edu/exact/core_pages/intro.html.

A Statement and proof of Lemma 19

Lemma 19. *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 19 is the special case $m = 3$ of the following lemma. A proof was given in Ebbens’s master’s thesis [25], but for completeness we have included it here as well.

Lemma 20. *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 [44]. We use the same approach to prove Lemma 20.

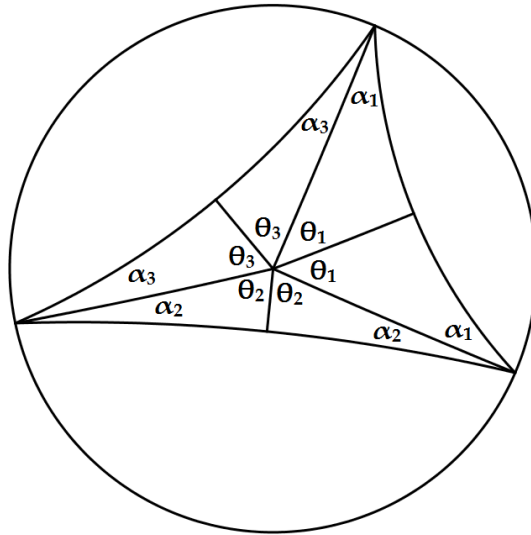
Consider $m = 3$. Divide P into three pairs of right-angled triangles with angles θ_j at the center of the circumscribed circle, angles α_j at the vertices and right angles at the edges of P (see Figure 20).

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

$$f(\theta_1, \theta_2, \theta_3) = \sum_{j=1}^3 \operatorname{arccot}(\cosh R \tan \theta_j) \tag{15}$$

Figure 20: Division of P into three pairs of right-angled triangles

subject to the constraints $\sum_{j=1}^3 \theta_j = \pi$ and $0 \leq \theta_j < \pi$, i.e., minimizing (15) over the triangle in \mathbb{R}^3 with vertices $(\pi, 0, 0)$, $(0, \pi, 0)$, $(0, 0, \pi)$. We parametrize this triangle as follows

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

for $0 < s < \pi/2$ and $|t| \leq s$. By (15), 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 Proofs omitted in Section 4

First, we give the proof of Lemma 11.

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

1. Consider a segment γ_j of separation $k \geq 4$. By symmetry of D_g , we can assume that γ_j is a segment between \mathbf{s}_0 and \mathbf{s}_k . The length of γ_j is greater than or equal to the distance between \mathbf{s}_0 and \mathbf{s}_k , which is given as the length of the common orthogonal line segment γ_j^\perp between \mathbf{s}_0 and \mathbf{s}_k (see Figure 21).

To find an expression for $\text{length}(\gamma_j^\perp)$, we draw line segments between the origin O and \mathbf{s}_0 and between O and \mathbf{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 \mathbf{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 21. Therefore, [8, Theorem 7.11.3(i)]

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

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 [8, Theorem 7.17.1(ii)]

$$\cosh(\text{length}(\gamma_j^\perp)) = \cosh(d(O, \mathbf{s}_0)) \sin(\frac{k\pi}{4g}) = \cot(\frac{\pi}{4g}) \sin(\frac{k\pi}{4g}) \quad (16)$$

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

2. Now, consider a segment γ_j of separation $k \geq 2$. Using the same argument as in Part 1, we see that formula (16) still holds. The lower bound for the length of γ_j follows from $\sin(\frac{k\pi}{4g}) \geq \sin(\frac{\pi}{2g})$ and a direct computation using properties of trigonometric functions.

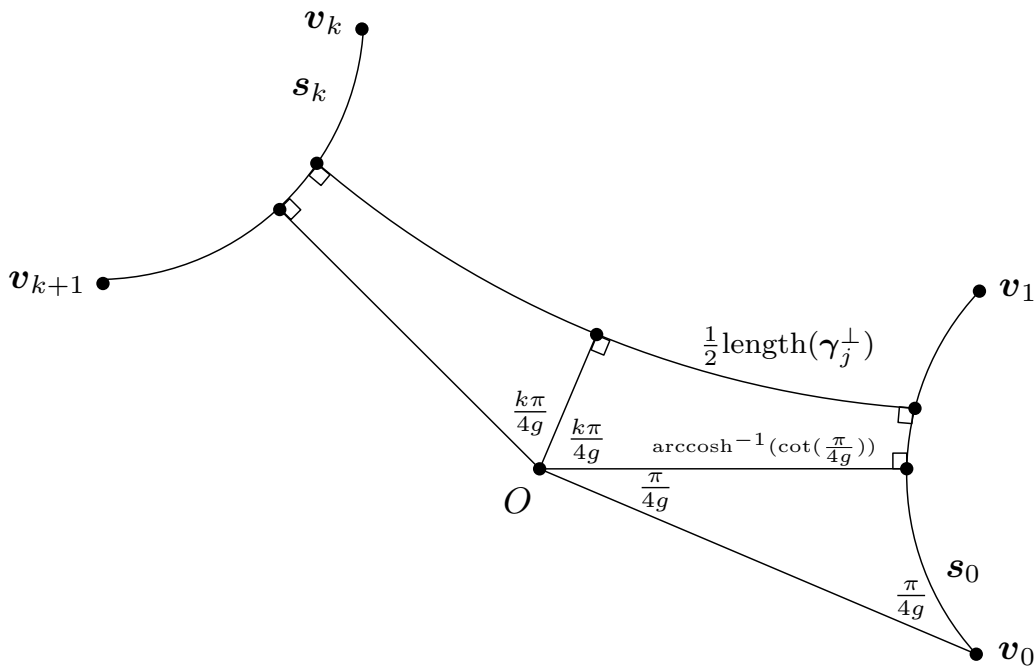


Figure 21: Computing the length of γ_j^\perp .

- Denote the two segments by γ_1 and γ_2 . We can assume without loss of generality that γ_1 is a 1-segment between s_0 and s_1 and γ_2 is a $(4g - 1)$ -segment between s_{2g+1} and s_{2g} (see Figure 22). Let x be the distance between p_1 and v_1 and let α_1 be the angle between γ_1 and s_0 . 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 s_{2g} .

By the hyperbolic sine rule,

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

so

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

In a similar way, we obtain

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

We minimize

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

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

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

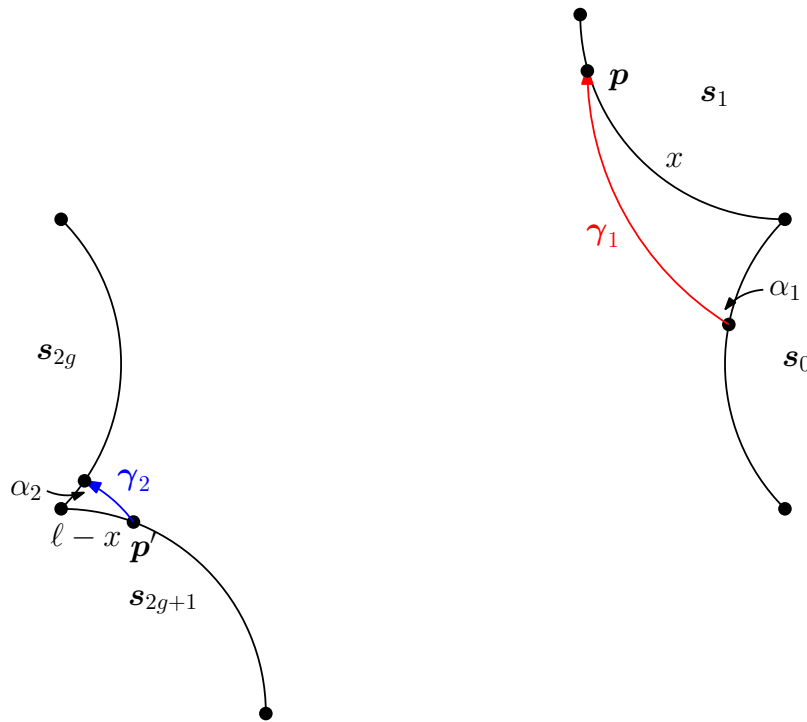


Figure 22: Construction in the proof of Part 3 of Lemma 11.

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

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

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

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

Then

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

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

4. Now, using the notation from Part 3, assume that $\gamma = \gamma_1 \cup \gamma_2$. By the argument in Part 3, the length of γ is minimal when p_1 is the midpoint of s_1 and p'_1 is the

midpoint of \mathbf{s}_{2g+1} , given any location of the starting point of γ_1 and the endpoint of γ_2 . By symmetry of the argument, it follows that $\text{length}(\gamma)$ is minimal when the starting point of γ_1 is the midpoint of \mathbf{s}_0 and the endpoint of γ_2 is the midpoint of \mathbf{s}_{2g} . It can be seen that the resulting curve is the curve constructed in the proof of Lemma 9, the length of which is ς_g . This finishes the proof. □

Second, we give the statement and proof of Lemma 21.

Lemma 21. *Let $m \in \{3, 2g - 2, 2g + 2, 4g - 3\}$ and define*

$$Y_m = \prod_{j=0}^m f_{j(2g+1)}.$$

Then

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

Proof. As mentioned in Section 2.4, the side pairing transformation f_k is represented by the matrix $\pm A_k$ with

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

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

$$\frac{1}{2} |\text{Tr}(Y_m)| = \frac{1}{2} \text{Tr} \left(\prod_{j=0}^m A_{j(2g+1)} \right),$$

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

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

$$\prod_{j=0}^{4g-3} A_{j(2g-1)} = \prod_{j=0}^{4g-1} A_{j(2g-1)} \cdot A_{(4g-1)(2g-1)}^{-1} A_{(4g-2)(2g-1)}^{-1} = A_{(4g-1)(2g-1)}^{-1} A_{(4g-2)(2g-1)}^{-1}$$

by the group relation. Hence, we can find $\frac{1}{2} |\text{Tr}(Y_m)|$ for $m = 4g - 3$ in a similar way by computing the trace of a product of two 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 $\chi(t) = t^2 - t \cdot \text{Tr}(A_k) + 1$. Because $\text{Tr}(A_k) = 2 \cot \alpha > 2$, 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

$$Y_m = S_0 \Lambda S_0^{-1} \cdot S_{2g-1} \Lambda S_{2g-1}^{-1} \cdots S_{(m-1)(2g-1)} \Lambda S_{(m-1)(2g-1)}^{-1} \cdot S_{m(2g-1)} \Lambda S_{m(2g-1)}^{-1}.$$

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,

$$\begin{aligned} Y_m &= S_0 (\Lambda T_{2g-1})^{m+1} T_{2g-1}^{-1} S_{m(2g-1)}^{-1}, \\ &= S_0 (\Lambda T_{2g-1})^{m+1} T_{-(m+1)(2g-1)} S_0^{-1}, \end{aligned}$$

where $T_{2g-1}^{-1} S_{m(2g-1)}^{-1} = T_{-(m+1)(2g-1)} S_0^{-1}$ in the second step can be verified with a direct computation. Writing $X = \Lambda T_{2g-1}$ and $Z_m = T_{-(m+1)(2g-1)}$, it follows that

$$\operatorname{Tr}(Y_m) = \operatorname{Tr}(X^{m+1} Z_m),$$

because Y_m and $X^{m+1} Z_m$ are conjugate matrices. It is known that [35, Equation (6)]

$$\operatorname{Tr}(X^{m+1} Z_m) = \operatorname{Tr}(X Z_m) \cdot \frac{\sin(m+1)\theta}{\sin \theta} - \operatorname{Tr}(Z_m) \cdot \frac{\sin m\theta}{\sin \theta}, \quad (17)$$

where θ is the angle such that $\operatorname{Tr}(X) = 2 \cos \theta$. To simplify Equation (17), we compute $\operatorname{Tr}(X Z_m)$, $\operatorname{Tr}(Z_m)$ and θ . Note that $T_k T_\ell = T_{k+\ell}$ for all k and ℓ , so

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

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

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

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

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

Similarly,

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

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

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

Finally,

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

so $\theta = \alpha$. Substituting the expressions for $\operatorname{Tr}(X Z_m)$, $\operatorname{Tr}(Z_m)$ and θ into Equation 17, we obtain after some simplifications

$$\frac{1}{2} \operatorname{Tr}(Y_m) = (-1)^{m/2} \cot \alpha \frac{\sin 2(m+1)\alpha}{\sin 2\alpha}.$$

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

C Proofs omitted in Section 5

First, we give the proof of Proposition 13, 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 the following lemma.

Lemma 22. *Let C_q be a Euclidean disk centered at O and passing through a point \mathbf{q} (Figure 23). Let H_1 and H_2 denote the two open half-planes bounded by the Euclidean line through O and \mathbf{q} . Let H_0 be a half-plane that contains \mathbf{q} , bounded by another Euclidean line passing through O but not through \mathbf{q} . Let $R_j = (H_0 \cap H_j) \setminus C_q, j = 1, 2$, and let $\mathbf{p}_j \in R_j, i = 1, 2$. The disk $C(\mathbf{p}_1, \mathbf{p}_2)$ through O, \mathbf{p}_1 , and \mathbf{p}_2 contains \mathbf{q} .*

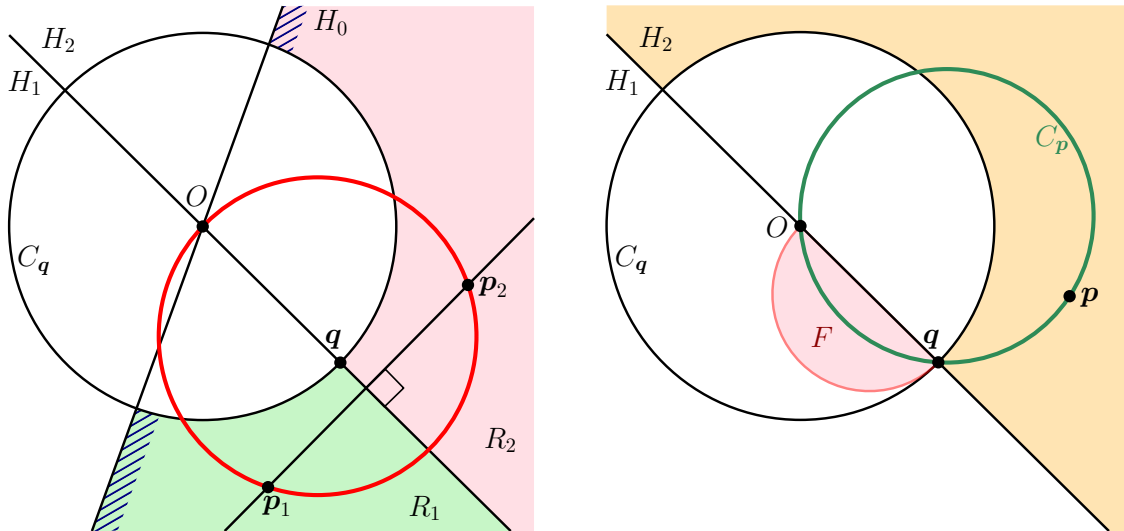


Figure 23: Illustrations for the proof of Lemma 22.

Proof. It is easy to verify that there exist pairs of points $(\mathbf{p}_1, \mathbf{p}_2) \in R_1 \times R_2$ for which the point \mathbf{q} lies inside the disk $C(\mathbf{p}_1, \mathbf{p}_2)$. For instance, consider a line perpendicular to the line through O and \mathbf{q} so that \mathbf{q} is closer to O than to their intersection point, as shown in Figure 23 - Left. If \mathbf{p}_1 lies on this perpendicular line and \mathbf{p}_2 is the reflection of \mathbf{p}_1 in the line through O and \mathbf{q} , then the disk $C(\mathbf{p}_1, \mathbf{p}_2)$ contains \mathbf{q} . Since this disk varies continuously when $(\mathbf{p}_1, \mathbf{p}_2)$ ranges over $R_1 \times R_2$, it is sufficient to prove that there are no pairs $(\mathbf{p}_1^*, \mathbf{p}_2^*) \in R_1 \times R_2$ for which \mathbf{q} lies on the boundary of $C(\mathbf{p}_1^*, \mathbf{p}_2^*)$.

Suppose, for a contradiction, that there exists a pair $(\mathbf{p}_1^*, \mathbf{p}_2^*) \in R_1 \times R_2$ for which $C(\mathbf{p}_1^*, \mathbf{p}_2^*)$ is a disk with \mathbf{q} and O on its boundary. Consider the disk C_q centered at O and passing through \mathbf{q} . Let F be the intersection of the disk with diameter $[O, \mathbf{q}]$ with the half-plane H_1 , as shown in Figure 23 - Right. For any point $\mathbf{p} \in H_2 \setminus C_q$, the circle C_p through O, \mathbf{q} , and \mathbf{p} has a non-empty intersection with H_1 , which is completely included in F , so in particular $C(\mathbf{p}_1^*, \mathbf{p}_2^*)$ intersects H_1 inside the disk with diameter $[O, \mathbf{q}]$. By a symmetric

observation, $C(\mathbf{p}_1^*, \mathbf{p}_2^*)$ also intersects H_2 inside the same disk. Therefore, $C(\mathbf{p}_1^*, \mathbf{p}_2^*)$ is the disk with diameter $[O, \mathbf{q}]$. This implies that both \mathbf{p}_1 and \mathbf{p}_2 lie in the disk $C_{\mathbf{q}}$, which is a contradiction. Therefore, there exists no pair $(\mathbf{p}_1^*, \mathbf{p}_2^*) \in R_1 \times R_2$ for which $C(\mathbf{p}_1^*, \mathbf{p}_2^*)$ has \mathbf{q} on its boundary. This finishes the proof. \square

Note that Lemma 22 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 13.

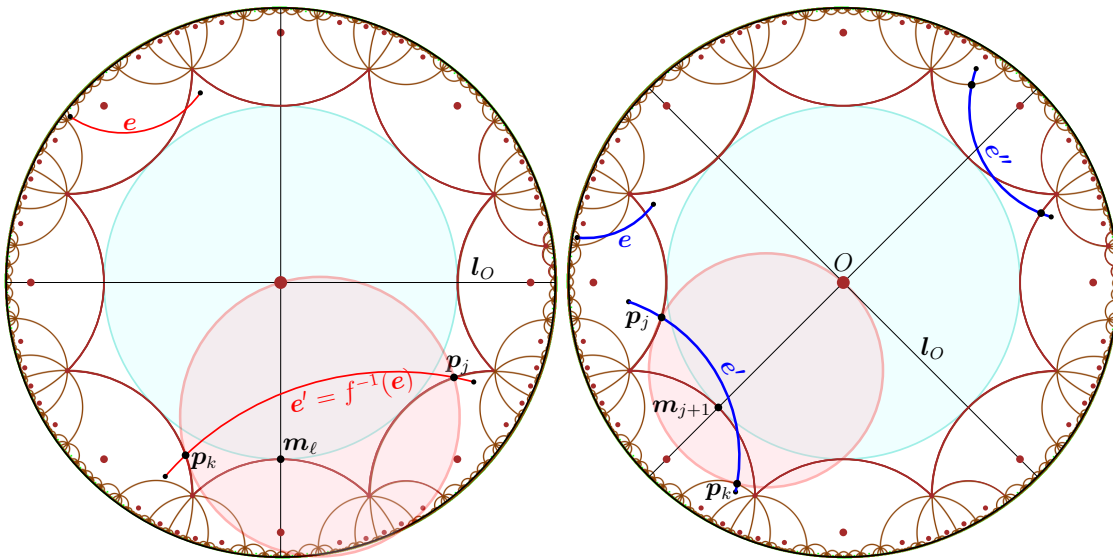


Figure 24: Illustration for the proof of Proposition 13 for $g = 2$.

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

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

Case A: The edge e crosses only one image of \tilde{D}_g before leaving $D_{\mathcal{N}_g}$. This case is illustrated in Figure 24 - Left. Let $f(\tilde{D}_g), f \in \mathcal{N}_g \setminus \{\mathbb{1}\}$ be the Dirichlet region that e crosses. The image $e' = f^{-1}(e)$ of e then crosses \tilde{D}_g , intersecting two of its non-adjacent sides s_j and s_k in the points \mathbf{p}_j and \mathbf{p}_k , respectively. We can assume without loss of generality that the hyperbolic segment $[\mathbf{p}_j, \mathbf{p}_k]$ does not contain the origin, since in that case any disk through \mathbf{p}_j and \mathbf{p}_k clearly contains the origin. Then, there exists a line l_O through O such that \mathbf{p}_j and \mathbf{p}_k are contained in the same half-space. Let \mathbf{m}_ℓ be the midpoint of a side between s_j and s_k in the same half-space as \mathbf{p}_j and \mathbf{p}_k . Consider the disk centered at O that passes

through \mathbf{m}_ℓ (and, of course, through all the other midpoints \mathbf{m}_k as well), and consider also the line through O and \mathbf{m}_ℓ . By Lemma 22, the disk $C(\mathbf{p}_j, \mathbf{p}_k)$ passing through O, \mathbf{p}_j , and \mathbf{p}_k contains \mathbf{m}_ℓ . Since O and \mathbf{m}_ℓ 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}_ℓ or O , therefore there is no empty disk that passes through \mathbf{p}_j and \mathbf{p}_k . Assume now that there is an empty disk that passes through the endpoints of e' . This empty disk can then be shrunk continuously so that it passes through \mathbf{p}_j and \mathbf{p}_k . The shrunk version of the disk must be also empty, which is a contradiction. Therefore, there is no empty disk passing through the endpoints of 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 24 - 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 \mathbf{s}_j and \mathbf{s}_{j+1} of \tilde{D}_g so that the hyperbolic line supporting e' separates O and the midpoint \mathbf{m}_{j+1} . Note that such an image of e exists always: e either separates O and the midpoint \mathbf{m}_{j+1} , or it separates an image of O under some translation f of Γ_g and \mathbf{m}_{j+1} ; in the second case, $f^{-1}(e)$ separates O and the midpoint $f^{-1}(\mathbf{m}_{j+1})$. The edge e' intersects also the side \mathbf{s}_k adjacent to \mathbf{s}_{j+1} in the Dirichlet region that shares the side \mathbf{s}_{j+1} with \tilde{D}_g (see Figure 24 - Right). Let \mathbf{p}_j and \mathbf{p}_k be the intersection points of e' with \mathbf{s}_j and \mathbf{s}_k , respectively. Consider the circle centered at the origin that passes through \mathbf{m}_{j+1} . Consider also the line through O and \mathbf{m}_{j+1} and the line \mathbf{l}_0 through O perpendicular to it. By Lemma 22, 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