Implementing Delaunay Triangulations of the Bolza Surface

Iordan Iordanov, Monique Teillaud

Implementing Delaunay Triangulations of the Bolza Surface

Iordan Iordanov[1] **and Monique Teillaud**[2]

1 **Loria, Inria Centre de recherche Nancy – Grand Est, Université de Lorraine, CNRS UMR 7503, Villers-lès-Nancy, France**
2 **Loria, Inria Centre de recherche Nancy – Grand Est, Université de Lorraine, CNRS UMR 7503, Villers-lès-Nancy, France**

Abstract

The CGAL library offers software packages to compute Delaunay triangulations of the (flat) torus of genus one in two and three dimensions. To the best of our knowledge, there is no available software for the simplest possible extension, i.e., the Bolza surface, a hyperbolic manifold homeomorphic to a torus of genus two.

In this paper, we present an implementation based on the theoretical results and the incremental algorithm proposed last year [2]. We describe the representation of the triangulation, we detail the different steps of the algorithm, we study predicates, and report experimental results.

1998 ACM Subject Classification I.3.5 Computational Geometry and Object Modeling

Keywords and phrases hyperbolic surface, Fuchsian group, arithmetic issues, Dehn's algorithm, CGAL

Digital Object Identifier 10.4230/LIPIcs.SoCG.2017.44

1 Introduction

Motivated by applications in various fields, some packages to compute periodic Delaunay triangulations in the Euclidean spaces $\mathbb{E}^2$ and $\mathbb{E}^3$ have been introduced in the CGAL library [4, 12] and have attracted a number of users. To the best of our knowledge, no software is available to compute periodic triangulations in a hyperbolic space. This would be a natural extension: periodic triangulations in $\mathbb{E}^2$ can be seen as triangulations of the two-dimensional (flat) torus of genus one; similarly, periodic triangulations in the hyperbolic plane $\mathbb{H}^2$ can be seen as triangulations of hyperbolic surfaces. The Bolza surface is a hyperbolic surface with the simplest possible topology, as it is homeomorphic to a genus-two torus. First steps in computing Delaunay triangulations of hyperbolic surfaces have recently been made [2]. Due to lack of space, we refer the reader to that paper for examples of applications.

All previous work mentioned above is generalizing the well-known incremental algorithm introduced by Bowyer [3], which has proved to be reasonably easy to implement and very efficient in practice. For each new point p, the set of conflicting simplices, i.e., simplices whose circumscribing ball contains p, are removed; then their union is triangulated by simply filling it with new simplices with apex p. This simple update operation relies on the fact that the union of conflicting simplices is always a topological ball. As proved earlier [2], for an input set S on a closed hyperbolic surface M, this property is ensured as soon as

$$\mathrm{sys}(M) > 2\delta_S, \tag{1}$$

where $\mathrm{sys}(M)$ denotes the *systole* of M, i.e., the length of a shortest non-contractible loop on M, and δ_S denotes the diameter of the largest disks that do not contain any point of S

in their interior. This condition ensures that there is no cycle of length one or two in the 1-skeleton of the Delaunay triangulation.

Two ideas have been proposed to fulfill this condition [2]. The first one consists in increasing the systole by using covering spaces of M, but it was shown to require at least 32 sheets for the Bolza surface, so this does not lead to a practical method. A more practical idea was quickly sketched in the last section of the same paper; it consists in initializing the triangulation with a set of "dummy" vertices that ensure that largest empty disks are small enough so that inequality (1) holds. As the diameter of largest empty disks cannot increase when new points are inserted, the condition will still be fulfilled when inserting points. If sufficiently many reasonably well-distributed points are inserted, then the dummy vertices can be removed from the triangulation without violating condition (1). In this paper, we elaborate on this approach and propose a first implementation.

We recall some background for the Bolza surface in Section 2. In Section 3 we propose a representation for Delaunay triangulations of the Bolza surface. Then we present the various steps of the construction in Section 4. We investigate the algebraic complexity of the algorithm in Section 5. Finally, we present some results of our implementation.

The source code, all Maple sheets, and the appendix are publicly available at

`https://members.loria.fr/Monique.Teillaud/DT_Bolza_SoCG17/`.

The software package will be submitted for integration in CGAL as soon as the documentation is completed.

2 The Bolza surface

Details and references for the background given in this section can be found in [2].

As the *Poincaré disk model* of the hyperbolic plane $\mathbb{H}^2$ is conformal, it is often used in applications. The hyperbolic plane is represented as the open unit disk $\mathcal{B}$ of the Euclidean plane $\mathbb{E}^2$. The boundary of $\mathcal{B}$ represents the set of points at infinity, denoted as $\mathcal{H}_\infty$. Hyperbolic lines, or geodesics, are represented as diameters of $\mathcal{B}$ or as arcs of circles orthogonal to $\mathcal{H}_\infty$. A hyperbolic circle is represented as a Euclidean circle contained in $\mathcal{B}$ and its hyperbolic center is the limit point of the pencil of circles that it generates with $\mathcal{H}_\infty$.

We denote the group of orientation-preserving isometries on $\mathbb{H}^2$ as $\mathrm{Isom}^+(\mathbb{H}^2)$. By identifying $\mathbb{E}^2$ with the complex plane $\mathbb{C}$, each $g \in \mathrm{Isom}^+(\mathbb{H}^2)$ is a mapping in the form $g(z) = \frac{\alpha z+\beta}{\overline{\beta} z+\overline{\alpha}}$, $z \in \mathbb{C}$ with matrix $g = \begin{bmatrix} \alpha & \beta \\ \overline{\beta} & \overline{\alpha} \end{bmatrix}$, where $\alpha, \beta \in \mathbb{C}$ and $|\alpha|^2 - |\beta|^2 = 1$. We are only interested here in one type of orientation-preserving isometries: the *hyperbolic* isometries, also called *translations*. A hyperbolic translation fixes two points at infinity and no point inside $\mathcal{B}$. The geodesic X_g through the two fixed points of a translation g is called the *axis* of g. Points lying on X_g are all translated along X_g by the same fixed distance $\ell(g)$ called the translation *length*. The length can be computed from the matrix as $\ell(g) = 2 \cdot \mathrm{arcosh}\left(\frac{1}{2}\mathrm{Tr}(g)\right)$, where $\mathrm{Tr}(g)$ denotes the trace of the matrix of g. A point that does not lie on X_g is translated by a distance greater than $\ell(g)$ along a curve equidistant from X_g (of course the distance between a point and its image is measured on the geodesic containing them). See Figure 1-Left.

A *hyperbolic surface* is a connected 2-dimensional manifold such that every point has a neighborhood isometric to a disk of $\mathbb{H}^2$. A closed (i.e., compact) and orientable hyperbolic surface is isometric to a quotient of $\mathbb{H}^2$ under the action of a Fuchsian group Γ (i.e., a discrete subgroup of $\mathrm{Isom}^+(\mathbb{H}^2)$) that contains only translations (and the identity). The *Bolza surface* is the simplest possible closed orientable hyperbolic surface. Consider the

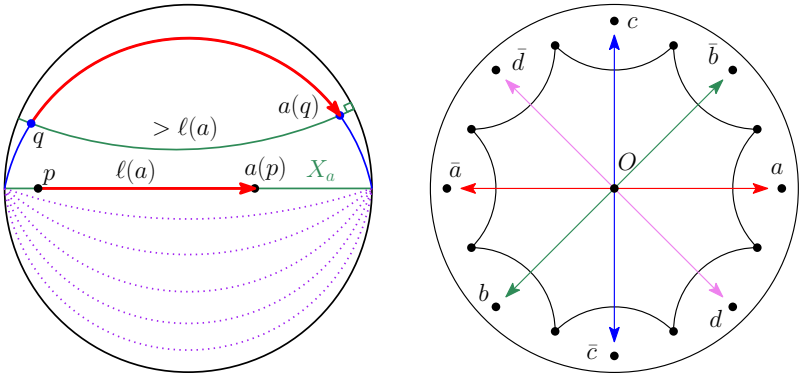

Figure 1 Left: Action of translation a on $\mathbb{H}^2$. **Right:** Regular octagon $\mathcal{D}_O$ and generators of $\mathcal{G}$.

regular hyperbolic octagon $\mathcal{D}_O$ centered at the origin O, with angles equal to $\pi/4$. The four hyperbolic translations a, b, c, and d that identify opposite sides of $\mathcal{D}_O$ generate a Fuchsian group denoted as $\mathcal{G}$. See Figure 1-Right.[1] For simplicity, we also denote as g the image gO of the origin by a translation g of $\mathcal{G}$.

The Bolza surface is defined as the quotient of $\mathbb{H}^2$ under the action of the group $\mathcal{G}$:

$$\mathcal{M} = \mathbb{H}^2/\mathcal{G}.$$

The projection map $\pi : \mathbb{H}^2 \to \mathcal{M} = \mathbb{H}^2/\mathcal{G}$ is a local isometry and a covering projection. The *Dirichlet region* $\mathcal{D}_p(\mathcal{G})$ for $\mathcal{G}$ centered at p is defined as the the closure of the open cell of p in the Voronoi diagram $VD_{\mathbb{H}}(\mathcal{G}p)$ of the infinite set of points $\mathcal{G}p$ in $\mathbb{H}^2$. From the compactness of $\mathcal{M}$, $\mathcal{D}_p(\mathcal{G})$ is a compact convex hyperbolic polygon with finitely many sides. The fact that $\mathcal{G}$ is non-Abelian leads to interesting difficulties. Among other properties, the Dirichlet regions $\mathcal{D}_p(\mathcal{G})$ and $\mathcal{D}_q(\mathcal{G})$ of two different points p and q do not always have the same combinatorics. The set of points $\mathcal{G}O$ is quite degenerate: all vertices of $VD_{\mathbb{H}}(\mathcal{G}O)$ have degree eight. See Figure 2-Left. The octagon $\mathcal{D}_O$ is in fact the Dirichlet region $\mathcal{D}_O(\mathcal{G})$ of the origin. Figure 2-Right illustrates notation that will be used throughout this paper: the vertices of $\mathcal{D}_O$ are denoted as $V_0, \ldots, V_7$ and their associated Delaunay circles are denoted as $C_0, \ldots, C_k$.

Each Dirichlet region $\mathcal{D}_p(\mathcal{G})$ is a *fundamental domain* for the action of $\mathcal{G}$ on $\mathbb{H}^2$, i.e., *(i)* $\mathcal{D}_p(\mathcal{G})$ contains at least one point of the preimage by π of any point of $\mathcal{M}$, and *(ii)* if it contains more than one point of the same preimage, then all these points lie on its boundary.

The group $\mathcal{G}$ has the following *finite presentation*:

$$\mathcal{G} = \left\langle a, b, c, d \mid abcd\bar{a}\bar{b}\bar{c}\bar{d} \right\rangle,$$

which denotes the quotient of the group $\langle a, b, c, d\rangle$ generated by a, b, c, and d, by the normal closure (i.e., the smallest normal subgroup) in $\langle a, b, c, d\rangle$ of the element $\mathcal{R}_{\mathcal{G}} = abcd\bar{a}\bar{b}\bar{c}\bar{d}$, called the *relation* of $\mathcal{G}$. Here and throughout the paper, $\overline{g}$ denotes the inverse of an element $g \in \mathcal{G}$. We use $\mathbb{1}$ to denote the identity of $\mathcal{G}$: $\forall\, g \in \mathcal{G},\ g\overline{g} = \overline{g}g = \mathbb{1}$, and $\mathcal{R}_{\mathcal{G}} = \mathbb{1}$ in $\mathcal{G}$.

The Bolza surface $\mathcal{M}$ is homeomorphic to a double torus. Its area (which is also the area of $\mathcal{D}_O$) is equal to $4\pi(\text{genus}(\mathcal{M}) - 1) = 4\pi$. The generators of $\mathcal{G}$ are naturally ordered around

[1] The octagon is rotated compared to [2]. The notation adopted now seems to be more standard in the literature, see for instance [1].

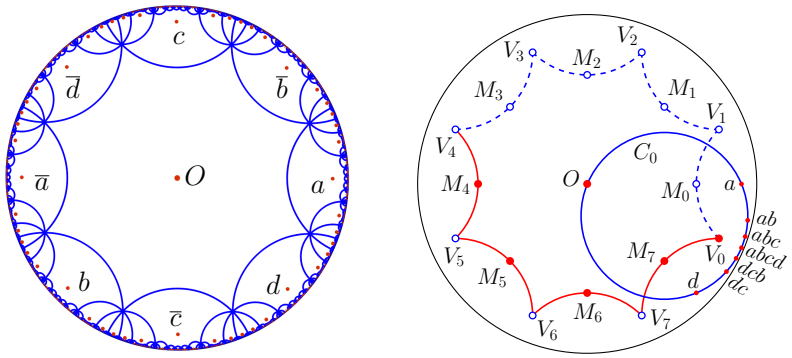

Figure 2 Left: Voronoi diagram of the infinite set of points $\mathcal{G}O$. **Right:** The original domain $\mathcal{D}$ and the images of O around the vertex V_0 of $\mathcal{D}_O$.

the octagon $\mathcal{D}_O$ as an ordered cyclical sequence $\mathcal{A} = \left[a, \overline{b}, c, \overline{d}, \overline{a}, b, \overline{c}, d\right] = [g_0, g_1, \ldots, g_7]$. The matrices of the elements $g_k,\ k = 0, 1, \ldots, 7$, are

$$g_k = \begin{bmatrix} \xi^2 & e^{ik\pi/4}\sqrt{2}\,\xi \\ e^{-ik\pi/4}\sqrt{2}\,\xi & \xi^2 \end{bmatrix}, \text{ where } \xi = \sqrt{1+\sqrt{2}}. \tag{2}$$

The translations g_k all have the same length, which is the systole of $\mathcal{M}$:

$$\mathrm{sys}(\mathcal{M}) = \ell(g_k) = 2 \cdot \mathrm{arcosh}\left(1+\sqrt{2}\right) \approx 3.05714,\ k = 0, 1, \ldots, 7.$$

3 Representation of the triangulation

As mentioned in the introduction, the use of dummy points allows us to always assume that **the set $\mathcal{P}$ of input points satisfies inequality** (1) for the Bolza surface $\mathcal{M}$.

We introduce the *original domain* $\mathcal{D} \subset \mathcal{D}_O$ for $\mathcal{M}$, which contains exactly one point of the fiber under π of each point on the surface $\mathcal{M}$. See Figure 2-Right: $\mathcal{D}$ consists of the interior of $\mathcal{D}_O$, its four "solid" sides, and one vertex of the octagon (chosen to be V_0).[2]

We can consider that all points of $\mathcal{P}$ lie in $\mathcal{D}$. Similarly, we will now define a unique representative of each face of the Delaunay triangulation $DT_{\mathcal{M}}(\mathcal{P})$ of $\mathcal{M}$ defined by $\mathcal{P}$.

3.1 Canonical representative of a face

The definition of the canonical representative of a face will rely on Theorem 2, which is reminiscent of the result proved for the flat torus by Dolbilin and Huson [9] and recalled in [5, Lemma 6.3].

We denote the hyperbolic distance between two points p and q in $\mathbb{H}^2$ as $\mathrm{dist}_{\mathbb{H}}(p, q)$ and the (hyperbolic) segment with endpoints p and q as $[p, q]$. Let us recall our abuse of notation: g denotes both a translation and the point gO. The points $M_k,\ k = 0, \ldots, 7$ visible on Figure 2-Right are defined as the midpoints of V_k and V_{k+1} (indices are meant modulo 8).

Let $\mathcal{U}_{\mathcal{D}}$ be the union of the disks bounded by the circles $C_k, k = 0, 1, \ldots, 7$, and let $\mathcal{C}_{\mathcal{D}}$ be the boundary of $\mathcal{U}_{\mathcal{D}}$. See Figure 3-Left.

[2] $\pi(V_k) = \pi(V_0),\ k = 1, \ldots, 7$: $V_5 = \overline{a}V_0, V_2 = \overline{b}V_5, V_7 = \overline{c}V_2, V_4 = \overline{d}V_2, V_1 = aV_4, V_6 = bV_1, V_3 = cV_6, V_0 = dV_3$.

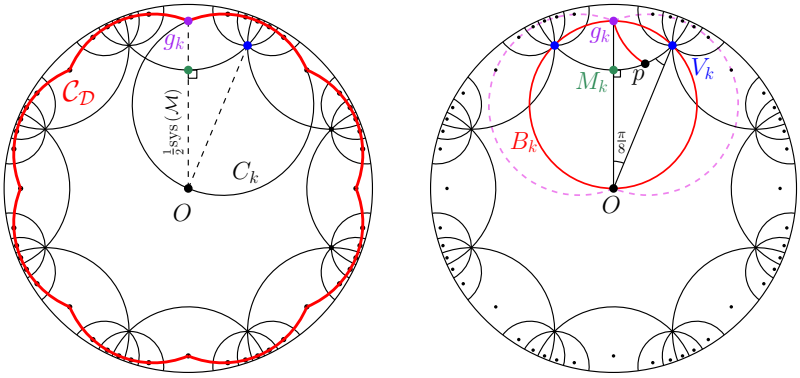

Figure 3 Left: Curve $\mathcal{C}_\mathcal{D}$ (in bold). **Right:** The distance between $\mathcal{C}_\mathcal{D}$ and $\partial\mathcal{D}_O$ is realized as the distance of the points g_k and M_k.

▶ **Lemma 1.** *The distance between $\mathcal{C}_\mathcal{D}$ and $\partial\mathcal{D}_O$ is equal to* $\mathrm{dist}_\mathbb{H}(M_k, g_k) = \frac{1}{2}\mathrm{sys}(\mathcal{M})$, $k = 0, 1, \ldots, 7$.

Proof. Using symmetries, we get (see Figure 3-Right):

$$\mathrm{dist}_\mathbb{H}(\mathcal{C}_\mathcal{D}, \partial\mathcal{D}_O) = \min_{p\in[V_k,M_k],\ q\in C_k\cap\mathcal{C}_\mathcal{D}} \mathrm{dist}_\mathbb{H}(p, q)\,.$$

The hyperbolic circle B_k centered at M_k and passing through O also contains the points V_k, g_k and V_{k+1}: indeed, by definition of the Dirichlet region of O, segment $[V_k, V_{k+1}]$ lies on the bisecting line of O and g_k, moreover $[O, g_k]$ lies on the bisecting line of V_k and V_{k+1}; in addition, $\mathrm{dist}_\mathbb{H}(M_k, V_k) = \mathrm{dist}_\mathbb{H}(O, M_k)$ since the angles of the triangle (O, V_k, M_k) at O and at V_k are both equal to $\pi/8$.

The points O and g_k are the intersections of C_k and C_{k+1}, and the segment $[O, g_k]$ is a diameter of B_k, so B_k is contained in the union of the disks bounded by C_k and C_{k+1}. The segment $[V_k, M_k]$ is a radius of B_k, so for any point $q \in C_k \cap \mathcal{C}_\mathcal{D}$ and for any $p \in [M_k, V_k]$, $\mathrm{dist}_\mathbb{H}(p, B_k) \leq \mathrm{dist}_\mathbb{H}(p, q)$. Equality is attained when $q = g_k$, so: $\mathrm{dist}_\mathbb{H}(\mathcal{C}_\mathcal{D}, \partial\mathcal{D}_O) = \min_{p\in[V_k,M_k]} \mathrm{dist}_\mathbb{H}(p, g_k)$. For every point $p \in [V_k, M_k]$, $\mathrm{dist}_\mathbb{H}(g_k, p) \geq \mathrm{dist}_\mathbb{H}(g_k, M_k)$ because the angle $g_k M_k p$ is right. The result follows. ◀

Let $\mathcal{D}_g$ denote the closure of the region of g in $VD_\mathbb{H}(\mathcal{G}O)$; $\mathcal{D}_g$ is the image of $\mathcal{D}_O$ by the translation g. The infinite set of regions $\mathcal{D}_g$, for $g \in \mathcal{G}$, form a tiling of the plane $\mathbb{H}^2$ (it was shown on Figure 2-Left.) We define $\mathcal{N}$ as the set of translations g in $\mathcal{G}$ for which $\mathcal{D}_g \cap \mathcal{D}_O \neq \emptyset$. The set $\mathcal{N}$ has 48 elements; it is naturally ordered counterclockwise around O, following the boundary of $\mathcal{D}_O$. Each element ν of $\mathcal{N}$ has an index $\mathrm{index}_\mathcal{N}(\nu)$ in this sequence. We choose *abcd* as the first element for the sequence $\mathcal{N}$, i.e., $\mathrm{index}_\mathcal{N}(abcd) = 0$. See Figure 4-Left. We define $\mathcal{D}_\mathcal{N}$ as

$$\mathcal{D}_\mathcal{N} = \bigcup_{g\in\mathcal{N}} \mathcal{D}_g.$$

▶ **Theorem 2.** *Let $\mathcal{P} \subset \mathbb{H}^2$ be a set of points such that inequality* (1) *holds for $\mathcal{M}$. If a 2-face σ of $\mathrm{DT}_\mathbb{H}(\mathcal{GP})$ has at least one of its vertices in $\mathcal{D}_O$, then σ is contained in $\mathcal{D}_\mathcal{N}$.*

From now on, 2-faces will simply be named *faces*, as done in CGAL.

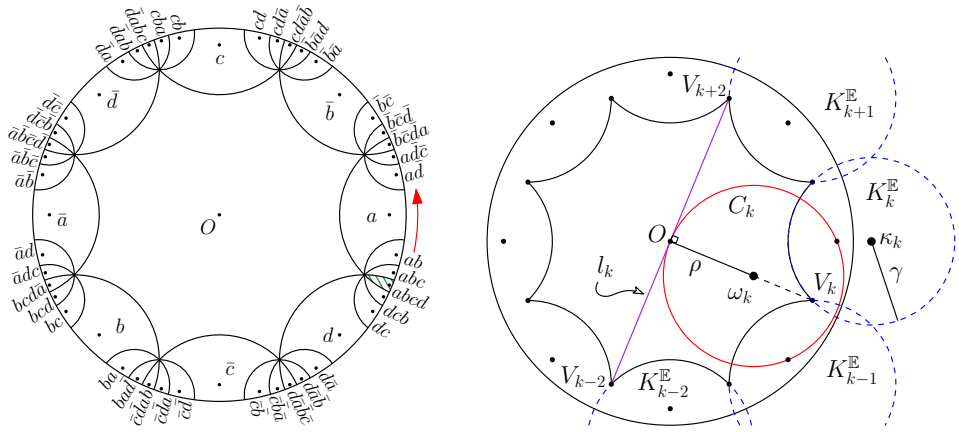

Figure 4 Left: The translations in $\mathcal{N}$. **Right:** Proof of Theorem 2.

Proof. Let σ be a face in $DT_{\mathbb{H}}(\mathcal{GP})$ with at least one vertex in $\mathcal{D}_O$. By definition of $\delta_{\mathcal{P}}$, the circumscribing disk of σ has diameter smaller than $\delta_{\mathcal{P}}$, which is smaller than $\frac{1}{2}\text{sys}(\mathcal{M})$ by inequality (1). Lemma 1 allows us to conclude that this disk is contained in $\mathcal{U}_{\mathcal{D}}$.

We will now prove that $\mathcal{U}_{\mathcal{D}}$ is contained in $\mathcal{D}_{\mathcal{N}}$, by proving that each circle C_k, for $k \in \{0, 1, \ldots, 7\}$ is contained in $\mathcal{D}_{\mathcal{N}}$. A circle C_k is centered at the Voronoi vertex V_k; it passes through the origin O and its images under the action of seven consecutive elements of $\mathcal{N}$. Rotating C_k around V_k by $\pi/4$ maps each of these eight points (and its Voronoi region) to the next one along C_k. This rotational symmetry shows that in order to prove that $C_k \subset \mathcal{D}_{\mathcal{N}}$, it is enough to prove that C_k intersects only the two sides of $\mathcal{D}_O$ that are incident to its hyperbolic center V_k.

Indices below are again taken modulo eight, e.g., we write V_{k+1} instead of $V_{k+1 \bmod 8}$. Let us first show that C_k intersects the sides $[V_{k-1}, V_k]$ and $[V_k, V_{k+1}]$ of $\mathcal{D}_O$. Consider a hyperbolic triangle (O, V_k, V_{k+1}). Its angle at O is $\pi/4$, while the angles at the vertices V_k and V_{k+1} are $\pi/8$. From the Hyperbolic law of sines,[3] we conclude that the length of $[V_k, V_{k+1}]$ is larger than the length of $[O, V_k]$. The result follows, since the segment $[O, V_k]$ is a radius of C_k.

Consider now the line segment $l_k = [V_{k-2}, V_{k+2}], k = 0, 1, \ldots, 7$, which cuts the octagon into two halves. See Figure 4-Right. Both l_k and C_k contain O; moreover l_k is perpendicular to the segment $[O, V_k]$, which is supported by a diameter of C_k. So l_k and C_k are tangent at O and l_k separates C_k from the other half of the octagon, thus C_k cannot intersect any side $[V_{k+j}, V_{k+j+1}]$ of $\mathcal{D}_O$ for $j = 2, 3, 4, 5$.

Using the fact that hyperbolic circles in the Poincaré disk model are Euclidean circles (see Section 2), we continue the proof and the computations in the Euclidean plane $\mathbb{E}^2$. The sides of $\mathcal{D}_O$ are supported by the Euclidean circles $K^{\mathbb{E}}_j = (\kappa_j, \gamma), j = 0, 1, \ldots, 7$ shown on Figure 4-Right. The centers and radii of $K^{\mathbb{E}}_j$, as well as the Euclidean centers ω_k and radii ρ of C_k are given in Table 1 and computed with Maple.

What is left to do is to show that C_k does not intersect either $K^{\mathbb{E}}_{k+1}$ or $K^{\mathbb{E}}_{k-2}$. By symmetry, it suffices to consider $K^{\mathbb{E}}_{k+1}$. The signed Euclidean distance of the circles C_k and $K^{\mathbb{E}}_{k+1}$ is

[3] If A, B, C are the sides of a hyperbolic triangle and $\vartheta_A, \vartheta_B, \vartheta_C$ the angles opposite to each side, then

$$\frac{\sin(\vartheta_A)}{\sinh(A)} = \frac{\sin(\vartheta_B)}{\sinh(B)} = \frac{\sin(\vartheta_C)}{\sinh(C)}.$$

Table 1 Expressions for the Euclidean radii and centers of $K_j^{\mathbb{E}}$ and C_k, $i, k = 0, 1, \ldots, 7$.

Quantity	Notation	Expression	Approximation
radius of $K_j^{\mathbb{E}}$	γ	$\sqrt{\frac{\sqrt{2}-1}{2}}$	0.4551
center of $K_j^{\mathbb{E}}$	κ_j	$e^{ij\pi/4}\sqrt{\frac{\sqrt{2}+1}{2}}$	–
radius of C_k	ρ	$\sqrt{(2-\sqrt{2})(\sqrt{2}-1)}$	0.4926
center of C_k	ω_k	$e^{i(2k+7)\pi/8}\sqrt{3\sqrt{2}-4}$	–

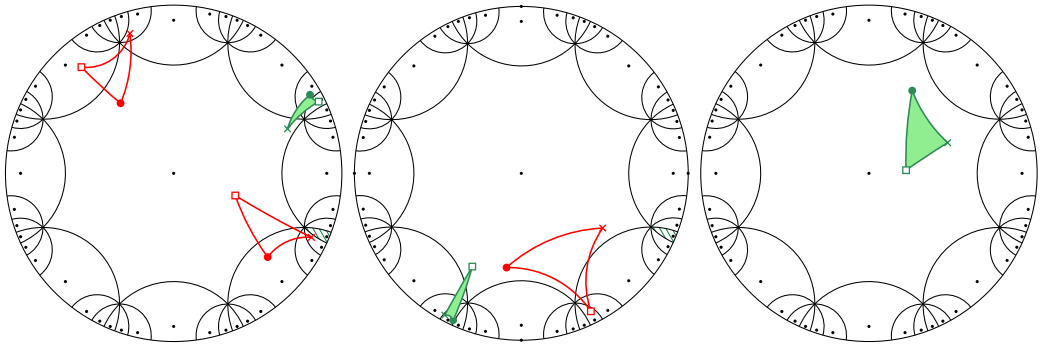

Figure 5 Examples of faces of $DT_{\mathbb{H}}(\mathcal{GP})$ with one, two and three vertices in $\mathcal{D}$, that project to the same face on $\mathcal{M}$. Their respective vertices drawn as a dot project to the same vertex on $\mathcal{M}$ (same for cross and square). The canonical representative is the shaded face.

$\text{dist}_{\mathbb{E}}\left(C_k, K_{k+1}^{\mathbb{E}}\right) = \text{dist}_{\mathbb{E}}\left(\omega_k, \kappa_{k+1}\right) - \rho - \gamma$. Maple calculations yield: $\text{dist}_{\mathbb{E}}\left(C_k, K_{k+1}^{\mathbb{E}}\right) = \frac{\sqrt{\sqrt{2}-1}}{2}\left(3\sqrt{2} - \sqrt{2}\sqrt{4-2\sqrt{2}} - 1\right)$. The last factor is positive: $\left(3\sqrt{2}-1\right)^2 = 19 - 6\sqrt{2}$, $\left(\sqrt{2}\sqrt{4-2\sqrt{2}}\right)^2 = 8 - 4\sqrt{2}$, and clearly $19 - 6\sqrt{2} > 8 - 4\sqrt{2} > 0$. This shows that C_k and $K_{k+1}^{\mathbb{E}}$ do not intersect. ◀

Let $\mathcal{P} \subset \mathcal{D}$ be a set of points satisfying inequality (1) for $\mathcal{M}$. The rest of this section is dedicated to the choice of a unique *canonical representative* σ^c in $DT_{\mathbb{H}}(\mathcal{GP})$ for each face $\boldsymbol{\sigma}$ in $DT_{\mathcal{M}}(\mathcal{P})$.

Let $\boldsymbol{\sigma}$ be a face in $DT_{\mathcal{M}}(\mathcal{P})$. By definition of $\mathcal{D}$, each vertex of $\boldsymbol{\sigma}$ has a unique preimage by π in $\mathcal{D}$, so, the set

$$\Sigma = \left\{\sigma \in \pi^{-1}(\boldsymbol{\sigma}) \;\middle|\; \sigma \text{ has at least one vertex in } \mathcal{D}\right\} \tag{3}$$

contains at most three faces. See Figure 5. When Σ contains only one face, then this face is completely included in $\mathcal{D}$, and we naturally choose it to be σ^c. Let us now assume that Σ contains two or three faces. From Theorem 2, each face $\sigma \in \Sigma$ is contained in $\mathcal{D}_{\mathcal{N}}$. So, for each vertex v of σ, there is a unique translation $\nu(v, \sigma)$ in $\mathcal{N} \cup \{\mathbb{1}\}$ such that v lies in $\nu(v, \sigma)\mathcal{D}$.

We consider all faces in $DT_{\mathbb{H}}(\mathcal{GP})$ oriented counterclockwise. For $\sigma \in \Sigma$, we denote as $v_\sigma^{\text{first\_out}}$ the first vertex of σ (in the counterclockwise order) that is not lying in $\mathcal{D}$. Using the indexing on $\mathcal{N}$ defined above, we can now choose σ^c as the face of Σ whose first vertex lying outside $\mathcal{D}$ is "closest" to the region $abcd\mathcal{D}$ in the counterclockwise order around O:

▶ **Definition 3** (Canonical representative)**.** With the notation defined above, the canonical representative of a face $\boldsymbol{\sigma}$ of $DT_{\mathcal{M}}(\mathcal{P})$ is the face $\sigma^c \in \Sigma$ such that

$$\text{index}_{\mathcal{N}}(\nu(v_{\sigma^c}^{\text{first\_out}}, \sigma^c)) = \min_{\sigma \in \Sigma} \text{index}_{\mathcal{N}}(\nu(v_{\sigma}^{\text{first\_out}}, \sigma)).$$

3.2 Data structure in CGAL

General two-dimensional triangulation data structures in CGAL store the vertices and faces of the triangulation. Each vertex stores a point and a pointer to one of its incident faces. Each face stores three pointers to its vertices v_0, v_1, and v_2, as well as three pointers to its three adjacent faces. Edges are not explicitly stored.

As mentioned above, we can assume that all input points of $\mathcal{P}$ lie in $\mathcal{D}$. We adapt the CGAL structure to store a triangulation of the Bolza surface. Each vertex v of $DT_{\mathcal{M}}(\mathcal{P})$ represents an orbit under the action of $\mathcal{G}$; it stores the point of $\pi^{-1}(v)$ that belongs to $\mathcal{D}$. Faces of $DT_{\mathcal{M}}(\mathcal{P})$ are stored through their canonical representative in $DT_{\mathbb{H}}(\mathcal{GP})$. Concretely, in addition to the pointers to vertices and neighbors, each face σ^c stores the three translations $\nu(v_i, \sigma^c) \in \mathcal{N},\ i = 0, 1, 2$ defined at the end of Section 3.1. In this way, for a given face σ^c in the structure, the corresponding canonical representative is the triangle in $\mathbb{H}^2$ whose vertices are the images by $\nu(v_i, \sigma^c)$ of the point in $\mathcal{D}$ stored in v_i for $i = 0, 1, 2$. The translations $\nu(v_i, \sigma^c)$ play a similar role as the so-called "offsets" of the CGAL Euclidean periodic triangulations.

We choose to represent translations in the faces of the triangulation data structure as words. This is detailed below.

Translations as words. We consider the cyclical sequence $\mathcal{A}$ formed by generators of $\mathcal{G}$ (see Section 2) as an alphabet, and we denote the set of words on $\mathcal{A}$ as $\mathcal{A}^*$. Each translation g in $\mathcal{G}$ can be seen as a word in $\mathcal{A}^*$, also denoted as g. For two translations $g, g' \in \mathcal{G}$, the composition (or multiplication) gg' corresponds to the concatenation of the two words g and g'. Recall that composition is not commutative. We have seen in the two previous sections that we only need to store translations in $\mathcal{N}$. Let us note here that $\mathcal{N}$ is closed under inversion, but not under composition.

The finite presentation of $\mathcal{G}$ captures the fact that a translation $g \in \mathcal{G}$ does not have a unique representation in terms of the generators (see Section 2). To obtain a unique representation of the translations that are involved in our algorithm, we slightly modify Dehn's algorithm. Dehn's algorithm solves the *word problem* (i.e., the problem of deciding whether a given word on the generators of a group is equal to the group identity) in the case of fundamental groups of closed orientable surfaces of genus at least 2 [6, 11].[4]

Let us present our implementation, tailored to our specific case.

We encode each element $g_k,\ k = 0, 1, \ldots, 7$ of $\mathcal{A}$ as its index k. By concatenation, each word of $\mathcal{A}^*$ is encoded as a sequence of integers.

Let w be a non-trivial word in $\mathcal{A}^*$. The first step of the reduction consists in freely reducing w, i.e., removing all sub-words of the form $g\overline{g}$ or $\overline{g}g$ for $g \in \mathcal{A}$. With our encoding, two elements g_i and g_j of $\mathcal{A}$ are inverses in $\mathcal{G}$ if $i = (j + 4) \bmod 8$.

The relation $\mathcal{R}_{\mathcal{G}} = abcd\overline{a}\overline{b}\overline{c}\overline{d}$ is encoded as 05274163. Let us note that any cyclical permutation of $\mathcal{R}_{\mathcal{G}}$ or of its inverse $\overline{\mathcal{R}_{\mathcal{G}}}$ is equal to $\mathbb{1}$ in $\mathcal{G}$. This can be viewed in another

[4] For interesting historical facts on this topic, see [14]. Software solving the word problem can be found for instance in [10, 13].

way by considering $\mathcal{R}_{\mathcal{G}}^{\infty}$, the infinite word formed by infinitely many concatenations of $\mathcal{R}_{\mathcal{G}}$: any subsequence $\mathcal{R}$ of $\mathcal{R}_{\mathcal{G}}^{\infty}$ or $\overline{\mathcal{R}_{\mathcal{G}}^{\infty}}$ with $|\mathcal{R}| = |\mathcal{R}_{\mathcal{G}}|$ is a relation in $\mathcal{G}$, i.e., it reduces to $\mathbb{1}$. (Here $|\cdot|$ denotes the length of a word.) The next step of the reduction consists in detecting a factorization of the (now freely-reduced) word w of the form $w = w_\lambda w_\mu w_\kappa$, where $w_\mu t$ is a relation $\mathcal{R}$ for some $t \in \mathcal{A}^*$ with $|t| < |w_\mu|$. Then $|w_\mu| > |\mathcal{R}_{\mathcal{G}}|/2 = 4$ and w_μ can be substituted in w by $\bar{t}$, which yields the word $w_\lambda \bar{t} w_\kappa$ with length shorter than $|w|$.

In our implementation, to find the sub-word w_μ, we use the fact that a sequence of letters $(g_{k_j})_{j=0,1,\ldots,n},\ g_{k_j} \in \mathcal{A}$, is a sub-word of $\mathcal{R}_{\mathcal{G}}^{\infty}$ of length n if, for every j from 0 to $n-1$, $k_{j+1} = (k_j + 5) \bmod 8$. Similarly, (g_{k_j}) is a sub-word of $\overline{\mathcal{R}_{\mathcal{G}}^{\infty}}$ of length n if for every j from 0 to $n-1$, $k_{j+1} = (k_j - 5) \bmod 8$. It holds that $|\mathcal{R}| < 2|w|$, so all words in $\mathcal{A}^*$ with length less than $2|w|$ can be listed in order to find such a word $\mathcal{R}$.

The two steps are repeated until $w = \mathbb{1}$ or until w cannot be further reduced. In the original algorithm by Dehn, words of length $|\mathcal{R}_{\mathcal{G}}|/2$ are not reduced. In order to have a unique representations of words of length four, we introduce a small modification to the algorithm: whenever we get an irreducible word w with $|w| = 4$, we check whether w is a sub-word of $\overline{\mathcal{R}_{\mathcal{G}}^{\infty}}$. If so, we return $\overline{w}$; in all other cases, we return w.

Dehn's algorithm terminates in a finite number of steps and its time complexity is polynomial in the length of the input word. Note that we reduce words that are formed by the concatenation of two or three words in $\mathcal{N}$; this will become clear in Section 4.2. Since the longest word in $\mathcal{N}$ has four letters, the longest words that we reduce have length 12.

4 Constructing the triangulation

Let us now describe the steps of our implementation of the incremental algorithm that was quickly recalled in the introduction.

4.1 Initialization

The set $\mathcal{Q}$ of 14 dummy points proposed in [2, Section 4.2] is as follows:

- the origin O;
- the eight midpoints P_k of the hyperbolic segments $[O, V_k],\ k = 0, 1, \ldots 7$;
- the midpoints $M_k,\ k = 4, 5, 6, 7$ of the closed sides of $\mathcal{D}$;
- the vertex V_0 of $\mathcal{D}$.

The canonical representatives of the 32 faces forming the Delaunay triangulation of $\mathcal{Q}$ are shown in Figure 6-Left. They can be constructed in four iterations ($i = 0, 1, 2, 3$) by using the numbering shown in Figure 6-Right (but faces are not numbered in the code).

The coordinates of the dummy points are algebraic numbers, as reported in Table 2. They have been computed using Maple. These exact coordinates would increase the algebraic degree of the predicates (studied in Section 5) in an artificial way; therefore, we introduce a set $\mathcal{Q}'$ of rational approximations of the points in $\mathcal{Q}$. See the third column of Table 2. We have verified that $DT_{\mathcal{M}}(\mathcal{Q})$ and $DT_{\mathcal{M}}(\mathcal{Q}')$ have identical combinatorial structures. We initialize the triangulation as $DT_{\mathcal{M}}(\mathcal{Q}')$.[5]

4.2 Finding faces in conflict with a new point

Let $p \in \mathcal{P} \subset \mathcal{D}$ be a new point to be inserted in the Delaunay triangulation. Consider $\boldsymbol{\sigma}$ a face in $DT_{\mathcal{M}}(\mathcal{P})$, and the set Σ defined in (3). We say that σ^c is *in conflict* with the input

[5] Note that any other set of points that satisfies condition (1) could be used to initialize the triangulation.

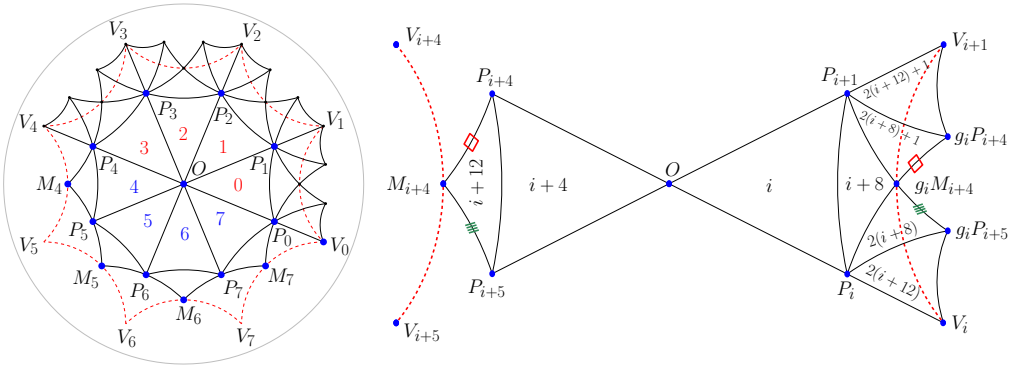

Figure 6 Left: Delaunay triangulation of $\mathcal{M}$ defined by the dummy points. **Right:** Zooming in on the faces created in iteration i. Note the identification of the marked edges.

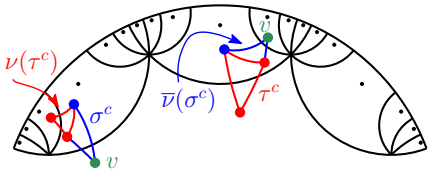

Figure 7 Translating τ^c by $\nu = \nu_{nbr}(\sigma^c, \tau^c)$ gives a face adjacent to σ^c.

point p if there exists a face $\sigma \in \Sigma$ whose circumscribing disk contains p.

Recall that, since hyperbolic circles are Euclidean circles, a Delaunay triangulation in $\mathbb{H}^2$ has exactly the same combinatorics as the Euclidean Delaunay triangulation of the same points. Consequently, the Euclidean Delaunay triangle containing p gives us a hyperbolic Delaunay face in conflict with p; the Euclidean and hyperbolic faces will both be denoted as σ_p, which should not introduce any confusion. To find this triangle, we adapt the so-called *visibility walk* [7]. This walk starts from an arbitrary face, then, for each visited face, it visits one of its neighbors, until a face containing p is found. Before specifying how the neighbor to be visited is specified in the case of the Bolza surface, we introduce the notion of *neighbor translation*.

▶ **Definition 4** (Neighbor translation). Let $\boldsymbol{\sigma}, \boldsymbol{\tau}$ be two adjacent faces in $DT_{\mathcal{M}}(\mathcal{P})$ and σ, τ two of their preimages by π in $DT_{\mathbb{H}}(\mathcal{GP})$. We define the neighbor translation $\nu_{nbr}(\sigma, \tau)$ from σ to τ as the translation of $\mathcal{G}$ such that $\nu_{nbr}(\sigma, \tau)\,\tau$ is adjacent to σ in $DT_{\mathbb{H}}(\mathcal{GP})$.

Let v be a vertex common to $\boldsymbol{\sigma}$ and $\boldsymbol{\tau}$, and let v_σ and v_τ the vertices of σ and τ that project on v by π. We can compute the neighbor translation from σ to τ as

$$\nu_{nbr}(\sigma, \tau) = \nu(v_\tau, \tau)\,\overline{\nu(v_\sigma, \sigma)}.$$

Figure 7 illustrates the neighbor translation of the canonical representatives of $\boldsymbol{\sigma}$ and $\boldsymbol{\tau}$. It can be easily seen that $\nu_{nbr}(\sigma, \tau) = \nu(v, \tau)\,\overline{\nu(v, \sigma)} = \overline{\nu(v, \sigma)}\,\overline{\nu(v, \tau)} = \overline{\nu_{nbr}(\tau, \sigma)}$.

We define the *location translation* ν_{loc} as follows: let σ_p be the Euclidean Delaunay triangle containing p. ν_{loc} is the translation that moves σ_p^c to σ_p.

The location procedure starts from a face incident to O. Then, for each visited face σ of $DT_{\mathbb{H}}(\mathcal{GP})$, we consider the Euclidean edge e defined by two of the vertices of σ. With a simple orientation test, we can check whether the Euclidean line supporting e separates p from the vertex of σ opposite to e. If this is the case, the next visited face is the neighbor τ of σ through e, and we repeat the process, until

Table 2 Exact and rational expressions for the dummy points.

Point	Expression	Rational approximation
V_0	$\left(\frac{2^{3/4}\sqrt{2+\sqrt{2}}}{4}, -\frac{2^{3/4}\sqrt{2-\sqrt{2}}}{4}\right)$	$(97/125, -26/81)$
M_4	$\left(-\sqrt{\sqrt{2}-1}, 0\right)$	$(-9/14, 0)$
M_5	$\left(-\frac{\sqrt{2}\sqrt{\sqrt{2}-1}}{2}, -\frac{\sqrt{2}\sqrt{\sqrt{2}-1}}{2}\right)$	$(-5/11, -5/11)$
M_6	$\left(0, -\sqrt{\sqrt{2}-1}\right)$	$(0, -9/14)$
M_7	$\left(\frac{\sqrt{2}\sqrt{\sqrt{2}-1}}{2}, -\frac{\sqrt{2}\sqrt{\sqrt{2}-1}}{2}\right)$	$(5/11, -5/11)$
P_0	$\left(\frac{2^{1/4}\sqrt{2+\sqrt{2}}}{2\sqrt{2}+2\sqrt{2-\sqrt{2}}}, -\frac{2^{1/4}\sqrt{2-\sqrt{2}}}{2\sqrt{2}+2\sqrt{2-\sqrt{2}}}\right)$	$(1/2, -4/19)$
P_1	$\left(\frac{2^{3/4}\left(\sqrt{2+\sqrt{2}}+\sqrt{2-\sqrt{2}}\right)}{4\sqrt{2}+4\sqrt{2-\sqrt{2}}}, \frac{2^{3/4}\left(\sqrt{2+\sqrt{2}}-\sqrt{2-\sqrt{2}}\right)}{4\sqrt{2}+4\sqrt{2-\sqrt{2}}}\right)$	$(1/2, 4/19)$
P_2	$\left(\frac{2^{1/4}\sqrt{2-\sqrt{2}}}{2\sqrt{2}+2\sqrt{2-\sqrt{2}}}, \frac{2^{1/4}\sqrt{2+\sqrt{2}}}{2\sqrt{2}+2\sqrt{2-\sqrt{2}}}\right)$	$(4/19, 1/2)$
P_3	$\left(\frac{2^{3/4}\left(\sqrt{2-\sqrt{2}}-\sqrt{2+\sqrt{2}}\right)}{4\sqrt{2}+4\sqrt{2-\sqrt{2}}}, \frac{2^{3/4}\left(\sqrt{2+\sqrt{2}}+\sqrt{2-\sqrt{2}}\right)}{4\sqrt{2}+4\sqrt{2-\sqrt{2}}}\right)$	$(-4/19, 1/2)$
P_4	$\left(-\frac{2^{1/4}\sqrt{2+\sqrt{2}}}{2\sqrt{2}+2\sqrt{2-\sqrt{2}}}, \frac{2^{1/4}\sqrt{2-\sqrt{2}}}{2\sqrt{2}+2\sqrt{2-\sqrt{2}}}\right)$	$(-1/2, 4/19)$
P_5	$\left(-\frac{2^{3/4}\left(\sqrt{2+\sqrt{2}}+\sqrt{2-\sqrt{2}}\right)}{4\sqrt{2}+4\sqrt{2-\sqrt{2}}}, \frac{2^{3/4}\left(\sqrt{2-\sqrt{2}}-\sqrt{2+\sqrt{2}}\right)}{4\sqrt{2}+4\sqrt{2-\sqrt{2}}}\right)$	$(-1/2, -4/19)$
P_6	$\left(-\frac{2^{1/4}\sqrt{2-\sqrt{2}}}{2\sqrt{2}+2\sqrt{2-\sqrt{2}}}, -\frac{2^{1/4}\sqrt{2+\sqrt{2}}}{2\sqrt{2}+2\sqrt{2-\sqrt{2}}}\right)$	$(-4/19, -1/2)$
P_7	$\left(\frac{2^{3/4}\left(\sqrt{2+\sqrt{2}}-\sqrt{2-\sqrt{2}}\right)}{4\sqrt{2}+4\sqrt{2-\sqrt{2}}}, -\frac{2^{3/4}\left(\sqrt{2-\sqrt{2}}+\sqrt{2+\sqrt{2}}\right)}{4\sqrt{2}+4\sqrt{2-\sqrt{2}}}\right)$	$(4/19, -1/2)$

- either we find the Euclidean Delaunay face σ_p containing p by visiting only faces that do not cross the border of $\mathcal{D}$; then σ_p is a (canonical) face of $DT_{\mathbb{H}}(\mathcal{GP})$ in conflict with p, and $\nu_{loc} = \mathbb{1}$.
- or, at some point, we visit a (canonical) face $\sigma_{\mathcal{D}}$ included in $\mathcal{D}$ and its (non-canonical) neighbor τ that crosses the border of $\mathcal{D}$. Then the walk continues in non-canonical faces, until we find the Euclidean triangle σ_p containing p. Then ν_{loc} is $\nu_{nbr}(\sigma_{\mathcal{D}}, \tau^c)$ and the canonical face in conflict with p is $\sigma_p^c = \overline{\nu_{loc}}\sigma_p$.

If a (Euclidean) face with edges e_1, e_2, and e_3 is entered through e_1 during the walk, and if none of e_2 and e_3 separates its opposite vertex from p, then the face contains p. So, two orientation tests are enough to conclude that a face contains p (except for the starting face).

The location translation ν_{loc} is also used when looking for all other faces in conflict with p. Starting from σ_p^c and for each face in conflict with p, we recursively examine the translated image under ν_{loc} of each neighbor (obtained with a neighbor translation) that has not yet been visited. We store the set Z^c of canonical faces in conflict with p. Note that Z^c is not necessarily a connected region.

4.3 Insertion

It remains to create the new faces and delete the faces in conflict. The translation ν_{loc} computed in the previous step will again be used. We know that p lies in $\nu_{loc}\sigma_p^c$. We first create a new vertex v_{new} and store p in it.

By construction, the union of all translated faces $\nu_{loc}\nu_{nbr}\left(\sigma_p^c, \tau^c\right)\tau^c$, $\tau^c \in Z^c$ is a topological disk Z in $\mathbb{H}^2$. We identify the sequence of edges E on the border of Z; each edge e is incident to one face in Z and one face that is not in Z. For each face τ^c in Z^c,

we temporarily store the translations $\nu_{loc}\nu_{nbr}(\sigma^c, \tau^c)\,\nu(v_i, \tau^c), i = 0, 1, 2$ directly in its three vertices (not in τ^c, since it will be deleted). Since Z is a topological disk, the result for a given vertex v is independent of the face of Z incident to v that is considered. We store $\mathbb{1}$ in vertex v_{new}.

For each edge $e \in E$, we create a new face τ_e having e as an edge and v_{new} as third vertex. The neighbor of τ_e outside Z^c is the neighbor through e of the face in Z^c incident to e. Two new faces consecutive along E are adjacent. We can now delete all faces in Z.

All that is left to do now is to compute the translations to be stored in the new faces. Let τ_{new} be a newly created face. We retrieve the translations temporarily stored in its vertices v_0, v_1, v_2 and we store them in τ_{new}. Equipped with these translations, τ_{new} is not necessarily canonical. If all translations stored in τ^c_{new} are equal to $\mathbb{1}$, then τ_{new} is contained in $\mathcal{D}$, so it is actually canonical. Otherwise, one of the vertices of τ_{new} is v_{new}; without loss of generality, $v_0 = v_{new}$, and $\nu(v_0, \tau_{new}) = \mathbb{1}$. For $i = 0, 1, 2$ we can easily compute $\Delta_i = \mathrm{index}_{\mathcal{N}}(\nu(v_{\tau_i}^{\mathrm{first\_out}}, \tau_i))$, where τ_i is the image of τ_{new} under $\overline{\nu(v_i, \tau_{new})}$: in each face, $v_{\tau_i}^{\mathrm{first\_out}}$ is the first vertex of τ_i such that $\nu(v_{\tau_i}^{\mathrm{first\_out}}, \tau_i) \neq \mathbb{1}$. Note that we do not actually compute the images of τ_{new}, we only compute translations (as words). We then find the index k for which Δ_k is minimal, and in τ_{new} we store the translations $\overline{\nu(v_k, \tau_{new})}\nu(v_i, \tau_{new})$, $i = 0, 1, 2$. The face τ_{new} has now been canonicalized. Once this is done for all new faces, temporary translations can be removed from the vertices.

5 Algebraic complexity

We follow the so-called *Exact Geometric Computation* paradigm pioneered by Chee Yap [15]. As can be seen in Section 4, the correctness of the combinatorial structure $DT_{\mathcal{M}}(\mathcal{P})$ relies on the exact evaluation of three predicates:

- *SideOfOctagon*, which checks whether an input point lies inside $\mathcal{D}$. This predicate is used as a precondition for the insertion of each point.
- *Orientation*, which checks whether an input point p in $\mathcal{D}$ lies on the right side, the left side, or on an oriented Euclidean segment. This predicate is used when looking for the Euclidean triangle containing an input point.
- *InCircle*, which checks whether an input point p in $\mathcal{D}$ lies inside, outside, or on the boundary of the disk circumscribing an oriented triangle. It is used when looking for all faces in conflict with an input point.

Let the coordinates of a point $p_i \in \mathbb{H}^2$ be denoted as x_i and y_i. The last two predicates can be expressed as signs of determinants:

$$\mathit{Orientation}(p_1, p_2, p_3) = \mathrm{sign}\begin{vmatrix} x_1 & y_1 & 1 \\ x_2 & y_2 & 1 \\ x_3 & y_3 & 1 \end{vmatrix}, \quad \mathit{InCircle}(p_1, p_2, p_3, p_4) = \mathrm{sign}\begin{vmatrix} x_1 & y_1 & x_1^2+y_1^2 & 1 \\ x_2 & y_2 & x_2^2+y_2^2 & 1 \\ x_3 & y_3 & x_3^2+y_3^2 & 1 \\ x_4 & y_4 & x_4^2+y_4^2 & 1 \end{vmatrix}. \quad (4)$$

We assume that all input points (which lie in $\mathcal{D}$) have rational coordinates (recall that this holds for the initial dummy points, see Section 4.1). So, in the above determinants, at least one point (x_i, y_i) is rational. However, the points against which the predicates are testing the new input point are vertices of some face of $DT_{\mathbb{H}}(\mathcal{GP})$ contained in $\mathcal{U}_{\mathcal{D}}$, so they are images of some input points by translations in $\mathcal{N} \cup \{\mathbb{1}\}$. Therefore, the evaluation of the two predicates (4) boils down to determining the sign, considered as an element of $\{-1, 0, 1\}$ of polynomial expressions in rational variables, whose coefficients are lying in some extension field of the rationals, as made precise below.

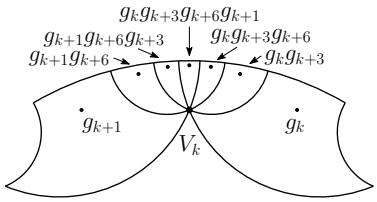

Figure 8 Elements of $\mathcal{N}$ around a vertex V_k.

Table 3 Matrices of translations around a vertex V_k, k. Recall that $\xi = \sqrt{1+\sqrt{2}}$.

$$
\begin{aligned}
g_k &= \begin{bmatrix} 1+\sqrt{2} & e^{ik\pi/4}\sqrt{2}\xi \\ e^{-ik\pi/4}\sqrt{2}\xi & 1+\sqrt{2} \end{bmatrix} \quad \text{(see (2))} \\
g_k g_{k+3} &= \begin{bmatrix} \left(1+\sqrt{2}\right)\left(1-i\sqrt{2}\right) & e^{ik\pi/4}\left(1+i\left(1+\sqrt{2}\right)\right)\xi \\ e^{-ik\pi/4}\left(1-i\left(1+\sqrt{2}\right)\right)\xi & \left(1+\sqrt{2}\right)\left(1+i\sqrt{2}\right) \end{bmatrix} \\
g_k g_{k+3} g_{k+6} &= \begin{bmatrix} -\left(1+\sqrt{2}\right)(1+2i) & e^{ik\pi/4}\left(1+\sqrt{2}\right)(-1+i)\,\xi \\ -e^{-ik\pi/4}\left(1+\sqrt{2}\right)(1+i)\,\xi & -\left(1+\sqrt{2}\right)(1-2i) \end{bmatrix} \\
g_k g_{k+3} g_{k+6} g_{k+1} &= \begin{bmatrix} -2\sqrt{2}-3 & -e^{ik\pi/4}\left(2+\sqrt{2}+i\sqrt{2}\right)\xi \\ -e^{-ik\pi/4}\left(2+\sqrt{2}-i\sqrt{2}\right)\xi & -2\sqrt{2}-3 \end{bmatrix} \\
g_{k+1} g_{k+6} g_{k+3} &= \begin{bmatrix} \left(1+\sqrt{2}\right)(-1+2i) & -e^{ik\pi/4}\left(2+\sqrt{2}\right)i\xi \\ e^{-ik\pi/4}\left(2+\sqrt{2}\right)i\xi & -\left(1+\sqrt{2}\right)(1+2i) \end{bmatrix} \\
g_{k+1} g_{k+6} &= \begin{bmatrix} \left(1+\sqrt{2}\right)+\left(2+\sqrt{2}\right)i & e^{ik\pi/4}\left(1+\sqrt{2}-i\right)\xi \\ e^{-ik\pi/4}\left(1+\sqrt{2}+i\right)\xi & \left(1+\sqrt{2}\right)-\left(2+\sqrt{2}\right)i \end{bmatrix} \\
g_{k+1} &= \begin{bmatrix} 1+\sqrt{2} & e^{ik\pi/4}(1+i)\,\xi \\ e^{-ik\pi/4}(1-i)\,\xi & 1+\sqrt{2} \end{bmatrix}
\end{aligned}
$$

The evaluation of the degree of the predicates requires to perform a case analysis on the different possible positions of the faces in $\mathcal{U}_\mathcal{D}$, i.e., on the possible translations of $\mathcal{N}$ that can be involved in each predicate. The following property shows how we can take symmetries of $\mathcal{D}$ into account to reduce the number of possible cases.

▶ **Lemma 5.** *Let $\boldsymbol{\sigma}$ be a face in $\mathrm{DT}_\mathcal{M}(\mathcal{P})$. Then, for any edge uv of its canonical representative σ^c, such that $\nu(u,\sigma^c) \neq \mathbb{1}$ and $\nu(v,\sigma^c) \neq \mathbb{1}$,*

$$\left|\mathrm{index}_\mathcal{N}(\nu(u,\sigma^c)) - \mathrm{index}_\mathcal{N}(\nu(v,\sigma^c))\right| \leq 7.$$

Proof. We can assume that $\sigma^c \not\subset \mathcal{D}$, otherwise all its three translations are equal to $\mathbb{1}$. Reusing the proof of of Lemma 1 and the notation therein, we see that σ^c is either contained in the disk bounded by B_k, or in the disk bounded by C_k, for some $k \in \{0,\ldots,7\}$. So σ^c can only intersect $\mathcal{D}$ and the seven octagons around some V_k. The result follows. ◀

Figure 8 shows the possible translations of $\mathcal{N}$ involved in a given canonical representative, for some $k \in \{0,\ldots,7\}$. Their matrices are given in Table 3.

For k even, the sine and cosine of $\pm k\pi/4$ have values in $\{-1,0,1\}$, while for k odd they are both equal to $\pm\sqrt{2}/2$. Therefore, up to sign, the above matrices are divided into two "classes". Due to the symmetries of $\mathcal{D}_O$, we actually only need to examine one case in each class, therefore we can focus on the two cases $k=0$ and $k=1$.

▶ **Proposition 1.** *All predicates can be evaluated by determining the sign of rational polynomial expressions of total degree at most 72 in the coordinates of input points.*

Proof. We examine the complexity of the *Orientation* predicate and refer the reader to the appendix for other ones. As mentioned above, at least one point is inside $\mathcal{D}$. Without loss of generality, we assume that $p_3 \in \mathcal{D}$. Let us consider the possible cases for the other two points.

- All three points are inside $\mathcal{D}$. In this case, all the arguments of the predicate are rational, so from (4) we get a polynomial with rational coefficients of total degree 2 in the coordinates of the input points.
- Point p_2 is also in $\mathcal{D}$, and p_1 is outside $\mathcal{D}$. In this case, p_1 can be the image of an input point by 14 possible different translations in $\mathcal{N}$ (seven around V_0 and seven around V_1).
- Only p_3 is inside $\mathcal{D}$. In this case, both p_1 and p_2 can be images of input points under the translations around V_0 and V_1. Of course, we avoid redundancies: if we examine the case $\mathit{Orientation}(g_i p'_1, g_j p'_2, p_3)$, $p'_i, p'_j \in \mathcal{D}$, we do not examine the case $\mathit{Orientation}(g_j p'_1, g_i p'_2, p_3)$ since it would have the same degree. This amounts to 56 cases in total – 28 cases around V_0 and another 28 around V_1.

We have found with Maple that in all cases, the expressions produced by (4) have denominators that are strictly positive and numerators that can be brought into the form

$$\left(A\sqrt{2}+B\right)\xi + C\sqrt{2} + D, \qquad \xi = \sqrt{1+\sqrt{2}}, \tag{5}$$

where A, B, C, D are rational polynomial expressions in the input coordinates. Moreover, the maximum total degree of A, B, C, D is 5. By squaring twice (to eliminate square roots coming from ξ), we get a rational polynomial of degree 20 in rational variables. ◀

The degree itself, as well as the high number of cases (in spite of the reduction) that would need to be considered, show that giving a complete implementation for all polynomial expressions involved in the predicates is hardly feasible. Therefore, we use the type `CORE::Expr` [16] included in the CGAL distribution to compute the coordinates of translated points and directly evaluate the signs of determinants (4). This number type guarantees that predicates are exact.

Our implementation handles degeneracies using symbolic perturbations [8]. Note that there are no degeneracies in the initial triangulation $DT_{\mathcal{M}}(\mathcal{Q}')$.

6 Experimental results

Experiments are run on a MacBook Pro with CPU Intel Core i5 @ 2.9 GHz, 16 GB RAM @ 1867 MHz running the `master` version of CGAL from GitHub, compiled in release mode with clang-700.1.81. We insert random points uniformly distributed with respect to the hyperbolic metric in $\mathcal{D}$. As mentioned in introduction, dummy points are removed after the insertion of new points. Averaged over 10 executions, the running time is 34 seconds for one million points. This is slower than the computation of 2D Euclidean Delaunay triangulations with CGAL, which takes around 12 seconds on average for the same sets of points, using `CORE::Expr` as number type (and about one second with `double` number type). This is due in particular to the much higher arithmetic demand in our case (Proposition 1), as confirmed by preliminary profiling, which shows that almost two thirds of the running time is spent in computations of predicates. For the insertion of one million points, only 0.76% calls to predicates involve images of rational points under translations in $\mathcal{N}$, but these calls account for 36% of the total time spent in predicates.

We have also executed tests in which we insert random points in the triangulation and progressively remove dummy points whenever doing so does not violate condition (1). Over

100 executions, all dummy points are removed with the insertion of at least 17 and at most 72 random points.

Due to lack of space, pictures showing some Delaunay triangulations are shown on the web page.

References

1. 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.
2. 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.
3. A. Bowyer. Computing Dirichlet tessellations. *The Computer Journal*, 24(2):162–166, 1981. doi:10.1093/comjnl/24.2.162.
4. Manuel Caroli and Monique Teillaud. 3D periodic triangulations. In *CGAL User and Reference Manual*. CGAL Editorial Board, 3.5 (and further) edition, 2009-. URL: http://doc.cgal.org/latest/Manual/packages.html#PkgPeriodic3Triangulation3Summary.
5. Manuel Caroli and Monique Teillaud. Delaunay triangulations of closed Euclidean d-orbifolds. *Discrete & Computational Geometry*, 55(4):827–853, 2016. doi:10.1007/s00454-016-9782-6.
6. M. Dehn. Transformation der Kurven auf zweiseitigen Flächen. *Mathematische Annalen*, 72(3):413–421, 1912. doi:10.1007/BF01456725.
7. 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.
8. Olivier Devillers and Monique Teillaud. Perturbations for Delaunay and weighted Delaunay 3D Triangulations. *Computational Geometry: Theory and Applications*, 44:160–168, 2011. doi:10.1016/j.comgeo.2010.09.010.
9. Nikolai P. Dolbilin and Daniel H. Huson. Periodic Delone tilings. *Periodica Mathematica Hungarica*, 34:1-2:57–64, 1997.
10. The GAP Group. *GAP – Groups, Algorithms, and Programming, Version 4.8.6*, 2016. URL: http://www.gap-system.org.
11. 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.
12. Nico Kruithof. 2D periodic triangulations. In *CGAL User and Reference Manual*. CGAL Editorial Board, 4.3 (and further) edition, 2013-. URL: http://doc.cgal.org/latest/Manual/packages.html#PkgPeriodic2Triangulation2Summary.
13. The Magma Development Team. *Magma Computational Algebra System*. URL: http://magma.maths.usyd.edu.au/magma/.
14. John Joseph O'Connor and Edmund Frederick Robertson. The MacTutor History of Mathematics archive, 2003. URL: http://www-history.mcs.st-andrews.ac.uk/HistTopics/Word_problems.html.
15. 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.
16. Chee Yap *et al.* The CORE Library Project. URL: http://cs.nyu.edu/exact/core_pages/intro.html.