

Delaunay triangulations on orientable surfaces of low genus *

Mikhail Bogdanov¹, Monique Teillaud², and Gert Vegter³

1 Inria, Centre de recherche Sophia Antipolis - Méditerranée, France

2 Inria, Centre de recherche Nancy - Grand Est. CNRS, Loria. Université de Lorraine. France

3 Johann Bernoulli Institute for Mathematics and Computer Science, University of Groningen, The Netherlands

Abstract

Earlier work on Delaunay triangulation of point sets on the 2D flat torus, which is locally isometric to the Euclidean plane, was based on lifting the point set to a locally isometric 9-sheeted covering space of the torus [28, 21, 13, 12]. Under mild conditions the Delaunay triangulation of the lifted point set, consisting of 9 copies of the input set, projects to the Delaunay triangulation of the input set.

We improve and generalize this work. First we present a new construction based on an 8-sheeted covering space, which shows that eight copies suffice for the standard flat torus. Then we generalize this construction to the context of compact orientable surfaces of higher genus, which are locally isometric to the hyperbolic plane. We investigate more thoroughly the Bolza surface, homeomorphic to a sphere with two handles, both because it is the hyperbolic surface with lowest genus, and because triangulations on the Bolza surface have applications in various fields such as neuromathematics and cosmological models [32, 4, 16]. While the general properties (existence results of appropriate covering spaces) show similarities with the results for the flat case, explicit constructions and their proofs are much more complex, even in the case of the apparently simple Bolza surface. One of the main reasons is the fact that two hyperbolic translations do not commute in general.

To the best of our knowledge, the results in this paper are the first ones of this kind. The interest of our contribution lies not only in the results, but most of all in the construction of covering spaces itself and the study of their properties.

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

Keywords and phrases covering spaces, hyperbolic surfaces, finitely presented groups, Fuchsian groups, systole

Digital Object Identifier 10.4230/LIPIcs.SoCG.2016.20

1 Introduction

Delaunay triangulations of the flat cubic torus [39], i.e., the flat torus for which the fundamental domain is a cube, are used in many fields of science, ranging from the infinitely small, with “nuclear pasta” [33] to the infinitely large, with cosmology [35, 36, 40, 24] (more references can be found in [13, 14, 15, 11]). They are needed either because the objects studied are periodic in space, or because they allow for modeling periodic boundary conditions, e.g., in fluid mechanics. Periodic Delaunay triangulations in the hyperbolic plane are also used

* Work partially supported by the Inria OrbiCG associate team <http://www.loria.fr/~teillaud/collab/OrbiCG>



© Mikhail Bogdanov and Monique Teillaud and Gert Vegter;
licensed under Creative Commons License CC-BY

32nd International Symposium on Computational Geometry (SoCG 2016).

Editors: Sándor Fekete and Anna Lubiw; Article No. 20; pp. 20:1–20:15

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

in diverse fields, such as physics, solid modeling, cosmological models, neuromathematics [32, 31, 4, 16]; they can also be seen as triangulations on hyperbolic surfaces.

A closed orientable flat (resp. hyperbolic) surface M is isometric to a quotient \mathbb{E}^2/Γ (resp. \mathbb{H}^2/Γ), where Γ is a discrete group of isometries acting on the Euclidean plane \mathbb{E}^2 (resp. hyperbolic plane \mathbb{H}^2), with properties made precise later in the paper.

As in [15], we stick to the standard definition of a *triangulation* [30, Section 9.1, condition (3)] of a topological space \mathbb{Y} as a geometric simplicial complex K such that $K = \bigcup_{\sigma \in K} \sigma$ is homeomorphic to \mathbb{Y} . A triangulation defined by a point set \mathcal{P} is a triangulation whose set of vertices is \mathcal{P} . The Delaunay triangulation of a point set \mathcal{P} in \mathbb{E}^2 (resp. \mathbb{H}^2) is a triangulation of the convex hull (resp. hyperbolic convex hull) of \mathcal{P} that can be defined uniquely even in degenerate cases [20].

The Delaunay triangulation of points on a 2D flat torus was usually obtained from a triangulation of the convex hull of nine periodic copies of the points in the plane \mathbb{E}^2 , laid in a 3x3 pattern [28, 22] (or 27 copies in a 3x3x3 pattern in the 3D case [21]). A more recent algorithm uses covering spaces instead, and computes the Delaunay triangulation directly in the torus whenever possible [13], thus providing all adjacency relations between simplices. It led to CGAL packages [27, 12], which have already been used in various fields.¹ In that work, the Delaunay triangulation of a surface M is defined as follows:

► **Definition 1.** The *Delaunay triangulation* $DT_M(\mathcal{P})$ of M defined by a finite point set \mathcal{P} in \mathbb{E}^2 (resp. \mathbb{H}^2) is the projection onto M of the Delaunay triangulation $DT_{\mathbb{E}}(\Gamma\mathcal{P})$ (resp. $DT_{\mathbb{H}}(\Gamma\mathcal{P})$) of the infinite point set $\Gamma\mathcal{P}$, if it is a triangulation.

Note that the fact that Γ is discrete guarantees that the triangulation of $\Gamma\mathcal{P}$ in \mathbb{E}^2 (resp. \mathbb{H}^2) is locally finite [15]. We elaborate a bit on the projection here and check that it is well-defined. Let the projection be denoted as $\pi : \mathbb{E}^2$ (resp. \mathbb{H}^2) $\rightarrow M$. By definition, for any given point $p \in \mathcal{P}$, all elements of the orbit $\Gamma p = \{g(p) \mid g \in \Gamma\}$ project by π to the same point $\pi(p)$ of M . We can assume that no two points of \mathcal{P} lie on the same Γ -orbit: otherwise we can always choose a proper subset $\mathcal{P}' \subset \mathcal{P}$, so that $\Gamma\mathcal{P}' = \Gamma\mathcal{P}$. The Delaunay triangulation $DT_{\mathbb{E}}(\Gamma\mathcal{P})$ (resp. $DT_{\mathbb{H}}(\Gamma\mathcal{P})$) is a partition of \mathbb{E}^2 (resp. \mathbb{H}^2) by vertices, edges, and triangular faces. Each $g \in \Gamma$ is an isometry, so the empty disk circumscribing a Delaunay triangle (p, q, r) , $p, q, r \in \mathbb{E}^2$ (resp. \mathbb{H}^2) is transformed by g into the disk circumscribing triangle $(g(p), g(q), g(r))$, which is also empty. If the set $\Gamma\mathcal{P}$ contains degeneracies, the Delaunay triangulation of a polygon formed by cocircular points $P = (p_0, p_1, \dots, p_k = p_0)$, $k \geq 4$ is not uniquely defined. For any $g \in \Gamma$, the polygon $g(P) = (g(p_0), g(p_1), \dots, g(p_k) = g(p_0))$ is also formed by cocircular points. Once a triangulation of P is chosen, the polygon $g(P)$ can be triangulated by the image by g of the triangulation of P . Then, these two triangulations project under π to the same set of triangles on M . Since π is surjective, the projection $\pi(DT_{\mathbb{E}}(\Gamma\mathcal{P}))$ (resp. $\pi(DT_{\mathbb{H}}(\Gamma\mathcal{P}))$) forms a partition of M by vertices, edges, and triangular faces.² This discussion allows us to ignore degeneracies in the sequel.

The algorithm follows the classic incremental algorithm in \mathbb{E}^d [8]: for each new point p , the conflicting simplices (i.e., simplices whose circumscribing balls contain p) are deleted, and the region that they formed is then triangulated with new simplices with apex p . That algorithm relies on the fact that the union of the set of conflicting simplices is always a topological ball, which subsumes that the graph of edges of the triangulation has no cycle

¹ See <http://www.cgal.org/projects.html> for some identified users.

² In the case of the flat torus, this can automatically be achieved by means of a symbolic perturbation scheme [20] that ensures that the triangulation in the universal covering space is uniquely defined and invariant under the action of the group.

of length one (*1-cycles*) or two (*2-cycles*). However, the graph of edges of the projection of the Delaunay triangulation on a manifold can have non-trivial cycles: for instance, in the extreme case when the input consists of only one point p , then the set Γp is the orbit of p under Γ , so all Delaunay vertices project to the same point, i.e., Delaunay edges project to a small set of 1-cycles on the manifold, all with the same vertex. The existence of such 1- or 2-cycles can be checked using a geometric criterion. The algorithm computes in some finite-sheeted covering manifold (i.e., roughly speaking, using “copies” of the input points) satisfying the criterion: a 9- (resp 27-) sheeted covering manifold is sufficient for the 2D (resp. 3D) flat torus. When sufficiently many points have been inserted so that the criterion is satisfied on the initial manifold, the computation continues in this manifold, i.e., without using more copies of points.

To the best of our knowledge, there are no known similar results on hyperbolic surfaces. Intuition is challenged there, in particular because the group Γ is non-Abelian in general.

In this paper, we first propose a construction of a family of 2^k sheeted-covering spaces, $k > 0$, of the flat cubic torus (Section 2). We show how this construction allows us to improve the computation of a Delaunay triangulation of that torus by reducing the number of sheets of the covering spaces.

For any closed hyperbolic surface M , we prove in Section 3 the existence of a finite-sheeted covering space \widetilde{M} of M so that the projection of the Delaunay triangulation in \mathbb{H}^2 onto \widetilde{M} has no 1- or 2-cycle (Proposition 9). We also present a geometric criterion to checks that 1- or 2-cycles don't exist (Proposition 8). When such a covering space \widetilde{M} can be constructed, then the algorithm previously proposed for closed Euclidean manifolds [15] naturally extends to a set of points \mathcal{P} on the hyperbolic surface M .

The Bolza surface \mathcal{M} is the simplest possible surface after the 2D Euclidean torus: it is homeomorphic to a double torus. The main contribution of this paper (Section 4) is a construction of some finite-sheeted covering spaces of \mathcal{M} , which allows us to concretely exhibit a 4-sheeted covering space in which the projection of a Delaunay triangulation never has a 1-cycle, and a 128-sheeted covering space in which it has no 2-cycle. These results mostly are of theoretical interest for the incremental algorithm above, due to the large number of points that would need to be considered for the computations; we propose a more practical algorithm in the last section.

2 Some finitely-sheeted covering spaces of the flat torus

Let a, b denote two orthogonal translations of the same length. We denote as \mathcal{G}_1 the group $\langle a, b \rangle$ generated by a and b . Since \mathcal{G}_1 is Abelian, any $x \in \mathcal{G}_1$ can be uniquely written as $x = a^\alpha b^\beta$, where $\alpha, \beta \in \mathbb{Z}$. The *length* $\lambda_1(x)$ of x can thus be defined as the sum $|\alpha| + |\beta|$. We are going to study some covering spaces of the flat torus $\mathbb{T} = \mathbb{T}_1^2 = \mathbb{E}^2/\mathcal{G}_1$.

The minimal distance by which \mathcal{G}_1 moves a point of \mathbb{E}^2 is $\delta(\mathcal{G}_1) = \|a\| = \|b\|$, where $\|\cdot\|$ denotes the Euclidean norm. Any closed square of edge length $l = \delta(\mathcal{G}_1)$ whose edges are parallel to a and b is a fundamental domain of \mathcal{G}_1 .

In the sequel, the inverse x^{-1} of an element x of \mathcal{G}_1 will alternatively be denoted by \bar{x} . All subgroups are normal, since \mathcal{G}_1 is Abelian. Let us consider the subgroup \mathcal{G}_2 consisting of elements of \mathcal{G}_1 of even length. It is easy to see that $\mathcal{G}_2 = \langle \bar{a}b, ab \rangle$.

► **Lemma 1.** \mathcal{G}_2 is a subgroup of index 2 in \mathcal{G}_1 .

Proof. Let $\varphi : \mathcal{G}_1 \rightarrow \mathbb{Z}_2$ be the group homomorphism defined by $\varphi(x) = \lambda_1(x) \pmod 2$. The subgroup \mathcal{G}_2 is the kernel of φ . According to the First Isomorphism Theorem (see, e.g., [2])

$\ker \varphi$ is a normal subgroup of \mathcal{G}_1 , and $\mathcal{G}_1/\ker \varphi \cong \varphi(\mathcal{G}_1)$. Therefore, $\mathcal{G}_1/\mathcal{G}_2 \cong \mathbb{Z}_2$. ◀

For $k > 1$, if $\mathcal{G}_{2^{k-1}}$ is a subgroup of index 2^{k-1} of \mathcal{G}_1 generated by two elements, i.e., $\mathcal{G}_{2^{k-1}} = \langle g, h \rangle$, $g, h \in \mathcal{G}_1$, we construct the subgroup \mathcal{G}_{2^k} of $\mathcal{G}_{2^{k-1}}$ as follows: $\mathcal{G}_{2^k} = \langle \bar{g}h, gh \rangle$. \mathcal{G}_{2^k} consists of elements of even length of $\mathcal{G}_{2^{k-1}}$. The proof of the following lemma is similar to that of Lemma 1.

► **Lemma 2.** For any $k > 0$, \mathcal{G}_{2^k} is a subgroup of index 2^k in \mathcal{G}_1 .

By construction, for any $k \geq 0$, the generators of \mathcal{G}_{2^k} are two translations whose vectors a_k and b_k are orthogonal and of equal length $\delta(\mathcal{G}_{2^k}) = \sqrt{2}^k l$. The orbits of \mathcal{G}_{2^k} are isomorphic to \mathbb{Z}^2 . The half-open square $\{t_a a_k + t_b b_k \mid t_a, t_b \in [-\frac{1}{2}, \frac{1}{2})\}$ contains exactly one representative of each orbit. We call it the *original domain* of \mathcal{G}_{2^k} and denote it as \mathcal{D}_{2^k} . See Figure 1; for simplicity, x denotes both an element x of \mathcal{G}_1 and the image xO of the origin O by x . The closure of the original domain \mathcal{D}_{2^k} is a fundamental domain of \mathcal{G}_{2^k} ; it is also the cell of O in the Voronoi diagram of its orbit $\mathcal{G}_{2^k}O$.




Figure 1 The original domains \mathcal{D}_1 and \mathcal{D}_2 of \mathcal{G}_1 and \mathcal{G}_2 respectively.

For any $k > 0$, $\mathbb{T}_{2^k}^2 = \mathbb{E}^2/\mathcal{G}_{2^k}$ is a flat torus, with the corresponding projection map $\pi_{2^k} : \mathbb{E}^2 \rightarrow \mathbb{T}_{2^k}^2$. The torus $\mathbb{T}_{2^k}^2$ is a covering space of \mathbb{T}^2 together with the covering map $\rho_k = \pi_1 \circ \pi_{2^k}^{-1}$. Since the index of \mathcal{G}_{2^k} in \mathcal{G}_1 is 2^k , we get the following result:

► **Proposition 3.** For any $k > 0$, $\mathbb{T}_{2^k}^2$ is a 2^k -sheeted covering space of \mathbb{T}^2 .

Delaunay triangulations via 2^k -sheeted covering spaces, $k > 0$

As shown in [13] the Delaunay triangulation $DT_{\mathbb{T}}(\mathcal{P})$ as defined in introduction (Def. 1) does not exist for each set of points $\mathcal{P} \in \mathbb{E}^2$, because $\pi(DT(\mathcal{G}\mathcal{P}))$ can have 1- or 2-cycles. However, there are always some covering spaces of the torus in which the Delaunay triangulation exists and can be computed.

Let $\Delta(\mathcal{S})$ denote the diameter of the largest disk in \mathbb{E}^2 that does not contain any point of a set \mathcal{S} in its interior. Note that for any $p \in \mathbb{E}^2$, $\Delta(\mathcal{G}_{2^k}p) = \sqrt{2}^{k+1}l$.

► **Proposition 4** ([13]). If $\Delta(\mathcal{G}\mathcal{P}) < \frac{\delta(\mathcal{G})}{2}$, then for any finite $\mathcal{P}' \supseteq \mathcal{P}$, $\pi(DT(\mathcal{G}\mathcal{P}'))$ is a triangulation of \mathbb{T}^2 .

► **Proposition 5.** If $\Delta(\mathcal{G}_1\mathcal{P}) < \frac{1}{2}\delta(\mathcal{G}_{2^k})$, then $\pi_{2^k}(DT(\mathcal{G}_1\mathcal{P} \cup \mathcal{G}_{2^k}\mathcal{T}))$ is a triangulation of $\mathbb{T}_{2^k}^2$ for any finite point set \mathcal{T} in \mathbb{E}^2 . In particular, then $\pi_{2^k}(DT(\mathcal{G}_1\mathcal{P}'))$ is a triangulation of $\mathbb{T}_{2^k}^2$ for any finite $\mathcal{P}' \supseteq \mathcal{P}$, $\mathcal{P}' \in \mathbb{E}^2$.

Proof. Let \mathcal{P}_{2^k} denote the set $\mathcal{G}_1\mathcal{P} \cap \mathcal{D}_{2^k}$. By Proposition 4, if $\Delta(\mathcal{G}_{2^k}\mathcal{P}_{2^k}) < \frac{1}{2}\delta(\mathcal{G}_{2^k})$, then $\pi_{2^k}(DT(\mathcal{G}_{2^k}\mathcal{P}_{2^k} \cup \mathcal{G}_{2^k}\mathcal{T}))$ is a triangulation. Let us note that $\mathcal{G}_{2^k}\mathcal{P}_{2^k} = \mathcal{G}_1\mathcal{P}$.


For $\mathcal{P}' \supseteq \mathcal{P}$, $\mathcal{G}_1\mathcal{P}' = \mathcal{G}_1\mathcal{P} \cup \mathcal{G}_1(\mathcal{P}' \setminus \mathcal{P}) = \mathcal{G}_1\mathcal{P} \cup \mathcal{G}_{2^k}(\mathcal{G}_1(\mathcal{P}' \setminus \mathcal{P}) \cap \mathcal{D}_{2^k})$. It remains to take $\mathcal{T} = \mathcal{G}((\mathcal{P}' \setminus \mathcal{P}) \cap \mathcal{D}_{2^k})$. ◀

By Proposition 5, if the maximum empty disk diameter $\Delta(\mathcal{G}_1\mathcal{P})$ is smaller than $\frac{1}{2}\delta(\mathcal{G}_1) = l$, then $\pi_4(DT(\mathcal{G}_1\mathcal{P} \cup \mathcal{G}_4\mathcal{T}))$ is a simplicial complex for any finite \mathcal{T} in \mathbb{E}^2 . If it is smaller than $\frac{1}{2}\delta(\mathcal{G}_2) = \frac{\sqrt{2}}{2}l$, then $\pi_2(DT(\mathcal{G}_1\mathcal{P} \cup \mathcal{G}_2\mathcal{T}))$ is a simplicial complex.

Note that, for $k > 3$, $\frac{1}{2}\delta(\mathcal{G}_{2^k}) = \frac{1}{2}\sqrt{2}^k l > \sqrt{2}l = \Delta(\mathcal{G}_1p)$, for any $p \in \mathcal{P}$. Therefore, by Proposition 5, $\pi_{2^k}(DT(\mathcal{G}_1\mathcal{P}))$ is a simplicial complex. Let us now consider the case $k = 3$.

► **Corollary 6.** For any finite set $\mathcal{P} \subset \mathcal{D}_1$ of $n > 1$ points, $\pi_8(DT(\mathcal{G}_1\mathcal{P}))$ is a simplicial complex.

Proof. The maximum empty disk diameter $\Delta(\mathcal{G}_1p)$, for any $p \in \mathcal{P}$, is $\sqrt{2}l$, that is equal to $\frac{1}{2}\delta(\mathcal{G}_8)$ (See Figure 2). We are going to show that $\Delta(\mathcal{G}_1\mathcal{P})$ is strictly less than $\sqrt{2}l$. Let us suppose that B is a disk with diameter $\sqrt{2}l$ centered at some point $c \in \mathbb{E}^2$ and containing no point of $\mathcal{G}_1\mathcal{P}$ in its interior (See Figure 2). B contains $\mathcal{D}_1(c)$, the original domain centered at c . Since $\mathcal{D}_1(c)$ is a half-open square, there is only one point of $\mathcal{D}_1(c)$ on the boundary of B . For p and $q \neq p, q \in \mathcal{P}$, there is a representative of \mathcal{G}_1p and a representative of \mathcal{G}_1q in $\mathcal{D}_1(c)$. At least one of these representatives lies in the interior of B . This contradicts that the interior of B is empty. If we add more points, the diameter of the largest empty disk cannot become larger.



■ **Figure 2** The original domain \mathcal{D}_8 . Maximum empty disk with diameter $\Delta(\mathcal{G}_1O)$.

By Proposition 5, $\pi_8(DT(\mathcal{G}_1\mathcal{P}))$ is a simplicial complex. ◀

Algorithm. For a finite set of points $\mathcal{P} \subset \mathcal{D}_1$, the algorithm of [13], inspired by the standard incremental algorithm [8], computes $DT_{\mathbb{T}}(\mathcal{G}_1\mathcal{P})$ via a 9-sheeted covering space. As soon as the condition of Proposition 4 is fulfilled upon insertion of a point, it switches to computing in \mathbb{T}^2 . Assuming that \mathcal{P} contains at least two points, we modify this algorithm and use the covering spaces $\mathbb{T}_8^2, \mathbb{T}_4^2$, and \mathbb{T}_2^2 using Proposition 5 and Corollary 6.

- The algorithm starts by precomputing $\pi_8(DT(\mathcal{G}_1\{p, q\}))$, for any $\{p, q\} \subset \mathcal{P}$. Then it adds points one by one in the Delaunay triangulation of the 8-sheeted covering space \mathbb{T}_8^2 .
- If after insertion of some set of points $\mathcal{S} \subseteq \mathcal{P}$, the maximum empty disk diameter becomes less than l , then $\pi_4(DT(\mathcal{G}_1\mathcal{S}'))$ is guaranteed to be a triangulation for any finite $\mathcal{S}' \supseteq \mathcal{S}$. So, we can discard all redundant periodic copies of simplices of $\pi_8(DT(\mathcal{G}_1\mathcal{S}))$ and switch to incrementally computing $\pi_4(DT(\mathcal{G}_1\mathcal{S}'))$ in the 4-sheeted covering space \mathbb{T}_4^2 for $\mathcal{S} \subseteq \mathcal{S}' \subseteq \mathcal{P}$.
- Similarly, if after some more insertions the maximum empty disk diameter becomes less than $\frac{\sqrt{2}}{2}l$, then we switch to incrementally computing triangulations in the 2-sheeted covering space \mathbb{T}_2^2 .
- If it becomes less than $\frac{1}{2}l$, then we switch to computing $\pi(DT(\mathcal{G}_1\mathcal{P}))$ in \mathbb{T}^2 .


For any $\mathcal{S} \subset \mathcal{P}$, $\pi_{2^k}(DT(\mathcal{G}_1\mathcal{S}))$ contains 2^k periodic copies of each point of \mathcal{S} . Hence, using covering spaces whose number of sheets is as small as possible improves the efficiency of the algorithm, even though the asymptotic complexity is unchanged: it stays randomized worst-case time and space optimal [13].

3 Delaunay triangulations on hyperbolic closed surfaces

3.1 Background

As the *Poincaré disk* is a conformal model of the hyperbolic plane (i.e., it preserves angles), it is widely used in applications [25, 32, 31, 4, 16]. In this model [10] [5, Chapter 19] the hyperbolic plane \mathbb{H}^2 is represented as the open unit disk \mathcal{B} of the Euclidean plane \mathbb{E}^2 . The points on the boundary of \mathcal{B} are the *points at infinity*; their set is denoted as $\mathcal{H}_\infty = \{p \in \mathbb{E}^2 : \|p\|=1\}$. Hyperbolic lines, or geodesics, are represented either as arcs of Euclidean circles orthogonal to \mathcal{H}_∞ or as diameters of \mathcal{B} ; 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}_∞ (Fig. 3(left)). The group of isometries of \mathbb{H}^2



■ **Figure 3** (left) Lines and circles. (right) Translation g of axis X_g .

is denoted as $\text{Isom}(\mathbb{H}^2)$, and the group of orientation-preserving isometries is denoted as $\text{Isom}^+(\mathbb{H}^2)$. Identifying \mathbb{E}^2 with the complex plane \mathbb{C} , each $g \in \text{Isom}^+(\mathbb{H}^2)$ is of the form $g(z) = \frac{\alpha z + \beta}{\beta z + \alpha}$, with matrix $m_g = \begin{bmatrix} \alpha & \beta \\ \beta & \alpha \end{bmatrix}$, $\alpha, \beta \in \mathbb{C}$, $|\alpha|^2 - |\beta|^2 = 1$. There are three types of orientation-preserving isometries of \mathbb{H}^2 ; in this paper we only consider *hyperbolic* isometries, also called *translations*. A translation g fixes no point of \mathcal{B} and fixes two points on \mathcal{H}_∞ . The *axis* X_g of g is the hyperbolic line ending at the two fixed points (Fig. 3(right)). Any point is translated by g along a curve at constant distance to the axis. All such curves have the fixed points of g as infinite points. The distance by which g moves all points of X_g is the same, it is called the *translation length* of g , $l(g) = 2 \cdot \text{acosh}(\frac{1}{2} \text{trace}(m_g))$. If $p \notin X_g$, then the distance by which g moves p is strictly greater than $l(g)$. It is worth noticing that

— *two translations do not commute in general.* —

A *Fuchsian group* is a discrete subgroup of $\text{Isom}^+(\mathbb{H}^2)$ [26]. A *hyperbolic surface* is a connected 2-dimensional manifold, such that every point has a neighborhood isometric to a disk of \mathbb{H}^2 . In this paper all hyperbolic surfaces are *closed* surfaces, i.e., they are compact.

— *A closed orientable hyperbolic surface M is of the form \mathbb{H}^2/Γ ,* —

— *where Γ is a Fuchsian group only containing hyperbolic translations (and the identity)* — (See [30, 37, 38, 41].) The projection map $\pi : \mathbb{H}^2 \rightarrow M = \mathbb{H}^2/\Gamma$ is a local isometry and a covering projection. The *Dirichlet region* $\mathcal{D}_p(\Gamma)$ for Γ centered at p is defined as the closure $Cl(V_p(\Gamma p))$ of the open cell $V_p(\Gamma p)$ of p in the Voronoi diagram $\text{VD}_{\mathbb{H}}(\Gamma p)$ of Γp in \mathbb{H}^2 . From compactness of M follows that $\mathcal{D}_p(\Gamma)$ is a compact hyperbolic convex polygon with finitely many sides. The area of M is finite: $\text{area}(M) = 4\pi(\text{genus}(M) - 1)$. Each Dirichlet region $\mathcal{D}_p(\Gamma)$ is a *fundamental domain* for the action of Γ on \mathbb{H}^2 , i.e., (i) $\bigcup_{g \in \Gamma} g(\mathcal{D}_p(\Gamma)) = \mathbb{H}^2$, and (ii) $\forall g_1, g_2 \in \Gamma, g_1 \neq g_2, g_1(\text{Int}(\mathcal{D}_p(\Gamma))) \cap g_2(\text{Int}(\mathcal{D}_p(\Gamma))) = \emptyset$.

The *systole* $\text{sys}(M)$ is the least length of a non-contractible loop on M . In the sequel, it will play a role similar to the role played by the edge length of the fundamental domain above. There are upper bounds depending on the area and the genus of M [23, 9, 17]. The shortest closed curves on M are simple, they are closed geodesics and they are in one-to-one correspondence with the conjugacy classes of elements of Γ . The length of the closed geodesic γ corresponding to a conjugacy class $[h], h \in \Gamma$, is the translation length $l(h)$ [30]. The axis X_h projects to γ by π ; it can be thought of as “winding” X_h on γ infinitely many times. Any gX_h , where $g \in \Gamma$, projects onto the same closed geodesic γ by π ; the hyperbolic line gX_h is the axis of the translation ghg^{-1} . Note that the conjugacy classes $[h]$ and $[h^2]$ are different; the closed geodesic corresponding to $[h^2]$ traverses twice the closed geodesic corresponding to $[h]$. Also, the closed geodesics corresponding to $[h]$ and $[h^{-1}]$ are different.

3.2 Projecting $DT_{\mathbb{H}}(\Gamma\mathcal{P})$ on a closed hyperbolic surface

A triangulation defined by a point set \mathcal{P} in \mathbb{H}^2 is a *Delaunay triangulation* iff the circumscribing disk of each triangle does not intersect \mathcal{H}_{∞} and is *empty*, i.e., it does not contain any point of \mathcal{P} in its interior.

Let $M = \mathbb{H}^2/\Gamma$ be a closed hyperbolic surface, with projection map π . The images by π of two vertices of a triangle σ of $DT_{\mathbb{H}}(\Gamma\mathcal{P})$ may coincide; in this case at least one of the edges of $\pi(\sigma)$ is a 1-cycle. Any 1-cycle in the graph of edges of $\pi(DT_{\mathbb{H}}(\Gamma\mathcal{P}))$ is freely homotopic to a simple closed geodesic. A 1-cycle e with a vertex v is a geodesic loop. Note that e is not necessarily a closed geodesic: it may form a non-flat angle at vertex v . We refer the reader to the preprint [7] for the more technical characterization of 2-cycles.

Let us consider disks that do not contain any element of the Γ -orbit of a given point p . The Voronoi diagram of Γp is the partition of \mathbb{H}^2 by the images by all elements of Γ of Dirichlet regions $\mathcal{D}_p(\Gamma)$. Largest empty disks are centered at images of the vertices of $\mathcal{D}_p(\Gamma)$ that are furthest to p ; δ_p denotes their diameter. We define δ_M as $\delta_M = \sup\{\delta_z \mid z \in \mathcal{D}_O(\Gamma)\}$, where O denotes the origin. δ_M is completely determined by the group Γ . The following inequality holds:

► **Proposition 7.** $\delta_O \leq \delta_M \leq 2\delta_O$. (proof in preprint [7])

For a set of points \mathcal{P} , we denote as $\delta_{\mathcal{P}}$ the largest diameter of disks that do not contain any point of \mathcal{P} . Clearly, for $p \in \mathcal{P}$, $\delta_{\mathcal{P}} \leq \delta_p \leq \delta_M$.

► **Proposition 8.** If $sys(M) > \delta_{\mathcal{P}}$ (resp. $sys(M) > 2\delta_{\mathcal{P}}$) then the graph of edges of $\pi(DT_{\mathbb{H}}(\Gamma\mathcal{P}))$ has no 1-cycle (resp. no 2-cycle). (proof in [7])

Section 4.1 shows a covering of the Bolza surface in which the graph $\pi(DT_{\mathbb{H}}(\Gamma\mathcal{P}))$ has a 2-cycle formed by edges of length $\delta_{\mathcal{P}}$, i.e., $sys(M) = 2\delta_{\mathcal{P}}$, so the criterion is sharp. Let now $\tilde{\Gamma}$ be a subgroup of finite index k in Γ . The quotient space $\tilde{M} = \mathbb{H}^2/\tilde{\Gamma}$ is a closed hyperbolic surface with projection map $\tilde{\pi} : \mathbb{H}^2 \rightarrow \tilde{M}$. The surface \tilde{M} is a k -sheeted covering space of M with covering map and local isometry $f : \tilde{M} \rightarrow M$ such that $\pi = f \circ \tilde{\pi}$. Note that $Area(\tilde{M}) = k \cdot Area(M)$, $genus(\tilde{M}) = k(genus(M) - 1) + 1$, $sys(\tilde{M}) \geq sys(M)$, and $\delta_{\tilde{M}} = \delta_M$.

► **Proposition 9.** If $sys(\tilde{M}) > \delta_M$ (resp. $sys(\tilde{M}) > 2\delta_M$), then for any \mathcal{P} , the graph of edges of $\tilde{\pi}(DT_{\mathbb{H}}(\Gamma\mathcal{P}))$ has no 1-cycle (resp. no 2-cycle). There exists a finite-sheeted covering space \tilde{M}_1 (resp. \tilde{M}_2) of M such that $sys(\tilde{M}_1) > \delta_M$ (resp. $sys(\tilde{M}_2) > 2\delta_M$). (proof in [7])

Let us now focus on the specific case of the Bolza surface.

4 The Bolza surface

Consider a regular octagon in \mathbb{H}^2 centered at the origin O with angle $\pi/4$ at each vertex, and the four hyperbolic translations a, b, c, d that pairwise identify opposite sides of the octagon (Fig. 4). They generate a Fuchsian group \mathcal{G} for which the octagon is a fundamental domain. The inverse of an element $g \in \mathcal{G}$ is denoted as \bar{g} . The group \mathcal{G} can be written as a *finitely presented group*

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

i.e., the quotient of the group $\langle a, b, c, d \rangle$ generated by a, b, c, d , by the normal closure (i.e., the smallest normal subgroup) of the element $abcd\bar{a}\bar{b}\bar{c}\bar{d}$ in $\langle a, b, c, d \rangle$ [18, Chapter 5.5].

The *Bolza surface* is $\mathcal{M} = \mathbb{H}^2/\mathcal{G}$. It is a closed orientable hyperbolic surface homeomorphic to a torus having two handles, i.e., a double torus.





Figure 4 Translations a, b, c, d and their inverses.

The group \mathcal{G} shows interesting properties in relation to some triangle groups. For integers $p, q, r \geq 2$, $\frac{1}{p} + \frac{1}{q} + \frac{1}{r} < 1$, the *triangle group* $T(p, q, r)$ is defined as the group generated by reflections in the edges of a triangle $\Delta(p, q, r)$ whose angles are $\frac{\pi}{p}, \frac{\pi}{q}, \frac{\pi}{r}$ [30]. The set of triangles $\{g(\Delta(p, q, r)), g \in T(p, q, r)\}$ partitions the plane \mathbb{H}^2 . \mathcal{G} is a normal subgroup of index 16 without fixed points of $T(2, 8, 8)$. \mathcal{G} is a normal subgroup of index 96 without fixed points of $T(2, 3, 8)$. These properties of \mathcal{G} are illustrated geometrically on Fig. 5. For $T(2, 3, 8)$ (right): repeated applications of reflections in the edges of $\Delta(2, 3, 8)$ form the fundamental domain of \mathcal{G} ; the projection map from \mathbb{H}^2 to \mathcal{M} induces the partition of the surface \mathcal{M} into triangles $\Delta(2, 3, 8)$. $T(2, 3, 8)$ is actually the largest subgroup of $\text{Isom}(\mathbb{H}^2)$ in which \mathcal{G} is normal, i.e., $T(2, 3, 8) = \{t \in \text{Isom}(\mathbb{H}^2) \mid t\mathcal{G}t^{-1} = \mathcal{G}\}$.



■ **Figure 5** Partition of $\mathcal{D}_O(\mathcal{G})$ into fundamental domains for (left) $T(2, 8, 8)$ and (right) $T(2, 3, 8)$ (circles show the elements of the orbit $T(2, 3, 8)O$ in $\mathcal{D}_O(\mathcal{G})$).

► **Proposition 10.** If two points p and q of \mathbb{H}^2 lie in the same orbit under $T(2, 3, 8)$, then the Dirichlet regions $\mathcal{D}_p(\mathcal{G})$ and $\mathcal{D}_q(\mathcal{G})$ are isometric.

As will be seen in the sequel, this is not true otherwise, which is one of the striking differences with the Euclidean case.

Proof. $q = r(p)$ for $r \in T(2, 3, 8)$. Since \mathcal{G} is normal in $T(2, 3, 8)$, $r\mathcal{G}r^{-1} = \mathcal{G}$. Then the orbit $\mathcal{G}q = r\mathcal{G}r^{-1}q = r\mathcal{G}p$ is the image of the orbit $\mathcal{G}p$ under r . $\text{VD}_{\mathbb{H}}(\mathcal{G}q)$ is the image of $\text{VD}_{\mathbb{H}}(\mathcal{G}p)$ by r . The result follows since $\mathcal{D}_q(\mathcal{G}) = \text{Cl}(V_q(\mathcal{G}q))$. ◀

Notation. For simplicity, in the sequel we denote as g the image $g(O)$ of O by $g \in \mathcal{G}$. We also label each side of $\mathcal{D}_p(\mathcal{G})$ by the translation of \mathcal{G} that maps $\mathcal{D}_p(\mathcal{G})$ to the Dirichlet region adjacent through this side. Let $\rho \in T(2, 3, 8)$ be the rotation around the origin O by $\pi/4$, and let

$$\mathcal{A} = (a, \bar{b}, c, \bar{d}, \bar{a}, b, \bar{c}, d)$$

be the alphabet formed by the generators of \mathcal{G} and their inverses, considered as an ordered sequence (Fig. 4). Conjugation by ρ acts as a permutation on \mathcal{A} .

The Dirichlet regions for \mathcal{G} have been studied in [29]. The rest of this section presents a more intuitive description.

For $x \in \mathcal{A}$, we denote as v_x the vertex of $\mathcal{D}_O(\mathcal{G})$ that is equidistant to O, x , and $\rho^{-1}x\rho$. The eight elements of \mathcal{A} are the elements of the orbit $\mathcal{G}O$ that are closest to O . In the same way, the eight elements $xz, z \in \mathcal{A}$ are closest to x . Any $z \in \mathcal{A}$ can be written as $\rho^k \bar{x} \rho^{-k}$, where $0 \leq k < 8$, so the above $xz, z \in \mathcal{A}$ can be written as $xz = x \cdot \rho^k \bar{x} \rho^{-k}(O)$. A careful but straightforward observation of the Dirichlet regions of $O, x, x \cdot \rho \bar{x} \rho^{-1}$, etc., shows:

► **Observation 11.** Each vertex v_x is incident to eight isometric Dirichlet regions. See Fig. 6, where $\omega_4(x) = x \cdot \rho \bar{x} \rho^{-1} \cdot \rho^2 x \rho^{-2} \cdot \rho^3 \bar{x} \rho^{-3} = y \cdot \rho^{-1} \bar{y} \rho \cdot \rho^{-2} y \rho^2 \cdot \rho^{-3} \bar{y} \rho^3$.

Without loss of generality, we now examine more closely the vertex v of $\mathcal{D}_O(\mathcal{G})$ that is incident to $\mathcal{D}_a(\mathcal{G})$ and $\mathcal{D}_d(\mathcal{G})$ (see Fig. 4). The eight Dirichlet regions incident to v

are centered at the images of O by the following translations, in clockwise order: $F = (\mathbb{K}, a, ab, abc, abcd = dcba, dcb, dc, d)$ (Fig. 7). The eight Dirichlet regions $\rho^k \mathcal{D}_v(\mathcal{G}), 0 \leq k < 8$, share the common vertex O . The Dirichlet region $\mathcal{D}_v(\mathcal{G})$ is isometric to $\mathcal{D}_O(\mathcal{G})$ since v lies in the $T(2, 3, 8)$ -orbit of O (Fig. 5); $\mathcal{D}_v(\mathcal{G})$ is the regular octagon with vertices in $F(O)$.




Figure 6 Dirichlet regions around v_x .




Figure 7 Schematic representation of $\mathcal{D}_v(\mathcal{G})$.

For a point $p \in \mathbb{H}^2$, we define $C(p) = \{g \mid g \in \mathcal{G}, \mathcal{D}_p(\mathcal{G}) \cap \mathcal{D}_{g(p)}(\mathcal{G}) \neq \emptyset\}$. Notation C can be remembered as standing for “crown”. For conveniency, we will sometimes use the same notation $C(p)$ for the set of points $\{g(p) \mid g \in C(p)\}$ (the context will disambiguate). It is easy to check that for $q = t(p), t \in T(2, 3, 8)$ we have $C(q) = \{tgt^{-1} \mid g \in C(p)\}$.

Let p be a point in triangle SOQ shown in Fig. 8. Let us consider the Dirichlet regions that are adjacent to $\mathcal{D}_p(\mathcal{G})$.




Figure 8 SOQ .

- If p is in the interior of SOQ , the Dirichlet region $\mathcal{D}_p(\mathcal{G})$ is a polygon with 18 sides. Each vertex of $\mathcal{D}_p(\mathcal{G})$ has degree 3. The sides of $\mathcal{D}_p(\mathcal{G})$ have the following labels in counterclockwise order: $a, a\bar{d}, \bar{b}c\bar{d}, \bar{b}, \bar{b}\bar{a}, c, \bar{d}, \bar{a}, b, \bar{c}\bar{d}, \bar{c}, \bar{c}\bar{b}\bar{a}, \bar{d}\bar{a}, d, dc, dcb, abc, ab$.
- If p lies in the interior of edge OQ , the sides dc and abc degenerate to a vertex of $\mathcal{D}_p(\mathcal{G})$. On SO , the sides ab and dcb degenerate to a vertex, and on QS b and c degenerate to a vertex. Such a new vertex has degree 4. The Dirichlet region $\mathcal{D}_p(\mathcal{G})$ is a polygon with 14 sides.
- If $p \in \{O, Q, S\}$, $\mathcal{D}_p(\mathcal{G})$ degenerates to an octagon (Proposition 10). Its adjacent octagons are centered at some points in $C(p)$. In particular: $\mathcal{D}_p(\mathcal{G})$ shares a common vertex with $\mathcal{D}_{abcd(p)}(\mathcal{G})$ iff $p = O$.

Fig. 9 shows the Dirichlet regions of three different points q , obtained by computing the Delaunay triangulation in \mathbb{H}^2 of sufficiently many points of $\mathcal{G}q$, with the software described in [6].




Figure 9 Dirichlet region for \mathcal{G} centered at different points q (rightmost picture: $q = O$).

We define $C^-(O) = C(O) \setminus \{\rho^k(abcd)\rho^{-k}(O) \mid 0 \leq k < 8\}$; it excludes the elements of $C(O)$ corresponding to octagons that meet $\mathcal{D}_O(\mathcal{G})$ opposite a vertex. For $O' = rO, r \in T(2, 3, 8)$, define $C^-(O') = \{rgr^{-1} \mid g \in C^-(O)\}$. From the description above:

► **Proposition 12.** For any $p \in \mathbb{H}^2$ there exists $O' \in T(2, 3, 8)O$, such that $C(p) \subset C(O')$. Moreover, if $p \notin T(2, 3, 8)O$, then $C(p) \subset C^-(O')$.

► **Proposition 13.** The largest empty disk diameter satisfies $4.89 < \delta_{\mathcal{M}} < 6.62$. (*proof in [7]*) There exist closed geodesics that are $< \delta_{\mathcal{M}}$ [3] (see preprint [7]).

4.1 Covering spaces of the Bolza surface

The function `LowIndexSubgroupsFpGroup(Gamma, H, k, excluded)` of GAP [1, 34] provides in principle a way to compute the two covering spaces of Prop. 9. However, its huge memory consumption does not allow us to use it for the Bolza surface. We propose a construction of coverings spaces of \mathcal{M} in which, for any set \mathcal{P} , the projection $\pi(DT_{\mathbb{H}}(\Gamma\mathcal{P}))$ does not contain any 1- or 2-cycle. In a similar spirit as for the flat torus, here we look for subgroups $\tilde{\mathcal{G}}$ of \mathcal{G} not containing elements corresponding to short loops, and whose associated Dirichlet regions keep some symmetry.

\mathcal{G} is a normal subgroup of the triangle groups $T(2, 3, 8), T(2, 4, 8), T(2, 8, 8)$. They contain the rotation ρ defined earlier, and the reflection τ with respect to the axis X_a of translation a . Obviously, if $\tilde{\mathcal{G}}$ is a normal subgroup of one of these three triangle groups, then $\tilde{\mathcal{G}} = \rho\tilde{\mathcal{G}}\rho^{-1}$ and $\tilde{\mathcal{G}} = \tau\tilde{\mathcal{G}}\tau^{-1}$ (which holds for $\tilde{\mathcal{G}} = \mathcal{G}$).

Let us introduce the subset $N_O(\tilde{\mathcal{G}}) \subset \tilde{\mathcal{G}}$ of elements g such that the Dirichlet regions $\mathcal{D}_O(\tilde{\mathcal{G}})$ and $\mathcal{D}_g(\tilde{\mathcal{G}})$ are adjacent (i.e., they share a common edge). For $\tilde{\mathcal{G}} = \mathcal{G}$, $N_O(\mathcal{G}) = \mathcal{A}$. More generally, $N_O(\tilde{\mathcal{G}})$ generates $\tilde{\mathcal{G}}$ [26].

► **Observation 14.** The Dirichlet region $\mathcal{D}_O(\tilde{\mathcal{G}})$ for $\tilde{\mathcal{G}}$ centered at O is invariant by ρ (resp. τ) if and only if $\tilde{\mathcal{G}} = \rho\tilde{\mathcal{G}}\rho^{-1}$ (resp. $\tilde{\mathcal{G}} = \tau\tilde{\mathcal{G}}\tau^{-1}$).

Proof. If $\tilde{\mathcal{G}} = \rho\tilde{\mathcal{G}}\rho^{-1}$, then by definition of ρ , $\rho\tilde{\mathcal{G}}O = \rho\tilde{\mathcal{G}}\rho^{-1}O = \tilde{\mathcal{G}}O$, and the Dirichlet region $\mathcal{D}_O(\tilde{\mathcal{G}})$, as a cell in $VD_{\mathbb{H}}(\tilde{\mathcal{G}}O)$, is invariant by ρ (and similarly for τ).

Now assume that $\mathcal{D}_O(\tilde{\mathcal{G}})$ is invariant by ρ . Then in particular $N_O(\tilde{\mathcal{G}})(O) = \rho N_O(\tilde{\mathcal{G}})(O)$, and we can rewrite $N_O(\tilde{\mathcal{G}})(O)$ as $\rho N_O(\tilde{\mathcal{G}})\rho^{-1}(O)$ (since $\rho(O) = O$). Thus for each $g \in N_O(\tilde{\mathcal{G}})$, there is some $h \in N_O(\tilde{\mathcal{G}})$ such that $g(O) = \rho h \rho^{-1}(O)$, and since $\mathcal{G} = \rho\mathcal{G}\rho^{-1}$, so $\rho h \rho^{-1}$ is in \mathcal{G} . Then the translation $\rho h \rho^{-1} \cdot \bar{g}$ is in \mathcal{G} , and it fixes O . Since \mathcal{G} has no fixed point, it must be the identity, and $\rho h \rho^{-1} = g$. This is true for each g in $N_O(\tilde{\mathcal{G}})$, so $N_O(\tilde{\mathcal{G}}) = \rho N_O(\tilde{\mathcal{G}})\rho^{-1}$. Since $N_O(\tilde{\mathcal{G}})$ generates $\tilde{\mathcal{G}}$, we have $\tilde{\mathcal{G}} = \rho\tilde{\mathcal{G}}\rho^{-1}$. The same proof works for τ : the only property of ρ that we used is that $\rho(O) = O$, which is also true for τ . ◀

It follows as a corollary that if $\tilde{\mathcal{G}}$ is normal in a triangle group, then $\mathcal{D}_O(\tilde{\mathcal{G}})$ is invariant under ρ and τ . The three triangle groups $T(2, m, 8)$ mentioned above can be generated by ρ, τ and some other rotation r_m . If $\mathcal{D}_O(\tilde{\mathcal{G}})$ is invariant under ρ and τ , then $\tilde{\mathcal{G}}$ is normal in $T(2, m, 8)$ if and only if $\tilde{\mathcal{G}} = r_m\tilde{\mathcal{G}}r_m^{-1}$.

Before explicitly constructing subgroups of \mathcal{G} , let us derive a condition that will be used later (Prop. 22):

► **Proposition 15.** The condition that $\tilde{\mathcal{G}}$ contains no element of $C^-(O')$ for any $O' \in T(2, 3, 8)O$ is necessary and sufficient for the graph of edges of the projection of $DT_{\mathbb{H}}(\mathcal{G}\mathcal{P})$ onto $\tilde{\mathcal{M}} = \mathbb{H}^2/\tilde{\mathcal{G}}$ to have no 1-cycle, for any \mathcal{P} .

Proof. Let e be an edge of $DT_{\mathbb{H}}(\mathcal{G}\mathcal{P})$ such that $\pi(e)$ is a geodesic loop on \mathcal{M} , i.e., $e = [p, g(p)]$ for some $p \in \mathbb{H}^2, g \in \mathcal{G}$. Since $e \in DT_{\mathbb{H}}(\mathcal{G}\mathcal{P})$, the cells $Cl(V_p(\mathcal{G}\mathcal{P}))$ and $Cl(V_{g(p)}(\mathcal{G}\mathcal{P}))$

intersect. Since $Cl(V_p(\mathcal{GP})) \subset Cl(V_p(\mathcal{G}_p)) = \mathcal{D}_p(\mathcal{G})$ and similarly $Cl(V_{g(p)}(\mathcal{GP})) \subset \mathcal{D}_{g(p)}(\mathcal{G})$, then $\mathcal{D}_p(\mathcal{G}) \cap \mathcal{D}_{g(p)}(\mathcal{G}) \neq \emptyset$, i.e., $g \in C(p)$.

By Proposition 12, $g \in C(p) \subset C(O')$ for some $O' \in T(2, 3, 8)O$. So if $\tilde{\mathcal{G}}$ is a subgroup of \mathcal{G} that contains no elements of $C(O')$, then the projection of $DT_{\mathbb{H}}(\mathcal{GP})$ to $\tilde{\mathcal{M}} = \mathbb{H}^2/\tilde{\mathcal{G}}$ has no 1-cycle.

Assume that $g \notin C^-(O')$. Then Proposition 12 implies that $p \in T(2, 3, 8)O$. Without loss of generality assume $p = O$ and $g = abcd$. If \mathcal{P} contains more than one point, then $[O, abcd]$ cannot be a Delaunay edge. If $\mathcal{P} = \{O\}$, then as shown on Fig. 10, we can choose a triangulation of the octagon $\mathcal{D}_v(\mathcal{G})$ that does not contain edge $[O, abcd]$ (see also the treatment of degenerate cases in the introduction section). Then the projection of $DT_{\mathbb{H}}(\mathcal{GP})$ has no 1-cycle. Remark that for any $g \in C^-(O)$, we can choose $DT_{\mathbb{H}}(\mathcal{GO})$ so that the geodesic segment with endpoints O and g is an edge of $DT_{\mathbb{H}}(\mathcal{GO})$.




Figure 10 Avoiding edge $[O, abcd]$.

In the sequel, we use the standard notation $[g, h]$ for the commutator of two isometries g and h , i.e., $[g, h] = ghg^{-1}h^{-1}$.

We first define the subgroup \mathcal{G}_2 of \mathcal{G} generated by words of even size on \mathcal{A} (Fig. 11). \mathcal{G}_2 is a subgroup of index 2 in \mathcal{G} , hence it is normal in \mathcal{G} .

► Proposition 16. \mathcal{G}_2 is generated by $\{[x, \rho] \mid x \in \mathcal{A}\}$.

Proof. It is sufficient to prove that any word of size 2 on \mathcal{A} can be generated by this subset. Consider $yz, y, z \in \mathcal{A}$. As mentioned in Section 4, ρ acts as a permutation on \mathcal{A} , i.e., $y = \rho^{k_1}x\rho^{-k_1}$ and $z = \rho^{k_2}\bar{x}\rho^{-k_2}$ for some $k_1, k_2, 0 \leq k_1, k_2 < 8$. Then $yz = (\rho^{k_1}x\rho^{-k_1} \cdot \rho^{k_1+1}\bar{x}\rho^{-k_1-1}) \cdot (\rho^{k_1+1}x\rho^{-k_1-1} \cdot \rho^{k_1+2}\bar{x}\rho^{-k_1-2}) \cdot \dots \cdot (\rho^{k_2-1}x\rho^{-k_2+1} \cdot \rho^{k_2}\bar{x}\rho^{-k_2})$.

The subgroup \mathcal{G}_2 can be written in a finitely presented form:

$$\mathcal{G}_2 = \langle ab, cd, \bar{a}\bar{b}, \bar{c}\bar{d}, bc, \bar{d}\bar{a}, \bar{b}\bar{c}, \bar{d}\bar{a} \mid abcd\bar{a}\bar{b}\bar{c}\bar{d} \rangle.$$

As a consequence of Observation 11, $[x, \rho]$, $[\rho^{-1}x\rho, \rho]$, and O are equidistant to v_x , and the Dirichlet region of O for \mathcal{G}_2 can be completely described (Fig. 11):

► Observation 17. The fundamental region $\mathcal{D}_O(\mathcal{G}_2)$ has 16 vertices that form the circular sequence $(v_x, x \mid x \in \mathcal{A})$, in counterclockwise order. The set $N_O(\mathcal{G}_2)$ is the circular sequence $([x, \rho], [x, \rho^{-1}] \mid x \in \mathcal{A})$ in counterclockwise order.




Figure 11 $\mathcal{D}_O(\mathcal{G}_2)$.

The Dirichlet region of O is clearly invariant under ρ and τ . Indeed, conjugation by ρ or τ leaves \mathcal{A} invariant, so it preserves the size of a word in \mathcal{G} . By Observation 14, $\mathcal{G}_2 = \rho\mathcal{G}_2\rho^{-1} = \tau\mathcal{G}_2\tau^{-1}$. We can actually show:

► Proposition 18. \mathcal{G}_2 is the only subgroup of index 2 of \mathcal{G} for which the Dirichlet region centered at O is invariant by ρ . (proof in preprint [7])

The generators of \mathcal{G}_2 all have the same translation length $l(a) = l(b) = l(c) = l(d)$. As a result, $sys(\mathcal{M}_2) = sys(\mathcal{M})$. However, there are fewer shortest loops on $\tilde{\mathcal{M}}$ than on \mathcal{M} , since the elements of \mathcal{A} are not in \mathcal{G}_2 .

As for the flat case, we can continue and define inductively the group \mathcal{G}_{2^k} consisting of elements of even size in the alphabet of generators of $\mathcal{G}_{2^{k-1}}$, for $k \geq 1$. Recall notation $\omega_4(x) = x \cdot \rho\bar{x}\rho^{-1} \cdot \rho^2x\rho^{-2} \cdot \rho^3\bar{x}\rho^{-3}$ (Fig. 6).


► **Observation 19.** The domain $\mathcal{D}_O(\mathcal{G}_4)$ has 24 vertices that form the circular sequence $([x, \rho], x, [x, \rho^{-1}] \mid x \in \mathcal{A})$, in counterclockwise order. The set $N_O(\mathcal{G}_4)$ is the circular sequence $(\omega_4(x), x \cdot \rho^2 \bar{x} \rho^{-2}, x \cdot \rho^{-2} \bar{x} \rho^2 \mid x \in \mathcal{A})$, in counterclockwise order.

Proof. Let Π be the polygon with vertices $([x, \rho], x, [x, \rho^{-1}] \mid x \in \mathcal{A})$, in counterclockwise order. Consider a vertex $x \in \mathcal{A}$ of Π . Recall that the elements of the set $\{x \cdot \rho^k \bar{x} \rho^{-k} \mid 0 \leq k < 8\}$ are equidistant from x . Vertices $(\rho \cdot x \cdot \rho^{-1}) \cdot (\rho^2 \cdot \bar{x} \cdot \rho^{-2})$, $(\rho^2 \cdot x \cdot \rho^{-2}) \cdot (\rho^3 \cdot \bar{x} \cdot \rho^{-3})$, and $\rho \bar{x} \rho^{-1} \cdot \bar{x}$ of $\mathcal{D}_O(\mathcal{G}_2)$ are equidistant to O . Thus $[x, \rho] \cdot (\rho \cdot x \cdot \rho^{-1}) \cdot (\rho^2 \cdot \bar{x} \cdot \rho^{-2})$, $[x, \rho] \cdot (\rho^2 \cdot x \cdot \rho^{-2}) \cdot (\rho^3 \cdot \bar{x} \cdot \rho^{-3})$, and $[x, \rho] \rho \bar{x} \rho^{-1} \cdot \bar{x}$ are equidistant from the vertex $[x, \rho]$ of Π , i.e., $x \cdot \rho^2 \bar{x} \rho^{-2}$, $\omega_4(x)$, and O are equidistant from the vertex $[x, \rho]$ of Π . The proof is similar for $[x, \rho^{-1}]$. ◀

Using the fact that $N_O(\mathcal{G}_4)$ generates \mathcal{G}_4 , we get:

► **Corollary 20.** The set $\{\omega_4(x), x \cdot \rho^2 \bar{x} \rho^{-2}, x \cdot \rho^{-2} \bar{x} \rho^2 \mid x \in \mathcal{A}\}$ generates \mathcal{G}_4 .

In other words, the generators of \mathcal{G}_4 are $ac, abcd, d\bar{b}$, and their conjugates by ρ^k , $1 < k < 8$ (Fig. 12). In a similar way as done for \mathcal{G}_2 , we can show that \mathcal{G}_4 is the only subgroup of index 2 of \mathcal{G}_2 for which the Dirichlet region centered at O is invariant by ρ .



► **Proposition 21.** The systole of $\mathcal{M}_4 = \mathbb{H}^2/\mathcal{G}_4$ is equal to $l(abcd)$.

■ **Figure 12** $\mathcal{D}_O(\mathcal{G}_4)$.

Proof. $sys(\mathcal{M})$ is equal to the translation length $l(a)$ of the translation a . The translations of \mathcal{G} that have the same translation length as a are conjugates hah^{-1} , where $h \in T(2, 3, 8)$. As $a \notin \mathcal{G}_4$ and \mathcal{G}_4 is normal in $T(2, 3, 8)$, it does not contain hah^{-1} for any $h \in T(2, 3, 8)$. Thus $sys(\mathbb{H}^2/\mathcal{G}_4) > l(a)$. The second shortest length of a non-contractible loop corresponds to geodesics of length $l(abcd)$ [3] (see preprint [7]). Since $abcd \in \mathcal{G}_4$ (violet point in Fig. 12), the result follows. ◀

The systole of \mathcal{M}_4 is greater than $sys(\mathcal{M})$, but still smaller than $\delta_{\mathcal{M}}$. However, as a consequence of Proposition 15:

► **Proposition 22.** For any \mathcal{P} , the projection of $DT_{\mathbb{H}}(\mathcal{G}\mathcal{P})$ on \mathcal{M}_4 has no 1-cycle.

Due to lack of space, we only mention the following result (see preprint [7]):

► **Proposition 23.** The index of a subgroup $\tilde{\mathcal{G}}$ for which $sys(\mathbb{H}^2/\tilde{\mathcal{G}}) > 2\delta_{\mathcal{M}}$ is greater than 32.

There is a unique subgroup \mathcal{G}_{128} of \mathcal{G} of index 128 that is normal in $T(2, 3, 8)$. For any set of points \mathcal{P} , the graph of edges of the projection of $DT_{\mathbb{H}}(\mathcal{G}\mathcal{P})$ to $\mathbb{H}^2/\mathcal{G}_{128}$ has no 2-cycle.

4.2 Triangulating the Bolza surface in practice

As mentioned in introduction, the incremental algorithm for computing a Delaunay triangulation can be extended to \mathcal{M} , using a covering space of \mathcal{M} in which the graph of its edges has no 2-cycles, similarly to what was done for Euclidean manifolds [15]. However, computing in $\mathbb{H}^2/\mathcal{G}_{128}$ (Proposition 23) would not be efficient, since we would have to consider 128 images of \mathcal{P} .

In this section, we propose a heuristic method instead, inspired from [11, 12]. It consists in first initializing the triangulation with a well chosen set \mathcal{Q} of 14 points for which $sys(\mathcal{M}) > 2\delta_{\mathcal{Q}}$. Inserting new points in the triangulation cannot increase the size of empty disks, i.e., $\delta_{\mathcal{P}' \cup \mathcal{Q}} \leq \delta_{\mathcal{Q}}$ for any set \mathcal{P}' , so by Proposition 8, $\pi(DT_{\mathbb{H}}(\mathcal{G}(\mathcal{P}' \cup \mathcal{Q})))$ has no 2-cycle. Thus, the incremental algorithm can be used to compute the Delaunay triangulation of $\mathcal{P} \cup \mathcal{Q}$ directly

in \mathcal{M} . In a lot of applications, keeping these 14 extra vertices in the triangulation would not be an issue. However, there may be cases when this should be avoided. If \mathcal{P} contains sufficiently many points, and if they are reasonably well distributed, which is likely to be the case in applications,³ then the final triangulation $\pi(\text{DT}_{\mathbb{H}}(\mathcal{G}\mathcal{P}))$ would have no 1- or 2-cycles. In such cases, all points of \mathcal{Q} can be removed from $\pi(\text{DT}_{\mathbb{H}}(\mathcal{G}(\mathcal{P} \cup \mathcal{Q})))$ (basically, as in a 2D Euclidean triangulation [19] —practical details are omitted). However, for badly distributed point sets, such short cycles will appear during the process, which can be checked after each removal using the geometric criterion; in such cases, points of \mathcal{Q} cannot be removed further.

The rest of this section is devoted to presenting the construction of \mathcal{Q} .

The group \mathcal{G} is a normal subgroup without fixed point of index 48 in the triangle group $T(2, 4, 8)$. The octagon $\mathcal{D}_O(\mathcal{G})$ can be partitioned into 48 triangles that are fundamental domains for $T(2, 4, 8)$ (Fig. 13). Let Δ be such a triangle incident to O . If \mathcal{Q}_T denotes the set of vertices of Δ , we define $\mathcal{Q} = T(2, 4, 8)\mathcal{Q}_T \cap \mathcal{D}_O(\mathcal{G})$. Note that $T(2, 4, 8)\mathcal{Q}_T = \mathcal{G}\mathcal{Q}$. Finally, we flip an edge of Δ as shown on Fig. 14.




Figure 14 Flip.

Figure 13 A fundamental domain Δ for $T(2, 4, 8)$.

The disks circumscribing Δ_1 and Δ_2 contain no point of $T(2, 4, 8)\mathcal{Q}_T$ in their interiors. $\Delta_1 \cup \Delta_2$ is a fundamental domain for $T^+(2, 4, 8) = T(2, 4, 8) \cap \text{Isom}^+(\mathbb{H}^2)$. This partition of \mathbb{H}^2 is $\text{DT}_{\mathbb{H}}(T(2, 4, 8)\mathcal{Q}_T) = \text{DT}_{\mathbb{H}}(\mathcal{G}\mathcal{Q})$ since $T(2, 4, 8)\mathcal{Q}_T = \mathcal{G}\mathcal{Q}$.

Let $\delta_{\mathcal{Q}}$ denote the largest empty disk diameter for $\mathcal{G}\mathcal{Q}$, and d_1 (resp. d_2) denote the diameters of the disks circumscribing Δ_1 (resp. Δ_2). There is a triangle in $\text{DT}_{\mathbb{H}}(\mathcal{G}\mathcal{Q})$ whose circumscribing disk has diameter $\delta_{\mathcal{Q}}$, and since any triangle in $\text{DT}_{\mathbb{H}}(\mathcal{G}\mathcal{Q})$ is the image of Δ_1 or Δ_2 under some isometry of $T^+(2, 4, 8)$, either d_1 or d_2 is equal to $\delta_{\mathcal{Q}}$.

Both d_1 and d_2 are less than $\frac{1}{2}\text{sys}(\mathcal{M})$, which actually is the radius of the circle inscribed in $\mathcal{D}_O(\mathcal{G})$. So $\delta_{\mathcal{Q}} < \frac{1}{2}\text{sys}(\mathcal{M})$.

Fig. 15 shows unique representatives in \mathbb{H}^2 of the triangles of $\pi(\text{DT}_{\mathbb{H}}(\mathcal{G}\mathcal{Q}))$. The point set \mathcal{Q} contains 9 points in the interior of \mathcal{D}_O (origin and red points), 4 points on the boundary of \mathcal{D}_O (green points) and 1 point at its vertex (blue point).




Figure 15 The set \mathcal{Q} .

In total \mathcal{Q} consists of 14 points.

5 Open problems

Further work directions are twofold. On a theoretical side, we are planning to reduce the gap between the lower bound 32 and the upper bound 128 in Proposition 23; we conjecture that 32 sheets are always sufficient to avoid 2-cycles. We will also investigate generalizations of the results given here to closed hyperbolic surfaces of higher genus. On a practical side,

³ this was observed and analyzed in the 3D Euclidean case [11, 13]

we are working on arithmetic issues arising in the heuristic method presented in Section 4.2, and on its implementation for CGAL.

References

- 1 Gap - groups, algorithms, programming. URL: <http://www.gap-system.org>.
- 2 M. Artin. *Algebra*. Prentice-Hall, 2002.
- 3 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.
- 4 A. Bachelot-Motet. Wave computation on the hyperbolic double doughnut. *Journal of Computational Mathematics*, 28:1–17, 2010. doi:10.4208/jcm.1004.m3120.
- 5 M. Berger. *Geometry (vols. 1-2)*. Springer-Verlag, 1987.
- 6 M. Bogdanov, O. Devillers, and M. Teillaud. Hyperbolic Delaunay complexes and Voronoi diagrams made practical. *Journal of Computational Geometry*, 5:56–85, 2014. URL: <http://jocg.org/index.php/jocg/article/view/141>.
- 7 M. Bogdanov and M. Teillaud. Delaunay triangulations and cycles on closed hyperbolic surfaces. Research Report 8434, INRIA, December 2013. URL: <http://hal.inria.fr/hal-00921157>.
- 8 A. Bowyer. Computing Dirichlet tessellations. *The Computer Journal*, 24:162–166, 1981.
- 9 P. Buser and P. Sarnak. On the period matrix of a Riemann surface of large genus. *Inventiones mathematicae*, 117:27–56, 1994. doi:10.1007/BF01232233.
- 10 J. W. Cannon, W. J. Floyd, R. Kenyon, and W. R. Parry. Hyperbolic geometry. *Flavors of geometry*, 31:59–115, 1997. URL: <http://www.math.uwo.ca/~shafikov/teaching/winter2010/4156/hyperbolic.pdf>.
- 11 M. Caroli. *Triangulating Point Sets in Orbit Spaces*. Phd thesis, Université de Nice-Sophia Antipolis, France, 2010. URL: <http://tel.archives-ouvertes.fr/tel-00552215/>.
- 12 M. Caroli and M. Teillaud. 3D periodic triangulations. In *CGAL Manual*. CGAL Editorial Board, 3.5 edition, 2009. URL: <http://doc.cgal.org/latest/Manual/packages.html#PkgPeriodic3Triangulation3Summary>.
- 13 M. Caroli and M. Teillaud. Computing 3d periodic triangulations. In *European Symposium on Algorithms*, volume 5757 of *LNCS*, pages 59–70, 2009. To appear in [15]. doi:10.1007/978-3-642-04128-0_6.
- 14 M. Caroli and M. Teillaud. Delaunay triangulations of point sets in closed Euclidean d -manifolds. In *Proc. 27th Symposium on Computational Geometry*, pages 274–282, 2011. To appear in [15]. doi:10.1145/1998196.1998236.
- 15 M. Caroli and M. Teillaud. Delaunay triangulations of closed Euclidean d -orbifolds. *Discrete & Computational Geometry*, To Appear. See also [13] and [14].
- 16 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.
- 17 É. Colin de Verdière, A. Hubard, and A. de Mesmay. Discrete systolic inequalities and decompositions of triangulated surfaces. In *Proc. 30th Symposium on Computational Geometry*, pages 335–344, 2014. To appear in *Discrete & Computational Geometry*.
- 18 H. S. M. Coxeter and W. O. J. Moser. *Generators and Relations for Discrete Groups*. Springer-Verlag, Berlin, Heidelberg, New York, Tokyo, 1957.
- 19 O. Devillers. Vertex removal in two-dimensional Delaunay triangulation: Speed-up by low degrees optimization. *CGTA*, 44:169–177, 2011. doi:10.1016/j.comgeo.2010.10.001.
- 20 O. Devillers and M. Teillaud. Perturbations for Delaunay and weighted Delaunay 3D triangulations. *CGTA*, 44:160–168, 2011. doi:10.1016/j.comgeo.2010.09.010.
- 21 N.P. Dolbilin and D.H. Huson. Periodic Delone tilings. *Periodica Mathematica Hungarica*, 34:1-2:57–64, 1997.

- 22 C. I. Grima and A. Márquez. *Computational Geometry on Surfaces*. Kluwer Academic Publishers, 2001.
- 23 M. Gromov. Systoles and intersystolic inequalities. In *Actes de la Table Ronde de Géométrie Différentielle (Luminy, 1992)*, volume 1 of *Sémin. Congr.*, pages 291–362. Soc. Math. France, 1996.
- 24 J. Hidding, R. vd Weygaert, G. Vegter, B.J.T. Jones, and M. Teillaud. Video: The sticky geometry of the cosmic web. In *Proceedings of the Twenty-eighth Annual Symposium on Computational Geometry*, pages 421–422, 2012. URL: <http://www.computational-geometry.org/SoCG-videos/socg12video/>.
- 25 M.K. Hurdal and K. Stephenson. Cortical cartography using the discrete conformal approach of circle packings. *Neuroimage*, 23:119–128, 2004.
- 26 S. Katok. *Fuchsian Groups*. Chicago Lectures in Mathematics. Uni. Chicago Press, 1992.
- 27 N. Kruithof. 2D periodic triangulations. In *CGAL User and Reference Manual*. CGAL Editorial Board, 4.4 edition, 2014. URL: <http://doc.cgal.org/latest/Manual/packages.html#PkgPeriodic2Triangulation2Summary>.
- 28 M. Mazón and T. Recio. Voronoi diagrams on orbifolds. *Computational Geometry: Theory and Applications*, 8:219–230, 1997. doi:10.1016/S0925-7721(96)00017-X.
- 29 M. Naatanen. On the stability of identification patterns for Dirichlet regions. *Annales Academiæ Scientiarum Fennicæ, Series A.I. Mathematica*, 10:411–417, 1985.
- 30 J. Ratcliffe. *Foundations of Hyperbolic Manifolds*. Graduate Texts in Mathematics. Springer, 2006.
- 31 G. Ron, M. Jin, and X. Guo. Hyperbolic centroidal Voronoi tessellation. In *Proc. ACM Symp. on Solid and Physical Modeling*, pages 117–126, 2010. doi:10.1145/1839778.1839795.
- 32 F. Sausset, G. Tarjus, and P. 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.
- 33 B. Schuetrumpf, M. A. Klatt, K. Iida, G. E. Schröder-Turk, J. A. Maruhn, K. Mecke, and P.-G. Reinhard. Appearance of the single gyroid network phase in “nuclear pasta” matter. *Physical Review C*, 91(025801), 2015. doi:10.1103/PhysRevC.91.025801.
- 34 C.C. Sims. *Computation with Finitely Presented Groups*. Encyclopedia of Mathematics and its Applications. Cambridge University Press, 1994.
- 35 T. Sousbie. The persistent cosmic web and its filament structure I. *Monthly Notices of the Royal Astronomical Soc.*, 414:350–383, 2011. doi:10.1111/j.1365-2966.2011.18394.x.
- 36 T. Sousbie, C. Pichon, and H. Kawahara. The persistent cosmic web and its filament structure II. *Monthly Notices of the Royal Astronomical Society*, 414:384–403, 2011. doi:10.1111/j.1365-2966.2011.18395.x.
- 37 J. Stillwell. *Geometry of Surfaces*. Springer-Verlag, New York, 1992.
- 38 W. P. Thurston. Three dimensional manifolds, Kleinian groups, and hyperbolic geometry. *Bull. Amer. Math. Soc.*, 6(3):357–381, 1982. URL: <http://www.ams.org/mathscinet-getitem?mr=648524>.
- 39 W. P. Thurston. *Three Dimensional Geometry and Topology, vol. I*. Princeton University Press, Princeton, New Jersey, 1997.
- 40 R. vd Weygaert, G. Vegter, H. Edelsbrunner, B.J.T. Jones, P. Pranav, C. Park, W. A. Hellwing, B. Eldering, N. Kruithof, E.G.P. Bos, J. Hidding, J. Feldbrugge, E. ten Have, M. v Engelen, M. Caroli, and M. Teillaud. Alpha, Betti and the megaparsec universe: on the homology and topology of the cosmic web. In *Trans. on Comp. Science XIV*, volume 6970 of *LNCS*, pages 60–101. Springer-Verlag, 2011. doi:10.1007/978-3-642-25249-5_3.
- 41 P M. H. Wilson. *Curved Spaces*. Cambridge University Press, Cambridge, 2008.