



HAL
open science

Representing infinite hyperbolic periodic Delaunay triangulations using finitely many Dirichlet domains

Vincent Despré, Benedikt Kolbe, Monique Teillaud

► **To cite this version:**

Vincent Despré, Benedikt Kolbe, Monique Teillaud. Representing infinite hyperbolic periodic Delaunay triangulations using finitely many Dirichlet domains. 2020. hal-03045921v1

HAL Id: hal-03045921

<https://hal.science/hal-03045921v1>

Preprint submitted on 8 Dec 2020 (v1), last revised 6 Dec 2021 (v3)

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

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

Half-minimizers and Delaunay triangulations on closed hyperbolic surfaces

Vincent Despre^{*†}

Benedikt Kolbe^{*‡}

Monique Teillaud^{*§}

Abstract

Given a finite point cloud \mathcal{P} in a Dirichlet fundamental domain \mathcal{F} with respect to a fundamental group Γ of a closed hyperbolic surface S , we derive an explicit way to construct a set of copies of \mathcal{F} that accounts for all points incident to points of \mathcal{P} in the Delaunay triangulation of $\Gamma\mathcal{P}$. We also compute precise bounds on the size of this set that only depend on the genus of S and are thus independent of the hyperbolic metric under consideration. The results in this paper lay the foundations for a practical algorithm to compute Delaunay triangulations on an arbitrary hyperbolic surface, akin to existing implementations for periodic sets of points in Euclidean space.

1 Introduction and motivation

1.1 Closed hyperbolic surfaces, covering spaces, and symmetry groups

The aim of this paper is to lay the foundations for a practical algorithm for the construction of Delaunay triangulations on arbitrary hyperbolic surfaces. We define a hyperbolic surface to be a *closed* orientable topological surface equipped with a hyperbolic metric of constant curvature -1 . It is well-known that a hyperbolic surface S is covered universally by the hyperbolic plane \mathbb{H}^2 , such that the *projection map* $\pi : \mathbb{H}^2 \rightarrow S$ is a local isometry. Including flat tori, all surfaces of interest to us are isometric to the compact quotient space X/Γ with induced metric from X , where Γ is a *symmetry group* of X , i.e. a discrete subgroup of the group of isometries of X . The space X is either \mathbb{E}^2 , or \mathbb{H}^2 , and we call the surface Euclidean or hyperbolic, accordingly. In case $X = \mathbb{H}^2$, the symmetry group is also known as a Non-Euclidean Crystallographic (NEC) group. That S is a hyperbolic surface implies that Γ contains only orientation preserving isometries and has no fixed points in X . The group Γ can be naturally identified with the fundamental group $\pi_1(S)$ of S (after choosing base points appropriately).

1.2 Hyperbolic surfaces in other sciences and nature

One of the motivations for this paper, and also the guiding example we use to illustrate our results and approach is the hyperbolic surface associated to the family of triply-periodic minimal surfaces (TPMS) that contains the gyroid, the primitive, and the diamond surface. A TPMS is a minimal surface in \mathbb{R}^3 that is moreover invariant under three linearly independent translations, i.e. a rank 3 lattice L_3 [43]. To associate to a TPMS a closed surface, one considers the closed surface S_g of genus g embedded in the 3-torus $\mathbb{T}^3 = \mathbb{R}^3/L_3$, which corresponds to glueing the equivalent faces of a unit cell of the TPMS, i.e. the smallest block out of which the TPMS can be reassembled through translations. It turns out that such a surface is always intrinsically hyperbolic [32]. The gyroid, the primitive, and the diamond TPMS are arguably the most prominent and simple [39] examples of TPMS and have received considerable attention in the mathematical, physical, chemical and biological as well as interdisciplinary literature, see, for example, [20, 25, 22, 23, 12]. More recently, TPMS have also found a role in the materials sciences

^{*}Université de Lorraine, CNRS, Inria, LORIA, F-54000 Nancy, France

[†]vincent.despre@loria.fr

[‡]benedikt.kolbe@inria.fr

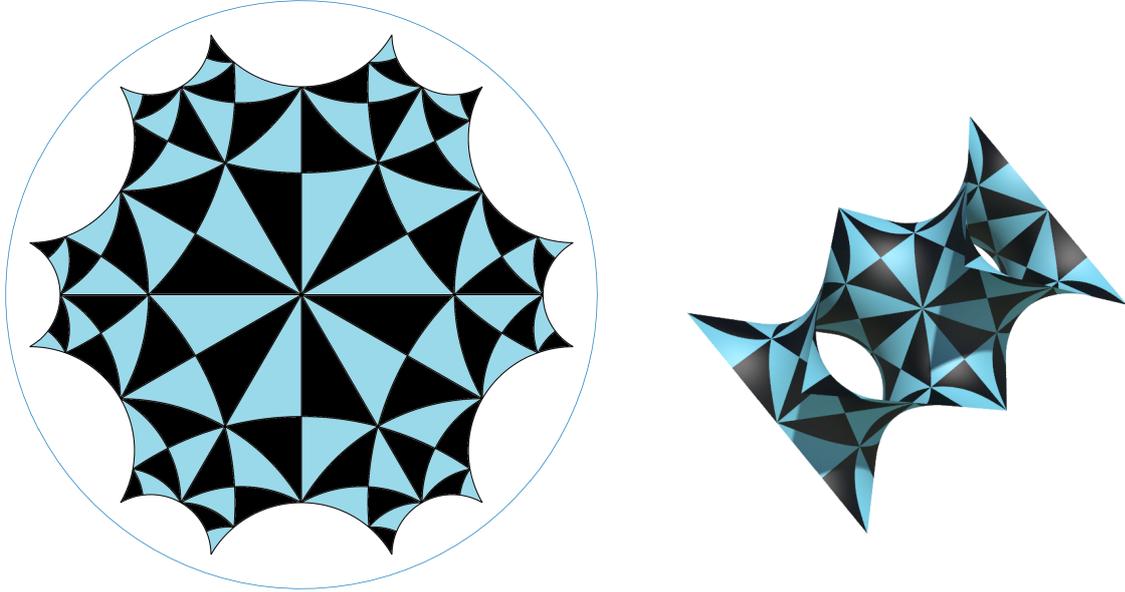
[§]monique.teillaud@inria.fr

The authors were supported by grant ANR-17-CE40-0033 of the French National Research Agency ANR (project SoS) (<https://members.loria.fr/Monique.Teillaud/collab/SoS/>).

December 8, 2020

as a scaffold for crystallographic structures [15, 16], leading to both new mathematical formalisms [27, 26] and a database of such structures [1].

Figure 1 shows a region of \mathbb{H}^2 in the Poincaré disk model and a portion of the diamond surface in \mathbb{R}^3 , illustrating how TPMS are covered by \mathbb{H}^2 . Notice that the angles at which triangles meet are the same in \mathbb{R}^3 as they are in \mathbb{H}^2 , owing to the fact that the covering is conformal.



(a) Fundamental domain for the fundamental group of S_3 in the Poincaré disk model.

(b) A section of the D -surface in \mathbb{R}^3 , together with its smallest asymmetric triangle patches.

Figure 1: The covering of the diamond TPMS by \mathbb{H}^2 .

Very recently, the edge flip algorithm for the computation of a Delaunay triangulation was shown to work for hyperbolic surface [13]. However, robust and efficient software to compute Delaunay triangulations on hyperbolic surfaces, and particularly TPMS, does not exist to date, as far as we know. This paper attempts to help fill this gap by establishing fundamental theoretical results in the hyperbolic case similar to those that have led to such software in the Euclidean setting.

1.3 Set-up and motivation

Given a finite point cloud \mathcal{P} on a hyperbolic surface $S = \mathbb{H}^2/\Gamma$, \mathcal{P} lifts to a locally finite¹ point cloud $\tilde{\mathcal{P}}$ in the covering space \mathbb{H}^2 . The Delaunay triangulation $DT_{\tilde{\mathcal{P}}}$ defined by $\tilde{\mathcal{P}}$ in \mathbb{H}^2 projects to a triangulation $DT_{\mathcal{P}}$ on S , which serves as a definition for the Delaunay triangulation of \mathcal{P} on S [6, 11, 24]. We do not assume triangulations to be embedded simplicial complexes in this paper, in contrast to some previous work [6, 24]. In our setting, every finite point cloud on a hyperbolic surface has an associated locally finite Delaunay triangulation, see [9, Corollary 5.2], [13, Proposition 8].

Using the Poincaré disk model of \mathbb{H}^2 , where Euclidean circles correspond to hyperbolic circles, one sees that the combinatorial structure of a Delaunay triangulation is equivalent to the Euclidean Delaunay triangulation defined by the same set of points in the unit disk. For algorithms to determine the Delaunay triangulation $DT_{\mathcal{P}}$, it is essential to limit oneself to a finite set of points $\tilde{\mathcal{P}}_f \subset \tilde{\mathcal{P}}$ that contains all the information needed to construct the triangulation on S by projecting part of the triangulation of $\tilde{\mathcal{P}}_f$ in \mathbb{H}^2 . Our goal is to deduce bounds for the size of $\tilde{\mathcal{P}}_f$, for a practical algorithm to compute the Delaunay triangulation on an arbitrary hyperbolic surface. Our results are the first of their kind for general hyperbolic surfaces, but Euclidean analogs exist [14], have received attention in the literature in different fields [8, 46, 18], and provide the foundation for the recent and, as far as we know, only implementations of algorithms that compute the Delaunay triangulation of general infinite, locally finite,

¹A collection of points P in a topological space S is locally finite if every point $s \in S$ admits a neighborhood U_s such that $P \cap U_s$ is finite, which implies that if D is a compact disk in S , $P \cap D$ is finite.

periodic point sets in the Euclidean plane [35] (although solutions to special cases do exist). In [35], a set $\tilde{\mathcal{P}}_f$ like above is found for a periodic, infinite set of points $\tilde{\mathcal{P}} \subset \mathbb{E}^2$ with the help of bounds on $\tilde{\mathcal{P}}_f$ that depend only on the structure of the group of translations that generates the periodicity of $\tilde{\mathcal{P}}$. After having found $\tilde{\mathcal{P}}_f$, classical algorithms for the construction of the Delaunay triangulation are used to compute the triangulation $\text{DT}_{\tilde{\mathcal{P}}_f}$ of $\tilde{\mathcal{P}}_f$ in \mathbb{E}^2 , which is then projected to the quotient torus. Note for this that the edges of the convex hull of $\tilde{\mathcal{P}}_f$, contained in $\text{DT}_{\tilde{\mathcal{P}}_f}$, are the only edges of $\text{DT}_{\tilde{\mathcal{P}}_f}$ that are not necessarily a part of $\text{DT}_{\tilde{\mathcal{P}}}$.

Our focus is on generalizing the results of [35, 14] to the more complicated case of hyperbolic surfaces. The proofs in [14] that make possible the algorithms of [35] depend crucially on the Abelian nature of the symmetry groups involved. Since hyperbolic symmetry groups are inherently non-Abelian we need new tools to tackle the problem.

We picture the setup more concretely as follows. Recall that a tessellation or tiling is a countable collection of closed topological disks in a metric space that cover the whole space and whose interiors are disjoint. One can think of the projection map $\pi : \mathbb{H}^2 \rightarrow \mathbb{H}^2/\Gamma = S$ as giving rise to tilings of \mathbb{H}^2 by copies of some fundamental domain for Γ . A *fundamental domain* \mathcal{V} for the action of a symmetry group Γ is a connected subset $\mathcal{V} \subset X$, equal to the closure $\bar{\mathcal{V}}_O$ of a connected set \mathcal{V}_O that intersects each orbit in exactly one point, or, equivalently, such that the restriction of π to \mathcal{V}_O is a bijection from \mathcal{V}_O to S [31]. In $\text{DT}_{\tilde{\mathcal{P}}}$, some points in $\tilde{\mathcal{P}} \cap \mathcal{V}_O$ become adjacent to points inside a translate of \mathcal{V}_O under an element of Γ . We will define a number, the *combinatorial length* of an edge, which relates to the number of copies of \mathcal{V} an edge traverses. The objective of this paper is to find a bound, independent of the hyperbolic metric on S , on the combinatorial distance between any point in \mathcal{V} and its neighbors in $\text{DT}_{\tilde{\mathcal{P}}}$, given a reasonable choice for the tessellation.

We remark that even in the Euclidean case there cannot exist a bound as discussed in the previous paragraph for an arbitrary choice of fundamental domain, even if the Delaunay triangulation is defined by just one point on S . Intuitively speaking, this is because there is no bound on how stretched a fundamental domain can appear for an Abelian symmetry group of \mathbb{E}^2 , necessarily isomorphic to \mathbb{Z}^2 . Assume for simplicity that the group is \mathbb{Z}^2 with its natural action on \mathbb{E}^2 by translations. For ever more long and thin parallelograms that serve as a fundamental domain \mathcal{V} , an edge in a Delaunay triangulation may traverse an unbounded number of copies of \mathcal{V} . We illustrate the situation in Figure 2, which shows one edge \tilde{e} , thick and in cyan, connecting \tilde{a} to \tilde{b} , of the Delaunay triangulation defined by $\mathbb{Z}^2 \cdot \tilde{a}$ in two different situations. The Figure shows \tilde{e} in a portion of two different tilings of \mathbb{E}^2 by parallelograms that serve as fundamental domains for \mathbb{Z}^2 with their boundaries in green. Proposition 5 shows that a similar situation occurs for hyperbolic surfaces. We note that in Euclidean spaces, the celebrated LLL algorithm [28] can be applied to eliminate such problems by finding well-behaved generators of lattices in polynomial time, but this does not work for NEC groups.

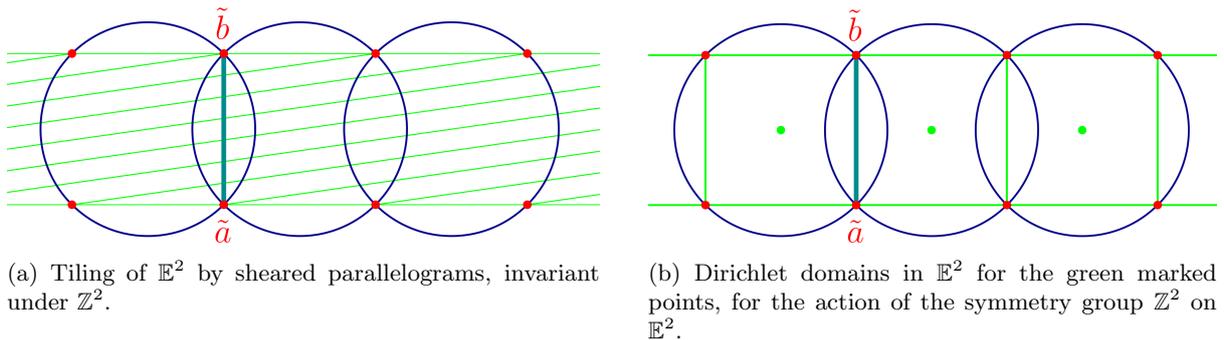


Figure 2: Fundamental domains and the boundary edges, in green, of parts of their associated tessellations in \mathbb{E}^2 . The thick cyan edge is an edge in the Delaunay triangulation defined by $\mathbb{Z}^2 \cdot \tilde{a}$, the set of red points.

Another way to resolve the issue for \mathbb{E}^2 is to choose for \mathcal{V} a Dirichlet fundamental domain $\mathcal{V}_{\tilde{x}}$ for a point \tilde{x} , which corresponds to the Voronoi cell of \tilde{x} in the Voronoi diagram of $\mathbb{Z}^2 \tilde{x}$, as in Figure 2b. In fact, in this case, the shape of the Dirichlet fundamental domain $\mathcal{V}_{\tilde{x}}$ does not depend on the point \tilde{x} and is a unit square, by the Abelianity of the group of translations in \mathbb{E}^2 and the general relation

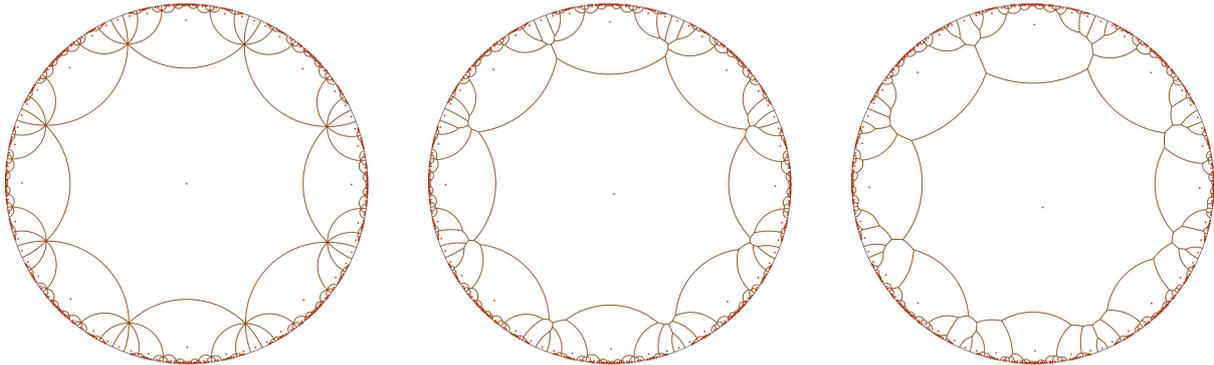


Figure 3: Dirichlet domain w.r.t. different base points. Images from [6, Figure 9].

$\mathcal{V}_{T(\bar{x})}^\Gamma = T(\mathcal{V}_{\bar{x}}^{T^{-1}\Gamma T})$ [3, Section 9.4], where we make explicit the dependence of a Dirichlet domain \mathcal{V}^Γ on the symmetry group Γ , and T is any isometry. One readily sees that for a Delaunay triangulation defined by a single arbitrary point $\tilde{y} \in \mathcal{V}_{\bar{x}}$, it is sufficient to use all copies of $\mathcal{V}_{\bar{x}}$ with nonempty intersection with $\mathcal{V}_{\bar{x}}$ to account for all vertices adjacent to \tilde{y} . Indeed, choosing for \tilde{y} the center of the unit square results in the tiling by squares, shown in Figure 2b, with a choice of diagonals and any other choice of \tilde{y} is a translated version. See [14] for a complete discussion of the case of Euclidean symmetry groups generated by translations.

The main result of this paper is that there is a bound on the number of domains one has to consider to account for all adjacent vertices of a point in a Dirichlet domain in a Delaunay triangulation, for closed hyperbolic surfaces S , depending only on the genus, when we restrict our choice of fundamental domain to Dirichlet fundamental domains; also we explicitly compute a bound. Note that it is far from clear that this works, *a priori*, because the shape, combinatorial structure, and group presentation a Dirichlet fundamental domain for Γ gives rise to depends on the chosen point in its construction in \mathbb{H}^2 , since NEC groups are non-Abelian, in contrast to the above situation in \mathbb{E}^2 . See also Figure 3, which depicts different Dirichlet domains for the same group.

We present a result valid for hyperbolic surfaces that presupposes a metric, but will remain valid for all hyperbolic metrics. Note that the diameter of a Dirichlet fundamental domain for a hyperbolic surfaces can become arbitrarily large, even when the genus is fixed, so there is no bound on the number of copies of Dirichlet fundamental domains an empty disk in a Delaunay triangulation intersects. Together with algorithms that compute a Dirichlet fundamental domain in \mathbb{H}^2 [44], this lays the foundation for a practical algorithm to compute the Delaunay triangulation of a finite set of points on a hyperbolic surface, in the same way as [35].

A hallmark of our approach is that some of its techniques are actually valid in greater generality, for many Riemannian manifolds. In particular, our ideas and approach based on half-minimizers remain valid in a more general setting.

2 Dirichlet fundamental domains and Delaunay triangulations

For the remainder of this paper, S is a hyperbolic surface, with universal cover $\pi : \mathbb{H}^2 \rightarrow S$, and fundamental group Γ , as described in the introduction. To simplify the notation in the following, we denote objects in \mathbb{H}^2 with a tilde, and those on $S = \mathbb{H}^2/\Gamma$ without, \mathcal{P} will always denote a finite set of points on the hyperbolic surface S , and $\tilde{\mathcal{P}}$ its lifted point set to \mathbb{H}^2 . We will again (implicitly) treat \mathbb{H}^2 as a subset of \mathbb{E}^2 , which makes sense when making use of the Poincaré disk model for the hyperbolic plane [4]. This is a conformal model for \mathbb{H}^2 obtained by biconformally mapping \mathbb{H}^2 to the interior of the unit disk in \mathbb{E}^2 such that the conformal mappings of the unit disk correspond exactly to the orientation preserving isometries of \mathbb{H}^2 . This model is well suited for the study of Delaunay triangulations in hyperbolic spaces, because in it, circles in \mathbb{H}^2 correspond to circles in \mathbb{E}^2 .

We briefly recall the relevant definitions for our setup. For more details, see, for example, [7] for the Euclidean case, [11] or [6, 5] for Delaunay triangulations of hyperbolic spaces and [3] for Dirichlet domains. See [37] or [42] for a general introduction to hyperbolic geometry. Let X be a metric space

with distance function d_X . We consider Voronoi diagrams where every cell is topologically a closed disk. For a locally finite point set $\mathcal{P} \subset X$ and $x \in \mathcal{P}$, we denote the Voronoi cell of x by $\mathcal{V}_x^{\mathcal{P}}$ and the whole Voronoi diagram by $\mathcal{V}^{\mathcal{P}}$. In the cases $X = \mathbb{E}^2, \mathbb{H}^2$, the Voronoi diagram is a locally finite collection of convex subsets of X [11, Lemma 5.2]. In particular, if a Voronoi cell is compact, then it only has a finite number of edges.

Definition 1 ([3]). For $X = \mathbb{E}^2, \mathbb{H}^2$, let Γ be an NEC group with no fixed points in X and $\tilde{x} \in X$. The cell $\mathcal{V}_{\tilde{x}}^{\Gamma\tilde{x}}$ of \tilde{x} in the Voronoi diagram of the orbit $\Gamma\tilde{x}$ in X is called the (closed) *Dirichlet fundamental domain* of \tilde{x} , and is denoted as $\mathcal{V}_{\tilde{x}}$ for short.

The Dirichlet fundamental domain $\mathcal{V}_{\tilde{x}}$ can also be defined equivalently as the sets

$$\mathcal{V}_{\tilde{x}} = \{ y \in X \mid d_X(x, y) \leq d_X(x, \Gamma y) \} = \{ y \in X \mid d_X(x, y) \leq d_X(\Gamma x, y) \}$$

where the equality is true since Γ acts as isometries w.r.t. d_X . In particular, we see that

$$\tilde{z} \in \mathcal{V}_{\tilde{x}} \iff \tilde{x} \in \mathcal{V}_{\tilde{z}}. \quad (2.1)$$

Dirichlet domains and more generally Voronoi cells in \mathbb{H}^2 and \mathbb{E}^2 are bounded by geodesics, which is why they are also known as Dirichlet and Voronoi polygons, respectively. A Dirichlet domain is a fundamental domain for Γ . Since S is compact, $\mathcal{V}_{\tilde{x}}$ is compact in \mathbb{H}^2 and therefore, as above, it has only a finite number of edges and the tessellation $\Gamma\mathcal{V}_{\tilde{x}}$ associated to $\mathcal{V}_{\tilde{x}}$, called the *Dirichlet tessellation* w.r.t. \tilde{x} , is a locally finite tessellation.

We now collect relevant definitions and results on Delaunay triangulations. A Delaunay triangulation $\text{DT}_{\tilde{\mathcal{P}}}$ of a locally finite point set $\tilde{\mathcal{P}} \subset \mathbb{H}^2$ is combinatorially a Euclidean Delaunay triangulation with vertex set $\tilde{\mathcal{P}}$, but with geodesic edges. Though we use the term *triangulation*, we consider that more than three cocircular points form a non-triangulated polygonal face. Such faces can always be triangulated in any way when needed. To preserve the duality between Voronoi diagrams and Delaunay triangulations, one removes edges that lead to triangles in $\text{DT}_{\tilde{\mathcal{P}}}$ with non-compact circumcircles, i.e. those not contained in \mathbb{H}^2 . We will see below that this technicality is not necessary for our purposes because all triangulations we consider have compact circumcircles in \mathbb{H}^2 .

The Delaunay triangulation of a finite point set $\mathcal{P} \subset S$ on a surface S is defined as the projection, to S , of the Delaunay triangulation in \mathbb{H}^2 of the lifted point set $\tilde{\mathcal{P}} := \pi^{-1}(\mathcal{P})$, where $\tilde{\mathcal{P}}$ is necessarily locally finite.

A Delaunay edge in \mathbb{H}^2 cannot have a hyperbolic length greater than $2 \text{dia}(S)$, where $\text{dia}(M)$ denotes the diameter of the set M . Indeed, the circumcircle C of a triangle with such an edge has a diameter greater than $2 \text{dia}(S)$, which would cover all of S , since the length of a geodesic in \mathbb{H}^2 and its projection in S are equal [19, Proposition 2.109]. Therefore, there exists at least two points inside C that are equivalent under the action of Γ . In particular, each circumcircle of a cyclic polygon, i.e. a polygon with geodesic edges and concircular vertices in $\text{DT}_{\tilde{\mathcal{P}}}$ is compact [11, 13].

Proposition 2. ([13, Proposition 8],[9, Corollary 5.2]) *The 1-skeleton of the Delaunay triangulation $\text{DT}_{\mathcal{P}}$ on S is an embedded graph on S .*

Definition 3. Let $\mathcal{F} \subset \mathbb{H}^2$ be a fundamental domain for Γ . The edges of \mathcal{F} are identified pairwise under the action of Γ . We fix one representative of each equivalence class of open edges of \mathcal{F} under the action of Γ to obtain a set E of edges and also one representative of each vertex orbit to obtain a set V of vertices. An *original fundamental domain* \mathcal{F}_O associated to \mathcal{F} is a subset of \mathcal{F} consisting of $\text{int}(\mathcal{F}) \cup E \cup V$, where $\text{int}(M)$ denotes the interior of a set M , such that $\overline{E} \cup V$ is connected.

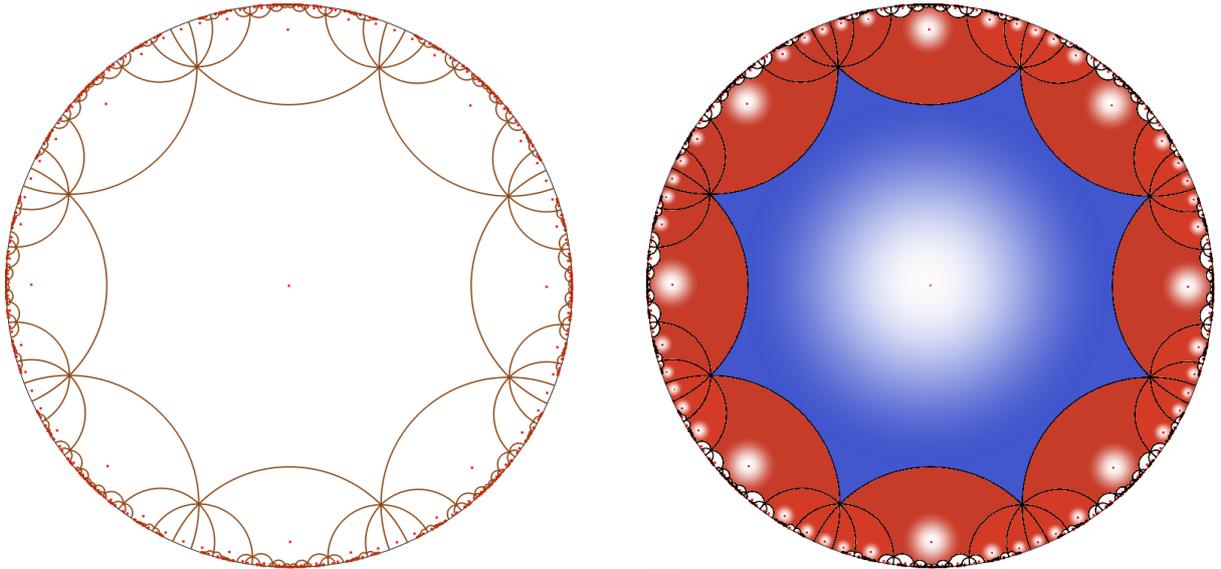
Let $x \in \mathcal{F}_O$ and $y \in \mathbb{H}^2$. Consider the set $\{\mathcal{F}_1^i\}_i$ of copies of \mathcal{F} with $\mathcal{F}_1^i \cap \mathcal{F} \neq \emptyset$ and $\gamma_1^i \in \Gamma$ such that $\gamma_1^i \mathcal{F}_O \subset \mathcal{F}_1^i$. We call $\mathcal{N}_1 := \bigcup_i \gamma_1^i \mathcal{F}_O$ the *first neighborhood layer* of \mathcal{F}_O . If $y \in \mathcal{N}_1$ for some i and $y \notin \mathcal{F}_O$, then y has combinatorial distance from x equal to 1. Repeating this process inductively, to find the n th neighborhood layer \mathcal{N}_n , we find all copies $\{\mathcal{F}_n^i\}_i$ of \mathcal{F} such that $\mathcal{F}_n^i \cap \mathcal{F}_{n-1}^j \neq \emptyset$ for some i, j but such that \mathcal{F}_n^i is not contained in any m th neighborhood layer for $m \leq n-1$. We furthermore find group elements γ_n^i such that $\gamma_n^i \mathcal{F}_O \subset \mathcal{F}_n^i$. Then, if $y \in \mathcal{N}_n$ and $y \notin \mathcal{N}_m$ for $m \leq n-1$, y has combinatorial distance to x equal to $d_{\mathcal{F}_O}^C(x, y) = n$.

The *combinatorial length* of a geodesic segment designates the combinatorial distance between its endpoints. The (maximal) combinatorial length of a triangulation w.r.t. \mathcal{F}_O is $d_{\mathcal{F}_O}^C = \max_{x \sim y} d_{\mathcal{F}_O}^C(x, y)$, where $x \sim y$ if they are adjacent in the triangulation.

Remark 4. Original domains always exist, for any NEC-group, by [29, Theorem 4.1]. Assume that the tessellation induced by \mathcal{F} and Γ features only vertices of degree 3, which is the generic case [3, Theorem 9.4.5]. Then the combinatorial distance between points is equal to the graph distance in the 1-skeleton of the dual tessellation, between the copies of the original domain \mathcal{F}_O that the points lie in.

Special cases of the above definition of original fundamental domains have been considered in the literature [8, 24]. Using an original fundamental domain \mathcal{F}_O , $\text{DT}_{\tilde{\mathcal{P}}}$ can be reconstructed by considering $\tilde{\mathcal{P}} \cap \mathcal{F}_O$ and edges in $\text{DT}_{\tilde{\mathcal{P}}}$ that connect to these.

See Figure 4b for an illustration of the definition of neighborhood layers, which shows a blue colored fundamental domain surrounded by its red colored closed first neighborhood layer. The translations that induce opposite edge identifications of the hyperbolic octagon in the center generate a group Γ_B , the fundamental group of the Bolza surface [24].



(a) Translations that identify opposite edges generate a group Γ_B , under which all vertices of the Dirichlet fundamental domain are equivalent.

(b) The first neighboring layer built around a blue colored Dirichlet fundamental domain in the center.

Figure 4: Dirichlet fundamental domain and tessellation of an NEC group, along with the second neighborhood layer.

Since \mathcal{F} is compact, the combinatorial length of a geodesic segment \tilde{e} is finite if and only if \tilde{e} crosses only a finite number of copies of the fundamental domain \mathcal{F} . In particular, since a Dirichlet tessellation is locally finite, the combinatorial length of a triangulation with locally finite point set and finite vertex degrees, such as the Delaunay triangulation, is always finite.

For a proof of the following proposition, reminiscent of the situation in \mathbb{E}^2 described in the introduction 1 and which motivates the use of Dirichlet fundamental domains for our purposes, see Proposition 25 in the appendix.

Proposition 5. *Let $n \in \mathbb{N}$. There is a fundamental domain for Γ in \mathbb{H}^2 such that there exists an edge in $\text{DT}_{\tilde{\mathcal{P}}}$ the combinatorial length of which is greater than n .*

3 Half-minimizers

This section is dedicated to introducing one of the central concepts in this paper.

Definition 6. A *distance minimizer* between two points on the surface S is a curve that has the minimum length out of all curves on S with the same endpoints.

We call a closed curve c based at a point x in a surface S a *half-minimizer* if the distance in S between x and any point on c is realized as the length of the shortest arc in c that connects these two points.

More generally, a curve c from x to $y \in S$ is a *half-minimizer* if the distance from any point $z \in c$ to x or y is realized by the length of a subarc of c that connects z to x or y .

We also call a curve \tilde{c} in \mathbb{H}^2 a *half-minimizer* if it is the lift of a half-minimizer c in S .

Remark 7. A distance minimizer is necessarily a geodesic on S , but in general geodesics are not distance minimizers on S , since geodesics only locally minimize distances. Furthermore, a distance minimizer contains an arc minimizing the length of all curves that join any two points lying on it. A half-minimizer c between x and y on S necessarily has at least one distinguished point $m \in c$, such that the two distinct arcs between x and m , and y and m are distance minimizers. Indeed, observe that if c_z^x is the subarc of c realizing the distance of $z \in c$ to x , then c_z^x is contained in $c_{z'}^x$ for $z' \in c_z^x$. Since a similar relation holds for subarcs of c that realize the distance to y , c can be expressed as the union of subarcs of the form c_m^x and c_m^y for some point $m \in c$. Moreover, if c is not a distance minimizer between x and y , then, by Lemma 10 below, the point m is uniquely determined. We shall call m a *half-point*. If c is a closed half-minimizer, the half-point m coincides with the midpoint of c . A half-minimizer is smooth except that it may have a kink at m and, if closed, also at the base point, so it makes sense to call the two geodesic parts of a half-minimizer *halves*. See the violet curve from x_1 to y_1 in Figure 5 for an example of a non-closed half-minimizer and the blue curves for closed half-minimizers, where the curve based at x_2 has two kinks, one at x_2 and the other at m_2 . Note that the curve based at x_4 has no kinks and one can easily imagine the base point in this case to not be significant, as the curve is a half-minimizer based at all points lying on it. However, the closed geodesic based at x_3 illustrates that there may be half-minimizers that are smooth everywhere but do not allow the base point to be chosen freely, if, for example, the green curve between x'_3 and m'_3 depicts a distance minimizer.

The concept of closed half-minimizers has been studied before in the literature [2], in different contexts and using a more restrictive definition, under the name half-geodesic.

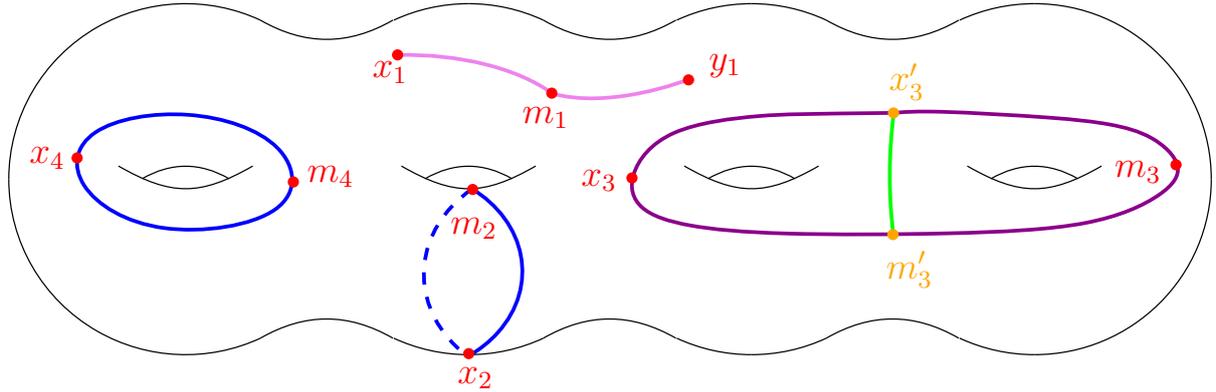


Figure 5: Examples and counter-examples of half-minimizers.

Lemma 8. (Lemma 28 in appendix) Let $\mathcal{V}_{\tilde{x}}$ be a Dirichlet domain. Then the geodesic segment \tilde{g} that joins \tilde{x} to any $\tilde{y} \in \mathcal{V}_{\tilde{x}}$ projects to a distance minimizer in S . Conversely, if g from $x = \pi(\tilde{x})$ to y is a distance minimizer in S , then the lift \tilde{g} of g based at \tilde{x} satisfies $\tilde{g} \subset \mathcal{V}_{\tilde{x}}$.

As a consequence of Lemma 8, we see that two geodesics c_1, c_2 in $\mathcal{V}_{\tilde{x}}$, both incident to \tilde{x} , give rise to a half-minimizer c by concatenation.

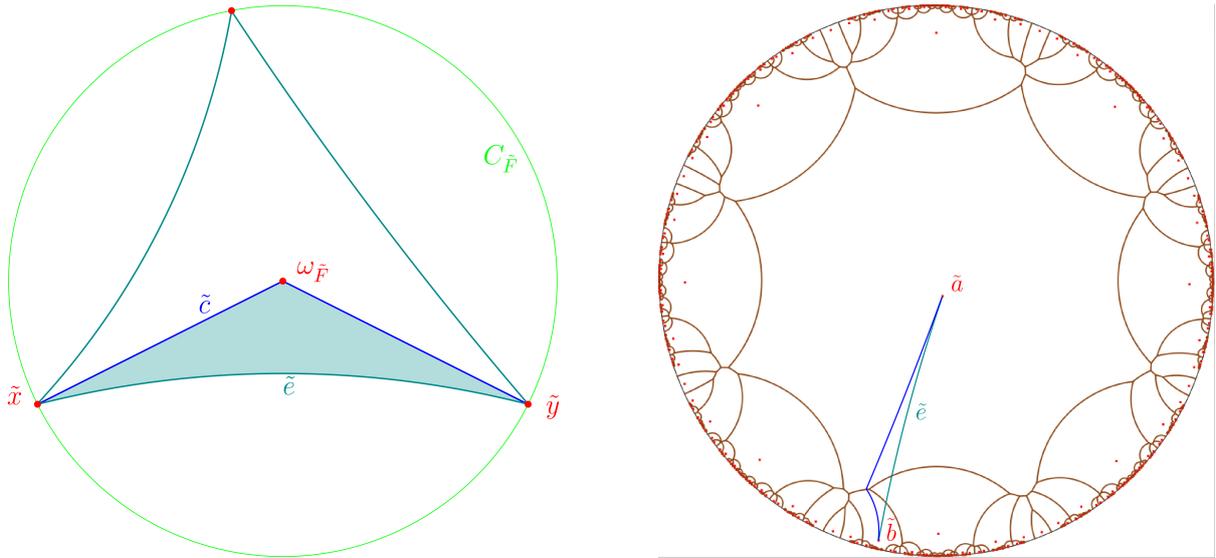
Recall that two curves $c_1, c_2 \subset \mathbb{H}^2$ are Γ -equivariantly isotopic if and only if their projections are isotopic in $S = \mathbb{H}^2/\Gamma$. The relevance of half-minimizers for Delaunay triangulations comes from the following proposition.

Proposition 9. Let $\tilde{x} \in \tilde{\mathcal{P}}$. Then any edge \tilde{e} in $\text{DT}_{\tilde{\mathcal{P}}}$ based at \tilde{x} is Γ -equivariantly isotopic, with fixed endpoints, to a curve \tilde{c} that is a half-minimizer based at \tilde{x} . The projection c of \tilde{c} to S is simple and, if closed, nontrivial.

Before we present a proof, we discuss the situation. Note that it is possible in general that an edge \tilde{e} in $\text{DT}_{\tilde{\mathcal{P}}}$ is not itself a half-minimizer. Figure 6 provides an example of this. The figure shows a Dirichlet

tessellation from a Voronoi diagram, constructed from the red points, which are all equivalent under the action of the group Γ generated by the translations that induce opposite edge identifications of the hyperbolic octagon in Figure 4a above. Since every vertex has degree 3, the dual tessellation $\text{DT}_{\tilde{\mathcal{P}}}$ is a triangulation and the geodesic edge \tilde{e} shown in Figure 6b in cyan connecting the points \tilde{a} and \tilde{b} is an edge of $\text{DT}_{\tilde{\mathcal{P}}}$, with a combinatorial length of 1. Moreover, \tilde{e} intersects two edges in the Voronoi diagram in their interiors.

In general, if an edge \tilde{e}' in a Delaunay tessellation incident to \tilde{x} and some copy of \tilde{x} is a closed half-minimizer, then it follows straight from the definition of $\mathcal{V}_{\tilde{x}}$ and Lemma 8 that the midpoint of \tilde{e}' must be contained in $\mathcal{V}_{\tilde{x}}$. Said in another way, such a half-minimizer intersects the interior of exactly two Dirichlet domains. Therefore, the edge \tilde{e} in $\text{DT}_{\tilde{\mathcal{P}}}$ cannot be a half-minimizer. On the other hand, the curve in blue in Figure 6b, connecting the points \tilde{a} and \tilde{b} , is a half-minimizer, isotopic to the cyan curve, by Proposition 9.



(a) The isotopy in the proof of Proposition 9 moves the edge \tilde{e} over the shaded region to the curve \tilde{c} .

(b) A generic Dirichlet fundamental domain and a geodesic, in cyan, connecting points \tilde{a} and \tilde{b} , isotopic to the half-minimizer in blue.

Figure 6: Dirichlet domains and tessellations for the fundamental group of the Bolza surface.

With that, we turn to the proof of Proposition 9.

Proof. Consider the curve \tilde{c} that connects the two vertices \tilde{x} and \tilde{y} of a triangular face \tilde{F} in $\text{DT}_{\tilde{\mathcal{P}}}$ consisting of the two geodesic segments from the center $\omega_{\tilde{F}}$ of the circumcircle $C_{\tilde{F}}$ of \tilde{F} to either endpoint. The situation is illustrated in Figure 6a. The point $\omega_{\tilde{F}}$ lies in $\mathcal{V}_{\tilde{x}}$ (by Lemma 26 in Appendix B). Together with Lemma 8, this implies that \tilde{c} is a half-minimizer based at either endpoint, with half-point $\omega_{\tilde{F}}$. Consider the Dirichlet fundamental domain $\mathcal{V}_{\omega_{\tilde{F}}}$, which contains \tilde{x} and \tilde{y} , by (2.1). Since $\mathcal{V}_{\omega_{\tilde{F}}}$ is convex, it also contains the geodesic edge \tilde{e} in \tilde{F} joining \tilde{x} and \tilde{y} . Now, \tilde{c} and \tilde{e} form a bigon B completely contained in $\mathcal{V}_{\omega_{\tilde{F}}}$, so there is an isotopy inside every copy of $\mathcal{V}_{\omega_{\tilde{F}}}$ from the copies of \tilde{e} to \tilde{c} , which fixes points equivalent to \tilde{x} and \tilde{y} . Note also that B does not contain any points of $\tilde{\mathcal{P}}$ by definition of $\text{DT}_{\tilde{\mathcal{P}}}$ and convexity of $C_{\tilde{F}}$. Since $\mathcal{V}_{\omega_{\tilde{F}}}$ contains only one representative of every point moved in the isotopy and the tessellation of \mathbb{H}^2 by copies of $\mathcal{V}_{\omega_{\tilde{F}}}$ is Γ -invariant, one obtains a Γ -invariant isotopy of \mathbb{H}^2 .

That c , if closed, is nontrivial is clear, and simple follows from Proposition 2 for $\text{DT}_{\tilde{\mathcal{P}}}$ and the fact that c is isotopic to the projection of \tilde{c} . \square

The situation illustrated in Figure 6 motivates the following definition. The Voronoi diagram dual to a Delaunay tessellation $\text{DT}_{\tilde{\mathcal{P}}}$ has edges joining the circumcenters of the polygons in $\text{DT}_{\tilde{\mathcal{P}}}$. We call an edge of either tessellation that intersects its dual edge a *centered edge*. This makes sense because the dual of the dual of $\text{DT}_{\tilde{\mathcal{P}}}$ is the same tessellation of \mathbb{H}^2 as $\text{DT}_{\tilde{\mathcal{P}}}$, and the same is true for $\mathcal{V}_{\tilde{\mathcal{P}}}$, so an edge is

centered if and only if its dual is. The concept of centered edges has been studied before [9], restricted to edges of a Voronoi diagram.

In the remainder of this section we study the number of intersections that two half-minimizers can have.

Lemma 10. *(Lemma 29 in appendix) Two distinct distance minimizers that are not subarcs of one another on S cannot intersect each other more than once in their interior. Moreover, if an intersection occurs at an endpoint, then there cannot be an intersection in the interior.*

Remark 11. Note that it is possible for two distance minimizers to intersect more than once, at endpoints, as the existence of closed half-minimizer shows.

Lemma 12. *(Lemma 30 in appendix) A half-minimizer can intersect a distance minimizer at most 2 times, or at least half of the half-minimizer is a subarc of the distance minimizer.*

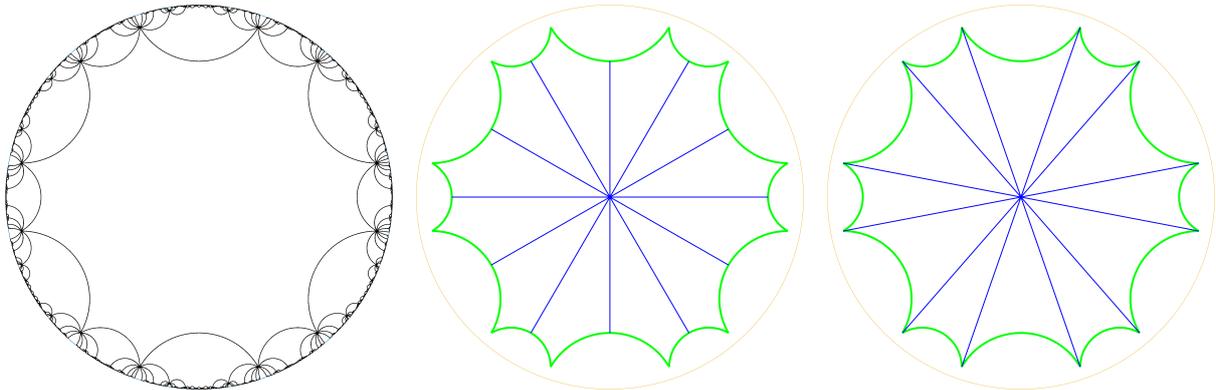
Lemma 13. *(Lemma 31 in appendix) Two half-minimizers cannot intersect each other more than 4 times, as long as their halves are distinct.*

We illustrate the machinery developed thus far in the context of Delaunay triangulations of the primitive, diamond, and gyroid triply-periodic minimal surface (TPMS) family discussed in the introduction.

Let S_3 denote the closed hyperbolic surface that gives rise to the gyroid TPMS family by lifting, to \mathbb{R}^3 , embeddings of S_3 into the three-torus \mathbb{T}^3 . Consider the fundamental polygon \mathcal{D} containing the origin 0 and its tessellation in Figure 7a, under the symmetry group Γ_{S_3} generated by the translations that identify opposite edges of \mathcal{D} . It is well-known that \mathcal{D} is a fundamental domain for Γ_{S_3} and that $\mathbb{H}^2 \rightarrow \mathbb{H}^2/\Gamma_{S_3} = S_3$ corresponds to the universal (conformal) covering of S_3 [38].

Lemma 14. *(Lemma 32 in appendix) The dodecagon \mathcal{D} depicted in Figure 1a is the Dirichlet fundamental domain $\mathcal{V}_0 \subset \mathbb{H}^2$ for Γ_{S_3} .*

It is straightforward to see that the boundary of \mathcal{D} consists of 6 inequivalent edges and that all corners are equivalent, under Γ_{S_3} , since Γ_{S_3} identifies opposite edges of \mathcal{D} . Therefore, by Lemma 14 and Lemma 8, the straight lines through the origin depicted in Figure 7c in blue also give rise to closed half-minimizers based at the origin. Moreover, by Lemma 8, one can combine any two of the line segments l_i from the origin to the corners, to obtain many more half-minimizers based at the origin. Denote the set of these line segments by \mathcal{L} and the set of all possible half-minimizers that can be formed from \mathcal{L} by \mathcal{L}_H .



(a) Tessellation of \mathbb{H}^2 by Dirichlet dodecagons.

(b) The Dirichlet dodecagon along with lines, in blue, through the origin, representing the invariant axes of hyperbolic translations.

(c) The Dirichlet dodecagon, in green, along with line segments that are distance minimizers emanating from the origin, in blue.

Figure 7: Half-minimizers in the Dirichlet dodecagon, corresponding to the Dirichlet fundamental domain of the genus 3 hyperbolic surface that gives rise to the gyroid TPMS family.

By Proposition 9 and Lemma 13, no Delaunay edge can project to a curve in S_3 that is homotopic to a half-minimizer that intersects any other half-minimizer more than 4 times. So, to find a bound on

the combinatorial length of Delaunay edges, we can count intersections with half-minimizers in \mathcal{L}_H . To see if a geodesic segment \tilde{e} emanating from \mathcal{D} connecting two points of some $\tilde{\mathcal{P}}$ is potentially a Delaunay edge, we assign to each line segment in \mathcal{L} the numbers of intersections, in its interior, with \tilde{e} . Then, the sum σ of any two of these values has to be less than five if the intersections are nontrivial, i.e. cannot be eliminated by homotopies. Note that the origin and corners of \mathcal{D} can only be joint intersection points of the curves in \mathcal{L} with \tilde{e} , so can only add 1 each to each sum σ . When counting nontrivial intersections, it is well-known that geodesic representatives of two curves on a hyperbolic surface minimize the number of intersections these two curves can have in their homotopy classes (with fixed endpoints), see [17, Propositions 1.3 and 1.9] and [17, Section 1.2.7]. Observe that the blue curves in Figure 7c cut S_3 into a set of topological disks $\{D_i\}_i$. Therefore, for every curve emanating from \mathcal{D} , there is a maximal combinatorial length it can attain before violating the restriction on the sum.

There is a maximal number of edges of a Dirichlet fundamental domain for the fundamental group of a surface of given genus [3, Theorem 10.5.1] and therefore also for the number of disks $\{D_i\}$. Therefore, with the above, it is clear that half-minimizers on hyperbolic surfaces lead to a bound on the combinatorial length of Delaunay edges, independent of the hyperbolic metric on the surface. To improve the results that the above method would generate, our strategy is to find as many distance minimizers within a given Dirichlet fundamental domain as possible.

Proposition 15. *(Proposition 37 in appendix) Each centered edge of $\mathcal{V}_{\tilde{x}}$ is a half-minimizer and each non-centered edge a distance minimizer.*

For the following, we recall an extended definition of convex to apply to regions on S .

Definition 16. We call a subset $K \subset S$ *convex* if for every two points $x, y \in K$, there is a unique distance minimizer $c \subset K$ with endpoints x and y .

Note that not all convex subsets of \mathbb{H}^2 project to convex subsets of S . Furthermore, a subset A of S with the property that there is a unique geodesic in A that joins any two of its points is also not a convex set in S in general. In fact, the interior of every Dirichlet fundamental domain projected to S provides an example of such behavior.

Proposition 17. *(Proposition 42 in appendix) A (closed) triangle in S with edges that are distance minimizers is convex.*

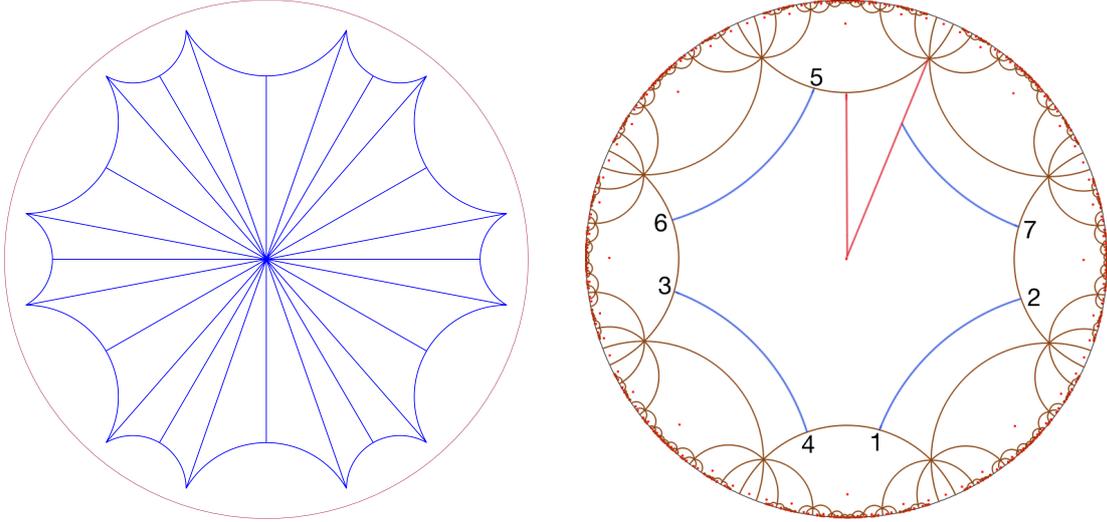
Remark 18. The generalization of Proposition 17 to polygons with more edges does not hold. For a counter-example, consider a Dirichlet fundamental domain $\mathcal{V}_{\tilde{x}}$ for some NEC group Γ with only one vertex orbit and connect \tilde{x} with two vertices incident to the same edge \tilde{e} of $\mathcal{V}_{\tilde{x}}$ by geodesics. By splitting \tilde{e} along its half-point, which exists by Proposition 15, we obtain a convex quadrilateral \tilde{Q} in \mathbb{H}^2 with edges that are distance minimizers w.r.t. $\tilde{\mathcal{P}} = \Gamma\tilde{x}$. However, as two of its vertices project to the same point on $S = \mathbb{H}^2/\Gamma$, by construction, so Q cannot be convex in S . It is easy to see that shifting Q around a bit on S , one can also obtain examples of this type without equivalent vertices.

4 A bound on the combinatorial length of Delaunay edges

In this section, we prove the main results of this paper: bounds on the combinatorial length of Delaunay triangulations of point sets $\tilde{\mathcal{P}} = \pi^{-1}(\mathcal{P})$, where \mathcal{P} is finite.

We start by summarizing the situation thus far. By Lemma 8, Proposition 15 and Proposition 17, every Dirichlet fundamental domain for Γ can be partitioned into a set of triangles each of whose projection to $S = \mathbb{H}^2/\Gamma$ is convex, as illustrated for a particular Dirichlet fundamental domain for the origin, in Figure 8a. For other symmetry groups Γ or other Dirichlet domains, the situation may be slightly different as triangles incident to one edge of the Dirichlet polygon are not necessarily congruent. Furthermore, Figure 8a shows the worst possible case in terms of the number of triangles, because in case of the presence of non-centered edges, the triangle formed by the edge and the origin is already convex. By Lemma 10, the intersection of a distance minimizer on S and each triangle is necessarily connected.

We first examine the recently considered, more restricted case [24, 6, 35], where the edges in the Delaunay triangulation are distance minimizers on S . The condition adopted in these works is that no two points in \mathcal{P} that are connected by an edge in $\text{DT}_{\tilde{\mathcal{P}}}$ are further apart than a distance that is smaller



(a) Every blue arc in the Dirichlet domain projects to a distance minimizer. (b) A distance minimizer that maximizes the combinatorial distance illustrated for the Bolza surface.

Figure 8: Distance minimizers in Dirichlet domains.

than half the length of the smallest noncontractible loop on S . This means that any path in S joining points x and y in \mathcal{P} not homotopic to the edge e in $\text{DT}_{\tilde{\mathcal{P}}}$ joining these points is strictly longer than e . This implies that the edges in $\text{DT}_{\tilde{\mathcal{P}}}$ are distance minimizers.

Theorem 19. (Theorem 45 in appendix) Let Γ be the fundamental group of a closed hyperbolic surface S . Let $\mathcal{V}_{\bar{x}}$ be a Dirichlet fundamental domain for Γ , with k edges and $(\mathcal{V}_{\bar{x}})_O \subset \mathcal{V}_{\bar{x}}$ any original fundamental domain. Let \mathcal{P} be a finite set of points on S and $\text{DT}_{\tilde{\mathcal{P}}}$ the Delaunay triangulation of the lifted set of points $\tilde{\mathcal{P}} \subset \mathbb{H}^2$, with edges that are distance minimizers. Then the maximal combinatorial length of a Delaunay edge is bounded by $k/2$.

A rigorous and detailed proof can be found in the appendix, Section C. Here, we merely present a sketch and explain the idea.

By the discussion before the statement of the theorem, a Dirichlet domain $\mathcal{V}_{\bar{x}}$ can be decomposed into triangles $\{\tilde{\Delta}_j\}_{j=1}^{2k}$ as shown in Figure 8a. Since the projection of each of the triangles is convex in S , the intersection of a distance minimizer $e \subset S$ and each Δ_j has to be connected. The edge e starts at some point in $\mathcal{V}_{\bar{x}}$ and may exit it. We consider the part that lies in $\mathcal{V}_{\bar{x}}$ as the first geodesic segment \tilde{c}_1 . By using an appropriate element of $\gamma \in \Gamma$, one can map the first part of e that lies outside $\mathcal{V}_{\bar{x}}$ back into $\mathcal{V}_{\bar{x}}$ to obtain another geodesic segment, which may again exit $\mathcal{V}_{\bar{x}}$. Repeating this process of mapping the beginning part of the geodesic segment that lies outside $\mathcal{V}_{\bar{x}}$ back into $\mathcal{V}_{\bar{x}}$ yields a collection $\{\tilde{c}_i\}_{i=1}^r$ of pairwise disjoint geodesic arcs in $\mathcal{V}_{\bar{x}}$ that project to e , as illustrated in Figure 8b. To see what the maximal possible combinatorial length is, we try to maximize it using a maximal collection of \tilde{c}_i that, when added iteratively to the others in order, each increase the combinatorial length by the maximum possible value of 1. For this, \tilde{c}_i cannot join two neighboring edges of $\mathcal{V}_{\bar{x}}$. Therefore, each \tilde{c}_i has to intersect at least 4 of the $\tilde{\Delta}_j$ triangles. Thus, ignoring possible issues with the first and last geodesic segments and the corners of $\mathcal{V}_{\bar{x}}$, the number r of arcs \tilde{c}_i , and therefore the combinatorial distance, is at most $2k/4 = k/2$.

Corollary 20. Let $\mathcal{V}_{\bar{x}}$ be a Dirichlet fundamental domain for an NEC group $\Gamma = \pi_1(S)$, with S a genus g hyperbolic surface. Let $\text{DT}_{\tilde{\mathcal{P}}}$ be a Delaunay triangulation with edges that are distance minimizers, invariant under Γ . Then the maximal combinatorial length is bounded by $6g - 3$.

Proof. The number of edges of a Dirichlet fundamental domain is bounded by $12g - 6$ from above [3, Theorem 10.5.1]. Note also that the case of $12g - 6$ edges corresponds to the generic situation where the Delaunay tessellation dual to $\mathcal{V}^{\Gamma x}$ is a triangulation. Theorem 19 completes the proof. \square

We now treat the case of general Delaunay triangulations.

Corollary 21. *Let $\mathcal{V}_{\tilde{x}}$ be a Dirichlet fundamental domain, with k edges, for an NEC group $\Gamma = \pi_1(S)$. Let $\text{DT}_{\tilde{\mathcal{P}}}$ be an arbitrary Delaunay triangulation, invariant under Γ . Then the maximal combinatorial length is bounded by k .*

Proof. Since a Delaunay edge \tilde{e} in \mathbb{H}^2 is Γ -equivariantly homotopic to a half-minimizer \tilde{e}_H by Proposition 9, we will use Theorem 19 as follows. Given a starting point $\tilde{a} \in \mathcal{V}_{\tilde{x}}$ for a curve, its homotopy class on the surface is uniquely determined by which copy of the original fundamental domain $(\mathcal{V}_{\tilde{x}})_O$ contains its endpoint \tilde{b} . This is a consequence of the fact that the geodesic from \tilde{a} and \tilde{b} in \mathbb{H}^2 is unique and this projects to a geodesic on the surface which is also unique in its homotopy class [17, Section 1.2.7]. This means that if \tilde{b} lies in a copy of $(\mathcal{V}_{\tilde{x}})_O$ outside of where the concatenation of two distance minimizers can reach, then the geodesic between \tilde{a} and \tilde{b} cannot be homotopic to a half-minimizer. \square

Corollary 22. *Let $\mathcal{V}_{\tilde{x}}$ be a Dirichlet fundamental domain for an NEC group $\Gamma = \pi_1(S)$, with S a genus g hyperbolic surface. Let $\text{DT}_{\tilde{\mathcal{P}}}$ be an arbitrary Delaunay triangulation, invariant under Γ . Then the maximal combinatorial distance is bounded by $12g - 6$.*

Remark 23. For a Dirichlet tessellation of \mathbb{E}^2 of the group \mathbb{Z}^2 , with squares, Corollary 21 yields a bound on the combinatorial length of 4, whereas the optimal bound is 2.

Together with the algorithm from [44] to compute a Dirichlet fundamental domain for hyperbolic symmetry groups, the results of this section lay the foundations for implementations of practical algorithms to construct the Delaunay triangulation $\text{DT}_{\mathcal{P}}$ for a finite set of points on an arbitrary hyperbolic surface. We exhibited the first known bound for the set of points \mathcal{P}_f in \mathbb{H}^2 one has to consider to guarantee that there is at least one representative edge in $\text{DT}_{\tilde{\mathcal{P}}_f}$ for every edge in $\text{DT}_{\tilde{\mathcal{P}}}$. To find $\tilde{\mathcal{P}}_f$, it suffices to take into account as many copies of a given Dirichlet fundamental domain as dictated by the bounds on the combinatorial length found in this section. Note that the bound of n on the combinatorial length of a Dirichlet domain of a genus g surface translates to a bound of order $O((12g)^n)$ for the number of copies of a generic Dirichlet fundamental domain with $12g - 6$ edges.

References

- [1] Epinet. Accessed: 2020-10-23. URL: <http://epinet.anu.edu.au>.
- [2] Ian M. Adelstein and Benjamin Schmidt. Characterizing round spheres using half-geodesics. *Proceedings of the National Academy of Sciences of the United States of America*, 116(29):14501–14504, 2019. doi:10.1073/pnas.1907226116.
- [3] A.F. Beardon. *The Geometry of Discrete Groups*. 3Island Press, 1983.
- [4] Marcel Berger. *Geometry 2*. Springer, 2009.
- [5] Mikhail Bogdanov, Olivier Devillers, and Monique Teillaud. Hyperbolic Delaunay complexes and Voronoi diagrams made practical. *Journal of Computational Geometry*, 5:56–85, 2014. doi:10.20382/jocg.v5i1a4.
- [6] Mikhail Bogdanov, Monique Teillaud, and Gert Vegter. Delaunay Triangulations on Orientable Surfaces of Low Genus. In Sándor Fekete and Anna Lubiw, editors, *32nd International Symposium on Computational Geometry (SoCG 2016)*, volume 51 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 20:1–20:17, Dagstuhl, Germany, 2016. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik. doi:10.4230/LIPIcs.SoCG.2016.20.
- [7] Jean-Daniel Boissonnat and Mariette Yvinec. *Algorithmic Geometry*. Cambridge University Press, 1998. doi:10.1017/CB09781139172998.
- [8] Manuel Caroli and Monique Teillaud. Delaunay Triangulations of Closed Euclidean d-Orbifolds. *Discrete & Computational Geometry*, 55(4):827–853, jun 2016. URL: <https://hal.inria.fr/hal-01294409>, doi:10.1007/s00454-016-9782-6.
- [9] Jason DeBlois. The centered dual and the maximal injectivity radius of hyperbolic surfaces. *Geometry and Topology*, 19(2):953–1014, 2015. doi:10.2140/gt.2015.19.953.

- [10] Jason DeBlois. The geometry of cyclic hyperbolic polygons. *Rocky Mountain J. Math.*, 46(3):801–862, 06 2016. doi:10.1216/RMJ-2016-46-3-801.
- [11] Jason DeBlois. The Delaunay tessellation in hyperbolic space. *Mathematical Proceedings of the Cambridge Philosophical Society*, 164(1):15–46, 2018. doi:10.1017/S0305004116000827.
- [12] Yuru Deng and Mark Mieczkowski. Three-dimensional periodic cubic membrane structure in the mitochondria of amoebae *Chaetos carolinensis*. *Protoplasma*, 203:16–25, 1998. doi:10.1007/BF01280583.
- [13] Vincent Despré, Jean-Marc Schlenker, and Monique Teillaud. Flipping Geometric Triangulations on Hyperbolic Surfaces. In Sergio Cabello and Danny Z. Chen, editors, *36th International Symposium on Computational Geometry (SoCG 2020)*, volume 164 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 35:1–35:16, Dagstuhl, Germany, 2020. Schloss Dagstuhl–Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.SoCG.2020.35.
- [14] Nikolai P. Dolbilin and Daniel H. Huson. Periodic Delone Tilings. In *Periodica Mathematica Hungarica*, volume 34, pages 57–64, 1997. doi:10.1023/A:1004272423695.
- [15] Myfanwy E. Evans, Vanessa Robins, and Stephen T. Hyde. Periodic entanglement I: Networks from hyperbolic reticulations. *Acta Crystallographica Section A: Foundations of Crystallography*, 69(3):241–261, 2013. doi:10.1107/S0108767313001670.
- [16] Myfanwy E. Evans, Vanessa Robins, and Stephen T. Hyde. Periodic entanglement II: Weavings from hyperbolic line patterns. *Acta Crystallographica Section A: Foundations of Crystallography*, 69(3):262–275, 2013. doi:10.1107/S0108767313001682.
- [17] Benson Farb and Dan Margalit. *A Primer on Mapping Class Groups (PMS-49)*. Princeton University Press, 2012. URL: <http://www.jstor.org/stable/j.ctt7rkjw>.
- [18] Eric L. First, Chrysanthos E. Gounaris, James Wei, and Christodoulos A. Floudas. Computational characterization of zeolite porous networks: an automated approach. *Phys. Chem. Chem. Phys.*, 13:17339–17358, 2011. doi:10.1039/C1CP21731C.
- [19] Sylvestre Gallot, Dominique Hulin, and Jacques Lafontaine. *Riemannian geometry*. Springer, 3 edition, 2004.
- [20] Karsten Große-Brauckmann and Meinhard Wohlgemuth. The gyroid is embedded and has constant mean curvature companions. *Calculus of Variations and Partial Differential Equations*, 1996. doi:10.1007/BF01261761.
- [21] Allen Hatcher. *Algebraic Topology*, volume 44. Cambridge University Press, 2002.
- [22] S. T. Hyde, S. Ramsden, T. Di Matteo, and J. J. Longdell. Ab-initio construction of some crystalline 3D Euclidean networks. *Solid State Sciences*, 5(1):35–45, 2003. doi:10.1016/S1293-2558(02)00079-1.
- [23] Stephen T. Hyde, T. Landh, S. Lidin, B.W. Ninham, K. Larsson, and S. Andersson. *The Language of Shape*. Elsevier Science, 1996.
- [24] Jordan Jordanov and Monique Teillaud. Implementing Delaunay Triangulations of the Bolza Surface. In *33rd International Symposium on Computational Geometry (SoCG 2017)*, pages 44:1 – 44:15, Brisbane, Australia, July 2017. doi:10.4230/LIPIcs.SoCG.2017.44.
- [25] Jacob J K Kirkensgaard, Myfanwy E Evans, Liliana de Campo, and Stephen T Hyde. Hierarchical self-assembly of a striped gyroid formed by threaded chiral mesoscale networks. *Proceedings of the National Academy of Sciences of the United States of America*, 111(4):1271–6, 2014. doi:10.1073/pnas.1316348111.
- [26] Benedikt Kolbe and Myfanwy Evans. Enumerating isotopy classes of tilings guided by the symmetry of triply-periodic minimal surfaces. 2020. Under review.

- [27] Benedikt Kolbe and Myfanwy Evans. Isotopic tiling theory for hyperbolic surfaces. *Geometriae Dedicata*, 2020. doi:10.1007/s10711-020-00554-2.
- [28] A. K. Lenstra, H. W. Lenstra, and L. Lovász. Factoring polynomials with rational coefficients. *Mathematische Annalen*, 261(4):515–534, 1982. doi:10.1007/BF01457454.
- [29] Z. Lučić and E. Molnár. Fundamental domains for planar discontinuous groups and uniform tilings. *Geometriae Dedicata*, 40(2):125–143, Nov 1991. doi:10.1007/BF00145910.
- [30] A. M. Macbeath. The classification of non-euclidean plane crystallographic groups. *Canadian Journal of Mathematics*, 19:1192–1205, 1967. doi:10.4153/CJM-1967-108-5.
- [31] William S. Massey. *A Basic Course in Algebraic Topology*. Graduate Texts in Mathematics. Springer, 1991.
- [32] William H. Meeks. *The Theory of Triply Periodic Minimal Surfaces*, 1990. doi:10.1512/iumj.1990.39.39043.
- [33] John Meier. *Groups, Graphs and Trees: An Introduction to the Geometry of Infinite Groups*. London Mathematical Society Student Texts. Cambridge University Press, 2008. doi:10.1017/CB09781139167505.
- [34] Barrett O’Neill. *Semi-Riemannian geometry : with applications to relativity*. Academic Press, 1983.
- [35] Georg Osang, Mael Rouxel-Labbé, and Monique Teillaud. Generalizing CGAL Periodic Delaunay Triangulations. In Fabrizio Grandoni, Grzegorz Herman, and Peter Sanders, editors, *28th Annual European Symposium on Algorithms (ESA 2020)*, volume 173 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 75:1–75:17, Dagstuhl, Germany, 2020. Schloss Dagstuhl–Leibniz-Zentrum für Informatik. Best paper award (Track B: Engineering and Applications). doi:10.4230/LIPIcs.ESA.2020.75.
- [36] S. J. Ramsden, V. Robins, and S. T. Hyde. Three-dimensional Euclidean nets from two-dimensional hyperbolic tilings: Kaleidoscopic examples. *Acta Crystallographica Section A: Foundations of Crystallography*, 65(2):81–108, 2009. doi:10.1107/S0108767308040592.
- [37] J. Ratcliffe. *Foundations of Hyperbolic Manifolds*. Graduate Texts in Mathematics. Springer New York, 2006.
- [38] J. -F Sadoc and J. Charvolin. Infinite periodic minimal surfaces and their crystallography in the hyperbolic plane. *Acta Crystallographica Section A*, 45(1):10–20, 1989. doi:10.1107/S0108767388008438.
- [39] G. E. Schröder-Turk, A. Fogden, and S. T. Hyde. Bicontinuous geometries and molecular self-assembly: Comparison of local curvature and global packing variations in genus-three cubic, tetragonal and rhombohedral surfaces. *European Physical Journal B*, 54(4):509–524, 2006. doi:10.1140/epjb/e2007-00025-7.
- [40] George Springer. *Introduction to Riemann Surfaces*. Addison-Wesley Publishing Company, Inc., 1957.
- [41] A. Szczepański. *Geometry of Crystallographic Groups*. Algebra and discrete mathematics. World Scientific, 2012.
- [42] William Thurston. *Geometry and Topology of Three-Manifolds*. Princeton lecture notes, 1980.
- [43] Dierkes Ulrich, Stefan Hildebrandt, and Friedrich Sauvigny. *Minimal Surfaces*. Springer-Verlag Berlin Heidelberg, 2 edition, 2010.
- [44] John Voight. Computing fundamental domains for Fuchsian groups. *Journal de Theorie des Nombres de Bordeaux*, 21(2):467–489, 2009. doi:10.5802/jtnb.683.
- [45] H C Wilkie. On non-Euclidean crystallographic groups. *Math. Zeitschr.*, 91:87–102, 1966.

- [46] Qingyuan Yang, Dahuan Liu, Chongli Zhong, and Jian-Rong Li. Development of computational methodologies for metal–organic frameworks and their application in gas separations. *Chemical Reviews*, 113(10):8261–8323, 2013. PMID: 23826973. doi:10.1021/cr400005f.

A Unbounded combinatorial distance

This section is dedicated to the proof of Proposition 25 below. We explain the necessary set-up. For NEC groups, Wilkie established a connection between tessellations with fundamental domains and group presentations [45], which lends itself well to applications and, in particular, motivates the kind of bounds we establish. Actually, Wilkie’s treatment is restricted to Dirichlet fundamental domains but the results apply to the general case because any isomorphism of symmetry groups is realized as a homeomorphism of \mathbb{H}^2 [30, Theorem 3] and for every fundamental domain there is a combinatorially equivalent Dirichlet fundamental domain of an isomorphic symmetry group [29, Theorem 3.3]. A Dirichlet domain $\mathcal{V}_{\tilde{x}}$ for an NEC group Γ induces the Dirichlet tessellation $\Gamma\mathcal{V}_{\tilde{x}}$, which induces a presentation of the group by introducing a *Wilkie generator* for every edge of $\mathcal{V}_{\tilde{x}}$. For \tilde{e} an edge of $\mathcal{V}_{\tilde{x}}$, the group element $[\tilde{e}]$ corresponds to traversing the edge \tilde{e} to get to the neighboring copy of $\mathcal{V}_{\tilde{x}}$ on the other side of \tilde{e} . Each vertex of $\mathcal{V}_{\tilde{x}}$ gives rise to a relation, corresponding to the sequence of edge traversals, counterclockwise, that lead back to the original $\mathcal{V}_{\tilde{x}}$. By connecting two points $\tilde{x}_1, \tilde{x}_2 \in \Gamma\tilde{x}$ if and only if the copies of $\mathcal{V}_{\tilde{x}}$ containing \tilde{x}_1 , resp. \tilde{x}_2 share an edge, we obtain an embedding into \mathbb{H}^2 of the Cayley graph corresponding to the presentation corresponding to the fundamental domain $\mathcal{V}_{\tilde{x}}$. Recall that a Cayley graph is a graph representing a presentation of a group Γ with labeled edges corresponding to a set \mathcal{G} of generators together with their inverses, and vertices corresponding to group elements [33]. Two elements $a, b \in \Gamma$ are joined by an edge labelled $g \in \mathcal{G}$ if $ag = b$. One usually regards the edges labelled by inverse elements of \mathcal{G} that connect the same two vertices as a single edge with two orientations, in which case the edge is considered as undirected. Notice that as abstract graphs, the Cayley graph and the 1-skeleton of the Delaunay tessellation are the same. By Proposition 2, the embedding of the Cayley graph can thus be assumed to feature geodesic edges.

Remark 24. One could hope to find restrictions for the number of edges of a Dirichlet fundamental domain for Γ , or the structure of the group presentation induced by the Wilkie generators, as opposed to taking any other fundamental domain, possibly simplifying our problem when restricting to Dirichlet fundamental domains. However, it turns out that there are no obstructions of this kind that limit the freedom of a Dirichlet fundamental domain for NEC groups [29, Theorem 3.3].

We will use the notion of Dehn twists on S in the following, which are homeomorphisms defined up to isotopy of the surface it acts on [17, Chapter 3]. Each nontrivial closed curve α on a surface has an associated Dehn twist, defined by taking a regular neighborhood of α , which is topologically an annulus A . A Dehn twist around α is defined as the isotopy class of the homeomorphism of the surface that twists A by 2π , is supported in A and the identity on ∂A . Recently, a theory linking groups containing Dehn twists and isotopy classes of tilings of the underlying surface has been developed [27], containing ideas which we will use in the following.

Before presenting a general result that illustrates the difficulty of our problem, we first revisit the example from the introduction: The Euclidean torus \mathbb{T} , with fundamental group Γ isomorphic to \mathbb{Z}^2 , acting on \mathbb{E}^2 . After an affine transformation of \mathbb{E}^2 , we can assume that $\Gamma = \mathbb{Z}^2$ [Theorem 2.1][41]. The Dehn twists of \mathbb{T} generate a group isomorphic to $\text{SL}(2, \mathbb{Z})$ [17, Theorem 2.5], so the idea from the introduction of shearing the square fundamental domain of \mathbb{Z}^2 with matrices in $\text{SL}(2, \mathbb{Z})$ can be seen as a consequence of this group having infinitely many elements.

Proposition 25. *Let $n \in \mathbb{N}$. There is a fundamental domain for Γ in \mathbb{H}^2 such that there exists an edge in $\text{DT}_{\tilde{\mathcal{P}}}$ the combinatorial length of which is greater than n .*

Proof. By compactness of S , the number of isotopy classes of geodesic edges connecting two points on S that are shorter than a given constant is finite. Fix an edge \tilde{e} of $\text{DT}_{\tilde{\mathcal{P}}}$ whose projection e to S is not closed, or, if closed, is nonseparating. To see that such an edge exists, observe that the only case where all edges are closed is the case where $\tilde{\mathcal{P}} = \Gamma\tilde{x}$. In this case, since $\text{DT}_{\tilde{\mathcal{P}}}$ is dual to $\mathcal{V}^{\tilde{\mathcal{P}}}$, $\text{DT}_{\tilde{\mathcal{P}}}$ represents the embedding of a Cayley graph for Γ w.r.t. the Wilkie generators for $\mathcal{V}^{\tilde{\mathcal{P}}}$ and therefore induces a

presentation of the fundamental group of S , given by the closed curves represented by edges in $\text{DT}_{\mathcal{P}}$. By Hurewicz' theorem [21, Theorem 2A.1], the curves in $\text{DT}_{\mathcal{P}}$ are a basis of the homology group of S , and therefore cannot all be separating, because separating curves are homologically trivial.

Given \tilde{e} as above, we construct a fundamental domain \mathcal{V} with one vertex orbit for Γ such that there is an edge $\tilde{e}_{\mathcal{V}}$ of \mathcal{V} that is the only edge of \mathcal{V} that intersects \tilde{e} , and does so transversally, both in their interiors, as we now explain. By Proposition 2, $e := \pi(\tilde{e})$ is a simple curve on S , we can now choose a way to complete the curve e to obtain a simple nontrivial closed curve $e_C \subset S$. Now, since e_C is nontrivial, we can find another simple closed nontrivial curve $e_{\mathcal{V}} \subset S$ that intersects e_C exactly once, somewhere in the interior of e . By applying a homeomorphism f of S , we map $e_C \cup e_{\mathcal{V}}$ to the curves on S shown in Figure 9a [17, Section 1.3.3] and consider the graph $f(e_{\mathcal{V}}) \cup v$, with vertex $v \in f(e_{\mathcal{V}} - e_C)$. Augmenting this graph by attaching $2g - 1$, where g is the genus of S , closed curves e'_2, \dots, e'_{2g} based at v as illustrated in Figure 9b, we obtain a graph G' . Here, we understand all of the closed curves to

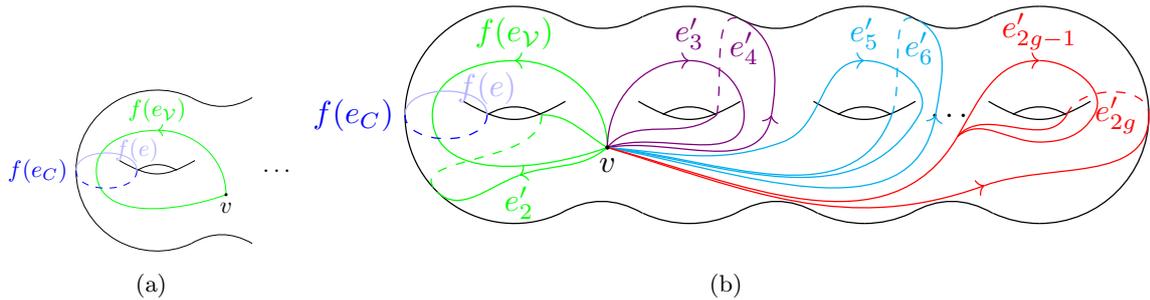


Figure 9: The graph G' in the proof of Proposition 25.

be disjoint except at v . Cutting open S along the graph G' produces a disk, as G' yields the standard presentation of the fundamental group of a surface [21, p. 5]. Any graph \hat{G} such that $S - \hat{G}$ is a disk gives rise to a fundamental domain of Γ [29, Theorem 5.1]. Therefore, $G := f^{-1}(G')$ gives rise to the sought for fundamental domain $\mathcal{V} \subset \mathbb{H}^2$, with $G \cap e = e_{\mathcal{V}} \cap e$ and therefore $\tilde{e} \cap \partial \mathcal{V} = \tilde{e} \cap \tilde{e}_{\mathcal{V}}$, for an appropriate lift $\tilde{e}_{\mathcal{V}}$ of $e_{\mathcal{V}}$. Consider the Dehn twist t about e_C , an infinite order element, which we can interpret as acting on the space of fundamental domains for Γ [27]. Here, it suffices to note that t has a representative that is a homeomorphism that maps G to a similar graph and therefore gives rise to another fundamental domain. The Dehn twist t only changes the edges of G it intersects, so by construction, it leaves invariant all other edges but $e_{\mathcal{V}}$. The number of intersections of $t^M(e_{\mathcal{V}})$ with e_C , where $t^M = t \circ t \circ \dots \circ t$ (M -times), is equal to M [17, Proposition 3.2]. Therefore, there is a representative in the isotopy class of $t^M(e_{\mathcal{V}})$ that intersects e M times, without forming any bigons which by [17, Proposition 1.7] and [17, Section 1.2.7] means that these intersections cannot be eliminated by using homotopies. Thus, the combinatorial length of the edge \tilde{e} can be made arbitrarily large. \square

B A collection of lemmata and remarks

Lemma 26. *Let $\tilde{x} \in \tilde{\mathcal{P}}$. Then for every triangle $\tilde{\Delta}$ in $\text{DT}_{\tilde{\mathcal{P}}}$ with vertex \tilde{x} , the center $\omega_{\tilde{\Delta}}$ of the circumcircle $C_{\tilde{\Delta}}$ of $\tilde{\Delta}$ satisfies $\omega_{\tilde{\Delta}} \in \mathcal{V}_{\tilde{x}}$.*

Proof. Assume that $\omega_{\tilde{\Delta}}$ is contained in some other cell of the Voronoi diagram of $\tilde{\mathcal{P}}$. We have $\tilde{\mathcal{P}} \cap \text{int}(D_{\tilde{\Delta}}) = \emptyset$, where $D_{\tilde{\Delta}}$ denotes the disk bounded by $C_{\tilde{\Delta}}$. However, if $\omega_{\tilde{\Delta}}$ is contained in $\text{int}(\mathcal{V}_{\tilde{y}})$ for some $\tilde{y} \neq \tilde{x}$, then $d_{\mathbb{H}^2}(\omega_{\tilde{\Delta}}, \tilde{y}) < d_{\mathbb{H}^2}(\omega_{\tilde{\Delta}}, \tilde{x})$, i.e. $\tilde{y} \in \text{int}(D_{\tilde{\Delta}})$, in contradiction to $\tilde{\mathcal{P}} \cap \text{int}(D_{\tilde{\Delta}}) = \emptyset$. \square

Remark 27. In case $\tilde{\mathcal{P}} = \Gamma \tilde{x}$, Lemma 26 asserts that the circumcenter of a triangle with vertex \tilde{x} is contained in the Dirichlet domain $\mathcal{V}_{\tilde{x}}$, generalizing a similar statement for Euclidean symmetry groups in the proof of Lemma 3.2 in [14].

Lemma 28. *Let $\mathcal{V}_{\tilde{x}}$ be a Dirichlet domain. Then the geodesic segment \tilde{g} that joins \tilde{x} to any $\tilde{y} \in \mathcal{V}_{\tilde{x}}$ projects to a distance minimizer in S . Conversely, if g from $x = \pi(\tilde{x})$ to y is a distance minimizer in S , then the lift \tilde{g} of g based at \tilde{x} satisfies $\tilde{g} \subset \mathcal{V}_{\tilde{x}}$.*

Proof. By definition, $\mathcal{V}_{\tilde{x}}$ consists of the points in \mathbb{H}^2 that are at least as close to \tilde{x} as to any other point in $\Gamma\tilde{x}$. Since $\mathcal{V}_{\tilde{x}}$ is convex in \mathbb{H}^2 , $\tilde{g} \subset \mathcal{V}_{\tilde{x}}$. If the statement of the lemma were not true, then there is a simple geodesic g' in S joining $x = \pi(\tilde{x})$ and $y = \pi(\tilde{y})$, which is shorter than $g = \pi(\tilde{g})$. Since geodesics connecting points are uniquely determined by their endpoints in their homotopy class in S [17, Section 1.2.7], g' is not homotopic to g , so the lift \tilde{g}' of g' based at \tilde{x} is not equal to \tilde{g} and therefore joins \tilde{x} to a point $\tilde{y}' \neq \tilde{y}$, possibly equivalent, under the action of Γ , to \tilde{y} . Moreover, \tilde{g}' is shorter than \tilde{g} because π is a local isometry, a contradiction. For the converse statement, simply observe that if the endpoint \tilde{y} of \tilde{g} was not contained in $\mathcal{V}_{\tilde{x}}$, then it is strictly closer to another point $\tilde{x}_1 \in \Gamma\tilde{x}$, which by the first part of the proof contradicts the minimality of g . \square

Lemma 29. *Two distinct distance minimizers that are not subarcs of one another on S cannot intersect each other more than once in their interior. Moreover, if an intersection occurs at an endpoint, then there cannot be an intersection in the interior.*

Proof. Suppose there exist 2 distance minimizers c_1 and c_2 , connecting the points x_1 and y_1 , and x_2 and y_2 , respectively. Assume further that c_1 and c_2 intersect each other at least 2 times in their interiors, at points z and z' .

Note that the intersection of geodesics has to be transversal, meaning that their tangent vectors cannot be parallel at the point of intersection. This can be seen using a general argument concerning geodesics in Riemannian manifolds. If two geodesics g_1, g_2 intersected nontransversally at some point p , then $g_1 = g_2$, after reparametrization. This is because geodesics are the solution of a second-order differential equation, and these are uniquely determined by their initial values, i.e. the point p and their tangent vectors at p , so after rescaling the tangent vector if necessary, the geodesics agree [34, Chapter 3].

Since both c_1 and c_2 minimize distances along their paths, the portion of both c_1 and c_2 connecting z to z' realize the distance between these points as geodesic arcs. Therefore, we can connect x_1 to z along c_1 , then connect z to z' along c_2 , and z' to y_1 along c_1 to obtain a curve c_3 connecting x_1 to y_1 with the same length as c_1 . However, since the intersections of c_1 and c_2 are transversal, c_3 is not a geodesic, as it features kinks at z and z' . This implies, by the Hopf-Rinow theorem [34, Theorem 5.21], that we can further shorten the curve c_3 by applying a homotopy with fixed endpoints to find a curve connecting x_1 and y_1 homotopic to c_3 with length shorter than c_1 , a contradiction.

The second statement of the lemma corresponds to the situation where either $x_1 = z$ or $y_1 = z'$. \square

Lemma 30. *A half-minimizer can intersect a distance minimizer at most 2 times, or at least half of the half-minimizer is a subarc of the distance minimizer.*

Proof. The only case that presents difficulty and is not immediate from Lemma 29 is the case where a distance minimizer c intersects a half-minimizer c_H at both endpoints and at the half-point y of c_H . Let c join x to z , passing through y , and let $(c_H)_1, (c_H)_2$ be the two geodesic ‘halves’ of c_H , connecting x to y and y to z , respectively. Now, if $(c_H)_1$ intersected c at y nontransversally, then $(c_H)_1$ would be included in c , similarly to the proof of Lemma 29 above. If we exclude such cases, then, similarly to the above proof of Lemma 29, we find a contradiction to the distance minimality of c by considering the geodesic in the homotopy class with fixed endpoints of the curve that follows first $(c_H)_1$ to y and then c from y to z . \square

Lemma 31. *Two half-minimizers cannot intersect each other more than 4 times, as long as their halves are distinct.*

Proof. Since the two halves of a half-minimizer consist of distance minimizers by Remark 7, Lemma 30 implies the result. \square

Lemma 32. *The dodecagon \mathcal{D} depicted in Figure 1a is the Dirichlet fundamental domain $\mathcal{V}_0 \subset \mathbb{H}^2$ for Γ_{S_3} .*

Proof. First note that \mathcal{D} is symmetric w.r.t. reflections across the axes through the origin shown in Figure 7b and that \mathcal{D} is the result of glueing together 96 congruent hyperbolic triangles (with geodesic sides), each with angles $\pi/6, \pi/2$, and $\pi/4$, such that the triangles meet only along edges or at corners with equal angles [36, 38]. Moreover, the tessellation of \mathbb{H}^2 by these triangles is invariant under hyperbolic reflections in the geodesics in \mathbb{H}^2 formed by concatenating the boundaries of the triangles to obtain

geodesics. Therefore, the tessellation of \mathbb{H}^2 by copies of \mathcal{D} is invariant under reflections across the boundary of \mathcal{D} , because the boundary is made up of geodesics that correspond to such a mirror axis. See Figure 1a for an illustration of the decomposition of \mathcal{D} into such triangles. By symmetry, the lines connecting the origin 0 to the corners of \mathcal{D} have the same length. Moreover, the lines l_i from 0 to the midpoints of edges intersect the edges perpendicularly, by symmetry and the values of the angles of the triangles that decompose \mathcal{D} . Since the edges of \mathcal{D} are axes of reflection symmetries of the tessellation of \mathbb{H}^2 by \mathcal{D} , they correspond to perpendicular bisectors of 0 and points equivalent to 0. If some l_i were not a distance minimizer, then $\text{int}(\mathcal{V}_0)$ would be strictly contained in \mathcal{D} . This would imply that $\text{int}(\mathcal{D})$ contains points equivalent under the action of Γ_{S_3} , in contradiction to \mathcal{D} being a fundamental domain. \square

Remark 33. As a further consequence of the symmetries of \mathcal{D} used in the proof of Lemma 32, we see that the blue straight lines through the origin shown in Figure 7b correspond to closed geodesics on S_3 , or, equivalently, show the intersections of \mathcal{D} with the invariant axes of the translations corresponding to the side pairings that generate Γ_{S_3} .

We analyse the situation when Delaunay edges correspond to half-minimizers.

Lemma 34. *An edge \tilde{e} of $\text{DT}_{\tilde{\mathcal{P}}}$ is a half-minimizer joining its endpoints if the two circumcenters ω_1, ω_2 of the circles C_1, C_2 of the two Delaunay polygons on both sides of \tilde{e} lie on opposite sides of \tilde{e} . The circumcenters of circles containing the endpoints of \tilde{e} lie on the perpendicular bisector of \tilde{e} .*

If either the center of C_1 (or C_2) lies on \tilde{e} , necessarily a diameter of C_1 (or C_2), \tilde{e} is still a half-minimizer, but there may be multiple distance minimizers connecting a vertex of $\text{DT}_{\tilde{\mathcal{P}}}$ to the midpoint of \tilde{e} .

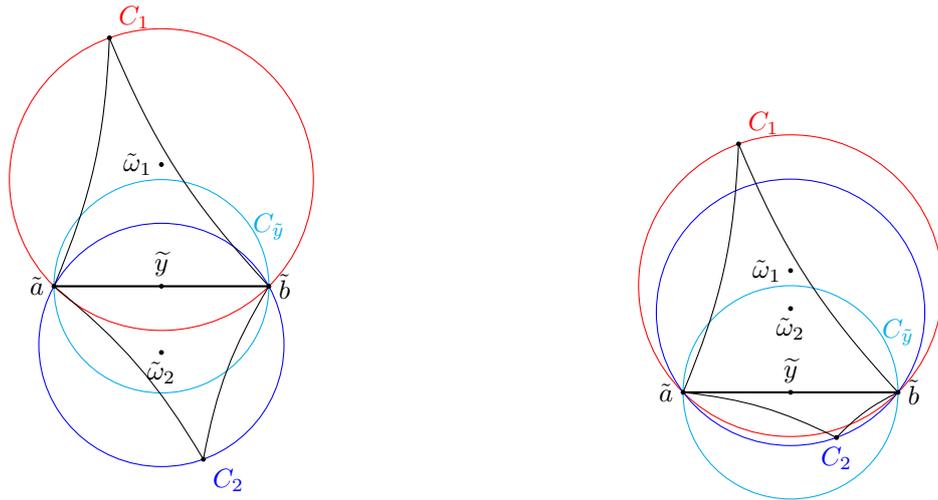
Proof. Let \tilde{a} and \tilde{b} be the endpoints of \tilde{e} , and \tilde{y} its midpoint. After applying an isometry of \mathbb{H}^2 we can assume that \tilde{e} lies on the real axis and that \tilde{y} lies on the origin. Indeed, it is well-known that the group of conformal transformations of the unit disk acts transitively on triples on the boundary of the unit disk² and therefore we can choose to map the geodesic in \mathbb{H}^2 containing \tilde{e} to the real axis by mapping its intersections with the boundary to the points 1 and -1 on the real axis. We obtain the desired situation by furthermore mapping one of the intersections of the perpendicular bisector of \tilde{a} and \tilde{b} with the boundary of the unit circle to the point i on the imaginary axis. The resulting situation is illustrated in Figure 10. By definition, the interiors of the two circles C_1 and C_2 do not contain a vertex of $\text{DT}_{\tilde{\mathcal{P}}}$. The fact that the figure depicts the situation accurately can be seen as follows. The hyperbolic distance from the origin of a point with Euclidean distance r from the origin is equal to $\text{arctanh}(r)$ [40, Section 9 – 2, eq. (1)], which is a monotonous function in r . In particular, points that have a greater Euclidean distance from the origin in \mathbb{H}^2 also have a greater hyperbolic distance. Thus, as depicted in Figure 10a, if we can show that the light blue circle $C_{\tilde{y}}$ with center \tilde{y} is contained in the union of the two closed disks bounded by the circles C_1 and C_2 , then \tilde{e} is a half-minimizer based at either endpoint. We will do this using Euclidean geometry. Consider the set, or *pencil*, \mathcal{C} of circles containing the vertices \tilde{a} and \tilde{b} . The set \mathcal{C} can also be described as the set of circles with center on the line through \tilde{y} , perpendicular to \tilde{e} . It is clear that if the Euclidean center of a circle $C_u \in \mathcal{C}$ lies above \tilde{e} , then C_u contains the upper half-disk of $C_{\tilde{y}}$, bounded by \tilde{e} . Similarly, a circle $C_d \in \mathcal{C}$ with center below \tilde{e} contains the lower half-disk of $C_{\tilde{y}}$, bounded by \tilde{e} , so \tilde{e} is a half-minimizer.

The last statement of the lemma is the limiting case where $C_{\tilde{y}}$ equals C_2 . \square

Remark 35. Figure 10b illustrates the situation where the two circumcenters lie on the same side of \tilde{e} . If $\text{DT}_{\tilde{\mathcal{P}}}$ is defined by $\Gamma\tilde{a}$, with $\tilde{a} \in \Gamma\tilde{b}$, then, using the same set-up as in the proof of Lemma 34, one sees that the edge between \tilde{a} and \tilde{b} cannot be a half-minimizer.

Remark 36. It is worth discussing the differences between the concepts of centered edges and half-minimizers in the context of the tessellations defined by $\tilde{\mathcal{P}}$. First and foremost, there is the conceptual difference that half-minimizers are intrinsically defined on a surface without reference to a particular tiling or a point set. Furthermore, the property of being a half-minimizer is hereditary, since it follows straight from their definition that subarcs of half-minimizers inherit their length-minimizing properties and thus are half-minimizers as well. On the other hand, in the presence of a point set, by Lemma 34,

²This is also sometimes paraphrased as saying that there is only one triangle with only ideal vertices in \mathbb{H}^2 , up to isometry.



(a) The situation when the circle centers lie on both sides of the edge from \tilde{a} and \tilde{b} .

(b) The situation when the circle centers lie on the same side of the edge from \tilde{a} and \tilde{b} .

Figure 10: The two situations of Lemma 34.

a centered edge of a Delaunay tessellation is a half-minimizer. Moreover, by Proposition 37 below, the centered edges of Dirichlet domains are also half-minimizers. However, it is not true in general that half-minimizers give rise to centered edges, even for closed half-minimizers that are geodesics everywhere. For an example, recall the NEC group Γ_B and Dirichlet fundamental domain from Figure 4a and consider the line segments from the origin to the corners of the Dirichlet polygon in Figure 11(left), where the line segments are shown in blue. These are equivalent, under Γ_B , to the edges shown in Figure 11(right). By Lemma 28, analogous to the discussion of Figure 7 above, these line segments can be combined arbitrarily to yield closed geodesics without kinks that one can produce on \mathbb{H}^2/Γ_B in this way, and none of these are edges of the Delaunay tessellation dual to the shown Voronoi diagram. Moreover, it is impossible to make all of these part of a single Delaunay triangulation with the red points as vertices. Thus, there are half-minimizers that are closed geodesics based at a point that are not edges in a given Delaunay triangulation with that point as vertex set.

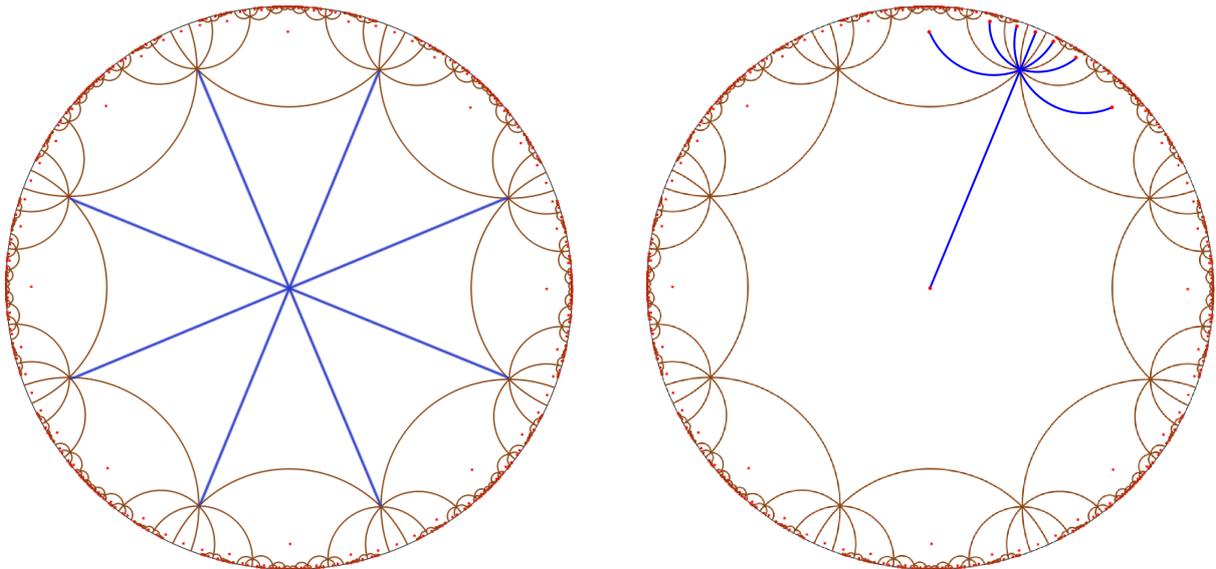
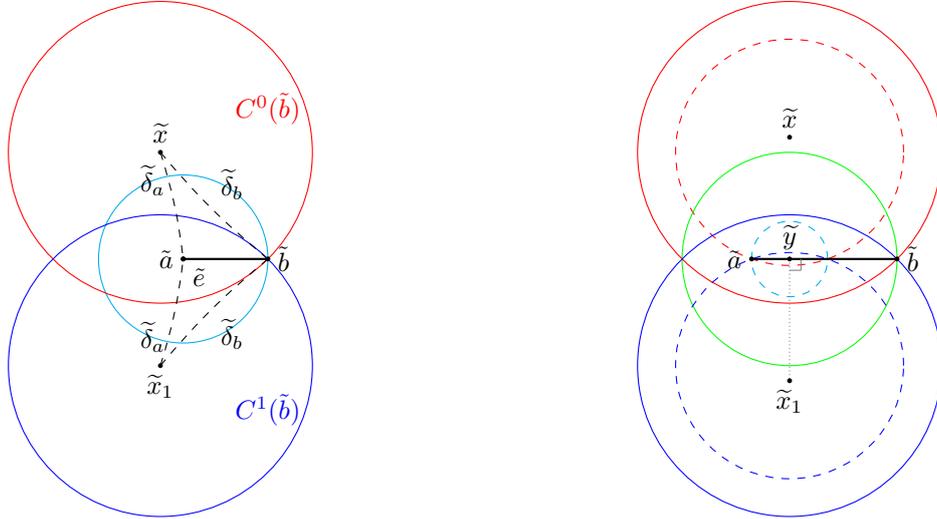


Figure 11: Closed geodesics on the Bolza surface, illustrated in a fundamental domain in \mathbb{H}^2 , from line segments through the origin.

Proposition 37. *Each centered edge of $\mathcal{V}_{\tilde{x}}$ is a half-minimizer and each non-centered edge a distance minimizer.*

Proof. Given an edge \tilde{e} of $\mathcal{V}_{\tilde{x}}$, denote its endpoints by \tilde{a} and \tilde{b} and consider the other domain incident to \tilde{e} , $\mathcal{V}_{\tilde{x}_1}$. The geodesic segments $\tilde{\delta}_z \subset \mathcal{V}_{\tilde{x}}$ and $(\tilde{\delta}_z)_1 \subset \mathcal{V}_{\tilde{x}_1}$, from \tilde{x} and \tilde{x}_1 to a point $\tilde{z} \in \tilde{e}$, respectively, together form a half-minimizer $(\tilde{\delta}_z)_h$, by Lemma 28.

We consider first the case that \tilde{e} is non-centered and assume w.l.o.g. that $d_{\mathbb{H}^2}(\tilde{b}, \tilde{x}) \geq d_{\mathbb{H}^2}(\tilde{a}, \tilde{x})$. Notice that the circle $C^0(\tilde{b})$ in \mathbb{H}^2 with $\tilde{b} \in C^0(\tilde{b})$ and center \tilde{x} bounds a disk $D^0(\tilde{b})$ the interior of which does not contain any points in $\tilde{\mathcal{P}}$, by Lemma 28. There is a similar disk $D^1(\tilde{b})$, bounded by the circle $C^1(\tilde{b}) \ni \tilde{b}$ with center \tilde{x}_1 . Since \tilde{e} is the perpendicular bisector of \tilde{x} and \tilde{x}_1 , the radii of $C^0(\tilde{b})$ and $C^1(\tilde{b})$ agree. By applying an appropriate isometry, similarly to the proof of Lemma 34, we can map \tilde{e} to the real axis in such a way that the situation is that of figure 12a. Note that we simplified the situation by mapping the geodesic from \tilde{x} to \tilde{x}_1 to the imaginary axis. The edge \tilde{e} is necessarily shorter than $\tilde{\delta}_b$, by non-centeredness. To see this, simply note that the circles in Figure 12 are symmetric w.r.t. a reflection along the geodesic from \tilde{x} to \tilde{x}_1 . Therefore, by the triangle inequality in \mathbb{H}^2 , $2l(\tilde{e}) < 2l(\tilde{\delta}_b)$, where $l(c)$ denotes the hyperbolic length of a geodesic segment c . Whence, one easily sees that the light blue circle, centered at \tilde{a} with radius $d_{\mathbb{H}^2}(\tilde{b}, \tilde{a})$, does not contain any point equivalent to \tilde{b} , and therefore, \tilde{e} is a distance minimizer.



(a) The edge from \tilde{a} to \tilde{b} is a distance minimizer in case the edge is not centered.

(b) The arcs from \tilde{y} to \tilde{a} and \tilde{b} are distance minimizers when the edge between \tilde{a} and \tilde{b} is centered.

Figure 12: The two cases of the proof of proposition 37.

Consider now the case where \tilde{e} is centered. By mapping \tilde{e} to the real axis in \mathbb{H}^2 and the points \tilde{x} and \tilde{x}_1 to the imaginary axis, we can assume we are in the situation of Figure 12b. Projecting \tilde{x} (and \tilde{x}_1) to its nearest point on \tilde{e} , we obtain a point $\tilde{y} \in \tilde{e}$ between \tilde{a} and \tilde{b} . We claim that \tilde{y} is a half-point of the half-minimizer \tilde{e} , which will finish the proof. For this, similarly to above, we consider the open disks $D^0(\tilde{b})$ and $D^1(\tilde{b})$ around \tilde{x} and \tilde{x}_1 , respectively, that do not contain any point equivalent to \tilde{b} . Their boundary circles are shown as solid arcs in Figure 12b along with a disk $D(\tilde{b})$ bounded by the solid green circle centered at \tilde{y} , with $D(\tilde{b}) \subset D^0(\tilde{b}) \cup D^1(\tilde{b})$. This shows that the subarc of \tilde{e} from \tilde{y} to \tilde{b} is a distance minimizer. Consider now the disks bounded instead by the dashed circles, where the red and blue ones both have a radius of $d_{\mathbb{H}^2}(\tilde{a}, \tilde{x})$ and centers \tilde{x} and \tilde{x}_1 , respectively. One deduces, similarly to above, that the light blue dashed circle centered at \tilde{y} with radius $d_{\mathbb{H}^2}(\tilde{a}, \tilde{y})$ bounds a disk not containing any points equivalent to \tilde{a} . Therefore, the subarc of \tilde{e} from \tilde{y} to \tilde{a} is a distance minimizer. \square

Remark 38. Notice that the statement of Proposition 37 follows immediately from Lemma 34 in case the Dirichlet fundamental domain has only one vertex orbit, since it is then a polygon inscribed into a circle. Therefore, Proposition 37 can be seen as a generalization of an observation for cyclic polygons. In

view of this, we consider other possible generalizations of observations for cyclic Dirichlet fundamental domains. A natural question is whether the edges of a maximal cyclic polygon defined by equivalent points on the boundary are half-minimizers. Note that the equivalent vertices of a Dirichlet fundamental domain $\mathcal{V}_{\tilde{x}}$ have the same distance to \tilde{x} [3, Theorem 9.4.3].

The following is a similar statement to lemma 2.5 in [9].

Proposition 39. *Let $\widetilde{CP}_{\tilde{x}} \subset \mathcal{V}_{\tilde{x}}$ be the cyclic polygon obtained by joining vertices of $\mathcal{V}_{\tilde{x}}$ equivalent under Γ . If all edges of $\widetilde{CP}_{\tilde{x}}$ are half-minimizers, then $\widetilde{CP}_{\tilde{x}}$ is centered, i.e. the circumcenter \tilde{x} is contained in $\widetilde{CP}_{\tilde{x}}$.*

Proof. Note that by Lemma 28 and the uniqueness of the Delaunay tessellation, the edges of $\widetilde{CP}_{\tilde{x}}$ are part of a Delaunay triangulation $DT_{\tilde{\mathcal{P}}}$ with vertex set $\tilde{\mathcal{P}} = \Gamma\tilde{y}$, where \tilde{y} is a vertex of $\widetilde{CP}_{\tilde{x}}$. Since $\widetilde{CP}_{\tilde{x}}$ has embedded geodesic edges $\{\tilde{e}_i\}_{i=1}^n$ in \mathbb{H}^2 since $\mathcal{V}_{\tilde{x}}$ is convex, $\widetilde{CP}_{\tilde{x}}$ is contained in the half-spaces of \mathbb{H}^2 that lie to one side of each of the full geodesics \widehat{e}_i containing each edge \tilde{e}_i , respectively. Note for this that since $\widetilde{CP}_{\tilde{x}}$ is cyclic, all interior angles must be smaller than π , so $\widetilde{CP}_{\tilde{x}}$ is convex in \mathbb{H}^2 . Although one can give a direct proof of the proposition, we will proceed with a proof by contradiction, because we will employ its reasoning subsequently. If \tilde{x} lies outside of $\widetilde{CP}_{\tilde{x}}$, it must lie on the other side of some \widehat{e}_k than $\widetilde{CP}_{\tilde{x}}$. We claim that \tilde{e}_k with endpoints \tilde{a} and \tilde{b} is a non-centered edge in this case. Indeed, by Lemma 34, the possible circles that contain \tilde{a} and \tilde{b} have centers lying on the perpendicular bisector of \tilde{a} and \tilde{b} . Therefore, the center of any circle not containing any other vertex of $\widetilde{CP}_{\tilde{x}}$ must be further away from \tilde{e}_k than \tilde{x} , see also Figure 10b. Therefore, by Lemma 34, \tilde{e}_k is not a half-minimizer. \square

Remark 40. Proposition 39 is also valid for cyclic polygons without reference to a symmetry group Γ , as the proof shows. If a cyclic polygon $\widetilde{CP}_{\tilde{x}}$ does not contain its center \tilde{x} , then the non-centered edge \tilde{e}_k with vertices \tilde{a} and \tilde{b} found in the proof, with the property that \tilde{x} and $\widetilde{CP}_{\tilde{x}}$ lie to different sides, is uniquely determined. Indeed, by arguing just like in the proof above, we see that \tilde{x} lies on the perpendicular bisector \tilde{e}_i^\perp of each edge \tilde{e}_i of $\widetilde{CP}_{\tilde{x}}$. However, because subsequent edges of $\widetilde{CP}_{\tilde{x}}$ enclose an angle smaller than π in $\widetilde{CP}_{\tilde{x}}$, for \tilde{e}_i^\perp for $i \neq k$ to meet \tilde{x} , $\text{int}(\widetilde{CP}_{\tilde{x}}) \cap \tilde{e}_i^\perp \neq \emptyset$ for $i \neq k$. Therefore, $\widetilde{CP}_{\tilde{x}}$ lies on the same side of each \tilde{e}_i^\perp for $i \neq k$ as \tilde{x} . In particular, it is straightforward to see³ that \tilde{e}_k is the unique longest edge of $\widetilde{CP}_{\tilde{x}}$, see also [10, Proposition 2.2] and Figure 13b below.

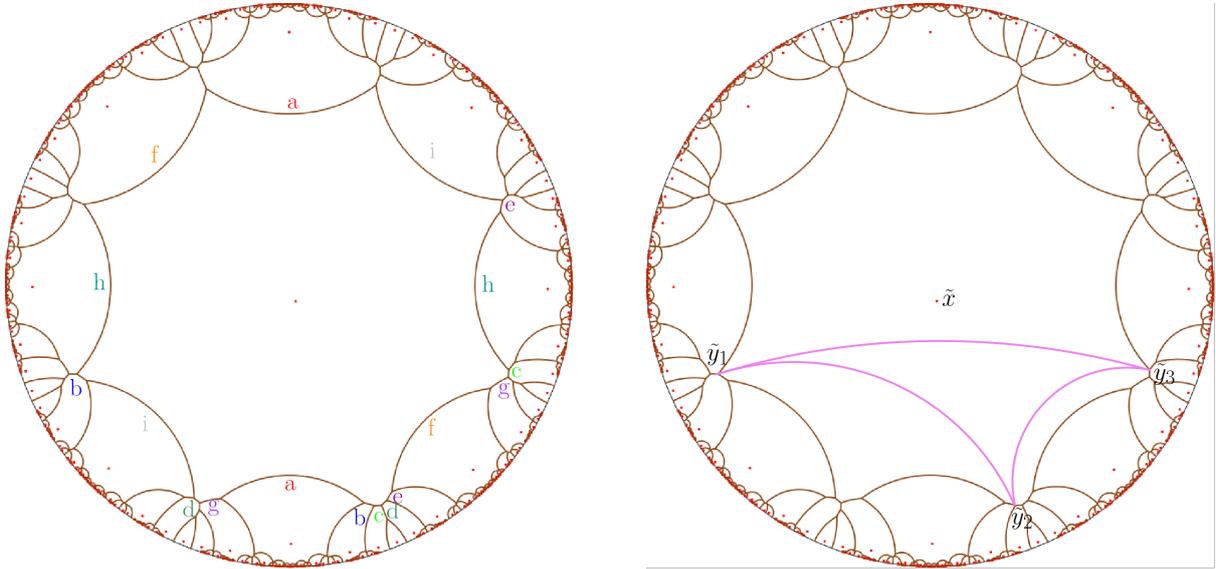
Remark 41. Note that it is generally not true that a centered polygon with only equivalent vertices has only half-minimizers as edges, since if the circumcenters of two polygons sharing an edge \tilde{e} lie on the same side of \tilde{e} , then \tilde{e} is not a half-minimizer, see Remark 35. On the other hand, it follows straight from Lemma 34 that a sufficient condition for all edges of such a centered polygon to be half-minimizers is for it to only share an edge with other centered polygons.

Using Proposition 39, it is straightforward to find a counter-example to the proposed question of Remark 38. We again take our example of the group Γ_B from Figure 4a above and the orbit of any of the red points and consider the Voronoi diagram. Using this definition of Γ_B , one can find the edge identifications of the Dirichlet fundamental 18-gon depicted, together with its tessellation, in Figure 13a. The vertex orbit on the boundary of $\mathcal{V}_{\tilde{x}}$ containing the three vertices, marked by \tilde{y}_1 , \tilde{y}_2 and \tilde{y}_3 in Figure 13b yields a cyclic triangle $\tilde{\Delta}_B$ with vertices $\{\tilde{y}_1, \tilde{y}_2, \tilde{y}_3\}$. However, since \tilde{x} is not contained in $\tilde{\Delta}_B$, Proposition 39 (and Remark 40) implies that the geodesic between \tilde{y}_1 and \tilde{y}_3 in \mathbb{H}^2 is not a half-minimizer.

Proposition 42. *A (closed) triangle in S with edges that are distance minimizers is convex.*

Proof. Triangles in S have the property that their lift to \mathbb{H}^2 is convex and therefore contain a unique geodesic joining any two points in them. We will prove that this geodesic is a distance minimizer, which also shows, in particular, that no two points of $\tilde{\Delta} \subset \mathbb{H}^2$, a lift of a triangle Δ with distance minimizers edges in S , project to the same point on S . Denote the geodesic connecting two points \tilde{z}_1 and \tilde{z}_2 in \mathbb{H}^2 by $\tilde{g}_{\tilde{z}_1\tilde{z}_2}$. We first prove that the geodesics from corners to arbitrary points on the edges are distance minimizers. Denote the vertices of $\tilde{\Delta}$ by \tilde{x}, \tilde{y} , and \tilde{z} . Consider w.l.o.g. the edge $\tilde{g}_{\tilde{x}\tilde{z}}$. The geodesics $\tilde{g}_{\tilde{y}\tilde{z}}$ and $\tilde{g}_{\tilde{y}\tilde{x}}$ are distance minimizers so $\tilde{z}, \tilde{x} \in \mathcal{V}_{\tilde{y}}^{\tilde{\mathcal{P}}}$, by Lemma 28. Since $\mathcal{V}_{\tilde{y}}$ is convex, $\tilde{g}_{\tilde{x}\tilde{z}} \subset \mathcal{V}_{\tilde{y}}$ and therefore, again by Lemma 28, there is a distance minimizer from \tilde{y} to any point on the edge $\tilde{g}_{\tilde{x}\tilde{z}}$.

³Simply observe that the angle subtended by \tilde{e}_k from \tilde{x} is the maximal angle among all chords that represent edges of $\widetilde{CP}_{\tilde{x}}$ and is $\leq \pi$.



(a) A Dirichlet fundamental domain with 18 sides for the symmetry group Γ_B from Figure 4a above, with edge identifications. Translations that identify edges with the same label generate Γ_B and the Dirichlet tessellation.

(b) The Dirichlet fundamental domain $\mathcal{V}_{\tilde{x}}$ and tessellation from (a), with $\{\tilde{y}_1, \tilde{y}_2, \tilde{y}_3\}$ representing a complete set of equivalent points on the boundary. The edge between \tilde{y}_1 and \tilde{y}_3 is not a half-minimizer.

Figure 13: Voronoi diagrams from the points marked in red, the orbit of \tilde{x} under the symmetry group Γ_B .

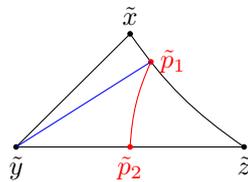


Figure 14: The hyperbolic triangles $\tilde{\Delta}$ and $\tilde{\Delta}'$.

If there were points $\tilde{p}_1, \tilde{p}_2 \in \tilde{\Delta}$ such that $\tilde{g}_{\tilde{p}_1\tilde{p}_2}$ was not a distance minimizer then we may assume that these lie on the boundary of $\tilde{\Delta}$. Indeed, extending $\tilde{g}_{\tilde{p}_1\tilde{p}_2}$ until it crosses the boundary $\partial\tilde{\Delta}$ would yield a curve that is not the distance minimizer between the points on the boundary if there is a curve from \tilde{p}_1 to \tilde{p}_2 not contained in $\tilde{\Delta}$. So, assume that $\tilde{p}_1, \tilde{p}_2 \in \partial\tilde{\Delta}$ and further that $\tilde{p}_1 \in \tilde{g}_{\tilde{x}\tilde{z}}, \tilde{p}_2 \in \tilde{g}_{\tilde{y}\tilde{z}}$. Figure 14 illustrates the situation. Then, consider the geodesic $\tilde{g}_{\tilde{p}_1\tilde{y}} \subset \tilde{\Delta}$. By the first part of the proof, $\tilde{g}_{\tilde{p}_1\tilde{y}}$ is a distance minimizer and therefore the triangle $\tilde{\Delta}' \subset \tilde{\Delta}$ with vertices $\{\tilde{p}_1, \tilde{y}, \tilde{z}\}$ is again a triangle with distance minimizers as edges. By the above reasoning, the geodesic $\tilde{g}_{\tilde{p}_1\tilde{p}_2}$ must be a distance minimizer in $\tilde{\Delta}'$ since it connects a corner of $\tilde{\Delta}'$ to a point on its edge, finishing the proof. \square

Remark 43. The generalization of Proposition 42 to polygons with more edges does not hold. For a counter-example, consider a Dirichlet fundamental domain $\mathcal{V}_{\tilde{x}}$ for some NEC group Γ with only one vertex orbit and connect \tilde{x} with two vertices incident to the same edge \tilde{e} of $\mathcal{V}_{\tilde{x}}$ by geodesics. By splitting \tilde{e} along its half-point, which exists by Proposition 37, we obtain a convex quadrilateral \tilde{Q} in \mathbb{H}^2 with edges that are distance minimizers w.r.t. $\tilde{\mathcal{P}} = \Gamma\tilde{x}$. However, as two of its vertices project to the same point on $S = \mathbb{H}^2/\Gamma$, by construction, so Q cannot be convex in S . It is easy to see that shifting Q around a bit on S , one can also obtain examples of this type without equivalent vertices.

Remark 44. Notice that a closed half-minimizer c based at $x \in S$ with the property that c is a half-minimizer based at y for all $y \in c$ has special properties, which can be used to further reduce the number of allowed intersections of it with a half-minimizer given by Lemma 31. Such a curve is necessarily smooth everywhere because the half-point of a closed half-minimizer is located at the point furthest away from the base point and the base point and the half-point are the only points where a closed half-minimizer can have kinks, see Remark 7. The converse is, however, not true for general half-minimizers, as illustrated by Remark 7 and Figure 5.

C Proof of the main theorem

Theorem 45. *Let Γ be the fundamental group of a closed hyperbolic surface S . Let $\mathcal{V}_{\tilde{x}}$ be a Dirichlet fundamental domain for Γ , with k edges and $(\mathcal{V}_{\tilde{x}})_O \subset \mathcal{V}_{\tilde{x}}$ any original fundamental domain. Let $\tilde{\mathcal{P}}$ be a finite set of points on S and $\text{DT}_{\tilde{\mathcal{P}}}$ the Delaunay triangulation of the lifted set of points $\tilde{\mathcal{P}} \subset \mathbb{H}^2$, with edges that are distance minimizers. Then the maximal combinatorial length of a Delaunay edge is bounded by $k/2$.*

Proof. We first set up the necessary notation. Denote the counter-clockwise cyclically ordered edges of $\mathcal{V}_{\tilde{x}}$ by $\{\tilde{e}_i\}_{i=1}^k$, with corners $\{\tilde{v}_i\}_{i=1}^k$, where \tilde{v}_i and \tilde{v}_{i+1} (where $\tilde{v}_{k+1} := \tilde{v}_1$) are incident to \tilde{e}_i and a half-point \tilde{m}_i on every edge \tilde{e}_i , using Proposition 37. In case an edge of $\mathcal{V}_{\tilde{x}}$ is non-centered for $\Gamma\tilde{x}$, we can choose an arbitrary half-point in the interior of the edge. Note that $\tilde{e}_i \setminus \tilde{m}_i$ has two components, whose closures $(\tilde{e}_i)_1$ and $(\tilde{e}_i)_2$ are distance minimizers. Denote the geodesic segments from \tilde{x} to the corners \tilde{v}_i by \tilde{c}_{v_i} and those from \tilde{x} to \tilde{m}_i by \tilde{c}_{m_i} . By Lemma 28, \tilde{c}_{v_i} and \tilde{c}_{m_i} are also distance minimizers, so the projection Δ_i^j of each closed triangle $\tilde{\Delta}_i^j$ with vertices $\{\tilde{x}, \tilde{v}_{i+j}, \tilde{m}_i\}$ for $i \in \{1, \dots, k\}$ and $j \in \{0, 1\}$ as in Figure 8a is convex in $S = \mathbb{H}^2/\Gamma$, by Proposition 42. Let \tilde{c}_p be a Delaunay edge that is a distance minimizer incident to \tilde{p} with projection $p \in S$. Consider first the case where $\tilde{p} \in \tilde{\mathcal{P}} \cap \partial\mathcal{V}_{\tilde{x}}$.

Since \tilde{c}_p is simple by Lemma 29, we can picture the projection c_p of \tilde{c}_p as a collection of nonintersecting geodesic segments $\{\tilde{c}_p^n\}_{n=1}^r$ in $\mathcal{V}_{\tilde{x}}$, each connecting different points in $\mathcal{V}_{\tilde{x}}$ by considering piece-wise lifts of c_p to $\mathcal{V}_{\tilde{x}}$ as illustrated in Figure 8b for the Bolza surface. We assume that the ordering of $\{\tilde{c}_p^n\}_{n=1}^r$ is such that the projection $c_p^{\leq m}$ to S of the first m segments of $\{\tilde{c}_p^n\}_n$ is a curve on S for all m and \tilde{c}_p^1 emanates from \tilde{p} . We denote with $\tilde{c}_p^{\leq m}$ the lift, starting at \tilde{p} , of $c_p^{\leq m}$ to \mathbb{H}^2 . Assuming that \tilde{c}_p^1 consists of more than one point, every \tilde{c}_p^n is incident to two points in $\partial\mathcal{V}_{\tilde{x}}$, except possibly \tilde{c}_p^r . Observe that $\tilde{c}_p^{\leq m}$ is contained in \tilde{c}_p and $\tilde{c}_p^{\leq r} = \tilde{c}_p$ and that each \tilde{c}_p^n has an orientation such that its endpoint in $\partial\mathcal{V}_{\tilde{x}}$ maps to the starting point of \tilde{c}_p^{n+1} under the pairwise identification of edges under Γ for $n < r$.

Since \tilde{c}_p is a distance minimizer, by Lemma 29 and the convexity of Δ_i^j in S , $c_p \cap \Delta_i^j$ has to be connected for all i, j . Therefore, if \tilde{c}_p^n intersects $\tilde{\Delta}_i^j$ for some n , then \tilde{c}_p^m cannot intersect $\tilde{\Delta}_i^j$ for $m \neq n$. We will explain how to bound the combinatorial length of \tilde{c}_p w.r.t. $(\mathcal{V}_{\tilde{x}})_O$ using the geodesic segments in $\{\tilde{c}_p^n\}_n$.

If $\tilde{c}_p^{\leq m}$ has a combinatorial length of l , then the combinatorial length of $\tilde{c}_p^{\leq m+1}$ is at most $l+1$, since \tilde{c}_p^{m+1} is contained in $\mathcal{V}_{\tilde{x}}$. We will now maximize the combinatorial distance that \tilde{c}_p can have finding a \tilde{c}_p that has as many segments \tilde{c}_p^n as possible. In particular, we can assume that $r > 1$.

For clarity of the following arguments, we have to adjust the above set-up slightly, in the following technical way: Assume first that no \tilde{c}_p^n is incident to a vertex \tilde{v}_i . We slide the segments \tilde{c}_p^n for $1 \neq n \neq r$ along appropriate lifts of \tilde{c}_p slightly so they both start and end in the interior of a copy of $\mathcal{V}_{\tilde{x}}$, so that they intersect exactly one edge of $\mathcal{V}_{\tilde{x}}$, the edge on which its former endpoint lay. We also extend \tilde{c}_p^1 slightly accordingly so it crosses over the edge containing its endpoint, and shorten \tilde{c}_p^r , so its starting point lies in the interior of $\mathcal{V}_{\tilde{x}}$. These technical adjustments mean that we do not have to consider different cases for the combinatorial lengths of the \tilde{c}_p^n depending on which edges of $\mathcal{V}_{\tilde{x}}$ belong to $(\mathcal{V}_{\tilde{x}})_O$. We call the edge of $\mathcal{V}_{\tilde{x}}$ where the original \tilde{c}_p^n started its starting edge in the following.

With the above adjustment, if \tilde{c}_p^n for $n > 1$ intersects no edge or a neighboring edge of its starting edge in $\partial\mathcal{V}_{\tilde{x}}$, then the combinatorial length of $\tilde{c}_p^{\leq n-1}$ is equal to that of $\tilde{c}_p^{\leq n}$, since then the endpoint of \tilde{c}_p^n is contained in the (interior of the) same neighborhood layer as the endpoint of $\tilde{c}_p^{\leq n-1}$. This implies that to increase the combinatorial distance of $\tilde{c}_p^{\leq n-1}$ by concatenation with a copy of the segment \tilde{c}_p^n , \tilde{c}_p^n has to intersect an edge of $\mathcal{V}_{\tilde{x}}$ that is not neighboring its starting edge. In particular, such an edge must intersect at least 4 distinct $\tilde{\Delta}_i^j$. Note that if a segment did intersect a corner, contrary to our simplifying assumption above, then its projection to S would intersect at least 6 of the $\tilde{\Delta}_i^j$. This is because the degree of a vertex in the Dirichlet tessellation associated to $\mathcal{V}_{\tilde{x}}$ is at least 3 and at least 3 points in $\mathcal{V}_{\tilde{x}}$ are equivalent to every corner under Γ [3, Section 9.3]. We therefore see that the assumption that each \tilde{c}_p^n is disjoint from the $\{\tilde{v}_i\}$ is without loss of generality.

The rest of the proof is now straightforward. The assumption that the first segment \tilde{c}_p^1 consists of more than one point means that it has a combinatorial distance of at most 2, in which case it intersects at least $5+1$ triangles, because p is incident to at least 2 triangles on S . For each segment \tilde{c}_p^n with $n < r$, increasing the combinatorial distance of $\tilde{c}_p^{\leq n-1}$ by 1 by concatenation with a copy of \tilde{c}_p^n uses up a further 4 of the $2k$ triangles in $\{\tilde{\Delta}_i^j\}_{i,j}$. Inductively, after concatenation of the first n segments, we obtain $\tilde{c}_p^{\leq n}$ with a combinatorial length of $\leq 1+n$, and non-empty intersection with at least $6+4(n-1)$ of the triangles. Since k is always an even number [3], $n \leq k/2 - 1$ and therefore the maximum value for the combinatorial length is $k/2$.

Note that the last segment \tilde{c}_p^r can terminate in the interior of $\mathcal{V}_{\tilde{x}}$, without adding to the combinatorial length of $\tilde{c}_p^{\leq r-1}$. In particular, this means that the maximum value of r is $k/2$. See Figure 8b for the piece-wise lifts $\{\tilde{c}_p^n\}_{n=1}^4$ of a curve that maximizes the combinatorial distance on the Bolza surface.

Lastly, observe that if there is no point on the boundary of $\partial\mathcal{V}_{\tilde{x}}$ in $\tilde{\mathcal{P}}$, then the combinatorial length of the first segment \tilde{c}_p^1 is at most 1. On the other hand, the last segment may add 1 to the combinatorial length of $\tilde{c}_p^{\leq r-1}$ in at least one situation where this was previously impossible, so $k/2$ remains an upper bound for the maximal combinatorial distance. \square

Remark 46. Given the situation of Theorem 45 and the notation of the proof, the combinatorial distance between two points in $\tilde{\mathcal{P}}$ depends on the direction the points lie in. To see an example of this, consider that a distance minimizer \tilde{c}_p incident to a point \tilde{p} on the edge \tilde{e}_i of $\mathcal{V}_{\tilde{x}}$ cannot intersect more than two copies of \tilde{e} , because then its projection would have to intersect either of the projections of $(\tilde{e}_i)_1$ and $(\tilde{e}_i)_2$ to S more than once, which is impossible by Lemma 29.