



HAL
open science

Representing Infinite Periodic Hyperbolic 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 Periodic Hyperbolic Delaunay Triangulations Using Finitely Many Dirichlet Domains. *Discrete and Computational Geometry*, 2024, 10.1007/s00454-024-00653-x . hal-04577851

HAL Id: hal-04577851

<https://hal.science/hal-04577851v1>

Submitted on 16 May 2024

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.

Representing infinite periodic hyperbolic Delaunay triangulations using finitely many Dirichlet domains

Vincent Despré¹, Benedikt Kolbe², and Monique Teillaud¹

¹Inria, Université de Lorraine, CNRS, LORIA, Nancy, F-54000, France. E-mails: vincent.despre@loria.fr, monique.teillaud@inria.fr

²Corresponding author. Hausdorff Center for Mathematics, University of Bonn, Friedrich-Hirzebruch-Allee 5, Bonn, 53115, Germany. E-mail: benedikt.kolbe@physik.hu-berlin.de. This work was done while this author was working at Université de Lorraine, CNRS, Inria, LORIA, F-54000 Nancy, France.

May 14, 2024

Abstract

The Delaunay triangulation of a set of points P on a hyperbolic surface is the projection of the Delaunay triangulation of the set \tilde{P} of lifted points in the hyperbolic plane. Since \tilde{P} is infinite, the algorithms to compute Delaunay triangulations in the plane do not generalize naturally. Using a Dirichlet domain, we exhibit a finite set of points that captures the full triangulation.

We prove that an edge of a Delaunay triangulation has a combinatorial length (a notion we define in the paper) smaller than $12g - 6$ with respect to a Dirichlet domain. To achieve this, we introduce new tools, of intrinsic interest, that capture the properties of length-minimizing curves in the context of closed curves. We then use these to derive structural results on Delaunay triangulations and exhibit certain distance minimizing properties of both the edges of a Delaunay triangulation and of a Dirichlet domain.

The bounds produced in this paper depend only on the topology of the surface. They provide mathematical foundations for hyperbolic analogs of the algorithms to compute periodic Delaunay triangulations in Euclidean space.

Keywords— Delaunay triangulations, Hyperbolic surfaces, Fuchsian group, Moduli space, Dirichlet domain .

Mathematics Subject Classification 68U05 · 52C45 · 57K20 · 32G15

Acknowledgements. We want to thank Hugo Parlier, Jean-Marc Schlenker and Gert Vegter for fruitful discussions on the topic of Delaunay triangulations on hyperbolic surfaces over many years. We also thank the anonymous reviewers for their valuable and insightful comments that have led to an improvement of the paper.

1 Introduction

A hyperbolic surface is a *closed* and *orientable* topological surface equipped with some hyperbolic metric of constant curvature -1 . Recently, motivated in part by applications in other sciences and its ubiquity, there has been an increased effort to understand the hyperbolic geometry of surfaces from a computational geometry point of view. A fundamental question addresses the computation of Delaunay triangulations on hyperbolic surfaces. The classic edge flip algorithm of Lawson [33] computing Delaunay triangulations in the Euclidean plane was recently extended to hyperbolic surfaces [14]. However, robust and efficient software to compute Delaunay triangulations on hyperbolic surfaces, and particularly on the triply-periodic minimal surfaces presented below, does not exist to date. A primary motivation for the work in this paper is to help fill this gap by establishing fundamental theoretical results in the hyperbolic case, similar to those that have led to the only such software for flat quotient spaces [7, 32, 6, 39]. Our results yield structural insights of independent interest into the relationship between different representations of hyperbolic surfaces and Delaunay triangulations on them. These structural results are a consequence of an investigation of a natural class of paths, called half-minimizers, that also seems useful in other contexts.

Motivation - Hyperbolic surfaces in other sciences and nature. One of the motivations for this paper 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 invariant under three linearly independent translations, i.e. a rank 3 lattice L_3 [43]. To associate a closed surface to a TPMS, one considers the TPMS as the lift to \mathbb{R}^3 of the closed surface S_g of genus g in the 3-torus $\mathbb{T}^3 = \mathbb{R}^3/L_3$. This corresponds to gluing the TPMS along the equivalent faces of a translational domain for L_3 . It turns out that the surface S_g is always intrinsically hyperbolic [37]. The gyroid, the primitive, and the diamond TPMS are arguably the most prominent and simple examples of TPMS [41] and have received considerable attention in the mathematical, physical, chemical and biological, as well as interdisciplinary literature [44, 27, 23, 25, 12, 31]. More recently, TPMS have also found a role in the materials sciences as a scaffold for crystallographic structures [17, 18], leading to both new mathematical formalisms [29, 30, 31, 28] and a database of such structures [24]. These three TPMS are closely related to each other [43, Section 3.1][21] and have the same underlying hyperbolic surface S_3 , of genus 3, embedded in \mathbb{T}^3 .

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 , also known as D-surface, illustrating how TPMS are covered by \mathbb{H}^2 . 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.

Another motivation comes from previous research in the field. A similar algorithm for the computation of certain types of Delaunay triangulations on the Bolza surface has been implemented in CGAL [26] to accommodate the push for practical algorithms in other sciences. Our results add to the growing body of both theoretical and applied work surrounding periodic hyperbolic Delaunay triangulations [4, 10, 11, 16].

The flat case. The computation of Delaunay triangulations in flat tori, which can be seen equivalently as periodic triangulations in Euclidean space, was addressed by Dolbilin and Huson [15], who provided a first cornerstone for the algorithms and CGAL packages handling the square/cubic periodicity [8, 7, 32]; the algorithms were generalized later [9]. Their work was used again for the recent and, as far as we know, only implementations for general periodic point sets in the Euclidean plane (or three-dimensional space) [39]. The idea is as follows.

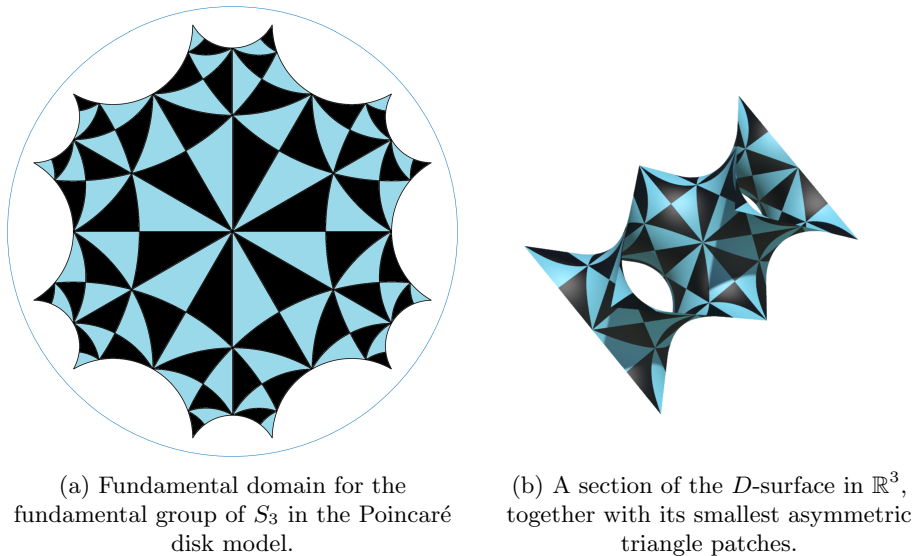


Figure 1: The covering of the diamond TPMS by \mathbb{H}^2 .

Let $\tilde{\mathcal{P}}$ be a locally finite periodic point set in the plane \mathbb{E}^2 . A finite set $\tilde{\mathcal{P}}_f$ can be defined, such that the infinite Delaunay triangulation $DT_{\tilde{\mathcal{P}}}$ can be deduced from the Delaunay triangulation $DT_{\tilde{\mathcal{P}}_f}$; as $\tilde{\mathcal{P}}_f$ is finite, $DT_{\tilde{\mathcal{P}}_f}$ can be computed by any classical algorithm. In practice, the periodic set $\tilde{\mathcal{P}}$ is generated from a finite set of points in a fundamental domain (defined properly below) of a lattice, which one can assume to be given as a parallelogram. The periodicity is obtained by the action of the group of translations, isomorphic to \mathbb{Z}^2 , generated by the two vectors corresponding to the sides of the parallelogram.

Figure 2 shows the special case when the infinite set $\tilde{\mathcal{P}}$ is obtained from a unique point \tilde{a} ; all blue points are images of \tilde{a} under the action of the group of translations. We observe that there cannot exist a general bound on the size of $\tilde{\mathcal{P}}_f$ for an arbitrary choice of parallelogram/translations. Intuitively speaking, this is because there is no bound on how stretched a fundamental domain \mathcal{F} can appear for the same lattice points. For ever more long and thin parallelograms, an edge \tilde{e} in the Delaunay triangulation of $\tilde{\mathcal{P}}$ may traverse an unbounded number of copies of \mathcal{F} , see Figure 3.

However, if choosing as fundamental domain the Dirichlet domain $\mathcal{D}_{\tilde{x}}$ of an arbitrary point \tilde{x} , i.e. the Voronoi cell of \tilde{x} in the Voronoi diagram of $\mathbb{Z}^2\tilde{x}$ (Figure 4), then a bound can be proved on the number of copies of $\mathcal{D}_{\tilde{x}}$ necessary to account for an edge in a Delaunay triangulation [15, 39]. Figure 4 depicts 19 shaded copies of $\mathcal{D}_{\tilde{x}}$ that are sufficient and form layers around the domain containing \tilde{x} , illustrating the general situation for hexagonal Dirichlet domains. Note that the shape of the Dirichlet domain $\mathcal{D}_{\tilde{x}}$ does not depend on the chosen point \tilde{x} : for another point \tilde{y} , $\mathcal{D}_{\tilde{y}}$ is a translated version of $\mathcal{D}_{\tilde{x}}$.

The hyperbolic case. A hyperbolic surface S is homeomorphic to the quotient \mathbb{H}^2/Γ of the hyperbolic plane \mathbb{H}^2 under the action of a *symmetry group* Γ of \mathbb{H}^2 , i.e. a discrete subgroup of the group of isometries of \mathbb{H}^2 . That S is a hyperbolic surface implies that Γ contains only orientation preserving isometries and has no fixed points in \mathbb{H}^2 . The group Γ can be naturally identified with the fundamental group $\pi_1(S)$ of S (after choosing base points appropriately). The symmetry group is also known as a Non-Euclidean Crystallographic (NEC) group. The universal covering space of S is




Figure 2: Periodic point set $\tilde{\mathcal{P}}$ given by the orbit of \tilde{a} under the action of $\mathbb{Z} \cdot u + \mathbb{Z} \cdot v$ on \mathbb{E}^2 . A fundamental domain is shaded. The (infinite) Delaunay triangulation of $\tilde{\mathcal{P}}$ is shown in green.




Figure 3: The same (black) lattice points correspond to fundamental domains that can be arbitrarily elongated. The number of fundamental domains traversed by e is unbounded.




Figure 4: Tiling of \mathbb{H}^2 by translated images of the Dirichlet domain of \tilde{x} .

\mathbb{H}^2 , and the *projection map* $\pi : \mathbb{H}^2 \rightarrow S$ is a local isometry.

The projection π induces tilings of \mathbb{H}^2 by copies of some fundamental domain for Γ . A *fundamental domain* \mathcal{F} for the action of Γ is defined as a closed domain such that $\Gamma\mathcal{F} = \mathbb{H}^2$ and the interiors of different copies of \mathcal{F} under Γ are disjoint. We also define an *original domain* as a (connected) subset \mathcal{F}_o of a fundamental domain that contains exactly one point of each orbit; then the closure $\overline{\mathcal{F}_o}$ of an original domain \mathcal{F}_o is a fundamental domain. The restriction of π to \mathcal{F}_o is then a bijection from \mathcal{F}_o to S [36].

We use the Poincaré disk model for the hyperbolic plane \mathbb{H}^2 [2]. 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 biconformal 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 \mathbb{H}^2 as hyperbolic circles correspond to Euclidean circles, so that the combinatorial structure of a Delaunay triangulation is equivalent to the Euclidean Delaunay triangulation defined by the same set of points [3].

Given a finite point cloud \mathcal{P} on S , \mathcal{P} lifts to a locally finite¹ point cloud $\tilde{\mathcal{P}}$ in the covering space \mathbb{H}^2 . The Delaunay triangulation $\text{DT}_{\tilde{\mathcal{P}}}$ defined by $\tilde{\mathcal{P}}$ in \mathbb{H}^2 projects to a triangulation $\text{DT}_{\mathcal{P}}$ on S , which serves as a definition for the Delaunay triangulation of \mathcal{P} on S [10, 14]. We do not assume triangulations to be simplicial complexes in this paper, in contrast to some previous work [4, 26]. In our setting, every finite point cloud on a hyperbolic surface has an associated locally finite Delaunay triangulation [10, Cor 5.2][14, Prop 8].

The Dirichlet domain $\mathcal{D}_{\tilde{x}}^\Gamma$ of a point \tilde{x} can be defined as in the flat case; we will simply denote it as $\mathcal{D}_{\tilde{x}}$, unless there is an ambiguity. Note that, unlike the flat case, the shape and even the combinatorial structure of a Dirichlet domain depends on the chosen point \tilde{x} (see Figure 5). This is because NEC groups are non-Abelian, in contrast to the above situation in \mathbb{E}^2 . Indeed, for any isometry f and two points \tilde{x} and $f(\tilde{x})$, there is a relation between the Dirichlet domains: $\mathcal{D}_{f(\tilde{x})}^\Gamma = f(\mathcal{D}_{\tilde{x}}^{f^{-1}\Gamma f})$ [1, Section 9.4]. Methods used to treat the Euclidean case depend crucially on the fact

¹A collection of points P in a topological space X is locally finite if every point $x \in X$ admits a neighborhood U_x such that $P \cap U_x$ is finite; if K is compact in X , then $P \cap K$ is finite.




Figure 5: Dirichlet domains of different points for the group of the genus 2 Bolza surface [4, Fig 9].

that the involved groups are Abelian, so we need a new approach and tools to tackle the problem. Note that there has been recent progress on the development of an efficient algorithm to compute a Dirichlet domain for a hyperbolic surface [13].

Notation. Throughout the paper, we use the same notation as above: $S = \mathbb{H}^2/\Gamma$ is a (closed orientable) hyperbolic surface; the group Γ is an NEC group with no fixed point; the projection map is $\pi : \mathbb{H}^2 \rightarrow S$. We denote objects in \mathbb{H}^2 with a tilde, and those on $S = \mathbb{H}^2/\Gamma$ without; \mathcal{P} always denotes a finite set of points on S , and $\tilde{\mathcal{P}}$ the corresponding lifted point set in \mathbb{H}^2 .

Results. Let \mathcal{F}_o denote a fundamental domain (more precisely an original domain, as defined in Section 2). Consider the Delaunay triangulation $DT_{\tilde{\mathcal{P}}}$ of the lifted point set $\tilde{\mathcal{P}}$. Some edges are incident to a point in $\tilde{\mathcal{P}} \cap \mathcal{F}_o$ and a point lying in a translate of \mathcal{F}_o under an element of Γ . In Section 3, we define the *combinatorial length* of an edge, which relates to the number of translates of \mathcal{F}_o an edge traverses. For a general fundamental domain, the combinatorial length is unbounded.

Our main result, stated as Theorem 6 in Section 5, is an explicit upper bound when a Dirichlet domain is chosen as a fundamental domain. The bound is purely topological and depends linearly on the genus of S . Furthermore, we provide a discussion for why an optimal upper bound on the combinatorial length should depend linearly on the genus. We also give bounds on the number of copies of a domain within a given combinatorial distance of that domain and show that this number increases exponentially with the distance.

Our results rely on intersection properties of edges of Delaunay triangulations and Dirichlet domains, studied in Section 4, and on a decomposition of a S into convex subsets, see Section 5. In Section 6, we give bounds on the number of domains within a given combinatorial distance of another domain.

To the best of our knowledge, our results are the first of their kind for general hyperbolic surfaces.

2 Dirichlet domains and Delaunay triangulations

Let us briefly recall a few definitions and basic properties. We refer the reader to textbooks for the background on hyperbolic geometry [40, 42].

We denote by $d_{\mathbb{H}^2}$ the hyperbolic metric on \mathbb{H}^2 . For a locally finite point set $\mathcal{P} \subset \mathbb{H}^2$ and $\tilde{y} \in \mathcal{P}$, we denote the closed Voronoi cell of \tilde{y} by $\mathcal{V}_{\tilde{y}}^{\mathcal{P}}$ and the whole Voronoi diagram by $\mathcal{V}^{\mathcal{P}}$. The Voronoi diagram is a locally finite collection of convex subsets of \mathbb{H}^2 [11, Lemma 5.2].

Definition 1 ([1]). *The (closed) Dirichlet domain $\mathcal{D}_{\tilde{x}}$ of a point \tilde{x} in \mathbb{H}^2 is the cell $\mathcal{V}_{\tilde{x}}^{\Gamma\tilde{x}}$ of \tilde{x} in the Voronoi diagram of the orbit $\Gamma\tilde{x}$.*

The Dirichlet domain $\mathcal{D}_{\tilde{x}}$ can also be defined equivalently as

$$\mathcal{D}_{\tilde{x}} = \{ \tilde{y} \in \mathbb{H}^2 \mid d_{\mathbb{H}^2}(\tilde{x}, \tilde{y}) \leq d_{\mathbb{H}^2}(\tilde{x}, \Gamma\tilde{y}) \} = \{ \tilde{y} \in \mathbb{H}^2 \mid d_{\mathbb{H}^2}(\tilde{x}, \tilde{y}) \leq d_{\mathbb{H}^2}(\Gamma\tilde{x}, \tilde{y}) \}.$$

The equality is true since Γ acts as isometries w.r.t. $d_{\mathbb{H}^2}$. In particular, we see that

$$\tilde{z} \in \mathcal{D}_{\tilde{x}} \iff \tilde{x} \in \mathcal{D}_{\tilde{z}}. \quad (1)$$

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 $\mathcal{D}_{\tilde{x}}$ is a fundamental domain for Γ and since S is compact, $\mathcal{D}_{\tilde{x}}$ is also compact. Therefore, $\mathcal{D}_{\tilde{x}}$ has a finite number of edges and the tessellation $\Gamma\mathcal{D}_{\tilde{x}}$ associated to $\mathcal{D}_{\tilde{x}}$, called the *Dirichlet tessellation* w.r.t. \tilde{x} , is a locally finite tessellation.

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 [3]. Note that the circumcircles of faces in $\text{DT}_{\tilde{\mathcal{P}}}$ are all compact [11, 14]. Though we use the term *triangulation*, we also consider the case where more than three concyclic points form a non-triangulated polygonal face. In such a case, we observe that such a face has vertices on a circle and is thus geodesically convex in \mathbb{H}^2 (since any interior angle is less than π). We can thus triangulate it arbitrarily in a Γ -invariant way. The Delaunay triangulation $\text{DT}_{\tilde{\mathcal{P}}}$ 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 $\tilde{\mathcal{P}} := \pi^{-1}(\mathcal{P})$. Let us also recall the following result:

Proposition 1. ([14, Proposition 8],[10, Corollary 5.2]) *The 1-skeleton of the Delaunay triangulation $\text{DT}_{\mathcal{P}}$ on S is an embedded graph on S .*

When a fundamental domain \mathcal{F} is a polygon, its edges 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. We also choose one representative of each vertex orbit to obtain a set V of vertices. Let $\text{int}(M)$ denote the interior of a set M .

Definition 2. *Let $\mathcal{F} \subset \mathbb{H}^2$ be a polygonal fundamental domain for Γ . An original domain \mathcal{F}_o associated to \mathcal{F} is defined as a subset of \mathcal{F} consisting of $\text{int}(\mathcal{F}) \cup E \cup V$.*

Special cases of Definition 2 have been considered in the literature [9, 26].

3 Combinatorial length

The tiling of \mathbb{H}^2 formed by fundamental domains and its copies under Γ can be decomposed into layers, giving rise to a combinatorial notion of distance associated to a tiling:

Definition 3. *Let $\mathcal{F} \subset \mathbb{H}^2$ be a polygonal fundamental domain for Γ . Let \tilde{x} be a point in a fixed original domain \mathcal{F}_o .*

Consider the set $\{\mathcal{F}_1^i\}_i$ of nontrivial copies, under the action of Γ , of \mathcal{F} with $\mathcal{F}_1^i \cap \mathcal{F} \neq \emptyset$. Each \mathcal{F}_1^i corresponds to a nontrivial $f_1^i \in \Gamma$ such that $f_1^i\mathcal{F}_o \subset \mathcal{F}_1^i$. We call $\mathcal{N}_1 := \bigcup_i f_1^i\mathcal{F}_o$ the first neighborhood layer of \mathcal{F}_o . For a point $\tilde{y} \in \mathcal{N}_1$, the combinatorial distance $d_{\mathcal{F}_o}^c(\tilde{x}, \tilde{y})$ from \tilde{x} to \tilde{y} is equal to 1. We repeat this process inductively. Consider

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 \mathcal{N}_m for $m \leq n-1$. Each \mathcal{F}_n^i corresponds to a nontrivial $f_n^i \in \Gamma$ such that $f_n^i \mathcal{F}_o \subset \mathcal{F}_n^i$. The n th neighborhood layer is defined as $\mathcal{N}_n = \bigcup_i f_n^i \mathcal{F}_o$. The combinatorial distance from \tilde{x} to a point $\tilde{y} \in \mathcal{N}_n$ is $d_{\mathcal{F}_o}^c(\tilde{x}, \tilde{y}) = n$.


The combinatorial length $d_{\mathcal{F}_o}^c(\tilde{e})$ of a geodesic segment \tilde{e} designates the combinatorial distance between its endpoints.

The (maximal) combinatorial length of a triangulation T w.r.t. \mathcal{F}_o is $d_{\mathcal{F}_o}^c(T) = \max_{\tilde{x} \sim \tilde{y}} d_{\mathcal{F}_o}^c(\tilde{x}, \tilde{y})$, where $\tilde{x} \sim \tilde{y}$ if they are joined by an edge in the triangulation.

Remark that two different neighborhood layers have empty intersection.

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.

Figure 6 illustrates the definition of neighborhood layers, in a non-generic case. In this example, the translations that identify opposite edges of the Dirichlet domain of the origin generate the group Γ_B of the Bolza surface. All vertices of the Dirichlet fundamental domain are equivalent under Γ_B .



(a) Tesselation of \mathbb{H}^2 formed by copies of the Dirichlet domain of the origin for Γ_B [4, Fig 9].

(b) The first neighborhood layer (red) built around the Dirichlet domain of the origin (blue).

Figure 6: Dirichlet fundamental domain and tessellation for the Bolza surface.

Assume that the tessellation induced by \mathcal{F} and Γ features only vertices of degree 3, which is the generic case [1, Theorem 9.4.5], in the sense that for almost every point \tilde{x} , $\mathcal{D}_{\tilde{x}}$ has only vertices of degree 3. Then the combinatorial distance between points is equal to the graph distance in the 1-skeleton of the dual triangulation, between the copies of the original domain \mathcal{F}_o that the points lie in.

Let us look back at our objective, as described in the introduction. Using an original domain \mathcal{F}_o , $\text{DT}_{\tilde{\mathcal{P}}}$ can be reconstructed by considering the points $\tilde{\mathcal{P}} \cap \mathcal{F}_o$ and the edges in $\text{DT}_{\tilde{\mathcal{P}}}$ that connect to these. One question related to the combinatorial length of an edge is how many fundamental domains it can intersect. Unfortunately, as in the flat case above, there is no bound on this number of intersections, by Corollary 1 below, which motivates the use of Dirichlet domains. Intuitively, what can happen is that a fundamental domain arises by cutting S along a set of paths that each intersect a given edge arbitrarily many times.

Proposition 2. *Let $n \in \mathbb{N}$ and c a simple nonseparating path on S . There is a fundamental domain for Γ in \mathbb{H}^2 such that the combinatorial length of the lift \tilde{c} of c is greater than n .*

Proof. We construct a fundamental domain \mathcal{F}_1 with one vertex orbit for Γ such that there is a single edge \tilde{c}_1 of \mathcal{F}_1 that intersects \tilde{c} , and does so transversally, both in their interiors, as we now explain. By assumption c is a simple path on S . Extend c arbitrarily to obtain a simple nontrivial nonseparating closed path $c_C \subset S$ if necessary. Since c_C is nontrivial, we can find another simple closed nontrivial nonseparating path $c_1 \subset S$ that intersects c_C exactly once, somewhere in the interior of c . By applying a homeomorphism f of S , we map $c_C \cup c_1$ to the paths on S shown in Figure 7a [19, Section 1.3.3] and consider the graph $f(c_1) \cup v$, with vertex $v \in f(c_1 - c_C)$. Augmenting this graph by attaching $2g - 1$, where g is the genus of S , closed paths c'_2, \dots, c'_{2g} as illustrated in Figure 7b, we obtain a graph G' whose edges only intersect in v . Cutting




Figure 7: The graph G' in the proof of Proposition 2.

open S along the graph G' produces a disk, as G' yields the standard presentation of the fundamental group of a surface [22, p. 5]. Any graph \hat{G} such that $S - \hat{G}$ is a disk gives rise to a fundamental domain in \mathbb{H}^2 for Γ , by taking as interior of the domain a connected component of the preimage $\pi^{-1}(S - \hat{G})$ [34, Theorem 5.1]. Therefore, $G := f^{-1}(G')$ gives rise to the sought-for fundamental domain $\mathcal{F}_1 \subset \mathbb{H}^2$, with $G \cap c = c_1 \cap c$. Thus, $\tilde{c} \cap \partial \mathcal{F}_1 = \tilde{c} \cap \tilde{c}_1$, for an appropriate lift \tilde{c}_1 of c_1 . Consider the Dehn twist t about c_1 . For our purposes, it suffices to note that t can be represented as a homeomorphism of infinite order that maps G to a similar graph and therefore yields another fundamental domain. The Dehn twist t only changes the edges of G with nontrivial intersection number with c_1 , so, by construction, it leaves invariant all edges except for c_2 . The number of intersections of $t^M(c_2)$ with c_C , where $t^M = t \circ t \circ \dots \circ t$ (M -times), is equal to M [19, Proposition 3.4]. Therefore, there is a representative in the isotopy class of $t^M(c_2)$ that intersects c_C M times, without forming any bigons, which by [19, Proposition 1.7] and [19, Section 1.2.7] means that these intersections cannot be eliminated by using homotopies. In \mathbb{H}^2 , \tilde{c} then intersects M distinct copies of $t^M(c_2)$. Therefore, the combinatorial length of the edge \tilde{c} can be made arbitrarily large by successive application of the Dehn twist t . \square

Remark 1. *The statement of Proposition 2 remains valid if one drops the assumption that the simple path is nonseparating, as long as the path is nontrivial. The proof of this statement with the above method would entail different cases according to the genera of the two parts of the surface resulting from cutting it along the given path. For our purposes, the more restricted version of the proposition suffices.*

Corollary 1. *There is no bound on the combinatorial length of an edge in $\text{DT}_{\bar{p}}$ for arbitrary fundamental domains.*

Proof. Fix an edge \tilde{e} of $\text{DT}_{\bar{p}}$ 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 $\text{DT}_{\mathcal{P}}$ is a triangulation, the closed paths in $\text{DT}_{\mathcal{P}}$ generate the fundamental group Γ . By Hurewicz' theorem [22, Theorem 2A.1], the paths in $\text{DT}_{\mathcal{P}}$ are a basis of the homology group of S , and therefore cannot all be separating, because separating paths are homologically trivial. By Proposition 1, \tilde{e} is simple, and Proposition 2 concludes the proof. \square

4 Half-minimizers

At the heart, our approach is based on the fact that two paths that minimize the distance between points cannot intersect too many times. Edges of a Delaunay triangulation or of a Voronoi diagram do not generally have this feature, but they do have related properties.

Definition 4. A distance path between two points on a surface S is a path, i.e. the image of a continuous map from $[0, 1]$ to S , that has the minimum length out of all paths on S with the same endpoints. We also call a lift \tilde{c} in \mathbb{H}^2 of a distance path c in S a distance path.

A distance path γ is necessarily a geodesic on S , but, in general, not all geodesics are distance paths as they only locally minimize distances. Furthermore, the property of being a distance path is inherited by subarcs. A distance path is also necessarily simple.

Definition 5. A path c from x to $y \in S$ is a half-minimizer if it is the concatenation of at most two distance paths. We call the point m where the two distance paths join a half-point of c . We also call a half-minimizer in \mathbb{H}^2 the lift of a half-minimizer in S .

A half-minimizer is smooth except that it may have a kink at the half-point m . If there is such a kink then m is uniquely defined, otherwise this is not necessarily the case. Half-minimizers provide a natural class of paths with distance minimizing properties. In contrast to distance paths, half-minimizers can be closed, and indeed the systole of a hyperbolic surface is an example of a nontrivial half-minimizer.

Let us study intersection properties of distance paths or half-minimizers. The proofs of the following two lemmas are based on general considerations concerning the structure of distance paths in Riemannian manifolds.

Lemma 1. Two distinct distance paths on S that do not have a subarc in common cannot intersect each other more than once in their interior. If an intersection occurs at an endpoint, then there cannot be an intersection in the interior. Moreover, a distance path is necessarily simple.

Note that two distance paths can still share the same two endpoints.

Proof. Let c_1 and c_2 be distance paths between the points x_1 and y_1 , and x_2 and y_2 , respectively. Assume 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 not sharing a subarc 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 γ_1, γ_2 intersected nontransversally at some point p , then $\gamma_1 = \gamma_2$ locally, 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 [38, Chapter 3].

Since both c_1 and c_2 are distance 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 path 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' . By the Hopf-Rinow theorem [38, Theorem 5.21], we can then further shorten the path c_3 by applying a homotopy with fixed endpoints and obtain a path connecting x_1 and y_1 , homotopic to c_3 and shorter than c_1 , a contradiction.

The second statement of the lemma corresponds to the situation where either $x_1 = z$ or $y_1 = z'$. The same argument also shows that a distance path is simple. \square

Lemma 2. *A half-minimizer can intersect a distance path at most 2 times, or the half-minimizer and the distance path have a common subarc.*

Proof. The only case that is not immediate from Lemma 1 is the case where a distance path γ intersects a half-minimizer c at both endpoints x and y and at the half-point m of c . Let γ join x to y , passing through m , and let c_m^x, c_m^y be the two distance paths on c , connecting x to m and m to y , respectively. Now, if c_m^x intersected γ at m nontransversally, then c_m^x would be included in γ , similarly to the proof of Lemma 1 above. If we exclude such cases, then we find a contradiction to γ being a distance path by considering the geodesic in the homotopy class with fixed endpoints of the path that follows first c_m^x from x to m and then γ from m to y . \square

Lemma 2 implies that half-minimizers are either simple or they backtrack from the half-point, in which case their image is contained in the image of a distance path. We now show that the edges of Delaunay triangulations and Dirichlet domains are either half-minimizers or isotopic to a half-minimizer.

Lemma 3. *Let \tilde{x} be a point in \mathbb{H}^2 . For any point \tilde{y} in $\mathcal{D}_{\tilde{x}}$, the geodesic segment $\tilde{\gamma}$ that joins \tilde{x} to \tilde{y} projects to a distance path in S . Conversely, if a geodesic γ from $x = \pi(\tilde{x})$ to y is a distance path on S , then the lift $\tilde{\gamma}$ of γ based at \tilde{x} is contained in $\mathcal{D}_{\tilde{x}}$.*

Proof. As $\mathcal{D}_{\tilde{x}}$ is convex, a geodesic $\tilde{\gamma}$ between \tilde{x} and $\tilde{y} \in \mathcal{D}_{\tilde{x}}$ is contained in $\mathcal{D}_{\tilde{x}}$. Assume, for the sake of contradiction, that there is a simple geodesic γ' on S joining $x = \pi(\tilde{x})$ and $y = \pi(\tilde{y})$, shorter than $\gamma = \pi(\tilde{\gamma})$. Since geodesics connecting points are uniquely determined by their endpoints in their homotopy class in S [5, Theorem 1.5.3], γ' is not homotopic to γ . So, the lift $\tilde{\gamma}'$ of γ' based at \tilde{x} is not equal to $\tilde{\gamma}$ and therefore joins \tilde{x} to a point $\tilde{y}' \neq \tilde{y}$, equivalent to \tilde{y} under the action of Γ . Moreover, $\tilde{\gamma}'$ is shorter than $\tilde{\gamma}$ because π is a local isometry and therefore the length of a geodesic in \mathbb{H}^2 and its projection in S are equal [20, Proposition 2.109], a contradiction.

For the converse statement, simply observe that if the endpoint \tilde{y} of $\tilde{\gamma}$ was not contained in $\mathcal{D}_{\tilde{x}}$, then it would be strictly closer to another point $x' \in \Gamma\tilde{x}$. This contradicts the minimality of γ , by the first part of the proof, since then there is a distance path $\tilde{\gamma}'$ in $\mathcal{D}_{\tilde{x}'}$ that projects to a distance path γ' shorter than γ . \square

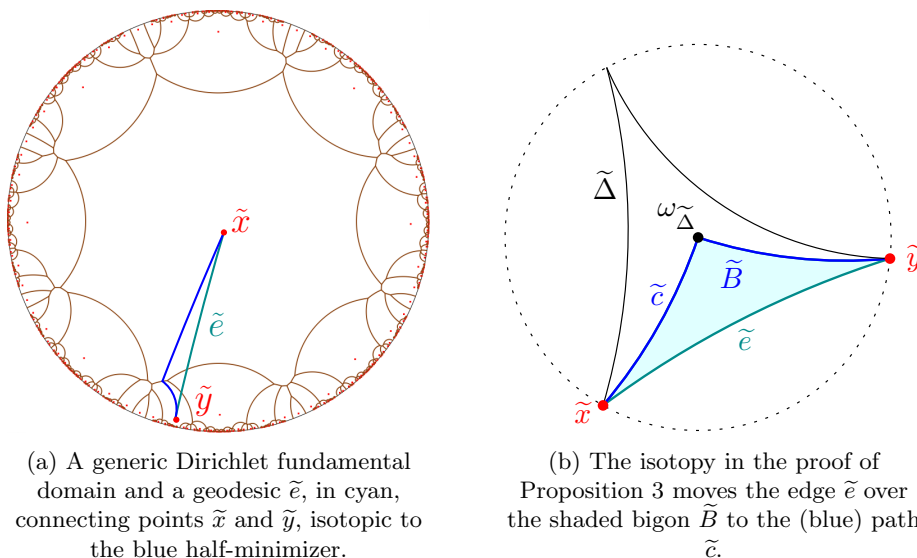
The following lemma generalizes a result known in the flat case [15]. In particular, when $\tilde{\mathcal{P}} = \Gamma\tilde{x}$, the circumcenter lies in the Dirichlet domain $\mathcal{D}_{\tilde{x}}$, as in the flat case [15, Lemma 3.2].

Lemma 4. *Let \tilde{x} be a point in $\tilde{\mathcal{P}}$ and $\tilde{\Delta}$ a triangle of $\text{DT}_{\tilde{\mathcal{P}}}$ with vertex \tilde{x} . The (hyperbolic) circumcenter of $\tilde{\Delta}$ lies in $\mathcal{V}_{\tilde{x}}^{\tilde{\mathcal{P}}}$.*

Proof. Let $\omega_{\tilde{\Delta}}$ denote the circumcenter of $\tilde{\Delta}$ and $C_{\tilde{\Delta}}$ its circumcircle. 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 $\text{int}(\mathcal{V}_{\tilde{y}}^{\tilde{\mathcal{P}}})$

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 2. A Delaunay edge \tilde{e} is not generally a half-minimizer, as Figure 8a illustrates. Indeed, consider the set $\Gamma\tilde{x}$, for a point $\tilde{x} \in \mathbb{H}^2$. If a Delaunay edge \tilde{e} incident to \tilde{x} and some $\tilde{x}' \in \Gamma\tilde{x}$ projects to a (closed) half-minimizer in S , then the midpoint of \tilde{e} must lie on the boundary $\partial\mathcal{D}_{\tilde{x}}$ by Lemma 3. Said in another way, such a half-minimizer intersects the interior of exactly two Dirichlet domains.



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

(b) The isotopy in the proof of Proposition 3 moves the edge \tilde{e} over the shaded bigon \tilde{B} to the (blue) path \tilde{c} .

Figure 8: Illustrations for Remark 2 (left) and Proposition 3 (right).

Proposition 3. Let $x \in \mathcal{P}$ and e an edge of $\text{DT}_{\mathcal{P}}$ incident to x . The edge e is isotopic with fixed endpoints to a half-minimizer c on S , based at x . The path c is simple and, if closed, nontrivial.

Proof. In \mathbb{H}^2 , let $\tilde{\Delta}$ be a triangular face of $\text{DT}_{\tilde{\mathcal{P}}}$, and $\omega_{\tilde{\Delta}}$ its circumcenter. Consider the path \tilde{c} that connects $\omega_{\tilde{\Delta}}$ to two vertices \tilde{x} and \tilde{y} of $\tilde{\Delta}$, as shown in Figure 8b. The point $\omega_{\tilde{\Delta}}$ lies in $\mathcal{D}_{\tilde{x}}$, by Lemma 4 and because $\Gamma\tilde{x} \subset \tilde{\mathcal{P}}$ implies that $\mathcal{V}_{\tilde{x}}^{\tilde{\mathcal{P}}} \subset \mathcal{D}_{\tilde{x}}$. Similarly, $\omega_{\tilde{\Delta}}$ lies in $\mathcal{D}_{\tilde{y}}$. So, by Lemma 3, \tilde{c} is a half-minimizer based at either endpoint, with half-point $\omega_{\tilde{\Delta}}$.

Consider now the Dirichlet domain $\mathcal{D}_{\omega_{\tilde{\Delta}}}$, which contains \tilde{x} and \tilde{y} , by equivalence (1) (page 7) and Lemma 4. Since $\mathcal{D}_{\omega_{\tilde{\Delta}}}$ is convex, it also contains the geodesic edge \tilde{e} of $\tilde{\Delta}$ with the same endpoints, \tilde{x} and \tilde{y} , as \tilde{c} . The bigon \tilde{B} formed by \tilde{c} and \tilde{e} is then completely contained in $\mathcal{D}_{\omega_{\tilde{\Delta}}}$. Therefore, there is an isotopy inside $\mathcal{D}_{\omega_{\tilde{\Delta}}}$ from \tilde{e} to \tilde{c} , which fixes \tilde{x} and \tilde{y} . Note that \tilde{B} does not contain any points of $\tilde{\mathcal{P}}$ (except for \tilde{x} and \tilde{y}), by definition of $\text{DT}_{\tilde{\mathcal{P}}}$ and convexity of the disk circumscribing $\tilde{\Delta}$. The tessellation of \mathbb{H}^2 by copies of $\mathcal{D}_{\omega_{\tilde{\Delta}}}$ is of course Γ -invariant. Since $\mathcal{D}_{\omega_{\tilde{\Delta}}}$ contains only one representative of every point moved in the isotopy, one obtains a Γ -invariant isotopy of the whole plane \mathbb{H}^2 by using the same constructed isotopy in every copy of $\mathcal{D}_{\omega_{\tilde{\Delta}}}$.

That $c = \pi(\tilde{c})$, if closed, is nontrivial is clear, and simple follows from Proposition 1 for $\text{DT}_{\mathcal{P}}$ and the fact that c is isotopic to the projection of \tilde{e} . \square

Definition 6. Let \tilde{e} be an edge of either $\text{DT}_{\tilde{\mathcal{P}}}$ or $\mathcal{V}^{\tilde{\mathcal{P}}}$. The edge \tilde{e} is said to be centered if it intersects its dual edge.

The concept of centered edges has been studied before [10], restricted to Voronoi edges.

Proposition 4. Each non-centered edge of $\mathcal{D}_{\tilde{x}}$ is a distance path and each centered edge is a half-minimizer.

Proof. Let \tilde{e} be an edge of $\mathcal{D}_{\tilde{x}}$, with endpoints by \tilde{a} and \tilde{b} , and let $\mathcal{D}_{\tilde{x}_1}$ be the other domain in $\Gamma\mathcal{D}_{\tilde{x}}$ that is also incident to \tilde{e} . The group of conformal transformations of the unit disk acts transitively on triples on the boundary of the unit disk, so 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 this axis. Furthermore, the points \tilde{x} and \tilde{x}_1 lie on a geodesic perpendicular to the real axis, so by mapping the intersection above the real axis of this geodesic with the boundary to the point $i \in \mathbb{C}$, we can assume that the points \tilde{x} and \tilde{x}_1 lie on the imaginary axis.

For any point $\tilde{z} \in \tilde{e}$, denote by $\tilde{\gamma}_z$ ($(\tilde{\gamma}_z)_1$) the geodesic segment from \tilde{x} (\tilde{x}_1) to \tilde{z} , respectively. Then the concatenation of $\tilde{\gamma}_z \subset \mathcal{D}_{\tilde{x}}$ with $(\tilde{\gamma}_z)_1 \subset \mathcal{D}_{\tilde{x}_1}$ forms a half-minimizer from \tilde{x} to \tilde{x}_1 , by Lemma 3.

We consider first the case that \tilde{e} is non-centered and assume, without loss of generality, that $d_{\mathbb{H}^2}(\tilde{a}, \tilde{x}) \leq d_{\mathbb{H}^2}(\tilde{b}, \tilde{x})$. See Figure 9(Left). Note that the circle $C^0(\tilde{b})$ passing through \tilde{b} and centered at \tilde{x} bounds a disk $D^0(\tilde{b})$ the interior of which does not contain any points in $\Gamma\tilde{b}$, by Lemma 3. 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. The edge \tilde{e} is necessarily shorter than $\tilde{\gamma}_b$, by non-centeredness. Indeed the circles 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 , $2l(\tilde{e}) < 2l(\tilde{\gamma}_b)$, where $l(c)$ denotes the hyperbolic length of a geodesic segment c . Whence, one easily sees that the cyan circle, centered at \tilde{a} and passing through \tilde{b} , is contained in the union of $D^0(\tilde{b})$ and $D^1(\tilde{b})$ and consequently does not contain any point of $\Gamma\tilde{b}$. Therefore, \tilde{e} is a distance path.




Figure 9: Proof of Proposition 4: (Left) non-centered case (Right) centered case.

Consider now the case where \tilde{e} is centered. See Figure 9(Right). Projecting \tilde{x} (and \tilde{x}_1) to its nearest point on \tilde{e} , we obtain a point $\tilde{m} \in \tilde{e}$ between \tilde{a} and \tilde{b} . We claim that \tilde{m} 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, which do not contain any point equivalent to \tilde{b} . Their boundary circles are shown in the figure along with the green circle $C(\tilde{b})$ centered at \tilde{m} and passing through \tilde{b} . The




Figure 10: The two triangles $\tilde{\Delta}$ and $\tilde{\Delta}'$ in the proof of Proposition 5.

interior disk $D(\tilde{b})$ of $C(\tilde{b})$ is included in the union $D^0(\tilde{b}) \cup D^1(\tilde{b})$, which shows that the subarc of \tilde{e} from \tilde{m} to \tilde{b} is a distance path. Consider now the disks bounded instead by the dashed circles, centered at \tilde{x} and \tilde{x}_1 , respectively, and passing through \tilde{a} . One deduces, similarly to above, that the cyan dashed circle centered at \tilde{m} passing through \tilde{a} bounds a disk not containing any points of $\Gamma\tilde{a}$. Therefore, the subarc of \tilde{e} from \tilde{m} to \tilde{a} is a distance path. \square

5 Main result

There are several definitions of convexity on a surface in the literature. We choose the following definition, adapted to our case study.

Definition 7. *A simply connected subset $K \subset S$ is convex if for every two points $x, y \in K$, there is a unique distance path in K with endpoints x and y .*

Proposition 5. *A (closed) triangle in S whose edges are distance paths is convex.*

Proof. For a triangle Δ in S , fix a lift $\tilde{\Delta} \subset \mathbb{H}^2$. The triangle $\tilde{\Delta}$ is obviously convex in \mathbb{H}^2 ; it contains a unique geodesic joining any two points in it. We will prove that, if the edges of Δ are distance paths, any such geodesic is also a distance path. Denote the geodesic connecting two points \tilde{z}_1 and \tilde{z}_2 in \mathbb{H}^2 by $\tilde{\gamma}_{\tilde{z}_1\tilde{z}_2}$.

We first prove that the (unique) geodesics in $\tilde{\Delta}$ from corners to arbitrary points on its edges are distance paths. Denote the vertices of $\tilde{\Delta}$ by \tilde{x}, \tilde{y} , and \tilde{z} . Consider w.l.o.g. the corner \tilde{y} . The edges $\tilde{\gamma}_{\tilde{y}\tilde{z}}$ and $\tilde{\gamma}_{\tilde{y}\tilde{x}}$ of $\tilde{\Delta}$ are distance paths so $\tilde{z}, \tilde{x} \in \mathcal{D}_{\tilde{y}}$ by Lemma 3. Since $\mathcal{D}_{\tilde{y}}$ is convex, $\tilde{\gamma}_{\tilde{x}\tilde{z}} \subset \mathcal{D}_{\tilde{y}}$ and therefore, again by Lemma 3, there is a distance path from \tilde{y} to any point on the edge $\tilde{\gamma}_{\tilde{x}\tilde{z}}$.

Assume that there are two points $\tilde{p}'_1, \tilde{p}'_2 \in \tilde{\Delta}$ such that $\tilde{\gamma}_{\tilde{p}'_1\tilde{p}'_2}$ is not a distance path. The complete geodesic through \tilde{p}'_1 and \tilde{p}'_2 that is inside $\tilde{\Delta}$ crosses the boundary of $\tilde{\Delta}$ in two points \tilde{p}_1 and \tilde{p}_2 such that the part of the geodesic between \tilde{p}_1 and \tilde{p}_2 contains \tilde{p}'_1 and \tilde{p}'_2 . This geodesic cannot be a distance path between \tilde{p}_1 and \tilde{p}_2 if its subpath between \tilde{p}'_1 and \tilde{p}'_2 is not a distance path. For the proof of the proposition, it is therefore enough to show that geodesics in $\tilde{\Delta}$ that connect two points of the boundary are distance paths.

So, assume that $\tilde{p}_1 \in \tilde{\gamma}_{\tilde{x}\tilde{z}}$ and $\tilde{p}_2 \in \tilde{\gamma}_{\tilde{y}\tilde{z}}$, as in Figure 10, and consider the geodesic $\tilde{\gamma}_{\tilde{p}_1\tilde{y}} \subset \tilde{\Delta}$. Since \tilde{y} is a corner of $\tilde{\Delta}$, $\tilde{\gamma}_{\tilde{p}_1\tilde{y}}$ is a distance path and therefore the triangle $\tilde{\Delta}' \subset \tilde{\Delta}$ with vertices \tilde{p}_1, \tilde{y} , and \tilde{z} is a triangle with distance paths as edges. The geodesic $\tilde{\gamma}_{\tilde{p}_1\tilde{p}_2}$ connects a corner of $\tilde{\Delta}'$ to a point on its edge, so it is a distance path by the first part of the proof. \square

We can partition a Dirichlet domain into triangles whose projections to S are convex. The notation is illustrated in Figure 11a in the case of a dodecagon. Let $\{\tilde{v}_i\}_{i=0}^{k-1}$, with $\tilde{v}_k = \tilde{v}_0$, denote the k corners of $\mathcal{D}_{\tilde{x}}$, indexed counter-clockwise, and

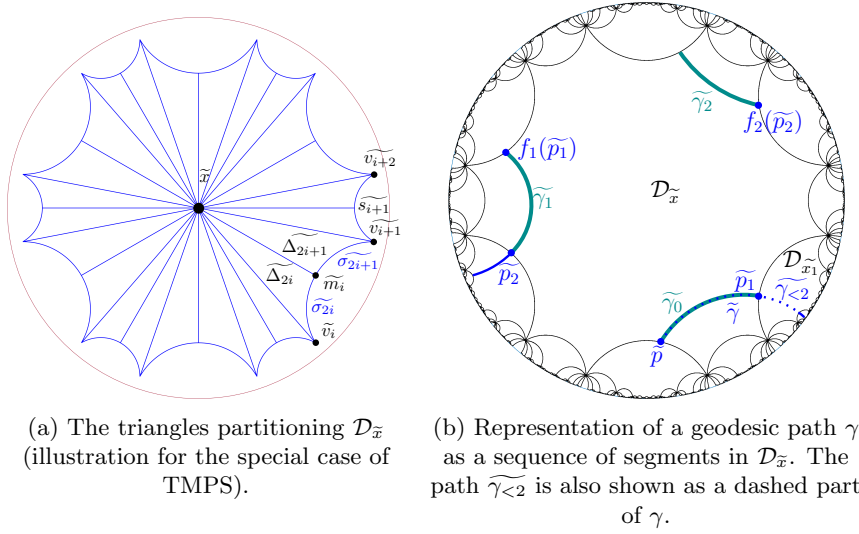


Figure 11: Partitioning into triangles and paths on the surface represented as a sequence of paths in a fundamental domain. The fundamental group identifies opposite edges of the dodecagon.

let \tilde{s}_i denote the side with endpoints \tilde{v}_i and \tilde{v}_{i+1} , for $i = 0, \dots, k-1$. We know by Proposition 4 that \tilde{s}_i is a half-minimizer or a distance path; in the first case we denote as \tilde{m}_i a half-point of \tilde{s}_i ; in the second case we choose an arbitrary point \tilde{m}_i in the interior of \tilde{s}_i . The two subarcs of \tilde{s}_i obtained in this way are distance paths; let us denote them as $\tilde{\sigma}_{2i}$, with endpoints \tilde{v}_i and \tilde{m}_i , and $\tilde{\sigma}_{2i+1}$, with endpoints \tilde{m}_i and \tilde{v}_{i+1} , respectively. Denote now as $\tilde{\gamma}_z$ the geodesic segment between \tilde{x} and \tilde{z} for any point $z \in \partial\mathcal{D}_{\tilde{x}}$; by Lemma 3, $\tilde{\gamma}_{v_i}$ and $\tilde{\gamma}_{m_i}$ are distance paths, for $i = 0, \dots, k-1$. Let us denote the closed triangle formed by \tilde{x} and $\tilde{\sigma}_{2i}$ as Δ_{2i} and the triangle formed by \tilde{x} and $\tilde{\sigma}_{2i+1}$ as Δ_{2i+1} ; the triangles Δ_j , $j = 0, \dots, 2k-1$ partition $\mathcal{D}_{\tilde{x}}$. By Proposition 5, the projection Δ_j of each Δ_j on S is convex.

Lemma 5. *Let $\mathcal{D}_{\tilde{x}}$ be a Dirichlet domain with k edges. The combinatorial length of a distance path with an endpoint in the original domain $(\mathcal{D}_{\tilde{x}})_o$ is at most $k/2$.*

Before we prove the result, let us observe that a geodesic path $\tilde{\gamma}$ in \mathbb{H}^2 that projects to a path γ on S can be represented as a sequence of geodesic segments in a Dirichlet domain, as illustrated in Figure 11b. Let $\tilde{p} \in (\mathcal{D}_{\tilde{x}})_o$ be an endpoint of $\tilde{\gamma}$. If $\tilde{\gamma} \subset (\mathcal{D}_{\tilde{x}})_o$ then the sequence is reduced to $\{\tilde{\gamma}\}$. Otherwise, denote as $\tilde{\gamma}_0$ the intersection $\tilde{\gamma} \cap (\mathcal{D}_{\tilde{x}})_o$. The path $\tilde{\gamma}$ exits $\mathcal{D}_{\tilde{x}}$ at the intersection point \tilde{p}_1 : $\{\tilde{p}_1\} = \tilde{\gamma} \cap \partial\mathcal{D}_{\tilde{x}}$. The path γ continues in a Dirichlet domain $\mathcal{D}_{\tilde{x}_1}$. By using the appropriate element of $f_1 \in \Gamma$, one can map $\tilde{\gamma} \cap (\mathcal{D}_{\tilde{x}_1})_o$ back into $\mathcal{D}_{\tilde{x}}$ and obtain another geodesic segment $\tilde{\gamma}_1$, which may again exit $\mathcal{D}_{\tilde{x}}$. Repeating this process until we reach the other endpoint of γ yields a collection $\{\tilde{\gamma}_n\}_{n=0}^{r-1}$ of geodesic arcs in $\mathcal{D}_{\tilde{x}}$. If we keep track of the order of the paths, then they project to a sequence $\{\gamma_n\}_{n=0}^{r-1}$ of paths in S whose concatenation corresponds to γ . Note that the geodesic arcs are disjoint in $\mathcal{D}_{\tilde{x}}$ if and only if $\tilde{\gamma}$ projects to a simple path on S . We denote with $\gamma_{<l}$ the concatenation of the l first arcs $\gamma_0, \dots, \gamma_{l-1}$ and with $\tilde{\gamma}_{<l}$ the lift of $\gamma_{<l}$ to \mathbb{H}^2 , starting at \tilde{p} . Observe that $\tilde{\gamma}_{<r} = \tilde{\gamma}$.

Proof. Let $\tilde{\gamma}$ be a distance path with endpoint $\tilde{p} \in (\mathcal{D}_{\tilde{x}})_o$. We can represent $\tilde{\gamma}$ by a sequence of subarcs $\{\tilde{\gamma}_n\}_{n=0}^{r-1}$ in $\mathcal{D}_{\tilde{x}}$ as above. Assuming that $\tilde{\gamma}$ consists of more than one segment, every $\tilde{\gamma}_n$ is incident to two points on $\partial\mathcal{D}_{\tilde{x}}$, except possibly $\tilde{\gamma}_0$ and $\tilde{\gamma}_{r-1}$.

Observe that since $\tilde{\gamma}$ is a distance path, the intersection $\gamma \cap \text{int}(\Delta_j)$ is connected for all j . Indeed, if it were not connected, then there is a distance path that connects two connected components of the intersection by the convexity of the triangles Δ_j in S (Proposition 5). This distance path would intersect the distance path γ more than once, which is impossible by Lemma 1. We similarly see that the only way $\gamma \cap \Delta_j$ is not connected is if γ intersects one triangle in only two points on the boundary, at the start and endpoint of γ . Therefore, aside from the start and endpoint, if $\tilde{\gamma}_n$ intersects a triangle Δ_j for some n , then $\tilde{\gamma}_l$ cannot intersect the same Δ_j for $l \neq n$.

Recall the definition of neighborhood layers (Definition 3) and Figure 6b: if after leaving the (blue) central Dirichlet domain, a path intersects consecutive sides of Dirichlet domains sharing the same corner of this blue domain, then it stays in the red layer. More generally, if the two endpoints of a segment $\tilde{\gamma}_l$ lie on two adjacent sides of $\mathcal{D}_{\tilde{x}}$, then the combinatorial length $d^c(\tilde{\gamma}_{<l+1})$ is the same as the combinatorial length $d^c(\tilde{\gamma}_{<l})$. So, as we want to find an upper bound on the combinatorial length of a Delaunay edge, we will assume that the two endpoints of each segment $\tilde{\gamma}_n$, $n = 1, \dots, r-2$ lie on non-consecutive sides of $\mathcal{D}_{\tilde{x}}$, so that the combinatorial length of $\tilde{\gamma}$ is $r-1$.

Remark now that such a segment $\tilde{\gamma}_n$, $n \in \{1, \dots, r-2\}$ must intersect at least four triangles $\Delta_{2j-1}, \Delta_{2j}, \Delta_{2j+1}$, and $\Delta_{2(j+1)}$ for some j , where the indices are taken modulo $2k$. (Note that in the special case when $\tilde{\gamma}_n$ passes through a corner of a Dirichlet domain, the number of triangle orbits intersected by $\tilde{\gamma}_n$ is at least 6 as the triangles Δ_j are closed.) There are $2k$ triangles Δ_j in total, and a triangle cannot be intersected by more than one segment $\tilde{\gamma}_n$, $n \in \{0, \dots, r-2\}$. Moreover, $\tilde{\gamma}_0$ intersects at least one triangle. Thus, $r-2 \leq (2k-1)/4 = k' - 1/4$, as $k = 2k'$ is even. Since r and k' are integers, in fact $r-2 \leq k' - 1$, and the combinatorial length $r-1$ is thus at most $k/2$. \square

The previous lemma gives an upper bound, depending only on the genus, in the restricted case where all its edges are distance paths on S . Note that this case corresponds to the framework of previous work [26, 4, 39]. Indeed, the condition there is that each Delaunay edge is smaller than half the length of the smallest noncontractible loop on S ; then any path in S joining two points of \mathcal{P} not homotopic to a Delaunay edge e between these points is strictly longer than e , which implies that all edges in $\text{DT}_{\mathcal{P}}$ are distance paths.

Corollary 2. *Let $S = \mathbb{H}^2/\Gamma$ be a genus g hyperbolic surface and $\mathcal{D}_{\tilde{x}}$ be a Dirichlet domain for the NEC group Γ . For a set of points \mathcal{P} on S , if the edges of $\text{DT}_{\mathcal{P}}$ are distance paths, then the maximal combinatorial length of edges of $\text{DT}_{\mathcal{P}}$ in \mathbb{H}^2 is at most $6g - 3$.*

Proof. The number of edges of a Dirichlet fundamental domain is at most $12g - 6$, by a direct use of the Euler characteristic [1, Theorem 10.5.1]. The result follows from Lemma 5. \square

We now consider the general case where there are no restrictions on the Delaunay edges of the triangulations.

Theorem 6. *Let $\mathcal{D}_{\tilde{x}}$ be a Dirichlet domain, with k edges. 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 $\tilde{\mathcal{P}} \subset \mathbb{H}^2$. The maximal combinatorial length of a Delaunay edge \tilde{e} is at most k .*

Proof. Consider any path \tilde{c} joining $\tilde{a} \in (\mathcal{D}_{\tilde{x}})_o$ and $\tilde{b} \in \mathbb{H}^2$. The geodesic from \tilde{a} and \tilde{b} in \mathbb{H}^2 is unique; it projects to a geodesic on the surface which is also unique in its homotopy class (with fixed endpoints) [5, Theorem 1.5.3]. So, the homotopy class of c is uniquely determined by which copy of the original domain $(\mathcal{D}_{\tilde{x}})_o$ contains \tilde{b} .

As a consequence, if the combinatorial length of \tilde{e} were larger than the sum of the combinatorial lengths of two distance paths, then it could not be isotopic to the

concatenation of two such paths. However, the projection to S of a Delaunay edge \tilde{e} in \mathbb{H}^2 is isotopic to a half-minimizer by Proposition 3. The result now follows directly from Lemma 5. \square

Corollary 3. *Let $S = \mathbb{H}^2/\Gamma$ be a genus g hyperbolic surface and $\mathcal{D}_{\tilde{x}}$ be a Dirichlet domain for the NEC group Γ . For a set of points \mathcal{P} on S , the maximal combinatorial length of edges of $\text{DT}_{\tilde{\mathcal{P}}}$ in \mathbb{H}^2 is at most $12g - 6$.*

In light of the discussion below, it seems likely that our linear bound on the combinatorial distance is of the optimal order.

Using the above results, we can produce a natural set of lifts of points on the surface that are sufficient for the construction of all Delaunay edges: the union of all domains in the first $12g - 6$ neighborhood layers. However, this set would also contain points leading to nonsimple geodesics and the number of domains is exponential in the genus. Precise expressions for bounds are deduced in Section 6 below.

Remark 3. *An interesting observation worth mentioning arises in the context of similar problems for more general metric spaces. Since our approach is based on general geometric and topological considerations as well as the notion of half-minimizers, the methods of this section can also be applied to more general situations than that of hyperbolic surfaces. For example, it appears that half-minimizers can also be exploited to find similar upper bounds on the combinatorial length of Delaunay edges w.r.t. Dirichlet domains for higher dimensional hyperbolic manifolds. Moreover, the approach also lends itself to the study of singular spaces, such as orbifolds.*

Discussion of the optimality of our bounds

We were unable to exhibit an explicit sequence of Dirichlet domains for which the combinatorial length of Delaunay edges is linear in the genus and indeed, there are many open questions surrounding the structure of Dirichlet domains for surfaces. However, we can give arguments why such domains should exist.

The boundary of a fundamental domain projects to a graph G on S . General results on the structure of G state that it may be obtained through the following process. Select $2g$ simple closed curves that cut the surface into a disk and are pairwise disjoint except at a point $p \in S$. Choose a neighborhood of p that intersects the $2g$ curves in $4g$ points on the boundary and replace the part of the curves in the neighborhood with an embedded tree T that spans these $4g$ points. The graph G is the union of a choice of $2g$ curves and T [34, Theorem 5.1]. It is further known that any such embedded graph appears, up to homeomorphisms of S , as the edge-graph of some Dirichlet polygon [35, Theorem 8.1] for a group of isometries isomorphic to Γ . There are thus no restrictions on the embedded graph, up to homeomorphisms of S , when restricting to fundamental domains that appear as Dirichlet domains for Γ . The most common case is that where every vertex of the Dirichlet tessellation has degree 3, in which case T has $4g - 2$ vertices (outside of the $4g$ vertices on the boundary of the disk), which follows from a computation using the Euler characteristic.

The structure of G is closely related to the arrangement of neighborhood layers around a central domain. Figure 12 shows an example of a Dirichlet domain and a neighborhood of a vertex containing the tree T associated to the fundamental domain. The domain $\mathcal{D}_{\tilde{x}}$ shown in Figure 12(left) is a slight perturbation of the domain in Figure 6a. The length of the edges of a Dirichlet domain $\mathcal{D}_{\tilde{x}}$ depend continuously on \tilde{x} , so that the edges of the tree T can be made arbitrarily small. In particular, the geodesic shown in Figure 12(right) can be made arbitrarily small, while maintaining a combinatorial length of 3. For a related phenomenon occurring in the flat case, see Figure 4.

We consider the operation of adding two vertices of degree 3 to each of two branches of T and the impact on the combinatorial distance between $\mathcal{D}_{\tilde{x}}$ and domains with nonempty intersection with T , as illustrated for one branch in Figure 13(left). The depicted operation realizes the transition from the tree in Figure 12(right) to that in Figure 13(right), with the effect of adding one to the genus of the hyperbolic surface (here, this corresponds to transitioning from a genus 2 to a genus 3 surface). Topologically, the operation consists of gluing a handle to S in such a way that the 4 new vertices on the boundary of the disk (2 vertices are added to each branch) are connected by two new disjoint curves traversing the added handle, with one end of the handle being attached to S to the left of the shown disk, and the other to the right. The new curves are chosen such that the collection of curves again cut the new hyperbolic surface into a disk. We observe that the operation increases the maximal combinatorial distance between two points in the shown neighborhood, as exemplified with the points \tilde{v}_1 and \tilde{v}_2 in Figure 13(right). To see this, assume, by induction, that the combinatorial distances are as indicated in Figure 13(left) before the insertion, so that the combinatorial length of a path between a point \tilde{v}_2 in the region marked by $d + 2$ and a point \tilde{v}_1 in the original domain is at least $d + 1$. Recall that, since all vertices are of degree 3, the combinatorial distance between two points located inside two copies of $\mathcal{D}_{\tilde{x}}$ corresponds to the distance between the vertices in the dual graph associated to the two copies of $\mathcal{D}_{\tilde{x}}$. If there was another path in the dual graph from \tilde{v}_1 to \tilde{v}_2 that is shorter than $d + 2$ then it would have to avoid copies of $\mathcal{D}_{\tilde{x}}$ with a vertex incident to the tree, by the induction hypothesis. However, since each domain has $12g - 6$ edges, any such path must be longer than the path along copies incident to the tree, since the tree has $4g - 2$ vertices. We have thus illustrated a process that yields a fundamental domain for a surface of genus g that admits geodesics contained in a (small) neighborhood of a vertex with a combinatorial length of $g + 1$.




Figure 12: The tree around a vertex of a Dirichlet polygon for the genus 2 Bolza surface.

To summarize, if the edges of T can be made to be sufficiently small in a Dirichlet domain, then the combinatorial distance between two points that are very close generally depends at least linearly on the genus. Furthermore, this phenomenon occurs even for Dirichlet domains of the most well-behaved surfaces, such as the Bolza surface in Figure 12(left), by a slight shift of the point \tilde{x} for the construction of $\mathcal{D}_{\tilde{x}}$. A good candidate for further investigations would be the class of hyperbolic surfaces of genus g defined by a regular polygon with $4g$ edges centered around the origin, with Dirichlet domains of points close to the origin.

The above discussion shows one way in which the combinatorial length of short




Figure 13: Splitting a branch of the tree to increase the maximal combinatorial distance of geodesics contained in the neighborhood.

geodesics can depend linearly on the genus. What is missing from this is a demonstration that the combinatorial structure of the tree can be obtained as a tree with very short edges, as in the genus 1 and genus 2 examples above. In light of the above analysis, we conjecture that the combinatorial length of Delaunay edges depends linearly on the genus g . Note that the reasoning remains true in the case where we require the point \tilde{x} to be part of $\tilde{\mathcal{P}}$, as the distance from \tilde{v}_1 to \tilde{x} is much larger than the distance between \tilde{v}_1 and \tilde{v}_2 . In particular, we observe that adding points to a Delaunay triangulation can increase the combinatorial length of the resulting Delaunay triangulation.

6 On the number of copies within a given combinatorial distance

To conclude our investigations, we deduce explicit bounds on the number of domains within a given combinatorial distance from a central original domain. The following proposition demonstrates the result based on our upper bound on the combinatorial distance from the previous section. It will follow directly from the results we prove in the remainder of this section.

Proposition 7. *The number N of copies of $\mathcal{D}_{\tilde{x}}$ that have a combinatorial distance of at most $12g - 6$ from $\mathcal{D}_{\tilde{x}}$ satisfies, for $g \geq 2$,*

$$\frac{g^{12g-5} - 1}{g - 1} \leq N \leq 4g \frac{(4g - 2)^{24g-10} - 1}{4g - 3}.$$

We first prove lemmas for two different cases - the worst case, where all vertices of the fundamental domain are equivalent, and the generic case, where all vertices have degree 3.

Lemma 6. *In the generic case, where all vertices of $\Gamma\mathcal{D}_{\tilde{x}}$ have degree 3, the number $|\mathcal{N}_{\leq n}|$ of copies of $\mathcal{D}_{\tilde{x}}$ that have a combinatorial distance of at most $n \geq 1$ from $\mathcal{D}_{\tilde{x}}$ is $1 + 3n(n + 1)$ for $g = 1$ and satisfies*

$$\frac{g^{n+1} - 1}{g - 1} \leq |\mathcal{N}_{\leq n}| \leq \frac{(12g - 10)^{n+1} - 1}{12g - 11}$$

for $g \geq 2$.

Proof. We inductively count the copies in the first few layers until a pattern emerges.

Observe that the combinatorial distance between two points in the interior of two copies of $\mathcal{D}_{\tilde{x}}$ is equal to the graph distance between the vertices representing these two copies in the dual tessellation, as all vertices have degree 3 in the generic case. The zeroth layer is just the original domain $(\mathcal{D}_{\tilde{x}})_o$. The first neighborhood layer \mathcal{N}_1 consists of $12g - 6$ copies of $(\mathcal{D}_{\tilde{x}})_o$, since there is one copy for each of the $12g - 6$ edges. Observe that every copy of $(\mathcal{D}_{\tilde{x}})_o$ in \mathcal{N}_1 has two edges that are incident to \mathcal{N}_1 itself and one edge incident to $(\mathcal{D}_{\tilde{x}})_o$. In general, all edges incident to a copy of $(\mathcal{D}_{\tilde{x}})_o$ in the previous or current neighborhood layer do not lead to a new copy of $(\mathcal{D}_{\tilde{x}})_o$ in the next layer and all other edges do. Consider now an edge \tilde{e} incident to two copies of $(\mathcal{D}_{\tilde{x}})_o$ in \mathcal{N}_1 , as illustrated in Figure 14 in the Euclidean plane.

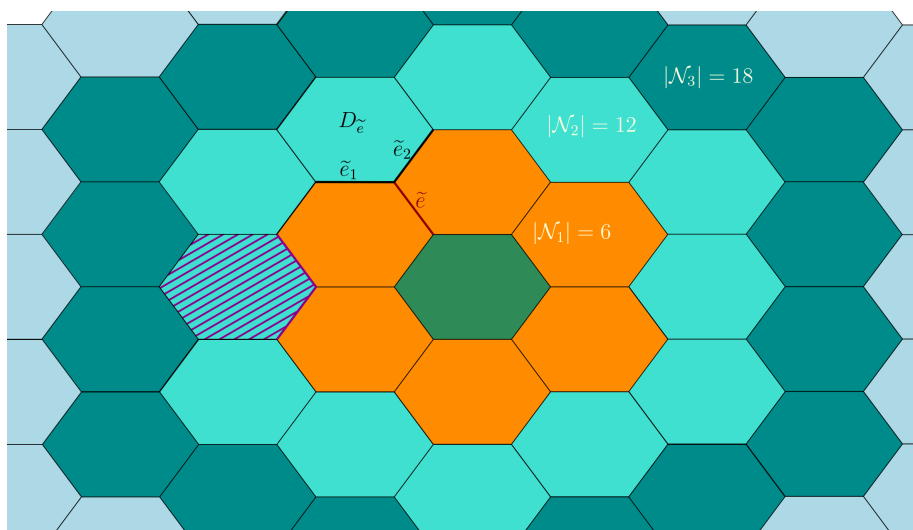


Figure 14: Counting neighborhood layers in the Euclidean plane.

There is a copy $D_{\tilde{e}}$ of $(\mathcal{D}_{\tilde{x}})_o$ in \mathcal{N}_2 whose closure has nonempty intersection with the closed edge \tilde{e} . There are then exactly two edges incident to both $\overline{\mathcal{N}_1}$ and $D_{\tilde{e}}$. Therefore, if we add the number $12g - 6 - 2 - 1$ of edges incident to copies of domains in \mathcal{N}_1 and \mathcal{N}_2 to account for the number of domains in \mathcal{N}_2 , we will have counted $12g - 6$ too many, one for each edge like \tilde{e} in \mathcal{N}_1 , or, in other words, one for each copy of $(\mathcal{D}_{\tilde{x}})_o$ in \mathcal{N}_1 . In total, in the second layer \mathcal{N}_2 , there are then

$$(12g - 6)(12g - 6 - 2 - 1) - (12g - 6) = (12g - 6)(12g - 10)$$

copies of $(\mathcal{D}_{\tilde{x}})_o$. The computation of the number of copies in \mathcal{N}_3 works in largely the same way, but introduces a minor change to the above: Instead of each copy in \mathcal{N}_2 having exactly one edge incident to \mathcal{N}_1 , as was the case with each copy in \mathcal{N}_1 and $(\mathcal{D}_{\tilde{x}})_o$, there are some with two edges incident to \mathcal{N}_1 . Figure 14 shows an example of a such a copy as a striped domain. By the above discussion, there is one of these for every copy of $(\mathcal{D}_{\tilde{x}})_o$ in \mathcal{N}_1 . To find the edges incident to both \mathcal{N}_2 and \mathcal{N}_3 , we therefore have to subtract $|\mathcal{N}_1|$ from $(12g - 6 - 2 - 1)|\mathcal{N}_2|$. All in all, similarly to above, we therefore obtain $|\mathcal{N}_3| = |\mathcal{N}_2|(12g - 6 - 2 - 1) - |\mathcal{N}_1| - |\mathcal{N}_2|$. More generally and by the same reasoning, for $n \geq 2$,

$$|\mathcal{N}_{n+1}| = |\mathcal{N}_n|(12g - 10) - |\mathcal{N}_{n-1}|. \quad (2)$$

For $g = 1$ we readily see that $|\mathcal{N}_n| = 6n$, whence $|\mathcal{N}_{\leq n}| = 1 + 3n(n+1)$. For $g \geq 2$, a crude estimation yields $g^n \leq |\mathcal{N}_n| \leq (12g - 10)|\mathcal{N}_{n-1}| \leq (12g - 10)^n$, so that

$$\frac{g^{n+1} - 1}{g - 1} \leq |\mathcal{N}_{\leq n}| \leq \frac{(12g - 10)^{n+1} - 1}{12g - 11}.$$

For the sake of completeness we give the exact solution of the recursion relation (2) with constant coefficients. The roots of the characteristic polynomial of the recurrence relation (2) are given by

$$r_{1/2} = (6g - 5) \pm \sqrt{(6g - 5)^2 - 1},$$

so the solution is given as ($n \geq 1$)

$$|\mathcal{N}_n| = 6n, \text{ for } g = 1,$$

and

$$|\mathcal{N}_n| = \frac{\sqrt{3}(2g - 1)((2\sqrt{9g^2 - 15g + 6} + 6g - 5)^n) - (-2\sqrt{9g^2 - 15g + 6} + 6g - 5)^n}{2\sqrt{3g^2 - 5g + 2}},$$

for $g > 1$. □

Lemma 7. *In the degenerate case, where all vertices of $\Gamma\mathcal{D}_{\bar{x}}$ are equivalent, the number $|\mathcal{N}_{\leq n}|$ of copies of $\mathcal{D}_{\bar{x}}$ that have a combinatorial distance of at most n from $\mathcal{D}_{\bar{x}}$ satisfies*

$$|\mathcal{N}_{\leq n}| \leq 4g \frac{(4g - 2)^{2n+2} - 1}{4g - 3}.$$

Proof. Observe first that there are $4g(4g - 2)$ copies of $\mathcal{D}_{\bar{x}}$ in the first neighborhood layer, as each of the $4g$ vertices of $\mathcal{D}_{\bar{x}}$ has degree $4g$. For \mathcal{N}_{n+1} , note that a vertex incident to \mathcal{N}_n accounts for either $4g - 3$ copies of $\mathcal{D}_{\bar{x}}$ in \mathcal{N}_{n+1} , or $4g - 2$, depending on whether or not it is incident to one or two domains in \mathcal{N}_n , respectively. A vertex v incident to \mathcal{N}_n is incident to two copies of $\mathcal{D}_{\bar{x}}$ in \mathcal{N}_n if and only if it is incident to an edge in \mathcal{N}_n that is itself incident to two domains in \mathcal{N}_n . Any such edge corresponds to a copy of $\mathcal{D}_{\bar{x}}$ (incident to that edge and in counter-clockwise direction when viewed from v) in \mathcal{N}_n . Therefore, if v_n denotes the set of vertices incident to both $\overline{\mathcal{N}_{n+1}}$ and $\overline{\mathcal{N}_n}$, then

$$|\mathcal{N}_{n+1}| = |v_n|(4g - 2) - |\mathcal{N}_n|.$$

To find $|v_{n+1}|$ using $|v_n|$, recall that any vertex $v \in v_n$ has degree $4g$, with (at least) two edges incident to \mathcal{N}_n . As a result, if v is incident to only one domain in \mathcal{N}_n , the number $(4g - 2)(4g - 2) - 1$ counts the number of vertices, incident to domains incident to v , that belong to $\overline{\mathcal{N}_{n+1}}$ and not to $\overline{\mathcal{N}_n}$. If we count these many vertices for every vertex in v_n , then we will have overcounted the vertices in v_{n+1} , by exactly $|\mathcal{N}_n|(4g - 2)$. Indeed, for every vertex incident to two domains in \mathcal{N}_n , the procedure counts $(4g - 2)$ vertices too many. Moreover, each domain in \mathcal{N}_n can be assigned one unique such vertex, by choosing (say) the counter-clockwise such vertex incident to the domain for $n > 1$. We therefore find

$$|v_{n+1}| = ((4g - 2)(4g - 2) - 1)|v_n| - |\mathcal{N}_n|(4g - 2).$$

We similarly find $|v_0| = 4g$ and $|v_1| = ((4g - 2)(4g - 2) - 1)4g$. Using the above relations and induction, one readily sees that we obtain the crude estimate $|v_n| \leq (4g - 2)^2 |v_{n-1}| \leq (4g - 2)^{2n} 4g$ and subsequently $|\mathcal{N}_n| \leq (4g - 2)^{2n+1} 4g$, whence

$$|\mathcal{N}_{\leq n}| \leq 4g(4g - 2) \frac{(4g - 2)^{2n+2} - 1}{(4g - 2)^2 - 1} \leq 4g \frac{(4g - 2)^{2n+2} - 1}{4g - 3}.$$

□

Proposition 8. *The number $|\mathcal{N}_{\leq n}|$ of copies of $\mathcal{D}_{\bar{x}}$ that have a combinatorial distance of at most n from $\mathcal{D}_{\bar{x}}$ satisfies, for $g > 1$,*

$$\frac{g^{n+1} - 1}{g - 1} \leq |\mathcal{N}_{\leq n}| \leq 4g \frac{(4g - 2)^{2n+2} - 1}{4g - 3}.$$

Proof. The graph of boundary edges of any fundamental domain \mathcal{F} projects to a graph G on the surface S . Combinatorially, any fundamental domain can be obtained from a series of splitting operations of the vertices of this graph, starting from the case where G has only one vertex. A splitting operation corresponds to splitting a vertex v of G into two vertices v_1 and v_2 joined by a new edge e_v to obtain a new graph G' , so that $\deg(v_1) + \deg(v_2) - 2 = \deg(v)$. The preimage $\pi^{-1}(G')$ of the canonical projection gives rise to a fundamental domain \mathcal{F}_{e_v} for Γ . In the associated tessellation $\Gamma\mathcal{F}_{e_v}$, on the boundary of \mathcal{F}_{e_v} , there are $\deg(v_1)$ vertices of degree $\deg(v_1)$ and $\deg(v_2)$ vertices of degree $\deg(v_2)$, in place of the $\deg(v)$ vertices of degree $\deg(v)$, while all other vertices stay invariant. The difference of the number of copies of \mathcal{F}_0 in the first neighborhood layer associated to some choice of original fundamental domain \mathcal{F}_o and the same number for a choice of $(\mathcal{F}_{e_v})_o$ is then

$$D = (\deg(v) - 2) \deg(v) - (\deg(v_1)(\deg(v_1) - 2) + \deg(v_2)(\deg(v_2) - 2)). \quad (3)$$

We show that (3) is always positive. We have, with $\deg(v_1) = a$ and $\deg(v_2) = b$ for better readability,

$$\begin{aligned} D &= (a + b)^2 - 2(a + b) - 4(a + b) + 8 - a^2 + 2a - b^2 + 2b \\ &= 2(ab - 2a - 2b + 4) \geq 2(3(b - 2) - 2b + 4) = 2(b - 2) \geq 2 \end{aligned}$$

for $a = \deg(v_1) \geq 3 \leq b$. In particular, this shows that the splitting operation decreases the number of copies of domains in the first neighborhood layer \mathcal{N}_1 . To see that this trend continues in subsequent neighborhood layers, one can argue similarly to the proofs of Lemmas 6 and 7, but we sketch a somewhat different argument in the following. To this end, we imagine the original domain as having only long edges and picture the splitting operation as inserting a very small edge e_v into G . Consider looking at the tiling $\Gamma\mathcal{F}_{e_v}$ from far away and coloring domains within a given combinatorial distance from the central domain \mathcal{F} . Because e_v is short, $\Gamma\mathcal{F}_{e_v}$ looks like the tiling $\Gamma\mathcal{F}$, with domains incident to a lift of v , situated at the interface of a current neighborhood layer and the next, corresponding to domains incident to an appropriate lift of e_v after splitting. We observe that the insertion of e_v can only decrease the number of these domains that belong to the next neighborhood layer. As a result, we can map the colored domains in the split tiling injectively to those of the unsplit tiling, where we retain only the domains of $\Gamma\mathcal{F}$ that correspond to those which were colored in the split tiling.

All in all, we see that splitting vertices to obtain new fundamental domains decreases the number of domains in the neighborhood layers, so that the size of the neighborhood layers for an arbitrary fundamental domain lies between the two extremal cases treated in the previous two lemmas 6 and 7, from which the proposition follows. □

Proof of Proposition 7

Proof. The proof follows from Proposition 8 by setting $n = 12g - 6$, the maximal value for combinatorial length of a Delaunay edge, by Corollary 3. □

References

- [1] A.F. Beardon. *The Geometry of Discrete Groups*. 3Island Press, 1983.
- [2] Marcel Berger. *Geometry 2*. Springer, 2009.
- [3] 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.
- [4] 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 für Informatik. doi:10.4230/LIPIcs.SoCG.2016.20.
- [5] Peter Buser. *Geometry and Spectra of Compact Riemann Surfaces*. Birkhäuser, Boston, Mass, 2nd ed. edition, 2010. doi:10.1007/978-0-8176-4992-0.
- [6] Manuel Caroli, Aymeric Pellé, Mael Rouxel-Labbé, and Monique Teillaud. 3D periodic triangulations. In *CGAL User and Reference Manual*. CGAL Editorial Board, 4.11 edition, 2017. URL: <http://doc.cgal.org/latest/Manual/packages.html#PkgPeriodic3Triangulation3>.
- [7] Manuel Caroli and Monique Teillaud. 3D periodic triangulations. In *CGAL User and Reference Manual*. CGAL Editorial Board, 3.5 edition, 2009. URL: <http://doc.cgal.org/latest/Manual/packages.html#PkgPeriodic3Triangulation3>.
- [8] Manuel Caroli and Monique Teillaud. Computing 3D periodic triangulations. In *Proceedings 17th European Symposium on Algorithms*, volume 5757 of *Lecture Notes in Computer Science*, pages 59–70, 2009. URL: <https://hal.inria.fr/hal-02954152>, doi:10.1007/978-3-642-04128-0_6.
- [9] 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.
- [10] Jason DeBlois. The centered dual and the maximal injectivity radius of hyperbolic surfaces. *Geometry & Topology*, 19(2):953–1014, 2015. doi:10.2140/gt.2015.19.953.
- [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 *Chaos carolinensis*. *Protoplasma*, 203:16–25, 1998. doi:10.1007/BF01280583.
- [13] Vincent Despré, Benedikt Kolbe, Hugo Parlier, and Monique Teillaud. Computing a Dirichlet Domain for a Hyperbolic Surface. In Erin W. Chambers and Joachim Gudmundsson, editors, *39th International Symposium on Computational Geometry (SoCG 2023)*, volume 258 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 27:1–27:15, Dagstuhl, Germany, 2023. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. URL: <https://drops.dagstuhl.de/opus/volltexte/2023/17877>, doi:10.4230/LIPIcs.SoCG.2023.27.
- [14] 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.

- [15] 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.
- [16] Matthijs Ebbens, Hugo Parlier, and Gert Vegter. Minimal Delaunay Triangulations of Hyperbolic Surfaces. In Kevin Buchin and Éric Colin de Verdière, editors, *37th International Symposium on Computational Geometry (SoCG 2021)*, volume 189 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 31:1–31:16, Dagstuhl, Germany, 2021. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.SoCG.2021.31.
- [17] 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.
- [18] 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.
- [19] 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>.
- [20] Sylvestre Gallot, Dominique Hulin, and Jacques Lafontaine. *Riemannian geometry*. Springer, 3 edition, 2004.
- [21] 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.
- [22] Allen Hatcher. *Algebraic Topology*, volume 44. Cambridge University Press, 2002.
- [23] 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.
- [24] Stephen Hyde, Stuart Ramsden, and Vanessa Robins. Epinet. Accessed: 2020-10-23. URL: <http://epinet.anu.edu.au>.
- [25] Stephen T. Hyde, T. Landh, S. Lidin, B.W. Ninham, K. Larsson, and S. Andersson. *The Language of Shape*. Elsevier Science, 1996.
- [26] Jordan Iordanov 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.
- [27] 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 (PNAS)*, 111(4):1271–6, 2014. doi:10.1073/pnas.1316348111.
- [28] Benedikt Kolbe. Towards a combinatorial algorithm for the enumeration of isotopy classes of symmetric cellular embeddings of graphs on hyperbolic surfaces. Under review, December 2020. URL: <https://inria.hal.science/hal-03047231>.
- [29] Benedikt Kolbe and Myfanwy E. Evans. Isotopic tiling theory for hyperbolic surfaces. *Geometriae Dedicata*, 2020. doi:10.1007/s10711-020-00554-2.

- [30] Benedikt Kolbe and Myfanwy E. Evans. Enumerating Isotopy Classes of Tilings Guided by the Symmetry of Triply Periodic Minimal Surfaces. *SIAM Journal on Applied Algebra and Geometry*, 6(1):1–40, mar 2022. doi:10.1137/20M1358943.
- [31] Benedikt Kolbe and Vanessa Robins. Tile-transitive tilings of the Euclidean and hyperbolic planes by ribbons. *Research in Computational Topology 2*, 2022. URL: <https://hal.inria.fr/hal-03046798v1>.
- [32] Nico Kruithof. 2D periodic triangulations. In *CGAL User and Reference Manual*. CGAL Editorial Board, 4.4 edition, 2014. URL: <http://doc.cgal.org/latest/Manual/packages.html#PkgPeriodic2Triangulation2Summary>.
- [33] C.L. Lawson. *Software for C^1 Surface Interpolation*, pages 161–194. ACADEMIC PRESS, INC., 1977. doi:10.1016/b978-0-12-587260-7.50011-x.
- [34] 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.
- [35] A. M. Macbeath. Generic Dirichlet polygons and the modular group. *Glasgow Mathematical Journal*, 27:129–141, 1985. doi:10.1017/S0017089500006133.
- [36] William S. Massey. *A Basic Course in Algebraic Topology*. Graduate Texts in Mathematics. Springer, 1991.
- [37] William H. Meeks. The theory of triply periodic minimal surfaces, 1990. doi:10.1512/iumj.1990.39.39043.
- [38] Barrett O’Neill. *Semi-Riemannian geometry : with applications to relativity*. Academic Press, 1983.
- [39] 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.
- [40] J. Ratcliffe. *Foundations of Hyperbolic Manifolds*. Graduate Texts in Mathematics. Springer New York, 2006.
- [41] 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.
- [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] Adam G. Weyhaupt. Deformations of the gyroid and lidinoid minimal surfaces. *Pacific Journal of Mathematics*, 235(1):137–171, mar 2008. doi:10.2140/pjm.2008.235.137.