



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. 2021. hal-03045921v2

HAL Id: hal-03045921

<https://hal.science/hal-03045921v2>

Preprint submitted on 13 Jul 2021 (v2), 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.

Representing infinite hyperbolic periodic Delaunay triangulations using finitely many Dirichlet domains

Vincent Despré^{*†}

Benedikt Kolbe^{*‡}

Monique Teillaud ^{*§}

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. Assuming that the surface comes with a Dirichlet domain, we exhibit a finite set of points that captures the full triangulation.

Indeed, 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. On the way, we prove that both the edges of a Delaunay triangulation and of a Dirichlet domain have some kind of distance minimizing properties that are of intrinsic interest.

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.

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

[†]vincent.despre@loria.fr

[‡]benedikt.kolbe@inria.fr

[§]monique.teillaud@inria.fr

This work was partially supported by the grant(s) ANR-17-CE40-0033 of the French National Research Agency ANR (project SoS) and INTER/ANR/16/11554412/SoS of the Luxembourg National Research fund FNR.
<https://members.loria.fr/Monique.Teillaud/collab/SoS/>.

July 13, 2021

1 Introduction

The classic edge flip algorithm of Lawson [Law77] computing Delaunay triangulations in the Euclidean plane was recently extended to hyperbolic surfaces [DST20]. However, robust and efficient software to compute Delaunay triangulations on hyperbolic surfaces, and particularly triply-periodic minimal surfaces presented below, does not exist to date, as far as we know. A primary motivation for the work in this paper is to attempt 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 [CT09a, Kru14, CPRT17, ORT20]. Moreover, our results yield structural insights of independent interest into the relationship between different representations of hyperbolic surfaces and Delaunay triangulation on them.

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 [UHS10]. 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 [Mee90]. The gyroid, the primitive, and the diamond TPMS are arguably the most prominent and simple examples of TPMS [STFH06] and have received considerable attention in the mathematical, physical, chemical and biological as well as interdisciplinary literature [Wey08, KEdCH14, HRDL03, HLL⁺96, DM98]. More recently, TPMS have also found a role in the materials sciences as a scaffold for crystallographic structures [ERH13a, ERH13b], leading to both new mathematical formalisms [KE20b, KE20a] and a database of such structures [epi]. These three TPMS are closely related to each other [UHS10, Section 3.1][GW96] and give rise to the same 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.



(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 .

The flat case. The computation of Delaunay triangulations in flat tori, which can be seen equivalently as periodic triangulations in the Euclidean space, was addressed by Dolbilin and Huson [DH97], who provided a first cornerstone for the algorithms and CGAL packages handling the

square/cubic periodicity [CT09b, CT09a, Kru14]; the algorithms were generalized later [CT16]. 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) [ORT20]. The idea is as follows. 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 $\text{DT}_{\tilde{\mathcal{P}}}$ can be deduced from the Delaunay triangulation $\text{DT}_{\tilde{\mathcal{P}}_f}$; as $\tilde{\mathcal{P}}_f$ is finite, $\text{DT}_{\tilde{\mathcal{P}}_f}$ 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. It is well-known that the size of $\tilde{\mathcal{P}}_f$ is unbounded for an arbitrary choice of parallelogram/translations. 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 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).




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.

However, if choosing as fundamental domain the Dirichlet domain $\mathcal{D}_{\tilde{x}}$ of a 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 domains an edge of a Delaunay triangulation can traverse [DH97, ORT20]. 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. In our setting, a hyperbolic surface is a *closed* and *orientable* topological surface equipped with some hyperbolic metric of constant curvature -1 . It is well-known that 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




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




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




Figure 5: Dirichlet domains of different points for the same group, which defines the genus 2 Bolza surface [BTV16, Figure 9].

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

One can think of the projection map π as giving rise to 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. It is convenient to 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 [Mas91].

We use the Poincaré disk model for the hyperbolic plane \mathbb{H}^2 [Ber09]. 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 [BDT14].

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 [DeB15, DST20]. We do not assume triangulations to be simplicial complexes in this paper, in contrast to some previous work [BTV16, IT17]. In our setting, every finite point cloud on a hyperbolic surface has an associated locally finite Delaunay triangulation [DeB15, Corollary 5.2][DST20, Proposition 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, however, 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: $\mathcal{D}_{f(\tilde{x})}^\Gamma = f(\mathcal{D}_{\tilde{x}}^{f^{-1}\Gamma f})$ [Bea83, Section 9.4]. The methods to treat the Euclidean case depend crucially on the fact that the involved groups are Abelian, so we need new tools to tackle the problem.

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 a 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

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, which implies that if K is compact in X , then $P \cap K$ is finite.

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 $\text{DT}_{\tilde{\mathcal{P}}}$ 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 22 in Section 5, is an explicit bound when a Dirichlet domain is chosen as a fundamental domain. The bound only depends on the genus of S .

The proof of the theorem relies on intersection properties of edges of Delaunay triangulations and Dirichlet domains, studied in Section 4.

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 [Rat06, Thu80].

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 $\tilde{x} \in \mathcal{P}$, we denote the closed Voronoi cell of \tilde{x} by $\mathcal{V}_{\tilde{x}}^{\mathcal{P}}$ 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 [DeB18, Lemma 5.2].

Definition 1 ([Bea83]). 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

$$\tilde{z} \in \mathcal{D}_{\tilde{x}} \iff \tilde{x} \in \mathcal{D}_{\tilde{z}}. \tag{1}$$

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 only 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 [BDT14]. Note that the circumcircles of faces in $\text{DT}_{\tilde{\mathcal{P}}}$ are all compact [DeB18, DST20]. Though we use the term *triangulation*, we also consider the case where more than three cocircular 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 property:

Proposition 2. ([DST20, Proposition 8],[DeB15, Corollary 5.2]) *The 1-skeleton of the Delaunay triangulation $DT_{\mathcal{P}}$ 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 3. 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 3 have been considered in the literature [CT16, IT17].

3 Combinatorial length

We observe that the lattice of \mathbb{H}^2 formed by a polygonal fundamental domain and its copies, i.e. its images under Γ , can be decomposed into layers:

Definition 4. 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 .

Assume that the tessellation induced by \mathcal{F} and Γ features only vertices of degree 3, which is the generic case [Bea83, Theorem 9.4.5]. Then the combinatorial distance between points is equal to



(a) Tessellation of \mathbb{H}^2 formed by copies of the Dirichlet domain of the origin for Γ_B [BTV16, Figure 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.

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 to be solved is how many fundamental domains an edge can intersect. Unfortunately, as in the flat case above, there is no bound on this number of intersections, as we now show.

Proposition 5. *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 non-separating path $c_1 \subset S$ that intersects c_C exactly once, somewhere in the interior of e . By applying a homeomorphism f of S , we map $c_C \cup c_1$ to the paths on S shown in Figure 7a [FM12, 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 open S




Figure 7: The graph G' in the proof of Proposition 5.

along the graph G' produces a disk, as G' yields the standard presentation of the fundamental group of a surface [Hat02, 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

$\pi^{-1}(S - \hat{G})$ [LM91, 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_C . 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 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 c_1 . The number of intersections of $t^M(c_1)$ with c_C , where $t^M = t \circ t \circ \dots \circ t$ (M -times), is equal to M [FM12, Proposition 3.2]. Therefore, there is a representative in the isotopy class of $t^M(c_1)$ that intersects c_C M times, without forming any bigons which by [FM12, Proposition 1.7] and [FM12, 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 $\widetilde{t^M(c_1)}$. Therefore, the combinatorial length of the edge \tilde{c} can be made arbitrarily large by successive application of the Dehn twist t . \square

Remark 6. The statement of Proposition 5 remains valid if one drops the assumption that the simple path is nonseparating, as long as the path is non-trivial. 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 7. *There is no bound on the combinatorial length of an edge in $\text{DT}_{\tilde{\mathcal{P}}}$ for arbitrary fundamental domains.*

Proof. 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 $\text{DT}_{\mathcal{P}}$ is a triangulation, the closed paths in $\text{DT}_{\mathcal{P}}$ generate the fundamental group Γ . By Hurewicz' theorem [Hat02, 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 2, \tilde{e} is simple, and thus Proposition 5 concludes the proof. \square

This motivates the use of Dirichlet domains for our purposes. The following lemma generalizes a result known in the flat case [DH97]. 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 [DH97, Lemma 3.2].

Lemma 8. *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 $\tilde{\Delta}$. 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

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. This motivates the following definitions.

Definition 9. 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 path \tilde{c} in \mathbb{H}^2 a *distance path* if it is the lift of a distance path c in S .

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

Definition 10. 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.

Let us study intersection properties of distance paths or half-minimizers.

Lemma 11. *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 [O'N83, 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 [O'N83, 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 12. *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 11 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 11 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

We now show that the edges of Delaunay triangulations and Dirichlet domains are either half-minimizers or isotopic to a half-minimizer.


Lemma 13. *Let \tilde{x} be a point in \mathbb{H}^2 . For any point \tilde{y} in the Dirichlet domain $\mathcal{D}_{\tilde{x}}$, 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 [Bus10, 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 [GHL04, 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 $\tilde{x}' \in \Gamma\tilde{x}$. 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

Remark 14. A Delaunay edge \tilde{e} is not always a half-minimizer (Figure 8a). 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 13. 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 15 moves the edge \tilde{e} over the shaded bigon B to the (blue) path \tilde{c} .

Figure 8: Illustrations for Remark 14 (left) and Proposition 15 (right).

Proposition 15. *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 the covering space \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 8 and because $\Gamma\tilde{x} \subset \tilde{\mathcal{P}}$ implies that $\mathcal{V}_x^{\tilde{\mathcal{P}}} \subset \mathcal{D}_{\tilde{x}}$. Similarly, $\omega_{\tilde{\Delta}}$ lies in $\mathcal{D}_{\tilde{y}}$. So, by Lemma 13, \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 5) and Lemma 8. 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 2 for $\text{DT}_{\mathcal{P}}$ and the fact that c is isotopic to the projection of \tilde{e} . \square

We now consider the edges of a Dirichlet domain.

Definition 16. 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 [DeB15], restricted to Voronoi edges.

Proposition 17. *Each centered edge of $\mathcal{D}_{\tilde{x}}$ is a half-minimizer and each non-centered edge is a distance path.*

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, we can assume that the points \tilde{x} and \tilde{x}_1 lie on the imaginary axis.

For any point $\tilde{z} \in \tilde{e}$, the geodesic segments $\tilde{\gamma}_z \subset \mathcal{D}_{\tilde{x}}$ and $(\tilde{\gamma}_z)_1 \subset \mathcal{D}_{\tilde{x}_1}$, from \tilde{x} and \tilde{x}_1 to \tilde{z} respectively, together form a half-minimizer, by Lemma 13.

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 13. 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} , does not contain any point of $\Gamma\tilde{b}$, so, \tilde{e} is a distance path.




Figure 9: Proof of Proposition 17: (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 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 18. 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 19. 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 $\tilde{\gamma}_{\tilde{z}_1\tilde{z}_2}$.

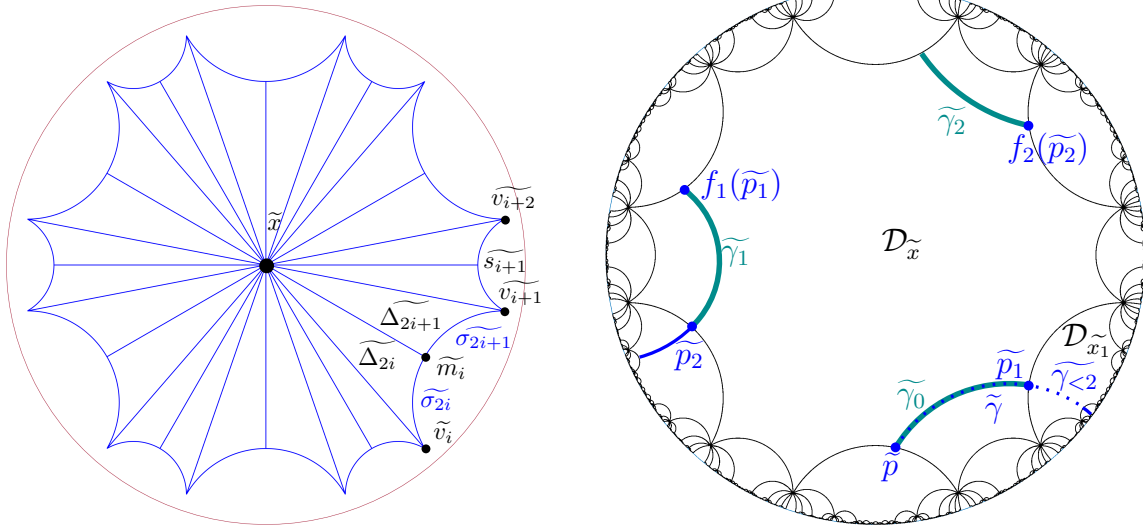
We first prove that the (unique) geodesics in $\tilde{\Delta}$ from corners to arbitrary points on its edges are distance paths. Denote the vertices of $\tilde{\Delta}$ by \tilde{x}, \tilde{y} , and \tilde{z} . Consider w.l.o.g. the corner \tilde{y} . The edges $\tilde{\gamma}_{\tilde{y}\tilde{z}}$ and $\tilde{\gamma}_{\tilde{y}\tilde{x}}$ of $\tilde{\Delta}$ are distance paths so $\tilde{z}, \tilde{x} \in \mathcal{D}_{\tilde{y}}$ by Lemma 13. Since $\mathcal{D}_{\tilde{y}}$ is convex, $\tilde{\gamma}_{\tilde{x}\tilde{z}} \subset \mathcal{D}_{\tilde{y}}$ and therefore, again by Lemma 13, 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 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 $\tilde{p}_2 \in \tilde{\gamma}_{\tilde{y}\tilde{z}}$, and consider the geodesic $\tilde{\gamma}_{\tilde{p}_1\tilde{p}_2} \subset \tilde{\Delta}$. 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 10a in the case of a dodecahedron. 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 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 17 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.



(a) The triangles partitioning $\mathcal{D}_{\tilde{x}}$ (illustration for the special case of TMPS).

(b) Representation of a geodesic path γ as a sequence of segments in $\mathcal{D}_{\tilde{x}}$. The path $\widetilde{\gamma}_{<2}$ is also shown as a dashed part of γ .

Figure 10: 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.

Denote now as $\widetilde{\gamma}_z$ the geodesic segment between \tilde{x} and \tilde{z} for any point $z \in \partial\mathcal{D}_{\tilde{x}}$; by Lemma 13, $\widetilde{\gamma}_{v_i}$ and $\widetilde{\gamma}_{m_i}$ are distance paths, for $i = 0, \dots, k-1$. Let us denote the closed triangle formed by \tilde{x} and $\widetilde{\sigma}_{2i}$ as $\widetilde{\Delta}_{2i}$ and the triangle formed by \tilde{x} and $\widetilde{\sigma}_{2i+1}$ as $\widetilde{\Delta}_{2i+1}$; the triangles $\widetilde{\Delta}_j$, $j = 0, \dots, 2k-1$ partition $\mathcal{D}_{\tilde{x}}$. By Proposition 19, the projection Δ_j of each $\widetilde{\Delta}_j$ on S is convex.

Lemma 20. *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 bounded by $k/2$.*

Before we prove the result, let us observe that a geodesic path $\widetilde{\gamma}$ 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 10b. Let $\tilde{p} \in (\mathcal{D}_{\tilde{x}})_o$ be an endpoint of $\widetilde{\gamma}$. If $\widetilde{\gamma} \subset (\mathcal{D}_{\tilde{x}})_o$ then the sequence is reduced to $\{\widetilde{\gamma}\}$. Otherwise, denote as $\widetilde{\gamma}_0$ the intersection $\gamma \cap (\mathcal{D}_{\tilde{x}})_o$. The path $\widetilde{\gamma}$ exits $\mathcal{D}_{\tilde{x}}$ at the intersection point \tilde{p}_1 : $\{\tilde{p}_1\} = \widetilde{\gamma} \cap \partial\mathcal{D}_{\tilde{x}}$. 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 $\widetilde{\gamma} \cap (\mathcal{D}_{\tilde{x}_1})_o$ back into $\mathcal{D}_{\tilde{x}}$ and obtain another geodesic segment $\widetilde{\gamma}_1$, which may again exit $\mathcal{D}_{\tilde{x}}$. Repeating this process until we reach the other endpoint of γ yields a collection $\{\widetilde{\gamma}_n\}_{n=0}^{r-1}$ 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 $\widetilde{\gamma}$ 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 $\widetilde{\gamma}_{<l}$ the lift of $\gamma_{<l}$ to \mathbb{H}^2 , starting at \tilde{p} . Observe that $\widetilde{\gamma}_{<r} = \widetilde{\gamma}$.

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

Observe that since $\widetilde{\gamma}$ 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 19). This distance path would intersect the distance path γ more than once, which is impossible by Lemma 11. 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 $\widetilde{\gamma}_n$ intersects a triangle $\widetilde{\Delta}_j$ for some n , then $\widetilde{\gamma}_l$ cannot intersect the same $\widetilde{\Delta}_j$ for $l \neq n$.

Recall the definition of neighborhood layers (Definition 4) and Figure 6b: if after leaving the (red) central Dirichlet domain, a path intersects consecutive sides of Dirichlet domains sharing the same corner of this red domain, then it stays in the blue layer. More generally, if the two endpoints of a segment $\widetilde{\gamma}_l$ lie on two adjacent sides of $\mathcal{D}_{\widetilde{x}}$, then the combinatorial length $d^c(\widetilde{\gamma}_{<l+1})$ is the same as the combinatorial length $d^c(\widetilde{\gamma}_{<l})$. 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 $\widetilde{\gamma}_n, n = 1, \dots, r-2$ lie on non-consecutive sides of $\mathcal{D}_{\widetilde{x}}$, so that the combinatorial length of $\widetilde{\gamma}$ is $r-1$.

Remark now that such a segment $\widetilde{\gamma}_n, n \in \{1, \dots, r-2\}$ 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 $\widetilde{\gamma}_n$ passes through a corner of a Dirichlet domain, the number of triangle orbits intersected by $\widetilde{\gamma}_n$ 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 $\widetilde{\gamma}_n, n \in \{0, \dots, r-2\}$. Moreover, $\widetilde{\gamma}_0$ 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 a bound, depending only on the genus, for the maximal combinatorial length of a Delaunay triangulation in the restricted case where all its edges are distance paths on S . Note that this case corresponds to the framework of previous work [IT17, BTV16, ORT20]. 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 $\text{DT}_{\widetilde{\mathcal{P}}}$ are distance paths.

Corollary 21. *Let $S = \mathbb{H}^2/\Gamma$ be a genus g hyperbolic surface and $\mathcal{D}_{\widetilde{x}}$ 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 $\text{DT}_{\widetilde{\mathcal{P}}}$ in \mathbb{H}^2 is bounded by $6g - 3$.*

Proof. The number of edges of a Dirichlet fundamental domain is at most $12g - 6$, by a direct use of the Euler characteristic [Bea83, Theorem 10.5.1]. The result follows from Lemma 20. \square

Note that the upper bound can only be attained in the generic situation where all vertices of $\Gamma\mathcal{D}_{\widetilde{x}}$ have degree 3.

Theorem 22. *Let $\mathcal{D}_{\widetilde{x}}$ be a Dirichlet domain, with k edges. Let \mathcal{P} be a finite set of points on S and $\text{DT}_{\widetilde{\mathcal{P}}}$ the Delaunay triangulation of the lifted set of points $\widetilde{\mathcal{P}} \subset \mathbb{H}^2$. The maximal combinatorial length of a Delaunay edge is bounded by k .*

Proof. Consider a path \widetilde{c} joining $\widetilde{a} \in \mathcal{D}_{\widetilde{x}}$ and $\widetilde{b} \in \mathbb{H}^2$. The geodesic from \widetilde{a} and \widetilde{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) [Bus10, Theorem 1.5.3]. So, the homotopy class of c is uniquely determined by which copy of the original domain $(\mathcal{D}_{\widetilde{x}})_o$ contains \widetilde{b} .

As a consequence, if the combinatorial length of \widetilde{c} 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 \widetilde{c} in \mathbb{H}^2 is isotopic to a half-minimizer by Proposition 15. The result now follows directly from Lemma 20. \square

Corollary 23. *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 bounded by $12g - 6$.*

Again, the upper bound can only be attained in the generic situation.

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.

References

- [BDT14] 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.
- [Bea83] A.F. Beardon. *The Geometry of Discrete Groups*. 3Island Press, 1983.
- [Ber09] Marcel Berger. *Geometry 2*. Springer, 2009.
- [BTV16] 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.
- [Bus10] Peter Buser. *Geometry and Spectra of Compact Riemann Surfaces*. Birkhäuser, Boston, Mass, 2nd ed. edition, 2010. doi:10.1007/978-0-8176-4992-0.
- [CPRT17] 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>.
- [CT09a] 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>.
- [CT09b] 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.
- [CT16] 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.
- [DeB15] 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.
- [DeB18] Jason DeBlois. The Delaunay tessellation in hyperbolic space. *Mathematical Proceedings of the Cambridge Philosophical Society*, 164(1):15–46, 2018. doi:10.1017/S0305004116000827.
- [DH97] 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.
- [DM98] 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.

- [DST20] 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.
- [epi] Epinet. Accessed: 2020-10-23. URL: <http://epinet.anu.edu.au>.
- [ERH13a] 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.
- [ERH13b] 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.
- [FM12] 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>.
- [GHL04] Sylvestre Gallot, Dominique Hulin, and Jacques Lafontaine. *Riemannian geometry*. Springer, 3 edition, 2004.
- [GW96] 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.
- [Hat02] Allen Hatcher. *Algebraic Topology*, volume 44. Cambridge University Press, 2002.
- [HLL⁺96] Stephen T. Hyde, T. Landh, S. Lidin, B.W. Ninham, K. Larsson, and S. Andersson. *The Language of Shape*. Elsevier Science, 1996.
- [HRDL03] 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.
- [IT17] 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.
- [KE20a] Benedikt Kolbe and Myfanwy Evans. Enumerating isotopy classes of tilings guided by the symmetry of triply-periodic minimal surfaces, 2020. Forthcoming in *SIAM Journal on Applied Algebra and Geometry*.
- [KE20b] Benedikt Kolbe and Myfanwy Evans. Isotopic tiling theory for hyperbolic surfaces. *Geometriae Dedicata*, 2020. doi:10.1007/s10711-020-00554-2.
- [KEdCH14] 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.

- [Kru14] 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>.
- [Law77] Charles L. Lawson. Software for C^1 surface interpolation. In *Symposium on Mathematical Software*, 1977. NASA Technical Report JPL-PUBL-77-30. URL: <https://ntrs.nasa.gov/citations/19770025881>.
- [LM91] 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.
- [Mas91] William S. Massey. *A Basic Course in Algebraic Topology*. Graduate Texts in Mathematics. Springer, 1991.
- [Mee90] William H. Meeks. The theory of triply periodic minimal surfaces, 1990. doi:10.1512/iumj.1990.39.39043.
- [O’N83] Barrett O’Neill. *Semi-Riemannian geometry : with applications to relativity*. Academic Press, 1983.
- [ORT20] 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.
- [Rat06] J. Ratcliffe. *Foundations of Hyperbolic Manifolds*. Graduate Texts in Mathematics. Springer New York, 2006.
- [STFH06] 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.
- [Thu80] William Thurston. *Geometry and Topology of Three-Manifolds*. Princeton lecture notes, 1980.
- [UHS10] Dierkes Ulrich, Stefan Hildebrandt, and Friedrich Sauvigny. *Minimal Surfaces*. Springer-Verlag Berlin Heidelberg, 2 edition, 2010.
- [Wey08] Adam G. Weyhaupt. Deformations of the gyroid and lidinoid minimal surfaces. *Pacific J. Math.*, 235(1):137–171, 2008.