



HAL
open science

Flipping geometric triangulations on hyperbolic surfaces

Vincent Despré, Jean-Marc Schlenker, Monique Teillaud

► **To cite this version:**

Vincent Despré, Jean-Marc Schlenker, Monique Teillaud. Flipping geometric triangulations on hyperbolic surfaces. *Journal of Computational Geometry*, 2024, 10.20382/jocg.v15i1a8 . hal-04857057

HAL Id: hal-04857057

<https://hal.science/hal-04857057v1>

Submitted on 27 Dec 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.



Distributed under a Creative Commons Attribution 4.0 International License

FLIPPING GEOMETRIC TRIANGULATIONS ON HYPERBOLIC SURFACES*

Vincent Despré,[†] Jean-Marc Schlenker,[‡] and Monique Teillaud [§]

December 27, 2024

Abstract

We consider geometric triangulations of surfaces, i.e., triangulations whose edges can be realized by disjoint geodesic segments. We prove that the flip graph of geometric triangulations with fixed vertices of a flat torus or a closed hyperbolic surface is connected. We prove that any Delaunay triangulation is geometric, and give upper bounds on the number of edge flips that are necessary to transform any geometric triangulation on such a surface into a Delaunay triangulation.

Keywords: Hyperbolic surface, Topology, Delaunay triangulation, Algorithm, Flip graph

1 Introduction

We investigate triangulations of two categories of surfaces: flat tori, i.e., orientable surfaces of genus 1 with a locally Euclidean metric, and hyperbolic surfaces, i.e., surfaces of genus at least 2 with a locally hyperbolic metric (these surfaces will be introduced more formally in Section 2.1). We do not consider the sphere, which can be equipped with a spherical metric but for which there is no uniqueness of geodesic segments between two points in a homotopy class.

Triangulations of surfaces can be considered in a purely topological manner: a triangulation of a surface is an isotopy class of graphs whose vertices, edges and faces partition the surface and whose faces have three (non-necessarily distinct) vertices. However, when the surface is equipped with a Euclidean or hyperbolic structure, it is possible to consider *geometric* triangulations, i.e., triangulations whose edges can be realized as geodesic segments that can only intersect at common endpoints (Definition 2.1). Note that a geometric triangulation can

*The authors were partially supported by the ANR/FNR SoS grant, ANR-17-CE40-0033 of the French National Research Agency ANR and INTER/ANR/16/11554412/SoS of the Luxembourg National Research fund FNR <https://SoS.loria.fr/> and by the ANR grant, ANR-23-CE48-0017 of the French National Research Agency ANR <https://homepages.loria.fr/VDespre/Abysm/>.

[†]Université de Lorraine, CNRS, LORIA, F-54000 Nancy, France, vincent.despre@loria.fr, <https://members.loria.fr/VDespre/>

[‡]Department of Mathematics, University of Luxembourg, Luxembourg jean-marc.schlenker@uni.lu, <http://math.uni.lu/schlenker/> Jean-Marc Schlenker was partially supported by FNR grant CoSH, O20/14766753

[§]Université de Lorraine, CNRS, Inria, LORIA, F-54000 Nancy, France monique.teillaud@inria.fr <https://members.loria.fr/Monique.Teillaud/>

still have loops and multiple edges, but no contractible loop and no contractible cycle formed of two edges. We will prove that any Delaunay triangulation (Definition 2.2) of the considered surfaces is geometric (Proposition 3.3).

The flip graph of triangulations of the Euclidean plane is known to be connected; moreover the number of edge flips that are needed to transform any given triangulation with n vertices in the plane into the Delaunay triangulation behaves as $\Theta(n^2)$ [18]. We are interested in generalizations of this result to surfaces. Flips in triangulations of surfaces will be defined precisely later (Definition 2.4), for now we can just think of them as similar to edge flips in triangulations of the Euclidean plane. Geodesics only locally minimize the length, so the edges of a geometric triangulation are generally not shortest paths. We will prove that the number of geometric triangulations on a set of points can be infinite, whereas the flip graph of “shortest path” triangulations is small but not connected in most situations [8].

Definition 1.1. Let (\mathcal{M}_2, h) be either a torus (\mathbb{T}^2, h) equipped with a Euclidean structure h or a closed oriented surface (S, h) equipped with a hyperbolic structure h . Let $V \subset \mathcal{M}_2$ be a set of n points. The *geometric flip graph* $\mathcal{F}_{\mathcal{M}_2, h, V}$ of (\mathcal{M}_2, h, V) is the graph whose vertices are the geometric triangulations of (\mathcal{M}_2, h) with vertex set V and where two vertices are connected by an edge if and only if the corresponding triangulations are related by a flip.

Our results are mainly interesting in the hyperbolic setting, which is richer than the flat setting. However, to help the readers’ intuition, we also present them for flat tori, where they are slightly simpler to prove and might even be considered as folklore.

The main results of this paper are:

- The geometric flip graph of (\mathcal{M}_2, h, V) is connected (Theorems 4.4 and 4.6).
- The Delaunay triangulation can be reached from any geometric triangulation by a path in the geometric flip graph $\mathcal{F}_{\mathcal{M}_2, h, V}$ whose length is bounded by n^2 times a quantity measuring the *quality* of the input triangulation (Theorem 5.4).

The connectivity of the geometric flip graph was actually already known in the case of flat surfaces [25], however no algorithmic result was given.

An immediate consequence of our results is the extension of the algorithm based on flips [20], which is standard in the Euclidean plane, to the case of a flat or hyperbolic surface, when an initial triangulation only having one vertex is given: For each new point, the triangle containing it is split into three triangles, then the Delaunay property is restored by propagating flips. This approach can handle triangulations of a surface with loops and multiarcs, which is not the case for the approach based on Bowyer’s incremental algorithm [6, 5].

Our results also provide a tool to tackle other questions. For instance, they were recently used as a step to compute a Dirichlet domain for a hyperbolic surface [10].

2 Background and notation

2.1 Surfaces

In this section, we first recall a few notions, then we illustrate them for the two classes of surfaces (flat tori and hyperbolic surfaces) that we are interested in.

Let \mathcal{M}_2 be a closed oriented surface, i.e., a compact connected oriented 2-manifold without boundary. There is a unique simply connected surface $\widetilde{\mathcal{M}}_2$, called the *universal cover* of \mathcal{M}_2 , equipped with a projection $\rho : \widetilde{\mathcal{M}}_2 \rightarrow \mathcal{M}_2$ that is a local diffeomorphism. There is a natural action on $\widetilde{\mathcal{M}}_2$ of the fundamental group $\pi_1(\mathcal{M}_2)$ of \mathcal{M}_2 so that for all $p \in \mathcal{M}_2$, $\rho^{-1}(p)$ is an orbit under the action of $\pi_1(\mathcal{M}_2)$. We will denote as \tilde{p} a *lift* of p , i.e., one of the elements of the orbit $\rho^{-1}(p)$. A *fundamental domain* in $\widetilde{\mathcal{M}}_2$ for the action of $\pi_1(\mathcal{M}_2)$ on $\widetilde{\mathcal{M}}_2$ is a connected subset Ω of $\widetilde{\mathcal{M}}_2$ that intersects each orbit in exactly one point, or, equivalently, such that the restriction of ρ to Ω is a bijection from Ω to \mathcal{M}_2 [22]. The genus g of \mathcal{M}_2 is its number of handles. In this paper, we consider surfaces with constant curvature (0 or -1). The value of the curvature is given by the Gauss-Bonnet Theorem and thus only depends on the genus: an orientable surface of genus 0 only admits spherical structures (not considered here); a flat torus is a surface of genus 1 and admits Euclidean structures; an orientable surface of genus 2 and above admits only hyperbolic structures (see below).

From now on, \mathcal{M}_2 will denote either a flat torus or a closed hyperbolic surface.

Flat tori.

We denote by \mathbb{T}^2 the topological torus, that is, the product $\mathbb{T}^2 = \mathbb{S}^1 \times \mathbb{S}^1$ of two copies of the circle. Flat tori are obtained by taking the quotient of the Euclidean plane by an Abelian group generated by two independent translations. There are in fact many different Euclidean structures on \mathbb{T}^2 ; if one considers Euclidean structures up to homothety – which is sufficient for our purposes here – a Euclidean structure is uniquely determined by a vector u in the upper half-plane $\mathbb{R} \times \mathbb{R}_{>0}$: to such a vector u is associated the Euclidean structure $(\mathbb{T}^2, h_u) \sim \mathbb{R}^2 / (\mathbb{Z}e_1 + \mathbb{Z}u)$, where $e_1 = (1, 0)$ and $u = (u_x, u_y) \in \mathbb{R}^2$ is linearly independent from e_1 . The orbit of a point of the plane is a lattice. The area A_h of the surface is $|u_y|$. The plane \mathbb{R}^2 , equipped with the Euclidean metric, is then isometric to the universal cover of the corresponding quotient surface.

The space of flat structures on a torus, considered up to diffeomorphism isotopic to the identity and up to scaling, can in fact be identified to upper half-plane and thus to the hyperbolic space, see *e.g.* [19, Chapter 1].

Orientable hyperbolic surfaces.

We now consider a closed oriented surface S (a compact oriented surface without boundary) of genus $g \geq 2$. Such a surface does not admit any Euclidean structure, but it admits many *hyperbolic* structures, corresponding to metrics of constant curvature -1 , locally modeled on the hyperbolic plane \mathbb{H}^2 . Given a hyperbolic structure h on S , the surface (S, h) is isometric to the quotient \mathbb{H}^2/G , where G is a (non-Abelian) discrete subgroup of the isometry group $PSL(2, \mathbb{R})$ of \mathbb{H}^2 isomorphic to the fundamental group $\pi_1(S)$. The universal cover \tilde{S} is isometric to the hyperbolic plane \mathbb{H}^2 .

The space of hyperbolic metrics on S , considered up to diffeomorphisms isotopic to the identity, is the *Teichmüller space* of S , denote by \mathcal{T}_S . It is a manifold of dimension $6g - 6$, diffeomorphic to a ball, see *e.g.* [19]. An easy way to visualize this space is to think of pants decompositions, a crucial notion that will use throughout this paper. Indeed, on a pair of pants, i.e., a sphere with three geodesic holes, there is a unique hyperbolic metric that induces prescribed length for the boundaries. Then, a surface of genus g can be decomposed into $2g - 2$

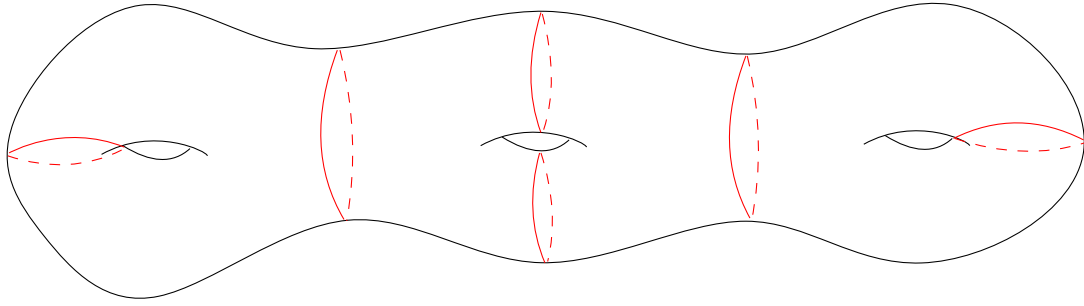


Figure 1: A pants decomposition of a surface of genus 3 with 6 curves and 4 pairs of pants.

pairs of pants for a total of $3g - 3$ cutting curves (see Figure 2.1). Giving two parameters by curve (its length and the way to glue the two hyperbolic pairs of pants together as a twist parameter) fix a unique hyperbolic metric. This set of $6g - 6$ parameters are called the Fenchel-Nielsen coordinates and gives an explicit parameterization of the Teichmüller space.

2.2 The Poincaré disk model of the hyperbolic plane

In the Poincaré disk model, the hyperbolic plane is represented as the open unit disk \mathbb{D}^2 of \mathbb{C} .

The geodesic lines consist of circular arcs contained in the disk \mathbb{D}^2 and that are orthogonal to its boundary (Figure 2 (left)). The model is conformal, i.e., the Euclidean angles measured in the plane are equal to the hyperbolic angles.

We will not need the exact expression of the hyperbolic metric here. However, the notion of a hyperbolic circle is relevant to us. Three non-collinear points in the hyperbolic plane \mathbb{H}^2 determine a *circle*, which is the restriction to the Poincaré disk of a Euclidean circle or line. If C is a Euclidean circle or line and $\phi : \mathbb{D}^2 \rightarrow \mathbb{D}^2$ is an isometry of the hyperbolic plane, then $\phi(C \cap \mathbb{D}^2)$ is still the intersection with \mathbb{D}^2 of a Euclidean circle or a line.

A key difference with the Euclidean case is that the “circle” defined by 3 non-collinear points in \mathbb{H}^2 is generally not compact (i.e., it is not included in the Poincaré disk). The compact circles are sets of points at constant (hyperbolic) distance from a point. Non-compact circles are either hypercycles, i.e., connected components of the set of points at constant (hyperbolic) distance from a hyperbolic line, or horocycles (Figure 2 (right)) [17].¹

Therefore, the relatively elementary tools that can be used for flat tori must be refined for hyperbolic surfaces. Still, some basic properties of circles still hold for non-compact circles. A non-compact circle splits the hyperbolic plane into two connected regions; when it is not a geodesic, we will call a *disk* the region of the Poincaré disk that is convex, in the hyperbolic sense (Figure 3). When a non-compact circle is determined by the three vertices of a triangle, it bounds a disk, and this disk contains the whole triangle.

Triangulations of hyperbolic spaces have been studied [3] and implemented in CGAL in 2D [4]. Note that that previous work was not considering non-compact circles as circles.

¹A synthetic presentation can be found at [http://en.wikipedia.org/wiki/Hypercycle_\(geometry\)](http://en.wikipedia.org/wiki/Hypercycle_(geometry))

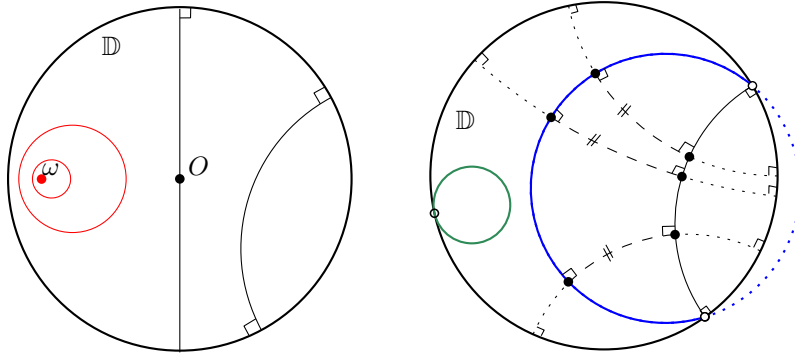


Figure 2: The Poincaré disk. Left: Geodesic lines (black) and compact circles (red) centered at point ω . Right: A horocycle (green). A hypercycle (blue), whose points have constant distance from the black geodesic line.

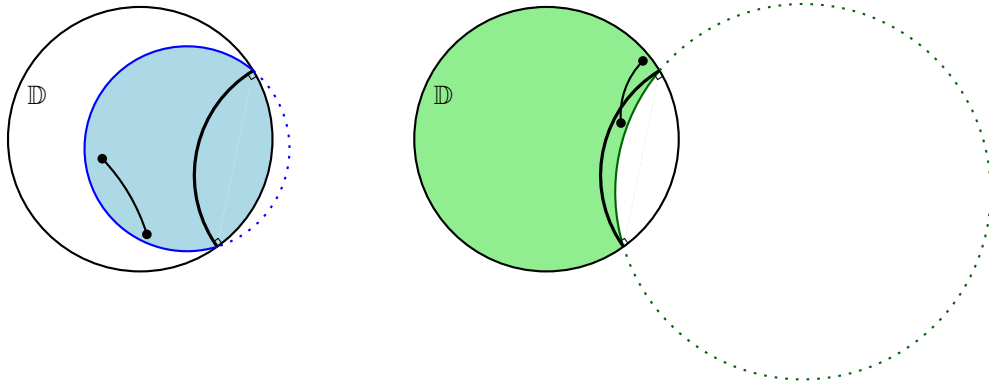


Figure 3: Shaded: Two (convex) non-compact disks.

2.3 Triangulations on surfaces

Let (\mathcal{M}_2, h) be either a torus (\mathbb{T}^2, h) equipped with a Euclidean structure h or a closed surface (S, h) equipped with a hyperbolic structure h .

For a given finite set of points $V \subset \mathcal{M}_2$, we will consider any two topological triangulations T and T' of \mathcal{M}_2 with vertex set V as *equivalent* if for any two vertices u and v in V , the edges of T with vertices u and v are in one-to-one correspondence with the edges of T' with the same vertices u and v through homotopies with fixed points.

Recall that given two distinct points $v, w \in \mathcal{M}_2$, any homotopy class of paths on \mathcal{M}_2 with endpoints v and w contains a unique geodesic segment [14, Chapter 1]. So, any triangulation is equivalent to a unique *geodesic* triangulation, i.e., a triangulation whose edges are geodesic segments. Note that the edges of a geodesic triangulation can intersect in their interiors.

We can now recall the following simple notion of geometric triangulation.

Definition 2.1. A triangulation T on \mathcal{M}_2 is said to be *geometric* for h if the edges of its equivalent geodesic triangulation do not intersect except at common endpoints.

If T is a triangulation of \mathcal{M}_2 , its inverse image $\rho^{-1}(T)$ is the (infinite) triangulation of $\widetilde{\mathcal{M}}_2$ whose vertices, edges and faces are the lifted images by ρ^{-1} of those of T .

Definition 2.2. We say that a triangulation T of \mathcal{M}_2 is a *Delaunay* triangulation if for each face \tilde{f} of $\rho^{-1}(T)$, the open disk in $\widetilde{\mathcal{M}}_2$ that is bounded by the circle passing through the three vertices of \tilde{f} is *empty*, i.e., it contains no vertex of $\rho^{-1}(T)$.

It follows that if T is a geodesic Delaunay triangulation of \mathcal{M}_2 with vertex set V , then $\rho^{-1}(T)$ is the Delaunay triangulation in $\widetilde{\mathcal{M}}_2$ of $\rho^{-1}(V)$. So, for a non-degenerate set of points on \mathcal{M}_2 (i.e., a set whose lift does not contain four cocyclic points), since the Delaunay triangulation of their lifts in $\widetilde{\mathcal{M}}_2$ is unique, there is also a unique geodesic Delaunay triangulation on \mathcal{M}_2 . However its edges may *a priori* intersect.

For a degenerate set V of points, at least two adjacent triangles in the possible Delaunay triangulations of $\rho^{-1}(V)$ in $\widetilde{\mathcal{M}}_2$ have cocircular vertices. Any triangulation of the subset \mathcal{C} of $\rho^{-1}(V)$ consisting of c cocircular points is a Delaunay triangulation. Any of these triangulations can be transformed in any other by $O(c)$ flips [18]. From now on, we can thus assume that the set of points V on the surfaces that we consider is always non-degenerate.

We will see in Section 3 that any Delaunay triangulation of \mathcal{M}_2 is in fact geometric.

Remark 2.3. For a hyperbolic surface, the closure of every empty disk in the universal cover \mathbb{H}^2 is compact.

Proof. Let us define the *diameter* $\Delta(T)$ of a geodesic triangulation T as the smallest diameter of a fundamental domain that is the union of lifts of the triangles of T in $\widetilde{\mathcal{M}}_2$. The diameter $\Delta(T)$ is not smaller than the diameter of (S, h) . It is unclear how to compute $\Delta(T)$ algorithmically and the problem looks difficult. However bounds are easy to obtain: $\Delta(T)$ is at least equal to the maximum of the diameters of the triangles of $\rho^{-1}(T)$ in $\widetilde{\mathcal{M}}_2$ and is at most the sum of the diameters of these triangles.

Any non-compact disk contains at least one disk of any diameter, so, at least one disk of diameter $\Delta(T)$, thus it contains a fundamental domain (actually, infinitely many fundamental domains) and its interior cannot be empty. \square

Let us now give a natural definition for flips in triangulations of surfaces. It is based on the usual notion of flips in the Euclidean plane.

Definition 2.4. Let T be a geometric triangulation of \mathcal{M}_2 . Let (v_1, v_2, v_3) and (v_2, v_1, v_4) be two adjacent triangles in T , sharing the edge $e = (v_1, v_2)$. Let us lift the quadrilateral (v_1, v_2, v_3, v_4) to a quadrilateral $(\tilde{v}_1, \tilde{v}_2, \tilde{v}_3, \tilde{v}_4)$ in $\widetilde{\mathcal{M}}_2$ so that $(\tilde{v}_1, \tilde{v}_2, \tilde{v}_3)$ and $(\tilde{v}_2, \tilde{v}_1, \tilde{v}_4)$ form two adjacent triangles of $\rho^{-1}(T)$ sharing the edge $\tilde{e} = (\tilde{v}_1, \tilde{v}_2)$.

Flipping e in T consists of replacing the diagonal \tilde{e} in the quadrilateral $(\tilde{v}_1, \tilde{v}_2, \tilde{v}_3, \tilde{v}_4)$ (which lies in $\widetilde{\mathcal{M}}_2$, i.e., \mathbb{R}^2 or \mathbb{H}^2) by the geodesic segment $(\tilde{v}_3, \tilde{v}_4)$, then projecting the two new triangles $(\tilde{v}_3, \tilde{v}_4, \tilde{v}_2)$ and $(\tilde{v}_4, \tilde{v}_3, \tilde{v}_1)$ to \mathcal{M}_2 by ρ .

We say that the flip of T along e is *Delaunay* if the triangulation is *locally Delaunay* in the quadrilateral after the flip, i.e., the disk inscribing $(\tilde{v}_3, \tilde{v}_4, \tilde{v}_2)$ does not contain \tilde{v}_1 (and the disk inscribing $(\tilde{v}_4, \tilde{v}_3, \tilde{v}_1)$ does not contain \tilde{v}_2).

An edge e is said to be *Delaunay flippable* if the flip along e is Delaunay.

Note that even though T is geometric in this definition, the triangulation after a flip is not necessarily geometric. We will prove later (Lemma 4.1) that a Delaunay flip transforms a geometric triangulation into a geometric triangulation.

Triangulations and polyhedral surfaces.

The complex plane can be identified with the plane ($z = 1$) in \mathbb{R}^3 , while the Poincaré model of the hyperbolic plane can be identified with the unit disk in that plane. We can now use the stereographic projection $\sigma : \mathbb{S}^2 \setminus \{s_0\} \rightarrow \mathbb{R}^2$ to send the unit sphere \mathbb{S}^2 to this plane ($z = 1$), where $s_0 = (0, 0, -1)$ is the south pole. In this projection, each point $p \neq s_0$ on the sphere is sent to the unique intersection with the plane ($z = 1$) of the line going through s_0 and p . The inverse image of the plane ($z = 1$) is $\mathbb{S}^2 \setminus \{s_0\}$, while the inverse image of the disk containing the Poincaré model of the hyperbolic plane is a disk, which is the set of points of \mathbb{S}^2 above a horizontal plane.

Let T^* be a triangulation of the Euclidean or the hyperbolic plane – for instance, T^* could be the inverse image $\rho^{-1}(T)$ of a triangulation T of a surface (\mathcal{M}_2, h) , in which case T^* has infinitely many vertices. We associate to T^* a polyhedral surface Σ in \mathbb{R}^3 , constructed as follows. The construction is similar to the classic duality originally presented with a paraboloid in the case of (finite) triangulations in a Euclidean space [13]. It can also be seen as a simpler version, sufficient for our purpose, of the construction presented for triangulations in hyperbolic spaces using the space of spheres [3].

- The vertices of Σ are the inverse images on \mathbb{S}^2 by σ of the vertices of T^* .
- The edges of Σ are line segments in \mathbb{R}^3 corresponding to the edges of T^* and the faces of Σ are triangles in \mathbb{R}^3 corresponding to the faces of T^* .

Note that Σ is not necessarily convex. We can make the following well-known remarks. Let t_1 and t_2 be two triangles of T^* sharing an edge e , and let t_1^Σ and t_2^Σ be the corresponding faces of the polyhedral surface Σ , sharing the edge e^Σ . Then Σ is concave at e^Σ if and only if e is Delaunay flippable. Flipping e in the triangulation T^* in the plane corresponds to replacing the two faces t_1^Σ and t_2^Σ of Σ by the two other faces of the tetrahedron formed by their vertices. That tetrahedron lies between Σ and \mathbb{S}^2 . We obtain a new edge $e^{\Sigma'}$ at which the new polyhedral surface Σ' is convex, and which is strictly closer to \mathbb{S}^2 than Σ . By an abuse of language, we will say that Σ' *contains* Σ , which we will denote as $\Sigma \subset \Sigma'$.

As a consequence, Σ is convex if and only if T^* is Delaunay.

There is a direct corollary of this statement: Given a (non-degenerate, see above) discrete set V of points in \mathbb{R}^2 or \mathbb{H}^2 , there is a unique Delaunay triangulation with this set of vertices.

However we are going to see in the next two sections that there can be infinitely many geometric (non-Delaunay) triangulations on a surface, with the same given finite vertex set.

Polyhedral surfaces and geometric triangulations

A first, geometric way to understand geometric and non-geometric triangulations is through the polyhedral surfaces associated to a triangulation, as seen in the previous section.

It is helpful in this respect to change slightly the construction of the previous section, by scaling the disk containing the Poincaré model, so that it becomes a disk of radius 2. The inverse image of this disk by the stereographic projection is then precisely the upper hemisphere in \mathbb{S}^2 . Since the stereographic projection sends lines and circles to circles in \mathbb{S}^2 , the hyperbolic geodesics are sent to circles in \mathbb{S}^2 . Since hyperbolic geodesics correspond to circles (and lines) orthogonal to the boundary, and since the stereographic projection

preserves angles, the hyperbolic geodesics correspond precisely to the half-circles orthogonal to the equator in the upper hemisphere.

As a consequence, those hyperbolic geodesics are represented exactly by the intersections with the upper hemisphere of the vertical planes. This means that the vertical projection of the upper hemisphere to the horizontal plane ($z = 0$) projects the Poincaré model to the Klein model, where hyperbolic geodesics correspond to intersections of lines with the disk.

A triangulation is geometric if and only if, once each edge is realized as a geodesic segment, those edges are not intersecting in their interiors. The following remark follows directly.

Remark 2.5. Let T^* be a triangulation of the hyperbolic plane, and let Σ be the polyhedral surface corresponding to T^* constructed in the previous section. Then T^* is geometric if and only if Σ is a graph above the unit disk, that is, if and only if the vertical projection of Σ to the horizontal plane is injective.

3 Geometric triangulations of surfaces

We now consider Dehn twists, which are usually considered as acting on the space of metrics on a surface [7], but are defined here equivalently, for simplicity, as acting on triangulations of a closed oriented surface (\mathcal{M}_2, h) equipped with a fixed Euclidean or hyperbolic structure (figures in this section illustrate the flat case, but the results are proved for both flat and hyperbolic cases). Let T be a triangulation of (\mathcal{M}_2, h) , with vertex set V , and let c be an oriented homotopically non-trivial simple closed curve on $\mathcal{M}_2 \setminus V$. We define a new triangulation $\tau_c(T)$ of \mathcal{M}_2 by performing a *Dehn twist* along c : whenever an edge e of T intersects c at a point p , we orient e so that the unit vectors of the tangent plane along e and c form a positively oriented basis (see Figure 4 (left)), and then replace e by the oriented path following e until p , then following c until it comes back to p , then following e until its endpoint (see Figure 4 (right)). This defines a map τ_c from the space of triangulations of \mathbb{T}^2 with vertex set V to itself. Note that, even if T is a geometric triangulation, $\tau_c(T)$ is not necessarily geometric. If we denote by $-c$ the curve c with the opposite orientation, then one easily checks that $\tau_{-c} = \tau_c^{-1}$.

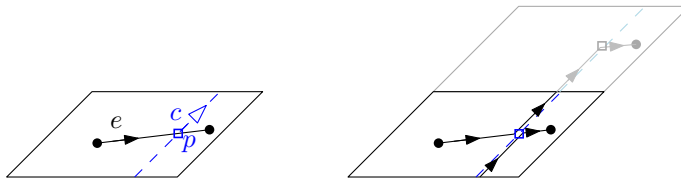


Figure 4: Transformation of an edge e by the Dehn twist along c on a flat torus \mathbb{T}^2 . Here the black parallelepiped is a fundamental domain, and the gray one, used for the construction of the image of e by τ_c , is another fundamental domain, image through an element of the group Γ of isometries.

Lemma 3.1. *There exists a geometric triangulation T of (\mathcal{M}_2, h) and a simple closed curve $c \subset \mathcal{M}_2$ such that for all $k \in \mathbb{Z}$, $\tau_c^k(T)$ is geometric.*

Proof. Let us focus on the hyperbolic case (the construction is easier in the flat case). Consider a pants decomposition of \mathcal{M}_2 and denote as \mathcal{C} the set of its boundary curves, which are simple

closed geodesics. Let us choose c in \mathcal{C} and $\varepsilon > 0$. We denote by c_-, c_+ the two hypercycles at distance ε from c on both sides of c . The value of ε must be sufficiently small so that the region between c_- and c_+ is an annulus drawn on \mathcal{M}_2 that does not intersect any curve in $\mathcal{C} \setminus \{c\}$. Each curve in $\mathcal{C} \setminus \{c\}$ is split into two geodesic segments by putting two points on it; let us add those two segments as edges of T . Let us put two points on c_- (resp. c_+) and add as edges of T the two geodesic segments between them, whose union forms a curve homotopic to c . Each pair of pants not bounded by c , as well as the two “shortened pants” bounded by c_- and c_+ (i.e., the pants obtained from the initial pants by removing the annulus between c_- and c_+), can be decomposed into two hexagons, which can easily be triangulated with geodesic edges. All these edges are left unchanged by τ_c (or τ_{-c}) as they do not intersect c . The annulus between c_- and c_+ can be triangulated with four edges each intersecting c exactly once. We realize the image by τ_c of each of these four edges as a geodesic segment – there is a unique choice in the homotopy class of the path described above (Figure 5). The annulus is convex, as the projection onto \mathcal{M}_2 of the intersection of two (convex) disks (Figure 3), so, the geodesic segment is completely contained in it. Let e, e' be two edges of T . If neither

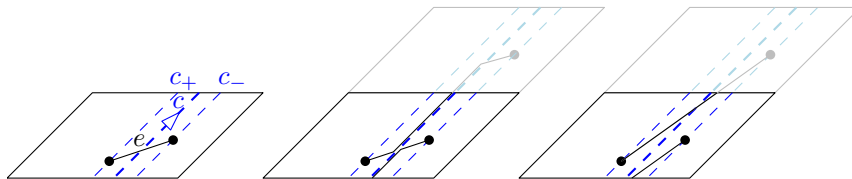


Figure 5: Image of e by a Dehn twist (middle), realized as a geodesic edge (right).

e nor e' intersect c , then they are not changed by τ_c , so they remain disjoint. If either e or e' intersects c , but not both, then their images by τ_c (or τ_{-c}) remain disjoint, as they lie in different regions separated by c_- and c_+ . If e and e' intersect c , then again their images by τ_c (or τ_{-c}) remain disjoint, as their endpoints appear in the same order on c_- and c_+ and two geodesic lines cannot intersect more than once (Figure 6). As a consequence, $\tau_c(T)$ and

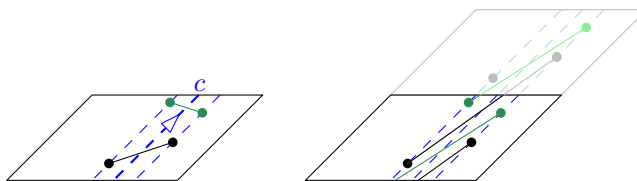


Figure 6: The Dehn twist of two edges along c for two edges intersecting c .

$\tau_{-c}(T)$ are geometric. They are not equivalent as each edge e crossing c is replaced by an edge that does not lie in the same homotopy class as e . The same result follows by induction for $\tau_c^k(T)$ for any $k \in \mathbb{Z}$. \square

Corollary 3.2. *For any closed oriented surface (\mathcal{M}_2, h) , there exists a finite set of points $V \subset \mathcal{M}_2$ such that the graph of geometric triangulations with vertex set V is infinite.*

Proposition 3.3. *Any Delaunay triangulation of a closed oriented surface (\mathcal{M}_2, h) is geometric.*

Proof. Let V be a finite set of points on \mathcal{M}_2 , and let T be a Delaunay triangulation of (\mathcal{M}_2, h) with vertex set V . Realize every edge of T as the unique geodesic segment in its homotopy class, so that T is geodesic. We argue by contradiction and suppose that T is not geometric, so that there are two edges e_1 and e_2 that intersect in their interiors. We then lift e_1 and e_2 to edges \tilde{e}_1 and \tilde{e}_2 of $\rho^{-1}(T)$ whose interiors still intersect.

We can find two distinct faces \tilde{f}_1 and \tilde{f}_2 of $\rho^{-1}(T)$ such that \tilde{e}_1 is an edge of \tilde{f}_1 and \tilde{e}_2 is an edge of \tilde{f}_2 . Let \tilde{C}_1 and \tilde{C}_2 be the circles inscribing \tilde{f}_1 and \tilde{f}_2 , respectively. Since $\rho^{-1}(T)$ is Delaunay, \tilde{C}_1 and \tilde{C}_2 bound empty disks \tilde{D}_1 and \tilde{D}_2 , i.e., open disks not containing any point of $\rho^{-1}(V)$. Recall that, as mentioned in Section 2.3, the closures of empty disks are compact even in the hyperbolic case, and that $\tilde{e}_1 \subset \tilde{D}_1$ and $\tilde{e}_2 \subset \tilde{D}_2$ (edges are considered as open). The two circles \tilde{C}_1 and \tilde{C}_2 intersect twice as the intersection point of \tilde{e}_1 and \tilde{e}_2 lies in $\tilde{D}_1 \cap \tilde{D}_2$. Let \tilde{L} be the geodesic line through the intersection points. The endpoints of \tilde{e}_1 are on $\tilde{C}_1 \setminus \tilde{D}_2$ and those of \tilde{e}_2 are on $\tilde{C}_2 \setminus \tilde{D}_1$, so the two pairs of endpoints are on opposite sides of \tilde{L} . As a consequence, \tilde{e}_1 and \tilde{e}_2 are on opposite sides of \tilde{L} , so they cannot intersect. This leads to a contradiction. \square

4 The flip algorithm

Let us consider a closed oriented surface (\mathcal{M}_2, h) . The flip algorithm consists of performing Delaunay flips in any order, starting from a given input geometric triangulation of \mathcal{M}_2 , until there is no more Delaunay flippable edge.

The following statement is a key starting point for proving the correctness of the algorithm.

Lemma 4.1. *Let T be a geometric triangulation of (\mathcal{M}_2, h) , and let T' be obtained from T by a Delaunay flip. Then T' is still geometric.*

Proof. Let e be a Delaunay flippable edge and \tilde{e} a lift in $\widetilde{\mathcal{M}}_2$. Denote the vertices of \tilde{e} by \tilde{v} and \tilde{v}' . Let \tilde{t}_1 and \tilde{t}_2 be the triangles of $\rho^{-1}(T)$ incident to \tilde{e} . To prove that T' is geometric, it is sufficient to prove that $\tilde{t}_1 \cup \tilde{t}_2$ is a strictly convex quadrilateral.

Let \tilde{C}_1 (resp. \tilde{C}_2) be the circle through the three vertices of \tilde{t}_1 (resp. \tilde{t}_2). Note that \tilde{C}_1 and \tilde{C}_2 may be non-compact. Let \tilde{D}_1 and \tilde{D}_2 be the corresponding disks (as defined in Section 2.2 in the case of non-compact circles). The disk \tilde{D}_1 (resp. \tilde{D}_2) is convex (in the Euclidean plane if \mathcal{M}_2 is a flat torus, or in the sense of hyperbolic geometry if \mathcal{M}_2 is a hyperbolic surface) and contains \tilde{t}_1 (resp. \tilde{t}_2). The fact that e is Delaunay flippable then implies that \tilde{t}_1 and \tilde{t}_2 are contained in $\tilde{D}_1 \cap \tilde{D}_2$ (see Figure 7). As a consequence, the sum of angles of \tilde{t}_1 and \tilde{t}_2 at \tilde{v} is smaller than the interior angle at \tilde{v} of $\tilde{D}_1 \cap \tilde{D}_2$, which is at most π , and similarly at \tilde{v}' . As a consequence, the quadrilateral $\tilde{t}_1 \cup \tilde{t}_2$ is strictly convex at \tilde{v} and \tilde{v}' . Since it is strictly convex at its other two vertices (as each of these vertices is a vertex of a triangle), it is strictly convex, and the statement follows. \square

The following lemma is central in the proof of the termination of the algorithm (Theorem 4.6) for hyperbolic surfaces and in its analysis for both flat tori and hyperbolic surfaces (Section 5).

Lemma 4.2. *Let T be a geometric triangulation of (\mathcal{M}_2, h) . Then, the flip algorithm starting from T will never insert an edge longer than $2\Lambda(T)$, where $\Lambda(T)$ denotes the length of the longest edge(s) in T .*

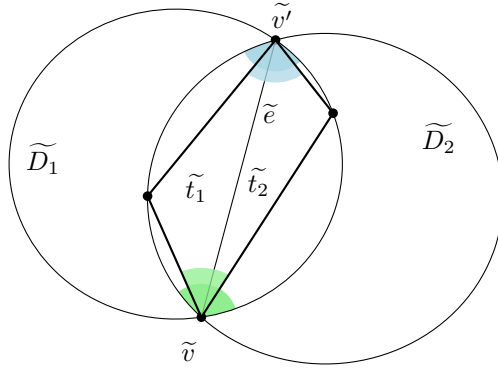


Figure 7: The quadrilateral is convex (edges are represented schematically as straight line segments).

Note that the length of an edge can be measured on any of its lifts in the universal covering space $\widetilde{\mathcal{M}}_2$.

Proof. Let T_k be the triangulation obtained from $T = T_0$ after k flips and let Σ_k be the corresponding polyhedral surface of \mathbb{R}^3 as defined in Section 2.3. Since we perform only Delaunay flips, $\Sigma_0 \subset \dots \subset \Sigma_k \subset \Sigma_{k+1}$.

We will prove the result by contradiction. Let us assume that T_k has an edge e of length larger than $2\Lambda(T)$. Let v be the midpoint of e and t be the triangle in $T = T_0$ that contains v (if v lies on an edge of T , any incident triangle can be chosen as t). Let us lift e, v , and t in a consistent way, so that \tilde{v} is the midpoint of \tilde{e} and lies in \tilde{t} . Denote the endpoints of \tilde{e} as \tilde{v}_1 and \tilde{v}_2 . The triangle \tilde{t} is strictly included in the disk \tilde{D} of radius $\Lambda(T)$ and centered at \tilde{v} , by definition of $\Lambda(T)$ (see Figure 8 (left)).

Let P_D denote the plane in \mathbb{R}^3 containing the circle on \mathbb{S}^2 that is the boundary of $\sigma^{-1}(D)$ (recall that σ denotes the stereographic projection, see Section 2.3), and let p denote the point $\sigma^{-1}(\tilde{v})$ on \mathbb{S}^2 . As $p \in \sigma^{-1}(t) \subset \sigma^{-1}(\tilde{D})$, the projection p^{Σ_0} of p onto Σ_0 lies above P_D (Figure 8 (right)).

Now, denote the edge $\sigma^{-1}(\tilde{e})$ on \mathbb{S}^2 as (p_1, p_2) . The points p_1 and p_2 lie outside $\sigma^{-1}(D)$ since \tilde{e} is longer than $2\Lambda(T)$. So, the corresponding edge $e^\Sigma = [p_1, p_2]$ of Σ_k lies below the plane P_D , thus the projection $p^{\Sigma_k} \in [p_1, p_2]$ of p onto Σ_k lies below P_D .

From what we have shown, p^{Σ_k} is a point of Σ_k that lies strictly between the pole s_0 and the point p^{Σ_0} of Σ_0 , which contradicts the inclusion $\Sigma_0 \subset \Sigma_k$. \square

We will now show that, for any order, the flip algorithm terminates and returns the Delaunay triangulation of the surface. The proof given for the hyperbolic case would also work for the flat case. However we propose a more elementary proof for the flat case.

Flat tori

The case of flat tori is easy, and might be considered as folklore. However, as we have not found a reference, we give the details here for completeness.

We define the weight of a triangle t of a geometric triangulation T of \mathbb{T}^2 as the number of vertices of $\rho^{-1}(T)$ that lie in the open circumdisk of a lift of t . The weight $w(T)$ of T is

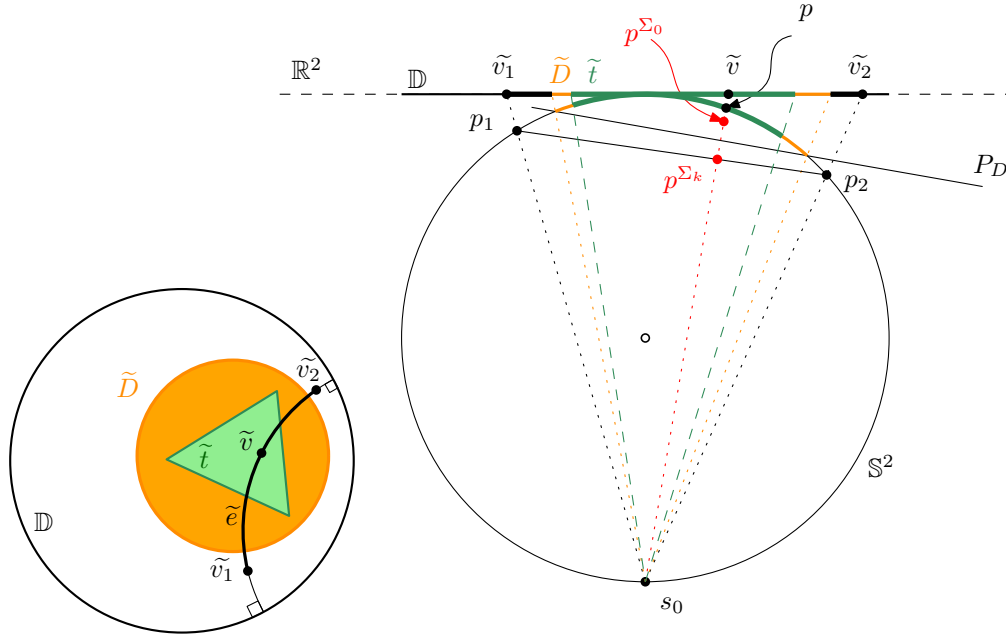


Figure 8: Illustration for the proof of Lemma 4.2 (for a hyperbolic surface). Left: notation in \mathbb{H}^2 . Right: contradiction seen in a cutting plane in \mathbb{R}^3 .

defined as the sum of the weights of its triangles.

Lemma 4.3. *The weight $w(T)$ of a triangulation T of a flat torus (\mathbb{T}^2, h) is finite. Let T' be the triangulation obtained from a geometric triangulation T after performing a Delaunay flip. Then $w(T') \leq w(T) - 2$.*

Proof. The closed circumdisk of any triangle in \mathbb{R}^2 is compact, so, it can only contain a finite number of vertices of $\rho^{-1}(T)$. The sum $w(T)$ of these numbers over triangles of T is clearly finite as the number of triangles of T is finite. Let us now focus on a quadrilateral in \mathbb{R}^2 that is a lift of the quadrilateral on \mathbb{T}^2 whose diagonal e is flipped. Let \tilde{D}_1 and \tilde{D}_2 denote the two open circumdisks in \mathbb{R}^2 before the flip and \tilde{D}'_1 and \tilde{D}'_2 denote the two open circumdisks after the flip, then $\tilde{D}'_1 \cup \tilde{D}'_2 \subset \tilde{D}_1 \cup \tilde{D}_2$ and $\tilde{D}'_1 \cap \tilde{D}'_2 \subset \tilde{D}_1 \cap \tilde{D}_2$ (see Figure 9). Moreover, by definition of a Delaunay flip, the union $\tilde{D}'_1 \cup \tilde{D}'_2$ contains at least two fewer vertices of $\rho^{-1}(T)$ than $\tilde{D}_1 \cup \tilde{D}_2$, which are the two vertices of the quadrilateral that are not vertices of \tilde{e} . This concludes the proof. \square

The result follows trivially:

Theorem 4.4. *Let T be a geometric triangulation of a flat torus with finite vertex set V . The flip algorithm terminates and outputs the Delaunay triangulation of V .*

Corollary 4.5. *The geometric flip graph $\mathcal{F}_{\mathbb{T}^2, h, V}$ is connected.*

Hyperbolic surfaces

To show that the flip algorithm terminates in the hyperbolic case, we cannot mimic the proof presented for the flat tori since the circumcircle of a hyperbolic triangle can be non-compact

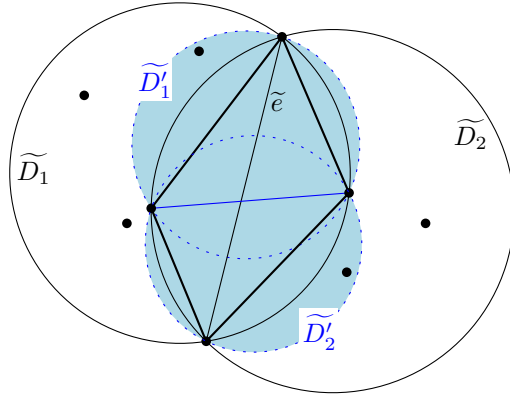


Figure 9: Circumdisks \widetilde{D}_1 and \widetilde{D}_2 before flipping \widetilde{e} and \widetilde{D}'_1 and \widetilde{D}'_2 after the Delaunay flip.

(see Section 2.2) and thus can have an infinite weight. Note also that the proof cannot use a property on the angles of the Delaunay triangulation similar to what holds in the Euclidean case: in \mathbb{H}^2 , the locus of points seeing a segment with a given angle is not a circular arc, and thus the Delaunay triangulation of a set of points in \mathbb{H}^2 does not maximize the smallest angle of triangles (Figure 10) The proof relies on Lemma 4.2.

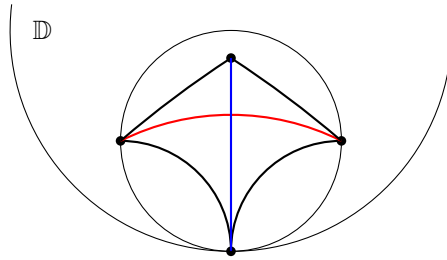


Figure 10: The Delaunay triangulation contains the blue edge, which subdivides the smallest angle.

Theorem 4.6. *Let T be a geometric triangulation of a closed hyperbolic surface with finite vertex set V . The flip algorithm terminates and outputs the Delaunay triangulation of V .*

Proof. We use the same notation as in the proof Lemma 4.2. Once an edge of T_k is flipped, it can never reappear in the triangulation, as the corresponding segment in \mathbb{R}^3 becomes interior to the polyhedral surface Σ_{k+1} (see Section 2.3) and further surfaces $\Sigma_{k'}, k' \geq k + 1$. In addition, all the introduced edges have length smaller than $2\Lambda(T)$ by Lemma 4.2. Moreover, there is only a finite number of edges with vertices in V that are shorter than $2\Lambda(T)$ on S , as a circle given by a center and a bounded radius is compact. So, the flip algorithm terminates. The output does not have any Delaunay flippable edge, so, it is the Delaunay triangulation. \square

Corollary 4.7. *The geometric flip graph $\mathcal{F}_{S,h,V}$ is connected.*

5 Algorithm analysis

In this section, we first describe a data structure that supports the flip algorithm.

A 2-dimensional combinatorial map is a natural data structure for storing a graph cellularly embedded on a surface. We refer to the literature for a formal definition [24, Section 3.3]. In a few words, similarly to a halfedge data structure, a combinatorial map stores the graph of edges. The *dart* (or *flag*) gives access to incidence relations in the graph (Figure 11) through three permutations $\beta_0, \beta_1,$ and β_2 of the set of darts. Permutations β_0 and β_1 are inverse of

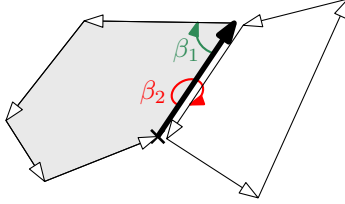


Figure 11: A dart (bold) in a combinatorial map.

each other; the permutation β_2 is an involution. A face is given by a cycle of darts related by powers of β_1 (or β_0). Faces are glued along their boundaries: two paired darts are related through β_2 .

Each triangulation that we handle is geometric. The underlying combinatorial triangulation is stored as a combinatorial map, while the geometry is given by cross-ratios on each edge.

Recall the definition of a cross-ratio in \mathbb{H}^2 [1]: for four pairwise-distinct points $z_1, z_2, z_3, z_4 \in \mathbb{H}^2$, it is the complex number

$$[z_1, z_2, z_3, z_4] = \frac{(z_4 - z_2)(z_3 - z_1)}{(z_4 - z_1)(z_3 - z_2)}.$$

Cross-ratios encode the condition for a Delaunay flip: If z_1, z_2, z_3, z_4 are oriented counterclockwise, then $\text{Im}[z_1, z_2, z_3, z_4] > 0$ if and only if z_4 lies inside the open disk circumscribing the triangle (z_1, z_2, z_3) (here $\text{Im}[z]$ denotes the imaginary part of a complex).

Given an edge e of a triangulation of \mathcal{M}_2 we consider a lift $\tilde{e} = (\tilde{u}_1, \tilde{u}_3)$ of e in \mathbb{D}^2 . Denote as \tilde{u}_2 and \tilde{u}_4 the remaining vertices of the two faces incident to \tilde{e} in the lifted triangulation, where vertices are numbered counterclockwise. The cross-ratio of e is defined as the cross-ratio of $\tilde{u}_1, \tilde{u}_2, \tilde{u}_3,$ and \tilde{u}_4 in \mathbb{D}^2 ; it is independent of the choice of the lift of e , as the cross-ratio is invariant under orientation preserving isometries of \mathbb{D}^2 . An edge e of a triangulation is Delaunay-flippable if and only if the imaginary part of its cross-ratio is positive.

When an edge is flipped in a triangulation, updating the combinatorial map that represents it is straightforward. Updating the cross-ratios of the five involved edges of the quadrilateral is also easy; details can be found in [9].

With this data structure, a flip can be performed in constant time.

For a triangulation on n vertices in the Euclidean plane, counting the weights of triangulations leads to the optimal $O(n^2)$ bound. However the same argument does not yield a bound even for the flat torus, since points must be counted in the universal cover.

Our previous complexity analysis for the flat torus [11] was recently improved:

Theorem 5.1 ([12]). *For any triangulation T with n vertices of a torus (\mathbb{T}^2, h) , there is a sequence of flips of length $C_h \cdot \Lambda(T) \cdot n^2$ connecting T to a Delaunay triangulation of (\mathbb{T}^2, h) , where C_h only depends on h .*

The rest of this section is devoted to estimating the number of edges not longer than $2\Lambda(T)$ between two fixed points v_1 and v_2 on a hyperbolic surface (S, h) . Counting the number of points in a disk of fixed radius would give an exponential bound because the area of a circle in \mathbb{H}^2 is exponential in its radius [21]. Note that we only consider geodesic edges, so we only need to count homotopy classes of simple paths. The behavior of the number N_l of simple closed curves smaller than a fixed length l is well understood: N_l/l^{6g-6} converges to a positive constant depending “continuously” on the structure h [23]. However, we need a result for geodesic segments instead of geodesic closed curves, and Mirzakhani’s proof is too deep and relies on too sophisticated structures to easily be generalized. So, we will only prove an upper bound on the number of segments. Such an upper bound could be derived from the theory of measured laminations of Thurston, which is also quite intricate. Fortunately, a more comprehensible proof, specific to simple closed geodesic curves on hyperbolic structures, can be found in [16, 4.III, p.61-67] [15]. While recalling the main steps of the proof, we show how to extend it to geodesic segments.

Let $\Gamma = \{\gamma_i, i = 1, \dots, 3g - 3\}$ be a set of $3g - 3$ simple disjoint closed geodesics on (S, h) not containing v_1 and v_2 that forms a pants decomposition on S , where each γ_i belongs to two different pairs of pants. A set $\{\bar{\gamma}_i, i = 1, \dots, 3g - 3\}$ of disjoint closed annuli is defined on S , where each $\bar{\gamma}_i$ is a tubular neighborhood of γ_i containing none of v_1, v_2 . This yields a decomposition of S into $3g - 3$ annuli $\bar{\gamma}_i$ ($i = 1, \dots, 3g - 3$) and $2g - 2$ pairs of “short pants” P_j ($j = 1, \dots, 2g - 2$). For $i = 1, \dots, 3g - 3$, let us denote as $\partial\bar{\gamma}_i$ any one of the two curves bounding the annulus $\bar{\gamma}_i$ (this is an abuse of notation but should not introduce any confusion). In each pair of pants $P_j, j = 1, \dots, 2g - 2$, for each of its boundary components $\partial\bar{\gamma}$, an arc $J_j^{\bar{\gamma}}$ is drawn in P_j between any two points on the boundary of $\bar{\gamma}$, such that it separates the other two boundary components of P_j and it has minimal length.

Two curves γ' and γ'' are associated to each $\gamma \in \Gamma$ in the following way (Figure 12). The annulus $\bar{\gamma}$ is glued with the two pairs of pants P_i and P_j between which it is lying, which yields a sphere with four boundary components: $\partial\bar{\gamma}_{i,1}$ and $\partial\bar{\gamma}_{i,2}$ bounding P_i and $\partial\bar{\gamma}_{j,1}$ and $\partial\bar{\gamma}_{j,2}$ bounding P_j . A curve γ' is then defined: it coincides with $J_i^{\bar{\gamma}}$ in P_i and $J_j^{\bar{\gamma}}$ in P_j , it separates $\partial\bar{\gamma}_{i,1}$ and $\partial\bar{\gamma}_{j,1}$ from $\partial\bar{\gamma}_{i,2}$ and $\partial\bar{\gamma}_{j,2}$, and it has exactly 2 crossings with γ . The curve γ'' is defined in the same way, separating $\partial\bar{\gamma}_{i,1}$ and $\partial\bar{\gamma}_{j,2}$ from $\partial\bar{\gamma}_{i,2}$ and $\partial\bar{\gamma}_{j,1}$.

For each P_i and $m_{i,1}, m_{i,2}, m_{i,3} \in \mathbb{N}$, a model multiarc is fixed in P_i , having $m_{i,1}, m_{i,2}$, and $m_{i,3}$ intersections with the three boundaries $\partial\bar{\gamma}_{i,1}, \partial\bar{\gamma}_{i,2}$, and $\partial\bar{\gamma}_{i,3}$ of P_i . The model is chosen among all the possible multiarcs as the one that has a minimal number of intersections with the three arcs $J_i^{\bar{\gamma},j}$ ($j = 1, 2, 3$) of P_i ; its endpoints lie on the curves $\gamma_{i,1}, \gamma_{i,2}$ and $\gamma_{i,3}$. The model is unique, up to homeomorphisms of the pair of pants, and those homeomorphisms are rather simple to understand since they can be decomposed into three Dehn twists around curves homotopic to the three boundaries of the pair of pants. It means that any homotopy class of multiarcs differs from the fixed corresponding model multiarc only by twisting along the three boundaries. Since these twists are independent, we can construct the chosen homotopy class as the fixed corresponding model on $P_i \setminus \{\bar{\gamma}_{i,1}, \bar{\gamma}_{i,2}, \bar{\gamma}_{i,3}\}$ and some paths winding around $\gamma_{i,1}, \gamma_{i,2}$ or $\gamma_{i,3}$ in the cylinders.

Let now f be a path between v_1 and v_2 on S . We decompose f into three parts: (v_1, w_1) ,

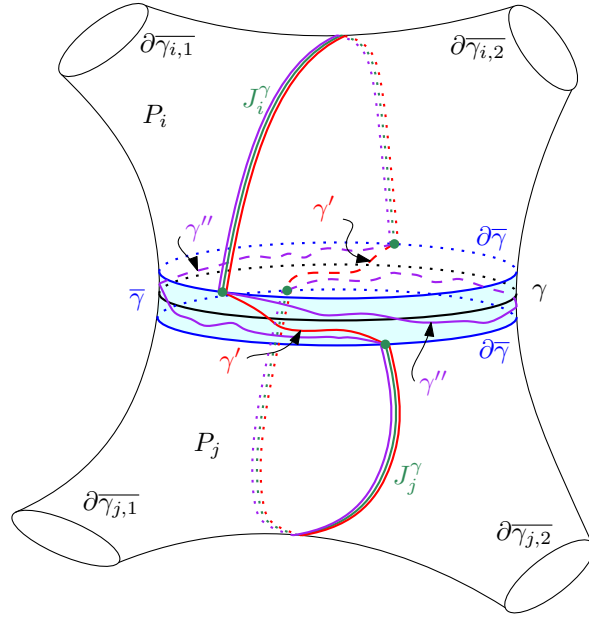


Figure 12: Two adjacent pairs of pants P_i and P_j .

$f^w = (w_1, w_2)$ and (w_2, v_2) where w_1 and w_2 are the first and the last points of f on an annulus boundary. We “push” all the twists of f^w into the annuli $\overline{\gamma}, \gamma \in \Gamma$, and obtain a *normal form* homotopic to f , whose definition adapts the definition given in the book [16] for closed curves:

1. It is simple.
2. It has a minimal number m_i of intersections with each $\gamma_i, i = 1, \dots, 3g - 3$.
3. In each $P_j, j = 1, \dots, 2g - 2$, it is homotopic with fixed endpoints to the model that corresponds to the number of intersections with its boundaries. For P_{j_1} (resp. P_{j_2}) containing v_1 (resp. v_2), only the intersections different from w_1 (resp. w_2) are counted.
4. Between v_1 and w_1 (resp. w_2 and v_2), it has a minimal number of intersections with the three arcs $J_{j_1}^{\gamma_{j_1, k}}$ ($k = 1, 2, 3$) in P_{j_1} containing v_1 (resp. $J_{j_2}^{\gamma_{j_2, k}}$ in P_{j_2} containing v_2).
5. It has a minimal number t_i of intersections with γ'_i inside $\overline{\gamma}_i$, for any $i = 1, \dots, 3g - 3$.
6. It has a minimal number s_i of intersections with γ''_i inside $\overline{\gamma}_i$, for any $i = 1, \dots, 3g - 3$.

The existence of a normal form is clear. The two forms of the path f are used to define two notions of complexity: its geodesic form is used to define its length, which can be seen as a geometric complexity, whereas its *normal coordinates* m_i, s_i and t_i can be seen as a combinatorial complexity. In the case of simple multicurves, these coordinates are known as Dehn-Thurston coordinates. Lemma 5.3 shows some equivalence between the two notions of complexity. We first show that a fixed set of coordinates corresponds to a finite number of possible non-homotopic paths.

Lemma 5.2. *For any set of coordinates $m_i, t_i, s_i, i = 1, \dots, 3g - 3$, there are at most $9(\max_{\{i=1, \dots, 3g-3\}}(m_i))^{2g-3}$ non-homotopic normal forms.*

Proof. Let f be a path, decomposed as above into (v_1, w_1) , $f^w = (w_1, w_2)$ and (w_2, v_2) . In each pair of pants not containing any endpoint v_1 or v_2 , fixing the m_i, s_i and t_i leads to a unique homotopy class of models [16, Lemma 5, p.63]. For the two (not necessarily different) pairs of pants P_{j_1} and P_{j_2} containing v_1 and v_2 , w_1 and w_2 are in fact fixing unique models (see Figure 13). There are three possible annulus boundaries $\partial\overline{\gamma_{j,i}}, i = 1, 2, 3$ for w_1 in the

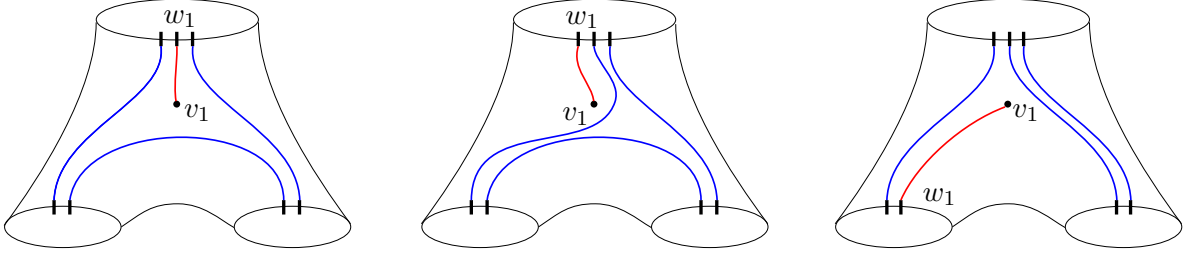


Figure 13: Three possible choices for w_1 . The two left choices correspond to the same model, but the orderings on the upper boundary lead to non-homotopic paths. The right choice leads to different models.

pair of pants P_j that contains v_1 (resp. $\overline{\gamma_{j,i}}$ for w_2), so, at most $3 \max_i(m_i)$ possibilities for each of them, and the result follows. \square

Lemma 5.3. *Let f be a geodesic segment of length l , then there exists a constant c_h such that the coordinates m_i, t_i , and $s_i, i = 1, \dots, 3g - 3$ of the normal form of f are smaller than $c_h \cdot l$.*

Proof. For any simple closed geodesic δ on S , the geodesic form of f intersects δ in a minimal number k_δ of points, since they are both geodesics. If ε_δ is the width of a tubular neighborhood of δ , then $l \geq \varepsilon_\delta(k_\delta - 1)$ [2, Lemma 3.1]. Each coordinate m_i, t_i , and s_i of f corresponds to the minimal number of intersections with a curve. The number m_i corresponds to γ_i . The number t_i is actually not larger than the number of intersections of f with the geodesic curve that is homotopic to γ'_i (γ'_i is generally not geodesic), and similarly s_i is not larger than the number of intersections of f with the geodesic homotopic to γ''_i . These curves $\gamma_i, \gamma'_i, \gamma''_i$ only depend on (S, h) , so, we can take ε_h to be the largest of all the $9g - 9$ widths $\varepsilon_{\gamma_i}, \varepsilon_{\gamma'_i}, \varepsilon_{\gamma''_i}$ and we obtain $l \geq \varepsilon_h \cdot \max(m_i, t_i, s_i)$ and thus $\max(m_i, t_i, s_i) \leq l/\varepsilon_h$. \square

Theorem 5.4. *For any hyperbolic structure h on S and any triangulation T of (S, h) , there is a sequence of flips of length at most $C_h \cdot \Lambda(T)^{6g-4} \cdot n^2$ in the geometric flip graph connecting T to a Delaunay triangulation of (S, h) .*

Proof. Let N_{v_1, v_2} be the number of segments from v_1 to v_2 shorter than $l = 2 \cdot \Lambda(T)$. From the previous lemma, we obtain that the $9g - 9$ coordinates m_i, t_i , and s_i of any such segment f are smaller than $c_h \cdot 2\Lambda(T)$. It appears that, $\forall i, m_i = t_i + s_i, t_i = m_i + s_i$ or $s_i = m_i + t_i$ [16, Lemma 6, p.64 & Fig.5, p.65]. So, if we fix m_i and t_i there are at most 3 possible s_i . Lemma 5.2 and 5.3 proves that there are $9(c_h \cdot 2\Lambda(T))^2$ potential segments for each coordinate set. We obtain a bound for N_{v_1, v_2} : $N_{v_1, v_2} \leq 9(c_h \cdot 2\Lambda(T))^2 \cdot 3(c_h \cdot 2\Lambda(T))^{6g-6}$ and thus, there is a

constant C'_h such that $N_{v_1, v_2} \leq C'_h \cdot \Lambda(T)^{6g-4}$. Since there are $1/2 \cdot n^2$ possible sets $\{v_1, v_2\}$, we obtain the bound on the number of edges. \square

References

- [1] M. Berger. *Geometry (vols. 1-2)*. Springer-Verlag, 1987.
- [2] Joan S Birman and Caroline Series. Geodesics with bounded intersection number on surfaces are sparsely distributed. *Topology*, 24(2):217–225, 1985.
- [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, Iordan Jordanov, and Monique Teillaud. 2D hyperbolic Delaunay triangulations. In *CGAL User and Reference Manual*. CGAL Editorial Board, 4.14 edition, 2019. URL: <https://doc.cgal.org/latest/Manual/packages.html#PkgHyperbolicTriangulation2>.
- [5] Mikhail Bogdanov, Monique Teillaud, and Gert Vegter. Delaunay triangulations on orientable surfaces of low genus. In *Proceedings of the 32nd Annual Symposium on Computational Geometry*, pages 20:1–20:17, 2016. doi:10.4230/LIPIcs.SoCG.2016.20.
- [6] Manuel Caroli and Monique Teillaud. Delaunay triangulations of closed Euclidean d-orbifolds. *Discrete & Computational Geometry*, 55(4):827–853, 2016. URL: <http://hal.inria.fr/hal-01294409>, doi:10.1007/s00454-016-9782-6.
- [7] Andrew J. Casson and Steven A. Bleiler. *Automorphisms of surfaces after Nielsen and Thurston*, volume 9 of *London Mathematical Society Student Texts*. Cambridge University Press, Cambridge, 1988. doi:10.1017/CB09780511623912.
- [8] Carmen Cortés, Clara I Grima, Ferran Hurtado, Alberto Márquez, Francisco Santos, and Jesus Valenzuela. Transforming triangulations on nonplanar surfaces. *SIAM Journal on Discrete Mathematics*, 24(3):821–840, 2010. doi:10.1137/070697987.
- [9] Vincent Despré, Loïc Dubois, Benedikt Kolbe, and Monique Teillaud. Experimental analysis of Delaunay flip algorithms on genus two hyperbolic surfaces (abstract). In *Abstracts 38th European Workshop on Computational Geometry*, pages 33:1–33:7, Perugia, Italy, 2022. Full version available as preprint <https://hal.inria.fr/hal-03462834>. URL: <https://hal.inria.fr/hal-03665888>.
- [10] 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. doi:10.4230/LIPIcs.SoCG.2023.27.

- [11] Vincent Despré, Jean-Marc Schlenker, and Monique Teillaud. Flipping geometric triangulations on hyperbolic surfaces. In *Proceedings of the 36th International Symposium on Computational Geometry*, pages 35:1–35:16, 2020. doi:10.4230/LIPIcs.SoCG.2020.35.
- [12] Loïc Dubois. A bound for Delaunay flip algorithms on flat tori. In *34th Canadian Conference on Computational Geometry*, pages 105–112, 2022. Best student paper. URL: <https://inria.hal.science/hal-03666488>.
- [13] H. Edelsbrunner and R. Seidel. Voronoi diagrams and arrangements. *Discrete & Computational Geometry*, 1(1):25–44, 1986. doi:10.1007/BF02187681.
- [14] 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>.
- [15] Albert Fathi, François Laudenbach, and Valentin Poénaru. *Thurston’s Work on Surfaces (MN-48)*. Princeton University Press, 2012.
- [16] Albert Fathi, François Laudenbach, Valentin Poénaru, et al. *Travaux de Thurston sur les surfaces*, volume 66–67 of *Astérisque*. Société Mathématique de France, Paris, 1979.
- [17] Martin Gardner. Chapter 19: Non-Euclidean geometry. In *The Last Recreations*. Springer, 1997.
- [18] F. Hurtado, M. Noy, and J. Urrutia. Flipping edges in triangulations. *Discrete & Computational Geometry*, 3(22):333–346, 1999. doi:s10.1007/PL00009464.
- [19] Y. Iwayoshi and M. Taniguchi. *An introduction to Teichmüller spaces*. Springer-Verlag, Tokyo, 1992. Translated and revised from the Japanese by the authors. doi:10.1007/978-4-431-68174-8.
- [20] C.L. Lawson. Software for C^1 Surface Interpolation. In *Mathematical Software*, pages 161–194. Academic Press, 1977. URL: <https://ntrs.nasa.gov/citations/19770025881>, doi:10.1016/b978-0-12-587260-7.50011-x.
- [21] Gregorii A Margulis. Applications of ergodic theory to the investigation of manifolds of negative curvature. *Functional analysis and its applications*, 3(4):335–336, 1969.
- [22] William S. Massey. *A basic course in algebraic topology*, volume 127 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1991.
- [23] Maryam Mirzakhani. Growth of the number of simple closed geodesics on hyperbolic surfaces. *Annals of Mathematics*, 168(1):97–125, 2008.
- [24] Bojan Mohar and Carsten Thomassen. *Graphs on Surfaces*. Johns Hopkins University Press, Baltimore, 2001.
- [25] Guillaume Tahar. Geometric triangulations and flips. *C. R. Acad. Sci. Paris, Ser. I*, 357:620–623, 2019. doi:10.1016/j.crma.2019.07.001.