

Minimal Delaunay triangulations of hyperbolic surfaces

Matthijs Ebbens¹, Hugo Parlier², and Gert Vegter¹

¹Bernoulli Institute for Mathematics, Computer Science and Artificial Intelligence, University of Groningen, Netherlands

²Department of Mathematics, University of Luxembourg, Luxembourg

Abstract

Motivated by recent work on Delaunay triangulations of hyperbolic surfaces, we consider the minimal number of vertices of such triangulations. First, we will show that every hyperbolic surface of genus g has a simplicial Delaunay triangulation with $O(g)$ vertices, where edges are given by distance paths. Then, we will construct a class of hyperbolic surfaces for which the order of this bound is optimal. Finally, to give a general lower bound, we will show that the $\Omega(\sqrt{g})$ lower bound for the number of vertices of a simplicial triangulation of a topological surface of genus g is tight for hyperbolic surfaces as well.

2012 ACM Subject Classification. Computational Geometry, Graphs and Surfaces.

Keywords and phrases. Delaunay triangulations, hyperbolic surfaces, metric graph embeddings, moduli spaces.

Funding. The second author was partially supported by ANR/FNR project SoS, INTER/ANR/16/11554412/SoS, ANR-17-CE40-0033.

1 Introduction

The classical topic of Delaunay triangulations has recently been studied in the context of hyperbolic surfaces. Bowyer’s incremental algorithm for computing simplicial Delaunay triangulations in the Euclidean plane [5] has been generalized to orientable hyperbolic surfaces and implemented for some specific cases [4, 11]. Moreover, it has been shown that the flip graph of geometric (but not necessarily simplicial) Delaunay triangulations on a hyperbolic surface is connected [7].

In this work, we consider the minimal number of vertices of a simplicial Delaunay triangulation of a closed hyperbolic surface of genus g . Motivated by the interest in embeddings where edges are shortest paths between their endpoints [8, 10], which have applications in for example the field of graph drawing [18], we restrict ourselves to *distance* Delaunay triangulations, where edges are distance paths.

Our main result is the upper bound on the number of vertices with sharp order of growth:

Theorem 1. *An orientable closed hyperbolic surface of genus $g \geq 2$ has a distance Delaunay triangulation with at most $O(g)$ vertices. Furthermore, there exists a family of surfaces, X_g , $g \geq 2$, such that the number of vertices of any distance Delaunay triangulation grows like $\Omega(g)$.*

The above result is a compilation of Theorems 4 and 19 where explicit upper and lower bounds are given.

Another reason to study triangulations whose edges are distance paths, comes from the study of moduli spaces \mathcal{M}_g , which we can think of as a space of all hyperbolic surfaces of genus $g \geq 2$ up to isometry. These spaces admit natural coordinates associated to pants decompositions (the so-called Fenchel-Nielsen coordinates, see Section 2 for details). It is a classical theorem of Bers [2] that any surface admits a short pants decomposition, meaning each of its simple closed geodesics is bounded by a function that only depends on the topology of the surface (but not its geometry). As these curves provide a local description of the surface, one might hope that they are also geodesically convex, meaning that the shortest distance path between any two points of a given curve is contained in the curve. It is perhaps surprising that *most* surfaces admit no such pants decompositions. Indeed it is known that *any* pants decomposition of a random surface (chosen with respect to a natural probability measure on \mathcal{M}_g) have at least one curve of length on the order of $g^{\frac{1}{6}-\varepsilon}$ as g grows (for any fixed $\varepsilon > 0$) [9]. And it is a theorem of Mirzakhani that these same random surfaces are also of *diameter* on the order of $\log(g)$ [13]. Hence the longest curve of any pants decomposition of a random surface is not convex.

The lengths of edges in a given triangulation are another parameter set for \mathcal{M}_g . By the theorem above, such a parameter set can be chosen with a reasonable number of vertices such that the edges are all convex. Using the moduli space point of view, one has a function $\omega : \mathcal{M}_g \rightarrow \mathbb{N}$ which associates to a surface the minimal number of vertices of any of its distance Delaunay triangulations. The above result implies that

$$\limsup_{g \rightarrow \infty} \max_{X \in \mathcal{M}_g} \frac{\omega(X)}{g}$$

is finite and strictly positive, but for instance we do not know whether the actual limit exists.

The examples we exhibit are geometrically quite simple, as they are made by gluing hyperbolic pants, with bounded cuff lengths, in something that resembles a line as the genus grows. One might wonder whether all surfaces have this property, but we show this is not the case by exploring the quantity $\min_{X \in \mathcal{M}_g} \omega(X)$. This quantity has a precise lower bound on the order of $\Theta(\sqrt{g})$ because we ask that our triangulations be simplicial [12]. We show how to use the celebrated Ringel-Youngs construction [16] to construct a family of hyperbolic surfaces that attain this bound for infinitely many genus (Theorem 27), showing that one cannot hope for better than the simplicial lower bound in general.

Although our results provide a good understanding on the extremal values of ω , there are still plenty of unexplored questions. For example, what is the behavior of ω for a random surface (using Mirzakhani's notion of randomness [13] alluded to above)?

This paper is structured as follows. In Section 2, we will introduce our notation and give some preliminaries on hyperbolic surface theory and triangulations. In Section 3, we will prove our linear upper bound for the number of vertices of a minimal distance Delaunay triangulation. In Section 4, we will construct classes of hyperbolic surfaces attaining the order of this linear upper bound. Finally, in Section 5, we will construct a family of hyperbolic surfaces attaining the general $\Theta(\sqrt{g})$ lower bound. Proofs of two technical lemmas appear in the appendices.

Acknowledgements. The authors warmly thank Monique Teillaud and Vincent Depré for fruitful discussions.

2 Preliminaries

We will start by recalling some hyperbolic geometry. There are several models for the hyperbolic plane [1]. In the Poincaré disk model, the hyperbolic plane is represented by the unit disk \mathbb{D} in the complex plane equipped with a specific Riemannian metric of constant Gaussian curvature -1 . With respect to this metric, hyperbolic lines, i.e., geodesics are given by diameters of \mathbb{D} or circle segments intersecting $\partial\mathbb{D}$ orthogonally. A hyperbolic circle is a Euclidean circle contained in \mathbb{D} . However, in general the centre and radius of a hyperbolic circle are different from the Euclidean centre and radius.

We refer to [6] for all of the following facts. A *hyperbolic surface* is a 2-dimensional Riemannian manifold that is locally isometric to an open subset of the hyperbolic plane [17], thus of constant curvature -1 . Our surfaces are assumed throughout to be *closed* and *orientable*, and because they are hyperbolic, via Gauss-Bonnet, their genus g satisfies $g \geq 2$ and their area is $4\pi(g - 1)$. Note that we will frequently be interested in subsurfaces of a closed surface which we think of as compact surfaces with boundary consisting of a collection of simple closed geodesics. The signature of such a subsurface is (g', k) where g' is its genus and k is the number of boundary geodesics.

Via the uniformization theorem, any hyperbolic surface X can be written as a quotient space $X = \mathbb{D}/\Gamma$ of the hyperbolic plane under the action of a Fuchsian group Γ (a discrete subgroup of the group of orientation-preserving isometries of \mathbb{D}). The hyperbolic plane \mathbb{D} is the *universal cover* of X and is equipped with a projection $\pi : \mathbb{D} \rightarrow \mathbb{D}/\Gamma$.

In the free homotopy class of any non-contractible closed curve lies a unique closed geodesic. If the curve is simple, then the corresponding geodesic is simple, and hence it is a straightforward topological exercise to decompose a hyperbolic surface into $2g - 2$ pairs of pants by cutting along $3g - 3$ disjoint simple closed geodesics. A pair of pants is a surface homeomorphic to a three times punctured sphere but we generally think of its closure, and thus of a hyperbolic pair of pants as being a surface of genus 0 with three simple closed geodesics as boundary. Its closure is thus a surface of signature $(0, 3)$. See Figure 1).

It is a short but useful exercise in hyperbolic trigonometry to show that a hyperbolic pair of pants is determined by its three boundary lengths. This is done by cutting the pair of pants along the three geodesic paths, orthogonal to the boundary, which realize the distance between the different boundary geodesics and then arguing on the resulting right angled hexagons. Hence, the lengths of the $3g - 3$ geodesics determine the geometry of each of the $2g - 2$ pairs of pants, but to determine X , one needs to add twist parameters that control how the pants are pasted together. How one computes the twist coordinate is at least partially a matter of taste, and although we will not make much use of it, for completeness we follow [6], where the twist is the signed distance between the points lying on the basepoints of the orthogeodesics mentioned above.

The length and twist parameters determine X and are called Fenchel-Nielsen coordinates. These parameters can be chosen freely in the set $(\mathbb{R}^{>0})^{3g-3} \times \mathbb{R}^{3g-3}$. What they determine is more than just an isometry class of surface: they determine a *marked* hyperbolic surface, homeomorphic to a base topological surface Σ . As the lengths and twists change, the marked surface changes, and the Fenchel-Nielsen coordinates provide a parameter set for the space of marked hyperbolic surfaces of genus g , called

Teichmüller space \mathcal{T}_g . The underlying *moduli space* \mathcal{M}_g can be thought of as the space of hyperbolic surfaces up to isometry, obtained from \mathcal{T}_g by “forgetting” the marking.

Throughout the paper, lengths of closed geodesics will play an important role. As mentioned above, in the free homotopy class of a non-contractible closed curve lies a unique geodesic representative, and as the metric changes, the free homotopy class is well defined, but the length of the geodesic changes. Generally we will be dealing with a fixed surface $X \in \mathcal{T}_g$, and the length of a geodesic γ will be denoted by $\ell(\gamma)$. Nonetheless, it is sometimes useful to think of the length of the corresponding homotopy class as a function over \mathcal{T}_g which associates to X the length of the geodesic corresponding to γ . To a pair of pants decomposition, we can associate a 3-regular graph. In this 3-regular

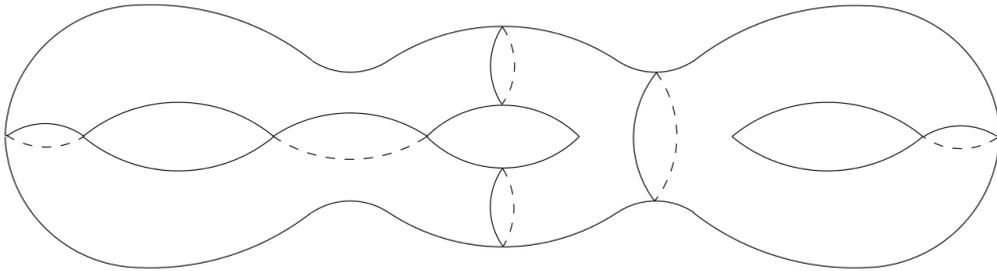


Figure 1: Decomposition of a genus 3 surface into 4 pair of pants using 6 disjoint simple closed geodesics.

graph, each pair of pants is represented by a vertex and two vertices share an edge if the corresponding pairs of pants share a boundary geodesic. For example, Figure 2 shows the 3-regular graph corresponding to the pair of pants decomposition in Figure 1. As our parametrization of \mathcal{T}_g depends on a choice of pants decomposition, one can think of the Fenchel-Nielsen coordinates associating a length and a twist to each edge.

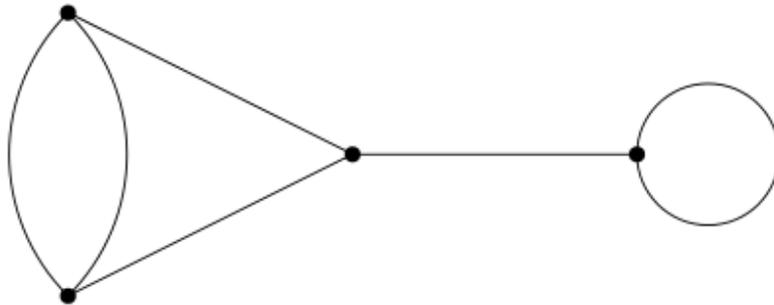


Figure 2: 3-regular graph corresponding to the pair of pants decomposition shown in Figure 1.

Around a simple closed geodesic γ , the local geometry of a surface is given by its so-called collar. Roughly speaking, for small enough r , the set

$$C_\gamma(r) = \{x \in X \mid d(x, \gamma) \leq r\}$$

is an embedded cylinder. A bound on how large one can take the r to be while retaining the cylinder topology is given by the Collar Lemma:

Lemma 2 ([6, Theorem 4.1.1]). *Let γ be a simple closed geodesic on a closed hyperbolic surface X . The collar $C_\gamma(w(\gamma))$ of width $w(\gamma)$ given by*

$$w(\gamma) = \operatorname{arcsinh} \left(\frac{1}{\sinh(\frac{1}{2}\ell(\gamma))} \right) \tag{1}$$

is an embedded hyperbolic cylinder isometric to $[-w(\gamma), w(\gamma)] \times \Gamma^1$ with the Riemannian metric $ds^2 = d\rho^2 + \ell^2(\gamma) \cosh^2(\rho) dt^2$ at (ρ, t) . Furthermore, if two simple closed geodesics γ and γ' are disjoint, then the collars $C_\gamma(w(\gamma))$ and $C_{\gamma'}(w(\gamma'))$ are disjoint as well.

This paper is about *distance Delaunay triangulations* on closed hyperbolic surfaces.

Definition 3. A *distance Delaunay triangulation* is a triangulation satisfying the following three properties:

1. it is a simplicial complex,
2. it is a Delaunay triangulation,
3. its edges are distance paths.

The set of all distance Delaunay triangulations of a closed hyperbolic surface X is denoted by $\mathcal{D}(X)$.

We will describe each of the three properties of distance Delaunay triangulations in more detail below.

Simplicial complexes. We will use the standard definition of a simplicial complex. In our case, an embedding of a graph into a surface is a simplicial complex if and only if it does not contain any 1- or 2-cycles. In particular, a geodesic triangulation of a point set in the Euclidean or hyperbolic planes is always a simplicial complex. This is because there are no geodesic monogons or bigons.

Delaunay triangulations. We recall that given a set of vertices in the Euclidean plane a triangle is called a *Delaunay triangle* if its circumscribed disk does not contain any vertex in its interior. A triangulation of a set of vertices in the Euclidean plane is a *Delaunay triangulation* if all triangles are Delaunay triangles.

We can define Delaunay triangulations in the hyperbolic plane in the same way as in the Euclidean plane, where we use the fact that hyperbolic circles are Euclidean circles. Delaunay triangulations on hyperbolic surfaces can be defined by lifting vertices on a hyperbolic surface X to the universal cover \mathbb{D} [3, 7]. More specifically, let \mathcal{P} be a set of vertices on X and let $\pi : \mathbb{D} \rightarrow \mathbb{D}/\Gamma$ be the projection of the hyperbolic plane \mathbb{D} to the hyperbolic surface $X = \mathbb{D}/\Gamma$. A triangle (v_1, v_2, v_3) with $v_i \in \mathcal{P}$ is called a *Delaunay triangle* if there exist pre-images $v'_i \in \pi^{-1}(\{v_i\})$ such that the circumscribed disk of the triangle (v'_1, v'_2, v'_3) in the hyperbolic plane does not contain any point of $\pi^{-1}(\mathcal{P})$ in its interior. A triangulation of \mathcal{P} on X is a *Delaunay triangulation* if all triangles are Delaunay triangles.

A Delaunay triangulation of a point set on a hyperbolic surface X is related to a Delaunay triangulation in \mathbb{D} as follows [3]. Given a point set \mathcal{P} on X , we consider a Delaunay triangulation T' of the infinite point set $\pi^{-1}(\mathcal{P})$. Then, we let $T = \pi(T')$. By definition, T is a Delaunay triangulation. Moreover, because every triangulation in \mathbb{D} is a simplicial complex, T' is a simplicial complex. However, T is not necessarily a simplicial complex, because projecting T' to X might introduce 1- or 2-cycles. We will use the correspondence between Delaunay triangulations in \mathbb{D} and in X in Definition 11 and the proof of Theorem 4 and show explicitly that in these cases the result after projecting to X is simplicial.

To make sure that $T = \pi(T')$ is a well-defined triangulation, we will assume without loss of generality that T' is Γ -invariant, i.e., the image of any Delaunay triangle in T'

under an element of Γ is a Delaunay triangle. Otherwise, it is possible that in so-called *degenerate cases* T contains edges that intersect in a point that is not a vertex [4]. Namely, suppose that T' contains a polygon $P = \{p_1, p_2, \dots, p_k\}$ consisting of $k \geq 4$ concircular vertices and let T_P be the Delaunay triangulation of P in T' . Because the Delaunay triangulation of a set of at least four concircular vertices is not uniquely defined, assume that there exists $A \in \Gamma$ such that the Delaunay triangulation $T_{A(P)}$ of $A(P)$ in T' is not equal to $A(T_P)$. Because $\pi(P) = \pi(A(P))$, there exists an edge of $\pi(T_{A(P)})$ and an edge of $\pi(A(T_P))$ that intersect in a point that is not a vertex.

Distance paths. Suppose we are given an edge (u, v) in a triangulation of a hyperbolic surface X . Because (u, v) is embedded in X , there exists a geodesic $\gamma : [0, 1] \rightarrow X$ with $\gamma(0) = u$ and $\gamma(1) = v$. We say that (u, v) is a *distance path* if $\ell(\gamma) = d(u, v)$, where $d(u, v)$ is the infimum of the lengths of all curves joining u to v .

3 Linear upper bound for the number of vertices of a minimal distance Delaunay triangulation

As our first result, we prove that for every hyperbolic surface there exists a distance Delaunay triangulation whose cardinality grows linearly as a function of the genus. Note that the constant 151 is certainly not optimal.

Theorem 4. *For every closed hyperbolic surface X of genus g there exists $T \in \mathcal{D}(X)$ with at most $151g$ vertices.*

The idea of the proof is the following. Given a hyperbolic surface X , we construct a vertex set \mathcal{P} on X consisting of at most $151g$ vertices such that the projection T of a Delaunay triangulation of $\pi^{-1}(\mathcal{P})$ in \mathbb{D} to X is a distance Delaunay triangulation of X .

It is known that T is a simplicial complex if \mathcal{P} is sufficiently dense and well-distributed [3]. More precisely, there are no 1- or 2-cycles in T if the diameter of the largest disk in \mathbb{D} not containing any points of $\pi^{-1}(\mathcal{P})$ is less than $\frac{1}{2} \text{sys}(X)$, where $\text{sys}(X)$ is the systole of X , i.e. the length of the shortest homotopically non-trivial closed curve. However, the systole of a hyperbolic surface can be arbitrarily close to zero, which means that we would need an arbitrarily dense set \mathcal{P} to satisfy this condition.

Instead, for a constant $\varepsilon > 0$ we subdivide X into its ε -*thick* part

$$X_{\text{thick}}^\varepsilon = \{x \in X \mid \text{inrad}(x) > \varepsilon\}$$

and its ε -*thin* part $X_{\text{thin}}^\varepsilon = X \setminus X_{\text{thick}}^\varepsilon$, where $\text{inrad}(x)$ is the injectivity radius at x , i.e., the radius of the largest embedded open disk centered at x . Note that the minimum of $\text{inrad}(x)$ over all $x \in X$ is given by $\frac{1}{2} \text{sys}(X)$. We see in Section 3.1 that $X_{\text{thin}}^\varepsilon$ is a collection of hyperbolic cylinders for sufficiently small ε (see Figure 3). In these hyperbolic cylinders we want to construct a set of vertices of which the cardinality does not depend on $\text{sys}(X)$. To do this, we put three vertices on the “waist” and each of the two boundary components of the cylinders that are “long and narrow”. In the cylinders that are not “long and narrow” it suffices to place three vertices on its waist only. The notions of “waist” and “long and narrow” will be specified in Section 3.1. Because $\text{inrad}(x) > \varepsilon$ for all $x \in X_{\text{thick}}^\varepsilon$, we can construct a sufficiently dense and well-distributed point set in $X_{\text{thick}}^\varepsilon$ whose cardinality does not depend on $\text{sys}(X)$ but only on ε . In Section 3.2 we will describe how we combine the vertices placed in the hyperbolic cylinders with the dense and well-distributed set of vertices in $X_{\text{thick}}^\varepsilon$. Finally, the proof of Theorem 4 is given in Section 3.3.

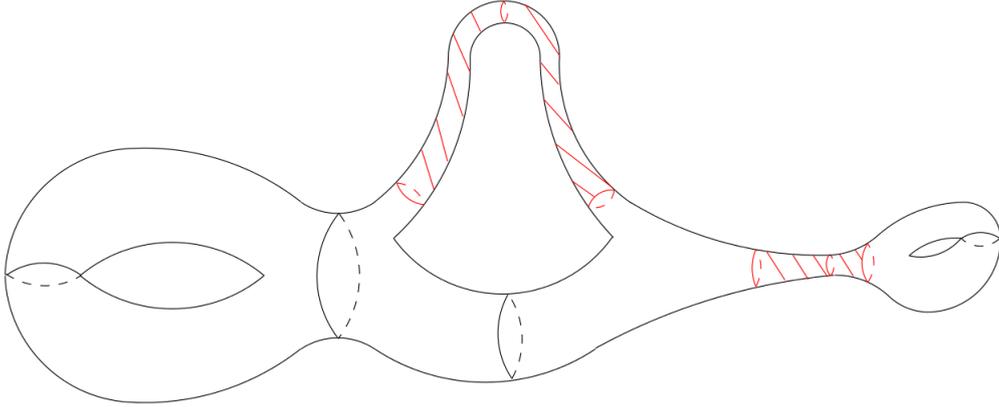


Figure 3: Decomposition of a hyperbolic surface into a thick part consisting of two connected components and two narrow hyperbolic cylinders (in red).

3.1 Distance Delaunay triangulations of hyperbolic cylinders

We now describe our construction of a set of vertices for the ε -thin part $X_{\text{thin}}^\varepsilon$. The following lemma describes $X_{\text{thin}}^\varepsilon$ in more detail.

Lemma 5 ([6, Theorem 4.1.6]). *If $\varepsilon < \operatorname{arcsinh}(1)$ then $X_{\text{thin}}^\varepsilon$ is a collection of at most $3g - 3$ pairwise disjoint hyperbolic cylinders.*

The following description of the geometry of the hyperbolic cylinders in $X_{\text{thin}}^\varepsilon$ is based primarily on a similar description in the context of colourings of hyperbolic surfaces [15]. Each hyperbolic cylinder C in $X_{\text{thin}}^\varepsilon$ consists of points with injectivity radius at most ε and the boundary curves γ^+ and γ^- consist of all points with injectivity radius equal to ε . Every point on the boundary curves is the base point of an embedded geodesic loop of length 2ε (Figure 4), which is completely contained in the hyperbolic cylinder. All points on the boundary curves have the same distance K_C to γ , which only depends on ε and the length $\ell(\gamma)$ of γ . To see this, fix a point p on γ^+ and consider a distance path ξ from p to γ (Figure 4). Cutting along γ, ξ and the loop of length 2ε with base point p yields a hyperbolic quadrilateral. The common orthogonal of γ and the geodesic loop subdivides this quadrilateral into two congruent quadrilaterals, each with three right angles. Applying a standard result from hyperbolic trigonometry yields [6, Formula Glossary 2.3.1(v)]

$$\sinh(\varepsilon) = \sinh\left(\frac{1}{2}\ell(\gamma)\right) \cosh(\ell(\xi)).$$

Because $K_C = \ell(\xi)$, it follows that

$$K_C = \operatorname{arccosh}\left(\frac{\sinh(\varepsilon)}{\sinh\left(\frac{1}{2}\ell(\gamma)\right)}\right). \quad (2)$$

We see that γ^+ consists of points that are equidistant to γ . By symmetry, the distance between a point on γ^- and γ is equal to K_C as well. Moreover, γ^+ and γ^- are smooth.

Recall the notion of a collar from Section 2. In particular, each hyperbolic cylinder C in $X_{\text{thin}}^\varepsilon$ is a collar of width K_C . Comparing equation (2) for K_C with equation (1) in the statement of the Collar Lemma, we see that $w(\gamma) > K_C$, because $\sinh \varepsilon < 1$. This inequality will be used in the proof of Lemma 7 to give a lower bound for the distance between distinct hyperbolic cylinders in $X_{\text{thin}}^\varepsilon$.

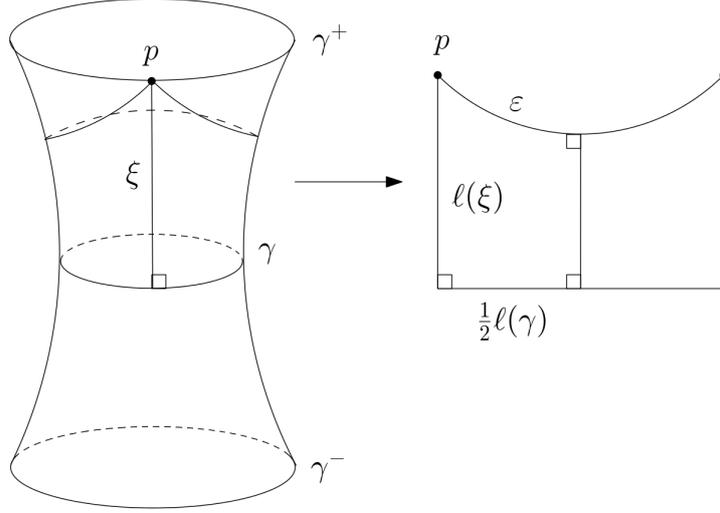


Figure 4: Computing K_C .

We distinguish between two kinds of hyperbolic cylinders in $X_{\text{thin}}^\varepsilon$, namely ε' -thin cylinders and ε' -thick cylinders, where $\varepsilon' = 0.99\varepsilon$. An ε' -thick cylinder with waist γ satisfies $2\varepsilon' \leq \ell(\gamma) \leq 2\varepsilon$, since γ is contained in the ε -thin part. An ε' -thin cylinder satisfies $\ell(\gamma) < 2\varepsilon'$.

Lemma 14 in Section 3.2 states that the triangulation depicted in Figure 5 is a Delaunay triangulation for ε' -thin cylinders. We call this triangulation a *standard triangulation* and describe it in more detail in the following definition. For ε' -thick cylinders we use a different construction defined in Definition 10.

Definition 6. Let X be a closed hyperbolic surface. Let C be an ε' -thin hyperbolic cylinder in $X_{\text{thin}}^\varepsilon$ with waist γ and boundary curves γ^+, γ^- . Place three equally-spaced points $x_i, i = 1, 2, 3$ on γ (see Figure 5). Then, place three points $x_i^+, i = 1, 2, 3$ on γ^+ and three points $x_i^-, i = 1, 2, 3$ on γ^- such that the projection of x_i^\pm on γ is equal to x_i for $i = 1, 2, 3$. Let V be the set consisting of x_i, x_i^- and x_i^+ for $i = 1, 2, 3$. Let E be the set of edges of one of the forms

$$(x_i^-, x_{i+1}^-), (x_i^-, x_i), (x_i^-, x_{i+1}^+), (x_i, x_{i+1}), (x_i, x_i^+), (x_i, x_{i+1}^+), (x_i^+, x_{i+1}^+)$$

for $i = 1, 2, 3$ (counting modulo 3), where the embedding of an edge in C is as shown in Figure 5. We call (V, E) a *standard triangulation* of C .

We not only have to prove that a standard triangulation of an ε' -thin cylinder is a Delaunay triangulation, we also have to show that its edges are distance paths. Corollary 9 states that all edges in a standard triangulation are distance paths if $\varepsilon \leq 0.72$. Before we can prove Corollary 9, we first need the following lemma.

Lemma 7. *Let X be a closed hyperbolic surface and let $\varepsilon \leq 0.72$. For each pair of distinct closed geodesics γ_1 and γ_2 in $X_{\text{thin}}^\varepsilon$ the collars $C_{\gamma_1}(K_{C_1} + \frac{1}{3}\varepsilon)$ and $C_{\gamma_2}(K_{C_2} + \frac{1}{3}\varepsilon)$ are embedded and disjoint.*

Remark 8. The value 0.72 was found experimentally and is optimal up to two decimal digits, i.e., the statement is not true for $\varepsilon = 0.73$. More specifically, if $\varepsilon \geq 0.73$ then there exists a closed hyperbolic surface X with disjoint closed geodesics γ_1 and γ_2 in $X_{\text{thin}}^\varepsilon$ such that $C_{\gamma_1}(K_{C_1} + \frac{1}{3}\varepsilon)$ and $C_{\gamma_2}(K_{C_2} + \frac{1}{3}\varepsilon)$ are not disjoint.

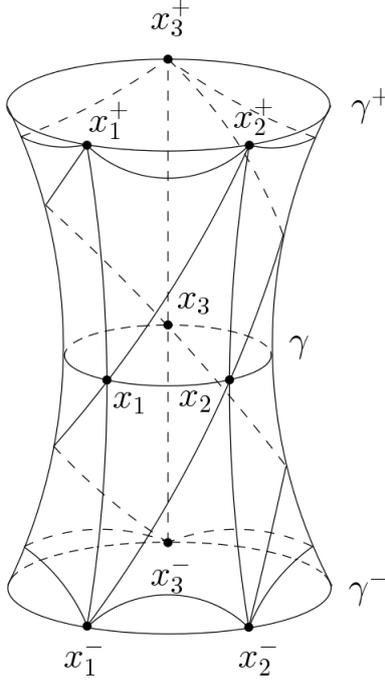


Figure 5: Standard triangulation of an ε' -thin cylinder.

Proof. See Figure 6. We will show that $w(\gamma_i) - K_{C_i} \geq \frac{1}{3}\varepsilon$ for $i = 1, 2$. Namely, this implies that $C_{\gamma_i}(K_{C_i} + \frac{1}{3}\varepsilon) \subseteq C_{\gamma_i}(w(\gamma_i))$. Because $C_{\gamma_1}(w(\gamma_1))$ and $C_{\gamma_2}(w(\gamma_2))$ are embedded and disjoint by the Collar Lemma, it follows that $C_{\gamma_1}(K_{C_1} + \frac{1}{3}\varepsilon)$ and $C_{\gamma_2}(K_{C_2} + \frac{1}{3}\varepsilon)$ are embedded and disjoint as well.

Comparing expression (2) for K_{C_i} and expression (1) for $w(\gamma_i)$, we see that $w(\gamma_i) - K_{C_i}$ is a positive number, with infimum when $\ell(\gamma_i) \rightarrow 0$ [15]. A straightforward computation shows that for $\varepsilon = 0.72$ this infimum is equal to $0.24\dots > \frac{1}{3}\varepsilon$. Since $w(\gamma_i) - K_{C_i}$ is decreasing as a function of ε , it follows that $w(\gamma_i) - K_{C_i} \geq \frac{1}{2}\varepsilon$ for all $\varepsilon \leq 0.72$. \square

Corollary 9. *Let X be a closed hyperbolic surface and let $\varepsilon \leq 0.72$. All edges in a standard triangulation of an ε' -thin cylinder in $X_{\text{thin}}^\varepsilon$ are distance paths.*

Proof. It is clear that edges of the form (x_i^-, x_{i+1}^-) , (x_i, x_{i+1}) , (x_i^+, x_{i+1}^+) for $i = 1, 2, 3$ are distance paths. Now, consider the edge of length K_C between x_i and x_i^+ . Because we know the metric of the cylinder, it can be shown explicitly that there are no shorter paths completely contained in the cylinder. Furthermore, because the collar $C_\gamma(K_C + \frac{1}{3}\varepsilon)$ is embedded by Lemma 7, any path that leaves the top half of the cylinder and returns through the bottom half has length at least $K_C + \frac{2}{3}\varepsilon$. It follows that the edges of the form (x_i, x_i^+) are distance paths. By symmetry, the edges of the form (x_i^-, x_i) are distance paths as well.

Finally, consider the edge between x_i and x_{i+1}^+ . Because $d(x_i, x_{i+1}) = \frac{1}{3}\ell(\gamma) < \frac{2}{3}\varepsilon$ and $d(x_{i+1}, x_{i+1}^+) = K_C$, we see from the triangle inequality that $d(x_i, x_{i+1}^+) < K_C + \frac{2}{3}\varepsilon$. Because any path that leaves the top half of the cylinder and returns through the bottom part of the cylinder has length at least $K_C + \frac{2}{3}\varepsilon$ by the same reasoning as above, it follows that edges of the form (x_i, x_{i+1}^+) are distance paths. By symmetry, edges of the form (x_i^-, x_{i+1}) are distance paths as well. \square

For ε' -thick cylinders, we see from Equation (2) for K_C that the width K_C is close

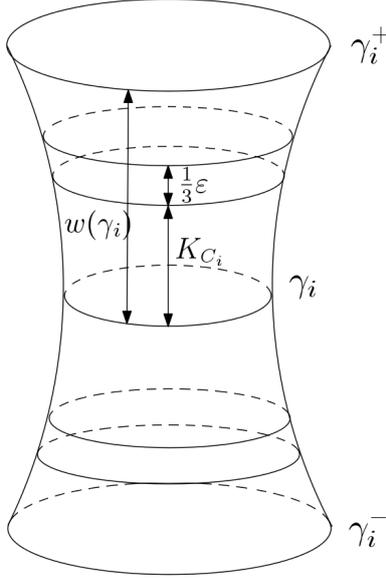


Figure 6: Illustration of the collars $C_{\gamma_i}(K_{C_i}) \subset C_{\gamma_i}(K_{C_i} + \frac{1}{4}\varepsilon) \subseteq C_{\gamma_i}(w(\gamma_i))$.

to zero. It turns out that we do not need to place three points on its waist and on each of its two boundary curves. Instead, three vertices on its waist suffice.

Definition 10. Let X be a closed hyperbolic surface. Let C be a ε' -thick hyperbolic cylinder in $X_{\text{thin}}^\varepsilon$ with waist γ . Place three equally-spaced points $x_i, i = 1, 2, 3$ on γ . Let $V = \{x_i \mid i = 1, 2, 3\}$ and $E = \{(x_1, x_2), (x_2, x_3), (x_3, x_1)\}$. We call (V, E) a *standard cycle* of C .

3.2 Constructing a distance Delaunay triangulation of X with few vertices

After constructing sets of vertices in the cylinders in the ε -thin part $X_{\text{thin}}^\varepsilon$, we construct a sufficiently dense and well-distributed set of vertices in the remainder of the surface. The following definition shows more precisely how we construct a set of vertices in $X_{\text{thick}}^\varepsilon$ and a corresponding Delaunay triangulation.

Definition 11. Set $\varepsilon = 0.72$ and $\varepsilon' = 0.99\varepsilon$. Let X be a closed hyperbolic surface. Let \mathcal{P}_1 be the set consisting of the vertices of a standard triangulation of every ε' -thin cylinder in $X_{\text{thin}}^\varepsilon$ together with the vertices of a standard cycle for every ε' -thick cylinder in $X_{\text{thin}}^\varepsilon$. Let T_j be the union of triangles in a standard triangulation (V_j, E_j) of an ε' -thin cylinder C_j . For every ε' -thick cylinder C_j , set $T_j = \emptyset$. Define \mathcal{P}_2 to be a maximal set in $X \setminus \cup_j T_j$ such that $d(p, q) \geq \frac{1}{2}\varepsilon$ for all distinct $p \in \mathcal{P}_1 \cup \mathcal{P}_2, q \in \mathcal{P}_2$. Denote the union $\mathcal{P}_1 \cup \mathcal{P}_2$ by \mathcal{P} and let T be the Delaunay triangulation of \mathcal{P} on X obtained after projecting a Delaunay triangulation of $\pi^{-1}(\mathcal{P})$ in \mathbb{D} to X . We call T a *thick-thin Delaunay triangulation* of X . The vertices in \mathcal{P}_1 and \mathcal{P}_2 are called the *cylinder vertices* and *non-cylinder vertices* of T , respectively.

Remark 12. Because by Corollary 9 all edges in a standard triangulation of any ε' -thin cylinder are distance paths if we choose $\varepsilon \leq 0.72$, we have chosen $\varepsilon = 0.72$ in Definition 11. Namely, we will see in the proof of Theorem 4 that the larger we choose ε , the smaller the constant (in our case 151) in the upper bound for the number of vertices. As in Section 3.1 we will fix $\varepsilon = 0.72$ and $\varepsilon' = 0.99\varepsilon$ throughout this subsection.

The edges between vertices on the same boundary curve of C_j are not equal to the boundary curves of C_j (because the latter are not geodesics), so T_j is strictly contained in C_j . We define \mathcal{P}_2 as a point set in $X \setminus \cup_j T_j$ instead of in $X \setminus \cup_j C_j$ to simplify our proof of Lemma 17, where we show that a thick-thin Delaunay triangulation of X is a simplicial complex.

The definition of \mathcal{P} does not explicitly forbid placing vertices of \mathcal{P}_2 in ε' -thick cylinders. However, we will see in the next lemma that there are no vertices of \mathcal{P}_2 in ε' -thick cylinders, because then they would be too close to the vertices of a standard cycle.

Lemma 13. *Let X be a closed hyperbolic surface and let T be a thick-thin Delaunay triangulation of X . Every vertex of T contained in an ε' -thick cylinder in $X_{\text{thin}}^\varepsilon$ is a cylinder vertex.*

Proof. Let \mathcal{P}_1 be the set of cylinder vertices and \mathcal{P}_2 the set of non-cylinder vertices. Let C be an arbitrary ε' -thick cylinder with waist γ and standard cycle (V, E) . We will show that the union U of the disks of radius $\frac{1}{2}\varepsilon$ centered at the vertices of V covers C completely. Namely, this implies that every point of C has distance at most $\frac{1}{2}\varepsilon$ to a vertex of V . Because $d(p, q) \geq \frac{1}{2}\varepsilon$ for all $p \in \mathcal{P}_1$ and $q \in \mathcal{P}_2$, it follows that there are no vertices of \mathcal{P}_2 contained in C .

To prove that U covers C completely, first observe that $d(x_i, x_{i+1}) = \frac{1}{3}\ell(\gamma) < \frac{2}{3}\varepsilon$ for all $i = 1, 2, 3$ (counting modulo 3). Therefore, the circles of radius $\frac{1}{2}\varepsilon$ centered at x_i and x_{i+1} intersect in two points, of which we call one p . Since the collar $C_\gamma(d(\gamma, p))$ is contained in U , it suffices to show that $K_C < d(\gamma, p)$, because then $C = C_\gamma(K_C) \subset C_\gamma(d(\gamma, p)) \subset U$. From equation (2) for K_C we know that

$$\cosh(K_C) = \frac{\sinh(\varepsilon)}{\sinh(\frac{1}{2}\ell(\gamma))} \leq \frac{\sinh(\varepsilon)}{\sinh(\varepsilon')} \leq 1.02,$$

where we substituted $\varepsilon' = 0.99\varepsilon$ and $\varepsilon = 0.72$ in the last step. On the other hand, the hyperbolic Pythagorean theorem yields

$$\cosh(d(\gamma, p)) = \frac{\cosh(\frac{1}{2}\varepsilon)}{\cosh(\frac{1}{6}\ell(\gamma))} \geq \frac{\cosh(\frac{1}{2}\varepsilon)}{\cosh(\frac{1}{3}\varepsilon)} \geq 1.03,$$

(see Figure 7) where again we substituted $\varepsilon = 0.72$ in the last step. We conclude that $K_C < d(\gamma, p)$, which finishes the proof.

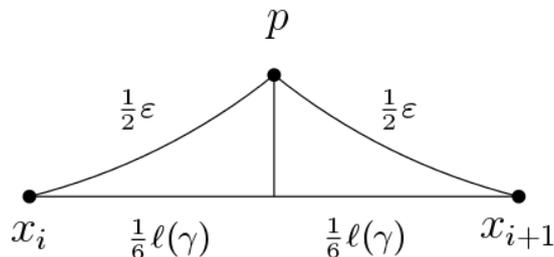


Figure 7: Computing $d(\gamma, p)$.

□

Even though the set of vertices of a thick-thin Delaunay triangulation of X contains the vertices of a standard triangulation (V_j, E_j) for every ε' -thin cylinder C_j , a priori it is not clear that the edges in E_j are edges in T as well. In the next lemma, we will show

that for every ε' -thin cylinder the triangles in a standard triangulation are Delaunay triangles with respect to the set of vertices of any thick-thin Delaunay triangulation of X . Namely, if this holds, then there exists a Delaunay triangulation of \mathcal{P} on X containing a standard triangulation of every ε' -thin cylinder in $X_{\text{thin}}^\varepsilon$.

Lemma 14. *Let X be a closed hyperbolic surface. Let T be a thick-thin Delaunay triangulation of X with vertex set \mathcal{P} and let C be an ε' -thin cylinder in $X_{\text{thin}}^\varepsilon$ with waist γ . Let (V, E) be a standard triangulation of C such that $V \subset \mathcal{P}$. Then all triangles of (V, E) are Delaunay triangles with respect to the point set \mathcal{P} .*

Remark 15. The proof of Lemma 14 is given in Appendix A. Even though we do not give the proof here, we note that in the proof it is shown as an intermediate step that $d(x_i^\pm, x_{i+1}^\pm) < \varepsilon$ for all $i = 1, 2, 3$. This inequality is used once more in the proof of Lemma 16.

Henceforth, we will assume that for each ε' -thin cylinder the vertices and edges of a standard triangulation are contained in a thick-thin Delaunay triangulation of X . To show that $T \in \mathcal{D}(X)$, we must show that T is a simplicial complex, i.e. it does not contain any 1- or 2-cycles, and that its edges are distance paths.

In the next lemma, we show that any edge that intersects $X_{\text{thick}}^\varepsilon$ has length smaller than ε . Moreover, we show that it follows that all edges that intersect $X_{\text{thick}}^\varepsilon$ are distance paths and that there are no 1- and 2-cycles consisting of edges intersecting $X_{\text{thick}}^\varepsilon$.

Lemma 16. *Let X be a closed hyperbolic surface and let T be a thick-thin Delaunay triangulation of X . Any edge of T that intersects $X_{\text{thick}}^\varepsilon$ has length smaller than ε and is a distance path. Moreover, there are no 1- or 2-cycles that intersect $X_{\text{thick}}^\varepsilon$ and consist of edges of length smaller than ε .*

Proof. Let (u, v) be an edge of T with non-empty intersection with $X_{\text{thick}}^\varepsilon$. Assume that (u, v) is contained in a triangle (u, v, w) in T with circumradius r and circumcenter c . We will first show that $\ell(u, v) < \varepsilon$. We consider two cases, depending on which set c is contained in. First, if $c \in T_j$ for some ε' -thin cylinder C_j , then at least one of u and v is contained in \mathcal{P}_1 . If both are contained in \mathcal{P}_1 , then (u, v) is contained in T_j , because the edges of a standard triangulation of an ε' -thin cylinder are distance paths by Corollary 9. This contradicts $(u, v) \cap X_{\text{thick}}^\varepsilon \neq \emptyset$, so we can assume that only one of u and v , say v , is contained in \mathcal{P}_1 . Then without loss of generality the situation is as depicted in Figure 8, where $\{v, w\} = \{x_i^+, x_{i+1}^+\}$.

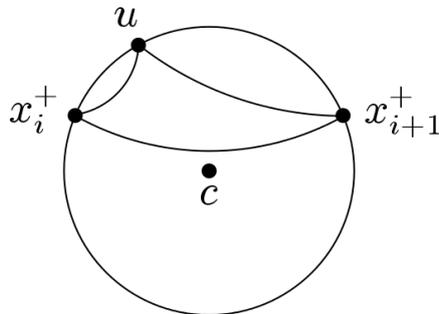


Figure 8: Circumscribed disk of a triangle (u, x_i^+, x_{i+1}^+) .

It follows that $\ell(u, v) < d(x_i^+, x_{i+1}^+)$. Because $d(x_i^+, x_{i+1}^+) < \varepsilon$ by Remark 15, it follows that $\ell(u, v) < \varepsilon$. Second, if $c \in X \setminus \cup_{j \in I} T_j$, then we can deduce that $r < \frac{1}{2}\varepsilon$.

Namely, if we suppose for a contradiction that $r \geq \frac{1}{2}\varepsilon$, then $d(c, p) \geq \frac{1}{2}\varepsilon$ for all $p \in \mathcal{P}$, because the circumcircle of (u, v, w) is empty. Then we could add c to \mathcal{P}_2 , which contradicts its maximality. We conclude that $r < \frac{1}{2}\varepsilon$. Because (u, v) is contained in a circle of radius $r < \frac{1}{2}\varepsilon$, it follows that $\ell(u, v) < \varepsilon$. Because $\ell(u, v) < \varepsilon$ in both cases, the first claim of the lemma follows.

To show that (u, v) is a shortest distant path between its endpoints, suppose for a contradiction that it is not. Then there exists a geodesic γ from u to v , such that $\ell(\gamma) < \ell(u, v)$. This means that $(u, v) \cup \gamma$ is a homotopically non-trivial closed curve of length smaller than $2\ell(e) < 2\varepsilon$. However, because $\text{inrad}(x) > \varepsilon$ for all $x \in X_{\text{thick}}^\varepsilon$, every homotopically non-trivial closed curve γ intersecting $X_{\text{thick}}^\varepsilon$ has length at least 2ε , which contradicts $\ell((u, v) \cup \gamma) < 2\varepsilon$. We conclude that (u, v) is a distance path between its endpoints.

A 1- or 2-cycle in G corresponds to a homotopically non-trivial closed curve on X [4]. By the same argument as before, the length of a 1- or 2-cycle σ intersecting $X_{\text{thick}}^\varepsilon$ is at least 2ε . Therefore, there are no 1- or 2-cycles that intersect $X_{\text{thick}}^\varepsilon$ and consist of edges of length smaller than ε . \square

Using the previous lemma, we show that a thick-thin Delaunay triangulation of X is a distance Delaunay triangulation.

Lemma 17. *Every thick-thin Delaunay triangulation of a closed hyperbolic surface is a distance Delaunay triangulation.*

Proof. Let X be a closed hyperbolic surface and let T be a thick-thin Delaunay triangulation of X . By definition, T is a Delaunay triangulation. We will show that T does not contain any 1- or 2-cycles to prove that it is a simplicial complex. We know from Lemma 16 that any edge (u, v) such that $(u, v) \cap X_{\text{thick}}^\varepsilon \neq \emptyset$ is not a 1-cycle. Because by construction there are no 1-cycles in a standard triangulation or standard cycle in $X_{\text{thin}}^\varepsilon$ as well, we conclude that T contains no 1-cycles.

To prove that T does not contain any 2-cycles, consider two distinct edges (u, v) and (v, w) of T with at least one shared endpoint. There are three cases, depending on whether two, one or zero of the edges (u, v) and (v, w) intersect $X_{\text{thick}}^\varepsilon$.

First, if (u, v) and (v, w) both intersect $X_{\text{thick}}^\varepsilon$, then they do not form a 2-cycle by Lemma 16.

Second, if precisely one of (u, v) and (v, w) , say (u, v) , intersects $X_{\text{thick}}^\varepsilon$, then $\ell(u, v) < \varepsilon$ and (v, w) is an edge contained in a hyperbolic cylinder of $X_{\text{thin}}^\varepsilon$. If (v, w) is an edge contained in an ε' -thick cylinder C with waist γ , then (v, w) is one of the edges of the standard cycle of C , because there are no other vertices in C by Lemma 13. Then $\ell(v, w) = \frac{1}{3}\ell(\gamma) < \frac{2}{3}\varepsilon$, so (u, v) and (v, w) do not form a 2-cycle by Lemma 16. Next, assume that (v, w) is an edge in an ε' -thin cylinder with waist γ . Then either w lies on γ and v lies on one of the boundary curves of C or v and w both lie on the same boundary curve of C . If w lies on γ and v on a boundary curve of C , then (u, v) and (v, w) do not form a 2-cycle, because u does not lie on γ . If v and w both lie on the same boundary geodesic, then $\ell(v, w) < \varepsilon$ by Remark 15, so (u, v) and (v, w) do not form a 2-cycle by Lemma 16.

Third, if neither (u, v) nor (v, w) intersects $X_{\text{thick}}^\varepsilon$, then (u, v) and (v, w) are both contained in a hyperbolic cylinder in $X_{\text{thin}}^\varepsilon$. They are contained in the same cylinder, because different cylinders are separated by $X_{\text{thick}}^\varepsilon$. Because by construction standard triangulations and standard cycles do not contain any 2-cycle, (u, v) and (v, w) do not form a 2-cycle.

This finishes the case analysis and we conclude that T is a simplicial complex.

To prove that all edges of T are distance paths, we know from Lemma 16 that any edge that intersects $X_{\text{thick}}^\varepsilon$ is a distance path. Because all edges in a standard triangulation are distance paths by Corollary 9 and because all edges in a standard cycle are distance paths by construction, we conclude that all edges in T are distance paths. \square

3.3 Proof of Theorem 4

Proof. (Theorem 4)

Let X be an arbitrary hyperbolic surface of genus g and let T be a thick-thin Delaunay triangulation of X . By definition, T is a Delaunay triangulation. By Lemma 17, T is a simplicial complex and all edges of T are distance paths. Hence, $T \in \mathcal{D}(X)$.

We will show here that the number of vertices of T is smaller than $151g$. By Lemma 5, $X_{\text{thin}}^\varepsilon$ consists of at most $3g - 3$ cylinders and each of these cylinders contains either 9 vertices (if it is ε' -thin) or 3 vertices (if it is ε' -thick). Therefore, $|\mathcal{P}_1| \leq 27g - 27$.

To find an upper bound for the cardinality of \mathcal{P}_2 , observe that for distinct $p, q \in \mathcal{P}_2$ the disks $B_p(\frac{1}{4}\varepsilon)$ and $B_q(\frac{1}{4}\varepsilon)$ of radius $\frac{1}{4}\varepsilon$ centered at p and q , respectively, are embedded and disjoint. Therefore, the cardinality of \mathcal{P}_2 is bounded above by the number of disjoint, embedded disks of radius $\frac{1}{4}\varepsilon$ that we can fit in X . Because the area of a hyperbolic disk of radius $\frac{1}{4}\varepsilon$ is $2\pi(\cosh(\frac{1}{4}\varepsilon) - 1)$ [1] and because the area of X is $4\pi(g - 1)$ [17], we obtain

$$|\mathcal{P}_2| \leq \frac{2(g - 1)}{\cosh(\frac{1}{4}\varepsilon) - 1}.$$

Therefore,

$$|\mathcal{P}| \leq 27g - 27 + \frac{2(g - 1)}{\cosh(\frac{1}{4}\varepsilon) - 1} \leq 151g.$$

This finishes the proof. \square

Remark 18. The constant 151 is not optimal. We can obtain the stronger upper bound $|\mathcal{P}| \leq 124g$ by looking more precisely at the upper bounds of $|\mathcal{P}_1|$ and $|\mathcal{P}_2|$ but because we are mainly interested in the the order of growth, we will not provide any details.

4 Classes of hyperbolic surfaces attaining the order of the upper bound

As our second result, we show that there exists a class of hyperbolic surfaces which attains the order of the upper bound presented in Theorem 4. We will first introduce this class of hyperbolic surfaces and then state the precise result in Theorem 19.

Recall from the Preliminaries that cutting a hyperbolic surface along $3g - 3$ disjoint simple closed geodesics decomposes the surface into $2g - 2$ pairs of pants and that each pair of pants decomposition has an associated 3-regular graph. Conversely, define L_g as the trivalent graph depicted in Figure 9. Here, every vertex v_i corresponds to a pair of pants Y_i . There is one edge from v_1 to itself and similarly from v_{2g-2} to itself. Moreover, for $1 \leq i \leq 2g - 3$ there is one edge between v_i and v_{i+1} if i is odd and there are two edges if i is even.

Now, fix some interval $[a, b] \subset \mathbb{R}$ with $0 < a < b$. Let $S_g(a, b)$ be the subset of \mathcal{T}_g with underlying graph L_g such that all length parameters are contained in $[a, b]$. In particular, $S_g(a, b)$ contains an open subset of \mathcal{T}_g , showing that having a linear number

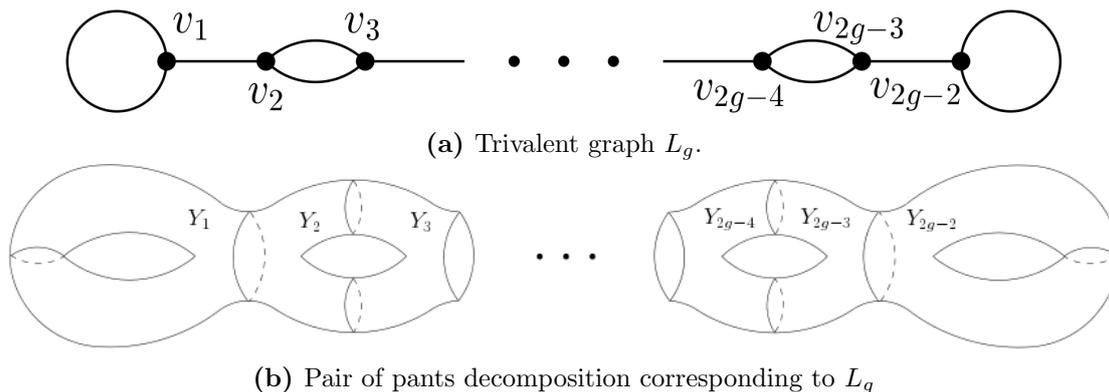


Figure 9: Trivalent graph L_g with corresponding pair of pants decomposition.

of vertices in terms of genus is relatively stable in this part of Teichmüller space. We will now state the result of this section.

Theorem 19. *There exists a constant $B > 0$ depending only on a, b such that a minimal distance Delaunay triangulation of any hyperbolic surface in $S_g(a, b)$ has at least Bg vertices.*

The idea of the proof is the following. Let a hyperbolic surface $X \in S_g(a, b)$ and $T \in \mathcal{D}(X)$ be given. Euler's formula implies $v - \frac{1}{3}e = 2 - 2g$ for triangulations of a surface of genus g , where v and e are the number of vertices and edges of the triangulation. We prove that $e \leq B'v$ for some constant $B' > 3$ only depending on a, b , which implies that

$$v \geq \frac{6g - 6}{B' - 3}.$$

This implies the result of Theorem 19. Hence, the argument consists mostly in finding an upper bound for the number of edges in terms of the number of vertices.

Before we continue with the proof of Theorem 19, we will look at our construction of $S_g(a, b)$ in more detail. By definition, every boundary geodesic of a pair of pants in the pair of pants decomposition of $X \in S_g(a, b)$ with respect to L_g has length in $[a, b]$. As explained in Section 2, the geometry of a pair of pants depends continuously on the lengths of its three boundary geodesics. In particular, the diameter $\text{diam}(Y)$ of a pair of pants Y as well as the minimal distance $\text{mindist}(Y)$ between any two of its boundary geodesics depend continuously on the lengths of its boundary geodesics. Because $[a, b]$ is a compact set, we obtain as an immediate consequence the following lemma.

Lemma 20. *There exist positive numbers $m(a, b)$ and $M(a, b)$ depending on a and b such that $m(a, b) \leq \text{mindist}(Y) < \text{diam}(Y) \leq M(a, b)$ for every pair of pants Y whose boundary geodesics have length in $[a, b]$.*

Remark 21. It is not too difficult to compute bounds for $\text{mindist}(Y)$ and $\text{diam}(Y)$ in terms of the lengths of the boundary geodesics of Y . This would give explicit expressions for $m(a, b)$ and $M(a, b)$ in terms of a and b . As we are only interested in the order of growth, to avoid further technical details, we do not provide details.

From now on, the numbers $m = m(a, b)$ and $M = M(a, b)$ will be fixed. Furthermore, the way in which we constructed the graph L_g allows us to speak of *consecutive* pairs of pants. A *cluster* in a hyperbolic surface X is a subset of X consisting of a number

of consecutive pairs of pants. Consider $T \in \mathcal{D}(X)$. An m -gap is a cluster consisting of m consecutive empty pairs of pants, where empty means that the pairs of pants do not contain any vertices of T . We say that an edge of T *crosses* a cluster if the pairs of pants containing its endpoints are separated by the cluster. Note that the cluster need not contain all pairs of pants which separate the two endpoints.

The next lemma states that if an edge of a distance Delaunay triangulation crosses many pairs of pants, then “many” of these pairs of pants are empty.

Lemma 22. *Let $X \in S_g(a, b)$ and define $N = N(a, b)$ as*

$$N(a, b) := \left\lceil \frac{M(a, b)}{m(a, b)} \right\rceil + 1.$$

Then, for every $T \in \mathcal{D}(X)$:

1. *If an edge of T crosses a cluster consisting of at least $3N$ pairs of pants, this cluster contains an N -gap.*
2. *If an edge of T crosses a cluster in which the first N and the last N pairs of pants are empty, then all pairs of pants in the cluster are empty.*

Proof. Let $T \in \mathcal{D}(X)$ and let (u, v) be an edge of T with $u \in Y_i$ and $v \in Y_j$. We will show that the cluster consisting of the union of all Y_k with $i + N + 1 \leq k \leq j - N - 1$ is empty. In other words, only the first N and last N pairs of pants are possibly non-empty. In particular, this implies that the two properties of the lemma hold.

Now, because (u, v) is a Delaunay edge, it is contained in some empty disk D passing through u and v . Consider Y_k with $i + N + 1 \leq k \leq j - N - 1$ and take $p \in Y_k$ arbitrarily (see Figure 10). We will show that the distance between p and the center c of D satisfies $\text{dist}(p, c) < \text{dist}(u, c)$. This implies that p is contained in the interior of D , so it cannot be a vertex of T . Therefore, Y_k is empty.

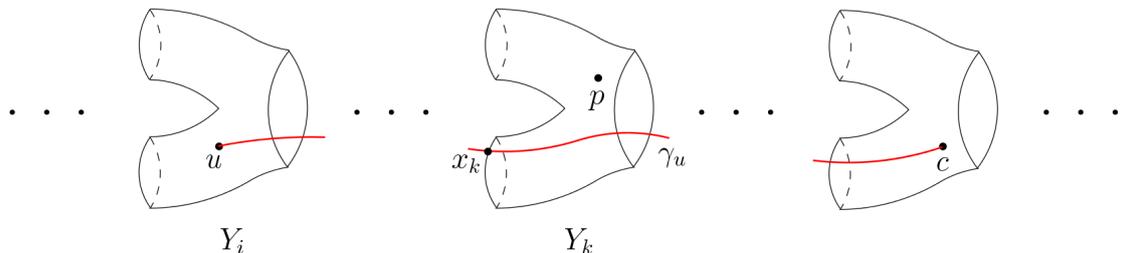


Figure 10: Depiction of the construction in the proof of Lemma 22.

Let γ_u be a distance path from c to u . First assume that $\gamma_u \cap Y_k \neq \emptyset$. Let x_k be the intersection of γ_u with one of the boundary geodesics of Y_k . By the triangle inequality and the definition of M , we know that

$$\text{dist}(c, p) \leq \text{dist}(c, x_k) + \text{dist}(x_k, p) \leq \text{dist}(c, x_k) + M.$$

To give an upper bound for $\text{dist}(c, x_k)$, observe that the part of γ_u from x_k to u passes through Y_{k-1}, \dots, Y_{i+1} before reaching $u \in Y_i$. By definition of m , the length of the part of γ_u within each of these $k-1-i$ pair of pants is at least m , so $\text{dist}(x_k, u) \geq (k-1-i)m$. This means that

$$\text{dist}(c, x_k) = \text{dist}(c, u) - \text{dist}(x_k, u) \leq \text{dist}(c, u) - (k-1-i)m.$$

It follows that

$$\text{dist}(c, p) \leq \text{dist}(c, u) + M - (k - 1 - i)m.$$

Because $k - 1 - i \geq N$ and $Nm > M$, we see that $M - (k - 1 - i)m < 0$, so

$$\text{dist}(c, p) < \text{dist}(c, u),$$

so $Y_k \in D$, hence Y_k does not contain any vertices of T . Note that we assumed that $\gamma_u \cap Y_k \neq \emptyset$. If $\gamma_u \cap Y_k = \emptyset$, then we consider a distance path γ_v from c to v instead of γ_u . The rest of the proof is analogous. We conclude that Y_k is empty for $i + N + 1 \leq k \leq j - N - 1$. This finishes the proof. \square

The following lemma states that we can construct a set of clusters which has as one of its properties that every edge of the distance Delaunay triangulation has its endpoints in the same cluster, or in two consecutive clusters.

Lemma 23. *Let $X \in S_g(a, b)$ be a hyperbolic surface and let $N = N(a, b)$ be as defined in Lemma 22. Let $T \in \mathcal{D}(X)$. There are interior-disjoint clusters with the following properties:*

1. *Each cluster consists of at most $6N$ consecutive pairs of pants;*
2. *Every cluster contains at least one vertex of \mathcal{T} , and every vertex of \mathcal{T} belongs to at least one cluster;*
3. *Every edge of \mathcal{T} has its endpoints in the same cluster, or in two consecutive clusters.*

Proof. A *wide gap* in the sequence of $2g - 2$ pairs of pants is a maximal subsequence consisting of at least N empty pairs of pants, where maximality is defined with respect to inclusion. The complement of the collection of wide gaps of X consists of a number of sequences of consecutive pairs of pants, that we call *superclusters* (see Figure 11). To obtain clusters that satisfy the properties of the lemma, each supercluster will be chopped up into one or more subsequences of length at most $6N$ in the following way:

- Each supercluster consisting of at most $3N$ pairs of pants is a cluster. Such a cluster is said to be a *short cluster*.
- Each supercluster consisting of more than $3N$ pairs of pants is chopped up into non-overlapping subsequences of length $3N$, followed by a subsequence of length between $3N$ and $6N - 1$.

More precisely, let m be the length of such a supercluster. Since $m > 3N$ there are integers k and r , with $0 \leq r < 3N$ and $k \geq 1$, such that $m = 3kN + r$. Let $m_1 = \dots = m_{k-1} = 3N$ and $m_k = 3N + r$, then $m_1 + \dots + m_k = m$ and $3N \leq m_i < 6N$. Therefore, the supercluster is the concatenation of k subsequences of length m_1, \dots, m_k .

This construction enforces Property 1.

Proof of Property 2. Each supercluster contains at least one vertex. Therefore, short clusters contain at least one vertex. The clusters obtained by chopping up a supercluster have length at least $3N$. Since a supercluster contains no N -gap, each of these clusters contains at least one vertex.

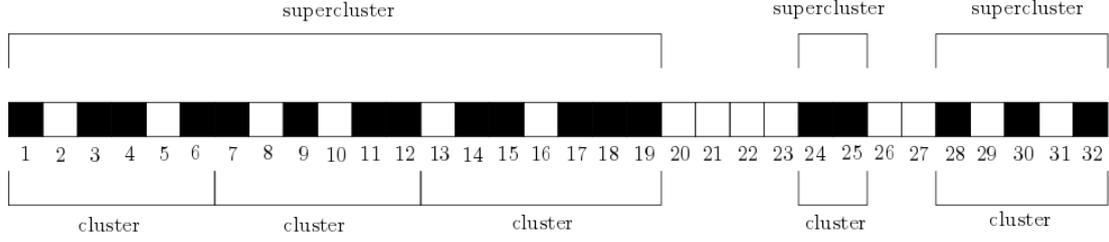


Figure 11: Illustration of the construction in the proof of Lemma 23. Each square represents a pair of pants in the pair of pants decomposition of a hyperbolic surface of genus 17. A pair of pants corresponding to a black square contains vertices of T , whereas a pair of pants corresponding to a white square does not contain vertices of T . We assume that $N = 2$. The superclusters and clusters that are defined during the construction are indicated.

Since all vertices belong to some supercluster, they belong to at least one cluster. This completes the proof of Property 2.

Proof of Property 3. Suppose the property does not hold. Then there is an edge of \mathcal{T} with vertices in non-adjacent clusters. In other words, this edge crosses an other cluster, say C . The construction of clusters implies that C does not contain an N -gap. By Part 1 of Corollary 22 the cluster C consists of less than $3N$ pairs of pants. Therefore, C is a short cluster. Since this short cluster is neither the first nor the last in the sequence of superclusters, it is preceded by a wide gap and succeeded by a wide gap. Since a wide gap contains an N -gap, Part 2 of Corollary 22 implies that cluster C is empty. This contradicts Property 2. Therefore, Property 3 holds. \square

In the following corollary, we denote the number of vertices of $T \in \mathcal{D}(X)$ contained in a subset U of X by $v(U)$. Likewise, let $e(U, W)$ be the number of edges of T with one endpoint in $U \subset X$ and one endpoint in $W \subset X$.

Corollary 24. *Let $X \in S_g(a, b)$ be a hyperbolic surface and let $T \in \mathcal{D}(X)$. Let $\{\Gamma_i \mid i = 1, \dots, n\}$ be a collection of clusters satisfying the properties of Lemma 23 for some $n \in \mathbb{N}$. If v and e are the number of vertices and edges of T , respectively, then*

$$\begin{aligned}
 n &\leq v, \\
 v &= \sum_{i=1}^n v(\Gamma_i), \\
 e &= \sum_{i=1}^n e(\Gamma_i, \Gamma_i) + \sum_{i=1}^{n-1} e(\Gamma_i, \Gamma_{i+1}).
 \end{aligned}$$

Proof. Because every cluster contains at least one vertex, the number of clusters is at most the number of vertices, which proves the first equation. The second equation follows from the property that every vertex is contained in a cluster. Because every edge has its endpoints in the same cluster, or in two consecutive clusters, the third equation holds. \square

Recall that we want to find a linear upper bound for the number of edges of a distance Delaunay triangulation in terms of the number of vertices. By Corollary 24, it suffices to find upper bounds for $e(\Gamma_i, \Gamma_i)$ and $e(\Gamma_i, \Gamma_{i+1})$ for clusters Γ_i satisfying the properties of Lemma 23. We will do this in the next lemma.

Lemma 25. *Let the notation be as in Corollary 24. Then, the following upper bounds hold:*

1. $e(\Gamma_i, \Gamma_i) \leq 3v(\Gamma_i) + 18N(N + 1)$ for all $i = 1, \dots, n$,
2. $e(\Gamma_i, \Gamma_{i+1}) \leq 18v(\Gamma_i \cup \Gamma_{i+1}) + 216N(N + 1)$ for all $i = 1, \dots, i - 1$.

The proof is given in Appendix B. We can now commence with the proof of Theorem 19.

Proof. (Theorem 19)

Take $X \in S_g(a, b)$ arbitrary and let $T \in \mathcal{D}(X)$ be arbitrary. Let $\{\Gamma_i \mid i = 1, \dots, n\}$ be a collection of clusters satisfying the properties of Lemma 23. By Corollary 24,

$$e = \sum_{i=1}^n e(\Gamma_i, \Gamma_i) + \sum_{i=1}^{n-1} e(\Gamma_i, \Gamma_{i+1}).$$

Substituting the upper bounds for $e(\Gamma_i, \Gamma_i)$ and $e(\Gamma_i, \Gamma_{i+1})$ from Lemma 25, we obtain

$$\begin{aligned} e &\leq \sum_{i=1}^n (3v(\Gamma_i) + 18N(N + 1)) + \sum_{i=1}^{n-1} (18v(\Gamma_i \cup \Gamma_{i+1}) + 216N(N + 1)), \\ &\leq 39 \sum_{i=1}^n (v(\Gamma_i) + 6N(N + 1)). \end{aligned}$$

From Corollary 24, we know that $\sum_{i=1}^n v(\Gamma_i) = v$ and $n \leq v$. Hence,

$$e \leq 39(1 + 6N(N + 1))v.$$

Euler's formula for triangulations $v - \frac{1}{3}e = 2 - 2g$ implies that

$$\begin{aligned} 39(1 + 6N(N + 1))v &\geq e = 3v + 6g - 6, \\ v &\geq \frac{g - 1}{6 + 39N(N + 1)}, \end{aligned}$$

which finishes the proof. □

5 Lower bound

In this section, we will look at a general lower bound for the minimal number of vertices of a distance Delaunay triangulation of a hyperbolic surface of genus g .

In the more general situation of a simplicial triangulation of a topological surface of genus g , one has an immediate lower bound on the minimal number of vertices. The fact that this lower bound is sharp is the following classical theorem of Jungerman and Ringel:

Theorem 26. *[12, Theorem 1.1] The minimal number of vertices of a simplicial triangulation of a topological surface of genus g is*

$$\left\lceil \frac{7 + \sqrt{1 + 48g}}{2} \right\rceil.$$

We show that the same result holds for the minimal number of vertices of a distance Delaunay triangulation of a hyperbolic surface of genus g for infinitely many values of g .

Theorem 27. *For any $g \geq 2$ of the form*

$$g = \frac{(n-3)(n-4)}{12}$$

for some $n \equiv 0 \pmod{12}$, the minimal number of vertices of a distance Delaunay triangulation of a hyperbolic surface of genus g is

$$n = \frac{7 + \sqrt{1 + 48g}}{2}.$$

Proof. Because every distance Delaunay triangulation of a hyperbolic surface is a simplicial triangulation of the corresponding topological surface, it follows from Theorem 26 that the minimal number of vertices is at least

$$\left\lceil \frac{7 + \sqrt{1 + 48g}}{2} \right\rceil.$$

In the remainder of the proof, we will construct for a given hyperbolic surface a distance Delaunay triangulation with the required number of vertices, inspired by a similar construction in the context of the chromatic number of hyperbolic surfaces [15].

Let $n \equiv 0 \pmod{12}$ and assume that $n \neq 0$. The complete graph K_n on n vertices can be embedded in a topological surface S_g of genus

$$g = \frac{(n-3)(n-4)}{12},$$

which is the smallest possible genus [16]. Because we have assumed that $n \equiv 0 \pmod{12}$, we know that the embedding of K_n into S_g is a triangulation T [19]. To turn T into a distance Delaunay triangulation, we will add a hyperbolic metric to the topological surface as follows. Every triangle in T is replaced by the unique equilateral hyperbolic triangle with all three angles equal to $\frac{2\pi}{n-1}$. In the complete graph K_n every vertex has $n-1$ neighbouring vertices. This means that in every vertex $n-1$ equilateral triangles meet, so the total angle at each vertex is 2π . Therefore, the result after replacing all triangles in T by hyperbolic triangles is a smooth hyperbolic surface Z_g .

It remains to be shown that $T \in \mathcal{D}(Z)$. By construction, T is a simplicial complex. It has also been shown that all edges are distance paths [15]. We will show here that T is a Delaunay triangulation of Z_g . Consider an arbitrary triangle (u, v, w) in T with circumcenter c and let $p \notin \{u, v, w\}$ be an arbitrary vertex of T (Figure 12). Consider a distance path γ from c to p . We can regard γ as the concatenation of simple segments that each pass through an individual triangle.

The first of these simple segments starts from c and leaves the triangle (u, v, w) , so its length is at least the distance between c and a side of (u, v, w) . Therefore, denoting by x the projection of c on one of the edges as shown in Figure 12, the length of the first segment is at least $d(c, x)$. The last of the simple segments passes through a triangle, say Δ , before arriving at p , so it has to pass through the side of Δ opposite to p . Therefore, its length is at least the distance between p and the opposite side of Δ . It is known that the distance between a vertex and the opposite side of an equilateral triangle is at least $\frac{1}{2}\ell$, where ℓ denotes the length of the sides of the equilateral triangle [15]. Hence,

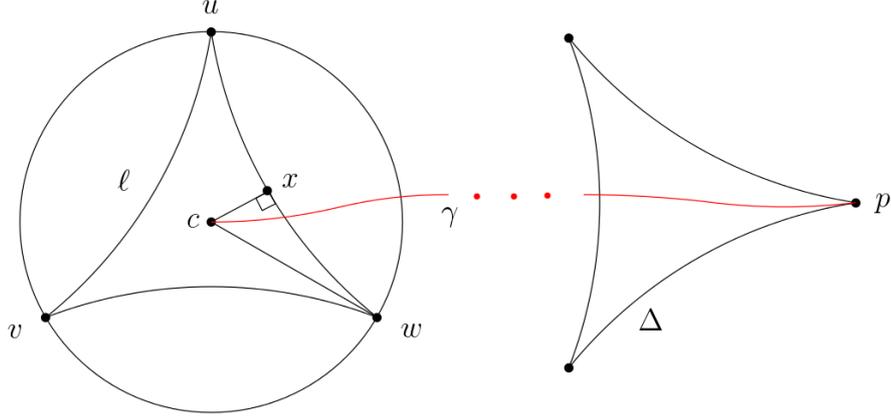


Figure 12: Schematic overview of the proof of T being a Delaunay triangulation.

$d(c, p) = \ell(\gamma) \geq d(c, x) + \frac{1}{2}\ell$. By the triangle inequality in triangle (c, w, x) we see that $d(c, w) \leq d(c, x) + d(x, w) = d(c, x) + \frac{1}{2}\ell$, so we conclude that $d(c, p) \geq d(c, w)$. This means that p is not contained in the interior of the circumcircle of (u, v, w) , which shows that (u, v, w) is a Delaunay triangle. By symmetry, it follows that all triangles are Delaunay triangles, which finishes the proof. \square

A Proof of Lemma 14

Proof. (**Lemma 14**)

To prove that the triangles of (V, E) are Delaunay triangles, we will show that every circumscribed disk does not contain any point of \mathcal{P} in its interior. By symmetry, it is sufficient to consider the top half of the cylinder. Let $i = 1, 2, 3$ be arbitrary and denote the disk passing through $x_i^+, x_{i+1}^+, x_i, x_{i+1}$ by D_i . That D_i does not contain any $p \in V$ in its interior is clear. The remainder of the proof consists of showing that p is not contained in the interior of D_i for all $p \in \mathcal{P} \setminus V$. Take $p \in \mathcal{P} \setminus V$ arbitrarily. Let c_i be the center of D_i . If $d(c_i, p) > d(c_i, x_i)$, then p is not contained in the interior of D_i .

Observe that $d(p, x_i^\pm) \geq \frac{1}{2}\varepsilon$ for $i = 1, 2, 3$. Namely, if $p \in \mathcal{P}_2$, where \mathcal{P}_2 is the subset of \mathcal{P} constructed in $X_{\text{thick}}^\varepsilon$, then by definition $d(p, x_i^\pm) \geq \frac{1}{2}\varepsilon$ for $i = 1, 2, 3$. On the other hand, if $p \in \mathcal{P}_1$, then p is a vertex in some hyperbolic cylinder $C' \neq C$ with waist γ' in $X_{\text{thin}}^\varepsilon$. By Lemma 7, the collars $C_\gamma(K_C + \frac{1}{3}\varepsilon)$ and $C_{\gamma'}(K_{C'} + \frac{1}{2}\varepsilon)$ are disjoint, so the distance between C and C' is at least $\frac{2}{3}\varepsilon$. Hence, $d(p, x_i) \geq \frac{1}{2}\varepsilon$ for $i = 1, 2, 3$.

Now, we claim that we can assume without loss of generality that $d(x_i^+, p) = d(x_{i+1}^+, p) = \frac{1}{2}\varepsilon$. See Figure 13. Consider the curve γ_c consisting of points of distance $d(c_i, \gamma)$ from γ and let p_c be the point of γ_c closest to p . Because $c_i \in \gamma_c$, we know that $d(p, c_i) \geq d(p, p_c)$. Now, note that a standard result from hyperbolic trigonometry in the quadrilateral (x_i, x_i^+, m_i^+, m_i) with three right angles [6, Formula Glossary 2.3.1(v)] states that

$$\sinh\left(\frac{1}{2}d(x_i^+, x_{i+1}^+)\right) = \sinh\left(\frac{1}{6}\ell(\gamma)\right) \cosh(K_C) = \frac{\sinh\left(\frac{1}{6}\ell(\gamma)\right) \sinh(\varepsilon)}{\sinh\left(\frac{1}{2}\ell(\gamma)\right)},$$

where the last equality follows from expression (2) for K_C . It can be deduced that $d(x_i^+, x_{i+1}^+) < \varepsilon$. Because $d(x, p) \geq \frac{1}{2}\varepsilon$ for all $x \in V$, the circles of radius $\frac{1}{2}\varepsilon$ centered at x_i^+ and x_{i+1}^+ intersect in two points. We see that $d(p, p_c)$ is minimized when $d(x_i^+, p) = d(x_{i+1}^+, p) = \frac{1}{2}\varepsilon$. Furthermore, if $d(x_i^+, p) = d(x_{i+1}^+, p) = \frac{1}{2}\varepsilon$, then p lies on the geodesic

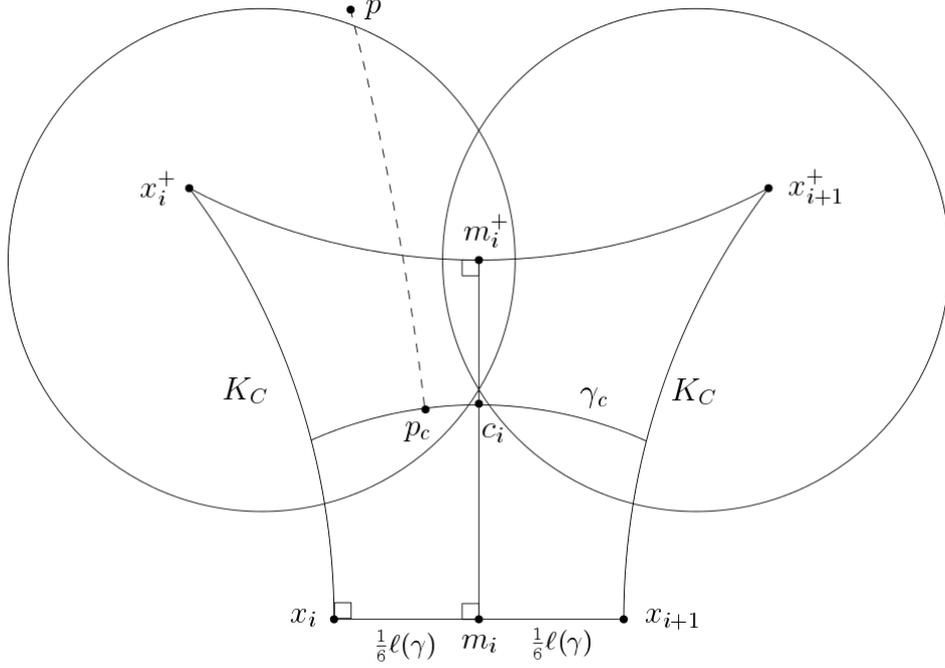


Figure 13: Construction to show that we can assume without loss of generality that $d(x_i^+, p) = d(x_{i+1}^+, p) = \frac{1}{2}\varepsilon$. 14.

passing through m_i and m_i^+ , so $d(p, p_c) = d(p, c_i)$, which means that $d(p, c)$ is minimized as well. We conclude that we can assume without loss of generality that $d(x_i^+, p) = d(x_{i+1}^+, p) = \frac{1}{2}\varepsilon$.

See Figure 14, where c'_i is the projection of c_i on (x_i, x_{i+1}^+) . To prove that

$$d(c_i, p) > d(c_i, x_i),$$

observe that $d(c_i, p) = d(m_i, m_i^+) - d(m_i, c_i) + d(m_i^+, p)$, where $d(m_i, m_i^+)$, $d(m_i, c_i)$ and $d(m_i^+, p)$ satisfy the equations

$$\coth(d(m_i, m_i^+)) = \frac{\cosh(\frac{1}{6}\ell(\gamma))}{\tanh(K_C)}, \quad (3)$$

$$\tanh(d(m_i, c_i)) = \frac{\cosh(\frac{1}{6}\ell(\gamma))}{\coth(\frac{1}{2}K_C)}, \quad (4)$$

$$\cosh(d(m_i^+, p)) = \frac{\cosh(\frac{1}{2}\varepsilon)}{\cosh(\frac{1}{2}d(x_i^+, x_{i+1}^+))}. \quad (5)$$

Here, equation (3) follows from applying a standard formula in hyperbolic trigonometry [6, Formula Glossary 2.3.1(iv)] in quadrilateral (x_i, x_i^+, m_i^+, m_i) . Equation (4) follows from applying the same formula in quadrilateral (x_i, c'_i, c_i, m_i) . Equation (5) follows from the hyperbolic Pythagorean theorem in triangle (x_i^+, p, m_i^+) . Moreover, applying the hyperbolic Pythagorean theorem in triangle (x_i, c_i, m_i) yields

$$\begin{aligned} \cosh(d(c_i, x_i)) &= \cosh(\frac{1}{6}\ell(\gamma)) \cosh(d(c, m_i)), \\ &= \cosh(\frac{1}{6}\ell(\gamma)) \cosh\left(\operatorname{arctanh}\left(\frac{\cosh(\frac{1}{6}\ell(\gamma))}{\coth(\frac{1}{2}K_C)}\right)\right), \end{aligned} \quad (6)$$

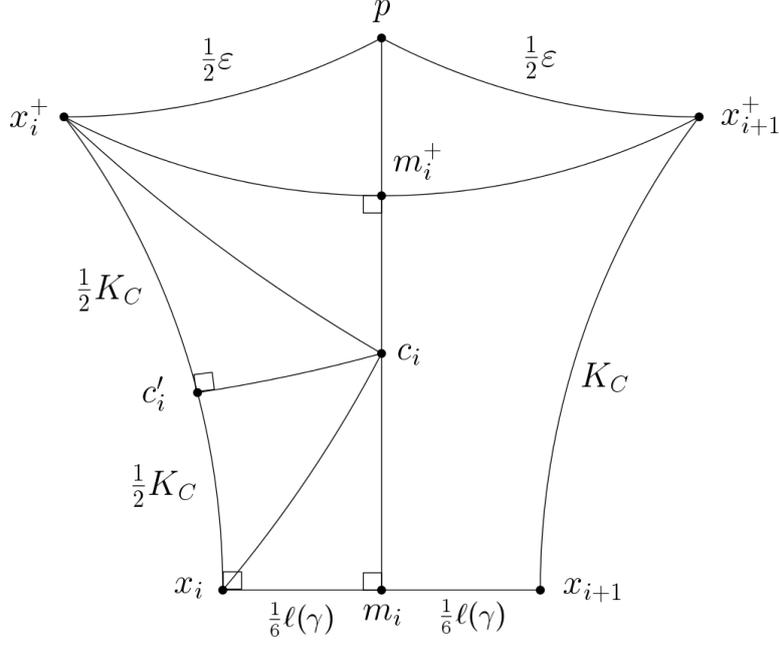


Figure 14: Schematic overview of the trigonometry in Lemma 14.

where we used equation (4) in the second line.

When we substitute the expressions for K_C and $d(x_i^+, x_{i+1}^+)$ into equations (3),(4),(5) and (6), we find expressions for $d(m_i, m_i^+)$, $d(m_i, c_i)$, $d(m_i^+, p)$ and $d(c_i, x_i)$ in terms of ε and $\ell(\gamma)$. As $\varepsilon = 0.72$ is fixed, we can treat these as functions of $\ell(\gamma)$. By a straightforward (but tedious) computation, it can be shown that $d(m_i, m_i^+) - d(m_i, c_i) - d(c_i, x_i)$ is strictly decreasing as a function of $\ell(\gamma)$ with minimum $-0.180\dots$ for $\ell(\gamma) = 2\varepsilon'$. By a similar computation, $d(m_i^+, p)$ is strictly increasing as a function of $\ell(\gamma)$ with minimum $0.247\dots$ for $\ell(\gamma) \rightarrow 0$. We conclude that

$$d(m_i, m_i^+) - d(m_i, c_i) - d(c_i, x_i) + d(m_i^+, p) \geq -0.180\dots + 0.247\dots > 0,$$

from which it follows that $d(c_i, p) = d(m_i, m_i^+) - d(m_i, c_i) + d(m_i^+, p) > d(c_i, p)$. Hence, p is not contained in D_i . This finishes the proof. \square

B Proof of Lemma 25

Proof. (Lemma 25)

Throughout the proof, we denote the set of vertices of T contained in a subset U of X by $V(U)$. Likewise, let $E(U, W)$ be the set of edges with one endpoint in $U \subset X$ and one endpoint in $W \subset X$.

Part 1. Consider the graph $G_i = (V(\Gamma_i), E(\Gamma_i, \Gamma_i))$. Let g_i be the genus of G_i , i.e., the minimal genus of a surface onto which G_i can be embedded. It is known that [14, Proposition 4.4.4]

$$g_i \geq \left\lceil \frac{e(\Gamma_i, \Gamma_i)}{6} - \frac{v(\Gamma_i)}{2} + 1 \right\rceil,$$

or, equivalently, that

$$e(\Gamma_i, \Gamma_i) \leq 6g_i + 3v(\Gamma_i) - 6. \quad (7)$$

We will show that the embedding of G_i into X intersects at most $6N(N+1)+2$ pairs of pants, which implies that $g_i \leq 3N(N+1)+1$. Certainly, $V(\Gamma_i)$ is contained in Γ_i , which consists of at most $6N$ consecutive pairs of pants. Now, let $e \in E(\Gamma_i, \Gamma_i)$. Because the diameter of Γ_i is at most $6NM$, we know that there exists a path of length at most $6NM$ between the endpoints of e . Because e is a distance path, it follows that

$$\ell(e) \leq 6NM < 6N^2m.$$

Suppose that e intersects exactly k pairs of pants that are not contained in Γ_i and denote the farthest pair of pants that it intersects by Y^* . Here, farthest is defined with respect to the distance along the trivalent graph L_g . Because e is an edge between vertices in Γ_i , e has to traverse at least $k-1$ pairs of pants to reach Y^* and similarly at least $k-1$ pairs of pants to return to Γ_i . As the length of e within each of these pairs of pants is at least m , we know

$$\ell(e) \geq 2(k-1)m.$$

It follows that

$$2(k-1)m < 6N^2m,$$

which implies that $k < 3N^2+1$. We conclude that G_i is embedded in a surface consisting of at most $(3N^2+1)+6N+(3N^2+1) = 6N(N+1)+2$ pairs of pants. It follows that $g_i \leq 3N(N+1)+1$. Hence

$$e(\Gamma_i, \Gamma_i) \leq 6(3N(N+1)+1) + 3v(\Gamma_i) - 6 = 3v(\Gamma_i) + 18N(N+1),$$

which finishes the proof.

Part 2. We consider two cases.

Case 1: there are at most $6N^2+2$ pairs of pants between Γ_i and Γ_{i+1} .

Consider the graph $(V(\Gamma_i \cup \Gamma_{i+1}), E(\Gamma_i \cup \Gamma_{i+1}, \Gamma_i \cup \Gamma_{i+1}))$. We have shown in Part 1 that edges in $E(\Gamma_i, \Gamma_i)$ and $E(\Gamma_{i+1}, \Gamma_{i+1})$ can traverse at most $3N^2-1$ pairs of pants that are not contained in Γ_i and Γ_{i+1} respectively. Therefore, $E(\Gamma_i, \Gamma_i) \cup E(\Gamma_{i+1}, \Gamma_{i+1})$ is contained in a surface consisting of at most

$$(3N^2+1) + 6N + (3N^2+1) + (3N^2+1) + 6N + (3N^2+1) = 12N^2 + 12N + 4$$

pairs of pants. With a similar argument it can be shown that $E(\Gamma_i, \Gamma_{i+1})$ is contained in this surface as well. Therefore, $(V(\Gamma_i \cup \Gamma_{i+1}), E(\Gamma_i \cup \Gamma_{i+1}, \Gamma_i \cup \Gamma_{i+1}))$ is embedded in a surface consisting of at most $12N^2+12N+4$ pairs of pants. Replacing Γ_i by $\Gamma_i \cup \Gamma_{i+1}$ in Inequality (7) yields

$$\begin{aligned} e(\Gamma_i \cup \Gamma_{i+1}, \Gamma_i \cup \Gamma_{i+1}) &\leq 6(6N^2+6N+4) + 3v(\Gamma_i \cup \Gamma_{i+1}) - 6, \\ &= 3v(\Gamma_i \cup \Gamma_{i+1}) + 36N(N+1) + 18. \end{aligned}$$

Because $e(\Gamma_i, \Gamma_{i+1}) \leq e(\Gamma_i \cup \Gamma_{i+1}, \Gamma_i \cup \Gamma_{i+1})$, the desired inequality follows.

Case 2: there are more than $6N^2+2$ pairs of pants between Γ_i and Γ_{i+1} .

We show that there are integers $g_{i,1}$ and $g_{i,2}$ with

$$g_{i,1} \leq g_{i,2}, \tag{8}$$

such that

$$e(\Gamma_i, \Gamma_{i+1}) \leq 2v(\Gamma_i) + 2v(\Gamma_{i+1}) + 4g_{i,1} - 4 \quad (9)$$

and

$$e(\Gamma_i, \Gamma_{i+1}) \geq 3v(\Gamma_i) + 3v(\Gamma_{i+1}) - 3e(\Gamma_i, \Gamma_i) - 3e(\Gamma_{i+1}, \Gamma_{i+1}) + 6g_{i,2} - 6. \quad (10)$$

Combining (9) and (10) yields

$$e(\Gamma_i, \Gamma_{i+1}) \leq 6e(\Gamma_i, \Gamma_i) + 6e(\Gamma_{i+1}, \Gamma_{i+1}) + 12g_{i,1} - 12g_{i,2}.$$

Using the upper bound $e(\Gamma_j, \Gamma_j) \leq 3v(\Gamma_j) + 18N(N+1)$ from Part 1 for $j = i$ and $j = i+1$, together with (8) yields

$$e(\Gamma_i, \Gamma_{i+1}) \leq 18v(\Gamma_i \cup \Gamma_{i+1}) + 216N(N+1),$$

which is the desired inequality.

The number $g_{i,1}$ is the genus of the graph $G_{i,1}$ with edge set $E(\Gamma_i, \Gamma_{i+1})$ and vertex set $V_{i,1}$ consisting of all vertices incident to some edge in $E(\Gamma_i, \Gamma_{i+1})$. The number $g_{i,2}$ is the genus of the graph $G_{i,2}$ with edge set $E(\Gamma_i \cup \Gamma_{i+1}, \Gamma_i \cup \Gamma_{i+1})$ and vertex set $V(\Gamma_{i+1} \cup \Gamma_{i+1})$. Since $G_{i,1}$ is a subgraph of $G_{i,2}$, inequality (8) holds. Therefore, it remains to be proved that (9) holds for this value of $g_{i,1}$, and that (10) holds for this value of $g_{i,2}$.

To prove that (9) holds, we apply a result in graph theory about bipartite graphs to $G_{i,1}$. By construction, $G_{i,1}$ contains no cycle of length 3. We claim that $G_{i,1}$ is connected. Then, [14, Prop. 4.4.4, eq. 4.13]

$$e(\Gamma_i, \Gamma_{i+1}) \leq 2|V_{i,1}| + 4g_{i,1} - 4.$$

Observing that $|V_{i,1}| \leq v(\Gamma_i) + v(\Gamma_{i+1})$ yields the desired inequality.

To prove that $G_{i,1}$ is connected, consider a pair of pants Y between Γ_i and Γ_{i+1} such that there are at least $3N^2 + 1$ pairs of pants between Y and Γ_i and between Y and Γ_{i+1} . Such a pair of pants Y exists because there are more than $6N^2 + 2$ pairs of pants between Γ_i and Γ_{i+1} . Let γ be the boundary geodesic of Y such that $X \setminus \gamma$ consists of exactly two connected components. Every edge in $E(\Gamma_i, \Gamma_{i+1})$ intersects γ , because γ separates Γ_i and Γ_{i+1} . Furthermore, every edge in $E(\Gamma_i, \Gamma_{i+1})$ intersects γ *exactly once*, because the edges in $E(\Gamma_i, \Gamma_{i+1})$ and γ are geodesics and there are no hyperbolic bigons. No edge in $E(\Gamma_i, \Gamma_i) \cup E(\Gamma_{i+1}, \Gamma_{i+1})$ intersects γ , because by the reasoning in Part 1 edges in $E(\Gamma_j, \Gamma_j)$ for $j = i$ and $j = i+1$ intersect fewer than $3N^2 + 1$ pairs of pants that are not contained in Γ_j . Because T has no edges between non-adjacent clusters, it follows that the only edges of T that intersect γ are edges in $E(\Gamma_i, \Gamma_{i+1})$ and that each edge intersects γ in exactly one point. Therefore, we can write $E(\Gamma_i, \Gamma_{i+1}) = \{e_1, \dots, e_k\}$ for $k = |E(\Gamma_i, \Gamma_{i+1})|$, where the indices of the edges correspond to the order in which they intersect γ . For every $i = 1, \dots, k$, the edges e_i and e_{i+1} are contained in a triangle of T consisting of e_i, e_{i+1} and an edge in either $E(\Gamma_i, \Gamma_i)$ or $E(\Gamma_{i+1}, \Gamma_{i+1})$. Hence, e_i and e_{i+1} share an endpoint. Therefore, there is a path in $G_{i,1}$ between any endpoint of e_i and any endpoint of e_{i+1} . This implies that there is a path in $G_{i,1}$ between any endpoint of e_i and e_j for all $i, j = 1, \dots, k$. Because $V_{i,1}$ consists of the vertices incident to some edge in $E(\Gamma_i, \Gamma_{i+1})$, it follows that $G_{i,1}$ is connected. This concludes the proof of Inequality (9).

We continue with the proof of Inequality (10). Recall that the graph $G_{i,2}$ has edge set $E(\Gamma_i \cup \Gamma_{i+1}, \Gamma_i \cup \Gamma_{i+1})$ and vertex set $V(\Gamma_{i+1} \cup \Gamma_{i+1})$. The union of the triangles of T

with edges and vertices in $G_{i,2}$ define a topological surface with boundary components¹. Each boundary component consists of a finite number of edges of $G_{i,2}$. By adding a face to each boundary component, we obtain an embedding of $G_{i,2}$ in a closed surface S_i . When we speak of faces of $G_{i,2}$, we will always refer to faces with respect to the embedding of $G_{i,2}$ in S_i .

Denote the number of edges that are contained in exactly zero, one or two triangles in $G_{i,2}$ by δ_0 , δ_1 and δ_2 , respectively. The total number of edges is given by $e_{i,2} = \delta_0 + \delta_1 + \delta_2$. Now, let f_Δ be the number of triangular faces of $G_{i,2}$. As the total number of faces $f_{i,2}$ is at least the number of triangular faces, we know that $f_{i,2} \geq f_\Delta$. Since $3f_\Delta = 2\delta_2 + \delta_1$, it follows that

$$3f_{i,2} \geq 2\delta_2 + \delta_1.$$

As $e_{i,2} = \delta_0 + \delta_1 + \delta_2$, we obtain

$$3f_{i,2} \geq 2(e_{i,2} - \delta_1 - \delta_0) + \delta_1 = 2e_{i,2} - \delta_1 - 2\delta_0.$$

Because

$$e_{i,2} = e(\Gamma_i, \Gamma_i) + e(\Gamma_i, \Gamma_{i+1}) + e(\Gamma_{i+1}, \Gamma_{i+1}), \quad (11)$$

we see

$$3f_{i,2} \geq 2e(\Gamma_i, \Gamma_i) + 2e(\Gamma_i, \Gamma_{i+1}) + 2e(\Gamma_{i+1}, \Gamma_{i+1}) - \delta_1 - 2\delta_0. \quad (12)$$

We will now bound the right-hand side from below by proving that

$$\delta_1 + 2\delta_0 \leq 2e(\Gamma_i, \Gamma_i) + 2e(\Gamma_{i+1}, \Gamma_{i+1}).$$

We claim that every edge in $E(\Gamma_i, \Gamma_{i+1})$ is part of two triangles. To prove this, let $e = (u, v) \in E(\Gamma_i, \Gamma_{i+1})$ and consider a triangle (u, v, w) in T containing e . Because every edge has its endpoints in either the same cluster or consecutive clusters, there is only an edge between u and w if $w \in \Gamma_j$ for $j = i - 1, i, i + 1$ and there is only an edge between v and w if $w \in \Gamma_j$ for $j = i, i + 1, i + 2$. Since (u, v, w) is a triangle in T , it follows that $w \in \Gamma_i \cup \Gamma_{i+1}$, so $w \in G_{i,2}$. This means that (u, v, w) is a triangle in $G_{i,2}$ and since we have chosen it arbitrarily, both triangles in T containing e are triangles in $G_{i,2}$. It follows that the edges that are contained in no triangles or exactly one triangle in $G_{i,2}$ are contained in $E(\Gamma_i, \Gamma_i) \cup E(\Gamma_{i+1}, \Gamma_{i+1})$. This means that

$$\delta_0 + \delta_1 \leq e(\Gamma_i, \Gamma_i) + e(\Gamma_{i+1}, \Gamma_{i+1}),$$

so

$$2\delta_0 + \delta_1 \leq 2e(\Gamma_i, \Gamma_i) + 2e(\Gamma_{i+1}, \Gamma_{i+1}).$$

Combining this upper bound with Equation (12) we obtain

$$3f_{i,2} \geq 2e(\Gamma_i, \Gamma_{i+1}). \quad (13)$$

To conclude, we will look at Euler's formula for the graph $G_{i,2}$, which is given by

$$v_{i,2} - e_{i,2} + f_{i,2} = 2 - 2g'_{i,2}, \quad (14)$$

where $g'_{i,2}$ is the genus of the embedding of $G_{i,2}$ in S_i . Because $g_{i,2}$ is the minimal genus of a surface onto which $G_{i,2}$ can be embedded, in particular $g_{i,2} \leq g'_{i,2}$. Substituting Equation (11) and Inequality (13) into Euler's formula (14), we obtain after some simplifications

$$e(\Gamma_i, \Gamma_{i+1}) \geq 3v(\Gamma_i \cup \Gamma_{i+1}) - 3e(\Gamma_i, \Gamma_i) - 3e(\Gamma_{i+1}, \Gamma_{i+1}) + 6g_{i,2} - 6.$$

This finishes the proof. □

¹In fact, these triangles define a hyperbolic surface with boundary components consisting of a finite number of geodesic segments. However, for our argument we do not use any metric properties.

References

- [1] Alan F. Beardon. *The geometry of discrete groups*, volume 91 of *Graduate Texts in Mathematics*. Springer-Verlag, 2012.
- [2] Lipman Bers. An inequality for Riemann surfaces. In *Differential geometry and complex analysis*, pages 87–93. Springer, Berlin, 1985.
- [3] Mikhail Bogdanov and Monique Teillaud. Delaunay triangulations and cycles on closed hyperbolic surfaces. Technical Report RR-8434, INRIA, 2013.
- [4] Mikhail Bogdanov, Monique Teillaud, and Gert Vegter. Delaunay triangulations on orientable surfaces of low genus. In Leibniz International Proceedings in Informatics, editor, *32nd International Symposium on Computational Geometry (SoCG 2016)*, pages 20:1–20:17, 2016.
- [5] Adrian Bowyer. Computing Dirichlet tessellations. *The Computer Journal*, 24(2):162–166, 1981.
- [6] Peter Buser. *Geometry and spectra of compact Riemann surfaces*. Springer-Verlag, 2010.
- [7] Vincent Despré, Jean-Marc Schlenker, and Monique Teillaud. Flipping geometric triangulations on hyperbolic surfaces. *arXiv preprint arXiv:1912.04640*, 2019.
- [8] István Fáry. On straight-line representation of planar graphs. *Acta scientiarum mathematicarum*, 11(229-233):2, 1948.
- [9] Larry Guth, Hugo Parlier, and Robert Young. Pants decompositions of random surfaces. *Geom. Funct. Anal.*, 21(5):1069–1090, 2011.
- [10] Alfredo Hubard, Vojtěch Kaluža, Arnaud De Mesmay, and Martin Tancer. Shortest path embeddings of graphs on surfaces. *Discrete & Computational Geometry*, 58(4):921–945, 2017.
- [11] Jordan Iordanov and Monique Teillaud. Implementing Delaunay triangulations of the Bolza surface. In *Proceedings of the Thirty-third International Symposium on Computational Geometry*, pages 44:1–44:15, 2017.
- [12] Mark Jungerman and Gerhard Ringel. Minimal triangulations on orientable surfaces. *Acta Mathematica*, 145(1):121–154, 1980.
- [13] Maryam Mirzakhani. Growth of Weil-Petersson volumes and random hyperbolic surfaces of large genus. *J. Differential Geom.*, 94(2):267–300, 2013.
- [14] Bojan Mohar and Carsten Thomassen. *Graphs on surfaces*, volume 10. JHU Press, 2001.
- [15] Hugo Parlier and Camille Petit. Chromatic numbers of hyperbolic surfaces. *Indiana University Mathematics Journal*, pages 1401–1423, 2016.
- [16] Gerhard Ringel and John W.T. Youngs. Solution of the Heawood map-coloring problem. *Proceedings of the National Academy of Sciences*, 60(2):438–445, 1968.
- [17] John Stillwell. *Geometry of surfaces*. Springer-Verlag, 1992.

- [18] Roberto Tamassia. *Handbook of graph drawing and visualization*. Chapman and Hall/CRC, 2013.
- [19] Charles M. Terry, Lloyd R. Welch, and John W.T. Youngs. The genus of K_{12s} . *Journal of Combinatorial Theory*, 2(1):43–60, 1967.