



Minimal Delaunay Triangulations of Hyperbolic Surfaces

Matthijs Ebbens¹ · Hugo Parlier² · Gert Vegter¹

Received: 17 December 2020 / Revised: 24 August 2021 / Accepted: 19 October 2021 /
Published online: 21 February 2022
© The Author(s) 2022

Abstract

Motivated by recent work on Delaunay triangulations of hyperbolic surfaces, we consider the minimal number of vertices of such triangulations. First, we 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 construct a class of hyperbolic surfaces for which the order of this bound is optimal. Finally, to give a general lower bound, we 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.

Keywords Delaunay triangulations · Hyperbolic surfaces · Metric graph embeddings · Moduli spaces

Mathematics Subject Classification 32G15 · 52C45 · 57K20 · 57M15

Editor in Charge: Csaba D. Tóth

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

Matthijs Ebbens
ymebbens@gmail.com

Hugo Parlier
hugo.parlier@uni.lu

Gert Vegter
g.vegter@rug.nl

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

² Mathematics Research Unit, University of Luxembourg, Esch-sur-Alzette, Luxembourg

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 [17], 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 1.1 *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 of them grows like $\Omega(g)$.*

The above result is a compilation of Theorems 3.1 and 4.1 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 Sect. 2 for details). It is a classical theorem of Bers [2] that any surface admits a short pants decomposition, meaning that the length of 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 short pants decompositions with geodesically convex curves. Indeed it is known that *any* pants decomposition of a random surface (chosen with respect to a natural probability measure on \mathcal{M}_g) has at least one curve of length on the order of $g^{1/6-\varepsilon}$ 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.

Apart from the Fenchel–Nielsen coordinates, the lengths of edges in a given triangulation are also a parameter set for \mathcal{M}_g . By Theorem 1.1, 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.

Theorem 1.1 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 [15] to construct a family of hyperbolic surfaces that attain this bound for infinitely many genera (Theorem 5.2), 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 Sect. 2, we introduce our notation and give some preliminaries on hyperbolic surface theory and triangulations. In Sect. 3, we prove our linear upper bound for the number of vertices of a minimal distance Delaunay triangulation. In Sect. 4, we construct classes of hyperbolic surfaces attaining the order of this linear upper bound. Finally, in Sect. 5, we construct a family of hyperbolic surfaces attaining the general $\Theta(\sqrt{g})$ lower bound. The proof of a technical lemma appears in the appendix.

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 center and radius of a hyperbolic circle are different from the Euclidean center 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 [16], 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 on a hyperbolic surface 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 (Fig. 1). 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, i.e., a surface of signature $(0, 3)$.

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 marked points on the boundary curves.

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 a 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 length of the geodesic changes, but the free homotopy class does not. 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 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, Fig. 2 shows the 3-regular graph corresponding to the pair of pants decomposition in Fig. 1. As our parametrization of \mathcal{T}_g depends on a choice of pair of pants decomposition, one




Fig. 1 Decomposition of a genus 3 surface into four pairs of pants using six disjoint simple closed geodesics




Fig. 2 3-Regular graph corresponding to the pair of pants decomposition shown in Fig. 1

can think of the Fenchel–Nielsen coordinates associating a length and a twist to each edge.

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 2.1 ([6, Thm. 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*

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

is an embedded hyperbolic cylinder isometric to $[-w(\gamma), w(\gamma)] \times \mathbb{S}^1$ 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 2.2 *A distance Delaunay triangulation is a triangulation satisfying the following three properties:*

- it is a simplicial complex,
- it is a Delaunay triangulation,
- 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. 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. Using the correspondence between hyperbolic and Euclidean circles, we define Delaunay triangulations in the hyperbolic plane similarly.

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 3.7 and the proof of Theorem 3.1 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 segment γ mapping bijectively to (u, 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 150 is certainly not optimal.

Theorem 3.1 *For every closed hyperbolic surface X of genus g there exists a distance Delaunay triangulation $T \in \mathcal{D}(X)$ with at most 150g 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 150g 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 $\text{sys}(X)/2$, 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

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

and its ε -*thin* part $X_{\text{thin}}^\varepsilon = X \setminus X_{\text{thick}}^\varepsilon$, where $\text{injrads}(x)$ is the injectivity radius at x , i.e., the radius of the largest embedded open disk centered at x . Note that the minimum of $\text{injrads}(x)$ over all $x \in X$ is given by $\text{sys}(X)/2$. We will see in Sect. 3.1 that, for a sufficiently small $\varepsilon > 0$, $X_{\text{thin}}^\varepsilon$ is a collection of hyperbolic cylinders (see Fig. 3). In these hyperbolic cylinders we want to construct a set of vertices the cardinality of which 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 Sect. 3.1. Because $\text{injrads}(x) > \varepsilon$ 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 Sect. 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 3.1 is given in Sect. 3.3.




Fig. 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 $X_{\text{thin}}^\varepsilon$ of the hyperbolic surface X . The following lemma describes $X_{\text{thin}}^\varepsilon$ in more detail.

Lemma 3.2 ([6, Thm. 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 [14]. 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 (Fig. 4), which is completely contained in the hyperbolic cylinder. All points on the boundary curves have the same distance K_C to a closed geodesic γ (called the *waist* of C), where K_C 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 (Fig. 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)]

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

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

$$K_C = \operatorname{arccosh} \frac{\sinh \varepsilon}{\sinh(\ell(\gamma)/2)}. \tag{2}$$

We see that γ^+ consists of points that are equidistant to γ . Moreover, γ^+ and γ^- are smooth.




Fig. 4 Computing K_C

Recall the notion of a collar from Sect. 2. In particular, each hyperbolic cylinder C in $X_{\text{thin}}^\varepsilon$ is a collar of width K_C , i.e., $C = C_\gamma(K_C)$. 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 3.4 to give a lower bound for the distance between distinct hyperbolic cylinders in $X_{\text{thin}}^\varepsilon$.

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 $2\varepsilon' < \ell(\gamma) \leq 2\varepsilon$, where the first inequality follows from γ being contained in an ε' -thick cylinder and the second inequality from γ being contained in the ε -thin part $X_{\text{thin}}^\varepsilon$. An ε' -thin cylinder satisfies $\ell(\gamma) \leq 2\varepsilon'$.

Lemma 3.10 in Sect. 3.2 states that the triangulation depicted in Fig. 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 3.6.

Definition 3.3 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 Fig. 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 Fig. 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 3.5 states that all edges in a standard triangulation are distance paths if




Fig. 5 Standard triangulation of an ε' -thin cylinder

$\varepsilon \leq (3 \log \phi)/2$, where $\phi = (1 + \sqrt{5})/2$ is the golden ratio. Before we can prove Corollary 3.5, we first need the following lemma.

Lemma 3.4 *Let X be a closed hyperbolic surface and let $\varepsilon \leq (3 \log \phi)/2$, where ϕ is the golden ratio. For each pair of distinct closed geodesics γ_1 and γ_2 in $X_{\text{thin}}^\varepsilon$ the collars $C_{\gamma_1}(K_{C_1} + \varepsilon/3)$ and $C_{\gamma_2}(K_{C_2} + \varepsilon/3)$ are embedded and disjoint.*

Proof See Fig. 6. We will show that $w(\gamma_i) - K_{C_i} \geq \varepsilon/3$ for $i = 1, 2$. Namely, this implies that $C_{\gamma_i}(K_{C_i} + \varepsilon/3) \subseteq C_{\gamma_i}(w(\gamma_i))$. 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 $C_{\gamma_1}(K_{C_1} + \varepsilon/3)$ and $C_{\gamma_2}(K_{C_2} + \varepsilon/3)$ 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. We write $w(\gamma_i) - K_{C_i} = F(1/\sinh(\ell(\gamma_i)/2))$, where

$$F(x) = \operatorname{arcsinh} x - \operatorname{arccosh}(x \sinh \varepsilon).$$

It can be easily seen that the derivative of F is negative, so its infimum is obtained in the limit $x \rightarrow \infty$, i.e., when $\ell(\gamma_i) \rightarrow 0$. Straightforward computations show that this limit is equal to $-\log(\sinh \varepsilon)$ and that $-\log(\sinh \varepsilon) \geq \varepsilon/3$ for $\varepsilon = (3 \log \phi)/2$. Since $w(\gamma_i) - K_{C_i}$ is decreasing as a function of ε , it follows that $w(\gamma_i) - K_{C_i} \geq \varepsilon/3$ for all $\varepsilon \leq (3 \log \phi)/2$. □

Corollary 3.5 *Let X be a closed hyperbolic surface and let $\varepsilon \leq (3 \log \phi)/2$, where ϕ is the golden ratio. All edges in a standard triangulation of an ε' -thin cylinder in $X_{\text{thin}}^\varepsilon$ are distance paths.*




Fig. 6 Illustration of the collars $C_{\gamma_i}(K_{C_i}) \subset C_{\gamma_i}(K_{C_i} + \varepsilon/3) \subseteq C_{\gamma_i}(w(\gamma_i))$

Proof Because the convex hull of each boundary geodesic γ^+ and γ^- is contained in C [14, pp. 1407–1408], 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+1}^+ . 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 $C_\gamma(K_C + \varepsilon/3)$ is embedded by Lemma 3.4, any path that leaves the top half of the cylinder and returns through the bottom half has length at least $K_C + 2\varepsilon/3$. It follows that the edges of the form (x_i, x_{i+1}^+) 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 $d(x_i, x_{i+1}) = \ell(\gamma)/3 < 2\varepsilon/3$ and $d(x_{i+1}, x_{i+1}^+) = K_C$, we see from the triangle inequality that $d(x_i, x_{i+1}^+) < K_C + 2\varepsilon/3$. Because any path that leaves the top half of the cylinder and returns through the bottom part of the cylinder has length at least $K_C + 2\varepsilon/3$ 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 (2) for K_C that the width K_C is close 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 3.6 Let X be a closed hyperbolic surface. Let C be an 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 3.7 Set $\varepsilon = (3 \log \phi)/2$, where ϕ is the golden ratio, 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 $X \setminus \bigcup_j T_j$ such that $d(p, q) \geq \varepsilon/2$ 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 3.8 Because by Corollary 3.5 all edges in a standard triangulation of any ε' -thin cylinder are distance paths if we choose $\varepsilon \leq (3 \log \phi)/2$, we have chosen $\varepsilon = (3 \log \phi)/2$ in Definition 3.7. Namely, we will see in the proof of Theorem 3.1 that the larger we choose ε , the smaller the constant (in our case 150) in the upper bound for the number of vertices. As in Sect. 3.1 we will fix $\varepsilon = (3 \log \phi)/2$ 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 $X \setminus \bigcup_j T_j$ instead of in $X \setminus \bigcup_j C_j$ to simplify our proof of Lemma 3.13, 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 3.9 *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 $\varepsilon/2$ centered at the vertices of V covers C completely. Namely, this implies that every point of C has distance at most $\varepsilon/2$ to a vertex of V . Because $d(p, q) \geq \varepsilon/2$ 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 $d(x_i, x_{i+1}) = \ell(\gamma)/3 < 2\varepsilon/3$ for all $i = 1, 2, 3$ (counting modulo 3). Therefore, the circles of radius $\varepsilon/2$ centered at x_i and x_{i+1} intersect in two points, of which we call one p . Since the collar




Fig. 7 Computing $d(\gamma, p)$

$C_\gamma(d(\gamma, p))$ 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 (2) for K_C we know that

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

where we substituted $\varepsilon' = 0.99\varepsilon$ and $\varepsilon = (3 \log \phi)/2$ in the last step. On the other hand, the hyperbolic Pythagorean theorem yields

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

(see Fig. 7) where again we substituted $\varepsilon = (3 \log \phi)/2$ in the last step. We conclude that $K_C < d(\gamma, p)$, which finishes the proof. □

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 3.10 *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 3.11 The proof of Lemma 3.10 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 3.12.

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 $X_{\text{thick}}^\varepsilon$.




Fig. 8 Circumscribed disk of a triangle (u, v, w) of which the center is contained in T_j for some j

Lemma 3.12 *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, assume that $c \in T_j$ for some ε' -thin cylinder C_j . It is not possible that both u and v are vertices of a standard triangulation of C_j , because this would contradict (u, v) having non-empty intersection with $X_{\text{thick}}^\varepsilon$. Therefore, we can assume without loss of generality that the circumcircle of (u, v, w) intersects some edge $[x_i^+, x_{i+1}^+]$ (Fig. 8). Denote the intersection points of this circumcircle with $[x_i^+, x_{i+1}^+]$ by \tilde{x}_i^+ and \tilde{x}_{i+1}^+ . Note that in general x_i^+ (resp. x_{i+1}^+) may coincide with \tilde{x}_i^+ (resp. \tilde{x}_{i+1}^+), and one of u, v , and w may also coincide with one of x_i^+ and x_{i+1}^+ , but this does not affect the argument.

The distance between two points on the shortest arc of the circumscribed circle between \tilde{x}_i^+ and \tilde{x}_{i+1}^+ is smaller than the distance between \tilde{x}_i^+ and \tilde{x}_{i+1}^+ , which is smaller than the distance between x_i^+ and x_{i+1}^+ . Therefore, $\ell(u, v) < d(x_i^+, x_{i+1}^+)$. Because $d(x_i^+, x_{i+1}^+) < \varepsilon$ by Remark 3.11, it follows that $\ell(u, v) < \varepsilon$. Second, if $c \in X \setminus \bigcup_{j \in I} T_j$, then we can deduce that $r < \varepsilon/2$. Namely, if we suppose for a contradiction that $r \geq \varepsilon/2$, then $d(c, p) \geq \varepsilon/2$ 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 $r < \varepsilon/2$. Because (u, v) is contained in a circle of radius $r < \varepsilon/2$, 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 $2\ell(u, v) < 2\varepsilon$. However, because $\text{injrad}(x) > \varepsilon$ 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 T 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 3.13 *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 3.12 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 3.12.

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 3.9. Then $\ell(v, w) = \ell(\gamma)/3 < 2\varepsilon/3$, so (u, v) and (v, w) do not form a 2-cycle by Lemma 3.12. 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 3.11, so (u, v) and (v, w) do not form a 2-cycle by Lemma 3.12.

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 3.12 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 3.5 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 3.1

Proof of Theorem 3.1 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 3.13, 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 $150g$. By Lemma 3.2, $X_{\text{thin}}^\varepsilon$ consists of at most $3g - 3$ cylinders and each of these cylinders contains either nine vertices (if it is ε' -thin) or three 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 $B_p(\varepsilon/4)$ and $B_q(\varepsilon/4)$ of radius $\varepsilon/4$ 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 $\varepsilon/4$ that we can fit in X . Because the area of a hyperbolic disk of radius $\varepsilon/4$ is $2\pi(\cosh(\varepsilon/4) - 1)$ [1] and because the area of X is $4\pi(g - 1)$ [16], we obtain

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

Therefore, substituting $\varepsilon = (3 \log \phi)/2$, we obtain

$$|\mathcal{P}| \leq 27g - 27 + \frac{2(g - 1)}{\cosh(\varepsilon/4) - 1} \leq 150g.$$

This finishes the proof. □

Remark 3.14 The constant 150 is not optimal. We can obtain the stronger upper bound $|\mathcal{P}| \leq 123g$ 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 A Class 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 3.1. We will first introduce this class of hyperbolic surfaces and then state the precise result in Theorem 4.1.

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 Fig. 9 with corresponding pair of pants decomposition depicted in Fig. 10. 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.



Fig. 9 Trivalent graph L_g




Fig. 10 Pair of pants decomposition corresponding to L_g

Now, fix some positive real number M . Let $\mathcal{T}_g(M)$ be the subset of \mathcal{T}_g with underlying graph L_g such that for all odd $i = 1, 3, \dots, 2g - 3$ the length of the boundary geodesic $Y_i \cap Y_{i+1}$ is at most M . In particular, $\mathcal{T}_g(M)$ contains an open subset of \mathcal{T}_g , showing that having a linear number 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 4.1 *A distance triangulation of any hyperbolic surface in $\mathcal{T}_g(M)$ where*

$$M = 4 \operatorname{arccosh} \left(\frac{1}{6} \sqrt[3]{54 - 6\sqrt{33}} + \frac{1}{6} \sqrt[3]{54 + 6\sqrt{33}} \right) \approx 2.4375 \dots$$

has at least g vertices.

We emphasize that the property that all edges are distance paths is sufficient; we do not require the triangulation to be a Delaunay triangulation.

The idea of the proof is to show that, given any distance triangulation of a hyperbolic surface $X \in \mathcal{T}_g(M)$, the union $Y_i \cup Y_{i+1}$ contains at least one vertex for each even $i = 2, 4, \dots, 2g - 4$. To show this, we need the following lemma.

Lemma 4.2 *Let X be a closed hyperbolic surface. Let γ be a simple closed geodesic on X with length at most M , where M is as defined in Theorem 4.1. Any shortest path between two points on X that is not a proper subset of γ intersects γ in at most one point.*

Proof Suppose, for a contradiction, that there exists a shortest path β between two points p and s on X that intersects γ in at least two points, say q and r . We subdivide β into the shortest paths $[p, q]$, $[q, r]$, and $[r, s]$ and will show that $\ell([q, r])$ is larger than $\ell(\gamma)/2$ (Fig. 11). Since the length of one of the two segments of γ between q and r is at most $\ell(\gamma)/2$, this contradicts the fact that $[q, r]$ is a distance path.

Consider the collar $C_\gamma(w(\gamma))$ around γ of width

$$w(\gamma) = \operatorname{arcsinh} \frac{1}{\sinh(\ell(\gamma)/2)}.$$




Fig. 11 Construction used in the proof of Lemma 4.2

If $[q, r]$ is contained in $C_\gamma(w(\gamma))$, then it has the same homotopy class of paths as one of the two segments of γ between q and r . Because there are no hyperbolic bigons, it follows that $[q, r]$ is contained in γ , which, by assumption, is not possible.

Therefore, $[q, r]$ is not contained in $C_\gamma(w(\gamma))$. Then we can subdivide $[q, r]$ into a path from q to a point $q^* \notin C_\gamma(w(\gamma))$ and a path from q^* to r , the length of each of which is larger than $w(\gamma)$. Therefore, $\ell([q, r]) > 2w(\gamma)$. A straightforward computation shows that if $\ell(\gamma) \leq M$, then

$$\sinh \frac{\ell(\gamma)}{4} \sinh \frac{\ell(\gamma)}{2} \leq 1,$$

i.e.,

$$w(\gamma) = \operatorname{arcsinh} \frac{1}{\sinh(\ell(\gamma)/2)} \geq \frac{\ell(\gamma)}{4}.$$

Hence, $\ell([q, r]) > 2w(\gamma) \geq \ell(\gamma)/2$. This finishes the proof. □

We proceed with the proof of Theorem 4.1.

Proof of Theorem 4.1 Let X be a hyperbolic surface in $\mathcal{T}_g(M)$ and let T be a distance triangulation of X . We first show that the union $Y_i \cup Y_{i+1}$ contains at least one vertex for each even $i = 2, 4, \dots, 2g - 4$ and then that Y_1 and Y_{2g-2} also contain at least one vertex each, leading to a total of at least g vertices.

Take an arbitrary even $i = 2, 4, \dots, 2g - 4$. Suppose, for a contradiction, that $Y_i \cup Y_{i+1}$ does not contain any vertices of T . Consider a triangle Δ that intersects $Y_i \cup Y_{i+1}$. An edge of Δ with one of its endpoints in $\bigcup_{j=1}^{i-1} Y_j$ (resp. $\bigcup_{j=i+1}^{2g-2} Y_j$) that intersects $Y_i \cup Y_{i+1}$ has its other endpoint in $\bigcup_{j=i+1}^{2g-2} Y_j$ (resp. $\bigcup_{j=1}^{i-1} Y_j$), since the edge intersects each of the boundary geodesics $Y_{i-1} \cap Y_i$ and $Y_{i+1} \cap Y_{i+2}$ at most once by Lemma 4.2. In particular, two of the vertices of Δ are contained in $\bigcup_{j=1}^{i-1} Y_j$ and the other in $\bigcup_{j=i+1}^{2g-2} Y_j$ (or reversely), and precisely two of the edges of Δ intersect $Y_i \cup Y_{i+1}$.

Intuitively, the situation is as follows. If a triangle Δ of T intersects $Y_i \cup Y_{i+1}$, then the intersection is a quadrilateral. In this quadrilateral, two of the sides are given by the intersections of two of the edges of Δ with $Y_i \cup Y_{i+1}$ and the other two are

given by the segments of the boundary geodesics $Y_{i-1} \cap Y_i$ and $Y_{i+1} \cap Y_{i+2}$ contained in Δ . Every edge of T that intersects $Y_i \cup Y_{i+1}$ is contained in two triangles of T , so the intersections of triangles of T with $Y_i \cup Y_{i+1}$ glue together to form a cylinder. The boundary geodesics of this cylinder are the boundary geodesics $Y_{i-1} \cap Y_i$ and $Y_{i+1} \cap Y_{i+2}$, formed by gluing together the segments of these boundary geodesics contained in the triangles of T . However, this contradicts the fact that $Y_i \cup Y_{i+1}$ is a torus.

To make this argument more precise, we will look at the embedding of a graph into $Y_i \cup Y_{i+1}$ and apply Euler's formula $v - e + f = 2 - 2g$. We place a vertex on each of the intersection points of an edge of T with the boundary geodesics $Y_{i-1} \cap Y_i$ and $Y_{i+1} \cap Y_{i+2}$ and shorten the corresponding edges so that their endpoints are now the added vertices. For each boundary geodesic we add edges between consecutive vertices, where the cyclic order on the vertices is naturally induced from the boundary geodesic, and we add a face that is incident to all vertices on the boundary geodesic. In this way, we obtain an embedding of a graph G into a closed topological surface of genus 1. Let k be the number of triangles of T that intersect $Y_i \cup Y_{i+1}$. Precisely two edges of each triangle intersect $Y_i \cup Y_{i+1}$ and each edge is contained in two triangles, so k edges of T intersect $Y_i \cup Y_{i+1}$. By Lemma 4.2, each of these edges intersects each of the boundary geodesics $Y_{i-1} \cap Y_i$ and $Y_{i+1} \cap Y_{i+2}$ precisely once, so G has $2k$ vertices. We shortened the k edges of T that intersect $Y_i \cup Y_{i+1}$ and we added edges between consecutive vertices on the boundary geodesics, so G has $3k$ edges in total. We added two faces to the k faces obtained from the triangles of T intersecting $Y_i \cup Y_{i+1}$, so G has $k + 2$ faces. Now, substituting v , e and f into Euler's formula yields $g = 0$, which is a contradiction. We conclude that $Y_i \cup Y_{i+1}$ contains at least one vertex for each even $i = 2, 4, \dots, 2g - 4$.

To show that Y_1 contains at least one vertex, suppose, for a contradiction, that it does not. There exists a triangle Δ of T that intersects Y_1 and, by assumption, Δ has its vertices in $\bigcup_{j=2}^{2g-2} Y_j$. However, edges of Δ intersect the boundary geodesic $Y_1 \cap Y_2$ at most once, which leads to a contradiction. Hence, Y_1 contains at least one vertex. It can be shown in a similar way that Y_{2g-2} contains at least one vertex. This finishes the proof. \square

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 5.1 ([12, Thm. 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 5.2 *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 5.1 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 [14].

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 [15]. 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 [18]. 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 $2\pi/(n - 1)$. In the complete graph K_n every vertex has $n - 1$ neighboring 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 .




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

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 [14]. 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 (Fig. 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 Fig. 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 $\ell/2$, where ℓ denotes the length of the sides of the equilateral triangle [14]. Hence, $d(c, p) = \ell(\gamma) \geq d(c, x) + \ell/2$. By the triangle inequality in triangle (c, w, x) we see that $d(c, w) \leq d(c, x) + d(x, w) = d(c, x) + \ell/2$, 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

Acknowledgements The authors warmly thank Monique Teillaud and Vincent Despré for fruitful discussions. We would also like to thank the referees for their helpful comments, one of which allowed us to shorten Sect. 4 significantly.

Declarations

Conflict of interest The authors declared that they have no conflict of interest.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included

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

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

where the last equality follows from expression (2) for K_C . Interpreting the right-hand side as a function of $\ell(\gamma)$ we see that it attains its supremum for $\ell(\gamma) \rightarrow 0$, for which

$$\frac{\sinh(\ell(\gamma)/6) \sinh \varepsilon}{\sinh(\ell(\gamma)/2)} \rightarrow \frac{\sinh \varepsilon}{3}.$$

Because $(\sinh \varepsilon)/3 < \sinh(\varepsilon/2)$ for $\varepsilon = (3 \log \phi)/2$, it follows $\sinh(d(x_i^+, x_{i+1}^+)/2) < \sinh(\varepsilon/2)$, so $d(x_i^+, x_{i+1}^+) < \varepsilon$. Because $d(x_i^+, x_{i+1}^+) < \varepsilon$, the circles of radius $\varepsilon/2$ centered at x_i^+ and x_{i+1}^+ intersect in two points.

Now, to show that we can assume without loss of generality that $d(x_i^+, p) = d(x_{i+1}^+, p) = \varepsilon/2$, we show that $d(p, c_i)$ is minimal when these equations hold, i.e., when p is one of the intersection points of the circles of radius $\varepsilon/2$ centered at x_i^+ and x_{i+1}^+ . 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 (see Fig. 13). Because all points on γ_c have the same distance to γ , the distance $d(p, p_c)$ is minimal for some p if and only if the distance $d(p, \gamma)$ is minimal for this p . Since $d(p, x_j^+) \geq \varepsilon/2$ for all $j = 1, 2, 3$ and since the three circles of radius $\varepsilon/2$ centered at x_j^+ for $j = 1, 2, 3$ intersect pairwise, $d(p, \gamma)$ (and hence $d(p, p_c)$) is minimal if p is an intersection point of a pair of these circles, say, of the circles centered at x_k^+ and x_{k+1}^+ for some $k \in \{1, 2, 3\}$. Note that $d(p, c_i) \geq d(p, p_c)$ by definition of p_c , with equality if and only if p lies on the geodesic passing through m_i and m_i^+ . A point p such that $d(p, p_c)$ is minimal lies on this geodesic if and only if $k = i$, i.e., if and only if $d(x_i^+, p) = d(x_{i+1}^+, p) = \varepsilon/2$. Therefore, $d(p, c_i)$ is minimal when $d(x_i^+, p) = d(x_{i+1}^+, p) = \varepsilon/2$. We conclude that we can assume without loss of generality that $d(x_i^+, p) = d(x_{i+1}^+, p) = \varepsilon/2$.

Let c'_i be the projection of c_i on (x_i, x_i^+) , as in Fig. 14. 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

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

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

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

Here, (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




Fig. 14 Schematic overview of the trigonometry in Lemma 3.10

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

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

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 (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 $\varepsilon = (3 \log \phi)/2$ 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 $d(c_i, p) = d(m_i, m_i^+) - d(m_i, c_i) + d(m_i^+, p) > d(c_i, x_i)$. Hence, p is not contained in D_i . This finishes the proof.

References

1. Beardon, A.F.: The Geometry of Discrete Groups. Graduate Texts in Mathematics, vol. 91. Springer, New York (1995)

2. Bers, L.: An inequality for Riemann surfaces. In: *Differential Geometry and Complex Analysis*, pp. 87–93. Springer, Berlin (1985)
3. Bogdanov, M., Teillaud, M.: Delaunay triangulations and cycles on closed hyperbolic surfaces (2014). <https://hal.archives-ouvertes.fr/hal-00921157>
4. Bogdanov, M., Teillaud, M., Vegter, G.: Delaunay triangulations on orientable surfaces of low genus. In: *32nd International Symposium on Computational Geometry (Boston 2016)*. Leibniz Int. Proc. Inform., vol. 51, # 20. Leibniz-Zent. Inform., Wadern (2016)
5. Bowyer, A.: Computing Dirichlet tessellations. *Comput. J.* **24**(2), 162–166 (1981)
6. Buser, P.: *Geometry and Spectra of Compact Riemann Surfaces*. Birkhäuser, Boston (2010)
7. Despré, V., Schlenker, J.M., Teillaud, M.: Flipping geometric triangulations on hyperbolic surfaces (2019). [arXiv:1912.04640](https://arxiv.org/abs/1912.04640)
8. Fáry, I.: On straight line representation of planar graphs. *Acta Univ. Szeged. Sect. Sci. Math.* **11**(4), 229–233 (1948)
9. Guth, L., Parlier, H., Young, R.: Pants decompositions of random surfaces. *Geom. Funct. Anal.* **21**(5), 1069–1090 (2011)
10. Hubbard, A., Kaluža, V., de Mesmay, A., Tancer, M.: Shortest path embeddings of graphs on surfaces. *Discret. Comput. Geom.* **58**(4), 921–945 (2017)
11. Iordanov, I., Teillaud, M.: Implementing Delaunay triangulations of the Bolza surface. In: *33rd International Symposium on Computational Geometry (Brisbane 2017)*. Leibniz Int. Proc. Inform., vol. 77, # 44. Leibniz-Zent. Inform., Wadern (2017)
12. Jungerman, M., Ringel, G.: Minimal triangulations on orientable surfaces. *Acta Math.* **145**(1–2), 121–154 (1980)
13. Mirzakhani, M.: Growth of Weil–Petersson volumes and random hyperbolic surfaces of large genus. *J. Differ. Geom.* **94**(2), 267–300 (2013)
14. Parlier, H., Petit, C.: Chromatic numbers of hyperbolic surfaces. *Indiana Univ. Math. J.* **65**(4), 1401–1423 (2016)
15. Ringel, G., Youngs, J.W.T.: Solution of the Heawood map-coloring problem. *Proc. Natl Acad. Sci. U.S.A.* **60**, 438–445 (1968)
16. Stillwell, J.: *Geometry of Surfaces*. Universitext. Springer, New York (1992)
17. Tamassia, R. (ed.): *Handbook of Graph Drawing and Visualization*. Discrete Mathematics and its Applications (Boca Raton). CRC Press, Boca Raton (2014)
18. Terry, C.M., Welch, L.R., Youngs, J.W.T.: The genus of $K_{12\mathfrak{S}}$. *J. Combin. Theory* **2**, 43–60 (1967)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.