Counting essential surfaces
in a closed hyperbolic three-manifold

JEREMY KAHN

VLADIMIR MARKOVIĆ

Let $\mathbf{M}^3$ be a closed hyperbolic three-manifold. We show that the number of genus g surface subgroups of $\pi_1(\mathbf{M}^3)$ grows like g^{2g}.

1 Introduction

Let $\mathbf{M}^3$ be a closed hyperbolic 3–manifold and let S_g denote a closed surface of genus g. Given a continuous mapping $f\colon S_g \to \mathbf{M}^3$ we let $f_*\colon \pi_1(S_g) \to \pi_1(\mathbf{M}^3)$ denote the induced homomorphism.

Definition 1.1 *We say that $G < \pi_1(\mathbf{M}^3)$ is a surface subgroup of genus $g \geq 2$ if there exists a continuous map $f\colon S_g \to \mathbf{M}^3$ such that the induced homomorphism f_* is injective and $f_*(\pi_1(S_g)) = G$. Moreover, the subsurface $f(S_g) \subset \mathbf{M}^3$ is said to be an essential subsurface.*

Recently, we showed [3] that every closed hyperbolic 3–manifold $\mathbf{M}^3$ contains an essential subsurface and consequently $\pi_1(\mathbf{M}^3)$ contains a surface subgroup. It is therefore natural to consider the question: How many conjugacy classes of surface subgroups of genus g there are in $\pi_1(\mathbf{M}^3)$? This has already been considered by Masters [5], and our approach to this question builds on our previous work and improves on the work by Masters.

Let $s_2(\mathbf{M}^3, g)$ denote the number of conjugacy classes of surface subgroups of genus at most g. We say that two surface subgroups G_1 and G_2 of $\pi_1(\mathbf{M}^3)$ are commensurable if $G_1 \cap G_2$ has a finite index in both G_1 and G_2. Let $s_1(\mathbf{M}^3, g)$ denote the number surface subgroups of genus at most g, modulo the equivalence relation of commensurability. Then clearly $s_1(\mathbf{M}^3, g) \leq s_2(\mathbf{M}^3, g)$. The main result of this paper is the following theorem.

Theorem 1.1 *Let* $\mathbf{M}^3$ *be a closed hyperbolic* 3*–manifold. There exist two constants* $c_1, c_2 > 0$ *such that*

$$(c_1 g)^{2g} \leq s_1(\mathbf{M}^3, g) \leq s_2(\mathbf{M}^3, g) \leq (c_2 g)^{2g},$$

for g *large enough. The constant* c_2 *depends only on the injectivity radius of* $\mathbf{M}^3$.

In fact, Masters shows that

$$s_2(g, \mathbf{M}^3) < g^{c_2 g}$$

for some $c_2 \equiv c_2(\mathbf{M}^3)$, and likewise for some $c_1 \equiv c_1(\mathbf{M}^3)$

$$g^{c_1 g} < s_1(g, \mathbf{M}^3)$$

when $\mathbf{M}^3$ has a self-transverse totally geodesic subsurface. We follow Masters' approach to the upper bound, improving it from $g^{c_2 g}$ to $(c_2 g)^{2g}$ by more carefully counting the number of suitable triangulations of a genus g surface. Using our previous work [3] we replace Masters' conditional lower bound with an unconditional one, and we improve it from g^{cg} to $(c_1 g)^{2g}$ with the work of Muller and Puchta [6] counting number of maximal surface subgroups of a given surface group. We then make a new subgroup from old in the spirit of Masters' construction, but taking the nearly geodesic subgroup from [3] as our starting point.

The above theorem enables us to determine the order of the number of surface subgroups up to genus g. We have the following corollary.

Corollary 1.1 *We have*

$$\lim_{g \to \infty} \frac{\log s_1(\mathbf{M}^3, g)}{2g \log g} = \lim_{g \to \infty} \frac{\log s_2(\mathbf{M}^3, g)}{2g \log g} = 1.$$

We make the following conjecture.

Conjecture 1.1 *For a given closed hyperbolic* 3*–manifold* $\mathbf{M}^3$, *there exists a constant* $c(M) > 0$ *such that*

$$\lim_{g \to \infty} \frac{1}{g} \sqrt[2g]{s_i(\mathbf{M}^3, g)} = c(M), \quad i = 1, 2.$$

2 The upper bound

Fix a closed hyperbolic 3–manifold $\mathbf{M}^3$. In this section we prove the upper bound in Theorem 1.1, that is we show

$$(1) \qquad\qquad s_2(\mathbf{M}^3, g) \leq (c_2 g)^{2g}$$

for some constant $c_2 > 0$.

2.1 Genus g triangulations

We have the following definition.

Definition 2.1 *Let S_g denote a closed surface of genus g. We say that a connected graph τ is a triangulation of genus g if it can be embedded into the surface S_g such that every component of the set $S_g \setminus \tau$ is a triangle. The set of genus g triangulations is denoted by $\mathcal{T}(g)$. We say that $\tau \in \mathcal{T}(k, g) \subset \mathcal{T}(g)$ if*

- *each vertex of τ has the degree at most k,*
- *the graph τ has at most kg vertices and at most kg edges.*

We observe that any given genus g triangulation τ, can be in a unique way (up to a homeomorphism of S_g) be embedded in S_g.

We say that Riemann surface is s–thick is its injectivity radius is bounded below by $s > 0$. Every thick Riemann surface has a good triangulation in the sense of the following lemma.

Lemma 2.1 *Let S be an s–thick Riemann surface of genus $g \geq 2$. Then there exists $k = k(s) > 0$ and a triangulation $\tau \in \mathcal{T}(k, g)$ that embeds in S, such that every edge of τ is a geodesic arc of length at most s.*

Proof Choose a maximal collection of disjoint open balls in S of radius $s/4$. Let V denote the set of centers of the balls from the collection. We may assume that no four points from V lie on a round circle (we always reduce the radius of the balls by a small amount and move them into a general position). We construct the Delaunay triangulation associated to the set V as follows. We connect two points from V with the shortest geodesic arc between them, providing they belong to the boundary of a closed ball in S that does not contain any other point from V. This gives an embedded graph τ. Since no four points from V lie on the same circle the graph τ is a triangulation. It is elementary to check that τ has the stated properties, and we leave it to the reader. $\qquad\square$

Given any π_1–injective immersion of $g\colon S_g \to \mathbf{M}^3$, we can find a genus g hyperbolic surface S, and a map $f\colon S \to \mathbf{M}^3$ homotopic to g, such that $f(S)$ is a pleated surface. Then f does not increase the hyperbolic distance. Let s denote the injectivity radius of $\mathbf{M}^3$. It follows that the injectivity radius of S is bounded below by s. We choose a triangulation $\tau(S)$ of S that satisfies the conditions in Lemma 2.1.

Let $\mathcal{C} = \{C_1, \ldots, C_m\}$ be a finite collection of balls of radius $s/4$ that covers $\mathbf{M}^3$. We may assume that $\mathcal{C}$ is a minimal collection, that is, if we remove a ball from $\mathcal{C}$, the new collection of balls does not cover $\mathbf{M}^3$. Let $f_i\colon S_i \to \mathbf{M}^3$, $i = 1, 2$, be two pleated maps, and denote by $\tau(S_1)$ and $\tau(S_2)$ the corresponding triangulations from Lemma 2.1 of genus g surfaces S_1 and S_2. If the genus g triangulations $\tau(S_1)$ and $\tau(S_2)$ are identical, there exists a homeomorphism $h\colon S_1 \to S_2$ such that $h(\tau(S_1)) = \tau(S_2)$. Assume in addition that for every vertex v of $\tau(S_1)$, the points $f_1(v)$ and $f_2(h(v))$ belong to the same ball $C_i \in \mathcal{C}$. Then by [5, Lemma 2.4], the maps f_1 and $f_2 \circ h$ are homotopic.

Since the set $\mathcal{C}$ has m elements, there are at most m ways of mapping a given vertex of τ to the set $\mathcal{C}$. Choose a vertex v_1 of τ and choose an image of v_1 in $\mathcal{C}$, say v_1 is mapped to C_1. Let v_2 be a vertex of τ, such that v_2 and v_1 are the endpoints of the same edge.

Each edge of τ has the length at most s, and the balls from $\mathcal{C}$ have the radius $s/4$. Since f does not increase the distance, and $\mathcal{C}$ is a minimal cover of $\mathbf{M}^3$, it follows that v_2 can be mapped to at most K elements of $\mathcal{C}$, where K is a constant that depends only on s. Repeating this analysis to the remaining vertices of τ yields the estimate

$$(2) \qquad\qquad \widetilde{s}_2(\mathbf{M}^3, g) \leq m K^{kg-1} |\mathcal{T}(k, g)|,$$

where $\widetilde{s}_2(\mathbf{M}^3, g)$ denotes the number of conjugacy classes of surface subgroups of genus equal to g.

Let $\nu(k, n)$ denote the set of all graphs on n vertices so that each vertex has the degree at most k. Then $|\mathcal{T}(k, g)| \leq |\nu(k, kg)|$.

Remark Observing the estimate

$$|\nu(k, n)| \leq n^{kn},$$

Masters showed

$$\widetilde{s}_2(\mathbf{M}^3, g) \leq g^{Dg},$$

for some constant $D > 0$. However, the set $\nu(k, kg)$ has many more elements than the set $\mathcal{T}(k, g)$.

The following lemma will be proved in the next subsection.

Lemma 2.2 *There exists a constant $C > 0$ that depends only on k, such that for g large we have*

$$|\mathcal{T}(k, g)| \leq (Cg)^{2g}.$$

Given this lemma we now prove estimate (1). It follows from the Lemma 2.2 that for every g large we have

$$|\mathcal{T}(k, g)| \leq (Cg)^{2g}.$$

Combining this with (2) we get that

$$\tilde{s}_2(\mathbf{M}^3, g) \leq mK^{kg-1}(Cg)^{2g} \leq (C_1 g)^{2g}$$

holds for every $g \geq 2$ for some constant C_1. Then

$$
\begin{aligned}
s_2(\mathbf{M}^3, g) &= \sum_{r=2}^{g} \tilde{s}_2(\mathbf{M}^3, r) \\
&= \sum_{r=2}^{g} (C_1 r)^{2r} \\
&\leq (c_2 g)^{2g}
\end{aligned}
$$

for some constant c_2. This proves the estimate (1).

2.2 The proof of Lemma 2.2

Fix a triangulation $\tau \in \mathcal{T}(k, g)$ and denote the set of oriented edges by $E(\tau)$. Let $\mathbb{Q}E(\tau)$ denote the vector space of all formal sums (with rational coefficients) of edges from $E(\tau)$.

Choose a spanning tree T (a spanning tree of a connected graph is a connected tree that contains all of its vertices) for τ. Let $H_1(S_g)$ denote the first homology with rational coefficients of the surface S_g. We define the linear map $\phi \colon \mathbb{Q}E(\tau) \to H_1(S_g)$ as follows. Let $e \in (E(\tau) \setminus T)$. Then the union $e \cup T$ is homotopic (on S_g) to a unique (up to homotopy) simple closed curve $\gamma_e \subset S_g$. We let $\phi(e)$ denote the homology class of the curve γ_e in $H_1(S_g)$. We extend the map ϕ to $\mathbb{Q}E(\tau)$ by linearity.

Denote the kernel of ϕ by $K(\phi)$ and set

$$H_1(\tau, T) = \frac{\mathbb{Q}E(\tau)}{K(\phi)}.$$

Then the quotient map (also denoted by) $\phi\colon H_1(\tau, T) \to H_1(S_g)$ is injective, and in fact is an isomorphism. Since τ is a genus g triangulation, the embedding of the triangulation τ to S_g induces the surjective map of the fundamental group of τ to the fundamental group of S_g. Then the induced map ϕ between the corresponding homology groups is injective.

Let $e_1, \ldots, e_{2g} \in E(\tau)$ denote a set of $2g$ edges whose equivalence classes generate $H_1(\tau, T)$.

Lemma 2.3 Let $X = T \cup \{e_1, \ldots, e_{2g}\}$. Then every component of the set $S_g \setminus X$ is simply connected.

Proof The set X is connected (since it contains the spanning tree T, and the tree T contains all the vertices). Suppose that there exists a component of the set $S_g \setminus X$ that is not simply connected. Then there exists a simple closed curve $\gamma \subset S_g$ that is not homotopic to a point, and such that

$$\gamma \cap X = \varnothing.$$

If γ is a nonseparating curve then the homology class of γ is nontrivial in $H_1(S_g)$. Therefore, there exists a nonseparating simple closed $\alpha \subset S_g$ that intersects the curve γ exactly once. Let $q_1, \ldots, q_{2g} \in \mathbb{Q}$ be such that

$$\phi(q_1 e_1 + \cdots + q_{2g} e_{2g}) = [\alpha],$$

where $[\alpha] \in H_1(S_g)$ denotes the homology class of α. Since the intersection pairing between $[\alpha]$ and $[\gamma]$ is nonzero, and $\phi(e_1), \ldots, \phi(e_{2g})$ is a basis for $H_1(S_g)$, we conclude that for some $i \in \{1, \ldots, 2g\}$, the curve γ intersects $e_i \cup T$, which is a contradiction.

Suppose that γ is a separating curve and denote by A_1 and A_2 the two components of the set $S_g \setminus \gamma$. The set X is connected, and by the assumption it does not intersect γ. This implies that X is contained in one of the two subsurfaces A_i, say $X \subset A_1$. Then $X \cap A_2 = \varnothing$.

Since γ is not homotopic to a point, each A_i is a nonplanar surface with one boundary component. Therefore, the subsurface A_2 contains a nonseparating simple closed curve γ_2. Then γ_2 is a nonseparating simple closed curve in S_g by the above argument we have that γ_2 intersects the set X. This is a contradiction since $X \cap A_2 = \varnothing$. □

Let $P_1, \ldots, P_l$ denote the components of the set $S_g \setminus X$. Each P_i is a polygon and we let m_i denote the number of sides of the polygon P_i. Since each edge in X can

appear as a side in at most two such polygons, we have the inequality

$$(3) \qquad \sum_{i=1}^{l} m_i \leq 2kg,$$

since by definition the triangulation τ has at most kg edges.

We proceed to prove Lemma 2.2. We can obtain every triangulation $\tau \in \mathcal{T}(k, g)$ as follows. We first choose a spanning tree T, which is a tree that has at most kg vertices. Then to the tree T we add $2g$ edges $e_1, \ldots, e_{2g}$ in an arbitrary way. After adding the edges, at each vertex of the graph $T \cup \{e_1, \ldots, e_{2g}\}$ we choose a cyclic ordering. We thicken the edges of the graph $T \cup \{e_1, \ldots, e_{2g}\}$ to obtain the ribbon graph and the corresponding surface R with boundary (if this surface does not have genus g we discard this graph). The boundary components of the surface R are polygonal curves P_i, $i = 1, .., l$, made out of the edges from $T \cup \{e_1, \ldots, e_{2g}\}$. We then choose a triangulation of each polygon P_i.

It follows from this description that we can bound the number of triangulations from $\mathcal{T}(k, g)$ by $|\mathcal{T}(k, g)| \leq abcd$, where

      $a = \{\text{number of unlabeled trees } T \text{ with } n \leq kg \text{ vertices}\},$

      $b = \{\text{number of ways of adding } 2g \text{ unlabeled edges } e_1, \ldots, e_{2g} \text{ to } T\},$

      $c = \{\text{number of cyclic orderings of edges of } T \cup \{e_1, \ldots, e_{2g}\}\},$

      $d = \{\text{number of triangulations of the polygons } P_i\}.$

Let $t(n)$ denote the number of different unlabeled trees on n vertices. By [1] we have $t(n) \leq C12^n$, for some universal constant $C > 0$. It follows that $a \leq 2C12^{kg}$. The tree T has at most kg edges, so there are at most $(kg)^2$ ways of adding a labeled edge to T. All together there are at most $(kg)^{4g}$ ways of adding a labeled collection of $2g$ edges to T. To obtain the number of ways of adding unlabeled collection of $2g$ edges we need to divide this number by $(2g)!$. This yields the estimate

$$b \leq \frac{(kg)^{4g}}{(2g)!} < (k^2 g)^{2g}$$

for g large.

Since each vertex of τ has the degree at most k, and τ has at most kg edges, we obtain the estimate

$$c \leq (k!)^{kg}.$$

Let $p(m)$ denote the number of triangulations of a polygon with m sides. Then $p(m)$ is the $(m-2)$–th Catalan number and we have $p(m) < 2^{2m}$. As above, let $P_1, \ldots, P_l$

denote the polygons that we need to triangulate and let m_i denote the number of sides of the polygon P_i. Then

$$d \leq \max \prod_{i=1}^{l} p(m_i) \leq 4^{m_1 + \cdots + m_l},$$

where the maximum is taken over all possible vectors $(m_1, \ldots, m_l)$, $1 \leq l \leq 2kg$, such that $m_1 + \cdots + m_l \leq 2kg$ (see estimate (3) above). But since $m_1 + \cdots + m_l \leq 2kg$ we have $d \leq 4^{2kg}$.

Putting the estimates for a, b, c, d together we prove the lemma.

Remark If we are given a tree on a surface S, along with $2g$ edges connecting the vertices of the tree (and satisfying the hypothesis of Lemma 2.3) and a map of the resulting graph into $\mathbf{M}^3$, the we can determine the map of S into $\mathbf{M}^3$, up to homotopy. Thus we need only bound $|\mathcal{T}'(k, g)|$, where $\mathcal{T}'(k, g)$ is the set of trees of size at most kg, with $2g$ more edges added; we observe that $|\mathcal{T}'(k, g)| < ab$.

3 Quasifuchsian representations of surface groups

3.1 Generalized pants decomposition and the complex Fenchel–Nielsen coordinates

For background on complex Fenchel–Nielsen coordinates, see Series [7], Kourouniotis [4], Tan [8] and our previous work [3]. The exposition and notation we use here is in line with [3, Section 2].

Let X be a compact topological surface (possibly with boundary) and let $\rho\colon \pi_1(X) \to \mathrm{PSL}(2, \mathbb{C})$ be a representation (a homomorphism). We say that ρ is a K–quasifuchsian representation if the group $\rho(\pi_1(X))$ is K–quasifuchsian, in which case we can equip X with a complex structure $X = \mathbb{H}^2/F$, for some Fuchsian group F, such that $f_* = \rho \circ \iota$. Here $\iota\colon F \to \pi_1(X)$ is an isomorphism, and $f_*\colon F \to fFf^{-1}$ is the conjugation homomorphism, induced by an equivariant K–quasiconformal map $f\colon \partial\mathbb{H}^3 \to \partial\mathbb{H}^3$.

We will also say that a quasisymmetric map $f\colon \partial\mathbb{H}^2 \to \partial\mathbb{H}^3$ is K–quasiconformal if it has a K–quasiconformal extension to $\partial\mathbb{H}^3$.

By Π we denote a topological pair of pants with cuffs C_i, $i = 1, 2, 3$. Recall that to every representation $\rho\colon \pi_1(\Pi) \to \mathrm{PSL}(2, \mathbb{C})$, we associate the three half-lengths $\mathbf{hl}(C_i) \in \mathbb{C}_+/2i\pi\mathbb{Z}$, where $\mathbb{C}_+ = \{z \in \mathbb{C} : \mathrm{Re}(z) > 0\}$. If ρ is quasifuchsian then it

is uniquely determined by the half-lengths. The conjugacy class $[\rho]$ of a quasifuchsian representation ρ is called a skew pair of pants.

We let Π and Π' denote two pairs of pants and let $\rho\colon \pi_1(\Pi) \to \mathrm{PSL}(2, \mathbb{C})$ and $\rho'\colon \pi_1(\Pi') \to \mathrm{PSL}(2, \mathbb{C})$ denote two representations. Suppose that for some $c_1 \in \pi_1(\Pi)$ and $c_1' \in \pi_1(\Pi')$, that belong to the conjugacy classes of C_1 and C_1' respectively, we have $\rho(c_1) = \rho'(c_1')$, and $\mathbf{hl}(C_1) = \mathbf{hl}(C_1')$. By $s(C) \in \mathbb{C}/(\mathbf{hl}(C)\mathbb{Z} + 2\pi i \mathbb{Z})$ we denote the reduced twist-bend parameter, which measures how the two skew pairs of pants $[\rho]$ and $[\rho']$ align together along the axis of the loxodromic transformation $\rho(c_1) = \rho'(c_1')$.

A pair $(\widetilde{\Pi}, \chi)$ is a generalized pair of pants if $\widetilde{\Pi}$ is a compact surface with boundary and χ is a finite degree covering map $\chi\colon \widetilde{\Pi} \to \Pi$, where Π is a pair of pants. (We will also call $\widetilde{\Pi}$ a generalized pair of pants if χ is understood.) By $\chi_*\colon \pi_1(\widetilde{\Pi}) \to \pi_1(\Pi)$ we denote an induced homomorphism.

Definition 3.1 *Let $(\widetilde{\Pi}, \chi)$ be a generalized pair of pants and*

$$\widetilde{\rho}\colon \pi_1(\widetilde{\Pi}) \to \mathrm{PSL}(2, \mathbb{C})$$

be a representation. We say that $\widetilde{\rho}$ is admissible with respect to χ if it factors through χ_*, that is, there exists $\rho\colon \pi_1(\Pi) \to \mathrm{PSL}(2, \mathbb{C})$ such that $\widetilde{\rho} = \rho \circ \chi_*$.

Let $\widetilde{C}_j$, $j = 1, \ldots, k$, denote the cuffs (the boundary curves) of the surface $\widetilde{\Pi}$, and let C_1, C_2, C_3 continue to denote the cuffs of Π. Then χ maps each $\widetilde{C}_j$ onto some C_i with some degree $m_j \in \mathbb{N}$. We say that such a curve $\widetilde{C}_j$ is a degree m_j curve. For every admissible $\widetilde{\rho}$ we define the half-length $\mathbf{hl}(\widetilde{C}_j)$ as $\mathbf{hl}(\widetilde{C}_j) = \mathbf{hl}(C_i)$. Let $\widetilde{c}_j \in \pi_1(\widetilde{\Pi}^0)$ be in the conjugacy class that corresponds to the cuff $\widetilde{C}_j$. Then

$$\mathbf{l}(\widetilde{\rho}(c_i)) = 2m_j \, \mathbf{hl}(C_i) \ (\mathrm{mod}(2\pi i \mathbb{Z})).$$

Let S be an oriented closed topological surface with a generalized pants decomposition. By this we mean that we are given a collection $\mathcal{C}$ of disjoint simple closed curves on S, such that for every component $\widetilde{\Pi}$ of $S \setminus \mathcal{C}$ there is an associated finite cover $\chi\colon \widetilde{\Pi} \to \Pi$. Let

$$\widetilde{\rho}\colon \pi_1(S) \to \mathrm{PSL}(2, \mathbb{C})$$

be a representation. We make the following assumptions on ρ:

(1) Given a curve $C \in \mathcal{C}$ there exists two (not necessarily different) generalized pairs of pants $\widetilde{\Pi}_1$ and $\widetilde{\Pi}_2$ that both contain C as a cuff, and that lie on different sides of C. Let $\chi_1\colon \widetilde{\Pi}_1 \to \Pi_1$ and $\chi_2\colon \widetilde{\Pi}_2 \to \Pi_2$ be the corresponding finite covers, where Π_1 and Π_2 are two pairs of pants. We assume that the restrictions of χ_1 and χ_2 on the curve C are of the same degree.

(2) For every generalized pair of pants $\widetilde{\Pi}$ from the above decomposition of S, the
     restriction $\rho\colon \pi_1(\widetilde{\Pi}) \to \mathrm{PSL}(2,\mathbb{C})$ is admissible with respect to the covering
     map $\chi\colon \widetilde{\Pi} \to \Pi$ (in the sense of Definition 3.1).

(3) For every $C \in \mathcal{C}$, the half-lengths of C coming from the representations
     $\rho\colon \pi_1(\widetilde{\Pi}_1) \to \mathrm{PSL}(2,\mathbb{C})$ and $\rho\colon \pi_1(\widetilde{\Pi}_2) \to \mathrm{PSL}(2,\mathbb{C})$ are one and the same.

Continuing with the above notation, let $C_i \subset \Pi_i$ denote the cuff such that $\chi_i(C) = C_i$.
Let $\rho_i\colon \pi_1(\Pi_i) \to \mathrm{PSL}(2,\mathbb{C})$, $i = 1,2$, be the representations such that the restriction
of ρ to $\pi_1(\widetilde{\Pi}_i)$ is equal to $\rho_i \circ (\chi_i)_*$. We define the reduced twist bend parameter $s(C)$
associated to ρ to be equal to the reduced twist-bend parameter for the representations
ρ_1 and ρ_2.

So given a closed surface S with a generalized pants decomposition $\mathcal{C}$, and a represen-
tation $\rho\colon \pi_1(S) \to \mathrm{PSL}(2,\mathbb{C})$, we have defined the parameters $\mathbf{hl}(C) \in \mathbb{C}_+/2k\pi\mathbb{Z}$
and $s(C) \in \mathbb{C}/(\mathbf{hl}(C)\mathbb{Z} + 2\pi i \mathbb{Z})$. The collection of pairs $(\mathbf{hl}(C), s(C))$, $C \in \mathcal{C}$, is
called the reduced Fenchel–Nielsen coordinates. We observe that a representation
$\rho\colon \pi_1(S) \to \mathrm{PSL}(2,\mathbb{C})$ is Fuchsian if and only if all the coordinates $(\mathbf{hl}(C), s(C))$ are
real. This is well known (see [3]) when $\mathcal{C}$ is a pants decomposition. The same is true
when $\mathcal{C}$ is a generalized pants decomposition. This follows directly from Definition 3.1
and the above three conditions we impose on ρ.

The following elementary proposition (see [3]) states that although a representation
$\rho\colon \pi_1(S) \to \mathrm{PSL}(2,\mathbb{C})$ is not uniquely determined by its reduced Fenchel–Nielsen
coordinates, it can be in a unique way embedded in a holomorphic family of represen-
tations (uniquely means that there is a unique holomorphic family of representations
such that ρ can embedded in this family as described in the following lemma).

Proposition 3.1 *Fix a closed topological surface S with a generalized pants decom-
position $\mathcal{C}$. Let $z \in \mathbb{C}_+^{\mathcal{C}}$ and $w \in \mathbb{C}^{\mathcal{C}}$ denote complex parameters. Then there exists a
holomorphic (in (z,w)) family of representations*

$$\rho_{z,w}\colon \pi_1(S) \to \mathrm{PSL}(2,\mathbb{C}),$$

*such that $\mathbf{hl}(C) = z(C)$ (mod $2\pi i \mathbb{Z}$) and $s(C) = w(C)$ (mod $\mathbf{hl}(C)\mathbb{Z} + 2\pi i \mathbb{Z}$).
Moreover, for any $(z_0, w_0) \in \mathbb{C}_+^{\mathcal{C}} \times \mathbb{C}^{\mathcal{C}}$, the family of representations $\rho_{z,w}$ is uniquely
determined by the representation ρ_{z_0,w_0}.*

The representation $\rho_{z,w}$ is Fuchsian if and only if both z and w are real, that is $z \in \mathbb{R}_+^{\mathcal{C}}$
and $w \in \mathbb{R}^{\mathcal{C}}$. In this case the group $\rho_{z,w}(\pi_1(S))$ is of course discrete. Moreover,
in [4] it has been proved that all quasifuchsian representations (up to conjugation in
$\mathrm{PSL}(2,\mathbb{C})$) of $\pi_1(S)$ correspond to some neighborhood of the set $\mathbb{R}_+^{\mathcal{C}}$ and $\mathbb{R}^{\mathcal{C}}$. But in

general, little is known for which choice of parameters z, w the group $\rho_{z,w}(\pi_1(S))$ will be discrete. In the next subsection we prove the following result in this direction. Start with a nearly Fuchsian group $G < \mathrm{PSL}(2, \mathbb{C})$. We obtain a new group $G_1 < \mathrm{PSL}(2, \mathbb{C})$ from G by bending (by some definite angles) along some sparse equivariant collection of geodesics whose endpoints are in the limit set of G. Then the new group G_1 is also quasifuchsian (although it is not nearly Fuchsian anymore).

3.2 Small deformations of a sparsely bent pleated surface

We let S continue to denote a closed surface with a generalized pants decomposition $\mathcal{C}$, and we fix a holomorphic family of representations $\rho_{z,w}$ as in Proposition 3.1. We set $G(z, w) = \rho_{z,w}(\pi_1(S))$.

Let $\mathcal{C}_0 \subset \mathcal{C}$ denote a subcollection of curves. For $z \in \mathbb{R}_+^{\mathcal{C}}$ and $w \in \mathbb{R}^{\mathcal{C}}$, we let $S_{z,w}$ denote the Riemann surface isomorphic to $\mathbb{H}^2/G(z, w)$, and on $S_{z,w}$ we identify the curves from $\mathcal{C}$ with the corresponding geodesics representatives. By $\mathcal{K}(S_{z,w})$ we denote the largest number so that the collection of collars (of width $\mathcal{K}(S_{z,w})$) around the curves from $\mathcal{C}_0$ is disjoint on $S_{z,w}$. For each $C \in \mathcal{C}_0$, we choose a number $-(3/4)\pi < \theta_C < (3/4)\pi$ (for each curve $C \in (\mathcal{C} \setminus \mathcal{C}_0)$ we set $\theta_C = 0$).

The purpose of this subsection is to prove the following theorem.

Theorem 3.1 *There exist constants $K > 1$ and $D > 0$ such that the following holds. Let $z_0 \in \mathbb{R}_+^{\mathcal{C}}$ and $w_0 \in \mathbb{R}^{\mathcal{C}}$, and $z_1 \in \mathbb{C}_+^{\mathcal{C}}$ and $w_1 \in \mathbb{C}^{\mathcal{C}}$ be such that the representation $\rho = \rho_{z_1,w_1} \circ \rho_{z_0,w_0}^{-1} \colon G(z_0, w_0) \to G(z_1, w_1)$, is K–quasifuchsian. Set $z_2 = z_1$ and $w_2 = w_1 + i\theta_C$. If $\mathcal{K}(S_{z_0,w_0}) \geq D$, then the representation $\rho_{z_2,w_2} \colon \pi_1(S) \to \mathrm{PSL}(2, \mathbb{C})$ is K_1–quasifuchsian, where K_1 depends only on K and D.*

The following lemma is elementary.

Lemma 3.1 *Let $0 \leq \theta_0 < \pi$ and $B_0 \geq 1$. There exist constants $L(\theta_0, B_0) > 0$ and $D(\theta_0, B_0) > 0$ such that the following holds. Let $I \subset \mathbb{R}$ be an interval that is partitioned into intervals I_j, $j = 1, \ldots, k$. Let $\psi \colon I \to \mathbb{H}^3$ be a continuous map, such that ψ maps each I_j onto a geodesic segment and the restriction of ψ on I_j is B_0–bilipschitz. Assume in addition that the bending angle between two consecutive geodesic intervals $\psi(I_j)$ and $\psi(I_{j+1})$ is at most θ_0. If the length of every I_j is at least $D(\theta_0, B_0)$ then ψ is $L(\theta_0, B_0)$–bilipschitz.*

Let $\psi \colon I \to \mathbb{H}^3$ be a C^1 map, where $I \subset \mathbb{R}$ is a closed interval. For $x \in I$ let $v(x) \in T^1 I$ denote the unit vector that points toward $+\infty$. Let $\delta > 0$. We say that the

map ψ is δ–nearly geodesic if for every $x, y \in I$ such that $x < y \le x + 1$, we have that the angle between the vector $\psi_*(v(x))$ and the oriented geodesic segment from $\psi(x)$ to $\psi(y)$ is at most δ.

Clearly, every 0–nearly geodesic map is an isometry, and a sequence of normalized δ_n–nearly geodesic maps converges (uniformly on compact sets) in the C^1 sense to an isometry, when $\delta_n \to 0$. The following lemma is a generalization of the previous one.

Lemma 3.2   *There exist universal constants $L, D, \delta > 0$, such that the following holds. Suppose that I is partitioned into intervals I_j, $j = 1, \ldots, k$, and let $\psi : I \to \mathbb{H}^3$ be a continuous map, whose restriction on every closed subinterval I_j is C^1 and δ–nearly geodesic. Assume that the bending angle between two consecutive curves $\psi(I_j)$ and $\psi(I_{j+1})$ is at most $3\pi/4$ (by the bending angle between two C^1 curves we mean the appropriate angle determined by the two tangent vectors at the point where the two curves meet). If the length of every I_j is at least D then ψ is L–bilipschitz.*

Proof   Choose any two numbers $3\pi/4 < \theta_0 < \pi$ and $B_0 > 1$. Assuming that $D > D(\theta_0, B_0)$ we can partition each I_j into subintervals of length between $D(\theta_0, B_0)$ and $2D(\theta_0, B_0)$. Replacing each I_j with these new intervals we obtain the new partition of I into intervals J_i, where each J_i has the length between $D(\theta_0, B_0)$ and $2D(\theta_0, B_0)$. Let $\phi : I \to \mathbb{H}^3$ be the continuous map that agrees with ψ at the endpoints of all intervals J_i, and such that the restriction of ϕ to each J_i maps J_i onto a geodesic segment in $\mathbb{H}^3$, and is affine (the map ϕ either stretches or contracts distances by a constant factor on a given J_i).

Next, since we have the upper bound $2D(\theta_0, B_0)$ on the length of each interval J_i, we can choose $\delta > 0$ small enough such that the bending angle between two consecutive geodesic segments $\phi(J_i)$ and $\phi(J_{i+1})$ is at most θ_0. Also, by choosing δ small we can arrange that the map $\phi \circ \psi^{-1}$ is 2–bilipschitz (the same statement holds if we replace 2 by any other number greater than 1). By the previous lemma the map ϕ is $L(\theta_0, B_0)$–bilipschitz. Then the map ψ is $2L(\theta_0, B_0)$–bilipschitz. We take $L = 2L(\theta_0, B_0)$, and $D = 2D(\theta_0, B_0)$, and the lemma is proved. $\square$

We are now ready to prove Theorem 3.1.

Proof   Recall that $f : \partial\mathbb{H}^2 \to \partial\mathbb{H}^3$ is a K–quasiconformal map that conjugates $G(z_0, w_0)$ to $G(z_1, w_1)$. Let $\tilde{f} : \mathbb{H}^2 \to \mathbb{H}^3$ denote the Douady–Earle extension of f.

Remark   Usually the Douady–Earle extension refers to the barycentric extensions of a homeomorphism $f : \partial S^1 \to \partial S^1$ (see [2]). In the same paper (see [2, Section 11])

Douady and Earle have shown that in a similar vein one defines the barycentric extension of any homeomorphism $f\colon \partial S^k \to \partial S^k$, for any $k > 0$. Similarly one can define the barycentric extension $\widetilde{f}\colon \mathbb{H}^m \to \mathbb{H}^n$ of any continuous map $f\colon \partial S^{m-1} \to \partial S^{n-1}$. Alternatively, given a K-quasiconformal map $f\colon \partial \mathbb{H}^2 \to \partial \mathbb{H}^3$ we first extend f to an equivariant K–quasiconformal map $\overline{f}\colon \partial \mathbb{H}^3 \to \partial \mathbb{H}^3$ and then take the corresponding Douady–Earle extension of $\widetilde{f}\colon \mathbb{H}^3 \to \mathbb{H}^3$ of $\overline{f}$ (as defined by Douady and Earle [2]). The restriction of $\widetilde{f}$ to $\mathbb{H}^2$ is also called $\widetilde{f}$.

Then $\widetilde{f}$ is δ–nearly geodesic (this means that the restriction of $\widetilde{f}$ to every geodesic segment is δ–nearly geodesic in the sense of the above definition) for some $\delta = \delta(K)$, and $\delta(K) \to 0$, when $K \to 1$.

If we assume that $\mathcal{K}(S_{z_0,w_0})$ is large enough, by adjusting $\widetilde{f}$, we can arrange that $\widetilde{f}$ is then C^∞ mapping that maps the geodesics in $\mathbb{H}^2$ that are lifts of the geodesics from $\mathcal{C}_0$ onto the corresponding geodesics in $\mathbb{H}^3$, and ensure that $\widetilde{f}$ is 2δ–nearly geodesic. Moreover, we can arrange that $\widetilde{f}$ is conformal at every point of every geodesic γ that is a lift of a curve from $\mathcal{C}_0$.

We construct the map $\widetilde{g}\colon \mathbb{H}^2 \to \mathbb{H}^3$ that conjugates $G(z_0, w_0)$ to $G(z_2, w_2)$ as follows. Let M be a component of the set $S_{z_0,w_0} \setminus \mathcal{C}_0$, and let $\widetilde{M} \subset \mathbb{H}^2$ denote its universal cover, that is, $\widetilde{M}$ is an ideal polygon with infinitely many sides in $\mathbb{H}^2$, whose sides are lifts of the geodesics from $\mathcal{C}_0$ that bound M. We set $\widetilde{g} = \widetilde{f}$ on $\widetilde{M}$.

Let $\widetilde{M}_1 \subset \mathbb{H}^2$ be the universal cover of some other component M_1 of the set $S_{z_0,w_0} \setminus \mathcal{C}_0$. Let γ denote a lift of a geodesic $C \in \mathcal{C}_0$, and assume that the polygons $\widetilde{M}$ and $\widetilde{M}_1$ are glued to each other along γ (that is, C is in the boundary of both M and M_1). Let $R(\theta_C) \in \mathrm{PSL}(2, \mathbb{C})$, denote the rotation about $\widetilde{g}(\gamma)$ for the angle θ_C. We define $\widetilde{g}$ on $\widetilde{M}_1$ by letting $\widetilde{g} = R(\theta_C) \circ \widetilde{f}$. We then define $\widetilde{g}$ inductively on the rest of $\mathbb{H}^2$.

Clearly $\widetilde{g}$ conjugates $G(z_0, w_0)$ to $G(z_2, w_2)$. Let $x \in \gamma$, and $v(x)$ a nonzero vector that is orthogonal to γ. Since $|\theta_C| \leq (3/4)\pi$, and since $\widetilde{f}$ is differentiable at x, it follows that the bending angle between the vectors $\widetilde{g}_*(v(x))$ and $\widetilde{g}_*(-v(x))$ is at most $(3/4)\pi$. If $u(x)$ is any other vector at x, since $\widetilde{f}$ is conformal at x, it follows that the bending angle between the vectors $\widetilde{g}_*(u(x))$ and $\widetilde{g}_*(-u(x))$ is at most as big as the bending angle between the vectors $\widetilde{g}_*(v(x))$ and $\widetilde{g}_*(-v(x))$. Therefore, the restriction of the map $\widetilde{g}$ on every geodesic segment satisfies the assumptions of Lemma 3.2. It follows that $\widehat{g}$ is L–bilipschitz, where L depends only on K and D. Therefore the representation $\rho_{z_2,w_2}\colon \pi_1(S) \to \mathrm{PSL}(2, \mathbb{C})$ is K_1–quasifuchsian, where K_1 depends only on K and D. $\qquad\square$

3.3 Convex hulls and pleated surfaces

In this subsection we digress from the notions of generalized pants decompositions and Fenchel–Nielsen coordinates, to prove a preliminary lemma about hyperbolic convex hulls of quasicircles.

Let λ be a discrete geodesic lamination in $\mathbb{H}^2$, and let $\mathcal{K}(\lambda)$ denote the largest number such that for every small $\epsilon > 0$, the collection of collars (crescent in $\mathbb{H}^2$) of width $\mathcal{K}(\lambda) - \epsilon$ around the leafs of λ is disjoint in $\mathbb{H}^2$. Let μ denote a real valued measure on λ. By $\iota_{\lambda,\mu} = \iota \colon \mathbb{H}^2 \to \mathbb{H}^3$, we denote the corresponding pleating map. As usual, by $\iota(\lambda)$ we denote the collection of geodesics in $\mathbb{H}^3$ that are images of geodesics from λ under ι. If the map ι is L–bilipschitz then ι extends continuously to a K–quasiconformal map $f \colon \partial \mathbb{H}^2 \to \partial \mathbb{H}^3$, for some $K = K(L)$. In this case, let $W \subset \mathbb{H}^3$ denote the convex hull of the quasicircle $\iota(\partial \mathbb{H}^2)$. The convex hull W has two boundary components which we denote by $\partial_1 W$ and $\partial_2 W$. We prove the following lemma.

Lemma 3.3 *There exist universal constants $C_1, \delta_1 > 0$, with the following properties. Assume that $\mathcal{K}(\lambda) > C_1$, and that $\pi/4 \le |\mu(l)| \le 3\pi/4$, for every $l \in \lambda$. Then for every geodesic $\gamma \subset W$ the following holds:*

(1) *If $\gamma \in \iota(\lambda)$, then for every point $p \in \gamma$, the inequality*

$$\max_{i=1,2} d(p, \partial_i W) > \delta_1$$

   holds.

(2) *If γ does not belong to $\iota(\lambda)$, then for some point $p \in \gamma$, the inequality $\max_{i=1,2} d(p, \partial_i W) < \delta_1/3$ holds.*

Compare this lemma with [5, Lemma 4.2].

Proof It follows from Lemma 3.1 that for C_1 large enough, the pleating map ι is L–bilipschitz for some universal constant $L > 1$. Observe that $\iota(\mathbb{H}^2) \subset W$. Moreover, there is a constant $M_0 > 0$, that depends only on L, such that for every $p \in W$ we have $d(p, \iota(\mathbb{H}^2)) < M_0$

We choose $\delta_1 > 0$ as follows. Let P_0 be the pleated surface in $\mathbb{H}^3$ that has a single bending line γ_0, and with the bending angle equal to $\pi/4$. Then P_0 is bounded by a quasicircle at $\partial \mathbb{H}^3$. Denote by W_0 the convex hull of this quasicircle and let $\partial_i(W_0)$, $i = 1, 2$, denote the two boundary components of W_0. Then there exists $\delta_1 > 0$ such that for every point $p \in \gamma_0$, we have $\max_{i=1,2} d(p, \partial_i W_0) > 2\delta_1$. Observe that γ_0 belongs to exactly one of the convex hull boundaries $\partial_1 W_0$ and $\partial_2 W_0$, so one of the numbers $d(p, \partial_1 W_0)$ and $d(p, \partial_2 W_0)$ is zero and the other one is larger than $2\delta_1$.

Assume that the first statement of the lemma is false. Then there exists a sequence of measured laminations (λ_n, μ_n) with the property $\mathcal{K}(\lambda_n) \to \infty$, and there are geodesics $l_n \in \lambda_n$, and points $p_n \in \gamma_n = \iota_n(l_n)$, such that the inequality

$$(4) \qquad \max_{i=1,2} d(p_n, \partial_i W_n) \leq \delta_1,$$

holds. We may assume that $p_n = p$, and $\gamma_n = \gamma$, for every n, where p and γ are fixed. Since ι_n is L–bilipschitz, after passing to a subsequence if necessary, the sequence ι_n converges (uniformly on compact sets) to a pleating map ι_∞. The pleating map ι_∞ corresponds to the pleating surface P_∞, that has a single bending line γ_∞, with the bending angle at least $\pi/4$. Then W_n converges to W_∞ uniformly on compact sets in $\mathbb{H}^3$, where W_∞ is the convex hull of the quasicircle that bounds P_∞. It follows that $d(p_n, \partial_i W_n) \to d(p, \partial_i W_\infty)$. We may assume that $\gamma_\infty = \gamma_0$, where γ_0 is the bending line of the pleated surface P_0 defined above. Then we have $\max_{i=1,2} d(p, \partial_i W_\infty) \geq \max_{i=1,2} d(p, \partial_i W_0) > 2\delta_1$. But this contradicts (4).

We now prove the second statement of the lemma. Let γ be a geodesic in W that is not in $\iota(\lambda)$. Then we can find a point $p \in \gamma$, such that $d(p, \iota(\lambda)) > \mathcal{K}(\lambda)$. Assuming that the second statement is false, we again produce a sequence λ_n with $\mathcal{K}(\lambda_n) \to \infty$, and such that for some sequence of geodesics $\gamma_n \subset W_n$, that do not belong to $\iota(\lambda_n)$, and all the points $p \in \gamma_n$, the inequality

$$(5) \qquad \max_{i=1,2} d(p, \partial_i W_n) \geq \delta_1/3,$$

holds for n large enough. By the previous discussion, there exists a sequence of points $p_n \in \gamma_n$, such that $d(p_n, \iota_n(\lambda_n)) > \mathcal{K}(\lambda_n)$.

Let $q_n \in \iota_n(\mathbb{H}^2)$ be points such that $d(p_n, q_n) < M_0$, where M_0 is the constant defined at the beginning of the proof. Let $z_n \in \mathbb{H}^2$, such that $q_n = \iota(z_n)$. We may assume that $z_n = i \in \mathbb{H}^2$ and $q_n = q$, for some point q that we fix. Then $p_n \to p$, where $d(p, q) \leq M_0$. Moreover, since $\mathcal{K}(\lambda_n) \to \infty$, the pleating maps $\iota(\lambda_n)$ converge to an isometry uniformly on compact sets in $\mathbb{H}^2$. In particular, the sequence of convex hulls W_n converges to a geodesic plane uniformly on compact sets, and therefore $d(p_n, \partial_i W_n) \to 0$. But this contradicts (5), and thus we have completed the proof of the lemma. $\qquad\square$

3.4 (ϵ, R) Skew pants

We let S continue to denote a closed surface with a generalized pants decomposition $\mathcal{C}$, and we fix a holomorphic representations $\rho_{z,w}$ as in Proposition 3.1.

Let $\mathcal{C}_0 \subset \mathcal{C}$ denote a subcollection of curves, and for each $C \in \mathcal{C}_0$ we choose a number $-(3/4)\pi < \theta_C < (3/4)\pi$ (for each curve $C \in (\mathcal{C} \setminus \mathcal{C}_0)$ we set $\theta_C = 0$).

For $C \in \mathcal{C}$, let $\zeta_C, \eta_C \in \mathbb{D}$, where $\mathbb{D}$ denotes the unit disc in the complex plane. Let $\tau \in \mathbb{D}$ denote a complex parameter and let $t \in \{0, 1\}$. Fix $R > 1$, and let $z \colon \mathbb{D} \to \mathbb{C}_+^{\mathcal{C}}$ and $w \colon \mathbb{D} \to \mathbb{C}^{\mathcal{C}}$ be the mappings given by

$$z(C)(\tau) = \frac{R}{2} + \frac{\tau \zeta_C}{2},$$

$$w(C)(\tau, t) = 1 + it\theta_C + \frac{\tau \eta_C}{R}.$$

The maps $z(\tau)$ and $w(\tau, t)$ are complex linear, and therefore holomorphic in τ and t. Therefore the induced family of representations $\rho_{\tau,t} = \rho_{z(\tau), w(\tau, t)}$ is holomorphic in τ and t. Note that $\rho_{\tau,t}$ depends on R, ζ_C, η_C and θ_C, but we suppress this.

The representation $\rho_{0,0}$ is Fuchsian. Let S_0 denote the Riemann surface isomorphic to the quotient $\mathbb{H}^2/\rho_{0,0}(\pi_1(S))$ (we also equip S_0 with the corresponding hyperbolic metric). Let $\mathcal{K}(\rho_{0,0})$ denote the largest number so that the collection of collars (of width $\mathcal{K}(\rho_{0,0})$) around the curves from $\mathcal{C}_0$ is disjoint on S_0.

The representation $\rho_{0,1}$ is not Fuchsian (unless $\theta(\mathcal{C}_0) = 0$), and the following theorem gives a sufficient condition for it to be quasifuchsian.

We adopt the following notation. Let $G(\tau, t) = \rho_{\tau,t}(\pi_1(S))$. If $G(\tau, t)$ is a quasifuchsian group we let $f_{\tau,t} \colon \partial \mathbb{H}^2 \to \partial \mathbb{H}^3$, denote the quasiconformal map that conjugates $G(0,0)$ to $G(\tau, t)$. The following theorem is a generalization of [3, Theorem 2.2] (see Theorem 3.4 below). Assuming the above notation, we have:

Theorem 3.2 *There exist universal constants $\widehat{R}, \widehat{\epsilon}, M > 0$, such that the following holds. If $\mathcal{K}(\rho_{0,0}) > M$, then for every $R \geq \widehat{R}$ and $|\tau| < \widehat{\epsilon}$, and any choice of constants $\eta_C, \zeta_C \in \mathbb{D}$, and $-(3/4)\pi < \theta_C < (3/4)\pi$, for $C \in \mathcal{C}_0$, the group $G(\tau, 1)$ is quasifuchsian and the induced quasiconformal map $f_{\tau,1} \circ (f_{0,1})^{-1}$ (that conjugates $G(0,1)$ to $G(\tau, 1)$), is $K(\tau)$–quasiconformal, where*

$$K(\tau) = \frac{\widehat{\epsilon} + |\tau|}{\widehat{\epsilon} - |\tau|}.$$

Let $\mathcal{C}_0(\tau, t)$ denote the collection of axes of elements of the form $\rho_{\tau,t}(c)$, where $c \in \pi_1(S)$ and c belongs to the conjugacy class of some curve $C \in \mathcal{C}_0$. Then by definition, the set $\mathcal{C}_0(\tau, t)$ is invariant under the group $G(\tau, 1)$. Next, we prove that $\mathcal{C}_0(\tau, 1)$ is invariant under any Möbius transformation from $\mathrm{PSL}(2, \mathbb{C})$ that preserves the limit set of $G(\tau, 1)$. The following theorem is the main result of this section.

Theorem 3.3 *There exist constants $\hat{\epsilon}_1$, $M_1 > 0$, with the following properties. Assume that $\mathcal{K}(\rho_{0,0}) > M_1$ and let $|\tau| < \hat{\epsilon}_1$. If $T \in \mathrm{PSL}(2, \mathbb{C})$, is a Möbius transformation that preserves the limit set of $G(\tau, 1)$, then the set of geodesics $\mathcal{C}_0(\tau, 1)$ is invariant under T.*

Compare this theorem with [5, Lemma 4.2].

Proof Let $W(\tau, t)$ denote the convex hull of the limit set of $G(\tau, t)$. It follows from Lemma 3.3 that for $\mathcal{K}(\rho_{0,0})$ large enough, the following holds:

(1) For every $\gamma \in \mathcal{C}_0(0, 1)$ and $p \in \gamma$, the inequality

$$\max_{i=1,2} d(p, \partial_i W(0, t)) > \delta_1$$

   holds.

(2) For every $\gamma \subset W(0, 1)$, there exists $p \in \gamma$ such that

$$\max_{i=1,2} d(p, \partial_i W(0, 1)) < \frac{\delta_1}{2}.$$

Then by Theorem 3.2 we can choose $\hat{\epsilon}_1$ small enough so that for $|\tau| < \hat{\epsilon}_1$, the constant $K(\tau)$ (from Theorem 3.2) is close enough to 1, so that the following holds:

(1) For every $\gamma \in \mathcal{C}_0(\tau, 1)$ and $p \in \gamma$, the inequality

$$\max_{i=1,2} d(p, \partial_i W(0, t)) > \frac{4\delta_1}{5}$$

   holds,

(2) For every $\gamma \subset W(0, 1)$, there exists $p \in \gamma$ such that

$$\max_{i=1,2} d(p, \partial_i W(0, 1)) < \frac{2\delta_1}{3}.$$

Then any Möbius transformation $A \in \mathrm{PSL}(2, \mathbb{C})$ that preserves $W(\tau, 1)$ will also preserve the set $\mathcal{C}(\tau, 1)$. This proves the theorem. □

3.5 A proof of Theorem 3.2

We need to prove that $G(\tau, 1)$ is a quasifuchsian group. The last estimate in Theorem 3.2 then follows from the fact that a holomorphic map from the unit disc into the Teichmüller space of a Riemann surface is a contraction with respect to the hyperbolic metric on the unit disc and the Teichmüller metric.

Theorem 3.4 [3, Theorem 2.2]   *There exist universal constants $\widehat{R}, \widehat{\epsilon}$, such that the following holds. For every $R \geq \widehat{R}$ and $|\tau| < \widehat{\epsilon}$, and any choice of constants $\eta_C, \zeta_C \in \mathbb{D}$, the group $G(\tau, 0)$ is quasifuchsian, and the induced quasiconformal map $f_{\tau,0}$ that conjugates $G(0,0)$ to $G(\tau, 0)$, is $K(\tau)$–quasiconformal, where*

$$K(\tau) = \frac{\widehat{\epsilon} + |\tau|}{\widehat{\epsilon} - |\tau|}.$$

The group $G(\tau, 1)$ is obtained from the group $G(\tau, 0)$, by bending along the lifts of curves $C \in \mathcal{C}_0$, for the angle θ_C. It follows from Theorem 3.1 that the group $G(\tau, 1)$ is quasifuchsian if $\mathcal{K}(\rho_{0,0}) > C$, and if the map $f_{\tau,0}$ is K–quasiconformal, where K is close enough to 1. But it follows from Theorem 3.4 that for $|\tau|$ small enough this will be the case. This proves Theorem 3.2.

4   The lower bound

4.1   Amalgamating two representations

Let S denote a closed surfaces with generalized pants decompositions $\mathcal{C}$, and let $\rho \colon \pi_1(S) \to \mathrm{PSL}(2, \mathbb{C})$ denote an admissible (in sense of Definition 3.1) representation with the reduced Fenchel–Nielsen coordinates satisfying the inequalities

$$\left| \mathbf{hl}(C) - \frac{R}{2} \right| \leq \epsilon,$$

$$|s(C) - 1| \leq \frac{\epsilon}{R},$$

for some $\epsilon, R > 0$, and $C \in \mathcal{C}$. We say that such a representation is (ϵ, R)–good.

Let $\mathbf{M}^3$ denote a closed hyperbolic manifold such that $\mathbf{M}^3 = \mathbb{H}^3 / \Gamma$ for some Kleinian group Γ. In [3] we proved that one can find many (ϵ, R)–good representations $\rho \colon \pi_1(S) \to \Gamma$, for a given $\epsilon > 0$ and R large enough. Moreover, if $A \in \Gamma$ has the translation length $\mathbf{l}(A)$ satisfying the inequality $|\mathbf{l}(A) - R| \leq \epsilon/2$, then we can find such ρ so that A is in the image of ρ. From now on we assume that such $A \in \Gamma$ is primitive, that is A is not equal to an integer power of another element of Γ.

In particular, it follows from [3, Section 4] (the statements about the equidistribution of (ϵ, R)–good pairs of skew pants around a given closed curve in $\mathbf{M}^3$ whose length is ϵ close to R) that we can find two (ϵ, R)–good representations $\rho(i) \colon \pi_1(S(i)) \to \Gamma$, $i = 1, 2$, where $S(1)$ and $S(2)$ are two closed surfaces with pants decompositions $\mathcal{C}(i)$, and two pairs of pants Π_i^+ and Π_i^- with the following properties:

- There are two oriented, degree one curves $C(i) \in \mathcal{C}(i)$, and $c(i) \in \pi_1(S(i))$ in the conjugacy classes of $C(1)$ and $C(2)$ respectively, such that $\rho(1)(C(1)) = \rho(2)(C(2)) = [A]$, where $[A]$ is the conjugacy class of a given primitive element $A \in \Gamma$, whose translation length $\mathbf{l}(A)$ satisfies the inequality $|\mathbf{l}(A) - R| \le \epsilon/2$.

- Let γ denote the closed geodesic corresponding to A. There exist two pars of skew pants Π_i^+ and Π_i^- in $\rho(i)(\pi_1(S(i)))$ such that γ is positively oriented boundary component of Π_i^+ and negatively oriented for Π_i^-, and recalling the notation from [3] we have the inequality

(6)
$$\left| \mathrm{foot}_\gamma(\Pi_2^+) - \mathrm{foot}_\gamma(\Pi_1^-) - \frac{\pi}{2} \right| \le \frac{\epsilon}{R}.$$

After replacing $S(1)$ and $S(2)$ with appropriate finite degree covers if necessary, we may assume in addition to the above two conditions the following also hold:

- The curves $C(1)$ and $C(2)$ are nonseparating simple closed curves in $S(1)$ and $S(2)$ respectively.

- The surfaces $S(1)$ and $S(2)$ have the same genus.

- By Proposition 3.1 the representation $\rho(i)$ can be embedded in the holomorphic family of representations $\rho_{\tau,t}(i)$. We may assume that $\mathcal{K}(\rho_{0,0}(S(i))) > C_1$, $i = 1, 2$, where C_1 is the constant from Theorem 3.3.

We now fix such two representations $\rho(1)$ and $\rho(2)$, surfaces $S(1)$ and $S(2)$, and the two oriented curves $C(1)$ and $C(2)$ (we also fix the corresponding primitive element $A \in \Gamma$).

Let $i \in \{1, 2\}$. For $n > 1$, let $S_n(1)$ and $S_n(2)$ denote two primitive degree n covers of $S(1)$ and $S(2)$ respectively (a finite cover of a surface is primitive if it does not factor through an intermediate cover), such that for some $1 \le k \le (n-1)$, the curves $C(1)$ and $C(2)$ have two degree k lifts $C_n(1)$ and $C_n(2)$. Then $C_n(1)$ and $C_n(2)$ are two oriented, nonseparating simple closed curves in $S_n(1)$ and $S_n(2)$ respectively. We then have the two induced representations $\rho_n(i)\colon \pi_1(S_n(i)) \to \Gamma$, that also satisfy the above five conditions, except that

$$\rho_n(1)(\pi_1(S_n(1))) \cap \rho_n(2)(\pi_1(S_n(2))) = \{A^k\}.$$

We amalgamate them as follows. Cut the surface $S_n(i)$ along $C_n(i)$, to get two topological surfaces $\bar{S}_n(i)$, $i = 1, 2$, each having two boundary components $C_n^1(i)$ and $C_n^2(i)$. We glue together the surfaces $\bar{S}_n(1)$ and $\bar{S}_n(2)$ by gluing $C_n^j(1)$ to $C_n^j(2)$, $j = 1, 2$, and obtain a closed topological surface S_n (this is well defined up to a twist by $\Re(\mathbf{l}(A))$ which has a period k). The surface S_n has the induced generalized pants

decomposition $\mathcal{C}_n$. The pair of curves $C_n^1(1)$ and $C_n^1(2)$ that were glued together produce a closed curve C_n^1 in S_n. Similarly, the pair of curves $C_n^2(1)$ and $C_n^2(2)$ that were glued together produce a closed curve C_n^2 in S_n. We set $\mathcal{C}_{0,n} = \{C_n^1, C_n^2\}$.

Then there is the induced representation $\rho_n\colon \pi_1(S_n) \to \Gamma$. We orient the curves C_n^1 and C_n^2 such that for any choice of $c_i \in \pi_1(S_n)$, where c_i is in the conjugacy class of C_n^i, we have that both $\rho_n(c_1)$ and $\rho_n(c_2)$ are in the conjugacy class of A^k in Γ.

The representation ρ_n has the reduced Fenchel–Nielsen coordinates satisfying the relations

$$\left| \mathbf{hl}(C) - \frac{R}{2} \right| \leq \epsilon,$$

$$|s(C) - 1| \leq \frac{\epsilon}{R},$$

if C does not belong to $\mathcal{C}_{0,n}$, and

$$\left| s(C) - \left(1 + i\,\frac{\pi}{2}\right) \right| \leq \frac{\epsilon}{R},$$

if $C \in \mathcal{C}_{0,n}$.

It follows from Theorem 3.2 that for ϵ small enough and R large enough, the group $\rho_n(\pi_1(S_n))$ is quasifuchsian. In the remainder of this subsection we prove that the group $\rho_n(\pi_1(S_n))$ is a maximal subgroup of Γ.

First we prove a preliminary lemma. Let $\overline{S}$ be a surface with boundary components C_+ and C_-, oriented so that $\overline{S}$ is on the left of C_+ and on the right of C_-. We say that $f\colon \overline{S} \to \mathbf{M}^3$ is rejoinable if the restrictions of f to C^+ and C_- respectively are freely homotopic in $\mathbf{M}^3$. We say $(f, \overline{S})$ is geodesically rejoinable if $f|_{C_+}$ and $f|_{C_-}$ map to the same closed geodesic in $\mathbf{M}^3$. In this case we say a rejoining of $(f, \overline{S})$ is a homeomorphism $h\colon C_+ \to C_-$ such that $f \circ h = f$, and we say $(f, \overline{S}/h)$ is $\overline{S}$ rejoined by h.

Lemma 4.1 *If $(f, \overline{S})$, and $(g, \overline{T})$ are (geodesically) rejoinable surfaces, and $\pi\colon \overline{S} \to \overline{T}$ is a finite cover such that $g \circ \pi$ is homotopic to f, then for any rejoining h of $(f, \overline{S})$ we can find a rejoining k of $(g, \overline{T})$ such that $(f, \overline{S})$ rejoined by h covers $(g, \overline{T})$ rejoined by k.*

Proof Left to the reader.   □

The following theorem is a corollary of Theorem 3.3. We adopt the following definition. Let $f\colon S \to \mathbf{M}^3$ be a map such that $f(S)$ is a quasifuchsian surface in $\mathbf{M}^3$, and

let $\mathcal{C}_0$ denote a collection of disjoint simple closed curves on S. We say that f is bent along each curve of $\mathcal{C}_0$ and nearly locally isometric on $S \setminus \mathcal{C}_0$ if the induced map $f_*: \pi_1(S) \to \Gamma$ is of the form $\rho_{\tau,1}$ for some $|\tau| \le \hat{\epsilon}$.

Theorem 4.1 *Let S be a closed surface. Suppose that $f: S \to \mathbf{M}^3$ is a π_1–injective and quasifuchsian, and $\mathcal{C}_0$ is a collection of disjoint simple closed curves on S, such that f is bent along each curve of $\mathcal{C}_0$ and nearly locally isometric on $S \setminus \mathcal{C}_0$. Suppose that $f = g \circ \pi$, where $\pi: S \to Q$ is a covering, and $g: Q \to \mathbf{M}^3$ is π_1–injective and quasifuchsian. Then we can find a collection of simple closed curves $\hat{\mathcal{C}}_0$ on Q such that $\mathcal{C}_0 = \pi^{-1}(\hat{\mathcal{C}}_0)$.*

Proof We get a discrete lamination $\tilde{\mathcal{C}}_0$ on $\mathbb{H}^2$, which we push forward by $\tilde{f} = \tilde{g}$ to $\mathbb{H}^3$. We find a homomorphism $\sigma: \mathrm{Deck}(\mathbb{H}^2/Q) \to \Gamma$ such that $\tilde{f}(\gamma(x)) = \sigma(\gamma)(\tilde{f}(x))$ for every $x \in \mathbb{H}^2$ and $\gamma \in \mathrm{Deck}(\mathbb{H}^2/Q)$.

We let $G = \sigma(\mathrm{Deck}(\mathbb{H}^2/Q))$, and $H = \sigma(\mathrm{Deck}(\mathbb{H}^2/S)) < G$. Then $[G : H] < \infty$, and G and H are quasifuchsian groups, and they have the same limit set, so by Theorem 3.3 every element of G maps $\tilde{g}(\tilde{\mathcal{C}}_0)$ to itself. Hence $\mathrm{Deck}(\mathbb{H}^2/Q)$ maps $\tilde{\mathcal{C}}_0$ to itself, so $\tilde{\mathcal{C}}_0$ is a lift of $\hat{\mathcal{C}}_0$ on Q, and hence $\mathcal{C}_0$ is. $\square$

Theorem 4.2 *The quasifuchsian group $\rho_n(\pi_1(S_n)) < \Gamma$ is a maximal surface subgroup of Γ, that is, if $\rho_n(\pi_1(S_n)) < G$ for a surface subgroup $G < \Gamma$, then $G = \rho_n(\pi_1(S_n))$.*

Proof For simplicity let $G_n = \rho_n(\pi_1(S_n))$ and $G(1) = \rho(1)(\pi_1(S(1)))$. Also set $G_n(1) = \rho_n(\pi_1(\bar{S}_n(1)))$, where we consider $\pi_1(\bar{S}_n(1))$ as a subgroup of $\pi_1(S_n)$.

Let $f_n: S_n \to \mathbf{M}^3$ denote the continuous map that corresponds to the representation ρ_n. We claim that $f_n: S_n \to \mathbf{M}^3$ is primitive. If not, we can find a Riemann surface Q and $\pi: S_n \to Q$ and $g: Q \to \mathbf{M}^3$ such that $g \circ \pi = f_n$ and $d > 1$ where d is the degree of the cover π. We recall that f_n is bent along C_n^1 and C_n^2, and nearly isometric on the complement. So by Theorem 4.1, $\{C_n^1, C_n^2\}$ are the lifts by π of some set $\mathcal{C}_Q$ of simple closed curves on Q. So $|\mathcal{C}_Q| = 1$ or $|\mathcal{C}_Q| = 2$.

If $|\mathcal{C}_Q| = 2$, then each component of $S_n \setminus \bigcup C_n^i$ maps by degree d to a component of $Q \setminus \mathcal{C}_Q$. We can then write $Q \setminus \mathcal{C}_Q = \bar{Q}(1) \cup \bar{Q}(2)$ such that $\pi: \bar{S}_n(i) \to \bar{Q}(i)$ is a degree d cover, and then by Lemma 4.1 we can rejoin the boundary curves of $\bar{Q}(1)$ to form $Q'(1)$ such that $S_n(1)$ is a cover of $Q'(1)$. But then we get a subgroup $G_{Q'}$ of $G_n(1)$ ($G_{Q'} = \pi_1(Q'(1))$), and $G_n(1) < G_{Q'} \cap G(1) < G(1)$, where both inclusions are proper. The first inclusion is proper because $A^{k/d} \in G_{Q'} \cap G(1) \setminus G_n(1)$, and the

second is proper because $k < n$. This contradicts the assumption on the maximality of $G_n(1)$.

If $|\mathcal{C}_Q| = 1$, we let $\mathcal{C}_Q = \{C_Q\}$. First suppose that C_Q is nonseparating. Then writing $Q \setminus C_Q = \bar{Q}$ we find that $\bar{S}_n(1)$ and $\bar{S}_n(2)$ are both degree $d/2$ covers of $\bar{Q}$. But then we can reassemble $\bar{Q}$ to make Q' (by Lemma 4.1) such that $S_n(1)$ is a degree $d/2$ cover of Q', when $d/2 \le k$. Then we arrive at a contradiction by the same reasoning as before.

Finally, suppose that C_Q is separating. Then we can write $Q \setminus C_Q = \bar{Q}(1) \cup \bar{Q}(2)$ so that the restriction of π to $\bar{S}_n(i)$ is a cover of $\bar{Q}(i)$. Then the conjugacy classes for C_n^1 and C_n^2, oriented as curves covered by the axis of A, are both in $[A^k]$, but C_n^1 and C_n^2 both cover C_Q with opposite orientations, so the conjugacy class for C_Q must be both $[A^l]$ and $[A^{-l}]$, where $l = 2k/d$. But then $B^{-1}A^l B = A^{-l}$ for some $B \in \Gamma$, which means that B preserves the axis of A and reverses its orientation; such B would have a fixed point in $\mathbb{H}^3$, which is a contradiction. $\qquad\square$

4.2 The lower bound

We now proceed to prove the lower bound

$$(7) \qquad\qquad (c_1 g)^{2g} \le s_1(\mathbf{M}^3, g),$$

for g large enough, from Theorem 1.1.

By the above theorem the representation $\rho_n\colon \pi_1(S_n) \to \Gamma$, is maximal. It remains to count the number of such representations. Let g_n denote the genus of the surface S_n. If g_0 denotes the genus of the surfaces $S(1)$ and $S(2)$, we have

$$g_n = n(2g_0 - 1).$$

Given a closed surface S_0, Let $m_n(S_0)$ denote the number of maximal degree n covers of S_0. Let C_0 denote a simple closed and nonseparating curve in S_0. For $1 \le k \le n$, by $m_n(S_0, C_0, k)$ we denote the number of maximal n degree covers of S_0 such that the curve C_0 has at least one lift of degree k. Clearly the number $m_n(S_0, C_0, k)$ does not depend on the choice of the simple closed and nonseparating curve C_0, so we sometimes write $m_n(S_0, k) = m_n(S_0, C_0, k)$.

Theorem 4.3 *Let g_0 denote the genus of S_0. Then for n large we have*

$$m_n(S_0) = (n!)^{g_0 - 2}(1 + o(1)),$$

where $o(1) \to 0$ when $n \to \infty$. Moreover, for some $1 \le k \le (n-1)$, $k = k(n, g_0)$, we have

$$m_n(S_0, k) > ((n-1)!)^{g_0 - 2}(1 + o(1)).$$

Proof The first equality directly follows from Corollary 3 and the formula in Section 4.4 in [6], which shows that a random finite cover of a closed surface is maximal. It remains to prove the second inequality.

Since

$$\sum_{k=1}^{n} m_n(S_0, k) \geq m_n(S_0),$$

it follows that for some $1 \leq k \leq n$, the second inequality in the statement of the theorem holds. The following lemma implies that this inequality holds for some $1 \leq k \leq (n-1)$.

Lemma 4.2 *The inequality $m_n(S_0, 1) \geq m_n(S_0, n)$, holds for every n.*

Proof Let C_0 and D_0 be two simple closed and nonseparating curves on S_0, that intersect exactly once. Let S_n be a degree n cover of S_0, such that the curve C_0 has a degree n lift which we denote by C_n. Then C_n is the only lift of C_0. We show that in this case, every lift of the curve D_0 is a degree one lift. Let $\widetilde{S}_0 = S_0 \setminus C_0$ and $\widetilde{S}_n = S_n \setminus C_n$, denote the two surfaces each having exactly two boundary components. Then $\widetilde{S}_n$ covers $\widetilde{S}_0$, because C_n is the only lift of C_0 to S_n. After removing the curve C_0 from S_0, the closed curve D_0 becomes an interval $I_0 \subset \widetilde{S}_0$, whose endpoints lie on different boundary components of $\widetilde{S}_0$. Therefore, every lift of I_0 to $\widetilde{S}_n$ is a degree one lift. This proves the statement.

Restricting to when S_n is a maximal cover yields the inequality $m_n(S_0, C_0, n) \leq m_n(S_0, D_0, 1)$. Since $m_n(S_0, C_0, k) = m_n(S_0, D_0, k) = m_n(S_0, k)$, it follows that $m_n(S_0, 1) \geq m_n(S_0, n)$, and we have proved the lemma. □

This proves the theorem. □

Now fix a large n and choose $1 \leq k \leq (n-1)$ so that the second inequality in Theorem 4.3 holds. We then amalgamate any two maximal covers $S_n(1)$ and $S_n(2)$ along the curves $C_n(1)$ and $C_n(2)$ that are both k degree lifts of the curves $C(1)$ and $C(2)$ respectively (there may be more than one such k degree lift, but we choose arbitrarily). Then the corresponding group $\rho_n(\pi_1(S_n)) < \Gamma$ is maximal surface subgroup of Γ. Combining the above formula for g_n with the Theorem 4.3, we derive the estimate (7) for some $c_1 > 0$.

References

[1] **E A Bender**, **E R Canfield**, *The asymptotic number of rooted maps on a surface*, J. Combin. Theory Ser. A 43 (1986) 244–257   MR867650

[2]   **A Douady**, **C J Earle**, *Conformally natural extension of homeomorphisms of the circle*, Acta Math. 157 (1986) 23–48   MR857678

[3]   **J Kahn**, **V Markovic**, *Immersing almost geodesic surfaces in a closed hyperbolic three manifold*   arXiv:0910.5501

[4]   **C Kourouniotis**, *Complex length coordinates for quasi-Fuchsian groups*, Mathematika 41 (1994) 173–188   MR1288062

[5]   **J D Masters**, *Counting immersed surfaces in hyperbolic* 3–*manifolds*, Algebr. Geom. Topol. 5 (2005) 835–864   MR2153105

[6]   **T W Müller**, **J-C Puchta**, *Character theory of symmetric groups and subgroup growth of surface groups*, J. London Math. Soc. 66 (2002) 623–640   MR1934296

[7]   **C Series**, *An extension of Wolpert's derivative formula*, Pacific J. Math. 197 (2001) 223–239   MR1810217

[8]   **S P Tan**, *Complex Fenchel–Nielsen coordinates for quasi-Fuchsian structures*, Internat. J. Math. 5 (1994) 239–251   MR1266284

Department of Mathematics, Brown University
Providence RI 02912, USA

Department of Mathematics, California Institute of Technology
Pasadena CA 91105, USA

`kahn@math.brown.edu`, `markovic@caltech.edu`

`http://www.math.brown.edu/~kahn/`, `http://www.its.caltech.edu/~markovic`