



Nir Lazarovich · Arie Levit · Yair Minsky

Surface groups are flexibly stable

Received August 10, 2021; revised June 12, 2022

Abstract. We show that surface groups are flexibly stable in permutations. This is the first non-trivial example of a non-amenable flexibly stable group. Our method is purely geometric and relies on an analysis of branched covers of hyperbolic surfaces. Along the way we establish a quantitative variant of the LERF property for surface groups which may be of independent interest.

Keywords. Stability in permutations, flexible stability, surface groups, LERF, $CAT(-1)$, combinatorial group theory

1. Introduction

Stability in group theory studies when almost-actions of groups are close to honest actions. This recent notion has been investigated for various types of actions and groups. Of particular interest is the stability of actions by permutations on finite sets, as motivated by the outstanding open problem of whether every group is sofic. For more on the relation to soficity and existing results we refer to the end of this introduction.

Flexible stability

Let G be a finitely presented group with set of generators Σ . Let F_Σ be a free group with basis Σ and $\pi_\Sigma : F_\Sigma \rightarrow G$ be the natural quotient map. Consider the space $\text{Hom}(F_\Sigma, S_N)$ of homomorphisms from F_Σ into the finite symmetric group S_N for some $N \in \mathbb{N}$.

Roughly speaking, the group G is flexibly stable if any $\rho \in \text{Hom}(F_\Sigma, S_N)$ that almost factors through π_Σ is close to some $\rho' \in \text{Hom}(F_\Sigma, S_M)$ that actually factors through π_Σ for some $M \in \mathbb{N}$ potentially slightly larger than N .

Nir Lazarovich: Department of Mathematics, Technion–Israel Institute of Technology, 2000 Haifa, Israel; lazarovich@technion.ac.il

Arie Levit: School of Mathematical Sciences, Department of Pure Mathematics, Tel Aviv University, 6997801 Tel Aviv, Israel; arielevit@tauex.tau.ac.il

Yair Minsky: Department of Mathematics, Yale University, New Haven, CT 06511, USA; yair.minsky@yale.edu

Mathematics Subject Classification (2020): Primary 20B07; Secondary 20F67, 20F69, 57K20, 15A24

To make the above notion precise equip the symmetric group S_N with the bi-invariant *normalized Hamming metric* d_N given by

$$d_N(\sigma_1, \sigma_2) = \frac{1}{N} |\{i \in \{1, \dots, N\} : \sigma_1(i) \neq \sigma_2(i)\}|$$

for any pair of elements $\sigma_1, \sigma_2 \in S_N$. The space $\text{Hom}(F_\Sigma, S_N)$ then admits the corresponding metric

$$d_N(\rho_1, \rho_2) = \sum_{\sigma \in \Sigma} d_N(\rho_1(\sigma), \rho_2(\sigma))$$

for any pair of homomorphisms $\rho_1, \rho_2 \in \text{Hom}(F_\Sigma, S_N)$. Consider the map

$$\mathcal{E}_N^M : \text{Hom}(F_\Sigma, S_N) \rightarrow \text{Hom}(F_\Sigma, S_M)$$

defined for $N, M \in \mathbb{N}$ with $N \leq M$ by extending a given permutation representation of F_Σ to act trivially on the extra $M - N$ points.

Definition. The group G is *flexibly stable in permutations* if it admits a finite presentation $G = \langle \Sigma | R \rangle$ with the following property: for any $\varepsilon > 0$ there is some $\delta = \delta(\varepsilon) > 0$ such that for any $N \in \mathbb{N}$ and $\rho \in \text{Hom}(F_\Sigma, S_N)$ satisfying

$$d_N(\rho(r), \text{id}_N) < \delta \quad \forall r \in R$$

there exists $M \in \mathbb{N}$ with $N \leq M \leq (1 + \varepsilon)N$ and $\rho' \in \text{Hom}(F_\Sigma, S_M)$ that factors through π_Σ and satisfies

$$d_M(\mathcal{E}_N^M(\rho), \rho') < \varepsilon.$$

We remark that if the above property holds with respect to one finite presentation of the group G then indeed it holds for all such presentations [1].

The following is our main result.

Theorem 1.1. *Let S be a closed orientable surface of genus $g \geq 2$. Then the fundamental group $\pi_1(S)$ is flexibly stable in permutations.*

In fact, our proof of Theorem 1.1 gives an explicit relation between ε and δ : up to constants, $\varepsilon = \delta \ln(1/\delta)$. This means that $\delta = o(\varepsilon)$ and $\varepsilon^2 = o(\delta)$ in the limit as both ε and δ go to 0.

On our method

Our approach to proving flexible stability is purely geometric.

We use covering space theory to reformulate the problem in terms of certain branched covers of S . In this language the goal is to convert a given branched cover of S into an unramified one by performing an amount of changes controlled by the total branching degree. For more information we refer the reader to Theorem 2.1 and the discussion leading to it.

We dissect a given branched cover X of S along a carefully constructed embedded graph Γ with vertices at singular points of X and with geodesic edges. In fact, Γ is taken to be a subgraph of the Delaunay graph with respect to the singular points. It is chosen in such a way that every resulting connected component of $X \setminus \Gamma$ has locally convex boundary and that the total length of all boundary curves is controlled by the total branching degree of X (see Theorem 4.1).

The desired unramified cover of S is obtained by individually embedding every connected component of $X \setminus \Gamma$ into an unramified cover of S with controlled area. This relies on the following result, which seems to be of independent interest.

Theorem 1.2. *Let S be a closed hyperbolic surface. Let R be a surface with boundary which is isometrically embedded in some cover of S . If the boundary ∂R is locally convex then R can be isometrically embedded in a cover Q of S such that the diagram*

$$\begin{array}{ccc}
 R & \hookrightarrow & Q \\
 & \searrow & \swarrow \\
 & S &
 \end{array}
 \tag{A}$$

commutes and

$$\text{Area}(Q \setminus R) \leq b_S l(\partial R)$$

where $b_S > 0$ is a constant depending only on the surface S .

Surface groups are locally extended residually finite (LERF). This fact was established by Scott [14, 15]. It implies that any surface R as above isometrically embeds in some finite cover of S . Therefore Theorem 1.2 can be regarded as a certain quantitative variant of the LERF property for surface groups. As such it is closely related to and is inspired by Patel’s work [13]. It is however crucial for our purposes that the upper bound depends on $l(\partial R)$ rather than on $\text{Area}(R)$.

A finitely presented group is *strictly stable in permutations* if it satisfies the definition of flexible stability with M being exactly equal to N . The methods we use towards proving Theorem 1.2 will in general increase the total area. For this reason we are only able to establish flexible rather than strict stability.

This note includes an appendix on Voronoi cells and Delaunay graphs constructed on “hyperbolic planes with singularities”.

Soficity, stability and related works

There has been a significant recent interest in the notion of sofic groups. Roughly speaking, a group is sofic if it can be approximated by almost-actions on finite sets; see e.g. [8, 16] for formal definitions. Groups known to be sofic include amenable and residually finite ones. An outstanding open problem asks: is every group sofic?

A part of the interest in stability stems from the observation made in [9] that a non-residually-finite group which is stable in permutations cannot be sofic. This observation extends to flexible stability.

Groups that are currently known to be strictly stable in permutations include finite groups (Glebsky and Rivera [9]), finitely generated abelian groups (Arzhantseva and Păunescu [1]) and polycyclic and Baumslag–Solitar $BS(1, n)$ groups (Becker, Lubotzky and Thom [4]).

Becker and Lubotzky [3] proved that a group G with Kazhdan’s property (T) is not strictly stable by removing a single point from an action of G on a finite set. This strategy has led Becker and Lubotzky to introduce the flexible notion of stability in permutations and ask whether some Kazhdan groups might still be stable in the flexible sense. The question of strict and flexible stability for surface groups is discussed in [3] as well.

The importance of the question concerning flexible stability for some Kazhdan groups is highlighted by the recent work of Bowen and Burton [6]. They prove the existence of a non-sofic group, conditioned on the assumption that the group $PSL_d(\mathbb{Z})$ is flexibly stable for some $d \geq 5$.

The fundamental group of a closed orientable surface with constant non-negative curvature is either finite or abelian and as such stable respectively by [9] and [1]. Our Theorem 1.1 completes the picture for all closed orientable surfaces, at least as far as flexible stability is concerned.

Free non-abelian groups are clearly stable in a void sense. On the other hand, there exist small cancellation hyperbolic groups that are not even flexibly stable [3]. Interestingly, while hyperbolic surface groups are possibly the simplest example of non-free hyperbolic groups, the question of stability for these groups is not trivial. To the best of the authors’ knowledge it is not known whether surface groups are stable in permutations in the strict sense. The problem of adapting our present approach to deal with this question seems challenging.

Finally, we point out that it was recently shown by Becker and Mosheiff [5] that for the free abelian group \mathbb{Z}^d the parameter $\delta = \delta(\varepsilon)$ goes to zero at least as fast as a polynomial of degree d in ε . So in some quantitative sense hyperbolic surface groups are “more stable” than free abelian ones.

2. $*$ -covers and geometric stability

We develop a geometric framework for the study of flexible stability of surface groups in terms of branched covers of surfaces.

**-covers of surfaces*

Let S be a closed surface. Endow S with a metric of constant sectional curvature. Recall that a finite *branched cover* of S is a continuous surjection $p : X \rightarrow S$ where every point $b \in S$ admits a neighborhood $b \in V_b \subset S$ with $p^{-1}(V_b) = U_1 \sqcup \cdots \sqcup U_{n_b}$ and such that $p : U_i \rightarrow V_b$ is topologically conjugate to the complex map $z \mapsto z^d$ of some degree $d \in \mathbb{N}$ depending on i . The degree is equal to 1 unless b belongs to a finite subset of S called the *branch set*.

Definition. A $*$ -cover of S is a compact surface X admitting a branched cover $p : X \rightarrow S$ with branch set consisting of a single branch point $* \in S$.

Let $p : X \rightarrow S$ be a $*$ -cover. In general X is not required to be connected. We pull back the metric from S to every connected component of X so that p is a local isometry away from $p^{-1}(*)$.

The total angle α_x locally at any point $x \in X$ is an integer multiple of 2π . In fact, α_x is equal to $2\pi d_x$ where $d_x \in \mathbb{N}$ is the *degree* (or *index*) of the point $x \in X$.

The *singular set* of X is $s(X) = \{x \in X : d_x > 1\}$. Clearly $s(X) \subset p^{-1}(*)$ and in particular $s(X)$ is discrete. The *branching degree* of X is $\beta(X) = \sum_{x \in s(X)} (d_x - 1)$. We say that the $*$ -cover X is *unramified* if $s(X) = \emptyset$, or equivalently $\beta(X) = 0$. This happens if and only if p is a covering map. We emphasize that covers are not required to be connected.

The *degree* $|X|$ of the $*$ -cover $p : X \rightarrow S$ is equal to $|p^{-1}(x)|$ for any $x \in S \setminus \{*\}$.

Geometric stability

The geometric notion of stability is defined in terms of certain graphs embedded into $*$ -covers. Recall that a *graph* is a simplicial 1-complex. We let $\Gamma^{(0)}$ denote the vertex set of the graph Γ .

Definition. A $*$ -graph on X is an embedded graph Γ such that $\Gamma^{(0)} \subset p^{-1}(*)$ and the edges of Γ are geodesic arcs.

The fact that Γ is embedded means that two edges of Γ may intersect only at the vertex set $\Gamma^{(0)}$. If Γ is a $*$ -graph on X then any closed curve γ contained in Γ is a piecewise geodesic closed curve in X . Given a $*$ -graph Γ let $l(\Gamma)$ denote the total length of all of its edges.

Definition. The $*$ -cover $p : X \rightarrow S$ is ε -reparable if X admits a $*$ -graph Γ with $l(\Gamma) \leq \varepsilon \text{Area}(X)$ such that the complement $X \setminus \Gamma$ isometrically embeds into a cover C of S such that the diagram

$$\begin{array}{ccc}
 X \setminus \Gamma & \hookrightarrow & C \\
 & \searrow p & \swarrow \\
 & & S
 \end{array} \tag{B}$$

commutes and

$$\text{Area}(X) \leq \text{Area}(C) \leq (1 + \varepsilon) \text{Area}(X).$$

We emphasize that the complement $X \setminus \Gamma$ as well as the cover C are in general allowed to be disconnected.

Definition. The surface S is *flexibly geometrically stable* if for every $\varepsilon > 0$ there is some $\delta = \delta(\varepsilon) > 0$ such that any $*$ -cover X with $\beta(X) < \delta \text{Area}(X)$ is ε -reparable.

Our main result can be reformulated in the language of $*$ -covers.

Theorem 2.1. *Every closed orientable hyperbolic surface is flexibly geometrically stable.*

In the remaining part of this section we show that geometric stability implies algebraic stability in permutations; see Proposition 2.4 below.

Permutation representations and covering theory

Consider the punctured surface $S \setminus \{*\}$. Fix a base point $x_0 \in S \setminus \{*\}$ and let $F = \pi_1(S \setminus \{*\}, x_0)$ so that F is a free group. Note that $X \setminus p^{-1}(*) \rightarrow S \setminus \{*\}$ is a covering map of degree $|X|$. Enumerate the fiber $p^{-1}(x_0)$ by fixing an arbitrary identification with the finite set $\{1, \dots, |X|\}$. Covering space theory [10, p. 68] shows how to associate a natural permutation representation $\rho_X : F \rightarrow S_{|X|}$ to the covering $X \setminus p^{-1}(*) \rightarrow S \setminus \{*\}$.

In what follows it is convenient to fix a specific presentation for the fundamental group of S . Since the notion of flexible stability is known to be independent of the chosen presentation we may do so without any loss of generality.

Let \mathcal{P} be a compact fundamental domain for the action of the fundamental group $\pi_1(S, x_0)$ on the universal cover \tilde{S} . Assume that \mathcal{P} is a polygon with finitely many geodesic sides and that the vertices of \mathcal{P} are all lifts of the point $* \in S$. Moreover, assume that the lift of the point x_0 lies in the interior of \mathcal{P} .

Let Σ be the finite generating set of $\pi_1(S, x_0)$ consisting of a single element σ from any pair $\{\sigma, \sigma^{-1}\} \subset \pi_1(S, x_0)$ such that $\sigma\mathcal{P} \cap \mathcal{P}$ is a geodesic side of \mathcal{P} . We may identify F with the free group F_Σ on the generators Σ . There is a natural quotient homomorphism $\pi_\Sigma : F_\Sigma \rightarrow \pi_1(S, x_0)$. The surface group $\pi_1(S, x_0)$ admits a presentation with generating set Σ and a single relation $r \in F_\Sigma$.

In particular, we may let the fundamental domain \mathcal{P} be a $4g$ -sided polygon so that $\Sigma = \{a_1, b_1, \dots, a_g, b_g\}$ and $r = \prod_{i=1}^g [a_i, b_i]$.

Proposition 2.2. *Given any $\rho \in \text{Hom}(F_\Sigma, S_N)$ with $N \in \mathbb{N}$ there exists a $*$ -cover $X_\rho \rightarrow S$ with $|X_\rho| = N$ and $\rho_{X_\rho} = \rho$. Moreover, $\beta(X_\rho)/N \leq d_N(\rho(r), \text{id}_N)$ where r is the defining relation of $\pi_1(S, x_0)$ as above.*

Note that $d_N(\rho(r), \text{id}_N)$ is equal to $1 - |\text{Fix}(\rho(r))|/N$ where $\text{Fix}(\rho(r))$ is the set of fixed points of the permutation $\rho(r) \in S_N$.

Proof of Proposition 2.2. Let X'_ρ be the punctured surface covering $S \setminus \{*\}$ that corresponds to the permutation representation $\rho : F_\Sigma \rightarrow S_N$. Let $p : X_\rho \rightarrow S$ be the $*$ -cover obtained by completing X'_ρ at the punctures. It is clear that $|X_\rho| = N$ and $\rho_{X_\rho} = \rho$.

Consider the cycle decomposition $o_1 \cdots o_m$ of the permutation $\rho(r)$ in its action on the set $p^{-1}(x_0) \cong \{1, \dots, N\}$. Let l_i denote the length of the cycle o_i . We claim that without loss of generality $p^{-1}(*) = \{y_1, \dots, y_m\}$ and $d_{y_i} = l_i$. This would imply the required upper bound on $\beta(X_\rho)$ because

$$\frac{\beta(X_\rho)}{N} = \frac{1}{N} \sum_{y \in p^{-1}(*)} (d_y - 1) = \frac{1}{N} \sum_{i=1}^m (l_i - 1) \leq d_N(\rho(r), \text{id}_N).$$

Let γ be a simple closed curve in $S \setminus \{*\}$ based at x_0 and representing the element r of F_Σ . The preimage of the curve γ in X'_ρ is a disjoint union of simple closed curves. Every such curve corresponds to an orbit o_i of $\rho(r)$ in its action on $p^{-1}(x_0)$ and bounds a disc in X_ρ that contains a single point y_i from $p^{-1}(*)$ in its interior. Moreover, the degree d_{y_i} is equal to the size l_i of the orbit o_i . The claim follows. ■

Geometric stability implies algebraic

We show that geometric stability in the sense of Theorem 2.1 implies our main result Theorem 1.1. We continue using the presentation $\pi_1(S, x_0) \cong \langle \Sigma | r \rangle$ constructed above and in particular the polygon \mathcal{P} .

Proposition 2.3. *Let Γ be a $*$ -graph on X with $N = |X|$. Assume that the complement $X \setminus \Gamma$ isometrically embeds into some cover C of S of degree $M = |C|$ such that diagram (B) commutes. Regard the two permutation representations ρ_X and ρ_C as mapping respectively into the symmetric groups S_N and S_M . Then*

$$d_M(\mathcal{E}_N^M(\rho_X), \rho_C) \leq a_S \left(\frac{l(\Gamma) + (\text{Area}(C) - \text{Area}(X))}{\text{Area}(C)} \right)$$

where $a_S > 0$ is a constant depending only on the surface S .

Recall that the map $\mathcal{E}_N^M : \text{Hom}(F_\Sigma, S_N) \rightarrow \text{Hom}(F_\Sigma, S_M)$ is defined by extending a permutation representation to act trivially on the extra $M - N$ points.

Proof of Proposition 2.3. The map $p : X \rightarrow S$ induces a tessellation $\mathcal{T} = \mathcal{T}(X, \mathcal{P})$ of the $*$ -cover X by $|X|$ isometric copies of the polygon \mathcal{P} . We claim that every edge α of the graph Γ satisfies

$$|\{\mathcal{D} \in \mathcal{T} : \mathcal{D} \cap \alpha \not\subseteq p^{-1}(*)\}| \leq a'_S l(\alpha)$$

where $a'_S > 0$ is a constant depending only on the surface S , the choice of the fundamental domain \mathcal{P} and the particular generating set Σ .

To prove the claim fix an edge α of the graph Γ . This edge is a geodesic segment in X such that $\alpha \cap p^{-1}(*) = \partial\alpha$. We want to bound the number of polygons \mathcal{D} of the tessellation \mathcal{T} that α meets in its interior.

If α has a subsegment that coincides with a subsegment of an edge of the tessellation \mathcal{T} , then α equals that edge (since both are geodesics which end in $p^{-1}(*)$). In this case α meets the two polygons incident to this edge, and its length is at least the length ℓ of the shortest edge of \mathcal{P} . Therefore, α satisfies the desired inequality if $a'_S \geq 2/\ell$.

Otherwise, α meets the edges of the tessellation \mathcal{T} transversely. Consider how α is cut by the boundary edges of the tessellation \mathcal{T} into geodesic segments $\alpha_1, \dots, \alpha_k$. The number of polygons that α meets is bounded by k .

Note that since α does not pass through the singular points of X , $p \circ \alpha$ is a geodesic in S (which is not necessarily embedded). Each $p \circ \alpha_i$ is an embedded geodesic in S

whose boundary is on the edges of \mathcal{Q} , where \mathcal{Q} is the polygon in S to which \mathcal{P} projects under the universal covering map.

Fix $\varepsilon > 0$ smaller than half the injectivity radius of S . There exists $\delta > 0$ (depending on \mathcal{P} and ε) such that each $p \circ \alpha_i$ either has length at least $\delta > 0$ or is contained in the ε -neighborhood of $*$. This shows in particular that $l(\alpha) \geq \delta$. Moreover, since the edges of \mathcal{Q} are geodesic, there exists m such that any geodesic segment in S of length 2ε meets the edges of \mathcal{Q} at most m times. It follows that any $m + 1$ consecutive segments in the sequence $p \circ \alpha_1, \dots, p \circ \alpha_k$ must leave the ε -neighborhood of $*$, and therefore must have length at least δ altogether. This shows that $\delta \frac{k}{2m+2} \leq \max\{\delta, \delta \lfloor \frac{k}{m+1} \rfloor\} \leq l(\alpha)$. So the edge α satisfies the desired inequality if $a'_S \geq \frac{2m+2}{\delta}$. Choosing $a'_S = \max\{\frac{2}{\ell}, \frac{2m+2}{\delta}\}$ completes the proof of the claim.

Taking into account the claim and summing over all edges of the graph Γ gives the estimate

$$|\{\mathcal{D} \in \mathcal{T} : \mathcal{D} \cap \Gamma \not\subseteq p^{-1}(*)\}| \leq a'_S l(\Gamma).$$

Consider a point $x \in p^{-1}(x_0)$ and let $\mathcal{D} \in \mathcal{T}$ be the polygon containing x in its interior. The fiber $p^{-1}(x_0)$ is clearly a subset of the fiber $q^{-1}(x_0)$ so that $x \in q^{-1}(x_0)$. Given a particular generator $\sigma \in \Sigma$ let $\mathcal{D}_\sigma \in \mathcal{T}$ be the polygon sharing with \mathcal{D} the geodesic side corresponding to σ . Then

$$\mathcal{E}_N^M(\rho_X)(\sigma)(x) = \rho_C(\sigma)(x)$$

provided that the $*$ -graph Γ does not intersect $(\mathcal{D} \cup \mathcal{D}_\sigma) \setminus p^{-1}(*)$. The Hamming metric between $\mathcal{E}_N^M(\rho_X)$ and ρ_C is therefore bounded above by

$$d_M(\mathcal{E}_N^M(\rho_X), \rho_C) \leq \frac{2|\Sigma|}{M}(a'_S l(\Gamma) + |q^{-1}(x_0) \setminus p^{-1}(x_0)|)$$

and the conclusion follows for an appropriate choice of the constant $a_S > 0$. ■

Proposition 2.4. *Let S be a closed hyperbolic surface. If S is flexibly geometrically stable then its fundamental group $\pi_1(S, x_0)$ is flexibly stable in permutations.*

Proof. The algebraic property of flexible stability in permutations does not depend on the choice of a particular presentation. This is proved in [1] for the strict notion of stability using only the fact that any two finite presentations for G are related by a finite sequence of Tietze transformations [11, II.2.1]. Exactly the same argument goes through in the flexible case as well. It will be convenient to verify the definition with respect to the presentation $\pi_1(S, x_0) \cong \langle \Sigma | r \rangle$.

Let $\varepsilon > 0$ be given. Denote $\varepsilon' = \min\{\varepsilon, \frac{\varepsilon}{2a_S}\}$. According to our assumption there is some $\delta = \delta(\varepsilon')$ such that any $*$ -cover X with $\beta(X) < \delta \text{Area}(X)$ is ε' -reparable. We claim that the condition of flexible stability in permutations for the group $\pi_1(S, x_0)$ is satisfied with respect to the given $\varepsilon > 0$ and this particular δ .

Consider some $N \in \mathbb{N}$ and $\rho \in \text{Hom}(F_\Sigma, S_N)$ with $d_N(\rho(r), \text{id}_N) < \delta$. Making use of Proposition 2.2 we find a $*$ -cover $X_\rho \rightarrow S$ with $\beta(X_\rho) < \delta N$ and $\rho_{X_\rho} = \rho$. Therefore X_ρ is ε' -reparable, which means it admits an embedded $*$ -graph Γ with $l(\Gamma) \leq \varepsilon' \text{Area}(X_\rho)$ and

such that the complement $X_\rho \setminus \Gamma$ isometrically embeds into some cover C of S making diagram (B) commute and satisfying

$$\text{Area}(X_\rho) \leq \text{Area}(C) \leq (1 + \varepsilon') \text{Area}(X_\rho).$$

Consider the permutation representation $\rho' = \rho_C$ so that $\rho' \in \text{Hom}(F_\Sigma, S_M)$ where $M = |C|$. Since C is unramified, the homomorphism ρ' factors through $\pi_\Sigma : F_\Sigma \rightarrow \pi_1(S, x_0)$. Roughly speaking, the two permutation representations ρ and ρ' agree on points lying on the subsurface $X \setminus \Gamma$ of C away from the $*$ -graph Γ . To be precise, it follows from Proposition 2.3 that

$$d_M(\mathcal{E}_N^M(\rho), \rho') \leq a_S \left(\frac{l(\Gamma) + (\text{Area}(C) - \text{Area}(X_\rho))}{\text{Area}(C)} \right) \leq 2a_S \varepsilon' \leq \varepsilon.$$

Finally, the degree M satisfies

$$N \leq M \leq (1 + \varepsilon')N \leq (1 + \varepsilon)N$$

as required. ■

The remainder of this work is dedicated to establishing Theorem 2.1.

3. Hyperbolic planes with singularities

Let S be a compact hyperbolic surface and consider a fixed $*$ -cover $p : X \rightarrow S$. We discuss the geometry of the universal cover of X . We then discuss the Voronoi tessellation and its dual, the Delaunay graph. This graph will be used in Section 4 to construct the $*$ -graph as required in the definition of flexible geometric stability.

The geometry of the universal cover of X

Let $q : \tilde{X} \rightarrow X$ denote the universal cover of X equipped with the pullback length metric $d_{\tilde{X}}$. Topologically speaking \tilde{X} is homeomorphic to the plane.

The singular set of \tilde{X} is $s(\tilde{X}) = q^{-1}(s(X))$. The space \tilde{X} is locally isometric to the hyperbolic plane \mathbb{H} away from its singular set. The group $\pi_1(X)$ acts freely on \tilde{X} admitting the surface X as a quotient. The singular set $s(\tilde{X})$ is discrete in \tilde{X} . In fact, $s(\tilde{X})$ is co-bounded provided $s(X)$ is not empty.

Recall that the Cartan–Hadamard theorem admits a generalization due to Gromov to complete geodesic metric spaces. See [7, Theorem II.4.1] for reference. This result implies that \tilde{X} is a CAT(−1)-space. In particular, \tilde{X} is *uniquely geodesic* and every *local geodesic in \tilde{X} is a geodesic*.

Let γ be a continuous path in \tilde{X} . The path γ is a *local geodesic* provided it is a local geodesic in the sense of hyperbolic geometry away from the singular set $s(\tilde{X})$ and its angle θ_x at every singular point $x \in s(\tilde{X})$ along γ satisfies $\pi \leq \theta_x \leq \alpha_x - \pi$ where α_x is the local angle at x . This local characterization implies that \tilde{X} is *geodesically complete*

in the sense that any geodesic segment can be extended to a bi-infinite geodesic line. The geodesic extension need not be unique since a geodesic segment terminating at a singular point admits many extensions.

Convex subsets of the universal cover \tilde{X}

Recall the following useful consequence of the Cartan–Hadamard theorem [7, Theorems II.4.13 and II.4.14].

Lemma 3.1. *Let N_1 and N_2 be connected complete non-positively curved metric spaces and $f : N_1 \rightarrow N_2$ be a local isometry. Then $f_* : \pi_1(N_1) \rightarrow \pi_1(N_2)$ is injective and any lift $F : \tilde{N}_1 \rightarrow \tilde{N}_2$ of the map f is an isometric embedding.*

Here is an example of a straightforward application of Lemma 3.1.

Corollary 3.2. *Let $C \subset \tilde{X}$ be a convex subset such that $\mathring{C} \cap s(\tilde{X}) = \emptyset$. Then C is isometric to a contractible subset of the hyperbolic plane \mathbb{H} .*

Proof. Since C is a convex subset of \tilde{X} , it is a CAT(−1)-space in its own right. The corollary follows by applying Lemma 3.1 with $N_1 = C$, $N_2 = S$ and $f = p \circ q$. The lift F gives the required isometric embedding into \mathbb{H} . ■

Let γ be a bi-infinite geodesic path in \tilde{X} . A *half-space* in \tilde{X} is the closure of a connected component of $\tilde{X} \setminus \gamma$. Note that a half-space is convex.

Voronoi cells and the Delaunay graph

Assume that the singular set $s(X)$ is non-empty. We define the Voronoi cells and the Delaunay graph with respect to the set of points $s(\tilde{X})$. These are natural generalizations of the parallel notions in the classical Euclidean and hyperbolic cases.

Definition. The *Voronoi cell* A_v at the singular point $v \in s(\tilde{X})$ is

$$A_v = \{x \in \tilde{X} : d_{\tilde{X}}(x, v) \leq d_{\tilde{X}}(x, u) \ \forall u \in s(\tilde{X})\}.$$

The family of Voronoi cells is equivariant in the sense that $gA_v = A_{gv}$ for every $g \in \pi_1(X)$ and $v \in s(\tilde{X})$.

We remark that it is not a priori clear that the Voronoi cells form a tessellation. This turns out to be true and will be established as a consequence of Proposition 3.4 below. The difficulty has to do with the fact that for two singular points $v, u \in s(\tilde{X})$ the intersection of the two sets $\{x \in \tilde{X} : d_{\tilde{X}}(x, v) \leq d_{\tilde{X}}(x, u)\}$ and $\{x \in \tilde{X} : d_{\tilde{X}}(x, v) \geq d_{\tilde{X}}(x, u)\}$ might have a non-empty interior in general.

Definition. The vertex set of the *Delaunay graph* D is the singular set $s(\tilde{X})$. Two vertices v_1 and v_2 of the Delaunay graph D span an edge whenever there is a closed metric ball $B \subset \tilde{X}$ with $\mathring{B} \cap s(\tilde{X}) = \emptyset$ and $B \cap s(\tilde{X}) = \{v_1, v_2\}$.

We summarize a few basic properties of Voronoi cells and their relationship with the Delaunay graph. These are quite elementary in the classical Euclidean or hyperbolic cases. Extra caution is required in the presence of singular points – some of the following statements are false unless the Voronoi cells are taken with respect to a set of points containing all singularities, which is always the case here.

The proofs of Propositions 3.3, 3.4, 3.6 and 3.7 are postponed to Appendix A dedicated to this topic. The proofs are similar to the classical case by repeatedly relying on Corollary 3.2 to embed the relevant local picture into the hyperbolic plane.

Proposition 3.3. *Let $v \in s(\tilde{X})$ be a vertex of the Delaunay graph D . Consider the Voronoi cell A_v . Then*

- (1) A_v is homeomorphic to a closed disc and $A_v \cap s(\tilde{X}) = \{v\}$,
- (2) A_v is convex and the boundary ∂A_v is piecewise geodesic,
- (3) A_v is equal to the intersection of the sets $\{x \in \tilde{X} : d_{\tilde{X}}(x, v) \leq d_{\tilde{X}}(x, u)\}$ where $u \in s(\tilde{X})$ ranges over the vertices adjacent to v in the graph D .

The following proposition describes the intersection of two Voronoi cells.

Proposition 3.4. *Let $v, u \in s(\tilde{X})$ be a pair of distinct vertices of the Delaunay graph. Consider the two Voronoi cells A_v and A_u .*

- If $A_v \cap A_u \neq \emptyset$ then $A_v \cap A_u$ is either a single point or a common geodesic side.
- A_v and A_u have a common geodesic side if and only if v and u span an edge in the Delaunay graph.

As a consequence of Proposition 3.4 we deduce that the family of Voronoi cells A_v for $v \in s(\tilde{X})$ forms a tessellation of \tilde{X} called the *Voronoi tessellation*. Moreover, the Delaunay graph D is dual to this tessellation. These facts are well-known in the classical Euclidean and hyperbolic situations.

Corollary 3.5. *Let $v \in s(\tilde{X})$ be a vertex of the Delaunay graph. Then the interior of the Voronoi cell A_v embeds into X via the restriction to A_v of the covering map $q : \tilde{X} \rightarrow X$.*

Proof. It suffices to observe that $g\mathring{A}_v \cap \mathring{A}_v = \mathring{A}_{gv} \cap \mathring{A}_v = \emptyset$ for every $g \in \pi_1(X)$ with $g \neq \text{id}$. ■

The geometric realization of the Delaunay graph

From now on we regard the Delaunay graph D as being *geometrically realized* in \tilde{X} . More precisely, we identify the vertices of D with the singular set $s(\tilde{X})$ and realize every edge of D by the corresponding geodesic arc in \tilde{X} .

Proposition 3.6. *The Delaunay graph \tilde{D} is embedded in \tilde{X} . The projection $q(D)$ is embedded in the surface X .*

The above statement means that geodesic arcs realizing two distinct edges of D may intersect only at a vertex incident to both. In particular, D is planar. Likewise two distinct edges of $q(D)$ may intersect only at a common singular point of $s(X)$.

Proposition 3.7. *Any connected component of the complement of the Delaunay graph in \tilde{X} has locally convex boundary.*

An estimate for the area of Voronoi cells

Recall that a subset L of a metric space M is r -separated if $d_M(x_1, x_2) \geq r$ for any two distinct points $x_1, x_2 \in L$. Moreover, recall that we defined a half-space in \tilde{X} to be the closure of a connected component of the complement of some bi-infinite geodesic line in \tilde{X} .

Lemma 3.8. *Let A_v be the Voronoi cell associated to the singular point $v \in s(\tilde{X})$. Let \mathfrak{S} be a half-space in \tilde{X} with $v \in \partial\mathfrak{S}$. Assume that $s(\tilde{X})$ is r -separated and that $d_{\tilde{X}}(v, u) \geq R$ for every singular point $u \in \mathfrak{S} \cap s(\tilde{X})$ distinct from v . Then*

$$\text{Area}(\mathfrak{S} \cap A_v) \geq 2\varphi \sinh^2(R/4)$$

where $\varphi > 0$ is a constant depending only on the distance r .

The expression appearing on the right-hand side of the above estimate is the area of the hyperbolic sector with central angle φ and of radius $R/2$. Throughout the following proof and given a point $p \in \tilde{X}$ it is convenient to introduce the notation

$$\mathfrak{F}(p) = \{x \in \tilde{X} : d_{\tilde{X}}(x, v) \leq d_{\tilde{X}}(x, p)\}.$$

Proof of Lemma 3.8. Let us first determine the angle $\varphi > 0$ as follows. Let l_1 and l_2 be a pair of geodesic lines in the hyperbolic plane \mathbb{H} such that $d_{\mathbb{H}}(l_1, l_2) = r$. Let m be the midpoint of the geodesic arc perpendicular to both l_1 and l_2 . Then φ is the angle between the two geodesic rays emanating from the point m towards the ideal points $l_1(\infty)$ and $l_2(\infty)$.

Consider the bi-infinite geodesic $\rho : \mathbb{R} \rightarrow \tilde{X}$ parametrized by arc length and such that $\rho(\mathbb{R}) = \partial\mathfrak{S}$. Assume that $v = \rho(0)$ and denote $v_t = \rho(t)$ for all $t \in \mathbb{R}$.

Let \mathcal{W}_r^R be the subset of \tilde{X} given by

$$\mathcal{W}_r^R = \mathfrak{S} \cap \mathfrak{F}(v_{-r}) \cap \mathfrak{F}(v_r) \cap B_{\tilde{X}}(v, R).$$

As a consequence of Corollary 3.2, the two convex sets $B_{\tilde{X}}(v, r)$ and $\mathfrak{S} \cap B_{\tilde{X}}(v, R)$ are isometric to a hyperbolic ball of radius r and a hyperbolic sector of radius R and angle π , respectively. Let $Q_{\varphi, R/2}$ be the hyperbolic sector of angle φ and radius $R/2$ which is based at the vertex v and contained in \mathcal{W}_r^R . The lemma will be established by showing that \mathcal{W}_r^R and therefore $Q_{\varphi, R/2}$ is contained in $\mathfrak{S} \cap A_v$.

Let u_1, \dots, u_n with $n \in \mathbb{N}$ be the vertices of the Delaunay graph D adjacent to the vertex v . For every $i \in \{1, \dots, n\}$ let w_i be the point along the geodesic arc from v to u_i




Fig. 1. The Voronoi cell A_v based at the vertex v .

such that $d_{\tilde{X}}(v, w_i) = r$. The convexity of the metric [7, Proposition II.2.2] implies that $\mathfrak{F}(w_i) \subset \mathfrak{F}(u_i)$ and likewise $\mathfrak{F}(v_{\pm r}) \subset \mathfrak{F}(v_{\pm 2r})$.

Claim. *Let u_i be one of the Delaunay vertices adjacent to v . If $u_i \notin \mathfrak{S}$ then*

$$\mathcal{W}_r^R \subset \mathfrak{F}(w_i) \subset \mathfrak{F}(u_i).$$

Proof of Claim. Consider a point $x \in \mathcal{W}_r^R$. We need to show that $x \in \mathfrak{F}(w_i)$. Note that w_i and x lie on opposite sides of the bi-infinite geodesic ρ . Both w_i and x belong to the following locally convex and contractible and hence *convex* set:

$$B_{\tilde{X}}(v, r) \cup (\mathfrak{F}(v_{-2r}) \cap \mathfrak{F}(v_{2r}) \cap \mathfrak{S} \cap B_{\tilde{X}}(v, R)).$$

Therefore the intersection point z of ρ with the geodesic arc from w_i to x lies along the geodesic arc from v_{-r} to v_r . This relies on the fact that $r \leq R$ (which follows from the assumptions) and on the observation that $\mathfrak{F}(v_{-2r}) \cap \mathfrak{F}(v_{2r}) \cap \rho$ is equal to the geodesic segment from v_{-r} to v_r .

Assume without loss of generality that in fact the intersection point z lies along the geodesic arc from v to v_r . Consider the two geodesic triangles

$$T_1 = \Delta(v, v_r, w_i) \quad \text{and} \quad T_2 = \Delta(x, v_r, w_i).$$

The triangles T_1 and T_2 have no singular points in their interior. Therefore T_1 and T_2 are isometric to their respective comparison hyperbolic triangles according to Corollary 3.2. The triangle T_1 is isosceles with $d_{\tilde{X}}(v, w_i) = d_{\tilde{X}}(v, v_r) = r$. Observe that the respective angles of T_1 and T_2 at the vertices w_i and v_r satisfy

$$\angle_{T_1} w_i \geq \angle_{T_2} w_i \quad \text{and} \quad \angle_{T_1} v_r \leq \angle_{T_2} v_r.$$

The hyperbolic law of sines [2, Section 7.12] implies

$$d_{\tilde{X}}(x, v_r) \leq d_{\tilde{X}}(x, w_i).$$

Finally, since $x \in \mathfrak{F}(v_r)$ we have $d_{\tilde{X}}(x, v) \leq d_{\tilde{X}}(x, v_r)$ and so $x \in \mathfrak{F}(w_i)$. ■

Denote $I = \{i \in \{1, \dots, n\} : u_i \in \mathfrak{S}\}$ and $J = \{1, \dots, n\} \setminus I$. The above Claim implies that $\mathcal{W}_r^R \subset \bigcap_{j \in J} \mathfrak{F}(u_j)$. On the other hand, the assumptions of Lemma 3.8 imply that $B_{\tilde{X}}(v, R/2) \subset \bigcap_{i \in I} \mathfrak{F}(u_i)$. Recall that the Voronoi cell A_v in question is equal to $\bigcap_{i=1}^n \mathfrak{F}(u_i)$. Combining the above statements gives $\mathcal{W}_r^R \subset \mathfrak{S} \cap A_v$. ■

4. Cut graphs on branched covers

Let S be a compact hyperbolic surface. Consider an arbitrary $*$ -cover X of S . To prove that S is flexibly geometrically stable we need to construct a $*$ -graph Γ on X so that the complement $X \setminus \Gamma$ embeds into some cover of S . In addition the total length of Γ should be controlled by the branching degree of X when both quantities are normalized relative to the size of X . This motivates the following.

Definition. Let X be a $*$ -cover of S . A c -cut graph for some $c > 0$ is a $*$ -graph Γ on X satisfying

- (1) the singular set $s(X)$ is equal to the vertex set $\Gamma^{(0)}$,
- (2) if e_1 and e_2 are two edges incident at the vertex $x \in \Gamma^{(0)}$ and consecutive in the cyclic order on the link of the graph Γ at the vertex x induced by its embedding in X , then the angle between e_1 and e_2 at x is at most π ,
- (3) the total edge length $l(\Gamma)$ is bounded above by $c \text{Area}(X)$.

A cut graph on X is a c -cut graph for some $c > 0$.

The above condition (2) is equivalent to saying that every connected component of $X \setminus \Gamma$ has locally convex boundary.

The main goal of the current section is the following.

Theorem 4.1. *For every $c > 0$ there is a $\delta = \delta(c) > 0$ such that any $*$ -cover X with $\beta(X) < \delta \text{Area}(X)$ admits a c -cut graph.*

The graph Γ will be constructed by considering the Delaunay graph and then carefully removing some of its edges until the graph-theoretical degree at each vertex is roughly proportional to the local branching degree. The function c will depend on the topology and the metric of the surface S .

Proof of Theorem 4.1. Let $p : X \rightarrow S$ be a $*$ -cover whose branching degree satisfies $\beta(X) < \delta |X|$ for some sufficiently small $\delta > 0$ to be determined below. Recall that the branching degree $\beta(X)$ is defined to be $\sum_{v \in s(X)} (d_v - 1)$ where d_v is the index of the point v .

We may assume that the singular set $s(X)$ is non-empty for otherwise the empty graph is a c -cut graph for any $c > 0$. Consider the singular set $s(\tilde{X}) = q^{-1}(s(X))$. Since $s(\tilde{X}) \subset (p \circ q)^{-1}(*)$, this set is r -separated for some constant $r > 0$ depending only on the injectivity radius of the compact surface S .

Consider the family of Voronoi cells in \tilde{X} with respect to the vertex set $s(\tilde{X})$. The interior of every cell $A_{\tilde{v}}$ with $\tilde{v} \in s(\tilde{X})$ embeds into X according to Corollary 3.5. Denote

$A_v = q(A_{\tilde{v}})$ where $v \in s(X)$ and \tilde{v} is any vertex in $q^{-1}(v)$. The family of the A_v 's with $v \in s(X)$ forms a tessellation of X and satisfies

$$\sum_{v \in s(X)} \text{Area}(A_v) = \text{Area}(X). \tag{I}$$

Recall that a half-space in \tilde{X} is the closure of a connected component of the complement of some bi-infinite geodesic line in \tilde{X} . A *half-cell* \mathcal{C} at the vertex $v \in s(X)$ is $q(A_{\tilde{v}} \cap \mathfrak{H})$ for some half-space \mathfrak{H} at any vertex $\tilde{v} \in q^{-1}(v)$.

Let D be the geometric realization of the Delaunay graph in \tilde{X} with respect to the vertex set $s(\tilde{X})$. This notion is discussed in Section 3 above. Let E denote the projection $q(D)$ of the Delaunay graph to the $*$ -cover X . It follows from Proposition 3.6 that E is a $*$ -graph in X . This $*$ -graph satisfies condition (1) by its construction and condition (2) according to Proposition 3.7. The required cut graph will be obtained by discarding some of the edges of E while making sure condition (2) continues to hold.

Let L_v denote the set of edges of the graph E incident at the vertex $v \in s(X)$. Our next step is to find a subset M_v of L_v of size $|M_v| \leq 4d_v$ that continues to satisfy condition (2) of cut graphs and such that

$$\varphi \sum_{e \in M_v} \sinh^2(l(e)/4) \leq 2 \text{Area}(A_v) \tag{II}$$

where $\varphi > 0$ is the constant given in Lemma 3.8 with respect to the parameter r .

Consider the subset K_v of L_v consisting of all edges e such that

$$2\varphi \sinh^2(l(e)/4) \leq \text{Area}(\mathcal{C}_e)$$

for *some* half-cell \mathcal{C}_e at v containing the edge e in its interior. We claim that K_v satisfies condition (2) of cut graphs. If this was not the case then there would exist a half-cell \mathcal{C} at v which does not contain any edge of K_v in its interior. However, at least one edge $e \in L_v$ is contained in \mathcal{C} and therefore no such edge is in K_v . Hence $2\varphi \sinh^2(l(e)/4) > \text{Area}(\mathcal{C})$ for every edge $e \in L_v$ contained in the half-cell \mathcal{C} . This contradicts Lemma 3.8.

Take M_v to be a subset of K_v minimal with respect to inclusion that still satisfies condition (2) of cut graphs. The minimality of M_v implies that any half-cell \mathcal{C} at v contains at most two edges from M_v . Since A_v is the union of $2d_v$ half-cells at v it follows that $|M_v| \leq 4d_v$. Therefore any point x of the Voronoi cell A_v belongs to at most four half-cells from the family $\{\mathcal{C}_e\}_{e \in M_v}$, and inequality (II) follows.

Consider the sub- $*$ -graph Γ of the $*$ -graph E embedded in X with vertex set $s(X)$ and edge set $\bigcup_{v \in s(X)} M_v$. Observe that conditions (1) and (2) of cut graphs hold true. It remains to bound the total length $l(\Gamma)$ from above as in condition (3) and to determine the precise value of δ .

Let \mathcal{E} denote the set of edges of Γ . Since $\frac{1}{2} \sum_{v \in V} |M_v| \leq |\mathcal{E}| \leq \sum_{v \in V} |M_v|$ and $2d_v \leq |M_v| \leq 4d_v$ for every $v \in V$, we have

$$\beta(X) \leq \beta(X) + |s(X)| \leq |\mathcal{E}| \leq 4(\beta(X) + |s(X)|) \leq 8\beta(X).$$

Making use of the area formula (I) and summing (II) over all vertices $v \in s(X)$ gives

$$\varphi \sum_{e \in \mathcal{E}} \sinh^2(l(e)/4) \leq 2 \text{Area}(X).$$

The function $\sinh^2(\cdot/4)$ is convex. Jensen's inequality gives

$$\varphi \sinh^2\left(\frac{l(\Gamma)}{4|\mathcal{E}|}\right) \leq \frac{2 \text{Area}(X)}{|\mathcal{E}|}$$

and after rearranging

$$l(\Gamma) \leq 4|\mathcal{E}| \operatorname{arcsinh}\left(\sqrt{\frac{2}{\varphi} \frac{\text{Area}(X)}{|\mathcal{E}|}}\right).$$

To conclude the proof define the function $c : \delta \mapsto c_\delta$ by

$$c_\delta = 32\delta \operatorname{arcsinh}\left(\sqrt{\frac{2}{\varphi\delta}}\right).$$

Note that c is monotone and that $\lim_{\delta \rightarrow 0} c_\delta = 0$. Finally, observe that

$$l(\Gamma) \leq c \frac{\beta(X)}{\text{Area}(X)} \text{Area}(X) \leq c_\delta \text{Area}(X).$$

Therefore Γ is a c -cut graph provided δ is so small that $c_\delta < c$. ■

5. Boundaries and non-separating closed curves

Let S be a compact hyperbolic surface with a non-separating family of disjoint simple curves. We study a necessary and a sufficient condition due to Walter D. Neumann [12, Lemma 3.2] for the existence of a cover for S such that the preimages of the given family of curves have prescribed degrees.

1-manifolds

A connected closed 1-manifold is homeomorphic to S^1 . A closed compact 1-manifold is a disjoint union of finitely many homeomorphic copies of S^1 . The set of connected components of a 1-manifold B is denoted $\pi_0(B)$. 1-manifolds are orientable and admit $2^{|\pi_0(B)|}$ orientations.

The boundary ∂F of a compact surface F is a closed 1-manifold. If the surface F is oriented then ∂F inherits an induced orientation.

The Neumann lemma

We use $\chi(F)$ to denote the Euler characteristic of the surface F . Fix a degree $N \in \mathbb{N}$.

Lemma 5.1 ([12, Lemma 3.2]). *Let F be a compact connected orientable surface of positive genus. Let $p : B \rightarrow \partial F$ be a covering map for some closed 1-manifold B . The following two statements are equivalent:*

- (1) *There exist a connected surface R with $\partial R \cong B$ and a covering map $r : R \rightarrow F$ of degree N with $r|_B = p$.*
- (2) *$|\pi_0(B)|$ has the same parity as $\chi(F)N$ and for every connected component γ of ∂F the total degree of the covering map p on $p^{-1}(\gamma)$ is N .*

The implication (1) \Rightarrow (2) holds true even without assuming that F has positive genus, or that R is connected.

We remark that in [12] the equivalence of (1) and (2) is proved under the assumption of positive genus. However, the proof of the direction (1) \Rightarrow (2) clearly applies in any genus.

We obtain the following consequence of Lemma 5.1. It will be used towards the proof of Theorem 1.2 presented in Section 6.

Corollary 5.2. *Let S be a closed hyperbolic surface. Let $p : B \rightarrow A$ be a covering map of closed compact and oriented 1-manifolds where A is embedded and is non-separating in S . Then there exist*

- *a covering map $r : C \rightarrow S$,*
- *an oriented subsurface R of C such that $\partial R \cong B$ as an oriented 1-manifold with the induced orientation and with $r|_B = p$*

if and only if $|\pi_0(B)|$ is even and $B = B^+ \cup B^-$ where

- *$p^+ = p|_{B^+}$ is orientation preserving,*
- *$p^- = p|_{B^-}$ is orientation reversing,*
- *for every connected component α of A the total covering degrees of p^+ and of p^- on $p^{-1}(\alpha)$ are equal to each other.*

If the latter conditions are satisfied then the cover r can be taken of degree equal to four times the maximum of the total degree of p over any connected component of A .

We consider a topological reduction needed to apply Neumann’s lemma towards the proof of Corollary 5.2. The main issue is to allow for the degree of p to vary over the different connected components α of A .

Let \bar{S} denote the completion of $S \setminus A$. In particular, \bar{S} is a compact connected orientable surface with boundary. Consider the quotient map $f : \bar{S} \rightarrow S$ with

$$f(\partial\bar{S}) = A.$$

Choose orientations for the two surfaces S and \bar{S} such that f is orientation preserving on the interior of \bar{S} . Equip $\partial\bar{S}$ with the boundary orientation induced from that of \bar{S} . In particular,

$$\partial\bar{S} = \bar{A}^+ \amalg \bar{A}^-$$

in such a way that the restriction of f to \bar{A}^+ and to \bar{A}^- is, respectively, orientation preserving and orientation reversing.

Observe that if $q : \bar{C} \rightarrow \bar{S}$ is any orientation preserving cover and $\bar{E} \subset \partial\bar{C}$ is any connected closed 1-manifold then the composition $f \circ q : \bar{E} \rightarrow A$ is orientation preserving if and only if $q(\bar{E}) \subset \bar{A}^+$. Finally, let $\tau : \partial\bar{S} \rightarrow \partial\bar{S}$ be the homeomorphism determined by

$$\tau(\bar{A}^+) = \bar{A}^- \quad \text{and} \quad f \circ \tau = f.$$

The main ideas of the following proof are summarized in Figure 2.

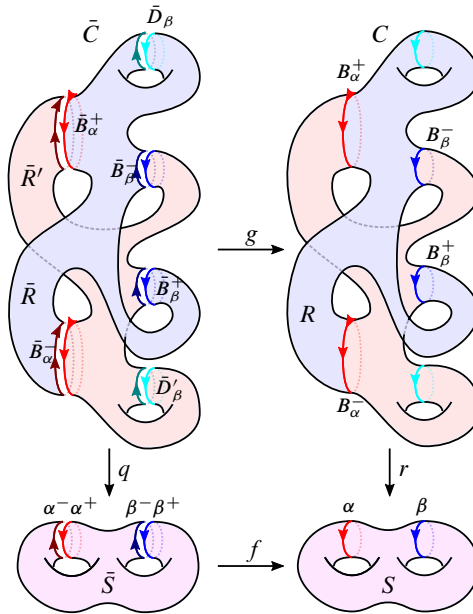


Fig. 2. The maps r and q are covers, the maps f and g are quotients, and the diagram is commutative. In this example $A = \{\alpha, \beta\}$ and $\bar{A}^\pm = \{\alpha^\pm, \beta^\pm\}$. The curve α satisfies $N_\alpha = M$ and the curve β satisfies $N_\beta < M$.

Proof of Corollary 5.2. The “only if” direction. Consider a cover $r : C \rightarrow S$ and a subsurface R of C with $\partial R \cong B$ as in the statement of Corollary 5.2.

Let \bar{C} denote the compact surface with boundary obtained as the completion of $C \setminus r^{-1}(A)$. Note that \bar{C} is disconnected. There is a natural quotient map $g : \bar{C} \rightarrow C$ as well as a covering map $q : \bar{C} \rightarrow \bar{S}$ satisfying

$$f \circ q = r \circ g.$$

Denote $\bar{E} = \partial\bar{C}$ so that \bar{E} is a compact closed 1-manifold and $f \circ q : \bar{E} \rightarrow A$ is a covering map of 1-manifolds. In particular, $\partial R \cong B$ is a 1-submanifold of $g(\bar{E})$.

Let \bar{R} be the subsurface of \bar{C} corresponding to R . We may apply Lemma 5.1 to the restricted covering map $q : \bar{R} \rightarrow \bar{S}$ and to the boundary $\partial\bar{R} = \bar{R} \cap \bar{E}$. The result follows

in this direction by letting

$$B^+ = g(q^{-1}(\bar{A}^+)) \cap B \quad \text{and} \quad B^- = g(q^{-1}(\bar{A}^-)) \cap B.$$

Observe that the restrictions of r to B^+ and to B^- are orientation preserving and orientation reversing respectively. The connected components of the 1-manifold $\bar{D} = \bar{E} \setminus g^{-1}(\partial R)$ occur in matching pairs δ^+ and δ^- with $g(\delta^+) = g(\delta^-)$. Therefore the total degrees of $p^+ = p|_{B^+}$ and of $p^- = p|_{B^-}$ agree over every connected component of A , and $|\pi_0(B)|$ is even.

The “if” direction. Assume first that the surface \bar{S} has positive genus. Assume that $|\pi_0(B)|$ is even and that there exists a decomposition $B = B^+ \cup B^-$ as in the statement. For each connected component α of A , let $N_\alpha \in \mathbb{N}$ denote the common total degree of p^+ and of p^- over α . Denote $M = \max N_\alpha$ taken over all such α 's.

We first enlarge B to a 1-manifold \bar{E} with a covering map $t : \bar{E} \rightarrow \partial \bar{S}$ such that over every component of A the degree is the same. If α is a component of A satisfying $N_\alpha < M$, let D_α denote the disjoint union of two circles D_α^+ and D_α^- , equipped with covering maps $t_\alpha^\pm : D_\alpha^\pm \rightarrow \bar{\alpha}^\pm$ of degree $M - N_\alpha$ (where $\bar{\alpha}^\pm$ are the two components of $f^{-1}(\alpha)$). Denote $D = D^+ \cup D^-$ where $D^\pm = \bigcup_\alpha D_\alpha^\pm$ with the union being taken over all components α of A satisfying $N_\alpha < M$.

Let \bar{E} be the disjoint union of B and D , and let $t : \bar{E} \rightarrow \partial \bar{S}$ be the covering map defined by t_α^\pm on each component D_α^\pm of D , and on B by

$$t(B^+) = \bar{A}^+, \quad t(B^-) = \bar{A}^-, \quad \text{and} \quad f \circ t = p.$$

The covering map t from \bar{E} to $\partial \bar{S}$ has total degree M over every connected component of $\partial \bar{S}$. We may apply Lemma 5.1 with respect to the surface \bar{S} and the covering map t to obtain a connected surface \bar{R} with $\partial \bar{R} \cong \bar{E}$ and a covering map $q : \bar{R} \rightarrow \bar{S}$ of degree M with $q|_{\bar{E}} = t$. Note that the parity hypothesis of Lemma 5.1 is satisfied since both $|\pi_0(B)|$ and $\chi(\bar{S})$ are even.

Let $\bar{E}' = B' \cup D'$ be a homeomorphic copy of the 1-manifold \bar{E} considered with the twisted covering map $t' = \tau \circ t$. This covering map t' has degree M . We may apply Lemma 5.1 again with respect to the covering map t' to obtain a corresponding surface cover $q' : \bar{R}' \rightarrow \bar{S}$.

Denote $\bar{C} = \bar{R} \cup \bar{R}'$. Allowing for a slight abuse of notation, let $q : \bar{C} \rightarrow \bar{S}$ denote the union of the two covering maps $q : \bar{R} \rightarrow \bar{S}$ and $q' : \bar{R}' \rightarrow \bar{S}$. There is a quotient map $g : \bar{C} \rightarrow C$ defined by identifying boundary components of \bar{C} as follows. Identify D^+ with D^- and D'^+ with D'^- . Moreover, identify B^+ with B'^- and B'^+ with B^- . The identification is performed respecting the covering map q . In particular, there exists a resulting covering map $r : C \rightarrow S$ with $r \circ g = f \circ q$.

The cover C has degree $2M$, which is equal to the maximal total degree of p . Let $R = g(\bar{R})$. This completes the proof of this direction under the assumption that \bar{S} has positive genus.

Now, assume that the surface \bar{S} has genus 0. Let S' be the surface obtained by splitting S along a curve $\alpha \in A$ and gluing two copies of the resulting surface along their

boundaries. Let $q' : S' \rightarrow S$ be the obvious covering map mapping the two copies back to S . Each curve in A lifts to two curves in S' . Let $A' \subset q'^{-1}(A)$ be a collection of curves containing exactly one lift for each curve in A . Note that $q'_{|A'} : A' \rightarrow A$ is a homeomorphism as each curve is covered with degree 1. The completion \bar{S}' of $S' \setminus A'$ is connected and has positive genus.

Denote by $p' : B \rightarrow A'$ the map $p' = (q'_{|A'})^{-1} \circ p$. The surface \bar{S}' has positive genus and so the previous argument can be applied to get a covering map $r' : C \rightarrow S'$ of degree $2M$ and an oriented subsurface R of C such that $\partial R \cong B$ and $r'_{|B} = p'$. The composition $r = q' \circ r' : C \rightarrow S$ is as required, and has degree $4M$. ■

6. Quantitatively capping off surfaces with boundary

Let S be a closed hyperbolic surface of genus $g \geq 2$. The goal of the current section is to prove Theorem 1.2, restated below for the reader's convenience.

Theorem. *Let R be a surface with boundary which is isometrically embedded in some cover of S . If the boundary ∂R is locally convex then R can be isometrically embedded in a cover Q of S making diagram (A) commute and such that*

$$\text{Area}(Q \setminus R) \leq b_S l(\partial R)$$

where $b_S > 0$ is a constant depending only on S .

We will deal with the problem of capping off a boundary component of R in a controlled way by working in a certain combinatorial framework. More precisely, we will consider surfaces tessellated by isometric copies of a particularly nice fundamental domain. The generic situation as in the above theorem can be easily reduced to this combinatorial framework.

Surfaces with locally convex boundary

We point out the following consequence of Lemma 3.1 and of the Cartan–Hadamard theorem.

Lemma 6.1. *Let R be a surface with boundary. Then the boundary ∂R is locally convex and R is isometrically embedded in some cover of S if and only if R admits a local isometry to S .*

Here we use the term *local isometry* in the metric space rather than Riemannian geometry sense.

Proof of Lemma 6.1. If R is isometrically embedded in some cover $p : C \rightarrow S$ and its boundary ∂R is locally convex then the restriction of p to R is a local isometry.

Conversely, assume that R admits a local isometry f into S . This implies that the boundary ∂R is locally convex. The map $f_* : \pi_1(R) \rightarrow \pi_1(C)$ is injective and f lifts to

an isometric embedding $F : \tilde{R} \rightarrow \tilde{S} \cong \mathbb{H}$ of the universal covers according to Lemma 3.1. The map F descends to an isometric embedding of R into the cover of S that corresponds to the subgroup $f_*\pi_1(R) \leq \pi_1(S)$. ■

The family of closed geodesic curves $\lambda_0, \dots, \lambda_{2g-1}$

Lemma 6.2. *The surface S admits a family $\lambda_0, \dots, \lambda_{2g-1}$ of non-separating simple closed geodesic curves such that*

- (1) *the geodesic curves λ_i and λ_j with $i \neq j$ are disjoint unless $|i - j| = 1$, in which case they intersect at a single point,*
- (2) *the complement $S \setminus (\lambda_0 \cup \dots \cup \lambda_{2g-1})$ is a topological disc.*

Proof. Consider a family $\gamma_0, \dots, \gamma_{2g-1}$ of non-separating simple closed curves positioned on the surface S as described in Figure 3 below.




Fig. 3. The simple closed curves $\gamma_0, \dots, \gamma_{2g-1}$ are given by $\gamma_0 = e_0$, $\gamma_{2g-1} = e_{2g-1}$, and $\gamma_i = e_i \bar{e}'_i$ for any other i .

Let λ_i be the geodesic representative of the simple closed curve γ_i for every i in $\{0, \dots, 2g - 1\}$. Since geodesic curves are in minimal position we have

$$|\lambda_i \cap \lambda_j| = \begin{cases} 1 & \text{if } |i - j| = 1, \\ 0 & \text{if } |i - j| > 1, \end{cases}$$

as required for statement (1) of the lemma. The completion \mathcal{P} of the complement $S \setminus (\lambda_0 \cup \dots \cup \lambda_{2g-1})$ is a connected surface with boundary admitting $4(2g - 1)$ vertices, $2 + 2 + 4(2g - 2) = 8g - 4$ edges and a single face. The Euler characteristic of \mathcal{P} is equal to

$$\chi(\mathcal{P}) = 4(2g - 1) - (8g - 4) + 1 = 1.$$

Therefore \mathcal{P} is a topological disc as required for statement (2). ■

Let \mathcal{P} denote the completion of the complement $S \setminus (\lambda_0 \cup \dots \cup \lambda_{2g-1})$. Then \mathcal{P} is a compact convex hyperbolic $(8g - 4)$ -sided polygon. Moreover, \mathcal{P} is isometric to a fundamental domain for the action of the fundamental group $\pi_1(S)$ on the hyperbolic plane. We introduce the following notations for the edges of \mathcal{P} (see Figure 3):

- e_0 and \bar{e}_0 are the two geodesic sides of \mathcal{P} that correspond to the curve λ_0 ,
- e_{2g-1} and \bar{e}_{2g-1} are the two geodesic sides of \mathcal{P} that correspond to the curve λ_{2g-1} ,

- e_i, e'_i, \bar{e}_i and \bar{e}'_i are the four sides of \mathcal{P} that correspond to the curve λ_i for any other $i \in \{1, \dots, 2g - 2\}$.

The surface S can be recovered by identifying every geodesic side e of the fundamental polygon \mathcal{P} with the geodesic side of opposite orientation denoted \bar{e} .

We will keep the above notations throughout the remainder of Section 6.

Tessellations and \mathcal{P} -surfaces

Let $\mathcal{T}(C, \mathcal{P})$ denote the induced tessellation of a given cover C of the surface S by isometric copies of the fundamental polygon \mathcal{P} .

The link of every vertex of the tessellation $\mathcal{T}(C, \mathcal{P})$ is a cycle graph of size 4. In other words, locally, exactly four polygons of $\mathcal{T}(C, \mathcal{P})$ meet at every vertex. Another crucial property of the tessellation $\mathcal{T}(C, \mathcal{P})$ is that the concatenation of any pair of edges incident at a given vertex v and *not* consecutive in the cyclic ordering determined by the link of v is a local geodesic. Such a concatenation of two edges is, up to edge orientation, of the form $e_0e_0, e_{2g-1}e_{2g-1}, e_i e'_i$ or $e'_i e_i$ for some $i \in \{1, \dots, 2g - 2\}$.

Definition. A \mathcal{P} -surface R is a subsurface of some cover C of the surface S tiled by polygons from the tessellation $\mathcal{T}(C, \mathcal{P})$. The induced tessellation of R will be denoted $\mathcal{T}(R, \mathcal{P})$.

We remark that a \mathcal{P} -surface could in general be disconnected.

Proposition 6.3. *Let R be a \mathcal{P} -surface. Then every geodesic arc e of $\partial\mathcal{P}$ appears along the boundary ∂R the same number of times with each orientation e and \bar{e} .*

Proof. Every geodesic arc e appears along the boundary of the polygon \mathcal{P} once with each orientation. Moreover, in the tessellation of R by isometric copies of \mathcal{P} every interior edge is accounted for once with each orientation. The result follows. ■

Given a \mathcal{P} -surface R it follows from the properties of the tessellation $\mathcal{T}(R, \mathcal{P})$ discussed above that ∂R is locally convex if and only if the link of any vertex v of $\mathcal{T}(R, \mathcal{P})$ that lies on the boundary ∂R is a path graph of length at most 2. In other words, locally, at most two polygons meet at a vertex in the boundary, forming together an angle of at most π . This observation motivates the following.

Proposition 6.4. *Let R be a \mathcal{P} -surface with locally convex boundary. Let α and β be two totally geodesic closed curves/maximal geodesic arcs along the boundary ∂R that map to the same geodesic curve/arc on S with opposite orientations. Then the surface R' obtained by identifying α and β is a \mathcal{P} -surface with locally convex boundary.*

Proof. The fact that R' has a locally convex boundary is clear when α and β are totally geodesic boundary curves.

Consider the case where α and β are maximal geodesic arcs. The vertices representing the end-points of the arcs α and β in the tessellation $\mathcal{T}(R, \mathcal{P})$ have links of size 1 in R .

Therefore the boundary components of R' that have been modified by the identification of α and β have links of size at most 2 in R and so remain locally convex as required.

The \mathcal{P} -surface structure of R gives rise to a local isometry $f : R \rightarrow S$. It is compatible with the identification of α and β and descends to a local isometry $f' : R' \rightarrow S$. It follows from Lemma 6.1 that R' isometrically embeds into some cover of S . It is clear that R' is tiled by polygons from the tessellation $\mathcal{T}(C, \mathcal{P})$. Therefore R' is a \mathcal{P} -surface. ■

The notion of \mathcal{P} -surfaces is useful for our purposes since a surface with boundary as in Theorem 1.2 can always be embedded inside a \mathcal{P} -surface in an efficient way.

Proposition 6.5. *Let R be a surface with boundary which is isometrically embedded in some cover of S . If the boundary ∂R is locally convex then R can be isometrically embedded in a \mathcal{P} -surface Q with locally convex boundary making diagram (A) commute and such that*

$$\text{Area}(Q \setminus R) \leq d_S l(\partial R) \quad \text{and} \quad l(\partial Q) \leq d_S l(\partial R)$$

where $d = d_S > 0$ is a constant depending only on the surface S .

Moreover, we may assume, if desired, that any geodesic subarc of ∂Q has at most two edges and that no boundary component of ∂Q is totally geodesic.

The following proof is inspired by and relies on ideas of Patel’s work [13].

Proof of Proposition 6.5. We assume that R has non-empty boundary for otherwise there is nothing to prove. Since R is embedded in some cover C of S , we may identify $\pi_1(R)$ with a certain infinite index subgroup of $\pi_1(S)$. Let R' be the cover of S corresponding to that subgroup. In particular, R is an embedded subsurface of R' and every connected component of $R' \setminus R$ retracts to a boundary component of ∂R . Let Q' be the \mathcal{P} -surface consisting of all polygons in $\mathcal{T}(R', \mathcal{P})$ that intersect R non-trivially.

The boundary of the \mathcal{P} -surface Q' need not yet be locally convex, as it may have *bad* vertices whose links are paths of length 3. As explained in [13, Theorem 3.1], since R is locally convex, there are no consecutive bad vertices in $\partial Q'$. One can now form the locally convex \mathcal{P} -surface Q'' by attaching polygons along the geodesic subarcs of $\partial Q'$ emanating from bad vertices. For the full details we refer the reader to Patel’s argument in [13, Theorem 3.1].

Remark. We emphasize that, at least formally speaking, there are a few minor differences between our situation and that of [13]:

- A right-angled pentagon is used in [13] while our polygon \mathcal{P} is $(8g - 4)$ -sided and is not assumed to be right-angled.¹
- The quotient of the hyperbolic plane by the action of a single hyperbolic element is considered in the context of [13, Theorem 3.1] while we consider locally convex boundary components of \mathcal{P} -surfaces.

¹In fact, \mathcal{P} can be chosen to be right-angled in our case as well. However, making this choice is not essential to the proof.

Patel’s arguments rely only on certain properties of links in the tessellation $\mathcal{T}(R', \mathcal{P})$ that were mentioned above and hold true in our case as well. Namely, all links have size 4, and the concatenation of any two non-consecutive edges is a local geodesic. The same proof goes through.

We conclude that the \mathcal{P} -surface Q'' has locally convex boundary and satisfies $Q' \subset Q'' \subset R'$. Proceeding with the proof of Proposition 6.5, the boundary of the \mathcal{P} -surface Q'' may contain geodesic arcs longer than two edges or totally geodesic boundary components. If desired, this can be overcome simply by attaching an additional layer of polygons to Q'' . Namely, let Q consist of all polygons in $\mathcal{T}(R', \mathcal{P})$ that intersect Q'' or its boundary non-trivially. The boundary of Q remains locally convex and satisfies the requirements as in the statement.

An area estimate analogous to [13, Theorem 4.3] shows that $\text{Area}(Q' \setminus R) \leq d'_S l(\partial R)$ for some constant $d'_S > 0$. The total boundary length of Q' is bounded above by the number of polygons meeting $Q' \setminus R$ times the perimeter of the polygon \mathcal{P} . In particular, $l(\partial Q') \leq d''_S l(\partial R)$ for some other constant $d''_S > 0$. Repeating this argument twice for the pair of \mathcal{P} -surfaces Q'' and Q gives the required linear upper bounds on $\text{Area}(Q \setminus R)$ and $l(\partial Q)$ in terms of some positive constant $d_S > 0$. ■

Capping off boundary components

We now have all the required machinery to complete the proof of Theorem 1.2 of the introduction.

Proof of Theorem 1.2. Let R be a surface with locally convex boundary which is embedded in some cover of S . Making use of Proposition 6.5 it is possible to find a \mathcal{P} -surface Q_1 containing an embedded copy of R such that $\text{Area}(Q_1 \setminus R)$ as well as $l(\partial Q_1)$ are bounded above by $d_S l(\partial R)$ where $d_S > 0$ is a positive constant. Moreover, any geodesic subarc of ∂Q_1 has at most two edges and no component of ∂Q_1 is geodesic. Recall that \mathcal{P} is the convex polygonal fundamental domain obtained as the complement of the system of curves $\lambda_0, \dots, \lambda_{2g-1}$ on the surface S .

The boundary components of Q_1 are piecewise geodesic and map to the arcs e_0, \dots, e_{2g-1} . Each geodesic boundary arc of length 1 is labeled by either e_i, e'_i or their inverses. Disregarding edge orientations, each arc of length 2 is labeled by either $e_0 e_0, e_{2g-1} e_{2g-1}, e_i e'_i$, or $e'_i e_i$. We say that a geodesic boundary arc has *odd label* or *even label* if i is odd or even, respectively.

The proof proceeds in three steps.

(1) *Identifying length 2 geodesic boundary arcs with even labels.* Consider a geodesic boundary arc α of length 2 and even label. We will assume that with the induced orientation, α reads $e_i e'_i$ with i even. The other cases are treated analogously.

The midpoint of the length 2 arc α has an incident edge contained in the interior of Q_1 and labeled e_{i+1} . Trace a geodesic arc γ on the surface Q_1 starting at this edge with labels

alternating between e_{i+1} and e'_{i+1} . It runs transverse to geodesic arcs of the two forms $e_i e'_i$ and $e_{i+2} e'_{i+2}$ until it reaches the boundary of Q_1 again, necessarily at the midpoint of some length 2 geodesic arc.

There are two possibilities to consider. If γ reaches the boundary of Q_1 at the midpoint of a geodesic arc β labeled $e_i e'_i$ then we may simply identify the two arcs α and β to reduce the number of even boundary components of length 2.

The other possibility is that γ reaches the midpoint of a geodesic arc β labeled $e_{i+2} e'_{i+2}$. Without loss of generality, assume that the last edge along γ is labeled e_{i+1} . If this is the case, let D be the \mathcal{P} -surface obtained by gluing two copies of the polygon \mathcal{P} along the edge e'_{i+1} so that D has two boundary geodesic arcs labeled $e_i e'_i$ and $e_{i+2} e'_{i+2}$. Attach the boundary labeled $e_{i+2} e'_{i+2}$ in D to the arc β in Q_1 . The resulting \mathcal{P} -surface is locally convex by Proposition 6.4. The geodesic γ can now be extended by one edge, so that it reaches the midpoint of the boundary geodesic arc labeled $e_i e'_i$ in D . We can now proceed as before.

While this operation might increase the total length of geodesic boundary arcs with odd labels or the total number of edges with even labels, it will reduce the number of geodesic boundary arcs of length 2 and even label.

Let Q_2 denote the resulting surface. We point out that Q_2 is again a \mathcal{P} -surface with locally convex boundary by Proposition 6.4.

(2) *Identifying the remaining geodesic boundary arcs of length 1 with even labels.* Every geodesic boundary arc of Q_2 with an even label has length 1. Moreover, Q_2 has the same number of edges labeled e_i and \bar{e}_i for every even i by Proposition 6.3. Choose any bijection between these two sets and identify edges in pairs. Let Q_3 denote the resulting surface. Once again Proposition 6.4 implies that Q_3 is a \mathcal{P} -surface with locally convex boundary.

(3) *Attaching a surface along the totally geodesic boundary curves with odd labels.* Every boundary geodesic arc of Q_3 is locally convex and has an odd label. This implies that Q_3 has totally geodesic boundary and every boundary component maps to a power of one of the non-separating simple closed geodesic curves λ_i with some orientation and i odd.

One direction of Corollary 5.2 to Neumann’s lemma implies that the number of boundary components of Q_3 is even and that the total degree over each curve λ_i with one orientation is equal to the total degree over λ_i in the opposite orientation. This fact and the other direction of Corollary 5.2 allow us to construct a \mathcal{P} -surface T with $\text{Area}(T)$ linearly bounded in $l(Q_3)$ whose boundary components are the same as those of Q_3 but with opposite orientation. Let Q be the cover of S obtained by identifying the boundary of Q_3 with that of T in the obvious way.

To conclude observe that R embeds in the cover Q making diagram (A) commute and that $\text{Area}(Q \setminus R)$ is bounded above linearly in $l(\partial Q_1)$ and hence in $l(\partial R)$, as required. ■

7. Proof of geometric flexible stability

Let S be a compact hyperbolic surface. Consider some $*$ -cover $p : X \rightarrow S$. We rely on Theorem 1.2 established in the previous section to complete the proof of Theorem 2.1 and therefore of our main result, Theorem 1.1.

Our strategy is to construct a cut graph Γ on X and then cap off the complement $X \setminus \Gamma$ to obtain an unramified cover by making use of Theorem 1.2.

Since X need not be a cover in the usual sense, it might contain certain pathological closed curves that cannot exist on a cover of S . For example, X might admit a simple closed curve γ such that a lift of $p \circ \gamma$ to \tilde{S} admits self-intersections. It is clear that if X is ε -reparable then such a curve has to be eliminated. This motivates the following.

Proposition 7.1. *Let Γ be a cut graph on X , and C be any connected component of $X \setminus \Gamma$. Then C embeds in some cover of S as a subsurface with locally convex boundary.*

Proof. Let C be a connected component of $X \setminus \Gamma$. It follows from the definition of cut graphs that C has no singular points of $s(X)$ in its interior and that its boundary is locally convex. In particular, C is non-positively curved and the map $f = p|_C : C \rightarrow S$ is a local isometry. The result follows from Lemma 6.1. ■

We are now ready to prove that closed hyperbolic surfaces are flexibly geometrically stable.

Proof of Theorem 2.1. Let the constant $\varepsilon > 0$ be given. Take $c > 0$ to be sufficiently small so that $c \max\{1, b_S\} \leq \varepsilon$ where b_S is the constant as in Theorem 1.2. Let $\delta = \delta(c)$ be the constant provided by Theorem 4.1 such that any $*$ -cover X with $\beta(X) < \delta \text{Area}(X)$ admits a c -cut graph.

We claim that δ is as required in the definition of flexible geometric stability with respect to the given $\varepsilon > 0$. To see this consider a $*$ -cover $p : X \rightarrow S$ with $\beta(X) < \delta \text{Area}(X)$. We need to show that X is ε -reparable. Let Γ be a c -cut graph on X . In particular,

$$l(\Gamma) \leq c \text{Area}(X) \leq \varepsilon \text{Area}(X).$$

Consider the complement $C = X \setminus \Gamma$. Every connected component of C isometrically embeds into some cover of S according to Proposition 7.1. Relying on Theorem 1.2 we find a cover $q : Q \rightarrow S$ admitting an isometrically embedded copy of C on which the restriction of q agrees with the map p , so diagram (B) commutes. Moreover

$$\text{Area}(Q) - \text{Area}(C) \leq b_S l(\partial C) \leq b_S c \text{Area}(X) \leq \varepsilon \text{Area}(X).$$

This completes the verification that X is indeed ε -reparable. ■

The fact that Theorem 2.1 implies our main result Theorem 1.1 is contained in Proposition 2.4.

Appendix A. Voronoi and Delaunay on singular planes

We generalize some of the basic properties of the Voronoi tessellation and the Delaunay graph from the classical Euclidean and hyperbolic cases to the framework of a “hyperbolic plane with singularities”.

In particular, we present the detailed proofs of Propositions 3.3, 3.4, 3.6 and 3.7 that were merely stated without proof in Section 3.

Hyperbolic planes with singularities

Let S be a compact hyperbolic surface and $p : X \rightarrow S$ a fixed $*$ -cover. Let $q : \tilde{X} \rightarrow X$ denote the universal cover of X equipped with the pullback length metric $d_{\tilde{X}}$. The singular set of \tilde{X} is given by $s(\tilde{X}) = q^{-1}(s(X))$. The *Voronoi cell* A_v at any vertex $v \in s(\tilde{X})$ and the *Delaunay graph* D were defined in Section 3.

We find it useful to introduce the following additional notation. Let $m(v, u)$ denote the set of *midpoints*

$$m(v, u) = \{x \in \tilde{X} : d_{\tilde{X}}(x, v) = d_{\tilde{X}}(x, u)\}$$

between any two given points $v, u \in \tilde{X}$. It is well-known [2, Section 7.21] that the set of midpoints of any pair of points in the hyperbolic plane is the perpendicular bisector to the geodesic arc connecting these two points.

Two vertices v_1 and v_2 of the Delaunay graph are connected by an edge if and only if there exists a midpoint $x \in m(v_1, v_2)$ such that $d(x, u) > d(x, v_1) = d(x, v_2)$ for any other vertex $u \in s(\tilde{X}) \setminus \{v_1, v_2\}$.

Lemma A.1. *Let $x_1, x_2, x_3 \in \tilde{X}$ be three distinct points. Then there is at most a single closed metric ball $B \subset \tilde{X}$ with $\overset{\circ}{B} \cap s(\tilde{X}) = \emptyset$ and $\{x_1, x_2, x_3\} \subset \partial B$.*

Proof. Assume that $B \subset \tilde{X}$ is a closed ball as in the statement of the lemma. Let $T \subset \tilde{X}$ be the geodesic triangle with vertices x_1, x_2, x_3 . The convexity of the ball B implies that $T \subset B$. Making use of Corollary 3.2 we may regard B as being isometrically embedded in the hyperbolic plane. The center x_0 of the ball B is determined by the triangle T . More precisely, x_0 is the mutual intersection point of the three midpoint bisectors to the edges of T . The same reasoning shows that any other ball $B' \subset \tilde{X}$ as in the statement of the lemma necessarily has the same center as B and must therefore agree with B . ■

Proposition 3.3 says that every Voronoi cell is homeomorphic to a disc, is convex with a piecewise geodesic boundary and is determined by its Delaunay neighbors.

Proof of Proposition 3.3. Let $v \in s(\tilde{X})$ be a vertex of the Delaunay graph. The distance function $d_{\tilde{X}}(\cdot, u)$ is continuous for every point $u \in \tilde{X}$. Therefore the Voronoi cell A_v is closed. It is clear from the definition that $A_v \cap s(\tilde{X}) = \{v\}$. Since the singular set $s(\tilde{X})$ is co-bounded it follows that A_v is bounded and hence compact. Therefore there is a finite

subset of vertices $\{u_1, \dots, u_n\} \subset s(\tilde{X})$ for some $n \in \mathbb{N}$ such that the Voronoi cell A_v is given by

$$A_v = \{x \in \tilde{X} : d_{\tilde{X}}(x, v) \leq d_{\tilde{X}}(x, u_i) \forall i \in \{1, \dots, n\}\}.$$

If $x \in \partial A_v$ is a boundary point of A_v then at least one of the above inequalities must be an equality. In other words, $\partial A_v \subset A_v \cap \bigcup_{i=1}^n m(v, u_i)$.

For every boundary point $x \in \partial A_v$ let $F_x \subset \{u_1, \dots, u_n\}$ be the subset given by $u_i \in F_x$ if and only if $x \in m(v, u_i)$. Note that F_x is non-empty for every $x \in \partial A_v$. Every boundary point $x \in \partial A_v$ admits a closed ball $B_x \subset \tilde{X}$ centered at x with $\overset{\circ}{B}_x \cap s(\tilde{X}) = \emptyset$ and $\partial B_x \cap s(\tilde{X}) = F_x \cup \{v\}$. This relies on the fact that $x \in A_v$ and therefore $d(x, v) \leq d(x, u)$ for every vertex $u \in s(\tilde{X})$.

Every boundary point $x \in \partial A_v$ admits an open neighborhood U_x such that

$$A_v \cap U_x = \{y \in U_x : d_{\tilde{X}}(y, v) \leq d_{\tilde{X}}(y, u_i) \forall i \in F_x\}.$$

Assume without loss of generality that U_x is sufficiently small so that $U_x \subset B_x$. The ball B_x is convex and can be regarded as being isometrically embedded in the hyperbolic plane by Corollary 3.2. We conclude that $A_v \cap U_x$ can be isometrically identified with a neighborhood of the point x inside a *hyperbolic Voronoi cell*.

Note that $|F_x| = 1$ for all but finitely many boundary points $x \in \partial A_v$ according to Lemma A.1. This condition is open in the sense that given $x \in \partial A_v$ with $|F_x| = 1$ we have $F_x = F_y$ for every $y \in A_v \cap U_x$. We conclude that ∂A_v is locally isometric to a geodesic segment away from finitely many points where $|F_x| > 1$. Moreover, ∂A_v is locally convex at those points as well. In particular, the boundary of A_v is piecewise geodesic.

The above description of the boundary of A_v in terms of the local hyperbolic picture with respect to the open cover U_x shows that A_v is locally convex. The Cartan–Hadamard theorem implies that A_v is convex (see e.g. Lemma 3.1). The exponential map at v sets up a homeomorphism of A_v to a disc. This concludes the proof of items (1) and (2).

Note that a vertex u is adjacent to v in the Delaunay graph if and only if $u = u_i$ for some $i \in \{1, \dots, n\}$ and there is a boundary point $x \in \partial A_v$ with $F_x = \{u_i\}$. In other words, the Delaunay neighbors of v correspond to the geodesic sides of the cell A_v , and item (3) of Proposition 3.3 follows. ■

Proposition 3.4 says that the interiors of distinct Voronoi cells are disjoint.

Proof of Proposition 3.4. Assume that the intersection $A_{v,u} = A_v \cap A_u$ is non-empty. Every point $x \in A_{v,u}$ admits a closed metric ball $B_x \subset \tilde{X}$ centered at x with $\overset{\circ}{B}_x \cap s(\tilde{X}) = \emptyset$ and $\{v, u\} \subset \partial B_x$. The ball B_x is convex and so can be regarded as being isometrically embedded in the hyperbolic plane by Corollary 3.2.

Hyperbolic geometry implies that $m(v, u) \cap B_x$ is a geodesic segment for every point $x \in A_{v,u}$. Since $A_{v,u} \subset m(v, u)$, the convexity of $A_{v,u}$ implies that $A_{v,u} \cap B_x$ is a geodesic segment as well. The compactness of $A_{v,u}$ allows us to extract a finite cover using such metric balls. It follows that $A_{v,u}$ is a finite union of geodesic segments. Since $A_{v,u}$ is convex, it must be a single point or a single geodesic segment. In particular, $A_{v,u}$ has

empty interior and $\overset{\circ}{A}_v \cap \overset{\circ}{A}_u = \emptyset$. Since both A_v and A_u are topological discs, this implies $\partial A_v \cap \overset{\circ}{A}_u = \overset{\circ}{A}_v \cap \partial A_u = \emptyset$. We conclude that $A_{v,u} \subset \partial A_v \cap \partial A_u$ and that $A_{v,u}$ is a common geodesic side of the Voronoi cells A_v and A_u .

The second statement of Proposition 3.4 concerning adjacency in the Delaunay graph follows from the last paragraph of the proof of Proposition 3.3. ■

Proposition 3.6 deals with the Delaunay graph and shows it is embedded.

Proof of Proposition 3.6. Let v_1, u_1 and v_2, u_2 be vertices of the Delaunay graph such that v_i is adjacent to u_i . Let x_1 and x_2 be points in \tilde{X} such that for $i \in \{1, 2\}$ there is a closed metric ball $B_i \subset \tilde{X}$ centered at x_i and with $\overset{\circ}{B}_i \cap s(\tilde{X}) = \emptyset$ and $\partial B_i \cap s(\tilde{X}) = \{v_i, u_i\}$. Let γ_i with $i \in \{1, 2\}$ be the geodesic arc realizing the Delaunay edge between v_i and u_i and connecting these two points in \tilde{X} . In particular, $\gamma_i \subset B_i$.

Assume towards a contradiction that the two chords γ_1 and γ_2 intersect non-trivially along their interior. Recall that a pair of geodesics in a CAT(−1)-space admit at most a single intersection point. An examination of the resulting planar diagram shows that $|\partial B_1 \cap \partial B_2| > 2$. Therefore Lemma A.1 applied with respect to any three distinct points of the intersection $\partial B_1 \cap \partial B_2$ shows that the balls B_1 and B_2 must coincide. This is a contradiction.

The fact that the projection $q(D)$ is embedded in X follows immediately from the above discussion combined with the $\pi_1(X)$ -invariance of the Delaunay graph. ■

Proposition 3.7 shows that the connected components of the complement of the Delaunay graph have locally convex boundary.

Proof of Proposition 3.7. Assume towards a contradiction that the boundary of some connected component of $\tilde{X} \setminus D$ is not locally convex. Equivalently, there is some Delaunay vertex $v \in s(\tilde{X})$ and some half-space \mathfrak{H} with $v \in \partial \mathfrak{H}$ such that no Delaunay edge incident at v is contained in \mathfrak{H} . The proof of Lemma 3.8 and in particular the Claim contained in that proof show that the Voronoi cell A_v has infinite area.² This contradicts the fact that A_v is compact. ■

Acknowledgements. The authors would like to thank Goulnara Arzhantseva, Nir Avni, Oren Becker, Tsachik Gelander and Alex Lubotzky for sharing their inspiring ideas and for useful discussions. The authors would like to thank the anonymous referee for their useful corrections and suggestions.

Funding. The third named author was partially supported by NSF grant DMS-1610827.

References

- [1] Arzhantseva, G., Păunescu, L.: [Almost commuting permutations are near commuting permutations](#). *J. Funct. Anal.* **269**, 745–757 (2015) Zbl [1368.20025](#) MR [3350728](#)
- [2] Beardon, A. F.: [The geometry of discrete groups](#). Graduate Texts in Mathematics 91, Springer, New York (1983) Zbl [0528.30001](#) MR [698777](#)

²While Lemma 3.8 appears below Proposition 3.7 in the text, the proof of that lemma is independent of this proposition.

- [3] Becker, O., Lubotzky, A.: [Group stability and Property \(T\)](#). J. Funct. Anal. **278**, art. 108298, 20 pp. (2020) Zbl [1516.20100](#) MR [4027744](#)
- [4] Becker, O., Lubotzky, A., Thom, A.: [Stability and invariant random subgroups](#). Duke Math. J. **168**, 2207–2234 (2019) Zbl [1516.20007](#) MR [3999445](#)
- [5] Becker, O., Mosheiff, J.: [Abelian groups are polynomially stable](#). Int. Math. Res. Notices **2021**, 15574–15632 Zbl [07456010](#) MR [4329877](#)
- [6] Bowen, L., Burton, P.: [Flexible stability and nonsficity](#). Trans. Amer. Math. Soc. **373**, 4469–4481 (2020) Zbl [1476.20046](#) MR [4105530](#)
- [7] Bridson, M. R., Haefliger, A.: [Metric spaces of non-positive curvature](#). Grundlehren der mathematischen Wissenschaften 319, Springer, Berlin (1999) Zbl [0988.53001](#) MR [1744486](#)
- [8] Elek, G., Szabó, E.: [On sofic groups](#). J. Group Theory **9**, 161–171 (2006) Zbl [1153.20040](#) MR [2220572](#)
- [9] Glebsky, L., Rivera, L. M.: [Almost solutions of equations in permutations](#). Taiwanese J. Math. **13**, 493–500 (2009) Zbl [1203.20003](#) MR [2500002](#)
- [10] Hatcher, A.: [Algebraic topology](#). Cambridge University Press, Cambridge (2002) Zbl [1044.55001](#) MR [1867354](#)
- [11] Lyndon, R. C., Schupp, P. E.: [Combinatorial group theory](#). Classics in Mathematics, Springer (2001) Zbl [0997.20037](#) MR [1812024](#)
- [12] Neumann, W. D.: [Immersed and virtually embedded \$\pi_1\$ -injective surfaces in graph manifolds](#). Algebr. Geom. Topol. **1**, 411–426 (2001) Zbl [0979.57007](#) MR [1852764](#)
- [13] Patel, P.: [On a theorem of Peter Scott](#). Proc. Amer. Math. Soc. **142**, 2891–2906 (2014) Zbl [1311.57003](#) MR [3209342](#)
- [14] Scott, P.: [Subgroups of surface groups are almost geometric](#). J. London Math. Soc. (2) **17**, 555–565 (1978) Zbl [0412.57006](#) MR [494062](#)
- [15] Scott, P.: [Correction to: “Subgroups of surface groups are almost geometric”](#). J. London Math. Soc. (2) **32**, 217–220 (1985) Zbl [0581.57005](#) MR [811778](#)
- [16] Weiss, B.: [Sofic groups and dynamical systems](#). Sankhyā Ser. A **62**, 350–359 (2000) Zbl [1148.37302](#) MR [1803462](#)