Algebraic & Geometric Topology 17 (2017) 1041–1057

Infima of length functions and dual cube complexes

JONAH GASTER

In the presence of certain topological conditions, we provide lower bounds for the infimum of the length function associated to a collection of curves on Teichmüller space that depend on the dual cube complex associated to the collection, a concept due to Sageev. As an application of our bounds, we obtain estimates for the "longest" curve with k self-intersections, complementing work of Basmajian [J. Topol. 6 (2013) 513–524].

51M10; 51M16

Let Σ be an oriented topological surface of finite type. We denote the Teichmüller space of Σ by $\mathcal{T}(\Sigma)$, which we interpret as the deformation space of marked hyperbolic structures on Σ. Given $X \in \mathcal{T}(\Sigma)$ and a free homotopy class (or *closed curve*) γ on Σ, we denote by $\ell(\gamma, X)$ the length of the geodesic representative of γ in the hyperbolic structure determined by X. If $\Gamma = \{\gamma_i\}$ is a collection of closed curves, then we define $\ell(\Gamma, X) = \sum \ell(\gamma_i, X)$.

In this note, we are concerned with translating topological information of Γ into quantitative information about the length function $\ell(\Gamma, \cdot)\colon \mathcal{T}(\Sigma) \to \mathbb{R}$. In particular, we develop tools to estimate the infimum of $\ell(\Gamma, \cdot)$ over $\mathcal{T}(\Sigma)$. This work naturally complements that of Basmajian [3], where such estimates are obtained that depend on the number of self-intersections of Γ. Here we consider a finer topological invariant than the self-intersection number.

A construction of Sageev [7] associates to a curve system Γ an isometric action of $\pi_1\Sigma$ on a finite-dimensional cube complex, or the *dual cube complex* $\mathcal{C}(\Gamma)$ of Γ. In what follows we connect geometric properties of the dual cube complex $\mathcal{C}(\Gamma)$ to the length of the collection of curves Γ on any hyperbolic surface. Indeed, [1, Theorem 3] of Aougab and Gaster suggests that any such information is implicitly contained in the combinatorics of $\mathcal{C}(\Gamma)$. We have:

Theorem A *Suppose that the action of $\pi_1\Sigma$ on $\mathcal{C}(\Gamma)$ has a set of cubes $C_1, \dots, C_m$, of dimensions $n_1, \dots, n_m$, respectively, in distinct $\pi_1\Sigma$–orbits, such that the union of orbits $\bigcup_i \pi_1\Sigma \cdot C_i$ is hyperplane separated. Then*

$$\inf_{X \in \mathcal{T}(\Sigma)} \ell(\Gamma, X) \geq \sum_{i=1}^{m} n_i \log\left(\frac{1+\cos\pi/n_i}{1-\cos\pi/n_i}\right).$$

Published: 14 March 2017 DOI: 10.2140/agt.2017.17.1041

The definition of *hyperplane separated* can be found in Section 1. The main use of this idea is that it allows one to conclude that large chunks of the preimage of the curves Γ in the universal cover embed on the surface under the covering map. The proof of Theorem A proceeds by minimizing the length of these large chunks.

Remark The contribution to the bound above from a cube C_i is useless when C_i is a 2–cube. On the other hand, it still seems reasonable to expect a lower bound on the length function in the presence of many maximal 2–cubes. Note that the presence of m 2–cubes contributes m to the self-intersection number, so that Basmajian's bounds immediately imply a lower bound for the length function that is logarithmic in m. While Basmajian's lower bounds are sharp, the examples that demonstrate sharpness have high-dimensional dual complexes. We expect a positive answer to the following.

Question If $\mathcal{C}(\Gamma)$ contains m maximal 2–cubes, is there a lower bound for the length function $\ell(\Gamma, X)$ that is linear in m?

Remark The bounds in Theorem A are sharp in the following sense: for each $n \in \mathbb{N}$, there exists a set of curves Γ_n on the $(n+1)$–holed sphere $\Sigma_{0,n+1}$, and hyperbolic structures $X_n \in \mathcal{T}(\Sigma_{0,n+1})$ such that

(1) the dual cube complex $\mathcal{C}(\Gamma_n)$ has a hyperplane separated n–cube, and
(2) the hyperbolic length $\ell(\Gamma_n, X_n)$ is asymptotic to $n \log n$.

The problem remains of determining when a collection of curves gives rise to hyperplane separated orbits of cubes in the action of $\pi_1\Sigma$ on the dual cube complex. We offer a sufficient condition below which applies in many cases, toward which we fix some terminology. Recall that a *ribbon graph* is a graph with a cyclic order given to the oriented edges incident to each vertex. A ribbon graph G is *even* if the valence of each vertex is even. When an even ribbon graph G is embedded on a surface Σ, a collection of homotopy classes of curves is determined by G by going straight at each vertex. See Section 6 for a more precise description.

Theorem B *Suppose that $G \hookrightarrow \Sigma$ is an embedding of an even ribbon graph G into Σ with vertices of valence $n_1, \dots, n_m$, such that the complement $\Sigma \setminus G$ contains no monogons, bigons, or triangles. Let Γ indicate the union of the closed curves determined by G. Then G is a minimal position realization of Γ, the self-intersection of Γ is given by $\binom{n_1}{2} + \cdots + \binom{n_m}{2}$, and $\mathcal{C}(\Gamma)$ contains cubes $C_1, \dots, C_m$ of dimensions $n_1, \dots, n_m$, respectively, in distinct $\pi_1\Sigma$–orbits, whose union is hyperplane separated.*

This provides a general method to construct curves with definite self-intersection number and definite hyperplane separated cubes in their dual cube complexes. For example:

Example Consider the curve in Figure 4. Theorem B implies that the curve has six hyperplane separated 3–cubes. The estimate in Theorem A now applies, so that the length of the pictured curve is at least $18\log 3$ in any hyperbolic metric on Σ_6.

Let $\mathfrak{C}_k(\Sigma)$ indicate curves on Σ with self-intersection number k. Basmajian examined the following quantities, showing that they both are asymptotic to $\log k$:

$$\begin{aligned} m_k(\Sigma) &:= \min_{\gamma\in\mathfrak{C}_k(\Sigma)} \inf\{\ell(\gamma, X) \mid X \in \mathcal{T}(\Sigma)\}, \\ M_k &:= \inf\{m_k(\Sigma) \mid \Sigma \text{ is a finite-type surface with } \chi(\Sigma) < 0\}. \end{aligned}$$

Note that, for each k and Σ, there are finitely many mapping class group orbits among $\mathfrak{C}_k(\Sigma)$. This justifies the use of minimum in the definition of $m_k(\Sigma)$ above. One may define analogously

$$\begin{aligned} \overline{m}_k(\Sigma) &:= \max_{\gamma\in\mathfrak{C}_k(\Sigma)} \inf\{\ell(\gamma, X) \mid X \in \mathcal{T}(\Sigma)\}, \\ \overline{M}_k &:= \sup\{\overline{m}_k(\Sigma) \mid \Sigma \text{ is a finite-type surface with } \chi(\Sigma) < 0\}. \end{aligned}$$

The curves that realize the minima $m_k(\Sigma)$ and M_k manage to gain a lot of self-intersection while remaining quite short, which they achieve by winding many times around a very short curve. By constructing explicit families of curves that behave quite differently — namely, they return many times to a fixed small compact set on the surface — we provide a lower bound for $\overline{M}_k$ that grows faster than Basmajian's bounds for the "shortest" curves with k self-intersections.

Theorem C *We have the estimate*

$$\limsup_{k\to\infty} \frac{\overline{M}_k}{k} \geq \frac{\log 3}{3}.$$

Remark It is not hard to observe that

$$\limsup_{k\to\infty} \frac{\overline{m}_k(\Sigma)}{\sqrt{k}} > 0.$$

Indeed, given any k–curve $\gamma \in \mathfrak{C}_k(\Sigma)$, consider the closed curve γ^n given by wrapping n times around γ. The infimum of the length function of γ^n will grow linearly in n, while the self-intersection number will grow quadratically in n. Performing this calculation with a curve with one self-intersection, one finds that

$$\limsup_{k\to\infty} \frac{\overline{m}_k(\Sigma)}{\sqrt{k}} \geq 4\log(1+\sqrt{2}).$$

The problem of sharpness for our examples, namely good upper bounds for $\overline{m}_k(\Sigma)$ and $\overline{M}_k$, seems subtle. In particular, such upper bounds would imply an asymptotically good answer to the following question.

Question Given a curve $\gamma \in \mathfrak{C}_k(\Sigma)$, what is an explicit function $C(k,\Sigma)$ such that there is a point $X \in \mathcal{T}(\Sigma)$ with $\ell(\gamma, X) \le C(k,\Sigma)$?[1]

Note that one could also ask for an upper bound that is independent of Σ, towards which the lower bound in Theorem C is more relevant.

Organization In Section 1 we briefly recall Sageev's construction, and define hyperplane separation. In Section 2 we lay out the necessary tools for the proof of Theorem A, and in Section 3 and Section 4 we prove these tools. Section 5 describes a straightforward method of detecting self-intersection and hyperplane separation, and Section 6 introduces even ribbon graphs and the proof of Theorem B. Finally, Section 7 describes a family of examples to which these tools apply, and contains a proof of Theorem C.

Acknowledgements The author thanks Ara Basmajian, Martin Bridgeman, Spencer Dowdall, and David Dumas for helpful conversations, and Ian Biringer for a correction and reference in regard to Lemma 5.1.

1 Dual cube complexes and hyperplane separation

We recall Sageev's construction. A collection Γ of homotopy classes of curves on Σ gives rise to an isometric action of $\pi_1\Sigma$ on a CAT(0)-cube complex $\mathcal{C}(\Gamma)$. This action is obtained roughly as follows: choose a minimal position realization λ of the curves in Γ, and consider the preimage $\widetilde{\lambda}$ of λ in the universal cover $\widetilde{\Sigma}$. In the language of [9], the set $\widetilde{\lambda}$ decomposes into a union[2] of *elevations*. Each elevation splits $\widetilde{\Sigma}$ into two connected components. A *labeling* of $\widetilde{\lambda}$ is a choice of half-space in the complement of each of the elevations. The one skeleton of the cube complex $\mathcal{C}(\Gamma)$ is built from *admissible* labelings of $\widetilde{\lambda}$, ie choices of half-spaces in the complement of the elevations so that any pair intersect.

[1] While this paper was under review, this question has been given an answer by Aougab et al [2, Theorem 1.4].

[2] An illustrative example is provided by the case that λ is given by the geodesic representatives of Γ relative to a chosen hyperbolic structure, in which case $\widetilde{\lambda} \subset \mathbb{H}^2$ is evidently a union of complete geodesics.

Two such admissible labelings are connected by an edge when they differ on precisely one elevation of $\widetilde{\lambda}$ (in analogy with the dual graph to the set $\widetilde{\lambda}$). Finally, $\mathcal{C}(\Gamma)$ is given by the unique nonpositively curved cube complex with the prescribed 1–skeleton. The action of $\pi_1\Sigma$ on the elevations comprising $\widetilde{\lambda}$ naturally induces a permutation of the labelings, which induces an isometry of $\mathcal{C}(\Gamma)$. See [7; 8; 4] for details.

We collect this information conveniently.

Theorem (Sageev) *The action of $\pi_1\Sigma$ on the* CAT(0) *cube complex $\mathcal{C}(\Gamma)$ is independent of realization. There is a $\pi_1\Sigma$–equivariant incidence-preserving correspondence of the hyperplanes of $\mathcal{C}(\Gamma)$ with the elevations in $\widetilde{\lambda}$, so that maximal n–cubes are in correspondence with maximal collections of n pairwise intersecting elevations of curves in Γ.*

In light of Sageev's theorem we may sometimes identify the elevations in $\widetilde{\lambda}$ with the hyperplanes of the cube complex $\mathcal{C}(\Gamma)$.

Given a cube C in a cube complex $\mathcal{C}$, we denote the set of hyperplanes of $\mathcal{C}$ by $\mathcal{H}(\mathcal{C})$, and the set of hyperplanes of C by $\mathcal{H}(C) \subset \mathcal{H}(\mathcal{C})$.

Definition Suppose C and D are two cubes in a cube complex. We say that C and D are *hyperplane separated* if either $\mathcal{H}(C) \cap \mathcal{H}(D) = \varnothing$, or $|\mathcal{H}(C) \cap \mathcal{H}(D)| = 1$, and, for any $c \in \mathcal{H}(C)$ and $d \in \mathcal{H}(D)$ with $c \neq d$, the hyperplanes c and d are disjoint.

A union of cubes $\bigcup_i C_i$ is hyperplane separated when every pair is hyperplane separated.

2 Proof of Theorem A

Consider an n–cube $C \subset \mathcal{C}(\Gamma)$. The orientation of $\widetilde{\Sigma}$ induces a counterclockwise cyclic ordering of the n elevations of curves from Γ that correspond to the hyperplanes of C. In what follows, we fix a hyperbolic surface $X \in \mathcal{T}(\Sigma)$ and identify the universal cover $\widetilde{\Sigma}$ with $\mathbb{H}^2$. We will work with the Poincaré disk model for $\mathbb{H}^2$, with conformal boundary S^1.

Enumerate the n geodesic representatives $(\gamma_1, \dots, \gamma_n)$ of the elevations of curves that correspond to C, respecting the cyclic order. Each geodesic γ_i has two endpoints $p_i, q_i \in S^1$. Choose these labels so that $p_1, \dots, p_n, q_1, \dots, q_n$ is consistent with the cyclic order of S^1.

Given a trio of elevations $\gamma_{i-1}, \gamma_i, \gamma_{i+1} \subset \mathbb{H}^2$, consider the pair of distinct disjoint geodesics (p_{i-1}, p_{i+1}) and (q_{i-1}, q_{i+1}). We will refer to this pair as the *separators*

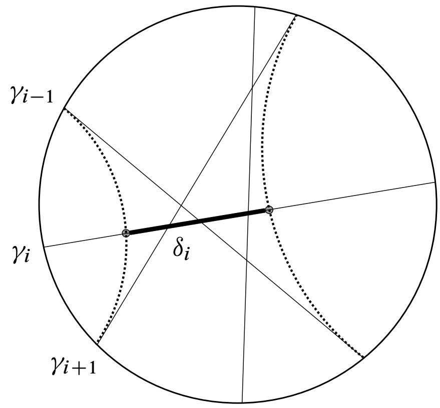

Figure 1: The separators of the hyperplane γ_i and the diagonal δ_i

of the hyperplane γ_i in C, and we will denote the pair by $\operatorname{sep}(\gamma_i, C)$. (When $i = 1$ or $i = n$, the separators of γ_i are the geodesics (p_{i-1}, q_{i+1}) and (q_{i-1}, p_{i+1}), with indices read modulo n.) Let δ_i indicate the portion of γ_i between the separators. We will refer to the arcs $\{\delta_1, \dots, \delta_n\}$ as the *diagonals* of the cube C. See Figure 1 for a schematic picture.

Lemma 2.1 *For each i, we have*

$$\ell(\delta_i) = \log\left|\frac{(p_i - q_{i-1})(p_i - q_{i+1})(q_i - p_{i+1})(q_i - p_{i-1})}{(p_i - p_{i+1})(p_i - p_{i-1})(q_i - q_{i+1})(q_i - q_{i-1})}\right|.$$

Proof The proof is a calculation in $\mathbb{H}^2$. □

Towards Theorem A, we suppose below that $C_1, \dots, C_m$ are cubes of $\mathcal{C}(\Gamma)$ in distinct $\pi_1\Sigma$–orbits. Let $\delta^i_1, \dots, \delta^i_{n_i}$ be the diagonals of the cube C_i, let

$$\mathcal{D}_i := \bigcup_k \delta^i_k$$

indicate the union of the diagonals of C_i, and $\mathcal{D} := \mathcal{D}_1 \cup \cdots \cup \mathcal{D}_m$.

For ease of exposition, we postpone the proof of the following proposition.

Proposition 2.2 *If the union of orbits $\bigcup \pi_1\Sigma \cdot C_i$ is hyperplane separated, then the covering map $\pi\colon \mathbb{H}^2 \to \Sigma$ is injective on the union $\mathcal{D}$ minus a finite set of points.*

Finally, we will need the solution to the following optimization problem, whose proof we also postpone: given $2n$ distinct points $x_1, \dots, x_{2n} \in S^1$, for notational convenience we adopt the natural convention that subscripts should be read modulo $2n$, so that $x_{2n+1} = x_1$ and $x_0 = x_{2n}$. Let $F(x_1, \dots, x_{2n})$ be defined by

$$F(x_1, \dots, x_{2n}) = \log \prod_{j=1}^{2n} \left| \frac{(x_j - x_{j+n+1})(x_j - x_{j+n-1})}{(x_j - x_{j+1})(x_j - x_{j-1})} \right|.$$

Lemma 2.3 *When $(x_1, \dots, x_{2n})$ are cyclically ordered in S^1, we have*

$$F(x_1, \dots, x_{2n}) \geq n \log\left(\frac{1 + \cos \pi/n}{1 - \cos \pi/n} \right).$$

Assuming for now Proposition 2.2 and Lemma 2.3, we are ready to prove Theorem A.

Proof of Theorem A We bound from below the sum of lengths of the curves from Γ in the hyperbolic structure determined by $X \in \mathcal{T}(\Sigma)$. Pull Γ tight to geodesics, and consider the preimage under the covering transformation. As described above, each cube C_i, of dimension n_i, has n_i hyperplanes with n_i corresponding elevations of mutually intersecting geodesics in $\mathbb{H}^2$. These curves determine $2n_i$ cyclically ordered distinct points

$$p_1^i, p_2^i, \dots, p_{n_i}^i, q_1^i, q_2^i, \dots, q_{n_i}^i$$

on S^1, the diagonals $\mathcal{D}_i$, and $\mathcal{D}$, the union of $\mathcal{D}_i$. We estimate

$$\begin{aligned}
\ell(\Gamma, X) &\geq \ell(\mathcal{D}) \\
&= \sum_{i=1}^{m} \ell(\mathcal{D}_i) \\
&= \sum_{i=1}^{m} \sum_{j=1}^{n_i} \ell(\delta_j^i) \\
&= \sum_{i=1}^{m} \log \prod_{j=1}^{n_i} \left| \frac{(p_j^i - q_{j-1}^i)(p_j^i - q_{j+1}^i)(q_j^i - p_{j+1}^i)(q_j^i - p_{j-1}^i)}{(p_j^i - p_{j+1}^i)(p_j^i - p_{j-1}^i)(q_j^i - q_{j+1}^i)(q_j^i - q_{j-1}^i)} \right| \\
&\geq \sum_{i=1}^{m} n_i \log\left(\frac{1 + \cos \pi/n_i}{1 - \cos \pi/n_i} \right),
\end{aligned}$$

where the first, fourth and fifth lines follow from Proposition 2.2, Lemma 2.1 and Lemma 2.3, respectively. □

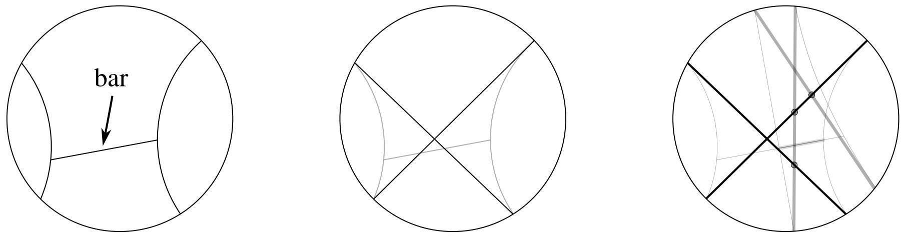

Figure 2: The notion of an H in $\mathbb{H}^2$. From left to right: an H and its bar, the cross of an H, and H's with overlapping bars.

3 Proof of Proposition 2.2

Proposition 2.2 is the sole motivation for the definition of hyperplane separated. We turn to the proof. To aid our exposition, we will say that an H in $\mathbb{H}^2$ is a pair of disjoint geodesics, and a geodesic arc connecting them. Associated to an H is a *cross*, a pair of intersecting geodesics with the same limit points as the H. See Figure 2 left and middle for a schematic. For example, the union of the diagonal δ_i and the separators $\mathrm{sep}(\gamma_i, C)$ form an H, with associated cross $\{\gamma_{i-1}, \gamma_{i+1}\}$.

Lemma 3.1 *Suppose* $\mathrm{H}_1, \mathrm{H}_2 \subset \mathbb{H}^2$ *are distinct* H*'s whose bars overlap in an interval. Then the crosses of* H_1 *and* H_2 *intersect.*

See Figure 2 right for a schematic.

Proof The convex hull of an H is an ideal quadrilateral. By assumption, the convex hulls of H_1 and H_2 intersect. The lemma follows from the following simple observation: if two ideal quadrilaterals intersect, then their crosses intersect. We demonstrate this below. Note that the ideal points of an ideal quadrilateral are cyclically ordered. We say that two such points are *opposite* if they are not neighbors in the cyclic order.

Let P and Q be intersecting ideal quadrilaterals, with cyclically ordered ideal points ∂P and ∂Q in $\partial_\infty \mathbb{H}^2$. As P and Q intersect, there are two points $q, q' \in \partial Q$ lying in distinct components of $\partial_\infty \mathbb{H}^2 \setminus \partial P$. Suppose that q and q' are not opposite vertices. The vertex that follows q' in the cyclic order is either in the same component of $\partial_\infty \mathbb{H}^2 \setminus \partial P$ as q, in which case there are three vertices in the same component as q, or it is in a distinct component from q. Thus if P and Q intersect, there are a pair of opposite vertices of ∂Q in distinct components of $\partial_\infty \mathbb{H}^2 \setminus \partial P$.

Opposite vertices of an ideal quadrilateral are boundary points of the cross of the quadrilateral. Thus there is a geodesic of the cross of Q that runs between distinct components of $\partial_\infty \mathbb{H}^2 \setminus \partial P$, so that the crosses of P and Q intersect. $\square$

Proof of Proposition 2.2 As the union $\mathcal{D}$ is compact, properness of the action of $\pi_1\Sigma$ ensures that there are finitely many elements $g \in \pi_1\Sigma$ such that $g\cdot\mathcal{D}\cap\mathcal{D} \neq \varnothing$. Let $\mathcal{D}'$ indicate the complement in $\mathcal{D}$ of the finitely many points that are transversal intersections of $\mathcal{D}$ with $g\cdot\mathcal{D}$. If π is not injective on $\mathcal{D}'$, then there is an element $1 \neq g \in \pi_1\Sigma$ such that $g\cdot\delta_i^k$ and δ_j^l overlap in an open interval. In particular, note that g sends the hyperplane containing δ_i^k to the hyperplane containing δ_j^l.

If $k = l$ and $g\cdot C_k = C_k$, then $g\cdot\mathcal{D}_k = \mathcal{D}_k$, and by Brouwer's fixed point theorem there would be a fixed point of g, violating freeness of the action of $\pi_1\Sigma$. Since $g\cdot C_k \neq C_l$ for $k \neq l$ (recall that the cubes $\{C_i\}$ are in distinct $\pi_1\Sigma$–orbits), we may thus assume that $g\cdot C_k$ and C_l are distinct cubes that share the common hyperplane containing the diagonals $g\cdot\delta_i^k$ and δ_j^l. Let the hyperplanes of C_k be given by $\{\gamma_1,\dots,\gamma_{n_k}\}$, and those of C_l by $\{\eta_1,\dots,\eta_{n_l}\}$, so that $g\cdot\gamma_i = \eta_j$.

Observe that a trivial consequence of separatedness is that the separators $\mathrm{sep}(\eta_j, g\cdot C_k)$ are not the same pair of geodesics as the separators $\mathrm{sep}(\eta_j, C_l)$: if they were identical, then $g\cdot C_k$ and C_l would be two distinct cubes in the union $\bigcup \pi_1\Sigma\cdot C_i$ that share the hyperplanes corresponding to γ_{j-1}, γ_j, and γ_{j+1}.

Consider then the two H's formed by $g\cdot\delta_i^k$ and $\mathrm{sep}(\eta_j, g\cdot C_k)$ on the one hand, and δ_j^l and $\mathrm{sep}(\eta_j, C_l)$ on the other. By assumption these two H's have overlapping bars, so that by Lemma 3.1 their crosses intersect. Namely, one of $g\cdot\gamma_{i-1}$ and $g\cdot\gamma_{i+1}$ intersects one of η_{j-1} and η_{j+1}. This contradicts separatedness of C. We conclude that π is injective on $\mathcal{D}'$, the union of the diagonals of $C_1,\dots,C_m$ minus finitely many points, as desired. □

4 Proof of Lemma 2.3

We solve the necessary optimization problem.

Proof of Lemma 2.3 Note that F has several useful invariance properties: first it is clear that F is invariant under rotations of S^1. More generally, the conformal automorphisms of the disk $\mathrm{Aut}(\mathbb{D})$ act diagonally on $(S^1)^{2n}$, and for any $\sigma \in \mathrm{Aut}(\mathbb{D})$, $F\circ\sigma = F$. As well, it is immediate from the definition that $F(x_1, x_2, \dots, x_{2n}) = F(x_2, \dots, x_{2n}, x_1)$.

For each $j = 1,\dots,2n$, let $x_j = e^{i\theta_j}$. Applying a rotation of S^1 if necessary, we assume that $0 \le \theta_1 < \cdots < \theta_{2n} < 2\pi$.

The identity $|e^{i\alpha} - e^{i\beta}| = \sqrt{2 - 2\cos(\alpha-\beta)}$ implies that

$$\log\left|\frac{e^{i\theta_j} - e^{i\theta_k}}{e^{i\theta_j} - e^{i\theta_l}}\right| = \tfrac{1}{2}\log\frac{1-\cos(\theta_j-\theta_k)}{1-\cos(\theta_j-\theta_l)}.$$

Taking a derivative we find

$$\frac{\partial F}{\partial \theta_j} = \frac{\sin(\theta_j - \theta_{j+n-1})}{1-\cos(\theta_j - \theta_{j+n-1})} + \frac{\sin(\theta_j - \theta_{j+n+1})}{1-\cos(\theta_j - \theta_{j+n+1})} - \frac{\sin(\theta_j - \theta_{j-1})}{1-\cos(\theta_j - \theta_{j-1})} - \frac{\sin(\theta_j - \theta_{j+1})}{1-\cos(\theta_j - \theta_{j+1})}.$$

Since $\sin\theta/(1-\cos\theta) = \cot\frac{\theta}{2}$, we may write the above as

$$\frac{\partial F}{\partial \theta_j} = \cot\frac{\theta_j - \theta_{j+n-1}}{2} + \cot\frac{\theta_j - \theta_{j+n+1}}{2} - \cot\frac{\theta_j - \theta_{j+1}}{2} - \cot\frac{\theta_j - \theta_{j-1}}{2}.$$

Towards candidates for absolute minima of F, we seek solutions to the system of equations $\{\partial F/\partial\theta_j = 0\}$. Given the invariance properties of F, any such solution is far from unique, even locally. In order to characterize the unique $\mathrm{Aut}(\mathbb{D})$–orbit of a solution, we pick a j, and fix the choices $\theta_{n+j} - \theta_j = \pi$, and $\theta_{n+j+1} - \theta_{j+1} = \pi$.

With $\theta_j + \pi$ substituted for θ_{n+j}, the equations $\partial F/\partial\theta_j = 0$ and $\partial F/\partial\theta_{n+j} = 0$ now yield

$$\begin{aligned}
\cot\frac{\theta_j - \theta_{j+n-1}}{2} + \cot\frac{\theta_j - \theta_{j+n+1}}{2} &= \cot\frac{\theta_j - \theta_{j+1}}{2} + \cot\frac{\theta_j - \theta_{j-1}}{2},\\
\tan\frac{\theta_j - \theta_{j-1}}{2} + \tan\frac{\theta_j - \theta_{j+1}}{2} &= \tan\frac{\theta_j - \theta_{j+n-1}}{2} + \tan\frac{\theta_j - \theta_{j+n+1}}{2},
\end{aligned}$$

respectively. Eliminating $\tan\frac{\theta_j - \theta_{j+1}}{2}$, we find

$$\cot\frac{\theta_j - \theta_{j+n-1}}{2} + \cot\frac{\theta_j - \theta_{j+n+1}}{2} - \cot\frac{\theta_j - \theta_{j-1}}{2} = \left(\tan\frac{\theta_j - \theta_{j+n-1}}{2} + \tan\frac{\theta_j - \theta_{j+n+1}}{2} - \tan\frac{\theta_j - \theta_{j-1}}{2}\right)^{-1}.$$

Recall the remarkable fact that the solutions of the equation

$$\frac{1}{x} + \frac{1}{y} + \frac{1}{z} = \frac{1}{x+y+z}$$

are precisely the equations $x = -y$, $x = -z$, or $y = -z$. As a consequence, we have one of the following:

$$\begin{aligned}
\tan\frac{\theta_j - \theta_{j+n-1}}{2} &= -\tan\frac{\theta_j - \theta_{j+n+1}}{2},\\
\tan\frac{\theta_j - \theta_{j+n-1}}{2} &= \tan\frac{\theta_j - \theta_{j-1}}{2}, \quad \text{or}\\
\tan\frac{\theta_j - \theta_{j+n+1}}{2} &= \tan\frac{\theta_j - \theta_{j-1}}{2}.
\end{aligned}$$

By assumption,

$$0 < \theta_{j+n-1} - \theta_j < \pi < \theta_{j+n+1} - \theta_j < \theta_{j-1} - \theta_j < 2\pi,$$

so that

$$-\pi < \frac{\theta_j - \theta_{j-1}}{2} < \frac{\theta_j - \theta_{j+n+1}}{2} < -\frac{\pi}{2} < \frac{\theta_j - \theta_{j+n-1}}{2} < 0.$$

The only possibility above is thus the first equation, so that

$$\frac{\theta_{n+j+1} - \theta_j}{2} - \pi = \frac{\theta_j - \theta_{j+n-1}}{2},$$

or $\theta_{j+n+1} + \theta_{j+n-1} = 2\theta_j + 2\pi$. The equation $\partial F/\partial\theta_j = 0$ now implies as well that $\theta_{j-1} + \theta_{j+1} = 2\theta_j + 2\pi$.

On the other hand, we have also assumed that $\theta_{j+n+1} - \theta_{j+1} = \pi$, so

$$\theta_{j+n-1} + \theta_{j+1} = 2\theta_j + 2\pi - \theta_{j+n+1} + \theta_{j+1} = 2\theta_j + \pi.$$

This implies that

$$\begin{aligned}\theta_{j-1} - \theta_{j+n-1} &= \theta_{j-1} - (2\theta_j + \pi - \theta_{j+1}) \\ &= \theta_{j-1} - \theta_2 - (2\theta_j + \pi) \\ &= (2\theta_j + 2\pi) - (2\theta_j + \pi) \\ &= \pi.\end{aligned}$$

We now know that if we make the normalizing assumptions $\theta_{j+n} = \theta_j + \pi$ and $\theta_{j+n+1} = \theta_{j+1} + \pi$, then the equations $\{\partial F/\partial\theta_j = 0, \partial F/\partial\theta_{j+n} = 0\}$ ensure $\theta_{j-1} = \theta_{j+n-1} + \pi$. Using all the equations $\{\partial F/\partial\theta_j = 0\}$, it is now evident that $\theta_{n+k} = \theta_k + \pi$, for each $k = 1, \ldots, n$.

We apply this understanding to the equation $\partial F/\partial\theta_j = 0$:

$$\begin{aligned}\cot\frac{\theta_j - \theta_{j-1}}{2} + \cot\frac{\theta_j - \theta_{j+1}}{2} &= \cot\frac{\theta_j - \theta_{j+n-1}}{2} + \cot\frac{\theta_j - \theta_{j+n+1}}{2} \\ &= \cot\frac{\theta_j - \theta_{j-1} - \pi}{2} + \cot\frac{\theta_j - \theta_{j+1} - \pi}{2} \\ &= -\tan\frac{\theta_j - \theta_{j-1}}{2} - \tan\frac{\theta_j - \theta_{j+1}}{2},\end{aligned}$$

so that

$$-\tan\frac{\theta_j - \theta_{j-1}}{2} - \cot\frac{\theta_j - \theta_{j-1}}{2} = \tan\frac{\theta_j - \theta_{j+1}}{2} + \cot\frac{\theta_j - \theta_{j+1}}{2}.$$

Since $\tan x + \cot x = 2/\sin 2x$, we obtain

$$\sin(\theta_j - \theta_{j-1}) = \sin(\theta_{j+1} - \theta_j).$$

If $\theta_j - \theta_{j-1} = \pi - (\theta_{j+1} - \theta_j)$, then $\theta_{j+1} - \theta_{j-1} = \pi$. However, $\theta_{j+n} = \theta_j + \pi$, and the θ_j are distinct. Thus $\theta_j - \theta_{j-1} = \theta_{j+1} - \theta_j$, for each $j = 1, \dots, 2n$. Set $\theta_1 = 0$, and we see that $\left(1, e^{\pi i/n}, e^{2\pi i/n}, \dots, e^{(2n-1)\pi i/n}\right)$ is the unique $\mathrm{Aut}(\mathbb{D})$–orbit for which the partial derivatives simultaneously vanish. As it is evident that $F(x_1, \dots, x_{2n})$ goes to $+\infty$ as points x_j and x_{j+1} collide, the absolute minimum of F must occur at a simultaneous zero of its partial derivatives. Evaluating $F\left(1, e^{\pi i/n}, e^{2\pi i/n}, \dots, e^{(2n-1)\pi i/n}\right)$ achieves the result. □

5 Bigons and triangles

Towards Theorem B, for the computation of the self-intersection number of a self-intersecting closed curve, we require a slight generalization of the "bigon criterion" of [5]. Recall that a representative λ of a collection of closed curves Γ is in *minimal position* if its intersection points are transverse, and the number of intersections of λ, counted with multiplicity, is minimal among representatives of Γ. A *monogon* is a polygon with one side and a *bigon* is a polygon with two sides.

Definition A representative λ of a collection of closed curves $\Gamma \subset \Sigma$ has an *immersed monogon* if there is an immersion of a monogon whose boundary arc is contained in λ, and it has an *immersed bigon* if there is such an immersion of a bigon.

Lemma 5.1 *If the representative λ of a collection of closed curves on Σ is without immersed monogons and without immersed bigons, then it is in minimal position.*

Remark An error in a previous version of this lemma was pointed out by Ian Biringer, as well as a reference to a very similar statement due to Hass and Scott. The corrected statement is above (see [6, Theorems 3.5 and 4.2]). As they note, a nonprimitive curve demonstrates that the converse is false [6, p. 94].

Proof Suppose λ is without immersed bigons or monogons, and has n transverse self-intersections. Let $G_\lambda \subset \Sigma$ indicate the graph determined by λ, choose a spanning tree T_λ for G_λ, a lift of T_λ to the universal cover, and a representative of each of the n intersection points. At each of these representative intersection points, the preimage of λ in $\widetilde{\Sigma}$ consists of a pair of linked curves: if there was only one curve the covering map would produce an immersed monogon for λ on Σ, and if the pair of curves at this intersection point were not linked the covering map would produce an immersed bigon for λ on Σ. The self-intersection number of Γ is equal to the number of $\pi_1\Sigma$-orbits of linked elevations of curves from Γ in the universal cover $\widetilde{\Sigma}$, so we are done. $\square$

In order to recognize the presence of hyperplane separated cubes in the dual cube complex of Γ, we will employ a lemma.

Lemma 5.2 *Suppose that Γ has a minimal position realization $\lambda \subset \Sigma$ such that λ has k points of transverse self-intersection of orders $n_1, \dots, n_k$, listed with multiplicity. Then $\mathcal{C}(\Gamma)$ has cubes of dimensions $n_1, \dots, n_k$, with multiplicity. Moreover, the $\pi_1\Sigma$–orbit of the union of these cubes is hyperplane separated for the action of $\pi_1\Sigma$ on $\mathcal{C}(\Gamma)$ if and only if the complement $\Sigma \setminus \lambda$ has no triangles.*

Proof Consider the preimage $\widetilde{\lambda} := \pi^{-1}\lambda \subset \widetilde{\Sigma}$, and choose lifts $p_1, \dots, p_k$ of the self-intersection points of λ, where p_i has order n_i. By hypothesis, there are n_i linked elevations from $\widetilde{\lambda}$ through p_i, so that there is an n_i–cube in $\mathcal{C}(\Gamma)$. We denote this n_i–cube corresponding to the choice of lift p_i by C_i.

If the complement $\Sigma \setminus \lambda$ had a triangle, then this triangle would lift to $\widetilde{\Sigma}$, so that the curves corresponding to C_i, for some i, would contain two sides of the lifted triangle. As a consequence, there would be a different lift p' of one of the intersection points, so that p' would also abut this triangle. Let C' indicate the maximal cube corresponding to the lift p'. By construction, C' shares a hyperplane with C_i, while there are two other hyperplanes, one of C_i and one of C', that intersect. As C' is in the same $\pi_1\Sigma$–orbit as C_j, for some j, the union of orbits $\bigcup \pi_1\Sigma \cdot C_i$ is not hyperplane separated.

Finally, suppose the union of orbits is not hyperplane separated. Then there is $g \in \pi_1\Sigma$ such that C_i and $g \cdot C_i$ share a hyperplane γ, and have a pair of other intersecting hyperplanes, say γ_1 and γ_2. In this case, there is a triangle $T \subset \widetilde{\Sigma}$ formed by γ, γ_1 and γ_2. While this triangle may not embed under the covering map, it contains an innermost triangle, namely a triangle in the complement of $\widetilde{\Sigma} \setminus \widetilde{\lambda}$. This triangle must embed under the covering map, so $\Sigma \setminus \lambda$ contains a triangle. $\square$

6 Closed curves from ribbon graphs

We now prove Theorem B, thus obtaining explicit constructions of curves to which Theorem A applies. Recall that a *ribbon graph* is a graph with a cyclic order given to the oriented edges incident to each vertex, and a ribbon graph is *even* if the valence of each edge is even. In what follows, we introduce notation for even ribbon graphs and analyze the closed curves that they determine.

Let $S(n)$ indicate the union of the n line segments

$$\{t\exp(\pi i m/n) \mid t \in [-1,1]\},$$

for $m = 1,\dots,n$, and label the endpoints $\exp(\pi i m/n)$ and $-\exp(\pi i m/n)$ by a_m and a'_m, respectively. Fix a permutation of endpoints μ by $\mu(a_m) = a'_m = \mu^{-1}(a_m)$, for $m = 1,\dots,n$. We refer to $S(n)$ as a *star*, and μ as the *switch* map of the star.

Let $\underline{n}$ be the tuple $(n_1,\dots,n_k)$, and consider the union $S(\underline{n}) := S(n_1) \sqcup \cdots \sqcup S(n_k)$. Let σ be a fixed-point-free, order-two permutation (that is, a pairing) of the set

$$\{a_{j,i}, a'_{j,i} \mid 1 \le j \le n_i, 1 \le i \le k\}.$$

Let $\Gamma(\underline{n},\sigma)$ indicate the graph given by

$$\Gamma(\underline{n},\sigma) := S(\underline{n})/\sim,$$

where $a_{j,i} \sim \sigma(a_{j,i})$. The vertices of $\Gamma(\underline{n},\sigma)$ are in bijection with the stars $S(n_j)$, and the orientation of $\mathbb{C}$ at 0 induces a cyclic order to the vertex contained in each $S(n_j)$. These orientations give $\Gamma(\underline{n},\sigma)$ the structure of an even ribbon graph. Moreover, it is clear that every even ribbon graph can be constructed in this way.

Let $\Sigma(\underline{n},\sigma)$ be the surface with boundary associated to the ribbon graph $\Gamma(\underline{n},\sigma)$. We identify $\Gamma(\underline{n},\sigma)$ as smoothly[3] and incompressibly embedded in $\Sigma(\underline{n},\sigma)$, so that the embedding induces isomorphisms on the level of fundamental groups. By a *closed curve* in $\Gamma(\underline{n},\sigma)$, we mean the free homotopy class of the image of a smooth immersion of S^1 into $\Gamma(\underline{n},\sigma)$.

Lemma 6.1 *Closed curves in $\Gamma(\underline{n},\sigma)$ are in correspondence with fixed cycles of $(\mu\sigma)^l$, for $l > 0$. The closed curves in $\Gamma(\underline{n},\sigma)$ are in minimal position in $\Sigma(\underline{n},\sigma)$, and the total intersection number of these closed curves is given by $\binom{n_1}{2} + \binom{n_2}{2} + \cdots + \binom{n_k}{2}$.*

Proof The first statement is evident. The second follows from Lemma 5.1, since the complement $\Sigma(\underline{n},\sigma) \setminus \Gamma(\underline{n},\sigma)$ contains no disks, and hence no immersed bigons or monogons. □

To obtain closed curves on closed surfaces, one may glue together $\Sigma(\underline{n},\sigma)$ and another (possibly disconnected) surface along its boundary. Some of the components that are glued may be disks, so it is possible that the curves from $\Gamma(\underline{n},\sigma)$ are no longer in minimal position. While Section 5 can be used to build an algorithm that can be applied on a case-by-case basis, a more straightforward control on this phenomenon can be obtained in many cases.

[3]Note that the smooth structure of $\Gamma(\underline{n},\sigma)$ in a neighborhood of its vertices is induced by viewing $S(n_j)$ as an immersed submanifold of $\mathbb{C}$.

Lemma 6.2 *Suppose that $\widehat{\Sigma}$ is a surface obtained by a gluing of $\Sigma(\underline{n},\sigma)$, so that there is a natural inclusion $\Sigma(\underline{n},\sigma) \hookrightarrow \widehat{\Sigma}$. If the complement $\widehat{\Sigma} \setminus \Gamma(\underline{n},\sigma)$ contains no monogons, bigons, or triangles, then the closed curve $\Gamma(\underline{n},\sigma) \subset \widehat{\Sigma}$ is in minimal position.*

Proof Suppose that $\Gamma(\underline{n},\sigma) \subset \widehat{\Sigma}$ is not in minimal position, so by Lemma 5.1 it has either an immersed monogon or bigon. Suppose that $\phi\colon B \hookrightarrow \widehat{\Sigma}$ is an example of the latter. By assumption, the bigon is not embedded. Thus $\phi^{-1}(\Gamma(\underline{n},\sigma))$ consists of the two sides of B, together with some connected arcs that connect opposite sides of the bigon B. It is easy to see by induction on the number of such arcs that the complement in B must contain either a bigon or a triangle. This triangle embeds under ϕ, violating the assumption that $\widehat{\Sigma} \setminus \Gamma(\underline{n},\sigma)$ contains no triangles. The case of an immersed monogon is straightforwardly similar. □

Lemmas 6.1, 6.2 and 5.2 imply Theorem B directly.

7 A family of examples and Theorem C

Towards Theorem C, for each k let $\underline{\tau}_k$ indicate the sequence $(3,\dots,3)$ with k terms. The vertices of $S(\underline{\tau}_k)$ are given by

$$\{a_{1,j}, a_{2,j}, a_{3,j}, a'_{1,j}, a'_{2,j}, a'_{3,j} \mid 1 \le j \le k\}.$$

Let σ indicate the following pairing:

$$\begin{aligned} a_{2,j} &\leftrightarrow a'_{1,j}, \\ a_{3,j} &\leftrightarrow a'_{2,j} && \text{for } j = 1,\dots,k, \\ a_{1,j} &\leftrightarrow a'_{3,j+1} && \text{for } j = 1,\dots,k-1, \\ a_{1,k} &\leftrightarrow a'_{3,1}. \end{aligned}$$

See Figure 3 for a schematic picture of $\Gamma(\underline{\tau}_k,\sigma)$, and Figure 4 for a gluing of $\Sigma(\underline{\tau}_6,\sigma)$ to which Lemma 5.2 and Lemma 6.2 apply.

Proof of Theorem C Let Σ indicate the surface $\Sigma(\underline{\tau}_k,\sigma)$, so that Σ contains an embedded minimal position copy of the curve $\Gamma(\underline{\tau}_k,\sigma)$, with the self-intersection number $3k$ by Lemma 6.1.

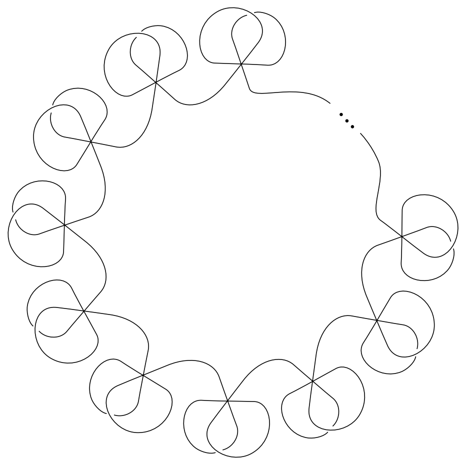

Figure 3: The ribbon graph $\Gamma(\underline{\tau}_k, \sigma)$

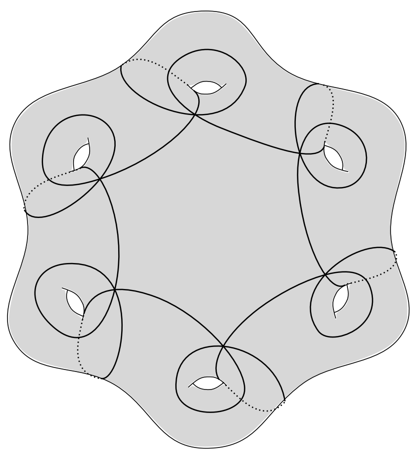

Figure 4: A gluing of $\Sigma(\underline{\tau}_6, \sigma)$ without triangles or bigons in the complement of $\Gamma(\underline{\tau}_6, \sigma)$, such that the given closed curve has six hyperplane separated 3–cubes

By Lemma 5.2, the dual cube complex of $\Gamma(\underline{\tau}_k, \sigma)$ in Σ contains k hyperplane separated 3–cubes. Using Theorem A, we may estimate

$$\begin{aligned}\limsup_{k\to\infty} \frac{\overline{M}_k}{k} &\geq \limsup_{k\to\infty} \frac{\overline{M}_{3k}}{3k} \\ &\geq \limsup_{k\to\infty} \frac{1}{3k} \inf\{\ell(\Gamma(\underline{\tau}_k, \sigma), X) \mid X \in \mathcal{T}(\Sigma)\} \\ &\geq \limsup_{k\to\infty} \frac{1}{3k} k \log\Big(\frac{1+\cos\pi/3}{1-\cos\pi/3}\Big) = \tfrac{1}{3}\log 3. \qquad \square\end{aligned}$$

Note that in the construction above, it is evident that the genus of $\Sigma(\underline{\tau}_k, \sigma)$ will grow with the self-intersection number of $\Gamma(\underline{\tau}_k, \sigma)$. As a consequence, these lower bounds are not applicable to $\overline{m}_k(\Sigma)$ for a fixed surface Σ. It seems likely[4] that $\overline{m}_k(\Sigma)$ grows as $\sqrt{k}$.

References

[1] **T Aougab**, **J Gaster**, *Curves intersecting exactly once and their dual cube complexes* (2016) arXiv To appear in Groups Geom. Dyn.

[2] **T Aougab**, **J Gaster**, **P Patel**, **J Sapir**, *Building hyperbolic metrics suited to closed curves and applications to lifting simply*, preprint (2016) arXiv

[3] **A Basmajian**, *Universal length bounds for non-simple closed geodesics on hyperbolic surfaces*, J. Topol. 6 (2013) 513–524 MR

[4] **I Chatterji**, **G Niblo**, *From wall spaces to* CAT(0) *cube complexes*, Internat. J. Algebra Comput. 15 (2005) 875–885 MR

[5] **B Farb**, **D Margalit**, *A primer on mapping class groups*, Princeton Mathematical Series 49, Princeton University Press (2012) MR

[6] **J Hass**, **P Scott**, *Intersections of curves on surfaces*, Israel J. Math. 51 (1985) 90–120 MR

[7] **M Sageev**, *Ends of group pairs and non-positively curved cube complexes*, Proc. London Math. Soc. 71 (1995) 585–617 MR

[8] **M Sageev**, CAT(0) *cube complexes and groups*, from "Geometric group theory" (M Bestvina, M Sageev, K Vogtmann, editors), IAS/Park City Math. Ser. 21, Amer. Math. Soc., Providence, RI (2014) 7–54 MR

[9] **D T Wise**, *Subgroup separability of graphs of free groups with cyclic edge groups*, Q. J. Math. 51 (2000) 107–129 MR

Department of Mathematics, Boston College
140 Commonwealth Avenue, Chestnut Hill, MA 02467, United States

gaster@bc.edu

http://www2.bc.edu/jonah-gaster

Received: 11 January 2016 Revised: 21 June 2016

[4] While this paper was under review, this has been shown in [2, Theorem 1.4].