

COVERS OF SURFACES, KLEINIAN GROUPS, AND THE CURVE COMPLEX

TARIK AOUGAB, PRIYAM PATEL, SAMUEL J. TAYLOR

ABSTRACT. We prove an effective version of a theorem relating curve complex distance to electric distance in hyperbolic 3-manifolds, up to errors that are polynomial in the complexity of the underlying surface. We use this to give an effective proof of a result regarding maps between curve complexes of surfaces induced by finite covers. As applications, we effectively relate the electric circumference of a fibered manifold to the curve complex translation length of its monodromy, and we give quantitative bounds on virtual specialness for cube complexes dual to curves on surfaces.

1. INTRODUCTION

Let S be an orientable surface of finite type with negative Euler characteristic. The *curve graph* $\mathcal{C}(S)$ of S is the graph whose vertices are homotopy classes of essential simple closed curves and whose edges correspond to pairs of such homotopy classes that admit disjoint representatives. A finite-sheeted cover $p: \tilde{S} \rightarrow S$ induces a (coarsely well-defined) map $p^*: \mathcal{C}(S) \rightarrow \mathcal{C}(\tilde{S})$ sending a vertex γ of $\mathcal{C}(S)$ to its full preimage $p^{-1}(\gamma) \subset \tilde{S}$. In [RS09], Rafi–Schleimer show that the map p^* is a C -quasi-isometric embedding, with C depending only on $\deg(p)$, the degree of p , and on $\chi(S)$. Their result roughly implies that “pairs of simple closed curves do not detangle very much under pull-back by finite covers of small degree,” leading us to pose the following question:

Question 1. Given simple closed curves α and β on S , what is the minimal degree of a cover \tilde{S} of S such that there exist *disjoint* components $\tilde{\alpha}, \tilde{\beta}$ of the pre-images of α, β , respectively?

Unfortunately, this question cannot be answered using [RS09] because their techniques do not pin down how the constant C depends on $\deg(p)$ and $\chi(S)$. Therefore, our approach to the question is to prove the following theorem:

Theorem (7.1). *Let $p: \tilde{S} \rightarrow S$ be a finite covering map between non-sporadic surfaces \tilde{S}, S . Then for any α, β distinct essential simple closed curves on S ,*

$$\frac{d_{\mathcal{C}(S)}(\alpha, \beta)}{\deg(p) \cdot A_3(|\chi(S)|)} \leq d_{\mathcal{C}(\tilde{S})}(p^*(\alpha), p^*(\beta)) \leq d_{\mathcal{C}(S)}(\alpha, \beta),$$

where A_3 is the polynomial $A_3(x) = 80 e^{54} \pi x^{13}$ when S is closed.

When S has punctures,

$$\frac{d_{C(S)}(\alpha, \beta)}{\deg(p)^4 \cdot A_3(|\chi(S)|)} \leq d_{C(\tilde{S})}(p^*(\alpha), p^*(\beta)) \leq d_{C(S)}(\alpha, \beta),$$

where A_3 is the polynomial $A_3(x) = 416 e^{114} \pi^{37} x^{37}$.

Moreover, the linear dependence on $\deg(p)$ in the first bound of Theorem 7.1 is sharp—see Remark 7.2 for more details.

The polynomial A_3 is a product $A_1 \cdot A_2$ of polynomials

$$(1) \quad \begin{aligned} A_1(x) &= \begin{cases} 20 e^{44} x^{10} & \text{for } S \text{ closed,} \\ 104 e^{94} \pi^{30} x^{30} & \text{otherwise} \end{cases} \\ &\text{and} \\ A_2(x) &= \begin{cases} 4 e^{10} \pi x^3 & \text{for } S \text{ closed,} \\ 4 e^{20} \pi^7 x^7 & \text{otherwise,} \end{cases} \end{aligned}$$

which arise independently of one another (e.g., see Theorem 4.1 below) and we will refer to them often in what follows.

The main ingredient in proving Theorem 7.1 is the following theorem regarding the relationship between curve graph distance and electric distance in hyperbolic 3-manifolds. Throughout the paper, we use the same notation for a simple closed curve, its corresponding vertex of the curve graph, and its geodesic representative in a 3-manifold whenever it is clear through context which of the three we are referring to. In Section 3 we define a constant ϵ_S which is bounded from below by a polynomial of degree 2 in $\frac{1}{|\chi(S)|}$ when S is closed and degree 6 in $\frac{1}{|\chi(S)|}$ in general. This is used in the following theorem.

Theorem (4.1). *Let α and β be curves in S and let M be a complete hyperbolic structure on $S \times \mathbb{R}$ such that $\ell_M(\alpha), \ell_M(\beta) \leq \epsilon_S$. Then*

$$\frac{1}{A_1(|\chi(S)|)} \cdot d_{C(S)}(\alpha, \beta) \leq d_M^{\epsilon_S}(\alpha, \beta) \leq A_2(|\chi(S)|) \cdot d_{C(S)}(\alpha, \beta),$$

where the polynomials A_1 and A_2 are as in Equation (1), and $d_M^{\epsilon_S}$ is the metric obtained from the hyperbolic metric d_M by electrifying the ϵ_S -thin part of M .

In [Tan12], Tang used the original, non-effective (i.e. where the dependence on $|\chi(S)|$ was not explicit) version of Theorem 4.1 to reprove the Rafi-Schleimer result, and we follow his argument to obtain Theorem 7.1 from Theorem 4.1.

The non-effective version of Theorem 4.1 is originally due to Bowditch [Bow11, Theorem 5.4]. (See also the statement of Theorem 4.1 in Biringer–Souto [BS15].) As Biringer–Souto state, in reference to Theorem 4.1, “credit should also be given to Yair Minsky, since [the theorem] is implicit in the development of the model manifolds of [Min10], and to Brock–Bromberg [BB11], who prove a closely related result.” However, it is important to note that all of the proofs of these results rely on compactness arguments, which cannot be made effective in a way that is necessary for our applications. Thus, the main contribution of Theorem 4.1 is that it gives the explicit

relationship between curve complex and electric distance. Instead of relying on compactness arguments using geometric limits, we argue using *1-Lipschitz sweepouts in M* (see Theorem 2.1 in Section 2).

Indeed, although there are by now many results relating geometric invariants of hyperbolic manifolds to combinatorial invariants of curves on surfaces, almost none of these can quantify or estimate the exact dependence on the complexity of the underlying surface. Moreover, even when such dependence has been estimated, it is usually (at least) exponential in $|\chi(S)|$. For example, Brock's theorems relating volumes of hyperbolic manifolds to distances in the pants graph [Bro03a, Bro03b] are prime examples of important results relating geometry to combinatorics where dependence on the surface was left completely undetermined. However, in forthcoming work of the first and third author with Webb [ATW18], this dependence is bounded using a careful analysis of Masur–Minsky hierarchies [MM00], but the bound produced is on the order of $|\chi(S)|^{|\chi(S)|}$. Hence, one major novelty of Theorem 4.1 is that our error terms are explicit and depend polynomially on $|\chi(S)|$. To our knowledge, the only other such results are due to Futer–Schleimer [FS14] who estimate the cusp area of a hyperbolic manifold in terms of translation length in the arc complex.

Using Theorem 7.1 we address Question 1 by giving a lower bound on $\deg(\alpha, \beta)$, the minimal degree of a cover necessary to have disjoint components $\tilde{\alpha}, \tilde{\beta}$ of the preimages of α, β , respectively. We emphasize again that this application requires the effective statement of Theorem 7.1 proven here.

Corollary 1.1. *For two simple closed curves α and β on a surface S ,*

$$\frac{d_{C(S)}(\alpha, \beta)}{C(|\chi(S)|)} \leq \deg(\alpha, \beta)^a,$$

where $C(x) = 3A_1(x)A_2(x)$ is a polynomial of degree 13 and $a = 1$ when S is closed. In general, $C(x)$ is a polynomial of degree 37 and $a = 4$.

In Section 8, we provide an application of this corollary to certain virtually special cube complexes. Given a sufficiently complicated collection Γ of curves on a closed surface S , Sageev's construction [Sag95] gives rise to a dual CAT(0) cube complex on which $\pi_1(S)$ acts freely and properly discontinuously. The quotient of the complex under this action is a non-positively curved cube complex \mathfrak{C}_Γ . It is well known by the work of Haglund–Wise [HW08] that \mathfrak{C}_Γ is *virtually special*, meaning that it has a finite degree cover whose fundamental group embeds nicely into a right-angled Artin group. The following theorem quantifies this statement by estimating the degree of the required cover in the case that Γ is a pair of curves:

Theorem (8.3). *Suppose that α and β are two simple closed curves that together fill a closed surface S . Let $\deg \mathfrak{C}_\Gamma$ be the minimal degree of a special cover of the dual cube complex \mathfrak{C}_Γ to the curve system $\Gamma = \alpha \cup \beta$. Then*

$$\frac{d_{C(S)}(\alpha, \beta)}{C(|\chi(S)|)} \leq \deg \mathfrak{C}_\Gamma,$$

where $C(x) = 3A_1(x)A_2(x)$ is a polynomial of degree 13.

Theorem 8.3 is related to work of M. Chu [Chu17] and J. DeBlois, N. Miller, and the second author [DMP18] on quantifying virtual specialness for various hyperbolic manifolds.

As a second application, we use Theorem 4.1 to effectively relate the electric circumference of a fibered manifold M_ϕ to the curve graph translation length $\ell_S(\phi)$ of its monodromy $\phi: S \rightarrow S$. (For definitions, see Section 9.)

Theorem (9.1). *If $\phi: S \rightarrow S$ is a pseudo-Anosov homeomorphism of a closed surface S , then*

$$\frac{1}{A_1(|\chi(S)|)} \cdot \ell_S(\phi) \leq \text{circ}_{\epsilon_S}(M_\phi) \leq A_2(|\chi(S)|) \cdot (\ell_S(\phi) + 2),$$

where the polynomials A_1 and A_2 are as in Equation (1)

The outline of the paper is as follows. In Section 2 we give the necessary background on curve graphs, Margulis tubes in hyperbolic manifolds, pleated surfaces and sweepouts. We then prove various lemmas regarding curves on surfaces and Margulis tubes in 3-manifolds in Section 3 before proving Theorem 4.1 in Sections 4, 5, and 6 and Theorem 7.1 in Section 7. In Section 8 we prove the application regarding cube complexes and in Section 9 we give the application to fibered manifolds.

ACKNOWLEDGEMENTS

The authors would like to thank Darren Long and Dave Futer for helpful discussions. The authors acknowledge support from U.S. National Science Foundation grants DMS 1107452, 1107263, 1107367 "RNMS: GEometric structures And Representation varieties" (the GEAR Network). The first author was partially supported by NSF Grants DMS-1502623 and DMS-1807319. The second author was partially supported by NSF Grant DMS-1812014. The third author was partially supported by NSF Grant DMS-1744551.

2. BACKGROUND

Given an orientable surface S of genus $g \geq 0$, $n \geq 0$ punctures, and without boundary, define $\omega(S) := 3g + n - 4$. We call an orientable, boundary-less surface of finite type S *non-sporadic* if $\omega(S) > 0$. In all that follows, in order to avoid trivial cases we will assume that all surfaces are non-sporadic.

2.1. Curves on surfaces. Recall that a simple curve is *essential* if it is neither null-homotopic nor peripheral (i.e. it doesn't bound a disk or once punctured disk on S). As is usual in the subject, we will generally refer to a vertex $\alpha \in \mathcal{C}(S)^{(0)}$ as a curve. We reserve the term *loop* to refer to an embedded circle in S . For example, with these conventions, a curve is represented by a loop.

Given curves $\alpha, \beta \in \mathcal{C}(S)^{(0)}$, their *geometric intersection number*, denoted $i(\alpha, \beta)$, is defined as

$$(2) \quad i(\alpha, \beta) = \min \{|a \cap b|\},$$

where the minimum is taken over loop representatives a and b of α and β , respectively. A surgery argument due to Hempel [Hem01] yields the following upper bound on distance in $\mathcal{C}(S)$ in terms of geometric intersection number:

$$(3) \quad d_{\mathcal{C}(S)}(\alpha, \beta) \leq 2 \log_2(i(\alpha, \beta)) + 2.$$

In particular, this shows that $\mathcal{C}(S)$ is connected.

More recently, Bowditch [Bow14] proved a stronger version of Equation (3) which is sensitive to the topology of the underlying surface, and which we will need in subsequent sections. We reformulate Corollary 2.2 of [Bow14] as follows:

$$(4) \quad d_{\mathcal{C}(S)}(\alpha, \beta) < 2 + 2 \cdot \frac{\log(i(\alpha, \beta)/2)}{\log((|\chi(S)| - 2)/2)},$$

so long as the denominator is well-defined and positive which is the case for all S with $|\chi(S)| \geq 5$.

2.2. Hyperbolic manifolds and Margulis tubes. Here we review some basic background on hyperbolic 2- and 3-manifolds. For additional details, see [Bus10].

Let S be a finite area hyperbolic surface and let p be a puncture of S . A peripheral curve about p corresponds to a parabolic element of $\pi_1(S)$. A horodisk in \mathbb{H}^2 based at a lift \tilde{p} of the puncture p will project to a neighborhood of p in S . There exists a horocycle \tilde{Q}_p such that the quotient of \tilde{Q}_p by $\text{stab}(\tilde{p}) < \pi_1(S)$ is not embedded, and that for any horocycle \tilde{H} based at \tilde{p} separating \tilde{Q}_p from \tilde{p} , $\tilde{H}/\text{stab}(\tilde{p})$ is embedded. We call \tilde{Q}_p the *maximal horocycle* for \tilde{p} . The region of S facing p and bounded by the quotient Q_p of \tilde{Q}_p is called the *maximal cusp neighborhood*.

There is another horocycle \tilde{H}_p based at \tilde{p} which projects to an embedded loop of length 2. The region bounded between this loop and p is called a *standard cusp neighborhood*. The standard cusp neighborhood is isometric to the cylinder $(-\infty, \log 2) \times S^1$ equipped with the metric

$$(5) \quad dx^2 + e^{2x} d\theta^2,$$

where $-\infty \leq x \leq \log(2)$ and $\theta \in S^1 = [0, 1]/(0 \sim 1)$ (see pages 110-112 of [Bus10]).

A key feature of hyperbolic geometry is that the volume of an n -dimensional ball of radius r , denoted $\text{Vol}_n(r)$, grows exponentially as a function of r . In subsequent sections we will need explicit formulae for this volume in low dimensions, so we record that information here:

$$(6) \quad \text{Vol}_2(r) = 4\pi \sinh^2(r/2), \quad \text{Vol}_3(r) = \pi(\sinh(2r) - 2r).$$

In particular,

$$(7) \quad \text{Vol}_2(r) = O(e^r), \quad \text{Vol}_3(r) = O(e^{2r}),$$

and the following limits exist:

$$(8) \quad \lim_{r \rightarrow 0} \frac{\text{Vol}_2(r)}{r^2}, \quad \frac{\text{Vol}_3(r)}{r^3}.$$

Given a hyperbolic manifold M and $\delta > 0$, the δ -thin part of M , denoted by $M_{(0, \delta]}$, is the set of points in M with injectivity radius at most $\delta/2$. Similarly, the δ -thick

part, $M_{[\delta, \infty)}$, consists of all points with injectivity radius at least $\delta/2$. Any hyperbolic manifold M is of course the union of its δ -thick and -thin parts.

For $n \geq 2$, there exists $\epsilon_n > 0$ called the n -dimensional Margulis constant so that the ϵ_n -thin part of any hyperbolic n -manifold M decomposes into a disjoint union of cusps and subsets of the form $\mathbb{T}_{\alpha_1}, \mathbb{T}_{\alpha_2}, \dots, \mathbb{T}_{\alpha_n}$ where \mathbb{T}_{α_i} is a tubular neighborhood of the closed geodesic α_i .

Thus the ϵ_2 -thin part of a hyperbolic surface is homeomorphic to a disjoint union of annuli, and the ϵ_3 -thin part of a hyperbolic 3-manifold decomposes as a disjoint union of solid tori and cusps.

Meyerhoff [Mey87] demonstrated the following lower bound for ϵ_3 , which we will subsequently need:

$$(9) \quad \epsilon_3 > 0.104.$$

Given $\delta \leq \epsilon_3$, we denote by $\mathbb{T}_\alpha(\delta)$ the component of $M_{(0, \delta]}$ containing the geodesic α . This is called the *Margulis tube* for α . We note that $\mathbb{T}_\alpha(\delta)$ can be empty if the length of α is greater than δ . We note here that there is a definite distance between $\mathbb{T}_\alpha(\delta)$ and $\partial\mathbb{T}_\alpha(\epsilon_3)$, which goes to infinity as $\delta \rightarrow 0$. A concrete estimate of this growth was obtained recently by Futer–Purcell–Schleimer [FPS18]. We will require this estimate in Section 3 and so we record it there in detail.

We now briefly discuss hyperbolic 3-manifolds $M = \mathbb{H}^3/\Gamma$ homeomorphic to $S \times \mathbb{R}$, as these manifolds are the focus of this paper. Here and throughout, we always consider such a manifold with a fixed homotopy equivalence $\iota: S \rightarrow M$, called a marking, which allows us to identify homotopy classes of curves in S with homotopy classes in M . Hence, for any essential curve α in S it makes sense to speak of its length $\ell_M(\alpha)$ in M , which is defined to be the length of the geodesic in M homotopic to $\iota(\alpha) \subset M$.

Associated to any such manifold M without accidental parabolics¹ is a pair of *end invariants* (λ^-, λ^+) , each of which is either:

- (1) (*non-degenerate*) a point in the Teichmüller space of S , namely a pair (f, σ) where σ is a complete hyperbolic metric on S and $f: S \rightarrow \sigma$ is a (homotopy class of) homeomorphism;
- (2) (*degenerate*) a lamination on S .

End invariants describe the behavior of the geometry of $M = \mathbb{H}^3/\Gamma$ as one exits out of each of the two topological ends $\mathcal{E}^-, \mathcal{E}^+$ of M . In the non-degenerate case, an end \mathcal{E} of M is foliated by surfaces S_t equipped with induced metrics that converge (after a re-scaling) to a hyperbolic metric on S .

In the degenerate case, Thurston [Thu78] proved that there exists a sequence of simple closed curves on S whose geodesic representatives exit \mathcal{E} , and which converge, in the proper sense, to a lamination on S . That an end, in general, behaves in exactly one of the above two ways follows from work of Bonahon [Bon86] and Canary [Can93], and ultimately the proof of the Tameness conjecture by Agol [Ago04] and Calegari–Gabai [CG06].

¹That is, where parabolics of Γ all come from peripheral loops in $\pi_1(S)$.

The celebrated Ending Lamination Theorem, proved by Minsky [Min10] and Brock-Canary-Minsky [BCM12], asserts that the end invariants (λ^+, λ^-) associated to $\mathcal{E}^+, \mathcal{E}^-$, respectively, determines the hyperbolic manifold M .

2.3. Pleated surfaces and sweepouts. Fixing a hyperbolic 3-manifold M , a topological surface S , and a lamination λ on S , a λ -pleated surface is a map $F : S \rightarrow M$ so that:

- (1) F is proper, and hence sends cusps to cusps,
- (2) for each leaf l of λ , $F(l) \subset M$ is a geodesic,
- (3) for each component R of $S \setminus \lambda$, $F(R)$ is totally geodesic.

The map F induces via pull-back a complete hyperbolic metric on the surface S . With respect to this metric, F is a 1-Lipschitz map of hyperbolic manifolds. Pleated surfaces arise very naturally in the study of hyperbolic 3-manifolds. For example, the convex core of a quasi-Fuchsian hyperbolic 3-manifold is always bounded by the image of a pair of pleated surfaces. Moreover, if M is homeomorphic to $S \times \mathbb{R}$, work of Thurston [Thu78] implies that if all leaves of λ can be realized as geodesics in M , there exists a λ -pleated surface into M .

Given a surface S and a hyperbolic 3-manifold M , a *1-Lipschitz sweepout* is a homotopy $f_t : X_t = (S, g_t) \rightarrow M$, where g_t is a hyperbolic metric on S and f_t is a 1-Lipschitz map for each t . Whenever M is homeomorphic to $S \times \mathbb{R}$, we only consider 1-Lipschitz maps homotopic to our fixed marking. An important result of Canary [Can96] yields the existence of 1-Lipschitz sweepouts interpolating between geodesics in M .

Theorem 2.1 (Canary). *Let $M = \mathbb{H}^3/\Gamma$ be a hyperbolic 3-manifold homeomorphic to $S \times \mathbb{R}$, so that $\rho \in \Gamma$ is parabolic if and only if it corresponds to a peripheral curve on S . Let α, β be a pair of simple closed curves on S with geodesic representatives α^*, β^* in M . Then there exists a 1-Lipschitz sweepout $f_t : X_t \rightarrow M, 0 \leq t \leq 1$, so that*

$$\alpha^* \subset f_0(X_0) \text{ and } \beta^* \subset f_1(X_1).$$

Moreover, there exists a 1-Lipschitz sweepout $g_t : S_t \rightarrow M, -\infty < t < \infty$, surjecting onto M and so that $g_t(X_t)$ exits out of \mathcal{E}^+ (resp. \mathcal{E}^-) as $t \rightarrow \infty$ (resp. $-\infty$).

Theorem 2.1 follows from Canary's work on *simplicial hyperbolic surfaces* [Can96]. These are path-isometric mappings into M of singular hyperbolic surfaces with cone points coinciding with vertices of a geodesic triangulation. Concretely, Canary shows the existence of 1-Lipschitz sweepouts where g_t is a simplicial hyperbolic surface. Brock [Bro03a] reformulates Canary's construction by uniformizing and replacing each g_t with the unique non-singular hyperbolic metric in its conformal class. For additional details, see the proof of Lemma 4.2 in [Bro03a].

3. HYPERBOLIC SURFACES AND 3-MANIFOLDS

In this section we cover some fairly basic results in hyperbolic geometry. While nothing here will be surprising to experts, care must be taken in order to keep track of how the constants involved depend on the underlying parameters.

Let \hat{S}_δ denote the compact subset of S obtained by deleting neighborhoods of each cusp consisting of points with injectivity radius at most $\delta/2$.

Lemma 3.1. *Let S be an orientable surface with $\chi(S) < 0$ which is not a 3-times punctured sphere. Then for any $\delta \geq 0$ there is a constant $L_{S,\delta}$ such that for any finite area hyperbolic structure on S and any x in \hat{S}_δ , there is an essential loop in S through x of length less than $L_{S,\delta}$.*

Proof. Let \tilde{x} be any lift of x to the universal cover $\tilde{S} = \mathbb{H}^2$, and let $\pi : \mathbb{H}^2 \rightarrow S$ denote the universal covering map. Let \tilde{B} be a lift of a maximal embedded open ball centered at x . By maximality, there must be a pair of points z, y on the boundary of \tilde{B} which lie in the same fiber over S . Then if $[a, b]$ represents the geodesic segment with endpoints $a, b \in \mathbb{H}^2$,

$$\tilde{\rho} := \pi([\tilde{x}, z]) * \pi([y, \tilde{x}])$$

is a loop ρ representing a non-trivial element of $\pi_1(S, p)$. If ρ is not simple, there will be a representative for an essential simple loop through x which is contained in the image of ρ (and whose length is therefore bounded above by that of ρ), and we replace ρ with this.

Recall that the area of a hyperbolic disk of radius $r > 0$ is $4\pi \sinh^2(r/2)$, and therefore by the Gauss-Bonnet theorem the radius of \tilde{B} is at most

$$2 \sinh^{-1} \left(\sqrt{|\chi(S)|/2} \right).$$

Hence the length of $\tilde{\rho}$ is at most

$$(10) \quad \ell_S := 4 \sinh^{-1} \left(\sqrt{|\chi(S)|/2} \right) = 4 \log \left(\sqrt{|\chi(S)|/2} + \sqrt{1 + |\chi(S)|/2} \right).$$

When S is closed, this concludes the proof. When S has cusps, the above argument gives us the desired loop unless the element of $\pi_1(S, x)$ represented by ρ is peripheral about a puncture, p . In this case, recall that Q_p (resp. H_p) denotes the quotient of a maximal horocycle \tilde{Q}_p (resp. a standard horocycle \tilde{H}_p).

Suppose first that x lies in the standard cusp neighborhood. Since $x \in \hat{S}_\delta$, (5) implies that the distance $d_S(x, H_p)$ between x and H_p satisfies

$$e^{-d_S(x, H_p)} \geq \delta,$$

and therefore

$$(11) \quad d_S(x, H_p) \leq \log(1/\delta).$$

Let N_p be the subset of the maximal cusp neighborhood bounded by Q_p and H_p . Since the area of the neighborhood of a cusp is equal to the length of its boundary, by the Gauss-Bonnet theorem we have that Q_p has length at most $2\pi|\chi(S)|$. The region N_p can be lifted to a rectangle \tilde{N}_p in the upper half-plane model which is (up to isometry) of the form

$$\tilde{N}_p = \{(y, z) \in \mathbb{H}^2 : 0 \leq y < a, 0 < r \leq z \leq b\},$$

for some positive a and b . Then H_p lifts to the top edge of \tilde{N}_p and Q_p lifts to the bottom edge. Therefore,

$$\ell(H_p) = \frac{a}{b} = 2; \ell(Q_p) = \frac{a}{r} \leq 2\pi|\chi(S)|.$$

Hence $a = 2b$ and so $r \geq b/\pi|\chi(S)|$, and so

$$(12) \quad d_S(H_p, Q_p) \leq \log(\pi|\chi(S)|).$$

By maximality, a fundamental lift \tilde{Q}_p of Q_p will project to the concatenation of (at least) two loops ρ_1 and ρ_2 on S . It cannot be the case that *both* ρ_i are peripheral since otherwise S would be a thrice punctured sphere. Moreover, there is a representative of ρ_i of length at most $2\log(\pi|\chi(S)|) + 2$; indeed, if $[\tilde{y}_i, \tilde{z}_i]$ is a geodesic segment projecting to ρ_i , by (12) there is a path from \tilde{y}_i to \tilde{H}_p of length at most $\log(\pi|\chi(S)|)$. Concatenating this with a segment of \tilde{H}_p as necessary and then (applying (12) again) with a segment back to \tilde{z}_i yields the desired bound. Abusing notation slightly, we refer to these representatives as ρ_1, ρ_2 and we note also that ρ_i is contained completely within $\overline{N_p}$ and ρ_i touches H_p .

Since x lies within the standard cusp neighborhood, x must be within a distance of at most $2 + \log(1/\delta)$ from each ρ_i , and therefore there is an essential loop through x (which, using the same argument as above we can assume is simple) of length at most

$$2(2 + \log(1/\delta)) + 2\log(\pi|\chi(S)|) + 2.$$

If $x \in N_c$, it can be at most $\log(\pi|\chi(S)|) + 2$ from each ρ_i and thus there is an essential loop through x of length at most

$$2(\log(\pi|\chi(S)|) + 2) + 2\log(\pi|\chi(S)|) + 2 = 4\log(\pi|\chi(S)|) + 6.$$

It remains to consider the case that x is separated from the puncture by Q_p . Recall the simple loop ρ constructed in the first part of the argument, and that we are assuming that ρ is peripheral. We claim that ρ must touch Q_p . Indeed, let $\tilde{p} \in \partial_\infty(\mathbb{H}^2)$ denote a lift of the puncture p and let \tilde{Q}_p be the horocycle based at \tilde{p} projecting to Q_p . Since ρ is peripheral, there is a lift $\tilde{\rho}$ bounded by lifts \tilde{x}_1, \tilde{x}_2 of x so that \tilde{x}_1 and \tilde{x}_2 are on the same horocycle R based at \tilde{p} . By maximality of \tilde{Q}_p , there is another lift \tilde{Q}_p' of Q_p that is tangent to \tilde{Q}_p and which intersects R at two points (see Figure 1).

Letting $g \in \pi_1(S)$ be the parabolic element corresponding to the peripheral loop ρ , all translates of \tilde{x}_1 under the action of g on \mathbb{H}^2 lie along R and lie outside of all lifts of the horoball $\overline{Q_p}$ bounded by Q_p since x is separated from p by Q_p . Therefore, there exists a lift $\tilde{\rho}'$ of ρ with endpoints \tilde{x}_3 and \tilde{x}_4 that lie along R such that $\overline{\tilde{Q}_p} \cup \overline{\tilde{Q}_p}'$ separate \tilde{x}_3 and \tilde{x}_4 (again see Figure 1). Thus, $\tilde{\rho}'$ must intersect $\tilde{Q}_p \cup \tilde{Q}_p'$ so that ρ touches Q_p .

Since the length of ρ is at most ℓ_S , it follows that x must be a distance of at most $\ell_S/2$ from Q_p , and hence from the region N_p . Thus, there is an essential simple loop through x of length at most $\ell_S + 4\log(\pi|\chi(S)|) + 6$, and so in all three cases the loop constructed has length at most

$$L_{S,\delta} = 2\log(1/\delta) + \ell_S + 4\log(\pi|\chi(S)|) + 6.$$

□

Remark 3.2. We use the proof of Lemma 3.1 to give an upper bound on $L_{S,\delta}$ that will be useful in subsequent sections. First we use (10) to give the following upper bound

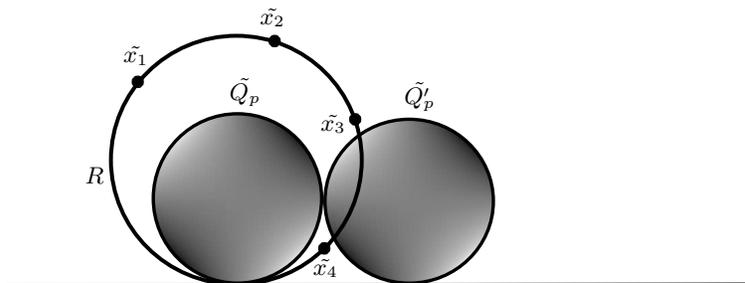


FIGURE 1. When ρ is separated from p by Q_p , it must touch Q_p .

on ℓ_S :

$$(13) \quad \ell_S \leq 4 \log(2\sqrt{|\chi(S)|}) = 4 \log(2) + 2 \log(|\chi(S)|) \leq 4 + 2 \log(|\chi(S)|).$$

Note the proof of Lemma 3.1 is much simpler when S is closed and we can use ℓ_S in place of $L_{S,\delta}$ in this case. Combining (13) and the definition of $L_{S,\delta}$ in the general setting we have:

$$(14) \quad L_{S,\delta} \leq 2 \log(1/\delta) + 6 \log(\pi|\chi(S)|) + 9.$$

Additionally, when S has punctures, we let L_S denote L_{S,ϵ_3} , which by (9) is at most

$$6 \log(\pi|\chi(S)|) + 14.$$

When S is closed we set $L_S = \ell_S$, which by (13) is at most

$$2 \log(|\chi(S)|) + 4.$$

Lemma 3.3. *Let M be a hyperbolic manifold with $M \cong S \times \mathbb{R}$, and let α be an essential curve on S . There is a positive constant $\epsilon_S < \epsilon_3$ such that if $f: S \rightarrow M$ is a π_1 -injective 1-Lipschitz map such that $f(S) \cap \mathbb{T}_\alpha(\epsilon_S) \neq \emptyset$, then $\ell_S(\alpha) \leq L_S$.*

Further, there is a loop in the isotopy class of α whose length is less than L_S in S and whose image in M is contained in $\mathbb{T}_\alpha(\epsilon_3)$.

Proof. Given a positive $\mu < \epsilon_3$ and a non-empty μ -tube $\mathbb{T}_\alpha(\mu)$, let $\mathcal{F}_\alpha(\mu)$ denote the distance between the boundary of the Margulis tube and $\mathbb{T}_\alpha(\mu)$. The function \mathcal{F}_α is decreasing in μ , and Theorem 1.1 of Futer–Purcell–Schleimer [FPS18] states that

$$(15) \quad \mathcal{F}_\alpha(\mu) \geq \mathcal{F}(\mu) := \operatorname{arccosh} \frac{\epsilon_3}{\sqrt{7.256\mu}} - .0424.$$

Let $\epsilon_S = \mathcal{F}^{-1}(L_S/2)$. Hence (15) implies

$$\mathcal{F}(\epsilon_S) = L_S/2 = \operatorname{arccosh} \frac{\epsilon_3}{\sqrt{7.256\epsilon_S}} - .0424,$$

and thus,

$$\epsilon_S = \frac{\epsilon_3^2}{7.256 \cdot \cosh^2(L_S/2 + .0424)}.$$

Applying the upper bounds on L_S obtained at the end of Remark 3.2 and (9), we see that

$$(16) \quad \epsilon_S \geq \frac{\epsilon_3^2}{8 \cdot \cosh^2(\log(|\chi(S)|) + 2)} \geq \frac{1}{e^{10} \cdot |\chi(S)|^2}$$

when S is closed, and

$$(17) \quad \epsilon_S \geq \frac{\epsilon_3^2}{8 \cdot \cosh^2(3 \log(\pi|\chi(S)|) + 7)} \geq \frac{1}{e^{20\pi^6} \cdot |\chi(S)|^6},$$

when S has punctures. The lower bound on ϵ_S obtained above does not play a direct role in the proof of Lemma 3.3, but we will refer to this bound in later sections.

If $f(x) \in \mathbb{T}_\alpha(\mu)$ for $\mu < \epsilon_3$, then $x \in \hat{S}_{\epsilon_3}$ since the ϵ_3 -thin part of any cusp neighborhood must map via f into the ϵ_3 -thin part of a cusp neighborhood of M , and any Margulis tube in M is disjoint from all ϵ_3 -thin cusp neighborhoods in M . Thus by Lemma 3.1 there is an essential simple loop ρ through x of length at most L_S .

Since the map is 1-Lipschitz, $f(\rho)$ has length less than L_S and meets $\mathbb{T}_\alpha(\epsilon_S)$. Hence, any point on $f(\rho)$ has distance at most $L_S/2$ from $\mathbb{T}_\alpha(\epsilon_S)$, and so by construction $f(\rho) \subset \mathbb{T}_\alpha(\epsilon_3)$.

As f induces an isomorphism on π_1 , $f(\rho)$ is homotopic to some power of α . But, ρ is a simple curve on S and so we must have that ρ and α are isotopic curves on S . Set $\alpha = \rho$ and note that $\ell_S(\alpha) = \ell_S(\rho) \leq L_S$. \square

Remark 3.4. Notice in the proof that ϵ_S is chosen so that it is bigger than the quantities in Equations (16) and (17), but also small enough in comparison to ϵ_3 so that the distance between the ϵ_S -tube and the boundary of the ϵ_3 -tube is at least $\mathcal{F}(\epsilon_S) \geq 2$.

Lemma 3.5. *There is a universal constant D so that if γ_1 and γ_2 are essential loops on S with length less than L_S , then $d_{\mathcal{C}(S)}(\gamma_1, \gamma_2) \leq D$.*

Proof. By the collar lemma, γ_1 has an embedded annular neighborhood of width at least

$$\log(\coth(\ell(\gamma_1))/4) =: c(\gamma_1).$$

Since $\log(\coth(x/4)) \rightarrow \infty$ as $x \rightarrow 0$ and decays to 0 as $x \rightarrow \infty$, there is some positive constant c so that

$$c/2 = \log(\coth(c/4)).$$

By inspection we see that $c < 2$.

We first present a proof in the special case that S is closed, as in this setting the argument is more conceptual:

S is closed: Assume that γ_1 is the shortest curve on S , and let $\mathcal{N}(\gamma_1)$ denote a maximal collar neighborhood of γ_1 . Let $\tilde{\mathcal{N}}$ denote a lift of the boundary of $\mathcal{N}(\gamma_1)$ to the universal cover. By maximality, there is a pair of points $\tilde{x}, \tilde{y} \in \tilde{\mathcal{N}}$ which project to the same point on the boundary of $\mathcal{N}(\gamma_1)$. If \tilde{x}, \tilde{y} lie on opposite boundary components, the geodesic segment connecting them must project to an essential loop ρ , since it intersects γ_1 exactly once. Moreover, in this case the width of $\mathcal{N}(\gamma_1)$ must be at least $\ell(\gamma_1)/2$, for

if not there will be a representative of ρ with length less than $\ell(\gamma_1)$. Indeed, ρ is the loop which lifts to $\tilde{\rho}$ which begins at \tilde{x} with the perpendicular between \tilde{x} and $\tilde{\gamma}_1$, then travels at most half of $\tilde{\gamma}_1$ (wrapping around if necessary), and then concludes with the perpendicular to $\tilde{\gamma}_1$ containing \tilde{y} .

If \tilde{x}, \tilde{y} lie on the same boundary component \mathcal{E} of $\tilde{\mathcal{N}}$, a loop ρ is constructed in the exact same fashion as at the end of the previous paragraph, although it is now no longer the case that ρ intersects γ_1 once. However ρ must still be essential, for if not, ρ would lift to the boundary of a geodesic triangle with three non-ideal vertices and two right angles. Since S is closed, ρ is non-peripheral and γ_1 must admit a collar neighborhood of width at least $\ell(\gamma_1)/2$ as in the previous paragraph.

Thus, γ_1 must admit a collar neighborhood of width at least $\ell(\gamma_1)/2$. Therefore,

$$\begin{aligned} i(\gamma_2, \gamma_1) &\leq \min[2\ell(\gamma_2)/\ell(\gamma_1), \ell(\gamma_2)/c(\gamma_1)] \\ &\leq L_S \cdot \min[2/\ell(\gamma_1), 1/c(\gamma_1)]. \end{aligned}$$

Thus, if $\ell(\gamma_1) > 2$, $i(\gamma_1, \gamma_2) \leq L_S$, and if $\ell(\gamma_1) < 2$, $c(\gamma_1) > 1/2$, so we have $i(\gamma_1, \gamma_2) \leq 2L_S$.

Consider the case where $|\chi(S)| \geq 5$. Using (4), we note that if $d_{\mathcal{C}(S)}(\gamma_1, \gamma_2) \geq 6$, then γ_1 and γ_2 must intersect at least $\frac{(|\chi(S)|-2)^2}{2}$ times. Using the fact that $L_S \leq 2 \log(|\chi(S)|) + 4$ for closed surfaces, it follows that $d_{\mathcal{C}(S)}(\gamma_1, \gamma_2) \leq 5$ so long as $|\chi(S)| \geq 8$ since

$$x \geq 8 \Rightarrow \frac{(x-2)^2}{2} > 2 \cdot (2 \log(x) + 4).$$

For the finite list of remaining surfaces, we use the fact that on any surface S ,

$$d_{\mathcal{C}(S)}(\alpha, \beta) \leq 2 \log_2(i(\alpha, \beta)) + 2.$$

Note that if $|\chi(S)| < 8$, $i(\gamma_1, \gamma_2) \leq 2L_S < 16.32$, so we must have

$$d_{\mathcal{C}(S)}(\gamma_1, \gamma_2) \leq 10.$$

In general, let α represent the systole of S ; γ_1 need not coincide with α , but the above argument shows that

$$d_{\mathcal{C}(S)}(\gamma_1, \alpha) \leq \begin{cases} 10 & |\chi(S)| < 8 \\ 5 & |\chi(S)| \geq 8 \end{cases}$$

and so by applying the same argument to γ_2 and then using the triangle inequality,

$$d_{\mathcal{C}(S)}(\gamma_1, \gamma_2) \leq \begin{cases} 20 & |\chi(S)| < 8 \\ 10 & |\chi(S)| \geq 8 \end{cases}$$

The non-closed case:

As for the general case where S is not necessarily closed, we use the collar lemma and argue that γ_1, γ_2 each have embedded collar neighborhoods of width at least

$$\log(\coth(L_S/4)),$$

which, applying (14), is at least

$$\log \left(\coth \left(\frac{3}{2} \log(\pi|\chi(S)|) + \frac{7}{2} \right) \right).$$

It follows that

$$\begin{aligned} i(\gamma_1, \gamma_2) &\leq \frac{6 \log(\pi|\chi(S)|) + 14}{\log \left(\coth \left(\frac{3}{2} \log(\pi|\chi(S)|) + \frac{7}{2} \right) \right)} \\ (18) \quad &= \frac{6 \log(\pi|\chi(S)|) + 14}{\log \left(\frac{e^7 \pi^3 |\chi(S)|^3 + 1}{e^7 \pi^3 |\chi(S)|^3 - 1} \right)} =: W_S. \end{aligned}$$

We compute directly that

$$W_S \leq 2^{40} \cdot |\chi(S)|^3 \log(|\chi(S)|).$$

Assuming that $|\chi(S)| \geq 10$ and using (4), we conclude that

$$\begin{aligned} d_{\mathcal{C}(S)}(\gamma_1, \gamma_2) &\leq 2 + 2 \cdot \frac{\log(2^{39} |\chi(S)|^3 \log(|\chi(S)|))}{\log((|\chi(S)| - 2)/2)} \\ &\leq 2 + \frac{78}{\log(|\chi(S)| - 2) - 1} + \frac{8 \cdot \log |\chi(S)|}{\log(|\chi(S)| - 2) - 1} =: 2 + A + B. \end{aligned}$$

As we are assuming that $|\chi(S)| \geq 10$, this is in turn at most

$$2 + 72.3 + 17.06 < 92.$$

On the other hand, if $|\chi(S)| < 10$, W_S is bounded from above and applying (3) yields

$$d_{\mathcal{C}(S)}(\gamma_1, \gamma_2) \leq 104.$$

In conclusion,

$$d_{\mathcal{C}(S)}(\gamma_1, \gamma_2) \leq \begin{cases} 104 & |\chi(S)| < 10 \\ 92 & |\chi(S)| \geq 10 \end{cases}$$

We note that as $|\chi(S)| \rightarrow \infty$, $A \rightarrow 0$, $B \rightarrow 8$ and thus for sufficiently large Euler characteristic, we obtain the much smaller bound of 11. Moreover, using a stronger version of (4) due to the first author [Aou13], one can conclude that for all S with $|\chi(S)|$ sufficiently large,

$$d_{\mathcal{C}(S)}(\gamma_1, \gamma_2) < 6.$$

□

Lemma 3.6. *Let $0 \leq \delta \leq 1$ and $L \geq 1$. Fix $x \in M_{[\delta, \infty)}$. Then the number of homotopy classes of loops of length less than L based at x is less than*

$$P(L, \delta) := \frac{\text{Vol}_3(L + \delta)}{\text{Vol}_3(\delta)}.$$

Proof. The argument is standard, but we provide it for the reader's convenience.

Let $\mathbb{H}^3 \rightarrow M$ be the universal covering and let \tilde{x} be a fixed lift of x . Let B' be the ball of radius L about \tilde{x} so that the based homotopy classes of loops of length less than L at x in M correspond to the translates of \tilde{x} in \mathbb{H} contained in B' . Since $x \in M_{[\delta, \infty)}$, the δ -balls about these translates are all disjoint, and since they are contained in the ball B of radius $L + \delta$ about \tilde{x} , we see that the number of such points is bounded by $\frac{\text{Vol}_3(L+\delta)}{\text{Vol}_3(\delta)}$, as required. \square

Remark 3.7. Using (6), (7), and (8), we have that

$$(19) \quad P(L, \delta) = \frac{\text{Vol}_3(L + \delta)}{\text{Vol}_3(\delta)} = \frac{\sinh(2(L + \delta)) - 2(L + \delta)}{\sinh(2\delta) - 2\delta},$$

For large L and small δ ,

$$(20) \quad P(L, \delta) = \frac{\sinh(2(L + \delta)) - 2(L + \delta)}{(2\delta + \frac{1}{6}(2\delta)^3 + \frac{1}{120}(2\delta)^5 + \dots) - 2\delta} \leq \frac{\sinh(2(L + \delta))}{\delta^3}.$$

4. ELECTRIC DISTANCE

For a hyperbolic manifold M , let d_M denote distance in the hyperbolic metric. Fixing $0 < \delta \leq \epsilon$, let \check{M}_δ denote the manifold obtained from M by removing δ -thin cusps. Of course, when M has no cusps, $M = \check{M}_\delta$. For two points $x, y \in \check{M}_\delta$, their δ -electric distance is defined as

$$d_M^\delta(x, y) = \inf\{\text{length}(p \cap M_{[\delta, \infty)})\}$$

where p varies over all paths with image contained in \check{M}_δ , joining x and y . When M has no cusps, this is the length of the portion of the shortest hyperbolic geodesic joining x and y that occurs outside of the δ -tubes of M . Our main technical result is an explicit inequality relating distance in the curve graph of S with electric distance in M .

Theorem 4.1. *Let α and β be curves in S and let $M \cong S \times \mathbb{R}$ be a hyperbolic manifold such that $\ell_M(\alpha), \ell_M(\beta) \leq \epsilon_S$. Then*

$$1/A_1(|\chi(S)|) \cdot d_{C(S)}(\alpha, \beta) \leq d_M^{\epsilon_S}(\alpha, \beta) \leq A_2(|\chi(S)|) \cdot d_{C(S)}(\alpha, \beta),$$

where the polynomials A_1 and A_2 are as in Equation (1).

The proof will be completed over the next several sections.

Fix $0 < \delta \leq \epsilon_3$ and let α be a curve in the hyperbolic manifold M . Let $\text{tube}_M(\alpha)$ denote the δ -tube $\mathbb{T}_\alpha(\delta)$ about the geodesic representative of α in M in the case where α is δ -short. Otherwise, set $\text{tube}_M(\alpha)$ to be the geodesic representative of α .

The idea behind the following proposition is simple and well-known to experts.

Proposition 4.2. *Let $0 < \eta < \frac{\epsilon_3}{e^{6(\pi|\chi(S)|)^3}}$. Then for any curves α and β in S ,*

$$d_M^\eta(\alpha, \beta) \leq \frac{4\pi|\chi(S)|}{\eta} \cdot d_{C(S)}(\alpha, \beta).$$

Remark 4.3. When S is closed, η need only be less than the Margulis constant ϵ_3 . This fact will be used in the proof of Theorem 7.1.

Proof. Let $\alpha = \alpha_0, \alpha_1, \dots, \alpha_n = \beta$ be a curve graph geodesic from α to β . For each i , let $f_i: X_i = (S, g_i) \rightarrow M$ be a pleated surface through $\alpha_i \cup \alpha_{i+1}$. In particular, f_i maps each of the geodesic representatives of α_i and α_{i+1} in X_i to its geodesic representative in M .

Let p_i be the shortest path in X_i from the geodesic representative of α_i to the geodesic representative of α_{i+1} . The bounded diameter lemma of Thurston and Bonahon gives that $\text{length}(p_i \cap (X_i)_{[\eta, \infty)}) \leq \frac{4\pi|\chi(S)|}{\eta}$. Indeed, if $p_i^\eta = p_i \cap (X_i)_{[\eta, \infty)}$, then the $\eta/2$ -neighborhood C_i of p_i^η is embedded in X_i and so

$$\eta/2 \cdot \ell_{X_i}(p_i^\eta) \leq \text{Area}(C_i) \leq 2\pi|\chi(S)|.$$

Since f_i is 1-Lipschitz, it maps η -thin parts of X_i to η -thin parts of M and so $\text{length}(f_i(p_i) \cap M_{[\eta, \infty)}) \leq \ell_{X_i}(p_i^\eta)$. When M has no cusps, this immediately gives that $d_M^\eta(\text{tube}_M(\alpha_i), \text{tube}_M(\alpha_{i+1})) \leq \ell_{X_i}(p_i^\eta)$.

In the presence of cusps, we argue as follows: First, we claim that p_i cannot enter any horocyclic cusp neighborhood in X_i whose boundary has length $2/e$. To see this, begin with the standard fact that simple closed geodesics on X_i do not enter any standard cusp neighborhoods. So for any cusp of X_i , the endpoints of p_i lie outside of its standard cusp neighborhood. Since p_i is embedded, the length of any component of its intersection with a standard cusp neighborhood is no more than 2. Hence, its deepest point in the standard cusp neighborhood has distance no more than 1 from the horocycle boundary. This means that it does not cross the horocycle for that cusp of length $2/e$.

Now suppose that there is some $z \in p_i$ such that $f_i(z)$ lies in an η -cusp of M . Then any nontrivial loop based at z whose length is less than $2 \log(\epsilon_3/\eta)$ must be peripheral. This is because the image of such a loop is entirely contained in the corresponding ϵ_3 -cusp of M and so the loop must represent a peripheral element of $\pi_1 S$. But since

$$\eta < \frac{\epsilon_3}{e^{6(\pi|\chi(S)|)^3}},$$

we see that every loop of length no more than $6 \log(\pi|\chi(S)|) + 12$ based at z is peripheral. However, the fact that p_i does not enter any horocyclic cusp neighborhood in X_i whose boundary has length $2/e$, together with Lemma 3.1 and Equation (14), implies that every point along p_i is the basepoint of some essential (i.e. nonperipheral) loop of length no more than $6 \log(\pi|\chi(S)|) + 11$, a contradiction. Here we are using the fact that the injectivity radius along p_i is at least $1/e$ so we set $\delta = 1/e$ in (14).

We conclude that $f_i(p_i)$ does not enter any η -cusps of M . Hence, just as in the case without cusps, we conclude that $d_M^\eta(\text{tube}_M(\alpha_i), \text{tube}_M(\alpha_{i+1})) \leq \ell_{X_i}(p_i^\eta)$.

Finally, using the fact that f_i maps the η -thin part of X_i to the η -thin part of M , we obtain

$$\begin{aligned} d_M^\eta(\alpha, \beta) &\leq \sum_{i=0}^{n-1} d_M^\eta(\text{tube}_M(\alpha_i), \text{tube}_M(\alpha_{i+1})) \\ &\leq \sum_{i=0}^{n-1} \ell_{X_i}(p_i^\eta) \\ &\leq \frac{4\pi|\chi(S)|}{\eta} \cdot d_{\mathcal{C}(S)}(\alpha, \beta). \end{aligned}$$

□

We label the coefficient at the end of the proof above by

$$(21) \quad \mathcal{A}(x, \eta) = \frac{4\pi x}{\eta}.$$

Thus, (21) and (16) yield the inequality

$$(22) \quad \mathcal{A}(|\chi(S)|, \epsilon_S) \leq A_2(|\chi(S)|),$$

for

$$(23) \quad A_2(x) = \begin{cases} 4 e^{10} \pi x^3 & \text{for } S \text{ closed} \\ 4 e^{20} \pi^7 x^7 & \text{otherwise,} \end{cases}$$

as in Equation (1).

Hence, we can complete the proof of the upper bound in Theorem 4.1 using Proposition 4.2 with $\eta = \epsilon_S$ so that

$$(24) \quad \begin{aligned} d_M^{\epsilon_S}(\alpha, \beta) &\leq \mathcal{A}(|\chi(S)|, \epsilon_S) \cdot d_{\mathcal{C}(S)}(\alpha, \beta) \\ &\leq A_2(|\chi(S)|) \cdot d_{\mathcal{C}(S)}(\alpha, \beta) \end{aligned}$$

Note that for this upper bound there is no requirement on the lengths of α, β .

The main idea for the other direction is contained in the following lemma. Roughly, the lemma says that as long as we can find a sweepout between the curves α and β which separates α from β at all times, then we obtain the desired bound on curve graph distance in terms of electric distance in M . The fact that we can find such a sweepout will be proved in the next section.

Lemma 4.4. *Let α and β be curves in S and $M \cong S \times \mathbb{R}$ a hyperbolic manifold such that $\ell_M(\alpha), \ell_M(\beta) \leq \epsilon_S$. Let p be a path in M joining $\mathbb{T}_\alpha(\epsilon_S)$ and $\mathbb{T}_\beta(\epsilon_S)$ and suppose that*

- (1) p is contained in $M_{[\epsilon_S, \infty)}$,
- (2) there is a 1-Lipschitz sweepout $(f_t: X_t = (S, g_t) \rightarrow M)_{t \in [0, 1]}$ such that $f_t(S) \cap p \neq \emptyset$ for all $t \in [0, 1]$, and
- (3) $f_0(S) \cap \mathbb{T}_\alpha(\epsilon_S) \neq \emptyset$ and $f_1(S) \cap \mathbb{T}_\beta(\epsilon_S) \neq \emptyset$.

Then

$$d_{\mathcal{C}(S)}(\alpha, \beta) \leq A_1(|\chi(S)|) \cdot \ell_M(p)$$

for $A_1(x)$ as in Equation (1).

Proof. Note that by Lemma 3.3, $\ell_{X_a}(\alpha) \leq L_S$ and $\ell_{X_b}(\beta) \leq L_S$. For each curve γ on S , set

$$I(\gamma) = \{t \in [a, b] : \ell_{X_t}(\gamma_p) \leq L_S\}.$$

Here γ_p is the shortest loop over all representatives of γ on S with the property that $f_t(\gamma_p)$ passes through the geodesic p . By (2) (and Lemma 3.1), these intervals cover $[a, b]$.

Now break p up into N segments p_1, \dots, p_N , the first $N - 1$ of which have length 1 and the last of which has length less than 2 so that $N = \lfloor \ell_M(p) \rfloor$. Let I_i be the set of all γ such that there is a $t \in I(\gamma)$ with $f_t(\gamma_p) \cap p_i \neq \emptyset$. That is, γ can be realized as a loop of length less than L_S starting at a point along p_i . We note that it is possible that I_i is empty since we are not requiring that each point of p is in the image of a slice of the sweepout.

By criterion (1), Lemma 3.6, and Remark 3.7,

$$\#I_i \leq \frac{\text{Vol}_3(L_S + \epsilon_S + 2)}{\text{Vol}_3(\epsilon_S)} \leq \frac{\sinh(2(L_S + \epsilon_S + 2))}{\epsilon_S^3} \leq \frac{e^{2(L_S + \epsilon_S + 2)}}{\epsilon_S^3}.$$

Furthermore,

$$\begin{aligned} e^{2(L_S + \epsilon_S + 2)} &\leq e^{2L_S + 6} \leq e^6 \cdot e^{4 \log(|\chi(S)|) + 8} \\ &= e^{14} |\chi(S)|^4, \end{aligned}$$

when S is closed, and

$$\begin{aligned} e^{2(L_S + \epsilon_S + 2)} &\leq e^6 \cdot e^{12 \log(\pi |\chi(S)|) + 28} \\ &= e^{34} \pi^{12} |\chi(S)|^{12}, \end{aligned}$$

when S has punctures. Along with the lower bounds on ϵ_S established in (16) and (17) this implies,

$$\#I_i \leq \begin{cases} e^{44} |\chi(S)|^{10} & \text{for } S \text{ closed} \\ e^{94} \pi^{30} |\chi(S)|^{30} & \text{otherwise.} \end{cases}$$

By Lemma 3.5, if $I(\gamma_1) \cap I(\gamma_2) \neq \emptyset$, $d_{\mathcal{C}(S)}(\gamma_1, \gamma_2) \leq D$, where D is the universal constant from Lemma 3.5. Further, by the above paragraph, the total number of such curves seen along p is no more than $\#I_i$. Since the sweepout passes through this set of curves with jumps of size no more than D ,

$$\begin{aligned} d_{\mathcal{C}(S)}(\alpha, \beta) &\leq D \cdot \#I_i \cdot N \\ &\leq D \cdot \#I_i \cdot \ell_M(p) \end{aligned}$$

By the proof of Lemma 3.5, $D \leq 20$ when S is closed and $D \leq 104$ in general, so that setting

$$(25) \quad A_1(x) = \begin{cases} 20 e^{44x^{10}} & \text{for } S \text{ closed} \\ 104 e^{94\pi^{30}x^{30}} & \text{otherwise} \end{cases}$$

as in Equation (1) gives us

$$d_{C(S)}(\alpha, \beta) \leq A_1(|\chi(S)|) \cdot \ell_M(p).$$

□

5. SEPARATING SWEEPOUTS

In what follows, let \mathbb{T}_α be shorthand for the tube $\mathbb{T}_\alpha(\epsilon_S)$. In order to find sweepouts satisfying the conditions of Lemma 4.4, we require the following:

Proposition 5.1. *Let α, β be intersecting curves on S whose lengths in M are no more than ϵ_S . Let $f_t: S \rightarrow M$, $t \in [a, b]$ be a 1-Lipschitz sweepout such that Σ_a lies to the left of α, β and Σ_b lies to the right of α, β , where $\Sigma_t = f_t(S)$. Then there is a subinterval $[c, d] \subset [a, b]$ such that*

- (1) Both \mathbb{T}_α and \mathbb{T}_β meet $\Sigma_c \cup \Sigma_d$,
- (2) Neither \mathbb{T}_α nor \mathbb{T}_β meet Σ_t for $t \in (c, d)$, and
- (3) Σ_t separates \mathbb{T}_α from \mathbb{T}_β for each $t \in (c, d)$.

The proof requires some notation. Let $m_\alpha \subset [a, b]$ be the set of times the sweepout meets \mathbb{T}_α :

$$m_\alpha = \{t \in [a, b] : \Sigma_t \cap \mathbb{T}_\alpha \neq \emptyset\}.$$

Define m_β similarly, and note that m_α and m_β are *disjoint* closed subsets of $[a, b]$, since no 1-Lipschitz map can meet both \mathbb{T}_α and \mathbb{T}_β . This follows from the fact that if Σ_t meets both \mathbb{T}_α and \mathbb{T}_β , then by Lemma 3.3 there are representative loops a and b on S such that $f_t(a) \subset \mathbb{T}_\alpha$ and $f_t(b) \subset \mathbb{T}_\beta$, and so a and b are disjoint. This contradicts the assumption that α and β intersect.

The components of $[a, b] \setminus m_\alpha$ are open in the interval $[a, b]$, and each is a subset of one of three disjoint subsets of $[a, b]$, denoted $l_\alpha, r_\alpha, b_\alpha$ and defined as follows. By definition l_α consists of those times when Σ_t is to the left of \mathbb{T}_α . This means that \mathbb{T}_α lies in the component of $M \setminus \Sigma_t$ containing the λ^+ end of M . Similarly, let r_α be those times for which Σ_t lies to the right of \mathbb{T}_α , and let b_α be those time when \mathbb{T}_α lies in a bounded component of $M \setminus \Sigma_t$. Since Σ_t always separates M , $[a, b] \setminus m_\alpha = l_\alpha \cup b_\alpha \cup r_\alpha$.

Define $l_\beta, b_\beta, r_\beta$ in the analogous way, and note that $a \in l_\alpha \cap l_\beta$ and $b \in r_\alpha \cap r_\beta$ by hypothesis. We will think of each point in $[a, b]$ as being colored by the subsets they are in – each point gets an α color and a β color.

With this terminology, we claim that the following lemma immediately proves Proposition 5.1.

Lemma 5.2. *There is a closed interval $I \subset [a, b]$ whose interior is a component of $[a, b] \setminus (m_\alpha \cup m_\beta)$ such that*

- (1) I has one endpoint in m_α and one endpoint in m_β , and

(2) for each t in the interior of I , its α color is different from its β color.

Proposition 5.1 follows from the fact that if t gets a different α color and β color (the colors being either l, r, b) then \mathbb{T}_α and \mathbb{T}_β lie in different components of $M \setminus \Sigma_t$.

We now turn to finding the desired subinterval of $[a, b]$. Let us begin by making a few observations. First, m_α and m_β are closed and disjoint, so components of one cannot accumulate onto a component of the other. Hence, if we are at a component of (say) m_α it makes sense to talk about the component of m_β immediately after or before it in the time interval. Similarly, any monotone (with respect to the order on $[a, b]$) sequence of components of $m_\alpha \cup m_\beta$ alternating between components of m_α and m_β must terminate after finitely many steps. Second, in what follows we only consider components of $m_\alpha \cup m_\beta$ which have nonempty interior. We call such components thick. Note that by continuity of the sweepout, the α color can change only across a thick m_α component. More accurately, if two points in $[a, b]$ are not separated by a thick component of m_α , then they have the same α color. Finally, call an interval in $[a, b] \setminus (m_\alpha \cup m_\beta)$ switching if it has one endpoint in m_α and one endpoint in m_β . It is clear that a switching interval must exist: otherwise we can construct a sequence of nested intervals $I_0 \supset I_1 \supset \dots$ each with one endpoint in m_α and one endpoint in m_β such that $\cap I_k = \{x\}$. Since we would necessarily have that $x \in m_\alpha \cap m_\beta$, this is a contradiction.

Proof of Lemma 5.2. By an m_α component we mean a thick connected component of m_α , and the same goes for β . Let us look at some time in the interval $[a, b]$ that we see an m_α component followed by an m_β component or vice versa. This exists since there is a switching interval. Up to reversing the time parameter, we assume that the m_α component comes first. Let m_{α_0} denote that m_α component, and m_{β_0} the m_β component.

Let m_{α_1} be the first m_α component after m_{β_0} and let blue be the α -color of the interval in between m_{α_0} and m_{α_1} . We assume that the β -color to the left of m_{β_0} is also blue, otherwise the interval we are looking for is $[m_{\alpha_0}, m_{\beta_0}]$. (This notation means the interval between the endpoints of these components not containing their interior.)

Let m_{β_1} be the last m_β component before m_{α_1} and let m_{β_2} be the next m_β component after m_{β_1} (and thus the first one after m_{α_1} , so it is well-defined). If the interval between m_{β_1} and m_{β_2} is not blue, say it is green, then the interval we are looking for is the one between m_{β_1} and m_{α_1} . It is colored α -blue but β -green. Thus, we assume for contradiction that the interval between m_{β_1} and m_{β_2} is in fact blue (see Figure 2).

We will in general argue that if we have not found our desired time interval, and we see two m_α components, m'_α and m''_α with the color of the interval $[m'_\alpha, m''_\alpha]$ colored blue, then the color to the left of the m_β component directly after m''_α is also blue.

Now let m_{α_2} and m_{α_3} be the α component directly before and after m_{β_2} . If $[m_{\alpha_2}, m_{\alpha_3}]$ is not blue, say it is green, then the interval we are looking for is $[m_{\alpha_2}, m_{\beta_2}]$; it is α -green but β -blue (see Figure 3). Assume that it is blue, and let m_{β_3} and m_{β_4} be the m_β components directly before and after m_{α_3} , respectively. If $[m_{\beta_3}, m_{\beta_4}]$ is not blue, then we are again done. The interval we want is $[m_{\beta_3}, m_{\alpha_3}]$. Thus we can assume that both are blue, and we have succeeded in demonstrating the set up outlined in the previous

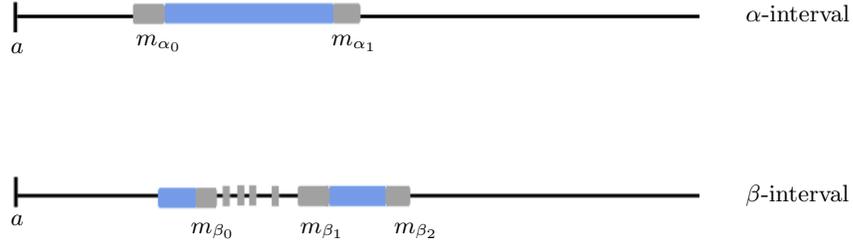


FIGURE 2. We represent the time interval $[a, b]$ as two separate intervals to keep track of relevant α and β information.

paragraph since the color to the left of m_{β_4} is blue (see Figure 4). This perpetuates to the right.

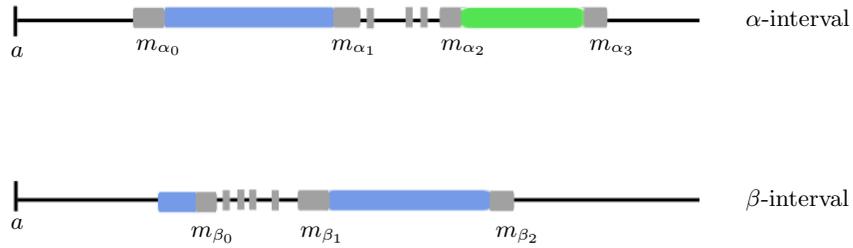


FIGURE 3.

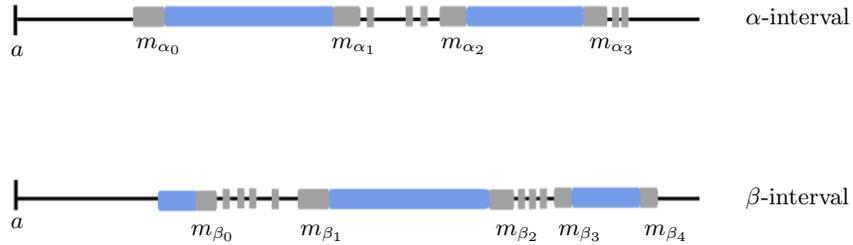


FIGURE 4.

Recall that the last α interval and the last β interval must be the same color (these are the intervals ending in b). Suppose that the last α interval and the last β interval are not blue, say they are red. We now show that, in this case, we will always find the interval we are looking for.

Let m_{β_n} be the last m_β component before b resulting from the argument outlined above; m_{β_n} exists otherwise there would be an infinite alternating sequence of m_α and m_β components, which we said above cannot happen. Let m_{α_n} be the m_α component directly before m_{β_n} , and note that the color to the left of m_{β_n} is blue. We claim that m_{β_n} is the last m_β component, or m_{α_n} is the last m_α component. Otherwise the argument outlined above applies again, and we either find our desired interval or we have obtained a contradiction to the choice of m_{β_n} .

Suppose that m_{β_n} is the last m_β component of the interval $[a, b]$ and that m_{α_n} is not the last m_α component. Let $m_{\alpha_{n+1}}$ be the m_α component directly after m_{β_n} and consider the interval $[m_{\alpha_n}, m_{\alpha_{n+1}}]$. If it is blue, the interval we are looking for is $[m_{\beta_n}, m_{\alpha_{n+1}}]$ (see Figure 5) and if it is not, then the interval we are looking for is $[m_{\alpha_n}, m_{\beta_n}]$.

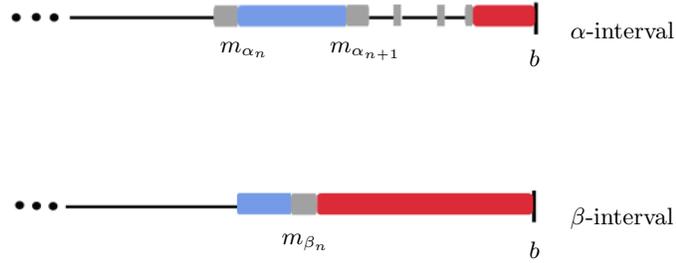


FIGURE 5.

Next suppose that m_{α_n} is the last m_α component. Then by our hypothesis, the color to the right of m_{α_n} is red. Thus, the interval we are looking for is $[m_{\alpha_n}, m_{\beta_n}]$ which is α -red but β -blue (see Figure 6).

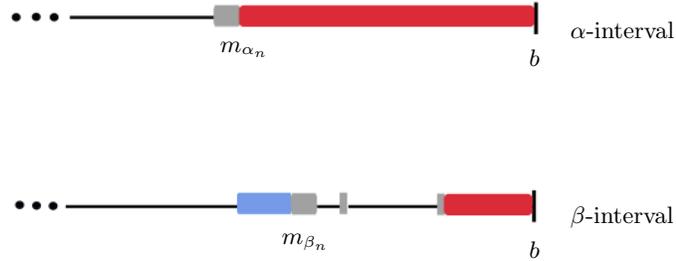


FIGURE 6.

So we assume for now that the last α - and β - intervals are blue. A similar perpetuation argument applies as we move left in the interval $[a, b]$ so that either we find our desired interval, or the first α - and β - intervals, which begin at a , are also blue. This contradicts the fact that the sweepout starts with Σ_a to the left of α and β , and ends with Σ_b to the right of α and β . \square

Another method for proving Lemma 5.1 was suggested to the authors by Dave Futer. In short, one uses a result of Otal [Ota95, Ota03], which guarantees that short curves in M are unlinked, to topologically order the short γ_i and the 1-Lipschitz surfaces they meet. Rather than attempt to make effective this technique, we chose instead to employ the direct combinatorial argument found here.

6. FINISHING THE PROOF OF THEOREM 4.1

Recall that $A_2(|\chi(S)|)$ is obtained by setting $\eta = \epsilon_S$ in (21), giving us the upper bound in Theorem 4.1

$$d_M^{\epsilon_S}(\alpha, \beta) \leq \frac{2\pi|\chi(S)|}{\epsilon_S} \cdot d_{\mathcal{C}(S)}(\alpha, \beta) = A_2(|\chi(S)|) \cdot d_{\mathcal{C}(S)}(\alpha, \beta).$$

For the lower bound, suppose that α and β are given and let p be a geodesic from α to β in M . Let \mathcal{S} be the set of curves γ in S such that p meets $\mathbb{T}_\gamma = \mathbb{T}_\gamma(\epsilon_S)$ in M and index $\mathcal{S} = \{\gamma_i\}_{i=1}^N$ according to the order in which these tubes are met by p . (Set $\gamma_0 = \alpha$ and $\gamma_{N+1} = \beta$.)

Let p_i be the subarc of p between the last point of $p \cap \mathbb{T}_{\gamma_i}$ and the first point of $p \cap \mathbb{T}_{\gamma_{i+1}}$. Then these subarcs are disjoint and $p \cap M_{[\epsilon_S, \infty)} = \bigcup_i p_i$. Additionally, let f_t , $t \in [a, b]$ be a sweepout of 1-Lipschitz maps such that Σ_a lies to the left of α, β, γ_i and Σ_b lies to the right of α, β, γ_i for $1 \leq i \leq N$. Such a sweepout always exists by Theorem 2.1.

Our analysis breaks into two cases, depending on whether γ_i and γ_{i+1} intersect as curves on S . If not, then $d_{\mathcal{C}(S)}(\gamma_i, \gamma_{i+1}) \leq 1$ and we can only say that $\ell_M(p_i) \geq 4$ by Remark 3.4.

Now suppose that γ_i and γ_{i+1} are such that $d_{\mathcal{C}(S)}(\gamma_i, \gamma_{i+1}) \geq 2$. In this case, apply Proposition 5.1 to obtain (up to reversing the time parameter) a sub-sweepout $(f_t: X_t = (S, g_t) \rightarrow M)_{t \in [a_i, b_i]}$ with the following properties:

- (1) $\Sigma_{a_i} \cap \mathbb{T}_{\gamma_i} \neq \emptyset$ and $\Sigma_{b_i} \cap \mathbb{T}_{\gamma_{i+1}} \neq \emptyset$,
- (2) $\Sigma_t \cap p_i \neq \emptyset$ for all $t \in [a_i, b_i]$,
- (3) p_i is contained in $M_{[\epsilon_S, \infty)}$.

Note that we are using that Σ_t cannot meet both \mathbb{T}_{γ_i} and $\mathbb{T}_{\gamma_{i+1}}$ and that since Σ_t separates γ_i from γ_{i+1} , it must meet p_i . Hence, we may apply Lemma 4.4 to conclude that

$$d_{\mathcal{C}(S)}(\gamma_i, \gamma_{i+1}) \leq A_1(|\chi(S)|) \cdot \ell_M(p_i),$$

and thus,

$$\begin{aligned} d_{\mathcal{C}(S)}(\alpha, \beta) &\leq \sum_i d_{\mathcal{C}(S)}(\gamma_i, \gamma_{i+1}) \\ &\leq A_1(|\chi(S)|) \cdot \sum_i \ell_M(p_i) \\ &\leq A_1(|\chi(S)|) \cdot d_M^{\epsilon_S}(\alpha, \beta) \end{aligned}$$

as wanted. This completes the proof of Theorem 4.1.

7. COVERS AND THE CURVE COMPLEX

In this section, we follow Tang [Tan12] and apply Theorem 4.1 to analyze maps between curve graphs induced by covering maps of surfaces.

If $p : \tilde{S} \rightarrow S$ is a covering map, there is a coarsely well-defined map $p^* : \mathcal{C}(S) \rightarrow \mathcal{C}(\tilde{S})$ induced by p ; given an essential simple closed curve γ on S , define $p^*(\gamma)$ to be the full pre-image $p^{-1}(\gamma) \subseteq \tilde{S}$. This will be a multi-curve on \tilde{S} corresponding to a complete subgraph of $\mathcal{C}(\tilde{S})$. Given α and β vertices of $\mathcal{C}(S)$, we can then define the distance in $\mathcal{C}(\tilde{S})$ between $p^*(\alpha)$ and $p^*(\beta)$ to be the diameter of their union:

$$d_{\mathcal{C}(\tilde{S})}(p^*(\alpha), p^*(\beta)) := \text{diam}(p^*(\alpha) \cup p^*(\beta)).$$

With this setup, we prove the following:

Theorem 7.1. *Let $p : \tilde{S} \rightarrow S$ be a finite covering map between non-sporadic surfaces \tilde{S}, S . Then for any α, β distinct essential simple closed curves on S ,*

$$\frac{d_{\mathcal{C}(S)}(\alpha, \beta)}{\deg(p) \cdot A_3(|\chi(S)|)} \leq d_{\mathcal{C}(\tilde{S})}(p^*(\alpha), p^*(\beta)) \leq d_{\mathcal{C}(S)}(\alpha, \beta),$$

where A_3 is the polynomial $A_3(x) = 80 e^{54} \pi x^{13}$ when S is closed.

When S has punctures,

$$\frac{d_{\mathcal{C}(S)}(\alpha, \beta)}{\deg(p)^4 \cdot A_3(|\chi(S)|)} \leq d_{\mathcal{C}(\tilde{S})}(p^*(\alpha), p^*(\beta)) \leq d_{\mathcal{C}(S)}(\alpha, \beta),$$

where A_3 is the polynomial $A_3(x) = 416 e^{114} \pi^{37} x^{37}$.

Recall that $A_3(x) = A_1(x) \cdot A_2(x)$ for A_1, A_2 as in Equation 1

Proof. Given γ_1, γ_2 disjoint essential simple closed curves on S , $p^*(\gamma_1)$ will be disjoint from $p^*(\gamma_2)$. This proves the upper bound on $d_{\mathcal{C}(\tilde{S})}(p^*(\alpha), p^*(\beta))$ in Theorem 7.1.

For the lower bound, we choose a hyperbolic manifold $M \cong S \times \mathbb{R}$ so that $\ell_M(\alpha)$ and $\ell_M(\beta)$ are at most ϵ_S . Constructing such a manifold is standard, see [Kap10, Chapter 8]. Thus, the first inequality of Theorem 4.1 implies that

$$(26) \quad d_{\mathcal{C}(S)}(\alpha, \beta) \leq A_1(|\chi(S)|) \cdot d_M^{\epsilon_S}(\alpha, \beta).$$

The covering map p gives rise to a covering of 3-manifolds between p^*M and M . Let $p^*\alpha, p^*\beta$ also denote the geodesic representatives in p^*M of the lifts $p^{-1}(\alpha), p^{-1}(\beta)$, respectively, and let γ be a path in p^*M from any component of $p^*\alpha$ to any component of $p^*\beta$. Then γ maps to a path in M from α to β .

Since a covering map is distance non-increasing and sends the thin part into the thin part, it follows that

$$d_M^{\epsilon_S}(\alpha, \beta) \leq d_{p^*M}^{\epsilon_S}(p^*\alpha, p^*\beta),$$

where the left hand side is defined to be the minimum electric distance between a tube about any component of $p^*\alpha$ and a tube about any component of $p^*\beta$. Combining this observation with (26) yields

$$(27) \quad d_{\mathcal{C}(S)}(\alpha, \beta) \leq A_1(|\chi(S)|) \cdot d_{p^*M}^{\epsilon_S}(p^*\alpha, p^*\beta).$$

When S is closed, the upper bound on η in Proposition 4.2 is simply ϵ_3 , the Margulis constant. Thus, applying Proposition 4.2 to the right hand side of (27) with $\eta = \epsilon_S$ we obtain

$$d_{\mathcal{C}(S)}(\alpha, \beta) \leq A_1(|\chi(S)|) \cdot \mathcal{A}(|\chi(\tilde{S})|, \epsilon_S) \cdot d_{\mathcal{C}(\tilde{S})}(p^*\alpha, p^*\beta).$$

Recall that $\mathcal{A}(|\chi(\tilde{S})|, \epsilon_S) = \deg(p) \cdot \mathcal{A}(|\chi(S)|, \epsilon_S) \leq \deg(p) \cdot A_2(|\chi(S)|)$ by (22), which yields the lower bound in the closed case.

In order to apply Proposition 4.2 to the right hand side of (27) when S is not closed, it is necessary to choose

$$\eta < \frac{\epsilon_3}{e^{6(\pi|\chi(\tilde{S})|)^3}},$$

and when the degree of $p: \tilde{S} \rightarrow S$ is large we note that ϵ_S is not small enough to satisfy the above inequality. Hence, we set $\eta = \min\{\epsilon_S, \frac{\epsilon_3}{e^{6(\pi|\chi(\tilde{S})|)^3}}\}$. We first note that since $\eta \leq \epsilon_S$,

$$d_{p^*M}^{\epsilon_S}(p^*\alpha, p^*\beta) \leq d_{p^*M}^{\eta}(p^*\alpha, p^*\beta).$$

Now starting with (27) and applying Proposition 4.2 with this η yields

$$d_{\mathcal{C}(S)}(\alpha, \beta) \leq A_1(|\chi(S)|) \cdot \mathcal{A}(|\chi(\tilde{S})|, \eta) \cdot d_{\mathcal{C}(\tilde{S})}(p^*\alpha, p^*\beta),$$

where

$$\mathcal{A}(|\chi(\tilde{S})|, \eta) \leq \max\{\deg(p) \cdot A_2(|\chi(S)|), \deg(p)^4 \cdot 4e^9 \pi^4 (|\chi(S)|)^4\},$$

which in either case is less than $\deg(p)^4 \cdot A_2(|\chi(S)|)$. This completes the proof. \square

Corollary 1.1 is immediate from Theorem 7.1 after noting that if

$$d_{\mathcal{C}(\tilde{S})}(p^*\alpha, p^*\beta) \geq 4,$$

then every lift of α intersects every lift of β .

Remark 7.2. We conclude this section by showing that the linear dependence on $\deg(p)$ in Theorem 7.1 is sharp in the closed case.

Let $S = S_g$ ($g \geq 2$) be a fixed closed surface with curves α and β such that α is nonseparating, β is separating, and α and β fill S . Set $I = i(\alpha, \beta)$ and let \tilde{S}_n be the n -fold cyclic cover of S built as follows: Take n copies of $X = S \setminus \alpha$, $X_0 \dots X_{n-1}$ and glue them cyclically along their boundaries. That is, if we let $\partial X = \alpha^l \cup \alpha^r$, then we glue α_i^r to α_{i+1}^l for each $i \bmod n$. Rename the resulting curve $\alpha_i^r = \alpha_{i+1}^l$ by $\tilde{\alpha}_i$. These are the preimages of α in \tilde{S}_n .

We note that \tilde{S}_n is the cover corresponding to the kernel of the homomorphism $\phi: \pi_1(S) \rightarrow \mathbb{Z}/n\mathbb{Z}$ taking a loop to its algebraic intersection number with α , mod n . Hence, β has n lifts to the cover \tilde{S}_n and any such lift intersects no more than I of the $\tilde{\alpha}_i$. In particular, each lift of β lives in $X_i \cup X_{i+1} \cup \dots \cup X_{i+I}$ for some $0 \leq i \leq n$.

Now set $f = \tau_\beta^{-1} \tau_\alpha$, which is pseudo-Anosov by Thurston's criterion ([Thu88]). Hence, there is a $\kappa > 0$ (depending only on S) such that $d_{\mathcal{C}(S)}(\alpha, f^j(\alpha)) \geq \kappa j$ ([MM99], [GT11]). Since both τ_α and τ_β fix the homology class of α (recall that β is a separating curve), so does f . Thus, f fixes the kernel of ϕ and hence lifts to a map $\tilde{f}: \tilde{S}_n \rightarrow \tilde{S}_n$.

Indeed, if we denote by $\tilde{\alpha}$ and $\tilde{\beta}$ the full preimage of α and β , one such lift is $\tilde{f} = \tau_{\tilde{\beta}}^{-1} \tau_{\tilde{\alpha}}$, a composition of multitwists.

But then, we must have

$$\tilde{f}(\tilde{\alpha}_1) = \tau_{\tilde{\beta}}^{-1}(\tilde{\alpha}_1) = \tau_{\tilde{\beta}}^{-1}(\tilde{\alpha}_1),$$

where $\bar{\beta}$ is a multicurve consisting of components of $\tilde{\beta}$ that meet $\tilde{\alpha}_1$. Hence, from our observation above, $\bar{\beta}$, and therefore $\tilde{f}(\tilde{\alpha}_1)$, is supported in $X_{-I} \cup X_{-I+1} \cup \dots \cup X_I$. Therefore, so long as $n \geq 2I + 1$, we have that $d_{\mathcal{C}(\tilde{S}_n)}(\tilde{\alpha}_1, \tilde{f}(\tilde{\alpha}_1)) \leq 2$.

This construction can be iterated by choosing $n \geq 2jI + 1$, and considering $\tilde{f}^j(\tilde{\alpha}_1)$. This curve is contained in $Y_j := X_{-jI} \cup X_{-jI+1} \cup \dots \cup X_{jI}$, which is a proper subsurface of \tilde{S}_n . Indeed the Euler characteristic of Y_j is $2j\chi(S)$, which, in absolute value, is strictly less than $|\chi(\tilde{S}_n)|$ under the assumption that $n \geq 2jI + 1$. Hence, $d_{\mathcal{C}(\tilde{S}_n)}(\tilde{\alpha}_1, \tilde{f}^j(\tilde{\alpha}_1)) \leq 2$ for $n \geq 2jI + 1$.

Now since $\tilde{f}^j(\tilde{\alpha}_1) \subset \tilde{f}^j(\tilde{\alpha}) = \widetilde{f^j(\alpha)}$, we can set $\gamma_j = f^j(\alpha)$ to see that we have produced curves α and γ_j on S that have distance at least κj and a degree $n = 2jI + 1$ cover \tilde{S}_n such that $d_{\mathcal{C}(\tilde{S}_n)}(\tilde{\alpha}, \tilde{\gamma}_j) \leq 2$.

8. APPLICATION TO QUANTIFIED VIRTUAL SPECIALNESS

In this section we give an application of Theorem 7.1 to dual cube complexes for collections of curves on closed surfaces and their special covers.

8.1. Dual cube complexes and Sageev's construction. Given a finite and filling collection Γ of simple closed curves on a closed surface S , Sageev's construction [Sag95] gives rise to a dual CAT(0) cube complex $\tilde{\mathcal{C}}_\Gamma$, on which $\pi_1 S$ acts freely, properly discontinuously, and cocompactly. The quotient of $\tilde{\mathcal{C}}_\Gamma$ by this action is a non-positively curved cube complex \mathcal{C}_Γ , which can be thought of as a cubulation of the surface S since $\pi_1 S \cong \pi_1 \mathcal{C}_\Gamma$.

The construction of $\tilde{\mathcal{C}}_\Gamma$ roughly goes as follows. In the language of Wise [Wis00], the full preimage $\tilde{\Gamma}$ of Γ in the universal cover \tilde{S} of S is a union of *elevations*, which each split \tilde{S} into two half-spaces. A *labelling* of $\tilde{\Gamma}$ is a choice of half-space for each elevation in $\tilde{\Gamma}$, and the admissible labelings form the vertex set for $\tilde{\mathcal{C}}_\Gamma$. (For more details on admissible labellings see [Sag95].) Two labellings are joined by an edge when they differ on the choice of a half-space for exactly one elevation. The unique CAT(0) cube complex defined by this 1-skeleton is $\tilde{\mathcal{C}}_\Gamma$, and there is an intersection preserving identification of the curves in the system Γ with the hyperplanes of $\tilde{\mathcal{C}}_\Gamma$. The action of $\pi_1 S$ on \tilde{S} permutes the elevations, inducing an isometry of $\tilde{\mathcal{C}}_\Gamma$. We note that this construction of cube complexes works in a far more general setting. We summarize Sageev's construction with the following theorem:

Theorem 8.1 (Sageev). *Suppose Γ is a finite, filling collection of curves on S . Then the dual cube complex \mathfrak{C}_Γ is $CAT(0)$ and there is an intersection preserving identification of the curves in Γ with the hyperplanes of $\widetilde{\mathfrak{C}}_\Gamma$. The group $\pi_1 S$ acts freely, properly discontinuously, and cocompactly on $\widetilde{\mathfrak{C}}_\Gamma$.*

8.2. Virtual specialness. It is well known that there exists a finite cover $\overline{\mathfrak{C}}_\Gamma$ of \mathfrak{C}_Γ which is special [HW08]. Here $\overline{\mathfrak{C}}_\Gamma$ is called special because its hyperplanes avoid three key pathologies (self-intersection, direct self-osculation, and inter-osculation). There is an algebraic characterization of specialness [HW08]: that $\pi_1 \overline{\mathfrak{C}}_\Gamma$ embeds in a particular right-angled Artin group (RAAG). The defining graph of that RAAG is the *crossing graph* of $\overline{\mathfrak{C}}_\Gamma$. The crossing graph of $\overline{\mathfrak{C}}_\Gamma$ is the simplicial graph whose vertices are hyperplanes of $\overline{\mathfrak{C}}_\Gamma$ and whose edges connect distinct, intersecting hyperplanes. Thus, Theorem 8.1 implies that the specialness of a cube complex dual to a collection of curves on a surface is determined by the intersection pattern of the underlying curves.

Suppose that Γ consists of two simple closed curves, α and β , with nontrivial geometric self-intersection number and that together fill the surface S . Consider a finite-degree covering map $p : \widetilde{S} \rightarrow S$, and as in Section 7 let $p^* : \mathcal{C}(S) \rightarrow \mathcal{C}(\widetilde{S})$ be the induced map between their curve complexes.

There is also an induced covering map on the level of dual cube complexes $p_* : \mathfrak{C}_{\Gamma'} \rightarrow \mathfrak{C}_\Gamma$ where $\mathfrak{C}_{\Gamma'}$ is the dual complex to the curve system $\Gamma' = p^{-1}(\alpha) \cup p^{-1}(\beta)$ on \widetilde{S} and is also the cover of \mathfrak{C}_Γ corresponding to the subgroup $\pi_1 \widetilde{S} < \pi_1 S \cong \pi_1 \mathfrak{C}_\Gamma$. We record the following lemma as an obstruction to the specialness of $\mathfrak{C}_{\Gamma'}$.

Lemma 8.2. *Suppose that α and β are two simple closed curves that nontrivially intersect and together fill a surface S , and that $p : \widetilde{S} \rightarrow S$ is a finite degree covering map. If every lift of α to \widetilde{S} intersects every lift of β to \widetilde{S} , then the cover $\mathfrak{C}_{\Gamma'}$ of \mathfrak{C}_Γ corresponding to $\pi_1 \widetilde{S} < \pi_1 S \cong \pi_1 \mathfrak{C}_\Gamma$ cannot be special.*

Proof. If every lift of α intersects every lift of β , then the underlying graph for the right-angled Artin group A in which $\mathfrak{C}_{\Gamma'}$ should embed is the join of two sets of non-adjacent vertices. Thus, $A = F_n \times F_m$ is the product of two free groups. However, $\pi_1 \mathfrak{C}_{\Gamma'}$ is a surface group, which cannot embed in the product of two free groups [BR84]. \square

Note that if $d_{\mathcal{C}(\widetilde{S})}(p^*\alpha, p^*\beta) \geq 4$, then every lift of α intersects every lift of β . Thus, Theorem 7.1 gives us the following:

Theorem 8.3. *Suppose that α and β are two simple closed curves that together fill a closed surface S . Let $\deg \mathfrak{C}_\Gamma$ be the minimal degree of a special cover of the dual cube complex \mathfrak{C}_Γ to the curve system $\Gamma = \alpha \cup \beta$. Then*

$$\frac{d_{\mathcal{C}(S)}(\alpha, \beta)}{C(S)} \leq \deg \mathfrak{C}_\Gamma,$$

where $C(S)$ is a polynomial in $|\chi(S)|$ of degree 13.

Proof. Suppose that $p : \widetilde{S} \rightarrow S$ is a finite degree cover of the surface S and that $p_* : \mathfrak{C}_{\Gamma'} \rightarrow \mathfrak{C}_\Gamma$ is the induced cover of cube complexes. Additionally, assume that $\mathfrak{C}_{\Gamma'}$ is

special. Theorem 7.1 gives us that

$$\frac{d_{\mathcal{C}(S)}(\alpha, \beta)}{\deg(p) \cdot A_1(|\chi(S)|) \cdot A_2(|\chi(S)|)} \leq d_{\mathcal{C}(\tilde{S})}(p^*(\alpha), p^*(\beta)).$$

Given that S is closed, $A_1(|\chi(S)|)$ is a polynomial of degree 10 in $|\chi(S)|$ and $A_2(|\chi(S)|)$ is a polynomial of degree 3 in $|\chi(S)|$ (see Equation (1)). Lemma 8.2 shows that $\mathfrak{C}_{\Gamma'}$ cannot be special unless $d_{\mathcal{C}(\tilde{S})}(p^*(\alpha), p^*(\beta)) \leq 3$. Combining these results and solving for $\deg(p)$ gives

$$\frac{d_{\mathcal{C}(S)}(\alpha, \beta)}{C(S)} \leq \deg(p),$$

where $C(S) = 3 \cdot A_1(|\chi(S)|) \cdot A_2(|\chi(S)|)$. \square

9. THE CIRCUMFERENCE OF A FIBERED MANIFOLD

The methods developed above generalize to effectively relate the electric circumference of a fibered manifold to the curve graph translation length of its monodromy. The noneffective version of this relation has proven useful, for example, in work of Biringer–Souto on the rank of the fundamental group of such manifolds [BS15]. As in the previous section, we restrict to the case where S is closed.

Let $\phi \in \text{Mod}(S)$ be pseudo-Anosov and denote its mapping torus by M_ϕ . For $0 < \delta < \epsilon_3$, denote the hyperbolic *circumference* and δ -*electric circumference* of M_ϕ by $\text{circ}(M_\phi)$ and $\text{circ}_\delta(M_\phi)$, respectively. That is, $\text{circ}(M_\phi)$ is the minimum geodesic length of a loop in M which is not in the kernel of the associated map $\pi_1(M_\phi) \rightarrow \mathbb{Z}$, and similarly $\text{circ}_\delta(M_\phi)$ is the minimum δ -electric length of a loop in M which is not in the kernel of the map. Let $\ell_S(\phi)$ be the *stable translation length* of ϕ in $\mathcal{C}(S)$,

$$\ell_S(\phi) = \lim_{n \rightarrow \infty} \frac{d_{\mathcal{C}(S)}(\alpha, \phi^n \alpha)}{n}.$$

Theorem 9.1. *If $\phi: S \rightarrow S$ is a pseudo-Anosov homeomorphism of a closed surface S , then*

$$\frac{1}{A_1(|\chi(S)|)} \cdot \ell_S(\phi) \leq \text{circ}_{\epsilon_S}(M_\phi) \leq A_2(|\chi(S)|) \cdot (\ell_S(\phi) + 2),$$

where the polynomials A_1 and A_2 are as in Equation (1).

Our argument follows the outline from Brock in [Bro03b]. There, Brock extends his theorem on volumes of quasi-fuchsian manifolds to volumes of hyperbolic mapping tori. Similarly, we deduce Theorem 9.1 from the tools we used to prove Theorem 4.1.

Proof. Let $M = M_\phi$ and let N be the infinite cyclic cover of M corresponding to S . The inclusion $\iota: S \rightarrow M$ lifts to a marking $\tilde{\iota}: S \rightarrow N$. Let Φ denote the (isometric) deck transformation of N such that $\tilde{\iota} \circ \phi$ is homotopic to $\Phi \circ \tilde{\iota}$. Following the proof of [Bro03b, Theorem 1.1] there is a 1-Lipschitz map $f: X = (S, g) \rightarrow N$ homotopic to $\tilde{\iota}$ and a 1-Lipschitz sweepout $f_t: X_t = (S, g_t) \rightarrow N$ from $f_0 = f$ to $f_1 = \Phi \circ f \circ \phi^{-1}$. (The hyperbolic structure X_1 on S agrees with that of X under ϕ , up to isotopy.) As in Theorem 2.1, this sweepout has the property that there is some curve α in S such

that the geodesic representative of α in N is in the image of f . Hence, the geodesic representative of $\phi(\alpha)$ lies in the image of f_1 .

Let $H: S \times [0, 1] \rightarrow N$ be the homotopy given by $H(x, t) = f_t(x)$ and set Σ_t to be the image of f_t . Finally, fix an embedding $h: S \rightarrow N$ homotopic to $\tilde{\iota}$ which lies to the left of the image of H . Note that there is some $n_0 \geq 1$ such that $\Phi^{n_0}h(S)$ lies to the right of the image of H .

For $n > 0$, define a function $\mathfrak{s}_n: [n, n+1] \rightarrow [0, 1]$ by $\mathfrak{s}_n(x) = x - n$ and let $H^n: S \times [0, n] \rightarrow N$ denote the homotopy formed by gluing together

$$H, \quad \Phi \circ H \circ (\phi^{-1} \times \mathfrak{s}_1), \quad \dots, \quad \Phi^{n-1} \circ H \circ (\phi^{-(n-1)} \times \mathfrak{s}_{n-1})$$

to form a sweepout from f to $\Phi^n \circ f \circ \phi^{-n}$. (Note that H^n is indeed continuous since the functions agree on their overlap.) Also, extend the definition of Σ_t for $t \in [0, n]$ to be the image of $H^n(\cdot, t)$, so that in particular $\Sigma_n = \Phi^n(\Sigma_0)$. Note that the image of H^n is contained in the compact region between $h(S)$ and $\Phi^{n+n_0}(h(S))$, which we name C_n .

To prove the first inequality, let $\rho: [0, l] \rightarrow M$ be the shortest loop in M which realizes $\text{circ}_{\epsilon_S}(M)$. Note that ρ cannot be ϵ_S -short itself. Otherwise, since the image of f under the covering $N \rightarrow M$ necessarily meets ρ , the argument in Lemma 3.3 would produce an essential loop in S which is mapped into the Margulis tube about ρ . This would imply that ρ represents an element of the kernel of $\pi_1(M_\phi) \rightarrow \mathbb{Z}$, a contradiction.

Denote by $\tilde{\rho}$ the preimage of ρ in N (joining the ends of N) and let $\tilde{\rho}_n = \tilde{\rho} \cap C_n$. Since C_n is the union of $n + n_0$ fundamental domains of Φ ,

$$\ell_N^{\epsilon_S}(\tilde{\rho}_n) = (n + n_0) \cdot \ell_M^{\epsilon_S}(\rho).$$

By choice of $h(S)$ and n_0 , each Σ_t separates the boundary components of C_n (which are $h(S)$ and $\Phi^{n+n_0}h(S)$) for $t \in [0, n]$. Hence, each such Σ_t intersects $\tilde{\rho}_n$. Now pick any curve β that is L_S -short on X and observe that $\phi^n(\beta)$ is L_S -short on $X_n = \phi^n X$. Then, using Proposition 5.1 and Lemma 4.4 as in the proof of Theorem 4.1, we conclude that

$$\begin{aligned} d_{C(S)}(\beta, \phi^n(\beta)) &\leq A_1(|\chi(S)|) \cdot \ell_{\epsilon_S}(\tilde{\rho}_n) \\ &\leq A_1(|\chi(S)|) \cdot (n + n_0) \cdot \ell_N^{\epsilon_S}(\rho). \end{aligned}$$

Hence, diving both sides by n and taking $n \rightarrow \infty$ shows that

$$\ell_S(\phi) \leq A_1(|\chi(S)|) \cdot \ell_M^{\epsilon_S}(\rho),$$

proving the first inequality.

For the second inequality, let ξ_n be the shortest electric geodesic in N joining the geodesic representatives of α and $\phi^n(\alpha)$, where α is as above. Apply Proposition 4.2 to these curves to obtain

$$\ell_N^{\epsilon_S}(\xi_n) \leq A_2(|\chi(S)|) \cdot d_{C(S)}(\alpha, \phi^n(\alpha)).$$

Alter ξ_n to a new path ω_n as follows: for $0 < j < n$, choose some $x_j \in \xi_n \cap \Sigma_j$, and connect x_j to $\Phi^j \tilde{\rho}(0) \in \Sigma_j$ by a shortest electric path γ_j in Σ_j . For $j = 0$ and $j = n$, define γ_j to be a shortest electric path in Σ_j starting at the initial and terminal points of ξ_n and ending at lifts x_0 and x_n of $\rho(0)$ in Σ_0 and Σ_n , respectively. Then define ω_n to be the path obtained from ξ_n by inserting $\gamma_j * \gamma_j^{-1}$ after x_j for each $0 < j < n$,

and by inserting γ_0^{-1} at the beginning and γ_n at the end. Using the bounded diameter lemma, we have that

$$\begin{aligned} \ell_N^{\epsilon_S}(\omega_n) &\leq \ell_N^{\epsilon_S}(\xi_n) + 2n \cdot \frac{4\pi|\chi(S)|}{\epsilon_S} \\ &\leq \ell_N^{\epsilon_S}(\xi_n) + 2n \cdot A_2(|\chi(S)|). \end{aligned}$$

Let $\omega_n[j-1, j]$ denote the portion of ω_n between $\Phi^{j-1}\tilde{\rho}(0)$ and $\Phi^j\tilde{\rho}(0)$. Since $\omega_n[j-1, j]$ descends to a loop in M which is not in the kernel of $\pi_1(M) \rightarrow \mathbb{Z}$, we have

$$\ell_M^{\epsilon_S}(\rho) \leq \ell_N^{\epsilon_S}(\omega_n[j-1, j]) \quad \forall j,$$

hence

$$\begin{aligned} n \cdot \ell_M^{\epsilon_S}(\rho) &\leq \ell_N^{\epsilon_S}(\xi_n) + 2n \cdot A_2(|\chi(S)|) \\ &\leq A_2(|\chi(S)|) \cdot d_{\mathcal{C}(S)}(\alpha, \phi^n(\alpha)) + 2n \cdot A_2(|\chi(S)|). \end{aligned}$$

Dividing through by n and taking a limit as $n \rightarrow \infty$ produces the second inequality. \square

REFERENCES

- [Ago04] Ian Agol. Tameness of hyperbolic 3-manifolds. *arXiv preprint math/0405568*, 2004.
- [Aou13] Tarik Aougab. Uniform hyperbolicity of the graphs of curves. *Geometry & Topology*, 17(5):2855–2875, 2013.
- [ATW18] Tarik Aougab, Samuel J. Taylor, and Richard C.H. Webb. Effective Masur–Minsky distance formulas and applications to hyperbolic 3-manifolds. August 2018.
- [BB11] Jeffrey Brock and Kenneth Bromberg. Geometric inflexibility and 3-manifolds that fiber over the circle. *Journal of Topology*, 4(1):1–38, 2011.
- [BCM12] Jeffrey F Brock, Richard D Canary, and Yair N Minsky. The classification of kleinian surface groups, II: The ending lamination conjecture. *Annals of Mathematics*, pages 1–149, 2012.
- [Bon86] Francis Bonahon. Bouts des variétés hyperboliques de dimension 3. *Annals of Mathematics*, 124(1):71–158, 1986.
- [Bow11] Brian H Bowditch. The ending lamination theorem. *Preprint Warwick, September*, 2011.
- [Bow14] Brian Bowditch. Uniform hyperbolicity of the curve graphs. *Pacific Journal of Mathematics*, 269(2):269–280, 2014.
- [BR84] Gilbert Baumslag and James E. Roseblade. Subgroups of direct products of free groups. *Journal of the London Mathematical Society*, s2-30(1):44–52, 1984.
- [Bro03a] Jeffrey Brock. The Weil-Petersson metric and volumes of 3-dimensional hyperbolic convex cores. *Journal of the American Mathematical Society*, 16(3):495–535, 2003.
- [Bro03b] Jeffrey F Brock. Weil-Petersson translation distance and volumes of mapping tori. *Comm. Anal. Geom.*, 11:987–999, 2003.
- [BS15] Ian Biringer and Juan Souto. Ranks of mapping tori via the curve complex. *Journal für die reine und angewandte Mathematik (Crelles Journal)*, 2015.
- [Bus10] Peter Buser. *Geometry and spectra of compact Riemann surfaces*. Springer Science & Business Media, 2010.
- [Can93] Richard D Canary. Ends of hyperbolic 3-manifolds. *Journal of the American Mathematical Society*, 6(1):1–35, 1993.
- [Can96] Richard D Canary. A covering theorem for hyperbolic 3-manifolds and its applications. *Topology*, 35(3):751–778, 1996.
- [CG06] Danny Calegari and David Gabai. Shrinkwrapping and the taming of hyperbolic 3-manifolds. *Journal of the American Mathematical Society*, 19(2):385–446, 2006.
- [Chu17] Michelle Chu. Special subgroups of Bianchi groups. *arXiv preprint arXiv:1709.10503*, 2017.

- [DMP18] Jason DeBlois, Nicholas Miller, and Priyam Patel. Effective virtual and residual properties of some arithmetic hyperbolic 3-manifolds. *arXiv preprint arXiv:1806.02360*, 2018.
- [FPS18] David Futer, Jessica S Purcell, and Saul Schleimer. Effective distance between nested margulis tubes. *arXiv preprint arXiv:1801.05342*, 2018.
- [FS14] David Futer and Saul Schleimer. Cusp geometry of fibered 3-manifolds. *American Journal of Mathematics*, 136(2):309–356, 2014.
- [GT11] Vaibhav Gadre and Chia-Yen Tsai. Minimal pseudo-anosov translation lengths on the complex of curves. *Geometry & Topology*, 15(3):1297–1312, 2011.
- [Hem01] John Hempel. 3-manifolds as viewed from the curve complex. *Topology*, 40(3):631–657, 2001.
- [HW08] Frédéric Haglund and Daniel T Wise. Special cube complexes. *Geometric and Functional Analysis*, 17(5):1551–1620, 2008.
- [Kap10] Michael Kapovich. *Hyperbolic Manifolds and Discrete Groups*. Birkhäuser Basel, 2010.
- [Mey87] Robert Meyerhoff. A lower bound for the volume of hyperbolic 3-manifolds. *Canad. J. Math*, 39(5):1038–1056, 1987.
- [Min10] Yair Minsky. The classification of kleinian surface groups, I: Models and bounds. *Annals of Mathematics*, pages 1–107, 2010.
- [MM99] Howard A Masur and Yair N Minsky. Geometry of the complex of curves I: Hyperbolicity. *Inventiones mathematicae*, 138(1):103–149, 1999.
- [MM00] Howard A. Masur and Yair N. Minsky. Geometry of the complex of curves. II. Hierarchical structure. *Geom. Funct. Anal.*, 10(4):902–974, 2000.
- [Ota95] Jean-Pierre Otal. Sur le nouage des géodésiques dans les variétés hyperboliques. *Comptes Rendus de l'Académie des Sciences - Series I - Mathematics*, 320:847–852, 1995.
- [Ota03] Jean-Pierre Otal. *Les géodésiques fermées d'une variété hyperbolique en tant que nœuds*, pages 95–104. London Mathematical Society Lecture Note Series. Cambridge University Press, 2003.
- [RS09] Kasra Rafi and Saul Schleimer. Covers and the curve complex. *Geometry & Topology*, 13(4):2141–2162, 2009.
- [Sag95] Michah Sageev. Ends of group pairs and non-positively curved cube complexes. *Proceedings of the London Mathematical Society*, 3(3):585–617, 1995.
- [Tan12] Robert Tang. The curve complex and covers via hyperbolic 3-manifolds. *Geometriae Dedicata*, 161(1):233–237, 2012.
- [Thu78] William Thurston. The geometry and topology of 3-manifolds. *Lecture notes*, 1978.
- [Thu88] William P Thurston. On the geometry and dynamics of diffeomorphisms of surfaces. *Bulletin of the American Mathematical Society*, 19(2):417–431, 1988.
- [Wis00] Daniel T. Wise. Subgroup separability of graphs of free groups with cyclic edge groups. *Quarterly Journal of Mathematics*, 50(1):107–129, 2000.

DEPARTMENT OF MATHEMATICS, BROWN UNIVERSITY,
E-mail address: `tarik_aougab@brown.edu`

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF CALIFORNIA, SANTA BARBARA,
E-mail address: `patel@math.ucsb.edu`

DEPARTMENT OF MATHEMATICS, TEMPLE UNIVERSITY,
E-mail address: `samuel.taylor@temple.edu`