

AN EFFECTIVE ALGEBRAIC DETECTION OF THE NIELSEN–THURSTON CLASSIFICATION OF MAPPING CLASSES

THOMAS KOBERDA AND JOHANNA MANGAHAS

June 25, 2014

ABSTRACT. In this article, we propose two algorithms for determining the Nielsen–Thurston classification of a mapping class ψ on a surface S . We start with a finite generating set X for the mapping class group and a word ψ in $\langle X \rangle$. We show that if ψ represents a reducible mapping class in $\text{Mod}(S)$ then ψ admits a canonical reduction system whose total length is exponential in the word length of ψ . We use this fact to find the canonical reduction system of ψ . We also prove an effective conjugacy separability result for $\pi_1(S)$ which allows us to lift the action of ψ to a finite cover \tilde{S} of S whose degree depends computably on the word length of ψ , and to use the homology action of ψ on $H_1(\tilde{S}, \mathbb{C})$ to determine the Nielsen–Thurston classification of ψ .

1. INTRODUCTION

In this paper we develop two algorithms for determining the Nielsen–Thurston classification of an element of the mapping class group $\text{Mod}(S)$. Both pivot on our ability to compute an upper bound on the length of the *canonical reduction system* $\sigma(f)$ of a given mapping class f , which we state as our first theorem. We define $\sigma(f)$ in Section 2, noting here that it consists of nonperipheral curves fixed by powers of f , and is empty exactly when f has finite order or is pseudo-Anosov. We clarify what constitutes a *reasonable definition of curve length* just after Proposition 1.4 below.

Theorem 1.1. *Let $X \subset \text{Mod}(S)$ be a finite generating set. There exists computable $C = C(X, S, \ell)$ such that, for any $f \in \text{Mod}(S)$ and $\alpha \in \sigma(f)$, we have $\ell(\alpha) \leq C^{\ell_X(f)}$, where ℓ_X is word length with respect to X and $\ell(\alpha)$ is any reasonable definition of curve length.*

The impetus for Theorem 1.1 came from a challenge to make effective the result in [Kob12] that, given $f \in \text{Mod}(S)$, one may find a finite cover \tilde{S} such that the action on $H_1(\tilde{S}, \mathbb{C})$ induced by lifts of f to \tilde{S} determines the Nielsen–Thurston classification of f . Call such \tilde{S} a *classifying cover* for f . Our first algorithm can be stated as a theorem in the following way:

Theorem 1.2 (Cover Algorithm). *Suppose $f \in \text{Mod}(S)$ is given, and \tilde{S} is its smallest classifying cover. The degree of \tilde{S} over S is bounded above by a computable function of the word length of f .*

Our second algorithm uses Theorem 1.1 to directly determine $\sigma(f)$:

Theorem 1.3 (List Algorithm). *There exists an algorithm, exponential-time when S is punctured, which takes as input a mapping class $f \in \text{Mod}(S)$, presented as a word in X , and outputs $\sigma(f)$. When $\sigma(f)$ is empty the algorithm determines whether f is finite-order or pseudo-Anosov.*

Existence of C in Theorem 1.1 follows from Theorem E in Tao’s solution to the conjugacy problem for $\text{Mod}(S)$, which relies on heavy machinery related to the curve complex [Tao13]. Now that key curve complex constants have been determined (the smallest in [HPW, Web], but see also [Aou13, Bow, CRS]), it may be of interest to try to compute C by this approach. In contrast, we use relatively elementary methods to show:

Proposition 1.4. *Suppose S has $g \geq 0$ genus and $n \geq 0$ punctures. There exists a set X generating $\text{Mod}(S)$ and a filling set of curves $\{\gamma_j\}$ so that, if curve length is defined by*

$$\ell(\alpha) = \sum_j i(\alpha, \gamma_j),$$

where $i(\cdot, \cdot)$ denotes geometric intersection number, then for any $\alpha \in \sigma(f)$, we have

$$\ell(\alpha) \leq A \cdot B^{217\chi(S)^2\ell_X(f)}$$

with $A = 2g + 1$ and $B = 2$ if $n \leq 1$, otherwise with $A = 2g + 2n - 1$ and $B = 3$.

By reasonable definition of curve length we mean a length estimate which is quasi-isometric to $\ell(\alpha)$ as defined above, which we note includes length according to any hyperbolic metric on S or, fixing any base point p and choice of word length on $\pi_1(S, p)$, minimal length of elements in the conjugacy class in $\pi_1(S, p)$ represented by α (for non-peripheral curves only). Theorem 1.1 follows easily from Proposition 1.4, which is proved in Section 3 (see Theorem 3.1).

When S has boundary, the List Algorithm indicated by Theorem 1.3 can be implemented in exponential time, as described in Section 4.2, using compressed representations of curves on S such as intersection with a fixed triangulation, or Dehn–Thurston coordinates. The main points of interest in the results here are: the conceptual simplicity of the List Algorithm, the connection to homology classification afforded by the Cover Algorithm, and the elementary nature of our proof of Theorem 1.1, which relies mainly on the pigeon-hole principle and easy observations about intersection number, given Ivanov’s description of canonical reduction systems of mapping classes.

Both our algorithms are independent of the famous algorithm of Bestvina and Handel, which uses train tracks and Perron–Frobenius theory to prove and implement Nielsen–Thurston classification for $\text{Mod}(S)$ [BH95]. This algorithm outputs a *train track* for the input element, which is more information than $\sigma(f)$. While this algorithm runs quickly in practice, its theoretical complexity is unknown, and discerning it appears to be a hard problem. Recently, Mark Bell has developed an exponential-time classification algorithm based on triangulations of punctured surfaces [Bela], and has implemented it in his program FLIPPER [Belb]. Hamidi-Tehrani and Chen obtained an exponential-time classification of $\text{Mod}(S)$ using its piecewise-linear action on measured train tracks [HTC96]. Polynomial-time classification is available for braid groups [Ca14], which are closely related to mapping class groups of punctured disks and admit the same Nielsen–Thurston trichotomy. Another closely related group, $\text{Out}(F_n)$, admits a similar, albeit more complicated classification (including an analogous train track construction in [BH92]); an algorithmic approach modeled on Theorems 1.1 and 1.3 appears in [CMP].

1.1. Acknowledgements. The authors would like to thank J. Malestein for asking whether there exists an effective method of determining the Nielsen–Thurston classification of a mapping class via examination of the homology of finite covers. They would also like to thank C. Atkinson, D. Canary, M. Clay, W. Dicks, D. Futer, D. Margalit, H. Namazi, A. Pettet, I. Rivin, S. Schleimer, J. Souto, and D. Thurston for productive conversations. The first author was supported by an NSF Graduate Student Research Fellowship for part of the time this research was carried out, and by NSF Grant DMS-1203964. The second author acknowledges support from NSF DMS-1204592 and AMS-Simons. The authors are grateful to the Banff International Research Station, CRM Barcelona, and the Polish Academy of Sciences for hospitality while this work was completed. An abstract of this project will appear in the series, Research Perspectives CRM Barcelona. Finally, the authors thank an anonymous referee for numerous helpful comments.

2. BACKGROUND

We consider orientable surfaces S with g genus, n punctures, and negative, finite Euler characteristic $\chi(S) = 2 - 2g - n$. By a *curve* on S , we mean an isotopy class of homotopically nontrivial,

non-boundary-parallel simple loops. A *filling set* of curves $\{\gamma_i\}$ is such that any curve α has nonzero geometric intersection $i(\alpha, \gamma_j)$ for some j . As in Proposition 1.4, we measure the length of a curve α is by its intersection with a fixed filling set of curves: $\ell(\alpha) = \sum_j \ell_i(\alpha, \gamma_j)$. Note that, if ℓ_h is length with respect to some fixed hyperbolic metric h , then

$$(C_h/2) \cdot \ell(\alpha) \leq \ell_h(\alpha) \leq D_h \cdot \ell(\alpha),$$

where D_h is the maximum diameter of the components of $S - \bigcup_i \gamma_i$, and C_h is the minimal thickness of collar neighborhoods around the $\{\gamma_i\}$ chosen thinly enough so that no more than two collar neighborhoods of distinct elements of $\{\gamma_i\}$ cover any one point on S .

2.1. The Nielsen–Thurston classification of mapping classes. Let $f \in \text{Mod}(S)$. Since f is a homeomorphism of S up to isotopy, we have that f acts on curves on S . We have the following trichotomy:

- (1) The mapping class f has *finite order* as an element of $\text{Mod}(S)$. Thus, some power of f is isotopic to the identity.
- (2) The mapping class f has infinite order in $\text{Mod}(S)$ and there is a curve on S with a finite orbit under the action of f . Such a mapping class is called *reducible*.
- (3) For any curve α , the f -orbit of α is infinite. Such a mapping class is called *pseudo-Anosov*.

The trichotomy above is called the *Nielsen–Thurston classification*. The following are standard results about mapping classes, and proofs can be found in [BLM83], [FM12] and [FLP79], for instance.

Theorem 2.1 (Nielsen). *Let f be a finite order mapping class. Then there exists a hyperbolic metric on S and a representative of f which is an isometry of the corresponding Riemann surface. In particular, the representative has finite order as a homeomorphism of S .*

Theorem 2.2. *Let f be a reducible mapping class. Then there exists a multicurve (i.e. a finite set of distinct, disjoint curves) whose isotopy class is setwise stabilized by f .*

Such a multicurve is called a *reduction system*. We will be interested in a mapping class's *canonical reduction system* $\sigma(f)$, which may be defined as the set of curves common to all maximal reduction systems [BLM83]. It is detailed in [Iva92] that, for some k depending only on the surface, the power f^k sends each component of $S \setminus \sigma(f)$ to itself, inducing on each either a trivial or pseudo-Anosov mapping class. Furthermore, any curve in $\sigma(f)$ that bounds only components on which f^k acts trivially, must be the core curve of an annulus where f^k acts locally as a Dehn twist.

2.2. Artin–Dehn–Humphries–Lickorish generators. For a surface S with genus $g \geq 0$ and punctures $n \geq 0$, we can generate $\text{Mod}(S)$ with $2g + n$ Dehn twists ($2g + 1$ when $n = 0$) and, when $n > 1$, an additional $n - 1$ *half-twists*, along the curves shown in Figure 1. Precise definitions and proof can be found in [FM12]. We remark that a smaller, similar generating set appeared first in [LP01] (Corollary 2.11 there), and in [BD11] one can find an alternate, algebraic proof that these generate, along with a detailed history of the project of generating the mapping class group.

Here we note our labeling convention. For $1 \leq i \leq \max\{2g + n, 2g + 1\}$, we let γ_i be the base curves of the Dehn twists, denoted T_{γ_i} . When $n > 1$, we have $2g + n < i \leq 2g + 2n - 1$ for which γ_i denotes curves corresponding to half-twists in the following way: one component of S cut along γ_i is a disk with two punctures, and the half-twist H_{γ_i} permutes these two punctures and is supported on this disk. In this latter case, we call γ_i the base curve for the half-twist H_{γ_i} .




FIGURE 1. The mapping class group is generated by Dehn twists and half-twists about the curves shown. The curves bounding punctures are the base curves for half-twists. Graphic based on Figures 4.10 and 9.5 in [FM12].

3. THE MAIN TECHNICAL TOOL

Here we show how to compute C in Theorem 1.1. Consider the special case that the generating set X consists of Dehn twists or half-twists about separating curves, an example of which is described in Section 2.2. Let Γ be the set of base curves for the Dehn twists and half-twists in X . Define $M = M(X)$ by

$$M(X) = \max\{i(\gamma_i, \gamma_j) : \gamma_i, \gamma_j \in \Gamma\}.$$

The generators described in Section 2.2 give optimal $M = 1$ when S has no more than one puncture, and $M = 2$ otherwise. Let $\ell(f) = \ell_X(f)$ denote word length in the mapping class group with respect to X . In this section we prove

Theorem 3.1. *For any component α of $\sigma(f)$ and any $\gamma \in \{\gamma_i\}$,*

$$i(\alpha, \gamma) \leq (M + 1)^{217\chi(S)^2\ell(f)}$$

3.1. Proof of Proposition 1.4 using Theorem 3.1. This follows immediately by using the generators described in Section 2.2 for X .

3.2. Proof of Theorem 1.1 using Proposition 1.4. Different generating sets for $\text{Mod}(S)$ change word lengths by a quasi-isometry, and reasonable definitions of curve length also differ from ℓ by a quasi-isometry. The proposition and the quasi-isometry constants allow one to compute C .

3.3. Proof of Theorem 3.1.

Lemma 3.2. *For any γ_i and γ_j , we have $i(f(\gamma_i), \gamma_j) \leq M(M + 1)^{\ell(f)}$.*

Proof. This follows from induction, using the well-known fact that

$$i(T_\gamma(\alpha), \beta) \leq i(\alpha, \gamma)i(\beta, \gamma) + i(\alpha, \beta)$$

and the easily verified fact that

$$i(H_\gamma(\alpha), \beta) \leq \frac{i(\alpha, \gamma)}{2} \cdot \frac{i(\beta, \gamma)}{2} \cdot 2 + i(\alpha, \beta)$$

where H_γ is the half-twist with base curve γ . □

Now suppose α is a component of $\sigma(f)$. There are two possibilities: either α is in the boundary of a component S' of $S \setminus \sigma(f)$ on which some power of f acts as a pseudo-Anosov, or α is the core curve around which some power of f acts as a twist. We deal with each case separately.

Case with pseudo-Anosov component. Suppose α is a component of $\partial S'$. The maximum number of distinct isotopy classes of disjoint arcs on S' is $3|\chi(S')|$, the number of edges in an ideal triangulation of S' with its boundary shrunk to punctures. Because every point of intersection of γ with α has a component of $\gamma \cap S'$ on at least one side, there are at least $i(\alpha, \gamma)/2$ components of $\gamma \cap S'$. Thus at least $i(\alpha, \gamma)/(6|\chi(S')|)$ of these run parallel. Suppose we can find a power P such that, for any essential arc a through S' , we have $f^p(a)$ intersects a for some $p \leq P$ (in particular, f^p also fixes S'). If n arcs of γ run parallel to a , these guarantee n^2 points of intersection between $f^p(\gamma)$ and γ . Putting these estimates together, we have

$$(M+1)^{P\ell(f)+1} \geq M(M+1)^{\ell(f^p)} \geq i(f^p(\gamma), \gamma) \geq [i(\alpha, \gamma)/(6|\chi(S')|)]^2.$$

For this case, it remains to compute P . Let a be the arc through S' of interest. Observe that there are at most $|\chi(S)/\chi(S')|$ copies of S' possibly permuted by f , so for some $p_1 \leq |\chi(S)/\chi(S')|$, f^{p_1} preserves S' . Only $3|\chi(S')|$ distinct arc isotopy classes can be simultaneously disjoint, so for some $p_2 \leq 3|\chi(S')|$, $f^{p_1 p_2}(a)$ intersects a . Thus we can take $P = 3|\chi(S)|$ in the inequality above, and deduce

$$(1) \quad 6|\chi(S)|(M+1)^{(3|\chi(S)|\ell(f)+1)/2} \geq 6|\chi(S')|(M+1)^{(3|\chi(S')|\ell(f)+1)/2} \geq i(\alpha, \gamma).$$

Dehn twist case. For the second case, we assume that α is not the boundary of a pseudo-Anosov component for f , so some power of f acts as Dehn twisting around α . That is to say, let $\sigma' = \sigma(f) \setminus \{\alpha\}$, and let S' be the component of $S \setminus \sigma'$ containing α . Some power of f induces in $\text{Mod}(S')$ a power of a Dehn twist about α . Suppose p is this power, and consider f^{3p} . We may decompose this as $f^{3p} = g^k h = h g^k$, where g is a single Dehn twist about α , its power k is a nonzero multiple of three, and finally h fixes each component of $\partial S'$ while inducing the identity map in $\text{Mod}(S')$.

We claim that $i(f^{3p}\gamma, \gamma) \geq i(\alpha, \gamma)^2/2$. This can be seen by lifting to the universal cover. Consider the pre-image P of geodesic representatives of $\partial S' \cup \alpha$ in the universal cover, and choose a particular lift $\tilde{\alpha}$ of α . Letting C_- and C_+ be the components of $\mathbb{H}^2 \setminus P$ on each side of $\tilde{\alpha}$, we may choose a lift \tilde{g} of g which fixes C_- and acts on C_+ as the deck transformation D_α corresponding to α . Let \tilde{h} be a lift of h that also fixes C_- , so that we may lift f^{3p} to $\tilde{f}^{3p} = \tilde{h}\tilde{g}^k$.

We choose $n = i(\alpha, \gamma)$ consecutive lifts $\{\tilde{\gamma}_i\}_{i=1}^n$ of γ intersecting $\tilde{\alpha}$ on the restriction of a fundamental domain. By examining the endpoints at infinity, one can see that the geodesic representative of $\tilde{g}^k(\tilde{\gamma}_i)$ intersects either $D_\alpha^{k/3}(\tilde{\gamma}_j)$ or $D_\alpha^{2k/3}(\tilde{\gamma}_j)$ for each j (see Figure 2). Let us verify that the same is true replacing \tilde{g}^k with $\tilde{f}^{3p} = \tilde{h}\tilde{g}^k$. Let \tilde{x} be the point where $\tilde{\gamma}_i$ intersects $\tilde{\alpha}$. Starting at \tilde{x} and tracing the C_+ and C_- sides of $\tilde{\gamma}_i$, let $\tilde{\sigma}_+$ and $\tilde{\sigma}_-$ be the first lifts of components of $\partial S' \cup \alpha$ encountered on each side respectively. Note that, because h induces the identity map on S' , the lift \tilde{h} preserves all the lifts of $\partial S' \cup \alpha$ bordering C_+ and C_- , and therefore the half-spaces of $\tilde{\sigma}_-$ and $D_\alpha^k(\tilde{\sigma}_+)$ which do not contain α (and do contain the endpoints of $\tilde{g}^k(\tilde{\gamma}_i)$). Therefore these half-spaces contain the endpoints of $\tilde{f}^{3p}(\tilde{\gamma}_i)$. By their arrangement on the boundary at infinity, any geodesic between and not intersecting $D_\alpha^{k/3}(\tilde{\gamma}_i)$ and $D_\alpha^{2k/3}(\tilde{\gamma}_i)$ is forced to essentially intersect $\tilde{f}^{3p}(\tilde{\gamma}_i)$; such geodesics include a translate of each γ_j by either $D_\alpha^{k/3}$ or $D_\alpha^{2k/3}$. Clearly, $\tilde{f}^{3p}(\tilde{\gamma}_i)$ and $\tilde{f}^{3p}(\tilde{\gamma}_j)$ are disjoint for $i \neq j$, and by equivariance, the n^2 points of intersection can be pulled back to $(\cup_i \gamma_i) \cap (C_+ \cup C_-)$. It is not hard to argue that any point on the quotient surface has at most two pre-image points on this union of segments. Therefore at least $n^2/2$ unique intersections with γ and $f^{3p}(\gamma)$ appear on the surface, as claimed.

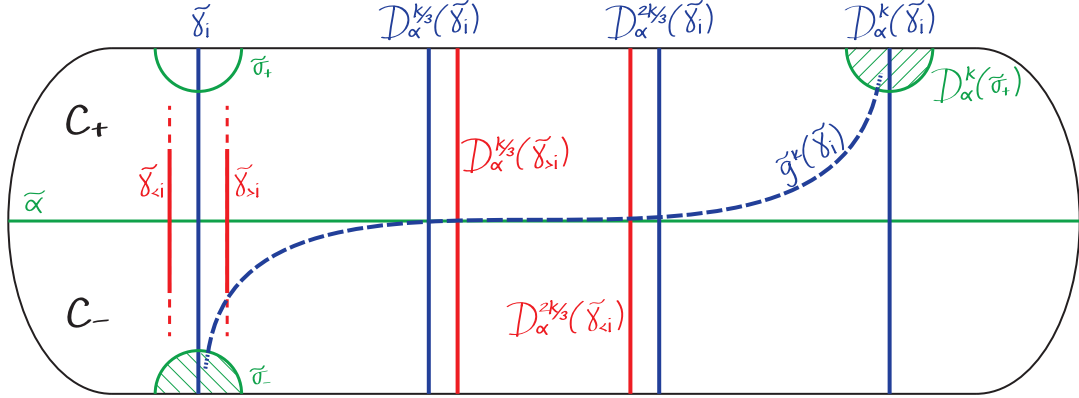


FIGURE 2. Schematic of \mathbb{H}^2 distorted to emphasize action by D_α . The endpoints of $\tilde{g}^k(\tilde{\gamma}_i)$ lie within the shaded half-spaces, which are preserved by the map \tilde{h} . Note that $\tilde{\gamma}_j$ for $j \neq i$ may or may not intersect $\tilde{\sigma}_\pm$.

So far we have

$$(M+1)^{3p \cdot \ell(f)+1} \geq M(M+1)^{\ell(f^{3p})} \geq i(f^{3p}(\gamma), \gamma) \geq i(\alpha, \gamma)^2/2,$$

assuming f^p restricts on a subsurface S' to Dehn twisting about α ; let us now bound the power p . Let S'_1 be one of the at most two components of $S \setminus \sigma(f)$ whose boundary contains α . As before, for some $p_0 \leq |\chi(S)/\chi(S'_1)|$, f^{p_0} preserves S'_1 . Some power p_1 of f^{p_0} induces an identity map on S'_1 . If S'_1 is the only component of $S \setminus \sigma(f)$ whose boundary contains α (so that S' is S'_1 glued along α), then $f^{p_0 p_1}$ Dehn twists around a neighborhood of α . The other possibility is that $S' \setminus \alpha$ consists of two components S'_1 and S'_2 ; then we need an additional power p_2 so that $f^{p_0 p_1 p_2}$ induces identity maps in both $\text{Mod}(S'_1)$ and $\text{Mod}(S'_2)$, and can therefore only be Dehn twisting on S' . For $i \in \{1, 2\}$, the powers p_i are bounded above by the maximal order of a mapping class of a surface S'_i with boundary. An upper bound is often given as $42|\chi(S'_i)|$ based on the fact that the smallest hyperbolic 2-orbifold has area $\pi/21$; however, the smallest cusped hyperbolic 2-orbifold has area $\pi/6$ [Sie45]. For this application we can thus use an upper bound for p_i of $12|\chi(S'_i)|$. We have

$$(2) \quad (M+1)^{(3 \cdot 144 \chi(S)^2 \ell(f)+2)/2} \geq i(\alpha, \gamma).$$

Theorem 3.1 follows from comparing inequalities (1) and (2). \square

4. THE LIST ALGORITHM

In order to find the canonical reducing system $\sigma(f)$ for a given mapping class f , we make use of Ivanov's elegant characterization, which generalizes to subgroups of $\text{Mod}(S)$ [Iva92]. Under this expanded definition, $\sigma(f)$ is really a property of $\langle f \rangle$. Recall that there is a power N , universal in S , such that for any $f \in \text{Mod}(S)$, the mapping class f^N is guaranteed to fix individually, rather than to permute, the components of $\sigma(f)$. In particular, one may take $N = \xi(S)!$, where $\xi(S)$ is the maximal number of mutually disjoint, isotopically distinct curves on S , given by $3g + p - 3$ for a surface with g genus and p punctures. Then the canonical reducing system $\sigma(f)$ is the set of simple closed curves α such that

- $f^N(\alpha) = \alpha$, and

- if $i(\alpha, \beta) > 0$, then $f^N(\beta) \neq \beta$.

For our algorithm, we need to know we can find a “test” curve β of short distance that will disqualify a candidate for $\sigma(f)$ that fulfills the first bullet but not the second. We use the following:

Lemma 4.1. *For $f \in \text{Mod}(S)$, write*

$$F = F(\ell_X(f)) = C^{\ell_X(f)},$$

where X is the Artin–Dehn–Humphries–Lickorish generating set for $\text{Mod}(S)$ and C is the constant furnished by Theorem 1.1, so that every $\alpha \in \sigma(f)$ has $\ell(\alpha) \leq F$, where ℓ is the length function defined by geometric intersection number with Γ , the set of base curves for the generating twists.

Suppose that δ is a simple closed curve with $f^N(\delta) = \delta$ and $\ell(\delta) \leq F$. If δ is not a component of $\sigma(f)$, then there exists β such that $i(\delta, \beta) > 0$ but $f^N(\beta) = \beta$, and $\ell(\beta) \leq 2F + 4$.

Proof. First observe that $\ell(\gamma) \leq 2$ for each $\gamma \in \Gamma$. Since Γ fills S , some $\gamma \in \Gamma$ must intersect δ . If this γ is fixed by f^N , we may take $\beta = \gamma$. So suppose γ is not fixed by f^N . It cannot be that γ lies entirely in a pseudo-Anosov component of f , since then δ intersects that component, contradicting that f^N fixes δ . Therefore γ intersects some component of $\sigma(f)$. Consider a point of intersection of γ and δ . Let a be the arc of γ one obtains by tracing that curve in both directions until it first hits $\sigma(f)$ on each side. This gives an essential arc of $\gamma \cap S'$ intersecting δ at least once, where S' is the component of $S \setminus \sigma(f)$ to which f restricts as the identity, and which contains δ . Let α consist of the one or two components of $\sigma(f)$ intersecting a at its endpoints. Since a is an essential arc, at least one component of the regular neighborhood of $a \cup \alpha$ is an essential curve of S lying completely in S' and thus fixed by f^N ; let this component be β . By construction, the length of β is less than the length of α plus twice the length of γ , which in turn is less than $2F + 4$. Also by construction, δ intersects β , and $f^N(\beta) = \beta$; thus β is the desired disqualifying curve. \square

4.1. The list-and-check algorithm. We are given a word f in Artin–Dehn–Humphries–Lickorish generators $X = \{T_{\gamma_i}\}$. Let $F = F(\ell_X(f))$ be the upper bound on length of components of $\sigma(f)$ given by Theorem 1.1 and referenced in Lemma 4.1. The algorithm proceeds as follows:

- (1) Populate the set LIST with the words representing all simple closed curves of length at most $2F + 4$.
- (2) For each curve $\gamma \in \text{LIST}$, test if $f^N(\gamma) = \gamma$. If this is true, add γ to the set TEST-LIST.
- (3) Let SHORTLIST be the subset of curves in TEST-LIST with length at most F .
- (4) For each curve $\gamma \in \text{SHORTLIST}$, test for intersection with curves in TEST-LIST. If there are no intersections, then $\gamma \in \sigma(f)$.
- (5) If $\sigma(f)$ is not empty, then f is reducible and its canonical reducing system is given. If $\sigma(f)$ is empty, but all the curves γ_i appear in SHORTLIST, then f is finite-order. Otherwise f is pseudo-Anosov.

4.2. Implementation and complexity. We need to encode simple closed curves so that we may compute their intersection numbers, and also their images under our fixed set of generators of $\text{Mod}(S)$. Here we describe an exponential-time implementation of our algorithm using a triangulation of the surface for the encoding, and computational methods described by Schaefer, Sedgwick, and Stefankovic [SSŠ02, SSŠ08]. Alternatively, our algorithm may be implemented using Dehn–Thurston coordinates as described in [Thu]. However, both approaches are currently limited to mapping class groups of surfaces with punctures. While either choice of coordinates certainly makes sense for curves on closed surfaces, we do not know how to compute intersection number in polynomial time.

A fixed triangulation of the surface allows one to describe a simple closed curve in terms of *normal coordinates*: one takes a representative of the simple closed curve which intersects the edges of the triangulation a minimal number of times, and then one produces a list with entries for each edge of the triangulation, where the entries record how many times a representative of the curve intersects

that edge. We require that the representative avoids vertices and enters and exits each triangle through different edges, and we require that the vertices of the triangulation lie on the boundary of the surface; then two sets of normal coordinates represent the same curve if and only if the coordinates are equal [SSŠ02]. In this case, Schaefer-Sedgwick-Stefankovic have polynomial-time algorithms that compute the coordinates of a curve after a Dehn twist about another curve, and the intersection number between curves [SSŠ08].

Our list-and-check algorithm starts with a function $F(\ell_X(f))$ for identifying sufficiently short curves. To be compatible with the algorithms of Schaefer-Sedgwick-Stefankovic, we use a new notion of length, which we call *triangulation length*: the number of times a curve intersects the triangulation. That is, where T is a fixed minimal triangulation, for any curve γ , let $\ell_T(\gamma)$ be the sum of all entries in the (unique) normal coordinates for γ . Now we must relate $\ell(\gamma)$ to $\ell_T(\gamma)$. Recall that X consists of twist and half-twist generators whose base curves fill the surface. We may assume all the vertices of our triangulation lie on the boundary punctures, and that all edges of the triangulation are geodesics in some hyperbolic metric, so that minimal intersection is realized between any edge of the triangulation and any geodesic curve. Furthermore, since all vertices are on the boundary, and minimal intersection allows no bigons, the count of intersections of any geodesic curve with each edge of the triangulation gives its normal coordinates. The filling curves cut S into a collection of disks and punctured disks. Let K be the maximum number of arcs of the triangulation through any of these components. Observe that, for any curve γ , we have $K \cdot \ell(\gamma) > \ell_T(\gamma)$. Therefore the task of listing all curves with $\ell < L$ is fulfilled by listing all curves with $\ell_T < LK$. We can create LIST in Step (1) of the algorithm using this encoding, in time polynomial in F , and thus exponential in $\ell_X(f)$. We have the necessary algorithms for Steps (2) and (4) by Theorem 2 of [SSŠ08], again in time polynomial in F . The complexity of the entire list-and-check algorithm implemented this way is exponential in word length.

We know of a third approach to curve-related computations, by encoding curves as words in the fundamental group of the surface. Then it is straightforward to compute images of a curve under products of a finite list of generators, and [BS84, BS87, CL87] give polynomial-time algorithms to count geometric intersection number, whether or not the surface has boundary. However, we do not know how to enumerate only words that represent simple curves, short of enumerating all words and checking for self-intersection. This way, the list-and-check algorithm requires time doubly exponential in word length.

5. THE COVER ALGORITHM

In this section, we prove Theorem 1.2. Let $X \subset \text{Mod}(S)$ be a finite generating set for $\text{Mod}(S)$. Fixing a basepoint on S , we obtain $\text{Mod}(S)$ as a quotient of $\text{Mod}^1(S)$, which is the group of based homotopy classes of orientation-preserving homeomorphisms of S . Viewing the chosen basepoint as a basepoint for $\pi_1(S)$, we obtain an action of $\text{Mod}^1(S)$ on $\pi_1(S)$. We fix a finite generating set Y for $\pi_1(S)$. Given Y , we compute conjugacy representatives $\{p_1, \dots, p_n\} \subset \pi_1(S)$ of small loops about each puncture of S . In all of our calculations, we disregard the conjugacy classes of these elements and their powers. This is a reasonable assumption, since once we identify conjugacy representatives of small loops about the punctures of S , the solution to the conjugacy problem in $\pi_1(S)$ gives us an effective way of determining if a homotopy class of loops on S is boundary parallel (see Section 6).

In what follows, $\ell_X(\cdot)$ denotes word length in $\text{Mod}(S)$ with respect to X , and likewise $\ell_Y(\cdot)$ is length in $\pi_1(S)$ using Y . We can lift X to $\tilde{X} \subset \text{Mod}^1(S)$ by choosing arbitrary preimages of the mapping classes in X . Such a choice gives us a set-theoretic splitting $\text{Mod}(S) \rightarrow \text{Mod}^1(S)$. For $x \in X$, the splitting $x \mapsto \tilde{x}$ allows us to determine a (maximum) *expansion factor* λ for mapping classes in $\text{Mod}(S)$, which we define by

$$\lambda = \max_{x \in X} \max_{y \in Y} \ell_Y(\tilde{x}(y)).$$

It is immediate from the definition of the expansion factor that if $f \in \text{Mod}(S)$ satisfies $\ell_X(f) = n$ and $w \in \pi_1(S)$, we have

$$(3) \quad \ell_Y(\tilde{f}(w)) \leq \lambda^n \cdot \ell_Y(w)$$

where \tilde{f} is any lift of f in $\langle \tilde{X} \rangle$ of length n . Thus, if $[w]$ is the conjugacy class of w in $\pi_1(S)$, then the shortest representative for $f([w])$ in $\pi_1(S)$ has length at most $\lambda^n \cdot \ell_Y(w)$.

Given a mapping class $f \in \text{Mod}(S)$ presented as a word in the generating set X , we would like to find a finite cover $\tilde{S} \rightarrow S$ to which f lifts, whose degree is controlled by a computable function in $\ell_X(f)$, and such that the action of any lift of f on $H_1(\tilde{S}, \mathbb{C})$ determines the Nielsen–Thurston classification (compare with Theorem 1.2). The idea is to build \tilde{S} as a regular cover of S and to exploit the interplay between the action of the deck group and the action of a lift of f on $H_1(\tilde{S}, \mathbb{C})$. The proof of Theorem 1.2 is fairly quick after we gather the necessary facts.

Lemma 5.1 (Lemma 3.1 in [Kob12]). *Let $\tilde{S} \rightarrow S$ be a regular finite cover with deck group Q . Then the action of Q on $H_1(\tilde{S}, \mathbb{C})$ is faithful. Moreover, we have an isomorphism of Q -representations:*

$$H_1(\tilde{S}, \mathbb{C}) \cong (\mathbb{C}[Q]/\mathbb{C})^{|\chi(S)|} \oplus \mathbb{C}^{b_1(S)}.$$

Here, $\mathbb{C}[Q]/\mathbb{C}$ denotes the regular representation of Q modulo its trivial summand.

In what follows, we identify each essential, nonperipheral closed curve c on S with the pair of nontrivial conjugacy classes $\{[c], [c^{-1}]\} \subset \pi_1(S)$ it determines. It is well-known that a mapping class $f \in \text{Mod}(S)$ fixes the isotopy class of c if and only if it preserves $\{[c], [c^{-1}]\}$ set-wise.

Lemma 5.2. *Given $\{c_1, \dots, c_s\}$ a finite list of essential, nonperipheral homotopy classes in $\pi_1(S)$, and $f \in \text{Mod}(S)$, there exists a finite characteristic quotient Q of $\pi_1(S)$ such that if $f(c_i) \neq c_i$ as isotopy classes of curves, then the images in Q of the pairs $\{[c_i], [c_i^{-1}]\}$ and $\{f([c_i]), f([c_i^{-1}])\}$ belong to distinct (conjugacy class, inverse conjugacy class) pairs in Q . Furthermore, $|Q| < N$, where N may be computed from $\ell_X(f)$, the length s of the list of curves, and an upper bound for $\ell_Y(c_i)$.*

Proof. When $s = 1$, the lemma follows from the effective conjugacy separability in free groups and surface groups (see Section 6), and the expansion factor observation noted in (3) above. In general, one can produce a quotient Q_i for each conjugacy class and then intersect the kernels. The intersection of the kernels will again be a characteristic subgroup of $\pi_1(S)$, and the index of the intersection will be controlled by the product of the orders of the quotients. \square

Lemma 5.3. *Let \tilde{S} be the connected cover of S determined by a finite characteristic quotient Q of $\pi_1(S)$. Suppose that $g \in \pi_1(S)$ and its image under $f \in \text{Mod}(S)$ have non-conjugate images in Q . Then any lift \tilde{f} of f acts nontrivially on $H_1(\tilde{S}, \mathbb{C})$; in particular, the images of g and $f(g)$ in Q induce different characters on $H_1(\tilde{S}, \mathbb{C})$ viewed as a representation of Q .*

Proof. Since Q is a characteristic quotient of $\pi_1(S)$, the mapping class f lifts to \tilde{S} , and we choose a fixed lift \tilde{f} . The induced isomorphism of vector spaces $\tilde{f}_* : H_1(\tilde{S}, \mathbb{C}) \rightarrow H_1(\tilde{S}, \mathbb{C})$ generally depends on the choice of lift, but we claim that its nontriviality is independent of the choice of lift. To see this, we fix a representative of $f \in \text{Aut}(\pi_1(S))$, which we will also call f . This way, we get an unambiguous action of f on $H_1(\tilde{S}, \mathbb{C})$ as well as on the quotient group Q . We will show that no matter what representative of f we choose in $\text{Aut}(\pi_1(S))$, the action of f on $H_1(\tilde{S}, \mathbb{C})$ is nontrivial.

Note that the naturality of the f - and Q -actions on $H_1(\tilde{S}, \mathbb{C})$ imply that $q \cdot f(v) = f(f^{-1}(q) \cdot v)$, for any $q \in Q$ and $v \in H_1(\tilde{S}, \mathbb{C})$. If q is the image of g in Q , we have that $f(q)$ is the image of $f(g)$ in Q . The hypothesis that q and $f(q)$ are not conjugate implies that there is an irreducible character χ of Q such that $\chi(q) \neq \chi(f^{-1}(q))$; this stems from the fact that the irreducible complex characters of Q span the ring of class functions on Q . Note in particular that the computation $\chi(q) \neq \chi(f^{-1}(q))$ is independent of the choice of representative of f in $\text{Aut}(\pi_1(S))$, because characters are invariant

under conjugation. Here, we precompose the character χ with f^{-1} instead of f in order to get a left action. In the decomposition of $H_1(\tilde{S}, \mathbb{C})$ into a direct sum of irreducible representations of Q , the summands corresponding to χ are mapped by f to the summands corresponding to $\chi \circ f^{-1}$, and these are distinct because their corresponding characters evaluate q differently. \square

Lemma 5.4. *There exists an L such that the (finite) set F of elements of $\pi_1(S)$ of length at most L represents a collection of closed curves on S which fills S .*

Proof. This is obvious after fixing any finite, filling set of curves on S . \square

Proof of Theorem 1.2. Let f be given of length n . We first check if f is nontrivial and whether or not it has finite order. There is a k which is independent of f such that f^k is trivial in $\text{Mod}(S)$ if and only if f has finite order. We have that f^k is trivial if and only if f^k preserves the conjugacy classes of each element in F from Lemma 5.4. We lift f to \tilde{f} in $\text{Mod}^1(S)$ and apply both \tilde{f} and \tilde{f}^k to each element of F . This process gives us a finite set of words of $\pi_1(S)$ with lengths at most $\lambda^{nk} \cdot L$, where λ is the expansion factor determined by our generating sets for $\text{Mod}(S)$ and $\pi_1(S)$.

By Lemma 5.2, there is a finite characteristic quotient Q of $\pi_1(S)$ such that any two non-conjugate elements of $\pi_1(S)$ of length at most $\lambda^{nk} \cdot L$ have non-conjugate images in Q . Fix $g \in \pi_1(S)$ such that g and $f(g)$ both have length at most $\lambda^{nk} \cdot L$, and write q and $f(q)$ for the images of g and $f(g)$ respectively in Q . Let \tilde{S} be the connected cover of S determined by the quotient Q ; the degree of \tilde{S} is bounded above by the N from Lemma 5.2.

By Lemma 5.3, we may check the values of the various characters of Q on the images of the elements of F , and determine whether or not f is the identity mapping class. In particular, we can determine whether or not f is the identity mapping class by examining its action on $H_1(\tilde{S}, \mathbb{C})$. A similar observation concerning the triviality of f^k shows that \tilde{S} is a classifying cover for finite-order mapping classes on S .

When f^k is nontrivial, it remains to determine whether f^k (and hence f) is reducible or pseudo-Anosov. We may assume k is such that f^k is reducible if and only if it fixes a curve. Using Theorem 1.1, we enumerate a list of pairs of conjugacy classes

$$\{[g_1^{\pm 1}], \dots, [g_s^{\pm 1}]\}$$

such that if f^k is reducible then f^k stabilizes at least one of these pairs of conjugacy classes. If λ is the expansion factor for our generating set for $\text{Mod}(S)$ then there are representatives for the conjugacy classes

$$\{f^k([g_1^{\pm 1}]), \dots, f^k([g_s^{\pm 1}])\}$$

whose lengths are at most $(C\lambda)^{nk}$.

The rest of the proof is identical to the finite order case. By Lemma 5.2, there is a finite characteristic quotient Q of $\pi_1(S)$ such that any two non-conjugate elements of $\pi_1(S)$ of length at most $(C\lambda)^{nk}$ have non-conjugate images in Q . As before, we take \tilde{S} to be the connected cover of S determined by the quotient Q . The mapping class f is pseudo-Anosov if and only if the characters induced on $H_1(\tilde{S}, \mathbb{C})$ by the images of $[g]$ and $f([g])$ in Q are not equal for each $g \in \pi_1(S)$ of length at most C^{nk} . \square

6. EFFECTIVE CONJUGACY SEPARABILITY

In this section, we will give a geometric proof that free groups and surface groups are effectively conjugacy separable. That is to say, given any non-conjugate pair of elements $g, h \in G$ where G is finitely generated by a set \mathcal{S} , there is a finite quotient $Q_{g,h}$ of G in which g and h are not conjugate, and such that the order of $Q_{g,h}$ is controlled by a computable function of $\ell_{\mathcal{S}}(g) + \ell_{\mathcal{S}}(h)$, where $\ell_{\mathcal{S}}(\cdot)$ denotes word length with respect to \mathcal{S} . We will handle the free group and closed surface group cases separately, though the arguments are essentially the same.

6.1. From quotients to characteristic quotients. We first gather some general facts about p -groups and characteristic quotients. The first is a trivial observation.

Lemma 6.1. *Let $g, h \in G$, let $G \rightarrow Q$ be a quotient of G in which the images of g and h are not conjugate, and let $G \rightarrow \overline{G}$ be a quotient of G which factors the quotient map $G \rightarrow Q$. Then the images of g and h are not conjugate in \overline{G} .*

Let p be a fixed prime.

Lemma 6.2. *Let G be a group and let H, K be a normal subgroups of G of p -power index. Then $H \cap K$ has p -power index in G .*

Proof. This is an easy consequence of the Second Isomorphism Theorem for groups. \square

We remark that in the previous lemma, the normality of H and K is essential.

Corollary 6.3. *Let G be an m -generated group and let $H < G$ be a normal subgroup of index $q = p^n$. Then there exists a subgroup $K < H < G$ such that K is characteristic in G and has index at most q^{q^m} .*

Proof. Write $Q = G/H$, which is a finite group of order q . There are at most q^m homomorphisms from G to Q , and the intersection K of the kernels of these homomorphisms has index at most q^{q^m} . \square

6.2. Free groups. The argument we give here is essentially a geometric version of the argument given in [LS01].

Proposition 6.4. *Finitely generated free groups are effectively conjugacy separable.*

Proof. Let F be a free group on a finite set \mathcal{S} . We may assume \mathcal{S} has at least two elements, since it is easy to see that, when F is cyclic, we may use $Q_{g,h}$ of order

$$p^{\max\{\ell_{\mathcal{S}}(g), \ell_{\mathcal{S}}(h)\}}.$$

Let X be a wedge of circles with vertex x , whose edges are labelled by the elements of \mathcal{S} . We choose orientations on the edges arbitrarily and choose a midpoint of each edge of X . If γ is an oriented based loop in X , we can record the sequence of midpoints traversed by γ while keeping track of orientations. This gives us an identification of $\pi_1(X, x)$ with F . If γ does not backtrack, its *combinatorial length* is defined to be the absolute number of midpoints traversed by γ . The combinatorial length of γ coincides with $\ell_{\mathcal{S}}(g)$ for g the element of F represented by γ .

We will use induction on the combinatorial length N to take two elements $w, v \in F$ with $\ell_{\mathcal{S}}(w) + \ell_{\mathcal{S}}(v) = N$, and produce a p -group quotient of F which has order controlled by a computable function of N , and in which the images of w and v are conjugate if and only if they are conjugate in F .

So, fix a prime $p > 2$, and let $w, v \in F$ be reduced and cyclically reduced words, represented by non-backtracking unbased loops in X with $\ell_{\mathcal{S}}(w) + \ell_{\mathcal{S}}(v) = N$. Throughout the proof, we allow for the case that one of v or w is trivial, so that the corresponding loop is just a point.

If $N \leq 2$ then either only one of v and w is nontrivial, or $\ell_{\mathcal{S}}(v) = \ell_{\mathcal{S}}(w) = 1$, in which case v and w are conjugate if and only if they are equal, which happens if and only if their images in $H_1(F, \mathbb{Z}/p\mathbb{Z})$ are equal (since $p > 2$). If only one of v and w is nontrivial, say v , then the image of v in $H_1(F, \mathbb{Z}/p\mathbb{Z})$ is nontrivial whereas the image of w is trivial. Indeed, v is a reduced word in F of combinatorial length exactly two and is therefore a product of at most two generators of F which are not inverses of each other. Again, since $p \neq 2$, we have that the image of v in $H_1(F, \mathbb{Z}/p\mathbb{Z})$ must be nontrivial.

If loops representing v and w are concentrated in a single wedge summand of X , so that these loops only intersect one midpoint of one edge of X , then v and w generate a cyclic subgroup of F .

In this case, we may look at the images of v and w in $H_1(F, \mathbb{Z}/p^N\mathbb{Z})$. It is clear that v and w will have distinct images in this quotient if and only if $v \neq w$.

For the general inductive step, we can make several simplifying assumptions. The first is that the span of the images of v and w in $H_1(F, \mathbb{Z}/p\mathbb{Z})$ has dimension at most one, after replacing v or w by an appropriate conjugate. Indeed, otherwise v and w are not conjugate since their homology classes are linearly independent after applying any conjugacy. Secondly, we may assume that v and w fill X , in the sense that they are not simultaneously homotopic into a proper subcomplex of X , because in that case we may restrict our attention to the proper subgraph of X filled by v and w .

We have that the loops representing v and w fill all of X . Note that the subgroup generated by v and w may still be cyclic even though all of X is filled. By the first assumption, there is a homomorphism $\phi : \pi_1(X) \rightarrow \mathbb{Z}/p\mathbb{Z}$ for which v and w are both in the kernel. Thus, loops representing v and w lift to the cover $Y \rightarrow X$ determined by ϕ . We pull back that edge-midpoint structure on X to Y . Choose lifts of the two loops representing v and w , which we call v_Y and w_Y respectively. Different choices of v_Y and w_Y differ by the action of the deck group, so choosing a particular lift does not affect conjugacy considerations. Therefore, we choose v_Y and w_Y so that they span a one-dimensional subspace of $H_1(Y, \mathbb{Z}/p\mathbb{Z})$, if possible.

Observe that the combinatorial lengths of v_Y and w_Y are equal to those of v and w respectively, with respect to the pulled-back edge-midpoint structure on Y . By the second assumption, namely the assumption that v and w fill X , there is an edge of Y which is not a closed loop and which is traversed by at least one of v and w . This edge can be extended to a maximal tree T in Y . A free basis for $\pi_1(Y)$ is identified with $Y \setminus T$.

Now consider the images $\overline{v_Y}$ and $\overline{w_Y}$ of v_Y and w_Y in the wedge of circles Y/T . The midpoints of edges in $Y \setminus T$ give rise to new combinatorial lengths for $\overline{v_Y}$ and $\overline{w_Y}$ as loops in Y/T . By the choice of T , we have that the sum of the combinatorial lengths of $\overline{v_Y}$ and $\overline{w_Y}$ is at most $N - 1$. The conclusion follows by induction on N . \square

We remark that the exact rank of F or of $\pi_1(Y/T)$ in the proof above is irrelevant for the induction. Furthermore, we remark that if the conjugacy classes of elements v and w of F are separated in a p -group quotient \tilde{Q} of F with kernel K , then by Subsection 6.1 we may produce a characteristic p -group quotient \tilde{Q} of F which admits Q as a further quotient, and whose size is controlled by the order of Q .

6.3. Closed surface groups. We will restrict our attention to orientable hyperbolic-type surfaces.

Proposition 6.5. *Let S be a closed surface. Then $\pi_1(S)$ is effectively conjugacy separable.*

Conjugacy separability of surface groups has been studied by various authors. H. Wilton [Wil10] gives an effective proof which relies on the work of Niblo [Nib92].

Proof of Proposition 6.5. Let S be a closed surface of genus at least two. We fix a hyperbolic structure on S such that a fundamental domain F for the action of $\pi_1(S)$ on \mathbb{H}^2 is given by a convex right-angled polygon. By taking finite connected unions of F along its faces, we can build fundamental domains for finite (possibly non-regular) covers of S .

Write $p : \mathbb{H}^2 \rightarrow S$ for the universal covering map. We assume that the boundary of the polygon F descends to a set of non-separating geodesic curves on S with linearly independent homology classes. Call this collection of non-separating curves \mathcal{X} . We tile \mathbb{H}^2 by $\pi_1(S)$ -translates of F , whose edges fit together into the collection $\mathcal{W} = p^{-1}(\mathcal{X})$ of bi-infinite geodesics of \mathbb{H}^2 . We orient the curves of \mathcal{X} arbitrarily and pull back these orientations to \mathcal{W} .

Let γ be an oriented geodesic representative of a primitive free homotopy class of closed curves in $\pi_1(S)$, and write $\ell_{hyp}(\gamma)$ for its hyperbolic length. The preimage $\tilde{\gamma} = p^{-1}(\gamma)$ in \mathbb{H}^2 is a union of geodesics, which generally will not be pairwise disjoint. Assume for now that γ is not a curve of \mathcal{X} . Choose a point x in the interior of $\tilde{\gamma}$ which is disjoint from \mathcal{W} and which sits on exactly

one bi-infinite geodesic $\delta \subset \tilde{\gamma}$. Following δ in the positive direction for a distance of $\ell_{hyp}(\gamma)$ and counting (without multiplicity) the number of intersection points with \mathcal{W} gives a well-defined notion of length which we call the *combinatorial length* of γ , and denote by $\ell_c(\gamma)$. It is easy to see that this number is independent of the choice of x . The combinatorial length of γ is also the number of interiors of distinct fundamental domains (from our tiling by translates of F) which a subsegment of $\tilde{\gamma}$ of length $\ell_{hyp}(\gamma)$ intersects, minus one. Observe that the combinatorial length of γ can be computed on S by counting the total number of intersections with \mathcal{X} . In the case $\gamma \in \mathcal{X}$, we make the convention that $\ell_c(\gamma) = 1$.

If $g = h^n \in \pi_1(S)$ is the n^{th} power of a primitive free homotopy class then we think of the corresponding geodesic γ_g representing g to be the same as the geodesic γ_h representing h , except that g traverses γ_h a total of n times. We call the geodesic γ_g the *underlying geodesic* for g , and we define the combinatorial length of g by $\ell_c(g) = n \cdot \ell_c(\gamma_h)$. Because $\ell_c(g)$ can be understood as the minimal number of times a curve representing g intersects \mathcal{X} , we have $\ell_c(g) \leq K\ell_S(g)$, where \mathcal{S} is a fixed generating set for $\pi_1(S)$ and $K = \max_{s \in \mathcal{S}} \{\ell_c(s)\}$. By construction, a nontrivial free homotopy class in $\pi_1(S)$ has nonzero combinatorial length.

For $S' \rightarrow S$ a finite cover, we may choose a fundamental domain F' for S' such that F' is a connected union of copies of F , and translates of F' tile \mathbb{H}^2 . Specifically, consider the cell decomposition of S' given by taking the pre-image of \mathcal{X} . We may delete edges of adjoining faces until the cell decomposition has a single 2-cell, whose pre-images in \mathbb{H}^2 we use to define the boundaries of the copies of F' . Let \mathcal{X}' be the 1-skeleton on S' of this new cell decomposition. If $\gamma \notin \mathcal{X}$ is a closed geodesic on S which lifts to $\tilde{\gamma}$ on S' , we define the combinatorial length of $\tilde{\gamma}$ to be the number of tiles of F' in \mathbb{H}^2 traversed by a subsegment of $\tilde{\gamma}$ of length $\ell_{hyp}(\gamma)$ which starts in the interior of a tile, minus one. We make the convention that, if $\gamma \in \mathcal{X}$ lifts to S' , the combinatorial length of its lift is again one. By construction, the combinatorial length of a geodesic on S' is bounded above by the combinatorial length of its quotient on S . It is also precisely the number of points of transverse intersection of the geodesic representative γ with \mathcal{X}' , a piecewise-geodesic graph on S' . If

$$S^i \rightarrow S^{i-1} \rightarrow \dots \rightarrow S$$

is a tower of covers, we may construct the corresponding fundamental domains compatibly so that F^i is a concatenation of copies of F^{i-1} , and the preimage of \mathcal{X}^i in \mathbb{H}^2 is a proper subset of the preimage of \mathcal{X}^{i-1} .

As before, for $g = h^n \in \pi_1(S') \subset \pi_1(S)$ the n^{th} power of a primitive free homotopy class, we make a new definition of combinatorial length in $\pi_1(S')$ by considering the underlying geodesic $\tilde{\gamma}_h$ in S' , and setting $\ell_{c'}(g) = n \cdot \ell_{c'}(\gamma_h)$.

As in the case of free groups, we proceed by induction on the sum of the combinatorial lengths of two free homotopy classes in $\pi_1(S)$. There are two different base cases to consider: the case where both g_1 and g_2 are nontrivial, and the case where only g_1 is nontrivial. Let γ_1 and γ_2 be closed, oriented geodesics on S underlying two (possibly trivial) free homotopy classes g_1 and g_2 in $\pi_1(S)$, and let $p > 2$ be a fixed prime. Suppose

$$\ell_c(g_1) + \ell_c(g_2) \leq 1.$$

In this case, we have that only g_1 is nontrivial. We see that γ_1 intersects exactly one curve in \mathcal{X} exactly once and therefore represents a primitive homology class of S which survives in $H_1(S, \mathbb{Z}/p\mathbb{Z})$. For the second base case, we have g_1 and g_2 are nontrivial and therefore

$$\ell_c(g_1) + \ell_c(g_2) = 2.$$

Since both g_1 and g_2 are nontrivial, we have $\ell_c(g_1) = \ell_c(g_2) = 1$ and hence γ_1 and γ_2 each intersect some curve in \mathcal{X} exactly once. We thus obtain that g_1 and g_2 are either equal and therefore conjugate, or they are distinct in $H_1(S, \mathbb{Z}/p\mathbb{Z})$ and therefore not conjugate in $\pi_1(S)$ (here, we are using the fact that $p \neq 2$).

Now suppose that $\ell_c(g_1) + \ell_c(g_2) = N > 2$. Let us first consider the case that γ_1 and γ_2 fill an annulus A of S . If A is nonseparating then we look at the images of g_1 and g_2 in $H_1(S, \mathbb{Z}/p^N\mathbb{Z})$ and note that they will be equal if and only if $g_1 = g_2$ as free homotopy classes. If A is separating, then we consider the cover $S_p \rightarrow S$ classified by the quotient $H_1(S, \mathbb{Z}/p\mathbb{Z})$. Clearly A lifts to S_p , and each lift of A to S_p is nonseparating. Choose lifts of g_1 and g_2 whose underlying geodesics lie within a fixed lift of A to S_p . Again, we have that the the images of these lifts in $H_1(S_p, \mathbb{Z}/p^N\mathbb{Z})$ will be equal if and only if $g_1 = g_2$ as free homotopy classes.

Now assume that γ_1 and γ_2 fill a subsurface $S_0 \subset S$ which is not an annulus. In particular, at least two distinct elements of \mathcal{X} intersect S_0 essentially. As in the free group case, we may assume that the span of the homology classes of γ_1 and γ_2 in $H_1(S, \mathbb{Z})$ has rank at most one, by replacing g_1 by an appropriate conjugate in $\pi_1(S)$. Finally, we may assume that under the inclusion $S_0 \subset S$, the image of $H_1(S_0, \mathbb{Z}) \rightarrow H_1(S, \mathbb{Z})$ has rank at least two.

To see this last claim, if S_0 has positive genus then there is nothing to do. Otherwise, we may assume that S_0 has genus zero and at least three boundary components. We describe the case where S_0 has exactly three boundary components, with the general case being a straightforward generalization. If D is a boundary component of S_0 which is separating in S , then it is easy to see that there is a degree p cover of S to which S_0 lifts and wherein every lift of D becomes nonseparating. Thus, we can replace S by an abelian p -cover of S to which S_0 lifts and in which each boundary component of each lift of S_0 is nonseparating. In this case, we have that the rank of the map $H_1(S_0, \mathbb{Z}) \rightarrow H_1(S, \mathbb{Z})$ is exactly two. Indeed, suppose the contrary. We have that each boundary component of S_0 represents a primitive homology class in $H_1(S, \mathbb{Z})$. Choosing a pair of boundary components $b_1, b_2 \subset S_0$, we have that b_1 and b_2 together separate S . But then the third boundary component $b_3 \subset S_0$ must separate S as well, which is a contradiction since the homology class of b_3 in S is nonzero. This establishes the claim.

We choose a homomorphism $\phi : \pi_1(S) \rightarrow \mathbb{Z}/p\mathbb{Z}$ such that the homology classes of γ_1 and γ_2 lie in the kernel of ϕ , so that the free homotopy classes g_1 and g_2 also lie in the kernel of ϕ , and we choose a component of the preimage of S_0 in the corresponding cover. Since S_0 is not an annulus and since γ_1 and γ_2 span a one-dimensional subspace of $H_1(S, \mathbb{Z})$, we can choose ϕ so that the homology class of some homologically nontrivial simple closed curve $\alpha \subset S_0$ does not lie in the kernel of ϕ . Let S_ϕ be the cover of S determined by the homomorphism ϕ . A fundamental domain F_ϕ for S_ϕ in \mathbb{H}^2 can be obtained by concatenating p copies of the fundamental domain F for S together, and the $\pi_1(S_\phi)$ -translates of ∂F_ϕ are a proper subset of \mathcal{W} , which descends to a 1-skeleton \mathcal{X}_ϕ on S_ϕ .

Choose any lifts $\tilde{\gamma}_1$ and $\tilde{\gamma}_2$ of γ_1 and γ_2 . Note that if the total combinatorial length

$$\ell_{c_\phi}(\tilde{\gamma}_1) + \ell_{c_\phi}(\tilde{\gamma}_2)$$

on S_ϕ is not strictly smaller than

$$\ell_c(\gamma_1) + \ell_c(\gamma_2)$$

on S , then the subsurface S_0 filled by γ_1 and γ_2 lifts to S_ϕ as well. This can be seen by considering a regular neighborhood of the segment of the lift of γ_i to \mathbb{H}^2 used to compute combinatorial length, for each i ; each time γ_i crosses to a new tile of F , it also crosses to a new tile of F_ϕ . However, we produced a curve $\alpha \subset S_0$ above which does not lift, which is a contradiction. The induction follows. \square

REFERENCES

- [Aou13] Tarik Aougab, *Uniform hyperbolicity of the graphs of curves*, Geom. Topol. **17** (2013), no. 5, 2855–2875.
- [BD11] Lluís Bacardit and Warren Dicks, *The Zieschang-McCool method for generating algebraic mapping-class groups*, Groups Complex. Cryptol. **3** (2011), no. 2, 187–220.
- [Bela] Mark C. Bell, *An algorithm for deciding reducibility*, preprint.
- [Belb] ———, *Flipper*, https://bitbucket.org/Mark_Bell/flipper.
- [BH92] Mladen Bestvina and Michael Handel, *Train tracks and automorphisms of free groups*, Ann. of Math. (2) **135** (1992), no. 1, 1–51.

- [BH95] M. Bestvina and M. Handel, *Train-tracks for surface homeomorphisms*, *Topology* **34** (1995), no. 1, 109–140.
- [BLM83] Joan S. Birman, Alex Lubotzky, and John McCarthy, *Abelian and solvable subgroups of the mapping class groups*, *Duke Math. J.* **50** (1983), no. 4, 1107–1120.
- [Bow] Brian H. Bowditch, *Uniform hyperbolicity the curve graphs*, *Pacific J. Math*, to appear.
- [BS84] Joan S. Birman and Caroline Series, *An algorithm for simple curves on surfaces*, *J. London Math. Soc.* (2) **29** (1984), no. 2, 331–342. MR 744104 (85m:57002)
- [BS87] ———, *Dehn’s algorithm revisited, with applications to simple curves on surfaces*, *Combinatorial group theory and topology* (Alta, Utah, 1984), *Ann. of Math. Stud.*, vol. 111, Princeton Univ. Press, Princeton, NJ, 1987, pp. 451–478. MR 895628 (88k:20059)
- [Cal14] Matthieu Calvez, *Fast Nielsen–Thurston classification of braids*, *Algebr. Geom. Topol.* **14** (2014), no. 3, 1745–1758.
- [CL87] Marshall Cohen and Martin Lustig, *Paths of geodesics and geometric intersection numbers. I*, *Combinatorial group theory and topology* (Alta, Utah, 1984), *Ann. of Math. Stud.*, vol. 111, Princeton Univ. Press, Princeton, NJ, 1987, pp. 479–500. MR 895629 (88m:57016)
- [CMP] Matt Clay, Johanna Mangahas, and Alexandra Pettet, *An algorithm to detect full irreducibility by bounding the volume of periodic free factors*, preprint, arXiv:1402.7342.
- [CRS] Matt Clay, Kasra Rafi, and Saul Schleimer, *Uniform hyperbolicity of the curve graph via surgery sequences*, preprint, arXiv:1302.5519.
- [FLP79] A. Fathi, F. Laudenbach, and V. Poenaru, *Travaux de Thurston sur les surfaces*, *Astérisque*, vol. 66, Société Mathématique de France, Paris, 1979, Séminaire Orsay, With an English summary.
- [FM12] Benson Farb and Dan Margalit, *A primer on mapping class groups*, *Princeton Mathematical Series*, vol. 49, Princeton University Press, Princeton, NJ, 2012.
- [HPW] Sebastian Hensel, Piotr Przytycki, and Richard C. H. Webb, *Slim unicorns and uniform hyperbolicity for arc graphs and curve graphs*, preprint, arXiv:1301.5577.
- [HTC96] Hessam Hamidi-Tehrani and Zong-He Chen, *Surface diffeomorphisms via train-tracks*, *Topology Appl.* **73** (1996), no. 2, 141–167.
- [Iva92] Nikolai V. Ivanov, *Subgroups of Teichmüller modular groups*, *Translations of Mathematical Monographs*, vol. 115, American Mathematical Society, Providence, RI, 1992, Translated from the Russian by E. J. F. Primrose and revised by the author.
- [Kob12] Thomas Koberda, *Asymptotic linearity of the mapping class group and a homological version of the Nielsen–Thurston classification*, *Geom. Dedicata* **156** (2012), 13–30.
- [LP01] Catherine Labruère and Luis Paris, *Presentations for the punctured mapping class groups in terms of Artin groups*, *Algebr. Geom. Topol.* **1** (2001), 73–114 (electronic).
- [LS01] Roger C. Lyndon and Paul E. Schupp, *Combinatorial group theory*, *Classics in Mathematics*, Springer-Verlag, Berlin, 2001, Reprint of the 1977 edition. MR 1812024 (2001i:20064)
- [Nib92] Graham Niblo, *Separability properties of free groups and surface groups*, *J. Pure Appl. Alg.* **78** (1992).
- [Sie45] Carl Ludwig Siegel, *Some remarks on discontinuous groups*, *Ann. of Math.* (2) **46** (1945), 708–718.
- [SSŠ02] Marcus Schaefer, Eric Sedgwick, and Daniel Štefankovič, *Algorithms for normal curves and surfaces*, *Computing and combinatorics, Lecture Notes in Comput. Sci.*, vol. 2387, Springer, Berlin, 2002, pp. 370–380.
- [SSŠ08] ———, *Computing dehn twists and geometric intersection numbers in polynomial time*, *20th Canadian Conference in Computational Geometry (CCCG 2008)*, 2008, pp. 111–114.
- [Tao13] Jing Tao, *Linearly bounded conjugator property for mapping class groups*, *Geom. Funct. Anal.* **23** (2013), no. 1, 415–466.
- [Thu] Dylan P. Thurston, *Geometric intersection of curves on surfaces*, preprint.
- [Web] Richard C. H. Webb, *A short proof of the bounded geodesic image theorem*, preprint, arXiv:1301.6187.
- [Wil10] Henry Wilton, *Virtual retractions, conjugacy separability*, *J. Algebra* **323** (2010).

DEPARTMENT OF MATHEMATICS, YALE UNIVERSITY, P.O. BOX 208283, NEW HAVEN, CT 06520-8283
E-mail address: thomas.koberda@gmail.com

DEPARTMENT OF MATHEMATICS, BROWN UNIVERSITY, 151 THAYER ST., PROVIDENCE, RI 02912
E-mail address: mangahas@math.brown.edu