Michigan Math. J. 64 (2015), 279–292

An Algorithm to Detect Full Irreducibility by Bounding the Volume of Periodic Free Factors

Matt Clay, Johanna Mangahas, & Alexandra Pettet

Abstract. We provide an effective algorithm for determining whether an element ϕ of the outer automorphism group of a free group is fully irreducible. Our method produces a finite list that can be checked for periodic proper free factors.

1. Introduction

Let F be a finitely generated nonabelian free group of rank at least 2. An outer automorphism ϕ is *reducible* if there exists a free factorization $F = A_1 * \cdots * A_k * B$ such that ϕ permutes the conjugacy classes of the A_i; else it is *irreducible.* Although irreducible elements have nice properties, for example, they are known to possess irreducible train-track representatives, irreducibility is not preserved under iteration. Thus, one often considers elements that are *irreducible with irreducible powers* (*iwip*), or *fully irreducible.* These are precisely the outer automorphisms ϕ for which there does not exist a proper free factor $A < F$ whose conjugacy class $[A]$ satisfies $\phi^p([A]) = [A]$ for any $p > 0$. If $\phi^p([A]) = [A]$ for some proper free factor $A < F$ and for some $p > 0$, then we say that $[A]$ is ϕ-periodic and, to avoid cumbersome language, also that the free factor A is ϕ-periodic. Fully irreducible elements are considered analogous to pseudo-Anosov mapping classes of hyperbolic surfaces. As such, they play an important role in the geometry and dynamics of the outer automorphism group $\mathrm{Out}(F)$ of F.

Although considered in some sense a "generic" property in $\mathrm{Out}(F)$, full irreducibility is not generally easy to detect. Kapovich [16] gave an algorithm for determining whether a given $\phi \in \mathrm{Out}(F)$ is fully irreducible, inspired by Pfaff's criterion for full irreducibility in [21]. At points in his algorithm, two processes run simultaneously, and although it is known that one of these must terminate, it is not a priori known which will; it thus seems unclear that the complexity of Kapovich's algorithm can be found without running the algorithm itself.

For mapping class groups and braid groups, there exist algorithms for determining whether or not a given element is pseudo-Anosov [6; 8; 3; 4; 20; 7]. Recently, Koberda and the second author [17] provided an elementary algorithm for determining whether or not a given mapping class is pseudo-Anosov, using a method of "list and check". They show that if a mapping class f is reducible,

Received March 12, 2014. Revision received July 22, 2014.

The first author is partially supported by NSF grant DMS-1006898. The second author is partially supported by NSF award DMS-1204592. The third author is partially supported by an NSERC Discovery grant.

that is, has an invariant multicurve, then the curves in its reduction system have length bounded by an exponential function in terms of the number of generators needed to write f. Therefore, given a mapping class f, a *list* is produced of all multicurves whose curves are sufficiently short. The action of f is then *checked* on these finitely many multicurves. If f fixes a multicurve from the list, it is reducible; otherwise, it is necessarily pseudo-Anosov.

In this article, we provide, in essence, a method of "list and check" for elements of $\mathrm{Out}(F)$, akin to that of Koberda and the second author. That is, we provide an algorithm that, given an element ϕ expressed as a product of generators from a finite generating set of $\mathrm{Out}(F)$, produces a finite list of free factors and checks each for ϕ-periodicity. The algorithm effectively determines whether or not the given element ϕ is fully irreducible. By *effective* we understand that there is a computable function that bounds the number of steps in terms of the size of the input and that does not utilize the algorithm. In particular, we avoid the use of dual processes, one of which must terminate.

2. Statement of Results

By $\mathrm{rk}(F)$ we denote the rank of the free group F. Let $\xi(F) = 3\,\mathrm{rk}(F) - 3$. This is the maximum number of edges in a finite graph with fundamental group F and without degree one or two vertices. This is also the maximum number of isotopy classes of disjoint, essential (not bounding a ball) spheres in the double of the handlebody of genus $\mathrm{rk}(F)$. An element $\phi \in \mathrm{Out}(F)$ that is not fully irreducible is *cyclically reducible* if there exists a ϕ-periodic rank 1 free factor; else it is *noncyclically reducible.*

Our algorithm to determine full irreducibility of an element $\phi \in \mathrm{Out}(F)$ consists of two effective processes. Process I determines (in the absence of an obvious reduction) if ϕ is cyclically reducible. As we shall see in Section 3, this will exploit algorithms that are already well known. Our main contribution to the algorithm is in process II. For this, we construct a finite list of conjugacy classes of proper free factors that contains a ϕ-periodic free factor if ϕ is noncyclically reducible. The length of this list is controlled by the word length of ϕ; this is the content of Theorem 1. A systematic check of the list then determines whether or not ϕ is fully irreducible.

To state our main theorem, we start by fixing a basis $\mathcal{X}$ for the free group F. Let $T = T_{\mathcal{X}}$ denote the Cayley graph for F with respect to $\mathcal{X}$. Given a subgroup $A \leq F$, the *volume* $\|A\|_{\mathcal{X}}$ of A is the number of edges in the *Stallings core* of the graph T/F. Recall that the Stallings core is the graph T_A/A, where T_A is the minimal subtree of T with respect to the action of A; or, equivalently, the Stallings core is the smallest subgraph of the cover of T/F associated to A that contains every embedded cycle (see [22] for details). Note that the volume function $\|\cdot\|_{\mathcal{X}}$ is constant on conjugacy classes of subgroups. The quantity $\|A\|_{\mathcal{X}}$ gives some measure of the complexity of the subgroup A in terms of the basis $\mathcal{X}$. For instance, if $A = \langle a \rangle$ is a cyclic subgroup, then the volume $\|\langle a \rangle\|_{\mathcal{X}}$ is the cyclic length of the element a as a word in the basis $\mathcal{X}$.

Now fix a finite generating set $\mathcal{S}$ for $\mathrm{Out}(F)$. Denote by $|\phi|_{\mathcal{S}}$ the word length of $\phi \in \mathrm{Out}(F)$ with respect to $\mathcal{S}$. Our main theorem describes a relation between the word length of a noncyclically reducible element of $\mathrm{Out}(F)$ and the volume of one of its periodic free factors.

THEOREM 1. *There is a computable constant $C = C(\mathcal{X}, \mathcal{S})$ such that, for any $\phi \in \mathrm{Out}(F)$, either*

(i) *ϕ is fully irreducible, or*
(ii) *ϕ has a periodic rank-1 free factor, or*
(iii) *ϕ has a periodic proper free factor A such that $\|A\|_{\mathcal{X}} \le C^{|\phi|_{\mathcal{S}}}$.*

In other words, if ϕ is noncyclically reducible, then $C^{|\phi|_{\mathcal{S}}}$ bounds the volume of some proper ϕ-periodic free factor. An exact formula for C is given at the end of Section 7.

As there are a finite number of conjugacy classes of free factors A of F for which $\|A\|_{\mathcal{X}}$ is bounded, the theorem provides a bound for the size of a list of conjugacy classes of free factors that can be used to conclusively determine whether or not an element $\phi \in \mathrm{Out}(F)$ of length $|\phi|_{\mathcal{S}}$ is fully irreducible if ϕ is not cyclically reducible.

To prove Theorem 1, we utilize the notion of *intersection number $i(S, T)$* defined between a pair of trees S and T equipped with an isometric action by F, as defined by Guirardel [11]. Horbez [15] related the intersection number $i(T, T\phi)$ to the word length of $\phi \in \mathrm{Out}(F)$ (Section 5, Theorem 5). We thus need only bound the volume of a ϕ-periodic proper free factor by $i(T, T\phi)$ (Section 7, Proposition 13).

Before embarking on the details of the proof of Theorem 1, we will first describe the procedure used in our algorithm for detecting fully irreducible elements of $\mathrm{Out}(F)$. This is contained in the next section, where we establish the following:

THEOREM 2. *There exists an effective algorithm for determining if an outer automorphism is fully irreducible.*

3. List and Check Algorithm

The input of our algorithm is an element $\phi_0 \in \mathrm{Out}(F)$. Recall that $\phi_0 \in \mathrm{Out}(F)$ is not fully irreducible if there exists a periodic proper free factor and note that the periodic free factors of ϕ_0 are exactly the periodic free factors of each of its powers. Feighn and Handel [10] showed that there is a power Q, depending on the rank of F but not on the element ϕ_0, so that any periodic free factor of ϕ_0^Q is in fact invariant. An explicit function for Q depending only on $\mathrm{rk}(F)$ can be found in [12] and [9]. For instance, Handel and Mosher show that this property is shared by all elements in $\ker(\mathrm{Out}(F) \to \mathrm{GL}(\mathrm{rk}(F), \mathbb{Z}_3))$ and hence $Q = \prod_{j=1}^{\mathrm{rk}(F)} (3^{\mathrm{rk}(F)} - 3^{j-1})$ suffices. This is analogous to the fact that the mapping class group has a finite index subgroup all of whose elements are *pure;* that is, any invariant multicurve is curve-wise fixed. As a preliminary step to our algorithm, we replace the element ϕ_0 by $\phi = \phi_0^Q$, so that henceforth we need only look for invariant free factors. Note that ϕ is irreducible if and only if it is fully irreducible if and only if ϕ_0 is fully irreducible.

PROCESS I. To begin process I, we apply an effective algorithm due to Bestvina and Handel [5], which finds a *relative train track*[1] representative $f : \Gamma \to \Gamma$ of ϕ. At its conclusion, if Γ has a nontrivial f-invariant subgraph, then ϕ fixes a proper free factor and is therefore reducible. Otherwise, the algorithm gives us an honest train track map representing ϕ. Recall that Bestvina and Handel [5] proved that the fixed subgroup of an automorphism whose outer class is irreducible is at most rank 1. Thus, we next want to check for loops homotopically fixed by f, which correspond to a fixed conjugacy classes of ϕ, and then see whether their corresponding elements generate a higher-rank subgroup of F.

For this, we make use of an algorithm of Turner in [23]. For an outer automorphism ϕ with train track map $f : \Gamma \to \Gamma$, Turner begins by describing a graph D_f equipped with a graph map $D_f \to \Gamma$. The components of D_f are in one-to-one correspondence with the fixed subgroups of the automorphisms in the outer class of ϕ, so that, restricted to a component of D_f, the map $D_f \to \Gamma$ is the covering map corresponding to the fixed subgroup of one of the elements of the outer class of ϕ. The algorithm provides an effective procedure for obtaining a finite subgraph C_f of D_f that carries the fundamental group of D_f. If any component of C_f has rank greater than 1, then ϕ is reducible. Otherwise, Whitehead's algorithm provides an effective method for determining whether any component of C_f corresponds to a primitive element. If one does, then ϕ is cyclically reducible. This marks the end of process I. At this point, we stop if we have found that ϕ is reducible, and we continue to process II if we have only managed to determine that ϕ is noncyclically reducible or fully irreducible.

PROCESS II. Theorem 1 gives an upper bound $V = C^{|\phi|_S}$ on the volume of the smallest ϕ-invariant free factor if ϕ is noncyclically reducible. (Recall that we have replaced our original input ϕ_0 by $\phi = \phi_0^Q$ for which invariance and periodicity are the same.) There are a finite number of conjugacy classes of subgroups H with volume less than this bound, and these can be systematically listed since they correspond to core graphs made from at most V edges, where each edge is oriented and labeled by an element of $\mathcal{X}$. For a gross overestimate of the number of these, one has $V \cdot (2\,\mathrm{rk}(F))^V \cdot B_{2V}$, where B_n, known as the nth *Bell number,* counts the number of partitions of n objects. In our case this is equivalent to the number of ways one can glue the $2V$ endpoints of V edges to obtain a graph. In particular, the number of conjugacy classes is less than $V(8V^2\,\mathrm{rk}(F))^V$ since $B_{2V} \le (2V)^{2V}$. Whitehead's algorithm is then used to eliminate conjugacy classes that are not free factors. We obtain a *list* of conjugacy classes of free factors that are *checked* (using, say, Stallings's graph pull backs [22]) one-by-one for

[1]Loosely speaking, a relative train track representative is akin to a Jordan form for a linear transformation; we will not make use of any properties of relative train track representatives and refer the reader to the references for details.

ϕ-invariance. This process, and hence the algorithm, stops once either an invariant free factor is identified, concluding with ϕ reducible, or once every item on the list is checked and found not to be invariant, determining that ϕ is fully irreducible.

This completes the proof of Theorem 2, with the assumption of Theorem 1. Now we proceed with the proof of Theorem 1.

4. Outer Space, Trees, and Morphisms

For mapping class groups, the intersection number between curves on the surface is in various contexts useful in comparison to distances in, for instance, Teichmüller space or the complex of curves. Similar methods have been emerging for $\mathrm{Out}(F)$ and its associated spaces. Culler and Vogtmann's *outer space* is the space cv consisting of metric simplicial trees T equipped with simplicial, free F-actions that are *minimal* (meaning that they leave no proper subtree invariant) up to isometry that commutes with the action. The action of $\mathrm{Out}(F)$ on cv is defined by precomposing the free group action with the outer automorphism; this action is therefore on the right. In some contexts, it is convenient to consider the projectivized outer space CV in which the sum of the lengths of the edges of the quotient T/F is 1. Outer space is treated as the analogue for $\mathrm{Out}(F)$ of Teichmüller space; we refer the reader to Vogtmann's survey [24] for a more detailed description.

For a tree $T \in cv$, we use $d_T(\cdot,\cdot)$ to denote the metric on T, $\ell_T(\cdot)$ to denote the length of edges or paths, and $\mathcal{E}(T)$ to denote the set of edges. We may consider edges oriented, depending on the context. In the special case that T has a single vertex orbit and unit length on every edge, we call T a *unit rose.*

Any pair of unit roses S, T are related by a *morphism* $f : S \to T$, by which we mean a cellular F-equivariant map that linearly expands every edge of S over a nonbacktracking edge path of T. The *length of a morphism* $f : S \to T$ is

$$\ell(f) = \max\{\ell_T(f(s)) \mid s \in \mathcal{E}(S)\},$$

and the *length of S in T* is

$$\ell_T(S) = \min\{\ell(f) \mid f : S \to T \text{ is a morphism}\}.$$

We use $\ell_T(S)$ instead of Lipschitz distance to simplify computations in the next section, but we remark that the two values are easily related by [15, Lemma 2.4]. If $f : S \to T$ satisfies $\ell(f) = \ell_T(S)$, then we say that f is *length minimizing.* In general, $\ell_T(S)$ and $\ell_S(T)$ are not equal, but it is known that the ratio of their logarithms is bounded away from zero, independently of S and T, when both trees are unit roses (or more generally, in the "thick part" of cv) [1; 13]. We do not require this fact in what follows. Instead, it is convenient to define

$$\lambda(S,T) = \max\{\ell_T(S), \ell_S(T)\}.$$

5. Intersection and the Guirardel Core

The utility of the intersection number between curves on a surface is carried over to free groups via the so-called *Guirardel core* $\mathcal{C}(S \times T)$: a certain closed, F-

invariant (with the diagonal action) cellular subset of the product $S \times T$ of trees in $S, T \in cv$. The *intersection number* $i(S, T)$ is the covolume of $\mathcal{C}(S \times T)$, that is, the sum of the areas of the 2-cells in $\mathcal{C}(S \times T)/F$. Often we may assume that S and T are unit roses, in which case $i(S, T)$ simply counts the squares in $\mathcal{C}(S \times T)/F$.

For our purpose, we do not need the full definition of $\mathcal{C}(S \times T)$, for which we refer the reader to [11; 15]. Rather, we make use of two approaches to computing $i(S, T)$. In one of these, intersection numbers are interpreted as the geometric intersection between sphere systems in the doubled handlebody. This connection is recalled in the proof of Lemma 11, where it is used.

The other approach is a simple criterion, given by Behrstock, Bestvina, and the first author [2], for when two edges $s \in \mathcal{E}(S)$, $t \in \mathcal{E}(T)$ determine a square $s \times t$ in the core $\mathcal{C}(S \times T)$. For a tree $T \in cv$, we let ∂T denote its boundary, that is, equivalence classes of geodesic rays where two rays are equivalent if their images lie in a bounded neighborhood of one another. An oriented edge $t \in \mathcal{E}(T)$ determines a subset $\mathrm{Cyl}_T^+(t) \subset \partial T$, its (*forward*) *one-sided cylinder,* which consists of equivalence classes of geodesics that contain a representative whose image contains t with the correct orientation. The complement of $\mathrm{Cyl}_T^+(t)$ in ∂T will be denoted by $\mathrm{Cyl}_T^-(t)$; clearly $\mathrm{Cyl}_T^-(t) = \mathrm{Cyl}_T^+(\bar{t})$, where $\bar{t}$ is t with the reverse orientation. We will typically not bother with specifying an orientation since we will consider both one-sided cylinders simultaneously. For $S, T \in cv$, there exists a canonical F-equivariant homeomorphism $\partial : \partial S \to \partial T$, which is induced by any morphism $f : S \to T$.

LEMMA 3 [2, Lemma 2.3]. *Let $S, T \in cv$, and let $\partial : \partial S \to \partial T$ denote the canonical F-equivariant homeomorphism. Given two edges $s \in \mathcal{E}(S)$ and $t \in \mathcal{E}(T)$, the rectangle $s \times t$ is in the core $\mathcal{C}(S \times T)$ if and only if each of the four subsets $\partial(\mathrm{Cyl}_S^{(\pm)}(s)) \cap \mathrm{Cyl}_T^{(\pm)}(t)$ is nonempty.*

Let $S, T \in cv$ and $t \in \mathcal{E}(T)$. The *slice* of the core $\mathcal{C}(S \times T)$ above t is the set

$$\mathcal{C}_t = \{s \in \mathcal{E}(S) \mid s \times t \subset \mathcal{C}(S \times T)\}.$$

Similarly, define the slice $\mathcal{C}_s = \{t \in \mathcal{E}(T) \mid s \times t \subset \mathcal{C}(S \times T)\}$ for $s \in \mathcal{E}(S)$. A simple application of Lemma 3 can be used to describe the slice.

LEMMA 4 [2, Lemma 3.7]. *Let $S, T \in cv$ and suppose $f : S \to T$ is a morphism. Given an edge $t \in \mathcal{E}(T)$ and a point y in the interior of t, the slice $\mathcal{C}_t \subset S$ of the core $\mathcal{C}(S \times T)$ is contained in the subtree spanned by $f^{-1}(y)$.*

Since F acts freely on the edges of T, for any point y that is in the interior of t, the subtree $\mathcal{C}_t \times \{y\}$ embeds in the quotient $\mathcal{C}(S \times T)/F$. Similarly, for a point x in the interior of s, the subtree $\{x\} \times \mathcal{C}_s$ embeds in the quotient. Therefore, the intersection number $i(S, T)$ can be expressed as

$$i(S, T) = \sum_{e \in \mathcal{E}(T/F)} \ell_T(\tilde{e}) \operatorname{vol}(\mathcal{C}_{\tilde{e}}) = \sum_{e \in \mathcal{E}(S/F)} \ell_S(\tilde{e}) \operatorname{vol}(\mathcal{C}_{\tilde{e}}), \tag{1}$$

where by $\tilde{e}$ we denote any lift of the edge e to T or S, respectively, and by $\mathrm{vol}(\cdot)$ we denote the sum of the lengths of the edges in the respective slice.

As mentioned in Section 2, Horbez [15] has recently given, for two trees in cv, a bound on their Guirardel intersection number based on F-equivariant maps between them, which in turn we can relate to the geometry of $\mathrm{Out}(F)$. We require a more precise formulation of his result than what is stated in [15], and we need only consider intersection between unit roses:

THEOREM 5 (Horbez [15]). *Let $S, T \in cv$ be unit roses. Then*

$$i(S,T) \leq 2\,\mathrm{rk}(F)^3 \lambda(S,T)^4.$$

For the remainder of this section, we derive this statement by a variation on the arguments in [15, Section 2.2]. Using trees rather than marked graphs, and keeping track of the precise dependence on the F-equivariant maps, we obtain inequalities not stated directly in [15] but needed for our applications.

REMARK 6. If $f : S \to T$ is a morphism between unit roses, then f is a bijection between the vertices of S and the vertices of T. Indeed, this follows since F acts freely and transitively on the vertex sets.

LEMMA 7 (cf. Lemmas 2.3 and 2.5 in [15]). *Suppose that $S, T \in cv$ are unit roses and $f : S \to T$ is a length minimizing morphism. If v_0, v_1 are vertices of T, then there exist vertices u_0, u_1 of S such that $f(u_0) = v_0$, $f(u_1) = v_1$, and*

$$d_S(u_0, u_1) \leq \lambda(S,T) d_T(v_0, v_1) + \lambda(S,T)^2.$$

Proof. Since the existence of u_0 and u_1 is clear by Remark 6, we need only prove the inequality. Let $g : T \to S$ be a length minimizing morphism. For any vertices u, u' in S, it is clear by concatenating the images of edges that

$$d_S(gf(u), gf(u')) \leq \ell(g) d_T(f(u), f(u')) \leq \ell(g)\ell(f) d_S(u, u').$$

Now consider the geodesic edge path q from u to $u'' = gf(u)$. There is a $w \in F$ for which $d_S(u, wu) = 1$ and whose axis intersects q only at u. Thus, $d_S(u'', wu'') = 2d_S(u, u'') + 1$. On the other hand,

$$d_S(u'', wu'') = d_S(gf(u), wgf(u)) = d_S(gf(u), gf(wu)) \leq \ell(g)\ell(f)$$

by our first observation and our choice of w. These last two statements together imply $2d_S(u, gf(u)) \leq \ell(g)\ell(f)$. Since u was arbitrary,

$$\begin{aligned} d_S(u_0, u_1) &\leq d_S(gf(u_0), gf(u_1)) + d_S(u_0, gf(u_0)) + d_S(u_1, gf(u_1)) \\ &\leq \ell(g) d_T(v_0, v_1) + \ell(g)\ell(f), \end{aligned}$$

from which the lemma follows. □

LEMMA 8 (cf. Proposition 2.8 in [15]). *Suppose that $S, T \in cv$ are unit roses and that $f : S \to T$ is a length minimizing morphism. Given any edge $t \in \mathcal{E}(T)$ and a point y in the interior of t, for any $x, x' \in f^{-1}(y)$, we have*

$$d_S(x, x') \leq 4\lambda(S,T)^2 + 2.$$

Proof. Fix v_0 a vertex on one end of t, and let u_0 be the vertex of S such that $f(u_0) = v_0$. Consider $x \in f^{-1}(y)$. Let $s \in \mathcal{E}(S)$ be the edge that contains x. Let u_1 be a vertex on one end of s, and set $v_1 = f(u_1)$. Observe that $f(s)$ is a geodesic of length at most $\ell(f)$. Because $f(s)$ contains both v_0 and v_1, $d_T(v_0, v_1) \leq \ell(f)$. By Lemma 7,

$$d_S(u_0, u_1) \leq \lambda(S, T)\ell(f) + \lambda(S, T)^2 \leq 2\lambda(S, T)^2.$$

Because $d_S(u_1, x) \leq 1$, $d_S(u_0, x) \leq 2\lambda(S, T)^2 + 1$. The same is true replacing x with x', hence the conclusion. □

COROLLARY 9. *Suppose that $S, T \in cv$ are unit roses. Given an edge $t \in \mathcal{E}(T)$, the diameter of the slice $\mathcal{C}_t \subset S$ of the core $\mathcal{C}(S \times T)$ is at most* $4\lambda(S, T)^2$.

Proof. Let $f : S \to T$ be a length minimizing morphism. Suppose that y is a point in the interior of the edge t. By Lemma 4 any two points in the slice $\mathcal{C}_t$ are contained in a geodesic between points in $f^{-1}(y)$, which by Lemma 8 has length at most $4\lambda(S, T)^2 + 2$. The endpoints of this geodesic are interior to edges, whereas the slice is a union of closed edges; in particular, the slice must exclude the partial edges at each end of the geodesic. Thus, the two points in the slice have distance at most $4\lambda(S, T)^2$. □

Proof of Theorem 5 (cf. Proposition 2.8 in [15]*).* Fix an edge $t \in \mathcal{E}(T)$ and a point y in the interior of t. Let $f : S \to T$ be a length minimizing morphism. The cardinality of $f^{-1}(y)$ is at most $\mathrm{rk}(F)\ell(f)$. By Lemma 4 and Corollary 9, the slice $\mathcal{C}_t$ is covered by the union of $\frac{1}{2}(\mathrm{rk}(F)\ell(f))^2$ edge paths of length at most $4\lambda(S, T)^2$. Thus, $\mathrm{vol}(\mathcal{C}_t) \leq 2\,\mathrm{rk}(F)^2\lambda(S, T)^4$. Hence, by (1) we have $i(S, T) \leq 2\,\mathrm{rk}(F)^3\lambda(S, T)^4$ as claimed. □

6. Subgroups and Volume Bounds

Given trees $S, T \in cv$ and a nontrivial finitely generated subgroup $A \leq F$, there exist nonempty subtrees $S_A \subset S$, $T_A \subset T$, on each of which A acts minimally. We can thus consider the Guirardel core $\mathcal{C}(S_A \times T_A)$ for these minimal subtrees with respect to the action of A. We might hope that, if A is a free factor of F, then $\mathcal{C}(S_A \times T_A)$ embeds into $\mathcal{C}(S \times T)$, so that $i(S_A, T_A)$ is always dominated by $i(S, T)$. Unfortunately this appears to be too much to expect, but we do achieve the following.

PROPOSITION 10. *Let A be a noncyclic finitely generated subgroup of F. Suppose that $S, T \in cv$ are unit roses and let $S_A \subset S$, $T_A \subset T$ be the minimal subtrees with respect to A. Then*

$$i(S_A, T_A) \leq 6\xi(A) \cdot \lambda(S, T)^3 \cdot i(S, T).$$

Proof. By equation (1) we have:

$$i(S,T) = \sum_{e \in \mathcal{E}(T/F)} \mathrm{vol}(\mathcal{C}_{\tilde{e}}),$$
$$i(S_A,T_A) = \sum_{e \in \mathcal{E}(T_A/A)} \ell_{T_A}(\tilde{e})\,\mathrm{vol}(\mathcal{A}_{\tilde{e}}), \tag{2}$$

where $\mathcal{C}_{\tilde{e}} \subset S$ and $\mathcal{A}_{\tilde{e}} \subset S_A \subset S$ are the slices in the respective cores. We denote by $\partial : \partial S \to \partial T$ the canonical F-equivariant homeomorphism and by $\partial_A : \partial S_A \to \partial T_A$ the canonical A-equivariant homeomorphism. Observe that $\partial|_{\partial S_A} = \partial_A$.

First, we claim that for each edge $\tilde{e} \subset T_A \subset T$, we have that $\mathcal{A}_{\tilde{e}} \subseteq \mathcal{C}_{\tilde{e}}$. Indeed, let s be an edge in $\mathcal{A}_{\tilde{e}}$. By Lemma 3, each of the four sets $\partial_A(\mathrm{Cyl}^{(\pm)}_{S_A}(s)) \cap \mathrm{Cyl}^{(\pm)}_{T_A}(\tilde{e})$ is nonempty. Since $\partial_A(\mathrm{Cyl}^{(\pm)}_{S_A}(s)) = \partial(\mathrm{Cyl}^{(\pm)}_{S_A}(s)) \subset \partial(\mathrm{Cyl}^{(\pm)}_{S}(s))$ and $\mathrm{Cyl}^{(\pm)}_{T_A}(\tilde{e}) \subset \mathrm{Cyl}^{(\pm)}_{T}(\tilde{e})$, each of the four sets $\partial(\mathrm{Cyl}^{(\pm)}_{S}(s)) \cap \mathrm{Cyl}^{(\pm)}_{T}(\tilde{e})$ is nonempty. Hence, s is an edge in $\mathcal{C}_{\tilde{e}}$.

By a *natural edge* of T_A we mean an edge path $\tilde{\mathrm{E}} = \tilde{e}_1, \tilde{e}_2, \ldots, \tilde{e}_n$ that is a connected component of $T_A - \mathcal{V}_{\geq 3}(T_A)$, where $\mathcal{V}_{\geq 3}(T_A)$ is the collection of vertices of degree at least three. A natural edge in T_A/A is the image of a natural edge in T_A; the set of all natural edges is denoted $\mathcal{E}_N(T_A/A)$.

Suppose that $\tilde{\mathrm{E}}$ is a natural edge of T_A consisting of the edge path $\tilde{e}_1, \tilde{e}_2, \ldots, \tilde{e}_n$. Then since $\mathrm{Cyl}^{(\pm)}_{T_A}(\tilde{e}_i) = \mathrm{Cyl}^{(\pm)}_{T_A}(\tilde{e}_j)$ for all i, j, from Lemma 3 we see that $\mathcal{A}_{\tilde{e}_i} = \mathcal{A}_{\tilde{e}_j}$. Therefore, we are justified in writing $\mathcal{A}_{\tilde{\mathrm{E}}}$ to denote any of the slices $\mathcal{A}_{\tilde{e}_i}$. Applying the observation that $\mathcal{A}_{\tilde{e}_i} \subseteq \mathcal{C}_{\tilde{e}_i}$, we have that

$$\mathcal{A}_{\tilde{\mathrm{E}}} \subseteq \bigcap_{i=1}^{n} \mathcal{C}_{\tilde{e}_i}. \tag{3}$$

Next, let us bound $\ell_{T_A}(\tilde{\mathrm{E}})$ whenever $\mathcal{A}_{\tilde{\mathrm{E}}}$ is not empty. Let $f : S \to T$ be a length minimizing morphism, and consider interior points $x_1 \in \tilde{e}_1$ and $x_n \in \tilde{e}_n$. If there exists a point $p \in \mathcal{A}_{\tilde{\mathrm{E}}} \subseteq \mathcal{C}_{\tilde{e}_1} \cap \mathcal{C}_{\tilde{e}_n}$, then it lives in a geodesic between a pair of points in $f^{-1}(x_1)$ by Lemma 4. By Lemma 8, this path has length at most $4\lambda(S,T)^2 + 2$. The same is true replacing x_1 with x_n. Because p is in the intersection of these two bounded geodesics, we may choose $y_1 \in f^{-1}(x_1)$ and $y_n \in f^{-1}(x_n)$ so that $d(y_1, y_n) \leq 4\lambda(S,T)^2 + 2$. Furthermore, $d(x_1, x_n) \leq \ell(f)d(y_1, y_n)$, and we may choose x_1 and x_n so that $d(x_1, x_n)$ is arbitrarily close to $\ell_{T_A}(\tilde{e})$. Thus, we conclude that, when $\mathrm{vol}(\mathcal{A}_{\tilde{\mathrm{E}}}) > 0$,

$$\ell_{T_A}(\tilde{\mathrm{E}}) \leq \ell_T(S)(4\lambda(S,T)^2 + 2) < 6\lambda(S,T)^3.$$

Rewriting (2), we get:

$$i(S_A,T_A) = \sum_{e \in \mathcal{E}(T_A/A)} \ell_{T_A}(\tilde{e})\,\mathrm{vol}(\mathcal{A}_{\tilde{e}}) = \sum_{\mathrm{E} \in \mathcal{E}_N(T_A/A)} \ell_{T_A}(\tilde{\mathrm{E}})\,\mathrm{vol}(\mathcal{A}_{\tilde{\mathrm{E}}})$$
$$\leq \sum_{\mathrm{E} \in \mathcal{E}_N(T_A/A)} 6\lambda(S,T)^3\,\mathrm{vol}(\mathcal{A}_{\tilde{\mathrm{E}}}).$$

By (3), $\mathrm{vol}(\mathcal{A}_{\tilde{\mathrm{E}}}) \leq i(S,T)$ for every natural edge $\mathrm{E} \in \mathcal{E}_N(T_A/A)$. Since T_A/A has at most $\xi(A)$ natural edges, the proof is completed. □

In order to eventually relate intersection number to the volume of an invariant free factor, we find an effective lower bound for intersection under a bounded iterate of a fully irreducible automorphism.

LEMMA 11. *Let ϕ be a fully irreducible element of* $\mathrm{Out}(F)$ *and consider a tree $T \in cv$ with edge lengths at least* 1. *Then, for some* $1 \leq P \leq \xi(F)$,

$$\mathrm{vol}(T/F) \leq \xi(F) \cdot i(T, T\phi^P).$$

Proof. The lemma will be proved once we establish that, for some $1 \leq P \leq \xi(F)$, the slice of the core $\mathcal{C}(T \times T\phi^P)$ above the longest edge of T contains at least one edge of $T\phi^P$ and hence has volume at least 1. This is because, if the longest edge is $\tilde{e}$, by equation (1) we would have

$$\mathrm{vol}(T/F) \leq \xi(F) \cdot \ell_T(\tilde{e}) \leq \xi(F) \cdot \ell_T(\tilde{e}) \cdot \mathrm{vol}(\mathcal{C}_{\tilde{e}}) \leq \xi(F) \cdot i(T, T\phi^P).$$

To prove the claim above, we compute the intersection number using *sphere systems* in the doubled handlebody with fundamental group F. Briefly, let M be the connect sum of as many copies of $S^1 \times S^2$ as the rank of F. By $\mathbb{S}$ we denote the simplicial complex whose n-simplices correspond to $n+1$ isotopy classes of disjoint essential spheres in M, and by $\mathbb{S}^\infty$ we denote the subcomplex of $\mathbb{S}$ consisting of simplices where the complement of the corresponding sphere system in M has a non-simply-connected component. By work of Laudenbach [18; 19], there is a well-defined simplicial action of $\mathrm{Out}(F)$ on $\mathbb{S}$ that leaves $\mathbb{S}^\infty$ invariant. In this action, fully irreducible elements of $\mathrm{Out}(F)$ act on $\mathbb{S}$ without periodic orbits. Hatcher [14] established an $\mathrm{Out}(F)$-equivariant isomorphism between projectivized outer space CV and $\mathbb{S} - \mathbb{S}^\infty$. Under this isomorphism, edges of a marked graph T/F correspond bijectively to spheres in some sphere system.

Horbez details the correspondence between geometric intersection of the sphere systems and the volume of the Guirardel core [15]. In particular, he shows that, if T_0, T_1 are trees in CV, then for the corresponding sphere systems $\Sigma_0, \Sigma_1 \in \mathbb{S} - \mathbb{S}^\infty$, we have $i(T_0, T_1) = i(\Sigma_0, \Sigma_1)$, where the latter counts the minimal number of circles common to each sphere system, weighted appropriately. This minimum is achieved by representative sphere systems in a notion of normal position first described by Hatcher [14]. Whereas not stated explicitly in [15], it can be verified that each circle of intersection occurring on a given component $\sigma_0 \in \Sigma_0$ corresponds to an edge in the slice of the core $\mathcal{C}(T_0 \times T_1)$ above an edge in T_0 corresponding to the lift of the edge in T_0/F dual to σ_0. This is because $\mathcal{C}(T_0 \times T_1)$ can be built as the 2-complex dual to preimages of Σ_1 and Σ_2 in the universal cover of M, where Σ_1 and Σ_2 are assumed to be in normal position. For our considerations, the weights on the spheres do not matter since we are only concerned with showing that some slice is nonempty, that is, that the corresponding sphere has nontrivial intersection with another sphere.

Now, given $T \in cv$, we scale its edges equally by $1/\mathrm{vol}(T/F)$ to get a point $\overline{T} \in CV$. The slice over the longest edge of T in $\mathcal{C}(T \times T\phi^P)$ is nonempty if and only if the slice over the longest edge of $\overline{T}$ in $\mathcal{C}(\overline{T} \times \overline{T}\phi^P)$ is nonempty. Suppose that Σ is the sphere system dual to $\overline{T}$ and that $\sigma \in \Sigma$ is dual to the longest edge of $\overline{T}$. Since the maximum number of isotopy classes of disjoint essential spheres in M is $\xi(F)$ and since ϕ is fully irreducible, at least two of the spheres $\sigma, \phi(\sigma), \ldots, \phi^{\xi(F)}(\sigma)$ have essential intersection, so in particular σ essentially intersects $\phi^P(\sigma)$ for some $1 \leq P \leq \xi(F)$. By the foregoing discussion, this means that the slice above the longest edge in T/F contains at least one edge, proving the lemma. □

Remark 12. We use sphere systems in the proof to potentially give intuition on how the reasoning parallels that for its mapping class group analogue in [17]. Alternatively, Lemma 11 can be proved using trees instead; we give a sketch of that argument here. As in Lemma 11, we will show that the slice of the core above the longest edge in T contains at least one edge.

To this end, let $\tilde{e}$ be the longest edge in T and consider $T_1 = X$, the tree obtained by collapsing every edge other than the ones in the orbit of $\tilde{e}$. If $i(T_1, X\phi) \neq 0$, then the slice $\mathcal{C}_{\tilde{e}} \subset X\phi$ is nonempty since $i(T_1, X\phi) = \ell_{T_1}(\tilde{e})\,\mathrm{vol}(\mathcal{C}_{\tilde{e}})$. Since $T\phi$ collapses to $X\phi$, we see that the slice of the core above $\tilde{e}$ in $\mathcal{C}(T \times T\phi)$ is also nonempty.

If $i(T_1, X\phi) = 0$, then by [11, Theorem 6.1] the core is tree, denote it T_2. Moreover, T_2 is a *common refinement* for both T_1 and $X\phi$. That is, there are edge collapse maps $T_1 \leftarrow T_2 \rightarrow X\phi$. There is a unique edge in T_2 that is mapped homeomorphically to $\tilde{e} \subset T_1$. Abusing notation, we denote this edge by $\tilde{e}$ as well.

As before, if $i(T_2, X\phi^2) \neq 0$, then we see that the slice $\mathcal{C}_{\tilde{e}} \subset X\phi^2$ is nonempty, and therefore the slice above $\tilde{e}$ in $\mathcal{C}(T \times T\phi^2)$ is also nonempty. If $i(T_2, X\phi^2) = 0$, then the core is a common refinement for the two trees T_2 and $X\phi^2$, denote it T_3.

Continuing in this fashion, if $i(T_k, X\phi^k) = 0$, denote the common refinement by T_{k+1}. Since ϕ is fully irreducible, at each step T_k/F has k edges. Thus, for some $1 \leq P \leq \xi(F)$, we must have that $i(T_P, X\phi^P) \neq 0$. Hence, the slice of the core above $\tilde{e}$ in $\mathcal{C}(T \times T\phi^P)$ contains at least one edge, proving the lemma.

We apply the previous two results to prove the following.

Proposition 13. *Let $T = T_{\mathcal{X}}$ be the Cayley graph with respect to the basis $\mathcal{X}$, with all edges of unit length. If $\phi \in \mathrm{Out}(F)$ acts fully irreducibly on a proper free factor A of rank at least 2, then for some $1 \leq P \leq \xi(F)$,*

$$\|A\|_{\mathcal{X}} \leq 6\xi(F)^2 \cdot \lambda(T, T\phi^P)^3 \cdot i(T, T\phi^P).$$

Proof. The minimal tree $T_A \subset T$ of A has natural edge-lengths at least 1, and as such can be thought of as an element of $cv(A)$, the unprojectivized outer space for A. We can apply Lemma 11 to T_A with its free A-action to obtain $P \leq \xi(A)$ for which

$$\|A\|_{\mathcal{X}} = \mathrm{vol}(T_A/A) \leq \xi(A) \cdot i(T_A, T_A\phi^P).$$

The conclusion follows by applying Proposition 10, noting that $\xi(A) \leq \xi(F)$. □

7. Proof of Theorem 1

In this section we prove the key new result for our algorithm, which is Theorem 1. We wish to show that if $\phi \in \text{Out}(F)$ is noncyclically reducible, then there is a ϕ-periodic free factor whose volume is bounded above by an exponential function in terms of the word length $|\phi|_{\mathcal{S}}$. Since ϕ is noncyclically reducible, there is a ϕ^Q-invariant free factor A for which $1 < \text{rk}(A) < \text{rk}(F)$, where $Q = Q(\text{rk}(F))$ is the constant power mentioned at the beginning of Section 3. We can assume that $\phi^Q|_A$ is fully irreducible.

Now let us complete the proof of Theorem 1. Let $T = T_{\mathcal{X}}$ be the Cayley graph with respect to the basis $\mathcal{X}$, with all edges of unit length. Combining with Theorem 5 with Proposition 13, we have some $1 \leq P \leq \xi(F)$ for which

$$\begin{aligned} \|A\|_{\mathcal{X}} &\leq 6\xi(F)^2 \cdot \lambda(T, T\phi^{QP})^3 \cdot i(T, T\phi^{QP}) \\ &\leq 12\xi(F)^5 \cdot \lambda(T, T\phi^{QP})^7. \end{aligned}$$

Let $\lambda_{\mathcal{X}}(\mathcal{S}) = \max\{\lambda(T, T\psi) \mid \psi \in \mathcal{S}\}$. Since the length of a composition of morphisms is bounded by the product of their lengths, we have $\lambda(T, T\phi^{QP}) \leq \lambda_{\mathcal{X}}(\mathcal{S})^{QP|\phi|_{\mathcal{S}}}$.

The proof of Theorem 1 is complete with

$$\begin{aligned} \|A\|_{\mathcal{X}} &\leq 12\xi(F)^5 \cdot \lambda_{\mathcal{X}}(\mathcal{S})^{7QP|\phi|_{\mathcal{S}}} \\ &\leq C^{|\phi|_{\mathcal{S}}}, \end{aligned}$$

where $C = 12\xi(F)^5 \cdot \lambda_{\mathcal{S}}(\mathcal{X})^{7Q\xi(F)}$ depends only on $\mathcal{X}$ and $\mathcal{S}$.

ACKNOWLEDGMENTS. We would like to thank the Centre de Recerca Matemàtica for its hospitality during its research program, *Automorphisms of Free Groups: Algorithms, Geometry, and Dynamics,* in Fall, 2012. The authors also thank Sam Taylor for comments on an earlier version of this work.

References

[1] Y. Algom-Kfir and M. Bestvina, *Asymmetry of outer space,* Geom. Dedicata 156 (2012), 81–92.

[2] J. Behrstock, M. Bestvina, and M. Clay, *Growth of intersection numbers for free group automorphisms,* J. Topol. 3 (2010), no. 2, 280–310.

[3] M. Bell, *An algorithm for deciding reducibility,* preprint, arXiv:1403.2997.

[4] D. Benardete, M. Gutiérrez, and Z. Nitecki, *A combinatorial approach to reducibility of mapping classes,* Mapping class groups and moduli spaces of Riemann surfaces, Göttingen, 1991/Seattle, WA, 1991, Contemp. Math., 150, pp. 1–31, Amer. Math. Soc., Providence, RI, 1993.

[5] M. Bestvina and M. Handel, *Train tracks and automorphisms of free groups,* Ann. of Math. (2) 135 (1992), no. 1, 1–51.

[6] ———, *Train-tracks for surface homeomorphisms,* Topology 34 (1995), no. 1, 109–140.

[7] M. Calvez, *Fast Nielsen–Thurston classification of braids,* preprint, arXiv:1112.0165.

[8] Z.-H. Chen and H. Hamidi-Tehrani, *Surface diffeomorphisms via train-tracks,* Topology Appl. 73 (1996), no. 2, 141–167.

[9] M. Feighn and M. Handel, in preparation.

[10] ______, *The recognition theorem for* Out(F_n)*,* Groups Geom. Dyn. 5 (2011), no. 1, 39–106.

[11] V. Guirardel, *Cœur et nombre d'intersection pour les actions de groupes sur les arbres,* Ann. Sci. Éc. Norm. Supér. (4) 38 (2005), no. 6, 847–888.

[12] M. Handel and L. Mosher, *Subgroup decomposition in* Out(F_n) *part II: a relative Kolchin theorem,* preprint, arXiv:1302.2379.

[13] ______, *The expansion factors of an outer automorphism and its inverse,* Trans. Amer. Math. Soc. 359 (2007), no. 7, 3185–3208 (electronic).

[14] A. Hatcher, *Homological stability for automorphism groups of free groups,* Comment. Math. Helv. 70 (1995), no. 1, 39–62.

[15] C. Horbez, *Sphere paths in outer space,* Algebr. Geom. Topol. 12 (2012), no. 4, 2493–2517.

[16] I. Kapovich, *Algorithmic detectability of iwip automorphisms,* Bull. Lond. Math. Soc. 46 (2014), no. 2, 279–290.

[17] T. Koberda and J. Mangahas, *Effective algebraic detection of the Nielsen–Thurston classification of mapping classes,* preprint, arXiv:1312.6141.

[18] F. Laudenbach, *Sur les* 2*-sphères d'une variété de dimension* 3*,* Ann. of Math. (2) 97 (1973), 57–81.

[19] ______, *Topologie de la dimension trois: homotopie et isotopie,* Astérisque, 12, Société Mathématique de France, Paris, 1974, with an English summary and table of contents.

[20] J. E. Los, *Pseudo-Anosov maps and invariant train tracks in the disc: a finite algorithm,* Proc. Lond. Math. Soc. (3) 66 (1993), no. 2, 400–430.

[21] C. Pfaff, *Ideal Whitehead graphs in $Out(F_r)$ II: the complete graph in each rank,* J. Homotopy Relat. Struct. (2013).

[22] J. R. Stallings, *Topology of finite graphs,* Invent. Math. 71 (1983), no. 3, 551–565.

[23] E. C. Turner, *Finding indivisible Nielsen paths for a train track map,* Combinatorial and geometric group theory, Edinburgh, 1993, London Math. Soc. Lecture Note Ser., 204, pp. 300–313, Cambridge Univ. Press, Cambridge, 1995.

[24] K. Vogtmann, *Automorphisms of free groups and outer space,* Proceedings of the conference on geometric and combinatorial group theory, part I, Haifa, 2000, Geom. Dedicata 94 (2002), 1–31.

M. Clay
Department of Mathematics
University of Arkansas
Fayetteville, AR 72701
USA

mattclay@uark.edu

J. Mangahas
Department of Mathematics
University at Buffalo
Buffalo, NY 14260
USA

mangahas@buffalo.edu

A. Pettet
Department of Mathematics
University of British Columbia
Vancouver, BC V6T 1Z2
Canada

alexandra@math.ubc.ca