

Virtually abelian subgroups of $\text{IA}_n(\mathbb{Z}/3)$ are abelian

Michael Handel* and Lee Mosher†

August 8, 2021

Abstract

When studying subgroups of $\text{Out}(F_n)$, one often replaces a given subgroup \mathcal{H} with one of its finite index subgroups \mathcal{H}_0 so that virtual properties of \mathcal{H} become actual properties of \mathcal{H}_0 . In many cases, the finite index subgroup is $\mathcal{H}_0 = \mathcal{H} \cap \text{IA}_n(\mathbb{Z}/3)$. For which properties is this a good choice? Our main theorem states that being abelian is such a property. Namely, every virtually abelian subgroup of $\text{IA}_n(\mathbb{Z}/3)$ is abelian.

Contents

1	Introduction	1
2	Reduction to Proposition 1.3	3
3	Background	4
3.1	Basics	4
3.2	UPG elements	7
3.3	Abelian UPG subgroups	13
4	Proof of Proposition 1.3	15

1 Introduction

It is common, when studying elements of $\text{Out}(F_n)$, to replace the given element by an iterate in order to improve its invariance properties. For example, each $\theta \in \text{Out}(F_n)$ has an iterate $\phi = \theta^k$ satisfying the following properties.

*The first author was supported by National Science Foundation grant DMS-1308710 and by a PSC-CUNY grant in Program Year 47.

†The second author is supported by National Science Foundation grant DMS-1708361.

- (1) If some iterate of ϕ fixes a conjugacy class $[a]$ then ϕ fixes $[a]$.
- (2) If some iterate of ϕ fixes the conjugacy class $[F]$ of a free factor F then ϕ fixes $[F]$.
- (3) ϕ fixes each element in its set $\mathcal{L}(\phi)$ of attracting laminations.
- (4) ϕ fixes each element in its set of singular rays and eigenrays.

If θ is *rotationless* in the sense of [FH11] then iteration is not necessary: each of the above properties is automatically satisfied by $\phi = \theta$ [FH11, Lemma 3.30 and Definition 3.13]. Every θ has a rotationless iterate and the number of iterates required is uniformly bounded [FH11, Lemma 4.42].

The subgroup analog of replacing an individual element θ with a rotationless iterate θ^k is to replace a given subgroup \mathcal{H} with its finite index subgroup $\mathcal{H} \cap \text{IA}_n(\mathbb{Z}/3)$ where $\text{IA}_n(\mathbb{Z}/3) < \text{Out}(F_n)$ is the finite index subgroup consisting of elements that act trivially on $\mathbb{Z}/3$ - homology. This was done, for example, in the proof of the Tits Alternative for $\text{Out}(F_n)$ [BFH00], [BFH05], [BFH04].

In [HM, Section 2] (see also [BFH05, Propositions 3.16 and 4.41]) we proved that all elements of $\text{IA}_n(\mathbb{Z}/3)$ satisfy (1) - (3) above. (If an element of $\text{IA}_n(\mathbb{Z}/3)$ satisfies (4) then it is rotationless [FH, Lemma 3.12].) These invariance properties played a significant role in our series of papers [HM] establishing the ‘subgroup decomposition’ theorem for $\text{Out}(F_n)$ and then again in ([HM15], [HM17a]), in which the H_b^2 -alternative for $\text{Out}(F_n)$ is established: for every finitely generated subgroup $\mathcal{H} < \text{Out}(F_n)$ either \mathcal{H} is virtually abelian or $H_b^2(\mathcal{H}; \mathbb{R})$ has uncountably infinite dimension.

The main result of this paper is motivated in part by [HM15] and [HM17a], in which virtually abelian subgroups appear naturally and in which information is lost when one passes to finite index subgroups, and in part by our appreciation of the importance of $\text{IA}_n(\mathbb{Z}/3)$. Having seen that elements of $\text{IA}_n(\mathbb{Z}/3)$ satisfy (1) - (3) without iteration, one can ask analogously, which virtual properties of arbitrary subgroups of $\text{Out}(F_n)$ are true for subgroups of $\text{IA}_n(\mathbb{Z}/3)$ without passing to a subgroup of finite index? Our main theorem in this paper is one such property.

Theorem 1.1. *Each virtually abelian subgroup $\mathcal{H} < \text{IA}_n(\mathbb{Z}/3)$ is abelian.*

Abelian subgroups of $\text{Out}(F_n)$ are finitely generated and $\text{IA}_n(\mathbb{Z}/3)$ is torsion free; the former is contained in [BL94] and the latter follows from [BFH00, Corollary 5.7.6]. Thus,

Corollary 1.2. *Every virtually abelian subgroup of $\text{Out}(F_n)$ has a finitely generated, free abelian subgroup of index at most $|GL(n, \mathbb{Z}_3)| < 3^{n^2}$.*

In Section 2, after a brief review of PG and UPG subgroups, we reduce Theorem 1.1 to the following proposition.

Proposition 1.3. *Suppose that $K < \text{IA}_n(\mathbb{Z}/3)$ is an abelian UPG subgroup. Then the normalizer of K in $\text{IA}_n(\mathbb{Z}/3)$ equals the centralizer of K in $\text{IA}_n(\mathbb{Z}/3)$.*

All PG elements of $\text{IA}_n(\mathbb{Z}/3)$ are UPG [BFH00, Corollary 5.7.6] and, in fact, rotationless (Lemma 3.12). Therefore Proposition 1.3 may be equivalently restated using PG in place of UPG. The proof of Proposition 1.3 appears in Section 4.

Continuing with the theme of studying $\text{IA}_n(\mathbb{Z}/3)$, we pose the following question, the answer to which is yes if ϕ and ψ are rotationless by an easy application of [FH11, Theorem 5.3].

Question 1.4. *Are roots unique in the group $\text{IA}_n(\mathbb{Z}/3)$? That is, if $\phi, \psi \in \text{IA}_n(\mathbb{Z}/3)$ and $\phi^k = \psi^k$ for some $k \geq 1$, is $\phi = \psi$?*

Section 3 contains background material including subsections on UPG elements and UPG subgroups.

2 Reduction to Proposition 1.3

Each $\psi \in \text{Out}(F_n)$ has an associated finite set $\mathcal{L}(\psi)$ of attracting laminations, each of which is invariant under some iterate of ψ ; see Section 3.1, and in particular Definitions 3.15, of [BFH00]. (The role of attracting laminations in this paper is limited to the first paragraph of the proof below.) For a subgroup $\mathcal{H} < \text{Out}(F_n)$, we let $\mathcal{L}(\mathcal{H}) = \cup_{\psi \in \mathcal{H}} \mathcal{L}(\psi)$. If $\mathcal{L}(\psi) = \emptyset$, then we say that ψ has *polynomial growth* and write $\psi \in PG(F_n)$ or simply $\psi \in PG$. If in addition, the image of ψ in $\text{GL}(n, \mathbb{Z})$ is unipotent then we write $\psi \in UPG(F_n)$ or simply $\psi \in UPG$ [BFH00], [BFH05].

Proof of Theorem 1.1 assuming Proposition 1.3: Let $\mathcal{H} < \text{IA}_n(\mathbb{Z}/3)$ be virtually abelian. We first follow the proof of [BFH00, Theorem 7.0.1] to show that there is an exact sequence

$$1 \rightarrow K \rightarrow \mathcal{H} \rightarrow \mathbb{Z}^k \rightarrow 1$$

for some k and some abelian subgroup $K < UPG$. By [HM15, Lemma 4.8], $\mathcal{L}(\mathcal{H})$ is a finite collection $\{\Lambda_1, \dots, \Lambda_k\}$ of \mathcal{H} -invariant laminations. For each $1 \leq i \leq k$, let $\text{PF}_{\Lambda_i} : \text{Stab}(\Lambda_i) \rightarrow \mathbb{Z}$ be the expansion factor homomorphism for Λ_i as defined in [BFH00, Section 3.3]. Let $\text{PF} = \oplus_{i=1}^k \text{PF}_{\Lambda_i} : \mathcal{H} \rightarrow \mathbb{Z}^k$ be the direct sum of the restrictions to \mathcal{H} of the PF_{Λ_i} 's. If θ is an element of the kernel K of PF then $\mathcal{L}(\theta) \cap \mathcal{L}(\mathcal{H}) = \emptyset$ by [BFH00, Corollary 3.3.1]. Thus $\mathcal{L}(\theta) = \emptyset$ and K is PG. Applying our assumption that $\mathcal{H} < \text{IA}_n(\mathbb{Z}/3)$, we have that K is UPG by [BFH00, Corollary 5.7.6]. It then follows that K is solvable [BFH05, Corollary 1.3]. Since K is virtually abelian, it is finitely generated by [BL94] (see also [BFH04]). We can therefore apply [BFH04, Corollary 3.11] to conclude that K is abelian.

Proposition 1.3 implies that K is in the center of \mathcal{H} . In particular $[\psi_1, \psi_2]$ (which is an element of K) commutes with ψ_1 and ψ_2 for all $\psi_1, \psi_2 \in \mathcal{H}$. For all $p \geq 1$ and

all $\psi_1, \psi_2 \in \mathcal{H}$ we have $[\psi_1, \psi_2]^p = [\psi_1^p, \psi_2]$ and similarly $[\psi_1, \psi_2]^p = [\psi_1, \psi_2^p]$. For the first of these equations the inductive step is:

$$\begin{aligned} [\psi_1, \psi_2]^p &= [\psi_1, \psi_2]^{p-1} \psi_1 \psi_2 \psi_1^{-1} \psi_2^{-1} = \psi_1 [\psi_1, \psi_2]^{p-1} \psi_2 \psi_1^{-1} \psi_2^{-1} \\ &= \psi_1 [\psi_1^{p-1}, \psi_2] \psi_2 \psi_1^{-1} \psi_2^{-1} \\ &= \psi_1 \psi_1^{p-1} \psi_2 \psi_1^{1-p} \psi_2^{-1} \psi_2 \psi_1^{-1} \psi_2^{-1} = [\psi_1^p, \psi_2] \end{aligned}$$

Since \mathcal{H} is virtually abelian, there exists $p \geq 1$ such that $[\psi_1^p, \psi_2^p]$ is trivial. It follows that $[\psi_1, \psi_2]^{p^2}$ is trivial. Since finite order UPG elements are trivial [BFH05, Lemma 4.47], we conclude that $[\psi_1, \psi_2]$ is trivial for all $\psi_1, \psi_2 \in \mathcal{H}$. \square

3 Background

3.1 Basics

Much of the material in this subsection is standard and is included to establish notation and for convenient reference. Further details can be found in [HM, Section 2.2], [FH11, Section 2] or [HM, Section 1.1].

Marked graphs The free group F_n of rank n is identified with $\pi_1(R_n)$ where R_n is the graph with one vertex and n edges. A *marked n -graph* is a connected finite graph G of rank n that has no valence one vertices and is equipped with a homotopy equivalence $R_n \rightarrow G$ called a *marking* of G . The marking provides an identification of F_n with $\pi_1(G)$ that is well defined up to inner automorphism. A homotopy equivalence $f : G \rightarrow G$ determines an outer automorphism of $\pi_1(G)$ and hence an element $\phi \in \text{Out}(F_n)$ that we say is *represented by* $f : G \rightarrow G$.

Edges of G are assumed to be oriented with $\bar{E} = E^{-1}$ denoting the edge E with its orientation reversed. All of the $f : G \rightarrow G$ that we consider will take vertices to vertices and restrict to an immersion on each edge. A *direction* d at a vertex $v \in G$ is the germ of an oriented edge E with initial vertex v . Define the action of f on directions by $d \mapsto d'$ where d' is the direction determined by the first edge in $f(E)$.

We denote the universal cover of G by \tilde{G} and the set of ends of \tilde{G} by $\partial\tilde{G}$.

Fact 3.1. *Suppose that $f : G \rightarrow G$ is a homotopy equivalence of a marked graph. Then each lift $\tilde{f} : \tilde{G} \rightarrow \tilde{G}$ extends continuously over $\partial\tilde{G}$ by a homeomorphism $\hat{f} : \partial\tilde{G} \rightarrow \partial\tilde{G}$.*

For each $a \in F_n$, the inner automorphism $i_a : F_n \rightarrow F_n$ is defined by $i_a(g) = aga^{-1}$; the conjugacy class of a is denoted by $[a]$.

Fact 3.2. *Each $\Theta \in \text{Aut}(F_n)$ extends continuously to a homeomorphism $\widehat{\Theta} : \partial F_n \rightarrow \partial F_n$. For each non-trivial inner automorphism i_a , its boundary extension $\widehat{i_a}$ fixes two points, a source a^- and a sink a^+ .*

Fact 3.3. *For each marked graph G , the identification of $\pi_1(G)$ with F_n induces*

- (1) *an identification of the group of covering translations of \tilde{G} with F_n and*
- (2) *an identification of $\partial\tilde{G}$ with ∂F_n*

so that for each non-trivial $a \in F_n$, if T_a is the covering translation of \tilde{G} identified by (1) with a , and if A_a is the axis of T_a then the following hold: a^- and a^+ are identified by (2) with the repelling and attracting endpoints of A_a respectively; and the projection of A_a into G is a circuit that represents the conjugacy class $[a]$ of a . More generally, for any $\theta \in \text{Out}(F_n)$ and homotopy equivalence $f : G \rightarrow G$ representing θ , there is a bijection

$$\tilde{f} \leftrightarrow \Theta$$

between the set of lifts $\tilde{f} : \tilde{G} \rightarrow \tilde{G}$ and the set of $\Theta \in \text{Aut}(F_n)$ representing θ such that the homeomorphisms $\tilde{f} : \partial\tilde{G} \rightarrow \partial\tilde{G}$ and $\widehat{\Theta} : \partial F_n \rightarrow \partial F_n$ defined in Fact 3.1 and Fact 3.2 agree under identification (2).

Fact 3.4. *[BFH04, Lemma 2.4] Suppose that $\Phi \in \text{Aut}(F_n)$ and $a \in F_n$. If $\partial\Phi$ fixes either a^+ or a^- then $a \in \text{Fix}(\Phi)$ and $\partial\Phi$ fixes both a^+ and a^- .*

Principal lifts, rotationless outer automorphisms and rotationless maps
[FH11] We will only be interested in principal lifts and principal vertices in the UPG setting and so we can give simplified versions of their definitions.

Fact 3.5. *[GJLL98, Proposition I.1] For $\Theta \in \text{Aut}(F_n)$, we denote the fixed subgroup of Θ by $\text{Fix}(\Theta)$. There is a disjoint union*

$$\text{Fix}(\widehat{\Theta}) = \text{Fix}_-(\widehat{\Theta}) \cup \text{Fix}_+(\widehat{\Theta}) \cup \partial \text{Fix}(\Theta)$$

where $\text{Fix}_-(\widehat{\Theta}) \subset \partial F_n$ is a finite union of $\text{Fix}(\Theta)$ -orbits of repellers and $\text{Fix}_+(\widehat{\Theta}) \subset \partial F_n$ is a finite union of $\text{Fix}(\Theta)$ -orbits of attractors. $\text{Fix}_N(\Theta) \subset \partial F_n$ is defined to be $\partial \text{Fix}(\Theta) \cup \text{Fix}_+(\Theta)$.

Remark 3.6. In the special case that $\text{Fix}(\Theta) = \langle a \rangle$ and $\text{Fix}(\widehat{\Theta}) = \partial \text{Fix}(\Theta) = \{a^\pm\}$, it may happen that a^+ or a^- has an attracting neighborhood for the action of $\widehat{\Theta}$. This happens for example if $\Theta = i_a$. In all other cases, $\text{Fix}_+(\widehat{\Theta})$ is exactly the set of isolated attractors and $\text{Fix}_-(\widehat{\Theta})$ is exactly the set of isolated repellers.

Notation 3.7. The F_n -orbit of an element of $\text{Fix}_+(\widehat{\Theta})$ is called an *eigenray* for θ . The F_n -orbit of an element of $\text{Fix}_-(\widehat{\Theta})$ is an eigenray for θ^{-1} .

Definition 3.8. An automorphism Θ representing a UPG $\theta \in \text{Out}(F_n)$ is *principal* [FH11, Definition 3.1] if either $\text{Fix}_N(\Theta)$ contains at least three points or if $\text{Fix}_N(\Theta)$ is a two point set that is not $\{a^\pm\}$ for some $a \in F_n$ on ∂F_n . The set of principal Φ

representing θ is denoted by $P(\theta)$. An element $\phi \in \text{Out}(F_n)$ is *rotationless* [FH11, Definition 3.13] if: (i) $\Phi \mapsto \Phi^k$ defines a bijection between $P(\phi)$ and $P(\phi^k)$ for all $k \geq 1$; and (ii) $\text{Fix}_N(\Phi) = \text{Fix}_N(\Phi^k)$ for all $\Phi \in P(\phi)$ and all $k \geq 1$.

If $f : G \rightarrow G$ represents θ and $\tilde{f} : \tilde{G} \rightarrow \tilde{G}$ corresponds to $\Theta \in P(\theta)$ as in Fact 3.3 then we say that \tilde{f} is *principal*. An element x of the set $\text{Per}(f)$ of f -periodic points is *principal* [FH11, Definition 3.18] unless it is contained in a component C of $\text{Per}(f)$ that is topologically a circle and each point in C has exactly two periodic directions. If each principal vertex and periodic direction at a principal vertex has period one then we say that $f : G \rightarrow G$ is *rotationless*.

Paths An *edge path* σ in a marked graph G is a concatenation of edges $\sigma = \dots E_i E_{i+1} \dots$ of G where the terminal endpoint of E_i equals the initial endpoint of E_{i+1} for all i . If there is no backtracking, i.e. if $E_{i+1} \neq E_i^{-1}$ for all i , then we say that σ is a *path*. If a path $\sigma \subset G$ is a bi-infinite concatenation then we say that σ is a *line* in G . (All of the lines in this paper are oriented.) If a path $\sigma \subset G$ is a singly infinite concatenation then we say that σ is *ray*. We also allow the *trivial path* which is just a single vertex. Concatenation $\sigma = \sigma_1 \sigma_2$ of edge paths σ_1 and σ_2 is defined if σ_1 has a terminal vertex, σ_2 has an initial vertex and if these vertices are equal. The concatenation of paths need not be path.

Paths and edge paths in \tilde{G} are defined similarly. Edge paths in G lift to edge paths in \tilde{G} with paths lifting to paths. A line in G lifts to a line in \tilde{G} with well defined distinct ideal endpoints in $\partial \tilde{G}$. Conversely, every ordered pair of distinct points in $\partial \tilde{G}$ is the ideal endpoint pair for a unique line in \tilde{G} . A ray in G lifts to a ray in \tilde{G} with one endpoint at a vertex and the other an ideal endpoint in $\partial \tilde{G}$.

Suppose that $f : G \rightarrow G$ is a homotopy equivalence and $\tilde{f} : \tilde{G} \rightarrow \tilde{G}$ is a lift. For any finite path $\tilde{\sigma} \subset \tilde{G}$ with endpoints \tilde{x}, \tilde{y} , we define $\tilde{f}_\#(\tilde{\sigma})$ to be the unique path with endpoints $\tilde{f}(\tilde{x}), \tilde{f}(\tilde{y})$. We define $\tilde{f}_\#(\tilde{\sigma})$ for rays and lines similarly using \hat{f} if one or both endpoint is ideal. This descends to a well defined action $\sigma \mapsto f_\#(\sigma)$ of f on the set of paths in G .

A *circuit* in G is a cyclic concatenation of edges without backtracking and so can be viewed as an immersion of a circle. A circuit in G lifts to a line in \tilde{G} and we can extend the definition of $f_\#$ to include circuits. A closed path σ determines a circuit if the initial edges of σ and $\bar{\sigma}$ are distinct. If a circuit σ represents the conjugacy class $[a]$ of $a \in F_n$ and if $f : G \rightarrow G$ represents $\theta \in \text{Out}(F_n)$ then $f_\#(\sigma)$ represents $\theta([a])$.

A decomposition $\sigma = \dots \sigma_i \sigma_{i+1} \dots$ into subpaths is a *splitting* if

$$f_\#^k(\sigma) = \dots f_\#^k(\sigma_i) f_\#^k(\sigma_{i+1}) \dots$$

is a decomposition into subpaths for all $k \geq 1$. When the decomposition into σ_i 's is a splitting we write $\sigma = \dots \cdot \sigma_i \cdot \sigma_{i+1} \cdot \dots$

An *abstract line* is the F_n -orbit of an ordered pair of distinct points in ∂F_n . If G is any marked graph then the identification of ∂G with ∂F_n (Fact 3.3) defines a bijection between abstract lines and F_n -orbits of lines in \tilde{G} and so also a bijection

between abstract lines and lines in G . An *abstract ray* is an F_n -orbit of a point in ∂F_n . There is a bijection between abstract rays and equivalence classes of rays in G , where two rays in G are equivalent if they have a common infinite subray.

Free factor systems [BFH00, Section 2.6] If A_1, \dots, A_k are non-trivial free factors and $A_1 * \dots * A_k$ is a free factor of F_n then the set of conjugacy classes $\{[A_1], \dots, [A_k]\}$ is a *free factor system*. We write

$$\{[B_1], \dots, [B_l]\} \sqsubset \{[A_1], \dots, [A_k]\}$$

and say that $\{[B_1], \dots, [B_l]\}$ is contained in $\{[A_1], \dots, [A_k]\}$ if for each B_i there exists A_j so that some conjugate of B_i is a subgroup of A_j .

For every inclusion $H \subset G$ of a subgraph in a marked graph, there is an associated free factor system $\mathcal{F}(H) = \{[\pi_1(C_1)], \dots, [\pi_1(C_k)]\}$ where $\{C_1, \dots, C_k\}$ is the set of non-contractible components of G ; see [BFH00, Example 2.6.1] for details. We say that $H \subset G$ *realizes* $\mathcal{F}(H)$. Every free factor system is realized by some $H \subset G$ and every nested sequence $\mathcal{F}_1 \subset \mathcal{F}_2 \subset \dots \subset \mathcal{F}_l$ is realized by some nested sequence of subgraphs $H_1 \subset H_2 \subset \dots \subset H_l \subset G$. One may assume without loss that the H_i 's are *core* subgraphs, meaning that all vertices have valence at least two. If $\mathcal{F} \sqsubset \mathcal{F}'$ can be realized by core subgraphs $H \subset H'$ such that $H' \setminus H$ is a single edge then we say that $\mathcal{F} \sqsubset \mathcal{F}'$ is a *one-edge extension*; otherwise, $\mathcal{F} \sqsubset \mathcal{F}'$ is a *multi-edge extension*.

Fact 3.9. [BFH00, Section 2.6] Suppose that $[F]$ is a θ -invariant free factor conjugacy class and that $\Theta \in \text{Aut}(F_n)$ represents θ and preserves F . Then the element $\theta \mid F$ of $\text{Out}(F)$ determined by the restriction $\theta \mid F$ is independent of the choice of Θ .

A conjugacy class is *carried* by $[F]$ if some representative of it is an element of F . An abstract ray is *carried* by $[F]$ if it is represented by a point in ∂F . An abstract line is *carried* by $[F]$ if it is represented by an ordered pair of points, both of which are contained in ∂F . A conjugacy class, abstract ray or abstract line is carried by a free factor system \mathcal{F} if it is carried by a component of \mathcal{F} . If H is a subgraph of a marked graph G then a conjugacy class [resp. abstract line] is carried by $\mathcal{F}(H)$ if and only if the corresponding circuit [resp. line] in G is contained in H .

Fact 3.10. [HM17b, Fact 1.10] (see also [BH92, Section 2.6]) For any set X of abstract lines, abstract rays and conjugacy classes there is a unique minimal (with respect to \sqsubset) free factor system $\mathcal{F}_{\text{supp}}(X)$ that carries each element of X . If $\theta \in \text{Out}(F_n)$ and X is θ -invariant then $\mathcal{F}_{\text{supp}}(X)$ is θ -invariant.

3.2 UPG elements

In this section we review some facts about individual UPG elements of $\text{Out}(F_n)$.

A CT $f : G \rightarrow G$ is a particularly nice kind of topological representative of $\theta \in \text{Out}(F_n)$. The complete definition of a CT is given on [FH11, page 47]. Since

we will only use CT representatives in the special case when θ is UPG, the definition can be simplified considerably. Fact 3.11 and the proof of Lemma 3.12 give a pretty complete picture of CTs in this context.

We delay the proof that every UPG element is rotationless, and hence represented by a CT $f : G \rightarrow G$ [FH11, Theorem 4.28], until we have listed some properties enjoyed by such CTs.

A CT $f : G \rightarrow G$ is equipped with a filtration $\emptyset = G_0 \subset G_1 \subset \dots \subset G_N = G$ by f -invariant subgraphs. The subgraphs $H_r = G_r \setminus G_{r-1}$ are called the *strata*. A path has *height* r if it is contained in G_r and crosses at least one edge in H_r . The set of fixed points for f and the set of periodic points for f are denoted by $\text{Fix}(f)$ and $\text{Per}(f)$ respectively. The set of vertices of G is V . Recall that V is invariant under each $f : G \rightarrow G$ that we consider. .

Fact 3.11. *Each CT $f : G \rightarrow G$ representing $\theta \in \text{UPG}$ satisfies the following properties.*

- (1) *Each stratum H_i is a single edge E_i . If E_i is not fixed then there is a non-trivial closed path $u_i \subset G_{i-1}$ such that $f(E_i) = E_i u_i$.*
- (2) *$\text{Fix}(f) = \text{Per}(f)$ is the union of V with the set of fixed edges.*
- (3) *A direction based at a vertex is fixed if and only if it is periodic if and only if it is not the terminal direction of a non-fixed edge.*
- (4) *Each vertex is principal.*
- (5) *A lift $\tilde{f} : \tilde{G} \rightarrow \tilde{G}$ is principal if and only if $\text{Fix}(\tilde{f}) \neq \emptyset$.*

Proof. The strata of $f : G \rightarrow G$ are classified into three types: EG, NEG, and zero strata. From [BFH00, Lemma 3.1.9] and our assumption that $\mathcal{L}(\theta) = \emptyset$, it follows that $f : G \rightarrow G$ has no EG strata. The (Zero Strata) property of a CT therefore implies that $f : G \rightarrow G$ has no zero strata. Thus, every stratum of $f : G \rightarrow G$ is NEG. Item (1) therefore follows from [FH11, Lemma 4.21]. Items (2) and (3) follow from (1). A vertex that is incident to a fixed edge or is the terminal endpoint of a non-fixed edge is principal by the (Periodic Edges) and (Vertices) properties of a CT respectively. All other vertices are the initial endpoints of at least two non-fixed edges and so are principal by Definition 3.8. This proves (4). Item (5) follows from (4) and [FH11, Remark 4.8, Corollaries 3.17 and 3.27]. \square

A finite path $\sigma \subset G$ is a *Nielsen path* if $f_{\#}(\sigma) = \sigma$ and is an *indivisible Nielsen path* if there is no non-trivial decomposition of σ into Nielsen subpaths. Note that by Lemma 3.11(2), we would have the same set of indivisible Nielsen paths if we allowed paths to have endpoints that are not vertices.

In order to apply CT theory to UPG elements we must prove that they are rotationless. We will do this indirectly by using a result from [BFH00] to find a pretty

good relative train track map, namely one that satisfies various of the conclusions of Fact 3.11, and then we will quote [FH11, Proposition 3.29].

Lemma 3.12. *Each $\theta \in UPG$ is rotationless.*

Proof. By [BFH00, Proposition 5.7.5], θ is represented by a relative train track map $f : G \rightarrow G$ and filtration $\emptyset = G_0 \subset G_1 \subset \dots \subset G_N = G$ with a subsequence of invariant core subgraphs $\emptyset = G_0 = G_{r(0)} \subset G_{r(1)} \subset \dots \subset G_{r(m)} = G$ such that $G_{r(j+1)}$ is obtained from $G_{r(j)}$ in one of the following ways.

- (a) Adding a single fixed edge that is either a loop or has both endpoints in $G_{r(j)}$; $r(j+1) = r(j) + 1$.
- (b) Adding a single non-fixed edge satisfying Fact 3.11(1) with both endpoints in $G_{r(j)}$; $r(j+1) = r(j) + 1$.
- (c) Adding two non-fixed edges satisfying Fact 3.11(1) with a common initial vertex not in $G_{r(j)}$ and both terminal endpoints in $G_{r(j)}$; $r(j+1) = r(j) + 2$.

It is obvious from (a) - (c) that Fact 3.11(1) is satisfied. This implies items (2) and (3) of Fact 3.11, which in turn prove that $f : G \rightarrow G$ is rotationless (Definition 3.8). By [FH11, Proposition 3.29], we are reduced to showing that $f : G \rightarrow G$ satisfies the five properties listed in [FH11, Theorem 2.19]. Property (Z) applies only to zero strata and so is vacuous in this context. Properties (F) and (NEG) are immediate from (a) - (c). The endpoints of an indivisible Nielsen path are not contained in the interior of a fixed edge and so are vertices by Fact 3.11(1). This verifies the (V) property of [FH11, Theorem 2.19]. If a stratum H_m is a forest in $\text{Per}(f)$ then it is a single fixed edge E_m with endpoints in a core subgraph G_{m-1} by (a) - (c). It follows that the free factor support of G_{m-1} is not equal to the free factor support of $G_l \cup E_m$ for any filtration element G_l . This verifies the (P) property of [FH11, Theorem 2.19] and we are done. \square

We assume for the rest of this subsection that

- $\theta \in UPG$ and that $f : G \rightarrow G$ is a CT representing θ , hence f satisfies the conclusion of Lemma 3.11.

If a root-free $a \in F_n$ is fixed by $m \geq 2$ elements of $\mathcal{P}(\theta)$ then its unoriented conjugacy class $[a]_u$ is called an *axis* or *twistor* for θ with *multiplicity* $m-1$. An edge E in a stratum H_i is *linear* if there is a Nielsen path $u \subset G_{i-1}$ such that $f(E) = Eu$. Recall from Fact 3.3 that each non-trivial $a \in F_n$ corresponds to a covering translation $T_a : \tilde{G} \rightarrow \tilde{G}$ with axis A_a .

Fact 3.13. *For each root-free $a \in F_n$, if $[a]_u$ is a twistor for θ of multiplicity $m \geq 2$ then :*

- (1) *There is a closed path w that determines a circuit representing $[a]_u$.*

- (2) There are exactly $m - 1$ linear edges E^1, \dots, E^{m-1} such that $f(E^l) = E^l w^{d_l}$ for some $d_l \neq 0$. Furthermore, the values d_l are pairwise distinct. We say that w is the twist path for E^l .
- (3) For each $l = 1, \dots, m - 1$, there is a lift \tilde{f}_l of f such that for each lift \tilde{E}^l , if the terminal endpoint of \tilde{E}^l is contained in A_a then the initial endpoint of \tilde{E}^l is fixed by \tilde{f}_l . These lifts are pairwise distinct, preserve A_a and each acts without fixed points on A_a .
- (4) For each $l = 1, \dots, m - 1$, the lift \tilde{f}_l corresponds to an element $\Theta_l \in \mathcal{P}(\theta)$ that fixes a . These automorphisms Θ_l account for all but one element $\Theta_0 \in \mathcal{P}(\theta)$ that fixes a . The lift \tilde{f}_0 of f that corresponds to Θ_0 fixes points in A_a .

Proof. The (NEG Nielsen Paths) property of a CT implies that $\text{Fix}(\tilde{f}_i) \cap A_a = \emptyset$. The rest of the fact follows from the (Linear Edges) property of a CT and [FH11, Lemma 4.40]. \square

Fact 3.14. *If E is a linear edge with twist path w then every occurrence of E in a Nielsen path ρ for f is contained in a subpath of ρ of the form $Ew^p\bar{E}$.*

Proof. We may assume that ρ is indivisible. Let r be the height of ρ and s the height of E . The $r < s$ case is vacuous and the $r = s$ case follows from the (NEG Nielsen Paths) property for a CT. If $r > s$ then the (NEG Nielsen Paths) property implies that $\sigma = E'w'^p\bar{E}'$ for some $p \neq 0$, where E' is a linear edge of height r and its twist path w' has height less than r . Each occurrence of E in σ is contained in w'^p and we are done by induction. \square

Notation 3.15. Letting E be a non-fixed edge of height r with $f(E) = Eu$ for some closed path u of height $< r$ [Fact 3.11(1)], its iterates split as $f^k(E) = E \cdot u \cdot f_\#(u) \cdot \dots \cdot f_\#^{k-1}(u)$. In this case, the nested sequence $E \subset f(E) \subset f^2(E) \subset \dots$ converges to a ray R_E that we say is determined by E . If E is a linear edge with twist path w then $R_E = Ew^{\pm\infty}$. If E is non-linear then the set of terminal endpoints of lifts of R_E to \tilde{G} is an F_n -orbit in ∂F_n that we denote $[\partial R_E]$. Note that $R_E \setminus E$ is a ray in G_{r-1} and so $[\partial R_E]$ is carried by the free factor system determined by G_{r-1} .

Fact 3.16. *The assignment $E \mapsto [\partial R_E]$ defines a bijection between the set \mathcal{E} of non-linear, non-fixed edges of G and the set of eigenrays of θ (Notation 3.7).*

Proof. This is contained in [HM17b, Fact 1.49]; see also [FH11, Lemma 4.36]. \square

Corollary 3.17. *If $\mathcal{F} \sqsubset \{[F_n]\}$ is a θ -invariant one-edge extension then \mathcal{F} carries every twistor and eigenray for θ .*

Proof. By [FH11, Theorem 4.5], there exists a CT $f : G \rightarrow G$ representing θ in which \mathcal{F} is represented by a filtration element G_s . Each non-fixed edge E above G_s satisfies $f(E) = E \cdot u$ for some non-trivial path $u \subset G_s$. The corollary therefore follows from Fact 3.13, Notation 3.15 and Fact 3.16. \square

Fact 3.18. $\text{Fix}_N(\Theta) \neq \emptyset$ for all $\Theta \in \text{Aut}(F_n)$ representing θ .

Proof. Let $\tilde{f} : \tilde{G} \rightarrow \tilde{G}$ be the lift of f corresponding to Θ . If Θ is principal then $\text{Fix}_N(\Theta) \neq \emptyset$ by definition. We may therefore assume that Θ is not principal and hence by Fact 3.11(5) that \tilde{f} is fixed point-free. In this case there is a path $\tilde{\sigma} \subset \tilde{G}$ such that $\tilde{\sigma} \cdot \tilde{f}_\#(\tilde{\sigma}) \cdot \tilde{f}_\#^2(\tilde{\sigma}) \cdot \dots$ converges to a point in $\text{Fix}_N(\Theta)$. The construction of $\tilde{\sigma}$ is carried out in the proof of [BFH00, Proposition 5.4.3]. A more directly quotable reference is [FH, Lemma 6.4]. \square

Fact 3.19. [BFH05, Proposition 4.44] If the conjugacy class of the free factor F is θ -invariant then $\theta|F$ is UPG.

The following lemma is not known for elements of $\text{Out}(F_n)$ that are not UPG.

Lemma 3.20. $\Theta \in \mathcal{P}(\theta) \iff \Theta^{-1} \in \mathcal{P}(\theta^{-1})$.

Proof. By symmetry, it suffices to assume that $\Theta \in \mathcal{P}(\theta)$ and prove that $\Theta^{-1} \in \mathcal{P}(\theta^{-1})$. If the rank of $\text{Fix}(\Theta)$ is at least two then this follows from the Definition 3.8 and the fact that $\text{Fix}(\Theta) = \text{Fix}(\Theta^{-1})$. We may therefore assume that $\text{Fix}(\Theta)$ has rank one or zero. We show below that there is an injective map $\text{Fix}_+(\widehat{\Theta}) \rightarrow \text{Fix}_+(\widehat{\Theta^{-1}})$. Assuming this for now, we complete the proof as follows. The cardinality of $\text{Fix}_+(\widehat{\Theta^{-1}})$ is at least one in the rank one case and at least two in the rank zero case. It follows that $\text{Fix}_N(\widehat{\Theta^{-1}})$ contains at least three points unless $\text{Fix}(\Theta^{-1})$ has rank zero and $\text{Fix}_+(\widehat{\Theta^{-1}})$ contains exactly two points. In this case, Fact 3.4 implies that $\text{Fix}_+(\widehat{\Theta^{-1}}) \neq \{a^\pm\}$ for any non-trivial $a \in F_n$ so $\Theta^{-1} \in \mathcal{P}(\theta^{-1})$ in this case as well.

It remains to show that there is an injective map $\text{Fix}_+(\widehat{\Theta}) \rightarrow \text{Fix}_+(\widehat{\Theta^{-1}})$. The lift $\tilde{f} : \tilde{G} \rightarrow \tilde{G}$ corresponding to Θ satisfies $\hat{f} = \widehat{\Theta}$. For each $P \in \text{Fix}_+(\hat{f})$ there is (Fact 3.16) a non-fixed non-linear edge E , a lift \tilde{E} of E and a lift \tilde{R}_E of R_E (Notation 3.15) with initial edge \tilde{E} , such that \tilde{R}_E converges to P and intersects $\text{Fix}(\hat{f})$ only in its initial endpoint \tilde{v} [FH11, Lemma 3.36]. Let r be the height of E . By Lemma 3.11(1) there is a component C of G_{r-1} that contains the terminal endpoint of E and hence contains all of R_E but its first edge. Let $\Gamma \subset \tilde{G}$ be the component of the full pre-image C that contains the terminal endpoint of the initial edge \tilde{E} of \tilde{R}_E . Then Γ is \hat{f} -invariant and $\partial\Gamma$ contains P . The (NEG Nielsen Paths) property of a CT implies that $\hat{f} \mid \Gamma$ is fixed point free. [FH11, Lemma 3.16] therefore implies that P is the only element of $\text{Fix}_+(\hat{f})$ contained in $\partial\Gamma$. Since $\text{Fix}_+(\hat{f})$ is \hat{i}_a -invariant for each $a \in \text{Fix}(\Theta)$, it follows that $\partial\text{Fix}(\Theta) \cap \partial\Gamma = \emptyset$.

Let F be the free factor that represents the unique element of $\mathcal{F}(C)$ and satisfies $\partial F = \partial\Gamma$. The automorphism $\Psi := \Theta^{-1} \mid F$ represents the restriction $\psi = \theta^{-1} \mid F$, which is UPG by Fact 3.19. By Fact 3.18 there exists at least one point $Q \in \text{Fix}_N(\Psi)$. Since $\partial\text{Fix}(\Psi) = \partial\text{Fix}(\Theta) \cap \partial\Gamma = \emptyset$, $Q \in \text{Fix}_+(\widehat{\Psi})$.

To see that $P \mapsto Q$ is injective, suppose that $P' \neq P$ is a point in $\text{Fix}_+(\widehat{\Theta})$ and that C', Γ', F' and Q' are defined as above with P replaced by P' . Since C and C' are

components of filtration elements of G , either they are disjoint or one is contained in the other. It follows [HM17b, Fact 1.2] that either ∂F and $\partial F'$ are disjoint or one is contained in the other. The latter is ruled out by the fact that $P = \text{Fix}_+(\Theta) \cap \partial F$ and $P' = \text{Fix}_+(\Theta) \cap \partial F'$. Thus ∂F and $\partial F'$ are disjoint and $Q \neq Q'$. \square

Definition 3.21. Every path $\tilde{\sigma} \subset \tilde{G}$ with endpoints, if any, at vertices has a *highest edge splitting* $\tilde{\sigma} = \dots \tilde{\sigma}_{-1} \cdot \tilde{\sigma}_0 \cdot \tilde{\sigma}_1 \dots$ defined as follows. If r is the height of σ and E_r is not fixed then this splitting is defined by taking the splitting vertices (i.e. the endpoints of the terms) to be exactly those vertices that are either the initial endpoint of an edge in $\tilde{\sigma}$ that projects to E_r or the terminal endpoint of an edge that projects to \tilde{E}_r . If E_r is fixed then both endpoints of an edge that projects to E_r or \tilde{E}_r are splitting vertices. The projected splitting $\sigma = \dots \sigma_{-1} \cdot \sigma_0 \cdot \sigma_1 \dots$ is the *highest edge splitting* of σ .

Fact 3.22. (1) *The highest edge splitting $\sigma = \dots \sigma_{-1} \cdot \sigma_0 \cdot \sigma_1 \dots$ of σ is in fact a splitting.*

(2) *For any lift \tilde{f} ,*

$$\tilde{f}_\#(\tilde{\sigma}) = \dots \tilde{f}_\#(\tilde{\sigma}_{-1}) \cdot \tilde{f}_\#(\tilde{\sigma}_0) \cdot \tilde{f}_\#(\tilde{\sigma}_1) \dots$$

is the highest edge splitting of $\tilde{f}_\#(\tilde{\sigma})$.

Proof. Item (1) is contained in the statement and proof of [BFH00, Lemma 4.1.4].

For (2), let $\mathcal{V}_{\tilde{\sigma}}$ and $\mathcal{V}_{\tilde{f}_\#(\tilde{\sigma})}$ be the highest edge splitting vertices of $\tilde{\sigma}$ and $\tilde{f}_\#(\tilde{\sigma})$ respectively. Assuming at first that E_r is not fixed, each term $\tilde{\sigma}_j$ in the highest edge splitting of $\tilde{\sigma}$ have the form $\tilde{E}_r \tilde{\gamma} \tilde{E}_r^{-1}$, $\tilde{E}_r \tilde{\gamma}$, $\tilde{\gamma} \tilde{E}_r^{-1}$, $\tilde{\gamma}$, \tilde{E}_r or \tilde{E}_r^{-1} for some non-trivial path $\tilde{\gamma}$ that projects into G_{r-1} . Since $f(E_r) = E_r u_r$ for some path $u_r \subset G_{r-1}$, the $\tilde{f}_\#$ -image of each of these types is another path of the same type. It follows that $\mathcal{V}_{\tilde{f}_\#(\tilde{\sigma})} \subset \tilde{f}(\mathcal{V}_{\tilde{\sigma}})$. It also follows that if $\tilde{\sigma}_j$ ends with \tilde{E}_r^{-1} [respectively begins with \tilde{E}_r] then $\tilde{f}_\#(\tilde{\sigma}_j)$ ends with \tilde{E}_r^{-1} [respectively begins with \tilde{E}_r]. This implies that $\tilde{f}(\mathcal{V}_{\tilde{\sigma}}) \subset \mathcal{V}_{\tilde{f}_\#(\tilde{\sigma})}$. This completes the proof in the case that E_r is not fixed. The remaining case is similar and is left to the reader. \square

Lemma 3.23. *Suppose that \tilde{f} is a lift of $f : G \rightarrow G$, that $\tilde{\mu}$ is an $\tilde{f}_\#$ -invariant line that is disjoint from $\text{Fix}(\tilde{f})$ and that an endpoint of $\tilde{\mu}$ is fixed by a covering translation T . Then $\tilde{\mu}$ is the axis A_T of T .*

Proof. Let $\tilde{\mu} = \dots \tilde{\mu}_{-1} \cdot \tilde{\mu}_0 \cdot \tilde{\mu}_1 \dots$ be the highest edge splitting of μ . Fact 3.22 implies that there exists $p \in \mathbb{Z}$ such that $(\tilde{f})_\#(\tilde{\mu}_i) = \tilde{\mu}_{i+p}$ for all $\tilde{\mu}_i$. From the assumption that $\text{Fix}(\tilde{f}) \cap \tilde{\mu} = \emptyset$, it follows that $p \neq 0$ and so the splitting is bi-infinite. The highest edge splitting $A_T = \dots \tilde{\alpha}_{-1} \cdot \tilde{\alpha}_0 \cdot \tilde{\alpha}_1 \dots$ of A_T is also bi-infinite. Since $\tilde{\mu}$ and A_T have a common ray, they must have the same height. After re-indexing the $\tilde{\alpha}_j$'s, we may assume that $\tilde{\alpha}_j = \tilde{\mu}_j$ for all sufficiently large j . It follows that $(\tilde{f})_\#(\tilde{\alpha}_j) = \tilde{\alpha}_{j+p}$ for all $\tilde{\alpha}_j$ and hence that for all j there exists $k > 0$ such that $(\tilde{f})_\#^k(\tilde{\mu}_j) = (\tilde{f})_\#^k(\tilde{\alpha}_j)$. Since μ_j and α_j are paths in G with the same endpoints, it follows that $\tilde{\alpha}_j = \tilde{\mu}_j$ for all j and $\tilde{\sigma} = A_T$. \square

3.3 Abelian UPG subgroups

We assume throughout this section that K is an abelian UPG subgroup of $\text{Out}(F_n)$. Lemma 3.12 implies that each element of K is rotationless and so K is a rotationless abelian UPG subgroup.

The main definitions in [FH09] make use of

$$\mathcal{P}^\pm(\theta) := \mathcal{P}(\theta) \cup (\mathcal{P}(\theta^{-1}))^{-1}$$

In the UPG case, Lemma 3.20 implies that

$$\mathcal{P}^\pm(\theta) = \mathcal{P}(\theta) = \mathcal{P}(\theta^{-1})^{-1}$$

We have simplified the definitions in this subsection accordingly.

Remark 3.24. As noted in [HM17a, Section 6.1.2], the definition of $\mathcal{P}^\pm(\theta)$ was misstated in [FH11] as $\mathcal{P}^\pm(\theta) := \mathcal{P}(\theta) \cup \mathcal{P}(\theta^{-1})$.

Definition 3.25. [FH09, Definition 3.9] A set $X \subset \partial F_n$ with at least three points is a *principal set* for K if for each $\phi \in K$ there exists $\Phi \in \mathcal{P}(\phi)$ such that $X \subset \text{Fix}(\Phi)$. For each such X , the assignment $\phi \rightarrow \Phi$ defines a lift s_X of K into $\text{Aut}(F_n)$ called the *principal lift* determined by X .

If X is a principal set and $X' \subset X$ contains at least three points then X' is a principal set and $s_{X'} = s_X$. For any principal set X' , the maximal (with respect to inclusion) principal set containing X' is given by

$$X = \cap_{\phi \in K} \text{Fix}(\widehat{s_{X'}(\phi)})$$

See [FH09, Remark 3.10].

Automorphisms that differ by conjugation by an inner automorphism are said to be *isogredient*. If X is a maximal principal set and $c \in F_n$ then $i_c(X)$ is also a maximal principal set and $s_{i_c(X)}(\phi) = i_c s_X(\phi) i_c^{-1}$ for each $\phi \in K$. Thus F_n -orbits of maximal principal sets correspond to isogredience classes of principal lifts.

The following definition generalizes to the setting of UPG abelian subgroups the concepts that were defined just preceding Fact 3.13.

Definition 3.26. [FH09, Definition 4.1] We say that the unoriented conjugacy class $[a]_u$ of $a \in F_n$ is an *axis* or *twistor* of multiplicity $m - 1 \geq 1$ for K and write $[a]_u \in \mathcal{A}(K)$ if $\{a^\pm\}$ is contained in m distinct maximal principal sets. The maximal principal sets that contain $\{a^\pm\}$ are called *linear principal sets* or more specifically *a -linear principal sets*. If X_1 and X_2 are distinct a -linear principal sets then for each $\theta \in K$ there exists an integer $d(\theta)$ such that $s_{X_2}(\phi) = i_a^{d(\theta)} s_{X_1(\theta)}$. The assignment $\theta \mapsto d(\theta)$ defines a homomorphism $\omega : K \rightarrow \mathbb{Z}$ called the *comparison homomorphism determined by X_1 and X_2* . Note that ω depends only on the F_n -orbit of the pair (X_1, X_2) ; i.e. (X_1, X_2) and $i_c(X_1, X_2) := (i_c X_1, i_c X_2)$ determine the same comparison homomorphism.

Fact 3.27. [FH09, Lemma 4.3] *There are only finitely many comparison homomorphisms for K .*

Fact 3.28. [FH09, Lemma 4.6] *If $\theta, \phi \in K$ and $\omega(\theta) = \omega(\phi)$ for all comparison homomorphisms ω then $\theta = \phi$.*

Definition 3.29. [FH09, Definition 4.7] $\phi \in K$ is generic if $\omega(\phi) \neq 0$ for each comparison homomorphism ω .

Fact 3.30. [FH09, Lemma 4.10] *If $\theta \in K$ is generic then $\{\text{Fix}(\widehat{\Theta}) : \Theta \in \mathcal{P}(\theta)\}$ is the set of maximal principal sets for K .*

Corollary 3.31. *If $\phi, \theta \in K$ are generic then $\{\text{Fix}(\widehat{\Theta}) : \Theta \in \mathcal{P}(\theta)\} = \{\text{Fix}(\widehat{\Phi}) : \Phi \in \mathcal{P}(\phi)\}$.*

Corollary 3.32. *$[a]_u \in F_n$ is a twistor for K of multiplicity $m-1$ if and only if $[a]_u$ is a twistor of multiplicity $m-1$ for some, and hence every, generic element of K .*

Fact 3.33. [FH09, Lemma 4.9] *K has a basis of generic elements.*

Fact 3.34. [FH09, Lemma 2.6] *Suppose that $\theta, \psi \in \text{Out}(F_n)$, that $\phi := \theta^\psi = \psi\theta\psi^{-1}$ and that $\Psi \in \text{Aut}(F_n)$ represents ψ . Then*

- (1) $\text{Fix}(\widehat{\Psi\Theta\Psi^{-1}}) = \widehat{\Psi}(\text{Fix}(\widehat{\Theta}))$ for all $\Theta \in \text{Aut}(F_n)$ representing θ .
- (2) $\text{Fix}_+(\widehat{\Psi\Theta\Psi^{-1}}) = \widehat{\Psi}(\text{Fix}_+(\widehat{\Theta}))$ and $\text{Fix}_-(\widehat{\Psi\Theta\Psi^{-1}}) = \widehat{\Psi}(\text{Fix}_-(\widehat{\Theta}))$ for all $\Theta \in \text{Aut}(F_n)$ representing θ .
- (3) $\Theta \mapsto \widehat{\Psi\Theta\Psi^{-1}}$ defines a bijection between $\mathcal{P}(\theta)$ and $\mathcal{P}(\phi)$ that preserves isogredience classes.

Lemma 3.35. *If θ is generic in K and $\psi \in \text{Out}(F_n)$ then θ^ψ is generic in $K^\psi = \{\psi\phi\psi^{-1} : \phi \in K\}$. Moreover, for any $\Psi \in \text{Aut}(F_n)$ representing ψ , $\widehat{\Psi}$ induces a bijection between $[a\text{-linear}]$ principal sets in K and $[\Psi(a)\text{-linear}]$ principal sets in K^ψ .*

Proof. Choose $\Psi \in \text{Aut}(F_n)$ representing ψ . The following are easy consequences of Fact 3.34:

- $\mathcal{A}(K^\psi) = \psi(\mathcal{A}(K))$.
- If X_i is a $[a\text{-linear}]$ principal set for K then $\widehat{\Psi}(X_i)$ is a $[\Psi(a)\text{-linear}]$ principal set for K^ψ .
- If $\omega : K \rightarrow \mathbb{Z}$ is the comparison homomorphism determined by X_1 and X_2 then the comparison homomorphism $\omega^\psi : K^\psi \rightarrow \mathbb{Z}$ determined by $\widehat{\Psi}(X_1)$ and $\widehat{\Psi}(X_2)$ satisfies $\omega^\psi(\theta^\psi) = \omega(\theta)$.

The lemma now follows from the definition of genericity. \square

Lemma 3.36. *If $\theta \in K$ is generic in K and $[F]$ is a K -invariant free factor conjugacy class then $K \mid F$ (see Fact 3.9) is an abelian UPG subgroup and $\theta \mid F$ is generic in $K \mid F$.*

Proof. Fact 3.19 implies that $K \mid F$ is an abelian UPG subgroup. Each maximal principal set for $K \mid F$ extends uniquely to a maximal principal set for K . It follows that each coordinate homomorphism ω_K for $K \mid F$ is the restriction of a coordinate homomorphism ω for K . Thus each $\omega_K(\theta \mid K) = \omega(\theta) \neq 0$ proving that $\theta \mid K$ is generic. \square

4 Proof of Proposition 1.3

Assuming that $K < \text{IA}_n(\mathbb{Z}/3)$ is an abelian UPG subgroup and that $\psi \in \text{IA}_n(\mathbb{Z}/3)$ normalizes K , our goal is to show that ψ commutes with each $\theta \in K$. By Facts 3.12 and 3.33, K is rotationless and has a basis of generic elements. We may therefore assume that θ is generic. Letting $\phi = \theta^\psi = \psi\theta\psi^{-1} \in K$, our goal is to show that $\phi = \theta$.

Suppose that $[a]_u \in \mathcal{A}(K)$ (see Definition 3.26) and that Ψ is a representative of ψ . Lemma 3.35 implies that $[\Psi(a)]_u \in \mathcal{A}(K^\psi) = \mathcal{A}(K)$ and hence that ψ permutes the elements of $\mathcal{A}(K)$. It follows that ψ fixes each element of $\mathcal{A}(K)$ by [HM17c, Theorem 4.1]. We can therefore choose Ψ_a representing ψ that fixes a . Lemma 3.35 implies that if X_i is an a -linear principal set for K then $\widehat{\Psi_a}(X_i)$ is an a -linear principal set for $K^\psi = K$. Thus

- (*) For each $a \in F_n$ with $[a]_u \in \mathcal{A}(K)$ there exists $\Psi_a \in \text{Aut}(F_n)$ representing ψ such that $\Psi_a(a) = a$. Letting $\{X_0, \dots, X_{m-1}\}$ be the a -linear principal sets, the automorphism Ψ_a induces a permutation π_a of $\{0, \dots, m-1\}$ such that $\widehat{\Psi_a}(X_i) = X_{\pi_a(i)}$. Corollary 3.32 implies that $[a]_u$ is twistor for θ with multiplicity $m-1$.

The main work of the proof is to show that each π_a is the identity. Assuming this fact for the moment, we complete the proof of the proposition.

Fix $[a]_u \in \mathcal{A}(K)$ and let $\{X_0, \dots, X_{m-1}\}$, Ψ_a and π_a be as in (*). Let $\Theta_i = s_{X_i}(\theta) \in \mathcal{P}(\theta)$ (see Definition 3.25) and $\Phi_i = \Psi_a \Theta_i \Psi_a^{-1} \in \mathcal{P}(\phi)$. Then

$$\text{Fix}(\widehat{\Phi}_i) = \widehat{\Psi_a}(\text{Fix}(\widehat{\Theta}_i)) = \widehat{\Psi_a}(X_i) = X_{\pi_a(i)} = X_i$$

with the middle equality following from Fact 3.30 and the genericity of θ . Thus $s_{X_i}(\phi) = \Phi_i$. If X_i and X_j are distinct a -linear principal sets then

$$\Theta_j = i_a^d \Theta_i$$

for some $d \neq 0$ and

$$\Phi_j = \Psi_a \Theta_j \Psi_a^{-1} = \Psi_a i_a^d \Theta_i \Psi_a^{-1} = i_{\Psi_a(a)}^d \Psi_a \Theta_i \Psi_a^{-1} = i_a^d \Phi_i$$

This proves that $\omega(\theta) = \omega(\phi)$ where ω is the comparison homomorphism determined by X_i and X_j . Since a, i and j are arbitrary, Fact 3.28 completes the proof of Proposition 1.3.

Fixing a as in $(*)$, it remains to prove that π_a is the identity.

Lemma 4.1. *Suppose that $\emptyset = \mathcal{F}_0 \sqsubset \mathcal{F}_1 \sqsubset \dots \sqsubset \mathcal{F}_J = \{[F_n]\}$ is a maximal nested sequence of free factor systems that are invariant by both K and ψ and that $[F]$ is a component of \mathcal{F}_j for some $1 \leq j \leq J$. Then $[F]$ is invariant by both K and ψ ; moreover, for all $\eta \in K$, each axis and eigenray for $\eta \mid F$ is carried by \mathcal{F}_{j-1} .*

Proof. Since $K, \psi \subset \text{IA}_n(\mathbb{Z}/3)$, [HM17c, Lemma 4.2] implies that $[F]$ is invariant by both K and ψ . By Fact 3.19, $K_F := K \mid F$ is a UPG subgroup that is obviously normalized by $\psi_F := \psi \mid F$. Each component of \mathcal{F}_{j-1} is contained in a unique component of \mathcal{F}_j . The union of the components of \mathcal{F}_{j-1} that are contained in $[F]$ define a free factor system of F that we denote by \mathcal{F}' . Since $\mathcal{F}_{j-1} \sqsubset \mathcal{F}_j$ is invariant by both K and ψ and is maximal with respect to these properties, it follows that $\mathcal{F}' \sqsubset \{[F]\}$ is invariant by both K_F and ψ_F and is maximal with respect to these properties. By [BFH05, Theorem 5.1] there is a K_F -invariant free factor system \mathcal{F}'' of F such that $\mathcal{F}' \sqsubset \mathcal{F}''$ and such that $\mathcal{F}'' \sqsubset \{[F]\}$ is a one-edge extension. By Corollary 3.17, \mathcal{F}'' carries each axis and eigenray for all $\eta_F \in K_F$. It follows from Fact 3.34 and the definitions that ψ maps the axes and eigenrays of $\eta_F \in K_F$ to the axes and eigenrays of $\eta_F^{\psi_F} \in K_F$ and so the set X of axes and eigenrays that occur for some element of K_F is ψ_F -invariant. For the same reason, X is K -invariant. By Fact 3.10, $\mathcal{F}_{\text{supp}}(X \cup \mathcal{F}')$ is invariant by both K and ψ . Moreover, $\mathcal{F}_{\text{supp}}(X \cup \mathcal{F}')$ contains \mathcal{F}' by construction and is properly contained in $\{[F]\}$ because it is contained in \mathcal{F}'' . It follows that $\mathcal{F}_{\text{supp}}(X \cup \mathcal{F}') = \mathcal{F}'$ and so $\mathcal{F}_{\text{supp}}(X) \sqsubset \mathcal{F}' \sqsubset \mathcal{F}_{j-1}$. \square

Having fixed a as in $(*)$, the notations of $(*)$ remain in force for the rest of the paper, as do Notations 4.2 and 4.4 below. Recall also that θ is a generic element of K , that ψ normalizes K and that $\phi = \theta^\psi$.

Notation 4.2. Let $\emptyset = \mathcal{F}_0 \sqsubset \mathcal{F}_1 \sqsubset \dots \sqsubset \mathcal{F}_J = \{[F_n]\}$ be a maximal nested sequence of free factor systems that are invariant by both K and ψ . By [FH, Theorem 1.1] there is a CT $f : G \rightarrow G$ representing θ with filtration $\emptyset = G_0 \subset G_1 \subset \dots \subset G_N = G$ and a subfiltration by core graphs $\emptyset = G_0 = G_{r(0)} \subset G_{r(1)} \subset \dots \subset G_{r(J)} = G$ such each \mathcal{F}_j is realized by $G_{r(j)}$; moreover, for each j and each component C of $G_{r(j)}$, $f|C$ is a CT representing $\theta \mid [\pi_1(C)]$. For each $1 \leq j \leq J$, define the j^{th} -stratum of the subfiltration to be the subgraph $S_j = G_{r(j)} \setminus G_{r(j-1)}$. Choose $h : G \rightarrow G$ representing ψ such that each $G_{r(j)}$ is h -invariant.

Corollary 4.3. *If E_i is a non-fixed edge in S_j then $f(E_i) = E_i \cdot u_i$ for some non-trivial closed path $u_i \subset G_{r(j-1)}$.*

Proof. Let C be the component of $G_{r(j)}$ that contains E_i and let F be a free factor representing $[\pi_1(C)]$. By construction, $f|C$ is a CT representing $\theta|F$. If E_i is linear then $u_i = w_i^{p_i}$ where w_i is the twist path for E_i and $[w_i]_u$ is an axis for $\theta|F$. In this case, Lemma 4.1 completes the proof. If E_i is non-linear then ∂R_{E_i} is an eigenray for $\theta|F$ by Fact 3.16. Lemma 4.1 therefore implies that $f_{\#}^k(u_i) \subset G_{r(j-1)}$ for some $k > 0$ which implies that $u_i \subset G_{r(j-1)}$. \square

Notation 4.4. For $0 \leq i \leq m-1$, let $\Theta_i = s_{X_i}(\theta)$, let $\tilde{f}_i : \tilde{G} \rightarrow \tilde{G}$ be the lift corresponding to Θ_i and let $\hat{f}_i : \partial F_n \rightarrow \partial F_n$ be the extension of \tilde{f}_i given by Fact 3.1. Then $\Theta_0, \dots, \Theta_{m-1}$ are the elements of $\mathcal{P}(\theta)$ that fix a and $\text{Fix}(\hat{f}_i) = X_i$ by Facts 3.3 and 3.30. After possibly reindexing the X_i 's, we may assume that $\Theta_1, \dots, \Theta_{m-1}$ correspond to linear edges E^1, \dots, E^{m-1} as in Fact 3.13. We may also assume that the twist path for $[a]_u$ represents $[a]$. Superscript indices, primarily i, j, k or l , take values in the set $\{1, \dots, m-1\}$. Edges indexed by subscripts like E_p or E_q can be any edge at all, perhaps even an element of $\{E^1, \dots, E^{m-1}\}$.

The twist path for E^1, \dots, E^{m-1} is denoted by w . We define rays by iterating w and w^{-1} in the positive and negative directions as follows:

$$R_+(w) = (www\dots) \quad R_+(\bar{w}) = (\bar{w}\bar{w}\bar{w}\dots)$$

$$R_-(w) = (\dots www) \quad R_-(\bar{w}) = (\dots \bar{w}\bar{w}\bar{w})$$

Given a line α , we say that α *ends with w^∞* if there is a concatenation expression $\alpha = \beta R_+(w)$. Similarly, α *begins with w^∞* if $\alpha = R_-(w)\beta$; α *ends with \bar{w}^∞* if $\alpha = \beta R_+(\bar{w})$; and α *begins with \bar{w}^∞* if $\alpha = R_-(\bar{w})\beta$. Since the notation should make the context clear, we will usually abuse notation by ignoring 'R' and writing βw^∞ instead of $\beta R_+(w)$, and similarly for the other three possibilities.

Definition 4.5. Recall that all lines in this paper are oriented. Since X_i contains a^+, a^- and at least one other point, there exist lines in \tilde{G} with initial endpoint in $X_i \setminus \{a^\pm\}$ and terminal endpoint in $\{a^\pm\}$. For $0 \leq i \leq m-1$, let Σ_i be the set of such lines of minimum 'subfiltration height' $j(i) \in \{1, \dots, J\}$. To be more precise, let $1 \leq j(i) \leq J$ be the minimum value for which there exists an $(\tilde{f}_i)_{\#}$ -invariant line $\tilde{\sigma} \neq A_a$ (equivalently, a line $\tilde{\sigma} \neq A_a$ with endpoints in X_i) that terminates at either a^+ or a^- and whose projection σ is contained in $G_{r(j(i))}$. The set of all such $\tilde{\sigma}$ is denoted by $\tilde{\Sigma}_i$ and the set of projections σ is denoted by Σ_i . Note that every $\sigma \in \Sigma_i$ ends with w^∞ or $w^{-\infty}$.

Remark 4.6. If $i \neq 0$ then by Lemma 4.9(3) below, $j(i)$ is the minimum value for which $G_{r(j(i))}$ contains E^i .

Remark 4.7. If $\sigma \in \Sigma_i$ decomposes as $\sigma = \beta w^{-\infty}$ then there is a lift $\tilde{\sigma}$ with terminal endpoint a^- and initial endpoint, say $P \neq a^+$, in X_i . The line $\tilde{\sigma}'$ that has initial endpoint P and terminal endpoint a^+ projects to $\sigma' \in \Sigma_i$ that decomposes as $\sigma' = \beta' w^\infty$.

Lemma 4.8. *If $\pi_a(k) = i$ then:*

- (1) $j(k) = j(i)$.
- (2) $\psi_{\#}(\Sigma_k) = \Sigma_i$.
- (3) $\mathcal{F}_{\text{supp}}(\Sigma_k) = \psi(\mathcal{F}_{\text{supp}}(\Sigma_k)) = \mathcal{F}_{\text{supp}}(\Sigma_i)$.

Proof. Recall that $h : G \rightarrow G$ represents ψ and that each $G_{r(j)}$ is h -invariant. Let $\tilde{h} : \tilde{G} \rightarrow \tilde{G}$ be the lift corresponding to Ψ_a . In particular, $\tilde{h} = \hat{\Psi}_a$ fixes both a^- and a^+ . If $\tilde{\sigma} \in \Sigma_k$ then $\tilde{h}_{\#}(\tilde{\sigma})$ has endpoints in $\hat{\Psi}_a(X_k) = X_i$ and projects to $h_{\#}(\sigma) \subset G_{r(j(k))}$. Thus $j(i) \leq j(k)$. This can be repeated to show that $j(k) \geq j(\pi_a(k)) \geq j(\pi_a^2(k)) \geq \dots$. Since π_a has finite order it follows that $j(\pi_a^l(k))$ is independent of l and in particular that $j(k) = j(i)$. Since $\hat{\Psi}_a(X_k \cap \partial G_{r(j(k))}) = X_i \cap \partial G_{r(j(i))}$, it follows that $h_{\#}(\Sigma_k) = \Sigma_i$ and hence that $\psi(\mathcal{F}_{\text{supp}}(\Sigma_k)) = \mathcal{F}_{\text{supp}}(\Sigma_i)$. Iterating this argument, shows that $\psi^l(\mathcal{F}_{\text{supp}}(\Sigma_k)) = \mathcal{F}_{\text{supp}}(\Sigma_{\pi_a^l(k)})$ for all $l \geq 1$. Since π_a has finite order, $\mathcal{F}_{\text{supp}}(\Sigma_k)$ is preserved by an iterate of ψ and so [HM17c, Lemma 4.2] is also preserved by ψ . \square

Our strategy for proving that π_a is the identity is to assume that this is not the case, and to produce a closed path $\delta \subset G$ representing a homology class in $H_1(G; \mathbb{Z}/3)$ that is not fixed by ϕ . Typically δ will occur as a subpath of some nonperiodic line τ in some Σ_i having the form $\tau = w^\infty \delta w^\infty$. Homology information of such paths δ will be extracted from the algebraic crossing number of each line in each Σ_i with certain edges E_p . The information about lines in the sets Σ_i that we need for these purposes is contained in the following lemma.

Lemma 4.9. *For $0 \leq i \leq m-1$ and $\tilde{\sigma} \in \tilde{\Sigma}_i$, let $\tilde{\sigma} = \dots \cdot \tilde{\sigma}_{-1} \cdot \tilde{\sigma}_0 \cdot \tilde{\sigma}_1 \cdot \dots$ be the decomposition where the endpoints of the $\tilde{\sigma}_m$'s are exactly the vertices of $\text{Fix}(\tilde{f}_i) \cap \tilde{\sigma}$.*

- (1) *The decomposition is a non-trivial splitting whose finite terms are fixed edges and indivisible Nielsen paths for \tilde{f}_i .*
- (2) *If a finite $\tilde{\sigma}_s$ crosses E_p or \bar{E}_p for some non-fixed edge $E_p \subset S_{j(i)}$ then E_p is linear and $\sigma_s = E_p w_p^* \bar{E}_p$ where w_p is the twist path for E_p . Moreover, $w_p^{\pm\infty} \bar{E}_p \cdot \sigma_{s+1} \sigma_{s+2} \dots$ is an element of Σ_i .*
- (3) *If $i \neq 0$ then the splitting has a last term and it projects to $E^i w^{\pm\infty}$.*
- (4) *If $i = 0$ then the splitting has no last term.*

(5) If the splitting has a first term $\tilde{\sigma}_0$ and if σ_0 crosses E^t or \bar{E}^t where $E^t \subset S_{j(i)}$ and $1 \leq t \leq m-1$ then σ_0 has the form $w^{\pm\infty}\bar{E}^t$.

Proof. Non-triviality of the decomposition follows from Lemma 3.23. It is a splitting because each of its terms is fixed by $\tilde{f}_{i\#}$. For the same reason, each finite $\tilde{\sigma}_l$ is a Nielsen path. If $\tilde{\sigma}_l$ is neither a single edge nor an indivisible Nielsen path then it would contain a fixed point in its interior; the (Vertices) property of a CT would then imply that $\tilde{\sigma}_l$ contains a fixed vertex in its interior which is not the case. This proves (1).

If $\tilde{\sigma}_s$ is finite but not a fixed edge then $\sigma_s = E_q w_q^* \bar{E}_q$ for some linear edge $E_q \subset G_{r(j(i))}$ with twist path w_q by (1) and the (NEG Nielsen path) property of a CT. Corollary 4.3 implies that $w_q \subset G_{r(j(i-1))}$ and so does not cross any edge in $S_{j(i)}$. Thus $E_q = E_p$ and the main statement of (2) is satisfied. The moreover part of (2) follows from the following observations: the turn (E_p, σ_{s+1}) is legal; $\tilde{\sigma}_{s+1} \cdot \tilde{\sigma}_{s+2} \cdots$ and $\tilde{E}_p \tilde{w}_p^{\pm\infty}$ are $\tilde{f}_{i\#}$ -invariant rays; $\tilde{\sigma}_{s+1} \cdot \tilde{\sigma}_{s+2} \cdots$ and $\tilde{E}_p \tilde{w}_p^{\pm\infty}$ project into $S_{j(i)}$.

For (3) we assume that $i \neq 0$. Lemma 3.13 implies that $\text{Fix}(\tilde{f}_i) \cap A_a = \emptyset$ so there is a last fixed point in $\tilde{\sigma}$ and a last term, say $\tilde{\sigma}_b$. Since $\tilde{\sigma}_b$ is an $\tilde{f}_{i\#}$ -invariant ray, Fact 3.22 implies that each term in the highest edge splitting of $\tilde{\sigma}_b$ is $\tilde{f}_{i\#}$ -invariant. If this splitting is non-trivial, the terminal endpoint of its first term would be fixed in contradiction to the fact that the interior of $\tilde{\sigma}_b$ is disjoint from $\text{Fix}(\tilde{f}_i)$. The highest edge splitting of $\tilde{\sigma}_b$ is therefore trivial which means that $\tilde{\sigma}_b = \tilde{E}_q \tilde{\rho}$ where \tilde{E}_q is a non-fixed edge with fixed initial direction and $\tilde{\rho}$ has height strictly less than that of E_q . Let R_{E_q} be the ray determined by E_q (Notation 3.15). The lift \tilde{R}_{E_q} that begins with \tilde{E}_q terminates at some $Q \in \partial \text{Fix}(\Theta_i) \cup \text{Fix}_+(\hat{\Theta}_i) \subset X_i$ (see Fact 3.5) and intersects $\text{Fix}(\tilde{f}_i)$ only in its initial endpoint by (NEG Nielsen Paths). If Q is not equal to the terminal endpoint P of $\tilde{\sigma}$ then the line \tilde{L} connecting Q to P is disjoint from $\text{Fix}(\tilde{f}_i)$ and so equals A_a by Lemma 3.23. In this case, Q is either a^+ or a^- . The same is true if $Q = P$ because P is either a^+ or a^- . If E_q is non-linear then $Q \in \text{Fix}_+(\tilde{f}_i)$ by Facts 3.16 and 3.5 therefore imply that E_q is linear. It follows that $\tilde{R}_{E_q} \setminus \tilde{E}_q$ is contained in the axis of a covering translation that shares a terminal ray with A_a and so equals A_a . Since $\tilde{R}_{E_q} \setminus \tilde{E}_q$ projects to $w_q^{\pm\infty}$, $[w_q]_u = [w]_u$ and so $w_q = w$ and the terminal endpoint of \tilde{E}_q is in A_a . Combining this with Fact 3.13 and the fact that the initial endpoint of \tilde{E}_q is fixed, we see that that $E_q = E^i$. This completes the proof of (3).

If $i = 0$ then $A_a \cap \text{Fix}(\tilde{f}_i)$ is non-empty by Lemma 3.13 and is invariant under the covering translation T_a associated to a because \tilde{f}_i fixes a^+ and a^- . It follows that a^+ and a^- are in the closure of $A_a \cap \text{Fix}(\tilde{f}_i)$. This proves (4).

The proof of (5) is similar to that of (3). Assuming that the splitting has a first term $\tilde{\sigma}_0$ crossing E^t or $\bar{E}^t \subset S_{j(i)}$, let $\tilde{R} = \tilde{\sigma}_0^{-1}$ and let P be its terminal endpoint. As in the proof of (3), the highest edge splitting of \tilde{R} must be trivial so $\tilde{R} = \tilde{E}_q \tilde{\rho}$ where \tilde{E}_q is a non-fixed edge with fixed initial direction and $\tilde{\rho}$ has height strictly less than that of E_q . Let R_{E_q} be the ray determined by E_q , let \tilde{R}_{E_q} be the lift that begins with

\tilde{E}_q and let $Q \in X_i$ be the terminal endpoint of \tilde{R}_{E_q} . If $P = Q$ then $\tilde{R} = \widetilde{R}_{E_q}$ and $\widetilde{R}_{E_q} \setminus \tilde{E}_q \subset G_{r(j(i)-1)}$ by Corollary 4.3. In this case, $E_q = E^t$ so $\sigma_0^{-1} = R_{E^t}$ and (5) is satisfied. Suppose then that $P \neq Q$. The line \tilde{L} connecting Q to P is disjoint from $\text{Fix}(\tilde{f}_i)$. Its highest edge splitting must be bi-infinite for otherwise the endpoints of its first or last term would be fixed. It follows that the height of $\tilde{\rho}$ equals the height of $\widetilde{R}_{E_q} \setminus \tilde{E}_q$ and so is at most $r(j(i)-1)$ by Corollary 4.3. It then follows that $E_q = E^t$ and $R_{E^t} = E^t w^{\pm\infty}$. There is a conjugate a' of a such that $A_{a'}$ shares an endpoint with \widetilde{R}_{E_q} . Lemma 3.23 implies that $\tilde{L} = A_{a'}$. We conclude that $\tilde{\sigma}_0^{-1}$ begins with $\tilde{E}_q = \tilde{E}^t$ and is otherwise contained in $A_{a'}$. This completes the proof of (5). \square

Definition 4.10. If a path τ crosses E^i and \bar{E}^i a finite number of times then we define the *algebraic crossing number* $c_i(\tau)$ of τ with E^i to be the number of times that τ crosses E^i minus the number of times that τ crosses \bar{E}^i . For closed paths $\delta \subset G$, the formula

$$\delta \mapsto c_i(\delta) \pmod{3}$$

defines a homomorphism

$$H_1(G; \mathbb{Z}/3) \mapsto \mathbb{Z}/3$$

and this homomorphism is nontrivial if and only if E^i is nonseparating. For infinite paths of the form $\tau = w^\infty \delta w^\infty$, since w does not cross E^i we have $c_i(\tau) = c_i(\delta)$.

We sometimes use edge path notation in describing lines and rays. In particular, a line or ray σ *ends with* w^∞ if the ray $R(w) := w^\infty$ is a terminal subray of σ and σ *ends with* $w^{-\infty}$ if the ray $R(\bar{w}) := \bar{w}^\infty$ is a terminal subray of σ . Analogously σ *begins with* w^∞ [resp. $w^{-\infty}$] if σ^{-1} ends with $R(\bar{w})$ [resp. $R(w)$].

Corollary 4.11. For all $1 \leq i \leq m-1$ and all $\sigma \in \Sigma_i$, $c_i(\sigma) = 0$ or 1.

Proof. Let $\sigma = \dots \cdot \sigma_{-1} \cdot \sigma_0 \cdot \sigma_1 \cdot \dots$ be the splitting given by Lemma 4.9. The corollary follows from

- $c_i(\sigma_s) = 0$ for each finite σ_s .
- There is a last terminal term and its contribution to $c_i(\sigma)$ is 1.
- If there is a first term then its contribution to $c_i(\sigma)$ is 0 or -1 .

each of which is an immediate consequence of Lemma 4.9. \square

Verification that π_a (as defined in $(*)$) is the identity: To prove that π_a is the identity, it suffices to show that $\pi_a(l) = l$ for all $l > 0$; the $l = 0$ case then follows from the fact that π_a is a permutation. Suppose to the contrary that $l > 0$ and that $\pi_a(k) = l$ for some $k \neq l$.

We claim that Σ_k contains a line τ that begins with w^∞ , ends with w^∞ and satisfies $c_l(\tau) = -1$. If each line in Σ_k is contained in $G \setminus E^l$ then the same is true for the realization in G of each line in $\mathcal{F}(\Sigma_k)$ in contradiction to the fact (Lemma 4.9-(3))

that every line in Σ_l crosses E^l and the fact (Lemma 4.8-(3)) that $\mathcal{F}(\Sigma_l) = \mathcal{F}(\Sigma_k)$. We may therefore choose a line $\mu \in \Sigma_k$ that crosses E^l or \bar{E}^l . By Remark 4.7, we may assume that μ ends with w^∞ . Let $\mu = \dots \cdot \mu_{-1} \cdot \mu_0 \cdot \mu_1 \cdot \dots$ be the splitting given by Lemma 4.9. The last term in the splitting, if it exists, does not cross E_l or \bar{E}_l by items (3) and (4) of Lemma 4.9. If some finite term μ_s crosses E^l or \bar{E}^l then $\mu_s = E^l w^* \bar{E}^l$ by Lemma 4.9-(2). In this case, $\tau = w^\infty \bar{E}^l \mu_{s+1} \mu_{s+2} \dots$ is contained in Σ_k by Lemma 4.9-(2) and $c_l(\tau) = -1$. If no finite term of the splitting crosses E^l or \bar{E}^l then there must be a first term μ_b and it must cross E^l or \bar{E}^l so $\mu_b = w^{\pm\infty} \bar{E}^l$ by Lemma 4.9(5). In this case, $\tau = w^\infty \bar{E}^l \mu_{b+1} \mu_{b+2} \dots$ satisfies $c_l(\tau) = -1$. If $\tau \neq \mu$ then τ and μ have lifts $\tilde{\tau}$ and $\tilde{\mu}$ with terminal endpoint a^+ and with initial endpoints bounding A_b some $b \in F_n$ satisfying $[b] = [w]$. Since one of the endpoints of A_b is contained in X_k the other is also so $\tau \in \Sigma_k$ as desired. This completes the proof of the claim.

There is a closed path δ such that $\tau = w^\infty \delta w^\infty$. The line $\tau' = h_\#(\tau) \in \Sigma_l$ is obtained from $w^\infty h_\#(\delta) w^\infty$ by tightening. No copies of E_l or \bar{E}_l are cancelled during the tightening process, so $\tau' = w^\infty \delta' w^\infty$ where δ' is a closed path satisfying $c_l(\delta') = c_l(h_\#(\delta))$. Applying Definition 4.10, we have

$$c_l(\delta) = c_l(\tau) = -1 \quad \text{and} \quad c_l(h_\#(\delta)) = c_l(\delta') = c_l(\tau')$$

Since $\psi \in \text{IA}_n(\mathbb{Z}/3)$ and h is a topological representative of ψ ,

$$c_l(\tau') \bmod 3 = c_l(h_\#(\delta)) \bmod 3 = c_l(\delta) \bmod 3 = -1$$

This contradiction to Corollary 4.11 completes the verification of $(*)$ and hence the proof of Proposition 1.3. \square

References

- [BFH00] M. Bestvina, M. Feighn, and M. Handel, *The Tits alternative for $\text{Out}(F_n)$. I. Dynamics of exponentially-growing automorphisms.*, Ann. of Math. **151** (2000), no. 2, 517–623.
- [BFH04] ———, *Solvable subgroups of $\text{Out}(F_n)$ are virtually Abelian*, Geometriae Dedicata **104** (2004), 71–96.
- [BFH05] ———, *The Tits alternative for $\text{Out}(F_n)$. II. A Kolchin type theorem*, Ann. of Math. **161** (2005), no. 1, 1–59.
- [BH92] M. Bestvina and M. Handel, *Train tracks and automorphisms of free groups*, Ann. of Math. **135** (1992), 1–51.
- [BL94] H. Bass and A. Lubotzky, *Linear-central filtrations on groups*, The mathematical legacy of Wilhelm Magnus: groups, geometry and special functions (Brooklyn, NY, 1992), Cont. Math., vol. 169, AMS, 1994.

- [FH] M. Feighn and M. Handel, *Algorithmic constructions of relative train track maps and cts*, to appear.
- [FH09] M. Feighn and H. Handel, *Abelian subgroups of $\text{Out}(F_n)$* , Geometry and Topology **13** (2009), 1657–1727.
- [FH11] _____, *The recognition theorem for $\text{Out}(F_n)$* , Groups Geom. Dyn. **5** (2011), 39–106.
- [GJLL98] D. Gaboriau, A. Jaeger, G. Levitt, and M. Lustig, *An index for counting fixed points of automorphisms of free groups*, Duke Math. J. **93** (1998), no. 3, 425–452.
- [HM] M. Handel and L. Mosher, *Subgroup decomposition in $\text{Out}(F_n)$* , to appear in Memoirs of Amer. Math. Soc.
- [HM15] _____, *Hyperbolic actions and 2nd bounded cohomology of subgroups of $\text{Out}(F_n)$ Part I: Infinite lamination subgroups*, arXiv:1511.06913, Nov. 2015.
- [HM17a] _____, *Hyperbolic actions and 2nd bounded cohomology of subgroups of $\text{Out}(F_n)$ Part II: Finite lamination subgroups*, arXiv:1702.08050, Jun 2017.
- [HM17b] _____, *Subgroup decomposition in $\text{Out}(F_n)$, Part I: Geometric models*, Memoirs AMS (2017), to appear. arXiv:1302.2378.
- [HM17c] _____, *Subgroup decomposition in $\text{Out}(F_n)$, Part II: A relative Kolchin theorem*, Memoirs AMS (2017), to appear. arXiv:1302.2379.