

Abelian subgroups of $\text{Out}(F_n)$

MARK FEIGHN
MICHAEL HANDEL

We classify abelian subgroups of $\text{Out}(F_n)$ up to finite index in an algorithmic and computationally friendly way. A process called disintegration is used to canonically decompose a single rotationless element ϕ into a composition of finitely many elements and then these elements are used to generate an abelian subgroup $\mathcal{A}(\phi)$ that contains ϕ . The main theorem is that up to finite index every abelian subgroup is realized by this construction. As an application we give an explicit description, in terms of relative train track maps and up to finite index, of all maximal rank abelian subgroups of $\text{Out}(F_n)$ and of IA_n .

20F65; 20F28

1 Introduction

In this paper we classify abelian subgroups of $\text{Out}(F_n)$ up to finite index in an algorithmic and computationally friendly way. There are two steps. The first is to construct an abelian subgroup $\mathcal{D}(\phi)$ from a given $\phi \in \text{Out}(F_n)$ by a process that we call *disintegration*. The subgroup $\mathcal{D}(\phi)$ is very well understood in terms of relative train track maps and has natural coordinates that embed it into some \mathbb{Z}^M . The second step is to prove the following theorem.

Theorem 7.2 *For every abelian subgroup A of $\text{Out}(F_n)$ there exists $\phi \in A$ such that $A \cap \mathcal{D}(\phi)$ has finite index in A .*

To motivate the disintegration process, consider an element μ of the mapping class group $\text{MCG}(S)$ of a compact oriented surface S . After possibly replacing μ by an iterate, there is, by the Thurston classification theorem [17; 9], a decomposition of S into subsurfaces S_I , some of which are annuli and the rest of which have negative Euler characteristic, and there is a homeomorphism $h: S \rightarrow S$ representing μ , called a *normal form* for μ , that preserves each S_I . If S_I is an annulus then $h|_{S_I}$ is a nontrivial Dehn twist. If S_I has negative Euler characteristic then $h|_{S_I}$ is either the identity or pseudo-Anosov. In all cases, $h|_{\partial S_I}$ is the identity.

We may assume that the S_l 's are numbered so that $h|_{S_l}$ is the identity if and only if $l > M$ for some M . For each M -tuple of integers $\mathbf{a} = (a_1, \dots, a_M)$ let $h_{\mathbf{a}}: S \rightarrow S$ be the homeomorphism that agrees with h^{a_l} on S_l for $1 \leq l \leq M$ and is the identity on the remaining S_l 's. Then $h_{\mathbf{a}}$ is a normal form for an element $\mu_{\mathbf{a}} \in \text{MCG}(S)$ and we define $\mathcal{D}(\mu)$ to be the subgroup consisting of all such $\mu_{\mathbf{a}}$. It is easy to check that $\mu_{\mathbf{a}} \mapsto \mathbf{a}$ defines an isomorphism between $\mathcal{D}(\mu)$ and \mathbb{Z}^M .

An element ϕ of $\text{Out}(F_n)$ has finite sets of natural invariants on which it acts by permutation. If these actions are trivial then we say that ϕ is *rotationless*; complete details can be found in Section 3. Suppose that ϕ is a rotationless element of $\text{Out}(F_n)$. The analog of a normal form $h: S \rightarrow S$ is a relative train track map $f: G \rightarrow G$ which is a particularly nice homotopy equivalence of a marked graph that represents ϕ in the sense that the outer automorphism of $\pi_1(G)$ that it induces is identified with ϕ by the marking. There is an associated maximal filtration $\emptyset = G_0 \subset G_1 \subset \dots \subset G_N = G$ by f -invariant subgraphs. The i -th stratum H_i is the closure of $G_i \setminus G_{i-1}$. The exact properties satisfied by $f: G \rightarrow G$ and $\emptyset = G_0 \subset G_1 \subset \dots \subset G_N = G$ are detailed in Section 2.

As a first attempt to mimic the construction of $\mathcal{D}(\mu)$, let X_1, \dots, X_M be the strata that are not pointwise fixed by f , let $\mathbf{a} = (a_1, \dots, a_M)$ be an M -tuple of nonnegative integers and define $f_{\mathbf{a}}$ to agree with f^{a_l} on X_l and to be the identity on the subgraph of edges fixed by f . Although it is not obvious, $f_{\mathbf{a}}: G \rightarrow G$ is a homotopy equivalence (see Lemma 6.7) and so defines an element $\phi_{\mathbf{a}} \in \text{Out}(F_n)$.

Without some restrictions on \mathbf{a} however, the subgroup generated by the $\phi_{\mathbf{a}}$'s need not be abelian. In the following examples, we do not distinguish between a homotopy equivalence of the rose and the outer automorphism that it represents.

Example 1.1 Let G be the graph with one vertex and with edges labelled A , B and C . Define $f: G \rightarrow G$ by

$$A \mapsto A \quad B \mapsto BA \quad C \mapsto CB.$$

Let $X_1 = \{B\}$ and $X_2 = \{C\}$ and $\mathbf{a} = (m, n)$. Then

$$f_{(m,n)} \circ f(C) = f_{(m,n)}(CB) = f^n(C)f^m(B)$$

$$\text{and } f \circ f_{(m,n)}(C) = f(f^n(C)) = f^n(f(C)) = f^n(CB) = f^n(C)f^n(B).$$

This shows that $f_{(m,n)}$ commutes with $f = f_{(1,1)}$ if and only if $m = n$.

The underlying problem is that strata are not invariant. It does not matter that the path $f(B)$ crosses A since A is fixed by f . The lack of commutativity stems from the fact that $f(C)$ crosses B .

To address this problem we enlarge the X_i 's to be unions of strata. It is not necessary to choose the X_i 's to be fully invariant (ie to satisfy $f(X_i) \subset X_i$) but they must be *almost invariant* as made precise in Definition 6.3.

The next example illustrates a more subtle relation on the coordinates of \mathbf{a} that is needed to insure that the $f_{\mathbf{a}}$'s commute.

Example 1.2 Let G be the graph with one vertex and with edges labelled A , B , C and D . Define $f: G \rightarrow G$ by

$$A \mapsto A \quad B \mapsto BA^2 \quad C \mapsto CA^5 \quad D \mapsto DC\bar{B}$$

where \bar{B} is B with its orientation reversed. Let $X_1 = \{B\}$, $X_2 = \{C\}$ and $X_3 = \{D\}$ and let $\mathbf{a} = (m, n, p)$. Then

$$f \circ f_{\mathbf{a}}(D) = f(f^p(D)) = f^p(f(D)) = f^p(DC\bar{B}) = f^p(D)f^p(C\bar{B})$$

$$\text{and } f_{\mathbf{a}} \circ f(D) = f_{\mathbf{a}}(DC\bar{B}) = f^p(D)f_{\mathbf{a}}(C\bar{B}).$$

If f commutes with $f_{\mathbf{a}}$ then

$$f^p(C\bar{B}) = f_{\mathbf{a}}(C\bar{B}).$$

Thus $CA^{3p}\bar{B} = CA^{5n-2m}\bar{B}$ and $3p = 5n - 2m$. One can check that the converse holds as well. Namely if we require that \mathbf{a} be an element of the linear subspace of $\mathbb{Z}^3 = \{(m, n, p)\}$ defined by $3p = 5n - 2m$ then the $\phi_{\mathbf{a}}$'s commute.

The path $C\bar{B}$ of Example 1.2 is *quasi-exceptional* as defined in Section 6. When the image of an edge in X_k contains a quasi-exceptional path with initial edge in X_i and terminal edge in X_j then there is an induced relation between the i -th, j -th and k -th coefficients of \mathbf{a} . These define a subgroup of \mathbb{Z}^M . The nonnegative M -tuples that lie in this subspace are said to be *admissible*. The map $\mathbf{a} \rightarrow \phi_{\mathbf{a}}$ on admissible M -tuples extends to an injective homomorphism of this subgroup of \mathbb{Z}^M and we define the image of this subspace to be $\mathcal{D}(\phi)$.

The mapping class group version of Theorem 7.2 is a straightforward consequence of two easily proved, well known facts. The first (see for example Corollary 5.2 of Franks, Handel and Parwani [11]) is that the subsurfaces S_l can be chosen independently of $\mu \in A$. The second (see for example Lemma 2.10 of Franks, Handel and Parwani [12]) is that an abelian subgroup containing a pseudo-Anosov element is virtually cyclic.

The proof for $\text{Out}(F_n)$ is considerably harder. This is due, in part, to the fact that disintegration in $\text{Out}(F_n)$ is a more complicated operation, as illustrated by the examples, than it is $\text{MCG}(S)$. Another factor is that, unlike normal forms in the mapping class

group, relative train track maps representing an element $\phi \in \text{Out}(F_n)$ are not unique. No matter how canonical a construction is with respect to a particular $f: G \rightarrow G$, one must still check the extent to which it is independent of the choice of $f: G \rightarrow G$. The most technically difficult argument in this paper (Section 7) is a proof that the rank of the admissible linear subspace of \mathbb{Z}^M described above depends only on ϕ and not the choice of $f: G \rightarrow G$.

Recall that IA_n is the subgroup of $\text{Out}(F_n)$ consisting of elements that act as the identity on $H_1(F_n)$. As an application of Theorem 7.2 we classify, up to finite index, maximal rank abelian subgroups of $\text{Out}(F_n)$ and of IA_n . The exact statements appear as Proposition 8.9 and Proposition 8.10. Roughly speaking, we prove that if $\mathcal{D}(\phi)$ has maximal rank in $\text{Out}(F_n)$ then $f: G \rightarrow G$ has $2n - 3$ strata, each of which is either a single linear edge or is exponentially growing and is closely related to a pseudo-Anosov homeomorphism of a four times punctured sphere. If $\mathcal{D}(\phi)$ has maximal rank in IA_n then $f: G \rightarrow G$ has $2n - 4$ such strata and pointwise fixes a rank two subgraph.

From an algebraic point of view, the natural abelian subgroup associated to an element $\phi \in \text{Out}(F_n)$ is the center $Z(C(\phi))$ of the centralizer $C(\phi)$ of ϕ which can also be described as the intersection of all maximal (with respect to inclusion) abelian subgroups that contain ϕ . In our context it is natural to look at the weak center $WZ(C(\phi))$ of $C(\phi)$ defined as the subgroup of elements that commute with an iterate of each element of $C(\phi)$. The following result is a step toward an algorithmic construction of $Z(C(\phi))$.

Theorem 6.21 $\mathcal{D}(\phi) \subset WZ(C(\phi))$ for all rotationless ϕ .

In Section 9 we apply this theorem to give algebraic characterizations of certain maximal rank abelian subgroups of $\text{Out}(F_n)$ and IA_n . This characterization is needed in the calculation of the commensurator group of $\text{Out}(F_n)$ by the authors [8].

In Section 3 we define what it means for $\phi \in \text{Out}(F_n)$ to be *rotationless*, prove that the rotationless elements of any abelian subgroup A form a finite index subgroup A_R and consider lifts of A_R from $\text{Out}(F_n)$ to $\text{Aut}(F_n)$. These lifts are essential to our approach and are similar to ones used in Bestvina, Feighn and Handel [3].

In Section 4 we define a natural embedding of A_R into a lattice in Euclidean space and say what it means for an element of A_R to be *generic* with respect to this embedding.

In Section 5 we associate an abelian subgroup $\mathcal{A}(\phi)$ to each rotationless ϕ and prove that if ϕ is generic in A_R then $A_R \subset \mathcal{A}(\phi)$. We also prove (Corollary 5.6) that $\mathcal{A}(\phi) \subset WZ(C(\phi))$.

In Section 6 we define $\mathcal{D}(\phi)$ and prove (Corollary 6.20) that $\mathcal{D}(\phi) \subset \mathcal{A}(\phi)$, thereby completing the proof of Theorem 6.21.

In Section 7 we prove (Theorem 7.1) that $\mathcal{D}(\phi)$ has finite index in $\mathcal{A}(\phi)$ by reconciling the normal forms point of view used to define $\mathcal{D}(\phi)$ with the “action on ∂F_n ” point of view used to define $\mathcal{A}(\phi)$. Theorem 7.2 is an immediate consequence of this result and the fact, mentioned above, that $A_R \subset \mathcal{A}(\phi)$ for generic $\phi \in A$.

We make use of several important results from our paper [10], including the Recognition Theorem and the existence of relative train track maps that are especially well suited to disintegrating an element and forming $\mathcal{D}(\phi)$. Section 2 reviews this and other relevant material and sets notation for the paper.

Acknowledgements We thank Gilbert Levitt and the referee for many helpful comments. This material is based upon work of the first author supported by the National Science Foundation under Grant No. DMS0805440 and work of the second author supported by the National Science Foundation under Grant No. DMS0405814.

2 Background

Fix $n \geq 2$ and let F_n be the free group of rank n . Denote the automorphism group of F_n by $\text{Aut}(F_n)$, the group of inner automorphisms of F_n by $\text{Inn}(F_n)$ and the group of outer automorphisms of F_n by $\text{Out}(F_n) = \text{Aut}(F_n)/\text{Inn}(F_n)$. We follow the convention that elements of $\text{Aut}(F_n)$ are denoted by upper case Greek letters and that the same Greek letter in lower case denotes the corresponding element of $\text{Out}(F_n)$. Thus $\Phi \in \text{Aut}(F_n)$ represents $\phi \in \text{Out}(F_n)$.

Marked graphs and outer automorphisms Identify F_n with $\pi_1(R_n, *)$ where R_n is the rose with one vertex $*$ and n edges. A *marked graph* G is a graph of rank n without valence one vertices, equipped with a homotopy equivalence $m: R_n \rightarrow G$ called a *marking*. Letting $b = m(*) \in G$, the marking determines an identification of F_n with $\pi_1(G, b)$.

A homotopy equivalence $f: G \rightarrow G$ and a path σ from b to $f(b)$ determines an automorphism of $\pi_1(G, b)$ and hence an element of $\text{Aut}(F_n)$. As the homotopy class of σ varies, the automorphism ranges over all representatives of the associated outer automorphism ϕ . We say that $f: G \rightarrow G$ *represents* ϕ . We always assume that f maps vertices to vertices and that the restriction of f to any edge is an immersion.

Paths, circuits and edge paths Let Γ be the universal cover of a marked graph G and let $\text{pr}: \Gamma \rightarrow G$ be the covering projection. A proper map $\tilde{\sigma}: J \rightarrow \Gamma$ with domain a (possibly infinite) closed interval J will be called a *path in* Γ if it is an embedding or if J is finite and the image is a single point; in the latter case we say that $\tilde{\sigma}$ is a

trivial path. If J is finite, then any map $\tilde{\sigma}: J \rightarrow \Gamma$ is homotopic rel endpoints to a unique (possibly trivial) path $[\tilde{\sigma}]$; we say that $[\tilde{\sigma}]$ is obtained from $\tilde{\sigma}$ by *tightening*. If $f: \Gamma \rightarrow \Gamma$ is a lift of a homotopy equivalence $f: G \rightarrow G$, we denote $[f(\tilde{\sigma})]$ by $f_{\#}(\tilde{\sigma})$.

We will not distinguish between paths in Γ that differ only by an orientation preserving change of parametrization. Thus we are interested in the oriented image of $\tilde{\sigma}$ and not $\tilde{\sigma}$ itself. If the domain of $\tilde{\sigma}$ is finite, then the image of $\tilde{\sigma}$ has a natural decomposition as a concatenation $\tilde{E}_1 \tilde{E}_2 \cdots \tilde{E}_{k-1} \tilde{E}_k$ where \tilde{E}_i , $1 < i < k$, is an edge of Γ , \tilde{E}_1 is the terminal segment of an edge and \tilde{E}_k is the initial segment of an edge. If the endpoints of the image of $\tilde{\sigma}$ are vertices, then \tilde{E}_1 and \tilde{E}_k are full edges. The sequence $\tilde{E}_1 \tilde{E}_2 \cdots \tilde{E}_k$ is called *the edge path associated to $\tilde{\sigma}$* . This notation extends naturally to the case that the interval of domain is half-infinite or bi-infinite. In the former case, an edge path has the form $\tilde{E}_1 \tilde{E}_2 \cdots$ or $\cdots \tilde{E}_{-2} \tilde{E}_{-1}$ and in the latter case has the form $\cdots \tilde{E}_{-1} \tilde{E}_0 \tilde{E}_1 \tilde{E}_2 \cdots$.

A *path in G* is the composition of the projection map pr with a path in Γ . Thus a map $\sigma: J \rightarrow G$ with domain a (possibly infinite) closed interval will be called a path if it is an immersion or if J is finite and the image is a single point; paths of the latter type are said to be *trivial*. If J is finite, then any map $\sigma: J \rightarrow G$ is homotopic rel endpoints to a unique (possibly trivial) path $[\sigma]$; we say that $[\sigma]$ is obtained from σ by *tightening*. For any lift $\tilde{\sigma}: J \rightarrow \Gamma$ of σ , $[\sigma] = \text{pr}[\tilde{\sigma}]$. We denote $[f(\sigma)]$ by $f_{\#}(\sigma)$. We do not distinguish between paths in G that differ by an orientation preserving change of parametrization. The *edge path associated to σ* is the projected image of the edge path associated to a lift $\tilde{\sigma}$. Thus the edge path associated to a path with finite domain has the form $E_1 E_2 \cdots E_{k-1} E_k$ where E_i , $1 < i < k$, is an edge of G , E_1 is the terminal segment of an edge and E_k is the initial segment of an edge. We will identify paths with their associated edge paths whenever it is convenient.

We reserve the word *circuit* for an immersion $\sigma: S^1 \rightarrow G$. Any homotopically nontrivial map $\sigma: S^1 \rightarrow G$ is homotopic to a unique circuit $[\sigma]$. As was the case with paths, we do not distinguish between circuits that differ only by an orientation preserving change in parametrization and we identify a circuit σ with a *cyclically ordered edge path* $E_1 E_2 \cdots E_k$.

A path or circuit *crosses* or *contains* an edge if that edge occurs in the associated edge path. For any path σ in G define $\bar{\sigma}$ to be “ σ with its orientation reversed”. For notational simplicity, we sometimes refer to the inverse of $\tilde{\sigma}$ by $\tilde{\sigma}^{-1}$.

A decomposition of a path or circuit into subpaths is a *splitting* for $f: G \rightarrow G$ and is denoted $\sigma = \cdots \sigma_1 \cdot \sigma_2 \cdots$ if $f_{\#}^k(\sigma) = \cdots f_{\#}^k(\sigma_1) f_{\#}^k(\sigma_2) \cdots$ for all $k \geq 0$. In other

words, a decomposition of σ into subpaths σ_i is a splitting if one can tighten the image of σ under any iterate of $f_{\#}$ by tightening the images of the σ_i 's.

A path σ is a *periodic Nielsen path* if $f_{\#}^k(\sigma) = \sigma$ for some $k \geq 1$. The minimal such k is the *period* of σ and if $k = 1$ then σ is a *Nielsen path*. Two elements of $\text{Fix}(f)$ are in the same *Nielsen class* if they are the endpoints of a Nielsen path. A (periodic) Nielsen path is *indivisible* if it does not decompose as a concatenation of nontrivial (periodic) Nielsen subpaths. A path or circuit is *root-free* if it is not a multiple of a simpler path or circuit.

Automorphisms and lifts Section 1 of Gaboriau et al [13] and Section 2.1 of Bestvina, Feighn and Handel [3] are good sources for facts that we record below without specific references. The universal cover Γ of a marked graph G with marking $m: R_n \rightarrow G$ is a simplicial tree. We always assume that a base point $\tilde{b} \in \Gamma$ projecting to $b = m(*) \in G$ has been chosen, thereby defining an action of F_n on Γ . The set of ends $\mathcal{E}(\Gamma)$ of Γ is naturally identified with the boundary ∂F_n of F_n and we make implicit use of this identification throughout the paper.

Each nontrivial $c \in F_n$ acts by a nontrivial *covering translation* $T_c: \Gamma \rightarrow \Gamma$ and each T_c induces a homeomorphism $\hat{T}_c: \partial F_n \rightarrow \partial F_n$ that fixes two points, a sink T_c^+ and a source T_c^- . The line in Γ whose ends converge to T_c^- and T_c^+ is called the *axis* of T_c and is denoted A_c . The image of A_c in G is the circuit corresponding to the conjugacy class $[c]$ of c . We say that c is *root-free* if it is not a multiple of some other element of F_n . In that case T_c is not a multiple of some other covering translation and we say that T_c is *root-free*.

If $f: G \rightarrow G$ represents $\phi \in \text{Out}(F_n)$ then there is a bijection, defined by $\tilde{f}T_c = T_{\Phi(c)}\tilde{f}$ for all $c \in F_n$, between the set of lifts $\tilde{f}: \Gamma \rightarrow \Gamma$ of $f: G \rightarrow G$ and the set of automorphisms $\Phi: F_n \rightarrow F_n$ representing ϕ . We say that \tilde{f} *corresponds to* Φ or *is determined by* Φ and vice versa. Under the identification of $\mathcal{E}(\Gamma)$ with ∂F_n , a lift \tilde{f} determines a homeomorphism \hat{f} of ∂F_n . An automorphism Φ also determines a homeomorphism $\hat{\Phi}$ of ∂F_n and $\hat{f} = \hat{\Phi}$ if and only if \tilde{f} corresponds to Φ . In particular, $\hat{i}_c = \hat{T}_c$ for all $c \in F_n$ where $i_c(w) = cwc^{-1}$ is the inner automorphism of F_n determined by c . We use the notation \hat{f} and $\hat{\Phi}$ interchangeably depending on the context.

We are particularly interested in the dynamics of $\hat{f} = \hat{\Phi}$. We denote the fixed point set of $\hat{\Phi}$ by $\text{Fix}(\hat{\Phi})$ and the fixed subgroup of Φ by $\text{Fix}(\Phi)$. The following two lemmas are contained in Lemma 2.3 and Lemma 2.4 of [3] and in Proposition 1.1 of [13].

Lemma 2.1 Assume that $\tilde{f}: \Gamma \rightarrow \Gamma$ corresponds to $\Phi \in \text{Aut}(F_n)$. Then the following are equivalent:

- (i) $c \in \text{Fix}(\Phi)$.
- (ii) T_c commutes with \tilde{f} .
- (iii) \hat{T}_c commutes with \hat{f} .
- (iv) $\text{Fix}(\hat{T}_c) \subset \text{Fix}(\hat{f}) = \text{Fix}(\hat{\Phi})$.
- (v) $\text{Fix}(\hat{f}) = \text{Fix}(\hat{\Phi})$ is \hat{T}_c -invariant.

A point $P \in \partial F_n$ is an *attractor* for $\hat{\Phi}$ if it has a neighborhood $U \subset \partial F_n$ such that $\hat{\Phi}(U) \subset U$ and such that $\bigcap_{n=1}^{\infty} \hat{\Phi}^n(U) = P$. If Q is an attractor for $\hat{\Phi}^{-1}$ then we say that it is a *repeller* for $\hat{\Phi}$.

Lemma 2.2 Assume that $\tilde{f}: \Gamma \rightarrow \Gamma$ corresponds to $\Phi \in \text{Aut}(F_n)$ and that $\text{Fix}(\hat{\Phi}) \subset \partial F_n$ contains at least three points. Denote $\text{Fix}(\Phi)$ by \mathbb{F} and the corresponding subgroup of covering translations of Γ by $\mathbb{T}(\Phi)$. Then:

- (i) $\partial \mathbb{F}$ is naturally identified with the closure of $\{T_c^{\pm} : T_c \in \mathbb{T}(\Phi)\}$ in ∂F_n . None of these points is isolated in $\text{Fix}(\hat{\Phi})$.
- (ii) Each point in $\text{Fix}(\hat{\Phi}) \setminus \partial \mathbb{F}$ is isolated and is either an attractor or a repeller for the action of $\hat{\Phi}$.
- (iii) There are only finitely many $\mathbb{T}(\Phi)$ -orbits in $\text{Fix}(\hat{\Phi}) \setminus \partial \mathbb{F}$.

Lines and laminations Suppose that Γ is the universal cover of a marked graph G . An unoriented bi-infinite path in Γ is called a *line* in Γ . The *space of lines* in Γ is denoted $\tilde{\mathcal{B}}(\Gamma)$ and is equipped with what amounts to the compact-open topology. Namely, for any finite path $\tilde{\alpha}_0 \subset \Gamma$ (with endpoints at vertices if desired), define $N(\tilde{\alpha}_0) \subset \tilde{\mathcal{B}}(\Gamma)$ to be the set of lines in Γ that contain $\tilde{\alpha}_0$ as a subpath. The sets $N(\tilde{\alpha}_0)$ define a basis for the topology on $\tilde{\mathcal{B}}(\Gamma)$.

An unoriented bi-infinite path in G is called a *line* in G . The *space of lines* in G is denoted $\mathcal{B}(G)$. There is a natural projection map from $\tilde{\mathcal{B}}(\Gamma)$ to $\mathcal{B}(G)$ and we equip $\mathcal{B}(G)$ with the quotient topology.

A line in Γ is determined by the unordered pair of its endpoints (P, Q) , so it corresponds to a point in the *space of abstract lines* defined to be $\tilde{\mathcal{B}} := ((\partial F_n \times \partial F_n) \setminus \Delta) / \mathbb{Z}_2$, where Δ is the diagonal and where \mathbb{Z}_2 acts on $\partial F_n \times \partial F_n$ by interchanging the factors. The action of F_n on ∂F_n induces an action of F_n on $\tilde{\mathcal{B}}$ whose quotient space is denoted \mathcal{B} . The “endpoint map” defines a homeomorphism between $\tilde{\mathcal{B}}$ and $\tilde{\mathcal{B}}(\Gamma)$ and

we use this implicitly to identify $\tilde{\mathcal{B}}$ with $\tilde{\mathcal{B}}(\Gamma)$ and hence $\tilde{\mathcal{B}}(\Gamma)$ with $\tilde{\mathcal{B}}(\Gamma')$ where Γ' is the universal cover of any other marked graph G' . There is a similar identification of $\mathcal{B}(G)$ with \mathcal{B} and with $\mathcal{B}(G')$. We sometimes say that the line in G or Γ corresponding to an abstract line is *the realization* of that abstract line in G or Γ .

A closed set of lines in G or a closed F_n -invariant set of lines in Γ is called a *lamination* and the lines that compose it are called *leaves*. If Λ is a lamination in G then we denote its preimage in Γ by $\tilde{\Lambda}$ and vice-versa.

Suppose that $f: G \rightarrow G$ represents ϕ and that \tilde{f} is a lift of f . If $\tilde{\gamma}$ is a line in Γ with endpoints P and Q , then there is a bounded homotopy from $\tilde{f}(\tilde{\gamma})$ to the line $\tilde{f}_\#(\tilde{\gamma})$ with endpoints $\tilde{f}(P)$ and $\tilde{f}(Q)$. This defines an action $\tilde{f}_\#$ of \tilde{f} on lines in Γ . If $\Phi \in \text{Aut}(F_n)$ corresponds to \tilde{f} then $\Phi_\# = \tilde{f}_\#$ is described on abstract lines by $(P, Q) \mapsto (\hat{\Phi}(P), \hat{\Phi}(Q))$. There is an induced action $\phi_\#$ of ϕ on lines in G and in particular on laminations in G .

To each $\phi \in \text{Out}(F_n)$ is associated a finite ϕ -invariant set of laminations $\mathcal{L}(\phi)$ called the set of *attracting laminations* for ϕ . The individual laminations need not be ϕ -invariant. By definition (see Definition 3.1.5 of [2]) $\mathcal{L}(\phi) = \mathcal{L}(\phi^k)$ for all $k \geq 1$ and each $\Lambda \in \mathcal{L}(\phi)$ contains birecurrent leaves, called *generic leaves*, whose weak closure is all of Λ . Complete details on $\mathcal{L}(\phi)$ can be found in Section 3 of [2].

A point $P \in \partial F_n$ determines a lamination $\Lambda(P)$, called *the accumulation set of P* , as follows. Let Γ be the universal cover of a marked graph G and let \tilde{R} be any ray in Γ converging to P . A line $\tilde{\sigma} \subset \Gamma$ belongs to $\Lambda(P)$ if every finite subpath of $\tilde{\sigma}$ is contained in some translate of \tilde{R} . Since any two rays converging to P have a common infinite end, this definition is independent of the choice of \tilde{R} . The bounded cancellation lemma of Cooper [5] implies (cf Lemma 3.1.4 of [2]) that this definition is independent of the choice of G and Γ and that

$$\hat{\Phi}_\#(\Lambda(P)) = \Lambda(\hat{\Phi}(P)).$$

In particular, if $P \in \text{Fix}(\hat{\Phi})$ then $\Lambda(P)$ is $\phi_\#$ -invariant.

Free factor systems The conjugacy class of a free factor F^i of F_n is denoted $[[F^i]]$. If F^1, \dots, F^k are nontrivial free factors and if $F^1 * \dots * F^k$ is a free factor then we say that the collection $\{[[F^1]], \dots, [[F^k]]\}$ is a *free factor system*. For example, if G is a marked graph and $G_r \subset G$ is a subgraph with noncontractible components C_1, \dots, C_k then the conjugacy class $[[\pi_1(C_i)]]$ of the fundamental group of C_i is well defined and the collection of these conjugacy classes is a free factor system denoted $\mathcal{F}(G_r)$; we say that G_r *realizes* $\mathcal{F}(G_r)$.

The image of a free factor F under an element of $\text{Aut}(F_n)$ is a free factor. This induces an action of $\text{Out}(F_n)$ on the set of free factor systems. We sometimes say that a free factor is ϕ -invariant when we really mean that its conjugacy class is ϕ -invariant. If $[[F]]$ is ϕ -invariant then F is Φ -invariant for some automorphism Φ representing ϕ and $\Phi|_F$ determines a well defined element $\phi|_F$ of $\text{Out}(F)$.

We say that the conjugacy class $[a]$ of $a \in F_n$ is *carried by* $[[F^i]]$ if F^i contains a representative of $[a]$ and that an abstract line ℓ is carried by $[[F^i]]$ if it is the limit of periodic lines corresponding to conjugacy classes $[a_i]$ carried by $[[F^i]]$. A collection W of abstract lines and conjugacy classes in F_n is carried by a free factor system $\mathcal{F} = \{[[F^1]], \dots, [[F^k]]\}$ if each element of W is carried by some $[[F^i]]$. Sometimes we say that a is carried by F^i when we really mean that $[a]$ is carried by $[[F^i]]$. If G_r is a subgraph of a marked graph G then $[a]$ (resp. ℓ) is carried by $\mathcal{F}(G_r)$ if and only if the circuit (resp. line) in G that represents $[a]$ (resp. ℓ) is contained in G_r .

There is a partial order \sqsubset on conjugacy classes of free factors and on free factor systems generated by inclusion. More precisely, $[[F^1]] \sqsubset [[F^2]]$ if F^1 is conjugate to a free factor of F^2 and $\mathcal{F}_1 \sqsubset \mathcal{F}_2$ if for each $[[F^i]] \in \mathcal{F}_1$ there exists $[[F^j]] \in \mathcal{F}_2$ such that $[[F^i]] \sqsubset [[F^j]]$.

The *complexity* of a free factor system is defined on page 531 of [2]. We include the following results for easy reference. The first is [2, Corollary 2.6.5]. The second is an immediate consequence of the uniqueness of $\mathcal{F}(W)$.

Lemma 2.3 *For any collection W of abstract lines there is a unique free factor system $\mathcal{F}(W)$ of minimal complexity that carries every element of W . If W is a single element then $\mathcal{F}(W)$ has a single element.*

Corollary 2.4 *If a collection W of abstract lines and conjugacy classes in F_n is ϕ -invariant then $\mathcal{F}(W)$ is ϕ -invariant.*

Further details on free factor systems can be found in section 2.6 of [2].

Forward rotationless elements of $\text{Out}(F_n)$ and the Recognition Theorem In this section we recall a key definition and the main theorem of [10].

Definition 2.5 For $\Phi \in \text{Aut}(F_n)$ representing ϕ , let $\text{Fix}_N(\widehat{\Phi}) \subset \text{Fix}(\widehat{\Phi})$ be the set of nonrepelling fixed points of $\widehat{\Phi}$. We say that Φ is a *principal automorphism* and write $\Phi \in \mathcal{P}(\phi)$ if either of the following hold.

- $\text{Fix}_N(\widehat{\Phi})$ contains at least three points.
- $\text{Fix}_N(\widehat{\Phi})$ is a two point set that is neither the set of endpoints of an axis A_c nor the set of endpoints of a lift $\tilde{\lambda}$ of a generic leaf of an element of $\mathcal{L}(\phi)$.

The corresponding lift $\tilde{f}: \Gamma \rightarrow \Gamma$ is a *principal lift*.

There is an equivalence relation, called *isogredience*, on automorphisms defined by $\Phi_1 \sim i_c \Phi_2 i_c^{-1}$ for some $c \in F_n$. There are only finitely many isogredience classes of principal automorphisms. In fact by Levitt and Lustig [14], for all but finitely many isogredience classes, the only fixed points of $\widehat{\Phi}$ are a source and a sink.

We include the next lemma for easy reference.

Lemma 2.6 *The following properties hold for all Φ representing ϕ and Ψ representing ψ .*

- (1) $\text{Fix}(\widehat{\Psi\Phi\Psi^{-1}}) = \widehat{\Psi}(\text{Fix}(\widehat{\Phi}))$ and $\text{Fix}_N(\widehat{\Psi\Phi\Psi^{-1}}) = \widehat{\Psi}(\text{Fix}_N(\widehat{\Phi}))$.
- (2) *Conjugation by Ψ defines a bijection $i_\Psi: P(\phi) \mapsto P(\psi\phi\psi^{-1})$ that preserves isogredience classes. The induced bijection on the set of isogredience classes depends only on ψ and not on the choice of Ψ .*

Proof (1) is standard and easily checked; it implies that $i_\Psi: P(\phi) \mapsto P(\psi\phi\psi^{-1})$ is a bijection. The rest of (2) follows from $\Psi(i_c \Phi i_c^{-1})\Psi^{-1} = i_{\Psi(c)} \Psi\Phi\Psi^{-1} i_{\Psi(c)}^{-1}$ and $(i_d \Psi)(\Phi)(i_d \Psi)^{-1} = i_d(\Psi\Phi\Psi^{-1})i_d^{-1}$. \square

Definition 2.7 For $\Phi \in \text{Aut}(F_n)$ representing ϕ , let $\text{Per}_N(\widehat{\Phi})$ be the set of nonrepelling periodic points of $\widehat{\Phi}$. An outer automorphism ϕ is *forward rotationless* if $\text{Fix}_N(\widehat{\Phi}) = \text{Per}_N(\widehat{\Phi})$ for all $\Phi \in P(\phi)$ and if for each $k \geq 1$, $\Phi \mapsto \Phi^k$ defines a bijection between $P(\phi)$ and $P(\phi^k)$. Our standing assumption is that $n \geq 2$. For notational convenience we say that the identity element of $\text{Out}(F_1)$ is forward rotationless.

Remark 2.8 By [10, Lemma 4.43], there is a constant K , depending only on n , such that ϕ^K is forward rotationless for each $\phi \in \text{Out}(F_n)$.

As an illustration of the utility of being forward rotationless, and for convenient reference, we recall [10, Corollary 3.30].

Lemma 2.9 *The following properties hold for each forward rotationless $\phi \in \text{Out}(F_n)$.*

- (1) *Each periodic conjugacy class is fixed and each representative of that conjugacy class is fixed by some principal automorphism representing ϕ .*
- (2) *Each $\Lambda \in \mathcal{L}(\phi)$ is ϕ -invariant.*
- (3) *A free factor that is invariant under an iterate of ϕ is ϕ -invariant.*

The following theorem motivates the construction in Section 5 of a certain subgroup $\mathcal{A}(\phi)$ associated to ϕ and is applied in the proof that $\mathcal{A}(\phi)$ is abelian.

Theorem 2.10 (Recognition Theorem [10]) *Suppose that $\phi, \psi \in \text{Out}(F_n)$ are forward rotationless and that:*

- (1) $\text{PF}_\Lambda(\phi) = \text{PF}_\Lambda(\psi)$, for all $\Lambda \in \mathcal{L}(\phi) = \mathcal{L}(\psi)$.
- (2) *There is bijection $h: \text{P}(\phi) \rightarrow \text{P}(\psi)$ such that:*
 - (i) (Fixed sets preserved) $\text{Fix}_N(\widehat{\Phi}) = \text{Fix}_N(\widehat{h(\Phi)})$.
 - (ii) (Twist coordinates preserved) *If $u \in \text{Fix}(\Phi)$ and $\Phi, i_u\Phi \in \text{P}(\phi)$, then $h(i_u\Phi) = i_u h(\Phi)$.*

Then $\phi = \psi$.

Remark 2.11 In the special case that ϕ is realized as an element of $\text{MCG}(S)$, a u that occurs in item (2)(ii) has the form $u = w^d$ where w is root free and represents a reducing curve and where d is the degree of Dehn twisting about that reducing curve. See also the discussion of “axes” at the end of this section.

Relative train track maps We assume some familiarity with the basic definitions of relative train track maps. Complete details can be found in [10] and [2].

Suppose that $f: G \rightarrow G$ is a relative train track map defined with respect to a maximal filtration $\emptyset = G_0 \subset G_1 \subset \dots \subset G_N = G$. A path or circuit has *height* r if it is contained in G_r but not G_{r-1} . A lamination has height r if each leaf in its realization in G has height at most r and some leaf has height r . The r -th stratum H_r is defined to be the closure of $G_r \setminus G_{r-1}$. If $f(H_r) \subset G_{r-1}$ then H_r is called a *zero stratum*; all other strata have irreducible transition matrices and are said to be *irreducible*. If H_r is irreducible and if the Perron–Frobenius eigenvalue of the transition matrix for H_r is greater than one, then H_r is *exponentially growing* or simply EG. All other irreducible strata are *non-EG* or simply NEG.

A *direction* d at $x \in G$ is the germ of an initial segment of an oriented edge (or partial edge if x is not a vertex) based at x . There is an f -induced map Df on directions and we say that d is a *periodic direction* if it is periodic under the action of Df ; if the period is one then d is a *fixed direction*. Thus the direction determined by an oriented edge E is fixed if and only if E is the initial edge of $f(E)$.

A *turn* is an unordered pair of directions with a common base point. The turn is *nondegenerate* if it is defined by distinct directions and is *degenerate* otherwise. A turn is *illegal* with respect to $f: G \rightarrow G$ if its image under some iterate of Df is degenerate; a turn is *legal* if it is not illegal. If (d_1, d_2) is an illegal turn then the directions d_1 and d_2 are said to belong to the same *gate*. If $E_1 E_2 \dots E_{k-1} E_k$ is the edge path associated to a path σ , then we say that σ *contains the turns* (\bar{E}_i, E_{i+1}) for $1 \leq i \leq k-1$. A *path*

or circuit $\sigma \subset G$ is legal if it contains only legal turns. If $\sigma \subset G_r$ does not contain any illegal turns in H_r , meaning that both directions correspond to edges of H_r , then σ is r -legal. It is immediate from the definitions that Df maps legal turns to legal turns and that the restriction of f to a legal path is an immersion.

If H_r is EG then a nontrivial path in G_{r-1} with endpoints in $H_r \cap G_{r-1}$ is called a *connecting path*. As discussed below, connecting paths that are contained in zero strata play a special role.

For $a \in F_n$, we let $[a]_u$ be the *unoriented conjugacy class determined by a* . Thus, $[a]_u = [b]_u$ if and only if a is conjugate to either of a or \bar{a} . If σ is a closed path then we let $[\sigma]_u$ be the *unoriented conjugacy class determined by σ* , thought of as a circuit.

If an NEG stratum H_i is a single edge E_i satisfying $f(E_i) = E_i u_i$ for a nontrivial closed Nielsen path u_i then we say that E_i is a *linear edge* and we define the *axis or twistor* for E_i to be $[w_i]_u$ where w_i is root-free and $u_i = w_i^{d_i}$ for some $d_i \neq 0$. If E_i and E_j are distinct linear edges such that $w_i = w_j$ and such that d_i and d_j have the same sign then a path of the form $E_i w^p \bar{E}_j$ where $p \in \mathbb{Z}$, is called an *exceptional path of height $\max(i, j)$* or just an *exceptional path* if the height is not relevant. The set of exceptional paths of height i is invariant under the action of $f_\#$.

Notation 2.12 Suppose that $u < r$ and that:

- (1) H_u is irreducible.
- (2) H_r is EG and each component of G_r is noncontractible.
- (3) For each $u < i < r$, H_i is a zero stratum that is a component of G_{r-1} and each vertex of H_i has valence at least two in G_r .

We say that each H_i is *enveloped by H_r* and write $H_r^z = \bigcup_{k=u+1}^r H_k$. We say that H_r^z is the *extended EG stratum determined by H_r* .

Definition 2.13 If E is an edge in an irreducible stratum H_r and $k > 0$ then a maximal subpath σ of $f_\#^k(E)$ in a zero stratum H_i is said to be r -*taken* or just *taken* if r is irrelevant. Note that if H_i is enveloped by an EG stratum H_s then σ has endpoints in H_s and so is a connecting path. A nontrivial path or circuit σ is *completely split* if it has a splitting, called a *complete splitting*, into subpaths, each of which is either a single edge in an irreducible stratum, an indivisible Nielsen path, an exceptional path or a connecting path in a zero stratum H_i that is both maximal (meaning that it is not contained in a larger subpath of σ in H_i) and taken.

Definition 2.14 A relative train track map is *completely split* if:

- (1) $f(E)$ is completely split for each edge E in each irreducible stratum.
- (2) If σ is a taken connecting path in a zero stratum then $f_\#(\sigma)$ is completely split.

Remark 2.15 If $f: G \rightarrow G$ is a completely split relative train track map and σ is a completely split path or circuit then $f_{\#}(\sigma)$ is completely split. This is immediate from the definitions and the fact that $f_{\#}$ carries exceptional paths to exceptional paths.

Remark 2.16 If $f: G \rightarrow G$ is a CT (see Definition 2.20) then each completely split path or circuit has a unique complete splitting by [10, Lemma 4.12].

Definition 2.17 A periodic vertex w that does not satisfy one of the following two conditions is *principal*.

- w is the only element of $\text{Per}(f)$ in its Nielsen class and there are exactly two periodic directions at w , both of which are contained in the same EG stratum.
- w 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.

We also say that a lift of a principal vertex to the universal cover is a principal vertex. If each principal vertex and each periodic direction at a principal vertex has period one then we say that $f: G \rightarrow G$ is *rotationless*.

Remark 2.18 It is immediate from the definition that the initial endpoint of an NEG edge is a principal vertex. By [10, Lemma 3.19] every EG stratum H_r contains a principal vertex that is the basepoint for a periodic direction in H_r .

Complete details on principal vertices and rotationless relative train track maps, including the relationship between principal lifts and principal vertices and the relationship between forward rotationless outer automorphisms and rotationless relative train track maps can be found in [10, Section 3].

For any finite graph K , the *core of K* is the subgraph of K consisting of edges that are crossed by some circuit in K . The core of K contains no valence one vertices.

Definition 2.19 A filtration $\emptyset = G_0 \subset G_1 \subset \cdots \subset G_N = G$ that satisfies the following property is said to be *reduced (with respect to ϕ)*: if a free factor system \mathcal{F}' is ϕ^k -invariant for some $k > 0$ and if $\mathcal{F}(G_{r-1}) \sqsubset \mathcal{F}' \sqsubset \mathcal{F}(G_r)$ then either $\mathcal{F}' = \mathcal{F}(G_{r-1})$ or $\mathcal{F}' = \mathcal{F}(G_r)$.

We now recall the properties of a very useful kind of relative train track map and the existence theorem for relative train track maps with these properties.

Definition 2.20 A relative train track map $f: G \rightarrow G$ and filtration \mathcal{F} given by $\emptyset = G_0 \subset G_1 \subset \cdots \subset G_N = G$ is said to be a CT (for completely split improved relative train track map) if it satisfies the following properties.

- (1) (Rotationless) $f: G \rightarrow G$ is rotationless.
- (2) (Completely split) $f: G \rightarrow G$ is completely split.
- (3) (Filtration) \mathcal{F} is reduced. The core of each filtration element is a filtration element.
- (4) (Vertices) The endpoints of all indivisible periodic (necessarily fixed) Nielsen paths are (necessarily principal) vertices. The terminal endpoint of each nonfixed NEG edge is principal (and hence fixed).
- (5) (Periodic edges) Each periodic edge is fixed and each endpoint of a fixed edge is principal. If the unique edge E_r in a fixed stratum H_r is not a loop then G_{r-1} is a core graph and both ends of E_r are contained in G_{r-1} .
- (6) (Zero strata) If H_i is a zero stratum, then H_i is enveloped by an EG stratum H_r , each edge in H_i is r -taken and each vertex in H_i is contained in H_r and has link contained in $H_i \cup H_r$.
- (7) (Linear edges) For each linear E_i there is a closed root-free Nielsen path w_i such that $f(E_i) = E_i w_i^{d_i}$ for some $d_i \neq 0$. If E_i and E_j are distinct linear edges with the same axes then $w_i = w_j$ and $d_i \neq d_j$.
- (8) (NEG Nielsen paths) If the highest edges in an indivisible Nielsen path σ belong to an NEG stratum then there is a linear edge E_i with w_i as in (Linear edges) and there exists $k \neq 0$ such that $\sigma = E_i w_i^k \bar{E}_i$.
- (9) (EG Nielsen paths) If H_r is EG and ρ is an indivisible Nielsen path of height r , then $f|G_r = \theta \circ f_{r-1} \circ f_r$ where :
 - (a) $f_r: G_r \rightarrow G^1$ is a composition of proper extended folds defined by iteratively folding ρ .
 - (b) $f_{r-1}: G^1 \rightarrow G^2$ is a composition of folds involving edges in G_{r-1} .
 - (c) $\theta: G^2 \rightarrow G_r$ is a homeomorphism.

Theorem 2.21 [10, Theorem 4.29] Suppose that $\phi \in \text{Out}(F_n)$ is forward rotationless and that \mathcal{C} is a nested sequence of ϕ -invariant free factor systems. Then ϕ is represented by a CT $f: G \rightarrow G$ and filtration $\emptyset = G_0 \subset G_1 \subset \cdots \subset G_N = G$ such that the nested sequence of ϕ -invariant free factor systems defined by the G_j 's contains \mathcal{C} .

Iterating an edge We make frequent use of isolated points in $\text{Fix}_N(\hat{f})$ for principal lifts \tilde{f} . For the reader's convenience we quote three results that we refer to several times.

Lemma 2.22 *If $f: G \rightarrow G$ is a CT then \tilde{f} is a principal lift if and only if some (and hence every) element of $\text{Fix}(\tilde{f})$ is a principal vertex.*

Proof This follows from Remark 4.8, Corollary 3.22 and Corollary 3.27 of [10]. \square

Lemma 2.23 *The following properties hold for every principal lift $\tilde{f}: \Gamma \rightarrow \Gamma$ of a CT $f: G \rightarrow G$.*

- (1) *If $\tilde{v} \in \text{Fix}(\tilde{f})$ and a nonfixed edge \tilde{E} determines a fixed direction at \tilde{v} , then $\tilde{E} \subset \tilde{f}_\#(\tilde{E}) \subset \tilde{f}_\#^2(\tilde{E}) \subset \dots$ is an increasing sequence of paths whose union is a ray \tilde{R} that converges to some $P \in \text{Fix}_N(\hat{f})$ and whose interior is fixed point free. If \tilde{E} is a lift of an edge in an EG stratum then the accumulation set of P is the element in $\mathcal{L}(\phi)$ corresponding to that stratum.*
- (2) *For every isolated $P \in \text{Fix}_N(\hat{f})$ there exists \tilde{E} and \tilde{R} as in (1) that converges to P .*

Proof This is a combination of [10, Lemma 3.26] and [10, Lemma 4.37]. \square

If \tilde{E} and P are as in Lemma 2.23 then we say that \tilde{E} iterates to P and that P is associated to \tilde{E} .

Lemma 2.24 *Suppose that $\psi \in \text{Out}(F_n)$ is forward rotationless and that $P \in \text{Fix}_N(\hat{\Psi})$ for some $\Psi \in \mathcal{P}(\psi)$. Suppose further that Λ is an attracting lamination for some element of $\text{Out}(F_n)$, that Λ is ψ -invariant and that Λ is contained in the accumulation set of P . Then $\text{PF}_\Lambda(\psi) \geq 0$ and $\text{PF}_\Lambda(\psi) > 0$ if and only if P is isolated in $\text{Fix}_N(\hat{\Psi})$.*

Proof This is [10, Lemma 4.39]. \square

Axes Assume that ϕ is forward rotationless and that $f: G \rightarrow G$ is a CT. Following the notation of [3] we say that an unoriented conjugacy class μ of a root-free element of F_n is an *axis for ϕ* if for some (and hence any) representative $c \in F_n$ there exist distinct $\Phi_1, \Phi_2 \in \mathcal{P}(\phi)$ that fix c . Equivalently $\text{Fix}_N(\hat{\Phi}_1) \cap \text{Fix}_N(\hat{\Phi}_2)$ is the endpoint set of the axis A_c for T_c . The number of distinct elements of $\mathcal{P}(\phi)$ that fix c is called the *multiplicity* of μ . It is a consequence of Lemma 2.25 below that both the number of axes and the multiplicity of each axis is finite.

Lemma 2.9 implies that the oriented conjugacy class of c is ϕ -invariant. By Lemmas 4.1.4 and 4.2.6 of [2], the circuit γ representing c splits into a concatenation of subpaths α_i , each of which is either a fixed edge or an indivisible Nielsen path. (NEG Nielsen paths) and [10, Corollary 4.20] imply that each turn $(\bar{\alpha}_i, \alpha_{i+1})$ is legal. [10, Lemma 4.12(1)] therefore implies that this splitting is the complete splitting of γ .

There is an induced complete splitting of A_c into subpaths $\tilde{\alpha}_i$ that project to either fixed edges or indivisible Nielsen paths. The lift $\tilde{f}_0: \Gamma \rightarrow \Gamma$ that fixes the endpoints of each $\tilde{\alpha}_i$ is a principal lift by Lemma 2.22 and commutes with T_c . We say that \tilde{f}_0 and the corresponding $\Phi_0 \in P(\phi)$ are the *base lift* and *base principal automorphism* associated to μ and the choices of T_c and $f: G \rightarrow G$. By [10, Lemma 4.12(2)], for each $\tilde{\alpha}_i$ and for each $\tilde{x} \in \tilde{\alpha}_i$, the nearest point to $\tilde{f}_0(\tilde{x})$ in A_c is contained in $\tilde{\alpha}_i$. It follows that $\text{Fix}(T_c^j \tilde{f}_0) = \emptyset$ for all $j \neq 0$ and hence that \tilde{f}_0 is the only lift that commutes with T_c and has fixed points in A_c .

We recall [10, Lemma 4.14].

Lemma 2.25 *Suppose that ϕ is forward rotationless and that the unoriented conjugacy class μ is an axis for ϕ . Assume notation as above. There is a bijection between the set of principal lifts [principal automorphisms] $\tilde{f}_j \neq \tilde{f}_0$ [respectively $\Phi_j \neq \Phi_0 \in P(\phi)$] that commute with T_c [fix c] and the set of linear edges $\{E_j\}$ with axis equal to μ . Moreover, if $f(E_j) = E_j w_j^{d_j}$ then $\tilde{f}_j = T_c^{d_j} \tilde{f}_0$ [$\Phi_j = i_c^{d_j} \Phi_0$].*

3 Rotationless abelian subgroups

The Recognition Theorem is stated purely in terms of ϕ and its forward iterates. No condition on ϕ^{-1} is required. In the context of abelian subgroups, it is more natural to give ϕ and ϕ^{-1} equal footing.

Definition 3.1 $P^\pm(\phi) = P(\phi) \cup P(\phi^{-1})$. An outer automorphism ϕ is *rotationless* if it satisfies the following two conditions.

- (1) $\text{Fix}(\hat{\Phi}) = \text{Per}(\hat{\Phi})$ for all $\Phi \in P^\pm(\phi)$.
- (2) For each $k \geq 1$, $\Phi \mapsto \Phi^k$ defines a bijection (see Remark 3.2) between $P^\pm(\phi)$ and $P^\pm(\phi^k)$.

A subgroup of $\text{Out}(F_n)$ is *rotationless* if each of its elements is. Our standing assumption is that $n \geq 2$. For notational convenience we say that the identity element of $\text{Out}(F_1)$ is rotationless.

Remark 3.2 Assuming that the first item in Definition 3.1 is satisfied, the assignment $\Phi \mapsto \Phi^k$ defines an injection of $P^\pm(\phi)$ into $P^\pm(\phi^k)$. Indeed if $\Phi \mapsto \Phi^k$ is not injective then there exist distinct $\Phi_1, \Phi_2 \in P^\pm(\phi)$ and $k \geq 1$ such that $\text{Fix}(\widehat{\Phi}_1) = \text{Fix}(\widehat{\Phi}_1^k) = \text{Fix}(\widehat{\Phi}_2^k) = \text{Fix}(\widehat{\Phi}_2)$, which contains at least two points and is not the endpoint set of an axis, is contained in $\text{Fix}(\widehat{\Phi}_2 \widehat{\Phi}_1^{-1})$ in contradiction to the fact that $\Phi_2 \Phi_1^{-1}$ is a nontrivial covering translation. We may therefore replace the assumption in the second item of Definition 3.1 that the assignment $\Phi \mapsto \Phi^k$ defines a bijection with a priori weaker assumption that $\Phi \mapsto \Phi^k$ defines a surjection.

Remark 3.3 It is an immediate consequence of the definition that if ϕ is rotationless and Φ' is a principal lift of ϕ^k for $k \neq 0$ then ϕ has a principal lift Φ such that $\text{Fix}(\widehat{\Phi}) = \text{Fix}(\widehat{\Phi}')$.

The natural guess is that ϕ is rotationless if and only if ϕ and ϕ^{-1} are forward rotationless. The following lemma and corollary fall short of proving this (imagine Φ such that $\text{Fix}(\widehat{\Phi})$ consists of three fixed attractors and a repelling orbit of period two) but is sufficient for our needs.

Lemma 3.4 (1) *If ϕ is rotationless then ϕ and ϕ^{-1} are forward rotationless.*

(2) *If ϕ and ϕ^{-1} are forward rotationless and $(*)$ is satisfied for $\theta = \phi$ and $\theta = \phi^{-1}$ then ϕ is rotationless.*

(*) *For all $\Theta \in P(\theta)$, the set of repelling periodic points for $\widehat{\Theta}$ is not a period two orbit that is the endpoint set of a lift of a generic leaf γ of an element of $\mathcal{L}(\theta^{-1})$.*

Proof Assume that ϕ is rotationless. For $k > 0$, each element of $P(\phi^k)$ has the form Φ^k where $\text{Fix}(\widehat{\Phi}) = \text{Per}(\widehat{\Phi})$ and hence $\text{Fix}_N(\widehat{\Phi}) = \text{Per}_N(\widehat{\Phi}^k)$. Thus $\Phi \in P(\phi)$ proving that ϕ is forward rotationless. The symmetric argument showing that ϕ^{-1} is forward rotationless completes the proof of (1).

Assume now that the hypotheses of (2) are satisfied, that $k \geq 1$ and that $\Phi_k \in P^\pm(\phi^k)$. The plus and minus cases are symmetric so we may assume that $\Phi_k \in P(\phi^k)$. Since ϕ is forward rotationless, $\Phi_k = \Phi^k$ for some $\Phi \in P(\phi)$ satisfying $\text{Fix}_N(\widehat{\Phi}) = \text{Per}_N(\widehat{\Phi})$. To prove that $\text{Fix}(\widehat{\Phi}) = \text{Per}(\widehat{\Phi})$ it suffices to show that all periodic repelling points for $\widehat{\Phi}$ have period one. Since ϕ^{-1} is forward rotationless, the only way this could fail would be if the repelling set is a period two orbit and if $\Phi^2 \notin P(\phi^{-1})$. This possibility is ruled out by (*). \square

Corollary 3.5 *If ϕ and ϕ^{-1} are forward rotationless then ϕ^2 is rotationless. There exists $k > 0$, depending only on n , so that ϕ^{2k} is rotationless for every $\phi \in \text{Out}(F_n)$.*

Proof The first statement follows from Lemma 3.4 and the second statement from Remark 2.8. \square

Example 3.6 Let G be the graph with one vertex v and edges labelled A , B and C . Let $f: G \rightarrow G$ be the homotopy equivalence defined by

$$A \mapsto B^3 A \quad B \mapsto C^3 B \quad C \mapsto (B^3 A)^3 C.$$

The directions at v determined by \bar{A} , \bar{B} and \bar{C} are fixed by Df and those determined by B and C are interchanged by Df . Thus f is not rotationless and the outer automorphism ϕ that it determines is neither forward rotationless nor rotationless. The map f factors as $f_3 f_2 f_1$ where f_1 fixes A and B and $f_1(C) = A^3 C$, f_2 fixes A and C and $f_2(B) = C^3 B$ and f_3 fixes B and C and $f_3(A) = B^3 A$. It is easy to check that each of these homotopy equivalence determines a rotationless element of $\text{Out}(F_n)$. This shows that the composition of rotationless elements need not be rotationless. Obviously, ϕ induces the identity on $H_1(G, \mathbb{Z}_3)$ and so illustrates that not every such element is rotationless. We will see (Corollary 3.13) that the composition of commuting rotationless elements is rotationless.

Lemma 3.7 *If ϕ is rotationless and if Φ represents ϕ , then $\text{Fix}_N(\hat{\Phi}^2) \neq \emptyset$. If, in addition, $\Phi \in P^\pm(\phi)$ then $\text{Fix}_N(\hat{\Phi})$ and $\text{Fix}_N(\hat{\Phi}^{-1})$ are nonempty.*

Proof The second statement follows from the first and the assumption that ϕ is rotationless.

Choose $f: G \rightarrow G$ representing ϕ and let $\tilde{f}: \Gamma \rightarrow \Gamma$ be the lift corresponding to Φ . It suffices to show that $\text{Fix}_N(\hat{f}^2) \neq \emptyset$. If $\text{Fix}(\tilde{f}) = \emptyset$ then $\text{Fix}_N(\hat{f}) \neq \emptyset$ by [10, Lemmas 3.23 and 3.15] and we are done. If $\tilde{x} \in \text{Fix}(\tilde{f})$ and \tilde{f} fixes a direction at \tilde{x} then $\text{Fix}_N(\hat{f}) \neq \emptyset$ by [10, Lemma 3.26] and again we are done. Since ϕ is rotationless, the only remaining case is that there are exactly two \tilde{f} periodic directions at \tilde{x} . These directions are fixed by \tilde{f}^2 so a second application of [10, Lemma 3.26] completes the proof. \square

Abelian subgroups of $\text{Out}(F_n)$ are finitely generated [1]. Thus given any generating set for an abelian subgroup, there is a finite subset which also generates. At the end of this section (Corollary 3.13) we prove that an abelian subgroup A of $\text{Out}(F_n)$ that is generated by rotationless elements, is rotationless.

Many of our arguments proceed by induction on the cardinality of a given set of rotationless generators.

Lemma 3.8 *If ϕ is rotationless and F is a ϕ -invariant free factor then $\phi|F$ is rotationless.*

Proof If F has rank one then this follows from the first item of Lemma 2.9 and our convention that the identity element of $\text{Out}(F_1)$ is rotationless. If F has rank at least two then every automorphism Φ_F representing $\phi|F$ extends uniquely to an automorphism $\hat{\Phi}$ representing ϕ because no nontrivial covering translation restricts to the identity on F . Since $\text{Fix}_N(\widehat{\Phi}_F) \subset \text{Fix}_N(\hat{\Phi})$, we have that $\hat{\Phi}$ is principal if Φ_F is principal and $\phi|F$ is rotationless if ϕ is rotationless. \square

We study lifts of an abelian subgroup of $\text{Out}(F_n)$ to $\text{Aut}(F_n)$ that is generated by rotationless elements via the following definition and lemma.

Definition 3.9 A set $\mathcal{X} \subset \partial F_n$ with at least three points is a *principal set* for an abelian subgroup A of $\text{Out}(F_n)$ if each $\psi \in A$ is represented by $\Psi \in \text{Aut}(F_n)$ satisfying $\mathcal{X} \subset \text{Fix}(\hat{\Psi})$ and if this necessarily unique Ψ is an element of $P^\pm(\psi)$. The assignment $\psi \mapsto \Psi$ is a lift of A from $\text{Out}(F_n)$ to $\text{Aut}(F_n)$.

Remark 3.10 If \mathcal{X} a principal set for A and $\mathcal{X}_0 \subset \mathcal{X}$ contains at least three points then \mathcal{X}_0 is a principal set for A and \mathcal{X}_0 and \mathcal{X} determine the same lift of A to $\text{Aut}(F_n)$. If $\psi \mapsto \Psi$ is the lift of A determined by \mathcal{X} then $\bigcap_{\psi \in A} \text{Fix}(\hat{\Psi})$ is the unique maximal principal set containing \mathcal{X} .

Lemma 3.11 *Suppose that A is an abelian subgroup of $\text{Out}(F_n)$ that is generated by rotationless elements, that $\phi \in A$ is rotationless and that $\Phi \in P^\pm(\phi)$. Let $\mathbb{F} = \text{Fix}(\Phi)$.*

- (1) *If \mathbb{F} has rank zero then $\text{Fix}(\hat{\Phi})$ is a principal set for A .*
- (2) *If \mathbb{F} has rank one with generator c and if P is an isolated point in $\text{Fix}(\hat{\Phi})$ then $\{P, T_c^\pm\}$ is a principal set for A .*
- (3) *If \mathbb{F} has rank at least two then $\partial \text{Fix}(\Phi)$ contains at least one principal set \mathcal{X} for A and one can choose X to contain T_c^\pm for any given A -invariant $[c]$ with $c \in \mathbb{F}$. Moreover, for every isolated point P in $\text{Fix}(\hat{\Phi})$ there is a principal set \mathcal{Y} for A that contains P and at least two elements of $\partial \mathbb{F}$.*

In particular, $\text{Fix}(\hat{\Phi})$ contains at least one principal set for A and every isolated point in $\text{Fix}(\hat{\Phi})$ is contained in such a principal set. If $s: A \rightarrow \text{Aut}(F_n)$ is the lift determined by a principal set contained in $\text{Fix}(\hat{\Phi})$ then $s(\phi) = \Phi$.

Proof Let S be a finite rotationless generating set for A . For each $\psi \in S$, we choose an automorphism Ψ_k that commutes with Φ and that representing ψ^k for some $k \geq 1$ as follows. Begin with any Ψ representing ψ . Lemma 2.6 implies that conjugation by Ψ defines a permutation of the finite set of isogredience classes in $P^\pm(\phi)$. Choose $k > 0$ so that the permutation induced by Ψ^k is trivial. Then $\Psi^k \Phi \Psi^{-k} = i_b \Phi i_b^{-1}$ for some $b \in F_n$ and $\Psi_k := i_b^{-1} \Psi^k$ commutes with Φ . In particular, \mathbb{F} is Ψ_k -invariant.

Assume at first that \mathbb{F} has rank zero. By Lemma 2.2, $\text{Fix}(\widehat{\Phi})$ is a finite union of attractors and repellers and by Lemma 3.7 there is at least one of each. Since $\Phi \in P^\pm(\phi)$, there are at least three points in $\text{Fix}(\widehat{\Phi})$.

We claim that if Θ represents $\theta \in A$ and if $\text{Fix}(\widehat{\Phi}) \subset \text{Fix}(\widehat{\Theta})$ then $\Theta \in P^\pm(\theta)$. If $\text{Fix}(\widehat{\Theta})$ contains at least five points then this is obvious. After replacing θ with its inverse if necessary, there are two potentially bad cases. The first is that $\text{Fix}(\widehat{\Theta})$ has exactly one repelling point and exactly two attracting points and that the attractors bound a lift $\tilde{\gamma}$ of a generic leaf of some $\Lambda \in \mathcal{L}(\theta)$. Since the endpoints of $\tilde{\gamma}$ are isolated fixed points of $\widehat{\Phi}$, $\Lambda \in \mathcal{L}(\phi) \cup \mathcal{L}(\phi^{-1})$ by Lemma 2.24. After replacing ϕ with its inverse if necessary, we may assume that $\Lambda \in \mathcal{L}(\phi)$ and that the endpoints of $\tilde{\gamma}$ are attractors for Φ . Since $\text{Fix}(\widehat{\Phi})$ contains only three points and by Lemma 3.7 has at least one $\widehat{\Phi}$ -repeller, this contradicts the assumption that $\Phi \in P^\pm(\phi)$.

The other bad possibility is that $\text{Fix}(\widehat{\Theta})$ is a four point set with two repelling points that bound a lift of a leaf of an element of $\mathcal{L}(\theta^{-1})$ and two attracting points that bound a lift of a leaf of an element of $\mathcal{L}(\theta)$. As in the previous case, this description also applies to Φ in contradiction to the assumption that $\Phi \in P^\pm(\phi)$. This completes the proof that $\Theta \in P^\pm(\theta)$.

After replacing Ψ_k with an iterate, we may assume that $\text{Fix}(\widehat{\Phi}) \subset \text{Fix}(\widehat{\Psi}_k)$ and hence that $\Psi_k \in P^\pm(\psi^k)$. Since ψ is rotationless, there exists $\Psi \in P^\pm(\psi)$ with $\text{Fix}(\widehat{\Phi}) \subset \text{Fix}(\widehat{\Psi})$. As this holds for every element of S , we have proved (1).

Suppose next that \mathbb{F} has rank one with generator c and that P is an isolated point in $\text{Fix}(\widehat{\Phi})$. Lemma 2.2 implies that there are only finitely many i_c -orbits of isolated points in $\text{Fix}(\widehat{\Phi})$. After increasing k if necessary, we may assume that $c \in \text{Fix}(\Psi_k)$ and that Ψ_k preserves each such i_c -orbit. In particular, $\widehat{\Psi}_k(P) = \widehat{T}_c^q(P)$ for some q . Let $\Psi'_k := i_c^{-q} \Psi_k$. Then $\{T_c^\pm, P\} \subset \text{Fix}(\widehat{\Psi}'_k)$ and $\Psi'_k \in P^\pm(\psi)$. Since ψ is rotationless, there exists $\Psi \in P^\pm(\psi)$ such that $\{T_c^\pm, P\} \subset \text{Fix}(\widehat{\Psi})$. As this holds for every element of S , it follows that for each $\theta \in A$ there exists Θ such that $\{T_c^\pm, P\} \subset \text{Fix}(\widehat{\Theta})$. In this case it is obvious that $\Theta \in P^\pm(\theta)$. This completes the proof of (2).

We turn next to the moreover part of (3). Assume that P is an isolated point in $\text{Fix}(\widehat{\Phi})$. As in the rank one case, the fact that there are only finitely many \mathbb{F} -orbits of isolated

points in $\text{Fix}(\Phi)$ allows us to choose Ψ_k^* representing an iterate ψ^k of ψ such that $P \in \text{Fix}(\widehat{\Psi}_k^*)$ and such that \mathbb{F} is Ψ_k^* -invariant. We claim that $(\Psi_k^*)^2 \in P^\pm(\psi^{2k})$. Lemmas 3.8 and 3.7 together imply that $(\Psi_k^*)^2|_{\mathbb{F}}$ has at least one fixed nonattractor Q_- and symmetrically one fixed nonrepeller Q_+ . Lemma 2.24 implies that Q_+ and Q_- do not cobound a lift of a generic leaf of an attracting lamination. (This method for proving that a pair of points do not cobound a lift of a generic leaf of an attracting lamination is used implicitly throughout the rest of the proof.) Generic leaves of an attracting lamination are birecurrent and so either have both endpoints in $\partial\mathbb{F}$ or neither endpoint in $\partial\mathbb{F}$. Thus P and Q_\pm do not cobound a lift of a generic leaf of an attracting lamination. This verifies our claim. Since ψ is rotationless, there exists $\Psi \in P^\pm(\psi)$ with $\{P, Q_+, Q_-\} \subset \text{Fix}(\widehat{\Psi})$. These three points are also in $\text{Fix}(\widehat{\Phi})$. It follows that Ψ commutes with Φ and hence that \mathbb{F} is Ψ -invariant.

We have shown that if $S = \{\psi_1, \dots, \psi_K\}$ then for all $1 \leq j \leq K$ there exists $\Psi_j \in P^\pm(\psi_j)$ such that $P \in \text{Fix}(\widehat{\Psi}_j)$ and such that \mathbb{F} is Ψ_j -invariant. Item (i) of Lemma 2.2 implies that P is not fixed by any covering translation and hence that the Ψ_j 's commute.

We produce the desired principal set \mathcal{Y} by induction on j . To this end, let

$$\mathcal{Y}_j = \left(\bigcap_{i=1}^j \text{Fix}(\widehat{\Psi}_i) \right) \cap \partial\mathbb{F} = \bigcap_{i=1}^j \text{Fix}(\widehat{\Psi}_i|_{\mathbb{F}}),$$

$$\mathbb{F}_j = \bigcap_{i=1}^j \text{Fix}(\Psi_i|_{\mathbb{F}}),$$

and let I_j be the statement that \mathbb{F}_j is finitely generated and that \mathcal{Y}_j either contains at least three points or contains two points that do not cobound a lift of a generic leaf of any attracting lamination. If \mathcal{Y}_j contains the endpoint set of an axis then $\bigcap_{i=1}^j \text{Fix}(\widehat{\Psi}_i)$ is infinite. As noted above, P and an element of $\partial\mathbb{F}$ can not cobound a generic leaf of an attracting lamination or any axis. Thus I_K completes the proof of the moreover part of (3).

I_1 follows from Lemma 3.7 applied to $\Psi_1|_{\mathbb{F}}$. Assume that I_{j-1} holds. \mathcal{Y}_{j-1} is $\widehat{\Psi}_j$ -invariant and \mathbb{F} is Ψ_j -invariant. If \mathcal{Y}_{j-1} is finite then it is fixed by an iterate of $\widehat{\Psi}_j$ and hence by $\widehat{\Psi}_j$; in this case \mathbb{F}_j has rank zero. If \mathcal{Y}_{j-1} contains T_b^\pm for some unique root-free unoriented b then T_b^\pm is fixed by an iterate of $\widehat{\Psi}_j$ and hence by $\widehat{\Psi}_j$; in this case \mathbb{F}_j has rank one. In either case I_j holds. In the remaining case \mathbb{F}_{j-1} has rank at least two and I_j follows from $\mathbb{F}_j = \text{Fix}(\Psi_j|_{\mathbb{F}_{j-1}})$ and from Lemma 3.7 applied to $\widehat{\Psi}_j|_{\mathbb{F}_{j-1}}$, keeping in mind that $\text{Fix}(\widehat{\Psi}_j|_{\mathbb{F}_{j-1}}) \subset \mathcal{Y}_j$. This completes the induction step and so proves I_K .

It remains to prove the main statement of (3). We argue by induction on the cardinality K of our given rotationless generating set S for A . If $K = 1$ and $S = \{\psi\}$ then there

exists $\Psi \in P^\pm(\psi)$ such that $\text{Fix}(\widehat{\Psi}) = \text{Fix}(\widehat{\Phi})$ and $\text{Fix}(\widehat{\Psi})$ is obviously a principal set for A . We now assume that $K \geq 2$ and that (3) holds for subgroups that are generated by fewer than K rotationless elements.

The defining property of Ψ_k is that it commutes with Φ . We may therefore replace our current Ψ_k with any lift of any iterate of ψ that preserves \mathbb{F} . By Lemma 5.2 of [3] or Proposition 9.4 of [15], there is such a lift, still called Ψ_k , such that $\Psi_k|_{\mathbb{F}} \in P^\pm(\psi^k|_{\mathbb{F}})$; moreover if $c \in \text{Fix}(\Phi)$ is A -invariant then we may choose Ψ_k so that $c \in \text{Fix}(\Psi_k)$. Since ψ is rotationless, there exists $\Psi \in P^\pm(\psi)$ such that $\text{Fix}(\widehat{\Psi}) = \text{Fix}(\widehat{\Psi}_k)$. Thus $\text{Fix}(\widehat{\Psi}) \cap \text{Fix}(\widehat{\Phi})$ contains at least three points which implies that Ψ commutes with Φ . To summarize, we have $\Psi \in P^\pm(\psi)$ that preserves \mathbb{F} and such that $\Psi|_{\mathbb{F}} \in P^\pm(\psi|_{\mathbb{F}})$; if $c \in \text{Fix}(\Phi)$ is A -invariant then we may assume that $c \in \text{Fix}(\Psi)$. As each Ψ preserves \mathbb{F} , it follows that $[\mathbb{F}]$ is A -invariant.

Let $A' = A|_{\mathbb{F}}$, let $\psi' = \psi|_{\mathbb{F}}$ and let $\Psi' = \Psi|_{\mathbb{F}}$. As noted in the proof of Lemma 3.8, a principal set for A' is also a principal set for A . To prove the existence of a principal set \mathcal{X} (containing T_c^\pm) for A it suffices to prove the existence of a principal set \mathcal{X}' (containing T_c^\pm) for A' . If $\text{Fix}(\Psi')$ has rank less than two then the existence of \mathcal{X}' follows from (1) and (2) applied to $\Psi' \in A'$. Suppose then that $\text{Fix}(\Psi')$ has rank at least two. By the same logic, it is sufficient to find a principal set \mathcal{X}'' (containing T_c^\pm) for $A'|_{\text{Fix}(\Psi')}$ and this exists by the inductive hypothesis and the fact that $A'|_{\text{Fix}(\Psi')}$ has a rotationless (by Lemma 3.8) generating set with fewer than K elements. \square

Lemma 3.12 *An abelian subgroup A that is generated by rotationless elements is torsion free.*

Proof If $\theta \in A$ is a torsion element then [6] it is represented by a finite order homeomorphism $f': G' \rightarrow G'$ of a marked graph G' . Suppose that \mathcal{X} is a principal set for A and that $P_1, P_2, P_3 \in \mathcal{X}$. There is a lift $\tilde{f}': \Gamma' \rightarrow \Gamma'$ such that each $P_i \in \text{Fix}(\tilde{f}')$. The line L_{12} with endpoints P_1 and P_2 and the line L_{13} with endpoints P_1 and P_3 are $\tilde{f}'_\#$ -invariant and since \tilde{f}' is a homeomorphism they are \tilde{f}' -invariant. The intersection $L_{12} \cap L_{13}$ is an \tilde{f}' -invariant ray and so is contained in $\text{Fix}(\tilde{f}')$. It follows that $L_{12} \subset \text{Fix}(\tilde{f}')$ and that the image of L_{12} in G' is contained in $\text{Fix}(f')$. It therefore suffices to show that every edge of G' is crossed by at least one such line.

For any set $\mathcal{Y} \subset \partial F_n$, let $C_{\mathcal{Y}}$ be the set of bi-infinite lines cobounded by pairs of elements of \mathcal{Y} . Let $W_A = \cup C_{\mathcal{X}}$ where the union is over all principal sets \mathcal{X} for A and let \mathcal{F} be the smallest free factor system that carries W_A . It suffices to show that $\mathcal{F} = \{[F_n]\}$. The proof of this assertion is by induction on the cardinality K of a given rotationless generating set S for A .

Assume to the contrary that \mathcal{F} is proper and choose $\psi \in S$. Choose a CT $f: G \rightarrow G$ representing ψ in which \mathcal{F} is realized as a filtration element G_r . Lemma 2.23(1) implies that each $\Lambda \in \mathcal{L}(\psi)$ is the accumulation set of an isolated point in $\text{Fix}_N(\widehat{\Psi})$ for some $\Psi \in P(\psi)$. By Lemma 3.11, P is contained in some principal set \mathcal{X} . This implies that \mathcal{F} carries a line that accumulates on Λ and so carries Λ . Thus each stratum above G_r is NEG. Items (Rotationless) and (Periodic edges) of Definition 2.20 and Remark 2.18 imply that every edge E of $G \setminus G_r$ has an orientation so that its initial vertex is principal and so that its initial direction is fixed. Choose a lift \tilde{E} of E and let $\tilde{f}: \Gamma \rightarrow \Gamma$ be the principal lift that fixes the initial direction determined by \tilde{E} . There is a ray that begins with \tilde{E} and converges to a point in $\text{Fix}_N(\widehat{f})$. This follows from Lemma 2.23 if E is not a fixed edge and from Lemma 3.26 of [10] otherwise. Let Ψ be the principal automorphism corresponding to \tilde{f} . By Lemma 4.15 of [10] there is at least one other fixed direction based at the initial vertex of \tilde{E} . Applying the same argument to this direction, we see that some element of $C_{\text{Fix}(\widehat{\Psi})}$ crosses \tilde{E} . It therefore suffices to show that each element of $C_{\text{Fix}(\widehat{\Psi})}$ is carried by \mathcal{F} . This is obvious if $K = 1$. We have now proved the basis step of our induction argument and may assume that $K > 1$ and that $\mathcal{F} = \{[F_n]\}$ when A has a rotationless generating set with fewer than K elements.

If $\text{Fix}(\Psi)$ has rank zero then $\text{Fix}(\widehat{\Psi})$ is a principal set for A by Lemma 3.11(1) and $C_{\text{Fix}(\widehat{\Psi})}$ is carried by \mathcal{F} . If $\text{Fix}(\Psi)$ has rank one with generator c then Lemma 3.11(2) implies that the line connecting P to T_c^+ is carried by \mathcal{F} for each $P \in \text{Fix}(\widehat{\Psi})$. It follows that the line connecting any two points of $\text{Fix}(\widehat{\Psi})$ is carried by \mathcal{F} .

We may therefore assume that $\text{Fix}(\Psi)$ has rank at least two. Let us show that $\text{Fix}(\Psi)$ is carried by \mathcal{F} . Lemma 3.8 implies that $A|_{\text{Fix}(\Psi)}$ has a rotationless generating set with fewer than K elements. The inductive hypothesis therefore implies that no proper free factor system of $\text{Fix}(\Psi)$ carries $W_{A|_{\text{Fix}(\Psi)}}$. The Kurosh subgroup theorem therefore implies that any free factor system of F_n that carries $W_{A|_{\text{Fix}(\Psi)}}$ also carries all of $\text{Fix}(\Psi)$. Since $W_{A|_{\text{Fix}(\Psi)}} \subset W_A$ we conclude that $\text{Fix}(\Psi)$ is carried by \mathcal{F} .

Lemma 3.11(3) implies that for each $P \in \text{Fix}(\widehat{\Psi})$ there exists $Q \in \partial \text{Fix}(\Psi)$ so that the line connecting P to Q is carried by \mathcal{F} . Since the line connecting any two points in $\partial \text{Fix}(\Psi)$ is carried by \mathcal{F} it follows that the line connecting any two points in $\text{Fix}(\widehat{\Psi})$ is carried by \mathcal{F} . \square

Corollary 3.13 *An abelian subgroup A that is generated by rotationless elements is rotationless.*

Proof Suppose that $\phi \in A$, that $k > 1$ and that $\Phi_k \in P^\pm(\phi^k)$. Choose $m \geq 1$ so that ϕ^{km} is rotationless. By Lemma 3.11 there is a principal set \mathcal{X} for A with

$\mathcal{X} \subset \text{Fix}(\widehat{\Phi}_k^m)$. Let $s: A \rightarrow \text{Aut}(F_n)$ be the lift determined by \mathcal{X} and let $\Phi = s(\phi)$. Then $\Phi^{km} = s(\phi^k)^m = \Phi_k^m$ and so $\Phi^k = s(\phi_k) = \Phi_k$ by Lemma 3.12. To complete the proof it suffices by Remark 3.2 to show that $\Phi \in P^\pm(\phi)$ and for this it suffices to show that $\text{Fix}(\widehat{\Phi}^{km}) \subset \text{Fix}(\widehat{\Phi})$.

Since $s(A)$ is abelian, $\mathbb{F} := \text{Fix}(\Phi^{km})$ is $s(A)$ -invariant. Lemma 3.12 implies that Φ is uniquely characterized by $\Phi^{km} = \Phi_k^m$ and hence that Φ is independent of the choice of the principal set $\mathcal{X} \subset \text{Fix}(\widehat{\Phi}_k^m)$ for A . Thus each $\mathcal{X} \subset \text{Fix}(\widehat{\Phi})$. Lemma 3.11 implies that $\text{Fix}(\widehat{\Phi})$ contains each isolated point in $\text{Fix}(\widehat{\Phi}_k^m)$ so it remains to show that $\partial\mathbb{F} \subset \text{Fix}(\widehat{\Phi})$. This follows from (1) and (2) of Lemma 3.11 if \mathbb{F} has rank less than two and from Lemma 3.12 applied to $A|\mathbb{F}$ if \mathbb{F} has rank at least two. \square

Corollary 3.14 *For each abelian subgroup A of $\text{Out}(F_n)$, the set of rotationless elements is a rotationless subgroup A_R that has finite index in A .*

Proof This follows immediately from Corollaries 3.5 and 3.13 and the fact that A is finitely generated. \square

4 Generic elements of rotationless abelian subgroups

In this section we define an embedding of a given rotationless abelian subgroup A into an integer lattice \mathbb{Z}^N and say what it means for an element of A to be generic with respect to this embedding.

Definition 4.1 Suppose that \mathcal{X}_1 and \mathcal{X}_2 are principal sets for A that define distinct lifts s_1 and s_2 of A to $\text{Aut}(F_n)$ and that $T_c^\pm \in \mathcal{X}_1 \cap \mathcal{X}_2$. Then i_c commutes with $s_1(\psi)$ and with $s_2(\psi)$ and $s_2(\psi) = i_c^{d(\psi)} s_1(\psi)$ for all $\psi \in A$ and some $d(\psi) \in \mathbb{Z}$; the assignment $\psi \mapsto d(\psi)$ defines a homomorphism that we call the *comparison homomorphism* $\omega: A \rightarrow \mathbb{Z}$ determined by \mathcal{X}_1 and \mathcal{X}_2 .

Remark 4.2 Principal sets \mathcal{X}_1 and \mathcal{X}_2 for A define distinct lifts of A to $\text{Aut}(F_n)$ if and only if $\mathcal{X}_1 \cup \mathcal{X}_2$ is not a principal set for A .

Lemma 4.3 *For any rotationless abelian subgroup A there are only finitely many comparison homomorphisms $\omega: A \rightarrow \mathbb{Z}$.*

Proof Distinct comparison homomorphisms must disagree on some element of each basis of A so we can restrict attention to those comparison homomorphisms that disagree on a single element $\psi \in A$. If ω is defined with respect to \mathcal{X}_1 , \mathcal{X}_2 and c

then $[c]_u$, the unoriented conjugacy class of c , is an axis of ψ . By Lemma 2.25, ψ has only finitely many axes. We may therefore restrict attention to those comparison homomorphisms that are defined with respect to the same $[c]_u$. If $a \in F_n$ and if ω' is defined with respect to $\hat{i}_a \mathcal{X}_1, \hat{i}_a \mathcal{X}_2$ and $i_a(c)$ then $\omega' = \omega$. We may therefore restrict attention to comparison homomorphisms that are defined with respect to the same c . The number of such comparison homomorphisms is bounded by the multiplicity of $[c]_u$ as an axis for ψ by Lemma 2.25. \square

Lemma 4.4 *If A is a rotationless abelian subgroup then $\mathcal{L}(A) = \bigcup_{\phi \in A} \mathcal{L}(\phi)$ is a finite collection of A -invariant laminations.*

Proof Let $\{\psi_1, \dots, \psi_K\}$ be a basis for A . If $\mathcal{L}(\phi) = \{\Lambda_1, \dots, \Lambda_q\}$ and $F(\Lambda_i)$ is the smallest free factor that carries Λ_i then the $F(\Lambda_i)$'s are distinct by Lemma 3.2.4 of [2]. Each ψ_j permutes the Λ_i 's by Lemma 3.1.6 of [2] and so permutes the $F(\Lambda_i)$'s by Corollary 2.4. Since ψ_j is rotationless, each $F(\Lambda_i)$, and hence each Λ_i , is ψ_j -invariant by Lemma 2.9. This proves that Λ_i is A -invariant and hence that PF_{Λ_i} is defined on A . Each PF_{Λ_i} must be nonzero when applied to some ψ_j and by Corollary 3.3.1 of [2] this is equivalent to $\Lambda_i \in \mathcal{L}(\psi_j) \cup \mathcal{L}(\psi_j^{-1})$, which is a finite set. \square

Definition 4.5 For each $\Lambda \in \mathcal{L}(A)$, we say that $PF_{\Lambda}|_A$ is the *expansion factor homomorphism for A* determined by Λ . Let N be the number of distinct comparison and expansion factor homomorphisms for A . Define $\Omega: A \rightarrow \mathbb{Z}^N$ to be the product of these homomorphisms. We say that Ω is the *coordinate homomorphism for A* and that each comparison homomorphism and expansion factor homomorphism is a *coordinate* of Ω .

Lemma 4.6 *If A is a rotationless abelian subgroup then $\Omega: A \rightarrow \mathbb{Z}^N$ is injective.*

Proof Given nontrivial $\theta \in A$, choose a CT $f: G \rightarrow G$ and filtration $\emptyset = G_0 \subset G_1 \subset \dots \subset G_N = G$ representing θ and let H_l be the lowest nonfixed irreducible stratum. If H_l is EG then $PF_{\Lambda}(\theta) \neq 0$ for the attracting lamination $\Lambda \in \mathcal{L}(\theta)$ associated to H_l . Otherwise H_l is a single edge E and $f(E) = E \cdot u$ where $u \subset G_{l-1}$ is a loop that is fixed by f .

Choose a lift $\tilde{E} \subset \Gamma$, let \tilde{u} be the lift of u whose initial endpoint is the terminal endpoint of \tilde{E} and let T_c be the covering translation that carries the initial endpoint of \tilde{u} to the terminal endpoint of \tilde{u} . The initial and terminal endpoints of \tilde{E} are principal; the former by Remark 2.18 and the latter by property (Vertices) in the definition of CT. Lemma 2.22 implies that the lifts $\tilde{f}_1: \Gamma \rightarrow \Gamma$ and $\tilde{f}_2: \Gamma \rightarrow \Gamma$ of f that fix the

initial and terminal endpoints of \tilde{E} respectively are principal. By construction, \tilde{f}_1 and \tilde{f}_2 are distinct and commute with T_c . By Lemma 3.11 there exist principal sets $\mathcal{X}_1 \subset \text{Fix}(\hat{f}_1)$ and $\mathcal{X}_2 \subset \text{Fix}(\hat{f}_2)$ that contain T_c^\pm . Since $\tilde{f}_1 \neq \tilde{f}_2$, θ is not contained in the kernel of the comparison homomorphism determined by \mathcal{X}_1 and \mathcal{X}_2 . We have shown that some coordinate of $\Omega(\theta) \neq 0$ and since θ was arbitrary, Ω is injective. \square

Definition 4.7 Assume that A is a rotationless abelian subgroup and that $\Omega: A \rightarrow \mathbb{Z}^N$ is its coordinate homomorphism. Then $\phi \in A$ is *generic* if all coordinates of $\Omega(\phi)$ are nonzero.

Remark 4.8 For $\Lambda \in \mathcal{L}(A)$ and $\phi \in A$, Corollary 3.3.1 of [2] implies that $PF_\Lambda(\phi) \neq \emptyset$ if and only if $\Lambda \in \mathcal{L}(\phi) \cup \mathcal{L}(\phi^{-1})$. Thus ϕ is generic in A if and only if “ ϕ has the same axes and multiplicity as A ” and $\mathcal{L}(A) = \mathcal{L}(\phi) \cup \mathcal{L}(\phi^{-1})$.

Lemma 4.9 Every rotationless abelian subgroup A has a basis of generic elements.

Proof Given a basis ψ_1, \dots, ψ_K for A and $\theta \in A$ let $\text{NZ}(\theta) \subset \{1, \dots, N\}$ be the nonzero coordinates of $\Omega(\theta)$. For all but finitely many positive integers a_2 , $\text{NZ}(\psi_1 \psi_2^{a_2}) = \text{NZ}(\psi_1) \cup \text{NZ}(\psi_2)$. Inductively choose positive integers a_i for $i > 1$ so that $\Psi'_1 := \Psi_1 \Psi_2^{a_2} \cdots \Psi_K^{a_K}$ satisfies $\text{NZ}(\Psi'_1) = \bigcup_{i=1}^K \text{NZ}(\psi_i) = \{1, \dots, N\}$. Replacing ψ_1 with Ψ'_1 produces a new basis in which the first element is generic. For all but finitely many positive integers m , $\psi_1, \psi_2 \psi_1^m, \psi_3 \psi_1^m \cdots, \psi_K \psi_1^m$ is a basis of generic elements. \square

Lemma 4.10 If $\phi \in A$ is generic then $\{\text{Fix}(\hat{\Phi}) : \Phi \in P^\pm(\phi)\}$ is the set of maximal (with respect to inclusion) principal sets for A .

Proof Each principal set \mathcal{X}' for A determines a lift $s: A \rightarrow \text{Aut}(F_n)$. If $\Phi \in P^\pm(\phi)$ and $\text{Fix}(\hat{\Phi}) \subset \mathcal{X}'$ then $s(\phi) = \Phi$ and $\mathcal{X}' \subset \text{Fix}(\hat{\Phi})$. This proves that $\text{Fix}(\hat{\Phi})$ is a maximal principal set if it is a principal set. It therefore suffices to show that each $\text{Fix}(\hat{\Phi})$ is a principal set.

If $\mathbb{F} := \text{Fix}(\Phi)$ has rank zero then $\text{Fix}(\hat{\Phi})$ is a principal set by Lemma 3.11(1). If \mathbb{F} has rank one with generator c and with isolated points $P, Q \in \text{Fix}(\hat{\Phi})$ then by Lemma 3.11(2) there is a maximal principal set \mathcal{X}_P that contains P and T_c^\pm and a maximal principal set \mathcal{X}_Q that contains Q and T_c^\pm . Let s_Q and s_P be the lifts of A to $\text{Aut}(F_n)$ determined by \mathcal{X}_P and \mathcal{X}_Q respectively. If $\mathcal{X}_P \neq \mathcal{X}_Q$ then the comparison homomorphism that they determine evaluates to zero on ϕ since $s_P(\phi) = s_Q(\phi) = \Phi$ in contradiction to the assumption that ϕ is generic. Thus $\mathcal{X}_P = \mathcal{X}_Q$. Since P and Q are arbitrary, $\mathcal{X}_P = \text{Fix}(\hat{\Phi})$.

Suppose finally that \mathbb{F} has rank at least two. We claim that $A|\mathbb{F}$ is trivial. If not, let Ω' be the homomorphism defined on $A|\mathbb{F}$ as the product of expansion factor and comparison homomorphisms that occur for $A|\mathbb{F}$. Each coordinate ω' of Ω' extends to a coordinate ω of Ω . Since $\phi|\mathbb{F}$ is the identity, $\omega(\phi) = 0$ in contradiction to the assumption that ϕ is generic. Thus $A|\mathbb{F}$ is trivial and $\partial\mathbb{F}$ is contained in a maximal principal set \mathcal{X} for A .

By Lemma 3.11(3), each isolated point P in $\text{Fix}(\widehat{\Phi})$ is contained in a maximal principal set \mathcal{X}_P whose intersection Y with $\partial\mathbb{F}$ contains at least two points. If $\mathcal{X}_P \neq \mathcal{X}$ then Y has exactly two points and in fact equals $\{T_b^\pm\}$ for some $b \in \mathbb{F}$ since every lift of the identity outer automorphism is an inner automorphism. The comparison homomorphism ω determined by \mathcal{X}_P and \mathcal{X} evaluates to 0 on ϕ in contradiction to the assumption that ϕ is generic. Thus $\mathcal{X}_P = \mathcal{X}$ for all isolated points P and $\text{Fix}(\widehat{\Phi}) = \mathcal{X}$ as desired. \square

It is an immediate corollary, that from the point of view of fixed points of principal lifts, generic elements are indistinguishable.

Corollary 4.11 *For any generic $\phi, \psi \in A$ there is a bijection $h: P^\pm(\phi) \rightarrow P^\pm(\psi)$ such that $\text{Fix}(\widehat{\Phi}) = \text{Fix}(h(\widehat{\Phi}))$ for all $\Phi \in P^\pm(\phi)$.*

5 $\mathcal{A}(\phi)$

The data required in the Recognition Theorem (Theorem 2.10) has both qualitative and quantitative components. If we fix the qualitative part and allow the quantitative part to vary then we generate an abelian group that is naturally associated to the outer automorphism being considered. This section contains a formal treatment of this observation. A more computational friendly approach in terms of relative train track maps is given in the next section.

Definition 5.1 Assume that ϕ is rotationless. $\mathcal{A}(\phi)$ is the subgroup of $\text{Out}(F_n)$ generated by rotationless elements θ for which there is a bijection $h: P^\pm(\phi) \rightarrow P^\pm(\theta)$ satisfying $\text{Fix}(h(\widehat{\Phi})) = \text{Fix}(\widehat{\Phi})$ for all $\widehat{\Phi} \in P(\phi)$.

Remark 5.2 It is an immediate consequence of the definitions that $\mathcal{A}(\phi) = \mathcal{A}(\phi^k)$ for all $k \neq 0$ and for all rotationless ϕ .

Remark 5.3 If A is a rotationless abelian subgroup and ϕ and ψ are generic in A then Corollary 4.11 implies that $\mathcal{A}(\phi) = \mathcal{A}(\psi)$. One can therefore define $\mathcal{A}(A)$ to be $\mathcal{A}(\phi)$ for any generic ϕ in A .

Lemma 5.4 *If A is a rotationless abelian subgroup and ϕ is generic in A , then $A \subset \mathcal{A}(\phi)$.*

Proof Lemma 4.9 and Corollary 4.11 imply that there is a generating set of A that is contained in $\mathcal{A}(\phi)$. \square

To prove that $\mathcal{A}(\phi)$ is abelian we appeal to the following characterization of the rotationless elements in the centralizer $C(\phi)$ of ϕ .

Lemma 5.5 *If $\phi, \psi \in \text{Out}(F_n)$ are rotationless, then $\psi \in C(\phi)$ if and only if the following three properties are satisfied for all $\Phi \in P^\pm(\phi)$:*

- (Φ -1) *There exists $\Psi \in P^\pm(\psi)$ such that $\text{Fix}(\widehat{\Phi})$ is $\widehat{\Psi}$ -invariant.*
- (Φ -2) *If $P \in \text{Fix}(\widehat{\Phi})$ is isolated then one may choose Ψ in (Φ -1) such that $P \in \text{Fix}(\widehat{\Psi})$.*
- (Φ -3) *If $a \in \text{Fix}(\Phi)$ and $[a]_u$ is an axis of ϕ then one may choose Ψ in (Φ -1) such that $a \in \text{Fix}(\Psi)$.*

Moreover, if $\psi \in C(\phi)$ and Ψ is as in (Φ -1) then Ψ commutes with Φ .

Proof If $\psi \in C(\phi)$, let $A = \langle \phi, \psi \rangle$. Lemma 3.11 implies that for each $\Phi \in P^\pm(\phi)$, there is a principal set \mathcal{X} for A whose associated lift $s: A \rightarrow \text{Aut}(F_n)$ satisfies $s(\phi) = \Phi$. Then $s(\psi) \in P^\pm(\psi)$ commutes with Φ and (Φ -1) is satisfied. (Φ -2) follows from Lemma 3.11. If $[a]_u$ is an axis of ϕ then $[a]_u$ is ψ^k -invariant for some $k > 0$ and so is ψ -invariant by Lemma 2.9. Items (2) and (3) of Lemma 3.11 allow us to choose \mathcal{X} to contain T_a^\pm which implies (Φ -3). This completes the only if direction of the lemma.

For the if direction, we assume that ψ satisfies the three items, define $\phi' := \psi\phi\psi^{-1}$ and prove that $\phi' = \phi$ by applying the Recognition Theorem.

For each $\Phi \in P(\phi)$ choose Ψ_1 satisfying (Φ -1) and define $\Phi' = \Psi_1\Phi\Psi_1^{-1} \in P(\phi')$. If Ψ_2 also satisfies (Φ -1) then $\Psi_2 = \Psi_1 i_x$ where $\text{Fix}(\widehat{\Phi})$ is \widehat{i}_x -invariant. By Lemma 2.1, $x \in \text{Fix}(\Phi)$. Thus $\Psi_2\Phi\Psi_2^{-1} = \Psi_1 i_x \Phi i_x^{-1} \Psi_1^{-1} = \Psi_1\Phi\Psi_1^{-1}$ and Φ' is independent of the choice of Ψ_1 . We denote $\Phi \mapsto \Phi'$ by $h: P(\phi) \rightarrow P(\phi')$ and note that $\text{Fix}(\widehat{h(\Phi)}) = \widehat{\Psi}_1(\text{Fix}(\widehat{\Phi})) = \text{Fix}(\widehat{\Phi})$ and that $\text{Fix}_N(\widehat{h(\Phi)}) = \text{Fix}_N(\widehat{\Phi})$. In particular, h is injective. If Φ is replaced by $i_c\Phi i_c^{-1}$ then Ψ_1 can be replaced by $i_c\Psi_1 i_c^{-1}$ and Φ' is replaced by $i_c\Phi' i_c^{-1}$. Thus the restriction of h to an equivalence class in $P(\phi)$ is a bijection onto an equivalence class in $P(\phi')$. Lemma 2.6(2) implies that $P(\phi)$ and $P(\phi')$ have the same number of equivalence classes and hence that h is a bijection.

Suppose that $\Phi_1 \in P(\phi)$, that $a \in \text{Fix}(\Phi_1)$ is root-free and that $\Phi_2 := i_a^d \Phi_1 \in P(\phi)$ for some $d \neq 0$. Then $[a]_u$ is an axis for ϕ and by (Φ_1-3) and (Φ_2-3) we may choose Ψ_1 for Φ_1 and Ψ_2 for Φ_2 to fix a . Thus $\Psi_2 = i_a^m \Psi_1$ for some m and $\Phi'_2 = i_a^m \Psi_1 i_a^d \Phi_1 \Psi_1^{-1} i_a^{-m} = i_a^d \Psi_1 \Phi_1 \Psi_1^{-1} = i_a^d \Phi'_1$ which proves that h satisfies Theorem 2.10(2)(ii).

By Lemma 2.23, for each $\Lambda \in \mathcal{L}(\phi)$ there exists $\Phi \in P(\phi)$ and an isolated point $P \in \text{Fix}_N(\hat{\Phi})$ whose accumulation set equals Λ . By $(\Phi-2)$, we may assume that P is $\hat{\Psi}_1$ -invariant and hence that Λ is ψ -invariant. It follows that Λ is ϕ' -invariant and that $\text{PF}_\Lambda(\phi') = \text{PF}_\Lambda(\phi)$. Theorem 2.10 implies that $\phi = \phi'$ and since $\text{Fix}_N(\hat{\Phi}') = \text{Fix}_N(\hat{\Phi})$, $\Phi = \Phi'$, which proves that Ψ commutes with Φ . \square

We denote the *center* of a group H by $Z(H)$ and define the *weak center* $WZ(H)$ to be the subgroup of H consisting of elements that commute with some iterate of each element of H .

Corollary 5.6 *If $\phi \in \text{Out}(F_n)$ is rotationless then $\mathcal{A}(\phi)$ is an abelian subgroup of $C(\phi)$. Moreover, each element of $\mathcal{A}(\phi)$ commutes with each rotationless element of $C(\phi)$ and so $\mathcal{A}(\phi) \subset WZ(C(\phi))$.*

Proof Lemma 5.5 implies that $\theta \in C(\phi)$ for each θ in the defining generating set of $\mathcal{A}(\phi)$ and that $C(\phi)$ and $C(\theta)$ contains the same rotationless elements. The corollary follows. \square

Remark 5.7 In general, $\mathcal{A}(\phi)$ is not contained in the center of $C(\phi)$. For example, if $n = 2k$ and $\Phi \in P^\pm(\phi)$ commutes with an order two automorphism Θ that interchanges the free factor generated by the first k elements in a basis with the free factor generated by the last k elements of that basis, then $\mathcal{A}(\phi)$ will contain elements that do not commute with θ .

It is natural to ask if ϕ is generic in $\mathcal{A}(\phi)$.

Lemma 5.8 *If ϕ is rotationless then ϕ is generic in $\mathcal{A}(\phi)$.*

Proof We must show that if ω is a coordinate of $\Omega: \mathcal{A}(\phi) \rightarrow \mathbb{Z}^N$ then $\omega(\phi) \neq 0$. Choose an element θ of the defining generating set for $\mathcal{A}(\phi)$ such that $\omega(\theta) \neq 0$. If $\omega = \text{PF}_\Lambda$ then, after replacing θ with θ^{-1} if necessary, $\Lambda \in \mathcal{L}(\theta)$. By Remark 2.18 and Lemma 2.23, there exist $\Theta \in P(\theta)$ and an isolated point $P \in \text{Fix}_N(\Theta)$ whose accumulation set is Λ . After replacing ϕ with ϕ^{-1} if necessary, there exists $\Phi \in P(\phi)$

such that $\text{Fix}(\widehat{\Phi}) = \text{Fix}(\widehat{\Theta})$ and such that P is an isolated point in $\text{Fix}_N(\Psi)$. Lemma 2.24 implies that $\omega(\phi) \neq 0$.

If ω is a comparison homomorphism determined by lifts $s, t: \mathcal{A}(\phi) \rightarrow \text{Aut}(F_n)$ then $s(\theta) \neq t(\theta)$. Thus $\text{Fix}(s(\widehat{\phi})) = \text{Fix}(s(\widehat{\theta})) \neq \text{Fix}(t(\widehat{\theta})) = \text{Fix}(t(\widehat{\phi}))$ which implies that $\omega(\phi) \neq 0$. \square

The following characterization of $\mathcal{A}(\phi)$ is an immediate corollary of Lemma 5.4 and Lemma 5.8.

Lemma 5.9 *Suppose that ϕ is rotationless. Then $\mathcal{A}(\phi)$ is the maximal rotationless abelian subgroup in which ϕ is generic.*

6 Disintegrating ϕ

We have reduced the study of rotationless abelian subgroups of $\text{Out}(F_n)$, and so of abelian subgroups of $\text{Out}(F_n)$ up to finite index, to the study of $\mathcal{A}(\phi)$ for rotationless $\phi \in \text{Out}(F_n)$. In this section we construct the subgroup $\mathcal{D}(\phi)$ of $\mathcal{A}(\phi)$ described in the introduction. In Section 7 we show that $\mathcal{D}(\phi)$ has finite index in $\mathcal{A}(\phi)$.

Let $f: G \rightarrow G$ be a CT representing ϕ . We will need a coarsening of the complete splitting of a path. If $\{E_i\}$ is the set of linear edges associated to an axis μ for ϕ then by (Linear edges) there is a root-free closed Nielsen path w and there are distinct nonzero integers d_i such that $f(E_i) = E_i \cdot w^{d_i}$; we say that d_i is the *exponent of E_i* . For distinct E_i and E_j and for $l \in \mathbb{Z}$, the path $E_i w^l \bar{E}_j$ is said to be *quasi-exceptional*. The paths obtained by varying l but keeping i and j fixed are said to *belong to the same quasi-exceptional family*. When l is unimportant we write $E_i w^* \bar{E}_j$. If d_i and d_j have the same sign then $E_i w^* \bar{E}_j$ is an exceptional path but otherwise it is not. Note also that since E_i and E_j are distinct, no Nielsen path is quasi-exceptional.

Assume that $\sigma = \sigma_1 \cdots \sigma_s$ is the (necessarily unique) complete splitting of a path σ . If $a \leq b$ and $\sigma_{ab} := \sigma_a \cdots \sigma_b$ is quasi-exceptional then we say that σ_{ab} is a *QE-subpath* of σ .

Lemma 6.1 *For any completely split path σ , distinct QE-subpaths of σ have disjoint interiors.*

Proof Suppose that $\sigma = \sigma_1 \cdots \sigma_s$ is the complete splitting of σ and that there exist $1 \leq a \leq b \leq s$ and $1 \leq a \leq c \leq d \leq s$ such that $\sigma_{ab} := \sigma_a \cdots \sigma_b$ and $\sigma_{cd} := \sigma_c \cdots \sigma_d$ are distinct quasi-exceptional paths. We must show that $c > b$.

If $a < b$ then $\sigma_a = E_i$ is a linear edge, $\sigma_b = \bar{E}_j$ is the inverse of a linear edge and each σ_l , $a < l < b$, is a Nielsen path. None of these terms is quasi-exceptional so we may assume that $c < d$. The initial edge σ_c of σ_{cd} is a linear edge so is either equal to a or greater than b . In the latter case we are done. In the former case, the terminal edge of σ_{cd} must be σ_b since it is the first term after σ_c in the complete splitting of σ that is not a Nielsen path. This contradicts the assumption that $\sigma_{ab} \neq \sigma_{cd}$ and so completes the proof if $a < b$. The case that $c < d$ is proved similarly and the case that both $a = b$ and $c = d$ is obvious. \square

Definition 6.2 The *QE-splitting* of a completely split path σ is the coarsening of the complete splitting of σ obtained by declaring each QE-subpath to be a single element. Thus the QE-splitting is a splitting into single edges in irreducible strata, connecting subpaths in zero strata, Nielsen paths and quasi-exceptional paths. These subpaths are the *terms of the QE-splitting*.

Definition 6.3 Define a finite directed graph B as follows. There is one vertex v_i^B for each nonfixed irreducible stratum H_i . If H_i is NEG then a v_i^B -path is the unique edge in H_i ; if H_i is EG then a v_i^B -path is either an edge in H_i or a taken connecting path in a zero stratum contained in H_i^z . There is a directed edge from v_i^B to v_j^B if there is a v_i^B -path κ_i such that some term in the QE-splitting of $f_\#(\kappa_i)$ is an edge in H_j . (Note that edges in B are defined with regard to $f_\#$ rather than an iterate of $f_\#$.) The components of B are labelled B_1, \dots, B_M . For each B_s , define X_s to be the minimal subgraph of G that contains H_i if $v_i^B \in B_s$ and H_i is NEG and contains H_i^z if $v_i^B \in B_s$ and H_i is EG. We say that X_1, \dots, X_M are the *almost invariant subgraphs* associated to $f: G \rightarrow G$.

Remark 6.4 If a vertex v belongs to distinct almost invariant subgraphs then v is a principal vertex by [10, Remark 4.9] and is hence fixed by f .

We could construct a directed graph with the same vertices as B by having a directed edge from v_i^B to v_j^B if there is a v_i^B -path κ_i such that some term in the QE-splitting of $f_\#(\kappa_i)$ is a v_j^B -path. The following lemma shows that this produces the same graph B .

Lemma 6.5 *If $i \neq j$ and there is a v_i^B -path κ_i such that some term in the QE-splitting of $f_\#(\kappa_i)$ is a v_j^B -path κ_j then some term in the QE-splitting of $f_\#(\kappa_i)$ is an edge in H_j ; in particular there is a directed edge in B from v_i^B to v_j^B .*

Proof We may assume without loss that H_j is EG and that some term σ_k in the QE-splitting of $f_\#(\kappa_i)$ is a connecting path in some zero stratum in H_j^z . If H_i is NEG

then the endpoints of κ_i , and hence the endpoints of $f_{\#}(\kappa_i)$, are contained in $\text{Fix}(f)$ by Remark 2.18 and (Vertices). If H_i is EG then the endpoints of $f_{\#}(\kappa_i)$ belong to both H_i and H_j and so belong to $\text{Fix}(f)$ by Remark 6.4. Since the endpoints of σ_k are not fixed by f , σ_k is neither the first nor last term in the QE-splitting of $f_{\#}(\kappa_i)$. The terms adjacent to σ_k in the QE-splitting of $f_{\#}(\kappa_i)$ must be edges in H_j . \square

Definition 6.6 For each M -tuple \mathbf{a} of nonnegative integers, define $f_{\mathbf{a}}: G \rightarrow G$ by

$$f_{\mathbf{a}}(E) = \begin{cases} f_{\#}^{a_i}(E) & \text{for each edge } E \subset X_i, \\ E & \text{for each edge } E \text{ that is fixed by } f. \end{cases}$$

Lemma 6.7 $f_{\mathbf{a}}: G \rightarrow G$ is a homotopy equivalence for all \mathbf{a} .

Proof Let NI be the number of irreducible strata in the filtration and for each $0 \leq m \leq \text{NI}$, let $G_{i(m)}$ be the smallest filtration element containing the first m irreducible strata. We will prove by induction that each $f_{\mathbf{a}}|_{G_{i(m)}}$ is a homotopy equivalence.

Since H_1 is never a zero stratum, $i(1) = 1$. If G_1 is not a single edge fixed by f , then every edge in G_1 is contained in a single almost invariant subgraph X_i . Thus $f_{\mathbf{a}}|_{G_1}$ is either the identity or is homotopic to $f^{a_i}|_{G_1}$; in either case it is a homotopy equivalence.

We assume now that $f_{\mathbf{a}}|_{G_{i(m)}}$ is a homotopy equivalence. Define $g_1: G_{i(m+1)} \rightarrow G_{i(m+1)}$ on edges by

$$g_1(E) = \begin{cases} f_{\mathbf{a}}(E) & \text{if } E \subset G_{i(m)}, \\ E & \text{if } E \subset G_{i(m+1)} \setminus G_{i(m)}. \end{cases}$$

Remark 6.4 guarantees that g_1 is well defined. It is easy to check that g_1 is a homotopy equivalence. If the edges of $H_{i(m+1)}$ are fixed by f , then $g_1 = f_{\mathbf{a}}|_{G_{i(m+1)}}$ and we are done.

If $f|_{H_{i(m+1)}}$ is not the identity, then the edges in $G_{i(m+1)} \setminus G_{i(m)}$ are contained in a single almost invariant subgraph, say X_k . Define $g_2: G_{i(m+1)} \rightarrow G_{i(m+1)}$ on edges by

$$g_2(E) = \begin{cases} f_{\#}^{a_k}(E) & \text{if } E \subset G_{i(m)}, \\ E & \text{if } E \subset G_{i(m+1)} \setminus G_{i(m)}, \end{cases}$$

and $g_3: G_{i(m+1)} \rightarrow G_{i(m+1)}$ on edges by

$$g_3(E) = \begin{cases} E & \text{if } E \subset G_{i(m)}, \\ f_{\#}^{a_k}(E) & \text{if } E \subset G_{i(m+1)} \setminus G_{i(m)}. \end{cases}$$

Then g_2 is a homotopy equivalence and $f^{a_k}|_{G_{i(m+1)}} = g_3 g_2$. Each component of $G_{i(m+1)}$ is noncontractible by [10, Lemma 4.16], so $f^{a_k}|_{G_{i(m+1)}}$ is a homotopy

equivalence. It follows that g_3 , and hence also $f_{\mathbf{a}}|_{G_{i(m+1)}} = g_3 g_1$ is a homotopy equivalence. \square

Almost invariant subgraphs are defined without reference to the quasi-exceptional paths in the QE-splitting of edge images. The next definition brings these into the discussion.

Definition 6.8 Suppose $\{X_1, \dots, X_M\}$ are the almost invariant subgraphs of $f: G \rightarrow G$. An M -tuple $\mathbf{a} = (a_1, \dots, a_M)$ of nonnegative integers is *admissible* if for all axes μ , whenever

- X_s contains a linear edge E_i associated to μ with exponent d_i ,
- X_t contains a linear edge E_j associated to μ with exponent d_j ,
- there exists a vertex v^B of B and a v^B -path $\kappa \subset X_r$ such that some element in the quasi-exceptional family determined by $E_i \bar{E}_j$ is a term in the QE-splitting of $f_{\#}(\kappa)$,

then $a_r(d_i - d_j) = a_s d_i - a_t d_j$.

Example 6.9 Suppose that G is the rose with edges E_1, E_2, E_3 and E_4 and that $f: G \rightarrow G$ is defined by $E_1 \mapsto E_1$, $E_2 \mapsto E_2 E_1^2$, $E_3 \mapsto E_3 E_1$ and $E_4 \mapsto E_4 E_3 E_3 \bar{E}_2$. Then $M = 2$ with X_1 having the single edge E_2 and X_2 consisting of E_3 and E_4 . In the notation of Definition 6.8, $s = 1$, $i = 2$, $d_i = 2$, $t = 2$, $j = 3$, $d_j = 1$, $r = 2$ and $\kappa = E_4$. The pair (a_1, a_2) is admissible if $a_2 = 2a_1 - a_2$ or equivalently $a_2 = a_1$. Thus $f_{\mathbf{a}} = f^{a_1}$ for each admissible \mathbf{a} .

Definition 6.10 Each $f_{\mathbf{a}}$ determines an element $\phi_{\mathbf{a}} \in \text{Out}(F_n)$ and also an element $[f_{\mathbf{a}}]$ in the semigroup of homotopy equivalences of G that respect the filtration modulo homotopy relative to the set of vertices of G . Define $\mathcal{D}(\phi) = \langle \phi_{\mathbf{a}} : \mathbf{a} \text{ is admissible} \rangle$. Both $\phi_{\mathbf{a}}$ and $\mathcal{D}(\phi)$ depend on the choice of $f: G \rightarrow G$; see Example 6.11 below. Since we work with a single $f: G \rightarrow G$ throughout the paper and since $\mathcal{D}(\phi)$ is well defined up to finite index by Theorem 7.1, we suppress this dependence in the notation.

Example 6.11 Let G be the rose with edges E_1, E_2 and E_3 . Subdivide E_3 into $E_3 = \bar{D}_1 D_2$. Define $f_1: G \rightarrow G$ by

$$E_1 \mapsto E_1 \quad E_2 \mapsto E_1 E_2 \quad D_1 \mapsto D_1 \bar{E}_1^2 \quad D_2 \mapsto D_2 E_1$$

and $f_2: G \rightarrow G$ by

$$E_1 \mapsto E_1 \quad E_2 \mapsto E_2 E_1 \quad D_1 \mapsto D_1 \bar{E}_1 \quad D_2 \mapsto D_2 E_1^2.$$

The automorphisms of F_n determined by f_1 and f_2 differ by i_{E_1} and so determine the same element $\phi \in \text{Out}(F_n)$. The homotopy equivalence of G that fixes E_1 , D_1 and D_2 and satisfies $E_2 \mapsto E_2 E_1$ represents an element $\eta \in \text{Out}(F_n)$ that is contained in $\mathcal{D}(\phi)$ if f_2 is used but not if f_1 is used. Note that η^2 is contained in $\mathcal{D}(\phi)$ if either f_2 or f_1 is used.

Notation 6.12 The set of Nielsen paths for f (resp. $f_{\mathbf{a}}$) with endpoints at vertices is denoted $\mathcal{N}(f)$ (resp. $\mathcal{N}(f_{\mathbf{a}})$). For each $1 \leq s \leq M$, let \mathcal{K}_s be the set of v^B -paths for $v^B \in B_s$. Equivalently, \mathcal{K}_s consists of all edges in irreducible nonfixed strata in X_s and all taken connecting paths in zero strata in X_s . Let \mathcal{Q}_s be the set of quasi-exceptional subpaths for f that belong to the same quasi-exceptional family as a quasi-exceptional subpath in the QE-splitting of $f_{\#}(\kappa)$ for some $\kappa \in \mathcal{K}_s$. Finally, let \mathcal{P}_s be the set of paths that have complete splittings with respect to f each of whose terms is an element of $\mathcal{N}(f)$, \mathcal{Q}_s or \mathcal{K}_s .

Lemma 6.13 *The following hold for all admissible \mathbf{a} .*

- (1) *If $\sigma \in \mathcal{N}(f)$ then $\sigma \in \mathcal{N}(f_{\mathbf{a}})$.*
- (2) *If $\sigma \in \mathcal{Q}_s$ then $(f_{\mathbf{a}})_{\#}(\sigma) = f_{\#}^{a_s}(\sigma)$. In particular, \mathcal{Q}_s is $(f_{\mathbf{a}})_{\#}$ -invariant.*

Proof Our proof is by induction on the height r of σ . In the context of (1), we may assume that σ is indivisible.

G_1 is either a single fixed edge or is contained in a single almost invariant subgraph. Thus $f_{\mathbf{a}}|_{G_1}$ is either the identity or an iterate of $f|_{G_1}$. In either case (1) is obvious for $\sigma \subset G_1$. Since G_1 does not contain any quasi-exceptional paths, the lemma holds for $\sigma \subset G_1$. We assume now that $r \geq 2$, that the lemma holds for paths in G_{r-1} and that σ has height r and is either an element of $\mathcal{N}(f)$ or an element of \mathcal{Q}_s . Since f satisfies (NEG Nielsen paths), H_r is either EG or linear.

Let X_u be the almost invariant subgraph containing H_r . Suppose at first that H_r is linear and is hence a single edge E_r such that $f(E_r) = E_r w^{d_r}$ for some nontrivial root-free Nielsen path w and $d_r \neq 0$. If $\sigma \in \mathcal{N}(f)$, then $\sigma = E_r w^p \bar{E}_r$ for some integer p . By the inductive hypothesis, $(f_{\mathbf{a}})_{\#}(w) = w$ so

$$(f_{\mathbf{a}})_{\#}(\sigma) = [(E_r w^{a_u d_r}) w^p (\bar{w}^{a_u d_r} \bar{E}_r)] = E_r w^p \bar{E}_r = \sigma.$$

If $\sigma \in \mathcal{Q}_s$, then up to a reversal of orientation, $\sigma = E_r w^p \bar{E}_j$ where $f(E_j) = E_j w^{d_j}$. Let X_t be the almost invariant subgraph containing E_j . Since \mathbf{a} is admissible, $a_s(d_r - d_j) = a_u d_r - a_t d_j$.

Thus

$$\begin{aligned}
 (f_{\mathbf{a}})_{\#}(\sigma) &= [f_{\#}^{a_u}(E_r)(f_{\mathbf{a}}(w))^p f_{\#}^{a_t}(\bar{E}_j)] \\
 &= [E_r w^{a_u d_r} w^p \bar{w}^{a_t d_j} \bar{E}_j] \\
 &= [E_r w^{a_u d_r - a_t d_j + p} \bar{E}_j] \\
 &= [E_r w^{a_s(d_r - d_j) + p} E_j] \\
 &= [E_r w^{a_s d_r} w^p \bar{w}^{a_s d_j} \bar{E}_j] \\
 &= [f_{\#}^{a_s}(E_r)(f_{\#}^{a_s}(w))^p f_{\#}^{a_s}(\bar{E}_j)] \\
 &= f_{\#}^{a_s}(\sigma).
 \end{aligned}$$

Suppose now that H_r is EG. There are no quasi-exceptional paths of height r so σ is an indivisible Nielsen path of height r . By [10, Lemma 4.25], σ decomposes as a concatenation of edges in H_r and Nielsen paths in G_{r-1} . By definition and by the inductive hypothesis, $(f_{\mathbf{a}})_{\#}$ equals $f_{\#}^{a_s}$ on all terms in this decomposition and hence on σ . \square

The next two corollaries are immediate consequences of Lemma 6.13, the definition of X_s and the definition of $f_{\mathbf{a}}$.

Corollary 6.14 For $1 \leq i \leq M$, \mathcal{P}_s is preserved by both $f_{\#}$ and $(f_{\mathbf{a}})_{\#}$ and moreover $(f_{\mathbf{a}})_{\#}(\sigma) = f_{\#}^{a_s}(\sigma)$ for all $\sigma \in \mathcal{P}_s$. Thus $(f_{\mathbf{a}}^k)_{\#}(\sigma) = f_{\#}^{ka_s}(\sigma)$ for all $\sigma \in \mathcal{P}_s$ and all $k \geq 1$.

Corollary 6.15 Suppose that E_i and E_j are linear edges with the same axis and that w, d_i and d_j are as in (Linear edges). Suppose further that $E_i \subset X_s$ and $E_j \subset X_t$. Then $(f_{\mathbf{a}})_{\#}(E_r w^p \bar{E}_j) = E_r w^{a_s + p - a_t} \bar{E}_j$.

Corollary 6.16 For each admissible \mathbf{a} and \mathbf{b} , $[f_{\mathbf{a}}][f_{\mathbf{b}}] = [f_{\mathbf{b}}][f_{\mathbf{a}}] = [f_{\mathbf{a}+\mathbf{b}}]$. In particular, $\mathcal{D}(\phi)$ is abelian.

Proof Let $\gamma_1, \dots, \gamma_n$ be closed paths based at a vertex $v \in G$ that represent a basis for $\pi_1(G, v)$. Choose K so large that $\beta_i = f_{\#}^K(\gamma_i)$ is completely split for all i . After increasing K if necessary we may also assume that each connecting path in a zero stratum that is a term in the complete splitting of β_i is an element of \mathcal{K}_s for some s . Thus β_1, \dots, β_n represent a basis for $\pi_1(G, v)$ and each term in the QE-splitting of β_i is either an element of some \mathcal{P}_s or a quasi-exceptional path (remember that not every quasi-exceptional path belongs to some \mathcal{Q}_s). It suffices to show that $(f_{\mathbf{a}+\mathbf{b}})_{\#} = (f_{\mathbf{a}})_{\#}(f_{\mathbf{b}})_{\#}$ on each such term and by Corollary 6.14 we are reduced to showing that $(f_{\mathbf{a}+\mathbf{b}})_{\#}(\sigma) = (f_{\mathbf{a}})_{\#}(f_{\mathbf{b}})_{\#}(\sigma)$ for each quasi-exceptional path σ .

Let $\sigma = E_r w^p \bar{E}_j$ where $w \in \mathcal{N}(f)$, $f(E_r) = E_r^{d_r}$ and $f(E_j) = E_r^{d_j}$. Let X_u be the almost invariant subgraph containing H_r and let X_t be the almost invariant subgraph containing H_j . Then

$$\begin{aligned} (f_{\mathbf{a}+\mathbf{b}})_{\#}(E_r w^p \bar{E}_j) &= f_{\#}^{a_u+b_u}(E_r) w^p f_{\#}^{a_t+b_t}(\bar{E}_j) \\ &= [E_r w^{a_u d_r + b_u d_r} w^p \bar{w}^{a_t d_r + b_t d_r} \bar{E}_j] \\ &= [E_r w^{a_u d_r} w^{b_u d_r} w^p \bar{w}^{b_t d_r} \bar{w}^{a_t d_r} \bar{E}_j] \\ &= [f_{\#}^{a_u}(E_r) w^{b_u d_r + p - b_t d_r} f_{\#}^{a_t}(\bar{E}_j)] \\ &= (f_{\mathbf{a}})_{\#}([E_r w^{b_u d_r + p - b_t d_r} \bar{E}_j]) \\ &= (f_{\mathbf{a}})_{\#}(f_{\mathbf{b}})_{\#}(E_r w^p \bar{E}_j). \end{aligned} \quad \square$$

Definition 6.17 An admissible \mathbf{a} is *generic* if each $a_i > 0$ and if whenever $E_i \in X_r$ and $E_j \in X_s$ are distinct linear edges associated to the same axis, then $a_r d_i \neq a_s d_j$ where d_i and d_j are the exponents of E_i and E_j respectively.

Lemma 6.18 If \mathbf{a} is generic then $f_{\mathbf{a}}: G \rightarrow G$ is a CT and has the same principal vertices and Nielsen paths as f . In particular, $f_{\mathbf{a}}$ is rotationless.

Proof We first note (justification below) that:

- (1) $f_{\mathbf{a}}$ has the same periodic edges and the same periodic directions at vertices as f .
- (2) $f_{\mathbf{a}}$ preserves the filtration $\emptyset = G_0 \subset G_1 \subset \dots \subset G_N = G$ and each stratum H_i has the same type (zero, EG, NEG, linear) for $f_{\mathbf{a}}$ as it does for f .
- (3) $f_{\mathbf{a}}: G \rightarrow G$ is a relative train track map; ie the following hold for each EG stratum H_r .
 - (a) $Df_{\mathbf{a}}$ preserves the set of directions that are based at vertices and determined by edges in H_r .
 - (b) A path $\sigma \subset G_r$ is r -legal for f if and only if it is r -legal for $f_{\mathbf{a}}$.
 - (c) If σ is a connecting path for H_r then $(f_{\mathbf{a}})_{\#}(\sigma)$ is nontrivial.

Properties (1), (2) and (3)(a) follow from Corollary 6.14 and the assumption that each $a_s > 0$. Suppose that σ is a connecting path σ for $H_r \subset X_s$. If σ is contained in a zero stratum, then it decomposes as a concatenation (not necessarily a splitting) of edges in \mathcal{P}_s by (Zero strata) and so $(f_{\mathbf{a}})_{\#}(\sigma) = f_{\#}^{a_s}(\sigma)$ by Corollary 6.14. The nontriviality of $(f_{\mathbf{a}})_{\#}(\sigma)$ therefore follows from the fact that (3)(c) holds for f . In the remaining case σ is contained in a noncontractible component of G_{r-1} and (3)(c) is equivalent to the endpoints of σ being fixed points by [10, Remark 2.8]. Since $\text{Fix}(f) \subset \text{Fix}(f_{\mathbf{a}})$,

(3)(c) for $f_{\mathbf{a}}$ follows from (3)(c) for f . Now that (3)(c) is verified, (3)(b) follows from Corollary 6.14.

If $E_j \in X_S$ is a linear edge for f and $f(E_j) = E_j w_j^{d_j}$ then $f_{\mathbf{a}}(E_j) = E_j w_j^{a_s d_j}$. From this and the fact that \mathbf{a} is generic, it follows that $f_{\mathbf{a}}$ satisfies (Linear edges).

Properties (Completely split), (Filtration), (Periodic edges) and (Zero strata) for $f_{\mathbf{a}}$ follow from items (1) and (2), Corollary 6.14 and the corresponding property for f .

By Lemma 6.13, every Nielsen path for f is a Nielsen path for $f_{\mathbf{a}}$. We prove the converse below. Assuming for now that f and $f_{\mathbf{a}}$ have the same Nielsen paths, we complete the proof of the lemma. Property (NEG Nielsen paths) for $f_{\mathbf{a}}$ follows from the corresponding property for f as do (Rotationless) and (Vertices) for $f_{\mathbf{a}}$ by applying item (1).

If H_r is EG and ρ is an indivisible Nielsen path of height r , then ρ splits into a concatenation of edges in H_r and Nielsen paths in G_{r-1} by [10, Lemma 4.25]. It follows that the extended fold determined by ρ is the same with respect to $f_{\mathbf{a}}$ as it is with respect to f and that this remains true as one iteratively folds ρ . Property (EG Nielsen paths) for $f_{\mathbf{a}}$ follows from the corresponding property for f and [10, Corollary 4.34] which states that (EG Nielsen paths) holds if and only if the illegal turn at each indivisible Nielsen path obtained by iteratively folds ρ is proper.

It remains to assume that ρ is an indivisible Nielsen path for $f_{\mathbf{a}}$ and prove that it is a Nielsen path for f . If an endpoint of ρ is not a vertex then it is contained in an EG stratum. Subdividing at this point and declaring both new edges to be in the same stratum as the original edge preserves all the properties of a CT. We may therefore assume that the endpoints of ρ are vertices. Once we have established that ρ is a Nielsen path for f it will follow that this subdivision was unnecessary.

Let i be the height of ρ and let X_r be the almost invariant subgraph that contains H_i . We consider first the case that H_i , which is necessarily irreducible, is EG. By Lemma 5.11 of [4], $\rho = \alpha\beta$ where α and β are i -legal paths for $f_{\mathbf{a}}$ that begin and end in H_i . Let E_{α} be the initial edge of α . By Corollary 6.14, there exists $k \geq 1$ so that $f_{\#}^k(E_{\alpha})$ contains α . Since both E_{α} and the terminal edge of α are edges of height i , the QE-splitting of $f_{\#}^k(E_{\alpha})$ restricts to a QE-splitting of α . Corollary 6.14 therefore implies that $(f_{\mathbf{a}})_{\#}(\alpha) = f_{\#}^{a_r}(\alpha)$. The analogous argument applies to β and we conclude that ρ is a Nielsen path for f^{a_r} . By [10, Lemma 4.14], every periodic Nielsen path for f has period one. In particular, ρ is a Nielsen path for f .

Suppose next that H_i is a single NEG edge E_i . After reversing the orientation on ρ if necessary, we may assume by Lemma 4.1.4 of [4] applied to $f_{\mathbf{a}}$ that E_i is the initial edge of ρ and that E_i is not fixed by $f_{\mathbf{a}}$ and hence not fixed by f . Choose lifts $\tilde{\rho} \subset \Gamma$

and $\tilde{f}_a: \Gamma \rightarrow \Gamma$ such that \tilde{f}_a fixes the endpoints of $\tilde{\rho}$. Let $\tilde{f}: \Gamma \rightarrow \Gamma$ be the lift of f that fixes the initial endpoint and direction of $\tilde{\rho}$. By Lemma 2.23, there is a ray \tilde{R}_1 with the same initial vertex and direction as $\tilde{\rho}_1$ and satisfying the following properties.

- (i) $\text{Fix}(\tilde{f}) \cap \tilde{R}_1$ is the initial endpoint of \tilde{R}_1 .
- (ii) If $\tilde{R}_1 = \tilde{\tau}_1 \cdot \tilde{\tau}_2 \cdots$ is the QE-splitting of \tilde{R}_1 and if \tilde{x}_l is the terminal endpoint of $\tilde{\tau}_l$ then $\tilde{f}(\tilde{x}_l) = \tilde{x}_k$ for some $k > l$ and Df maps the turn taken by \tilde{R}_1 at \tilde{x}_l to the turn taken by \tilde{R}_1 at \tilde{x}_k .

Corollary 6.14 implies that $(\tilde{f}_a)_\# \tilde{R}_1 = f_\#^{ar} \tilde{R}_1$ and hence that (ii) holds with \tilde{f} replaced by \tilde{f}_a . If (i) fails with \tilde{f} replaced by \tilde{f}_a then there is a fixed point for \tilde{f}_a in the interior of some $\tilde{\tau}_l$ and so by (ii) for \tilde{f}_a there exists an initial subpath $\tilde{\mu}$ of $\tilde{\tau}_l$ such that $(\tilde{f}_a)_\#(\tilde{\mu})$ is trivial. But no such $\tilde{\mu}$ can exist. This follows from Corollary 6.14 if τ_l is a single edge and is easy to check by inspection (see [10, Lemma 4.12]) if τ_l is either a quasi-exceptional path or a Nielsen path for f . Since τ_l contains a fixed point for \tilde{f}_a it is not a connecting subpath in a zero stratum. This completes the proof that (i) and (ii) hold with \tilde{f} replaced by \tilde{f}_a . In particular, \tilde{R}_1 does not contain the terminal endpoint of $\tilde{\rho}$.

Let $P_1 \in \partial\Gamma$ be the terminal endpoint of \tilde{R}_1 , let $\tilde{\rho}_1$ be the common initial segment of \tilde{R}_1 and $\tilde{\rho}$ and let $\tilde{\rho}_2$ be the terminal segment of $\tilde{\rho}$ such that $\tilde{\rho} = \tilde{\rho}_1 \tilde{\rho}_2$. By an argument exactly analogous to the one in the previous paragraph, \tilde{f}_a moves the terminal endpoint \tilde{w} of $\tilde{\rho}_1$ toward P_1 ; more precisely, the ray from $\tilde{f}_a(\tilde{w})$ to P_1 does not contain \tilde{w} . Since \tilde{w} is the initial endpoint of $\tilde{\rho}_2$ and since the interior of $\tilde{\rho}_2$ is disjoint from $\text{Fix}(\tilde{f}_a)$, Lemma 3.16 of [10] (see also Section 2 of [4]) implies that \tilde{f}_a moves each point in the interior of $\tilde{\rho}_2$ toward P_1 . In particular, the initial direction of ρ^{-1} is fixed by f_a .

Let \tilde{E}_j be the initial edge of $\tilde{\rho}_2^{-1}$. There is a ray \tilde{R}_2 with initial edge \tilde{E}_j that satisfies (i) and (ii) with \tilde{R}_1 replaced by \tilde{R}_2 and \tilde{f} replaced by \tilde{f}_a . If \tilde{E}_j is NEG, or more generally if the initial endpoint of E_j is principal, then the existence of \tilde{R}_2 follows from Lemma 2.23 as above. If \tilde{E}_j is contained in an EG stratum but the initial endpoint of E_j is not principal then Lemma 2.23 does not apply. In this case we define \tilde{R}_2 to be the increasing union $\tilde{E}_j \subset \tilde{f}_a(\tilde{E}_j) \subset (\tilde{f}_a^2)_\#(\tilde{E}_j) \subset (\tilde{f}_a^3)_\#(\tilde{E}_j) \subset \cdots$. It is shown in (the proof of) [10, Lemma 2.13] that (i) and (ii) with \tilde{R}_1 replaced by \tilde{R}_2 and \tilde{f} replaced by \tilde{f}_a are satisfied,

Let $P_2 \in \partial\Gamma$ be the terminal endpoint of \tilde{R}_2 . If $P_1 \neq P_2$, let \tilde{L}_{12} be the line connecting P_1 to P_2 . Then \tilde{L}_{12} is contained in $\tilde{R}_1 \cup \tilde{\rho} \cup \tilde{R}_2$ and does not contain the endpoints of $\tilde{\rho}$, which are also the endpoints of \tilde{R}_1 and \tilde{R}_2 . It follows that $\tilde{L}_{12} \cap \text{Fix}(\tilde{f}_a) = \emptyset$

which contradicts the fact that there are points arbitrarily close to P_i that are moved toward P_i by \tilde{f}_a . We conclude that $P_1 = P_2$.

The lift of f^{ar} that fixes the initial endpoint of $\tilde{\rho}$ and the lift of f^{ar} that fixes the terminal endpoint of $\tilde{\rho}$ both fix $P_1 = P_2$. If these lifts are the same then $\tilde{\rho}$ is a Nielsen path for \tilde{f}^r , ρ is a Nielsen path for f^r and by [10, Lemma 4.14], ρ is a Nielsen path for f . We may therefore assume that these lifts are distinct, in which case $P_1 = P_2 = T_b^\pm$ for some $b \in F_n$ and E_i and E_j are distinct linear edges for f associated to the axis $[b]_u$ (with associated root-free w) and $\rho = E_i w^p \bar{E}_j$. But this contradicts Corollary 6.15 and the assumption that ρ is a Nielsen path for f_a . This completes the proof that f and f_a have the same Nielsen paths and so the proof of the lemma. \square

We now relate $\mathcal{D}(\phi)$ to $\mathcal{A}(\phi)$, using the correspondence between principal lifts of relative train track maps and principal automorphisms.

Corollary 6.19 *For each generic \mathbf{a} there is a bijection $h: P(\phi) \rightarrow P(\phi_a)$ such that $\text{Fix}_N(h(\hat{\Phi})) = \text{Fix}_N(\hat{\Phi})$ for all $\hat{\Phi} \in P(\phi)$. If \tilde{f} corresponds to $\hat{\Phi}$ and \tilde{f}_a corresponds to $h(\hat{\Phi})$ then $\text{Fix}(\tilde{f}) = \text{Fix}(\tilde{f}_a)$.*

Proof By Lemma 6.18, f and f_a have the same Nielsen classes of principal vertices. There is an induced bijection h between principal lifts of f_a and principal lifts of f ; if $\tilde{f}_a = h(\tilde{f})$ then $\text{Fix}(\tilde{f}) = \text{Fix}(\tilde{f}_a)$. Lemma 2.2 implies that $\text{Fix}_N(\tilde{f})$ and $\text{Fix}_N(\tilde{f}_a)$ have the same nonisolated points. Lemma 2.23 and Corollary 6.14 imply that $\text{Fix}_N(\tilde{f})$ and $\text{Fix}_N(\tilde{f}_a)$ have the same isolated points. \square

Corollary 6.20 *$\mathcal{D}(\phi)$ is contained in $\mathcal{A}(\phi)$ and is generated by elements of the form ϕ_a with \mathbf{a} generic.*

Proof Let $S' = \{\phi_{\mathbf{b}} : \mathbf{b} \in \mathcal{B}\}$ be any generating set for $\mathcal{D}(\phi)$. If \mathbf{I} is the M -tuple with 1's in each coordinate then \mathbf{I} is generic and $\phi_{\mathbf{I}} = \phi$ is represented by $f_{\mathbf{I}} = f$. There exists $k > 0$ so that $\mathbf{a} = \mathbf{b} + k\mathbf{I}$ is generic (because it is projectively close to \mathbf{I}) for each $\mathbf{b} \in \mathcal{B}$. Corollary 6.16 implies that $\phi_a = \phi^k \phi_{\mathbf{b}}$ and Corollary 6.19 that $\phi^k \phi_{\mathbf{b}} \in \mathcal{A}(\phi)$. Thus $S = \{\phi, \phi^k \phi_{\mathbf{b}} : \mathbf{b} \in \mathcal{B}\} \subset \mathcal{A}(\phi)$ is a generating set for $\mathcal{D}(\phi)$. \square

Theorem 6.21 *$\mathcal{D}(\phi) \subset \text{WZ}(C(\phi))$ for all rotationless ϕ .*

Proof $\mathcal{D}(\phi) \subset \mathcal{A}(\phi) \subset \text{WZ}(C(\phi))$ by Corollary 6.20 and Corollary 5.6. \square

7 Finite index

Our goal in this section is to prove:

Theorem 7.1 $\mathcal{D}(\phi)$ has finite index in $\mathcal{A}(\phi)$ for all rotationless ϕ .

Before turning to the proof of Theorem 7.1 we use it to prove one of our main results.

Theorem 7.2 For every abelian subgroup A of $\text{Out}(F_n)$ there exists $\phi \in A$ such that $A \cap \mathcal{D}(\phi)$ has finite index in A .

Proof Corollary 3.14 and Lemma 5.4 imply that $A \cap \mathcal{A}(\phi)$ has finite index in A for each generic $\phi \in A_R$. Theorem 7.1 therefore completes the proof. \square

Choose once and for all a CT $f: G \rightarrow G$ representing ϕ .

We set notation for the linear edges associated to an axis $[c]_u$ of ϕ following (Linear edges). If $[c]_u$ has multiplicity $m + 1$ then there is a root-free closed path w whose circuit represents $[c]_u$ and for $1 \leq j \leq m$, there are linear edges E_j and distinct nonzero integers d_j such that $f(E_j) = E_j \cdot w^{d_j}$. Choose a lift \tilde{w} of w that is contained in the axis $A_c \subset \Gamma$ and let \tilde{E}_j be the lift of E_j whose terminal endpoint is the initial endpoint of \tilde{w} . The lift \tilde{f}_j of f that fixes the initial endpoint of \tilde{E}_j is principal; the associated principal automorphism is denoted Φ_j . Both \tilde{f}_j and Φ_j are independent of the choice of \tilde{w} . By Lemma 4.10 and Lemma 5.8, $\text{Fix}(\hat{\Phi}_j)$ is a maximal principal set for $\mathcal{A}(\phi)$ that we denote \mathcal{X}_j . The lift s_j of $\mathcal{A}(\phi)$ to $\text{Aut}(F_n)$ determined by \mathcal{X}_j satisfies $s_j(\phi) = \Phi_j$. The principal lift of f that fixes the terminal endpoint of \tilde{E}_j is denoted \tilde{f}_0 , its associated principal automorphism is denoted Φ_0 , the maximal principal set $\text{Fix}(\hat{\Phi}_0)$ is denoted \mathcal{X}_0 and the lift to $\text{Aut}(F_n)$ determined by \mathcal{X}_0 is denoted s_0 . The automorphisms Φ_0, \dots, Φ_m are the only elements of $\text{P}(\phi)$ that commute with T_c (Lemma 2.25).

Recall that in Definition 6.10, the notation $\phi_{\mathbf{a}}$ was introduced and $\mathcal{D}(\phi)$ was defined as $\langle \phi_{\mathbf{a}} \mid \mathbf{a} \text{ is admissible} \rangle$. In particular, we only write $\phi_{\mathbf{a}}$ if \mathbf{a} is admissible. We saw that $\mathcal{D}(\phi) \subset \mathcal{A}(\phi)$ in Corollary 6.20.

For $1 \leq j \neq k \leq m$, let $\omega_{c,j}$ be the comparison homomorphism determined by X_0 and X_j and let $\omega_{c,j,k}$ be the comparison homomorphism determined by X_j and X_k . Thus $\omega_{c,j,k} = \omega_{c,j} - \omega_{c,k}$. There is an obvious bijection between the $\omega_{c,j}$'s and the linear edges E_j associated to c . There is also a bijection between the $\omega_{c,j,k}$'s and the families of quasi-exceptional paths $E_j w^* \bar{E}_k$ associated to c . We make use of these bijections without further notice.

For each $\Lambda \in \mathcal{L}(\phi)$ let $\omega_\Lambda = \text{PF}_\Lambda | \mathcal{A}(\phi)$. We also identify Λ with ω_Λ when convenient.

We define a new homomorphism $\Omega^\phi: \mathcal{A}(\phi) \rightarrow \mathbb{Z}^K$ whose coordinates are in one to one correspondence with the linear and EG strata of $f: G \rightarrow G$ by removing extraneous coordinates from $\Omega: \mathcal{A}(\phi) \rightarrow \mathbb{Z}^N$.

Definition 7.3 $\Omega^\phi: \mathcal{A}(\phi) \rightarrow \mathbb{Z}^K$ is the product of the $\omega_{c,j}$'s and the ω_Λ 's as $[c]_u$ varies over the axes of ϕ and as Λ varies over $\mathcal{L}(\phi)$.

Lemma 7.4 $\Omega^\phi: \mathcal{A}(\phi) \rightarrow \mathbb{Z}^K$ is injective.

Proof The coordinates of Ω^ϕ are coordinates of the injective homomorphism Ω . It therefore suffices to assume that $\omega(\psi) \neq 0$ for a coordinate ω of Ω and prove that the image of ψ under some coordinate of Ω^ϕ is nonzero. There is no loss in assuming that ω is not a coordinate of Ω^ϕ and so by Lemma 5.8 and Remark 4.8 is either some $\omega_{c,j,k}$ or ω_Λ for some $\Lambda \in \mathcal{L}(\phi^{-1})$. In the former case, $\omega_{c,j}(\psi) \neq 0$ or $\omega_{c,k}(\psi) \neq 0$ and we are done. In the latter case, Corollary 3.3.1 of [2] implies that $\Lambda \in \mathcal{L}(\psi) \cup \mathcal{L}(\psi^{-1})$. By Lemma 3.2.4 of [2] there is a unique $\Lambda' \in \mathcal{L}(\phi) \cup \mathcal{L}(\phi^{-1})$ such that $\Lambda' \neq \Lambda$ and such that Λ and Λ' are carried by the same minimal rank free factor; moreover, $\Lambda' \in \mathcal{L}(\phi)$. Similarly there is a unique $\Lambda'' \in \mathcal{L}(\psi) \cup \mathcal{L}(\psi^{-1})$ such that $\Lambda'' \neq \Lambda$ and such that Λ and Λ'' are carried by the same minimal rank free factor. Lemma 5.8 and Remark 4.8 imply that $\Lambda'' \in \mathcal{L}(\phi) \cup \mathcal{L}(\phi^{-1})$ and hence that $\Lambda'' = \Lambda' \in \mathcal{L}(\phi)$. Thus $\omega_{\Lambda''}$ is a coordinate of Ω^ϕ and $\omega_{\Lambda''}(\psi) \neq 0$. \square

Lemma 7.5 If $\phi_{\mathbf{a}} \in \mathcal{A}(\phi)$ and if a coordinate ω of Ω^ϕ corresponds to a stratum in the almost invariant subgraph X_s then $\omega(\phi_{\mathbf{a}}) = a_s \omega(\phi)$.

Proof We may assume by Corollary 6.20 that \mathbf{a} is generic. If $\omega = \omega_\Lambda$ then the lemma follows from Corollary 6.14 and the definition of the expansion factor homomorphism. Suppose then that $\omega = \omega_{c,j}$. Lemma 2.25 implies that $s_j(\phi_{\mathbf{a}})$ corresponds to the principal lift of $f_{\mathbf{a}}$ that fixes the initial endpoint of \tilde{E}_j and $s_0(\phi_{\mathbf{a}})$ corresponds to the principal lift of $f_{\mathbf{a}}$ that fixes the terminal endpoint of \tilde{E}_j . Since $f_{\mathbf{a}}(E_j) = E_j \cdot w^{a_s d_j}$ we have $\omega_{c,j}(\phi_{\mathbf{a}}) = a_s d_j$. \square

Corollary 7.6 The rank of $\mathcal{D}(\phi)$ is equal to the rank of the subgroup L of \mathbb{Z}^M generated by the admissible M -tuples for $f: G \rightarrow G$.

Proof By Corollaries 6.16 and 6.20, $\mathbf{a} \mapsto \phi_{\mathbf{a}}$ determines a homomorphism $\rho: L \rightarrow \mathcal{A}(\phi)$. It suffices to show that ρ is injective. The subgroup L contains the M -tuple \mathbf{I} , all of whose coordinates are 1. Given distinct $\mathbf{x}, \mathbf{y} \in L$ there exists $k \geq 0$ so that

$\mathbf{a} = \mathbf{x} + k\mathbf{I}$ and $\mathbf{b} = \mathbf{y} + k\mathbf{I}$ are admissible. Lemma 7.5 implies that $\Omega^\phi(\phi_{\mathbf{a}}) \neq \Omega^\phi(\phi_{\mathbf{b}})$ and hence by Lemma 7.4 that $\rho(\mathbf{x} + k\mathbf{I}) = \phi_{\mathbf{a}} \neq \phi_{\mathbf{b}} = \rho(\mathbf{y} + k\mathbf{I})$. Since ρ is a homomorphism $\rho(\mathbf{x}) \neq \rho(\mathbf{y})$. \square

We now come to our main technical proposition, a generalization of Lemma 2.24. (The process of iterating an edge is discussed in Section 2. Coordinate homomorphisms are reviewed at the beginning of this section.)

Proposition 7.7 *Suppose that $\tilde{f}: \Gamma \rightarrow \Gamma$ is a principal lift of f , that \tilde{E} is an oriented edge whose initial direction is fixed by $D\tilde{f}$ and that the ray \tilde{R} determined by iterating \tilde{E} converges to $P \in \text{Fix}_N(\hat{f})$. Let $s: \mathcal{A}(\phi) \rightarrow \text{Aut}(F_n)$ be the lift determined by the maximal principal set $\mathcal{X}_s := \text{Fix}(\hat{f})$. Suppose further that μ is a term in the QE-splitting of R that is either an edge in an EG or linear stratum or a quasi-exceptional path. Let $\omega: \mathcal{A}(\phi) \rightarrow \mathbb{Z}$ be the coordinate homomorphism associated to μ . Then the following are equivalent for all $\psi \in \mathcal{A}(\phi)$.*

- (1) P is isolated in $\text{Fix}(s(\widehat{\psi}))$.
- (2) $\omega(\psi) \neq 0$.

Before proving Proposition 7.7 we derive a corollary and use that corollary to prove Theorem 7.1. The set \mathcal{Q}_s is defined in Notation 6.12.

Corollary 7.8 *Suppose that X_s is an almost invariant subgraph and that \mathcal{W}_s is the set of coordinate homomorphisms $\omega: \mathcal{A}(\phi) \rightarrow \mathbb{Z}$ associated to either an edge in an irreducible stratum in X_s or to an element of \mathcal{Q}_s . Then for all $\psi \in \mathcal{A}(\phi)$ either $\omega(\psi) = 0$ for all $\omega \in \mathcal{W}_s$ or $\omega(\psi) \neq 0$ for all $\omega \in \mathcal{W}_s$.*

Proof Recall from Definition 6.3 that B_s is a connected directed graph with one vertex for each nonfixed irreducible stratum in X_s . Define a new directed graph C_s with the same set of vertices and with an edge from the vertex w_i corresponding to H_i to the vertex w_j corresponding to H_j if $i \neq j$ and if for some (hence every) edge E_i in H_i and some (hence every) edge E_j in H_j there exists $k > 0$ so that E_j occurs as a term in the QE-splitting of $f_{\#}^k(E_i)$.

If there is a directed edge from w_i to w_j in B_s but not in C_s then H_i is EG and there is a taken connecting path κ_i in a zero stratum of H_i^z such that some term in the QE-splitting of $f_{\#}(\kappa_i)$ is an edge in H_j . Since κ_i is taken it occurs as a term in the QE-splitting of $f_{\#}^l(E_m)$ for some edge E_m in an irreducible stratum H_m . There are directed edges in C_s from w_m to w_i and from w_m to w_j ; the existence of the latter

is immediate from the definition of C_s and the existence of the former follows from Lemma 6.5 applied to f^l . This proves that C_s is connected.

Enlarge C_s to C'_s by adding a vertex for each quasi-exceptional family α in \mathcal{Q}_s and a directed edge from w_i to the vertex corresponding to α if some element of α occurs as a term in the QE-splitting of $f_{\#}^k(E_i)$ for some (hence every) edge E_i in H_i and some $k > 0$. The graph C'_s is still connected. Note also that if τ is a directed edge path in C'_s then there is a directed edge from the initial endpoint of τ to the terminal endpoint of τ .

For each vertex $w_i \in C_s$ let $Y_s(i)$ be the subgraph of C'_s consisting of all directed edges with initial vertex v_i and let $\mathcal{W}_s(i)$ be the set of $\omega_p \in \mathcal{W}_s$ whose associated vertex is contained in $Y_s(i)$. Note that $Y_s(i)$ contains the terminal endpoint of every edge path in $Y_s(i)$ starting at v_i . We claim that for all $\psi \in \mathcal{A}(\phi)$, either $\omega_p(\psi)$ is zero for all $\omega_p \in \mathcal{W}_s(i)$ or $\omega_p(\psi)$ is nonzero for all $\omega_p \in \mathcal{W}_s(i)$.

The claim is obvious if $\mathcal{W}_s(i)$ contains only one element so we may assume that H_i is either EG or nonlinear NEG. If H_i is NEG then it is a single edge E_i whose initial vertex is principal and whose initial direction is fixed. If H_i is EG then we can choose such an E_i by Remark 2.18. Choose a lift \tilde{E}_i , let \tilde{f} be the principal lift that fixes the initial endpoint of \tilde{E}_i , let $P \in \text{Fix}(\tilde{f})$ be the terminal endpoint of the ray \tilde{R} obtained by iterating \tilde{E}_i by \tilde{f} and let $s: \mathcal{A}(\phi) \rightarrow \text{Aut}(F_n)$ be the lift determined by the maximal principal set $\text{Fix}(\tilde{f})$. For each $\omega_p \in \mathcal{W}_s(i)$, there is a term in the QE-splitting of \tilde{R} that corresponds to ω_p . The claim therefore follows from Proposition 7.7 since P being isolated in $\text{Fix}(\tilde{f})$ is independent of ω_p .

To complete the proof of the corollary it suffices to show that if $Y_s(i) \cap Y_s(j) \neq \emptyset$ then $\mathcal{W}_s(i) \cap \mathcal{W}_s(j) \neq \emptyset$. This could only fail if every vertex in $Y_s(i) \cap Y_s(j)$ corresponds to a nonlinear NEG stratum. But this is impossible since every such vertex has at least one outgoing edge. \square

Proof of Theorem 7.1 For each coordinate ω_i of Ω^ϕ and each $\psi \in \mathcal{A}(\phi)$, define $a_i(\psi) = \omega_i(\psi)/\omega_i(\phi) \in \mathbb{Q}$. Since $a_i(\phi\psi) = a_i(\psi) + 1$ there is a finite generating set of elements ψ with the property that each $a_i(\psi) > 0$. It suffices to show that under this hypothesis, $\psi^K \in \mathcal{D}(\phi)$ for some $K > 0$.

As we are now working with a single ψ , we refer to $a_i(\psi)$ simply as a_i . After replacing ψ with an iterate, we may assume that each a_i is a positive integer. Define $\theta_i = \psi\phi^{-a_i}$ and note that

$$\omega_i(\theta_i) = \omega_i(\psi) - a_i\omega_i(\phi) = 0.$$

Let X_s be the almost invariant subgraph that contains the stratum associated to ω_i and let ω_j be another coordinate of Ω^ϕ that is associated to a stratum in X_s . Corollary 7.8 implies that

$$\omega_j(\psi) - a_i\omega_j(\phi) = \omega_j(\theta_i) = 0$$

and hence that

$$\omega_j(\psi) = a_i\omega_j(\phi).$$

This shows that $a_i = a_j$ so the a_i 's determine a well defined M -tuple $\hat{\mathbf{a}} = (\hat{a}_1, \dots, \hat{a}_M)$ with one \hat{a}_s for each almost invariant subgraph X_s .

Suppose that $\omega \in \mathcal{Q}_s$ corresponds to the quasi-exceptional family containing $E_\alpha \bar{E}_\beta$ where $E_\alpha \subset X_\alpha$ and $E_\beta \subset X_\beta$ are linear edges with exponent d_α and d_β . As above, Corollary 7.8 implies that

$$\omega(\psi) = \hat{a}_s\omega(\phi).$$

Letting ω_α and ω_β be the coordinates of Ω^ϕ associated to E_α and E_β we have

$$\hat{a}_s(d_\alpha - d_\beta) = \hat{a}_s(\omega_\alpha(\phi) - \omega_\beta(\phi)) = \hat{a}_s\omega(\phi)$$

and $\omega(\psi) = \omega_\alpha(\psi) - \omega_\beta(\psi) = \hat{a}_\alpha\omega_\alpha(\phi) - \hat{a}_\beta\omega_\beta(\phi) = \hat{a}_\alpha d_\alpha - \hat{a}_\beta d_\beta$.

The last three displayed equations show that $\hat{\mathbf{a}}$ is admissible.

Corollary 6.20 implies that $\phi_{\hat{\mathbf{a}}} \in \mathcal{A}(\phi)$. Lemmas 7.5 and 7.4 then imply that $\psi = \phi_{\hat{\mathbf{a}}} \in \mathcal{D}(\phi)$ as desired. \square

The remainder of the section is devoted to the proof of Proposition 7.7. For motivation we consider the proof as it applies to a simple example.

Example 7.9 Suppose that $G = R_3$ with edges A , B and C and that $f: G \rightarrow G$ representing ϕ is defined by $A \mapsto A$, $B \mapsto BA$ and $C \mapsto CB$. In the notation of Proposition 7.7, C plays the role of E and B plays the role of μ .

Let T_A be the covering translation corresponding to A and let \tilde{B} be a lift of B with terminal endpoint in the axis of T_A . Denote the principal lifts of f that fix the initial and terminal endpoints of \tilde{B} by \tilde{f}_- and \tilde{f}_+ respectively. The fixed point sets \mathcal{X}_\pm of \tilde{f}_\pm are maximal principal sets for $\mathcal{A}(\phi)$ and so determine lifts $s_\pm: \mathcal{A}(\phi) \rightarrow \text{Aut}(F_n)$ such that $X_\pm \subset \text{Fix}(s_\pm(\psi))$ for all $\psi \in \mathcal{A}(\phi)$. The coordinate homomorphism ω corresponding to B satisfies $\omega(\psi) = 0$ if and only if $s_+(\psi) = s_-(\psi)$. Note that T_A^\pm is contained in both X_+ and X_- .

Choose a lift \tilde{C} of C and let \tilde{f} be the principal lift that fixes its initial endpoint. Iterating \tilde{C} by \tilde{f} produces a ray \tilde{R} that converges to some $P \in \text{Fix}(\tilde{f})$ and that projects to an f -invariant ray $R = CBBA^2 \dots BA^l BA^{l+1} BA^{l+2} \dots$. The maximal principal

set $\text{Fix}(\hat{f})$ determines a lift $s: \mathcal{A}(\phi) \rightarrow \text{Aut}(F_n)$. Denote the subpath $BBABA^2$ of R that follows the initial C by σ_0 and the subpath $f_{\#}^l(\sigma_0) = BA^l BA^{l+1} BA^{l+2}$ of R by σ_l . There are lifts $\tilde{\sigma}_l \subset \tilde{R}$ of σ_l , $l \rightarrow \infty$, that are cofinal in \tilde{R} and so limit on P .

There are also lifts $\tilde{\delta}_l$ of σ_l for which \tilde{B} is the edge that projects to the middle B in σ_l . The endpoints of $\tilde{\delta}_l$ are denoted \tilde{x}_l and \tilde{y}_l . The path connecting \tilde{x}_l to the initial endpoint of \tilde{B} is a lift of BA^l and the path connecting the terminal endpoint of \tilde{B} to $T_A^{-l} \tilde{y}_l$ is a lift of ABA^{l+2} . Thus $\tilde{x}_l \rightarrow Q_- \in X_- \setminus T_A^{\pm}$ and $T_A^{-l} \tilde{y}_l \rightarrow Q_+ \in X_+ \setminus T_A^{\pm}$. The line connecting Q_- to Q_+ projects to $A^{\infty} B A B A^{\infty}$.

Choose a CT $g: G' \rightarrow G'$ representing ψ . For simplicity, we assume that $G' = G$. The lift $\tilde{g}: \Gamma \rightarrow \Gamma$ corresponding to $s(\psi)$ satisfies $P \in \text{Fix}(\hat{g})$.

If P is not isolated in $\text{Fix}(\hat{g})$ then Lemma 2.2 implies that \tilde{g} moves the endpoints of $\tilde{\sigma}_l$ by an amount D that is bounded above independently of l . Since $\tilde{\delta}_l$ is a translate of $\tilde{\sigma}_l$ there is a lift \tilde{g}_l of g that moves \tilde{x}_l and \tilde{y}_l by at most D . In Lemma 7.11 below we show that under these circumstances, \tilde{g}_l commutes with T_A for all sufficiently large l . The lift $s_-(\psi)$ of g commutes with T_A and fixes Q_- . Since \tilde{g}_l also commutes with T_A there exists d_l such that $\tilde{g}_l = T_A^{d_l} s_-(\psi)$. In particular, $\tilde{g}_l(Q_-) = T_A^{d_l}(Q_-)$. If $d_l \neq 0$ and \tilde{x}_l is sufficiently close to Q_- then the distance between \tilde{x}_l and $\tilde{g}_l(\tilde{x}_l)$ would be greater than D which is a contradiction. Thus $d_l = 0$ and $Q_- \in \text{Fix}(\hat{g}_l)$ for all sufficiently large l . A second consequence of the fact that \tilde{g}_l commutes with T_A is that \tilde{g}_l moves $T_A^{-l} \tilde{y}_l$ by a uniformly bounded amount. Arguing as in the previous case we conclude that $Q_+ \in \text{Fix}(\hat{g}_l)$ for all sufficiently large l . For these l , $\text{Fix}(\hat{g}_l)$ intersects both X_+ and X_- in at least three points which implies that \tilde{g}_l is the lift associated to both $s_-(\psi)$ and $s_+(\psi)$ and hence that $\omega(\psi) = 0$.

If P is isolated in $\text{Fix}(\hat{g})$ then by Lemma 2.23 there is an edge \tilde{E}' of Γ' that iterates toward P under the action of \tilde{g} . The ray \tilde{R}' connecting \tilde{E}' to P eventually agrees with \tilde{R} and so contains $\tilde{\sigma}_l$ for large l . Lemma 7.13 below states, roughly speaking, that since iterating E' by g produces segments of the form $BA^l B$ for arbitrarily large l , it must be that $g_{\#}(BAB) = BA^k B$ for some $k > 1$. This implies that $A^{\infty} B A B A^{\infty}$ is not $g_{\#}$ -invariant and hence that the lifts of g' corresponding to $s_-(\psi)$ and to $s_+(\psi)$ are distinct. Equivalently, $\omega(\psi) \neq 0$.

We now turn to the formal proof.

Remark 7.10 For the following lemmas it is useful to recall that if the circuits representing the conjugacy classes $[b]$ and $[c]$ of root-free elements $b, c \in F_n$ have edge length L_b and L_c and if $A_b \cap A_c$ has edge length at least $L_b + L_c$ then T_c commutes with T_b because the initial endpoint \tilde{x} of $A_b \cap A_c$ satisfies $T_b T_c(\tilde{x}) = T_c T_b(\tilde{x})$. It follows that $A_b = A_c$ and that $T_b = T_c^{\pm}$.

Lemma 7.11 *Suppose that $\psi \in \text{Out}(F_n)$ is rotationless and that $g: G' \rightarrow G'$ is a CT representing ψ . Then for any root-free covering translation T_c of the universal cover Γ' of G' , there exists $K > 0$ with the following property. If $\tau \subset G'$ is a Nielsen path for g and $\tilde{\tau} \subset \Gamma'$ is a lift whose intersection with the axis A_c of T_c contains at least K edges, then the lift \tilde{g} that fixes the endpoints of $\tilde{\tau}$ commutes with T_c .*

Proof Choose L greater than the number of edges in each of the following:

- (1) the loop in G' that represents c
- (2) each of the loops in G' representing an axis of ψ
- (3) any indivisible Nielsen path associated to an EG stratum for $g: G' \rightarrow G'$.

There is a decomposition $\tilde{\tau} = \tilde{\tau}_1 \cdots \tilde{\tau}_N$ into subpaths $\tilde{\tau}_i$ that are either fixed edges or indivisible Nielsen paths. The endpoints of the $\tilde{\tau}_i$'s are fixed by \tilde{g} . There is no loss in assuming that each $\tilde{\tau}_i$ intersects A_c in at least an edge.

If $N \geq L + 1$ then by (1), there exist $\tilde{\tau}_i$ with initial endpoint \tilde{x} and $\tilde{\tau}_j$ with initial endpoint $T_c^l(\tilde{x})$ for some $l \neq 0$. Thus $\tilde{g}T_c^l(\tilde{x}) = T_c^l(\tilde{x}) = T_c^l\tilde{g}(\tilde{x})$. Since lifts of a map that agree on a point are identical, $\tilde{g}T_c^l = T_c^l\tilde{g}$. It follows that \tilde{g} fixes T_c^\pm which then implies that \tilde{g} commutes with T_c .

We may therefore assume $N < L$. In fact we may assume that $N = 1$: if K works in this case then $(L + 2)K$ works in the general case. If τ is a fixed edge then $K = 2$ vacuously works. We may therefore assume that τ is indivisible.

Let $K = 2L + 2$. We may assume by (3) that τ is not associated to an EG stratum. By the (NEG Nielsen paths) property for g , $\tilde{\tau} = \tilde{E}_i \tilde{w}^p \tilde{E}_i^{-1}$ for some linear edge E_i satisfying $f(E_i) = E_i w^{d_i}$ where w represents an axis μ of ψ and therefore has fewer than L edges. There is an axis A_b for a root-free $b \in F_n$ that contains \tilde{w}^p and whose projection into G' is the loop determined by w . Remark 7.10 and our choice of K imply that $T_b = T_c^\pm$. Both $T_b^p \tilde{g}$ and $\tilde{g} T_b^p$ take the initial endpoint of $\tilde{\tau}$ to the terminal endpoint of $\tilde{\tau}$. Since these are both lifts of g they must be equal. This proves that \tilde{g} commutes with T_b and so also commutes with T_c . \square

Suppose that E_i is a linear edge and that $f(E_i) = E_i w^{d_i}$. If either E_i or a quasi-exceptional path $E_i w^* \bar{E}_j$ occurs as a term in the quasi-exceptional splitting of some $f_\#^l(\sigma)$ then $f_\#^m(\sigma)$ contains subpaths of the form w^k where $k \rightarrow \pm\infty$ as $m \rightarrow \infty$. This is essentially the only way that such paths develop under iteration. Lemma 7.13 below is an application of this observation stated in the way that it is applied in the proof of Proposition 7.7.

We use $\text{EL}(\cdot)$ to denote edge length of a path or circuit. By extension, for $c \in F_n$, we use $\text{EL}(c)$ to denote the edge length of the circuit representing $[c]$.

Directions based at nonprincipal vertices of a CT $f: G \rightarrow G$ need not stabilize under iteration by Df . It is sometimes convenient to pass to a power of $f: G \rightarrow G$ so that every direction d based at a vertex of G is *pre-fixed*, meaning that $Df^i(d)$ is fixed by Df for some $i > 0$.

Lemma 7.12 *Suppose that $g: G' \rightarrow G'$ is a CT, that every direction based at a vertex of g is pre-fixed, and that $\tau \subset G'$ is a completely split path such that $\text{EL}(g_{\#}^m(\tau))$ is not uniformly bounded from above. Then for all $L > 0$ there exists $M > 0$ so that for all $m \geq M$, $\text{EL}(g_{\#}^m(\tau)) > 2L$ and the initial and terminal subpaths of $g_{\#}^m(\tau)$ with edge length L are independent of m .*

Proof The proof is by induction on the height r of τ . The $r = 0$ case is vacuous so we may assume that the lemma holds for paths of height less than r . By symmetry it is sufficient to show that $\text{EL}(g_{\#}^m(\tau)) \rightarrow \infty$ and that initial segment of $g_{\#}^m(\tau)$ with edge length L stabilizes under iteration.

Let $\tau = \tau_1 \cdots \tau_s$ be the complete splitting of τ and let τ_i be the first term such that $\text{EL}(g_{\#}^m(\tau_i))$ is not uniformly bounded from above. The terms preceding τ_i , if any, are Nielsen paths or pre-Nielsen (meaning that they are mapped by some iterate of $g_{\#}$ to a Nielsen path) connecting paths in zero strata. Their iterates stabilize so there is no loss in truncating τ by removing them. We may therefore assume that $i = 1$. It now suffices to show that $\text{EL}(g_{\#}^m(\tau_1)) \rightarrow \infty$ and that initial segment of $g_{\#}^m(\tau_1)$ with edge length L stabilizes under iteration. If τ_1 is a connecting path in a zero stratum then this follows by induction on r . The remaining cases are that τ_1 is a nonfixed edge in an irreducible stratum or a quasi-exceptional path and the result is clear in both these cases. \square

The following lemma is a case-by-case analysis of the occurrence of long periodic segments in iterates of a single path. The basic observation is that once a periodic segment reaches a certain length it continues to get longer under further iteration.

Lemma 7.13 *Suppose that $g: G' \rightarrow G'$ is a CT, that every direction based at a vertex of g is pre-fixed, that $c \in F_n$ is root-free and that $\tilde{g}: \Gamma' \rightarrow \Gamma'$ is a lift of g that commutes with T_c . Then for all completely split paths $\sigma \subset G'$, there exists $L_{\sigma} > 0$ so that if $m \geq 0$ and $\tilde{\rho}_m$ is a lift of $\rho_m = g_{\#}^m(\sigma)$ such that $\text{EL}(\tilde{\rho}_m \cap A_c) > L_{\sigma}$ then $\text{EL}(\tilde{g}_{\#}(\tilde{\rho}_m) \cap A_c) > \text{EL}(\tilde{\rho}_m \cap A_c)$.*

Proof It suffices to show that the lemma holds for all sufficiently large m so there is no loss in replacing σ by $g_{\#}(\sigma)$ when this is useful.

Lemma 2.1 implies that the circuit μ corresponding to c is $g_{\#}$ -invariant. Since some $g_{\#}$ -iterate of μ has a complete splitting [10, Lemma 4.26], μ has a complete splitting; each term in this complete splitting is either a g -fixed edge or an indivisible Nielsen path for g . There is an induced complete splitting of A_c with respect to \tilde{g} . There is a lift of g that fixes the endpoints of each term in this splitting and that commutes with T_c , and so equals $T_c^k \tilde{g}$ for some k . After replacing \tilde{g} by $T_c^k \tilde{g}$, we may assume that all the terms in the complete splitting of A_c are \tilde{g} -Nielsen paths. The endpoints of these Nielsen paths are called *splitting vertices*. Note that the set of splitting vertices coincides with the set of \tilde{g} -fixed vertices in A_c .

The proof is by induction on the height r of σ . The induction statement is enhanced to include the following property: if $\text{EL}(\tilde{\rho}_m \cap A_c) > L_{\sigma}$ and if $\tilde{\rho}_m \cap A_c$ contains an endpoint \tilde{v} of $\tilde{\rho}_m$ then \tilde{v} is a splitting vertex.

In certain cases we will show that $\text{EL}(A_c \cap \tilde{\rho}_m)$ is uniformly bounded, meaning that it is bounded independently of m . One then chooses L_{σ} greater than that bound. The $r = 0$ case is vacuously true so we may assume that the inductive statement holds for all paths of height less than r .

Assume for now that there is only one term in the QE-splitting of σ . There are five cases, two of which are immediate. If σ is a Nielsen path then $\text{EL}(A_c \cap \tilde{\rho}_m)$ is uniformly bounded and we are done. If σ is a connecting path in a zero stratum then we let $L_{\sigma} = L_{g_{\#}(\sigma)}$ where the latter exists by the inductive hypothesis and the fact that $g_{\#}(\sigma)$ has height less than r .

If σ is a linear edge E then $\rho_m = Ew^{dm}$ for some Nielsen path w that forms a root-free circuit and some $d > 0$. Let $L_{\sigma} = \text{EL}(c) + \text{EL}(w)$. If $\text{EL}(\tilde{\rho}_m \cap A_c) > L_{\sigma}$ then by Remark 7.10 there is a lift \tilde{w} of w such that $\tilde{\rho}_m \cap A_c = \tilde{w}^{dm}$ contains all of $\tilde{\rho}_m$ but the initial edge, and $\tilde{g}_{\#}(\tilde{\rho}_m) \cap A_c = \tilde{w}^{d(m+1)}$. Since w is a Nielsen path and \tilde{w} is a fundamental domain of A_c the endpoints of \tilde{w} are splitting vertices.

If σ is a quasi-exceptional path $E_i w^p \bar{E}_j$ where $g(E_i) = E_i w^{d_i}$ and $g(E_j) = E_j w^{d_j}$, then the proof is similar to the linear case and we can use the same value of L_{σ} . If $\text{EL}(\tilde{\rho}_m \cap A_c) > L_{\sigma}$ then there is a lift \tilde{w} of w such that $\tilde{\rho}_m \cap A_c = \tilde{w}^{m(d_i - d_j) + p}$ contains all of $\tilde{\rho}_m$ but the initial and terminal edges, and $\tilde{g}_{\#}(\tilde{\rho}_m) \cap A_c = \tilde{w}^{(m+1)(d_i - d_j) + p}$. In this case the endpoints of $\tilde{\rho}_m$ are not contained in A_c .

The fifth and hardest case is that σ is a single edge E in a nonlinear irreducible stratum H_r . If the height of A_c is greater than r then $A_c \cap \tilde{\rho}_m$ has uniformly bounded

length. We may therefore assume that A_c has height at most r . We consider the EG and NEG subcases separately.

If H_r is EG then ρ_m is r -legal and so does not contain an indivisible Nielsen path of height r . If A_c contains an indivisible Nielsen path of height r then $\text{EL}(A_c \cap \tilde{\rho}_m) < \text{EL}(c)$ and we are done. We may therefore assume that A_c has height less than r . In particular, the endpoints of $\tilde{\rho}_m$ are not contained in A_c . There are no quasi-exceptional paths and no fixed edges of height r . Thus the terms in the QE-splitting of $g(E)$ are either single edges in H_r or are contained in G_{r-1} . After amalgamating terms we have a splitting of $g(E)$ into r -legal subpaths in H_r and completely split subpaths in G_{r-1} . There is a similar splitting of $g(E')$ for each edge E' of H_r . Let $\{\mu_j\}$ be the set of completely split paths of G_{r-1} that occur in this way as E' varies over all edges of H_r . An easy induction argument shows that $\rho_m = g_{\#}^m(E)$ has a splitting into r -legal subpaths in H_r and completely split subpaths in G_{r-1} ; each of the subpaths in G_{r-1} equals $g_{\#}^l(\mu_j)$ for some μ_j and some $0 \leq l \leq m$. We may therefore choose $L_{\sigma} = \max\{L_{\mu_j}\}$.

Finally, suppose that H_r is nonlinear and NEG. There is a path $u \subset G_{r-1}$ such that $g^m(E) = E \cdot u \cdot g_{\#}(u) \cdots g_{\#}^{m-1}(u)$ for all m and such that $\text{EL}(g_{\#}^j(u)) \rightarrow \infty$. We may assume without loss that A_c has height less than r and hence that $\tilde{\rho}_m \cap A_c$ projects into $u \cdot g_{\#}(u) \cdots g_{\#}^m(u)$. In particular, the initial endpoint of $\tilde{\rho}_m$ is not contained in A_c . We claim that if r is sufficiently large, say $r > R$, then the projection of $\tilde{\rho}_m \cap A_c$ does not contain $g_{\#}^r(u)$ for any m . Assume the claim for now. If $\text{EL}(\tilde{\rho}_m \cap A_c) > \text{EL}(u \cdot g_{\#}(u) \cdots g_{\#}^{R+1}(u))$ then the projection of $\tilde{\rho}_m \cap A_c$ is contained in $g_{\#}^{q-1}(u) \cdot g_{\#}^q(u) = g_{\#}^{q-1}(u \cdot g_{\#}(u))$ for some q . We may therefore choose L_{σ} to be the maximum of $\text{EL}(u \cdot g_{\#}(u) \cdots g_{\#}^{R+1}(u))$ and $L_{u \cdot g_{\#}(u)}$. If $\text{EL}(\tilde{\rho}_m \cap A_c) > L_{\sigma}$ and if the terminal vertex \tilde{v} of $\tilde{\rho}_m$ is contained in A_c then $\tilde{\rho}_m \cap A_c$ is a terminal segment of a lift of $g_{\#}^{m-1}(u \cdot g_{\#}(u))$ and \tilde{v} is a splitting vertex of A_c by the inductive hypothesis.

The claim is obvious unless u and A_c have the same height, say t , so assume that this is the case. The claim is also obvious if the maximal length of a subpath of $g_{\#}^q(u)$ with height less than t goes to infinity with q . We may therefore assume that the number of height t edges in $g_{\#}^q(u)$ goes to ∞ with q . Thus H_t is EG and $g_{\#}^r(u)$ contains t -legal subpaths of length greater than $\text{EL}(c)$ for all sufficiently large r . Since no such subpath is contained in A_c this completes the proof of the claim and so also the induction step when there is only one term in the QE-splitting of σ .

Assume now that $\sigma = \sigma_1 \cdots \sigma_s$ is the QE-splitting of σ and that $s > 1$. Let $L_1 = \max\{L_{\sigma_i}\}$. By Lemma 7.12 there exists $M > 0$ so that for all $m > M$ and all σ_i , either $g_{\#}^m(\sigma_i)$ is independent of m or $\text{EL}(g_{\#}^m(\sigma_i)) > 2L_1$ and the initial and terminal segments

of $g_{\#}^m(\sigma_i)$ with edge length L_1 are independent of m . The former corresponds to σ_i being a Nielsen path or a pre-Nielsen connecting path in a zero stratum and the latter to all remaining cases. Choose $L_{\sigma} > sL_1$ so that $\text{EL}(g_{\#}^m(\sigma)) < L_{\sigma}$ for all $m \leq M$.

Denote $g_{\#}^m(\sigma_i)$ by $\rho_{i,m}$ and write $\tilde{\rho}_m = \tilde{\rho}_{1,m} \cdots \tilde{\rho}_{s,m}$. If $\text{EL}(\tilde{\rho} \cap A_c) \geq L_{\sigma}$ then $\text{EL}(\tilde{\rho}_{i,m} \cap A_c) \geq L_{\sigma_i}$ for some $1 \leq i \leq s$. Thus $\text{EL}(\tilde{g}_{\#}(\tilde{\rho}_{i,m}) \cap A_c) > \text{EL}(\tilde{\rho}_{i,m} \cap A_c)$. If $\tilde{\rho}_m \cap A_c \subset \tilde{\rho}_{i,m}$ we are done. Otherwise we may assume that $\tilde{\rho}_{i+1,m} \cap A_c$ is a nontrivial initial segment of $\tilde{\rho}_{i+1,m}$ that begins at a splitting vertex of A_c . (This is where the enhanced induction hypothesis is used.) If $\rho_{i+1,m}$ is a Nielsen path then $\tilde{g}_{\#}(\tilde{\rho}_{i+1,m}) = \tilde{\rho}_{i+1,m}$ so $g_{\#}(\tilde{\rho}_{i+1,m}) \cap A_c = \tilde{\rho}_{i+1,m} \cap A_c$. This same equality holds if $\text{EL}(\tilde{\rho}_{i+1,m} \cap A_c) \leq L_1$ by our choice of M . Finally, if $\text{EL}(\tilde{\rho}_{i+1,m} \cap A_c) > L_1$ then $\text{EL}(g_{\#}(\tilde{\rho}_{i+1,m}) \cap A_c) > \text{EL}(\tilde{\rho}_{i+1,m} \cap A_c)$. This completes the proof if $\tilde{\rho}_m \cap A_c \subset \tilde{\rho}_{i,m} \tilde{\rho}_{i+1,m}$. Iterating this argument completes the proof in general. \square

We need one more lemma before proving the main proposition.

Lemma 7.14 *Suppose that $g': G' \rightarrow G'$ is a CT, that σ is a completely split non-Nielsen path for g and that $\tilde{\sigma} \subset \Gamma'$ is a lift of σ with endpoints at vertices \tilde{x} and \tilde{y} . If $\tilde{g}': \Gamma' \rightarrow \Gamma'$ is a principal lift that fixes \tilde{x} then $\lim_{k \rightarrow \infty} \tilde{g}'^k(\tilde{y}) \rightarrow Q$ for some $Q \in \text{Fix}_N(\hat{g})$.*

Proof There is no loss in assuming that σ is either a single nonfixed edge or an exceptional path $E\tau^l\bar{E}'$. In the former case the lemma follows from Lemma 2.23. In the latter case, Q is an endpoint of the axis of a covering translation corresponding to τ . \square

Proof of Proposition 7.7 Without loss we may replace f by a power and so may assume that all directions based at vertices are pre-fixed.

The case that μ is an EG edge follows from Lemma 2.24. In the remaining cases there is an axis $[c]_{\mu}$ associated to μ and we let $T_c, \Phi_0, \{\Phi_i\}, \{E_i\}$ and $\{d_i\}$ be as chosen at the beginning of this section; see also Lemma 2.25. Thus μ is either E_j for some j or an element of the quasi-exceptional family determined by $E_j\bar{E}_{j'}$ for some j and j' .

Letting \tilde{u} be the path such that $\tilde{f}(\tilde{E}) = \tilde{E} \cdot \tilde{u}$, we have $\tilde{R} = \tilde{E} \cdot \tilde{R}_0$ where $\tilde{R}_0 = \tilde{u} \cdot \tilde{f}_{\#}(\tilde{u}) \cdot \tilde{f}_{\#}^2(\tilde{u}) \cdots$. Since \tilde{E} is not linear, $\tilde{\mu}$ occurs infinitely often as a term in the QE-splitting of \tilde{R}_0 , where we do not distinguish between elements of the same quasi-exceptional family of subpaths. There is a completely split subpath $\tilde{\sigma}_0$ of \tilde{R}_0 and a coarsening $\tilde{\sigma}_0 = \tilde{\tau}_1 \cdot \tilde{\mu} \cdot \tilde{\tau}_2$ of the QE-splitting of σ_0 where $\tilde{\mu}$ is a lift of μ and where $\tilde{\tau}_1$ and $\tilde{\tau}_2$ are not Nielsen paths. Denote the initial and terminal endpoints of $\tilde{\sigma}_0$ by \tilde{a}_0 and \tilde{b}_0 and for $l \geq 1$, let $\tilde{\sigma}_l = \tilde{f}_{\#}^l(\tilde{\sigma}_0)$, $\tilde{a}_l = \tilde{f}^l(\tilde{a}_0)$, and $\tilde{b}_l = \tilde{f}^l(\tilde{b}_0)$. Then:

- (1) $\tilde{\sigma}_l \subset \tilde{R}_0$ and $\tilde{\sigma}_l \rightarrow P$.

Let \tilde{f}_j be the lift of f corresponding to Φ_j and let \tilde{E}_j be a lift of E_j whose initial endpoint is fixed by \tilde{f}_j and whose terminal endpoint is contained in A_c . There is a covering translation $S_0: \Gamma \rightarrow \Gamma$ such that \tilde{E}_j is the initial edge of $S_0(\tilde{\mu})$. Let $\tilde{\delta}_0 = S_0(\tilde{\sigma}_0)$. For $l \geq 1$, let $S_l: \Gamma \rightarrow \Gamma$ be the covering translation such that $\tilde{f}_j^l S_0 = S_l \tilde{f}^l$, let $\tilde{\delta}_l = S_l(\tilde{\sigma}_l)$ and let \tilde{x}_l and \tilde{y}_l be the endpoints of $\tilde{\delta}_l$. It is immediate that:

- (2) $\tilde{E}_j \subset \tilde{\delta}_l$.
- (3) $\tilde{\delta}_l = \tilde{f}_j^l \#(\tilde{\delta}_0)$.
- (4) the length of $\tilde{\delta}_l \cap A_c$ goes to infinity with l .

Lemma 7.14 applied to \tilde{f}_j and $S_0(\tilde{\tau}_1)$ implies that:

- (5) $\tilde{x}_l \rightarrow Q_- \in \text{Fix}_N(\hat{\Phi}_j) \setminus \{T_c^\pm\}$.

If μ corresponds to E_j , let $m = d_j$ and $t = 0$. If μ corresponds to $E_j \bar{E}_{j'}$, let $m = d_j - d_{j'}$ and $t = j'$. Thus $T_c^{-m} \tilde{f}_j = \tilde{f}_t$ and the terminal endpoint of $S_0(\tilde{\mu})$ is fixed by $T_c^{-m} \tilde{f}_j$. Lemma 7.14 applied to $T_c^{-m} \tilde{f}_j$ and $S_0(\tilde{\tau}_2)$ implies that:

- (6) $T_c^{-ml} \tilde{y}_l \rightarrow Q_+ \in \text{Fix}_N(\hat{\Phi}_t) \setminus \{T_c^\pm\}$.

The maximal principal sets $\mathcal{X}_j = \text{Fix}(\Phi_j)$ and $\mathcal{X}_t = \text{Fix}(\Phi_t)$ contain T_c^\pm and determine lifts $s_j, s_t: \mathcal{A}(\phi) \rightarrow \text{Aut}(F_n)$.

We have so far only focused on ϕ . We now bring in ψ . Let $g: G' \rightarrow G'$ be a CT representing ψ and let \tilde{g}, \tilde{g}_j and \tilde{g}_t be lifts of g to the universal cover Γ' corresponding to $\Psi = s(\psi)$, $\Psi_j = s_j(\psi)$ and $\Psi_t = s_t(\psi)$ respectively. The following are equivalent.

- $\omega(\psi) = 0$.
- $\Psi_j = \Psi_t$.
- $Q_+ \in \text{Fix}(\hat{\Psi}_j)$.

It suffices to show that P is isolated in $\text{Fix}(\hat{\Psi})$ if and only if $Q_+ \notin \text{Fix}(\hat{\Psi}_j)$.

To compare points in Γ and Γ' , choose an equivariant map $h: \Gamma \rightarrow \Gamma'$; equivalently, when $\partial\Gamma$ and $\partial\Gamma'$ are identified with ∂F_n then $\hat{h}: \partial\Gamma \rightarrow \partial\Gamma'$ is the identity. Let C be the bounded cancellation constant [5] (see also Lemma 2.3.1 of [2]) for $h: \Gamma \rightarrow \Gamma'$ and let $\tilde{R}' = h_\#(\tilde{R})$. We use prime notation for covering translations and axes of Γ' . Thus $S'_l: \Gamma' \rightarrow \Gamma'$ is the covering translation such that $S'_l h = h S_l$. Denote $h(\tilde{a}_l)$, $h(\tilde{b}_l)$ and the path that they bound by \tilde{a}'_l , \tilde{b}'_l and $\tilde{\sigma}'_l$. Let $\tilde{x}'_l = S'_l(\tilde{a}'_l) = S'_l h(\tilde{a}_l) = h S_l(\tilde{a}_l) = h(\tilde{x}_l)$, let $\tilde{y}'_l = S'_l(\tilde{b}'_l) = h(\tilde{y}_l)$ and let $\tilde{\delta}'_l = S'_l(\tilde{\sigma}'_l) = h_\#(\tilde{\delta}_l)$ be the path connecting \tilde{x}'_l to \tilde{y}'_l . We have:

- (1') $\tilde{\sigma}'_l$ is C -close to R' and $\tilde{\sigma}'_l \rightarrow P$.
- (4') the length of $\tilde{\delta}'_l \cap A'_c$ goes to infinity with l .
- (5') $\tilde{x}'_l \rightarrow Q_- \in \text{Fix}(\hat{\Psi}_j) \setminus \{T'^{\pm}_c\}$.
- (6') $T'^{-ml}_c \tilde{y}'_l \rightarrow Q_+ \in \text{Fix}(\hat{\Psi}_t) \setminus \{T'^{\pm}_c\}$.

If P is not isolated in $\text{Fix}(\hat{\Psi})$ then Lemma 2.2 implies, after increasing C if necessary, that \tilde{a}'_l and \tilde{b}'_l are C -close to $\text{Fix}(\tilde{g})$ for all sufficiently large l . After replacing \tilde{a}'_l and \tilde{b}'_l with C -close elements of $\text{Fix}(\tilde{g})$, replacing $\tilde{\sigma}'_l$ with the path connecting the new values of \tilde{a}'_l and \tilde{b}'_l , and replacing C by $2C$, properties (1'), (4'), (5') and (6') still hold and each σ'_l is a Nielsen path for g . Since $\tilde{\delta}'_l$ is a lift of σ'_l , Lemma 7.11 implies that for all sufficiently large l , the lift of g that fixes \tilde{x}'_l and \tilde{y}'_l commutes with T'_c and so equals $T'^{d_l}_c \tilde{g}_j$ for some d_l . Since $Q_- \in \text{Fix}(\tilde{g}_j)$ there is a neighborhood of Q_- in Γ' that is disjoint from $\text{Fix}(T'^m_c \tilde{g}_j)$ for all $m \neq 0$. Since $\tilde{x}'_l \rightarrow Q_-$, it follows that $d_l = 0$ and hence that $\tilde{y}'_l \in \text{Fix}(\tilde{g}_j)$ for all sufficiently large l . Since $\text{Fix}(\tilde{g}_j)$ is T'_c -invariant, $T'^{-ml}_c \tilde{y}'_l \in \text{Fix}(\tilde{g}_j)$ and so $Q_+ \in \text{Fix}(\tilde{g}_j)$ as desired.

Suppose then that P is isolated in $\text{Fix}(\hat{\Psi})$. After replacing \tilde{a}'_l and \tilde{b}'_l by their nearest points in \tilde{R}' , we may assume that $\tilde{\sigma}'_l \subset R'$ and that properties (1'), (4'), (5') and (6') still hold. Lemma 2.23 implies that there is a nonlinear edge \tilde{E}' that iterates toward P under the action of \tilde{g} . Denoting $g_{\#}^m(E')$ by ρ_m we have that for all sufficiently large l there exists $m > 0$ such that σ'_l is a subpath of ρ_m . There is a lift $\tilde{\rho}_m$ of ρ_m that contains $\tilde{\delta}'_l$ and so has endpoints $\partial_{\pm} \tilde{\rho}_m$ such that $\partial_- \tilde{\rho}_m \rightarrow Q_-$ and $T'^{-ml}_c \partial_+ \tilde{\rho}_m \rightarrow Q_+$. The former implies that for sufficiently large m , the initial endpoints of $\tilde{\rho}_m \cap A'_c$ and $\tilde{g}_{j\#}(\tilde{\rho}_m) \cap A'_c$ are equal and the latter implies that if $Q_+ \in \text{Fix}(\tilde{g}_j) = \text{Fix}(\hat{\Psi}_j)$ then the terminal endpoints of $\tilde{\rho}_m \cap A'_c$ and $\tilde{g}_{j\#}(\tilde{\rho}_m) \cap A'_c$ are equal. On the other hand, $\tilde{\rho}_m \cap A'_c$ and $\tilde{g}_{j\#}(\tilde{\rho}_m) \cap A'_c$ have different lengths by Lemma 7.13 so we conclude that $Q_+ \notin \text{Fix}(\hat{\Psi}_j)$. \square

8 Abelian subgroups of maximal rank

By Theorem 7.2, all abelian subgroups are realized, up to finite index, as subgroups of some $\mathcal{D}(\phi)$. In this section we describe those ϕ for which $\mathcal{D}(\phi)$ has maximal rank. As usual, ϕ is represented by a CT $f: G \rightarrow G$ with filtration $\emptyset = G_0 \subset G_1 \subset \cdots \subset G_N = G$.

For the simplest example, start with G_2 having one vertex v_1 , two edges E_1 and E_2 and with f defined by $f(E_1) = E_1$ and $f(E_2) = E_2 E_1^{m_1}$ for some $m_1 \in \mathbb{Z}$. For $k = 1, \dots, n-2$, add pairs of linear edges, E_{2k+1} and E_{2k+2} , initiating at a

new common vertex v_{k+1} , terminating at v_1 and satisfying $f(E_j) = E_j E_1^{m_j}$ for distinct m_j . Thus G has $2n - 3$ linear edges and the resulting $\mathcal{D}(\phi)$ has rank $2n - 3$, which is known [7] to be maximal. In this example all edges terminate at the same vertex and there is only one axis, but this is just for simplicity. One could, for example, take the terminal vertex of E_5 equal to v_2 and define $f(E_5) = E_5 w_5$ where w_5 is a closed Nielsen path based at v_2 . Similar modifications can be done to the other edges as well.

Another simple modification is to redefine $f|_{G_2}$ so that G_2 is a single EG stratum with Nielsen path ρ and redefine f on the other edges to be linear with axis represented by ρ . We may view the original example as being built over a Dehn twist of the punctured torus and this modification as being built over a pseudo-Anosov homeomorphism of the punctured torus.

A perhaps more surprising example of a maximal rank abelian subgroup is constructed as follows. Let S be the genus zero surface with four boundary components β_1, \dots, β_4 and let $h: S \rightarrow S$ be a homeomorphism that represents a pseudo-Anosov mapping class and that pointwise fixes each β_m . Let A be an annulus with boundary components α_1 and α_2 and with its central circle labeled α_3 . Define $D_{jk}: A \rightarrow A$ to be the homeomorphism that restricts to a Dehn twist of order j on the subannulus bounded by α_1 and α_3 and to a Dehn twist of order k on the subannulus bounded by α_2 and α_3 . Finally, define $Y = S \cup A / \sim$ where \sim identifies α_m to β_m for $1 \leq m \leq 3$. The homeomorphisms $g_{ijk}: Y \rightarrow Y$ induced by h^i and D_{jk} for $i, j, k \in \mathbb{Z}$ define a rank three abelian subgroup \mathcal{A}' . The fundamental group of Y is a free group of rank three and the image of \mathcal{A}' in $\text{Out}(F_3)$ is an abelian subgroup \mathcal{A} of maximal rank.

We present a slight generalization of this example in terms of relative train track maps as follows.

Example 8.1 Suppose that G is a rank three marked graph with vertices v_1, \dots, v_4 , that $\emptyset = G_0 \subset G_1 \subset \dots \subset G_4 = G$ is a filtration and that $f: G \rightarrow G$ is a relative train track map such that:

- G_1 is a single fixed edge E_1 with both ends attached to v_1 .
- For $m = 2, 3$, H_m is a single edge E_m with terminal endpoint v_1 and initial endpoint v_m ; $f(E_m) = E_m E_1^{d_m}$ where d_2 and d_3 are distinct nonzero integers.
- H_4 is an EG stratum with three edges, one connecting v_4 to v_l for each $l = 1, 2, 3$; for each edge E of H_4 , $f(E)$ is a concatenation of edges in H_4 and Nielsen paths in G_3 . (The Nielsen paths are iterates of E_1 , $E_2 E_1 \bar{E}_2$ and $E_3 E_1 \bar{E}_3$ and their inverses.)

Then f determines an element $\phi \in \text{Out}(F_3)$ such that $\mathcal{D}(\phi)$ has rank three. The example described above using a four times punctured sphere is a special case of this construction. In general, H_4 is not a geometric stratum in the sense of [2].

We think of the strata $H_2 \cup H_3 \cup H_4$ in Example 8.1 as being a single unit added on to the lower filtration element, which in this case is a single circle. To this end we choose notation as follows.

Notation 8.2 Recall from (Filtration) that the core of each filtration element is a filtration element. The *core filtration* $\emptyset = G_0 \subset G_{l_1} \subset G_{l_2} \subset \cdots \subset G_{l_K} = G_N = G$ is defined to be the coarsening of the full filtration obtained by restricting to those elements that are their own cores or equivalently have no valence one vertices. Note that $l_1 = 1$ by (Periodic edges). For each G_{l_i} , let $H_{l_i}^c$ be the i -th stratum of the core filtration. Namely $H_{l_i}^c = \bigcup_{j=l_{i-1}+1}^{l_i} H_j$. The change in Euler characteristic $\chi(G_{l_{i-1}}) - \chi(G_{l_i})$ is denoted $\Delta_i \chi$. We will also use the notation G_{u_i} defined in item (2) of Lemma 8.3.

We also make use of the following notation.

Lemma 8.3 (1) If $H_{l_i}^c$ does not contain any EG stratum then one of the following holds.

- (a) $l_i = l_{i-1} + 1$ and the unique edge in $H_{l_i}^c$ is a fixed loop that is disjoint from $G_{l_{i-1}}$.
- (b) $l_i = l_{i-1} + 1$ and both endpoints of the unique edge in $H_{l_i}^c$ are contained in $G_{l_{i-1}}$.
- (c) $l_i = l_{i-1} + 2$ and the two edges in $H_{l_i}^c$ are nonfixed and have a common initial endpoint that is not in $H_{l_{i-1}}$ and terminal endpoints in $G_{l_{i-1}}$.

In case (a), $\Delta_i \chi = 0$; in cases (b) and (c), $\Delta_i \chi = 1$.

(2) If $H_{l_i}^c$ contains an EG stratum then H_{l_i} is the unique EG stratum in $H_{l_i}^c$ and there exists $l_{i-1} \leq u_i < l_i$ such that both of the following hold.

- (a) For $l_{i-1} < j \leq u_i$, H_j is a single nonfixed edge E_j whose terminal vertex is in $G_{l_{i-1}}$ and whose initial vertex has valence one in G_{u_i} . In particular, G_{u_i} deformation retracts to $H_{l_{i-1}}$ and $\chi(G_{u_i}) = \chi(G_{l_{i-1}})$.
- (b) For $u_i < j < l_i$, H_j is a zero stratum. In other words, the closure of $G_{l_i} \setminus G_{u_i}$ is the extended EG stratum $H_{l_i}^z$.

If some component of $H_{l_i}^c$ is disjoint from G_{u_i} then $H_{l_i}^c = H_{l_i}$ is a component of G_{l_i} and $\Delta_i \chi \geq 1$; otherwise $\Delta_i \chi \geq 2$.

Proof Suppose at first that $H_{l_i}^c$ does not contain any EG stratum and hence does not contain any zero strata. Then H_j is a single edge E_j for each $l_{i-1} + 1 \leq j \leq l_i$ and if some E_j is fixed then either (1)(a) or (1)(b) is satisfied by (Periodic edges). We may therefore assume that each E_j is nonfixed. The terminal endpoint of E_j must have valence at least two in G_{j-1} by [10, Lemma 4.22]. Thus E_j adds a valence one vertex to $G_{l_{i-1}}$ for $j < l_i$, and all such vertices must be endpoints of E_{l_i} . It follows that either (1)(b) or (1)(c) holds. The Euler characteristic statements are obvious.

We now consider the case that $H_{l_i}^c$ contains an EG stratum H_s . Since the core of each filtration element is a filtration element and G_{s-1} does not carry the attracting lamination associated to H_s , G_s is its own core. This proves that H_{l_i} is the unique EG strata in $H_{l_i}^c$. The existence of u_i satisfying (a) and (b) follows from (Zero strata) and [10, Lemma 4.22]. Since f is rotationless, $H_{l_i}^z$ is contained in a single f -invariant component M of G_{l_i} . The lowest stratum in M can not be a zero stratum, so if $M \cap G_{u_i} = \emptyset$ then $M = H_{l_i} = H_{l_i}^c$. By Corollary 3.2.2 of [2], G_{l_i} is not homotopy equivalent to a graph obtained from $G_{l_{i-1}}$ by adding a single edge. This proves that $\Delta_i \chi \geq 1$ if H_{l_i} is a component of G_{l_i} and $\Delta_i \chi \geq 2$ otherwise. \square

Returning now to our examples of maximal rank abelian subgroups, we formalize the class to which Example 8.1 belongs as follows, where the acronym FPS is chosen to remind the reader of the four times punctured sphere.

Notation 8.4 Assume that $H_{l_i}^c$ is a core filtration element that contains an EG stratum and that u_i is as in Lemma 8.3. We say that $H_{l_i}^c$ is a *partial FPS core stratum* if:

- (1) $u_i = l_{i-1} + 2$ and both edges in $G_{u_i} \setminus H_{l_{i-1}}$ are linear.
- (2) $\Delta_i \chi = 2$ and $H_{l_i}^z$ is a tree.
- (3) Each zero stratum in $H_{l_i}^z$ is a single edge.
- (4) For each edge E of $H_{l_i}^z$, $f(E)$ has a complete splitting each of whose terms is either an edge in $H_{l_i}^z$ or a Nielsen path in G_{u_i} .

There is also the option of adding an additional linear edge. In the geometric case this amounts to Dehn twisting on three boundary components of the four times punctured sphere instead of just two. We formalize this as follows.

Notation 8.5 Assume that $H_{l_i}^c$ is a core filtration element that contains an EG stratum and that u_i is as in Lemma 8.3. We say that $H_{l_i}^c$ is a *FPS core stratum* if:

- (1) $u_i = l_{i-1} + 3$ and all three edges in $G_{u_i} \setminus H_{l_{i-1}}$ are linear.

- (2) $\Delta_i \chi = 2$ and $H_{l_i}^z$ is a tree.
- (3) Each zero stratum in $H_{l_i}^z$ is a single edge.
- (4) For each edge E of $H_{l_i}^z$, $f(E)$ has a complete splitting each of whose terms is either an edge in $H_{l_i}^z$ or a Nielsen path in G_{u_i} .

Remark 8.6 The second items of Notation 8.4 and Notation 8.5 imply that $G_u \cap H_{l_i}^z$ is a three point set. In the context of Notation 8.5 this intersection equals the initial endpoints of the linear edges in $H_{l_i}^c$; in the context of Notation 8.4 it equals the union of the initial endpoints of the linear edges in $H_{l_i}^c$ and one vertex in $G_{l_{i-1}}$.

Remark 8.7 We allow zero strata as in the third items of Notation 8.4 and Notation 8.5 because it is not worth modifying the CTs that occur in our proofs to remove them.

Remark 8.8 The core filtration of the CT $f: G \rightarrow G$ of Example 8.1 has two strata. The first is a fixed loop and the second is a partial FPS core stratum. If one replaces the partial FPS core stratum with an FPS core stratum then the resulting map is not a CT because the fixed loop is a component of $\text{Fix}(f)$ that has no fixed outgoing directions in violation of the fact (Periodic edges) that endpoints of fixed edges are principal. On the other hand, adding a second partial FPS core stratum will result in $\mathcal{D}(\phi)$ not having maximal rank.

We can now state the main results of this section.

Proposition 8.9 *Suppose that $\phi \in \text{Out}(F_n)$ is rotationless and that $\mathcal{D}_R(\phi)$ has rank $2n - 3$. Then ϕ is represented by a CT $f: G \rightarrow G$ and filtration $\emptyset = G_0 \subset G_1 \subset \dots \subset G_N = G$ with the following properties.*

- (A) *One of the following holds.*
 - (a) G_{l_1} has rank two and is a single EG stratum.
 - (b) G_1 is a fixed loop and $H_{l_2}^c$ is a single linear edge E_2 which both endpoints in G_1 . In particular, $l_2 = 2$ and G_{l_2} has rank two.
 - (c) G_1 is a fixed circle and $H_{l_2}^c$ is a partial FPS core stratum. In particular, G_{l_2} has rank three.
- (B) *In case (1) above let $m = 1$; for cases (2) and (3) let $m = 2$. Then for all $i > m$,*
 - $H_{l_i}^c$ *is either*
 - (a) *a pair of linear edges with a common initial vertex that is not contained in $G_{l_{i-1}}$ or*
 - (b) *an FPS core stratum.*

There is an analogous result for abelian subgroups of the subgroup IA_n of $\text{Out}(F_n)$ consisting of elements that act trivially in homology.

Proposition 8.10 *Suppose that $\phi \in \text{Out}(F_n)$ is rotationless and that $\mathcal{D}(\phi) \subset \text{IA}_n$ has rank $2n - 4$. Then ϕ is represented by a CT $f: G \rightarrow G$ and filtration $\emptyset = G_0 \subset G_1 \subset \dots \subset G_N = G$ with the following properties.*

- (A) $l_2 = 2$ and G_2 is connected, has rank two and is contained in $\text{Fix}(f)$.
- (B) For $i > 2$, $H_{l_i}^c$ is either
 - (a) a pair of linear edges with homologically trivial axes and with a common initial vertex that is not contained in G_{l_i-1} or
 - (b) an FPS core stratum whose linear edges have homologically trivial axes.

Recall that one uses the QE-splitting of the f -image of edges of G to define almost invariant subgraphs X_1, \dots, X_M of G and that if a_i is a nonnegative integer assigned to X_i then (a_1, \dots, a_M) is admissible (Definition 6.8) if it satisfies certain linear relations involving two or three of the a_i 's. The rank of $\mathcal{D}(\phi)$ is equal (Corollary 7.6) to the rank of the subspace of \mathbb{R}^M generated by the admissible M -tuples for $f: G \rightarrow G$.

After renumbering the X_i 's we may assume that there exist $0 \leq M_1 \leq M_2 \leq \dots \leq M_N = M$ such that X_1, \dots, X_{M_j} is the smallest set of almost invariant subgraphs that contain all the nonfixed strata of G_j . Let R_j be the rank of the subspace of \mathbb{R}^{M_j} generated by the restriction of admissible M -tuples to the first M_j coordinates.

Lemma 8.11 (1) $R_j \geq R_{j-1}$ for all j .

- (2) If H_j is a fixed edge then $R_j = R_{j-1}$.
- (3) If H_j is a linear edge then $R_j \leq R_{j-1} + 1$.
- (4) If H_j is a nonfixed nonlinear NEG edge then $R_j = R_{j-1}$.
- (5) If $H_{l_i}^c$ is a core filtration element that contains an EG stratum then $R_{l_i} \leq R_{u_i} + 1$ with equality holding only if the following condition is satisfied.
 - (a) If σ is either an edge in H_{l_i} or a taken connecting path in a zero stratum of $H_{l_i}^z$ then the terms in the QE-splitting of $f_{\#}(\sigma)$ are either edges in H_{l_i} , taken connecting paths in a zero stratum of $H_{l_i}^z$ or Nielsen paths in G_{u_i} .
 - (b) The almost invariant subgraph X_q that contains $H_{l_i}^z$ is otherwise disjoint from G_{u_i} . Moreover, a_q is not part of any relation that involves only $(a_1, \dots, a_{M_{l_i}})$.

Proof The first item is immediate from the definitions. If H_j is a fixed edge then $M_j = M_{j-1}$. If H_j is a linear edge then $M_j \leq M_{j-1} + 1$. This proves (2) and (3). If H_j is a nonfixed nonlinear NEG edge E_j then at least one term σ in the QE-splitting of $f(E_j)$ is contained in G_{j-1} and is not a Nielsen path. If σ is a single nonfixed edge in an irreducible stratum or a connecting path in a zero stratum then $E_j \subset X_i$ for some $i \leq M_{j-1}$ so $M_j = M_{j-1}$ and (4) follows. If σ is a quasi-exceptional path then, assuming without loss that $M_j = M_{j-1} + 1$, there is a relation involving a_{M_j} and one or two a_i 's with $i \leq M_{j-1}$. Thus a_{M_j} is determined by $(a_1, \dots, a_{M_{j-1}})$ and $R_j = R_{j-1}$. This completes the proof of (4).

If $H_{l_i}^c$ is a core filtration element that contains an EG stratum then $M_{l_i} \leq M_{u_i} + 1$ and hence $R_{l_i} \leq R_{u_i} + 1$. Suppose that $R_{l_i} = R_{u_i} + 1$. Item (b) is an immediate consequence of the definitions and (a) follows from (b) by the argument used to prove (4). \square

Notation 8.12 Suppose that $H_{l_i}^c$ is a core stratum for a CT $f: G \rightarrow G$. Let $\Delta_i R = R_{l_{i-1}} - R_{l_i}$ and let δ_i be the number of components of $G_{l_{i-1}}$ that contain the base point of an edge in $H_{l_i}^c$ that determines a fixed direction.

Corollary 8.13 Suppose that $H_{l_i}^c$ is a core stratum for a CT $f: G \rightarrow G$ and that $H_{l_i}^c$ does not contain an EG stratum. Then

$$\Delta_i R \leq 2\Delta_i \chi - \delta_i$$

and if equality holds then one of the following is satisfied.

- (1) $\Delta_i R = \Delta_i \chi = \delta_i = 0$ and $H_{l_i}^c \subset \text{Fix}(f)$ is a component of G_{l_i} .
- (2) $\Delta_i R = \Delta_i \chi = \delta_i = 1$ and $H_{l_i}^c$ is a single linear edge.
- (3) $\Delta_i R = 2$, $\Delta_i \chi = 1$, $\delta_i = 0$ and $H_{l_i}^c$ is a pair of linear edges with a common initial vertex.

Proof This follows immediately from Lemma 8.3 and Lemma 8.11. \square

The analog of Corollary 8.13 for the case that $H_{l_i}^c$ contains an EG stratum is the main step in the proofs of Proposition 8.9 and Proposition 8.10.

Proposition 8.14 Suppose that $H_{l_i}^c$ is a core stratum for a CT $f: G \rightarrow G$ and that $H_{l_i}^c$ contains an EG stratum. Then

$$\Delta_i R \leq 2\Delta_i \chi - \delta_i$$

and if equality holds then the following are satisfied.

- (1) H_p is a single linear edge for each $l_{i-1} < p \leq u_i$.
- (2) $\Delta_i R = V_L + 1$ where $V_L = u_i - l_{i-1}$ is the number of linear edges in $H_{l_i}^c$.
- (3) If an almost invariant subgraph X_q contains either H_p for some $l_{i-1} < p \leq u_i$ or contains $H_{l_i}^z$, then X_q is otherwise disjoint from G_{u_i} . Moreover a_q is not part of any relation that involves only (a_1, \dots, a_{M_i}) .
- (4) If $\delta_i \leq 1$ then:
 - (a) $H_{l_i}^c$ is an FPS core stratum and $\delta_i = 0$.
 - (b) $H_{l_i}^c$ is a partial FPS core stratum and $\delta_i = 1$.

Proof If some component X of $H_{l_i}^z$ is disjoint from G_{u_i} then Lemma 8.3 implies that $H_{l_i}^c = H_{l_i}$ is a component of G_{l_i} and that $\Delta_i \chi \geq 1$. In this case $\Delta_i R = 1$ and $\delta_i = 0$ so the lemma is clear. We assume for the remainder of the proof that each component of $H_{l_i}^z$ has nonempty intersection with G_{u_i} .

Denote $G_{u_i} \cap H_{l_i}^z$ by \mathcal{V} , the cardinality of \mathcal{V} by V and the number of components in $H_{l_i}^z$ by C_i . Adding a component X of $H_{l_i}^z$ to G_{u_i} and then collapsing a maximal tree in X to a point is the same as identifying all the elements of the nonempty set $\mathcal{V} \cap X$ to a single point and possibly adding some loops. If $\mathcal{V} \cap X$ is a single point then there must be at least one loop because each vertex of G_{l_i} has valence at least two. This proves

$$(1) \quad \Delta_i \chi \geq V - C_i$$

with equality if and only if each component of $H_{l_i}^z$ is a tree (in other words, no loops are added after the elements of $\mathcal{V} \cap X$ are identified) and

$$\Delta_i \chi \geq C_i$$

with equality if and only if each component of $H_{l_i}^z$ is topologically either an arc that intersects G_{u_i} in exactly two points or a loop that intersects G_{u_i} in a single point. Adding these inequalities we get

$$2\Delta_i \chi \geq V$$

with equality if and only if each component of $H_{l_i}^z$ is topologically an arc that intersects G_{u_i} in exactly two points.

On the other hand, there must be at least one illegal turn in H_{l_i} . (If there were no illegal turns in H_{l_i} there would be $m > 0$ so that for any loop $\gamma \subset G_{l_i}$ that intersects H_{l_i} nontrivially, the number of edges of H_{l_i} in $f_{\#}^m(\gamma)$ would be strictly larger than the number of edges of H_{l_i} in γ . This can not be true as one easily sees by considering loops γ_{-m} satisfying $f_{\#}^m(\gamma_{-m}) = \gamma$.) This rules out the possibility

that each component of $H_{l_i}^z$ is topologically an arc that intersects G_{u_i} in exactly two points and we conclude that

$$(2) \quad 2\Delta_i\chi \geq V + 1.$$

For $l_{i-1} < j \leq u_i$, the stratum H_j is a single edge E_j . We write $E_j \in \mathcal{E}_L$ if E_j is linear. The initial endpoints \mathcal{V}_L of the edges in \mathcal{E}_L have valence one in G_{u_i} . We denote the cardinality of \mathcal{V}_L by V_L . Lemma 8.11 implies that

$$\Delta_i R \leq V_L + 1.$$

Note also $V_L \leq V - \delta_i$. Thus

$$(3) \quad \Delta_i R \leq V_L + 1 \leq V + 1 - \delta_i \leq 2\Delta_i\chi - \delta_i$$

which completes the proof of the main inequality.

We assume now that all the inequalities in Equation (3) are equalities. From $V_L + 1 = V + 1 - \delta_i$ it follows that $V - V_L = \delta_i$ which implies item (1). Item (2) follows from $\Delta_i R = V_L + 1$ and implies item (3). We now assume that $\delta_i \leq 1$ and prove that either (4)(a) or (4)(b) holds.

Suppose that $C_i > 1$. Since $V - V_L = \delta_i \leq 1$ there is a component Y of $H_{l_i}^z$ whose intersection with G_{u_i} is contained in \mathcal{V}_L . By (NEG Nielsen paths) each Nielsen path in G_{u_i} with an endpoint in $\mathcal{V}_L \cap Y$ is a closed path and in particular has both endpoints in $\mathcal{V}_L \cap Y$. Choose an edge E in H_{l_i} and $k \geq 1$ so that $f_{\#}^k(E)$ intersects each component of $H_{l_i}^z$. Since $\Delta_i R = V_L + 1$, Lemma 8.11(5) implies (by an obvious induction argument) that the terms in the QE-splitting of $f_{\#}^k(E)$ are either edges in H_{l_i} , connecting paths in zero strata of $H_{l_i}^z$ or Nielsen paths in G_{u_i} . But this contradicts the fact that some maximal subpath of $f_{\#}^k(E)$ in G_{u_i} must have one endpoint in Y and the other in a different component of H_{l_i} . We conclude that $C_i = 1$.

Recall from Lemma 8.3 that $\Delta_i\chi \geq 2$. Combining this with $\Delta_i\chi \geq V - 1$ from Equation (1) and with $2\Delta_i\chi = V + 1$ we see that $\Delta_i\chi = 2$ and $V = 3$. It follows that $\Delta_i R + \delta_i = 2\Delta_i\chi = 4$ and $V_L = V - \delta_i$ is either 2 or 3. Equation (1) implies that $H_{l_i}^z$ is a tree. Lemma 8.11 will complete the proof that $H_{l_i}^c$ is an FPS core stratum when $\delta_i = 0$ and is a partial FPS core stratum when $\delta_i = 1$ once we show that each zero stratum in $H_{l_i}^z$ is an edge.

Topologically (meaning that we ignore valence two vertices whose link in G_{l_i} is contained in $H_{l_i}^z$) there are two possibilities for $H_{l_i}^z$. One is that $H_{l_i}^z$ has one valence three vertex that is disjoint from G_{u_i} and three valence one vertices that are contained in G_{u_i} . The other is that $H_{l_i}^z$ has one valence two vertex and two valence one vertices, all of which are contained in G_{u_i} . In both cases the illegal turn in $H_{l_i}^z$ is based at

the unique vertex with valence greater than one. (Zero strata) implies that each zero stratum in H_i^z is a single edge as desired. \square

Proof of Proposition 8.9 Choose a CT $f: G \rightarrow G$ and filtration $\emptyset = G_0 \subset G_1 \subset \cdots \subset G_N = G$ representing ϕ . Let $\emptyset = G_0 \subset G_{l_1} \subset \cdots \subset G_{l_K} = G$ be the associated core filtration. Recall that $G_{l_1} = G_1$. Let $m = -\chi(G_1)$ and let $d = \sum_{i=2}^K \delta_i$. Corollary 8.13 and Proposition 8.14 implies that

$$\begin{aligned} 2n - 3 - R_1 &= R_N - R_1 = \sum_{i=2}^K \Delta_i R \leq \sum_{i=2}^K (2\Delta_i \chi - \delta_i) \\ &= 2\chi(G_{l_1}) - 2\chi(G) - d = 2n - 2 - 2m - d \end{aligned}$$

which implies that

$$2m + d \leq R_1 + 1$$

with equality if and only if $\Delta_i R = 2\Delta_i \chi - \delta_i$ for each $i \geq 2$.

If G_1 is EG then $R_1 = 1$ and $m \geq 1$ by Lemma 8.3. Thus $m = 1$, $d = 0$ and $\Delta_i R = 2\Delta_i \chi - \delta_i$ with $\delta_i = 0$ for each $i \geq 2$. Corollary 8.13, Proposition 8.14 and (Periodic edges), which (since $d = 0$) implies that case (1) of Corollary 8.13 does not happen, complete the proof.

If G_1 is NEG then $G_1 \subset \text{Fix}(f)$, $R_1 = 0$ and $m = 0$. Thus $d \leq 1$ and (Periodic edges) implies that $d \geq 1$. It follows that $d = 1$ and $\Delta_i R = 2\Delta_i \chi - \delta_i$ for each $i \geq 2$. If $\delta_2 = 1$ then Corollary 8.13, Proposition 8.14 and (Periodic edges) complete the proof.

It remains to show that if G_1 is a fixed loop and $\delta_2 = 0$ then we can modify $f: G \rightarrow G$ and the filtration, without changing G_1 , to arrange that $\delta_2 = 1$. Let v and E_1 be the unique vertex and edge in G_1 and let E_i be an edge in H_i^c that determines a fixed direction at v pointing out of G_1 . Inspecting the possibilities in Corollary 8.13 and Proposition 8.14 we see that H_{l_i} is the only stratum containing an edge that determines a fixed direction at v pointing out of G_1 and that the link $L(G, v)$ consists of E_1, \bar{E}_1 , edges in H_{l_i} and the terminal ends of some linear edges.

As a first case suppose that H_{l_i} is EG. By (NEG Nielsen paths), (EG Nielsen paths) and [10, Remark 4.21] every closed Nielsen path based at v is an multiple of E_1 or its inverse. Thus each linear edge E_k whose terminal endpoint is v satisfies $f(E_k) = E_k E_1^{d_k}$ for some $d_k \neq 0$. Create a new graph G' by replacing v with a pair of vertices v_1 and v_2 , attaching the edges in $L(G, v)$ coming from H_{l_i} to v_2 , attaching all the remaining edges in $L(G, v)$ to v_1 and by adding an oriented edge E' with initial endpoint v_2 and terminal endpoint v_1 . There is an induced map $f': G' \rightarrow G'$ that fixes E' . This process is the inverse of collapsing an edge to a point and it is

straightforward to check that $f': G' \rightarrow G'$ satisfies all of the properties of a CT except for (Periodic edges).

The link $L(G', v_1)$ consists of E_1, \bar{E}_1, \bar{E}' and the terminal ends of linear edges E_k satisfying $f(E_k) = E_k E_1^{d_k}$. Choose an E_k , say E_2 , that is contained in $H_{l_2}^c$. Define a new homotopy equivalence $g: G' \rightarrow G'$ by replacing d_k with $d_k - d_2$ and by replacing $f'(E') = E'$ with $g(E') = E' E_1^{-d_2}$. Note that f and g are freely homotopic (the homotopy can be chosen to have support in a small neighborhood of E_1 and to restrict to a homotopy on E_1 from the identity to rotation by $-2d_2\pi$) and so represent the same element of $\text{Out}(F_n)$. We have changed the fixed edge from E' to E_2 . Finally, modify g by collapsing E_2 to a point. The resulting map is a CT with $\delta_2 = 1$ as desired.

The remaining case is that H_{l_i} is NEG and so is a single linear edge E_{l_i} satisfying $f(E_{l_i}) = E_{l_i} \sigma^d$ for some closed Nielsen path σ and some $d \neq 0$. In this case $E_{l_i} \sigma \bar{E}_{l_i}$ is a closed Nielsen path based at v so there may be linear edges E_k with $f(E_k) = E_k \sigma_k$ where σ_k not an iterate of E_1 or its inverse. To take this into account we attach the terminal end of E_k to v_1 if σ_k is a multiple of E_1 or its inverse and to v_2 otherwise. As above E_1, \bar{E}_1 and \bar{E}' are attached to v_1 and E'_1 is attached to v_2 . The rest of the construction is the same as in the previous case. We leave the details to the reader. \square

Proof of Proposition 8.10 Choose a CT $f: G \rightarrow G$ and filtration $\emptyset = G_0 \subset G_1 \subset \dots \subset G_N = G$ representing ϕ . Let $\emptyset = G_0 \subset G_{l_1} \subset \dots \subset G_{l_K} = G$ be the associated core filtration. Since IA_2 is trivial [16], either there are no ϕ -invariant free factors of rank two or there is a ϕ -invariant free factor of rank two on which ϕ acts trivially. In the latter case we may assume by [10, Theorem, 4.29 and Remark 4.42] that $G_{l_2} = G_2 \subset \text{Fix}(f)$. For $0 \leq j \leq K$, let $d_j = \sum_{i=j}^K \delta_i$ and let $m_j = -\chi(G_{l_j})$. Proposition 8.14 implies that

$$\begin{aligned} 2n - 4 - R_{l_j} &= \sum_{i=j+1}^K \Delta_i R \leq \sum_{i=j+1}^K (2\Delta_i \chi - \delta_i) \\ &= 2\chi(G_{l_j}) - 2\chi(G) - d_{j+1} = 2n - 2 - 2m_j - d_{j+1} \end{aligned}$$

which implies that

$$(4) \quad R_{l_j} \geq 2m_j + d_{j+1} - 2$$

and that if Equation (4) is an equality then $\Delta_i R = 2\Delta_i \chi - \delta_i$ for each $i \geq j + 1$. Equation (4) with $j = 0$ implies that $d_1 \leq 2$.

Suppose at first that $G_{l_2} = G_2 \subset \text{Fix}(f)$. Then $m_2 = 1$ and $R_{l_2} = R_2 = 0$. Equation (4) implies that $d_3 = 0$ and that $\Delta_i R = 2\Delta_i \chi - \delta_i$ with $\delta_i = 0$ for each $i \geq 3$. Corollary 8.13, Proposition 8.14 and (Periodic edges), which implies that case (1) of Corollary 8.13 does not happen, reduce the proof to showing that the axis associated to each linear edge E_j is homologically trivial. From the explicit descriptions given in Corollary 8.13 and Proposition 8.14 we see that E_j does not disconnect G and that the almost invariant subgraph containing E_j contains no other stratum and is not part of any relation. Thus the homotopy equivalence $g: G \rightarrow G$ that fixes $G \setminus E_j$ and satisfies $g|_{E_j} = f|_{E_j}$ represents an element of $\mathcal{D}(\phi)$ and so acts by the identity on homology. Since there are loops in G that cross E exactly once, the axis associated to E_j must be homologically trivial.

We assume now that $G_2 \not\subset \text{Fix}(f)$ and argue to a contradiction.

As noted above, no rank two free factor is ϕ -invariant. Since $R_1 \leq 1$ and $m_1 \neq 1$, Equation (4) implies that $m_1 = 0$. Thus G_1 is a fixed loop, $R_1 = 0$ and $1 \leq d_2 \leq 2$. Suppose that $H_{l_2}^c$ is disjoint from G_1 . Since we can switch the order of H_1 and H_{l_2} in this case, G_{l_2} is a fixed loop, $l_2 = 2$, $m_2 = \Delta_2 R = 0$ and $d_3 = 2$. By the same logic, if $H_{l_3}^c$ is disjoint from $G_1 \cup G_2$ then d_4 would be at least three in contradiction to $d_3 \leq d_1 \leq 2$. It follows that G_{l_3} has at most two components. After switching the order of H_1 and H_2 if necessary, we may assume that either G_{l_2} or G_{l_3} is connected; define $q = 2$ if G_{l_2} is connected and $q = 3$ otherwise. Since G_{l_q} does not have rank two, H_{l_q} is EG. Let V_L be the number of linear edges in G_{u_q} .

If $q = 2$ then $d_{q+1} \geq 1 - \delta_q$ and if $q = 3$ then $d_{q+1} \geq 2 - \delta_q$. Equation (4) therefore implies that

$$(5) \quad \Delta_q R = R_{l_q} \geq 2m_q - \delta_q - 1 \geq 3 - \delta_q$$

if $q = 2$ and

$$\Delta_q R = R_{l_q} \geq 2m_q - \delta_q \geq 4 - \delta_q$$

if $q = 3$. Combining this with Equation (3) we conclude that $R_{l_q} = V_L$ or $R_{l_q} = V_L + 1$ if $q = 2$ and $R_{l_q} = V_L + 1$ if $q = 3$. The proof now divides into cases.

Case 1 ($R_{l_q} = V_L + 1$.) Choose an edge E in H_{l_q} and $k \geq 1$ so that $f_{\#}^k(E)$ crosses every edge in $H_{l_q}^z$. Item (5)(a) of Lemma 8.11 implies (by an obvious induction argument) that the terms in the QE-splitting of $f_{\#}^k(E)$ are either edges in H_{l_q} , connecting paths in a zero stratum of $H_{l_q}^z$ or Nielsen paths in G_{u_i} . Since all Nielsen paths in G_{u_i} are closed paths there is a path in $H_{l_q}^z$ that contains every edge in $H_{l_q}^z$. This proves that $H_{l_q}^z$ is connected.

The number of edges in G_{u_q} that separate G_{l_q} is at most $1 - \delta_q$ if $q = 2$ and at most $2 - \delta_2$ if $q = 3$. We may therefore choose a nonseparating linear edge E_j

in G_{l_q} . By item (5)(b) of Lemma 8.11 there is a homotopy equivalence $g: G \rightarrow G$ that represents an element of $\mathcal{D}(\phi)$ and whose restriction to G_{l_q} fixes $G_{l_q} \setminus E_j$ and satisfies $g|_{E_j} = f|_{E_j}$. Since $g|_{G_{l_q}}$ acts by the identity on homology and since there are loops in G_{l_q} that cross E_j exactly once, the axis associated to E_j must be homologically trivial in contradiction to the fact that the axis of E_j is represented by the basis element E_1 .

Case 2 ($R_{l_q} = V_L$.) In this case we have $q = 2$. Choose an edge E in H_{l_2} and $k \geq 1$ so that $f_{\#}^k(E)$ crosses every edge of $H_{l_2}^z$. If each term σ in the QE-splitting of $f_{\#}^k(E)$ is either an edge in H_{l_2} , a connecting path in a zero stratum of $H_{l_2}^z$ or a Nielsen path in G_{u_2} then $H_{l_2}^z$ is connected by the argument used in case 1. Otherwise some σ is a linear edge or a quasi-exceptional path.

If σ is a linear edge E_i then E_i and $H_{l_2}^z$ belong to the same almost invariant subgraph, the other almost invariant subgraphs that contain strata in G_{l_2} contain a unique stratum in G_{l_2} and there are no relations between any of these almost invariant subgraphs. The terms in the QE-splitting of $f_{\#}^k(E)$ that intersect G_{u_2} are closed paths or E_i or its inverse. It follows that each component of $H_{l_2}^z$ contains either the initial endpoint of E_i or the unique vertex in G_1 . These components can be the same or different.

If σ is a quasi-exceptional path with initial and terminal edges say E_j and E_k then the almost invariant subgraphs that contain strata in G_{l_2} contain a unique stratum in G_{l_2} and the only relation between them is the one determined by σ . The terms in the QE-splitting of $f_{\#}^k(E)$ that intersect G_{u_2} are closed paths or a quasi-exceptional path in the same family as $E_j \bar{E}_k$. It follows that each component of $H_{l_2}^z$ contains the initial endpoint of either E_j or E_k . These components can be the same or different.

Case 2a ($H_{l_2}^z$ is connected.) If $H_{l_2}^z$ is connected then

$$m_2 \geq V_L - 1 + \delta_2.$$

Combining this with Equation (5) and the assumption that $R_{l_q} = V_L$ we have

$$V_L \geq 2m_2 - \delta_2 - 1 \geq 2V_L - 2 + 2\delta_2 - \delta_2 - 1 = 2V_L - 3 + \delta_2$$

which implies that

$$V_L \leq 3 - \delta_2$$

and hence by Equation (5) that

$$V_L = 3 - \delta_2.$$

If $\delta_2 = 0$ then $H_{l_2}^z$ and the three linear edges in G_{u_2} are contained in either three or four almost invariant subgraphs; in the former case there are no relations between them and in the latter case there is one relation between two or three of them. In either case

there exist admissible M -tuples \mathbf{a} and \mathbf{b} such that $\mathbf{a}_\alpha = \mathbf{b}_\alpha$, $\mathbf{a}_\beta = \mathbf{b}_\beta$ and $\mathbf{a}_\gamma \neq \mathbf{b}_\gamma$ where X_α contains $H_{l_2}^z$ and where X_β and X_γ contain linear edges E_β and E_γ in G_{u_2} . Choose a path ρ in $H_{l_2}^z$ connecting the initial vertex of E_β to the initial vertex of E_γ . Then $\sigma = \bar{E}_\beta \rho E_\gamma$ is a loop such that $(f_{\mathbf{a}})_\#(\sigma)$ and $(f_{\mathbf{b}})_\#(\sigma)$ determine different homology classes in contradiction to the assumption that $\mathcal{D}(\phi) \subset \text{IA}_n$.

If $\delta_2 = 1$ then $H_{l_2}^z$ and the two linear edges in G_{u_2} are contained in either two or three almost invariant subgraphs; in the former case there are no relations between them and in the latter case there is one relation between two or three of them. In either case there exist admissible M -tuples \mathbf{a} and \mathbf{b} such that $\mathbf{a}_\alpha = \mathbf{b}_\alpha$ and $\mathbf{a}_\beta \neq \mathbf{b}_\beta$ where X_α contains $H_{l_2}^z$ and where X_β contains a linear edge E_β in G_{l_2} . Choose a path ρ in $H_{l_2}^z$ connecting the initial vertex of E_β to the unique vertex in G_1 . Then $\sigma = \bar{E}_j \rho$ is a loop such that $(f_{\mathbf{a}})_\#(\sigma)$ and $(f_{\mathbf{b}})_\#(\sigma)$ determine different homology classes in contradiction to the assumption that $\mathcal{D}(\phi) \subset \text{IA}_n$.

Case 2b ($H_{l_2}^z$ is not connected.) If $H_{l_2}^z$ is not connected then it has two components. Equation (5) implies that $V_L \geq 3 - \delta_2$.

If $\delta_2 = 0$ then neither component of $H_{l_2}^z$ contains the unique vertex of G_1 so some term in the QE-splitting of $f_\#^k(E)$ is quasi-exceptional with one endpoint in each component of $H_{l_2}^z$. All the linear edges in G_{u_2} and $H_{l_2}^z$ are contained in distinct almost invariant subgraphs and there is a relation between the almost invariant subgraph containing $H_{l_2}^z$ and two of the almost invariant subgraphs containing linear edges whose initial edges are contained in distinct components of $H_{l_2}^z$. In this case the proof concludes as when $\delta_2 = 0$ for case (2)(a) where E_β and E_γ have initial endpoints in the same component of $H_{l_2}^z$.

If $\delta_2 = 1$ then $H_{l_2}^z$ and a linear edges in G_{u_2} with initial endpoint in the component of $H_{l_2}^z$ that does not contain the unique vertex of G_1 belong to the same almost invariant subgraph, the other linear edges are in distinct almost invariant subgraphs and there are no relations between any of these almost invariant subgraphs. In this case the proof concludes as when $\delta_2 = 1$ for case (2)(a) where the initial endpoint of E_β is contained in the component of $H_{l_2}^z$ that contains the unique vertex of G_1 . \square

9 Two families of abelian subgroups

We now return to the simplest examples of maximal rank abelian subgroups: those that are rotationless, have linear growth and have only one axis. We prove that these subgroups and their standard generators can be characterized using only algebraic (as opposed to dynamical systems) properties. These results are needed in the calculation [8] of the commensurator of $\text{Out}(F_n)$.

We begin by relating the rank of $\mathcal{A}(\psi)$ to the dynamical properties of ψ in a special case. Recall that $\mathcal{L}(\psi)$ is the set of attracting laminations for ψ .

Lemma 9.1 *Suppose that A is a maximal rank rotationless abelian subgroup of $\text{Out}(F_n)$ or IA_n , that $\psi \in A$ and that $\mathcal{A}(\psi)$ has rank one. Then either $\mathcal{L}(\psi)$ has exactly one element and ψ has no axes or $\mathcal{L}(\psi) = \emptyset$ and ψ has exactly one axis and that axis has multiplicity one.*

Proof We show below that each $\Lambda \in \mathcal{L}(\psi)$ is minimal, meaning that every line in Λ is dense in Λ , and that if $g: G' \rightarrow G'$ is a CT representing ψ then each nonfixed NEG edge is linear. The former implies that if E' is an edge of an EG stratum H'_t and $k \geq 0$ then the terms in the QE-splitting of $g_{\#}^k(E')$ are edges in H'_t , connecting paths in the zero strata that are enveloped by H'_t and Nielsen paths. It follows that each almost invariant subgraph is a single core stratum and that there are no relations between the almost invariant subgraphs. This implies that the rank of $\mathcal{D}(\psi)$, and hence the rank of $\mathcal{A}(\psi)$, is equal to the number of nonfixed core strata. The lemma follows immediately.

By Lemma 5.4 there exists a rotationless $\phi \in A$ so that $A \subset \mathcal{A}(\phi)$. Choose a CT $f: G \rightarrow G$ and filtration $\emptyset = G_0 \subset G_1 \subset \cdots \subset G_N = G$ representing ϕ . Let $\emptyset = G_0^c \subset G_{I_1}^c \subset \cdots \subset G_{I_K}^c = G$ be the associated core filtration. Proposition 8.9 and Proposition 8.10 imply that if $\Lambda \in \mathcal{L}(\phi)$ corresponds to an EG stratum H_r then both ends of every leaf of Λ intersect H_r infinitely often. By Lemma 3.1.15 of [2] each leaf of Λ is dense in Λ . In other words Λ is minimal. The symmetric argument applied to ϕ^{-1} shows that every element of $\mathcal{L}(\phi) \cup \mathcal{L}(\phi^{-1})$ is minimal. Lemma 5.8 and Remark 4.8 therefore imply that every element of $\mathcal{L}(\psi)$ is minimal.

The proof that nonfixed NEG edges of $g: G' \rightarrow G'$ are linear is less direct. The first step is to show that there does not exist a proper free factor system \mathcal{F} that carries each element of $\mathcal{L}(\phi) \cup \mathcal{L}(\phi^{-1})$ and each ϕ -invariant conjugacy class. We do this by assuming that \mathcal{F} exists and arguing to a contradiction.

Suppose that A is a subgroup of IA_n . Proposition 8.10 implies that there is a ϕ -invariant free factor of rank two on which ϕ acts trivially. This free factor is carried by \mathcal{F} so by Theorem 2.21 and [10, Remark 4.42] we may assume that $\mathcal{F} = \mathcal{F}(G_r)$ for some G_r and that $f|_{G_r}$ is the identity. As shown in the proof of Proposition 8.10, $f: G \rightarrow G$ satisfies the conclusions of Proposition 8.10. In particular, $H_{I_K}^c$ either contains an EG stratum or is a pair of linear edges with a common initial vertex not in G_r . In the former case there is an element of $\mathcal{L}(\phi)$ not carried by G_r and in the latter case there is a ϕ -invariant conjugacy class not carried by G_r .

In the case that A is not a subgroup of IA_n , we may still assume that $\mathcal{F} = \mathcal{F}(G_r)$ for some G_r . As shown in the proof of Proposition 8.9, there is one additional

possibility for $H_{I_k}^c$. Namely, $H_{I_k}^c$ can be a single linear edge with initial base point in $G_1 \subset \text{Fix}(f)$. But this also contradicts the assumption that G_r carries every ϕ -invariant conjugacy class. This completes the proof that \mathcal{F} as above does not exist.

After replacing ψ with an iterate if necessary we may assume that $\psi \in \mathcal{D}(\phi)$. Lemma 6.13 and Corollary 6.20 imply that each ϕ -invariant conjugacy class is ψ -invariant.

We can now complete the proof by assuming that there is a nonfixed nonlinear NEG edge for $g: G' \rightarrow G'$ and arguing to a contradiction. The highest such edge E_j can not be a term in the QE-splitting of $f^k(E)$ for any edge E in a linear or EG stratum. This is obvious for linear strata and follows from the minimality of $\Lambda \in \mathcal{L}(\phi)$ for EG strata. In conjunction with (Zero strata), this proves that E_j is not in the image of any edge above it. We may therefore assume that E_j is the top stratum. But then G_{j-1} carries each element of $\mathcal{L}(\phi) \cup \mathcal{L}(\phi^{-1})$ and each ϕ -invariant conjugacy class. This contradiction completes the proof. \square

Let G be the rose with $n - 2$ of its n edges subdivided into two edges. Thus there are edges E_1, \dots, E_{2n-2} and vertices v_1, \dots, v_{n-1} with v_1 the terminal vertex of all edges and the initial edges of E_1 and E_2 and with v_k the initial vertex of E_{2k-1} and E_{2k} for $2 \leq k \leq n - 1$.

For $1 \leq i \leq 2n - 3$, define $f_i: G \rightarrow G$ by $E_{i+1} \mapsto E_{i+1}E_1$ and all other edges fixed. Choose a basis x_1, \dots, x_n for F_n and a marking on G that identifies x_j with the j -th loop of G . The elements $\eta_i \in \text{Out}(F_n)$ determined by f_i are a basis for an abelian subgroup A_1 of rank $2n - 3$. If $i = 2k - 2$ for $k \geq 2$ then $\hat{\eta}_i$ is defined by $x_k \mapsto x_k x_1$. If $i = 2k - 1$ for $k \geq 2$ then $\hat{\eta}_i$ is defined by $x_k \mapsto \bar{x}_1 x_k$. The remaining element $\hat{\eta}_1$ is defined by $x_2 \mapsto x_2 x_1$. Borrowing notation from [8] we say that A_1 is the *type E subgroup associated to the basis* x_1, \dots, x_n and that $\eta_1, \dots, \eta_{2n-3}$ are its *standard generators*.

Remark 9.2 It is not hard to check (see for example [8, Lemma 2.14]) that η_1 is conjugate to each η_j and to $\eta_j \eta_l$ if $\{j, l\} \neq \{2k - 2, 2k - 1\}$ for some $k \geq 2$. Corollary 5.6 and [8, Lemma 4.4] imply that $\mathcal{A}(\eta_1)$ has rank one. This explains the hypothesis in the next lemma.

Lemma 9.3 *Suppose that $\phi_1, \dots, \phi_{2n-3}$ form a rotationless basis for an abelian subgroup A of $\text{Out}(F_n)$, $n \geq 3$, that each $\mathcal{A}(\phi_j)$ has rank one and that $\mathcal{A}(\phi_j \phi_l)$ has rank one if $\{j, l\} \neq \{2k - 2, 2k - 1\}$ for some $k \geq 2$. Then there is a basis x_1, \dots, x_n for F_n , standard generators η_j of the type E subgroup associated to this basis and $t > 0$ so that $\phi_j = \eta_j^t$ for all j .*

Proof Corollary 3.13 implies that A is rotationless. By Lemma 5.4 there exists $\theta \in A$ such that each $A \subset \mathcal{A}(\theta)$. Choose a CT $f: G \rightarrow G$ representing θ as in Proposition 8.9. In particular, each nonfixed NEG edge is linear. The coordinates of $\Omega^\theta: \mathcal{A}(\theta) \rightarrow \mathbb{Z}^{2n-3}$ (Definition 7.3) are in one to one correspondence with the linear edges and EG strata.

By hypothesis, for each ϕ_j there exists $\phi_{j'} \neq \phi_j$ such that $\mathcal{A}(\phi_j \phi_{j'})$ has rank one.

Suppose that $\psi \in \mathcal{A}(\theta)$, that ω_i is a coordinate of Ω^θ and that $\omega_i(\psi) \neq 0$. If $\omega_i = \text{PF}_\Lambda$ then $\Lambda \in \mathcal{L}(\psi) \cup \mathcal{L}(\psi^{-1})$ by Remark 4.8. If ω_i corresponds to a linear edge with associated axis $[c]_u$ then $[c]_u$ is an axis for ψ ; if ω_r also corresponds to a linear edge with associated axis $[c]_u$ and if $\omega_r(\psi) \neq \omega_i(\psi)$ and $\omega_r(\psi) \neq 0$ then $[c]_u$ is an axis for ψ with multiplicity greater than one. Lemma 9.1 therefore implies that for each ϕ_j the coordinates of $\Omega^\theta(\phi_j)$ takes on a single nonzero value and that if more than one coordinate takes this value then all such coordinates come from linear edges associated to the same axis. The same holds true for the coordinates of $\Omega^\theta(\phi_j \phi_{j'})$.

Suppose that $\omega_i = \text{PF}_\Lambda$ and that $\omega_i(\phi_j) \neq 0$. At least one of $\omega_i(\phi_{j'})$ or $\omega_i(\phi_j \phi_{j'})$ is nonzero, say $\omega_i(\phi_{j'})$. Then $\Omega^\theta(\phi_j)$ and $\Omega^\theta(\phi_{j'})$ are contained in a cyclic subgroup of \mathbb{Z}^{2n-3} in contradiction to the fact that ϕ_j and $\phi_{j'}$ generate a rank two subgroup and the injectivity of Ω^θ . We conclude that each coordinate of Ω^θ corresponds to a linear edge of $f: G \rightarrow G$.

Since $\Omega^\theta(\phi_j)$ and $\Omega^\theta(\phi_{j'})$ are not contained in a cyclic subgroup of \mathbb{Z}^{2n-3} , the coordinates on which they are nonzero can not be identical. It follows that these coordinates are disjoint and correspond to the same axis of θ ; moreover, the unique nonzero values taken by $\Omega^\theta(\phi_j)$ and $\Omega^\theta(\phi_{j'})$ are the same. For each i this applies to all but one j . It follows that all linear edges correspond to the same axis, that only one coordinate of $\Omega^\theta(\phi_j)$ can be nonzero and that the nonzero value t that is taken is independent of j . The lemma now follows from the explicit description of $f: G \rightarrow G$ given by Proposition 8.9 and the definition of Ω^θ . \square

There is an analogous result for IA_n . For the model subgroup, we use the same marked graph G as in the definition of type E subgroups. Choose a closed path in G_2 based at v_1 that forms a circuit and determines a trivial element of homology. For $1 \leq i \leq 2n-4$ define $f_i: G \rightarrow G$ by $E_{i+2} \mapsto E_{i+2}w$. The elements $\mu_{i,w} \in \text{Out}(F_n)$ determined by f_i are a basis for an abelian subgroup A_w of IA_n with rank $2n-4$. We think of w as both a path in G_2 and an element of the free factor $\langle x_1, x_2 \rangle$. If $i = 2k-5$ then $\hat{\mu}_{i,w}$ is defined by $x_k \mapsto x_k w$ and if $i = 2k-4$ then $\hat{\eta}_i$ is defined by $x_k \mapsto \bar{w}x_k$. Borrowing notation from [8] we say that A_w is the *type C subgroup associated to w and to the basis x_1, \dots, x_n* and that $\mu_{1,w}, \dots, \mu_{2n-4,w}$ are its *standard generators*.

Lemma 9.4 Suppose that $\phi_1, \dots, \phi_{2n-4}$ are a rotationless basis for an abelian subgroup of IA_n , $n \geq 4$, that each $\mathcal{A}(\phi_j)$ has rank one and that $\mathcal{A}(\phi_j \phi_l)$ has rank one if $\{j, l\} \neq \{2k, 2k + 1\}$. Then there exists a basis x_1, \dots, x_n for F_n , a homologically trivial element $w \in \langle x_1, x_2 \rangle$ and standard generators η_j of the type C subgroup associated to w and this basis, and $t > 0$ so that $\phi_j = \eta_j^t$

Proof We have assumed that $n \geq 4$ so that for all ϕ_j there exists $\phi_{j'}$ such that $\mathcal{A}(\phi_j \phi_{j'})$ has rank one. Otherwise the proof of Lemma 9.3 carries over to this context without modification, w representing the unique axis of the elements ϕ_j . \square

References

- [1] **H Bass, A Lubotzky**, *Linear-central filtrations on groups*, from: “The mathematical legacy of Wilhelm Magnus: groups, geometry and special functions (Brooklyn, NY, 1992)”, (W Abikoff, JS Birman, K Kuiken, editors), Contemp. Math. 169, Amer. Math. Soc. (1994) 45–98 MR1292897
- [2] **M Bestvina, M Feighn, M Handel**, *The Tits alternative for $\text{Out}(F_n)$. I. Dynamics of exponentially-growing automorphisms*, Ann. of Math. (2) 151 (2000) 517–623 MR1765705
- [3] **M Bestvina, M Feighn, M Handel**, *Solvable subgroups of $\text{Out}(F_n)$ are virtually Abelian*, Geom. Dedicata 104 (2004) 71–96 MR2043955
- [4] **M Bestvina, M Handel**, *Train tracks and automorphisms of free groups*, Ann. of Math. (2) 135 (1992) 1–51 MR1147956
- [5] **D Cooper**, *Automorphisms of free groups have finitely generated fixed point sets*, J. Algebra 111 (1987) 453–456 MR916179
- [6] **M Culler**, *Finite groups of outer automorphisms of a free group*, from: “Contributions to group theory”, (KI Appel, JG Ratcliffe, PE Schupp, editors), Contemp. Math. 33, Amer. Math. Soc. (1984) 197–207 MR767107
- [7] **M Culler, K Vogtmann**, *Moduli of graphs and automorphisms of free groups*, Invent. Math. 84 (1986) 91–119 MR830040
- [8] **B Farb, M Handel**, *Commensurations of $\text{Out}(F_n)$* , Publ. Math. Inst. Hautes Études Sci. (2007) 1–48 MR2354204
- [9] **A Fathi, L Laudenbach, V Poénaru, editors**, *Travaux de Thurston sur les surfaces*, Astérisque 66, Soc. Math. France, Paris (1979) MR568308 Séminaire Orsay, with an English summary
- [10] **M Feighn, M Handel**, *The recognition theorem for $\text{Out}(F_n)$* , Preprint
- [11] **J Franks, M Handel, K Parwani**, *Fixed points of abelian actions*, J. Mod. Dyn. 1 (2007) 443–464 MR2318498

- [12] **J Franks, M Handel, K Parwani**, *Fixed points of abelian actions on S^2* , Ergodic Theory Dynam. Systems 27 (2007) 1557–1581 MR2358978
- [13] **D Gaboriau, A Jaeger, G Levitt, M Lustig**, *An index for counting fixed points of automorphisms of free groups*, Duke Math. J. 93 (1998) 425–452 MR1626723
- [14] **G Levitt, M Lustig**, *Most automorphisms of a hyperbolic group have very simple dynamics*, Ann. Sci. École Norm. Sup. (4) 33 (2000) 507–517 MR1832822
- [15] **G Levitt, M Lustig**, *Automorphisms of free groups have asymptotically periodic dynamics*, J. Reine Angew. Math. 619 (2008) 1–36 MR2414945
- [16] **J Nielsen**, *Die Isomorphismengruppe der freien Gruppen*, Math. Ann. 91 (1924) 169–209 MR1512188
- [17] **W P Thurston**, *On the geometry and dynamics of diffeomorphisms of surfaces*, Bull. Amer. Math. Soc. (N.S.) 19 (1988) 417–431 MR956596

*Math Department, Rutgers University
Newark, NJ 07102, USA*

*Math Department, Lehman College
Bronx, NY 10468, USA*

`feighn@newark.rutgers.edu`, `michael.handel@lehman.cuny.edu`

Proposed: Benson Farb

Seconded: Martin Bridson, Joan Birman

Received: 21 February 2007

Revised: 25 February 2009

