Geometry & Topology

Volume 29 (2025)

The conjugacy problem for UPG elements of $\mathrm{Out}(F_n)$

MARK FEIGHN
MICHAEL HANDEL

msp

Geometry & Topology **29**:4 (2025) 1693–1817
DOI: 10.2140/gt.2025.29.1693
Published: 27 June 2025

The conjugacy problem for UPG elements of Out(F_n)

MARK FEIGHN
MICHAEL HANDEL

An element ϕ of the outer automorphism group $\mathsf{Out}(F_n)$ of the rank n free group F_n is *polynomially growing* if the word lengths of conjugacy classes in F_n grow at most polynomially under iteration by ϕ. It is *unipotent* if, additionally, its action on the first homology of F_n with integer coefficients is unipotent. In particular, if ϕ is polynomially growing and acts trivially on first homology with coefficients the integers mod 3, then ϕ is unipotent and also every polynomially growing element has a positive power that is unipotent. We solve the conjugacy problem in $\mathsf{Out}(F_n)$ for the subset of unipotent elements. Specifically, there is an algorithm that decides if two such are conjugate in $\mathsf{Out}(F_n)$.

20F65, 57M07

1 Introduction

In this paper, we consider the conjugacy problem for $\mathsf{Out}(F_n)$, the group of outer automorphisms of the free group of rank n. Namely, given $\phi, \psi \in \mathsf{Out}(F_n)$, find an algorithm that decides if ϕ and ψ are conjugate in $\mathsf{Out}(F_n)$.

The case in which ϕ is fully irreducible, also known as iwip, was first solved by Sela [1995] using his solution to the isomorphism problem for torsion-free word hyperbolic groups. This was recently generalized, using a similar approach, by Dahmani [2016] to the case that ϕ is hyperbolic, or equivalently, that every nontrivial element of F_n has exponential growth under iteration by ϕ. See also [Dahmani 2017]. An alternate approach to the fully irreducible case takes advantage of the fact that the finite set of (unmarked) train track maps that represent a fully irreducible ϕ is a complete invariant for the conjugacy class of ϕ. Los [1996] and Lustig [2007] (see also [Handel and Mosher 2011]) solved the conjugacy problem for fully irreducible ϕ by algorithmically constructing the set of (unmarked) train track maps for ϕ.

On the other end of the growth spectrum, the conjugacy problem for Dehn twists (equivalently rotationless, linearly growing ϕ) was solved by Cohen and Lustig [1999] using, among other things, Whitehead's algorithm (see below). Krstić, Lustig and Vogtmann [Krstić et al. 2001] proved an equivariant Whitehead algorithm and used that to solve the conjugacy problem for all elements with linear growth.

Building on the approach of Sela mentioned above, Dahmani and Touikan [2021] reduce the conjugacy problem for $\mathsf{Out}(F_n)$ to a list of problems about mapping tori of polynomial growing elements. This is applied in their solution to the conjugacy problem for outer automorphisms of free groups whose polynomially growing part is unipotent linear [Dahmani and Touikan 2023].

Dahmani, Francaviglia, Martino and Touikan [Dahmani et al. 2025] solve the conjugacy problem for $\mathsf{Out}(F_3)$.

Lustig [2000; 2001] posted preprints addressing the general case of the conjugacy problem but these have never been published.

Our main theorem addresses the case that ϕ is polynomially growing and rotationless, equivalently ϕ is polynomially growing and induces a unipotent action on $H_1(F_n, \mathbb{Z})$; we write $\phi \in \mathsf{UPG}(F_n)$. Being an element of $\mathsf{UPG}(F_n)$ is a conjugacy invariant and can be checked algorithmically.

It is often the case, when studying $\mathsf{Out}(F_n)$, that the techniques required to treat the $\mathsf{UPG}(F_n)$ case are very different from those needed for the cases in which there is exponential growth. For example, the polynomially growing and exponentially growing cases of the Tits alternative for $\mathsf{Out}(F_n)$ are proved in separate papers; see Bestvina, Feighn and Handel [2005; 2000].

Theorem 1.1 *There is an algorithm that takes as input $\phi, \psi \in \mathsf{UPG}(F_n)$ and outputs* YES *or* NO *depending on whether or not there exists $\theta \in \mathsf{Out}(F_n)$ such that $\phi = \psi^\theta := \theta\psi\theta^{-1}$. Further, if* YES *then the algorithm also outputs such a θ.*

Remark 1.2 If one knows that ϕ and ψ are conjugate, then a conjugator θ can be produced by searching a list of the elements of $\mathsf{Out}(F_n)$. This is not what we do. Rather, the construction of a conjugator, when one exists, is an integral part of the proof of the main statement of Theorem 1.1.

Remark 1.3 Theorem 1.1 is not an abstract existence theorem. It is proved by constructing an explicit algorithm satisfying the conclusions of the theorem. The same is true for other results in this paper that begin with, "There is an algorithm".

A detailed description of the algorithm is given in Section 2 so we restrict ourselves here to four results/observations that underlie our proof.

- Each $\phi \in \mathsf{UPG}$ is rotationless (Lemma 3.18) and so can be represented by a particularly nice relative train track map $f : G \to G$ call a CT; see Section 3.6. There is an algorithm (Theorem 3.20) to construct one such $f : G \to G$ and from this we can compute all of the invariants used in this paper.

- A set equipped with an action by a group G is a *G-set*. A G-set X *satisfies property* W (for Whitehead) if it comes equipped with an algorithm that takes as input $x, y \in X$ and outputs `YES` or `NO` depending on whether or not there exists $\theta \in G$ such that $\theta(x) = y$ together with such a θ if `YES`. We call such an algorithm a W-*algorithm.* The Whitehead/Gersten algorithm is a W-algorithm for the $\mathsf{Out}(F_n)$-set of finite lists of conjugacy classes of finitely generated subgroups of F_n; see [Gersten 1984, Theorems W&M], and also [Kalajdžievski 1992] and [Bestvina et al. 2023].

This can be applied directly to our problem by finding subgroups associated to elements of UPG. For example, there is a free factor system $\mathcal{F}_0(\phi)$ characterized by the fact that a conjugacy class in F_n is carried by $\mathcal{F}_0(\phi)$ if and only if it grows linearly under iteration by ϕ. Since a free factor system is an unordered list of conjugacy classes of free factors, we can check if there exists $\theta \in \mathsf{Out}(F_n)$ such that $\mathcal{F}_0(\psi) = \theta(\mathcal{F}_0(\phi))$. If no such θ exists then ϕ and ψ are not conjugate. If there is such a θ then after replacing ψ by $\psi^{\theta^{-1}}$, we may assume, as far as the conjugacy problem is concerned, that $\mathcal{F}_0(\phi) = \mathcal{F}_0(\psi)$. Moreover, any conjugator will preserve $\mathcal{F}_0(\phi) = \mathcal{F}_0(\psi)$.

In Sections 10–15 we show that the Whitehead/Gersten algorithm can be used as the platform on which to build other useful $\mathsf{Out}(F_n)$-sets that satisfy property W. The $\mathsf{Out}(F_n)$-set of finite lists of finitely generated subgroups of F_n also satisfies property M (for McCool). Namely, it is equipped with an algorithm that takes as input $x \in X$ and outputs a finite presentation for $G_x := \{\theta \in G \mid \theta(x) = x\}$. Although it is not strictly necessary for solving the conjugacy problem, property M is important in its own right and we show that all of the $\mathsf{Out}(F_n)$-sets constructed in Sections 10–15 satisfy property M.

- Lemma 4.21, an adaptation of the recognition theorem [Feighn and Handel 2011, Theorem 5.3], gives necessary and sufficient conditions for $\theta \in \mathsf{Out}(F_n)$ to conjugate $\phi \in \mathsf{UPG}$ to $\psi \in \mathsf{UPG}$. The nonnumerical condition is that $\theta(\mathcal{L}(\phi)) = \mathcal{L}(\psi)$, where $\mathcal{L}(\phi)$ is a certain set of lines associated to ϕ and similarly for $\mathcal{L}(\psi)$. If $f : G \to G$ is a CT representing ϕ then $\mathcal{L}(\phi)$ is the set of lines carried by a finite type Stallings graph $\Gamma(f)$ called the *eigengraph* for f. $\Gamma(f)$ depends on f but the set of lines carried by $\Gamma(f)$ depends only on ϕ. The numerical condition of Lemma 4.21 concerns the "twist coordinates" associated to the linear parts of ϕ and ψ and is relatively easy to handle; see Lemmas 17.1 and 17.8. Almost all of the paper is concerned with the existence or not of θ satisfying $\theta(\mathcal{L}(\phi)) = \mathcal{L}(\psi)$.

- A CT $f : G \to G$ comes equipped with a filtration $G_{i_0} \subset G_{i_1} \subset \cdots \subset G_{i_t}$, where for $j > 0$, each G_{i_j} is an f-invariant core subgraph which is obtained from $G_{i_{j-1}}$ by adding a single topological arc, possibly divided into two edges. Edges of $G_{i_j} \setminus G_{i_{j-1}}$ are said to have height j. $\Gamma(f)$ has a compact core to which finitely many rays $\{R_E\}$ are added, one for each nonfixed nonlinear edge E of G. Understanding the structure of rays is an important step in understanding $\mathcal{L}(\phi)$. Each R_E has initial edge E and $R_E \setminus E$ is a ray that crosses only edges with height strictly less than that of E. (This is most definitely a UPG phenomenon. If E belongs to an exponentially growing stratum then E occurs infinitely often in R_E.) Thus R_E can be studied inductively, working up through the filtration. This is carried out in Section 5 and Sections 15–17.

Example 3.1 gives an illustrative element of $\mathsf{UPG}(F_n)$ and is further developed as we progress through the text.

Acknowledgements We are indebted to a referee for an exceptionally thorough and useful report.

Feighn was supported by the National Science Foundation under grant DMS-1406167 and also under grant DMS-14401040 while in residence at the Mathematical Sciences Research Institute in Berkeley, California, during the Fall 2016 semester. Handel was supported by National Science Foundation grant DMS-1308710 and by PSC-CUNY grants in Program Years 47 and 49.

2 The algorithm

The logical structure of our proof of Theorem 1.1 is a series of reductions

$$\text{Theorem 1.1} \Longleftarrow \text{Proposition 14.7} \Longleftarrow \text{Proposition 16.4} \Longleftarrow \text{Proposition 17.2},$$

and a proof of Proposition 17.2. The above theorem and propositions produce algorithms that we denote by $\mathrm{ALG}_{1.1}$, $\mathrm{ALG}_{14.7}$, $\mathrm{ALG}_{16.4}$ and $\mathrm{ALG}_{17.2}$, respectively. The proof of Theorem 1.1 $\Longleftarrow$ Proposition 14.7 shows how to use $\mathrm{ALG}_{14.7}$ to construct $\mathrm{ALG}_{1.1}$, and similarly for the other implications. Thus $\mathrm{ALG}_{1.1}$ calls $\mathrm{ALG}_{14.7}$, which calls $\mathrm{ALG}_{16.4}$, which calls $\mathrm{ALG}_{17.2}$.

2.1 Theorem 1.1 $\Longleftarrow$ Proposition 14.7

One way to make progress on the conjugacy problem for UPG is to find *W-invariants* for UPG; ie find $\mathsf{Out}(F_n)$-equivariant maps $J : \mathsf{UPG} \to X$, where X is an $\mathsf{Out}(F_n)$-set with a W-algorithm $W_X(\cdot,\cdot)$. If $W_X(J(\phi), J(\psi)) = \mathtt{NO}$, then ϕ is not conjugate to ψ in $\mathsf{Out}(F_n)$. If $W_X(J(\phi), J(\psi)) = (\mathtt{YES}, \xi)$, then $J(\psi^{\xi^{-1}}) = J(\phi)$. Replacing ψ by $\psi^{\xi^{-1}}$, we may assume that $J(\phi) = J(\psi)$. In this case, any θ conjugating ϕ to ψ is contained in the subgroup of $\mathsf{Out}(F_n)$ that fixes $J(\phi)$.

In Sections 5–14 we construct seven such W-invariants and bundle them into a single invariant $\mathsf{I}_{\mathfrak{c}}(\phi)$. Once this is done, it is easy to use an algorithm satisfying the conclusions of Proposition 14.7 to produce an algorithm satisfying the conclusions of Theorem 1.1. The details are given in the proof of Lemma 14.8.

Items (1)–(4) below outline how our ultimate W-invariant $\mathsf{I}_{\mathfrak{c}} \colon \mathsf{UPG} \to \overline{\mathsf{IS}}(\mathbb{A}_\bullet)$ is chosen. Item (5) refers to shrinking the set of potential conjugators from the stabilizer of $\mathsf{I}_{\mathfrak{c}}(\phi)$ to one of its finite-index subgroups $\mathcal{X}_{\mathfrak{c}}(\phi)$.

(1) **Dynamical invariants of $\phi \in \mathsf{UPG}$**
- The finite multiset $\mathsf{Fix}(\phi)$ of conjugacy classes of *fixed subgroups of* ϕ; see Definition 3.14.
- The *linear free factor system* $\mathcal{F}_0(\phi)$; see Definition 6.5.
- The finite set $\{\mathfrak{c}\}$ of *special* ϕ*-chains*; see Section 6.1 and in particular Notation 6.8.
- The finite set $\mathsf{A}_{\mathrm{or}}(\phi)$ of *axes* for ϕ; see Section 4.2.
- The finite set $\mathsf{SA}(\phi)$ of *strong axes* for ϕ; see Section 4.2.
- The finite set $\Omega_{\mathsf{NP}}(\phi)$ of all nonperiodic *limit lines* for all eigenrays of ϕ; see Section 5.
- For each one-edge extension $\mathfrak{e}$ of each $\mathfrak{c}$, the set $\mathsf{L}_{\mathfrak{e}}(\phi)$ of *added lines with respect to* $\mathfrak{e}$; see Definition 6.14.

The invariants in the first four items are *algebraic* in that they take values in $\mathsf{Out}(F_n)$-sets that can be expressed in terms of conjugacy classes of finitely generated subgroups of F_n or more generally are *iterated sets* (Section 10.1). In particular, they take values in $\mathsf{Out}(F_n)$-sets with W-algorithms and so can be used as they are. The others must be modified.

(2) **Algebraic versions of dynamical invariants** For the last three dynamical invariants, define corresponding (but weaker) algebraic invariants. The last two depend on a choice of special chain $\mathfrak{c}$; see Section 13.1.
- The finite set of *algebraic strong axes*; see Section 13.6.
- The finite set $\{\mathsf{H}_{\mathfrak{c}}(L) \mid L \in \Omega_{\mathsf{NP}}(\phi)\}$ of *algebraic limit lines*; see Section 13.8.
- For each one-edge extension $\mathfrak{e}$ in $\mathfrak{c}$, the finite set $\mathsf{H}_{\mathfrak{e}\in\mathfrak{c}}(\phi)$ of *algebraic added lines with respect to* $\mathfrak{e}$; see Section 13.7.

Remark 2.1 If the seven dynamical invariants in (1) take the same values on ϕ and ψ then, using Lemma 4.21, it is easy to check if ϕ and ψ are conjugate. The same is not true for the seven algebraic invariants in (1) and (2). Too much information was lost in translation.

(3) **W-invariants** Iterated sets, and in particular $\overline{\mathsf{IS}}(\mathbb{A}_\bullet)$, are defined in Sections 10 and 11. By construction, all of our algebraic invariants take values in the iterated set $\overline{\mathsf{IS}}(\mathbb{A}_\bullet)$. We construct a W-algorithm for $\overline{\mathsf{IS}}(\mathbb{A}_\bullet)$ (and all other iterated sets).

(4) **The total invariant $\mathsf{I}_{\mathfrak{c}}(\phi)$** This is defined by combining the algebraic invariants in (1) and (2) into a single algebraic invariant that takes values in $\overline{\mathsf{IS}}(\mathbb{A}_\bullet)$; see Definition 13.13.

(5) **Reduce potential conjugators** Elements of $\mathcal{X}_{\mathfrak{c}}(\phi) < \mathsf{Out}(F_n)$ not only stabilize the algebraic invariants in (1) and (2) making up $\mathsf{I}_{\mathfrak{c}}(\phi)$, they also induce trivial permutations on those invariants that are finite sets; see Definition 14.1.

As mentioned above Lemma 14.8 is proved by constructing $\mathrm{ALG}_{1.1}$ using $\mathrm{ALG}_{14.7}$ and properties of $\mathsf{I}_{\mathfrak{c}}(\phi)$. Hence to prove Theorem 1.1, we are reduced to proving:

Proposition 14.7 *There is an algorithm that takes as input $\phi, \psi \in \mathsf{UPG}(F_n)$ and a chain $\mathfrak{c}$ such that*

- *$\mathfrak{c}$ is special for both ϕ and ψ, and*
- *$\mathsf{I}_{\mathfrak{c}}(\phi) = \mathsf{I}_{\mathfrak{c}}(\psi)$,*

and that outputs `YES` *or* `NO` *depending whether or not there is $\theta \in \mathcal{X}_{\mathfrak{c}}(\phi)$ conjugating ϕ to ψ. Further, if* `YES` *then such a θ is produced.*

2.2 Proposition 14.7 $\Longleftarrow$ Proposition 16.4

$\mathrm{ALG}_{14.7}$ and $\mathrm{ALG}_{16.4}$ differ only in the subgroup of potential conjugators that must be considered. In Proposition 14.7 it is $\mathcal{X}_{\mathfrak{c}}(\phi)$ and in Proposition 16.4 it is an infinite-index subgroup $\mathsf{Ker}(\bar{Q}^{\phi}) < \mathcal{X}_{\mathfrak{c}}(\phi)$ defined in Definition 16.3. See statement of Proposition 14.7 below.

The set of (eigen)rays $\mathcal{R}(\phi)$ (Definition 3.14) is a fundamental dynamical invariant of ϕ. Each $r \in \mathcal{R}(\phi)$ is the conjugacy class $[\tilde{r}]$ of a point $\tilde{r} \in \partial F_n$. There is no W-algorithm for ∂F_n so we work with a weaker algebraic invariant, the conjugacy class $F_{\mathfrak{c}}(r)$ of a free factor determined by r and a special chain $\mathfrak{c}$; see Section 13.4. We do not list this in (2) because it is built into the set of algebraic lines and the set of algebraic added lines. The great advantage of $\mathsf{Ker}(\bar{Q}^{\phi})$ over $\mathcal{X}_{\mathfrak{c}}(\phi)$ is that in the proof of Proposition 17.2 we need only consider conjugating elements that preserve r. (See Lemmas 15.45 and 17.9.) Instead of having to check if two rays are conjugate, we need only check if they are equal.

The definition of $\bar{Q}^{\phi}(\xi)$ for $\xi \in \mathcal{X}_{\mathfrak{c}}$ is given in Definition 16.3. The key result, from the algorithmic point of view, is:

Proposition 16.6 *There is an algorithm that produces a finite set $\{\eta_i\} \subset \mathcal{X}$ so that the union of the cosets of $\mathsf{Ker}(\bar{Q}^{\phi})$ determined by the η_i contains each $\theta \in \mathcal{X}$ that conjugates ϕ to ψ.*

The proof of Proposition 16.6 requires a detailed understanding of the structure of eigenrays and is the most technical part of the paper. The proof of Lemma 16.5 shows how to quickly construct $\mathrm{ALG}_{14.7}$ using $\mathrm{ALG}_{16.4}$ and the coset representatives produced by the algorithm of Proposition 16.6. In other words, to prove Proposition 14.7, we are reduced to proving:

Proposition 16.4 *There is an algorithm that takes as input $\phi, \psi \in \mathsf{UPG}(F_n)$ and a chain $\mathfrak{c}$ such that*

- *$\mathfrak{c}$ is a special chain for ϕ and ψ, and*
- *$\mathsf{I}_{\mathfrak{c}}(\phi) = \mathsf{I}_{\mathfrak{c}}(\psi)$,*

and that outputs `YES` *or* `NO` *depending on whether or not there is $\theta \in \mathsf{Ker}(\bar{Q}^{\phi})$ conjugating ϕ to ψ. Further, if* `YES` *then such a θ is produced.*

2.3 Proposition 16.4 $\Longleftarrow$ Proposition 17.2

This is an easy step. The details are given in the proof of Proposition 16.4 (assuming Lemma 17.1 and Proposition 17.2) following the statement of Proposition 17.2. After this step, we may assume that the restrictions of ϕ and ψ to the linear free factor system $\mathcal{F}_0(\phi) = \mathcal{F}_0(\psi)$ are equal. This provides the basis for an inductive argument completed in the next step.

2.4 Proof of Proposition 17.2

Proposition 17.2 is the inductive step of an argument up the filtration induced by $\mathfrak{c}$. There are six items labeled (1)–(5), (7) that are sequentially checked. If any of these is false then return `NO`. Otherwise, construct the desired conjugator following pages 1809–1811.

3 Background

3.1 Standard notation

The free group on n generators is denoted by F_n. For $a \in F_n$, conjugation by a is denoted by i_a, ie $i_a(x) = axa^{-1}$ for $x \in F_n$. The group of automorphisms of F_n, the group of inner automorphisms of F_n and the group of outer automorphisms of F_n are denoted by $\mathsf{Aut}(F_n)$, $\mathsf{Inn}(F_n) := \{i_a \mid a \in F_n\}$ and $\mathsf{Out}(F_n) = \mathsf{Aut}(F_n)/\mathsf{Inn}(F_n)$, respectively.

For subgroups $H < F_n$, $[H]$ denotes the conjugacy class of H and, for elements $a \in F_n$, $[a]$ denotes the conjugacy class of a.

An outer automorphism $\phi \in \mathsf{Out}(F_n)$ has *polynomial growth*, written $\phi \in \mathsf{PG}$, if for each $a \in F_n$ there is a polynomial P such that reduced word length of $\phi^k([a])$ with respect to a fixed set of generators of F_n is bounded above by $P(k)$. Equivalently, the set of attracting laminations for ϕ [Bestvina et al. 2000, Section 3] is empty. The set $\mathsf{UPG}(F_n)$ of *unipotent outer automorphisms* is the subset of $\mathsf{Out}(F_n)$ consisting of polynomially growing ϕ whose induced action on $H_1(F_n, \mathbb{Z})$ is unipotent. We sometimes write $\phi \in \mathsf{UPG}$ instead of $\phi \in \mathsf{UPG}(F_n)$. In Section 3.5 we show that $\phi \in \mathsf{UPG}$ if and only if $\phi \in \mathsf{PG}$ and ϕ is rotationless in the sense of [Feighn and Handel 2011, Definition 3.13] (where it is called forward rotationless). There is $K_n > 0$ such that if $\phi \in \mathsf{PG}$ then $\phi^{K_n} \in \mathsf{UPG}$ [Feighn and Handel 2018, Corollary 3.14].

The graph with one vertex $*$ and with n edges is the *rose* R_n. Making use of the standard identification of $\pi_1(R_n, *)$ with F_n, there are bijections between $\mathsf{Aut}(F_n)$ and the group of pointed homotopy classes of homotopy equivalences $f\colon (R_n, *) \to (R_n, *)$ and between $\mathsf{Out}(F_n)$ and the group of free homotopy classes of homotopy equivalences $f\colon R_n \to R_n$.

If G is a graph without valence one vertices then a homotopy equivalence $\mu\colon R_n \to G$ is called a *marking* and G, equipped with a marking, is called a *marked graph*. A marking μ induces an identification,

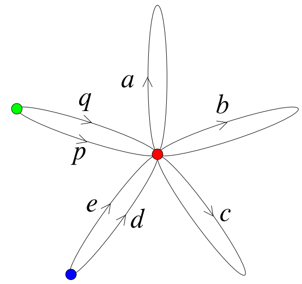

Figure 1: The graph G of Example 3.1.

well-defined up to inner automorphism, of the fundamental group of G with the fundamental group of R_n and hence with F_n. This in turn induces an identification of the group of homotopy classes of homotopy equivalences $f: G \to G$ with $\mathsf{Out}(F_n)$. If $\phi \in \mathsf{Out}(F_n)$ corresponds to the homotopy class of $f: G \to G$ then we say that $f: G \to G$ *represents* ϕ. In Section 3.6 we recall the existence of very well-behaved homotopy equivalences $f: G \to G$ representing an element of $\mathsf{UPG}(F_n)$.

Example 3.1 Here is an example of a homotopy equivalence $f: G \to G$ of a marked graph that represents an element ϕ of $\mathsf{UPG}(F_5)$. Let F_5 be represented as the fundamental group of the rose R_5, let G be the subdivision of R_5 pictured in Figure 1, and let $f: G \to G$ be given by $a \mapsto a$, $b \mapsto ba$, $c \mapsto cb$, $d \mapsto db^2$, $e \mapsto eb^3$, $p \mapsto pa^2$, $q \mapsto qc$. To see that ϕ has polynomial growth, note that the edge q has cubic growth in that

$$|f^k(q)| = |q \cdot c \cdot f(c) \cdot f^2(c) \cdot f^3(c) \cdot \ldots \cdot f^{k-1}(c)| = |q \cdot c \cdot cb \cdot cbba \cdot \ldots \cdot cbba \ldots ba^{k-1}| = \frac{k^3+5k+6}{6}$$

and that no edge grows at a higher rate. In particular conjugacy classes of F_n have at most cubic growth. As we progress through this paper, we will expand upon this example.

3.2 Paths, circuits and lines

A *path* in a marked graph G is a proper immersion of a closed interval into G. In this paper, we will assume that the endpoints of a path, if any, are at vertices. If the interval is degenerate then the path is *trivial*; if the interval is infinite or bi-infinite then the path is a *ray* or a *line*, respectively. We do not distinguish between paths that differ only by a reparametrization of the domain interval. Thus, every nontrivial path has a description as a concatenation of oriented edges and we will use this *edge path* formulation without further mention. Reversing the orientation on a path σ produces a path denoted by either $\overline{\sigma}$ or σ^{-1}. A *circuit* is an immersion of S^1 into G. Unless otherwise stated, a circuit is assumed to have an orientation. Circuits have cyclic edge decompositions. Each conjugacy class in F_n is represented by a unique circuit in G. The conjugacy class in F_n represented by the circuit σ is denoted by $[\sigma]$.

Notation 3.2 Each $\Phi \in \mathsf{Aut}(F_n)$ induces an equivariant homeomorphism of ∂F_n. To simplify notation somewhat, we refer to this extension as Φ rather than, say, $\partial\Phi$. In situations where this might cause confusion, we write $\Phi|\partial F_n$ for the induced homeomorphism of ∂F_n. For example, $\mathsf{Fix}(\Phi)$ is the subgroup of F_n fixed by Φ, and $\mathsf{Fix}(\Phi|\partial F_n)$ is the set of points in ∂F_n fixed by the induced homeomorphism.

The action of F_n on ∂F_n is by conjugation, ie by $a \cdot P = i_a(P)$ for $P \in \partial F_n$. For each nontrivial $a \in F_n$, i_a fixes two points in ∂F_n: a repeller a^- and an attractor a^+.

A marking μ induces an identification, well-defined up to inner automorphism, of the set of ends of $\widetilde{G}$ with ∂F_n, and likewise the group of covering translations of $\widetilde{G}$ with $\mathsf{Inn}(F_n)$. We choose such an identification once and for all. The covering translation corresponding to i_a is denoted by T_a as is the extension of T_a to a homeomorphism of ∂F_n. We have $T_a|\partial F_n = i_a|F_n$. If $f: G \to G$ represents ϕ then each lift $\tilde{f}: \widetilde{G} \to \widetilde{G}$ induces an equivariant homeomorphism, still called $\tilde{f}$, of ∂F_n; see, for example, Section 2.3 of [Feighn and Handel 2011]. There is a bijection between the set of lifts $\tilde{f}$ of $f: G \to G$ and the set of automorphisms Φ representing ϕ defined by $\tilde{f} \leftrightarrow \Phi$ if $\tilde{f}|\partial F_n = \Phi|\partial F_n$.

A line $\widetilde{L}$ in the universal cover $\widetilde{G}$ of a marked graph G is a bi-infinite edge path. The ends of $\widetilde{L}$ determine ends of $\widetilde{G}$ and hence points in ∂F_n. In this way, the *space of oriented lines in the tree* $\widetilde{G}$ can be identified with the space $\widetilde{\mathcal{B}}$ of ordered pairs of distinct elements of ∂F_n. The *space of oriented lines in* G is then identified with the space $\mathcal{B}$ of F_n-orbits of elements of $\widetilde{\mathcal{B}}$. The topology on ∂F_n induces a topology on $\widetilde{\mathcal{B}}$ and hence a topology on $\mathcal{B}$ called the *weak topology*.

3.3 Free factor systems

The subgroup system $\mathcal{F} = \{[A_1], \dots, [A_m]\}$ is a *free factor system* if $A_1, \dots, A_m$ are nontrivial free factors of F_n and either $F_n = A_1 * \cdots * A_m$ or $F_n = A_1 * \cdots * A_m * B$ for some nontrivial free factor B. The $[A_i]$ are the *components* of $\mathcal{F}$. If G is a marked graph and K is a subgraph whose noncontractible components are $K_1, \dots, K_m$ then $\mathcal{F}(K, G) := \{[\pi_1(K_1)], \dots, [\pi_1(K_m)]\}$ is a free factor system that is *realized by* $K \subset G$. Every free factor system $\mathcal{A}$ can be realized by $K \subset G$ for some marked graph G and some core subgraph $K \subset G$. Recall that a graph is *core* if through every edge there is an immersed circuit and that the *core of a graph* is the union of the images of its immersed circuits.

We write $\mathcal{F}_1 \sqsubset \mathcal{F}_2$ and say that $\mathcal{F}_1$ *is contained in* $\mathcal{F}_2$ if for each component $[A_i]$ of $\mathcal{F}_1$ there is a component $[B_j]$ of $\mathcal{F}_2$ so that A_i is conjugate to a subgroup of B_j. Equivalently, there is a marked graph G with core subgraphs $K_1 \subset K_2$ such that $\mathcal{F}_1 = \mathcal{F}(K_1, G)$ and $\mathcal{F}_2 = \mathcal{F}(K_2, G)$. If one can choose K_1 and K_2 so that $K_2 \setminus K_1$ is a single edge then we say that $\mathcal{F}_1 \sqsubset \mathcal{F}_2$ is a *one-edge extension*. For example, $\{[A]\} \sqsubset \{[B]\}$ is a one-edge extension if and only if $\mathrm{rank}(B) = \mathrm{rank}(A) + 1$ and $\{[A_1], [A_2]\} \sqsubset \{[B]\}$ is a one-edge extension if and only if $\mathrm{rank}(B) = \mathrm{rank}(A_1) + \mathrm{rank}(A_2)$.

Example 3.3 Suppose that H_1 is a subgraph of a marked graph G, that $H_2 = H_1 \cup E_2 \subset G$, where E_2 is an edge that forms a loop that is disjoint from H_1, and that $H_3 = H_2 \cup E_3$, where $E_3 \subset G$ is an edge with one endpoint in H_1 and the other at the unique endpoint of E_2. Then $\mathcal{F}(H_1, G) \sqsubset \mathcal{F}(H_2, G)$ and $\mathcal{F}(H_2, G) \sqsubset \mathcal{F}(H_3, G)$ are proper inclusions and $\mathcal{F}(H_1, G) \sqsubset \mathcal{F}(H_3, G)$ is a one-edge extension.

This is essentially the only way in which a one-edge extension can be "reducible". We record a specific consequence of this in the following lemma.

Lemma 3.4 *Suppose that $\mathcal{F}(H_1, G) \sqsubset \mathcal{F}(H_2, G)$ and $\mathcal{F}(H_2, G) \sqsubset \mathcal{F}(H_3, G)$ are proper inclusions and that $\mathcal{F}(H_1, G)$ and $\mathcal{F}(H_2, G)$ have the same number of components. Then $\mathcal{F}(H_1, G) \sqsubset \mathcal{F}(H_3, G)$ is not a one-edge extension.*

Proof This follows from [Handel and Mosher 2020, Part 2, Definition 2.4 and Lemma 2.5]. □

If $\mathcal{F} = \{[A_1], \dots, [A_m]\}$ and $a \in F_n$ is conjugate into some A_i then $[a]$ is *carried* by $\mathcal{F}$. A line $L \in \mathcal{B}$ is *carried* by $\mathcal{F}$ if it is a limit of periodic lines corresponding to conjugacy classes that are carried by $\mathcal{F}$. Equivalently, $[a]$ or L is *carried* by $\mathcal{F}$ if for some, and hence every, $K \subset G$ realizing $\mathcal{F}$, the realization of $[a]$ or L in G is contained in K. For every collection of conjugacy classes and lines there is a unique minimal (with respect to $\sqsubset$) free factor system that carries each element of the collection [Bestvina et al. 2000, Corollary 2.6.5].

Notation 3.5 $\mathsf{Out}(F_n)$ acts on the set of conjugacy classes $[F]$ of free factors F. If $\phi \in \mathsf{Out}(F_n)$ fixes $[F]$ then we say that $[F]$ *is ϕ-invariant* and write $\phi|[F]$ for the *restriction of ϕ to* $[F]$ (which is well-defined because F is its own normalizer in F_n). We often say that F is ϕ-invariant and write $\phi|F$ just to simplify notation. [Bestvina et al. 2005, Proposition 4.44] implies that if ϕ is UPG then $\phi|F$ is UPG. If $\mathcal{F} = \{[A_1], \dots, [A_m]\}$ is a free factor system and each $[A_i]$ is ϕ-invariant then we say that $\mathcal{F}$ is ϕ-invariant and denote $\{\phi|A_1, \dots, \phi|A_m\}$ by $\phi|\mathcal{F}$.

3.4 $\mathsf{Fix}_{\mathsf{N}}(\phi)$, principal lifts and $\mathcal{R}(\phi)$

We continue with Notation 3.2. If $P \in \mathsf{Fix}(\Phi|\partial F_n)$ and if there is a neighborhood U of P in ∂F_n such that $\Phi(U) \subset U$ and such that $\bigcap_{i=1}^{\infty} \Phi^i(U) = P$ then P is *attracting*. If P is an attracting fixed point for $\Phi^{-1}|\partial F_n$ then it is a *repelling* fixed point for $\Phi|\partial F_n$. By $\mathsf{Fix}_+(\Phi)$, $\mathsf{Fix}_-(\Phi)$ and $\mathsf{Fix}_{\mathsf{N}}(\Phi)$ we denote the *set of attracting fixed points* for $\Phi|\partial F_n$, the *set of repelling fixed points* for $\Phi|\partial F_n$ and the *set of nonrepelling fixed points* for $\Phi|\partial F_n$, respectively; thus $\mathsf{Fix}_{\mathsf{N}}(\Phi) = \mathsf{Fix}(\Phi|\partial F_n) \setminus \mathsf{Fix}_-(\Phi)$. Note that all of these sets are contained in ∂F_n.

If $A < F_n$ is a finitely generated subgroup then the inclusion of A into F_n is a quasi-isometric embedding and so extends to an inclusion of ∂A into ∂F_n with the property that $\{a^{\pm} \mid \text{nontrivial } a \in A\}$ is dense in ∂A. In particular, since the subgroup $\mathsf{Fix}(\Phi)$ consisting of elements in F_n that are fixed by Φ is finitely generated [Gersten 1987] (see also [Bestvina and Handel 1992]), we have $\partial\,\mathsf{Fix}(\Phi) \subset \partial F_n$. The following lemma implies that $\partial\,\mathsf{Fix}(\Phi) \subset \mathsf{Fix}(\Phi|\partial F_n)$, and that $\mathsf{Fix}_+(\Phi)$, $\mathsf{Fix}_-(\Phi)$ and $\mathsf{Fix}_{\mathsf{N}}(\Phi)$ are $\mathsf{Fix}(\Phi)$-invariant.

Lemma 3.6 *Let $\Phi \in \mathsf{Aut}(F_n)$ and $0 \neq a \in F_n$. The following are equivalent*:

- $a \in \mathsf{Fix}(\Phi)$.
- *Either a^- or a^+ is contained in $\partial\,\mathsf{Fix}(\Phi)$.*
- *Both a^- and a^+ are contained in $\partial\,\mathsf{Fix}(\Phi)$.*

- *The automorphism* i_a *commutes with* Φ.
- *The automorphism* $i_a|\partial F_n$ *commutes with* $\Phi|\partial F_n$.

Proof This is well known; see eg Lemmas 2.3 and 2.4 of [Bestvina et al. 2004] and Proposition I.1 of [Gaboriau et al. 1998]. □

Lemma 3.7 *If* $P \in \partial F_n$ *is fixed by automorphisms* $\Phi \neq \Phi'$ *representing* $\phi \in \mathsf{Out}(F_n)$ *then* $P = a^{\pm}$ *for some nontrivial* $a \in F_n$.

Proof There exists a nontrivial $a \in F_n$ such that $i_a = \Phi^{-1}\Phi'$ fixes P. □

Definition 3.8 An automorphism Φ representing $\phi \in \mathsf{UPG}(F_n)$ is *principal* if $\mathsf{Fix}_\mathsf{N}(\Phi)$ contains at least two points and if $\mathsf{Fix}_\mathsf{N}(\Phi) \neq \{a^-, a^+\}$ for any nontrivial $a \in F_n$. The *set of principal automorphisms representing* ϕ is denoted by $\mathcal{P}(\phi)$. See Section 3.2 of [Feighn and Handel 2011] for complete details.

Lemma 3.9 *If* Φ *is principal then* $\mathsf{Fix}_\mathsf{N}(\Phi)$ *is the disjoint union of* $\partial\,\mathsf{Fix}(\Phi)$ *and* $\mathsf{Fix}_+(\Phi)$. *Moreover,* $\mathsf{Fix}_+(\Phi)$ *is a union of finitely many* $\mathsf{Fix}(\Phi)$ *orbits.*

Proof The first assertion follows from Proposition I.1 of [Gaboriau et al. 1998]. The second is obvious if $\mathsf{Fix}_+(\Phi)$ is finite and follows from Lemma 2.5 of [Bestvina et al. 2004] if $\mathsf{Fix}_+(\Phi)$ is infinite. □

Remark 3.10 We sometimes say that $P \in \partial F_n$ is periodic if it is fixed by i_a for some nontrivial $a \in \mathcal{F}_n$. Nonperiodic points are dense in $\mathsf{Fix}_\mathsf{N}(\Phi)$ for each $\Phi \in \mathcal{P}(\phi)$. Lemmas 3.6 and 3.9 imply that no element of $\mathsf{Fix}_+(\Phi)$ is periodic. If $\mathsf{Fix}_+(\Phi) \neq \varnothing$ then $\mathsf{Fix}_+(\phi)$ is dense in $\mathsf{Fix}_\mathsf{N}(\phi)$ and we are done. Otherwise, $\mathsf{Fix}(\Phi)$ has rank at least two and $\mathsf{Fix}_\mathsf{N}(\Phi) = \partial\,\mathsf{Fix}(\Phi)$.

Definition 3.11 Two automorphisms Φ_1 and Φ_2 are in the same *isogredience class* if there exists $a \in F_n$ such that $\Phi_2 = i_a \Phi_1 i_a^{-1}$, in which case $\mathsf{Fix}_\mathsf{N}(\Phi_2) = i_a\,\mathsf{Fix}_\mathsf{N}(\Phi_1)$ and similarly for $\mathsf{Fix}_-(\Phi_2)$, $\mathsf{Fix}_+(\Phi_2)$ and $\mathsf{Fix}(\Phi_2)$. It follows that if Φ_1 and Φ_2 are isogredient then $[\mathsf{Fix}(\Phi_1)] = [\mathsf{Fix}(\Phi_2)]$ and $[\mathsf{Fix}_\mathsf{N}(\Phi_1)] = [\mathsf{Fix}_\mathsf{N}(\Phi_2)]$, where $[\ ]$ denotes the orbit under the action of F_n on sets of points in ∂F_n. It is easy to see that isogredience defines an equivalence relation on $\mathcal{P}(\phi)$. The *set of isogredience classes* of $\mathcal{P}(\phi)$ is denoted by $[\mathcal{P}(\phi)]$.

We recall the following result from Remark 3.9 of [Feighn and Handel 2011]; see also Lemma 3.25 of this paper.

Lemma 3.12 $\mathcal{P}(\phi)$ *is a finite union of isogredience classes.*

Our next lemma states that $[\mathsf{Fix}_\mathsf{N}(\Phi)]$ determines the isogredience class of $\Phi \in \mathcal{P}(\phi)$.

Lemma 3.13 *Suppose that* $\Phi_1, \Phi_2 \in \mathcal{P}(\phi)$. *Then* Φ_1 *and* Φ_2 *are isogredient if and only if* $[\mathsf{Fix}_\mathsf{N}(\Phi_1)] = [\mathsf{Fix}_\mathsf{N}(\Phi_2)]$. *More precisely,* $\Phi_2 = i_a \Phi_1 i_a^{-1}$ *if and only if* $\mathsf{Fix}_\mathsf{N}(\Phi_2) = i_a\,\mathsf{Fix}_\mathsf{N}(\Phi_1)$.

Proof It is obvious that if $\Phi_2 = i_a \Phi_1 i_a^{-1}$ then $\mathsf{Fix}_\mathsf{N}(\Phi_2) = i_a\,\mathsf{Fix}_\mathsf{N}(\Phi_1)$. For the converse note that if $\mathsf{Fix}_\mathsf{N}(\Phi_2) = i_a\,\mathsf{Fix}_\mathsf{N}(\Phi_1) = \mathsf{Fix}_\mathsf{N}(i_a \Phi_1 i_a^{-1})$ then $\Phi_2^{-1} i_a \Phi_1 i_a^{-1}$ is an inner automorphism whose induced action on ∂F_n fixes $\mathsf{Fix}_\mathsf{N}(\Phi_2)$ and so is not equal to $\{a^-, a^+\}$ for any nontrivial a. This proves that $\Phi_2^{-1} i_a \Phi_1 i_a^{-1}$ is trivial and so $\Phi_2 = i_a \Phi_1 i_a^{-1}$. □

Definition 3.14 Define sets

$$\mathsf{Fix}_\mathsf{N}(\phi) := \{[\mathsf{Fix}_\mathsf{N}(\Phi_1)], \dots, [\mathsf{Fix}_\mathsf{N}(\Phi_m)]\} \quad \text{and} \quad \mathcal{R}(\phi) := \left\{ [P] \;\middle|\; P \in \bigcup_{i=1}^{m} \mathsf{Fix}_+(\Phi_i) \right\} \subset \partial F_n / F_n\,,$$

and a multiset (repeated elements allowed)

$$\mathsf{Fix}(\phi) := \{[\mathsf{Fix}(\Phi_1)], \dots, [\mathsf{Fix}(\Phi_m)]\},$$

where the Φ_i are representatives of the isogredience classes in $\mathcal{P}(\phi)$. Thus $\mathsf{Fix}_\mathsf{N}(\phi)$ is a finite set of F_n-orbits of subsets of ∂F_n, and $\mathcal{R}(\phi)$ is a finite set of F_n-orbits of points in ∂F_n.

Definition 3.15 For us a *natural invariant of a group G* is a map $I : G \to X$, where X is a G-set and, for all $\phi, \theta \in G$, we have $I(\phi^\theta) = \theta(I(\phi))$.

The following lemma says that $[\mathcal{P}(\phi)]$, $\mathsf{Fix}(\phi)$, $\mathsf{Fix}_\mathsf{N}(\phi)$ and $\mathcal{R}(\phi)$ are natural invariants of $\mathsf{Out}(F_n)$.

Lemma 3.16 *Suppose that $\Theta \in \mathsf{Aut}(F_n)$ represents $\theta \in \mathsf{Out}(F_n)$ and that $\psi = \theta\phi\theta^{-1}$. Then*:

1. *The map $\Phi \mapsto \Psi := \Theta\Phi\Theta^{-1}$ defines a bijection between $\mathcal{P}(\phi)$ and $\mathcal{P}(\psi)$, and induces a bijection $[\mathcal{P}(\phi)] \leftrightarrow [\mathcal{P}(\psi)]$.*
2. $\mathsf{Fix}(\Psi) = \Theta(\mathsf{Fix}(\Phi))$, $\mathsf{Fix}_\mathsf{N}(\Psi) = \Theta(\mathsf{Fix}_\mathsf{N}(\Phi))$ *and* $\mathsf{Fix}_+(\Psi) = \Theta(\mathsf{Fix}_+(\Phi))$.
3. $\mathsf{Fix}(\psi) = \theta(\mathsf{Fix}(\phi))$, $\mathsf{Fix}_\mathsf{N}(\psi) = \theta(\mathsf{Fix}_\mathsf{N}(\phi))$ *and* $\mathcal{R}(\psi) = \theta(\mathcal{R}(\phi))$.

Proof The automorphism Ψ represents $\theta\phi\theta^{-1} = \psi \in \mathsf{Out}(F_n)$. If $\Phi' = i_c \Phi i_c^{-1}$ then $\Psi' := \Theta\Phi'\Theta^{-1} = i_{\Theta(c)}(\Theta\Phi\Theta^{-1}) i_{\Theta(c)}^{-1} = i_{\Theta(c)} \Psi i_{\Theta(c)}^{-1}$ so conjugation by Θ maps isogredience classes of ϕ to isogredience classes of ψ. The items in (2) are easy standard facts about conjugation. Since $\Theta(a^\pm) = (\Theta(a))^\pm$, it follows that Ψ is principal if Φ is principal. The induced map $\mathcal{P}(\phi) \to \mathcal{P}(\psi)$ is obviously invertible and is hence a bijection. This completes the proof of (1). If Θ is replaced by $i_a\Theta$ then Ψ is replaced by $i_a \Psi i_a^{-1}$ and $\mathsf{Fix}(\Psi)$, $\mathsf{Fix}_\mathsf{N}(\Psi)$ and $\mathsf{Fix}_+(\Psi)$ are replaced by $i_a(\mathsf{Fix}(\Psi))$, $i_a(\mathsf{Fix}_\mathsf{N}(\Psi))$ and $i_a(\mathsf{Fix}_+(\Psi))$ respectively. Thus $\theta([\mathsf{Fix}(\Phi)]) = [\mathsf{Fix}(\Psi)]$, $\theta([\mathsf{Fix}_\mathsf{N}(\Phi)]) = [\mathsf{Fix}_\mathsf{N}(\Psi)]$ and $\theta([\mathsf{Fix}_+(\Phi)]) = [\mathsf{Fix}_+(\Psi)]$. This verifies (3). □

The following lemma is used implicitly throughout the paper.

Lemma 3.17 *If A is a ϕ-invariant free factor then the inclusion of ∂A into ∂F_n induces an inclusion of $\mathcal{R}(\phi|A)$ into $\mathcal{R}(\phi)$.*

Proof An automorphism $\Phi'\colon A \to A$ representing $\phi|A$ extends to an automorphism $\Phi\colon F_n \to F_n$ representing ϕ. We claim that if $P \in \partial A$ then $P \in \mathsf{Fix}_+(\Phi')$ if and only if $P \in \mathsf{Fix}_+(\Phi)$. Symmetrically, $P \in \mathsf{Fix}_-(\Phi')$ if and only if $P \in \mathsf{Fix}_-(\Phi)$. It follows that $\mathsf{Fix}_{\mathsf{N}}(\Phi') \subset \mathsf{Fix}_{\mathsf{N}}(\Phi)$ and hence that Φ is principal if Φ' is principal. This will complete the proof of the lemma.

To prove the claim, extend a basis $\mathcal{A}$ for A to a basis $\mathcal{B}$ for F_n. Following [Gaboriau et al. 1998], we view $P \in \partial F_n$ as an infinite word $P = x_1x_2x_3\cdots$ with each $x_i \in \mathcal{B}$. For each $i \in \mathbb{N}$, let $x_1, \dots, x_{k(i)}$ be the common initial segment of $\Phi(x_1 \dots x_i)$ and P. Then $P \in \mathsf{Fix}_+(\Phi)$ if and only if $k(i) - i \to \infty$ [Gaboriau et al. 1998, Proposition I.1]. If $P \in \partial A$ then each $x_i \in \mathcal{A}$ and each $\Phi(x_1 \dots x_i) = \Phi'(x_1 \dots x_i) \in A$ so $k(i)$ is the same whether we compute using Φ or Φ'. □

3.5 UPG is rotationless

Relative train track theory is most effective when applied to elements of $\mathsf{Out}(F_n)$ that are rotationless as defined in [Feighn and Handel 2011, Definition 3.13 and Remark 3.14]. In this section, we show that for PG elements, ϕ is rotationless if and only if ϕ is UPG. The exact definition of rotationless plays no role in this paper so is not repeated here.

Lemma 3.18 *Each $\phi \in \mathsf{UPG}$ is rotationless.*

Proof By [Bestvina et al. 2000, Proposition 5.7.5], there is a sequence $\mathcal{F}_0 \sqsubset \mathcal{F}_1 \sqsubset \cdots \sqsubset \mathcal{F}_K$ of ϕ-invariant one-edge extensions where $\mathcal{F}_0$ is trivial and $\mathcal{F}_K = \{[F_n]\}$. We may assume without loss of generality that $\mathcal{F}_0 \sqsubset \mathcal{F}_1 \sqsubset \cdots \sqsubset \mathcal{F}_K$ is a maximal such chain.

[Feighn and Handel 2011, Theorem 2.19], which makes no assumptions on ϕ, proves the existence of a relative train track map $f\colon G \to G$ and filtration $\varnothing = G_0 \subset G_1 \subset \cdots \subset G_N = G$ representing ϕ and satisfying a certain list of five properties, two of which are denoted by (P) and (NEG). Additionally, the filtration realizes $\mathcal{F}_1 \sqsubset \mathcal{F}_2 \sqsubset \cdots \sqsubset \mathcal{F}_K$ in the sense that each $\mathcal{F}_k$ is represented by a core filtration element G_{i_k}. Since $\mathcal{F}_0 \sqsubset \mathcal{F}_1 \sqsubset \cdots \sqsubset \mathcal{F}_K$ is maximal, G_{i_k} is obtained from $G_{i_{k-1}}$ by adding either a topological circle that is disjoint from $G_{i_{k-1}}$ or a topological arc with both endpoints in $G_{i_{k-1}}$ [Handel and Mosher 2020, Part II, Lemma 2.5]. We denote the closure of $G_{i_k} \setminus G_{i_{k-1}}$, equipped with the simplicial structure inherited from G_{i_k}, by $\widehat{H}_k$. Since ϕ is PG, there are no EG strata.

We use the following consequences of properties (P) and (NEG).

(1) The terminal endpoint of a nonperiodic edge in $\widehat{H}_k$ is contained in $G_{i_{k-1}}$.

(2) If $\widehat{H}_k$ contains a periodic edge then it is a single periodic stratum [Feighn and Handel 2011, Lemma 2.20(1)].

If $\widehat{H}_k$ is a circle that is disjoint from $G_{i_{k-1}}$, then its conjugacy class is fixed by some iterate of ϕ and so is fixed by ϕ [Bestvina et al. 2005, Proposition 3.16]. There are two possibilities; $\widehat{H}_k$ is a single fixed edge; or $\widehat{H}_k$ has more than one edge and $f|\widehat{H}_k$ is a nontrivial rotation with one orbit of edges. In the latter case, we say that $\widehat{H}_k$ is a *rotating circle*.

If $\widehat{H}_k$ intersects $G_{i_{k-1}}$ then it is a topological arc E_k whose ends may or may not be identified. Either $f(E_k) = v_k E_k u_k$ or $f(E_k) = v_k \bar{E}_k u_k$ for some paths $u_k, v_k \subset G_{i_{k-1}}$ [Bestvina et al. 2000, Corollary 3.2.2]. Since ϕ is UPG, the latter is ruled out by [Bestvina et al. 2005, Proposition 5.7.5(2); see the second paragraph on page 595]. If both u_k and v_k are trivial then E_k is a single fixed edge. If exactly one of u_k and v_k is trivial then E_k is a single nonperiodic NEG edge. If neither u_k and v_k are trivial then E_k consists of two nonperiodic NEG edges with a common fixed initial endpoint. In all three cases, the directions determined by E_k and $\bar{E}_k$ are either nonperiodic or fixed.

An easy induction argument on k shows that:

(a) If a vertex v is not contained in a rotating circle then v is fixed by f and each periodic direction based at v is fixed by f.

(b) Each rotating circle is a component of $\mathrm{Per}(f)$, the set of periodic points for f, and each point in a rotating circle has exactly two periodic directions.

These are exactly the conditions needed to verify that $f : G \to G$ is rotationless in the sense of [Feighn and Handel 2011, Definition 3.18]. Proposition 3.19 of [loc. cit.] states that the existence of a rotationless $f : G \to G$ satisfying the conclusions of [Feighn and Handel 2011, Theorem 2.19] is equivalent to ϕ being rotationless. □

Remark 3.19 The converse of Lemma 3.18, that every rotationless PG ϕ is UPG, is also true. We make no use of this fact but include a proof for completeness. See Section 3.6 for a review of CTs. Since ϕ is rotationless and PG, it is represented by a CT $f : G \to G$ without EG or zero strata. For any such $f : G \to G$, there is a filtration $\varnothing = G_0 \subset G_1 \subset \cdots \subset G_N = G$ by f-invariant core subgraphs such that G_i is obtained from G_{i-1} by adding a single topological edge E_i whose image $f(E_i) \subset G_i$ crosses E_i exactly once and crosses $\bar{E}_i$ not at all. [Bestvina et al. 2000, Proposition 5.7.5] therefore implies that ϕ is UPG.

3.6 CTs

A CT is a particularly nice kind of homotopy equivalence $f : G \to G$ of a marked directed graph. Every rotationless ϕ, and in particular every $\phi \in \mathsf{UPG}(F_n)$, is represented by a CT; see [Feighn and Handel 2011, Theorem 4.28] or [Bestvina et al. 2000, Theorem 5.1.8]. Moreover, CTs are considerably simpler in the $\mathsf{UPG}(F_n)$-case than in the general case.

For the remainder of the section we assume that $f : G \to G$ is a CT representing an element of $\mathsf{UPG}(F_n)$ and review its properties. Complete details can be found in [Feighn and Handel 2011] (see in particular Section 4.1) and in [Feighn and Handel 2018]. The latter introduces the (Inheritance) property for a CT $f : G \to G$, which states that the restriction of f to each component of each core filtration element is also a CT, and contains an algorithm to produce CTs satisfying (Inheritance). We say that $f : G \to G$ *realizes* a chain $\mathcal{F}_0 \sqsubset \mathcal{F}_1 \sqsubset \cdots \sqsubset \mathcal{F}_K$ of ϕ-invariant free factor systems if each F_j is realized by an f-invariant core subgraph of G; see Section 3.3.

Theorem 3.20 [Feighn and Handel 2018, Theorem 1.1] *There is an algorithm whose input is a rotationless $\phi \in \mathsf{Out}(F_n)$ and whose output is a* CT $f : G \to G$ *that represents ϕ and satisfies* (Inheritance). *Moreover, for any chain $\mathcal{C}$ of ϕ-invariant free factor systems, one can choose $f : G \to G$ to realize $\mathcal{C}$.*

We assume throughout this paper that our chosen CTs satisfy (Inheritance).

The marked graph G comes equipped with an f-invariant filtration $\varnothing = G_0 \subset G_1 \subset \cdots \subset G_N = G$ by subgraphs G_i in which each G_i is obtained from G_{i-1} by adding a single oriented edge E_i. For each E_i there is a (possibly trivial) closed path $u_i \subset G_{i-1}$ such that $f(E_i) = E_i u_i$; if u_i is nontrivial then it forms a circuit. A path or circuit has *height i* if it crosses E_i, meaning that either E_i or $\overline{E}_i$ occurs in its edge decomposition, but does not cross E_j for any $j > i$.

Example 3.1 (continued) The homotopy equivalence $f : G \to G$ is a CT with f-invariant filtration $\varnothing = G_0 \subset G_1 \subset \cdots \subset G_7 = G$ given by adding one edge at a time in alphabetical order.

Every map α into G with domain a closed interval or S^1 and with endpoints, if any, at vertices is properly homotopic rel endpoints to a path or circuit $[\alpha]$; we say that $[\alpha]$ is obtained from α by *tightening*. If σ is a path or circuit then we usually denote $[f(\sigma)]$ by $f_\#(\sigma)$. A decomposition into subpaths $\sigma = \sigma_1 \cdot \sigma_2 \cdot \ldots$ is a *splitting* if $f_\#^k(\sigma) = f_\#^k(\sigma_1) \cdot f_\#^k(\sigma_2) \cdot \ldots$ for all $k \geq 1$. In other words, $f^k(\sigma)$ can be tightened by tightening each $f^k(\sigma_i)$.

A finite path σ is a *Nielsen path* if $f_\#(\sigma) = \sigma$; it is an *indivisible Nielsen path* if it is not a fixed edge and does not split into a nontrivial concatenation of Nielsen paths. Every Nielsen path has a splitting into fixed edges and indivisible Nielsen paths. If $f_\#^k(\sigma) = \sigma$ for some $k \geq 1$ then σ is a *periodic Nielsen path*. In a CT, every periodic Nielsen path is a Nielsen path.

An edge E_i is *linear* if u_i is a nontrivial Nielsen path. The *set of oriented linear edges* is denoted by $\mathsf{Lin}(f)$ and the set obtained from $\mathsf{Lin}(f)$ by reversing orientation is denoted by $\mathsf{Lin}^{-1}(f)$. In our example, $\mathsf{Lin}(f) = \{b, p\}$.

Associated to a CT $f : G \to G$ is a finite set of nontrivial closed Nielsen paths called *twist paths*. This set is well-defined up to a change of orientation on each path. In the remainder of this paragraph we recall some useful properties of twist paths. Each twist path w determines a circuit $[w]$ in G representing a root-free[1] conjugacy class in F_n and distinct twist paths determine distinct unoriented circuits; ie circuits whose cyclic edge decompositions differ by more than a change of orientation. For each twist path w, the set $\mathsf{Lin}_w(f)$ of (necessarily linear) edges E_i such that $f(E_i) = E_i w^{d_i}$ for some $d_i \neq 0$ is nonempty and is called the *linear family associated to* w; note that $f_\#^k(E_i) = E_i w^{k d_i}$ grows linearly in k. Every linear edge belongs to one of these linear families. If $E_i \in \mathsf{Lin}_w(f)$ and $p \neq 0$ then $E_i w^p \overline{E}_i$ is an indivisible Nielsen path. All indivisible Nielsen paths have this form. If E_i and E_j are distinct edges in $\mathsf{Lin}_w(f)$ then $d_i \neq d_j$; if d_i and d_j have the same sign, then paths of the form $E_i w^p \overline{E}_j$ are *exceptional paths*

[1] Nontrivial $a \in F_n$ is *root-free* if $x \in F_n$ and $x^k = a$ implies $k = \pm 1$.

associated to w. Note that $f_\#^k(E_i w^p \overline{E}_j) = E_i w^{p+k(d_i-d_j)} \overline{E}_j$ so these paths also grow linearly under iteration. Exceptional paths have no nontrivial splittings (which would not be true if we allowed d_i and d_j to have the opposite sign).

Example 3.1 (continued) In our example, we choose our set of twist paths to be $\{a\}$, as opposed to $\{a^{-1}\}$.

A splitting $\sigma = \sigma_1 \cdot \sigma_2 \cdot \ldots$ is a *complete splitting* if each σ_i is either a single edge or an indivisible Nielsen path or an exceptional path. If σ_i is not a Nielsen path then it is *a growing term*; if at least one σ_i is growing then σ is growing. We say that σ_i is *a linear term* if it is exceptional or equal to E or $\overline{E}$ for some $E \in \mathsf{Lin}(f)$. Complete splittings are unique when they exist [Feighn and Handel 2011, Lemma 4.11]. A path with a complete splitting is said to be *completely split*. For each edge E_i there is a complete splitting of $f(E_i)$ whose first term is E_i and whose remaining terms define a complete splitting of u_i. The image under $f_\#$ of a completely split path or circuit is completely split. For each path or circuit σ, the image $f_\#^k(\sigma)$ is completely split for all sufficiently large k [Feighn and Handel 2011, Lemma 4.25].

The *set of oriented nonfixed nonlinear edges* is denoted by $\mathcal{E}_f$ and the set obtained from $\mathcal{E}_f$ by reversing orientation is denoted by $\mathcal{E}_f^{-1}$. We say that an edge in $\mathcal{E}_f$ or $\mathcal{E}_f^{-1}$ has *higher order*. An easy induction argument shows that, for each $E_i \in \mathcal{E}_f$ and $k \geq 1$, $f_\#^k(E_i)$ is completely split and

$$f_\#^k(E_i) = E_i \cdot u_i \cdot f_\#(u_i) \cdot \ldots \cdot f_\#^{k-1}(u_i).$$

Thus $f_\#^{k-1}(E_i)$ is an initial segment of $f_\#^k(E_i)$ and the union

$$R_{E_i} = E \cdot u_i \cdot f_\#(u_i) \cdot f_\#^2(u_i) \cdot \ldots$$

of this nested sequence is an $f_\#$-invariant ray called the *eigenray associated to* E_i. The complete splittings of the individual $f_\#^k(u_i)$'s define a complete splitting of R_{E_i}.

Example 3.1 (continued) In our example, the edge a is fixed, the edges b and p are linear, and the other edges have higher order, ie $\mathcal{E}_f = \{c, d, e, q\}$. As an example of an eigenray,

$$R_q = q \cdot c \cdot cb \cdot cbba \cdot \ldots \cdot cbba \ldots ba^{k-1} \cdot \ldots.$$

Lemma 3.21 *If $E \in \mathcal{E}_f \cup \mathcal{E}_f^{-1}$ then E is not crossed by any Nielsen path or exceptional path. In particular, each crossing of E by a completely split path is a term in the complete splitting of that path.*

Proof Suppose that some Nielsen path σ crosses E. Since σ is a concatenation of fixed edges and indivisible Nielsen paths and since every indivisible Nielsen path has the form $E_i w^p \overline{E}_i$ for some linear edge and twist path w, E must be crossed by some Nielsen path w with height lower than σ. The obvious induction argument completes the proof. □

Remark 3.22 One can define R_E for a linear edge E in the same way that one does for a higher-order edge. If $f(E) = Ew^d$ for some twist path w then $R_E = Ew^\infty$ if $d > 0$ and $R_E = Ew^{-\infty}$ if $d < 0$. These rays play a different role in the theory than eigenrays.

3.7 Principal lifts from the CT point of view

Suppose that $f: G \to G$ is a CT representing ϕ. If Φ is a principal lift for ϕ then we say that the corresponding $\tilde{f}$ *is a principal lift of* $f: G \to G$.

Lemma 3.23 *A lift $\tilde{f}: \tilde{G} \to \tilde{G}$ is principal if and only if $\mathsf{Fix}(\tilde{f}) \neq \varnothing$ in which case $\mathsf{Fix}(\tilde{f})$ contains a vertex.*

Proof This follows from Corollary 3.17, Corollary 3.27 and Remark 4.9 of [Feighn and Handel 2011] and the fact that $\mathsf{Fix}(f)$ is a union of vertices and fixed edges. □

Lemma 3.24 *Suppose that $\tilde{f}: \tilde{G} \to \tilde{G}$ is the lift of $f: G \to G$ corresponding to $\Phi \in \mathcal{P}(\phi)$ and that $a \in F_n$. The following are equivalent*:

- $a \in \mathsf{Fix}(\Phi)$;
- $i_a|\partial F_n = T_a|\partial F_n$ *commutes with* $\Phi|\partial F_n = \tilde{f}|\partial F_n$;
- $T_a|\tilde{G}$ *commutes with* $\tilde{f}|\tilde{G}$;
- $T_a(\mathsf{Fix}(\tilde{f}|\tilde{G})) = \mathsf{Fix}(\tilde{f}|\tilde{G})$.

Proof This is well known. All but the equivalence of the third and fourth bullets can be found in Lemma 2.4 of [Bestvina et al. 2004]. If $T_a|\tilde{G}$ commutes with $\tilde{f}|\tilde{G}$ then it preserves the fixed-point set of $\tilde{f}|\tilde{G}$. Conversely, if $\tilde{x}, T_a(\tilde{x}) \in \mathsf{Fix}(\tilde{f}|\tilde{G})$ then $\tilde{f}|\tilde{G}$ and $T_a \tilde{f} T_a^{-1}|\tilde{G}$ both fix $T_a(\tilde{x})$ and so must be equal. This proves the equivalence of the third and fourth bullets. □

We say that lifts $\tilde{f}_1$ and $\tilde{f}_2$ are *isogredient* if they correspond to isogredient automorphisms Φ_1 and Φ_2. Equivalently, $\tilde{f}_2 = T_a \tilde{f}_1 T_a^{-1}$ for some covering translation T_a. Recall that $x, y \in \mathsf{Fix}(f)$ are *Nielsen equivalent* if they are the endpoints of a Nielsen path or equivalently, if for each lift $\tilde{x}$, the unique lift $\tilde{f}$ that fixes $\tilde{x}$ also fixes some lift $\tilde{y}$ of y.

Lemma 3.25 *The map which assigns to each principal lift $\tilde{f}: \tilde{G} \to \tilde{G}$ the projection into G of $\mathsf{Fix}(\tilde{f})$ induces a bijection between the set of isogredience classes of principal lifts and the set of Nielsen classes for f. In particular, there are only finitely many isogredience classes of principal lifts.*

Proof This follows from [Feighn and Handel 2011, Lemma 3.8] and Lemma 3.23. □

Recall from Section 3.6 that for each $E \in \mathcal{E}_f$ there is a closed completely split path u such that $f(E) = E \cdot u$ is a splitting and such that the eigenray $R_E = E \cdot u \cdot f_\#(u) \cdot f_\#^2(u) \cdot \ldots$ is $f_\#$-invariant.

The following lemma is similar to [Feighn and Handel 2018, Lemma 3.10], which applies more generally but has a weaker conclusion.

Lemma 3.26 *Suppose that $\tilde{f}$ corresponds to $\Phi \in \mathcal{P}(\phi)$. If $\tilde{E}$ is a lift of $E \in \mathcal{E}_f$ and if the initial endpoint of $\tilde{E}$ is contained in $\mathsf{Fix}(\tilde{f})$ then the lift $\tilde{R}_{\tilde{E}}$ of R_E that begins with $\tilde{E}$ converges to a point in $\mathsf{Fix}_+(\Phi)$. This defines a bijection between $\mathsf{Fix}_+(\Phi)$ and the set of all such $\tilde{E}$ and also a bijection between $\mathcal{R}(\phi)$ and $\mathcal{E}_f$.*

Proof $\widetilde{R}_{\widetilde{E}}$ converges to some $P \in \mathsf{Fix}_\mathsf{N}(\Phi)$ by [Feighn and Handel 2011, Lemma 4.36(1)]. Since E is not linear, u is not a Nielsen path and hence not a periodic Nielsen path. The length of $f_\#^k(u)$ therefore goes to infinity with k. Proposition I.1 of [Gaboriau et al. 1998] implies that $P \in \mathsf{Fix}_+(\Phi)$.

Suppose that $\widetilde{E}_1$ and $\widetilde{E}_2$ are distinct edges that project into $\mathcal{E}_f$, that the initial endpoint $\tilde{x}_i$ of $\widetilde{E}_i$ is fixed by $\tilde{f}$ and that, for $i = 1, 2$, $\widetilde{R}_i$ is the lift of R_{E_i} with initial edge $\widetilde{E}_i$. The path that connects $\tilde{x}_1$ to $\tilde{x}_2$ projects to a Nielsen path $\sigma \subset G$. If $\widetilde{R}_1$ and $\widetilde{R}_2$ converge to the same point in $\mathsf{Fix}_+(\Phi)$ then σ crosses E_1 or E_2 in contradiction to Lemma 3.21. This proves that the map $\{\widetilde{E}\} \mapsto \mathsf{Fix}_+(\Phi)$ is injective; surjectivity follows from [Feighn and Handel 2011, Lemma 4.36(2)]. The second bijection is obtained from the first by projecting to the sets of F_n-orbits. □

Example 3.1 (continued) Since $\mathcal{E}_f = \{c, d, e, q\}$, $\mathcal{R}(\phi)$ has four elements, denoted by $\{r_c, r_d, r_e, r_q\}$; see Figure 2.

4 Recognizing a conjugator

Associated to each CT $f: G \to G$ representing a rotationless element $\phi \in \mathsf{Out}(F_n)$ is a finite type labeled graph $\Gamma(f)$ that realizes $\mathsf{Fix}_\mathsf{N}(\phi)$. We refer to $\Gamma(f)$ as the *eigengraph* for $f: G \to G$. In Section 4.1 we recall the construction and relevant properties of $\Gamma(f)$ in the case that $\phi \in \mathsf{UPG}(F_n)$. Every ϕ-invariant conjugacy class $[a]$ is represented by an oriented circuit in $\Gamma(f)$. There is a finite set $\mathcal{A}_{\mathrm{or}}(\phi)$ of such $[a]$ that are root-free and that are represented by more than one oriented circuit in $\Gamma(f)$. The "extra" circuits correspond to the linear edges in $f: G \to G$. In Section 4.2, we describe how the extra circuits can be incorporated into an invariant $\mathsf{SA}(\phi)$ of ϕ that is independent of the choice of $f: G \to G$. Moreover, certain pairs of elements of $\mathsf{SA}(\phi)$ have *twist coordinates* that can be read off from the twist coordinates on linear edges in $f: G \to G$. Section 4.3 is an application of the recognition theorem of [Feighn and Handel 2011]. Assuming that $\phi, \psi \in \mathsf{UPG}(F_n)$, we use eigengraphs, $\mathsf{SA}(\phi)$ and twist coordinates to give necessary and sufficient conditions for a given $\theta \in \mathsf{Out}(F_n)$ to conjugate ϕ to ψ.

4.1 The eigengraph $\Gamma(f)$

In this section, we recall a finite type labeled graph that captures many of the invariants of ϕ that are essential to our algorithm. For further details and more examples, see [Feighn and Handel 2018, Sections 9, 10 and 12].

A graph Γ without valence one vertices and equipped with a simplicial immersion $p: \Gamma \to G$ to a marked graph G will be called a *Stallings graph*. We label the vertices and edges of Γ by their p-images in G. Two Stallings graphs $p_1: \Gamma_1 \to G$ and $p_2: \Gamma_2 \to G$ are equivalent if there is a label-preserving simplicial homeomorphism $h: \Gamma_1 \to \Gamma_2$. We will not distinguish between equivalent Stallings graphs. Since all vertices have valence at least two, every edge in Γ is crossed by a line. We say that Γ has *finite type* if its core is finite and if the complement of the core is a finite union of rays.

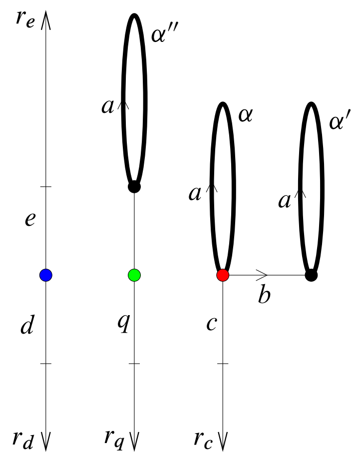

Figure 2: The eigengraph $\Gamma(f)$ of our example. Only the first edge of the eigenrays is labeled here. For example, the eigenray $R_c = cbbaba^2ba^3\ldots$ starts at the red vertex and only its first edge is labeled (by c). Other aspects of this figure are explained later.

Given a CT $f\colon G \to G$ representing ϕ, we construct a finite type Stallings graph $\Gamma(f)$, called the *eigengraph for* $f\colon G \to G$, as follows. Let $\Gamma^0(f)$ be G with the interiors of all nonfixed edges removed. In particular, $\Gamma^0(f)$ may contain isolated vertices. The labeling on $\Gamma^0(f)$ is the obvious one. For each $E \in \mathsf{Lin}(f)$, first attach an edge, say E', to $\Gamma^0(f)$ by identifying the initial endpoint of E' with the initial endpoint of E, thought of as a vertex in $\Gamma^0(f)$. The label on E' is E. Then add a path ω by attaching both of its endpoints to the terminal endpoint of E', which now has valence three. If w denotes the twist path associated to E then we label ω by w, thought of as an edge path, and subdivide ω so that each edge in ω is labeled by a single edge in G. The net effect is to add a *lollipop* to $\Gamma^0(f)$ for each edge in $\mathsf{Lin}(f)$. This labeling defines an immersion because w determines a circuit that does not cross the edge E. Finally, for each $E \in \mathcal{E}_f$, attach a ray labeled R_E (as defined in Section 3.6) by identifying the initial endpoint of this ray with the initial endpoint of E, thought of as a vertex in $\Gamma^0(f)$. We will also use the term *eigenray* for this added ray. The resulting graph is denoted by $\Gamma(f)$. This labeling maintains the immersion property because R_E is an immersed ray in G and because no other edge labeled E has initial vertex in $\Gamma^0(f)$.

Example 3.1 (continued) The eigengraph for our running example is pictured in Figure 2.

The vertices of $\Gamma(f)$ that are not in $\Gamma^0(f)$ have valence either two or three by construction. The valence of $v \in \Gamma^0(f)$ in $\Gamma(f)$ is equal to the number of fixed directions based at v in G. If v, thought of as a vertex in G, is not the terminal endpoint of an edge in $\mathcal{E}_f \cup \mathsf{Lin}(f)$, then v has the same valence in $\Gamma(f)$ that it does in G. If v is the terminal endpoint of an edge in $\mathcal{E}_f \cup \mathsf{Lin}(f)$, let E_i be the lowest such edge. Then $f(E_i) = E_i \cdot u_i$, where u_i is a closed path based at v whose ends determine distinct fixed directions at v by [Feighn and Handel 2011, Lemma 4.21]. This proves that v has valence at least two in $\Gamma(f)$ and hence that $\Gamma(f)$ is a Stallings graph. It has finite type by construction.

As noted in Section 3.6, each Nielsen path in G decomposes as a concatenation of fixed edges and indivisible Nielsen paths and each indivisible Nielsen path is a closed path. It follows that two vertices in $\mathsf{Fix}(f)$ are in the same Nielsen class if and only if they are connected by a sequence of fixed edges. In particular, the vertices in each component of $\Gamma^0(f)$ form exactly one Nielsen class in $\mathsf{Fix}(f)$. Since the inclusion of $\Gamma^0(f)$ into $\Gamma(f)$ induces a bijection of components, there is a bijection between the set of components of $\Gamma(f)$ and the set of Nielsen classes in $\mathsf{Fix}(f)$ and hence (Lemma 3.25) a bijection between the set of components of $\Gamma(f)$ and the set of isogredience classes in $\mathcal{P}(\phi)$. We denote the component of $\Gamma(f)$ corresponding to the isogredience class $[\Phi]$ by $\Gamma_{[\Phi]}(f)$ or by $\Gamma(\tilde{f})$, where $\tilde{f}$ is the lift of f that corresponds to Φ.

We say that *a line is carried by* $\Gamma_{[\Phi]}(f)$ if its realization $L \subset G$ lifts into $\Gamma_{[\Phi]}(f)$ and is *carried by* $\Gamma(f)$ if it is carried by some component $\Gamma_{[\Phi]}(f)$. The following lemma shows that the set of lines carried by $\Gamma(f)$ is independent of the choice of $f: G \to G$. We will sometimes refer to these as *principal lines.*

Lemma 4.1 *The following are equivalent for any* CT $f: G \to G$ *representing* ϕ, *any* $\Phi \in \mathcal{P}(\phi)$ *and any line* $L \subset G$.

(1) L *is carried by* $\Gamma_{[\Phi]}(f)$ *(resp. the core of* $\Gamma_{[\Phi]}(f)$*).*

(2) *There is a lift* $\tilde{L} \subset \tilde{G}$ *such that* $\{\partial_\pm \tilde{L}\} \subset \mathsf{Fix}_{\mathsf{N}}(\Phi)$ *(resp.* $\{\partial_\pm \tilde{L}\} \subset \partial\,\mathsf{Fix}(\Phi)$*).*

Proof It suffices to prove the unbracketed statement. Let $q: \tilde{G} \to G$ and $q_\Gamma: \tilde{\Gamma}_{[\Phi]}(f) \to \Gamma_{[\Phi]}(f)$ be the universal covering maps. The labeling map $p: \Gamma_{[\Phi]}(f) \to G$ is an immersion and so lifts to an embedding $\tilde{p}: \tilde{\Gamma}_{[\Phi]}(f) \hookrightarrow \tilde{G}$. If a line $L \subset G$ lifts to a line $L_\Gamma \subset \Gamma_{[\Phi]}(f)$ and if $\tilde{L}_\Gamma \subset \tilde{\Gamma}_{[\Phi]}(f)$ is a lift of L_Γ then $\tilde{L} := \tilde{p}(\tilde{L}_\Gamma) \subset \tilde{G}$ is a lift of L. Conversely, if $L \subset G$ lifts to $\tilde{L} \subset \tilde{G}$ and there exists $\tilde{L}_\Gamma \subset \tilde{\Gamma}_{[\Phi]}(f)$ with $\tilde{p}(\tilde{L}_\Gamma) = \tilde{L}$, then $L_\Gamma := q_\Gamma(\tilde{L}_\Gamma)$ is a lift of L. The lemma therefore follows from [Feighn and Handel 2018, Lemma 12.4], which states that $L \subset G$ lifts to $\tilde{p}(\tilde{L}_\Gamma)$ if and only if (2) is satisfied. □

An end of an immersed line in $\Gamma(f)$ can either be an end of $\Gamma(f)$ or can wrap infinitely around one of the lollipop circuits or can cross a vertex in $\Gamma^0(f)$ infinitely often. This gives the following description of lines that lift into $\Gamma(f)$.

Lemma 4.2 *A line* $\sigma \subset G$ *lifts into* $\Gamma(f)$ *if and only if it contains a (possibly trivial) subpath* β *that is a concatenation of fixed edges and indivisible Nielsen paths and such that the complement of* β *is* 0, 1 *or* 2 *rays, each of which is either* R_e *for some higher-order edge* e, *or* $Ew^{\pm\infty}$ *for some twist path* w *and some linear edge* $E \in \mathsf{Lin}_w(f)$. □

The following lemma, in conjunction with Lemma 4.1, implies that the set of conjugacy classes determined by twist paths and their inverses is an invariant of ϕ. This set is explored further in Section 4.2.

Lemma 4.3 *Suppose that $f : G \to G$ is a* CT *representing ϕ and that $[a]$ is a root-free conjugacy class that is fixed by ϕ and that σ_a is the circuit in G representing $[a]$.*

(1) *If $[a] = [w]$ (resp. $[a] = [\overline{w}]$) for some twist path w, then for each edge $E \in \mathsf{Lin}_w(f)$ there is a lift of σ_a to the loop ω (resp. $\overline{\omega}$) in the lollipop associated to E. Additionally there is a unique lift of σ_a to a circuit in $\Gamma(f)$ that is not contained in any lollipop.*

(2) *Otherwise, there is a unique lift of σ_a to a circuit in $\Gamma(f)$.*

Proof We claim that if τ' is a path in the core of $\Gamma(f)$ that projects to a Nielsen path τ in G and if the initial vertex v'' of τ' is not in $\Gamma^0(f)$, then τ' is contained in the loop ω of some lollipop. We may assume without loss that the claim holds for paths with height less than that of τ and that τ is either a single fixed edge or an indivisible Nielsen path. Since the core of $\Gamma(f)$ is contained in the union of $\Gamma^0(f)$ with the lollipops associated to the linear edges of G, there is a lollipop composed of an edge E_1' projecting to a linear edge $E_1 \subset G$ and a loop ω projecting to its twist path w_1 such that $v'' \in \omega$. If τ is a fixed edge then τ' is disjoint from the interior of E_1' and so is contained in ω. If τ is an indivisible Nielsen path then it has the form $E_2 {w_2}^p \overline{E}_2$ for some linear edge E_2 with twist path w_2. There is an induced decomposition $\tau' = X' w_2' Y'$, where X', w_2' and Y' project to E_2, w_2^p and $\overline{E}_2$, respectively. Since $E_2 \neq \overline{E}_1$, we have $X' \subset \omega$ and the initial vertex of ${w_2'}^p$ is contained in ω. Since w_2 has height less than E_2 and so height less than τ, ${w_2'}^p \subset \omega$. Finally, since X' is contained in ω, E_2 has height less than E_1 and in particular, $\overline{E}_2 \neq \overline{E}_1$. Thus $Y' \subset \omega$. This completes the proof of the claim.

Each $\sigma = \sigma_a$ as in the statement of the lemma has a cyclic splitting $\sigma = \sigma_1 \cdot \ldots \cdot \sigma_m$ into fixed edges and indivisible Nielsen paths σ_i. The above claim shows that if $\sigma' = \sigma_1' \cdot \ldots \cdot \sigma_m'$ is a lift to $\Gamma(f)$ in which an endpoint of some σ_i' is not contained in $\Gamma(f)^0$ then σ' is entirely contained in the loop ω associated to one of the lollipops.

To complete the proof we need only show each σ has a unique lift in which the endpoints of each σ_i lift into $\Gamma^0(f)$. Since the vertices of G have unique lifts into $\Gamma^0(f)$, it suffices to show that each σ_i has a unique lift with endpoints in $\Gamma^0(f)$ and this is immediate from the construction of $\Gamma^0(f)$. □

4.2 Strong axes and twist coordinates

The following lemma describes the extent to which fixed subgroups fail to be malnormal.

Lemma 4.4 *For distinct automorphisms Φ_1 and Φ_2 representing the same outer automorphism and for any $c \in F_n$:*

(1) $\mathsf{Fix}(\Phi_1) \cap \mathsf{Fix}(\Phi_2)$ *is either trivial or a maximal cyclic subgroup.*

(2) *If $c \notin \mathsf{Fix}(\Phi_1)$ then* $\mathsf{Fix}(\Phi_1) \cap (\mathsf{Fix}(\Phi_1))^c = \mathsf{Fix}(\Phi_1) \cap \mathsf{Fix}(i_c \Phi_1 i_c^{-1})$ *is either trivial or a maximal cyclic subgroup.*

(3) $\mathsf{Fix}(\Phi_1)$ *is its own normalizer.*

Proof If Φ_1 and Φ_2 are distinct automorphisms representing the same outer automorphism then $\Phi_1^{-1}\Phi_2$ is a nontrivial inner automorphism and $\mathsf{Fix}(\Phi_1) \cap \mathsf{Fix}(\Phi_2)$ is a subgroup of the cyclic group $\mathsf{Fix}(\Phi_1^{-1}\Phi_2)$. Maximality of $\mathsf{Fix}(\Phi_1) \cap \mathsf{Fix}(\Phi_2)$ follows from Lemma 3.6 and the fact that $(a^k)^{\pm} = a^{\pm}$ for all nontrivial $a \in F_n$ and all $k \geq 1$. This proves (1).

For (2) note that if $c \notin \mathsf{Fix}(\Phi_1)$ then $i_c \Phi_1 i_c^{-1} = i_{c\Phi_1(c^{-1})}\Phi_1 \neq \Phi_1$. Note also that $\mathsf{Fix}(i_c \Phi_1 i_c^{-1}) = (\mathsf{Fix}(\Phi_1))^c$. Item (2) therefore follows from (1) applied with $\Phi_2 = i_c \Phi_1 i_c^{-1}$. In proving (3) we may assume by (2) that $\mathsf{Fix}(\Phi_1) = \langle a \rangle$ for some root-free $a \in F_n$ and in this case (3) is obvious. □

The conjugacy class of a cyclic subgroup is determined by the conjugacy class of either of its generators. As we have no way to canonically choose a generator, we work, for now, with unoriented conjugacy classes. The following definition appeared as [Bestvina et al. 2004, Definition 4.6] under slightly different hypotheses and in the paragraph before [Feighn and Handel 2011, Remark 4.39] in the CT context.

Definition 4.5 Elements $a, b \in F_n$ are in the same *unoriented conjugacy class* if $a = i_c(b)$ or $a = i_c(b^{-1})$ for some $c \in F_n$. An unoriented conjugacy class $[a]_u$ of a nontrivial root-free $a \in F_n$ is an *axis* for ϕ if $\langle a \rangle = \mathsf{Fix}(\Phi_1) \cap \mathsf{Fix}(\Phi_2)$ for distinct $\Phi_1, \Phi_2 \in \mathcal{P}(\phi)$. The *multiplicity of an axis* $[a]_u$ is the number of distinct $\Phi_i \in \mathcal{P}(\phi)$ that fix a. The *set of axes for* ϕ is denoted by $\mathcal{A}(\phi)$. The set $\{[a] \mid [a]_u \in \mathcal{A}(\phi)\}$ is denoted by $\mathcal{A}_{\mathrm{or}}(\phi)$.

There is a very useful description of $\mathcal{A}(\phi)$ in terms of a CT $f : G \to G$.

Lemma 4.6 *If $f : G \to G$ is a* CT *representing ϕ and $\{w_i\}$ is the set of twist paths for f, then $\mathcal{A}(\phi) = \{[w_i]_u\}$. In particular, $\mathcal{A}(\phi)$ is finite.*

Proof This follows from [Feighn and Handel 2011, Lemma 4.40]. □

Notation 4.7 If $[a]_u = [w]_u$ for some twist path w then, up to a reversal of orientation, the axis of the covering translation $T_a : \widetilde{G} \to \widetilde{G}$ can be viewed as an infinite concatenation $\ldots \widetilde{w}_{-1}\widetilde{w}_0\widetilde{w}_1 \ldots$ of paths that project to w. There is a principal lift $\widetilde{f}_{a,0} : \widetilde{G} \to \widetilde{G}$, called the *base principal lift for* a, that fixes the endpoints of each $\widetilde{w}_l$. The principal automorphism $\Phi_{a,0}$ corresponding to $\widetilde{f}_{a,0}$ is called the *base principal automorphism* for a. If a is an element of some basis for F_n then the base principal lift for a depends on the choice of $f : G \to G$, and not just on ϕ.

For each edge $E^j \in \mathsf{Lin}_w(f)$, there is a principal lift $\widetilde{f}_{a,j} : \widetilde{G} \to \widetilde{G}$ that fixes the initial endpoint of each lift $\widetilde{E}^j$ with terminal endpoint equal to the initial endpoint of some $\widetilde{w}_l$. (We write E^j rather than E_j to emphasize that j is not an indicator of height in G.) The principal automorphism corresponding to $\widetilde{f}_{a,j}$ is denoted by $\Phi_{a,j}$. Note that $\Phi_{a,0} = \Phi_{a^{-1},0}$ and $\Phi_{a,j} = \Phi_{a^{-1},j}$. Further details can be found in [Feighn and Handel 2011, Lemma 4.40] and the paragraph that precedes it.

Lemma 4.8 *Suppose that $f : G \to G$ is a* CT *representing ϕ, that w is a twist path for f and that $a \in F_n$ satisfies $[a]_u = [w]_u$. Suppose also that $E^1, \dots, E^{m-1}$ are the edges in $\mathsf{Lin}_w(f)$.*

(1) *$\{\Phi_{a,0}, \Phi_{a,1}, \dots, \Phi_{a,m-1}\}$ is the set of principal automorphisms that fix a. In particular, the multiplicity of each element of $\mathcal{A}(\phi)$ is finite.*

(2) *If $f(E^j) = E^j w^{d_j}$ then $\Phi_{a,j} = i_a^{d_j} \Phi_{a,0}$ if $[a] = [w]$, and $\Phi_{a,j} = i_a^{-d_j} \Phi_{a,0}$ if $[a] = [\overline{w}]$.*

Proof This follows from Lemma 4.40 of [Feighn and Handel 2011]. □

Definition 4.9 Suppose that the group G acts on the sets X_i for $i = 1, \dots, k$, and that $x_i \in X_i$. The orbit of $(x_1, \dots, x_k)$ under the diagonal action of G on $\prod_{i=1}^k X_i$, denoted by $[x_1, \dots, x_k]_G$, is a *conjugacy k-tuple*. If $k = 2$ then we say $[x_1, x_2]_G$ is a *conjugacy pair*. We sometimes suppress the subscript, in which case $G = F_n$.

Examples 4.10 Here are some examples of conjugacy pairs where $G = F_n$.

- We will often take X_i to be the set of finitely generated subgroups of F_n or F_n itself with the action of F_n given by conjugation. If $H < F_n$ (resp. $x \in F_n$) then $[H]_{F_n}$ (resp. $[x]_{F_n}$) is the conjugacy class of H in F_n (resp. x in F_n). Conjugacy pairs formed with these X_i will play an important role in this paper, especially in Section 10.3.
- If $X = \partial F_n$ and if $x \neq y \in X$, then $(x, y) \in X \times X$ is an oriented line. The conjugacy pair $[x, y]_{F_n}$ is represents an oriented line in any marked graph.
- If X is the power set of ∂F_n and A and B are disjoint subsets of ∂F_n, then $(A, B) \in X \times X$ denotes the set of lines L with $\partial_- L \in A$ and $\partial_+ L \in B$. The conjugacy pair $[A, B]_{F_n}$ represents a set of oriented lines in any marked graph.

We now define strong axes, the first of our invariants that is expressed as a conjugacy pair.

Definition 4.11 Let $\mathcal{A}_{\mathrm{or}}(\phi)$ be the set of conjugacy classes representing elements of $\mathcal{A}(\phi)$, ie $[a] \in \mathcal{A}_{\mathrm{or}}(\phi)$ if $\{[a], [a^{-1}]\} \in \mathcal{A}(\phi)$. F_n acts on pairs (Φ, a) where $\Phi \in \mathcal{P}(\phi)$, $a \in \mathsf{Fix}(\Phi)$, and $[a] \in \mathcal{A}_{\mathrm{or}}(\phi)$ via $(\Phi, a)^g = (\Phi^{i_g}, a^g)$. The F_n-orbit, equivalently conjugacy pair, $[\Phi, a]$, is a *strong axis for ϕ*. If $\alpha_s = [\Phi, a]$ then we let $\alpha_s^{-1} := [\Phi, a^{-1}]$. The set of all strong axes for ϕ is denoted by $\mathsf{SA}(\phi)$. $\mathsf{Aut}(F_n)$ acts on pairs (Φ, a) by $\Theta \cdot (\Phi, a) = (\Theta\Phi\Theta^{-1}, \Theta(a))$. This descends to an action of $\mathsf{Out}(F_n)$ on $\mathsf{SA}(\phi)$.

We can partition $\mathsf{SA}(\phi)$ according to the second coordinate: for each $\mu \in \mathcal{A}_{\mathrm{or}}(\phi)$ let $\mathsf{SA}(\phi, \mu)$ be the subset of $\mathsf{SA}(\phi)$ consisting of elements in which some, and hence every, representative (Φ, a) satisfies $[a] = \mu$.

Lemma 4.12 *Suppose that $a \in F_n$, that $[a] \in \mathcal{A}_{\mathrm{or}}(\phi)$ and that $\Phi_{a,0}, \dots, \Phi_{a,m-1}$ are as in Notation 4.7. For each $\alpha_s \in \mathsf{SA}(\phi, [a])$, there is a unique $\Phi_{a,j}$ such that $\alpha_s = [\Phi_{a,j}, a]$. Thus*

$$\mathsf{SA}(\phi, [a]) = \{[\Phi_{a,0}, a], [\Phi_{a,1}, a], \dots, [\Phi_{a,m-1}, a]\}.$$

Proof Each $\alpha_s \in \mathsf{SA}(\phi, [a])$ is represented by (Φ, a^c) and hence by $(i_c \Phi i_c^{-1}, a)$, for some $c \in F_n$ and some $\Phi \in \mathcal{P}(\phi)$. Since $a \in \mathsf{Fix}(i_c \Phi i_c^{-1})$, there exists j such that $i_c \Phi i_c^{-1} = \Phi_j$.

For uniqueness, note that if $(\Phi_j, a) = c \cdot (\Phi_i, a)$ for some $c \in F_n$ then $c = a^p$ for some p so i_c commutes with Φ_j and $c \cdot (\Phi_i, a) = (\Phi_i, a)$. □

Remark 4.13 There is another useful description of $[\Phi_{a,j}, a] \in \mathsf{SA}(\phi, [a])$ in terms of a CT $f : G \to G$. Let w be the twist path satisfying $[a]_u = [w]_u$ and let v be the initial vertex of w. There is an automorphism $f_{v\#} \colon \pi_1(G, v) \to \pi_1(G, v)$ that sends the homotopy class of the closed path σ with basepoint v to the homotopy class of the closed path $f(\sigma)$ with basepoint v. Let τ be the element of $\pi_1(G, v)$ determined by w if $[a] = [w]$ and by $\overline{w}$ if $[a] = [\overline{w}]$. In both cases, τ is fixed by $f_{v\#}$. There is an isomorphism from $\pi_1(G, v)$ to F_n that is well-defined up to postcomposition with an inner automorphism of F_n. The pair $(f_{v\#}, \tau)$ determines a well-defined element, namely $[\Phi_{a,0}, a]$, of $\mathsf{SA}(\phi, [a])$. Similarly if v_j is the initial endpoint of $E^j \in \mathsf{Lin}_w(f)$, let τ_j be the element of $\pi_1(G, v_j)$ determined by $E^j w \overline{E}^j$ if $[a] = [w]$ and by $E^j \overline{w} \overline{E}^j$ if $[a] = [\overline{w}]$. Then $(f_{v_j\#}, \tau_j)$ determines $[\Phi_{a,j}, a]$.

Continuing with this notation, we can relate $\mathsf{SA}(\phi, [a])$ to circuits in the eigengraph $\Gamma(f)$ that are lifts of $[w]_u$. For $j \neq 0$, $[\Phi_{a,j}, a]$ corresponds to the loop at the end of the lollipop in $\Gamma(f)$ determined by E^j. By Lemma 4.3 there is one more lift of $[w]_u$ into $\Gamma(f)$, and this corresponds to $[\Phi_{a,0}, a]$.

Definition 4.14 Suppose that $\mu \in \mathcal{A}_{\mathrm{or}}(\phi)$ and that $\alpha_s, \alpha'_s \in \mathsf{SA}(\phi, \mu)$. Choose $a \in F_n$ such that $[a] = \mu$ and let $\Phi, \Phi' \in \mathcal{P}(\phi)$ be the unique elements such that $\alpha_s = [\Phi, a]$ and $\alpha'_s = [\Phi', a]$. Since Φ and Φ' both fix a there exists $\tau \in \mathbb{Z}$ such that $\Phi' = i_a^\tau \Phi$; equivalently, $\Phi' \Phi^{-1} = i_a^\tau$. We say that $\tau = \tau(\alpha'_s, \alpha_s)$ is the *twist coordinate* associated to α'_s and α_s.

Example 3.1 (continued) In our example, $\mathsf{SA}(\phi, [a])$ is represented in Figure 2 by the three circles α, α', and α'' labeled a and drawn with thicker lines. $\mathsf{SA}(\phi) = \mathsf{SA}(\phi, [a]) \cup \mathsf{SA}(\phi, [a^{-1}])$. We have for example $\tau(\alpha', \alpha) = 1$.

Lemma 4.15 *Twist coordinates are well-defined.*

Proof We have to show that $\tau(\alpha'_s, \alpha_s)$ is independent of the choice of a representing μ. If a is replaced by a^c then Φ and Φ' are replaced by $i_c \Phi i_c^{-1}$ and $i_c \Phi' i_c^{-1}$, respectively, and so $i_a^\tau = \Phi' \Phi^{-1}$ is replaced by $i_c \Phi' \Phi^{-1} i_c^{-1} = i_c i_a^\tau i_c^{-1} = i_{a^c}{}^\tau$. □

The following lemma allows us to compute twist coordinates for strong axes from a CT $f : G \to G$. It is an immediate consequence of Lemma 4.8 and the definitions.

Lemma 4.16 (1) *If $[a] = [w]$ and $E^j \in \mathsf{Lin}_w(f)$ satisfies $f(E^j) = E^j w^{d_j}$, then*

$$\tau([\Phi_{a,j}, a], [\Phi_{a,0}, a]) = d_j.$$

(2) *Suppose that* $\mu \in \mathcal{A}_{\mathrm{or}}(\phi)$ *and that* $\alpha_s, \beta_s, \gamma_s \in \mathsf{SA}(\phi, \mu)$. *Then*

(a) $\tau(\alpha_s, \beta_s) = -\tau(\beta_s, \alpha_s)$,

(b) $\tau(\alpha_s, \gamma_s) = \tau(\alpha_s, \beta_s) + \tau(\beta_s, \gamma_s)$,

(c) $\tau(\alpha_s, \beta_s) = -\tau(\alpha_s^{-1}, \beta_s^{-1})$.

The next lemma shows that $\mathcal{A}(\phi)$, $\mathsf{SA}(\phi, [a])$ and $\mathsf{SA}(\phi)$ are natural invariants.

Lemma 4.17 *Assume that* $\psi = \theta\phi\theta^{-1}$ *and that* Θ *represents* θ.

(1) *The correspondence* $[a]_u \leftrightarrow (\theta[a])_u$ *defines a bijection* $\mathcal{A}(\phi) \leftrightarrow \mathcal{A}(\psi)$.

(2) *The correspondence* $(\Phi, a) \leftrightarrow (\Theta\Phi\Theta^{-1}, \Theta(a))$ *induces a bijection* $\mathsf{SA}(\phi, [a]) \leftrightarrow \mathsf{SA}(\psi, \theta([a]))$ *that preserves twist coordinates.*

Proof If $\Phi_1, \Phi_2 \in \mathcal{P}(\phi)$ fix $a \in F_n$, then $\Psi_1 := \Theta\Phi_1\Theta^{-1}$ and $\Psi_2 := \Theta\Phi_2\Theta^{-1}$ in $\mathcal{P}(\psi)$ fix $\Theta(a)$. This proves (1).

For (2), let $\Psi = \Theta\Phi\Theta^{-1}$ and note that if $\Phi \in \mathcal{P}(\phi)$ and $a \in \mathsf{Fix}(\Phi)$ then $\Psi \in \mathcal{P}(\psi)$ and $\Theta(a) \in \mathsf{Fix}(\Psi)$. Moreover,

$$c \cdot (\Phi, a) = (i_c \Phi i_c^{-1}, a^c) \mapsto (i_{\Theta(c)} \Psi i_{\Theta(c)}^{-1}, \Theta(a)^{\Theta(c)}) = \Theta(c) \cdot (\Psi, \Theta(a)).$$

This proves that $(\Phi, a) \mapsto (\Theta\Phi\Theta^{-1}, \Theta(a))$ induces a well-defined map $\mathsf{SA}(\phi, [a]) \to \mathsf{SA}(\psi, \theta([a]))$ that is obviously invertible and is hence a bijection. If Φ_i and Ψ_i are as in the proof of (1) and if $\Phi_2 = i_a^\tau \Phi_1$, then $\Psi_2 = i_{\Theta(a)}^\tau \Psi_1$. This proves that twist coordinates are preserved. □

We conclude this section with a conjugacy class of pairs construction that is better suited to the techniques in Section 10 than the one in Definition 4.11 but is only applicable when the fixed subgroups in question have rank at least two.

Definition 4.18 Given $\phi \in \mathsf{Out}(F_n)$, consider pairs $(\mathsf{Fix}(\Phi), a)$ where $\Phi \in \mathcal{P}(\phi)$, $a \in \mathsf{Fix}(\Phi)$ and $[a] \in \mathcal{A}_{\mathrm{or}}(\phi)$. Using $\mathsf{Fix}(i_c \Phi i_c^{-1}) = i_c(\mathsf{Fix}(\Phi))$, the action of F_n on such pairs is given by $c \cdot (\mathsf{Fix}(\Phi), a) = (i_c(\mathsf{Fix}(\Phi)), i_c(a))$, giving a conjugacy pair $[\mathsf{Fix}(\Phi), a]$. Similarly, $\mathsf{Aut}(F_n)$ acts on pairs $(\mathsf{Fix}(\Phi), a)$ by $\Theta \cdot (\mathsf{Fix}(\Phi), a) = (\Theta(\mathsf{Fix}(\Phi)), \Theta(a))$. This descends to an action of $\mathsf{Out}(F_n)$ on the set of such conjugacy pairs.

Remark 4.19 Since $\mathsf{Fix}(\Phi)$ is its own normalizer (Lemma 4.4(3)), $[\mathsf{Fix}(\Phi), a] = [\mathsf{Fix}(\Phi), a']$ if and only if $a' = i_c(a)$ for some $c \in \mathsf{Fix}(\Phi)$; equivalently, a and a' are conjugate as elements of $\mathsf{Fix}(\Phi)$.

Lemma 4.20 *Suppose that* $\mathsf{Fix}(\Phi)$ *and* $\mathsf{Fix}(\Phi')$ *have rank at least two. Then*

$$[\Phi, a] = [\Phi', a'] \iff [\mathsf{Fix}(\Phi), a] = [\mathsf{Fix}(\Phi'), a'].$$

Proof By definition, $[\Phi, a] = [\Phi', a']$ if and only if there exists $c \in F_n$ such that

$$\Phi' = i_c \Phi i_c^{-1} \quad \text{and} \quad a' = i_c(a).$$

Similarly, $[\mathsf{Fix}(\Phi), a] = [\mathsf{Fix}(\Phi'), a']$ if and only if there exists $c \in F_n$ such that $\mathsf{Fix}(\Phi') = i_c \mathsf{Fix}(\Phi)$ and $a' = i_c(a)$. As we are assuming that $\mathsf{Fix}(\Phi)$ and $\mathsf{Fix}(\Phi')$ have rank at least two, $\Phi' = i_c \Phi i_c^{-1}$ if and only if $\mathsf{Fix}(\Phi') = i_c \mathsf{Fix}(\Phi)$. □

4.3 Applying the recognition theorem

The recognition theorem [Feighn and Handel 2011, Theorem 5.1] gives invariants that completely determine rotationless elements of $\mathsf{Out}(F_n)$. In this paper, via the following lemma, we use it to give a sufficient condition for two elements of $\mathsf{UPG}(F_n)$ to be conjugate in $\mathsf{Out}(F_n)$.

Lemma 4.21 *Suppose that $f: G \to G$ and $g: G' \to G'$ are* CT*s representing ϕ and ψ respectively, that $\theta \in \mathsf{Out}(F_n)$ and that a line L lifts into $\Gamma(f)$ (meaning that the realization of L in G is the image of a line in $\Gamma(f)$) if and only if $\theta(L)$ lifts into $\Gamma(g)$. Then for each $\Theta \in \mathsf{Aut}(F_n)$ representing θ:*

1. *There is a bijection $B_{\mathcal{P}}: \mathcal{P}(\phi) \to \mathcal{P}(\psi)$ such that $\mathsf{Fix}_\mathsf{N}(B_{\mathcal{P}}(\Phi)) = \Theta(\mathsf{Fix}_\mathsf{N}(\Phi)) = \mathsf{Fix}_\mathsf{N}(\Phi^\Theta)$. In particular, $\mathsf{Fix}(B_{\mathcal{P}}(\Phi)) = \Theta\, \mathsf{Fix}(\Phi) = \mathsf{Fix}(\Phi^\Theta)$.*
2. *The map $[\Phi, a] \mapsto [B_{\mathcal{P}}(\Phi), \Theta(a)]$ defines a bijection $B_{\mathsf{SA}}: \mathsf{SA}(\phi) \to \mathsf{SA}(\psi)$, independent of the choice of Θ, such that $\phi^\theta = \psi$ if and only if B_{SA} preserves twist coordinates.*

Proof Given $\Phi \in \mathcal{P}(\phi)$, choose a line $\widetilde{L}_1 \subset \widetilde{G}$ with both ends nonperiodic and both ends in $\mathsf{Fix}_\mathsf{N}(\Phi)$. (This is possible by Remark 3.10.) By Lemma 4.1, the projection $L \subset G$ lifts to the component $\Gamma_{[\Phi]}(f)$ of $\Gamma(f)$ that corresponds to $[\Phi]$. By hypothesis, the line $L_1' \subset G'$ corresponding to $\theta(L)$ lifts to a component of $\Gamma(g)$ and so by a second application of Lemma 4.1 there is a unique $\Psi \in \mathcal{P}(\psi)$ such that $\mathsf{Fix}_\mathsf{N}(\Psi)$ contains the endpoints $\{\Theta(\partial_\pm \widetilde{L}_1)\}$ of $\Theta(\widetilde{L}_1)$; moreover, L_1' lifts into $\Gamma_{[\Psi]}(g)$. To see that Ψ is independent of the choice of $\widetilde{L}_1$, suppose that we are given some other $\widetilde{L}_2$ with both ends nonperiodic and both ends in $\mathsf{Fix}_\mathsf{N}(\Phi)$. Let $\widetilde{L}_3$ be the line connecting the terminal endpoint of $\widetilde{L}_1$ to the initial endpoint of $\widetilde{L}_2$. Since $\widetilde{L}_1$ and $\widetilde{L}_3$ have a common endpoint, replacing $\widetilde{L}_1$ with $\widetilde{L}_3$ does not change Ψ. For the same reason, replacing $\widetilde{L}_3$ with $\widetilde{L}_2$ does not change Ψ. We conclude that $B_{\mathcal{P}}(\Phi) = \Psi$ is well-defined. This argument also shows that Θ maps each nonperiodic element of $\mathsf{Fix}_\mathsf{N}(\Phi)$ to a nonperiodic element of $\mathsf{Fix}_\mathsf{N}(\Psi)$. Since nonperiodic points in $\mathsf{Fix}_\mathsf{N}(\Phi)$ are dense in $\mathsf{Fix}_\mathsf{N}(\Phi)$, $\Theta(\mathsf{Fix}_\mathsf{N}(\Phi)) \subset \mathsf{Fix}_\mathsf{N}(\Psi)$. Reversing the roles of ϕ and ψ and replacing θ with θ^{-1}, we see that $\Theta^{-1}(\mathsf{Fix}_\mathsf{N}(\Psi)) \subset \mathsf{Fix}_\mathsf{N}(\Phi)$, which completes the proof of (1). Note that if $\Psi = B_{\mathcal{P}}(\Phi)$ then for all $c \in F_n$,

$$B_{\mathcal{P}}(i_c \Phi i_c^{-1}) = i_{\Theta(c)} \Psi i_{\Theta(c)}^{-1}$$

because

$$\Theta(\mathsf{Fix}_\mathsf{N}(i_c \Phi i_c^{-1})) = \Theta(i_c\, \mathsf{Fix}_\mathsf{N}(\Phi)) = i_{\Theta(c)} \Theta(\mathsf{Fix}_\mathsf{N}(\Phi)) = i_{\Theta(c)}\, \mathsf{Fix}_\mathsf{N}(\Psi) = \mathsf{Fix}_\mathsf{N}(i_{\Theta(c)} \Psi i_{\Theta(c)}^{-1}).$$

For (2), suppose that $[a] \in \mathcal{A}_{\text{or}}(\phi)$, that $\Phi \in \mathcal{P}(\phi)$ fixes a and that $\Psi = B_{\mathcal{P}}(\Phi)$. Define

$$B_{\mathsf{SA}}([\Phi, a]) = [B_{\mathcal{P}}(\Phi), \Theta(a)] = [\Psi, \Theta(a)].$$

Then for all $c \in F_n$,

$$B_{\mathsf{SA}}([\Phi, a]^c) = B_{\mathsf{SA}}([i_c \Phi i_c^{-1}, i_c(a)]) = [i_{\Theta(c)} \Psi i_{\Theta(c)}^{-1}, i_{\Theta(c)}(\Theta(a))] = [\Psi, a]^{\Theta(c)},$$

so B_{SA} is well defined. By symmetry, B_{SA} is a bijection. If Θ is replaced by $i_b\Theta$ for some $b \in F_n$ then $(\Psi, \Theta(a))$ is replaced by $(i_b \Psi i_b^{-1}, i_b\Theta(a)) = b \cdot (\Psi, \Theta(a))$. This shows that B_{SA} is independent of the choice of Θ. It remains to show that $\phi^\theta = \psi$ if and only if B_{SA} preserves twist coordinates.

Let $\upsilon = \phi^\theta$. By Lemmas 3.16 and 4.17, conjugation by Θ induces

- a bijection $B'_{\mathcal{P}} \colon \mathcal{P}(\phi) \to \mathcal{P}(\upsilon)$ defined by $\Phi \mapsto \Theta\Phi\Theta^{-1}$ and satisfying $\mathsf{Fix}_{\mathsf{N}}(B'_{\mathcal{P}}(\Phi)) = \Theta\, \mathsf{Fix}_{\mathsf{N}}(\Phi)$,
- a bijection $B'_{\mathsf{SA}} \colon \mathsf{SA}(\phi) \to \mathsf{SA}(\upsilon)$ defined by $[\Phi, a] \mapsto [\Theta\Phi\Theta^{-1}, \Theta(a)]$ that preserves twist coordinates.

The bijections $B''_{\mathcal{P}} = B_{\mathcal{P}} {B'}_{\mathcal{P}}^{-1} \colon \mathcal{P}(\upsilon) \to \mathcal{P}(\psi)$ and $B''_{\mathsf{SA}} = B_{\mathsf{SA}} {B'}_{\mathsf{SA}}^{-1} \colon \mathsf{SA}(\upsilon) \to \mathsf{SA}(\psi)$ satisfy:

(a) $\mathsf{Fix}_{\mathsf{N}}(B''_{\mathcal{P}}(\Upsilon)) = \mathsf{Fix}_{\mathsf{N}}(\Upsilon)$ for all $\Upsilon \in \mathcal{P}(\upsilon)$.

(b) B''_{SA} preserves twist coordinates if and only if B_{SA} does.

Applying (b), it suffices to show that $\phi^\theta = \psi$ if and only if B''_{SA} preserves twist coordinates.

Suppose that $[b] \in \mathcal{A}_{\text{or}}(\upsilon)$, that $b \in \mathsf{Fix}(\Upsilon)$ and that $\Upsilon, i_{b^d}\Upsilon \in \mathcal{P}(\upsilon)$. Let $a = \Theta^{-1}(b)$ and $\Phi = \Theta^{-1}\Upsilon\Theta$. Then

$$B''_{\mathsf{SA}}[\Upsilon, b] = B_{\mathsf{SA}}[\Phi, a] = [B_{\mathcal{P}}(\Phi), b],$$

and likewise

$$B''_{\mathsf{SA}}[i_b^d\Upsilon, b] = B_{\mathsf{SA}}[i_a^d\Phi, a] = [B_{\mathcal{P}}(i_a^d\Phi), b].$$

By definition, the twist coordinate for $[i_{b^d}\Upsilon, b]$ and $[\Upsilon, b]$ is d. It follows that B''_{SA} preserves twist coordinates if and only if

$$B_{\mathcal{P}}(i_a^d\Phi) = i_{b^d} B_{\mathcal{P}}(\Phi).$$

Since

$$B_{\mathcal{P}}(i_a^d\Phi) = B''_{\mathcal{P}}(i_b^d\Upsilon) \quad \text{and} \quad B_{\mathcal{P}}(\Phi) = B''_{\mathcal{P}}(\Upsilon),$$

we conclude that B''_{SA} preserves twist coordinates if and only if

(c) $B''_{\mathcal{P}}(i_b^d\Upsilon) = i_b^d B''_{\mathcal{P}}(\Upsilon)$.

By the recognition theorem [Feighn and Handel 2011, Theorem 5.3], (a) and (c) are equivalent to $\upsilon = \psi$. □

5 Limit lines $\Omega(r) \subset \mathcal{B}$

Each point $P \in \partial F_n$ determines a closed set of lines; see eg [Feighn and Handel 2011, Section 2.4], where the closed set of lines is called the accumulation set of P. In this section we focus on the case that $P \in \mathcal{R}(\phi)$ and analyze these lines using CTs.

Definition 5.1 For each $r \in \partial F_n/F_n$, we define the set $\Omega(r) \subset \mathcal{B}$ of *limit lines of* r as follows. Choose a lift $\tilde{r} \in \partial F_n$, a marked graph K and a ray $\tilde{R} \subset \tilde{K}$ with terminal end $\tilde{r}$. Let $R \subset K$ be the projected image of $\tilde{R}$. Then $L \in \Omega(r)$ (thought of as a line in K) if and only if the following equivalent conditions are satisfied.

(1) Each finite subpath of L occurs as a subpath of R.

(2) For each lift $\tilde{L} \subset \tilde{K}$ of $L \subset K$ there are translates $\tilde{R}_j$ of $\tilde{R}$ such that the initial endpoints of $\tilde{R}_j$ converge to the initial endpoint of $\tilde{L}$ and the terminal endpoints of $\tilde{R}_j$ converge to the terminal endpoint of $\tilde{L}$.

Let $\Omega_{\mathsf{NP}}(r)$ be the set of nonperiodic elements of $\Omega(r)$.

Lemma 5.2 *$\Omega(r)$ and $\Omega_{\mathsf{NP}}(r)$ are well-defined. Moreover, for each $\theta \in \mathsf{Out}(F_n)$, $\theta(\Omega(r)) = \Omega(\theta(r))$ and $\theta(\Omega_{\mathsf{NP}}(r)) = \Omega_{\mathsf{NP}}(\theta(r))$.*

Proof If R' is another ray with terminal end r, then R and R' have a common terminal subray R''. Let $R = \alpha R''$ and $R' = \alpha' R''$. Given a finite subpath $\tau_2 \subset K$ of a line ℓ, extend it to a finite subpath $\tau_1\tau_2 \subset K$ of ℓ, where τ_1 is longer than both α and α'. If $\tau_1\tau_2$ occurs in R then τ_2 occurs in R''. Since τ_2 was arbitrary, every finite subpath of ℓ occur in R if and only if every finite subpath of ℓ occurs in R''. The same holds for R' and R''. This proves that $\Omega(r)$ is independent of the choice of R. Independence of the choice of $\tilde{R}$ is obvious, as is the equivalence of (1) and (2).

Suppose that K' is another marked graph and that $g\colon K \to K'$ is a homotopy equivalence that preserves markings and so represents the identity outer automorphism. Let $\tilde{g}\colon \tilde{K} \to \tilde{K}'$ be a lift of g. If $\tilde{L} \subset \tilde{K}$ is a lift of L and $\tilde{R}_j \subset \tilde{K}$ is a sequence of translates of ray $\tilde{R}$ such that the initial and terminal endpoints of $\tilde{R}_j$ converge to those of $\tilde{L}$, then the same is true of $\tilde{L}' = \tilde{g}_\#(\tilde{L}) \subset \tilde{K}'$ and $\tilde{R}'_j = \tilde{g}_\#(\tilde{R}_j) \subset \tilde{K}'$. This proves that $\Omega(r)$ is independent of the choice of K.

For the moreover statement, choose a homotopy equivalence $h\colon K \to K$ that represents θ and lifts $\tilde{L} \subset \tilde{K}$ and $\tilde{h}\colon \tilde{K} \to \tilde{K}$. If $\tilde{R}_j \subset \tilde{K}$ is a sequence of translates of $\tilde{R}$ whose initial and terminal endpoints converge to those of $\tilde{L}$, then the initial and terminal endpoints of $\tilde{h}_\#(\tilde{R}_j)$ converge to those of $\tilde{h}_\#(\tilde{L}) \subset \tilde{K}$. This proves that $\theta(\Omega_{\mathsf{NP}}(r)) \subset \Omega_{\mathsf{NP}}(\theta(r))$. The reverse inclusion follows by symmetry. □

We now specialize to $r \in \mathcal{R}(\phi)$.

Notation 5.3 For $\phi \in \mathsf{UPG}(F_n)$, let

$$\Omega(\phi) = \bigcup_{r \in \mathcal{R}(\phi)} \Omega(r) \quad \text{and} \quad \Omega_{\mathsf{NP}}(\phi) = \bigcup_{r \in \mathcal{R}(\phi)} \Omega_{\mathsf{NP}}(r).$$

As an immediate consequence of Lemma 3.16 and the moreover statement of Lemma 5.2 we have:

Corollary 5.4 *Suppose that $\theta \in \mathsf{Out}(F_n)$ and that $\psi = \theta\phi\theta^{-1} \in \mathsf{UPG}(F_n)$. Then $\theta(\Omega(\phi)) = \Omega(\psi)$ and $\theta(\Omega_{\mathsf{NP}}(\phi)) = \Omega_{\mathsf{NP}}(\psi)$.* □

For the remainder of the section we assume that $f : G \to G$ is a CT *representing $\phi \in \mathsf{UPG}(F_n)$.* Our goal is to describe $\Omega(\phi)$ and $\Omega_{\mathsf{NP}}(\phi)$ in terms of $f : G \to G$. See in particular Corollary 5.17.

One advantage of working in a CT is that we can work with finite paths and not just with lines and rays.

Definition 5.5 Given a path $\sigma \subset G$, we say that a line $L \subset G$ is contained in the *accumulation set* $\mathsf{Acc}(\sigma)$ *of σ with respect to f* if every finite subpath of L occurs as a subpath of $f_\#^k(\sigma)$ for arbitrarily large k.

Notation 5.6 For each twist path w, we write w^∞ for both the ray that is an infinite concatenation of copies of w and the line that is a bi-infinite concatenation of copies of w, using context to distinguish between the two. We use either $\bar{w}^\infty$ or $w^{-\infty}$ for the ray or line obtained from w^∞ by reversing orientation on w.

Examples 5.7 (1) If σ is a Nielsen path, then $\mathsf{Acc}(\sigma) = \varnothing$.

(2) Suppose that $E \in \mathsf{Lin}(f)$ and $f(E) = Ew^d$.
   (a) If $d > 0$ then $\mathsf{Acc}(E) = \{w^\infty\}$ and $\mathsf{Acc}(\bar{E}) = \{\bar{w}^\infty\}$.
   (b) If $d < 0$ then $\mathsf{Acc}(E) = \{\bar{w}^\infty\}$ and $\mathsf{Acc}(\bar{E}) = \{w^\infty\}$.

(3) If $E_i, E_j \in \mathsf{Lin}(f)$ satisfy $f(E_i) = E_i w^{d_i}$ and $f(E_j) = E_j w^{d_j}$ for $d_i \neq d_j$ then for all $p \in \mathbb{Z}$, $\mathsf{Acc}(E_i w^p \bar{E}_j) = \{w^\infty\}$ if $d_i > d_j$, and $\mathsf{Acc}(E_i w^p \bar{E}_j) = \{\bar{w}^\infty\}$ if $d_i < d_j$.

Recall from Lemma 3.26 that there is a bijection between $\mathcal{R}(\phi)$ and the set $\mathcal{E}_f$ of nonfixed nonlinear edges of G and that if $r \in \mathcal{R}(\phi)$ corresponds to $E \in \mathcal{E}_f$ then the eigenray $R_E = E \cdot u_E \cdot [f(u_E)] \cdot [f^2(u_E)] \cdot \ldots \subset G$ has terminal end r. Thus, a line $L \subset G$ is an element of $\Omega(r)$ if and only if each finite subpath of L occurs as a subpath of R_E.

Limit lines of eigenrays are connected to accumulation sets as follows.

Lemma 5.8 *If $r \in \mathcal{R}(\phi)$ corresponds to $E \in \mathcal{E}_f$, and $f(E) = E \cdot u_E$, then*

$$\Omega(r) = \mathsf{Acc}(E) = \mathsf{Acc}(u_E \cdot f_\#(u_E)) = \mathsf{Acc}\big(u_E \cdot f_\#(u_E) \cdot \ldots \cdot f_\#^k(u_E)\big)$$

for any $k \geq 1$.

Proof The first equality is an immediate consequence of the definitions and the fact that $E \subset f(E) \subset f_\#^2(E) \subset \cdots$ is an increasing sequence whose union is R_E. Likewise,

$$\mathsf{Acc}(u_E \cdot f_\#(u_E)) \subset \mathsf{Acc}(u_E \cdot f_\#(u_E) \cdot \ldots \cdot f_\#^k(u_E)) \subset \mathsf{Acc}(f_\#^{k+1}(E)) = \mathsf{Acc}(E)$$

is immediate. It therefore suffices to show that $\Omega(r) \subset \mathsf{Acc}(u_E \cdot f_\#(u_E))$.

If $L \in \Omega(r)$ then every finite subpath σ of L occurs as a subpath of every subray of R_E. Since the length of $f_\#^k(u_E)$ tends to infinity with k, each occurrence of σ that is sufficiently far away from the initial endpoint of R_E is contained in some $f_\#^k(u_E) \cdot f_\#^{k+1}(u_E) = f_\#^k(u_E \cdot f_\#(u_E))$. As the occurrence of σ moves farther down the ray, $k \to \infty$. □

Notation 5.9 Define a partial order on the set $\mathcal{E}_f \cup \mathcal{E}_f^{-1}$ by $E_1 \gg E_2$ if $E_1 \neq E_2$ and if, for some $k \geq 0$, E_2 is crossed by $f_\#^k(E_1)$ and so by Lemma 3.21 is a term in the complete splitting of $f_\#^k(E_1)$. (In Notation 6.1 we define a partial order $>$ on $\mathcal{E}_f$ that does not distinguish between E and $\bar{E}$.)

As an immediate consequence of the definition, we have:

Lemma 5.10 *If $E, E' \in \mathcal{E}_f \cup \mathcal{E}_f^{-1}$ and $E \gg E'$, then the height of E' is less than the height of E, and $\mathsf{Acc}(E') \subset \mathsf{Acc}(E)$.* □

The terms μ_i in the complete splitting of $u_E \cdot f_\#(u_E)$ are Nielsen paths, exceptional paths and single edges with height strictly less than that of E. Each $\mathsf{Acc}(\mu_i)$ is a subset of $\mathsf{Acc}(E) = \Omega(r)$. If $\mu_i \in \mathcal{E}_f \cup \mathcal{E}_f^{-1}$ then $\mathsf{Acc}(\mu_i)$ can be understood inductively. The remaining $\mathsf{Acc}(\mu_i)$ are given in Examples 5.7. The work in identifying $\Omega(r) = \mathsf{Acc}(u_E \cdot f_\#(u_E))$ is to determine what additional lines must be added to $\bigcup \mathsf{Acc}(\mu_i)$.

Notation 5.11 For a path $\alpha \subset G$, we say that $f_\#^k(\alpha)$ *converges to a ray* $R \subset G$ if for all m there exists K such that the initial m-length segments of $f_\#^k(\alpha)$ and of R are equal for all $k \geq K$. Note that R is necessarily unique and $f_\#$-invariant. We sometimes write $R = f_\#^\infty(\alpha)$.

Examples 5.12 (1) Suppose that $E \in \mathsf{Lin}(f)$ and that $f(E) = Ew^d$.

- (a) If $d > 0$ then $f_\#^k(E)$ converges to Ew^∞ and $f_\#^k(\bar{E})$ converges to $\bar{w}^\infty$.
- (b) If $d < 0$ then $f_\#^k(E)$ converges to $E\bar{w}^\infty$ and $f_\#^k(\bar{E})$ converges to w^∞.

(2) If $E \in \mathcal{E}_f$ then $f_\#^k(E)$ converges to R_E.

(3) If $E_i, E_j \in \mathsf{Lin}(f)$ satisfy $f(E_i) = E_i w^{d_i}$ and $f(E_j) = E_j w^{d_j}$ for $d_i \neq d_j$ then for all $p \in \mathbb{Z}$, $f_\#^k(E_i w^p \bar{E}_j)$ converges to $E_i w^\infty$ if $d_i > d_j$ and to $E_i \bar{w}^\infty$ if $d_i < d_j$.

Notation 5.13 If $E \in \mathcal{E}_f$, then the first growing term of $f(\bar{E})$ has height less than that of E. It follows that there exists $M > 1$ such that if σ_i is a growing term in the complete splitting of a path σ and if $m \geq M$, then the first growing term in the complete splitting of $f_\#^m(\sigma_i)$ is not an element of $\mathcal{E}_f^{-1}$, and the last growing term in the complete splitting of $f_\#^m(\sigma_i)$ is not an element of $\mathcal{E}_f$. We refer to M as the *stabilization constant* for f.

Lemma 5.14 *Let M be the stabilization constant for f. If σ is a completely split growing path, then $f_\#^k(\sigma)$ converges to a ray $f_\#^\infty(\sigma) = \rho R$, where*

(1) *ρ is a (possibly trivial) Nielsen path and one of the following holds:*

- (a) *$R = R_E$ for some $E \in \mathcal{E}_f$.*
- (b) *$R = Ew^{\pm\infty}$ for some $E \in \mathsf{Lin}_w(f)$.*
- (c) *$R = w^{\pm\infty}$ for some twist path w.*

(2) *If $\sigma = \mu_1 \cdot \nu_1 \cdot \mu_2 \cdot \ldots$ is the coarsening of the complete splitting of σ into maximal (possibly trivial) Nielsen paths μ_i and single growing terms ν_i, then $f_\#^\infty(\sigma) = \mu_1 f_\#^\infty(\nu_1)$.*

(3) *In case (1c) there exists $E \in \mathsf{Lin}_w(f)$ and a smallest $k_\sigma \le M$ such that the first growing term in the coarsened complete splitting of $f_\#^k(\sigma)$ is $\bar{E}$ for all $k \ge k_\sigma$. Moreover, if the first growing term in the coarsened complete splitting of σ is not an edge in $\mathcal{E}_f^{-1}$ then $k_\sigma = 1$.*

Proof There is no loss in replacing σ with its first growing term. The only case that does not follow from Examples 5.12 is that $\sigma = \bar{E} \in \mathcal{E}_f^{-1}$. This case follows from the definition of M and the obvious induction argument. □

Remark 5.15 The rays in Lemma 5.14 are finitely determined: in case (a) R is determined by the edge E, in case (b) R is determined by E, w and a choice of $\pm$, and in case (c) R is determined by w and a choice of $\pm$. From this data one can write down any finite initial subpath of R.

Lemma 5.16 *Suppose that $\sigma \subset G$ is a completely split path and that $\sigma = \alpha \cdot \beta$ is a coarsening of the complete splitting in which both α and β are growing. Let $R^- = f_\#^\infty(\bar{\alpha})$, let $R^+ = f_\#^\infty(\beta)$ and let $\ell = (R^-)^{-1} R^+$. Then $\mathsf{Acc}(\sigma) = \mathsf{Acc}(\alpha) \cup \mathsf{Acc}(\beta) \cup \{\ell\}$.*

Proof The inclusion $\mathsf{Acc}(\alpha) \cup \mathsf{Acc}(\beta) \subset \mathsf{Acc}(\sigma)$ follows from the fact that α and β occur as concatenation of terms in a splitting of σ. It is an immediate consequence of the definitions that $\ell \in \mathsf{Acc}(\gamma)$. It therefore suffices to assume that $L \in \mathsf{Acc}(\sigma)$ is not contained in $\mathsf{Acc}(\alpha) \cup \mathsf{Acc}(\beta)$ and prove that $L = \ell$.

Choose a finite subpath L_1 of L and $K > 0$ so that L_1 does not occur as a subpath of $f_\#^k(\alpha)$ or of $f_\#^k(\beta)$ for $k \ge K$. Extend L_1 to an increasing sequence $L_1 \subset L_2 \subset \cdots$ of finite subpaths of L whose union is L. For each $j \ge 1$, let C_j be the length of L_j. There exist arbitrarily large k so that L_j includes as a subpath of $f_\#^k(\sigma) = f_\#^k(\alpha) \cdot f_\#^k(\beta)$. The induced inclusion of L_1 in $f_\#^k(\sigma)$ must intersect both $f_\#^k(\alpha)$ and $f_\#^k(\beta)$ and so L_j is included as a subpath of the concatenation of the terminal segment of $f_\#^k(\alpha)$ of length C_j with the initial segment of $f_\#^k(\beta)$ of length C_j. If k is sufficiently large then the length C_j initial segments of R^- and of $f_\#^k(\bar{\alpha})$ agree and the length C_j initial segments of R^+ and of $f_\#^k(\beta)$ agree. Thus each L_j can be included as a subpath of ℓ. Since the induced inclusion of L_1 contains the juncture point between R^- and R^+, we may pass to a subsequence of L_j's and choose inclusions of L_j into ℓ so that induced inclusion of L_1 in ℓ is independent of j. It follows that if $i < j$ then the inclusion of L_i into ℓ is the restriction of the inclusion of L_j into ℓ and hence that there is a well-defined inclusion of L into ℓ. This inclusion is necessarily onto and so $L = \ell$. □

Corollary 5.17 *For each $r \in \mathcal{R}(\phi)$:*

(1) *Each $L \in \Omega(r)$ decomposes as $L = (R^-)^{-1} \rho R^+$ where ρ is a (possibly trivial) Nielsen path and R^+ and R^- satisfy (1a), (1b) or (1c) of Lemma 5.14. In particular, each L is ϕ-invariant and is finitely determined in the sense of Remark 5.15, and each periodic L equals $w^{\pm\infty}$ for some twist path w.*

(2) $\Omega(r)$ *is a finite set and the finite data that determines each of its elements can be read off from* $f: G \to G$.

(3) $\Omega_{\mathsf{NP}}(r) \neq \varnothing$.

(4) *For each* $L \in \Omega(r)$ *and each lift* $\widetilde{L}$, *there exists* $\Phi \in \mathcal{P}(\phi)$ *such that* $\partial_-\widetilde{L}, \partial_+\widetilde{L} \in \mathsf{Fix}_{\mathsf{N}}(\Phi)$. *Equivalently,* L *lifts into* $\Gamma(f)$.

Proof By Lemma 5.8 we can replace $\Omega(r)$ with $\mathsf{Acc}(E)$, where $E \in \mathcal{E}_f$ corresponds to r. Lemma 5.8 also implies that $\mathsf{Acc}(E) = \mathsf{Acc}(u \cdot f_\#(u))$, where $f(E) = E \cdot u$. Let

$$u \cdot f_\#(u) = \rho_0 \cdot \sigma_1 \cdot \rho_1 \cdot \sigma_2 \ldots \sigma_q \cdot \rho_q$$

be a coarsening of the complete splitting of $u \cdot f_\#(u)$ so that each σ_i is a single growing term and so that the ρ_i are (possibly trivial) Nielsen paths. For $1 \le i \le q-1$, let $R_i^- = f_\#^\infty(\bar{\sigma}_i)$ and for $2 \le i \le q$, let $R_i^+ = f_\#^\infty(\sigma_i)$. For $1 \le i \le q-1$, define $\ell_i = (R_i^-)^{-1} \rho_i R_{i+1}^+$. Lemma 5.16 and the obvious induction argument imply that

$$\Omega(r) = \mathsf{Acc}(u \cdot f_\#(u)) = \mathsf{Acc}(\sigma_1) \cup \ell_1 \cup \mathsf{Acc}(\sigma_2) \cup \cdots \cup \ell_{q-1} \cup \mathsf{Acc}(\sigma_q).$$

Lemma 5.14(1) implies that each ℓ_i satisfies (1). If σ_i is linear then $\mathsf{Acc}(\sigma_i) = w^{\pm\infty}$ for some twist path w by Examples 5.7. The remaining σ_i have the form E' or $\bar{E}'$ for some $E' \in \mathcal{E}_f$ with height less than that of E. Downward induction on the height of E completes the proof of (1) and (2).

We now turn to (3), assuming at first that $q > 2$. If σ_2 is exceptional or an element of $\mathcal{E}_f \cup \mathsf{Lin}(f)$ then ℓ_1 is nonperiodic. Otherwise, $\bar{\sigma}_2 \in \mathcal{E}_f \cup \mathsf{Lin}(f)$ and ℓ_2 is nonperiodic. Both of these statements follow from Lemma 5.14. If $q = 2$ then σ_1 is linear. One easily checks that ℓ_1 is nonperiodic in the various cases that can occur. For example if $\sigma_1 = E_1 \in \mathsf{Lin}(f)$ then $\sigma_2 = E_1$ and $\ell_1 = \bar{w}^{\pm\infty} \rho_1 E_1 w^{\pm\infty}$. The remaining cases are left to the reader.

The equivalence of the two conditions in (4) follows from Lemma 4.1. To prove that L lifts into $\Gamma(f)$, we make use of the fact that each vertex in G lifts uniquely to $\Gamma^0(f)$ and the fact that each Nielsen path in G lifts uniquely into $\Gamma(f)$ with one, and hence both, endpoints in $\Gamma^0(f)$. These facts follow immediately from the construction of $\Gamma(f)$ and the fact that every Nielsen path is a concatenation of fixed edges and (necessarily closed) indivisible Nielsen paths. Given these facts, we may assume that $L = (R^-)^{-1} \rho R^+$ is not a concatenation of Nielsen paths and hence that the initial edge E_j of either R^- or R^+ is an element of $\mathsf{Lin}(f) \cup \mathcal{E}_f$. The two cases are symmetric so we may assume that E_j is the initial edge of R^-. Let $E_j' \subset \Gamma(f)$ be the unique lift of E_j with initial vertex $v' \in \Gamma^0(f)$ and then extend this to a lift of R^- into $\Gamma(f)$. The Nielsen path ρ lifts to a path $\rho' \subset \Gamma^0(f)$ with initial vertex v' and terminal vertex, say w'. If R^+ is a concatenation of Nielsen paths then it lifts into $\Gamma^0(f)$ with initial vertex w'. Otherwise we lift R^+ in the same way that we lifted R^-. □

Example 3.1 (continued) Recall $R_q = q \cdot c \cdot cb \cdot cbba \cdot \ldots \cdot cbba \ldots ba^{k-1} \cdot \ldots$ and so $\Omega(r_q) = \{a^\infty R_c, a^\infty ba^\infty, a^\infty\}$ and $\Omega_{\mathsf{NP}}(r_q) = \{a^\infty R_c, a^\infty ba^\infty\}$.

6 Special free factor systems

6.1 A canonical collection of free factor systems

In this section, we define a canonical partial order $<$ on $\mathcal{R}(\phi)$ and then associate a nested sequence $\vec{\mathcal{F}}(\phi, <_T) = \mathcal{F}_0 \sqsubset \mathcal{F}_1 \sqsubset \cdots \sqsubset \mathcal{F}_t$ of ϕ-invariant free factor systems to each total order $<_T$ on $\mathcal{R}(\phi)$ that extends $<$. The bottom free factor system $\mathcal{F}_0$ is the smallest free factor system that carries all conjugacy classes that grow at most linearly and is independent of $<_T$. The inclusions $\mathcal{F}_{i-1} \sqsubset \mathcal{F}_i$ are all one-edge extensions. The CTs that represent ϕ with filtrations that realize $\vec{\mathcal{F}}(\phi, <_T)$ are easier to work with than generic CTs; see Lemma 6.9.

Notation 6.1 Suppose that $f: G \to G$ is a CT representing ϕ and that E_1 and E_2 are distinct elements of $\mathcal{E}_f$. If E_1 or $\bar{E}_1$ is a term of the complete splitting of $f_\#^k(E_2)$ for some $k \geq 1$, then we write $E_1 < E_2$. Lemma 3.21 implies that $<$ is a partial order on $\mathcal{E}_f$. If $E_1 < E_2$ are consecutive elements in the partial order then we write $E_1 <_c E_2$. Note that if we define $E_1 <' E_2$ to mean E_1 or $\bar{E}_1$ is a term of the complete splitting of $f(E_2)$ then $<$ is the partial order determined from $<'$ by extending transitively. Thus $<$ can be computed.

If $r_1, r_2 \in \mathcal{R}(\phi)$ and r_1 is an end of some element of $\Omega_{\mathsf{NP}}(r_2)$ then we write $r_1 < r_2$. Lemma 6.2 below implies that $<$ defines a partial order on $\mathcal{R}(\phi)$. If $r_1 < r_2$ are consecutive elements in the partial order then we write $r_1 <_c r_2$.

Example 3.1 (continued) In our example, the only relation is $r_c < r_q$.

Recall from Lemma 3.26 that the map that sends E to the end of R_E defines a bijection between $\mathcal{E}_f$ and $\mathcal{R}(\phi)$.

Lemma 6.2 *For any* CT $f: G \to G$, *the bijection between* $\mathcal{E}_f$ *and* $\mathcal{R}(\phi)$ *preserves* $<$.

Proof Suppose that $E_1, E_2 \in \mathcal{E}_f$ correspond to $r_1, r_2 \in \mathcal{R}(\phi)$, respectively.

If $E_1 < E_2$ and $f(E_2) = E_2 \cdot u_2$, then E_1 or $\bar{E}_1$ is a term in the complete splitting of $f_\#^k(u_2)$ for some, and hence all sufficiently large, k. By Lemma 5.8, there exists a completely split path γ such that $\Omega(r_2) = \mathsf{Acc}(\gamma)$ and such that the complete splitting of γ has a coarsening $\gamma = \gamma_1 \cdot \gamma_2 \cdot \gamma_3$ into three growing terms with γ_2 equal to either $\bar{E}_1$ or E_1. Lemma 5.16 therefore implies that R_{E_1} is a terminal ray of L or L^{-1} for some $L \in \Omega(r_2)$. Thus $r_1 < r_2$.

If $r_1 < r_2$ then R_{E_1} is a terminal ray of L or L^{-1} for some $L \in \Omega(r_2)$ by Corollary 5.17(1). It follows that $f_\#^k(E_2)$ crosses E_1 or $\bar{E}_1$ for all sufficiently large k. Lemma 3.21 implies that $E_1 < E_2$. □

Lemma 6.3 *If* $\psi = \theta\phi\theta^{-1}$ *then the bijection* $\mathcal{R}(\phi) \to \mathcal{R}(\psi)$ *induced by* θ *(see Lemma 3.16) preserves partial orders.*

Proof By Lemmas 3.16 and 5.4, we have $\theta(\mathcal{R}(\phi)) = \mathcal{R}(\psi)$ and $\theta(\Omega_{\mathsf{NP}}(r)) = \Omega_{\mathsf{NP}}(\theta(r))$ for each $r \in \mathcal{R}(\phi)$. The fact that θ preserves the partial order now follows from the definition of the partial order. □

Notation 6.4 Extend the partial order $<$ on $\mathcal{R}(\phi)$ to a total order $<_T$ and write $\mathcal{R}(\phi) = \{r_1, \dots, r_s\}$, where the elements are listed in increasing order. Given a CT $f : G \to G$ representing ϕ, transfer the total order $<_T$ on $\mathcal{R}(\phi)$ to a total order (also called) $<_T$ on $\mathcal{E}_f = \{E_1, \dots, E_s\}$ using the bijection between $\mathcal{E}_f$ and $\mathcal{R}(\phi)$ given in Lemma 3.26.

Recall from Section 4.1 that each component C of the eigengraph $\Gamma(f)$ is constructed from a component C_0 of $\mathsf{Fix}(f)$ by first adding "lollipops", one for each $E \in \mathsf{Lin}(f)$ with initial vertex in C_0, to form C_1, and then adding rays labeled R_E, one for each $E \in \mathcal{E}_f$ with initial vertex in C_0. Each $E \in \mathcal{E}_f$ contributes exactly one ray to $\Gamma(f)$ and we identify that ray with the eigenray R_E; it is the unique lift of R_E to $\Gamma(f)$. Each $E \in \mathsf{Lin}(f)$ contributes exactly one lollipop to $\Gamma(f)$. Note that C is contractible if and only if C_1 is contractible if and only if C_0 is contractible and there are no $E \in \mathsf{Lin}(f)$ with initial vertex in C_0. In this contractible case, C is obtained from a (possibly trivial) tree in $\mathsf{Fix}(f)$ by adding eigenrays and we single out the ray $R_E \subset C$ whose associated edge E is lowest with respect to $<_T$. These edges define subsets $\mathcal{R}^*(\phi) \subset \mathcal{R}(\phi)$ and $\mathcal{E}^*_f \subset \mathcal{E}_f$ that correspond under the bijection between $\mathcal{R}(\phi)$ and $\mathcal{E}_f$.

Definition 6.5 A conjugacy class $[a]$ *grows at most linearly* under iteration by ϕ if for some, and hence every, set of generators there is a linear function P such that word length of $\phi^k([a])$ with respect to those generators is bounded by $P(k)$. If $f : G \to G$ represents ϕ, then word length of $\phi^k([a])$ can be replaced by edge length of $f^k_\#(\sigma)$ in G, where $\sigma \subset G$ is the circuit representing $[a]$. The *linear growth free factor system* $\mathcal{F}_0(\phi)$ is the minimal free factor system that carries all conjugacy classes that grow at most linearly under iteration by ϕ.

Lemma 6.6 *Suppose that $f : G \to G$ is a* CT *representing ϕ and that $<_T$ and $\mathcal{E}_f = \{E_1, \dots, E_s\}$ are as in Notation 6.4. Let $K_0 \subset G$ be the subgraph consisting of all fixed and linear edges for $f : G \to G$. For $1 \le j \le s$, inductively define $K_j = K_{j-1} \cup E_j$. Then:*

(1) *$\mathcal{F}_0(\phi) = \mathcal{F}(K_0, G)$ (as defined at the beginning of Section 3.3).*

(2) *Each K_j is f-invariant.*

(3) *If $E_j \in \mathcal{E}^*_f$, then $\mathcal{F}(K_j, G) = \mathcal{F}(K_{j-1}, G)$; otherwise $\mathcal{F}(K_{j-1}, G) \sqsubset \mathcal{F}(K_j, G)$ is a proper one-edge extension.*

Proof In proving (1), we work with circuits $\sigma \subset G$ and edge length in G rather than conjugacy classes $[a]$ and word length with respect to a set of generators of F_n. If E is an edge of K_0, then $f(E) = E \cdot u$ for some (possibly trivial) closed Nielsen path u. Lemma 3.21 implies that $u \subset K_0$ and hence that K_0 is f-invariant. Each circuit in K_0 grows at most linearly under iteration by $f_\#$ since every edge in K_0 does. Thus $\mathcal{F}(K_0, G) \sqsubset \mathcal{F}_0(\phi)$.

After replacing σ with $f^m_\#(\sigma)$ for some $m \ge 0$, we may assume by [Feighn and Handel 2011, Lemma 4.25] that σ is completely split. Lemma 3.21 implies that if σ is not contained in K_0 then, up to reversal of orientation, some term in the complete splitting of σ is an edge $E \in \mathcal{E}_f$. In this case, σ grows at least as fast as E does. If $f(E) = E \cdot u$ then the length of $f_\#^k(u)$ goes to infinity with k and so the length of $f_\#^k(E)$ grows faster than any linear function. This proves that K_0 contains every circuit that grows at most linearly so $\mathcal{F}_0(\phi) \sqsubset \mathcal{F}(K_0, G)$. This completes the proof of (1).

For the remainder of the proof we may assume that $j \geq 1$. For $E_j \in \mathcal{E}_f$, the terms in the complete splitting of $f(E_j)$, other than E_j itself, are exceptional paths, Nielsen paths and single edges E_i or $\overline{E}_i$ that are either linear or satisfy $E_i < E_j$. Lemma 3.21 implies that the exceptional paths and Nielsen paths are contained in K_0. The single edge terms other than E_j are contained in K_{j-1} by construction. Thus $f(E_j) \subset K_j$ and K_j is f-invariant. This proves (2).

The terminal endpoint of each $E_j \in \mathcal{E}_f$ is contained in a noncontractible component of K_{j-1} because $f(E_j) = E_j u_j$ for a nontrivial closed path $u_j \subset K_{j-1}$. If $E_j \in \mathcal{E}_f^*$ with initial vertex v_j then the component of K_{j-1} that contains v_j is a contractible component of $\mathsf{Fix}(f)$. In this case every line in K_j is contained in K_{j-1} so $\mathcal{F}(K_j, G) = \mathcal{F}(K_{j-1}, G)$. Otherwise, v_j is contained in a noncontractible component of K_{j-1} so $\mathcal{F}(K_{j-1}, G) \sqsubset \mathcal{F}(K_j, G)$ is a proper inclusion. Obviously K_j is obtained from K_{j-1} by adding a single edge. □

Recall from Lemma 4.1 that the set of lines that lift to $\Gamma(f)$ is independent of the choice of CT $f : G \to G$ representing ϕ. The next lemma shows that the $\mathcal{F}(K_j, G)$ defined in Lemma 6.6 depend only on ϕ and $<_T$ and not on the choice of CT $f : G \to G$.

Lemma 6.7 *Continue with the notation of Lemma 6.6. For each $r_j \notin \mathcal{R}^*(\phi)$, there exists at least one line $L(r_j)$ that lifts to $\Gamma(f)$, whose terminal end is r_j and whose initial end is not r_l for any $l \geq j$. Moreover, for any such choice of lines, $\mathcal{F}(K_j, G)$ is the smallest free factor system that contains $\mathcal{F}_0(\phi)$ and carries $\{L(r_l) \mid l \leq j \text{ and } r_l \notin \mathcal{R}^*(\phi)\}$.*

Proof Let $C = C(j)$ be the component of $\Gamma(f)$ that contains R_{E_j}, let $C_0 \subset C_1 \subset C$ be as in Notation 6.4 and, for each $1 \leq q \leq s$, let $A_q \subset C$ be the union of C_1 with the rays R_{E_l} in C with $E_l \leq E_q$. By construction, and by Lemma 6.6, R_{E_l} is included in A_q if and only if $R_{E_l} \subset K_q$. Since $r_j \notin \mathcal{R}^*(\phi)$, either C_1 is noncontractible or A_{j-1} contains at least one ray R_{E_l}. In both cases, the ray $R_{E_j} \subset C$ extends by a ray in A_{j-1} to a line in A_j. The projection $L(r_j)$ of this line into K_j satisfies the conclusions of the main statement of the lemma.

The "moreover" part of the lemma is proved by induction on j, with the base case $j = 0$ following from Lemma 6.6(1). For the inductive case, let $\mathcal{F}_j'$ be the smallest free factor system that carries K_{j-1} and $L(r_j)$. Then $\mathcal{F}(K_{j-1}, G) \sqsubset \mathcal{F}_j' \sqsubset \mathcal{F}(K_j, G)$ with the first inclusion being proper and $\mathcal{F}_j'$ does not have more components than $\mathcal{F}(K_{j-1}, G)$. Lemma 3.4 implies that $\mathcal{F}_j' = \mathcal{F}(K_j, G)$. □

Notation 6.8 Let $K_0 \subset K_1 \subset \cdots \subset K_s = G$ be as in Lemma 6.6 and let $\vec{\mathcal{F}}(\phi, <_T) = \mathcal{F}_0 \sqsubset \mathcal{F}_1 \sqsubset \cdots \sqsubset \mathcal{F}_t$ be the increasing sequence of distinct free factor systems determined by the K_j. (Equivalently, $\vec{\mathcal{F}}(\phi, <_T)$ is the sequence determined by those K_j with $r_j \notin \mathcal{R}^*(\phi)$.) We say that $\vec{\mathcal{F}}(\phi, <_T)$ is the *sequence of free factor systems determined by ϕ and $<_T$*. Lemma 6.7 justifies this description by showing that $\vec{\mathcal{F}}(\phi, <_T)$

depends only on ϕ and $<_T$. To simplify notation a bit, we write L_k for $L(r(j))$ where $r(j)$ is the k^{th}-lowest element of $\mathcal{R}(\phi) \setminus \mathcal{R}^*(\phi)$. Thus $\mathcal{F}_k$ is filled by $\mathcal{F}_0$ and $L_1, \dots, L_k$.

We sometimes refer to a nested sequence of free factor systems as a *chain*. A chain $\mathfrak{c} = (\mathcal{F}_0 \sqsubset \cdots \sqsubset \mathcal{F}_t)$ is *special for* ϕ if $\mathfrak{c} = \vec{\mathcal{F}}(\phi, <_T)$ for some extension $<_T$ of $<$ to a total order on $\mathcal{R}(\phi)$. A free factor system $\mathcal{F}$ is *special for* ϕ if $\mathcal{F}$ is an element of some special chain for ϕ. The set of special free factor systems for ϕ is denoted by $\mathfrak{L}(\phi)$. A free factor F or its conjugacy class is *special for* ϕ if $[F]$ is an element of some special free factor system for ϕ. A pair $\mathfrak{e} = (\mathcal{F}^- \sqsubset \mathcal{F}^+)$ of free factor systems is a *special one-edge extension for* ϕ if its appears as consecutive elements of some special chain for ϕ.

By applying the existence theorem for CTs given in [Feighn and Handel 2018, Theorem 1.1], we can choose a CT whose filtration realizes $\vec{\mathcal{F}}(\phi, <_T)$ for any given $<_T$. The following lemma shows that the case analysis for a CT with this property is simpler than that of a random CT.

Lemma 6.9 *Suppose that $\vec{\mathcal{F}}(\phi, <_T) = \mathcal{F}_0 \sqsubset \mathcal{F}_1 \sqsubset \cdots \sqsubset \mathcal{F}_t$ and that $f : G \to G$ and $\varnothing = G_0 \subset G_1 \subset \cdots \subset G_N = G$ are a* CT *and filtration representing ϕ and realizing $\vec{\mathcal{F}}(\phi, <_T)$; ie for all $0 \le k \le t$ there is an f-invariant core subgraph G_{i_k} such that $\mathcal{F}_k = \mathcal{F}(G_{i_k}, G)$. Then $G_{i_k} \setminus G_{i_{k-1}}$ is a single topological arc A_k with both endpoints in $G_{i_{k-1}}$. Moreover, letting D_k be the element of $\mathcal{E}_f$ corresponding to $\partial_+ L_k \in \mathcal{R}(\phi)$ (as in Lemma 6.7), A_k can be oriented so that one of the following is satisfied*:

- [HH] $A_k = \bar{C}_k D_k$, *where* $C_k \in \mathcal{E}_f$.
- [LH] $A_k = \bar{C}_k D_k$, *where* $C_k \in \mathsf{Lin}(f)$.
- [H] $A_k = D_k$.

Proof By Lemma 6.6, each $\mathcal{F}_{j-1} \sqsubset \mathcal{F}_j$ is a one-edge extension. [Handel and Mosher 2020, Part II, Lemma 2.5] therefore implies that G_{i_k} is constructed from $G_{i_{k-1}}$ in one of three ways: add a single topological edge with both endpoints in $G_{i_{k-1}}$; add a single topological edge that forms a circuit that is disjoint from $G_{i_{k-1}}$; add an edge forming a disjoint circuit and then add an edge connecting that circuit to $G_{i_{k-1}}$. In the second and third cases the circuit is f-invariant in contradiction to the fact that K_0 contains all ϕ-invariant conjugacy classes. Thus G_{i_k} is obtained from $G_{i_{k-1}}$ by adding a single topological arc A_k with both endpoints in $G_{i_{k-1}}$.

The arc A_k consists of either one or two edges of G. Indeed, a "middle" edge cannot be fixed by the (Periodic Edge) property [Feighn and Handel 2011, Definition 4.7(5)] of a CT and cannot be nonfixed because in that case its terminal end would be contained in a core subgraph of $G_{i_{k-1}}$ by [Feighn and Handel 2011, Lemma 4.21]. The (Periodic Edge) property also implies that if A_k consists of two edges, then neither is fixed. To complete the proof, it suffices to show that A_k crosses D_k. We will do so by showing that D_k is not contained in $G_{i_{k-1}}$ and is contained in G_{i_k}.

A line $L \subset G$ that lifts to $\Gamma(f)$ but is not contained in K_0 either decomposes as the concatenation of a ray in K_0 and an eigenray $R_{E'}$, or decomposes as the concatenation of a finite path in K_0 and a pair of eigenrays $R_{E'}$ and $R_{E''}$. In the former case, each $E \in \mathcal{E}_f$ crossed by L satisfies $E \le E'$, and in the latter case each $E \in \mathcal{E}_f$ satisfies $E \le E'$ or $E \le E''$. It follows that every edge $E \in \mathcal{E}_f$ crossed by $\bigcup_{q=1}^{k-1} L_q$ satisfies $E <_T D_k$. Since K_0 and $L_1, \dots, L_{k-1}$ fill $\mathcal{F}_{k-1}$, we conclude that D_k is not contained in $G_{i_{k-1}}$. Since L_k lifts to $\Gamma(f)$ and $\partial_+ L_k = \partial R_{D_k}$, it follows that R_{D_k} is a terminal ray of L_k. In particular, L_k crosses D_k. Lemma 6.7 implies that $L_k \subset G_{i_k}$ and we are done. □

Lemma 6.10 *Let $\mathfrak{e} = (\mathcal{F}^- \sqsubset \mathcal{F}^+)$ be special for ϕ.*

(1) *The types* HH, LH *or* H *of $\mathfrak{e}$ as in Lemma 6.9 are mutually exclusive and independent of the special chain $\vec{\mathcal{F}}(\phi, <_T)$ containing $\mathfrak{e}$ and the choice of* CT $f : G \to G$ *realizing $\vec{\mathcal{F}}(\phi, <_T)$.*

(2) *Suppose that $\mathfrak{e}$ appears as consecutive elements $\mathcal{F}_{k-1} \sqsubset \mathcal{F}_k$ in $\vec{\mathcal{F}}(\phi, <_T)$, which is realized by the* CT $f : G \to G$. *Using terminology as in Lemma 6.9, say that $\mathfrak{e}$ is, respectively,* contractible, infinite cyclic, *or* large *depending on whether the component of the eigengraph $\Gamma_{f|\mathcal{F}_k}$ containing the eigenray R_{D_k} is contractible, has infinite cyclic fundamental group, or has fundamental group with rank at least two. The types* contractible, infinite cyclic, *or* large *of $\mathfrak{e}$ are mutually exclusive and independent of the choices of $\vec{\mathcal{F}}(\phi, <_T)$ and f.*

Proof (1) Suppose $\mathfrak{e} = (\mathcal{F}_{k-1} \sqsubset \mathcal{F}_k)$ in $\vec{\mathcal{F}}(\phi, <_T)$. The difference between the cardinality of $\mathcal{R}(\phi|\mathcal{F}_k)$ and the cardinality of $\mathcal{R}(\phi|\mathcal{F}_{k-1})$ is 2 in the [HH] case and 1 in the [LH] and [H] cases. In case [LH], either the number of axes for $\phi|\mathcal{F}_k$ is strictly larger than the number of axes for $\phi|\mathcal{F}_{k-1}$ or there is a common axis of $\phi|\mathcal{F}_k$ and $\phi|\mathcal{F}_{k-1}$ whose multiplicity in the former is strictly larger than in the latter. Neither of these happens in case [H].

(2) Here is an invariant description. Let $r \in \Delta := \mathcal{R}(\phi|\mathcal{F}_k) \setminus \mathcal{R}(\phi|\mathcal{F}_{k-1})$, for example we could take r to be determined by R_{D_k}. Either $\Delta = \{r\}$ or $\Delta = \{r, s\}$ and there is a $\phi|\mathcal{F}_k$-fixed line L whose ends represent r and s. Let $\tilde{r}$ be a lift of r to ∂F_n and let $\tilde{L} = [\tilde{r}, \tilde{s}]$ be a lift of L if $\Delta = \{r, s\}$. By definition, $\mathfrak{e}$ is contractible, infinite cyclic or large if and only if $\mathrm{Fix}(\Phi_{\tilde{r}})$ is trivial, infinite cyclic, or of rank at least two, where $\Phi_{\tilde{r}}$ is the unique representative of ϕ fixing $\tilde{r}$. We are done by noting that $\Phi_{\tilde{r}} = \Phi_{\tilde{s}}$ if $\Delta = \{r, s\}$. □

Example 3.1 (continued) If we extend the partial order $r_c < r_q$ on $\mathcal{R}(\phi)$ to the total order $r_c <_T r_d <_T r_e <_T r_q$ we get the special chain $\mathfrak{c}$ represented by the sequence of graphs in Figure 3. See the notation in the examples on pages 1700 and 1707.

Example 6.11 Consider the CT $f : G \to G$ given as follows: start with a rose with edges a and b. Define $f(a) = a$ and $f(b) = ba$. Add a new vertex v with adjacent edges c, d and define $f(c) = cb$ and $f(d) = db^2$. Add another new vertex v' with adjacent edges c' and d' with $f(c') = c'b^3$ and $f(d') = d'b^4$ finally add an f-fixed edge e with endpoints v and v'. The ϕ-fixed free factor system $\mathcal{F}$ represented by the complement of e in G is not in $\mathfrak{L}(\phi)$. Indeed, $\mathcal{F} \sqsubset \{F_n\}$ is 1-edge, but not of type H, HH or LH, contradicting Lemma 6.9.

Example 6.12 Suppose $f: G \to G$ is a CT containing a circle C with only one vertex x and such that x is the initial endpoint of an H-edge, the terminal endpoint of a linear edge in an LH extension (so that C is an axis), and there are no other edges containing x. Then $\Gamma(f)$ has no components of rank at least two containing an axis corresponding to C.

Lemma 6.13 *Referring to Notation 6.8, suppose $\mathcal{F}$, $[F]$, $\mathfrak{c}$ and $\mathfrak{e}$ are special for ϕ. If $\theta \in \mathsf{Out}(F_n)$, then*

- $\theta(\mathcal{F}), \theta([F]), \theta(\mathfrak{c})$, *and* $\theta(\mathfrak{e})$ *are special for* ϕ^θ, *and*
- *the types* H, HH *or* LH *and the types* contractible, infinite cyclic *or* large *of $\mathfrak{e}$ and $\theta(\mathfrak{e})$ are the same.*

Proof This is an immediate consequence of the fact (Lemma 6.3) that conjugation preserves partial orders and the invariant description of types given in the proof of Lemma 6.10. □

Definition 6.14 (added lines) Suppose that $\mathfrak{c}$ is a special chain for ϕ that is realized by $f: G \to G$ and that $\mathfrak{e} = (\mathcal{F}^- \sqsubset \mathcal{F}^+) \in \mathfrak{c}$. Then $\mathcal{R}(\phi|\mathcal{F}^+) \setminus \mathcal{R}(\phi|\mathcal{F}^-)$ contains two elements if $\mathfrak{e}$ has type HH and one element otherwise. These elements are said to be *new* with respect to $\mathfrak{e}$. Similarly $\Gamma(f|\mathcal{F}^+)$ carries more lines than $\Gamma(f|\mathcal{F}^-)$. The set of *added lines with respect to* $\mathfrak{e}$, denoted by $\mathsf{L}_\mathfrak{e}(\phi)$, is a ϕ-invariant subset of these lines. In case $\mathfrak{e}$ is contractible, $\mathsf{L}_\mathfrak{e}(\phi)$ consists of all lines L in $\Gamma(f|\mathcal{F}^+)$ with $\partial_+ L$ new. If $\mathfrak{e}$ is not contractible then we also require that $\partial_- L$ is not in $\mathcal{R}(\phi)$. $\mathsf{L}_\mathfrak{e}(\phi)$ has an equivalent invariant description as follows. Set $\Phi := \Phi_{\tilde{r}^+}|F^+$, where r^+ is new, $\tilde{r}^+ \in \partial F^+ \subset \partial F_n$, and $[F^+]$ is the component of $\mathcal{F}^+$ carrying r^+. Define $\mathsf{L}_\mathfrak{e}(\phi)$ to be $[\partial \operatorname{Fix}(\Phi), \tilde{r}^+]$ if $\operatorname{Fix}(\Phi)$ is nontrivial; else $[\operatorname{Fix}_{\mathsf{N}}(\Phi) \setminus \{\tilde{r}^+\}, \tilde{r}^+]$ if there is only one new eigenray; else the set consisting of the two lines with lifts with endpoints in $\operatorname{Fix}_{\mathsf{N}}(\Phi)$. This invariant description shows that $\mathsf{L}_\mathfrak{e}(\phi)$ is independent of the special chain $\mathfrak{e} \in \mathfrak{c}$, which is why $\mathfrak{c}$ does not appear in the notation.

Example 3.1 (continued) Referring to Figures 1, 2 and 3, if $\mathfrak{e}_1 := (\{[G_2]\} \sqsubset \{[G_3]\})$ then $\mathsf{L}_{\mathfrak{e}_1}(\phi)$ consists of the infinitely many lines L in the third listed component of $\Gamma(f)$ in Figure 2 that cross the oriented edge c exactly once and c^{-1} not at all. If $\mathfrak{e}_2 := (\{[G_3]\} \sqsubset \{[G_5]\})$ then $\mathsf{L}_{\mathfrak{e}_2}(\phi)$ consists of two lines; they are represented by $(R_d)^{-1} R_e$ and its inverse. If $\mathfrak{e}_3 := (\{[G_5]\} \sqsubset \{[G_7]\})$ then $\mathsf{L}_{\mathfrak{e}_3}(\phi)$ consists of two lines; they are represented by $a^\infty p^{-1} R_q$ and $a^{-\infty} p^{-1} R_q$.

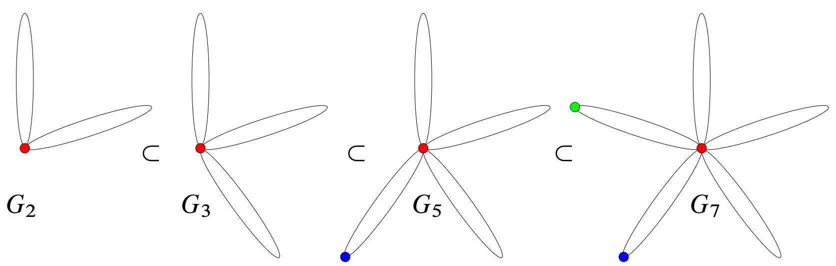

Figure 3: The special chain $\mathfrak{c} = \{[G_2]\} \sqsubset \{[G_3]\} \sqsubset \{[G_5]\} \sqsubset \{[G_7]\}$.

Lemma 6.15 *For* $\theta \in \mathsf{Out}(F_n)$, $\theta(\mathsf{L}_{\mathfrak{e}}(\phi)) = \mathsf{L}_{\theta(\mathfrak{e})}(\phi^\theta)$.

Proof By Lemma 6.13, the $\theta(\mathfrak{e}) \in \theta(\mathfrak{c})$ are special for ϕ^θ. The set of new elements of $\mathcal{R}(\phi)$ with respect to $\mathfrak{e} = (\mathcal{F}^+ \sqsubset \mathcal{F}^-)$ has the invariant description $\mathcal{R}(\phi|\mathcal{F}^+) \setminus \mathcal{R}(\phi|\mathcal{F}^-)$. In particular, θ takes the new elements with respect to ϕ to those with respect to ϕ^θ; see Lemma 3.16. The equation in the lemma then follows from the invariant definition of added lines in Definition 6.14 and the naturality results of Lemma 3.16. □

In the next lemma, we record some consequences of Lemmas 6.3, 6.6 and 6.7.

Lemma 6.16 (1) *A conjugacy class grows at most linearly under iteration by* ϕ *if and only if it is carried by* $\mathcal{F}_0(\phi)$.

(2) $\mathcal{F}_0(\phi) = \mathcal{F}(\mathsf{Fix}(\phi))$, *ie* $\mathcal{F}_0(\phi)$ *is the smallest free factor system carrying* $\mathsf{Fix}(\phi)$.

Proof (1) By definition $\mathcal{F}_0(\phi)$ carries all conjugacy classes that grow at most linearly. Conversely, by Lemma 6.6, $\mathcal{F}_0(\phi)$ is represented a graph K_0 consisting of linear and fixed edges. Hence every conjugacy class carried by $\mathcal{F}_0(\phi)$ grows at most linearly.

(2) By (1), $\mathcal{F}(\mathsf{Fix}(\phi)) \sqsubset \mathcal{F}_0(\phi)$. Suppose $\sqsubset$ is proper. By [Feighn and Handel 2018, Theorem 1.1], there is a CT $f : G \to G$ realizing ϕ with f-invariant core subgraphs $G_k \subsetneq G_l$ representing these two free factor systems and such that $f|G_l$ is a CT. By Lemma 6.6, every edge of G_l is fixed or linear. Let E be an edge of $G_l \setminus G_k$. There is a Nielsen circuit ρ in G_l containing E. Indeed, by the construction of eigengraphs in Section 4.1,

- every edge of a CT is the label of some edge in its eigengraph,
- the eigengraph of a linear growth CT is a compact core graph, and
- every circuit in an eigengraph is Nielsen.

The existence of ρ now follows from the defining property of a core graph that there is a circuit through every edge. The fixed conjugacy class represented by ρ is not in $\mathcal{F}(\mathsf{Fix}(\phi))$, a contradiction. □

6.2 The lattice of special free factor systems

This section is not needed for the rest of the paper and so could be skipped by the reader. Recall (second paragraph of Notation 6.8) that $\mathfrak{L}(\phi)$ denotes the set of special free factor systems for ϕ. The main results, Lemmas 6.18 and 6.20, are that $(\mathfrak{L}(\phi), \sqsubset)$ is a lattice that is natural with respect to $\mathsf{Out}(F_n)$ in the sense that, for $\theta \in \mathsf{Out}(F_n)$,

$$\theta(\mathfrak{L}(\phi), \sqsubset) = (\mathfrak{L}(\phi^\theta), \sqsubset).$$

In this section, $f : G \to G$ will always denote a CT for ϕ. We will conflate an element of $\mathcal{R}(\phi)$ and its image in $\mathcal{E}_f$ under the bijection $\mathcal{R}(\phi) \leftrightarrow \mathcal{E}_f$; see Lemma 3.26. A subset S of $\mathcal{R}(\phi)$ is *admissible* if it satisfies

$$(q \in S) \wedge (r < q) \implies r \in S.$$

If $S \subset \mathcal{R}(\phi)$, then mimicking Lemma 6.6 we let $K(S)$ denote the union of K_0 and the edges in S. Recall that K_0 is the union of the fixed and linear edges of G.

Lemma 6.17 *The following are equivalent*:

(1) *$\mathcal{F}$ is special for ϕ.*

(2) *$\mathcal{F} = \mathcal{F}(K(S), G)$ for some admissible $S \subset \mathcal{R}(\phi)$.*

(3) *$\mathcal{F} = \mathcal{F}(H, G)$ for some f-invariant $H \subset G$ containing K_0.*

Proof (1) $\implies$ (2) By definition, a free factor system $\mathcal{F}$ is special if and only if there is a total order $<_T$ extending $<$ and an initial interval $[r_1, \dots, r_k]$ of $(\mathcal{R}, <_T)$ such that $\mathcal{F} = \mathcal{F}(K(\{r_1, \dots, r_k\}), G)$. Since an initial interval is admissible, we may take $S = \{r_1, \dots, r_k\}$.

(2) $\implies$ (3) We may take $H = K(S)$.

(3) $\implies$ (2) Let S be the set of edges in H that are not in K_0. It is enough to show that S is admissible. Let $q \in S$ and let $r \in \mathcal{R}(\phi)$ satisfy $r < q$. By definition of $<$, there is $k > 1$ so that the edge r or its inverse is a term in the complete splitting of $f_\#^k(q)$. Since the edge q is in H and H is f-invariant, the edge r is also in H.

(2) $\implies$ (1) We claim that if S is admissible then there is an extension $<_T$ of $<$ such that S is an initial segment of $(\mathcal{R}(\phi), <_T)$. Indeed, start with any total order $<_T$ extending $<$ and iteratively interchange r and s if $r <_T s$ are consecutive, $s \in S$ and $r \notin S$. For such a $<_T$, S represents an element of $\vec{\mathcal{F}}(\phi, <_T)$. □

Lemma 6.18 *Let $\mathcal{F}$ be a special free factor system for ϕ.*

(1) *The set of admissible subsets of $\mathcal{R}(\phi)$ is a sublattice of $2^{\mathcal{R}(\phi)}$.*

(2) *There is a minimal admissible $S \subset \mathcal{R}(\phi)$ such that $\mathcal{F} = \mathcal{F}(K(S), G)$. We say that such an admissible S is* ***efficient for $\mathcal{F}$***. *In fact, if S' is admissible then the set of edges of $\mathrm{core}(K(S'))$ not in K_0 is efficient for $\mathcal{F}(K(S'), G)$.*

(3) *$S = \bigcap\{S' \mid \mathcal{F} = \mathcal{F}(K(S'), G)\}$ is efficient for $\mathcal{F}$.*

(4) *S is efficient for $\mathcal{F}$ if and only if $\mathcal{F} = \mathcal{F}(K(S), G)$ and through every edge representing an element of S there is a circuit in $K(S)$.*

(5) *If S_1 and S_2 are efficient admissible and $\mathcal{F}(K(S_1), G) \sqsubset \mathcal{F}(K(S_2), G)$, then $S_1 \subset S_2$.*

(6) *$(\mathfrak{L}(\phi), \sqsubset)$ is a lattice.*

(7) *Every maximal chain in $\mathfrak{L}(\phi)$ is special.*

(8) *Every minimal[2] pair $\mathfrak{e} = (\mathcal{F}^- \sqsubset \mathcal{F}^+)$ in $\mathfrak{L}(\phi)$ is special.*

Proof (1) This follows directly from the definition of admissible.

[2]That is, if $\mathcal{F}$ is special and $\mathcal{F}^- \sqsubset \mathcal{F} \sqsubset \mathcal{F}^+$, then $\mathcal{F}^- = \mathcal{F}$ or $\mathcal{F} = \mathcal{F}^+$.

(2) Suppose S_1, S_2 are admissible and $K(S_1), K(S_2)$ each represent $\mathcal{F}$. Hence $C := \mathsf{core}(K(S_1)) = \mathsf{core}(K(S_2))$ and, since G is a CT, C is f-invariant. (Indeed, by [Feighn and Handel 2011, Lemma 4.21], the removal of an edge with a valence one vertex from an f-invariant subgraph results in an f-invariant subgraph.) It follows that $K_0 \cup C$ is the minimal f-invariant subgraph of G representing $\mathcal{F}$ and containing K_0; see Lemma 6.17. Hence S is the set of edges of $K_0 \cup C$ not in K_0.

(3), (4), (5) These follow easily from (2).

(6) Suppose S_1 and S_2 are efficient. Then using (4), $S_1 \cup S_2$ is efficient. It follows that $\mathcal{F}(K(S_1 \cup S_2), G)$ is the smallest (with respect to $\sqsubset$) special free factor system for ϕ containing $\mathcal{F}_1 := \mathcal{F}(K(S_1), G)$ and $\mathcal{F}_2 := \mathcal{F}(K(S_2), G)$. Suppose S is efficient and $\mathcal{F}(K(S), G) \sqsubset F_i$ for $i = 1, 2$. By (5), $S \subset S_1 \cap S_2$. Since $K(S_1 \cap S_2) = K(S_1) \cap K(S_2)$, the largest special free factor system $\mathcal{F}$ for ϕ in each of $\mathcal{F}_1$ and $\mathcal{F}_2$ is represented by $K(S_1 \cap S_2)$, ie by $K(S)$, where S is efficient for $\mathcal{F}(K(S_1 \cap S_2), G)$.

(7) Let $\mathfrak{c}$ be represented by $K(S_1) \subset \cdots \subset K(S_N)$ with each S_i efficient. By (5), $S_i \subset S_{i+1}$. An argument similar to that in the proof of Lemma 6.17, (2) $\Longrightarrow$ (1), shows that there is $<_T$ extending $<$ such that each S_i is an initial interval in $(\mathcal{R}(\phi), <_T)$. Hence $\mathfrak{c}$ is special.

(8) This follows from (7) by enlarging $\mathfrak{c}$ to a maximal chain. □

Remark 6.19 $\mathfrak{L}(F_n)$ is not a sublattice of the lattice of all ϕ-invariant free factor systems. For example, reconsider Example 6.11. $S = \{c, d\}$ and $S' = \{c', d'\}$ are efficient. If $\mathcal{F} = \mathcal{F}(K(S), G)$ and $\mathcal{F}' = \mathcal{F}(K(S'), G)$ then the smallest ϕ-invariant free factor system containing $\mathcal{F}$ and $\mathcal{F}'$ is represented by the complement of the fixed edge e whereas the smallest element of $\mathfrak{L}(\phi)$ containing $\mathcal{F}$ and $\mathcal{F}'$ is F_n.

In the proof of Lemma 6.18 we noted that the union of efficient sets is efficient. The intersection need not be efficient. For example, suppose highest-order edges a, b and c share an initial vertex of valence three. Consider the complement S of a and the complement S' of b. The edge c is in $S \cap S'$ and has initial vertex of valence one in $K(S \cap S')$.

Lemma 6.20
- *If $\mathfrak{e} = (\mathcal{F}^- \sqsubset \mathcal{F}^+)$ is minimal in $\mathfrak{L}(\phi)$, then $\mathfrak{e}$ has a well-defined type* H, HH *or* LH, *and a well-defined type* contractible, infinite cyclic *or* large.
- *For $\theta \in \mathsf{Out}(F_n)$, the map $\mathcal{F} \mapsto \theta(\mathcal{F})$ induces a lattice isomorphism*

$$(\mathfrak{L}(\phi), \sqsubset) \to (\mathfrak{L}(\phi^\theta), \sqsubset)$$

that preserves the above types.

Proof This follows Lemmas 6.18 and 6.13. □

7 More on conjugacy pairs

Recall that conjugacy pairs were introduced in Definition 4.9. In this section we define some conjugacy pairs that will be used to define invariants of elements of $\mathsf{UPG}(F_n)$ and describe their properties.

7.1 $[\partial H, \partial K]$

We will want to compare conjugacy pairs of subgroups $[H, K]$ with the set of lines $[\partial H, \partial K]$; see Examples 4.10. For this we will use the next lemma, which is a corollary of [Kapovich and Short 1996, Lemma 3.9].

Lemma 7.1 *Suppose that $H < F_n$ is finitely generated. Then the stabilizer G in F_n of $\partial H \subset \partial F_n$ is the maximal $K < F_n$ in which H has finite index.*

Corollary 7.2 *Suppose that finitely generated $H < F_n$ is root-closed, ie $a^k \in H$, $k \neq 0$ implies $a \in H$. Then H is the stabilizer in F_n of ∂H.*

Proof If $H < G$ has finite index and $H \neq G$, then H is not root-closed. □

Corollary 7.3 (1) *If $H < F_n$ is a free factor, then H is the stabilizer of ∂H.*

(2) *If $H = \mathsf{Fix}(\Phi)$ for $\Phi \in \mathsf{Aut}(F_n)$, then H is the stabilizer of ∂H.*

(3) *If $a \in F_n$ is root-free, then $A = \langle a \rangle$ is the stabilizer of the two-point set ∂A.*

Proof Free factors and the group generated by a root-free element are clearly root-closed. For (2), $\Phi(a^k) = a^k$ for $k \neq 0$ implies that $\Phi(a)$ is a k^{th} root of a^k and so equals a. □

Remark 7.4 Corollary 7.3, which contains the only cases that we need in this paper, does not require the generality of Lemma 7.1. Item (3) is elementary. Items (1) and (2) follow from (3), and:

- For $H, K < F_n$ finitely generated, $\partial H \cap \partial K = \partial(H \cap K)$. See [Kapovich and Benakli 2002, Theorem 12.2(9)] in the setting of hyperbolic groups or, for the case at hand, [Handel and Mosher 2020, Fact 1.2].
- If H is a nontrivial free factor, then $H \cap H^g = \{1\}$ unless $g \in H$.
- $\mathsf{Fix}(\Phi) \cap \mathsf{Fix}(\Phi)^g$ is cyclic unless $g \in \mathsf{Fix}(\Phi)$; see Lemma 4.4(1).

Corollary 7.5 *Suppose that $H, K < F_n$ are finitely generated and root-closed. Then $[H, K]$ determines $[\partial H, \partial K]$, and vice versa.*

Proof Suppose that $H', K' < F_n$ are finitely generated and root-closed.

If $[\partial H, \partial K] = [\partial H', \partial K']$ then there is $x \in F_n$ such that $(x\partial H, x\partial K) = (\partial H^x, \partial K^x) = (\partial H', \partial K')$. Hence

$$(H^x, K^x) = (\mathsf{Stab}(\partial H^x), \mathsf{Stab}(\partial K^x)) = (\mathsf{Stab}(\partial H'), \mathsf{Stab}(\partial K')) = (H', K').$$

So $[H, K] = [H', K']$.

Conversely, if $[H, K] = [H', K']$ then there is $x \in F_n$ such that $(H^x, K^x) = (H', K')$. Hence

$$(\partial H^x, \partial K^x) = (x\partial H, x\partial K) = (\partial H', \partial K'),$$

and so $[\partial H, \partial K] = [\partial H', \partial K']$. □

Remark 7.6 If we are in the setting of Corollary 7.5 and $\partial H \cap \partial K = \varnothing$, we will sometimes abuse notation and think of $[H, K]$ as the set of lines $[\partial H, \partial K]$ and vice versa.

7.2 Some Stallings graph algorithms

In this section we assume that G is a marked graph with marking $m\colon (R_n, *) \to (G, b)$, where $*$ is the unique vertex of the rose R_n and $b = m(*)$ is the basepoint for G. There is an induced identification of F_n with $\pi_1(G, b)$.

For each finitely generated subgroup $H < F_n$, Stallings [1983, 5.4] constructs a finite graph $\Sigma_b(H)$ with basepoint b_H and an immersion $p_H\colon (\Sigma_b(H), b_H) \to (G, b)$ such that the image of the injection $\pi_1(\Sigma_b(H), b_H) \to \pi_1(G, b)$ induced by p_H equals H. The basepoint b_H may have valence one but all other vertices of $\Sigma_b(H)$ have valence at least two. We equip $\Sigma_b(H)$ with the CW-structure whose vertex set is the preimage of the vertex set of G. The resulting edges of $\Sigma_b(H)$, sometimes called *edgelets*, are labeled by their image edges in G. The core of $\Sigma_b(H)$ is denoted by $\Sigma(H)$. The minimal edgelet-path from b_H to $\Sigma(H)$ is denoted by β_H. The terminal endpoint of β_H is denoted by $c_H \in \Sigma(H)$.

For finitely generated subgroups $K, H < F_n$, let $\mathrm{Imm}(K, H)$ be the set of immersions $J\colon \Sigma(K) \to \Sigma(H)$ that maps edgelets to edgelets and preserves labels; we say that J *preserves labels*. We do not distinguish between elements of $\mathrm{Imm}(K, H)$ that induce the same map on the set of edgelets. Thus $\mathrm{Imm}(K, H)$ is finite and can be computed by inspection. An *equivalence* is an element of $\mathrm{Imm}(K, H)$ that is a homeomorphism. Note that elements of $\mathrm{Imm}(K, H)$ that agree on a vertex of K are equal.

Lemma 7.7 *If $K < H$ are finitely generated subgroups of F_n then there is a (necessarily unique) label-preserving immersion $J_{K,H}\colon (\Sigma_b(K), b_K) \to (\Sigma_b(H), b_H)$.*

Proof We recall Stallings' construction [1983, 5.4] of $\Sigma_b(K)$. Choose closed paths $\rho_1, \dots, \rho_m \subset G$ based at b that represent generators of $K < \pi_1(G, b)$. Define Γ_1 to be a rose of rank m with unique vertex b' and define $p_1\colon (\Gamma_1, b'_1) \to (G, b)$ to be an immersion on edges, mapping the i^{th} edge to ρ_i. Subdivide Γ_1 into edgelets labeled by edges of G to obtain $p_2\colon (\Gamma_2, b'_2) \to (G, b)$. The map p_2 factors into a sequence of edgelet folds $(\Gamma_2, b'_2) \to (\Gamma_3, b'_3) \to \cdots \to (\Gamma_k, b'_k)$ followed by an immersion $p_k\colon (\Gamma_k, b'_k) \to (G, b)$. Define $(\Sigma_b(K), b_K) = (\Gamma_k, b'_k)$ and $p_H = p_k$.

Since $K < H$, each ρ_i lifts to a closed edgelet-path ρ'_i in $\Sigma_b(H)$ based at b_H. Since the i^{th} edge of Γ_2 and ρ'_i agree as labeled edgelet-paths, there is an induced label-preserving map $q_2\colon (\Gamma_2, b') \to (\Sigma_b(H), b_H)$ satisfying $p_2 = p_H q_2$. Since p_H is an immersion, the edgelets that are identified by the folding maps $\Gamma_2 \to \Gamma_3 \to \cdots \to \Gamma_k$ are also identified by q_2. Thus, there exists a map $J_{K,H}\colon (\Gamma_k, b'_k) \to (\Sigma_b(H), b_H)$ such that $p_K = p_H q_k$. Since p_H and p_K are immersions, the same is true for $J_{K,H}$. □

Note that if $a \in F_n$ and $K^a < H$ then $K^{ha} < H$ for all $h \in H$. Let $\mathrm{RC}(K, H)$ be the set of right cosets of H in F_n such that $K^a < H$ for some (each) a representing that coset.

Lemma 7.8 *There is an algorithm with output a bijection* $\mathrm{Imm}(K,H) \leftrightarrow \mathrm{RC}(K,H)$*. In particular, there is an algorithm that produces coset representatives for the elements of* $\mathrm{RC}(K,H)$*.*

Proof $(\rightarrow)$ We associate a coset $Ha \in \mathrm{RC}(K,H)$ to $J \in \mathrm{Imm}(K,H)$ as follows. Choose a path $\xi \subset \Sigma(H)$ from c_H to $J(c_K)$ and note that $p_K(c_K) = p_H(J(c_K))$. Let $a \in \pi_1(G,b)$ be represented by the closed path $[p_H(\beta_H\xi)p_K(\overline{\beta}_K)] \subset G$, where $[\cdot]$ indicates tightening. Each $x \in K$ is represented in $\pi_1(G,b)$ by $p_K(\beta_K\gamma\overline{\beta}_K)$ for some closed path $\gamma \subset \Sigma(K)$ based at c_K. It follows that x^a is represented in $\pi_1(G,b)$ by

$$[p_H(\beta_H\xi)p_K(\gamma)p_H(\overline{\xi}\overline{\beta}_H)] = [p_H(\beta_H\xi)p_H(J(\gamma))p_H(\overline{\xi}\overline{\beta}_H)] = p_H(\beta_H[\xi J(\gamma)\overline{\xi}]\overline{\beta}_H),$$

which represents an element in H. This proves that $K^a < H$. If ξ is replaced by another path ξ' connecting c_H to $J(c_K)$ then a is replaced by ha, where $h \in H$ is represented by $[p_H(\beta_H\xi'\overline{\xi}\overline{\beta}_H)]$. Thus, Ha is independent of the choice of ξ. If J is replaced with $J' \neq J$ and if $\eta \subset \Sigma(H)$ is a path connecting $J(c_K)$ to $J'(c_K)$, then ξ is replaced with $\xi\eta$ and a is replaced with $a' = da$, where d is represented in $\pi_1(G,b)$ by $[p_H(\beta_H\xi\eta\overline{\xi}\overline{\beta}_H)]$. Since η is not a closed path, d does not lift into $\Sigma(H)$ and a' does not belong to the same right coset of H as a. This shows that $J \mapsto Ha$ defines an injection from $\mathrm{Imm}(K,H)$ to $\mathrm{RC}(K,H)$.

$(\leftarrow)$ We begin the proof of surjectivity by constructing $\Sigma_b(K^a)$ from $\Sigma_b(K)$. Represent a in $\pi_1(G,b)$ by a closed edge-path $\alpha \subset G$ based at b and let β' be the edgelet path labeled by the path in G obtained by tightening $\alpha p_K(\beta_K)$. Define Σ' from the disjoint union of β' and a copy $(\Sigma'(K), c')$ of $(\Sigma(K), c_K)$ by identifying the terminal endpoint of β' with c'. The labeling on edgelets induces $p' \colon (\Sigma', b') \to (G,b)$, where b' is the initial vertex of β'. The image of the injection $\pi_1(\Sigma',b') \to \pi_1(G,b)$ induced by p' equals K^a. If p' is an immersion then $(\Sigma(K^a), \beta_{K^a}, c_{K^a}) = (\Sigma', \beta', c')$. Otherwise, $\Sigma(K^a)$ is obtained from Σ' by folding a maximal initial edgelet-subpath of $\overline{\beta}'$ with an edgelet-subpath $\mu \subset \Sigma'(K)$ that begins at c'. In this case, b_{K^a} is the folded image of b' and c_{K^a} is the terminal endpoint of μ.

Continuing with the above notation, define the equivalence $J_{K,a} \in \mathrm{Imm}(K, K^a)$ to be the identifying homeomorphism from $(\Sigma(K), c_K)$ to $(\Sigma'(K), c')$. Assuming that $K^a < H$, apply Lemma 7.7 and define $J_{a,K,H} = J_{K^a,H}|\Sigma(K^a) \circ J_{K,a} \in \mathrm{Imm}(K,H)$. By construction, $[\alpha p_K(\beta_K)] \subset G$ lifts to a the path in $\Sigma_b(H)$ from b_H to $J_{a,K,H}(c_K)$. Writing this path as $\beta_H\xi$, we have that $[p_H(\beta_H\xi)p_K(\overline{\beta}_K)] = \alpha$ and hence that (in the notation of the first paragraph of this proof) $J_{a,K,H} \mapsto Ha$. □

We will need the following well-known result.

Corollary 7.9 *If* $H < F_n$ *is a finitely generated and* $H^a < H$ *for* $a \in F_n$*, then* $H^a = H$*.*

Proof The obvious induction argument shows that $H^{a^p} < H^{a^{p-1}} < \cdots < H^a < H$ for all $p \geq 1$. Each Ha^s for $s \geq 1$ is therefore an element $\mathrm{RC}(H,H)$, which is finite by Lemma 7.8. It follows that $Ha^s = Ha^t$ for some $s \neq t$ and hence that $a^p \in H$ for some $p \geq 1$. Thus $H^{a^p} = H$, which implies that $H < H^a < H$ and hence that $H = H^a$. □

The following three algorithms are easy consequences of Lemma 7.8.

Lemma 7.10 *There is an algorithm that decides if a given pair H and K of finitely generated subgroups of F_n are conjugate, and if so produces an element $a \in F_n$ satisfying $K^a = H$.*

Proof We continue with notation from the proof of Lemma 7.8. If $K^a = H$ then $J_{K^a,H}$ is the identity and hence $J_{a,K,H} \in \mathrm{Imm}(K, H)$ is an equivalence. This shows that if $\mathrm{Imm}(K, H)$ does not contain an equivalence then K and H are not conjugate. If $\mathrm{Imm}(K, H)$ does contain an equivalence J, apply Lemma 7.8 to J and $J^{-1} \in \mathrm{Imm}(H, K)$ to produce $a, b \in F_n$ such that $K^a < H$ and $H^b < K$. From $H^{ab} < H$ and Corollary 7.9 it follows that $H^{ab} = H$ and hence that $H = (H^b)^a < K^a < H$, which implies that $K^a = H$. □

Lemma 7.11 *The normalizer $N(H)$ of a finitely generated subgroup $H < F_n$ is finitely generated. We have an algorithm that produces coset representatives $\{a_1, \dots, a_p\}$ of H in $N(H)$.*

Proof Corollary 7.9 implies that $N(H)/H = \mathrm{RC}(H, H)$. Lemma 7.8 therefore completes the proof. □

Lemma 7.12 *If $K < H < F_n$ are finitely generated subgroups then the set of subgroups of H that are F_n-conjugate to K determine finitely many H-conjugacy classes. There is an algorithm that produces representatives $K_1, \dots, K_p$ of these H-conjugacy classes.*

Proof If $K^{a_1}, K^{a_2} < H$ then K^{a_1} and K^{a_2} determine the same H-conjugacy class if and only if $ha_1 = a_2$ for some $h \in H$. Lemma 7.8, which produces representatives of the elements of $\mathrm{RC}(K, H)$, therefore completes the proof. □

7.3 Good conjugacy pairs

In addition to conjugacy classes of finitely generated subgroups of F_n, our adaptation of Gersten's algorithm will also take conjugacy pairs as input; see Notation 11.1. If H_1 and H_2 are subgroups of H and the natural map $H_1 * H_2 \to H$ is an isomorphism then we say that H is the *internal free product* of H_1 and H_2. If $A < B < F_n$ then $[A]_B$ denotes the conjugacy class of A in B. If $B = F_n$ then we sometimes suppress the subscript.

Definition 7.13 (good conjugacy pairs) For $H_1, H_2 < F_n$, the conjugacy pair $[H_1, H_2]_{F_n}$ is *good* if $\langle H_1, H_2 \rangle$ is the internal free product of H_1 and H_2.

The next lemma collects some facts about good pairs.

Lemma 7.14 *Let $H_1, H_2 < F_n$ be finitely generated.*

(1) *$[H_1, H_2]$ is good if and only if* $\mathrm{rank}(\langle H_1, H_2 \rangle) = \mathrm{rank}(H_1) + \mathrm{rank}(H_2)$.

(2) *If $[H_1, H_2]$ is good, then ∂H_1 and ∂H_2 are disjoint.*

Proof The natural map $H_1 * H_2 \to \langle H_1, H_2 \rangle$ is surjective. Since finitely generated free groups are Hopfian (surjective endomorphisms are isomorphisms) [Magnus et al. 1966, Theorem 2.13], the "only if" direction of (1) follows. The "if" direction of (1) is obvious.

(2) This follows from $\partial H_1 \cap \partial H_2 = \partial(H_1 \cap H_2)$; see the first item in Remark 7.4. $\square$

Our next goal is necessary and sufficient conditions for two good conjugacy pairs to be equal. We begin with an important special case.

Lemma 7.15 *Suppose that $K_1, K_2, L_1, L_2 < F_n$ are finitely generated, that $[K_1, K_2]_{F_n}$ and $[L_1, L_2]_{F_n}$ are good conjugacy pairs, and that $\langle K_1, K_2 \rangle = \langle L_1, L_2 \rangle = F_n$. Then the following are equivalent.*

(1) $[K_1, K_2]_{F_n} = [L_1, L_2]_{F_n}$.

(2) $[K_1]_{F_n} = [L_1]_{F_n}$ *and* $[K_2]_{F_n} = [L_2]_{F_n}$.

Proof (1) $\Longrightarrow$ (2) If $[K_1, K_2]_{F_n} = [L_1, L_2]_{F_n}$ then by definition there is $g \in F_n$ such that $(K_1^g, K_2^g) = (L_1, L_2)$.

(2) $\Longrightarrow$ (1) By hypothesis there are $g_i \in F_n$ such that $K_i^{g_i} = L_i$. In particular, $\Delta := (i_{g_1} | K_1) * (i_{g_2} | K_2) \in \mathsf{Aut}(F_n)$ represents an element $\delta \in \mathsf{Out}(F_n)$. Let r_i be the rank of K_i, let A_i be a rose with rank r_i whose petals are labeled by a basis for K_i and let A be the rose of rank n obtained from A_1 and A_2 by identifying their unique vertices v_1 and v_2 to a single vertex v. Blow up v to an arc. More precisely, let X be the graph obtained from the disjoint union of A_1, A_2 and a vertex w by adding oriented edges E_1 and E_2 connecting w to v_1 and v_2, respectively. Denote the arc $\bar{E}_1 E_2 \subset X$ by E and the subgraph $A_i \cup E_i \subset X$ by X_i. Identify $\pi_1(X_i, w)$ with $\pi_1(A_i, v_i) = K_i$ via the map $q_i : X_i \to A_i$ that collapses E_i to v. Let $q : X \to A$ be the map that collapses E to v. If $\alpha_i \subset A$ is the closed path based at v that represents g_i then there is a unique closed path $\beta_i \subset X$ based at w that satisfies $q_\#(\beta_i) = \alpha_i$. The map $f : X \to X$ defined by $f | A_i =$ identity and $f(E_i) = [\beta_i E_i]$ induces the automorphism Δ and so is a homotopy equivalence. Homotop f rel $A_1 \cup A_2$ to a map $f' : X \to X$ whose restriction to E is an immersion. Then f' is a topological representative of δ and [Bestvina et al. 2000, Corollary 3.2.2] implies that $f'(E) = \bar{\gamma}_1 E \gamma_2$ for some (necessarily closed) paths $\gamma_i \subset A_i$. If $k_i \in K_i$ is represented by the homotopy class of γ_i, then f' induces the automorphism $\Delta' := (i_{k_1} | K_1) * (i_{k_2} | K_2)$. There exists $h \in F_n$ such that $\Delta = i_h \Delta'$. We have $i_{g_i} = i_{h k_i}$ and hence $i_{g_1 k_1^{-1}} = i_{g_2 k_2^{-1}}$. Thus

$$\begin{aligned}[L_1, L_2]_{F_n} = [K_1^{g_1}, K_2^{g_2}]_{F_n} = [K_1^{g_1}, (K_2^{g_1^{-1} g_2})^{g_1}]_{F_n} = [K_1, K_2^{g_1^{-1} g_2}]_{F_n} \\ = [K_1, K_2^{k_1^{-1} k_2}]_{F_n} = [K_1^{k_1}, K_2^{k_2}]_{F_n} = [K_1, K_2]_{F_n} \qquad .\end{aligned}$$

$\square$

For each finitely generated $L < F_n$, we define a function f_L as follows. The domain of f_L is the set of good conjugacy pairs $[K_1, K_2]$ with $K := \langle K_1, K_2 \rangle$ conjugate to L. Any $g \in F_n$ such that $K^g = L$ is well-defined up to the normalizer $N(L)$ of L in F_n. That is, if $L = K^g = K^{g'}$ then $g' g^{-1} \in N(L)$.

Hence $([K_1^g]_L, [K_2^g]_L)$ is well-defined up to the diagonal action of $N(L)$ (equivalently $N(L)/L$) on the set of pairs of conjugacy classes of subgroups of L. We define $f_L([K_1, K_2])$ to be the $N(L)/L$ orbit of $([K_1^g]_L, [K_2^g]_L)$. Note that if $K = L$ and $\xi_1, \dots, \xi_r$ are coset representatives of L in $N(L)$ then $f_L([K_1, K_2]) = \{([K_1^{\xi_1}]_L, [K_2^{\xi_1}]_L), \dots, ([K_1^{\xi_r}]_L, [K_2^{\xi_r}]_L)\}$.

Remark 7.16 Suppose $x \in F_n$ and $L^x = L'$. Then $\mathsf{Domain}(f_L) = \mathsf{Domain}(f_{L'})$ and conjugation i_x by x induces a bijection (which we give the same name) $i_x\colon \mathsf{Codomain}(f_L) \to \mathsf{Codomain}(f_{L'})$ given by mapping the $N(L)$-orbit of $([L_1]_L, [L_2]_L)$ to the $N(L')$-orbit of $([L_1^x]_{L'}, [L_2^x]_{L'})$. It is an easy check that $f_{L'} = i_x \circ f_L$.

Lemma 7.17 *Suppose $K_1, K_2, L_1, L_2 < F_n$ are finitely generated and that $[K_1, K_2]_{F_n}$ and $[L_1, L_2]_{F_n}$ are good conjugacy pairs. Set $K := \langle K_1, K_2\rangle$ and $L := \langle L_1, L_2\rangle$. Then the following are equivalent*:

(1) $[K_1, K_2]_{F_n} = [L_1, L_2]_{F_n}$.

(2) *There is $g \in F_n$ such that $K^g = L$, $[K_1^g]_L = [L_1]_L$ and $[K_2^g]_L = [L_2]_L$.*

(3) $[K] = [L]$ *and* $f_L([K_1, K_2]) = f_L([L_1, L_2])$.

(4) $[K] = [L]$ *and, for some (any) $H < F_n$ with $[H] = [L] = [K]$, $f_H([K_1, K_2]) = f_H([L_1, L_2])$.*

Proof (1) $\Longrightarrow$ (2) If $[K_1, K_2]_{F_n} = [L_1, L_2]_{F_n}$ then by definition there is $g \in F_n$ such that $(K_1^g, K_2^g) = (L_1, L_2)$.

(2) $\Longrightarrow$ (1) By Lemma 7.15 applied to K_i^g and L_i with L playing the role of F_n we have $[K_1^g, K_2^g]_L = [L_1, L_2]_L$. In particular, $[K_1, K_2]_{F_n} = [L_1, L_2]_{F_n}$.

(2) $\Longrightarrow$ (3) This is clear from the definition of f_L.

(3) $\Longrightarrow$ (2) Suppose $[K] = [L]$ and $f_L([K_1, K_2]) = f_L([L_1, L_2])$. By the former there is $g' \in F_n$ such that $K^{g'} = L$ and by the latter there is $n \in N(L)$ such that $([K_1^{g'}]_L, [K_2^{g'}]_L) = ([L_1^n]_L, [L_2^n]_L)$. Take $g = n^{-1} g'$.

(3) $\Longleftrightarrow$ (4) This follows directly from Remark 7.16. □

Corollary 7.18 *There is an algorithm with input two good conjugacy pairs $[K_1, K_2]$ and $[L_1, L_2]$ of finitely generated subgroups of F_n, and output* `YES` *or* `NO` *depending on whether or not $[K_1, K_2] = [L_1, L_2]$.*

Proof Apply Lemma 7.10 to decide if K and L are F_n-conjugate. If not, then output `NO`. Otherwise, Lemma 7.10 gives $x \in F_n$ such that $K^x = L$ and we replace K_1 and K_2 by K_1^x and K_2^x so that now $K = L$. Apply Lemma 7.11 to produce coset representatives $\xi_1, \dots, \xi_r$ of L in $N(L)$. According to Lemma 7.17, $[K_1, K_2] = [L_1, L_2]$ if and only if $[K_1^{\xi_i}]_L = [L_1]_L$ and $[K_2^{\xi_i}]_L = [L_2]_L$ for some $1 \le i \le r$. This can be checked by applying Lemma 7.10 with F_n replaced by L. □

The following lemma is used in Lemma 7.21 to determine which pairs of conjugacy classes correspond to good conjugacy pairs.

Lemma 7.19 *There is algorithm with input two finitely generated subgroups $K_1, K_2 < F_n$ and output* YES *or* NO *depending on whether or not there exist $K_i' < F_n$ such that $[K_i'] = [K_i]$ and such that F_n is the internal free product of K_1' and K_2'. If* YES *then one such K_1' and K_2' are produced.*

Proof Choose any finitely generated subgroups A_i such $\operatorname{rank}(A_i) = \operatorname{rank}(K_i)$ and such that F_n is the internal free product of A_1 and A_2. K_1' and K_2' exist if and only if there is a $\theta \in \mathsf{Out}(F_n)$ such that $\theta([K_i]) = [A_i]$. The existence of such a θ can be checked using Gersten's generalization [1984] of Whitehead's theorem [Bestvina et al. 2023], which appears as Theorem 10.2 in this paper. Additionally, the algorithm produces such a θ if one exists; we take $K_i' = \Theta^{-1}(A_i)$, where $\Theta \in \theta$. □

Notation 7.20 $\mathcal{C}(F_n)$ denotes the set of conjugacy classes of finitely generated subgroups of F_n.

To aid in working with good conjugacy pairs, we relate them to ordered triples in $\mathcal{C}(F_n)$. Consider the following map from good conjugacy pairs to ordered triples in $\mathcal{C}(F_n)$:

$$[H_1, H_2] \mapsto ([H_1], [H_2], [H]), \quad \text{where } H := \langle H_1, H_2 \rangle. \tag{7-1}$$

Lemma 7.21 *We have an algorithm with input a good conjugacy pair $[H_1, H_2]$ and output a finite enumeration of the fiber F of the above map* (7-1) *containing $[H_1, H_2]$.*

Proof Consider the map induced by f_H from F to $\{([K_1]_H, [K_2]_H)\}/N(H)$, where K_i ranges over finitely generated subgroups of H such that H_i and K_i are conjugate in F_n. By Lemma 7.17, this map is injective. So, it remains to produce an element of F for each element of the image. The $[K_i]_H$ can be finitely enumerated by Lemma 7.12. By Lemma 7.19 we can decide if $([K_1]_H, [K_2]_H)$ represents an element of the image. Applying Lemma 7.11 and then Lemma 7.10 (with F_n replaced by H) we can decide if two pairs $([K_1]_H, [K_2]_H)$ and $([K_1']_H, [K_2']_H)$ are in the same $N(H)$-orbit and so remove redundancy from our list. □

Lemma 7.22 *We have an algorithm with input an ordered triple $([H_1], [H_2], [H])$ of elements of $\mathcal{C}(F_n)$ and output* YES *or* NO *depending on whether or not the fiber F of the above map* (7-1) *is empty. Further, if* NO*, the algorithm also outputs an element of F.*

Proof Our goal is to either find subgroups K_i in the same F_n conjugacy class as H_i such that $K_i < H$ and such that H is the internal free product of K_1 and K_2, or to conclude that no such K_i exist.

By Lemma 7.8, we can compute coset representatives for the elements of $\mathrm{RC}(H_i, H)$. If $\mathrm{RC}(H_1, H) = \varnothing$, then no element of the F_n-conjugacy of H_1 is a subgroup of H and we return YES. Similarly, return YES if $\mathrm{RC}(H_2, H) = \varnothing$. Otherwise, choose b_i representing a coset in $\mathrm{RC}(H_i, H)$. Replacing H_i by $H_i^{b_i}$ we may assume that $H_i < H$.

Lemma 7.11 produces coset representatives $\{a_1, \dots, a_p\}$ of H in $N(H)$. Thus a subgroup $K_i < H$ is in the same F_n conjugacy class as H_i if and only if it is in the same H-conjugacy class as $H_i^{a_j}$ for some $1 \le j \le p$. Order the pairs $(H_1^{a_j}, H_2^{a_k})$ lexicographically on $1 \le j, k \le p$. Apply Lemma 7.19 with F_n replaced by H and with (K_1, K_2) replaced by the first pair on the list $(H_1^{a_1}, H_2^{a_1})$ to either produce $K_1, K_2 < H$ such that

- K_1 is in the same H conjugacy class as $H_1^{a_1}$,
- K_2 is in the same H conjugacy class as $H_2^{a_1}$,
- H is the internal free product of K_1 and K_2,

or to conclude that no such K_1 and K_2 exist. In the former case return `NO` and $[K_1, K_2]$. In the latter case proceed on to the next pair on the list. Continue until you either return `NO` and the desired $[K_1, K_2]$, or reach the end of the list, in which case return `YES`. □

Corollary 7.23 *We have an algorithm with input an ordered triple* $([H_1], [H_2], [H])$ *of elements of* $\mathcal{C}(F_n)$ *and output a finite enumeration of the fiber* F *of the above map* (7-1) *over* $([H_1], [H_2], [H])$.

Proof Use Lemma 7.22 to determine if F is empty or not and obtain an element $[H_1', H_2'] \in F$ if not. Input $[H_1', H_2']$ into the algorithm of Lemma 7.21 to enumerate F. □

We will also use conjugacy pairs that aren't necessarily good.

Lemma 7.24 *Consider the set of conjugacy pairs of the form* $[H, A]$ *with* $A < H < F_n$ *all finitely generated and nontrivial.* (*In particular this pair is not good.*)

1. (1) *Two such* $[H, A]$ *and* $[H', A']$ *are equal if and only if there is* $g \in F_n$ *such that* $H^g = H'$ *and* $[A^g]_{H'} = [A']_{H'}$. *In particular,* $[H, A] = [H, A']$ *if and only if* A *and* A' *are in the same orbit of the action of* $N(H)$ *on* H.
2. (2) *The map from the set of such pairs to ordered sequences in* $\mathcal{C}(F_n)$ *given by* $[H, A] \mapsto ([H], [A])$ *has fibers that can be finitely enumerated.*

Proof (1) The "only if" direction is obvious. The "if" direction follows from the fact that if $[A^g]_{H'} = [A']_{H'}$ then there exists $h' \in H'$ such that $A^{h'g} = A'$ for some $h' \in H'$.

(2) Given finitely generated nontrivial subgroups $K, L < F_n$, compute $\mathrm{RC}(L, K)$ by applying Lemma 7.8. If $\mathrm{RC}(L, K) = \varnothing$, then L is not conjugate into K so the fiber over $([K], [L])$ is empty. Otherwise, choose b representing a coset in $\mathrm{RC}(L, K)$. Replacing L by L^b we may assume that $L < K$. Apply Lemma 7.11 to produce coset representatives $\xi_1, \dots, \xi_p$ of K in $N(K)$. The fiber containing $[K, L]$ equals $\{[K, L^{\xi_1}], \dots, [K, L^{\xi_p}]\}$ by (1). □

8 Computable G-sets

The ultimate goal of this paper is to provide an algorithm solving the conjugacy problem for $\mathsf{UPG}(F_n)$, ie Theorem 1.1. We will need other algorithms as part of our solution. In this section and the following two (Sections 8, 9 and 10), we formalize some of the algorithmic aspects present in the $\mathsf{Out}(F_n)$-setting. In particular, we provide what could be viewed as a "data structure" for the input and output of our algorithms. These sections require no knowledge of F_n and are independent of the rest of the paper.

Definition 8.1 (computable)

- A function $f: X \to Y$ is *computable* if it comes equipped with an algorithm with input $x \in X$ and output $f(x) \in Y$.
- An *enumeration* of a set X is a surjection $\mathbb{N} \to X$. A *finite enumeration* of X is a surjection $\{1, 2, \dots, N\} \to X$. The *index* of $x \in X$ is the minimal n such that $n \mapsto x$.
- A set X is *computable* if it comes equipped with a computable enumeration $\mathbb{N} \to X$ and an algorithm with input $x, x' \in X$ and output YES or NO depending on whether or not $x = x'$. By default, the empty set is computable. We sometimes write $X = (x_1, x_2, \dots)$ to indicate the enumeration. See Lemma 8.2.
- A group G is *computable* if the underlying set is computable and it comes equipped with a third algorithm with input $\theta, \theta', \theta'' \in G$ and output YES or NO depending on whether or not $\theta\theta' = \theta''$.
- A G-set X is *computable* if G and X are computable and it comes equipped with yet another algorithm with input $\theta \in G$, $x, x' \in X$ and output YES or NO depending on whether or not $\theta(x) = x'$.

Lemma 8.2 *If $X = (x_1, x_2, \dots)$ is a computable set then we have an algorithm with input $x \in X$ and output the index of x.*

Proof Starting with $i = 1$, iteratively check if $x = x_i$. □

To see how $\mathsf{Out}(F_n)$ and our $\mathsf{Out}(F_n)$-sets are enumerated and that they are computable, see Section 11.

A set Y of interest is often the quotient of a computable set X, ie $Y = X/{\sim}$ for some equivalence relation $\sim$. We want to use X to give Y the structure of a computable set. We view elements of X as *representatives* of elements of Y and always give elements $y \in Y$ as $y = [x]$, where $x \in X$ and $[x]$ is the equivalence class of x.

Lemma 8.3 *Suppose X is a computable set and $Y = X/{\sim}$ is a quotient of X. If we have an algorithm with input $x, x' \in X$ and output* YES *or* NO *depending on whether or not $[x] = [x']$, then Y is computable.*

There are the obvious generalizations for groups, etc.

Proof The computable enumeration of Y maps $i \in \mathbb{N}$ to $[x_i] \in Y$. Given input $y = [x]$, $y' = [x']$, we can use the algorithm in the hypothesis to output YES or NO depending on whether or not $[x] = [x']$, ie whether or not $y = y'$. □

Example 8.4 Suppose we are given a finite generating set for a group G. Elements of G are represented as finite words in the generators and their inverses. This set X of finite words can be computably enumerated, say using length, and X is computable. The composition of the enumeration for X and the evaluation map $e: X \to G$ computably enumerates G. If we have an algorithm with input $x, x' \in X$ and output `YES` or `NO` depending on whether or not $e(x) = e(x')$ then G is computable. This is the case, for example, if we are given a finite presentation for G and an algorithm solving the word problem for this presentation.

Lemma 8.5 (1) *Suppose Z is a subset of the computable set X. If we have an algorithm with input $x \in X$ and output* `YES` *or* `NO` *depending on whether or not $x \in Z$, then Z is computable.*

(2) *If X and Y are computable sets then $X \times Y$ is a computable set.*

There are the obvious generalizations for groups, etc.

Proof (1) If Z is empty then it is computable by definition. Suppose $Z \neq \varnothing$. The computable enumeration $f: \mathbb{N} \to Z$ is given as follows. Applying the algorithm in the hypothesis a finite number of times, we can find the minimal $i \in \mathbb{N}$ such that $x_i \in Z$. Define $f(j) = x_i$, for $1 \leq j \leq i$. If $n > i$, then $f(n) = x_n$ if $x_n \in Z$ and $f(n) = f(n-1)$ otherwise.

The proof of (2) is left to the reader. □

In the setting of Lemma 8.5(1), we view elements of Z as given to us as elements of X that are in Z. One reason for the rather odd looking enumeration f in the proof is that we have to make sure that f is defined on all of $\mathbb{N}$. (Consider the case where Z is finite.)

We now collect some basic properties of computable groups.

Lemma 8.6 *Let $G = (g_1, g_2, \ldots)$ and $G' = (g'_1, g'_2, \ldots)$ be computable groups.*

(1) *We have algorithms*
   (a) *with input $g \in G$ and output the index of G,*
   (b) *with input $g, h \in G$ and output the index of gh,*
   (c) *with output the index of* 1, *and*
   (d) *with input g and output the index of g^{-1}.*

(2) *We have an algorithm with input a finite word w in $\{g_1, g_2, \ldots\}^{\pm 1}$ and output the index of w in G. In particular, we have an algorithm to solve the word problem in G.*

(3) *Suppose we are given a finite generating set $\mathcal{G} = \{h_1, \ldots, h_N\} \subset \{g_1, g_2, \ldots\}$ for G. Then we have an algorithm with input $g \in G$ and output a word w with letters in $\mathcal{G}$ such that $g = w$ in G.*

(4) *Suppose $f: G \to G'$ is a homomorphism. If we are given a finite generating set $\mathcal{G} \subset \{g_1, g_2, \ldots\}$ for G and $f(\mathcal{G})$, then f is computable (with algorithm given in the proof).*

(5) *If $f: G \to G'$ is a computable homomorphism, then* $\mathrm{Ker}(f)$ *is computable.*

Proof (1) For (a), see Lemma 8.2. For (b), (c) and (d): starting with $i = 1$, iteratively use the algorithm that comes with a computable group to respectively check: (b) if $gh = g_i$; (c) if $g_i g_i = g_i$; and (d): if $gg_i = 1$.

(2) Use (d) to remove all negative exponents in w. Then use (b) to iteratively reduce the length of w by replacing consecutive letters $g_i g_j$ with a single letter g_k.

(3) Enumerate the words in $\mathcal{G}$ (say using length). Iteratively check if g is the N^{th} word.

(4) To compute $f(g_i)$, use (3) to write g_i as a word w in $\mathcal{G}$. Then $f(w)$ is a word in $f(\mathcal{G})$. Finally, use (2) to find the index of $f(w)$.

(5) This follows from Lemma 8.5 since we can algorithmically check if $f(g_i) = 1$ in G'. □

Remark 8.7 Using Lemma 8.6(1), we may rewrite a given finite subset of $\{g_1, g_2, \dots\}^{\pm 1}$ as a finite subset of $\{g_1, g_2, \dots\}$. In particular, if $\mathcal{G}$ generates G then we may algorithmically compute a finite subset of $\{g_1, g_2, \dots\}$ that is a symmetric generating set.

Example 8.8 A computable group need not be finitely generated; a finitely generated computable group need not be finitely presented; etc. Indeed, the kernel of $f : (F_2)^n \to \mathbb{Z}$ which sends each basis element to $1 \in \mathbb{Z}$ has varying finiteness properties depending on n; see [Bestvina and Brady 1997]. By Lemma 8.6(5), $\mathsf{Ker}(f)$ is computable.

Lemma 8.9 *If the G-set X is computable and $x \in X$, then the stabilizer G_x in $G = (g_1, g_2, \dots)$ of x is computable.*

Proof This follows from Lemma 8.5 applied to $G_x < G$, since we can algorithmically check whether $g_i(x) = x$. □

9 Finite presentations and finite-index subgroups

The following lemma is useful for finding presentations of finite-index subgroups of a finitely presented group. It is well-known — see eg [Lyndon and Schupp 1977, Chapter 2, Section 4, The Reidemeister–Schreier method] — but for completeness we provide a proof.

Lemma 9.1 *There is an algorithm that takes as input*

- *a finite presentation for a computable group $G = (g_1, g_2, \dots)$,*
- *the multiplication table for a finite group Q,*
- *a computable surjection $P : G \to Q$, and*
- *a subgroup Q' of Q given as a finite list of elements of Q,*

and outputs

(1) *a finite presentation for the subgroup* $G' := P^{-1}(Q')$ *of* G, *and*

(2) *finite sets of left and right coset representatives for* $G' < G$.

Remark 9.2 In some applications, G will act on a finite set S and so we have a homomorphism $G \to \mathsf{Perm}(S)$ to the permutation group of S. The group Q will be the image of this map and $P: G \to Q$ the induced map. This will allow us to compute the multiplication table for Q.

Proof of Lemma 9.1 (1) Let $G = \langle h_1, \dots, h_i \mid r_1, \dots, r_j \rangle$ be the given finite presentation for G, where the generators are in $\{g_1, \dots\}$ and the relations are words in the generators; see Lemma 8.6(2). Let X_G denote its presentation 2-complex. We assume the reader is familiar with obtaining a finite presentation for a group H from a finite based 2-complex with fundamental group H. Hence (1) is reduced to constructing the finite based cover $X_{G'}$ of X_G whose fundamental group has image in the fundamental group of X_G equal to G'.

Set $k := [G : G'] = [Q : Q']$ and note that $k \le |Q|$. Then every index k based cover of X_G has $k \cdot i$ 1-cells and $k \cdot j$ 2-cells. Further, if $|r_q|$ denotes the length of r_q as a word in $\{h_1, \dots, h_i\}$, then each 2-cell in the cover has boundary of length at most $k \cdot \max\{|r_q| \mid 1 \le q \le j\}$. Hence we can construct all based covers of X_G of index k. Examine each in turn to check whether the image K of its fundamental group in X_G is G'. This can be done by checking whether the image in X_G of a generating set for the fundamental group of the cover has P-image contained in Q'. (Indeed, if so then $K < G'$ but both K and G' have the same index in G.) Since G' has index k in G, we are guaranteed that one of these covers satisfies $K = G'$. This completes the proof of (1).

(2) We find left coset representatives, the other case being symmetric. Using the hypotheses on Q, choose a set S_Q of left cosets representatives for $Q' < Q$. Then $S_G \subset G$ is a set of left coset representatives for $G' < G$ if the restriction $P|S_G$ is injective with image S_Q. To find $g \in G$ with P-image $q \in Q$, the P-image of a set of generators for G generates Q and in Q we may write q in terms of these generators for Q. □

Given a short exact sequence $1 \to N \to G \xrightarrow{f} Q \to 1$, we are interested in finding a finite presentation for G from finite presentations for N and Q.

Lemma 9.3 *Suppose* $f: G \to Q$ *is a computable surjection between computable groups* G *and* Q.

Suppose we are given

- *a finite presentation for* $N := \mathsf{Ker}(f)$, *and*
- *a finite presentation for* Q (*for example if* Q *is finite and we are given the multiplication table for* Q *then this item is satisfied*).

Then we may find a finite presentation for G. (*In fact, one is constructed in the proof.*)

Proof Suppose that the finite presentation for N has generating set $\{g_{N,i} \mid i \in I\} \subset G$ and set of relators $\{r_{N,j} \mid j \in J\}$ and suppose that the finite presentation for Q has generating set $\{g_{Q,k} \mid k \in K\}$ and set of relators $\{r_{Q,l} \mid l \in L\}$.

For each $g_{Q,k}$, find an element $\hat{g}_{Q,k} \in G$ with image $g_{Q,k}$ in Q. This can be done algorithmically by iteratively searching for g_i such that $f(g_i) = g_{Q,k}$.

Set $\hat{r}_{Q,l} = r_{Q,l}(\hat{g}_{Q,k} \mid k \in K)$, ie $r_{Q,l}$ is a word in $\{g_{Q,k} \mid k \in K\}$ and $\hat{r}_{Q,l}$ denotes the same word in $\{\hat{g}_{Q,k} \mid k \in K\}$. The image of $\hat{r}_{Q,l}$ in Q represents 1_Q and so there is a word $n_{Q,l}$ in $\{g_{N,i} \mid i \in I\}$ such that $\hat{r}_{Q,l} = n_{Q,l}$ in G. Since N is normal, $\hat{g}_{Q,k}(g_{N,i})\hat{g}_{Q,k}^{-1} = n_{N,i,Q,k}$ for some word $n_{N,i,Q,k}$ in $\{g_{N,i} \mid i \in I\}$. Since G is computable, $n_{Q,l}$ and $n_{N,i,Q,k}$ can be found algorithmically; see Lemma 8.6(3). By [Bridson and Wilton 2011, Lemma 2.1], there is a finite presentation for G with

- generating set $\{g_{N,i}, \hat{g}_{Q,k} \mid i \in I, k \in K\}$, and
- set of relators that is the union of
  - $\{r_{N,j} \mid j \in J\}$,
  - $\{\hat{r}_{Q,l} = n_{Q,l} \mid l \in L\}$, and
  - $\{\hat{g}_{Q,k}(g_{N,i})\hat{g}_{Q,k}^{-1} = n_{N,i,Q,k} \mid i \in I, k \in K\}$. □

10 MW-algorithms

Our solution of the conjugacy problem for $\mathsf{UPG}(F_n)$ in $\mathsf{Out}(F_n)$ will use a generalization of an algorithm of Gersten that in turn generalizes algorithms of Whitehead and McCool. This section is devoted to describing these generalizations.

A set equipped with an action by a group G is a *G-set*. We will only consider *computable* G-sets; see Definition 8.1.

Definition 10.1 (property MW) A computable G-set X *satisfies property* MW (for McCool and Whitehead) if it comes equipped with an algorithm that takes as input $x, y \in X$ and outputs

(M) finite presentations for $G_x := \{\theta \in G \mid \theta(x) = x\}$ and for $G_y := \{\theta \in G \mid \theta(y) = y\}$, and

(W) `YES` or `NO` depending on whether or not there exists $\theta \in G$ such that $\theta(x) = y$ together with such a θ if `YES`.

We call such an algorithm an MW-*algorithm.* Sometimes we refer to an algorithm that satisfies item M (resp. item W) as an M-*algorithm* (resp. W-*algorithm*). Recall that G_x and G_y are computable by Lemma 8.9.

Of course our interest here is in the case $G = \mathsf{Out}(F_n)$ where the second item is associated with J H C Whitehead [1936a; 1936b] whose algorithm decides if there is θ taking one finite ordered set of conjugacy classes in F_n to another and produces such a θ if one exists. The first item is associated with McCool [1975], whose algorithm produces a finite presentation for the stabilizer of a finite ordered set of conjugacy classes of elements of F_n. S Gersten generalized the algorithms of Whitehead and McCool to finite sequences in the $\mathsf{Out}(F_n)$-set $\mathcal{C}(F_n)$; see Notation 7.20. ($\mathcal{C}(F_n)$ is shown to be a computable $\mathsf{Out}(F_n)$-set at the beginning of Section 11.) We state Gersten's result in a slightly weakened form.

Theorem 10.2 ([Gersten 1984, Theorems W&M], see also [Kalajdžievski 1992] and [Bestvina et al. 2023]) *The action of* $\mathsf{Out}(F_n)$ *on the set of finite ordered sets in* $\mathcal{C}(F_n)$ *satisfies property MW.*

We will define algebraic invariants for elements of $\mathsf{UPG}(F_n)$; see Section 13. An obstruction to $\phi, \psi \in \mathsf{UPG}(F_n)$ being conjugate in $\mathsf{Out}(F_n)$ is the existence of a $\theta \in \mathsf{Out}(F_n)$ taking the algebraic invariants of ϕ to those of ψ. More specifically, if such a θ does not exist then ϕ and ψ are not conjugate. If such a θ does exist, then we replace ϕ by ϕ^θ (or ψ by $\psi^{(\theta^{-1})}$) and reduce to the case where the algebraic invariants of ϕ and ψ agree. One step in our algorithm for the conjugacy problem for $\mathsf{UPG}(F_n)$ in $\mathsf{Out}(F_n)$ will be to check whether such a θ exists and to produce one if so.

Some of our invariants are elements of $\mathcal{C}(F_n)$ and so fit nicely into the setting of Gersten's theorem. We will extend Gersten's theorem so that it applies to all our invariants. These invariants are best described in terms of *iterated sets* of elements of $\mathcal{C}(F_n)$, or more generally in terms of sets with finite-to-one maps to iterated sets. Roughly, an iterated set in a G-set $\mathbb{A}$ is a finite set consisting of elements of $\mathbb{A}$ and previously produced iterated sets. The set may be ordered or not. We commonly take $\mathbb{A}$ to be $\mathcal{C}(F_n)$.

There are two main results. The first, Proposition 10.14, promotes MW-algorithms for finite ordered sets in $\mathbb{A}$ to MW-algorithms for the set $\overline{\mathsf{IS}}(\mathbb{A})$ (of equivalence classes) of iterated sets in $\mathbb{A}$. More specifically, it states that if the G-action on finite ordered subsets of $\mathbb{A}$ satisfies property MW then so does the G-action on $\overline{\mathsf{IS}}(\mathbb{A})$. The second, Corollary 10.22, is a method for enlarging $\mathbb{A}$.

After reviewing our invariants in Section 12 and defining the algebraic invariants in Section 13, we apply our generalized Gersten's algorithm to obtain a reduction of the conjugacy problem for $\mathsf{UPG}(F_n)$ in $\mathsf{Out}(F_n)$ to Proposition 14.7 in Lemma 14.8.

10.1 Iterated sets and their equivalence classes

Definition 10.3 A *rooted tree* $(T, *)$ is a finite, simplicial, directed tree T with a basepoint $*$ called the *root*. A valence 0 vertex (ie $T = \{*\}$) or a valence one vertex that is not the root is a *leaf*. The set of leaves in T is denoted by $\mathcal{L}(T)$. All edges are oriented away from the root. The set of edges exiting a vertex $x \in T$ is denoted by $\mathcal{E}_T(x)$. Paths are directed. We also may give some vertices an extra structure: a vertex x that is not a leaf is *ordered* if an order has been imposed on $\mathcal{E}_T(x)$. A vertex that is not a leaf and that is not ordered is *unordered*.

We view rooted trees $(T, *)$ as combinatorial objects. In particular, edges are specified by ordered pairs of vertices. For technical reasons having to do with computability, we require that all vertices of the trees we consider lie in a set V that we fix once and for all. (For our purposes, one can take V to be $\mathbb{N}$.)

Example 10.4 We will draw rooted trees with the root at the top. Ordered vertices are indicated by using dashed lines for its exiting edges. The imposed ordering is displayed from left to right. In the rooted tree below only the root is ordered. There are 4 leaves:

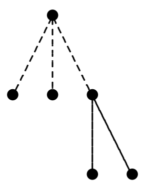

Definition 10.5 (iterated set) An *iterated set in a set* $\mathbb{A}$ is a rooted tree such that each leaf is labeled by an element of $\mathbb{A}$. Specifically, an iterated set in $\mathbb{A}$ is a pair $((T,*),\chi)$ where $(T,*)$ is a rooted tree and $\chi\colon \mathcal{L}(T)\to\mathbb{A}$ is a function. We do not assume that χ is one-to-one. We will often use sans serif capital letters for iterated sets and write, for example, $\mathsf{X}=((T,*),\chi)$. The *set of atoms of* X is $\chi(\mathcal{L}(T))$. We sometimes refer to $\mathbb{A}$ as *the set of atoms*. For $l\in\mathcal{L}(T)$, we sometimes refer to $\chi(l)$ as the *label* or *atom* of l. The set of iterated sets in $\mathbb{A}$ is denoted by $\mathsf{IS}(\mathbb{A})$.

Example 10.6 If we take $\mathbb{A}=\mathcal{C}(F_n)$, then a nested sequence $\vec{\mathcal{F}}=\mathcal{F}_1\sqsubset\cdots\sqsubset\mathcal{F}_N$ of free factor systems determines an iterated set in $\mathbb{A}$ as follows. First we identify each free factor system $\mathcal{F}_i=\{[F_{i,1}],\dots,[F_{i,m_i}]\}$ with an iterated set:

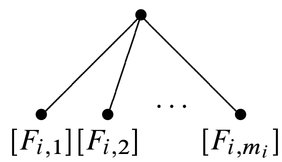

Then $\vec{\mathcal{F}}$ determines the ordered set $\{\mathcal{F}_1,\dots,\mathcal{F}_N\}$:

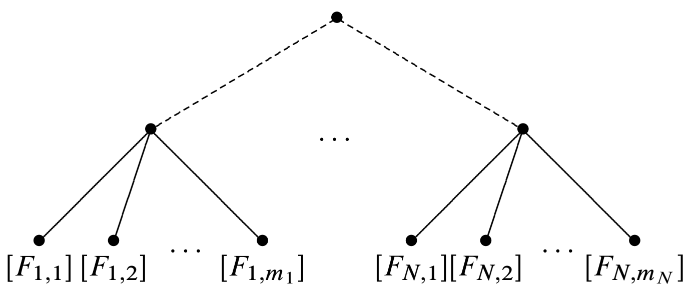

Definition 10.7 Let $\mathsf{X}=((T,*),\chi)$ and $\mathsf{X}'=((T',*'),\chi')$ be iterated sets in $\mathbb{A}$.

- An *order-preserving simplicial isomorphism* $f:(T,*)\to(T',*')$ is a simplicial isomorphism that satisfies:
  (1) A vertex x of T is ordered if and only if $f(x)$ is ordered and the induced map $\mathcal{E}_T(x)\to\mathcal{E}_{T'}(f(x))$ is order-preserving.

- An *equivalence* $\mathsf{f}\colon \mathsf{X} \to \mathsf{X}'$ is an order-preserving simplicial isomorphism $f : (T, *) \to (T', *')$ that additionally satisfies:

  (2) For $l \in \mathcal{L}(T)$, $\chi(l) = \chi'(f(l))$.

  Clearly equivalence induces an equivalence relation on $\mathsf{IS}(\mathbb{A})$.

- $\overline{\mathsf{IS}}(\mathbb{A})$ denotes the set of equivalence classes of iterated sets.

Remark 10.8 We will not need this, but if $\mathbb{A}$ is the set of objects of a category $\widehat{\mathbb{A}}$, then naturally so are $\mathsf{IS}(\mathbb{A})$ and $\overline{\mathsf{IS}}(\mathbb{A})$. A *morphism* $((T, *), \chi) \to ((T', *'), \chi')$ is an order-preserving simplicial isomorphism $f : (T, *) \to (T', *')$ together with a function $m\colon \chi(\mathcal{L}(T))$ into the morphisms of $\widehat{\mathbb{A}}$ such that, for $l \in \mathcal{L}(T)$, $m(\chi(l)) \in \mathsf{Hom}(\chi(l), \chi'(f(l)))$. An earlier version of this paper used a simplified, but more restrictive variant of this category, which was ultimately not needed.

10.2 Promoting property MW

Definition 10.9 Suppose G is a group and $\mathbb{A}$ is a G-set. Then $\mathsf{IS}(\mathbb{A})$ and $\overline{\mathsf{IS}}(\mathbb{A})$ are G-sets with actions given as follows. If $\theta \in G$ and $\mathsf{X} = ((T, *), \chi)$, then $\theta(\mathsf{X}) := ((T, *), \theta \circ \chi)$, ie $\theta(\mathsf{X})$ is obtained by relabeling $\mathcal{L}(T)$ according to θ. The G-action descends to $\overline{\mathsf{IS}}(\mathbb{A})$.

We want $\mathsf{IS}(\mathbb{A})$ and $\overline{\mathsf{IS}}(\mathbb{A})$ to be computable. This is the case if our set V of vertices and $\mathbb{A}$ are computable.

Lemma 10.10 (1) *If V is a computable set, then the set of rooted trees with vertices in V is computable.*

(2) *If additionally $\mathbb{A}$ is a computable set, then the sets $\mathsf{IS}(\mathbb{A})$ and $\overline{\mathsf{IS}}(\mathbb{A})$ are computable.*

(3) *If additionally $\mathbb{A}$ is a computable G-set, then the G-sets $\mathsf{IS}(\mathbb{A})$ and $\overline{\mathsf{IS}}(\mathbb{A})$ are computable.*

Proof (1) We view rooted trees $(T, *)$ as combinatorial objects. In particular, edges are specified by ordered pairs of vertices. To completely specify $(T, *)$ we also choose a root vertex, designate some vertices as ordered and choose an order on exiting edges of those vertices. Rooted trees T with vertices in V can be computably enumerated using, say, the sum $|T|$ of the number of vertices and largest index among the vertices of T. That is, to enumerate the set of rooted trees, first list all those with $|T| = 1$, then 2, etc. Two rooted trees are equal if and only if the they have the same vertices, edges, root, ordered vertices and same order on outgoing edges of ordered vertices.

(2) For an iterated set $\mathsf{X} = ((T, *), \chi)$, let $|\mathsf{X}|$ denote the sum of the number of vertices of $|T|$ and the largest index of a label (an element of $\chi(\mathcal{L}(T))$). $\mathsf{IS}(\mathbb{A})$ may be countably enumerated using $|\mathsf{X}|$. To enumerate $\mathsf{IS}(\mathbb{A})$, for example, list all X with $|\mathsf{X}| = 1$, then 2, etc. Two iterated sets are equal if and only if the underlying rooted trees are equal and the functions on the leaves are equal. The first condition can be checked by (1) and the second can be checked since $\mathbb{A}$ is computable. We can use the same

enumeration for $\overline{\mathsf{IS}}(\mathbb{A})$. Here if $\mathsf{X}_1 := ((T_1, *_1), \chi_1)$ and $\mathsf{X}_2 := ((T_2, *_2), \chi_2)$ represent respectively Q_1 and Q_2 in $\overline{\mathsf{IS}}(\mathbb{A})$, then $\mathsf{Q}_1 = \mathsf{Q}_2$ if and only if X_1 and X_2 are equivalent and this is a finite check. Indeed, finitely enumerate the set S of order-preserving simplicial isomorphisms $f : (T_1, *_1) \to (T_2, *_2)$. If S is empty then $\mathsf{Q}_1 \neq \mathsf{Q}_2$. Otherwise, if some $f \in S$ is an equivalence then $\mathsf{Q}_1 = \mathsf{Q}_2$ and if not then $\mathsf{Q}_1 \neq \mathsf{Q}_2$.

(3) Since $\mathbb{A}$ is a computable G-set, it is a finite check whether $\theta(\mathsf{X}_1) = \mathsf{X}_2$ and also whether $\theta(\mathsf{X}_1)$ and X_2 are equivalent. □

Assumption 10.11 Going forward, we assume that our fixed set V of vertices is computable; see Definition 10.3. In all applications, $\mathbb{A}$ will be computable.

Notation 10.12 Unless otherwise specified,

- $\mathbb{A}$ denotes a computable G-set,
- $\mathsf{X}, \mathsf{X}', \dots$ denote elements $((T, *), \chi), ((T', *'), \chi'), \dots$ of $\mathsf{IS}(\mathbb{A})$,
- $\mathsf{Q}, \mathsf{Q}', \dots$ denote elements of $\overline{\mathsf{IS}}(\mathbb{A})$ and are represented by $\mathsf{X}, \mathsf{X}', \dots$, and
- an equivalence $\mathsf{f} : \mathsf{X} \to \mathsf{X}'$ is given by $f : (T, *) \to (T', *')$.

Notation 10.13
- The map $\mathbb{A} \to \overline{\mathsf{IS}}(\mathbb{A})$ determined by $a \mapsto ((*, *), * \mapsto a)$ is a G-equivariant inclusion. In other words, map a to the trivial tree with vertex labeled a. Thus we may think of $\mathbb{A}$ as a subset of $\overline{\mathsf{IS}}(\mathbb{A})$.
- $\mathcal{S}_{\mathrm{or}}(\mathbb{A})$ denotes the subset of $\overline{\mathsf{IS}}(\mathbb{A})$ represented by iterated sets in which $*$ is ordered and $*$ is the initial endpoint of every edge of T. $\mathcal{S}_{\mathrm{or}}(\mathbb{A})$ is G-invariant. There is an obvious G-invariant bijection between the set of nonempty, finite, ordered sequences in $\mathbb{A}$ and $\mathcal{S}_{\mathrm{or}}(\mathbb{A})$. We pass back and forth between these two points of view whenever convenient.
- The G-set $\mathcal{S}_{\mathrm{un}}(\mathbb{A})$ is defined analogously where $*$ is unordered. There is an obvious G-invariant bijection between the set of nonempty, finite, multisets in $\mathbb{A}$ and $\mathcal{S}_{\mathrm{un}}(\mathbb{A})$.

Proposition 10.14 (promoting MW) *Let $\mathbb{A}$ be a computable G-set. If $\mathcal{S}_{\mathrm{or}}(\mathbb{A})$ satisfies property MW, then so does $\overline{\mathsf{IS}}(\mathbb{A})$.*

Proof We follow the conventions in Notation 10.12. First we provide a W-algorithm for $\overline{\mathsf{IS}}(\mathbb{A})$; ie an algorithm that either finds $\theta \in G$ satisfying $\theta(\mathsf{Q}) = \mathsf{Q}'$ or concludes that there is no such θ. Finitely enumerate the set S of order-preserving simplicial isomorphisms $f : (T, *) \to (T', *')$. If S is empty then return `NO`. Otherwise, start with the first element f of S. By hypothesis there is a W-algorithm that either finds $\theta \in G$ such that $\theta(\chi(l)) = \chi'(f(l))$ for each $l \in \mathcal{L}(T)$ or concludes that no such θ exists. If θ is found then f gives an equivalence $\theta(\mathsf{X}) \to \mathsf{X}'$ and our W-algorithm returns `YES` and θ. If no such θ exists, move on to the next element of S and try again. If after considering each element of the finite set S we have not returned `YES`, then return `NO`. (Equivalently, we could make the queries indexed by S in parallel. Note that a different choice of representatives X, X' for Q, Q' would give the same queries.)

The stabilizer G_{Q} in G of Q is computable by Lemma 8.9. For the M-algorithm, we will produce a finite presentation for G_{Q} by applying Lemma 9.3 to the short exact sequence induced by $\pi\colon G_{\mathsf{Q}} \to \mathsf{Perm}(A)$, where A denotes the set $\chi(\mathcal{L}(\mathsf{X}))$ of atoms of X. Since the kernel of π is the subgroup of G fixing each element of A, we can use the M-algorithm for $\mathcal{S}_{\mathrm{or}}(\mathbb{A})$ to produce a finite presentation for this kernel.

To apply Lemma 9.3, it remains to produce an element of G_{Q} realizing each element of the image of π. This is done as follows. Given $\sigma \in \mathsf{Perm}(A)$, use the W-algorithm for $\mathcal{S}_{\mathrm{or}}(\mathbb{A})$ to produce $\theta \in G$ realizing σ if such exists. If there is no such θ then σ is not in the image of π. Finally, use the computability of $\overline{\mathsf{IS}}(\mathbb{A})$ (Lemma 10.10(3)) to check if $\theta(\mathsf{X})$ is equivalent to X. If it is then σ is in the image of π (and is realized by θ) and otherwise it is not. (As above with the W-algorithm, it is easy to see that the choice of representative X for Q is immaterial. Also, the choice of θ does not matter.) □

For reference we record the following consequence of the previous proof (really just definitions).

Corollary 10.15 *If* $\mathsf{X} = ((T,*),\chi)$ *represents* $\mathsf{Q} \in \overline{\mathsf{IS}}(\mathbb{A})$, *then the subgroup of* G *fixing each* $\chi(l)$, $l \in \mathcal{L}(T)$, *has finite index in the stabilizer* G_{Q} *of* Q. □

10.3 More atoms

Proposition 10.14 concludes, under conditions on $\mathbb{A}$, that $\overline{\mathsf{IS}}(\mathbb{A})$ has property MW. In this section, conclusions have the form $\overline{\mathsf{IS}}(\mathbb{A}')$ satisfies property MW, where $\mathbb{A}'$ is a G-set constructed from $\mathbb{A}$ in various ways. Intuitively, we are enlarging our collection of useful sets of atoms.

Notation 10.16 Suppose $p\colon \hat{Y} \to Y$ is an equivariant map of G-sets. For $y \in Y$ [resp. $\hat{y} \in \hat{Y}$], let G_y (resp. $G_{\hat{y}}$) denote the stabilizer of y (resp. $\hat{y}$) with respect to the action of G on Y (resp. $\hat{Y}$). Let F_y denote the fiber $p^{-1}(y)$. If $p(\hat{y}) = y$ then by p-equivariance $G_{\hat{y}} < G_y$ and G_y acts on F_y inducing a homomorphism $\rho_y\colon G_y \to \mathsf{Perm}(F_y)$. (We declare the permutation group of the empty set to be trivial.)

In this setting, we say that p *has explicit finite fibers* if the G-sets Y and $\hat{Y}$ are computable and p comes equipped with an algorithm with input $y \in Y$ and output a finite enumeration of F_y.

Lemma 10.17 *Suppose the* G*-map* $p\colon \hat{Y} \to Y$ *has explicit finite fibers and that* Y *satisfies property M. Then*:

(1) *There is an algorithm with input* $y \in Y$, $\theta \in G_y$ *and output* $\rho_y(\theta)$.

(2) *There is an algorithm with input* $y \in Y$ *and output the multiplication table for* $\rho_y(G_y)$.

(3) $\hat{Y}$ *satisfies property M.*

Proof Let $y \in Y$. By Lemma 8.9, G_y is computable. Since p has explicit finite fibers, we can compute the finite list of elements of F_y. Since Y satisfies property M we can produce a finite presentation for G_y.

(1) Since $\hat{Y}$ is computable, we can compute the action of θ on $F_y \subset \hat{Y}$.

(2) This follows by applying (1) to our generators of G_y.

(3) Since $G_{\hat{y}}$ is the ρ_y-preimage of the stabilizer S of $\hat{y}$ in $\mathsf{Perm}(F_y)$, we can find a finite presentation for $G_{\hat{y}}$ by applying Lemma 9.1, taking P to be the surjective homomorphism $G_y \to \rho_y(G_y)$ and $Q' := \rho_y(G_y) \cap S$. □

Lemma 10.18 *Suppose $f: Z \to Y$ and $g: Y \to X$ each has explicit finite fibers. Then the composition $h = g \circ f: Z \to X$ has explicit finite fibers.*

Proof Since each one of f and g has explicit finite fibers, the G-sets X, Y and Z are each computable. Given $x \in X$, since g has explicit finite fibers, we can list the elements of $g^{-1}(x)$. Since f has explicit finite fibers, we can list the elements of $f^{-1}(y)$ for each $y \in g^{-1}(x)$. We are done by noting that $h^{-1}(x) = \bigsqcup\{f^{-1}(y) \mid y \in g^{-1}(x)\}$. □

Lemma 10.19 *Suppose that $p: \hat{Y} \to Y$ is a G-equivariant map of G-sets such that*

- *p has explicit finite fibers, and*
- *Y satisfies property MW.*

Then $\hat{Y}$ satisfies property MW.

Proof $\hat{Y}$ satisfies property M by Lemma 10.17(3).

For the W-algorithm, we use Notation 10.16. Suppose that $\hat{z} \in \hat{Y}$ and that $z = p(\hat{z})$. Since Y satisfies property W we can check whether or not there is θ_0 such that $\theta_0(y) = z$ and we can compute such a θ_0 if it exists. If not, then return `NO`. If yes, then $\hat{y}$ and $\theta_0^{-1}(\hat{z})$ are in F_y and there is an element of G taking $\hat{y}$ to $\hat{z}$ if and only if there is an element $\theta \in G$ (necessarily in G_y) taking $\hat{y}$ to $\theta_0^{-1}(\hat{z})$. Our goal becomes to check whether there is $\theta \in G_y$ such that $\rho_y(\theta) \in \mathsf{Perm}(F_y)$ takes $\hat{y}$ to $\theta_0^{-1}(\hat{z})$, and to produce such a θ if so. This can be done using our finite generating set for G_y and its action on F_y. If there is no such θ then return `NO`. Otherwise return `YES` and $\theta_0 \cdot \theta$. □

Construction 10.20 (canonical extension) Suppose that $\mathbb{A}$ and $\mathbb{A}'$ are G-sets and that $I: \mathbb{A}' \to \overline{\mathsf{IS}}(\mathbb{A})$ is G-equivariant. We now define a G-equivariant map $I_{\overline{\mathsf{IS}}}: \overline{\mathsf{IS}}(\mathbb{A}') \to \overline{\mathsf{IS}}(\mathbb{A})$ that restricts to I on $\mathbb{A}'$; see Notation 10.13. We call $I_{\overline{\mathsf{IS}}}$ the *canonical extension of I*:

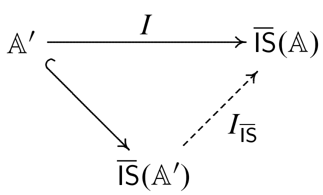

From $\mathsf{Q}' \in \overline{\mathsf{IS}}(\mathbb{A}')$, we construct $\mathsf{Q} := I_{\overline{\mathsf{IS}}}(\mathsf{Q}') \in \overline{\mathsf{IS}}(\mathbb{A})$. We do this by constructing a representative $\mathsf{X} = ((T, *), \chi)$ for Q from a representative $\mathsf{X}' = ((T', *'), \chi')$ for Q' and, for each $l' \in \mathcal{L}(T')$, a representative $\mathsf{X}_{l'} = ((T_{l'}, *_{l'}), \chi_{l'})$ for $\mathsf{Q}_{l'} := I(\chi'(l')) \in \overline{\mathsf{IS}}(\mathbb{A})$. Let T be the tree obtained from

$T' \sqcup \big(\bigsqcup\{T_{l'} \mid l' \in \mathcal{L}(T')\}\big)$ by identifying $l' \in T'$ and $*_{l'} \in T_{l'}$. We declare the image of $*'$ in T to be the root $*$ of T. The leaves of T biject naturally with $\bigsqcup\{\mathcal{L}(T_{l'}) \mid l' \in \mathcal{L}(T')\}$ and we define $\chi\colon \mathcal{L}(T) \to \mathbb{A}$ by $\chi|\mathcal{L}(T_{l'}) := \chi_{l'}$.

We next show that Q is independent of our choices of representatives, ie that $\mathsf{Q}' \mapsto \mathsf{Q}$ is well-defined. Let $\mathsf{Y}' = ((S', \star'), \eta')$ also represent Q' and so we have a simplicial isomorphism $f'\colon T' \to S'$ inducing an equivalence $\mathsf{X}' \to \mathsf{Y}'$. In particular, $I(\chi'(l')) = I(\eta'(f(l')))$ in $\overline{\mathsf{IS}}(\mathbb{A})$. Thus if $\mathsf{Y}_{f(l')}$ is a representative of $I(\eta'(f(l')))$, then we have equivalences $\mathsf{X}_{l'} \to \mathsf{Y}_{f(l')}$ induced by simplicial isomorphisms $f_{l'}\colon (T_{l'}, *_{l'}) \to (S_{f(l')}, \star_{f(l')})$ between the underlying trees. The map f' and the $f_{l'}$ induce a map

$$T' \sqcup \Big(\bigsqcup\{T_{l'} \mid l' \in \mathcal{L}(T')\}\Big) \to S' \sqcup \Big(\bigsqcup\{S_{f_{l'}(l')} \mid l' \in \mathcal{L}(T')\}\Big),$$

which descends to a simplicial isomorphism $T \to S$ that induces an equivalence $\mathsf{X} \to \mathsf{Y}$. Hence the map $I_{\overline{\mathsf{IS}}}\colon \overline{\mathsf{IS}}(\mathbb{A}') \to \overline{\mathsf{IS}}(\mathbb{A})$ given by $\mathsf{Q}' \mapsto \mathsf{Q}$ is well-defined.

If we start the above construction with $\theta(\mathsf{X}')$ instead of X', the only difference is that χ is replaced by $\theta \circ \chi$, ie $I_{\overline{\mathsf{IS}}}$ is G-equivariant. Recall (Notation 10.13) that we identify $\mathbb{A}'$ with the subset of elements of $\overline{\mathsf{IS}}(\mathbb{A}')$ with underlying tree consisting of only the root. Thus I and $I_{\overline{\mathsf{IS}}}$ agree on $\mathbb{A}'$.

Lemma 10.21 *Let* $\mathbb{A}$ *and* $\mathbb{A}'$ *be* G*-sets and suppose the* G*-map* $I\colon \mathbb{A}' \to \overline{\mathsf{IS}}(\mathbb{A})$ *has explicit finite fibers. Then* $I_{\overline{\mathsf{IS}}}\colon \overline{\mathsf{IS}}(\mathbb{A}') \to \overline{\mathsf{IS}}(\mathbb{A})$ *has explicit finite fibers. If additionally* $\mathcal{S}_{\mathrm{or}}(\mathbb{A})$ *satisfies property MW, then* $\overline{\mathsf{IS}}(\mathbb{A}')$ *satisfies property MW.*

Proof Since I has explicit finite fibers, the G-set $\mathbb{A}'$ is computable. The G-set $\overline{\mathsf{IS}}(\mathbb{A}')$ is therefore computable by Lemma 10.10(3). Let $\mathsf{Q} \in \overline{\mathsf{IS}}(\mathbb{A})$ be given and let $\mathsf{X} = ((T, *), \chi) \in \mathsf{IS}(\mathbb{A})$ represent Q. Using notation as in Construction 10.20, each Q' in the fiber F of $I_{\overline{\mathsf{IS}}}$ over Q has a representative of the form $((T', *), \chi')$, where $(T', *)$ is a rooted subtree of $(T, *)$. Further, each leaf l' of T' then determines a rooted tree $(T_{l'}, l')$, where $T_{l'}$ is the subtree of T spanned by l' and all leaves l of T with a directed path from l' to l. We then get a representative $((T_{l'}, l'), \chi_{l'})$ for an element $\mathsf{Q}_{l'}$ of $\overline{\mathsf{IS}}(\mathbb{A})$, where $\chi_{l'}$ is the restriction of χ to the leaves of $T_{l'}$. If there are elements $a'_{l'} \in \mathbb{A}'$ such that $I(a'_{l'}) = \mathsf{Q}_{l'}$ and if we define $\chi'(l') := a'_{l'}$ then $((T', *), \chi')$ represents an element of F and all elements of F have this form (for some choice of T'). Since I has explicit finite fibers, we can finitely enumerate the fiber of I over $\mathsf{Q}_{l'}$ and so also find a finite list in $\mathsf{IS}(\mathbb{A}')$ of representatives for the elements of F. (It is easy to see that a different choice of representative X for Q produces F with a perhaps different enumeration.)

If additionally $\mathcal{S}_{\mathrm{or}}(\mathbb{A})$ satisfies property MW, then by Proposition 10.14 so does $\overline{\mathsf{IS}}(\mathbb{A})$. That $\overline{\mathsf{IS}}(\mathbb{A}')$ satisfies property MW is now a direct consequence of Lemma 10.19. □

Corollary 10.22 *Let* $\mathbb{A}$ *and* $\mathbb{A}_i$ *for* $i = 1, \dots, k$ *be* G*-sets with* $\mathbb{A}$ *computable. Suppose that* $\mathcal{S}_{\mathrm{or}}(\mathbb{A})$ *satisfies property MW and that* $I_i\colon \mathbb{A}_i \to \overline{\mathsf{IS}}(\mathbb{A})$ *is* G*-equivariant and has explicit finite fibers. Then the induced map* $\bigsqcup_i \mathbb{A}_i \to \overline{\mathsf{IS}}(\mathbb{A})$ *has explicit finite fibers, as does* $\overline{\mathsf{IS}}\big(\bigsqcup_i \mathbb{A}_i\big) \to \overline{\mathsf{IS}}(\mathbb{A})$*, and* $\overline{\mathsf{IS}}\big(\bigsqcup_i \mathbb{A}_i\big)$ *satisfies property MW.*

Proof It is apparent that, since I_i has explicit finite fibers, so does $\bigsqcup_i \mathbb{A}'_i \to \overline{\mathsf{IS}}(\mathbb{A})$. The rest of the corollary then follows directly from Lemma 10.21. $\square$

Corollary 10.23 *Suppose $\mathbb{A}'$ is a computable G-set and $\mathcal{S}_{\mathrm{or}}(\mathbb{A}')$ satisfies property MW. Then:*

(1) *For $k = 2, 3, \dots$, inductively define $\overline{\mathsf{IS}}_k(\mathbb{A}') := \overline{\mathsf{IS}}(\overline{\mathsf{IS}}_{k-1}(\mathbb{A}'))$, where $\overline{\mathsf{IS}}_1(\mathbb{A}') := \overline{\mathsf{IS}}(\mathbb{A}')$. The G-set $\overline{\mathsf{IS}}_k(\mathbb{A}')$ satisfies property MW.*

(2) *Here we use Notation 10.13. For $i = 1, 2, \dots$, let $s_i \in \{\mathcal{S}_{\mathrm{or}}, \mathcal{S}_{\mathrm{un}}\}$. Set $S_1(\mathbb{A}') := s_1(\mathbb{A}') \subset \overline{\mathsf{IS}}(\mathbb{A}')$ and inductively define $S_k(\mathbb{A}') := s_k(S_{k-1}(\mathbb{A}')) \subset \overline{\mathsf{IS}}_k(\mathbb{A}')$. The G-set $S_k(\mathbb{A}')$ satisfies property MW. The natural map $S_k(\mathbb{A}') \subset \overline{\mathsf{IS}}_k(\mathbb{A}') \to \overline{\mathsf{IS}}_{k-1}(\mathbb{A}') \to \dots \to \overline{\mathsf{IS}}(\mathbb{A}')$ is injective, where $\overline{\mathsf{IS}}_i(\mathbb{A}') \to \overline{\mathsf{IS}}_{i-1}(\mathbb{A}')$ is the canonical extension (Construction 10.20) of the identity map of $\overline{\mathsf{IS}}_{i-1}(\mathbb{A}')$.*

Proof (1) An application of Proposition 10.14 gives that $\overline{\mathsf{IS}}(\mathbb{A}')$ satisfies property MW. Suppose that, for $k \geq 2$, we have that $\overline{\mathsf{IS}}_{k-1}(\mathbb{A}')$ satisfies property MW. Since the identity map of $\overline{\mathsf{IS}}_{k-1}(\mathbb{A}')$ has explicit finite fibers, so does $\overline{\mathsf{IS}}_k(\mathbb{A}') \to \overline{\mathsf{IS}}_{k-1}(\mathbb{A}')$ by Lemma 10.21. By Lemma 10.19, $\overline{\mathsf{IS}}_k(\mathbb{A}')$ also satisfies property MW.

(2) The first conclusion follows from (1) since $S_k(\mathbb{A}') \subset \overline{\mathsf{IS}}_k(\mathbb{A}')$, and an MW-algorithm for $\overline{\mathsf{IS}}_k(\mathbb{A}')$ provides an MW-algorithm for $S_k(\mathbb{A}')$. To prove the injectivity statement by induction, we show that if $\mathbb{A}''$ is a G-set and $f : \mathbb{A}'' \to \overline{\mathsf{IS}}(\mathbb{A}')$ is injective, then the restriction of the canonical extension $s_i(\mathbb{A}'') \to \overline{\mathsf{IS}}(\mathbb{A}')$ is also injective for $s_i \in \{\mathcal{S}_{\mathrm{or}}, \mathcal{S}_{\mathrm{un}}\}$.

Suppose that $s_i = \mathcal{S}_{\mathrm{or}}$ and that $(a_1, \dots, a_r), (b_1, \dots, b_s) \in s_i(\mathbb{A}'')$. If the images $(f(a_1), \dots, f(a_r))$ and $(f(b_1), \dots f(b_s))$ under the canonical map are equal, then $f(a_k) = f(b_k)$, $k = 1, \dots, m = n$. Since f is injective, $a_k = b_k$. The case where $s_i = \mathcal{S}_{\mathrm{un}}$ is similar. The inductive proof is left to the reader. $\square$

Remark 10.24 Via the natural map in Corollary 10.23(2), we sometimes view $S_k(\mathbb{A}')$ as a subset of $\overline{\mathsf{IS}}(\mathbb{A}')$.

Example 10.25 Let $\mathbb{A} = \mathcal{C}(F_n)$. A free factor system is a finite unordered set of free factors and so may be interpreted as an element of $\mathcal{S}_{\mathrm{un}}(\mathbb{A})$; cf Example 10.6. A special chain is an ordered set of free factor systems and so gives an element of $\mathcal{S}_{\mathrm{or}}\mathcal{S}_{\mathrm{un}}(\mathbb{A})$. The set of special chains associated to ϕ gives an element of $\mathcal{S}_{\mathrm{un}}\mathcal{S}_{\mathrm{or}}\mathcal{S}_{\mathrm{un}}(\mathbb{A})$. We may view $\mathcal{S}_{\mathrm{un}}(\mathbb{A})$, $\mathcal{S}_{\mathrm{or}}\mathcal{S}_{\mathrm{un}}(\mathbb{A})$, and $\mathcal{S}_{\mathrm{un}}\mathcal{S}_{\mathrm{or}}\mathcal{S}_{\mathrm{un}}(\mathbb{A})$ all as subsets of $\overline{\mathsf{IS}}(\mathbb{A})$ (Remark 10.24).

11 Our atoms

In this section we apply Section 10.3 to enlarge the set $\mathcal{C}(F_n)$ (Notation 7.20) to the sets of atoms that we will need for the remainder of the paper. See Lemma 11.2 and also Section 14.

To start, we need to know that some familiar sets are computable. Some proofs refer to the Stallings graph associated to a finitely generated subgroup of F_n; see Stallings [1983, 5.4] and Section 7.2. The proofs also use Lemmas 8.3 and 8.5, sometimes without explicit mention. Let $\mathcal{B}$ be a basis for F_n. The following objects are computable.

- **The group F_n** Elements of F_n are represented by words in $\mathcal{B} \sqcup \mathcal{B}^{-1}$ and can be enumerated using length. Multiplication is given by concatenation. An element is represented uniquely by a reduced word, which can be found by iteratively canceling a letter and its inverse. See Example 8.4.
- **The set of conjugacy classes of elements of F_n** The enumeration of F_n serves also as an enumeration of the set of conjugacy classes. An element is represented uniquely by a cyclically reduced word.
- **The set of finitely generated subgroups of F_n** An enumeration is given by an enumeration of finite sets of representatives of elements of F_n. Two such sets are equal if and only if they determine the same based Stallings graph with labels in $\mathcal{B}$, which can be found using the Stallings folding algorithm.
- **$\mathcal{C}(F_n)$** The enumeration of the set of finitely generated subgroups of F_n also serves as an iteration for $\mathcal{C}(F_n)$. Two representatives are equal in $\mathcal{C}(F_n)$ if and only if they determine the same (unbased) Stallings graph with labels in $\mathcal{B}$; see Lemma 7.10.
- **The group $\mathsf{Aut}(F_n)$** The set of endomorphisms of F_n is bijective with $(F_n)^{\mathcal{B}}$ under the map $\Theta \mapsto (b \mapsto \Theta(b))$. Using the Hopf property of F_n [Magnus et al. 1966, Theorem 2.13], an endomorphism Θ is an isomorphism if and only if it is surjective if and only if the based Stallings graph of $\langle \Theta(\mathcal{B}) \rangle$ is the rose with petals labeled by the elements of $\mathcal{B}$. Thus an enumeration of $\mathsf{Aut}(F_n)$ can be obtained using the enumeration for endomorphisms.
- **The group $\mathsf{Out}(F_n)$** Our enumeration of automorphisms serves also as an enumeration of outer automorphisms. Two automorphisms represent the same outer automorphisms if and only if they have the same action on conjugacy classes of words of length at most two in $\mathcal{B} \sqcup \mathcal{B}^{-1}$; see [Serre 1980].
- **The $\mathsf{Aut}(F_n)$-set F_n** Representative endomorphisms act on representative words.
- **The $\mathsf{Out}(F_n)$-set of conjugacy classes of elements of F_n** Representative endomorphisms act on representative words.
- **The $\mathsf{Aut}(F_n)$-set of finitely generated subgroups of F_n** Representative endomorphisms act on representative finite subsets of F_n.
- **The $\mathsf{Out}(F_n)$-set $\mathcal{C}(F_n)$** Representative endomorphisms act on representative finite subsets of F_n.
- **The $\mathsf{Out}(F_n)$-set $\overline{\mathsf{IS}}(\mathcal{C}(F_n))$** See Lemma 10.10(3).
- **The $\mathsf{Out}(F_n)$-sets $\mathcal{S}_{\mathrm{or}}(\mathcal{C}(F_n))$ and $\mathcal{S}_{\mathrm{un}}(\mathcal{C}(F_n))$** See Corollary 10.23.

Notation 11.1 In the remainder of the paper, $G = \mathsf{Out}(F_n)$. We now define $\mathsf{Out}(F_n)$-sets $\mathbb{A}_0, \mathbb{A}_1, \dots, \mathbb{A}_6$ that will be used to express our algebraic invariants for elements of $\mathsf{UPG}(F_n)$. At the same time we show each $\mathbb{A}_i$ is computable and admits a map that has explicit finite fibers to a previously defined set.

- $\mathbb{A}_0$ denotes $\mathcal{C}(F_n)$, is computable (see the above itemized list), and has been identified (Notation 10.13) as an $\mathsf{Out}(F_n)$-subset of $\overline{\mathsf{IS}}(\mathbb{A}_0)$.

- $\mathbb{A}_1$ denotes the set of good conjugacy pairs of nontrivial finitely generated subgroups of F_n. We saw above that the set of finitely generated subgroups of F_n is computable, hence (Lemma 8.5(2)) so is its square. By Lemma 7.14(1), a pair (K_1, K_2) represents a good $[K_1, K_2]$ if and only if $\mathrm{rank}(\langle K_1, K_2\rangle) = \mathrm{rank}(K_1) + \mathrm{rank}(K_2)$, and this can be checked using the Stallings graphs of $[K_1]$, $[K_2]$ and $[\langle K_1, K_2\rangle]$. Hence (Lemma 8.5(1)) the subset of pairs representing good conjugacy pairs is computable. By Corollary 7.18, there is an algorithm deciding whether good conjugacy pairs are equal. Hence, by Lemma 8.3, $\mathbb{A}_1$ is computable.

By Corollary 7.23, the map $\mathbb{A}_1 \to \mathcal{S}_{\mathrm{or}}(\mathbb{A}_0) \hookrightarrow \overline{\mathsf{IS}}(\mathbb{A}_0)$ has explicit finite fibers where $\mathbb{A}_1 \to \mathcal{S}_{\mathrm{or}}(\mathbb{A}_0)$ is given by $[H_1, H_2] \mapsto ([H_1], [H_2], [\langle H_1, H_2\rangle])$.

- $\mathbb{A}_2$ denotes the set of conjugacy pairs $[H, a]$ where H is a nontrivial finitely generated subgroup of F_n, $a \in F_n$ is nontrivial, and $[H, \langle a\rangle]$ is good. It is clear that $[H, a] = [H', a']$ if and only if $[H, \langle a\rangle] = [H', \langle a'\rangle]$ and a and a' are conjugate. In particular, $\mathbb{A}_2$ is computable and the map $\mathbb{A}_2 \to \mathbb{A}_1$ given by $[H, a] \mapsto [H, \langle a\rangle]$ has explicit finite fibers with fibers of size zero or two.

- $\mathbb{A}_3$ denotes the set of conjugacy pairs $[a, H]$ where H is a nontrivial finitely generated subgroup of F_n, $a \in F_n$ is nontrivial, and $[H, \langle a\rangle]$ is good. That $\mathbb{A}_3$ is computable and $\mathbb{A}_3 \to \mathbb{A}_1$ given by $[a, H] \mapsto [\langle a\rangle, H]$ has explicit finite fibers with fibers of size zero or two follows exactly as with $\mathbb{A}_2$.

- $\mathbb{A}_4$ is the set of conjugacy pairs $[a, b]$ of elements of F_n where $\langle a, b\rangle$ has rank 2. We have $[a, b] = [a', b']$ if and only if $[\langle a\rangle, \langle b\rangle] = [\langle a'\rangle, \langle b'\rangle]$, a and a' are conjugate, and b and b' are conjugate. It follows that $\mathbb{A}_4$ is computable and $\mathbb{A}_4 \to \mathbb{A}_1$ given by $[a, b] \mapsto [\langle a\rangle, \langle b\rangle]$ has explicit finite fibers with fibers of size zero or four.

- $\mathbb{A}_5$ is the set of conjugacy classes $[a]$ of nontrivial elements $a \in F_n$. We saw earlier in this subsection that $\mathbb{A}_5$ is computable. The map $\mathbb{A}_5 \to \mathbb{A}_0$ given by $[a] \mapsto [\langle a\rangle]$ has explicit finite fibers with fibers of size zero or two.

- $\mathbb{A}'_6$ is the set of conjugacy pairs $[H, A]$ with $A < H < F_n$ all finitely generated and nontrivial. (In particular, $[H, A]$ is not good.) Using Lemma 7.24(1), the proof that $\mathbb{A}'_6$ is computable is similar to the proof that $\mathbb{A}_1$ is computable. By Lemma 7.24(2), $\mathbb{A}'_6 \to \mathcal{S}_{\mathrm{or}}(\mathbb{A}_0) \hookrightarrow \overline{\mathsf{IS}}(\mathbb{A}_0)$ given by $[H, A] \mapsto ([H], [A])$ has explicit finite fibers. $\mathbb{A}'_6$ is only used to define $\mathbb{A}_6$.

- $\mathbb{A}_6$ is the set of conjugacy pairs $[H, a]$ where H is a nontrivial finitely generated subgroup of F_n and $a \neq 1$ is in H. (In particular, $[H, \langle a\rangle]$ is not good.) $[H, a] = [H', a']$ if and only if $[H, \langle a\rangle] = [H', \langle a'\rangle]$ and a is conjugate to a'. Hence $\mathbb{A}_6 \to \mathbb{A}'_6$ has explicit finite fibers with fibers of size zero or two.

- $\mathbb{A}_\bullet := \mathbb{A}_0 \sqcup \mathbb{A}_1 \sqcup \mathbb{A}_2 \sqcup \mathbb{A}_3 \sqcup \mathbb{A}_4 \sqcup \mathbb{A}_5 \sqcup \mathbb{A}_6$.

Lemma 11.2 *Using Notation 11.1, the* $\mathsf{Out}(F_n)$*-set* $\overline{\mathsf{IS}}(\mathbb{A}_\bullet)$ *satisfies property MW.*

Proof Since $\mathcal{S}_{\mathrm{or}}(\mathbb{A}_0)$ satisfies property MW by Theorem 10.2, it follows from Corollary 10.22 that it is enough to show that, for each i, we have that $\mathbb{A}_i$ admits G-equivariant map to $\overline{\mathsf{IS}}(\mathbb{A}_0)$ that has explicit finite fibers. Using Notation 11.1, we see that each $\mathbb{A}_i$ admits a map to $\overline{\mathsf{IS}}(\mathbb{A}_0)$ that is a composition of two maps, each of which has explicit finite fibers. We are done by Lemma 10.18. □

12 List of dynamical invariants

In Section 13 we define algebraic invariants of $\phi \in \mathsf{UPG}(F_n)$ that are derived from the dynamical invariants of ϕ established in the first five sections of this paper. For the convenience of the reader, we list those dynamical invariants here and provide pointers to the relevant sections of the paper. Here $f : G \to G$ always denotes a CT for ϕ and $\Gamma(f)$ its eigengraph; see Section 4.1. We also use the notation of conjugacy pairs; recall Definition 4.9, Examples 4.10 and Section 7.

- $\mathcal{P}(\phi)$ denotes the set of principal automorphisms for ϕ (Definition 3.8) and $[\mathcal{P}(\phi)]$ denotes the set of isogredience classes in $\mathcal{P}(\phi)$ (Definition 3.11). $[\mathcal{P}(\phi)]$ parametrizes the components of $\Gamma(f)$. $\mathsf{Fix}(\phi) = \{[\mathsf{Fix}(\Phi)] \mid [\Phi] \in [\mathcal{P}(\phi)]\}$. Since $[\mathcal{P}(\phi)]$ is finite, $\mathsf{Fix}(\phi)$ is a finite multiset of (possibly trivial) conjugacy classes of finitely generated subgroups of F_n. Geometrically it is the core of $\Gamma(f)$. $\mathsf{Fix}_{\geq 2}(\phi) := \{[\mathsf{Fix}(\Phi)] \mid [\Phi] \in [\mathcal{P}(\phi)], \mathrm{rank}(\mathsf{Fix}(\Phi)) \geq 2\}$. See Sections 3.4, 3.7 and 4.2.

- We use $\mathfrak{c} = \vec{\mathcal{F}}(\phi, <_T)$ to denote a special chain for ϕ as in Notation 6.8. It is a set of free factor systems naturally ordered by $\sqsubset$. We usually work with a prechosen $\mathfrak{c}$. For example, the filtration of our CT $f : G \to G$ will usually realize $\mathfrak{c}$. If $\mathcal{F} \in \mathfrak{c}$ (resp. $[F] \in \mathcal{F} \in \mathfrak{c}$) and if the filtration of $f : G \to G$ realizes $\mathfrak{c}$ then $f|\mathcal{F}$ (resp. $f|[F]$) denotes the restriction of f to the core filtration element representing $\mathcal{F}$ (resp. the component of the core filtration element representing $[F]$). The corresponding eigengraph is denoted by $\Gamma(f|\mathcal{F})$ (resp. $\Gamma(f|[F])$).

- A free factor system is special if it is in some special chain. $\mathfrak{L}(\phi)$ denotes the set of special free factor systems of ϕ; see Notation 6.8. Each element of $\mathfrak{L}(\phi)$ is a free factor system and so is a set of conjugacy classes of free factors in F_n. If $[F] \in \mathcal{F} \in \mathfrak{L}(\phi)$ then F and $[F]$ are also said to be special. The unique minimal (with respect to $\sqsubset$) element of $\mathfrak{L}(\phi)$, denoted by $\mathcal{F}_0(\phi)$, is the linear free factor system of ϕ. It is represented by the core of the subgraph of G that is the union of fixed and linear edges. An invariant description of $\mathcal{F}_0(\phi)$ is $\mathcal{F}(\mathsf{Fix}(\phi))$, ie the smallest free factor system carrying $\mathsf{Fix}(\phi)$; see Lemma 6.16.

- Define
$$\mathcal{R}(\phi) := \left\{ [P] \;\middle|\; P \in \bigcup_{i=1}^{m} \mathsf{Fix}_+(\Phi_i) \right\} \subset \partial F_n / F_n,$$
where the Φ_i are representatives of the isogredience classes in $\mathcal{P}(\phi)$. In other words, $\mathcal{R}(\phi)$ is the set of conjugacy classes of points in ∂F_n that are isolated fixed points for some principal lift of ϕ. See Section 3.4. In any CT $f : G \to G$ representing ϕ there is a bijection $r \leftrightarrow E$ between $\mathcal{R}(\phi)$ and the set $\mathcal{E}_f$ of higher-order edges of G. The eigenray R_E has terminal end r.

- Let $\mathfrak{e} \in \mathfrak{c}$ denote a special 1-edge extension in $\mathfrak{c}$, ie $\mathfrak{e} = (\mathcal{F}^- \sqsubset \mathcal{F}^+)$ is a pair of consecutive elements of $\mathfrak{c}$. Suppose $f : G \to G$ realizes $\mathfrak{c}$. $\Gamma(f|\mathcal{F}^+) \setminus \Gamma(f|\mathcal{F}^-)$ has one or two ends; these represent the *new* (*with respect to* $\mathfrak{e}$) elements of $\mathcal{R}(\phi)$. The 1-edge extension $\mathfrak{e}$ has type H, HH or LH. There are two new elements if and only if $\mathfrak{e}$ has type HH. A new element is often denoted by r^+. Further, $\mathfrak{e}$ can be contractible, infinite cyclic, or large. See Section 6.1.

- Continuing the previous bullet point, if the filtration of $f : G \to G$ realizes $\mathfrak{c}$, then $\Gamma(f|\mathcal{F}^+)$ carries more lines than $\Gamma(f|\mathcal{F}^-)$. The set of *added lines with respect to* $\mathfrak{e}$, denoted by $\mathsf{L}_{\mathfrak{e}}(\phi)$, is a ϕ-invariant subset of these lines. See Definition 6.14.

- $\Omega(\phi) = \bigcup_{r \in \mathcal{R}(\phi)} \Omega(r)$ denotes the finite set of limit lines for ϕ. See Section 5. Here $\Omega(r)$ denotes the accumulation set of r or equivalently of the eigenray in $\Gamma(f)$ representing r. The elements of $\Omega(\phi)$ are all represented as lines in $\Gamma(f)$. $\Omega_{\mathsf{NP}}(\phi) \subset \Omega(\phi)$ is the subset of nonperiodic lines.

- $\mathcal{A}_{\mathrm{or}}(\phi)$ denotes the set of oriented axes of ϕ, where a root-free conjugacy class $[a]$ of an element of F_n is an axis if it has more than one representation in $\Gamma(f)$. An axis has an invariant description: $[a]$ is an axis if there are $\Phi_1, \Phi_2 \in \mathcal{P}(\phi)$ such that $a \in \mathsf{Fix}(\Phi_1) \cap \mathsf{Fix}(\Phi_2)$ and $\Phi_1 \neq \Phi_2$. In this case, we say the conjugacy pair $[\Phi, a]$ is a strong axis; see Definition 4.11. It is represented geometrically as a lift to $\Gamma(f)$ of $[a]$. The set of strong axes is denoted by $\mathsf{SA}(\phi)$. Associated to each pair of strong axes $\alpha_1 = [\Phi_1, a], \alpha_2 = [\Phi_2, a]$ is a twist coordinate $\tau(\alpha_1, \alpha_2)$ in $\mathbb{Z}$. See Definition 4.14.

13 Algebraic data associated to invariants

In this section we define algebraic versions of some of our dynamical invariants. We also explain how the algebraic versions can be computed and viewed as an element of $\overline{\mathsf{IS}}(\mathbb{A}_\bullet)$; see Notation 11.1. The algebraic invariants are typically weaker than their dynamic versions. However they have the advantage that they are iterated sets and so fit into the framework of Section 10. Some of our invariants, for example chains, are already algebraic in nature and so need no modification.

All of our algebraic invariants for $\phi \in \mathsf{UPG}$ will be computed using a CT $f : G \to G$ for ϕ; see Section 3.6. Additionally, the core of the eigengraph $\Gamma(f)$ can be computed from $f : G \to G$; see Section 4.1. In fact, since $\Gamma(f)$ is obtained from its core by adding the eigenrays of f and the eigenrays have a simple form (Section 3.6), we can compute arbitrarily large neighborhoods of the core in $\Gamma(f)$.

13.1 Special chains

Recall from Notation 6.1 and Lemma 6.2 that there is a canonical partial order $(\mathcal{R}(\phi), <)$ that can be computed from any CT for ϕ. Hence all extensions of $<$ to a total order $<_T$ can also be computed. The special chain $\vec{\mathcal{F}}(\phi, <_T)$ for ϕ can also be computed from any CT for ϕ; see Notation 6.8. A special chain is an element of $\mathcal{S}_{\mathrm{or}}\mathcal{S}_{\mathrm{un}}(\mathbb{A}_0) \subset \overline{\mathsf{IS}}(\mathbb{A}_0) \subset \overline{\mathsf{IS}}(\mathbb{A}_\bullet)$; see Example 10.25 and Corollary 10.23(2). Similarly, the set of all special chains for ϕ and the set $\mathfrak{L}(\phi)$ of all special free factor systems for ϕ can be computed from any CT for ϕ. Note that the former set is in $\mathcal{S}_{\text{un}}\mathcal{S}_{\text{or}}\mathcal{S}_{\text{un}}(\mathbb{A}_0) \subset \overline{\mathsf{IS}}(\mathbb{A}_\bullet)$ and $\mathfrak{L}(\phi) \in \mathcal{S}_{\text{un}}\mathcal{S}_{\text{un}}(\mathbb{A}_0) \subset \overline{\mathsf{IS}}(\mathbb{A}_\bullet)$. We will tacitly use Corollary 10.23(2) throughout the rest of Section 13.

- Throughout the rest of Section 13, $\phi \in \mathsf{UPG}$, $\mathfrak{c}$ denotes a special chain for ϕ, and $f : G \to G$ denotes a CT that represents ϕ, satisfies (Inheritance), and realizes $\mathfrak{c}$.

13.2 Fix(ϕ)

The multiset $\mathsf{Fix}(\phi) := \{[\mathsf{Fix}(\Phi)] \mid [\Phi] \in [\mathcal{P}(\phi)]\}$ is already algebraic and is an element of $\mathcal{S}_{\text{un}}(\mathbb{A}_0) \subset \overline{\mathsf{IS}}(\mathbb{A}_\bullet)$. As reviewed in Section 12, $\mathsf{Fix}(\phi)$ is represented by the core of $\Gamma(f)$ and so can be computed.

13.3 Axes

The set $\mathcal{A}_{\text{or}}(\phi)$ of oriented axes of ϕ (Definition 4.5) is already algebraic and is an element of $\mathcal{S}_{\text{un}}(\mathbb{A}_5) \subset \overline{\mathsf{IS}}(\mathbb{A}_\bullet)$. In terms of $f : G \to G$, $[a] \in \mathcal{A}_{\text{or}}(\phi)$ if and only if either $[a]$ or $[a^{-1}]$ is represented by a twist path, which can be found by inspecting the linear edges of G; see Section 3.6. In terms of $\Gamma(f)$, $[a] \in \mathcal{A}_{\text{or}}(\phi)$ if and only if a is root-free and represented by more than one circuit in the core of $\Gamma(f)$ with at least one representative embedded.

13.4 Algebraic rays

Remark 13.1 If F is a free factor and $\tilde{r} \in \partial F$ then we say that $[F]$ *carries* r. Equivalently, if G is a marked graph and $H \subset G$ is a core subgraph representing $[F]$, then there is a ray in $\widetilde{G}$ that converges to $\tilde{r}$ and projects into H. In the case that concerns us, $r \in \mathcal{R}(\phi)$ corresponds to some $E \in \mathcal{E}_f$ (see Lemma 3.26) and $[F]$ is a component of a free factor system in $\mathfrak{c}$. If C is the component of the core filtration element of G corresponding to $[F]$ then $[F]$ carries r if and only if some subray of R_E is contained in C. By construction, $R_E = E \cdot u \cdot f_\#(u) \cdot \ldots$, where the closed path u satisfies $f(E) = E \cdot u$. The height of $f_\#^k(u)$ is independent of k, so r is carried by $[F]$ if and only if $u \subset C$.

- (algebraic rays) For $r \in \mathcal{R}(\phi)$, $F_{\mathfrak{c}}(r)$ denotes the minimal special free factor $[F] \in \mathcal{F} \in \mathfrak{c}$ carrying r. If $\tilde{r} \in \partial F_n$ is a lift of r then we also write $F_{\mathfrak{c}}(\tilde{r}) := F$, where F is the unique representative of $F_{\mathfrak{c}}(r)$ that contains $\tilde{r}$. An algebraic ray $F_{\mathfrak{c}}(r)$ is an element of $\mathbb{A}_0 \subset \overline{\mathsf{IS}}(\mathbb{A}_\bullet)$.

Remark 13.2 Continuing Remark 13.1, $F_{\mathfrak{c}}(r)$ is represented by the minimal component C of a core filtration element of G containing u. In particular, we can compute $F_{\mathfrak{c}}(r)$ from our CT f. In our running example (see pages 1700, 1707, 1725 and 1729), $\mathcal{R}(\phi) = \{r_c, r_d, r_e, r_q\}$, the only relation is $r_c < r_q$, and the choice of total order is $r_c <_T r_d <_T r_e <_T r_q$. We have $[F_{\mathfrak{c}}(r_q)] = [\langle a, b, c\rangle]$.

Remark 13.3 We could work with all chains and define $F(r)$ to be the minimal special free factor $[F] \in \mathcal{F} \in \mathfrak{L}(\phi)$ carrying r. This would cause some extra work later in Lemma 17.19.

13.5 Algebraic lines

Recall that $[\cdot,\cdot]$ denotes a conjugacy pair (see Definition 4.9, Examples 4.10 and Section 7) and, for nontrivial $a \in F_n$, $a^+ \in \partial F_n$ (resp. a^-) denotes the attractor (resp. repeller) of $i_a|\partial F_n$; see the beginning of Section 3.4. Recall also that a line L is principal with respect to ϕ if there is a lift $\widetilde{L}$ whose endpoints are contained in $\mathsf{Fix}_{\mathsf{N}}(\Phi)$ for some $\Phi \in \mathcal{P}(\phi)$. Equivalently, L lifts into the eigengraph $\Gamma(f)$; see Lemma 4.1. In this section, we define an algebraic version $\mathsf{H}_{\phi,\mathfrak{c}}(L)$ for a certain principal line L and associate to $\mathsf{H}_{\phi,\mathfrak{c}}(L)$ a set of lines containing L that in turn determines $\mathsf{H}_{\phi,\mathfrak{c}}(L)$.

Definition 13.4 (algebraic lines) Suppose that L is a nonperiodic principal line for ϕ and that the nonperiodic ends of L are contained in $\mathcal{R}(\phi)$. There are four possibilities.

[P-P] L has type P-P if some (hence every) lift $\widetilde{L}$ has the form (a^-, b^+) for some root-free $a, b \in F_n$ with $a \neq b^{\pm 1}$ in F_n. In particular, a and b are nontrivial. $\mathsf{H}_{\phi,\mathfrak{c}}(L) := [a,b]$. To $[a,b]$ we associate $\{L\}$. We also define $\mathsf{H}_{\phi,\mathfrak{c}}(\widetilde{L}) := (a,b)$ and associate to it $\{(a^-, b^+)\}$. In this case $\mathsf{H}_{\phi,\mathfrak{c}}(L)$ determines L.

[P-NP] L has type P-NP if some (hence every) lift $\widetilde{L}$ has the form $(a^-, \widetilde{r})$ for some root-free $a \in F_n$ and a lift $\widetilde{r}$ of some $r \in \mathcal{R}(\phi)$. $\mathsf{H}_{\phi,\mathfrak{c}}(L) := [a, F_{\mathfrak{c}}(\widetilde{r})]$. To $\mathsf{H}_{\phi,\mathfrak{c}}(L)$ we associate the set of lines $[a^-, \partial F_{\mathfrak{c}}(\widetilde{r})]$. $\mathsf{H}_{\phi,\mathfrak{c}}(\widetilde{L}) := (a, F_{\mathfrak{c}}(\widetilde{r}))$ and has the associated set of lines $(a^-, \partial F_{\mathfrak{c}}(\widetilde{r}))$.

[NP-P] L has type NP-P if some (hence every) lift $\widetilde{L}$ has the form $(\widetilde{r}, b^+)$ for some root-free $b \in F_n$ and a lift $\widetilde{r}$ of some $r \in \mathcal{R}(\phi)$. $\mathsf{H}_{\phi,\mathfrak{c}}(L) := [F_{\mathfrak{c}}(\widetilde{r}), b]$. To $\mathsf{H}_{\phi,\mathfrak{c}}(L)$ we associate the set of lines $[\partial F_{\mathfrak{c}}(\widetilde{r}), b^+]$. $\mathsf{H}_{\phi,\mathfrak{c}}(\widetilde{L}) := (F_{\mathfrak{c}}(\widetilde{r}), b)$ with associated set of lines $(\partial F_{\mathfrak{c}}(\widetilde{r}), b^+)$.

[NP-NP] L has type NP-NP if some (hence every) lift $\widetilde{L}$ has the form $(\widetilde{r}, \widetilde{s})$ for lifts $\widetilde{r}$ of $r \in \mathcal{R}(\phi)$ and $\widetilde{s}$ of $s \in \mathcal{R}(\phi)$. $\mathsf{H}_{\phi,\mathfrak{c}}(\widetilde{L}) := [F_{\mathfrak{c}}(\widetilde{r}), F_{\mathfrak{c}}(\widetilde{s})]$. To $\mathsf{H}_{\phi,\mathfrak{c}}(\widetilde{L})$ we associate the set of lines $[\partial F_{\mathfrak{c}}(\widetilde{r}), \partial F_{\mathfrak{c}}(\widetilde{s})]$. $\mathsf{H}_{\phi,\mathfrak{c}}(\widetilde{L}) := (F_{\mathfrak{c}}(\widetilde{r}), F_{\mathfrak{c}}(\widetilde{s}))$ with associated set of lines $(\partial F_{\mathfrak{c}}(\widetilde{r}), \partial F_{\mathfrak{c}}(\widetilde{s}))$.

Lemma 13.5 *Suppose that L lifts to $\Gamma(f)$ and has one of the types* P-P, P-NP, NP-P, *or* NP-NP. *Then with notation as above*:

[P-P] $[\langle a\rangle, \langle b\rangle]$ *is a good conjugacy pair.*

[P-NP] $[\langle a\rangle, F_{\mathfrak{c}}(\widetilde{r})]$ *is a good conjugacy pair.*

[NP-P] $[F_{\mathfrak{c}}(\widetilde{r}), \langle b\rangle]$ *is a good conjugacy pair.*

[NP-NP] $[F_{\mathfrak{c}}(\widetilde{r}), F_{\mathfrak{c}}(\widetilde{s})]$ *is a good conjugacy pair.*

Proof [P-P] Since $a \neq b^{\pm 1}$ and a and b are root-free, $\langle a, b\rangle$ is a free group of rank 2. In particular, $[\langle a\rangle, \langle b\rangle]$ is good by Lemma 7.14(1).

[P-NP] Suppose L has the lift $\widetilde{L} = (a^-, \widetilde{r})$. Set $A = \langle a\rangle$. Since L lifts to $\Gamma(f)$, $L = \alpha^\infty \sigma R_E$ for some $E \in \mathcal{E}_f$, where α is a circuit in the core of $\Gamma(f)$ representing $[a]$ (Lemma 4.2). We choose σ to have minimal length. Let $\star$ be the terminal vertex of E and let $\widetilde{\star}$ be the terminal vertex of the unique lift $\widetilde{E}$ of E in $\widetilde{L}$.

The based labeled graph $(C,\star)$ as in Remark 13.2 immerses to $(G,\star)$, and similarly the based labeled graph $(G_A,\star)$ that is a lollipop formed by the union of a circle labeled α and a segment labeled σE immerses to $(G,\star)$. If we define $(H,\star)$ to be the one-point union of $(G_A,\star)$ and $(C,\star)$ then by construction, the immersions of $(G_A,\star)$ and $(C,\star)$ to $(G_A,\star)$ induce a map of $H \to G$ that does not admit any Stallings folds and so, by [Stallings 1983, Proposition 5.3], induces an injection on the level of fundamental groups. We now have an identification of $(A, F_{\mathfrak{c}}(\widetilde{r}), \langle A, F_{\mathfrak{c}}(\widetilde{r})\rangle)$ and $(\pi_1(G_A,\star), \pi_1(C,\star), \pi_1(H,\star))$. By Van Kampen, we see that $\langle A, F_{\mathfrak{c}}(\widetilde{r})\rangle$ is the internal free product of A and $F_{\mathfrak{c}}(\widetilde{r})$.

The cases [NP-P] and [NP-NP] are similar. □

Remark 13.6 $\mathsf{H}_{\phi,\mathfrak{c}}(L)$ is an element of $\mathbb{A}_1 \sqcup \mathbb{A}_2 \sqcup \mathbb{A}_3 \sqcup \mathbb{A}_4 \subset \overline{\mathsf{IS}}(\mathbb{A}_\bullet)$. Not all lines that lift to $\Gamma(f)$ are assigned a type. For each L that has a type, $\mathsf{H}_{\phi,\mathfrak{c}}(L)$ can be recovered from its associated set of lines. In the NP-NP case, this is a direct application of Corollary 7.5, Remark 7.6 and Lemma 7.14(2). The obvious modification needed for the other cases where $\langle a\rangle$ is replaced by a is left to the reader. We often conflate $\mathsf{H}_{\phi,\mathfrak{c}}(L)$ with its associated set of lines.

Example 3.1 (continued) If L is the upward line represented by the contractible component of $\Gamma(f)$ in Figure 2 then L has type NP-NP and $\mathsf{H}_{\phi,\mathfrak{c}}(L) = [\langle a,b\rangle, \langle a,b\rangle^{d^{-1}e}]$ consists of the set of lines in the graph in Figure 4 that cross $d^{-1}e$ once and $e^{-1}d$ not at all.

Lemma 13.7 *Suppose that $\mathfrak{c}$ is a special chain for ϕ, that L is a nonperiodic principal line for ϕ whose nonperiodic ends are contained in $\mathcal{R}(\phi)$ and that $\theta \in \mathsf{Out}(F_n)$. Then $\theta(\mathfrak{c})$ is a special chain for ϕ^θ, $\theta(L)$ is a nonperiodic principal line for ϕ^θ whose nonperiodic ends are contained in $\mathcal{R}(\phi^\theta)$ and $\theta(\mathsf{H}_{\phi,\mathfrak{c}}(L)) = \mathsf{H}_{\phi^\theta,\theta(\mathfrak{c})}(\theta(L))$.*

Proof We will do the case P-NP; the others are similar. Lemmas 6.13 and 3.16 imply that $\theta(\mathfrak{c})$ is a special chain for ϕ^θ and that $\theta(L)$ is a principal line for ϕ^θ. If $\Theta \in \theta$ and $\widetilde{L} = (a^-, \widetilde{r})$, then $\Theta(\widetilde{L}) = (\Theta(a), \Theta(\widetilde{r}))$ and

$$\Theta(\mathsf{H}_{\phi,\mathfrak{c}}(\widetilde{L})) = [\Theta(a), \Theta(F_{\mathfrak{c}}(\widetilde{r}))] = [\Theta(a), F_{\Theta(\mathfrak{c})}(\Theta(\widetilde{r}))] = \mathsf{H}_{\phi^\theta,\theta(\mathfrak{c})}(\Theta(\widetilde{L})).$$ □

Remark 13.8 For applications, it follows from definitions that if $\theta(\mathfrak{c}) = \mathfrak{c}$ and $\mathfrak{e} \in \mathfrak{c}$ then $\theta(\mathfrak{e}) = \mathfrak{e}$.

In the next lemma we abuse notation and identify $\mathsf{H}_{\phi,\mathfrak{c}}(L)$ with its associated set of lines; see Remark 13.6.

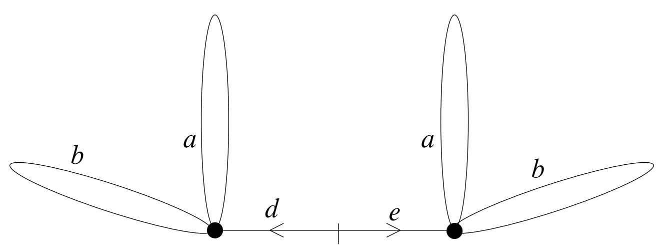

Figure 4

Lemma 13.9 *Suppose $f: G \to G$ is a* CT *for ϕ that realizes $\mathfrak{c}$. If L is either an element of $\Omega_{\mathsf{NP}}(\phi)$ or an element of $\mathsf{L}_{\mathfrak{c}}(\phi)$ such that $\mathfrak{c}$ is not* large, *then L is the only line in $\mathsf{H}_{\phi,\mathfrak{c}}(L)$ that lifts to $\Gamma(f)$.*

Proof We assume at first that $L \in \Omega_{\mathsf{NP}}(\phi)$. By Corollary 5.17(4), all elements of $\Omega_{\mathsf{NP}}(\phi)$ lift to $\Gamma(f)$ and further, by Corollary 5.17(1), have one of the types P-P, P-NP, NP-P or NP-NP. In particular, $\mathsf{H}_{\phi,\mathfrak{c}}(L)$ is defined. We will do the case that $L = [a^-, \tilde{r}]$ is P-NP, the others being similar. By Corollary 5.17(1), L has the form $(R_1)^{-1} R_E$, where R_1 consists of only linear and fixed edges, and where $E \in \mathcal{E}_f$. In particular, E is the highest edge of L.

Every line $L' \in \mathsf{H}_{\phi,\mathfrak{c}}(L)$ has a lift of the form $\tilde{L}' = (a^-, \tilde{s})$, where $\tilde{s} \in F_{\mathfrak{c}}(\tilde{r})$. Since $f: G \to G$ realizes $\mathfrak{c}$, $F_{\mathfrak{c}}(r)$ is represented by a subgraph of G whose edges are all lower than E. If $\tilde{r} = \tilde{s}$ then $L' = L$ and we are done. We may therefore assume $\tilde{L}'' := (\tilde{r}, \tilde{s})$ is a line with both endpoints in $\partial F_{\mathfrak{c}}(\tilde{r})$. Thus L'' only crosses lifts of edges that are lower than E. It follows that $L' = (R_1)^{-1} E R_2$, where the ray R_2 only crosses edges that are lower than E. If L' lifts to $\Gamma(f)$ then E is the first higher-order edge it crosses and so $ER_2 = R_E$ and $L' = L$. This completes the proof when $L \in \Omega_{\mathsf{NP}}(\phi)$.

We now assume that $L \in \mathsf{L}_{\mathfrak{c}}(\phi)$ and that $\mathfrak{c}$ is not large. If $\mathfrak{c}$ is contractible then, because L lifts to $\Gamma(f)$, L has the form $[\tilde{r}, \tilde{s}] = R_{E_1}^{-1} \rho R_{E_2}$, where ρ is a Nielsen path. By definition $L' \in \mathsf{H}_{\phi,\mathfrak{c}}(L)$ has a representative $\tilde{L}' = (\tilde{r}_1, \tilde{r}_2)$ such that either $\tilde{r}_1 = \tilde{r}$ or $(\tilde{r}_1, \tilde{r})$ is a line with both endpoints in $\partial F_{\mathfrak{c}}(\tilde{r})$ and such that either $\tilde{r}_2 = \tilde{s}$ or $(\tilde{s}, \tilde{r}_2)$ is a line with both endpoints in $\partial F_{\mathfrak{c}}(\tilde{s})$. We argue as above to conclude that $L' = L$ if L' lifts to $\Gamma(f)$. The case where $\mathfrak{c}$ is infinite cyclic is similar. □

Lemma 13.10 *Suppose that L has one of the types* P-P, P-NP, NP-P *or* NP-NP *and that $\theta(\mathsf{H}_{\phi,\mathfrak{c}}(L)) = \mathsf{H}_{\phi,\mathfrak{c}}(L)$. If $\tilde{L}$ is a lift of L then there is a unique $\Theta \in \theta$ such that $\Theta(\mathsf{H}_{\phi,\mathfrak{c}}(\tilde{L})) = \mathsf{H}_{\phi,\mathfrak{c}}(\tilde{L})$.*

Proof The existence of at least one such Θ follows from the definitions.

(P-P) Suppose $\tilde{L} = (a^-, b^+)$. If $\Theta_1 \neq \Theta_2$ represent θ and leave (a^-, b^+) invariant then the difference $\Theta_1 \Theta_2^{-1}$ has the form i_x for some $x \neq 1$ in F_n and leaves (a^-, b^+) invariant. It follows that a, b and x share a power. This is impossible since L is not periodic.

(NP-P) Suppose $\tilde{L} = (\tilde{r}, b^+)$. If $\Theta_1 \neq \Theta_2$ in θ leave $(F_{\mathfrak{c}}(\tilde{r}), b^+)$ invariant, then $\Theta_1 \Theta_2^{-1} = i_x$ for some $x \neq 1$ in F_n, $i_x\big(F_{\mathfrak{c}}(\tilde{r})\big) = F_{\mathfrak{c}}(\tilde{r})$, and $i_x(b) = b$. Hence $x \in F_{\mathfrak{c}}(\tilde{r}) \cap \langle b \rangle$, which is impossible since $\partial F_{\mathfrak{c}}(\tilde{r}) \cap \partial \langle b \rangle = \varnothing$ by Lemma 13.5.

The cases P-NP and NP-NP are similar. □

13.6 Algebraic strong axes

- (algebraic strong axes) If $[\Phi, a]$ is a strong axis and $\mathsf{Fix}(\Phi)$ is not cyclic (equivalently $\mathrm{rank}(\mathsf{Fix}(\Phi)) > 1$), then $[\mathsf{Fix}(\Phi), a]$ is the associated algebraic strong axis.

Remark 13.11 In (algebraic strong axes) above we are focusing only on strong axes for the restriction of ϕ to its linear free factor system $\mathcal{F}_0(\phi)$; see the proof of Lemma 17.1. By definition of strong axes, $a \in \mathsf{Fix}(\Phi)$. In particular, $[\mathsf{Fix}(\Phi), a]$ is an element of $\mathbb{A}_6 \subset \overline{\mathsf{IS}}(\mathbb{A}_\bullet)$ and is not good. The set of algebraic strong axes of ϕ is an element of $\mathcal{S}_{\mathrm{un}}(\mathbb{A}_6) \subset \overline{\mathsf{IS}}(\mathbb{A}_\bullet)$. By Lemma 4.20, $[\mathsf{Fix}(\Phi), a]$ determines $[\Phi, a]$.

Remark 4.13 can be used to compute the algebraic strong axes from our CT $f : G \to G$. Using the notation there, the set $\mathsf{SA}(\phi)$ of strong axes is in one-to-one correspondence with the set of nontrivial circuits in the core of $\Gamma(f)$ representing elements of $\mathcal{A}_{\mathrm{or}}(\phi)$. The strong axis $[\Phi_{a,0}, a]$ corresponding to $(f_{v\#}, \tau)$ has algebraic invariant represented by $[\mathsf{Fix}(f_{v\#}), a]$. $\mathsf{Fix}(f_{v\#})$ is the image in $\pi_1(G, v)$ of $\pi_1(\Gamma(f), v)$, where $v \in \Gamma^0(f)$ is viewed as an element of G; see the construction of $\Gamma(f)$ in Section 4.1. The case for $(f_{v_j\#}, v_j)$ is similar.

In our running example (see page 1716), the algebraic invariants of α and α' are respectively represented by the pairs $(\langle a, a^b\rangle, a)$ and $(\langle a, a^b\rangle, a^b)$. The strong axis α'' doesn't have an algebraic invariant since the corresponding fix set is cyclic. This ends Remark 13.11.

13.7 Algebraic added lines

- (algebraic added lines with respect to $\mathfrak{e}$) We use notation as in Definition 6.14. $\mathsf{H}_{\mathfrak{e}\in\mathfrak{c}}(\phi)$ is defined to be $\{\mathsf{H}_{\phi,\mathfrak{c}}(L) \mid L \in \mathsf{L}_{\mathfrak{e}}(\phi)\}$ if $\mathsf{L}_{\mathfrak{e}}(\phi)$ is finite, and is defined to be the singleton $\{[\mathsf{Fix}(\Phi), F_{\mathfrak{c}}(\tilde{r}^+)]\}$ otherwise. The proof that $[\mathsf{Fix}(\Phi), F_{\mathfrak{c}}(\tilde{r}^+)]$ is good is similar to the proof of Lemma 13.5. $\mathsf{H}_{\mathfrak{e}\in\mathfrak{c}}(\phi)$ is an element of $\mathcal{S}_{\mathrm{un}}(\mathbb{A}_1 \sqcup \mathbb{A}_2 \sqcup \mathbb{A}_3) \subset \overline{\mathsf{IS}}(\mathbb{A}_\bullet)$.

In terms of our CT $f : G \to G$, algebraic added lines can be computed as follows. We use the notation of Definition 6.14. Since we already know how to compute algebraic lines, in the cases where the number of added lines is finite, it is enough to describe the added lines. This is done in Definition 6.14 using the eigengraph $\Gamma(f|F^+)$. The case where there are infinitely many added lines is the intersection of [H] and large in Lemma 6.10. Let v be the initial vertex of the edge $E \in \mathcal{E}_f$ in $\Gamma(f|F^+)$ corresponding to r^+. Then $[\mathsf{Fix}(\Phi), F_{\mathfrak{c}}(\tilde{r}^+)]$ is represented by $(\mathsf{Fix}(f_{v\#}|F^+), \pi_1(C^E, v))$, where C is the connected, core subgraph of G representing $F_{\mathfrak{c}}(r^+)$ and C^E denotes the one-point union at the terminal endpoint of E of C and an edge labeled E (so that C^E immerses to G). See our running example on page 1730, where $\mathfrak{e}_2$ is contractible, $\mathfrak{e}_3$ is infinite cyclic, $\mathfrak{e}_1$ is large and

$$\mathsf{H}_{\mathfrak{e}_1\in\mathfrak{c}}(\phi) = \{[\langle a, a^b\rangle, \langle a, b\rangle^c]\},$$
$$\mathsf{H}_{\mathfrak{e}_2\in\mathfrak{c}}(\phi) = \{[\langle a, b\rangle, \langle a, b\rangle^{d^{-1}e}], [\langle a, b\rangle, \langle a, b\rangle^{e^{-1}d}]\},$$
$$\mathsf{H}_{\mathfrak{e}_3\in\mathfrak{c}}(\phi) = \{[a^{-1}, \langle a, b, c\rangle^{p^{-1}q}], [a, \langle a, b, c\rangle^{p^{-1}q}]\}.$$

13.8 Algebraic limit lines

- $\{\mathsf{H}_{\mathfrak{c}}(L) \mid L \in \Omega_{\mathsf{NP}}(\phi)\}$ is the set of algebraic limit lines.

The set of algebraic limit lines is an element of $\mathcal{S}_{\text{un}}(\mathbb{A}_1 \sqcup \mathbb{A}_2 \sqcup \mathbb{A}_3 \sqcup \mathbb{A}_4) \subset \overline{\mathsf{IS}}(\mathbb{A}_\bullet)$. As in the case of added lines, to compute algebraic limit lines, we only need to compute $\Omega_{\mathsf{NP}}(\phi)$ from our CT $f: G \to G$. This is done in Section 5; see Corollary 5.17. Referring to our running example (see page 1724), $\Omega_{\mathsf{NP}}(\phi) = \Omega_{\mathsf{NP}}(r_q) = \{a^\infty R_c, a^\infty b a^\infty\}$ and so $\mathsf{H}_{\mathfrak{c}}(\phi) = \{[a^{-1}, \langle a, b\rangle^c], [a^{-1}, a^b]\}$.

13.9 Naturality

We will need the following naturality statements.

Lemma 13.12 *Suppose* $\mathfrak{e} \in \mathfrak{c}$ *and* $\Theta \in \theta \in \mathsf{Out}(F_n)$. *Then*:

(1) $\theta(\{\mathsf{H}_{\phi,\mathfrak{c}}(L) \mid L \in \Omega_{\mathsf{NP}}(\phi)\}) = \{\mathsf{H}_{\phi^\theta,\theta(\mathfrak{c})}(L') \mid L' \in \Omega_{\mathsf{NP}}(\phi^\theta)\}$.

(2) $\theta(\mathsf{H}_{\mathfrak{e}\in\mathfrak{c}}(\phi)) = \mathsf{H}_{\theta(\mathfrak{e})\in\theta(\mathfrak{c})}(\phi^\theta)$.

(3) $[\mathsf{Fix}(\Phi), a] \leftrightarrow \theta([\mathsf{Fix}(\Phi), a]) = ([\mathsf{Fix}(\Phi^\Theta), \Theta(a)]$ *defines a bijection between the algebraic strong axes for* ϕ *and the algebraic strong axes for* ϕ^θ.

Proof (1) We have

$$\begin{aligned}\theta(\{\mathsf{H}_{\mathfrak{c}}(L) \mid L \in \Omega_{\mathsf{NP}}(\phi)\}) &\overset{\text{def}}{=} \{\theta(\mathsf{H}_{\mathfrak{c}}(L)) \mid L \in \Omega_{\mathsf{NP}}(\phi)\} \\ &= \{\mathsf{H}_{\theta(\mathfrak{c})}(\theta(L)) \mid L \in \Omega_{\mathsf{NP}}(\phi)\} \quad \text{by Lemma 13.7,} \\ &= \{\mathsf{H}_{\theta(\mathfrak{c})}(L') \mid L' \in \Omega_{\mathsf{NP}}(\phi^\theta)\} \quad \text{by Corollary 5.4.}\end{aligned}$$

(2) If $\mathsf{L}_{\mathfrak{e}}(\phi)$ is finite,

$$\begin{aligned}\theta(\mathsf{H}_{\mathfrak{e}\in\mathfrak{c}}(\phi)) &\overset{\text{def}}{=} \theta(\{\mathsf{H}_{\mathfrak{c}}(L) \mid L \in \mathsf{L}_{\mathfrak{e}}(\phi)\}) \\ &= \{\theta(\mathsf{H}_{\mathfrak{c}}(L)) \mid L \in \mathsf{L}_{\mathfrak{e}}(\phi)\}) \\ &= \{\mathsf{H}_{\theta(\mathfrak{c})}(\theta(L)) \mid L \in \mathsf{L}_{\mathfrak{e}}(\phi)\} \quad \text{by Lemma 13.7,} \\ &= \{\mathsf{H}_{\theta(\mathfrak{c})}(L') \mid L' \in \mathsf{L}_{\theta(\mathfrak{e})}(\phi^\theta)\} \quad \text{by Lemma 6.15,} \\ &\overset{\text{def}}{=} \mathsf{H}_{\theta(\mathfrak{e})\in\theta(\mathfrak{c})}(\phi^\theta).\end{aligned}$$

If $\mathsf{L}_{\mathfrak{e}}(\phi)$ is infinite,

$$\begin{aligned}\theta(\mathsf{H}_{\mathfrak{e}\in\mathfrak{c}}(\phi)) &\overset{\text{def}}{=} \theta(\{[\mathsf{Fix}(\Phi), F_{\mathfrak{c}}(\tilde{r}^+)]\}) \\ &= \{[\Theta(\mathsf{Fix}(\Phi)), \Theta(F_{\mathfrak{c}}(\tilde{r}^+))]\} \\ &= \{[\mathsf{Fix}(\Phi^\Theta), F_{\theta(\mathfrak{c})}(\Theta(\tilde{r}^+))\} \quad \text{by Lemma 6.13,} \\ &\overset{\text{def}}{=} \mathsf{H}_{\theta(\mathfrak{e})\in\theta(\mathfrak{c})}(\phi^\theta).\end{aligned}$$

(3) Lemmas 4.17(2) and 3.16(2) imply that $(\Phi, a) \leftrightarrow (\Phi^\Theta, \Theta(a))$ induces a bijection $\mathsf{SA}(\phi, [a]) \leftrightarrow \mathsf{SA}(\psi, \theta([a]))$ and that the ranks of $\mathsf{Fix}(\Phi)$ and $\mathsf{Fix}(\Phi^\Theta)$ are equal. (3) therefore follows from

$$\theta([\mathsf{Fix}(\Phi), a]) \overset{\text{def}}{=} [\Theta(\mathsf{Fix}(\Phi)), \Theta(a)] = [\mathsf{Fix}(\Phi^\Theta), \Theta(a)]. \qquad \square$$

13.10 The algebraic invariant of ϕ rel $\mathfrak{c}$

In this subsection we collect our algebraic invariants into a single master algebraic invariant.

Definition 13.13 Fix a special chain $\mathfrak{c}$ for $\phi \in \mathsf{UPG}$.

(1) The *algebraic invariant of* ϕ *rel* $\mathfrak{c}$ is the element of $\overline{\mathsf{IS}}(\mathbb{A}_\bullet)$ that is the ordered set $\mathsf{I}_\mathfrak{c}(\phi)$ consisting of
  - $\mathfrak{c}$,
  - $\mathsf{Fix}(\phi)$,
  - $(\mathsf{H}_{\mathfrak{e}\in\mathfrak{c}}(\phi) \mid \mathfrak{e} \in \mathfrak{c})$, where the special 1-edge extensions $\mathfrak{e}$ are ordered using $\mathfrak{c}$,
  - $\{\mathsf{H}_{\phi,\mathfrak{c}}(L) \mid L \in \Omega_{\mathsf{NP}}(\phi)\}$,
  - $\mathcal{A}_{\mathrm{or}}(\phi)$, and
  - the set of algebraic strong axes for ϕ.

  The six elements in the ordered list $\mathsf{I}_\mathfrak{c}(\phi)$ are the *components of* $\mathsf{I}_\mathfrak{c}(\phi)$.

(2) Order (noncanonically) the elements of the union of the six sets defining $\mathsf{I}_\mathfrak{c}(\phi)$. The resulting element of $\overline{\mathsf{IS}}(\mathbb{A}_\bullet)$ is denoted by J.

Remark 13.14 In light of Theorem 3.20, we have seen in this section that the set of special chains, $\mathsf{I}_\mathfrak{c}(\phi)$, and J can be computed. We stress that they take values in $\overline{\mathsf{IS}}(\mathbb{A}_\bullet)$, which satisfies property MW by Lemma 11.2.

14 Stabilizers of algebraic invariants

At this point, it is reasonably straightforward to reduce the conjugacy problem for $\mathsf{UPG}(F_n)$ in $\mathsf{Out}(F_n)$ to the problem of deciding whether $\phi, \psi \in \mathsf{UPG}(F_n)$ with $\mathsf{I}_\mathfrak{c}(\phi) = \mathsf{I}_\mathfrak{c}(\psi)$ are conjugate by some θ in the stabilizer of $\mathsf{I}_\mathfrak{c}(\phi)$. We want a little more. Namely, we want to restrict the set of potential conjugators to those elements that stabilize $\mathsf{I}_\mathfrak{c}(\phi)$ and induce trivial permutations on the components of $\mathsf{I}_\mathfrak{c}(\phi)$. Continuing with the notation of the previous section, we make this precise as follows.

Definition 14.1 ($\mathcal{X}_\mathfrak{c}(\phi)$) Let $\phi \in \mathsf{UPG}(F_n)$, $\mathfrak{c}$ be a special chain for ϕ, and let J be in Definition 13.13(2). Let $\mathcal{X}_\mathfrak{c}(\phi)$ denote the stabilizer $\mathsf{Out}_\mathsf{J}(F_n)$ of J in $\mathsf{Out}(F_n)$.

Unraveling definitions, $\mathcal{X}_\mathfrak{c}(\phi)$ also has a description as the subgroup of $\mathsf{Out}(F_n)$ fixing each element in the union of the following six sets:

(1) $\{[F] \mid [F] \in \mathcal{F} \in \mathfrak{c}\}$,
(2) $\mathsf{Fix}(\phi)$,
(3) $\bigcup_{\mathfrak{e}\in\mathfrak{c}} \mathsf{H}_{\mathfrak{e}\in\mathfrak{c}}(\phi)$,
(4) $\{\mathsf{H}_{\phi,\mathfrak{c}}(L) \mid L \in \Omega_{\mathsf{NP}}(\phi)\}$,
(5) $\mathcal{A}_{\mathrm{or}}(\phi)$, and
(6) $\{[\mathsf{Fix}(\Phi), a] \mid [\Phi, a] \in \mathsf{SA}(\phi), \operatorname{rank} \mathsf{Fix}(\Phi) \geq 2\}$.

Remark 14.2 As noted in Definition 13.13(2), the construction of J was noncanonical. That is, there were choices in its construction. Every choice has the same stabilizer and so $\mathcal{X}_{\mathfrak{c}}(\phi)$ is independent of choice.

In passing and for future use, we have the expected:

Lemma 14.3 $$\phi \in \mathcal{X}_{\mathfrak{c}}(\phi).$$

Proof We have to check that ϕ fixes each of the sets (1)–(6) above elementwise. In (1), (2), (5) and (6) this is immediate from definitions. For set (4), this is because $\phi(L) = L$ for all $L \in \Omega_{\mathsf{NP}}(\phi)$ (Corollary 5.17(4)) and Lemma 13.7. That the elements of set (3) are fixed follows from definitions and Lemmas 6.15 and 13.12(2). □

Definition 14.4 A group G *is of type* F if it has a finite Eilenberg–Mac Lane space. G *is of type* VF if it has a finite-index subgroup of type F.

Proposition 14.5 *The stabilizer* $\mathsf{Out}_{\mathsf{Y}}(F_n)$ *of an element* $\mathsf{Y} \in \overline{\mathsf{IS}}(\mathbb{A}_\bullet)$ *is of type* VF.

Proof We saw in the proof of Lemma 11.2 that there is a map $\overline{\mathsf{IS}}(\mathbb{A}_\bullet) \to \overline{\mathsf{IS}}(\mathbb{A}_0)$ that has explicit finite fibers (Notation 10.16). If $\overline{\mathsf{Y}}$ is the image of Y, then $\mathsf{Out}_{\mathsf{Y}}(F_n)$ has finite index in $\mathsf{Out}_{\overline{\mathsf{Y}}}(F_n)$. By Corollary 10.15, the subgroup G of $\mathsf{Out}(F_n)$ fixing each label of $\overline{\mathsf{Y}}$ has finite index in $\mathsf{Out}_{\overline{\mathsf{Y}}}(F_n)$. Also, the subgroup G has type VF by [Bestvina et al. 2023, Theorem 1.1]. $\mathsf{Out}_{\mathsf{Y}}(F_n)$, being commensurate with a group of type VF, also has type VF. □

As usual, naturality will be important.

Lemma 14.6 *Suppose* $\xi \in \mathsf{Out}(F_n)$.

(1) $(\mathcal{X}_{\mathfrak{c}}(\phi))^{\xi} = \mathcal{X}_{\xi(\mathfrak{c})}(\phi^{\xi})$.

(2) $\xi(\mathsf{I}_{\mathfrak{c}}(\phi)) = \mathsf{I}_{\xi(\mathfrak{c})}(\phi^{\xi})$.

Proof (1) By Lemma 6.13, $\xi(\mathfrak{c})$ is special for ϕ^{ξ} and so the statement makes sense. The lemma follows easily from the naturality of the quantities appearing in Definition 14.1. For example, we verify that if $\theta \in \mathcal{X}_{\mathfrak{c}}(\phi)$, then $\theta^{\xi}(\mathsf{H}_{\phi^{\xi},\xi(\mathfrak{c})}(L)) = \mathsf{H}_{\phi^{\xi},\xi(\mathfrak{c})}(L)$ for all $L \in \Omega_{\mathsf{NP}}(\phi^{\xi})$. Indeed,

$$\xi\theta\xi^{-1}(\mathsf{H}_{\phi^{\xi},\xi(\mathfrak{c})}(L)) = \xi\theta(\mathsf{H}_{\phi,\mathfrak{c}}(\xi^{-1}(L))) = \xi(\mathsf{H}_{\phi,\mathfrak{c}}(\xi^{-1}(L))) = \mathsf{H}_{\phi^{\xi},\xi(\mathfrak{c})}(L),$$

where the first and third equalities use Lemma 13.7 and the second uses Corollary 5.4. The remainder of the proof consists of similar checks and is left to the reader.

The proof of (2) is similar. □

We next reduce the proof of the main result (Theorem 1.1) of this paper, ie the conjugacy problem for $\mathsf{UPG}(F_n)$ in $\mathsf{Out}(F_n)$, to the proof of Proposition 14.7 stated immediately below.

Proposition 14.7 *There is an algorithm that takes as input $\phi, \psi \in \mathsf{UPG}(F_n)$ and a chain $\mathfrak{c}$ such that*

- *$\mathfrak{c}$ is special for both ϕ and ψ, and*
- *$\mathsf{I}_{\mathfrak{c}}(\phi) = \mathsf{I}_{\mathfrak{c}}(\psi)$,*

and that outputs `YES` *or* `NO` *depending whether or not there is a $\theta \in \mathcal{X}_{\mathfrak{c}}(\phi)$ conjugating ϕ to ψ. Further, if* `YES` *then such a θ is produced.*

Proposition 14.7 is proved in Sections 16 and 17 below.

Lemma 14.8 *Proposition 14.7 implies Theorem 1.1. That is, an algorithm that satisfies the conclusions of Proposition 14.7 can be used to produce an algorithm that satisfies the conclusions of Theorem 1.1.*

Proof Assume Proposition 14.7 holds and $\phi, \psi \in \mathsf{UPG}(F_n)$.

View the multiset $\mathsf{I}(\phi) := \{\mathsf{I}_{\mathfrak{c}}(\phi) \mid \mathfrak{c} \text{ is a special chain for } \phi\}$ as an element of $\overline{\mathsf{IS}}(\mathbb{A}_\bullet)$ as described in Section 13.1. By Lemma 14.6(2), $\mathsf{I}(\phi)$ is a conjugacy invariant of ϕ. That is, if $\phi^\theta = \psi$ then $\theta(\mathsf{I}(\phi)) = \mathsf{I}(\cdot)$. We may compute $\mathsf{I}(\phi)$ and $\mathsf{I}(\psi)$; see Remark 13.14.

Since $\overline{\mathsf{IS}}(\mathbb{A}_\bullet)$ satisfies property MW (Lemma 11.2), we can algorithmically check if there is $\theta' \in \mathrm{Out}(F_n)$ such that $\theta'(\mathsf{I}(\phi)) = \mathsf{I}(\psi)$. If there is no such θ' then ϕ and ψ are not conjugate; return `NO`. If there is such a θ', then one is produced by the M-algorithm for $\overline{\mathsf{IS}}(\mathbb{A}_\bullet)$. Note that ϕ and ψ are conjugate in $\mathrm{Out}(F_n)$ if and only if ϕ and $\psi' := \theta'^{-1}\psi\theta'$ are conjugate in $\mathrm{Out}(F_n)$ if and only if ϕ and ψ' are conjugate in the stabilizer $G := \mathrm{Out}_{\mathsf{I}(\phi)}(F_n)$ of $\mathsf{I}(\phi)$.

G acts by permutation on the set of labels of $\mathsf{I}(\phi)$. Let $G' < G$ denote the subgroup fixing each label. By the M-algorithm for $\overline{\mathsf{IS}}(\mathbb{A}_\bullet)$, we may construct a finite presentation for G. Using our finite set of generators for G, we may construct the image Q of G in our permutation group. By Lemma 9.1, we can compute a finite set θ_i such that $G = \bigsqcup_i \theta_i G'$. Hence, ϕ and ψ' are conjugate in $\mathrm{Out}(F_n)$ if and only if ϕ and some $\theta_i^{-1}\psi'\theta_i$ are conjugate in $\mathrm{Out}(F_n)$ if and only if ϕ and some $\theta_i^{-1}\psi'\theta_i$ are conjugate in G'. Since $G' < \mathcal{X}_{\mathfrak{c}}(\phi) < \mathrm{Out}(F_n)$ for all $\mathfrak{c}$, ϕ and ψ' are conjugate in $\mathrm{Out}(F_n)$ if and only if ϕ and some $\theta_i^{-1}\psi'\theta_i$ are conjugate in $\mathcal{X}_{\mathfrak{c}}(\phi)$. We may use the supposed algorithm of Proposition 14.7 to decide whether or not this is the case and return a conjugator if it is. The returned conjugator allows us to compute a conjugator for ϕ and ψ. □

15 Staple pairs

15.1 Limit lines $\Omega_{\mathsf{NP}}(\phi, \tilde{r}) \subset \tilde{\mathcal{B}}$

In Section 5, we associated a finite set $\Omega_{\mathsf{NP}}(r) \subset \mathcal{B}$ of ϕ-invariant nonperiodic lines to each $r \in \mathcal{R}(\phi)$. In this section we associate, to each lift $\tilde{r}$ of r, a subset $\Omega_{\mathsf{NP}}(\phi, \tilde{r}) \subset \tilde{\mathcal{B}}$ of the full preimage of $\Omega_{\mathsf{NP}}(r)$ and then establish properties of $\Omega_{\mathsf{NP}}(\phi, \tilde{r})$ that will be needed later in the paper.

Definition 15.1 Choose a marked graph K. For each lift $\tilde r \in \partial F_n$ of $r \in \mathcal{R}(\phi)$, let $\Phi_{\tilde r}$ be the unique lift of ϕ that fixes $\tilde r$ and let $\tilde R \subset \tilde K$ be a ray with terminal end $\tilde r$. If $\tilde L$ is a lift of $L \in \Omega_{\mathsf{NP}}(r)$ then L is ϕ-invariant by Corollary 5.17(1) and so each $\Phi_{\tilde r}^j(\tilde L)$ is a translate of $\tilde L$, say $\Phi_{\tilde r}^j(\tilde L) = T_j(\tilde L)$ for some unique T_j. Define $\tilde L$ to be in $\Omega_{\mathsf{NP}}(\phi, \tilde r)$ if for every finite subpath $\tilde \beta$ of $\tilde L$ there exists $J(\tilde\beta)$ such that $T_j(\tilde\beta) \subset \tilde R$ — equivalently, $\tilde\beta \subset T_j^{-1}(\tilde R)$ — for all $j \ge J(\tilde\beta)$.

Remark 15.2 As defined, $\Omega_{\mathsf{NP}}(\phi, \tilde r)$ depends on $\Phi_{\tilde r}$ and hence on ϕ, in contrast to $\Omega_{\mathsf{NP}}(r)$, which is independent of ϕ.

Lemma 15.3 *$\Omega_{\mathsf{NP}}(\phi, \tilde r)$ is well-defined and $\Phi_{\tilde r}$-invariant. Moreover, if $\psi = \theta\phi\theta^{-1}$ for some $\theta \in \mathsf{Out}(F_n)$ and if Θ is a lift of θ, then $\Theta(\Omega_{\mathsf{NP}}(\phi, \tilde r)) = \Omega_{\mathsf{NP}}(\psi, \partial\Theta(\tilde r))$.*

Proof Replacing $\tilde R$ by a subray does not change $\Omega_{\mathsf{NP}}(\phi, \tilde r)$. Since any two rays with terminal end $\tilde r$ share a common subray, $\Omega_{\mathsf{NP}}(\phi, \tilde r)$ is independent of the choice of $\tilde R$.

As defined above, $\Omega_{\mathsf{NP}}(\phi, \tilde r)$ depends on the marked graph K so we write $\Omega_{\mathsf{NP}}(\phi, \tilde r, K)$ to make this explicit. We will prove:

(∗) $\Theta(\Omega_{\mathsf{NP}}(\phi, \tilde r, K)) = \Omega_{\mathsf{NP}}(\theta\phi\theta^{-1}, \Theta(\tilde r), K')$ *for any marked graphs K and K' and any $\Theta \in \mathsf{Aut}(F_n)$ representing any $\theta \in \mathsf{Out}(F_n)$.*

Applied with $\Theta =$ identity, (∗) proves that $\Omega_{\mathsf{NP}}(\phi, \tilde r, K)$ is independent of K and hence that $\Omega_{\mathsf{NP}}(\phi, \tilde r)$ is well defined. The moreover statement is equivalent to (∗) and $\Phi_{\tilde r}$-invariance of $\Omega_{\mathsf{NP}}(\phi, \tilde r)$ is an immediate consequence of the definitions. Thus the proof of the lemma will be complete once we prove (∗).

Assume the notation of Definition 15.1. Let $\tilde r' = \Theta(\tilde r)$ and $\psi = \theta\phi\theta^{-1}$; note that $\Psi_{\tilde r'} = \Theta\Phi_{\tilde r}\Theta^{-1}$. Choose a homotopy equivalence $g\colon K \to K'$ of marked graphs that represents θ when $\pi_1(K)$ and $\pi_1(K')$ are identified with F_n via their markings. Let $\tilde g\colon \tilde K \to \tilde K'$ be the lift satisfying $\tilde g|\partial F_n = \Theta|\partial F_n$, let $\tilde R' = \tilde g_\#(\tilde R) \subset \tilde K'$, let $\tilde L' = \Theta(\tilde L) = \tilde g_\#(\tilde L)$ and let $T_j'\colon \tilde K' \to \tilde K'$ be the covering translation satisfying $T_j'|\partial F_n = (\Theta T_j \Theta^{-1})|\partial F_n$. Then

$$\Psi_{\tilde r'}^j(\tilde L') \cap \tilde R' = \Psi_{\tilde r'}^j(\tilde g_\#(\tilde L)) \cap \tilde g_\#(\tilde R) = \tilde g_\#(\Phi_{\tilde r}^j(\tilde L)) \cap \tilde g_\#(\tilde R).$$

By [Cooper 1987] (see also [Bestvina et al. 1997, Lemma 3.1]), there is a constant C, depending only on g, such that $\tilde g_\#(\Phi_{\tilde r}^j(\tilde L)) \cap \tilde g_\#(\tilde R)$ contains the subpath of $\tilde g_\#(\Phi_{\tilde r}^j(\tilde L) \cap \tilde R)$ obtained by C-trimming (ie removing the first and last C edges) and so contains the subpath of $\tilde g_\#(T_j(\tilde\beta)) = T_j'\tilde g_\#(\tilde\beta)$ obtained by C-trimming for any chosen $\tilde\beta$ and all $j \ge J(\tilde\beta)$. Given a finite subpath $\tilde\beta'$ of $\tilde L'$ choose a finite subpath $\tilde\beta$ of $\tilde L$ such that the C-trimmed subpath of $\tilde g_\#(\tilde\beta)$ contains $\tilde\beta'$. Then $\Psi_{\tilde r'}^j(\tilde L') \cap \tilde R' \supset T_j'(\tilde\beta')$ for all $j \ge J(\tilde\beta)$. Letting $J(\tilde\beta') = J(\tilde\beta)$, we conclude that $\tilde L' \in \Omega_{\mathsf{NP}}(\phi, \tilde r)$. By symmetry, we have proved (∗). □

Our goal in the remainder of this subsection is to understand $\Omega_{\mathsf{NP}}(\phi, \tilde r)$ from the CT point of view.

Notation 15.4 Choose $r \in \mathcal{R}(\phi)$ and a CT $f: G \to G$ representing ϕ; let $E \in \mathcal{E}_f$ correspond to r as in Lemma 3.26. Following the proof of Corollary 5.17, let

$$R_E = E \cdot \rho_0 \cdot \sigma_1 \cdot \rho_1 \cdot \sigma_2 \cdot \ldots$$

be the *coarsened complete splitting* of R_E, where each σ_i is a single growing term in the complete splitting of R_E and each ρ_i is a (possibly trivial) Nielsen path. For future reference, note that if $f(E) = E \cdot u$ then Eu is an initial subpath of R_E whose terminal endpoint is a splitting vertex in the complete splitting of R_E and hence is contained in some ρ_p.

Following Notation 5.11 and Lemma 5.14, define, for all $i \geq 1$,

$$R_i^- = f_\#^\infty(\overline{\sigma}_i), \quad R_i^+ = f_\#^\infty(\sigma_i), \quad \ell_i = (R_i^-)^{-1}\rho_i(R_{i+1}^+).$$

Choose a lift $\widetilde{r}$ of r, let $\Phi_{\widetilde{r}}$ be the automorphism representing ϕ that fixes $\widetilde{r}$ and let $\widetilde{f}: \widetilde{G} \to \widetilde{G}$ be the lift corresponding to $\Phi_{\widetilde{r}}$. Let $\widetilde{R}_{\widetilde{E}}$ be the lift of R_E whose terminal end converges to $\widetilde{r}$ and whose initial edge is denoted by $\widetilde{E}$, let

$$\widetilde{R}_{\widetilde{E}} = \widetilde{E} \cdot \widetilde{\rho}_0 \cdot \widetilde{\sigma}_1 \cdot \widetilde{\rho}_1 \cdot \widetilde{\sigma}_2 \cdot \ldots$$

be the induced decomposition and let $\widetilde{\ell}_i$ be the lift of ℓ_i in which ρ_i lifts to $\widetilde{\rho}_i$. Thus

$$\widetilde{\ell}_i = (\widetilde{R}_i^{\,-})^{-1}\widetilde{\rho}_i \widetilde{R}_{i+1}^+,$$

where $\widetilde{\sigma}_{i+1}$ and $\widetilde{R}_{i+1}^+$ have the same initial endpoint and likewise for $\widetilde{\sigma}_i^{-1}$ and $\widetilde{R}_i^-$. We say that lines $\widetilde{\ell}_1, \widetilde{\ell}_2, \ldots$ are *visible* in $\widetilde{R}_{\widetilde{E}}$.

Lemma 15.5 *Assume Notation 15.4.*

1. *Each ℓ_i is an element of $\Omega(r)$; see Definition 5.1.*
2. *If $\ell_i \in \Omega_{\mathsf{NP}}(r)$, then $\widetilde{\ell}_i \in \Omega_{\mathsf{NP}}(\phi, \widetilde{r})$.*

Proof Item (1) follows from Lemmas 5.8 and 5.16 applied with $\alpha = \sigma_i$ and $\beta = \rho_i\sigma_{i+1}$.

When verifying that $\widetilde{\ell}_i$ satisfies Definition 15.1, it suffices to consider finite subpaths $\widetilde{\beta} = \widetilde{\mu}^{-1}\widetilde{\rho}_i\widetilde{\nu}$ of $\widetilde{\ell}_i$ with projections $\beta = \mu^{-1}\rho_i\nu$, where μ is an initial segment of $R_i^- = f_\#^\infty(\overline{\sigma}_i)$ that is a concatenation of terms in the coarsened complete splitting of R_i^- and ν is an initial segment of $R_{i+1}^+ = f_\#^\infty(\sigma_{i+1})$ that is a concatenation of terms in the coarsened complete splitting of R_{i+1}^+. It follows from the definition of $f_\#^\infty$ (Notation 5.11) that for all sufficiently large j, the lift of ρ_i to $\widetilde{f}_\#^j(\widetilde{\rho}_i)$ extends to a lift of β to a path

$$\widetilde{\beta}_j \subset \widetilde{f}_\#^j(\widetilde{\sigma}_i) \cdot \widetilde{f}_\#^j(\widetilde{\rho}_i) \cdot \widetilde{f}_\#^j(\widetilde{\sigma}_{i+1}) = \widetilde{f}_\#^j(\widetilde{\sigma}_i \cdot \widetilde{\rho}_i \cdot \widetilde{\sigma}_{i+1}) \subset \widetilde{R}_{\widetilde{E}}.$$

Since $f_\#^j$ preserves ρ_i, R_i^- and R_{i+1}^+, there is a covering translation T_j such that

$$T_j(\widetilde{\rho}_i) = \widetilde{f}_\#^j(\widetilde{\rho}_i), \quad T_j(\widetilde{R}_i^-) = \widetilde{f}_\#^j(\widetilde{R}_i^-), \quad T_j(\widetilde{R}_{i+1}^+) = \widetilde{f}_\#^j(\widetilde{R}_{i+1}^+),$$

and so

$$T_j(\widetilde{\ell}_i) = \widetilde{f}^j_\#(\widetilde{\ell}_i).$$

From $T_j(\widetilde{\rho}_i) = \widetilde{f}^j_\#(\widetilde{\rho}_i)$ we conclude that $T_j(\widetilde{\beta}) = \widetilde{\beta}_j$ and so $T_j(\widetilde{\beta}) \subset \widetilde{R}_{\widetilde{E}}$. This completes the proof of (2). □

Our next result is a weak converse of Lemma 15.5(2), namely that if $\widetilde{L} \in \Omega_{\mathsf{NP}}(\phi, \widetilde{r})$ then $\widetilde{L}$ is in the $\Phi_{\widetilde{r}}$-orbit of some $\widetilde{\ell}_i$.

Proposition 15.6 *Assume Notation 15.4.*

(1) *For each $\widetilde{L} \in \Omega_{\mathsf{NP}}(\phi, \widetilde{r})$ there exists K such that $\widetilde{f}^k_\#(\widetilde{L}) \in \{\widetilde{\ell}_i\}$ for all $k \geq K$. Moreover, $\Omega_{\mathsf{NP}}(\phi, \widetilde{r}) = \bigcup \Phi^m_{\widetilde{r}}(\widetilde{\ell}_i)$, where the union varies over all $\widetilde{\ell}_i \in \Omega_{\mathsf{NP}}(\phi, \widetilde{r})$ and all $m \in \mathbb{Z}$.*

(2) *For each $L \in \Omega_{\mathsf{NP}}(r)$ there is a lift $\widetilde{L} \in \Omega_{\mathsf{NP}}(\phi, \widetilde{r})$.*

We delay the proof of Proposition 15.6 for two needed lemmas.

Lemma 15.7 *Given a* CT $f\colon G \to G$, *there exists $M \geq 1$ so that the following holds for each twist path w, each nonfixed edge E and each $k \geq 0$: If $|m| \geq M$ and $\alpha_0 = w^m$ is a subpath of $f^k_\#(E)$ then α_0 extends to a subpath α_1 of $f^k_\#(E)$ satisfying the following two properties*:

(1) *$\alpha_1 = E' w^q$ or $\alpha_1 = w^q \bar{E}'$ for some $E' \in \mathsf{Lin}_w(f)$.*

(2) *α_1 is not contained in any Nielsen subpath of $f^k_\#(E)$.*

Proof Let us first note that if the conclusions of the lemma hold for a subpath of $\alpha_0 = w^m$ of the form w^t then they also hold for α_0. We may therefore shorten α_0 whenever it is convenient. After replacing E by $\bar{E}$ if necessary, we may assume that $E \in \mathcal{E}_f \cup \mathsf{Lin}(f)$.

Choose $M'' > 0$ so that if $w_1 \neq w_2$ are twist paths then $w_1^{M''}$ is not a subpath of w_2^m for any $m \in \mathbb{Z}$. Items (1) and (2) hold for $M = M''$ and $E \in \mathsf{Lin}(f)$. We may therefore assume that $E \in \mathcal{E}_f$ and that there exists $M' \geq M''$ such that (1) and (2) hold for $M = M'$ and for all edges $E' \in \mathcal{E}_f$ with height less than that of E.

There is a path u with height less than that of E and a complete splitting

$$u = \tau_1 \cdot \ldots \cdot \tau_s \quad \text{such that} \quad f^k(E) = E \cdot u \cdot f_\#(u) \cdot \ldots \cdot f^{k-1}_\#(u)$$

for all $k \geq 1$. Assuming without loss that M' is greater than the length of any τ_j, choose $M_1 \geq sM'$.

As a special case, we prove the lemma when $|m| \geq M_1$ and when $\alpha_0 = w^m$ is contained in some $f^l_\#(u)$. In this case there exists $1 \leq j \leq s$ and $|m'| \geq M'$ and a subpath $\alpha'_0 = w^{m'}$ of α_0 such that $\alpha'_0 \subset f^l_\#(\tau_j)$ for some $1 \leq l \leq k-1$. As observed above, we can replace α_0 with α'_0. Since the length of τ_j is less than M' and the length of $f^l_\#(\tau_j)$ is at least M', τ_j is not a Nielsen path and so is either exceptional or an edge E' with height less than that of E. If τ_j is exceptional then its linear edges must be in the family determined by w and we take α_1 to be all of τ_j except for the terminal edge. If $\tau_j = E'$, then the inductive hypothesis implies that α_0 extends to a subpath α_1 of $f_\#^l(E')$ that is not contained in a Nielsen subpath of $f_\#^l(E')$ and that satisfies (1). The hard splitting property of a complete splitting (Lemma 4.11(2) of [Feighn and Handel 2011]) implies that an indivisible Nielsen path in a completely split path is contained in a single term of that splitting. Thus α_1 is not contained in a Nielsen subpath of $f_\#^k(E)$ and so (2) is satisfied and we have completed the proof of the special case.

Now choose M so large that if $|m| \geq M$ and $\alpha_0 = w^m$ is a subpath of $f_\#^k(E)$ then there is a subpath $\alpha_0' = w^{m'}$ of α_0 with $m' \geq M_1$ so that $\alpha_0' \subset f_\#^l(u)$ for some l. The existence of M follows from the fact that the length of $f_\#^l(u)$ goes to infinity with l. Replacing α_0 with α_0', we are reduced to the special case. □

We choose a "central" subpath τ_L of $L \in \Omega_{\mathsf{NP}}(r)$ as follows. By Corollary 5.17 $L = (R^-)^{-1} \cdot \rho \cdot R^+$, where $R^\pm$ satisfy (1a), (1b) or (1c) of Lemma 5.14. In all three cases we will choose $\tau = \tau_L = \tau_-^{-1} \rho \tau_+$, where $\tau_\pm$ is an initial segment of $R^\pm$. Let M be the constant from Lemma 15.7.

- In the case (1a), $R^+ = R_{E'}$ for some $E' \in \mathcal{E}_f$ and we take $\tau_+ = E'$.
- In the case (1b), $R^+ = E' w^{\pm\infty}$ for some $E' \in \mathsf{Lin}_w(f)$ and we take $\tau_+ = E' w^{\pm M}$.
- In the case (1c), $R^+ = w^{\pm\infty}$ and we take $\tau_+ = w^{\pm M}$.

The subpath τ_- is defined symmetrically.

Lemma 15.8 *Assume the notation of Notation 15.4 and of the previous paragraph. Suppose that $\widetilde{\tau}_L \subset \widetilde{L}$ is a lift of $\tau_L \subset L$ and that $\widetilde{\tau}_L \subset \widetilde{R}_{\widetilde{E}}$. Then $\widetilde{L} = \widetilde{\ell}_i$ for some i.*

Proof As a first case, suppose that $\tau_+ = E' \in \mathcal{E}_f$ and so $R^+ = R_{E'}$. Then $\widetilde{\tau}_+$ is a term $\widetilde{\sigma}_{i+1}$ in the coarsened complete splitting of $\widetilde{R}_{\widetilde{E}}$ by Lemma 3.21 and $R^+ = f_\#^\infty(\sigma_{i+1})$ by Examples 5.12. There are three subcases to consider, the first being that $\tau_- = E'' \in \mathcal{E}_f$. In this subcase, $\widetilde{\tau}_-^{-1}$ is also a term $\widetilde{\sigma}_j$ in coarsened complete splitting. Since $\widetilde{\tau}_-$ is separated from $\widetilde{\sigma}_{i+1} = \widetilde{\tau}_+$ by the Nielsen subpath $\widetilde{\rho}$, we have $\widetilde{\tau}_-^{-1} = \widetilde{\sigma}_i$ and $\widetilde{\rho} = \widetilde{\rho}_i$. Thus $f_\#^\infty(\overline{\sigma}_i) = R^-$ and $\widetilde{L} = \widetilde{\ell}_i$.

The second subcase is that $\tau_- = E'' w^{\pm M}$, where $E'' \in \mathsf{Lin}_w(f)$. By Lemma 15.7(2), τ_-^{-1} is not contained in a Nielsen subpath of R_E. It follows that the terminal edge $\widetilde{E''}^{-1}$ of $\widetilde{\tau}_-^{-1}$ is contained in a $\widetilde{\sigma}_j$ that is either a single edge or an exceptional path. As in the previous subcase, $j = i$. Also as in the previous subcase, $\widetilde{\rho} = \widetilde{\rho}_i$, $f_\#^\infty(\overline{\sigma}_i) = R^-$ and $\widetilde{\ell}_i = \widetilde{L}$.

The third and final subcase is that $\tau_- = w^{\pm M}$. Since $\widetilde{\tau}_-^{-1}$ is followed in $\widetilde{R}_{\widetilde{E}}$ by $\widetilde{\rho}\widetilde{E}'$, it is not contained in a subpath of $\widetilde{R}_{\widetilde{E}}$ of the form $\widetilde{w}^m \widetilde{E}_1^{-1}$, where $E_1 \in \mathsf{Lin}_w(f)$. Items (1) and (2) of Lemma 15.7 imply that $\widetilde{\sigma}_i = \widetilde{E}_1$, where $E_1 \in \mathsf{Lin}_w(f)$, and that $\widetilde{\rho}_i = \widetilde{w}^t \widetilde{\rho}$ for some t. Examples 5.12 implies that $f_\#^\infty(\overline{\sigma}_i) = w^{\pm\infty} = R^-$ and so $\widetilde{\ell}_i = \widetilde{L}$. We have now completed the proof in the case that $\tau_+ = E' \in \mathcal{E}_f$. Symmetric arguments apply in the case that $\tau_- = E' \in \mathcal{E}_f$.

Our next case is that $\tau_+ = E' w^{\pm M}$, where $E' \in \mathsf{Lin}_w(f)$ and $R^+ = E' w^{\pm\infty}$. Lemma 15.7(2) implies that the initial edge $\widetilde{E}'$ of $\widetilde{\tau}_+$ is not contained in a Nielsen subpath of $\widetilde{R}_{\widetilde{E}}$ and so is either equal to some $\widetilde{\sigma}_{i+1}$ or is the first edge in some $\widetilde{\sigma}_{i+1}$ that projects to an exceptional path. In either case $f_\#^\infty(\sigma_{i+1}) = E' w^{\pm\infty} = R^+$. The remainder of the proof in this second case is exactly the same as in the first case. Symmetric arguments apply in the case that $\tau_- = E' w^{\pm M}$ with $E' \in \mathsf{Lin}_w(f)$.

We are now reduced to the case that $\tau_+ = w_2^{\pm M}$, $R_+ = w_2^{\pm\infty}, \tau_- = w_1^{\pm M}$ and $R_- = w_1^{\pm\infty}$. Thus $\widetilde{\tau} = \widetilde{w}_1^{\mp M} \widetilde{\rho} \widetilde{w}_2^{\pm M}$. If $w_1 = w_2$ then ρ is not an iterate of $w_1 = w_2$ because L is not periodic. Lemma 15.7(1) implies that $\widetilde{\tau}$ extends to a subpath $\widetilde{E}_1 \widetilde{w}_1^p \widetilde{\rho} \widetilde{w}_2^q \widetilde{E}_2^{-1}$, where $E_i \in \mathsf{Lin}_{w_i}(f)$ and $p, q \in \mathbb{Z}$. It follows that $\widetilde{E}_1 = \widetilde{\sigma}_i$, $\widetilde{\rho}_i = \widetilde{w}_1^p \widetilde{\rho} \widetilde{w}_2^q$ and $\widetilde{E}_2^{-1} = \widetilde{\sigma}_{i+1}$ for some choice of i. As in the previous cases, $\widetilde{\ell}_i = \widetilde{L}$. □

Proof of Proposition 15.6 The first statement of (1) follows from Lemma 15.8 and the definition of $\Omega_{\mathsf{NP}}(\phi, \widetilde{r})$. The moreover statement of (1) follows from the first statement and $\Phi_{\widetilde{r}}$-invariance of $\Omega_{\mathsf{NP}}(\phi, \widetilde{r})$.

For (2), let $E \in \mathcal{E}_f$ correspond to r. Let $\tau_L \subset L$ be as in Lemma 15.8. Since $L \in \Omega_{\mathsf{NP}}(r)$, τ_L lifts to a subpath $\widetilde{\tau}_L \subset \widetilde{R}_{\widetilde{E}}$. Lemma 15.8 implies that the lift of τ_L to $\widetilde{\tau}_L$ extends to a lift of L to an element of $\Omega_{\mathsf{NP}}(\phi, \widetilde{r})$. □

Lemma 15.9 *Continue with Notation 15.4.*

(1) *For all $i \geq 1$, there exists $j = j(i) > i$ such that $\widetilde{f}_\#(\widetilde{\ell}_i) = \widetilde{\ell}_j$. More precisely, there exists $j > i$ such that $\widetilde{f}_\#(\widetilde{\rho}_i) \subset \widetilde{\rho}_j$ and $\widetilde{f}_\#(\widetilde{\ell}_i) = \widetilde{\ell}_j$ and there is a covering translation T such that $T(\widetilde{\rho}_i) = \widetilde{f}_\#(\widetilde{\rho}_i) \subset \widetilde{\rho}_j$ and $T(\widetilde{\ell}_i) = \widetilde{\ell}_j$.*

(2) *The assignment $i \mapsto j(i)$ is order-preserving and $j(1) > p$.*

Proof It suffices to prove (2) and the "more precisely" statement of (1). We begin with the latter. Since $\widetilde{f}_\#(\widetilde{\rho}_i)$ is a Nielsen path that is a concatenation of terms in the complete splitting of $\widetilde{R}_{\widetilde{E}}$, there exists j such that $\widetilde{f}_\#(\widetilde{\rho}_i) \subset \widetilde{\rho}_j$. Let T be the unique covering translation satisfying $T(\widetilde{\rho}_i) = \widetilde{f}_\#(\widetilde{\rho}_i)$. It suffices to prove that $\widetilde{f}_\#(\widetilde{\ell}_i) = \widetilde{\ell}_j$ and $T(\widetilde{\ell}_i) = \widetilde{\ell}_j$.

From $\widetilde{f}_\#(\widetilde{\rho}_i) \subset \widetilde{\rho}_j$ it follows that

$$\widetilde{\rho}_j = \widetilde{\alpha} \cdot \widetilde{f}_\#(\widetilde{\rho}_i) \cdot \widetilde{\beta},$$

where $\widetilde{\alpha}$ and $\widetilde{\beta}$ are, possibly trivial, Nielsen paths. Since $\widetilde{\sigma}_i$ and $\widetilde{\sigma}_{i+1}$ are growing,

$$\widetilde{\sigma}_j \cdot \widetilde{\rho}_j \cdot \widetilde{\sigma}_{j+1} = \widetilde{\sigma}_j \cdot \widetilde{\alpha} \cdot \widetilde{f}_\#(\widetilde{\rho}_i) \cdot \widetilde{\beta} \cdot \widetilde{\sigma}_{j+1} \subset \widetilde{f}_\#(\widetilde{\sigma}_i) \cdot \widetilde{f}_\#(\widetilde{\rho}_i) \cdot \widetilde{f}_\#(\widetilde{\sigma}_{i+1}).$$

By Lemma 5.14(2),

$$R_{i+1}^+ = f_\#^\infty(\sigma_{i+1}) = \beta f_\#^\infty(\sigma_{j+1}) = \beta \cdot R_{j+1}^+ \quad \text{and} \quad R_i^- = f_\#^\infty(\overline{\sigma}_i) = \overline{\alpha} f_\#^\infty(\overline{\sigma}_j) = \overline{\alpha} \cdot R_j^-,$$

which implies that

$$\widetilde{f}_\#(\widetilde{\ell}_i) = \widetilde{f}_\#((\widetilde{R}_i^-)^{-1} \cdot \widetilde{\rho}_i \cdot \widetilde{R}_{i+1}^+) = (\widetilde{R}_j^-)^{-1} \cdot \widetilde{\alpha} \cdot \widetilde{f}_\#(\widetilde{\rho}_i) \cdot \widetilde{\beta} \cdot \widetilde{R}_{j+1}^+ = (\widetilde{R}_j^-)^{-1} \cdot \widetilde{\rho}_j \cdot \widetilde{R}_{j+1}^+ = \widetilde{\ell}_j.$$

This completes the proof that $\widetilde{f}_\#(\widetilde{\ell}_i) = \widetilde{\ell}_j$.

The covering translation T that carries $\widetilde{\rho}_i$ to $\widetilde{f}_\#(\widetilde{\rho}_i) \subset \widetilde{\rho}_j$ also carries $\widetilde{R}^+_{i+1}$ to $\widetilde{\beta} \cdot \widetilde{R}^+_{j+1}$ and R^-_i to $\widetilde{\alpha}^{-1} \cdot \widetilde{R}^-_j$. Thus $T(\widetilde{\ell}_i) = \widetilde{\ell}_j$.

Finally, note that $j(i+1) - j(i)$ is equal to the number of growing terms in $\widetilde{f}_\#(\widetilde{\sigma}_i)$. This implies (2) and hence also $j(i) > j$. Since $\widetilde{E} \cdot \widetilde{f}_\#(\widetilde{u}) \cdot \widetilde{f}_\#(\widetilde{\rho}_0) \cdot f_\#(\widetilde{\sigma}_1) \cdot \widetilde{f}_\#(\widetilde{\rho}_1)$ is an initial segment of $\widetilde{R}_E$ and since $\widetilde{f}_\#(\widetilde{\rho}_0) \subset \widetilde{\rho}_p$, it follows that $j(1) > p$. □

We conclude this subsection by defining a total order on $\Omega_{\mathsf{NP}}(\phi, \widetilde{r})$.

Definition 15.10 Continue with Notation 15.4. Given distinct $\widetilde{L}_1, \widetilde{L}_2 \in \Omega_{\mathsf{NP}}(\phi, \widetilde{r})$, choose $k \geq 0$ so that $\widetilde{f}^k_\#(\widetilde{L}_1) = \widetilde{\ell}_i$ and $\widetilde{f}^k_\#(\widetilde{L}_2) = \widetilde{\ell}_j$ for some $i \neq j$. (The existence of k is guaranteed by Proposition 15.6.) Define $\widetilde{L}_1 \prec \widetilde{L}_2$ if $i < j$.

Lemma 15.11 *The relation $\prec$ is a well-defined, $\Phi_{\widetilde{r}}$-invariant total order on $\Omega_{\mathsf{NP}}(\phi, \widetilde{r})$ that is independent of the choice of $f: G \to G$ representing ϕ. Moreover, if $\psi = \theta\phi\theta^{-1}$ for some $\theta \in \mathsf{Out}(F_n)$ and if Θ is a lift of θ, then $\Theta: \Omega_{\mathsf{NP}}(\phi, \widetilde{r}) \to \Omega_{\mathsf{NP}}(\psi, \Theta(\widetilde{r}))$ preserves $\prec$.*

We delay the proof of Lemma 15.11 to state and prove a technical lemma that allows us to redefine $\prec$ with less dependence on the location of $\widetilde{\rho}_i$ and $\widetilde{\rho}_j$ in $\widetilde{R}_{\widetilde{E}}$.

Lemma 15.12 *Continue with Notation 15.4. Suppose that $\widetilde{\ell}_i$ and $\widetilde{\ell}_j$ are distinct nonperiodic visible lines. For all $k \geq 0$, let $\widetilde{\ell}_{i_k} = \widetilde{f}^k_\#(\widetilde{\ell}_i)$ and $\widetilde{\ell}_{j_k} = \widetilde{f}^k_\#(\widetilde{\ell}_j)$, and let $\widetilde{y}_{i,k}$ and $\widetilde{y}_{j,k}$ be the terminal endpoints of $\widetilde{\ell}_{i_k} \cap \widetilde{R}_{\widetilde{E}}$ and $\widetilde{\ell}_{j_k} \cap \widetilde{R}_{\widetilde{E}}$, respectively. Then $i < j$ if and only if the following two conditions are satisfied*:

(1) $\widetilde{\ell}_i \notin \Omega_{\mathsf{NP}}(\phi, \partial_+\widetilde{\ell}_j)$.

(2) *One of the following is satisfied*:

  (a) $\widetilde{\ell}_j \in \Omega_{\mathsf{NP}}(\phi, \partial_+\widetilde{\ell}_i)$.

  (b) $y_{j,k} - y_{i,k} \to \infty$, *where $y_{j,k} - y_{i,k} = \pm$ the number of edges in the subpath connecting $\widetilde{y}_{i,k}$ to $\widetilde{y}_{j,k}$, and the sign is $+$ if and only if $\widetilde{y}_{,jk} > \widetilde{y}_{i,k}$ in the orientation on $\widetilde{R}_{\widetilde{E}}$.*

Proof For the "only if" direction, assume that $i < j$. Lemma 15.9, and an obvious induction argument imply that $i_k < j_k$ and that there are unique covering translations T_k satisfying

$$T_k(\widetilde{\rho}_i) \subset \widetilde{\rho}_{i_k} \quad \text{and} \quad T_k(\widetilde{\ell}_i) = \widetilde{f}^k_\#(\widetilde{\ell}_i) = \widetilde{\ell}_{i_k},$$

and S_k satisfying

$$S_k(\widetilde{\rho}_j) \subset \widetilde{\rho}_{j_k} \quad \text{and} \quad S_k(\widetilde{\ell}_j) = \widetilde{f}^k_\#(\widetilde{\ell}_j) = \widetilde{\ell}_{j_k}.$$

Note that $\widetilde{h}_k := S_k^{-1}\widetilde{f}^k$ is the lift of f^k that preserves $\widetilde{\ell}_j$ and so corresponds to the automorphism $\Phi^k_{\partial_+\widetilde{\ell}_j}$. Note also that $S_k^{-1}T_k(\widetilde{\ell}_i) = S_k^{-1}\widetilde{f}^k_\#(\widetilde{\ell}_i) = (\widetilde{h}_k)_\#(\widetilde{\ell}_i)$, and $T_k(\widetilde{\rho}_i) \subset \widetilde{\rho}_{i_k}$ is disjoint from $\widetilde{\rho}_{j_k}\widetilde{R}^+_{j_k+1}$. The latter implies that $S_k^{-1}T_k(\widetilde{\rho}_i)$ is disjoint from $\widetilde{\rho}_j\widetilde{R}^+_{j+1}$. It now follows from the definition that $\widetilde{\ell}_i \notin \Omega_{\mathsf{NP}}(\phi, \partial_+\widetilde{\ell}_j)$. This completes the proof of (1).

For (2), we assume that $\widetilde{\ell}_j \notin \Omega_{\mathsf{NP}}(\phi, \partial_+\widetilde{\ell}_i)$ and prove that $\widetilde{y}_{j,k} - \widetilde{y}_{i,k} \to \infty$. Continuing with the above notation, $\widetilde{g}_k := T_k^{-1} \widetilde{f}^k$ corresponds to $\Phi^k_{\partial_+\widetilde{\ell}_i}$ and $T_k^{-1} S_k(\widetilde{\ell}_j) = \widetilde{g}_k(\widetilde{\ell}_j)$. We claim that there is a finite subpath $\widetilde{\beta} \subset \widetilde{\ell}_j$ so that for all $k \geq 0$, $T_k^{-1} S_k(\widetilde{\beta}) \not\subset \widetilde{R}^+_{i+1}$ and hence $S_k(\widetilde{\beta}) \not\subset T_k(\widetilde{R}^+_{i+1})$. If $\ell_j \notin \Omega_{\mathsf{NP}}(\partial_+\ell_i)$ then this follows from Definition 5.1 and the fact that $T_k^{-1} S_k(\widetilde{\ell}_j)$ is a lift of ℓ_j. If $\ell_j \in \Omega_{\mathsf{NP}}(\partial_+\ell_i)$ then this follows from Lemmas 15.8 and 15.5.

On the other hand, $S_k(\widetilde{\beta}) \subset \widetilde{R}_{\widetilde{E}}$ for all sufficiently large k. It follows that $\widetilde{y}_{i,k}$ precedes the terminal endpoint of $S_k(\widetilde{\beta})$ in $\widetilde{R}_{\widetilde{E}}$. Since the number of edges in $S_k(\widetilde{R}^+_{j+1}) \cap \widetilde{R}_{\widetilde{E}}$ goes to infinity with k, $\widetilde{y}_{j,k} - \widetilde{y}_{i,k} \to \infty$.

For the "if" direction, we assume that $j < i$ and we prove that either (1) or (2) fails. From the "only if" direction we know that (1) with i and j reversed is satisfied. Thus (2a) fails. Similarly, either (2a) or (2b), with the roles of i and j reversed, is satisfied. If the former holds then (1) fails and we are done. Suppose then that (2b) with the roles of i and j reversed is satisfied. Then $\widetilde{y}_{i,k} - \widetilde{y}_{j,k} \to \infty$ so (2b) fails. □

Proof of Lemma 15.11 To make the dependence of $\prec$ on $f\colon G \to G$ explicit we will write $\prec_f$. Lemma 15.9 implies that $\prec_f$ is a well-defined, $\Phi_{\widetilde{r}}$-invariant total order on $\Omega_{\mathsf{NP}}(\phi, \widetilde{r})$.

Suppose that θ, ψ and Θ are as in the "moreover" statement, that $f'\colon G' \to G'$ is a CT representing ψ, that $g\colon G \to G'$ is a homotopy equivalence representing θ and that $\widetilde{g}\colon \widetilde{G} \to \widetilde{G}'$ is the lift corresponding to Θ. Letting $\widetilde{r}' = \Theta(\widetilde{r})$, we have $\Theta\Phi_{\widetilde{r}} = \Psi_{\widetilde{r}'}\Theta$. By Lemma 15.3, $\Theta(\Omega_{\mathsf{NP}}(\phi, \widetilde{r})) = \Omega_{\mathsf{NP}}(\psi, \widetilde{r}')$. Let $\widetilde{f}'$ be the lift of f' corresponding to $\Psi_{\widetilde{r}'}$.

Given $\widetilde{L}_1, \widetilde{L}_2 \in \Omega_{\mathsf{NP}}(\phi, \widetilde{r})$ such that $\widetilde{L}_1 \prec_f \widetilde{L}_2$, we must show that $\widetilde{L}'_1 \prec_{f'} \widetilde{L}'_2$, where $\widetilde{L}'_1 = \Theta(\widetilde{L}_1) = \widetilde{g}_\#(\widetilde{L}_1)$ and $\widetilde{L}'_2 = \Theta(\widetilde{L}_2) = \widetilde{g}_\#(\widetilde{L}_2)$. We may replace $\widetilde{L}_1$ and $\widetilde{L}_2$ with $\Phi^k_{\widetilde{r}}\widetilde{L}_1$ and $\Phi^k_{\widetilde{r}}\widetilde{L}_2$ for any $k \geq 1$. This follows from the $\Phi_{\widetilde{r}}$-invariance of $\prec_f$, the $\Psi_{\widetilde{r}'}$-invariance of $\prec_{f'}$ and the fact that $\Theta\Phi_{\widetilde{r}} = \Psi_{\widetilde{r}'}\Theta$. In particular, we may assume that there exists $i < j$ and $i' \neq j'$ such that $\widetilde{L}_1 = \widetilde{\ell}_i$, $\widetilde{L}_2 = \widetilde{\ell}_j$, $\widetilde{L}'_1 = \widetilde{\ell}'_i$ and $\widetilde{L}'_2 = \widetilde{\ell}'_j$, where the $\widetilde{\ell}'_i$ and $\widetilde{\ell}'_j$ are visible lines determined for $\Omega_{\mathsf{NP}}(\psi, \widetilde{r}')$ defined with respect to f'.

To prove that $i' < j'$, and thereby complete the proof of the lemma, we will verify items (1) and (2) of Lemma 15.12 in the prime system, which we will call (1)$'$ and (2)$'$. Items (1)$'$ and (2a)$'$ follow from (1), (2a) and Lemma 15.3. Item (2b)$'$ follows from (2b) and the bounded cancellation lemma applied to g. □

We conclude this section with a result that will be used in Lemma 15.45.

Lemma 15.13 *Suppose that F is a free factor, that $\widetilde{r} \in \partial F$ is a lift of $r \in \mathcal{R}(\phi)$ and that $[F]$ is ϕ-invariant. Then each endpoint of each $\widetilde{\ell} \in \Omega_{\mathsf{NP}}(\phi, \widetilde{r})$ is contained in ∂F.*

Proof Choose a CT $f\colon G \to G$ representing ϕ in which F is realized by a component C of a core filtration element H. By assumption, there is a ray R in C with terminal end r. For each $\ell \in \Omega_{\mathsf{NP}}(r)$, each finite subpath of ℓ is contained in C and hence ℓ is contained in C. Let $\widetilde{C}$ be the unique lift of C whose boundary contains $\widetilde{r}$ and note that $\partial\widetilde{C} = \partial F$. Let $\widetilde{R} \subset \widetilde{C}$ be the lift of R with terminal endpoint $\widetilde{r}$

and let $\Phi_{\widetilde{r}}$ be the automorphism representing ϕ that fixes $\widetilde{r}$. From uniqueness of $\widetilde{C}$, it follows that $\partial\widetilde{C}$ is $\Phi_{\widetilde{r}}$-invariant. For all sufficiently large j, $\Phi_{\widetilde{r}}^j(\widetilde{\ell}) \cap \widetilde{R} \neq \varnothing$. Since $\ell \subset C$ and distinct lifts of C are disjoint, $\Phi_{\widetilde{r}}^j(\widetilde{\ell}) \subset \widetilde{C}$. It follows that the endpoints of $\Phi_{\widetilde{r}}^j(\widetilde{\ell})$, and hence the endpoints of $\widetilde{\ell}$, are contained in ∂F. □

15.2 Topmost lines, translation numbers and offset numbers

We continue with Notation 15.4 and with the partial orders $<$ on $\mathcal{R}(\phi)$ and $\mathcal{E}_f$ given in Notation 6.1 and Lemma 6.2.

Definition 15.14 An element $L \in \Omega_{\mathsf{NP}}(r)$ is *ϕ-topmost* if one of the following mutually exclusive properties is satisfied for the partial order $<$ on $\mathcal{R}(\phi)$:

(1) r is minimal in the partial order $<$.

(2) L has an end $r_1 \in \mathcal{R}(\phi)$ such that $r_1 <_c r$.

If $\widetilde{L} \in \Omega_{\mathsf{NP}}(\phi, \widetilde{r})$ projects to a ϕ-topmost element of $\Omega_{\mathsf{NP}}(r)$, then $\widetilde{L}$ is a *topmost* element of $\Omega_{\mathsf{NP}}(\phi, \widetilde{r})$. Let $\mathcal{T}_{\phi,\widetilde{r}}$ be the *set of topmost elements of* $\Omega_{\mathsf{NP}}(\phi, \widetilde{r})$.

Lemma 15.15 *$\mathcal{T}_{\phi,\widetilde{r}}$ is nonempty and $\Phi_{\widetilde{r}}$-invariant.*

Proof Lemma 15.3 implies that $\Omega_{\mathsf{NP}}(\phi, \widetilde{r})$ is $\Phi_{\widetilde{r}}$-invariant. Since each element of $\mathcal{R}(\phi)$ is ϕ-invariant, $\Phi_{\widetilde{r}}$-invariance of $\mathcal{T}_{\phi,\widetilde{r}}$ follows from the definitions. If r is minimal with respect to $<$ then every element of $\Omega_{\mathsf{NP}}(\phi, \widetilde{r})$ is topmost and we are done. Otherwise, apply Lemma 6.2 to choose $E' \in \mathcal{E}_f$ such that $E' <_c E$. Either E' or $\overline{E}'$ occurs as a term σ_j in the coarsened complete splitting of R_E. In the former case, $\widetilde{\ell}_{j-1}$ is topmost in $\Omega_{\mathsf{NP}}(\phi, \widetilde{r})$; in the latter case $\widetilde{\ell}_j$ is topmost in $\Omega_{\mathsf{NP}}(\phi, \widetilde{r})$. □

Lemma 15.16 *There is an algorithm that lists the ϕ-topmost elements of $\Omega_{\mathsf{NP}}(r)$.*

Proof Recall that the elements of $\Omega_{\mathsf{NP}}(r)$ can be enumerated by Corollary 5.17(2) and that the partial order on $\Omega_{\mathsf{NP}}(r)$ can be computed by Notation 6.1. If r is minimal then every element of $\Omega_{\mathsf{NP}}(r)$ is topmost. Otherwise inspect the elements of $\Omega_{\mathsf{NP}}(r)$ to see which ones satisfy 15.14(2). □

Recall from Notation 15.4 that p is chosen so that $\widetilde{f}_\#(\rho_0) \subset \widetilde{\rho}_p$.

Lemma 15.17 *Each $\widetilde{L} \in \mathcal{T}_{\phi,\widetilde{r}}$ is in the $\Phi_{\widetilde{r}}$-orbit of $\widetilde{\ell}_j$ for some $1 \le j \le p$.*

Proof By Proposition 15.6 and Lemma 15.15, we may assume that $\widetilde{L} = \widetilde{\ell}_i$ for some $i > p$. By Lemma 15.9, it suffices to show that there exist $1 \le j \le p$ and $k \ge 1$ such that $\widetilde{f}_\#^k(\widetilde{\rho}_j) \subset \widetilde{\rho}_i$. If this fails then there exists $1 \le j' \le p$ and $k' \ge 1$ such that $\widetilde{\rho}_i$ separates $\widetilde{f}_\#^{k'}(\widetilde{\rho}_{j'-1})$ from $\widetilde{f}_\#^{k'}(\widetilde{\rho}_{j'})$. Assuming this we argue towards a contradiction by showing that neither (1) nor (2) in Definition 15.14 is satisfied. First note that $\widetilde{\sigma}_i\widetilde{\rho}_i\widetilde{\sigma}_{i+1} \subset \widetilde{f}_\#^{k'}(\widetilde{\sigma}_{j'})$. It follows that $\sigma_{j'}$ is not a linear term and so $\sigma_{j'} = E'$ or $\overline{E}'$ for some $E' \in \mathcal{E}_f$. Since $E' < E$, (1) is not satisfied. If an end r'' of $\widetilde{\ell}_i$ corresponds to an element $E'' \in \mathcal{E}_f$ then $E'' < E' < E$ and so (2) is not satisfied. □

Notation 15.18 The total order $\prec$ on $\Omega_{\mathsf{NP}}(\phi,\tilde{r})$ given in Definition 15.10 induces a total order (also called) $\prec$ on $\mathcal{T}_{\phi,\tilde{r}}$. Let $\tilde{L}_1,\dots,\tilde{L}_{\tau(\phi,\tilde{r})}$ be, in order, the elements of $\{\tilde{\ell}_1,\dots,\tilde{\ell}_p\}\cap\mathcal{T}_{\phi,\tilde{r}}$. For $k\in\mathbb{Z}$ and $1\le j\le\tau(\phi,\tilde{r})$, define $\tilde{L}_{j+k\tau(\phi,\tilde{r})}=\Phi^k_{\tilde{r}}(\tilde{L}_j)$.

The following lemma allows us to change our notation from $\tau(\phi,\tilde{r})$ to $\tau(\phi,r)$.

Lemma 15.19 *The number $\tau(\phi,\tilde{r})$ depends only on ϕ and r and not on the choice of $\tilde{r}$.*

Proof The definition of $\tau(\phi,\tilde{r})$ uses the lift $\tilde{f}\colon\tilde{G}\to\tilde{G}$ corresponding to $\Phi_{\tilde{r}}$, the lift $\tilde{R}_{\tilde{E}}$ of R_E whose terminal endpoint is $\tilde{r}$, the lines $\{\tilde{\ell}_i\}$ determined by $\tilde{R}_{\tilde{E}}$ as described in Notation 15.4 and the integer p, which depends only on E and f. If $a\in F_n$ and $T_a\colon\tilde{G}\to\tilde{G}$ is the corresponding covering translation, then the data associated to $\tilde{r}'=a\tilde{r}$ is $\tilde{f}'=T_a\tilde{f}T_a^{-1}$, $\tilde{R}_{\tilde{E}'}=T_aR_{\tilde{E}}$, $\tilde{\ell}'_i=T_a\tilde{\ell}_i$ and p. Since $\tilde{\ell}_i$ and $\tilde{\ell}'_i$ are lifts of the same line, $\tilde{\ell}_i\in\mathcal{T}_{\phi,\tilde{r}}\iff\tilde{\ell}'_i\in\mathcal{T}_{\phi,\tilde{r}'}$. This proves that $\tau(\phi,\tilde{r})=\tau(\phi,\tilde{r}')$, as desired. □

Lemma 15.20 *With notation as above,*

1. *$s\mapsto\tilde{L}_s$ defines an order-preserving bijection between $\mathbb{Z}$ and $\mathcal{T}_{\phi,\tilde{r}}$,*
2. *$\Phi_{\tilde{r}}(\tilde{L}_s)=\tilde{L}_{s+\tau(\phi,r)}$ for all s, and*
3. *$\tilde{L}_s$ is visible if and only if $s\ge 1$.*

Proof The map $s\mapsto\tilde{L}_s$ is surjective by Lemma 15.17 and is order-preserving (and hence injective) because $\tilde{f}_\#$ preserves $\prec$ and because $\tilde{L}_1\prec\tilde{L}_2\prec\cdots\prec\tilde{L}_{\tau(\phi,r)}\prec\tilde{f}_\#(\tilde{L}_1)$, where the last inequality follows Lemma 15.9, which implies that $\tilde{f}_\#(\tilde{L}_1)=\tilde{\ell}_j$ for some $j>p$. Item (2) follows from the definitions. Item (3) follows from Lemma 15.9. □

For the next lemma, we must choose a CT $f'\colon G'\to G'$ representing ψ and then define $\mathcal{T}_{\psi,\tilde{r}'}$ and $\tau(\psi,r')$ with respect to $f'\colon G'\to G'$.

Lemma 15.21 *Suppose that $\theta\in\mathsf{Out}(F_n)$ conjugates ϕ to ψ, that $\theta(r)=r'\in\mathcal{R}(\psi)$, that $\tilde{r},\tilde{r}'\in\partial F_n$ represent r and r', respectively, and that Θ is the lift of θ such that $\Theta(\tilde{r})=\tilde{r}'$. Then*

1. *$\tau(\phi,r)=\tau(\psi,r')$, and*
2. *there is an integer $\mathsf{offset}(\theta,r)$ such that $\Theta(\tilde{L}_s)=\tilde{L}'_{s+\mathsf{offset}(\theta,r)}$ for all s.*

Proof Lemmas 15.3 and 15.11 imply that Θ induces a $\prec$-preserving bijection between $\Omega_{\mathsf{NP}}(\phi,\tilde{r})$ and $\Omega_{\mathsf{NP}}(\psi,\tilde{r}')$. Lemma 6.3 implies that this bijection restricts to a bijection between $\mathcal{T}_{\phi,\tilde{r}}$ and $\mathcal{T}_{\psi,\tilde{r}'}$. Since the only order-preserving bijections of $\mathbb{Z}$ are translations, there is an integer $\mathsf{offset}(\theta,\tilde{r},\tilde{r}')$ such that $\Theta(\tilde{L}_s)=\tilde{L}'_{s+\mathsf{offset}(\theta,\tilde{r},\tilde{r}')}$ for all s. If we replace $\tilde{r}$ by another lift $\tilde{r}^*=i_a\tilde{r}$ then Θ is replaced by $\Theta^*=\Theta i_a^{-1}$ and $\tilde{L}_i\in\mathcal{T}_{\phi,\tilde{r}}$ is replaced by $\tilde{L}^*_i=i_a\tilde{L}_i\in\mathcal{T}_{\phi,\tilde{r}^*}$; see the proof of Lemma 15.19. It follows that $\Theta^*(L^*_i)=\Theta(\tilde{L}_i)$ and hence that $\mathsf{offset}(\theta,\tilde{r},\tilde{r}')$ is independent of the choice of lift $\tilde{r}$. The symmetric argument implies that $\mathsf{offset}(\theta,\tilde{r},\tilde{r}')$ is also independent of the choice of $\tilde{r}'$. This completes the proof of (2).

Item (1) therefore follows from

$$\widetilde{L}'_{s+\tau(\phi,r)+\mathsf{offset}(\theta,r)} = \Theta\Phi_{\widetilde{r}}\widetilde{L}_s = \Psi_{\widetilde{r}'}\Theta\widetilde{L}_s = \widetilde{L}'_{s+\mathsf{offset}(\theta,r)+\tau'(\psi,r')}. \qquad \square$$

Remark 15.22 The bijection between $\mathbb{Z}$ and $\mathcal{T}_{\phi,\widetilde{r}}$ depends on the notion of visible lines and so depends on the choice of CT. On the other hand, Lemma 15.20(2) implies that $\tau(\phi,r)$ depends only on ϕ and r and not on the choice of a CT. As such it can be computed from any CT for ϕ. The integer $\mathsf{offset}(\theta,r)$ depends on the choices of CTs.

15.3 Staple pairs

We continue with Notation 15.4. We set further notation as follows.

Notation 15.23 If E_i and E_j are distinct elements of $\mathsf{Lin}_w(f)$ then there exist nonzero $d_i \neq d_j$ such that $f(E_i) = E_i w^{d_i}$ and $f(E_j) = E_j w^{d_j}$. Recall that a path of the form $E_i w^p \overline{E}_j$ is called *exceptional* if d_i and d_j have the same sign. If d_i and d_j have different signs then we say $E_i w^p \overline{E}_j$ is *quasi-exceptional*.

Notation 15.24 We write $L \in \mathcal{S}(\phi)$ and say that L is a *staple* if $L \in \Omega_{\mathsf{NP}}(\phi)$ has at least one periodic end; if both ends of L are periodic then L is a *linear staple*. If $\widetilde{L} \in \Omega_{\mathsf{NP}}(\phi,\widetilde{r})$ projects to an element of $\mathcal{S}(\phi)$ for $r \in \mathcal{R}(\phi)$ and lift $\widetilde{r}$, then we write $\widetilde{L} \in \mathcal{S}(\phi,\widetilde{r})$ and $L \in \mathcal{S}(\phi,r)$ and we say that L and $\widetilde{L}$ *occur* in r and $\widetilde{r}$, respectively.

For each $r \in \mathcal{R}(\phi)$, an ordered pair $b = (L_1, L_2)$ of elements of $\mathcal{S}(\phi,r)$ is a *staple pair* if there are lifts $\widetilde{L}_1, \widetilde{L}_2 \in \mathcal{S}(\phi,\widetilde{r})$ and a periodic line $\widetilde{A}$ such that $\{\partial_+\widetilde{L}_1, \partial_-\widetilde{L}_2\} \subset \{\partial_-\widetilde{A}, \partial_+\widetilde{A}\}$. We write $b \in \mathcal{S}_2(\phi,r)$ and $\widetilde{b} = (\widetilde{L}_1, \widetilde{L}_2) \in \mathcal{S}_2(\phi,\widetilde{r})$ and say that b and $\widetilde{b}$ *occur* in r and $\widetilde{r}$ respectively and that $\widetilde{A}$ is the *common axis* of $\widetilde{b}$. By Corollary 5.17, $\widetilde{A}$ corresponds to an element of $\mathcal{A}(\phi)$. Define $\mathcal{S}_2(\phi) = \cup\mathcal{S}_2(\phi,r)$, where the union is taken over all $r \in \mathcal{R}(\phi)$.

Lemma 15.25 *Each $b \in \mathcal{S}_2(\phi,r)$ is ϕ-invariant. The set $\mathcal{S}_2(\phi,\widetilde{r})$ is $\Phi_{\widetilde{r}}$-invariant.*

Proof The first statement follows from the second and the fact (Corollary 5.17(1)) that each element of $\Omega_{\mathsf{NP}}(\phi)$ is ϕ-invariant. The second follows from the $\Phi_{\widetilde{r}}$-invariance of $\Omega_{\mathsf{NP}}(\phi,\widetilde{r})$ (Lemma 15.3) and the definition of $\mathcal{S}_2(\phi,\widetilde{r})$. $\square$

Example 3.1 (continued) In our example,

$$\mathcal{S}(\phi) = \{a^\infty R_c, a^\infty b a^\infty\} \quad \text{and} \quad \mathcal{S}_2(\phi) = \{(a^\infty b a^\infty, a^\infty b a^\infty), (a^\infty b a^\infty, a^\infty R_c)\}.$$

Throughout this section, M is the stabilization constant defined in Notation 5.13.

Our next lemma explains how staple pairs occur in an eigenray $\widetilde{R}_{\widetilde{E}}$.

Lemma 15.26 *Assume Notation 15.4.*

(1) *If* $\sigma_i\rho_i\sigma_{i+1}$ *is quasi-exceptional, then* $(\widetilde{\ell}_{i-1},\widetilde{\ell}_{i+1})\in\mathcal{S}_2(\phi,\widetilde{r})$ *with common axis* $\widetilde{\ell}_i$.

(2) *If one of the following holds:*
  (a) σ_i *is exceptional, or*
  (b) $\sigma_i\in\mathsf{Lin}_w(f)$ *and* ℓ_i *is not periodic, or*
  (c) $\overline{\sigma}_i\in\mathsf{Lin}_w(f)$ *and* ℓ_{i-1} *is not periodic;*
  then $(\widetilde{\ell}_{i-1},\widetilde{\ell}_i)\in\mathcal{S}_2(\phi,\widetilde{r})$.

(3) *If* $\widetilde{\ell}_i$ *is periodic and neither* σ_i *nor* $\overline{\sigma}_{i+1}$ *is in* $\mathcal{E}_f$, *then* $\sigma_i\rho_i\sigma_{i+1}$ *is quasi-exceptional and so* $(\widetilde{\ell}_{i-1},\widetilde{\ell}_{i+1})\in\mathcal{S}_2(\phi,\widetilde{r})$ *with common axis* $\widetilde{\ell}_i$. *See also Remark 15.43.*

(4) *For each* $\widetilde{b}\in\mathcal{S}_2(\phi,\widetilde{r})$, *there exists* $K=K(\widetilde{b})$ *such that* $\Phi^k_{\widetilde{r}}(\widetilde{b})$ *is as in* (1) *or* (2) *for all* $k\geq K$. *Moreover, in case* (b), $R_i^-=w^{\pm\infty}$ *and in case* (c), $R_{i-1}^+=w^{\pm\infty}$.

Proof If $\sigma_i\rho_i\sigma_{i+1}$ is quasi-exceptional then there are a twist curve w and edges $E',E''\in\mathsf{Lin}_w(f)$ such that $\sigma_i=E'$, $\rho_i=w^q$ for some $q\in\mathbb{Z}$ and $\sigma_{i+1}=\overline{E}''$. Moreover, $f(E')=E'w^{d'}$ and $f(E'')=E''w^{d''}$, where d' and d'' have opposite signs. If $\widetilde{\sigma}_i=\widetilde{E}'$, let $\widetilde{w}$ be the lift of w that begins with the terminal endpoint $\widetilde{x}$ of $\widetilde{E}'$. Extend $\widetilde{w}$ to a periodic line $\widetilde{A}$ that projects bi-infinitely to the circuit determined by w and is oriented consistently with $\widetilde{w}$. Let $\widetilde{y}\in\widetilde{A}$ be the terminal endpoint of the lift of w^q that begins at $\widetilde{x}$ and let $\widetilde{E}''$ be the lift of E'' that ends at $\widetilde{y}$. Then $\widetilde{R}_i^+$ is the concatenation of $\widetilde{E}'$ and a ray in $\widetilde{A}$ beginning at $\widetilde{x}$ and terminating at $\partial_+\widetilde{A}$ if $d'>0$ and at $\partial_-\widetilde{A}$ if $d'<0$. Similarly, $\widetilde{R}_{i+1}^-$ is the concatenation of $\widetilde{E}''$ and a ray in $\widetilde{A}$ beginning at $\widetilde{y}$ and terminating at $\partial_+\widetilde{A}$ if $d''>0$ and at $\partial_-\widetilde{A}$ if $d''<0$. Neither $\widetilde{\ell}_{i-1}$ nor $\widetilde{\ell}_{i+1}$ is periodic. Up to a change of orientation, $\widetilde{\ell}_i=\widetilde{A}$. Thus $(\widetilde{\ell}_{i-1},\widetilde{\ell}_{i+1})\in\mathcal{S}(\phi,\widetilde{r})$ with common axis $\widetilde{\ell}_i$, and (1) is proved.

If σ_i is exceptional, then $\sigma_i=E'w^q\overline{E}''$, where E',w and E'' are as above except that d' and d'' have the same sign. Following the above notation, $\widetilde{R}_i^+$ begins with $\widetilde{E}'$, $\widetilde{R}_i^-$ begins with $\widetilde{E}''$ and both rays terminate at the same endpoint of $\widetilde{A}$. Neither $\widetilde{\ell}_{i-1}$ nor $\widetilde{\ell}_i$ is periodic. This completes the proof of (2a).

If $\sigma_i=E'\in\mathsf{Lin}_w(f)$, then following the above notation, $\widetilde{\ell}_{i-1}$ is nonperiodic (because it crosses $\widetilde{E}'$) with terminal endpoint in $\{\partial_-\widetilde{A},\partial_+\widetilde{A}\}$ and $\widetilde{R}_i^-$ has terminal endpoint in $\{\partial_-\widetilde{A},\partial_+\widetilde{A}\}$. If $\widetilde{\ell}_i$ is nonperiodic then $(\widetilde{\ell}_{i-1},\widetilde{\ell}_i)\in\mathcal{S}_2(\phi,\widetilde{r})$. This completes the proof of (2b). The proof of (2c) is similar.

Suppose that $\widetilde{\ell}_i$ is as in (3). If $\overline{\sigma}_i\in\mathcal{E}_f$ then $\widetilde{R}_i^-$ is not asymptotic to a periodic line, in contradiction to the assumption that $\widetilde{\ell}_i$ is periodic. If $\overline{\sigma}_i\in\mathsf{Lin}(f)$ or if σ_i is exceptional then $R_i^-=E_iw^{\pm\infty}$, where $E_i\in\mathsf{Lin}_w(f)$, again in contradiction to the assumption that $\widetilde{\ell}_i$ is periodic. We conclude that $\sigma_i=E'\in\mathsf{Lin}(f)$. The symmetric argument shows that $\sigma_{i+1}=\overline{E}''$ for some $E''\in\mathsf{Lin}(f)$. Thus ℓ_i has the form $(w')^{\pm\infty}\rho_i(w'')^{\pm\infty}$, where w' is the twist path for E' and w'' is the twist path for E''. Since $\widetilde{\ell}_i$ is a periodic line, $w'=w''$ and $\rho_i=(w')^q$ for some $q\in\mathbb{Z}$. This proves that $\sigma_i\rho_i\sigma_{i+1}$ is quasi-exceptional, which in conjunction with (1), completes the proof of (3).

For (4), suppose that $\widetilde{b} \in \mathcal{S}_2(\phi, \widetilde{r})$. After replacing $\widetilde{b}$ with some $\Phi_{\widetilde{r}}^k(\widetilde{b})$, we may assume by Proposition 15.6 that $\widetilde{b} = (\widetilde{\ell}_{i-1}, \widetilde{\ell}_j)$ for some $i-1 \neq j$. After replacing $\widetilde{b}$ with $\Phi_{\widetilde{r}}^M(\widetilde{b})$, we may assume that $\overline{\sigma}_i \notin \mathcal{E}_f$. (To see this note that if $\widetilde{f}_\#^M(\widetilde{\rho}_{i-1}) \subset \widetilde{\rho}_{s-1}$ then $\widetilde{f}_\#^M(\widetilde{\ell}_{i-1}) = \widetilde{\ell}_{s-1}$ and $\widetilde{\sigma}_s$ is the first growing term of $\widetilde{f}_\#^M(\widetilde{\sigma}_i)$.) By assumption, $\partial_+ \widetilde{R}_i^+$ is an endpoint of the common axis $\widetilde{A}$ of $\widetilde{b}$. Lemma 5.14 therefore implies that $\sigma_i \notin \mathcal{E}_f$ and hence that σ_i is linear. In other words, $\sigma_i = E_i$ or $\sigma_i = \overline{E}_i$ or $\sigma_i = E_i w_i^* \overline{E}_l$ for some twist path w_i and for some $E_i, E_l \in \mathsf{Lin}_{w_i}(f)$. In all three cases, the terminal endpoint of $\widetilde{E}_i$ is contained in $\widetilde{A}$. For the same reasons, we may assume that $\sigma_j = E_j$ or $\sigma_j = \overline{E}_j$ or $\sigma_j = E_m w_j^* \overline{E}_j$ for some twist path w_j and for some $E_j, E_m \in \mathsf{Lin}_{w_j}(f)$; moreover, the terminal endpoint of $\widetilde{E}_j$ is in $\widetilde{A}$.

The proof now proceeds by a case analysis. If $\sigma_i = E_i w_i^* \overline{E}_l$ then the midpoint of $\widetilde{E}_i$ (resp. $\widetilde{E}_l^{-1}$) separates $\widetilde{A}$ from $\widetilde{\sigma}_q$ for all $q < i$ (resp. $q > i$) so $j = i$ and we are in case (a). The same argument, with the same conclusion, applies if $\sigma_j = E_m w_j^* \overline{E}_j$. We may now assume that σ_i is either E_i or $\overline{E}_i$ and that σ_j is either E_j or $\overline{E}_j$. By considering the midpoints of σ_i and σ_j as in the previous case we see that:

(a) If $\sigma_i = E_i$, then $j \geq i$.

(b) If $\sigma_i = \overline{E}_i$, then $j \leq i$.

(c) If $\sigma_j = E_j$, then $i \geq j$.

(d) If $\sigma_j = \overline{E}_j$, then $i \leq j$.

If (a) and (c) are satisfied then $j = i$, we are in case (b) and $R_i^- = w^{\pm\infty}$. Similarly, if (b) and (d) are satisfied then $j = i$, we are in case (c) and $R_{i-1}^+ = w^{\pm\infty}$. Suppose next that (a) and (d) are satisfied. In this case, $j \geq i$, $w_i = w_j$ and the interval $\widetilde{\tau}$ of $\widetilde{A}$ bounded by the terminal endpoints of $\widetilde{E}_i$ and $\widetilde{E}_j$ equals $\widetilde{w}_i^q$ for some $q \in \mathbb{Z}$; in particular, τ is a Nielsen path. It must be that $\tau = \rho_i$ and $j = i+1$, which is (1). Finally, suppose that (b) and (c) are satisfied. Then $j \leq i$, $w_i = w_j$ and the interval $\widetilde{\tau}$ of $\widetilde{A}$ bounded by the terminal endpoints of $\widetilde{E}_j$ and $\widetilde{E}_i$ equals $\widetilde{w}_i^q$ for some $q \in \mathbb{Z}$; in particular, τ is a Nielsen path. It must be that $\tau = \rho_{i-1}$ and $j = i-1$, which contradicts the fact that $i - 1 \neq j$. Thus this last case does not happen, and we are done. □

Notation 15.27 We say that the staple pairs $(\widetilde{\ell}_{i-1}, \widetilde{\ell}_{i+1})$ and $(\widetilde{\ell}_{i-1}, \widetilde{\ell}_i)$ that occur in items (1) and (2) of Lemma 15.26 are *visible with index* i or just *visible* if the index is not explicitly given. Note that if $\widetilde{b}$ is visible then $\Phi_{\widetilde{r}}^k(\widetilde{b})$ is visible for all $k \geq 0$.

Corollary 15.28 *The set of visible elements of* $\mathcal{S}_2(\phi, \widetilde{r})$ *is infinite.*

Proof From $\Phi_{\widetilde{r}}$-invariance of $\mathcal{S}_2(\phi, \widetilde{r})$ (Lemma 15.25) and Lemma 15.9, we need only show that $\mathcal{S}_2(\phi, \widetilde{r})$ contains a visible element. There are always linear edges crossed by R_E. We are therefore reduced, by Lemma 15.26, to the case that some $\widetilde{\ell}_i$ is periodic. If $\widetilde{f}_\#^M(\widetilde{\rho}_i) \subset \widetilde{\rho}_j$ then $\widetilde{\sigma}_j$ is the last growing term in $\widetilde{f}_\#^M(\widetilde{\sigma}_i)$ and so $\sigma_j \notin \mathcal{E}_f$. Similarly, $\widetilde{\sigma}_{j+1}$ is the first growing term in $\widetilde{f}_\#^M(\widetilde{\sigma}_{i+1})$ and so $\overline{\sigma}_{j+1} \notin \mathcal{E}_f$. Lemma 15.26(3) implies that $(\widetilde{\ell}_{j-1}, \widetilde{\ell}_{j+1}) \in \mathcal{S}_2(\phi, \widetilde{r})$ and we are done. □

Recall from Notation 15.4 that the $\widetilde{\ell}_i$ are said to be *visible*.

Lemma 15.29 *Suppose that $\widetilde{b} = (\widetilde{L}_1, \widetilde{L}_2) \in \mathcal{S}_2(\phi, \widetilde{r})$ with common axis $\widetilde{A}$ and that one of the following two conditions are satisfied.*

(a) *Either $\widetilde{L}_1$ or $\widetilde{L}_2$ is visible and there exist $k \ge 0$ such that $\Phi^k_{\widetilde{r}}(\widetilde{L}_1, \widetilde{L}_2) = (\widetilde{\ell}_{i-1}, \widetilde{\ell}_i)$ for some i.*

(b) *Either $\widetilde{L}_1, \widetilde{L}_2$ or the common axis of $\widetilde{A}$ is visible and there exist $k \ge 0$ such that $\Phi^k_{\widetilde{r}}(\widetilde{L}_1, \widetilde{L}_2) = (\widetilde{\ell}_{i-1}, \widetilde{\ell}_{i+1})$ with common axis $\widetilde{\ell}_i$ for some i.*

Then $\Phi^{2M}_{\widetilde{r}}(\widetilde{b})$ is visible.

Proof We begin by establishing the following properties for each visible line $\widetilde{\ell}_j$.

(1) Suppose that $\partial_+\widetilde{\ell}_j$ is periodic and that $\widetilde{\ell}_s = \Phi^M_{\widetilde{r}}(\widetilde{\ell}_j)$. Then $\Phi^m_{\widetilde{r}}(\widetilde{\ell}_s)$ and $\Phi^m_{\widetilde{r}}(\widetilde{\ell}_{s+1})$ are consecutive (ie their indices differ by 1) for all $m \ge 0$.

(2) Suppose that $\partial_-\widetilde{\ell}_j$ is periodic and that $\widetilde{\ell}_s = \Phi^M_{\widetilde{r}}(\widetilde{\ell}_j)$. Then $\Phi^m_{\widetilde{r}}(\widetilde{\ell}_{s-1})$ and $\Phi^m_{\widetilde{r}}(\widetilde{\ell}_s)$ are consecutive for all $m \ge 0$.

For (1), Lemma 15.9 implies that $\widetilde{f}^M_\#(\widetilde{\rho}_j) \subset \widetilde{\rho}_s$ and our choice of M implies that $\widetilde{\sigma}_{s+1} \notin \mathcal{E}_f^{-1}$. Since $\partial_+\widetilde{\ell}_j$ is periodic, the same is true for $\partial_+\widetilde{\ell}_s$ and so $\widetilde{\sigma}_{s+1} \notin \mathcal{E}_f$. We conclude that σ_{s+1} is linear. In particular, $\widetilde{f}^m_\#(\widetilde{\sigma}_{s+1})$ has exactly one growing term. If $\widetilde{f}^m_\#(\widetilde{\rho}_s) \subset \widetilde{\rho}_a$ then $\widetilde{f}^m_\#(\widetilde{\rho}_{s+1}) \subset \widetilde{\rho}_{a+1}$. Lemma 15.9 implies that $\Phi^m_{\widetilde{r}}(\widetilde{\ell}_s) = \widetilde{\ell}_a$ and $\Phi^m_{\widetilde{r}}(\widetilde{\ell}_{s+1}) = \widetilde{\ell}_{a+1}$. This completes the proof of (1). Item (2) is proved by the symmetric argument.

We now apply (1) and (2) to prove the lemma, assuming without loss that $k > M$. In case (a), we will show that $\Phi^M_{\widetilde{r}}(\widetilde{b})$ is visible. If $\widetilde{L}_1$ is visible let $\widetilde{\ell}_{s-1} = \Phi^M_{\widetilde{r}}(\widetilde{L}_1)$. Since $\Phi^{k-M}_{\widetilde{r}}(\widetilde{\ell}_{s-1}) = \Phi^k_{\widetilde{r}}(\widetilde{L}_1) = \widetilde{\ell}_{i-1}$, property (1), applied with $m = k - M$, implies that $\Phi^{k-M}_{\widetilde{r}}(\widetilde{\ell}_s) = \widetilde{\ell}_i$. Since $\Phi^{k-M}_{\widetilde{r}}(\Phi^M_{\widetilde{r}}(\widetilde{L}_2)) = \widetilde{\ell}_i$, we have $\Phi^M_{\widetilde{r}}(\widetilde{L}_2) = \widetilde{\ell}_s$. Thus $\Phi^M_{\widetilde{r}}(\widetilde{b}) = (\Phi^M_{\widetilde{r}}(\widetilde{L}_1), \Phi^M_{\widetilde{r}}(\widetilde{L}_2)) = (\widetilde{\ell}_{s-1}, \widetilde{\ell}_s)$ is visible. This completes the proof when $\widetilde{L}_1$ is visible. When $\widetilde{L}_2$ is visible, a symmetric argument, using (2) instead of (1), shows that $\Phi^M_{\widetilde{r}}(\widetilde{L}_1)$ and hence $\Phi^M_{\widetilde{r}}(\widetilde{b})$ is visible.

In case (b), note that $\partial_+\widetilde{L}_1, \partial_-\widetilde{L}_2$ and both ends of $\widetilde{A}$ are periodic. If $\widetilde{L}_1$ is visible then the above argument shows that the common axis of $\Phi^M_{\widetilde{r}}(\widetilde{b})$ is visible and a second application shows that $\Phi^{2M}_{\widetilde{r}}(\widetilde{L}_2)$ is visible. The other cases are similar. □

Notation 15.30 Suppose that $\widetilde{b} = (\widetilde{L}_1, \widetilde{L}_2) \in \mathcal{S}_2(\phi, \widetilde{r})$ projects to $b \in \mathcal{S}_2(\phi, r)$. If $\widetilde{\ell}_j \prec \widetilde{L}_1$ (see Definition 15.10) then we write $\widetilde{\ell}_j \prec \widetilde{b}$. We say that b and $\widetilde{b}$ are *topmost* elements of $\mathcal{S}_2(\phi, r)$ and $\mathcal{S}_2(\phi, \widetilde{r})$, respectively, if for all $r_1 < r$ (see Notation 6.1) neither b nor $b^{-1} := (L_2^{-1}, L_1^{-1})$ is an element of $\mathcal{S}_2(\phi, r_1)$. Since $\widetilde{b}$ and $\Phi_{\widetilde{r}}(\widetilde{b})$ project to the same element of $\mathcal{S}_2(\phi, r)$ and since $\mathcal{S}_2(\phi, r)$ is $\Phi_{\widetilde{r}}$-invariant, it follows that the set of topmost element of $\mathcal{S}_2(\phi, \widetilde{r})$ is $\Phi_{\widetilde{r}}$-invariant.

Lemma 15.31 *The set of topmost elements of $\mathcal{S}_2(\phi, \widetilde{r})$ is the union of a finite number of $\Phi_{\widetilde{r}}$-orbits. Moreover, there exists a computable $B(r) > 0$ such that each of these orbits has a visible representative with index at most $B(r)$.*

Proof Define $B(r) > 2M$ by $\Phi^{2M}_{\tilde{r}}(\tilde{\ell}_p) = \tilde{\ell}_{B(r)}$.

Suppose that $\tilde{b}$ is a topmost element of $\mathcal{S}_2(\phi, \tilde{r})$. After replacing $\tilde{b}$ with some $\Phi^k_{\tilde{r}}(\tilde{b})$, we may assume by Lemma 15.26(4) that $\tilde{b} = (\tilde{\ell}_{i-1}, \tilde{\ell}_i)$ or $\tilde{b} = (\tilde{\ell}_{i-1}, \tilde{\ell}_{i+1})$ with common axis $\tilde{l}_i$. We consider the $\tilde{b} = (\tilde{\ell}_{i-1}, \tilde{\ell}_i)$ case first, assuming without loss that $i > 2M$. The proof below is similar to that of Lemma 15.17.

Suppose that there exists $1 \le j' \le p$ and $k' > 0$ so that $\tilde{\sigma}_{i-1}\tilde{\rho}_{i-1}\tilde{\sigma}_i\tilde{\rho}_i\tilde{\sigma}_{i+1} \subset \tilde{f}^{k'}_{\#}(\tilde{\sigma}_{j'})$. Then $\sigma_{j'} = E'$ or $\bar{E}'$ for some $E' \in \mathcal{E}_f$ so either $\sigma_{i-1}\rho_{i-1}\sigma_i\rho_i\sigma_{i+1}$ or $\bar{\sigma}_{i+1}\bar{\rho}_i\bar{\sigma}_i\bar{\rho}_{i-1}\bar{\sigma}_{i-1}$ occurs as a concatenation of terms in the coarsening of the complete splitting of $R_{E'}$. Letting $r' \in \mathcal{R}(\phi)$ correspond to E', Lemma 15.26 implies that either b or b^{-1} is an element of $\mathcal{S}_2(\phi, r')$ in contradiction to the assumption that b is topmost in $\mathcal{S}_2(\phi, r)$. Thus no such j' and k' exist. It follows that there exists $1 \le j \le p$ and $k > 0$ such that $\tilde{f}^k_{\#}(\tilde{\rho}_j)$ is contained in either $\tilde{\rho}_{i-1}$ or $\tilde{\rho}_i$. Equivalently, $\Phi^k_{\tilde{r}}(\tilde{\ell}_j)$ is equal to either $\tilde{\ell}_{i-1}$ or $\tilde{\ell}_i$. Since $\tilde{\ell}_j$ is one of the lines comprising the pair $\Phi^{-k}_{\tilde{r}}(\tilde{b})$, Lemma 15.29 implies that $\Phi^{2M-k}_{\tilde{r}}(\tilde{b}) = \Phi^{2M}_{\tilde{r}}(\Phi^{-k}_{\tilde{r}}(\tilde{b}))$ is visible with index at most $B(r)$.

In the remaining case, $\tilde{b} = (\tilde{\ell}_{i-1}, \tilde{\ell}_{i+1})$ with common axis $\tilde{\ell}_i$. Arguing as in the first case, we conclude that there exists $k \ge 0$ and $1 \le j \le p$ such that $\Phi^k_{\tilde{r}}(\tilde{\ell}_j)$ is equal to either $\tilde{\ell}_{i-1}$ or $\tilde{\ell}_i$ or $\tilde{\ell}_{i+1}$. The proof then concludes as in the first case. □

Remark 15.32 The $\Phi_{\tilde{r}}$-image of a visible topmost staple pair is a visible topmost staple pair. It follows that if a topmost staple pair $\tilde{b}$ occurs in $\tilde{r}$ and if $\tilde{\ell}_{B(r)} \prec \tilde{b}$ then $\tilde{b}$ is visible.

Remark 15.33 The set of topmost elements of $\mathcal{S}_2(\phi, r)$ could be empty.

Lemma 15.34 *If $\hat{r} < r$ and $b \in \mathcal{S}_2(\phi, \hat{r})$, then either $b \in \mathcal{S}_2(\phi, r)$ or $b^{-1} \in \mathcal{S}_2(\phi, r)$.*

Proof The proof is similar to that of Lemma 15.31. Let $\hat{E}$ be the higher-order edge corresponding to $\hat{r}$, let $R_{\hat{E}} = \hat{\rho}_0 \cdot \hat{\sigma}_1 \cdot \hat{\rho}_1 \cdot \ldots$ be the coarsening of the complete splitting into single growing terms and Nielsen paths and let $\hat{\ell}_1, \hat{\ell}_2, \ldots$ be the associated visible lines. By Lemmas 15.26(4) and 15.25, there exists $i \ge 1$ such that $b = (\hat{\ell}_{i-1}, \hat{\ell}_i)$ or $b = (\hat{\ell}_{i-1}, \hat{\ell}_{i+1})$. Since $\hat{r} < r$, there exists $j > 1$ such that either $\sigma_j = \hat{E}$ or $\sigma_j = \hat{E}^{-1}$. The cases are symmetric so we assume that $\sigma_j = \hat{E}^{-1}$ and leave the $\sigma_j = \hat{E}$ case to the reader. Since $\ell_j \in \Omega_{\mathsf{NP}}(r)$, the inverse of every finite subpath of $R_{\hat{E}}$ occurs as subpath of R_E. In particular, the inverse of $\hat{\rho}_{i-2} \cdot \hat{\sigma}_{i-1} \ldots \cdot \hat{\sigma}_{i+2}\hat{\rho}_{i+2}$ occurs as a concatenation of terms in R_E. Lemma 15.26 therefore implies that $b^{-1} \in \mathcal{S}_2(\phi, r)$. □

Lemma 15.35 *Suppose that $\theta \in \mathsf{Out}(F_n)$ conjugates ϕ to ψ, that $\theta(r) = r' \in \mathcal{R}(\psi)$, that $\tilde{r}, \tilde{r}' \in \partial F_n$ represent r and r', respectively, and that Θ is the lift of θ such that $\Theta(\tilde{r}) = \tilde{r}'$. Then Θ induces a bijection $\mathcal{S}_2(\phi, \tilde{r}) \mapsto \mathcal{S}_2(\psi, \tilde{r}')$ that restricts to a bijection on topmost elements.*

Proof This follows from Lemma 15.3, which provides a bijection between $\Omega_{\mathsf{NP}}(\phi, \tilde{r})$ and $\Omega_{\mathsf{NP}}(\psi, \tilde{r}')$, and the definitions. □

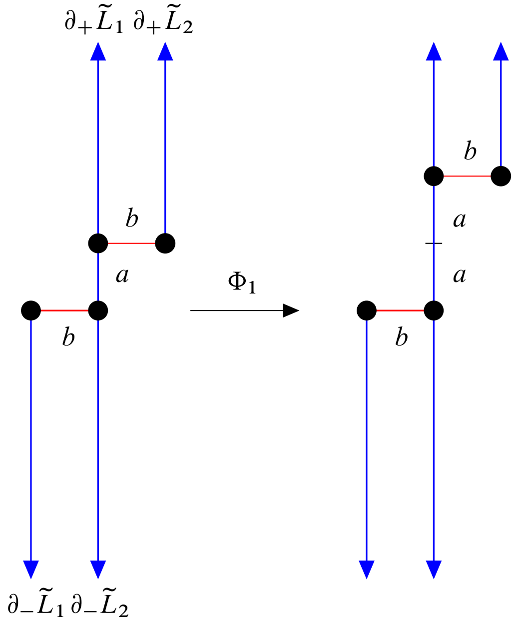

Figure 5: The graphs here are parts of $\widetilde{G}$ with each horizontal segment a lift of the edge b and where vertical segments project into a^∞. $\widetilde{L}_1$ and $\widetilde{L}_2$ are lifts of the staple $a^\infty b a^\infty$ and (L_1, L_2) is a staple pair. Φ_1 is the lift of ϕ that fixes $\widetilde{L}_1$. Intuitively $\widetilde{L}_2$ slides away from $\widetilde{L}_1$ under the action of Φ_1 by 1 period and so $m_{(L_1,L_2)}(\phi) = 1$

.

Definition 15.36 Given $b = (L_1, L_2) \in \mathcal{S}_2(\phi, r)$, choose lifts $\widetilde{L}_1, \widetilde{L}_2$ and a periodic line $\widetilde{A}$ such that $\{\partial_+\widetilde{L}_1, \partial_-\widetilde{L}_2\} \subset \{\partial_-\widetilde{A}, \partial_+\widetilde{A}\}$. Orient $\widetilde{A}$ to be consistent with the twist path w to which it projects and let $a \in F_n$ be the root-free element of F_n that stabilizes $\widetilde{A}$ and satisfies $a^+ = \partial_+\widetilde{A}$. Each $\theta \in \mathcal{X}_\mathfrak{c}(\phi)$ (Definition 14.1) satisfies $\theta(\mathsf{H}_{\phi,\mathfrak{c}}(L_i)) = \mathsf{H}_{\phi,\mathfrak{c}}(L_i)$ for $i = 1, 2$. Lemma 13.10 therefore implies that there are unique $\Theta_i \in \theta$ such that $\Theta_i(\mathsf{H}_{\phi,\mathfrak{c}}(\widetilde{L}_i)) = \mathsf{H}_{\phi,\mathfrak{c}}(\widetilde{L}_i)$. Since both Θ_1 and Θ_2 represent θ and fix a there exists $m_b(\theta) \in \mathbb{Z}$ such that $\Theta_1 = i_a^{m_b(\theta)}\Theta_2$.

Example 3.1 (continued) See Figure 5.

Lemma 15.37 *For each $b = (L_1, L_2) \in \mathcal{S}_2(\phi)$, the map $m_b : \mathcal{X}_\mathfrak{c}(\phi) \to \mathbb{Z}$ is a well-defined homomorphism.*

Proof We first check that $m_b(\theta)$ is independent of the choice of $\widetilde{L}_1$ and $\widetilde{L}_2$ and so is well-defined. Suppose that $\widetilde{L}'_1$ and $\widetilde{L}'_2$ are another choice with corresponding $\widetilde{A}'$, a' and Θ'_i. Choose $c \in F_n$ such that $i_c(a) = a'$. For $i = 1, 2$, $\widetilde{L}'_i$ and $i_c(\widetilde{L}_i)$ are lifts of L_i with an endpoint in $\{a'^-, a'^+\}$ and so there exists an n_i such that

$$\widetilde{L}'_i = i_{a'}^{n_i} i_c(\widetilde{L}_i) = i_c i_a^{n_i}(\widetilde{L}_i).$$

By uniqueness,

$$\Theta'_i = (i_c i_a^{n_i})\Theta_i(i_c i_a^{n_i})^{-1} = i_c i_a^{n_i}\Theta_i i_a^{-n_i} i_c^{-1} = i_c\Theta_i i_c^{-1},$$

so

$$\Theta_1'\Theta_2'^{-1} = i_c\Theta_1\Theta_2^{-1}i_c^{-1} = i_c i_a^{m_b(\theta)} i_c^{-1} = i_{a'}^{m_b(\theta)},$$

as desired.

To prove that $m_b(\theta)$ defines a homomorphism, suppose that $\psi \in \mathcal{X}_{\mathfrak{c}}(\phi)$ and Ψ_i satisfies $\Psi_i(\mathsf{H}_{\phi,\mathfrak{c}}(\widetilde{L}_i)) = \mathsf{H}_{\phi,\mathfrak{c}}(\widetilde{L}_i)$. Then $\Psi_i\Theta_i(\mathsf{H}_{\phi,\mathfrak{c}}(\widetilde{L}_i)) = \mathsf{H}_{\phi,\mathfrak{c}}(\widetilde{L}_i)$ and

$$i_a^{m_b(\psi\theta)} = \Psi_1\Theta_1\Theta_2^{-1}\Psi_1^{-1} = \Psi_1 i_a^{m_b(\theta)}\Psi_2^{-1} = \Psi_1\Psi_2^{-1} i_a^{m_b(\theta)} = i_a^{m_b(\psi)} i_a^{m_b(\theta)} = i_a^{m_b(\psi)+m_b(\theta)},$$

so $m_b(\psi\theta) = m_b(\psi) + m_b(\theta)$. □

Remark 15.38 The same proof shows that m_b defines a homomorphism on both

$$\{\theta \in \mathsf{Out}(F_n) \mid \theta(L_i) = L_i \text{ for } i = 1,2\} \quad \text{and} \quad \{\theta \in \mathsf{Out}(F_n) \mid \theta(\mathsf{H}_{\phi,\mathfrak{c}}(L_i)) = \mathsf{H}_{\phi,\mathfrak{c}}(L_i) \text{ for } i = 1,2\}.$$

The former is the stabilizer of b and the latter can be thought of as the "weak stabilizer" of b.

The next lemma relates $m_b(\phi)$ to the twist coefficients of ϕ.

Lemma 15.39 *Suppose that $b = (L_1, L_2) \in \mathcal{S}_2(\phi, r)$, where $L_1 = (R_1^-)^{-1}\rho_1 R_1^+$ and $L_2 = (R_2^-)^{-1}\rho_2 R_2^+$ are the decompositions of Corollary 5.17(1). Suppose also that w is a twist path and that $E', E'' \in \mathsf{Lin}_w(f)$ satisfy $f(E') = E'w^{d'}$ and $f(E'') = E''w^{d''}$.*

(1) *If $R_1^+ = E'w^{\pm\infty}$ and $R_2^- = E''w^{\pm\infty}$, then $m_b(\phi) = d' - d''$.*
(2) *If $R_1^+ = E'w^{\pm\infty}$ and $R_2^- = w^{\pm\infty}$, then $m_b(\phi) = d'$.*
(3) *If $R_1^+ = w^{\pm\infty}$ and $R_2^- = E''w^{\pm\infty}$, then $m_b(\phi) = -d''$.*

In particular, $m_b(\phi) \neq 0$ for all $b \in \mathcal{S}_2(\phi)$.

Proof Choose lifts $\widetilde{L}_1 = (\widetilde{R}_1^-)^{-1}\widetilde{\rho}_1\widetilde{R}_1^+$, $\widetilde{L}_2 = (\widetilde{R}_2^-)^{-1}\widetilde{\rho}_2\widetilde{R}_2^+$ and $\widetilde{A} = \widetilde{w}^\infty$ so that $\partial\widetilde{R}_1^+, \partial\widetilde{R}_2^- \in \{\partial_-\widetilde{A}, \partial_+\widetilde{A}\}$. Denote the initial endpoints of $\widetilde{R}_1^+$ and $\widetilde{R}_2^-$ by $\widetilde{x}$ and $\widetilde{y}$, respectively. There exist $\Phi_1, \Phi_2 \in \phi$ such that Φ_1 fixes the endpoints of $\widetilde{L}_1$ and Φ_2 fixes the endpoints of $\widetilde{L}_2$. The corresponding lifts $\widetilde{f}_1$ and $\widetilde{f}_2$ fix $\widetilde{x}$ and $\widetilde{y}$, respectively. In particular, $\widetilde{f}_1(\widetilde{y}) = i_a^{m_b(\phi)}\widetilde{f}_2(\widetilde{y}) = i_a^{m_b(\phi)}\widetilde{y}$. In case (1), the path $\widetilde{\tau}$ connecting $\widetilde{x}$ to $\widetilde{y}$ equals $\widetilde{E}'\widetilde{w}^p(\widetilde{E}'')^{-1}$ for some $p \in \mathbb{Z}$ and $(\widetilde{f}_1)_\#(\widetilde{E}'\widetilde{w}^p(\widetilde{E}'')^{-1}) = \widetilde{E}'\widetilde{w}^{p+d'-d''}(\widetilde{E}'')^{-1}$. It follows that $\widetilde{f}_1(\widetilde{y}) = i_a^{d'-d''}\widetilde{y}$ and hence that $m_b(\phi) = d' - d''$. In case (2), $\widetilde{\tau} = \widetilde{E}'\widetilde{w}^p$ and $(\widetilde{f}_1)_\#(\widetilde{E}'\widetilde{w}^p) = \widetilde{E}'\widetilde{w}^{p+d'}$. Thus, $\widetilde{f}_1(\widetilde{y}) = i_a^{d'}\widetilde{y}$ and $m_b(\phi) = d'$. Case (3) is proved similarly.

Lemma 15.26 implies that if b is as in case (1) then either $b = (\ell_{i-1}, \ell_{i+1})$, where $\sigma_i\rho_i\sigma_{i+1}$ is quasi-exceptional, or $b = (\ell_{i-1}, \ell_i)$, where σ_i is exceptional. In either case, $E' \neq E''$ so $d' \neq d''$. This completes the proof that $m_b(\phi) \neq 0$ and hence the proof of the lemma. □

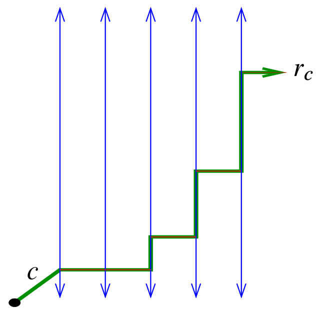

Figure 6: An end of $R_c = cbbaba^2ba^3\dots$ is in a union of staple pairs, cf Figure 5. The relation between staple pairs and R_E for E of higher than quadratic growth is more complicated.

15.4 Spanning staple pairs

We continue with the notation of the preceding subsections; in particular, see Notation 15.4. In addition, we let $\tilde{\mu}_0, \tilde{\mu}_1 \dots$ be the sequence of lines obtained from $\tilde{\ell}_0, \tilde{\ell}_1, \dots$ by removing all periodic lines. In other words, $\tilde{\mu}_0, \tilde{\mu}_1 \dots$ is the set of visible lines in $\Omega_{\mathsf{NP}}(\phi, \tilde{r})$.

If E has quadratic growth (equivalently, each σ_i is linear) then every $(\tilde{\mu}_t, \tilde{\mu}_{t+1})$ is an element of $\mathcal{S}_2(\phi, \tilde{r})$ by Lemma 15.26; see Figure 6. This is not the case when E has higher order. We now define a related but weaker property that does hold for every $(\tilde{\mu}_t, \tilde{\mu}_{t+1})$. Its utility is illustrated in the proof of Lemma 15.45, which is applied in the proof of Lemma 17.9.

Definition 15.40 We say that an ordered pair $(\tilde{\eta}_1, \tilde{\eta}_2)$ of elements of $\Omega_{\mathsf{NP}}(\phi, \tilde{r})$ is *spanned by a staple pair* $\tilde{b} = (\tilde{L}_1, \tilde{L}_2) \in \mathcal{S}_2(\phi, \tilde{r})$ if the following two conditions are satisfied:

- Either $\tilde{L}_1 = \tilde{\eta}_1$ or $\tilde{L}_1 \in \Omega_{\mathsf{NP}}(\phi, \partial_+\tilde{\eta}_1)$.
- Either $\tilde{L}_2 = \tilde{\eta}_2$ or $\tilde{L}_2^{-1} \in \Omega_{\mathsf{NP}}(\phi, \partial_-\tilde{\eta}_2)$.

Note that if $(\tilde{\eta}_1, \tilde{\eta}_2)$ is spanned by a staple pair then $(\Phi_{\tilde{r}}^k\tilde{\eta}_1, \Phi_{\tilde{r}}^k\tilde{\eta}_2)$ is spanned by a staple pair for all $k \in \mathbb{Z}$.

Our next result uses techniques from the proofs of Lemmas 15.12 and 15.31.

Lemma 15.41 (1) *Suppose that $\sigma_i \in \mathcal{E}_f$ and that $\tilde{r}_i = \partial_+\tilde{\ell}_{i-1} = \partial\tilde{R}_i^+$. If $\tilde{\sigma}_j\tilde{\rho}_j\tilde{\sigma}_{j+1}$ is a subpath of $\tilde{f}_\#^m(\tilde{\sigma}_i)$ for $m > 0$ and if $\tilde{\ell}_j$ is nonperiodic, then $\Phi_{\tilde{r}}^{-m}(\tilde{\ell}_j) \in \Omega_{\mathsf{NP}}(\phi, \tilde{r}_i)$.*

(2) *Suppose that $\sigma_i \in \mathcal{E}_f^{-1}$ and that $\tilde{r}_i = \partial_-\tilde{\ell}_i = \partial\tilde{R}_i^-$. If $\tilde{\sigma}_j\tilde{\rho}_j\tilde{\sigma}_{j+1}$ is a subpath of $\tilde{f}_\#^m(\tilde{\sigma}_i)$ for $m > 0$ and if $\tilde{\ell}_j$ is nonperiodic, then $\Phi_{\tilde{r}}^{-m}(\tilde{\ell}_j^{-1}) \in \Omega_{\mathsf{NP}}(\phi, \tilde{r}_i)$.*

Proof The two cases are symmetric so we prove (1) and leave (2) to the reader. Assuming that $\sigma_i \in \mathcal{E}_f$, let $T : \tilde{G} \to \tilde{G}$ be the covering translation that carries $\tilde{\sigma}_i$ to the initial edge of $\tilde{f}_\#^m(\tilde{\sigma}_i)$ and hence satisfies

$T(\widetilde{R}_i^+) = \widetilde{f}_\#^m(\widetilde{R}_i^+) = \Phi_{\widetilde{r}}^m(\widetilde{R}_i^+)$. Then $T^{-1}\widetilde{f}^m$ is the lift of f^m that preserves the terminal endpoint $\widetilde{r}_i$ of $\widetilde{R}_i^+$ and so $i_c^{-1}\Phi_{\widetilde{r}}^m = \Phi_{\widetilde{r}_i}^m$, where i_c is the inner automorphism corresponding to T. Note that $\widetilde{f}_\#^m(\widetilde{\sigma}_i)$ is a concatenation of terms in the complete splitting of $\widetilde{R}_{\widetilde{E}}$ whose first edge $T(\widetilde{\sigma}_i)$ projects into $\mathcal{E}_f$. Thus $\widetilde{f}_\#^m(\widetilde{\sigma}_i) = \widetilde{\sigma}_a \cdot \widetilde{\rho}_a \cdot \ldots \cdot \widetilde{\rho}_{b-1} \cdot \widetilde{\sigma}_b \cdot \widetilde{\tau}$ for some $a \le j < b$ and some (possibly trivial) Nielsen path $\widetilde{\tau}$ that is an initial segment of $\widetilde{\rho}_b$. Note also that $T^{-1}\widetilde{f}_\#^m(\widetilde{\sigma}_i)$ is a concatenation of terms in the complete splitting of $\widetilde{R}_i^+$. It follows that $T^{-1}\widetilde{\sigma}_a \cdot T^{-1}\widetilde{\rho}_a \cdot \ldots \cdot T^{-1}\widetilde{\rho}_{b-1} \cdot T^{-1}\widetilde{\sigma}_b$ is a concatenation of terms in the coarsened complete splitting of $\widetilde{R}_i^+$. In particular, $T^{-1}\widetilde{\sigma}_j \cdot T^{-1}\widetilde{\rho}_j \cdot T^{-1}\widetilde{\sigma}_{j+1}$ is a concatenation of terms in the coarsened complete splitting of $\widetilde{R}_i^+$. Since $\widetilde{\ell}_j$ is nonperiodic, the same is true for $T^{-1}(\widetilde{\ell}_j)$ and we conclude that $T^{-1}(\widetilde{\ell}_j) \in \Omega_{\mathsf{NP}}(\phi, \widetilde{r}_i)$. Proposition 15.6 implies that $\Phi_{\widetilde{r}}^{-m}(\widetilde{\ell}_j) = \Phi_{\widetilde{r}_i}^{-m} i_c^{-1}(\widetilde{\ell}_j) = \Phi_{\widetilde{r}_i}^{-m}(T^{-1}(\widetilde{\ell}_j))$ is an element of $\Omega_{\mathsf{NP}}(\phi, \widetilde{r}_i)$. □

Proposition 15.42 *Each ordered pair* $(\widetilde{\mu}_t, \widetilde{\mu}_{t+1})$ *is spanned by an element* $(\widetilde{L}_1, \widetilde{L}_2) \in \mathcal{S}_2(\phi, \widetilde{r})$. *If* $\partial_+\widetilde{\mu}_t$ (*resp.* $\partial_-\widetilde{\mu}_{t+1}$) *is periodic, then* $\widetilde{L}_1 = \widetilde{\mu}_t$ (*resp.* $L_2 = \widetilde{\mu}_{t+1}$).

Proof Set $\Phi = \Phi_{\widetilde{r}}$. We first show that if $\widetilde{\ell}_i$ is periodic then $\widetilde{\ell}_{i-1}$ and $\widetilde{\ell}_{i+1}$ are nonperiodic and $(\widetilde{\mu}_t, \widetilde{\mu}_{t+1}) := (\widetilde{\ell}_{i-1}, \widetilde{\ell}_{i+1})$ satisfies the conclusions of the lemma. Let M be the stabilization constant for f (Notation 5.13). If $\widetilde{f}_\#^M(\widetilde{\rho}_i) \subset \widetilde{\rho}_j$, then $\widetilde{\sigma}_j$ is the last growing term in $\widetilde{f}_\#^M(\widetilde{\sigma}_i)$ and $\widetilde{\sigma}_{j+1}$ is the first growing term in $\widetilde{f}_\#^M(\widetilde{\sigma}_{i+1})$. Moreover, Lemma 15.9 implies that $\widetilde{\ell}_j = \widetilde{f}_\#^M(\widetilde{\ell}_i)$ and so $\widetilde{\ell}_j$ is periodic. By our choice of M, $\widetilde{\sigma}_j \notin \mathcal{E}_f$ and $\widetilde{\sigma}_{j+1} \notin \mathcal{E}_f^{-1}$. Lemma 15.26(3) implies that $\widetilde{\sigma}_j\widetilde{\rho}_j\widetilde{\sigma}_{j+1}$ is quasi-exceptional and $(\widetilde{\ell}_{j-1}, \widetilde{\ell}_{j+1}) \in \mathcal{S}_2(\phi, \widetilde{r})$. Since $\sigma_j \in \mathsf{Lin}(f)$, it follows that either $\sigma_i \in \mathsf{Lin}(f)$ or $\sigma_i \in \mathcal{E}_f$. In the former case, $\widetilde{f}_\#^M(\widetilde{\rho}_{i-1}) \subset \widetilde{\rho}_{j-1}$ and $\widetilde{f}_\#^M(\widetilde{\ell}_{i-1}) = \widetilde{\ell}_{j-1}$ so $\Phi^{-M}(\widetilde{\ell}_{j-1}) = \widetilde{\ell}_{i-1}$; in particular, $\widetilde{\ell}_{i-1}$ is nonperiodic and $\partial_+\widetilde{\ell}_{i-1}$ is periodic. In the latter case, $\widetilde{\sigma}_{j-1}\widetilde{\rho}_{j-1}\widetilde{\sigma}_j$ is a terminal subpath of $\widetilde{f}_\#^M(\widetilde{\sigma}_i)$ so Lemma 15.41 implies that $\Phi^{-M}(\widetilde{\ell}_{j-1}) \in \Omega_{\mathsf{NP}}(\phi, \partial_+\widetilde{\ell}_{i-1})$; note that in this latter case, $\partial_+\widetilde{\mu}_t = \partial_+\widetilde{\ell}_{i-1}$ is not periodic. A symmetric argument shows that either $\Phi^{-M}(\widetilde{\ell}_{j+1}) = \widetilde{\ell}_{i+1}$ or $\Phi^{-M}(\widetilde{\ell}_{j+1}^{-1}) \in \Omega_{\mathsf{NP}}(\phi, \partial_-\widetilde{\ell}_{i+1})$. Thus $(\widetilde{\ell}_{i-1}, \widetilde{\ell}_{i+1})$ is spanned by $\Phi^{-M}(\widetilde{\ell}_{j-1}, \widetilde{\ell}_{j+1})$, and the "if" statement of the lemma is satisfied.

Remark 15.43 The above argument includes a proof that if $\widetilde{\ell}_i$ is periodic and if $\widetilde{\ell}_j = \widetilde{f}_\#^M(\widetilde{\ell}_i)$, then $\widetilde{\sigma}_j\widetilde{\rho}_j\widetilde{\sigma}_{j+1}$ is quasi-exceptional and $\widetilde{\ell}_j$ is the common axis of $(\widetilde{\ell}_{j-1}, \widetilde{\ell}_{j+1}) \in \mathcal{S}_2(\phi, \widetilde{r})$.

Continuing with the proof, we now know that for each $(\widetilde{\mu}_t, \widetilde{\mu}_{t+1})$ there exists i such that $(\widetilde{\mu}_t, \widetilde{\mu}_{t+1})$ is equal to either $(\widetilde{\ell}_{i-1}, \widetilde{\ell}_i)$ or $(\widetilde{\ell}_{i-1}, \widetilde{\ell}_{i+1})$. Moreover, the conclusions of the lemma hold in the latter case so we may assume that $(\widetilde{\mu}_t, \widetilde{\mu}_{t+1}) = (\widetilde{\ell}_{i-1}, \widetilde{\ell}_i)$. Lemma 15.26(1) implies that $\sigma_i\rho_i\sigma_{i+1}$ is not quasi-exceptional. If σ_i is linear (ie σ_i is exceptional or $\sigma_i \in \mathsf{Lin}(f)$ or $\overline{\sigma}_i \in \mathsf{Lin}(f)$) then $(\widetilde{\mu}_t, \widetilde{\mu}_{t+1}) = (\widetilde{\ell}_{i-1}, \widetilde{\ell}_i) \in \mathcal{S}_2(\phi, \widetilde{r})$ by Lemma 15.26(2).

It remains to consider the $\sigma_i \in \mathcal{E}_f$ and $\sigma_i \in \mathcal{E}_f^{-1}$ cases. These are symmetric so we assume that $\sigma_i \in \mathcal{E}_f$, and leave the $\sigma_i \in \mathcal{E}_f^{-1}$ case to the reader. For the "if" statement, note that $\partial_+\widetilde{\ell}_{i-1}$ is nonperiodic. As above, there exists $j > i$ such that $\widetilde{f}_\#^M(\widetilde{\rho}_i) \subset \widetilde{\rho}_j$ and $\widetilde{\ell}_j = \widetilde{f}_\#^M(\widetilde{\ell}_i)$; in particular, $\widetilde{\ell}_j$ is not periodic. Since $\widetilde{\sigma}_j$ is the last growing term in $\widetilde{f}_\#^M(\widetilde{\sigma}_i)$ and since $\widetilde{f}_\#^M(\widetilde{\sigma}_i)$ contains at least two growing terms, Lemma 15.41 implies that if $\widetilde{\ell}_{j-1}$ is nonperiodic then $\Phi^{-M}(\widetilde{\ell}_{j-1}) \in \Omega_{\mathsf{NP}}(\phi, \partial_+\widetilde{\ell}_{i-1})$. If σ_j is exceptional or $\sigma_j \in \mathsf{Lin}(f)$ then another application of Lemma 15.26(2) shows that $(\widetilde{\ell}_{j-1}, \widetilde{\ell}_j) \in \mathcal{S}_2(\phi, \widetilde{r})$. In this case, $(\widetilde{\ell}_{i-1}, \widetilde{\ell}_i)$ is spanned by $\Phi^{-M}(\widetilde{\ell}_{j-1}, \widetilde{\ell}_j) = (\Phi^{-M}(\widetilde{\ell}_{j-1}), \widetilde{\ell}_i)$ and we are done. The same argument works if $\overline{\sigma}_j \in \mathsf{Lin}(f)$ and $\widetilde{\ell}_{j-1}$ is not periodic. Suppose then that $\overline{\sigma}_j \in \mathsf{Lin}(f)$ and $\widetilde{\ell}_{j-1}$ is periodic. There exists $s > j$ such that $\widetilde{f}^M_\#(\widetilde{\rho}_j) \subset \widetilde{\rho}_s$ and $\widetilde{\ell}_s = \widetilde{f}^M_\#(\widetilde{\ell}_j) = \widetilde{f}^{2M}_\#(\widetilde{\ell}_i)$. Since $\widetilde{\sigma}_j$ is linear, $\widetilde{f}^M_\#(\widetilde{\sigma}_j)$ contains a single growing term and so $\widetilde{f}^M_\#(\widetilde{\rho}_{j-1}) \subset \widetilde{\rho}_{s-1}$. Thus $\widetilde{\ell}_{s-1} = \widetilde{f}^M_\#(\widetilde{\ell}_{j-1})$ is periodic. Remark 15.43 implies that $\widetilde{\sigma}_{s-1}\widetilde{\rho}_{s-1}\widetilde{\sigma}_s$ is quasi-exceptional and $\widetilde{\ell}_{s-1}$ is the common axis of $(\widetilde{\ell}_{s-2}, \widetilde{\ell}_s) \in \mathcal{S}_2(\phi, \widetilde{r})$. Moreover, $\widetilde{\sigma}_{s-2}\widetilde{\rho}_{s-2}\widetilde{\sigma}_{s-1}\widetilde{\rho}_{s-1}\widetilde{\sigma}_s \subset \widetilde{f}^{2M}_\#(\widetilde{\sigma}_i)$. The proof now concludes as in the previous case with $(\widetilde{\ell}_{i-1}, \widetilde{\ell}_i)$ spanned by $\Phi^{-2M}(\widetilde{\ell}_{s-2}, \widetilde{\ell}_s) = (\Phi^{-2M}(\widetilde{\ell}_{s-2}), \widetilde{\ell}_i)$.

We may now assume that σ_j has higher order and so $\sigma_j \in \mathcal{E}_f^{-1}$ by our choice of M. In particular, $\partial_-\widetilde{\ell}_j$, and hence $\partial_-\widetilde{\ell}_i$, is nonperiodic.

Choose $k > 0$ so that the coarsened complete splitting of $\widetilde{f}^k_\#(\widetilde{\sigma}_j)$ has at least one linear term $\widetilde{\sigma}_s$ that is neither the first nor second term nor the last or next to last growing term in that splitting. Thus $\widetilde{\sigma}_{s-2} \cdot \widetilde{\rho}_{s-2} \cdot \ldots \cdot \widetilde{\rho}_{s+1} \cdot \widetilde{\sigma}_{s+2}$ is a subpath of $\widetilde{f}^k_\#(\widetilde{\sigma}_j)$ and hence also a subpath of $\widetilde{f}^{M+k}_\#(\widetilde{\sigma}_i)$. By Lemma 15.26, two of the three lines $\widetilde{\ell}_{s-1}, \widetilde{\ell}_s, \widetilde{\ell}_{s+1}$ form an element of $\mathcal{S}_2(\phi, \widetilde{r})$ that we denote by $(\widetilde{L}_1, \widetilde{L}_2)$. Lemma 15.41 implies that

$$\Phi^{-M-k}(\widetilde{L}_1) \in \Omega_{\mathsf{NP}}(\phi, \partial_+\widetilde{\ell}_{i-1}) \quad \text{and} \quad \Phi^{-k}(\widetilde{L}_2^{-1}) \in \Omega_{\mathsf{NP}}(\phi, \partial_-\widetilde{\ell}_j).$$

It follows that

$$\Phi^{-M-k}(\widetilde{L}_2^{-1}) \in \Omega_{\mathsf{NP}}(\phi, \partial_-\widetilde{\ell}_i),$$

and hence that $\Phi^{-M-k}((\widetilde{L}_1, \widetilde{L}_2))$ spans $(\widetilde{\ell}_{i-1}, \widetilde{\ell}_i)$. □

In Lemma 15.45 below we use Proposition 15.42 to give conditions on $\upsilon \in \mathsf{Out}(F_n)$ which imply that υ fixes r. The proof of Lemma 15.45 is inductive and it is useful in the induction step to know that υ strongly fixes r in the following sense.

Definition 15.44 We say that $\upsilon \in \mathsf{Out}(F_n)$ *strongly fixes* $r \in \mathcal{R}(\phi)$ if for some (and hence every) lift $\widetilde{r}$ there is a lift $\Upsilon \in \upsilon$ that fixes each element of $\Omega_{\mathsf{NP}}(\phi, \widetilde{r})$.

Lemma 15.45 *Suppose that $[F]$ is a ϕ-invariant free factor conjugacy class, that $r \in \mathcal{R}(\phi)$ is carried by $[F]$ and that $\upsilon \in \mathsf{Out}(F_n)$ is such that*

(1) *$[F]$ is υ-invariant,*

(2) *υ fixes each element of $\Omega_{\mathsf{NP}}(r)$ and each $r' < r$ (as defined in Notation 6.1),*

(3) *the restriction $\upsilon|F$ commutes with $\phi|F$, and*

(4) *$m_b(\upsilon) = 0$ for all $b \in \mathcal{S}_2(\phi, r)$ (see Definition 15.36 and Remark 15.38),*

then υ strongly fixes r.

Proof Given a lift $\widetilde{r}$, we continue with the $\widetilde{\sigma}_i$, $\widetilde{\rho}_i$, $\widetilde{\ell}_i$, $\widetilde{\mu}_t$ and $\Phi_{\widetilde{r}}$ notation. We may assume without loss of generality that $\widetilde{r} \in \partial F$. By uniqueness, F is $\Phi_{\widetilde{r}}$-invariant. By Lemma 15.13 each $\widetilde{\mu}_t$ has endpoints in ∂F.

For each $t \geq 1$, item (2) implies the existence of a (necessarily unique) lift Υ_t of υ that fixes $\widetilde{\mu}_t$. We show below that Υ_t is independent of t, say $\Upsilon_t = \Upsilon$ for all t. Assuming this for now, the proof concludes as follows. Since the endpoints of the $\widetilde{\mu}_t$ limit on $\widetilde{r}$, we have $\Upsilon(\widetilde{r}) = \widetilde{r}$. From this and (1) it follows that $\Upsilon(F) = F$. Item (3) implies that the commutator $[\Phi_{\widetilde{r}}|F, \Upsilon|F]$ is inner. Since the commutator $[\Phi_{\widetilde{r}}|F, \Upsilon|F]$ fixes $\widetilde{r}$, it must be trivial. Thus, $\Phi_{\widetilde{r}}|F$ and $\Upsilon|F$ commute and the same is true for $\Phi_{\widetilde{r}}|\partial F$ and $\Upsilon|\partial F$. Given $\widetilde{L} \in \Omega_{\mathsf{NP}}(\phi, \widetilde{r})$, there exist $m \geq 0$ and $t \geq 1$ such that $\Phi_{\widetilde{r}}^m(\widetilde{L}) = \widetilde{\mu}_t$ by Proposition 15.6. Since $\widetilde{\mu}_t$ has endpoints in ∂F, the same is true for $\widetilde{L}$. Thus

$$\Upsilon(\widetilde{L}) = (\Phi_{\widetilde{r}}^{-m} \Upsilon \Phi_{\widetilde{r}}^m)(\widetilde{L}) = \Phi_{\widetilde{r}}^{-m} \Upsilon(\Phi_{\widetilde{r}}^m(\widetilde{L})) = \Phi_{\widetilde{r}}^{-m} \Upsilon(\widetilde{\mu}_t) = \Phi_{\widetilde{r}}^{-m}(\widetilde{\mu}_t) = \widetilde{L},$$

as desired.

It remains to prove that $\Upsilon_t = \Upsilon_{t+1}$ for all $t \geq 1$.

The proof is by induction on the height of r in the partial order $<$ on $\mathcal{R}(\phi)$. In the base case, r is a minimal element of $\mathcal{R}(\phi)$ so each σ_i is linear by Lemmas 6.2 and 3.21. In this case, each $\widetilde{\mu}_t \in \mathcal{S}(\phi, \widetilde{r})$ and each $(\widetilde{\mu}_t, \widetilde{\mu}_{t+1}) \in \mathcal{S}_2(\phi, \widetilde{r})$ by Lemma 15.26. Item (4) completes the proof.

For the inductive step, we use Proposition 15.42. Let $\mu_t = \ell_i$ and $\mu_{t+1} = \ell_j$. As a first case, suppose that $\partial_+ \widetilde{\ell}_i = \widetilde{r}'$ for some $r' \in \mathcal{R}(\phi)$. Let $E' \in \mathcal{E}_f$ be the higher-order edge corresponding to r'. Then $r > r'$ and Lemmas 6.2 and 3.21 imply that either E' or $\bar{E}'$ occurs as a term in the complete splitting of some $f_\#^k(E)$. Thus either $\mathsf{Acc}(E') \subset \mathsf{Acc}(E)$ or $\mathsf{Acc}(\bar{E}') \subset \mathsf{Acc}(E)$. Lemma 5.8 therefore implies that if $L' \in \Omega_{\mathsf{NP}}(r')$ then either $L' \in \Omega_{\mathsf{NP}}(r)$ or ${L'}^{-1} \in \Omega_{\mathsf{NP}}(r)$. From (2) we see that υ fixes each element of $\Omega_{\mathsf{NP}}(r')$. Lemma 15.34 and (4) imply that $m_b(\upsilon) = 0$ for all $b \in \mathcal{S}_2(\phi, r')$ so r' and $[F]$ satisfy the hypotheses of this lemma. By the inductive hypothesis, there is a (necessarily unique) lift Υ' that fixes each element of $\Omega_{\mathsf{NP}}(\phi, \widetilde{r}')$. As noted above, it follows that Υ' fixes $\widetilde{r}'$ and so $\Upsilon' = \Upsilon_t$.

There are two subcases. The first is that $\partial_- \widetilde{\ell}_j = \widetilde{r}''$ for some $r'' \in \mathcal{E}_f$. Arguing as in the previous paragraph we see that Υ_{t+1} fixes each element of $\Omega_{\mathsf{NP}}(\phi, \widetilde{r}'')$. By Proposition 15.42, there exist $\widetilde{L}' \in \Omega_{\mathsf{NP}}(\phi, \widetilde{r}')$ and $\widetilde{L''}^{-1} \in \Omega_{\mathsf{NP}}(\phi, \widetilde{r}'')$ such that $(\widetilde{L}', \widetilde{L}'') \in \mathcal{S}_2(\phi, \widetilde{r})$. Item (4) implies that $\Upsilon_t = \Upsilon_{t+1}$.

The second subcase is that $\partial_- \widetilde{\ell}_j$ is the end of a periodic line and hence $\ell_j \in \mathcal{S}(\phi, r)$. By Proposition 15.42, there exists $\widetilde{L}' \in \Omega_{\mathsf{NP}}(\phi, \widetilde{r}')$ such that $(\widetilde{L}', \widetilde{\ell}_j) \in \mathcal{S}_2(\phi, \widetilde{r})$. Once again, (4) implies that $\Upsilon_t = \Upsilon_{t+1}$. This completes the proof when $\partial_+ \widetilde{\ell}_i$ projects into $\mathcal{R}(\phi)$.

A symmetric argument handles the case that $\partial_- \widetilde{\ell}_j$ projects into $\mathcal{R}(\phi)$. The remaining case is that both $\partial_+ \widetilde{\ell}_i$ and $\partial_- \widetilde{\ell}_j$ are endpoints of periodic lines. Proposition 15.42 implies that $(\widetilde{\ell}_i, \widetilde{\ell}_j) \in \mathcal{S}_2(\phi, \widetilde{r})$ so (4) completes the proof as in previous cases. □

16 The homomorphism $\bar{Q}$

Recall that our main theorem is reduced to Proposition 14.7 by Lemma 14.8.

For the rest of the paper, we assume the hypotheses of Proposition 14.7, ie $\phi, \psi \in \mathsf{UPG}(F_n)$, $\mathfrak{c}$ is a special chain for ϕ and ψ, and $\mathsf{I}_\mathfrak{c}(\phi) = \mathsf{I}_\mathfrak{c}(\psi)$. Our goal is to find a conjugator $\theta \in \mathcal{X}_\mathfrak{c}(\phi)$ or prove that no such conjugator exists.

Because ϕ and $\mathfrak{c}$ are fixed for the rest of the paper, we will often write $\mathcal{X}$ for $\mathcal{X}_\mathfrak{c}(\phi)$. In fact, we will often suppress $\mathfrak{c}$ when it appears as a decoration.

Lemma 16.1 *For each $\mathcal{F}_i \in \mathfrak{c}$ and each $r \in \mathcal{R}(\phi|\mathcal{F}_i)$ there exists $r' \in \mathcal{R}(\psi|\mathcal{F}_i)$ such that $\theta(r) = r'$ for each $\theta \in \mathcal{X}$ whose restriction $\theta|\mathcal{F}_i$ conjugates $\phi|\mathcal{F}_i$ to $\psi|\mathcal{F}_i$.*

Proof Let $\mathfrak{e} = \mathcal{F}^- \sqsubset \mathcal{F}^+ \in \mathfrak{c}$ be the one-edge extension with respect to which r is new (Definition 6.14). In other words, $r \in \mathcal{R}^+(\mathfrak{e}, \phi) := \mathcal{R}(\phi|\mathcal{F}^+) \setminus \mathcal{R}(\phi|\mathcal{F}^-)$. Since θ fixes $\mathfrak{c}$, it follows that $\theta|\mathcal{F}^\pm$ conjugates $\phi|\mathcal{F}^\pm$ to $\psi|\mathcal{F}^\pm$ and so θ induces a bijection between $\mathcal{R}^+(\mathfrak{e}, \phi)$ and $\mathcal{R}^+(\mathfrak{e}, \psi)$. This completes the proof if r is the only element of $\mathcal{R}^+(\mathfrak{e}, \phi)$. Otherwise $\mathcal{R}^+(\mathfrak{e}, \phi) = \{r, s\}$ and $\mathcal{R}^+(\mathfrak{e}, \psi) = \{r', s'\}$ and we are in case HH. By definition, $\mathsf{L}_\mathfrak{e}(\phi) = \{L, L^{-1}\}$, where $\partial_- L = r$ and $\partial_+ L = s$. By Lemma 6.15, $\theta(L) \in \mathsf{L}_\mathfrak{e}(\psi)$. Since $\theta \in \mathcal{X}$, we get $\theta(\mathsf{H}_{\phi,\mathfrak{c}}(L)) = \mathsf{H}_{\phi,\mathfrak{c}}(L) = \mathsf{H}_{\psi,\mathfrak{c}}(\theta(L))$. By Lemma 13.9, there is a unique $L' \in \mathsf{L}_\mathfrak{e}(\psi)$ that is in $\mathsf{H}_{\psi,\mathfrak{c}}(\theta(L))$. Hence $\theta(L) = L'$ and $\theta(r) = r' := \partial_- L'$. □

We continue with the notation of Section 15 and also assume that a CT $f' \colon G' \to G'$ representing ψ has been chosen that realizes $\mathfrak{c}$. We use prime notation when working with ψ and r'; for example, $E' \in \mathcal{E}_{f'}$ is the edge corresponding to r' and $\tilde{\ell}'_1, \tilde{\ell}'_2, \dots$ are the visible lines in $\tilde{R}'_{\tilde{E}'}$ and $\Psi_{\tilde{r}'}$ is the lift $\Psi \in \psi$ that fixes $\tilde{r}'$.

Definition 16.2 Recall from Corollary 5.17, Definition 15.36 and Lemma 15.39 that $\mathcal{S}_2(\phi)$ is finite, that for all $b \in \mathcal{S}_2(\phi)$ there is a homomorphism $m_b \colon \mathcal{X} \to \mathbb{Z}$ and that $m_b(\phi) \neq 0$. Define a homomorphism $Q^\phi \colon \mathcal{X} \to \mathbb{Q}^{\mathcal{S}_2(\phi)}$ by letting the b-coordinate of $Q^\phi(\theta)$ be $Q^\phi_b(\theta) = m_b(\theta)/m_b(\phi)$.

Definition 16.3 Let $\sim$ be the equivalence relation on $\mathcal{S}_2(\phi)$ generated by $b \sim b'$ if b and b' occur in the same $r \in \mathcal{R}(\phi)$ (as defined in Notation 15.24), and let

$$\mathcal{S}_2(\phi) = \mathcal{S}_2^1(\phi) \sqcup \mathcal{S}_2^2(\phi) \sqcup \cdots$$

be the decomposition of $\mathcal{S}_2(\phi)$ into $\sim$-equivalence classes. For each i, consider the diagonal action of $\mathbb{Z}$ on $\mathbb{Q}^{\mathcal{S}_2^i(\phi)}$, ie $k\vec{s} = \vec{s} + k(1, 1, \dots)$. Let $\bar{Q}^\phi$ denote the homomorphism

$$\mathcal{X}_\mathfrak{c}(\phi) \xrightarrow{Q^\phi} \mathbb{Q}^{\mathcal{S}_2(\phi)} \to \bar{\mathbb{Q}}^{\mathcal{S}_2(\phi)} := (\mathbb{Q}^{\mathcal{S}_2^1(\phi)}/\mathbb{Z}) \oplus (\mathbb{Q}^{\mathcal{S}_2^2(\phi)}/\mathbb{Z}) \oplus \cdots.$$

For the rest of the paper, Q and $\bar{Q}$ will always denote Q^ϕ and $\bar{Q}^\phi$.

We can now state the second reduction of the conjugacy problem for $\mathsf{UPG}(F_n)$ in $\mathsf{Out}(F_n)$.

Proposition 16.4 *There is an algorithm that takes as input $\phi, \psi \in \mathsf{UPG}(F_n)$ and a chain $\mathfrak{c}$ such that*

- *$\mathfrak{c}$ is a special chain for ϕ and ψ, and*
- $\mathsf{I}_{\mathfrak{c}}(\phi) = \mathsf{I}_{\mathfrak{c}}(\psi)$,

and that outputs `YES` *or* `NO` *depending on whether or not there is a $\theta \in \mathsf{Ker}(\bar{Q}^{\phi})$ conjugating ϕ to ψ. Further, if* `YES`*, then such a θ is produced.*

Lemma 16.5 *Proposition 16.4 implies Proposition 14.7 and hence Theorem 1.1.*

Proposition 16.4 is proved in Section 17.

Lemma 16.5 is proved by applying the following technical proposition, whose proof takes up the rest of this section.

Proposition 16.6 *We have an algorithm that produces a finite set $\{\eta_i\} \subset \mathcal{X}$ so that the union of the cosets of $\mathsf{Ker}(\bar{Q})$ determined by the η_i contains each $\theta \in \mathcal{X}$ that conjugates ϕ to ψ.*

Proof of Lemma 16.5, assuming Proposition 16.6 Let $\{\eta_i\}$ be the finite set produced by Proposition 16.6 and let $\psi_i = \psi^{(\eta_i^{-1})}$. It follows that $\phi^\theta = \psi$ if and only if $\phi^{\theta_i'} = \psi_i$, where $\theta_i' = \eta_i^{-1}\theta$, and that θ is in the coset represented by η_i if and only if $\theta_i' \in \mathsf{Ker}(\bar{Q})$. Thus, by applying Proposition 16.4 to ϕ and ψ_1, we can decide if there exists θ in the coset represented by η_1 that conjugates ϕ to ψ. If `YES` then return `YES` and one such θ. Otherwise move on to η_2 and repeat. If `NO` for each η_i, then return `NO`. □

The following two lemmas are proved in Sections 16.1 and 16.3, respectively. In the remainder of this section we use them to prove Proposition 16.6. The definition of topmost staple pair appears in Notation 15.30. The definition of $\mathsf{offset}(\theta, r)$ is given in Lemma 15.21(2). The partial order $<$ on $\mathcal{R}(\phi)$ is defined in Notation 6.1.

Lemma 16.7 *Suppose that $b \in \mathcal{S}_2(\phi, r)$ is topmost and that $\theta \in \mathcal{X}$ conjugates ϕ to ψ. Then given an upper bound for $|\mathsf{offset}(\theta, r)|$ one can compute an upper bound for $|m_\theta(b)|$.*

Lemma 16.8 *Suppose that $\theta \in \mathcal{X}$ conjugates ϕ to ψ and that $r, r_1 \in \mathcal{R}(\phi)$ satisfy $r_1 < r$. Then given an upper bound for $|\mathsf{offset}(\theta, r)|$ one can compute an upper bound for $|\mathsf{offset}(\theta, r_1)|$.*

Proof of Proposition 16.6, assuming Lemmas 16.7 and 16.8 We begin by computing $D = D(\phi, \psi)$ so that $|Q_{b_1}(\theta) - Q_{b_2}(\theta)| < D$ for all $\theta \in \mathcal{X}$ that conjugate ϕ to ψ and all $b_1, b_2 \in \mathcal{S}_2(\phi)$ satisfying $b_1 \sim b_2$.

Given $r \in \mathcal{R}(\phi)$ we will find D_r such that $|Q_{b_1}(\theta) - Q_{b_2}(\theta)| < D_r$ for all $\theta \in \mathcal{X}$ that conjugate ϕ to ψ and all $b_1, b_2 \in \mathcal{S}_2(\phi, r)$. We then take $D = |\mathcal{R}(\phi)| \max\{D_r\}$, where the $|\mathcal{R}(\phi)|$ factor allows us to consider equivalent staple pairs that do not occur in the same ray.

For all $s \in \mathbb{Z}$, $\theta\phi^s$ is an element of $\mathcal{X}$ and conjugates ϕ to ψ; see Lemma 14.3. The translation number $\tau(\phi, r)$ is defined in Notation 15.18. By definition and by Lemma 15.21 we have

$$\mathsf{offset}(\theta\phi^s, r) = \mathsf{offset}(\theta, r) + \tau(\phi^s, r) = \mathsf{offset}(\theta, r) + s\tau(\phi, r).$$

Since

$$Q_{b_1}(\theta\phi^s) - Q_{b_2}(\theta\phi^s) = (Q_{b_1}(\theta) + s) - (Q_{b_2}(\theta) + s) = Q_{b_1}(\theta) - Q_{b_2}(\theta),$$

we may assume without loss of generality that

$$0 \le \mathsf{offset}(\theta, r) \le \tau(\phi, r).$$

Using only this inequality we will produce an upper bound D_0 for $|m_\theta(b)|$ when $b \in \mathcal{S}_2(\phi, r)$. This determines an upper bound for $|Q_b(\theta)|$ when $b \in \mathcal{S}_2(\phi, r)$, which when doubled gives the desired upper bound D_r for $|Q_{b_1}(\theta) - Q_{b_2}(\theta)|$ when $b_1, b_2 \in \mathcal{S}_2(\phi, r)$.

If b is topmost in r then Lemma 16.7 gives us D_0. Otherwise, choose $r_1 < r$ so that b is topmost in r_1. Apply Lemma 16.8 to find an upper bound for $|\mathsf{offset}(\theta, r_1)|$ and then apply Lemma 16.7 to b and r_1 to produce D_0 and hence D.

To complete the proof of Proposition 16.6, define

$$\mathcal{X}(D) := \{\theta \in \mathcal{X} \mid |Q_{b_1}(\theta) - Q_{b_2}(\theta)| < D \text{ for all } b_1 \sim b_2 \in \mathcal{S}_2(\phi)\}.$$

Our choice of D guarantees that $\mathcal{X}(D)$ contains all $\theta \in \mathcal{X}$ that conjugate ϕ to ψ. For each i, the image of $\mathcal{X}(D)$ by

$$Q^i : \mathcal{X} \xrightarrow{Q} \mathbb{Q}^{\mathcal{S}_2(\phi)} \to \mathbb{Q}^{\mathcal{S}_2^i(\phi)}$$

is discrete, $\mathbb{Z}$-invariant, and contained in a bounded neighborhood of the diagonal in $\mathbb{Q}^{\mathcal{S}_2^i(\phi)}$. Hence the image of $\mathcal{X}(D)$ by

$$\bar{Q}^i : \mathcal{X} \xrightarrow{Q^i} \mathbb{Q}^{\mathcal{S}_2^i(\phi)} \to \mathbb{Q}^{\mathcal{S}_2^i(\phi)}/\mathbb{Z}$$

is finite and $\mathcal{X}(D)$ is contained in finitely many cosets of $\mathsf{Ker}(\bar{Q}^i)$ and so also in finitely many cosets of $\mathsf{Ker}(\bar{Q})$.

To get representatives of these cosets we must find, for each $\bar{q} \in \bar{Q}(\mathcal{X}(D))$, an element of $\mathcal{X} \cap \bar{Q}^{-1}(\bar{q})$. For this, it is enough to express $\bar{q}$ as a word in the $\bar{Q}$-image of the finite generating set $\mathcal{G}_{\mathcal{X}}$ for $\mathcal{X} = \mathsf{Out}_{\mathsf{J}}(F_n)$ supplied by Lemma 11.2. To accomplish this, we find a finite subset $S \subset \mathbb{Q}^{\mathcal{S}_2(\phi)}$ whose image in $\bar{\mathbb{Q}}^{\mathcal{S}_2(\phi)}$ covers $\bar{Q}(\mathcal{X}(D))$ and then express the elements of S in terms of the $Q(\mathcal{G}_{\mathcal{X}})$. To find S, we first find finite $S^i \subset \mathbb{Q}^{\mathcal{S}_2^i(\phi)}$ whose image in $\mathbb{Q}^{\mathcal{S}_2^i(\phi)}/\mathbb{Z}$ covers $\bar{Q}^i(\mathcal{X}(D))$ and then take for S the direct sum of the S^i, ie

$$S := \{q \in \mathbb{Q}^{\mathcal{S}_2(\phi)} \mid \text{the projection of } q \text{ to } \mathbb{Q}^{\mathcal{S}_2^i(\phi)} \text{ is in } S^i\}.$$

We now find S^i. By definition of Q, the denominators of the coordinates of the image of Q are bounded by $\max\{m_b(\phi) \mid b \in \mathcal{S}_2(\phi)\}$. For convenience, we assume we have cleared denominators and all coordinates in the image of Q are integers. Each $\bar{q}_i$ in the image of $\bar{Q}^i$ is represented by $q_i \in \mathbb{Q}^{\mathcal{S}_2^i(\phi)}$ with first coordinate equal to 0. Hence we may then take S^i to be the set of vectors in $\mathbb{Q}^{\mathcal{S}_2^i(\phi)}$ with integer coordinates of absolute value at most D and S to be the set of vectors in $\mathbb{Q}^{\mathcal{S}_2(\phi)}$ with integer coordinates of absolute value at most D.

The desired set of coset representatives can then be taken to be $\{\theta_s \mid s \in S \cap Q(\mathcal{X}(\phi))\}$ where by definition θ_s is a choice of element of $\mathcal{X}(\phi)$ satisfying $Q(\theta_s) = s$. We compute $S \cap Q(\mathcal{X}(\phi))$ and θ_s as follows. First compute $Q(\mathcal{G}_{\mathcal{X}})$. It remains to check which elements of S can be expressed as $\mathbb{Z}$-linear combinations of elements of this $Q(\mathcal{G}_{\mathcal{X}})$ and to produce such a $\mathbb{Z}$-linear combination if it exists. For this, recall that given a finite set of vectors in $\mathbb{Z}^N$ it is standard (see for example [Veblen and Franklin 1921]) to find compatible bases B_0 for the free $\mathbb{Z}$-submodule they generate and B for $\mathbb{Z}^N$. (B_0 and B are *compatible* if there is a subset $\{b_m\}$ of B and integers n_m such that $B_0 = \{n_m \cdot b_m\}$.) Without concern for efficiency, write each element of S in terms of B and check using divisibility of coordinates if it can be written in terms of B_0. □

16.1 Proof of Lemma 16.7

Lemma 16.9 *Assume that*:

1. $\theta \in \mathcal{X}$ *conjugates* ϕ *to* ψ.
2. $b \in \mathcal{S}_2(\phi, r)$ *and* $b' = \theta(b) \in \mathcal{S}_2(\psi, r')$, *where* $r' = \theta(r)$.
3. *We are given*
   - (a) *a lift* $\tilde{r}$ *of* r *and a lift* $\tilde{b}$ *of* b *that is visible in* $\tilde{r}$ *with index* i, *and*
   - (b) *a lift* $\tilde{r}'$ *of* r' *and a lift* $\tilde{b}'$ *of* b' *that is visible in* $\tilde{r}'$ *with index* i'

   such that $\Theta(\tilde{b}) = \tilde{b}'$, *where* Θ *is the unique automorphism representing* θ *and satisfying* $\Theta(\tilde{r}) = \tilde{r}'$.

Then one can compute $m_b(\theta)$ *up to an additive constant that is independent of* θ.

Proof We give a formula for $m_b(\theta)$ up to an error of at most one in terms of quantities s and s' (defined below) and then show how to compute s and s', up to a uniform additive constant, from i and i'.

Let $\tilde{b} = (\tilde{L}_1, \tilde{L}_2)$, where $\tilde{L}_1 = \tilde{\ell}_{i-1}$ and $\tilde{L}_2 = \tilde{\ell}_i$ or $\tilde{\ell}_{i+1}$ and let $\tilde{A}$ be the common axis for $\tilde{b}$. By Corollary 5.17, $\tilde{A}$ projects to a twist path w and we assume that the orientation on $\tilde{A}$ agrees with that of w. Similarly, $\tilde{b}' = (\tilde{L}'_1, \tilde{L}'_2)$, where $\tilde{L}'_1 = \tilde{\ell}'_{i'-1}$ and $\tilde{L}'_2 = \tilde{\ell}'_{i'}$ or $\tilde{\ell}'_{i'+1}$, and $\tilde{A}'$ is the common axis for $\tilde{b}'$. Let $\tilde{x}_1$ be the nearest point on $\tilde{A}$ to the initial end $\partial_-\tilde{L}_1$. (See Figure 7.) If L_1 is not a linear staple then $\mathsf{H}_{\phi,\mathfrak{c}}(L_1) = [F(\partial_-\tilde{L}_1), \partial_+\tilde{L}_1]$. In this case, the ray from $\tilde{x}_1$ to $\partial_-\tilde{L}_1$ contains an edge $\tilde{\sigma}_{i-1}$ of height greater than that of $F(\partial_-\tilde{L}_1)$ and so $\tilde{x}_1$ is the nearest point on $\tilde{A}$ to any point in $F(\partial_-\tilde{L}_1)$.

By hypothesis, $\theta(L_1) = L'_1$. Since $\theta \in \mathcal{X}$, it follows that $L'_1 \in \theta(\mathsf{H}_{\phi,\mathfrak{c}}(L_1)) = \mathsf{H}_{\phi,\mathfrak{c}}(L_1)$. Choose a homotopy equivalence $h\colon G' \to G$ that preserves markings. If L_1 is linear then $\mathsf{H}_{\phi,\mathfrak{c}}(L_1) = \{L_1\}$ so $L'_1 = L_1$. In this case, we let $\tilde{h}^1\colon \tilde{G}' \to \tilde{G}$ be the lift of h satisfying $\tilde{h}^1(\tilde{L}'_1) = \tilde{L}_1$. If L_1 is not linear then $\mathsf{H}_{\phi,\mathfrak{c}}(L_1) = [F_{\mathfrak{c}}(\partial_-\tilde{L}_1), \partial_+\tilde{L}_1]$. In this case, we let $\tilde{h}^1\colon \tilde{G}' \to \tilde{G}$ be the lift of h satisfying $\tilde{h}^1_\#(\tilde{L}'_1) \in (\partial F_{\mathfrak{c}}(\partial_-\tilde{L}_1), \partial_+\tilde{L}_1)$. Let Θ_1 be the unique lift of θ satisfying $\Theta_1(\tilde{L}_1) = \tilde{h}^1_\#(\tilde{L}'_1) \in \mathsf{H}_{\phi,\mathfrak{c}}(\tilde{L}_1)$.

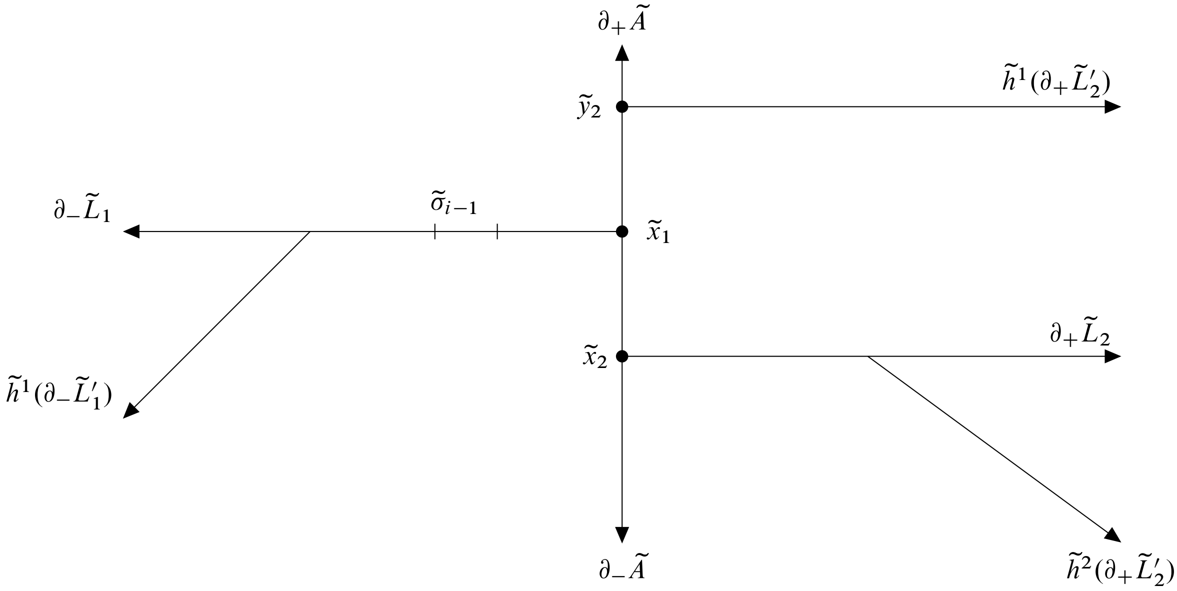

Figure 7

Let $\widetilde{x}_2$ and $\widetilde{y}_2$ be the nearest points on $\widetilde{A}$ to the terminal ends $\partial_+\widetilde{L}_2$ and $\widetilde{h}^1_\#(\partial_+\widetilde{L}'_2)$, respectively. Arguing as above, there is a lift $\widetilde{h}^2\colon \widetilde{G}' \to \widetilde{G}$ such that $\widetilde{h}^2_\#(\widetilde{A}') = \widetilde{A}$ and such that $\widetilde{x}_2$ is the nearest point to $\widetilde{h}^2(\partial_+\widetilde{L}'_2)$. Moreover, there is a lift Θ_2 of θ such that $\widetilde{h}^2_\#(\widetilde{L}'_2) = \Theta_2(\widetilde{L}_2) \in (\partial_- L_2, \partial F_{\mathrm{c}}(\partial_+\widetilde{L}_2))$. It follows from Definition 15.36 that the oriented path $\widetilde{\alpha} \subset \widetilde{A}$ from $\widetilde{x}_2$ to $\widetilde{y}_2$ has the form $\widetilde{w}^{m_b(\theta)}$.

Let $\widetilde{\beta}$ and $\widetilde{\beta}'$ be the paths in $\widetilde{A}$ connecting $\widetilde{x}_2$ to $\widetilde{x}_1$ and $\widetilde{x}_1$ to $\widetilde{y}_2$, respectively. Let s and s' be the number of complete copies of $\widetilde{w}$ (counted with orientation) crossed by the paths $\widetilde{\beta}$ and $\widetilde{\beta}'$, respectively. Then $|m_b(\theta) - (s' + s)| \le 1$.

Determining s from the index i is straightforward. We consider the cases of Lemma 15.26. In case (1), $\sigma_i \rho_i \sigma_{i+1}$ is quasi-exceptional and $\rho_i = w^s$, where w is the twist path for σ_i and $\overline{\sigma}_{i+1}$. In case (a), $\sigma_i = E' w^s \overline{E}''$ for some $E', E'' \in \mathsf{Lin}_w(f)$. In case (b), $\sigma_i = E'$ is linear with twist path w and ℓ_i is not periodic. If ρ_i has an initial segment of the form w^j for some $j > 0$ then s is the maximal such j; otherwise $-s$ is the maximal $j \ge 0$ such that ρ_i has an initial segment of the form w^{-j}. In case (c), $\overline{\sigma}_{i+1} = E'$ is linear with twist path w and ℓ_{i-1} is not periodic. In this case s is determined by the maximal initial segment of $\overline{\rho}_i$ of the form $w^{\pm j}$ as in the case (b).

Let $\widetilde{x}'_1$ and $\widetilde{x}'_2$ be the nearest points on $\widetilde{A}'$ to $\partial_-\widetilde{L}'_1$ and $\partial_+\widetilde{L}'_1$, respectively. Let $\widetilde{w}'$ be a fundamental domain for the natural action of $\mathbb{Z}$ on $\widetilde{A}'$. Arguing as above, using G' in place of G, we can compute the number t' of complete copies of $\widetilde{w}'$ (counted with orientation) crossed by the path connecting $\widetilde{x}'_2$ to $\widetilde{x}'_1$. We can also compute the bounded cancellation constant C' for h; see [Cooper 1987], also [Bestvina et al. 1997, Lemma 3.1]. Since $|s' - t'| < 2C'$, $m_b(\theta) = t' + s$ up to the additive constant $C = 2C' + 1$. □

Proof of Lemma 16.7 By Lemma 15.31 and Remark 15.32 applied to ψ, $b' = \theta(b)$ and $r' = \theta(r)$ we can find B' so that each lift $\widetilde{b}' \in \mathcal{S}_2(\psi, \widetilde{r}')$ of b' that satisfies $\widetilde{\ell}'_{B'} \prec \widetilde{b}'$ is visible in $\widetilde{r}'$. After increasing B' if necessary, we may assume that $\widetilde{\ell}'_{B'}$ is topmost in $\widetilde{r}'$. Now apply Lemma 15.31 to find a lift $\widetilde{b}_0 \in \mathcal{S}_2(\phi, \widetilde{r})$ of b. Using the given upper bound C on $|\mathsf{offset}(\theta, r)|$, choose $q \geq 0$ so that $\widetilde{\ell}'_{B'} \prec \theta(\Phi^q \widetilde{b}_0)$. Let $\widetilde{b} = \Phi^q \widetilde{b}_0$. From C and q we can compute an upper bound I' for the index of $\theta(\widetilde{b})$. By Lemma 15.31, we can list all visible $\widetilde{b}'$ with index at most I' and so have finitely many candidates for $\theta(\widetilde{b})$. Applying Lemma 16.9 to b and each of these candidates gives us the desired upper bound for $m_b(\theta)$. □

16.2 Stabilizing a ray

Suppose that E_i is the unique edge of height $i > 0$ and that $\sigma \subset G$ is a path with height i whose endpoints, if any, are not contained in the interior of E_i. Recall from Definition 4.1.3 and Lemma 4.1.4 of [Bestvina et al. 2000] that σ has a unique splitting, called the *highest-edge splitting of* σ, whose splitting vertices are the initial endpoints of each occurrence of E_i in σ and the terminal vertices of each occurrence of $\overline{E}_i$ in σ. In particular, each term in the splitting has the form $E_i \gamma \overline{E}_i$, $E_i \gamma$, $\gamma \overline{E}_i$ or γ for some $\gamma \subset G_{i-1}$.

The following lemma is used in the proof of Lemma 16.8. We make implicit use of [Feighn and Handel 2011, Lemma 4.6] which states that if $f: G \to G$ is completely split and a path $\sigma \subset G$ is completely split then $f^k_\#(\sigma)$ is completely split for all $k \geq 0$.

Lemma 16.10 *Suppose that* $f: G \to G$ *is a* CT *representing* ϕ, *that the edge* E *corresponds to some* $r \in \mathcal{R}(\phi)$, *that* ξ *is a finite subpath with endpoints at vertices and that* $R = [\xi R_E]$. *Equivalently,* $R = \tau R_1$ *for some finite path* τ *with endpoints at vertices and some subray* R_1 *of* R_E. *Then there exists a computable* $k \geq 0$ *such that* $f^k_\#(R)$ *is completely split.*

Proof The proof is by induction on the height h of R, with the base case being vacuous because the lowest stratum in the filtration is a fixed loop.

We are free to replace R by an iterate $f^l_\#(R)$ whenever it is convenient. We also have a less obvious replacement move.

(1) If there is a splitting $R = \nu \cdot R'$ where ν has endpoints at vertices then we may replace R by R'.

This follows from:

- [Feighn and Handel 2011, Lemma 4.25] For any finite path ν with endpoints at vertices, $f^k_\#(\nu)$ is completely split for all sufficiently large k.
- [Feighn and Handel 2011, Lemma 4.11] If a path σ has a decomposition $\sigma = \sigma_1 \sigma_2$ with σ_1 and σ_2 completely split and the turn $(\overline{\sigma}_1, \sigma_2)$ legal then $\sigma = \sigma_1 \sigma_2$ is a complete splitting.
- One can check if a given finite path with endpoints at vertices has a complete splitting (because there are only finitely many candidate decompositions).

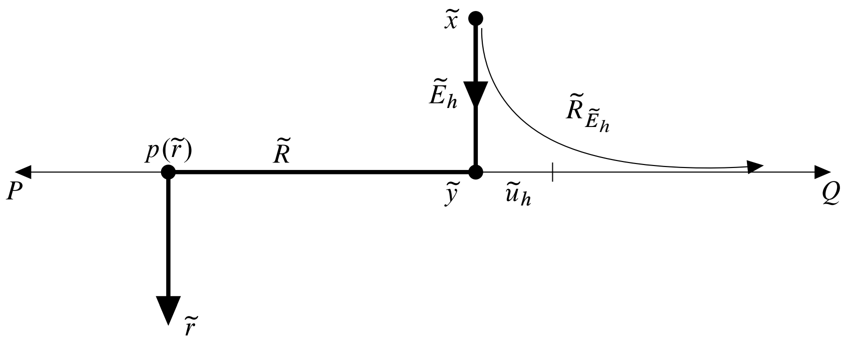

Figure 8

Let h_1 be the height of R_1. Each splitting vertex v for the highest-edge splitting of R_1 is also a splitting vertex for the complete splitting of R_E and so determines a splitting of R_1 into a finite initial subpath followed by a completely split terminal ray γ. If $h = h_1$, then v determines a splitting of R into a finite initial subpath followed by γ. In this case, an application of (1) completes the proof.

We may therefore assume that $h_1 < h$ and so the highest-edge splitting of R is finite. Applying (1), we may assume that the highest-edge splitting of R has just one term. Thus $R = E_h \mu R_1$, where E_h is the unique edge with height h and μ has height less than h. Let $h_2 < h$ be the height of $R_2 = \mu R_1$. (At various stages of the proof, we will let R_2 be the ray obtained from R by removing its initial edge. The exact edge description of R_2 will vary with the context.)

Let u_h be the path satisfying $f(E_h) = E_h \cdot u_h$. If the height of u_h is $> h_2$ then $f_\#(R) = E_h \cdot [u_h f_\#(R_2)]$ is a splitting so we may replace R by $[u_h f_\#(R_2)]$, which has height less than h. In this case the induction hypothesis completes the proof. If the height of u_h is $< h_2$ and $R_2 = \sigma_1 \cdot \sigma_2 \cdot \ldots$ is the highest-edge splitting of R_2, then $R = [E_h u_h \sigma_1] \cdot \sigma_2 \cdot \ldots$ is a splitting and the same argument completes the proof. We are now reduced to the case that the height of u_h is h_2, and we make this assumption for the rest of the proof.

We claim that there exists $k \geq 0$ so that $E_h u_h f_\#(u_h)$ is an initial segment of $f_\#^k(R)$. (Note that for any given k, one can check if $E_h u_h f_\#(u_h)$ is an initial segment of $f_\#^k(R)$ and so k with this property can be computed once one knows that it exists.) Choose a lift $\widetilde{E}_h \subset \widetilde{G}$ of E_h, let Γ be the component of the full preimage of G_{h_2} that contains the terminal $\widetilde{y}$ endpoint of $\widetilde{E}_h$ and let $\widetilde{f} \colon \widetilde{G} \to \widetilde{G}$ be the lift of f that fixes the initial endpoint $\widetilde{x}$ of $\widetilde{E}_h$. Then Γ is $\widetilde{f}$-invariant and the lift of R whose first edge is $\widetilde{E}_h$ decomposes as $\widetilde{R} = \widetilde{E}_h \widetilde{R}_2$, where $\widetilde{R}_2 \subset \Gamma$ is a lift of R_2. Let $\widetilde{u}_h$ be the lift of u_h with initial endpoint $\widetilde{y}$. Then $\widetilde{f}(\widetilde{E}_h) = \widetilde{E}_h \cdot \widetilde{u}_h$ and $\widetilde{R}_{\widetilde{E}_h} \setminus \widetilde{E}_h = \widetilde{u}_h \cdot f_\#(\widetilde{u}_h) \cdot f_\#^2(\widetilde{u}_h) \cdot \ldots$ is a ray of height h_2 that converges to an attracting fixed point $Q \in \partial\Gamma$ for the action of $\widetilde{f}$ on $\partial\Gamma$. By Lemma 2.8(ii) of [Bestvina et al. 2004] there is another fixed point $P \neq Q \in \partial\Gamma$ for the action of $\partial\widetilde{f}$. The line $\overrightarrow{PQ} \subset \Gamma$ from P to Q is $\widetilde{f}_\#$-invariant and has height h_2. Let $\mathcal{V}$ be the set of highest-edge splitting vertices of $\overrightarrow{PQ}$ with the order induced by the orientation on $\overrightarrow{PQ}$. Then $\widetilde{f}_\#$ preserves the highest-edge splitting of $\overrightarrow{PQ}$ and so $\widetilde{f}$ induces an order-preserving bijection of $\mathcal{V}$. Our choice of Q guarantees that $\widetilde{f}$ moves points in $\mathcal{V}$ away from P and towards Q. Since $\widetilde{f}$ induces an order-preserving injection of the set $\mathcal{V}'$ of highest-edge splitting vertices of $\widetilde{R}_{\widetilde{E}_h} \setminus \widetilde{E}_h$ into itself, it follows that $\mathcal{V}' \subset \mathcal{V}$. To see this, note that for each $\widetilde{v}' \in \mathcal{V}'$ and all sufficiently large m, $\widetilde{f}^m(\widetilde{v}')$ is a highest-edge splitting vertex for the common terminal ray of $\overrightarrow{PQ}$ and $R_{\widetilde{E}_h} \setminus \widetilde{E}_h$ and so $\widetilde{f}^m(\widetilde{v}') \in \mathcal{V}$. Since the restriction of $\widetilde{f}^m$ to the vertex set of Γ and the restriction of $\widetilde{f}^m$ to $\mathcal{V}$ are bijections, $\widetilde{v}' \in \mathcal{V}$.

Since $\widetilde{r}$ is an attractor for $\Phi_{\widetilde{r}}$, we get $\widetilde{r} \neq P$. If $\widetilde{r} = Q$ then the lemma is obvious so we may assume that the nearest-point projection $p(\widetilde{r})$ of $\widetilde{r}$ to $\overrightarrow{PQ}$ is well-defined. The line $\overrightarrow{\widetilde{r}Q}$ intersects $\overrightarrow{PQ}$ in the ray $\overrightarrow{p(\widetilde{r})Q}$. The set of highest-edge splitting vertices of $\overrightarrow{p(\widetilde{r})Q}$ equals the intersection of the set of highest-edge splitting vertices of $\overrightarrow{PQ}$ and the set of highest-edge splitting vertices of $\overrightarrow{\widetilde{r}Q}$. It follows that the set of highest-edge splitting vertices of $\overrightarrow{p(\widetilde{f}_\#(\widetilde{r}))Q}$ is the $\widetilde{f}_\#$-image of the set of highest-edge splitting vertices of $\overrightarrow{p(\widetilde{r})Q}$. Thus $p(f_\#^k(\widetilde{r})) \to Q$ and, after replacing R by some $\widetilde{f}_\#^k(R)$, we may assume that $p(\widetilde{r})$ is contained in $f_\#^2(\widetilde{u}_h) \cdot f_\#^3(\widetilde{u}_h) \cdot \ldots$. This completes the proof of the claim.

We now fix k satisfying the conclusions of the above claim and replace R by $f_\#^k(R)$. Thus $R = E_h u_h f_\#(u_h) \cdots$ and we let $R_2 = u_h f_\#(u_h) \cdots$ be the terminal ray of R obtained by removing its initial edge. We will prove that the decomposition of R determined by the highest-edge splitting vertices of R_2 is a splitting of R. The proof then concludes as in previous cases.

We continue with the notation established in the proof of the claim. Choose $\widetilde{v} \in \mathcal{V} \cap \widetilde{u}_h$ and decompose $\widetilde{R}$ as $\widetilde{R} = \widetilde{\alpha}\widetilde{\beta}\widetilde{\gamma}$, where

$$\widetilde{\alpha} = \overrightarrow{\widetilde{x}\widetilde{v}}, \quad \widetilde{\beta} = \overrightarrow{\widetilde{v}\widetilde{f}(\widetilde{v})} \quad \text{and} \quad \widetilde{\gamma} = \overrightarrow{\widetilde{f}(\widetilde{v})\widetilde{r}}.$$

Since $\widetilde{\alpha}\widetilde{\beta}$ is a subpath of $\widetilde{E}_h \cdot \widetilde{u}_h \cdot f_\#(\widetilde{u}_h) \cdot f_\#^2(\widetilde{u}_h) \cdot \ldots$, no edges of height h_2 are canceled when $\widetilde{f}(\widetilde{\alpha}\widetilde{\beta})$ is tightened to $\widetilde{f}_\#(\widetilde{\alpha}\widetilde{\beta})$. Similarly, no edges of height h_2 are canceled when $\widetilde{f}(\widetilde{\beta}\widetilde{\gamma})$ is tightened to $\widetilde{f}_\#(\widetilde{\beta}\widetilde{\gamma})$ because $\widetilde{\beta}\widetilde{\gamma}$ is a concatenation of terms in the highest-edge splitting of $\widetilde{R}_2$. Since $\widetilde{\beta}$ contains at least one edge of height h_2, it follows that no edges of height h_2 are canceled when $\widetilde{f}(\widetilde{R}) = \widetilde{f}(\widetilde{\alpha}\widetilde{\beta}\widetilde{\gamma})$ is tightened to $\widetilde{f}_\#(R)$. This proves that the highest-edge splitting of $\widetilde{R}_2$ is a splitting of $\widetilde{R}$, as desired. □

16.3 Proof of Lemma 16.8

Recall from Notation 15.18 and Lemma 16.1 that $\mathcal{T}_{\phi,\widetilde{r}}$ is the set of topmost elements of $\Omega_{\mathsf{NP}}(\phi,\widetilde{r})$ and that $r' = \theta(r)$ and $r_1' = \theta(r_1)$ are independent of $\theta \in \mathcal{X}$ that conjugates ϕ to ψ.

Suppose that $\widetilde{r}_1$ and $\widetilde{r}_1'$ are lifts of r_1 and r_1', respectively, and that Θ is the lift of θ satisfying $\Theta(\widetilde{r}_1) = \widetilde{r}_1'$. If $\Theta(\widetilde{L}) = \widetilde{L}'$, where $\widetilde{L} \in \mathcal{T}_{\phi,\widetilde{r}_1}$ has index s and $\widetilde{L}' \in \mathcal{T}_{\psi,\widetilde{r}_1'}$ has index s', then $\mathsf{offset}(\theta, r_1) = s' - s$. We will not be able to find $\widetilde{L}$ and $\widetilde{L}'$ whose indices we know exactly but we will be able to find $\widetilde{L}$ and $\widetilde{L}'$ whose indices we know up to a uniform bound, and this is sufficient.

Before beginning the formal proof, we introduce a way to find distinguished elements of $\mathcal{T}_{\phi,\widetilde{r}_1}$.

Notation 16.11 Suppose $r_1 <_c r$ (Notation 6.1) and that $\widetilde{r}_1$ and $\widetilde{r}$ are lifts such that $\mathcal{T}_{\phi,\widetilde{r}_1} \cap \Omega_{\mathsf{NP}}(\phi, \widetilde{r}) \neq \varnothing$. The $(\widetilde{r}, \widetilde{r}_1)$-*extreme* line is the element of $\mathcal{T}_{\phi,\widetilde{r}_1} \cap \Omega_{\mathsf{NP}}(\phi, \widetilde{r})$ that is maximal in the order on $\mathcal{T}_{\phi,\widetilde{r}_1}$.

The next lemma states that extreme lines behave well with respect to conjugation.

Lemma 16.12 *Suppose that θ conjugates ϕ to ψ, that $\Theta \in \theta$, that $\widetilde{r}$ and $\widetilde{r}_1$ are lifts of $r >_c r_1$ and that $\widetilde{L}_2 \in \mathcal{T}_{\phi,\widetilde{r}_1}$ is $(\widetilde{r}, \widetilde{r}_1)$-extreme. Then $\Theta(\widetilde{L}_2)$ is $(\Theta(\widetilde{r}), \Theta(\widetilde{r}_1))$-extreme.*

Proof This follows from Lemmas 15.21 and 15.3, which imply that Θ maps $\mathcal{T}_{\phi,\widetilde{r}_1}$ to $\mathcal{T}_{\psi,\Theta(\widetilde{r}_1)}$ preserving order, and maps $\Omega_{\mathsf{NP}}(\phi, \widetilde{r})$ to $\Omega_{\mathsf{NP}}(\psi, \Theta(\widetilde{r}))$. □

Proof of Lemma 16.8 If C is an upper bound for $|\mathsf{offset}(\theta, r)|$, it suffices to find, for each $|c| \le C$, an upper bound $C_{1,c}$ for $|\mathsf{offset}(\theta, r_1)|$ assuming that $\mathsf{offset}(\theta, r) = c$. The desired upper bound C_1 for $|\mathsf{offset}(\theta, r_1)|$ is then $\max\{C_{1,c}\}$. Going forward we may therefore assume that we know $|\mathsf{offset}(\theta, r)|$ exactly.

There is no loss of generality in assuming $r_1 <_c r$. Let E and E_1 be the elements of $\mathcal{E}_f$ corresponding to r and r_1, respectively. We will assume that E_1 occurs in R_E; the remaining case, in which $\overline{E}_1$ but not E_1 itself occurs in R_E, is similar and is left to the reader. Recall from Notation 15.18 that the visible elements of $\mathcal{T}_{\phi,\widetilde{r}}$ are enumerated $\widetilde{L}_1, \widetilde{L}_2, \dots$. For $j \ge 0$, define $q_j \ge 0$ by $\widetilde{L}_j = \widetilde{\ell}_{q_j}$ and so $\widetilde{L}_j = (\widetilde{R}^-_{q_j})^{-1} \widetilde{\rho}_{q_j} \widetilde{R}^+_{q_j+1}$.

The first step of the proof is to show that:

(a) There is a computable $J > 0$ so that if $j \ge J$ and if $\widetilde{r}_{1,j} := \partial_+ \widetilde{L}_j$ is a lift of r_1 (equivalently, $\sigma_{q_j+1} = E_1$ and $R^+_{q_j+1} = R_{E_1}$), then the line S_j connecting $\widetilde{r} = \partial_+ \widetilde{R}_{\widetilde{E}}$ to $\widetilde{r}_{1,j}$ is completely split. See Figure 9.

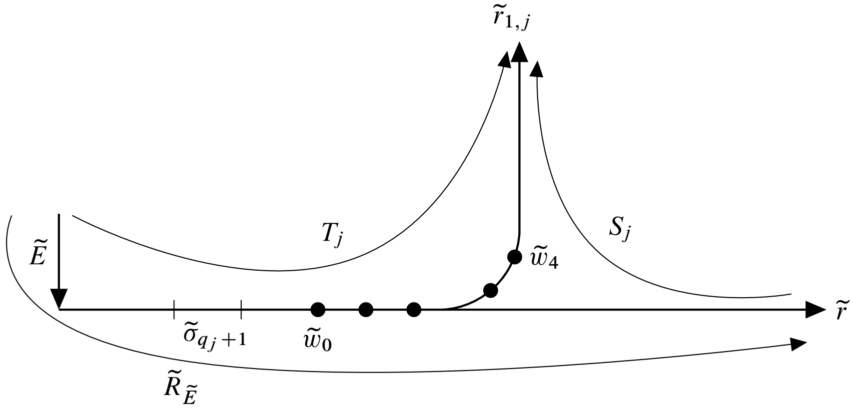

Figure 9

Lemma 15.20 implies that $\widetilde{f}^k_\#(\widetilde{L}_j) = \widetilde{L}_{j+k\tau(\phi,r)}$ and hence $\widetilde{f}^k_\#(S_j) = S_{j+k\tau(\phi,r)}$. It therefore suffices to show that for each $0 \le j \le \tau(\phi,r)$, there is a computable $K \ge 0$ so that $\widetilde{f}^K_\#(S_j)$ (and hence $\widetilde{f}^k_\#(S_j)$ for all $k \ge K$) is completely split.

The line S_j decomposes as a concatenation of (the inverse of) a ray in $\widetilde{R}_{\widetilde{E}} \setminus \widetilde{E}$ and a proper subray of a lift of R_{E_1}. The height of the former is at least that of E_1 and the height of the latter is at most that of E_1. Moreover, $\widetilde{R}_{\widetilde{E}_1} \setminus \widetilde{E}_1$ has height less than that of E_1. It follows that $\widetilde{R}_{\widetilde{E}} \setminus \widetilde{E}$ and S_j have the same height and that each splitting vertex $\widetilde{v}$ for the highest-edge splitting of S_j is contained in $\widetilde{R}_{\widetilde{E}} \setminus \widetilde{E}$ and is a splitting vertex for $\widetilde{R}_{\widetilde{E}}$. Splitting S_j at one such $\widetilde{v}$ we write $S_j = \widetilde{A}_j^{-1} \cdot \widetilde{B}_j$, where A_j is a concatenation of terms in the complete splitting of R_E, and B_j has a subray in common with R_{E_1}. For all $k \ge 0$, $S_{j+k\tau(\phi,r)} = \widetilde{f}^k_\#(S_j) = \widetilde{f}^k_\#(\widetilde{A}_j^{-1}) \cdot \widetilde{f}^k_\#(\widetilde{B}_j)$. By Lemma 16.10, we can find K so that $f^K_\#(B_j)$ is completely split. It follows — see the second bullet point in the proof of Lemma 16.10 — that $\widetilde{f}^K_\#(S_j)$ is completely split. This completes the first step.

Let $\mathcal{T}_{\psi,\widetilde{r}'} = \{\widetilde{L}'_1, \widetilde{L}'_2, \dots\}$ be the set of topmost elements of $\Omega_{\mathsf{NP}}(\psi, \widetilde{r}')$. By definition, $\Theta(\widetilde{L}_j) = \widetilde{L}'_{j+\mathsf{offset}(\theta,r)}$. The following ψ and r' analogue of (a) is verified by the same arguments given for (a):

- (b) There is a computable $J' > 0$ such that if $j \ge J'$ and if $\widetilde{r}'_{1,j} := \partial_+ \widetilde{L}'_j$ projects to r'_1, then the line S'_j connecting $\widetilde{r}'$ to $\widetilde{r}'_{1,j}$ is completely split.

Note also that:

- (c) For all $j \ge 1$, the line T_j connecting the initial vertex of $\widetilde{R}_{\widetilde{E}}$ to $\widetilde{r}_{1,j}$ is completely split, and similarly for the line T'_j connecting the initial vertex of $\widetilde{R}'_{\widetilde{E}'}$ to $\partial_+ \widetilde{L}'_j$.

For $j \ge 0$, let $\mathcal{V}_j$ be the set of highest-edge splitting vertices of $\widetilde{R}^+_{q_j+1} \setminus \widetilde{\sigma}_{q_j+1}$ (which is a terminal ray of $\widetilde{L}_j$) and $\mathcal{V}'_j$ be the set of highest-edge splitting vertices of $\widetilde{R}'^+_{q'_j+1} \setminus \widetilde{\sigma}'_{q'_j+1}$. The second step of the proof is to choose an index j so that the following four properties are satisfied:

- (i) S_j is completely split.
- (ii) There exist $\widetilde{w} \in \mathcal{V}_j$ such that $\widetilde{w}, \widetilde{f}(\widetilde{w}), \widetilde{f}^2(\widetilde{w}) \in \widetilde{R}^+_{q_j+1} \cap \widetilde{R}_{\widetilde{E}}$.
- (iii) Letting $j' = j + \mathsf{offset}(\theta, r)$, the line $S'_{j'}$ is completely split.
- (iv) There exist $\widetilde{w}' \in \mathcal{V}'_{j'}$ such that $\widetilde{w}', \widetilde{g}(\widetilde{w}'), \widetilde{g}^2(\widetilde{w}') \in \widetilde{R}'^+_{q'_{j'}+1} \cap \widetilde{R}'_{\widetilde{E}'}$.

Items (i) and (iii) hold for all $j \ge \max\{J, J' - \mathsf{offset}(\theta, r)\}$. For (ii), write $j = a_j + c_j\tau(\phi,r)$, where $0 \le a_j < \tau(\phi,r)$. Then $\widetilde{L}_j = \widetilde{f}^{c_j}_\#(\widetilde{L}_{a_j})$, and $\widetilde{L}_j \cap \widetilde{R}_{\widetilde{E}}$ contains an initial segment of $\widetilde{R}^+_{q_j+1}$ whose length goes to infinity with j. If c_j is sufficiently large then (ii) is satisfied. Item (iv) is established in the same way, completing the second step.

We have $\Theta(\widetilde{r}) = \widetilde{r}'$ and $\Theta(\widetilde{L}_j) = \widetilde{L}'_{j'}$. The latter implies that $\Theta(\widetilde{r}_{1,j}) = \widetilde{r}'_{1,j'}$. Lemma 16.12 implies that Θ maps the $(\widetilde{r}, \widetilde{r}_{1,j})$-extreme line to the $(\widetilde{r}', \widetilde{r}'_{1,j'})$-extreme line. Let s_j be the index of the $(\widetilde{r}, \widetilde{r}_{1,j})$-extreme line (as an element of $\mathcal{T}_{\widetilde{r}_{1,j}}$) and let $s'_{j'}$ be the index of the $(\widetilde{r}', \widetilde{r}'_{1,j'})$-extreme line (as an element of $\mathcal{T}_{\widetilde{r}'_{1,j'}}$). Then $\mathsf{offset}(\theta, r_1) = s'_{j'} - s_j$. We will complete the proof by finding $a_j \le s_j \le b_j$ such that

$b_j - a_j \le 3\tau(\phi, r_1)$, and $a'_{j'} \le s'_{j'} \le b'_{j'}$ such that $b'_{j'} - a'_{j'} \le 3\tau(\psi, r'_1) = 3\tau(\phi, r_1)$. These allow us to compute $\mathsf{offset}(\theta, r_1)$ with an error at most $6\tau(\phi, r_1)$ and hence compute an upper bound for $\mathsf{offset}(\theta, r_1)$.

Let h_2 be the height of $R_{E_1} \setminus E_1$ (which is the same as the height of $\widetilde{R}^+_{q_j+1} \setminus \widetilde{\sigma}_{q_j+1}$) and let E_2 be the unique edge of height h_2. We claim that:

(d) Each $\widetilde{w} \in \mathcal{V}_j \cap \widetilde{R}_{\widetilde{E}}$ is a splitting vertex for the complete splittings of T_j and $\widetilde{R}_{\widetilde{E}}$.

It suffices to show that $\widetilde{w}$ is not contained in the interior of a term $\widetilde{\mu}$ in one of those splittings. Such a $\widetilde{\mu}$ would be an indivisible Nielsen path or exceptional path with height $\ge h_2$ and whose first edge is contained in $\widetilde{R}^+_{q_j+1} \setminus \widetilde{\sigma}_{q_j+1}$ (because $\widetilde{\sigma}_{q_j+1}$ is a term in both splittings) and so has height at most h_2. Thus E_2 would be a linear edge with twist path w_2 and μ would have one of the following forms: $E_2 w_2^p \bar{E}_2$, $E_2 w_2^p \bar{E}_3$ or $E_3 w_2^p \bar{E}_2$, where $p \ne 0$ and where $E_3 \ne E_2$ is a linear edge of height $< h_2$ with twist path w_2. In none of these cases does the interior of μ contain a vertex that is the initial endpoint of E_2 or the terminal endpoint of $\bar{E}_2$. This completes the proof of (d).

A similar analysis shows that:

(e) Each $\widetilde{w} \in \mathcal{V}_j$ that is disjoint from $\widetilde{R}_{\widetilde{E}}$ is a splitting vertex for the complete splittings of S_j and T_j.

Let $\widetilde{w}_0$ be the last element of $\mathcal{V}_j$ such that $\widetilde{w}_1 = \widetilde{f}(\widetilde{w}_0)$ and $\widetilde{w}_2 = \widetilde{f}^2(\widetilde{w}_0)$ are contained in $\widetilde{R}_{\widetilde{E}}$ (and hence contained in $\mathcal{V}_j \cap \widetilde{R}_{\widetilde{E}}$). Item (d) implies that the path $\widetilde{\alpha}$ connecting $\widetilde{w}_0$ to $\widetilde{w}_2$ inherits the same complete splitting from $\widetilde{R}_{\widetilde{E}}$ and from $\widetilde{R}^+_{q_j+1}$. Thus the lift $\widetilde{\sigma}_a$ of E_2 or $\bar{E}_2$ with endpoint $\widetilde{w}_1$ determines an element $\widetilde{L}^1$ of $\mathcal{T}_{\phi,\widetilde{r}_{1,j}} \cap \Omega_{\mathsf{NP}}(\phi, \widetilde{r})$. (If $\widetilde{\sigma}_a$ is a lift of E_2 then $\widetilde{w}_1$ is the initial endpoint of $\widetilde{\sigma}_a$ and $\widetilde{L}^1 = \widetilde{\ell}_{a-1}$; if $\widetilde{\sigma}_a$ is a lift of $\bar{E}_2$ then $\widetilde{w}_1$ is the terminal endpoint of $\widetilde{\sigma}_a$ and then $\widetilde{L}^1 = \widetilde{\ell}_{a_1}$.) In particular, the index s_j of the $(\widetilde{r}, \widetilde{r}_{1,j})$-extreme line (as an element of $\mathcal{T}_{\phi,\widetilde{r}_{1,j}} \cap \Omega_{\mathsf{NP}}(\phi, \widetilde{r})$) is at least as big as that of $\widetilde{L}^1$.

Let $\widetilde{w}_3 = \widetilde{f}^3(\widetilde{w}_0)$ and $\widetilde{w}_4 = \widetilde{f}^4(\widetilde{w}_0)$, neither of which is contained in $\widetilde{R}_{\widetilde{E}}$. The lift of E_2 or $\bar{E}_2$ with endpoint $\widetilde{w}_4$ determines an element $\widetilde{L}^4$ in $\mathcal{T}_{\phi,\widetilde{r}_{1,j}}$. Item (e) and the hard splitting property of a complete splitting (Lemma 4.11 of [Feighn and Handel 2011]) implies that no point in the terminal ray of $\widetilde{R}^+_{q_j+1}$ that begins with $\widetilde{w}_4$ is ever identified, under iteration by $\widetilde{f}$, with a point in $\widetilde{R}_{\widetilde{E}}$. It follows that $\widetilde{L}^4$ is not an element of $\Omega_{\mathsf{NP}}(\phi, \widetilde{r})$ and so s_j is less than the index of $\widetilde{L}^4$.

Combining the inequalities established in the preceding two paragraphs we are able to compute s_j with an error of at most $3\tau(\phi, r_1)$. The parallel argument allows us to compute the index s'_j of the $(\widetilde{r}', \widetilde{r}'_{1,j'})$-extreme line (as an element of $\mathcal{T}_{\phi,\widetilde{r}_{1,j}}$) with an error of at most $3\tau(\psi, r'_1) = 3\tau(\phi, r_1)$. As noted above, this completes the proof. □

17 Proof of Proposition 16.4

Some of our arguments are by induction up through the elements $\mathcal{F}_k$ of the chain $\mathfrak{c}$. We write $\phi|\mathcal{F}_k = \psi|\mathcal{F}_k$ if $\phi|[F] = \psi|[F]$ for each component $[F]$ of $\mathcal{F}_k$. Similarly, we say $\theta|\mathcal{F}_k$ conjugates $\phi|\mathcal{F}_k$ to $\psi|\mathcal{F}_k$

if $\phi^\theta|\mathcal{F}_k = \psi|\mathcal{F}_k$. If G_s is the core filtration element corresponding to $\mathcal{F}_k$ and if C is a component of G_s with rank one, then $[C]$ is a component of $\mathcal{F}_0$ and we define $\Gamma(f|C) = C$. With this convention, $\Gamma(f|G_s)$ is the disjoint union $\bigsqcup \Gamma(f|C_i)$ as C_i varies over the components of G_s. (See Section 4.1.)

We show below that Proposition 16.4 is a consequence of the following lemma and proposition. The former addresses the restrictions to $\mathcal{F}_0$ and the latter provides the inductive step for the higher-order one-edge extensions.

Lemma 17.1 *Suppose that $\phi, \psi \in \mathsf{UPG}(F_n)$ share the special chain $\mathfrak{c}$ and satisfy $\mathsf{I}_\mathfrak{c}(\phi) = \mathsf{I}_\mathfrak{c}(\psi)$. Let $\mathcal{F}_0 = \mathcal{F}_0(\phi) = \mathcal{F}_0(\psi)$. Then*

1. *$\theta(L) = L$ for each $\theta \in \mathcal{X}$ and each $L \in \Omega(\phi)$ that is carried by $\mathcal{F}_0$, and*
2. *if there exists $\theta_0 \in \mathcal{X}$ such that $\phi^{\theta_0}|\mathcal{F}_0 = \psi|\mathcal{F}_0$, then $\phi|\mathcal{F}_0 = \psi|\mathcal{F}_0$ and $\phi^\theta|\mathcal{F}_0 = \psi|\mathcal{F}_0$ for all $\theta \in \mathcal{X}$.*

Proposition 17.2 *Suppose that $\phi, \psi \in \mathsf{UPG}(F_n)$ share the special chain $\mathfrak{c}$ and satisfy $\mathsf{I}_\mathfrak{c}(\phi) = \mathsf{I}_\mathfrak{c}(\psi)$, and that the special one-edge extension $\mathfrak{e} = (\mathcal{F}^- \sqsubset \mathcal{F}^+)$ in $\mathfrak{c}$ satisfies*

1. $\phi|\mathcal{F}^- = \psi|\mathcal{F}^-$,
2. $\{L \in \Omega(\phi) \mid L \subset \mathcal{F}^-\} = \{L' \in \Omega(\psi) \mid L' \subset \mathcal{F}^-\}$.

Then there is an algorithm to decide if there exists $\theta \in \mathsf{Ker}(\bar{Q}) < \mathcal{X}$ such that the following are satisfied:

3. $\phi^\theta|\mathcal{F}^+ = \psi|\mathcal{F}^+$,
4. $\theta(\{L \in \Omega(\phi) \mid L \subset \mathcal{F}^+\}) = \{L' \in \Omega(\psi) \mid L' \subset \mathcal{F}^+\}$.

Moreover, *if such an element θ exists*, *then one is produced.*

Before proving Lemma 17.1 and Proposition 17.2, we use them to prove Proposition 16.4.

Proof of Proposition 16.4, assuming Lemma 17.1 and Proposition 17.2 If $\phi|\mathcal{F}_0 \neq \psi|\mathcal{F}_0$, then no element of $\mathcal{X}$ conjugates ϕ to ψ by Lemma 17.1(2) so we return `NO` and `STOP`. Otherwise, $\phi^\theta|\mathcal{F}_0 = \psi|\mathcal{F}_0$ for all $\theta \in \mathcal{X}(\phi)$, and we define θ_0 =identity and $\psi_0 = \psi$.

Suppose $\mathfrak{c} = (\mathcal{F}_0 \sqsubset \mathcal{F}_1 \sqsubset \cdots \sqsubset \mathcal{F}_t)$. Apply Proposition 17.2 with $(\phi, \psi_0, \mathcal{F}_0, \mathcal{F}_1)$ in place of $(\phi, \psi, \mathcal{F}^-, \mathcal{F}^+)$. If the 17.2-algorithm returns `NO` then there is no θ as in the conclusion of Proposition 16.4 because any such θ would satisfy items (3) and (4) of Proposition 17.2; we return `NO` and `STOP`. Otherwise, Proposition 17.2 gives us an element $\theta_1 \in \mathsf{Ker}(\bar{Q})$. Letting $\psi_1 = \psi_0^{(\theta_1^{-1})}$ we have that $\phi|\mathcal{F}_1 = \psi_1|\mathcal{F}_1$ and $\{L \in \Omega(\phi) \mid L \subset \mathcal{F}_1\} = \{L' \in \Omega(\psi_1) \mid L' \subset \mathcal{F}_1\}$. From $\theta_1 \in \mathcal{X}$ and Lemma 14.6, it follows that $\mathsf{I}(\phi) = \mathsf{I}(\psi_1)$.

Apply Proposition 17.2 with $(\phi, \psi_1, \mathcal{F}_1, \mathcal{F}_2)$ in place of $(\phi, \psi, \mathcal{F}^-, \mathcal{F}^+)$. Suppose that the 17.2-algorithm returns `NO`. Then there are no elements of $\mathsf{Ker}(\bar{Q})$ that conjugate $\phi|\mathcal{F}_2$ to $\psi_1|\mathcal{F}_2$, and so also no elements of $\mathsf{Ker}(\bar{Q})$ that conjugate ϕ to ψ_1. It follows also then that there are no elements θ of $\mathsf{Ker}(\bar{Q})$ that conjugate ϕ to ψ. Indeed for such a θ, $\theta_1^{-1}\theta$ would conjugate ϕ to ψ_1. We therefore return `NO` and `STOP`. Otherwise, Proposition 17.2 gives us an element $\theta_2 \in \mathsf{Ker}(\bar{Q})$. Letting $\psi_2 = \psi_1^{(\theta_2)^{-1}}$ we have that $\phi|\mathcal{F}_2 = \psi_2|\mathcal{F}_2$ and $\{L \in \Omega(\phi) : L \subset \mathcal{F}_2\} = \{L' \in \Omega(\psi_2) : L' \subset \mathcal{F}_2\}$. As in the previous case, $\mathsf{I}(\phi) = \mathsf{I}(\psi_2)$. Repeat this until either some application of Proposition 17.2 returns `NO` or until we reach $\psi_t = \psi^{(\theta_1 \dots \theta_t)^{-1}}$ satisfying $\phi = \phi|\mathcal{F}_t = \psi_t|\mathcal{F}_t = \psi_t$. In the former case there is no θ as in the conclusion of Proposition 16.4 and we return `NO` and `STOP`. In the latter case $\theta = \theta_1 \dots \theta_t$ conjugates ϕ to ψ and is an element of $\mathsf{Ker}(\bar{Q})$; we return `YES` and θ and then `STOP`. □

Proof of Lemma 17.1 If $L \in \Omega(\phi)$ is carried by $\mathcal{F}_0$, then the ends of L are periodic. If L is periodic then $\widetilde{L} = a^\infty$ for some $[a] \in \mathcal{A}(\phi)$; see Corollary 5.17(1). By definition of $\mathcal{X}$, $\theta([a]) = [a]$ and so $\theta(L) = L$. Otherwise, $L \in \Omega_{\mathsf{NP}}(\phi)$ has type P-P, in which case $\mathsf{H}(L)$ determines L; see Section 13. Again by definition of $\mathcal{X}$, $\theta(\mathsf{H}(L)) = \mathsf{H}(L)$ and so $\theta(L) = L$. This verifies (1).

It suffices to show that if a free factor F represents a component of $\mathcal{F}_0$ then either $\phi^\theta|F = \psi|F$ for all $\theta \in \mathcal{X}(\phi)$ (and in particular for $\theta =$ identity) or $\phi^\theta|F = \psi|F$ is satisfied by no element of $\mathcal{X}(\phi)$.

Let $\phi_F = \phi|F$ and $\psi_F = \psi|F$. If F has rank one, then ϕ_F and ψ_F are both the identity because ϕ and ψ are rotationless. We may therefore assume that F has rank at least two. Since $\mathcal{R}(\phi_F) = \varnothing$, Lemma 3.9 implies that $\mathsf{Fix}_{\mathsf{N}}(\Phi_F) = \partial\,\mathsf{Fix}(\Phi_F)$ for each $\Phi_F \in \mathcal{P}(\phi_F)$. Also, $\mathsf{Fix}_{\mathsf{N}}(\Phi_F)$ contains at least three points, so $\mathsf{Fix}(\Phi_F)$ has rank at least two and $\mathsf{Fix}(\Phi_F) \ne \mathsf{Fix}(\Phi'_F)$ for $\Phi_F \ne \Phi'_F \in \mathcal{P}(\phi_F)$ by Lemma 4.4.

There is a unique $\Phi \in \mathcal{P}(\phi)$ such that $\Phi_F = \Phi|F$. From $\mathsf{I}(\phi) = \mathsf{I}(\psi)$, it follows that there exists $\Psi \in \mathcal{P}(\psi)$ such that $\mathsf{Fix}(\Phi) = \mathsf{Fix}(\Psi)$. Since $\theta \in \mathcal{X}$, there exists Θ representing θ such that $\mathsf{Fix}(\Phi)$ is Θ-invariant. It follows that $\Theta(F) \cap F$ is nontrivial and hence that $\Theta(F) = F$ (because F is a free factor and θ preserves $[F]$). Letting $\Psi_F = \Psi|F$ and $\Theta_F = \Theta|F$, we have that $\mathsf{Fix}_{\mathsf{N}}(\Phi_F) = \partial\,\mathsf{Fix}(\Phi_F) = \partial\,\mathsf{Fix}(\Psi_F) = \mathsf{Fix}_{\mathsf{N}}(\Psi_F)$ is Θ_F-invariant.Lemma 4.1 implies that the eigengraphs for ϕ_F and for ψ_F carry the same lines and that θ preserves this set of lines. Thus ϕ_F, ψ_F and θ_F satisfy the hypotheses of Lemma 4.21.

If $a \in F$ is fixed by distinct $\Phi_F, \Phi'_F \in \mathcal{P}(\phi_F)$ then $[\Phi_F, a]$ is an element of $\mathsf{SA}(\phi_F)$ and $[\Phi, a]$ is an element of $\mathsf{SA}(\phi)$. Lemma 4.21 implies that

$$[\Phi_F, a] \mapsto [\Psi_F, \Theta_F(a)]$$

defines a bijection $\mathcal{B}_{\mathsf{SA},F} : \mathsf{SA}(\phi_F) \to \mathsf{SA}(\psi_F)$ that is independent of the choice of Θ_F representing θ_F and preserving $\mathsf{Fix}_{\mathsf{N}}(\Phi_F)$. Since $\theta \in \mathcal{X}$, by Definition 14.1(6) we have $[\mathsf{Fix}(\Phi), a] = \theta([\mathsf{Fix}(\Phi), a]) = [\Theta(\mathsf{Fix}(\Phi)), \Theta(a)] = [\mathsf{Fix}(\Phi), \Theta(a)]$. Equivalently, there exists $c \in F_n$ such that $i_c(\mathsf{Fix}(\Phi)) = \mathsf{Fix}(\Phi)$ and $i_c\Theta(a) = a$. Thus $c \in \mathsf{Fix}(\Phi)$ and after replacing Θ by $i_c\Theta$ we may assume that $\Theta(a) = a$ and hence that $\Theta_F(a) = a$. We conclude that $\mathcal{B}_{\mathsf{SA},F}$ is independent of θ.

Check by inspection if $\mathcal{B}_{\mathsf{SA},F}$ preserves twist coordinates. If it does then Lemma 4.21 implies that each $\theta \in \mathcal{X}(\phi)$ conjugates ϕ_F to ψ_F; if not, then no element of $\mathcal{X}(\phi)$ conjugates ϕ_F to ψ_F. □

The rest of the paper is dedicated to the proof of Proposition 17.2.

Set $\mathfrak{c} = (\mathcal{F}_0 \sqsubset \mathcal{F}_1 \sqsubset \cdots \sqsubset \mathcal{F}_t)$ and thus $\mathfrak{e} \in \mathfrak{c}$ has the form $\mathcal{F}^- \sqsubset \mathcal{F}^+$, where $\mathcal{F}^- = \mathcal{F}_{k-1}$ and $\mathcal{F}^+ = \mathcal{F}_k$ for some $1 \le k \le t$. (We will use these notations interchangeably depending on the context.)

Definition 17.3 For $\epsilon = \pm$, $\mathcal{X}^\epsilon$ is the set of $\theta \in \mathcal{X}$ such that

(a) $\theta \in \mathsf{Ker}(\bar{Q})$, and

(b) $\theta|\mathcal{F}^\epsilon$ conjugates $\phi|\mathcal{F}^\epsilon$ to $\psi|\mathcal{F}^\epsilon$.

By the next lemma, our goal is to produce an element of $\mathcal{X}^+$ or deduce that $\mathcal{X}^+$ is empty.

Lemma 17.4 *An element $\theta \in \mathsf{Ker}(\bar{Q})$ satisfies items* (3) *and* (4) *of Proposition 17.2 if and only if $\theta \in \mathcal{X}^+$.*

Proof Comparing the definitions, it suffices to show that each $\theta \in \mathcal{X}^+$ satisfies Proposition 17.2(4); namely, $\theta(\{L \in \Omega(\phi) \mid L \subset \mathcal{F}^+\}) = \{L' \in \Omega(\psi) \mid L' \subset \mathcal{F}^+\}$. By symmetry, it suffices to show that if $L \in \Omega(\phi)$ is carried by $\mathcal{F}^+$ then $\theta(L) \in \Omega(\psi)$ is carried by $\mathcal{F}^+$. Since $\mathcal{F}^+$ is θ-invariant, it suffices to show that $\theta(L) \in \Omega(\psi)$. If L is periodic then $\widetilde{L} = a^\infty$ for some $[a] \in \mathcal{A}(\phi)$ by Corollary 5.17(1). Since $\theta \in \mathcal{X}$, one has that $\theta([a]) = [a]$ and $\theta(L) = L \in \Omega(\psi)$. Otherwise $L \in \Omega_{\mathsf{NP}}(\phi)$ and, as $\mathsf{I}_{\mathfrak{c}}(\phi) = \mathsf{I}_{\mathfrak{c}}(\psi)$, there exists $L' \in \Omega_{\mathsf{NP}}(\psi)$ such that $\mathsf{H}(L) = \mathsf{H}(L')$. Since $\theta|\mathcal{F}^+$ conjugates $\phi|\mathcal{F}^+$ to $\psi|\mathcal{F}^+$, Corollary 5.17(4) and Lemmas 4.1 and 3.16 imply that $\theta(L)$ lifts into $\Gamma(g_{u'})$. Also, $\theta(L) \in \theta\big(\mathsf{H}(L)\big) = \mathsf{H}(L) = \mathsf{H}(L')$ because $\theta \in \mathcal{X}$. Lemma 13.9 implies that $\theta(L) = L' \in \Omega_{\mathsf{NP}}(\psi)$, and we are done. □

By [Feighn and Handel 2018, Theorem 7.4] we can choose CTs $f : G \to G$ and $g : G' \to G'$ representing ϕ and ψ, respectively, such that each $\mathcal{F}_i$ is realized by a core filtration element and such that the core filtration elements of G and G' realizing $\mathcal{F}^- = \mathcal{F}_{k-1}$ are identical as marked graphs and that after identifying them to a common subgraph G_s, the restrictions $f_s = f|G_s$ and $g_s = g|G'_s$ are equal. In particular,

(1) $\Gamma(f_s) = \Gamma(g_s)$.

Before describing f_s and g_s in more detail, we record some useful properties of $\mathcal{X}^-$. We define $\mathcal{R}(\phi|\mathcal{F}_-) = \cup \mathcal{R}(\phi|[F])$ as $[F]$ varies over the components of $\mathcal{F}_-$.

Lemma 17.5 *Each $\theta \in \mathcal{X}^-$ satisfies the following properties*:

(1) *$\theta|\mathcal{F}^-$ commutes with $\phi|\mathcal{F}^- = \psi|\mathcal{F}^-$.*

(2) *θ fixes each element of $\Omega(\phi)$ that is carried by $\mathcal{F}^-$.*

(3) *θ fixes each element of $\mathcal{R}(\phi|\mathcal{F}^-)$.*

Proof Property (1) follows from Definition 17.3(b) and the hypothesis that $\phi|\mathcal{F}^- = \psi|\mathcal{F}^-$. For (2), note that if $L \in \Omega(\phi)$ is carried by $\mathcal{F}^-$ then L lifts to $\Gamma(f_s)$ by Corollary 5.17(4) so (1) and Lemmas 4.1 and 3.16 imply that $\theta(L) \in \Gamma(g_s)$. Since $\theta(L) = L$ if L is periodic and otherwise $\theta(L) \in \theta(\mathsf{H}(L)) = \mathsf{H}(L)$, (2) follows from (1) and Lemma 13.9. By (1) and Lemma 16.1, $\theta(r)$ is independent of $\theta \in \mathcal{X}^-$. Item (3) therefore follows from the fact that $\mathcal{X}^-$ contains the identity. □

Suppose that $G_u \subset G$ and $G'_{u'} \subset G'$ are the core filtration elements realizing $\mathcal{F}^+ = \mathcal{F}_k$. Let $f_u = f|G_u$ and $g_{u'} = g|G'_{u'}$. Since $\mathcal{F}^- \sqsubset \mathcal{F}^+$ is a special one-edge extension, G_u is obtained from G_s by adding a single topological arc E which is either a single edge D or is the union $E = \bar{C}D$ of a pair of edges C and D with common initial endpoint not in G_s. (We have previously denoted edges in G by E and now we are using C and D instead and using E for a topological arc. This is more convenient for the current argument and should not cause confusion.) By Lemma 6.9, there are three possibilities. In each case, there is one component $\Gamma_*(f_u)$ of $\Gamma(f_u)$ that is not a component of $\Gamma(f_s)$.

- [HH] ($E = \bar{C}D$ consists of two higher-order edges) $\Gamma(f_u)$ is obtained from $\Gamma(f_s)$ by adding a new component $\Gamma_*(f_u)$ which is a line labeled $R_C^{-1} \cdot R_D$.
- [LH] ($E = \bar{C}D$ where C is linear and D is higher-order) $\Gamma(f_u)$ is obtained from $\Gamma(f_s)$ by adding a new component $\Gamma_*(f_u)$ which is a one-point union of a lollipop corresponding to C and a ray labeled R_D.
- [H] ($E = D$, a higher-order edge) $\Gamma(f_u)$ is a one-point union of $\Gamma(f_s)$ and a ray labeled R_D. $\Gamma_*(f_u)$ is the one-point union of a component $\Gamma_*(f_s)$ of $\Gamma(f_s)$ and a ray labeled R_D.

Similarly, we can orient the topological arc E' that is added to $G_s = G'_s$ to form $G'_{u'}$ so that $\Gamma(g_{u'})$ is obtained from $\Gamma(g_s)$ in one of these three ways.

By Lemmas 6.10 and 6.13, we may assume:

(2) The extensions $\Gamma(f_s) \subset \Gamma(f_u)$ and $\Gamma(g_s) \subset \Gamma(g_{u'})$ have the same type HH, LH or H.

For if not, then $\phi|\mathcal{F}^+$ and $\psi|\mathcal{F}^+$ are not conjugate by an element of $\mathcal{X}$ so we return `NO` and `STOP`.

Remark 17.6 A vertex v in G that is new in an HH extension, is not incident to any fixed or linear edge. It therefore follows from the construction of $\Gamma(f)$ given at the beginning of Section 4.1 that the component $\Gamma(f, v)$ of $\Gamma(f)$ corresponding to v is obtained from the disjoint union of eigenrays R_E, one for each $E \in \mathcal{E}(f)$ with initial vertex v, by identifying their initial vertices. Similarly, if v is new in an LH extension, then $\Gamma(f, v)$ is the one-point union of the lollipop associated to the unique linear edge with v as initial vertex and the eigenrays R_E associated to $E \in \mathcal{E}_f$ with v as initial vertex.

Lemma 17.7 *Suppose that $\mathfrak{e}$ has type* H *and that $\mathcal{X}^+ \neq \varnothing$. Then $\Gamma_*(f_s) = \Gamma_*(g_s)$.*

Proof Assume that $\theta \in \mathcal{X}^+$. Denote the set of lines that lift into a Stallings graph Γ by $\Lambda(\Gamma)$. Lemmas 3.16 and 4.1 imply that $\theta(\Lambda(\Gamma_*(f_u))) = \Lambda(\Gamma_*(f_{u'}))$. By construction,

$$\Lambda(\Gamma_*(f_s)) = \{L \in \Lambda(\Gamma_*(f_u)) \mid L \subset G_s\} \quad \text{and} \quad \Lambda(\Gamma_*(g_s)) = \{L \in \Lambda(\Gamma_*(g_u)) \mid L \subset G_s\}.$$

Thus $\theta(\Lambda(\Gamma_*(f_s))) = \Lambda(\Gamma_*(f_s))$.

The proof now divides into cases. If $\Gamma_*(f_s)$ contains a ray corresponding to some $r \in \mathcal{R}(\phi)$ then $\Lambda(\Gamma_*(f_s))$ contains a line that ends at r. Lemma 17.5(3) then implies that $\Lambda(\Gamma_*(g_s))$ contains a line that ends at r and hence that $\Gamma_*(f_s)$ contains a ray corresponding to $r \in \mathcal{R}(\phi)$. This proves that $\Gamma_*(f_s) = \Gamma_*(g_s)$.

We may now assume that $\Gamma_*(f_s)$ is compact. If $\Gamma_*(f_s)$ has rank at least two then $\pi_1(\Gamma_*(f_s))$ is a component of $\mathsf{Fix}(\phi)$ and is hence θ-invariant. In this case, $\pi_1(\Gamma_*(f_s)) = \pi_1(\Gamma_*(g_s))$. Lemma 4.4(1) implies that $\Gamma_*(f_s) = \Gamma_*(g_s)$. The final case is that $\Gamma_*(f_s)$ has rank one and so is a topological circle labeled by a component Y of G_0 consisting of a single edge e. In this case, $\Lambda(\Gamma_*(f_s)) = \{e^\infty, e^{-\infty}\}$, which is θ-invariant. It follows that $\Lambda(\Gamma_*(g_s)) = \{e^\infty, e^{-\infty}\}$ and hence that $\Gamma_*(f_s) = \Gamma_*(g_s)$. □

We may therefore assume that:

(3) In the case H, $\Gamma_*(f_s) = \Gamma_*(g_s)$.

We next apply the recognition theorem to give criteria for an element in $\mathcal{X}^-$ to be in $\mathcal{X}^+$.

Lemma 17.8 *The following are equivalent for each* $\theta \in \mathcal{X}^-$*:*

(1) $\theta \in \mathcal{X}^+$*; equivalently,* $\theta|\mathcal{F}^+$ *conjugates* $\phi|\mathcal{F}^+$ *to* $\psi|\mathcal{F}^+$*.*

(2) (a) *A line* L *lifts into* $\Gamma(f_u)$ *if and only if* $\theta(L)$ *lifts into* $\Gamma(g_{u'})$*.*

(b) *If* $\mathcal{F}^- \sqsubset \mathcal{F}^+$ *has type* LH *then the twist index for* C *with respect to* f *equals the twist index for* C' *with respect to* g*. Equivalently, if* $f(C) = Cw^d$*, then* $g(C') = C'w^d$*.*

Proof (1) implies (a) by Lemmas 4.1 and 3.16. We may therefore assume that (a) is satisfied and prove that (1) is equivalent to (b).

If $[F]$ is a component of $\mathcal{F}^+$ that is also a component of $\mathcal{F}^-$ then $\theta|F$ conjugates $\phi|F$ to $\psi|F$ because $\theta \in \mathcal{X}^-$. We may therefore restrict our attention to the unique component of $\mathcal{F}^+$ that is not also a component of $\mathcal{F}^-$. In other words, we may assume that G_u is connected and so may assume that $G_u = G$ and $\mathcal{F}^+ = \{[F_n]\}$.

By Lemma 4.21, there is a bijection $B: \mathsf{SA}(\phi) \to \mathsf{SA}(\psi)$ that preserves twist coordinates if and only if θ conjugates ϕ to ψ. By definition of $\mathcal{X}$, θ preserves each element of $\mathcal{A}_{\mathrm{or}}(\phi)$. We are therefore reduced to showing that (b) is satisfied if and only if the following is satisfied for each $[a] \in \mathcal{A}_{\mathrm{or}}(\phi)$:

(∗) The restricted bijection $B: \mathsf{SA}(\phi, [a]) \to \mathsf{SA}(\psi, [a])$ preserves twist coordinates.

Since $(*)$ is satisfied for $[a]$ if and only if it is satisfied for $[\overline{a}]$, we may assume that the twist path w for $[a]_u$ satisfies $[a]=[w]$. Extending Notation 4.7, we define

$$\mathcal{P}(\phi, a) := \{\Phi_{a,0}, \dots, \Phi_{a,m-1}\}.$$

In particular, $\Phi_{a,0}$ is the base principal lift for a (with respect to f) and there is an order-preserving bijection between the set $\{E^1, \dots, E^{m-1}\}$ of linear edges with axis $[a]$ and $\{\Phi_{a,1}, \dots, \Phi_{a,m-1}\}$. For $1 \le j \le m-1$, there exist distinct twist indices $d_j \ne 0$ such that $f(E^j) = E^j w^{d_j}$. Define $d_0 = 0$. Lemmas 4.8 and 4.12 imply that

$$\mathsf{SA}(\phi, [a]) = \{[\Phi_{a,0}, a], \dots, [\Phi_{a,m-1}, a]\}$$

and that the twist coefficient for the pair $([\Phi_{a,i}, a], [\Phi_{a,j}, a])$ is $d_i - d_j$.

We consider two cases. In the first, we assume that either

- $\mathcal{F}^- \sqsubset \mathcal{F}^+$ has type LH and $C \notin \{E^1, \dots, E^{m-1}\}$, or
- $\mathcal{F}^- \sqsubset \mathcal{F}^+$ does not have type LH,

and we prove that $(*)$ is satisfied.

In this case,

$$\mathcal{P}(\psi, a) = \{\Psi_{a,0}, \dots, \Psi_{a,m-1}\} \quad \text{and} \quad \mathsf{SA}(\psi, [a]) = \{[\Psi_{a,0}, a], \dots, [\Psi_{a,m-1}, a]\},$$

with the same sequence of linear edges $\{E^1, \dots, E^{m-1}\}$ and the same sequence of twist indices $\{d_0, \dots, d_{m-1}\}$. The bijection $B \colon \mathsf{SA}(\phi) \to \mathsf{SA}(\psi)$ induces a permutation π of $\{0, \dots, m-1\}$ satisfying $B([\Phi_{a,i}, a]) = [\Psi_{a,\pi(i)}, a]$. We will show that π is the identity and hence that $B \colon \mathsf{SA}(\phi, [a]) \to \mathsf{SA}(\psi, [a])$ preserves twist coordinates.

Choose an automorphism Θ representing θ and fixing a. By Lemma 4.21,

$$\Theta(\mathsf{Fix}_{\mathsf{N}}(\Phi_{a,i})) = \mathsf{Fix}_{\mathsf{N}}(\Psi_{a,\pi(i)}).$$

Let C_s be the component of G_s that contains w, and hence contains each E^i, and let $[F]$ be the corresponding component of $\mathcal{F}^-$; we may assume without loss of generality that $a \in F$. Applying Notation 4.7 to $\phi|F = \psi|F$ represented by the CT $f|C_s$, we see that

$$\begin{aligned}\mathcal{P}(\psi|F, a) = \mathcal{P}(\phi|F, a) &= \{\Phi_{a,0}|F, \dots, \Phi_{a,m-1}|F\},\\ \mathsf{SA}(\psi|F, [a]) = \mathsf{SA}(\phi|F, [a]) &= \{[\Phi_{a,0}|F, a], \dots, [\Phi_{a,m-1}|F, a]\},\end{aligned}$$

with the same sequence of linear edges $\{E^1, \dots, E^{m-1}\}$ and the same sequence of twist indices $\{d_0, \dots, d_{m-1}\}$. Since C_s is f-invariant and $[F]$ is θ-invariant, (a) implies that the set of lines that lift to $\Gamma(f_s)$ is θ-invariant. Applying Lemma 4.21 produces a permutation B_F of $\mathsf{SA}(\phi|F, [a])$ and an induced permutation π_F of $\{0,\dots,m-1\}$. Since $\theta|F$ commutes with $\phi|F$, B_F preserves twist-coordinates. Thus, $d_i - d_j = d_{\pi_F(i)} - d_{\pi_F(j)}$ for all i and j. The only possibility is that π_F is the identity and so

$$\mathsf{Fix}_\mathsf{N}(\Psi_{a,\pi(i)}) \cap \partial F = \Theta(\mathsf{Fix}_\mathsf{N}(\Phi_{a,i})) \cap \partial F = (\Theta|F)(\mathsf{Fix}_\mathsf{N}(\Phi_{a,i}|F)) = \mathsf{Fix}_\mathsf{N}(\Psi_{a,i}|F) = \mathsf{Fix}_\mathsf{N}(\Psi_{a,i}) \cap \partial F.$$

It follows that $\mathsf{Fix}_\mathsf{N}(\Psi_{a,\pi(i)}) \cap \mathsf{Fix}_\mathsf{N}(\Psi_{a,i})$ contains $\mathsf{Fix}_\mathsf{N}(\Psi_{a,i}) \cap \partial F$, which has cardinality at least three. Lemma 3.7 implies that $\pi(i) = i$, as desired. This completes the first case.

For the second case, we assume that:

- $\mathcal{F}^- \sqsubset \mathcal{F}^+$ has type LH and $C \in \{E^1,\dots,E^{m-1}\}$,

and prove that $(*)$ is equivalent to (b).

Assuming without loss of generality that $C = E^{m-1}$, the sequence of linear edges for ψ is given by $\{E^1,\dots,E^{m-2},C'\}$ with twist indices $\{d_0,\dots,d_{m-2},d'_{m-1}\}$. Thus (b) is the statement that $d_{m-1} = d'_{m-1}$ and we are reduced to showing that π is the identity.

If $m > 2$ then $\mathcal{P}(\phi|F,a) = \mathcal{P}(\psi|F,a)$ is indexed by $\{E^1,\dots,E^{m-2}\}$ and the above analysis applies to show that π restricts to the trivial permutation of $\{0,\dots,m-2\}$. It then follows that π must fix the one remaining element $m-1$ of $\{0,\dots,m-1\}$.

We are now reduced to the case that $m = 2$. In particular, $[a] \notin \mathcal{A}(\phi|F)$. By construction, the base lift $\Phi_{a,0}$ restricts to an element of $\mathcal{P}(\phi|F,a)$. It follows that $\Phi_{a,1}$ does not restrict to an element of $\mathcal{P}(\phi|F,a)$. The same holds for $\Psi_{a,0}$ and $\Psi_{a,1}$. Since conjugation by $\theta|F$ preserves $\mathcal{P}(\phi|F,a)$, it must be that $\Phi^\Theta_{a,0} = \Psi_{a,0}$ and $\Phi^\Theta_{a,1} = \Psi_{a,1}$. This completes the proof of the lemma. □

The next step in the algorithm is to check if the following condition is satisfied:

(4) If $\mathcal{F}^- \sqsubset \mathcal{F}^+$ has type LH then the twist index for C with respect to f equals the twist index for C' with respect to g.

If not, return `NO` and `STOP`. This is justified by Lemma 17.8.

Lemma 17.9 *It holds that $\theta(r) = r$ for all $\theta \in \mathcal{X}^-$ and $r \in \Delta\mathcal{R}(\phi) = \mathcal{R}(\phi|\mathcal{F}^+) \setminus \mathcal{R}(\phi|\mathcal{F}^-)$.*

Proof Since $\theta \in \mathsf{Ker}(\bar{Q})$, there exists $p \in \mathbb{Z}$ such that $Q_b(\theta) = p$ for all $b \in \mathcal{S}_2(\phi)$ that occur in r. Letting $\upsilon = \theta^{-1}\phi^p$, it follows that $Q_b(\upsilon) = 0$ for all $b \in \mathcal{S}_2(\phi)$ that occur in r. Thus υ satisfies Lemma 15.45(4). Lemma 15.45(1) is obvious and the two remaining items in the hypotheses of that lemma follow from Lemma 17.5. We may therefore apply Lemma 15.45 to conclude that $\upsilon(r) = r$ and hence that $\theta(r) = \theta\upsilon(r) = \phi^p(r) = r$. □

Corollary 17.10 *If $\mathcal{X}^+ \neq \varnothing$, then $\Delta\mathcal{R}(\phi) = \Delta\mathcal{R}(\psi)$.*

Proof If $\theta \in \mathcal{X}^+$ then $\Delta\mathcal{R}(\phi) = \theta(\Delta\mathcal{R}(\phi)) = \Delta\mathcal{R}(\psi)$, where the first equality follows from $\mathcal{X}^+ \subset \mathcal{X}^-$ and Lemma 17.9 and the second equality follows from Definition 17.3(b) and Lemma 3.16(3). □

Notation 17.11 One has that $f(D) = D \cdot \sigma$ for some completely split path $\sigma \subset G_s$, and letting $S_D = \sigma \cdot f_\#(\sigma) \cdot \ldots \cdot f_\#^j(\sigma) \cdot \ldots$, the eigenray R_D determined by D decomposes as $R_D = DS_D$. In the HH case, S_C is defined analogously and $R_C = CS_C$. The rays $R'_{D'}$, $S'_{D'}$, $R'_{C'}$ and $S'_{C'}$ are defined similarly using $g\colon G' \to G'$ in place of $f\colon G \to G$.

Each element of $\Delta\mathcal{R}(\phi)$ is represented by $R_D = DS_D$ or $R_C = CR_C$ and similarly for each element of $\Delta\mathcal{R}(\psi)$. [Feighn and Handel 2018, Lemma 6.3] therefore supplies an algorithm to decide if a given $r \in \Delta\mathcal{R}(\phi)$ and a given $r' \in \Delta\mathcal{R}(\psi)$ are equal. Applying this up to three times, we can decide if $\Delta\mathcal{R}(\phi) = \Delta\mathcal{R}(\psi)$. If $\Delta\mathcal{R}(\phi) \neq \Delta\mathcal{R}(\psi)$ then $\mathcal{X}^+ = \varnothing$ by Corollary 17.10; return `NO` and `STOP`. We may therefore assume that:

(5) $\Delta\mathcal{R}(\phi) = \Delta\mathcal{R}(\psi)$. Denote this common set by $\Delta\mathcal{R}$. In the H and LH cases, the unique element of $\Delta\mathcal{R}$ corresponds to D and D' and is denoted by r_D. In the HH case, $\Delta\mathcal{R} = \{r_C, r_D\}$, where r_C corresponds to C and C', and r_D corresponds to D and D'; this may require reversing the orientation on E'. Remark 13.1 implies that S_D and $S'_{D'}$ are contained in the core filtration element G_p that realizes $F(r_D)$ and that, in the HH case, S_C and $S'_{C'}$ are contained in the core filtration element G_q realizing $F(r_C)$.

[Feighn and Handel 2018, Lemma 6.3] also gives us initial subpaths of S_D and $S'_{D'}$ whose terminal complements are equal. We may therefore assume:

(6) There is a finite path $\kappa_D \subset G_p$ such that $S'_{D'}$ is obtained from $\kappa_D S_D$ by tightening. Similarly, in the case HH, there is a finite path $\kappa_C \subset G_q$ such that $S'_{C'}$ is obtained from $\kappa_C S_C$ by tightening.

We record the following for convenient referencing.

Lemma 17.12 *Suppose that $f\colon G \to G$ is a* CT *representing ϕ and realizing $\mathfrak{c}$ and that either L is an element of $\Omega(\phi)$ or L is an element of $\mathsf{L}_\mathfrak{e}(\phi)$, where $\mathfrak{e} \in \mathfrak{c}$ is not* large. *Let σ be a line in $\mathsf{H}(L)$. Then one of the following (mutually exclusive) properties is satisfied:*

- *L does not cross any higher-order edges; $\sigma = L$.*
- *$L = \beta^{-1}R_e$ (resp. $R_e^{-1}\beta$) for some higher-order edge e and ray β that does not cross any higher-order edges; $\sigma = \beta^{-1}e\tau$ (resp. $\tau^{-1}\overline{e}\beta$), where τ is a ray in the core filtration element that realizes $F(r_e)$.*
- *$L = R_{e_1}^{-1}\rho R_{e_2}$, where e_1, e_2 are higher-order edges and ρ is a Nielsen path; $\sigma = \tau_1^{-1}e_1^{-1}\rho e_2\tau_2$, where τ_1 is a ray in the core filtration element that realizes $F(r_{e_1})$ and τ_2 is a ray in the core filtration element that realizes $F(r_{e_2})$.*

Proof The description of L comes from Lemma 4.2 and the fact (Lemma 13.9) that each such L is carried by $\Gamma(f)$. The description of σ is immediate from the definitions of $\mathsf{H}(L)$. □

Notation 17.13 Represent the trivial element of $\mathsf{Out}(F_n)$ by a homotopy equivalence $h\colon G \to G'$ that restricts to the identity on G_s. We may assume that $h(G_u) = G'_{u'}$ because G_u and $G'_{u'}$ are core graphs that represent $\mathcal{F}^+$. Recall from (5) that G_p is the core filtration element that realizes $F(r_D)$ and that in the HH case, G_q is the core filtration element realizing $F(r_C)$.

Remark 17.14 If the endpoint set of E is equal to the endpoint set of E', then G and G' differ only by a marking change so one can view h, combinatorially, as a homotopy equivalence from G to G (that does not preserve markings). In this case, [Bestvina et al. 2000, Corollary 3.2.2] implies that $h(E) = \bar\mu E'\nu$ or $h(E) = \bar\mu \bar E'\nu$ for some paths $\mu, \nu \subset G_s$. The same conclusion holds if the endpoint sets of E and E' are not equal, because one can fold initial and terminal segments of E' into G_s to arrange that the endpoint set of E is equal to the endpoint set of E'.

The next step in the algorithm is to check if the following statement is satisfied:

(7) If $\mathfrak{e}$ has type HH, then $h(E) = \bar\mu E'\nu$ for some paths $\mu \subset G_q$ and $\nu \subset G_p$.

If (7) fails then we return `NO` and `STOP`. We justify this by the following lemma.

Lemma 17.15 *If $\mathfrak{e}$ has type* HH *and $\mathcal{X}^+ \neq \varnothing$, then $h(E) = \bar\mu E\nu$ for some paths $\mu, \nu \subset G_s$. Moreover, $\mu \subset G_q$ and $\nu \subset G_p$.*

Proof By Remark 17.14, $h(E) = \bar\mu E'\nu$ or $h(E) = \bar\mu \bar E'\nu$ for some paths $\mu, \nu \subset G_s$, so for the main statement, we just want to rule out the latter possibility. Let $L = R_C^{-1} R_D$ and $L' = R_{C'}^{-1} R_{D'}$. Then $\mathsf{L}_{\mathfrak{e}}(\phi) = \{L, L^{-1}\}$ and $\mathsf{L}_{\mathfrak{e}}(\psi) = \{L', L'^{-1}\}$. Since $\mathcal{X}^+ \neq \varnothing$, there exists $\theta_0 \in \mathcal{X}^-$ such that $\theta_0|\mathcal{F}^+$ conjugates $\phi|\mathcal{F}^+$ to $\psi|\mathcal{F}^+$. Thus $\theta_0(\mathsf{L}_{\mathfrak{e}}(\phi)) = \mathsf{L}_{\mathfrak{e}}(\psi)$ and, by Lemma 17.9, L and $\theta_0(L)$ have the same initial ends and the same terminal ends. It follows that $\theta_0(L) = L'$. Since h represents an element of $\mathcal{X}^-$, $h_\#(L) \in \mathsf{H}(L')$. In particular, $h_\#(R_C^{-1} R_D)$ does not cross $\bar E'$, which implies that $h(E)$ does not cross $\bar E'$. This completes the proof of the main statement.

From $h_\#(L) = h_\#(S_C^{-1} \bar C D S_D) = [\bar S_C \bar\mu] \bar C' D' [\nu S_D]$ it follows that $[\nu S_D] \subset G_p$ and $[\mu S_C] \subset G_q$. Thus $\nu \subset G_p$ and $\mu \subset G_q$. □

The remainder of the proof of Lemma 16.5 is the construction of an element $\theta^+ \in \mathcal{X}^+$. By Lemma 17.8 and (4), it suffices to find $\theta \in \mathcal{X}^-$ that induces a bijection between lines that lift to $\Gamma(f_u)$ and lines that lift to $\Gamma(g_{u'})$.

The next lemma states that the conclusions of Lemma 17.15 are satisfied in the H and LH cases without the assumption that $\mathcal{X}^+ \neq \varnothing$.

Lemma 17.16 *It holds that $h(E) = \bar\mu E'\nu$ for some $\mu \subset G_s$ and $\nu \subset G_p$. In the case* HH, *$\mu \subset G_q$.*

Proof The HH case follows from (7) so we consider only the H and LH cases.

By Remark 17.14, $h(E) = \bar{\mu} E' \nu$ or $h(E) = \bar{\mu} \bar{E}' \nu$ for some paths $\mu, \nu \subset G_s$. Each $L \in \mathsf{H}_{\mathfrak{e}}(\phi)$ (realized in G) decomposes as $L = \bar{\alpha} E \beta$ for some rays $\alpha \subset G_s$ and $\beta \subset G_p$. Likewise each $L' \in \mathsf{H}_{\mathfrak{e}}(\psi)$ (realized in G') decomposes as $L' = \bar{\alpha}' E' \beta'$ for some rays $\alpha' \subset G_s$ and some $\beta' \subset G_p$. Since h represents an element of $\mathcal{X}$, Lemma 13.12 implies that $h_\#(L) \in \mathsf{H}_{\mathfrak{e}}(\psi)$. It follows that $h_\#(L)$ does not cross $\bar{E}'$ and hence that $h(E)$ does not cross $\bar{E}'$. This proves that $h(E) \neq \bar{\mu} \bar{E}' \nu$ and so $h(E) = \bar{\mu} E' \nu$. Note also that $h_\#(L) = [\bar{\alpha} \bar{\mu}] E' [\nu \beta]$, which implies that $[\nu \beta] \subset G_p$ and hence that $\nu \subset G_p$. □

Lemma 17.17 *In the case* H, *the path* μ *is a Nielsen path for* $f|G_s = g|G_s$.

Proof There are three cases to consider, depending on the rank of $\Gamma_*(f_u)$, the component of $\Gamma(f_u)$ containing the ray labeled R_D. Let x be the initial vertex of D. Recall that $\mathcal{F}^- = \mathcal{F}_{k-1}$ and $\mathcal{F}^+ = \mathcal{F}_k$.

In the first case, $\operatorname{rank}(\Gamma_*(f_u)) = 0$ and we will show that μ is trivial. By Remark 17.6, there exist edges $E_1, \dots, E_m \in \mathcal{E}_f$ with $m \geq 2$ such that $\Gamma_*(f_u)$ is obtained from $R_D \sqcup R_{E_1} \sqcup \cdots \sqcup R_{D_m}$ by identifying the initial endpoints of all of these rays. By construction, $\mathsf{L}_{\mathfrak{e}}(\phi) = \{L_1, \dots, L_m\}$, where $L_i = R_{E_i}^{-1} R_D$. The description of $\Gamma_*(g_{u'})$ and $\mathsf{L}_{\mathfrak{e}}(\psi)$ is similar with R_D replaced by $R'_{D'}$. There is a permutation π of $\{1, \dots, m\}$ such that $\mathsf{H}(L_i) = \mathsf{H}(L'_{\pi(i)})$. Writing $R_{E_i} = E_i S_i$ for some ray S_i with height less than that of E_i, we have

$$h_\#(L_i) = [S_i^{-1} \bar{E}_i \bar{\mu} D' \nu S_{D'}] = [S_i^{-1} \bar{E}_i \bar{\mu}] D' [\nu S_{D'}],$$

where $[\cdot]$ is the tightening operation. On the other hand, letting $l = \pi(i)$, we have $h_\#(L_i) = \alpha^{-1} \bar{E}_l D' \beta$, where α (resp. β) has height less than that of E_l (resp. D') and so

$$[\mu E_i S_i] = E_l \alpha.$$

Note that S_i has a subray that is disjoint from E_l. Since $S_i = u_i \cdot f_\#(u_i) \cdot f_\#^2(u_i) \cdot \dots$ is a coarsening of the complete splitting of S_i, it follows that S_i is disjoint form E_l; see Remark 13.1. If $i \neq l$ then E_l is the first edge of μ. If $i = l$ then μ is either trivial or has the form $E_i \sigma \bar{E}_i$. In either case, μ is either trivial or begins with E_l. As this is true for all $1 \leq i \leq m$, we conclude that μ is trivial.

If $\operatorname{rank}(\Gamma_*(f_u)) = 1$ then either x is contained in a circle component B of the core filtration element realizing $\mathcal{F}_0$, or there exists $j < k$ such that $\mathcal{F}_{j-1} \sqsubset \mathcal{F}_j$ is an LH extension realized by adding a linear edge C_j and a higher-order edge D_j with "new" initial endpoint x; in the former case, we say that D is type (i) and in the latter case we say that D is type (ii). If D is type (i) then B is a single edge e by the (Periodic Edges) property of a CT, and $\mathsf{L}_{\mathfrak{e}}(\phi) = \{e^\infty D S_D, e^{-\infty} D S_D\}$ and likewise $\mathsf{L}_{\mathfrak{e}}(\psi) = \{e^\infty D' S_{D'}, e^{-\infty} D' S_{D'}\}$. Lemma 17.12 implies that $h_\#(e^\infty D S_D) = e^{\pm\infty} D' \beta'$ for some ray $\beta' \subset G(r_D)$. We also have $h_\#(e^\infty D S_D) = [e^\infty \bar{\mu}] D' [\nu S_D]$. We conclude that $\mu = e^m$ for some m and in particular μ is a Nielsen path. If D is type (ii) then $\mathsf{L}_{\mathfrak{e}}(\phi) = \{w_j^{-\infty} \bar{C}_j D S_D, w_j^{\infty} \bar{C}_j D S_D\}$ and similarly for $\mathsf{L}_{\mathfrak{e}}(\psi)$. As in the previous case, $h_\#(w_j^\infty \bar{C}_j D S_D)$ is equal to both $w_j^{\pm\infty} \bar{C}_j D' \beta'$ and $[w_j^\infty \bar{C}_j \bar{\mu}] D' [\nu S_D]$.

It follows that $\mu C_j w_j^{\infty} = C_j w_j^{\pm\infty}$, which implies that $\mu = C_j w_j^m \bar{C}_j$. This completes the proof if $\mathrm{rank}(\Gamma_*(f_u)) = 1$.

For the final case, assume that $\mathrm{rank}(\Gamma_*(f_u)) \geq 2$ and hence that $x \in G_s$. Given a lift $\tilde{x} \in \tilde{G}$ of the initial endpoint x of D, we set notation as follows: $\tilde{f}\colon \tilde{G} \to \tilde{G}$ is the principal lift that fixes $\tilde{x}$; $\Phi \in \mathcal{P}(\phi)$ is the principal automorphism satisfying $\Phi|\partial F_n = \tilde{f}|\partial F_n$; $\tilde{D}$ is the lift of D with initial endpoint $\tilde{x}$; $\tilde{R}_{\tilde{D}}$ is the lift of R_D whose first edge is $\tilde{D}$; $\tilde{r}_{\tilde{D}}$ is the terminal endpoint of $\tilde{R}_{\tilde{D}}$; and $\tilde{N}$ is the set of lines $(\mathsf{Fix}(\Phi), F(\tilde{r}_D))$ and so is a lift of $\mathsf{H}_{\mathfrak{c}}(\phi) = \{[\mathsf{Fix}(\Phi), F(\tilde{r}_D)]\}$. Similarly, given a lift $\tilde{y} \in \tilde{G}'$ of the initial endpoint y of D', we have: $\tilde{g}\colon \tilde{G}' \to \tilde{G}'$, Ψ, $\tilde{D}'$, $\tilde{R}_{\tilde{D}'}$, $\tilde{r}'_{\tilde{D}'}$ and $\tilde{N}'$. By Lemma 13.12, $\mathsf{H}_{\mathfrak{c}}(\phi) = \mathsf{H}_{\mathfrak{c}}(\psi)$ so we may choose $\tilde{y}$ so that $\mathsf{Fix}(\Phi) = \mathsf{Fix}(\Psi)$ and $F(\tilde{r}_{\tilde{D}}) = F(\tilde{r}'_{\tilde{D}'})$. In particular, $\tilde{N} = \tilde{N}'$.

Let C_s be the component of G_s that contains both x and y — which is possible because they are the endpoints of $\mu \subset G_s$ — and let $\tilde{C}_s \subset \tilde{G}$ be the lift that contains $\tilde{x}$. Then $\tilde{C}_s$ is $\tilde{f}$-invariant and $\mathsf{Fix}(\tilde{f}) \subset \tilde{C}_s$ because G_s contains all Nielsen paths in G with an endpoint at x. There is a free factor F representing $[C_s]$ such that $\partial F = \partial \tilde{C}_s$. Since $\partial\,\mathsf{Fix}(\Phi)$ is contained in the closure of $\mathsf{Fix}(\tilde{f})$, we have $\partial\,\mathsf{Fix}(\Phi) \subset \partial \tilde{C}_s = \partial F$. Letting $\tilde{C}'_s \subset \tilde{G}'$ be the lift of C_s that contains $\tilde{y}$ and F' the free factor satisfying $\partial F' = \partial \tilde{C}'_s$, the same argument shows that $\partial\,\mathsf{Fix}(\Psi) \subset \partial F'$. Since $\mathsf{Fix}(\Phi) = \mathsf{Fix}(\Psi)$ is nontrivial, $F = F'$. Since $\phi|F = \psi|F$ and $\mathsf{Fix}(\Phi) = \mathsf{Fix}(\Psi)$ has rank at least two, $\Phi|F = \Psi|F$.

Let $\tilde{h}\colon \tilde{G} \to \tilde{G}'$ be the lift of $h\colon G \to G'$ that acts as the identity on ∂F_n and let $p\colon \tilde{G} \to G$ and $p'\colon \tilde{G}' \to G'$ be the covering projections. Then $\tilde{h}|\tilde{C}_s\colon \tilde{C}_s \to \tilde{C}'_s$ is a homeomorphism satisfying $p'\tilde{h}(\tilde{z}) = p(\tilde{z})$ for all $\tilde{z} \in \tilde{C}_s$ where, as usual, we are viewing G_s as a subgraph of both G and G'. Moreover, $\tilde{h}\tilde{f}\tilde{h}^{-1}|\tilde{C}'_s = \tilde{g}|\tilde{C}'_s$ because they both project to $g|C'_s$ and induce $\Psi|F'$. In particular $\tilde{g}$ fixes $\tilde{h}(\tilde{x})$. Choose $\tilde{L} \in \tilde{N}$ that decomposes as $\tilde{L} = \tilde{\alpha}\tilde{R}_{\tilde{D}}$. Then $h_\#(L) = [h(\alpha)\mu^{-1}]D'\tau'$ for some ray τ' with height lower than that of $\mathcal{D}'$. Since $\tilde{h}_\#(\tilde{L}) \in \tilde{N}'$ we have $\tilde{h}_\#(\tilde{L}) = [\tilde{h}(\tilde{\alpha})\tilde{\mu}^{-1}]\tilde{D}'\tilde{\tau}'$ for some lift $\tilde{\mu} \subset C'_s$ and some lift $\tilde{\tau}'$ with height lower than that of $\tilde{D}'$. In particular, $\tilde{\mu}$ connects $\tilde{y}$ to $\tilde{h}(\tilde{x})$ and so is a Nielsen path for $\tilde{g}$. Thus μ is a Nielsen path for $g|G_s = f|G_s$. □

We will complete the proof by constructing a homotopy equivalence $d\colon G' \to G'$ such that the outer automorphism θ^+ determined by $dh\colon G \to G'$ is an element of $\mathcal{X}^+$. By construction, d will be the identity on the complement of E' and satisfy $d(E') = \bar{\mu}' E' \nu'$, where $\mu', \nu' \subset G_s$ are closed paths. In the cases H and LH, μ' will be trivial.

Definition 17.18 We define ν', which always corresponds to D', as follows. By (6), there is a finite path $\kappa_D \subset G_p$ such that $S'_{D'}$ is obtained from $\kappa_D S_D$ by tightening. By construction, $h_\#(ES_D)$ is obtained from $\bar{\mu} E' \nu S_D$ by tightening. Letting

$$\nu' = [\kappa_D \bar{\nu}],$$

it follows that $(dh)_\#(ES_D) = [\bar{\mu}\bar{\mu}']E'[\kappa_D \bar{\nu}\nu S_D] = [\bar{\mu}\bar{\mu}']E' S'_{D'}$. Thus, in the cases HH and LH we have

$$(dh)_\#(R_D) = R'_{D'},$$

and in the case H we have
$$(dh)_\#(R_D) = [\bar{\mu}\bar{\mu}']R'_{D'} = \bar{\mu}R'_{D'},$$
where the second equality comes from the fact that μ' is trivial (see below) in the case H. Since $\nu, \kappa_D \subset G_p$, we have:

- (control of ν') $\nu' \subset G_p$.

In the cases LH and H, μ' is defined to be trivial. In the case HH we choose μ' as we did ν' replacing D with C. The result is that in the case HH,
$$(dh)_\#(R_C) = R'_{C'},$$
and so
$$(dh)_\#(R_C^{-1} R_D) = {R'_{C'}}^{-1} R'_{D'}.$$
Also,

- (control of μ') In the case HH, $\mu' \subset G_q$.

This completes the definition of d.

Lemma 17.19 *The outer automorphism δ represented by $d : G' \to G'$ is an element of* $\mathsf{Ker}(Q) \subset \mathcal{X}$.

Proof We use the following properties of $d : G' \to G'$ to prove that $\delta \in \mathcal{X}$:

(a) The map d preserves every component of every filtration element of G'. In particular, δ preserves $\mathfrak{c}$ and every $[F] \in \mathcal{F} \in \mathfrak{c}$.

(b) If e' is not a higher-order edge in E', then $d(e') = e'$.

These are both obvious from the definition.

(c) The map $d_\#$ fixes each Nielsen path of g.

This follows from (b) and the fact that Nielsen paths do not cross higher-order edges.

- If e' is a higher-order edge and $G'_{p'}$ is the filtration element that realizes $F(r'_{e'})$, then the set of rays of the form $e'\beta'$ with $\beta' \subset G'_{p'}$ is mapped into itself by $d_\#$.

If e' is neither C nor D then this follows from (a) and (b). Otherwise it follows from (a), (control of ν') and (control of μ').

Item (c) implies that δ fixes each component of $\mathsf{Fix}(\phi)$ and every element of $\mathcal{A}(\phi) = \mathcal{A}(\psi)$. Item (b) implies that $d|G'_s$ =identity and hence that δ fixes each element of $\mathsf{SA}(\phi|\mathcal{F}_0)$ and so satisfies defining property (6) of $\mathcal{X}$. Suppose that either L' is an element of $\Omega(\psi)$ or L' is an element of $\mathsf{L}_{\mathfrak{c}'}(\psi)$, where $\mathfrak{c}' \in \mathfrak{c}$ is not large. Then (b), (d) and Lemma 17.12 imply that $\mathsf{H}(L')$ is δ-invariant. Similarly, (c) and (d) imply that if $\mathfrak{c}'$ is large then $\mathsf{H}_{\mathfrak{c}'}(\phi)$ is δ-invariant. We have now shown that $\delta \in \mathcal{X}$. In particular, δ is in the domain of $\bar{Q}$.

To prove that $\delta \in \mathsf{Ker}(Q)$, suppose that $\widetilde{b}' = (\widetilde{L}'_1, \widetilde{L}'_2)$ is a staple pair for ψ with common axis $\widetilde{A}'$. By (b), there is a lift $\widetilde{d}\colon \widetilde{G}' \to \widetilde{G}'$ of d that pointwise fixes $\widetilde{A}'$. We claim that $\widetilde{d}$ preserves both $\mathsf{H}(\widetilde{L}'_1)$ and $\mathsf{H}(\widetilde{L}'_2)$. The $\widetilde{L}'_1$ and $\widetilde{L}'_2$ cases are symmetric so we will consider L'_2. If both ends of L'_2 are periodic then $\mathsf{H}(L'_2) = \{L'_2\}$ by Lemma 17.12. Moreover, $\widetilde{L}'_2$ does not cross any higher-order edges and so is pointwise fixed by $\widetilde{d}$. We may therefore assume that there is a decomposition $\widetilde{L}'_2 = \widetilde{\alpha}\widetilde{R}'_{\widetilde{e}'_2}$, where $\widetilde{e}'_2$ is a higher-order edge and $\widetilde{\alpha}$ does not cross any higher-order edges. It follows that $\widetilde{d}$ pointwise fixes $\widetilde{\alpha}$ and that $\widetilde{e}_2$ is the initial edge of $\widetilde{d}_\#(\widetilde{R}'_{\widetilde{e}'_2})$. Item (d) now implies that $\widetilde{d}$ preserves $\mathsf{H}(\widetilde{L}'_2)$. This completes the proof of the claim. It now follows from the definitions that $m_{b'}(\delta) = 0$ and hence $Q_{b'}(\delta) = 0$. Since b' is arbitrary, $\delta \in \mathsf{Ker}(Q)$. □

The final step in the algorithm that proves Proposition 16.4 is to return YES and the outer automorphism θ^+ represented by $dh\colon G \to G'$. In conjunction with Lemma 17.4, the following lemma justifies this step.

Lemma 17.20 *The map $dh\colon G \to G'$ represents an element $\theta^+ \in \mathcal{X}^+$.*

Proof Since h represents an element of $\mathsf{Ker}(\bar{Q})$, Lemma 17.19 implies that $\theta^+ \in \mathsf{Ker}(\bar{Q}) \subset \mathcal{X}$. We are therefore reduced to showing that $\theta^+|\mathcal{F}^+$ conjugates $\phi|\mathcal{F}^+$ to $\psi|\mathcal{F}^+$. By Lemma 17.8 and (4), it suffices to prove that:

(b1) A line $L \subset G$ lifts to $\Gamma(f_u)$ if and only if $\theta^+(L) \subset G'$ lifts to $\Gamma(g_{u'})$.

Recall that $\Gamma(f_u)$ is obtained from $\Gamma(f_s)$ by either adding a single new component (the HH and LH cases) or by adding a ray to one of the components $\Gamma_*(f_s)$ of $\Gamma(f_s)$ (the H case). The same is true for $\Gamma(g_{u'})$. The component of $\Gamma(f_u)$ that is not a component of $\Gamma(f_s)$ is denoted by $\Gamma_*(f_u)$; it is the unique component that contains a ray labeled R_D. Likewise, the "new" component $\Gamma_*(g_{u'})$ of $\Gamma(g_{u'})$ is the one that contains a ray labeled $R'_{D'}$. Recall also that $f_s = g_s$, that $\Gamma(f_s) = \Gamma(g_s)$, and that $\Gamma_*(f_s) = \Gamma_*(g_s)$.

Item (b1) is obvious if L lifts into a component of $\Gamma(f_s)$ so we may assume, after reversing orientation on L if necessary, that R_D is a terminal ray of L.

In the case HH, $L = R_C^{-1} R_D$ so (b1) follows from $(dh)_\#(R_C^{-1} R_D) = {R'_{C'}}^{-1} R'_{D'}$; see Definition 17.18.

In the case H, $(dh)_\#(R_D) = \bar{\mu} R'_{D'}$ by Definition 17.18. Let $\hat{x}$ be the initial endpoint of the lift of R_D into $\Gamma(f_u)$ and let x be its projection into G_s. Define $\hat{y}$ and y similarly using $R'_{D'}$ in place of R_D. If R_D is also a terminal ray of L^{-1}, then $L = R_D^{-1}\xi R_D$ for some Nielsen path ξ and $(dh)_\#(L) = R'^{-1}_{D'}[\mu\xi\bar{\mu}]R'_{D'}$, which lifts into $\Gamma(g_{u'})$ because $[\mu\xi\bar{\mu}]$ is a Nielsen path by Lemma 17.17. The remaining case is that $L = \beta^{-1} R_D$ for some ray $\beta \subset G_s$ that lifts to a ray in $\Gamma(f_s)$ based at $\hat{x}$. In this case, $(dh)_\#(L) = \beta^{-1}\bar{\mu} R'_{D'}$, which lifts into $\Gamma^0(g_{u'})$ because μ is a Nielsen path that lifts to a path in $\Gamma(g_s)$ connecting $\hat{y}$ to $\hat{x}$.

In the case LH, $\mathsf{L}_\mathfrak{c}(\phi) = \{L_+, L_-\}$ and $\mathsf{L}_\mathfrak{c}(\psi) = \{L'_+, L'_-\}$, where $L_\pm = w^{\pm\infty}\bar{C} R_D$ and $L'_\pm = w^{\pm\infty}\bar{C}' R'_{D'}$. Since $\mathsf{I}_\mathfrak{c}(\phi) = \mathsf{I}_\mathfrak{c}(\psi)$ and $\theta^+ \in \mathcal{X}$, we have that $\theta^+(L_+)$ is contained in either $\mathsf{H}(L'_+)$ or $\mathsf{H}(L'_-)$. By construction, $\theta^+(L_+) = (dh)_\#(L_+) = [w^\infty\bar{\mu}]\bar{C}' R'_{D'}$. Thus, $\theta^+(L_+) \in \mathsf{H}(L'_+)$ and $[w^\infty\bar{\mu}] = w^\infty$. It follows that $\mu = w^p$ for some $p \in \mathbb{Z}$ and hence that $\theta^+(L_\pm) = L'_\pm$ and $\theta^p(R_D^{-1} w^k R_D) = R'^{-1}_{D'} w^k R'_{D'}$ for all k. This completes the proof of (b1) and hence the proof of the lemma. □

Appendix More on $\mathsf{Ker}\,\bar{Q}$

The main results of this appendix are that $m_b(\theta)$ can be computed for $\theta \in \mathcal{X}$, that $\mathsf{Ker}\,\bar{Q}$ is of type VF (Definition 14.4), and that a finite presentation for $\mathsf{Ker}\,\bar{Q}$ can be computed. This section is needed for future work and is not used in the proof of the main theorem of this paper.

A.1 A Stallings graph for $\mathsf{H}_{\phi,\mathsf{c}}(\tilde{b})$

We will need the following remark.

Remark A.1 Suppose that G is a marked graph and that for $i = 1, 2$, $\tilde{A}_i$ is the axis of a covering translation $T_i\colon \tilde{G} \to \tilde{G}$ and that the number of edges in a fundamental domain for $\tilde{A}_i$ is s_i. If $\tilde{A}_1 \cap \tilde{A}_2$ contains at least $s_1 + s_2 + 1$ edges then $\tilde{A}_1 = \tilde{A}_2$. To see this, decompose $\tilde{A}_1 \cap \tilde{A}_2 = e_1 e_2 \dots$ into edges and note that $T_1T_2(e_1) = e_{s_1+s_2+1} = T_2T_1(e_1)$. Since T_1T_2 and T_2T_1 agree on an edge they are equal, and so T_1 and T_2 have the same axes.

Notation A.2 Let $f\colon G \to G$ be a CT for ϕ with $\tilde{f}\colon \tilde{G} \to \tilde{G}$ a lift to the universal cover. Assume notation as in the definition of m_b (Definition 15.36). In particular, $b = (L_1, L_2) \in \mathcal{S}_2(\phi)$, $\tilde{b} = (\tilde{L}_1, \tilde{L}_2)$ are lifts of (L_1, L_2) such that $\tilde{L}_1^+, \tilde{L}_2^- \in \{\tilde{A}^-, \tilde{A}^+\}$, where $\tilde{A}$ is the common axis of $\tilde{b}$ and $a \in F_n$ is a root-free element with axis $\tilde{A}$ and orientation chosen as in the definition. For $\theta \in \mathcal{X}$, $\Theta_i \in \theta$ is defined uniquely by $\Theta_i(\mathsf{H}_{\phi,\mathsf{c}}(\tilde{L}_i)) = \mathsf{H}_{\phi,\mathsf{c}}(\tilde{L}_i)$ and $m_b(\theta)$ is defined so that $\Theta_1 = i_a^{m_b(\theta)}\Theta_2$.

If $\tilde{L}_2^+ = \tilde{r}_2$ for some $r_2 \in \mathcal{R}(\phi)$, then define $\mathsf{H}^2_{\phi,\mathsf{c}}(\tilde{b}) = F_{\mathsf{c}}(\tilde{r}_2)$. Otherwise, $\tilde{L}_2^+ \in \{c_2^-, c_2^+\}$ for some root-free $c_2 \in F_n$ representing an element of $\mathcal{A}(\phi)$ and $\mathsf{H}^2_{\phi,\mathsf{c}}(\tilde{b}) := \langle c_2 \rangle$. Define $\mathsf{H}^1_{\phi,\mathsf{c}}(\tilde{b})$ similarly using $\tilde{L}_1^-$ in place of $\tilde{L}_2^+$. Finally, define

$$\mathsf{H}_{\phi,\mathsf{c}}(\tilde{b}) = \langle \mathsf{H}^1_{\phi,\mathsf{c}}(\tilde{b}), \mathsf{H}^2_{\phi,\mathsf{c}}(\tilde{b}) \rangle.$$

The covering transformation corresponding to a is denoted by τ. Additionally, $H^i := \mathsf{H}^i_{\phi,\mathsf{c}}(\tilde{b})$, T^i denotes the minimal subtree for H^i, Γ^i denotes the Stallings graph for H^i, and $\tilde{L}(k)$ denotes the line $[\tilde{L}_1^-, \tau^k(\tilde{L}_2^+)]$.

Remark A.3 Comparing definitions of $\mathsf{H}_{\phi,\mathsf{c}}(\tilde{L}_i)$ and H^i, Θ_i is the unique $\Theta \in \theta$ fixing $\langle a \rangle$ and H^i.

Lemma A.4 *There is a $k \ge 0$ such that*

- $T^1 \cap \tau^k(T^2) = \varnothing$, *and*
- *the arc $\tilde{\mu}$ spanning between T^1 and $\tau^k(T^2)$ contains more than two fundamental domains of $\tilde{A}$ with orientation agreeing with that of $\tilde{\mu}$.*

Proof The ends $\tilde{A}$ are not ends of T^i. Indeed, if $\tilde{r}_i$ is ray, then the associated higher-order edge separates T^i and the ends of $\tilde{A}$. If not, then T^i is the axis corresponding to the end of $\tilde{L}_i$ that is not an end of $\tilde{A}$. Hence there is a neighborhood of $\tilde{A}^+$ that is disjoint from T_i. Therefore, the conclusion holds for all large k. □

Corollary A.5 *We may compute*:

(1) *For all $l \in \mathbb{Z}$, the Stallings graph for $\langle H^1, i_a^l(H^2)\rangle$.*

(2) *An integer $k \geq 0$ as in Lemma A.4.*

(3) *For all $\theta \in \mathcal{X}$, $m_b(\theta)$.*

Proof (1) By Bass–Serre theory, for k as in Lemma A.4, the Stallings graph for $\langle H^1, i_a^k(H^2)\rangle$ is obtained by attaching at its endpoints a copy of the arc spanning between T^1 and $\tau^k(T^2)$ to the Stallings graphs for $[H^1]$ and $[i_a^k H^2] = [H^2]$.

These latter graphs can be computed. Indeed, if $\widetilde{L}_1^-$ is an eigenray, then, by definition, $[H^1]$ has as its Stallings graph a component of a stratum of G, and otherwise is a circle representing $\langle c_1\rangle$. There is a symmetric argument for $[H^2]$. $\widetilde{L}(k)$ spans between T^1 and $\tau^k(T^2)$. Hence the desired Stallings graph, for large k, is the result of immersing the ends of $\widetilde{L}(k)$ into the Stallings graphs and then performing any folding. By Lemma A.4, folding stops when the copy of $\widetilde{\mu}$ is the spanning arc.

We see that, for large k, $\langle H^1, i_a^k(H^2)\rangle$ is an internal free product. By Remark A.3, $\Phi_1^s(\langle H^1, i_a^k(H^2)\rangle) = \langle H^1, i_a^{sm_b(\phi)+k}(H^2)\rangle$ is also a free product and its Stallings graph can be computed as above (but perhaps the spanning arc is folded away). Recall (Lemma 15.39) that $m_b(\phi) \neq 0$. We note in passing that therefore $[H^1, i_a^l(H^2)]$ is good (Definition 7.13).

(2) For $l = 0, 1, 2, \dots$, iteratively start computing Stallings graphs for $\langle H^1, i_a^l(H^2)\rangle$. When, after folding in $\widetilde{L}(l)$, more than two correctly oriented fundamental domains of $\widetilde{A} \cap \widetilde{L}(l)$ remain in the spanning arc, stop and set $k = l$.

(3) Let $\theta \in \mathcal{X}$ and $m := m_b(\theta)$. If $\Theta_1(H^1) = H^1$, then by definition,

$$\Theta_1(H^1, i_a^k(H^2)) = (H^1, i_a^{k+m}(H^2)).$$

Hence, m can be read off by comparing the Stallings graphs for $[\langle H^1, i_a^k(H^2)\rangle]$ and $[\langle H^1, i_a^{k+m}(H^2)\rangle]$ for large enough k. The latter, being the Stallings graph for $\theta[\langle H^1, i_a^k(H^2)\rangle]$, can be computed by representing θ as a topological representative $g\colon G \to G$, applying g to the Stallings graph for $[\langle H^1, i_a^k(H^2)\rangle]$, tightening, and taking the core. □

Corollary A.6 (1) *$[H_1, i_a^l H_2]$ is good for all $l \in \mathbb{Z}$.*

(2) *No nontrivial power of a is in $\langle H_1, i_a^l H_2\rangle$.*

Proof (1) This was noted during the proof of Corollary A.5.

(2) Since $\Theta_1(a) = a$, for the second item it is enough to show that

$$a \notin \Theta_1^l(\langle H^1, i_a^l(H^2)\rangle) = \langle H^1, i_a^{l+m}(H^2)\rangle$$

for large $l + m$. So assume that $k > 0$ is as in Lemma A.4. If $a \in \langle H^1, i_a^k(H^2)\rangle$, then there is an immersion of $\widetilde{A}$ with image a closed loop into the Stallings graph for $\langle H^1, i_a^k(H^2)\rangle$ and that overlaps the copy of $\widetilde{\mu}$ in its intersection with $\widetilde{A}$. The immersion crosses this spanning arc at most once. Indeed, otherwise there would be a covering translation of $\widetilde{G}$ taking $\widetilde{A}$ to itself reversing orientation, but a and a^{-1} are not conjugate; see Remark A.1. Hence the image of $\widetilde{A}$ is not a closed loop. □

Corollary A.7 *For* $\theta \in \mathcal{X}$, $m_b(\theta) = 0 \iff [\mathsf{H}_{\phi,\mathfrak{c}}(\widetilde{b})]$ *is* θ*-invariant.*

Proof ($\Longrightarrow$) Suppose $m := m_b(\theta) = 0$. Using Remark A.3 we have

$$\Theta_1(\mathsf{H}_{\phi,\mathfrak{c}}(\widetilde{b})) = \Theta_1(\langle H^1, H^2\rangle) = \langle \Theta_1(H^1), \Theta_1(H^2)\rangle = \langle H^1, i_a^m(H^2)\rangle = \langle H^1, H^2\rangle.$$

($\Longleftarrow$) Suppose $m > 0$, the case for $m < 0$ being similar. Choose l so that $k := ml$ is as in Lemma A.4. We saw in Corollary A.5 that the Stallings graph for $[\langle H^1, i_a^{ml}(H^2)\rangle]$ has an arc spanning between Stallings graphs for $[H^1]$ and $[H^2]$ and that the Stallings graph for $[\Theta_1^{l+1}(\langle H^1, H^2\rangle)]$ is obtained by inserting m copies of a fundamental domain for $\widetilde{A}$ into this spanning arc. Since these Stallings graphs are not equal, $\theta([\langle H^1, H^2\rangle]) \neq [\langle H^1, H^2\rangle]$. □

A.2 Ker Q

Recall (Definition 16.2) the homomorphism $Q = Q^\phi : \mathcal{X} \to \mathbb{Q}^{\mathcal{S}_2(\phi)}$ given by setting the b^{th} coordinate of $Q(\theta)$ equal to $m_b(\theta)/m_b(\phi)$.

Proposition A.8 *A finite presentation for* $\mathsf{Ker}\, Q$ *can be computed.* $\mathsf{Ker}\, Q$ *is of type* VF.

Proof We use the notation as in Section A.1. For each $b \in \mathcal{S}_2(\phi)$, choose $\widetilde{b}$ and replace it with $\Phi_1^k(\widetilde{b})$ and compute $[\mathsf{H}_{\phi,\mathfrak{c}}(\widetilde{b})]$, where k is as in Lemma A.4. By Corollary A.7, $\theta \in \mathsf{Ker}\, Q$ if and only if $\theta \in \mathcal{X}$ and $\theta([\mathsf{H}_{\phi,\mathfrak{c}}(\widetilde{b})]) = [\mathsf{H}_{\phi,\mathfrak{c}}(\widetilde{b})]$ for each b. Hence $\theta \in \mathsf{Ker}\, Q$ if and only if θ fixes the concatenation of the sequences (J) and $([\mathsf{H}_{\phi,\mathfrak{c}}(\widetilde{b}_1)], \dots, [\mathsf{H}_{\phi,\mathfrak{c}}(\widetilde{b}_N)])$, where $(b_1, \dots b_N)$ is an ordering of $\mathcal{S}_2(\phi)$ and J is as in Definition 14.1. This concatenation is an element of $\overline{\mathsf{IS}}(\mathbb{A}_\bullet)$; see Notation 10.13. By Lemma 11.2, we can compute a finite presentation for $\mathsf{Ker}\, Q$ and, by Proposition 14.5, $\mathsf{Ker}\, Q$ is of type VF. □

A.3 Ker $\bar{Q}$

$\bar{Q}$ is defined in Definition 16.3.

Proposition A.9 *A finite presentation for* $\mathsf{Ker}\, \bar{Q}$ *can be computed.* $\mathsf{Ker}\, \bar{Q}$ *is of type* VF.

Proof Let π denote the quotient map $Q(\mathcal{X}) \to \bar{Q}(\mathcal{X})$. Since we have a finite generating set for $\mathcal{X}$, we can compute a finite presentation for the free abelian group $\mathsf{Ker}\, \pi$. Using Proposition A.8, Lemma 9.3 and

$$1 \to \mathsf{Ker}\, Q \to \mathsf{Ker}\, \bar{Q} \to \mathsf{Ker}\, \pi \to 1,$$

we can compute a finite presentation for $\mathsf{Ker}\, \bar{Q}$.

The above short exact sequence shows that $\mathsf{Ker}\, \bar{Q}$ is an extension of a group of type VF (Proposition A.8) by a finitely generated abelian group. It follows from the proposition below that $\mathsf{Ker}\, \bar{Q}$ is of type VF. □

The following proposition is from Moritz Rodenhausen's thesis. As far as we know, it is not published and so for the reader's convenience we copy his proof here.

Proposition A.10 [Rodenhausen 2013, Proposition 13.18] *Suppose that in the short exact sequence* $1 \to G' \xrightarrow{i} G \xrightarrow{\pi} G'' \to 1$, *the group* G' *is of type* VF *and* G'' *is finitely generated abelian. Then* G *is of type* VF.

Proof Suppose first that G'' is infinite cyclic. Let H' be a subgroup of some finite index d in G' such that H' is of type F. The intersection K' of all (finitely many) index d subgroups of G' also is of type F. Furthermore, the group K' is a characteristic subgroup of G'. Let now $t \in G$ be an element such that $\pi(t)$ generates G''. We denote by K the subgroup of G generated t and $i(K')$. It has finite index in G and fits into a short exact sequence

$$1 \to K' \to K \to G'' \to 1.$$

We see that K is an extension of groups of type F and so has type F; see [Geoghegan 2008, Theorem 7.3.4]. Hence G is of type VF.

The case where G'' is isomorphic to $\mathbb{Z}^m$ is proved by induction on m. Let A'' be an infinite cyclic summand of G'', so $G''/A'' \cong \mathbb{Z}^{m-1}$. The short exact sequences

$$1 \to G' \to \pi^{-1}(A'') \to A'' \to 1 \quad \text{and} \quad 1 \to \pi^{-1}(A'') \to G \to G''/A'' \to 1$$

together with the induction hypothesis completes the proof in the case $G'' \cong \mathbb{Z}^m$.

For the general case, let H'' be a free abelian subgroup of G'' of finite index and $H = \pi^{-1}(H'') \subset G$. We obtain a short exact sequence

$$1 \to G' \to H \to H'' \to 1.$$

Hence H, and so also G, is of type VF. □

References

[Bestvina and Brady 1997] **M Bestvina**, **N Brady**, *Morse theory and finiteness properties of groups*, Invent. Math. 129 (1997) 445–470 MR Zbl

[Bestvina and Handel 1992] **M Bestvina**, **M Handel**, *Train tracks and automorphisms of free groups*, Ann. of Math. 135 (1992) 1–51 MR Zbl

[Bestvina et al. 1997] **M Bestvina**, **M Feighn**, **M Handel**, *Laminations, trees, and irreducible automorphisms of free groups*, Geom. Funct. Anal. 7 (1997) 215–244 MR Zbl

[Bestvina et al. 2000] **M Bestvina**, **M Feighn**, **M Handel**, *The Tits alternative for* Out(F_n)*, I: Dynamics of exponentially-growing automorphisms*, Ann. of Math. 151 (2000) 517–623 MR Zbl

[Bestvina et al. 2004] **M Bestvina**, **M Feighn**, **M Handel**, *Solvable subgroups of* Out(F_n) *are virtually abelian*, Geom. Dedicata 104 (2004) 71–96 MR Zbl

[Bestvina et al. 2005] **M Bestvina**, **M Feighn**, **M Handel**, *The Tits alternative for* $\mathrm{Out}(F_n)$*, II: A Kolchin type theorem*, Ann. of Math. 161 (2005) 1–59 MR Zbl

[Bestvina et al. 2023] **M Bestvina**, **M Feighn**, **M Handel**, *A McCool Whitehead type theorem for finitely generated subgroups of* $\mathrm{Out}(F_n)$, Ann. H. Lebesgue 6 (2023) 65–94 MR Zbl

[Bridson and Wilton 2011] **M R Bridson**, **H Wilton**, *On the difficulty of presenting finitely presentable groups*, Groups Geom. Dyn. 5 (2011) 301–325 MR Zbl

[Cohen and Lustig 1999] **M M Cohen**, **M Lustig**, *The conjugacy problem for Dehn twist automorphisms of free groups*, Comment. Math. Helv. 74 (1999) 179–200 MR Zbl

[Cooper 1987] **D Cooper**, *Automorphisms of free groups have finitely generated fixed point sets*, J. Algebra 111 (1987) 453–456 MR Zbl

[Dahmani 2016] **F Dahmani**, *On suspensions and conjugacy of hyperbolic automorphisms*, Trans. Amer. Math. Soc. 368 (2016) 5565–5577 MR Zbl

[Dahmani 2017] **F Dahmani**, *On suspensions, and conjugacy of a few more automorphisms of free groups*, from "Hyperbolic geometry and geometric group theory", Adv. Stud. Pure Math. 73, Math. Soc. Japan, Tokyo (2017) 135–158 MR Zbl

[Dahmani and Touikan 2021] **F Dahmani**, **N Touikan**, *Reducing the conjugacy problem for relatively hyperbolic automorphisms to peripheral components*, preprint (2021) arXiv 2103.16602

[Dahmani and Touikan 2023] **F Dahmani**, **N Touikan**, *Unipotent linear suspensions of free groups*, preprint (2023) arXiv 2305.11274

[Dahmani et al. 2025] **F Dahmani**, **S Francaviglia**, **A Martino**, **N Touikan**, *The conjugacy problem for* $\mathrm{Out}(F_3)$, Forum Math. Sigma 13 (2025) art. id. e41 MR Zbl

[Feighn and Handel 2011] **M Feighn**, **M Handel**, *The recognition theorem for* $\mathrm{Out}(F_n)$, Groups Geom. Dyn. 5 (2011) 39–106 MR Zbl

[Feighn and Handel 2018] **M Feighn**, **M Handel**, *Algorithmic constructions of relative train track maps and CTs*, Groups Geom. Dyn. 12 (2018) 1159–1238 MR Zbl

[Gaboriau et al. 1998] **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 MR Zbl

[Geoghegan 2008] **R Geoghegan**, *Topological methods in group theory*, Grad. Texts in Math. 243, Springer (2008) MR Zbl

[Gersten 1984] **S M Gersten**, *On Whitehead's algorithm*, Bull. Amer. Math. Soc. 10 (1984) 281–284 MR Zbl

[Gersten 1987] **S M Gersten**, *Fixed points of automorphisms of free groups*, Adv. Math. 64 (1987) 51–85 MR Zbl

[Handel and Mosher 2011] **M Handel**, **L Mosher**, *Axes in outer space*, Mem. Amer. Math. Soc. 1004, Amer. Math. Soc., Providence, RI (2011) MR Zbl

[Handel and Mosher 2020] **M Handel**, **L Mosher**, *Subgroup decomposition in* $\mathrm{Out}(F_n)$, Mem. Amer. Math. Soc. 1280, Amer. Math. Soc., Providence, RI (2020) MR Zbl

[Kalajdževski 1992] **S Kalajdževski**, *Automorphism group of a free group: centralizers and stabilizers*, J. Algebra 150 (1992) 435–502 MR Zbl

[Kapovich and Benakli 2002] **I Kapovich**, **N Benakli**, *Boundaries of hyperbolic groups*, from "Combinatorial and geometric group theory", Contemp. Math. 296, Amer. Math. Soc., Providence, RI (2002) 39–93 MR Zbl

[Kapovich and Short 1996] **I Kapovich**, **H Short**, *Greenberg's theorem for quasiconvex subgroups of word hyperbolic groups*, Canad. J. Math. 48 (1996) 1224–1244 MR Zbl

[Krstić et al. 2001] **S Krstić**, **M Lustig**, **K Vogtmann**, *An equivariant Whitehead algorithm and conjugacy for roots of Dehn twist automorphisms*, Proc. Edinb. Math. Soc. 44 (2001) 117–141 MR Zbl

[Los 1996] **J E Los**, *On the conjugacy problem for automorphisms of free groups*, Topology 35 (1996) 779–808 MR Zbl

[Lustig 2000] **M Lustig**, *Structure and conjugacy for automorphisms of free groups*, *I*, preprint 2000-130, Max Planck Inst. (2000) Available at `https://www.mpim-bonn.mpg.de/preblob/834`

[Lustig 2001] **M Lustig**, *Structure and conjugacy for automorphisms of free groups*, *II*, preprint 2001-4, Max Planck Inst. (2001) Available at `https://www.mpim-bonn.mpg.de/preblob/854`

[Lustig 2007] **M Lustig**, *Conjugacy and centralizers for iwip automorphisms of free groups*, from "Geometric group theory", Birkhäuser, Basel (2007) 197–224 MR Zbl

[Lyndon and Schupp 1977] **R C Lyndon**, **P E Schupp**, *Combinatorial group theory*, Ergebnisse der Math. 89, Springer (1977) MR Zbl

[Magnus et al. 1966] **W Magnus**, **A Karrass**, **D Solitar**, *Combinatorial group theory*: *presentations of groups in terms of generators and relations*, Interscience, New York (1966) MR Zbl

[McCool 1975] **J McCool**, *Some finitely presented subgroups of the automorphism group of a free group*, J. Algebra 35 (1975) 205–213 MR Zbl

[Rodenhausen 2013] **M Rodenhausen**, *Centralisers of polynomially growing automorphisms of free groups*, PhD thesis, Universität Bonn (2013) Available at `https://hdl.handle.net/20.500.11811/5725`

[Sela 1995] **Z Sela**, *The isomorphism problem for hyperbolic groups, I*, Ann. of Math. 141 (1995) 217–283 MR Zbl

[Serre 1980] **J-P Serre**, *Trees*, Springer (1980) MR Zbl

[Stallings 1983] **J R Stallings**, *Topology of finite graphs*, Invent. Math. 71 (1983) 551–565 MR Zbl

[Veblen and Franklin 1921] **O Veblen**, **P Franklin**, *On matrices whose elements are integers*, Ann. of Math. 23 (1921) 1–15 MR Zbl

[Whitehead 1936a] **J H C Whitehead**, *On certain sets of elements in a free group*, Proc. Lond. Math. Soc. 41 (1936) 48–56 MR Zbl

[Whitehead 1936b] **J H C Whitehead**, *On equivalent sets of elements in a free group*, Ann. of Math. 37 (1936) 782–800 MR Zbl

Department of Mathematics and Computer Science, Rutgers University
Newark, NJ, United States

Mathematics and Computer Science Department, Herbert H Lehman College (CUNY)
Bronx, NY, United States

`feighn@rutgers.edu`, `michaelxhandel@gmail.com`

Proposed: Martin R Bridson | Received: 2 July 2019
Seconded: David Fisher, Bruce Kleiner | Revised: 7 March 2022

GEOMETRY & TOPOLOGY

msp.org/gt

See inside back cover or msp.org/gt for submission instructions.

The subscription price for 2025 is US $865/year for the electronic version, and $1210/year (+$75, if shipping outside the US) for print and electronic. Subscriptions, requests for back issues and changes of subscriber address should be sent to MSP. Geometry & Topology is indexed by Mathematical Reviews, Zentralblatt MATH, Current Mathematical Publications and the Science Citation Index.

Geometry & Topology (ISSN 1465-3060 printed, 1364-0380 electronic) is published 9 times per year and continuously online, by Mathematical Sciences Publishers, c/o Department of Mathematics, University of California, 798 Evans Hall #3840, Berkeley, CA 94720-3840. Periodical rate postage paid at Oakland, CA 94615-9651, and additional mailing offices. POSTMASTER: send address changes to Mathematical Sciences Publishers, c/o Department of Mathematics, University of California, 798 Evans Hall #3840, Berkeley, CA 94720-3840.

GT peer review and production are managed by EditFLOW® from MSP.

PUBLISHED BY

mathematical sciences publishers

nonprofit scientific publishing

http://msp.org/

GEOMETRY & TOPOLOGY

Volume 29 Issue 4 (pages 1693–2250) 2025