Groups Geom. Dyn. 12 (2018), 1159–1238
DOI 10.4171/GGD/466

Groups, Geometry, and Dynamics

Algorithmic constructions of relative train track maps and CTs

Mark Feighn[1] and Michael Handel[2]

Abstract. Building on [BH92, BFH00], we proved in [FH11] that every element ψ of the outer automorphism group of a finite rank free group is represented by a particularly useful relative train track map. In the case that ψ is rotationless (every outer automorphism has a rotationless power), we showed that there is a type of relative train track map, called a CT, satisfying additional properties. The main result of this paper is that the constructions of these relative train tracks can be made algorithmic. A key step in our argument is proving that it is algorithmic to check if an inclusion $\mathcal{F} \sqsubset \mathcal{F}'$ of ϕ-invariant free factor systems is reduced. We also give applications of the main result.

Mathematics Subject Classification (2010). 20F65, 20E36.

Keywords. Outer automorphisms of free groups, train tracks.

[1] This material is based upon work supported by the National Science Foundation under Grant No. DMS-1406167 and also under Grant No. DMS-14401040 while the author was in residence at the Mathematical Sciences Research Institute in Berkeley, California, during the Fall 2016 semester.

[2] This material is based upon work supported by the National Science Foundation under Grant No. DMS-1308710 and by PSC-CUNY grants in Program Years 46 and 47.

Contents

1. Introduction

An automorphism Φ of the rank n free group F_n is typically represented by giving its effect on a basis of F_n. Equivalently, if we identify the edges of the rose R_n (the graph with one vertex $*$ and n edges) with basis elements of F_n, then Φ may be represented as a self homotopy equivalence of R_n preserving $*$. In this paper, we are interested in outer automorphisms, that is we are interested in elements of the quotient

$$\mathsf{Out}(F_n) := \mathsf{Aut}(F_n)/\mathsf{Inn}(F_n)$$

of the automorphism group $\mathsf{Aut}(F_n)$ of F_n by its subgroup $\mathsf{Inn}(F_n)$ of inner automorphisms. Since outer automorphisms are defined only up to conjugation, the base point loses its special role and $\phi \in \mathsf{Out}(F_n)$ is typically represented as a self homotopy equivalence of R_n that is not required to fix $*$. More generally it is also typical to take advantage of the flexibility gained by representing ϕ as a self homotopy equivalence $f\colon G \to G$ of a marked graph G, i.e. a graph G equipped with a homotopy equivalence $\mu\colon R_n \to G$. The marking identifies the fundamental group of G with F_n, but only up to conjugation. In an analogy with linear maps, representing ϕ as $f\colon G \to G$ corresponds to writing a linear map in terms of a particular basis.

In [BH92], Bestvina–Handel showed that every element of $\mathsf{Out}(F_n)$ has a representation as a relative train track map, that is a representation $f: G \to G$ as above but with strong properties. In the analogy with linear maps, a relative train track map corresponds to a normal form. [BH92] goes on to use relative train track maps to solve the Scott conjecture: the rank of the fixed subgroup of an element of $\mathsf{Aut}(F_n)$ is at most n. Further applications spurred the further development of the theory of relative train tracks. See for example, [BFH00] where improved relative train tracks (IRTs) were used to show that $\mathsf{Out}(F_n)$ satisfies the Tits alternative or [FH09] where completely split relative train tracks (CTs) [FH11] were used to classify abelian subgroups of $\mathsf{Out}(F_n)$.

CTs are relative train tracks that were designed to satisfy the properties that have proven most useful (to us) for investigating elements of $\mathsf{Out}(F_n)$. By convention, the identity map on the rose with one petal is a CT representing the trivial element of $\mathsf{Out}(F_1)$. For the definition when $n \geq 2$, see Section 5. Not every $\phi \in \mathsf{Out}(F_n)$ is represented by a CT, but all rotationless (see Definition 3.3) elements are. This is not a big restriction since there is a specific $M > 0$ depending only on n (see Corollary 3.14) such that ϕ^M is rotationless. We believe that CTs will be of general use in approaching algorithmic questions about $\mathsf{Out}(F_n)$. As a first step in that process, our main theorem (Theorem 1.1) verifes that CTs can be constructed algorithmically.

Many arguments involving CTs go by induction up through the strata. It is therefore useful if a CT $f: G \to G$ satisfies the following axiom.

Inheritance. The restriction of f to each component of each core filtration element is a CT.

As Example 5.6 shows, not every CT satisfies (inheritance). By using an upward induction argument (see Section 7) instead of the downward induction arguments used in previous relative train track constructions, we prove that CTs satisfying (Inheritance) can be constructed algorithmically.

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 nested sequence $\mathcal{C}$ of ϕ-invariant free factor systems, one can choose $f: G \to G$ so that each non-empty element of $\mathcal{C}$ is realized by a core filtration element.*

A CT $f: G \to G$ representing rotationless $\phi \in \mathsf{Out}(F_n)$ is a graphical representation of ϕ, and may be used to find graphical representations of some important invariants of ϕ. For example, there is an algorithm to compute a finite core graph $S(f)$ immersing to G such that a closed path in G represents a fixed ϕ-conjugacy class if and only if it lifts to a closed path in $S(f)$. There is also a graph $S_N(f)$ immersing to G, obtained from $S(f)$ by attaching finitely many rays, that additionally records fixed points at infinity (see Sections 10 and 12). Arbitrarily large neighborhoods of $S(f)$ in $S_N(f)$ may be computed algorithmically.

The proof of Theorem 1.1 is in Sections 7 and 8. It relies heavily on our paper [FH11] and, more specifically, on the proof of Theorem 4.28 of [FH11] which states that every rotationless outer automorphism of F_n is represented by a CT. We strongly recommend that the reader have a copy of [FH11] handy while reading the current one. We will also rely on [FH11] for complete references. Much of the proof of [FH11, Theorem 4.28] is already algorithmic and the main work in the current paper is to make algorithmic the parts of that proof that are not explicitly algorithmic. In fact, there are only three places in the proof of Theorem 4.28 where non-algorithmic arguments are given. The first has to do with relative train track maps for general elements of $\mathtt{Out}(F_n)$ [BH92]. In this case there is an underlying algorithm to the arguments and we make it explicit in Sections 2 and 3.

The second part of the proof to be made algorithmic involves checking whether the filtration by invariant free factor systems induced by the given filtration by invariant subgraphs is reduced. (See Section 2.1 for a review of the relevant definitions.) This requires new arguments and is carried out in Section 4. The main result of Section 4 is captured in Corollary 1.2 below, which we believe to be of independent interest. Although we present this result as a corollary of Theorem 1.1, a stand-alone proof could be given using the methods in Section 4; see in particular Proposition 4.9. These methods are in turn key ingredients in the proof of the Theorem 1.1.

Recall that if $\psi \in \mathtt{Out}(F_n)$ and $\mathcal{F}_1 \sqsubset \mathcal{F}_2$ are ψ-invariant free factor systems then ψ is said to be *fully irreducible relative to* $\mathcal{F}_1 \sqsubset \mathcal{F}_2$ if for all $k \geq 1$ there are no ψ^k-invariant free factor systems properly contained between $\mathcal{F}_1$ and $\mathcal{F}_2$. Equivalently, if ϕ is a rotationless iterate of ψ then there are no ϕ-invariant free factor systems properly contained between $\mathcal{F}_1$ and $\mathcal{F}_2$; see [FH11, Lemma 3.30].

Corollary 1.2. *There is an algorithm with input $\psi \in \mathtt{Out}(F_n)$ and ψ-invariant free factor systems $\mathcal{F}_1 \sqsubset \mathcal{F}_2$ and output* YES *or* NO *depending on whether or not ψ is fully irreducible relative to $\mathcal{F}_1 \sqsubset \mathcal{F}_2$. In the case that ψ is not fully irreducible, $k \geq 1$ and a ψ^k-invariant free factor system that is properly contained between $\mathcal{F}_1$ and F_2 are found.*

Proof. Construct a CT $f : G \to G$ with filtration $\emptyset = G_0 \subset G_1 \subset \cdots \subset G_N = G$ for a rotationless $\phi = \psi^M$ with M as in Corollary 3.14 in which $\mathcal{F}_1 \sqsubset \mathcal{F}_2$ are realized by core subgraphs $G_r \subset G_t$. Then, by defining property (Filtration) of a CT, ψ is fully irreducible relative to $\mathcal{F}_1 \sqsubset \mathcal{F}_2$ if there are no core subgraphs G_s properly contained between G_r and G_t. If there is such a subgraph then its associated free factor system is properly contained between $\mathcal{F}_1$ and $\mathcal{F}_2$. □

Remark 1.3. In the special case that $\mathcal{F}_1 = \emptyset$, Corollary 1.2 is an algorithm for checking if ψ is fully irreducible. Our algorithm in this special case is different from the ones given in [Kap14] and [CMP15]. More recently, Kapovich [Kap] has produced a polynomial time algorithm to detect full irreducibility.

The third and final non-algorithmic part of the proof of [FH11, Theorem 4.28] requires a new fixed point result (Lemma 6.4) allowing us to properly attach the terminal endpoints of NEG edges; this takes place in Section 6.

We include some sample applications of Theorem 1.1 – most already known.

- In Proposition 9.10, we give another proof of the result of Bogopolski and Maslakova [BM16] that it is algorithmic to compute $\mathrm{Fix}(\Phi)$ for $\Phi \in \mathtt{Aut}(F_n)$.
- An outer automorphism ϕ is *primitively atoroidal* if it does not act periodically on the conjugacy class of a *primitive* element of F_n, i.e. an element of some basis. In Corollary 14.4 we give an algorithm to decide if ϕ is primitively atoroidal.
- In Section 15 we prove that the well-known index invariant $i(\phi)$ of [GJLL98] can be computed algorithmically and introduce a similar invariant $j(\phi) \geq i(\phi)$ that can also be computed algorithmically. In Proposition 15.14 we show that $j(\phi) \leq n-1$ and so the key inequality satisfied by $i(\phi)$ is still satisfied by $j(\phi)$.
- In Corollary 16.4 we reprove a result of Ilya Kapovich [Kap00] that it is algorithmic to tell if a given $\phi \in \mathtt{Out}(F_n)$ is hyperbolic. The Kapovich result is stronger in that ϕ is only assumed to be an injective endomorphism.

We thank the referee for suggesting applications to be added to this paper, specifically: algorithmically computing the index (Proposition 15.2); algorithmically deciding if an automorphism is primitively atoroidal (Corollary 14.4); and algorithmically finding the possibilities for $[\mathrm{Fix}(\Phi)]$ for a Φ representing a rotationless outer automorphism (Corollary 11.1).

2. Relative train track maps in the general case

In this section we revisit three existence theorems for relative train track maps representing arbitrary $\phi \in \mathtt{Out}(F_n)$, rotationless or not. The original statements of these results did not mention algorithms and the original proofs did not emphasize their algorithmic natures. In this section we give the algorithmic versions of two of these results; see Theorem 2.2 and Lemma 2.10. The third existence result that we revisit is [FH11, Theorem 2.19] which in the current paper is reproduced as Theorem 2.12. We will need the algorithmic version of Theorem 2.12 and its proof is postponed until Section 3.4 because the proof depends on consequences of Theorem 2.12 established in Section 3.

Rather than cut and paste arguments from [BH92], [BFH00], and [FH11] into this paper, we will point the reader to specific sections in those papers and explain how they fit together to give the desired results. In some cases, we refer to arguments that occur in lemmas whose hypotheses are not satisfied in our current context. Nonetheless the arguments that we refer to will apply.

2.1. Some standard notation and definitions. In this section we recall the basic definitions of relative train track theory, assuming that the reader has some familiarity with this material. Complete details can be found in any of [BH92, Sections 1 and 5], [BFH00, Sections 2 and 3], [FH11, Section 2], and Part I of [HM].

Identify F_n with $\pi_1(R_n, *)$ where the rose R_n is the graph with one vertex $*$ and n edges. A (not necessarily connected) graph is *core* if it is the union of its immersed circuits. A *marked graph* is a finite core graph G equipped with a homotopy equivalence $\mu\colon R_n \to G$ called the *marking* by which we identify $\mathsf{Out}(\pi_1(G))$ with $\mathsf{Out}(\pi_1(R_n))$ and hence with $\mathsf{Out}(F_n)$. In this way, a homotopy equivalence $f\colon G \to G$ determines an element $\phi \in \mathsf{Out}(F_n)$; we say that $f\colon G \to G$ *represents* ϕ or that $f\colon G \to G$ is a *topological representative of* ϕ. Conversely, each $\phi \in \mathsf{Out}(F_n)$ is represented (non-uniquely) by a homotopy equivalence of any marked graph G, cf. Lemma 7.2.

Convention 2.1. *Unless otherwise stated, we assume that f is an immersion when restricted to edges and that if J is an interval in the interior of an edge then some f-iterate of J either contains a vertex or is contained in a periodic edge. This can be arranged for example by assigning each edge length one and making the restriction of f to each edge linear.*

A (finite, infinite or bi-infinite) path in G is an immersion $\sigma\colon J \to G$ defined on an interval J such that the image of each end of J crosses infinitely many edges of G; equivalently, σ lifts to a proper map into the universal cover $\widetilde{G}$. We allow the possibility that J is a single point in which case we say that the path is *trivial*. Subdividing at the full pre-image of the set $\mathcal{V}$ of vertices of G, we view J as a simplicial complex and σ as a map whose restriction to an edge is one-to-one with image an edge or partial edge. In this way we view σ as an *edge path*; i.e. a concatenation of edges of G, where we allow the first and last to be partial edges if the endpoints are not at vertices. We do not distinguish between paths that have the same associated edge path and we often identify a path with its associated edge path and write $\sigma \subset G$. A path has *height* r if it is contained in G_r but not G_{r-1}.

A bi-infinite path σ is called a *line*. Each lift $\tilde{\sigma} \subset \widetilde{G}$ of a line $\sigma \subset G$ has well defined endpoints in the set ∂G of ends of G. A singly infinite path is called a *ray*. Each lift of a ray has one endpoint in $\widetilde{G}$ and one ideal endpoint in $\partial\widetilde{G}$. Conversely, each ordered pair of distinct points in $\partial\widetilde{G}$ is the endpoint set of a unique line in $\widetilde{G}$. For any marked graph G, one can identify ∂G with ∂F_n; see Section 3. In this way lines in G are identified with F_n-orbits of ordered pairs (P, Q) of distinct points in ∂F_n. We sometimes refer to (P, Q) as an *abstract line* or even just a *line*. Thus each line in G determines an F_n-orbit of abstract lines. Each lift $\tilde{f}\colon \widetilde{G} \to \widetilde{G}$ extends to a homeomorphism (also called $\tilde{f}$) of the set ∂G of the ends of $\widetilde{G}$.

We denote the conjugacy class of a subgroup A of F_n by $[A]$. If $A_1 * \cdots * A_m$ is a free factor of F_n and each A_i is non-trivial then $\{[A_1], \ldots, [A_m]\}$ is a *free factor system* and each $[A_i]$ is a *component* of that free factor system. We also allow the *trivial free factor system* $\emptyset$. For any marked graph G and subgraph C with non-contractible components $C_1, \ldots, C_m$, the fundamental group of each C_i determines a well defined conjugacy class that we denote $[C_i]$ and $[C] := \{[C_1], \ldots, [C_m]\}$ is a free factor system. We say that $C \subset G$ *realizes* $[C]$. Every free factor system is realized by some subgraph of some marked graph. There is a partial order $\sqsubset$ on free factor systems defined by $\{[A_1], \ldots, [A_k]\} \sqsubset \{[B_1], \ldots, [B_l]\}$ if each A_i is conjugate to a subgroup of some B_j. There is a natural action of $\phi \in \mathtt{Out}(F_n)$ on conjugacy classes of free factors. If each element of a free factor system $\mathcal{F}$ is ϕ-invariant then we say that $\mathcal{F}$ *is ϕ-invariant*.

More generally, a *subgroup system* is a finite collection $\{[A_1], \ldots, [A_k]\}$ of distinct conjugacy classes of finitely generated non-trivial subgroups of F_n. The conjugacy class $[c]$ of $c \in F_n$ is *carried by* the subgroup system $\{[A_1], \ldots, [A_k]\}$ if c is conjugate to an element of some A_i. Equivalently, the fixed points for the action of a conjugate of c on ∂F_n are contained in some ∂A_i. More generally if $P, Q \in \partial A_i$ then the F_n-orbit of *the abstract line* with endpoints P and Q *is carried by* $\{[A_1], \ldots, [A_k]\}$. If a subgraph C of a marked graph G realizes a free factor system $\mathcal{F}$ then a conjugacy class is carried by $\mathcal{F}$ if and only if the circuit representing it in G is contained in C; similarly, the F_n-orbit of an abstract line is carried by $\mathcal{F}$ if and only if the corresponding line in G is contained in C. By [BFH00, Lemma 2.6.5] (see also Lemma 4.2), for each collection of conjugacy classes of elements and F_n-orbits of abstract lines there is a unique minimal (with respect to the above partial order $\sqsubset$) free factor system that carries them all. If that minimal free factor system is $\{[F_n]\}$ then we say that the collection *fills*. In this context, we will treat a subgroup as the collection of conjugacy classes that it carries. In particular, it makes sense to talk about a subgroup system filling.

A *filtration* of a marked graph G is an increasing sequence of subgraphs $\emptyset = G_0 \subset G_1 \subset \cdots \subset G_N = G$. The r^{th}*stratum* H_r is the subgraph whose edges are contained in G_r but not G_{r-1}. A homotopy equivalence $f : G \to G$ *preserves the filtration* if $f(G_r) \subset G_r$ for each G_r. Assuming this to be the case and that the edges in H_r have been ordered, the *transition matrix* M_r associated to H_r is the square matrix with one row and column for each edge of H_r and whose ij^{th} coordinate is the number of times that the f-image of the i^{th} edge of H_r crosses (in either direction) the j^{th} edge of H_r. After enlarging the filtration if necessary, we may assume that each M_r is either the zero matrix or irreducible; we say that H_r is *a zero stratum* or *an irreducible stratum* respectively. In the irreducible case, each M_i has a *Perron-Frobenius eigenvalue* $\lambda_r \geq 1$. The stratum H_r is *EG* if $\lambda_r > 1$ and is *NEG* if $\lambda_r = 1$. If $f : G \to G$ is a topological representative of ϕ then we sometimes say that $f : G \to G$ *and* $\emptyset = G_0 \subset G_1 \subset \cdots \subset G_N = G$ *represent* ϕ.

If $\sigma \subset G$ is a path and $f: G \to G$ is a homotopy equivalence then $f(\sigma) := f \circ \sigma: J \to G$ need not be immersed and so need not be a path. If σ is a finite path, we define $f_\#(\sigma)$ to be the unique path that is homotopic to $f(\sigma)$ rel endpoints. For rays and lines, one defines $f_\#(\sigma)$ by choosing a lift $\tilde\sigma$, defining $\tilde f_\#(\tilde\sigma)$ to be the unique path that is homotopic to $\tilde f(\tilde\sigma)$ rel endpoints (including ideal endpoints) and then projecting $\tilde f_\#(\tilde\sigma)$ to a path $f_\#(\sigma) \subset G$.

If paths σ_1 and σ_2 can be concatenated then we denote the concatenation by $\sigma_1\sigma_2$. A decomposition $\sigma = \dots\sigma_1\sigma_2\dots\sigma_m\dots$ of a path into subpaths is a *splitting* if $f^k_\#(\sigma) = \dots f^k_\#(\sigma_1) f^k_\#(\sigma_2)\dots f^k_\#(\sigma_m)\dots$ for all $k \geq 1$; in this case we usually write $\sigma = \dots\cdot\sigma_1\cdot\sigma_2\cdot\dots\cdot\sigma_m\cdot\dots$.

If $f^k_\#(\sigma) = \sigma$ for some $k \geq 1$ and finite path σ then σ is a *periodic Nielsen path*. If $k = 1$ then σ is *a Nielsen path*. If a (periodic) Nielsen path σ can not be written as the concatenation of non-trivial (periodic) Nielsen subpaths then it is *indivisible*. Two points in $\mathrm{Fix}(f)$ are in the same *Nielsen class* if they bound a Nielsen path.

An edge E in a (necessarily NEG) stratum H_i is *linear* if $f(E) = Eu$ where $u \subset G_{i-1}$ is a non-trivial Nielsen path. If $u = w^d$ for some root-free w and some $d \neq 0$ then the unoriented conjugacy class of w is called the *axis* for E. All other edges in NEG strata are *non-linear*.

A *direction* at a point $x \in G$ is a germ of non-trivial finite paths with initial vertex x. If x is not a vertex then there are two directions at x. Otherwise there is one direction for each oriented edge based at x and we identify the direction with the oriented edge. A homotopy equivalence $f: G \to G$ induces a map Df from directions at x to directions at $f(x)$. A *turn* at x is an unordered pair (d_1, d_2) of directions based at x; it is *degenerate* if $d_1 = d_2$ and *non-degenerate* otherwise. Df induces a map on turns that we also denote by Df. A turn is *illegal* if its image under some iterate of Df is degenerate and is *legal* otherwise. If $\sigma = \dots E_i E_{i+1}\dots$ and if each turn $(\bar E_i, E_{i+1})$ is legal then σ *is legal*. Here $\bar E_i$ denotes E_i with the opposite orientation. We sometimes also use the exponent -1 to indicate the inverse of a path. If σ has height r then σ is r*-legal* if each turn $(\bar E_i, E_{i+1})$ for which both E_i and E_{i+1} are edges in H_r is legal.

The homotopy equivalence $f: G \to G$ is a *relative train track map* [BH92, p. 38] if it maps vertices to vertices and if the following conditions hold for every EG stratum H_r.

(RTT-i) Df maps directions in H_r to directions in H_r.

(RTT-ii) If $\sigma \subset G_{r-1}$ is a non-trivial path with endpoints in $H_r \cap G_{r-1}$ then $f_\#(\sigma)$ is a non-trivial path with endpoints in $H_r \cap G_{r-1}$.

(RTT-iii) If $\sigma \subset H_r$ is legal then $f(\sigma)$ is an r-legal path.

2.2. Algorithmic proofs. The following theorem is modeled on [BH92, Theorem 5.12] which proves the existence of relative train track maps that satisfy an additional condition called *stability*. This condition is algorithmically built into

CTs at a later stage of the argument (see [FH11, Step 1(EG Nielsen paths) in Section 4.5]) so we do not need it here. See also [DV96]. The proof of [BH92, Theorem 5.12] is mostly algorithmic. We have added Lemmas 2.4 and 2.7 to make the entire argument algorithmic.

Theorem 2.2. *There is an algorithm that produces for each $\phi \in \mathtt{Out}(F_n)$ a relative train track map $f : G \to G$ representing ϕ.*

The proof of Theorem 2.2 appears at the end of this section after we recall some definitions and results from [BH92].

Definition 2.3. Suppose that $f : G \to G$ and $\emptyset = G_0 \subset G_1 \subset \cdots \subset G_N = G$ are a topological representative and filtration representing ϕ. For each EG stratum H_r, let $\lambda_r \geq 1$ be the Perron–Frobenious eigenvalue of the transition matrix M_r associated to H_r. Let $\Lambda(f)$ be the set of (not necessarily distinct) λ_r's associated to EG strata H_r of $f : G \to G$, listed in non-increasing order. Say that $f : G \to G$ is *bounded* if there are at most $3n-3$ exponentially growing strata H_r and if each λ_r is the Perron–Frobenious eigenvalue of some irreducible matrix with at most $3n-3$ rows and columns. Note that if each vertex of G has valence at least 3 then $f : G \to G$ is bounded because G has at most $3n-3$ edges. Note also that if the set of all possible $\Lambda(f)$'s is ordered lexicographically, then any strictly decreasing sequence of $\Lambda(f)$'s associated to bounded f's is finite; see [BH92, p. 37].

Lemma 2.4. *There is an algorithm that checks if a given bounded topological representative $f : G \to G$ of ϕ is a relative train track map.*

Proof. Property (RTT-i) is obviously a finite property. We may therefore assume that each EG stratum satisfies (RTT-i). Suppose that H_j is an EG stratum and that C is a component of G_{j-1}. If C is contractible then it contains only finitely many paths with endpoints at vertices so checking (RTT-ii) for paths in C is a finite process. If C is not contractible then there is a smallest $p \geq 1$ such that $f^p(C) \subset C$. Since H_j satisfies (RTT-i), $f^p(H_j \cap C) \subset H_j \cap C$. If f^p induces a bijection of $H_j \cap C$ then (RTT-ii) is satisfied for all $\sigma \subset C$. Otherwise, there exist distinct $v, w \in H_j \cap C$ and a smallest $1 \leq q \leq p$ such that $f^q(v) = f^q(w)$. After replacing v and w by $f^{q-1}(v)$ and $f^{q-1}(w)$, we may assume that $q = 1$. In this case v and w are connected by a unique path $\sigma \subset C$ whose $f_\#$-image is trivial and (RTT-ii) fails. We have now proved that (RTT-ii) is a finite property. Finally, (RTT-iii) is equivalent to the statement that $f(E)$ is j-legal for each edge $E \subset H_j$ and so is a finite property. $\square$

Definition 2.5. [BH92, paragraph before Lemma 5.13] Suppose that $f : G \to G$ is a topological representative and that E is an edge in an EG stratum H_r. The *core of E* is defined to be the smallest closed subinterval of E such that each point in the complement of the core is eventually mapped into G_{r-1}. The set of endpoints of the cores of all the edges in H_j is mapped into itself by f. Subdivision at this finite set (and enlarging the original filtration to accommodate the new edges formed by the complements of the core) is called the *core subdivision of H_r*.

Lemma 2.6. [BH92, Lemma 5.13] *Suppose that $f: G \to G$ is a bounded topological representative of ϕ and that $f': G' \to G'$ is obtained from $f: G \to G$ by a core subdivision of the EG stratum H_r. Then,*

(1) $\Lambda(f) = \Lambda(f')$ *and*

(2) *there is a bijection $H_j \leftrightarrow H'_{j'}$ between the EG strata of f and the EG strata of f' such that*

   (a) *$H'_{j'}$ satisfies* (RTT-i);

   (b) *relative height is preserved*; *i.e. $j < k$ if and only if $j' < k'$*;

   (c) *if $j \neq r$ and H_j satisfies* (RTT-i) *or* (RTT-ii) *then $H'_{j'}$ satisfies* (RTT-i) *or* (RTT-ii), *respectively.*

Lemma 2.7. *Core subdivision is algorithmic.*

Proof. Suppose that $f: G \to G$ is a topological representative and that E is an edge in an EG stratum H_r. If the Df orbit of E (thought of as the initial direction of E) is contained in H_r then the core of E contains an initial segment of E and E does not contribute to the set of core subdivision points. Suppose then that some iterate of Df maps E into G_{r-1}. Define $D_r f(E)$ to be the first H_r-edge in the edge path $f(E)$. If E is $D_r f$-periodic of minimal period p, let $E_i = D_r f^i(E)$ for $i = 0, \dots, p-1$. Thus $f(E_i) = u_i E_{i+1} v_i$ where u_i is a possibly trivial path in G_{r-1} and indices are taken mod p. Consider the subintervals of E_i whose images under f^p are single edges. The first such subinterval e_i that is mapped into H_r satisfies $f^p(e_i) = E_i$ and we choose the core subdivision point corresponding to E_i to be the unique point in e_i that is fixed by f^p. If E is not $D_r f$-periodic then choose the minimal q such that $D_r f^q(E)$ is $D_r f$-periodic and take the core subdivision point for E to be the first point in E that is mapped to the core subdivision point of $D_r f^q(E)$. □

Lemma 2.8. *Given a bounded topological representative $f: G \to G$ of ϕ with an EG stratum that satisfies* (RTT-i), *but not* (RTT-iii) *there is an algorithm to construct a bounded topological representative $f'': G'' \to G''$ of ϕ such that $\Lambda(f'') < \Lambda(f)$.*

Proof. The proof of [BH92, Lemma 5.9] contains an algorithm that modifies $f: G \to G$ to produce a not necessarily bounded, topological representative $f': G' \to G'$ such that $\Lambda(f') < \Lambda(f)$. The proof of [BH92, Lemma 5.5] contains an algorithm that modifies $f': G' \to G'$ to produce a bounded topological representative $f'': G'' \to G''$ of ϕ such that $\Lambda(f'') \leq \Lambda(f') < \Lambda(f)$. □

We next recall [BH92, Lemma 5.14], making the algorithm used in its proof part of the statement of the lemma.

Lemma 2.9. [BH92, Lemma 5.14] *Suppose that $f\colon G \to G$ is a bounded topological representative of ϕ and that H_s is an EG stratum that does not satisfy* (RTT-ii). *Then there is an algorithm to construct a bounded topological representative $f'\colon G' \to G'$ of ϕ such that*

(1) $\Lambda(f) = \Lambda(f')$ *and*

(2) *there is a bijection $H_j \leftrightarrow H'_{j'}$ between the EG strata of f and the EG strata of f' such that*

  (a) *relative height is preserved*; *i.e. $j < k$ if and only if $j' < k'$*;

  (b) $|H'_{s'} \cap G'_{s'-1}| < |H_s \cap G_{s-1}|$;

  (c) *if $k > s$ and H_k satisfies* (RTT-i) *and* (RTT-ii) *then $H'_{k'}$ satisfies* (RTT-i) *and* (RTT-ii);

  (d) *if $k \geq s$ and H_k satisfies* (RTT-i) *then $H'_{k'}$* (RTT-i).

Proof of Theorem 2.2. Start with any bounded topological representative $f\colon G \to G$ of ϕ. For example, one can choose any homotopy equivalence of the rose that represents ϕ. If there are no EG strata then $f\colon G \to G$ is a relative train track map and we are done. Otherwise, apply Lemma 2.6 to produce a bounded topological representative (still called $f\colon G \to G$) of ϕ whose top EG stratum satisfies (RTT-i). If the top EG stratum does not also satisfy (RTT-ii), apply Lemma 2.9 to produce a new bounded topological representative (still called $f\colon G \to G$) of ϕ whose top EG stratum still satisfies (RTT-i). If the top EG stratum of the current $f\colon G \to G$ does not satisfy (RTT-ii) apply Lemma 2.9 again. Item (b) of that lemma guarantees that after finitely many applications of Lemma 2.9, we arrive at $f\colon G \to G$ whose top EG stratum satisfies (RTT-i) and (RTT-ii).

Repeat this procedure on the second highest EG stratum to produce a bounded topological representative of ϕ whose top two EG stratum satisfies (RTT-i) and (RTT-ii). After finitely many iterations, we have a bounded topological representative (still called $f\colon G \to G$) of ϕ, all of whose EG strata satisfy (RTT-i) and (RTT-ii).

Apply Lemma 2.4 to check if $f\colon G \to G$ is a relative train track map. If yes, we are done. Otherwise apply Lemma 2.8 to produce a bounded topological representative $f'\colon G' \to G'$ of ϕ with $\Lambda(f') < \Lambda(f)$. Then start over again with $f'\colon G' \to G'$ replacing the original $f\colon G \to G$. Since every decreasing sequence $\Lambda(f) > \Lambda(f') > \ldots$ is finite, this process produces a relative train track map in finite time. $\square$

Corollary 2.10. *There is an algorithm that takes $\phi \in \mathsf{Out}(F_n)$ and a nested sequence $\mathcal{C} = \mathcal{F}_1 \sqsubset \mathcal{F}_2 \sqsubset \cdots \sqsubset \mathcal{F}_m$ of ϕ-invariant free factor systems as input and produces a relative train track map $f : G \to G$ and filtration $\emptyset = G_0 \subset G_1 \subset \cdots \subset G_N = G$ representing ϕ and such that for each $\mathcal{F}_i$ there exists G_j satisfying $\mathcal{F}_i = [G_j]$.*

Proof. The proof of this corollary is explicitly contained in the proof of [BFH00, Lemma 2.6.7] (even though the statement of that lemma is weaker in that it assumes that $\mathcal{C}$ is a single free factor system). The first step of the proof of the lemma is to inductively construct a bounded topological representative $f : G \to G$ and filtration $\emptyset = G_0 \subset G_1 \subset \cdots \subset G_N = G$ representing ϕ such that for each $\mathcal{F}_i$ there exists G_j satisfying $\mathcal{F}_i = [G_j]$. Then one applies the relative train track algorithm of Theorem 2.2, checking that $\mathcal{C}$ is preserved, to promote $f : G \to G$ to a relative train track map. □

The third existence theorem that needs discussion is [FH11, Theorem 2.19]. We first recall some notation that is used in its statement.

Notation 2.11. If $u < r$ and

(1) H_u is irreducible;

(2) H_r is EG and each component of G_r is non-contractible; and

(3) for each $u < i < r$, H_i is a zero stratum that is a component of G_{r-1} and each vertex of H_i has valence at least two in G_r

then we say that each H_i is *enveloped by* H_r and write $H_r^z = \bigcup_{k=u+1}^{r} H_k$.

Theorem 2.12 ([FH11, Theorem 2.19]). *For each $\phi \in \mathsf{Out}(F_n)$ there is a relative train track map $f : G \to G$ and a filtration $\emptyset = G_0 \subset G_1 \subset \cdots \subset G_N = G$ representing ϕ and satisfying the following properties.*

(V) *The endpoints of all indivisible periodic Nielsen paths are vertices.*

(P) *If a stratum $H_m \subset \mathrm{Per}(f)$ is a forest then there exists a filtration element G_j such $[G_j] \neq [G_l \cup H_m]$ for any G_l.*

(Z) *Each zero stratum H_i is enveloped by an EG stratum H_r. Each vertex in H_i is contained in H_r and has link contained in $H_i \cup H_r$.*

(NEG) *The terminal endpoint of an edge in a non-periodic NEG stratum H_i is periodic and is contained in a filtration element of height less than i that is its own core.*

(F) *The core of each filtration element is a filtration element.*

Moreover, if $\mathcal{C}$ is a nested sequence of ϕ-invariant free factor systems then we may choose $f : G \to G$ so that for each $\mathcal{F}_i \in \mathcal{C}$ there exists G_j satisfying $\mathcal{F} = [G_j]$.

The proof that $f : G \to G$ as in Theorem 2.12 can be constructed algorithmically is contained in Section 3.4

3. Rotationless iterates

Every element of $\mathtt{Out}(F_n)$ has an iterate that is rotationless (Definition 3.3). Corollary 3.14 below gives an explicit bound on the size of the iterate; see also [FH11, Lemma 4.42] for a proof that such a bound exists.

3.1. More on markings. In this subsection we discuss markings in more detail and recall definitions and results from [FH11, Section 3]. We assume throughout this subsection that $f: G \to G$ is a relative train track map representing $\phi \in \mathtt{Out}(F_n)$.

Markings are used to translate the geometric properties of $f: G \to G$ into algebraic properties of ϕ. In this paper, we will focus on the geometric properties of the homotopy equivalences and only bring in markings at the last minute when necessary. Further details on the material presented in this section can be found in [FH11, Section 2.3].

Recall that the rose R_n denotes the rose with vertex $*$ and that we have once and for all identified $\pi_1(R_n, *)$ with F_n. A lift $\tilde{*} \in \widetilde{R}_n$ of $*$ to the universal cover $\widetilde{R}_n$ determines an isomorphism $J_{\tilde{*}}$ from $F_n = \pi_1(R_n, *)$ to the group $\mathcal{T}(\widetilde{R}_n)$ of covering translations of $\widetilde{R}_n$ given by $[\gamma]$ maps to the covering translation T of $\widetilde{R}_n$ that takes $\tilde{*}$ to the terminal endpoint of the lift of γ with initial endpoint $\tilde{*}$.

Let G be a finite graph equipped with a marking $\mu\colon R_n \to G$. Denoting $\mu(*)$ by $\star$, $\mu\colon (R_n, *) \to (G, \star)$ induces an isomorphism $\mu_\#\colon \pi_1(R_n, *) \to \pi_1(G, \star)$ that identifies F_n with $\pi_1(G, \star)$. Fix a lift $\tilde{\star}$ of $\star$ to $\widetilde{G}$. The lift $\tilde{\mu}\colon (\widetilde{R}_n, \tilde{*}) \to (\widetilde{G}, \tilde{\star})$ determines a homeomorphism $\partial\tilde{\mu}\colon \partial\widetilde{R}_n \to \partial\widetilde{G}$ of Gromov boundaries. In this way ∂F_n, $\partial\widetilde{R}_n$, and $\partial\widetilde{G}$ are all identified. Since covering translations are determined by their action on Gromov boundaries, there is an induced identification of $\mathcal{T}(\widetilde{R}_n)$ with $\mathcal{T}(\widetilde{G})$. For any $v \in G$ and lift $\tilde{v} \in \widetilde{G}$, there is an induced isomorphism $J_{\tilde{v}}\colon \pi_1(G, v) \to \mathcal{T}(\widetilde{G})$ defined exactly as $J_{\tilde{*}}$. It is straightforward to check that $\mu_\# = J_{\tilde{\star}}^{-1} J_{\tilde{*}}\colon \pi_1(R_n, *) \to \pi_1(G, \star)$.

We also have an identification of automorphisms representing $\phi \in \mathtt{Out}(F_n)$ with lifts $\tilde{f}: \widetilde{G} \to \widetilde{G}$ of $f: G \to G$ given by $\Phi \leftrightarrow \tilde{f}$ if the actions of Φ and $\tilde{f}$ on ∂F_n agree, i.e. if $\partial\Phi = \partial\tilde{f}$. We usually specify $\tilde{f}$ by specifying $\tilde{f}(\tilde{\star})$ or equivalently by specifying the path $\tilde{\rho} = [\tilde{\star}, \tilde{f}(\tilde{\star})]$ or its image ρ in G. We say that Φ or $\tilde{f}$ is *determined by* $\tilde{f}(\tilde{\star})$*,* $\tilde{\rho}$*, or* ρ. The action of Φ on $\pi_1(G, \star)$ is given by $\gamma \mapsto f(\gamma)^\rho := \rho f(\gamma)\bar{\rho}$. If $\tilde{f}$ is determined by ρ and $\tilde{f}'$ is determined by ρ' then $\Phi' = i_\gamma \Phi$ where $\Phi \leftrightarrow \tilde{f}$, $\Phi' \leftrightarrow \tilde{f}'$, and $\gamma \in F_n$ is represented by the loop $\rho'\bar{\rho}$. Working in the universal cover $\widetilde{G}$ is algorithmic in the sense that we can always compute the action of $\tilde{f}$ on arbitrarily large balls (in the graph metric) around $\tilde{\star}$. In particular, given Φ we may algorithmically find $\tilde{f}$ with $\Phi \leftrightarrow \tilde{f}$ and *vice versa*. If $\tilde{f}$ fixes $\tilde{v} \in \widetilde{G}$ then $\tilde{f} \leftrightarrow \Phi$ for Φ determined by ρ where $\tilde{\rho} = \tilde{\eta}\tilde{f}(\tilde{\eta}^{-1})$ and $\tilde{\eta} = [\tilde{\star}, \tilde{v}]$.

Definition 3.1. For $\tilde f:\tilde G \to \tilde G$ a lift of $f:G \to G$, we denote the subgroup of $\mathcal{T}(\tilde G)$ consisting of covering translations that commute with $\tilde f$ by $Z_{\mathcal{T}}(\tilde f)$.

We state the following well known fact (see for example [FH11, Lemma 2.1]) as a lemma for easy reference.

Lemma 3.2. *If $\tilde f$ corresponds to Φ as above then $Z_{\mathcal{T}}(\tilde f)$ and* $\mathrm{Fix}(\Phi)$ *are equal when viewed as subgroups of $\mathcal{T}(\tilde G)$.*

Automorphisms $\Phi_1, \Phi_2 \in \mathsf{Aut}(F_n)$ are *isogredient* if $\Phi_1 = i_a \Phi_2 i_a^{-1}$ for some inner automorphism i_a. Lifts $\tilde f_1$ and $\tilde f_2$ of f are *isogredient* if the corresponding automorphisms are isogredient. That is, $\tilde f_1$ and $\tilde f_2$ are isogredient if there exists a covering translation T of $\tilde G$ such that $\tilde f_2 = T \tilde f_1 T^{-1}$. The set of *attracting laminations* for $\phi \in \mathsf{Out}(F_n)$ is denoted $\mathcal{L}(\phi)$; see [BFH00, Section 3].

3.2. Principal automorphisms and principal points. Recall from [GJLL98] that for each $\Theta \in \mathsf{Aut}(F_n)$,

$$\mathrm{Fix}(\partial\Theta) = \mathrm{Fix}_-(\partial\Theta) \cup \mathrm{Fix}_+(\partial\Theta) \cup \partial\,\mathrm{Fix}(\Theta)$$

where $\mathrm{Fix}(\Theta)$ is the fixed subgroup for Φ, $\mathrm{Fix}_-(\partial\Theta) \subset \partial F_n$ is a finite union of $\mathrm{Fix}(\Theta)$-orbits of isolated repellers and $\mathrm{Fix}_+(\partial\Theta) \subset \partial F_n$ is a finite union of $\mathrm{Fix}(\Theta)$-orbits of isolated attractors.

Associated to each $\phi \in \mathsf{Out}(F_n)$ is a finite set $\mathcal{L}(\phi)$ of *attracting laminations*, each a closed subset of abstract lines which we refer to as *leaves* of the lamination. A leaf γ of $\Lambda \in \mathcal{L}(\phi)$ is *generic* in Λ if both of its ends are dense in Λ. See [BFH00, Section 3.1].

Definition 3.3 (Definition 3.1 in [FH11]). For $\Phi \in \mathsf{Aut}(F_n)$ representing ϕ, denote the set of non-repelling fixed points of $\partial\Phi$ by $\mathrm{Fix}_N(\partial\Phi)$. We say that Φ is a *principal automorphism* and write $\Phi \in \mathrm{P}(\phi)$ if either of the following hold.

- $\mathrm{Fix}_N(\partial\Phi)$ contains at least three points.
- $\mathrm{Fix}_N(\partial\Phi)$ is a two point set that is neither the set of fixed points for the action of some non-trivial $a \in F_n$ on ∂F_n nor the set of endpoints of a lift of a generic leaf of an element of $\mathcal{L}(\phi)$.

If $f:G \to G$ is a topological representative of ϕ and $\tilde f:\tilde G \to \tilde G$ is the lift corresponding to principal Φ then $\tilde f$ is a *principal lift*.

If $\Phi \in \mathrm{P}(\phi)$ and $k > 1$ then $\mathrm{Fix}_N(\partial\Phi) \subset \mathrm{Fix}_N(\partial\Phi^k)$ and $\Phi^k \in \mathrm{P}(\phi^k)$. It may be that the injection $\Phi \mapsto \Phi^k$ of $\mathrm{P}(\phi)$ into $\mathrm{P}(\phi^k)$ is not surjective. It may also be that $\mathrm{Fix}_N(\partial\Phi^k)$ properly contains $\mathrm{Fix}_N(\partial\Phi)$ for some principal Φ and some $k > 1$. If neither of these happen then we say that ϕ is *forward rotationless*. For a formal definition, see [FH11, Definition 3.13].

Remark 3.4. It is becoming common usage to suppress the word "forward" in "forward rotationless" and we will follow that convention in this paper. So, when we say that $\phi \in \mathsf{Out}(F_n)$ is *rotationless*, we mean that ϕ is forward rotationless. This convention was followed in the recent work of Handel and Mosher [HM]. Be aware though that the term "rotationless" has a slightly different meaning in [FH09].

Suppose that $f: G \to G$ is a topological representative of ϕ. By [FH11, Corollary 3.17], $\mathrm{Fix}(\tilde{f}) \neq \emptyset$ for each principal lift $\tilde{f}$. The projected image of $\mathrm{Fix}(\tilde{f})$ is exactly a Nielsen class in $\mathrm{Fix}(f)$ and a pair of principal lifts are isogredient if and only if they determine the same Nielsen class of $\mathrm{Fix}(f)$ [FH11, Lemma 3.8].

Definition 3.5. We say that $x \in \mathrm{Per}(f)$ is *principal* if neither of the following conditions are satisfied.

- x is not an endpoint of a non-trivial periodic Nielsen path and there are exactly two periodic directions at x, both of which are contained in the same EG stratum.
- x is contained in a component C of $\mathrm{Per}(f)$ that is topologically a circle and each point in C has exactly two periodic directions.

If each principal periodic vertex is fixed and if each periodic direction based at a principal periodic vertex is fixed then we say that f is *rotationless*.

Remark 3.6. By definition, a point is principal with respect to f if and only if it is principal with respect to f^k for all $k \geq 1$.

Remark 3.7. Definition 3.5 is a corrected version of Definition 3.18 of [FH11] in which 'x is not an endpoint of a non-trivial Nielsen path' in the first item of Definition 3.5 is replaced with the inequivalent condition 'x is the only point in its Nielsen class.' Our thanks to Lee Mosher who pointed this out to us. Fortunately, the definition we give here and not the one given in [FH11] is the one that is actually used in [FH11] so no further corrections to [FH11] are necessary.

We are mostly interested in the case of a CT, where characterizations of principal points are simpler. The next lemma gives two.

Lemma 3.8. *Suppose $f: G \to G$ is a* CT.

(1) *A point $x \in \mathrm{Per}(f)$ is principal if and only if $x \in \mathrm{Fix}(f)$ and the following condition is not satisfied*:

   - *x is not an endpoint of a non-trivial Nielsen path and there are exactly two periodic directions at x, both of which are contained in the same EG-stratum.*

(2) *The following are equivalent for a point* $x \in \mathrm{Fix}(f)$. *Let* $\tilde f : \widetilde G \to \widetilde G$ *be a lift of* f *fixing a lift* $\tilde x$ *of* x:

(a) x *is principal*;

(b) $\tilde f$ *is principal*;

(c) $\mathrm{Fix}_N(\partial \tilde f^2)$ *is not the set of endpoints of a generic leaf of an element of* $\mathcal{L}(\phi)$.

Proof. (1) Periodic Nielsen paths in a CT are fixed [FH11, Lemma 4.13] and so the bulleted item in the lemma is equivalent to the first bulleted item of the Definition 3.5. By definition, periodic edges of a CT are fixed and the endpoints of fixed edges are principal. Therefore the second item in Definition 3.5 never holds. To complete the proof it remains to show that all principal points of f are fixed. This holds for vertices because CTs are rotationless. If x is a periodic but not fixed point in the interior of an edge then (by definition of a CT) that edge must be in an EG stratum and so x is not principal.

(2) By [FH11, Corollaries 3.22 and 3.27], (2a) and (2b) are equivalent. If $\tilde f$ is principal, then by definition of rotationless and principal, $\mathrm{Fix}_N(\partial \tilde f) = \mathrm{Fix}_N(\partial \tilde f^2)$ is not contained in the set of endpoints of a generic leaf. We see (2b) implies (2c). If x is not principal for f then the bulleted item in (1) holds. In particular there are exactly two periodic directions at x, both of which are in the same EG-stratum. By [FH11, Lemma 2.13], $\mathrm{Fix}_N(\partial \tilde f^2)$ contains the set of endpoints of a generic leaf of an element of $\mathcal{L}(\phi)$. By Remark 3.6, x is not principal for f^2, and so $\tilde f^2$ is not a principal lift. Hence $|\mathrm{Fix}_N(\partial \tilde f^2)| < 3$. We conclude (2c) implies (2a). □

3.3. A sufficient condition to be rotationless and a uniform bound. Before turning to Lemma 3.12, which gives a sufficient condition for an outer automorphism to be rotationless, we recall the connection between edges in a CT and elements of $\mathrm{Fix}_+(\Phi)$.

Definition 3.9. Given a CT $f : G \to G$ representing ϕ, let $\mathcal{E}$ (or $\mathcal{E}_f$) be the set of oriented, non-fixed, and non-linear edges in G whose initial vertex is principal and whose initial direction is fixed by Df. For each $E \in \mathcal{E}$, there is a path u such that $f^k_\#(E) = E \cdot u \cdot f_\#(u) \cdot \dots \cdot f^{k-1}_\#(u)$ for all $k \geq 1$ and such that $|f^k_\#(u)| \to \infty$ with k. The union of the increasing sequence

$$E \subset f(E) \subset f^2_\#(E) \subset \cdots$$

of paths in G is a ray R_E. Each lift $\widetilde R_E$ of R_E to the universal cover of G has a well-defined terminal endpoint $\partial \widetilde R_E \in \partial F_n$ and so R_E determines an F_n-orbit ∂R_E in ∂F_n.

Lemma 3.10. *Suppose that $f: G \to G$ is a* CT *and that $E \in \mathcal{E}$. If $\widetilde{E}$ is a lift of E and $\tilde{f}$ is the lift of f that fixes the initial endpoint of $\widetilde{E}$ then the lift $\widetilde{R}_{\widetilde{E}}$ of R_E that begins with $\widetilde{E}$ converges to a point in* $\mathrm{Fix}_+(\partial \tilde{f})$. *Moreover, $E \mapsto \partial R_E$ defines a surjection $\mathcal{E} \to \big(\bigcup_{\Phi \in \mathrm{P}(\phi)} \mathrm{Fix}_+(\partial \Phi)\big)/F_n$.*

Proof. Suppose that x is the initial endpoint of $E \in \mathcal{E}$, that $\tilde{x}$ is a lift of x, that $R_{\widetilde{E}}$ is the lift of R_E that begins at $\tilde{x}$ and that $\tilde{f}: \widetilde{G} \to \widetilde{G}$ is the lift of f that fixes $\tilde{x}$. Lemma 3.8(2) implies that $\tilde{f}$ is a principal lift and [FH11, Lemma 4.36(1)] implies that $\widetilde{R}_{\widetilde{E}}$ converges to a point $\partial \widetilde{R}_E \in \mathrm{Fix}_N(\partial \Phi)$ where Φ is the principal automorphism corresponding to $\tilde{f}$. Since $|f^k_\#(u)| \to \infty$, it follows [GJLL98, Proposition I.1] that $\partial \widetilde{R}_{\widetilde{E}} \in \mathrm{Fix}_+(\partial \Phi)$. [FH11, Lemma 4.36(2)] implies that $E \mapsto \partial R_E$ is surjective. □

Remark 3.11. By [GJLL98, Proposition I.1], $P \in \mathrm{Fix}_+(\partial \Phi)$ is not fixed by any i_a and so is not fixed by $\partial \Phi'$ for any $\Phi' \neq \Phi$ representing ϕ. Thus $\bigcup_{\Phi \in \mathrm{P}(\phi)} \mathrm{Fix}_+(\partial \Phi)$ is a disjoint union.

Lemma 3.12. *Suppose that $\theta \in \mathsf{Out}(F_n)$ acts trivially on $H_1(F_n; \mathbb{Z}/3\mathbb{Z})$ and induces the trivial permutation on $\big(\bigcup_{\Phi \in \mathrm{P}(\phi)} \mathrm{Fix}_+(\partial \Phi)\big)/F_n$ for some (any) rotationless iterate $\phi = \theta^L$ of θ. Then θ is rotationless.*

Proof. We show below that

$*$ for any $\Phi \in \mathrm{P}(\phi)$, there is $\Theta \in \mathrm{P}(\theta)$ with the property that $\mathrm{Fix}_N(\partial \Phi) \subset \mathrm{Fix}_N(\partial \Theta)$.

To see why this is sufficient to prove the lemma, let $\Theta_k \in \mathrm{P}(\theta^k)$ for some $k \geq 1$. Since $\theta^{kL} = \phi^k$, $\Theta_k^L \in \mathrm{P}(\phi^k)$. Since ϕ is rotationless, there exists $\Phi \in \mathrm{P}(\phi)$ such that $\Theta_k^L = \Phi^k$ and $\mathrm{Fix}_N(\partial \Phi) = \mathrm{Fix}_N(\partial \Phi^k) = \mathrm{Fix}_N(\partial \Theta_k^L)$. By $(*)$, there is $\Theta \in \mathrm{P}(\theta)$ such that

$$\mathrm{Fix}_N(\partial \Theta_k) \subset \mathrm{Fix}_N(\partial \Theta_k^L) = \mathrm{Fix}_N(\partial \Phi) \subset \mathrm{Fix}_N(\partial \Theta) \subset \mathrm{Fix}_N(\partial \Theta^k)$$

It follows that $\Theta_k = \Theta^k$. We have now seen that $\mathrm{P}(\theta) \to \mathrm{P}(\theta^k)$ given by $\Theta \mapsto \Theta^k$ is surjective. By [FH11, Definition 3.13 and Remark 3.14], to show that θ is rotationless it remains to show that $\mathrm{Fix}_N(\partial \Theta^k) = \mathrm{Fix}_N(\partial \Theta)$ for all $\Theta \in \mathrm{P}(\theta)$ and $k \geq 1$. This follows from the above displayed sequence of inclusions by taking $\Theta_k := \Theta^k$.

We now turn to the proof of $(*)$. Set $\mathbb{F} := \mathrm{Fix}(\Phi)$. We claim that there exists Θ representing θ such that $\mathbb{F} \subset \mathrm{Fix}(\Theta)$. If the rank of $\mathbb{F}$ is < 2 then this follows from [HM, Part II Theorem 4.1], which implies that θ fixes each conjugacy class that is fixed by ϕ and in particular fixes each conjugacy class represented by an element of $\mathrm{Fix}(\Phi)$.

Suppose then that $\mathbb{F}$ has rank ≥ 2. We recall two facts.

- Each element of F_n is fixed by only finitely many elements of $\mathrm{P}(\phi)$ and the root-free ones that are fixed by at least two such automorphisms determine only finitely many conjugacy classes; see [FH11, Lemma 4.40].
- $\mathbb{F}$ is its own normalizer in F_n. Proof: Since $\mathbb{F}$ is finitely generated and has rank > 1, we can choose $x \in \partial\mathbb{F} \subset \partial F_n$ that is not fixed by any ∂i_a, $a \in F_n \setminus \{1\}$, and so is not fixed by $\partial\Phi'$ for any automorphism $\Phi' \neq \Phi$ representing ϕ. If y normalizes $\mathbb{F}$ then $\partial i_y(x) \in \partial\mathbb{F}$ and x is fixed by $\partial\Phi'$ where $\Phi' = i_{y^{-1}}\Phi i_y = i_{y^{-1}\Phi(y)}\Phi$. It follows that $\Phi(y) = y$ and hence $y \in \mathbb{F}$.

By the first bullet, we may choose a basis $\{b_j\}$ for $\mathbb{F}$ consisting of elements that are not fixed by any other element of $\mathrm{P}(\phi)$. Applying [HM, Part II Theorem 4.1] again, choose an automorphism Θ_j representing θ and fixing b_j. The automorphism $\Theta_j\Phi\Theta_j^{-1}$ fixes b_j by construction and belongs to $\mathrm{P}(\phi)$ by [FH09, Lemma 2.6] and the fact that θ and ϕ commute. By uniqueness, $\Theta_j\Phi\Theta_j^{-1} = \Phi$ and so Θ_j commutes with Φ. In particular, Θ_j preserves $\mathbb{F}$. Since $\mathbb{F}$ is its own normalizer, the outer automorphism $\theta|\mathbb{F}$ of $\mathbb{F}$ determined by Θ_j is independent of j. It follows that $\theta|\mathbb{F}$ acts trivially on $H_1(\mathbb{F};\mathbb{Z}/3\mathbb{Z})$. Since $\theta^L|\mathbb{F}$ is the identity and since the kernel of natural map $\mathtt{Out}(F_n) \to H_1(F_n;\mathbb{Z}/3\mathbb{Z})$ is torsion-free[1], $\theta|\mathbb{F}$ is the identity and the claim is proved.

Since θ acts trivially on $\big(\bigcup_{\Phi\in\mathrm{P}(\phi)} \mathrm{Fix}_+(\partial\Phi)\big)/F_n$, each $Q \in \mathrm{Fix}_+(\partial\Phi)$ is fixed by some Θ_Q representing θ. Since Θ_Q^L and Φ both fix Q and represent ϕ we have $\Theta_Q^L = \Phi$. As above, Θ_Q commutes with Φ and so preserves $\mathbb{F}$ and $\mathrm{Fix}_+(\partial\Phi)$. For any other $Q' \in \mathrm{Fix}_+(\partial\Phi)$ we have $\Theta_Q = i_a\Theta_{Q'}$ for some $a \in F_n$. Since both Θ_Q and $\Theta_{Q'}$ preserve $\mathbb{F}$, i_a does as well and so $a \in \mathbb{F}$.

It suffices to show that Θ_Q is independent of Q and $\mathbb{F} \subset \mathrm{Fix}(\Theta_Q)$. This is obvious if $\mathbb{F}$ is trivial. If $\mathbb{F}$ has rank one then $\mathbb{F} = \mathrm{Fix}(\Theta_Q)$ and

$$Q' = \Phi(Q') = \Theta_Q^L(Q') = (i_a\Theta_{Q'})^L(Q') = i_a^L\Theta_{Q'}^L(Q') = i_a^L(Q')$$

which implies that a must be trivial and we are done. If $\mathbb{F}$ has rank ≥ 2 then there is a unique Θ such that $\mathbb{F} \subset \mathrm{Fix}(\Theta)$ and Φ is the only automorphism representing ϕ such that $\mathbb{F} \subset \mathrm{Fix}(\Phi)$. Thus $\Theta^L = \Phi$. There exists $b \in \mathbb{F}$ such that $\Theta_Q = i_b\Theta$. We have $\Phi = \Theta_Q^L = i_b^L\Theta^L = i_b^L\Phi$ so b is trivial and the proof is complete. $\square$

To apply Lemma 3.12 we will need a bound on the cardinality of

$$\big(\textstyle\bigcup_{\Phi\in\mathrm{P}(\phi)} \mathrm{Fix}_+(\partial\Phi)\big)/F_n.$$

[1] This follows from the standard fact that the kernel of the natural map $GL_n(\mathbb{Z}) \to GL_n(\mathbb{Z}/3\mathbb{Z})$ is torsion-free and the result of Baumslag-Taylor [BT68] that the kernel of the natural map $\mathtt{Out}(F_n) \to GL_n(\mathbb{Z})$ is torsion-free.

Lemma 3.13. *If $\phi \in \mathtt{Out}(F_n)$ is rotationless then*

$$\big(\bigcup_{\Phi \in \mathrm{P}(\phi)} \mathrm{Fix}_+(\partial\Phi)\big)/F_n \leq 15(n-1).$$

Proof. Choose a CT $f : G \to G$ representing ϕ and assume the notation of Definition 3.9. By Lemma 3.10 it suffices to show that the cardinality of the image of $E \mapsto [R_E]$ is bounded by $15(n-1)$. By construction, the initial vertex of E is principal. If it has valence at least three then we say that it is *natural*.

There are at most $6(n-1)$ oriented edges based at natural vertices. Some of these are not fixed and so do not contribute to $\mathcal{E}$. For example, if E is NEG then (Lemma 4.21 of [FH11]) the terminal vertex of E is natural and the direction determined by $\bar{E}$ is not fixed. Similarly, if (E_1, E_2) is an illegal turn of EG height then the basepoint for this turn is natural and either E_1 or E_2 determines a non-fixed direction. It follows that $6(n-1)$ is an upper bound for the sum of the number of edges in $\mathcal{E}$ that are based at natural vertices, the number of non-fixed NEG edges and the number of EG stratum H_r with an illegal turn of height r.

It remains to account for those EG edges $E \in \mathcal{E}$ that are based at valence two vertices v. By our previous estimate there are at most $6(n-1)$ such v with the other edge incident to v being non-fixed NEG. The only other possibility is that both edges incident to v are EG. If the edges were in different strata, say H_r and $H_{r'}$ with $r < r'$ then v would have valence ≥ 2 in G_r (because G_r is a core subgraph), a contradiction. Thus both edges belong to the same stratum H_r. Since v is a principal vertex, it must be an endpoint of a Nielsen path ρ of height r. There are at most four edges incident to valence two vertices at the endpoints of ρ and these determine at most three points in $\big(\bigcup_{\Phi \in \mathrm{P}(\phi)} \mathrm{Fix}_+(\partial\Phi)\big)/F_n$ because the two directions pointing into ρ determine the same point. There is at most one such ρ for each EG stratum H_r and ρ has an illegal turn of height r so our initial bound of $6(n-1)$ counted each ρ once; we now have to count it two more times. In passing from the highest core G_s with $s < r$ to G_r, at least two natural edges are added. It follows that the number of EG strata is $\leq \frac{3}{2}(n-1)$. The total count then is $6(n-1) + 6(n-1) + 3(n-1) = 15(n-1)$. □

Corollary 3.14. *Let $h(n) = |GL(\mathbb{Z}/3/Z, n)| = 3^{(n^2-1)}$, let $g(m)$ be Landau's function, the maximum order of an element in the symmetric group S_m, and let $K_n = g(15(n-1))! \cdot h(n)$. If $\theta \in \mathtt{Out}(F_n)$ then θ^{K_n} is rotationless.*

Proof. Lemma 3.13 implies that $\theta^{g(15(n-1))!}$ induces the trivial permutation on

$$\big(\bigcup_{\Phi \in \mathrm{P}(\phi)} \mathrm{Fix}_+(\partial\Phi)\big)/F_n$$

for any rotationless iterate ϕ of θ. Hence $\theta^{K_n} = (\theta^{g(15(n-1))!})^{h(n)}$ satisfies the hypotheses of Lemma 3.12. □

Remark 3.15. In Corollary 15.18, we will see that

$$\left|\left(\bigcup_{\Phi\in \mathsf{P}(\phi)} \mathrm{Fix}_+(\partial\Phi)\right)/F_n\right| \leq 6(n-1)$$

and so we could take $K_n = g(6(n-1))! \cdot h(n)$ in Corollary 3.14.

3.4. Algorithmic Proof of Theorem 2.12. We review the proof of the existence of $f: G \to G$ for θ as given in [FH11, Theorem 2.19], altering it slightly to make it algorithmic.

If $\mathcal{C}$ is not specified, take it to be the single free factor system $[F_n]$. Apply Corollary 2.10 to construct a relative train track map $f: G \to G$ and $\emptyset = G_0 \subset G_1 \subset \cdots \subset G_N = G$ such that for each $\mathcal{F}_i \in \mathcal{C}$ there exists G_j satisfying $\mathcal{F} = [G_j]$. The modifications necessary to arrange all the properties but (V) are explicitly described in the original proof. These steps come after (V) has been established in that proof but make no use of (V) so there is no harm in our switching the order in which properties are arranged. For notational simplicity we continue to refer to the relative train track map as $f: G \to G$ even though it has been modified to satisfy all the properties except possibly (V).

For K_n as in Corollary 3.14, θ^{K_n} is rotationless. Subdivide $f: G \to G$ at the (finite) set S of isolated points in $\mathrm{Fix}(f^{K_n})$ that are not already vertices; these occur only in EG edges E and are in one to one correspondence with the occurrences of E or $\bar{E}$ in the edge path $f_\#^{K_n}(E)$. We claim that property (V) is satisfied. If not, then perform a further finite ([FH11, Lemma 2.12]) subdivision so that (V) is satisfied. [FH11, Proposition 3.29] and [FH11, Lemma 3.28] imply that every periodic Nielsen path of $f: G \to G$ (after the further subdivision) has period at most K_n. But then S contains the endpoints of all indivisible periodic Nielsen paths after all and no further subdivision was necessary. Since subdivision at S is algorithmic, we are done.

4. Reducibility

Given a relative train track map $f: G \to G$ and filtration $\emptyset = G_0 \subset G_1 \subset \cdots \subset G_N = G$ representing $\phi \in \mathsf{Out}(F_n)$, let $\emptyset = \mathcal{F}_0 \sqsubset \mathcal{F}_1 \sqsubset \cdots \sqsubset \mathcal{F}_K$ be the increasing sequence of distinct ϕ-invariant free factor systems that are realized by the G_i's. Assuming that $f: G \to G$ satisfies property (F) of Theorem 2.12, $\mathcal{F}_i$ is realized by a unique core filtration element for each $i \geq 1$ and $\mathcal{F}_0$ is realized by G_0. If $\mathcal{F}$ is a free factor system that is invariant by some iterate ϕ^k of ϕ and that is properly contained between $\mathcal{F}_i$ and $\mathcal{F}_{i+1}$ then we say that $\mathcal{F}$ is a *reduction for* $\mathcal{F}_i \sqsubset \mathcal{F}_{i+1}$ *with respect to* ϕ; if there is no such $\mathcal{F}$ then $\mathcal{F}_i \sqsubset \mathcal{F}_{i+1}$ is *reduced with respect to* ϕ. If each $\mathcal{F}_i \sqsubset \mathcal{F}_{i+1}$ is reduced with respect to ϕ then we say that *$f: G \to G$ is reduced*.

We assume for the rest of the section that ϕ is rotationless. In particular, a free factor system that is invariant by some iterate of ϕ is ϕ-invariant [FH11, Lemma 3.30].

The main results of this section are Proposition 4.9 and Lemma 4.12. The former, which assumes that $f : G \to G$ satisfies the conclusions of Theorem 2.12, provides an algorithm in the EG case for deciding if $\mathcal{F}_i \sqsubset \mathcal{F}_{i+1}$ is reduced and for finding a reduction if there is one. The latter has stronger requirements and easily leads to an algorithm that handles the NEG case. We save the final details of that algorithm for Section 8.

4.1. The EG case. Recall ([BFH00, Section 2] or [HM, Part I Fact 1.3]) that a pair of free factor systems $\mathcal{F}^1$ and $\mathcal{F}^2$ has a well-defined *meet* $\mathcal{F}^1 \wedge \mathcal{F}^2$ characterized by $[A] \in \mathcal{F}^1 \wedge \mathcal{F}^2$ if and only if there exist subgroups A^1, A^2 such that $[A^i] \in \mathcal{F}^i$ and $A^1 \cap A^2 = A$.

Let $\mathcal{B}$ be the basis of F_n corresponding to the edges of R_n (see Section 2.1). If A is a finitely generated subgroup of F_n then the *Stallings graph* R_A of the conjugacy class $[A]$ of A is the core of the cover of R_n corresponding to A. There is an immersion $R_A \to R_n$ and if we subdivide R_A at the pre-image of the vertex of R_n then we view the edges of R_A as *labeled* by their image edges in R_n and hence by elements of $\mathcal{B}$. The *complexity* of A is the number of edges in (subdivided) R_A. Stallings graphs are generalized in Section 9.2 and more discussion can be found there.

Lemma 4.1. *Given free factor systems $\mathcal{F}^1$ and $\mathcal{F}^2$ one can algorithmically construct $\mathcal{F}^1 \wedge \mathcal{F}^2$.*

Proof. We may assume without loss that $\mathcal{F}^1 = \{[A]\}$ and $\mathcal{F}^2 = \{[B]\}$ for given subgroups A, B. According to Stallings [Sta83, Theorem 5.5 and Section 5.7(b)], the conjugacy classes of the intersections of A with conjugates of B are all represented by components of the pullback of the diagram $R_A \to R \leftarrow R_B$. $\square$

Lemma 4.2. *Given a finite set $\{a_i\}$ of elements of F_n and a finite set $\{A_j\}$ of finitely generated subgroups of F_n there is an algorithm that finds the unique minimal free factor system that carries each $[a_i]$ and each conjugacy class carried by some $[A_j]$.*

Proof. By replacing $[a_i]$ by $[\langle a_i \rangle]$, we may assume that finite set $\{a_i\}$ is empty. The *complexity* of $\mathcal{A} = \{[A_j]\}$ is the sum of the complexities of the $[A_j]$. By Gersten [Ger84], there is an algorithm to find $\Theta \in \mathsf{Aut}(F_n)$ so that $\Theta(\mathcal{A}) = \{[\Theta(A_j)]\}$ has minimal complexity in the orbit of $\mathcal{A}$ under the action of $\mathsf{Aut}(F_n)$. Let $\mathcal{P}$ be the finest partition of $\mathcal{B}$ such that the labels of each Stallings graph $R_{\Theta(A_j)}$ are contained in some element of $\mathcal{P}$. The free factor system $\mathcal{F}(\mathcal{P})$ determined by $\mathcal{P}$ is the minimal free factor system carrying $\Theta(\mathcal{A})$; see [DF05, Lemma 9.19]. Hence $\Theta^{-1}(\mathcal{F}(\mathcal{P}))$ is the minimal free factor system carrying $\mathcal{A}$. $\square$

Corollary 4.3. *Suppose that $\phi \in \mathtt{Out}(F_n)$, that $\mathcal{F}^1$ is a proper free factor system and that $\mathcal{F}^0 \sqsubset \mathcal{F}^1$ is a (possibly trivial) ϕ-invariant free factor system. Then there is an algorithm that decides if there is a ϕ-invariant free factor system $\mathcal{F} \sqsubset \mathcal{F}^1$ that properly contains $\mathcal{F}^0$ and that finds such an $\mathcal{F}$ if one exists.*

Proof. First check if $\mathcal{F}^1$ is ϕ-invariant or if $\mathcal{F}^1 = \mathcal{F}^0$. If the latter is true the output of the algorithm is NO. If the latter is false and the former is true the output is YES. If neither is true apply Lemma 4.1 to compute $\mathcal{F}^1 \wedge \phi(\mathcal{F}^1)$ which is properly contained in $\mathcal{F}^1$, contains $\mathcal{F}^0$ and contains every ϕ-invariant free factor system that is contained in $\mathcal{F}^1$. Repeat these steps with $\mathcal{F}^1 \wedge \phi(\mathcal{F}^1)$ replacing $\mathcal{F}^1$. Since there is a uniform bound to the length of a strictly decreasing sequence of free factor systems [HM, Part I Fact 1.3] the process stops after finitely many steps. □

The following lemma is used in Step 1 of the proof of Proposition 4.9. Recall from [FH11, Remark 3.20] that if $f: G \to G$ represents a rotationless ϕ and satisfies the conclusions of Theorem 2.12, then for each $\Lambda \in \mathcal{L}(\phi)$ there is an EG stratum H_r such that Λ has height r and this defines a bijection between $\mathcal{L}(\phi)$ and the set of EG strata. The definition of a path being weakly attracted to Λ^+ appears as [BFH00, Definition 4.2.3].

Lemma 4.4. *Suppose that $f: G \to G$ and $\emptyset = G_0 \subset G_1 \subset \cdots \subset G_N = G$ are a relative train track map and filtration representing a rotationless ϕ and satisfying the conclusions of Theorem 2.12 and that H_r is an EG stratum with associated attracting lamination Λ^+. Then there is a computable constant C such that if $\sigma_0 \subset G_r$ is an r-legal path that crosses at least C edges in H_r then every path $\sigma \subset G_r$ that contains σ_0 as a subpath is weakly attracted to Λ^+.*

Proof. Choose l so that the f^l image of each edge in H_r crosses at least two edges in H_r. It is shown in the proof of [BFH00, Lemma 4.2.2] (see also [BFH00, Corollary 4.2.4]) that $C = 4lC_0 + 1$ satisfies the conclusions of our lemma for any constant C_0 that is greater than or equal to the bounded cancellation constant for f. Since the latter can be computed [BFH97, Lemma 3.1] using only the transition matrix for f, we are done. □

Lemma 4.5. *Suppose that $f: G \to G$ and $\emptyset = G_0 \subset G_1 \subset \cdots \subset G_N = G$ are a relative train track map and filtration representing a rotationless ϕ and satisfying the conclusions of Theorem 2.12. Suppose further that H_r is an EG stratum. Then every indivisible periodic Nielsen path with height r has period one and there is an algorithm that finds them all.*

Remark 4.6. For a more efficient method than the one described in the proof see [HM11, Section 3.4].

Proof. Suppose that ρ is an indivisible periodic Nielsen path of height r. Proposition 3.29 and Lemma 3.28 of [FH11] imply that ρ has period 1 and property (V) of Theorem 2.12 implies that the endpoints of σ are vertices. By [BH92, Lemma 5.11], ρ decomposes as a concatenation $\rho = \alpha\beta^{-1}$ of r-legal edge paths whose initial and terminal edges are in H_r. Let α_0 and β_0 be the initial edges of α and β respectively. By Lemma 4.4 we can bound the number of H_r edges crossed by α and β by some positive constant C. Since H_r is an EG stratum we can choose k so that $f_\#^k(E)$ crosses more than C edges in H_r for each edge E in H_r. Since $\rho = f_\#^k(\rho)$ is obtained from $f_\#^k(\alpha) f_\#^k(\beta^{-1})$ by canceling edges at the juncture point and since no edges in H_r are cancelled when $f^k(\alpha)$ and $f^k(\beta)$ are tightened to $f_\#^k(\alpha)$ and $f_\#^k(\beta)$, $\alpha \subset f_\#^k(\alpha_0)$ and $\beta \subset f_\#^k(\beta_0)$. In particular, we can compute an upper bound for the number of edges crossed by ρ, reducing us to testing a finite set of paths to decide which are indivisible Nielsen paths. $\square$

A subgroup system $\mathcal{A}$ is a *vertex group system* if there exists a real F_n-tree with trivial arc stabilizers such that $\mathcal{A}$ is the set of non-trivial vertex stabilizers.

Lemma 4.7. *Suppose that $f: G \to G$ and $\emptyset = G_0 \subset G_1 \subset \cdots \subset G_N = G$ are a relative train track map and filtration representing a rotationless ϕ and satisfying the conclusions of Theorem 2.12. Suppose also that H_N is an EG stratum with attracting lamination Λ^+, and that $[G_u] \cup \Lambda^+$ fills where $G_u = G \setminus H_N^z$. Then the following are satisfied.*

(1) *There is a unique vertex group system $\mathcal{A}$ such that a conjugacy class $[a]$ is not weakly attracted to Λ^+ if and only if $[a]$ is carried by an element of $\mathcal{A}$.*

(2) *A circuit $\sigma \subset G$ represents an element of $\mathcal{A}$ if and only if σ splits as a concatenation of subpaths each of which is either contained in G_u or is an indivisible Nielsen path of height N.*

(3) *There is a proper free factor system $\mathcal{F}' \sqsupset [G_u]$ such that one of the following holds*:

   (a) $\mathcal{A} = \mathcal{F}'$;

   (b) *$\mathcal{A}$ fills and $\mathcal{A} = \mathcal{F}' \cup \{[A]\}$ where $[A]$ has rank one.*

   Moreover, $[G_u] \sqsubset [G_N]$ is reduced with respect to ϕ if and only if $\mathcal{F}' = [G_u]$.

Remark 4.8. G_u is not necessarily a core subgraph. It deformation retracts to a core subgraph G_s and is obtained from G_s by adding NEG edges with terminal endpoints in G_s. We use G_u in this lemma rather than G_s because indivisible Nielsen paths of height N can have endpoints at the valence one vertices of G_u.

Proof. The existence of a vertex group system $\mathcal{A}$ as in (1) follows from [BFH00, Theorem 6.1] and [HM, Part III Proposition 1.4(1)]. Uniqueness of $\mathcal{A}$ follows from the fact [HM, Part I Lemma 3.1] that a vertex group system is determined by the conjugacy classes that it carries. For the rest of this proof we take (1) to be the defining property of $\mathcal{A}$. Note that $\mathcal{A}$ depends only on Λ^+ and ϕ and not on the choice of $f: G \to G$. In particular, $\mathcal{A}$ is ϕ-invariant.

If a circuit $\sigma \subset G$ splits into subpaths that are either contained in G_u or are Nielsen paths of height N then the number of H_N edges in $f^k_\#(\sigma)$ is independent of k and σ is not weakly attracted to Λ^+. This proves the if direction of (2).

The only if direction of (2) is more work. [BFH00, Lemmas 4.2.6 and 2.5.1] and Lemma 4.5 imply that there exists $k \geq 1$ such that $f^k_\#(\sigma)$ splits into subpaths that are either contained in G_{N-1}, are indivisible Nielsen paths of height N or are edges of height N. Assuming that σ, and hence $f^k_\#(\sigma)$, is not weakly attracted to Λ^+, [BFH00, Corollary 4.2.4] implies that no term in this splitting is an edge of height N. If $f^k_\#(\sigma) \subset G_u$ then $\sigma \subset G_u$ and we are done. If $f^k_\#(\sigma)$ is a closed Nielsen path of height N then σ and $f^k_\#(\sigma)$ have the same $f^k_\#$-image and so are equal. In particular, σ is a Nielsen path of height N. In the remaining case, there is a splitting

$$f^k_\#(\sigma) = \mu_1 \cdot \nu_1 \cdot \mu_2 \cdot \nu_2 \cdot \dots \cdot \mu_m \cdot \nu_m$$

into subpaths $\mu_i \subset G_u$ and Nielsen paths ν_i of height N. Since the endpoints of each μ_i are fixed by f and since the restriction of f to each f-invariant component of G_u is a homotopy equivalence, there exist paths $\mu'_i \subset G_u$ with the same endpoints as μ_i such that $f^k_\#(\mu'_i) = \mu_i$. Letting

$$\sigma' = \mu'_1 \cdot \nu_1 \cdot \mu'_2 \cdot \nu_2 \cdot \dots \cdot \mu'_m \cdot \nu_m$$

we have $f^k_\#(\sigma') = f^k_\#(\sigma)$ and hence $\sigma' = \sigma$. In particular, σ splits into subpaths of G_u and indivisible Nielsen paths of height N. This completes the proof of (2).

The main statement of (3) follows from [BFH00, Proposition 6.0.1 and Remark 6.0.2] (which applies because there is a CT representing ϕ in which Λ^+ corresponds to the highest stratum and $[G_u]$ is realized by a core filtration element). Since ϕ preserves $\mathcal{A}$, it acts periodically on the components of $\mathcal{A}$. [FH11, Lemma 3.30] therefore implies that ϕ preserves each rank one component of $\mathcal{A}$ and so also preserves $\mathcal{F}'$. If $\mathcal{F}' \neq [G_u]$ then $\mathcal{F}'$ is a reduction for $[G_u] \sqsubset [G_N]$. This proves the only if direction of the moreover statement.

Suppose then that $\mathcal{F}' = [G_u]$. If either (a) or (b) holds then any free factor system $\mathcal{F}$ that properly contains $[G_u]$ carries a conjugacy class not carried by $\mathcal{A}$. Item (1) implies that $\mathcal{F}$ carries a conjugacy classs that is weakly attracted to Λ^+. If $\mathcal{F}$ is ϕ-invariant then $\mathcal{F}$ carries Λ^+ in addition to containing G_u and so is improper. This completes the proof of the if direction of the moreover statement. $\square$

The next proposition shows how to reduce a relative train track map satisfying the conclusions of Theorem 2.12. One way to create an unreduced example is to identify a pair of distinct fixed points in a stratum H_r of a CT where H_r is both highest and EG.

Proposition 4.9. *Suppose that $f: G \to G$ and $\emptyset = G_0 \subset G_1 \subset \cdots \subset G_N = G$ are a relative train track map and filtration representing a rotationless ϕ and satisfying the conclusions of Theorem 2.12. Suppose also that H_r is an EG stratum and that G_s is the highest core filtration element below G_r. Then there is an algorithm to decide if $[G_s] \sqsubset [G_r]$ is reduced and if it is not to find a reduction.*

Proof. By [FH11, Lemma 3.30], the non-contractible components of G_r are f-invariant; in particular H_r is contained in a single component of G_r. By restricting to this component we may assume that H_r is the top stratum and hence that $r = N$.

Let $\Lambda^+ \in \mathcal{L}(\phi)$ be the lamination associated to H_N. By [BFH00, Lemma 3.2.4] there exists $\Lambda^- \in \mathcal{L}(\phi^{-1})$ such that the smallest free factor system that carries Λ^+ is the same as the smallest free factor system that carries Λ^- and we denote this by $\mathcal{F}_\Lambda$. It follows that the realizations of Λ^+ and Λ^- in any marked graph cross the same set of edges in that graph. It also follows that the smallest free factor system that carries $[G_s]$ and Λ^+ is the same as the smallest free factor system that carries $[G_s]$ and Λ^- and we denote this by $\mathcal{F}_{s,\Lambda}$. Since both $[G_s]$ and $\Lambda^{\pm}$ are ϕ-invariant, [BFH00, Corollary 2.6.5] implies that $\mathcal{F}_{s,\Lambda}$ is ϕ-invariant.

Choose constants as follows.

- $C_E = 6(n-1)$ is the maximal number of oriented natural edges in a marked core graph of rank n.
- $M = 2n$; if $\mathcal{F}_0 \sqsubset \mathcal{F}_1 \sqsubset \cdots \sqsubset \mathcal{F}_P$ is an increasing nested sequence of free factor systems of F_n then $P \le M - 1$. (This follows by induction on the rank n and the observation that if $\mathcal{F}_P = \{F_n\}$ with $n > 1$ then there exists $\mathcal{F}$ such that $\mathcal{F}_{P-1} \sqsubset \mathcal{F} \sqsubset \mathcal{F}_P$ and $\mathcal{F}$ consists of a pair of conjugacy classes whose ranks add to n.)
- $C_0 = C + 2$ where C satisfies the conclusion of Lemma 4.4.

Step 1: an existence result when $\mathcal{F}_{s,\Lambda}$ is proper. Consider the set $\mathcal{K}$ of marked graphs K containing a (possibly empty) core subgraph K_0 and equipped with a marking preserving homotopy equivalence $p: K \to G$ taking vertices to vertices such that

(a) $p \mid K_0: K_0 \to G_s$ is a homeomorphism (If $s = 0$ then $G_s = K_0 = \emptyset$);
(b) the restriction of p to each natural edge is either an immersion or constant;
(c) there is at most one natural edge on which p is constant.

Each natural edge E of $K \in \mathcal{K}$ is *labeled* by $p(E)$, thought of as a (possibly trivial) edge path in G. We do not distinguish between two elements of $\mathcal{K}$ if there is a label preserving homeomorphism between them. The length $|E|$ of a natural edge E is the number of edges in $p(E)$ and the total length $|K|$ of K is the sum of the length of its natural edges. The number of H_N edges in $p(E)$ is denoted $|E|_N$. Let Λ^+_K and Λ^-_K be the realizations of Λ^+ and Λ^- in K. As observed above, the set of edges crossed by Λ^+ is the same as the set of edges crossed by Λ^-. We denote this common core subgraph by K_Λ.

We claim that if $\mathcal{F}_{s,\Lambda}$ is proper then there exists an element $K \in \mathcal{K}$ with the following properties:

(1) $|E|_N \leq C_0$ for all natural edges E of K;

(2) the restriction of p to a leaf of Λ^+_K is an immersion;

(3) $K_0 \cup K_\Lambda$ is a proper core subgraph.

Note that $K_0 \cup K_\Lambda$ is a core subgraph because it is a union of core subgraphs so the content of (3) is that $K_0 \cup K_\Lambda$ is proper. Our proof of the claim makes use of an idea from the proof of Proposition 3.4 in Part IV of [HM].

Assuming that $\mathcal{F}_{s,\Lambda}$ is proper, there exists a marked graph K with proper core subgraphs $K_0 \subset K_1 \subset K$ and a marking preserving homotopy equivalence $p: K \to G$ satisfying (a) and (b) and $[K_1] = \mathcal{F}_{s,\Lambda}$. The last property implies that $K_1 = K_0 \cup K_\Lambda$ so (3) is satisfied. If there is at least one natural edge in $K \setminus K_1$ on which p is an immersion then collapse each natural edge in K on which p is constant to a point. Otherwise, collapse each natural edge in K_1 on which p is constant and all but one natural edge in $K \setminus K_1$ on which p is constant to a point. The resulting marked graph, which we continue to denote K, still satisfies (a) and (b) because $p|K_0$ is injective and so K_0 is unaffected by the collapsing and (c) is now satisfied so $K \in \mathcal{K}$. Replacing K_1 by its image under the collapse, it is still true that $K_1 = K_0 \cup K_\Lambda$ is a proper core subgraph and now the restriction of p to each natural edge of K_1 is an immersion. (It may now be that $[K_1]$ properly contains $\mathcal{F}_{s,\Lambda}$.)

If $p \bigm| K_1$ is not an immersion then there is a pair of natural edges E_1, E_2 in K_1 with the same initial vertex and such that the edge paths $p(E_1)$ and $p(E_2)$ have the same first edge, say e. Folding the initial segments of E_1 and E_2 that map to e produces an element $K' \in \mathcal{K}$ with subgraph K'_0 satisfying (a) and such that $|K'| < |K|$. Note that (3) is still satisfied because $K'_0 \cup K'_\Lambda$ is contained in the image $K'_1 \subset K'$ and $[K_1] = [K'_1]$. Replacing K with K' and repeating this finitely many times, we may assume that $p \bigm| K_1$ is an immersion and hence that p restricts to an immersion on leaves of Λ^+_K and Λ^-_K. In particular, (2) is satisfied. If (1) is satisfied then the proof of the claim is complete.

Suppose then that there is a natural edge E of K such that $|E|_r > C_0 = C + 2$. By [BFH00, Theorem 6.0.1], a leaf of Λ^- is not weakly attracted to a generic leaf of Λ^+. Lemma 4.4 therefore implies that at least one of the two laminations Λ^+_K and Λ^-_K does not cross E. It follows that $\mathcal{F}_\Lambda \sqsubset [K \setminus E]$ and hence that E is contained in the complement of K_Λ. Since E is obviously in the complement of K_0, we have that E is contained in the complement of $K_0 \cup K_\Lambda$. If there is a natural edge on which p is constant, it must be in the complement of $K_0 \cup K_\Lambda \cup E$ and we collapse it to a point. As above, the resulting marked graph is still in $\mathcal{K}$ and (2) and (3) are still satisfied. We may now assume that p is an immersion on each natural edge of K. The map $p: K \to G$ is not a homeomorphism so there is at least one pair of edges that can be folded. Perform the fold and carry the $K_0 \subset K$ notation to the new marked graph. It is still true that $K \in \mathcal{K}$ and that (2) holds. Folding reduces $|E|_r$ by at most 2 so there is still an edge with $|E|_r > C$. Arguing as above we see that E is contained in the complement of $K_0 \cup K_\Lambda$ so (3) is still satisfied. If (1) is satisfied then the proof of the claim is complete. Otherwise perform another fold. Conditions (2) and (3) are satisfied so check condition (1) again. Folding reduces $|K|$ so after finitely many folds, (1) is satisfied and the claim is proved.

Step 2: Part 1 of the algorithm. In this step we present an algorithm that either finds a reduction for $[G_s] \sqsubset [G_r]$ or concludes that $\mathcal{F}_{s,\Lambda}$ is improper, i.e. $\mathcal{F}_{s,\Lambda} = \{[F_n]\}$. In the former case we are done. In the latter case we move on to the second part of the algorithm in which we either find a reduction for $[G_s] \sqsubset [G_r]$ or we conclude that $[G_s] \sqsubset [G_r]$ is irreducible.

(A1) Choose a generic leaf $\gamma \subset G$ of Λ^+. One way to do this is to choose an edge e of H_r and $k > 0$ so that at least one occurrence of e in the edge path $f_\#^k(e)$ is neither the first nor last H_r edge. Then $e \subset f_\#^k(e) \subset f_\#^{2k}(e) \subset \cdots$ and the union of these paths is a generic leaf of Λ^+ by [BFH00, Corollary 3.1.11 and Lemma 3.1.15].

(A2) Let $C_1 = 1$. Choose a subpath γ_1 of γ that crosses at least

$$(C_E + 1)C_0 + C_E(C_1 + C_0) + 2C_0$$

edges in H_r and let L_1 be the number of edges in γ_1. (The choice of these constants will be clarified in Step 3.)

(A3) Enumerate all core graphs J of rank $< n$ satisfying: there is a core subgraph $J_0 \subset J$ and a map $p\colon J \to G$ taking vertices to vertices such that the restriction of p to each natural edge is an immersion onto a path of length at most L_1 and such that $p \mid J_0\colon J_0 \to G_s$ is a homeomorphism. Label the natural edges of J by their p images. We do not distinguish between labeled graphs that differ by a label preserving homeomorphism so there are only finitely many J and we consider them one at a time. If $\sigma \subset J$ is a path with endpoints at natural vertices and if p restricts to an immersion on σ then we let $|\sigma|$ be the number of edges crossed by $p(\sigma)$ and $|\sigma|_r$ the number of H_r edges crossed by $p(\sigma)$.

Let $p_\#$ be the homomorphism induced by $p\colon J \to G$ on fundamental groups. By Lemma 4.2 we can decide if $p_\#$ is an isomorphism to a free factor system $[J]$ of F_n. If not, then move on to the next candidate. If yes, then apply Lemma 4.3 with $\mathcal{F}_0 = [J_0] = [G_s]$ and $\mathcal{F}_1 = [J]$. If this produces a ϕ-invariant $\mathcal{F} \sqsubset [J]$ that properly contains $[G_s]$ then we have found a reduction and the algorithm stops. Otherwise we know that $[J]$ does not contain such an $\mathcal{F}$. In particular $[J]$ does not contain $\mathcal{F}_{s,\Lambda}$ and so does not carry Λ^+. Choose a finite subpath $\gamma_{1,J} \subset \gamma \subset G$ that does not lift to J.

By [BFH00, Lemma 3.1.10(4)] there exists an edge E in H_r and $k \geq 1$ such that $\gamma_{1,J}$ is a subpath of $\tilde{f}^k_\#(E)$. By [BFH00, Lemma 3.1.8(3)] there is a computable k_0 so that for any edge E' in H_r and any $l \geq k+k_0$, $\tilde{f}^k_\#(E)$, and hence $\gamma_{1,J}$, is a subpath of $f^l_\#(E')$. Finally, by [BFH00, Lemma 3.1.10(3)] there exists computable $C_{2,J} > 0$ so that if σ is a subpath of γ that crosses $\geq C_{2,J}$ edges of H_r then σ contains some $f^l_\#(E')$, and hence contains $\gamma_{1,J}$, as a subpath and so does not lift into J. Now move on to the next candidate.

At the end of the process we have either found a reduction and the algorithm stops or we have found a constant $C_2 = \max\{C_{2,J}\}$ so that if σ is a subpath of γ that crosses at least C_2 edges of H_r then σ does not lift into any J. Choose a subpath γ_2 of γ that crosses at least

$$(C_E + 1)C_0 + C_E(C_2 + C_0) + 2C_0$$

edges of H_r and let L_2 be the number of edges crossed by γ_2.

(A4) Repeat (A3) replacing L_1 with L_2. At the end of the process we have either found a reduction and the algorithm stops or we have found a constant C_3 so that if σ is a subpath of γ that crosses at least C_3 edges of H_r then σ does not lift into any J. Choose a subpath γ_3 of γ that crosses at least

$$(C_E + 1)C_0 + C_E(C_3 + C_0) + 2C_0$$

edges of H_r and let L_3 be the number of edges crossed by γ_3.

(A5) Iterate this process up to M times. If after M iterations the algorithm has not found a reduction and stopped, then stop and conclude that $\mathcal{F}_{s,\Lambda} = \{[F_n]\}$.

Step 3: justifying Part 1 of the algorithm. In this step we verify that if $\mathcal{F}_{s,\Lambda}$ is proper then the above algorithm finds a reduction in at most M steps.

Suppose $K \in \mathcal{K}$ satisfies (1)–(3) and let $K(L_1)$ be the subgraph of K consisting of natural edges E with $|E| \leq L_1$. Lift the path γ_1 into K_Λ. After removing initial and terminal subpaths contained in single natural edges of K_Λ, we have a natural (in K) edge path $\gamma_{1,K} \subset \Lambda^+_K \subset K_\Lambda$ that projects onto all of γ_1 except perhaps initial and terminal segments that cross at most C_0 edges of H_r. Thus $\gamma_{1,K} \subset K(L_1)$ and

$$|\gamma_{1,K}|_r \geq (C_E + 1)C_0 + C_E(C_1 + C_0)$$

Combining this inequality with (1), we see that $\gamma_{1,K}$ has a subpath that decomposes as

$$\alpha_1\beta_1\alpha_2 \ldots \beta_{C_E}\alpha_{C_E+1}$$

where each α_i is a single natural edge and each β_i is a natural subpath whose image under p crosses at least C_1 edges in H_r. Since there are at most C_E oriented natural edges in K, it follows that $\gamma_{1,K}$ contains a natural subpath that begins and ends with the same oriented natural edge and whose p-image crosses at least $C_1 = 1$ edges in H_r. This proves that that there is a circuit in $K(L_1)$ that is not in K_0 and hence that $[K(L_1)]$ properly contains $[K_0]$. By (3), $K_0 \cup K_\Lambda$ is proper and so if $K_0 \cup K_\Lambda \subset K(L_1)$ then $K_0 \cup K_\Lambda$ occurs as a J in the first iteration of the process and Lemma 4.3 finds a reduction in the first iteration of (A3) because $\mathcal{F}_{s,\Lambda} \sqsubset [K_0 \cup K_\Lambda]$.

If no reduction is found in the first iteration then proceed to the second iteration as described in (A4). By the same reasoning, there is a natural edge path $\gamma_{2,K} \subset K(L_2)$ such that

$$|\gamma_{2,K}|_r \geq (C_E + 1)C_0 + C_E(C_2 + C_0)$$

and there is a subpath of $\gamma_{1,K}$ that decomposes as

$$\alpha_1 \beta_1 \alpha_2 \ldots \beta_{C_E} \alpha_{C_E+1}$$

where each α_i is a single natural edge and each β_i is a natural subpath whose image under p crosses at least C_2 edges in H_r. It follows that there is a circuit in $K(L_2)$ that is not in $K(L_1)$ and hence that $[K(L_2)]$ properly contains $[K(L_1)]$. If $K_0 \cup K_\Lambda \subset K(L_2)$ then $K_0 \cup K_\Lambda$ occurs as a J in the second iteration of the process and our algorithm finds a reduction.

Continuing on, the iteration either produces a reduction within M steps or produces a properly nested sequence

$$[K_0] \sqsubset [K(L_1)] \sqsubset [K(L_2)] \sqsubset \cdots \sqsubset [K(L_M)]$$

of free factors. Since the latter contradicts the definition of M, a reduction must have been found.

Step 4: Part 2 of the algorithm. In this part of the algorithm we assume that $\mathcal{F}_{s,\Lambda}$ is improper. The filtration element $G_u = G_N \setminus H_N^z$ deformation retracts to G_s, see Remark 4.8. In particular, $[G_u] = [G_s]$.

Given a circuit $\sigma \subset G$ [resp. a subgroup A] let $\mathcal{F}_{u,\sigma}$ [resp. $\mathcal{F}_{u,A}$] be the smallest free factor system that carries $[\sigma]$ [resp. every conjugacy class in $[A]$] and every conjugacy class in $[G_u]$. By Lemma 4.2, $\mathcal{F}_{u,\sigma}$ and $\mathcal{F}_{u,A}$ can be algorithmically determined.

By Lemma 4.5, the set P of indivisible periodic Nielsen paths with height r is finite, can be determined algorithmically and each element of P has period one. For notational convenience we assume that P is closed under orientation reversal. Let Σ be the set of circuits in G that split into a concatenation of paths in G_u and elements of P. Lemma 4.7 implies that Σ is the set of circuits that are not weakly attracted to Λ^+ and that the conjugacy classes determined by Σ are exactly those carried by the subgroup system $\mathcal{A}$ of that lemma.

We consider several cases, each of which can be checked by inspection of G_u and the elements of P. If every circuit in Σ is contained in G_u then $\mathcal{A} = [G_u]$ and H_N is reduced by Lemma 4.7(3).

As a second case, suppose that there is a non-trivial path $\mu \subset G$ with endpoints $x, y \in G_u$ such that μ is homotopic rel endpoints to a concatenation of elements of P and so is a Nielsen path of height N. Let $[B]$ be the conjugacy class of the subgroup of F_n represented by closed paths based at x that decompose as a concatenation of subpaths, each of which is either μ, μ^{-1} or a path in G_u with endpoints in $\{x, y\}$. Then $[B]$ is ϕ-invariant, has rank ≥ 2, and each conjugacy class in $[B]$ is represented by a circuit $\sigma \in \Sigma$. By Lemma 4.7(3) there is a proper free factor system that carries $[G_u]$ and the conjugacy class of each element of B. It follows that $\mathcal{F}_{u,B}$ is proper. Since both $[G_u]$ and $[B]$ are ϕ-invariant, $\mathcal{F}_{u,B}$ is ϕ-invariant by [BFH00, Corollary 2.6.5] and we have found a reduction of H_N.

The final case is that there are no paths $\mu \subset G$ as in the second case and there is at least one element $\sigma \in \Sigma$ that is not contained in G_u. Each such σ is homotopic to a concatenation of elements of P. In particular, the conjugacy class determined by σ is ϕ-invariant and so $\mathcal{F}_{u,\sigma}$ is ϕ-invariant. Choose one such σ and check if $\mathcal{F}_{u,\sigma}$ is proper. If it is then we have found a reduction of H_N and we are done so suppose that it is not. Lemma 4.7(3) implies that $[\sigma]$ is carried by a rank one component $[A]$ of $\mathcal{A}$ and that $\mathcal{A} = \mathcal{F}' \cup [A]$ for some ϕ-invariant free factor system $\mathcal{F}'$. If there exists an element $\sigma' \in \Sigma$ that is not carried by $[G_u]$ and such that σ and σ' are not multiples of the same root-free circuit then $[\sigma']$ is not carried by $[A]$ so $\mathcal{F}_{u,\sigma'}$ is a reduction. Otherwise, $\mathcal{F}' = [G_u]$ and H_r is reduced. This completes the second step in the algorithm and so also the proof of the proposition. □

4.2. The NEG case. We now consider reducibility for NEG strata, beginning with a pair of examples.

Example 4.10. Suppose that $f\colon G \to G$ is a homotopy equivalence with filtration $\emptyset = G_0 \subset G_1 \subset \cdots \subset G_N = G$ representing $\phi \in \mathsf{Out}(F_n)$ and that $G_{r+2} = G_r \cup E_{r+1} \cup E_{r+2}$ where G_r is a connected core subgraph and where $H_{r+1} = E_{r+1}$ and $H_{r+2} = E_{r+2}$ are oriented edges with a common initial vertex not in G_r and a common terminal endpoint in G_r. Suppose also that $f(E_{r+1}) = E_{r+1}u$ and $f(E_{r+2}) = E_{r+2}u$ for some closed non-trivial path $u \subset G_r$ and that $f \mid G_r$ is a CT. Then the (NEG Nielsen paths) property of $f \mid G_{r+2}$ fails. The CT algorithm corrects this (see Section 8.1) by discovering that $E_{r+2}\overline{E}_{r+1}$ is a Nielsen path and then sliding the terminal endpoint of E_{r+2} along $\overline{E}_{r+1}$. In other words, E_{r+2} is replaced by a fixed loop E'_{r+2} based at the initial endpoint of E_{r+1}. Note that while establishing (NEG Nielsen paths) for $f \mid G_r$, we have discovered a reduction of $[G_r] \sqsubset [G_{r+2}]$. Namely, the ϕ-invariant free factor system $\{[G_r], [E'_{r+2}]\}$ is properly contained between $[G_{r+1}] = [G_r]$ and $[G_{r+2}]$.

Example 4.11. Suppose that $f: G \to G$ is a homotopy equivalence with filtration $\emptyset = G_0 \subset G_1 \subset \cdots \subset G_N = G$ representing $\phi \in \mathtt{Out}(F_n)$, that $G_{r+1} = G_r \cup E_{r+1}$ where G_r is a connected core subgraph such that $f \mid G_r$ is a CT and where $H_{r+1} = E_{r+1}$ is a fixed edge whose initial and terminal endpoints, x and y, belong to the same Nielsen class of $f \mid G_r$. Then $f \mid G_{r+1}$ satisfies all the properties of a CT except that $[G_r] \sqsubset [G_{r+1}]$ is not reduced. The CT algorithm corrects this in stages. First, a Nielsen path $\sigma \subset G_r$ connecting y to x is found. Then the terminal end of E_{r+1} is slid along σ so that its new terminal endpoint is x. Thus E_{r+1} is replaced by a fixed edge E'_{r+1} with both endpoints at x. Finally, x is blown up to a fixed edge E'_{r+2} with E'_{r+1} attached at the 'new' endpoint of E'_{r+2} and the remaining edges in the link of x still attached at x. Both $[G_r] \sqsubset [G_{r+1}]$ and $[G_{r+1}] \sqsubset [G_{r+2}]$ are reduced.

Lemma 4.12. *Suppose that $f: G \to G$ and $\emptyset = G_0 \subset G_1 \subset \cdots \subset G_N = G$ are a relative train track map and filtration satisfying the conclusions of Theorem 2.12 and representing a rotationless $\phi \in \mathtt{Out}(F_n)$. Suppose further that G_s is core, that $H_s = \{E_s\}$ is an NEG stratum, that C is the component of G_s that contains H_s and that $f \mid C$ satisfies (NEG Nielsen paths). If $[C \setminus E_s] \sqsubset [C]$ is reducible then E_s is fixed and its terminal endpoint is connected to its initial endpoint by a Nielsen path $\beta \subset C \setminus E_s$. In particular, $E_s\beta$ is a basis element and $\{[G_{s-1}], [E_s\beta]\}$ is a ϕ-invariant free factor system that is properly contained between $[G_{s-1}]$ and $[G_s]$.*

Proof. By restricting to C, we may assume that $G = C = G_s$ and hence that G_{s-1} has either a single component of rank $n-1$ or two components whose ranks add to n. If $[C \setminus E_s] \sqsubset [C]$ is reducible, there is a marked graph K with distinct proper core subgraphs $K_1 \subset K_2 \subset K$ such that $[K_1] = [G_{s-1}]$ and such that $[K_2]$ is ϕ-invariant. From

$$-\chi(K_1) \le -\chi(K_2) < -\chi(K) = -\chi(K_1) + 1$$

it follows that K_2 is obtained from K_1 by adding a disjoint loop α and that $K \setminus K_2$ is an edge E. In particular, K_1 is connected.

We claim that the circuit $\sigma \subset G$ representing $[\alpha]$ crosses E_s exactly once. The marked graph K' obtained from K by collapsing the components of a maximal forest is a rose with α as one of its edges. The marked graph G' obtained from G by collapsing the components of a maximal forest in G_{s-1} is a rose with E_s as one of its edges. Moreover $[K' \setminus \alpha] = [K_1] = [G_{s-1}] = [G' \setminus E_s]$. Let $h: K' \to G'$ be a homotopy equivalence that respects markings and that restricts to an immersion on each edge. Then $h(K' \setminus \alpha) = G' \setminus E_s$ and it suffices to show that $h(\alpha)$ crosses E_s exactly once. This follows from [BFH00, Corollary 3.2.2].

Having verified the claim, we can now complete the proof of the lemma. Since $[K_2]$ is ϕ-invariant and α is a component of K_2, α determines a ϕ-invariant conjugacy class. It follows that σ decomposes as a concatenation of indivisible Nielsen paths and fixed edges. Since σ crosses E_s exactly once, the (NEG Nielsen paths) property of $f: G \to G$ implies that E_s is a fixed edge. Thus σ decomposes as a circuit into $E_s\beta$ where β is a Nielsen path in G_{s-1}. □

5. Definition of a CT

The definition of a CT evolved over many years, growing out of improved relative trains [BFH00] which grew out of relative train tracks [BH92]. In addition to the nine items that make up the formal definition, there are numerous auxiliary consequences of these definitions that are used repeatedly. We have included the complete definition here for the reader's convenience but recommend consulting [FH11] (see also [HM, Part I Section 1.5]) for discussion and elaboration.

We have already discussed some of the technical terms in the definition of CT.

- For basics of relative train track theory, including the definitions of relative train track maps, filtrations, EG and NEG strata, linear edges and Nielsen paths see Section 2.1.
- For zero strata enveloped by EG strata see Notation 2.11.
- For principal vertices and rotationless $f: G \to G$ see Definition 3.3.
- For the definition of a filtration being reduced see the beginning of Section 4.

There are a few more terms that need defining.

Definition 5.1 (Definition 4.1 in [FH11]). If w is a closed root-free Nielsen path and E_i, E_j are linear edges satisfying $f(E_i) = E_i w^{d_i}$ and $f(E_j) = E_i w^{d_j}$ for distinct $d_i, d_j > 0$ then a path of the form $E_i w^* \bar{E}_j$ is called an *exceptional path*.

Definition 5.2 (Definition 4.3 in [FH11]). If H_r is EG and $\alpha \subset G_{r-1}$ is a non-trivial path with endpoints in $H_r \cap G_{r-1}$ then we say that α is a *connecting path for H_r*. If E is an edge in an irreducible stratum H_r and $k > 0$ then a maximal subpath σ of $f^k_\#(E)$ in a zero stratum H_i is said to be *r-taken* or just *taken* if r is irrelevant. A non-trivial path or circuit σ is *completely split* if it has a splitting, called a *complete splitting*, into subpaths, each of which is either a single edge in an irreducible stratum, an indivisible Nielsen path, an exceptional path or a connecting path in a zero stratum H_i that is both maximal (meaning that it is not contained in a larger subpath of σ in H_i) and taken. Note that the endpoints, if any, of a completely split path are at vertices.

Definition 5.3 (Definition 4.4 in [FH11]). A relative train track map is *completely split* if

(1) $f(E)$ is completely split for each edge E in each irreducible stratum;

(2) if σ is a taken connecting path in a zero stratum then $f_\#(\sigma)$ is completely split.

Proper extended folds, which are referred to in the (EG Nielsen paths) property of a CT are defined in [BFH00, Definition 5.3.2]. We will not review that here, in part because the actual definition is never used in this paper and in part because in applications one almost always refers to a consequence of this property (for example [FH11, Corollary 4.19]) rather than to the property itself.

Definition 5.4. A relative train track map $f: G \to G$ and filtration $\mathcal{F}$ given by $\emptyset = G_0 \subset G_1 \subset \cdots \subset G_N = G$ is said to be a CT (for completely split improved relative train track map) if it satisfies the following properties.

1. (*Rotationless*) $f: G \to G$ is rotationless.(Definition 3.3)
2. (*Completely split*) $f: G \to G$ is completely split.
3. (*Filtration*) $\mathcal{F}$ is reduced. (Section 4) The core of each filtration element is a filtration element.
4. (*Vertices*) The endpoints of all indivisible periodic (necessarily fixed) Nielsen paths are (necessarily principal) vertices. The terminal endpoint of each non-fixed NEG edge is principal (and hence fixed).
5. (*Periodic edges*) Each periodic edge is fixed and each endpoint of a fixed edge is principal. If the unique edge E_r in a fixed stratum H_r is not a loop then G_{r-1} is a core graph and both ends of E_r are contained in G_{r-1}.
6. (*Zero strata*) If H_i is a zero stratum, then H_i is enveloped by an EG stratum H_r, each edge in H_i is r-taken and each vertex in H_i is contained in H_r and has link contained in $H_i \cup H_r$.
7. (*Linear edges*) For each linear E_i there is a closed root-free Nielsen path w_i such that $f(E_i) = E_i w_i^{d_i}$ for some $d_i \neq 0$. If E_i and E_j are distinct linear edges with the same axes then $w_i = w_j$ and $d_i \neq d_j$.
8. *NEG* (*Nielsen paths*) If the highest edges in an indivisible Nielsen path σ belong to an NEG stratum then there is a linear edge E_i with w_i as in (Linear edges) and there exists $k \neq 0$ such that $\sigma = E_i w_i^k \bar{E}_i$.
9. *EG* (*Nielsen paths*) If H_r is EG and ρ is an indivisible Nielsen path of height r, then $f|G_r = \theta \circ f_{r-1} \circ f_r$ where
   (a) $f_r: G_r \to G^1$ is a composition of proper extended folds defined by iteratively folding ρ;
   (b) $f_{r-1}: G^1 \to G^2$ is a composition of folds involving edges in G_{r-1};
   (c) $\theta: G^2 \to G_r$ is a homeomorphism.

We include the following for future reference.

Lemma 5.5. *If $f: G \to G$ is a* CT *then $f^k: G \to G$ is a* CT *for all $k \geq 1$.*

Proof. By [FH11, Lemma 4.13], every periodic Nielsen path for f has period one. With this in hand, the first eight CT properties for f^k are easy to check. The remaining property (EG Nielsen paths) for f^k follows from [FH11, Corollary 4.33]. $\square$

The following example shows that the restriction of a CT f to a component of a core filtration element need not be a CT.

Example 5.6. We refer to Figure 1 for notation. The map $f: G \to G$ given by $e \mapsto e$, $b \mapsto be$, $c \mapsto c$, and $d \mapsto de^2$ is a CT with the filtration

$$\emptyset \subset \{e\} \subset \{e, b\} \subset \{e, b, c\} \subset G$$

and the restriction of f to each filtration element is a CT. If we however consider the new filtration

$$\emptyset \subset \{e\} \subset \{e, c\} \subset \{e, c, d\} \subset G$$

then f is still a CT with the new filtration but $f|\{e, c, d\}$ is not a CT because it does not satisfy (Vertices).

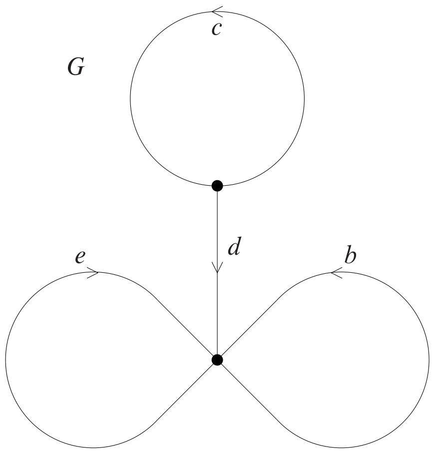

Figure 1. Example 5.6.

6. Sliding NEG edges

The key step for arranging that an NEG edge has good properties under iteration is to slide the terminal endpoint of the edge into an optimal position in the lower filtration element that contains it. This is carried out in [BFH00, Proposition 5.4.3]. In this section we make the algorithmic arguments needed to replace the non-algorithmic parts of the original proof.

6.1. Completely split rays. Recall from [FH11, Definition 4.4] that a splitting $\sigma = \sigma_1 \cdot \sigma_2 \ldots$ is a *complete splitting* if each σ_i is either a single edge in an irreducible stratum, an indivisible Nielsen path, an exceptional path or a maximal subpath in a zero stratum (with some additional features that we will not recall here.) A finite path or circuit has at most one complete splitting by [FH11, Lemma 4.11]. The first item of our next lemma states that the same is true for rays.

Recall from Section 2.1 that a ray σ with initial point a vertex is thought of as an edge path $\sigma = E_0 E_1 \ldots$. The initial and terminal endpoints, w_i and w_{i+1} of E_i are the *vertices of* σ. We view the set $\mathcal{W}$ of vertices of a path as being ordered by their subscripts. A decomposition of σ into subpaths is specified by a subset of $\mathcal{W}$; if w_i and w_j are consecutive elements of the subset then $E_i \ldots E_{j-1}$ is a term in the decomposition. If the decomposition is a splitting then we refer to these vertices as *splitting vertices*. A similar definition holds for circuits.

Lemma 6.1. *Suppose that* $f : G \to G$ *is a* CT, *that* $R \subset G$ *is a completely split ray and that* R_0 *is a subray of* R *that has a complete splitting.*

(1) *The complete splitting of* R *is unique.*

(2) *Let* v *be the first splitting vertex for* R *that is contained in* R_0 (*when the edge path* R_0 *is viewed as a subpath of the edge path* R). *Then each splitting vertex* w *for* R_0 *that comes after* v (*in the ordering of splitting vertices of* R_0) *is a splitting vertex for* R.

Proof. The second item implies the first by taking $R_0 = R$ so we need only prove the second. Let μ_0 be the term in the complete splitting of R_0 whose initial vertex is w. If μ_0 is either an indivisible Nielsen path or an exceptional path then the interior of μ_0 is an increasing union of pre-trivial paths by [FH11, Remark 4.2 and Lemma 2.11(2)] and so by [FH11, Lemma 4.11(2)] is contained in a single term μ of the complete splitting of R. Obviously μ is not a single edge and is not contained in a zero stratum so it must be either an indivisible Nielsen path or an exceptional path. Since v is the initial endpoint of some term in the complete splitting of R and w comes after v, it follows that v is not contained in the interior of μ and so $\mu \subset R_0$. The symmetric argument therefore applies to show that μ is contained in a term of the complete splitting of R_0 and hence that $\mu = \mu_0$ as desired. If μ_0 is either a single edge or is a maximal subpath in a zero stratum and μ_0 is not a term in the complete splitting of R then μ_0 is properly contained in a term μ of R that is an indivisible Nielsen path or an exceptional path. But this violates the hard splitting property [FH11, Lemma 4.11(2)] for the complete splitting of R_0 (applied to its finite completely split subpaths) and the fact that the interior of μ is the increasing union of pre-trivial paths. Thus μ_0 is a term in the complete splitting of R and we are done. $\square$

Definitions 6.2. Suppose that $f: G \to G$ is a CT, that $x \in G$, that $\sigma \subset G$ is a non-trivial completely split path connecting x to $f(x)$ and that the turn at $f(x)$ determined by $\bar{\sigma}$ and $f_\#(\sigma)$ is legal. The ray $R = \sigma \cdot f_\#(\sigma) \cdot f_\#^2(\sigma) \cdot \ldots$ satisfies $f_\#(R) \subset R$ and the given splitting of R has a refinement that is a complete splitting by [FH11, Lemma 4.11]; we say that R *is generated by* σ.

If E is a non-fixed edge of G whose initial direction is fixed then $f(E) = E \cdot \sigma$ for some σ as above. The ray $R_E = E \cdot \sigma \cdot f_\#(\sigma) \cdot f_\#^2(\sigma) \cdot \ldots$ is the *eigenray* determined by E. Note that we are not requiring that the initial vertex of E be principal (as we did in Section 3.3) or that E is non-linear and NEG so we are using the term eigenray a little more generally than is sometimes the case. We will need this inclusiveness in the proof of Lemma 6.4. For the same reason we assume that each isolated fixed point for f is a vertex.

Lemma 6.3. *Suppose that $f: G \to G$ is a* CT *and that σ and σ' are completely split non-Nielsen paths generating rays R and R' respectively. Then there is an algorithm to decide if the rays R and R' have a common terminal subray and if so to find initial subpaths $\tau \subset R$ and $\tau' \subset R'$ that terminate at splitting vertices of R and R' respectively and whose complementary terminal subrays are equal. Equivalently we find splitting vertices $v \in R$ and $v' \in R'$ such that terminal subrays of R and R' initiating at v and v' are equal* (*as edge paths*).

Proof. Let $\mathcal{V} = \{v_0, v_1, \ldots\}$ be the set of splitting vertices for R ordered so that v_{i-1} and v_i are the endpoints of the i^{th} term in the complete splitting. By construction, $f(\mathcal{V}) \subset \mathcal{V}$. For each $i \geq 0$, let $\sigma_i \subset R$ be the path connecting v_i to $f(v_i)$ and let $\ell_i = |\sigma_i|$ be the number of edges in σ_i; in particular, $\sigma = \sigma_0$. Note that σ_i generates the terminal subray of R that begins with v_i. Define V', σ'_j and ℓ'_j similarly using σ' and R' in place of σ and R. It is obvious that R and R' have a common terminal subray if $\sigma_i = \sigma'_j$ for some i and j. (Namely, the subrays of R and R' initiating at $v_i \in \mathcal{V}'$ and $v'_j \in \mathcal{V}'$ respectively.) The converse follows from Lemma 6.1. Our goal then is to either find i and j such that $\sigma_i = \sigma'_j$ or to conclude that no such i and j exist.

Let r [resp r'] be the maximal height of a term in the complete splitting of σ [resp. σ'] that is not a Nielsen path. Since $f(\sigma)$ and σ have a common endpoint, σ is not entirely contained in a zero stratum. Thus any term τ in the complete splitting of σ that is contained in a zero stratum is adjacent to a term that intersects the EG stratum that envelops τ (Notation 2.11). Since this adjacent edge has at least one non-fixed endpoint, it is neither an exceptional path nor a Nielsen path so must be a single edge in that EG stratum. We conclude that H_r is not a zero stratum. It follows that if μ is any height r term in the complete splitting of σ and if μ is not a Nielsen path then the length of $f_\#^i(\mu)$ goes to infinity with i; we say that μ is *growing*. Note that r is the maximal height of a growing term in the complete splitting of any σ_i and similarly for r' and σ'_j. Note also that for any given $L > 0$ one can find, by inspection, $M > 0$ such that $\ell_i, \ell'_j > L$ for all $i, j \geq M$.

The first step in the algorithm is to check if $r = r'$. If yes, then move on to step two. If not then there do not exist i and j such that $\sigma_i = \sigma'_j$ so the algorithm stops and outputs NO.

We may now assume that $r = r'$. If H_r is NEG define $K = 1$. Otherwise H_r is EG and we choose K so that for each edge E of H_r, $f_\#^K(E)$ contains at least C edges of H_r where C is the constant of Lemma 4.4. Now define I to be the number of terms in the complete splitting of $\sigma \cdot f_\#(\sigma) \cdot \dots \cdot f_\#^K(\sigma)$ or equivalently so that $i \le I$ if and only if $\tilde{v}_i \in \tilde{\sigma} \cdot \tilde{f}_\#(\tilde{\sigma}) \cdot \dots \cdot \tilde{f}_\#^K(\tilde{\sigma})$. Define J to be the number of terms in the complete splitting of $\sigma' \cdot f_\#(\sigma') \cdot \dots \cdot f_\#^K(\sigma')$ or equivalently so that $j \le J$ if and only if $\tilde{v}'_j \in \tilde{\sigma}' \cdot \tilde{f}_\#(\tilde{\sigma}') \cdot \dots \cdot \tilde{f}_\#^K(\tilde{\sigma}')$.

Fix j. To check if $\sigma'_j = \sigma_i$ for some i we need only consider $i < M$ where $\ell_i > \ell'_j$ for all $i \ge M$. The symmetric argument implies that for any fixed i we can check if $\sigma_i = \sigma'_j$ for some j. The second and final step of the algorithm is to decide if there exists $i \le I$ such that $\sigma_i = \sigma'_j$ for some j or if there exists $j \le J$ such that $\sigma_i = \sigma'_j$ for some i. If not then the algorithm outputs NO. If yes then the algorithm outputs YES and v_i, v'_j.

It remains to prove that if R and R' have a common subray then the algorithm outputs YES in the second step. Suppose that $R = \rho R''$ and $R' = \rho' R''$ where $\rho \subset R$ and $\rho' \subset R'$ are finite and R'' is a maximal common subray. Lift $R = \rho R'' \subset G$ to $\tilde{R} = \tilde{\rho} \tilde{R}'' \subset \tilde{G}$ and the initial segment $\sigma \subset R$ to an initial segment $\tilde{\sigma} \subset \tilde{R}$. Let $\tilde{f} : \tilde{G} \to \tilde{G}$ be the lift of f that takes the initial endpoint of $\tilde{\sigma}$ to the terminal endpoint of $\tilde{\sigma}$ and note that $\tilde{R} = \tilde{\sigma} \cdot \tilde{f}_\#(\tilde{\sigma}) \cdot \tilde{f}_\#^2(\tilde{\sigma}) \cdot \dots$. Since σ is not a Nielsen path, $|f_\#^k(\sigma)| \to \infty$. It follows that the terminal endpoint $P \in \partial F_n$ of $\tilde{R}$ is an attractor for the action of $\partial \tilde{f}$ and so is not fixed by any covering translation [GJLL98, Proposition I.1]. In particular, $\tilde{f}$ is the only lift of f that fixes P. Lift $R' = \rho' R''$ to $\tilde{R}' = \tilde{\rho}' \tilde{R}''$ and σ' to an initial segment $\tilde{\sigma}'$ of $\tilde{R}'$. The uniqueness of $\tilde{f}$ implies that $\tilde{R}' = \tilde{\sigma}' \cdot \tilde{f}_\#(\tilde{\sigma}') \cdot \tilde{f}_\#^2(\tilde{\sigma}') \cdot \dots$. Let $\tilde{E}''$ be the first height r edge crossed by $\tilde{R}''$. There exist unique $k, k' \ge 0$ such that $\tilde{E}''$ is crossed by $\tilde{f}_\#^k(\tilde{\sigma})$ and by $\tilde{f}_\#^{k'}(\tilde{\sigma}')$. We may assume without loss that $k' \ge k$. By Lemma 6.1, it suffices to show that $k \le K$ or equivalently that $\rho \subset \tilde{\sigma} \cdot \tilde{f}_\#(\tilde{\sigma}) \cdot \dots \cdot \tilde{f}^K(\tilde{\sigma})$.

If H_r is NEG then by the basic splitting property of NEG edges [BFH00, Lemma 4.1.4] there is a unique height r edge whose image under $\tilde{f}^k$ crosses $\tilde{E}''$. Since both $\tilde{\sigma}$ and $\tilde{f}_\#^{k'-k}(\tilde{\sigma}')$ cross such an edge, their intersection is non-empty. It follows that $\tilde{\rho} \subset \tilde{\sigma}_1$ so we are done.

If H_r is EG then there exist $\tilde{x} \in \tilde{\sigma}$ and $\tilde{x}' \in \tilde{f}_\#^{k'-k}(\tilde{\sigma}')$ such that $\tilde{f}^k(\tilde{x}) = \tilde{f}^k(\tilde{x}')$. The path $\tilde{\tau}$ from $\tilde{x}$ to $\tilde{x}'$ decomposes as a concatenation of subpaths $\tilde{\alpha}\tilde{\beta}^{-1}$ where $\tilde{\alpha} \subset \tilde{\rho}$ and $\tilde{\beta} \subset \tilde{\rho}'$. By construction $\tilde{f}_\#^k(\tilde{\tau})$ is trivial so in particular τ is not weakly attracted to Λ^+. Lemma 4.4 implies that $\tilde{\alpha}$ does not cross C edges that project to H_r and so does not contain $\tilde{f}_\#^K(\tilde{E})$ for any edge $\tilde{E}$ that projects to H_r. Since $\tilde{\sigma}$ crosses such an $\tilde{E}$ it follows that $\tilde{\alpha}$ and hence $\tilde{\rho}$ is contained in $\tilde{\sigma} \cdot \tilde{f}_\#(\tilde{\sigma}) \cdot \dots \cdot \tilde{f}^K(\tilde{\sigma})$ as desired. $\square$

6.2. Finding a Fixed Point. The proof of Theorem 1.1 begins with an arbitrary relative train track map $f: G \to G$ and filtration and modifies $f: G \to G$ and the filtration so that it satisfies more and more of the CT properties. Certain steps in this process are inductive and involve consideration of a, not necessarily core, component of a filtration element. Lemma 6.4 below is applied in that context. The reader will note that essentially all of the arguments take place in a core filtration element.

We state our next result in terms of a lift $\tilde{f}: \widetilde{G} \to \widetilde{G}$ of $f: G \to G$ and fixed points for $\tilde{f}$. We could just as easily have stated it in terms of finite paths $\rho \subset G$ as described in Section 3.1 but it seems more natural to work with lifts.

Following Definition 6.2 we say that a completely split path $\tilde{\sigma} \subset \widetilde{G}$ *generates a completely split ray* $\widetilde{R}$ if $\widetilde{R} = \tilde{\sigma} \cdot \tilde{f}_\#(\tilde{\sigma}) \cdot \tilde{f}^2_\#(\tilde{\sigma}) \cdot \dots$. As we have seen, $\tilde{f}_\#(\widetilde{R}) \subset \widetilde{R}$ and $\tilde{f}$ maps the set of splitting vertices for $\widetilde{R}$ into itself.

Lemma 6.4. *Suppose that $f: G \to G$ is a homotopy equivalence of a connected finite graph, that $\emptyset = G_0 \subset G_1 \subset \dots \subset G_K = G$ is an f-invariant filtration and that there is a connected core filtration element G_r such that $f \mid G_r$ is a* CT *and such that for each $r < s \le K$, H_s is a single non-fixed edge E_s satisfying the following properties.*

(1) *The terminal endpoint of E_s is contained in G_r and the initial vertex of E_s has valence one in G.*

(2) *$f(E_s) = E_s \cdot u_s$ where u_s is a completely split closed path whose endpoint is principal for $f \mid G_r$.*

(3) *$f: G \to G$ satisfies (Linear edges) and (NEG Nielsen paths).*

Then there is an algorithm that takes a lift $\tilde{f}: \widetilde{G} \to \widetilde{G}$ of $f: G \to G$ as input and determines if $\text{Fix}(\tilde{f})$ *is non-empty. If it is non-empty then the output of the algorithm is an element of* $\text{Fix}(\tilde{f})$. *If it is empty then the output of the algorithm is a completely split path $\tilde{\sigma} \subset \widetilde{G}_r$ that generates a completely split ray $\widetilde{R} \subset \widetilde{G}_r$. Moreover, if* $\text{Fix}(\tilde{f}) = \emptyset$ *and if the projection $\sigma \subset G_r$ of $\tilde{\sigma}$ is not a Nielsen path then* $\text{Fix}_N(\partial \tilde{f}) = \text{Fix}_N(\partial(\tilde{f} \mid \widetilde{G}_r)) = \{P\}$ *where P is the endpoint of $\widetilde{R}$ and P is not the endpoint of an axis of a covering translation.*

Proof. We dispense with the moreover statement first. Suppose that $\text{Fix}(\tilde{f}) = \emptyset$ and that σ is not a Nielsen path. Then $|f^k_\#(\sigma)| \to \infty$ and [GJLL98, Proposition I.1] implies that the terminal endpoint of $\widetilde{R} = \tilde{\sigma} \cdot \tilde{f}_\#(\tilde{\sigma}) \cdot \tilde{f}^2_\#(\tilde{\sigma}) \cdot \dots$, which is evidently fixed by $\partial \tilde{f}$, is an attractor for the action of $\partial \tilde{f}$, is contained in $\text{Fix}_N(\partial \tilde{f})$ and is not the endpoint of an axis of a covering translation. Since $\text{Fix}(\tilde{f}) = \emptyset$, [FH11, Corollary 3.16] implies that P is the only attractor in $\text{Fix}_N(\partial \tilde{f})$. If there were another point in $\text{Fix}_N(\partial \tilde{f})$ then it would be the endpoint of the axis of a covering translation that commuted with $\tilde{f}$ and the translates of P would be additional attractors in $\text{Fix}_N(\partial \tilde{f})$. Thus $\text{Fix}_N(\partial \tilde{f}) = \{P\}$.

We now turn to the algorithm. Following the proof of [BFH00, Proposition 5.4.3], we say that for each non-fixed vertex $\tilde{v} \in \widetilde{G}_r$, the initial edge of the path from $\tilde{v}$ to $\tilde{f}(\tilde{v})$ is *preferred* by $\tilde{v}$. If both $\widetilde{E}$ and $\widetilde{E}^{-1}$ are preferred by their initial vertices then some sub-interval of $\widetilde{E}$ is mapped over itself by $\tilde{f}$ and so contains a fixed point.

Consider the following (possibly infinite) method for finding either a fixed point in $\widetilde{G}$ or a ray whose terminal endpoint is an element of $\mathrm{Fix}_N(\partial \tilde{f})$. Choose any vertex $\tilde{v}_0 \in \widetilde{G}_r$. If $\tilde{v}_0$ is not fixed, let $\widetilde{E}_0$ be the edge preferred by $\tilde{v}$. If $\widetilde{E}_0^{-1}$ is preferred by the terminal vertex of $\widetilde{E}_0$ then $\widetilde{E}_0$ contains a fixed point that we can find by inspection (Section 3.4). Otherwise, let $\widetilde{E}_1$ be the edge preferred by the terminal vertex of $\widetilde{E}_0$. Repeat this to either find a fixed point in $\widetilde{E}_1$ or define $\widetilde{E}_2$ and so on. If this process does not terminate by finding a fixed point then the ray $\widetilde{R}_{\tilde{v}_0} = \widetilde{E}_1\widetilde{E}_2\ldots$ that it produces converges to a point in $\partial \widetilde{F}_n$ that is evidently fixed and not repelling so is contained in $\mathrm{Fix}_N(\partial \tilde{f})$. For each $m \ge 0$, let $\tilde{\sigma}_m$ be the path connecting the initial endpoint of $\widetilde{E}_m$ to its $\tilde{f}$-image.

Step 1 of the algorithm. Modify the above process by stopping not only if $\widetilde{E}_m$ contains a fixed point but also if $\tilde{\sigma}_m$ is completely split and the turn between $\tilde{\sigma}_m$ and $\tilde{f}_\#(\tilde{\sigma}_m)$ is legal.

To see that this modified process stops in finite time, it suffices to show that if the original process produces a ray $\widetilde{R}_{\tilde{v}_0}$ then at least one of the $\tilde{\sigma}_m$'s has the desired properties. We verify this by following (and tweaking) the proof of [BFH00, Proposition 5.4.3].

Consider the subsequence $\{\tilde{v}_i\}$ of the set of vertices of $\widetilde{R}_{\tilde{v}_0}$ starting with $\tilde{v}_0$ and inductively defined by letting $p \ge i$ be the largest integer such that the closest point to $\tilde{f}(\tilde{v}_i)$ in $\widetilde{E}_0\widetilde{E}_1\ldots\widetilde{E}_p$ is the terminal endpoint of $\widetilde{E}_p$ and then taking $\tilde{v}_{i+1}$ to be the terminal endpoint of $\widetilde{E}_p$. Equivalently, $\tilde{v}_{i+1}$ is the nearest point in $\widetilde{R}_{\tilde{v}_0}$ to $\tilde{f}(\tilde{v}_i)$.

Letting $[\tilde{v}_i, \tilde{v}_{i+1}]$ be the path connecting $\tilde{v}_i$ to $\tilde{v}_{i+1}$, the key property of the $\tilde{v}_i$'s is

$$\tilde{f}_\#([\tilde{v}_i, \tilde{v}_{i+1}]) \supset [\tilde{v}_{i+1}, \tilde{v}_{i+2}].$$

For $m \ge 1$, define

$$\widetilde{Y}_m = \{\tilde{y} \in [\tilde{v}_0, \tilde{v}_1] \colon \tilde{f}^i(\tilde{y}) \in [\tilde{v}_i, \tilde{v}_{i+1}] \text{ for all } 1 \le i \le m\}.$$

The obvious induction argument shows that $\tilde{f}(\widetilde{Y}_m) = [\tilde{v}_m, \tilde{v}_{m+1}]$ and in particular that $\widetilde{Y}_m$ is non-empty. The Y_m's are a nested sequence of closed non-empty subsets of $[\tilde{v}_0, \tilde{v}_1]$ and so their intersection $\bigcap_{m=0}^{\infty} \widetilde{Y}_m$ is non-empty. Each element of $\bigcap_{m=0}^{\infty} \widetilde{Y}_m$ is contained in

$$\widetilde{X} = \{\tilde{x} \colon \{\tilde{x}, \tilde{f}(\tilde{x}), \tilde{f}^2(\tilde{x}), \ldots\} \text{ is an ordered sequence of } \widetilde{R}_{\tilde{v}_0}\}.$$

In the first two paragraphs on p. 68 of [BFH00] it is shown that $\widetilde{X} \subset \widetilde{R}_{\tilde{v}_0}$ contains a vertex $\tilde{v}$ that is the initial vertex of an irreducible edge. For sufficiently large k, the paths $\tilde{\mu} := [\tilde{f}^k(\tilde{v}), \tilde{f}^{k+1}(\tilde{v})]$ and $\tilde{\nu} := [\tilde{f}^{k+1}(\tilde{v}), \tilde{f}^{k+2}(\tilde{v})]$ are completely split by [FH11, Lemma 4.25]. It follows, after increasing k if necessary, that the initial directions of μ^{-1} and ν are periodic by Df. Since $\tilde{v} \in \widetilde{X}$, these directions are distinct and so the turn they define is legal. Letting E_m be the edge in $\widetilde{R}_{\tilde{v}_0}$ that begins with $\tilde{f}^k(\tilde{v})$, we have found the desired $\tilde{\sigma}_m$. (The proof of [BFH00, Proposition 5.4.3] allows the possibility of subdividing at an endpoint of a periodic Nielsen path; in our context, these points are already fixed vertices so no subdivision is required.) This completes the proof that the first part of our algorithm stops in finite time.

If the first step of the algorithm produces a fixed point we are done and the algorithm stops. Suppose then that the first step produces a path $\tilde{\sigma} = \tilde{\sigma}_m$ as above. Let $P \in \mathrm{Fix}_N(\partial \tilde{f})$ be the terminal endpoint of the ray $\widetilde{R} = \tilde{\sigma} \cdot \tilde{f}_\#(\tilde{\sigma}) \cdot \cdots$ generated by $\tilde{\sigma}$. The hard splitting property [FH11, Lemma 4.11(2)] implies that $\widetilde{R}$ is fixed point free. If $\mathrm{Fix}(\tilde{f}) \neq \emptyset$ then there exists a ray $\widetilde{R}'$ with initial endpoint $\tilde{z} \in \mathrm{Fix}(\tilde{f})$, terminal endpoint P and with interior disjoint from $\mathrm{Fix}(\tilde{f})$. The initial edge $\widetilde{E}$ of $\widetilde{R}'$ determines a fixed direction; this follows from [FH11, Lemma 3.16] if $E \subset G_r$ and by hypothesis if $E \subset G \setminus G_r$. Let R_E be the eigenray determined by E (Definition 6.2), let $\widetilde{R}_{\widetilde{E}}$ be the lift of R_E with initial edge $\widetilde{E}$ and let $P' \in \mathrm{Fix}_N(\partial \tilde{f})$ be the terminal endpoint of $\widetilde{R}_{\widetilde{E}}$. If $P \neq P'$ then the line connecting P to P' would be a fixed point free line in $\widetilde{G}_r$ with endpoints in $\mathrm{Fix}_N(\tilde{f})$ in contradiction to [FH11, Lemma 3.16]. Thus $P' = P$ and the rays $\widetilde{R}$ and $\widetilde{R}_{\widetilde{E}} = \widetilde{R}'$ have a common terminal subray.

If σ is a Nielsen path then E is a linear edge, $f(E) = Ew^d$ for some root-free Nielsen path w that forms a circuit, $R_E = Ew^\infty$ and $\widetilde{R}$ is a ray in a line $\widetilde{L}$ that projects to w. It follows that that the terminal endpoint of $\widetilde{E}$ is contained in $\widetilde{L}$. The root-free covering translation that preserves L commutes with $\tilde{f}$. After translating by some iterate of this covering translation we may assume that the terminal endpoint of $\widetilde{E}$ is contained in any chosen lift $\tilde{w}$ of w in $\widetilde{R}$. This analysis justifies the next two steps of the algorithm.

Step 2 of the algorithm. Check if σ is a Nielsen path. If it is not, go to Step 3. If it is, then the algorithm ends as follows. Consider the finite set of points that are the initial vertex of a linear edge $\widetilde{E}$ with terminal endpoint in $\tilde{\sigma}$. If an element of this set is contained in $\mathrm{Fix}(\tilde{f})$ then that point is the output of the algorithm. Otherwise conclude that $\mathrm{Fix}(\tilde{f}) = \emptyset$.

We may now assume that σ is not a Nielsen path and hence that P is not fixed by any covering translation. It follows that $\tilde{f}$ is the only lift of f that fixes P and hence that $\tilde{f}$ fixes the initial endpoint of any eigenray that converges to P.

Step 3 of the algorithm. Apply Lemma 6.3 to the set of eigenrays (Definition 6.2) for f, one by one, to decide if there is an eigenray R_E that shares a terminal end with R. If there is no such eigenray, then $\mathrm{Fix}(\tilde{f}) = \emptyset$. Otherwise, we have an edge E, its eigenray R_E and decompositions $R = \tau R''$ and $R_E = \tau' R''$ for some ray R''. Let τ be the lift of τ that begins with $\tilde{\sigma}$ and let $\tilde{\tau}'$ be the lift of τ' that shares a terminal endpoint with $\tilde{\tau}$. Equivalently, $\widetilde{R} = \tilde{\tau}\widetilde{R}''$ and $\widetilde{R}_{\widetilde{E}} = \tilde{\tau}'\widetilde{R}''$. Then the initial endpoint of $\tilde{\tau}'$ is fixed by $\tilde{f}$. □

7. Upward induction and extension

The original construction of relative train track maps is by downward induction through the strata, making it difficult to prove extension statements. In this paper, we construct CTs using upward induction.

Notation 7.1. Let $\mathcal{F} = \{[F^i]\}$ be a free factor system in F_n. A core graph $K = \bigsqcup_i K_i$ is *$\mathcal{F}$-marked* if each K_i is marked by F^i, i.e. there is a rose $R(F^i)$ whose fundamental group is identified with F^i and a homotopy equivalence $R(F^i) \to K_i$.

Suppose $\phi \in \mathtt{Out}(F_n)$, K is $\mathcal{F}$-marked, $\mathcal{F}$ is ϕ-invariant and $h: K \to K$ is a homotopy equivalence that preserves components. We say that h is a *topological representative of* $\phi|\mathcal{F}$ if each induced map $h \mid K_i: K_i \to K_i$ is a topological representative of $\phi \mid F^i$. If each $h \mid K_i$ is a CT [and satisfying (Inheritance)] then we say that $h: K \to K$ is *a* CT *representing* $\phi \mid \mathcal{F}$ [and satisfying (Inheritance)]. See Section 1 for the definition of (Inheritance).

A topological representative $f: G \to G$ of ϕ is an *extension of* $h: K \to K$ if there is an embedding $K \to G$ respecting markings such that the following diagram commutes.

$$\begin{tikzcd} K \arrow[r, "h"] \arrow[d] & K \arrow[d] \\ G \arrow[r, "f"] & G \end{tikzcd}$$

In this situation we also say that f *extends* h and that h *extends to* F_n.

Interpreted in this language, the proof of Lemma 2.6.7 of [BFH00] shows that every restriction extends.

Lemma 7.2. *Every topological representative for $\phi \mid \mathcal{F}$ extends to F_n.*

Remark 7.3. The proof of Lemma 2.6.7 assumes the property that each $h \mid K_i$ fixes a point. This assumption is not necessary. However, we will only use Lemma 7.2 when $h|K$ is a CT and so this property holds.

As mentioned, we want a construction of CTs that proceeds by upward induction. Our main tool will be a relative version of Theorem 1.1.

Theorem 7.4 (**Extension).** *Suppose that $\phi \in \mathtt{Out}(F_n)$ is rotationless, that $\mathcal{C}$ is a nested sequence $(\mathcal{F} = \mathcal{F}_0) \sqsubset \mathcal{F}_1 \sqsubset \mathcal{F}_2 \sqsubset \cdots \sqsubset (\mathcal{F}_m = \{[F_n]\})$ of ϕ-invariant free factor systems and that $h\colon K \to K$ is a* CT *representing $\phi \mid \mathcal{F}$. Then there is an algorithm that produces a* CT *$f\colon G \to G$ and filtration that represents ϕ, extends $h\colon K \to K$ and such that each element of $\mathcal{C}$ is realized by a core filtration element*; *if $h\colon K \to K$ satisfies* (*Inheritance*) *then $f\colon G \to G$ satisfies* (*Inheritance*).

Theorem 1.1 is the special case of Theorem 7.4 in which $\mathcal{F} = \emptyset$.

The filtration we start with might not be reduced so our algorithm will have to discover reductions, if they exist, as it proceeds.

Theorem 7.5 (extend or reduce). *Suppose that $\phi \in \mathtt{Out}(F_n)$ is rotationless, that $\mathcal{F}$ is a ϕ-invariant free factor system and that $h\colon K \to K$ is a* CT *that represents $\phi \mid \mathcal{F}$* [*and satisfies* (*Inheritance*)]. *Then there is an algorithm that either produces a* CT *$f\colon G \to G$ that represents ϕ* [*satisfies* (*Inheritance*)] *and extends $h\colon K \to K$ or finds a ϕ-invariant proper free factor system $\mathcal{F}'$ that properly contains $\mathcal{F}$.*

Proof of Theorem 7.4 given Theorem 7.5. The *proper length of* $\mathcal{C}$ is the number of inclusions that are proper. If $\mathcal{C}'$ is a sequence of inclusions of ϕ-invariant free factor systems and if each element of $\mathcal{C}$ is an element of $\mathcal{C}'$ then we say that $\mathcal{C}'$ *is an extension of* $\mathcal{C}$. Define $L(\mathcal{C}) \geq 0$ to be the maximal proper length of some $\mathcal{C}'$ extending $\mathcal{C}$. The proof is an induction on $L(\mathcal{C})$. If $L(\mathcal{C}) = 0$ then $\mathcal{F} = \{[F_n]\}$ and the statement is vacuous. If $L(\mathcal{C}) = 1$ the statement follows from Theorem 7.5. Assume then that $L(\mathcal{C}) \geq 2$. We may assume that the inclusions $\mathcal{F} \sqsubset \mathcal{F}_{m-1} \sqsubset \{[F_n]\}$ are proper.

Step 1. Extend h to $\mathcal{F}_{m-1}$. Each component $[F]$ of $\mathcal{F}_{m-1}$ induces the $\phi \mid F$-invariant nested sequence $\mathcal{C}|F$ of free factor systems in F given by

$$(\mathcal{F}|F = \mathcal{F}_0|F) \sqsubset \mathcal{F}_1|F \sqsubset \mathcal{F}_2|F \sqsubset \cdots \sqsubset (\mathcal{F}_{m-1}|F = \{[F]\})$$

where $\mathcal{F}_i|F$ is the union of the components of $\mathcal{F}_i$ that conjugate into F. Clearly $L(\mathcal{C}|F) < L(\mathcal{C})$. Let $K(F)$ be the union of the components C of K such that $[C]$ is conjugate into F and so $[K(F)] = \mathcal{F}|F$. By induction we can algorithmically produce a CT and filtration that represents $\phi \mid F$, extends the restriction of h to $K(F)$ and such that each element of $\mathcal{C}|F$ is realized by a core filtration element; if $h\colon K \to K$ satisfies (Inheritance) then so does this CT. The disjoint union of these CTs is a CT [satisfying (Inheritance) if h does] representing $\phi \mid \mathcal{F}_{m-1}$ such that each element of $\mathcal{F} \sqsubset \mathcal{F}_1 \sqsubset \cdots \sqsubset \mathcal{F}_{m-1}$ is represented by a core filtration element.

Step 2. Further extend to F_n. The sequence $\mathcal{F}_{m-1} \sqsubset \{[F_n]\}$ also has proper length less than $L(\mathcal{C})$. Hence by induction we can algorithmically produce a CT that represents ϕ and extends the restriction of ϕ to $\mathcal{F}_{m-1}$ found in Step 1; if $h\colon K \to K$ satisfies (Inheritance) then so does this CT. $\square$

To prove Theorem 1.1, it remains to prove Theorem 7.5. The proof of Theorem 7.5 is given in Section 8.

8. Proof of Theorem 7.5

The proof of Theorem 7.5 is carried out in Sections 8.1 and 8.2. In both cases, ϕ, $\mathcal{F}$ and $h\colon K \to K$ are as in the statement of the theorem.

8.1. The one-edge extension case. In this section we assume that $\mathcal{F} \sqsubset \{[F_n]\}$ is a one-edge extension, meaning that either $\mathcal{F} = \{[F_1]\}$ where F_1 has rank $n-1$ or $\mathcal{F} = \{[F_1], [F_2]\}$ where the ranks of F_1 and F_2 add to n. In this case, there is a marked graph G obtained from K by adding a single topological edge E and there is a topological representative $g\colon G \to G$ of ϕ that agrees with h on K and satisfies $g(E) = \bar{u}Ev$ or $g(E) = \bar{u}\bar{E}v$ for some possibly trivial paths $u, v \subset K$ by [BFH00, Corollary 3.2.2]. We first complete the proof assuming that $g(E) = \bar{u}Ev$ and then apply CT theory to show that the $g(E) = \bar{u}\bar{E}v$ case does not happen.

If u and v are both trivial then $g(E) = E$ and g satisfies all the properties of a CT except that $K \subset G$ might not be reduced. Check (using Lemma 4.5 and (NEG Nielsen paths)) if there is a Nielsen path $\beta \subset K$ connecting the terminal endpoint of E to the initial endpoint of E. If not, then $K \subset G$ is reduced by Lemma 4.12 and g is a CT. If yes, then $E\beta$ is a closed Nielsen path and the free factor system $\{[K], [E\beta]\}$ is properly contained between $[K]$ and $[G]$.

Suppose next one of u and v is trivial and that the other is not. The cases are symmetric so we assume that u is trivial and v is non-trivial. Let C be the component of K that contains v. If C has rank one then C has a single fixed edge e and single vertex w and $g(E) = Ee^d$ for some $d \neq 0$. There are two subcases. If w is also the initial vertex of E then $n = 2$ and G has only two edges, E and e. (NEG Nielsen paths) follows from the basic splitting property for NEG edges [BFH00, Lemma 4.1.4]. Lemma 4.12 implies that $[K] \sqsubset F_n$ is reduced and so (Filtration) is satisfied. The remaining CT properties are clear so g is a CT. If w is not the initial vertex of e then g is not a CT because w is not principal. In this case we redefine g so that it fixes E and is unchanged on K. The resulting homotopy equivalence $f\colon G \to G$ is homotopic to g and so is still a topological representative of ϕ. We are now back in the case that u and v are trivial. Lemma 4.12 implies that $K \subset G$ is reduced so $f\colon G \to G$ is a CT.

We assume now that C has rank at least two and we apply [FH11, Step 5, pp. 91–93]. Choose a lift $\widetilde{E}$ of E, let $\tilde{g}\colon \widetilde{G} \to \widetilde{G}$ be the lift of g that fixes the initial endpoint of $\widetilde{E}$, let $\Gamma \subset \widetilde{G}$ be the component of the full pre-image of C that contains the terminal endpoint of $\widetilde{E}$ and let $\tilde{h} = \tilde{g} \mid \Gamma\colon \Gamma \to \Gamma$. To make this step algorithmic, we apply Lemma 6.4 to either find a fixed point for $\tilde{h}$ or to conclude that $\tilde{h}$ is fixed point free and find a completely split path $\tilde{\sigma}$ that generates a completely split ray $\widetilde{R}$. No algorithm for this was given in the original proof.

The remainder of [FH11, Step 5, pp. 91–93] can be applied as written in [FH11]. This is a sliding operation that changes the way that E is attached to C. It may be that E becomes a fixed edge for the new modified $g\colon G \to G$. In this case, we go back to the fixed edge case described above and proceed from there. If E is not fixed after the sliding operation then all of the CT properties are satisfied except that $K \subset G$ might not be reduced. The proof now concludes by applying Lemma 4.12 and [FH11, Step 5, p. 91] to conclude that $K \subset G$ is reduced.

The final case is that both u and v are non-trivial. Subdivide $E = \overline{E}_1 E_2$ where $g(E_1) = E_1 u$ and $g(E_2) = E_2 v$. Let C_i be the component of K that contains the terminal endpoint of E_i. If $C_1 = C_2$ has rank one, then $C_1 = C_2$ has a single fixed edge e and single vertex w and $g(E_i) = E_i e^{d_i}$ for some $d_1, d_2 \neq 0$. Define f by $f(e) = e$ and $f(E) = Ee^{d_2 - d_1}$ and note that f is homotopic to g and so is still a topological representative of ϕ. If $d_1 \neq d_2$ then f is a CT. If $d_1 = d_2$ then $\{[e], [E]\}$ is a ϕ-invariant free factor system that is properly contained between $[K]$ and $[F_n = F_2]$.

If $C_1 \neq C_2$ both have rank one then the identity map of G is homotopic to g and is a CT. We may now assume, after interchanging E_1 and E_2 if necessary that C_2 has rank at least two. If C_1 has rank one then g is homotopic to g' that agrees with h on K and that satisfies $g'(E) = Ev$ so we are reduced to a previous case.

The remaining case is that both C_1 and C_2 have rank at least two. The edges E_1 and E_2 are modified in the same way that E was modified in the case that C had rank at least two. Namely, follow [FH11, Step 5, pp. 91–93] and apply Lemma 6.4 to make the construction algorithmic. The edges E_1 and E_2 are considered one at a time so when E_2 is considered the subgraph it is being attached to is $K \cup E_1$ and not K (This accounts for the hypotheses of Lemma 6.4 being what they are.) It may be that after sliding E_2, both of its endpoints are attached to the initial vertex of E_1; see Example 4.10.

This completes the proof in the case that $g(E) = \bar{u}Ev$. Suppose now that $g(E) = \bar{u}\overline{E}v$. Applying the previous case to $h = g^2$ representing $\theta = \phi^2$, we see that there is a CT representing θ and extending $h^2\colon K \to K$. As shown in the first two paragraphs of [FH11, Section 5.1], there is a line $\gamma \subset G$ whose edge path contains one copy of E and zero copies of $\overline{E}$ and such that some (and hence every) lift $\tilde{\gamma}$ has endpoints in $\mathrm{Fix}_N(\Theta)$ for some $\Theta \in \mathrm{P}(\theta)$. Since ϕ is rotationless, there exists $\Phi \in \mathrm{P}(\phi)$ such that $\mathrm{Fix}_N(\Phi) = \mathrm{Fix}_N(\Theta)$. Thus γ is ϕ-invariant which contradicts the fact that the edge path for $h(\gamma)$ crosses $\overline{E}$ once and E zero times.

We conclude that this case does not happen.

8.2. The multi-edge extension case. We assume in this section that $\mathcal{F} \sqsubset \{F_n\}$ is not a one-edge extension. If $f: G \to G$ represents ϕ and extends $h: K \to K$ then the core filtration element that is identified with K will be denoted G_r. In particular, $[G_r] = \mathcal{F}$.

Step 1. There is an algorithm that either produces a relative train track map extending $h: K \to K$ or finds a proper ϕ-invariant free factor system $\mathcal{F}'$ that properly contains $\mathcal{F}$. Moreover, in the former case, either there are no EG strata above the core filtration element G_r that is identified with K or the top stratum H_N of G is the only EG stratum above G_r.

We can not simply apply the relative train track map algorithm of Theorem 2.2. That algorithm was by downward induction through the filtration and so made no effort to leave lower filtration elements untouched. We say that an algorithm is *K-safe* if it preserves the property of extending $h: K \to K$.

Lemma 8.1. *The algorithms of Lemma 2.6 and Lemma 2.8 are K-safe.*

Proof. The algorithms use only subdivision, folding, valence one homotopies and valence two homotopies applied to edges above K and so are K-safe. □

The algorithm of Lemma 2.9 (which is really [BH92, Lemma 5.14]) is not entirely K-safe. In order to isolate the part that is K-safe we introduce some notation and recall the steps in the algorithm.

If H_s is an EG stratum of a topological representative then a non-trivial path $\sigma \subset G_{s-1}$ with endpoints in $H_s \cap G_{s-1}$ is called an *inessential connecting path for H_s* if $f_\#(\sigma)$ is trivial. If the EG stratum H_s satisfies (RTT-i) then it satisfies (RTT-ii) if and only if there are no inessential connecting paths for H_s.

An inessential connecting path σ for H_s is 'collapsed' as follows. Choose a turn in σ whose two directions have the same image under Df. This exists because $f_\#(\sigma)$ is trivial. Then fold initial segments of these edges and tighten the new map if necessary to 'shorten' σ. After finitely many such moves σ is completely folded away and the endpoints of σ are identified to a single vertex thereby reducing the cardinality of $H_s \cap G_{s-1}$.

If $\sigma \subset G_{s-1}$ is disjoint from G_r, for example if σ is contained in a contractible component of G_{s-1}, then collapsing σ is K-safe. We record this fragment of Lemma 2.9 as Lemma 8.3 after adding one more piece of notation.

Notation 8.2. An EG stratum H_s satisfies "(partial RTT-ii)" if the contractible components of G_{s-1} do not contain any inessential connecting paths.

Lemma 8.3. *Suppose that $f: G \to G$ is a bounded topological representative of ϕ that extends $h: K \to K$ and that H_s is an EG stratum that does not satisfy* (partial RTT-ii). *Then there is a K-safe algorithm to construct a bounded topological representative $f': G' \to G'$ of ϕ such that*

(1) *$\Lambda(f) = \Lambda(f')$ and*

(2) *there is a bijection $H_j \leftrightarrow H'_{j'}$ between the EG strata of f and the EG strata of f' such that*

  (a) *relative height is preserved*; *i.e. $j < k$ if and only if $j' < k'$*;

  (b) $|H'_{s'} \cap G'_{s'-1}| < |H_s \cap G_{s-1}|$;

  (c) *if H_s satisfies* (RTT-i) *then $H'_{s'}$ satisfies* (RTT-i).

Remark 8.4. Note that Lemma 8.3 is not exactly parallel to Lemma 2.9. We will only apply Lemma 8.3 with H_s being the first EG stratum above K and we will not be concerned with preserving properties of higher EG strata.

The following corollary takes us as far as we can go using just the techniques of Theorem 2.2.

Corollary 8.5. *Given a bounded topological representative $f: G \to G$ of ϕ that extends $h: K \to K$, there is an algorithm that constructs a bounded topological representative $f': G' \to G'$ of ϕ that extends $h: K \to K$ such that $\Lambda(f') \le \Lambda(f)$ and such that either there are no EG strata above K or the first EG stratum above K satisfies* (RTT-i), (RTT-iii), *and* (partial RTT-ii).

Proof. If there are no EG strata above K then we are done. Otherwise, apply Lemma 2.6 and Lemma 8.1 to produce a bounded topological representative (still called $f: G \to G$) of ϕ that extends $h: K \to K$ whose first EG stratum above K satisfies (RTT-i). If that stratum does not also satisfy (partial RTT-ii), apply Lemma 8.3 to produce a new bounded topological representative (still called $f: G \to G$) of ϕ whose first EG stratum above K still satisfies (RTT-i). Item (b) of Lemma 8.3 guarantees that after finitely many applications of Lemma 8.3, we arrive at $f: G \to G$ whose first EG stratum above K satisfies (RTT-i) and (partial RTT-ii). Apply Lemma 2.4 to check if the first EG stratum above K satisfies (RTT-iii). If yes, we are done. Otherwise apply Lemma 2.8 and Lemma 8.1 to produce a bounded topological representative $f': G' \to G'$ of ϕ that extends $h: K \to K$ with $\Lambda(f') < \Lambda(f)$. Then start over again with $f': G' \to G'$ replacing the original $f: G \to G$. Since every decreasing sequence $\Lambda(f) > \Lambda(f') > \cdots$ is finite (Definition 2.3), this process produces the desired $f: G \to G$ in finite time. □

Before introducing the new move necessary to achieve (RTT-ii) in a K-safe manner, we prove a lemma that will simplify the situation in which the new move is needed.

Lemma 8.6. *Suppose that $f: G \to G$ is a bounded topological representative of ϕ that extends K and that H_s, $s > r$, is an EG stratum that satisfies* (RTT-i), (RTT-iii), *and* (partial RTT-ii). *Then $[G_r] = \mathcal{F}$ is properly contained in $[G_s]$. In particular, if $G_s \neq G$ then $[G_s]$ is a proper ϕ-invariant free factor system that properly contains $\mathcal{F}$.*

Proof. It suffices to show that there is a line in G_s that crosses an edge in H_s and for this it suffices to show that if a vertex $v \in H_s$ is either disjoint from G_{s-1} or is contained in a contractible component of G_{s-1} then there is path $\mu = \bar{E}_1 \tau E_2 \subset G_s$ where E_1, E_2 are edges in H_s and τ is a possibly trivial path that contains v and does not contain an edge in H_s. Since H_s is an EG stratum, the cardinality of $f^{-p}(v) \cap H_s$ goes to infinity with p. Choose $p \geq 1$ and a point x in the interior of an edge $E \subset H_s$ such that $f^p(x) = v$. We may assume without loss that either $f(x) \in G_{s-1}$ or that $f(x)$ is a vertex in H_s. Subdivide E into 'edgelets' that are mapped by f to single edges in G. There is an edgelet subpath $e_1 \dots e_t$ such that $f(e_j)$ is an edge in H_s if and only if $j = 1$ or $j = t$ and such that either $t > 2$ and $x \in e_2 \dots e_{t-1}$ or $t = 2$ and x is the common endpoint of e_1 and e_2. If $t = 2$ then $(f(\bar{e}_1), f(e_2))$ is a legal turn in H_s by (RTT-iii) and so $Df^{p-1}(f(\bar{e}_1))$ and $Df^{p-1}(f(e_2))$ are distinct edges in H_s based at v and we are done (with τ being trivial). If $t > 2$ then $\tau' = f(e_2) \dots f(e_{t-1})$ is a non-trivial path in G_{s-1} with endpoints in $H_s \cap G_{s-1}$. In particular, $f(x) \in G_{s-1}$. It follows that $v = f^p(x) \in G_{s-1}$ and hence (by hypothesis) v is contained in a contractible component of G_{s-1}. This in turn implies that each of $\tau', f_\#(\tau'), \dots, f_\#^{p-1}(\tau')$ is contained in a contractible component of G_{s-1}. Since τ' is a non-trivial path in G_{s-1} with endpoints in $H_s \cap G_{s-1}$, the same is true for $f_\#(\tau'), \dots, f_\#^{p-1}(\tau')$ by (RTT-i) and the (partial RTT-ii) property. Property (RTT-i) implies that $Df^{p-1}(f(\bar{e}_1))$ and $Df^{p-1}(f(e_t))$ are directions in H_s and again we are done. This completes the proof of the lemma. $\square$

Notation 8.7. Let $\mathcal{R}$ be the set of bounded topological representatives $f: G \to G$ of ϕ that extend $h: K \to K$ and such that the top stratum H_N

- is the only EG stratum above K;
- satisfies (RTT-i), (RTT-iii) and (partial RTT-ii).

The failure of $f: G \to G$ in $\mathcal{R}$ to be a relative train track map is measured by the number $\delta(f)$ of directions in H_N that are based at non-periodic vertices in non-contractible components of G_{N-1}.

We will only define the new move in this context.

Lemma 8.8. *Suppose that $f: G \to G$ is an element of $\mathcal{R}$ and that $\delta(f) > 0$. Then there is an algorithm to construct an element $\hat{f}: \widehat{G} \to \widehat{G}$ of $\mathcal{R}$ such that either*

(1) $\Lambda(\hat{f}) < \Lambda(f)$ *or*

(2) $\Lambda(\hat{f}) = \Lambda(f)$ *and* $\delta(\hat{f}) < \delta(f)$

Proof. Let A_p [resp. A_{np}] be the set of periodic [resp. non-periodic] vertices of H_N that are contained in non-contractible components of G_{N-1}. Thus f permutes the elements of A_p and each element of A_{np} is mapped by an iterate of f into A_p. By hypothesis, $A_{np} \neq \emptyset$. Choose an element $x \in A_{np}$ such that $f(x) \in A_p$. There is a unique $y \in A_p$ such that $f(x) = f(y)$ and there is a unique inessential connecting path ν connecting x to y. Choose an edge $E \subset H_N$ with initial endpoint x and slide its initial endpoint along ν to y to produce a new topological representative $f': G' \to G'$. The marked graph G' is obtained from G by replacing E with an edge E' with terminal endpoint equal to that of E and with initial endpoint y. Thus $G \setminus E$ can be viewed as a subgraph of both G and G'. For each edge $e \subset G \setminus E$, the edge path $f'(e)$ is obtained from the edge path $f(e)$ by replacing each copy of E with $\nu E'$ and each copy of $\bar{E}$ with $\bar{E}'\bar{\nu}$ and then tightening. The edge path $f'(E')$ is obtained from the edge path $f_\#(\bar{\nu}E) = f_\#(\bar{\nu})f(E) = f(E)$ by making the same replacements as in the previous case and then tightening. It is clear that $f \mid G_{N-1} = f' \mid G_{N-1}$ and that no edges in H_N are cancelled during the tightening operation. It follows that f' extends $h: K \to K$, that the top stratum H'_N of G' is the only EG stratum above K, that $\Lambda(f) = \Lambda(f')$ and that f' satisfies (partial RTT-ii). By construction, $\delta(f') = \delta(f) - 1$. If the initial direction determined by E is not in the image of Df then H'_N satisfies (RTT-i). If H'_N also satisfies (RTT-iii) then $\hat{f} = f'$ satisfies (2). Otherwise, we apply Lemma 2.8 and Lemma 8.1 to produce $\hat{f}$ satisfying (1).

If the initial direction determined by E is in the image of Df then H'_N does not satisfy (RTT-i). (For example, if $Df(e) = E$ then $f'(e)$ begins with $\nu E'$.) In this case, we perform a core subdivision producing a new map $f'': G'' \to G''$. (Continuing with the example, $e = e_1e_2$ where $f''(e_1) = \nu$ and $f''(e_2)$ begins with E'.) There is one subdivision point for each direction in H_N that is mapped by some iterate of Df to E. If an edge e is subdivided into e_1e_2 then e_1 is a zero stratum for f'' and e_2 replaces e as an edge in the top EG stratum. (It may be that the initial and terminal directions of e are both eventually mapped to E and so e is ultimately subdivided into two zero strata and one edge in the top EG stratum.) The contribution of e_2 to $\delta(f'')$ balances the contribution of e to $\delta(f')$ so $\delta(f'') = \delta(f') < \delta(f)$. If (RTT-iii) is satisfied then $\hat{f} = f''$ satisfies (2). Otherwise, we apply Lemma 2.8 and Lemma 8.1 to produce $\hat{f}$ satisfying (1). □

Step 1 can now be completed as follows. Start with any bounded topological representation of ϕ that extends $h: K \to K$ (Lemma 7.2). Apply Corollary 8.5 to produce a bounded topological representative $f': G' \to G'$ of ϕ that extends

$h: K \to K$ such that $\Lambda(f') \le \Lambda(f)$ and such that either there are no EG strata above K or the first EG stratum $H'_{s'}$ above K satisfies (RTT-i), (RTT-iii) and (partial RTT-ii). In the former case, f' is a relative train track map and we are done so assume the latter holds. If $H'_{s'}$ is not the top stratum then (Lemma 8.6) we have found a proper ϕ-invariant free factor system that properly contains $\mathcal{F}$ and we are done. We may therefore assume that $f': G' \to G'$ is an element of $\mathcal{R}$. If $\delta(f') = 0$ then $f': G' \to G'$ is a relative train track map and we are done. Otherwise, apply Lemma 8.8 to produce $f'': G'' \to G'$ with either $\Lambda(f'') < \Lambda(f') \le \Lambda(f)$ or $\delta(f'') < \delta(f')$. In the former case, we start all over. This can only happen a finite number of times so we may assume that we are in the case that $\delta(f'') < \delta(f')$. After applying Lemma 8.8 finitely many times, we arrive a relative train track map and are done.

Step 2. Modify $f: G \to G$ from Step 1 so that the conclusions of Theorem 2.12 are satisfied.

The algorithm for Step 2 is explicitly described in Subsection 3.4 and in [FH11, pp. 18–23]. Each move effects only strata above G_r and so is K-safe. For future reference we note that the number of edges in each EG stratum and the number of indivisible Nielsen paths of EG height are not increased in Step 2. This ends Step 2.

Going forward, we may now also assume (with justification below) that

(i) $f: G \to G$ is rotationless;

(ii) $\mathcal{F} \subset \{[F_n]\}$ is irreducible;

(iii) there are no periodic strata above G_r;

(iv) every indivisible periodic Nielsen path ρ of EG height has period one;

(v) H_N is EG and aperiodic, meaning that some iterate of its transition matrix is positive.

Let $f: G \to G$ be as in Step 2. Item (i) follows from [FH11, Proposition 3.29]. Applying the algorithm of Proposition 4.9, we can either find a reduction of $\mathcal{F} \subset \{[F_n]\}$, in which case we are done, or conclude that $\mathcal{F} \subset \{[F_n]\}$ is irreducible. We may therefore assume that (ii) is satisfied. Item (iii) follows from (ii) and property (P) of Theorem 2.12.

If there are no EG strata above G_r then each stratum H_j above G_r is non-periodic NEG by (iii) and property (Z) of Theorem 2.12. Item (i) and property (NEG) of Theorem 2.12 then imply that each H_j is a single edge E_j with terminal endpoint in a core filtration element of height less than j. By property (F) of Theorem 2.12, the core of each filtration element is a filtration element. Letting G_t be the first core filtration element above G_r, it follows that $[G_r] \sqsubset [G_t]$ is a one-edge extension. By hypothesis, $[G_r] \sqsubset [G_N]$ is not a one-edge extension.

We have therefore found a reduction in contradiction to (ii). We conclude that H_N is EG and is the only irreducible stratum above G_r (recall the output of Step 1).

We need only check (iv) for an indivisible periodic Nielsen path ρ of height N. Any such ρ begins and ends with edges in H_N by [BH92, Lemma 5.11]. It follows that the endpoints of ρ are incident to at least one periodic direction in H_N and so are principal by Definition 3.5. Item (iv) therefore follows from (i) and [FH11, Proposition 3.29]. Item (i) and [FH11, Lemma 3.19] imply that there are fixed directions in H_N which in turn implies that H_N is aperiodic so (v) is satisfied.

Step 3. Modify $f\colon G \to G$ from Step 2 to produce a relative train track map that satisfies (EG Nielsen paths).

An algorithm for Step 3 is given in [FH11, pp. 87–89] but it is not entirely K-safe because it makes use of the 'collapsing inessential connecting paths' move in the relative train track map algorithm. If we carry out this collapse as described in Step 1 above rather than as described in [BH92] then the algorithm becomes K-safe and we can use it to complete Step 3. For the readers convenience, we summarize this algorithm and point out the one place where the K-safe modification occurs.

Remark 8.9. The algorithm of [FH11] applies results from [BH92, Sections 3 and 5] and [BFH00, Sections 5.2 and 5.3]. Section 5.3 of [BFH00] has the global hypothesis that H_N is a geometric stratum. That hypothesis is not used in any of the results cited so our conclusions also hold in the non-geometric case. Most of the results cited ultimately derive from Section 5 of [BH92] where there is no assumption that H_N is geometric.

By [BH92, Lemma 5.11], every indivisible Nielsen path of height N has the form $\rho = \bar{\alpha}\beta$ where (α, β) is the only illegal turn of height N in ρ. Let E_1 and E_2 be the first edges of α and β respectively. Depending on the edge paths $f(E_1)$ and $f(E_2)$, there are three ways in which $f\colon G \to G$ and ρ can be modified to produce a new relative train track map $f'\colon G' \to G'$ and indivisible Nielsen path $\rho' \subset G'$. In all three cases, the number of edges in EG strata and the number of indivisible Nielsen paths of EG height do not increase. The first and third are K-safe as described in [BH92] and [BFH00]. To make the second K-safe we use Lemma 8.8 instead of Lemma 2.9.

If one of $f(E_1)$ and $f(E_2)$, say $f(E_2)$, is properly contained in the other then the fold is said to be *proper*. There is a maximal path $\delta \subset G_{N-1}$ such that $E_2\delta$ is an initial segment of β. In this case, $f'\colon G' \to G'$ is defined by folding an initial segment of E_1 with $E_2\delta$; see [BFH00, Definition 5.3.2 and Lemma 5.3.3]. The indivisible Nielsen path ρ' is the tightened image of ρ in G' under the folding map.

In the *improper* case, $f(E_1) = f(E_2)$ and one begins the process [BFH00, Definition 5.3.4] by folding E_1 and E_2 to form a single new edge. If the edge following E_1 in α or the edge following E_2 in β belongs to H_N then both of those edges belong to H_N and nothing more is required. Otherwise, the resulting map is not a relative train track map and one must perform core subdivisions and collapses of inessential connecting paths to restore the relative train track map properties. These should be done as in Step 1 so as to be K-safe. In this case, the the number of edges in the EG stratum decreases. See [BFH00, Lemma 5.3.5].

The third possibility, a *partial fold*, is that the maximal common subpath of $f(E_1)$ and $f(E_2)$ is proper in both edge paths. It follows [BH92, p. 25] that $\alpha = E_1$ and $\beta = E_2$. Items (ii) and (v) and [BFH00, Lemma 5.1.7] imply that the endpoints of ρ are distinct and that if both endpoints are contained in G_{N-1} then at least one of them is contained in a contractible component of G_{N-1} [BFH00, Lemma 5.1.7]. Item (iii) above implies that there are no invariant contractible components of G_{N-1} and we conclude that at least one of the endpoints of ρ is disjoint from G_{N-1}. In this case, do not perform a standard fold but rather entirely identify E_1 and E_2 to form a new graph G' with an induced map $f' \colon G' \to G'$. Since the turn (α, β) is not taken by $f(E)$ for any edge E, the image in G' of $f(E)$ is already tight and $f' \colon G' \to G'$ is a topological representative. It is immediate from the construction that (RTT-i) is preserved. Since at least one of the endpoints of ρ is disjoint from G_{N-1}, (RTT-ii) is also preserved. As argued in the proof of [BFH00, Lemma 5.2.4], $\Lambda' = \Lambda$ so Lemma 2.8 and Lemma 8.1 imply that (RTT-iii) is satisfied and $f' \colon G' \to G'$ is a relative train track map. By construction (see also the proof of [BFH00, Lemma 5.2.4]) the number of indivisible Nielsen paths with EG height has been decreased.

Since each of these operations begins and ends with a relative train track map and an indivisible Nielsen path, we can iteratively fold to produce a sequence of relative train track maps and indivisible Nielsen paths. Let $C(f)$ be the sum of the entries in the transition matrix for f. This gives an upper bound for the number of elementary folds in a Stallings factorization of $f \colon G \to G$. If the first $C(f)$ folds encountered are proper then f satisfies (EG Nielsen paths); see [FH11, Lemma 4.33 and Remark 4.34] and the proof of [BFH00, Lemma 5.3.6]. Otherwise, one encounters either a partial or an improper fold in which case either the number of edges in EG strata or the number of indivisible Nielsen paths of EG height decreases and we start over. Since these numbers never increase, this algorithm terminates in finite time.

At the end of Step 3, $f \colon G \to G$ satisfies (EG Nielsen paths). If the relative train track map that was the input to Step 3 did not satisfy (EG Nielsen paths) then either the number of edges in an EG stratum decreased or the number of indivisible Nielsen paths of EG height decreased during Step 3. Throughout the remaining steps, the number of edges in each EG stratum and the number of indivisible Nielsen paths of EG height are never increased. As a result, we need not verify the (EG Nielsen paths) after each step. If it fails at some point, go back to Step 3 and start again. This can only happen finitely many times so does not prevent the process from terminating.

Step 4. If the relative train track map produced by Step 3 does not satisfy the conclusions of Theorem 2.12 return to Step 2. As just noted, changes are required in Step 3 only a finite number of times so Step 4 is a finite process.

Step 5. (Rotationless), (Filtration), (Zero strata), and (Periodic edges). The first two properties follow from [FH11, Proposition 3.29] and Theorem 2.12 respectively. (Zero strata) is arranged by the K-safe tree replacement moves described in [FH11, Step 3]. (Periodic edges) follows from the assumption that $f \mid G_r$ is a CT and item (c) above.

Step 6. Modify $f: G \to G$ from Step 5 so that it satisfies (Vertices), (Linear edges) and (NEG Nielsen paths). Additionally, the modified $f: G \to G$ satisfies (Completely split) on all NEG edges.

By property (Z) of Theorem 2.12 there is an irreducible stratum H_p such that each stratum between H_p and H_N is a zero stratum that is a component of G_{N-1}. The edges $E_1, \dots, E_p$ of $G_p \setminus G_r$ are non-fixed NEG edges with terminal endpoints in G_r and initial endpoints that have valence one in G_{N-1}. The proof of this assertion is essentially the same as the proof of (v) and is left to the reader. The three conditions to be achieved depend only on the NEG edges E_i. Since $f: G \to G$ is rotationless and satisfies Theorem 2.12, $f(E_i) = E_i u_i$ for some $u_i \subset G_r$. By item (c) above, u_i is non-trivial.

Suppose that C is a rank one component of G_r and that the unique vertex w of C is the terminal endpoint of $E^1, \dots, E^q \subset \{E_1, \dots, E_p\}$. Then C has a single edge e and $f(E^j) = E^j e^{d_j}$ for some $d_j \neq 0$. If w is not the endpoint of an edge in H_N then redefine f on the E^j's by $f(E^j) = E^j e^{d_j - d_1}$. The new map still represents ϕ and none of our established properties are lost. The edge E^1 is now fixed and so can be collapsed. After these moves, w is the endpoint of at least one edge in H_N and so is principal for f. We assume now that the unique vertex of each rank one component of G_r is the endpoint of an edge in H_N.

With that special case out of the way, one just applies [FH11, Step 5 pp. 91–93], applying Lemma 6.4 as described in Section 8.1.

Step 7. Modify $f: G \to G$ from Step 6 so that it satisfies (Completely split). We can apply [FH11, Step 6] without change.

This completes the proof of Theorem 7.5.

9. Finding Fix(Φ)

The goal of this section is to give another proof of the result of Bogopolski-Maslakova [BM16] that there is an algorithm that, given $\Phi \in \mathsf{Aut}(F_n)$, computes $\mathrm{Fix}(\Phi)$.

9.1. The periodic case. In this section, we examine the special case that Φ is periodic. The analysis in this section will parallel that of the general case.

Recall from Section 3.1 that if G is a marked graph and $\tilde{v} \in \widetilde{G}$ is a lift to the universal cover of $v \in G$ then there is an isomorphism $J_{\tilde{v}}\colon \pi_1(G, v) \to \mathcal{T}(\widetilde{G})$ given by mapping $[\gamma]$ to the covering translation T of $\widetilde{G}$ that takes $\tilde{v}$ to the terminal endpoint of the lift $\tilde{\gamma}$ of γ with initial endpoint $\tilde{v}$.

Lemma 9.1. *Suppose that $h\colon G \to G$ is a periodic homeomorphism of a marked graph G and that $\tilde{h}\colon \widetilde{G} \to \widetilde{G}$ is a periodic lift of h to the universal cover $\widetilde{G}$. Then*

(1) $\mathrm{Fix}(\tilde{h}) \neq \emptyset$ *and a point* $\tilde{v} \in \mathrm{Fix}(\tilde{f})$ *can be found algorithmically*;

(2) *if* $\tilde{v} \in \mathrm{Fix}(\tilde{h})$ *projects to* $v \in G$ *then* $J_{\tilde{v}}(\pi_1(\mathrm{Fix}(h), v)) = Z_{\mathcal{T}}(\tilde{h})$. (*See Definition* 3.1.)

Proof. After subdividing if necessary we may assume that h, and hence $\tilde{h}$, pointwise fixes each edge that it preserves. We are now in the setting of Bass-Serre theory and we use its language. Note that $\tilde{h}$ is not hyperbolic, for otherwise $\tilde{h}$ has infinite order. Hence $\tilde{h}$ is elliptic; equivalently $\mathrm{Fix}(\tilde{h}) \neq \emptyset$. It is algorithmic to find a fixed point $\tilde{v}$. Indeed, if $\tilde{x} \in \widetilde{G}$ then the midpoint of $[\tilde{x}, \tilde{h}(\tilde{x})]$ is fixed. This completes the proof of (1).

For (2), let $\mathrm{Fix}(\tilde{h})$ denote the subtree of $\widetilde{G}$ consisting of $\tilde{h}$-fixed edges. Given $[\gamma] \in \pi_1(\mathrm{Fix}(h), v)$, let $\tilde{\gamma}$ be the lift of γ that begins at $\tilde{v}$ and let $\tilde{w}$ be the terminal endpoint of $\tilde{\gamma}$. Then $T = J_{\tilde{v}}(\gamma)$ is the covering translation that carries $\tilde{v}$ to $\tilde{w}$. Since $\tilde{h}$ fixes $\tilde{v}$ and $\gamma \subset \mathrm{Fix}(h)$, $\tilde{h}$ fixes $\tilde{w}$. In particular, $T \circ \tilde{h}(\tilde{v}) = T(\tilde{v}) = \tilde{w} = \tilde{h}(\tilde{w}) = \tilde{h} \circ T(\tilde{v})$ and so T and $\tilde{h}$ commute. We see that $J_{\tilde{v}}(\pi_1(\mathrm{Fix}(h), v))$ is contained in $Z_{\mathcal{T}}(\tilde{h})$. To see surjectivity, let $T \in Z_{\mathcal{T}}(\tilde{h})$. Then $T(\tilde{v}) \in \mathrm{Fix}(\tilde{h})$. Since $\mathrm{Fix}(\tilde{h})$ is a tree, $[\tilde{v}, T(\tilde{v})]$ descends to a closed path in $\mathrm{Fix}(h)$ based at v. $\square$

We could not find a reference for the following result so we have included a proof.

Lemma 9.2. *There is an algorithm that, given periodic $\Phi \in \mathtt{Aut}(F_n)$, computes* $\mathrm{Fix}(\Phi)$.

Proof. Let $\phi \in \mathtt{Out}(F_n)$ be represented by Φ. Since the only periodic automorphisms of $\mathbb{Z}$ are the identity and $x \mapsto -x$, we may assume that $n \geq 2$. The relative train track algorithm of [BH92] (see Theorem 2.2) finds a periodic homeomorphism $h\colon G \to G$ of a marked graph representing ϕ.

We recall some notation from Section 3.1. The marking homotopy equivalence is $\mu\colon (R_n, *) \to (G, \star)$. Via a lift of $\star$ to $\tilde{\star} \in \widetilde{G}$ we have an identification of $\mathcal{T}(\widetilde{G})$ with F_n, an isomorphism $J_{\tilde{\star}}\colon \pi_1(G, \star) \to \mathcal{T}(\widetilde{G})$ and a lift $\tilde{h}\colon \widetilde{G} \to \widetilde{G}$ that can be found algorithmically and that corresponds to Φ in a sense made precise in Section 3.1. The key points for us are that $\tilde{h}$ is periodic and (Lemma 3.2) that $Z_{\mathcal{T}}(\tilde{h})$ and $\mathrm{Fix}(\Phi)$ are equal when viewed as subgroups of $\mathcal{T}(\widetilde{G})$. Since F_n has been identified via $\mu_\#$ with $\pi_1(G, \star)$, our goal is to find $J_{\tilde{\star}}^{-1} Z_{\mathcal{T}}(\tilde{h}) < \pi_1(G, \star)$.

By Lemma 9.1 we can algorithmically find an element $\tilde{v} \in \mathrm{Fix}(\tilde{h})$. Moreover, letting $v \in G$ be the projection of $\tilde{v}$ and $H := \pi_1(\mathrm{Fix}(h), v) < \pi_1(G, v)$, we have $J_{\tilde{v}}(H) = Z_{\mathcal{T}}(\tilde{h})$. Let $\tilde{\eta}$ be the path in $\widetilde{G}$ from $\tilde{\star}$ to $\tilde{v}$ and let $\eta \subset G$ be its projection. A quick chase through the definitions shows that $J_{\tilde{\star}}^{-1} J_{\tilde{v}} \colon \pi_1(G, v) \to \pi_1(G, \star)$ is defined by $[\gamma] \to [\eta\gamma\eta^{-1}]$. Thus $\mathrm{Fix}(\Phi)$ is identified with $H^\eta < \pi_1(G, \star)$.

To recap, the algorithm is

$$\Phi \rightsquigarrow \tilde{h} \rightsquigarrow \tilde{v} \in \mathrm{Fix}(\tilde{h}) \rightsquigarrow H = \pi_1(\mathrm{Fix}(h), v) < \pi_1(G, v) \rightsquigarrow H^\eta < \pi_1(G, \star)$$

□

9.2. A G-graph of Nielsen paths. For the remainder of Section 9 we assume that $f\colon G \to G$ is a CT representing (a necessarily rotationless) $\phi \in \mathsf{Out}(F_n)$.

Let Σ be a (not necessarily connected) graph. It is often useful to work in the Stallings category of graphs labeled by Σ or Σ-graphs, i.e. an object is a graph H with a cellular immersion $H \to \Sigma$ and a morphism from $H \to \Sigma$ to $H' \to \Sigma$ is a cellular immersion $H \to H'$ making the following diagram commute:

$$\begin{tikzcd} H \arrow[r] \arrow[dr] & H' \arrow[d] \\ & \Sigma \end{tikzcd}$$

The map to Σ is often suppressed. In this section we assume that H is finite but in Section 12 we allow H to be infinite. If we give H the CW-structure whose vertex set is the preimage of the vertex set of Σ, then the resulting edges of H (often called *edgelets*) are *labeled* by their image edges in Σ and we refer to the oriented edges of Σ as the set of *labels*. An edge-path is *labeled* by its sequence of oriented edges. The *core of* H is the Σ-graph that is the union of all immersed circles in H. H is *core* if it is its own core.

Σ-graphs are useful because, on each component of H, the immersion $H \to \Sigma$ induces an injection on the level of π_1 and so H is a *geometric realization* of a collection of conjugacy classes $\pi_1(\Sigma)$ of subgroups of F_n indexed by the components of H. Our goal in this section is to construct a G-graph that is the geometric realization of the collection of conjugacy classes $[\mathrm{Fix}(\Phi)]$ as $[\Phi]$ varies over isogredience classes $\mathrm{P}(\phi)/\sim$ of principal automorphisms representing ϕ.

Review 9.3. We precede the construction with a quick review of Nielsen paths in a CT. Since Nielsen paths with endpoints at vertices split as products of fixed edges and indivisible Nielsen paths, we focus on indivisible Nielsen paths. There are only two sources of indivisible Nielsen paths μ. By the (Vertices) property of a CT, the endpoints of indivisible Nielsen paths are always vertices. If $E \in \mathrm{Lin}(f)$, the set of linear edges in G, then $f(E) = Ew_E^d$ for some twist path w_E and some $d \neq 0$ and $Ew_E^k\bar{E}$ is an indivisible Nielsen path for $k \neq 0$. By the (NEG Nielsen paths) property of a CT, all indivisible Nielsen paths of NEG-height have this form. To E we associate a G-graph Y_{μ_E} which is a lollipop. Specifically, Y_{μ_E} is the union of an edge labeled E and a circle labeled w_E attached to the terminal end of E. Note that each $Ew_E^k\bar{E} \subset G$ lifts to a path in Y_{μ_E} with both endpoints at the initial vertex of the edge labeled E. The other possibility is an indivisible Nielsen path μ of EG-height, say r. In this case, μ and $\bar{\mu}$ are the only indivisible Nielsen paths of height r and the initial edges of μ and $\bar{\mu}$ are distinct edges of H_r [FH11, Lemma 4.19]. Further, $\mu = \alpha\bar{\beta}$ where α and β are r-legal and the turn $(\bar{\alpha}, \bar{\beta})$ is illegal of height r [FH11, Lemma 2.11]. See [FH11, Section 4] for details on Nielsen paths in CTs.

Definition 9.4. We construct a G-graph $\widehat{S}(f)$ as follows. If $n = 1$, then $G = R_1$ (a rank 1 rose) and $\widehat{S}(f) := G$. Otherwise, start with the subgraph $\widehat{S}_1(f)$ of G consisting of all vertices in $\mathrm{Fix}(f)$ and all fixed edges. For each $E \in \mathrm{Lin}(f)$, attach the lollipop Y_{μ_E} to $\widehat{S}_1(f)$ at the initial vertex of E thought of as a vertex in $\widehat{S}_1(f)$. For each EG stratum with an indivisible Nielsen path of that height, choose one μ of that height (there are only two and they differ by orientation) and attach an edge, say E_μ, labeled by μ to $\widehat{S}_1(f)$ with initial and terminal endpoints equal to those of μ. This completes the construction of the graph $\widehat{S}(f)$. There is a natural map $h\colon \widehat{S}(f) \to G$ given by inclusion on $\widehat{S}_1(f)$ and by the G-graph structures on each Y_{μ_E} and E_μ. By construction, h is an immersion away from the attaching points in $\widehat{S}_1(f)$ and is a local homeomorphism at each vertex in $\widehat{S}_1(f)$. Thus $\widehat{S}(f)$ is a G-graph.

Let $v \in \mathrm{Fix}(f)$. We abuse notation slightly by also denoting the unique lift of v in $\widehat{S}_1(f)$ by v. Define $\widehat{S}(f, v)$ to be the component of $\widehat{S}(f)$ that contains v. It is possible that $\widehat{S}(f, v)$ is not core.

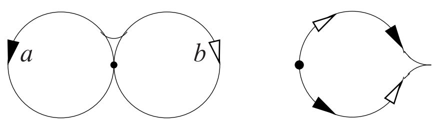

Figure 2. A CT $f\colon G \to G$ given by $a \mapsto ab, b \mapsto bab$ and the graph $\widehat{S}(f)$. $\widehat{S}_1(f)$ is the unique vertex of G. $\widehat{S}(f)$ is the closed Nielsen path $\mu = ab\bar{a}\bar{b}$.

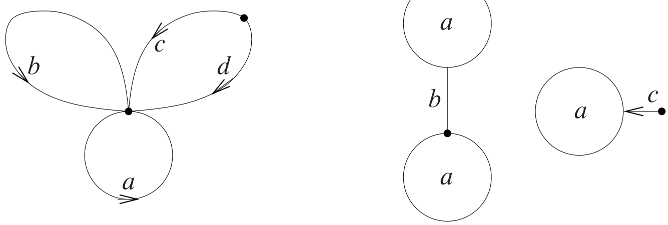

Figure 3. A CT $h\colon H \to H$ given by $a \mapsto a, b \mapsto ba^2, c \mapsto ca, d \mapsto db$ and the graph $\widehat{S}(h)$. $\widehat{S}_1(h)$ is the loop a and the common initial vertex of c and d. Add the lollipop associated to b to the former component and the lollipop associated to c to the latter to make $\widehat{S}(h)$.

Remark 9.5. Since the construction of $f: G \to G$ is algorithmic, finding the set of indivisible Nielsen paths of EG height, and finding $\mathrm{Fix}(f)$ are all algorithmic, it follows that the constructions of $\widehat{S}(f)$ and $\widehat{S}(f, v)$ are algorithmic.

Lemma 9.6. (1) *A path $\sigma \subset G$ with endpoints at vertices lifts to a (necessarily unique) path $\hat{\sigma} \subset \widehat{S}(f)$ with endpoints in $\widehat{S}_1(f)$ if and only if σ is a Nielsen path. Moreover, σ is closed if and only if $\hat{\sigma}$ is closed.*

(2) *If f represents $\phi \in \mathtt{Out}(F_n)$ then non-trivial ϕ-periodic (equivalently ϕ-fixed) conjugacy classes in F_n are characterized as those classes represented by circuits in $\widehat{S}(f)$.*

Proof. (1) The main statement is an immediate consequence of the construction of $\widehat{S}_1(f)$ and the fact that a Nielsen path in a CT is the concatenation of fixed edges and indivisible Nielsen paths with endpoints at vertices. The moreover part follows from the fact that each vertex in G has a unique lift into $\widehat{S}_1(f)$.

(2) Suppose that the circuit $\sigma \subset G$ represents a ϕ-fixed conjugacy class. [FH11, Lemmas 4.11 and 4.25] imply that $f_\#^k(\sigma)$ has a unique complete splitting for all sufficiently large k and hence σ has a unique complete splitting. It follows that each term in the splitting is a periodic, and hence fixed, Nielsen path. Viewing σ as a closed path with endpoint at one of the splitting vertices, we can lift σ to a path $\hat{\sigma} \subset \widehat{S}(f)$ with endpoints in $\widehat{S}_1(f)$. Since each vertex in G has a unique lift into $\widehat{S}_1(f)$, $\hat{\sigma}$ is a closed path and hence a circuit (because it projects to a circuit). Conversely, every circuit in $\widehat{S}(f)$ projects in G to a concatenation of fixed edges and indivisible Nielsen paths. □

Definition 9.7. We will also have need of the subgraph $S(f) := \bigcup_v S(f, v)$ of $\widehat{S}(f)$ where the union is over principal vertices v of G. Equivalently (see Lemma 3.8), $S(f)$ is obtained from $\widehat{S}(f)$ by removing components consisting of a single non-principal vertex. The construction of $S(f)$ goes exactly as in Definition 9.4 except we start with the subgraph $S_1(f)$ consisting of all principal vertices and fixed edges in $\mathrm{Fix}(f)$.

9.3. Fix(Φ) for rotationless ϕ. In this section we compute $\mathrm{Fix}(\Phi)$ for (not necessarily principal) $\Phi \in \mathtt{Aut}(F_n)$ representing rotationless $\phi \in \mathtt{Out}(F_n)$. We begin with the analog of Lemma 9.1(2). Recall that the unique lift of $v \in \mathrm{Fix}(f)$ in $\widehat{S}_1(f, v)$ is denoted v.

Lemma 9.8. *For each vertex $v \in \mathrm{Fix}(f)$ and lift $\tilde{v} \in \widetilde{G}$ let*

$$\widehat{J}_{\tilde{v}} = J_{\tilde{v}} h_\# \colon \pi_1(\widehat{S}(f, v), v) \longrightarrow \mathcal{T}(\widetilde{G}),$$

where

$$h_\# \colon \pi_1(\widehat{S}(f, v), v) \longrightarrow \pi_1(G, v)$$

is induced by the immersion $h\colon \widehat{S}(f,v) \to G$. *Let*

$$\tilde{f}\colon \widetilde{G} \longrightarrow \widetilde{G}$$

be the lift of f *that fixes* $\tilde{v}$. *Then* $\widehat{J}_{\tilde{v}}$ *is injective and has image equal to* $Z_{\mathcal{T}}(\tilde{f})$.

Proof. $\widehat{J}_{\tilde{v}}$ is injective because $h_\#$ is injective and $J_{\tilde{v}}$ is an isomorphism.

If $\hat{\gamma} \subset \widehat{S}(f,v)$ is a closed path based at v then $\gamma := h(\hat{\gamma}) \subset G$ is a closed Nielsen path based at v by Lemma 9.6. Lift γ to a Nielsen path $\tilde{\gamma} \subset \widetilde{G}$ for $\tilde{f}$ with initial endpoint $\tilde{v}$. By definition, $T := \widehat{J}_{\tilde{v}}([\gamma])$ is the covering translation that maps $\tilde{v}$ to the terminal endpoint $\tilde{w}$ of $\tilde{\gamma}$. Since $\tilde{w} \in \mathrm{Fix}(\tilde{f})$, we have $\tilde{f} \circ T(\tilde{v}) = \tilde{f}(\tilde{w}) = \tilde{w} = T(\tilde{v}) = T \circ \tilde{f}(\tilde{v})$ and so $\tilde{f} \circ T = T \circ \tilde{f}$, i.e. $T \in Z_{\mathcal{T}}(\tilde{f})$.

To see that $\widehat{J}_{\tilde{v}}$ is surjective, let $T \in Z_{\mathcal{T}}(\tilde{f})$. Then $T(\tilde{v})$ is fixed by $\tilde{f}$ and the path $\tilde{\gamma}$ from $\tilde{v}$ to $\tilde{w}$ is a Nielsen path for $\tilde{f}$ and so projects to a closed Nielsen path $\gamma \subset G$ based at v that lifts to a closed Nielsen path $\hat{\gamma} \subset \widehat{S}(f)$ based at $\hat{v}$. By construction, $\widehat{J}_{\tilde{v}}[\hat{\gamma}] = T$. □

Lemma 9.9. *There is an algorithm that, given rotationless* $\phi \in \mathtt{Out}(F_n)$, *computes* $\mathrm{Fix}(\Phi)$ *for* $\Phi \in \phi$.

Proof. Let $f\colon G \to G$ be a CT representing the element $\phi \in \mathtt{Out}(F_n)$ determined by Φ. Subdivide if necessary so that every isolated fixed point of f is a vertex. The setup is similar to that of Lemma 9.2. The marking homotopy equivalence is $\mu\colon (R_n, *) \to (G, \star)$. Via a lift of $\star$ to $\tilde{\star} \in \widetilde{G}$ we have an identification of $\mathcal{T}(\widetilde{G})$ with F_n, an isomorphism $J_{\tilde{\star}}\colon \pi_1(G,\star) \to \mathcal{T}(\widetilde{G})$ and a lift $\tilde{f}\colon \widetilde{G} \to \widetilde{G}$ that can be found algorithmically and that corresponds to Φ in a sense made precise in Section 3.1. The key point for us is Lemma 3.2 which states that $Z_{\mathcal{T}}(\tilde{f})$ and $\mathrm{Fix}(\Phi)$ are equal when viewed as subgroups of $\mathcal{T}(\widetilde{G})$.

Apply Lemma 6.4 to decide if $\mathrm{Fix}(\tilde{f}) = \emptyset$.

If $\mathrm{Fix}(\tilde{f}) \neq \emptyset$ then Lemma 6.4 finds $\tilde{v} \in \mathrm{Fix}(\tilde{f})$. Since isolated fixed points of f are vertices, we may assume that $\tilde{v}$ is a vertex. Indeed, by Convention 2.1 non-isolated fixed points only occur in fixed edges.

Let $v \in G$ be the projection of $\tilde{v}$, let $\eta \subset G$ be the projection of the path $\tilde{\eta} \subset \widetilde{G}$ from $\tilde{\star}$ to $\tilde{v}$, and let $h\colon \widehat{S}(f,v) \to G$ be the immersion given by the G-structure. Define $H = h_\#(\pi_1(\widehat{S}(f,v),v)) < \pi_1(G,v)$. Arguing as in Lemma 9.2, with Lemma 9.1(2) replaced by Lemma 9.8 we conclude that $\mathrm{Fix}(\Phi)$ is identified with $H^\eta < \pi_1(G,\star)$. Since F_n has been identified via $\mu_\#$ with $\pi_1(G,\star)$, we are done. In summary,

$$\begin{aligned}\Phi \rightsquigarrow \tilde{f} \rightsquigarrow \tilde{v} \in \mathrm{Fix}(\tilde{f}) \rightsquigarrow H &= h_\#(\pi_1(\widehat{S}(f,v),v)) \\ &< \pi_1(G,v) \rightsquigarrow H^\eta < \pi_1(G,\star).\end{aligned}$$

If $\mathrm{Fix}(\tilde f) = \emptyset$ then $\tilde f$ is not principal [FH11, Corollary 3.17] and so we have $\mathrm{rank}(\mathrm{Fix}(\Phi)) \le 1$ [FH11, Remark 3.3] and $\mathrm{Fix}_N(\partial \tilde f) = \mathrm{Fix}_N(\partial\Phi)$ has at most two points. Lemma 6.4 finds a completely split path $\tilde\sigma$ that generates a ray $\widetilde R$. There are two subcases. If the projected image $\sigma \subset G$ is not a Nielsen path then $|f^k_\#(\sigma)| \to \infty$ and [GJLL98, Proposition I.1] implies that the terminal endpoint P of $\widetilde R = \tilde\sigma \cdot \tilde f_\#(\tilde\sigma) \cdot \tilde f^2_\#(\tilde\sigma) \cdot \cdots$, which is evidently fixed by $\partial \tilde f$, is an attractor for the action of $\partial\tilde f$, is contained in $\mathrm{Fix}_N(\partial\tilde f)$ and is not the endpoint of an axis of a covering translation. If $Z_{\mathcal T}(\tilde f)$ contains a non-trivial element T then the T-orbit of P would be an infinite set in $\mathrm{Fix}_N(\partial\tilde f)$. This contradiction shows that $Z_{\mathcal T}(\tilde f)$, and hence $\mathrm{Fix}(\Phi)$, is trivial. The remaining subcase is that σ is a Nielsen path. Let $\tilde v$ and $\tilde w$ be the initial and terminal endpoints of $\tilde\sigma$ respectively and let T be the covering translation that carries $\tilde v$ to $\tilde w$. Since $\tilde f_\#(\tilde\sigma)$ is the unique lift of σ with initial endpoint at $\tilde w$, we have $\tilde f(\tilde v) = T(\tilde v) = \tilde w$ and $\tilde f(\tilde w) = T(\tilde w)$. Thus $\tilde f T(\tilde v) = T\tilde f(\tilde v)$ and $T \in Z_{\mathcal T}(\tilde f)$. Equivalently $J_{\tilde v}([\sigma]) \in Z_{\mathcal T}(\tilde f)$. Letting $H < \pi_1(G, v)$ be the maximal cyclic subgroup that contains $[\sigma]$, we have $J_{\tilde v}(H) = Z_{\mathcal T}(\tilde f)$ and the usual argument shows that $\mathrm{Fix}(\Phi)$ is identified with $H^\eta < \pi_1(G, \star)$. In summary,

$$\Phi \rightsquigarrow \tilde f \rightsquigarrow \tilde\sigma \rightsquigarrow \langle[\sigma]\rangle < H < \pi_1(G, v) \rightsquigarrow H^\eta < \pi_1(G, \star). \qquad \square$$

9.4. The general case

Proposition 9.10 (Bolgopolski-Maslakova [BM16]). *There is an algorithm that, given* $\Phi \in \mathsf{Aut}(F_n)$, *computes* $\mathrm{Fix}(\Phi)$.

Proof. The case of $\Phi \in \phi$ with ϕ rotationless was handled in Lemma 9.9.

Suppose then that ϕ is not rotationless. In Corollary 3.14 we computed M so that ϕ^M is rotationless. So quoting Lemma 9.9 again, $\mathrm{Fix}(\Phi^M)$ can be computed algorithmically. We are reduced to finding the fixed subgroup of the periodic action of Φ on $\mathrm{Fix}(\Phi^M)$. More generally, we are reduced to finding $\mathrm{Fix}(\Phi)$ for a finite order $\Phi \in \mathsf{Aut}(F_n)$, and this is the content of Lemma 9.2. $\square$

10. $S(f)$ and $\mathrm{Fix}(\phi)$

Let $f : G \to G$ be a CT representing $\phi \in \mathsf{Out}(F_n)$. In this section we characterize the components of $S(f)$ and show that $S(f)$ is the geometric realization of $\mathrm{Fix}(\phi)$ defined as follows.

Definition 10.1. For $\phi \in \mathsf{Out}(F_n)$, $\mathrm{Fix}(\phi)$ is defined to be the collection of $[\mathrm{Fix}(\Phi)]$ indexed by isogredience classes $[\Phi] \in \mathrm{P}(\phi)/\sim$.

Recall two facts from the review in Section 3.2.

- If $\tilde{f}$ is a principal lift of f then $\mathrm{Fix}(\tilde{f})$ is non-empty.
- If $\tilde{f}_1, \tilde{f}_2$ are principal lifts of f then $\tilde{f}_1$ and $\tilde{f}_2$ are isogredient if and only if $\mathrm{Fix}(\tilde{f}_1)$ and $\mathrm{Fix}(\tilde{f}_2)$ have equal projections in G.

By the definition of a CT, endpoints of indivisible Nielsen paths are vertices and so Lemma 3.8(1) implies that, for principal $\tilde{f}$, $\mathrm{Fix}(\tilde{f})$ consists of vertices and edges. It follows that there are only finitely many equivalence classes of principal lifts of f and that there is a 1–1 correspondence between isogredience classes $[\tilde{f}]$ of principal lifts $\tilde{f}$ of f and Nielsen classes $[v]$ of principal vertices v in G given by $[\tilde{f}] \leftrightarrow [v]$if and only if $\tilde{f}$ fixes some lift of v. It is algorithmic to tell if a vertex of G is principal and to find its Nielsen class (Lemmas 3.8(1) and 9.6).

By construction, for principal vertices v and v', $S(f, v) = S(f, v')$if and only if $[v] = [v']$. We write $S(f, v) = S(f, [v]) = S([\tilde{f}])$ where $[v] \leftrightarrow [\tilde{f}]$ and see that $S(f) = \bigsqcup_{[v]} S(f, [v]) = \bigsqcup_{[\tilde{f}]} S([\tilde{f}])$ where $[v]$ runs over Nielsen classes of principal vertices in G and $[\tilde{f}]$ runs over isogredience classes of principal lifts of f.

If $\tilde{f} \leftrightarrow \Phi$ and $\tilde{f}' \leftrightarrow \Phi'$ are isogredient then $\mathrm{Fix}(\Phi)$ and $\mathrm{Fix}(\Phi')$ are conjugate and $Z_{\mathcal{T}}(\tilde{f})$ and $Z_{\mathcal{T}}(\tilde{f}')$ are conjugate. Hence, the isogredience classes of Φ and $\tilde{f}$ determine a conjugacy class $\mathrm{Fix}([\Phi])$ of subgroup of F_n and a conjugacy class $Z_{\mathcal{T}}([\tilde{f}])$ of subgroup of $\mathcal{T}(\widetilde{G})$. The sets $\mathrm{Fix}([\Phi])$ and $Z_{\mathcal{T}}([\tilde{f}])$ correspond under the identifications of Section 3.1; we denote this by $\mathrm{Fix}([\Phi]) \leftrightarrow Z_{\mathcal{T}}([\tilde{f}])$. If $[\tilde{f}] \leftrightarrow [v]$ and $\tilde{f} \leftrightarrow \Phi$ then $S(f, [v])$ is the geometric realization of $\mathrm{Fix}([\Phi])$ and $S(f)$ is the geometric realization of $\mathrm{Fix}(\phi)$.

11. Possibilites for [Fix(Φ)]

In this section we algorithmically find the possibilities for $[\mathrm{Fix}(\Phi)] = \mathrm{Fix}([\Phi])$ for $\Phi \in \phi$ with ϕ rotationless (Corollary 11.1).

Corollary 11.1. *Let $\phi \in \mathsf{Out}(F_n)$ be rotationless.*

(1) *There are finitely many F_n-conjugacy classes in*

$$\{\mathrm{Fix}(\Phi) \mid \Phi \in \mathsf{Aut}(F_n) \text{ is principal and represents } \phi\}$$

These conjugacy classes are represented by the components of $S(f)$ where $f : G \to G$ is a CT *for ϕ. In particular, they can be computed algorithmically.*

(2) *For all $\Phi \in \phi$,* $\mathrm{Fix}(\Phi)$ *is root-free. Conversely, if $w \neq 1 \in F_n$ is root-free and $[w]$ is ϕ-invariant then*

  (a) *there is $\Phi \in \mathrm{P}(\phi)$ such that $w \in \mathrm{Fix}(\Phi)$ and*

  (b) *there is non-principal $\Phi' \in \phi$ such that* $\mathrm{Fix}(\Phi') = \langle w \rangle$.

Proof. (1) This is the content of Section 10.

(2) $\mathrm{Fix}(\Phi)$ is always root-free. The existence of Φ as in (a) follows from [FH11, Lemma 3.30 (1)]. The automorphism $\Phi_m := i_w^m \circ \Phi$ fixes w for all m and is principal for only finitely many $m \in \mathbb{Z}$; see [FH11, Lemma 4.40]. Take Φ' to be a non-principal Φ_m. Since Φ' is non-principal, $\mathrm{rank}(\mathrm{Fix}(\Phi')) < 2$ and we see that $\mathrm{Fix}(\Phi') = \langle w \rangle$. □

Lemma 11.3 below seems obvious but we could not find a reference for it in the literature so we are including it here with a proof for completeness. The following lemma is used in its proof.

Lemma 11.2. *For all non-trivial ϕ there exists a filling conjugacy class that is not fixed by ϕ.*

Proof. After replacing ϕ by an iterate, we may assume that ϕ is rotationless. Let $f: G \to G$ be a topological representative of ϕ and let $\beta \subset G$ be a circuit representing a conjugacy class $[c]$ that is not ϕ-invariant. Since ϕ is rotationless, $[c]$ is not fixed by any iterate of ϕ. We may therefore choose k so that the edge length of $f_\#^k(\beta)$ is greater than the edge length of β. Choose a path $\alpha \subset G$ such that any circuit containing α is filling (any path α whose Whitehead graph does not have a cut-point will do [Mar95]; see also [BF14, Proof of Lemma 3.2]). Choose a circuit τ that contains both α and β as disjoint subpaths. Decompose τ as a concatenation of subpaths $\tau = \mu\beta$ and let $\tau_N = \mu\beta^N$. The bounded cancellation lemma implies that the edge length of $f_\#^k(\tau_N)$ is greater than the edge length of τ_N for all sufficiently large N. Each such τ_N satisfies the conclusions of the lemma. □

Lemma 11.3. *Each infinite order $\phi \in \mathtt{Out}(F_n)$ is represented by $\Phi \in \mathtt{Aut}(F_n)$ with trivial fixed subgroup.*

Proof. Choose (Lemma 11.2) a filling conjugacy class $[a]$ that is not fixed by $\psi = \phi^l$ where l is chosen so that both ψ and ψ^{-1} are rotationless. Let $\{a^\pm\} \subset \partial F_n$ denote the endpoints of the axis of a. Choose Φ representing ϕ such that $\partial\Phi(\{a^\pm\}) \cap \{a^\pm\} = \emptyset$ and let $\Phi_m = i_a^m \circ \Phi$. We will prove that $\mathrm{Fix}(\Phi_m)$ is trivial for all sufficiently large m.

Assuming that Φ_m fixes some non-trivial b_m for arbitrarily large m, we will argue to a contradiction. For notational convenience we pass to a subsequence and assume that for all $m \geq 1$, b_m is fixed by Φ_m. After passing to a further subsequence we may assume that $b_m^+ \to P$ and $b_m^- \to Q$ for some $P, Q \in \partial F_n$. From $\partial\Phi_m(b_m^+) = b_m^+$ we see that $\partial\Phi(b_m^+) = \partial i_a^{-m}(b_m^+)$ and hence that P is either a^+ or $\partial\Phi^{-1}(a^-)$. Similarly, Q is either a^+ or $\partial\Phi^{-1}(a^-)$.

Let $f: G \to G$ be a CT representing ψ and let $\tilde{f}_m$ be the lift of $f: G \to G$ corresponding to $\Psi_m := \Phi_m^l$. Let T and S_m be the covering translations corresponding to a and b_m respectively. Thus $T^\pm = a^\pm$ and $S_m^\pm = b_m^\pm$. If $P \neq Q$ then the line connecting a^+ to $\partial\Phi^{-1}(a^-)$ is a weak limit of the axes for S_m. It follows that the axis of T is a weak limit of a sequence of Nielsen axes for $\tilde{f}_m$ and hence that the axis of T is an increasing union of arcs each lifting to $S(f)$. The axis of T therefore lifts to $S(f)$ and so represents a ψ-fixed conjugacy class by Lemma 9.6, contradiction.

Suppose then that $P = Q = \Phi^{-1}(T^-)$. (The case that $P = Q = T^+$ is argued symmetrically.) For all sufficiently large m, there is a neighborhood U_m^+ of T^+ in ∂F_n such that $\partial\Phi_m(U_m^+) \subset U_m^+$. By [LL08, Theorem I], U_m^+ contains a non-repelling periodic point for the action of $\partial\Phi_m$. Thus Ψ_m is a principal lift of ψ and U_m^+ contains an element B_m of $\mathrm{Fix}_N(\Psi_m)$. As $m \to \infty$ we may assume that $B_m \to T^+$.

By Lemma 3.10, there exists a lift $\widetilde{E}_m$ of some $E_m \in \mathcal{E}$ such that the lift $\widetilde{R}_{E_m}$ of the eigenray generated by E_m terminates at B_m. We may assume without loss that $E = E_m$ is independent of m. The ray from the initial endpoint $\tilde{x}_m$ of $\widetilde{E}_m$ to b_m^+ factors as a concatenation of Nielsen paths and so lifts into $S(f)$. After passing to yet another subsequence we may assume that $\tilde{x}_m$ converges to a point X. If $X = T^+$ then the axis of T is a weak limit of paths that lift into $S(f)$ and so lifts to $S(f)$. As above, this gives the desired contradiction. We may therefore assume that $X \neq T^+$. In this case the axis of T is a weak limit of lifts of R_E and so the periodic line that it projects to is a weak limit of R_E. We will complete the proof by showing that each such periodic line L is carried by G_{r-1} where H_r is the top stratum of G. There is no loss in assuming that E is an edge in H_r. If E is NEG then [FH11, Lemma 3.26(3)] completes the proof. If $E \subset H_r$ is EG then L is leaf in the attracting lamination Λ_r associated to H_r by [FH11, Lemma 3.26(2)]. But every leaf is either contained in G_{r-1} or is dense in Λ_r by [BFH00, Lemma 3.1.15] and the latter is impossible for a periodic line. □

12. A Stallings graph for $\mathrm{Fix}_N(\partial\Phi)$

As usual, throughout Section 12 $f: G \to G$ will denote a CT for ϕ. The goal of this section is to generalize Section 10 by describing a (not necessarily finite) G-graph $S_N(f)$ with core $S(f)$ that in a sense described below represents the collection of $\mathrm{Fix}_N(\partial\Phi)$ (Definition 3.3) indexed by $\Phi \in \mathrm{P}(\phi)$. This collection plays an important role in the main theorem of [FH11]. We will first define the graph $S_N(f)$ and then relate it to the collection of $\mathrm{Fix}_N(\partial\Phi)$. $\mathrm{Fix}_N(\partial\Phi)$ has been of considerable interest; see for example [Nie86, Nie29, IKT90, Coo87, GJLL98, BFH97].

12.1. Definition of $S_N(f)$. The idea of the construction of $S_N(f)$ is to start with $S(f)$ and add a copy of R_E for each $E \in \mathcal{E}_f$. This does not quite work since it is possible that $\partial R_E = \partial R_{E'}$ for $E \neq E'$. In the next paragraph we describe the ways that this can happen.

Suppose that $E \neq E' \in \mathcal{E}_f$ and that $\partial R_E = \partial R_{E'}$. Choose a lift $\widetilde{E}$ of E to $\widetilde{G}$, let $\widetilde{R}_E$ be the lift of R_E that begins at the initial endpoint $\tilde{v}$ of $\widetilde{E}$ and let P be the terminal endpoint of $\widetilde{R}_E$. By hypothesis, there is a lift $\widetilde{R}_{E'}$ of $R_{E'}$ that terminates at P. Let $\widetilde{E}'$ be the initial edge of $\widetilde{R}_{E'}$ and let $\tilde{w}$ be the initial vertex of $\widetilde{E}'$. If $\tilde{f}$ is the lift that fixes $\tilde{v}$ then $\partial\tilde{f}$ fixes P. Since $\tilde{f}$ is the only lift of f that fixes P (Remark 3.11), it follows that $\tilde{f}$ also fixes $\tilde{w}$. The path $\tilde{\rho}$ connecting $\tilde{v}$ to $\tilde{w}$ is therefore a Nielsen path for $\tilde{f}$ and its image $\rho \subset G$ is a Nielsen path for f. The edges E and E' are of the same EG height, say r, since the height of E (resp. E') is the height of R_E (resp. $R_{E'}$) and edges of this height occur infinitely often in R_E (resp. $R_{E'}$) and R_E and R'_E have a common tail. It follows from the facts in Review 9.3 that ρ is a concatenation of subpaths that are fixed edges or indivisible Nielsen paths, that E (resp. E') must be contained in an indivisible Nielsen subpath of height r, and that each of these indivisible Nielsen subpaths of height r contributes an illegal turn of height r. By construction, ρ has at most one illegal turn of height r and so we conclude that ρ is indivisible. Note that $\widetilde{R}_E$ and $\widetilde{R}_{E'}$ have a common terminal ray $\widetilde{R}_{E,E'}$ and that $\tilde{\rho} = \tilde{\alpha}\tilde{\beta}^{-1}$ where $\widetilde{R}_E = \tilde{\alpha}\widetilde{R}_{E,E'}$ and $\widetilde{R}_{E'} = \tilde{\beta}\widetilde{R}_{E,E'}$ giving ρ the form $\alpha\beta^{-1}$ as in Review 9.3.

Definition 12.1. We use the notation from the Definitions 9.4 and 9.7. Let $S_N(f)$ be the graph obtained from $S(f)$ as follows. For each NEG $E \in \mathcal{E}_f$ and for each EG $E \in \mathcal{E}_f$ that is not the initial edge of an indivisible Nielsen path, attach R_E to $S(f)$ by identifying the initial endpoint of R_E with the initial endpoint of E thought of as a vertex in $S_1(f)$. If $E \in \mathcal{E}_f$ belongs to an EG stratum H_r and E is the initial edge of an indivisible Nielsen path ρ of height r then the initial edge E' of ρ^{-1} is also an edge of $\mathcal{E}_f \cap H_r$. Subdivide the edge labeled ρ that was added to $S(f)$ during stage two so that it is now two edges, one labeled α and the other β^{-1}. Attach $R_{E,E'}$ to $S(f)$ at the newly created vertex.

Remark 12.2. Given that the construction of $S(f)$ is algorithmic and that any initial segment of a ray R_E of prescribed length can be explicitly computed, it follows that there is an algorithm that, given $d > 0$, constructs the d-neighborhood (in the graph metric) of $S(f)$ in $S_N(f)$.

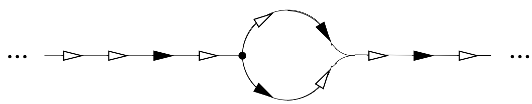

Figure 4. $S_N(f)$ for f as in Figure 2. $S_N(f)$ is obtained from $S(f)$ by adding the rays $R_{a,b} = babbab\ldots$ and $R_B = BABBABAB\ldots$.

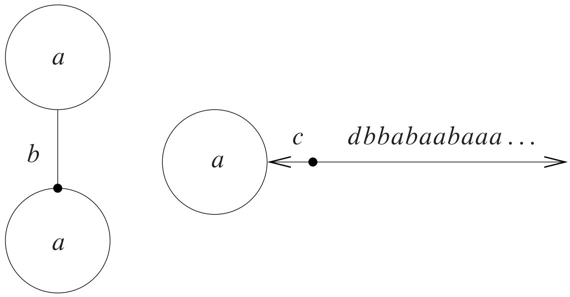

Figure 5. $S_N(h)$ for h as in Figure 3. $S_N(h)$ is obtained from $S(h)$ by adding the eigenray $dbb\ldots$.

By construction of $S_N(f)$, its components are in a 1-1 correspondence with the components of $S(f)$. That is, components of $S_N(f)$ are in a bijective correspondence with Nielsen classes $[v]$ of principal vertices v in G. Let $S_N(f,[v])$ denote the component of $S_N(f)$ containing v. We have $S_N(f) = \bigsqcup_{[v]} S_N(f,[v])$ where $[v]$ runs over Nielsen classes of principal vertices in G. We continue to identify the principal vertices of G with the vertices of $S_1(f)$. We call these vertices the *principal vertices of* $S_N(f)$.

12.2. Properties of $S_N(f)$. In this section we record some properties of $S_N(f)$. We defined the term *core* of a Σ-graph H in Section 9.2. The *weak core of* H is the union of all properly immersed lines in H. H is *weakly core* if it is its own weak core. We say that H has *finite type* if it is the union of a finite graph and finitely many rays.

Lemma 12.3. *$S_N(f)$ is a G-graph of finite type and all vertices have valence ≥ 2. In particular, $S_N(f)$ is weakly core.*

Proof. By construction, the labeling map $S_N(f) \to G$ is an immersion and $S_N(f)$ has finite type and its non-principal vertices have valence either two or three. By [FH11, Lemma 4.14], a principal vertex v has at least two fixed directions in G, each corresponding to a fixed edge, an edge in $\mathrm{Lin}(f)$, or an edge in $\mathcal{E}_f$. By construction, each of these edges contributes a direction at v in $S_N(f)$. Thus all principal vertices, and hence all vertices, have valence at least two. $\square$

Lemma 12.4. *Let $\tilde{f}: \widetilde{G} \to \widetilde{G}$ be principal and fix the vertex $\tilde{v} \in \widetilde{G}$. The labelling map $S_N(f,v) \to G$ induces an embedding $\widetilde{S} \to \widetilde{G}$ of universal covers. The induced map $\partial\widetilde{S} \to \partial\widetilde{G}$ has image* $\mathrm{Fix}_N(\partial \tilde{f})$.

Proof. The first conclusion follows since the labelling map is an immersion. To prove the second, note that every ray R with initial vertex v is either contained in $S(f,v)$ in which case the corresponding lift $\widetilde{R}$ has endpoint in $\partial(\mathrm{Fix}(\tilde{f}))$ or else is the concatenation of a Nielsen path and a ray R_E in which case $\tilde{R}$ has endpoint in $\mathrm{Fix}_+(\partial \tilde{f})$. Conversely, if $P \in \mathrm{Fix}_N(\partial \tilde{f})$ then the ray $[\tilde{v}, P)$ is a concatenation of a Nielsen path and the lift of some R_E by [FH11, Lemma 4.36(2)]. $\square$

Corollary 12.5. *In the notation of Lemma 12.4, $\tilde{S}$ is the convex hull of $\mathrm{Fix}_N(\partial \tilde{f})$, i.e. the union of all lines connected distinct points in $\mathrm{Fix}_N(\partial \tilde{f})$. In particular, if $n > 1$ then no component of $S_N(f)$ is a circle, an axis in $\tilde{G}$, or a generic leaf of an attracting lamination of f.*

Remark 12.6. In the main theorem of [FH11], a rotationless outer automorphism ϕ is characterized in terms of two invariants: one qualitative and the other quantitative. The collection of $\mathrm{Fix}_N(\partial\Phi)$ indexed by $\Phi \in \mathrm{P}(\phi)$ is the qualitative invariant. From the results of this section, we see that $S_N(f)$ represents this qualitative invariant in the following sense. If X is a component of $S_N(f)$ and $\tilde{X} \to \tilde{G}$ is a lift of the natural immersion $X \to G$, then the image of the induced map $\partial\tilde{X} \to \partial F_n$ is $\mathrm{Fix}_N(\partial\Phi)$ for some $\Phi \in \mathrm{P}(\phi)$. Conversely, if $\Phi \in \mathrm{P}(\phi)$ then $\mathrm{Fix}_N(\partial\Phi)$ is the image of $\partial\tilde{X} \to \partial F_n$ for some component X of $S_N(f)$ and some lift $\tilde{X} \to \tilde{G}$.

13. Moving up through the filtration

Several of our applications are verified by working up through the filtration of a CT $f: G \to G$. In this section we establish notation by recalling some notation and a lemma from [FH09, Notation 8.2 and Lemma 8.3]. The negative of the Euler characteristic will be important to us; we write χ^- for $-\chi$.

Notation 13.1. Recall from (Filtration) that the core of each filtration element is a filtration element. The core filtration

$$\emptyset = G_0 = G_{l_0} \subset G_{l_1} \subset G_{l_2} \subset \cdots \subset G_{l_K} = G_N = G$$

is defined to be the coarsening of the full filtration obtained by restricting to those elements that are their own cores or equivalently have no valence one vertices. Note that $l_1 = 1$ by (Periodic edges). For each G_{l_i}, let $H^c_{l_i}$ be the i-th stratum of the core filtration. Namely

$$H^c_{l_i} = \bigcup_{j=l_{i-1}+1}^{l_i} H_j.$$

The change in negative Euler characteristic is denoted

$$\Delta_i \chi^- := \chi^-(G_{l_i}) - \chi^-(G_{l_{i-1}}).$$

Referring to Notation 2.11, if H_{l_i} is EG then H_{u_i} denotes the highest irreducible stratum in $G_{l_i - 1}$.

Lemma 13.2 ([FH09, Lemma 8.3]). (1) *If $H^c_{l_i}$ does not contain any EG strata then one of the following holds.*

(a) $l_i = l_{i-1} + 1$ *and the unique edge in $H^c_{l_i}$ is a fixed loop that is disjoint from $G_{l_{i-1}}$.*

(b) $l_i = l_{i-1} + 1$ *and both endpoints of the unique edge in $H^c_{l_i}$ are contained in $G_{l_{i-1}}$.*

(c) $l_i = l_{i_1} + 2$ *and the two edges in $H^c_{l_i}$ are nonfixed and have a common initial endpoint that is not in $H_{l_{i-1}}$ and terminal endpoints in $G_{l_{i-1}}$.*

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

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

(a) *For $l_{i_1} < j \le u_i$, H_j is a single nonfixed edge E_j whose terminal vertex is in $G_{l_{i-1}}$ and whose initial vertex has valence one in G_{u_i}. In particular, G_{u_i} deformation retracts to $G_{l_{i-1}}$ and $\chi(G_{u_i}) = \chi(G_{l_{i-1}})$.*

(b) *For $u_i < j < l_i$, H_j is a zero stratum. In other words, the closure of $G_{l_i} \setminus G_{u_i}$ is the extended EG stratum $H^z_{l_i}$.*

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

14. Primitively atoroidal outer automorphisms

In this section we exhibit an algorithm to determine whether or not a $\phi \in \mathtt{Out}(F_n)$ is primitively atoroidal (Corollary 14.4).

Definition 14.1. A conjugacy class in F_n is *primitive* if it is represented by an element in some basis of F_n. An outer automorphism ϕ is *primitively atoroidal* if there does not exist a periodic conjugacy class for ϕ which is primitive.

Lemma 14.2. *Let $f : G \to G$ be a* CT. *Either some stratum is a fixed loop or all closed Nielsen paths are trivial in $H_1(G; \mathbb{Z}/2\mathbb{Z})$.*

Proof. We use the notation of Section 13. Suppose the statement holds for $G_{l_{l_i-1}}$. To get a contradiction, suppose the statement fails for G_{l_i}, that is suppose that $H^c_{l_i}$ is not a fixed loop and that some closed Nielsen path μ of height l_i is non-trivial in $H_1(G_{l_i}; \mathbb{Z}/2\mathbb{Z})$. We go through the cases of Lemma 13.2.

Recall that Nielsen paths are completely split and their complete splitting consists of fixed edges and indivisible Nielsen paths.

In (1a), H_{l_i} is a fixed loop and we are assuming this is not the case.

In (1b), by the (NEG Nielsen paths) property of a CT $E_{l_i} := H_{l_i}$ is either fixed or linear and if E_{l_i} is linear then μ or μ^{-1} has a term of the form $E_{l_i} w_{l_i}^k \bar{E}_{l_i}$ with notation as in (NEG Nielsen paths).

Suppose first that E_{l_i} is fixed.

- If μ has only one occurrence of E_{l_i} and no occurrences of its inverse (or the symmetric case) then μ is primitive. In this case, the union $\mathcal{F}$ of the free factor system $[G_{l_{i-1}}]$ with the conjugacy class of the cyclic subgroup generated by μ is a ϕ-invariant free factor system properly containing $[G_{l_{i-1}}]$ and contained in $[G_{l_i}]$; see Example 4.11. The (Filtration) property of a CT implies that $\mathcal{F} = [G_{l_i}]$ contradicting our assumption that $H_{l_i}^c$ is not a fixed loop.
- Consider the (cyclic) sequence of $E_{l_i}^{\pm 1}$'s that occurs in μ and suppose there are consecutive occurrences with the same orientation, say $\mu = \dots E_{l_i} x E_{l_i} \dots$. Then $\mu' = E_{l_i} x$ is a closed Nielsen path as in the preceding bullet, contradiction.
- The remaining case is that μ or μ^{-1} has the form $E_{l_i} x_1 E_{l_i}^{-1} x_2 E_{l_1} \dots E_{l_i}^{-1} x_N$, i.e. the orientations of the E_{l_i}'s alternate. Each x_i is a closed Nielsen path and by induction is trivial in $H_1(G; \mathbb{Z}/2\mathbb{Z})$. If follows that μ is also trivial in $H_1(G; \mathbb{Z}/2\mathbb{Z})$

Next suppose that E_{l_i} is linear and $f(E_{l_i}) = E_{l_i} \cdot w_{l_i}^{d_{l_i}}$ (with notation as in (Linear edges)). In particular, w_{l_i} is a closed Nielsen path of height $< l_i$ and so is trivial in $H_1(G; \mathbb{Z}/2\mathbb{Z})$. By (NEG Nielsen paths), μ a product of paths of the form $E_{l_i} u^k E_{l_k}^{-1}$ and closed Nielsen paths of height $< l_i$. In particular, μ is trivial in $H_1(G; \mathbb{Z}/2\mathbb{Z})$, contradiction. This completes the proof of (1b).

In (1c), by (NEG Nielsen paths) $H_{l_i}^c$ consists of two linear edges. This case is similar to the case that $H_{l_i}^c$ consists of one linear edge. We conclude μ is trivial in $H_1(G; \mathbb{Z}/2\mathbb{Z})$, again a contradiction.

In (2), by [FH11, Corollary 4.19 and Remark 4.20] either there are no closed Nielsen paths of height l_i or there is a Nielsen loop μ_{l_i} of height l_i such that every Nielsen loop of this height is a power of μ_{l_i} and further μ_{l_i} is trivial in $H_1(G; \mathbb{Z}/2\mathbb{Z})$. This is a contradiction. □

Corollary 14.3. *Let ϕ be rotationless. Either*

- *for some $\Phi \in \mathrm{P}(\phi)$, there is an element of* $\mathrm{Fix}(\Phi)$ *that is primitive in F_n; or*
- *for all $\Phi \in \mathrm{P}(\phi)$, every element of* $\mathrm{Fix}(\Phi)$ *is trivial in $H_1(F_n; \mathbb{Z}/2\mathbb{Z})$.*

There are more general purpose algorithms that check whether a finitely generated subgroup of F_n contains a primitive element; see [CG10, Dic14]. We thank the referee for pointing out these references to us.

Corollary 14.4. *There is an algorithm to tell whether or not a given $\phi \in \mathsf{Out}(F_n)$ is primitively atoroidal.*

Proof. Compute a CT $f : G \to G$ for a rotationless power of ϕ. According to Lemma 14.2, ϕ is primitively atoroidal if and only if some stratum of G is a fixed circle. $\square$

15. The index of an outer automorphism

Definition 15.1 ([GJLL98]). For $\Phi \in \mathsf{Aut}(F_n)$,

$$i(\Phi) := \max\left\{0, \operatorname{rank}(\operatorname{Fix}(\Phi)) + \frac{1}{2}a(\Phi) - 1\right\}$$

where $a(\Phi)$ is the number of $\operatorname{Fix}(\Phi)$-orbits of attracting fixed points of $\partial\Phi$. For $\phi \in \mathsf{Out}(F_n)$, $i(\phi) := \sum i(\Phi)$ where the sum ranges over representatives of isogredience classes of ϕ.

It is clear that $i(\Phi) > 0$ implies Φ is principal. In particular

$$i(\phi) := \sum \left(\operatorname{rank}(\operatorname{Fix}(\Phi)) + \frac{1}{2}a(\Phi) - 1\right)$$

where the sum $\sum i(\Phi)$ is over principal representatives Φ of ϕ only.

Proposition 15.2. *There is an algorithm to compute the index of rotationless $\phi \in \mathsf{Out}(F_n)$.*

See also [GJLL98, Section 6.1].

Proof. Using Theorem 1.1, construct a CT $f : G \to G$ for ϕ. Let $S_N(f) = \bigsqcup_{[v]} S_N(f, [v])$ be the graphs constructed from f in Section 12. If Φ is in the isogredience class of the principal representative of ϕ determined by the principal vertex $v \in G$ then $i(\Phi) = i(S_N(f, [v]))$ where the the index i of a finite type connected graph Γ is

$$i(\Gamma) := \max\left\{0, \operatorname{rank}(\Gamma) + \frac{1}{2}a(\Gamma) - 1\right\}$$

where $a(\Gamma)$ is the number of ends of Γ. $\square$

The index satisfies the following well-known inequality.

Theorem 15.3 ([GJLL98]). *For $\phi \in \mathsf{Out}(F_n)$, $i(\phi) \leq n - 1$.*

We take this opportunity to use CTs to somewhat strengthen this inequality; see Proposition 15.14. See Section 3.3 to recall notation. Also see Lemma 3.10.

From the CT point of view, the set of attracting fixed points of $\partial\Phi$ can be partitioned into those coming from EG strata and those coming from NEG strata. The following lemma states that this is independent of the choice of CT representing ϕ.

Lemma 15.4. [HM, Part I Definitions 2.9 and 2.10 and Lemma 2.11] *Suppose that ϕ is rotationless, that $\Phi \in \mathrm{P}(\phi)$ and that $P \in \mathrm{Fix}_+(\partial\Phi)$. Then the following are equivalent.*

(1) *For some* CT $f: G \to G$ *representing ϕ there is a non-linear NEG edge E with a lift $\tilde{E}$ to $\tilde{G}$ so that $\tilde{R}_{\tilde{E}}$ converges to P.*

(2) *For every* CT $f: G \to G$ *representing ϕ there is a non-linear NEG edge E with a lift $\tilde{E}$ to $\tilde{G}$ so that $\tilde{R}_{\tilde{E}}$ converges to P.*

Following Lemma 15.4 we make the following definition, with justification given in Remark 15.6.

Definition 15.5. Let $\Phi \in \mathrm{P}(\phi)$ with $\phi \in \mathtt{Out}(F_n)$ and $P \in \mathrm{Fix}_+(\partial\Phi)$. If ϕ is rotationless then P is an NEG-*ray for* Φ if the equivalent conditions of Lemma 15.4 are satisfied. If ϕ is not necessarily rotationless then P is an NEG-*ray for* Φ if P is an NEG-ray for Φ^K where ϕ^K is a rotationless iterate of ϕ. Let $\mathcal{R}_{\mathrm{NEG}}(\Phi)$ denote $\{R \mid R \text{ is an NEG-ray for } \Phi\}$, let $\mathcal{R}_{\mathrm{NEG}}(\phi)$ denote $\bigcup_{\Phi\in\mathrm{P}(\phi)} \mathcal{R}_{\mathrm{NEG}}(\Phi)$ and let $\mathcal{R}(\phi)$ denote $\bigcup_{\Phi\in\mathrm{P}(\phi)} \mathrm{Fix}_+(\Phi)$.

Remark 15.6. To see that Definition 15.5 is well defined (i.e. independent of the choice of K), note that if ϕ is rotationless and f is a CT representing ϕ then f^k is a CT representing ϕ^k for all $k \geq 1$ (Lemma 5.5). It follows that P is an NEG-ray for Φ if and only if it is a NEG-ray for each Φ^k. If ϕ is not rotationless but has rotationless iterates ϕ^K and ϕ^L then P is an NEG-ray for ϕ^K if and only if it is an NEG-ray for ϕ^{KL} if and only if it is an NEG-ray for ϕ^L.

Definition 15.7. Suppose $\phi \in \mathtt{Out}(F_n)$ and $\Phi \in \mathrm{P}(\phi)$. Define

$$j(\Phi) := i(\Phi) + \frac{1}{2}b(\Phi),$$

where $b(\Phi)$ is the number of $\mathrm{Fix}(\Phi)$-orbits of NEG-rays of Φ and define $j(\phi) := \sum j(\Phi)$ where the sum is over representatives of isogredience classes of principal representatives of ϕ.

Lemma 15.8. *Let* $\Phi, \Psi \in \mathrm{P}(\phi)$.

(1) *If* $\mathrm{Fix}_+(\Phi) \cap \mathrm{Fix}_+(\Psi) \neq \emptyset$ *then* $\Phi = \Psi$. *In particular,*

$$\mathcal{R}(\phi) = \bigsqcup_{\Phi \in \mathrm{P}(\phi)} \mathrm{Fix}_+(\Phi)$$

and

$$\mathcal{R}_{\mathrm{NEG}}(\phi) = \bigsqcup_{\Phi \in \mathrm{P}(\phi)} \mathcal{R}_{\mathrm{NEG}}(\Phi).$$

(2) *The stabilizer of* $\mathrm{Fix}_+(\Phi)$ [*resp.* $\mathcal{R}_{\mathrm{NEG}}(\Phi)$] *under the action of* F_n *on* $\mathcal{R}(\phi)$ [*resp.* $\mathcal{R}_{\mathrm{NEG}}(\phi)$] *is* $\mathrm{Fix}(\Phi)$.

(3) *The natural maps*

$$\bigsqcup_{i=1}^{N} \mathrm{Fix}_+(\Phi_i)/\mathrm{Fix}(\Phi_i) \longrightarrow \mathcal{R}(\phi)/F_n$$

and

$$\bigsqcup_{i=1}^{N} \mathcal{R}_{\mathrm{NEG}}(\Phi_i)/\mathrm{Fix}(\Phi_i) \longrightarrow \mathcal{R}_{\mathrm{NEG}}(\phi)/F_n$$

are bijective where $\{\Phi_i \mid i = 1, \dots, N\}$ *is a set of representatives of isogredience classes in* $\mathrm{P}(\phi)$.

Proof. (1) Since Φ and Ψ both represent ϕ, $\Phi\Psi^{-1} = i_a$ for some $a \in F_n$. If $R \in \mathrm{Fix}_+(\Phi) \cap \mathrm{Fix}_+(\Psi)$ then $aR = R$. Since R is not an endpoint of the axis of a, $a = 1$. (Here we are using the shorthand notation aR for $\partial i_a(R)$.)

(2) Suppose $a \in \mathrm{Fix}(\Phi)$ and $R \in \mathcal{R}_{\mathrm{NEG}}(\Phi)$. Then $\partial\Phi(aR) = \Phi(a)\partial\Phi(R) = aR$ and so $aR \in \mathcal{R}_{\mathrm{NEG}}(\Phi)$. Conversely, if $a \in F_n$, $R \in \mathcal{R}_{\mathrm{NEG}}(\Phi)$, and $\partial\Phi(aR) = aR$ then $aR = \partial\Phi(aR) = \Phi(a)\partial\Phi(R) = \Phi(a)R$. We conclude $a = \Phi(a)$ as in the proof of (1). The same argument applies with $\mathcal{R}_{\mathrm{NEG}}(\Phi)$ replaced by $\mathrm{Fix}_+(\Phi)$.

(3) This is a consequence of (1), (2), and the observation that F_n acts on $\mathrm{Fix}_+(\Phi)$ [resp. $\mathcal{R}_{\mathrm{NEG}}(\phi)$] by permuting the $\mathrm{Fix}_+(\Phi)$'s [resp. $\mathcal{R}_{\mathrm{NEG}}(\Phi)$'s]. $\square$

Lemma 15.9. (1) For $k > 0$ and $\phi \in \mathsf{Out}(F_n)$, $b(\phi^k) \geq b(\phi)$.

(2) For $k > 0$ and $\phi \in \mathsf{Out}(F_n)$, $a(\phi^k) \geq a(\phi)$.

Proof. (1) By definition of $b(\phi)$ and Lemma 15.8(3),

$$b(\phi) = \Big| \bigsqcup_{i=1}^{N} \mathcal{R}_{\mathrm{NEG}}(\Phi_i)/\mathrm{Fix}(\Phi_i) \Big| = |\mathcal{R}_{\mathrm{NEG}}(\phi)/F_n|$$

Also by definition, if R is an NEG-ray for $\Psi \in \mathrm{P}(\phi)$ then R is an NEG-ray for $\Psi^k \in \mathrm{P}(\phi^k)$ and so $\mathcal{R}_{\mathrm{NEG}}(\phi^k) \supseteq \mathcal{R}_{\mathrm{NEG}}(\phi)$. Hence

$$b(\phi^k) = |\mathcal{R}_{\mathrm{NEG}}(\phi^k)/F_n| \geq |\mathcal{R}_{\mathrm{NEG}}(\phi)/F_n| = b(\phi).$$

(2) The proof is the same as that of (1), replacing $\mathcal{R}_{\mathrm{NEG}}(\Phi)$ with $\mathrm{Fix}_+(\Phi)$ and $\mathcal{R}_{\mathrm{NEG}}(\phi)$ with $\mathcal{R}(\phi)$. $\square$

Notation 15.10. For $H < F_n$, $\hat{r}(H) := \max(0, \mathrm{rank}(H) - 1)$. For $\Phi \in \mathsf{Aut}(F_n)$, $\hat{r}(\Phi) := \hat{r}(\mathrm{Fix}(\Phi))$. For $\phi \in \mathsf{Out}(F_n)$, $\hat{r}(\phi) := \sum \hat{r}(\Phi)$ where the sum is over representatives of $\mathrm{P}(\phi)/\sim$.

Lemma 15.11. *Let $\phi \in \mathsf{Out}(F_n)$ and $k > 0$ such that $\psi := \phi^k$ is rotationless. Then $\hat{r}(\psi) \geq \hat{r}(\phi)$.*

Proof. Assuming without loss that $\hat{r}(\phi) > 0$, let $\Phi_1, \ldots, \Phi_s, \ldots, \Phi_m$ be representatives of $\mathrm{P}(\phi)/\sim$ where $\mathrm{Fix}(\Phi_i) \geq 2$ if and only if $i \leq s$. Let $\{\Psi_j\}$ be representatives of $\mathrm{P}(\psi)/\sim$. For each $1 \leq i \leq s$ there exists $j = p(i)$ such that $\Phi_i^k \sim \Psi_j$. If $j = p(i)$ then there exists $a \in F_n$ such that $\Psi_j = i_a \Phi_i^k i_{a^{-1}}$. Replacing Φ_i by $i_a \Phi i_{a^{-1}}$, we have $\Phi_i^k = \Psi_j$. Thus

$$\Phi_i^k = \Psi_{p(i)}$$

for $1 \leq i \leq s$. We may assume the Ψ_j's are ordered so that the function p, whose domain is $\{1, \ldots, s\}$, has image $\{1, \ldots, t\}$.

It suffices to show that

$$\sum_{j=1}^{t} \hat{r}(\Psi_j) \geq \sum_{i=1}^{s} \hat{r}(\Phi_i)$$

and so it also suffices to show that

$$\hat{r}(\Psi_j) \geq \sum_{i \in p^{-1}(j)} \hat{r}(\Phi_i)$$

for each $1 \leq j \leq t$.

Fix j and let $\mathbb{F} = \mathrm{Fix}(\Psi_j)$. For each $i \in p^{-1}(j)$, $\Phi_i^k = \Psi_j$ and so $\mathbb{F} = \mathrm{Fix}(\Phi_i^k)$. Thus, Φ_i preserves $\mathbb{F}$, $\mathrm{Fix}(\Phi_i) \subset \mathbb{F}$ and the restriction $\Phi_i|\mathbb{F}$ is a finite order automorphism of $\mathbb{F}$. Since $\mathbb{F}$ is its own normalizer (see the second bullet in the proof of Lemma 3.12) the restriction $\phi|\mathbb{F} \in \mathsf{Out}(\mathbb{F})$ is well-defined and has finite order.

We claim that if $i' \in p^{-1}(j)$ and $i' \neq i$ then the subgroups $\mathrm{Fix}(\Phi_i)$ and $\mathrm{Fix}(\Phi_{i'})$ are not conjugate in F_n and hence not conjugate in $\mathbb{F}$. Indeed, if $\mathrm{Fix}(\Phi_i) = h\,\mathrm{Fix}(\Phi_{i'})h^{-1}$ for some $h \in F_n$ then $\mathrm{Fix}(\Phi_i) = \mathrm{Fix}(i_h \Phi_{i'} i_{h^{-1}})$. Since these groups have rank at least two and both Φ_i and $i_h \Phi_{i'} i_{h^{-1}}$ represent ϕ, $\Phi_i = i_h \Phi_{i'} i_{h^{-1}}$, in contradiction to the assumption that Φ_i and $\Phi_{i'}$ represent distinct isogredience classes.

By Culler's Theorem 3.1 of [Cul84], applied to $\phi|\mathbb{F}$, the conjugacy classes in $\mathbb{F}$ determined by the subgroups $\{\mathrm{Fix}(\Phi_i) \colon i \in p^{-1}(j)\}$ form a free factor system of $\mathbb{F}$. Since $\hat{r}(\Psi_j) = \mathrm{rank}(\mathbb{F}) - 1$ and $\hat{r}(\Phi_i) = \mathrm{rank}(\mathrm{Fix}(\Phi_i)) - 1$, the desired displayed inequality holds. □

Lemma 15.12. *Suppose* $\phi \in \mathtt{Out}(F_n)$ *and* $k > 0$ *such that* ϕ^k *is rotationless. Then* $j(\phi^k) \geq j(\phi)$.

Proof. By definition of $j(\phi)$, the lemma is a direct consequence of Lemmas 15.9 and 15.11. □

Proposition 15.13. *There is an algorithm to compute* $j(\phi)$ *for rotationless* $\phi \in \mathtt{Out}(F_n)$.

Proof. This follows from Proposition 15.2 and the fact that $j(\phi)$ is the sum of $i(\phi)$ and one half the number of non-linear NEG edges in any CT representing ϕ. □

Proposition 15.14. *For* $\phi \in \mathtt{Out}(F_n)$, $j(\phi) \leq n - 1$.

Example 15.15. Start with the usual linear 2-rose $a_1 \mapsto a_1$ and $a_2 \mapsto a_2 a_1$. Further attach edges a_i, $2 < i \leq n$, with $a_i \mapsto a_2^{2i} a_i a_2^{2i+1}$. If we subdivide each a_i, $2 < i \leq n$ then the result $f: G \to G$ is a CT for a $\phi \in \mathtt{Out}(F_n)$. Here $S_N(f)$ is the disjoint union of a pair of eyeglasses and $n-2$ lines each with two NEG-rays. Hence $j(\phi) = n - 1$ and $i(\phi) = 1$.

Proof of Proposition 15.14. By Lemma 15.12, we may assume that ϕ is rotationless. Let $f: G \to G$ be a CT for ϕ. We use the terminology of Section 13. In particular,

$$\emptyset = G_0 = G_{l_0} \subset G_{l_1} \subset G_{l_2} \subset \cdots \subset G_{l_K} = G_N = G$$

is the core filtration and $\Delta_i \chi^- := \chi^-(G_{l_i}) - \chi^-(G_{l_{i-1}})$.

Recall that $S_N(f)$ is built up in three stages: starting with principal vertices and fixed edges, lollipops corresponding to linear edges, edges corresponding to EG Nielsen paths, and then rays corresponding (perhaps not bijectively) to oriented edges in $\mathcal{E}$ are added. For each $1 \leq 1 \leq K$, define $S_N(k)$ to be the subgraph of $S_N(f)$ corresponding to G_{l_k}. More precisely, start with principal vertices in G_{l_k} and fixed edges in G_{l_k}, then add lollipops corresponding to linear edges in G_{l_k} and edges corresponding to EG Nielsen paths in G_{l_k} and then add rays corresponding to oriented edges in $\mathcal{E} \cap G_{l_k}$.

For each component Γ of $S_N(k)$ define

$$j(\Gamma) := \operatorname{rank}(\Gamma) + \frac{1}{2}a(\Gamma) + \frac{1}{2}b(\Gamma) - 1$$

where $a(\Gamma)$ is the number of ends of Γ and $b(\Gamma)$ is the number of NEG-rays of Γ. Then define

$$j(k) = \sum_{\Gamma} j(\Gamma)$$

where the sum is taken over all components Γ of $S_N(k)$. Using Remark 12.6, $j(S_N(f)) = j(\phi) = j(K)$. Proposition 15.14 therefore follows from the more general inequality

$$j(k) \le \chi^-(G_{l_k}) \tag{$*_k$}$$

for $0 \le k \le K$, which we will prove by induction on k.

The $k = 0$ case holds because $j(\emptyset) = \chi^-(\emptyset) = 0$.

Suppose that $(*_{k-1})$ holds for some $k \in \{1, \dots, K\}$. Define $\Delta_k j := j(k) - j(k-1)$. We check that, in each of the cases of Lemma 13.2, $\Delta_k j \le \Delta_k \chi^-$. Once this has been done, $(*_k)$ holds and the proposition is proven.

We use the following formula to compute χ^- of a finite graph:

$$\chi^- = \Sigma_v \left(\frac{\mathrm{val}(v)}{2} - 1 \right)$$

where the sum is over the vertices v and $\mathrm{val}(v)$ is the number of directions at v. Thus, $\Delta_k \chi^-$ is the half number of "new" directions (two for each new edge) minus the number of "new" vertices. As we proceed through the process of verifying $\Delta_k j \le \Delta_k \chi^-$, we refer to vertices, directions, or rays of G_{l_k} or $S_N(k)$ previously considered as *old*; others are *new*.

Lemma 3.8 implies that in cases (1a), (1b) and (1c) the endpoints of all new edges are principal.

Case 1a. $S_N(k)$ is obtained from $S_N(k-1)$ by attaching a circle component and hence $\Delta_k j = \Delta_k \chi^- = 0$.

Case 1b. If the edge $H := H_{l_k}$ is fixed then $S_N(k)$ [resp. G_{l_k}] is obtained from $S_N(k-1)$ [resp. $G_{l_{k-1}}$] by adding a new edge to old vertices. Hence $\Delta_k j = \Delta_k \chi^- = 1$. We may therefore assume that the edge H is not fixed and so $f(H) = H \cdot u$ with u non-trivial. In this case, $S_N(k)$ is obtained from $S_N(k-1)$ by attaching to an old vertex either a lollipop (if u is a Nielsen path) or the NEG-ray $H \cdot u \cdot f(u) \cdot f^2(u) \dots$ (otherwise). In each case, $\Delta_k j = \Delta_k \chi^- = 1$.

Case 1c. There are three subcases depending on whether or not the added edges are linear. Their shared initial vertex is a new principal vertex in G_{l_k} and so $S_N(k)$ is obtained from $S_N(k-1)$ by adding a new component. The new component is a pair of eyeglasses if both edges are linear. It is a one point union of an NEG-ray and a lollipop if only one edge is linear, and it is a line with two NEG-rays if neither edge is linear. In all cases, $\Delta_k j = \Delta \chi^- = 1$.

Case 2. In Case 2, we will break the verification that $\Delta_k j \le \Delta_k \chi^-$ into two steps: first passing from $G_{l_{k-1}}$ to G_{u_k} and then passing from G_{u_k} to G_{l_k}.

Step 1. For the edge H_j with j as in Case (2a), the contribution to $S_N(k)$ is a new NEG-ray [or lollipop] with new initial vertex. (The terminal direction in H_j is not fixed and so does not contribute.) Thus, the contribution to $\Delta_k j$ is 0 ($\frac{1}{2}$ to the a-count since the ray is new, another $\frac{1}{2}$ to the b-count since it is NEG [or 1 to rank] and -1 for the new component). This is balanced with the contribution of H_j to $\Delta_k \chi^-$, namely $\frac{1}{2}$ for each of its directions and -1 for the new vertex.

Step 2. For each vertex $v \in H_{l_k}$, let $\Delta_k j(v)$ and $\Delta_k \chi^-(v)$ be the contributions to $\Delta_k j$ and $\Delta_k \chi^-$ coming from v and the H_{l_k}-directions that are incident to v. For each principal vertex $v \in H_{l_k}$ let $\kappa(v)$ be the number of H_{l_k}-directions that are incident to v and not fixed. If $v \in H_{l_k}$ is not principal let $\kappa(v) = \mathrm{val}(v) - 2$. Since each direction we encounter in this step is contained in H_{l_k}, and is hence EG, the b-count does not change.

As a first case assume that there are no indivisible Nielsen paths of height l_k. If v is not principal then $\Delta_k j(v) = 0$. By Lemma 3.8, v is new so the directions incident to v are all in H_{l_k} and $\Delta_k \chi^-(v) = \frac{1}{2}(\mathrm{val}(v) - 2) = \frac{1}{2}\kappa(v)$. Thus $\Delta_k \chi^-(v) - \Delta_k j(v) = \frac{1}{2}\kappa(v) \ge 0$.

We next consider principal v, letting $L(v)$ be the number of fixed H_{l_k}-directions that are based at v. Thus $L(v) + \kappa(v)$ is the number of H_{l_k}-directions that are based at v. If v is old then $\Delta_k j(v) = \frac{1}{2}L(v)$ and $\Delta_k \chi^-(v) = \frac{1}{2}(L(v) + \kappa(v))$. If v is new then $\Delta_k j(v) = \frac{1}{2}L(v) - 1$ and $\Delta_k \chi^-(v) = \frac{1}{2}(L(v) + \kappa(v)) - 1$. As in the previous case, $\Delta_k \chi^-(v) - \Delta_k j(v) = \frac{1}{2}\kappa(v) \ge 0$.

The second and final case is that there is an indivisible Nielsen path ρ of height l_k. By [FH11, Lemma 4.24], $l_k = u_k + 1$. From [BFH00, Lemma 5.1.7] and [FH11, Corollary 4.19] it follows that if the endpoints w, w' of ρ are distinct then at least one of w, w' is new and if $w = w'$ then w is new. Let d and d' be the directions determined by ρ at w and w' respectively.

Let $\mathcal{V}$ be the set of vertices of H_{l_k} that are not endpoints of ρ. Each $v \in \mathcal{V}$ is handled as in the no Nielsen path case. The conclusion is that

$$\sum_{v\in\mathcal{V}} \Delta_k \chi^-(v) - \sum_{v\in\mathcal{V}} \Delta_k j(v) = \sum_{v\in\mathcal{V}} \frac{1}{2}\kappa(v) \ge 0.$$

Let $\Delta_k j(\rho)$ and $\Delta_k \chi^-(\rho)$ be the contributions to $\Delta_k j$ and $\Delta_k \chi^-$ coming from the endpoints of ρ and the H_{l_k}-directions that are incident to the endpoints of ρ. The remaining analysis breaks up into subcases as follows.

- Suppose that ρ is a closed path based at a new vertex w. The change in $CS_N(k)$ corresponding to w is the addition of a new vertex w, an edge representing ρ with both endpoints at w, one ray for each fixed H_{l_k}-direction based at w other than d and d' and then one ray corresponding to d and d'.

(See Section 12.) We therefore have

$$\Delta_k j(\rho) = 1 + \frac{1}{2}(L(w) - 1) - 1 = \frac{1}{2}L(w) - \frac{1}{2}$$

and

$$\Delta_k \chi^-(\rho) = \frac{1}{2}(L(w) + \kappa(w)) - 1.$$

Thus

$$\Delta_k \chi^-(\rho) - \Delta_k j(\rho) = \frac{1}{2}(\kappa(w) - 1).$$

Since there is always at least one illegal turn between H_{l_k}-directions (for example, the one in ρ) there must be at least one vertex $v \in \mathcal{V} \cup \{w\}$ with $\kappa(v) \neq 0$. We conclude that $\Delta_k j \leq \Delta_k \chi^-$ as desired with equality if and only if $\sum \kappa(v) = 1$ where the sum is taken over all vertices in H_{i_k}.

- Suppose that w is new and w' is old. The change in $CS(k)$ corresponding to w and w' is the addition of a new vertex w, an edge representing ρ connecting w to w', one ray for each fixed H_{l_k}-direction based at w or w' other than d and d' and then one ray corresponding to d and d'. This yields

$$\Delta_k j(\rho) = \frac{1}{2}(L(w) + L(w')) - \frac{1}{2},$$

$$\Delta_k \chi^-(\rho) = \frac{1}{2}(L(w) + L(w') + \kappa(w) + \kappa(w')) - 1$$

and the proof concludes as in the previous case.

- Suppose that w and w' are distinct and both new. The change in $CS(k)$ corresponding to w and w' is the addition of two new vertices, an edge representing ρ connecting them, one ray for each fixed H_{l_k}-direction based at w or w' other than d and d' and then one ray corresponding to d and d'. The calculation is the same as in the previous case. □

We have the following two corollaries to the proof of Proposition 15.14.

Corollary 15.16. *If G has no EG-strata then $j(\phi) = n - 1$. If G has an EG-stratum without Nielsen paths, then $j(\phi) < n - 1$.*

For the next corollary, recall (Sections 3.2 and 12.1) that there are bijections between Nielsen classes of principal vertices of G and isogredience classes in $\mathrm{P}(\phi)$ and the set of components of $S_N(f)$.

Corollary 15.17. (1) *Suppose that $S_N(f, [v])$ is a line L. Then one of its rays is NEG.*

(2) *Each $S_N(f, [v])$ contributes at least $1/2$ to $j(\phi)$.*

(3) $|[\mathrm{P}(\phi)]| \leq 2j(\phi)$, *i.e. $S_N(f)$ has at most $2j(\phi)$ components.*

Proof. (3) follows from (2) We will prove (1) and (2) simultaneously.

Since $S_N(f,[v])$ is a weakly core graph we may assume that each vertex of $S_N(f,[v])$ has valence two. In particularly, if $[w]=[v]$ then w is not the endpoint of an indivisible Nielsen path (see Section 12) and w has only two gates in G. We view $S_N(f)$ as being built up as in the proof of Proposition 15.14. Let w be the lowest principal vertex satisfying $[w]=[v]$ and let G_{l_k} be the lowest core filtration element that contains w. We consider the cases enumerated in the proof of Proposition 15.14. Case (1a) is ruled out because w is principal and so would have at least three gates. Since w is new in H_{l_k} we are not in case (1b). If we are in case (1c) then both edges must be non-linear and $j(\Gamma)\geq 1$. It remains to rule out case (2). Since w is a new vertex in the EG strata H_{l_k} there are at least two gates in H_{l_k} based at w. Since w is principal but is is not the endpoint of an indivisible Nielsen path, there must be at least three gates at w. This completes the proof of (1) and (2). □

Corollary 15.18. $\left|\left(\bigcup_{\Phi\in \mathrm{P}(\phi)}\mathrm{Fix}_+(\partial\Phi)\right)/F_n\right|\leq 6(n-1)$.

Proof.

$$\begin{aligned}
j(\phi) &= \sum_{[v]} j(S_N(f,[v]))\\
&\geq \sum_{[v]}\left(\frac{a(S_N(f,[v]))}{2}-1\right)\\
&= \frac{|\text{ends of } S_N(f)|}{2} - |\text{components of } S_N(f)|\\
&\geq \frac{1}{2}\cdot\left|\left(\bigcup_{\Phi\in \mathrm{P}(\phi)}\mathrm{Fix}_+(\partial\Phi)\right)/F_n\}\right| - 2j(\phi)
\end{aligned}$$

Thus $\left|\left(\bigcup_{\Phi\in \mathrm{P}(\phi)}\mathrm{Fix}_+(\partial\Phi)\right)/F_n\right|\leq 6j(\phi)\leq 6(n-1)$ by Proposition 15.14. □

Remark 15.19. We commented in Remark 3.15 how Proposition 15.14 could be used to improve the bound in Lemma 3.13 and hence also Corollary 3.14. With more care, Proposition 15.14 could be used to further improve this bound.

16. Appendix: Hyperbolic and atoroidal automorphisms

In this appendix we reprove a result of Brinkmann (Lemma 16.2) and a result of Kapovich (Corollary 16.4).

Definition 16.1. An outer automorphism $\phi\in\mathtt{Out}(F_n)$ is *hyperbolic* if for some $N>0$ and $\lambda>1$,

$$\lambda\|\alpha\|\leq\max\{\|\phi^N(\alpha)\|,\|\phi^{-N}(\alpha)\|\}$$

for all non-trivial conjugacy classes α in F_n. We say ϕ is *atoroidal* if ϕ has no non-trivial periodic conjugacy classes. $\Phi \in \mathsf{Aut}(F_n)$ is *hyperbolic* if for some $N > 0$ and $\lambda > 1$,

$$\lambda|a| \leq \max\{|\Phi^N(a)|, |\Phi^{-N}(a)|\}$$

for all non-trivial a in F_n. Here $\|\cdot\|$ and $|\cdot|$ denote respectively reduced word length and word length with respect to a fixed basis for F_n. The *mapping torus M_Φ of* $\Phi \in \mathsf{Aut}(F_n)$ is the group with presentation

$$\langle F_n, t \mid tat^{-1} = \Phi(a) \text{ for each } a \in F_n\rangle.$$

Lemma 16.2 ([Bri00]). *Suppose that no conjugacy class in F_n is fixed by an iterate of ϕ. Then ϕ is hyperbolic.*

Proof. Suppose that $\Lambda_1^\pm, \dots, \Lambda_m^\pm$ are the lamination pairs for ϕ. After replacing ϕ by an iterate, we may assume that ϕ and ϕ^{-1} are rotationless. Since ϕ does not fix any conjugacy classes, each $\Lambda_i^\pm$ is non-geometric, and so ([HM, Part III Theorem F]; see also [BFH00, Definition 5.1.4 and Theorem 6.0.1]) there is a ϕ-invariant free factor system $\mathcal{A}_i$ for which the following are equivalent for each conjugacy class $[a]$ in F_n:

- $[a]$ is not weakly attracted to Λ_i^+ under iteration by ϕ;
- $[a]$ is not weakly attracted to Λ_i^- under iteration by ϕ^{-1};
- $[a]$ is carried by $\mathcal{A}_i$.

In particular, $\mathcal{A}_i$ does not carry $\Lambda_i^\pm$.

We claim that no line is carried by every $\mathcal{A}_i$. If this failed, then one could choose free factors A_i such that $[A_i]$ is a component of $\mathcal{A}_i$ and such that $A := A_1 \cap \cdots \cap A_m$ is non-trivial. Thus A is a non-trivial free factor that does not carry any $\Lambda_i^\pm$. There is a CT representing ϕ in which $[A]$ is represented by a filtration element. The lowest stratum cannot be EG and so must be a fixed loop. But that contradicts the assumption that there are no ϕ-invariant conjugacy classes and so completes the proof of the claim.

Choose CTs $f: G \to G$ representing ϕ and $f': G' \to G'$ representing ϕ^{-1}. Let λ_i the expansion factor ([BFH00, Definition 3.3.2]) for ϕ with respect to Λ_i^+, μ_i the expansion factor for ϕ^{-1} with respect to Λ_i^- and $\lambda = \min\{\lambda_i, \mu_i\} > 1$.

Let $H_i \subset G$ be the stratum corresponding to Λ_i^+. By [BFH00, Lemma 4.2.2] there exists a subpath $\delta_i \subset G$ of Λ_i^+ (any subpath of Λ_i^+ that crosses sufficiently many edges of H_i will do) with the following property: if $\sigma \subset G$ is a circuit or path and $\sigma_0 \subset \sigma$ is a copy of δ_i or its inverse $\bar{\delta}_i$, then there is a splitting of σ in which one of the terms is an edge of H_i contained in σ_0. Let $B^+(\sigma)$ be the maximum number of disjoint subpaths of σ that are copies of some δ_i or $\bar{\delta}_i$. Then σ has a splitting in which $B^+(\sigma)$ terms are single edges in EG strata. It follows that

$$|f_\#^k(\sigma)| > CB^+(\sigma)\lambda^k$$

for all $k \geq 1$ where C is a positive constant and $|\cdot|$ denotes length.

Define $\delta'_i \subset G'$ and $B^-(\sigma')$ symmetrically replacing Λ_i^+ and $f: G \to G$ with Λ_i^- with $f': G' \to G'$. Each path or circuit $\sigma' \subset G'$ has a splitting in which $B^-(\sigma')$ terms are single edges in EG strata of $f': G' \to G'$. After decreasing C if necessary, we have

$$|(f')^k_\#(\sigma')| > CB^-(\sigma')\lambda^k$$

for all $k \geq 1$.

Let $h: G \to G'$ be a homotopy equivalence that preserves markings. After replacing ϕ (and hence f and f') by an iterate, we may assume that for each i the neighborhood $V(\delta_i)$ of Λ_i^+ consisting of lines that contain either δ_i or $\bar{\delta}_i$ as a subpath is mapped into itself by ϕ and that the neighborhood $V(\delta'_i)$ of Λ_i^- determined by δ'_i is mapped into itself by ϕ^{-1}. By [HM, Part III Theorem H] there is a positive integer N_i so that if $\beta \subset G$ is a line that is not carried by $\mathcal{A}_i$ then either $h_\#(\beta)$ contains a copy of δ'_i or $\bar{\delta}'_i$ or $f_\#^{N_i}(\beta)$ contains a copy of δ_i or $\bar{\delta}_i$. Let $N := \max\{N_i\}$. Since $f_\#$ maps each $V(\delta_i)$ into itself and since no line is carried by $\mathcal{A}_i$ for all i, we have

$$B^+(f_\#^N(\beta)) + B^-(h_\#(\beta)) \geq 1$$

for all lines β.

Recall (see [HM, Part I Section 1.1.6]) that for all paths $\alpha \subset G$ there is a path $h_{\#\#}(\alpha) \subset G'$ obtained from $h_\#(\sigma)$ by removing initial and terminal segments of length at most the bounded cancellation constant of h and satisfying $h_{\#\#}(V(\alpha)) \subset V(h_\#(\alpha))$. We claim that there exists a positive integer L so that if $\sigma \subset G$ has length at least L then

$$B^+(f_\#^N(\sigma)) + B^-(h_{\#\#}(\sigma)) \geq 1.$$

Indeed, if this fails then there exists $L_j \to \infty$ and paths $\sigma_j \subset G$ with length $\geq L_j$ such that for all $1 \leq i \leq m$

- $h_{\#\#}(\sigma_j)$ does not contain δ'_i or $\bar{\delta}'_i$ as a subpath.
- $f_\#^N(\sigma_j)$ does not contain δ_i or $\bar{\delta}_i$ as a subpath.

By focusing on the 'middle' of each σ_j and passing to a subsequence we may assume that there are subpaths $\beta_j \subset \sigma_j$ such that $\beta_1 \subset \beta_2 \subset \cdots$ is an increasing sequence of paths whose union is a line $\beta \subset G$. As verified above, there exists $1 \leq i \leq m$ such that $\mathcal{A}_i$ does not carry β. By [HM, Part I Lemma 1.6(3)] $h_{\#\#}(\beta_j) \subset h_{\#\#}(\beta_{j+1}) \cap h_{\#\#}(\sigma_j)$ and $h_\#(\beta)$ is the union of the $h_{\#\#}(\beta_j)$'s. It follows that $h_\#(\beta)$ does not contain δ'_i or $\bar{\delta}'_i$ as a subpath and so $f_\#^N(\beta)$ must contain a copy of δ_i or $\bar{\delta}_i$. But then $f_{\#\#}^N(\beta_j)$ contains a copy of δ_i or $\bar{\delta}_i$ for all sufficiently large j and hence $f_\#^N(\sigma_j)$ contains a copy of δ_i or $\bar{\delta}_i$ for all sufficiently large j. This contradiction completes the proof of the claim.

We are now ready to complete the proof. After replacing ϕ with ϕ^N, we may assume that $N = 1$. Given a circuit $\sigma \subset G$ with $|\sigma| \geq 2L$ divide it into $[\frac{|\sigma|}{L}]$ subpaths σ_l of length at least L where $[\cdot]$ is the greatest integer function. Applying the preceding claim and the fact that the $f_{\#\#}(\sigma_l)$'s are disjoint subpaths of $f_\#(\sigma)$ we have

$$\max\{B^+(f_\#(\sigma)), B^-(\sigma')\} \geq \left[\frac{|\sigma|}{2L}\right]$$

In conjunction with the above displayed inequalities this completes the proof of the lemma if $|\hat{\sigma}| \geq 2L$. As there are only finitely many remaining $\hat{\sigma}$ and none of these is fixed by an iterate of ϕ we are done. □

Proposition 16.3. *Let* $\Phi \in \mathsf{Aut}(F_n)$ *represent* $\phi \in \mathsf{Out}(F_n)$. *The following are equivalent*:

(1) ϕ *is atoroidal*;

(2) ϕ *is hyperbolic*;

(3) Φ *is hyperbolic*;

(4) M_Φ *is hyperbolic.*

Proof. (1) $\Longrightarrow$ (2) is Lemma 16.2. (2) $\Longrightarrow$ (3) follows from the proof of Theorem 5.1 of [BFH97]. (3) $\Longrightarrow$ (4) is a consequence of the first corollary of Section 5 of [BF92]. (4) $\Longrightarrow$ (1) since otherwise M_Φ contains $\mathbb{Z}^2$. □

Corollary 16.4 ([Kap00]; see also [Dah16, p. 2]). *There is an algorithm with input* $\phi \in \mathsf{Out}(F_n)$ *that outputs* YES *or* NO *depending on whether or not* ϕ *is hyperbolic.*

Proof. Construct a CT $f : G \to G$ for a ϕ^M with M as in Corollary 3.14. By Proposition 16.3, ϕ is hyperbolic if and only if ϕ is atoroidal. By Lemma 9.6(2), ϕ is atoroidal if and only if $S(f)$ has no circuits and this can be checked algorithmically; see Remark 9.5. □

References

[BF92] M. Bestvina and M. Feighn, A combination theorem for negatively curved groups. *J. Differential Geom.* **35** (1992), no. 1, 85–101. Zbl 0724.57029 MR 1152226

[BF14] M. Bestvina and M. Feighn, Hyperbolicity of the complex of free factors. *Adv. Math.* **256** (2014), 104–155. Zbl 1348.20028 MR 3177291

[BFH97] M. Bestvina, M. Feighn, and M. Handel, Laminations, trees, and irreducible automorphisms of free groups. *Geom. Funct. Anal.* **7** (1997), no. 2, 215–244. Zbl 0884.57002 MR 1445386

[BFH00] M. Bestvina, M. Feighn, and M. Handel, The Tits alternative for Out(F_n). I. Dynamics of exponentially-growing automorphisms. *Ann. of Math.* (2) **151** (2000), no. 2, 517–623. Zbl 0984.20025 MR 1765705

[BH92] M. Bestvina and M. Handel, Train tracks and automorphisms of free groups. *Ann. of Math.* (2) **135** (1992), no. 1, 1–51. Zbl 0757.57004 MR 1147956

[BM16] O. Bogopolski and O. Maslakova, An algorithm for finding a basis of the fixed point subgroup of an automorphism of a free group. *Internat. J. Algebra Comput.* **26** (2016), no. 1, 29–67. Zbl 1348.20030 MR 3463201

[Bri00] P. Brinkmann, Hyperbolic automorphisms of free groups. *Geom. Funct. Anal.* **10** (2000), no. 5, 1071–1089. Zbl 0970.20018 MR 1800064

[BT68] G. Baumslag and T. Taylor, The centre of groups with one defining relator. *Math. Ann.* **175** (1968), 315–319. Zbl 0157.34901 MR 0222144

[CG10] A. Clifford and R. Z. Goldstein, Subgroups of free groups and primitive elements. *J. Group Theory* **13** (2010), no. 4, 601–611. Zbl 1206.20025 MR 2661660

[CMP15] M. Clay, J. Mangahas, and A. Pettet, An algorithm to detect full irreducibility by bounding the volume of periodic free factors. *Michigan Math. J.* **64** (2015), no. 2, 279–292. Zbl 1336.20037 MR 3359026

[Coo87] D. Cooper, Automorphisms of free groups have finitely generated fixed point sets. *J. Algebra* **111** (1987), no. 2, 453–456. Zbl 0628.20029 MR 0916179

[Cul84] M. Culler, Finite groups of outer automorphisms of a free group. In K. I. Appel, J. G. Ratcliffe and P. E. Schupp (eds.), *Contributions to group theory.* Contemporary Mathematics, 33. American Mathematical Society, Providence, R.I., 1984, 197–207. Zbl 0552.20024 MR 0767107

[Dah16] F. Dahmani, On suspensions and conjugacy of hyperbolic automorphisms. *Trans. Amer. Math. Soc.* **368** (2016), no. 8, 5565–5577. Zbl 06551574 MR 3458391

[DF05] G.-A. Diao and M. Feighn, The Grushko decomposition of a finite graph of finite rank free groups: an algorithm. *Geom. Topol.* **9** (2005), 1835–1880. Zbl 1093.20022 MR 2175158

[Dic14] W. Dicks, On free-group algorithms that sandwich a subgroup between free-product factors. *J. Group Theory* **17** (2014), no. 1, 13–28. Zbl 1301.20023 MR 3176649

[DV96] W. Dicks and E. Ventura, *The group fixed by a family of injective endomorphisms of a free group.* Contemporary Mathematics, 195, American Mathematical Society, Providence, R.I., 1996. Zbl 0845.20018 MR 1385923

[FH09] M. Feighn and M. Handel, Abelian subgroups of Out(F_n). *Geom. Topol.* **13** (2009), no. 3, 1657–1727. Zbl 1201.20031 MR MR2496054

[FH11] M. Feighn and M. Handel, The recognition theorem for Out(F_n), *Groups Geom. Dyn.* **5** (2011), no. 1, 39–106. Zbl 1239.20036 MR 2763779

[Ger84] S. M. Gersten, On Whitehead's algorithm. *Bull. Amer. Math. Soc.* (*N.S.*) **10** (1984), no. 2, 281–284. Zbl 0537.20015 MR 0733696

[GJLL98] D. Gaboriau, A. Jaeger, G. Levitt, and M. Lustig, An index for counting fixed points of automorphisms of free groups. *Duke Math. J.* **93** (1998), no. 3, 425–452. Zbl 0946.20010 MR 1626723

[HM] M. Handel and L. Mosher, Subgroup decomposition in Out(F_n). To appear in *Mem. Amer. Math. Soc.*

[HM11] M. Handel and L. Mosher, Axes in outer space. *Mem. Amer. Math. Soc.* **213** (2011), no. 1004. Zbl 1238.57002 MR 2858636

[IKT90] W. Imrich, S. Krstić, and E. C. Turner, On the rank of fixed point sets of automorphisms of free groups. In G. Hahn, G. Sabidussi, and R. E. Woodrow (eds.), *Cycles and rays.* (Montreal, 1987.) NATO Advanced Science Institutes Series C: Mathematical and Physical Sciences, 301. Kluwer, Dordrecht, 1990, 113–122. Zbl 0706.20025 MR 1096989

[Kap] I. Kapovich, Detecting fully irreducible automorphisms: a polynomial time algorithm. With an appendix by M. C. Bell. Preprint, 2016. arXiv:1609.03820 [math.GR]

[Kap00] I. Kapovich, Mapping tori of endomorphisms of free groups. *Comm. Algebra* **28** (2000), no. 6, 2895–2917. Zbl 0953.20035 MR 1757436

[Kap14] I. Kapovich, Algorithmic detectability of iwip automorphisms. *Bull. Lond. Math. Soc.* **46** (2014), no. 2, 279–290. Zbl 1319.20030 MR 3194747

[LL08] G. Levitt and M. Lustig, Automorphisms of free groups have asymptotically periodic dynamics. *J. Reine Angew. Math.* **619** (2008), 1–36. Zbl 1157.20017 MR 2414945

[Mar95] R. Martin, *Non-uniquely ergodic foliations of thin type, measured currents, and automorphisms of free groups.* Ph.D. thesis. University of California Los Angeles. Los Angeles, CA, 1995.

[Nie29] J. Nielsen, Untersuchungen zur Topologie der geschlossenen zweiseitigen Flächen. II. *Acta Math.* **53** (1929), no. 1, 1–76. JFM 55.0971.01 MR 1555290

[Nie86] J. Nielsen, *Collected mathematical papers.* Vol. 2. Edited and with a preface by V. Lundsgaard Hansen. Contemporary Mathematicians. Birkhäuser Boston, Boston, MA, 1986. Zbl 0609.01050 MR 865336

[Sta83] J. Stallings, Topology of finite graphs. *Invent. Math.* **71** (1983), no. 3, 551–565. Zbl 0521.20013 MR 0695906

Received January 27, 2017

Mark Feighn, Department of Mathematics and Computer Science, Rutgers University, Newark, NJ 07102, USA

e-mail: feighn@rutgers.edu

Michael Handel, Math Department, Lehman College, Bronx, NY 10468, USA

e-mail: michael.handel@lehman.cuny.edu